Network Working Group J.-L. Le Roux (Editor)
Internet Draft France Telecom

Category: Informational
Expires: April 2006
October 2005

PCE Communication Protocol (PCECP) specific requirements for Inter-Area (G)MPLS Traffic Engineering

draft-leroux-pce-pcecp-interarea-reqs-00.txt

Status of this Memo

Abstract

For scalability purposes a network may comprise multiple IGP areas. An inter-area TE-LSP is an LSP that transits through at least two IGP areas. In a multi-area network, topology visibility remains local to a given area, and a head-end LSR cannot compute alone an inter-area shortest constrained path. One key application of the Path Computation Element (PCE) architecture is the computation of inter-area TE-LSP paths. In this context, this document lists a detailed set of PCE Communication Protocol (PCECP) specific requirements for support of inter-area TE-LSP path computation. It complements generic requirements for a PCE Communication Protocol.

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.

Table of Contents

1. Contributors

The following are the authors that contributed to the present document:

Jerry Ash (AT&T)
Dean Cheng (Cisco)
Kenji Kumaki (KDDI)
J.L. Le Roux (France Telecom)
Eiji Oki (NTT)
Nabil Bitar (Verizon)
Raymond Zhang (BT Infonet)

2. Terminology

LSR: Label Switching Router

LSP: MPLS Label Switched Path

TE-LSP: Traffic Engineering Label Switched Path

IGP area: OSPF Area or IS-IS level

ABR: IGP Area Border Router, a router that is attached to more than one IGP areas (ABR in OSPF or L1/L2 router in IS-IS)

Inter-Area TE LSP: TE LSP that traverses more than one IGP area

CSPF: Constraint Shortest Path First

SRLG: Shared Risk Link Group

PCE: Path Computation Element, an entity that can compute path based on a network graph and applying computational constraints

PCC: Path Computation Client, any application that request path computation to be performed by a PCE

PCECP: PCE Communication Protocol, a protocol for communication between PCCs and PCEs

3. Introduction

IGP hierarchy consists of separating an IGP domain into multiple IGP areas, and limiting the topology visibility local to an area. This mechanism significantly improves IGP scalability.

[RFC4105] lists a set of motivations and requirements for setting up TE-LSPs across IGP area boundaries. These LSPs are called inter-area TE-LSPs. These requirements include the computation of inter-area shortest constrained paths with key guidelines being to respect the IGP hierarchy concept, and particularly the containment of topology information. The main challenge with inter-area MPLS-TE relies actually on path computation. The head-end LSR cannot compute an end-to-end shortest constrained path, as its topology visibility is limited to one area. Path computation can rely on loose hops with ERO expansion on each ABR, but this faces two issues: (1) it does not guarantee the computation of a shortest path that satisfies the TE-LSP constraints, and (2) it may result in several signalling crankback messages before it successfully sets up the path.

The Path Computation Element (PCE) Architecture, defined in [PCE-ARCH] can provide a suitable framework for computing an inter-area shortest path for a TE-LSP.
[PCE-ARCH] defines PCEs as entities that can compute paths based on a network graph and applying computational constraints. A PCE function can be located on a LSR or a network server. It defines a Path Computation Client (PCC) as an application requesting a path computation to be performed by a PCE. Typically a PCC can be a head-end LSR, a transit LSR requesting a TE-LSP path computation, or a PCE requesting a path computation of another PCE, in a collaborative mode.

One of the key applications of the PCE architecture is inter-domain path computation, where head-end LSRs have a limited topology view beyond its own domain. This includes both inter-area and inter-AS path computation.
Inter-area path computation requirements expressed in [RFC4105] may be achieved using the services of one or more PCEs.
PCE-based inter-area path computation could rely for instance on a single multi-area PCE that has the TE database of all the areas in the IGP domain and can directly compute an end-to-end shortest constrained path.
Alternatively, PCE-based inter-area path computation could rely on the cooperation between PCEs whereby each PCE covers one or more IGP areas and the full set of PCEs covers all areas.

