

Internet Engineering Task Force (IETF)
Request for Comments: 9603
Category: Standards Track
ISSN: 2070-1721

李呈 (C. Li), Ed.
技术有限公司 (Huawei Technologies)
P. Kaladharan
RtBrick Inc
S. Sivabalan
M. Koldychev
Ciena Corporation
Y. Zhu
China Telecom
July 2024

Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing

Abstract

Segment Routing (SR) can be used to steer packets through a network using the IPv6 or MPLS data plane, employing the source routing paradigm.

An SR 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 data planes, a PCE should be able to compute an SR Path for both MPLS and IPv6 data planes. The Path Computation Element Communication Protocol (PCEP) extension and mechanisms to support SR-MPLS have been defined. This document outlines the necessary extensions to support SR for the IPv6 data plane within PCEP.

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/rfc9603>.

Copyright Notice

Copyright (c) 2024 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction
 - 1.1. Requirements Language
 2. Terminology
 3. Overview of PCEP Operation in SRv6 Networks
 - 3.1. Operation Overview
 - 3.2. SRv6-Specific PCEP Message Extensions
 4. Object Formats
 - 4.1. The OPEN Object
 - 4.1.1. The SRv6 PCE Capability sub-TLV
 - 4.2. The RP/SRP Object
 - 4.3. ERO
 - 4.3.1. SRv6-ERO Subobject
 - 4.3.1.1. SID Structure
 - 4.3.1.2. Order of the Optional Fields
 - 4.4. RRO
 - 4.4.1. SRv6-RRO Subobject
 5. Procedures
 - 5.1. Exchanging the SRv6 Capability
 - 5.2. ERO Processing
 - 5.2.1. SRv6 ERO Validation
 - 5.2.2. Interpreting the SRv6-ERO
 - 5.3. RRO Processing
 6. Security Considerations
 7. Manageability Considerations
 - 7.1. Control of Function and Policy
 - 7.2. Information and Data Models
 - 7.3. Liveness Detection and Monitoring
 - 7.4. Verify Correct Operations
 - 7.5. Requirements on Other Protocols
 - 7.6. Impact on Network Operations
 8. IANA Considerations
 - 8.1. PCEP ERO and RRO Subobjects
 - 8.2. New SRv6-ERO NAI Type Registry
 - 8.3. New SRv6-ERO Flag Registry
 - 8.4. LSP-ERROR-CODE TLV
 - 8.5. PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators
 - 8.6. SRv6 PCE Capability Flags
 - 8.7. New Path Setup Type
 - 8.8. ERROR Objects
 9. References
 - 9.1. Normative References
 - 9.2. Informative References
- Acknowledgements
- Contributors
- Authors' Addresses

1. Introduction

As defined in [RFC8402], Segment Routing (SR) architecture allows the source node to steer a packet through a path indicated by an ordered list of instructions, called "segments". A segment can represent any instruction, topological or service based, and it can have a semantic local to an SR node or global within an SR domain.

[RFC5440] describes Path Computation Element Communication Protocol (PCEP) for communication between a Path Computation Client (PCC) and a PCE or between a pair of PCEs. A PCE or a PCC operating as a PCE (in a hierarchical PCE environment) computes paths for MPLS Traffic Engineering Label Switched Paths (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, and 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 computed, the stateful PCE can initiate an SR-TE path on a PCC using PCEP extensions specified in [RFC8281] and the SR-specific PCEP extensions specified in [RFC8664].

[RFC8664] specifies PCEP extensions for supporting an SR-TE LSP for the MPLS data plane. This document extends [RFC8664] to support SR for the IPv6 data plane. Additionally, using procedures described in this document, a PCC can request an SRv6 path from either a stateful or 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 [RFC9256], which applies to both SR-MPLS and SRv6.

1.1. Requirements Language

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

2. Terminology

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

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

Further, the following terms are used in the document:

MSD: Maximum SID Depth

PST: Path Setup Type

SR: Segment Routing

SID: Segment Identifier

SRv6: Segment Routing over IPv6 data plane

SRH: IPv6 Segment Routing Header [RFC8754]

SRv6 path: IPv6 Segment List (A list of IPv6 SIDs representing a path in IPv6 SR domain in the context of this document.)

