

PCE Working Group
Internet-Draft
Updates: 8231 (if approved)
Intended status: Standards Track
Expires: 6 October 2025

M. Koldychev
S. Sivabalan
Ciena Corporation
S. Sidor
Cisco Systems, Inc.
C. Barth
Juniper Networks, Inc.
S. Peng
Huawei Technologies
H. Bidgoli
Nokia
4 April 2025

Path Computation Element Communication Protocol (PCEP) Extensions for
Segment Routing (SR) Policy Candidate Paths
draft-ietf-pce-segment-routing-policy-cp-27

Abstract

A Segment Routing (SR) Policy is an ordered list of instructions, called "segments" that represent a source-routed policy. Packet flows are steered into an SR Policy on a node where it is instantiated. An SR Policy is made of one or more candidate paths.

This document specifies the Path Computation Element Communication Protocol (PCEP) extension to signal candidate paths of an SR Policy. Additionally, this document updates RFC 8231 to allow delegation and setup of an SR Label Switched Path (LSP), without using the path computation request and reply messages. This document is applicable to both Segment Routing over MPLS (SR-MPLS) and Segment Routing over IPv6 (SRv6).

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 6 October 2025.

Copyright Notice

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

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

Table of Contents

1. Introduction	3
1.1. Requirements Language	4
2. Terminology	4
3. Overview	5
4. SR Policy Association (SRPA)	6
4.1. SR Policy Identifier	7
4.2. SR Policy Candidate Path Identifier	7
4.3. SR Policy Candidate Path Attributes	7
4.4. Association Parameters	8
4.5. Association Information	9
4.5.1. SR Policy Name TLV	10
4.5.2. SR Policy Candidate Path Identifier TLV	10
4.5.3. SR Policy Candidate Path Name TLV	12
4.5.4. SR Policy Candidate Path Preference TLV	12
5. SR Policy Signaling Extensions	13
5.1. SR Policy Capability TLV	13
5.2. LSP Object TLVs	15
5.2.1. Computation Priority TLV	15
5.2.2. Explicit Null Label Policy (ENLP) TLV	15
5.2.3. Invalidation TLV	16
5.2.3.1. Drop-upon-invalid applies to SR Policy	18
5.3. Update to RFC 8231	18
6. IANA Considerations	19
6.1. Association Type	19
6.2. PCEP TLV Type Indicators	19
6.3. PCEP Errors	20
6.4. TE-PATH-BINDING TLV Flag field	21
6.5. SR Policy Invalidation Operational State	21
6.6. SR Policy Invalidation Configuration State	22
6.7. SR Policy Capability TLV Flag field	22

7. Implementation Status	23
7.1. Cisco	23
7.2. Juniper	24
8. Security Considerations	24
9. Manageability Considerations	24
9.1. Control of Function and Policy	25
9.2. Information and Data Models	25
9.3. Liveness Detection and Monitoring	25
9.4. Verify Correct Operations	25
9.5. Requirements On Other Protocols	25
9.6. Impact On Network Operations	25
10. Acknowledgement	26
11. References	26
11.1. Normative References	26
11.2. Informative References	28
Appendix A. Contributors	29
Authors' Addresses	30

1. Introduction

Segment Routing (SR) Policy Architecture [RFC9256] details the concepts of Segment Routing (SR) Policy [RFC8402] and approaches to steering traffic into an SR Policy.

Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing [RFC8664] specifies extensions to the PCEP that allow a stateful Path Computation Element (PCE) to compute and initiate Traffic Engineering (TE) paths, as well as a Path Computation Client (PCC) to request a path subject to certain constraints and optimization criteria in SR domain. Although PCEP extensions introduced in [RFC8664] enables the creation of SR-TE paths, these do not constitute SR Policies as defined in [RFC9256] and therefore lack support for:

- * Association of SR Policy Candidate Paths signaled via PCEP with Candidate Paths of the same SR Policy signaled via other sources (e.g., local configuration or BGP).
- * Association of SR Policy with an intent via color, enabling headend-based steering of BGP service routes over SR Policies provisioned via PCEP.

