Workgroup: PCE Working Group
Internet-Draft:
draft-ietf-pce-segment-routing-policy-cp-15
Updates: 8231 (if approved)
Published: 17 March 2024
Intended Status: Standards Track
Expires: 18 September 2024
Authors: M. Koldychev
Ciena Corporation
S. Sivabalan
Ciena Corporation
C. Barth
Juniper Networks, Inc.
S. Peng
Huawei Technologies
H. Bidgoli
Nokia

Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths

Abstract

Segment Routing (SR) allows a node to steer a packet flow along any path. SR Policy is an ordered list of segments (i.e., instructions) that represent a source-routed policy. Packet flows are steered into an SR Policy on a node where it is instantiated called a headend node. An SR Policy is made of one or more candidate paths.

This document specifies Path Computation Element Communication Protocol (PCEP) extension to associate candidate paths of the SR Policy. Additionally, this document updates [RFC8231] to allow stateful bringup of an SR LSP, without using PCReq/PCRep messages. This document is applicable to both Segment Routing over MPLS and to Segment Routing over IPv6 (SRv6).

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 18 September 2024.

Copyright Notice

Table of Contents

1. Introduction

[RFC8664] specifies extensions that allow PCEP to work with basic SR-TE paths. [RFC8697] introduces a generic mechanism to create a grouping of LSPs, called an Association. [RFC9256] introduces the SR Policy construct as a grouping of SR Candidate Paths.

This document extends [RFC8664] to support signaling SR Policy Candidate Paths and their attributes. SR Policy is modeled in PCEP as an Association, where the SR Candidate Paths are the members of that Association. 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.

Segment Routing Policy for Traffic Engineering [RFC9256] details the concepts of SR Policy and approaches to steering traffic into an SR Policy.

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.

This document extends [RFC8664] to support signaling SR Policy Candidate Paths and their attributes. SR Policy is modeled in PCEP as an Association, where the SR Candidate Paths are the members of that Association. 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 PCEP Association object with a new association type and several new TLVs.

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 [RFC9256].

SRPA: SR Policy Association. A new association type 'SR Policy Association' is used to group candidate paths belonging to the SR Policy. Depending on discussion context, it can refer to the PCEP ASSOCIATION object of SR Policy type or to a group of LSPs that belong to the association.

Association Parameters: As described in [RFC8697], refers to the key data, that uniquely identifies the Association.

Association Information: As described in [RFC8697], refers to the non-key information about the Association.

3. Overview

The SR Policy is represented by a new type of PCEP Association, called the SR Policy Association (SRPA). The SR Candidate Paths of an SR Policy are the PCEP LSPs within the same SRPA. The subject of encoding multiple Segment Lists within an SR Policy Candidate Path is described in [I-D.ietf-pce-multipath].

The SRPA carries three pieces of information: SR Policy Identifier, SR Policy Candidate Path Identifier, and SR Policy Candidate Path Attribute(s).

This document also specifies some additional information that is not encoded as part of SRPA: Computation Priority, Explicit Null Label Policy, Drop-upon-invalid behavior, and Specified-BSID-only.

3.1. SR Policy Identifier

SR Policy Identifier uniquely identifies the SR Policy [RFC9256] within the network. SR Policy Identifier MUST be the same for all SR Policy Candidate Paths in the same SRPA. SR Policy Identifier MUST NOT change for a given SR Policy Candidate Path during its lifetime. SR Policy Identifier MUST be different for different SRPAs. When these rules are not satisfied, the PCEP speaker MUST send a PCErr

message with Error-Type = 26 "Association Error", Error Value = 20 "SR Policy Identifier Mismatch". SR Policy Identifier consist of:

*Headend router where the SR Policy originates.

*Color of SR Policy.

*Endpoint of SR Policy.

3.2. SR Policy Candidate Path Identifier

SR Policy Candidate Path Identifier uniquely identifies the SR Policy Candidate Path within the context of an SR Policy. SR Policy Candidate Path Identifier MUST NOT change for a given LSP during its lifetime. SR Policy Candidate Path Identifier MUST be different for distinct Candidate Paths within the same SRPA. When these rules are not satisfied, the PCEP speaker MUST send a PCErr message with Error-Type = 26 "Association Error", Error Value = 21 "SR Policy Candidate Path Identifier Mismatch". SR Policy Candidate Path Identifier consist of:

*Protocol Origin.

