PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 5, 2015

U. Palle
D. Dhody
Huawei Technologies
Y. Tanaka
Y. Kamite
NTT Communications
Z. Ali
Cisco Systems

July 4, 2014

Path Computation Element (PCE) Protocol Extensions for Stateful PCE usage for Point-to-Multipoint Traffic Engineering Label Switched Paths
draft-palle-pce-stateful-pce-p2mp-04

Abstract

The Path Computation Element (PCE) has been identified as an appropriate technology for the determination of the paths of point-to-multipoint (P2MP) TE LSPs. [I-D.ietf-pce-stateful-pce-app] presents several use cases, demonstrating scenarios that benefit from the deployment of a stateful PCE. [I-D.ietf-pce-stateful-pce] provides the fundamental PCE communication Protocol (PCEP) extensions needed to support stateful PCE functions. This memo provides extensions required for PCEP so as to enable the usage of a stateful PCE capability in supporting point-to-multipoint (P2MP) TE LSPs.

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 January 5, 2015.

Table of Contents

1. Introduction

As per [RFC4655], the Path Computation Element (PCE) is an entity that is capable of computing a network path or route based on a network graph, and applying computational constraints. A Path Computation Client (PCC) may make requests to a PCE for paths to be computed.

[RFC4857] describes how to set up point-to-multipoint (P2MP) Traffic Engineering Label Switched Paths (TE LSPs) for use in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks. The PCE has been identified as a suitable application for the computation of paths for P2MP TE LSPs ([RFC5671]).

The PCEP is designed as a communication protocol between PCCs and PCEs for point-to-point (P2P) path computations and is defined in [RFC5440]. The extensions of PCEP to request path computation for P2MP TE LSPs are described in [RFC6006].

Stateful PCEs are shown to be helpful in many application scenarios, in both MPLS and GMPLS networks, as illustrated in [I-D.ietf-pce-stateful-pce-app]. These scenarios apply equally to P2P and P2MP TE LSPs. [I-D.ietf-pce-stateful-pce] provides the fundamental extensions needed for stateful PCE to support general functionality for P2P TE LSP. Complementarily, this document focuses on the extensions that are necessary in order for the deployment of stateful PCEs to support P2MP TE LSPs.

1.1. 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].

2. Terminology

Terminology used in this document is same as terminology used in [I-D.ietf-pce-stateful-pce] and [RFC6006].

3. Supporting P2MP TE LSP for Stateful PCE

3.1. Motivation

[I-D.ietf-pce-stateful-pce-app] presents several use cases, demonstrating scenarios that benefit from the deployment of a stateful PCE including optimization, recovery, etc which are equally applicable to P2MP TE LSPs. [I-D.ietf-pce-stateful-pce] defines the extensions to PCEP for P2P TE LSPs. Complementarily, this document focuses on the extensions that are necessary in order for the deployment of stateful PCEs to support P2MP TE LSPs.

In addition to that, the stateful nature of a PCE simplifies the information conveyed in PCEP messages since it is possible to refer to the LSPs via PLSP-ID. For P2MP this is an added advantage, where the size of message is much larger. Incase of stateless PCE, a modification of P2MP tree requires encoding of all leaves along with the paths in PCReq message, but using a stateful PCE with P2MP capability, the PCEP message can be used to convey only the modifications (the other information can be retrieved from the P2MP LSP identifier).

3.2. Objectives

The objectives for the protocol extensions to support P2MP TE LSP for stateful PCE are same as the objectives described in section 3.2 of [I-D.ietf-pce-stateful-pce].

4. Functions to Support P2MP TE LSPs for Stateful PCEs

[I-D.ietf-pce-stateful-pce] specifies new functions to support a stateful PCE. It also specifies that a function can be initiated either from a PCC towards a PCE (C-E) or from a PCE towards a PCC (E-C).

This document extends these functions to support P2MP TE LSPs.

Capability Advertisement (E-C,C-E): both the PCC and the PCE must announce during PCEP session establishment that they support PCEP Stateful PCE extensions for P2MP using mechanisms defined in Section 5.2.

