

Internet Engineering Task Force (IETF)
Request for Comments: 8664
Updates: 8408
Category: Standards Track
ISSN: 2070-1721

S. Sivabalan
C. Filsfils
Cisco Systems, Inc.
J. Tantsura
Apstra, Inc.
W. Henderickx
Nokia
J. Hardwick
Metaswitch Networks
December 2019

Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing

Abstract

Segment Routing (SR) enables any head-end node to select any path without relying on a hop-by-hop signaling technique (e.g., LDP or RSVP-TE). It depends only on "segments" that are advertised by link-state Interior Gateway Protocols (IGPs). An SR path can be derived from a variety of mechanisms, including an IGP Shortest Path Tree (SPT), an explicit configuration, or a Path Computation Element (PCE). This document specifies extensions to the Path Computation Element Communication Protocol (PCEP) that allow a stateful 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 networks.

This document updates RFC 8408.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8664>.

Copyright Notice

Copyright (c) 2019 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

1.	Introduction
2.	Terminology
2.1.	Requirements Language
3.	Overview of PCEP Operation in SR Networks
4.	Object Formats
4.1.	The OPEN Object
4.1.1.	The Path Setup Type Capability TLV
4.1.2.	The SR PCE Capability Sub-TLV
4.2.	The RP/SRP Object
4.3.	ERO
4.3.1.	SR-ERO Subobject
4.3.2.	NAI Associated with SID
4.4.	RRO
4.5.	METRIC Object
5.	Procedures
5.1.	Exchanging the SR PCE Capability
5.2.	ERO Processing
5.2.1.	SR-ERO Validation
5.2.2.	Interpreting the SR-ERO
5.3.	RRO Processing
6.	Management Considerations
6.1.	Controlling the Path Setup Type
6.2.	Migrating a Network to Use PCEP Segment-Routed Paths
6.3.	Verification of Network Operation
6.4.	Relationship to Existing Management Models
7.	Security Considerations
8.	IANA Considerations
8.1.	PCEP ERO and RRO Subobjects
8.2.	New NAI Type Registry
8.3.	New SR-ERO Flag Registry
8.4.	PCEP-Error Object
8.5.	PCEP TLV Type Indicators
8.6.	PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators
8.7.	New Path Setup Type
8.8.	New Metric Type
8.9.	SR PCE Capability Flags
9.	References
9.1.	Normative References
9.2.	Informative References
Appendix A. Compatibility with Early Implementations	
Acknowledgements	
Contributors	
Authors' Addresses	

1. Introduction

Segment Routing (SR) leverages the source-routing paradigm. Using SR, a source node steers a packet through a path without relying on hop-by-hop signaling protocols such as LDP or RSVP-TE. Each path is specified as an ordered list of instructions called "segments". Each segment is an instruction to route the packet to a specific place in the network or to perform a function on the packet. A database of segments can be distributed through the network using a routing protocol (such as IS-IS or OSPF) or by any other means. Several types of segments are defined. A node segment uniquely identifies a specific node in the SR domain. Each router in the SR domain associates a node segment with an ECMP-aware shortest path to the node that it identifies. An adjacency segment represents a unidirectional adjacency. An adjacency segment is local to the node that advertises it. Both node segments and adjacency segments can be used for SR.

[RFC8402] describes the SR architecture. The corresponding IS-IS and OSPF extensions are specified in [RFC8667] and [RFC8665], respectively.

The SR architecture can be implemented using either an MPLS forwarding plane [RFC8660] or an IPv6 forwarding plane [IPv6-SRH]. The MPLS forwarding plane can be applied to SR without any change; in which case, an SR path corresponds to an MPLS Label Switching Path (LSP). This document is relevant to the MPLS forwarding plane only. In this document, "Node-SID" and "Adj-SID" denote the Node Segment Identifier and Adjacency Segment Identifier, respectively.

An SR path can be derived from an IGP Shortest Path Tree (SPT). Segment Routing Traffic-Engineering (SR-TE) paths may not follow an IGP SPT. Such paths may be chosen by a suitable network planning tool and provisioned on the ingress node of the SR-TE path.

[RFC5440] describes the Path Computation Element Communication Protocol (PCEP) for communication between a Path Computation Client (PCC) and a Path Computation Element (PCE) or between a pair of PCEs. A PCE computes paths for MPLS Traffic-Engineering (MPLS-TE) LSPs based on various constraints and optimization criteria. [RFC8231] specifies extensions to PCEP that allow a stateful PCE to compute and recommend network paths in compliance with [RFC4657]. It also defines objects and TLVs for MPLS-TE LSPs. Stateful PCEP extensions provide synchronization of LSP state between a PCC and a PCE or between a pair of PCEs, delegation of LSP control, reporting of LSP state from a PCC to a PCE, and control of the setup and path routing of an LSP from a PCE to a PCC. Stateful PCEP extensions are intended for an operational model in which LSPs are configured on the PCC, and control over them is delegated to the PCE.

A mechanism to dynamically initiate LSPs on a PCC based on the requests from a stateful PCE or a controller using stateful PCE is specified in [RFC8281]. This mechanism is useful in Software-Defined Networking (SDN) applications, such as on-demand engineering or bandwidth calendaring [RFC8413].

It is possible to use a stateful PCE for computing one or more SR-TE paths, taking into account various constraints and objective functions. Once a path is chosen, the stateful PCE can initiate an SR-TE path on a PCC using the PCEP extensions specified in [RFC8281] and the SR-specific PCEP extensions specified in this document. Additionally, using procedures described in this document, a PCC can request an SR path from either a stateful or a stateless PCE.

This specification relies on the procedures specified in [RFC8408] to exchange the Segment Routing capability and to specify that the path setup type of an LSP is Segment Routing. This specification also updates [RFC8408] to clarify the use of sub-TLVs in the PATH-SETUP-TYPE-CAPABILITY TLV. See Section 4.1.1 for details.