PCEP Extensions for establishing relationships between sets of Label Switched Paths (LSPs) [RFC8697] introduces a generic mechanism to create a grouping of LSPs which is called an Association.

An SR Policy is associated with one or more candidate paths. A candidate path is the unit for signaling of an SR Policy to a headend as described in Section 2.2 of [RFC9256]. This document extends [RFC8664] to support signaling SR Policy Candidate Paths as LSPs and to signal Candidate Path membership in an SR Policy by means of the Association mechanism. A PCEP Association corresponds to a SR Policy and a LSP corresponds to a Candidate Path. The unit of signaling in PCEP is the LSP, thus all the information related to SR Policy is carried at the Candidate Path level.

Also, this document updates Section 5.8.2 of [RFC8231], making the use of Path Computation Request (PCReq) and Path Computation Reply (PCRep) messages optional for LSPs setup using Path Setup Type 1 (Segment Routing) [RFC8664] and Path Setup Type 3 (SRv6) [RFC9603] with the aim of reducing the PCEP message exchanges and simplifying implementation.

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.

2. Terminology

This document uses the following terms defined in [RFC5440]: ERO, PCC, PCE, PCEP Peer, and PCEP speaker.

This document uses the following term defined in [RFC3031]: LSP.

This document uses the following term defined in [RFC9552]: BGP-LS.

The following terms are used in this document:

Endpoint: The IPv4 or IPv6 endpoint address of an SR Policy, as described in Section 2.1 of [RFC9256].

Color: The 32-bit color of an SR Policy, as described in Section 2.1 of [RFC9256].

Protocol-Origin: The protocol that was used to create a Candidate Path, as described in Section 2.3 of [RFC9256].

Originator: A device that created a Candidate Path, as described in Section 2.4 of [RFC9256].

Discriminator: Distinguishes Candidate Paths created by the same device, as described in Section 2.5 of [RFC9256].

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

Association Information: As described in Section 6.1.4 of [RFC8697], refers to information related to Association Type.

SR Policy LSP: An LSP setup using Path Setup Type [RFC8408] 1 (Segment Routing) or 3 (SRv6).

SR Policy Association: A new association type used to group candidate paths belonging to same SR Policy. Depending on the 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.

The base PCEP specification [RFC4655] originally defined the use of the PCE architecture for MPLS and GMPLS networks with LSPs instantiated using the RSVP-TE signaling protocol. Over time, support for additional path setup types, such as SRv6, has been introduced [RFC9603]. The term "LSP" is used extensively in PCEP specifications and, in the context of this document, refers to a Candidate Path within an SR Policy, which may be an SRv6 path (still represented using the LSP Object as specified in [RFC8231]).

3. Overview

The SR Policy is represented by a new type of PCEP Association, called the SR Policy Association (SRPA) (see Section 4). The SR Policy Candidate Paths of specific SR Policy are the LSPs within the same SRPA. The extensions in this document specify the encoding of a single segment list within an SR Policy Candidate Path. Encoding of multiple segment lists is outside the scope of this document and specified in [I-D.ietf-pce-multipath].

An 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 an SRPA: Computation Priority of the LSP, Explicit Null Label Policy for the unlabeled IP packets and Drop-upon-invalid behavior for traffic steering when the LSP is operationally down (see Section 5).

4. SR Policy Association (SRPA)

Per [RFC8697], LSPs are associated with other LSPs with which they interact by adding them to a common association group. An association group is uniquely identified by the combination of the following fields in the ASSOCIATION object (Section 6.1 of [RFC8697]): Association Type, Association ID, Association Source, and (if present) Global Association Source, or Extended Association ID. These fields are referred to as Association Parameters (Section 4.4).

[RFC8697] specifies 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 Type MUST be done before using the SRPA. To that aim, a 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 the SRPA (Section 6.1).

