

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 2, 2020

I. Minei
Google, Inc.
E. Crabbe
Individual Contributor
S. Sivabalan
Cisco Systems, Inc.
H. Ananthakrishnan
Netflix
D. Dhody
Huawei
Y. Tanaka
NTT Communications Corporation
August 1, 2019

Path Computation Element Communication Protocol (PCEP) Extensions for
Establishing Relationships Between Sets of Label Switched Paths (LSPs)
draft-ietf-pce-association-group-10

Abstract

This document introduces a generic mechanism to create a grouping of Label Switched Paths (LSPs) in the context of a Path Computation Element (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 behaviours), and is equally applicable to the stateful PCE (active and passive modes) as well as the stateless PCE.

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.

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 February 2, 2020.

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
3.	Architectural Overview	4
3.1.	Motivation	4
3.2.	Relationship with the SVEC List	5
3.3.	Operation Overview	5
3.4.	Operator-configured Association Range	6
4.	Discovery of Supported Association Types	7
4.1.	ASSOC-Type-List TLV	7
4.1.1.	Procedure	8
5.	Operator-configured Association Range TLV	8
5.1.	Procedure	10
6.	ASSOCIATION Object	11
6.1.	Object Definition	11
6.1.1.	Global Association Source TLV	13
6.1.2.	Extended Association ID TLV	13
6.1.3.	Association Source	14
6.1.4.	Unique Identification for an Association Group	14
6.2.	Relationship with the RSVP ASSOCIATION	15
6.3.	Object Encoding in PCEP messages	15
6.3.1.	Stateful PCEP messages	15
6.3.2.	Request Message	17
6.3.3.	Reply Message	18
6.4.	Processing Rules	19

7.	IANA Considerations	21
7.1.	PCEP Object	21
7.2.	PCEP TLV	22
7.3.	Association Flags	22
7.4.	Association Type	23
7.5.	PCEP-Error Object	23
8.	Security Considerations	24
9.	Manageability Considerations	24
9.1.	Control of Function and Policy	24
9.2.	Information and Data Models	24
9.3.	Liveness Detection and Monitoring	25
9.4.	Verify Correct Operations	25
9.5.	Requirements on Other Protocols	25
9.6.	Impact on Network Operations	25
10.	Acknowledgments	25
11.	Contributors	25
12.	References	26
12.1.	Normative References	26
12.2.	Informative References	27
	Appendix A. Example for Operator-configured Association Range	28
	Authors' Addresses	29

1. Introduction

[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].

[RFC4872] defines the RSVP ASSOCIATION object, which was defined in the context of GMPLS-controlled Label Switched Paths (LSPs) to be used to associate recovery LSPs with the LSP they are protecting. This object also has broader applicability as a mechanism to associate RSVP state and [RFC6780] described how the ASSOCIATION object can be more generally applied by defining the Extended ASSOCIATION Object.

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 behaviours), and is equally applicable to stateful PCE (active and passive modes) and stateless PCE. The associations could be created dynamically and conveyed to a PCEP peer within PCEP, or it could be configured manually by an operator on the PCEP peers. Refer Section 3.3 for more details.

2. Terminology

This document uses the following terms defined in [RFC5440]: PCC, PCE, PCEP Peer, Path Computation Request (PCReq), Path Computation Reply (PCRep), and PCEP Error (PCErr).

This document uses the following terms defined in [RFC8051]: Stateful PCE, Active Stateful PCE, Passive Stateful PCE, and Delegation.

This document uses the following terms defined in [RFC8231]: LSP State Report (PCRpt), LSP Update Request (PCUpd), and State Timeout Interval.

This document uses the following terms defined in [RFC8281]: PCE-initiated LSP, and LSP Initiate Request (PCInitiate).

3. Architectural Overview

3.1. Motivation

Stateful PCE provides the ability to update existing LSPs and to instantiate new ones. There are various situations where several LSPs need to share common information. E.g., to support for PCE-controlled make-before-break, an association between the original and the re-optimized path is desired. Similarly, for end-to-end protection, the association between working and protection LSPs is required (see [I-D.ietf-pce-stateful-path-protection]). For diverse paths, an association between a group of LSPs could be used (see [I-D.ietf-pce-association-diversity]). Another use for the LSP grouping is for applying a common set of configuration parameters or behaviours 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 policy, configuration parameters or behaviours with the request.

