

Networking Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: November 12, 2007

JP. Vasseur, Ed.  
Cisco Systems, Inc  
JL. Le Roux  
France Telecom  
Y. Ikejiri  
NTT Communications Corporation  
May 11, 2007

**A set of monitoring tools for Path Computation Element based  
Architecture  
draft-vasseur-pce-monitoring-03.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](#).

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>.

This Internet-Draft will expire on November 12, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

A Path Computation Element (PCE) based architecture has been specified for the computation of Traffic Engineering (TE) Label Switched Paths (LSPs) in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks in the context of single or

multiple domains (where a domain is referred to as a collection of network elements within a common sphere of address management or path computational responsibility such as IGP areas and Autonomous Systems). In PCE-based environments it is thus critical to monitor the state of the path computation chain for troubleshooting and performance monitoring purposes: liveness of each element (PCE) involved in the PCE chain, detection of potential resource contention states, statistics in term of path computation times are examples of such metrics of interest. This document specifies procedures and extensions to the Path Computation Element Protocol (PCEP) in order to gather such information.

#### Requirements Language

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



## Table of Contents

<a href="#">1.</a>	<a href="#">Terminology</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Introduction</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Path Computation Monitoring messages</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">Path Computation Monitoring Request message (PCMonReq)</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">Path Monitoring Reply message (PCMonRep)</a>	<a href="#">8</a>
<a href="#">4.</a>	<a href="#">Path Computation Monitoring Objects</a>	<a href="#">9</a>
<a href="#">4.1.</a>	<a href="#">MONITORING Object</a>	<a href="#">9</a>
<a href="#">4.2.</a>	<a href="#">PCE-ID Object</a>	<a href="#">11</a>
<a href="#">4.3.</a>	<a href="#">PROC-TIME Object</a>	<a href="#">12</a>
<a href="#">4.4.</a>	<a href="#">CONGESTION Object</a>	<a href="#">14</a>
<a href="#">4.5.</a>	<a href="#">TIMESTAMP Object</a>	<a href="#">15</a>
<a href="#">5.</a>	<a href="#">Multi-destination monitoring</a>	<a href="#">15</a>
<a href="#">6.</a>	<a href="#">Policy</a>	<a href="#">15</a>
<a href="#">7.</a>	<a href="#">Elements of procedure</a>	<a href="#">15</a>
<a href="#">8.</a>	<a href="#">Manageability Considerations</a>	<a href="#">16</a>
<a href="#">9.</a>	<a href="#">To be considered in a further revision of this document</a>	<a href="#">16</a>
<a href="#">10.</a>	<a href="#">IANA Considerations</a>	<a href="#">16</a>
<a href="#">11.</a>	<a href="#">Security Considerations</a>	<a href="#">17</a>
<a href="#">12.</a>	<a href="#">Acknowledgements</a>	<a href="#">18</a>
<a href="#">13.</a>	<a href="#">References</a>	<a href="#">18</a>
<a href="#">13.1.</a>	<a href="#">Normative References</a>	<a href="#">18</a>
<a href="#">13.2.</a>	<a href="#">Informative References</a>	<a href="#">18</a>
	<a href="#">Authors' Addresses</a>	<a href="#">19</a>
	<a href="#">Intellectual Property and Copyright Statements</a>	<a href="#">20</a>



## 1. Terminology

AS: Autonomous System.

LSR: Label Switching Router.

PCC (Path Computation Client): any client application requesting a path computation to be performed by a Path Computation Element.

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.

TE LSP: Traffic Engineering Label Switched Path.

TED: Traffic Engineering Database.

## 2. Introduction

The Path Computation Element (PCE) based architecture has been specified in [[RFC4655](#)] for the computation of Traffic Engineering (TE) Label Switched Paths (LSPs) in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks in the context of single or multiple domains where a domain is referred to as a collection of network elements within a common sphere of address management or path computational responsibility such as IGP areas and Autonomous Systems.

In PCE-based environments, it is critical to monitor the state of the path computation chain for troubleshooting and performance monitoring purposes: liveness of each element (PCE) involved in the PCE chain, detection of potential resource contention states, statistics in term of path computation times are examples of such metrics of interest. This document specifies procedures and extensions to the Path Computation Element Protocol (PCEP) ([[I-D.ietf-pce-pcep](#)]) in order to monitor the path computation chain and gather various performance metrics.

