

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 7 January 2026

M. Mishra
Cisco Systems
Z. Zhang
Juniper Networks
H. Bidgoli
Nokia
6 July 2025

Multicast over SRv6 networks
draft-mankamana-pim-multicast-over-srv6-00

Abstract

This document presents solutions for deploying multicast in SRv6 networks. It explores the use of the native IPv6 multicast data plane for multicast distribution. The document discusses distributed control plane mechanisms, including PIM, and its integration with IGP Flex-Algo to optimize multicast delivery. The document also addresses overlay multicast solutions for both the Global Table Multicast (GTM) and Multicast VPNs (MVPNs), utilizing IP-in-IPv6 encapsulation without requiring additional shim layers.

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 7 January 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document.

Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. IPv6 Data plane	3
3.1. Multicast forwarding state	4
3.2. Route instantiation	4
4. Distributed Control Plane	4
4.1. PIM-SM	4
4.2. PIM-SM with Flex-Algo	5
5. Centralized Control plane	5
6. Virtualization and Overlay Solutions	6
7. Data plane IP in IP Encapsulations	7
8. IANA Considerations	8
9. Security Considerations	8
10. Acknowledgement	8
11. Contributors	8
12. Normative References	8
Authors' Addresses	9

1. Introduction

Segment Routing, as defined in [RFC8402], is a source-routing paradigm that simplifies network operations and enhances scalability. Segment routing is defined for unicast. The application of the source-route concept to Multicast is declared out-of-scope in [RFC8402].

Since multicast distribution is orthogonal to segment routing, classic multicast solutions are leveraged to provide multicast services in a segment routing network. In SR-MPLS networks, operators traditionally leveraged solutions based on IPv4 or MPLS data planes for multicast distribution, as both data planes were available in these networks.

However, with the advent of SRv6, network use a unifying IPv6 data plane end-to-end, replacing the need for MPLS or IPv4. While migrating the networks to SRv6, multicast services can still use the MPLS or IPv4 data planes, as ships-in-the-night. But the target is an IPv6-only SRv6 network.

This document focuses on deploying multicast in IPv6-only networks and describes possible solutions. Key Topics Addressed:

- * Leveraging IPv6's native multicast capabilities for SRv6 networks.
- * Distributed and centralized control planes for multicast.
- * Overlay solutions for supporting multicast in IPv6-only networks.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] .

- * PIM: Protocol Independent Multicast
- * PIM-SM: PIM Sparse Mode
- * PIM-SSM: PIM Source Specific Multicast
- * MVPN: Multicast Virtual Private Network
- * MBGP: Multi-protocol Border Gateway Protocol
- * IGMP: Internet Group Management Protocol
- * MLD: Multicast Listener Discovery
- * I-PMSI : Inclusive P-Multicast Service Interface
- * S-PMSI: Selective P-Multicast Service Interface
- * P(S,G) : Multicast flow (S,G) in provider network
- * C(S,G) : Multicast flow (S,G) in customer network
- * All the terminology from [RFC6514] is also applicable in this draft.

3. IPv6 Data plane

Segment Routing over IPv6 micro-segments (SRv6 uSID) enables end-to-end service delivery across domains using the IPv6 protocol without additional shim layers or gateways. This makes the IPv6 protocol fully self-sufficient for unicast services. And its built-in multicast capabilities makes IPv6 also self-sufficient for multicast distribution.

3.1. Multicast forwarding state

To forward multicast packets, multicast forwarding state must be set up on the nodes in the multicast distribution trees. A typical IPv6 multicast route comprises:

- * (Source Address, Group Address) or (S, G)
 - Identifying the distribution tree
 - Source (S): The IPv6 address of the streaming device, which can also be * (any source)
 - Group (G): The multicast IPv6 group address
- * Incoming interface : The interface receiving the multicast packets, also known as the Reverse Path Forwarding (RPF) interface
- * Downstream interfaces : The set of interfaces where the multicast packets will be forwarded

3.2. Route instantiation

Like IPv6 unicast routes, IPv6 multicast routes can be instantiated dynamically through control plane protocols (e.g., PIM) or north-bound APIs. This document focuses only on PIM instantiated multicast forwarding state and complimentary document will cover north bound API route instantiation.

