Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 22, 2014

T. Otani
K. Ogaki
KDDI
D. Caviglia
Ericsson
F. Zhang
Huawei Technologies
C. Margaria
Coriant R&D GmbH
July 21, 2013

Requirements for GMPLS applications of PCE
draft-ietf-pce-gmpls-aps-req-09.txt

Abstract

The initial effort of the PCE (Path computation element) WG was mainly focused on MPLS. As a next step, this draft describes functional requirements for GMPLS application of PCE.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The initial effort of the PCE (Path computation element) WG was mainly focused on solving the path computation problem within a domain or over different domains in MPLS networks. As the same case with MPLS, service providers (SPs) have also come up with requirements for path computation in GMPLS-controlled networks [RFC3945] such as wavelength, TDM-based or Ethernet-based networks as well.

[RFC4655] and [RFC4657] discuss the framework and requirements for PCE on both packet MPLS networks and GMPLS-controlled networks. This document complements these RFCs by providing some considerations of GMPLS applications in the intra-domain and inter-domain networking environments and indicating a set of requirements for the extended definition of PCE-related protocols.

Note that the requirements for inter-layer and inter-area traffic engineering described in [RFC6457] and [RFC4927] are outside of the scope of this document.

Constraint-based shortest path first (CSPF) computation within a domain or over domains for signaling GMPLS Label Switched Paths (LSPs) is usually more stringent than that of MPLS TE LSPs [RFC4216], because the additional constraints, e.g., interface switching capability, link encoding, link protection capability, SRLG (Shared risk link group) [RFC4202] and so forth need to be considered to establish GMPLS LSPs. GMPLS signaling protocol [RFC3473] is designed taking into account bi-directionality, switching type, encoding type and protection attributes of the TE links spanned by the path, as well as LSP encoding and switching type of the end points, appropriately.

This document provides requirements for GMPLS applications of PCE in support of GMPLS path computation, included are requirements for both intra-domain and inter-domain environments.

2. GMPLS applications of PCE

2.1. Path computation in GMPLS network

Figure 1 depicts a model GMPLS network, consisting of an ingress link, a transit link as well as an egress link. We will use this model to investigate consistent guidelines for GMPLS path computation. Each link at each interface has its own switching capability, encoding type and bandwidth.

```
          Ingress             Transit             Egress
+-----+   link1-2   +-----+   link2-3   +-----+   link3-4   +-----+
|Node1|------------>|Node2|------------>|Node3|------------>|Node4|
|     |<------------|     |<------------|     |<------------|     |
+-----+   link2-1   +-----+   link3-2   +-----+   link4-3   +-----+
```

Figure 1: Path computation in GMPLS networks

For the simplicity in consideration, the below basic assumptions are made when the LSP is created.

(1) Switching capabilities of outgoing links from the ingress and egress nodes (link1-2 and link4-3 in Figure 1) are consistent with each other.

(2) Switching capabilities of all transit links including incoming links to the ingress and egress nodes (link2-1 and link3-4) are consistent with switching type of a LSP to be created.

(3) Encoding-types of all transit links are consistent with encoding type of a LSP to be created.

GMPLS-controlled networks (e.g., GMPLS-based TDM networks) are usually responsible for transmitting data for the client layer.

These GMPLS-controlled networks can provide different types of connections for customer services based on different service bandwidth requests.

The applications and the corresponding additional requirements for applying PCE to, for example, GMPLS-based TDM networks, are described in Figure 2. In order to simplify the description, this document just discusses the scenario in SDH networks as an example. The scenarios in SONET or OTN are similar to this scenario.

```
                     N1                    N2
    +-----+       +------+              +------+
    |     |-------|      |--------------|      |       +-------+
    +-----+       |      |---|          |      |       |       |
       A1         +------+   |          +------+       |       |
                     |       |             |           +-------+
                     |       |             |              PCE
                     |       |             |
                     |      +------+       |
                     |      |      |       |
                     |      |      |-----| |
                     |      +------+     | |
                     |         N5        | |
                     |                   | |
                  +------+              +------+
                  |      |              |      |        +-----+
                  |      |--------------|      |--------|     |
                  +------+              +------+        +-----+
                     N3                    N4              A2
```

Figure 2: A simple TDM (SDH) network

Figure 2 shows a simple TDM (SDH) network topology, where N1, N2, N3, N4 and N5 are all SDH switches. Assume that one Ethernet service with 100M bandwidth is required from A1 to A2 over this network. The client Ethernet service could be provided by a VC4 container from N1 to N4, and it could also be provided by three concatenated VC3 containers (Contiguous or Virtual concatenation) from N1 to N4.

In this scenario, when the ingress node (e.g., N1) receives a client service transmitting request, the type of containers (one VC4 or three concatenated VC3) could be determined by PCC (Path computation client) (e.g., N1 or NMS), but could also be determined by PCE automatically based on policy [RFC5394]. If it is determined by PCC, PCC should be capable of specifying the ingress node and egress node, signal type, the type of the concatenation and the number of the concatenation in a PCReq (Path computation request) message. PCE should consider those parameters during path computation. The route information (co-route or separated-route) should be specified in a PCRep (Path computation reply) message if path computation is performed successfully.

As described above, PCC should be capable of specifying TE attributes defined in the next section and PCE should compute a path accordingly.

Where a GMPLS network is consisting of inter-domain (e.g., inter-AS or inter-area) GMPLS-controlled networks, requirements on the path computation follows [RFC5376] and [RFC4726].

2.2. Unnumbered Interface

GMPLS supports unnumbered interface ID that is defined in [RFC3477], which means that the endpoints of the path may be unnumbered. It should also be possible to request a path consisting of the mixture of numbered links and unnumbered links, or a P2MP (Point-to-multipoint) path with different types of endpoints. Therefore, the PCC should be capable of indicating the unnumbered interface ID of the endpoints in the PCReq message.

2.3. Asymmetric Bandwidth Path Computation

As per [RFC6387], GMPLS signaling can be used for setting up an asymmetric bandwidth bidirectional LSP. If a PCE is responsible for the path computation, the PCE should be capable of computing a path for the bidirectional LSP with asymmetric bandwidth. It means that the PCC should be able to indicate the asymmetric bandwidth requirements in forward and reverse directions in the PCReq message.

3. Requirements for GMPLS application of PCE

3.1. Requirements on Path Computation Request

As for path computation in GMPLS-controlled networks as discussed in section 2, the PCE should appropriately consider the GMPLS TE attributes listed below once a PCC or another PCE requests a path computation. The path calculation request message from the PCC or the PCE must contain the information specifying appropriate attributes. According to [RFC5440], [PCE-WSON-REQ] and to RSVP procedures like explicit label control(ELC),the additional attributes introduced are as follows:

(1) Switching capability/type: as defined in [RFC3471], [RFC4203] and, all current and future values.

(2) Encoding type: as defined in [RFC3471], [RFC4203] and, all current and future values.

(3) Signal Type: as defined in [RFC4606] and, all current and future values.

(4) Concatenation Type: In SDH/SONET and OTN, two kinds of concatenation modes are defined: contiguous concatenation which requires co-route for each member signal and requires all the interfaces along the path to support this capability, and virtual concatenation which allows diverse routes for the member signals and only requires the ingress and egress interfaces to support this capability. Note that for the virtual concatenation, it also may specify co-routed or separated-routed. See [RFC4606] and [RFC4328] about concatenation information.

(5) Concatenation Number: Indicates the number of signals that are requested to be contiguously or virtually concatenated. Also see [RFC4606] and [RFC4328].

(6) Technology-specific label(s) such as defined in [RFC4606], [RFC6060], [RFC6002] or [RFC6205].

(7) e2e Path protection type: as defined in [RFC4872], e.g., 1+1 protection, 1:1 protection, (pre-planned) rerouting, etc.

