

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 12, 2016

I. Minei
Google, Inc.
E. Crabbe

S. Sivabalan
Cisco Systems, Inc.
H. Ananthakrishnan
Packet Design
X. Zhang
Huawei Technologies
Y. Tanaka
NTT Communications Corporation
November 9, 2015

PCEP Extensions for Establishing Relationships Between Sets of LSPs
draft-minei-pce-association-group-04

Abstract

This document introduces a generic mechanism to create a grouping of LSPs in the context of a PCE. This grouping can then be used to define associations between sets of LSPs or between a set of LSPs and a set of attributes (such as configuration parameters or behaviors), and is equally applicable to the active and passive modes of a stateful PCE as well as a stateless PCE.

Requirements Language

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](#)].

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 May 12, 2016.

Copyright Notice

Copyright (c) 2015 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.	Architectural Overview	3
3.1.	Motivation	3
3.2.	Operation Overview	4
4.	ASSOCIATION Object	4
4.1.	Object Definition	4
4.1.1.	Global Association Source TLV	6
4.1.2.	Extended Association ID TLV	6
4.2.	Object Encoding in PCEP messages	7
4.3.	Processing Rules	9
5.	IANA Considerations	10
6.	Security Considerations	11
7.	Acknowledgements	11
8.	Contributors	11
9.	References	11
9.1.	Normative References	11
9.2.	Informative References	12
	Authors' Addresses	12

[1.](#) Introduction

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

Stateful pce [[I-D.ietf-pce-stateful-pce](#)] specifies a set of extensions to PCEP to enable stateful control of TE LSPs between and across PCEP sessions in compliance with [[RFC4657](#)] and focuses on a model where LSPs are configured on the PCC and control over them is delegated to the PCE. The model of operation where LSPs are initiated from the PCE is described in [[I-D.ietf-pce-pce-initiated-lsp](#)].

This document introduces a generic mechanism to create a grouping of LSPs. This grouping can then be used to define associations between sets of LSPs or between a set of LSPs and a set of attributes (such as configuration parameters or behaviors), and is equally applicable to the active and passive modes of a stateful PCE and a stateless PCE.

2. Terminology

This document uses the following terms defined in [[RFC5440](#)]: PCC, PCE, PCEP Peer.

The following term is defined in this document:

Association Timeout Interval: when a PCEP session is terminated, a PCC waits for this time period before deleting associations created by the PCEP peer.

3. Architectural Overview

3.1. Motivation

Stateful PCE provides the ability to update existing LSPs and to instantiate new ones. To enable support for PCE-controlled make-before-break and for protection, there is a need to define associations between LSPs. For example, the association between the original and the re-optimized path in the make-before break scenario, or between the working and protection path in end-to-end protection. Another use for LSP grouping is for applying a common set of configuration parameters or behaviors to a set of LSPs.

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.

Rather than creating separate mechanisms for each use case, this draft defines a generic mechanism that can be reused as needed.

3.2. Operation Overview

LSPs are associated with other LSPs with which they interact by adding them to a common association group. Association groups as defined in this document can be applied to LSPs originating at the same head end or different head ends. For LSPs originating at the same head end, the association can be initiated by either the PCC (head end) or by a PCE. Only a stateful PCE can initiate an association for LSPs originating at different head ends. For both cases, the association is uniquely identified by the combination of an association identifier and the address of the node that created the association.

Multiple types of groups can exist, each with their own identifiers space. The definition of the different association types and their behaviors is outside the scope of this document. The establishment and removal of the association relationship can be done on a per LSP basis. An LSP may join multiple association groups, of different or of the same type.

In the case of a stateless PCE, associations are created out of band, and PCEP peers should be aware of the association and its significance outside of the protocol.

Association groups can be created by both PCC and PCE. When a PCC's PCEP session with a PCE terminates unexpectedly, the PCC cleans up associations (as per the processing rules in this document).

4. ASSOCIATION Object

4.1. Object Definition

Creation of an association group and modifications to its membership can be initiated by either the PCE or the PCC. Association groups and their memberships are defined using the ASSOCIATION object for stateful PCE.

ASSOCIATION Object-Class is to be assigned by IANA (TBD).

