Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 14 March 2026

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

Multicast and Ethernet VPN with Segment Routing P2MP and Ingress Replication

draft-ietf-bess-mvpn-evpn-sr-p2mp-15

Abstract

A Point-to-Multipoint (P2MP) Tree in a Segment Routing domain carries traffic from a Root to a set of Leaves. This document describes 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

Copyright Notice

Table of Contents

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 Service Provider 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 via different technologies. A service provider network that uses Segment Routing can use a Point-to-Multipoint (SR P2MP) tree [I-D.ietf-pim-sr-p2mp-policy] or P2MP Ingress Replication to instantiate P-Tunnels for MVPN. SR P2MP P-Tunnels can be instantiated both for SR-MPLS [RFC8660] and SRv6 [RFC8986][RFC8754].

In a Segment Routing network, a P2MP tree allows efficient delivery of traffic from a Root to set of Leaf nodes. A SR P2MP tree is defined by a SR P2MP Policy [I-D.ietf-pim-sr-p2mp-policy] and instantiated via a controller such as a Path Computation Element (PCE). A 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 P2MP tree. A CP has zero or more P2MP tree instances (PTI).

This document describes extensions to BGP Auto-Discovery procedures specified in [RFC6514] for P-Tunnels constructed with SR P2MP tree instances. Use of PIM for Auto-Discovery is outside scope of this document. Support for customer BIDIR-PIM is outside the scope of this document.

For BGP MPLS Ethernet VPN specified in [RFC7432] and extensions to this document, 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 concepts and terminology of [RFC6513], [RFC6514] for MVPN procedures and terms like P-tunnel, Intra-AS I-PMSI, S-PMSI and Leaf Auto-Discovery route types. For EVPN procedures and terms like Inclusive Multicast Ethernet Tag route and Broadcast, Unknown Unicast and Multicast (BUM) traffic, the reader should refer to[RFC7432]. The reader is expected to be familiar with [RFC9524] for terms like Replication segment, Replication- SID, Root node, Leaf node, Bud node and Intermediate

Replication node. The reader is also expected to be familiar with [I-D.ietf-pim-sr-p2mp-policy] for teerms like SR P2MP policy, Candidate paths, P2MP tree instance (PTI) and Tree-SID.

2. SR P2MP P-Tunnels

For MVPN or EVPN, Provider Edge(PE) routers steer customer traffic into a P-Tunnel that can be instantiated by a SR-MPLS or SRv6 P2MP trees. An SR P2MP tree is defined by a Candidate path of an SR P2MP policy [I-D.ietf-pim-sr-p2mp-policy].

An Ingress PE can deliver payload to egress PEs of the service using Ingress Replication. This payload is encapsulated in SR-MPLS or SRv6 and replicated to each egress PE.

Given a Candidate Path of a SR P2MP policy, a controller computes and instantiates the SR P2MP tree instance on the nodes that are part of the tree by stitching Replication segments [RFC9524] at Root, Leaf and intermediate replication nodes. Tree-SID is an unique identifier for the tree. A Replication segment of a SR P2MP tree can be instantiated by various methods such as BGP, PCEP, NetConf etc., which are outside the scope of this document.

PE routers use the MVPN or EVPN auto-discovery procedures in this document to create, update and delete SR P2MP Policies on the controllers using various methods such BGP, PCEP, NetConf etc., which are outside the scope of this document.

The Root of a P2MP tree imposes the Tree-SID to steer the payload into the SR P2MP tree instance. Provider (P) routers replicate the encapsulated payload, using Replication segments, towards the Leaf nodes of the P2MP tree. Leaf nodes of the P2MP tree deliver the payload after disposing the Tree-SID.

3. PMSI Tunnel Attribute for SR P2MP

BGP PMSI Tunnel Attribute (PTA) is defined in Section 5 [RFC6514] with format consisting of Flags (1 octet), Tunnel Type (1 octet), MPLS Label (3 octets) and Tunnel Identifier (variable length) fields. The PTA identifies the P-Tunnel that is used to instantiate a Provider Multicast Service Interface (PMSI). The PTA is carried in Intra-AS I-PMSI, Inter-AS I-PMSI, Selective PMSI, and Leaf Auto-Discovery routes.

