

Network Working Group
Internet-Draft
Updates: 6514, 7988 (if approved)
Intended status: Standards Track
Expires: 1 May 2026

R. Parekh, Ed.
Arrcus
D. Voyer, Ed.
C. Filsfils
Cisco Systems, Inc.
H. Bidgoli
Nokia
Z. Zhang
Juniper Networks
28 October 2025

Multicast and Ethernet VPN with Segment Routing P2MP and Ingress
Replication
draft-ietf-bess-mvpn-evpn-sr-p2mp-16

Abstract

A Point-to-Multipoint (P2MP) Tree in a Segment Routing domain carries traffic from a Root to a set of Leaves. This document specifies extensions to BGP encodings and procedures for P2MP trees and Ingress Replication used in BGP/MPLS IP VPNs and Ethernet VPNs in a Segment Routing domain.

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.

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 1 May 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	3
1.1. Terminology	4
2. SR P2MP P-tunnel	5
3. MVPN with SR	7
3.1. SRv6 Multicast Endpoint Behaviors	7
3.1.1. End.DTMC4: Decapsulation and Specific IPv4 Multicast Table Lookup	7
3.1.2. End.DTMC6: Decapsulation and Specific IPv6 Multicast Table Lookup	8
3.1.3. End.DTMC46: Decapsulation and Specific IP Multicast Table Lookup	8
3.2. MVPN with SR P2MP P-tunnel	8
3.2.1. PMSI Tunnel Attribute for SR P2MP P-tunnel	8
3.2.2. Auto-Discovery procedures	11
3.3. MVPN with Ingress Replication over SR	12
3.3.1. SR-MPLS	13
3.3.2. SRv6	13
4. EVPN with SR	15
4.1. EVPN with SR P2MP P-tunnel	15
4.1.1. PMSI Tunnel Attribute for SR P2MP P-tunnel	15
4.1.2. Split Horizon for ES Multihoming	17
4.1.3. Auto-Discovery procedures	18
4.2. EVPN with Ingress Replication over SR	19
4.2.1. SR-MPLS	19
4.2.2. SRv6	20
5. Implementation Status	20
5.1. Cisco Implementation	21
5.2. Nokia Implementation	21
6. IANA Considerations	21
7. Security Considerations	22
8. Acknowledgements	22
9. Contributors	22

10. References	23
10.1. Normative References	23
10.2. Informative References	25
Authors' Addresses	26

1. Introduction

Multicast in MPLS/BGP IP VPNs [RFC6513] and BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs [RFC6514] specify procedures that allow a network operator to provide Multicast VPN (MVPN) service to its customers. Multicast traffic from a customer is tunneled across the service provider network over Provider Tunnels (P-tunnels). P-tunnels can be instantiated using different transport mechanisms. For example, a service provider network that uses Segment Routing (SR) can use a SR Point-to-Multipoint (SR P2MP) tree defined by an SR P2MP Policy [I-D.ietf-pim-sr-p2mp-policy] or SR P2MP Ingress Replication (IR) to instantiate P-tunnels for MVPN. This document refers to a P-tunnel realized by an SR P2MP Policy as SR P2MP P-tunnel. SR P2MP P-tunnels can be instantiated both for SR-MPLS [RFC8660] and SRv6 [RFC8986][RFC8754].

An SR P2MP tree is defined by an SR P2MP Policy [I-D.ietf-pim-sr-p2mp-policy] and instantiated via a controller such as a Path Computation Element (PCE). An SR P2MP Policy consists of a Root, a set of Leaf Nodes and a set of candidate paths (CPs) with optional set of constraints and/or optimization objectives to be satisfied by the SR P2MP tree. A CP has zero or more P2MP tree instances (PTI).

This document specifies extensions to BGP auto-discovery procedures specified in [RFC6514] for P-tunnels constructed with a PTI. Use of Protocol Independent Multicast (PIM) for auto-discovery is outside scope of this document. Support for customer Bidirectional PIM (BIDIR-PIM) is outside the scope of this document.

This document extends procedures for MVPN service with IR over SR-MPLS [RFC7988] data plane with an SLA. New procedures are defined for MVPN service with IR over SRv6 data plane with or without an SLA.

This document defines new SRv6 endpoint behaviors [RFC8986] used for MVPN with SR P2MP and IR P-tunnels.

For BGP MPLS Ethernet VPN specified in [RFC7432] and extensions [RFC9572], P-tunnels are advertised for handling multi-destination traffic. These P-tunnels can be instantiated by SR-MPLS or SRv6 P2MP trees.

The reader is expected to be familiar with [RFC6513], [RFC6514] for MVPN procedures. For EVPN procedures should refer to [RFC7432].

Scalability, operational and troubleshooting considerations for SR P2MP trees and Replication segments are described in [I-D.ietf-pim-sr-p2mp-policy] and [RFC9524] respectively. Operations, Administration, and Maintenance (OAM) ping and traceroute procedures for SR P2MP Policy using SR-MPLS data plane are specified in [I-D.ietf-pim-p2mp-policy-ping].

1.1. Terminology

The following terms are used as defined in [RFC8402].

- * Segment Routing (SR)
- * Segment Identifier (SID)
- * SR domain
- * SR-MPLS
- * SRv6
- * SRv6 SID

The following terms are used as defined in [RFC8754].

- * Segment Routing Header (SRH)

The following terms are used as defined in [RFC9524].

- * Replication segment
- * Replication-SID
- * Root, Leaf and Bud node
- * Intermediate Replication node

The following terms are used as defined in [I-D.ietf-pim-sr-p2mp-policy].

- * SR P2MP Policy
- * Tree-ID
- * Candidate Path (CP)

- * P2MP Tree instance (PTI)

- * Tree-SID

For MVPN, the following terms are used as defined in [RFC6513] and [RFC6514].

- * P-Multicast Service Interface (PMSI)

- * P-tunnel

- * PMSI Tunnel Attribute (PTA)

- * Auto-Discovery (A-D) routes

- * Intra-AS I-PMSI A-D, Inter-AS I-PMSI A-D, S-PMSI A-D, and Leaf A-D routes

For EVPN, the following terms are used as defined in [RFC7432], [RFC9251] and [RFC9572].

