

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 1, 2012

H. Pouyllau
Alcatel-Lucent
R. Theillaud
Marben Products
J. Meuric
France Telecom Orange
September 29, 2011

**Extension to the Path Computation Element Communication Protocol for
Enhanced Errors and Notifications
draft-pouyllau-pce-enhanced-errors-03.txt**

Abstract

This document defines new error and notification TLVs for the PCE Communication Protocol (PCEP) [[RFC5440](#)]. It identifies the possible PCEP behaviors in case of error or notification. Thus, this draft extends error and notification types in order to associate predefined PCEP behaviors.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 1, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Terminology	3
2.	Conventions used in this document	4
3.	Introduction	5
3.1.	Examples	5
3.1.1.	Error use-case	5
3.1.2.	Notification use-case	6
3.2.	PCEP Description	6
3.3.	PCEP Behaviors	6
3.3.1.	PCEP Behaviors in Case of Error	7
3.3.2.	PCEP Behaviors in Case of Notification	7
3.3.3.	PCE peer identification	8
4.	Error and Notification Handling in PCEP	9
4.1.	Propagation TLV	9
4.2.	Error-criticality TLV	9
4.3.	Notification type TLV	9
4.4.	Behaviors and TLV combinations	10
5.	Propagation Restrictions	12
5.1.	Time-To-Live object	12
5.2.	DIFFUSION-LIST Object (DLO)	12
5.3.	Extension rules applied to existing errors and notifications	14
6.	Error and Notification Scenarios	18
6.1.	Local Error with a low level of criticality	18
6.2.	Propagated Error with a low level of criticality	18
6.3.	Local Error with a medium level of criticality	19
6.4.	Propagated Error with a medium level of criticality	20
6.5.	Local request-specific notification	20
6.6.	Local non request-specific notification	21
6.7.	Propagated request-specific notification	21
6.8.	Propagated non request-specific notification	22
7.	Security Considerations	23
8.	IANA Considerations	24
8.1.	PCEP TLV Type Indicators	24
8.2.	New TTL object	24
8.3.	New DLO object	25
9.	References	26
9.1.	Normative References	26
9.2.	Informational References	26
	Authors' Addresses	28

1. Terminology

PCE terminology is defined in [[RFC4655](#)].

PCEP Peer: An element involved in a PCEP session (i.e. a PCC or a PCE).

Source PCC: the PCC which, for a given path computation query, initiates the 1st PCEP request, which may trigger a chain of successive requests.

Target PCE: the PCE that can compute a path to the destination without having to query any other PCE.

2. 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 [[RFC2119](#)].

3. Introduction

The PCE Communication Protocol [[RFC5440](#)] is designed to be flexible and extensible in order to allow future evolutions or specific constraint support such as proposed in [[I-D.ietf-pce-vendor-constraints](#)]. Crossing different PCE implementations (e.g. from different providers or due to different releases), a PCEP request may encounter unknown errors or notification messages. In such a case, the PCEP RFC [[RFC5440](#)] specifies to send a specific error code to the PCEP peer.

In the context of path computation crossing different routing domains or autonomous systems, the number of different PCE system specificities is potentially high, thus possibly leading to divergent and unstable situations. Such phenomenon can also occur in homogeneous cases since PCE systems have their own policies that can introduce differences in requests treatment even for requests having the same destination. Extending error and notification codes allow generalizing PCEP behaviors over heterogeneous PCE systems. Dealing with heterogeneity is a major challenge considering PCE applicability, particularly in multi-domain contexts. Thus, extending such error codes and PCEP behaviors accordingly would improve interoperability among different PCEP implementations and would solve some of these unstable issues. However, some of them would still remain (e.g. the divergences in request treatment introduced by different policies).

The purpose of this draft is to specify some PCEP error codes in order to generalize PCEP behaviors.

3.1. Examples

The two following scenarios underline the need for a normalization of the PCEP behaviors according to the error or notification type.

3.1.1. Error use-case

PCE(i-1) has sent a request to PCE(i) which has also sent a request to PCE(i+1). PCE(i-1) and PCE(i+1) have the same error semantic but not PCE(i). If PCE(i+1) throws an error type and value unknown by PCE(i). PCE(i) could then adopt any other behaviors and sends back to PCE(i-1) an error of type 2 (Capability not supported), 3 (Unknown Object) or 4 (Not supported Object) for instance. As a consequence, the path request would be cancelled but the error has no meaning for PCE(i-1) whereas if PCE(i) had simply forwarded the error sent by PCE(i+1), it would have been understood by PCE(i-1).

