

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 20, 2020

C. Li
M. Chen
Huawei Technologies
W. Cheng
China Mobile
J. Dong
Z. Li
Huawei Technologies
R. Gandhi
Cisco Systems, Inc.
Q. Xiong
ZTE Corporation
August 19, 2019

Path Computation Element Communication Protocol (PCEP) Extension for
Path Segment in Segment Routing (SR)
draft-li-pce-sr-path-segment-08

Abstract

The Path Computation Element (PCE) provides path computation functions in support of traffic engineering in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

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. 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).

Path identification is needed for several use cases such as performance measurement in Segment Routing (SR) network. This document specifies extensions to the Path Computation Element Communication Protocol (PCEP) to support requesting, replying, reporting and updating the Path Segment ID (Path SID) between PCEP speakers.

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 February 20, 2020.

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	3
2.	Terminology	4
2.1.	Requirements Language	4
3.	Overview of Path Segment Extensions in PCEP	4
4.	Objects and TLVs	5
4.1.	The OPEN Object	5
4.1.1.	The SR PCE Capability sub-TLV	5
4.1.2.	PCECC-CAPABILITY sub-TLV	6
4.2.	LSP Object	6
4.2.1.	Path Segment TLV	7
4.3.	FEC Object	8
4.4.	CCI Object	9
5.	Operations	10
5.1.	Stateful PCE Operation	10
5.1.1.	Ingress PCC-Initiated Path Segment Allocation	10
5.1.2.	PCE Initiated Path Segment Allocation	12
5.2.	PCECC Based Operation	13
5.2.1.	PCE Controlled Label Spaces Advertisement	13
5.2.2.	PCECC based Path Segment Allocation	13
6.	Dataplane Considerations	15
7.	Implementation Status	16
7.1.	Huawei's Commercial Delivery	16
7.2.	ZTE's Commercial Delivery	17
8.	IANA Considerations	17

8.1.	SR PCE Capability Flags	17
8.2.	New LSP Flag Registry	17
8.3.	New PCEP TLV	17
8.3.1.	Path Segment TLV	18
8.4.	New FEC Type Registry	18
8.5.	PCEP Error Type and Value	19
9.	Security Considerations	19
10.	Manageability Considerations	19
10.1.	Control of Function and Policy	19
10.2.	Information and Data Models	20
10.3.	Liveness Detection and Monitoring	20
10.4.	Verify Correct Operations	20
10.5.	Requirements On Other Protocols	20
10.6.	Impact On Network Operations	20
11.	Acknowledgments	20
12.	References	20
12.1.	Normative References	20
12.2.	Informative References	22
Appendix A.	Contributors	23
Appendix B.	SRv6 extensions	23
B.1.	The SRv6 PCE Capability sub-TLV	24
B.2.	SRv6 PCE Capability Flags	24
B.3.	Path Segment TLV	25
Authors' Addresses	25

1. Introduction

[RFC5440] describes the Path Computation Element (PCE) Communication Protocol (PCEP). PCEP enables the communication between a Path Computation Client (PCC) and a PCE, or between PCE and PCE, for the purpose of computation of Multiprotocol Label Switching (MPLS) as well as Generalized MPLS (GMPLS) Traffic Engineering Label Switched Path (TE LSP) characteristics.

[RFC8231] specifies a set of extensions to PCEP to enable stateful control of TE LSPs within and across PCEP sessions in compliance with [RFC4657]. It includes mechanisms to effect LSP State Synchronization between PCCs and PCEs, delegation of control over LSPs to PCEs, and PCE control of timing and sequence of path computations within and across PCEP sessions. The model of operation where LSPs are initiated from the PCE is described in [RFC8281].

[I-D.ietf-pce-pcep-extension-for-pce-controller] specify the procedures and PCEP protocol extensions for using the PCE as the central controller for static LSPs, where LSPs can be provisioned as explicit label instructions at each hop on the end-to-end path.

Segment routing (SR) [RFC8402] leverages the source routing and tunneling paradigms and supports steering packets into an explicit forwarding path at the ingress node.