- * EVPN Instance (EVI)

- * Ethernet Segment

- * Ethernet Segment Identifier (ESI)

- * Inclusive Multicast Ethernet Tag (IMET) route

- * Selective Multicast Ethernet Tag (SMET) route

- * Broadcast, Unknown Unicast and Multicast (BUM)

- * S-PMSI and Leaf A-D route

2. SR P2MP P-tunnel

For MVPN or EVPN, Provider Edge(PE) routers can steer customer traffic into a P-tunnel that can be instantiated by a SR-MPLS or SRv6 P2MP trees. An SR P2MP tree instance is by a Candidate path of an SR P2MP Policy [I-D.ietf-pim-sr-p2mp-policy]. A P-tunnel is signalled in MVPN or EVPN A-D routes using BGP PMSI Tunnel Attribute (PTA) [RFC6514]. The PTA identifies the P-tunnel that is used to instantiate a Provider Multicast Service Interface (PMSI).




Figure 1: MVPN/EVPN interaction with SR P2MP Policy

Above diagram show the interaction between the conceptual MVPN or EVPN module and SR P2MP Policy module on an ingress PE. The ingress PE participates in MVPN or EVPN autodiscovery procedures interacts with SR P2MP Policy module as shown above diagram to create and update SR P2MP Policy, Candidate path and the Leaf set of SR P2MP policies. The SR P2MP Policy module interacts with a controller using protocols such as PCEP [I-D.ietf-pce-sr-p2mp-policy], BGP, NetConf, etc., which are outside the scope of this document.

An ingress PE creates a CP of an SR P2MP Policy in the SR P2MP Policy module upon provisioning of the MVPN or EVPN module for SR P2MP P-tunnel, or based on events or policy, such as mapping MVPN or EVPN flow to a new SR P2MP P-tunnel. This CP can have optional traffic engineering constraints and/or optimization objective by provisioning or by a local policy. The SR P2MP Policy module signals the CP along with the constraints and optimization objective to the controller using a protocol mentioned above. The ingress PE deletes the CP of the SR P2MP Policy from the SR P2MP Policy module when the SR P2MP P-tunnel is not longer required. The SR P2MP Policy module signals deletion of the CP of the SR P2MP Policy to the controller.

An ingress PE participates in MVPN or EVPN auto-discovery procedures defined in this document to discover Leaf nodes of an SR P2MP P-tunnel, via various A-D routes. The MVPN or EVPN module on the ingress PE updates the Leaf set of the SR P2MP Policy corresponding to the P-tunnel in the SR P2MP Policy module. The SR P2MP Policy module signals the updated Leaf set to the controller using one of the protocols mentioned above.

An egress PE associates an MVPN or EVPN A-D route with MVPN or EVPN context and informs the SR P2MP Policy module to join the SR P2MP P-tunnel as a Leaf or a Bud node.

Given a Leaf set of an SR P2MP Policy and a CP with constraints and optimization objective, the controller computes and instantiates the PTI on the nodes that are part of the tree by stitching Replication segments [RFC9524] at Root (ingress PE) node, intermediate replication nodes and Leaf nodes (egress PEs). A Replication segment of an PTI can be instantiated by various methods such as BGP, PCEP, NetConf etc., which are outside the scope of this document. This PTI of the SR P2MP Policy instantiates the P-tunnel advertised by the MVPN or EVPN A-D routes.

The Tree-SID is the unique data plane identifier of the PTI. The Root node encapsulates the payload in Tree-SID to steer it into the PTI. The Provider (P) routers replicate the encapsulated payload, using Replication segments, towards the Leaf nodes of the PTI. The Leaf nodes of the PTI dispose the Tree-SID and deliver the payload. A Leaf node derives the MVPN or EVPN instance for delivering the payload either from the Tree-SID or other context encoded in the packet.

Note, an Ingress PE can deliver MVPN or EVPN payload to the egress PEs using IR over SR-MPLS or SRv6. The SR P2MP Policy module and controller do not participate in IR.

3. MVPN with SR

MVPN service can be provided using SR P2MP trees or IR over SR-MPLS or SRv6 data plane.

3.1. SRv6 Multicast Endpoint Behaviors

The following SRv6 Endpoint behaviors can be associated with the SRv6 Multicast Service SID used for MVPN with SR P2MP and IR P-tunnels.

3.1.1. End.DTMC4: Decapsulation and Specific IPv4 Multicast Table Lookup

The "Endpoint with Decapsulation and IPv4 Multicast Table Lookup" behavior (abbreviated End.DTMC4) is functionally identical to the End.DT4 behavior specified in [RFC8986], except that the forwarding lookup MUST be performed in the IPv4 multicast routing table rather than the unicast IPv4 routing table.

This behavior MUST only be associated with SRv6 Multicast Service SID carried in AFI/SAFI 1/129 (MVPN-IPv4) A-D routes.

3.1.2. End.DTMC6: Decapsulation and Specific IPv6 Multicast Table Lookup

The "Endpoint with Decapsulation and IPv6 Multicast Table Lookup" behavior (abbreviated End.DTMC6) is functionally identical to the End.DT6 behavior specified in [RFC8986], except that the forwarding lookup MUST be performed in the IPv6 multicast routing table rather than the unicast IPv6 routing table.

This behavior MUST only be associated with SRv6 Multicast Service SID carried in AFI/SAFI 2/129 (MVPN-IPv6) A-D routes.

3.1.3. End.DTMC46: Decapsulation and Specific IP Multicast Table Lookup

The "Endpoint with Decapsulation and IP Multicast Table Lookup" behavior (abbreviated End.DTMC46) is functionally identical to the End.DT4 and End.DT6 behaviors specified in [RFC8986], except that the forwarding lookup MUST be performed in the IP multicast routing table rather than in an IP unicast routing table.

This behavior MUST only be associated with SRv6 Multicast Service SID carried in AFI/SAFI 1/129 (MVPN-IPv4) or 1/129 (MVPN-IPv6) A-D routes.