SRPA MUST be assigned for all SR Policy LSPs by PCEP speaker originating the LSP if capability was advertised by both PCEP speakers. If the above condition is not satisfied, then the receiving PCEP speaker MUST send a PCERR message with Error-Type = 6 "Mandatory Object Missing", Error-Value = TBD1 "Missing SR Policy Association".

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".

The existing behavior for the use of Binding SID with SR Policy is already documented in [RFC9604]. If BSID value allocation failed, because of conflict with BSID used by another policy, then PCEP peer MUST send a PCERR message with Error-Type = 32 "Binding label/SID failure" and Error-value = 2 "Unable to allocate the specified binding value".

4.1. SR Policy Identifier

SR Policy Identifier uniquely identifies an SR Policy [RFC9256] within the SR domain. SR Policy Identifier is assigned by PCEP peer originating the LSP and MUST be uniform across all the PCEP sessions. Candidate Paths within an SR Policy MUST carry the same SR Policy Identifiers in their SRPAs. Candidate Paths within an SR Policy MUST NOT change their SR Policy Identifiers for the lifetime of the PCEP session. If the above conditions are not satisfied, the receiving PCEP speaker MUST send a PCEP Error (PCErr) message with Error-Type = 26 "Association Error" and Error Value = 20 "SR Policy Identifier Mismatch". SR Policy Identifier consists of:

- * Headend router where the SR Policy originates.
- * Color of the SR Policy ([RFC9256], Section 2.1).
- * Endpoint of the SR Policy ([RFC9256], Section 2.1).

4.2. 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 is assigned by PCEP peer originating the LSP. Candidate Paths within an SR Policy MUST NOT change their SR Policy Candidate Path Identifiers for the lifetime of the PCEP session. Two or more Candidate Paths within an SR Policy MUST NOT carry same SR Policy Candidate Path Identifiers in their SRPAs. If the above conditions are not satisfied, the PCEP speaker MUST send a PCErr message with Error-Type = 26 "Association Error" and Error Value = 21 "SR Policy Candidate Path Identifier Mismatch". SR Policy Candidate Path Identifier consists of:

- * Protocol Origin ([RFC9256], Section 2.3).
- * Originator ([RFC9256], Section 2.4).
- * Discriminator ([RFC9256], Section 2.5).

4.3. SR Policy Candidate Path Attributes

SR Policy Candidate Path Attributes carry optional, non-key information about a Candidate Path and MAY change during the lifetime of an LSP. SR Policy Candidate Path Attributes consists of:

- * Candidate Path preference ([RFC9256], Section 2.7).
- * Candidate Path name ([RFC9256], Section 2.6).

- * SR Policy name ([RFC9256], Section 2.1).

4.4. Association Parameters

Per Section 2.1 of [RFC9256], an SR Policy is identified through the <headend, color, endpoint> tuple.

The Association Parameters consists of:

- * Association Type: Set to 6 "SR Policy Association".
- * Association Source (IPv4/IPv6): Set to the headend value of the SR Policy, as defined in [RFC9256] Section 2.1.
- * Association ID (16-bit): Always set to the numeric value "1".
- * Extended Association ID TLV: Mandatory TLV for SR Policy Association. Encodes the Color and Endpoint of the SR Policy (Figure 1).




Figure 1: Extended Association ID TLV Format

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

Length: 8 octets if IPv4 address or 20 octets if IPv6 address is encoded in the Endpoint field.

Color: unsigned non-zero 32-bit integer value, SR Policy color per Section 2.1 of [RFC9256].

Endpoint: can be either IPv4 (4 octets) or IPv6 address (16 octets). 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 (Section 2.1 of [RFC9256]).

If a PCEP speaker receives an SRPA object whose Association Parameters do not follow the above specification, then the PCEP speaker MUST send a PCErr message with Error-Type = 26 "Association Error", Error-Value = 20 "SR Policy Identifier Mismatch".