ASSOCIATION Object-Type is 1 for IPv4 and its format is shown in Figure 1:

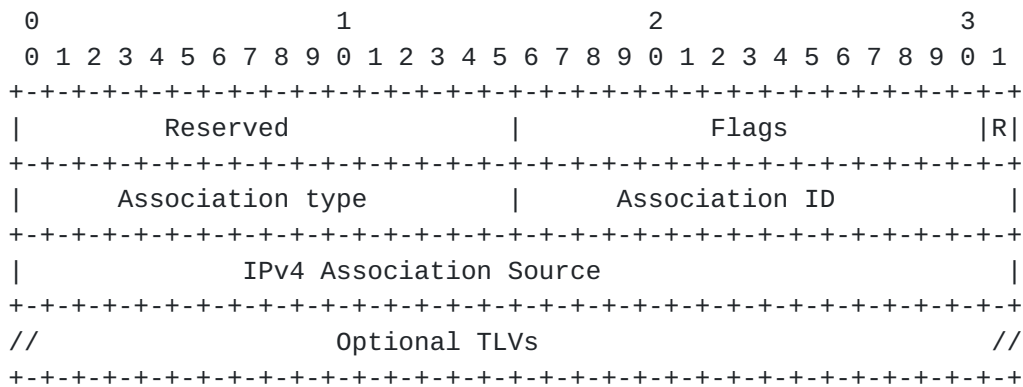


Figure 1: The IPv4 ASSOCIATION Object format

ASSOCIATION Object-Type is 2 for IPv6 and its format is shown in Figure 2:




Figure 2: The IPv6 ASSOCIATION Object format

Reserved: 16 bits - MUST be set to 0 and ignored upon receipt.

Flags: 16 bits - The following flags are currently defined:

R (Removal - 1 bit): when set, the requesting PCE peer requires the removal of an LSP from the association group.

Association type: 16 bits - the association type (for example protection). The association type will be defined in separate documents.

Association ID: 16 bits - the identifier of the association group. When combined with Type and Association Source, this value uniquely identifies an association group. The value 0xffff and 0x0 are reserved. The value 0xffff is used to indicate all association groups.

Association Source: 4 or 16 bytes - An IPv4 or IPv6 address, which is associated to the node that originated the association.

Optional TLVs: The optional TLVs follow the PCEP TLV format of [[RFC5440](#)]. This document defines two optional TLVs.

4.1.1. Global Association Source TLV

The Global Association Source TLV is an optional TLV for use in the Association Object.

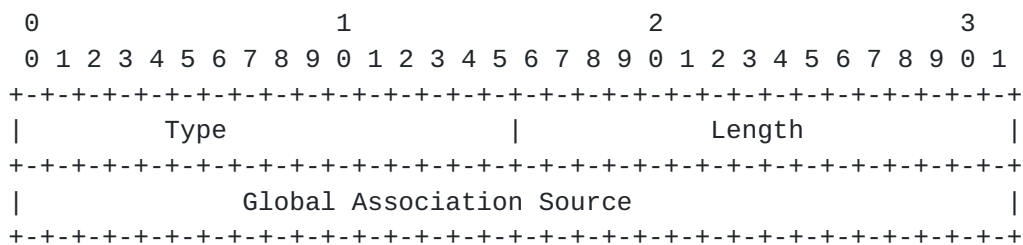


Figure 3: The Global Association Source TLV format

Type: To be allocated by IANA

Length: Fixed value of 4 bytes

Global Association Source: as defined in [[RFC6780](#)]

4.1.2. Extended Association ID TLV

The Extended Association ID TLV is an optional TLV for use in the Association Object.

