

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: April 24, 2014

H. Chen  
Huawei Technologies  
A. Liu  
Ericsson  
F. Xu  
Verizon  
M. Toy  
Comcast  
V. Liu  
China Mobile  
October 21, 2013

Extensions to PCEP for Distributing Label Cross Domains  
draft-chen-pce-label-x-domains-00.txt

Abstract

This document specifies extensions to PCEP for distributing labels crossing domains for an inter-domain Point-to-Point (P2P) or Point-to-Multipoint (P2MP) Traffic Engineering (TE) Label Switched Path (LSP).

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 April 24, 2014.

Copyright Notice

Copyright (c) 2013 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
2. Terminology . . . . .	3
3. Conventions Used in This Document . . . . .	4
4. Label Distribution . . . . .	4
4.1. An Exmaple . . . . .	4
5. Extensions to PCEP . . . . .	5
5.1. RP Object Extension . . . . .	5
5.2. Label Object . . . . .	6
5.3. LSP Tunnel Object . . . . .	7
5.4. Request Message Extension . . . . .	9
5.5. Reply Message Extension . . . . .	9
6. Procedures . . . . .	10
6.1. Distributing Label in Ordered Setup . . . . .	10
6.2. Distributing Label in Path Computation . . . . .	10
7. Security Considerations . . . . .	11
8. IANA Considerations . . . . .	11
8.1. Request Parameter Bit Flags . . . . .	11
9. Acknowledgement . . . . .	11
10. References . . . . .	11
10.1. Normative References . . . . .	11
10.2. Informative References . . . . .	12
Authors' Addresses . . . . .	12

## 1. Introduction

After a path crossing multiple domains is computed, an inter-domain Traffic Engineering (TE) Label Switched Path (LSP) tunnel may be set up along the path by a number of tunnel central controllers (TCCs). Each of the domains through which the path goes may be controlled by a tunnel central controller (TCC), which sets up the segment of the TE LSP tunnel in the domain. When the TCC sets up the segment of the TE LSP tunnel in its domain that is not a domain containing the tail end of the tunnel, it needs a label from a domain, which is next to it along the path.

This document specifies extensions to PCEP and various procedures for distributing a label from a domain to its previous domain along the path for the TE LSP tunnel crossing multiple domains.

## 2. Terminology

ABR: Area Border Router. Routers used to connect two IGP areas (areas in OSPF or levels in IS-IS).

ASBR: Autonomous System Border Router. Routers used to connect together ASes of the same or different service providers via one or more inter-AS links.

Boundary Node (BN): a boundary node is either an ABR in the context of inter-area Traffic Engineering or an ASBR in the context of inter-AS Traffic Engineering.

Entry BN of domain(n): a BN connecting domain(n-1) to domain(n) along a determined sequence of domains.

Exit BN of domain(n): a BN connecting domain(n) to domain(n+1) along a determined sequence of domains.

Inter-area TE LSP: A TE LSP that crosses an IGP area boundary.

Inter-AS TE LSP: A TE LSP that crosses an AS boundary.

LSP: Label Switched Path.

LSR: Label Switching Router.

PCC: Path Computation Client. Any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element. An entity (component, application, or

network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCE(i) is a PCE with the scope of domain(i).

TED: Traffic Engineering Database.

This document uses terminologies defined in RFC5440.

### 3. Conventions Used in This Document

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.

### 4. Label Distribution

The Label Distribution may be provided by the PCE-based path computation. A PCE responsible for a domain computes a path segment for the domain, which is from an entry boundary to an exit boundary (or an egress) node of the domain. The PCE gets an label from the entry boundary node and adds an label object containing the label in the reply message to be sent to the requesting PCC (or another PCE).

When a PCE or PCC receives a reply message containing an label object, it removes the object from the message. The PCE may store the information in the label object or send the information to another component such as a Tunnel Central Controller (TCC).

#### 4.1. An Exmaple

Figure 1 below illustrates a simple two-AS topology. There is a PCE responsible for the path computation in each AS. A path computation is requested from the Tunnel Central Controller (TCC), acting as the PCC, which sends the path computation request to PCE-1. PCE-1 is unable to compute an end-to-end path and invokes PCE-2 (possibly using the techniques described in [RFC5441]). PCE-2 computes a path segment from entry boundary node ASBR-2 of the right domain to the egress as {ASBR-2, C, D, Egress}. In addition to placing this path segment in the reply message to PCE-1, PCE-2 gets an label from the entry boundary node ASBR-2 and adds an label object containing the label and optionally the ASBR-2 into the reply message.