*Originator.

*Discriminator.

3.3. 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:

*Preference.

*Optionally, the SR Policy Candidate Path name.

*Optionally, the SR Policy name.

4. SR Policy Association

As per [RFC8697], LSPs are associated with other LSPs with which they interact by adding them to a common association group. As described in [RFC8697], the association group is uniquely identified by the combination of the following fields in the ASSOCIATION object: Association Type, Association ID, Association Source, and (if present) Global Association Source or Extended Association ID, referred to as Association Parameters.

[RFC8697] specify the ASSOCIATION Object with two Object-Types for IPv4 and IPv6 which includes the field "Association Type". This document defines a new Association type (6) "SR Policy Association" for SRPA.

[RFC8697] specifies the mechanism for the capability advertisement of the Association Types supported by a PCEP speaker by defining an ASSOC-Type-List TLV to be carried within an OPEN object. This capability exchange for the SR Policy Association Types MUST be done before using the SRPA. Thus, the PCEP speaker MUST include the SRPA Type (6) in the ASSOC-Type-List TLV and MUST receive the same from the PCEP peer before using SRPA.

A given LSP MUST belong to at most one SRPA, since an SR Policy Candidate Path cannot belong to multiple SR Policies. If a PCEP speaker receives a PCEP message requesting to join more than one SRPA 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".

4.1. Association Parameters

As per [RFC9256], an SR Policy is identified through the tuple <headend, color, endpoint>. The headend is encoded in the 'Association Source' field in the ASSOCIATION object and the color and endpoint are encoded as part of the Extended Association ID TLV.

The Association Parameters (see Section 2) consist of:

- *Association Type: set to 6 "SR Policy Association".
- *Association Source (IPv4/IPv6): set to the headend IP address.
- *Association ID (16-bit): set to "1" (this 16-bit field is not utilized, just set to a fixed value).
- *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 [RFC9256] Section 2.1.

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:

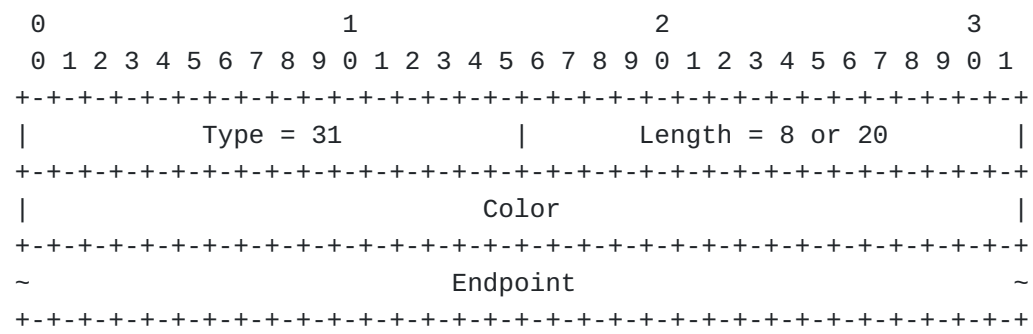

Figure 1: Extended Association ID TLV format

Type: Extended Association ID TLV, type = 31 [RFC8697].

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

Color: SR Policy color value, non-zero as per [RFC9256] Section 2.1.

Endpoint: can be either IPv4 or IPv6. This value MAY be different from the one contained in the Destination address field in the END-POINTS object, or in the Tunnel Endpoint Address field in the LSP-IDENTIFIERS TLV.

If the PCEP speaker receives an SRPA object 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 = 20 "SR Policy Identifier 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 associate 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 based on the SR Policy [RFC9256] parameters. Thus, there is no chance that different PCEP speakers will come up with different Association Parameters for the same SR Policy.

The last hop of the computed SR Policy Candidate Path MAY differ from the Endpoint contained in the <headend, color, endpoint> tuple. An example use case is to terminate the SR Policy before reaching the Endpoint and have decapsulated traffic go the rest of the way to the Endpoint node using the native IGP path(s). In this example, the destination of the SR Policy Candidate Paths will be some node before the Endpoint, but the Endpoint value is still used at the head-end to steer traffic with that Endpoint IP into the SR Policy. Destination of the SR Policy Candidate Path is signaled using the END-POINTS object and/or LSP-IDENTIFIERS TLV, as per the usual PCEP procedures. When neither END-POINTS object nor LSP-IDENTIFIERS TLV

is present, the PCEP speaker MUST extract the destination from the Endpoint field in the SRPA Extended Association ID TLV.

