Network Working Group   J. Ash, Ed.
Request for Comments: 4657   AT&T
Category: Informational   J.L. Le Roux, Ed.
                                                         France Telecom
                                                         September 2006

                Path Computation Element (PCE) Communication Protocol
                              Generic Requirements

Status of This Memo

   This memo provides information for the Internet community. It does
   not specify an Internet standard of any kind. Distribution of this
   memo is unlimited.

Copyright Notice

Abstract

   The PCE model is described in the "PCE Architecture" document and
   facilitates path computation requests from Path Computation Clients
   (PCCs) to Path Computation Elements (PCEs). This document specifies
   generic requirements for a communication protocol between PCCs and
   PCEs, and also between PCEs where cooperation between PCEs is
   desirable. Subsequent documents will specify application-specific
   requirements for the PCE communication protocol.

Table of Contents

1. Introduction

   A Path Computation Element (PCE) [RFC4655] supports requests for path
   computation issued by a Path Computation Client (PCC), which may be
   'composite' (co-located) or 'external' (remote) from a PCE. When the
   PCC is external from the PCE, a request/response communication
   protocol is required to carry the path computation request and return
   the response. In order for the PCC and PCE to communicate, the PCC
   must know the location of the PCE; PCE discovery is described in
   [PCE-DISC-REQ].

   The PCE operates on a network graph in order to compute paths based
   on the path computation request(s) issued by the PCC(s). The path
   computation request will include the source and destination of the
   paths to be computed and a set of constraints to be applied during
   the computation, and it may also include an objective function. The
   PCE response includes the computed paths or the reason for a failed
   computation.

This document lists a set of generic requirements for the PCE
Communication Protocol (PCECP). Application-specific requirements
are beyond the scope of this document, and will be addressed in
separate documents. For example, application-specific communication
protocol requirements are given in [PCECP-INTER-AREA] and
[PCECP-INTER-LAYER] for inter-area and inter-layer PCE applications,
respectively.

2. Conventions Used in This Document

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

3. Terminology

Domain: Any collection of network elements within a common sphere of
address management or path computational responsibility. Examples of
domains include Interior Gateway Protocol (IGP) areas, Autonomous
Systems (ASs), multiple ASs within a service provider network, or
multiple ASs across multiple service provider networks.

GMPLS: Generalized Multi-Protocol Label Switching

LSP: MPLS/GMPLS Label Switched Path

LSR: Label Switch Router

MPLS: Multi-Protocol Label Switching

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

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 (see
further description in [RFC4655]).

TED: Traffic Engineering Database, which contains the topology and
resource information of the network or network segment used by a PCE.

TE LSP: Traffic Engineering (G)MPLS Label Switched Path.

See [RFC4655] for further definitions of terms.

4. **Overview of PCE Communication Protocol (PCECP)**

   In the PCE model, path computation requests are issued by a PCC to a
   PCE that may be composite (co-located) or external (remote). If the
   PCC and PCE are not co-located, a request/response communication
   protocol is required to carry the request and return the response.
   If the PCC and PCE are co-located, a communication protocol is not
   required, but implementations may choose to utilize a protocol for
   exchanges between the components.

   In order for a PCC and PCE to communicate, the PCC must know the
   location of the PCE. This can be configured or discovered. The PCE
   discovery mechanism is out of scope of this document, but
   requirements are documented in [PCE-DISC-REQ].

   The PCE operates on a network graph built from the TED in order to
   compute paths. The mechanism by which the TED is populated is out of
   scope for the PCECP.

   A path computation request issued by the PCC includes a specification
   of the path(s) needed. The information supplied includes, at a
   minimum, the source and destination for the paths, but may also
   include a set of further requirements (known as constraints) as
   described in Section 5.

   The response from the PCE may be positive in which case it will
   include the paths that have been computed. If the computation fails
   or cannot be performed, a negative response is required with an
   indication of the type of failure.

   A request/response protocol is also required for a PCE to communicate
   path computation requests to another PCE and for that PCE to return
   the path computation response. As described in [RFC4655], there is
   no reason to assume that two different protocols are needed, and this
   document assumes that a single protocol will satisfy all requirements
   for PCC-PCE and PCE-PCE communication.

