

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 25 February 2026

H. Bidgoli, Ed.
Nokia
Z. Ali
Cisco System
Z. Zhang
Juniper Networks
A. Budhiraja
D. Voyer
Cisco System
24 August 2025

Segment Routing MPLS Point-to-Multipoint (P2MP) Policy Ping
draft-ietf-pim-p2mp-policy-ping-19

Abstract

Point-to-Multipoint (P2MP) Policies are used to define and manage explicit P2MP paths within a network. These policies are typically calculated via a controller-based mechanism and installed via, e.g., a Path Computation Element (PCE). In other cases these policies can be installed via using NETCONF/YANG or CLI. They are used to steer multicast traffic along optimized paths from a Root to a set of Leaf routers.

This document defines extensions to Ping and Traceroute mechanisms for SR-P2MP Policy with MPLS encapsulation to provide OAM (Operations, Administration, and Maintenance) capabilities. The extensions enable operators to verify connectivity, diagnose failures and troubleshoot forwarding issues within P2MP Policy multicast trees.

By introducing new mechanisms for detecting failures and validating path integrity, this document enhances the operational robustness of P2MP multicast deployments. Additionally, it ensures that existing MPLS and SR-based OAM tools can be effectively applied to networks utilizing P2MP Policies.

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 25 February 2026.

Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Conventions used in this document	3
3. Motivation	4
3.1. MPLS P2MP Policy Ping and Traceroute	4
3.1.1. Applicability of current RFC to SR P2MP Policies	4
3.1.2. Conformance to Existing Procedures and Additional Considerations	6
3.1.3. Considerations for Interworking with Unicast SR Domains	6
3.2. Packet format and new TLVs	6
3.2.1. Identifying a P2MP Policy	7
3.2.1.1. SR MPLS P2MP Policy Tree Instance FEC Stack Sub-TLVs	7
3.3. Limiting the Scope of Response	8
4. Implementation Status	8
4.1. Nokia Implementation	9
5. IANA Consideration	9
6. Security Considerations	9
7. Acknowledgments	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10
Authors' Addresses	10

1. Introduction

[draft-ietf-pim-sr-p2mp-policy] explains the concept of the P2MP Policy and its Candidate Paths (CPs). It also explains the concept of how a CP is selected to be the active CP. To enable seamless global optimization a CP may consist of multiple P2MP Tree Instances (PTIs), allowing for Make-Before-Break (MBB) procedures between an active PTI and a newly established, optimized PTI. A PTI is the actual P2MP tunnel set up from the Root to a set of Leaves via transit routers. A PTI is identified on the Root node by the PTI's instance ID.

To ensure reliable network operation, it is essential to verify end-to-end connectivity for both active and backup CPs, as well as all associated PTIs. This document specifies a mechanism for detecting data plane failures within a P2MP Policy CP and its associated PTIs, enabling operators to monitor and troubleshoot multicast path integrity.

This specification applies exclusively to Replication Segments (Replication SIDs) that use MPLS encapsulation for forwarding and does not cover Segment Routing over IPv6 (SRv6). The mechanisms described herein build upon the concepts established in [RFC6425] for P2MP MPLS Operations, Administration, and Maintenance (OAM). All considerations and limitations described in section 6 of [RFC6425] apply to this document as well.

1.1. Terminology

The readers of this document should familiarize themselves with the following documents and sections for terminology and details implementation of the P2MP SR Policy

[RFC9524] section 1.1 defines terms specific to SR Replication Segment and also explains the Node terminology in a Multicast domain, including the Root Node, Leaf Node and a Bud Node.

[draft-ietf-pim-sr-p2mp-policy] section 2, defines terms and concepts specific to SR P2MP Policy including the CP and the PTI.

2. Conventions used in this document

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.

3. Motivation

A P2MP Policy and its corresponding Replication Segments are typically provisioned via a centralized controller or configured using NETCONF/YANG or CLI. The root and the leaves are discovered in accordance with [draft-ietf-pim-sr-p2mp-policy] and the multicast tree is computed from the root to the leaves. However, there is no underlay signaling protocol to distribute the P2MP Policy from the root to the leaf routers. Consequently, when a P2MP tree fails to deliver user traffic, identifying the failure can be challenging without ping and traceroute mechanisms to isolate faults along the tree.

