

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 29, 2018

C. Barth
Juniper Networks
R. Gandhi, Ed.
Cisco Systems, Inc.
B. Wen
Comcast
April 27, 2018

**PCEP Extensions for
Associated Bidirectional Label Switched Paths (LSPs)
draft-ietf-pce-association-bidir-00**

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.

This document 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.

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 <http://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."

Copyright Notice

Copyright (c) 2018 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 (<http://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.	Conventions Used in This Document	4
2.1.	Key Word Definitions	4
2.2.	Terminology	4
3.	Overview	4
3.1.	Single-sided Initiation	5
3.2.	Double-sided Initiation	5
3.3.	Co-routed Associated Bidirectional LSP	6
4.	Protocol Extensions	6
4.1.	Association Object	6
4.2.	Bidirectional LSP Association Group TLV	7
5.	PCEP Procedure	8
5.1.	PCE Initiated LSPs	8
5.2.	PCC Initiated LSPs	8
5.3.	Stateless PCE	9
5.4.	State Synchronization	9
5.5.	Error Handling	9
6.	Security Considerations	9
7.	Manageability Considerations	10
7.1.	Control of Function and Policy	10
7.2.	Information and Data Models	10
7.3.	Liveness Detection and Monitoring	10
7.4.	Verify Correct Operations	10
7.5.	Requirements On Other Protocols	10
7.6.	Impact On Network Operations	10
8.	IANA Considerations	10
8.1.	Association Types	10
8.2.	Bidirectional LSP Association Group TLV	11
8.2.1.	Flag Fields in Bidirectional LSP Association Group TLV	11
8.3.	PCEP Errors	11
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	13
	Acknowledgments	15
	Authors' Addresses	15

1. Introduction

[RFC5440] describes the Path Computation Element Protocol (PCEP) as a communication mechanism between a Path Computation Client (PCC) and a Path Control Element (PCE), or between PCE and PCC, that enables computation of Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Label Switched Paths (LSPs).

[RFC8231] specifies extensions to PCEP to enable stateful control of MPLS TE LSPs. It describes two modes of operation - Passive Stateful PCE and Active Stateful PCE. In [RFC8231], the focus is on Active Stateful PCE where LSPs are provisioned on the PCC and control over them is delegated to a PCE. Further, [RFC8281] describes the setup, maintenance and teardown of PCE-Initiated LSPs for the Stateful PCE model.

[I-D.ietf-pce-association] 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].

The MPLS Transport Profile (MPLS-TP) requirements document [RFC5654] specifies that MPLS-TP MUST support associated bidirectional point-to-point LSPs. [RFC7551] specifies RSVP signaling extensions for binding two reverse unidirectional LSPs [RFC3209] into an associated bidirectional LSP. The fast reroute (FRR) procedures for associated bidirectional LSPs are described in [I-D.ietf-teas-assoc-corouted-bidir-frr].

This document 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 when using a Stateful (both active and passive modes) or Stateless PCE. The PCEP extensions cover the following cases:

- o A PCC initiates the forward and/ or reverse LSP of a single-sided or double-sided bidirectional LSP and retains the control of the LSP. The PCC computes the path itself or makes a request for path computation to a PCE. After the path setup, it reports the information and state of the path to the PCE. This includes the association group identifying the bidirectional LSP. This is the Passive Stateful mode defined in [RFC8051].
- o A PCC initiates the forward and/ or reverse LSP of a single-sided or double-sided bidirectional LSP and delegates the control of the LSP to a Stateful PCE. During delegation the association group

identifying the bidirectional LSP is included. The PCE computes the path of the LSP and updates the PCC with the information about the path as long as it controls the LSP. This is the Active Stateful mode defined in [RFC8051].

- o A PCE initiates the forward and/ or reverse LSP of a single-sided or double-sided bidirectional LSP on a PCC and retains the control of the LSP. The PCE is responsible for computing the path of the LSP and updating the PCC with the information about the path as well as the association group identifying the bidirectional LSP. This is the PCE-Initiated mode defined in [RFC8281].
- o A PCC requests co-routed or non co-routed paths for forward and reverse LSPs of a bidirectional LSP from a Stateless PCE [RFC5440].

2. Conventions Used in This Document

2.1. Key Word Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.2. Terminology

The reader is assumed to be familiar with the terminology defined in [RFC5440], [RFC7551], [RFC8231], and [I-D.ietf-pce-association].

3. Overview

As shown in Figure 1, two reverse unidirectional LSPs can be grouped to form an associated bidirectional LSP. There are two methods of initiating the bidirectional LSP association, single-sided and double-sided as described in the following sections.