3.2. MVPN with SR P2MP P-tunnel

[RFC6514] defines procedures for discovering PEs participating in a given MVPN and binding customer multicast flows to specific P-tunnels. This section specifies modifications to these procedures for SR P2MP P-tunnels. The Intra-AS I-PMSI, Inter-AS I-PMSI, and S-PMSI A-D routes have the PTA that specifies the SR P2MP tree P-tunnel. In this section, the term "SR P2MP" refers to both SR-MPLS and SRv6 data planes.

3.2.1. PMSI Tunnel Attribute for SR P2MP P-tunnel

A PTA for an SR P2MP P-tunnel is constructed as specified below.

- * Tunnel Type: One of below SR P2MP P-tunnel types from the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry additions defined in this document
 - 0x0C for SR-MPLS P2MP Tree
 - TBD for SRv6 P2MP Tree
- * Flags: See Section 3.2.2 for use of "Leaf Info Required bit".

- * MPLS Label: See Section 4.1.1.1
- * Tunnel Identifier: The SR P2MP P-tunnel identifies an SR P2MP Policy with the identifier <Tree-ID, Root> [I-D.ietf-pim-sr-p2mp-policy] in the order below,
 - Tree-ID: A 32-bit unsigned value that uniquely identifies an SR P2MP Policy at the Root.
 - Root: An IP address identifying the Root of the SR P2MP Policy. This can be either an IPv4 or IPv6 address. The address type can be inferred from the PTA length.

An PTA with SR P2MP P-tunnel is provisioned with Tunnel Type for SR-MPLS or SRv6 P2MP trees. An implementation SHOULD advertise PTA with SR P2MP P-tunnel only when the underlying SR-MPLS or SRv6 data plane is supported. The procedures for provisioning and determination of data plane support for SR-MPLS or SRv6 are outside scope of this document.

A P-tunnel can be segmented or non-segmented (see Section 8 of [RFC6513]). When a P-tunnel is non-segmented, the PTA is created by PE router at the Root of an SR P2MP tree. For segmented P-tunnels, each segment can be instantiated by a different technology. If a segment is instantiated using P2MP tree, the router at the root of an SR P2MP tree creates the PTA.

3.2.1.1. MPLS Label

When a SR P2MP P-tunnel is not shared across MVPNs i.e. there is one-to-one association between an MVPN instance and a SR P2MP P-tunnel, the "MPLS Label" field is set to zero as per [RFC6514] both for SR-MPLS and SRv6. In this case, the SR-MPLS or SRv6 Tree-SID of the PTI of the SR P2MP Policy advertised in the P-tunnel is sufficient to identify the MVPN instance for delivering the payload.

[RFC6514] allows a PE to aggregate two or more MVPNs onto one SR P2MP P-tunnel by advertising the same P-tunnel in PTA of A-D routes of different MVPNs. In this case, the "MPLS Label" field of PTA is filled to provide a context bound to a specific MVPN instance as described in below sections. The value in this field is used by egress PEs to identify the MVPN instance for delivering the payload encapsulated in SR-MPLS or SRv6.

3.2.1.1.1. SR-MPLS

When an SR P2MP P-tunnel is shared across two or more MVPNs in a SR-MPLS domain, the "MPLS Label" field of a PTA advertised in an A-D route MUST contain an upstream-assigned MPLS label [RFC5331] [RFC6513] or a label assigned from a global context such as "Domain-wide Common Block" (DCB) as specified in [RFC9573] that the advertising PE has bound to the MVPN,.

When an ingress PE steers the payload into a shared SR P2MP P-tunnel PTI, this MPLS label MUST be imposed before the MPLS label representing the Tree-SID. The trade-off of sharing an SR P2MP P-tunnel across MVPNs is two MPLS labels have to be imposed on ingress and disposed on egress.

The egress PEs of a shared SR P2MP P-tunnel use the MPLS label to determine the MVPN instance for delivering the payload.

3.2.1.1.2. SRv6

When an SR P2MP P-tunnel is shared across two or more MVPNs in a SRv6 domain [RFC8986], the "MPLS Label" field of a PTA advertised in an A-D route MUST contain an upstream-assigned SRv6 Multicast Service SID (Section 3.1) that the advertising PE has bound to the MVPN, or a SRv6 Multicast Service SID assigned from a global context; this follows same concept of "Domain-wide Common Block" (DCB) label as specified in [RFC9573]. The high order 20 bits of "MPLS Label" field carry the whole or a portion of the Function part of the SRv6 Multicast Service SID when Transposition Scheme of encoding as defined in [RFC9252] is used. When using the Transposition Scheme, the Transposition Length of SRv6 SID Structure Sub-Sub-TLV of SRv6 Prefix-SID attribute (see below) MUST be less than or equal to 20 and less than or equal to the Function Length. When Transposition scheme is not used, the label field MUST be set to zero as per [RFC6514] and Transposition Length MUST be zero.

The advertising ingress PE MUST attach a BGP Prefix-SID attribute [RFC8669] to Intra-AS I-PMSI, Inter-AS I-PMSI or S-PMSI A-D routes with SRv6 L3 Service TLV [RFC9252] to signal SRv6 Multicast Service SID. The SRv6 SID Information Sub-TLV carries the SRv6 Multicast Service SID in SRv6 SID Value field. The SRv6 Endpoint Behavior of the SRv6 SID Information Sub-TLV encodes one of End.DTMC4, End.DTMC6 or End.DTMC46 Section 3.1 code point values. The SRv6 SID Structure Sub-Sub-TLV encodes the structure of SRv6 Multicast Service SID. If Transposition scheme is used, the offset and length of SRv6 Multicast Endpoint function of SRv6 Multicast Service SID is set in Transposition Length and Transposition Offset fields of this sub-sub TLV. Otherwise, the Transposition Length and Offset fields MUST be

set to zero. The LOC of an SRv6 Multicast Service SID which is assigned from a global context, such as DCB, is outside the scope of this document.

The advertising ingress PE, which is the Root node of the shared SR P2MP P-tunnel, MUST encapsulate a payload in an outer IPv6 header with a SRH in which the SRv6 Multicast Service SID MUST be the last segment in the segment list (note the SRv6 Multicast Service SID may be the only segment in the SRH) . If Transposition scheme is used, ingress PE MUST merge Function in MPLS Label field of PTA with SRv6 SID in SID Information TLV using the Transposition Offset and Length fields from SID structure sub-sub TLV to create SRv6 Multicast Service SID.

