

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2019

C. Li
M. Chen
Huawei Technologies
W. Cheng
China Mobile
Z. Li
J. Dong
Huawei Technologies
R. Gandhi
Cisco Systems, Inc.
March 11, 2019

PCEP Extensions for Associated Bidirectional Segment Routing (SR) Paths
draft-li-pce-sr-bidir-path-04

Abstract

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Clients (PCCs) requests. The Stateful PCE extensions allow stateful control of Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Label Switched Paths (LSPs) using PCEP. Furthermore, PCEP can be used for computing paths in Segment Routing (SR) TE networks.

This document defines PCEP extensions for grouping two reverse unidirectional SR Paths into an Associated Bidirectional SR Path when using a Stateful PCE for both PCE-Initiated and PCC-Initiated LSPs as well as when using a Stateless PCE.

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 September 9, 2019.

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. PCEP Extension for Bidirectional SR Path 4
3.1. Double-sided Bidirectional SR Path Association Group Object 5
4. Bidirectional Flag 6
5. Procedures for Associated Bidirectional SR Path Computation . 6
5.1. PCE Initiated Associated Bidirectional SR Paths 7
5.2. PCC Initiated Associated Bidirectional SR Paths 8
5.3. Error Handling 9
6. IANA Considerations 10
6.1. Association Type 10
6.2. PCEP Errors 10
7. Security Considerations 10
8. Contributors 11
9. Acknowledgments 11
10. References 12
10.1. Normative References 12
10.2. Informative References 13
Authors' Addresses 14

1. Introduction

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

[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-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.ietf-pce-segment-routing-ipv6] extend PCEP to support SR for IPv6 data plane.

[I-D.ietf-pce-association-group] introduces a generic mechanism to create a grouping of LSPs which can then be used to define associations between a set of LSPs and/or a set of attributes, for example primary and secondary LSP associations, and is equally applicable to the active and passive modes of a Stateful PCE [RFC8231] or a stateless PCE [RFC5440].

Currently, SR networks only support unidirectional paths. However, bidirectional SR Paths are required in some networks, for example, in mobile backhaul transport networks. The requirement of bidirectional SR Path is specified in [I-D.ietf-spring-mpls-path-segment].

[I-D.ietf-pce-association-bidir] defines PCEP extensions for grouping two reverse unidirectional MPLS TE LSPs into an Associated Bidirectional LSP when using a Stateful PCE for both PCE-Initiated and PCC-Initiated LSPs as well as when using a Stateless PCE.

This document extends the bidirectional association to segment routing by specifying PCEP extensions for grouping two reverse unidirectional SR Paths into a bidirectional SR Path.

[I-D.ietf-pce-association-bidir] specifies the Double-sided Bidirectional LSP Association procedure, where the PCE creates the association and provisions at both endpoints, the RSVP-TE does the signaling to the egress the status of the forward LSP and the ingress about the reverse LSP. Thus, the both endpoints learn the reverse LSPs forming the bidirectional LSP association. In case of SR, to support the bidirectional path use-case, this is done using the PCEP protocol. This is done so that both endpoints are aware of the the unidirectional SR Path, as well as the status and other SR path related information.

[I-D.li-pce-sr-path-segment] defines a procedure for Path Segment Identifier (PSID) in PCEP for SR using PATH-SEGMENT TLV. The PSID can be a Path Segment Identifier in SR-MPLS [I-D.ietf-spring-mpls-path-segment], or a Path Segment Identifier in SRv6 [I-D.li-spring-srv6-path-segment]. The PSID can be used for an associated bidirectional SR Path for identifying the SR Path.

2. Terminology

This document makes use of the terms defined in [I-D.ietf-pce-segment-routing]. The reader is assumed to be familiar with the terminology defined in [RFC5440], [RFC8231], [RFC8281], [I-D.ietf-pce-association-group] and [I-D.ietf-pce-association-bidir].

