IDR WorkGroup
Internet-Draft
Intended status: Standards Track
Expires: 10 May 2026

S. Agrawal, Ed.
C. Filsfils
D. Rao
Cisco Systems
J. Dong
Huawei Technologies

6 November 2025

# BGP extensions for SRv6/MPLS Transport Interworking
draft-sa-idr-bgp-srv6-mpls-transport-iw-01

## Abstract

This document defines the BGP extensions required to provide transport interworking between SRv6 and MPLS in SRv6 deployment.

## Status of This Memo



## Copyright Notice

Table of Contents



# 1. Introduction

The deployment of SRv6 into existing transport network require SRv6 to interwork with MPLS. [I-D.ietf-spring-srv6-mpls-interworking] describes SRv6 and MPLS interworking architecture in multi domain network where each domain run SRv6 or MPLS data plane independently. To accomplish it, section 7.1.2 of [I-D.ietf-spring-srv6-mpls-interworking] details BGP inter-domain routing procedure in which domain border router set next hop to self when propagating transport routes (for example, locator or loopback prefix of a PE) across domains. Setting next hop to self results in allocation of label or SRv6 SID depending on dataplane type of the domain where route is propagated. Signaling of label is already specified in [RFC8277]. This document specifies method to signal SRv6 SID behaviors End.DTM, End.DTM46 and End.DPM specified in [I-D.ietf-spring-srv6-mpls-interworking] needed for SRv6 and MPLS interworking with BGP Labeled Unicast (BGP-LU) routes (SAFI value 4 routes as specified in [RFC8277]).

## 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. Signal SRv6 SID for BGP transport route

The Prefix-SID attribute is an optional, transitive path attribute defined in [RFC8669] to announce information about BGP Segment identifiers as list of TLVs. [RFC8669] defined TLV type 1 to encode label index and TLV type 3 to encode Originator SRGB to be attached to prefixes of BGP-LU SAFI([RFC8277]). [RFC9252] extended Prefix-SID attribute to carry SRv6 Segment Identifier that is encoded in SRv6 Service TLV and attached to BGP service routes.

### 2.1. SRv6 Transport TLV

This document defines a new TLV called "SRv6 Transport TLV" of the BGP Prefix-SID Attribute to announce SRv6 SID with prefixes of BGP-LU SAFI. SRv6 SID is encoded in "SRv6 Transport TLV" exactly like SRv6 Service SID TLV [RFC9252]. This document uses the TLV to signal End.DTM, End.DTM46 and End.DPM. The usage of this TLV for other SRv6 SID behaviors and SAFI’s is out of scope of this document and may be extended in future documents. Extension in this document lead to signaling of label in MPLS label field and SRv6 SID in SRv6 Transport TLV for BGP-LU route. Below section describes usage of both the information based on behavior of the SRv6 SID.

### 2.2. SRv6 encapsulation for MPLS transport

SRv6 SID is used to encapsulate the MPLS packet whose top label is the label value in MPLS field of the NLRI of the BGP-LU route. Draft [I-D.ietf-spring-srv6-mpls-interworking] defines End.DTM and End.DTM46 pseudo-code for this purpose.

Domain border router that set next hop to self, allocates label bound to each prefix of BGP-LU route. Label is encoded in MPLS label field as specified in [RFC8277]. SRv6 SID is encoded in "SRv6 Transport TLV" without transposition i.e. transposition length MUST be set to 0 in SRv6 SID Structure Sub-Sub-TLV. Behavior field MUST be set with SRv6 Endpoint Behavior codepoint value of End.DTM or End.DTM46. This behavior signals "AND" semantics i.e. push label signaled in NLRI and perform H.Encaps.M [I-D.ietf-spring-srv6-mpls-interworking] with DA as SRv6 SID signaled in the TLV. Same SRv6 SID can be attached to multiple BGP-LU routes by domain border router that sets next hop to self.

Section 7.1.2.2.1 of [I-D.ietf-spring-srv6-mpls-interworking] describe control and dataplane state using SRv6 Transport TLV.

## 2.3. SRv6 SID bound to prefix in NLRI

Domain border node that set next hop to self, allocates SRv6 SID and label for each prefix of BGP-LU SAFI. Draft [I-D.ietf-spring-srv6-mpls-interworking] defines End.DPM pseudo-code for this purpose. Label is encoded in MPLS label field as specified in [RFC8277]. SRv6 SID is encoded in SRv6 Transport TLV exactly like SRv6 Service SID TLV [RFC9252]. Behavior field MUST be set with SRv6 Endpoint Behavior codepoint value of End.DPM. Receiving node perform H.Encaps, where destination of IPv6 header is set to SRv6 SID for traffic destined to prefix in NLRI. Please refer to section 7.1.2.2.2 of [I-D.ietf-spring-srv6-mpls-interworking] for overall procedures when SRv6 SID is bound to prefix in NLRI.

