

Workgroup: Interdomain Routing Working Group

Internet-Draft: draft-wang-idr-cpr-00

Published: 25 July 2022

Intended Status: Informational

Expires: 26 January 2023

Authors: H. Wang J. Dong
 Huawei Technologies Huawei Technologies
 J. Xie X. Chen
 Huawei Technologies Huawei Technologies

BGP Colorful Prefix Routing (CPR) for SRV6 based Services

Abstract

This document describes a mechanism to advertise different IPv6 prefixes which are associated with different color attributes to establish end-to-end intent aware paths. Such IPv6 prefixes are called "colorful prefixes", and this mechanism is called Colorful Prefix Routing (CPR). The colorful prefixes are the SRV6 locator prefixes associated with different intent. SRV6 services (e.g. SRV6 VPN services) could be assigned with SRV6 SIDs under the SRV6 locator prefix with the required intent, so that the SRV6 service traffic can be steered to the end-to-end intent aware paths of the corresponding SRV6 locator prefix to meet the service requirements. The existing IPv6 unicast Address Family could be used for the advertisement of colorful prefixes, thus this mechanism is easy to interoperate and allows incremental deployment in multi-domain networks.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 26 January 2023.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Requirements Language](#)
- [2. BGP CPR](#)
 - [2.1. Colorful Prefix Allocation](#)
 - [2.2. Colorful Prefix Advertisement](#)
 - [2.3. CPR to Intra-domain Path Resolution](#)
 - [2.4. SRV6 Service Route Advertisement](#)
 - [2.5. SRV6 Service Steering](#)
- [3. Encapsulation and Forwarding Processes](#)
 - [3.1. CPR over SRV6 Intra-Domain Paths](#)
 - [3.2. CPR over MPLS Intra-Domain Paths](#)
- [4. Operational Considerations](#)
- [5. IANA Considerations](#)
- [6. Security Considerations](#)
- [7. Acknowledgements](#)
- [8. References](#)
 - [8.1. Normative References](#)
 - [8.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

With the trend of using one common network to carry multiple types of services, each service type can have different requirements on the network. Such requirements are usually considered as the "intent" of the service or customer, and is represented as an abstract notion called "color".

In network scenarios where the services are delivered across multiple network domains, there is need to provide the services with different end-to-end paths to meet the intent. [[I-D.hr-spring-intentaware-routing-using-color](#)] describes the problem statements and requirements for inter-domain intent aware routing.

The inter-domain path can be established using either MPLS or IP data plane. In MPLS based networks, the traditional inter-domain approach is to establish an end-to-end LSP based on the BGP-LU mechanisms as

defined in [\[RFC8277\]](#). Each domain or area border node needs to perform label swapping for the end-to-end BGP-LU LSP, and encapsulate the label stack which are used for the intra-domain LSP within the subsequent network domain or area.

While in IP based networks, IP reachability information can be advertised to network nodes in different domains using BGP, so that all the domain or area border nodes have the routes to the prefixes of the destination node in other domains. With the introduction of SRv6 [\[RFC8986\]](#), services are assigned with SRv6 Service SIDs [\[RFC9252\]](#), which are routable in the network according to its Locator prefix. Thus the inter-domain path can be established simply based on the inter-domain prefix routes, and the BGP-LU based LSP is not necessary in IPv6 and SRv6 based networks.

This document describes a mechanism to advertise different IPv6 prefixes of a node with different color attributes to establish end-to-end intent aware paths. Such IPv6 prefixes are called "colorful prefixes", and this mechanism is called Colorful Prefix Routing (CPR). The colorful prefixes are used as the SRv6 locators associated with different intent. SRv6 services (e.g. SRv6 VPN services) could be assigned with SRv6 SIDs under the SRv6 locator prefix according to the required intent, so that the SRv6 service traffic can be steered using the end-to-end intent aware paths of the corresponding SRv6 locator prefix to meet the service requirement. The existing IPv6 unicast Address Family could be used for the advertisement of colorful prefixes, which makes this mechanism easy to interoperate and helps the incremental deployment in multi domain networks.