4.2. Association Information

The SRPA object may carry the following TLVs:

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

Out of these TLVs, the SRPOLICY-CPATH-ID TLV is mandatory, all others are optional. When a mandatory TLV is missing from the SRPA object, the PCEP speaker MUST send a PCErr message with Error-Type = 6 "Mandatory Object Missing", Error-Value = 21 "Missing SR Policy Mandatory TLV".

This document specifies four new TLVs to be carried in the SRPA object. Only one TLV instance of each type can be carried, and only the first occurrence is processed. Any others MUST be ignored.

4.2.1. SR Policy Name TLV

The SRPOLICY-POL-NAME TLV is an optional TLV for the SRPA object.

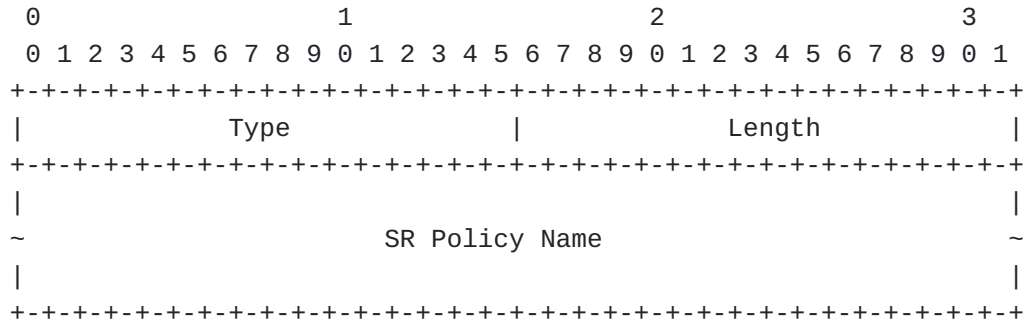

Figure 2: The SRPOLICY-POL-NAME TLV format

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 [RFC9256]. It SHOULD be a string of printable ASCII characters, without a NULL terminator.

4.2.2. SR Policy Candidate Path Identifier TLV

The SRPOLICY-CPATH-ID TLV is a mandatory TLV for the SRPA object.

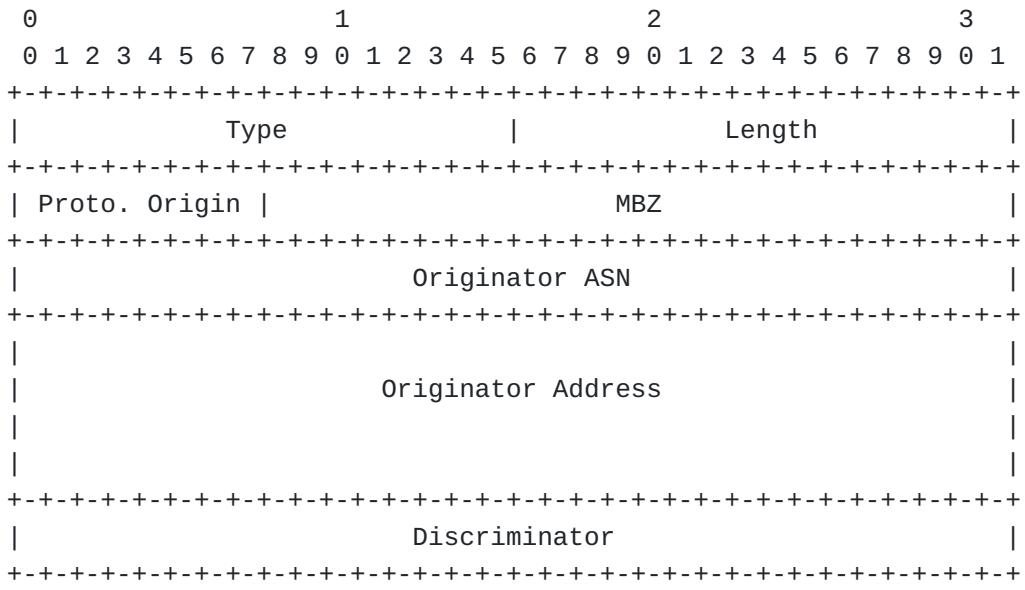

Figure 3: The SRPOLICY-CPATH-ID TLV format

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

Length: 28.

