

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 30, 2018

H. Ananthakrishnan
Packet Design
S. Sivabalan
Cisco
C. Barth
R. Torvi
Juniper Networks
I. Minei
Google, Inc
E. Crabbe
Individual Contributor
D. Dhody
Huawei Technologies
September 26, 2017

PCEP Extensions for MPLS-TE LSP Path Protection with stateful PCE
draft-ananthakrishnan-pce-stateful-path-protection-04

Abstract

A stateful Path Computation Element (PCE) is capable of computing as well as controlling via Path Computation Element Protocol (PCEP) Multiprotocol Label Switching Traffic Engineering Label Switched Paths (MPLS LSP). Furthermore, it is also possible for a stateful PCE to create, maintain, and delete LSPs. This document describes PCEP extension to associate two or more LSPs to provide end-to-end path protection.

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 March 30, 2018.

Copyright Notice

Copyright (c) 2017 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
1.1. Requirements Language	4
2. Terminology	4
3. PCEP Extensions	4
3.1. Path Protection Association Type	5
3.2. Path Protection Association TLV	6
4. Operation	7
4.1. PCC Initiated LSPs	7
4.2. PCE Initiated LSPs	7
4.3. State Synchronization	8
4.4. Session Termination	8
4.5. Error Handling	8
5. Other considerations	8
6. IANA considerations	9
6.1. Association Type	9
6.2. PPAG TLV	9
6.3. PCEP Errors	10
7. Security Considerations	10
8. Manageability Considerations	11
8.1. Control of Function and Policy	11
8.2. Information and Data Models	11
8.3. Liveness Detection and Monitoring	11
8.4. Verify Correct Operations	11
8.5. Requirements On Other Protocols	11
8.6. Impact On Network Operations	11
9. Acknowledgments	11
10. References	11
10.1. Normative References	12
10.2. Information References	13
Authors' Addresses	14

1. Introduction

[RFC5440] describes PCEP for communication between a Path Computation Client (PCC) and a PCE or between one a pair of PCEs as per [RFC4655]. A PCE computes paths for MPLS-TE LSPs based on various constraints and optimization criteria.

Stateful pce [RFC8231] specifies a set of extensions to PCEP to enable stateful control of paths such as MPLS TE LSPs between and across PCEP sessions in compliance with [RFC4657]. It includes mechanisms to effect LSP state synchronization between PCCs and PCEs, delegation of control of LSPs to PCEs, and PCE control of timing and sequence of path computations within and across PCEP sessions and focuses on a model where LSPs are configured on the PCC and control over them is delegated to the PCE. Furthermore, a mechanism to dynamically instantiate LSPs on a PCC based on the requests from a stateful PCE or a controller using stateful PCE, is specified in [I-D.ietf-pce-pce-initiated-lsp].

Path protection refers to a paradigm in which the working LSP is protected by one or more protection LSP(s). When the working LSP fails, protection LSP(s) is/are activated. When the working LSPs are computed and controlled by the PCE, there is benefit in a mode of operation where protection LSPs are as well.

This document specifies a stateful PCEP extension to associate two or more LSPs for the purpose of setting up path protection. The proposed extension covers the following scenarios:

- o A PCC initiates a protection 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 is the passive stateful mode [RFC8051].
- o A PCC initiates a protection LSP and delegates the control of the LSP to a stateful PCE. The PCE may compute the path for the LSP and update the PCC with the information about the path as long as it controls the LSP. This is the active stateful mode [RFC8051].
- o A protection LSP could be initiated by a stateful PCE, which retains the control of the LSP. The PCE is responsible for computing the path of the LSP and updating to the PCC with the information about the path. This is the PCE Initiated mode [I-D.ietf-pce-pce-initiated-lsp].

Note that protection LSP can be established prior to the failure (in which case the LSP is said to be in standby mode) or post failure of

the corresponding working LSP according to the operator choice/policy.

[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 that is equally applicable to stateful PCE (active and passive modes) and stateless PCE.