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. PCEP Extension for Bidirectional SR Path

As per [I-D.ietf-pce-association-group], LSPs are associated by adding them to a common association group.

[I-D.ietf-pce-association-bidir] specifies PCEP extensions for grouping two reverse unidirectional MPLS-TE LSPs into an Associated Bidirectional LSP for both single-sided and double-sided initiation cases by defining two new Bidirectional LSP Association Groups.

This document extends the procedure for associated bidirectional SR Paths by defining a new bidirectional association group (Double-sided Bidirectional SR Path Association Group). The document further

describes the mechanism for associating two unidirectional SR Paths into a bidirectional SR Path. [I-D.li-pce-sr-path-segment] defines a procedure for communicating Path Segment in PCEP for SR using PATH-SEGMENT TLV. The bidirectional SR Path can also use the PATH-SEGMENT TLV.

Note that an association group is defined in this document to define procedures specific to SR Paths (and the procedures are different than the RSVP-TE bidirectional association groups defined in [I-D.ietf-pce-association-bidir]).

3.1. Double-sided Bidirectional SR Path Association Group Object

As defined in [I-D.ietf-pce-association-bidir], two LSPs are associated as a bidirectional MPLS-TE LSP by a common bidirectional LSP association group. For associating two SR paths, this document defines a new association group called 'Double-sided Bidirectional SR Path Association Group' as follows:

- o Association Type (TBD1 to be assigned by IANA) = Double-sided Bidirectional SR Path Association Group

Similar to other bidirectional associations, this Association Type is operator-configured in nature and statically created by the operator on the PCEP peers. The paths belonging to this association is conveyed via PCEP messages to the PCEP peer. Operator-configured Association Range TLV [I-D.ietf-pce-association-group] MUST NOT be sent for these Association Types, and MUST be ignored, so that the entire range of association ID can be used for them. The handling of the Association ID, Association Source, optional Global Association Source and optional Extended Association ID in this association are set in the same way as [I-D.ietf-pce-association-bidir].

A member of the 'Double-sided Bidirectional SR Path Association Group' can take the role of a forward or reverse SR Path and follow the similar rules defined in [I-D.ietf-pce-association-bidir] for LSPs.

- o An SR Path (forward or reverse) can not be part of more than one 'Double-sided Bidirectional SR Path Association Group'.
- o The endpoints of the SR Paths in this associations cannot be different.

For describing the SR Paths in this association group, such as direction and co-routed information, this association group reuses the Bidirectional LSP Association Group TLV defined in [I-D.ietf-pce-association-bidir]. All fields and processing rules

are as per [I-D.ietf-pce-association-bidir].

4. Bidirectional Flag

As defined in [RFC5440], the B-flag in RP object MUST be set when the PCC specifies that the path computation request relates to a bidirectional TE LSP. In this document, the B-flag also MUST be set when the PCC specifies that the path computation request relates to a bidirectional SR Path. When a stateful PCE initiates or updates a bidirectional SR Paths including LSPs and SR paths, the B-flag in SRP object [I-D.ietf-pce-pcep-stateful-pce-gmpls] MAY be set as well.

5. Procedures for Associated Bidirectional SR Path Computation

Two unidirectional SR Paths can be associated by the association group object as specified in [I-D.ietf-pce-association-group]. A bidirectional LSP association group object is defined in [I-D.ietf-pce-association-bidir] (for MPLS-TE). This document extends these association mechanisms for bidirectional SR Paths. Two SR Paths can be associated together by using the Bidirectional SR Path Association Group defined in this document for PCEP messages. The PATH-SEGMENT TLV [I-D.li-pce-sr-path-segment] SHOULD also be included in the LSP object for these SR Paths to support required use-cases.

For bidirectional SR Paths, there is a need to know the reverse direction SR paths. The PCE SHOULD inform the reverse SR Paths to the ingress PCCs and vice versa. To achieve this, a PCInitiate message for the reverse SR Path is sent to the ingress PCC and a PCInitiate message for the forward SR Path is sent to the egress PCC (with the same association group). These PCInitiate message MUST NOT trigger initiation of SR Paths. The reverse direction SR Path can be used for several use-cases, such as directed BFD [I-D.ietf-mpls-bfd-directed].