Protocol Origin: 8-bit value that encodes the protocol origin, as specified in Section 6.5. Note that in PCInitiate message [RFC8281], the Protocol Origin is always set to 10 (PCEP).

MBZ: Must be zero.

Originator ASN: Represented as 4-byte number, part of the originator identifier, as specified in [RFC9256] Section 2.4. If 2-byte ASNs are in use, the low-order 16 bits is used, and the high-order bits are set to 0. When sending PCInitiate message [RFC8281], the PCE is acting as the originator and therefore MUST set this to an ASN that it belongs to.

Originator Address: Represented as 128-bit value where IPv4 address is encoded in lowest 32 bits and high-order bits are set to 0, part of the originator identifier, as specified in [RFC9256] Section 2.4. When sending PCInitiate message, the PCE is acting as the originator and therefore MUST set this to an address that it owns.

Discriminator: 32-bit value that encodes the Discriminator of the Candidate Path, as specified in [RFC9256] Section 2.5. This is the field that mainly distinguishes different SR Candidate Paths, coming from the same originator. It is allowed to be any number in the 32-bit range.

4.2.3. SR Policy Candidate Path Name TLV

The SRPOLICY-CPATH-NAME TLV is an optional TLV for the SRPA object.

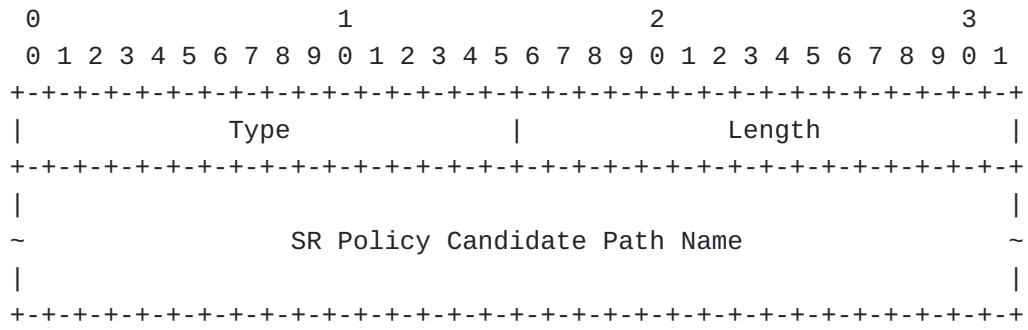

Figure 4: The SRPOLICY-CPATH-NAME TLV format

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 [RFC9256]. It SHOULD be a string of printable ASCII characters, without a NULL terminator.

4.2.4. SR Policy Candidate Path Preference TLV

The SRPOLICY-CPATH-PREFERENCE TLV is an optional TLV for the SRPA object. If the TLV is absent, then default Preference value is 100, as per Section 2.7 of [RFC9256].

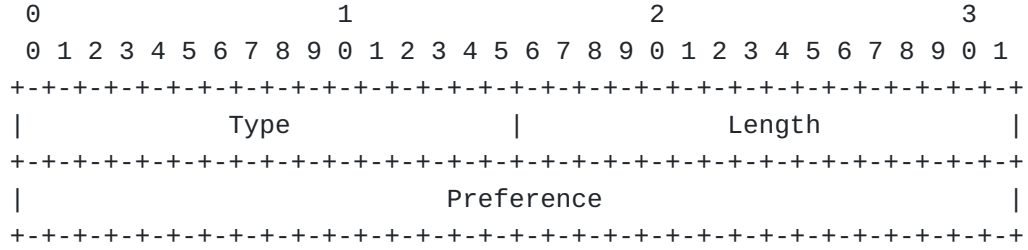

Figure 5: The SRPOLICY-CPATH-PREFERENCE TLV format

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

Length: 4.

Preference: Numerical preference of the Candidate Path as defined in Section 2.7 of [RFC9256].