This document specifies a PCEP extension to associate one working LSP with one or more protection LSPs using the generic association mechanism.

This document describes a PCEP extension to associate protection LSPs by creating Path Protection Association Group (PPAG) and encoding this association in PCEP messages for stateful PCEP sessions.

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

The following terminologies are used in this document:

ERO: Explicit Route Object.

LSP: Label Switched Path.

PCC: Path Computation Client.

PCE: Path Computation Element

PCEP: Path Computation Element Protocol.

PPAG: Path Protection Association Group.

TLV: Type, Length, and Value.

3. PCEP Extensions

3.1. Path Protection Association Type

LSPs are not associated by listing the other LSPs with which they interact, but rather by making them belong to an association group referred to as "Path Protection Association Group" (PPAG) in this document. All LSPs join a PPAG individually. PPAG is based on the generic Association object used to associate two or more LSPs specified in [I-D.ietf-pce-association-group]. A member of a PPAG can take the role of working or protection LSP. This document defines a new association type called "Path Protection Association Type" of value TBD1. A PPAG can have one working LSP and/or one or more protection LSPs. The source, destination and Tunnel ID (as carried in LSP-IDENTIFIERS TLV [RFC8231], with description as per [RFC3209]) of all LSPs within a PPAG MUST be the same. As per [RFC3209], TE tunnel is used to associate a set of LSPs during reroute or to spread a traffic trunk over multiple paths.

The format of the Association object used for PPAG is specified in [I-D.ietf-pce-association-group] and reproduced in this document for easy reference in Figure 1 and Figure 2.

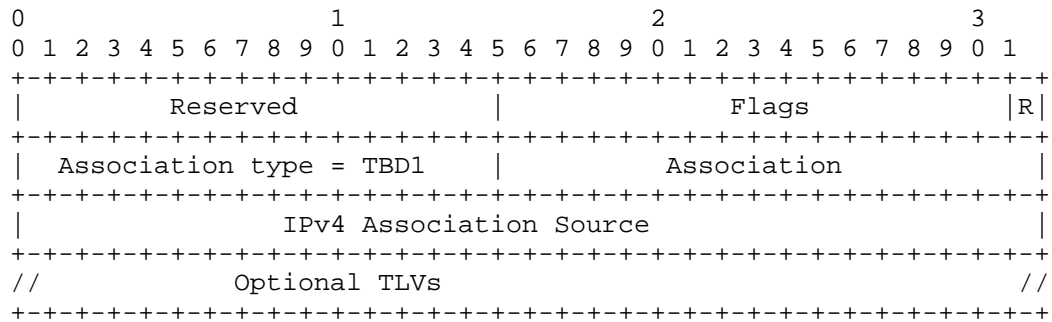


Figure 1: PPAG IPv4 ASSOCIATION Object format




Figure 2: PPAG IPv6 ASSOCIATION Object format

This document defines a new Association type, the Path Protection Association type, value will be assigned by IANA (TBD1).

This Association-Type is dynamic in nature and created by the PCC or PCE for the LSPs belonging to the same TE tunnel (as described in [RFC3209]) originating at the same head node and terminating at the same destination. These associations are conveyed via PCEP messages to the PCEP peer. Operator-configured Association Range SHOULD NOT be set for this association-type and MUST be ignored.

3.2. Path Protection Association TLV

The Path Protection Association TLV is an optional TLV for use with the Path Protection Association Object Type. The Path Protection Association 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 Path Protection Association TLV follows the PCEP TLV format of [RFC5440].

The type (16 bits) of the TLV is to be assigned by IANA. The length field is 16 bit-long and has a fixed value of 4.

The value comprises a single field, the Path Protection Association Flags (32 bits), where each bit represents a flag option.

The format of the Path Protection Association TLV (Figure 3) is as follows:




Figure 3: Path Protection Association TLV format

P (PROTECTION-LSP 1 bit) - Indicates whether the LSP associated with the PPAG is working or protection LSP. If this flag is set, the LSP is a protection LSP.

