

PCE Working Group
Internet-Draft
Updates: 8231, 8281 (if approved)
Intended status: Standards Track
Expires: 12 July 2026

M. Koldychev
S. Sivabalan
Ciena Corporation
T. Saad
Cisco Systems
V. Beeram
Juniper Networks, Inc.
H. Bidgoli
Nokia
S. Peng
Huawei Technologies
S. Sidor, Ed.
Cisco Systems.
8 January 2026

Path Computation Element Communication Protocol (PCEP) Extensions for
Signaling Multipath Information
draft-ietf-pce-multipath-18

Abstract

Certain traffic engineering path computation problems require solutions that consist of multiple traffic paths that together form a solution. Returning a single traffic path does not provide a valid solution. This document defines mechanisms to encode multiple paths for a single set of objectives and constraints. This allows encoding of multiple Segment Lists per Candidate Path within a Segment Routing Policy. The new Path Computation Element Communication Protocol (PCEP) mechanisms are designed to be generic, where possible, to allow for future re-use outside of SR Policy. The new PCEP mechanisms are applicable to both stateless and stateful PCEP. Additionally, this document updates RFC 8231 and RFC 8281 to allow encoding of multiple Segment Lists in PCEP.

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 12 July 2026.

Copyright Notice

Copyright (c) 2026 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. Requirements Language	4
1.2. Terminology	4
2. Motivation	4
2.1. Signaling Multiple Segment Lists of an SR Candidate Path	4
2.2. Splitting of Requested Bandwidth	4
2.3. Reverse Path Information	5
3. Protocol Extensions	5
3.1. PATH-ATTRIB Object	5
3.2. METRIC Object	6
3.3. MULTIPATH-WEIGHT TLV	6
3.4. MULTIPATH-BACKUP TLV	7
3.5. MULTIPATH-OPPDIR-PATH TLV	8
3.6. Composite Candidate Path	10
3.6.1. Per-Flow Candidate Path	10
4. Operation	11
4.1. Capability Negotiation	11
4.1.1. Multipath Capability TLV	11
4.2. Path ID	13
4.3. Signaling Multiple Paths for Loadbalancing	13
4.4. Signaling Multiple Paths for Protection	13
5. PCEP Message Extensions	14
6. Examples	15
6.1. SR Policy Candidate Path with Multiple Segment Lists	15

6.2.	Two Primary Paths Protected by One Backup Path	16
6.3.	Composite Candidate Path	17
6.4.	Opposite Direction Tunnels	18
7.	Implementation Status	20
7.1.	Cisco Systems	20
7.2.	Ciena Corp	20
7.3.	Huawei Technologies	20
8.	IANA Considerations	21
8.1.	PCEP Object	21
8.2.	PCEP TLV	21
8.3.	PCEP-Error Object	22
8.4.	Flags in the MULTIPATH-CAP TLV	22
8.5.	Flags in the PATH-ATTRIB Object	23
8.6.	Flags in the MULTIPATH-BACKUP TLV	23
8.7.	Flags in the MULTIPATH-OPPDIR-PATH TLV	24
9.	Security Considerations	24
10.	Manageability Considerations	24
10.1.	Control of Function and Policy	25
10.2.	Information and Data Models	25
10.3.	Liveness Detection and Monitoring	25
10.4.	Verify Correct Operations	25
10.5.	Requirements On Other Protocols	25
10.6.	Impact On Network Operations	25
11.	Acknowledgement	26
12.	Contributors	26
13.	References	26
13.1.	Normative References	26
13.2.	Informative References	28
	Authors' Addresses	29

1. Introduction

Segment Routing Policy for Traffic Engineering [RFC9256] details the concepts of Segment Routing (SR) Policy and approaches to steering traffic into an SR Policy. In particular, it describes the SR Candidate Path as a collection of one or more Segment Lists. The current PCEP standards only allow for signaling of one Segment List per Candidate Path. The PCEP extension to support Segment Routing Policy Candidate Paths [RFC9862] specifically avoids defining how to signal multiple Segment Lists.

This document defines the required extensions that allow the signaling of multipath information via PCEP. Although these extensions are motivated by the SR Policy use case, they are also applicable to other data plane types.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

The following terms are used in this document:

ECMP:

Equal Cost Multi Path, equally distributing traffic among multiple paths/links, where each path/link gets the same share of traffic as others.

W-ECMP:

Weighted ECMP, unequally distributing traffic among multiple paths/links, where some paths/links get more traffic than others.

2. Motivation

This extension is motivated by the use-cases described below.

2.1. Signaling Multiple Segment Lists of an SR Candidate Path

The Candidate Path of an SR Policy is the unit of signaling in PCEP, see [RFC9862]. Each Candidate Path can contain multiple Segment Lists and each Segment List is encoded by one Explicit Route Object (ERO). However, each PCEP Label Switched Path (LSP) can contain only a single ERO, which prevents the encoding of multiple Segment Lists within the same SR Candidate Path.

2.2. Splitting of Requested Bandwidth

A Path Computation Client (PCC) may request a path with 80 Gbps of bandwidth, but all links in the network have only 60 Gbps capacity. The Path Computation Element (PCE) can return two paths, that can together carry 80 Gbps. The PCC can then equally or unequally split the incoming 80 Gbps of traffic among the two paths. Section 3.3 introduces a new TLV that carries the path weight that facilitates control of load-balancing of traffic among the multiple paths.

2.3. Reverse Path Information

Path Computation Element Communication Protocol (PCEP) Extensions for Associated Bidirectional LSPs [RFC9059] defines a mechanism in PCEP to associate two opposite direction SR Policy Candidate Paths. However, within each Candidate Path there can be multiple Segment Lists, and [RFC9059] does not define a mechanism to specify mapping between Segment Lists of the forward and reverse Candidate Paths. Certain applications such as Circuit Style SR Policy [I-D.ietf-spring-cs-sr-policy], require the knowledge of reverse path(s) per Segment List, not just per Candidate Path. For example, when the headend knows the reverse Segment List for each forward Segment List, then Performance Measurement (PM)/Bidirectional Forwarding Detection (BFD) can run a separate session on every Segment List, by imposing a double stack (forward stack followed by reverse stack) onto the packet. If the reverse Segment List is co-routed with the forward Segment List, then the PM/BFD session would traverse the same links in the forward and reverse directions, thus allowing detection of link/node failures in both directions.