   [RFC4655] describes four models of PCE: composite, external, multiple
   PCE path computation, and multiple PCE path computation with inter-
   PCE communication. In all cases except the composite PCE model, a
   PCECP is required. The requirements defined in this document are
   applicable to all models described in [RFC4655].

5. PCE Communication Protocol Generic Requirements

5.1. Basic Protocol Requirements

5.1.1. Commonality of PCC-PCE and PCE-PCE Communication

   A single protocol MUST be defined for PCC-PCE and PCE-PCE
   communication. A PCE requesting a path from another PCE can be
   considered a PCC, and in the remainder of this document we refer to
   all communications as PCC-PCE regardless of whether they are PCC-PCE
   or PCE-PCE.

5.1.2. Client-Server Communication

   PCC-PCE communication is by nature client-server based. The PCECP
   MUST allow a PCC to send a request message to a PCE to request path
   computation, and for a PCE to reply with a response message to the
   requesting PCC once the path has been computed.

   In addition to this request-response mode, there are cases where
   there is unsolicited communication from the PCE to the PCC (see
   Section 5.1.11).

5.1.3. Transport

   The PCECP SHOULD utilize an existing transport protocol that supports
   congestion control. This transport protocol may also be used to
   satisfy some requirements in other sections of this document, such as
   reliability. The PCECP SHOULD be defined for one transport protocol
   only in order to ensure interoperability. The transport protocol
   MUST NOT limit the size of the message used by the PCECP.

5.1.4. Path Computation Requests

   The path computation request message MUST include at least the source
   and destination. Note that the path computation request is for an
   LSP or LSP segment, and the source and destination supplied are the
   start and end of the computation being requested (i.e., of the LSP
   segment).

The path computation request message MUST support the inclusion of a
set of one or more path constraints, including but not limited to the
requested bandwidth or resources (hops, affinities, etc.) to
include/exclude. For example, a PCC may request the PCE to exclude
points of failure in the computation of a new path if an LSP setup
fails. The actual inclusion of constraints is a choice for the PCC
issuing the request. A list of core constraints that must be
supported by the PCECP is supplied in Section 5.1.16. Specification
of constraints MUST be future-proofed as described in Section 5.1.14.

The requester MUST be allowed to select from or prefer an advertised
list or minimal subset of standard objective functions and functional
options. An objective function is used by the PCE to process
constraints to a path computation request when it computes a path in
order to select the "best" candidate paths (e.g., minimum hop path),
and corresponds to the optimization criteria used for the computation
of one path, or the synchronized computation of a set of paths. In
the case of unsynchronized path computation, this can be, for
example, the path cost or the residual bandwidth on the most loaded
path link. In the case of synchronized path computation, this can
be, for example, the global bandwidth consumption or the residual
bandwidth on the most loaded network link.

A list of core objective functions that MUST be supported by the
PCECP is supplied in Section 5.1.17. Specification of objective
functions MUST be future-proofed as described in Section 5.1.14.

The requester SHOULD also be able to select a vendor-specific or
experimental objective function or functional option. Furthermore,
the requester MUST be allowed to customize the function/options in
use. That is, individual objective functions will often have
parameters to be set in the request from PCC to PCE. Support for the
specification of objective functions and objective parameters is
required in the protocol extensibility specified in Section 5.1.14.

A request message MAY include TE parameters carried by the MPLS/GMPLS
LSP setup signaling protocol. Also, it MUST be possible for the PCE
to apply additional objective functions. This might include policy-
based routing path computation for load balancing instructed by the
management plane.

Shortest path selection may rely either on the TE metric or on the
IGP metric [METRIC]. Hence the PCECP request message MUST allow the
PCC to indicate the metric type (IGP or TE) to be used for shortest
path selection. Note that other metric types may be specified in the
future.

