PCE Working Group
Internet-Draft
Intended status: Experimental
Expires: October 18, 2013

D. Dhody
U. Palle
V. Kondreddy
Huawei Technologies India Pvt Ltd

April 16, 2013

Supporting explicit inclusion or exclusion of abstract nodes for a subset of P2MP destinations in Path Computation Element Communication Protocol (PCEP).

draft-dhody-pce-pcep-p2mp-per-destination-04

Abstract

The ability to determine paths of point-to-multipoint (P2MP) Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering Label Switched Paths (TE LSPs) is one the key requirements for Path Computation Element (PCE). [RFC6006] and [PCE-P2MP-PROCEDURES] describes these mechanisms for intra and inter domain path computation via PCE.

This document describes the motivation and PCE communication Protocol (PCEP) extension for explicitly specifying abstract nodes for inclusion or exclusion for a subset of destinations during the Point to Multipoint (P2MP) path computation via PCE.

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 October 18, 2013.

Copyright Notice

Table of Contents

1. Introduction

The Path Computation Element (PCE) architecture is defined in [RFC4655]. [RFC5862] lay out the requirements for PCEP to support Point-to-Multipoint (P2MP) path computation. [RFC6006] describe an extension to PCEP to compute optimal constrained intra-domain (G)MPLS P2MP TE LSPs. [PCE-P2MP-PROCEDURES] describes the mechanism for inter-domain P2MP path computation.

[RFC6006] describe a PCE-based path computation procedure to compute optimal constrained (G)MPLS P2MP TE LSPs. It describes mechanism to specify branch nodes that can or cannot be used via Branch Node Capability (BNC) object (which only supports IPv4 and IPv6 prefix sub-objects and are applied to all destinations). This document explains the need to add the capability to explicitly specify any abstract nodes (not just branch nodes) for inclusion or exclusion for a subset of destinations.

[PCE-P2MP-PROCEDURES] describes the core-tree procedure for computing inter-domain P2MP tree. It assumes that, due to deployment and commercial limitations, the sequence of domains for a path (the path domain tree) will be known in advance. For a group of destination which belong to a destination domain, the domain-sequence needs to be encoded separately as described in [DOMAIN-SEQ]. The mechanism, as described in this document, of explicitly specifying abstract nodes for inclusion or exclusion for a subset of destinations can be used for this purpose, where abstract nodes are domains.

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

The following terminology is used in this document.

IRO: Include Route Object.

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.

PCEP: Path Computation Element Protocol.

P2MP: Point-to-Multipoint

P2P: Point-to-Point

RRO: Record Route Object

RSVP: Resource Reservation Protocol

TE LSP: Traffic Engineering Label Switched Path.

XRO: Exclude Route Object.

3. Motivation

3.1. Domain Sequence Tree in Inter Domain P2MP Path Computation

[PCE-P2MP-PROCEDURES] describes the core-tree procedure for inter-domain path computation. The procedure assumes that the sequence of domains for a path (the path domain tree) will be known in advance due to deployment and commercial limitations (e.g., inter-AS peering agreements).

In the Figure 1 below, D1 is the root domain; D5 and D6 are the destination domains. The ingress is A in domain D1; egresses are X, Y in Domain D6 and Z in Domain D5.