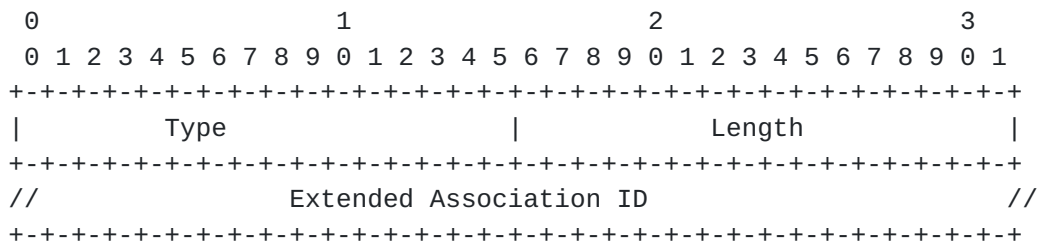


Figure 4: The Extended Association ID TLV format

Type: To be allocated by IANA

Length: variable

Extended Association ID: as defined in [[RFC6780](#)]

4.2. Object Encoding in PCEP messages

The ASSOCIATION Object is OPTIONAL and MAY be carried in the Path Computation Update (PCUpd), Path Computation Report (PCRpt) and Path Computation Initiate (PCinit) messages.

When carried in PCRpt message, it is used to report the association group membership information pertaining to a LSP to a stateful PCE. It can also be used to remove an LSP from one or more association groups by setting the R flag to 1. Unless, a PCE wants to delete an association from an LSP, it does not need to carry the ASSOCIATION object while updating other LSP attributes using the PCUpd message.

The PCRpt message is defined in [[I-D.ietf-pce-stateful-pce](#)] and updated as below:

```
<PCRpt Message> ::= <Common Header>
                        <state-report-list>
```

Where:

```
<state-report-list> ::= <state-report>[<state-report-list>]
```

```
<state-report> ::= [<SRP>
                     <LSP>
                     [<association-list>]
                     <path>
```

Where:

```
<association-list> ::= <ASSOCIATION> [<association-list>]
```


When an LSP is delegated to a stateful PCE, the stateful PCE can initiate a new association group for this LSP, or associate it with one or more existing association groups. This is done by including the ASSOCIATION Object in a PCUpd message or in a PCInit message. A stateful PCE can also remove a delegated LSP from one or more association groups by setting the R flag to 1.

The PCUpd message is defined in [[I-D.ietf-pce-stateful-pce](#)] and updated as below:

```
<PCUpd Message> ::= <Common Header>
                        <update-request-list>
```

Where:

```
<update-request-list> ::= <update-request>[<update-request-list>]
```

```
<update-request> ::= <SRP>
                    <LSP>
                    [<association-list>]
                    <path>
```

Where: <association-list> ::= <ASSOCIATION> [<association-list>]

The PCInitiate message is defined in [[I-D.ietf-pce-pce-initiated-lsp](#)] and updated as below:

```
<PCInitiate Message> ::= <Common Header>
                        <PCE-initiated-lsp-list>
```

Where:

```
<PCE-initiated-lsp-list> ::=
<PCE-initiated-lsp-request>[<PCE-initiated-lsp-list>]

<PCE-initiated-lsp-request> ::=
(<PCE-initiated-lsp-instantiation>|<PCE-initiated-lsp-deletion>)

<PCE-initiated-lsp-instantiation> ::= <SRP>
                                     <LSP>
                                     <END-POINTS>
                                     <ERO>
                                     [<association-list>]
                                     [<attribute-list>]
```

Where:

```
<association-list> ::= <ASSOCIATION> [<association-list>]
```

In case of passive stateful or stateless PCE, the ASSOCIATION Object is OPTIONAL and MAY be carried in the Path Computation Request (PCReq) message.

When carried in a PCReq message, the ASSOCIATION Object is used to associate the path computation request to an association group, the association might be further informed via PCRpt message in case of passive stateful PCE later or it might be created out of band in case of stateless PCE.

The PCReq message is defined in [[RFC5440](#)] and updated in [I-D.ietf-pce-stateful-pce], it is further updated below for association:

```
<PCReq Message> ::= <Common Header>
                    [<svec-list>]
                    <request-list>
```

Where:

```
<svec-list> ::= <SVEC> [<svec-list>]
<request-list> ::= <request> [<request-list>]

<request> ::= <RP>
              <END-POINTS>
              [<LSP>]
              [<LSPA>]
              [<BANDWIDTH>]
              [<metric-list>]
              [<association-list>]
              [<RRO> [<BANDWIDTH>]]
              [<IRO>]
              [<LOAD-BALANCING>]
```

Where:

```
<association-list> ::= <ASSOCIATION> [<association-list>]
```

Note that LSP object MAY be present for the passive stateful PCE.