This specification provides a mechanism for a network controller (acting as a PCE) to instantiate candidate paths for an SR Policy onto a head-end node (acting as a PCC) using PCEP. For more information on the SR Policy Architecture, see [SR-POLICY].

2. Terminology

The following terminology is used in this document:

ERO: Explicit Route Object

IGP: Interior Gateway Protocol

IS-IS: Intermediate System to Intermediate System

LSR: Label Switching Router

MSD: Base MPLS Imposition Maximum SID Depth, as defined in

[RFC8491]

NAI: Node or Adjacency Identifier

OSPF: Open Shortest Path First

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Communication Protocol

RRO: Record Route Object

SID: Segment Identifier

SR: Segment Routing

SR-DB: Segment Routing Database: the collection of SRGBs, SRLBs, and SIDs and the objects they map to, advertised by a link-state IGP

SR-TE: Segment Routing Traffic Engineering

SRGB: Segment Routing Global Block

SRLB: Segment Routing Local Block

2.1. 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.

3. Overview of PCEP Operation in SR Networks

In an SR network, the ingress node of an SR path prepends an SR header to all outgoing packets. The SR header consists of a list of SIDs (or MPLS labels in the context of this document). The header has all necessary information so that, in combination with the information distributed by the IGP, the packets can be guided from the ingress node to the egress node of the path; hence, there is no need for any signaling protocol.

In PCEP messages, LSP route information is carried in the Explicit Route Object (ERO), which consists of a sequence of subobjects. SR-TE paths computed by a PCE can be represented in an ERO in one of the following forms:

- * An ordered set of IP addresses representing network nodes/links.
- * An ordered set of SIDs, with or without the corresponding IP addresses.
- * An ordered set of MPLS labels, with or without corresponding IP addresses.

The PCC converts these into an MPLS label stack and next hop, as described in Section 5.2.2.

This document defines a new ERO subobject denoted by "SR-ERO subobject" that is capable of carrying a SID as well as the identity of the node/adjacency represented by the SID. SR-capable PCEP speakers should be able to generate and/or process such an ERO

subobject. An ERO containing SR-ERO subobjects can be included in the PCEP Path Computation Reply (PCRep) message defined in [RFC5440], the Path Computation LSP Initiate Request (PCInitiate) message defined in [RFC8281], and the Path Computation Update Request (PCUpd) and Path Computation State Report (PCRpt) messages for LSPs defined in [RFC8231].

When a PCEP session between a PCC and a PCE is established, both PCEP speakers exchange their capabilities to indicate their ability to support SR-specific functionality.

A PCE can update an LSP that is initially established via RSVP-TE signaling to use an SR-TE path by sending a PCUpd to the PCC that delegated the LSP to it [RFC8231]. A PCC can update an undelegated LSP that is initially established via RSVP-TE signaling to use an SR-TE path as follows. First, it requests an SR-TE path from a PCE by sending a Path Computation Request (PCReq) message. If it receives a suitable path, it establishes the path in the data plane and then tears down the original RSVP-TE path. If the PCE is stateful, then the PCC sends PCRpt messages indicating that the new path is set up and the old path is torn down, per [RFC8231].

Similarly, a PCE or PCC can update an LSP initially created with an SR-TE path to use RSVP-TE signaling, if necessary. This capability is useful for rolling back a change when a network is migrated from RSVP-TE to SR-TE technology.

A PCC MAY include a Record Route Object (RRO) containing the recorded LSP in PCReq and PCRpt messages as specified in [RFC5440] and [RFC8231], respectively. This document defines a new RRO subobject for SR networks. The methods used by a PCC to record the SR-TE LSP are outside the scope of this document.

In summary, this document:

- * Defines a new ERO subobject, a new RRO subobject, and new PCEP error codes.
- * Specifies how two PCEP speakers can establish a PCEP session that can carry information about SR-TE paths.
- * Specifies processing rules for the ERO subobject.
- * Defines a new path setup type to be used in the PATH-SETUP-TYPE and PATH-SETUP-TYPE-CAPABILITY TLVs [RFC8408].
- * Defines a new sub-TLV for the PATH-SETUP-TYPE-CAPABILITY TLV.

The extensions specified in this document complement the existing PCEP specifications to support SR-TE paths. As such, the PCEP messages (e.g., PCReq, PCRep, PCRpt, PCUpd, PCInitiate, etc.) are formatted according to [RFC5440], [RFC8231], [RFC8281], and any other applicable PCEP specifications.

4. Object Formats

4.1. The OPEN Object

4.1.1. The Path Setup Type Capability TLV

[RFC8408] defines the PATH-SETUP-TYPE-CAPABILITY TLV for use in the OPEN object. The PATH-SETUP-TYPE-CAPABILITY TLV contains an optional list of sub-TLVs, which are intended to convey parameters that are associated with the path setup types supported by a PCEP speaker.

This specification updates [RFC8408] as follows. It creates a new

registry that defines the valid type indicators of the sub-TLVs of the PATH-SETUP-TYPE-CAPABILITY TLV (see Section 8.6). A PCEP speaker MUST NOT include a sub-TLV in the PATH-SETUP-TYPE-CAPABILITY TLV unless it appears in this registry. If a PCEP speaker receives a sub-TLV whose type indicator does not match one of those from the registry or is not recognized by the speaker, then the speaker MUST ignore the sub-TLV.

4.1.2. The SR PCE Capability Sub-TLV

This document defines a new Path Setup Type (PST) for SR, as follows:

PST = 1: Traffic-engineering path is set up using Segment Routing.

A PCEP speaker SHOULD indicate its support of the function described in this document by sending a PATH-SETUP-TYPE-CAPABILITY TLV in the OPEN object with this new PST included in the PST list.

This document also defines the SR-PCE-CAPABILITY sub-TLV. PCEP speakers use this sub-TLV to exchange information about their SR capability. If a PCEP speaker includes PST=1 in the PST list of the PATH-SETUP-TYPE-CAPABILITY TLV, then it MUST also include the SR-PCE-CAPABILITY sub-TLV inside the PATH-SETUP-TYPE-CAPABILITY TLV.