An SR path needs to be identified in some use cases such as performance measurement. For identifying an SR path, [I-D.ietf-spring-mpls-path-segment] introduces a new segment that is referred to as Path Segment.

[I-D.ietf-pce-segment-routing] specifies extensions to the Path Computation Element Protocol (PCEP) [RFC5440] for SR networks, that allow a stateful PCE to compute and initiate SR-TE paths, as well as a PCC to request, report or delegate SR paths.

[I-D.zhao-pce-pcep-extension-pce-controller-sr] specifies the procedures and PCEP protocol extensions when a PCE-based controller is also responsible for configuring the forwarding actions on the routers (SR SID distribution in this case), in addition to computing the paths for packet flows in a segment routing network and telling the edge routers what instructions to attach to packets as they enter the network.

This document specifies a mechanism to carry the SR path identification information in PCEP messages [RFC5440] [RFC8231] [RFC8281]. The SR path identifier can be a Path Segment in SR-MPLS [I-D.ietf-spring-mpls-path-segment], or other IDs that can identify an SR path. This document also extends the PCECC-SR mechanism to inform the Path Segment to the egress PCC.

2. Terminology

This memo makes use of the terms defined in [RFC4655], [I-D.ietf-pce-segment-routing], and [RFC8402].

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 Path Segment Extensions in PCEP

This document specifies a mechanism of allocating Path Segment and extends PCEP to encode it in PCEP messages. For supporting Path Segment in PCEP, several TLVs and flags are defined. The formats of

the objects and TLVs are described in Section 4. The procedures of Path Segment allocation are described in Section 5.

There are various modes of operations, such as -

- o The Path Segment can be allocated by Egress PCC. The PCE should request the Path Segment from Egress PCC.
- o The PCE can allocate a Path Segment on its own accord and inform the ingress/egress PCC, useful for PCE-initiated LSPs.
- o Ingress PCC can also request PCE to allocate the Path Segment, in this case, the PCE would either allocate and inform the assigned Path Segment to the ingress/egress PCC using PCEP messages, or first request egress PCC for Path Segment and then inform it to the ingress PCC.

The path information to the ingress PCC and PCE is exchanged via an extension to [I-D.ietf-pce-segment-routing] and [I-D.ietf-pce-segment-routing-ipv6]. The Path Segment information (for SR-MPLS) to the egress PCC can be informed via an extension to the PCECC-SR procedures [I-D.zhao-pce-pcep-extension-pce-controller-sr].

For the PCE to allocate a Path Segment on its own, the PCE needs to be aware of the MPLS label space from the PCCs. This is done via mechanism as described in [I-D.li-pce-controlled-id-space]. Otherwise, the PCE should request the egress PCC for Path Segment allocation.

4. Objects and TLVs

4.1. The OPEN Object

4.1.1. The SR PCE Capability sub-TLV

[I-D.ietf-pce-segment-routing] defined a new Path Setup Type (PST) and SR-PCE-CAPABILITY sub-TLV for SR-MPLS. PCEP speakers use this sub-TLV to exchange information about their SR capability. The TLV defines a Flags field [I-D.ietf-pce-segment-routing].

This document adds an additional flag for Path Segment allocation, as follows -

- o P (Path Segment Identification bit): A PCEP speaker sets this flag to 1 to indicate that it has the capability to encode SR path identification (Path Segment, as per [I-D.ietf-spring-mpls-path-segment]).




Figure 1: P-flag in SR-PCE-CAPABILITY TLV

The figure is included for the ease of the reader and will be removed at the time of publication.

4.1.2. PCECC-CAPABILITY sub-TLV

Along with the SR sub-TLVs, the PCECC Capability as per [I-D.zhao-pce-pcep-extension-pce-controller-sr] should be advertised if the PCE allocates the Path Segment and acts as a Central Controller that manages the Label space.

The PCECC Capability should be advertised on the egress PCEP session, along with the SR sub-TLVs. This is needed to ensure that the PCE can use the PCECC objects/mechanism to request/inform the egress PCC of the Path Segment as described in Section 5.2.