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. BGP CPR

This section describes the BGP CPR mechanisms. More specifically, section 2.1 describes the allocation of the colorful IP prefixes, section 2.2 describes the advertisement of colorful prefixes in BGP, section 2.3 describes the resolution of CPR routes to the intra-domain paths, and section 2.4 describes the steering of SRv6 services to CPR routes.

2.1. Colorful Prefix Allocation

In SRv6 networks, an SRv6 Locator needs to be allocated for each node. In order to distinguish N different intent, a PE node needs to

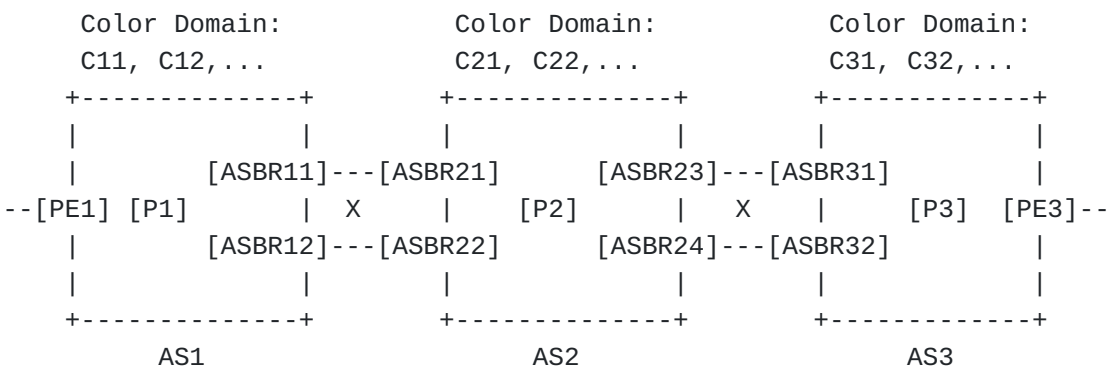
be allocated with N SRv6 Locators, each of which is associated a different intent. This can be achieved by splitting the base SRv6 locator of the node into N sub-locators, and these sub-locators are called colorful locators.

For example, node PE2 is allocated with the base SRv6 Locator 2001:db8:aaaa:1::/64. In order to provide 16 different intent, this base SRv6 Locator is split into 16 sub-locators from 2001:db8:aaaa:1:0000::/68 to 2001:db8:aaaa:1:F000::/68, each of these sub-locators is associated with a different intent, such as low-delay, high-bandwidth, etc.

2.2. Colorful Prefix Advertisement

After the allocation of colorful locator prefixes on a PE node, routes to these colorful locators need to be advertised both in the local domain and also to other domains using BGP, so that the SRv6 services route could be resolved using the corresponding CPR route.

BGP IPv6 unicast Address Family/Subsequent Address Family (AFI/SAFI = 2/1) is used for the advertisement of the colorful prefix routes. The procedure of colorful prefix advertisement is described using an example with the following topology:



Colorful Locator Prefixes of PE3:
 Low delay: 2001:db8:aaaa:1:1000::/68
 high bandwidth: 2001:db8:aaaa:1:2000::/68
 ...

Figure 1. Example Topology for CPR Route Illustration

Assume PE3 is provisioned with two different colorful locator prefixes CLP-1 and CLP-2 for two different intent such as "low-delay" and "high-bandwidth" respectively. The color for "low-delay" in AS1, AS2 and AS3 are C11, C21 and C31 respectively, and the color for "high-bandwidth" in AS1, AS2 and AS3 are C12, C22 and C32 respectively.

which are represented as color C11 and C12 respectively in domain AS1.

*PE3 originates BGP IPv6 unicast (AFI/SAFI=2/1) route for the colorful locator prefixes PE3:CL1:: and PE3:CL2::. Each route SHOULD carry the corresponding color extended community C31 or C32. PE3 also advertise a route for the base SRv6 Locator prefix PE3:BL, there is no color extended community carried with this route.

*ASBR31 and ASBR32 receive the CPR routes of PE3, and advertise the CPR routes further to ASBR23 and ASBR24 with next-hop set to itself.

*ASBR23 and ASBR24 receive the CPR routes of PE3. As the color-to-intent mapping in AS2 is different from AS3, the color in the received CPR routes are changed to the corresponding color in AS2, e.g. C21 and C22. ASBR23 and ASBR 24 advertise the CPR routes further in AS2 with the next-hop set to itself.