S (STANDBY 1 bit)- When the P flag is set, the S flag indicates whether the protection LSP associated with the PPAG is in standby mode. The S flag is ignored if the P flag is not set.

Unassigned bits are considered reserved. They MUST be set to 0 on transmission and MUST be ignored on receipt. If the Path Protection Association TLV is missing, it means the LSP is the working LSP.

4. Operation

LSPs are associated with other LSPs with which they interact by adding them to a common association group via ASSOCIATION object. All procedures and error-handling for the ASSOCIATION object is as per [I-D.ietf-pce-association-group].

4.1. PCC Initiated LSPs

A PCC can associate a set of LSPs under its control for path protection purpose. Similarly, the PCC can remove one or more LSPs under its control from the corresponding PPAG. In both cases, the PCC must report the change in association to PCE(s) via PCRpt message. A PCC can also delegate the working and protection LSPs to a stateful PCE, where PCE would control the LSPs. The stateful PCE could update the paths and attributes of the LSPs in the association group via PCUpd message. A PCE could also update the association to PCC via PCUpd message. The procedures are described in [I-D.ietf-pce-association-group].

4.2. PCE Initiated LSPs

A PCE can create/update working and protection LSPs independently. As specified in [I-D.ietf-pce-association-group], Association Groups can be created by both PCE and PCC.

Further, a PCE can remove a protection LSP from a PPAG as specified in [I-D.ietf-pce-association-group].

4.3. State Synchronization

During state synchronization, a PCC MUST report all the existing path protection association groups as well as any path protection flags to PCE(s). Following the state synchronization, the PCE would remove all stale information as per [I-D.ietf-pce-association-group].

4.4. Session Termination

As per [I-D.ietf-pce-association-group] the association information is cleared along with the LSP state information. When a PCEP session is terminated, after expiry of State Timeout Interval at PCC, the LSP state associated with that PCEP session is reverted to operator-defined default parameters or behaviors as per [RFC8231]. Same procedure is also followed for the association information. On session termination at the PCE, when the LSP state reported by PCC is cleared, the association information is also cleared. Where there are no LSPs in a association group, the association is considered to be deleted..

4.5. Error Handling

All LSPs (working or protection) within a PPAG MUST belong to the same TE Tunnel (as described in [RFC3209]) and have the same source and destination. If a PCE attempts to add an LSP to a PPAG and the Tunnel ID (as carried in LSP-IDENTIFIERS TLV [RFC8231], with description as per [RFC3209]) or source or destination of the LSP is different from the LSP(s) in the PPAG, the PCC MUST send PCErr with Error-Type= TBD (Association Error) [I-D.ietf-pce-association-group] and Error-Value = TBD3 (Tunnel ID or End points mismatch for Path Protection Association).

There MUST be only one working LSP within a PPAG. If a PCEP Speaker attempts to add another working LSP, the PCEP peer MUST send PCErr with Error-Type=TBD (Association Error) [I-D.ietf-pce-association-group] and Error-Value = TBD4 (Attempt to add another working LSP for Path Protection Association).

5. Other considerations

The diversity requirement for a group of LSPs is handled via another association type called "Disjointness Association", as described in [I-D.ietf-pce-association-diversity]. The diversity requirements for the the protection LSP are also handled by including both ASSOCIATION object for the group of LSPs.