As discussed in [[RFC4655](#)], a TE LSP may be computed by one PCE (referred to as single PCE path computation) or several PCEs (referred to as multiple PCE path computation). In the former case, the PCC may be able to use IGP extensions to check the liveness of the PCE (see [[I-D.ietf-pce-disco-ospf](#)] and [[I-D.ietf-pce-disco-isis](#)]) or PCEP using Keepalive messages. In contrast, when multiple PCEs are involved in the path computation chain an example of which is the BRPC procedure defined in [[I-D.ietf-pce-brpc](#)], the PCC's visibility may be limited to the first



PCE involved in the path computation chain. Thus, it is critical to define mechanisms in order to monitor the state of the path computation chain.

The aim of this document is to specify PCEP extensions in order to gather various state metrics along the path computation chain. In this document we call a "state metric" a metric that characterizes a PCE state. For example, such metric can have a form of a boolean (PCE is alive or not, PCE is congested or not) or a performance metric (path computation time at each PCE).

PCE state metrics collection can be gathered in two different contexts: in band or out of band. By "In band" we refer to the situation whereby a PCC requests to gather metrics in the context of a path computation request. For example, a PCC may send a path computation request to a PCE and may want to know the processing time of that request in addition to the computed path. Conversely, if the request is "out of band", PCE state metric collection is performed as a standalone request (e.g. check the liveness of a specific PCE chain, collect the average processing time computed over the last 5mn period on one or more PCE(s)).

In this document we define two monitoring request types: general and specific. A general monitoring request relates to the collection of a PCE state metric(s) that is not coupled to a particular path computation request (e.g. average CPU load on a PCE). Conversely, a specific monitoring request relates to a particular path computation request (processing time to complete the path computation for a TE LSP).

### **3. Path Computation Monitoring messages**

As defined in [[I-D.ietf-pce-pcep](#)], a PCEP message consists of a common header followed by a variable length body made of a set of objects that can either be mandatory or optional. As a reminder, an object is said to be mandatory in a PCEP message when the object must be included for the message to be considered as valid. The P flag (defined in [[I-D.ietf-pce-pcep](#)]) is located in the common header of each PCEP object and can be set by a PCEP peer to enforce a PCE to take into account the related information during the path computation. Because the P flag exclusively relates to a path computation request, it MUST be cleared in the two PCEP messages (PCMonReq and PCMonRep message) defined in this document.

For each PCEP message type a set of rules is defined that specify the set of objects that the message can carry. We use the Backus-Naur Form (BNF) to specify such rules. Square brackets refer to optional





sub-sequences. An implementation MUST form the PCEP messages using the object ordering specified in this document.

In this document we define two PCEP messages referred to as the Path Computation Monitoring request (PCMonReq) and Path Computation Monitoring Reply (PCMonRep) messages so as to handle "out of band" monitoring request. The aim of the PCMonReq message sent by a PCC to a PCE is to gather one or more PCE state metrics on a set of PCEs involved in a path computation chain. The PCMonRep message sent by a PCE to a PCC is used to provide such data.

### **3.1. Path Computation Monitoring Request message (PCMonReq)**

The Message-Type field of the PCEP common header for the PCMonReq message is set to 8 (To be confirmed by IANA).

There is one mandatory object that MUST be included within a PCMonReq message: the Monitoring object (see section [Section 4.1](#)). If the Monitoring object is missing, the receiving PCE MUST send an error message to the sender. Other objects are optional.



The format of a PCMonReq message is as follows:

```
<PCMonReq Message> ::= <Common Header>
                        <MONITORING>
                        [<pce-list>]
                        [<svec-list>]
                        [<request-list>]
```

where:

```
<svec-list> ::= <SVEC>
                [<OF>]
                [<svec-list>]
```

```
<request-list> ::= <request> [<request-list>]
```

```
<request> ::= <RP>
              <END-POINTS>
              [<LSPA>]
              [<BANDWIDTH>]
              [<metric-list>]
              [<RRO>]
              [<IRO>]
              [<LOAD-BALANCING>]
              [<XRO>]
```

```
<metric-list> ::= <METRIC> [<metric-list>]
```

```
<pce-list> ::= <PCE-ID> [<pce-list>]
```

The SVEC, RP, END-POINTS, LSPA, BANDWIDTH, METRIC, ERO, IRO and LOAD-BALANCING objects are defined in [[I-D.ietf-pce-pcep](#)]. The XRO object is defined in [[I-D.oki-pce-pcep-xro](#)] and the OF object is defined in [[I-D.leroux-pce-of](#)].