*The behavior of ASBR21 and ASBR22 are similar to the behavior of ASBR31 and ASBR32.

*The behavior of ASBR11 and ASBR12 are similar to behavior of ASBR31 and ASBR32. The color in the received CPR routes are changed to the corresponding color in AS1, e.g. C11 and C12.

2.3. CPR to Intra-domain Path Resolution

For a node which receives a CPR route, it SHOULD resolves the CPR routes to an intra-domain color-aware path based on the tuple (N, C), where N is the next-hop of the CPR route, and C is the color extended community of the CPR route. The intra-domain color aware path could be built with any of the following mechanisms:

*SRv6 or SR-MPLS Policy

*SRv6 or SR-MPLS Flex-Algo

*RSVP-TE

For example, PE1 receives a CPR route to PE3:CL1 with color C31 and next-hop ASBR11, it will resolve the CPR routes to an intra-domain SRv6 Policy based on the tuple (ASBR11, C31).

The intra-domain path resolution scheme could be based on any existing tunnel resolution policy, and new tunnel resolution mechanisms could also be introduced if needed.

2.4. SRV6 Service Route Advertisement

For an SRV6 service which is associated with a specific intent, the SRV6 Service SID MUST be allocated under the corresponding colorful locator prefix. For example, on PE3 in the example topology, an SRV6 VPN service with the low delay intent can be allocated with an SRV6 End.DT4 SID 2001:db8:aaaa:1:1000::0100, where 2001:db8:aaaa:1:1000::/68 is the SRV6 Colorful Locator for low delay service.

The SRV6 Service SIDs SHOULD be advertised using the mechanism defined in [[RFC9252](#)]. Inter-domain VPN Option C is used, which means the next-hop of the SRV6 service route is set to the originating PE and not changed. Since the intent of the service is embedded in the SRV6 service SID, the SRV6 service route does not need to carry color extended community.

2.5. SRV6 Service Steering

On an ingress PE node which receives a SRV6 service route, it SHOULD follow the behavior of SRV6 BE forwarding and use the SRV6 Service SID in the service route for route iteration. If the corresponding CPR route has been received and installed, the service route can be iterated and match with the to the CPR route, and the intra-domain color-aware path which the CPR route is resolved to will be used for the forwarding of the service traffic.

3. Encapsulation and Forwarding Processes

This section describes the encapsulation and forwarding process of data packets which are matched with the corresponding CPR route.

3.1. CPR over SRV6 Intra-Domain Paths

Following is an illustration of the packet encapsulation and forwarding process of CPR over SRV6 Policy. The abstract representation of IPv6 and SRH in section 6 of [[RFC8754](#)] is used.

PE3 is provisioned with a colorful locator prefix PE3:C1 for "low-delay".

In AS1, the SRV6 Policy for (ASBR11, C11) is represented with SID list (PE1, P1, BR11).

In AS2, the SRV6 Policy for (ASBR23, C21) is represented with the SID list (BR21, P2, BR23).

In AS3, the SRV6 Policy for (PE3, C31) is represented with the SID list (BR31, P3, PE3).

For packets which belong to an SRv6 VPN service associated with the SRv6 Service SID PE3:CL1.DT, the packet encapsulation and forwarding process is shown as below:

```
PE1 ->P1 : (PE1, P1)(PE3:CL1.DT, BR11; SL=2)(C-pkt)
P1 ->BR11: (PE1, BR11)(PE3:CL1.DT, BR11; SL=1)(C-pkt)
BR11->BR21: (PE1, PE3:CL1.DT)(C-pkt)
BR21->P2 : (PE1, P2)(PE3:CL1.DT, BR23; SL=2)(C-pkt)
P2 ->BR23: (PE1, BR23)(PE3:CL1.DT, BR23; SL=1)(C-pkt)
BR23->BR31: (PE1, PE3:CL1.DT)(C-pkt)
BR31->P3 : (PE1, P3)(PE3:CL1.DT, PE3; SL=2)(C-pkt)
P3 ->PE3 : (PE1, PE3)(PE3:CL1.DT, PE3; SL=1)(C-pkt)
```

In some network domains, SRv6 Flex-Algo may be used to provide intent aware intra-domain path. The encapsulation is similar to the case with SRv6 Policy.