An egress PE of the shared SR P2MP P-tunnel use the SRv6 Multicast Service SID in SRH to determine the MVPN instance in which the customer payload is to be delivered. The egress PE, in role of Leaf or Bud Node of Replication segment associated with shared SR P2MP P-tunnel tree, uses "look at next SID in SRH" [RFC9524] behavior to process the SRv6 Multicast Service SID. An egress PE MUST NOT install the SRv6 Multicast Service SID in its Forwarding Information Base (FIB) i.e. it MUST NOT forward packets based on the Locator portion of the SRv6 Multicast Service SID.

3.2.2. Auto-Discovery procedures

MVPN auto-discovery procedures specified in [RFC6514] are used to advertise an SR P2MP P-tunnel and discover the leaf nodes of the SR P2MP Policy associated with the P-tunnel. This section describes the processing of MVPN A-D routes to create, update and tear down an SR P2MP Policy of a SR P2MP P-tunnel as described in Section 2.

3.2.2.1. Creation of CP of SR P2MP Policy

A PE interacts with SR P2MP Policy module to create a CP, with optional traffic engineering constraints and optimization objective, of an SR P2MP Policy when it originates an Intra-AS I-PMSI A-D route [RFC6514], or an Inter-AS I-PMSI A-D route for intra-AS segment [RFC6514], or a S-PMSI A-D route [RFC6514], or "wildcard" S-PMSI A-D route [RFC6625], with a PTA having SR P2MP P-tunnel type.

The CP of the SR P2MP Policy associated with a SR P2MP P-tunnel is deleted from the SR P2MP Policy module when a PE withdraws the Intra-AS I-PMSI, Inter-AS I-PMSI or S-PMSI A-D route advertising that P-tunnel.

When a PE originates an Inter-AS I-PMSI or a S-PMSI A-D route with a PTA having SR P2MP P-tunnel type, it MUST set the "Leaf Information Required" flag in the PTA.

3.2.2.2. Discovery of Leaf nodes

An ingress PE that advertises an MVPN A-D route with PTA having SR P2MP P-tunnel type discovers a Leaf node of the SR P2MP Policy associated with the SR P2MP P-tunnel when it imports an Intra-AS I-PMSI A-D route [RFC6514], or Leaf A-D route [RFC6514] from an egress PE. The ingress PE adds the egress PE as a Leaf node in the SR P2MP Policy module.

An ingress PE removes an egress PE from the Leaf set of the SR P2MP Policy when the egress PE withdraws the Intra-AS I-PMSI or the Leaf A-D route.

An egress PE informs the SR P2MP Policy module to join the SR P2MP Policy as a Leaf or Bud node when it imports an Intra-AS I-PMSI A-D route, an Inter-AS I-PMSI A-D route, or a S-PMSI A-D route from an ingress PE having a PTA with SR P2MP P-tunnel type. The egress PE MUST originate a Leaf A-D route if the "Leaf Information Required" flag is set in the PTA.

An egress PE withdraws itself as a Leaf or Bud node of the SR P2MP Policy when the ingress PE withdraws a Intra-AS I-PMSI, Inter-AS I-PMSI, or a S-PMSI A-D route. The egress PE MUST withdraw the Leaf A-D it had originated earlier in response to the "Leaf Information Required" flag in the PTA.

3.3. MVPN with Ingress Replication over SR

A PE can provide MVPN service using IR over SR. The payload is encapsulated in SR-MPLS or SRv6 at an Ingress PE and replicated with each copy sent to an egress PE via unicast.

Ingress Replication Tunnels in Multicast VPN [RFC7988] specifies procedures that can be re-used to provide MVPN service with IR in an SR domain. A PE advertises Intra-AS I-PMSI A-D, Inter-AS I-PMSI A-D, or Selective PMSI A-D and Leaf A-D routes with PTA for IR. Egress PEs join as Leaf nodes using Intra-AS I-PMSI A-D or Leaf A-D routes. The procedures of [RFC7988] provide an MVPN IR service with best-effort unicast connectivity.

This document adds procedures for providing an MVPN IR service with a SLA from an ingress PE to an egress PE both for SR-MPLS and SRv6. This document extends BGP Update message of AFI/SAFI 1/129 (MVPN-IPv4) and 2/129 (MVPN-IPv6) to carry the Color Extended Community as

specified in [RFC9012] and Color-Only Type extension specified in Section 3 of [RFC9830]. An egress PE colors the Leaf or Intra-AS I-PMSI A-D route with a Color Extended Community. The ingress PE replicates MVPN customer payload to the egress PE by steering it into a SR-TE policy according to section 8 of [RFC9256]. The ingress PE encapsulates the payload packet into segment list of the matching SR-TE policy to the egress PE along with IR MPLS label or SRv6 Multicast Service SID received from the egress PE.

Note, the Color Extended Community is not used with MVPN SR P2MP P-tunnel. For MVPN with SR P2MP P-tunnel, the ingress PE dictates the traffic engineering treatment by specifying the constraints and the metric optimization in the CP of the SR P2MP policy corresponding to the P-tunnel on the ingress PE. This is necessary because packets are replicated in the PTI and therefore it is not possible to have differing traffic engineering treatment towards each egress PE (Leaf node) of the PTI.

3.3.1. SR-MPLS

The PTA carried in Intra-AS I-PMSI A-D, Inter-AS I-PMSI A-D, S-PMSI A-D and Leaf A-D routes is constructed as specified in [RFC7988].

MVPN IR service with SLA over SR-MPLS data plane can be provided by using the Color Extended Community as described above. Suppose an egress PE, say PE2, sends Leaf A-D route with Extended Color Community C1 with Color-Only Type 0 and IR label L10 to the ingress PE PE1. Assume the segment list of SR-TE policy (C1, PE2) at ingress PE1 is <L1, L2, L3>, PE1 will encapsulate MVPN payload into MPLS label stack <L1, L2, L3, L10> with L10 as BoS label.