(8) Administrative group: as defined in [RFC3630]

(9) Link Protection type: as defined in [RFC4203]

(10)Support for unnumbered interfaces: as defined in [RFC3477]

(11)Support for asymmetric bandwidth request: as defined in [RFC6387]

(12)Support for explicit label control during the path computation.

(13)Support of label restrictions in the requests/responses, similarly to RSVP-TE ERO (Explicit route object) and XRO (Exclude route object) as defined in [RFC3473] and [RFC4874].

3.2. Requirements on Path Computation Reply

As described above, a PCE should compute the path that satisfies the constraints which are specified in the PCReq message. Then the PCE should send a PCRep message including the computation result to the PCC. For Path Computation Reply message (PCRep) in GMPLS networks, there are some additional requirements. The PCEP (PCE communication protocol) PCRep message must be extended to meet the following requirements.

(1) Path computation with concatenation

In the case of path computation involving concatenation, when a PCE receives the PCReq message specifying the concatenation constraints described in section 3.1, the PCE should compute a path accordingly.

For path computation involving contiguous concatenation, a single route is required and all the interfaces along the route should support contiguous concatenation capability. Therefore, the PCE should compute a path based on the contiguous concatenation capability of each interface and only one ERO which should carry the route information for the response.

For path computation involving virtual concatenation, only the ingress/egress interfaces need to support virtual concatenation capability and there may be diverse routes for the different member signals. Therefore, multiple EROs may be needed for the response. Each ERO may represent the route of one or multiple member signals. In the case where one ERO represents several member signals among the total member signals, the number of member signals along the route of the ERO must be specified.

(2) Label constraint

In the case that a PCC does not specify the exact label(s) when requesting a label-restricted path and the PCE is capable of performing the route computation and label assignment computation procedure, the PCE needs to be able to specify the label of the path in a PCRep message.

Wavelength restriction is a typical case of label restriction. More generally in GMPLS-controlled networks label switching and selection constraints may apply and a PCC may request a PCE to take label constraint into account and return an ERO containing the label or set of label that fulfil the PCC request.

(3) Roles of the routes

When a PCC specifies the protection type of an LSP, the PCE should compute the working route and the corresponding protection route(s).

Therefore, the PCRep should allow to distinguish the working (nominal) and the protection routes. According to these routes, RSVP-TE procedure appropriately creates both the working and the protection LSPs for example with ASSOCIATION object [RFC6689].

3.3. GMPLS PCE Management

This document does not change any of the management or operational details for networks that utilise PCE. Please refer to [RFC4655] for an overview of this scenery. However, this document proposes the introduction of several PCEP objects and data for the better integration of PCE with GMPLS networks. Those protocol elements will need to be visible in any management tools that apply to the PCE, PCC, and PCEP. That includes, but is not limited to, adding appropriate objects to existing PCE MIB modules that are used for modelling and monitoring PCEP deployments [PCEP-MIB]. Ideas for what objects are needed may be guided by the relevant GMPLS extensions in GMPLS-TE-STD-MIB [RFC4802]."

4. Security Considerations

PCEP extensions to support GMPLS should be considered under the same security as current PCE work and this extension will not change the underlying security issues. Sec. 10 of [RFC5440] describes the list of security considerations in PCEP. At the time [RFC5440] was published, TCP Authentication Option (TCP-AO) had not been fully specified for securing the TCP connections that underlie PCEP sessions. TCP-AO [RFC5925] has now been published and PCEP implementations should fully support TCP-AO according to [RFC6952].

5. IANA Considerations

This document has no actions for IANA.

6. Acknowledgement

The author would like to express the thanks to Ramon Casellas, Julien Meuric, Adrian Farrel, Yaron Sheffer and Shuichi Okamoto for their comments.

7. References

7.1. Normative References

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

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

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4202] Kompella, K. and Y. Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.

[RFC4203] Kompella, K. and Y. Rekhter, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC4328] Papadimitriou, D., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control", RFC 4328, January 2006.

[RFC4606] Mannie, E. and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606, August 2006.

[RFC4802] Nadeau, T. and A. Farrel, "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802, February 2007.

[RFC4872] Lang, J., Rekhter, Y., and D. Papadimitriou, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.

[RFC4927] Le Roux, J., "Path Computation Element Communication Protocol (PCECP) Specific Requirements for Inter-Area MPLS and GMPLS Traffic Engineering", RFC 4927, June 2007.

[RFC5376] Bitar, N., Zhang, R., and K. Kumaki, "Inter-AS Requirements for the Path Computation Element Communication Protocol (PCECP)", RFC 5376, November 2008.

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

[RFC6002] Berger, L. and D. Fedyk, "Generalized MPLS (GMPLS) Data Channel Switching Capable (DCSC) and Channel Set Label Extensions", RFC 6002, October 2010.

[RFC6060] Fedyk, D., Shah, H., Bitar, N., and A. Takacs, "Generalized Multiprotocol Label Switching (GMPLS) Control of Ethernet Provider Backbone Traffic Engineering (PBB-TE)", RFC 6060, March 2011.

[RFC6205] Otani, T. and D. Li, "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, March 2011.

[RFC6387] Takacs, A., Berger, L., Caviglia, D., Fedyk, D., and J. Meuric, "GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)", RFC 6387, September 2011.

[RFC6689] Berger, L., "Usage of the RSVP ASSOCIATION Object", RFC 6689, July 2012.

7.2. Informative References

[PCE-WSON-REQ] Lee, Y., Bernstein, G., Martensson, J., Takeda, T., Tsuritani, T., and O. de Dios, "PCEP Requirements for WSON Routing and Wavelength Assignment", draft-ietf-pce-wson-routing-wavelength-09 (work in progress), June 2013.

[PCEP-MIB] Koushik, A., Emile, S., Zhao, Q., King, D., and J. Hardwick, "PCE communication protocol (PCEP) Management Information Base", draft-ietf-pce-pcep-mib-05 (work in progress), July 2013.

[RFC4216] Zhang, R. and J. Vasseur, "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, November 2005.

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

[RFC4657] Ash, J. and J. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.

[RFC4726] Farrel, A., Vasseur, J., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 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.

[RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, December 2008.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

[RFC6457] Takeda, T. and A. Farrel, "PCC-PCE Communication and PCE Discovery Requirements for Inter-Layer Traffic Engineering", RFC 6457, December 2011.

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, May 2013.

Authors' Addresses

Tomohiro Otani
KDDI Corporation
2-3-2 Nishi-shinjuku
Shinjuku-ku, Tokyo
Japan

Phone: +81-(3) 3347-6006
Email: tm-otani@kddi.com

Kenichi Ogaki
KDDI Corporation
3-10-10 Iidabashi
Chiyoda-ku, Tokyo
Japan

Phone: +81-(3) 6678-0284
Email: ke-oogaki@kddi.com

Diego Caviglia
Ericsson
16153 Genova Cornigliano
Italy

Phone: +390106003736
Email: diego.caviglia@ericsson.com

Fatai Zhang
Huawei Technologies Co., Ltd.
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District, Shenzhen 518129
P.R.China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

Cyril Margaria
Coriant R&D GmbH
St Martin Strasse 76
Munich, 81541
Germany

Phone: +49 89 5159 16934
Email: cyril.margaria@coriant.com

Draft Authors:
Daniel King<daniel@olddog.co.uk>
Adrian Farrel<adrian.farrel@huawei.com>
Quintin Zhao<qzhao@huawei.com>
Fatai Zhang<zhangfatai@huawei.com>

Network Working Group
Internet Draft
Intended status: Standard Track
Expires: June 16, 2018

Y. Lee
Huawei

G. Bernstein
Grotto Networking

Jonas Martensson
Acreo

T. Takeda
NTT

T. Tsuritani
KDDI

December 17, 2018

PCEP Extensions for WSON Impairments

draft-lee-pce-wson-impairments-08

Abstract

As an optical signal progresses along its path it may be altered by the various physical processes in the optical fibers and devices it encounters. When such alterations result in signal degradation, these processes are usually referred to as "impairments". These physical characteristics may be important constraints to consider in path computation process in wavelength switched optical networks.