The generic requirements for a PCE Communication Protocol (PCECP), allowing a PCC to send path computation requests to a PCE and the PCE to sent path computation response to a PCC are listed in [PCE-COM-REQ]. The use of a PCE-based approach, for inter-area path computation implies specific requirements on a PCE Communication

Protocol, in addition to the generic requirements already listed in [PCE-COM-REQ].
This document complements these generic requirements by listing a detailed set of PCECP requirements specific to the PCE-based computation of inter-area TE-LSPs.

The problem statement is discussed in section 4. Various PCE-based modes for inter-area path computation are described in section 5. Considerations for PCE location are provided in section 6. Finally detailed requirements are listed in section 7.

It is expected that a solution for a PCE Communication Protocol (PCECP) satisfies these requirements.

Note that PCE-based inter-area path computation may require a mechanism for an automatic PCE discovery across areas, which is out of the scope of this document. Detailed requirements for such mechanism are discussed in [PCE-DISCO-REQ].

4. Problem Statement

In intra-area MPLS-TE, a head-end LSR has complete topology visibility of the area and can compute an end-to-end shortest constrained path. IGP hierarchy allows improving IGP scalability, particularly in large networks with hundreds of nodes and thousands of links, by dividing the IGP domain into areas and limiting the flooding scope of topology information to area boundaries. A router in an area has full topology information for its own area but only reachability to routes in other areas._ Thus, a head-end LSR cannot compute an end-to-end constrained path that traverses more than one IGP area.

A solution for computing inter-area TE-LSP path relies on a per domain path computation ([PD-COMP]). It is based on loose hop routing with an ERO expansion on each ABR. This can allow setting up a constrained path, but faces two major limitations:

- -This does not allow computing an optimal constrained path
- -This may lead to several signalling crankback messages and hence delay the LSP setup, and invoke routing activities.

Note that, here, by optimal constrained path we mean the shortest constrained path across multiple areas, taking into account either the IGP or TE metric [METRIC]. In other words, such a path is the path that would have been computed by making use of some CSPF algorithm in the absence of multiple IGP areas.

The PCE architecture is well suited to inter-area path computation, as it allows overcoming the path computation limitations resulting from the limited topology visibility, by introducing path computation entities with more topology visibility, or by allowing cooperation between path computation entities in each area.

Several PCE-based path computation approaches can be used to compute inter-area optimal constrained paths, they are discussed in next section.

The use of a PCE-based approach, to perform inter-area path computation requires specific functions in a PCECP, in addition to the generic requirements listed in [PCE-COM-REQ]. Detailed requirements are discussed in section 7.

5. Various approaches for PCE-based inter-area path computation

There are various possible modes for PCE-Based inter-area path computation.
The computation of an inter-area optimal path could be done by:

- a single PCE, that has enough topology visibility and can alone compute an end-to-end optimal path,
- multiple PCEs, that have partial topology visibility and collaborate with each other so as to arrive at an end-to-end optimal path.

These two modes are referred as to "Single PCE computation" and "Multiple PCE computation with inter-PCE communication" in [PCE-ARCH]. Note that these two modes may co-exist in a given multi-area network.

Note that the per-area path computation mode relying on route expansion performed directly by ABRs on the path (which function has composite PCEs) , or on external PCEs contacted by the ABRs on the path, consists in fact of a simple concatenation of intra area paths. It actually only implies intra-area path computations and does not allow computing inter-area optimal paths. Hence this mode is not discussed in this document.

5.1. Single PCE Computation

In this mode the inter-area path computation is directly performed by a single PCE that has enough topology information to compute an end-to-end optimal path.
No inter-PCE communication is required in this mode.