3.3.2. SRv6

The procedures specified in [RFC7988], with the modifications defined in this section, are used to provide MVPN IR service over SRv6.

The PTA carried in Intra-AS I-PMSI A-D, Inter-AS I-PMSI A-D, Selective PMSI A-D and Leaf A-D routes is constructed as specified in [RFC7988] with modifications as below:

- * Tunnel Type: "Ingress Replication" as per [RFC6514].
- * MPLS Label: The high order 20 bits of this field carry the whole or a portion of the Function part of the SRv6 Multicast Service SID when the Transposition Scheme is used for encoding as defined in [RFC9252]. When using the Transposition Scheme, the Transposition Length of SRv6 SID Structure Sub-Sub-TLV of SRv6 Prefix-SID attribute (see below) MUST be less than or equal to 20

and less than or equal to the Function Length. When Transposition scheme is not used, the label field MUST be set to zero and Transposition Length MUST be zero.

Sections 6 and 7 of [RFC7988] describe considerations and procedures for allocating MPLS labels for IR P-tunnel. These considerations also apply to allocation of SRv6 Multicast Service SID for SRv6 IR.

To join a SRv6 IR P-tunnel advertised in PTA of Intra-AS I-PMSI A-D, Inter-AS I-PMSI A-D, or Selective S-PMSI A-D routes, an egress PE constructs a Leaf A-D or Intra-AS I-PMSI A-D route as described in [RFC7988] with modified PTA above. The egress PE MUST attach a BGP Prefix-SID attribute [RFC8669] with a Leaf A-D or Intra-AS I-PMSI A-D route with SRv6 L3 Service TLV [RFC9252] to signal SRv6 Multicast Service SID. The SRv6 SID Information Sub-TLV carries the SRv6 Multicast Service SID in SRv6 SID Value field. The SRv6 Endpoint Behavior of the SRv6 SID Information Sub-TLV MUST encode one of End.DTMC4, End.DTMC6 or End.DTMC46 Section 3.1 code point value. The SRv6 SID Structure Sub-Sub-TLV encodes the structure of SRv6 Multicast Service SID. If Transposition scheme is used, the offset and length of SRv6 Multicast Endpoint function of SRv6 Multicast Service SID is set in Transposition Length and Transposition Offset fields of this sub-sub TLV. Otherwise, the Transposition Length and Offset fields MUST be set to zero. The BGP Prefix SID attribute with SRv6 L3 Service TLV in Intra-AS I-PMSI or Leaf A-D route indicates to ingress PE that egress PE supports SRv6.

The SRv6 Multicast Service SID MUST be routable within the AS of the egress PE. As per [RFC7988], the Ingress PE uses the Tunnel Identifier of PTA to determine the unicast tunnel to use in order to send data to the egress PE. For SRv6 IR, the ingress PE MUST use the SRv6 Multicast Service SID to determine the unicast tunnel to be used. For best-effort MVPN IR service or SLA-based MVPN IR service using IGP Flexible Algorithm, the ingress PE MUST encapsulate the payload in an outer IPv6 header, with the SRv6 Multicast Service SID provided by the egress PE used as the destination address. If Transposition Scheme is used, ingress PE MUST merge Function in MPLS Label field of PTA with SRv6 SID in SID Information TLV using the Transposition Offset and Length fields from SID structure sub-sub TLV to create SRv6 Multicast Service SID.

MVPN IR service with SLA over SRv6 can be provided by using the Color Extended Community as described above. Suppose an egress PE, say PE2, sends Leaf A-D route with Extended color community C1 with Color-Only Type 0 and SRv6 Multicast Service SID S10 to the ingress PE PE1. Assume the segment list of the SR-TE policy (C1, PE2) at ingress PE1 is <S1, S2, S3>, PE1 will encapsulate a payload into an IPv6 header with SRH (PE1, S1) (S10, S3, S2; SL=3) (payload).

4. EVPN with SR

BGP MPLS Ethernet VPN specified in [RFC7432] specifies IMET route to support BUM traffic. This IMET route is the equivalent of MVPN Intra-AS I-PMSI route and is advertised with a PMSI Tunnel Attribute (PTA) as specified in [RFC6514] to advertise the inclusive P-tunnels. [RFC9252] specifies procedures for IMET routes with IR over SRv6.

[RFC9572] updates the EVPN BUM procedures to support selective P-tunnels. It defines new BGP route types which are advertised with a PTA, including the S-PMSI A-D and Leaf A-D route. The S-PMSI and Leaf A-D routes of [RFC9572] are analogous to MVPN S-PMSI and Leaf A-D routes [RFC6514]. Note, support of Inter-AS and Inter-Region segmentation procedures in [RFC9572] with SR P2MP P-tunnels are out of scope of this document.

Inclusive and Selective P-tunnels can be instantiated using SR P2MP trees or IR over SR.

4.1. EVPN with SR P2MP P-tunnel

This section specifies modifications to procedures of [RFC7432] and [RFC9572] for SR P2MP P-tunnels. The IMET, S-PMSI, and Leaf A-D routes have the PTA that specifies the SR P2MP tree P-tunnel. In this section, the term "SR P2MP" refers to both SR-MPLS and SRv6 data planes.

4.1.1. PMSI Tunnel Attribute for SR P2MP P-tunnel

A PTA for an SR P2MP P-tunnel is constructed as specified in Section 3.2.1 except the MPLS Label and Flags fields are filled as specified in Section 4.1.1.1 and Section 4.1.3 respectively.

4.1.1.1. MPLS Label

4.1.1.1.1. SR-MPLS

When a SR P2MP P-tunnel is not shared across MVPNs i.e. there is one-to-one association between a EVI and a SR P2MP P-tunnel, the "MPLS Label" field is set to zero as per [RFC6514]. In this case, the SR-MPLS Tree-SID of the PTI of the SR P2MP Policy advertised in the P-tunnel is sufficient to identify the MVPN instance for delivering the payload.

