

PCE Working Group
Internet Draft

J.L. Le Roux (France Telecom)
Paul Mabey (Qwest)
Raymond Zhang (BT Infonet)
Eiji Oki (NTT)
Ting Wo Chung (Bell Canada)
Richard Rabbat (Fujitsu)

Category: Informational
Expires: December 2005

June 2005

Requirements for Path Computation Element (PCE) Discovery

[draft-leroux-pce-discovery-reqs-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). 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>.

Abstract

This document presents a set of requirements for a Path Computation Element (PCE) discovery mechanism that would allow a Path Computation Client (PCC) to discover dynamically and automatically a set of PCEs along with their capabilities. It is intended that solutions that specify procedures and protocol extensions for such PCE discovery satisfy these requirements.

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.	Terminology.....	2
2.	Introduction.....	3
3.	Problem Statement and Requirements overview.....	4
3.1.	Problem Statement.....	4
3.2.	Requirements overview.....	5
4.	Example of application scenario.....	6
5.	Detailed Requirements.....	7
5.1.	PCE Information to be disclosed.....	7
5.1.1.	Discovery of PCE Location.....	7
5.1.2.	Discovery of PCE computation scopes and domain(s) under control.....	7
5.1.3.	Discovery of PCE Capabilities.....	8
5.1.4.	Discovery of Alternate PCEs.....	9
5.2.	Scope of PCE Discovery.....	9
5.3.	PCE Information Synchronization.....	9
5.4.	Detecting PCE Liveliness.....	10
5.5.	Discovery of PCE capacity and congestion.....	10
5.6.	Extensibility.....	10
5.7.	Scalability.....	10
5.8.	PCE Selection.....	10
6.	Security Considerations.....	11
7.	Acknowledgments.....	11
8.	References.....	11
9.	Authors' Addresses:.....	12
10.	Intellectual Property Statement.....	13

[1.](#) Terminology

Terminology used in this document

LSR: Label Switch Router

TE-LSP: Traffic Engineered Label Switched Path

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.

PCC: Path Computation Client: any client application requesting a

path computation to be performed by a Path Computation Element.

IGP Area: OSPF Area or ISIS level/area

Le Roux et al.

[Page 2]

ABR: IGP Area Border Router (OSPF ABR or ISIS L1L2 router)

Intra-area TE LSP: A TE LSP whose path does not cross IGP area boundaries.

Inter-area TE LSP: A TE LSP whose path transits through two or more IGP areas.

Inter-AS MPLS TE LSP: A TE LSP whose path transits through two or more ASes or sub-ASes (BGP confederations).

Domain: any collection of network elements within a common sphere of address management or path computational responsibility. Examples of domains include IGP areas and Autonomous Systems.

2. Introduction

The PCE Architecture [[PCE-ARCH](#)] defines a Path Computation Element (PCE) as an entity capable of computing TE-LSPs paths satisfying a set of constraints. The PCE function can be located on a router/LSR (composite PCE) or on a network server (external PCE).

A PCE serves TE-LSP path computation requests sent by Path Computation Clients (PCC).

A Path Computation Client (PCC) is a client application requesting a path computation to be performed by a PCE. This can be, for instance, an LSR requesting a path for a TE-LSP for which it is the head-end, or a PCE requesting a path computation of another PCE (inter-PCE communication). The communication between a PCC and a PCE requires a client-server protocol whose requirements are listed in [[PCE-COM-REQ](#)].

There are several motivations for the adoption of a PCE-based architecture to perform a TE-LSP path computation. They are listed in [[PCE-ARCH](#)]. This includes applications such as CPU intensive path computation, inter-domain path computation and backup path computation.

The PCE architecture requires, of course, that a PCC be aware of the location and capabilities of one or more PCEs in its domain, and also potentially of some relevant PCEs in other domains (in the context of inter-domain path computation).

In that context it would be highly desirable to define a mechanism for automatic and dynamic PCE discovery, which would allow PCCs to automatically discover a set of PCEs along with their capabilities, and to dynamically detect new PCEs or any modification of the PCE capabilities. This includes the discovery by a PCC of a set of one or more PCEs in its domain, and potentially in some other domains. The

latter is a desirable function in the case of inter-domain path computation for example.

