Network Working Group
Internet Draft
Intended status: Standards Track
Expires: December 09, 2025

C. Lin
M. Chen
New H3C Technologies
L. Gong
China Mobile

June 07, 2025

IGP Color-Aware Routing
draft-lin-lsr-igp-car-03

Abstract

This document describes an IGP based routing solution to establish end-to-end intent-aware paths across a multi-domain service provider transport network.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

[I-D.hr-spring-intentaware-routing-using-color] describes the scope, set of use-cases and requirements for a distributed routing based solution to establish end-to-end intent-aware paths spanning multi-domain packet networks. [I-D.ietf-idr-bgp-car] proposes a BGP based solution called "BGP Color-Aware Routing" (BGP CAR). [I-D.ietf-idr-bgp-ct] also proposes a BGP based solution called "BGP Classful Transport" (BGP CT).

Network operators often organize networks into multiple smaller network domains, and each network domain typically runs an IGP. In some scenarios, only the PE nodes run the BGP protocol, while the other nodes (including the ASBR nodes) run only the IGP protocol, as shown in Figure 1.

```
                           BGP
            .....................................
            :                                   :
            : +----------+----------+----------+ :
            : |          |          |          | :
            : |       ASBR121    ASBR231       | :
            : |          |          |          | :
           PE1+    D1    |    D2    |    D3    +PE2
              |          |          |          |
              |       ABSR122    ASBR232       |
              |          |          |          |
              +----------+----------+----------+

              |---IGP1---|---IGP2---|---IGP3---|
```

Figure 1: Multi IGP Domain Network Example

The BGP based solutions proposed by [I-D.ietf-idr-bgp-car] and [I-D.ietf-idr-bgp-ct] require the ASBR nodes to run BGP and signal BGP routes for CAR path. However, some network operators may not want to change existing routing protocol deployments.

This document describes an IGP based routing solution to establish end-to-end intent-aware paths, without requirement for the ASBR nodes to run BGP.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. IGP CAR Solution

The following sub-sections illustrate example scenarios of Colored Service Route Steering over E2E IGP CAR resolving over different intra-domain mechanisms.

The examples use MPLS-SR for the transport data plane. The main difference on SRv6 data plane is the replacement of CAR labels with CAR SIDs.

The OSPF and IS-IS extensions for advertising CAR route will be described in Section 3. The resolving of CAR route will be described in Section 4.

2.1. SR Policy based IGP CAR

Figure 2 illustrates an example scenario of Colored Service Route Steering over E2E IGP CAR based on SR Policy.

```
                                         RD:V/v via E2
                                 +--+    vpn label:30030
    .............................|RR|<.............................
    :                            +--+    Color C1                 :
    :                                                             :
   +:-------------------+--------------------+--------------------:+
   |:                   |                    |                    :|
   |:   Redist: Pfx=E2  |    Redist: Pfx=E2  |                    :|
   |:    (C1,LI=8002)<~~|~~~  (C1,LI=8002)<~~|~~~                 :|
   |:                 +---+                +---+ SR Policy(C1,E2):|
   |: |-------------->|121|--------------->|231|----------------| :|
   |: |   SR Policy   +---+    SR Policy   +---+                | :|
   |: |   (C1,121)      |      (C1,231)      |       CAR route: v :|
   |--+                 |                    |           Pfx=E2 +--|
   |E1|                 |                    |           C=C1   |E2|
   |--+ Redist: Pfx=E2  |    Redist: Pfx=E2  |           LI=8002+--|
   |  |  (C1,LI=8002)<~~|~~~  (C1,LI=8002)<~~|~~~               ^  |
   |  |               +---+                +---+                |  |
   |  |-------------->|122|--------------->|232|----------------|  |
   |      SR Policy   +---+    SR Policy   +---+     SR Policy     |
   |      (C1,122)      |      (C1,232)      |       (C1,E2)       |
   |                    |                    |                     |
   |       IGP 1        |       IGP 2        |       IGP 3         |
   +--------------------+--------------------+---------------------+
   iPE                iASBR                eASBR                ePE

   +------+             +------+             +------+
   |  S1  |             |  S2  |             |  S3  |
   +------+             +------+             +------+
   +------+             +------+             +------+
   |168002|             |168002|             |168002|
   +------+             +------+             +------+
   +------+             +------+             +------+
   |30030 |             |30030 |             |30030 |
   +------+             +------+             +------+
```

Figure 2: SR Policy based IGP CAR Path

With reference to the topology above:

- o SR Policy provides intent in each domain.

