

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-pce-sr-p2mp-policy-03
Published: 9 March 2023
Intended Status: Standards Track
Expires: 10 September 2023
Authors: H. Bidgoli, Ed. V. Voyer S. Rajarathinam
 Nokia Bell Canada Nokia
 A. Budhiraja T. Saad S. Sivabalan
 Cisco System Juniper Networks Ciena
PCEP extensions for p2mp sr policy

Abstract

SR P2MP policies are set of policies that enable architecture for P2MP service delivery. This document specifies extensions to the Path Computation Element Communication Protocol (PCEP) that allow a stateful PCE to compute and initiate P2MP paths from a Root to a set of Leaves.

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 10 September 2023.

Copyright Notice

Copyright (c) 2023 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. [Conventions used in this document](#)
3. [Overview of PCEP Operation in SR P2MP Network](#)
 - 3.1. [High level view of P2MP Policy Objects](#)
 - 3.2. [Existing drafts used for defining a P2MP Policy](#)
 - 3.2.1. [Existing Documents used by this draft](#)
 - 3.2.2. [P2MP Policy Identification](#)
 - 3.2.3. [Replication Segment Identification](#)
 - 3.2.4. [PCECC Use in Replication Segment](#)
 - 3.3. [High Level Procedures for P2MP SR LSP Instantiation](#)
 - 3.3.1. [PCE-Init Procedure](#)
 - 3.3.2. [PCC-Init Procedure](#)
 - 3.3.3. [Common Procedure](#)
 - 3.3.4. [Global Optimization of the Candidate Path](#)
 - 3.3.5. [Fast Reroute](#)
 - 3.3.6. [Connecting Replication Segment via Segment List](#)
 - 3.4. [SR P2MP Policy and Replication Segment TLVs and Objects](#)
 - 3.4.1. [SR P2MP Policy Objects](#)
 - 3.4.2. [Replication Segment Objects](#)
 - 3.4.3. [P2MP Policy and Replication Segment general considerations](#)
 - 3.4.3.1. [Path Attribute Object](#)
 - 3.4.3.2. [PCE as Central Controller](#)
4. [Object Format](#)
 - 4.1. [Open Message and Capability Exchange](#)
 - 4.2. [PCECC Path Setup Capability](#)
 - 4.3. [Association Type Capability](#)
 - 4.4. [Symbolic Name in PCInit Message from PCC](#)
 - 4.5. [P2MP Policy Specific Objects and TLVs](#)
 - 4.5.1. [P2MP Policy Association Group for P2MP Policy](#)
 - 4.5.1.1. [P2MP SR Policy Identifiers TLV](#)
 - 4.5.1.2. [P2MP SR Policy Name TLV](#)
 - 4.5.1.3. [P2MP SR Policy Candidate Path ID TLV](#)
 - 4.5.1.4. [P2MP SR Policy Candidate Path Name TLV](#)
 - 4.5.1.5. [P2MP SR Policy Candidate Path PREFERENCE Attributes TLV](#)
 - 4.5.2. [P2MP-END-POINTS Object](#)
 - 4.6. [P2MP Policy and Replication Segment Identifier Object and TLV](#)
 - 4.6.1. [Extension of the LSP Object, SR-P2MP-LSPID-TLV](#)
 - 4.7. [Replication Segment](#)
 - 4.7.1. [The format of the replication segment message](#)
 - 4.7.2. [PCECC](#)
 - 4.7.3. [Label action rules in replicating segment](#)

- [4.7.4. SR-ERO Rules](#)
 - [4.7.4.1. SR-ERO subobject changes](#)
- [5. Tree Deletion](#)
- [6. Fragmentation](#)
- [7. Some packet examples](#)
 - [7.1. Report for Leaf Add](#)
 - [7.2. P2MP Policy Candidate path Init](#)
 - [7.3. Replication segment PCE Initiated on Transit and LEaves](#)
- [8. Example Workflows](#)
- [9. IANA Consideration](#)
- [10. Security Considerations](#)
- [11. Acknowledgments](#)
- [12. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

The draft [[draft-ietf-pim-sr-p2mp-policy](#)] defines a variant of the SR Policy that uses [[RFC9256](#)] for constructing a P2MP segment to support multicast service delivery.

A Point-to-Multipoint (P2MP) Policy connects a Root node to a set of Leaf nodes, optionally through a set of intermediate replication nodes. A Replication segment [[draft-ietf-spring-sr-replication-segment](#)], corresponds to the state of a P2MP segment on a particular node and provide forwarding instructions for the segment.

A P2MP Policy is relevant on the root of the P2MP Tree and it contains candidate paths. The candidate paths are made of path-instances and each path-instance is constructed via replication segments. These replication segments are programmed on the root, leaves and optionally intermediate replication nodes.

Replication segments MAY be connected to each other directly, or they MAY be connected or steered via unicast SR segments or a segment list.

For a P2MP Tree, a controller may be used to compute paths from a Root node to a set of Leaf nodes, optionally via a set of replication nodes. A packet is replicated at the root node and optionally on Replication nodes towards each Leaf node.

There are two types of a P2MP Tree: Spray and Replication.

A Point-to-Multipoint service delivery could be via Ingress Replication, known as Spray. The root unicasts individual copies of traffic to each leaf. The corresponding P2MP Policy consists of replication segments only for the root and the leaves and they are connected via a unicast SR Segment.

A Point-to-Multipoint service delivery could also be via Downstream Replication, known as Replication Tree. The root and some downstream replication nodes replicate the traffic along the tree as it traverses closer to the leaves.

The PCE discovers the root and the leaves via different methods. As an example, the leaves and the root can be explicitly configured on the PCE or PCC can update the PCE with the identity of the root and the leaves when it discovers them via multicast protocols like MP-BGP and MVPN procedures [[RFC6513](#)] or PIM. The controller can calculate the P2MP Policy and any of its associated replication segments from the root to the leaves with these info and any additional Service Leave Agreements (SLAs) that is used to construct the tree.

This document defines PCEP objects, TLVs and the procedures to instantiate a P2MP Policy and Replication Segments.

2. Conventions used in this document

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](#)].

3. Overview of PCEP Operation in SR P2MP Network

After discovering the root and the leaves the PCE programs the PCCs with relevant information needed to create a P2MP Tree.

As per draft [[draft-ietf-pim-sr-p2mp-policy](#)] a P2MP Policy is defined by Root-ID, Tree-ID and a set of leaves. A P2MP policy is a variant of SR policy as such it uses the same concept as draft [[draft-ietf-pce-segment-routing-policy-cp](#)]. A P2MP policy is composed of a collection of SR P2mp Candidate Paths. Candidate paths are computed by the PCE and can be used for P2MP Tree redundancy. Only a single candidate path may be active at each time. Each candidate paths can be globally optimized, therefore it consists of multiple path-instances. A path-instance can be considered as a P2MP LSP. If a candidate path needs to be globally optimized two path-instances can be programmed from the root the leaves and via make before break procedures the candidate path can be switched from path-instance 1 to the 2nd path-instance. The forwarding states of these path-instances are build via replication segments, in short each path-instance initiated on the root has its own set of replication segments on the Root, Transit and Leaf nodes.

A replication segment is set of forwarding instructions on a specific node. Each instruction may be a PUSH or SWAP operation before forwarding out of an interface, or a POP action on bud and leaf nodes.

PCE could also calculate and download additional information for the replication segments, such as protections next-hops for link protection (FRR).

3.1. High level view of P2MP Policy Objects

*SR P2MP Policy

- Is only relevant on the Root of the P2MP tree and is a policy on PCE. It is downloaded only on the root node and is identified via <Root-ID, Tree-ID> It contains the following information:

 - oRoot node of the P2MP Segment

 - oSet of Leaf nodes of the P2MP Segment

 - oTree-ID, which is a unique identifier of the P2MP tree on the Root

 - oA set of Candidate paths belonging to the policy

 - oOptional Constraints used to build these candidate paths

*Candidate Path:

- Is used for P2MP Tree redundancy where the candidate path with the highest preference is the active path.

- Each Candidate Path can contain two path-instance for global optimization procedures (i.e. make before break)

- Contains information regarding originator, discriminator, preference, path-instances

*Path-instance:

- Is used for global optimization of the candidate path.
Multiple path-instance can be present under a candidate path but only a single path instance is active at a time.

- A path instance is identified via <Root-ID, Tree-ID, Instance-ID>

*Replication Segment:

- Is the forwarding information needed on each replication node for building the forwarding path for each path-instance of the P2MP Candidate path.

-Explained further in upcoming sections, there are 2 ways to identify the replication segment, depending on the type of replication segment (shared replication segment or non-shared replication segment)

- oIt is identified via Tree-ID and Root-ID and path-instance for non-shared replication segment.

- oIt is identified via Node-ID, Replication-ID, for shared replication segment. As per [[draft-ietf-pim-sr-p2mp-policy](#)] a shared replication segment is not associated to a tree and is used for constructing by-pass tunnels.