This document provides PCEP extensions to support Impairment Aware Routing and Wavelength Assignment (IA-RWA) in wavelength switched optical networks.

Status of this Memo

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on June 3, 2019.

Copyright Notice

Copyright (c) 2018 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.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Table of Contents

1. Introduction

[RFC4655] defines the PCE based architecture and explains how a Path Computation Element (PCE) may compute Label Switched Paths (LSP) in Multiprotocol Label Switching Traffic Engineering (MPLS-TE) and Generalized MPLS (GMPLS) networks at the request of Path Computation Clients (PCCs). A PCC is shown to be any network component that makes such a request and may be for instance an Optical Switching Element within a Wavelength Division Multiplexing (WDM) network. The PCE, itself, can be located anywhere within the network, and may be within an optical switching element, a Network Management System (NMS) or Operational Support System (OSS), or may be an independent network server.

The PCE communication Protocol (PCEP) is the communication protocol used between PCC and PCE, and may also be used between cooperating PCEs. [RFC4657] sets out the common protocol requirements for PCEP. Additional application-specific requirements for PCEP are deferred to separate documents.

This document provides a set of application-specific PCEP requirements for support of path computation in Wavelength Switched Optical Networks (WSON) with impairments. WSON refers to WDM based optical networks in which switching is performed selectively based on the wavelength of an optical signal.

The path in WSON is referred to as an optical path. An optical path may span multiple fiber links and the path should be assigned a wavelength for each link. A transparent optical network is made up of optical devices that can switch but not convert from one wavelength to another. In a transparent optical network, an optical path operates on the same wavelength across all fiber links that it traverses. In such case, the optical path is said to satisfy the wavelength-continuity constraint. Two optical paths that share a common fiber link can not be assigned the same wavelength. To do otherwise would result in both signals interfering with each other. Note that advanced additional multiplexing techniques such as polarization based multiplexing are not addressed in this document since the physical layer aspects are not currently standardized. Therefore, assigning the proper wavelength on an optical path is an essential requirement in the optical path computation process.

When a switching node has the ability to perform wavelength conversion the wavelength-continuity constraint can be relaxed, and a may use different wavelengths on different links along its route from origin to destination. It is, however, to be noted that wavelength converters may be limited due to their relatively high cost, while the number of WDM channels that can be supported in a fiber is also limited. As a WSON can be composed of network nodes that cannot perform wavelength conversion, nodes with limited wavelength conversion, and nodes with full wavelength conversion abilities, wavelength assignment is an additional routing constraint to be considered in all optical path computation.

One of the most basic questions in communications is whether one can successfully transmit information from a transmitter to a receiver within a prescribed error tolerance, usually specified as a maximum permissible bit error ratio (BER). This generally depends on the nature of the signal transmitted between the sender and receiver and the nature of the communications channel between the sender and receiver. The optical path utilized (along with the wavelength) determines the communications channel.

The optical impairments incurred by the signal along the fiber and at each optical network element along the path determine whether the BER performance or any other measure of signal quality can be met for this particular signal on this particular path. Given the existing standards covering optical characteristics (impairments) and the knowledge of how the impact of impairments may be estimated along a path, [RFC6566] provides a framework for impairment aware path computation and establishment utilizing GMPLS protocols and the PCE architecture.

Some transparent optical subnetworks are designed such that over any path the degradation to an optical signal due to impairments never exceeds prescribed bounds. This may be due to the limited geographic extent of the network, the network topology, and/or the quality of the fiber and devices employed. In such networks the path selection problem reduces to determining a continuous wavelength from source to destination (the Routing and Wavelength Assignment problem). These networks are discussed in [RFC6163]. In other optical networks, impairments are important and the path selection process must be impairment-aware.

In this document we first review the processes for routing and wavelength assignment (RWA) used when wavelength continuity constraints are present. We then review the processes for optical impairment aware RWA (IA-RWA). Based on selected process models we then specify requirements for PCEP to support IA-RWA. Note that requirements for PCEP to support RWA are specified in a separate document [RFC7449].

The remainder of this document uses terminology from [RFC4655].

1.1. WSON RWA Processes (no impairments)

In [RFC6163] three alternative process architectures were given for performing routing and wavelength assignment. These are shown schematically in Figure 1.

```
     +--------------------+
     |  +--------+  +--+  |    +--------+    +--+     +--------+    +---+
     |  |Routing|   |WA|  |    |Routing|--->|WA|      |Routing|--->|DWA|
     |  +--------+  +--+  |    +--------+    +--+     +--------+    +---+
     |   Combined         |     Separate Processes    Separate Processes
     |   Processes        |                           WA performed in a
     +--------------------+                           Distributed manner
           (a)                        (b)                     (c)
```

Figure 1 RWA process alternatives.

Detail description of each alternative can be found in [RFC6163].

1.2. WSON IA-RWA Processes

In [RFC6566] impairments were addressed by adding an "impairment validation" (IV) process. For approximate impairment validation three process alternatives were given in [RFC6566] and are shown in Figure 2. Since there are many possible alternative combinations, these are just three examples. Please note that the requirements for all possible architectures can be reduced to the cases in Figure 3 in section 2.

```
                  +-----------------------------------+
                  |   +--+     +-------+     +--+     |
                  |   |IV|     |Routing|     |WA|     |
                  |   +--+     +-------+     +--+     |
                  |                                   |
                  |        Combined Processes         |
                  +-----------------------------------+
                                 (a)

           +--------------+      +----------------------+
           | +----------+ |      | +-------+    +--+    |
           | |    IV    | |      | |Routing|    |WA|    |
           | |candidates| |----->| +-------+    +--+    |
           | +----------+ |      |  Combined Processes  |
           +--------------+      +----------------------+
                                 (b)

            +-----------+        +----------------------+
            | +-------+ |        |    +--+    +--+      |
            | |Routing| |------->|    |WA|    |IV|      |
            | +-------+ |        |    +--+    +--+      |
            +-----------+        | Distributed Processes|
                                 +----------------------+
                                 (c)
```

Figure 2 Process flows for the three main approximate impairment architectural alternatives.

These alternatives have the following properties and impact on PCEP requirements in this document.

1. Combined IV and RWA Process - Here the processes of impairment validation, routing and wavelength assignment are aggregated into a single PCE. The requirements for PCC-PCE interaction with such a combined IV-RWA process PCE is addressed in this document.

2. IV-Candidates + RWA Process - As explained in [RFC6566] separating the impairment validation process from the RWA process maybe necessary to deal with impairment sharing constraints. In this architecture one PCE computes impairment candidates and another PCE uses this information while performing RWA. The requirements for PCE-to-PCE interaction of this architecture will be addressed in this document.

3. Routing + Distributed WA and IV - Here a standard path computation (unaware of detailed wavelength availability or optical impairments) takes place, then wavelength assignment and impairment validation is performed along this path in a distributed manner via signaling (RSVP-TE). This alternative should be covered by existing or emerging GMPLS PCEP extensions and does not present new WSON specific requirements.

2. WSON PCE Architectures and Requirements

In the previous section we reviewed various process architectures for implementing RWA with and without regard for optical impairment. In Figure 3 we reduce these alternatives to two PCE based implementations. As specified in [RFC6566], the PCE in Figure 3(a) should be given the necessary information for RWA and impairment validation, including WSON topology, link wavelength utilization as well as impairment information such as the adjustment range of tunable parameters, etc. Similarly, RWA-PCE should be equipped with all the information other than impairment-related ones which is a necessity for IV-PCE.

In Figure 3(a) we show the three processes of routing, wavelength assignment and impairment validation accessed via a single PCE. The implementation details of the interactions of the processes are not subject to standardization; this document concerns only the PCC to PCE communications.