Some associations could be created dynamically, such as association between the working and protection LSPs of a tunnel, whereas some

associations could be created by the operator manually, such as policy-based association, where the LSP could join an operator-configured existing association.

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

3.2. Relationship with the SVEC List

Note that, [RFC5440] defines a mechanism for the synchronization of a set of path computation requests by using the SVEC (Synchronization VECTOR) object, that specifies the list of synchronized requests that can either be dependent or independent. The SVEC object identifies the relationship between the set of path computation requests, identified by 'Request-ID-number' in RP (Request Parameters) object. [RFC6007] further clarifies the use of the SVEC list for synchronized path computations when computing dependent requests as well as describes a number of usage scenarios for SVEC lists within single-domain and multi-domain environments.

The motivation behind the association group defined in this document and the SVEC object are quite different, though some use cases may overlap. PCEP extensions that define a new association type should clarify the relationship between the SVEC object and the association type, if any.

3.3. 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.

Some associations could be created dynamically by a PCEP speaker and the associations (along with the set of LSPs) are conveyed to a PCEP peer. Whereas some associations are configured by the operator on the PCEP peers involved beforehand, a PCEP speaker then could ask for a LSP to join the operator-configured association. Usage of dynamic and configured association is usually dependent on the type of the association.

For the operator-configured association, the association parameters such as the association identifier, association type, as well as the association source IP address, are manually configured by the operator. In case of dynamic association, the association parameters such as the association identifier, are allocated dynamically by the PCEP speaker, the association source is set as local PCEP speaker

address, unless local policy dictates otherwise, in which case association source is set based on the local policy.

The dynamically created association can be reported to the PCEP peer via the PCEP messages as per the stateful extensions. When the operator-configured association is known to the PCEP peer beforehand, a PCEP peer could ask for a LSP to join the operator-configured association via the stateful PCEP messages.

The associations are properties of the LSP and thus could be stored in the LSP state database. The dynamic association exists as long as the LSP state exists. In case of PCEP session termination, the LSP state clean-up MUST also take care of associations.

Multiple types of associations can exist, each with their own association identifier space. The definition of the different association types and their behaviours 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 association type.

3.4. Operator-configured Association Range

Some association types are dynamic, some are operator-configured and some could be both. For the association types that could be both dynamic and operator-configured and use the same association source, it is necessary to distinguish a range of association identifiers that are marked for operator-configured associations to avoid any association identifier clash within the scope of the association source. This document assumes that these two ranges are configured.

A range of association identifiers for each Association type (and Association Source) are kept for the operator-configured associations. Dynamic associations MUST NOT use the association identifier from this range.

This range as set at the PCEP speaker (PCC or PCE, as an association source) needs to be communicated to a PCEP peer in the Open Message. A new TLV is defined in this specification for this purpose (Section 5). See Appendix A for an example.

Association identifier range for sources other than the PCEP speaker (for example an NMS system) is not communicated in PCEP and the procedure for operator-configured association range setting is outside the scope of this document.

4. Discovery of Supported Association Types

This section defines PCEP extensions so as to support the capability advertisement of the association types supported by a PCEP speaker.

A new PCEP ASSOC-Type-List (Association Types list) TLV is defined. The PCEP ASSOC-Type-List TLV is carried within an OPEN object. This way, during PCEP session-setup phase, a PCEP speaker can advertise to a PCEP peer the list of supported Association types.

4.1. ASSOC-Type-List TLV

The PCEP ASSOC-Type-List TLV is OPTIONAL. It MAY be carried within an OPEN object sent by a PCEP speaker in an Open message to a PCEP peer so as to indicate the list of supported Association types.

The PCEP ASSOC-Type-List TLV format is compliant with the PCEP TLV format defined in [RFC5440]. That is, the TLV is composed of 2 octets for the type, 2 octets specifying the TLV length, and a Value field. The Length field defines the length of the value portion in octets. The TLV is padded to 4-octet alignment, and padding is not included in the Length field (e.g., a 3-octet value would have a length of three, but the total size of the TLV would be eight octets).