Figure 1: Example of Label Distribution

When PCE-1 receives the reply message containing the label object from PCE-2, it removes the object from the message. PCE-1 may store the information in the label object or send the information to another component such as a Tunnel Central Controller (TCC). TCC may set up the segment of the LSP tunnel from Ingress to ASBR-2 using the label in the label object from ASBR-2.

## 5. Extensions to PCEP

This section describes the extensions to PCEP for distributing labels crossing domains for an inter-domain Point-to-Point (P2P) or Point-to-Multipoint (P2MP) Traffic Engineering (TE) Label Switched Path (LSP). The extensions include the definition of a new flag in the RP object, tunnel information and label in a PCReq/PCRep message.

### 5.1. RP Object Extension

The following flags are added into the RP Object:

An L bit is added in the flag bits field of the RP object to tell a receiver of a PCReq/PCRep message that the message is for distributing labels crossing domains for an inter-domain LSP.

- o L (Label distribution bit - 1 bit):

- 0: This indicates that this is not a PCReq/PCRep message for distributing labels crossing domains.
- 1: This indicates that this is a PCReq or PCRep message for distributing labels crossing domains.

The IANA request is referenced in Section below (Request Parameter Bit Flags) of this document.

This L bit with the N bit defined in RFC6006 can indicate whether the PCReq/PCRep message is for distributing labels for an MPLS TE P2P LSP or an MPLS TE P2MP LSP.

- o L = 1 and N = 0: This indicates that this is a PCReq/PCRep message for distributing labels for a P2P LSP.
- o L = 1 and N = 1: This indicates that this is a PCReq/PCRep message for distributing labels for a P2MP LSP.


The C bit is added in the flag bits field of the RP object to tell the receiver of a PCReq/PCRep message that the message is for creating the segment of the LSP tunnel in a domain before distributing labels from this domain to its previous domain.

- o C (LSP tunnel Creation bit - 1 bit):
  - 0: This indicates that this is not a PCReq/PCRep message for creating the segment of the LSP tunnel.
  - 1: This indicates that this is a PCReq/PCRep message for creating the segment of the LSP tunnel in the domain before distributing labels to its previous domain.


The IANA request is referenced in Section below (Request Parameter Bit Flags) of this document.

## 5.2. Label Object

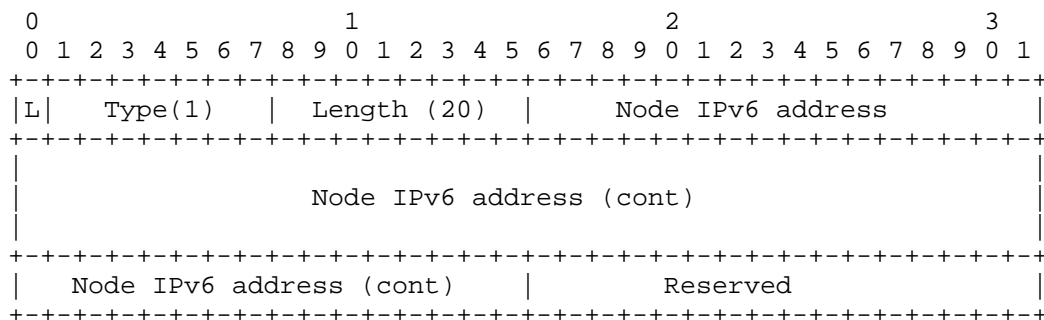
The format of a label object body (Object-Type=2) is illustrated below, which comprises a label and an optional node sub object. The node sub object contains a boundary node IP address, from which the label is allocated and distributed.



The format of the node IPv4 address sub object (Type=1) is as follows:



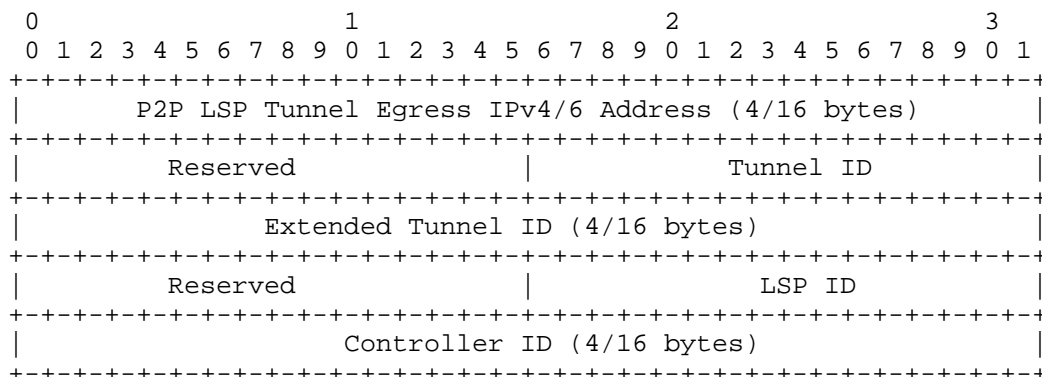
The format of the node IPv6 address sub object (Type=2) is illustrated below:



### 5.3. LSP Tunnel Object

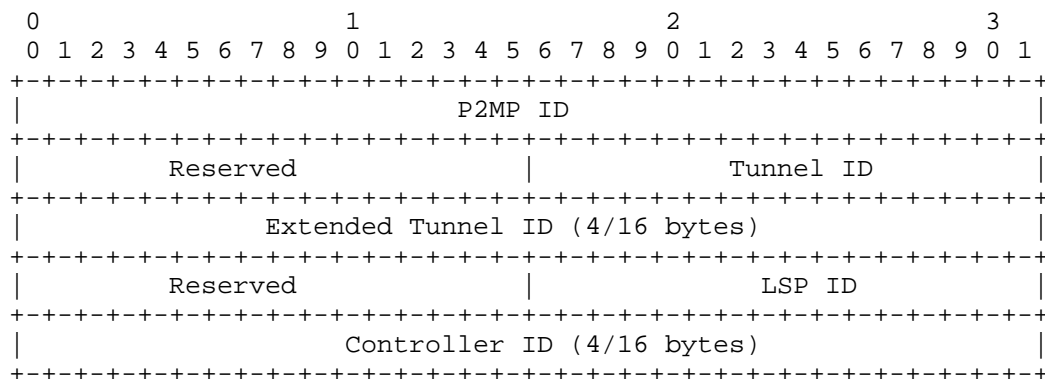
The LSP tunnel object contains the information that may be used to identify an LSP tunnel. An LSP tunnel may be a P2P or P2MP LSP tunnel. It may be an IPv4 or IPv6 LSP tunnel. Thus there are four types of LSP tunnels: 1) P2P LSP IPv4 tunnel, 2) P2P LSP IPv6 tunnel, 3) P2MP LSP IPv4 tunnel, and 4) P2MP LSP IPv6 tunnel.

The format of the P2P LSP IPv4/6 tunnel object body is as follows:



- o P2P LSP Tunnel Egress IPv4/6 Address:  
IPv4/6 address of the egress of the tunnel.
- o Tunnel ID:  
A 16-bit identifier that is constant over the life of the tunnel.
- o Extended Tunnel ID:  
A 4/16-byte identifier that is constant over the life of the tunnel.
- o LSP ID:  
A 16-bit identifier to allow resources sharing.
- o Controller ID:  
A 4/16-byte identifier for the controller responsible for the head segment of the tunnel.

The format of the P2MP LSP IPv4/6 tunnel object body is as follows:





- o P2MP ID:  
A 32-bit number unique within the ingress of LSP tunnel.
- o Tunnel ID:  
A 16-bit identifier that is constant over the life of the tunnel.
- o Extended Tunnel ID:  
A 4/16-byte identifier that is constant over the life of the tunnel.
- o LSP ID:  
A 16-bit identifier to allow resources sharing.
- o Controller ID:  
A 16-byte identifier for the controller responsible for the head segment of the tunnel.

#### 5.4. Request Message Extension

Figure below illustrates the format of a request message with a optional LSP tunnel object:

```

<PCReq Message> ::= <Common Header>
                    [<svec-list>]
                    <request-list>
<request-list> ::= <request> [<request-list>]
<request> ::= <RP> <END-POINTS> [<OF>] [<LSPA>] [<BANDWIDTH>]
               [<metric-list>] [<RRO> [<BANDWIDTH>]] [<IRO>]
               [<LOAD-BALANCING>]
               [<LSP-tunnel>]

```

Figure 2: Format for Request Message

#### 5.5. Reply Message Extension

Below is the format of a reply message with an optional Label object:

```

<PCReq Message> ::= <Common Header>
                    <response-list>
<response-list> ::= <response> [<response-list>]
<response> ::= <RP>
               [<NO-PATH>]
               [<attribute-list>]
               [<path-list>]
<path-list> ::= <path> [<path-list>]
<path> ::= <ERO> <attribute-list> [<LSP-tunnel>] [<Label>]

```

Figure 3: Format for Reply Message

## 6. Procedures

There may be a number of procedures for distributing labels crossing domains.