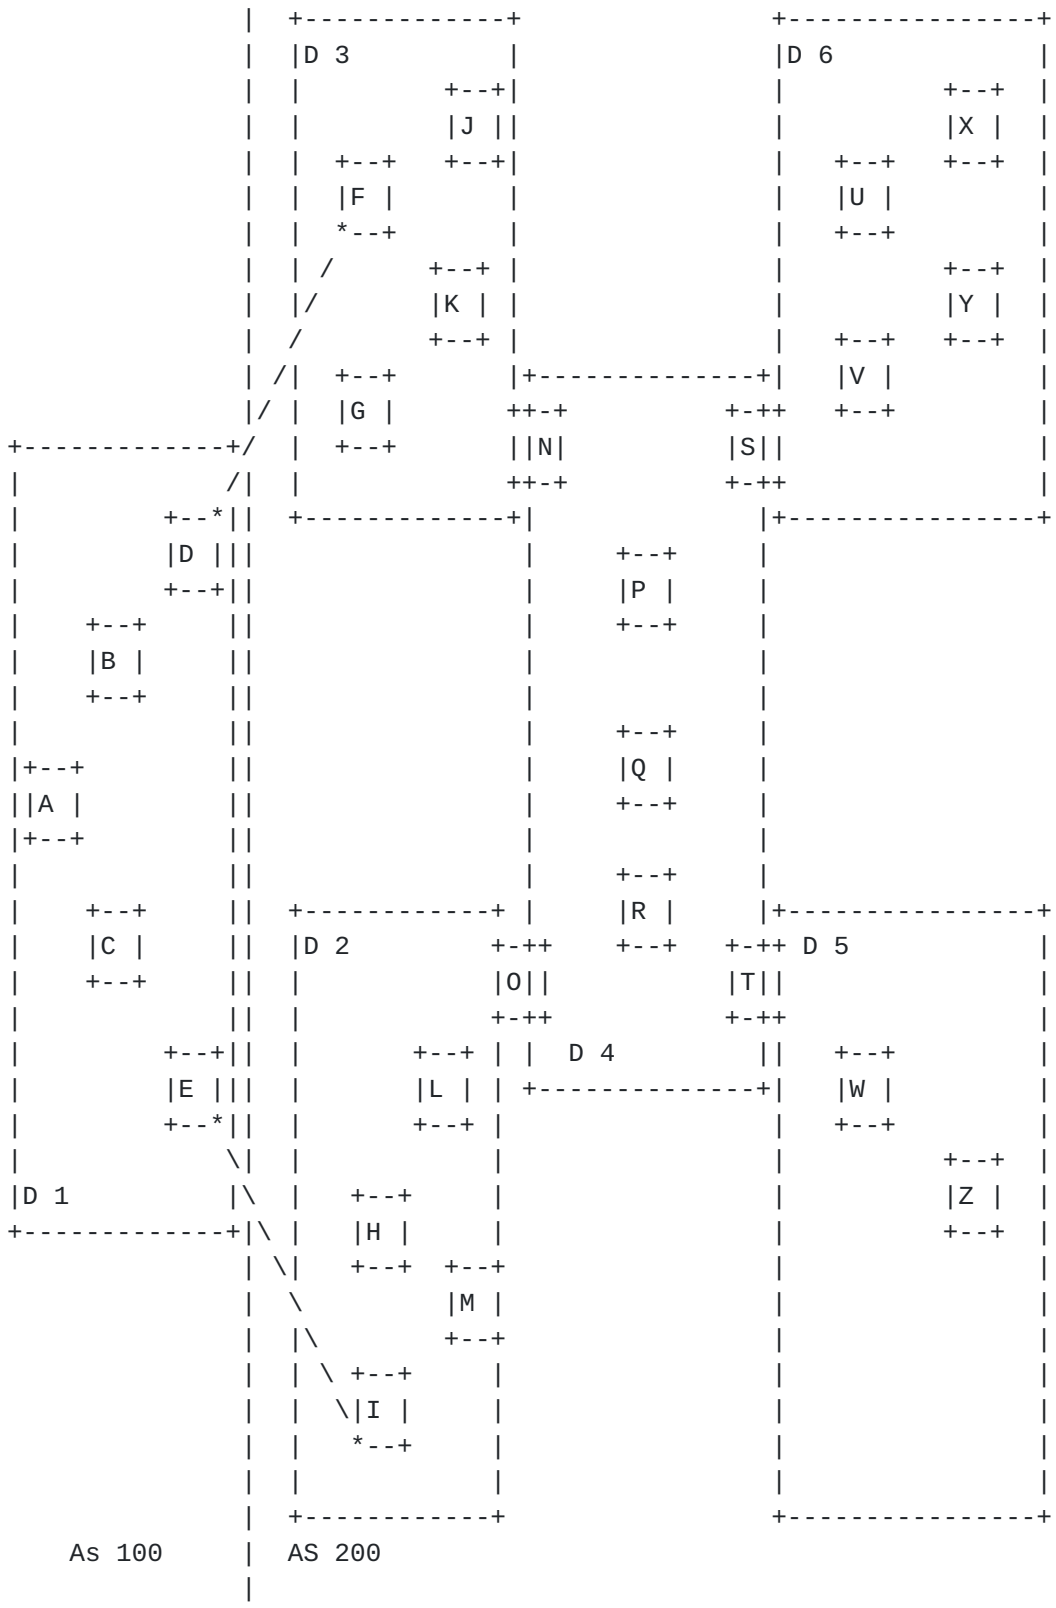

Figure 1: Domain Topology Example

In the Figure 2 below, the P2MP tree spans 5 domains. Destination in D6 (X & Y) would use the domain-sequence: D1-D3-D4-D6; and destination in D5 (Z) would use the domain-sequence: D1-D3-D4-D5.

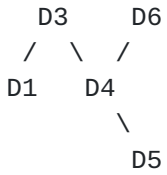

Figure 2: Domain Sequence Tree

Since destinations in different destination domain will have different domain sequence within the domain tree, it requires following encoding-

- Destination X and Y: D1-D3-D4-D6
- Destination Z : D1-D3-D4-D5

An extension in P2MP Path Computation request is needed to support this. (Refer Section 4.2)

The abstract nodes MAY include (but not limited to) domain subobjects AS number and IGP Area as described in [DOMAIN-SEQ].

[PCE-P2MP-PROCEDURES] also mentions PCE-sequence (i.e. PCE that serves each domain in the path domain tree); like domain-sequence as explained above, PCE-sequence will be different for different destinations and thus should be encoded as such.

3.2. Explicit inclusion or exclusion of abstract nodes

[RFC6006] describes four possible types of leaves in a P2MP request encoded in P2MP END-POINTS object.

- New leaves to add
- Old leaves to remove
- Old leaves whose path can be modified/reoptimized
- Old leaves whose path must be left unchanged

Currently [RFC6006] only allows a list of nodes that can be used as branch nodes or a list of nodes that cannot be used as branch nodes by using the Branch Node Capability (BNC) Object, which applies to