This document lists a set of functional requirements for such an automatic and dynamic PCE discovery mechanism. [Section 3](#) points out the problem statement. [Section 4](#) illustrates an application scenario. Finally [section 5](#) addresses detailed requirements.

It is intended that solutions that specify procedures and protocol extensions for such PCE discovery satisfy these requirements. There is no intent either to specify solution specific requirements or to make any assumption on the protocol(s) that could be used for the discovery.

Note that requirements listed in this document apply equally to MPLS-TE and GMPLS-capable PCEs.

It is also important to note that the notion of a PCC encompasses a PCE acting as PCC when requesting a TE-LSP path computation of another PCE (inter-PCE communication). Thus this document does not make the distinction between PCE discovery by PCCs and PCE discovery by PCEs.

[3. Problem Statement and Requirements overview](#)

[3.1. Problem Statement](#)

A routing domain may in practice be comprised of multiple PCEs:

- The path computation load may be balanced among a set of PCEs to improve scalability;
- For the purpose of redundancy, primary and backup PCEs may be used;
- PCEs may have distinct path computation capabilities (multi-constrained path computation, backup path computation...);
- In an inter-domain context there can be several PCEs with distinct path computation scopes (intra-area, inter-area, inter-AS, inter-layer), each PCE being responsible for path computation in one or more domains within its scope.

As an example, in a multi-area network made of one backbone area and N peripheral areas, and where inter-area MPLS-TE path computation relies on multiple-PCE path computation with ABRs acting as PCEs, the backbone area would comprise at least N PCEs. In existing multi-area networks, N can be quite large (e.g. beyond fifty).

In order to allow for efficient PCE selection by PCCs and efficient load balancing of requests, a PCC has to know the location of PCEs in its domain, along with their capabilities, and also potentially of some relevant PCEs in other domains (for inter-domain path computation purpose).

Such PCE information could be learnt through manual configuration, on each PCC, of the set of PCEs along with their capabilities and

scope(s). Such manual configuration approach may be sufficient, and even desired in some particular situations, but it obviously faces several limitations:

- This may imply a substantial configuration overhead (see the above example with N PCEs);
- This would not allow a PCC to dynamically detect that a new PCE is available, that an existing PCE is no longer available, or that there is a change in the PCE's capabilities.

Furthermore, as with any manual configuration approach, this may lead to undesirable configuration errors.

Hence, an automated PCE discovery mechanism allowing a PCC to dynamically discover a set of PCEs and their capabilities is highly desirable.

3.2. Requirements overview

A PCE discovery mechanism that satisfies the requirements set forth in this document MUST allow a PCC to automatically discover the location of one or more PCEs in its domain and also, potentially, of PCEs in other domains, of interest for inter-domain path computation purpose.

A PCE discovery mechanism MUST allow discovering the path computation scope(s) of a PCE. It MUST also allow a PCC to discover the set of one or more domains under the path computation responsibility of a PCE.

A PCE discovery mechanism MUST allow PCCs to dynamically discover that a new PCE has appeared or that there is a change in PCE information.

It MUST also allow PCCs to dynamically discover that a PCE is no longer available.

A PCE discovery mechanism SHOULD also allow PCCs to learn about a set of PCE path computation capabilities.

4. Example of application scenario




Figure 1.

Figure 1 above illustrates a network with several PCEs:

- The ABR R3 is a composite PCE that can take part in inter area path computation. It can compute paths in area 1 and area 0.
- The ABR R6 is a composite PCE that can take part in inter-area path computation. It can compute paths in area 0 and area2
- The ASBR R9 is a composite PCE that can take part in inter-AS path computation, responsible for path computation in AS1 towards AS2.
- The ASBR R12 is a composite PCE that can take part in inter-AS path computation, responsible for path computation in AS2 towards AS1.
- The server S1 is an external PCE that can be used to compute diverse paths and backup paths in area 1.

The PCE discovery mechanism will allow:

- each LSR in area 1 and 0 to dynamically discover R3, as a PCE for inter-area path computation as well as its path computation domains: area1 and area0.
- each LSR in area 0 and 2 to dynamically discover R6, as a PCE for inter-area path computation, as well as its path computation domains: area2 and area0.
- each LSR in AS1 and some PCEs in AS2 to dynamically discover R9 as a PCE for inter-AS path computation as well as its path computation domain: AS1
- each LSR in area 1 to dynamically discover S1, as a PCE for diverse path computation and backup path computation in area1.

