

Workgroup: Internet Engineering Task Force
Internet-Draft:
draft-chen-pce-compute-backup-egress-21
Published: 12 January 2023
Intended Status: Standards Track
Expires: 16 July 2023

A H. Chen
uFuturewei
t
h
o
r
s
:

Extensions to the Path Computation Element Communication Protocol (PCEP) for Backup Egress of a Traffic Engineering Label Switched Path

Abstract

This document presents extensions to the Path Computation Element Communication Protocol (PCEP) for a PCC to send a request for computing a backup egress for an MPLS TE P2MP LSP or an MPLS TE P2P LSP to a PCE and for a PCE to compute the backup egress and reply to the PCC with a computation result for the backup egress.

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 16 July 2023.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
- [3. Conventions Used in This Document](#)
- [4. Extensions to PCEP](#)
 - [4.1. Backup Egress Capability Advertisement](#)
 - [4.1.1. Capability TLV in Existing PCE Discovery Protocol](#)
 - [4.1.2. Open Message Extension](#)
 - [4.2. Request and Reply Message Extension](#)
 - [4.2.1. RP Object Extension](#)
 - [4.2.2. External Destination Nodes](#)
 - [4.2.3. Constraints between Egress and Backup Egress](#)
 - [4.2.4. Constraints for Backup Path](#)
 - [4.2.5. Backup Egress Node](#)
 - [4.2.6. Backup Egress PCEP Error Objects and Types](#)
 - [4.2.7. Request Message Format](#)
 - [4.2.8. Reply Message Format](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
 - [6.1. Backup Egress Capability Flag](#)
 - [6.2. Backup Egress Capability TLV](#)
 - [6.3. Request Parameter Bit Flags](#)
 - [6.4. PCEP Objects](#)
- [7. Acknowledgement](#)
- [8. References](#)
 - [8.1. Normative References](#)
 - [8.2. Informative References](#)
- [Author's Address](#)

1. Introduction

RFC 4655 "A Path Computation Element-(PCE) Based Architecture" describes a set of building blocks for constructing solutions to compute Point-to-Point (P2P) Traffic Engineering (TE) label switched paths across multiple areas or Autonomous System (AS) domains. A typical PCE-based system comprises one or more path computation servers, traffic engineering databases (TED), and a number of path computation clients (PCC). A routing protocol is used to exchange traffic engineering information from which the TED is constructed. A PCC sends a Point-to-Point traffic engineering Label Switched Path (LSP) computation request to the path computation server, which uses the TED to compute the path and responses to the PCC with the computed path. A path computation server is named as a PCE. The communications between a PCE and a PCC for Point-to-Point label switched path computations follow the PCE communication protocol (PCEP).

RFC6006 "Extensions to PCEP for Point-to-Multipoint Traffic Engineering Label Switched Paths" describes extensions to PCEP to handle requests and responses for the computation of paths for P2MP TE LSPs.

This document defines extensions to the Path Computation Element Communication Protocol (PCEP) for a PCC to send a request for computing a backup egress node for an MPLS TE P2MP LSP or an MPLS TE P2P LSP to a PCE and for a PCE to compute the backup egress node and

reply to the PCC with a computation result for the backup egress node.

2. Terminology

This document uses terminologies defined in RFC5440, and RFC4875.

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 RFC 2119.

4. Extensions to PCEP

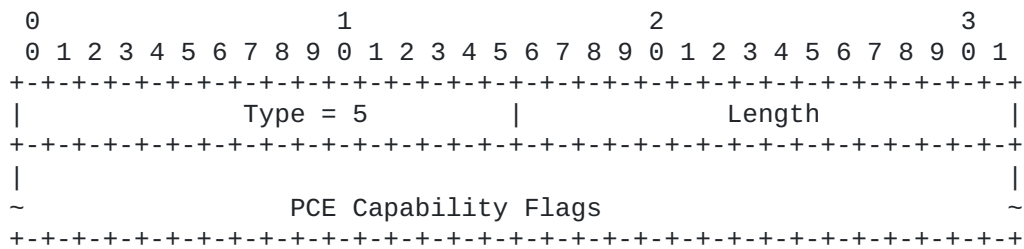
This section describes the extensions to PCEP for computing a backup egress of an MPLS TE P2MP LSP and an MPLS TE P2P LSP.