There may be cases where a single path cannot fit a given bandwidth
request, while a set of paths could be combined to fit the request.
Such path combination to serve a given request is called load-
balancing. The request message MUST allow the PCC to indicate if
load-balancing is allowed. It MUST also include the maximum number
of paths in a load-balancing path group, and the minimum path
bandwidth in a load-balancing path group. The request message MUST
allow specification of the degree of disjointness of the members of
the load-balancing group.

5.1.5. Path Computation Responses

The path computation response message MUST allow the PCE to return
various elements including, at least, the computed path(s).

The protocol MUST be capable of returning any explicit path that
would be acceptable for use for MPLS and GMPLS LSPs once converted to
an Explicit Route Object for use in RSVP-TE signaling. In addition,
anything that can be expressed in an Explicit Route Object MUST be
capable of being returned in the computed path. Note that the
resultant path(s) may be made up of a set of strict or loose hops, or
any combination of strict and loose hops. Moreover, a hop may have
the form of a non-simple abstract node. See [RFC3209] for the
definition of strict hop, loose hop, and abstract node.

A positive response from the PCE MUST include the paths that have
been computed. A positive PCECP computation response MUST support
the inclusion of a set of attributes of the computed path, such as
the path costs (e.g., cumulative link TE metrics and cumulative link
IGP metrics) and the computed bandwidth. The latter is useful when a
single path cannot serve the requested bandwidth and load balancing
is applied.

When a path satisfying the constraints cannot be found, or if the
computation fails or cannot be performed, a negative response MUST be
sent. This response MAY include further details of the reason(s) for
the failure and MAY include advice about which constraints might be
relaxed to be more likely to achieve a positive result.

The PCECP response message MUST support the inclusion of the set of
computed paths of a load-balancing path group, as well as their
respective bandwidths.

5.1.6. Cancellation of Pending Requests

A PCC MUST be able to cancel a pending request using an appropriate
message. A PCC that has sent a request to a PCE and no longer needs
a response, for instance, because it no longer wants to set up the

associated service, MUST be able to notify the PCE that it can clear
the request (i.e., stop the computation if already started, and clear
the context). The PCE may also wish to cancel a pending request
because of some congested state.

5.1.7. Multiple Requests and Responses

It MUST be possible to send multiple path computation requests within
the same request message. Such requests may be correlated (e.g.,
requesting disjoint paths) or uncorrelated (requesting paths for
unrelated services). It MUST be possible to limit by configuration
of both PCCs and PCEs the number of requests that can be carried
within a single message.

Similarly, it MUST be possible to return multiple computed paths
within the same response message, corresponding either to the same
request (e.g., multiple suited paths, paths of a load-balancing path
group) or to distinct requests, correlated or not, of the same
request message or distinct request messages.

It MUST be possible to provide "continuation correlation" where all
related requests or computed paths cannot fit within one message and
are carried in a sequence of correlated messages.

The PCE MUST inform the PCC of its capabilities. Maximum acceptable
message sizes and the maximum number of requests per message
supported by a PCE MAY form part of PCE capabilities advertisement
[PCE-DISC-REQ] or MAY be exchanged through information messages from
the PCE as part of the protocol described here.

It MUST be possible for a PCC to specify, in the request message, the
maximum acceptable response message sizes and the maximum number of
computed paths per response message it can support.

It MUST be possible to limit the message size by configuration on
PCCs and PCEs.

5.1.8. Reliable Message Exchange

The PCECP MUST support reliable transmission of PCECP packets. This
may form part of the protocol itself or may be achieved by the
selection of a suitable transport protocol (see Section 5.1.3).

In particular, it MUST allow for the detection and recovery of lost
messages to occur quickly and not impede the operation of the PCECP.

In some cases (e.g., after link failure), a large number of PCCs may
simultaneously send requests to a PCE, leading to a potential

saturation of the PCEs. The PCECP MUST support indication of
congestion state and rate limitation state. This should enable, for
example, a PCE to limit the rate of incoming request messages if the
request rate is too high.