A P2MP tree PTA is constructed as specified below.

* Tunnel Type: From the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry

   - 0x0C for SR-MPLS P2MP Tree

   - TBD for SRv6 P2MP Tree

* Flags: See Section 4 for use of "Leaf Info Required bit".

* MPLS Label: See Section 3.1

* Tunnel Identifier: The SR P2MP P-Tunnel is identified by <Tree-ID, Root> where,

   - Tree-ID is a 32-bit unsigned value that identifies a unique P2MP tree at a Root.

   - Root is an IP address identifying the Root of a P2MP tree. This can be either an IPv4 or IPv6 address. The address type can be inferred from the PTA length.

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 a 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 a P2MP tree creates the PTA.

3.1. MPLS Label

[RFC6514] allows a PE to aggregate two or more MVPNs onto one P-Tunnel by advertising the same P-Tunnel in PTA of Auto-Discovery routes of different MVPNs. This section specifies how the "MPLS Label" field of PTA is filled to provide a context bound to a specific MVPN. For EVPN considerations, see Section 7 section.

3.1.1. SR-MPLS

When a SR P2MP P-Tunnel is shared across two or more MVPNs in a SR MPLS domain [RFC8660], the "MPLS Label" field of a PTA advertised in an Auto-Discovery route MUST contain an upstream-assigned MPLS label [RFC5331] [RFC6513] that the advertising PE has bound to the MVPN, or a label assigned from a global context such as "Domain-wide Common Block" (DCB) as specified in [RFC9573].

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

3.1.2. SRv6

When a 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 Auto-Discovery route MUST contain an upstream-assigned SRv6 Multicast Service SID (Section 5.2.1) 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 and Transposition Length MUST be zero.

The advertising ingress PE attaches 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 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 locator (LOC) of SRv6 Multicast Service SID is the LOC of the advertising ingress PE. The advertising ingress PE MAY use a non-routable prefix as LOC of the SRv6 Multicast Service SID to prevent packets being routed to it based on the SID. 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.

The Egress PEs of a shared SR P2MP P-Tunnel use the SRv6 Multicast Service SID in SRH to determine the MVPN in which the customer payload is to be delivered. An 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 it’s Forwarding Information Base (FIB) i.e. it MUST NOT forward packets based on the Locator portion of the SRv6 Multicast Service SID.

4. MVPN Auto-Discovery and Binding Procedures for P2MP Trees

[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 tree P-Tunnels. In this section, the term "SR P2MP" refers to both SR-MPLS and SRv6 data planes.

4.1. Intra-AS I-PMSI

Intra-AS I-PMSI A-D routes are exchanged to discover PEs participating in a MVPN within an AS, or across different ASes when non-segmented P-Tunnels are used for inter-AS MVPNs.

4.1.1. Originating Intra-AS I-PMSI routes

RFC 6514 Section 9.1.1 (https://tools.ietf.org/html/rfc6514#section-9.1.1) describes procedures for originating Intra-AS I-PMSI A-D routes. For SR P2MP P-Tunnels, these procedures remain unchanged except as described in the following paragraphs.

When a PE originates an Intra-AS I-PMSI A-D route with a PTA having SR P2MP P-Tunnel Type, it MUST create a Candidate Path of SR P2MP policy on the controller. The Leaf nodes of P2MP tree are discovered using procedures described in Section 4.1.2.

For a PE in "Receiver Sites set", condition (c) in Section 9.1.1 (https://tools.ietf.org/html/rfc6514#section-9.1.1) is modified to include SR P2MP tree; such a PE MUST originate an Intra-AS I-PMSI A-D route when some PEs of the MVPN have VRFs that use SR P2MP tree but MUST NOT create a Candidate Path of SR P2MP policy as described above.