- o Egress PE E2 advertises a VPN route RD:V/v colored with (color extended community) C1 to steer traffic to SR Policy. VPN route propagates via RR(s) to ingress PE E1.

- o IGP CAR route of prefix E2 with color C1 and label 168002 (index is 2 and SRGB is 168000) is advertised by Egress PE E2, and redistributed by ASBRs.

- o On each ASBR and the Ingress PE E1, CAR label 168002 is resolved over SR Policy of the domain.

- o Ingress PE E1 steers colored VPN route RD:V/v into SR Policy according to color.

Packet forwarding:

```
@E1:  IPv4 VRF V/v => PUSH <S1, 168002, 30030>
@E1:  MPLS Table: S1 => forward via SR Policy to 121
@121: MPLS Table: 168002 => PUSH S2
@121: MPLS Table: S2 => forward via SR Policy to 231
@231: MPLS Table: 168002 => PUSH S1
@231: MPLS Table: S3 => forward via SR Policy to E2
@E2:  MPLS Table: 168002 => POP => MPLS Table: 30030 => POP and
lookup the IP DA in the VRF
```

2.2. Flex-Algo based IGP CAR

Figure 3 illustrates an example scenario of Colored Service Route Steering over E2E IGP CAR based on Flex-Algo (FA).

```
                                         RD:V/v via E2
                                 +--+    vpn label:30030
    .............................|RR|<............................
    :                            +--+    Color C1                :
    :                                                            :
   +:-------------------+--------------------+-------------------:+
   |:                   |                    |                   :|
   |:   Redist: Pfx=E2  |    Redist: Pfx=E2  |                   :|
   |:    (C1,LI=8002)<~~|~~~  (C1,LI=8002)<~~|~~~                :|
   |:                 +---+                +---+                 :|
   |:     Pfx-sid=8121|121|    Pfx-sid=8231|231|                 :|
   |:        Algo=128 +---+       Algo=128 +---+     CAR route:  :|
   |--+                 |                    |          Pfx=E2 +--|
   |E1|                 |                    |          C=C1   |E2|
   |--+ Redist: Pfx=E2  |    Redist: Pfx=E2  |          LI=8002+--|
   |     (C1,LI=8002)<~~|~~~  (C1,LI=8002)<~~|~~~                 |
   |                  +---+                +---+     Pfx-sid=8021 |
   |      Pfx-sid=8122|122|    Pfx-sid=8232|232|        Algo=128  |
   |         Algo=128 +---+       Algo=128 +---+                  |
   |                    |                    |                    |
   |       IGP 1        |       IGP 2        |       IGP 3        |
   | FA 128(Mapping C1) | FA 128(Mapping C1) | FA 128(Mapping C1) |
   +--------------------+--------------------+--------------------+
   iPE                iASBR                eASBR                ePE

   +------+             +------+             +------+
   |168121|             |168231|             |168021|
   +------+             +------+             +------+
   +------+             +------+             +------+
   |168002|             |168002|             |168002|
   +------+             +------+             +------+
   +------+             +------+             +------+
   |30030 |             |30030 |             |30030 |
   +------+             +------+             +------+
```

Figure 3: FA based IGP CAR Path

With reference to the topology above:

- o IGP FA 128 is running in each domain, and mapped to Color C1.

- o Egress PE E2 advertises a VPN route RD:V/v colored with (color extended community) C1 to steer traffic to IGP FA 128. VPN route propagates via RR(s) to ingress PE E1.

- o IGP CAR route of prefix E2 with color C1 and label 168002 (index is 8002 and SRGB is 160000) is advertised by Egress PE E2, and redistributed by ASBRs.

- o On each ASBR and the Ingress PE E1, label 168002 is resolved over IGP FA 128 of the domain.

- o Ingress PE E1 steers colored VPN route RD:V/v into IGP FA 128 according to the mapping relationship between FA and Color.