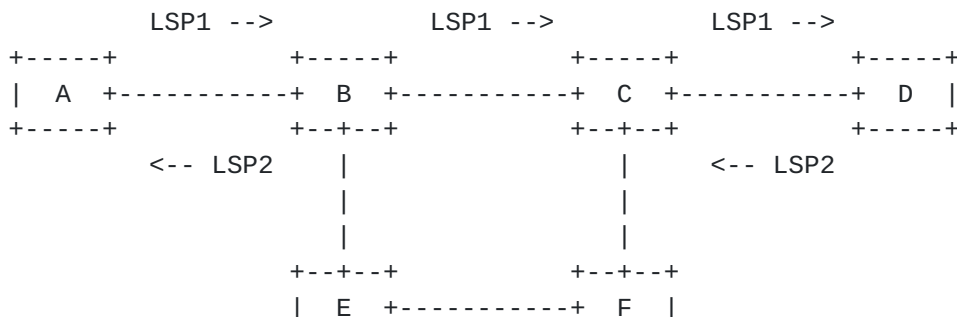




Figure 1: Example of Associated Bidirectional LSP

3.1. Single-sided Initiation

As specified in [RFC7551], in the single-sided case, the bidirectional tunnel is provisioned only on one endpoint node (PCC) of the tunnel. Both forward and reverse LSPs of this tunnel are initiated with the Association Type set to "Single-sided Bidirectional LSP Association" on the originating endpoint node. The forward and reverse LSPs are identified in the Bidirectional LSP Association Group TLV of their PCEP Association Objects.

The originating endpoint node signals the properties for the reverse LSP in the RSVP REVERSE_LSP Object [RFC7551] of the forward LSP Path message. The remote endpoint then creates the corresponding reverse tunnel and signals the reverse LSP in response to the received RSVP Path message.

The two unidirectional reverse LSPs on the originating endpoint node are grouped together using the PCEP Association Object and on the remote endpoint node by the RSVP signaled Association Object.

As shown in Figure 1, the forward tunnel and both the forward LSP LSP1 and the reverse LSP LSP2 are initiated on the originating endpoint node A, either by the PCE or the PCC. The creation of reverse tunnel and reverse LSP2 on the remote endpoint node D is triggered by the RSVP signaled LSP1.

As specified in [I-D.ietf-teas-assoc-corouted-bidir-frr], for fast-reroute bypass tunnel assignment, the LSP starting from the originating node is identified as the forward LSP of the single-sided initiated bidirectional LSP.

3.2. Double-sided Initiation

As specified in [RFC7551], in the double-sided case, the bidirectional tunnel is provisioned on both endpoint nodes (PCCs) of the tunnel. The forward and reverse LSPs of this tunnel are initiated with the Association Type set to "Double-sided Bidirectional LSP Association" on both endpoint nodes. The forward and reverse LSPs are identified in the Bidirectional LSP Association Group TLV of their Association Objects.

The two reverse unidirectional LSPs on both the endpoint nodes are grouped together by using the PCEP Association Object.

As shown in Figure 1, the forward tunnel and LSP1 are initiated on the endpoint node A and the reverse tunnel and LSP2 are initiated on the endpoint node D, either by the PCE or the PCCs.

As specified in [[I-D.ietf-teas-assoc-corouted-bidir-frr](#)], for fast-reroute bypass tunnel assignment, the LSP with the higher Source Address [[RFC3209](#)] is identified as the forward LSP of the double-sided initiated bidirectional LSP.

3.3. Co-routed Associated Bidirectional LSP

In both single-sided and double-sided initiation cases, forward and reverse LSPs may be co-routed as shown in Figure 2, where both forward and reverse LSPs follow the same congruent path in the forward and reverse directions, respectively.

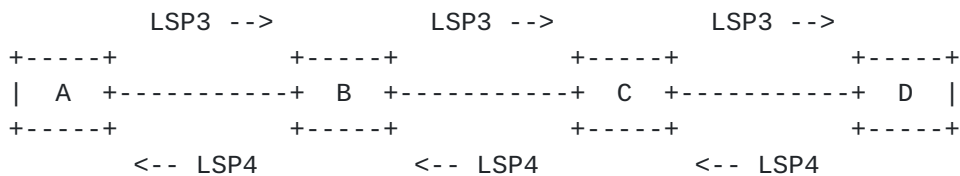


Figure 2: Example of Co-routed Associated Bidirectional LSP

4. Protocol Extensions

4.1. Association Object

As per [[I-D.ietf-pce-association](#)], LSPs are associated by adding them to a common association group. This document defines two new Bidirectional LSP Association Groups to be used by the associated bidirectional LSPs. A member of the Bidirectional LSP Association Group can take the role of a forward or reverse LSP and follows the following rules:

- o An LSP can not be part of more than one Bidirectional LSP Association Group.
- o The Tunnel (as defined in [[RFC3209](#)]) of forward and reverse LSPs of the single-sided bidirectional association MUST be the same.