4.2. LSP Object

The LSP Object is defined in Section 7.3 of [RFC8231]. This document adds a flag in the LSP Object:

- o P (PCE Allocation bit): If the bit is set to 1, it indicates that the PCC requests PCE to make allocations for this LSP. The TLV in LSP object identifies what should be allocated, such as Path Segment or Binding Segment. A PCC would set this bit to 1 and include a PATH-SEGMENT TLV in the LSP object to request for allocation of Path Segment by the PCE in the PCEP message. A PCE would also set this bit to 1 and include a PATH-SEGMENT TLV to indicate that the Path Segment is allocated by PCE and encoded in the PCEP message towards PCC. Further, a PCE would set this bit to 0 and include a PATH-SEGMENT TLV in the LSP object to indicate that the Path Segment should be allocated by the PCC as described in Section 5.1.1.




Figure 2: P-flag in LSP Object

The figure is included for the ease of the reader and will be removed at the time of publication.

4.2.1. Path Segment TLV

The PATH-SEGMENT TLV is an optional TLV for use in the LSP Object for Path Segment allocation. The type of this TLV is to be allocated by IANA (TBA4). The format is as shown below.




Figure 3: The PATH-SEGMENT TLV Format

The type (16-bit) of the TLV is TBA4 (to be allocated by IANA). The length (16-bit) has a variable length. The value contains the following fields:

- o ST (The Segment type - 8 bits): The ST field specifies the type of the Path Segment field, which carries a Path Segment corresponding to the SR path.
 - * 0: MPLS Path Segment, which is an MPLS label as defined in [I-D.ietf-spring-mpls-path-segment]. The PST type MUST be set to SR (MPLS).
 - * 1-255: Reserved for future use.
- o Flags (8 bits): One flag is currently defined:

- * L-Bit (Local/Global - 1 bit): If set, then the Path Segment carried by the PATH-SEGMENT TLV has local significance. If not set, then the Path Segment carried by this TLV has global significance (i.e. Path Segment is global within an SR domain).
- * The unassigned bits MUST be set to 0 and MUST be ignored at receipt.
- o Reserved (16 bits): MUST be set to 0 and MUST be ignored at receipt.
- o Path Segment: The Path Segment of an SR path. The Path Segment type is indicated by the ST field. When the ST is 0, it is a MPLS Path Segment [I-D.ietf-spring-mpls-path-segment] in the MPLS label format.

In general, only one instance of PATH-SEGMENT TLV will be included in LSP object. If more than one PATH-SEGMENT TLV is included, the first one is processed and others MUST be ignored. Multiple Path Segment allocation for use cases like alternate-making will be considered in future version of this draft.

When the Path Segment allocation is enabled, a PATH-SEGMENT TLV MUST be included in the LSP object.

If the label space is maintained by PCC itself, and the Path Segment is allocated by Egress PCC, then the PCE should request the Path Segment from Egress PCC as described in Section 5.1.1. In this case, the PCE should send a PCUpdate or PCInitiate message to the egress PCC to request the Path Segment. The P-flag in LSP should be unset in this case.

If a PCEP node does not recognize the PATH-SEGMENT TLV, it would behave in accordance with [RFC5440] and ignore the TLV. If a PCEP node recognizes the TLV but does not support the TLV, it MUST send PCErr with Error-Type = 2 (Capability not supported).

4.3. FEC Object

The FEC Object [I-D.zhao-pce-pcep-extension-pce-controller-sr] is used to specify the FEC information and carried within PCInitiate or PCRpt message for the PCECC-SR operations. The PCE MUST inform the Path Identification information to the Egress PCC. To do this, this document extends the procedures of [I-D.zhao-pce-pcep-extension-pce-controller-sr] by defining a new FEC object type for Path.

FEC Object-Type is TBA6 'Path'.