Packet forwarding:

```
@E1:  IPv4 VRF V/v => PUSH <168121, 168002, 30030>
@E1:  MPLS Table: 168121 => forward via FA path to 121
@121: MPLS Label: 168002 => PUSH 168231
@121: MPLS Label: 168231 => forward via FA path to 231
@231: MPLS Label: 168002 => PUSH 168021
@231: MPLS Label: 168021 => forward via FA path to E2
@E2:  MPLS Label: 168002 => POP => 30030 => POP and lookup the IP DA
in the VRF
```

2.3. Hybrid IGP CAR

Figure 4 illustrates an example where the same intent is provided by SR Policy in some domains but by Flex-Algo in some other domains.

```
                                         RD:V/v via E2
                                 +--+    vpn label:30030
    .............................|RR|<.............................
    :                            +--+    Color C1                 :
    :                                                             :
   +:-------------------+--------------------+--------------------:+
   |:                   |                    |                    :|
   |:   Redist: Pfx=E2  |     Redist: Pfx=E2 |                    :|
   |:    (C1,LI=8002)<~~|~~~   (C1,LI=8002)<~~|~~~                 :|
   |:                 +---+                +---+ SR Policy(C1,E2):|
   |: |-------------->|121|    Pfx-sid=8231|231|----------------| :|
   |: |   SR Policy   +---+       Algo=128 +---+                | :|
   |: |   (C1,121)      |                    |       CAR route: v :|
   |--+                 |                    |           Pfx=E2 +--|
   |E1|                 |                    |           C=C1   |E2|
   |--+ Redist: Pfx=E2  |     Redist: Pfx=E2 |           LI=8002+--|
   |  |  (C1,LI=8002)<~~|~~~  (C1,LI=8002)<~~|~~~               ^  |
   |  |               +---+                +---+                |  |
   |  |-------------->|122|    Pfx-sid=8232|232|----------------|  |
   |      SR Policy   +---+       Algo=128 +---+    SR Policy      |
   |      (C1,122)      |                    |      (C1,E2)        |
   |                    |                    |                     |
   |       IGP 1        |       IGP 2        |       IGP 3         |
   |       Algo 0       | FA 128(Mapping C1) |       Algo 0        |
   +--------------------+--------------------+---------------------+
   iPE                iASBR                eASBR                ePE

   +------+             +------+             +------+
   |  S1  |             |168231|             |  S3  |
   +------+             +------+             +------+
   +------+             +------+             +------+
   |168002|             |168002|             |168002|
   +------+             +------+             +------+
   +------+             +------+             +------+
   |30030 |             |30030 |             |30030 |
   +------+             +------+             +------+
```

Figure 4: Hybrid IGP CAR Path

With reference to the topology above:

- o SR Policy provides intent in domain 1 and domain 3. IGP FA 128 is running only in domain 2, and mapped to Color C1.

- o IGP CAR route of prefix E2 with color C1 and label 168002 (index is 2 and SRGB is 168000) is advertised by Egress PE E2, and redistributed by ASBRs.

o On ASBR 231 and the Ingress PE E1, label 168002 is resolved over SR Policy. On ASBR 121, label 168002 is resolved over IGP FA 128

Packet forwarding:

```
@E1:  IPv4 VRF V/v => PUSH <S1, 168002, 30030>
@E1:  MPLS Table: S1 => forward via SR Policy to 121
@121: MPLS Label: 168002 => PUSH 168231
@121: MPLS Label: 168231 => forward via FA path to 231
@231: MPLS Table: 168002 => PUSH S1
@231: MPLS Table: S3 => forward via SR Policy to E2
@E2:  MPLS Label: 168002 => POP => 30030 => POP and lookup the IP DA
in the VRF
```

3. Advertisement of IGP CAR Route

The advertisement of IGP CAR route is as following:

```
  Prefix TLV: E
    CAR Sub-TLV: C
      CAR Encapsulation Sub-sub-TLV: T
```

The new-defined CAR sub-TLV is attached to a prefix E, indicating a color-aware path of color C and encapsulation T towards that prefix. On the MPLS-SR data plane, T is a label or index. On the SRv6 data plane, T is an SRv6 SID.