The encoding choice of the Association Parameters in this way is meant 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 parameters [RFC9256].

The last hop of a 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 be forwarded the rest of the path to the Endpoint node using the native Interior Gateway Protocol (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 headend to steer traffic with that Endpoint IP address into the SR Policy. The Destination of the SR Policy Candidate Path is signaled using the END-POINTS object and/or LSP-IDENTIFIERS TLV, per the usual PCEP procedure. When neither the 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.

SR Policy with Color-Only steering is signaled with the Endpoint value set to unspecified, i.e., 0.0.0.0 for IPv4 or :: for IPv6, per Section 8.8. of [RFC9256].

4.5. Association Information

The SRPA object may carry the following TLVs:

- * SRPOLICY-POL-NAME TLV (Section 4.5.1): (optional) encodes the SR Policy Name string.
- * SRPOLICY-CPATH-ID TLV (Section 4.5.2): (mandatory) encodes the SR Policy Candidate Path Identifier.
- * SRPOLICY-CPATH-NAME TLV (Section 4.5.3): (optional) encodes the SR Policy Candidate Path string name.
- * SRPOLICY-CPATH-PREFERENCE TLV (Section 4.5.4): (optional) encodes the SR Policy Candidate Path preference value.

When a mandatory TLV is missing from an 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".

Only one TLV instance of each TLV type can be carried in an SRPA object, and only the first occurrence is processed. Any others MUST be silently ignored.

4.5.1. SR Policy Name TLV

The SRPOLICY-POL-NAME TLV (Figure 2) is an optional TLV for the SRPA object. It is RECOMMENDED that the size of the name for the SR Policy is limited to 255 bytes. Implementations MAY choose to truncate long names to 255 bytes to simplify interoperability with other protocols.

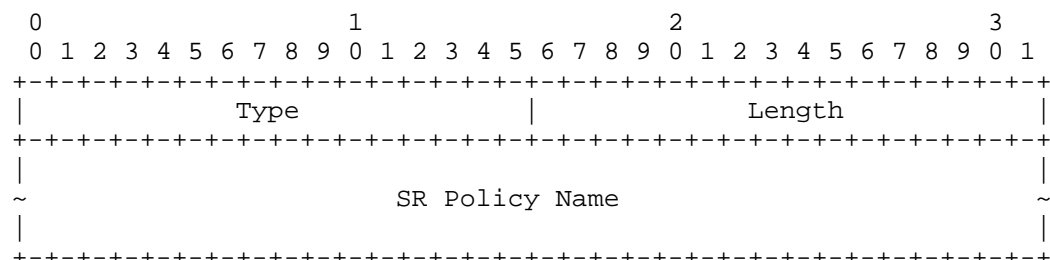


Figure 2: SRPOLICY-POL-NAME TLV Format

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. Padding is not included in the Length field.

SR Policy Name: SR Policy name, as defined in Section 2.1 of [RFC9256]. It MUST be a string of printable ASCII [RFC0020] characters, without a NULL terminator.

4.5.2. SR Policy Candidate Path Identifier TLV

The SRPOLICY-CPATH-ID TLV (Figure 3) is a mandatory TLV for the SRPA object.

