

PCE Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 27, 2023

M. Koldychev  
Cisco Systems, Inc.  
S. Sivabalan  
Ciena Corporation  
C. Barth  
Juniper Networks, Inc.  
S. Peng  
Huawei Technologies  
H. Bidgoli  
Nokia  
October 24, 2022

**PCEP extension to support Segment Routing Policy Candidate Paths  
draft-ietf-pce-segment-routing-policy-cp-08**

**Abstract**

This document introduces a mechanism to specify a Segment Routing (SR) policy, as a collection of SR candidate paths. An SR policy is identified by <headend, color, endpoint> tuple. An SR policy can contain one or more candidate paths where each candidate path is identified in PCEP by its uniquely assigned PLSP-ID. This document proposes extension to PCEP to support association among candidate paths of a given SR policy. The mechanism proposed in this document is applicable to both MPLS and IPv6 data planes of SR.

**Requirements Language**

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

**Status of This Memo**

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 April 27, 2023.

## Copyright Notice

Copyright (c) 2022 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Motivation</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">Group Candidate Paths belonging to the same SR policy</a>	<a href="#">5</a>
<a href="#">3.2.</a>	<a href="#">Instantiation of SR policy candidate paths</a>	<a href="#">5</a>
<a href="#">3.3.</a>	<a href="#">Avoid computing lower preference candidate paths</a>	<a href="#">5</a>
<a href="#">3.4.</a>	<a href="#">Minimal signaling overhead</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Procedure</a>	<a href="#">6</a>
<a href="#">4.1.</a>	<a href="#">Overview</a>	<a href="#">6</a>
<a href="#">4.1.1.</a>	<a href="#">SR Policy Identifiers</a>	<a href="#">7</a>
<a href="#">4.1.2.</a>	<a href="#">SR Policy Candidate Path Identifiers</a>	<a href="#">7</a>
<a href="#">4.1.3.</a>	<a href="#">SR Policy Candidate Path Attributes</a>	<a href="#">7</a>
<a href="#">4.2.</a>	<a href="#">Multiple Optimization Objectives and Constraints</a>	<a href="#">8</a>
<a href="#">5.</a>	<a href="#">SR Policy Association</a>	<a href="#">8</a>
<a href="#">5.1.</a>	<a href="#">Association Parameters</a>	<a href="#">8</a>
<a href="#">5.2.</a>	<a href="#">Association Information</a>	<a href="#">10</a>
<a href="#">5.2.1.</a>	<a href="#">SR Policy Name TLV</a>	<a href="#">10</a>
<a href="#">5.2.2.</a>	<a href="#">SR Policy Candidate Path Identifiers TLV</a>	<a href="#">11</a>
<a href="#">5.2.3.</a>	<a href="#">SR Policy Candidate Path Name TLV</a>	<a href="#">12</a>
<a href="#">5.2.4.</a>	<a href="#">SR Policy Candidate Path Preference TLV</a>	<a href="#">12</a>
<a href="#">6.</a>	<a href="#">Generic Mechanisms</a>	<a href="#">13</a>
<a href="#">6.1.</a>	<a href="#">Computation Priority TLV</a>	<a href="#">13</a>
<a href="#">6.2.</a>	<a href="#">Explicit Null Label Policy (ENLP) TLV</a>	<a href="#">13</a>
<a href="#">6.3.</a>	<a href="#">Invalidation TLV</a>	<a href="#">14</a>
<a href="#">6.4.</a>	<a href="#">Specified-BSID-only</a>	<a href="#">15</a>
<a href="#">7.</a>	<a href="#">Examples</a>	<a href="#">16</a>



<a href="#">7.1.</a>	<a href="#">PCC Initiated SR Policy with single candidate-path . . .</a>	<a href="#">16</a>
<a href="#">7.2.</a>	<a href="#">PCC Initiated SR Policy with multiple candidate-paths . .</a>	<a href="#">16</a>
<a href="#">7.3.</a>	<a href="#">PCE Initiated SR Policy with single candidate-path . . .</a>	<a href="#">17</a>
<a href="#">7.4.</a>	<a href="#">PCE Initiated SR Policy with multiple candidate-paths . .</a>	<a href="#">17</a>
<a href="#">8.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">18</a>
<a href="#">8.1.</a>	<a href="#">Association Type . . . . .</a>	<a href="#">18</a>
<a href="#">8.2.</a>	<a href="#">PCEP TLV Type Indicators . . . . .</a>	<a href="#">18</a>
<a href="#">8.3.</a>	<a href="#">PCEP Errors . . . . .</a>	<a href="#">19</a>
<a href="#">8.4.</a>	<a href="#">TE-PATH-BINDING TLV Flag field . . . . .</a>	<a href="#">19</a>
<a href="#">9.</a>	<a href="#">Implementation Status . . . . .</a>	<a href="#">19</a>
<a href="#">9.1.</a>	<a href="#">Cisco . . . . .</a>	<a href="#">20</a>
<a href="#">9.2.</a>	<a href="#">Juniper . . . . .</a>	<a href="#">20</a>
<a href="#">10.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">21</a>
<a href="#">11.</a>	<a href="#">Acknowledgement . . . . .</a>	<a href="#">21</a>
<a href="#">12.</a>	<a href="#">References . . . . .</a>	<a href="#">21</a>
<a href="#">12.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">21</a>
<a href="#">12.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">23</a>
<a href="#">Appendix A.</a>	<a href="#">Contributors . . . . .</a>	<a href="#">23</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">24</a>