### 6.1. Distributing Label in Ordered Setup

Suppose that a path for an MPLS TE LSP tunnel crossing multiple domains is computed by PCEs and a sequence of domains ( $D_1, D_2, \dots, D_n$ ) through which the path goes are controlled by a sequence of Tunnel Central Controllers TCCs ( $TCC_1, TCC_2, \dots, TCC_n$ ) respectively. The method or procedure for distributing a label in ordered setup may comprise the following steps:

Step 1:  $TCC_i$  ( $i = 1, \dots, n-1$ ) sends  $TCC_j$  ( $j = i + 1$ ) a request for establishing the TE LSP tunnel.

Step 2:  $TCC_n$  (e.g.,  $TCC_3$ ) allocates a label from the enter border node (e.g., border node R) of domain  $D_n$  (e.g.,  $D_3$ ) and sends  $TCC_{n-1}$  (e.g.,  $TCC_2$ ) a reply containing the label after establishing the TE LSP tunnel segment (e.g., from node R to U) in domain  $D_n$  (e.g.,  $D_3$ ).

Step 3:  $TCC_j$  ( $j = n-1, \dots, 2$ ) receives a reply containing a first label from  $TCC_{j+1}$ , allocates a second label from the enter border node of domain  $D_j$ , establishes the TE LSP tunnel segment in  $D_j$  and sends  $TCC_i$  ( $i = j - 1$ ) a reply containing the label.

Step 4:  $TCC_1$  receives a reply containing a label from  $TCC_2$  and establishes the TE LSP tunnel segment in  $D_1$ . At this point, the TE LSP tunnel crossing multiple domains is established.

### 6.2. Distributing Label in Path Computation

Suppose that a path for an MPLS TE LSP tunnel crossing multiple domains is computed by PCEs and a sequence of domains ( $D_1, D_2, \dots, D_n$ ) through which the path goes are controlled by a sequence of PCEs ( $PCE_1, PCE_2, \dots, PCE_n$ ) as TCCs respectively. The method or procedure for distributing a label in path computation may comprise the following steps:

Step 1: After  $PCE_n$  (e.g.,  $PCE_3$ ) receives a path request for computing the path and determines that the path segment of the path in domain  $D_n$  (e.g.,  $D_3$ ) is on the best path, it allocates a label from the enter border node (e.g., R) of domain  $D_n$  (e.g.,  $D_3$ ) on the path, establishes the TE LSP tunnel segment in domain  $D_n$  and sends  $PCE_{n-1}$  (e.g.,  $PCE_2$ ) a path reply containing the label.

Step 2: When PCEj ( $j = n-1, \dots, 2$ ) receives a path reply containing a first label from PCEj+1 and determines that the path segment of the path in domain Dj (e.g., D2) is on the best path, it allocates a second label from the enter border node of domain Dj, establishes the TE LSP tunnel segment in Dj and sends PCEi ( $i = j - 1$ ) a path reply containing the second label.

Step 3: After PCE1 receives a path reply containing a label from PCE2 and determines the path segment in domain D1, it establishes the TE LSP tunnel segment in D1. At this point, the TE LSP tunnel crossing multiple domains is established.

## 7. Security Considerations

The mechanism described in this document does not raise any new security issues for the PCEP protocols.

## 8. IANA Considerations

This section specifies requests for IANA allocation.

### 8.1. Request Parameter Bit Flags

A new RP Object Flag has been defined in this document. IANA is requested to make the following allocation from the "PCEP RP Object Flag Field" Sub-Registry:

Bit	Description	Reference
18	Label Distribution (L-bit)	This I-D
19	LSP tunnel Creation (C-bit)	This I-D

## 9. Acknowledgement

The author would like to thank people for their valuable comments on this draft.

## 10. References

### 10.1. Normative References

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

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

## 10.2. Informative References

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC5862] Yasukawa, S. and A. Farrel, "Path Computation Clients (PCC) - Path Computation Element (PCE) Requirements for Point-to-Multipoint MPLS-TE", RFC 5862, June 2010.

## Authors' Addresses

Huaimo Chen  
Huawei Technologies  
Boston, MA  
US

Email: [huaimo.chen@huawei.com](mailto:huaimo.chen@huawei.com)

Autumn Liu  
Ericsson  
CA  
USA

Email: [autumn.liu@ericsson.com](mailto:autumn.liu@ericsson.com)

Fengman Xu  
Verizon  
2400 N. Glenville Dr  
Richardson, TX 75082  
USA

Email: fengman.xu@verizon.com

Mehmet Toy  
Comcast  
1800 Bishops Gate Blvd.  
Mount Laurel, NJ 08054  
USA

Email: mehmet\_toy@cable.comcast.com

Vic Liu  
China Mobile  
No.32 Xuanwumen West Street, Xicheng District  
Beijing, 100053  
China

Email: liuzhiheng@chinamobile.com