This mode requires that the PCE have at least the TED of all the crossed areas for a given LSP. The actual distribution of PCEs may vary, i.e., a PCE may have TE database base from two, three or more IGP areas. If the head-end and tail-end LSRs are located in two peripheral areas, the PCE must have the TED of the source, backbone, and destination areas. In the particular case where the head-end/tail-End LSR is located in the backbone area and the tail-end/head-end LSR is located in a peripheral area, the PCE only needs the TED of the backbone area and the peripheral area to compute the path.

Figure 1 below illustrates an example of single PCE inter-area computation.

```
                    ------
                   | PCE0 |
                /   ------   \
               /      |       \
              /       |        \
             /        |         \
            /         |          \
  ------------------------------------------
 |            |               |             |
 |           ABR2            ABR3           |
R1    area1   |    area0      |    area2    R2
 |           ABR1            ABR4           |
 |            |               |             |
  ------------------------------------------
```

Figure 1: Example of single PCE computation.

In this multi area network PCE0 has topology visibility in area1, area0 and area2 and can compute and end-to-end path from area 1 to area 2. To setup an inter-area LSP from R1 in area 1 to R2 in area 2, R1 has to directly contact PCE0.

Note that this mode may rely on PCEs that have knowledge of topology in all areas. Such a PCE is called an "all-areas" PCE.
Particular attention should be given on the potential limitations of this "all-areas" PCE approach, in terms of scalability. Such all-area PCEs may have to maintain a large topology and this raises scalability issues both in terms of memory to maintain the TED and processing to synchronize TED information.
Also such all-area PCEs would potentially serve a large number of PCCs, and hence may face a huge path computation request overload during a network event such as link or node failure (that may impact a large number of TE-LSPs on a large number of head-end LSRs). This may significantly delay the TE-LSP recovery, and thus may diminish the benefits of such an approach.

5.2. Multiple PCE path computation with inter-PCE communication

In this mode the computation of an optimal inter-area TE-LSP path is distributed across multiple PCEs.
There is at least one PCE per area, and those PCEs do not have enough topology visibility to compute and inter-area optimal path.
PCEs in each area compute path segments in their respective areas and collaborate together to arrive at an end-to-end inter-area optimal path. Such collaboration is ensured thanks to inter-PCE communication.
The actual distribution of PCEs may vary, i.e. a PCE may have TE database from one, two, or more IGP areas, and the important thing is that the collection of topology and TE information maintained by a set of PCEs collectively must cover all the IGP areas where all inter-area LSPs traverse.

Figure 2 and 3 below illustrate two examples of multiple PCE inter-area computation

```
     ------          ------        ------
    | PCE1 |<------>| PCE0 |<---->| PCE2 |
     ------          ------        ------
        |               |              |
        |               |              |
  --------------------------------------------
 |           |                 |             |
 |          ABR2              ABR3           |
R1    area1  |    area0        |    area2    R2
 |          ABR1              ABR4           |
 |           |                 |             |
  --------------------------------------------
```

Figure 2: Cooperation between single-area PCEs

Figure 2 above illustrates a multi-area network with 3 areas. PCE0, PCE1 and PCE2 are PCEs responsible for path computation respectively in area 0, 1 and 2. These PCEs have topology visibility limited to one area and are called single-area PCEs.
To setup an inter-area LSP from R1 in area 1 to R2 in area 2, R1 has to contact PCE1. PCE1 then collaborates with PCE0, and PCE0 with PCE2 so as to compute an end-to-end shortest constrained path.

```
             ------              ------
            | PCE1 |   <---->  | PCE2 |
             ------              ------
            /      \             /    \
           /        \           /      \
     -------------------------------------------
    |            |                 |            |
    |           ABR2              ABR3          |
   R1    area1   |    area0        |    area2    R2
    |           ABR1              ABR4          |
    |            |                 |            |
     -------------------------------------------
```

Figure 3: Cooperation between multi-area PCEs

Figure 3 above illustrates a multi-area network with 3 areas. PCE1, and PCE2 are PCEs responsible for path computation respectively in area 0+1 and in area 0+2. This means that PCE1 and PCE2 have topology visibility in area0+area1 and area0+area2 respectively.
To setup an inter-area LSP from R1 in area 1 to R2 in area 2, R1 has to contact PCE1. PCE1 then collaborates with PCE2, so as to compute an end-to-end shortest constrained path.