## **[1.](#) Introduction**

Path Computation Element (PCE) Communication Protocol (PCEP) [[RFC5440](#)] enables the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between two PCEs based on the PCE architecture [[RFC4655](#)].

PCEP Extensions for the Stateful PCE Model [[RFC8231](#)] describes a set of extensions to PCEP to enable active control of Multiprotocol Label Switching Traffic Engineering (MPLS-TE) and Generalized MPLS (GMPLS) tunnels. [[RFC8281](#)] describes the setup and teardown of PCE-initiated LSPs under the active stateful PCE model, without the need for local configuration on the PCC, thus allowing for dynamic centralized control of a network.

PCEP Extensions for Segment Routing [[RFC8664](#)] specifies extensions to the Path Computation Element Protocol (PCEP) that allow a stateful PCE to compute and initiate Traffic Engineering (TE) paths, as well as a PCC to request a path subject to certain constraint(s) and optimization criteria in SR networks.

PCEP Extensions for Establishing Relationships Between Sets of LSPs [[RFC8697](#)] introduces a generic mechanism to create a grouping of LSPs which can then be used to define associations between a set of LSPs and a set of attributes (such as configuration parameters or behaviors) and is equally applicable to stateful PCE (active and passive modes) and stateless PCE.



Segment Routing Policy for Traffic Engineering

[[I-D.ietf-spring-segment-routing-policy](#)] details the concepts of SR Policy and approaches to steering traffic into an SR Policy.

An SR Policy contains one or more SR Policy Candidate Paths where one or more such paths can be computed via PCE. This document specifies PCEP extensions to signal additional information to map candidate paths to their SR policies. Each candidate path maps to a unique PLSP-ID in PCEP. By associating multiple candidate paths together, a PCE becomes aware of the hierarchical structure of an SR policy. Thus the PCE can take computation and control decisions about the candidate paths, with the additional knowledge that these candidate paths belong to the same SR policy. This is accomplished via the use of the existing PCEP Association object, by defining a new association type specifically for associating SR candidate paths into a single SR policy.

## 2. Terminology

The following terminologies are used in this document:

Endpoint: The IPv4 or IPv6 endpoint address of the SR policy in question, as described in [[I-D.ietf-spring-segment-routing-policy](#)].

Association Parameters: As described in [[RFC8697](#)], the combination of the mandatory fields Association Type, Association ID and Association Source in the ASSOCIATION object uniquely identify the association group. If the optional TLVs - Global Association Source or Extended Association ID are included, then they MUST be included in combination with mandatory fields to uniquely identify the association group.

Association Information: As described in [[RFC8697](#)], the ASSOCIATION object could also include other TLVs based on the association types, that provides non-key information.

SRPAG: SR Policy Association Group.

SRPAT: SR Policy Association Type.

SRPAT ASSOCIATION: ASSOCIATION object of type SR Policy Association.

PCC: Path Computation Client. Any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or



route based on a network graph and applying computational constraints.

PCEP: Path Computation Element Protocol.

PCEP Tunnel: The entity identified by the PLSP-ID, as per [\[I-D.koldychev-pce-operational\]](#).

### **3. Motivation**

The SR Policy Association and its TLVs, defined in this document, allow PCEP speakers to exchange additional information about SR Policy Candidate Paths and their container SR Policy.

#### **[3.1.](#) Group Candidate Paths belonging to the same SR policy**

Since each SR Policy Candidate Path appears as a different Tunnel (identified via a PLSP-ID) in PCEP, it is useful to group together all the SR Policy Candidate Paths that belong to the same SR Policy. Furthermore, it is useful for the PCE to have knowledge of the SR Policy related information such as color, endpoint, protocol origin, discriminator, and preference.

#### **[3.2.](#) Instantiation of SR policy candidate paths**

A PCE needs to instantiate one or more SR Policy Candidate Paths on the PCC, as specified in [\[RFC8281\]](#). Each SR Policy Candidate Path is identified by the tuple <headend, color, endpoint, originator, discriminator, preference>. This draft provides a mechanism to signal this information in PCEP.

#### **[3.3.](#) Avoid computing lower preference candidate paths**

When a PCE knows that a given set of SR Policy Candidate Paths all belong to the same SR Policy, then path computation MAY be done on only the highest preference candidate-path(s). Path computation for lower preference paths is not necessary if one or two higher preference paths are already computed. Since computing their paths will not affect traffic steering, it MAY be postponed until the higher preference paths become invalid.