Further, note that the term "LSP" used in the PCEP specifications would be equivalent to an 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 are built on [RFC8664].

In PCEP messages, route information is carried in the Explicit Route Object (ERO), which consists of a sequence of subobjects. [RFC8664] defined 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 for SR-MPLS. SR-capable PCEP speakers can 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 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]. [RFC8664] also defines a new Reported Route Object (RRO), called "SR-RRO", to represent the SID list that was applied by the PCC, which is the actual path taken by the LSP in SR-MPLS network.

The SRv6 paths computed by a PCE can be represented as an ordered list of SRv6 segments. This document defines new subobjects "SRv6-ERO" and "SRv6-RRO" in the ERO and the RRO, respectively, to carry the SRv6 SID. SRv6-capable PCEP speakers MUST be able to generate and/or process these subobjects.

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 as described in Section 4.1.1.

In summary, this document defines:

- * a new PCEP capability for SRv6,
- * a new subobject SRv6-ERO in ERO,
- * a new subobject SRv6-RRO in RRO, and
- * a new Path Setup type (PST) [RFC8408], carried in the PATH-SETUP-TYPE and PATH-SETUP-TYPE-CAPABILITY TLVs.

3.1. Operation Overview

In SR networks, an SR source node [RFC8754] steers a packet into an SR Policy resulting in a segment list.

When SR leverages the IPv6 data plane (i.e., SRv6), the PCEP procedures and mechanisms are extended in this document.

This document describes the extension to support SRv6 in PCEP. 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], the PCInitiate message specified in [RFC8281], or PCRpt and PCUpd messages specified in [RFC8231]. However, PCEP messages pertaining to SRv6 MUST include PATH-SETUP-TYPE TLV in the Request Parameters (RP) or Stateful PCE Request Parameters (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:

PST=3: Path is set up using SRv6.

A PCEP speaker indicates 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 (value 3) 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=3 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. For further error handling, please see Section 5.

The format of the SRv6-PCE-CAPABILITY sub-TLV is shown in Figure 1.




Figure 1: SRv6-PCE-CAPABILITY Sub-TLV Format

The code point for the TLV type is 27, and the 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. When set to 27, the Type field 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 sub-TLV is padded to 4-octet alignment, and padding is not included in the Length field. The (MSD-Type,MSD-Value) pairs are OPTIONAL. The number of (MSD-Type,MSD-Value) pairs can be determined by the Length field of the TLV.

The value is comprised of:

- * Reserved: 2 octets; this field MUST be set to 0 on transmission and ignored on receipt.
- * Flags: 2 octets; one bit is currently assigned in Section 8.6
 - N bit (bit position 14): A PCC sets this flag bit to 1 to indicate that it is capable of resolving a Node or Adjacency Identifier (NAI) to an SRv6-SID.
 - Unassigned bits MUST be set to 0 on transmission 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 [RFC9352], and where MSD-Value (1 octet) is as per [RFC8491].

The SRv6 MSD information advertised via SRv6-PCE-Capability sub-TLV conveys the SRv6 capabilities of the PCEP speaker alone. However, when it comes to the computation of an SR Policy for the SRv6 data plane, the SRv6 MSD capabilities of the intermediate SRv6 Endpoint node and the tail-end node also need to be considered to ensure those midpoints are able to correctly process their segments and for the tail-end to dispose of the SRv6 encapsulation. The SRv6 MSD capabilities of other nodes might be learned as part of the topology information via the Border Gateway Protocol - Link State (BGP-LS) [RFC9514] or via PCEP if the PCE also happens to have PCEP sessions with those nodes.

It is recommended that the SRv6 MSD information not be included in the SRv6-PCE-Capability sub-TLV in deployments where the PCE is able to obtain this via IGP/BGP-LS as part of the topology information.