This document defines two new Association Types for the Association Object as follows:

- o Association Type (TBD1) = Single-sided Bidirectional LSP Association Group

- o Association Type (TBD2) = Double-sided Bidirectional LSP Association Group

These Association Types are operator-configured associations in nature and statically created by the operator on the PCEP peers. The LSP belonging to these associations is conveyed via PCEP messages to the PCEP peer. Operator-configured Association Range TLV [I-D.ietf-pce-association] 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.

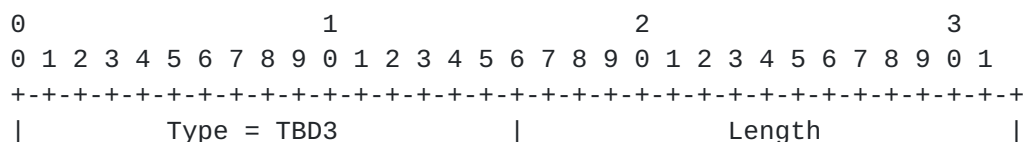
The Association ID, Association Source, optional Global Association Source and optional Extended Association ID in the Bidirectional LSP Association Group Object are initialized using the procedures defined in [RFC7551] and [I-D.ietf-pce-association].

4.2. Bidirectional LSP Association Group TLV

The Bidirectional LSP Association Group TLV is an optional TLV for use with the Single-sided and Double-sided Bidirectional LSP Association Group Object Types.

- o The Bidirectional LSP Association Group TLV follows the PCEP TLV format from [RFC5440].
- o The Type (16 bits) of the TLV is TBD3, to be assigned by IANA.
- o The Length is 4 Bytes.
- o The value comprises of a single field, the Bidirectional LSP Association Flags (32 bits), where each bit represents a flag option.
- o If the Bidirectional LSP Association Group TLV is missing, it means the LSP is the forward LSP.
- o The Bidirectional LSP Association Group TLV MUST NOT be present more than once. If it appears more than once, only the first occurrence is processed and any others MUST be ignored.

The format of the Bidirectional LSP Association Group TLV is shown in Figure 3:



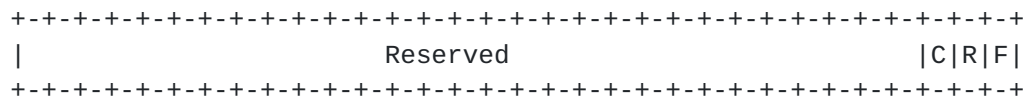


Figure 3: Bidirectional LSP Association Group TLV format

Bidirectional LSP Association Flags are defined as following.

F (Forward LSP, 1 bit) - Indicates whether the LSP associated is the forward LSP of the bidirectional LSP. If this flag is set, the LSP is a forward LSP.

R (Reverse LSP, 1 bit) - Indicates whether the LSP associated is the reverse LSP of the bidirectional LSP. If this flag is set, the LSP is a reverse LSP.

C (Co-routed LSP, 1 bit) - Indicates whether the bidirectional LSP is co-routed. This flag MUST be set for both the forward and reverse LSPs of a co-routed bidirectional LSP.

The C flag is used by the PCE (for both Stateful and Stateless) to compute bidirectional paths of the forward and reverse LSPs.

The Reserved flags MUST be set to 0 when sent and MUST be ignored when received.

5. PCEP Procedure

5.1. PCE Initiated LSPs

As specified in [[I-D.ietf-pce-association](#)], Association Groups can be created by both Stateful PCE and PCC.

A Stateful PCE can create and update the forward and reverse LSPs independently for both Single-sided and Double-sided bidirectional LSP association groups. The establishment and removal of the association relationship can be done on a per LSP basis. A PCE can create and update the association of the LSP on a PCC via PCInitiate and PCUpd messages, respectively, using the procedures described in [[I-D.ietf-pce-association](#)].

5.2. PCC Initiated LSPs

A PCC can create and update the forward and reverse LSPs independently for both Single-sided and Double-sided bidirectional LSP association groups. The establishment and removal of the association relationship can be done on a per LSP basis. In both

cases, the PCC must report the change in association to PCE(s) via PCRpt message. A PCC can also delegate the forward and reverse LSPs to a Stateful PCE, where PCE would control the LSPs. The Stateful PCE could update the LSPs in the association group via PCUpd message, using the procedures described in [[I-D.ietf-pce-association](#)].