To address this challenge, P2MP Policy ping and traceroute can be utilized to detect and localize faults within the P2MP tree and its associated Replication Segments, as defined in [RFC9524]. These OAM tools enable periodic ping operations to verify connectivity between the root and the leaves. In cases where a ping fails, a traceroute can be initiated to determine the point of failure along the tree. This diagnostic process can be initiated from the node responsible for establishing the P2MP Policy, ensuring proactive monitoring and fault detection.

3.1. MPLS P2MP Policy Ping and Traceroute

Ping/Traceroute packets are forwarded based upon the P2MP Policy, on a specific CP and its PTI toward the designated leaf routers. These packets are replicated at the replication point based on the Replication Segment forwarding information on the corresponding router.

Packets are processed based on the standard behavior when their Time-to-Live (TTL) expires or when they reach the egress (leaf) router. The appropriate response is sent back to the root node following the procedures outlined in [RFC6425].

3.1.1. Applicability of current RFC to SR P2MP Policies

The procedures in [RFC6425] define fault detection and isolation mechanisms for P2MP MPLS LSPs and extend the LSP ping techniques described in [RFC8029] such that they may be applied to P2MP MPLS LSPs, ensuring alignment with existing fault management tools. [RFC6425] emphasizes the reuse of existing LSP ping mechanisms designed for Point-to-Point P2P LSPs, adapting them to P2MP MPLS LSPs to facilitate seamless implementation and network operation.

The fault detection procedures specified in [RFC6425] are applicable to all P2MP MPLS protocols, including P2MP RSVP-TE and Multicast LDP and now P2MP SR Policy. While [RFC6425] highlights specific differences for P2MP RSVP-TE and Multicast LDP, this document introduces considerations unique to P2MP SR Policies, including:

1. Egress Address P2MP Responder Sub-TLVs: Multicast LDP, as per section 3.2.1 of [RFC6425], does not allow for the inclusion of Egress Address P2MP Responder Sub-TLVs, as upstream LSRs lack visibility into downstream leaf nodes. Similarly, P2MP SR Policies often rely on a Path Computation Element (PCE) for programming transit routers. This is why in P2MP SR domain, transit routers do not have knowledge of the leaf nodes. Only the Root node, where the P2MP SR Policy is programmed, may have visibility into the leaf nodes. Consequently, these Sub-TLVs SHOULD NOT be used when an echo request carries a P2MP Policy MPLS Candidate Path FEC. If a node receives the Egress Address P2MP Responder Sub-TLVs in an echo request, then it will not respond since it is unaware of whether it lies on the path to the address in the sub-TLV.
2. End of Processing for Traceroutes: In Multicast LDP LSPs, the initiating LSR may not always be aware of all egress nodes, unlike P2MP RSVP-TE. In the case of P2MP SR Policies, the Root of the tree may have full visibility into the egress nodes if the P2MP SR Policy is PCC-initiated. If the P2MP SR Policy is PCE-initiated, the Root may or may not have visibility into the egress nodes, as this depends on the specific implementation and configuration of the PCE. Based on this, a P2MP SR Policy SHOULD follow the recommendations in Section 4.3.1 of [RFC6425], depending on the level of visibility the Root has into the egress nodes. For example, in a PCC-initiated P2MP SR Policy, the Root can learn egress node identities through Next-Generation MVPN procedures and BGP, as described in [RFC6514]. In contrast, for a PCE-initiated P2MP SR Policy, the PCE may not provide the egress node information to the Root, making this process optional and implementation-specific.
3. Identification of the LSP under test: [RFC6425], in Section 3.1, defines distinct identifiers for P2MP RSVP-TE and Multicast LDP when identifying an LSP under test. As each protocol has its own identifier, this document introduces a new Target FEC Stack TLV specific to P2MP SR Policies to uniquely identify their Candidate Paths (CPs) and P2MP Tree Instances (PTIs). These modifications ensure that P2MP Policy OAM mechanisms are properly aligned with existing MPLS ping and traceroute tools while addressing the specific operational characteristics of P2MP SR Policies.

3.1.2. Conformance to Existing Procedures and Additional Considerations

In addition to major differences outlined in the previous section, P2MP SR Policies SHOULD follow to the common procedures specified in [RFC6425] for P2MP MPLS LSPs. Furthermore, this specification reuses the same destination UDP port as defined in [RFC8029] for consistency with existing MPLS OAM mechanism.