4.1. Backup Egress Capability Advertisement

4.1.1. Capability TLV in Existing PCE Discovery Protocol

An option for advertising a PCE capability for computing a backup egress for an MPLS TE P2MP LSP or an MPLS TE P2P LSP is to define two new flags. One new flag in the OSPF and IS-IS PCE Capability Flags indicates the capability that a PCE is capable to compute a backup egress for an MPLS TE P2MP LSP; and another new flag in the OSPF and IS-IS PCE Capability Flags indicates the capability that a PCE is capable to compute a backup egress for an MPLS TE P2P LSP.

The format of the PCE-CAP-FLAGS sub-TLV is as follows:



```

Type:      5
Length:    Multiple of 4 octets
Value:     This contains an array of units of 32-bit flags
           numbered from the most significant as bit zero, where
           each bit represents one PCE capability.

```

The following capability bits have been assigned by IANA:

Bit	Capabilities
0	Path computation with GMPLS link constraints
1	Bidirectional path computation
2	Diverse path computation
3	Load-balanced path computation
4	Synchronized path computation
5	Support for multiple objective functions
6	Support for additive path constraints (max hop count, etc.)
7	Support for request prioritization
8	Support for multiple requests per message
9	Global Concurrent Optimization (GCO)
10	P2MP path computation
11-31	Reserved for future assignments by IANA.

Reserved bits SHOULD be set to zero on transmission and MUST be ignored on receipt.

For the backup egress capabilities, one bit such as bit 13 may be assigned to indicate that a PCE is capable to compute a backup egress for an MPLS TE P2MP LSP and another bit such as bit 14 may be assigned to indicate that a PCE is capable to compute a backup egress for an MPLS TE P2P LSP as follows.

Bit	Capabilities
13	Backup egress computation for P2MP LSP
14	Backup egress computation for P2P LSP
15-31	Reserved for future assignments by IANA.

4.1.2. Open Message Extension

If a PCE does not advertise its backup egress computation capability during discovery, PCEP should be used to allow a PCC to discover, during the Open Message Exchange, which PCEs are capable of supporting backup egress computation.

To achieve this, we extend the PCEP OPEN object by defining a new optional TLV to indicate the PCE's capability to perform backup egress computation for an MPLS TE P2MP LSP and an MPLS TE P2P LSP.

We request IANA to allocate a value such as 8 from the "PCEP TLV Type Indicators" subregistry, as documented in Section below ("Backup Egress Capability TLV"). The description is "backup egress capable", and the length value is 2 bytes. The value field is set to indicate the capability of a PCE for backup egress computation for an MPLS TE LSP in details.

We can use flag bits in the value field in the same way as the PCE Capability Flags described in the previous section.

The inclusion of this TLV in an OPEN object indicates that the sender can perform backup egress computation for an MPLS TE P2MP LSP or an MPLS TE P2P LSP.

The capability TLV is meaningful only for a PCE, so it will typically appear only in one of the two Open messages during PCE session establishment. However, in case of PCE cooperation (e.g., inter-domain), when a PCE behaving as a PCC initiates a PCE session it SHOULD also indicate its path computation capabilities.

4.2. Request and Reply Message Extension

This section describes extensions to the existing RP (Request Parameters) object to allow a PCC to request a PCE for computing a backup egress of an MPLS TE P2MP LSP or an MPLS TE P2P LSP when the PCE receives the PCEP request.

4.2.1. RP Object Extension

The following flags are added into the RP Object:

The T bit is added in the flag bits field of the RP object to tell the receiver of the message that the request/reply is for computing a backup egress of an MPLS TE P2MP LSP and an MPLS TE P2P LSP.

- o T (Backup Egress bit - 1 bit):
 - 0: This indicates that this is not PCReq/PCRep for backup egress.
 - 1: This indicates that this is PCReq or PCRep message for backup egress.

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

This T bit with the N bit defined in RFC 6006 can indicate whether a request/reply is for a backup egress of an MPLS TE P2MP LSP or an MPLS TE P2P LSP.

- o T = 1 and N = 1: This indicates that this is a PCReq/PCRep message for backup egress of an MPLS TE P2MP LSP.
- o T = 1 and N = 0: This indicates that this is a PCReq/PCRep message for backup egress of an MPLS TE P2P LSP.

4.2.2. External Destination Nodes

In addition to the information about the path that an MPLS TE P2MP LSP or an MPLS TE P2P LSP traverses, a request message may comprise other information that may be used for computing the backup egress for the P2MP LSP or P2P LSP. For example, the information about an external destination node, to which data traffic is delivered from an egress node of the P2MP LSP or P2P LSP, is useful for computing a backup egress node.