The PCMonReq message is used to gather various PCE state metrics along a path computation chain. The path computation chain may be determined by the PCC (in the form of a series of a series of PCE-ID objects defined in [Section 4.2](#).) or may alternatively be determined by the path computation procedure. For example, if the BRPC procedure ([[I-D.ietf-pce-brpc](#)]) is used to compute an inter-domain TE LSP, the PCE chain may be determined dynamically. In that case, the PCC sends a PCMonReq message that contains the PCEP objects that characterize the TE LSP attributes along with the monitoring objects (see [Section 4.1](#)) that list the set of metric(s) of interest.

Several PCE state metrics may be requested that are specified by a set of objects defined in [Section 4](#). Note that this set of objects is by all means not limitative and may be extended in further revision of this document.



For the sake of illustration, consider the three following examples:

Example 1: PCC1 requests to check the path computation chain should a path computation be requested for a specific TE LSP named T1. A PCMonReq message is sent that contains a MONITORING object specifying a path computation check, along with the appropriate set of objects (e.g. RP, END-POINTS, ...) that would be included in a PCReq message for T1.

Example 2: PCC1 requests a path computation for a TE LSP and also request to gather the processing time along the path computation chain selected for the computation of T1. A PCReq message is sent that also contains a MONITORING object that specifies the performance metric of interest. The PCRep message also comprises a PROC-TIME object defined in section [Section 4.1](#) that reports the computed metrics.

Example 3: PCC2 requests to gather performance metrics along the specific path computation chain <pce1, pce2, pce3, pce7>. A PCMonreq message is sent to PCE1 that contains a set of PCE-ID objects that identify PCE1, PCE2, PCE3 and PCE7 respectively.

### **[3.2.](#) Path Monitoring Reply message (PCMonRep)**

The PCMonRep message is used to provide PCE state metrics back to the requester for "out of band" monitoring requests. The Message-Type field of the PCEP common header for the PCMonRep message is set to 9 (To be confirmed by IANA).

There is one mandatory object that MUST be included within a PCMonRep message: the Monitoring object (see [Section 4.1](#)). If the Monitoring object is missing, the receiving PCE MUST send an error message to the requesting PCC. Other objects are optional.



The format of a PCReq message is as follows:

```
<PCMonRep Message> ::= <Common Header>
                        <MONITORING>
                        [<RP>]
                        [<metric-pce-list>]
```

where:

```
<metric-pce-list> ::= <metric-pce> [<metric-pce-list>]
```

```
<metric-pce> ::= [<PCE-ID>]
                  [<PROC-TIME>]
                  [<TIME-STAMP>]
                  [<CONGESTION>]
```

#### **4. Path Computation Monitoring Objects**

The PCEP objects defined in the document are compliant with the PCEP object format defined in [[I-D.ietf-pce-pcep](#)], with the P flag and the I flag cleared since these flags are exclusively related to path computation request.

Several objects are defined in this section that can be carried within the PCEP PCReq or PCRep messages defined in [[I-D.ietf-pce-pcep](#)] in case of "in band" monitoring requests. In case of "out of band" monitoring requests, the objects defined in this section are carried within PCMonReq and PCMonRep messages. Conversely, if the PCC requests the computation of the TE LSP in addition to gathering PCE state metrics (i.e. "In band" requests), these objects are carried within PCReq and PCRep messages.

##### **4.1. MONITORING Object**

The MONITORING object MUST be present within PCMonReq and PCMonRep messages ("out of band" monitoring requests) and MAY be carried within PCERep and PCReq messages ("in band" monitoring requests). There MUST be exactly once instance of the MONITORING object: if more than one instance of the MONITORING object is present, the recipient MUST only process the first instance and ignore other instances. The MONITORING object is used to specify the set of requested PCE state metrics.

The MONITORING Object-Class is to be assigned by IANA (recommended value=19)

The MONITORING Object-Type is to be assigned by IANA (recommended value=1)





The format of the MONITORING object body is as follows:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      Reserved      |      Flags      |I|C|P|G|L|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                monitoring-id-number                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                                                |
//                                Optional TLV(s)                                //
|                                                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Flags: 18 bits

The following flags are currently defined:

L (Liveness) - 1 bit: when set, this indicates that the state metric of interest is the PCE's liveness and thus the PCE MUST include a PCE-ID object in the corresponding reply.

