

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 6, 2021

C. Li
Huawei Technologies
M. Negi
RtBrick Inc
M. Koldychev
Cisco Systems, Inc.
P. Kaladharan
RtBrick Inc
Y. Zhu
China Telecom
November 2, 2020

PCEP Extensions for Segment Routing leveraging the IPv6 data plane
draft-ietf-pce-segment-routing-ipv6-07

Abstract

The Source Packet Routing in Networking (SPRING) architecture describes how Segment Routing (SR) can be used to steer packets through an IPv6 or MPLS network using the source routing paradigm. 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 IGP. A Segment Routed Path can be derived from a variety of mechanisms, including an IGP Shortest Path Tree (SPT), explicit configuration, or a Path Computation Element (PCE).

Since SR can be applied to both MPLS and IPv6 forwarding plane, a PCE should be able to compute SR-Path for both MPLS and IPv6 forwarding plane. This document describes the extensions required for SR support for IPv6 data plane in Path Computation Element communication Protocol (PCEP). The PCEP extension and mechanism to support SR-MPLS is described in [RFC 8664](#). This document extends it to support SRv6 (SR over IPv6).

Requirements Language

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 6, 2021.

Copyright Notice

Copyright (c) 2020 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	3
2.	Terminology	5
3.	Overview of PCEP Operation in SRv6 Networks	5
3.1.	Operation Overview	6
3.2.	SRv6-Specific PCEP Message Extensions	6
4.	Object Formats	7
4.1.	The OPEN Object	7
4.1.1.	The SRv6 PCE Capability sub-TLV	7
4.2.	The RP/SRP Object	8
4.3.	ERO	8
4.3.1.	SRv6-ERO Subobject	9
4.3.1.1.	SID Structure	11
4.4.	RRO	12
4.4.1.	SRv6-RRO Subobject	12

5.	Procedures	13
5.1.	Exchanging the SRv6 Capability	13
5.2.	ERO Processing	15
5.2.1.	SRv6 ERO Validation	15
5.2.2.	Interpreting the SRv6-ERO	16
5.3.	RRO Processing	16
6.	Security Considerations	16
7.	IANA Considerations	17
7.1.	PCEP ERO and RRO subobjects	17
7.2.	New SRv6-ERO Flag Registry	17
7.3.	PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators	18
7.4.	SRv6 PCE Capability Flags	18
7.5.	New Path Setup Type	18
7.6.	ERROR Objects	19
8.	Acknowledgements	19
9.	References	19
9.1.	Normative References	19
9.2.	Informative References	21
Appendix A.	Contributor	23
	Authors' Addresses	23

[1.](#) Introduction

As per [\[RFC8402\]](#), with Segment Routing (SR), a node steers a packet through an ordered list of instructions, called segments. A segment can represent any instruction, topological or service-based. A segment can have a semantic local to an SR node or global within an SR domain. SR allows to enforce a flow through any path and service chain while maintaining per-flow state only at the ingress node of the SR domain. Segments can be derived from different components: IGP, BGP, Services, Contexts, Locator, etc. The list of segment forming the path is called the Segment List and is encoded in the packet header. Segment Routing can be applied to the IPv6 architecture with the Segment Routing Header (SRH) [\[RFC8754\]](#). A segment is encoded as an IPv6 address. An ordered list of segments is encoded as an ordered list of IPv6 addresses in the routing header. The active segment is indicated by the Destination Address of the packet. Upon completion of a segment, a pointer in the new routing header is incremented and indicates the next segment.

Segment Routing use cases are described in [\[RFC7855\]](#) and [\[RFC8354\]](#). Segment Routing protocol extensions are defined in [\[RFC8667\]](#), and [\[RFC8666\]](#).

As per [\[RFC8754\]](#), an SRv6 Segment is a 128-bit value. "SRv6 SID" or simply "SID" are often used as a shorter reference for "SRv6 Segment". Further details are in an illustration provided in [\[I-D.ietf-spring-srv6-network-programming\]](#).

The SR architecture can be applied to the MPLS forwarding plane without any change, in which case an SR path corresponds to an MPLS Label Switching Path (LSP). The SR is applied to IPV6 forwarding plane using SRH. A SR path can be derived from an IGP Shortest Path Tree (SPT), but SR-TE paths may not follow IGP SPT. Such paths may be chosen by a suitable network planning tool, or a PCE and provisioned on the ingress node.

