         Requirements for Path Computation Element (PCE) Discovery

                    draft-ietf-pce-discovery-reqs-01.txt

Status of this Memo

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 certain information relevant for PCE selection. It is
   intended that solutions that specify procedures and protocol(s) or
   protocol(s) 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. Contributors

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

    Jean-Louis Le Roux (France Telecom)
    Paul Mabey (Qwest Communications)
    Eiji Oki (NTT)
    Richard Rabbat (Fujitsu)
    Ting Wo Chung (Bell Canada)
    Raymond Zhang (BT Infonet)

2. 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

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

   AS: Autonomous System

   ASBR: AS Border 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.

3. 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. A PCE serves path computation requests sent by
Path Computation Clients (PCC).
A 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 generic requirements are listed in [PCE-COM-REQ].

There are several motivations for the adoption of a PCE-based
architecture to perform a 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 of one or more PCEs in its domain, and also potentially of
some PCEs in other domains, for 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, including information required
for PCE selection, and to dynamically detect new PCEs or any
modification of PCE's information. 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(s)
or protocol(s) 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 path computation of another PCE
(inter-PCE communication). Hence, this document does not make the
distinction between PCE discovery by PCCs and PCE discovery by PCEs.

4. Problem Statement and Requirements overview

4.1. 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 some information relevant for PCE selection,
   and also potentially of some 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. 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 information.

   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 is highly desirable.

4.2. 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 (intra-area, inter-area, inter-ASà). 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's
   information. It MUST also allow PCCs to dynamically discover that a
   PCE is no longer available.

   The PCE discovery MUST be secure. In particular, key consideration

MUST be given in terms of how to establish a trust model for PCE
   discovery.

   OPTIONALLY a PCE discovery mechanism MAY be used so as to disclose a
   set of PCE capabilities.

[5](5). Example of application scenario

```
   <---------------AS1-------------------->        <----AS2---
    Area 1            Area 0        Area 2
   R1---------R3-----R5-------R6-----------R9----------R11----R13
    |          |              |             |           |
    |          |              |             |           |
   R2---------R4-----R7-------R8-----------R10---------R12----R14


      S1
```

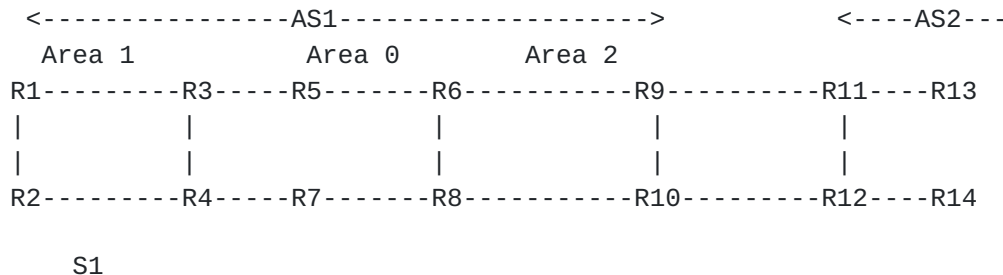   Figure 1

   Figure 1 above illustrates a network with several PCEs:
   -The ABR R3 is a 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 PCE that can take part in inter-area path
   computation. It can compute paths in area 0 and area2;
   -The ASBR R9 is a PCE that can take part in inter-AS path
   computation, responsible for path computation in AS1 towards AS2;
   -The ASBR R12 is a PCE that can take part in inter-AS path
   computation, responsible for path computation in AS2 towards AS1;
   -The server S1 is a 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 in AS1 towards AS2;
   -each LSR in area 1 to dynamically discover S1, as a PCE for diverse
   path computation and backup path computation in area1.

[6](6). Detailed Requirements

[6.1](6.1). PCE Information to be disclosed

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

   We distinguish two levels of information to be disclosed by the PCE
   discovery mechanism:
        -Mandatory information: This comprises discovery of PCE location,
         PCE computation scopes and domains under control
        -Optional information: This comprises discovery of PCE
         capabilities and alternate PCEs.

   Note that the latter information is optional in the context of the
   PCE discovery mechanism. It could also be obtained by other
   mechanisms, such as for instance the PCC-PCE communication protocol.

[6.1.1](6.1.1). 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.

[6.1.2](6.1.2). Discovery of PCE computation scope(s) 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 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.
It MUST also allow discovering the set of one or more domain(s)
towards which a PCE can compute paths.

6.1.3. 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 MAY be used so as to
   disclose 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 support for multi-layer
         path computation;
   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 discovery of PCE capabilities 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.

6.1.4. 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.

6.2. 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.

6.3. 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 information is expected to be fairly stable and not to
change frequently.

6.4. 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 detection under 60
seconds.

6.5. 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.
      -May the PCE discovery mechanism be used so 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.

6.6. Security Requirements

The PCE Discovery mechanism MUST address security issues across
multiple ASes.

Key consideration MUST be given in terms of how to establish a trust
model for PCE discovery. The PCE discovery mechanism MUST explicitly
support a specific set of one ore more trust model(s).

The PCE discovery mechanism MUST deliver the operational security
objectives where required. The overall security objectives of
privacy, authentication, and integrity may take on varying level of
importance. These objectives MAY be met by other established means
and protocols.

Particularly, mechanisms MUST be defined to ensure authentication,
integrity and privacy of PCE discovery information.

It MUST be possible to restrict the scope of discovery to a set of
authorized PCCs. In particular, the identity of any PCE MUST only be
learnt by authorized PCCs.

It MUST be possible for PCEs to authenticate PCCs and for PCCs to
authenticate PCEs. It MUST also be possible to encrypt discovery
information.

6.7. Extensibility

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

6.8. 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 discovery.

7. Security Considerations

   This document is a requirement document and hence does not raise by
   itself any particular security issue.

   A set of security requirements that MUST be addressed when
   considering the design and deployment of a PCE Discovery mechanism
   have been identified in section 6.6.

8. Acknowledgments

   We would like to thank Benoit Fondeviole, Thomas Morin, Emile
   Stephan, Jean-Philippe Vasseur, Dean Cheng, Adrian Farrel, Renhai
   Zhang, Mohamed Boucadair, Eric Gray, Igor Bryskin and Dimitri
   Papadimitriou, for their useful comments and suggestions.

9. 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, work in
   progress.

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

10. 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

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

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

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

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

11. 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.