The PCEP ASSOC-Type-List TLV has the following format:

TYPE: TBD
 LENGTH: N * 2 (where N is the number of association types)
 VALUE: list of 2-byte association type code points, identifying the association types supported by the sender of the Open message.




Figure 1: The ASSOC-Type-List TLV format

Assoc-type (2 bytes): Association type code point identifier. IANA manages the "ASSOCIATION Type Field" code point registry (see Section 7.4).

4.1.1. Procedure

An ASSOC-Type-List TLV within an OPEN object in an Open message is included by a PCEP speaker in order to advertise a set of one or more supported association types. The ASSOC-Type-List TLV MUST NOT appear more than once in an OPEN object. If it appears more than once, the PCEP session MUST be rejected with error type 1 and error value 1 (PCEP session establishment failure / Reception of an invalid Open message). As specified in [RFC5440], a PCEP peer that does not recognize the ASSOC-Type-List TLV will silently ignore it.

Association type (to be defined in other documents) can specify if the association type advertisement is mandatory for it. Thus, the ASSOC-Type-List TLV MUST be included if at least one mandatory association type needs to be advertised and the ASSOC-Type-List TLV MAY be included otherwise. For an association type that specifies that the advertisement is mandatory, a missing Assoc-type in the ASSOC-Type-List TLV (or missing ASSOC-Type-List TLV) is to be interpreted as the association type is not supported by the PCEP speaker.

The absence of the ASSOC-Type-List TLV in an OPEN object MUST be interpreted as an absence of information on the list of supported Association types (rather than the Association type is not supported). In this case, the PCEP speaker could still use the ASSOCIATION object: if the peer does not support the association, it will react as per the procedure described in Section 6.4.

In case the use of the ASSOC-Type-List TLV is triggered by support for a mandatory association type, then it is RECOMMENDED that the PCEP implementation include all supported Association types (including optional) to ease the operations of the PCEP peer.

5. Operator-configured Association Range TLV

This section defines PCEP extension to support the advertisement of the Operator-configured Association Range used for an Association type by the PCEP speaker (as an Association source).

A new PCEP OP-CONF-ASSOC-RANGE (Operator-configured Association Range) TLV is defined. The PCEP OP-CONF-ASSOC-RANGE TLV is carried within an OPEN object. This way, during PCEP session-setup phase, a PCEP speaker can advertise to a PCEP peer the Operator-configured Association Range for an association type.

The PCEP OP-CONF-ASSOC-RANGE TLV is OPTIONAL. It MAY be carried within an OPEN object sent by a PCEP speaker in an Open message to a PCEP peer. The OP-CONF-ASSOC-RANGE TLV format is compliant with the PCEP TLV format defined in [RFC5440]. That is, the TLV is composed of 2 bytes for the type, 2 bytes specifying the TLV length, and a Value field. The Length field defines the length of the value portion in bytes.

The PCEP OP-CONF-ASSOC-RANGE TLV has the following format:

TYPE: 29 (Early allocation by IANA)
 LENGTH: N * 8 (where N is the number of association types)
 VALUE:




Figure 2: The OP-CONF-ASSOC-RANGE TLV format

The Value portion includes the following fields, repeated for each association type:

Reserved (2 bytes): This field MUST be set to 0 on transmission and MUST be ignored on receipt.

Assoc-type (2 bytes): The association type (Section 7.4). The association types are defined in separate documents.

Start-Assoc-ID (2 bytes): The start association identifier for the Operator-configured Association Range for the particular association type. The values 0 and 0xffff MUST NOT be used and on receipt of these values in the TLV, the session is rejected with error message sent (see Section 5.1).

Range (2 bytes): The number of associations marked for the Operator-configured Associations. The Range MUST be greater than 0, and it MUST be such that (Start-Assoc-ID + Range) do not cross the association identifier range of 0xffff. In case this condition is not satisfied, the session is rejected with error message sent (see Section 5.1).