3.1.2. Notification use-case

PCE(i-1) has sent a request to PCE(i) which has also sent a request to PCE(i+1) but PCE(i+1) is overloaded. Without extensions, PCE(i+1) should send a notification of type 2 and a value flag giving its estimated congestion duration. PCE(i) can choose to stop the path computation and send a NO_PATH reply to PCE(i-1). Hence, PCE(i-1) ignores the congestion duration on PCE(i+1) and could seek it for further requests.

3.2. PCEP Description

One of the purposes of the PCE architecture is to compute paths across networks, but an added value is to compute such paths in inter-area/layer/domain environments. The PCE Communication Protocol [[RFC5440](#)] is based on the Transport Communication Protocol (TCP). Thus, to compute a path within the PCE architecture, several TCP/PCEP sessions have to be set up, in a peer-to-peer manner, along a set of identified PCEs.

When the PCEP session is up for 2 PCEP peers, the PCC of the first PCE System (the source PCC) sends a PCReq message. If the PCC does not receive any reply before the dead timer is out, then it goes back to the idle state. A PCC can expect two kinds of replies: a PCRep message containing one or more valid paths (EROs) or a negative PCRep message containing a NO-PATH object.

Beside PCReq and PCRep messages, notification and error messages, named respectively PCNtf and PCErr, can be sent. There are two types of notification messages: type 1 is for cancelling pending requests and type 2 for signaling a congestion of the PCE. Several error values are described in [[RFC5440](#)]. The error types concerning the session phase begin at 2, error type 1 values are dedicated to the initialization phase.

As the PCE Communication Protocol is built to work in a peer-to-peer manner (i.e. supported by a TCP Connection), it supposes that the 'deadtimer' of the source PCC is long enough to support the end-to-end distributed path computation process.

3.3. PCEP Behaviors

The exchange of messages in the PCE Communication Protocol is described in details when PCEP is in states OpenWait and KeepWait in [[RFC5440](#)]. When the session is up, message exchange is defined in [[RFC5440](#)] but detailed behavior is mostly let free to any specific implementation. [[RFC5441](#)] describes the Backward Recursive Path Computation (BRPC) procedure, and, because it considers an inter-

domain path computation, gives a bigger picture of the possible behaviors when the session is up.

3.3.1. PCEP Behaviors in Case of Error

[RFC5440] specifies that "a PCEP Error message is sent in several situations: when a protocol error condition is met or the request is not compliant with the PCEP specification". On this basis, and according to the other RFCs, the identified PCEP behaviors are the followings:

"Propagation": the received message requires to be propagated forwardly or backwardly (depending on which PCEP peer has sent the message) to a set of PCE peers;

Criticality level: in the different RFCs, error-types affects the state of the PCEP request or session in different manner; hence, different level of criticality can be observed:

Low-level of criticality: the received message does not affect the PCEP connection and further answer can still be expected;

Medium-level of criticality: the received message does not affect the PCEP connection but the request(s) is(are) cancelled;

High-level of criticality: the received message indicates that the PCE peer will close the session with its peer (and so pending requests associated by the error, if any, are cancelled)

The high-level of criticality has been extracted from [RFC5440] which associates such a behavior to error-type of 1 (errors raised during the PCEP session establishment). Hence, such errors are quite specific for the moment. For the sake of completeness, they have been included in this document.

3.3.2. PCEP Behaviors in Case of Notification

Notification messages can be employed in two different manners: during the treatment of a PCEP request, or independently from it to advertise information (in [RFC5440] the request id list within a PCNtf message is optional). Hence, three different behaviors can be identified:

- o "Local": the notification is or is not request-specific but does not imply any forward or backward propagation of the message;
- o "Request-specific Propagation": the received message requires to be propagated forwardly or backwardly (depending on which peer has sent the message) to the PCEP peers;
- o "Non request-specific Propagation": the received message must be propagated to any known peers (e.g. if PCE discovery is activated) or to a list of identified peers.