The format of the SR-PCE-CAPABILITY sub-TLV is shown in the following figure:




Figure 1: SR-PCE-CAPABILITY Sub-TLV Format

The codepoint for the TLV type is 26. The TLV length is 4 octets.

The 32-bit value is formatted as follows.

Reserved: MUST be set to zero by the sender and MUST be ignored by the receiver.

Flags: This document defines the following flag bits. The other bits MUST be set to zero by the sender and MUST be ignored by the receiver.

N: A PCC sets this flag bit to 1 to indicate that it is capable of resolving a Node or Adjacency Identifier (NAI) to a SID.

X: A PCC sets this flag bit to 1 to indicate that it does not impose any limit on the MSD.

Maximum SID Depth (MSD): specifies the maximum number of SIDs (MPLS label stack depth in the context of this document) that a PCC is capable of imposing on a packet. Section 5.1 explains the relationship between this field and the X-Flag.

4.2. The RP/SRP Object

To set up an SR-TE LSP using SR, the Request Parameter (RP) or Stateful PCE Request Parameter (SRP) object MUST include the PATH-SETUP-TYPE TLV, specified in [RFC8408], with the PST set to 1 (and

path setup using SR-TE).

The LSP-IDENTIFIERS TLV MAY be present for the above PST type.

4.3. ERO

An SR-TE path consists of one or more SIDs where each SID MAY be associated with the identifier that represents the node or adjacency corresponding to the SID. This identifier is referred to as the NAI. As described later, an NAI can be represented in various formats (e.g., IPv4 address, IPv6 address, etc). Furthermore, an NAI is used for troubleshooting purposes and, if necessary, to derive a SID value as described below.

The ERO specified in [RFC5440] is used to carry SR-TE path information. In order to carry a SID and/or NAI, this document defines a new ERO subobject referred to as the "SR-ERO subobject", whose format is specified in the following section. An ERO carrying an SR-TE path consists of one or more ERO subobjects, and it **MUST** carry only SR-ERO subobjects. Note that an SR-ERO subobject does not need to have both the SID and NAI. However, at least one of them **MUST** be present.

When building the MPLS label stack from ERO, a PCC MUST assume that SR-ERO subobjects are organized as a last-in-first-out stack. The first subobject relative to the beginning of ERO contains the information about the topmost label. The last subobject contains information about the bottommost label.

4.3.1. SR-ERO Subobject

An SR-ERO subobject is formatted as shown in the following diagram.

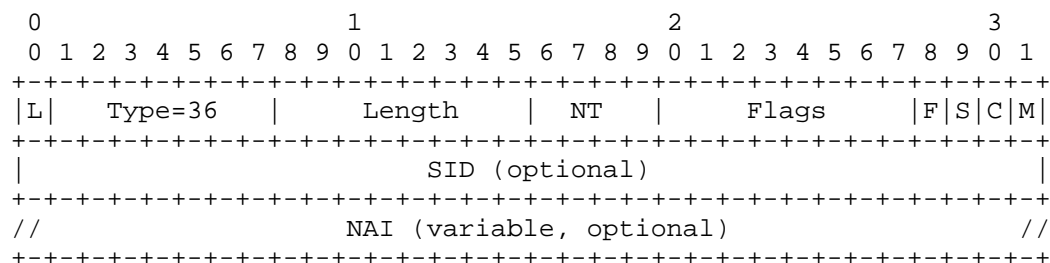


Figure 2: SR-ERO Subobject Format

The fields in the SR-ERO subobject are as follows:

The L-Flag: Indicates whether the subobject represents a loose hop in the LSP [RFC3209]. If this flag is set to zero, a PCC MUST NOT overwrite the SID value present in the SR-ERO subobject. Otherwise, a PCC MAY expand or replace one or more SID values in the received SR-ERO based on its local policy.

Type: Set to 36.

Length: Contains the total length of the subobject in octets. The Length MUST be at least 8 and MUST be a multiple of 4. An SR-ERO subobject MUST contain at least one SID or NAI. The flags described below indicate whether the SID or NAI fields are absent.

NAI Type (NT): Indicates the type and format of the NAI contained in the object body, if any is present. If the F bit is set to zero (see below), then the NT field has no meaning and MUST be ignored by the receiver. This document describes the following NT values:

NT=0 The NAI is absent.

NT=1 The NAI is an IPv4 node ID.

NT=2 The NAI is an IPv6 node ID.

NT=3 The NAI is an IPv4 adjacency.

NT=4 The NAI is an IPv6 adjacency with global IPv6 addresses.

NT=5 The NAI is an unnumbered adjacency with IPv4 node IDs.

NT=6 The NAI is an IPv6 adjacency with link-local IPv6 addresses.

Flags: Used to carry additional information pertaining to the SID.
This document defines the following flag bits. The other bits
MUST be set to zero by the sender and MUST be ignored by the
receiver.

M: If this bit is set to 1, the SID value represents an MPLS
label stack entry as specified in [RFC3032]. Otherwise, the
SID value is an administratively configured value that
represents an index into an MPLS label space (either SRGB or
SRLB) per [RFC8402].

C: If the M bit and the C bit are both set to 1, then the TC, S,
and TTL fields in the MPLS label stack entry are specified by
the PCE. However, a PCC MAY choose to override these values
according to its local policy and MPLS forwarding rules. If
the M bit is set to 1 but the C bit is set to zero, then the
TC, S, and TTL fields MUST be ignored by the PCC. The PCC
MUST set these fields according to its local policy and MPLS
forwarding rules. If the M bit is set to zero, then the C
bit MUST be set to zero.

S: When this bit is set to 1, the SID value in the subobject
body is absent. In this case, the PCC is responsible for
choosing the SID value, e.g., by looking it up in the SR-DB
using the NAI that, in this case, MUST be present in the
subobject. If the S bit is set to 1, then the M and C bits
MUST be set to zero.