In Figure 3(b) the impairment validation process is implemented in a separate PCE. Here the RWA-PCE acts as a coordinator and the PCC to RWA-PCE interface will be the same as in Figure 3(a), however in this case we have additional requirements for the RWA-PCE to IV-PCE interface.

```
                          +-----------------------------------+
            +-----+       |   +--+     +-------+     +--+     |
            |     |       |   |IV|     |Routing|     |WA|     |
            | PCC |<----->|   +--+     +-------+     +--+     |
            |     |       |                                   |
            +-----+       |               PCE                 |
                          +-----------------------------------+

                                   (a)


                       +----------------------+     +--------------+
         +-----+       | +-------+    +--+    |     |              |
         |     |       | |Routing|    |WA|    |     |      IV      |
         | PCC |<----->| +-------+    +--+    |<--->|  candidates  |
         |     |       |                      |     |              |
         +-----+       | RWA-PCE (coordinator)|     |    IV-PCE    |
                       +----------------------+     +--------------+

                                   (b)
```

Figure 3 PCE architectures for IA-RWA.

2.1. RWA PCC to PCE Interface

The PCC to PCE interface of Figure 3(a) and the PCC to RWA-PCE (coordinator) interface of Figure 3(b) are the same and we will cover both in this section. The following requirements for these interfaces are arranged by use cases:

2.1.1. A new RWA path request

The PCReq Message MUST include one or more specific measures of optical signal quality to which all feasible paths should conform:

- BER limit
- OSNR + Margin
- Power
- PMD
- Residual Dispersion (RD)
- Q factor
- TBD

(Editor's Note: this is not a complete list of optical signal quality measure and subject to further change.)

If the PCReq Message does not include the BER limit and no BER limit information related to the specific path request is provisioned at the PCE then the PCE will return an error specifying that a BER limit must be provided.

"Margin" means "insurance" (e.g. 3~6dB) for suppliers and operators which are set against unpredictable degradation and other degradation not included in the provided estimates such as that due to fiber nonlinearity.

In non-coherent WDM networks, PMD and CD should be carefully considered. However, coherent WDM networks usually have a high tolerance with these two optical signal quality measurements and thus it may not need to be considered.

2.1.1.1. Signal Quality Measure TLV

This TLV represents all impairment constraints that need to be considered by the PCE to calculate a path that passes the requested measure of signal quality for a signal for a given source and destination.

This TLV is repeated one after another until all signal quality types are specified.

The TLV type is TBD.

The TLV data is defined as follow:

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |P| Signal Quality Type         |    Reserved                   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                          Threshold                            |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The P bit (1 bit): Indicates if the associated impairment is a path level or not.

The P bit is set to 1 indicates that the associated impairment is a path level. This means that the impairment is associated with the end-to-end path and the threshold must be satisfied on a path level.

The P bit is set to 0 indicates that the associated impairment is a link level. This means the impairment is associated with the link and the threshold must be satisfied on every link of the end-to-end path.

The Signal Quality Type (15 bits): indicates the kind of optical signal quality of interest.

0: reserved

1: BER limit

2: OSNR+ Margin

3: Power

4: PMD

5: CD

6: Q factor

7-up: Reserved for future use

Threshold (32 bits) indicates the threshold (upper or lower) to which the specified signal quality measure must satisfy for the path/link (depending on the P bit).

The reserved bits MUST be set to 0 on transmit and MUST be ignored on receive.

2.1.2. A new RWA path reply

The PCRep Message MUST include the route, wavelengths assigned to the route, and an indicator that says if the path conforms to the required quality or not. Moreover, it should also be able to specify a list of impairment compensation information along the chosen route, i.e., the value or value range of optical signal quality parameter that needs to be adjusted, such as power level, in order to achieve the resultant measure of signal quality as given in Section 2.1.2.1. It is suggested to carry this information in the PCEP ERO object. According to [RFC5440], PCEP ERO object is identical to RSVP-TE ERO object. Therefore, it is suggested to modify the RSVP-TE ERO object to accommodate this need. This will be included in a separate draft in the future.

In the case where a valid path is not found, the PCRep Message MUST include why the path is not found (e.g., no route, wavelength not found, BER failure, etc.)

2.1.2.1. Signal Quality Measure TLV

This TLV represents the result of the requested measure of signal quality for a signal for a given source and destination.

This TLV is repeated one after another until all signal quality types are specified.

The TLV type is TBD.

The TLV data is defined as follow:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|P| Signal Quality Type         |    Reserved                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Signal Quality Value                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The P bit (1 bit): Indicates if the associated signal quality measure has passed the threshold or not.

The P bit is set to 1 indicates that the associated signal quality measure has passed the threshold.

The P bit is set to 0 indicates that the associated signal quality measure has failed the threshold.

The Signal Quality Type (15 bits): indicates the kind of optical signal quality of interest.

0: reserved

1: BER limit

2: OSNR_ Margin

3: Power

4: PMD

5: CD

6: Q factor

7-up: Reserved for future use

Signal Quality Value (32 bits) indicates the actual estimated value of the specified signal quality measure for the end-to-end path.

TBD: How to encode link based value needs to be determined in the revision.

The reserved bits MUST be set to 0 on transmit and MUST be ignored on reception.

2.2. RWA-PCE to IV-PCE Interface

In [RFC6566] a sequence diagram for the interaction of the PCC, RWA-PCE and IV-PCE of Figure 3(b) was given and is repeated here in Figure 4. The interface between the PCC and the RWA-PCE (acting as the coordinator) was covered in section 2.1.

```
     +---+                +-------------+          +-----------------+
     |PCC|                |RWA-Coord-PCE|          |IV-Candidates-PCE|
     +-+-+                +------+------+          +---------+-------+
       ...___     (a)          |                           |
       |     ````---...____      |                           |
       |                   ```-->|                           |
       |                         |                           |
       |                         |--..___     (b)            |
       |                         |       ```---...___        |
       |                         |                   ```---->|
       |                         |                           |
       |                         |           (c)       ___...|
       |                         |       ___....---''''      |
       |                         |<--''''                    |
       |                         |                           |
       |                         |                           |
       |          (d)       ___...|                           |
       |      ___....---'''      |                           |
       |<--'''                   |                           |
       |                         |                           |
       |                         |                           |
```

Figure 4 Sequence diagram for the interactions between PCC, RWA-Coordinating-PCE and the IV-Candidates-PCE.

The interface between the RWA-Coord-PCE and the IV-Candidates-PCE is specified by the following requirements:

1. The PCReq Message from the RWA-Coord-PCE to the IV-Candidate-PCE MUST include an indicator that more than one (candidate) path between source and destination is desired.

2. The PCReq message from the RWA-Coord-PCE to the IV-Candidates-PCE MUST include a limit on the number of optical impairment qualified paths to be returned by the IV-PCE.

3. The PCReq message from the RWA-Coord-PCE to the IV-Candidates-PCE MAY include wavelength constraints. Note that optical impairments are wavelength sensitive and hence specifying a wavelength constraint may help limit the search for valid paths. This requirement has been already covered in [RFC7449] and is presented here for an illustration purpose.

4. The PCRep Message from the IV-Candidates-PCE to RWA-Coord-PCE MUST include a set of optical impairment qualified paths along with any wavelength constraints on those paths.

5. The PCRep Message from the IV-Candidates-PCE to RWA-Coord-PCE MUST indicate "no path found" in case where a valid path is not found.

6. The PCReq Message from the RWA-Coord-PCE to the IV-Candidate-PCE MAY include one or more specified paths and wavelengths that is to be verified by the IV-PCE. This requirement is necessary when the IV-PCE is allowed to verify specific paths.

Note that once the RWA-Coord-PCE receives the resulting paths from the IV Candidates PCE, then the RWA-Coord-PCE computes RWA for the IV qualified candidate paths and sends the result back to the PCC.

2.2.1. A new impairment-validated (IV) path request

Details on encoding are TBD.