[RFC7432] and [RFC9572] allows a PE to aggregate two or more EVIss onto one SR P2MP P-tunnel by advertising the same P-tunnel in PTA of A-D routes of different EVIs. When an SR P2MP P-tunnel is shared across two or more EVIs in a SR-MPLS domain, the "MPLS Label" field

of a PTA advertised in an EVPN A-D route MUST contain an upstream-assigned MPLS label [RFC5331] [RFC6513] or a label assigned from a global context such as "Domain-wide Common Block" (DCB) as specified in [RFC9573] that the advertising PE has bound to the EVI. The egress PE uses this label as context to identify the specific EVI for delivering the EVPN payload encapsulated in SR-MPLS. When an ingress PE steers the payload into a shared SR P2MP P-tunnel PTI, this MPLS label MUST be imposed before the MPLS label representing the Tree-SID.

4.1.1.1.2. SRv6

For SRv6 P2MP, an EVI is identified by a SRv6 SID encoded in the SRH of an IPv6 encapsulated packet from the ingress PE as described later in this section. This is required even if SR P2MP P-tunnel is not shared across EVIs.

The advertising ingress PE MUST attach a BGP Prefix-SID attribute [RFC8669] to IMET [RFC9252] or S-PMSI [RFC9572] A-D routes with SRv6 L2 Service TLV [RFC9252]. The SRv6 SID Information Sub-TLV carries the SID in SRv6 SID Value field. The SRv6 Endpoint Behavior of the SRv6 SID Information Sub-TLV SHOULD be END.DT2M [RFC8986] as per Section 6.3 of [RFC9252]. The SRv6 SID Structure Sub-Sub-TLV encodes the structure of SID. If Transposition scheme is used, the offset and length of the SID is set in Transposition Length and Transposition Offset fields of this sub-sub TLV. Otherwise, the Transposition Length and Offset fields MUST be set to zero. The LOC of an SID which is assigned from a global context, such as DCB, is outside the scope of this document.

The "MPLS Label" field of a PTA advertised in an A-D route MUST contain an upstream-assigned SRv6 SID that the advertising PE has bound to the EVI, or a SRv6 SID assigned from a global context; this follows same concept of "Domain-wide Common Block" (DCB) label as specified in [RFC9573]. The high order 20 bits of "MPLS Label" field carry the whole or a portion of the Function part of the SRv6 SID when Transposition Scheme of encoding as defined in [RFC9252] is used. When using the Transposition Scheme, the Transposition Length of SRv6 SID Structure Sub-Sub-TLV of SRv6 Prefix-SID attribute MUST be less than or equal to 20 and less than or equal to the Function Length. When Transposition scheme is not used, the label field MUST be set to zero as per [RFC6514] and Transposition Length MUST be zero.

The advertising ingress PE, which is the Root node of the SR P2MP P-tunnel MUST encapsulate a payload in an outer IPv6 header with a SRH in which the SRv6 SID (advertised in the BGP Prefix-SID attribute) MUST be the last segment in the segment list (note the

SRv6 SID may be the only segment in the SRH) . If Transposition scheme is used, ingress PE MUST merge Function in MPLS Label field of PTA with SRv6 SID in SID Information TLV using the Transposition Offset and Length fields from SID structure sub-sub TLV to create the SRv6 SID. When ESI-based split horizon filtering is used, the Arg.FE2 (see Section 4.1.2.2) SHOULD be merged with the SRv6 SID by doing a bitwise logical OR operation to create the complete SRv6 SID encoded in the SRH on the ingress PE.

An egress PE of the SR P2MP P-tunnel use the SRv6 SID in SRH to determine the EVI in which the customer payload is to be delivered. The egress PE, in role of Leaf or Bud Node of Replication segment associated with SR P2MP P-tunnel tree, uses "look at next SID in SRH" [RFC9524] behavior to process the SRv6 SID. An egress PE MUST NOT install the SRv6 SID in its Forwarding Information Base (FIB) i.e. it MUST NOT forward packets based on the Locator portion of the SRv6 SID. The Arg.FE2 of the SID, if present, is used for ESI-based split horizon filtering as specified in Section 4.1.2.2.

4.1.2. Split Horizon for ES Multihoming

EVPN requires split-horizon filtering with ES multihoming in order to prevent duplicate BUM traffic as described in Section 8.3 of [RFC7432]. This section describes Split-horizon filtering procedures with SR P2MP P-tunnel.

4.1.2.1. SR-MPLS filtering

The split-horizon procedures specified in P2MP MPLS LSPs Section 8.3.1.2 of [RFC7432] are sufficient for SR P2MP P-tunnels. Note for shared SR P2MP P-tunnels, the ESI label is at the bottom of the label stack, followed by EVI context label and finally the Tree-SID label of the PTI associated with P-tunnel at the top in SR-MPLS encapsulation.

4.1.2.2. SRv6 filtering

For SRv6, "Arg.FE2" Argument of End.DT2M SRv6 Endpoint behavior [RFC8986] is used for ESI-based split-horizon filtering. For SR P2MP P-tunnel, the Arg.FE2 SID Argument is an upstream-assigned value assigned by an ingress PE and identifies an ES of origin. This SID Argument is analogous to the upstream-assigned ESI MPLS label specified in Section 8.3 of [RFC7432]. The Arg.FE2 is signalled in ESI Label field of the ESI Label extended community [RFC7432] and a BGP Prefix-SID attribute with a Ethernet A-D per ES route as specified in Section 6.1.1 of [RFC9252].

The advertised Arg.FE2 is encoded with SRv6 SID in SRH by the ingress PE as described above in Section 4.1.1.1.2. The egress PE uses the Arg.FE2 of the SRv6 SID, if present, to perform ESI-based split horizon filtering as specified in End.DT2M forwarding behavior in Section 4.12 of [RFC8986].

4.1.3. Auto-Discovery procedures

EVPN auto-discovery procedures specified in [RFC7432] and [RFC9572] are used to advertise an SR P2MP P-tunnel and discover the leaf nodes of the SR P2MP Policy associated with the P-tunnel. This section describes the processing of EVPN A-D routes to create, update and tear down an SR P2MP Policy of a SR P2MP P-tunnel as described in Section 2.