Implementations MUST account for the fact that a P2MP Policy may contain multiple CPs, and each CP may consist of multiple PTIs. As such, implementations SHOULD support the ability to individually test each CP and its corresponding PTI using ping and traceroute mechanisms. The ping and traceroute packets are forwarded along the specified CP and its PTI, traversing the associated Replication Segments. When a downstream node capable of understanding the replication SID receives a ping or traceroute packet, it MUST process the request and generate a response even if the CP and its PTI are not currently the active path.

3.1.3. Considerations for Interworking with Unicast SR Domains

In certain deployments, two Replication Segments may be interconnected via an intermediate Unicast SR domain, as an example ingress replication. In such scenarios, proper TTL handling is required based on the hierarchical MPLS TTL mode being used (e.g., Pipe Mode vs. Uniform Mode). For example, when a P2MP Policy Ping or Traceroute packet between two Replication Segment is transiting over an Unicast SR domain, it must be only processed on Replication Segments, based on the Replication SID's TTL value. The Unicast SR domain MUST be treated as a single hop for Replication SID's TTL value, meaning that the Replication SID TTL MUST be decremented by one before it is forwarded through the Unicast SR domain and the replication SID TTL should be examined on the next replication segment. The Unicast SR domain Failure detection is considered out of scope for this document. The operator can use the OAM mechanisms available for Unicast SR to farther troubleshoot the Unicast SR domain.

3.2. Packet format and new TLVs

The packet format used in this specification follow section 3 of [RFC8029]. However, additional TLVs and sub-TLVs are required to support the new functionality introduced for P2MP Policies. These extensions are described in the following sections.

3.2.1. Identifying a P2MP Policy

[RFC8029] defines a standardized mechanism for detecting data-plane failures in Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs). To correctly identify the Replication Segment associated with a given Candidate Path (CP) and P2MP Tree Instance (PTI), the Echo Request message MUST include a Target FEC Stack TLV that explicitly specifies the Candidate Path and P2MP Tree Instance under test.

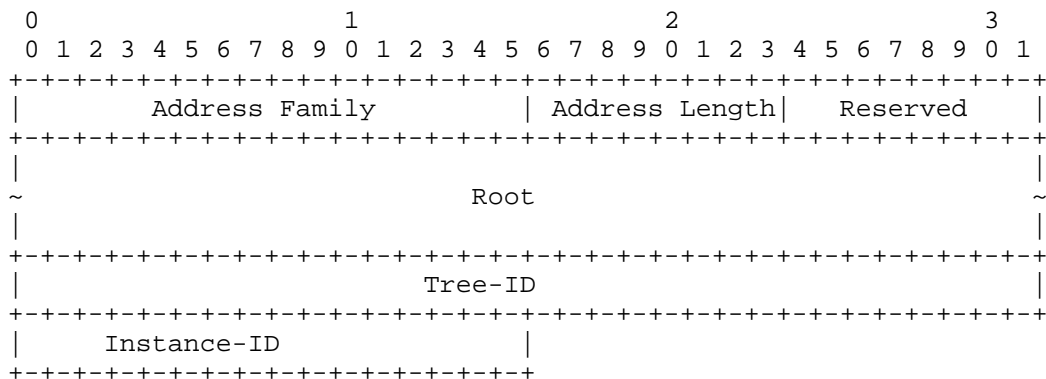
This document introduces a new sub-TLV, referred to as the SR MPLS P2MP Policy Tree Instance sub-TLV, which is defined as follows:

Sub-Type	Length	Value Field
-----	-----	-----
41	Variable	SR MPLS P2MP Policy Tree Instance

Further details regarding the structure and processing of this sub-TLV are provided in subsequent sections.

3.2.1.1. SR MPLS P2MP Policy Tree Instance FEC Stack Sub-TLVs

The SR MPLS P2MP Policy Tree Instance sub-TLV value field follows the format specified in Section 2.3 of [draft-ietf-pim-sr-p2mp-policy]. The structure of this sub-TLV is illustrated in the figure below. It should be noted that this sub-TLV is testing a specific PTI within a specific CP and it is not testing the CP.



- * Address Family: (2 octets) IPv4/IPv6 ADDRESS FAMILY NUMBERS as specified in [IANA-AF] , indicating the address family of the Root. Any other Address Family but IPv4/IPv6 is not supported by this draft.