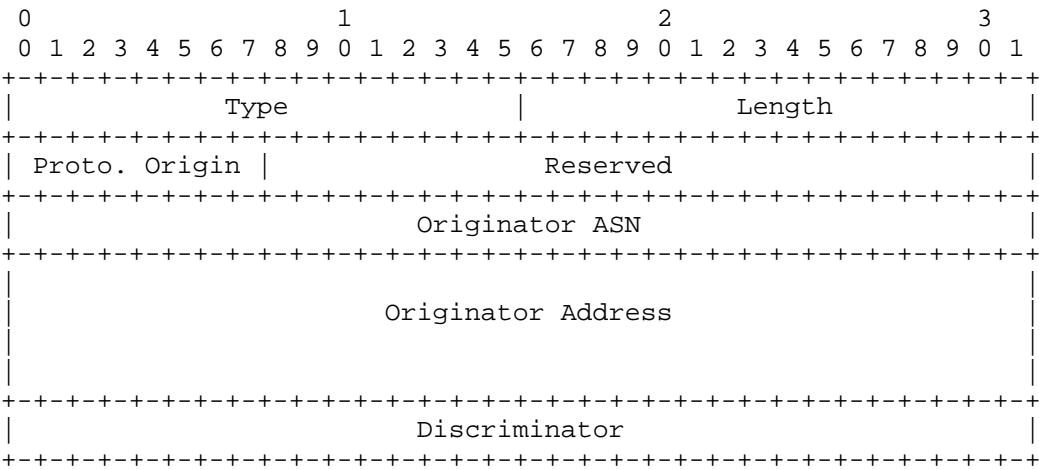


Figure 3: SRPOLICY-CPATH-ID TLV Format

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

Length: 28.

Protocol Origin: 8-bit unsigned integer value that encodes the protocol origin. The values of this field are specified in IANA registry "SR Policy Protocol Origin" under "Segment Routing" registry group, which was introduced in Section 8.4 of [I-D.ietf-idr-bgp-ls-sr-policy]. Note that in the PCInitiate message [RFC8281], the Protocol Origin is always set to 10 - "PCEP (In PCEP or when BGP-LS Producer is PCE)". The "SR Policy Protocol Origin" IANA registry includes a combination of values intended for use in PCEP and BGP-LS. When the registry contains two variants of values associated with the mechanism or protocol used for provisioning of the Candidate Path, for example 1 - "PCEP" and 10 - "PCEP (In PCEP or when BGP-LS Producer is PCE)", the "(In PCEP or when BGP-LS Producer is PCE)" variants MUST be used in PCEP.

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

Originator Autonomous System Number (ASN): Represented as a 32-bit unsigned integer value, part of the originator identifier, as specified in Section 2.4 of [RFC9256]. When sending a PCInitiate message [RFC8281], the PCE is the originator of the Candidate Path. If the PCE is configured with an ASN, then it MUST set it, otherwise the ASN is set to 0.

Originator Address: Represented as a 128-bit value as specified in Section 2.4 of [RFC9256]. When sending a PCInitiate message, the PCE is acting as the originator and therefore MAY set this to an address that it owns.

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

4.5.3. SR Policy Candidate Path Name TLV

The SRPOLICY-CPATH-NAME TLV (Figure 4) is an optional TLV for the SRPA object. It is RECOMMENDED that the size of the name for the SR Policy is limited to 255 bytes. Implementations MAY choose to truncate long names to 255 bytes to simplify interoperability with other protocols.




Figure 4: 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. Padding is not included in the Length field.

SR Policy Candidate Path Name: SR Policy Candidate Path Name, as defined in Section 2.6 of [RFC9256]. It MUST be a string of printable ASCII characters, without a NULL terminator.

4.5.4. SR Policy Candidate Path Preference TLV

The SRPOLICY-CPATH-PREFERENCE TLV (Figure 5) is an optional TLV for the SRPA object. If the TLV is absent, then default Preference value is 100, per Section 2.7 of [RFC9256].




Figure 5: SRPOLICY-CPATH-PREFERENCE TLV Format

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

Length: 4.

Preference: 32-bit unsigned integer value that encodes preference of the Candidate Path as defined in Section 2.7 of [RFC9256].

5. SR Policy Signaling Extensions

This section introduces mechanisms described for SR Policies in [RFC9256] to PCEP. These extensions do not make use of the SRPA for signaling in PCEP therefore cannot rely on the Association capability negotiation in ASSOC-Type-List TLV and separate capability negotiation is required.

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 (Figure 6) is a TLV for the OPEN object. It is used at session establishment to learn the peer's capabilities with respect to SR Policy. Implementations that support SR Policy MUST include SRPOLICY-CAPABILITY TLV in the OPEN object if the extension is enabled. In addition, the ASSOC-Type-List TLV containing SRPA Type (6) MUST be present in the OPEN object, as specified in Section 4.

