

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 11, 2016

H. Ananthakrishnan
Packet Design
S. Sivabalan
Cisco
C. Barth
R. Torvi
Juniper Networks
I. Minei
Google, Inc
E. Crabbe
March 10, 2016

PCEP Extensions for MPLS-TE LSP Path Protection with stateful PCE
draft-ananthakrishnan-pce-stateful-path-protection-01

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 <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."

This Internet-Draft will expire on September 11, 2016.

Copyright Notice

Copyright (c) 2016 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	2
2. Terminology	3
3. PCEP Extensions	4
3.1. Path Protection Association Type	4
3.2. Path Protection Association TLV	5
4. Operation	6
4.1. PCE Initiated LSPs	6
4.2. PCC Initiated LSPs	6
4.3. State Synchronization	7
4.4. Error Handling	7
5. IANA considerations	7
5.1. Association Type	7
5.2. PPAG TLV	7
5.3. PCEP Errors	8
6. Security Considerations	8
7. Acknowledgments	9
8. References	9
8.1. Normative References	9
8.2. Information References	10
Authors' Addresses	11

1. Introduction

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

Stateful pce [I-D.ietf-pce-stateful-pce] 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:

1. A protection LSP is initiated on a PCC by a stateful PCE which 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.
2. A PCC initiates a protection LSP and retains the control of the LSP. The PCC computes the path and updates the PCE with the information about the path as long as it controls the LSP.
3. 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.

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.

2. Terminology

The following terminologies are used in this document:

AGID: Association Group ID.

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 and destination of all LSPs within a PPAG MUST be the same.

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

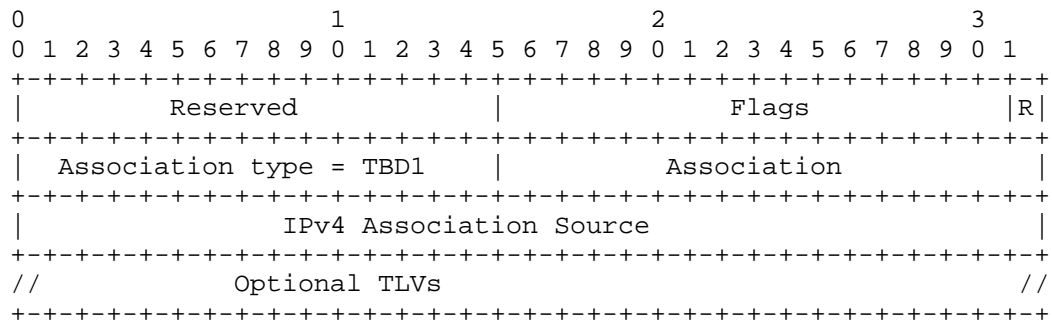


Figure 1: PPAG IPv4 ASSOCIATION Object format




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).

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.

If the Path Protection Association TLV is missing, it means the LSP is the working LSP.

4. Operation

4.1. 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.

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

4.2. PCC Initiated LSPs

A PCC can associate a set of LSPs under its control for path protection purpose. Similarly, the PCC can remove on 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 stateless PCC can request protection to a PCE thorough PCReq message.

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 MUST remove all stale path protection associations.

4.4. Error Handling

All LSPs (working or protection) within a PPAG MUST have the same source and destination. If a PCE attempts to add an LSP to a PPAG and the source and/or destination of the LSP is/are different from the LSP(s) in the PPAG, the PCC MUST send PCErr with Error-Type= TBD3 (Path Protection Association Error) and Error-Value = 1 (End points mismatch).

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=TBD3(Path Protection Association Error) and Error-Value = 2 (Attempt to add another working LSP).

5. IANA considerations

5.1. Association Type

This document defines a new association type for path protection as follows:

Association Type Value	Association Name	Reference
TBD1 (Suggested value - 1)	Path Protection Association	This document

5.2. PPAG TLV

This document defines a new TLV for carrying additional information of LSPs within a path protection association group as follows:

TLV Type Value	TLV Name	Reference
TBD2 (suggested Value - 29)	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 [RFC5226]. 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

5.3. PCEP Errors

This document defines new Error-Type and Error-Value related to path protection association as follows:

Error-Type	Meaning
TBD3 (suggested value - 25)	Path Protection Association error: Error-value=1: End-Points mismatch Error-value=2: Attempt to add another working LSP

6. Security Considerations

The same security considerations apply in head end as described in [I-D.ietf-pce-pce-initiated-lsp]

7. Acknowledgments

We would like to thank Jeff Tantsura, Dhruv Dhody and Zhangxian for their contributions to this document.

8. References

8.1. Normative References

- [I-D.ietf-pce-association-group]
Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Zhang, X., and Y. Tanaka, "PCEP Extensions for Establishing Relationships Between Sets of LSPs", draft-ietf-pce-association-group-00 (work in progress), November 2015.
- [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-05 (work in progress), October 2015.
- [I-D.ietf-pce-stateful-pce]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-13 (work in progress), December 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <<http://www.rfc-editor.org/info/rfc2205>>.
- [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, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <<http://www.rfc-editor.org/info/rfc4090>>.

- [RFC5088] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, DOI 10.17487/RFC5088, January 2008, <<http://www.rfc-editor.org/info/rfc5088>>.
- [RFC5089] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, DOI 10.17487/RFC5089, January 2008, <<http://www.rfc-editor.org/info/rfc5089>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<http://www.rfc-editor.org/info/rfc5440>>.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, DOI 10.17487/RFC5511, April 2009, <<http://www.rfc-editor.org/info/rfc5511>>.

8.2. Information References

- [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, DOI 10.17487/RFC2702, September 1999, <<http://www.rfc-editor.org/info/rfc2702>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<http://www.rfc-editor.org/info/rfc3031>>.
- [RFC3346] Boyle, J., Gill, V., Hannan, A., Cooper, D., Awduche, D., Christian, B., and W. Lai, "Applicability Statement for Traffic Engineering with MPLS", RFC 3346, DOI 10.17487/RFC3346, August 2002, <<http://www.rfc-editor.org/info/rfc3346>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<http://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, <<http://www.rfc-editor.org/info/rfc4657>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, DOI 10.17487/RFC5394, December 2008, <<http://www.rfc-editor.org/info/rfc5394>>.
- [RFC5557] Lee, Y., Le Roux, J.L., King, D., and E. Oki, "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", RFC 5557, DOI 10.17487/RFC5557, July 2009, <<http://www.rfc-editor.org/info/rfc5557>>.

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

EMail: edward.crabbe@gmail.com