3.2. CPR over MPLS Intra-Domain Paths

It is possible that some of the domains are still using MPLS based data plane. In these domains, A CPR route can be resolved over a color-aware intra-domain MPLS LSP. Such intra-domain MPLS LSP MAY be established using SR-MPLS Policy, SR-MPLS Flex-Algo or RSVP-TE.

The encapsulation and forwarding of SRv6 service packets over an intra-domain MPLS LSP is based on the MPLS mechanisms as defined in [\[RFC3031\]](#) [\[RFC3032\]](#) and [\[RFC8660\]](#).

For packets which belong to an SRv6 VPN service associated with the SRv6 Service SID PE3:CL1.DT, the packet encapsulation and forwarding process is shown as below:

```
PE1 ->P1 : (Label-stack for PE1 to BR11) (PE1, PE3:CL1.DT)(C-pkt)
P1 ->BR11: (Label-stack for PE1 to BR11) (PE1, PE3:CL1.DT)(C-pkt)
BR11->BR21: (PE1, PE3:CL1.DT)(C-pkt)
BR21->P2 : (Label-stack for BR21 to BR23)(PE1, PE3:CL1.DT)(C-pkt)
P2 ->BR23: (Label-stack for BR21 to BR23)(PE1, PE3:CL1.DT)(C-pkt)
BR23->BR31: (PE1, PE3:CL1.DT)(C-pkt)
BR31->P3 : (Label-stack for BR31 to PE3) (PE1, PE3:CL1.DT)(C-pkt)
P3 ->PE3 : (Label-stack for BR31 to PE3) (PE1, PE3:CL1.DT)(C-pkt)
```

4. Operational Considerations

Since the colorful locator prefixes are the sub-locators of the node's base SRv6 locator, the IPv6 unicast route of the base locator prefix is the covering prefix of all the colorful locator prefixes. To make sure the colorful locator prefixes can be distributed to the ingress PE nodes along the border nodes, it is required that route aggregation SHOULD be disabled for IPv6 unicast routes which carries the color extended community.

All the border nodes and the ingress PE nodes SHOULD install the colorful locator prefixes into the RIB and FIB. For transit domains which support the CPR mechanism, the border nodes SHOULD use the tuple (N, C) for the resolution of the CPR routes to intent aware intra-domain paths. For transit domains which do not support this mechanism, the border nodes MAY resolve the CPR routes over a best effort intra-domain path to the next-hop N, while the CPR route will be advertised further to the downstream domains with only the next-hop changed to itself. This allows the CPR routes to be resolved to intent aware intra-domain paths in the downstream domains which support the CPR mechanism.

5. IANA Considerations

This document makes no request of IANA.

6. Security Considerations

TBD

7. Acknowledgements

The authors would like to thank Shunwan Zhuang and Zhibo Hu for the review and discussion.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", DOI 10.17487/RFC2119, BCP 14, RFC 2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", DOI 10.17487/RFC8174, RFC 8174, BCP 14, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8754] Filshil, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", DOI 10.17487/RFC8754, RFC 8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filshil, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/

RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", DOI 10.17487/RFC9252, RFC 9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.

8.2. Informative 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 C. Bowers, "Problem statement for Inter-domain Intent-aware Routing using Color", Work in Progress, Internet-Draft, draft-hr-spring-intentaware-routing-using-color-00, 15 July 2022, <<https://www.ietf.org/archive/id/draft-hr-spring-intentaware-routing-using-color-00.txt>>.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", DOI 10.17487/RFC3031, RFC 3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", DOI 10.17487/RFC3032, RFC 3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.

[RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", DOI 10.17487/RFC8660, RFC 8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

Authors' Addresses

Haibo Wang
Huawei Technologies
China

Email: rainsword.wang@huawei.com

Jie Dong
Huawei Technologies
China

Email: jie.dong@huawei.com

Jingrong Xie
Huawei Technologies
China

Email: xiejingrong@huawei.com

Xinjun Chen
Huawei Technologies
China

Email: ifocus.chen@huawei.com