4.2. The RP/SRP Object

This document defines a new Path Setup Type (PST=3) for SRv6. In order to indicate that the path is for SRv6, any RP or SRP object MUST include the PATH-SETUP-TYPE TLV as specified in [RFC8408], where PST is set to 3.

4.3. ERO

In order to support SRv6, a new "SRv6-ERO" subobject is defined for inclusion in the ERO.

4.3.1. SRv6-ERO Subobject

An SRv6-ERO subobject is formatted as shown in Figure 2.

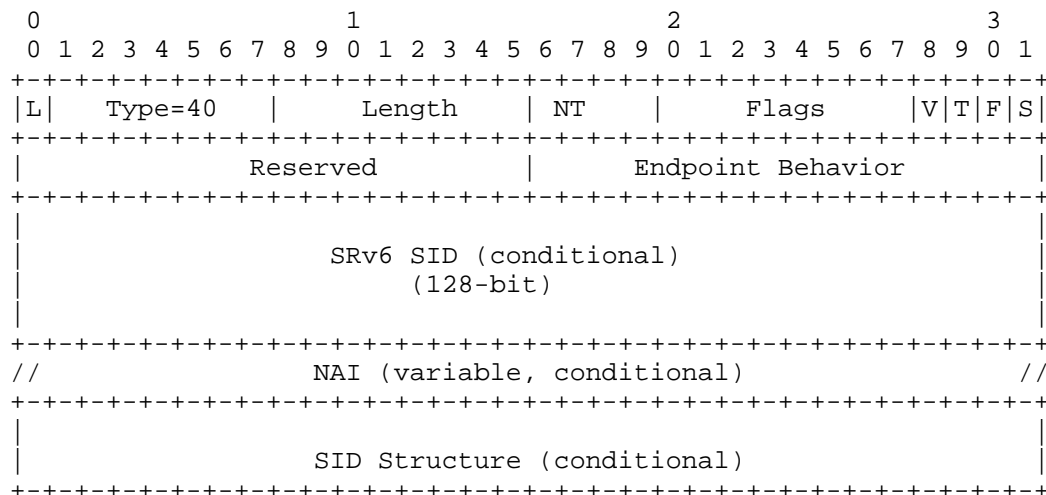


Figure 2: SRv6-ERO Subobject Format

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 40, the subobject is an SRv6-ERO subobject representing an 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 an SRv6-SID or an NAI. The S and F bits 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 are present. If the F bit is set to one (see below), then the NT field has no meaning and MUST be ignored by the receiver. This document creates a new PCEP SRv6-ERO NAI Types registry in Section 8.2 and allocates the following values:
 - 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.
- * 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. This document creates a new registry SRv6-ERO Flag Field registry in Section 8.3 and allocates the following values.
 - 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 that, in this case, MUST be present in the subobject. If the S bit is set to 1, then the F bit 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.
 - 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 the 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.
 - V: The "SID verification" bit usage is as per Section 5.1 of [RFC9256]. If a PCC "Verification fails" for a SID, it MUST report this error by including the LSP-ERROR-CODE TLV with LSP Error-value "SID Verification fails" in the LSP object in the PCRpt message to the PCE.
- * 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, but providing it is recommended whenever possible. It could be used for maintainability and diagnostic purposes. If behavior is not known, value "0xFFFF" as defined in the "SRv6 Endpoint Behaviors" registry is used [RFC8986].
- * SRv6 SID: SRv6 Identifier is a 128-bit value 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 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.

[RFC8986] defines an SRv6 SID as consisting of LOC:FUNCT:ARG, where a locator (LOC) is encoded in the L most significant bits of the SID, followed by F bits of function (FUNCT) and A bits of arguments (ARG). A locator may be represented as B:N where B is the SRv6 SID locator block (IPv6 prefix allocated for SRv6 SIDs by the operator) and N is the identifier of the parent node instantiating the SID called "locator node".

The SID Structure is formatted as shown in Figure 3.