If there are multiple color-aware path for different intents towards the same E. Multiple CAR sub-TLVs with different colors will be advertised.

```
  Prefix TLV: E
   CAR Sub-TLV: C1 + T1
   CAR Sub-TLV: C2 + T2
   CAR Sub-TLV: C3 + T3
   ...
```

3.1. IS-IS CAR Sub-TLV

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

The Sub-TLV has the following format:

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

o Type: TBD.

o Length: 1 octet. The length value is variable.

o Flags: 2 octets. No flags are defined in this document. Undefined flags MUST be set to 0 by the sender, and any unknown flags MUST be ignored by the receiver.

o Color: 4 octets. Contains color value associated with the prefix.

o Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs.

o Sub-sub-TLVs: Carrying the encapsulation information. IS-IS CAR MPLS-SR Encapsulation sub-sub-TLV and IS-IS CAR SRv6 Encapsulation sub-sub-TLV are defined in Section 3.1.1 and 3.1.2. Other encapsulations may be defined in the future.

3.1.1. IS-IS CAR MPLS-SR Encapsulation Sub-sub-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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |     Type      |     Length    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                          SID/Label (variable)                 |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

o Type: 1.

o Length: 3 or 4.

o SID/Label: If the length is set to 3, then the 20 rightmost bits represent an MPLS label. If the length is set to 4, then the value is a 32-bit index.

3.1.2. IS-IS CAR SRv6 Encapsulation Sub-sub-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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |     Type      |     Length    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | SID (128 bits) . . .                                          |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | SID (cont . . .)                                              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | SID (cont . . .)                                              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | SID (cont . . .)                                              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

o Type: 2.

o Length: 16.

o SID: 16 octets. This field encodes the SRv6 SID.

3.2. OSPF CAR Sub-TLV

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

The Sub-TLV has the following format:

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

o Type: TBD.

o Length: 1 octet. The length value is variable.

o Flags: 2 octets. No flags are defined in this document. Undefined flags MUST be set to 0 by the sender, and any unknown flags MUST be ignored by the receiver.

o Color: 4 octets. Contains color value associated with the prefix.

o Sub-sub-TLVs: Carrying the encapsulation information. OSPF CAR MPLS-SR Encapsulation sub-sub-TLV and OSPF CAR SRv6 Encapsulation sub-sub-TLV are defined in Section 3.2.1 and 3.2.2. Other encapsulations may be defined in the future.

3.2.1. OSPF CAR MPLS-SR Encapsulation Sub-sub-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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         SID/Label (variable)                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

o Type: 1.

o Length: 3 or 4.

o SID/Label: If the length is set to 3, then the 20 rightmost bits represent an MPLS label. If the length is set to 4, then the value is a 32-bit index.

OSPF CAR MPLS-SR Encapsulation Sub-sub-TLV is applicable to the OSPF CAR Sub-TLV carried in OSPFv2 Extended Prefix TLV.

3.2.2. OSPF CAR SRv6 Encapsulation Sub-sub-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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |            Length             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID (128 bits) . . .                                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID (cont . . .)                                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID (cont . . .)                                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID (cont . . .)                                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- o Type: 2.
- o Length: 16.
- o SID: 16 octets. This field encodes the SRv6 SID.

OSPF CAR SRv6 Encapsulation Sub-sub-TLV is applicable to the OSPF CAR Sub-TLV carried in OSPFv3 Inter-Area-Prefix TLV, OSPFv3 Intra-Area-Prefix TLV, and OSPFv3 External-Prefix TLV.

4. Resolving of IGP CAR Route

When the ASBR or Ingress PE receives the IGP CAR route, it will be resolved over the CAR paths, and then installed to CAR Route Database (CARDB). CARDB is a logical collection of resolved CAR routes. The ASBR 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.

```
          IGP Instance A
   +------------------------------+
   |IGP CAR Route                 |
   |     |                        |
   |     | Resolving over CAR path|
   |     | 1. SR Policy           |
   |     | 2. Flex-Algo           |
   |     | ...                    |
   +-----+------------------------+
         |
         |Installation            +---------------+
         v                     +->|IGP Instance B|
   +-----+-----+               |  +---------------+
   |    CAR    | Redistribution |  +---------------+
   |   Route   |-----------------+->|IGP Instance C|
   |  Database|                |  +---------------+
   +-----+-----+               |  +---------------+
         |                     +->|      ...      |
         |Next-hop Resolution     +---------------+
         v
   +-----+-----+
   |    BGP    |
   +-----------+
```