F: When this bit is set to 1, the NAI value in the subobject
body is absent. The F bit MUST be set to 1 if NT=0;
otherwise, it MUST be set to zero. The S and F bits MUST NOT
both be set to 1.

SID: The Segment Identifier. Depending on the M bit, it contains
either:

- * A 4-octet index defining the offset into an MPLS label space
per [RFC8402] or
- * A 4-octet MPLS label stack entry, where the 20 most significant
bits encode the label value per [RFC3032].

NAI: The NAI associated with the SID. The NAI's format depends on
the value in the NT field and is described in the following
section.

At least one SID and NAI MUST be included in the SR-ERO subobject,
and both MAY be included.

4.3.2. NAI Associated with SID

This document defines the following NAIs:


```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Local IPv6 address (16 octets)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Local Interface ID                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Remote IPv6 address (16 octets)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Remote Interface ID                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 6: NAI for IPv6 Link-Local Adjacency

4.4. RRO

A PCC reports an SR-TE LSP to a PCE by sending a PCRpt message, per [RFC8231]. The RRO on this message represents the SID list that was applied by the PCC, that is, the actual path taken by the LSP. The procedures of [RFC8231] with respect to the RRO apply equally to this specification without change.

An RRO contains one or more subobjects called "SR-RRO subobjects", whose format is shown below:

```

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

```

Figure 7: SR-RRO Subobject Format

The format of the SR-RRO subobject is the same as that of the SR-ERO subobject, but without the L-Flag.

A PCC MUST order the SR-RRO subobjects such that the first subobject relative to the beginning of the RRO identifies the first segment visited by the SR-TE LSP, and the last subobject identifies the final segment of the SR-TE LSP, that is, its endpoint.

4.5. METRIC Object

A PCC MAY request that PCE optimizes an individual path computation request to minimize the SID depth of the computed path by using the METRIC object defined in [RFC5440]. This document defines a new type for the METRIC object to be used for this purpose, as follows:

T = 11: Maximum SID Depth of the requested path.

If the PCC includes a METRIC object of this type on a path computation request, then the PCE minimizes the SID depth of the computed path. If the B (bound) bit is set to 1 in the METRIC object, then the PCE MUST NOT return a path whose SID depth exceeds the given metric value. If the PCC did not set the X-Flag in its SR-PCE-CAPABILITY TLV, then it MUST set the B bit to 1. If the PCC set the X-Flag in its SR-PCE-CAPABILITY TLV, then it MAY set the B bit to 1 or zero.

If a PCEP session is established with a non-zero default MSD value, then the PCC MUST NOT send an MSD METRIC object with an MSD greater than the session's default MSD. If the PCE receives a path

computation request with an MSD METRIC object on such a session that is greater than the session's default MSD, then it MUST consider the request invalid and send a PCEP Error (PCErr) with Error-Type = 10 ("Reception of an invalid object") and Error-value = 9 ("MSD exceeds the default for the PCEP session").

5. Procedures

5.1. Exchanging the SR PCE Capability

A PCC indicates that it is capable of supporting the head-end functions for SR-TE LSP by including the SR-PCE-CAPABILITY sub-TLV in the Open message that it sends to a PCE. A PCE indicates that it is capable of computing SR-TE paths by including the SR-PCE-CAPABILITY sub-TLV in the Open message that it sends to a PCC.

If a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV with a PST list containing PST=1, and supports that path setup type, then it checks for the presence of the SR-PCE-CAPABILITY sub-TLV. If that sub-TLV is absent, then the PCEP speaker MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 12 ("Missing PCE-SR-CAPABILITY sub-TLV") and MUST then close the PCEP session. If a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV with a SR-PCE-CAPABILITY sub-TLV, but the PST list does not contain PST=1, then the PCEP speaker MUST ignore the SR-PCE-CAPABILITY sub-TLV.

If a PCC sets the N-Flag to 1, then the PCE MAY send an SR-ERO subobject containing an NAI and no SID (see Section 5.2). Otherwise, the PCE MUST NOT send an SR-ERO subobject containing an NAI and no SID.

The number of SIDs that can be imposed on a packet depends on the PCC's data-plane capability. If a PCC sets the X-Flag to 1, then the MSD is not used and MUST be set to zero. If a PCE receives an SR-PCE-CAPABILITY sub-TLV with the X-Flag set to 1, then it MUST ignore the MSD field and assume that the sender can impose a SID stack of any depth. If a PCC sets the X-Flag to zero, then it sets the MSD field to the maximum number of SIDs that it can impose on a packet. In this case, the PCC MUST set the MSD to a number greater than zero. If a PCE receives an SR-PCE-CAPABILITY sub-TLV with the X-Flag and MSD both set to zero, then it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 21 ("Maximum SID depth must be non-zero") and MUST then close the PCEP session.

Note that the MSD value exchanged via the SR-PCE-CAPABILITY sub-TLV indicates the SID/label imposition limit for the PCC node. It is anticipated that, in many deployments, the PCCs will have network interfaces that are homogeneous with respect to MSD (that is, each interface has the same MSD). In such cases, having a per-node MSD on the PCEP session is sufficient; the PCE SHOULD interpret this to mean that all network interfaces on the PCC have the given MSD. However, the PCE MAY also learn a per-node MSD and a per-interface MSD from the routing protocols, as specified in [RFC8491], [RFC8476], and [MSD-BGP]. If the PCE learns the per-node MSD of a PCC from a routing protocol, then it MUST ignore the per-node MSD value in the SR-PCE-CAPABILITY sub-TLV and use the per-node MSD learned from the routing protocol instead. If the PCE learns the MSD of a network interface on a PCC from a routing protocol, then it MUST use the per-interface MSD instead of the MSD value in the SR-PCE-CAPABILITY sub-TLV when it computes a path that uses that interface.

Once an SR-capable PCEP session is established with a non-zero MSD value, the corresponding PCE MUST NOT send SR-TE paths with a number of SIDs exceeding that MSD value. If a PCC needs to modify the MSD

value, it MUST close the PCEP session and re-establish it with the new MSD value. If a PCEP session is established with a non-zero MSD value, and the PCC receives an SR-TE path containing more SIDs than specified in the MSD value, the PCC MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 3 ("Unsupported number of SR-ERO subobjects"). If a PCEP session is established with an MSD value of zero, then the PCC MAY specify an MSD for each path computation request that it sends to the PCE, by including a "maximum SID depth" METRIC object on the request, as defined in Section 4.5.

The N-Flag, X-Flag, and MSD value inside the SR-PCE-CAPABILITY sub-TLV are meaningful only in the Open message sent from a PCC to a PCE. As such, a PCE MUST set the N-Flag to zero, X-Flag to 1, and MSD value to zero in an outbound message to a PCC. Similarly, a PCC MUST ignore any MSD value received from a PCE. If a PCE receives multiple SR-PCE-CAPABILITY sub-TLVs in an Open message, it processes only the first sub-TLV received.

5.2. ERO Processing

5.2.1. SR-ERO Validation

If a PCC does not support the SR PCE Capability and thus cannot recognize the SR-ERO or SR-RRO subobjects, it will respond according to the rules for a malformed object per [RFC5440].

On receiving an SR-ERO, a PCC MUST validate that the Length field, S bit, F bit, and NT field are consistent, as follows.

- * If NT=0, the F bit MUST be 1, the S bit MUST be zero, and the Length MUST be 8.
- * If NT=1, the F bit MUST be zero. If the S bit is 1, the Length MUST be 8; otherwise, the Length MUST be 12.
- * If NT=2, the F bit MUST be zero. If the S bit is 1, the Length MUST be 20; otherwise, the Length MUST be 24.
- * If NT=3, the F bit MUST be zero. If the S bit is 1, the Length MUST be 12; otherwise, the Length MUST be 16.
- * If NT=4, the F bit MUST be zero. If the S bit is 1, the Length MUST be 36; otherwise, the Length MUST be 40.
- * If NT=5, the F bit MUST be zero. If the S bit is 1, the Length MUST be 20; otherwise, the Length MUST be 24.
- * If NT=6, the F bit MUST be zero. If the S bit is 1, the Length MUST be 44; otherwise, the Length MUST be 48.

If a PCC finds that the NT field, Length field, S bit, and F bit are not consistent, it MUST consider the entire ERO invalid and MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 11 ("Malformed object").

If a PCC does not recognize or support the value in the NT field, it MUST consider the entire ERO invalid and MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 13 ("Unsupported NAI Type in the SR-ERO/SR-RRO subobject").

If a PCC receives an SR-ERO subobject in which the S and F bits are both set to 1 (that is, both the SID and NAI are absent), it MUST consider the entire ERO invalid and send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 6 ("Both SID and NAI are absent in the SR-ERO subobject").