A PE MAY aggregate two or more Intra-AS I-PMSIs from different MVPNs onto the same SR P2MP P-Tunnel. When a PE withdraws the last Intra-AS I-PMSI A-D route, advertised with a PTA identifying a SR P2MP P-Tunnel , it MUST remove the Candidate Path of SR P2MP policy on the controller.

4.1.2. Receiving Intra-AS I-PMSI A-D routes

Procedure for receiving Intra-AS I-PMSI A-D routes, as described in RFC 6514 Section 9.1.2 (https://tools.ietf.org/html/rfc6514#section-9.1.2), remain unchanged for SR P2MP P-Tunnels except as described in the following paragraphs.

When a PE that advertises a SR P2MP P-Tunnel in the PTA of its Intra-AS I-PMSI A-D route, imports an Intra-AS I-PMSI A-D route from some PE, it MUST add that PE as a Leaf node to the SR P2MP Policy on the controller. The Originating IP Address of the Intra-AS i-PMSI A-D route is used as the Leaf Address. This procedure MUST also be followed for all Intra-AS I-PMSI routes that are already imported when the PE advertises a SR P2MP P-Tunnel in PTA of its Intra-AS I-PMSI A-D route.

A PE that imports and processes an Intra-AS I-PMSI A-D route from another PE with PTA having SR P2MP P-Tunnel MUST program the Tree-SID of the P2MP tree identified in the PTA of the route for disposition.

A PE MAY aggregate two or more Intra-AS I-PMSIs from different MVPNs onto the same SR P2MP P-Tunnel. When a remote PE withdraws an Intra-AS I-PMSI A-D route from a MVPN, and if this is the last MVPN sharing a SR P2MP P-Tunnel, a PE must remove the originating PE as a Leaf from the SR P2MP Policy on the controller.

4.2. Using S-PMSIs for binding customer flows to P2MP Segments

[RFC6514] specifies procedures for binding (C-S,C-G) customer flows to P-Tunnels using S-PMSI A-D routes. Wildcards in Multicast VPN Auto-Discovery Routes [RFC6625] specifies additional procedures to binding aggregate customer flows to P-Tunnels using "wildcard" S-PMSI A-D routes. This section describes modification to these procedures for SR P2MP P-Tunnels.

4.2.1. Originating S-PMSI A-D routes

RFC 6514 Section 12.1 (https://tools.ietf.org/html/rfc6514#section-12.1) describes procedures for originating S-PMSI A-D routes. For SR P2MP P-Tunnels, these procedures remain unchanged except as described in the following paragraphs.

When a PE originates S-PMSI A-D route with a PTA having SR P2MP P-Tunnel Type, it MUST set the "Leaf Info Required bit" in the PTA. The PE MUST create a Candidate Path of SR P2MP Policy on the controller.

The Leaf nodes of P2MP tree are discovered by Leaf A-D routes using procedures described in Section 4.4.2. When a PE originates S-PMSI A-D route with a PTA having SR P2MP P-Tunnel Type, it is possible the PE might have imported Leaf A-D routes whose route keys match the S-PMSI A-D route. The PE MUST re-apply procedures of Section 4.4.2 to these Leaf A-D routes.

A PE MAY aggregate two or more S-PMSIs onto the same SR P2MP P-Tunnel. When a PE withdraws the last S-PMSI A-D route, advertised with a PTA identifying a specific SR P2MP P-Tunnel , it MUST remove the the Candidate Path of SR P2MP Policy on the controller.

4.2.2. Receiving S-PMSI A-D routes

RFC 6514 Section 12.3 (https://tools.ietf.org/html/rfc6514#section-12.3) describes procedures for receiving S-PMSI A-D routes. For SR P2MP P-Tunnels, these procedures remain unchanged except as described in the following paragraphs.

The procedure for a PE to join SR P2MP P-Tunnel of S-PMSI A-D route by using a Leaf A-D route is described in Section 4.4.1. The PE MUST program the Tree-SID of the SR P2MP tree identified in the PTA of the route for disposition.