5. Detailed Requirements

5.1. PCE Information to be disclosed

The PCE discovery mechanism MUST disclose some PCE information that will allow PCCs to select appropriate PCEs.

This section details the kind of information that has to be disclosed.

5.1.1. Discovery of PCE Location

The PCE discovery mechanism MUST allow discovering, for a given PCE, the IPv4 and/or IPv6 address to be used to reach the PCE. This address will typically be a loop-back address that is always reachable, if there is any connectivity to the PCE.

This address will be used by PCCs to communicate with a PCE, thanks to a PCC-PCE communication protocol.

5.1.2. Discovery of PCE computation scopes and domain(s) under control

Inter-domain path computation is a key application of the PCE architecture. This can rely on a multiple-PCE path computation, where PCEs in each domain compute a part of the end-to-end path and collaborate with each other to find the end-to-end-path. This can also rely on a single-PCE path computation where a PCE has visibility inside multiple domains and can compute an inter-domain path.

Hence the PCE discovery mechanism MUST allow discovering the path computation scope of a PCE, i.e. if a PCE can be used to compute or to take part in the computation of intra-area, inter-area or inter-AS TE-LSP. Note that these path computation scopes are not mutually exclusive.

Also the PCE discovery mechanism MUST allow discovering the set of one or more domains under the path computation responsibility of a PCE, i.e. where a PCE has visibility and can compute TE-LSP paths. These domains can be identified using a domain identifier: For instance, an IGP area can be identified by the Area ID (OSPF or ISIS), and an AS can be identified by the AS number.

5.1.3. Discovery of PCE Capabilities

In the case where there are several PCEs with distinct capabilities, available, a PCC has to select one or more appropriate PCEs.

For that purpose the PCE discovery mechanism SHOULD allow the discovery of some PCE capabilities.

For the sake of illustration this could include for instance some path computation related capabilities:

- The capability to compute MPLS-TE and/or GMPLS paths;
- The type of link and path constraints supported: e.g. bandwidth, affinities, delay;
- The objective functions supported: e.g. shortest constrained path, shortest bounded delay path;
- The capability to compute multiple paths in a synchronized manner: e.g. diverse path computation, load balancing computation;
- Some GMPLS specific capabilities: e.g. the supported interface switching capabilities, the capability to compute multi-layer paths.

And this could also include some general capabilities:

- The capability to handle request prioritization;
- The capability to authenticate PCCs and to be authenticated.

Such information regarding PCE capabilities could then be used by a PCC to select an appropriate PCE from a list of candidate PCEs.

Note that the description of general and path computation specific PCE capabilities is out of the scope of this document. It is expected that this will be described in a separate document.

It is paramount that dynamic capability discovery MUST NOT generate an excessive amount of information and SHOULD be limited to a small set of generic capabilities.

If required, the exhaustive discovery of detailed capabilities could be ensured by means of the PCC-PCE communication protocol.

Actually a tradeoff should be found between capability discovery by the PCE discovery mechanism and by the PCC-PCE communication protocol. One of the objectives of the PCE discovery mechanism is to help PCCs to select appropriate PCEs and limit the likelihood of PCC-PCE communication rejections that may occur in case a PCE cannot support a given capability.

5.1.4. Discovery of Alternate PCEs

In the case of a PCE failure, a PCC has to select another PCE, if one is available. It could be useful in various situations, to indicate a set of one or more alternate PCEs that can be selected in case a given PCE fails.

Hence the PCE Discovery mechanism SHOULD allow the advertising, for a given PCE of the location of one or more assigned alternate PCEs.

5.2. Scope of PCE Discovery

The PCE Discovery mechanism MUST allow the control of the scope of the PCE information discovery (IGP Area, AS, set of AS) on a per PCE basis. In other words it MUST allow to control to which PCC or group of PCCs the information related to a PCE may be disclosed.

The choice for the discovery scope of a given PCE MUST include the followings:

- All PCCs in a single IGP area
- All PCCs in a set of adjacent IGP areas
- All PCCs in a single AS
- All PCCs in a set of ASes
- A set of one or more PCCs in a set of one or more ASes

Particularly this also implies that the PCE Discovery mechanism MUST allow for the discovery of PCE information across IGP areas and across AS boundaries.