4. Distributed Control Plane

4.1. PIM-SM

Protocol Independent Multicast Sparse Mode (PIM-SM) [RFC7761] is the most common multicast routing protocol, supporting both Any Source Multicast (ASM) and Source-Specific Multicast (SSM).

PIM-SM builds multicast distribution trees by leveraging the RPF interface derived from unicast routing tables. PIM-SSM is a subset of PIM-SM that does not use the Rendezvous Points (RPs) but instead requires that receivers know the (source, group) pair and signal that explicitly. Most current routing platforms support PIM-SM.

As its name implies, PIM is independent of the underlying unicast routing protocol and leverages the unicast routing tables to establish the multicast trees. It builds a distribution tree by joining the tree via the RPF interface, which is the interface that would be used to reach the source via unicast routing.

For a detailed introduction to PIM-SM and PIM-SSM, refer to Sections 3 and 4.8 of [RFC7761].

4.2. PIM-SM with Flex-Algo

By default, a multicast source advertised in the IGP is reachable via the IGP shortest-path (SPT algorithm 0). Therefore, by default, the distribution tree follows the IGP shortest-path tree.

Flex-Algo ([RFC9350], [RFC9502]) enables using customized algorithms to compute IGP shortest paths. When a multicast source (S1) is advertised with a specific Flex-Algo (e.g., FA128), PIM-SM builds the distribution trees rooted at S1 along the shortest paths defined by the Flex-Algo FA128. This effectively traffic-engineers the path of the multicast packets streamed by S1.

In an example use case, an operator uses FA128 to optimize the min-delay link metrics, providing low-delay paths. The multicast distribution trees using FA128 provide low-delay distribution, improving real-time applications like video streaming.

5. Centralized Control plane

Certain multicast use cases require explicit control over distribution trees, beyond what is possible with PIM combined with Flex-Algo. Very specific or complex traffic-engineered distribution trees may not be practical or possible using PIM. Dynamic distributed protocols may be undesirable in a network where strict determinism is required. Or the network operator wants to avoid the PIM protocol mechanisms (such as messages and RPF selection).

In these scenarios, a centralized controller can compute and instantiate multicast forwarding entries on the routers in the distribution trees. The controller maintains the IPv6 multicast forwarding entries and updates them when needed, e.g., following a topology change. These forwarding entries are explicitly instantiated native IPv6 multicast forwarding entries. They can be statically configured by the operator like static unicast routes or explicitly instantiated by a controller via a north-bound API.

The benefits of centralized control are that it enables complex traffic-engineered multicast trees and it removes the dependence on PIM protocol mechanisms.

On the data plane this solution still uses the well-know native IPv6 multicast forwarding entries.

The explicitly instantiated distribution trees and PIM-created trees can coexist in the network by using different multicast sources and groups for both mechanisms.

6. Virtualization and Overlay Solutions

Legacy IPv4 multicast implementations are often difficult or impossible to migrate to IPv6. Some services require multicast (IPv4 or IPv6) within a VPN, necessitating an overlay solution. These solutions leverage the native IPv6 multicast underlay by encapsulating the overlay IP multicast packets in IPv6 using IP-in-IPv6 encapsulation without intermediate shim layers.

The native IPv6 underlay transport distribution trees can be established using any of the mechanisms described in sections 6 and 7.

To deploy overlay IP multicast services in an IPv6-only network, existing solutions are used. There are no SRv6-specific requirements for these solutions. These solutions enable transporting overlay IP multicast (IPv4 and IPv6), in a Global Table Multicast (GTM) or MVPN context, over an IPv6-only network, encapsulated as IP-in-IPv6.

As existing native IPv6 capabilities are leveraged, there are no SRv6 specific procedures for this solution.