For a bidirectional LSP computation when using both direction LSPs on a node, the same LSP would need to be identified using 2 different PLSP-IDs based on the PCEP session to the ingress or the egress. In other words, the LSP will have a PLSP-ID A at the ingress node while it will have the PLSP-ID B at the egress node. The PCE will maintain the two PLSP-IDs for the same LSP. For instance, an ingress PCC requests a bidirectional SR Path computation, and the PCE computes a forward LSP1 with PLSP-ID say 100. The reverse LSP2 from the egress to the ingress with PLSP-ID say 200 is allocated by the egress PCC. Since the PLSP-ID space is independent at each PCC, the PLSP-ID allocated by the egress PCC can not be used for the LSP at the

ingress PCC (PLSP-ID conflict may occur). Hence, the PCE needs to allocate a PLSP-ID for LSP2 from the ingress PCC's PLSP-ID space, say 101. Similarly for LSP1, it has PLSP-ID 100 at the ingress, and may have say PLSP-ID 201 at the egress node.

5.1. PCE Initiated Associated Bidirectional SR Paths

As specified in [I-D.ietf-pce-association-group], Bidirectional SR Path Association Group can be created by a Stateful PCE.

- o Stateful PCE can create and update the forward and reverse SR Paths independently for a 'Double-sided Bidirectional SR Path Association Group'.
- o Stateful PCE can establish and remove the association relationship on a per SR Path basis.
- o Stateful PCE can create and update the SR Path and the association on a PCC via PCInitiate and PCUpd messages, respectively, using the procedures described in [I-D.ietf-pce-association-group].
- o The PATH-SEGMENT TLV SHOULD be included for each SR Path in the LSP object.
- o The reverse direction SR Path (LSP2(R) at node S, LSP1(R) at node D) SHOULD be informed by PCE via PCInitiate message with the matching association group.

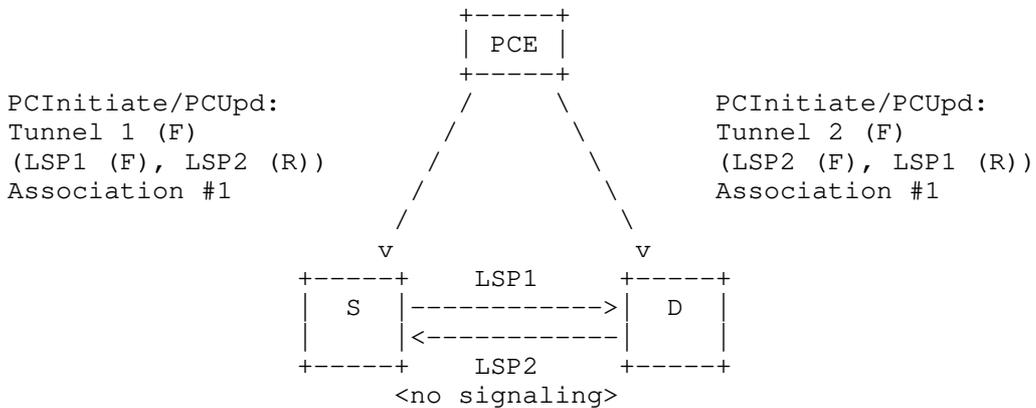


Figure 1: PCE-Initiated Double-sided Bidirectional SR Path with Forward and Reverse Direction SR Paths

5.2. PCC Initiated Associated Bidirectional SR Paths

As specified in [I-D.ietf-pce-association-group], Bidirectional SR Path Association Group can also be created by a PCC.