Figure 4: The path FEC object Format

One or more following TLV(s) are allowed in the 'path' FEC object -

- o SYMBOLIC-PATH-NAME TLV: As defined in [RFC8231], it is a human-readable string that identifies an LSP in the network.
- o LSP-IDENTIFIERS TLVs: As defined in [RFC8231], it is optional for SR, but could be used to encode the source, destination and other identification information for the path.
- o SPEAKER-ENTITY-ID TLV: As defined in [RFC8232], a unique identifier for the PCEP speaker, it is used to identify the Ingress PCC.

Either SYMBOLIC-PATH-NAME TLV or LSP-IDENTIFIERS TLV MUST be included. SPEAKER-ENTITY-ID TLV is optional. Only one instance of each TLV is processed, if more than one TLV of each type is included, the first one is processed and others MUST be ignored.

4.4. CCI Object

The Central Control Instructions (CCI) Object is used by the PCE to specify the forwarding instructions is defined in [I-D.ietf-pce-pcep-extension-for-pce-controller]. Further [I-D.zhao-pce-pcep-extension-pce-controller-sr] defined a CCI object type for SR.

The Path Segment information is encoded directly in the CCI SR object. The Path Segment TLV as described in the Section 4.2.1, MUST also be included in the CCI SR object as the TLV (as it includes additional information regarding the Path Segment identifier). The C flag in CCI object is used to indicate if the allocation needs to be done by the PCC.

5. Operations

The Path Segment allocation and encoding is as per the Stateful PCE operations for segment routing. The procedures are as per the corresponding extensions defined in [I-D.ietf-pce-segment-routing] and [I-D.ietf-pce-segment-routing-ipv6] (which are further based on [RFC8231] and [RFC8281]). The additional operations for Path Segment are defined in this section.

To notify (or request) the Path Segment to the Egress PCC, the procedures are as per the PCECC-SR [I-D.zhao-pce-pcep-extension-pce-controller-sr] (which is based on [I-D.ietf-pce-pcep-extension-for-pce-controller]). The additional operations are defined in this section.

5.1. Stateful PCE Operation

As defined in [I-D.ietf-spring-mpls-path-segment], a Path Segment can be allocated by the egress PCC. In this case, the label space is maintained on the PCC itself.

This section describes the mechanism of Path Segment allocation by using PCInitiate and PCUpd message in Stateful PCE model.

5.1.1. Ingress PCC-Initiated Path Segment Allocation

The ingress PCC could request the Path Segment to be allocated by the PCE via PCRpt message. The delegate flag (D-flag) MUST also be set for this LSP. Also, the P-flag in the LSP object MUST be set.

On receiving a delegation request with Path Segment allocation request from an ingress PCC, a stateful PCE requests the egress PCC to allocate a Path Segment.

The PATH-SEGMENT TLV MUST be included in an LSP object in the PCInitiate message sent from the PCE to the egress to request Path Segment allocation by the egress PCC. The P flag in LSP object MUST be set to 0. This PCInitiate message to egress PCC would be the similar to the one sent to ingress PCC as per [I-D.ietf-pce-segment-routing], but the egress PCC would only allocate the Path Segment and would not trigger the LSP initiation operation (as it would be the egress for this LSP).

If the value of Path Segment is 0x0, it indicates that the PCE is requesting a Path Segment for this LSP. If the Path Segment is set to a value 'n' and the P flag is unset in the LSP object, it indicates that the PCE requests a specific value 'n' of Path Segment. If the Path Segment is allocated successfully, the egress PCC reports

the Path Segment via PCRpt message with PATH-SEGMENT TLV in LSP object. Else, it MUST send a PCErr message with Error-Type = TBA7 ("Path SID failure") and Error Value = 1 ("Invalid SID"). If the value of Path Segment is valid, but the PCC is unable to allocate the Path Segment, it MUST send a PCErr message with Error-Type = TBA7 ("Path SID failure") and Error Value = 2 ("Unable to allocate the specified label/SID").

Once the PCE receives the PCRpt message, it can obtain the Path Segment information from the egress PCC and then update the path with Path Segment by sending PCUpd message to the ingress PCC.