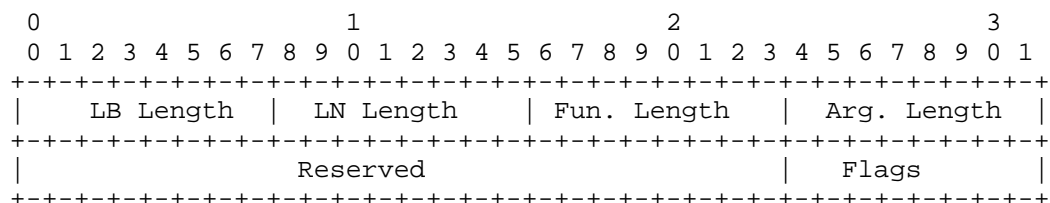


Figure 3: SID Structure Format

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 less than 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 invalid and a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 37 ("Invalid SRv6 SID Structure") is returned.

- * 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 provides the detailed encoding information of an SRv6 SID, which is helpful in the use cases that require the SRv6 SID structure to be known. When a PCEP speaker receives the SRv6 SID and its structure information, the SRv6 SID can be parsed based on the SRv6 SID Structure and/or possible local policies. The SRv6 SID Structure could be used by the PCE for ease of operations and monitoring. For example, this information could be used for validation of SRv6 SIDs being instantiated in the network and checked for conformance with the SRv6 SID allocation scheme chosen by the operator as described in Section 3.2 of [RFC8986]. In the future, PCE might also be utilized to verify and automate the security of the SRv6 domain by provisioning filtering rules at the domain boundaries as described in Section 5 of [RFC8754]. The details of these potential applications are outside the scope of this document.

4.3.1.2. Order of the Optional Fields

The optional elements in the SRv6-ERO subobject, i.e., SRv6 SID, NAI, and the SID Structure, MUST be encoded in the order as depicted in Figure 2. The presence or absence of each of them is indicated by the respective flags, i.e., S flag, F flag, and T flag.

In order to ensure future compatibility, any optional elements added to the SRv6-ERO subobject in the future must specify their order and request that the Internet Assigned Numbers Authority (IANA) allocate a flag to indicate their presence from the subregistry created in Section 8.3.

4.4. RRO

In order to support SRv6, a new "SRv6-RRO" subobject is defined for inclusion in the 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 in Figure 4.

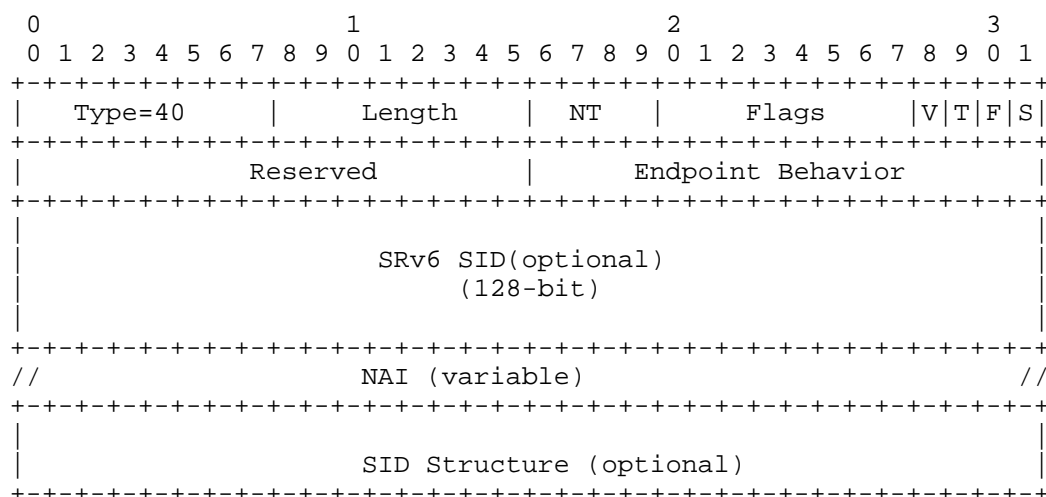


Figure 4: SRv6-RRO Subobject Format

The format of the SRv6-RRO 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-RRO and is ignored on receipt at the PCE.