6. Considerations on PCE location

As explained in [PCE-ARCH] a PCE can be a LSR or a network server.

But note that in the inter-area context, it may be quite efficient for the ABRs to act as PCEs. Indeed, ABRs have topology information of the backbone area and at least one peripheral area. An inter-area TE-LSP optimal path computation could rely on a single ABR, if the path crosses only two IGP areas, or on collaboration between two ABRs in case the path crosses three IGP areas.
For instance, in figure 2 above, ABR1 or ABR2 can play PCE1 role, and similarly ABR3 or ABR4 can play PCE2 role. Note that such ABRs are not necessarily transit LSRs on the computed inter-area TE LSP.

With such PCE distribution on ABRs, the PCECP would run directly between LSRs. Note that if N peripheral areas are connected to one backbone area, with at least N ABRs, inter-area path computation would potentially require a full mesh of N^2 PCE-PCE communications between ABRs. This reinforces the requirement for communication protocol overhead minimization, expressed in [PCE-COM-REQ].

7. Detailed Requirements on PCECP

This section lists a set of additional requirements for the PCE Communication Protocol that complement requirements listed in [PCE-COM-REQ] and are specific to inter-area (G)MPLS TE path computation.

7.1. Supported modes for PCE-based inter-area path computation

The PCECP MUST support the two PCE based inter-area path computation modes set forth in section 5.

Multiple PCE inter-area path computation requires cooperation between PCEs. Hence the PCECP MUST support cooperation between PCEs so as to arrive at an inter-area optimal path. It MUST allow requests and replies for cooperative inter-area path computation.

A simple cooperation may consists in exchanging intra or inter-area path Segments, and combine them to build an end-to-end optimal path. This is a basic cooperation level that allows building an inter-area optimal path in a recursive manner.
The path segment combination could be done in the backward direction, in which case an inter-PCE response message includes a set of computed intra or inter-area path segments from a set of downstream ABRs to the destination, along with their respective cost. These path segments have to be completed by upstream PCEs in a recursive manner so as to build an end-to-end optimal path across areas. To support this collaboration mode, a response message MUST allow the inclusion of multiple intra-area or inter-area path segments from a set of downstream ABRs, to the destination, along with their respective cost (see also 8.4).

Note that path segment combination in the forward direction is for further Study.

7.2. Control of area crossing

In addition to the path constraints specified in section 6.1.16 of [PCE-COM-REQ] the request message MUST allow indicating whether area crossing is allowed or not.
Indeed, for inter-area TE LSPs whose head-end and tail-end LSRs reside in the same IGP area, there may be intra-area and inter-area feasible paths, and, as set forth in [RFC4105], if the shortest path is an inter-area path, an operator either may want to avoid, as far as possible, crossing area and thus may prefer selecting a sub-optimal intra-area path or, conversely, may prefer to use a shortest path, even if it crosses areas.

7.3. Objective functions

Editorial note: This section will be moved to the generic requirement draft [PCE-COM-REQ] as this requirement applies to various PCE applications

As specified in [PCE-COM-REQ] an objective function corresponds to the optimization criteria used for the computation of one path, or the synchronized computation of a set of paths. In case of unsynchronized path computation, this can be, for example, the path cost or the residual bandwidth on the most loaded path link. In case of synchronized path computation, this can be, for example, the global bandwidth consumption or the residual bandwidth on the most loaded network link.

For the purpose of inter-area path computation the PCECP MUST support the following "unsynchronized" objective functions:

- -Minimum cost path (shortest path)
- -Least loaded path (widest path)
- -To be completed

Also the PCECP SHOULD support the following "synchronized" objective functions:

- -Minimize aggregate bandwidth consumption on all links
- -Maximize the residual bandwidth on the most loaded link.
- -Minimize the cumulative cost of a set of diverse paths.