If a PCC receives an SR-ERO subobject in which the S bit is set to 1 and the F bit is set to zero (that is, the SID is absent and the NAI is present), but the PCC does not support NAI resolution, it MUST consider the entire ERO invalid and send a PCErr message with Error-Type = 4 ("Not supported object") and Error-value = 4 ("Unsupported parameter").

If a PCC receives an SR-ERO subobject in which the S bit is set to 1 and either (or both) the M bit or the C bit is set to 1, it MUST consider the entire ERO invalid and send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 11 ("Malformed object").

If a PCC receives an SR-ERO subobject in which the S bit is set to zero and the M bit is set to 1, then the subobject contains an MPLS label. The PCC MAY choose not to accept a label provided by the PCE, based on its local policy. The PCC MUST NOT accept MPLS label value 3 (Implicit NULL), but it MAY accept other special-purpose MPLS label values. If the PCC decides not to accept an MPLS label value, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 2 ("Bad label value").

If both the M and C bits of an SR-ERO subobject are set to 1, and if a PCC finds an erroneous setting in one or more of the TC, S, and TTL fields, it MAY overwrite those fields with values chosen according to its own policy. If the PCC does not overwrite them, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 4 ("Bad label format").

If the M bit of an SR-ERO subobject is set to zero but the C bit is set to 1, then the PCC MUST consider the entire ERO invalid and MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 11 ("Malformed object").

If a PCC receives an SR-ERO subobject in which the S bit is set to zero and the M bit is set to zero, then the subobject contains a SID index value. If the SID is an Adj-SID, then the L-Flag MUST NOT be set. If the L-Flag is set for an Adj-SID, then the PCC MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 11 ("Malformed object").

If a PCC detects that the subobjects of an ERO are a mixture of SR-ERO subobjects and subobjects of other types, then it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 5 ("ERO mixes SR-ERO subobjects with other subobject types").

The SR-ERO subobjects can be classified according to whether they contain a SID representing an MPLS label value or an index value, or no SID. If a PCC detects that the SR-ERO subobjects are a mixture of more than one of these types, then it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 20 ("Inconsistent SIDs in SR-ERO/SR-RRO subobjects").

If an ERO specifies a new SR-TE path for an existing LSP and the PCC determines that the ERO contains SR-ERO subobjects that are not valid, then the PCC MUST NOT update the LSP.

5.2.2. Interpreting the SR-ERO

The SR-ERO contains a sequence of subobjects. Each SR-ERO subobject in the sequence identifies a segment that the traffic will be directed to, in the order given. That is, the first subobject identifies the first segment the traffic will be directed to, the second subobject represents the second segment, and so on.

The PCC interprets the SR-ERO by converting it to an MPLS label stack plus a next hop. The PCC sends packets along the segment-routed path by prepending the MPLS label stack onto the packets and sending the resulting, modified packet to the next hop.

The PCC uses a different procedure to do this conversion, depending on the information that the PCE has provided in the subobjects.

- * If the subobjects contain SID index values, then the PCC converts them into the corresponding MPLS labels by following the procedure defined in [RFC8660].
- * If the subobjects contain NAIs only, the PCC first converts each NAI into a SID index value and then proceeds as above. To convert an NAI to a SID index, the PCC looks for a fully specified prefix or adjacency matching the fields in the NAI. If the PCC finds a matching prefix/adjacency, and the matching prefix/adjacency has a SID associated with it, then the PCC uses that SID. If the PCC cannot find a matching prefix/adjacency, or if the matching prefix/adjacency has no SID associated with it, the PCC behaves as specified in Section 5.2.2.1.
- * If the subobjects contain MPLS labels, then the PCC looks up the offset of the first subobject's label in its SRGB or SRLB. This gives the first SID. The PCC pushes the labels in any remaining subobjects onto the packet (with the final subobject specifying the bottom-of-stack label).

For all cases above, after the PCC has imposed the label stack on the packet, it sends the packet to the segment identified by the first SID.

5.2.2.1. Handling Errors During SR-ERO Conversion

There are several errors that can occur during the process of converting an SR-ERO sequence to an MPLS label stack and a next hop. The PCC deals with them as follows.

- * If the PCC cannot find a SID index in the SR-DB, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 14 ("Unknown SID").
- * If the PCC cannot find an NAI in the SR-DB, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 15 ("NAI cannot be resolved to a SID").
- * If the PCC needs to convert a SID into an MPLS label value but cannot find the corresponding router's SRGB in the SR-DB, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 16 ("Could not find SRGB").
- * If the PCC finds that a router's SRGB is not large enough for a SID index value, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 17 ("SID index exceeds SRGB size").
- * If the PCC needs to convert a SID into an MPLS label value but cannot find the corresponding router's SRLB in the SR-DB, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 18 ("Could not find SRLB").
- * If the PCC finds that a router's SRLB is not large enough for a SID index value, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 19 ("SID index exceeds SRLB size").

- * If the number of labels in the computed label stack exceeds the maximum number of SIDs that the PCC can impose on the packet, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 3 ("Unsupported number of SR-ERO subobjects").

If an ERO specifies a new SR-TE path for an existing LSP and the PCC encounters an error while processing the ERO, then the PCC MUST NOT update the LSP.

5.3. RRO Processing

The syntax-checking rules that apply to the SR-RRO subobject are identical to those of the SR-ERO subobject, except as noted below.

If a PCEP speaker receives an SR-RRO subobject in which both SID and NAI are absent, it MUST consider the entire RRO invalid and send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 7 ("Both SID and NAI are absent in the SR-RRO subobject").

If a PCE detects that the subobjects of an RRO are a mixture of SR-RRO subobjects and subobjects of other types, then it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 10 ("RRO mixes SR-RRO subobjects with other subobject types").

The SR-RRO subobjects can be classified according to whether they contain a SID representing an MPLS label value or an index value, or no SID. If a PCE detects that the SR-RRO subobjects are a mixture of more than one of these types, then it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 20 ("Inconsistent SIDs in SR-ERO / SR-RRO subobjects").

6. Management Considerations

This document adds a new path setup type to PCEP to allow LSPs to be set up using Segment Routing techniques. This path setup type may be used with PCEP alongside other path setup types, such as RSVP-TE, or it may be used exclusively.

6.1. Controlling the Path Setup Type

The following factors control which path setup type is used for a given LSP.

- * The available path setup types are constrained to those that are supported by, or enabled on, the PCEP speakers. The PATH-SETUP-TYPE-CAPABILITY TLV indicates which path setup types a PCEP speaker supports. To use Segment Routing as a path setup type, it is a prerequisite that the PCC and PCE both include PST=1 in the list of supported path setup types in this TLV and also include the SR-PCE-CAPABILITY sub-TLV.
