

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: 23 February 2026

K. Vairavakkalai, Ed.  
N. Venkataraman, Ed.  
Juniper Networks, Inc.  
22 August 2025

BGP CT - Adaptation to SRv6 dataplane  
draft-ietf-idr-bgp-ct-srv6-07

## Abstract

This document specifies how the mechanisms for "Intent Driven Service Mapping" defined in [BGP-CT] are applied to SRv6 dataplane. The extensions needed for SRv6 dataplane operations are specified. Base procedures of BGP CT are followed unaltered.

Illustration of how BGP CT procedures work in SRv6 dataplane is provided in this document.

## 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 RFC 2119 [RFC2119] RFC 8174 [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 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 23 February 2026.

## Copyright Notice

Copyright (c) 2025 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 . . . . .	2
2. Terminology . . . . .	3
2.1. Definitions . . . . .	4
3. NLRI and Nexthop Encoding . . . . .	4
4. SRv6 Encapsulation Information . . . . .	5
5. BGP CT in SRv6 networks . . . . .	5
5.1. SID stacking approach . . . . .	6
5.1.1. Egress Node Procedures . . . . .	9
5.1.2. Ingress Border Node Procedures . . . . .	9
5.1.3. Transit Border Node Procedures . . . . .	10
5.1.4. Transport Layer Procedures at Service Ingress Node . . . . .	11
5.1.5. Service Layer Procedures . . . . .	12
5.2. Color-encoded Service SID (CPR) Approach . . . . .	13
5.2.1. Analysis of CPR Approach . . . . .	14
6. Error Handling Considerations . . . . .	15
7. IANA Considerations . . . . .	15
8. Security Considerations . . . . .	15
9. References . . . . .	16
9.1. Normative References . . . . .	16
9.2. Informative References . . . . .	17
Contributors . . . . .	18
Co-Authors . . . . .	18
Other Contributors . . . . .	19
Acknowledgements . . . . .	20
Authors' Addresses . . . . .	20

## 1. Introduction

This document specifies how the mechanisms for "Intent Driven Service Mapping" defined in [BGP-CT] are applied to SRv6 dataplane. The extensions needed for SRv6 dataplane operations are specified. Base procedures of BGP CT are followed unaltered.

The BGP CT family (e.g. AFI/SAFI 2/76) is used to set up inter-domain tunnels satisfying a certain Transport Class, when using Segment Routing over IPv6 (SRv6) data plane on the inter-AS links or as an intra-AS tunneling mechanism. Illustration of how BGP CT procedures work in these scenarios is provided in this document.

## 2. Terminology

AFI: Address Family Identifier

AS: Autonomous System

BGP CT: BGP Classful Transport family (AFI/SAFIs 1/76, 2/76)

BN: Border Node

EP: Endpoint, e.g. a loopback address in the network

MPLS: Multi Protocol Label Switching

NLRI: Network Layer Reachability Information

PE: Provider Edge

RD: Route Distinguisher

RT: Route Target extended community

SAFI: Subsequent Address Family Identifier

SID: SR Segment Identifier

SLA: Service Level Agreement

SN: Service Node

SR: Segment Routing

SRTE: Segment Routing Traffic Engineering

TC: Transport Class

TC ID: Transport Class Identifier

VRF: Virtual Router Forwarding Table

## 2.1. Definitions

**Import processing:** Receive side processing of an overlay route, including things like import policy application, resolution scheme selection and nexthop resolution.

**Intent:** A set of operational goals (that a network should meet) and outcomes (that a network is supposed to deliver) defined in a declarative manner without specifying how to achieve or implement them, as defined in Section 2 of [RFC9315].

**Mapping Community:** Any BGP Community/Extended Community on a BGP route that maps to a Resolution Scheme. e.g., color:0:100, transport-target:0:100.

**Resolution Scheme:** A construct comprising of an ordered set of TRDBs to resolve next hop reachability, for realizing a desired intent.

**Service Family:** A BGP address family used for advertising routes for "data traffic" as opposed to tunnels (e.g. AFI/SAFIs 1/1 or 1/128).