- oContains forwarding instructions, in the form of a list of outgoing segments each of which may be a segment list or a single replication segment with next-hop information.

- oOn the forwarding plane the Replication Segment is identified via the incoming Replication SID.

- oReplication segment information is downloaded on any node that is replicating the packet on the path of the tree including the Root, Transit and Leaf nodes respectively.

3.2. Existing drafts used for defining a P2MP Policy

This document attempts to leverage existing IETF draft and RFC documents which define PCEP objects, to update the PCE with Root and Leaves information when PCC Initiated method is used. Similarly, existing documents are utilized where feasible to update the PCC with relevant information to build the P2MP Policy and its Replication Segments. This document introduces new TLVs and Objects specific to a programming P2MP policy and its replication segment.

3.2.1. Existing Documents used by this draft

- *[[RFC8231](#)] The bases for a stateful PCE, and reuses the following objects or a variant of them

- <SRP Object>

- <LSP Object>

- A variation of the LSP identifier TLV is defined in this draft, to support P2MP LSP Identifier

- *[[RFC8306](#)] P2MP capabilities advertisement

- *[[draft-ietf-pce-segment-routing-policy-cp](#)] Candidate paths for P2MP Policy is used for Tree Redundancy. As an example, a P2MP

Policy can have multiple candidate paths. Each protecting the primary candidate path. The active path is chosen via the preference of the candidate path.

- *[[RFC3209](#)] Defines the instance-ID, instance-ID is used for global optimization of a candidate path with in a P2MP policy. Each Candidate path can have 2 path-instances. These path-instances are equivalent to sub-lsps (instance-IDs). There are used for MBB and global optimization procedures. instance-ID is equivalent to LSP ID
- *[[RFC9256](#)] Segment-list, used for connecting two non-adjacent replication policy via a unicast binding SID or Segment-list.
- *[[RFC8306](#)] P2MP End Point objects, used for the PCC to update the PCE with discovered Leaves.
- *[[RFC9050](#)] for programming and identifying the Replication Segment. A new PCE CC Capability sub Tlv is introduced to indicated the support to handle PCE CC based label download for SR P2MP.
- *[[draft-ietf-pce-multipath](#)] Forwarding instruction for a P2MP LSP is defined by a set of SR-ERO sub-objects in the ERO object, ERO-ATTRIBUTES object and MULTIPATH-BACKUP TLV as defined in this draft.
- *[[RFC8664](#)] SR-ERO Sub Object used in the multipath.

It should be noted that the [[draft-hb-spring-sr-p2mp-policy-yang](#)] can provide further details of the high level P2MP Policy Model.

3.2.2. P2MP Policy Identification

A P2MP Policy and its candidate path can be identified on the root via the P2MP LSP Object. This Object is a variation of the LSP ID Object defined in [[RFC8231](#)] and is as follow:

- *PLSP-ID: [[RFC8231](#)], is assigned by PCC and is unique per candidate path. It is constant for the lifetime of a PCEP session. Stand-by candidate paths will be assigned a new PLSP-ID by PCC. Stand-by candidate paths can co-exist with the active candidate path.
- *Root-ID: is equivalent to the first node on the P2MP path, as per [[RFC3209](#)], Section 4.6.2.1
- *Tree-ID: is equivalent to Tunnel Identifier color which identifies a unique P2MP Policy at a ROOT and is advertised via

the PTA in the BGP AD route or can be assigned manually on the root. Tree-ID needs to be unique on the root.

*Instance-ID: LSP ID Identifier as defined in RFC 3209, is the path-instance identifier and is assigned by the PCC. As it was mentioned the candidate path can have up to two path-instance for global optimization. Note that the Root-ID, Tree-ID and Instance-ID are part of a new SR- P2MP-LSP-IDENTIFIER TLV which will be identified in this draft. Instance-IDs are unique per P2MP Policy, as such path-instances assigned for each candidate path with in a P2MP Policy must have unique Instance-IDs.

3.2.3. Replication Segment Identification

The key to identify a replication segment is also a P2MP LSP Object. With varying encoding rules for the SR-P2MP-LSP- IDENTIFIER TLV which will be explained in later sections.

3.2.4. PCECC Use in Replication Segment

PCECC and a variant of CCI object is used in Replication Segment to identify a cross connect. A cross connect is a incoming SID and set of outgoing interfaces and their corresponding SID or SID List. The CCI objects contains the incoming SID and the outgoing interfaces which are presented via the ERO objects, which each may contain a segment list.

3.3. High Level Procedures for P2MP SR LSP Instantiation

A P2MP policy can be instantiated via the PCC or the PCE depending on how the root and the leaves are discovered. This document describes two way to discover the root and the leaves:

*They can be configured and identified on the controller and are considered PCE initiated.

*They can be discovered on the PCC via MVPN procedures [[RFC6513](#)] or legacy multicast protocols like PIM or IGMP etc... and are considered PCC initiated.

3.3.1. PCE-Init Procedure

*PCE is informed of the P2MP request through its API or configuration mechanism to instantiate a P2MP tunnel. PCE will be programmed with the Root and a set of leaf nodes.

*PCE will initiate the P2MP Policy for the request, by sending a PCInitiate message to the Root.

*Root in response to the PCInitiate message, will generate PLSP-ID for the candidate paths and an Instance-ID for the Path-Instance (LSP-ID) contained within the candidate path. The tree-id for the P2MP Policy is also filled. PCC will report back the PLSP-ID, Instance-ID and tree-id via PCRpt message

-Optionally, the Root can add any additional leaves that were discovered by multicast procedures in this PCRpt message.

*PCE will send a PCInitiate message to the Root, Transit and the Leaf nodes to download the Replication Segment information. These messages will utilize the CCI object to identify the p2mp cross connect and encode the forwarding instruction information.

*PCE will then send a PCUpdate to the root indicating the association information (Candidate path) , and implicitly indicate it to bind to the latest CCI information downloaded.

3.3.2. PCC-Init Procedure

Root node (PCC) discovers the leaves (as an example via MVPN Procedures or other mechanism), the following communication happens between the PCE and PCCs

*Root sends a PCRpt message for P2MP policy to PCE including the Root-ID, Tree-ID, PLSP-ID, Instance-ID, symbolic-path-name, and any leaves discovered until then.

*PCE on receiving of this report, will compute the P2MP Policy and its replication segments.

-PCE will send a PCInitiate message to the Root, Transit and the Leaf nodes to download the Replication Segment information. These messages will utilize the CCI object to encode the forwarding instruction information.

-PCE will then send a PCUpdate to the root indicating the association information (Candidate path) , and implicitly indicate it to bind to the latest CCI information downloaded.

3.3.3. Common Procedure

The following procedures are the same for PCE or PCC Init.

*PCE will download the replication segments for the Candidate-path's path-instances to all the leaves and transit nodes using PCInitiate message with PLSP-ID = 0, Instance-ID = 0, symbolic path name, Root-address, Tree-id (assigned by the root and obtained via P2MP Policy creation). This PCInitiate message includes the EROs needed for the replication segments. These messages will

utilize the CCI object to encode the forwarding instruction information.

*Any new candidate path for the P2MP Policy is downloaded by PCE to the Root by sending a PCInitiate message

- it should be noted, PLSP-ID, Path-Instance ID are generated by the PCC for these new candidate paths and their Path-instances

- Any update to the Candidate Paths or Replication Segments is done via the PCUpd message. Association object need to be present for Candidate Path updates and CCI object for the replication segment updates.

*The PCE will also download the necessary replication segment for the candidate path and its path-instances to the root, leaves and the transit nodes via a PCInit message

*New leaves can be discovered via Multicast procedures, and new replication segments can be instantiated or existing one updated to reach these leaves

- If these leaves reside on routers that are part of the P2MP LSP path, then PCUpd is sent from PCE to necessary PCCs (LEAVES, TRANSIT or ROOT) with the correct PLSP-ID, Instance-ID, Tree-ID and CC-ID.

- If the new leaves are residing on routers that are not part of the P2MP Tree yet, then a PCInitiate message is sent down with PLSP-ID=0 and Instance-ID=0 on the corresponding routers.

*The active candidate-path is indicated by the PCC through the operational bits(Up/Active) of the LSP object in the PCRpt message. If a candidate path needs to be removed, PCE sends PCInitiate message, setting the R-flag in the LSP object and R bit in the SRP-object.

*To remove the entire P2MP-LSP, PCE needs to send PCInitiate remove messages for every candidate path of the P2MP POLICY to the root and send PCInitiate remove messages for every Replication Segment on all the PCCs on the P2MP Tree. The R bit in the LSP Object as defined in [[RFC8231](#)], refers to the removal of the LSP as identified by the SR-P2MP-POLICY-ID-TLV (defined in this document). An all zero (SR-P2MP-LSP-ID-TLV defines to remove all the state of the corresponding PLSP-ID.