The ordering of SRv6-RRO subobjects by PCC in PCRpt message remains as per [RFC8664].

The ordering of optional elements in the SRv6-RRO subobject is the same as described in Section 4.3.1.2.

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=3, 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 = 34 ("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 an SRv6-PCE-CAPABILITY sub-TLV, but the PST list does not contain PST=3, then the PCEP speaker MUST ignore the SRv6-PCE-CAPABILITY sub-TLV.

In case the MSD-Type in the 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) pair exchanged via the SRv6-PCE-CAPABILITY sub-TLV indicates the SRv6 SID imposition limit for the sender PCC node only. However, if a PCE learns these via alternate mechanisms, e.g., routing protocols [RFC9352], then it ignores the values in the SRv6-PCE-CAPABILITY sub-TLV. Furthermore, whenever a PCE learns any other SRv6 MSD types that may be defined in the future via alternate mechanisms, it MUST use those values regardless of the values exchanged in the SRv6-PCE-CAPABILITY sub-TLV.

During path computation, a PCE must consider the MSD information of all the nodes along the path instead of only the MSD information of the ingress PCC since a packet may be dropped on any node in a forwarding path because of the SID depth exceeding the MSD of the node. The MSD capabilities of all SR nodes along the path can be learned as part of the topology information via IGP/BGP-LS or via PCEP if the PCE also happens to have PCEP sessions with those nodes.

A PCE MUST NOT send SRv6 paths that exceed the SRv6 MSD capabilities of the PCC. If a PCC needs to modify the SRv6 MSD value signaled via the Open message, it MUST close the PCEP session and re-establish it with the new value. If the PCC receives an SRv6 path that exceeds its SRv6 MSD capabilities, the PCC MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 40 ("Unsupported number of SRv6-ERO subobjects").

The N flag and (MSD-Type,MSD-Value) pair inside the SRv6-PCE-CAPABILITY sub-TLV are meaningful only in the Open message sent to a PCE. As such, the flags MUST be set to zero and a (MSD-Type,MSD-

Value) pair MUST NOT be present in the SRv6-PCE-CAPABILITY sub-TLV in an Open message sent to a PCC. Similarly, a PCC MUST ignore flags and any (MSD-Type,MSD-Value) pair in a received Open message. 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 processing of ERO remains unchanged in accordance with both [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-RRO subobjects, it should respond according to the rules for a malformed object as described in [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:

- * If NT=0, the F bit MUST be 1, the S bit MUST be zero, and the Length MUST be 24.
- * 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.
- * 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.
- * 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.
- * If the T bit is 1, then the 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 PCC does not recognize or support the value in the NT field, it MUST consider the entire ERO invalid and send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 41 ("Unsupported NAI Type in the SRv6-ERO/SRv6-RRO subobject").

If a PCC receives an SRv6-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 = 42 ("Both SID and NAI are absent in the SRv6-ERO subobject").

If a PCC receives an SRv6-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 detects that the subobjects of an ERO are a mixture of SRv6-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 = 43 ("ERO mixes SRv6-ERO subobjects with other subobject types").

In case a PCEP speaker receives an SRv6-ERO subobject, when the PST

is not set to 3 or SRv6-PCE-CAPABILITY sub-TLV was not exchanged, it MUST send a PCErr message with Error-Type = 19 ("Invalid Operation") and Error-value = 19 ("Attempted SRv6 when the capability was not advertised").

If a PCC receives an SRv6 path that exceeds the SRv6 MSD capabilities, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 40 ("Unsupported number of SRv6-ERO subobjects") as per [RFC8664].

5.2.2. Interpreting the SRv6-ERO

The SRv6-ERO contains a sequence of subobjects. According to [RFC9256], 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.

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 = 35 ("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 = 36 ("RRO mixes SRv6-RRO subobjects with other subobject types").

The mechanism by which the PCC learns the path is outside the scope of this document.

6. Security Considerations

The Security Considerations described in [RFC5440], Section 2.5 of [RFC6952], [RFC8231], [RFC8281], [RFC8253], and [RFC8664] are applicable to this specification.