[RFC5440] describes 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 or a PCC operating as a PCE (in hierarchical PCE environment) computes paths for MPLS Traffic Engineering LSPs (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] and 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, controlling 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]. As per [RFC8664], 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 PCEP extensions specified in [RFC8281] using the SR specific PCEP extensions specified in [RFC8664]. [RFC8664] specifies PCEP extensions for supporting a SR-TE LSP for MPLS data plane. This document extends [RFC8664] to support SR for IPV6 data plane. Additionally, using procedures described in this document, a PCC can request an SRv6 path from either stateful or a stateless PCE. This specification relies on the PATH-SETUP-TYPE TLV and procedures specified in [RFC8408].

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 [I-D.ietf-spring-segment-routing-policy].

2. Terminology

This document uses the following terms defined in [[RFC5440](#)]: PCC, PCE, PCEP Peer.

This document uses the following terms defined in [[RFC8051](#)]: Stateful PCE, Delegation.

The message formats in this document are specified using Routing Backus-Naur Format (RBNF) encoding as specified in [[RFC5511](#)].

NAI: Node or Adjacency Identifier.

PCC: Path Computation Client.

PCE: Path Computation Element.

PCEP: Path Computation Element Protocol.

SR: Segment Routing.

SID: Segment Identifier.

SRv6: Segment Routing for IPv6 forwarding plane.

SRH: IPv6 Segment Routing Header.

SR Path: IPv6 Segment List (List of IPv6 SIDs representing a path in IPv6 SR domain)

Further, note that the term LSP used in the PCEP specifications, would be equivalent to a SRv6 Path (represented as a list of SRv6 segments) in the context of supporting SRv6 in PCEP.

3. Overview of PCEP Operation in SRv6 Networks

Basic operations for PCEP speakers is as per [[RFC8664](#)]. SRv6 Paths computed by a PCE can be represented as an ordered list of SRv6 segments of 128-bit value. "SRv6 SID" or simply "SID" are often used as a shorter reference for "SRv6 Segment" in this document.

[[RFC8664](#)] defined a new Explicit Route Object (ERO) subobject denoted by "SR-ERO subobject" 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 ERO subobject. An ERO containing SR-ERO subobjects can be included in the PCEP Path Computation Reply (PCRep) message defined in [[RFC5440](#)], the PCEP LSP Initiate Request message (PCInitiate) defined in

[[RFC8281](#)], as well as in the PCEP LSP Update Request (PCUpd) and PCEP LSP State Report (PCRpt) messages defined in [\[RFC8231\]](#).

This document define new subobjects "SRv6-ER0" and "SRv6-RR0" in ER0 and RR0 respectively to carry SRv6 SID (IPv6 Address). SRv6-capable PCEP speakers MUST be able to generate and/or process this.

When a PCEP session between a PCC and a PCE is established, both PCEP speakers exchange their capabilities to indicate their ability to support SRv6 specific functionality.

In summary, this document:

- o Defines a new PCEP capability for SRv6.
- o Defines a new subobject SRv6-ER0 in ER0.
- o Defines a new subobject SRv6-RR0 in RR0.
- o Defines a new path setup type carried in the PATH-SETUP-TYPE TLV and the PATH-SETUP-TYPE-CAPABILITY TLV.

[3.1.](#) Operation Overview

In SR networks, an ingress node of an SR path appends all outgoing packets with an SR header consisting of a list of SIDs (IPv6 Prefix in case of SRv6). The header has all necessary information to guide the packets from the ingress node to the egress node of the path, and hence there is no need for any signaling protocol.

For IPv6 in control plane with MPLS data-plane, mechanism remains same as [\[RFC8664\]](#)

This document describes extensions to SR path for IPv6 data plane. SRv6 Path (i.e. ER0) consists of an ordered set of SRv6 SIDs(see details in Figure 2).

A PCC or PCE indicates its ability to support SRv6 during the PCEP session Initialization Phase via a new SRv6-PCE-CAPABILITY sub-TLV (see details in [Section 4.1.1](#)).

[3.2.](#) SRv6-Specific PCEP Message Extensions

As defined in [\[RFC5440\]](#), a PCEP message consists of a common header followed by a variable length body made up of mandatory and/or optional objects. This document does not require any changes in the format of PCReq and PCRep messages specified in [\[RFC5440\]](#), PCInitiate message specified in [\[RFC8281\]](#), and PCRpt and PCUpd messages

specified in [RFC8231]. However, PCEP messages pertaining to SRv6 MUST include PATH-SETUP-TYPE TLV in the RP or SRP object to clearly identify that SRv6 is intended.