5. Other Mechanisms

This section describes mechanisms that are standardized for SR Policies in [RFC9256], but do not make use of the SRPA for signaling in PCEP. Since SRPA is not used, there needs to be a separate capability negotiation.

This document specifies four new TLVs to be carried in the OPEN or LSP object. Only one TLV instance of each type can be carried, and only the first occurrence is processed. Any others MUST be ignored.

5.1. SR Policy Capability TLV

The SRPOLICY-CAPABILITY TLV is an optional TLV for the OPEN object. It is used at session establishment time to learn the other PCEP peer's capabilities with respect to SR Policy.

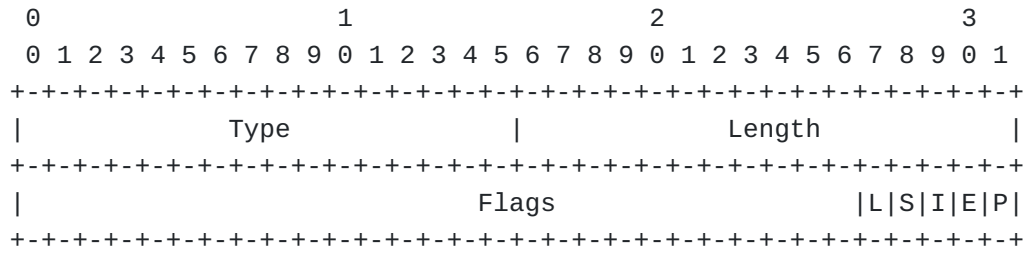

Figure 6: The SRPOLICY-CAPABILITY TLV format

Type: 71 for "SRPOLICY-CAPABILITY TLV.

Length: 4.

P-flag: If set to '1' by a PCEP speaker, the P flag indicates that the PCEP speaker supports the handling of COMPUTATION-PRIORITY TLV for the SR Policy, see Section 5.2. If this flag is not set, then the PCEP speaker MUST NOT send the COMPUTATION-PRIORITY TLV and SHOULD ignore it on receipt.

E-Flag: If set to '1' by a PCEP speaker, the E flag indicates that the PCEP speaker supports the handling of ENLP TLV for the SR Policy, see Section 5.3. If this flag is not set, then the PCEP speaker MUST NOT send the ENLP TLV and SHOULD ignore it on receipt.

I-Flag: If set to '1' by a PCEP speaker, the I flag indicates that the PCEP speaker supports the handling of INVALIDATION TLV for the SR Policy, see Section 5.4. If this flag is not set, then the PCEP speaker MUST NOT send the INVALIDATION TLV and SHOULD ignore it on receipt.

S-Flag: If set to '1' by a PCEP speaker, the S flag indicates that the PCEP speaker supports the handling of "Specified-BSID-only" behavior for the SR Policy, see Section 5.5. If this flag is not set, then the PCEP speaker MUST NOT set the Specified-BSID-only flag in the TE-PATH-BINDING TLV and SHOULD ignore it on receipt.

L-Flag: If set to '1' by a PCEP speaker, the L flag indicates that the PCEP speaker supports the stateless (PCReq/PCRep) operations for the SR Policy, see Section 5.6. If the PCE did not set this flag then the PCC SHOULD NOT send PCReq messages to this PCE for the SR Policy.

Unassigned bits MUST be set to '0' on transmission and MUST be ignored on receipt.

5.2. 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 [RFC9256]. If the TLV is absent from the LSP object, a default Priority value of 128 is used.

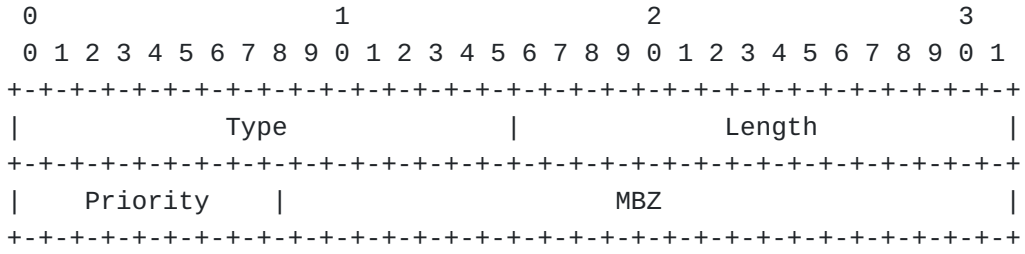

Figure 7: The COMPUTATION-PRIORITY TLV format

Type: 68 for "COMPUTATION-PRIORITY" TLV.