*A candidate path is made active based on the preference of the path. If the Root is programed with multiple candidate paths from different sources, as an example PCE and CLI, based on its tie-breaking rules, if it selects the CLI path, it will send a report

to PCE for the PCE path indicating the status of label-download and sets operational bit of the LSP object to UP and Not Active . At any instance, only one path will be active

3.3.4. Global Optimization of the Candidate Path

When a P2MP LSP needs to be optimized for any reason (i.e. it is taking a FRR tunnel or new routers are added to the network) a global optimization of the candidate path is possible.

Each Candidate Path can contain two Path-Instances. The current unoptimized Path-Instance is the active instance and its replication segments are forwarding the multicast PDUs from the root to the leaves. However the second optimized Path-Instance will be setup with its own unique replication segments throughout the network, from the Root to the leaves. These two Path-Instances can co-exist. The two Path-Instances are uniquely identified by their Instance-ID in the SR-P2MP-POLICY-ID-TLV (defined in this document). After the optimized LSP has been downloaded successfully PCC MUST send two reports, reporting UP of the new path indicating the new LSP-ID, and a second reporting the tear down of the old path with the R bit of the LSP Object SET with the old Instance-ID in the SR-P2MP-POLICY-ID-TLV. This MBB procedure will move the multicast PDUs to the optimized Path-Instance.

The leaf should be able to accept traffic from both Path-Instances to minimize the traffic outage by the Make Before Break process.

3.3.5. Fast Reroute

Currently this draft identifies the Facility FRR procedures. In addition, only LINK Protection procedures are defined. How the Facility Path is built and instantiated is beyond the scope of this document.




Figure 1




Figure 2

As an example, in figure 1 both R and T are configured with replication segments. There are two interface between R and T. One can be used as primary and second as a bypass in case the primary interface is down. There can be 2 method to protect the primary interface.

- *The two replication segments on R and T can take advantage of unicast SR to connect to each other. In this case the LFA of unicast SR can be utilize to protect the primary interface between R and T.

- *The replication segment provides protection nexthop, the protection nexthop can be programmed to take the alternate interface between R and T to protect the primary interface.

As a second example, in figure 2, R and T connected directly and via network F1..F2. In this example as per example 1 unicast SR can be used to connect the two replication segments and in this case the unicast SR LFA or R-LFA or TI-LFA can be used to protect the direct link between R-T via F1. That said if there is no unicast SR available with in the network, the PCE optionally can setup a shared replication point on F1 and F2 and protect all path-instances that are traversing R-T via this shared replication segment.

In addition, PHP procedure and implicit null label on the bypass path can be implemented to reduce the PCE programming on the MP PCC.

3.3.6. Connecting Replication Segment via Segment List

There could be nodes between two replication segment that do not support P2MP Policy or Replication segment. It is possible to connect two non-adjacent Replication segments via a unicast segment routing path and SID list. The SID list can be comprised of any IGP supported segment types (ex: Binding, Adjacency, Node). This information is encoded via the SR-ERO sub-objects and ERO-attributes objects. The last segment in an encoding SID list MUST be a replication segment

3.4. SR P2MP Policy and Replication Segment TLVs and Objects

3.4.1. SR P2MP Policy Objects

SR P2MP Policy can be constructed via the following objects

```
<P2mp Policy> ::= <Common Header>
                  <SRP>
                  <P2MP LSP>
                  <association-list>
```

optionally a list of end-point can be added.

This is true whether it is PCC initiated or PCE initiated

```
[<end-point-list>]
```

3.4.2. Replication Segment Objects

Replication segment can be constructed via the following objects

```
<Replication Segment> ::= <Common Header>
                          <SRP>
                          <P2MP LSP>
                          (<cci-list>|
                           (<CCI><intended-path>))
                          <cci-list> ::= <CCI>
                                          [<cci-list>]
                          <intended-path> ::=
                              ((<PATH-ATTRIB><ERO>)
                               [<intended-path>])
```

Path-attribute as per [[draft-ietf-pce-multipath](#)]

3.4.3. P2MP Policy and Replication Segment general considerations

The new objects and TLV's defined in this document can be included in PCRpt, PcInitiate and PCUpd messages.

It should be noted that every PCRpt, PcInitiate and PCUpd messages will contain full list of the Leaves and segment and forwarding information that is needed to build the Candidate path and its Replication segments. They will never send the delta information related to the new leaves or forwarding information that need to be added or updated. This is necessary to ensure that PCE or any new PCE is in sync with the PCC.

3.4.3.1. Path Attribute Object

This draft uses [[draft-ietf-pce-multipath](#)] to identify each out-going interface in the replication segment. In addition each out-going interface can be protected by a backup path. The Path Attributes Object is used to provide the relation between the primary path and its backup path as per draft [[draft-ietf-pce-multipath](#)].

When a replication segment is being updated or new out-going interfaces are added to a specific replication segment, the PCRpt, PcInitiate and PCUpd messages sent via PCEP maintains the previous ERO Path IDs and generates new Path IDs for new instructions. The PATH IDs are maintained for each specific forwarding instructions until the instructions are deleted. For example: When the first leaf is added, the PCE will update with Path ID 1 to the PCC. When the second leaf is added, according to the path calculated, PCE might just append the existing instruction Path ID 1 with a new Path ID 2 to construct the new PCUpd message.

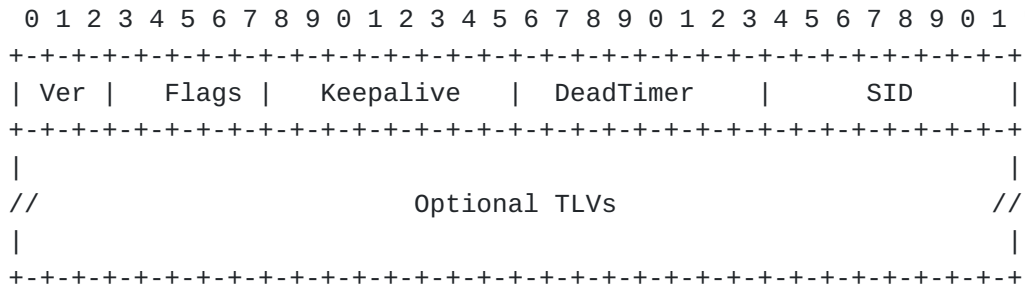
3.4.3.2. PCE as Central Controller

The CCI Object is used to identify the entire cross connect of incoming segment and the set of outgoing Interfaces and their corresponding SIDs/SIDList. Any modification to the cross connect should use this CCI ID to identify the cross connect uniquely. Leaves and their corresponding Path IDs can be removed from the cross connect identified via the CCI. The CC-ID is assigned by the PCE.

4. Object Format

4.1. Open Message and Capability Exchange

Format of the open Object:



All the nodes need to establish a PCEP connection with the PCE.

During PCEP Initialization Phase, PCEP Speakers need to set flags N, M, P in the STATEFUL-PCE-CAPABILITY TLV as defined in [\[draft-ietf-pce-stateful-pce-p2mp\]](#) section-5.2


This draft extends the PCEP OPEN object by defining an optional TLV to indicate the PCE's capability to perform SR-P2MP path computations with a new IANA capability type.

The inclusion of this TLV in an OPEN object indicates that the sender can perform SR-P2MP path computations. This will be similar to the P2MP-CAPABILITY defined in [\[RFC8306\]](#) section-3.1.2 and a new value needs to be defined for SR-P2MP.

4.2. PCECC Path Setup Capability

A Path Setup Type (PST) of PCECC is also added as per [\[RFC9050\]](#).

This document also introduces a new bit S in the SR PCECC capacity Sub TLV indicating the support to handle PCECC based label download for Replication segment.



Also, the N,M,P bits in STATEFUL-PCE-CAPABILITY TLV should be SET.

4.3. Association Type Capability

A Assoc-Type-List TLV as per [\[RFC8697\]](#) section 3.4 should be send via PCEP open object with following association type

Association Type Value	Association Name	Reference
TBD1	P2MP SR Policy Association	This document

OP-CONF-Assoc-RANGE (Operator-configured Association Range) should not be set for this association type and must be ignored.

The open message MUST include the MULTIPATH CAPABILITY TLV as defined in [[draft-ietf-pce-multipath](#)]

4.4. Symbolic Name in PCInit Message from PCC

As per [[RFC8231](#)] section 7.3.2. a Symbolic Path Name TLV should uniquely identify the P2MP path on the PCC. This symbolic path name is a human-readable string that identifies an P2MP LSP in the network. It needs to be constant through the lifetime of the P2MP path.

As an example in the case of P2MP LSP the symbolic name can be p2mp policy name + candidate path name of the LSP.