Specifically, the following RFCs are relevant in this context:

- * [RFC6513] specifies the architecture for Multicast Virtual Private Networks (MVPNs) in MPLS/BGP IP VPN environments. While it includes MPLS in its title, it also provides significant details relevant to IP-only solutions for MVPNs. The document describes how multicast traffic is transported over VPNs, how control plane signaling establishes the multicast forwarding paths, and the encapsulation mechanisms for data plane traffic.
- * [RFC6514] defines the specific BGP encodings and procedures for signaling multicast routing information in MVPNs as specified in [RFC6513].
- * [RFC6515] updates [RFC6514] to ensure that service providers can use either IPv4 or IPv6 addresses for internal infrastructure (e.g., PE-to-PE or PE-to-P router communication) in the context of multicast VPNs. This is particularly relevant in IPv6-only transport networks.

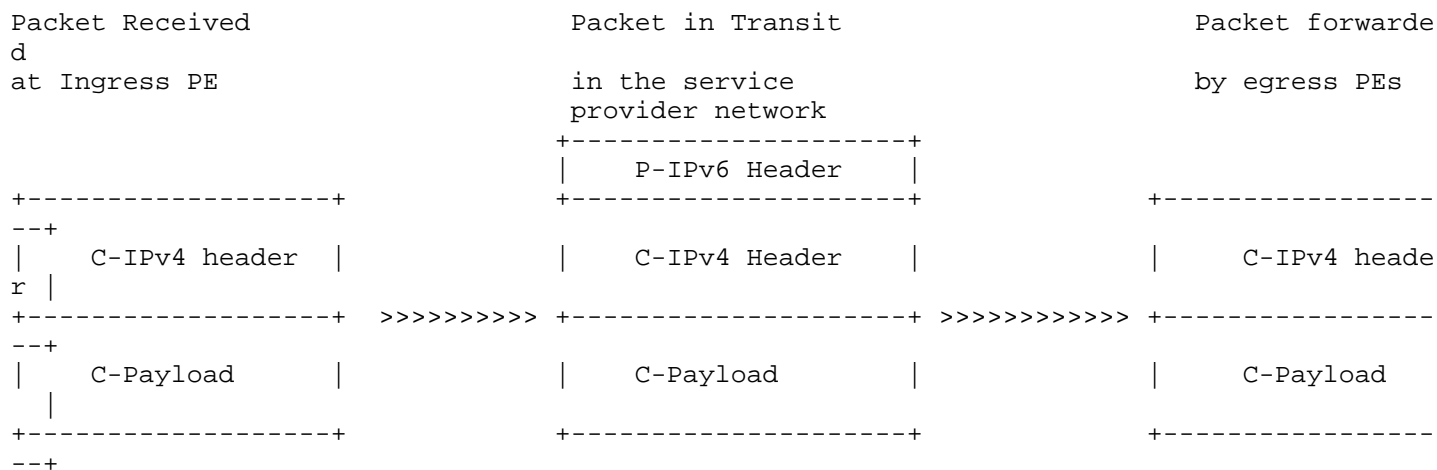
- * RFC 7716 extends the MVPN procedures defined in [RFC6514] to support Global Table Multicast (GTM), i.e., multicast routing in the global routing table rather than within the context of VPNs. This is particularly relevant in IPv6-only networks that require IPv4 multicast services without using VPNs.
- * [RFC8950] addresses the challenge of advertising IPv4 NLRI in BGP updates when the next hop is an IPv6 address. This is particularly relevant for enabling IPv4 routing in IPv6-only networks.

As an example, the following solution could be used to provide an MVPN service in an SRv6 network. Note that other solutions may equally be possible.