3. Protocol Extensions

3.1. PATH-ATTRIB Object

This document defines the PATH-ATTRIB object that is used to carry per-path information and to act as a separator between several ERO/Recorded Route Object (RRO) objects in the <intended-path>/<actual-path> Routing Backus-Naur Form (RBNF) [RFC5511] element. The PATH-ATTRIB object always precedes the ERO/RRO that it applies to. If multiple ERO/RRO objects are present, then each ERO/RRO object MUST be preceded by an PATH-ATTRIB object that describes it.

The PATH-ATTRIB Object-Class value is 45.

The PATH-ATTRIB Object-Type value is 1.

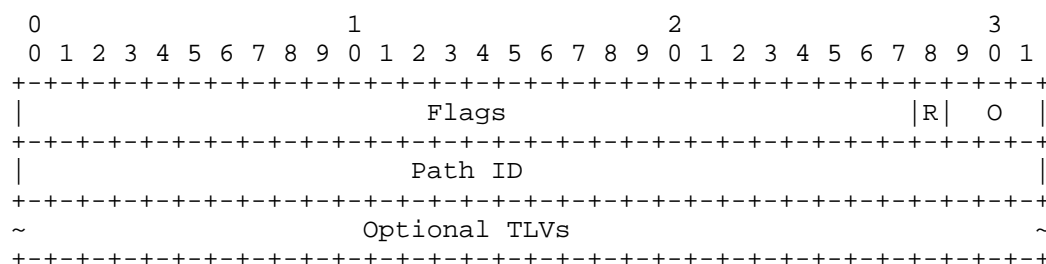


Figure 1: PATH-ATTRIB object format

Flags (32 bits):

- * O (Operational - 3 bits): operational state of the path, same values as the identically named field in the LSP object [RFC8231].
- * R (Reverse - 1 bit): Indicates this path is reverse, i.e., it originates on the LSP destination and terminates on the LSP source (usually the PCC headend itself). Paths with this flag set serve only informational purpose to the PCC.
- * Unassigned bits MUST be set to 0 on transmission and MUST be ignored on receipt.

Path ID (32 bits): 4-octet identifier that identifies a path (encoded in the ERO/RRO) within the set of multiple paths under the PCEP LSP. See Section 4.2 for details.

3.2. METRIC Object

The PCEP METRIC object can continue to be used at the LSP level. The metric value encoded into the LSP level METRIC object SHOULD be the maximum value of all the per PATH metrics. Per-path metrics are outside the scope of this document and would require further extensions.

3.3. MULTIPATH-WEIGHT TLV

New MULTIPATH-WEIGHT TLV is optional in the PATH-ATTRIB object.

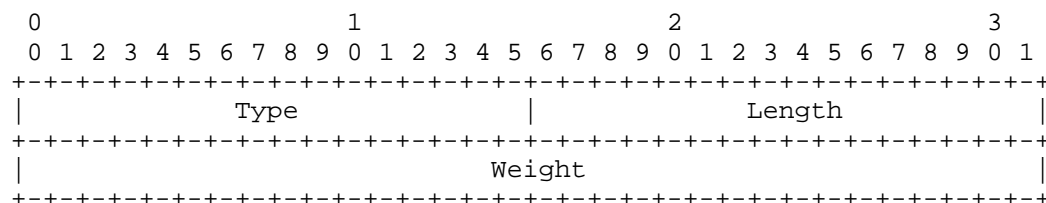


Figure 2: MULTIPATH-WEIGHT TLV format

Type (16 bits): 61 for "MULTIPATH-WEIGHT" TLV.

Length (16 bits): 4.

Weight (32 bits): weight of this path within the multipath, if W-ECMP is desired. The fraction of flows a specific ERO/RRO carries is derived from the ratio of its weight to the sum of the weights of all other paths.

When the MULTIPATH-WEIGHT TLV is absent from the PATH-ATTRIB object, or the PATH-ATTRIB object is absent from the <intended-path>/<actual-path>, then the Weight of the corresponding path is taken to be 1.

3.4. MULTIPATH-BACKUP TLV

New MULTIPATH-BACKUP TLV is optional in the PATH-ATTRIB object.

This TLV is used to specify protecting standby path(s), for each ECMP path within a PCEP LSP. This is similar to path protection, but works at the ECMP path level instead of at the PCEP LSP level.

This functionality is not part of the SR Policy Architecture [RFC9256], but is something optional that may be implemented for certain specialized use cases. One such use case is the Point-to-Multipoint (P2MP) SR Policy [I-D.draft-ietf-pce-sr-p2mp-policy].

Support for the MULTIPATH-BACKUP TLV is currently defined only for P2MP paths. Support for Point-to-Point (P2P) paths is out of scope for this document. If needed in the future, support for P2P paths using the MULTIPATH-BACKUP TLV can be defined in future documents. Future documents that extend this TLV to support P2P paths SHOULD also define explicit capability exchange mechanisms to allow PCEP peers to negotiate support for MULTIPATH-BACKUP with P2P paths.

