

Workgroup: Internet Engineering Task Force
Internet-Draft:
draft-chen-pce-compute-backup-ingress-19
Published: 10 January 2022
Intended Status: Standards Track
Expires: 14 July 2022
Authors: H. Chen
Futurewei

**Extensions to Path Computation Element Communication Protocol (PCEP)
for Backup Ingress 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 ingress for an MPLS TE P2MP LSP or an MPLS TE P2P LSP to a PCE and for a PCE to compute the backup ingress and reply to the PCC with a computation result for the backup ingress.

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 14 July 2022.

Copyright Notice

Copyright (c) 2022 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 Ingress 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 Source Node](#)
 - [4.2.3. Constraints between Ingress and Backup Ingress](#)
 - [4.2.4. Constraints for Backup Path](#)
 - [4.2.5. Backup Ingress Node](#)
 - [4.2.6. Backup Ingress 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 Ingress Capability Flag](#)
 - [6.2. Backup Ingress 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

RFC4090 "Fast Reroute Extensions to RSVP-TE for LSP Tunnels" describes two methods to protect P2P LSP tunnels or paths at local repair points. The local repair points may comprise a number of intermediate nodes between an ingress node and an egress node along the path. The first method is a one-to-one backup method, where a detour backup P2P LSP for each protected P2P LSP is created at each potential point of local repair. The second method is a facility bypass backup protection method, where a bypass backup P2P LSP tunnel is created using MPLS label stacking to protect a potential failure point for a set of P2P LSP tunnels. The bypass backup tunnel can protect a set of P2P LSPs that have similar backup constraints.

RFC4875 "Extensions to RSVP-TE for P2MP TE LSPs" describes how to use the one-to-one backup method and facility bypass backup method

to protect a link or intermediate node failure on the path of a P2MP LSP.

However, there is no mention of locally protecting an ingress node failure in a protected P2MP LSP or P2P LSP.

The methods for protecting an ingress node of a P2MP LSP or P2P LSP may be classified into two categories.

A first category uses a backup P2MP LSP that is from a backup ingress node to the number of destination nodes for the P2MP LSP, and a backup P2P LSP that is from a backup ingress node to the destination node for the P2P LSP. The disadvantages of this class of methods include more network resource such as computer power and link bandwidth consumption since the backup P2MP LSP or P2P LSP is from the backup ingress node to the number of destination nodes or the destination respectively.

A second category uses a local P2MP LSP or P2P LSP for protecting the ingress of a P2MP LSP or P2P LSP locally. The local P2MP LSP is from a backup ingress node to the next hop nodes of the ingress of the P2MP LSP. The local P2P LSP is from a backup ingress node to the next hop node of the ingress of the P2P LSP. It is desirable to let PCE compute these backup ingress nodes.

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

2. Terminology

This document uses terminologies defined in RFC5440, RFC4090, 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 RFC2119.

4. Extensions to PCEP

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

4.1. Backup Ingress Capability Advertisement

4.1.1. Capability TLV in Existing PCE Discovery Protocol

There are a couple of options for advertising a PCE capability for computing a backup ingress for an MPLS TE P2MP LSP or an MPLS TE P2P LSP.

The first option is to define a new flag in the OSPF and ISIS PCE Capability Flags to indicate the capability that a PCE is capable to compute both a backup ingress for an MPLS TE P2MP LSP and a backup ingress for an MPLS TE P2P LSP.

The second option is to define two new flags. One new flag in the OSPF and ISIS PCE Capability Flags indicates the capability that a PCE is capable to compute a backup ingress for an MPLS TE P2MP LSP; and another new flag in the OSPF and ISIS PCE Capability Flags indicates the capability that a PCE is capable to compute a backup ingress for an MPLS TE P2P LSP.

This second option is preferred now.

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

Bit	Capabilities
11	Backup ingress computation for P2MP LSP
12	Backup ingress computation for P2P LSP
13-31	Reserved for future assignments by IANA.

4.1.2. Open Message Extension

If a PCE does not advertise its backup ingress 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 ingress 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 ingress 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 Ingress Capability TLV"). The description is "backup ingress capable", and the length value is 2 bytes. The value field is set to indicate the capability of a PCE for backup ingress 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 ingress 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 ingress 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 I 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 ingress of an MPLS TE P2MP LSP and an MPLS TE P2P LSP.

- o I (Backup Ingress bit - 1 bit):
 - 0: This indicates that this is not PCReq/PCRep for backup ingress.
 - 1: This indicates that this is PCReq or PCRep message for backup ingress.

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

This I bit with the N bit defined in RFC6006 can indicate whether the request/reply is for a backup ingress of an MPLS TE P2MP LSP or an MPLS TE P2P LSP.

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

4.2.2. External Source Node

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 ingress for the P2MP LSP or P2P LSP. For example, the information about an external source node, from which data traffic is delivered to the ingress node of the P2MP LSP or P2P LSP and transported to the egress node(s) via the P2MP LSP or P2P LSP, is useful for computing a backup ingress node.

The PCC can specify an external source node (ESN) Object. The ESN Object has the same format as the IRO object defined in [RFC5440] except that it only supports IPv4 and IPv6 prefix sub-objects.

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

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

Alternatively, we may use END-POINTS to represent an external source node in a request message for computing a backup ingress node of MPLS LSP.

To represent an external source node efficiently, we define a new type of END-POINTS objects for computing a backup ingress node of

MPLS LSP. The format of the new END-POINTS object body for IPv4 (Object-Type 3) is 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Source Type (11)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               External Source IPv4 address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The new type of END-POINTS is External Source Node Type (11). The final value for the type will be assigned by IANA. This new type of END-POINTS object contains an external source node IPv4 address.

4.2.3. Constraints between Ingress and Backup Ingress

A request message sent to a PCE from a PCC for computing a backup ingress 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 ingress node to be computed to the ingress 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 ingress of a P2MP LSP or P2P LSP may comprise a constraint indicating that the backup ingress 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 ingress node.

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

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

The backup ingress node is the root or head node of the backup path computed.

4.2.6. Backup Ingress PCEP Error Objects and Types

In some cases, the PCE may not complete the backup ingress 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 ingress 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>]
              [<ESNO>]
              [<IRO>]
              [<LOAD-BALANCING>]
```

where:

<ESNO> is an external source node 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 and IS-IS protocols.

6. IANA Considerations

This section specifies requests for IANA allocation.

6.1. Backup Ingress Capability Flag

Two new OSPF Capability Flags are defined in this document to indicate the capabilities for computing a backup ingress 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
11	Backup ingress for P2MP LSP	This I-D
12	Backup ingress for P2P LSP	This I-D

6.2. Backup Ingress Capability TLV

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

Value	Description	Reference
8	Backup ingress 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
16	Backup ingress (I-bit)	This I-D

6.4. PCEP Objects

An External Source Node 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	33
Name	External Source Node
Object-Type	1: IPv4 2: IPv6 3-15: Unassigned
Reference	This I-D

7. Acknowledgement

The author would like to thank Cyril Margaria, Ramon Casellas, Dhruv Dhody and Quintin Zhao for their valuable comments and suggestions 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>>.

[RFC5440]

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

[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>>.

[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., and J. Ash, "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