- * Address Length: (1 octet) specifying the length of the Root Address in octets (4 octets for IPv4, 16 octets for IPv6).
- * Reserved: Must be set to zero.
- * Root: (variable length depending on the address family field) The root node of the P2MP Policy, as defined in [draft-ietf-pim-sr-p2mp-policy]
- * Tree-ID: (4 octets) A unique identifier for the P2MP tree, as defined in [draft-ietf-pim-sr-p2mp-policy]
- * Instance-ID: (2 octets) identifies the specific Path-Instance as defined in [draft-ietf-pim-sr-p2mp-policy]

3.3. Limiting the Scope of Response

As specified in section 3.2 of [RFC6425] , four sub-TLVs are used within the P2MP Responder Identifier TLV included in the echo request message.

The Sub-TLVs for IPv4 and IPv6 egress addresses of P2MP responder are aligned with section 3.2.1 of [RFC6425].

The sub-TLVs for IPv4 and IPv6 node addresses of the P2MP responder are aligned with Section 3.2.2 of [RFC6425]

These mechanisms ensure that responses are appropriately scoped to limit unnecessary processing and improve the efficiency of P2MP OAM procedures.

4. 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".

4.1. Nokia Implementation

Nokia has implemented [draft-ietf-pim-sr-p2mp-policy] and [RFC9524]. In addition, Nokia has implemented P2MP policy ping as defined in this draft to verify the end to end connectivity of a P2MP tree in 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.

5. IANA Consideration

IANA has assigned a TEMPORARY code point for the "SR MPLS P2MP Policy Tree Instance" Sub-TLV Name. This Sub-TLV is assigned from TLV type 1 (Target FEC Stack) from the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry group. The Sub-TLVs for TLV type 1 are listed under "Sub-TLVs for TLV Types 1, 16, and 21" sub-registry. This sub-type value is assigned from the standards Action range of 0-16383 from the "Sub-TLVs for TLV Types 1, 16, and 21" sub-registry.

Sub-Type	Sub-TLV Name
----------	--------------

41	SR MPLS P2MP Policy Tree Instance
----	-----------------------------------

6. Security Considerations

Overall, the security needs for P2MP policy ping are the same as [draft-ietf-pim-sr-p2mp-policy] and [RFC8029]. The P2MP policy ping is susceptible to the same three attack vectors as explained in RFC8029 section 5. The same procedures and recommendations explained in [RFC8029] section 5 should be taken and implemented to mitigate these attack vectors for P2MP policy Ping as well.

7. Acknowledgments

8. References

8.1. Normative References

- [draft-ietf-pim-sr-p2mp-policy]
"D. Yoyer, C. Filsfils, R.Prekh, H.bidgoli, Z. Zhang,
"draft-ietf-pim-sr-p2mp-policy"", July 2025.
- [RFC2119] "S. Brandner, "Key words for use in RFCs to Indicate
Requirement Levels"", March 1997.
- [RFC6425] "S. Saxena, G. Swallow, Z. Ali, A. Farrel, S. Yasukawa,
T.Nadeau "Detecting Data-Plane Failures in Point-to-
Multipoint MPLS"", November 2011.
- [RFC6514] "R.Aggarwal, E. Rosen, T. Morin, Y. Rekhter "BGP Encodings
and Procedures for Multicast in MPLS/BGP IP VPNs"",
February 2012.
- [RFC8029] "K. Kompella, G. Swallow, C. Pignataro, N. kumar, S. Aldrin
M. Chen, "Detecting Multiprotocol Label Switched (MPLS)
Data-Plane Failures.", February 2006.
- [RFC8174] "B. Leiba, "ambiguity of Uppercase vs Lowercase in RFC
2119 Key Words"", May 2017.
- [RFC9524] "D. Voyer, C. Filsfils, R. Parekh, H. Bidgoli, Z. Zhang,
"Segment Routing Replication for Multipoint Service
Delivery"", February 2024.

8.2. Informative References

- [IANA-AF] "IANA Assigned Port Numbers,
"http://www.iana.org/assignments/address-family-numbers"". "

Authors' Addresses

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

Zafar
Cisco System
San Jose,
United States of America
Email: zali@cisco.com

Zhaohui Zhang
Juniper Networks
Boston,
United States of America
Email: zzhang@juniper.net

Anuj Budhiraja
Cisco System
San Jose,
United States of America
Email: abudhira@cisco.com

Daniel Voyer
Cisco System
Montreal
Canada
Email: davoyer@cisco.com