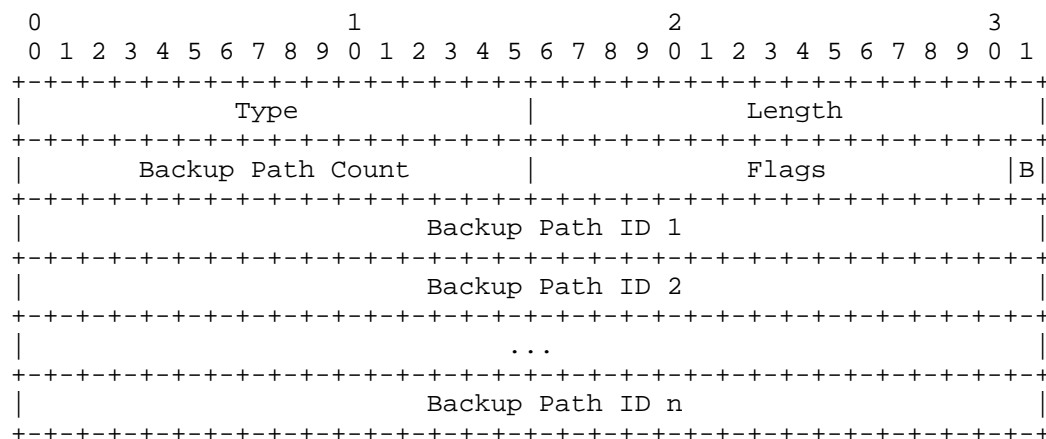


Figure 3: MULTIPATH-BACKUP TLV format

Type (16 bits): 62 for "MULTIPATH-BACKUP" TLV

Length (16 bits): 4 + (N * 4) (where N is the Backup Path Count)

Backup Path Count (16 bits): Number of backup path(s).

Flags (16 bits):

- * B: If set, indicates a pure backup path. This is a path that only carries rerouted traffic after the protected path fails. If this flag is not set, or if the MULTIPATH-BACKUP TLV is absent, then the path is assumed to be primary that carries normal traffic.
- * Unassigned bits MUST be set to 0 on transmission and MUST be ignored on receipt.

Backup Path ID(s): a series of 4-octet identifier(s) that identify the backup path(s) in the set that protect this primary path.

If a PCEP speaker receives a MULTIPATH-BACKUP TLV applied to a P2P path, it SHOULD reject the path and send a PCError message with Error-Type = 19 ("Invalid Operation") and Error-Value = 20 ("Not supported path backup").

3.5. MULTIPATH-OPPPATH TLV

New MULTIPATH-OPPPATH TLV is optional in the PATH-ATTRIB object. Multiple instances of the TLV are allowed in the same PATH-ATTRIB object. This TLV encodes a many-to-many mapping between forward and reverse paths.

Many-to-many mapping means that a single forward path MAY map to multiple reverse paths and conversely that a single reverse path MAY map to multiple forward paths. Many-to-many mapping can happen for an SR Policy, when a Segment List contains Node Segment(s) which traverse parallel links at the midpoint. The reverse of this Segment List may not be able to be expressed as a single Reverse Segment List, but requires multiple Reverse Segment Lists to cover all the parallel links at the midpoint.

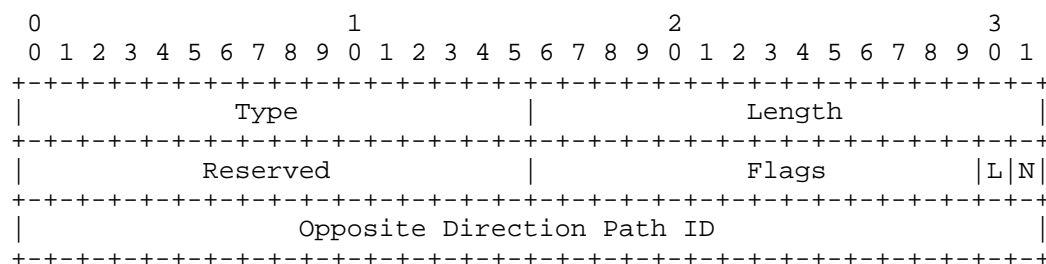


Figure 4: MULTIPATH-OPPPATH TLV format

Type (16 bits): 63 for "MULTIPATH-OPPPATH" TLV

Length (16 bits): 16.

Reserved: This field MUST be set to zero on transmission and MUST be ignored on receipt.

Flags (16 bits):

- * N (Node co-routed): If set, indicates this path is node co-routed with its opposite direction path, specified in this TLV. Two opposite direction paths are node co-routed if they traverse the same nodes, but MAY traverse different links.
- * L (Link co-routed): If set, indicates this path is link co-routed with its opposite directions path, specified in this TLV. Two opposite direction paths are link co-routed if they traverse the same links (but in opposite directions).
- * Unassigned bits MUST be set to 0 on transmission and MUST be ignored on receipt.

Opposite Direction Path ID (32 bits): Identifies a path that goes in the opposite direction to this path. If no such path exists, then this field MUST be set to 0, a value reserved to indicate the absence of a Path ID.

Multiple instances of this TLV present in the same PATH-ATTRIB object indicate that there are multiple opposite-direction paths corresponding to the given path. This allows for many-to-many relationship among the paths of two opposite direction LSPs.

Whenever path A references another path B as being the opposite-direction path, then path B MUST also reference path A as its own opposite-direction path. Furthermore, their values of the R-flag (Reverse) in the PATH-ATTRIB object MUST have opposite values. If a PCEP speaker receives an opposite-direction path mapping that is asymmetric or where the R-flags are inconsistent, it MUST treat this as an error. The PCEP speaker MUST send a PCErrror message with Error-Type = 19 ("Invalid Operation") and Error-Value = TBD4 ("Invalid opposite-direction path mapping").

See Section 6.4 for an example of usage.

3.6. Composite Candidate Path

SR Policy Architecture [RFC9256] defines the concept of a Composite Candidate Path. A regular SR Policy Candidate Path outputs traffic to a set of Segment Lists, while an SR Policy Composite Candidate Path outputs traffic recursively to a set of SR Policies on the same headend. In PCEP, the Composite Candidate Path still consists of PATH-ATTRIB objects, but ERO is replaced by Color of the recursively used SR Policy.

To signal the Composite Candidate Path, we make use of the COLOR TLV, defined in [RFC9863]. For a Composite Candidate Path, the COLOR TLV is included in the PATH-ATTRIB Object, thus allowing each Composite Candidate Path to do ECMP/W-ECMP among SR Policies identified by its constituent Colors. Only one COLOR TLV MUST be included into the PATH-ATTRIB object. If multiple COLOR TLVs are contained in the PATH-ATTRIB object, only the first one MUST be processed and the others MUST be ignored.