The PCECP or its transport protocol MUST provide the following:

- Detection and report of lost or corrupted messages
- Automatic attempts to retransmit lost messages without reference to
  the application
- Handling of out-of-order messages
- Handling of duplicate messages
- Flow control and back-pressure to enable throttling of requests and
  responses
- Rapid PCECP communication failure detection
- Distinction between partner failure and communication channel
  failure after the PCECP communication is recovered

If it is necessary to add functions to PCECP to overcome shortcomings
in the chosen transport mechanisms, these functions SHOULD be based
on and re-use where possible techniques developed in other protocols
to overcome the same shortcomings. Functionality MUST NOT be added
to the PCECP where the chosen transport protocol already provides it.

5.1.9. Secure Message Exchange

The PCC-PCE communication protocol MUST include provisions to ensure
the security of the exchanges between the entities. In particular,
it MUST support mechanisms to prevent spoofing (e.g.,
authentication), snooping (e.g., preservation of confidentiality of
information through techniques such as encryption), and Denial of
Service (DoS) attacks (e.g., packet filtering, rate limiting, no
promiscuous listening). Once a PCC is identified and authenticated,
it has the same privileges as all other PCCs.

To ensure confidentiality, the PCECP SHOULD allow local policy to be
configured on the PCE to not provide explicit path(s). If a PCC
requests an explicit path when this is not allowed, the PCE MUST
return an error message to the requesting PCC and the pending path
computation request MUST be discarded.

Authorization requirements [RFC3127] include reject capability,
reauthorization on demand, support for access rules and filters, and
unsolicited disconnect.

   IP addresses are used to identify PCCs and PCEs. Where the PCC-PCE
   communication takes place entirely within one limited domain, the use
   of a private address space that is not available to customer systems
   MAY be used to help protect the information exchange, but other
   mechanisms MUST also be available.

   These functions may be provided by the transport protocol or directly
   by the PCECP. See Section 6 for further discussion of security
   considerations.

5.1.10. Request Prioritization

   The PCECP MUST allow a PCC to specify the priority of a computation
   request.

   Implementation of priority-based activity within a PCE is subject to
   implementation and local policy. This application processing is out
   of scope of the PCECP.

5.1.11. Unsolicited Notifications

   The normal operational mode is for the PCC to make path computation
   requests to the PCE and for the PCE to respond.

   The PCECP MUST support unsolicited notifications from PCE to PCC, or
   PCC to PCE. This requirement facilitates the unsolicited
   communication of information and alerts between PCCs and PCEs. As
   specified in Section 5.1.8, these notification messages must be
   supported by a reliable transmission protocol. The PCECP MAY also
   support response messages to the unsolicited notification messages.

5.1.12. Asynchronous Communication

   The PCC-PCE protocol MUST allow for asynchronous communication. A
   PCC MUST NOT have to wait for a response to one request before it can
   make another request.

   It MUST also be possible to have the order of responses differ from
   the order of the corresponding requests. This may occur, for
   instance, when path request messages have different priorities (see
   Requirement 5.1.10). A consequent requirement is that path
   computation responses MUST include a direct correlation to the
   associated request.

5.1.13. Communication Overhead Minimization

   The request and response messages SHOULD be designed so that the
   communication overhead is minimized. In particular, the overhead per

message SHOULD be minimized, and the number of bytes exchanged to
arrive at a computation answer SHOULD be minimized. Other
considerations in overhead minimization include the following:

- the number of background messages used by the protocol or its
  transport protocol to keep alive any session or association
  between the PCE and PCC
- the processing cost at the PCE (or PCC) associated with
  request/response messages (as distinct from processing the
  computation requests themselves)

5.1.14. Extensibility

The PCECP MUST provide a way for the introduction of new path
computation constraints, diversity types, objective functions,
optimization methods and parameters, and so on, without requiring
major modifications in the protocol.

For example, the PCECP MUST be extensible to support various PCE-
based applications, such as the following:

- intra-area path computation
- inter-area path computation [PCECP-INTER-AREA]
- inter-AS intra provider and inter-AS inter-provider path
  computation [PCECP-INTER-AS]
- inter-layer path computation [PCECP-INTER-LAYER]

The PCECP MUST support the requirements specified in the
application-specific requirements documents. The PCECP MUST also
allow extensions as more PCE applications will be introduced in the
future.