If the Path Segment is updated successfully, the ingress PCC will acknowledge with a PCRpt message to the PCE. In case of error, an PCErr message with Error-Type = TBA7 ("Path SID failure") and Error Value = 1 ("Invalid SID") will be sent back to the PCE. The PCE MUST roll back the Path Segment value to the previous value (if any) by sending a PCUpd message to synchronize with the egress PCC.




Figure 5: Ingress PCC-Initiated Path Segment Allocation

If the ingress PCC wishes to withdraw or modify a previously reported Path Segment value, it MUST send a PCRpt message without any PATH-

SEGMENT TLV or with the PATH-SEGMENT TLV containing the new Path Segment respectively. In this case, the PCE should synchronize with egress PCC via PCUpd message.

The Path Segment MUST be withdrawn when the corresponding LSP is removed. When the LSP is deleted, the PCE MUST request the egress PCC to withdraw the LSP and associated Path Segment via PCInitiate message with the R flag is set in the SRP object.

If an egress PCC receives a valid Path Segment value from a PCE which is different than the current Path Segment, it MUST try to allocate the new value. If the new Path Segment is successfully allocated, the egress PCC MUST report the new value to the PCE. Otherwise, it MUST send a PCErr message with Error-Type = TBA7 ("Path label/SID failure") and Error Value = 2 ("Unable to allocate the specified label/SID").

5.1.2. PCE Initiated Path Segment Allocation

A stateful PCE also can initiate or update an LSP with Path Segment actively via requesting the egress PCC to allocate a Path Segment.

If a PCE wishes to modify a previously requested Path Segment value or allocate a Path Segment for an PCE-Initiated LSP, it MUST request the egress PCC to allocate a new value by sending a PCUpd message to the egress PCC with PATH-SEGMENT TLV containing the new Path Segment value. Also, the P flag in LSP object is unset. Absence of the PATH-SEGMENT TLV in PCUpd message means that the PCE wishes to withdraw the Path Segment.

The mechanism of requesting Path Segment is as per Section 5.1.1.

Once the PCE receives the PCRpt message, it can obtain the Path Segment information from the egress PCC and then update or initiate an LSP with Path Segment.

If the SR-Path is setup, the ingress PCC will acknowledge with a PCRpt message to the PCE. In case of error, as described in [I-D.ietf-pce-segment-routing], a PCErr message will be sent back to the PCE. The PCE MUST request the egress PCC to withdraw the LSP and associated Path Segment via PCInitiate message with the R flag is set in the SRP object.

If the Path Segment is updated successfully, the ingress PCC will acknowledge with a PCRpt message to the PCE. In case of error, a PCErr message with Error-Type = TBA7 ("Path SID failure") and Error Value = 1 ("Invalid SID") will be sent back to the PCE. The PCE MUST

roll back the Path Segment value to the previous value (if any) by sending a PCUpd message to synchronize with the egress PCC.




Figure 6: Stateful PCE-Initiated Path Segment Allocation

5.2. PCECC Based Operation

5.2.1. PCE Controlled Label Spaces Advertisement

For allocating the Path Segments to SR paths by the PCEs, the PCE controlled label space MUST be known at PCEs via configurations or any other mechanisms. The PCE controlled label spaces MAY be advertised as described in [I-D.li-pce-controlled-id-space].

5.2.2. PCECC based Path Segment Allocation

5.2.2.1. PCECC-Initiated

The PCE could allocate the Path Segment on its own for a PCE-Initiated (or delegated LSP). The allocated Path Segment needs to be informed to the Ingress and Egress PCC. The PCE would use the PCInitiate message [RFC8281] or PCUpd message [RFC8231] towards the Ingress PCC and MUST include the PATH-SEGMENT TLV in the LSP object.

The PCE would further inform the egress PCC about the Path Segment allocated by the PCE using the PCInitiate message as described in [I-D.zhao-pce-pcep-extension-pce-controller-sr].