- * When a PCE initiates an LSP, it proposes which path setup type to use by including it in the PATH-SETUP-TYPE TLV in the SRP object of the PCInitiate message. The PCE chooses the path setup type based on the capabilities of the network nodes on the path and on its local policy. The PCC MAY choose to accept the proposed path setup type or to reject the PCInitiate request, based on its local policy.
- * When a PCC requests a path for an LSP, it can nominate a preferred path setup type by including it in the PATH-SETUP-TYPE TLV in the RP object of the PCReq message. The PCE MAY choose to reply with

a path of the requested type, reply with a path of a different type, or reject the request, based on the capabilities of the network nodes on the path and on its local policy.

The operator can influence the path setup type as follows.

- * Implementations **MUST** allow the operator to enable and disable the Segment Routing path setup type on a PCEP-speaking device. Implementations **MAY** also allow the operator to enable and disable the RSVP-TE path setup type.
- * PCE implementations **MUST** allow the operator to specify that an LSP should be instantiated using Segment Routing or RSVP-TE as the proposed path setup type.
- * PCE implementations **MAY** allow the operator to configure a preference for the PCE to propose paths using Segment Routing or RSVP-TE in the absence of a specified path setup type.
- * PCC implementations **MUST** allow the operator to specify that a path requested for an LSP nominates Segment Routing or RSVP-TE as the path setup type.
- * PCC implementations **MAY** allow the operator to configure a preference for the PCC to nominate Segment Routing or RSVP-TE as the path setup type if none is specified for an LSP.
- * PCC implementations **SHOULD** allow the operator to configure a PCC to refuse to set up an LSP using an undesired path setup type.

6.2. Migrating a Network to Use PCEP Segment-Routed Paths

This section discusses the steps that the operator takes when migrating a network to enable PCEP to set up paths using Segment Routing as the path setup type.

- * The operator enables the Segment Routing PST on the PCE servers.
- * The operator enables the Segment Routing PST on the PCCs.
- * The operator resets each PCEP session. The PCEP sessions come back up with Segment Routing enabled.
- * If the operator detects a problem, they can roll the network back to its initial state by disabling the Segment Routing PST on the PCEP speakers and resetting the PCEP sessions.

Note that the data plane is unaffected if a PCEP session is reset. Any LSPs that were set up before the session reset will remain in place and will still be present after the session comes back up.

An implementation **SHOULD** allow the operator to manually trigger a PCEP session to be reset.

An implementation **MAY** automatically reset a PCEP session when an operator reconfigures the PCEP speaker's capabilities. However, note that if the capabilities at both ends of the PCEP session are not reconfigured simultaneously, then the session could be reset twice, which could lead to unnecessary network traffic. Therefore, such implementations **SHOULD** allow the operator to override this behavior and wait instead for a manual reset.

Once Segment Routing is enabled on a PCEP session, it can be used as the path setup type for future LSPs.

User traffic is not automatically migrated from existing LSPs onto

segment-routed LSPs just by enabling the Segment Routing PST in PCEP. The migration of user traffic from existing LSPs onto Segment Routing LSPs is beyond the scope of this document.

6.3. Verification of Network Operation

The operator needs the following information to verify that PCEP is operating correctly with respect to the Segment Routing path setup type.

- * An implementation SHOULD allow the operator to view whether the PCEP speaker sent the Segment Routing PST capability to its peer. If the PCEP speaker is a PCC, then the implementation SHOULD also allow the operator to view the values of the L-Flag and N-Flag that were sent and the value of the MSD field that was sent.
- * An implementation SHOULD allow the operator to view whether the peer sent the Segment Routing PST capability. If the peer is a PCC, then the implementation SHOULD also allow the operator to view the values of the L-Flag and N-Flag and MSD fields that the peer sent.
- * An implementation SHOULD allow the operator to view whether the Segment Routing PST is enabled on the PCEP session.
- * If one PCEP speaker advertises the Segment Routing PST capability, but the other does not, then the implementation SHOULD create a log to inform the operator of the capability mismatch.
- * An implementation SHOULD allow the operator to view the PST that was proposed, or requested, for an LSP and the PST that was actually used.
- * If a PCEP speaker decides to use a different PST to the one that was proposed, or requested, for an LSP, then the implementation SHOULD create a log to inform the operator that the expected PST has not been used. The log SHOULD give the reason for this choice (local policy, equipment capability, etc.).
- * If a PCEP speaker rejects a Segment Routing path, then it SHOULD create a log to inform the operator, giving the reason for the decision (local policy, MSD exceeded, etc.).

6.4. Relationship to Existing Management Models

The PCEP YANG module is defined in [PCE-PCEP-YANG]. In the future, this YANG module should be extended or augmented to provide the following additional information relating to Segment Routing:

- * The advertised PST capabilities and MSD per PCEP session.
- * The PST configured for, and used by, each LSP.

The PCEP MIB [RFC7420] could also be updated to include this information.

7. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8281], and [RFC8408] are applicable to this specification. No additional security measures are required.

Note that this specification enables a network controller to instantiate a path in the network without the use of a hop-by-hop signaling protocol (such as RSVP-TE). This creates an additional vulnerability if the security mechanisms of [RFC5440], [RFC8231], and

[RFC8281] are not used. If there is no integrity protection on the session, then an attacker could create a path that is not subjected to the further verification checks that would be performed by the signaling protocol.

Note that this specification adds the MSD field to the Open message (see Section 4.1.2), which discloses how many MPLS labels the sender can push onto packets that it forwards into the network. If the security mechanisms of [RFC8231] and [RFC8281] are not used with strong encryption, then an attacker could use this new field to gain intelligence about the capabilities of the edge devices in the network.

8. IANA Considerations

8.1. PCEP ERO and RRO Subobjects

This document defines a new subobject type for the PCEP ERO and a new subobject type for the PCEP RRO. The codepoints for subobject types of these objects are maintained in the "Resource Reservation Protocol (RSVP) Parameters" registry, under the EXPLICIT_ROUTE and ROUTE_RECORD objects, respectively.

Object	Subobject	Subobject Type
EXPLICIT_ROUTE	SR-ERO (PCEP specific)	36
ROUTE_RECORD	SR-RRO (PCEP specific)	36

Table 1

8.2. New NAI Type Registry

IANA has created a new sub-registry within the "Path Computation Element Protocol (PCEP) Numbers" registry called "PCEP SR-ERO NAI Types". The allocation policy for this new registry is by IETF Review [RFC8126]. The new registry contains the following values:

Value	Description	Reference
0	NAI is absent.	This document
1	NAI is an IPv4 node ID.	This document