6. IANA considerations

6.1. Association Type

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

Association Type Value	Association Name	Reference
TBD1	Path Protection Association	This document

6.2. PPAG TLV

This document defines a new TLV for carrying additional information of LSPs within a path protection association group. IANA is requested to make the assignment of a new value for the existing "PCEP TLV Type Indicators" registry as follows:

TLV Type Value	TLV Name	Reference
TBD2	Path Protection Association Group TLV	This document

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

Each bit should be tracked with the following qualities:

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

Bit Number	Name	Reference
31	P - PROTECTION-LSP	This document
30	S - STANDBY	This document

Table 1: PPAG TLV

6.3. PCEP Errors

This document defines new Error-Type and Error-Value related to path protection association. IANA is requested to allocate new error values within the "PCEP-ERROR Object Error Types and Values" sub-registry of the PCEP Numbers registry, as follows:

Error-Type	Meaning	Reference
TBD	Association error Error-value=TBD3: Tunnel ID or End points mismatch for Path Protection Association	[I-D.ietf-pce-association-group] This document
	Error-value=TBD4: Attempt to add another working LSP for Path Protection Association	This document

7. Security Considerations

The security considerations described in [RFC8231], [I-D.ietf-pce-pce-initiated-lsp], and [RFC5440] apply to the extensions described in this document as well. Additional considerations related to associations where a malicious PCEP speaker could be spoofed and could be used as an attack vector by creating associations is described in [I-D.ietf-pce-association-group]. Thus securing the PCEP session using Transport Layer Security (TLS) [I-D.ietf-pce-pceps], as per the recommendations and best current practices in [RFC7525], is RECOMMENDED.

8. Manageability Considerations

8.1. Control of Function and Policy

Mechanisms defined in this document do not imply any control or policy requirements in addition to those already listed in [RFC5440], [RFC8231], and [I-D.ietf-pce-pce-initiated-lsp].

8.2. Information and Data Models

[RFC7420] describes the PCEP MIB, there are no new MIB Objects for this document.

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

8.3. 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], [RFC8231], and [I-D.ietf-pce-pce-initiated-lsp].

8.4. 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-pce-initiated-lsp].

8.5. Requirements On Other Protocols

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

8.6. Impact On Network Operations

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [RFC5440], [RFC8231], and [I-D.ietf-pce-pce-initiated-lsp].

9. Acknowledgments

We would like to thank Jeff Tantsura and Xian Zhang for their contributions to this document.

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>>.
- [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>>.
- [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>>.
- [I-D.ietf-pce-pce-initiated-lsp]
Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", draft-ietf-pce-pce-initiated-lsp-10 (work in progress), June 2017.
- [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-04 (work in progress), August 2017.

10.2. Information 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>>.
- [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, DOI 10.17487/RFC7420, December 2014, <<https://www.rfc-editor.org/info/rfc7420>>.
- [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>>.
- [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>>.
- [I-D.ietf-pce-pceps]
Lopez, D., Dios, O., Wu, Q., and D. Dhody, "Secure Transport for PCEP", draft-ietf-pce-pceps-18 (work in progress), September 2017.
- [I-D.ietf-pce-pcep-yang]
Dhody, D., Hardwick, J., Beeram, V., and j. jeffrant@gmail.com, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", draft-ietf-pce-pcep-yang-05 (work in progress), June 2017.
- [I-D.ietf-pce-association-diversity]
Litkowski, S., Sivabalan, S., Barth, C., and D. Dhody, "Path Computation Element communication Protocol extension for signaling LSP diversity constraint", draft-ietf-pce-association-diversity-02 (work in progress), September 2017.

Authors' Addresses

Hariharan Ananthakrishnan
Packet Design
1 South Almaden Blvd, #1150,
San Jose, CA, 95113
USA

EMail: hari@packetdesign.com

Siva Sivabalan
Cisco
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

EMail: msiva@cisco.com

Colby Barth
Juniper Networks
1194 N Mathilda Ave,
Sunnyvale, CA, 94086
USA

EMail: cbarth@juniper.net

Raveendra Torvi
Juniper Networks
1194 N Mathilda Ave,
Sunnyvale, CA, 94086
USA

EMail: rtorvi@juniper.net

Ina Minei
Google, Inc
1600 Amphitheatre Parkway
Mountain View, CA, 94043
USA

EMail: inaminei@google.com

Edward Crabbe
Individual Contributor

EMail: edward.crabbe@gmail.com

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

EMail: dhruv.ietf@gmail.com