Length: 4.

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

5.3. Explicit Null Label Policy (ENLP) TLV

To steer an unlabeled IP packet into an SR policy, it is necessary to create a label stack for that packet, and push one or more labels

onto that stack. The Explicit NULL Label Policy (ENLP) TLV is an optional TLV used to indicate whether an Explicit NULL Label [RFC3032] must be pushed on an unlabeled IP packet before any other labels. The contents of this TLV are used by the SRPM as described in section 4.1 of [RFC9256]. If an ENLP TLV is not present, the decision of whether to push an Explicit NULL label on a given packet is a matter of local configuration. Note that Explicit Null is currently only defined for SR MPLS and not for SRv6. Therefore the PCEP speaker SHOULD ignore the presence of this TLV for SRv6 Policies.

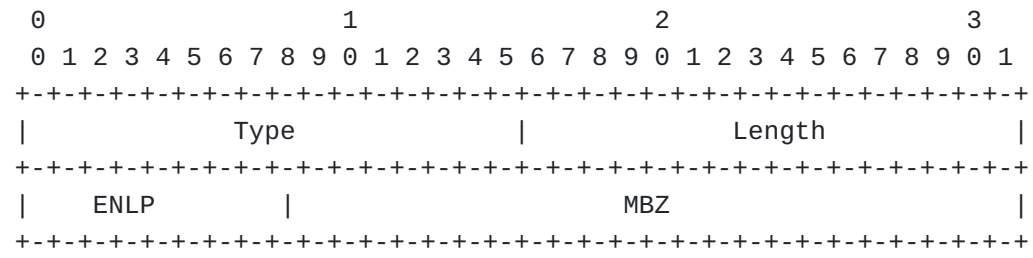

Figure 8: The Explicit Null Label Policy (ENLP) TLV format

Type: 69 for "ENLP" TLV.

Length: 4.

ENLP (Explicit NULL Label Policy): Indicates whether Explicit NULL labels are to be pushed on unlabeled IP packets that are being steered into a given SR policy. The values of this field are specified in section Section 6.6.

The ENLP reserved values may be used for future extensions and implementations SHOULD ignore the ENLP TLV with these values. The behavior signaled in this TLV MAY be overridden by local configuration. The section 4.1 of [RFC9256] describes the behavior on the headend for the handling of the explicit null label.

5.4. 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 facilitates the Drop-upon-invalid behavior, specified in Section 8.2 of [RFC9256]. Normally, if the LSP is down/invalid then it stops attracting traffic and traffic that would have been destined to that LSP is redirected somewhere else, such as via IGP or via another LSP. The Drop-upon-invalid behavior specifies that the LSP keeps attracting traffic and the traffic has to be dropped at the head-end. Such an LSP is said to be "in drop state". While in the drop state, the LSP operational state is "UP", as indicated by

the O-flag in the LSP object. However the ERO object MAY be empty, if no valid path has been computed.

The INVALIDATION TLV is used in both directions between PCEP peers:

- *PCE -> PCC: PCE specifies to the PCC whether to enable or disable Drop-upon-invalid (Config).

- *PCC -> PCE: PCC reports the current setting of the Drop-upon-invalid (Config) and also whether the LSP is currently in the drop state (Oper).

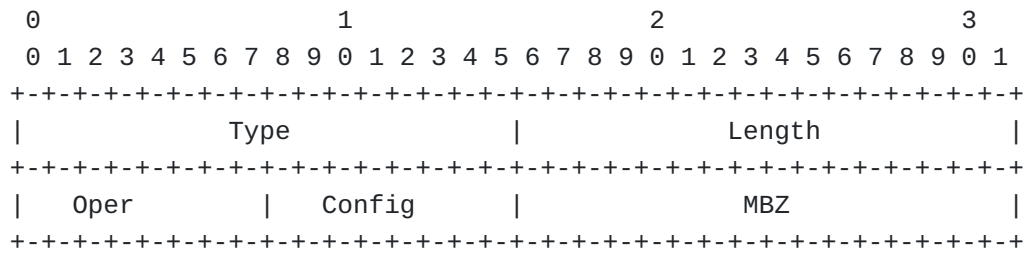

Figure 9: The INVALIDATION TLV format

Type: 70 for "INVALIDATION" TLV.