Note that the absence of an objective function in this list does not mean that it must not be supported. As per the extensibility requirement expressed in [PCE-COM-REQ], note that new objective functions can be added to this list without impacting the protocol.

7.4. TE metric / IGP metric

The shortest path selection may rely either on the TE metric or on the IGP metric (see [METRIC]). Hence the PCECP request message MUST allow indicating the metric type (IGP or TE) to be used for shortest path selection.

7.5. Recording path attributes

There are at least three aggregate path attributes defined in (G)MPLS-TE: the hop-count, the cumulated TE-metric, and the cumulated IGP-metric. The operator can actually give any semantic to the TE metric and IGP metric. As suggested in [METRIC], if the TE-metric encodes the link cost and the IGP metric the link delay, the cumulated TE-metric indicates the total cost of the LSP and the cumulated IGP metric the end-to-end propagation delay (provided that the LSR transit delay is neglected in a first approximation).

A PCC may need to know the aggregate path attributes of an LSR, for instance to select a preferred path among a set of computed paths.

In an inter-area context, a PCC may not be able to deduce this information from the supplied path.
Therefore the PCECP request message MUST allow indicating the set of aggregate path attributes (hop-count, cumulated TE-metric, cumulated IGP-Metric) that are required in the reply and the PCECP response message MUST support the inclusion of a set of aggregate path attributes.

Note that if new TE link attributes are defined in the future to encode specific link parameters, and allowing to define specific aggregate path constraints, such as, e.g. delay, distance or power loss, the PCEPC will have to be extended to support them.

Note that in case the computed path includes loose hops the PCE may not be able to give an accurate aggregate path attribute. Hence the response message MUST allow indicating that an aggregate path attribute is unknown.

7.6. Strict Explicit path and Loose Path

A Strict Explicit Path is defined as a set of strict hops.
A Loose Path is defined as a set of strict and loose hops.
An inter-area path may be strictly explicit or loose (e.g. a list of ABRs as loose hops)
It may be useful to indicate to the PCE if a Strict Explicit path is required or not.
Hence the PCECP request message MUST allow indicating if a Strict Explicit Path is required/desired.

7.7. PCE-list enforcement and recording in Multiple PCE Computation

In case of multiple-PCE path computation, a PCC may want to indicate a preferred list of PCEs to be used.
Hence the PCECP request message MUST support the inclusion of a list of preferred PCEs.
Note that this requires that a PCC in one area have knowledge of PCEs in other areas. This could rely on configuration or on a PCE discovery mechanism, allowing discovery across area boundaries (see [PCE-DISCO-REQ]).

Also it would be useful to know the list of PCEs which effectively participated in the computation.
Hence the request message MUST support requesting for PCE recording and the response message MUST support the recording of the set of one or more PCEs that took part into the computation.
It may also be useful to know the path segments computed by each PCE.
Hence the request message SHOULD allow requesting for the identification of path segment computed by a PCE, and the response

message SHOULD allow identifying the path segments computed by each PCE.

7.8. Inclusion of Area IDs in request

The knowledge of the area in which the source and destination lie would allow selection of appropriate cooperating PCEs.
A PCE may not be able to determine the location of the source and destination LSRs. Hence it would be useful that a PCC indicates the source area ID and destination area IDs.

For that purpose the request message MUST support the inclusion of source and destination area IDs.
Note that this information could be learned on the PCC by configuration.

7.9. Load-Balancing

Editorial note: This section will be moved to the generic requirement draft [PCE-COM-REQ] as this requirement applies to various PCE applications

In some cases a single inter-area path may not fit a TE-LSP bandwidth constraint. In this case it may be useful to setup a set of paths whose cumulated residual bandwidth fit the TE-LSP bandwidth request. This is what we call load balancing.
So as to avoid ending up with a huge number of paths for a given request, and/or with low bandwidth paths, it is required to control the number of computed paths and the minimum path bandwidth.