5.1. Procedure

A PCEP speaker MAY include an OP-CONF-ASSOC-RANGE TLV within an OPEN object in an Open message sent to a PCEP peer in order to advertise the Operator-configured Association Range for an association type. The OP-CONF-ASSOC-RANGE TLV MUST NOT appear more than once in an OPEN object. If it appears more than once, the PCEP session MUST be rejected with error type 1 and error value 1 (PCEP session establishment failure / Reception of an invalid Open message).

As specified in [RFC5440], a PCEP peer that does not recognize the OP-CONF-ASSOC-RANGE TLV will silently ignore it.

The Operator-configured Association Range SHOULD be included for each association type that could be both dynamic and operator-configured. For association types that are only dynamic or only operator-configured, this TLV MAY be skipped, in which case the full range of association identifier is considered dynamic or operator-configured respectively. Each association type (that are defined in separate documents) can specify the default value for the operator-configured association range for their respective association type.

The absence of the OP-CONF-ASSOC-RANGE TLV in an OPEN object MUST be interpreted as an absence of explicit Operator-configured Association Range at the PCEP peer. In this case, the default behavior as per each association type applies. If the association source is not a PCEP speaker, the default value for the operator-configured association range is used for the association source.

If the Assoc-type is not recognized or supported by the PCEP speaker, it MUST ignore that respective Start-Assoc-ID and Range. If the Assoc-type is recognized/supported but the Start-Assoc-ID or Range are set incorrectly, the PCEP session MUST be rejected with error type 1 and error value 1 (PCEP session establishment failure / Reception of an invalid Open message). The incorrect range include the case when the (Start-Assoc-ID + Range) crosses the association identifier range of 0xffff.

A given Assoc-type MAY appear more than once in the OP-CONF-ASSOC-RANGE TLV in the case of a non-contiguous Operator-configured Association Range. The PCEP speaker originating this TLV MUST NOT

ASSOCIATION Object-Type is 2 for IPv6 and its format is shown in Figure 4:




Figure 4: The IPv6 ASSOCIATION Object format

Reserved (2-byte): MUST be set to 0 and ignored upon receipt.

Flags (2-byte): The following flags are currently defined:

R (Removal - 1 bit): when set, the requesting PCEP peer requires the removal of an LSP from the association group. When unset, the PCEP peer indicates that the LSP is added or retained as part of the association group. This flag is used for the ASSOCIATION object in the PCRpt and the PCUpd message, the flag is ignored in other PCEP messages.

The unassigned flags MUST be set to zero on transmission and MUST be ignored on receipt.

Association type (2-byte): the association type (Section 7.4). The association types are defined in separate documents.

Association ID (2-byte): the identifier of the association group. When combined with other association parameters, such as Association Type and Association Source, this value uniquely identifies an association group. The values 0xffff and 0x0 are reserved. The value 0xffff is used to indicate all association groups and could be used with R flag to indicate removal for all associations for the LSP within the scope of association type and association source.

Association Source: Contains a valid IPv4 address (4 bytes) if the ASSOCIATION Object-Type is 1 or a valid IPv6 address (16 bytes) if the ASSOCIATION Object-Type is 2. The address provides scoping for the Association ID. See Section 6.1.3 for details.

Optional TLVs: The optional TLVs follow the PCEP TLV format of [RFC5440]. This document defines two optional TLVs. Other documents can define more TLVs in future.

6.1.1. Global Association Source TLV

The Global Association Source TLV is an optional TLV for use in the Association Object. The meaning and the usage of Global Association Source is as per Section 4 of [RFC6780].




Figure 5: The Global Association Source TLV format

Type: 30 (Early allocation by IANA).

Length: Fixed value of 4 bytes.

Global Association Source: as defined in Section 4 of [RFC6780].

6.1.2. Extended Association ID TLV

The Extended Association ID TLV is an optional TLV for use in the Association Object. The meaning and the usage of Extended Association ID is as per Section 4 of [RFC6780].




Figure 6: The Extended Association ID TLV format

Type: 31 (Early allocation by IANA).

Length: variable.

Extended Association ID: as defined in Section 4 of [RFC6780].