Length: 4.

Oper: encodes the current state of the LSP, i.e., whether it is actively dropping traffic right now. This field can be set to non-zero values only by the PCC, it MUST be set to 0 by the PCE and SHOULD be ignored by the PCC.

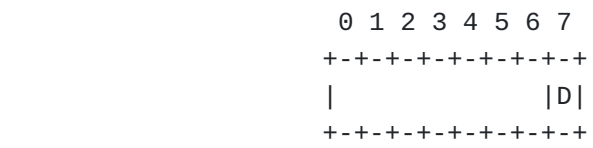

Figure 10: Oper state of Drop-upon-invalid feature

- *D: dropping - the LSP is currently attracting traffic and actively dropping it.

- *The unassigned bits in the Flag octet MUST be set to zero upon transmission and MUST be ignored upon receipt.

Config: encodes the current setting of the Drop-upon-invalid feature.

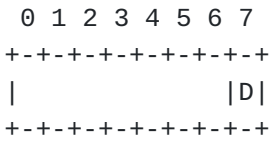

Figure 11: Config state of Drop-upon-invalid feature

*D: drop enabled - the Candidate Path has Drop-upon-invalid feature enabled.

*The unassigned bits in the Flag octet MUST be set to zero upon transmission and MUST be ignored upon receipt.

5.4.1. Drop-upon-invalid applies to SR Policy

The Drop-upon-invalid feature is somewhat special among the other SR Policy features in the way that it is enabled/disabled. This feature is enabled only on the whole SR Policy, not on a particular Candidate Path of that SR Policy, i.e., when ANY Candidate Path has Drop-upon-invalid enabled, it means that essentially the whole SR Policy has the feature enabled. As stated in [RFC9256] Section 8.1, the SR Policy is invalid when all its Candidate Paths are invalid, therefore all Candidate Paths MUST be attempted for bringup before the SR Policy is declared invalid.

Once all the Candidate Paths of the SR Policy have become invalid, then the SR Policy checks whether ANY of the Candidate Paths have Drop-upon-invalid enabled. If so, SR Policy enters the drop state and "activates" the highest preference Candidate Path which has the Drop-upon-invalid enabled. Note that only one Candidate Path needs to be reported to the PCE with the D (dropping) flag set.

5.5. Specified-BSID-only

Specified-BSID-only functionality is defined in Section 6.2.3 of [RFC9256]. 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 [I-D.ietf-pce-binding-label-sid]. 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.

5.6. Stateless Operation

[RFC8231] Section 5.8.2, allows delegation of an LSP in operationally down state, but at the same time mandates the use of PCReq before sending PCRpt. This document updates [RFC8231] Section 5.8.2, by making this section not applicable to SR Policy LSPs. Thus, when a PCC wants to delegate an SR Policy LSP, it MAY proceed directly to sending PCRpt, without first sending PCReq and waiting for PCRep. This has the advantage of reducing the number of PCEP messages and simplifying the implementation.

Furthermore, a PCEP speaker is not required to support PCReq/PCRep at all for SR Policies. The PCEP speaker can indicate support for PCReq/PCRep via the "L-Flag" in the SRPOLICY-CAPABILITY TLV (See Section 5.1). When this flag is cleared, or when the SRPOLICY-CAPABILITY TLV is absent, the given peer SHOULD NOT be sent PCReq/PCRep messages for SR Policy LSPs. Conversely when this flag is set, the peer can receive and process PCReq/PCRep messages for SR Policy LSPs.

The above applies only to SR Policy LSPs and does not affect other LSP types, such as RSVP-TE LSPs. For other LSP types, [RFC8231] Section 5.8.2 continues to apply.