#### **[3.4.](#) Minimal signaling overhead**

When an SR Policy contains multiple SR Policy Candidate Paths computed by a PCE, such candidate paths can be created, updated and deleted independently of each other. This is achieved by making each SR Policy Candidate Path correspond to a unique Tunnel (identified via PLSP-ID). For example, if an SR Policy has 4 SR Policy Candidate





Paths, then if the PCE wants to update one of those, only one set of PCUpd and PCRpt messages needs to be exchanged.

## **4. Procedure**

### **4.1. Overview**

As per [[RFC8697](#)], LSPs are placed into an association group. As per [[I-D.koldychev-pce-operational](#)], LSPs are contained in PCEP Tunnels and a PCEP Tunnel is contained in an Association if all of its LSPs are in that Association. PCEP Tunnels naturally map to SR Policy Candidate Paths and PCEP Associations naturally map to SR Policies.

The mapping between PCEP Associations and SR Policies is always one-to-one. However, the mapping between PCEP Tunnels and SR Policy Candidate Paths may be either one-to-one, or many-to-one, see [Section 4.2](#).

Each SR Policy Candidate Path contains one or more Segment Lists. The subject of encoding multiple Segment Lists within an SR Policy Candidate Path is described in [[I-D.koldychev-pce-multipath](#)].

This document defines a new Association Type called "SR Policy Association", of value 6 based on the generic ASSOCIATION object. The new Association Type is also called "SRPAT", for "SR Policy Association Type". We say "SRPAT ASSOCIATION" to mean "ASSOCIATION object of type SR Policy Association". The group of LSPs that are part of the SR Policy Association is called "SRPAG", for "SR Policy Association Group".

As per the processing rules specified in [section 6.4 of \[RFC8697\]](#), if a PCEP speaker does not support the SRPAT, it MUST return a PCErr message with Error-Type = 26 "Association Error", Error-Value = 1 "Association-type is not supported".

A given LSP MUST belong to at most one SRPAG, since an SR Policy Candidate Path cannot belong to multiple SR Policies. If a PCEP speaker receives a PCEP message with more than one SRPAT ASSOCIATION for the same LSP, then the PCEP speaker MUST send a PCErr message with Error-Type = 26 "Association Error", Error-Value = 7 "Cannot join the association group".

An SRPAT ASSOCIATION carries three pieces of information: SR Policy Identifiers, SR Policy Candidate Path Identifiers, and SR Policy Candidate Path Attributes.



#### **4.1.1. SR Policy Identifiers**

SR Policy Identifiers uniquely identify the SR policy within the context of the headend. SR Policy Identifiers MUST be the same for all SR Policy Candidate Paths in the same SRPAG. SR Policy Identifiers MUST NOT change for a given SR Policy Candidate Path during its lifetime. SR Policy Identifiers MUST be different for different SRPAGs. SR Policy Identifiers consist of:

- o Headend router where the SR Policy originates.
- o Color of SR Policy.
- o Endpoint of SR Policy.

#### **4.1.2. SR Policy Candidate Path Identifiers**

SR Policy Candidate Path Identifiers uniquely identify the SR Policy Candidate Path within the context of an SR Policy. SR Policy Candidate Path Identifiers MUST NOT change for a given LSP during its lifetime. SR Policy Candidate Path Identifiers MUST be different for different LSPs within the same SRPAG. When these rules are not satisfied, the PCE MUST send a PCErr message with Error-Type = 26 "Association Error", Error Value = TBD8 "SR Policy Candidate Path Identifiers Mismatch". SR Policy Candidate Path Identifiers consist of:

- o Protocol Origin.
- o Originator.
- o Discriminator.

#### **4.1.3. SR Policy Candidate Path Attributes**

SR Policy Candidate Path Attributes carry non-key information about the candidate path and MAY change during the lifetime of the LSP. SR Policy Candidate Path Attributes consist of:

- o Preference.
- o Optionally, the SR Policy Candidate Path name.
- o Optionally, the SR Policy name.



#### **4.2. Multiple Optimization Objectives and Constraints**

In certain scenarios, it is desired for each SR Policy Candidate Path to contain multiple sub-candidate paths, each of which has a different optimization objective and constraints. Traffic is then sent ECMP or UCMP among these sub-candidate paths.

This is represented in PCEP by a many-to-one mapping between PCEP Tunnels and SR Policy Candidate Paths. This means that multiple PCEP Tunnels are allocated for each SR Policy Candidate Path. Each PCEP Tunnel has its own optimization objective and constraints. When a single SR Policy Candidate Path contains multiple PCEP Tunnels, each of these PCEP Tunnels MUST have identical values of Candidate Path Identifiers, as encoded in SRPOLICY-CPATH-ID TLV, see [Section 5.2.2](#).