**Transport Family:** A BGP address family used for advertising tunnels, which are in turn used by service routes for resolution (e.g. AFI/SAFIs 1/4 or 1/76).

**Transport Class:** A construct to group transport tunnels offering the same SLA.

**Transport Class RT:** A Route Target Extended Community used to identify a specific Transport Class.

**Transport Plane:** An end-to-end plane consisting of transport tunnels belonging to the same Transport Class. Tunnels of the same Transport Class are stitched together by BGP CT route readvertisements with next hop self to enable Label-Swap forwarding across domain boundaries.

**Transport Route Database (TRDB):** At the SN and BN, a Transport Class has an associated Transport Route Database that collects its tunnel ingress routes.

## 3. NLRI and Nexthop Encoding

The procedures for encoding a BGP Classful Transport (BGP CT) family route are specified in sections 4,5,6 and 7 of [BGP-CT]. These are followed, with the addition of SRv6 encapsulation information.

A BGP CT node that supports SRv6 forwarding encodes the SID information for SR with respect to SRv6 Data Plane as specified in Section 4.

A BGP CT node that does not support MPLS forwarding advertises the special label 3 (Implicit NULL) in the [RFC8277] MPLS Label field. The Implicit NULL label carried in BGP CT route indicates to receiving node that it should not impose any BGP CT label for this route. Thus a pure SRv6 node carries Implicit NULL in the MPLS Label field in RFC8277 BGP CT NLRI.

Aspects regarding Interoperability between nodes supporting different forwarding technologies is discussed in Section 6.3 and Section 11.3.2 of [BGP-CT].

#### 4. SRv6 Encapsulation Information

[RFC8986] specifies the SRv6 Endpoint behaviors (End USD, End.B6.Encaps). [SRV6-INTER-DOMAIN] specifies the SRv6 Endpoint behaviors (END.REPLACE, END.REPLACEB6 and END.DB6). These are leveraged for BGP CT routes with SRv6 data plane.

The BGP Classful Transport route update for SRv6 MUST include an attribute containing SRv6 SID information, with Transposition scheme disabled. The BGP Prefix-SID attribute as specified in [RFC9252] is used for this purpose today.

If the [RFC9252] Prefix-SID attribute also contains a "SRv6 SID Structure Sub-Sub-TLV", the Transposition Length MUST be set to 0 and Transposition Offset MUST be set to 0. This indicates nothing is transposed and that the entire SRv6 SID value is encoded in the "SRv6 SID Information Sub-TLV".

It should be noted that prefixes carried in BGP CT family are transport layer end-points, e.g. PE loopback addresses. Thus, the SRv6 SID carried in a BGP CT route is also a transport layer identifier.

For an illustration of BGP CT deployment in SRv6 networks, refer following section Section 5 .

#### 5. BGP CT in SRv6 networks

This section describes BGP CT deployment in SRv6 multi-domain network using Inter-AS Option C architecture.

### 5.1. SID stacking approach

This approach uses stacking of service SRv6 SID over transport SRv6 SID. Transport layer SIDs of types End, End.B6.Encaps defined in [RFC8986], and type END.REPLACE\* defined in [SRV6-INTER-DOMAIN] are carried in AFI/SAFI 2/76. Service SID is carried in a Service Family like AFI/SAFI 2/1 or AFI/SAFI 2/128.

In this approach, the number of Service SIDs required at the egress SN is equal to service functions (e.g. Prefix, VRF or Next hop) and the number of Transport SIDs are equal to the number of transport classes.

This section describes the procedures of this approach with an illustration using an example topology.




Figure 1: BGP CT in SRv6 Only Data plane

In the topology shown in Figure 1, there are two AS domains, AS1 and AS2. These are pure IPv6 domains, with no MPLS enabled. Inter-AS links between AS1 and AS2 are also enabled with IPv6 forwarding.