When a S-PMSI A-D route, whose SR P2MP P-Tunnel has been joined by a PE, is withdrawn, or when conditions (see RFC 6514 Section 12.3 (https://tools.ietf.org/html/rfc6514#section-12.3)) required to join that P-Tunnel are no longer satisfied, the PE MUST leave the P-Tunnel. The PE MUST withdraw the Leaf A-D route it had originated.

4.3. Inter-AS P-tunnels using P2MP Segments

A segmented inter-AS P-Tunnel consists of one or more intra-AS segments, one in each AS, connected by inter-AS segments between ASBRs of different ASes https://tools.ietf.org/html/rfc6514#section-9.2. These segments are constructed by PEs/ASBRs originating or re-advertising Inter-AS I-PMSI A-D routes. This section describes procedures for instantiating intra-AS segments using SR P2MP trees.

4.3.1. Advertising Inter-AS I-PMSI routes into iBGP

RFC 6514 Section 9.2.3.2 (https://tools.ietf.org/html/rfc6514#section-9.2.3.2) specifies procedures for advertising an Inter-AS I-PMSI A-D route to construct an intra-AS segment. The PTA of the route identifies the type and identifier of the P-Tunnel instantiating the intra-AS segment. The procedure for creating SR P2MP P-Tunnel for intra-AS segment are same as specified in Section 4.2.1 except that instead of S-PMSI A-D routes, the procedures apply to Inter-AS I-PMSI A-D routes.

4.3.2. Receiving Inter-AS I-PMSI A-D routes in iBGP

RFC 6514 Section 9.2.3.2 (https://tools.ietf.org/html/rfc6514#section-9.2.3.2) specifies procedures for processing an Inter-AS I-PMSI A-D route received via iBGP. If the PTA of the Inter-AS I-PMSI A-D route has SR P2MP P-Tunnel Type, the procedures are same as specified in Section 4.2.2 except that instead of S-PMSI A-D routes, the procedures apply to Inter-AS I-PMSI A-D routes. If the receiving router is an ASBR, the Tree-SID is stitched to the inter-AS segments to ASBRs in other ASes.

4.4. Leaf A-D routes for P2MP Segment Leaf Discovery

This section describes procedures for originating and processing Leaf A-D routes used for Leaf discovery of SR P2MP trees.

4.4.1. Originating Leaf A-D routes

The procedures for originating Leaf A-D route in response to receiving a S-PMSI or Inter-AS I-PMSI A-D route with PTA having SR P2MP P-Tunnel Type are same as specified in RFC 6514 Section 9.2.3.4.1 (https://tools.ietf.org/html/rfc6514#section-9.2.3.4.1).

4.4.2. Receiving Leaf A-D routes

Procedures for processing a received Leaf A-D route are specified in RFC 6514 Section 9.2.3.5 (https://tools.ietf.org/html/rfc6514#section-9.2.3.5). These procedures remain unchanged for discovering Leaf nodes of SR P2MP Policy except for considerations described in following paragraphs. These procedures apply to Leaf A-D routes received in response to both S-PMSI and Inter-AS I-PMSI A-D routes, shortened to "A-D routes" in this section

A Root PE/ASBR MAY use the same SR P2MP P-Tunnel in PTA of two or more A-D routes. For such aggregated P2MP trees, the PE/ASBR may receive multiple Leaf A-D routes from a Leaf PE. The P2MP tree for which a Leaf A-D is received can be identified by examining the P2MP tunnel Identifier in the PTA of A-D route that matches "Route Key" field [RFC6514] of the Leaf A-D route. When the PE receives the first Leaf A-D route from a Leaf PE, identified by the Originating Router’s IP address field, it MUST add that PE as Leaf of the SR P2MP Policy on the controller.

When a Leaf PE withdraws the last Leaf A-D route for a given SR P2MP P-Tunnel, the Root PE MUST remove the Leaf PE node from the Leaf node set of the SR P2MP Policy on the controller. Note that Root PE MAY remove the Candidate path of the SR P2MP Policy from the controller, before the last Leaf A-D is withdrawn. In this case, the Root PE MAY not need to remove the Leaf PE node from Leaf node set of the SR P2MP Policy on the controller.