6. IANA Considerations

6.1. 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

6.2. PCEP TLV Type Indicators

This document defines eight 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" subregistry 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
68	COMPUTATION-PRIORITY	This.I-D
69	EXPLICIT-NULL-LABEL-POLICY	This.I-D
70	INVALIDATION	This.I-D
71	SRPOLICY-CAPABILITY	This.I-D

6.3. 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" subregistry of the PCEP Numbers registry, as follows:

Error-Type	Meaning	Error-value	Reference
6	Mandatory Object Missing		[RFC5440]
		21: Missing SR Policy Mandatory TLV	This.I-D
26	Association Error		[RFC8697]
		20: SR Policy Identifers Mismatch	This.I-D
		21: SR Policy Candidate Path Identifier Mismatch	This.I-D

6.4. TE-PATH-BINDING TLV Flag field

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

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

6.5. SR Policy Candidate Path Protocol Origin field

This document requests IANA to maintain a new registry under "Segment Routing" registry group. New values are to be assigned by "Standards Action" [RFC8126]. The new subregistry is requested to be created under it be called "SR Policy Protocol Origin". The subregistry contains the following codepoints, with initial values, to be assigned by IANA with the reference set to this document:

Value	Description
0	Reserved (not to be used)
1-9	Unassigned
10	PCEP
11-19	Unassigned
20	BGP SR Policy
21-29	Unassigned
30	Configuration (CLI, YANG model via NETCONF, etc.)
31-250	Unassigned
251 - 255	Private Use (not to be assigned by IANA)

6.6. SR Policy Explicit Null Label Policy field

This document requests IANA to maintain a new registry under "Segment Routing" registry group. New values are to be assigned by "Standards Action" [RFC8126]. The new subregistry is requested to be created under it be called "SR Policy Explicit Null Label Policy".

The subregistry contains the following codepoints, with initial values, to be assigned by IANA with the reference set to this document:

Value	Description
0	Reserved (not to be used).
1	Push an IPv4 Explicit NULL label on an unlabeled IPv4 packet, but do not push an IPv6 Explicit NULL label on an unlabeled IPv6 packet.
2	Push an IPv6 Explicit NULL label on an unlabeled IPv6 packet, but do not push an IPv4 Explicit NULL label on an unlabeled IPv4 packet.
3	Push an IPv4 Explicit NULL label on an unlabeled IPv4 packet, and push an IPv6 Explicit NULL label on an unlabeled IPv6 packet.
4	Do not push an Explicit NULL label.
5 - 255	Reserved.

7. 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".

7.1. Cisco

*Organization: Cisco Systems

*Implementation: IOS-XR PCC and PCE.

*Description: All features supported except Computation Priority, Explicit NULL and Invalidation Drop.

*Maturity Level: Production.

*Coverage: Full.

*Contact: ssidor@cisco.com

7.2. Juniper

*Organization: Juniper Networks

*Implementation: PCC and PCE.

*Description: Everything in -05 except SR Policy Name TLV and SR Policy Candidate Path Name TLV.

*Maturity Level: Production.

*Coverage: Partial.

*Contact: cbarth@juniper.net

8. Security Considerations

The information carried in the newly defined SRPA object and TLVs could provide an eavesdropper with additional information about the SR Policy. Thus securing the PCEP session using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525], is RECOMMENDED.

9. Acknowledgement

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

10. References

10.1. 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>.

[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/info/rfc9256>.

[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>.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <https://www.rfc-editor.org/info/rfc3032>.

[RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", 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>.

[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/info/rfc8126>.

[I-D.ietf-pce-binding-label-sid] Sivabalan, S., Filsfils, C., Tantsura, J., Previdi, S., and C. Li, "Carrying Binding Label/Segment Identifier (SID) in PCE-based Networks.", Work in Progress, Internet-Draft, draft-ietf-pce-binding-label-sid-16, 27 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-binding-label-sid-16>.

10.2. Informative References

[I-D.ietf-pce-multipath] Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., Peng, S., and G. S. Mishra, "PCEP Extensions for Signaling Multipath Information", Work in Progress, Internet-Draft, draft-ietf-pce-multipath-10, 16 January 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-10>.

Appendix A. Contributors

Dhruv Dhody
Huawei
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

Rajesh Melarcode
Cisco Systems, Inc.
2000 Innovation Dr.
Kanata, Ontario
Canada

Email: rmelarco@cisco.com

Authors' Addresses

Mike Koldychev
Ciena Corporation
385 Terry Fox Dr.
Kanata Ontario K2K 0L1
Canada

Email: mkoldych@proton.me

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