Intra-AS nodes in AS1 and AS2 speak IBGP CT (AFI: 2, SAFI: 76) and ISIS-SRv6 between them. The Inter-AS nodes ASBR1, ASBR2 speak EBGp CT (AFI: 2, SAFI:76) between them. Transport Classes Gold (100) and Bronze (200) are provisioned in all PEs and ASBRs. All BGP CT advertisements in this example carry a MPLS label value of 3 (Implicit NULL) in the NLRI encoding.

Reachability between PE1 and PE2 is formed using BGP CT family. Service families like IPv4 unicast (AFI: 1, SAFI: 1) and L3VPN (AFI: 2, SAFI: 128) is negotiated on multihop EBGp session between PE1 and PE2. These service routes carry service SID to identify service functions at the advertising PE, and mapping community to identify the desired Intent.

The following SRv6 locators are provisioned:

PE2-SRv6 : SRv6 Locator for PE2 best effort transport class

PE2-SRv6-gold-loc : SRv6 Locator for PE2 gold transport class

PE2-SRv6-bronze-loc : SRv6 Locator for PE2 bronze transport class

ASBR1-SRv6-loc : SRv6 Locator for ASBR1 best effort transport class

ASBR1-SRv6-gold-loc : SRv6 Locator for ASBR1 gold transport class

ASBR1-SRv6-bronze-loc : SRv6 Locator for ASBR1 bronze transport class

ASBR2-SRv6-loc : SRv6 Locator for ASBR2 best effort transport class

ASBR2-SRv6-gold-loc : SRv6 Locator for ASBR2 gold transport class

ASBR2-SRv6-bronze-loc : SRv6 Locator for ASBR2 bronze transport class

The following transport layer SRv6 End SIDs are provisioned or dynamically allocated on demand:

PE2-SRv6-gold : PE2 End SID from PE2-SRv6-gold-loc, for gold transport class.

PE2-SRv6-bronze : PE2 End SID from PE2-SRv6-bronze-loc, for bronze transport class.

ASBR2-SRv6-PE2-gold-Replace : at ASBR2 End.B6.Encaps SID for PE2, gold transport class.

ASBR2-SRv6-PE2-bronze-Replace : at ASBR2 End.B6.Encaps SID for PE2, bronze transport class.

ASBR1-SRv6-gold : ASBR1 End SID from ASBR1-SRv6-gold-loc, for gold transport class.

ASBR1-SRv6-PE2-gold-Replace : at ASBR1 End.REPLACE SID for PE2, gold transport class.

ASBR1-SRv6-bronze : ASBR1 End SID from ASBR1-SRv6-bronze-loc, for bronze transport class.

ASBR1-SRv6-PE2-bronze-Replace : at ASBR1 End.REPLACE SID for PE2, bronze transport class.

Architecturally, the forwarding semantic of End.REPLACE SID operation is similar to Label SWAP operation in MPLS data plane. When a route received with End SID (e.g. PE2-SRv6-gold or PE2-SRv6-bronze transport SIDs) is readvertised with next hop self, an IPv6 forwarding entry is emitted with a forwarding semantic of End.B6.Encaps operation, which means: Update IPv6 DA with Next Segment in SRH, and Encapsulate SRv6 SID corresponding to the correct transport class. This can be seen in IPv6 FIB of ASBR2 during "BGP CT processing at ASBR2" in the following illustration:

The following service layer SRv6 End.DT4 SIDs are provisioned:

PE2-SRv6-S1-DT4 : PE2 End.DT4 SID for service S1

The locators for the above provisioned SRv6 SIDs will be advertised via ISIS between Intra-AS nodes and the established SRv6 tunnel to the node's loopback will be installed into the corresponding TRDB based on color.