If a PCEP speaker receives SRPA but the SRPOLICY-CAPABILITY TLV is not exchanged, then the PCEP speaker MUST send a PCErr message with Error- Type = 10 ("Reception of an invalid object") and Error-Value = TBD2 ("Missing SRPOLICY-CAPABILITY TLV") and MUST then close the PCEP session.




Figure 6: SRPOLICY-CAPABILITY TLV Format

Type: 71 for "SRPOLICY-CAPABILITY" TLV.

Length: 4.

Flags (32 bits):

The following flags are currently defined:

- * P-Flag (Computation Priority): 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 (Section 5.2.1). If this flag is set to 0, then the receiving PCEP speaker MUST NOT send the COMPUTATION-PRIORITY TLV and MUST ignore it on receipt.
- * E-Flag (Explicit NULL Label Policy): If set to '1' by a PCEP speaker, the E flag indicates that the PCEP speaker supports the handling of Explicit Null Label Policy (ENLP) TLV for the SR Policy (Section 5.2.2). If this flag is set to 0, then the receiving PCEP speaker MUST NOT send the ENLP TLV and MUST ignore it on receipt.
- * I-Flag (Invalidation): 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 (Section 5.2.3). If this flag is set to 0, then the receiving PCEP speaker MUST NOT send the INVALIDATION TLV and MUST ignore it on receipt.
- * L-Flag (Stateless Operation): 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 (Section 5.3). If the PCE set this flag to 0, then the PCC MUST 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. More flags can be assigned in the future per (Section 6.7).

5.2. LSP Object TLVs

This section is introducing three new TLVs to be carried in LSP object introduced in Section 7.3 of [RFC8231].

5.2.1. Computation Priority TLV

The COMPUTATION-PRIORITY TLV (Figure 7) is an optional TLV. 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 and the P-flag in the SRPOLICY-CAPABILITY TLV is set to 1, a default Priority value of 128 is used.




Figure 7: COMPUTATION-PRIORITY TLV Format

Type: 68 for "COMPUTATION-PRIORITY" TLV.

Length: 4.

Priority: 8-bit unsigned integer value that encodes numerical priority with which this LSP is to be recomputed by the PCE upon topology change. Lowest value is the highest priority.

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

5.2.2. Explicit Null Label Policy (ENLP) TLV

To steer an unlabeled IP packet into an SR policy for the MPLS data plane, it is necessary to push a label stack of one or more labels on that packet. The Explicit NULL Label Policy (ENLP) TLV is an optional TLV for the LSP object 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 SR Policy Manager 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 receiving PCEP speaker MUST ignore the presence of this TLV for SRv6 Policies.




Figure 8: Explicit Null Label Policy (ENLP) TLV Format

Type: 69 for "ENLP" TLV.

Length: 4.

ENLP (Explicit NULL Label Policy): 8-bit unsigned integer value that 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 IANA registry "SR Policy ENLP Values" under "Segment Routing" registry group, which was introduced in Section 6.10 of [I-D.ietf-idr-sr-policy-safil].

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

The ENLP unassigned values may be used for future extensions and implementations MUST ignore the ENLP TLV with unrecognized values. The behavior signaled in this TLV MAY be overridden by local configuration by the network operator based on their deployment requirements. The Section 4.1 of [RFC9256] describes the behavior on the headend for the handling of the explicit null label.

5.2.3. Invalidation TLV

The INVALIDATION TLV (Figure 9) is an optional TLV. This TLV is used to control traffic steering into an LSP 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; traffic that would have been destined for that LSP is redirected somewhere else, such as via IGP or another LSP. The Drop-upon-invalid behavior specifies that the LSP keeps attracting traffic and the traffic has to be dropped at the headend. 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: INVALIDATION TLV Format

Type: 70 for "INVALIDATION" TLV.