Figure 7: PCE allocated Path Segment on its own

5.2.2.2. Ingress PCC-Initiated PCECC

The ingress PCC could request the Path Segment to be allocated by the PCE via PCRpt message as per [RFC8231]. The delegate flag (D-flag) MUST also be set for this LSP. Also, the P-flag in the LSP object MUST be set.

A PATH-SEGMENT TLV MUST be included in the LSP object. If the value of Path Segment is 0x0, it indicates that the Ingress PCC is requesting a Path Segment for this LSP. If the Path Segment is set to a value 'n', it indicates that the ingress PCC requests a specific value 'n' of Path Segment.

If the Path Segment is allocated successfully, the PCE would further respond to Ingress PCC with PCUpd message as per [RFC8231] and MUST include the PATH-SEGMENT TLV in a LSP object. Else, it MUST send a PCErr message with Error-Type = TBA7 ("Path SID failure") and Error Value = 1 ("Invalid SID"). If the value of Path Segment is valid, but the PCC is unable to allocate the Path Segment, it MUST send a PCErr message with Error-Type = TBA7 ("Path SID failure") and Error Value = 2 ("Unable to allocate the specified label/SID").

The active PCE would allocate the Path Segment as per the PATH-SEGMENT flags and in case PATH-SEGMENT is not included, the PCE MUST act based on the local policy.

The PCE would further inform the egress PCC about the Path Segment allocated by the PCE using the PCInitiate message as described in [I-D.zhao-pce-pcep-extension-pce-controller-sr].




Figure 8: Ingress PCC request Path Segment to PCE

6. Dataplane Considerations

As described in [I-D.ietf-spring-mpls-path-segment], in an SR-MPLS network, when a packet is transmitted along an SR path, the labels in the MPLS label stack will be swapped or popped. So that no label or only the last label may be left in the MPLS label stack when the packet reaches the egress node. Thus, the egress node cannot determine from which SR path the packet comes. For this reason, it introduces the Path Segment.

Apart from allocation and encoding of the Path Segment (described in this document) for the LSP, it would also be included in the SID/Label stack of the LSP (usually for processing by the egress). To support this, the Path Segment MAY also be a part of SR-ERO as prepared by the PCE as per [I-D.ietf-pce-segment-routing]. The PCC MAY also include the Path Segment while preparing the label stack based on the local policy and use-case.

It is important that the PCE learns the Maximum SID Depth (MSD) that can be imposed at each node/link of a given SR path to ensure that the SID stack depth does not exceed the number of SIDs the node is

capable of imposing. As a new type of segment, Path Segment will be inserted in the SID list just like other SIDs. Thus, the PCE needs to consider the affect of Path Segment when computing a LSP with Path Segment allocation.

7. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC7942].

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC7942], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

7.1. Huawei's Commercial Delivery

The feature is developing based on Huawei VRP8.

- o Organization: Huawei
- o Implementation: Huawei's Commercial Delivery implementation based on VRP8.
- o Description: The implementation is under development and follows the mechanism as defined in section-5.1.1.
- o Maturity Level: Product
- o Contact: tanren@huawei.com

7.2. ZTE's Commercial Delivery

- o Organization: ZTE
- o Implementation: ZTE's Commercial Delivery implementation based on Rosng v8.
- o Description: The implementation is under development and follows the mechanism as defined in section-5.1.1.
- o Maturity Level: Product
- o Contact: zhan.shuangping@zte.com.cn

8. IANA Considerations

8.1. SR PCE Capability Flags

SR PCE Capability TLV is defined in [I-D.ietf-pce-segment-routing], and the registry to manage the Flag field of the SR PCE Capability TLV is requested in [I-D.ietf-pce-segment-routing]. IANA is requested to make the following allocation in the "SR Capability Flag Field" sub-registry.