2	NAI is an IPv6 node ID.	This document
3	NAI is an IPv4 adjacency.	This document
4	NAI is an IPv6 adjacency with global IPv6 addresses.	This document
5	NAI is an unnumbered adjacency with IPv4 node IDs.	This document
6	NAI is an IPv6 adjacency with link-local IPv6 addresses.	This document
7-15	Unassigned	

Table 2

8.3. New SR-ERO Flag Registry

IANA has created a new sub-registry, named "SR-ERO Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field of the SR-ERO subobject. New values are to be assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit number (counting from bit 0 as the most significant bit)
- * Capability description
- * Defining RFC

The following values are defined in this document:

Bit	Description	Reference
0-7	Unassigned	
8	NAI is absent (F)	This document
9	SID is absent (S)	This document
10	SID specifies TC, S, and TTL in addition to an MPLS label (C)	This document
11	SID specifies an MPLS label (M)	This document

Table 3

8.4. PCEP-Error Object

IANA has allocated the following codepoints in the "PCEP-ERROR Object Error Types and Values" registry for the following new Error-values:

Error-Type	Meaning	Error-value
10	Reception of an invalid object	
		2: Bad label value
		3: Unsupported number of SR-ERO subobjects
		4: Bad label format
		5: ERO mixes SR-ERO subobjects with other subobject types
		6: Both SID and NAI are absent in the SR-ERO subobject
		7: Both SID and NAI are absent in the SR-RRO subobject
		9: MSD exceeds the default for the PCEP session
		10: RRO mixes SR-RRO subobjects with other subobject types
		12: Missing PCE-SR-CAPABILITY

		sub-TLV	
		13: Unsupported NAI Type in the SR-ERO/SR-RRO subobject	
		14: Unknown SID	
		15: NAI cannot be resolved to a SID	
		16: Could not find SRGB	
		17: SID index exceeds SRGB size	
		18: Could not find SRLB	
		19: SID index exceeds SRLB size	
		20: Inconsistent SIDs in SR-ERO / SR-RRO subobjects	
		21: MSD must be non-zero	

Table 4

8.5. PCEP TLV Type Indicators

IANA has allocated the following codepoint in the "PCEP TLV Type Indicators" registry. Note that this TLV type indicator is deprecated but retained in the registry to ensure compatibility with early implementations of this specification. See Appendix A for details.

Value	Meaning	Reference
26	SR-PCE-CAPABILITY (deprecated)	This document

Table 5

8.6. PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators

IANA has created a new sub-registry, named "PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the type indicator space for sub-TLVs of the PATH-SETUP-TYPE-CAPABILITY TLV. New values are to be assigned by Standards Action [RFC8126]. The valid range of values in the registry is 0-65535. IANA has initialized the registry with the following values. All other values in the registry should be marked as "Unassigned".

Value	Meaning	Reference
0	Reserved	This document
26	SR-PCE-CAPABILITY	This document

Table 6

8.7. New Path Setup Type

A sub-registry within the "Path Computation Element Protocol (PCEP)

Numbers" registry called "PCEP Path Setup Types" was created in [RFC8408]. IANA has allocated a new codepoint within this registry, as follows:

Value	Description	Reference
1	Traffic-engineering path is set up using Segment Routing.	This document

Table 7

8.8. New Metric Type

IANA has allocated the following codepoint in the PCEP "METRIC Object T Field" registry:

Value	Description	Reference
11	Segment-ID (SID) Depth.	This document

Table 8

8.9. SR PCE Capability Flags

IANA has created a new sub-registry, named "SR Capability Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field of the SR-PCE-CAPABILITY TLV. New values are to be assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit number (counting from bit 0 as the most significant bit)
- * Capability description
- * Defining RFC

The following values are defined in this document:

Bit	Description	Reference
0-5	Unassigned	
6	Node or Adjacency Identifier (NAI) is supported (N)	This document
7	Unlimited Maximum SID Depth (X)	This document

Table 9

9. References

9.1. 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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <<https://www.rfc-editor.org/info/rfc8491>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

9.2. Informative References

- [IPv6-SRH] Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", Work in Progress, Internet-Draft, draft-ietf-6man-segment-routing-header-26, 22 October 2019, <<https://tools.ietf.org/html/draft-ietf-6man-segment-routing-header-26>>.
- [MSD-BGP] Tantsura, J., Chunduri, U., Talaulikar, K., Mirsky, G., and N. Triantafyllis, "Signaling MSD (Maximum SID Depth) using Border Gateway Protocol Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-segment-routing-msd-09, 15 October 2019, <<https://tools.ietf.org/html/draft-ietf-idr-bgp-ls-segment-routing-msd-09>>.
- [PCE-PCEP-YANG] Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element

Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-13, 31 October 2019, <<https://tools.ietf.org/html/draft-ietf-pce-pcep-yang-13>>.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC4657] Ash, J., Ed. and J.L. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <<https://www.rfc-editor.org/info/rfc4657>>.
- [RFC7420] Koushik, A., Stephan, E., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module", RFC 7420, DOI 10.17487/RFC7420, December 2014, <<https://www.rfc-editor.org/info/rfc7420>>.
- [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>>.
- [RFC8413] Zhuang, Y., Wu, Q., Chen, H., and A. Farrel, "Framework for Scheduled Use of Resources", RFC 8413, DOI 10.17487/RFC8413, July 2018, <<https://www.rfc-editor.org/info/rfc8413>>.
- [RFC8476] Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak, "Signaling Maximum SID Depth (MSD) Using OSPF", RFC 8476, DOI 10.17487/RFC8476, December 2018, <<https://www.rfc-editor.org/info/rfc8476>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [SR-POLICY] Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-05, 17 November 2019, <<https://tools.ietf.org/html/draft-ietf-spring-segment-routing-policy-05>>.

Appendix A. Compatibility with Early Implementations

An early implementation of this specification will send the SR-CAPABILITY-TLV as a top-level TLV in the OPEN object instead of sending the PATH-SETUP-TYPE-CAPABILITY TLV in the OPEN object. Implementations that wish to interoperate with such early implementations should also send the SR-CAPABILITY-TLV as a top-level TLV in their OPEN object and should interpret receiving this top-level TLV as though the sender had sent a PATH-SETUP-TYPE-CAPABILITY TLV with a PST list of (0, 1) (that is, both RSVP-TE and SR-TE PSTs

are supported) with the SR-CAPABILITY-TLV as a sub-TLV. If a PCEP speaker receives an OPEN object in which both the SR-CAPABILITY-TLV and PATH-SETUP-TYPE-CAPABILITY TLV appear as top-level TLVs, then it should ignore the top-level SR-CAPABILITY-TLV and process only the PATH-SETUP-TYPE-CAPABILITY TLV.

Acknowledgements

We thank Ina Minei, George Swallow, Marek Zavodsky, Dhruv Dhody, Ing-Wher Chen, and Tomas Janciga for the valuable comments.

Contributors

The following people contributed to this document:

- * Lakshmi Sharma
- * Jan Medved
- * Edward Crabbe
- * Robert Raszuk
- * Victor Lopez

Authors' Addresses

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata Ontario K2K 3E8
Canada

Email: msiva@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Pegasus Parc
Brabant 1831 De kleetlaan 6a
Belgium

Email: cfilsfil@cisco.com

Jeff Tantsura
Apstra, Inc.
333 Middlefield Rd #200
Menlo Park, CA 94025
United States of America

Email: jefftant.ietf@gmail.com

Wim Henderickx
Nokia
Copernicuslaan 50
95134 Antwerp 2018
Belgium

Email: wim.henderickx@nokia.com

Jon Hardwick
Metaswitch Networks
100 Church Street
Enfield
United Kingdom

Email: jonathan.hardwick@metaswitch.com