The SRv6 tunnel ingress routes are published in the Gold and Bronze TRDBs at ASBR2 as follows:

Gold TRDB routes at ASBR2

```
[ISIS SRv6] PE2-LPBK
  NH: Encap "Gold-SRv6-Tunnel-to-PE2" tunnel
```

```
[ISIS SRv6] PE2-SRv6-gold
  NH: Encap "Gold-SRv6-Tunnel-to-PE2" tunnel
```

Bronze TRDB routes at ASBR2

```
[ISIS SRv6] PE2-LPBK
  NH: Encap "Bronze-SRv6-Tunnel-to-PE2" tunnel
```

```
[ISIS SRv6] PE2-SRv6-bronze:
  NH: Encap "Bronze-SRv6-Tunnel-to-PE2" tunnel
```

ASBR2: IPv6 FIB for SRv6

```
[ISIS SRv6] PE2-SRv6-gold,
  NH: Encap "Gold-SRv6-Tunnel-to-PE2"
```

```
[ISIS SRv6] PE2-SRv6-bronze,
  NH: Encap "Bronze-SRv6-Tunnel-to-PE2"
```



The illustration below shows the following BGP CT operations for the 'gold' Transport Plane.

- BGP CT route origination,
- Import processing for the route,
- BGP CT route propagation, and
- Service routes resolving over BGP CT route.

Similar processing is followed for the 'bronze' Transport Plane routes as well.

#### 5.1.1. Egress Node Procedures

Firstly, PE2 originates BGP CT route for its transport layer endpoints like Loopback address with SRv6 SID information to ASBR2 as follows:

IBGP CT routes from PE2 to ASBR2

```
RD1:PE2-LPBK,  
  transport-target:0:100,  
  Prefix-SID: PE2-SRv6-gold  
  NH: PE2-LPBK
```

```
RD2:PE2-LPBK,  
  transport-target:0:200,  
  Prefix-SID: PE2-SRv6-bronze  
  NH: PE2-LPBK
```

PE2: IPv6 FIB for SRv6

```
[BGP CT] PE2-SRv6-S1-DT4  
  NH: Decap, Perform service S1
```

#### 5.1.2. Ingress Border Node Procedures

When ASBR2 receives the IBGP CT advertisement for gold route from PE2, it performs import processing and next hop resolution for the endpoint PE2-LPBK in the gold TRDB based on its transport-target:0:100. This would resolve over the ISIS-SRv6 route in gold TRDB and pick "Gold-SRv6-Tunnel-to-PE2" tunnel.

On successful resolution, a IPv6 transit route for ASBR2-SRv6-PE2-gold-replace/128 is installed in the global IPv6 FIB with "Gold-SRv6-Tunnel-to-PE2" tunnel as next hop, enabling SRv6 forwarding for gold SLA.

### 5.1.3. Transit Border Node Procedures

The BGP CT routes for RD1:PE2-LPBK are advertised by ASBR2 further towards ASBR1 via EBGp CT. During this readvertisement, the next hop is set to self, and SID is rewritten to ASBR2-SRv6-gold-Replace. The following illustration captures the advertisements from ASBR2 to ASBR1 as well as the IPv6 FIB in ASBR2.

EBGP CT routes from ASBR2 to ASBR1

```
RD1:PE2-LPBK,  
  transport-target:0:100,  
  Prefix-SID: ASBR2-SRv6-PE2-gold-Replace,  
  NH: ASBR2_InterAS_Link  
  
RD2:PE2-LPBK,  
  transport-target:0:200,  
  Prefix-SID: ASBR2-SRv6-PE2-bronze-Replace,  
  NH: ASBR2_InterAS_Link
```

ASBR2: IPv6 FIB for SRv6

```
[BGP CT] ASBR2-SRv6-PE2-gold-Replace  
  NH: UpdateIPv6DA(SRH.NextSegment),  
      Encap "Gold-SRv6-Tunnel-to-PE2"  
  
[BGP CT] ASBR2-SRv6-PE2-bronze-Replace  
  NH: UpdateIPv6DA(SRH.NextSegment),  
      Encap "Bronze-SRv6-Tunnel-to-PE2"
```

When ASBR1 receives this EBGp CT advertisement from ASBR2, an IPv6 route for ASBR1-SRv6-gold-Replace/128 is installed with a next hop of ASBR1\_InterAS\_Link in the global IPv6 FIB, enabling SRv6 forwarding for gold SLA. The BGP CT route for RD1:PE2-LPBK is further advertised to PE1 via IBGP CT, with next hop set to self, and SID rewritten to ASBR1-SRv6-gold-Replace.