An ERO object MUST be included as per the existing RBNF, this ERO MUST contain no sub-objects. This empty ERO serves as a placeholder to maintain compatibility with existing implementations based on the RBNF defined in [RFC8231]. If the head-end receives a non-empty ERO for a Composite Candidate Path, it MUST send a PCError message with Error-Type = 19 ("Invalid Operation") and Error-Value = 21 ("Non-empty path").

See Section 6.3 for an example of the encoding.

3.6.1. Per-Flow Candidate Path

Per-Flow Candidate Path builds on top of the concept of the Composite Candidate Path. Each Path in a Per-Flow Candidate Path is assigned a 3-bit forward class value, which allows Quality of Service (QoS) classified traffic to be steered depending on the forward class.

New MULTIPATH-FORWARD-CLASS TLV is optional in the PATH-ATTRIB object.

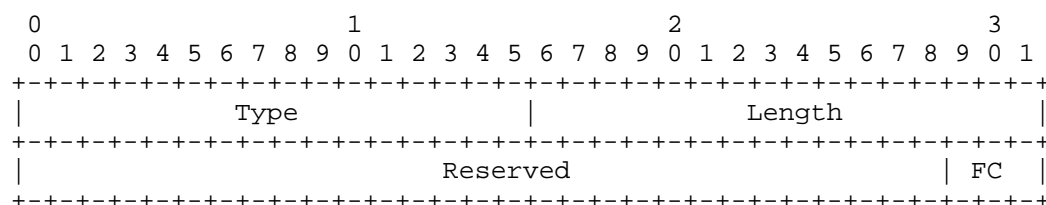


Figure 5: MULTIPATH-FORWARD-CLASS TLV format

Type (16 bits): TBD1 for "MULTIPATH-FORWARD-CLASS" TLV.

Length (16 bits): 4.

Reserved: This field MUST be set to zero on transmission and MUST be ignored on receipt.

FC (3 bits): Forward class value that is given by the QoS classifier to traffic entering the given Candidate Path. Different classes of traffic that enter the given Candidate Path can be differentially steered into different Colors.

4. Operation

4.1. Capability Negotiation

4.1.1. Multipath Capability TLV

New MULTIPATH-CAP TLV is defined. This TLV MAY be present in the OPEN object during PCEP session establishment.

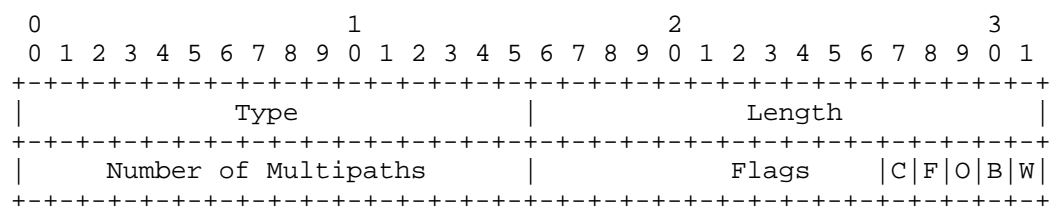


Figure 6: MULTIPATH-CAP TLV format

Type (16 bits): 60 for "MULTIPATH-CAP" TLV.

Length (16 bits): 4.

Number of Multipaths (16 bits): When sent from a PCC, it indicates how many multipaths the PCC can install in forwarding. From a PCE, it indicates how many multipaths the PCE can compute. The value 255 indicates an unlimited number. The value 0 is reserved.

Flags (16 bits):

- * W-flag: whether MULTIPATH-WEIGHT TLV is supported.
- * B-flag: whether MULTIPATH-BACKUP TLV is supported.

- * O-flag: whether MULTIPATH-OPPDIR-PATH TLV is supported and requested. If this flag is set, the PCE SHOULD tell the PCC the reverse path information, if it is able to.
- * F-flag: whether MULTIPATH-FORWARD-CLASS TLV is supported.
- * C-flag: whether Composite Candidate Path (Section 3.6) is supported.
- * Unassigned bits MUST be set to 0 on transmission and MUST be ignored on receipt.

Note that F-flag and C-flag can be set independently, i.e., F-flag can be set, but C-flag not set, etc.

When PCE computes the LSP path, it MUST NOT return more forward multipaths than the corresponding value of "Number of Multipaths" from the MULTIPATH-CAP TLV. If this TLV is absent (from both OPEN and LSP objects), then the "Number of Multipaths" is assumed to be 1.

From the PCC, the MULTIPATH-CAP TLV MAY also be present in the LSP object for each individual LSP, to specify per-LSP values. The PCC MUST NOT include this TLV in the LSP object if the TLV was not present in the OPEN objects of both PCEP peers. TLV values in the LSP object override the session default values in the OPEN object. If a PCEP speaker receives a PATH-ATTRIB object but the multipath capability was not successfully negotiated during session establishment, it MUST treat this as an error. The PCEP speaker MUST send a PCError message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD2 ("Unexpected PATH-ATTRIB object").

Additionally, if a PCEP speaker receives a TLV within the PATH-ATTRIB object (such as MULTIPATH-WEIGHT, MULTIPATH-BACKUP, MULTIPATH-OPPDIR-PATH, or MULTIPATH-FORWARD-CLASS) but the corresponding capability flag was not set in the negotiated MULTIPATH-CAP TLV, it MUST treat this as an error. The PCEP speaker MUST send a PCError message with Error-Type = 19 ("Invalid Operation") and Error-Value = TBD3 ("Unsupported multipath capability").

For example, the PCC includes this TLV in the OPEN object at session establishment, setting "Number of Multipaths" to 4 and "O-flag" to 0. The PCC also includes this TLV in the LSP object for a particular LSP, setting "Number of Multipaths" to 16 and "O-flag" to 1. This indicates that the PCC only wants to receive the reverse path information for that particular LSP and that this LSP can have up to 16 multipaths, while other LSPs can only have up to 4 multipaths.