5. MVPN with Ingress Replication over Segment Routing

A PE can provide MVPN service using Ingress Replication (IR) over Segment Routing. The payload is encapsulated in SR-MPLS or SRv6 at an Ingress PE. The encapsulated payload is replicated and a unicast copy is sent to each egress PE.

Ingress Replication Tunnels in Multicast VPN [RFC7988] specifies procedures that can be used to provide MVPN service with IR in a Segment Routing domain. A PE advertises Intra-AS I-PMSI A-D, Inter-AS I-PMSI A-D, Selective PMSI A-D and Leaf A-D routes with PTA for Ingress Replication. Egress PEs join as Leaf nodes using Intra-AS I-PMSI A-D or Leaf A-D routes.

RFC 7988 procedures allow an ingress PE to deliver MVPN traffic to egress PEs using best-effort unicast connectivity. For MVPN service with a SLA from ingress PE to an egress PE, the egress PE colors the Leaf Auto-Discovery route with a Color Extended Community as specified in [I-D.ietf-idr-sr-policy-safi]. The ingress PE replicates MVPN customer payload to that egress PE by steering traffic into a SR-TE policy (Color, egress PE) according to section 8 of [RFC9256].

5.1. SR-MPLS

Procedures of [RFC7988] are sufficient to create a SR-MPLS Ingress Replication for MVPN service without a SLA.

If an egress PE colors the Leaf A-D route with Color Extended Community, the ingress PE encapsulates the payload packet into segment list of (Color, egress PE) SR-TE policy along with Ingress Replication (IR) label received from the egress PE. Suppose the egress PE, say PE2, sends Leaf A-D route with extended color community C1 and IR label L10. 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.

5.2. SRv6

The procedures specified in [RFC7988], with the modifications defined in this section, MUST be used to construct an SRv6 Ingress Replication tree for MVPN service.

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 RFC 7988 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 ingress replication is used with the Transposition Scheme of 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.

Section 6 and 7 of RFC 7988 (https://datatracker.ietf.org/doc/html/rfc7988#section-6) describe considerations and procedures for allocating MPLS labels for IR P-Tunnel. These sections of [RFC7988] 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 RFC 7988 with modified PTA above. The egress PE attaches 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 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 RFC 7988, 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 both best-effort MVPN service and SLA-based MVPN 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.

If an egress PE colors a Leaf A-D route with Color Extended Community, the ingress PE SHOULD encapsulate the payload packet into outer IPv6 header with segment list of (Color, egress PE) SR-TE policy along with SRv6 Multicast Service SID received with Leaf A-D route from the egress PE using SRH. Suppose the egress PE, say PE2, sends Leaf A-D route with extended color community C1 and SRv6 Multicast Service SID S10. Assume the segment list of 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).

5.2.1. SRv6 Multicast Endpoint Behaviors

The following behaviors can be associated with SRv6 Multicast Service SID.

5.2.1.1. 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.

5.2.1.2. 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.

5.2.1.3. 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.

6. Dampening of MVPN routes

When P2MP trees are used as P-Tunnels for S-PMSI A-D routes, change in group membership of receivers connected to PEs has direct impact on the Leaf node set of a P2MP tree. If group membership changes frequently for a large number of groups with a lot of receivers across sites connected to different PEs, it can have an impact on the interaction between PEs and the controller.

Since Leaf A-D routes are used to discover Leaf PE of a P2MP tree, PEs SHOULD dampen Leaf A-D routes as described in Section 6.1 of RFC 7899 [RFC7899]. PEs MAY also implement procedures for damping other Auto-Discovery and BGP C-multicast Source Tree Join and Shared Tree Join routes as described in [RFC7899].

7. SR P2MP Trees for EVPN

BGP MPLS Ethernet VPN specified in [RFC7432] specifies Inclusive Multicast Ethernet Tag route to support Broadcast, Unknown Unicast and Multicast (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.