IBGP CT routes from ASBR1 to PE1

```
RD1:PE2-LPBK,  
  transport-target:0:100,  
  Prefix-SID: ASBR1-SRv6-PE2-gold-Replace,  
  NH: ASBR1-LPBK  
  
RD2:PE2-LPBK,  
  transport-target:0:200,  
  Prefix-SID: ASBR1-SRv6-PE2-bronze-Replace,  
  NH: ASBR1-LPBK
```

ASBR1: IPv6 FIB for SRv6

```
[BGP CT] ASBR1-SRv6-PE2-gold-Replace,  
  NH: ASBR2_InterAS_Link  
  SID op: ReplaceSID(ASBR2-SRv6-PE2-gold-Replace)  
  
[BGP CT] ASBR1-SRv6-PE2-bronze-Replace,  
  NH: ASBR2_InterAS_Link  
  SID op: ReplaceSID(ASBR2-SRv6-PE2-bronze-Replace)
```

#### 5.1.4. Transport Layer Procedures at Service Ingress Node

When PE1 receives this IBGP CT advertisement from ASBR1, it resolves the next hop ASBR1-LPBK in the gold TRDB based on its transport-target:0:100. This would resolve over the ISIS-SRv6 route in gold TRDB and pick "Gold-SRv6-Tunnel-to-ASBR1".

This forms the end-to-end Gold SLA path from PE1 to PE2. The gold BGP CT route for PE2-LPBK is installed in gold TRDB, and can be used for resolving service route next hops. The Transport layer SIDs are replaced at each border node, which reduces the number of SID decaps required at the egress PE.

Gold TRDB routes at PE1

```
[BGP CT] PE2-LPBK,  
  NH: ASBR1-SRv6-gold  
  SID op: EncapSID(ASBR1-SRv6-PE2-gold-Replace)
```

Bronze TRDB routes at PE1

```
[BGP CT] PE2-LPBK,  
  NH: ASBR1-SRv6-bronze  
  SID op: EncapSID(ASBR1-SRv6-PE2-bronze-Replace)
```

PE1: IPv6 FIB for SRv6

```
[BGP CT] PE2-LPBK,  
  NH: ASBR1-SRv6-gold  
  SID op: EncapSID(ASBR1-SRv6-PE2-gold-Replace)
```

```
[BGP CT] PE2-LPBK,  
  NH: ASBR1-SRv6-bronze  
  SID op: EncapSID(ASBR1-SRv6-PE2-bronze-Replace)
```

```
[ISIS SRv6] ASBR1-SRv6-gold,  
  NH: Encap "Gold-SRv6-Tunnel-to-ASBR1"
```

```
[ISIS SRv6] ASBR1-SRv6-bronze,  
  NH: Encap "Bronze-SRv6-Tunnel-to-ASBR1"
```

#### 5.1.5. Service Layer Procedures

At PE1, service routes that are received with next hop as PE2-LPBK and Mapping Community as Color:0:100 indicating Gold SLA will use the Resolution Scheme associated with its Mapping Community to resolve over the PE2-LPBK CT route installed in the gold TRDB, and push the SRv6-gold SID stack to reach PE2.

Similarly, any service routes received with next hop as PE2-LPBK and Mapping Community as Color:0:200 indicating Bronze SLA will use the Resolution Scheme associated with its Mapping Community to resolve over the PE2-LPBK CT route installed in the bronze TRDB, and push the SRv6-bronze SID stack to reach PE2. This is shown as follows:

BGP Service routes advertisement from PE2 to PE1:

```
SVC_PFX1,  
  color:0:100,  
  Prefix-SID: PE2-SRv6-S1-DT4,  
  NH: PE2-LPBK
```

```
SVC_PFX2,  
  color:0:200,  
  Prefix-SID: PE2-SRv6-S1-DT4,  
  NH: PE2-LPBK
```

PE1: Service routes FIB

```
[BGP INET] SVC_PFX1, color:0:100  
  NH: EncapSID "PE2-SRv6-S1-DT4,  
    ASBR1-SRv6-gold-Replace,  
    Gold-SRv6-Tunnel-to-ASBR1(outer)"  
  
[BGP INET] SVC_PFX2, color:0:200  
  NH: EncapSID "PE2-SRv6-S1-DT4,  
    ASBR1-SRv6-bronze-Replace,  
    Bronze-SRv6-Tunnel-to-ASBR1(outer)"
```

The operational, scaling and convergence aspects of this approach are similar to the aspects of applying BGP CT procedures to the MPLS data plane.