### **5. SR Policy Association**

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 (6) "SR Policy Association". This Association Type is dynamic in nature, thus operator-configured Association Range MUST NOT be set for this Association type and MUST be ignored.

#### **5.1. Association Parameters**

As per [\[I-D.ietf-spring-segment-routing-policy\]](#), an SR Policy is identified through the tuple <headend, color, endpoint>. the headend is encoded as the Association Source in the ASSOCIATION object and the color and endpoint are encoded as part of Extended Association ID TLV.

The Association Parameters (see [Section 2](#)) consist of:

- o Association Type: set to 6 "SR Policy Association".
- o Association Source (IPv4/IPv6): set to the headend IP address.
- o Association ID (16-bit): set to "1".
- o Extended Association ID TLV: encodes the Color and Endpoint of the SR Policy.

The Association Source MUST be set to the headend value of the SR Policy, as defined in [\[I-D.ietf-spring-segment-routing-policy\]](#) [Section 2.1](#). If the PCC receives a PCInit message for a non-existent SR Policy, where the Association Source is set not to the headend value but to some globally unique IP address that the PCC owns, then



the PCC SHOULD accept the PCInit message and create the SR Policy Association with the Association Source that was sent in the PCInit message.

The 16-bit Association ID field in the ASSOCIATION object MUST be set to the value of "1".

The Extended Association ID TLV MUST be included and it MUST be in the following format:

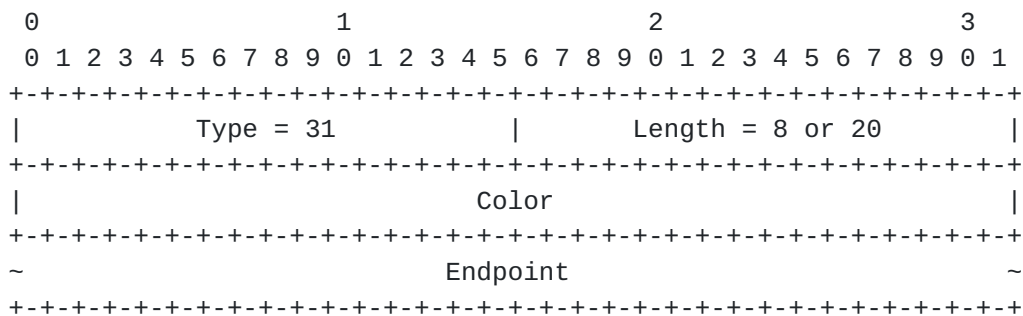


Figure 1: Extended Association ID TLV format

Type: Extended Association ID TLV, type = 31.

Length: Either 8 or 20, depending on whether IPv4 or IPv6 address is encoded in the Endpoint.

Color: SR Policy color value.

Endpoint: can be either IPv4 or IPv6, depending on whether the policy endpoint is IPv4 or IPv6. This value MAY be different from the one contained in the END-POINTS object, or in the LSP IDENTIFIERS TLV of the LSP object. This value is part of the tuple <color, endpoint> that identifies the SR Policy on a given headend.

If the PCEP speaker receives an SRPAT ASSOCIATION whose Association Parameters do not follow the above specification, then the PCEP speaker MUST send PCErr message with Error-Type = 26 "Association Error", Error-Value = TBD7 "SR Policy Identifiers Mismatch".

The purpose of choosing the Association Parameters in this way is to guarantee that there is no possibility of a race condition when multiple PCEP speakers want to create the same SR Policy at the same time. By adhering to this format, all PCEP speakers come up with the same Association Parameters independently of each other. Thus, there is no chance that different PCEP speakers will come up with different Association Parameters for the same SR Policy.





## 5.2. Association Information

The SRPAT ASSOCIATION contains the following TLVs:

- o SRPOLICY-POL-NAME TLV: (optional) encodes SR Policy Name string.
- o SRPOLICY-CPATH-ID TLV: (mandatory) encodes SR Policy Candidate Path Identifiers.
- o SRPOLICY-CPATH-NAME TLV: (optional) encodes SR Policy Candidate Path string name.
- o SRPOLICY-CPATH-PREFERENCE TLV: (optional) encodes SR Policy Candidate Path preference value.

Of these new TLVs, SRPOLICY-CPATH-ID TLV is mandatory. When a mandatory TLV is missing from the SRPAT ASSOCIATION object, the PCE MUST send a PCErr message with Error-Type = 6 "Mandatory Object Missing", Error-Value = TBD6 "Missing Mandatory TLV".

### 5.2.1. SR Policy Name TLV

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




Figure 2: The SRPOLICY-POL-NAME TLV format

Type: 56 for "SRPOLICY-POL-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.

SR Policy Name: SR Policy name, as defined in [\[I-D.ietf-spring-segment-routing-policy\]](#). It SHOULD be a string of printable ASCII characters, without a NULL terminator.