4. Object Formats

4.1. The OPEN Object

4.1.1. The SRv6 PCE Capability sub-TLV

This document defines a new Path Setup Type (PST) [RFC8408] for SRv6, as follows:

- o PST = TBD2: Path is setup using SRv6.

A PCEP speaker MUST 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 SRv6-PCE-CAPABILITY sub-TLV. PCEP speakers use this sub-TLV to exchange information about their SRv6 capability. If a PCEP speaker includes PST=TBD2 in the PST List of the PATH-SETUP-TYPE-CAPABILITY TLV then it MUST also include the SRv6-PCE-CAPABILITY sub-TLV inside the PATH-SETUP-TYPE-CAPABILITY TLV.

The format of the SRv6-PCE-CAPABILITY sub-TLV is shown in the following figure:




Figure 1: SRv6-PCE-CAPABILITY sub-TLV format

The code point for the TLV type (TBD1) is to be defined by IANA. The TLV length is variable.

The value comprises of -

Reserved: 2 octet, this field MUST be set to 0 on transmission, and ignored on receipt.

Flags: 2 octet, two bits are currently assigned in this document.

N bit: A PCC sets this flag bit to 1 to indicate that it is capable of resolving a Node or Adjacency Identifier (NAI) to a SRv6-SID.

X bit: A PCC sets this bit to 1 to indicate that it does not impose any limit on MSD (irrespective of the MSD-Type).

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

A pair of (MSD-Type, MSD-Value): Where MSD-Type (1 octet) is as per the IGP MSD Type registry created by [[RFC8491](#)] and populated with SRv6 MSD types as per [[I-D.ietf-lsr-isis-srv6-extensions](#)]; MSD-Value (1 octet) is as per [[RFC8491](#)].

This sub-TLV format is compliant with the PCEP TLV format defined in [[RFC5440](#)]. That is, the sub-TLV is composed of 2 octets for the type, 2 octets specifying the length, and a Value field. The Type field when set to TBD1 identifies the SRv6-PCE-CAPABILITY sub-TLV and the presence of the sub-TLV indicates the support for the SRv6 paths in PCEP. The Length field defines the length of the value portion in octets. The TLV is padded to 4-octet alignment, and padding is not included in the Length field. The number of (MSD-Type,MSD-Value) pairs can be determined from the Length field of the TLV.

[4.2.](#) The RP/SRP Object

In order to indicate the SRv6 path, RP or SRP object MUST include the PATH-SETUP-TYPE TLV specified in [[RFC8408](#)]. This document defines a new Path Setup Type (PST=TBD2) for SRv6.

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

[4.3.](#) ERO

In order to support SRv6, new subobject "SRv6-ERO" is defined in ERO.

4.3.1. SRv6-ERO Subobject

An SRv6-ERO subobject is formatted as shown in the following figure.

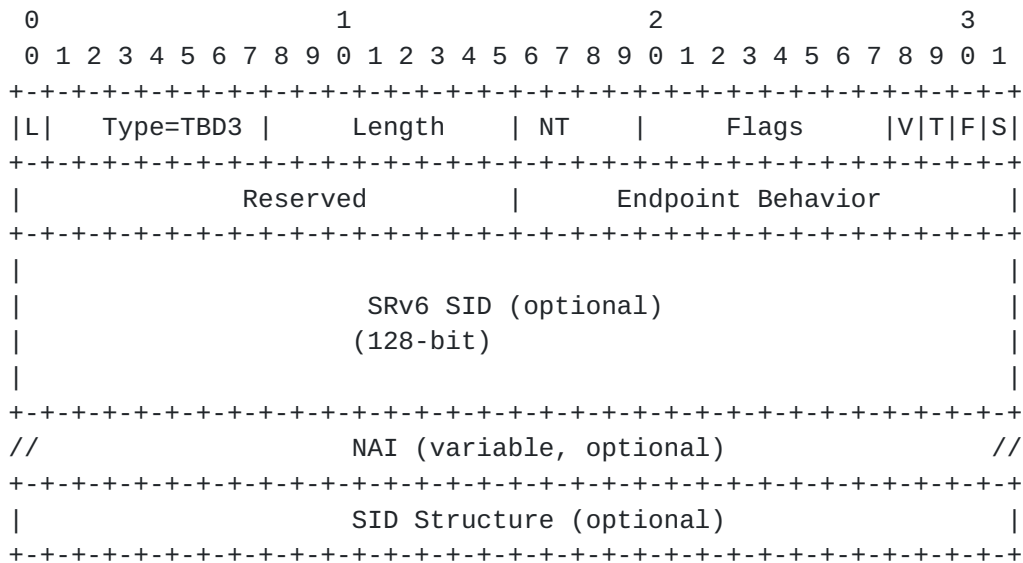


Figure 2: SRv6-ERO Subobject Format

The fields in the SRv6-ERO Subobject are as follows:

The 'L' Flag: Indicates whether the subobject represents a loose-hop (see [\[RFC3209\]](#)). If this flag is set to zero, a PCC MUST NOT overwrite the SID value present in the SRv6-ERO subobject. Otherwise, a PCC MAY expand or replace one or more SID values in the received SRv6-ERO based on its local policy.

Type: indicates the content of the subobject, i.e. when the field is set to TBD3, the subobject is a SRv6-ERO subobject representing a SRv6 SID.

Length: Contains the total length of the subobject in octets. The Length MUST be at least 24, and MUST be a multiple of 4. An SRv6-ERO subobject MUST contain at least one of a SRv6-SID or an NAI. The S and F bit in the Flags field indicates whether the SRv6-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 reuses NT types defined in [\[RFC8664\]](#):

If NT value is 0, the NAI MUST NOT be included.

When NT value is 2, the NAI is as per the 'IPv6 Node ID' format defined in [[RFC8664](#)], which specifies an IPv6 address. This is used to identify the owner of the SRv6 Identifier. This is optional, as the LOC (the locator portion) of the SRv6 SID serves a similar purpose (when present).

When NT value is 4, the NAI is as per the 'IPv6 Adjacency' format defined in [[RFC8664](#)], which specify a pair of IPv6 addresses. This is used to identify the IPv6 Adjacency and used with the SRv6 Adj-SID.

When NT value is 6, the NAI is as per the 'link-local IPv6 addresses' format defined in [[RFC8664](#)], which specify a pair of (global IPv6 address, interface ID) tuples. It is used to identify the IPv6 Adjacency and used with the SRv6 Adj-SID.

SR-MPLS specific NT types are not valid in SRv6-ERO.

Flags: Used to carry additional information pertaining to the SRv6-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.

- o S: When this bit is set to 1, the SRv6-SID value in the subobject body is absent. In this case, the PCC is responsible for choosing the SRv6-SID value, e.g., by looking up in the SR-DB using the NAI which, in this case, MUST be present in the subobject. If the S bit is set to 1 then F bit MUST be set to zero.
- o 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, and otherwise MUST be set to zero. The S and F bits MUST NOT both be set to 1.
- o T: When this bit is set to 1, the SID Structure value in the subobject body is present. The T bit MUST be set to 0 when S bit is set to 1. If the T bit is set when the S bit is set, the T bit MUST be ignored. Thus, the T bit indicates the presence of an optional 8-byte SID Structure when SRv6 SID is included. The SID Structure is defined in [Section 4.3.1.1](#).
- o V: The "SID verification" bit usage is as per Section 5.1 of [[I-D.ietf-spring-segment-routing-policy](#)].

[Editor's Note: Need to check if another flag - A flag for the SR Algorithm is required.]

Reserved: MUST be set to zero while sending and ignored on receipt.

Endpoint Behavior: A 16 bit field representing the behavior associated with the SRv6 SIDs. This information is optional and plays no role in the fields in SRH imposed on the packet. It could be used for maintainability and diagnostic purpose. If behavior is not known, 0 is used. The list of Endpoint behavior are defined in [[I-D.ietf-spring-srv6-network-programming](#)].

SRv6 SID: SRv6 Identifier is the 128 bit IPv6 addresses representing the SRv6 segment.

NAI: The NAI associated with the SRv6-SID. The NAI's format depends on the value in the NT field, and is described in [[RFC8664](#)].

At least one of the SRv6-SID or the NAI MUST be included in the SRv6-ERO subobject, and both MAY be included.

[4.3.1.1](#). SID Structure

The SID Structure is an optional part of the SR-ERO subobject, as described in [Section 4.3.1](#). It is formatted as shown in the following figure.