The PCECP SHOULD also be extensible to support future applications
not currently in the scope of the PCE working group, such as, for
instance, point-to-multipoint path computations, multi-hop pseudowire
path computation, etc.

Note that application specific requirements are out of the scope of
this document and will be addressed in separate requirements
documents.

5.1.15. Scalability

The PCECP MUST scale well, at least as good as linearly, with an
increase of any of the following parameters. Minimum order of
magnitude estimates of what the PCECP should support are given in
parenthesis (note: these are requirements on the PCECP, not on the
PCE):

        - number of PCCs (1000/domain)
        - number of PCEs (100/domain)
        - number of PCCs communicating with a single PCE (1000)
        - number of PCEs communicated to by a single PCC (100)
        - number of domains (20)
        - number of path request messages (average of 10/second/PCE)
        - handling bursts of requests (burst of 100/second/PCE within a 10-
          second interval).

     Note that path requests can be bundled in path request messages, for
     example, 10 PCECP request messages/second may correspond to 100 path
     requests/second.

     Bursts of requests may arise, for example, after a network outage
     when multiple recomputations are requested. The PCECP MUST handle
     the congestion in a graceful way so that it does not unduly impact
     the rest of the network, and so that it does not gate the ability of
     the PCE to perform computation.

5.1.16. Constraints

     This section provides a list of generic constraints that MUST be
     supported by the PCECP. Other constraints may be added to service
     specific applications as identified by separate application-specific
     requirements documents. Note that the provisions of Section 5.1.14
     mean that new constraints can be added to this list without impacting
     the protocol to a level that requires major protocol changes.

     The set of supported generic constraints MUST include at least the
     following:

     o MPLS-TE and GMPLS generic constraints:
        - Bandwidth
        - Affinities inclusion/exclusion
        - Link, Node, Shared Risk Link Group (SRLG) inclusion/exclusion
        - Maximum end-to-end IGP metric
        - Maximum hop count
        - Maximum end-to-end TE metric
        - Degree of paths disjointness (Link, Node, SRLG)

     o MPLS-TE specific constraints
        - Class-type
        - Local protection
        - Node protection
        - Bandwidth protection

     o GMPLS specific constraints
       - Switching type, encoding type
       - Link protection type

5.1.17. Objective Functions Supported

   This section provides a list of generic objective functions that MUST
   be supported by the PCECP. Other objective functions MAY be added to
   service specific applications as identified by separate application-
   specific requirements documents. Note that the provisions of Section
   5.1.14 mean that new objective functions MAY be added to this list
   without impacting the protocol.

   The PCECP MUST support at least the following "unsynchronized"
   functions:

   - Minimum cost path with respect to a specified metric
     (shortest path)
   - Least loaded path
   - Maximum available bandwidth path

   Also, the PCECP MUST support at least 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

5.2. Deployment Support Requirements

5.2.1. Support for Different Service Provider Environments

   The PCECP must at least support the following environments:

   - MPLS-TE and GMPLS networks
   - Packet and non-packet networks
   - Centralized and distributed PCE path computation
   - Single and multiple PCE path computation

   For example, PCECP is possibly applicable to packet networks (e.g.,
   IP networks), non-packet networks (e.g., time-division multiplexed
   (TDM) transport), and perhaps to multi-layer GMPLS control plane
   environments. Definitions of centralized, distributed, single, and
   multiple PCE path computation can be found in [RFC4655].

5.2.2. **Policy Support**

   The PCECP MUST allow for the use of policies to accept/reject
   requests. It MUST include the ability for a PCE to supply sufficient
   detail when it rejects a request for policy reasons to allow the PCC
   to determine the reason for rejection or failure. For example,
   filtering could be required for a PCE that serves one domain (perhaps
   an AS) such that all requests that come from another domain (AS) are
   rejected. However, specific policy details are left to application-
   specific PCECP requirements. Actual policies, configuration of
   policies, and applicability of policies are out of scope.

