Network Working Group Eiji Oki
Internet Draft NTT
Category: Informational Jean-Louis Le Roux
Expires: October 2006 France Telecom
Adrian Farrel
Old Dog Consulting
April 2006

Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering

draft-oki-pce-inter-layer-frwk-00.txt

Abstract

A network may comprise of multiple layers. It is important to globally optimize network resources utilization, taking into account all layers, rather than optimizing resource utilization at each layer independently. This allows better network efficiency to be achieved through a process that we call inter-layer traffic engineering. The Path Computation Element (PCE) can be a powerful tool to achieve inter-layer traffic engineering.

This document describes a framework for the PCE-based path computation architecture to inter-layer MPLS and GMPLS traffic engineering. It provides suggestions for the deployment of PCE in

support of multi-layer networks. This document also describes
network models where PCE performs inter-layer traffic engineering,

and the relationship between PCE and a functional component called the Virtual Network Topology Manager (VNTM).

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Terminology

This document uses terminology from the PCE-based path computation Architecture [PCE-ARCH] and also common terminology from Multi Protocol Label Switching (MPLS) [RFC3031], Generalized MPLS (GMPLS) [RFC3945] and Multi-Layer Networks [MLN-REQ].

2. Introduction

A network may comprise of multiple layers. These layers may represent separations of technologies (e.g., packet switch capable (PSC), time division multiplex (TDM) lambda switch capable (LSC)) [RFC3945], separation of data plane switching granularity levels (e.g. PSC-1, PSC-2, VC4, VC12) [MLN-REQ], or a distinction between client and server networking roles. In this multi-layer network, LSPs in a lower layer are used to carry higher-layer LSPs across

the lower-layer network. The network topology formed by lower-layer LSPs and advertised to the higher layer is called a Virtual Network Topology (VNT) [MLN-REQ].

It is important to optimize network resource utilization globally, i.e. taking into account all layers, rather than optimizing resource utilization at each layer independently. This allows better network efficiency to be achieved and is what we call inter-layer traffic engineering. This includes mechanisms allowing the computation of end-to-end paths across layers (known as inter-layer path computation), and mechanisms for control and management of the VNT by setting up and releasing LSPs in the lower layers [MLN-REQ].

Inter-layer traffic engineering is included in the scope of the PCE-based path computation architecture [PCE-ARCH], and PCE can provide a suitable mechanism for resolving inter-layer path computation issues.

PCE Communication Protocol requirements for inter-layer traffic engineering are set forth in [PCE-INTER-LAYER-REQ].

This document describes a framework for the PCE-based path computation Architecture to inter-layer traffic engineering. It provides suggestions for the deployment of PCE in support of multi-layer networks. This document also describes network models where PCE performs inter-layer traffic engineering, and the relationship between PCE and a functional component in charge of the control and management of the VNT, and called the Virtual Network Topology Manager (VNTM).

3. Inter-Layer Path Computation

This section describes key topics of inter-layer path computation in MPLS and GMPLS networks.

[RFC4206] defines a way to signal a higher-layer LSP, whose explicit route includes hops traversed by LSPs in lower layers. The computation of end-to-end paths across layers is called Inter-Layer Path Computation.

An LSR in the higher-layer may not have information on the lower-layer topology, particularly in an overlay or augmented model, and hence may not be able to compute an end-to-end path across layers.

PCE-based inter-layer path computation, consists of relying on one or more PCEs to compute an end-to-end path across layers. This could rely on a single PCE path computation where the PCE has topology information about multiple layers and can directly compute an end-to-end path across layers considering the topology of all of the layers. Alternatively, the inter-layer path computation could be performed as a multiple PCE computation where each member of a set of PCEs have information about the topology of one or more layers, but not all layers, and collaborate to compute an end-to-end path.

Consider a two-layer network where the higher-layer network is a packet-based IP/MPLS network or GMPLS network and the lower-layer network is a GMPLS optical network. An ingress LSR in the higher-layer network tries to set up an LSP to an egress LSR also in the higher-layer network across the lower-layer network, and needs a path in the higher-layer network. However, suppose that there is no TE link between border LSRs, which are located on the boundary between the higher-layer and lower-layer networks, and that the ingress LSR does not have topology visibility in the lower layer. If a single-layer path computation is applied for the higher-layer, the path computation fails. On the other hand, inter-layer path computation is able to provide a route in the higher-layer and a suggestion that a lower-layer LSP be setup between border LSRs, considering both layers' TE topologies.