3.3.3. PCE peer identification

The propagation of errors and notifications affects the state of the PCE peers along the chain. In some cases, for instance a notification that a PCE is overloaded, the identification of the PCE peer - or that the sender PCE is not the direct neighbor - might be an important information for the PCE peers receiving the message.

4. Error and Notification Handling in PCEP

This section describes extensions to support error and notification messages with respect to the PCEP behavior description defined in [Section 3.3.1](#). This document does not intend to modify errors and notification types previously defined in existing documents (e.g. [\[RFC5440\]](#), [\[RFC5441\]](#), etc.).

4.1. Propagation TLV

To support the propagation behavior mentioned in [Section 3.3.1](#), we extend the PCEP-ERROR and NOTIFICATION objects by creating a new optional TLV to indicate whether the message has to be propagated or not. The allocation from the "PCEP TLV Type Indicators" sub-registry will be assigned by IANA and the request is documented in [Section 8](#).

The description is "Propagation", the length value is 2 bytes. The value field is set to default value 0 meaning that the message MUST NOT be propagated. If the value field is set to 1, the message MUST be propagated. [Section 5](#) specifies the destination and to limit the number of messages.

4.2. Error-criticality TLV

To support the shutdown behavior mentioned in [Section 3.3.1](#), we extend the PCEP-ERROR object by creating a new optional TLV to indicate whether the error is recoverable or not. The allocation from the "PCEP TLV Type Indicators" sub-registry will be assigned by IANA and the request is documented in [Section 8](#).

The description is "Error-criticality", the length value is 2 bytes and the value field is 1 byte. The value field is set to default value 0 meaning that the error has a low-level of criticality (so further messages can be expected for this request). If the value field is set to 1, the error has a medium-level of criticality and requests whose the identifiers appear MUST be cancelled (so no further messages can be expected for these requests). If the value field is set to 2, the error has a high-level of criticality, the connection is closed by the sender PCE peer.

4.3. Notification type TLV

To support the request-specific behavior mentioned in [Section 3.3.1](#), we extend the NOTIFICATION object by creating a new optional TLV to indicate whether the notification is request-specific or not. The allocation from the "PCEP TLV Type Indicators" sub-registry will be assigned by IANA and the request is documented in [Section 8](#).

The description is "Notification Type", the length value is 2 bytes and the value field is 1 byte. The value field is set to default value 0 meaning that the notification is not request-specific. If the value field is set to 1, the notification is request-specific.

4.4. Behaviors and TLV combinations

The propagation behavior MIGHT be combined with all criticality levels, thus leading to 6 different behaviors. In the case of a criticality level of 2, the session is closed by the PCE peer which sends the message. Hence, the criticality level is purely informative for the PCE peer which receives the message. If it is combined with a propagation behavior, then the PCE propagating the message MUST indicate the same level of criticality if it closes the session. Otherwise, it MUST use a criticality level of 1 if it does not close the session.

The TLVs defined in the sections above allow to cover all the possible behaviors listed in [Section 3.3.1](#). Hence, for an error message, the behaviors are covered as ensued, with TLVs included in a PCEP-ERROR object:

- o "Local Error with a low level of criticality" : TLV "Propagation" with value 0 and TLV "Error-criticality" with value 0
- o "Local Error with a medium level of criticality": TLV "Propagation" with value 0 and TLV "Error-criticality" with value 1
- o "Local Error with a high level of criticality": TLV "Propagation" with value 0 and TLV "Error-criticality" with value 2
- o "Propagated Error with a low level of criticality": TLV "Propagation" with value 1 and TLV "Error-criticality" with value 0
- o "Propagated Error with a medium level of criticality": TLV "Propagation" with value 1 and TLV "Error-criticality" with value 1
- o "Propagated Error with a high level of criticality": TLV "Propagation" with value 1 and TLV "Error-criticality" with value 2

For a notification message, the behaviors are covered as ensued, with TLVs included in a NOTIFICATION object:

- o "Local request-specific": TLV "Propagation" with value 0 and TLV "Notification Type" with value 1
- o "Local non request-specific": TLV "Propagation" with value 0 and TLV "Notification Type" with value 0

- o "Request-specific Propagation": TLV "Propagation" with value 1 and TLV "Notification Type" with value 1
- o "Non request-specific Propagation": TLV "Propagation" with value 1 and TLV "Notification Type" with value 0

5. Propagation Restrictions

In order to limit the propagation of errors and notifications, the following mechanisms SHOULD be used:

A Time-To-Live object: to limit the number of PCEP peers that will recursively receive the message;

A DIFFUSION-LIST object (DLO) which specifies the PCEP peer addresses or domains of PCEP peers the message must be propagate to;

History mechanisms: if a PCEP peer keeps track of the messages it has relayed, it could avoid propagating an error or notification it already received.