Note that this specification enables a network controller to instantiate an SRv6 path in the network. 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 an SRv6 path that may not be subjected to the further verification checks. Further, the MSD field in the Open message could disclose node forwarding capabilities if suitable security mechanisms are not in place. Hence, securing the PCEP session using Transport Layer Security (TLS) [RFC8253] is RECOMMENDED.

7. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231], [RFC8281], and [RFC8664] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

7.1. Control of Function and Policy

A PCEP implementation SHOULD allow the operator to configure the SRv6

capability. Further, a policy to accept NAI only for the SRv6 SHOULD be allowed to be set.

7.2. Information and Data Models

The PCEP YANG module is out of the scope of this document; it is defined in other documents, for example, [PCEP-YANG]. An augmented YANG module for SRv6 is also specified in [PCEP-SRv6-YANG] that allows for SRv6 capability and MSD configurations as well as to monitor the SRv6 paths set in the network.

7.3. Liveness Detection and Monitoring

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

7.4. Verify Correct Operations

Verification of the mechanisms defined in this document can be built on those already listed in [RFC5440], [RFC8231], and [RFC8664].

7.5. Requirements on Other Protocols

Mechanisms defined in this document do not imply any new requirements on other protocols.

7.6. Impact on Network Operations

Mechanisms defined in [RFC5440], [RFC8231], and [RFC8664] also apply to PCEP extensions defined in this document.

8. IANA Considerations

8.1. 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 Reported Route Object (RRO). These have been registered in the "Resource Reservation Protocol (RSVP) Parameters" registry group as shown below.

IANA has allocated the following new subobject in the "Subobject type - 20 EXPLICIT_ROUTE - Type 1 Explicit Route" registry:

Value	Description
40	SRv6-ERO (PCEP-specific)

Table 1

IANA has allocated the following new subobject in the "Subobject type - 21 ROUTE_RECORD - Type 1 Route Record" registry:

Value	Description
40	SRv6-RRO (PCEP-specific)

Table 2

8.2. New SRv6-ERO NAI Type Registry

IANA has created the "PCEP SRv6-ERO NAI Types" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group to manage the 4-bit NT field in the SRv6-ERO subobject. The registration policy is IETF Review [RFC8126]. IANA has registered the values in Table 3.

Value	Description	Reference
0	NAI is absent.	RFC 9603
2	NAI is an IPv6 node ID.	RFC 9603
4	NAI is an IPv6 adjacency with global IPv6 addresses.	RFC 9603
6	NAI is an IPv6 adjacency with link-local IPv6 addresses.	RFC 9603

Table 3

8.3. New SRv6-ERO Flag Registry

IANA has created the "SRv6-ERO Flag Field" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group to manage the 12-bit Flag field of the SRv6-ERO subobject. New values are to be assigned by Standards Action [RFC8126]. Each registration should include the following information:

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

The following values are defined in this document:

Bit	Description	Reference
8	SID Verification (V)	RFC 9603
9	SID Structure is present (T)	RFC 9603
10	NAI is absent (F)	RFC 9603
11	SID is absent (S)	RFC 9603

Table 4

8.4. LSP-ERROR-CODE TLV

This document defines a new value in "LSP-ERROR-CODE TLV Error Code Field" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group.

Value	Meaning	Reference
10	SID Verification fails	RFC 9603

Table 5

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

IANA maintains the "PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group to manage the type indicator space for sub-TLVs of the PATH-SETUP-TYPE-CAPABILITY TLV. IANA has registered the following value:

Value	Meaning	Reference
27	SRv6-PCE-CAPABILITY	RFC 9603

Table 6

8.6. SRv6 PCE Capability Flags

IANA has created the "SRv6 Capability Flag Field" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group to manage the 16-bit Flag field of the SRv6-PCE-CAPABILITY sub-TLV. New values are to be assigned by Standards Action [RFC8126]. Each registration should include the following information:

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

The following value is defined in this document.