4.2. Path ID

The Path ID uniquely identifies a Path within the context of an LSP. Note that when the LSP is an SR Policy Candidate Path, the Paths within that LSP are the Segment Lists.

Value 0 indicates an unallocated Path ID. The value of 0 MAY be used when this Path is not referenced and the allocation of a Path ID is not necessary.

Path IDs are allocated by the PCEP peer that owns the LSP. If the LSP is delegated to the PCE, then the PCE allocates the Path IDs and sends them in the PCReply/PCUpd/PCInitiate messages. If the LSP is locally computed on the PCC, then the PCC allocates the Path IDs and sends them in the PCReq/PCRpT messages.

If a PCEP speaker detects that there are two Paths with the same Path ID, then the PCEP speaker MUST send PCErr message with Error-Type = 1 ("Reception of an invalid object") and Error-Value = 38 ("Conflicting Path ID").

4.3. Signaling Multiple Paths for Loadbalancing

The PATH-ATTRIB object can be used to signal multiple path(s) and indicate (un)equal loadbalancing amongst the set of multipaths. In this case, the PATH-ATTRIB is populated for each ERO as follows:

1. The PCE MAY assign a unique Path ID to each ERO path and populate it inside the PATH-ATTRIB object. The Path ID is unique within the context of a PLSP (when non-zero).
2. The MULTIPATH-WEIGHT TLV MAY be carried inside the PATH-ATTRIB object. A weight is populated to reflect the relative loadshare that is to be carried by the path. If the MULTIPATH-WEIGHT is not carried inside a PATH-ATTRIB object, the default weight 1 MUST be assumed when computing the loadshare.
3. The fraction of flows carried by a specific primary path is derived from the ratio of its weight to the sum of all other multipath weights.

4.4. Signaling Multiple Paths for Protection

The PATH-ATTRIB object can be used to describe a set of backup path(s) protecting a primary path within a PCEP LSP. This capability is currently defined only for P2MP paths. Support for P2P paths with the MULTIPATH-BACKUP TLV is out of scope for this document. In this case, the PATH-ATTRIB is populated for each ERO as follows:

1. The PCE assigns a unique Path ID to each ERO path and populates it inside the PATH-ATTRIB object. The Path ID is unique within the context of a PLSP.
2. The MULTIPATH-BACKUP TLV MAY be added inside the PATH-ATTRIB object for each ERO that is protected. The backup path ID(s) are populated in the MULTIPATH-BACKUP TLV to reflect the set of backup paths protecting the primary path. The Length field and Backup Path Count in the MULTIPATH-BACKUP are updated according to the number of backup path ID(s) included.
3. The MULTIPATH-BACKUP TLV MAY be added inside the PATH-ATTRIB object for each ERO that is unprotected. In this case, MULTIPATH-BACKUP does not carry any backup path IDs in the TLV. If the path acts as a pure backup i.e., the path only carries rerouted traffic after the protected path(s) fail then the B flag MUST be set.

Primary paths which do not include the MULTIPATH-BACKUP TLV are assumed to be protected by all the backup paths. I.e., omitting the TLV is equivalent to including the TLV with all the backup path IDs filled in.

Note that a given PCC may not support certain backup combinations, such as a backup path that is itself protected by another backup path, etc. If a PCC does not support a requested backup scenario, the PCC MUST send a PCErr message with Error-Type = 19 ("Invalid Operation") and Error-Value = 20 ("Not supported path backup"). Additionally, if a P2P path is sent with a MULTIPATH-BACKUP TLV, the PCC or PCE SHOULD reject it with the same PCErr as above.

5. PCEP Message Extensions

The RBNF of PCRpt and PCUpd messages, as defined in [RFC8231], use a combination of <intended-path> and/or <actual-path>. PCReq and PCRep messages, as defined in [RFC5440] and extended by [RFC8231], directly include ERO and RRO objects within their respective message structures rather than encapsulating them within <intended-path> or <actual-path> constructs. As specified in Section 6.1 of [RFC8231], within the context of messages that use these constructs, <intended-path> is represented by the ERO object and <actual-path> is represented by the RRO object:

<intended-path> ::= <ERO>

<actual-path> ::= <RRO>

This document updates [RFC8231] to allow multiple ERO/RRO objects to be present in the <intended-path>/<actual-path>:

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

```
<actual-path> ::= (<RRO>|
                  (<PATH-ATTRIB><RRO>)
                  [<actual-path>])
```

Similarly, this document updates [RFC8281] to allow multiple paths in the PCInitiate message by allowing multiple ERO objects with their associated path attributes. The PCE-initiated LSP instantiation format is updated to:

```
<PCE-initiated-lsp-instantiation> ::= <SRP>
                                     <LSP>
                                     [<END-POINTS>]
                                     <intended-path>
                                     [<attribute-list>]
```

where <intended-path> follows the recursive definition above, allowing multiple paths to be signaled in a single PCInitiate message. Each path is preceded by a PATH-ATTRIB object that describes it.

6. Examples

6.1. SR Policy Candidate Path with Multiple Segment Lists

Consider the following sample SR Policy, taken from [RFC9256].

```
SR policy POL1 <headend, color, endpoint>
  Candidate Path CP1 <protocol-origin = 20, originator =
    100:1.1.1.1, discriminator = 1>
    Preference 200
    Weight W1, SID-List1 <SID11...SID1i>
    Weight W2, SID-List2 <SID21...SID2j>
  Candidate Path CP2 <protocol-origin = 20, originator =
    100:2.2.2.2, discriminator = 2>
    Preference 100
    Weight W3, SID-List3 <SID31...SID3i>
    Weight W4, SID-List4 <SID41...SID4j>
```

As specified in [RFC9862], CP1 and CP2 are signaled as separate state-report elements and each has a unique PLSP-ID, assigned by the PCC. For this example, PLSP-ID 100 is assigned to CP1 and PLSP-ID 200 to CP2.