Note that it MUST be possible to deactivate PCE discovery on a per PCE basis.

5.3. PCE Information Synchronization

The PCE discovery mechanism MUST allow a PCC to detect any change in the information related to a PCE (e.g. capability modifications).

In addition it MUST be possible to dynamically detect new PCEs.

The PCE Discovery Mechanism SHOULD allow such detection under 60 seconds.

Note that PCE capabilities are expected to be fairly stable and not to change frequently.

5.4. Detecting PCE Liveliness

The PCE discovery mechanism MUST allow a PCC to detect when a PCE is no longer alive. This allows a PCC to rapidly switch to another PCE (for instance a predefined alternate PCE), and thus minimizes path computation service disruption.

The PCE discovery mechanism SHOULD allow such PCE liveliness detection under 60 seconds.

5.5. Discovery of PCE capacity and congestion

PCE WG feedback is requested on the following items:

- Is there a need for the discovery of PCE capacity in terms of computation power? This static parameter could be used to ensure weighted load balancing of requests in case PCEs do not have the same capacity.
- Would it be useful that a PCE report its status as "congested" in case it is too busy? PCCs may then use this dynamic information to prefer a different PCE.

5.6. Extensibility

The PCE discovery mechanism MUST be flexible and extensible so as to easily allow for the addition of some additional PCE information that could be defined in the future.

5.7. Scalability

The PCE discovery mechanism MUST be designed to scale well with an increase of any of the following parameters:

- Number of PCCs;
- Number of PCEs;
- Number of IGP Areas in the discovery scope;
- Number of ASs in the discovery scope.

Particularly, in case routing protocols (IGP, BGP) are extended to support PCE discovery, the extensions MUST NOT cause a degradation in routing protocol performance. The same applies to a signaling solution that could serve for this communication.

5.8. PCE Selection

A PCC may have to select among a set of candidate PCEs that have the same capabilities. A specific PCE selection procedure SHOULD be defined in order to ensure consistent behaviour when doing load balancing and avoid that all PCCs send the requests to only one PCE. The precise requirements and mechanisms for this function are out of the scope of this document. It is expected that this requirement will

be covered in another document.

Le Roux et al.

[Page 10]

6. Security Considerations

PCE discovery and particularly inter-AS PCE discovery may raise new security issues. PCE discovery procedures or protocol extensions MUST deliver the operational security objectives where required. The overall security objectives of confidentiality, integrity and availability may take on varying level of importance. These objectives MAY be met by other established means and protocols.

The PCE discovery mechanism MUST be able to restrict the scope of discovery to a set of authorized PCCs. The identity of any PCE MUST only be learnt by authorized PCCs. It MUST be possible to encrypt discovery information.

Note that a threat analysis of the proposed procedures and/or protocol extensions SHOULD address masquerade, eavesdropping, unauthorized access, loss or corruption of information (includes replay attacks), repudiation, forgery and denial of service attacks.

7. Acknowledgments

We would like to thank Benoit Fondeviolle, Thomas Morin, Emile Stephan, Jean-Philippe Vasseur, Dean Cheng and Mohamed Boucadair for their useful comments and suggestions.

8. 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.

[RFC3668] Bradner, S., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3668](#), February 2004.

[PCE-ARCH] Farrel, A., Vasseur, J.P., Ash, J., "Path Computation Element (PCE) Architecture", [draft-ietf-pce-architecture-00.txt](#), work in progress.

[PCE-COM-REQ] Ash, J., Le Roux, J.L., " PCE Communication Protocol Generic Requirements", [draft-ietf-pce-comm-protocol-gen-reqs-00.txt](#), work in progress.

9. Authors' Addresses:

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

Paul Mabey
Qwest Communications
950 17th Street,
Denver, CO 80202,
USA
Email: pmabey@qwest.com

Raymond Zhang
Infonet Services Corporation
2160 E. Grand Ave.
El Segundo, CA 90025
USA
Email: raymond_zhang@infonet.com

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

Ting Wo Chung
Bell Canada
181 Bay Street, Suite 350
Toronto, Ontario, M5J 2T3
CANADA,
Email: ting_wo.chung@bell.ca

Richard Rabbat
Fujitsu Laboratories of America
1240 East Arques Ave, MS 345
Sunnyvale, CA 94085
USA
Email: richard@us.fujitsu.com

10. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