Bit	Description	Reference
14	Node or Adjacency Identifier (NAI) is supported (N)	RFC 9603

Table 7

8.7. New Path Setup Type

[RFC8408] created the "PCEP Path Setup Types" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group. IANA has allocated the following value:

Value	Description	Reference
3	Traffic engineering path is set up using SRv6.	RFC 9603

Table 8

8.8. ERROR Objects

IANA has allocated the following Error-values in the "PCEP-ERROR Object Error Types and Values" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group:

Error-Type	Meaning	Error-value
10	Reception of an	34: Missing PCE-SRv6-CAPABILITY

	invalid object	sub-TLV
		35: Both SID and NAI are absent in SRv6-RRO subobject
		36: RRO mixes SRv6-RRO subobjects with other subobject types
		37: Invalid SRv6 SID Structure
		40: Unsupported number of SRv6-ERO subobjects
		41: Unsupported NAI Type in the SRv6-ERO/SRv6-RRO subobject
		42: Both SID and NAI are absent in the SRv6-ERO subobject
		43: ERO mixes SRv6-ERO subobjects with other subobject types
19	Invalid Operation	19: Attempted SRv6 when the capability was not advertised

Table 9

9. References

9.1. Normative References

- [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>>.
- [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>>.
- [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>>.
- [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>>.

- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <<https://www.rfc-editor.org/info/rfc8253>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", RFC 9514, DOI 10.17487/RFC9514, December 2023, <<https://www.rfc-editor.org/info/rfc9514>>.
- [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>>.
- [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>>.

9.2. Informative References

- [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>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <<https://www.rfc-editor.org/info/rfc6952>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture",

RFC 9256, DOI 10.17487/RFC9256, July 2022,
<<https://www.rfc-editor.org/info/rfc9256>>.

[RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B.,
and Z. Hu, "IS-IS Extensions to Support Segment Routing
over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352,
February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.

[PCEP-YANG]
Dhody, D., Ed., Beeram, V., Hardwick, J., and J. Tantsura,
"A YANG Data Model for Path Computation Element
Communications Protocol (PCEP)", Work in Progress,
Internet-Draft, draft-ietf-pce-pcep-yang-25, 21 May 2024,
<<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-yang-25>>.

[PCEP-SRV6-YANG]
Li, C., Sivabalan, S., Peng, S., Koldychev, M., and L.
Ndifor, "A YANG Data Model for Segment Routing (SR) Policy
and SR in IPv6 (SRv6) support in Path Computation Element
Communications Protocol (PCEP)", Work in Progress,
Internet-Draft, draft-ietf-pce-pcep-srv6-yang-05, 18 March
2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-srv6-yang-05>>.

Acknowledgements

The authors would like to thank Jeff Tantsura, Adrian Farrel, Aijun Wang, Khasanov Boris, Ketan Talaulikar, Martin Vigoureux, Hariharan Ananthakrishnan, Xinyue Zhang, John Scudder, Julien Meuric, and Robert Varga for valuable suggestions.

Thanks to Gunter Van de Velde, ric Vyncke, Jim Guichard, and Mahesh Jethanandani for their comments during the IESG review.

Contributors

Mahendra Singh Negi
RtBrick Inc
Bangalore
Karnataka
India
Email: mahend.ietf@gmail.com

Dhruv Dhody
Huawei
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

Ran Chen
ZTE Corporation
China
Email: chen.ran@zte.com.cn

Authors' Addresses

Cheng Li (editor)
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China
Email: c.l@huawei.com

Additional contact information:

李呈 (editor)
中国
100095
北京
北研所
技术有限公司

Prejeeth Kaladharan
RtBrick Inc
Bangalore
Karnataka
India
Email: prejeeth@rtbrick.com

Siva Sivabalan
Ciena Corporation
Email: msiva282@gmail.com

Mike Koldychev
Ciena Corporation
Canada
Email: mkoldych@ciena.com

Yongqing Zhu
China Telecom
109 West Zhongshan Ave, Tianhe District
Bangalore
Guangzhou,
China
Email: zhuyq8@chinatelecom.cn