G (General) - 1 bit: when set, this indicates that the monitoring request is a general monitoring request. When the requested performance metric is specific, the G bit MUST be cleared.

P (Processing Time) - 1 bit: the P bit of the MONITORING object carried in a PCMonReq or a PCReq message is set to indicate that the processing times is a metric of interest, in which case a PROC-TIME object MUST be inserted in the corresponding PCMonRep or PCRep message. The P bit MUST always be set in a PCMonRep message if also set in the corresponding PCMonReq message.

C (Congestion) - 1 bit: The C bit of the MONITORING object carried in a PCMonReq or a PCReq message is set to indicate that the congestion status is a metric of interest, in which case a CONGESTION object MUST be inserted in the corresponding PCMonRep or PCRep message. The C bit MUST always be set in a PCMonRep message if also set in the corresponding PCMonReq message.

I (Incomplete) - 1 bit: the I bit MUST be set by a PCE that supports the PCMonReq message, which does not trigger any policy violation but that cannot provide the set of requested performance metrics for unspecified reasons.

Monitoring-id-number (32 bits). The monitoring-id-number value combined with the source IP address of the PCC and the PCE address uniquely identify the monitoring request context. The monitoring-id-



number MUST be incremented each time a new monitoring is sent to a PCE. The value 0x00000000 is considered as invalid. If no reply to a monitoring request is received from the PCE, and the PCC wishes to resend its path computation monitoring request, the same monitoring-id-number MUST be used. Conversely, different monitoring-id-number MUST be used for different requests sent to a PCE. The same monitoring-id-number may be used for path computation monitoring requests sent to different PCEs. The path computation monitoring reply is unambiguously identified by the IP source address of the replying PCE.

Unassigned bits are considered as reserved and MUST be set to zero on transmission.

No optional TLVs are currently defined.

#### **4.2. PCE-ID Object**

The PCE-ID Object is used to specify a PCE's IP address.

A set of PCE-ID objects may be inserted within a PCReq or a PCMonReq message to specify the PCE for which PCE state metrics are requested and in a PCMonRep or a PCRep message to record the IP address of the PCE reporting PCE state metrics or that was involved in the path computation chain.

Two PCE-ID objects (for IPv4 and IPv6) are defined. PCE-ID Object-Class is to be assigned by IANA (recommended value=20) PCE-ID Object-Type is to be assigned by IANA (recommended value=1 for IPv4 and 2 for IPv6)

The format of the PCE-ID Object is as follows:



The format of the PCE-ID object body for IPv4 and IPv6 are as follows:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv4 Address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |
|                               |
|                               |
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The PCE-ID object body has a fixed length of 4 octets for IPv4 and 16 octets for IPv6.

A PCE MUST use the same IP address as the address used in the PCE-ADDRESS sub-TLV defined in [[I-D.ietf-pce-disco-proto-ospf](#)] and [[I-D.ietf-pce-disco-proto-isis](#)] should a dynamic discovery mechanism be used for PCE discovery.

### 4.3. PROC-TIME Object

The PROC-TIME object MUST be present within a PCMonRep or a PCRep message if the P bit of the MONITORING object carried within the corresponding PCMonReq or PCReq message is set. The PROC-TIME object is used to report various processing time related metrics.

#### 1) Case of general monitoring requests

A PCC may request processing time metrics for general monitoring requests (e.g. the PCC may want to know the minimum, maximum and average processing times on a particular PCE). In this case, general requests can only be made by using PCMonReq/PCMonRep messages. The processing-time field (as explained below) is exclusively used for specific monitoring requests and MUST be cleared for general monitoring requests. The algorithm(s) used by a PCE to compute the Min, Average, Max and Variance of the processing times are out of the scope of this document (A PCE may decide to compute the minimum processing time over a period of times, for the last N path computation requests, ...).

#### 2) Case of specific monitoring requests



In the case of a specific request, the algorithm(s) used by a PCE to compute the Processing-time metrics are out of the scope of this document but a flag is specified that is used to indicate to the requester whether the processing time value was estimated or computed. The PCE may either (1) estimate the processing time without performing an actual path computation or (2) effectively perform the computation to report the processing time. In the former case, the E bit of the PROC-TIME object MUST be set. The G bit MUST be cleared and the Min-processing-time, Max-processing-time, Average-processing-time and Variance-processing-time MUST be set to 0x00000000.