The request message MUST allow indicating if load-balancing is allowed or not. It MUST also include the number of paths in a load-balancing path group, and the minimum path bandwidth in a load-balancing path group. The response MUST support the inclusion of the set of computed paths of a load-balancing path group, as well as their respective bandwidth.

7.10. Diverse Path computation

For various reasons including protection and load balancing, the computation of diverse inter-area paths may be required.
There are various levels of diversity in an inter-area context:

- -Per area diversity (intra-area path segments are link, node or SRLG disjoint)
- -Inter-Area diversity (end-to-end inter-area paths are link, node or SRLG disjoint)

Note that two paths may be disjoint in the backbone area but shared in peripheral areas. Also two paths may be node disjoints within areas but may share ABRs.

The request message MUST allow requesting the computation of a set of diverse paths between a same couple of nodes or distinct couples of nodes. It MUST allow indicating the required level of intra-area diversity (link, node, SRLG) on a per area basis, as well as the level of inter-area diversity (shared ABRs or ABR disjointness).

The response message MUST allow indicating the level of diversity of a set of computed loose paths.

Note that specific objective function may be requested for diverse path computation, such as to minimize the cumulated cost of a set of diverse paths (see also 7.3).

7.11. LSP failure handling

7.11.1. LSP Rerouting

Editorial note: This section will be moved to the generic requirement draft [PCE-COM-REQ] as this requirement applies to various PCE applications

Upon LSP failure, due to link, node or SRLG failure, a head-end LSR may send a request to the PCE so as to reroute the LSP over an alternate path. So as to ease the computation such request should include the previous path and the failed element (if it can be identified).

Hence the request message MUST allow indicating if the computation is for an LSP restoration, and MUST support the inclusion of the previously computed path as well as the failed element.
Note that the old path is actually useful only if the old LSP is not torn down yet. This is up to the PCC to decide if it includes the old path or not.

Note that a network failure may impact a large number of LSPs. A potentially large number of PCCs, are going to simultaneously send a request to the PCE. Some jittering may be used on PCCs so as to delay a request to the PCE, under network failure condition.

The PCECP MAY support the inclusion, in a response message to a PCC, of an upper bound of the jitter to be used for further requests to the PCE (e.g. the PCC will wait for a random value between 0 and the upper bound before sending another request). This upper bound would depend on the level of congestion of the PCE.

7.11.2. Backup path computation

ABRs can be protected using Fast Reroute (FRR) node protection [MPLS-FRR]. This requires setting up inter-area FRR backup LSPs (bypass or detour).
The PCECP SHOULD support the computation of inter-area FRR backup LSPs (detour or bypass). Note that the objective function may be to

minimize overhaul backup bandwidth consumption, by maximizing bandwidth sharing among backup LSPs protecting independent elements.

Detailed requirements for intra and inter-area PCE-based backup path computation are for further study and will be addressed in a separate document.

7.12. Inter-Area policies

As already defined in Section 8.2 a request message MUST allow indicating whether area crossing is allowed or not.

A PCE may want to apply policies based on the initiating PCC. In a multiple-PCE computation the address of the initiating PCC may no longer be part of the request messages sent between PCEs. Hence, the request message MUST support the inclusion of the address of the originator PCC.

Note that in some case this is important to contain an inter-area path within a single AS. Hence the request message MUST allow indicating that AS crossing is not authorized.

7.13. Scalability

As already pointed out in [PCE-COM-REQ] the PCECP MUST scale well, at least as good as linearly, with an increase of any of the following parameters:

- number of PCCs communicating with a single PCE
- number of PCEs communicated to by a single PCC
- number of PCEs communicated to by another PCE
- number of request per PCE per second in steady state
- number of requests per PCE per second under emergency condition

Note that these numbers will depend on the level of PCE distribution and on the PCE approach used (Single PCE computation, Multiple PCEs computation)

For instance in a network that comprises I IGP areas, with P PCCs per area and A ABRs per area boundary then

- For single PCE computation with an all-areas PCE Server:
  - Number of PCCs communicating with a single PCE=I*P
  - Number of PCEs communicated to by a single PCC=1
  - Number of PCEs communicated to by another PCE=0

- For multiple PCE computation with ABRs acting as PCEs:
  - Number of PCCs communicating with a single PCE=P
  - Number of PCE communicated to by a single PCC=I*A
  - Number of PCEs communicated to by another PCE=I*A

Typical values for a large inter-area network can be: I=50, P=100, and A=2.