4.1.3.1. Creation of CP of SR P2MP Policy

A PE interacts with SR P2MP Policy module to create a CP, with optional traffic engineering constraints and optimization objective, of an SR P2MP Policy when it originates an IMET A-D route [RFC7432], or a S-PMSI A-D route [RFC9572], with a PTA having SR P2MP P-tunnel type.

The CP of the SR P2MP Policy associated with a SR P2MP P-tunnel is deleted from the SR P2MP Policy module when a PE withdraws the IMET or S-PMSI A-D route advertising that P-tunnel.

When a PE originates a S-PMSI A-D route with a PTA having SR P2MP P-tunnel type, it MUST set the "Leaf Information Required" flag in the PTA.

4.1.3.2. Discovery of Leaf nodes

An ingress PE that advertises an EVPN A-D route with PTA having SR P2MP P-tunnel type discovers a Leaf node of the SR P2MP Policy associated with the SR P2MP P-tunnel when it imports an IMET A-D route [RFC7432], or Leaf A-D route [RFC9572] from an egress PE. The ingress PE adds the egress PE as a Leaf node in the SR P2MP Policy module.

An ingress PE removes an egress PE from the Leaf set of the SR P2MP Policy when the egress PE withdraws the IMET A-D or the Leaf A-D route.

An egress PE informs the SR P2MP Policy module to join the SR P2MP Policy as a Leaf or Bud node when it imports an IMET A-D route or a S-PMSI A-D route from an ingress PE having a PTA with SR P2MP P-tunnel type. The egress PE MUST originate a Leaf A-D route if the "Leaf Information Required" flag is set in the PTA.

An egress PE withdraws itself as a Leaf or Bud node of the SR P2MP Policy when the ingress PE withdraws a IMET or a S-PMSI A-D route. The egress PE MUST withdraw the Leaf A-D it had originated earlier in response to the "Leaf Information Required" flag in the PTA.

4.2. EVPN with Ingress Replication over SR

A PE can provide EVPN service using IR over SR. The EVPN payload is encapsulated in SR-MPLS or SRv6 at an Ingress PE and replicated with each copy sent to an egress PE via unicast. IR procedures for MPLS are specified for IMET A-D route in [RFC7432] and SMET A-D route in [RFC9251]. Procedures for EVPN IR over SRv6 for IMET A-D route are specified in [RFC9252]. These procedures provide an EVPN IR service with best-effort unicast connectivity.

This document adds procedures for providing an EVPN IR service with a SLA from an ingress PE to an egress PE both for SR-MPLS and SRv6. This document extends BGP Update message of AFI/SAFI 25/70 (L2VPN-EVPN) to carry the Color Extended Community as specified in [RFC9012] and Color-Only Type extension specified in Section 3 of [RFC9830]. An egress PE colors the IMET with a Color Extended Community. The ingress PE replicates EVPN customer payload to the egress PE by steering it into a SR-TE policy according to section 8 of [RFC9256]. The ingress PE encapsulates the payload packet into segment list of the matching SR-TE policy to the egress PE along with IR MPLS label or SRv6 End.DT2M SID received from the egress PE.

Note, the Color Extended Community is not used with EVPN SR P2MP P-tunnel. For EVPN with SR P2MP P-tunnel, the ingress PE dictates the traffic engineering treatment by specifying the constraints and the metric optimization in the CP of the SR P2MP policy corresponding to the P-tunnel on the ingress PE. This is necessary because packets are replicated in the PTI and therefore it is not possible to have differing traffic engineering treatment towards each egress PE (Leaf node) of the PTI.

4.2.1. SR-MPLS

EVPN IR procedures for MPLS specified in [RFC7432] and [RFC9251] can be extended for EVPN IR over SR-MPLS.

EVPN IR service with SLA over SR-MPLS data plane can be provided by using the Color Extended Community as described above. Suppose an egress PE, say PE2, sends IMET A-D route with Extended Color Community C1 with Color-Only Type 0 and EVPN BUM label L10 to the ingress PE PE1. Assume the segment list of SR-TE policy (C1, PE2) at ingress PE1 is <L1, L2, L3>, PE1 will encapsulate MVPN payload into MPLS label stack <L1, L2, L3, L10> with L10 as BoS label. Note, the MPLS label stack may have ESI label at the bottom of label stack if ESI-based split-horizon is used.

4.2.2. SRv6

EVPN IR procedures for SRv6 are specified in Section 6.3 of [RFC9252] for IMET A-D route.

A PE can provide EVPN IR service with SLA over SRv6 using the Color Extended Community as described above. Suppose an egress PE, say PE2, sends IMET A-D route with Extended color community C1 with Color-Only Type 0 and SRv6 End.DT2M SID S10 to the ingress PE PE1. Assume the segment list of the SR-TE policy (C1, PE2) at ingress PE1 is <S1, S2, S3>, PE1 will encapsulate a payload into an IPv6 header with SRH (PE1, S1) (S10, S3, S2; SL=3) (payload). Note, the End.DT2M SID S10 may also have an Arg.FEs Argument if ESI-based split-horizon filter is used.

5. Implementation Status

Note to the RFC Editor: please remove this section before publication. This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in RFC7942. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist. According to RFC7942, "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

5.1. Cisco Implementation

Cisco has implemented MVPN procedures defined in this draft to provide MVPN service with SR P2MP policies in a segment routing domain. The implementation supports SR-MPLS encapsulation and has all the MUST and SHOULD clause in this draft. The implementation is at general availability maturity and is compliant with the latest version of the draft. The documentation for implementation can be found at Cisco website and the point of contact is abudhira@cisco.com.

5.2. Nokia Implementation

Nokia has implemented MVPN procedures specified in this draft to provide MVPN service with SR P2MP policies in a segment routing domain. The implementation supports SR-MPLS encapsulation and has all the MUST and SHOULD clause in this draft. The implementation is at general availability maturity and is compliant with the latest version of the draft. The documentation for implementation can be found at Nokia help and the point of contact is hooman.bidgoli@nokia.com.

6. IANA Considerations

IANA has assigned the value 0x0C for "SR-MPLS P2MP Tree" in the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry <https://www.iana.org/assignments/bgp-parameters/bgp-parameters.xhtml#pmsi-tunnel-types> [RFC7385] in the "Border Gateway Protocol (BGP) Parameters" registry.