2.2.2. A new impairment-validated (IV) path reply

Details on encoding are TBD.

3. Manageability Considerations

Manageability of WSON Routing and Wavelength Assignment (RWA) with PCE must address the following considerations:

3.1. Control of Function and Policy

In addition to the parameters already listed in Section 8.1 of [RFC5440], a PCEP implementation SHOULD allow configuring the following PCEP session parameters on a PCC:

- The ability to send a WSON IA-RWA request.

In addition to the parameters already listed in Section 8.1 of [RFC5440], a PCEP implementation SHOULD allow configuring the following PCEP session parameters on a PCE:

- The support for WSON IA-RWA.
- The maximum number of synchronized path requests associated with WSON IA-RWA per request message.
- A set of WSON IA-RWA specific policies (authorized sender, request rate limiter, etc).

These parameters may be configured as default parameters for any PCEP session the PCEP speaker participates in, or may apply to a specific session with a given PCEP peer or a specific group of sessions with a specific group of PCEP peers.

3.2. Information and Data Models, e.g. MIB module

Extensions to the PCEP MIB module defined in [PCEP-MIB] should be defined, so as to cover the WSON IA-RWA information introduced in this document. A future revision of this document will list the information that should be added to the MIB module.

3.3. 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 section 8.3 of [RFC5440].

3.4. Verifying Correct Operation

Mechanisms defined in this document do not imply any new verification requirements in addition to those already listed in section 8.4 of [RFC5440]

3.5. Requirements on Other Protocols and Functional Components

The PCE Discovery mechanisms ([RFC5089] and [RFC5088]) may be used to advertise WSON IA-RWA path computation capabilities to PCCs.

3.6. Impact on Network Operation

Mechanisms defined in this document do not imply any new network operation requirements in addition to those already listed in section 8.6 of [RFC5440].

4. Security Considerations

This document has no requirement for a change to the security models within PCEP [PCEP]. However the additional information distributed in order to address the RWA problem represents a disclosure of network capabilities that an operator may wish to keep private. Consideration should be given to securing this information.

5. IANA Considerations

A future revision of this document will present requests to IANA for codepoint allocation.

6. References

6.1. Normative References

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

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

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

6.2. Informative References

[RFC6163] Lee, Y. and Bernstein, G., W. Imajuku, "Framework for GMPLS and PCE Control of Wavelength Switched Optical Networks", RFC 6163, April 2011.

[RFC6566] Lee, Y. and Bernstein, G. (Editors), and D. Li, "Framework for GMPLS and PCE Control of Wavelength Switched Optical Networks", RFC 6566, March, 2012.

[RFC7449] Y. Lee, G. Bernstein, J. Martensson, T. Takeda and T. Otani, "PCEP Requirements for WSON Routing and Wavelength Assignment", RFC 7449, February 2015.

[RFC5088] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.

[RFC5089] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, January 2008.

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

[RFC4657] Ash, J. and J. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.

[RFC8174] B. Leiba, "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", RFC 8174, May 2017.

Authors' Addresses

Young Lee (Ed.)
Huawei Technologies
Email: leeyoung@huawei.com

Greg Bernstein (Ed.)
Grotto Networking
Fremont, CA, USA
Phone: (510) 573-2237
Email: gregb@grotto-networking.com

Jonas Martensson
Acreo
Email:Jonas.Martensson@acreo.se

Tomonori Takeda
NTT Corporation
3-9-11, Midori-Cho
Musashino-Shi, Tokyo 180-8585, Japan
Email: takeda.tomonori@lab.ntt.co.jp

Takehiro Tsuritani
2-1-15 Ohara, Fujimino, Saitama, 356-8502, JAPAN
KDDI R&D Laboratories Inc.
Phone: +81-49-278-7806
Email: tsuri@kddilabs.jp

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972913
Email: zhang.xian@huawei.com

7. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: June 21, 2011

Q. Zhao
Huawei Technology
Z. Ali
T. Saad
S. Sivabalan
Cisco Systems
D. King
Old Dog Consulting
R. Casellas
CTTC - Centre Tecnologic de
Telecomunicacions de Catalunya

January 21, 2011

PCE-based Computation Procedure To Compute Shortest Constrained P2MP Inter-domain Traffic Engineering Label Switched Paths

draft-zhao-pce-pcep-inter-domain-p2mp-procedures-07

Abstract

The ability to compute paths for constrained point-to-multipoint (P2MP) Traffic Engineering Label Switched Paths (TE LSPs) across multiple domains has been identified as a key requirement for the deployment of P2MP services in MPLS and GMPLS networks. The Path Computation Element (PCE) has been recognized as an appropriate technology for the determination of inter-domain paths of P2MP TE LSPs.

This document describes the procedures and extensions to the PCE communication Protocol (PCEP) to handle requests and responses for the computation of inter-domain paths for P2MP TE LSPs.

Status of this Memo

This Internet-Draft will expire on June 21, 2011.

Copyright Notice

Table of Contents

1. Introduction

Multicast services are increasingly in demand for high-capacity applications such as multicast Virtual Private Networks (VPNs), IP-television (IPTV) which may be on-demand or streamed, and content-rich media distribution (for example, software distribution, financial streaming, or data-sharing). The ability to compute constrained Traffic Engineering Label Switched Paths (TE LSPs) for point-to-multipoint (P2MP) LSPs in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks across multiple domains. A domain can be defined as a collection of network elements within a common sphere of address management or path computational responsibility such as an IGP area or an Autonomous Systems.

The applicability of the Path Computation Element (PCE) [RFC4655] for the computation of such paths is discussed in [RFC5671], and the requirements placed on the PCE communications Protocol (PCEP) for this are given in [RFC5862].

This document describes how multiple PCE techniques can be combined to address the requirements. These mechanisms include the use of the per-domain path computation technique specified in [RFC5152], extensions to the backward recursive path computation (BRPC) technique specified in [RFC5441] for P2MP LSP path computation in an inter-domain environment, and a new procedure for core-tree based path computation defined in this document. These three mechanisms are suitable for different environments (topologies, administrative domains, policies, service requirements, etc.) and can also be effectively combined.

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 RFC 2119 [RFC2119].

2. Terminology

Terminology used in this document is consistent with the related MPLS/GMPLS and PCE documents [RFC4461], [RFC4655], [RFC4875], [RFC5376], [RFC5440], [RFC5441]. [RFC5671], and [RFC5862].

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

ASBR: Autonomous System Border Router. Router 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.

Core Tree: the core tree is a P2MP tree where the root is the ingress LSR, the transit nodes and branch nodes are the BNs of the transit domains and the leaf nodes are the leaf BNs of the leaf domains.

Destination: The lead Nodes can be in Root Domain, Transit Domain and Leaf Domain.

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-AS TE LSP: a TE LSP that crosses an AS boundary.

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

Leaf Domain: a domain that does not have a downstream neighbor domain. Note that, with this definition, a domain with one or more leaf nodes is not necessarily a leaf domain.

Leaf Boundary Nodes: the entry boundary node in the leaf domain.

Leaf Nodes: the LSR which is the P2MP LSP’s final.

LSR: Label Switching Router.

LSP: Label Switched Path.

OF: Objective Function. A set of one or more optimization criterion (criteria) used for the computation of paths either for single or for synchronized requests (e.g. path cost minimization), or the synchronized computation of a set of paths (e.g. aggregate bandwidth consumption minimization, etc.). See [RFC4655] and [RFC5441].

P2MP LSP Path Tree: A set of LSRs and TE links that comprise the path of a P2MP TE LSP from its ingress LSR to all of its egress LSRs.

Path Domain Sequence: The known sequence of domains for a path between root and leaf.

Path Domain Tree: The tree formed by the domains that the P2MP path crosses, where the source (ingress) domain is the root domain.

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

P2MP LSP Path Tree: A set of LSRs and TE links that comprise the path of a P2MP TE LSP from its ingress LSR to all of its egress LSRs.