### 5.2.2. SR Policy Candidate Path Identifiers TLV

The SRPOLICY-CPATH-ID TLV is a mandatory TLV for the SRPAT ASSOCIATION. Only one SRPOLICY-CPATH-ID TLV SHOULD be encoded by the sender and only the first occurrence is processed and any others MUST be ignored.

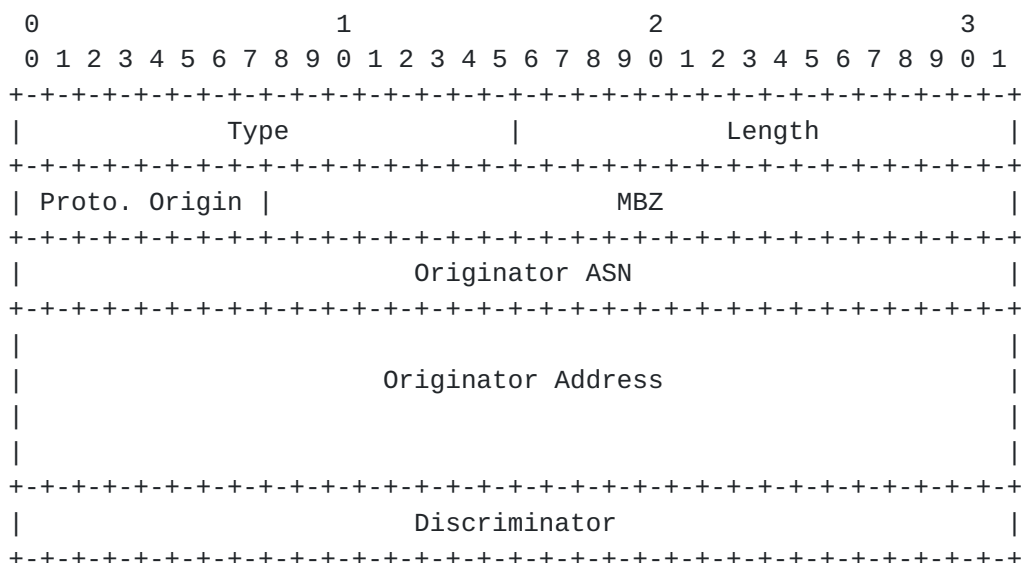


Figure 3: The SRPOLICY-CPATH-ID TLV format

Type: 57 for "SRPOLICY-CPATH-ID" TLV.

Length: 28.

Protocol Origin: 8-bit value that encodes the protocol origin, as specified in [[I-D.ietf-spring-segment-routing-policy](#)] [Section 2.3](#). Note that in PCInit messages, the Protocol Origin is always set to "PCEP".

Originator ASN: Represented as 4 byte number, part of the originator identifier, as specified in [[I-D.ietf-spring-segment-routing-policy](#)] [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 [[I-D.ietf-spring-segment-routing-policy](#)] [Section 2.4](#).

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



### 5.2.3. SR Policy Candidate Path Name TLV

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

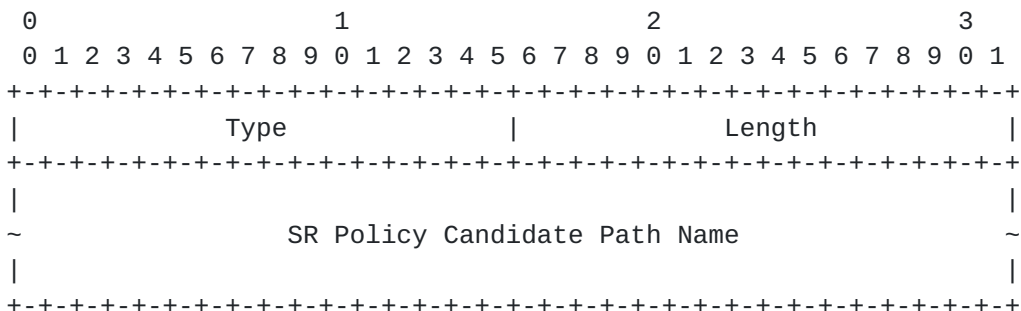


Figure 4: The SRPOLICY-CPATH-NAME TLV format

Type: 58 for "SRPOLICY-CPATH-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.

SR Policy Candidate Path Name: SR Policy Candidate Path Name, as defined in [[I-D.ietf-spring-segment-routing-policy](#)]. It SHOULD be a string of printable ASCII characters, without a NULL terminator.

### 5.2.4. SR Policy Candidate Path Preference TLV

The SRPOLICY-CPATH-PREFERENCE TLV is an optional TLV for the SRPAT ASSOCIATION. Only one SRPOLICY-CPATH-PREFERENCE TLV SHOULD be encoded by the sender and only the first occurrence is processed and any others MUST be ignored.




Figure 5: The SRPOLICY-CPATH-PREFERENCE TLV format

Type: 59 for "SRPOLICY-CPATH-PREFERENCE" TLV.



Length: 4.