- o PCC can create and update the forward and reverse SR Paths independently for a 'Double-sided Bidirectional SR Path Association Group'.
- o PCC can establish and remove the association relationship on a per SR Path basis.
- o PCC MUST report the change in the association group of an SR Path to PCE(s) via PCRpt message.
- o PCC can report the forward and reverse SR Paths independently to PCE(s) via PCRpt message.
- o PCC can delegate the forward and reverse SR Paths independently to a Stateful PCE, where PCE would control the SR Paths.
- o Stateful PCE can update the SR Paths in the 'Double-sided Bidirectional SR Path Association Group' via PCUpd message, using the procedures described in [I-D.ietf-pce-association-group].
- o The PATH-SEGMENT TLV MUST be handled as defined in [I-D.li-pce-sr-path-segment].
- o The reverse direction SR Path (LSP2(R) at node S, LSP1(R) at node D) SHOULD be informed by PCE via PCInitiate message with the matching association group.

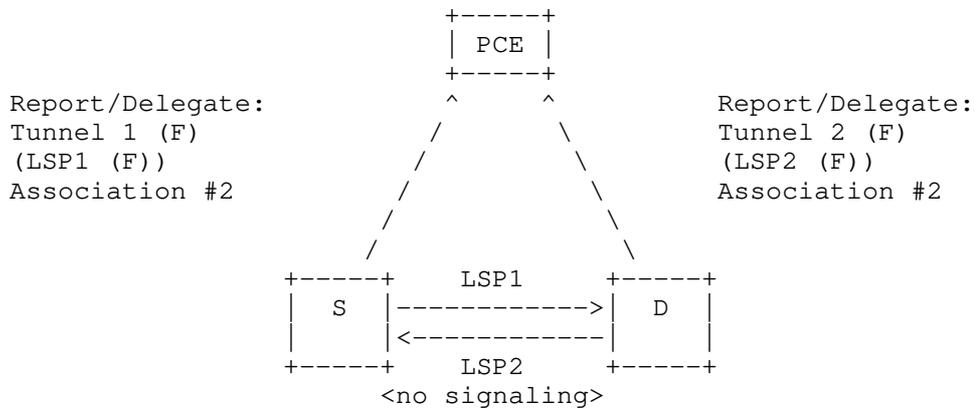


Figure 2a: Step 1: PCC-Initiated Double-sided Bidirectional SR Path with Forward Direction SR Paths

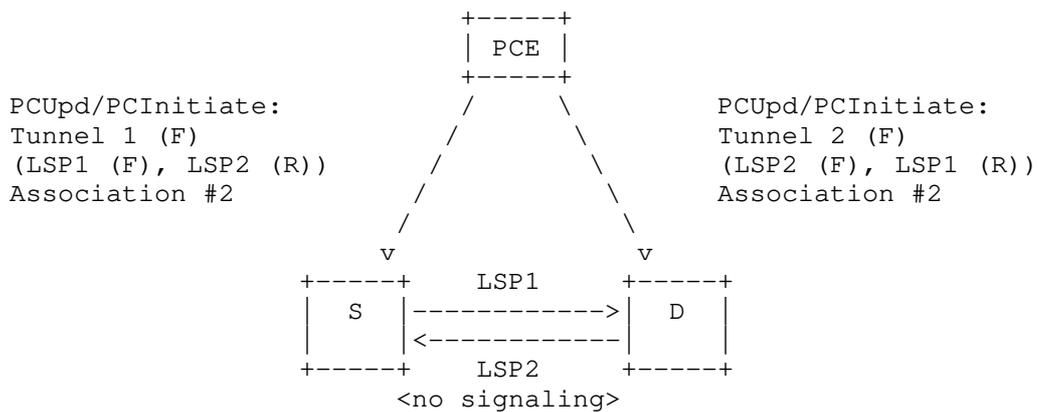


Figure 2b: Step 2: PCE-Upd/Initiated Double-sided Bidirectional Path Along with Reverse Direction SR Paths

5.3. Error Handling

The error handling as described in section 5.5 of [I-D.ietf-pce-association-bidir] continue to apply.