LSP State Synchronization (C-E): after the session between the PCC and a stateful PCE with P2MP capability is initialized, the PCE must learn the state of a PCC's P2MP TE LSPs before it can perform path computations or update LSP attributes in a PCC.

LSP Update Request (E-C): a stateful PCE with P2MP capability requests modification of attributes on a PCC's P2MP TE LSP.

LSP State Report (C-E): a PCC sends an LSP state report to a PCE whenever the state of a P2MP TE LSP changes.

LSP Control Delegation (C-E,E-C): a PCC grants to a PCE the right to update LSP attributes on one or more P2MP TE LSPs; the PCE becomes the authoritative source of the LSP's attributes as long as the delegation is in effect (See Section 5.5 of [I-D.ietf-pce-stateful-pce]); the PCC may withdraw the delegation or the PCE may give up the delegation at any time.

An update to [I-D.sivabalan-pce-disco-stateful] is needed to support autodiscovery of stateful PCEs with P2MP capability.

5. Architectural Overview of Protocol Extensions

5.1. Extension of PCEP Messages

New PCEP messages are defined in [I-D.ietf-pce-stateful-pce] to support stateful PCE for P2P TE LSPs. In this document these messages are extended to support P2MP TE LSPs.

Path Computation State Report (PCRpt): Each P2MP TE LSP State Report in a PCRpt message can contain actual P2MP TE LSP path attributes, LSP status, etc. An LSP State Report carried on a PCRpt message is also used in delegation or revocation of control of a P2MP TE LSP to/from a PCE. The extension of PCRpt message is described in Section 7.1.

Path Computation Update Request (PCUpd): Each P2MP TE LSP Update Request in a PCUpd message MUST contain all LSP parameters that a PCE wishes to set for a given P2MP TE LSP. An LSP Update Request carried on a PCUpd message is also used to return LSP delegations if at any point PCE no longer desires control of a P2MP TE LSP. The PCUpd message is described in Section 7.2.

5.2. Capability Advertisement

During PCEP Initialization Phase, as per Section 7.1.1 of [I-D.ietf-pce-stateful-pce], PCEP speakers advertises Stateful capability via Stateful PCE Capability TLV in open message. A new flag is defined for the STATEFUL-PCE-CAPABILITY TLV defined in [I-D.ietf-pce-stateful-pce]. Its format is shown in the following figure:

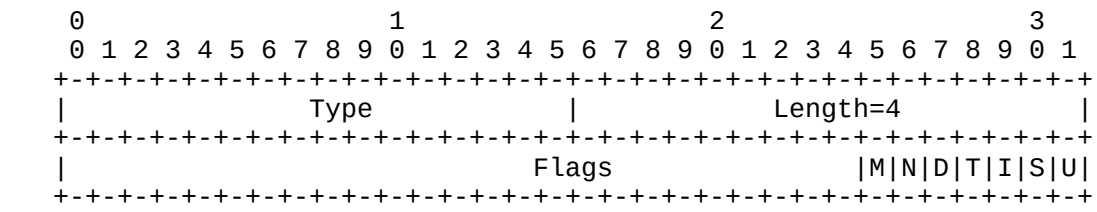

Figure 1: STATEFUL-PCE-CAPABILITY TLV Format

The U (LSP-UPDATE-CAPABILITY) bit is defined in [I-D.ietf-pce-stateful-pce]. The I (LSP-INSTANTIATION-CAPABILITY) bit is defined in [I-D.ietf-pce-pce-initiated-lsp]. The S (INCLUDE-DB-VERSION), T (TRIGGERED-SYNC) and D (DELTA-LSP-SYNC-CAPABILITY) bits are defined in [I-D.ietf-pce-stateful-sync-optimizations]. A new bit N (P2MP-CAPABILITY) and M (P2MP-LSP-UPDATE-CAPABILITY) are added in this document:

N (P2MP-CAPABILITY - 1 bit): if set to 1 by a PCC, the N Flag indicates that the PCC is willing to send P2MP LSP State Reports whenever P2MP LSP parameters or operational status changes.; if set to 1 by a PCE, the N Flag indicates that the PCE is interested in receiving LSP State Reports whenever LSP parameters or operational status changes. The P2MP-CAPABILITY Flag must be advertised by both a PCC and a PCE for PCRpt messages P2MP extension to be allowed on a PCEP session.