[RFC9572] updates the EVPN Broadcast, Unknown unicast, and Multicast (BUM) procedures to support selective P-tunnels and P-tunnel segmentation. That document defines new BGP route types that MUST be advertised with a PMSI Tunnel Attribute (PTA), including the Selective PMSI (S-PMSI) Auto-Discovery route.

Inclusive and Selective P-tunnels MAY be instantiated using SR P2MP trees. As with all other types of P2MP P-tunnels, the Ethernet Segment Identifier (ESI) label used for split-horizon MUST either be:

1. upstream-assigned by the PE that advertises the corresponding IMET or S-PMSI route, or

2. allocated from a global context, such as the Domain-wide Common Block (DCB), as specified in [RFC9573].

[RFC9625] specifies procedures to support Inter-Subnet Multicast. [I-D.ietf-bess-evpn-mvpn-seamless-interop] specifies how MVPN SAFI routes can be used to support Inter-Subnet Multicast. The P-Tunnels advertised in PTA of EVPN and MVPN routes as specified in these documents respectively MAY be instantiated by SR P2MP trees.

SRv6 P2MP trees can serve as an underlay multicast as described in RFC 8293 Section 3.4 (https://tools.ietf.org/html/rfc8293#section-3.4). A NVE encapsulates a tenant packet in an SRv6 header and deliver it over SRv6 P2MP trees to other NVEs.

SRv6 P2MP trees MAY be used as an underlay multicast mechanism, as described in RFC 8293 Section 3.4 (https://tools.ietf.org/html/rfc8293#section-3.4). In this case, a Network Virtualization Edge (NVE) MUST encapsulate tenant packets in an SRv6 header and deliver them over SRv6 P2MP trees to other NVEs.

EVPN ingress PEs discover the Leaf nodes of SR P2MP trees based on IMET routes or Leaf A-D routes (in response to S-PMSI A-D routes). The ingress PEs update the Leaf node set of SR P2MP policies on the controller based on this auto discovery. The controller then instantiates SR P2MP tree instances in SR domain to carry EVPN traffic mapped to the SR P2MP P-tunnels.

8. 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".

8.1. 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.

8.2. 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.

9. 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 of "Segment Routing Parameters" top-level registry.

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

Table 1: IETF - SRv6 Endpoint Behaviors

10. Security Considerations

The procedures in this document do not introduce any additional security considerations beyond those mentioned in [RFC6513] and [RFC6514]. 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.

11. Acknowledgements

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

12. 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 CanadaVancouver CA

Email: clayton.hassen@bell.ca

Mankamana Mishra Cisco Systems, Inc. US

Email: mankamis@cisco.com

13. References

13.1. 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>.

[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>.

[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>.

[RFC8660] Bashandy, A., Ed., Filsfils, 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., 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>.

[RFC8754] Filsfils, 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] 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>.

[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>.

13.2. Informative References

[I-D.ietf-bess-evpn-mvpn-seamless-interop]
Sajassi, A., Thiruvenkatasamy, K., Thoria, S., Gupta, A., and L. Jalil, "Seamless Multicast Interoperability between EVPN and MVPN PEs", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-mvpn-seamless-interop-08, 2 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-mvpn-seamless-interop-08>.

[I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>.

[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>.

[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>.

[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>.

[RFC7899] Morin, T., Ed., Litkowski, S., Patel, K., Zhang, Z., Kebler, R., and J. Haas, "Multicast VPN State Damping", RFC 7899, DOI 10.17487/RFC7899, June 2016, <https://www.rfc-editor.org/info/rfc7899>.

[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>.

[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>.

[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>.

[RFC9625] Lin, W., Zhang, Z., Drake, J., Rosen, E., Ed., Rabadan, J., and A. Sajassi, "EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding", RFC 9625, DOI 10.17487/RFC9625, August 2024, <https://www.rfc-editor.org/info/rfc9625>.

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