Preference: Numerical preference of the candidate path, as specified in Section 2.7 of [[I-D.ietf-spring-segment-routing-policy](#)].

If the TLV is missing, a default Preference value of 100 is used, as specified in Section 2.7 of [[I-D.ietf-spring-segment-routing-policy](#)].

## 6. Generic Mechanisms

This section describes various mechanisms that are standardized for SR Policies in [[I-D.ietf-spring-segment-routing-policy](#)], but are equally applicable to other tunnel types, such as RSVP-TE tunnels. Hence this section does not make use of the SRPAT ASSOCIATION.

### 6.1. Computation Priority TLV

The COMPUTATION-PRIORITY TLV is an optional TLV for the LSP object. It is used to signal the numerical computation priority, as specified in Section 2.12 of [[I-D.ietf-spring-segment-routing-policy](#)]. If the TLV is absent from the LSP object, a default Priority value of 128 is used.




Figure 6: The COMPUTATION-PRIORITY TLV format

Type: TBD1 for "COMPUTATION-PRIORITY" TLV.

Length: 4.

Priority: Numerical priority with which this LSP is to be recomputed by the PCE upon topology change.

### 6.2. Explicit Null Label Policy (ENLP) TLV

The ENLP TLV is an optional TLV for the LSP object. It is used to implement the "Explicit Null Label Policy", as specified in Section 2.4.5 of [[I-D.ietf-idr-segment-routing-te-policy](#)].








Figure 7: The Explicit Null Label Policy (ENLP) TLV format

Type: TBD2 for "ENLP" TLV.

Length: 4.

ENLP (Explicit NULL Label Policy): same values as in Section 2.4.5 of [\[I-D.ietf-idr-segment-routing-te-policy\]](#).

### 6.3. Invalidation TLV

The INVALIDATION TLV is an optional TLV for the LSP object. It is used to control traffic steering into the LSP during the time when the LSP is operationally down/invalid. In the context of SR Policy, this TLV facilitate the "Drop upon invalid" behavior, specified in Section 8.2 of [\[I-D.ietf-spring-segment-routing-policy\]](#). Normally, if the LSP is down/invalid then traffic that is originally destined for that LSP is steered somewhere else, such as via IGP or via another LSP. The "Drop upon invalid" behavior specifies that such traffic MUST NOT be re-routed and has to be dropped at the head-end. While in the "Drop upon invalid" state, the LSP operational state is "UP", as indicated by the O-flag in the LSP object. However the ERO object is empty, indicating that traffic is being dropped.

In addition to the above, this TLV can also be used by the PCC to report to the PCE various reasons for LSP being invalidated. Invalidation reasons are represented by a set of flags.




Figure 8: Invalidation Reasons Flags

- o U: Unknown - does not fit into any other categories below.
- o P: Path computation failure - no path was computed for the LSP.



- o F: First-hop resolution failure - head-end first hop resolution has failed.
- o V: Verification failure - OAM/PM/BFD path verification has indicated a breakage.




Figure 9: The INVALIDATION TLV format

Type: TBD3 for "INVALIDATION" TLV.

Length: 4.

Inval Reason: contains "Invalidation Reasons Flags" which encode the reason(s) why the LSP is currently invalidated. This field can be set to non-zero values only by the PCC, it MUST be set to 0 by the PCE and ignored by the PCC.

Drop Upon: contains "Invalidation Reasons Flags" for conditions that SHOULD cause the LSP to drop traffic. This field can be set to non-zero values by both PCC and PCE. This field MAY be set to all 1's (0xFF) to indicate that the LSP is to go into Drop upon invalid state for any reason. I.e., when the PCE does not wish to distinguish any reason for LSP invalidation and just simply wants it to always "Drop upon invalid" for any reason.

#### 6.4. Specified-BSID-only

Specified-BSID-only functionality is defined in Section 6.2.3 of [\[I-D.ietf-spring-segment-routing-policy\]](#). When specified-BSID-only is enabled for a particular binding SID, it means that the given binding SID is required to be allocated and programmed for the LSP to be operationally up. If the binding SID cannot be allocated or programmed for some reason, then the LSP must stay down.

To signal specified-BSID-only, a new bit: S (Specified-BSID-only) is allocated in the "TE-PATH-BINDING TLV Flag field" of the TE-PATH-BINDING TLV. When this bit is set for a particular BSID, it means that the BSID follows the Specified-BSID-only behavior. It is possible to have a mix of BSIDs for the same LSP: some with S=1 and some with S=0.



## **7. Examples**

### **7.1. PCC Initiated SR Policy with single candidate-path**

PCReq and PCRep messages are exchanged in the following sequence:

1. PCC sends PCReq message to the PCE, encoding the SRPAT ASSOCIATION and TLVs in the PCReq message.
2. PCE returns the path in PCRep message, and echoes back the SRPAT ASSOCIATION.