Signaling of label and SRv6 SID assist in migration by allowing receiving node to select relevant encapsulation. Similarly, it allows to advertise just label to legacy node. For example, in Figure 1 node 4 being SRv6 capable may select relevant encapsulation and node 44 being legacy continue MPLS encapsulation.

```
          +-----+                +-----+  RD:V/v via 10   +-----+
   .......|S-RR1|<...............|S-RR2|<.................|S-RR3| <..
   :      +-----+                +-----+                   +-----+   :
   :                                                                 :
   :                                                                 :
+--:--------------------+----------------------+---------------------:-+
|  :                    |                      |                     : |
|  :                  +---+                    |                     : |
|  :                  | 4 |                    |                     : |
|  :     SAFI 4       +---+    SAFI 4          |       SAFI 4        : |
|  : <-E10,Label via 4 |  <-E10,label via 7    |   <-E10,0x3 via E10 : |
|----+                  |  SRv6 SID=B7:E10:: +---+               +----|
| E1 |                  |                    | 7 |               |E10 |
|----+                +---+ <-E10,label via 7+---+               +----|
|                     |44 |                    |                      |
|                     +---+                    |                      |
|                       |                      |                      |
|                       |                      |                      |
|        MPLS           |       SRv6/MPLS      |       MPLS           |
+-----------------------+----------------------+----------------------+
iPE                    iBR                    eBR                    ePE

<----------LI---------><----------C---------><-----------LE---------->
```

Figure 1: SRv6 SID bound to NLRI of SAFI 4

## 2.4. Propagation of transport routes with SRv6 SID

A BGP speaker that receive Prefix-SID attribute with SRv6 Transport TLV with transport route observe following rules when advertising the route to other peers:

* If the nexthop is unchanged, the TLVs, including any unrecognized Types of Sub-TLV and Sub-Sub-TLV, SHOULD be propagated further. In addition, all Reserved fields in the TLV or Sub-TLV or Sub-Sub-TLV MUST be propagated unchanged.

* If the nexthop is modified, the TLV and associated sub-TLVs/Sub-Sub-TLVs SHOULD be updated based on local policy. For example, if upstream is MPLS domain, then TLVs carrying SRv6 SID should be removed and local MPLS label bound to address in NLRI is propagated further.

# 3. IANA Considerations

## 3.1. BGP Prefix-SID TLV Types registry

This document introduce a new TLV Type of the BGP Prefix-SID attribute. IANA is requested to assign Type value in the registry "BGP Prefix-SID TLV Types" as follows

| Value | Type | Reference |
|---|---|---|
| TBD | SRv6 Transport TLV | <this document> |

# 4. Security Considerations

This document defines new TLV of Prefix-SID attribute attached to transport route. Tranport route is signaled among the nodes under the single administrative domain and security considerations of the attribute described in [RFC8669] continues to apply.

# 5. Contributors

Zafar Ali
Cisco Systems
Email: zali@cisco.com

Daniel Voyer
Cisco Systems
Email: davoyer@cisco.com

## 6. Acknowledgements

The authors would like to acknowledge Stephane Litkowski and Ketan Talaulikar for review and comments.

## 7. References

### 7.1. Normative References

[I-D.ietf-spring-srv6-mpls-interworking]
Agrawal, S., Filsfils, C., Voyer, D., Dawra, G., Li, Z., and S. Hegde, "SRv6 and MPLS interworking", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-mpls-interworking-01, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-mpls-interworking-01>.

[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>.

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

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, DOI 10.17487/RFC4023, March 2005, <https://www.rfc-editor.org/info/rfc4023>.

[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>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <https://www.rfc-editor.org/info/rfc4364>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <https://www.rfc-editor.org/info/rfc4760>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <https://www.rfc-editor.org/info/rfc7432>.

[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>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <https://www.rfc-editor.org/info/rfc8664>.

[RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <https://www.rfc-editor.org/info/rfc8669>.

[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>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <https://www.rfc-editor.org/info/rfc9256>.

## 7.2. Informative References

[I-D.ietf-mpls-seamless-mpls]
Leymann, N., Decraene, B., Filsfils, C., Konstantynowicz, M., and D. Steinberg, "Seamless MPLS Architecture", Work in Progress, Internet-Draft, draft-ietf-mpls-seamless-mpls-07, 28 June 2014, <https://datatracker.ietf.org/doc/html/draft-ietf-mpls-seamless-mpls-07>.

Authors’ Addresses

Swadesh Agrawal (editor)
Cisco Systems
Email: swaagraw@cisco.com

Clarence Filsfils
Cisco Systems
Email: cfilsfil@cisco.com

Dhananjaya Rao
Cisco Systems
Email: dhrao@cisco.com

Jie Dong
Huawei Technologies
Email: jie.dong@huawei.com