4.5. P2MP Policy Specific Objects and TLVs

4.5.1. P2MP Policy Association Group for P2MP Policy

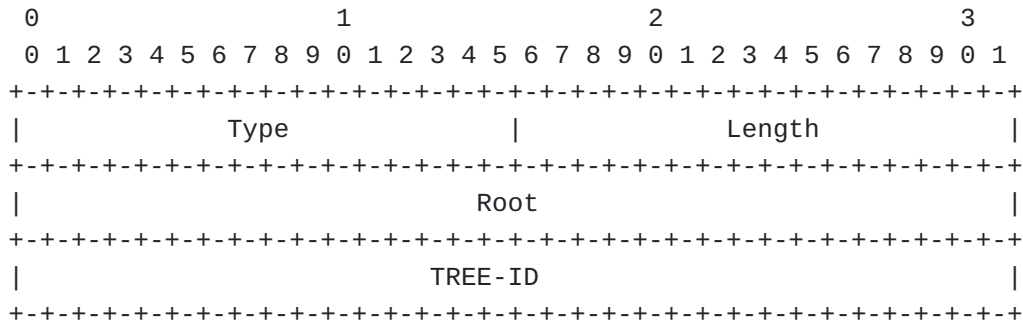
Two ASSOCIATION object types for IPv4 and IPv6 are defined in [[RFC8697](#)]. The ASSOCIATION object includes "Association type" indicating the type of the association group. This document adds a new Association type. Association type = TBD1 "P2MP SR Policy Association Type" for SR Policy Association Group (P2MP SRPAG). In par with [[draft-ietf-pce-segment-routing-policy-cp](#)] section 5.2, Five new TLVs are identified to carry association information:

1. P2MP-SRPOLICY-ID TLV: (mandatory) encodes P2MP SR Policy Identifier
2. P2MP-SRPOLICY-NAME TLV: (optional) encodes P2MP SR Policy Name
3. P2MP-SRPOLICY-CPATH-ID TLV: (mandatory) encodes P2MP SR Policy Candidate path Identifier
4. P2MP-SRPOLICY-CPATH-PREFERENCE TLV: (optional) encodes P2MP SR Policy Candidate path preference value.
5. P2MP-SRPOLICY-CPATH-NAME TLV: (optional) encodes P2MP SR Policy Candidate path name.

P2MP-SRPAG- POL-ID-TLV, P2MP-SRPAG-CPATH-ID-TLV, P2MP-SRPAG-CPATH-ATTR-TLV

4.5.1.1. P2MP SR Policy Identifiers TLV

The P2MP-SRPOLICY-ID TLV is a mandatory TLV for the P2MP SR Policy Association. Only one P2MP-SRPOLICY-ID TLV can be carried and only the first occurrence is processed and any others MUST be ignored.



Type: TBD1 for "P2MP-SRPolicy-ID" TLV.

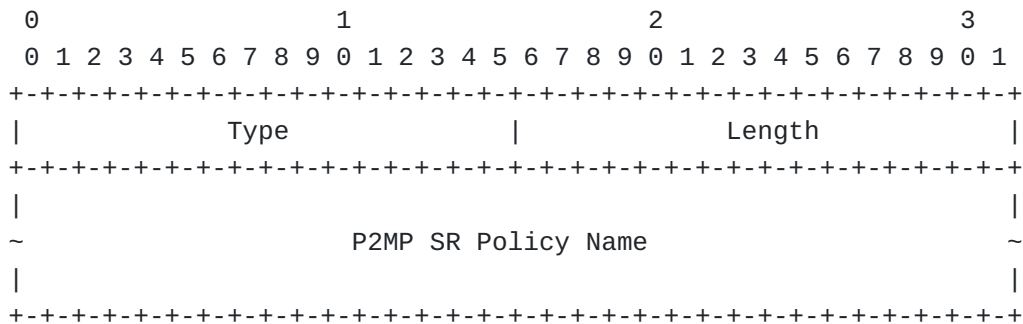
Length: 8 or 20, depending on length of End-point (IPv4 or IPv6)

Tunnel Sender Address : Can be either IPv4 or IPv6, this value is the value of the root loopback IP.

Tree-ID: Tree ID that the replication segment is part of as per draft-ietf-spring-sr-p2mp-policy

4.5.1.2. P2MP SR Policy Name TLV

The P2MP-SRPOLICY-NAME TLV is an optionally TLV for the P2MP SR Policy Association. Only one P2MP-SRPOLICY-NAME TLV can be carried and only the first occurrence is processed and any others **MUST** be ignored.



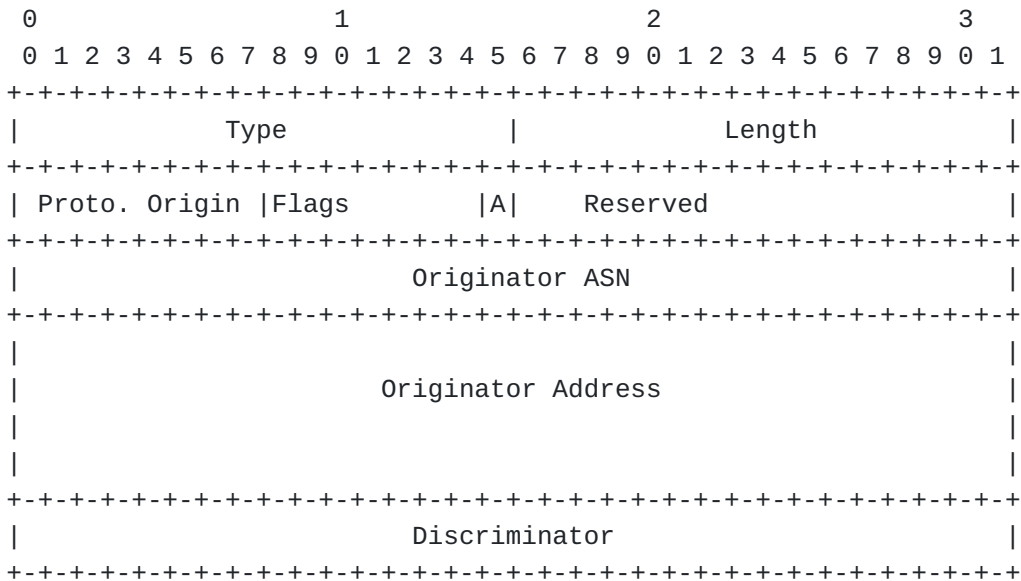
Type: TBD2 for "P2MP-SRPPOLICY-NAME" TLV.

Length: indicates the length of the value portion of the TLV in octets and MUST be greater than 0. The TLV MUST be zero-padded so that the TLV is 4-octet aligned.

P2MP SR POLICY Name : P2MP SR Policy name. It SHOULD be a string of printable ASCII characters, without a NULL terminator.

4.5.1.3. P2MP SR Policy Candidate Path ID TLV

The P2MP-SRPOLICY-CPATH-ID TLV is a mandatory TLV. Only one P2MP-SRPOLICY-CPATH-ID TLV can be carried and only the first occurrence is processed and any others MUST be ignored.



Type: TBD3 for "P2MP-SR-POLICY-CPATH-ID" TLV.

Length: 28.

Protocol Origin: 8-bit value that encodes the protocol origin, as specified in [[RFC9256](#)] Section 2.3.

Flags : A: This candidate path is active. At any instance only one candidate path can be active. PCC indicates the active candidate path to PCE through this bit. Reserved: MUST be set to zero on transmission and ignored on receipt.

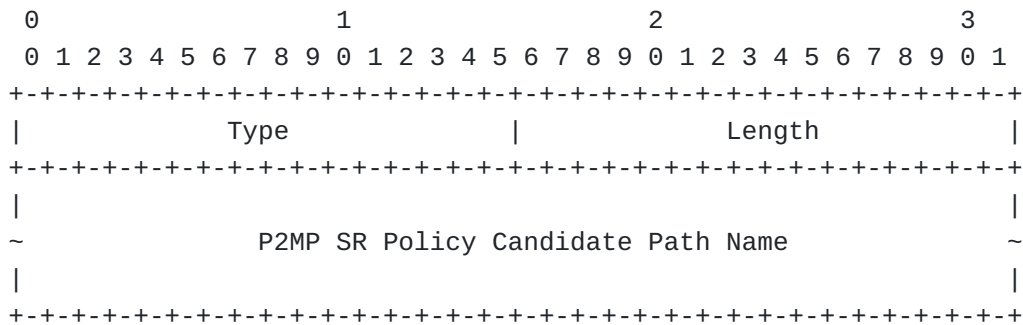
Originator ASN: Represented as 4 byte number, part of the originator identifier, as specified in [[RFC9256](#)] Section 2.4.

Originator Address: Represented as 128 bit value where IPv4 address are encoded in lowest 32 bits, part of the originator identifier, as specified in [[RFC9256](#)] Section 2.4.

Discriminator: 32-bit value that encodes the Discriminator of the candidate path.