Such mechanisms SHOULD be used jointly or independently depending the error or notification behaviors they are associated to. Note that, a non request-specific propagated notification (TLV "Propagation" at value 1 and TLV "Notification Type" at value 0) MUST include a DLO and SHOULD include a TTL. The conditions of use for the TTL and DIFFUSION-LIST object are described in sections below.

5.1. Time-To-Live object

The TTL value is set to any integer value to indicate the number of PCEP peers that will recursively receive the message. This TTL SHOULD be used with propagated errors or notifications (TLV "Propagation" at value 1 in PCEP-ERROR or NOTIFICATION objects). Each PCEP peer MUST decrement the TTL value before propagating the message. When the TTL value is at 0, the message is no more propagated.

If the message has to be propagated, is request-specific (TLV "Propagation" at value 1 in PCEP-ERROR or NOTIFICATION objects, and TLV "Notification Type" at value 1 in a NOTIFICATION object), and there is no TTL or DIFFUSION-LIST object included, the message MUST reach the source PCC (or alternatively the target PCE).

5.2. DIFFUSION-LIST Object (DLO)

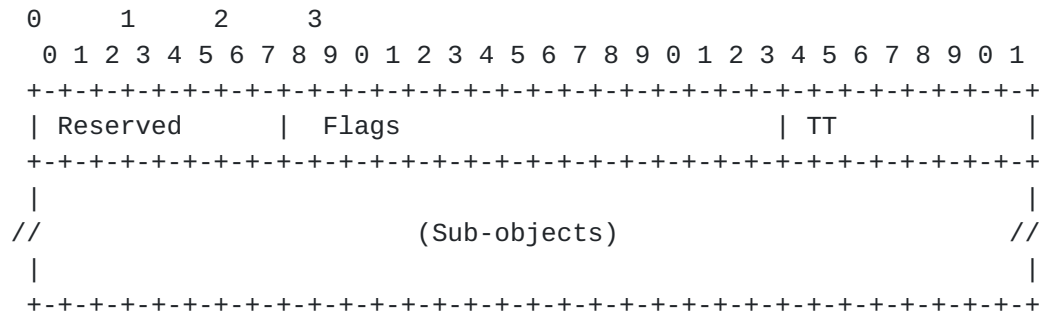
The DIFFUSION-LIST Object can be carried within a PCErr and a PCntf message and can either be used in a message sent by a PCC to a PCE or by a PCE to a PCC. The DLO MAY be used with propagated errors (TLV "Propagation" at value 1 in PCEP-ERROR object) and request-specific propagated notifications (TLV "Propagation" at value 1 and TLV "Notification Type" at value 1), and it MUST be used with non request-specific propagated notifications (TLV "Propagation" at value

1 and TLV "Notification Type" at value 0).

DIFFUSION-LIST Object-Class is 25.

DIFFUSION-LIST Object-Type is 1.

The format of the DIFFUSION-LIST body object is as follows:



Reserved (8 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

Flags (16 bits): No flags are currently defined. Unassigned flags MUST be set to zero on transmission and MUST be ignored on receipt.

TT (8 bits): The Target-type restricts the diffusion to certain peers. The following values are currently defined:

- 0: Any PCEP peer indicated in the list must be reached.
- 1: Only PCEs must be reached (and not PCC).
- 2: All PCEP peers with which a session is still opened must be reached

The DLO is made of sub-objects similar to the IRO defined in [\[RFC5440\]](#). The following sub-object types are supported.