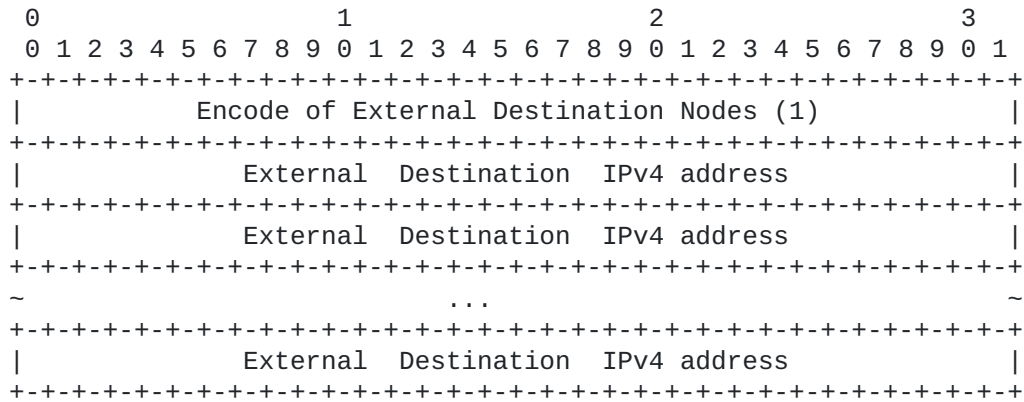
4.2.2.1. External Destination Nodes Object

The PCC can specify an external destination nodes (EDN) Object. In order to represent the external destination nodes efficiently, we define two types of encodes for the external destination nodes in the object.

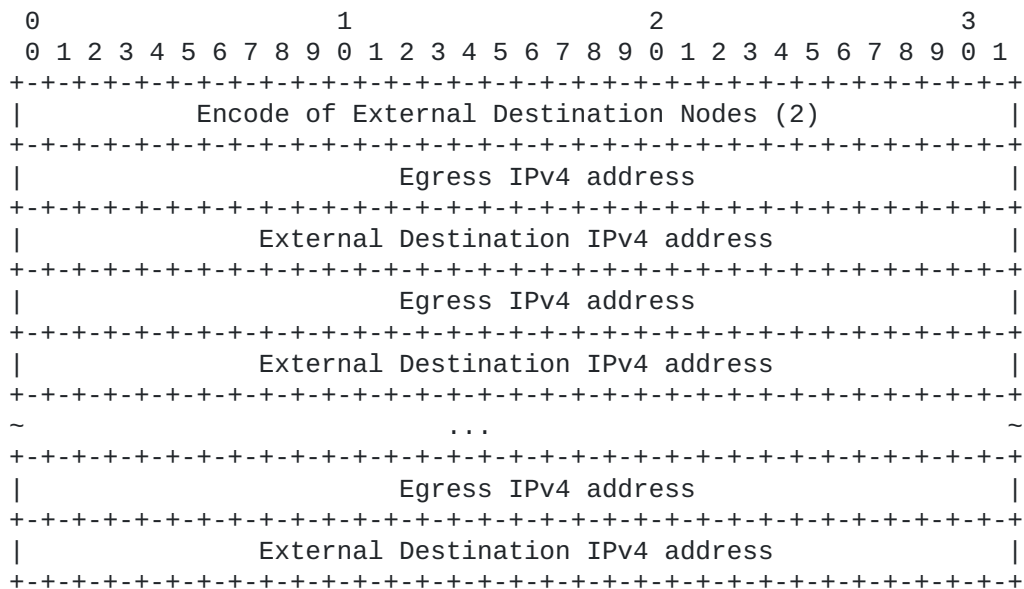
One encode indicates that the EDN object contains an external destination node for every egress node of an MPLS TE P2MP LSP or an MPLS TE P2P LSP. The order of the external destination nodes in the object is the same as the egress node(s) of the P2MP LSP or P2P LSP contained in the PCE messages.

Another encode indicates that the EDN object contains a list of egress node and external destination node pairs. For an egress node and external destination node pair, the data traffic is delivered to the external destination node from the egress node of the LSP.