4.5.1.4. P2MP SR Policy Candidate Path Name TLV

The P2MP -SRPOLICY-CPATH-NAME TLV is an optional TLV for the SRPAG ASSOCIATION. At most one P2MP -SRPOLICY-CPATH-NAME TLV SHOULD be encoded by the sender and only the first occurrence is processed and any others MUST be ignored.




Type: TBD4 for "P2MP-SRPOLICY-NAME" TLV.

Length: indicates the length of the value portion of the TLV in octets and MUST be greater than 0. The TLV MUST be zero-padded so that the TLV is 4-octet aligned.

P2MP SR POLICY Name : P2MP SR Policy Candidate Path name. It SHOULD be a string of printable ASCII characters, without a NULL terminator.

4.5.1.5. P2MP SR Policy Candidate Path PREFERENCE Attributes TLV

The P2MP-SRPOLICY-CPATH-PREFERENCE TLV is an optional TLV for the SRPAG Association. Only one P2MP-SRPOLICY-CPATH-PREFERENCE TLV can be carried and only the first occurrence is processed and any others MUST be ignored.



Type: TBD5 for "P2MP-SRPOLICY-CPATH-ATTR" TLV.

Length: 4.

Preference: Numerical preference of the candidate path, as specified in [[RFC9256](#)] Section 2.7.

If the TLV is missing, a default preference of 100 as specified in [[RFC9256](#)] is used.

4.5.2. P2MP-END-POINTS Object

In order for the Root to indicate operations of its leaves(Add/Remove/Modify/DoNotModify), the PC Report message is extended to include P2MP End Point <P2MP End-points> Object which is defined in [[RFC8306](#)]

The format of the PC Report message is as follow:

<Common Header>

[<SRP>]

<LSP>

[<association-list>]

[<end-points-list>]

IPV4-P2MP END-POINTS:

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Leaf type                             |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Source IPv4 address                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Destination IPv4 address               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~                                     ...                                     ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Destination IPv4 address               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

IPV6-P2MP END-POINTS:

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Leaf type                             |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     |                                     |
|                                     Source IPv6 address (16 bytes)         |
|                                     |                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     |                                     |
|                                     Destination IPv6 address (16 bytes)   |
|                                     |                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~                                     ...                                     ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     |                                     |
|                                     Destination IPv6 address (16 bytes)   |
|                                     |                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Leaf Types (derived from [[RFC8306](#)] section 3.3.2) :

1. New leaves to add (leaf type = 1)
2. Old leaves to remove (leaf type = 2)
3. Old leaves whose path can be modified/reoptimized (leaf type = 3), Future reserved not used for tree SID as of now.
4. Old leaves whose path must be left unchanged (leaf type = 4)
5. the entire pce leaf list is overwritten and replaced with the new leaf list (leaf type = 5)

A given P2MP END-POINTS object gathers the leaves of a given type. Note that a P2MP report can mix the different types of leaves by including several P2MP END-POINTS objects. The END-POINTS object body has a variable length. These are multiples of 4 bytes for IPv4, multiples of 16 bytes, plus 4 bytes, for IPv6.

4.6. P2MP Policy and Replication Segment Identifier Object and TLV

As it was mentioned previously both P2MP Policy and Replication Segment are identified via the LSP object and more precisely via the SR-P2MP-LSPID-TLV

The P2MP Policy uses the PLSP-ID to identify the Candidate Paths and the Instance-ID to identify a Path-Instance within the Candidate path.

On the other hand the Replication Segment uses the SR-P2MP-LSPID-TLV to identify and correlate a Replication Segment to a P2MP Policy

As it was noted previously on the Root, the P2MP Policy and the Replication Segment is downloaded via the same PCUpd message.

4.6.1. Extension of the LSP Object, SR-P2MP-LSPID-TLV

The LSP Object is defined in Section 7.3 of [[RFC8231](#)]. It specifies the PLSP-ID to uniquely identify an LSP that is constant for the life time of a PCEP session. Similarly, for a P2MP tunnel, the PLSP-ID identifies a Candidate Path uniquely within the P2MP policy.

The LSP Object MUST include the new SR-P2MP-POLICY-ID-TLV (IPv4/IPv6) defined in this document below. This is a variation to the P2MP object defined in [[draft-ietf-pce-stateful-pce-p2mp](#)]

SR-IPV4-P2MP-POLICY-ID TLV:

```

 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=TBD          |          Length=10          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Root              |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Tree-ID           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Instance-ID       |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

SR-IPV6-P2MP-POLICY-ID TLV :

```

 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=TBD          |          Length=22          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
+
|          Root              |
+
|          (16 octets)       |
+
|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Tree-ID           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Instance-ID       |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The type (16-bit) of the TLV is TBD (need allocation by IANA).

Root: Source Router IP Address

Tree-ID: Unique Identifier of this P2MP LSP on the Root.

Instance-ID : Contains 32 Bit instance ID.

4.7. Replication Segment

As per [[draft-ietf-spring-sr-replication-segment](#)] a replication segment has a next-hop-group which MAY contain a single outgoing replication SID or a list of SIDs (sr-policy-sid-list) In either case there needs to be a replication SID identifying the replication state on a downstream replication node. This means two replication segments can be directly connected or connected via a unicast SR domain.

4.7.1. The format of the replication segment message

The format of a Replication Segment message encoding is similar to P2MP Policy. However, the P2MP Policy contains the association object and the replication segment message does not contain the association object. In addition the replication segment uses the CCI object to identify a P2MP cross connect. The replication segment is downloaded individually to the root, transit and leaf nodes without the P2MP Policy. The P2MP Policy is a Root Concept. The replication segment uses SR-P2MP-LSPID-TLV as its identifier. The TLV is coded differently for shared and non-shared case.

*In the case of a replication segment being shared, the Tree-ID in the SR-P2MP-POLICY Identifier TLV is the replication-id of the Replication Segment and Root = 0, Instance-Id = 0. When downloading a shared replication segment from PCE through a PCInitiate message, the SR-P2MP-POLICY Identifier TLV is all 0, and on the report back from PCC, PCC generates PLSP-ID, Replication-id (Tree-id field will be populated with replication-id). Instance-id will be 0.

4.7.2. PCECC

The CCI Object as defined in [[RFC9050](#)] is used to identify a forwarding instruction in the Replication Segment. A forwarding instruction is incoming SID and a set of outgoing branches. The CCI Object-Type of 1 is used for the MPLS Label. The label in the CCI Object is the incoming SID. The outgoing SIDs are defined by the ERO Objects.

The CCI Object can be include in Reports, initiate and Update messages for Replication Segments.

The PCInitiate message defined in [[RFC8281](#)] and extended in [[RFC9050](#)] is further extended to support SR-P2MP replication segment based central control instructions.

The format of the extended PCInitiate message is as follows:

```
<PCInitiate Message> ::= <Common Header>
                           <PCE-initiated-lsp-list>
```

Where:

<Common Header> is defined in "RFC5440"

```
<PCE-initiated-lsp-list> ::= <PCE-initiated-lsp-request>
                              [<PCE-initiated-lsp-list>]
```

```
<PCE-initiated-lsp-request> ::=
    (<PCE-initiated-lsp-instantiation>|
     <PCE-initiated-lsp-deletion>|
     <PCE-initiated-lsp-central-control>)
```

```
<PCE-initiated-lsp-central-control> ::= <SRP>
                                         <LSP>
                                         (<cci-list>|
                                         (<CCI><intended-path>))
```

```
<cci-list> ::= <CCI>
               [<cci-list>]
```

```
<intended-path> ::= ((<PATH-ATTRIB><ERO>)
                     [<intended-path>])
```

Where:

<PCE-initiated-lsp-instantiation> and
<PCE-initiated-lsp-deletion> are as per
[RFC8281].

The LSP and SRP object is defined in [[RFC8231](#)]. The <intended-path> is as per [[RFC8281](#)] [[draft-ietf-pce-multipath](#)] (PATH-ATTRIB and ERO).

The format of the PCRpt message is as follows:

```
<PCRpt Message> ::= <Common Header>
                        <state-report-list>
```

Where:

```
<state-report-list> ::= <state-report>[<state-report-list>]
```

```
<state-report> ::= (<lsp-state-report>|
                    <central-control-report>)
```

```
<lsp-state-report> ::= [<SRP>]
                        <LSP>
                        <path>
```

```
<central-control-report> ::= [<SRP>]
                             <LSP>
                             (<cci-list>|
                              (<CCI><intended-path>))
```

```
<cci-list> ::= <CCI>
               [<cci-list>]
```

Where:

<path> is as per [RFC8231] and the LSP and SRP object are also defined in [RFC8231].

The <intended-path> is as per [draft-ietf-pce-multipath] (PATH- and ERO).