M (P2MP-LSP-UPDATE-CAPABILITY - 1 bit): if set to 1 by a PCC, the M Flag indicates that the PCC allows modification of P2MP LSP parameters; if set to 1 by a PCE, the M Flag indicates that the PCE is capable of updating P2MP LSP parameters. The P2MP-LSP-UPDATE-CAPABILITY Flag must be advertised by both a PCC and a PCE for PCUpd messages P2MP extension to be allowed on a PCEP session.

A PCEP speaker should continue to advertise the basic P2MP capability via mechanisms as described in [RFC6006].

5.3. State Synchronization

State Synchronization operations described in Section 5.4 of [I-D.ietf-pce-stateful-pce] are applicable for P2MP TE LSPs as well.

5.4. LSP Delegation

LSP delegation operations described in Section 5.5 of [I-D.ietf-pce-stateful-pce] are applicable for P2MP TE LSPs as well.

5.5. LSP Operations

5.5.1. Passive Stateful PCE

LSP operations for passive stateful PCE described in Section 5.6.1 of [I-D.ietf-pce-stateful-pce] are applicable for P2MP TE LSPs as well.

The Path Computation Request and Response message format for P2MP TE LSPs is described in Section 3.4 and Section 3.5 of [RFC6006] respectively.

The Request and Response message for P2MP TE LSPs are extended to support encoding of LSP object, so that it is possible to refer to a LSP with a unique identifier and simplify the PCEP message exchange. For example, incase of modification of one leaf in a P2MP tree, there should be no need to carry the full P2MP tree in PCReq message.

The extension for the Request and Response message for passive stateful operations on P2MP TE LSPs are described in Section 7.3 and Section 7.4.

5.5.2. Active Stateful PCE

LSP operations for active stateful PCE described in Section 5.6.2 of [I-D.ietf-pce-stateful-pce] are applicable for P2MP TE LSPs as well.

6. PCEP Object Extensions

The PCEP TLV defined in this document is compliant with the PCEP TLV format defined in [RFC5440].

6.1. Extension of LSP Object

LSP Object is defined in Section 7.3 of [I-D.ietf-pce-stateful-pce]. It specifies PLSP-ID to uniquely identify an LSP that is constant for the life time of a PCEP session. Similarly for P2MP tunnel, PLSP-ID identify a P2MP TE LSP uniquely. This document adds the following flags to the LSP Object:

N (P2MP bit): If the bit is set to 1, it specifies the message is for P2MP TE LSP which MUST be set in PCRpt or PCUpd message for a P2MP TE LSP.

F (Fragmentation bit): If the bit is set to 1, it specifies the message is fragmented.

If P2MP bit is set, the following P2MP-LSP-IDENTIFIER TLV MUST be present in LSP object.

6.2. P2MP-LSP-IDENTIFIER TLV

The P2MP LSP Identifier TLV MUST be included in the LSP object in PCRpt message for RSVP-signaled P2MP TE LSPs. If the TLV is missing, the PCE will generate an error with error-type 6 (mandatory object missing) and error-value TBD (12) (P2MP-LSP-IDENTIFIERS TLV missing) and close the PCEP session.

The P2MP LSP Identifier TLV MAY be included in the LSP object in PCUpd message for RSVP-signaled P2MP TE LSPs. The special value of all zeros for this TLV is used to refer to all paths pertaining to a particular PLSP-ID.

There are two P2MP LSP Identifier TLVs, one for IPv4 and one for IPv6.