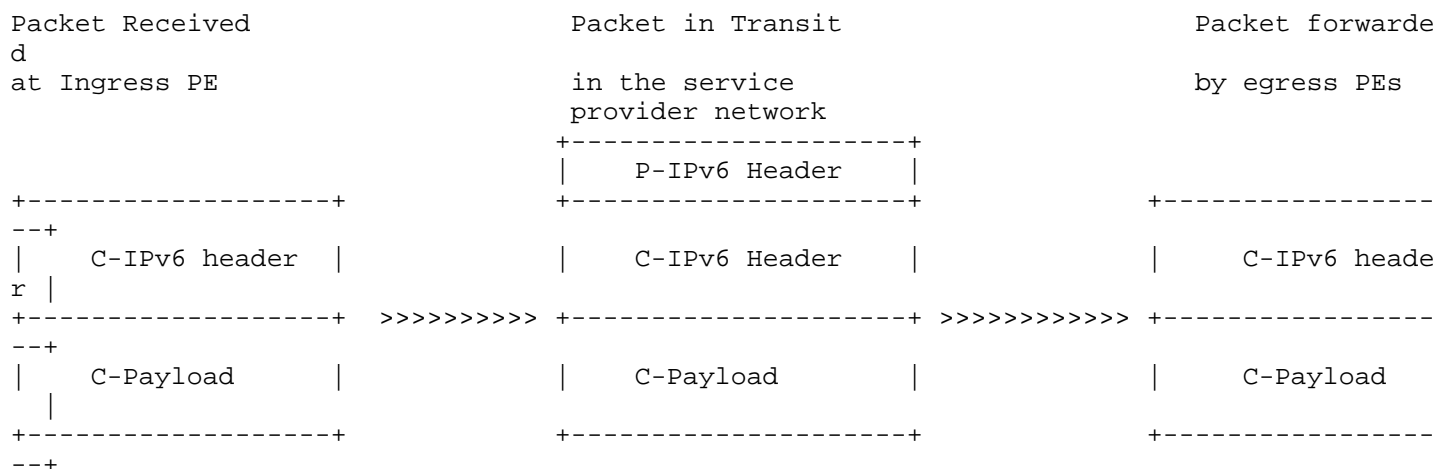
The native IPv6 multicast provider distribution trees (P-Multicast Service Interfaces (PMSIs) or P-tunnels) are build using PIM-SSM, signaling the PMSI A-D routes in MP-BGP. PIM is used to signal multicast participation between CE and PE. MP-BGP is used to signal the customer multicast routes between the MVPN PEs. The overlay customer multicast packets are transported over the network using IP-in-IPv6 encapsulation, traversing the underlay PMSI distribution trees.

7. Data plane IP in IP Encapsulations

IP in IP encapsulation is used when sending IPv4 multicast traffic though provider multicast tree. The following diagram shows the progression of the packet as it enters and leaves the service provider network.



The "Next Header" field in the P-IP IPv6 Header MUST be set to 4
Since the C-IP header is an IPv4 header.



The "Next Header" field in the P-IP IPv6 Header MUST be set to 41
Since the C-IP header is an IPv6 header.

8. IANA Considerations

Informational draft and nothing required from IANA.

9. Security Considerations

There is no new security issue imposed by this draft on top of
[RFC8950], [RFC6513]. [RFC6514], [RFC6515], [RFC7761].

10. Acknowledgement

11. Contributors

Kris Michielsen

Anuj Budhiraja

Zafar Ali

12. Normative References

- [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>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.

- [RFC6532] Yang, A., Steele, S., and N. Freed, "Internationalized Email Headers", RFC 6532, DOI 10.17487/RFC6532, February 2012, <<https://www.rfc-editor.org/info/rfc6532>>.
- [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>>.
- [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>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.
- [RFC9502] Britto, W., Hegde, S., Kaneriya, P., Shetty, R., Bonica, R., and P. Psenak, "IGP Flexible Algorithm in IP Networks", RFC 9502, DOI 10.17487/RFC9502, November 2023, <<https://www.rfc-editor.org/info/rfc9502>>.
- [RFC6515] Aggarwal, R. and E. Rosen, "IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPN", RFC 6515, DOI 10.17487/RFC6515, February 2012, <<https://www.rfc-editor.org/info/rfc6515>>.
- [RFC8950] Litkowski, S., Agrawal, S., Ananthamurthy, K., and K. Patel, "Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop", RFC 8950, DOI 10.17487/RFC8950, November 2020, <<https://www.rfc-editor.org/info/rfc8950>>.
- [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>>.

Authors' Addresses

Mankamana Mishra
Cisco Systems
821 Alder Drive,
MILPITAS, CALIFORNIA 95035
United States
Email: mankamis@cisco.com

Zhaohui Zhang
Juniper Networks
United States
Email: zzhang@juniper.net

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