all leaves (old and new) in the P2MP tree.

For an existing P2MP tree which may already have a branch node through which most of the leaves are connected, but when adding a set of new leaves, administrator may want to exclude that node (as it may soon be overloaded) and would like to balance the final P2MP tree. This cannot be achieved via the BNC object but by explicitly excluding a particular node or including a different node, for the P2MP END-POINTS object for new leaves only.

Administrator at the source can exert stronger control by providing explicit inclusion or exclusion of any abstract nodes (not limited to branch nodes) for a group (subset) of destinations and not all destinations.

4. Detailed Description

4.1. Objective

[RFC6006] defines Request Message Format and Objects, along with <end-point-rro-pair-list>. This section introduce the use of <IRO> and <XRO> which are added to the <end-point-rro-pair-list>.

To allow abstract nodes to be explicitly included or excluded for a subset of destinations (encoded in one <END-POINTS> object), changes are made as shown below.

The abstract node (encoded as subobject in <IRO> and <XRO>) MAY be an absolute hop, IP-Prefix, Autonomous system or IGP Area. The subobjects are described in [RFC3209], [RFC3477], [RFC4874] and [DOMAIN-SEQ].

Note that one P2MP Path request can have multiple <END-POINTS> objects and each P2MP <END-POINTS> object may have multiple destinations, the <IRO> and <XRO> is applied for all destinations in one such P2MP <END-POINTS> object.

