

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

PCEP Extensions for PCE-initiated Point-to-Multipoint LSP Setup in a  
Stateful PCE Model  
draft-palle-pce-stateful-pce-initiated-p2mp-lsp-03

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. The extensions described in [I-D.ietf-pce-stateful-pce] provide stateful control of Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSP) via PCE communication Protocol (PCEP), for a model where the Path Computation Client (PCC) delegates control over one or more locally configured LSPs to the PCE. Further [I-D.ietf-pce-pce-initiated-lsp] describes the creation and deletion of PCE-initiated LSPs under the stateful PCE model. This document provides extensions required for PCEP so as to enable the usage of a stateful PCE initiation capability in recommending point-to-multipoint (P2MP) TE LSP instantiation.

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.

## Copyright Notice

Copyright (c) 2014 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 . . . . .	3
1.1. Requirements Language . . . . .	3
2. Terminology . . . . .	4
3. Architectural Overview . . . . .	4
3.1. Motivation . . . . .	4
3.2. Operation Overview . . . . .	4
4. Support of PCE Initiated P2MP TE LSPs . . . . .	5
5. PCE-initiated P2MP TE LSP Operations . . . . .	5
5.1. The PCInitiate message . . . . .	5
5.2. P2MP TE LSP Instantiation . . . . .	6
5.3. P2MP TE LSP Deletion . . . . .	7
5.4. Adding and Pruning Leaves for the P2MP TE LSP . . . . .	7
5.5. P2MP TE LSP Delegation and Cleanup . . . . .	7
6. PCInitiate Message Fragmentation . . . . .	7
6.1. PCInitiate Fragmentation Procedure . . . . .	8
7. Non-Support of P2MP TE LSP Instantiation for Stateful PCE . . . . .	8
8. Security Considerations . . . . .	8
9. Manageability Considerations . . . . .	8
9.1. Control of Function and Policy . . . . .	9
9.2. Information and Data Models . . . . .	9
9.3. Liveness Detection and Monitoring . . . . .	9
9.4. Verify Correct Operations . . . . .	9
9.5. Requirements On Other Protocols . . . . .	9
9.6. Impact On Network Operations . . . . .	9
10. IANA Considerations . . . . .	9
10.1. STATEFUL-PCE-CAPABILITY TLV . . . . .	9
11. Acknowledgments . . . . .	10
12. References . . . . .	10
12.1. Normative References . . . . .	10
12.2. Informative References . . . . .	10

## 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. Further [I-D.palle-pce-stateful-pce-p2mp] focuses on the extensions that are necessary in order for the deployment of stateful PCEs to support P2MP TE LSPs. It includes mechanisms to effect P2MP LSP state synchronization between PCCs and PCEs, delegation of control of P2MP LSPs to PCEs, and PCE control of timing and sequence of P2MP path computations within and across PCEP sessions and focuses on a model where P2MP LSPs are configured on the PCC and control over them is delegated to the PCE.

[I-D.ietf-pce-pce-initiated-lsp] provides the fundamental extensions needed for stateful PCE-initiated P2P TE LSP recommended instantiation.

This document describes the setup, maintenance and teardown of PCE-initiated P2MP LSPs under the stateful PCE model, without the need for local configuration on the PCC, thus allowing for a dynamic network that is centrally controlled and deployed.

### 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], [I-D.ietf-pce-pce-initiated-lsp] and [RFC6006].

## 3. Architectural Overview

### 3.1. Motivation

[I-D.palle-pce-stateful-pce-p2mp] provides stateful control over P2MP TE LSPs that are locally configured on the PCC. This model relies on the Ingress taking an active role in delegating locally configured P2MP TE LSPs to the PCE, and is well suited in environments where the P2MP TE LSP placement is fairly static. However, in environments where the P2MP TE LSP placement needs to change in response to application demands, it is useful to support dynamic creation and tear down of P2MP TE LSPs. The ability for a PCE to trigger the creation of P2MP TE LSPs on demand can be seamlessly integrated into a controller-based network architecture, where intelligence in the controller can determine when and where to set up paths.