The state-report for CP1 can be encoded as:

```
<state-report> =  
  <LSP PLSP-ID=100>  
  <ASSOCIATION>  
  <END-POINT>  
  <PATH-ATTRIB Path ID=1 <WEIGHT-TLV Weight=W1>>  
  <ERO SID-List1>  
  <PATH-ATTRIB Path ID=2 <WEIGHT-TLV Weight=W2>>  
  <ERO SID-List2>
```

The state-report for CP2 can be encoded as:

```
<state-report> =  
  <LSP PLSP-ID=200>  
  <ASSOCIATION>  
  <END-POINT>  
  <PATH-ATTRIB Path ID=1 <WEIGHT-TLV Weight=W3>>  
  <ERO SID-List3>  
  <PATH-ATTRIB Path ID=2 <WEIGHT-TLV Weight=W4>>  
  <ERO SID-List4>
```

The above sample state-report elements only specify the minimum mandatory objects, of course other objects like SRP, LSPA, METRIC, etc., are allowed to be inserted.

Note that the syntax

```
<PATH-ATTRIB Path ID=1 <WEIGHT-TLV Weight=W1>>
```

means that this is PATH-ATTRIB object with Path ID field set to 1 and with a MULTIPATH-WEIGHT TLV carrying weight of "W1".

6.2. Two Primary Paths Protected by One Backup Path

Suppose there are 3 paths: A, B, C. Where A and B are primary and C is to be used only when A or B fail. Suppose the Path IDs for A, B, C are respectively 1, 2, 3. This would be encoded in a state-report as:

```

<state-report> =
  <LSP>
  <ASSOCIATION>
  <END-POINT>
  <PATH-ATTRIB Path ID=1 <BACKUP-TLV B=0, Backup_Paths=[3]>>
  <ERO A>
  <PATH-ATTRIB Path ID=2 <BACKUP-TLV B=0, Backup_Paths=[3]>>
  <ERO B>
  <PATH-ATTRIB Path ID=3 <BACKUP-TLV B=1, Backup_Paths=[]>>
  <ERO C>

```

Note that the syntax

```

<PATH-ATTRIB Path ID=1 <BACKUP-TLV B=0, Backup_Paths=[3]>>

```

means that this is PATH-ATTRIB object with Path ID field set to 1 and with a MULTIPATH-BACKUP TLV that has B-flag cleared and contains a single backup path with Backup Path ID of 3.

6.3. Composite Candidate Path

Consider the following Composite Candidate Path, taken from [RFC9256].

```

SR policy POL100 <headend = H1, color = 100, endpoint = E1>
  Candidate Path CP1 <protocol-origin = 20, originator =
    100:1.1.1.1, discriminator = 1>
    Preference 200
    Weight W1, SR policy <color = 1>
    Weight W2, SR policy <color = 2>

```

This is signaled in PCEP as:

```

<LSP PLSP-ID=100>
  <ASSOCIATION>
  <END-POINT>
  <PATH-ATTRIB Path ID=1
    <WEIGHT-TLV Weight=W1>
    <COLOR-TLV Color=1>>
  <ERO (empty)>
  <PATH-ATTRIB Path ID=2
    <WEIGHT-TLV Weight=W2>
    <COLOR-TLV Color=2>>
  <ERO (empty)>

```

6.4. Opposite Direction Tunnels

Consider the two opposite-direction SR Policies between endpoints H1 and E1.

SR policy POL1 <headend = H1, color, endpoint = E1>

```

Candidate Path CP1
  Preference 200
  Bidirectional Association = A1
  SID-List = <H1,M1,M2,E1>
  SID-List = <H1,M3,M4,E1>
Candidate Path CP2
  Preference 100
  Bidirectional Association = A2
  SID-List = <H1,M5,M6,E1>
  SID-List = <H1,M7,M8,E1>

```

SR policy POL2 <headend = E1, color, endpoint = H1>

```

Candidate Path CP1
  Preference 200
  Bidirectional Association = A1
  SID-List = <E1,M2,M1,H1>
  SID-List = <E1,M4,M3,H1>
Candidate Path CP2
  Preference 100
  Bidirectional Association = A2
  SID-List = <E1,M6,M5,H1>

```

The state-report for POL1, CP1 can be encoded as:

```

<state-report> =
  <LSP PLSP-ID=100>
  <BIDIRECTIONAL ASSOCIATION = A1>
  <PATH-ATTRIB Path ID=1 R-flag=0
    <OPPDIR-PATH-TLV OppositePath ID=3>>
  <ERO <H1,M1,M2,E1>>
  <PATH-ATTRIB Path ID=2 R-flag=0
    <OPPDIR-PATH-TLV OppositePath ID=4>>
  <ERO <H1,M3,M4,E1>>
  <PATH-ATTRIB Path ID=3 R-flag=1
    <OPPDIR-PATH-TLV OppositePath ID=1>>
  <ERO <E1,M2,M1,H1>>
  <PATH-ATTRIB Path ID=4 R-flag=1
    <OPPDIR-PATH-TLV OppositePath ID=2>>
  <ERO <E1,M4,M3,H1>>

```

The state-report for POL1, CP2 can be encoded as:

```
<state-report> =  
  <LSP PLSP-ID=200>  
  <BIDIRECTIONAL ASSOCIATION = A2>  
  <PATH-ATTRIB Path ID=1 R-flag=0  
    <OPPDIR-PATH-TLV OppositePath ID=3>>  
  <ERO <H1,M5,M6,E1>>  
  <PATH-ATTRIB Path ID=2 R-flag=0  
    <OPPDIR-PATH-TLV OppositePath ID=0>>  
  <ERO <H1,M7,M8,E1>>  
  <PATH-ATTRIB Path ID=3 R-flag=1  
    <OPPDIR-PATH-TLV OppositePath ID=1>>  
  <ERO <E1,M6,M5,H1>>
```