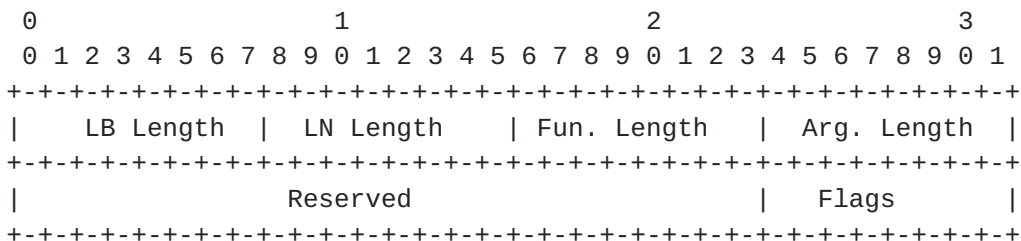


Figure 3: SID Structure Format

where:

LB Length: 1 octet. SRv6 SID Locator Block length in bits.

LN Length: 1 octet. SRv6 SID Locator Node length in bits.

Fun. Length: 1 octet. SRv6 SID Function length in bits.

Arg. Length: 1 octet. SRv6 SID Arguments length in bits.

The sum of all four sizes in the SID Structure must be lower or equal to 128 bits. If the sum of all four sizes advertised in the SID Structure is larger than 128 bits, the corresponding SRv6 SID MUST be considered as an Error. A PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("Invalid SRv6 SID Structure").

Reserved: MUST be set to zero while sending and ignored on receipt.

Flags: Currently no flags are defined. Unassigned bits must be set to zero while sending and ignored on receipt.

The SRv6 SID Structure TLV provides the detailed encoding information of an SRv6 SID, which is useful in the use cases that require to know the SRv6 SID structure. When a PCEP speaker receives the SRv6 SID and its structure information, the SRv6 SID can be parsed based on the SRv6 SID Structure TLV and/or possible local policies. The consumers of SRv6 SID structure MAY be other use cases, and the processing and usage of SRv6 SID structure TLV are different based on use cases. This is out of the scope of this document, and will be described in other documents.

4.4. RRO

In order to support SRv6, new subobject "SRv6-RRO" is defined in RRO.

4.4.1. SRv6-RRO Subobject

A PCC reports an SRv6 path 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. The procedures of [\[RFC8664\]](#) with respect to the RRO apply equally to this specification without change.

An RRO contains one or more subobjects called "SRv6-RRO subobjects" whose format is shown below:

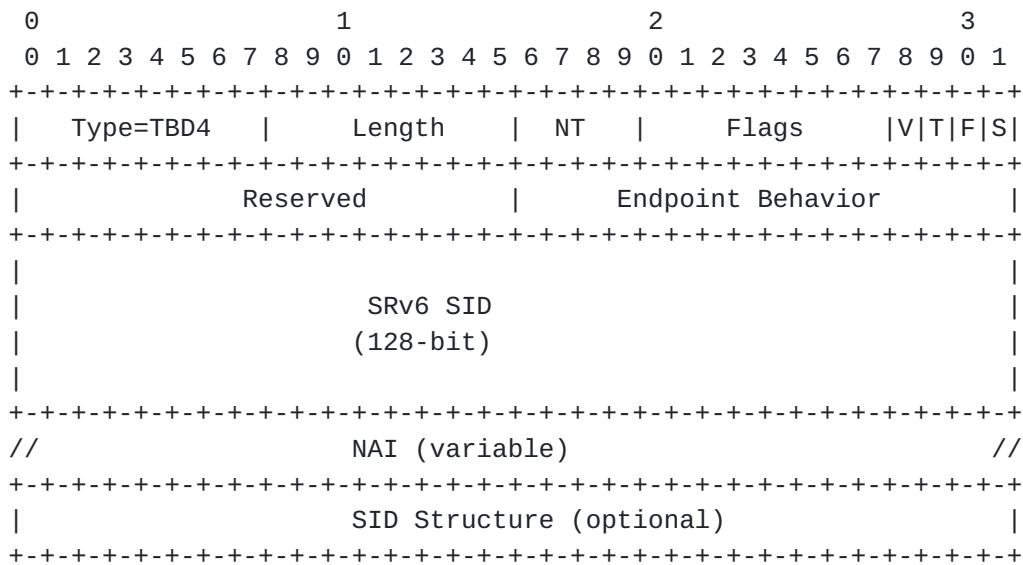


Figure 4: SRv6-RR0 Subobject Format

The format of the SRv6-RR0 subobject is the same as that of the SRv6-ERO subobject, but without the L flag.

The V flag has no meaning in the SRv6-RR0 and is ignored on receipt at the PCE.

Ordering of SRv6-RR0 subobjects by PCC in PCRpt message remains as per [\[RFC8664\]](#).