The format of the IPV4-P2MP-LSP-IDENTIFIER TLV is shown in the following figure:

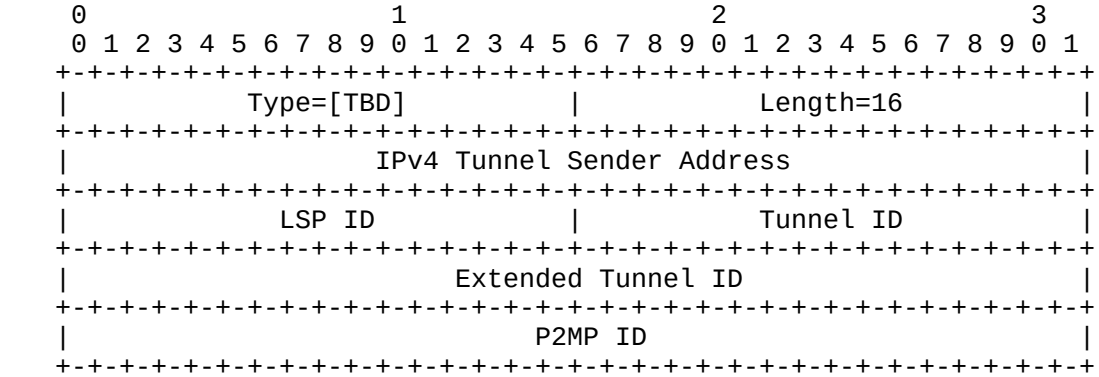

Figure 2: IPV4-P2MP-LSP-IDENTIFIER TLV format

The type of the TLV is [TBD] and it has a fixed length of 12 octets. The value contains the following fields:

IPv4 Tunnel Sender Address: contains the sender node's IPv4 address, as defined in [RFC3209], Section 4.6.2.1 for the LSP_TUNNEL_IPv4 Sender Template Object.

LSP ID: contains the 16-bit 'LSP ID' identifier defined in [RFC3209], Section 4.6.2.1 for the LSP_TUNNEL_IPv4 Sender Template Object.

Tunnel ID: contains the 16-bit 'Tunnel ID' identifier defined in [RFC3209], Section 4.6.1.1 for the LSP_TUNNEL_IPv4 Session Object. Tunnel ID remains constant over the life time of a tunnel.

Extended Tunnel ID: contains the 32-bit 'Extended Tunnel ID' identifier defined in [RFC3209], Section 4.6.1.1 for the LSP_TUNNEL_IPv4 Session Object.

P2MP ID: contains the 32-bit 'P2MP ID' identifier defined in Section 19.1.1 of [RFC4875]for the P2MP LSP Tunnel IPv4 SESSION Object.

The format of the IPV6-P2MP-LSP-IDENTIFIER TLV is shown in the following figure:

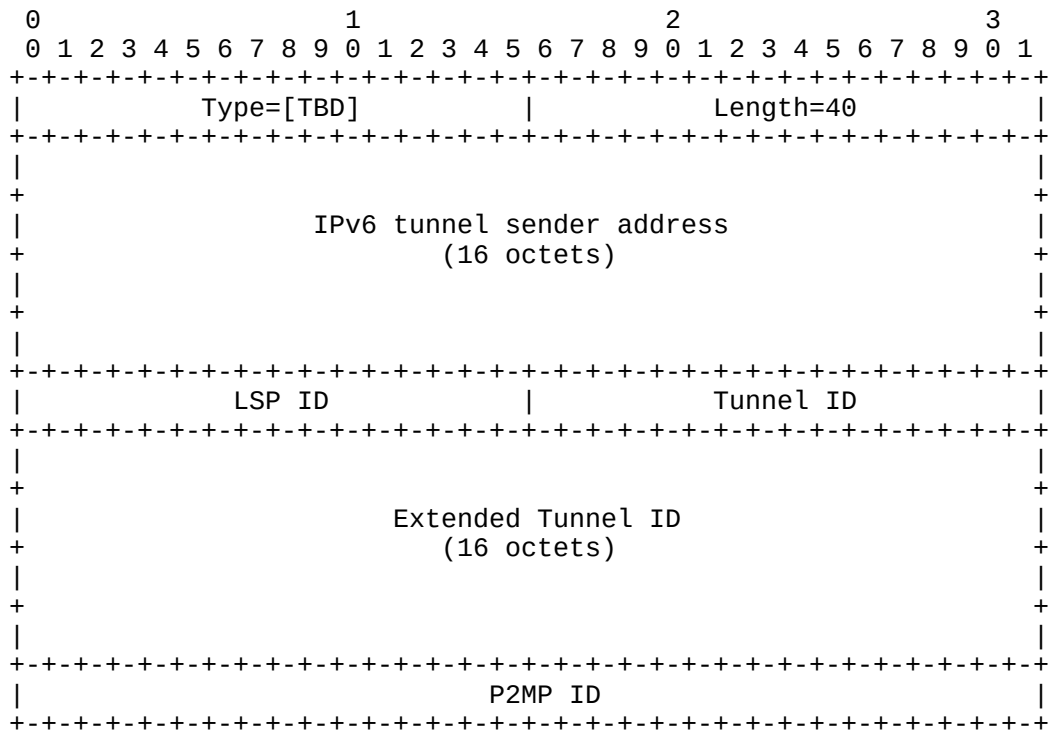

Figure 3: IPV6-P2MP-LSP-IDENTIFIER TLV format

The type of the TLV is [TBD] and it has a fixed length of 24 octets. The value contains the following fields:

IPv6 Tunnel Sender Address: contains the sender node's IPv6 address, as defined in [RFC3209], Section 4.6.2.2 for the LSP_TUNNEL_IPv6 Sender Template Object.

LSP ID: contains the 16-bit 'LSP ID' identifier defined in [RFC3209], Section 4.6.2.2 for the LSP_TUNNEL_IPv6 Sender Template Object.

Tunnel ID: contains the 16-bit 'Tunnel ID' identifier defined in [RFC3209], Section 4.6.1.2 for the LSP_TUNNEL_IPv6 Session Object. Tunnel ID remains constant over the life time of a tunnel.

Extended Tunnel ID: contains the 128-bit 'Extended Tunnel ID' identifier defined in [RFC3209], Section 4.6.1.2 for the LSP_TUNNEL_IPv6 Session Object.

P2MP ID: As defined above in IPV4-P2MP-LSP-IDENTIFIERS TLV.

6.3. S2LS Object

The S2LS (Source-to-Leaves) Object is used to report RSVP state of one or more destiantions (leaves) encoded within the END-POINTS object for a P2MP TE LSP. It MUST be carried in PCRpt message along with END-POINTS object when N bit is set in LSP object.

The format of the S2LS object is shown in the following figure:

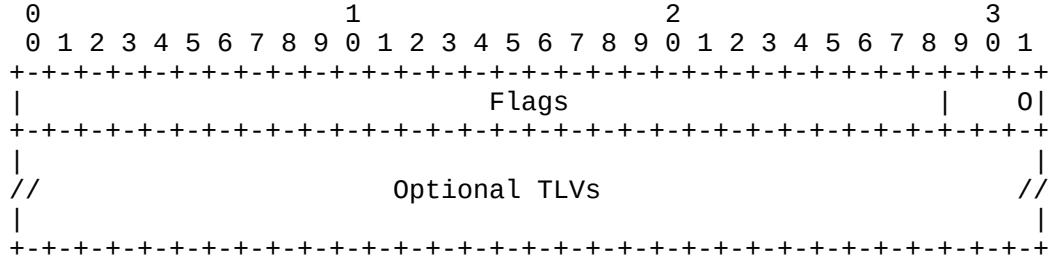

Figure 4: S2LS object format

Flags(32 bits):

O(Operational - 3 bits) the O Field represents the operational status of the group of destinations. The values are as per Operational field in LSP object defined in Section 7.3 of [I-D.ietf-pce-stateful-pce].

When N bit is set in LSP object then the O field in LSP object represents the operational status of the full P2MP TE LSP and the O field in S2LS object represents the operational status of a group of destinations encoded within the END-POINTS object.

Optional TLVs that may be included in the S2LS Object.

7. PCEP Message Extensions

7.1. The PCRpt Message

As per Section 6.1 of [I-D.ietf-pce-stateful-pce], PCRpt message is used to report the current state of a P2P TE LSP. This document extends the PCRpt message in reporting the status of P2MP TE LSP.