When the processing time is requested in addition to a path computation (case where the MONITORING object is carried within a PCReq message), the PROC-TIME object always report the actual processing time for that request and thus the E bits MUST be cleared.

The PROC-TIME Object-Class is to be assigned by IANA (recommended value=21)

The PROC-TIME Object-Type is to be assigned by IANA (recommended value=1)

The format of the PROC-TIME object body is as follows:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Reserved           |           Flags           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Processing-time           |E|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Min-processing-time           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Max-processing-time           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Average-processing time           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Variance-processing-time           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Flags: 18 bits - No Flags are currently defined:

E (Estimated) - 1 bit: when set, this indicates that the reported metric value is based on estimated processing time as opposed to actual computation(s).

Current-processing-time: This field indicates in milliseconds the





processing time for the path computation of interest characterized in the corresponding PCMonReq message.

Min-processing-time: This field indicates in milliseconds the minimum processing time.

Max-processing-time: This field indicates in milliseconds the maximum processing time.

Average-processing-time: This field indicates in milliseconds the average processing time.

Variance-processing-time: This field indicates in milliseconds the variance of the processing times.

Unassigned bits are considered as reserved and MUST be set to zero on transmission.

More granularity may be introduced in further revision of this document to get a monitoring metric for a general request of a particular class (e.g. all PCReq of priority X).

#### **4.4. CONGESTION Object**

The CONGESTION object MUST be present within a PCMonRep or a PCRep message if the C bit of the MONITORING object carried within the corresponding PCMonReq or PCReq message is set. The CONGESTION object is used to report a PCE processing congestion state. The CONGESTION Object-Class is to be assigned by IANA (recommended value=22) The CONGESTION Object-Type is to be assigned by IANA (recommended value=1)