The state-report for POL2, CP1 can be encoded as:

```
<state-report> =  
  <LSP PLSP-ID=100>  
  <BIDIRECTIONAL ASSOCIATION = A1>  
  <PATH-ATTRIB Path ID=1 R-flag=0  
    <OPPDIR-PATH-TLV OppositePath ID=3>>  
  <ERO <E1,M2,M1,H1>>  
  <PATH-ATTRIB Path ID=2 R-flag=0  
    <OPPDIR-PATH-TLV OppositePath ID=4>>  
  <ERO <E1,M4,M3,H1>>  
  <PATH-ATTRIB Path ID=3 R-flag=1  
    <OPPDIR-PATH-TLV OppositePath ID=1>>  
  <ERO <H1,M1,M2,E1>>  
  <PATH-ATTRIB Path ID=4 R-flag=1  
    <OPPDIR-PATH-TLV OppositePath ID=2>>  
  <ERO <H1,M3,M4,E1>>
```

The state-report for POL2, CP2 can be encoded as:

```
<state-report> =  
  <LSP PLSP-ID=200>  
  <BIDIRECTIONAL ASSOCIATION = A2>  
  <PATH-ATTRIB Path ID=1 R-flag=0  
    <OPPDIR-PATH-TLV OppositePath ID=3>>  
  <ERO <E1,M6,M5,H1>>  
  <PATH-ATTRIB Path ID=2 R-flag=1  
    <OPPDIR-PATH-TLV OppositePath ID=0>>  
  <ERO <H1,M7,M8,E1>>  
  <PATH-ATTRIB Path ID=3 R-flag=1  
    <OPPDIR-PATH-TLV OppositePath ID=1>>  
  <ERO <H1,M5,M6,E1>>
```

7. Implementation Status

Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC7942].

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".

7.1. Cisco Systems

Organization: Cisco Systems
Implementation: IOS-XR PCC and PCE
Description: Circuit-Style SR Policies
Maturity Level: Supported feature
Coverage: Multiple Segment Lists and reverse paths in SR Policy
Contact: mkoldych@cisco.com

7.2. Ciena Corp

Organization: Ciena Corp
Implementation: Head-end and controller
Maturity Level: Proof of concept
Coverage: Full
Contact: byadav@ciena.com

7.3. Huawei Technologies

Organization: Huawei Technologies Co.,Ltd.
 Implementation: Huawei's Router and Controller
 Maturity Level: Proof of concept
 Coverage: Partial
 Contact: tanren@huawei.com

8. IANA Considerations

8.1. PCEP Object

IANA is requested to confirm the following allocation in the "PCEP Objects" within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

Object-Class Value	Name	Object-Type Value	Reference
45	PATH-ATTRIB	1	This document

8.2. PCEP TLV

IANA is requested to confirm the following allocations within the "PCEP TLV Type Indicators" within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

TLV Type Value	TLV Name	Reference
60	MULTIPATH-CAP	This document
61	MULTIPATH-WEIGHT	This document
62	MULTIPATH-BACKUP	This document
63	MULTIPATH-OPPDIR-PATH	This document

IANA is requested to make new allocations within the "PCEP TLV Type Indicators" within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

TLV Type Value	TLV Name	Reference
TBD1	MULTIPATH-FORWARD-CLASS	This document

8.3. PCEP-Error Object

IANA is requested to confirm the following allocations within the "PCEP-ERROR Object Error Types and Values" within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

Error-Type	Error-Value	Reference
10	38 - Conflicting Path ID	This document
19	20 - Not supported path backup	This document
19	21 - Non-empty path	This document

IANA is requested to make new allocations within the "PCEP-ERROR Object Error Types and Values" within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

Error-Type	Error-Value	Reference
10	TBD2 - Unexpected PATH-ATTRIB Object	This document
19	TBD3 - Unsupported multipath capability	This document
19	TBD4 - Invalid opposite-direction path mapping	This document

8.4. Flags in the MULTIPATH-CAP TLV

IANA is requested to create a new sub-registry to manage the Flag field of the MULTIPATH-CAP TLV, called "Flags in MULTIPATH-CAP TLV" within the "Path Computation Element Protocol (PCEP) Numbers" registry group. New values are to be assigned by "IETF review" [RFC8126]

Bit	Description	Reference
0-10	Unassigned	This document
11	C-flag: support for Composite Candidate Path processing	This document
12	F-flag: support for processing MULTIPATH-FORWARD-CLASS TLV	This document
13	O-flag: support for processing MULTIPATH-OPPDIR-PATH TLV	This document
14	B-flag: support for processing MULTIPATH-BACKUP TLV	This document
15	W-flag: support for processing MULTIPATH-WEIGHT TLV	This document

8.5. Flags in the PATH-ATTRIB Object

IANA is requested to create a new sub-registry to manage the Flag field of the PATH-ATTRIB object, called "Flags in PATH-ATTRIB Object" within the "Path Computation Element Protocol (PCEP) Numbers" registry group. New values are to be assigned by "IETF review" [RFC8126]

Bit	Description	Reference
0-12	Unassigned	This document
13-15	O-flag: Operational state	This document

8.6. Flags in the MULTIPATH-BACKUP TLV

IANA is requested to create a new sub-registry to manage the Flag field of the MULTIPATH-BACKUP TLV, called "Flags in MULTIPATH-BACKUP TLV" within the "Path Computation Element Protocol (PCEP) Numbers" registry group. New values are to be assigned by "IETF review" [RFC8126]

Bit	Description	Reference
0-14	Unassigned	This document
15	B-flag: Pure backup	This document

8.7. Flags in the MULTIPATH-OPPDIR-PATH TLV

IANA is requested to create a new sub-registry to manage the flag fields of the MULTIPATH-OPPDIR-PATH TLV, called "Flags in the MULTIPATH-OPPDIR-PATH TLV" within the "Path Computation Element Protocol (PCEP) Numbers" registry group. New values are to be assigned by "IETF review" [RFC8126]

Bit	Description	Reference
0-12	Unassigned	This document
14	L-flag: Link co-routed	This document
15	N-flag: Node co-routed	This document

9. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8281], [RFC8664], [RFC9256], [RFC9862] and [RFC9863] are applicable to this specification.