5.3. Stateless PCE

For a stateless PCE, it might be useful to associate a path computation request to an association group, thus enabling it to associate a common set of configuration parameters or behaviors with the request. A PCC can request co-routed or non co-routed forward and reverse direction paths from a stateless PCE for the bidirectional LSP association group.

5.4. State Synchronization

During state synchronization, a PCC MUST report all the existing bidirectional LSP association groups to the Stateful PCE as per [[I-D.ietf-pce-association](#)]. After the state synchronization, the PCE MUST remove all stale bidirectional associations.

5.5. Error Handling

The LSPs (forward or reverse) in a single-sided bidirectional LSP association group MUST belong to the same TE Tunnel (as defined in [[RFC3209](#)]). If a PCE attempts to add an LSP in a single-sided bidirectional LSP association group for a different Tunnel, the PCC MUST send PCErr with Error-Type = 29 (Early allocation by IANA) (Association Error) and Error-Value = TBD4 (Bidirectional Association Tunnel Mismatch). Similarly, if a PCC attempts to add an LSP to a single-sided bidirectional LSP association group at PCE not complying to this rule, the PCE MUST send this PCErr.

6. Security Considerations

The security considerations described in [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)] apply to the extensions defined in this document as well.

Two new Association Types for the Association Object, Double-sided Bidirectional LSP Association Group and Single-sided Associated Bidirectional LSP Group are 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](#)] and apply to these Association Types. Hence, securing the PCEP session using Transport Layer Security (TLS) [[RFC8253](#)] is recommended.

7. Manageability Considerations

7.1. Control of Function and Policy

The mechanisms defined in this document do not imply any control or policy requirements in addition to those already listed in [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)].

7.2. Information and Data Models

[RFC7420] describes the PCEP MIB, there are no new MIB Objects defined for LSP associations.

The PCEP YANG module [[I-D.ietf-pce-pcep-yang](#)] supports LSP associations.

7.3. Liveness Detection and Monitoring

The mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)].

7.4. Verify Correct Operations

The mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)].

7.5. Requirements On Other Protocols

The mechanisms defined in this document do not add any new requirements on other protocols.

7.6. Impact On Network Operations

The mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)].

8. IANA Considerations

8.1. Association Types

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

Value Name	Reference

TBD1 Single-sided Bidirectional LSP Association Group	[This document]
TBD2 Double-sided Bidirectional LSP Association Group	[This document]

8.2. Bidirectional LSP Association Group TLV

This document defines a new TLV for carrying additional information of LSPs within a Bidirectional LSP Association Group. IANA is requested to add the assignment of a new value in the existing "PCEP TLV Type Indicators" registry as follows:

TLV-Type	Name	Reference

TBD3	Bidirectional LSP Association Group TLV	[This document]

8.2.1. Flag Fields in Bidirectional LSP Association Group TLV

This document requests that a new sub-registry, named "Bidirectional LSP Association Group TLV Flag Field", is created within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field in the Bidirectional LSP Association Group TLV. New values are to be assigned by Standards Action [[RFC8126](#)]. Each bit should be tracked with the following qualities:

- o Bit number (count from 0 as the most significant bit)
- o Description
- o Reference

The following values are defined in this document for the Flag field.

Bit No.	Description	Reference

31	F - Forward LSP	[This document]
30	R - Reverse LSP	[This document]
29	C - Co-routed LSP	[This document]

8.3. 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: TBD4 [This document]
Bidirectional Association Tunnel Mismatch

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997.
- [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.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [RFC7551] Zhang, F., Ed., Jing, R., and R. Gandhi, Ed., "RSVP-TE Extensions for Associated Bidirectional LSPs", [RFC 7551](#), May 2015.
- [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.
- [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.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", [RFC 8281](#), December 2017.
- [I-D.ietf-pce-association] 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](#) (work in progress).
- [I-D.ietf-teas-assoc-corouted-bidir-frr] Gandhi, R., Ed., Shah, H., and J. Whittaker, "Fast Reroute Procedures for Co-routed Associated Bidirectional Label Switched Paths (LSPs)", [draft-ietf-teas-assoc-corouted-bidir-frr](#) (work in progress).

9.2. Informative References

- [RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed.,

Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", [RFC 5654](#), September 2009.

- [RFC7420] Koushik, A., Stephan, E., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module", [RFC 7420](#), December 2014.
- [RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", [RFC 8051](#), January 2017.
- [RFC8253] Lopez, D., 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](#), October 2017.
- [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](#) (work in progress).

Acknowledgments

Authors would like to thank Dhruv Dhody for various discussions on association groups.

Authors' Addresses

Colby Barth
Juniper Networks

Email: cbarth@juniper.net

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

Bin Wen
Comcast

Email: Bin_Wen@cable.comcast.com