This document extends the use of PCUpd message with SR-P2MP CCI as follows:

```

<PCUpd Message> ::= <Common Header>
                    <update-request-list>

```

Where:

```

<update-request-list> ::= <update-request>[<update-request-list>]

```

```

<update-request> ::= (<lsp-update-request>|
                     <central-control-update>)

```

```

<lsp-update-request> ::= <SRP>
                        <LSP>
                        <path>

```

```

<central-control-update> ::= <SRP>
                             <LSP>
                             (<CCI><intended-path>)

```

Where:

<path> is as per [RFC8231] and the LSP and SRP object are also defined in [RFC8231].

The <intended-path> is as per [draft-ietf-pce-multipath] (PATH-ATT and ERO).

4.7.3. Label action rules in replicating segment

The node action and role of ingress, transit, leaf or bud, is indicated via a new Node Role TLV. This document introduces a new SR-P2MP-NODE-ROLE TLV (Type To be assigned by IANA) that will be present in the PATH-ATTRIB object.

```

  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=TBD          |          Length=4          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Role Type  |          Reserved          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

*ingress, role type = 1

*transit, role type = 2

*leaf, role type = 3

*bud, role type = 4

4.7.4. SR-ERO Rules

Forwarding information of a replication segment can be configured and steered via many different mechanisms.

As an example a replication SID can be steered via:

1. Replication SID steered with an IPv4/IPv6 directly connected nexthop

*In this case there will two SR-ERO in the ERO Object, with the Replication SID SR-ERO at the bottom and the IPv4/IPv6 SR-ERO on the top.

2. Replication SID steered with an IPv4/IPv6 loopback address that reside on the directly connected router.

*In this case there will two SR-ERO in the ERO Object, with the Replication SID SR-ERO at the bottom and the IPv4/IPv6 SR-ERO on the top.

*In addition a new flag D is added to the SR-ERO to signal that the Loopback nexthop is connected to the directly attached router.

3. Replication SID steered with unnumbered IPv4/IPv6 directly connected Interface

4. Replication SID steered via a SR adjacency or node SID

*In this case even a sid-list can be used to traffic engineer the path between two Replication Segment

*The Replication SID SR-ERO is at the bottom while the segments describing the path are on top in order.

4.7.4.1. SR-ERO subobject changes

SR-ERO from RFC 8664 is used to construct the forwarding information needed for Replication Segment.

A new D flag was added to indicate a loopback nexthop that is residing on the directly attached router. It should be noted that this flag should be set only for the loopback case and not for a local interface as a nexthop.

```

      0               1               2               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|  Type=36  |      Length      | NT  |      Flags  |D|F|S|C|M|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SID (optional)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                                     NAI (variable, optional)                                     //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Flags : F, S, C, M are already defined in rfc8664.

This document defines a new flag D: If the next-hop in NAI field is system IP or loopback, this bit indicates whether the system IP / loopback is directly connected router or not. If set indicates directly connected address. When this bit is set, F bit should be 0 (meaning NAI should be present)

5. Tree Deletion

To delete the entire tree (P2MP LSP), Root send a PCRpt message with the R bit of the LSP object set and all the fields of the SR-P2MP-LSP-ID TLV set to 0(indicating to remove all state associated with this P2MP tunnel). The PCE in response sends a PCInitiate message with R bit in the SRP object SET to all nodes along the path to indicate deletion of the entries.

6. Fragmentation

The Fragmentation bit in the LSP object (F bit) can be used to indicate a fragmented PCEP message

7. Some packet examples

7.1. Report for Leaf Add

This is an example of PCC initiated P2MP Policy. The PCC will send a Report message to the PCE to initiate a P2MP Policy with a set of leaves that are discovered via Next Generation MVPN procedures as per [\[RFC6513\]](#).

In addition, since the PCC is initiating the P2MP Policy, it does populate the PLSP-ID for the candidate path and the instance-id for the Path Instance.

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Flags = 0                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     SRP-ID-number   = 0                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| TLV Type = 28 (PathSetupType)| TLV Len = 4                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     | PST = TBD                                     |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                                <LSP OBJECT>
|          PLSP-ID = 1          |          A:1,D:1,N:1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=17          |          Length=<var>          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          symbolic path name          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=TBD          |          Length          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Root = A          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Tree-ID = Y          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Instance-ID =L1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

                                <P2MP END POINT OBJECT>
|          Leaf type = 5          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Source IPv4 address = A          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Destination IPv4 address = D          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Destination IPv4 address = E          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

7.2. P2MP Policy Candidate path Init

The following packet is the PCE creating the candidate path for the P2MP Policy and downloading the replication segment with the same message on the root.

It should be noted combining the P2MP Policy candidate path creation and replication segment only is possible on the root.

As such this message contains both association object and the CCI object. The CCI Object contains the incoming Binding SID and wraps all the Path Attribute messages and their corresponding EROs.

The PLSP-ID and Instance-ID are populated with the same ID as the previous PCC report message.

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                               Flags = 0                               |
|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                               SRP-ID-number   = 1                     |
|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| TLV Type = 28 (PathSetupType)| TLV Len = 4                           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                               | PST = TBD                             |
|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
<LSP OBJECT>
|
|           PLSP-ID = 1           |           A:1,D:1,N:1,C:0           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Type=17               |           Length=<var>               |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                               symbolic path name                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Type=TBD              |           Length                    |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Root =A               |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Tree-ID = Y           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Instance-ID = L1     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

<ASSOCIATION OBJECT>
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Reserved              |           Flags              |0|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Association type= SR-P2MP-PAG |           Association ID = z |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           IPv4 Association Source = <pce-address>                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Type=P2MP SR Policy ID   |           Length             |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Root   = A              |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           TREE-ID = 0              |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Type=P2MP SR Policy Name |           Length              |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|
~                               P2MP SR Policy Name                     ~
|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|Type=P2MP SR Pol Cand-path ID |           Length              |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|ProtOrigin 10 |           Reserved                                   |

```



```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Originator ASN                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |
|                               Originator Address = <pce-address>           |
|                               |
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Discriminator = 1                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Type=P2MP SR Pol Cand-path Name|                               Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |
~                               P2MP SR Policy Candidate Path Name           ~
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Type=P2MP SR Pol Cand-Path Pre |                               Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Preference = 100 <default>                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

                                <P2MP END POINT OBJECT>
|                               Leaf type = 5                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Source IPv4 address = A                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Destination IPv4 address = D                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Destination IPv4 address = E                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

                                <CCI OBJECT>
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               CC-ID = X                                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Reserved1           |           Flags           |0|0|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Label = 0           |           Reserved2       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