   Note that work on supported policy models and the corresponding
   requirements/implications is being undertaken as a separate work item
   in the PCE working group.

   PCECP messages MUST be able to carry transparent policy information.

5.3. **Aliveness Detection & Recovery Requirements**

5.3.1. **Aliveness Detection**

   The PCECP MUST allow a PCC/PCE to

   - check the liveliness of the PCC-PCE communication,
   - rapidly detect PCC-PCE communication failure (indifferently to
     partner failure or connectivity failure), and
   - distinguish PCC/PCE node failures from PCC-PCE connectivity
     failures, after the PCC-PCE communication is recovered.

   The aliveness detection mechanism MUST ensure reciprocal knowledge of
   PCE and PCC liveness.

5.3.2. **Protocol Recovery**

   In the event of the failure of a sender or of the communication
   channel, the PCECP, upon recovery, MUST support resynchronization of
   information (e.g., PCE congestion status) and requests between the
   sender and the receiver; this SHOULD be arranged so as to minimize
   repeat data transfer.

5.3.3. **LSP Rerouting & Reoptimization**

   If an LSP fails owing to the failure of a link or node that it
   traverses, a new computation request may be made to a PCE in order to
   repair the LSP. Since the PCC cannot know that the PCE's TED has
   been updated to reflect the failure network information, it is useful
   to include this information in the new path computation request.

Also, in order to re-use the resources used by the old LSP, it may be
advantageous to indicate the route of the old LSP as part of the new
path computation request.

Hence the path computation request message MUST allow an indication
of whether the computation is for LSP restoration, and it MUST
support the inclusion of the previously computed path as well as the
identity of the failed element. Note that the old path might only be
useful if the old LSP has not yet been torn down. The PCE MAY choose
to take failure indication information carried in a given request
into account when handling subsequent requests. This should be
driven by local policy decision.

Note that a network failure may impact a large number of LSPs. In
this case, a potentially large number of PCCs will simultaneously
send requests to the PCE. The PCECP MUST properly handle such
overload situations, such as, for instance, through throttling of
requests as set forth in Section 5.1.8.

The path computation request message MUST support TE LSP path
reoptimization and the inclusion of a previously computed path. This
will help ensure optimal routing of a reoptimized path, since it will
allow the PCE to avoid double bandwidth accounting and help reduce
blocking issues.

6. Security Considerations

Key management MUST be provided by the PCECP to provide for the
authenticity and integrity of PCECP messages. This will allow
protecting against PCE or PCC impersonation and also against message
content falsification.

The impact of the use of a PCECP MUST be considered in light of the
impact that it has on the security of the existing routing and
signaling protocols and techniques in use within the network.
Intra-domain security is impacted since there is a new interface,
protocol, and element in the network. Any host in the network could
impersonate a PCC and receive detailed information on network paths.
Any host could also impersonate a PCE, both gathering information
about the network before passing the request on to a real PCE and
spoofing responses. Some protection here depends on the security of
the PCE discovery process (see [PCE-DISC-REQ]). An increase in
inter-domain information flows may increase the vulnerability to
security attacks, and the facilitation of inter-domain paths may
increase the impact of these security attacks.

Of particular relevance are the implications for confidentiality
inherent in a PCECP for multi-domain networks. It is not necessarily

the case that a multi-domain PCE solution will compromise security,
but solutions MUST examine their impacts in this area.

Applicability statements for particular combinations of signaling,
routing, and path computation techniques are expected to contain
detailed security sections.

It should be observed that the use of an external PCE introduces
additional security issues. Most notable among these are the
following:

- Interception of PCE requests or responses
- Impersonation of PCE or PCC
- DoS attacks on PCEs or PCCs

The PCECP MUST address these issues in detail using authentication,
encryption, and DoS protection techniques. See also Section 5.1.9.

There are security implications of allowing arbitrary objective
functions, as discussed in Section 5.1.17, and the PCECP MUST allow
mitigating the risk of, for example, a PCC using complex objectives
to intentionally drive a PCE into resource exhaustion.