Length: 4.

Oper: An 8-bit flag field that encodes the operational state of the LSP. It MUST be set to 0 by the PCE when sending and MUST be ignored by the PCC upon receipt. See Section 6.5 for IANA information.




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

- * D: dropping - the LSP is actively dropping traffic as a result of Drop-upon-invalid behavior being activated.
- * The unassigned bits in the Flag octet MUST be set to zero upon transmission and MUST be ignored upon receipt.

Config: An 8-bit flag field that encodes the configuration of the LSP. See Section 6.6 for IANA information.




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.

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

5.2.3.1. 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 the whole SR Policy has the feature enabled. As stated in Section 8.1 of [RFC9256], an SR Policy is invalid when all its Candidate Paths are invalid.

Once all the Candidate Paths of an SR Policy have become invalid, then the SR Policy checks whether any of the Candidate Paths have Drop-upon-invalid enabled. If so, the 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.

5.3. Update to RFC 8231

Section 5.8.2 of [RFC8231], 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 Section 5.8.2 of [RFC8231], by making that section of [RFC8231] 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 MUST 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, Section 5.8.2 of [RFC8231] continues to apply.

6. IANA Considerations

IANA maintains the "Path Computation Element Protocol (PCEP) Numbers" registry at <<https://www.iana.org/assignments/pcep>>.

6.1. Association Type

This document defines a new association type: SR Policy Association. IANA is requested to confirm the following allocation in the "ASSOCIATION Type Field" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

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 Policy Candidate Paths. IANA is requested to confirm the following allocations in 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
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, two new Error-Values within the "Association Error" Error-Type and one new Error-Value within the "Reception of an invalid object".

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

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 Identifiers Mismatch	This.I-D
		21: SR Policy Candidate Path Identifier Mismatch	This.I-D

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

Error-Type	Meaning	Error-value	Reference
6	Mandatory Object Missing		[RFC5440]
		TBD1: Missing SR Policy Association	This.I-D
10	Reception of an invalid object		[RFC5440]
		TBD2: Missing SRPOLICY-CAPABILITY TLV	This.I-D

6.4. TE-PATH-BINDING TLV Flag field

An earlier version of this document added new bit within the "TE-PATH-BINDING TLV Flag field" registry of the "Path Computation Element Protocol (PCEP) Numbers" registry group, which was also early allocated by the IANA.

IANA is requested to mark the bit position as deprecated.

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

6.5. SR Policy Invalidation Operational State

This document requests IANA to maintain a new registry under "Path Computation Element Protocol (PCEP) Numbers" registry group. The new registry is called "SR Policy Invalidation Operational Flags". New values are to be assigned by "IETF review" [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit (counting from bit 0 as the most significant bit).
- * Description.
- * Reference.

Bit	Description	Reference
0 - 6	Unassigned	This.I-D
7	D: dropping - the LSP is currently attracting traffic and actively dropping it.	This.I-D

6.6. SR Policy Invalidation Configuration State

This document requests IANA to maintain a new registry under "Path Computation Element Protocol (PCEP) Numbers" registry group. The new registry is called "SR Policy Invalidation Configuration Flags". New values are to be assigned by "IETF review" [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit (counting from bit 0 as the most significant bit).
- * Description.
- * Reference.

Bit	Description	Reference
0 - 6	Unassigned.	This.I-D
7	D: drop enabled - the Drop-upon-invalid is enabled on the LSP.	This.I-D

6.7. SR Policy Capability TLV Flag field

This document requests IANA to maintain a new registry under "Path Computation Element Protocol (PCEP) Numbers" registry group. The new registry is called "SR Policy Capability TLV Flag Field". New values are to be assigned by "IETF review" [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit (counting from bit 0 as the most significant bit).
- * Description.
- * Reference.