Bit	Description	Reference
TBA1	Path Segment Allocation is supported(P)	This document

8.2. New LSP Flag Registry

[RFC8231] defines the LSP object; per that RFC, IANA created a registry to manage the value of the LSP object's Flag field. IANA has allocated a new bit in the "LSP Object Flag Field" sub-registry, as follows:

Bit	Description	Reference
TBA3	Request for Path Segment Allocation(P)	This document

8.3. New PCEP TLV

IANA is requested to add the assignment of a new allocation in the existing "PCEP TLV Type Indicators" sub-registry as follows:

Value	Description	Reference
TBA4	PATH-SEGMENT TLV	This document

8.3.1. Path Segment TLV

This document requests that a new sub-registry named "PATH-SEGMENT TLV Segment Type (ST) Field" to be created to manage the value of the ST field in the PATH-SEGMENT TLV.

Value	Description	Reference
0	MPLS Path Segment (MPLS label)	This document
1-255	Reserved for future use	This document

Further, this document also requests that a new sub-registry named "PATH-SEGMENT TLV Flag Field" to be created to manage the Flag field in the PATH-SEGMENT TLV. New values are 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

Bit	Description	Reference
7	Local Signification(L)	This document

8.4. New FEC Type Registry

A new PCEP object called FEC is defined in [I-D.zhao-pce-pcep-extension-pce-controller-sr]. IANA is requested to allocate a new Object-Type for FEC object in the "PCEP Objects" sub-registry.

Value	Description	Reference
TBA6	Path	This document

8.5. PCEP Error Type and Value

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

Error-Type	Meaning	Reference
TBA7	Path SID failure: Error-value = 1 Invalid SID	This document
	Error-value = 2 Unable to allocate Path SID	

9. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8281] and [I-D.ietf-pce-segment-routing] are applicable to this specification. No additional security measure is required.

As described [I-D.ietf-pce-segment-routing] and [I-D.ietf-pce-pcep-extension-for-pce-controller], SR allows a network controller to instantiate and control paths in the network. A rogue PCE can manipulate Path SID allocations to have impact based on the usage of Path SID such as accounting, bi-directional etc.

Thus, as per [RFC8231], it is RECOMMENDED that these PCEP extensions only be activated on authenticated and encrypted sessions across PCEs and PCCs belonging to the same administrative authority, using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525] (unless explicitly set aside in [RFC8253]).

10. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231], and [I-D.ietf-pce-segment-routing] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

10.1. Control of Function and Policy

A PCEP implementation SHOULD allow the operator to configure the policy based on which it allocates the Path SID. This includes the Path SID scope.

10.2. Information and Data Models

The PCEP YANG module is defined in [I-D.ietf-pce-pcep-yang]. In future, this YANG module should be extended or augmented to provide the following additional information relating to Path SID.

An implementation SHOULD allow the operator to view the Path SID allocated to the LSP as well as Path SID as part of the computed SID list for the SR path.

10.3. 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].

10.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440], [RFC8231], and [I-D.ietf-pce-segment-routing] .

10.5. Requirements On Other Protocols

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

10.6. Impact On Network Operations

Mechanisms defined in [RFC5440], [RFC8231], and [I-D.ietf-pce-segment-routing] also apply to PCEP extensions defined in this document. Further, the mechanism described in this document can help the operator to request control of the LSPs at a particular PCE.

11. Acknowledgments

TBA

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [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>>.
- [RFC8232] Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", RFC 8232, DOI 10.17487/RFC8232, September 2017, <<https://www.rfc-editor.org/info/rfc8232>>.
- [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>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [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>>.