Note also that the memory and CPU consumed to maintain and synchronize the TED on a PCE will directly depend on the number of areas under control of the PCE. This may diminish the benefits of "all area" PCEs, but this is beyond the scope of this document.

8. Manageability consideration

Manageability of inter-area PCEs must address the following consideration for section 7:

- need for a MIB module for control plane and monitoring
- need for built-in diagnostic tools
- configuration implications for the protocol

9. Security Considerations

IGP areas are administrated by the same entity. Hence the inter-area application does not imply new trust model, or new security issues beyond those already defined in [PCE-COM-REQ].

10. Acknowledgments

We would also like to thank Adrian Farrel, Jean-Philippe Vasseur, Bruno Decraene and Yannick Le Louedec for their useful comments and suggestions.

11. Informative References

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

[RFC3667] Bradner, S., "IETF Rights in Contributions", BCP 78, RFC 3667, February 2004.

[BCP79] Bradner, S., "Intellectual Property Rights in IETF Technology", RFC 3979, March 2005.

[RFC4105] Le Roux J.L., Vasseur J.P., Boyle, J., et al. "Requirements for inter-area MPLS-TE" RFC 4105, June 2005.

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

[PD-COMP] Vasseur, J.P., Ayyangar, A., Zhang, R., " A Per-domain path computation method for computing Inter-domain Traffic Engineering (TE) Label Switched Path (LSP)", draft-ietf-ccamp-inter-domain-pd-path-comp, work in progress

[METRIC] Le Faucheur, F., Uppili, R., Vedrenne, A., Merckx, P., and T. Telkamp, "Use of Interior Gateway Protocol(IGP) Metric as a second MPLS Traffic Engineering (TE) Metric", BCP 87, RFC 3785, May 2004.

[ID-RSVP] Ayyangar, A., Vasseur, J.P., "Inter domain GMPLS Traffic Engineering - RSVP-TE extensions", draft-ietf-ccamp-inter-domain-rsvp-te, work in progress.

12. Editor Address:

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: jeanlouis.leroux@francetelecom.com

13. Contributors' Addresses

Jerry Ash
AT&T
Room MT D5-2A01
200 Laurel Avenue
Middletown, NJ 07748, USA
Phone: +1-(732)-420-4578
Email: gash@att.com

Nabil Bitar
Verizon
40 Sylvan Road
Waltham, MA 02145
Email: nabil.bitar@verizon.com

Dean Cheng
Cisco Systems Inc.
3700 Cisco Way
San Jose CA 95134 USA
Phone: +1 408 527 0677
Email: dcheng@cisco.com

Kenji Kumaki
KDDI Corporation

Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460, JAPAN

Phone: +81-3-6678-3103
Email: ke-kumaki@kddi.com

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

Raymond Zhang
BT INFONET Services Corporation
2160 E. Grand Ave.
El Segundo, CA 90245 USA
Email: Raymond_zhang@bt.infonet.com

14. Intellectual Property Statement

Copyright Statement

R¥ 8 EÀ!÷ù¨7 óùJO®·f T¹¶0-draft-leroux-pce-pcecp-multiarea-reqs-00.txt@ 0 ³-
G ÐÅ@0Ðs;G ÐÅ77.txt7 7-draft-leroux-pce-pcecp-multiarea-reqs-00.txt
7ÿÿÿÿ7ú @û @Ý£WE³@ü @Ý£WE³ý
þ
ÿ µL