4.2. Request Message Format

The format of PCReq message is modified as follows:

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

where:

        <end-point-iro-xro-rro-pair-list>::=
                       <END-POINTS>
                       [<IRO>]
                       [<XRO>]
                       [<RRO-List>][<BANDWIDTH>]
                       [<end-point-iro-xro-rro-pair-list>]

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

From [RFC6006] usage of <end-point-rro-pair-list> is changed to <end-point-iro-xro-rro-pair-list> in this document.

[RFC6006] describes Branch Node Capability (BNC) Object which is different from the use of <IRO> and <XRO> to specify inclusion/exclusion of abstract nodes for a subset of destinations as described here.

4.3. Backward Compatibility

A legacy implementation that does not support explicit inclusion or exclusion of abstract nodes for a subset of P2MP destinations will act according to the procedures set out in [RFC5440], that is it will find the P2MP Path Request message out of order with respect to the format specified in [RFC6006].

5. IANA Considerations

There are no new IANA allocation in this document.

6. Security Considerations

PCEP security mechanisms as described in [RFC5440], [RFC6006] and [PCE-P2MP-PROCEDURES] are applicable for this document.

The new explicit inclusion or exclusion of abstract nodes for a subset of P2MP destination defined in this document allow finer and more specific control of the path computed by a PCE. Such control increases the risk if a PCEP message is intercepted, modified, or spoofed because it allows the attacker to exert control over the path that the PCE will compute or to make the path computation impossible. Therefore, the security techniques described in [RFC5440], [RFC6006] and [PCE-P2MP-PROCEDURES] are considered more important.

Note, however, that the route exclusion mechanisms also provide the operator with the ability to route around vulnerable parts of the network and may be used to increase overall network security.

7. Manageability Considerations

7.1. Control of Function and Policy

Mechanisms defined in this document do not add any new control function/policy requirements in addition to those already listed in [RFC6006].

7.2. Information and Data Models

Mechanisms defined in this document do not imply any new MIB requirements in addition to those already listed in [PCE-P2MP-MIB].

7.3. 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 [RFC6006].

7.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC6006].

7.5. Requirements On Other Protocols

Mechanisms defined in this document do not imply any requirements on other protocols in addition to those already listed in [RFC6006].

7.6. Impact On Network Operations

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [RFC6006].

8. Acknowledgments

We would like to thank Pradeep Shastry, Suresh babu, Quintin Zhao, Daniel King and Chen Huaimo for their useful comments and suggestions.

9. References

9.1. Normative References

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

9.2. 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.

[RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.

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

[RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.

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

[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.

[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.

[PCE-P2MP-PROCEDURES] Zhao, Q., Dhody, D., Ali, Z., Saad,, T., Sivabalan,, S., and R. Casellas, "PCE-based Computation Procedure To Compute Shortest Constrained P2MP Inter-domain Traffic Engineering Label Switched Paths (draft-ietf-pce-pcep-inter-domain-p2mp-procedures-03)", Feb 2013.

[PCE-P2MP-MIB] Zhao, Q., Dhody, D., Palle, U., and D. King, "Management Information Base for the PCE Communications Protocol (PCEP) When Requesting Point-to-Multipoint Services (draft-zhao-pce-pcep-p2mp-mib-05)", August 2012.

[DOMAIN-SEQ] Dhody, D., Palle, U., and R. Casellas, "Standard Representation Of Domain Sequence (draft-ietf-pce-pcep-domain-sequence-03)", March 2013.

Authors' Addresses

Dhruv Dhody
Huawei Technologies India Pvt Ltd
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: dhruv.dhody@huawei.com

Udayasree Palle
Huawei Technologies India Pvt Ltd
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: udayasree.palle@huawei.com

Venugopal Reddy Kondreddy
Huawei Technologies India Pvt Ltd
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: venugopalreddyk@huawei.com