Path Domain Sequence: the known sequence of domains for a path between the root node and a leaf node.

PCE Sequence: the known sequence of PCEs for calculating a path between the root node and a leaf node.

PCE Topology Tree: a list of PCE Sequences which has all the PCE Sequence for each path of the P2MP LSP path tree.

PCE(i): a PCE that performs path computations for domain(i).

Root Boundary Node: the egress LSR from the root domain on the path of the P2MP LSP.

Root Domain: the domain that includes the ingress (root) LSR.

TED: Traffic Engineering Database.

Transit/branch Domain: a domain that has an upstream and one or more downstream neighbour domain.

VSPT: Virtual Shortest Path Tree [RFC5441].

3. Problem Statement

The Path Computation Element (PCE) defined in [RFC4655] 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.

[RFC4875] describes how to set up P2MP TE LSPs for use in MPLS and GMPLS networks. The PCE is identified as a suitable application for the computation of paths for P2MP TE LSPs [RFC5671].

[RFC5441] specifies a procedure relying on the use of multiple PCEs to compute (P2P) inter-domain constrained shortest paths across a predetermined sequence of domains, using a backward recursive path computation technique. The technique can be combined with the use of path keys [RFC5520] to preserve confidentiality across domains, which is sometimes required when domains are managed by different Service Providers.

The PCE communication Protocol (PCEP) [RFC5440] is extended for point-to-multipoint(P2MP) path computation requests and in [RFC6006]. However, that specification does not provide all the necessary mechanisms to request the computation of inter-domain P2MP TE LSPs.

As discussed in [RFC4461], a P2MP tree is a graphical representation of all TE links that are committed for a particular P2MP LSP. In other words, a P2MP tree is a representation of the corresponding P2MP tunnel on the TE network topology. A sub-tree is a part of the P2MP tree describing how the root or an intermediate P2MP LSPs minimizes packet duplication when P2P TE sub-LSPs traverse common links. As described in [RFC5671] the computation of a P2MP tree requires three major pieces of information. The first is the path from the ingress LSR of a P2MP LSP to each of the egress LSRs, the second is the traffic engineering related parameters, and the third is the branch capability information.

Generally, an inter-domain P2MP tree (i.e., a P2MP tree with source and at least one destination residing in different domains) is particularly difficult to compute even for a distributed PCE architecture. For instance, while the BRPC recursive path computation may be well-suited for P2P paths, P2MP path computation involves multiple branching path segments from the source to the multiple destinations. As such, inter-domain P2MP path computation may result in a plurality of per-domain path options that may be difficult to coordinate efficiently and effectively between domains. That is, when one or more domains have multiple ingress and/or egress border nodes, there is currently no known technique for one domain to determine which border routers another domain will utilize for the inter-domain P2MP tree, and no way to limit the computation of the P2MP tree to those utilized border nodes.

A trivial solution to the computation of inter-domain P2MP tree would be to compute shortest inter-domain P2P paths from source to each destination and then combine them to generate an inter-domain, shortest-path-to-destination P2MP tree. This solution, however, cannot be used to trade cost to destination for overall tree cost (i.e., it cannot produce a MCT tree) and in the context of inter-domain P2MP LSPs it cannot be used to reduce the number of domain border nodes that are transited.

Computing P2P LSPs individually is not an acceptable solution for computing a P2MP tree. Even per domain path computation [RFC5152] can be used to compute P2P multi-domain paths, but it does not guarantee to find the optimal path which crosses multiple domains. Furthermore, constructing a P2MP tree from individual source to leaf P2P LSPs does not guarantee to produce a least-cost tree. This approach may also be considered to have scaling issues during LSP setup. That is, the LSP to each leaf is signaled separately, and each border node must perform path computation for each leaf.

P2MP Minimum Cost Tree (MCT), i.e. one which guarantees the least cost resulting tree, is an NP-complete problem. Moreover, adding and/or removing a single destination to/from the tree may result in an entirely different tree. In this case, frequent MCT path computation requests may prove computationally intensive, and the resulting frequent tunnel reconfiguration may even cause network destabilization. There are several heuristic algorithms presented in the literature that approximate the result within polynomial time that are applicable within the context of a single-domain.

This document presents a solution, and procedures and extensions to PCEP to support P2MP inter-domain path computation.

4. Assumptions

It is assumed that, due to deployment and commercial limitations (e.g., inter-AS peering agreements), the sequence of domains for a path (the path domain tree) will be known in advance.

In the figure below, the P2MP tree spans 6 domains, with D1 being the root domain. The corresponding domain sequences which are assumed known would be: D1-D3-D6, D1-D3-D5 and D1-D2-D4.

```
      D1
     / \
    D2  D3
   /   /  \
  D4  D5  D6
```

Figure 1: Domain Sequence Tree

The examples and scenarios used in this document are also based on the following assumptions:

o The PCE that serves each domain in the path domain tree is known, and the set of PCEs and their relationships is propagated to each PCE during the first exchange of path computation requests; [Editors note - this assumption needs to be more explicit.

- Each PCE knows about any leaf LSRs in the domain it serves;

- The boundary nodes to use on the LSP are pre-determined and are part of the path domain tree. [Editors Note - In this version of the document we do not consider multi-homed domains.]

Additional assumptions are documented in [RFC5441] and will not be repeated here.

5. Requirements

This section summarizes the requirements specific to computing inter-domain P2MP paths. In these requirements we note that the actual computation times by any PCE implementation are outside the scope of this document, but we observe that reducing the complexity of the required computations has a beneficial effect on the computation time regardless of implementation. Additionally, reducing the number of message exchanges and the amount of information exchanged will reduce the overall computation time for the entire P2MP tree. We refer to the "Complexity of the computation" as the impact on these aspects of path computation time as various parameters of the topology and the P2MP LSP are changed.

Its also important that the solution preserves confidentiality across domains, which is required when domains are managed by different Service Providers.

Other than the requirements specified in [RFC5376], a number of requirements specific to P2MP are detailed below:

1. The computed P2MP LSP should be optimal when only considering the paths among the BNs.

2. Grafting and pruning of multicast destinations in a domain should have no impact on other domains and on the paths among BNs.

3. The complexity of the computation for each sub-tree within each domain should be dependent only on the topology of the domain and it should be independent of the domain sequence.

4. The number of PCEP request and reply messages should be independent of the number of multicast destinations in each domain.

5. Specifying the domain entry and exit nodes.

6. Specifying which nodes should be used as branch nodes.

7. Reoptimization of existing sub-trees.

8. Computation of P2MP paths that need to be diverse from existing P2MP paths.

6. Objective Functions

For the computation of a single or a set of P2MP TE LSPs, a request to meet specific optimization criteria, called an Objective Function (OF) may be indicated.

The computation of one or more P2MP TE-LSPs may be subject to an OF in order to select the "best" candidate paths. A variety of objective functions have been identified as being important during the computation of inter-domain P2MP LSPs. These include:

1. The sub-tree within each domain should be optimized, which can be either the Minimum cost tree [RFC5862] or Shortest path tree [RFC5862].

2. The P2MP LSP path, formed by considering only the entry and exit nodes of the domains (the Core Tree) shold be optimal.

3. It should be possible to limit the number of entry points to a domain.

4. It should be possible to force the branches for all leaves within a domain to be in that domain.

7. P2MP Path Computation Procedures

The following sections describe the Core Tree based procedures to satisfy the requirements specified in the previous section. A core tree based solution provides an optimal inter-domain P2MP TE LSP.

7.1. Core Trees

A Core Tree is defined as a node tree, with nodes from the domains corresponding to the domain tree PCE topology, which satisfies the following conditions:

- The root of the core tree is the ingress LSR in the root domain;

- The leaves of the core tree are the entry nodes in the leaf domains;

- The transit and branch nodes of the core tree are from the entry and exit nodes from the transit and branch domains.