As per [RFC8231], it is RECOMMENDED that these PCEP extensions can only be activated on authenticated and encrypted sessions across PCEs and PCCs belonging to the same administrative authority, using Transport Layer Security (TLS) [RFC8253][I-D.ietf-pce-pceps-tls13] as per the recommendations and best current practices in [RFC9325].

10. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231], [RFC8664], and [RFC9256] apply to the PCEP protocol extensions defined in this document. In addition, the requirements and considerations listed in this section apply.

10.1. Control of Function and Policy

A PCEP speaker (PCC or PCE) implementation SHOULD allow an operator to enable or disable the multipath capabilities advertised in the MULTIPATH-CAP TLV (see Section 4).

10.2. Information and Data Models

It is expected that a future version of the PCEP YANG module [I-D.ietf-pce-pcep-yang] will be extended to include the PCEP extensions defined in this document.

10.3. Liveness Detection and Monitoring

The mechanisms defined in this document do not introduce any new liveness detection or monitoring requirements in addition to those already defined in [RFC5440] and [RFC8231].

10.4. Verify Correct Operations

In addition to the verification requirements in [RFC5440] and [RFC8231], the following considerations apply:

- * An implementation SHOULD allow an operator to view the capabilities advertised in the MULTIPATH-CAP TLV by each PCEP peer for a session and for individual LSPs.
- * An implementation SHOULD allow an operator to view the PATH-ATTRIB object and all its associated TLVs for each path within an LSP. This includes the Path ID, weight, backup information, and opposite-direction path associations.
- * An implementation SHOULD provide a mechanism to log and display the new PCEP errors defined in this document

10.5. Requirements On Other Protocols

The PCEP extensions defined in this document do not impose any new requirements on other protocols.

10.6. Impact On Network Operations

The mechanisms in this document allow for more complex LSP structures with multiple paths. Network operators should be aware of the potential increase in PCEP message sizes and the additional state that must be maintained by PCEP speakers. The "Number of Multipaths" field in the MULTIPATH-CAP TLV can be used to control the scale of multipath computations and state.

11. Acknowledgement

Thanks to Dhruv Dhody for ideas and discussion. Thanks to Yuan Yaping for review comments.

12. Contributors

Bhupendra Yadav
Ciena
Email: byadav@ciena.com

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com

Zafar Ali
Cisco Systems
Email: zali@cisco.com

Andrew Stone
Nokia
Email: andrew.stone@nokia.com

Chen Ran
ZTE
Email: chen.ran@zte.com.cn

13. References

13.1. Normative References

- [I-D.ietf-pce-pceps-tls13]
Dhody, D., Turner, S., and R. Housley, "Updates for PCEPS: TLS Connection Establishment Restrictions", Work in Progress, Internet-Draft, draft-ietf-pce-pceps-tls13-04, 9 January 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pceps-tls13-04>>.
- [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/rfc/rfc2119>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/rfc/rfc5440>>.

- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, DOI 10.17487/RFC5511, April 2009, <<https://www.rfc-editor.org/rfc/rfc5511>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/rfc/rfc8231>>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <<https://www.rfc-editor.org/rfc/rfc8253>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/rfc/rfc8281>>.
- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/rfc/rfc8664>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/rfc/rfc9256>>.
- [RFC9325] Sheffer, Y., Saint-Andre, P., and T. Fossati, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 9325, DOI 10.17487/RFC9325, November 2022, <<https://www.rfc-editor.org/rfc/rfc9325>>.

- [RFC9862] Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", RFC 9862, DOI 10.17487/RFC9862, October 2025, <<https://www.rfc-editor.org/rfc/rfc9862>>.
- [RFC9863] Rajagopalan, B., Beeram, V., Peng, S., Koldychev, M., and G. Mishra, "Path Computation Element Protocol (PCEP) Extension for Color", RFC 9863, DOI 10.17487/RFC9863, October 2025, <<https://www.rfc-editor.org/rfc/rfc9863>>.

13.2. Informative References

- [I-D.draft-ietf-pce-sr-p2mp-policy]
Bidgoli, H., Voyer, D., Budhiraja, A., Parekh, R., and S. Sivabalan, "PCEP extensions for SR P2MP Policy", Work in Progress, Internet-Draft, draft-ietf-pce-sr-p2mp-policy-13, 19 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-sr-p2mp-policy-13>>.
- [I-D.ietf-pce-pcep-yang]
Dhody, D., Beeram, V. P., Hardwick, J., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-30, 26 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-yang-30>>.
- [I-D.ietf-spring-cs-sr-policy]
Schmutzer, C., Ali, Z., Maheshwari, P., Rokui, R., and A. Stone, "Circuit Style Segment Routing Policy", Work in Progress, Internet-Draft, draft-ietf-spring-cs-sr-policy-13, 5 December 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-cs-sr-policy-13>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/rfc/rfc7942>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/rfc/rfc8126>>.

[RFC9059] Gandhi, R., Ed., Barth, C., and B. Wen, "Path Computation Element Communication Protocol (PCEP) Extensions for Associated Bidirectional Label Switched Paths (LSPs)", RFC 9059, DOI 10.17487/RFC9059, June 2021, <<https://www.rfc-editor.org/rfc/rfc9059>>.

Authors' Addresses

Mike Koldychev
Ciena Corporation
Email: mkoldych@ciena.com

Siva Sivabalan
Ciena Corporation
Email: ssivabal@ciena.com

Tarek Saad
Cisco Systems
Email: tsaad@cisco.com

Vishnu Pavan Beeram
Juniper Networks, Inc.
Email: vbeeram@juniper.net

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

Shuping Peng
Huawei Technologies
Email: pengshuping@huawei.com

Samuel Sidor (editor)
Cisco Systems.
Email: ssidor@cisco.com