Lower-layer LSPs form a Virtual Network Topology (VNT), which can be used for routing higher-layer LSPs or to carry IP traffic. Inter-layer path computation for end-to-end LSPs in the higher-layer network that span the lower-layer network may utilize the VNT, and PCE is a candidate for computing the paths of such higher-layer LSPs within the higher-layer network. The PCE-based path computation model can:

- Perform a single computation on behalf of the ingress LSR using information gathered from more than one layer. This mode is referred to as Single PCE Computation in [PCE-ARCH].

- Compute a path on behalf of the ingress LSR through cooperation between PCEs responsible for each layer. This mode is referred to as Multiple PCE Computation with inter-PCE communication in [PCE-ARCH].

- Perform separate path computations on behalf of the TE-LSP head-end and each transit LSR that is the entry point to a new layer.

This mode is referred to as Multiple PCE Computation (without inter-PCE communication) in [PCE-ARCH]. This option utilizes per-layer path computation performed independently by successive PCEs.

The PCE computes and returns a path to the PCC that the PCC can use to build an MPLS or GMPLS LSP once converted to an Explicit Route Object (ERO) for use in RSVP-TE signaling. There are two options.

- Option 1: Mono-layer path.
The PCE computes a "mono layer" path, i.e. a path that includes only TE-links from the same layer. There are two cases for this option. In the first case the PCE computes a path that includes already established lower-layer LSPs: that is the resulting ERO includes sub-object(s) corresponding to lower-layer hierarchical LSPs expressed as the TE link identifiers, which can be numbered or unnumbered ones, of the hierarchical LSPs when advertised as TE links in the higher-layer network. The TE link may be a regular TE link that is actually established, or a virtual TE link that is not established yet (see [MLN-REQ]). If it is a regular TE link, this

does not trigger new lower-layer LSP setup, but the utilization of existing lower-layer LSPs. If it is a virtual TE link, this triggers a new lower-layer LSP setup (provided that there are available resources in the lower layer). A transit LSR corresponding to the entry point of the virtual TE link is expected to trigger the new lower-layer LSP setup. Note that the path of a virtual TE link is not necessarily known in advance, and this may require path computation either on the entry point or on a PCE. The second case is that the PCE computes a path that includes loose hop(s). The higher layer would select which lower layers to use and would select the entry and exit points from those layers, but would not select the path across the layers. A transit LSR corresponding to the entry point is expected to expand the loose hop (either itself or relying on the services of a PCE). Path expansion process on border LSR may result either in the selection of an existing lower-layer LSP, or in the computation and setup of a new lower-layer LSP.

- Option 2: Multi-layer path. The PCE computes a "multi-layer" path, i.e. a path that includes TE links from distinct layers [RFC4206]. Such a path can include the complete path of one or more lower-layer LSPs that already exist or are not yet established. In the latter case, the signaling of the higher-layer LSP will trigger the establishment of the lower-layer LSPs.

4. Inter-layer Path Computation Models

As stated in Section 3, two PCE modes defined in the PCE architecture can be used to perform inter-layer path computation. They are discussed below.

4.1. Single PCE Inter-Layer Path Computation

In this model Inter-layer path computation is performed by a single PCE that has topology visibility in all layers. Such a PCE is called a multi-layer PCE.

In Figure 1, the network is comprised of two layers. LSR H1, H2, H3 and H4 belong to the higher layer, and LSRs L1 and L2 belong to the lower layer. The PCE is a multi-layer PCE that has visibility into both layers. It can perform end-to-end path computation across layers (single PCE path computation). For instance, it can compute an optimal path H2-L1-L2-H3-H4, for a higher layer LSP from H1 to H4. This path includes the path of a lower layer LSP from H2 to H3, already established or not.

```
                        -----
                       | PCE |
                        -----
     -----    -----                  -----    -----
    | LSR |--| LSR |................| LSR |--| LSR |
    | H1  |  | H2  |                | H3  |  | H4  |
     -----    -----\                /-----    -----
```



```
                    \-----    -----/
                    | LSR |--| LSR |
                    | L1  |  | L2  |
                     -----    -----
```

  Figure 1 : Multi-Layer PCE - A single PCE with multi-layer visibility

4.2. Multiple PCE Inter-Layer Path Computation

In this model there is at least one PCE per layer, and each PCE has topology visibility restricted to its own layer. These PCEs are called mono-layer PCEs. Mono-layer PCEs collaborate to compute an end-to-end optimal path across layers.