IANA is requested to assign code point for "SRv6 P2MP Tree" in the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry <https://www.iana.org/assignments/bgp-parameters/bgp-parameters.xhtml#pmsi-tunnel-types> [RFC7385] in the "Border Gateway Protocol (BGP) Parameters" registry.

This document requests IANA to allocate the following code points in "SRv6 Endpoint Behaviors" sub-registry <https://www.iana.org/assignments/segment-routing/segment-routing.xhtml#srv6-endpoint-behaviors> [RFC8986] of "Segment Routing Parameters" top-level registry.

Value	Hex	Endpoint behavior	Reference	Change Controller
76	0x004C	End.DTMC4	[This.ID]	[Change to IETF]
77	0x004D	End.DTMC6	[This.ID]	[Change to IETF]
78	0x004E	End.DTMC46	[This.ID]	[Change to IETF]

Table 1: Multicast Service SRv6 Endpoint Behaviors

7. Security Considerations

The procedures in this document do not introduce any additional security considerations beyond those mentioned in [RFC6513], [RFC6514], [RFC7432] and [RFC9572]. For general security considerations applicable to SR P2MP Policy and Replication segments, please refer to [I-D.ietf-pim-sr-p2mp-policy] and [RFC9524] respectively.

8. Acknowledgements

The authors would like to acknowledge Luc Andre Burdet reviewing the document..

9. Contributors

Zafar Ali Cisco Systems, Inc. US

Email: zali@cisco.com

Arvind Venkateswaran Cisco Systems, Inc. US

Email: arvvenka@cisco.com

Jayant Kotalwar Nokia Mountain View US

Email: jayant.kotalwar@nokia.com

Tanmoy Kundu Nokia Mountain View US

Email: tanmoy.kundu@nokia.com

Clayton Hassen Bell Canada Vancouver CA

Email: clayton.hassen@bell.ca

Mankamana Mishra Cisco Systems, Inc. US

Email: mankamis@cisco.com

10. References

10.1. Normative References

- [I-D.ietf-pim-sr-p2mp-policy]
Parekh, R., Voyer, D., Filsfils, C., Bidgoli, H., and Z. J. Zhang, "Segment Routing Point-to-Multipoint Policy", Work in Progress, Internet-Draft, draft-ietf-pim-sr-p2mp-policy-22, 4 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pim-sr-p2mp-policy-22>>.
- [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>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [RFC6625] Rosen, E., Ed., Rekhter, Y., Ed., Hendrickx, W., and R. Qiu, "Wildcards in Multicast VPN Auto-Discovery Routes", RFC 6625, DOI 10.17487/RFC6625, May 2012, <<https://www.rfc-editor.org/info/rfc6625>>.
- [RFC7385] Andersson, L. and G. Swallow, "IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points", RFC 7385, DOI 10.17487/RFC7385, October 2014, <<https://www.rfc-editor.org/info/rfc7385>>.
- [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>>.

- [RFC7988] Rosen, E., Ed., Subramanian, K., and Z. Zhang, "Ingress Replication Tunnels in Multicast VPN", RFC 7988, DOI 10.17487/RFC7988, October 2016, <<https://www.rfc-editor.org/info/rfc7988>>.
- [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>>.
- [RFC8402] Filtsfils, 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>>.
- [RFC8660] Bashandy, A., Ed., Filtsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.
- [RFC8669] Previdi, S., Filtsfils, 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>>.
- [RFC8754] Filtsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filtsfils, 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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [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>>.

- [RFC9524] Voyer, D., Ed., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Replication for Multipoint Service Delivery", RFC 9524, DOI 10.17487/RFC9524, February 2024, <<https://www.rfc-editor.org/info/rfc9524>>.
- [RFC9572] Zhang, Z., Lin, W., Rabadan, J., Patel, K., and A. Sajassi, "Updates to EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Procedures", RFC 9572, DOI 10.17487/RFC9572, May 2024, <<https://www.rfc-editor.org/info/rfc9572>>.

10.2. Informative References

- [I-D.ietf-pce-sr-p2mp-policy] Bidgoli, H., Voyer, D., Budhiraja, A., Parekh, R., and S. Sivabalan, "PCEP extensions for SR P2MP Policy", Work in Progress, Internet-Draft, draft-ietf-pce-sr-p2mp-policy-13, 19 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-sr-p2mp-policy-13>>.
- [I-D.ietf-pim-p2mp-policy-ping] Bidgoli, H., Ali, Z., Zhang, Z. J., Budhiraja, A., and D. Voyer, "Segment Routing MPLS Point-to-Multipoint (P2MP) Policy Ping", Work in Progress, Internet-Draft, draft-ietf-pim-p2mp-policy-ping-25, 9 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pim-p2mp-policy-ping-25>>.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, DOI 10.17487/RFC5331, August 2008, <<https://www.rfc-editor.org/info/rfc5331>>.
- [RFC9251] Sajassi, A., Thoria, S., Mishra, M., Patel, K., Drake, J., and W. Lin, "Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)", RFC 9251, DOI 10.17487/RFC9251, June 2022, <<https://www.rfc-editor.org/info/rfc9251>>.
- [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>>.

- [RFC9573] Zhang, Z., Rosen, E., Lin, W., Li, Z., and IJ. Wijnands,
"MVPN/EVPN Tunnel Aggregation with Common Labels",
RFC 9573, DOI 10.17487/RFC9573, May 2024,
<<https://www.rfc-editor.org/info/rfc9573>>.
- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Ed., Mattes,
P., and D. Jain, "Advertising Segment Routing Policies in
BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025,
<<https://www.rfc-editor.org/info/rfc9830>>.

Authors' Addresses

Rishabh Parekh (editor)
Arrcus
United States of America
Email: rishabh@arrcus.com

Daniel Voyer (editor)
Cisco Systems, Inc.
Montreal
Canada
Email: davoyer@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium
Email: cfilsfil@cisco.com

Hooman Bidgoli
Nokia
Ottawa
Canada
Email: hooman.bidgoli@nokia.com

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net