The PCEP Path Setup Type (PST) MUST be set to 'TE Path is Setup using Segment Routing' [I-D.ietf-pce-segment-routing] for the LSP belonging to the 'Double-sided Bidirectional SR Path Association Group'. In case a PCEP speaker receives a different PST value for this association group, it MUST send a PCErr message with Error-Type = 29

(Early allocation by IANA) (Association Error) and Error-Value = TBD2 (Bidirectional LSP Association - Path Setup Type Mismatch).

6. IANA Considerations

6.1. Association Type

This document defines a new Association Type for the Association Object defined [I-D.ietf-pce-association-group]. IANA is requested to make the assignment of a value for the sub-registry "ASSOCIATION Type Field" (to be created in [I-D.ietf-pce-association-group]), as follows:

Value	Name	Reference
TBD1	Double-sided Bidirectional SR Path Association Group	This document

6.2. PCEP Errors

This document defines new Error value for Error Type 29 (Association Error). IANA is requested to allocate new Error value within the "PCEP-ERROR Object Error Types and Values" sub-registry of the PCEP Numbers registry, as follows:

Error Type	Description	Reference
29	Association Error	
	Error value: TBD2 Bidirectional LSP Association - Path Setup Type Mismatch	This document

7. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8281], and [I-D.ietf-pce-segment-routing] apply to the extensions defined in this document as well.

A new Association Type for the Association Object, 'Double-sided Associated Bidirectional SR Path Association Group' is introduced in this document. Additional security considerations related to LSP associations due to a malicious PCEP speaker is described in [I-D.ietf-pce-association-group] and apply to this Association Type. Hence, securing the PCEP session using Transport Layer Security (TLS)

[RFC8253] is recommended.

8. 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

Quan Xiong
ZTE Corporation
No.6 Huashi Park Rd
Wuhan, Hubei 430223
China

Email: xiong.quan@zte.com.cn

9. Acknowledgments

Many thanks to Marina Fitzgeer for detailed review and comments.

10. References

10.1. 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>>.
- [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>>.
- [I-D.ietf-pce-association-group]
Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "PCEP Extensions for Establishing Relationships Between Sets of LSPs", draft-ietf-pce-association-group-08 (work in progress), March 2019.
- [I-D.ietf-pce-association-bidir]
Gandhi, R., Barth, C., and B. Wen, "PCEP Extensions for Associated Bidirectional Label Switched Paths (LSPs)", draft-ietf-pce-association-bidir (work in progress).
- [I-D.ietf-pce-pcep-stateful-pce-gmpls]
Lee, Y., Zhang, F., Casellas, R., Dios, O., and Z. Ali, "Path Computation Element (PCE) Protocol Extensions for Stateful PCE Usage in GMPLS-controlled Networks", draft-ietf-pce-pcep-stateful-pce-gmpls-10 (work in progress), March 2019.

- [I-D.li-pce-sr-path-segment]
Li, C., Chen, M., Dhody, D., Cheng, W., Dong, J., Li, Z.,
and R. Gandhi, "Path Computation Element Communication
Protocol (PCEP) Extension for Path Segment in Segment
Routing (SR)", draft-li-pce-sr-path-segment (work in
progress).
- [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 (work in progress).
- [I-D.li-spring-srv6-path-segment]
Li, C., 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-00 (work in progress),
October 2018.

10.2. 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>>.
- [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>>.
- [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>>.
- [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., and P. Kaladharan, "PCEP
Extensions for Segment Routing leveraging the IPv6 data
plane", draft-ietf-pce-segment-routing-ipv6 (work in
progress).

[I-D.ietf-mpls-bfd-directed]

Mirsky, G., Tantsura, J., Varlashkin, I., and M. Chen,
"Bidirectional Forwarding Detection (BFD) Directed Return
Path", draft-ietf-mpls-bfd-directed-10 (work in progress),
September 2018.

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

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

Email: lizhenbin@huawei.com

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

Email: jie.dong@huawei.com

Rakesh Gandhi
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com