PCRpt and PCUpd messages are exchanged in the following sequence:

1. PCC sends PCRpt message to the PCE, including the LSP object and the SRPAT ASSOCIATION.
2. PCE computes path, possibly making use of the Association Information from the SRPAT ASSOCIATION.
3. PCE updates the SR policy candidate path's ERO using PCUpd message.

### **7.2. PCC Initiated SR Policy with multiple candidate-paths**

PCRpt and PCUpd messages are exchanged in the following sequence:

1. For each candidate path of the SR Policy, the PCC generates a different PLSP-ID and symbolic-name and sends multiple PCRpt messages (or one message with multiple LSP objects) to the PCE. Each LSP object is followed by SRPAT ASSOCIATION with identical Color and Endpoint values. The Association Source is set to the IP address of the PCC and the Association ID is set to a number that PCC locally chose to represent the SR Policy.
2. PCE takes into account that all the LSPs belong to the same SR policy. PCE prioritizes computation for the highest preference LSP and sends PCUpd message(s) back to the PCC.
3. If a new candidate path is added on the PCC by the operator, then a new PLSP-ID and symbolic name is generated for that candidate path and a new PCRpt is sent to the PCE.
4. If an existing candidate path is removed from the PCC by the operator, then that PLSP-ID is deleted from the PCE by sending PCRpt with the R-flag in the LSP object set.



### **7.3. PCE Initiated SR Policy with single candidate-path**

A candidate-path is created using the following steps:

1. PCE sends PCInitiate message, containing the SRPAT ASSOCIATION. The Association Source and the Association ID are set as described in [Section 5.1](#).
2. PCC uses the color, endpoint and preference from the SRPAT ASSOCIATION to create a new candidate path. If no SR policy exists to hold the candidate path, then a new SR policy is created to hold the new candidate-path. The Originator of the candidate path is set to be the address of the PCE that is sending the PCInitiate message.
3. PCC sends a PCRpt message back to the PCE to report the newly created Candidate Path. The PCRpt message contains the SRPAT ASSOCIATION.

A candidate-path is deleted using the following steps:

1. PCE sends PCInitiate message, setting the R-flag in the LSP object.
2. PCC uses the PLSP-ID from the LSP object to find the candidate path and delete it. If this is the last candidate path under the SR policy, then the containing SR policy is deleted as well.

### **7.4. PCE Initiated SR Policy with multiple candidate-paths**

A candidate-path is created using the following steps:

1. PCE SHOULD send a separate PCInitiate message for every candidate path that it wants to create, or it MAY send multiple LSP objects within a single PCInitiate message. The SRPAT ASSOCIATION is sent for every LSP in the PCInitiate message. The Association Source and the Association ID are set as described in [Section 5.1](#).
2. PCC creates multiple candidate paths under the same SR policy, identified by Color and Endpoint.
3. PCC sends a PCRpt message back to the PCE to report the newly created Candidate Path. The PCRpt message contains the SRPAT ASSOCIATION. The Association Source and the Association ID are set as described in [Section 5.1](#).

A candidate path is deleted using the following steps:





1. PCE sends PCInitiate message, setting the R-flag in the LSP object.
2. PCC uses the PLSP-ID from the LSP object to find the candidate path and delete it.

## 8. IANA Considerations

### 8.1. Association Type

This document defines a new association type: SR Policy Association. IANA is requested to make the following codepoint assignment in the "ASSOCIATION Type Field" subregistry [[RFC8697](#)] within the "Path Computation Element Protocol (PCEP) Numbers" registry:

Type	Name	Reference
6	SR Policy Association	This.I-D

### 8.2. PCEP TLV Type Indicators

This document defines four new TLVs for carrying additional information about SR policy and SR candidate paths. IANA is requested to make the assignment of a new value for the existing "PCEP TLV Type Indicators" registry as follows:

Value	Description	Reference
56	SRPOLICY-POL-NAME	This.I-D
57	SRPOLICY-CPATH-ID	This.I-D
58	SRPOLICY-CPATH-NAME	This.I-D
59	SRPOLICY-CPATH-PREFERENCE	This.I-D
TBD1	COMPUTATION-PRIORITY	This.I-D
TBD2	EXPLICIT-NULL-LABEL-POLICY	This.I-D
TBD3	INVALIDATION	This.I-D



### 8.3. PCEP Errors

This document defines one new Error-Value within the "Mandatory Object Missing" Error-Type and two new Error-Values within the "Association Error" Error-Type. IANA is requested to allocate new error values within the "PCEP-ERROR Object Error Types and Values" sub-registry of the PCEP Numbers registry, as follows:

Error-Type	Meaning	Error-value	Reference
6	Mandatory Object Missing		[RFC5440]
		TBD6: SR Policy Missing Mandatory TLV	This.I-D
26	Association Error		[RFC8697]
		TBD7: SR Policy Identifiers Mismatch	This.I-D
		TBD8: SR Policy Candidate Path Identifiers Mismatch	This.I-D

### 8.4. TE-PATH-BINDING TLV Flag field

IANA is requested to allocate new bit within the "TE-PATH-BINDING TLV Flag field" sub-registry of the PCEP Numbers registry, as follows:

Bit position	Description	Reference
1	Specified-BSID-only	This.I-D

## 9. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC 7942](#).]

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".