Bit	Description	Reference
0 - 26	Unassigned	This.I-D
27	Stateless Operation (L-Flag)	This.I-D
28	Unassigned	This.I-D
29	Invalidation (I-Flag)	This.I-D
30	Explicit NULL Label Policy (E-Flag)	This.I-D
31	Computation Priority (P-flag)	This.I-D

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.

The security considerations described in [RFC5440], [RFC8231], [RFC8281], [RFC8664], [RFC8697], [RFC9256] and [RFC9603] 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] as per the recommendations and best current practices in [RFC9325].

9. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231], [RFC8664], [RFC9256], and [RFC9603] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

9.1. Control of Function and Policy

A PCE or PCC implementation MAY allow the capabilities specified in Section 5.1 and the capability for support of SRPA advertised in ASSOC-Type-List TLV to be enabled and disabled.

9.2. Information and Data Models

[I-D.ietf-pce-pcep-srv6-yang] defines YANG module with common building blocks for PCEP Extensions described in Section 4.

9.3. Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440], [RFC8664], and [RFC9256].

9.4. Verify Correct Operations

Operation verification requirements already listed in [RFC5440], [RFC8231], [RFC8664], [RFC9256], and [RFC9603] are applicable to mechanisms defined in this document.

An implementation MUST allow the operator to view SR Policy Identifier and SR Policy Candidate Path Identifier advertised in SRPA object.

An implementation SHOULD allow the operator to view the capabilities defined in this document advertised by each PCEP peer.

An implementation SHOULD allow the operator to view LSPs associated with specific SR Policy Identifier.

9.5. Requirements On Other Protocols

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

9.6. Impact On Network Operations

The mechanisms defined in [RFC5440], [RFC8231], [RFC9256] and [RFC9603] also apply to the PCEP extensions defined in this document.

10. Acknowledgement

We would like to thank Abdul Rehman, Andrew Stone, Boris Khasanov, Cheng Li, Dhruv Dhody, Gorrry Fairhurst, Gyan Mishra, Huaimo Chen, Ines Robles, Joseph Salowey, Ketan Talaulikar, Marina Fizgeer, Mike Bishopm, Praveen Kumar, Robert Sparks, Roman Danyliw, Stephane Litkowski, Tom Petch, Zoey Rose, Xiao Min, Xiong Quan for review and suggestions.

11. References

11.1. Normative References

- [RFC0020] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <<https://www.rfc-editor.org/info/rfc20>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8408] Sivabalan, S., Tantsura, J., Minei, I., Varga, R., and J. Hardwick, "Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages", RFC 8408, DOI 10.17487/RFC8408, July 2018, <<https://www.rfc-editor.org/info/rfc8408>>.
- [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>>.
- [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>>.
- [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>>.

- [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/info/rfc9325>>.
- [RFC9603] Li, C., Ed., Kaladharan, P., Sivabalan, S., Koldychev, M., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing", RFC 9603, DOI 10.17487/RFC9603, July 2024, <<https://www.rfc-editor.org/info/rfc9603>>.

11.2. Informative References

- [I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.
- [I-D.ietf-idr-bgp-ls-sr-policy]
Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.
- [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-12, 8 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-12>>.
- [I-D.ietf-pce-pcep-srv6-yang]
Li, C., Sivabalan, S., Peng, S., Koldychev, M., and L. Ndifor, "A YANG Data Model for Segment Routing (SR) Policy and SR in IPv6 (SRv6) support in Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-srv6-yang-06, 19 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-srv6-yang-06>>.

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., 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>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.
- [RFC9604] Sivabalan, S., Filsfils, C., Tantsura, J., Previdi, S., and C. Li, Ed., "Carrying Binding Label/SID in PCE-Based Networks", RFC 9604, DOI 10.17487/RFC9604, August 2024, <<https://www.rfc-editor.org/info/rfc9604>>.

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

Zafar Ali
Cisco Systems, Inc.

Email: zali@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

Samuel Sidor
Cisco Systems, Inc.
Eurovea Central 3.
811 09 Bratislava
Slovakia
Email: ssidor@cisco.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