## 5.2. Color-encoded Service SID (CPR) Approach

CPR is defined in the document: Colorful Prefix Routing for SRv6 based services [Colorful-Prefix-Routing-SRv6], and uses IPv6 Unicast (AFI/SAFI = 2/1) as a Transport Family. CPR mechanism does not use BGP CT (AFI/SAFI 2/76) address family.

CPR uses color encoded SRv6 service SIDs to determine the intent-aware transport paths for the service, without a separate transport SRv6 SID. It routes using "Colorful Prefix" locators in the transport layer, which are carried in the IPv6 Unicast BGP family.

A Next hop Resolution Scheme similar to that of BGP CT [BGP-CT] is used on IPv6 Unicast family to resolve "Colorful Prefix" locator routes that carry a mapping community to intent-aware paths in each domain.

By virtue of the CPR SID allocation scheme, the service SIDs inherit the Intent of the corresponding Colorful Prefix route just by performing longest prefix match in forwarding plane.

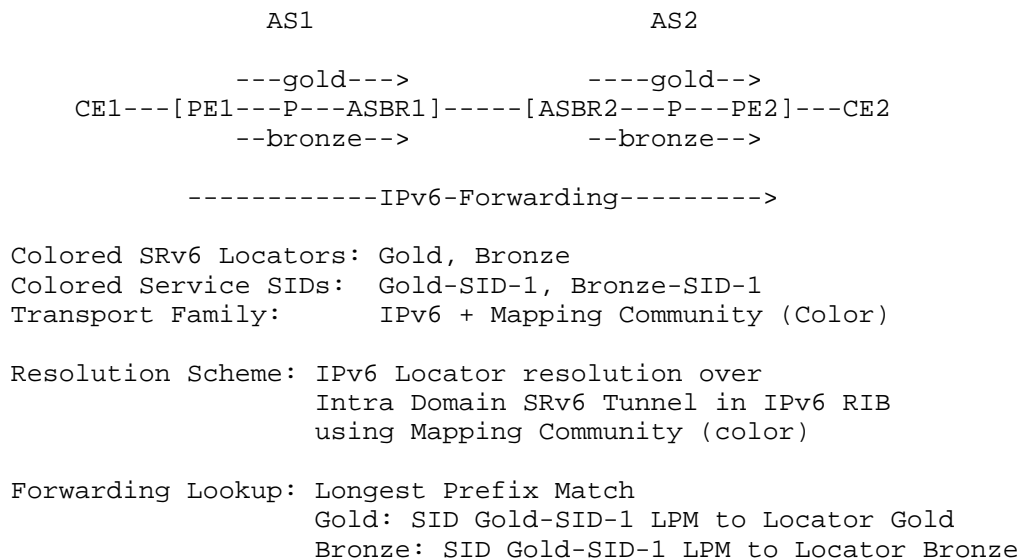


Figure 2: CT Interactions for CPR approach

#### 5.2.1. Analysis of CPR Approach

The CPR approach can be used to support intent driven routing while minimizing SRv6 encapsulation overhead, at the cost of careful SID numbering and planning. The state in the transport network is a function of total number of Colorful Prefixes.

In the CPR approach, typically one service SID is allocated for each service function (e.g. VRF) which is associated with a specific intent. In some special scenarios, for example, when different service routes in the same VRF are with different intents, a unique service SID would need to be allocated for each intent associated with the VRF.

However, the CPR mechanism preserves BGP PIC (Prefix scale Independent Convergence) for the egress SN failure scenario where only Colorful Prefix routes need to be withdrawn.