                                <PATH-ATTRIBUTES OBJECT>
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Flags                                         | Oper |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               ERO-path Id = 1                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Type                                           | Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Backup Path Count = 1                         | Flags |0|

```


7.3. Replication segment PCE Initiated on Transit and LEaves

The following packet examples shows the replication segment initiation via a PCE on transit nodes and/or leaves node.

Note:

1. This packet is generated from PCE to instantiate the replication segment, as such the PLSP-ID and Instance-ID is set to zero. PCC will assign these value and report them back to PCE.
2. This is a replication segment instantiation as such there is no association object.
3. The LSP Object Root A and Tree-ID Y associates this replication segment to the corresponding Candidate path and path instance on the root node.

there is no association object present in the packet.

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Flags = 0                                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     SRP-ID-number   = 1                                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  TLV Type = 28 (PathSetupType)|  TLV Len = 4                                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     |  PST = TBD                                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                                <LSP OBJECT>
|          PLSP-ID = 0          |          A:1,D:1,N:1,C:0          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=17          |          Length=<var>          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          symbolic path name          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=TBD          |          Length          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Root =A          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Tree-ID = Y          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Instance-ID = 0          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                                <CCI OBJECT>
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          CC-ID = X          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Reserved1          |          Flags          |0|0|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Label = d1          |          Reserved2          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                                <PATH-ATTRIBUTES OBJECT>
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Flags          | Oper|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          ERO-path Id = 1          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type          |          Length          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Backup Path Count = 1  |          Flags          |0|
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Backup Path ID 2          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Type=TBD          |          Length=4          |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```


8. Example Workflows

		+-----+		+-----+
		PCC		PCE
		Root		+-----+
		+-----+		
		PCC		
		Transit		
+-----+			---PCRpt, PLSP-ID=1, SRP-ID=1----->	PCECC LSP
PCC	+-----+		N=1, root-addr, tree-id=a,	SR-Policy
			instance-id =b,	Report to
Leaf			p2mp-end-points(LeafType=5)	PCE
+-----+				
			<--PCUpdate, PLSP-ID=1, SRP-ID =1,	Update
			root-addr, tree-id=a, instance-id=b,	CP
			p2mp-end-points, association-obj	
			-----PCRpt, PLSP-ID=1, SRP-ID = 1, ->	
			root-addr, tree-id=a, instance-id=b,	
			p2mp-end-points(LeafType=5)	
			association-object	
			<-----PCInitiate, PLSP-ID=0, -----	Download
			root-addr, tree-id=a, instance-id=0,	Leaf
			CC-ID=Z, C=0,	Replicatio
			0=0, L=z, path-attribute, ER0, SR-ER0	Segment(RS
			-----PCRpt, PLSP-ID=1----->	
			root-addr, tree-id=a, instance-id=b,	
			CC-ID=Z, Label=z, 0=0,	
			path-attribute, ER0, SR-ER0	
			<-----PCInitiate, PLSP-ID=0, -----	Download
			root-addr, tree-id=a, instance-id=0,	Transit
			CC-ID=Y, C=0,	RS
			0=0, L=y, path-attribute, ER0, SR-ER0	
			-----PCRpt, PLSP-ID=2----->	
			root-addr, tree-id=a, instance-id=c,	
			CC-ID=Y, Label=y, 0=0,	
			path-attribute, ER0, SR-ER0	
			<--PCInitiate, PLSP-ID=1,	Download
			root-addr, tree-id=a, instance-id=b,	Root
			CC-ID=X, C=0,	RS
			0=0, L=x, p2mp-end-	
			points(LeafType=5), path-attribute,	
			ER0, SR-ER0	
			-----PCRpt, PLSP-ID=1----->	

			root-addr,tree-id=a,instance-id=b,	
			CC-ID=X,Label=x,O=0,	
			p2mp-end-points(LeafType=5)	
			path-attriute,ERO,SR-ERO	
			<--PCUpdate,PLSP-ID=1, SRP-ID =2,	
			root-addr,tree-id=a,instance-id=b,	Activate
			p2mp-end-points	CP to last
				RS
			-----PCRpt,PLSP-ID=1, SRP-ID =2, ->	
			root-addr,tree-id=a,instance-id=b,	
			p2mp-end-points(LeafType=5)	

Note that on transit / leaf Initiate is with PLSP-ID = 0. Therefore PLSP-ID is locally unique to a node. It should be noted that the CC-ID does not need to be constant across all nodes that make up the path.

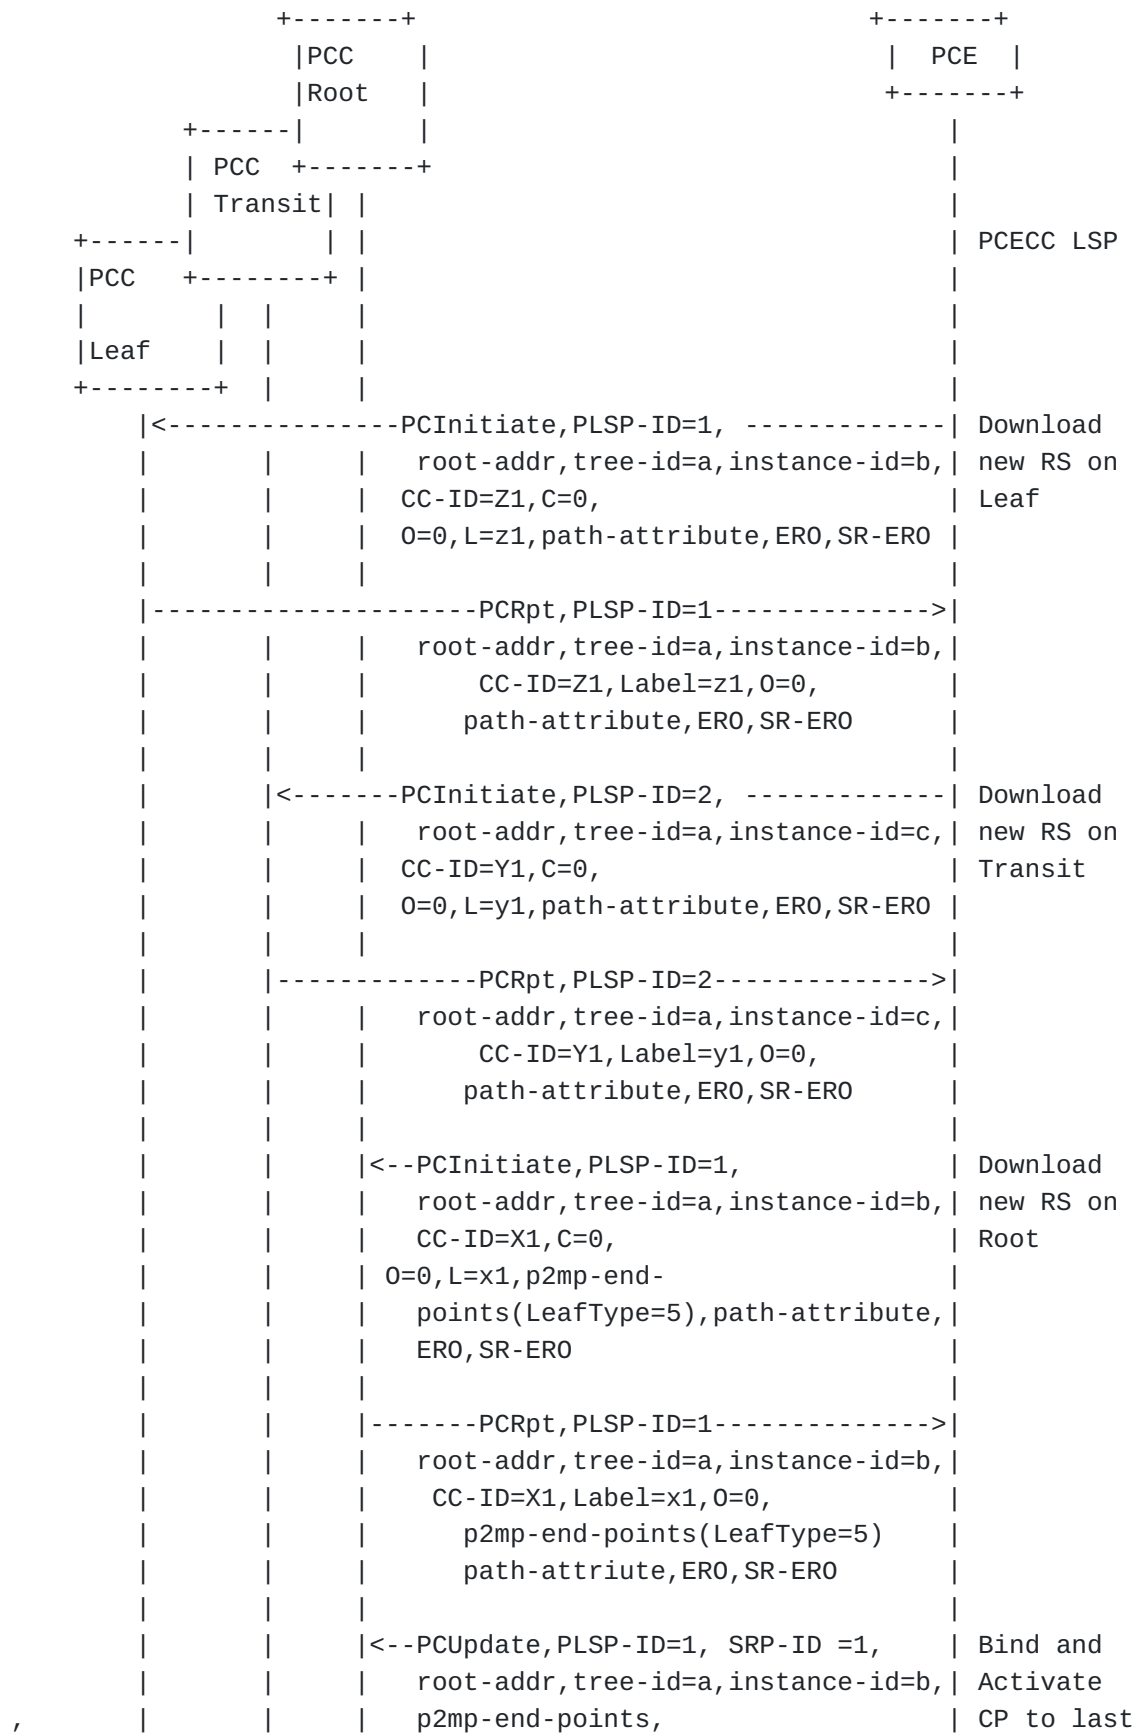
PCE-Initiated workflow

		+-----+		+-----+
		PCC		PCE
		Root		+-----+
		+-----		
		PCC +-----+		
		Transit		
+-----				PCECC LSP
PCC +-----+				
Leaf				
+-----+				
	<-----	PCInitiate, PLSP-ID=0, -----		Download
		root-addr, tree-id=a, instance-id=0,		Leaf RS
		CC-ID=Z, C=0,		
		0=0, L=z, path-attribute, ERO, SR-ERO		
	-----	PCRpt, PLSP-ID=1----->		
		root-addr, tree-id=a, instance-id=b,		
		CC-ID=Z, Label=z, 0=0,		
		path-attribute, ERO, SR-ERO		
	<-----	PCInitiate, PLSP-ID=0, -----		Download
		root-addr, tree-id=a, instance-id=0,		Transit RS
		CC-ID=Y, C=0,		
		0=0, L=y, path-attribute, ERO, SR-ERO		
	-----	PCRpt, PLSP-ID=2----->		
		root-addr, tree-id=a, instance-id=c,		
		CC-ID=Y, Label=y, 0=0,		
		path-attribute, ERO, SR-ERO		
	<--	PCInitiate, PLSP-ID=0,		Initiate