In Figure 2, there is one PCE in each layer. The PCEs from each layer collaborate to compute an end-to-end path across layers. PCE Hi is responsible for computations in the higher layer and may consult with PCE Lo to compute paths across the lower layer. PCE Lo is responsible for path computation in the lower layer. A simple example of cooperation between the PCEs could be: PCE Hi requests a path H2-H3 from PCE Lo. Of course more complex cooperation may be required if an end-to-end optimal path is desired.

```
                             -----
                            | PCE |
                            | Hi  |
                             --+--
                               |
       -----    -----          |           -----    -----
      | LSR |--| LSR |.........|..........| LSR |--| LSR |
      | H1  |  | H2  |         |          | H3  |  | H4  |
       -----    -----\       --+--        /-----    -----
                      \     | PCE |      /
                       \    | Lo  |     /
                        \    -----     /
                         \            /
                          \-----    -----/
                          | LSR |--| LSR |
                          | L1  |  | L2  |
                           -----    -----
```

Figure 2 : Cooperating Mono-Layer PCEs - Multiple PCEs with single-layer visibility

4.3. General observation

- Depending on implementation details, inter-layer path computation time in the Single PCE inter-layer path computation model may be less than that of the Multiple PCE model with cooperating mono-layer PCEs, because there is no requirement to exchange messages between cooperating PCEs.

- When TE topology for all layered networks is visible within one routing domain, the single PCE inter-layer path computation model may be adopted because a PCE is able to collect all layers' TE topologies by participating in only one routing domain.

- As the single PCE inter-layer path computation model uses more TE topology information than is used by PCEs in the Multiple PCE path computation model, it requires more computation power and memory.

5. Inter-Layer Path Control

5.1. VNT Management

As a result of inter-layer path computation, a PCE may determine that there is insufficient bandwidth available in the higher-layer network to support this or future higher-layer LSPs. The problem might be resolved if new LSPs are provisioned across the lower-layer network. Further, the modification, re-organization and new provisioning of lower-layer LSPs may enable better utilization of lower-layer network resources given the demands of the higher-layer network. In other words, the VNT needs to be controlled or managed in cooperation with inter-layer path computation.

A VNT Manager (VNTM) is defined as a network element that manages and controls the VNT. PCE and "VNT Management" are distinct functions that may or may not be co-located. To describe each function clearly, VNTM is considered as a functional element in this draft.

5.2. Inter-Layer Path Control Models

5.2.1. Cooperation model between PCE and VNTM

```
      -----      ------
     | PCE |--->| VNTM |
      -----      ------
        ^          :
        :          :
        :          :
        v          V
       -----      -----                  -----      -----
      | LSR |----| LSR |................| LSR |----| LSR |
      | H1  |    | H2  |                | H3  |    | H4  |
       -----      -----\               /-----      -----
                        \-----    -----/
                        | LSR |--| LSR |
                        | L1  |  | L2  |
                         -----    -----
```

Figure 3: Cooperation model between PCE and VNTM

A multi-layer network consists of higher-layer and lower-layer networks. LSRs H1, H2, H3, and H4 belong to the higher-layer network, LSRs H2, L1, L2, and H3 belong to the lower-layer network,

as shown in Figure 3. Consider that H1 requests PCE to compute an inter-layer path between H1 and H4. There is no TE link in the higher-layer between H2 and H3 before the path computation request.

The roles of PCE and VNTM are as follows. PCE performs inter-layer path computation and is unable to supply a path because there is not TE link between H2 and H3. The computation fails, but PCE suggests to VNTM that a lower-layer LSP (H2-H3) should be established to support future LSP requests. VNTM uses local policy and possibly management/configuration input to determine how to process the suggestion from PCE, and may request an ingress LSR (e.g. H2) to establish a lower-layer LSP. VNTM or the ingress LSR (H2) may use a PCE with visibility into the lower layer to compute the path of this new LSP.