5. Procedures

5.1. Exchanging the SRv6 Capability

A PCC indicates that it is capable of supporting the head-end functions for SRv6 by including the SRv6-PCE-CAPABILITY sub-TLV in the Open message that it sends to a PCE. A PCE indicates that it is capable of computing SRv6 paths by including the SRv6-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=TBD2, but the SRv6-PCE-CAPABILITY 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 TBD5 (to be assigned by IANA) (Missing PCE-SRv6-CAPABILITY sub-TLV) and MUST then close the PCEP session. If a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV with a SRv6-PCE-CAPABILITY sub-TLV, but the PST list does not contain PST=TBD2, then the PCEP speaker MUST ignore the SRv6-PCE-CAPABILITY sub-TLV.

The number of SRv6 SIDs that can be imposed on a packet depends on the PCC's IPv6 data plane's capability. If a PCC sets the X flag to 1 then the MSD is not used and MUST NOT be included. If a PCE receives an SRv6-PCE-CAPABILITY sub-TLV with the X flag set to 1 then it MUST ignore any MSD-Type, MSD-Value fields and MUST assume that the sender can impose any length of SRH. If a PCC sets the X flag to zero, then it sets the SRv6 MSD-Type, MSD-Value fields that it can impose on a packet. If a PCE receives an SRv6-PCE-CAPABILITY sub-TLV with the X flag and SRv6 MSD-Type, MSD-Value fields both set to zero then it is considered as an error and the PCE MUST respond with a PCErr message (Error-Type=1 "PCEP session establishment failure" and Error-Value=1 "reception of an invalid Open message or a non Open message."). In case the MSD-Type in SRv6-PCE-CAPABILITY sub-TLV received by the PCE does not correspond to one of the SRv6 MSD types, the PCE MUST respond with a PCErr message (Error-Type=1 "PCEP session establishment failure" and Error-Value=1 "reception of an invalid Open message or a non Open message.").

Note that the MSD-Type, MSD-Value exchanged via the SRv6-PCE-CAPABILITY sub-TLV indicates the SRv6 SID imposition limit for the PCC node. However, if a PCE learns these via different means, e.g routing protocols, as specified in:

[\[I-D.li-ospf-ospfv3-srv6-extensions\]](#);

[\[I-D.ietf-lsr-isis-srv6-extensions\]](#); [\[I-D.ietf-idr-bgpls-srv6-ext\]](#),

then it ignores the values in the SRv6-PCE-CAPABILITY sub-TLV.

Furthermore, whenever a PCE learns the other advanced SRv6 MSD via different means, it MUST use that value regardless of the values exchanged in the SRv6-PCE-CAPABILITY sub-TLV.

Once an SRv6-capable PCEP session is established with a non-zero SRv6 MSD value, the corresponding PCE MUST NOT send SRv6 paths with a number of SIDs exceeding that SRv6 MSD value (based on the SRv6 MSD Type). If a PCC needs to modify the SRv6 MSD value, it MUST close the PCEP session and re-establish it with the new value. If a PCEP session is established with a non-zero SRv6 MSD value, and the PCC receives an SRv6 path containing more SIDs than specified in the SRv6 MSD value (based on the SRv6 MSD type), the PCC MUST send a PCErr message with Error-Type 10 (Reception of an invalid object) and Error-Value 3 (Unsupported number of Segment ERO subobjects). If a PCEP session is established with an SRv6 MSD value of zero, then the PCC MAY specify an SRv6 MSD for each path computation request that it sends to the PCE, by including a "maximum SID depth" metric object on the request similar to [\[RFC8664\]](#).

The N flag, X flag and (MSD-Type,MSD-Value) pair inside the SRv6-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 flags to zero and not include any (MSD-Type,MSD-Value) pair in the SRv6-PCE-CAPABILITY sub-

TLV in an outbound message to a PCC. Similarly, a PCC MUST ignore N,X flag and any (MSD-Type,MSD-Value) pair received from a PCE. If a PCE receives multiple SRv6-PCE-CAPABILITY sub-TLVs in an Open message, it processes only the first sub-TLV received.

5.2. ERO Processing

The ERO processing remains as per [[RFC5440](#)] and [[RFC8664](#)].

5.2.1. SRv6 ERO Validation

If a PCC does not support the SRv6 PCE Capability and thus cannot recognize the SRv6-ERO or SRv6-RR0 subobjects, it will respond according to the rules for a malformed object per [[RFC5440](#)].

On receiving an SRv6-ERO, a PCC MUST validate that the Length field, the S bit, the F bit, the T bit, and the NT field are consistent, as follows.