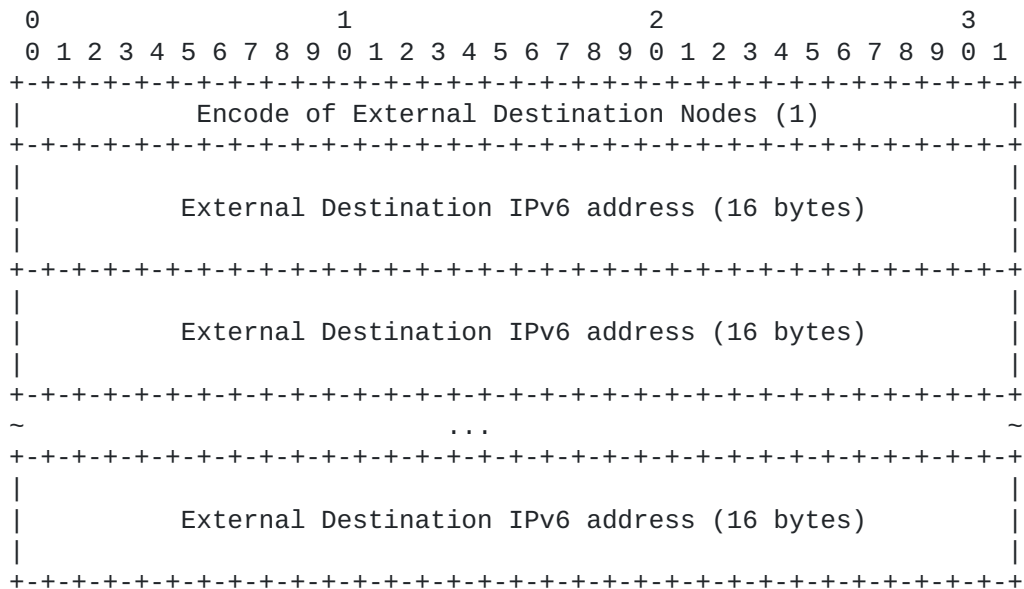
The format of the external destination nodes (EDN) object boby for IPv4 with the first type of encodes is illustrated as follows:




The format of the external destination nodes (EDN) object boby for IPv4 with the second type of encodes is illustrated below:



The format of the external destination nodes (EDN) object boby for IPv6 with the first type of encodes is illustrated as follows:



The format of the external destination nodes (EDN) object boby for IPv6 with the second type of encodes is illustrated below:



The object can only be carried in a PCReq message. A Path Request may carry at most one external destination nodes Object.

The Object-Class and Object-types will need to be allocated by IANA. The IANA request is documented in Section below (PCEP Objects).

4.2.2.2. New Type of END-POINTS Objects

Alternatively, we may use END-POINTS to represent external destination nodes in a request message for computing backup egress nodes of MPLS LSP.

The format of the external destination nodes (EDN) END-POINTS object body for IPv4 with the first type of encodes is illustrated as follows:


```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Compact External Destination Nodes Type (12)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               ...                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The new type of END-POINTS is Compact External Destination Nodes Type (12). The final value for the type will be assigned by IANA. The EDN END-POINTS object of type 12 contains an external destination node for every egress node of an MPLS TE P2MP LSP or an MPLS TE P2P LSP. The order of the external destination nodes in the object is the same as the egress node(s) of the P2MP LSP or P2P LSP contained in the PCE messages.

The format of the external destination nodes END-POINTS object boby for IPv4 with the second type of encodes is illustrated below:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination Nodes Type (13)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Egress IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Egress IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               ...                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Egress IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Destination IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The new type of END-POINTS is External Destination Nodes Type (13). The final value for the type will be assigned by IANA. The EDN END-POINTS object of type 13 contains a list of egress node and external destination node pairs. For an egress node and external destination node pair, the data traffic is delivered to the external destination node from the egress node of the LSP.

4.2.3. Constraints between Egress and Backup Egress

A request message sent to a PCE from a PCC for computing a backup egress of an MPLS TE P2MP LSP or an MPLS TE P2P LSP may comprise a

constraint indicating that there must be a path from the backup egress node to be computed to the egress node of the P2MP LSP or P2P LSP and that the length of the path is within a given hop limit such as one hop.

This constraint can be considered as default by a PCE or explicitly sent to the PCE by a PCC [TBD].

4.2.4. Constraints for Backup Path

A request message sent to a PCE from a PCC for computing a backup egress of a P2MP LSP or P2P LSP may comprise a constraint indicating that the backup egress node to be computed may not be a node on the P2MP LSP or P2P LSP. In addition, the request message may comprise a list of nodes, each of which is a candidate for the backup egress node.

A request message sent to a PCE from a PCC for computing a backup egress of a P2MP LSP or P2P LSP may comprise a constraint indicating that there must be a path from the previous hop node of the egress node of the P2MP LSP or P2P LSP to the backup egress node to be computed and that there is not an internal node of the path from the previous hop node of the egress node of the P2MP LSP or P2P LSP to the backup egress that is on the path of the P2MP LSP or P2P LSP.

Most of these constraints for the backup path can be considered as default by a PCE. The constraints for the backup path may be explicitly sent to the PCE by a PCC [TBD].

4.2.5. Backup Egress Node

The PCE may send a reply message to the PCC in return to the request message for computing a new backup egress node or a number of backup egress nodes. The reply message may comprise information about the computed backup egress node(s), which is contained in the path(s) from the previous-hop node of the egress node of the P2MP LSP or P2P LSP to the backup egress node(s) computed.

4.2.6. Backup Egress PCEP Error Objects and Types

In some cases, the PCE may not complete the backup egress computation as requested, for example based on a set of constraints. As such, the PCE may send a reply message to the PCC that indicates an unsuccessful backup egress computation attempt. The reply message may comprise a PCEP-error object, which may comprise an error-value, error-type and some detail information.

4.2.7. Request Message Format

The PCReq message is encoded as follows using RBNF as defined in [RFC5511].

Below is the message format for a request message:

```

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

```

where:

<EDNO> is an external destination nodes object.

The definitions for svec-list, RP, end-point-rro-pair-list, OF, LSPA, BANDWIDTH, metric-list, IRO, and LOAD-BALANCING are described in RFC5440 and RFC6006.

4.2.8. Reply Message Format

The PCRep message is encoded as follows using RBNF as defined in [RFC5511].

Below is the message format for a reply message:

```

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

```

The definitions for RP, NO-PATH, END-POINTS, OF, LSPA, BANDWIDTH, metric-list, IRO, and SERO are described in RFC5440, RFC6006 and RFC4875.

5. Security Considerations

The mechanism described in this document does not raise any new security issues for the PCEP, OSPF or IS-IS protocols.

6. IANA Considerations

This section specifies requests for IANA allocation.

6.1. Backup Egress Capability Flag

Two new OSPF Capability Flags are defined in this document to indicate the capabilities for computing a backup egress for an MPLS TE P2MP LSP and an MPLS TE P2P LSP. IANA is requested to make the assignment from the "OSPF Parameters Path Computation Element (PCE) Capability Flags" registry:

Bit	Description	Reference
13	Backup egress for P2MP LSP	This I-D
14	Backup egress for P2P LSP	This I-D

6.2. Backup Egress Capability TLV

A new backup egress capability TLV is defined in this document to allow a PCE to advertize its backup egress computation capability. IANA is requested to make the following allocation from the "PCEP TLV Type Indicators" sub-registry.

Value	Description	Reference
8	Backup egress capable	This I-D

6.3. 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
15	Backup egress (T-bit)	This I-D

6.4. PCEP Objects

An External Destination Nodes Object-Type is defined in this document. IANA is requested to make the following Object-Type allocation from the "PCEP Objects" sub-registry:

Object-Class Value	34
Name	External Destination Nodes
Object-Type	1: IPv4 2: IPv6 3-15: Unassigned
Reference	This I-D

7. Acknowledgement

The author would like to thank Ramon Casellas, Dhruv Dhody and Quintin Zhao for their valuable comments on this draft.

8. References

8.1. 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>>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.

[RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.

[RFC5440] Vasseur, JP., Ed., Le Roux, JL., Ed., and RFC Publisher, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.

[RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.

[RFC6006] Zhao, Q., Ed., King, D., Ed., 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, DOI 10.17487/RFC6006, September 2010, <<https://www.rfc-editor.org/info/rfc6006>>.

8.2. Informative References

[RFC4655] Farrel, A., Vasseur, J.-P., Ash, J., and RFC Publisher, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.

[RFC5862] Yasukawa, S. and A. Farrel, "Path Computation Clients (PCC) - Path Computation Element (PCE) Requirements for Point-to-Multipoint MPLS-TE", RFC 5862, DOI 10.17487/RFC5862, June 2010, <<https://www.rfc-editor.org/info/rfc5862>>.

Author's Address

Huaimo Chen
Futurewei
Boston, MA
United States of America

Email: Huaimo.chen@futurewei.com