If the PCE cannot compute a path for the higher-layer LSP without the establishment of a further lower-layer LSP, the PCE may notify VNTM and wait for the lower-layer LSP to be set up and advertised as a TE link. It can then compute the complete end-to-end path for the higher-layer LSP and return the result to the PCC. In this case, the PCC may be kept waiting some time, and it is important that the PCC understands this. It is also important that the PCE and VNTM have an agreement that the lower-layer LSP will be set up in a timely manner, the PCE operates a timeout, or the PCE will be notified by VNTM that no new LSP will become available. An example of such a cooperative procedure between PCE and VNTM is as follows.

Step 1: H1 (PCC) requests PCE to compute a path between H1 and H4. In the request, it indicates that inter-layer path computation is allowed.

Step 2: As a result of the inter-layer path computation, PCE judges that a new lower-layer LSP needs to be established.

Step 3: PCE suggests to VNTM that a new lower-layer LSP should be established if necessary and if acceptable within VNTM s policy constraints. The inter-layer path route computed by PCE may include one or more virtual TE links. If PCE knows the inclusion of the virtual TE link(s) in the inter-layer route, PCE may suggest VNTM that the corresponding new lower-layer LSP(s) should be established. Otherwise, new lower-layer LSP(s) may be setup according to the higher-layer signaling trigger model.

Step 4: VNTM requests an ingress LSR (e.g. H2) to establish a lower-layer LSP. The request message may include a pre-computed lower-layer LSP route obtained from the PCE responsible for the lower-layer network.

Step 5: The ingress LSR starts signaling to establish a lower-layer LSP.

Step 6: If the lower-layer LSP setup is completed, the ingress LSR notifies VNTM that the LSP is complete and supplies the tunnel information.

Step 7: VNTM replies to PCE to inform it that the lower-layer LSP is now established, and includes the lower-layer tunnel information. Alternatively, PCE may get to know about the existence of the lower-layer LSP when a new TE link in the higher-layer corresponding to the lower-layer LSP is advertised to PCE through the IGP.

Step 8: PCE replies to H1 (PCC) with a computed higher-layer LSP route. The computed path is categorized as a mono-layer path that includes the already-established lower layer-LSP. The higher-layer route is specified as H2-H3-H4, where all hops are strict.

Step 9: H1 initiates signaling with the computed path H2-H3-H4 to establish the higher-layer LSP.

5.2.2. Higher-Layer Signaling Trigger Model

```
   -----
  | PCE |
   -----
     ^
     :
     :
     v
    -----      -----                  -----    -----
   | LSR |----| LSR |................| LSR |--| LSR |
   | H1  |    | H2  |                | H3  |  | H4  |
    -----      -----\               /-----    -----
                     \-----    -----/
                     | LSR |--| LSR |
                     | L1  |  | L2  |
                      -----    -----
```

Figure 4: Higher-layer signaling trigger model

Figure 4 shows the higher-layer signaling trigger model. As in the case described in section 5.2.1, consider that H1 requests PCE to compute an inter-layer path between H1 and H4. There is no TE link in the higher-layer between H2 and H3 before the path computation request.

If PCE judges that a lower-layer LSP needs to be established based on the inter-layer path computation result, a lower-layer LSP is established during the higher-layer signaling procedure. After PCE completes inter-layer path computation, PCE sends a reply message including explicit route to the ingress LSR (PCC). There are two ways to express the higher-layer LSP route, which are a multi-layer path and a mono-layer path that includes loose hop(s).

In the higher-layer signaling trigger model with a multi-layer path, a high-layer LSP route includes a route for a lower-layer LSP that is not yet established. An LSR that is located at the boundary between the higher-layer and lower-layer networks, called a border

LSR, receives a higher-layer signaling message and then starts to setup the lower-layer LSP.

An example procedure of the signaling trigger model with a multi-layer path is as follows.

Step 1: H1 (PCC) requests PCE to compute a path between H1 and H4. The request indicates that inter-layer path computation is allowed.

Step 2: As a result of the inter-layer path computation, PCE judges that a new lower-layer LSP needs to be established.

Step 3: PCE replies to H1 (PCC) with a computed multi-layer route including higher-layer and lower-layer LSP routes. The route may be specified as H2-L1-L2-H3-H4, where all hops are strict.

Step 4: H1 initiates higher-layer signaling using the computed explicit router of H2-L1-L2-H3-H4.