6.1.3. Association Source

The Association Source field in the ASSOCIATION object is set to a valid IP address to identify the node that originates the association. In case of dynamic associations, the association source is usually set as the local PCEP speaker address, unless local policy dictates otherwise, in which case association source is set based on the local policy. In case of PCE redundancy, local policy could set the source as a virtual IP address which identifies all instances of the PCE. In case of operator-configured association, the association source is manually configured and it could be set as one of the PCEP speakers, Network Management System (NMS), or any other valid IP address that scopes the association identifier for the association type.

6.1.4. Unique Identification for an Association Group

The combination of the mandatory fields Association type, Association ID and Association Source in the ASSOCIATION object uniquely identify the association group. If the optional TLVs - Global Association Source or Extended Association ID are included, then they MUST be included in combination with mandatory fields to uniquely identify the association group. In this document, all these fields are collectively called 'association parameters'. Note that the ASSOCIATION object MAY include other optional TLVs (not defined in this document) based on the association types, that provides 'information' related to the association type, this document uses the term 'association information' for it.

6.2. Relationship with the RSVP ASSOCIATION

The format of PCEP ASSOCIATION Object defined in this document is aligned with the RSVP ASSOCIATION object ([RFC6780]). Various Association types related to RSVP association are defined in [RFC4872], [RFC4873], and [RFC7551]. The PCEP extensions that define new association types, should clarify how the PCEP associations would work with RSVP associations and vice-versa.

6.3. Object Encoding in PCEP messages

Message formats in this document are expressed using Reduced BNF (RBNF) as used in [RFC5440] and defined in [RFC5511].

6.3.1. Stateful PCEP messages

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

When carried in PCRpt message, it is used to report the association group membership pertaining to a LSP to a stateful PCE. The PCRpt message are used for both initial state synchronization operations (Section 5.6 of [RFC8231]) as well as whenever the state of the LSP changes. If the LSP belongs to an association group, then the associations MUST be included during the state synchronization operations.

The PCRpt message can also be used to remove an LSP from one or more association groups by setting the R flag to 1 in the ASSOCIATION object.

When an LSP is first reported to the PCE, the PCRpt message MUST include all the association groups that it belongs to. Any subsequent report message SHOULD include only the associations that are being modified or removed.

The PCRpt message is defined in [RFC8231] 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:

```
<path> ::= <intended-path>
           [<actual-attribute-list><actual-path>]
           <intended-attribute-list>
```

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

When an LSP is delegated to a stateful PCE, the stateful PCE can create 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. A stateful PCE can also remove a delegated LSP from one or more association groups by setting the R flag to 1 in the ASSOCIATION object.

The PCUpd message SHOULD include the association groups that are being modified or removed, there is no need to include associations that remains unchanged.

The PCUpd message is defined in [RFC8231] 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:

```
<path> ::= <intended-path><intended-attribute-list>
```

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

Unless a PCEP speaker wants to delete an association from an LSP or make changes to the association, it does not need to carry the ASSOCIATION object in future stateful messages.

A PCE initiating a new LSP can also include the association groups that this LSP belongs to. This is done by including the ASSOCIATION Object in a PCInitiate message. The PCInitiate message MUST include all the association groups that it belongs to. The PCInitiate message is defined in [RFC8281] 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>]
```

6.3.2. Request Message

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 (and the other LSPs) should be known to the PCE beforehand. These could be operator-configured or dynamically learned before via stateful PCEP messages. The R flag in ASSOCIATION object within PCReq message MUST be set to 0 while sending and ignored on receipt.

The PCReq message is defined in [RFC5440] and updated in [RFC8231], 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 the LSP object MAY be present for the passive stateful PCE mode.

6.3.3. Reply Message

In case of passive (stateful or stateless) PCE, the ASSOCIATION Object is OPTIONAL and MAY be carried in the Path Computation Reply (PCRep) message with the NO-PATH object. The ASSOCIATION object in PCRep message indicates the association group that cause the PCE to fail to find a path.

The PCRep message is defined in [RFC5440] and updated in [RFC8231] , it is further updated below for association:

```
<PCRep Message> ::= <Common Header>
                    <response-list>
```

Where:

```
<response-list> ::= <response> [<response-list>]
```

```
<response> ::= <RP>
               [<LSP>]
               [<NO-PATH>]
               [<association-list>]
               [<attribute-list>]
               [<path-list>]
```

Where:

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

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

6.4. Processing Rules

Association groups can be operator-configured on the necessary PCEP speakers and the PCEP speakers can join the existing association groups. In addition, a PCC or a PCE can create association groups dynamically and the PCEP speaker can also report the associations to its peer via PCEP messages. The operator-configured associations are created via configurations (where all association parameters are manually set) and exist until explicitly removed via configurations. The PCEP speaker can add LSPs to these configured associations and carry this via stateful PCEP messages. The dynamic associations are created dynamically by the PCEP speaker (where all association parameters are populated dynamically). The association group is attached to the LSP state, and the association group exists till there is at least one LSP as part of the association. As described in Section 6.1.4, the association parameters are the combination of Association type, Association ID and Association Source as well as optional global source and extended association identifier, that uniquely identifies an association group. The information related to the association types encoded via the TLVs of a particular association type (not described in this document) are the association information (Section 6.1.4).

If a PCEP speaker does not recognize the ASSOCIATION object in the stateful message, it will return a PCErr message with Error-Type "Unknown Object" as described in [RFC5440]. In case of PCReq message, the PCE would react based on the P flag as per [RFC5440]. If a PCEP speaker understands the ASSOCIATION object but does not support the Association type, it MUST return a PCErr message with

Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 1 "Association type is not supported". If any association parameters are invalid in the ASSOCIATION object, the PCEP speaker would consider this as malformed object and handle it as malformed message [RFC5440]. On receiving a PCEP message with ASSOCIATION, if a PCEP speaker finds that too many LSPs belong to the association group, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 2 "Too many LSPs in the association group". If a PCEP speaker cannot handle a new association, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 3 "Too many association groups". These numbers MAY be set by operator or decided based on a local policy.

If a PCE peer is unwilling or unable to process the ASSOCIATION object in the stateful message, it MUST return a PCErr message with the Error-Type "Not supported object" and follow the relevant procedures described in [RFC5440]. In case of PCReq message, the PCE would react based on the P flag as per [RFC5440]. On receiving a PCEP message with ASSOCIATION, if a PCEP speaker could not add the LSP to the association group for any reason, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 7 "Cannot join the association group".

If a PCEP speaker receives an ASSOCIATION object for an operator-configured association and the association identifier is not in the operator-configured association range for the Association type and Association Source, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 8 "Association identifier not in range".

If a PCEP speaker receives ASSOCIATION in PCReq message, and the association is not known (association is not configured, or created dynamically, or learned from a PCEP peer), it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 4 "Association unknown".

If the association information (related to the association group as a whole) received from the peer does not match with the local operator-configured information, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 5 "Operator-configured association information mismatch". On receiving association information (related to the association group as a whole) that does not match with the association information previously received about the same association from a peer, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 6 "Association information mismatch". Note that information related to each LSP within the

association as part of the association information TLVs could be different.

If a PCEP speaker receives an ASSOCIATION object with the R bit set for removal, and the association group (identified by association parameters) is not known, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 4 "Association unknown".

The dynamic associations are cleared along with the LSP state information as per the [RFC8231]. 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 behaviours. Same procedure is also followed for the association groups. On session termination at the PCE, when the LSP state reported by PCC is cleared, the association groups are also cleared. When there are no LSPs in an association group, the association is considered to be empty and thus deleted.