		root-addr, tree-id=0, instance-id=0,		CP
		p2mp-end-points, association-obj		
	-----	PCRpt, PLSP-ID=1, ----->		
		root-addr, tree-id=a, instance-id=b,		
		p2mp-end-points(LeafType=5)		
		association-object,		
	<--	PCInitiate, PLSP-ID=1,		Download
		root-addr, tree-id=a, instance-id=b,		Root RS
		CC-ID=X, C=0,		
		0=0, L=x, p2mp-end-		
		points(LeafType=5), path-attribute,		
		ERO, SR-ERO		
	-----	PCRpt, PLSP-ID=1----->		
		root-addr, tree-id=a, instance-id=b,		

		CC-ID=X, Label=x, O=0,	
		p2mp-end-points(LeafType=5)	
		path-attribute,ERO,SR-ERO	
		<-----PCInitiate,PLSP-ID=0, ----->	
		root-addr,tree-id=a,instance-id=0,	
		CC-ID=Y,C=0,	
		O=0,L=y,path-attribute,ERO,SR-ERO	
		-----PCRpt,PLSP-ID=2----->	
		root-addr,tree-id=a,instance-id=c,	
		CC-ID=Y,Label=y,O=0,	
		path-attribute,ERO,SR-ERO	
		<--PCUpdate,PLSP-ID=1, SRP-ID =1,	Bind and
		root-addr,tree-id=a,instance-id=b,	Activate
		p2mp-end-points,	CP to last
			RS
		-----PCRpt,PLSP-ID=1, SRP-ID = 1,->	
		root-addr,tree-id=a,instance-id=b,	
		p2mp-end-points(LeafType=5)	

MBB Workflow:

Common (PCE-INIT, PCC-INIT) MBB



				RS
			-----PCRpt, PLSP-ID=1, SRP-ID = 1, ->	
			root-addr, tree-id=a, instance-id=b,	
			p2mp-end-points(LeafType=5)	
			<--PCInitiate, PLSP-ID=1, R=1	Remove
			root-addr, tree-id=a, instance-id=b,	the old RS
			CC-ID=X1, C=0	from Leaf
			O=0, L=x1, p2mp-end-	
			points(LeafType=5), path-attribute,	
			ERO, SR-ERO	
			-----PCRpt, PLSP-ID=1, R=1----->	
			root-addr, tree-id=a, instance-id=b,	
			CC-ID=X1, Label=x1, O=0,	
			p2mp-end-points(LeafType=5)	
			path-attribute, ERO, SR-ERO	
			<-----PCInitiate, PLSP-ID=2, R=1----->	Remove the
			root-addr, tree-id=a, instance-id=c,	old RS fro
			CC-ID=Y1, C=0,	Transit
			O=0, L=y1, path-attribute, ERO, SR-ERO	
			-----PCRpt, PLSP-ID=2, R=1----->	
			root-addr, tree-id=a, instance-id=c,	
			CC-ID=Y1, Label=y1, O=0,	
			path-attribute, ERO, SR-ERO	
			<-----PCInitiate, PLSP-ID=1, R=1----->	Remove the
			root-addr, tree-id=a, instance-id=b,	old RS fro
			CC-ID=Z1, C=0,	Root
			O=0, L=z1, path-attribute, ERO, SR-ERO	
			-----PCRpt, PLSP-ID=1, R=1----->	
			root-addr, tree-id=a, instance-id=b,	
			CC-ID=Z1, Label=z1, O=0,	
			path-attribute, ERO, SR-ERO	

9. IANA Consideration

1. This draft extends the PCEP OPEN object by defining an optional TLV to indicate the PCE's capability to perform SR-P2MP path computations with a new IANA capability type (TBD).
2. PCEP open object with a new association type " P2MP SR Policy Association " value (TBD).
3. A new Association type. Association type = TBD1 "P2MP SR Policy Association Type" for SR Policy Association Group (P2MP SRPAG)
 1. Five new TLVs are identified to carry association information: P2MP-SRPOLICY-ID TLV, P2MP-SRPOLICY-NAME TLV, P2MP-SRPOLICY-CPATH-ID TLV, P2MP-SRPOLICY-CPATH-PREFERENCE TLV, P2MP-SRPOLICY-CPATH-NAME TLV
4. Two new TLVs for Identifying the P2MP Policy and the Replication segment SR-IPV4-P2MP-POLICY-ID TLV and SR-IPV6-P2MP-POLICY-ID TLV
5. A new SR-P2MP-NODE-ROLE TLV (Type To be assigned by IANA) that will be present in the PATH-ATTRIB object

10. Security Considerations

TBD

11. Acknowledgments

The authors would like to thank Tanmoy Kundu and Stone Andrew at Nokia for their feedback and major contribution to this draft.

12. Informative References

- [draft-hb-spring-sr-p2mp-policy-yang] "H.Bidgoli, D. Yoyer, R. Parekh, T.Saad, T.kundu "YANG Data Model for p2mp sr policy"", October 2020.
- [draft-ietf-pce-multipath] "M. Koldychev, S. Sivabalan, T. Saad, H. Bidgoli, B. Yadav, S. Peng, G. Mishra "PCEP Extensions for Signaling Multipath Information"", November 2022.
- [draft-ietf-pce-segment-routing-policy-cp] "M. Koldychev, S. Sivabalan, C. Barth, S. Peng, H. Bidgoli "PCEP extension to support Segment Routing Policy Candidate Paths", October 2022.
- [draft-ietf-pce-stateful-pce-p2mp] "U. Palle, D. Dhody, Y.Tanaka, V. Beeram "Protocol Extensions for Stateful PCE usage for

Point-to-Multipoint Traffic Engineering Label"", April 2019.

[draft-ietf-pim-sr-p2mp-policy] "D. Yoyer, C. Filsfils, R.Prekh, H.bidgoli, Z. Zhang, "Segment Routing Point-to-Multipoint Policy"", October 2019.

[draft-ietf-spring-sr-replication-segment] "D. Yoyer, C. Filsfils, R.Prekh, H.bidgoli, Z. Zhang, "SR Replication Segment for Multi-point Service Delivery"", July 2020.

[RFC2119] ""Key words for use in RFCs to Indicate Requirement Levels"", March 1997.

[RFC3209] "D. Awduche, L. Berger, T. Li, V. Srinivasan, G. Swallow "RSVP-TE: Extensions to RSVP for LSP Tunnels"", December 2001.

[RFC5440] "JP. Vasseur, JL. Le Roux "Path Computation Element Communication Protocol"", March 2009.

[RFC6513] "E. Rosen, R. Aggerwal "Multicast in MPLS/BGP IP VPNs"", February 2012.

[RFC8231] "E. Crabbe, I. Minei, J. Medved, R. Varga "PCEP Extensions for Stateful PCE"", September 2017.

[RFC8281] "E. Crabbe, I. Minei, S. Sivabalan, R. Varga "PCEP Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model"", December 2017.

[RFC8306] "Q. Zhao, D. Dhody, R. Palleti, D.King "PCEP for Point-to-Multipoint Traffic Engineering Label Switched Paths"", November 2017.

[RFC8664] "S. Sivabalan, C. Filsfils, J. Tantsura, W. Henderickx, J. Hardwick "PCEP Extensions for Segment Routing"", December 2019.

[RFC8697] "I. Minei, E. Crabbe, S. Sivabalan, H. Ananthakrishnan, D. Dhody, Y. Tanaka "PCEP Extensions for Establishing Relationships between Sets of LSPs"", January 2020.

[RFC9050] "Z. Li, S. Peng, M. Negi, Q. Zhao, C. Zhou "PCEP Procedures and Extensions for Using the PCECC"", July 2021.

[RFC9256] "C. Filsfils, K. Talaulikar, D. Voyer, A. Bogdanov, P. Mattes "Segment Routing Policy Architecture"", July 2022.

Authors' Addresses

Hooman Bidgoli (editor)
Nokia
Ottawa
Canada

Email: hooman.bidgoli@nokia.com

Daniel Voyer
Bell Canada
Montreal
Canada

Email: daniel.voyer@bell.ca

Saranya Rajarathinam
Nokia
Mountain View,
United States of America

Email: saranya.Rajarathinam@nokia.com

Anuj Budhiraja
Cisco System
San Jose,
United States of America

Email: abudhira@cisco.com

Tarek Saad
Juniper Networks
Ottawa
Canada

Email: tsaad@juniper.com

Siva Sivabalan
Ciena
Ottawa
Canada

Email: ssivabal@ciena.com