Step 5: The border LSR (H2) that receives the higher-layer signaling message starts lower-layer signaling to establish a lower-layer LSP along the specified lower-layer route of L1-L2-H3. That is, the border LSR recognizes the hops within the explicit route that apply to the lower-layer network, verifies with local policy that a new LSP is acceptable, and establishes the required lower-layer LSP. Note that it is possible that a suitable lower-layer LSP has been established (or become available) between the time that the computation was performed and the moment when the higher-layer signaling message reached the border LSR. In this case, the border LSR may select such a lower-layer LSP without the need to signal a new LSP provided that the lower-layer LSP satisfies the explicit route in the higher-layer signaling request.

Step 6: After the lower-layer LSP is established, the higher-layer signaling continues along the specified higher-layer route of H2-H3-H4.

On the other hand, in the signaling trigger model with mono-layer path, a higher-layer LSP route includes a loose or strict hop to traverse the lower-layer network between the two border LSRs. In the strict hop case, a virtual TE link may be advertised, but a lower-layer LSP is not setup. A border LSR that receives a higher-layer signaling message needs to determine a path for a new lower-layer LSP. It applies local policy to verify that a new LSP is acceptable and then either consults a PCE with responsibility for the lower-layer network or computes the path by itself, and initiates signaling to establish a lower-layer LSP. Again, it is possible that a suitable lower-layer LSP has been established (or become available) between the time that the higher-layer computation was performed and the moment when the higher-layer signaling message reached the border LSR. In this case, the border LSR may select such a lower-layer LSP without the need to signal a new LSP provided that the lower-layer LSP satisfies the explicit route in the higher-layer signaling request. Since the higher-layer

signaling request used a loose hop without specifying any specifics of the path within the lower-layer network, the border LSR has greater freedom to choose a lower-layer LSP than in the previous example.

The difference between procedures of the signaling trigger model with a multi-layer path and a mono-layer path is Step 5. Step 5 of the signaling trigger model with a mono layer path is as follows:

Step 5: The border LSR (H2) that receives the higher-layer signaling message applies local policy to verify that a new LSP is acceptable and then initiates establishment of a lower-layer LSP. It either consults a PCE with responsibility for the lower-layer network or computes the route by itself to expand the loose hop route in the higher-layer path.

5.2.3. Examples of multi-layer ERO

```
PCE
 ^
 :
 :
 V
H1--H2             H3--H4
     \                 /
      L1==L2==L3--L4--L5
               |
               |
              L6--L7
                    \
                     H5--H6
```

Figure 5 Example of multi-layer network

This section describes how lower-layer LSP setup is performed in the higher-layer signaling trigger model using an ERO that can include subobjects in both the higher and lower layers. It gives rise to several options for the ERO when it reaches the last LSR in the higher layer network (H2).
1. The next subobject is a loose hop to H3 (mono layer ERO).
2. The next subobject is a strict hop to L1 followed by a loose hop to H3.
3. The next subobjects are a series of hops (strict or loose) in the lower-layer network followed by H3. For example, {L1(strict), L3(loose), L5(loose), H3(strict)}

In the first, the lower layer can utilize any LSP tunnel that will deliver the end-to-end LSP to H3. In the third case, the lower layer must select an LSP tunnel that traverses L3 and L5. However, this does not mean that the lower layer can or should use an LSP from L1 to L3 and another from L3 to L5.

6. Choosing between inter-layer path control models

This section compares the cooperation model between PCE and VNTM, and the higher-layer signaling trigger model, in terms of VNTM functions, border LSR functions, and higher-layer signaling time.

VNTM functions:

In the cooperation model, VNTM functions are required. In this model, additional overhead communications between PCE and VNTM and between VNTM and a border LSR are required.

In the higher-layer signaling trigger model, no VNTM functions are required, and no such communications are required.

If VNTM functions are not supported in a multi-layer network, the higher-layer signaling trigger model has to be chosen.

The inclusion of VNTM functionality allows better coordination of cross-network LSP tunnels and application of network-wide policy that is not available in the trigger model.

Border LSR functions:

In the higher-layer signaling trigger model, a border LSR must have some additional functions. It needs to trigger lower-layer signaling when a higher-layer path message suggests that lower-layer LSP setup is necessary. The triggering signaling is also required in the cooperation case when the VNTM support virtual TE links. Note that, if only the cooperation model is applied, it is required that a PCE knows whether a link is a regular TE link or virtual TE link.