Section 3 of [I-D.ietf-pce-pce-initiated-lsp] further describes the motivation behind the PCE-Initiation capability, which are equally applicable for P2MP TE LSPs.

### 3.2. Operation Overview

A PCC or PCE indicates its ability to support PCE provisioned dynamic P2MP LSPs during the PCEP Initialization Phase via mechanism described in Section 4.

As per section 5.1 of [I-D.ietf-pce-pce-initiated-lsp], the PCE sends a Path Computation LSP Initiate Request (PCInitiate) message to the PCC to suggest instantiation or deletion of a P2P TE LSP. This document extends the PCInitiate message to support P2MP TE LSP (see details in Section 5.1).

P2MP TE LSP suggested instantiation and deletion operations are same as P2P LSP as described in section 5.3 and 5.4 of [I-D.ietf-pce-pce-initiated-lsp]. This document focuses on extensions needed for further handling of P2MP TE LSP (see details in Section 5.2).

4. Support of PCE Initiated P2MP TE LSPs

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:

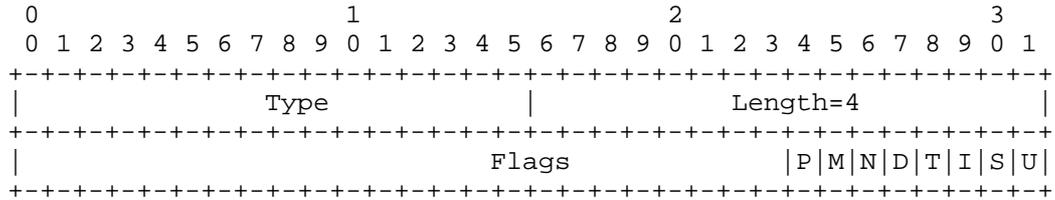


Figure 1: STATEFUL-PCE-CAPABILITY TLV Format

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]. The N (P2MP-CAPABILITY) and M (P2MP-LSP-UPDATE-CAPABILITY) bits are defined in [I-D.palle-pce-stateful-pce-p2mp]. A new bit P (P2MP-LSP-INSTANTIATION-CAPABILITY) is added in this document:

P (P2MP-LSP-INSTANTIATION-CAPABILITY - 1 bit): If set to 1 by a PCC, the P Flag indicates that the PCC allows suggested instantiation of an P2MP LSP by a PCE. If set to 1 by a PCE, the P flag indicates that the PCE will suggest P2MP LSP instantiation. The P2MP-LSP-INSTANTIATION-CAPABILITY flag must be set by both PCC and PCE in order to support PCE-initiated P2MP LSP instantiation.

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

5. PCE-initiated P2MP TE LSP Operations

5.1. The PCInitiate message

As defined in section 5.1 of [I-D.ietf-pce-pce-initiated-lsp], PCE sends a PCInitiate message to a PCC to recommend instantiation of a P2P TE LSP, this document extends the format of PCInitiate message for the creation of P2MP TE LSPs but the creation and deletion operations of P2MP TE LSP are same to the P2P TE LSP.

The format of PCInitiate message is as follows:

```
<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-point-path-pair-list>
                                       [<attribute-list>]
```

```
<PCE-initiated-lsp-deletion> ::= <SRP>
                                   <LSP>
```

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.

The PCInitiate message with an LSP object with N bit (P2MP) set is used to convey operation on a P2MP TE LSP. The SRP object is used to correlate between initiation requests sent by the PCE and the error reports and state reports sent by the PCC as described in [I-D.ietf-pce-stateful-pce].

## 5.2. P2MP TE LSP Instantiation

The Instantiation operation of P2MP TE LSP is same as defined in section 5.3 of [I-D.ietf-pce-pce-initiated-lsp] including handling of PLSP-ID, SYMBOLIC-PATH-NAME etc. Rules of processing and error codes remains unchanged. Further, as defined in section 6.1 of [I-D.palle-pce-stateful-pce-p2mp], N bit MUST be set in LSP object in

PCInitiate message by PCE to specify the instantiation is for P2MP TE LSP and the PCC or PCE MUST follow the mechanism defined in [I-D.palle-pce-stateful-pce-p2mp] for delegation and updation of P2MP TE LSPs.