The format of the CONGESTION object body is as follows:

```

0               1               2               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|C|           Reserved           | Congestion Duration          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

C (Congestion) - 1 bit: when set, this indicates that PCE is congested, in which case the congestion duration may be non nul. When cleared this indicates that the PCE is not congested.

Congestion duration - 16 bits: This field indicates in seconds the estimated congestion duration.



#### **4.5. TIMESTAMP Object**

A TIMESTAMP object will be specified in a further revision of this document that could be used to indicate when a PCMonReq message has been received by a PCE and when the PCMonReq message has been relayed to the next-hop PCE or the time at which a PCMonRep message has been sent to the requester.

#### **5. Multi-destination monitoring**

In a further revision of this document, a new object will be specified allowing a PCC or a user to gather PCE state metrics for a set of destinations using a single PCMonReq message. For example, using a single PCMonreq message originated by a PCC, PCE state metrics for the set of path computation chains involved in the computation of a set of TE LSPs will be gathered. Such set of destinations could be specified in the form of a subnets.

#### **6. Policy**

The receipt of a PCMonReq message may trigger a policy violation on some PCE in which case the PCE MUST send a PCErr message with Error-Type=5 and Error-value=3 (To be Confirmed by IANA).

#### **7. Elements of procedure**

I bit processing: as indicated in section [Section 4.1](#), the I bit MUST be set by a PCE that supports the PCMonReq message, which does not trigger any policy violation but that cannot provide the set of required performance metrics for unspecified reasons. Once set, the I bit MUST NOT be changed by a receiving PCE.

Reception of a PCMonReq message: upon receiving a PCMonReq message:

- 1) If the PCE does not support the PCMonReq message, the PCE MUST send a PCErr message with Error-type=14 and Error-value=1.
- 2) If the PCE supports the PCMonReq message but the request is prohibited by policy, the PCE MUST send a PCErr message with Error-Type=5 and Error-value=3.
- 3) If the PCE supports the PCMonReq and the monitoring request is not prohibited by policy, the receiving PCE MUST first determine whether it is the last PCE of the path computation chain. If the PCE is not the last element of the path computation chain, the PCMonReq message



is relayed to the next hop PCE: such next-hop may either be specified by means of a PCE-ID object present in the PCMonReq message or dynamically determined by means of a procedure outside of the scope of this document. Conversely, if the PCE is the last PCE of the path computation chain, the PCE originates a PCMonRep message that contains the requested objects according to the set of requested PCE states metrics listed in the MONITORING object carried in the corresponding PCMonReq message.

Reception of a PCMonRep message: upon receiving a PCMonRep message, the PCE processes the request, adds the relevant objects to the PCMonRep message and forwards the PCMonRep message to the upstream requesting PCE or PCC.

Special case of Multi-destination monitoring: monitoring request related to more than one destinations may involve a set of path computation chains. In that case, a PCE sends each copy of the PCMonReq message to each downstream PCE of each path computation chain.

## **8. Manageability Considerations**

To be addressed in a further revision of this document.

## **9. To be considered in a further revision of this document**

It might be desirable to modify the format of the PCMonReq and PCMonRep messages to support the bundling of multiple performance metrics collection for a set of TE LSPs.

## **10. IANA Considerations**

Two new PCEP (specified in [[I-D.ietf-pce-pcep](#)]) messages are defined in this document:

Value	Meaning
8	Path Computation Monitoring Request (PCMonReq)
9	Path Computation Monitoring Reply (PCMonRep)

The following new PCEP objects are defined in this document.



Object-Class	Name
19	MONITORING Object-Type 1
20	PCE-ID Object-Type 1: IPv4 addresses 2: IPv6 addresses
21	PROC-TIME Object-Type 1
22	CONGESTION Object-Type 1

A new Error type for the PCErr message (see [[I-D.ietf-pce-pcep](#)]) is defined in this document (Error-Type and Error-value to be assigned by IANA).

Error-type	Meaning
14	Performance Monitoring not supported Error-value 1: Monitoring message not supported by one of PCEs along the domain path 2: MONITORING object missing in a PCMonReq message

A new Error-value for the PCErr message Error-types=4 (see [[I-D.ietf-pce-pcep](#)]) is defined in this document (Error-Type and Error-value to be assigned by IANA).

Error-type	Meaning
5	Performance Monitoring Policy violation 3: Monitoring message supported but rejected due to policy violation

## **11. Security Considerations**

To be addressed in a further revision of this document.





## **12. Acknowledgements**

The authors would like to thank Eiji Oki, Mach Chen and Dimitri Papadimitriou for their useful comments.

## **13. References**

### **13.1. Normative References**

- [I-D.ietf-pce-pcep]  
Roux, J. and J. Vasseur, "Path Computation Element (PCE) communication Protocol (PCEP)", [draft-ietf-pce-pcep-07](#) (work in progress), March 2007.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.

### **13.2. Informative References**

- [I-D.ietf-pce-brpc]  
Vasseur, J., "A Backward Recursive PCE-based Computation (BRPC) procedure to compute shortest inter-domain Traffic Engineering Label Switched Paths", [draft-ietf-pce-brpc-04](#) (work in progress), March 2007.
- [I-D.ietf-pce-disco-proto-isis]  
Roux, J., "IS-IS protocol extensions for Path Computation Element (PCE) Discovery", [draft-ietf-pce-disco-proto-isis-05](#) (work in progress), May 2007.
- [I-D.ietf-pce-disco-proto-ospf]  
Roux, J., "OSPF protocol extensions for Path Computation Element (PCE) Discovery", [draft-ietf-pce-disco-proto-ospf-05](#) (work in progress), May 2007.
- [I-D.leroux-pce-of]  
Roux, J., "Encoding of Objective Functions in Path Computation Element (PCE) communication and discovery protocols", [draft-leroux-pce-of-00](#) (work in progress), February 2007.
- [I-D.oki-pce-pcep-xro]



Oki, E. and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", [draft-oki-pce-pcep-xro-00](#) (work in progress), January 2007.

#### Authors' Addresses

JP Vasseur (editor)  
Cisco Systems, Inc  
1414 Massachusetts Avenue  
Boxborough, MA 01719  
USA

Email: [jpv@cisco.com](mailto:jpv@cisco.com)

JL Le Roux  
France Telecom  
2, Avenue Pierre-Marzin  
Lannion, 22307  
FRANCE

Email: [jeanlouis.leroux@orange-ftgroup.com](mailto:jeanlouis.leroux@orange-ftgroup.com)

Yuichi Ikejiri  
NTT Communications Corporation  
1-1-6, Uchisaiwai-cho, Chiyoda-ku  
Tokyo, 100-8019  
Japan

Email: : [y.ikejiri@ntt.com](mailto:y.ikejiri@ntt.com)



## Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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.

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, THE IETF TRUST 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.

## Intellectual Property

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](mailto:ietf-ipr@ietf.org).

## Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