The format of PCRpt message is as follows:

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

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

<state-report> ::= [<SRP>]
                    <LSP>
                    <end-point-path-pair-list>
                    <attribute-list>
Where:

<end-point-path-pair-list>::=
                    [<END-POINTS>]
                    [<S2LS>]
                    <path>
                    [<end-point-path-pair-list>]

<path> ::= (<ERO>|<SERO>)
           [<RRO>]
           [<path>]
```

<attribute-list> is defined in [RFC5440] and extended by PCEP extensions.

The P2MP END-POINTS object defined in [RFC6006]is mandatory for specifying address of P2MP leaves grouped based on leaf types.

- New leaves to add (leaf type = 1)
- Old leaves to remove (leaf type = 2)
- Old leaves whose path can be modified/reoptimized (leaf type = 3)
- Old leaves whose path must be left unchanged (leaf type = 4)

When reporting the status of a P2MP TE LSP, the destinations are grouped in END-POINTS object based on the operational status (O field in S2LS object) and leaf type (in END-POINTS). This way the leaves that share the same operational status are grouped together. For reporing the status of delegated P2MP TE LSP, leaf-type = 3, where as for non-delegated P2MP TE LSP, leaf-type = 4 is used.

For delegated P2MP TE LSP configuration changes are reported via PCRpt message. For example, adding of new leaves END-POINTS (leaf-

type = 1) is used where as removing of old leaves (leaf-type = 2) is used.

Note that we preserve compatibility with the [I-D.ietf-pce-stateful-pce] definition of <state-report>. At least one instance of <END-POINTS> MUST be present in this message.

[Editor Note: suggest to add <END-POINTS> object mandatory in [I-D.ietf-pce-stateful-pce] document for <state-report>].

During state synchronization, the PCRpt message must report the status of the full P2MP TE LSP.

7.2. The PCUpd Message

As per Section 6.2 of [I-D.ietf-pce-stateful-pce], PCUpd message is used to update P2P TE LSP attributes. This document extends the PCUpd message in updating the attributes of P2MP TE LSP.

The format of a PCUpd message is as follows:

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

Where:

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

<update-request> ::= <SRP>
                     <LSP>
                     <end-point-path-pair-list>

<attribute-list>
Where:

<end-point-path-pair-list>::=
                [<END-POINTS>]
                <path>
                [<end-point-path-pair-list>]

<path> ::= (<ERO>|<SERO>)
           [<path>]

<attribute-list> is defined in [RFC5440] and
extended by PCEP extensions.
```

Note that we preserve compatibility with the [I-D.ietf-pce-stateful-pce] definition of <update-request>.

The PCC MAY use the make-before-break or sub-group-based procedures described in [RFC4875] based on a local policy decision.

7.3. The PCReq Message

As per Section 3.4 of [RFC6006], PCReq message is used for a P2MP path computation request. This document extends the PCReq message such that a PCC MAY include the LSP object in the PCReq message if the stateful PCE P2MP capability has been negotiated on a PCEP session between the PCC and a PCE.

The format of PCReq message is as follows:

```
<PCReq Message>::= <Common Header>
                   <request>
where:
<request>::= <RP>
             <end-point-rro-pair-list>
             [<LSP>]
             [<OF>]
             [<LSPA>]
             [<BANDWIDTH>]
             [<metric-list>]
             [<IRO>]
             [<LOAD-BALANCING>]

where:
<end-point-rro-pair-list>::=<END-POINTS>[<RRO-List>][<BANDWIDTH>]
                            [<end-point-rro-pair-list>]

<RRO-List>::=<RRO>[<BANDWIDTH>][<RRO-List>]
<metric-list>::=<METRIC>[<metric-list>]
```

7.4. The PCRep Message

As per Section 3.5 of [RFC6006], PCRep message is used for a P2MP path computation reply. This document extends the PCRep message such that a PCE MAY include the LSP object in the PCRep message if the stateful PCE P2MP capability has been negotiated on a PCEP session between the PCC and a PCE.