- o If NT=0, the F bit MUST be 1, the S bit MUST be zero and the Length MUST be 24.
- o If NT=2, the F bit MUST be zero. If the S bit is 1, the Length MUST be 24, otherwise the Length MUST be 40.
- o If NT=4, the F bit MUST be zero. If the S bit is 1, the Length MUST be 40, otherwise the Length MUST be 56.
- o If NT=6, the F bit MUST be zero. If the S bit is 1, the Length MUST be 48, otherwise the Length MUST be 64.
- o NT types (1,3, and 5) are not valid for SRv6.
- o If T bit is 1, then S bit MUST be zero.

If a PCC finds that the NT field, Length field, S bit, F bit, and T 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 PCEP speaker that does not recognize the NT value received in SRv6-ERO subobject, it would behave as per [[RFC8664](#)].

In case a PCEP speaker receives the SRv6-ERO subobject, when the PST is not set to TBD2 or SRv6-PCE-CAPABILITY sub-TLV was not exchanged, it MUST send a PCErr message with Error-Type = 19 ("Invalid Operation") and Error-Value = TBD5 ("Attempted SRv6 when the capability was not advertised").

If a PCC receives a list of SRv6 segments, and the number of SRv6 segments exceeds the SRv6 MSD that the PCC can impose on the packet (SRH), it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("Unsupported number of Segment ERO subobjects") as per [[RFC8664](#)].

When a PCEP speaker detects that all subobjects of ERO are not of type TBD3, and if it does not handle such ERO, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("Non-identical ERO subobjects") as per [[RFC8664](#)].

5.2.2. Interpreting the SRv6-ERO

The SRv6-ERO contains a sequence of subobjects. According to [[I-D.ietf-spring-segment-routing-policy](#)], each SRv6-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 SRv6-ERO subobject represents the second segment, and so on.

The PCC interprets the SRv6-ERO by converting it to an SRv6 SRH plus a next hop. The PCC sends packets along the segment routed path by prepending the SRH onto the packets and sending the resulting, modified packet to the next hop.

5.3. RRO Processing

The syntax checking rules that apply to the SRv6-RRO subobject are identical to those of the SRv6-ERO subobject, except as noted below.

If a PCEP speaker receives an SRv6-RRO subobject in which both SRv6 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 = TBD6 ("Both SID and NAI are absent in SRv6-RRO subobject").

If a PCE detects that the subobjects of an RRO are a mixture of SRv6-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 = TBD7 ("RRO mixes SRv6-RRO subobjects with other subobject types").

6. Security Considerations

The security considerations described in [[RFC5440](#)], [[RFC8231](#)] and [[RFC8281](#)], [[RFC8664](#)], are applicable to this specification. No additional security measure is required.

7. IANA Considerations

7.1. PCEP ERO and RRO subobjects

This document defines a new subobject type for the PCEP explicit route object (ERO), and a new subobject type for the PCEP record route object (RRO). The code points for subobject types of these objects is maintained in the RSVP parameters registry, under the EXPLICIT_ROUTE and ROUTE_RECORD objects. IANA is requested to allocate code-points in the RSVP Parameters registry for each of the new subobject types defined in this document.

Object	Subobject	Subobject Type
-----	-----	-----
EXPLICIT_ROUTE	SRv6-ERO (PCEP-specific)	TBD3
ROUTE_RECORD	SRv6-RRO (PCEP-specific)	TBD4

7.2. New SRv6-ERO Flag Registry

IANA is requested to create a new sub-registry, named "SRv6-ERO Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field of the SRv6-ERO subobject. New values are to be assigned by Standards Action [[RFC8126](https://tools.ietf.org/html/rfc8126)]. Each bit should be tracked with the following qualities:

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

The following values are defined in this document:

Bit	Description	Reference
-----	-----	-----
0-7	Unassigned	
8	SID Verification (V)	This document
9	SID Structure is present (T)	This document
10	NAI is absent (F)	This document
11	SID is absent (S)	This document

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

IANA maintains a 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. IANA is requested to make the following assignment:

Value	Meaning	Reference
-----	-----	-----
TBD1	SRv6-PCE-CAPABILITY	This Document

7.4. SRv6 PCE Capability Flags

IANA is requested to create a new sub-registry, named "SRv6 Capability Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field of the SRv6-PCE-CAPABILITY sub-TLV. New values are to be assigned by Standards Action [[RFC8126](#)]. Each bit should be tracked with the following qualities:

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

The following values are defined in this document:

Bit	Description	Reference
0-13	Unassigned	
14	Node or Adjacency Identifier (NAI) is supported (N)	This document
15	Unlimited Maximum SID Depth (X)	This document

7.5. New Path Setup Type

[RFC8408] created a sub-registry within the "Path Computation Element Protocol (PCEP) Numbers" registry called "PCEP Path Setup Types". IANA is requested to allocate a new code point within this registry, as follows:

Value	Description	Reference
-----	-----	-----
TBD2	Traffic engineering path is setup using SRv6.	This Document

7.6. ERROR Objects

IANA is requested to allocate code-points in the PCEP-ERROR Object Error Types and Values registry for the following new error-values:

Error-Type	Meaning
-----	-----
10	Reception of an invalid object Error-value = TBD5 (Missing PCE-SRv6-CAPABILITY sub-TLV) Error-value = TBD6 (Both SID and NAI are absent in SRv6-RR0 subobject) Error-value = TBD7 (RR0 mixes SRv6-RR0 subobjects with other subobject types) Error-value = TBD8 (Invalid SRv6 SID Structure)
19	Invalid Operation Error-value = TBD5 (Attempted SRv6 when the capability was not advertised)

8. Acknowledgements

The authors would like to thank Jeff Tentsura, Adrian Farrel, Aijun Wang and Khasanov Boris for valuable suggestions.

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>>.
- [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>>.

- [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>>.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", [RFC 5511](#), DOI 10.17487/RFC5511, April 2009, <<https://www.rfc-editor.org/info/rfc5511>>.
- [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>>.
- [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>>.

- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", [RFC 8664](#), DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extension to Support Segment Routing over IPv6 Dataplane", [draft-ietf-lsr-isis-srv6-extensions-11](#) (work in progress), October 2020.
- [I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-24](#) (work in progress), October 2020.

9.2. Informative References

- [RFC4657] Ash, J., Ed. and J. 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>>.
- [RFC7855] Previdi, S., Ed., Filsfils, C., Ed., Decraene, B., Litkowski, S., Horneffer, M., and R. Shakir, "Source Packet Routing in Networking (SPRING) Problem Statement and Requirements", [RFC 7855](#), DOI 10.17487/RFC7855, May 2016, <<https://www.rfc-editor.org/info/rfc7855>>.
- [RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", [RFC 8051](#), DOI 10.17487/RFC8051, January 2017, <<https://www.rfc-editor.org/info/rfc8051>>.
- [RFC8354] Brzozowski, J., Leddy, J., Filsfils, C., Maglione, R., Ed., and M. Townsley, "Use Cases for IPv6 Source Packet Routing in Networking (SPRING)", [RFC 8354](#), DOI 10.17487/RFC8354, March 2018, <<https://www.rfc-editor.org/info/rfc8354>>.
- [RFC8666] Psenak, P., Ed. and S. Previdi, Ed., "OSPFv3 Extensions for Segment Routing", [RFC 8666](#), DOI 10.17487/RFC8666, December 2019, <<https://www.rfc-editor.org/info/rfc8666>>.

- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [RFC 8754](#), DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-09](#) (work in progress), November 2020.
- [I-D.li-ospf-ospfv3-srv6-extensions]
Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", [draft-li-ospf-ospfv3-srv6-extensions-07](#) (work in progress), November 2019.
- [I-D.ietf-idr-bgpls-srv6-ext]
Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., daniel.bernier@bell.ca, d., and B. Decraene, "BGP Link State Extensions for SRv6", [draft-ietf-idr-bgpls-srv6-ext-04](#) (work in progress), November 2020.

Appendix A. Contributor

The following persons contributed to this document:

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

EMail: dhruv.ietf@gmail.com

Huang Wumin
Huawei Technologies
Huawei Building, No. 156 Beiqing Rd.
Beijing 100095
China

Email: huangwumin@huawei.com

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

Email: pengshuping@huawei.com

Authors' Addresses

Cheng Li(Editor)
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

EMail: c.l@huawei.com

Mahendra Singh Negi
RtBrick Inc
Bangalore, Karnataka
India

EMail: mahend.ietf@gmail.com

Mike Koldychev
Cisco Systems, Inc.
Canada

EMail: mkoldych@cisco.com

Prejeeth Kaladharan
RtBrick Inc
Bangalore, Karnataka
India

EMail: prejeeth@rtbrick.com

Yongqing Zhu
China Telecom
109 West Zhongshan Ave, Tianhe District
Bangalore, Guangzhou
P.R. China

EMail: zhuyq8@chinatelecom.cn