In case the LSP is delegated to another PCE on session failure, the associations (and association information) set by the PCE remains intact, unless updated by the new PCE that takes over.

Upon LSP delegation revocation, the PCC MAY clear the association created by the PCE, but in order to avoid traffic loss, it SHOULD perform this in a make-before-break fashion (same as [RFC8231]).

7. IANA Considerations

IANA maintains the "Path Computation Element Protocol (PCEP) Numbers" registry at <<http://www.iana.org/assignments/pcep>>.

7.1. PCEP Object

The "PCEP Numbers" registry contains a subregistry "PCEP Objects". IANA is requested to confirm the early allocation of the following code point in the PCEP Objects registry.

Object-Class Value	Name	Reference
40 (Early allocation by IANA)	Association	[This.I-D]
	Object-Type	
	0: Reserved	
	1: IPv4	
	2: IPv6	

7.2. PCEP TLV

IANA is requested to confirm the early allocation of the following code point in the "PCEP TLV Type Indicators" registry.

Value	Meaning	Reference
29 (Early allocation by IANA)	Operator-configured Association Range	[This.I-D]
30 (Early allocation by IANA)	Global Association Source	[This.I-D]
31 (Early allocation by IANA)	Extended Association ID	[This.I-D]

IANA is requested to fix the meaning for value 31 in the above registry to 'Extended Association ID', it is currently mentioned as 'Extended Association Id'.

IANA is also requested to make a new assignment for the existing "PCEP TLV Type Indicators" registry as follows:

Value	Meaning	Reference
TBD	ASSOC-Type-List	[This.I-D]

7.3. Association Flags

This document requests IANA to create a subregistry of the "PCEP Numbers" for the bits carried in the Flags field of the ASSOCIATION object. The subregistry is called "ASSOCIATION Flags Field". New values are assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- o Bit number (counting from bit 0 as the most significant bit)
- o Capability description
- o Defining RFC

Bit	Description	Reference
15	R (Removal)	[This.I-D]

7.4. Association Type

This document requests IANA to create a subregistry of the "PCEP Numbers" for the Association Type field of the the ASSOCIATION object. The subregistry is called "ASSOCIATION Type Field". New values are to be assigned by Standards Action [RFC8126]. Each value should be tracked with the following qualities:

- o Type
- o Name
- o Reference

There are no association types specified in this document, future documents should request the assignment of association types from this subregistry.

7.5. PCEP-Error Object

IANA is requested to confirm the early allocation of the following code points within the "PCEP-ERROR Object Error Types and Values" sub-registry of the "PCEP Numbers" registry, as follows:

Error-Type	Meaning
26	Association Error [This.I-D]
(early	Error-value=0:
alloc by	Unassigned
IANA)	Error-value=1:
	Association type is not supported
	Error-value=2:
	Too many LSPs in the association group
	Error-value=3:
	Too many association groups
	Error-value=4:
	Association unknown
	Error-value=5:
	Operator-configured association
	information mismatch
	Error-value=6:
	Association information mismatch
	Error-value=7:
	Cannot join the association group
	Error-value=8:
	Association identifier not in range

8. Security Considerations

The security considerations described in [RFC8231] and [RFC5440] apply to the extensions described in this document as well. Additional considerations related to a malicious PCEP speaker are introduced, as associations could be spoofed and could be used as an attack vector. An attacker could attempt to create too many associations in an attempt to load the PCEP peer. The PCEP peer responds with PCErr as described in Section 6.4. An attacker could impact LSP operations by creating bogus associations. Further, association groups could provide an adversary with the opportunity to eavesdrop on the relationship between the LSPs. Thus securing the PCEP session using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525], is RECOMMENDED.

Much of the information carried in the ASSOCIATION object, as per this document is not extra sensitive. It often reflects information that can also be derived from the LSP Database, but association provides a much easier grouping of related LSPs and messages. Implementations and operators can and should use indirect values in ASSOCIATION as a way to hide any sensitive business relationships.

9. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440] and [RFC8231] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

9.1. Control of Function and Policy

A PCE or PCC implementation MUST allow operator-configured associations and SHOULD allow setting of the operator-configured association range (Section 3.4) as described in this document.

9.2. Information and Data Models

The PCEP YANG module is defined in [I-D.ietf-pce-pcep-yang]. In future, this YANG module should be extended or augmented to provide the following additional information relating to association groups.