For example, consider the Domain Tree from the figure below, representing a domain tree of 5 domains, and part of the resulting Core Tree which satisfies the aforementioned conditions.

RN: Root Node
EN: Entry Border Node (domain, index)
XN: Exit Border Node (domain, index)

```
                        +----------------------+
                        | Domain1              |
                        |          (RN)        |
                        |                      |
                        |                      |
                        +--(XN1_1)---(XN1_2) --+
                            /               \
                           /                 \
                          /                   \
                +-------(EN2_1) -----+   +----(EN3_1)-------+
                | Domain2            |   | Domain3          |
                |                    |   |                  |
                |                    |   |                  |
                |                    |   |                  |
                +----(XN2_1)---------+   +-(XN3_1)-(XN3_2)--+
                      /                                 \
                     /                                   \
        +--------(EN4_1)-------+              +---------(EN5_1)-----+
        | Domain4              |              | Domain5             |
        |                      |              |                     |
        |                      |              |                     |
        |                      |              |                     |
        +----------------------+              +---------------------+
```

Figure 2: Domain Tree Example

```
                       (RN)
                      /    \
                 (XN1_1)   (XN3_1)
                    /        \
                 (EN2_1)   (EN3_1)
                  /              \
               (XN2_1)          (XN3_1)
                 /                  \
             (EN4_1)              (EN5_1)
```

Figure 3: Core Tree

7.2. Core Tree Computation Procedures

The algorithms to compute the optimal large core tree are outside scope of this document. The following extended BRPC based procedure can be used to compute the core tree.

BRPC Based Core Tree Path Computation Procedure:

1. Using the BRPC procedures to compute the VSPT(i) for each leaf BN(i), i=1 to n, where n is the total number of entry nodes for all the leaf domains. In each VSPT(i), there are a number of P(i) paths.

2. When the root PCE has computed all the VSPT(i), i=1 to n, take one path from each VSPT and form a set of paths, we call it a PathSet(j), j=1 to M, where M=P(1)xP(2)...xP(n);

3. For each PathSet(j), there are n S2L (Source to Leaf BN) paths and form these n paths into a Core Tree(j);

4. There will be M number of Core Trees computed from step3. Apply the OF to each of these M Core Trees and find the optimal Core Tree.

Note that the application of BRPC in the aforementioned procedure differs from the typical one since paths returned from a downstream PCE are not necessary pruned from the solution set by intermediate PCEs.

The reason for this is that if the PCE in a downstream domain does the prunning and returns the single optimal sub-path to its parent PCE, BRPC insures that the ingress PCE will get all the best optimal sub-paths for each LN (Leaf Border Nodes), but the combination of these single optimal sub-paths into a P2MP tree is not necessarily optimal even if each S2L (Source-to-Leaf) sub-path is optimal.

Without trimming, the ingress PCE will get all the possible S2L sub-paths set for LN, and eventually by looking through all the combinations, and taking one sub-path from each set to built one P2MP tree it finds the optimal tree.

The proposed method may present a scalability problem for the dynamic computation of the Core Tree (by iterative checking of all combinations of the solution space), specially with dense/meshed domains. Considering a domain sequence D1, D2, D3, D4, where the Leaf border node is at domain D4, PCE(4) will return 1 path. PCE(3) will return N paths, where N is E(3) x X(3), where E(k) x X(k) denotes the number of entry nodes times the number of exit nodes for that domain. PCE(2) will return M paths, where M = E(2) x X(2) x N = E(2) x X(2) x E(3) x X(3) x 1, etc. Generally speaking the number of potential paths at the ingress PCE Q = \prod E(k) x X(k).

Consequently, it is expected that the Core Path will be typically computed offline, without precluding the use of dynamic, online mechanisms such as the one presented here, in which case it SHOULD be possible to configure transit PCEs to control the number of paths sent upstream during BRPC (trading trimming for optimality at the point of trimming and downwards).

7.3. Sub Tree Computation Procedures

Once the core tree is built, the grafting of all the leaf nodes from each domain to the core tree can be achieved by a number of algorithms. One algorithm for doing this phase is that the root PCE will send the request with C bit set for the path computation to the destination(s) directly to the PCE where the destination(s) belong(s) along with the core tree computed from the phase 1.

This approach requires that the root PCE manage a potentially large number of adjacencies (either in persistent or non-persistent mode), including PCEP adjacencies to PCEs that are not within neighboring domains.

A first alternative would involve establishing PCEP adjacencies that correspond to the PCE domain tree. This would require that branch PCEs forward requests and responses from the root PCE towards the leaf PCEs and vice-versa.

Finally, another alternative would use a hierarchical PCE (H-PCE) architecture. The "hierarchically" parent would request sub tree path computations.

The algorithms to compute the optimal large sub tree are outside scope of this document. In the case that the number of destinations and the number of BNs within a domain are not big, the incremental procedure based on p2p path computation using the OSPF can be used.

7.4. PCEP Protocol Extensions

7.4.1. The Extension of RP Object

The extended format of the RP object body to include the C bit is as follows:

The C bit is added in the flag bits field of the RP object to signal the receiver of the message that the request/reply is for inter-domain P2MP Core Tree or not.

```
        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
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |    Reserved   | Flags                           |C|O|B|R| Pri |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                        Request-ID-number                      |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                                                               |
       //                      Optional TLV(s)                        //
       |                                                               |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: RP Object Body Format

The following flag is added in this draft:

C bit (P2MP Core Tree bit - 1 bit):

0: This indicates that this is normal PCReq/PCRrep for P2MP.

1: This indicates that this is PCReq or PCRep message for inter-domain Core Tree P2MP. When the C bit is set, then the request message should have the Core Tree passed along with the destinations which and then graphed to the tree.

7.4.2. The PCE Sequence Object

The PCE Sequence Object is added to the existing PCE protocol. A list of this objects will represent the PCE topology tree. A list of Sequence Objects can be exchanged between PCEs during the PCE capability exchange or on the first path computation request message between PCEs. In this case, the request message format needs to be changed to include the list of PCE Sequence Objects for the PCE inter-domain P2MP calculation request.

Each PCE Sequence can be obtained from the domain sequence for a specific path. All the PCE sequences for all the paths of P2MP inter-domain form the PCE Topology Tree of the P2MP LSP.

The format of the new PCE Sequence Object 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
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       | Object-Class  |   OT  |Res|P|I|   Object Length (bytes)       |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                    IPv4 address for root PCE                  |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                IPv4 address for the downstream PCE            |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                IPv4 address for the downstream PCE            |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |                         !!                                    |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       |   IPv4 address for the PCE corresponding to the leafDomain    |
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: The New PCE Sequence Object Body Format for IPv4

The format of the new PCE Sequence Object for IPv6 (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
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         | Object-Class  |   OT  |Res|P|I|   Object Length (bytes)       |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |                    IPv6 address for root PCE                  |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |                IPv6 address for the downstream PCE            |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |                IPv6 address for the downstream PCE            |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |                         !!                                    |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |  IPv6 address for the PCE corresponding to the leafDomain     |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: The New PCE Sequence Object Body Format for IPv6

7.5. Relationship with Hierarchical PCE

The actual grafting of subtrees into the Multi-Domain tree needs to be carried out by the source node. This means that the source node needs to get the computed sub-trees from all the involved domains. This requires that the source node either has a PCEP session with all the PCEs, or PCEP messages are routed via the PCEP sessions. This may mean an excessive number of sessions or an added complexity in implementations.

Alternatively, one may use an architecture based on the concept of hierarchical PCE [H-PCE]. The parent PCE would be responsible to request Intra-domain subtrees to the PCEs, combine them and return the overall P2MP tree.

7.6. Parallelism

In order to minimize latency in path computation in multi-domain networks, intra-domain path segments and intra-domain sub-trees SHOULD be computed in parallel when possible. The proposed procedures in this draft present opportunities for parallelism:

1. The BRPC procedure for each leaf node can be launched in parallel by the ingress/root PCE if the dynamic computation of the Core Tree is enabled.

2. Intra-domain P2MP paths can also be computed in parallel by the PCEs once the entry and exit nodes within a domain are known

One of the potential issues of parallelism is that the ingress PCE would require a potentially high number of PCEP adjacencies to "remote" PCEs and that may not be desirable, but a given PCE would only receive requests for the destinations that are in its domain (+ the core nodes), without PCEs forwarding requests.

8. Protection

It is envisaged that protection may be required when deploying and using inter-domain P2MP LSPs. The procedures and mechanisms defined in this document do not prohibit the use of existing and proposed types of protection, including: end-to-end protection [RFC4875] and domain protection schemes.

Segment or facility (link and node) protection is problematic in inter-domain environment due to the limit of Fast-reroute (FRR) [RFC4875] requiring knowledge of its next-hop across domain boundaries whilst maintaining domain confidentiality. Although the FRR protection might be implemented if manually provisioned if next-hop information was known in advance.

8.1 End-to-end Protection

8.2 Domain Protection

9. Manageability Considerations

[RFC5862] describes various manageability requirements in support of P2MP path computation when applying PCEP. This section describes how manageability requirements mentioned in [RFC5862] are supported in the context of PCEP extensions specified in this document.

Note that [RFC5440] describes various manageability considerations in PCEP, and most of manageability requirements mentioned in [PCE-P2MP P2MP] are already covered there.

9.1 Control of Function and Policy

In addition to PCE configuration parameters listed in [RFC5440], the following additional parameters might be required:

- The ability to enable or disable single domain P2MP path computations on the PCE.
- The ability to enable or disable multi-domain P2MP path computations on the PCE.
- The PCE may be configured to enable or disable the advertisement of its single domain and multi-domain P2MP path computation capability.

9.2. Information and Data Models

A number of MIB objects have been defined for general PCEP control and monitoring of P2P computations in [PCEP-MIB]. [RFC5862] specifies that MIB objects will be required to support the control and monitoring of the protocol extensions defined in this document. [PCEP-P2MP-MIB] describes managed objects for modeling of PCEP communications between a PCC and PCE, and PCE to PCE, P2MP path computation requests and responses.

Core tree management configuration. TBD.

9.3. Liveness Detection and Monitoring

No changes are necessary to the liveness detection and monitoring requirements as already embodied in [RFC4657].

It should be noted that multi-domain P2MP computations are likely to take longer than P2P computations, and single domain P2MP computatoins. The liveness detection and monitoring features of the PCECP SHOULD take this into account.

9.4. Verifying Correct Operation

There are no additional requirements beyond those expressed in [RFC4657] for verifying the correct operation of the PCECP. Note that verification of the correct operation of the PCE and its algorithms is out of scope for the protocol requirements, but a PCC MAY send the same request to more than one PCE and compare the results.

9.5. Requirements on Other Protocols and Functional Components

A PCE operates on a topology graph that may be built using information distributed by TE extensions to the routing protocol operating within the network. In order that the PCE can select a suitable path for the signaling protocol to use to install the P2MP LSP, the topology graph must include information about the P2MP signaling and branching capabilities of each LSR in the network.

Mechanisms for the knowledge of other domains, the discovery of corresponding PCEs and their capabilities should be provided and that this information MAY be collected by other mechanisms.

Whatever means is used to collect the information to build the topology graph, the graph MUST include the requisite information. If the TE extensions to the routing protocol are used, these SHOULD be as described in [RFC5073].

9.6. Impact on Network Operation

The use of a PCE to compute P2MP paths is not expected to have significant impact on network operations. However, it should be noted that the introduction of P2MP support to a PCE that already provides P2P path computation might change the loading of the PCE significantly, and that might have an impact on the network behavior, especially during recovery periods immediately after a network failure.

The dynamic computation of Core Trees might also have an impact on the load of the involved PCEs as well as path computation times.

9.7 Policy Control

TBD

10. Security Considerations

As described in [RFC5862], P2MP path computation requests are more CPU-intensive and also utilize more link bandwidth. In the event of an unauthorized P2MP path computation request, or a denial of service attack, the subsequent PCEP requests and processing may be disruptive to the network. Consequently, it is important that implementations conform to the relevant security requirements of [RFC5440] that specifically help to minimize or negate unauthorized P2MP path computation requests and denial of service attacks. These mechanisms include:

- Securing the PCEP session requests and responses using TCP security techniques (Section 10.2 of [RFC5440]).
- Authenticating the PCEP requests and responses to ensure the message is intact and sent from an authorized node (Section 10.3 of [RFC5440]).
- Providing policy control by explicitly defining which PCCs, via IP access-lists, are allowed to send P2MP path requests to the PCE (Section 10.6 of [RFC5440]).

PCEP operates over TCP, so it is also important to secure the PCE and PCC against TCP denial of service attacks. Section 10.7.1 of [RFC5440] outlines a number of mechanisms for minimizing the risk of TCP based denial of service attacks against PCEs and PCCs.

PCEP implementations SHOULD also consider the additional security provided by the TCP Authentication Option (TCP-AO) [RFC5925].

11. IANA Considerations

A new flag of the RP object (specified in [RFC5440]) is defined in this document.

Additional requirements for IANA will be included in future revisions of this document.

12. Acknowledgements

The authors would like to thank Adrian Farrel, Dan Tappan and Olufemi Komolafe for their valuable comments on this draft.

13. References

13.1. 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.

[RFC6006] Takeda, T., Chaitou M., Le Roux, J.L., Ali Z., Zhao, Q., King, D., "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", RFC6006, September 2010.

13.2. Informative References

[RFC4461] Yasukawa, S., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.

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

[RFC4657] Ash, J., Ed., and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.

[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.

[RFC5073] Vasseur, J., Ed., and J. Le Roux, Ed., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", RFC 5073, December 2007.

[RFC5152] Vasseur, JP., Ayyangar, A., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, February 2008.

[RFC5376] Bitar, N., Zhang, R., and K. Kumaki, "Inter-AS Requirements for the Path Computation Element Communication Protocol (PCECP)", RFC 5376, November 2008.

[RFC5441] Roux, J., Vasseur, J., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", RFC 5441, June 2009.

[RFC5520] Bradford, R., Ed., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.

[RFC5671] Yasukawa, S. and A. Farrel, "Applicability of the Path Computation Element (PCE) to Point-to-Multipoint (P2MP) Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Traffic Engineering (TE)", RFC 5671, August 2009.

[RFC5862] Yasukawa, S. and A. Farrel, "PCC-PCE Communication Requirements for Point to Multipoint Multiprotocol Label Switching Traffic Engineering (MPLS-TE)", RFC 5862, June 2010.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

[H-PCE] King, D. and A. Farrel, "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS & GMPLS", July 2010.

[PCEP-MIB] Koushik, K., Stephan, E., Zhao, Q., and D. King, "PCE communication protocol (PCEP) Management Information Base", Work in Progress, July 2010.

[PCEP-P2MP-MIB] Zhao, Q., Dhody, D., Palle, U., King, D., "Management Information Base for the PCE Communications Protocol (PCEP) When Requesting Point-to-Multipoint Services", Work in Progress, September 2010.

Authors’ Addresses

Quintin Zhao
Huawei Technology
125 Nagog Technology Park
Acton, MA 01719
US
Email: qzhao@huawei.com

Zafar Ali
Cisco Systems
US
Email: zali@cisco.com

Tarek Saad
Cisco Systems
US
Email: tsaad@cisco.com

Siva Sivabalan
Cisco Systems
Canada
Email: msiva@cisco.com

Daniel King
Old Dog Consulting
UK
Email: daniel@olddog.co.uk

Ramon Casellas
CTTC - Centre Tecnologic de Telecomunicacions de Catalunya
Spain
Email: ramon.casellas@cttc.es