The format of PCRep message is as follows:

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

<response>::=<RP>
             [<end-point-path-pair-list>]
             [<NO-PATH>]
             [<attribute-list>]

where:

<end-point-path-pair-list>::=
                [<END-POINTS>]<path>[<end-point-path-pair-list>]

<path> ::= (<ERO>|<SERO>) [<path>]

<attribute-list>::=[<LSP>]
                   [<OF>]
                   [<LSPA>]
                   [<BANDWIDTH>]
                   [<metric-list>]
                   [<IRO>]
```

7.5. Example

7.5.1. P2MP TE LSP Update Request

LSP Update Request message is sent by an active stateful PCE to update the P2MP TE LSP parameters or attributes. An example of a PCUpd message for P2MP TE LSP is described below:

```
Common Header
SRP
LSP with P2MP flag set
END-POINTS for leaf type 3
  ERO list
```

In this example, a stateful PCE request updation of path taken by some of the leaves in a P2MP tree. The update request uses the END-POINT type 3 (modified/reoptimized). The ERO list represents the S2LS path after modification. The update message does not need to encode the full P2MP tree in this case.

7.5.2. P2MP TE LSP Report

LSP State Report message is sent by a PCC to report or delegate the P2MP TE LSP. An example of a PCRpt message for a delegated P2MP TE LSP is described below to add new leaves to an existing P2MP TE LSP:

```
Common Header
LSP with P2MP flag set
END-POINTS for leaf type 1
  S2LS (O=DOWN)
  ERO list (empty)
```

An example of a PCRpt message for P2MP TE LSP is described below to prune leaves from an existing P2MP TE LSP:

```
Common Header
LSP with P2MP flag set
END-POINTS for leaf type 2
  S2LS (O=UP)
  ERO list
```

An example of a PCRpt message for a delegated P2MP TE LSP is described below to report status of leaves in an existing P2MP TE LSP:

```
Common Header
LSP with P2MP flag set
END-POINTS for leaf type 3
  S2LS (O=UP)
  ERO list
END-POINTS for leaf type 3
  S2LS (O=DOWN)
  ERO list
```

An example of a PCRpt message for a non-delegated P2MP TE LSP is described below to report status of leaves:

```
Common Header
LSP with P2MP flag set
END-POINTS for leaf type 4
  S2LS (O=ACTIVE)
  ERO list
END-POINTS for leaf type 4
  S2LS (O=DOWN)
  ERO list
```

7.6. Report and Update Message Fragmentation

The total PCEP message length, including the common header, is 16 bytes. In certain scenarios the P2MP report and update request may not fit into a single PCEP message (initial report or update). The

F-bit is used in the LSP object to signal that the initial report or update was too large to fit into a single message and will be fragmented into multiple messages. In order to identify the single report or update, each message will use the same PLSP-ID.

Fragmentation procedure described below for report or update message is similar to [RFC6006]which describes request and response message fragmentation.

7.6.1. Report Fragmentation Procedure

If the initial report is too large to fit into a single report message, the PCC will split the report over multiple messages. Each message sent to the PCE, except the last one, will have the F-bit set in the LSP object to signify that the report has been fragmented into multiple messages. In order to identify that a series of report messages represents a single report, each message will use the same PLSP-ID.

7.6.2. Update Fragmentation Procedure

Once the PCE computes and updates a path for some or all leaves in a P2MP TE LSP, an update message is sent to the PCC. If the update is too large to fit into a single update message, the PCE will split the update over multiple messages. Each update message sent by the PCE, except the last one, will have the F-bit set in the LSP object to signify that the update has been fragmented into multiple messages. In order to identify that a series of update messages represents a single update, each message will use the same PLSP-ID and SRP-ID-number.

[Editor Note: P2MP message fragmentation errors associated with a P2MP path report and update will be defined in future version].

8. Non-Support of P2MP TE LSPs for Stateful PCE

The PCEP protocol extensions described in this document for stateful PCEs with P2MP capability MUST NOT be used if PCE has not advertised its stateful capability with P2MP as per Section 5.2. If this is not the case and Stateful operations on P2MP TE LSPs are attempted, then a PCErr with error-type 19 (Invalid Operation) and error-value TBD needs to be generated.

If a Stateful PCE receives a P2MP TE LSP report message and it understands the P2MP flag in the LSP object, but the stateful PCE is not capable of P2MP computation, the PCE MUST send a PCErr message with error-type 19 (Invalid Operation) and error-value TBD.

If a Stateful PCE receives a P2MP TE LSP report message and the PCE does not understand the P2MP flag in the LSP object, and therefore the PCEP extensions described in this document, then the PCE SHOULD reject the request.

[Editor Note: more information on exact error value is needed]

9. Security Considerations

The stateful operations on P2MP TE LSP are more CPU-intensive and also utilize more link bandwidth. In the event of an unauthorized stateful P2MP operations, or a denial of service attack, the subsequent PCEP operations may be disruptive to the network. Consequently, it is important that implementations conform to the relevant security requirements of [RFC5440], [RFC6006] and [I-D.ietf-pce-stateful-pce].

10. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC6006] and [I-D.ietf-pce-stateful-pce] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

10.1. Control of Function and Policy

A PCE or PCC implementation MUST allow configuring the stateful PCEP capability and the LSP Update capability for P2MP LSPs.

10.2. Information and Data Models

The PCEP MIB module SHOULD be extended to include advertised P2MP stateful capabilities, P2MP synchronization status, and P2MP delegation status etc.

10.3. 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].

10.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440], [RFC6006] and [I-D.ietf-pce-stateful-pce].

10.5. Requirements On Other Protocols

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

10.6. 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], [RFC6006] and [I-D.ietf-pce-stateful-pce].

11. IANA Considerations

This document requests IANA actions to allocate code points for the protocol elements defined in this document. Values shown here are suggested for use by IANA.

11.1. STATEFUL-PCE-CAPABILITY TLV

The following values are defined in this document for the Flags field in the STATEFUL-PCE-CAPABILITY-TLV in the OPEN object:

Bit	Description	Reference
27	P2MP-CAPABILITY	This.I-D
26	P2MP-LSP-UPDATE-CAPABILITY	This.I-D

11.2. Extension of LSP Object

This document requests that a registry is created to manage the Flags field of the LSP object. New values are to be assigned by Standards Action [RFC5226]. Each bit should be tracked with the following qualities:

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

The following values are defined in this document:

Bit	Description	Reference
24	P2MP	This.I-D
23	Fragmentation	This.I-D

11.3. Extension of PCEP-Error Object

A new error types 6 and 19 defined in section 8.4 of [I-D.ietf-pce-stateful-pce]. This document extend the new Error-Values for those error types for the following error conditions:

Error-Type	Meaning
6	Mandatory Object missing Error-value=12: P2MP-LSP-IDENTIFIER TLV missing
19	Invalid Operation Error-value= TBD.

11.4. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs:

Value	Meaning	Reference
22	P2MP-IPV4-LSP-IDENTIFIERS	This.I-D
23	P2MP-IPV6-LSP-IDENTIFIERS	This.I-D

12. Acknowledgments

Thanks to Quintin Zhao and Venugopal Reddy for his comments.

13. References

13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[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-09 (work in progress), June 2014.

13.2. Informative References

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC4857] Fogelstroem, E., Jonsson, A., and C. Perkins, "Mobile IPv4 Regional Registration", RFC 4857, June 2007.

[RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5671] Yasukawa, S. and A. Farrel, "Applicability of the Path Computation Element (PCE) to Point-to-Multipoint (P2MP) MPLS and GMPLS Traffic Engineering (TE)", RFC 5671, October 2009.

[RFC6006] Zhao, Q., King, D., Verhaeghe, F., Takeda, T., Ali, Z., and J. Meuric, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", RFC 6006, September 2010.

[I-D.ietf-pce-stateful-pce-app]
Zhang, X. and I. Minei, "Applicability of a Stateful Path Computation Element (PCE)", draft-ietf-pce-stateful-pce-app-02 (work in progress), June 2014.

[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-01 (work in progress), June 2014.

[I-D.ietf-pce-stateful-sync-optimizations]
Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", draft-ietf-pce-stateful-sync-optimizations-01 (work in progress), June 2014.

[I-D.sivabalan-pce-disco-stateful]
Sivabalan, S., Medved, J., and X. Zhang, "IGP Extensions for Stateful PCE Discovery", draft-sivabalan-pce-disco-stateful-03 (work in progress), January 2014.

Authors' Addresses

Udayasree Palle
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: udayasree.palle@huawei.com

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
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

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

EMail: y.kamite@ntt.com

Zafar Ali
Cisco Systems

EMail: zali@cisco.com