#### **[9.1.](#) Cisco**

- o Organization: Cisco Systems
- o Implementation: IOS-XR PCC and PCE.
- o Description: An experimental code-point is currently used.
- o Maturity Level: Proof of concept.
- o Coverage: Full.
- o Contact: mkoldych@cisco.com

#### **[9.2.](#) Juniper**

- o Organization: Juniper Networks
- o Implementation: Head-end and controller.
- o Description: An experimental code-point is currently used.
- o Maturity Level: Proof of concept.
- o Coverage: Partial.
- o Contact: cbarth@juniper.net



## **10. Security Considerations**

This document defines one new type for association, which do not add any new security concerns beyond those discussed in [[RFC5440](#)], [[RFC8231](#)], [[RFC8664](#)], [[I-D.ietf-pce-segment-routing-ipv6](#)] and [[RFC8697](#)] in itself.

The information carried in the SRPAT ASSOCIATION, as per this document is related to SR Policy. It often reflects information that can also be derived from the SR Database, but association provides a much easier grouping of related LSPs and messages. The SRPAT ASSOCIATION could provide an adversary with the opportunity to eavesdrop on the relationship between the LSPs. Thus securing the PCEP session using Transport Layer Security (TLS) [[RFC8253](#)], as per the recommendations and best current practices in [[RFC7525](#)], is RECOMMENDED.

## **11. Acknowledgement**

Would like to thank Stephane Litkowski, Boris Khasanov, Praveen Kumar and Tom Petch for review and suggestions.

## **12. References**

### **12.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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/info/rfc5440>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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/info/rfc8231>>.





- [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/info/rfc8281>>.
- [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/info/rfc7942>>.
- [I-D.ietf-spring-segment-routing-policy] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-22](#) (work in progress), March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-22.txt>>.
- [I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-20](#) (work in progress), July 2022, <<https://www.ietf.org/archive/id/draft-ietf-idr-segment-routing-te-policy-20.txt>>.
- [RFC8697] Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)", [RFC 8697](#), DOI 10.17487/RFC8697, January 2020, <<https://www.rfc-editor.org/info/rfc8697>>.
- [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/info/rfc8664>>.
- [I-D.koldychev-pce-operational] Koldychev, M., Sivabalan, S., Peng, S., Achaval, D., and H. Kotni, "PCEP Operational Clarification", [draft-koldychev-pce-operational-06](#) (work in progress), July 2022, <<https://www.ietf.org/archive/id/draft-koldychev-pce-operational-06.txt>>.



[I-D.koldychev-pce-multipath]

Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., and S. Peng, "PCEP Extensions for Signaling Multipath Information", [draft-koldychev-pce-multipath-05](#) (work in progress), February 2021, <<https://www.ietf.org/archive/id/draft-koldychev-pce-multipath-05.txt>>.

## **12.2. Informative References**

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", [BCP 195](#), [RFC 7525](#), DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [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/info/rfc8253>>.
- [I-D.ietf-pce-segment-routing-ipv6] Li(Editor), C., Negi, M. S., Sivabalan, S., Koldychev, M., Kaladharan, P., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing leveraging the IPv6 dataplane", [draft-ietf-pce-segment-routing-ipv6-15](#) (work in progress), October 2022, <<https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-pce-segment-routing-ipv6/>>.

## **Appendix A. Contributors**



Dhruv Dhody  
Huawei Technologies  
Divyashree Techno Park, Whitefield  
Bangalore, Karnataka 560066  
India

Email: dhruv.ietf@gmail.com

Cheng Li  
Huawei Technologies  
Huawei Campus, No. 156 Beiqing Rd.  
Beijing, 10095  
China

Email: chengli13@huawei.com

Samuel Sidor  
Cisco Systems, Inc.  
Eurovea Central 3.  
Pribinova 10  
811 09 Bratislava  
Slovakia

Email: ssidor@cisco.com

#### Authors' Addresses

Mike Koldychev  
Cisco Systems, Inc.  
2000 Innovation Drive  
Kanata, Ontario K2K 3E8  
Canada

Email: mkoldych@cisco.com

Siva Sivabalan  
Ciena Corporation  
385 Terry Fox Dr.  
Kanata, Ontario K2K 0L1  
Canada

Email: ssivabal@ciena.com



Colby Barth  
Juniper Networks, Inc.

Email: cbarth@juniper.net

Shuping Peng  
Huawei Technologies  
Huawei Campus, No. 156 Beiqing Rd.  
Beijing 100095  
China

Email: pengshuping@huawei.com

Hooman Bidgoli  
Nokia

Email: hooman.bidgoli@nokia.com