Type Sub-object

- 1 IPv4 address
- 2 IPv6 address
- 4 Unnumbered Interface ID
- 5 OSPF area ID
- 32 Autonomous System number
- 33 Explicit eXclusion Route Sub-object (EXRS)

If the error or notification codes target specific PCEP peers, a DIFFUSION-LIST object avoids partially flooding all PCEP peers. Any PCEP peer receiving a PCerr or PCNTf message containing a PCEP-ERROR object, respectively a NOTIFICATION object, including a TLV "Propagation" at value 1, and where a DLO appears MUST remove from the DLO the addresses of the PCEP peers to whom it will propagate the message, before sending them the message. This is performed adding the PCEP peer addresses to the Explicit eXclusion Route Sub-object of the DLO. If a DIFFUSION-LIST object is empty, the PCEP peer MUST NOT propagate the message to any peer.

Note that, a Diffusion List Object could contain strict or loose addresses to refer to a network domain (e.g. an Autonomous System number, an OSPF area, an IP address). Hence, the PCEP peers targeted by the message would be the PCEP peers covering the corresponding domain. If an address is loose, each time a PCEP peer forwards a message to another PCEP peer of this address, it MUST add it to the Explicit eXclusion Route Sub-object (EXRS) of the DLO for any forwarded message. Hence, a PCE SHOULD avoid forwarding several times the same message to the same set of peers. Finally, when an address is loose, the forwarding SHOULD be restrained indicating what type of PCEP peers are targeted (i.e. PCE and/or PCC). Hence, a Target-Type is specified.

5.3. Extension rules applied to existing errors and notifications

Many existing normative references states on error definitions (see for instance [\[RFC5440\]](#), [\[RFC5441\]](#), [\[RFC5455\]](#), [\[RFC5521\]](#), [\[RFC5557\]](#), [\[RFC5886\]](#), [\[RFC6006\]](#)). According to the definitions provided in this document, the following rules are applicable:

Error-type 1, described in [\[RFC5440\]](#), relates to PCEP session establishment failures. All errors of this type are local (not to be propagated). Hence, if a TLV "Propagation" is added to the error message it MUST be set to value 0. Error-values 1,2,6,7 have a high level of criticality. Hence, if the TLV "Error-criticality" is included within a PCerr message of type 1 and value 1,2,6 or 7, it MUST have a value of 2.

Error-type 2,3,4, "Capability not supported", "Unknown object" and "Not supported object" respectively, described in [\[RFC5440\]](#): errors of this type MIGHT be propagated using the TLV "Propagation". Their level of criticality is defined as leading to cancel the path computation request (cf. [\[RFC5440\]](#)). Hence, if the TLV "Error-criticality" is included, it MUST have a value of 1. The error-value 4 of error-type 4 ("Unsupported parameter") associated to the BRPC procedure [\[RFC5441\]](#) SHOULD contain the TLV "Propagation" with a DIFFUSION-LIST object requesting a

propagation to the PCC at the origin of the request.

Error-type 5 refers to "Policy violation", error values for this type have been defined in [\[RFC5440\]](#), [\[RFC5541\]](#), [\[RFC5557\]](#), [\[RFC5886\]](#) and [\[RFC6006\]](#). In [\[RFC5440\]](#), it is specified that the path computation request MUST be cancelled when an error of type 5 occurs. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 1. As such errors might be conveyed to several PCEs, the TLV "Propagation" MIGHT be used.

Error-type 6 described as "Mandatory object missing" in [\[RFC5440\]](#), leads to the cancellation of the path computation request. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 1. The TLV "Propagation" MIGHT be used with such errors. The error-value defined in Monitoring object missing [\[RFC5886\]](#) is no exception to the rule.

Error-type 7 is described as "synchronized path computation request missing". In [\[RFC5440\]](#), it is specified that the referred synchronized path computation request MUST be cancelled when an error of type 5 occurs. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 1. The TLV "Propagation" MIGHT be used with such errors.

Error-type 8 is raised when a PCE receives a PCRep with an unknown request reference. If the TLV "Propagation" is used with error-type 8, it SHOULD be set at a value of 0. The TLV "Error-criticality" is not particularly relevant for error-type 8. Hence, if it used, it MUST have the value of 0.

Error-type 9 is raised when a PCE attempts to establish a second PCEP session. The existing session must be preserved. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 0. By definition, such an error message SHOULD NOT be propagated. Thus, if the TLV "Propagation" is used with error-type 9, it SHOULD be set at a value of 0.