CPR achieves strict Intent based forwarding for the service routes. Fallback to best effort transport class is achieved by numbering all SRv6 Colorful Prefix locators at the egress SN to fall in the same subnet as the SRv6 locator that uses best effort transport class. Customized intent fallback between different color transport classes may be achieved by allocating a CPR prefix for each such intent fallback policy, and advertising that CPR prefix with an appropriate mapping community, that maps to a customized resolution scheme. Alternatively, the intent fallback policy may be provisioned on the ingress nodes directly.

Furthermore, IPv6 Unicast family is widely deployed to carry Internet Service routes. Repurposing IPv6 Unicast family to carry Transport routes also may impact the operational complexity and security aspects in the network.

## 6. Error Handling Considerations

This document follows error handling procedures defined in [BGP-CT], and extends it further.

If a BGP CT route is received with a [RFC9252] BGP Prefix-SID attribute containing a "SRv6 SID Information Sub-TLV", and also contains a "SRv6 SID Structure Sub-Sub-TLV", the Transposition Length is not set to 0 or Transposition Offset is not set to 0. This indicates transposition is in use, which is not expected on BGP CT route. Treat-as-withdraw approach from [RFC7606] is used to handle this error. The route is kept as Unusable, with appropriate diagnostic information, to aid troubleshooting.

If a BGP speaker considers a received BGP CT route invalid for some reason, but is able to successfully parse the NLRI and attributes, Treat-as-withdraw approach from [RFC7606] is used. The route is kept as Unusable, with appropriate diagnostic information, to aid troubleshooting.

The error handling procedures specified in Section 7, [RFC9252] apply for the BGP Prefix-SID attribute carried in BGP CT routes also.

## 7. IANA Considerations

This document makes no new requests of IANA.

## 8. Security Considerations

This document does not change the underlying security issues inherent in the existing BGP protocol, such as those described in [RFC4271] and [RFC4272].

This document follows the security considerations described in [BGP-CT]. As mentioned there, the "Walled Garden" approach is followed to carry routes for loopback addresses in BGP CT family (AFI/SAFI: 1/76 or 2/76). Thus mitigating the risk of unintended route escapes.

BGP CT routes distribute SRv6 SIDs for SRv6 dataplanes and hence security considerations of Section 9.3 of [RFC9252] apply. Moreover, SRv6 SID transposition scheme is disabled in BGP CT, as described in Section 4, to mitigate the risk of misinterpreting transposed SRv6 SID information as an MPLS label.

As [RFC4272] discusses, BGP is vulnerable to traffic-diversion attacks. This SAFI routes adds a new means by which an attacker could cause the traffic to be diverted from its normal path. Potential consequences include "hijacking" of traffic (insertion of an undesired node in the path, which allows for inspection or modification of traffic, or avoidance of security controls) or denial of service (directing traffic to a node that doesn't desire to receive it).

In order to mitigate the risk of the diversion of traffic from its intended destination, BGPsec solutions ([RFC8205] and Origin Validation [RFC8210]) may be extended in future to work for non-Internet SAFIs (SAFIs other than 1).

The restriction of the applicability of the BGP CT SAFI 76 to its intended well-defined scope limits the likelihood of traffic diversions. Furthermore, as long as the filtering and appropriate configuration mechanisms discussed previously are applied diligently, risk of the diversion of the traffic is eliminated.

## 9. References

### 9.1. Normative References

- [BGP-CT] Vairavakkalai, Ed. and Venkataraman, Ed., "BGP Classful Transport Planes", 25 April 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ct-33>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.



- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.
- [RFC8986] Filsfils, 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)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.
- [SRV6-INTER-DOMAIN]  
M., Ed., "SRv6 inter-domain mapping SIDs", 8 October 2024, <<https://datatracker.ietf.org/doc/html/draft-salih-spring-srv6-inter-domain-sids-07>>.

## 9.2. Informative References

- [Colorful-Prefix-Routing-SRv6]  
Wang, Ed., "BGP Colored Prefix Routing (CPR) for SRv6 based Services", 1 July 2024, <<https://www.ietf.org/archive/id/draft-ietf-idr-cpr-04.html>>.

- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/info/rfc8205>>.
- [RFC8210] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1", RFC 8210, DOI 10.17487/RFC8210, September 2017, <<https://www.rfc-editor.org/info/rfc8210>>.
- [RFC9315] Clemm, A., Ciavaglia, L., Granville, L. Z., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", RFC 9315, DOI 10.17487/RFC9315, October 2022, <<https://www.rfc-editor.org/info/rfc9315>>.

#### Contributors

#### Co-Authors

Reshma Das  
Juniper Networks, Inc.  
1133 Innovation Way,  
Sunnyvale, CA 94089  
United States of America  
Email: [dreshma@juniper.net](mailto:dreshma@juniper.net)

Israel Means  
AT&T  
2212 Avenida Mara,  
Chula Vista, California 91914  
United States of America  
Email: [israel.means@att.com](mailto:israel.means@att.com)

Csaba Mate  
KIFU, Hungarian NREN  
Budapest  
35 Vaci street,  
1134  
Hungary  
Email: [ietf@nop.hu](mailto:ietf@nop.hu)

Deepak J Gowda  
Extreme Networks  
55 Commerce Valley Drive West, Suite 300,  
Thornhill, Toronto, Ontario L3T 7V9  
Canada

Email: dgowda@extremenetworks.com

Other Contributors

Balaji Rajagopalan  
Juniper Networks, Inc.  
Electra, Exora Business Park~Marathahalli - Sarjapur Outer Ring Road,  
Bangalore 560103  
KA  
India  
Email: balajir@juniper.net

Rajesh M  
Juniper Networks, Inc.  
Electra, Exora Business Park~Marathahalli - Sarjapur Outer Ring Road,  
Bangalore 560103  
KA  
India  
Email: mrajesh@juniper.net

Chaitanya Yadlapalli  
AT&T  
200 S Laurel Ave,  
Middletown,, NJ 07748  
United States of America  
Email: cy098d@att.com

Mazen Khaddam  
Cox Communications Inc.  
Atlanta, GA  
United States of America  
Email: mazen.khaddam@cox.com

Rafal Jan Szarecki  
Google.  
1160 N Mathilda Ave, Bldg 5,  
Sunnyvale,, CA 94089  
United States of America  
Email: szarecki@google.com

Xiaohu Xu  
China Mobile  
Beijing  
China  
Email: xuxiaohu@cmss.chinamobile.com

#### Acknowledgements

The authors thank Jeff Haas, John Scudder, Susan Hares, Dongjie (Jimmy), Moses Nagarajah, Jeffrey (Zhaohui) Zhang, Joel Harper, Jingrong Xie, Mohamed Boucadair, Greg Skinner, Simon Leinen, Navaneetha Krishnan, Ravi M R, Chandrasekar Ramachandran, Shradha Hegde, Colby Barth, Vishnu Pavan Beeram, Sunil Malali, William J Britto, R Shilpa, Ashish Kumar (FE), Sunil Kumar Rawat, Abhishek Chakraborty, Richard Roberts, Krzysztof Szarkowicz, John E Drake, Srihari Sangli, Jim Uttaro, Luay Jalil, Keyur Patel, Ketan Talaulikar, Dhananjaya Rao, Swadesh Agarwal, Robert Raszuk, Ahmed Darwish, Aravind Srinivas Srinivasa Prabhakar, Moshiko Nayman, Chris Tripp, Gyan Mishra, Vijay Kestur, Santosh Kolenchery for all the valuable discussions, constructive criticisms, and review comments.

The decision to not reuse SAFI 128 and create a new address-family to carry these transport-routes was based on suggestion made by Richard Roberts and Krzysztof Szarkowicz.

#### Authors' Addresses

Kaliraj Vairavakkalai (editor)  
Juniper Networks, Inc.  
1133 Innovation Way,  
Sunnyvale, CA 94089  
United States of America  
Email: kaliraj@juniper.net

Natrajan Venkataraman (editor)  
Juniper Networks, Inc.  
1133 Innovation Way,  
Sunnyvale, CA 94089  
United States of America  
Email: natv@juniper.net