4.3. Processing Rules

Both a PCC and a PCE can create one or more association groups for an LSP. But a PCE peer cannot add new members for association group created by another peer. If a PCE peer does not recognize the ASSOCIATION object, it MUST return a PCErr message with Error-Type "Unknown Object" as described in [[RFC5440](#)]. If a PCE peer is unwilling or unable to process the ASSOCIATION object, it MUST return a PCErr message with the Error-Type "Not supported object" and follow the relevant procedures described in [[RFC5440](#)].

The association timeout interval is as a PCC-local value that can be operator-configured or computed by the PCC based on local policy and is used in the context of cleaning up associations on session failure. The association timeout must be set to a value no larger

than the state timeout interval (defined in [\[I-D.ietf-pce-stateful-pce\]](#)) and larger than the delegation timeout interval (defined in [\[I-D.ietf-pce-stateful-pce\]](#)).

When a PCC's PCEP session with the PCE terminates unexpectedly, the PCC MUST wait for the association timeout interval before cleaning up the association. If this PCEP session can be re-established before the association timeout interval time expires, no action is taken to clean the association created by this PCE. During the time window of the redelegation timeout interval and the association timeout interval, the PCE, after re-establishing the session, can also ask for redelegation following the procedure defined in [\[I-D.ietf-pce-stateful-pce\]](#) and [\[I-D.ietf-pce-pce-initiated-lsp\]](#). When the association timeout interval timer expires, the PCC clears all the associations which are not delegated to any PCEs.

Upon LSP delegation revocation, the PCC MAY clear the association created by the related PCE, but in order to avoid traffic loss, it can perform this in a make-before-break fashion, which is the same as what is defined in Stateful pce [\[I-D.ietf-pce-stateful-pce\]](#) for handling LSP state cleanup.

Error handling for situations for multiple PCE scenarios will be included in future versions of this draft.

5. IANA Considerations

The "PCEP Parameters" registry contains a subregistry "PCEP Objects". This document request IANA to allocate the values from this registry.

Object-Class Value	Name	Reference
TBD	Association Object-Type 1: IPv4 2: IPv6	This document

This document defines the following new PCEP TLVs:

Value	Meaning	Reference
TBD	Global Association Source	This document
TBD	Extended Association Id	This document

This document requests IANA to create a subregistry of the "PCEP Parameters" for the bits carried in the Flags field of the ASSOCIATION object. The subregistry is called "ASSOCIATION Flags Field".

The field contains 12 bits numbered from bit 0 as the most significant bit.

Bit;	Name: Description	Reference
15	R: Removal	This document

6. Security Considerations

The security considerations described in [[I-D.ietf-pce-stateful-pce](#)] apply to the extensions described in this document. Additional considerations related to a malicious PCE are introduced, as the PCE may now create additional state on the PCC through the creation of association groups.

7. Acknowledgements

We would like to thank Yuji Kamite and Joshua George for their contributions to this document. Also Thank Venugopal Reddy and Cyril Margaria for their useful comments.

8. Contributors

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560037
India
Email: dhruv.ietf@gmail.com

9. References

9.1. Normative References

- [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-12](#) (work in progress), October 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>>.
- [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>>.
- [RFC6780] Berger, L., Le Faucheur, F., and A. Narayanan, "RSVP ASSOCIATION Object Extensions", [RFC 6780](#), DOI 10.17487/RFC6780, October 2012, <<http://www.rfc-editor.org/info/rfc6780>>.

9.2. 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, <<http://www.rfc-editor.org/info/rfc4657>>.

Authors' Addresses

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

Email: inaminei@google.com

Edward Crabbe

Email: edward.crabbe@gmail.com

Siva Sivabalan
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
US

Email: msiva@cisco.com

Hariharan Ananthakrishnan
Packet Design

Email: hari@packetdesign.com

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base Bantian, Longgang District
Shenzhen, Guangdong 518129
P.R.China

Email: zhang.xian@huawei.com

Yosuke Tanaka
NTT Communications Corporation
Granpark Tower 3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan

Email: yosuke.tanaka@ntt.com