If the ERO in the higher-layer Path message uses a mono-layer path or specifies loose hop, a border LSR receiving the Path message MUST obtain a lower-layer route either by consulting PCE or by using its own computation engine. If the ERO in the higher-layer Path message uses multi-layer path, the border LSR MUST judge whether lower-layer signaling is needed.

In the cooperation model, no additional function for triggered signaling in border LSRs is required except when virtual TE links are used. Therefore, if these additional functions are not supported in border LSRs, the cooperation model, where a border LSR is controlled by VNTM to set up a lower-layer LSP, has to be chosen.

Complete inter-layer LSP setup time:

Complete inter-layer LSP setup time includes inter-layer path computation, signaling, and communication time between PCC and PCE, PCE and VNTM, and VNTM and LSR. In the cooperation model, the additional communication steps are required compared with the higher-layer signaling trigger model. On the other hand, the cooperation model provides better control at the cost of a longer service setup time.

Note that, in terms of higher-layer signaling time, in the higher-layer signaling trigger model, the required time from when higher-layer signaling starts to when it is completed, is more than that of the cooperation model except when any virtual TE link is included. This is because the former model requires lower-layer signaling to take place during the higher-layer signaling. A higher-layer ingress LSR has to wait for more time until the higher-layer signaling is completed. A higher-layer ingress LSR is required to be tolerant of longer path setup times.

An appropriate model is chosen, taking into all of the above considerations.

7. Security Considerations

Inter-layer traffic engineering with PCE may raise new security issues in both inter-layer path control models.

In the cooperation model between PCE and VNTM, when PCE judges a new lower-layer LSP, communications between PCE and VNTM and between VNTM and a border LSR are needed. In this case, there are some security concerns that need to be addressed for these communications. These communications should have some security mechanisms to ensure authenticity, privacy and integrity.

In the higher-layer signaling trigger model, there are several security concerns. First, PCE may inform PCC, which is located in the higher-layer network, of multi-layer path information that includes an ERO in the lower-layer network, while the PCC may not have TE topology visibility into the lower-layer network. This raises a security concern, where lower-layer hop information is known to transit LSRs supporting a higher-layer LSP. Some security mechanisms to ensure authenticity, privacy and integrity may be used.

Security issues may also exist when a single PCE is granted full visibility of TE information that applies to multiple layers.

8. Acknowledgment

We would like to thank Kohei Shiomoto, Ichiro Inoue, Julien Meuric and Jean-Francois Peltier for their useful comments.

9. References

9.1. Normative Reference

[RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.

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

[RFC4206] Kompella, K., and Rekhter, Y., "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.

[RFC4208] G. Swallow et al., "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.

9.2. Informative Reference

[PCE-ARCH] A. Farrel, JP. Vasseur and J. Ash, "Path Computation Element (PCE) Architecture", draft-ietf-pce-architecture (work in progress).

[PCE-COM-REQ] J. Ash, J.L Le Roux et al., "PCE Communication Protocol Generic Requirements", draft-ietf-pce-comm-protocol-gen-reqs (work in progress).

[PCE-DISC-REQ] JL Le Roux et al., "Requirements for Path Computation Element (PCE) Discovery", draft-ietf-pce-discovery-reqs (work in progress).

[MLN-REQ] K. Shiomoto et al., "Requirements for GMPLS-based multi-region networks (MRN) ", draft-ietf-ccamp-gmpls-mln-reqs (work in progress).

[PCE-INTER-LAYER-REQ] E. Oki et al., "PCC-PCE Communication Requirements for Inter-Layer Traffic Engineering", draft-ietf-pce-inter-layer-req (work in progress).

[PCEP] JP. Vasseur et al, "Path Computation Element (PCE) communication Protocol (PCEP) - Version 1 - draft-ietf-pce-pcep (work in progress).

10. Authors' Addresses

Eiji Oki
NTT
3-9-11 Midori-cho,
Musashino-shi, Tokyo 180-8585, Japan
Email: oki.eiji@lab.ntt.co.jp

Jean-Louis Le Roux
France Telecom R&D,
Av Pierre Marzin,
22300 Lannion, France
Email: jeanlouis.leroux@francetelecom.com

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

11. Intellectual Property Statement