Though N bit is set in the LSP object, P2MP-LSP-IDENTIFIER TLV defined in section 6.2 of [I-D.palle-pce-stateful-pce-p2mp] MUST NOT be included in the LSP object in PCInitiate message as it SHOULD be generated by PCC and carried in PCRpt message.

### 5.3. P2MP TE LSP Deletion

The deletion operation of P2MP TE LSP is same as defined in section 5.4 of [I-D.ietf-pce-pce-initiated-lsp] by sending an LSP Initiate Message with an LSP object carrying the PLSP-ID of the LSP to be removed and an SRP object with the R flag set (LSP-REMOVE as per section 5.2 of [I-D.ietf-pce-pce-initiated-lsp]). Rules of processing and error codes remains unchanged.

### 5.4. Adding and Pruning Leaves for the P2MP TE LSP

Adding of new leaves and Pruning of old Leaves for the PCE initiated P2MP TE LSP MUST be carried in PCUpd message and SHOULD refer [I-D.palle-pce-stateful-pce-p2mp] for P2MP TE LSP extensions. As defined in [RFC6006], leaf type = 1 for adding of new leaves, leaf type = 2 for pruning of old leaves of P2MP END-POINTS Object are used in PCUpd message.

PCC MAY use the Incremental State Update mechanisms as described in [RFC4875] to signal adding and pruning of leaves.

### 5.5. P2MP TE LSP Delegation and Cleanup

P2MP TE LSP delegation and cleanup operations are same as defined in section 6 of [I-D.ietf-pce-pce-initiated-lsp]. Rules of processing and error codes remains unchanged.

## 6. PCInitiate Message Fragmentation

The total PCEP message length, including the common header, is 16 bytes. In certain scenarios the P2MP LSP Initiate may not fit into a single PCEP message (initial PCInitiate message). The F-bit is used in the LSP object to signal that the initial PCInitiate was too large to fit into a single message and will be fragmented into multiple messages.

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

#### 6.1. PCInitiate Fragmentation Procedure

Once the PCE initiates to set up the P2MP TE LSP, a PCInitiate message is sent to the PCC. If the PCInitiate is too large to fit into a single PCInitiate message, the PCE will split the PCInitiate over multiple messages. Each PCInitiate message sent by the PCE, except the last one, will have the F-bit set in the LSP object to signify that the PCInitiate has been fragmented into multiple messages. In order to identify that a series of PCInitiate messages represents a single Initiate, each message will use the same PLSP-ID (in this case 0) and SRP-ID-number.

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

#### 7. Non-Support of P2MP TE LSP Instantiation for Stateful PCE

The PCEP protocol extensions described in this document for PCC or PCE with instantiation capability for P2MP TE LSPs MUST NOT be used if PCC or PCE has not advertised its stateful capability with Instantiation and P2MP capability as per Section 4. If this is not the case and Stateful initiation operations on P2MP TE LSPs are attempted, then a PCErr with error-type 19 (Invalid Operation) and error-value TBD needs to be generated.

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

#### 8. 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], [I-D.ietf-pce-stateful-pce] and [I-D.ietf-pce-pce-initiated-lsp].

#### 9. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC6006], [I-D.ietf-pce-stateful-pce] and [I-D.ietf-pce-pce-initiated-lsp] 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 configuring the stateful Initiation capability for P2MP LSPs.

### 9.2. Information and Data Models

The PCEP MIB module SHOULD be extended to include advertised P2MP stateful PCE-Initiation capability etc.

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

### 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 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].

## 10. 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.

### 10.1. 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
25	P2MP-LSP- INSTANTIATION- CAPABILITY	This.I-D

## 11. Acknowledgments

Thanks to Quintin Zhao and Venugopal Reddy for his comments.

## 12. References

### 12.1. 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.
- [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.palle-pce-stateful-pce-p2mp]  
Palle, U., Dhody, D., Tanaka, Y., Kamite, Y., and Z. Ali, "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-03 (work in progress), June 2014.

### 12.2. Informative References

- [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.
- [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-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.

## 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