- [I-D.ietf-pce-segment-routing]
Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W.,
and J. Hardwick, "PCEP Extensions for Segment Routing",
draft-ietf-pce-segment-routing-16 (work in progress),
March 2019.
- [I-D.ietf-pce-segment-routing-ipv6]
Negi, M., Li, C., Sivabalan, S., Kaladharan, P., and Y.
Zhu, "PCEP Extensions for Segment Routing leveraging the
IPv6 data plane", draft-ietf-pce-segment-routing-ipv6-02
(work in progress), April 2019.
- [I-D.zhao-pce-pcep-extension-pce-controller-sr]
Zhao, Q., Li, Z., Negi, M., and C. Zhou, "PCEP Procedures
and Protocol Extensions for Using PCE as a Central
Controller (PCECC) of SR-LSPs", draft-zhao-pce-pcep-
extension-pce-controller-sr-05 (work in progress), July
2019.
- [I-D.ietf-pce-pcep-extension-for-pce-controller]
Zhao, Q., Li, Z., Negi, M., and C. Zhou, "PCEP Procedures
and Protocol Extensions for Using PCE as a Central
Controller (PCECC) of LSPs", draft-ietf-pce-pcep-
extension-for-pce-controller-02 (work in progress), July
2019.
- [I-D.li-spring-srv6-path-segment]
Li, C., Cheng, W., Chen, M., Dhody, D., Li, Z., Dong, J.,
and R. Gandhi, "Path Segment for SRv6 (Segment Routing in
IPv6)", draft-li-spring-srv6-path-segment-03 (work in
progress), August 2019.

12.2. Informative References

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation
Element (PCE)-Based Architecture", RFC 4655,
DOI 10.17487/RFC4655, August 2006,
<<https://www.rfc-editor.org/info/rfc4655>>.
- [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>>.
- [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>>.

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.

[I-D.li-pce-controlled-id-space]

Li, C., Chen, M., Dong, J., Li, Z., Wang, A., Cheng, W., and C. Zhou, "PCE Controlled ID Space", draft-li-pce-controlled-id-space-03 (work in progress), June 2019.

[I-D.ietf-spring-mpls-path-segment]

Cheng, W., Li, H., Chen, M., Gandhi, R., and R. Zigler, "Path Segment in MPLS Based Segment Routing Network", draft-ietf-spring-mpls-path-segment-00 (work in progress), March 2019.

[I-D.ietf-pce-pcep-yang]

Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", draft-ietf-pce-pcep-yang-12 (work in progress), July 2019.

Appendix A. Contributors

The following people have substantially contributed to this document:

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

Email: dhruv.ietf@gmail.com

Zafar Ali
Cisco Systems, Inc.

Email: zali@cisco.com

Appendix B. SRv6 extensions

This section would be rolled into the document once the SPRING WG adopts SRv6 path segment.

B.1. The SRv6 PCE Capability sub-TLV

[I-D.ietf-pce-segment-routing-ipv6] defined a new Path Setup Type (PST) and SRv6-PCE-CAPABILITY sub-TLV for SRv6. PCEP speakers use this sub-TLV to exchange information about their SRv6 capability. The TLV includes a Flags field and one bit (L-flag) was allocated in [I-D.ietf-pce-segment-routing-ipv6].

This document adds an additional flag for Path Segment allocation, as follows -

- o P (Path Segment Identification bit): A PCEP speaker sets this flag to 1 to indicate that it has the capability to encode SRv6 path identification. (Path Segment, as per [I-D.li-spring-srv6-path-segment]).




Figure 9: P-flag in SRv6-PCE-CAPABILITY TLV

The figure is included for the ease of the reader and can be removed at the time of publication.

B.2. SRv6 PCE Capability Flags

SRv6 PCE Capability TLV is defined in defined in [I-D.ietf-pce-segment-routing-ipv6], and the registry to manage the Flag field of the SRv6 PCE Capability Flags is requested in [I-D.ietf-pce-segment-routing-ipv6]. IANA is requested to make the following allocation in the aforementioned registry.

Bit	Description	Reference
TBA2	Path Segment Allocation is supported(P)	This document

B.3. Path Segment TLV

A new assignment should be done to the "PATH-SEGMENT TLV Segment Type (ST) Field" sub-registry for SRv6.

Value	Description	Reference
1	SRv6 Path Segment (IPv6 addr)	This document
2-255	Reserved for future use	This document

Authors' Addresses

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

Email: chengli13@huawei.com

Mach(Guoyi) Chen
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: Mach.chen@huawei.com

Weiqiang Cheng
China Mobile
China

Email: chengweiqiang@chinamobile.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Rakesh Gandhi
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

Quan Xiong
ZTE Corporation
China

Email: xiong.quan@zte.com.cn