An implementation SHOULD allow the operator to view the associations configured or created dynamically. Further implementation SHOULD allow to view associations reported by each peer, and the current set of LSPs in the association.

It might also be useful to find out how many associations for each association type currently exist and to know how many free association identifiers are available for a particular association type and source.

9.3. 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].

9.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440] and [RFC8231].

9.5. Requirements on Other Protocols

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

9.6. Impact on Network Operations

Mechanisms defined in [RFC5440] and [RFC8231] also apply to PCEP extensions defined in this document.

10. Acknowledgments

We would like to thank Yuji Kamite and Joshua George for their contributions to this document. Also thanks to Venugopal Reddy, Cyril Margaria, Rakesh Gandhi, Mike Koldychev, and Adrian Farrel for their useful comments.

We would like to thank Julien Meuric for shepherding this document and providing comments with text suggestions.

Thanks to Stig Venaas for the RTGDIR review comments.

Thanks to Alvaro Retana, Mirja Kuhlewind, Martin Vigoureux, Barry Leiba, Eric Vyncke, Suresh Krishnan, and Benjamin Kaduk for the IESG comments.

11. Contributors

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com

Xian Zhang
Huawei Technologies
F3-1-B RnD Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129
China

Email: zhang.xian@huawei.com

Mustapha Aissaoui
Nokia

Email: mustapha.aissaoui@nokia.com

12. References

12.1. 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>>.
- [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, <<https://www.rfc-editor.org/info/rfc5511>>.
- [RFC6780] Berger, L., Le Faucheur, F., and A. Narayanan, "RSVP ASSOCIATION Object Extensions", RFC 6780, DOI 10.17487/RFC6780, October 2012, <<https://www.rfc-editor.org/info/rfc6780>>.

- [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>>.
- [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>>.
- [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>>.

12.2. 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>>.
- [RFC4872] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, DOI 10.17487/RFC4872, May 2007, <<https://www.rfc-editor.org/info/rfc4872>>.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, DOI 10.17487/RFC4873, May 2007, <<https://www.rfc-editor.org/info/rfc4873>>.

- [RFC6007] Nishioka, I. and D. King, "Use of the Synchronization VECtor (SVEC) List for Synchronized Dependent Path Computations", RFC 6007, DOI 10.17487/RFC6007, September 2010, <<https://www.rfc-editor.org/info/rfc6007>>.
- [RFC7551] Zhang, F., Ed., Jing, R., and R. Gandhi, Ed., "RSVP-TE Extensions for Associated Bidirectional Label Switched Paths (LSPs)", RFC 7551, DOI 10.17487/RFC7551, May 2015, <<https://www.rfc-editor.org/info/rfc7551>>.
- [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>>.
- [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-12 (work in progress), July 2019.
- [I-D.ietf-pce-stateful-path-protection]
Ananthakrishnan, H., Sivabalan, S., Barth, C., Minei, I., and M. Negi, "PCEP Extensions for Associating Working and Protection LSPs with Stateful PCE", draft-ietf-pce-stateful-path-protection-07 (work in progress), June 2019.
- [I-D.ietf-pce-association-diversity]
Litkowski, S., Sivabalan, S., Barth, C., and M. Negi, "Path Computation Element Communication Protocol (PCEP) Extension for LSP Diversity Constraint Signaling", draft-ietf-pce-association-diversity-08 (work in progress), July 2019.

Appendix A. Example for Operator-configured Association Range

Consider an association type T1 (which allows both dynamic and operator-configured association with a default range of <0x1000, 0xffff>). Consider that, because of need of the network, the PCE needs to create more dynamic associations and would like to change the association range to <0xbffe, 0xffff> instead. During PCEP session establishment the PCE would advertise the new range, the PCC could skip advertising as the default values are used. If a PCC is creating a dynamic association (with PCC as association source) it needs to pick a free association identifier for type T1 in the range of <0x1, 0xffff> whereas if a PCE is creating a dynamic association (with PCE as association source) it needs to pick a free association

identifier from the range <0x1, 0xbffd>. Similarly if an operator-configured association is manually configured with the PCC as association source, it should be from the range <0x1000, 0xffff> whereas if the PCE is association source, it should be from <0xbffe, 0xffff>. In case the association source is not a PCEP peer (for example an NMS system), then the default range of <0x1000, 0xffff> is considered.

Authors' Addresses

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

Email: inaminei@google.com

Edward Crabbe
Individual Contributor

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
Netflix

Email: hari@netflix.com

Dhruv Dhody
Huawei
Divyashree Techno Park, Whitefield
Bangalore, KA 560066
India

Email: dhruv.ietf@gmail.com

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

Email: yosuke.tanaka@ntt.com