7. **Manageability Considerations**

Manageability of the PCECP MUST address the following considerations:

- The need for a MIB module for control and monitoring of PCECP
- The need for built-in diagnostic tools to test the operation of the
  protocol (e.g., partner failure detection, Operations
  Administration and Maintenance (OAM), etc.)
- Configuration implications for the protocol

PCECP operations MUST be modeled and controlled through appropriate
MIB modules. There are enough specific differences between PCCs and
PCEs to lead to the need of defining separate MIB modules.
Statistics gathering will form an important part of the operation of
the PCECP. The MIB modules MUST provide information that will allow
an operator to determine PCECP historical interactions and the
success rate of requests. Similarly, it is important for an operator
to be able to determine PCECP and PCE load and whether an individual
PCC is responsible for a disproportionate amount of the load. It
MUST be possible, through use of MIB modules, to record and inspect
statistics about the PCECP communications, including issues such as
malformed messages, unauthorized messages, and messages discarded
owing to congestion.

The new MIB modules should also be used to provide notifications
(traps) when thresholds are crossed or when important events occur.
For example, the MIB module may support indication of exceeding the
congestion state threshold or rate limitation state.

PCECP techniques must enable a PCC to determine the liveness of a PCE
both before it sends a request and in the period between sending a
request and receiving a response.

It is also important for a PCE to know about the liveness of PCCs to
gain a predictive view of the likely loading of a PCE in the future
and to allow a PCE to abandon processing of a received request.

The PCECP MUST support indication of congestion state and rate
limitation state, and MAY allow the operator to control such a
function.

[8](#). Contributors

This document is the result of the PCE Working Group PCECP
requirements design team joint effort. In addition to the
authors/editors listed in the "Authors' Addresses" section, the
following are the design team members who contributed to the
document:

Alia K. Atlas
Google Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043 USA
EMail: akatlas@alum.mit.edu

Arthi Ayyangar
Nuova Systems,
2600 San Tomas Expressway
Santa Clara, CA 95051
EMail: arthi@nuovasystems.com

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

Igor Bryskin
Independent Consultant
EMail: i_bryskin@yahoo.com

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

Durga Gangisetti
MCI
EMail: durga.gangisetti@mci.com

Kenji Kumaki
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460, JAPAN
Phone: 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

9. Acknowledgements

10. References

10.1. Normative References

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

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

10.2. Informative References

[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.

[PCE-DISC-REQ]   Le Roux, J.L., et al., "Requirements for Path
                    Computation Element (PCE) Discovery", Work in
                    Progress.

[PCECP-INTER-AREA] Le Roux, J.L., et al., "PCE Communication
                    Protocol (PCECP) specific requirements for
                    Inter-Area (G)MPLS Traffic Engineering", Work in
                    Progress.

[PCECP-INTER-LAYER] Oki, E., et al., "PCC-PCE Communication
                    Requirements for Inter-Layer Traffic
                    Engineering", Work in Progress.

[PCECP-INTER-AS]   Bitar, N., Zhang, R., Kumaki, K., "Inter-AS
                    Requirements for the Path Computation Element
                    Communication Protocol (PCECP)", Work in
                    Progress.

[RFC3209]   Awduche, D., Berger, L., Gan, D., Li, T.,
                    Srinivasan, V., and G. Swallow, "RSVP-TE:
                    Extensions to RSVP for LSP Tunnels", RFC 3209,
                    December 2001.

[RFC3127]   Mitton, D., St.Johns, M., Barkley, S., Nelson,
                    D., Patil, B., Stevens, M., and B. Wolff,
                    "Authentication, Authorization, and Accounting:
                    Protocol Evaluation", RFC 3127, June 2001.

Authors' Addresses

    Jerry Ash (Editor)
    AT&T
    Room MT D5-2A01
    200 Laurel Avenue
    Middletown, NJ 07748, USA

    Phone: (732)-420-4578
    EMail: gash@att.com

    Jean-Louis Le Roux (Editor)
    France Telecom
    2, avenue Pierre-Marzin
    22307 Lannion Cedex, FRANCE

    EMail: jeanlouis.leroux@orange-ft.com