Figure 5: CAR Route Database

Take the following IGP CAR Route on MPLS-SR data plane as an example:

```
  Advertiser: N
    Prefix TLV: E
     CAR Sub-TLV: C1 + L1
     CAR Sub-TLV: C2 + L2
     CAR Sub-TLV: C3 + L3
     CAR Sub-TLV: C4 + L4
```

Assume that:

C1 is mapped to FA 128, and N’s Prefix-SID in algorithm 128 is L128.
C2 is mapped to FA 129, and N’s Prefix-SID in algorithm 129 is L129.
C3 path is provided by SR Policy 1 (C3, N) with segment-list S1.
C4 path is provided by SR Policy 2 (C4, N) with segment-list S2.

When the ASBR or Ingress PE resolved it over the CAR paths, the example of MPLS forwarding entries is as following:

Index	In-Label	Out-Label	Next-Hop & Out-Intf
1	L1	<L128,L1>	FA 128 Path
2	L2	<L129,L2>	FA 129 Path
3	L3	<S1,L3>	SR Policy 1 Path
4	L4	<S2,L4>	SR Policy 2 Path

The example of CAR route entries installed in the CARDB is as following:

Key	Encap	Proto	Forwarding Path
E + C1	L1	IS-IS 1	MPLS Index 1
E + C2	L2	IS-IS 1	MPLS Index 2
E + C3	L3	IS-IS 1	MPLS Index 3
E + C4	L4	IS-IS 1	MPLS Index 4

The resolving of IGP CAR route SHOULD only be enabled on the ASBR and Ingress PE. The P nodes do not need to resolve IGP CAR routes, and only do SPF computation for the prefix of Egress PE, providing best-effort forwarding for traditional services.

4.1. Resolving Over SR Policy

Resolving of IGP CAR Route over SR Policy can use the enhanced IGP shortcut mechanism in [I-D.cheng-lsr-igp-shortcut-enhancement]. Briefly, the main point is to choose the SR Policy with the same color as the next-hop.

4.2. Resolving Over Flex-Algo

When resolving of IGP CAR Route over Flex-Algo, the node will determine the FA to which the color is mapped, and check if the advertiser node is reachable in the topology of that FA. If yes, use the FA path as next-hop, and add the Prefix-SID or SRv6 End SID associated with that FA into encapsulation.

The mapping relationship of FA and color should be pre-configured.

5. Security Considerations

TBD.

6. IANA Considerations

TBD.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017

[I-D.cheng-lsr-igp-shortcut-enhancement] Cheng, W., Gong, L., Lin, C., and M. Chen, "IGP Shortcut Enhancement", Work in Progress, Internet-Draft, draft-cheng-lsr-igp-shortcut-enhancement-03, 27 Feburary 2024, <https://datatracker.ietf.org/doc/html/draft-cheng-lsr-igp-shortcut-enhancement-03>.

7.2. Informational References

[I-D.hr-spring-intentaware-routing-using-color] Hegde, S., Rao, D., Sangli, S. R., Agrawal, S., Filsfils, C., Talaulikar, K., Patel, K., Uttaro, J., Decraene, B., Bogdanov, A., Jalil, L., Alston, A., Xu, X., Gulko, A., Khaddam, M., Contreras, L. M., Steinberg, D., Guichard, J., Henderickx, W., and Co-authors, "Problem statement for Inter-domain Intent-aware Routing using Color", Work in Progress, Internet-Draft, draft-hr-spring-intentaware-routing-using-color-03, 23 October 2023, <https://datatracker.ietf.org/doc/html/draft-hr-spring-intentaware-routing-using-color-03>.

[I-D.ietf-idr-bgp-car] Rao, D., Agrawal, S., and Co-authors, "BGP Color-Aware Routing (CAR)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-car-07, 3 April 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-07>.

[I-D.ietf-idr-bgp-ct] Vairavakkalai, K. and N. Venkataraman, "BGP Classful Transport Planes", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ct-31, 10 April 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ct-31>.

Authors' Addresses

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Mengxiao Chen
New H3C Technologies
China
Email: chen.mengxiao@h3c.com

Liyan Gong
China Mobile
China
Email: gongliyan@chinamobile.com