Error-type 10 which refers to the reception of an invalid object is described in [\[RFC5440\]](#) no indication is provide on the cancellation of the path computation request. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 0. The TLV "Propagation" MIGHT be used with such errors with any value depending on the expected behavior.

Error-type 11 relates to "Unrecognized EXRS subobject" and is described in [\[RFC5521\]](#). No path computation request cancellation is required by [\[RFC5521\]](#). Hence, if the TLV "Error-criticality" is included, it MUST have a value of 0. The TLV "Propagation" MIGHT

be used with such errors with any value depending on the expected behavior.

Error-type 12 refers to "Diffserv-aware TE error" and is described in [\[RFC5455\]](#). Such errors are raised when the CLASSTYPE object of a PCReq is recognized but not supported by a PCE. [\[RFC5455\]](#) does not state about the path computation request when such errors are met. Hence, both "Propagation" and "Error-criticality" TLVs could be used within such error-types' messages and set to any specified values.

Error-type 13 on "BRPC procedure completion failure" is described in [\[RFC5441\]](#). [\[RFC5441\]](#) states that in such cases, the PCerr message MUST be relayed to the PCC. Hence, such messages SHOULD contain a TLV "Propagation" and a DIFFUSION-LIST object with a Target-Type of 0 and corresponding addresses or with a Target-Type of 2. It is not specified in [\[RFC5441\]](#) whether the path computation request should be canceled or not. If the procedure is not supported, it does not necessarily imply to cancel the path computation request if another procedure is able to read and write VSPT objects. Thus, the TLV "Error-criticality" MIGHT be used with any value depending on the expected behavior.

Error-type 15 refers to "Global Concurrent Optimization Error" defined in [\[RFC5557\]](#). [\[RFC5557\]](#) states that the corresponding global concurrent path optimization MUST be cancelled at the PCC. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 1. The TLV "Propagation" MIGHT be used with such errors.

Error-type 16 relates to "P2MP Capability Error" defined in [\[RFC6006\]](#). Such errors lead to the cancellation of the path computation request. Hence, if the TLV "Error-criticality" is included, it MUST have a value of 1. The TLV "Propagation" MIGHT be used with such errors.

Error-type 17, titled "P2MP END-POINTS Error" is defined [\[RFC6006\]](#). Such errors are thrown when a PCE tries to add or prune nodes to or from a P2MP Tree. [\[RFC6006\]](#) does not specify if such errors lead to cancel the path computation request. Hence, TLVs "Error-criticality" and "Propagation" MIGHT be used with this type of errors with any value depending on the expected behavior.

Error-type 18 on "P2MP Fragmentation Error" is described [\[RFC6006\]](#) which does not specify whether the path computation request should be cancelled. But, as messages are fragmented, it is natural to think that the PCE should wait at least a bit for further messages. The TLV "Error-criticality" MIGHT be included in such error messages and is particularly adapted to differ the semantic

of the same error-type message: if it is included with a value of 0 then the PCE will still wait for further fragmented messages, when this waiting time ends, the TLV can be included with a value of 1 in order to finally cancel the request. The TLV "Propagation" MIGHT also be used with such errors.

Among the existing normative references, only the [[RFC5440](#)] defines some notification-types and values. The recommendations with respect to the TLVs definitions provided in this document are the followings:

Notification-type=1, Notification-value=1 or 2: a PCC, respectively a PCE, cancels a set of pending requests, such a notification SHOULD be propagated to the list of PCEs which were implied in the path computation requests. Hence, the NOTIFICATION object SHOULD contains the TLV "Propagation" with value 1 and the TLV "Notification Type" with value 1, together with a DIFFUSION-LIST object containing the list of PCEs.

Notification-type=2, Notification-value=1 or 2 : indicates to the PCC that the PCE is, respectively is no longer, in an overloaded state, such a notification can be propagated or stay local. It is therefore RECOMMENDED to specify this behavior using the TLV "Propagation" and associated restriction mechanisms.

6. Error and Notification Scenarios

This section provides some examples depicting how the error and notification types described above can be used in a PCEP session. The origin of the errors or notifications is only illustrative and has no normative purpose. Sometimes the PCE features behind may be implementation-specific (e.g. detection of flooding). This section does not provide scenarios for errors with a high-level of criticality since such errors are very specific and until now have been normalized only during the session establishment (error-type of 1).

6.1. Local Error with a low level of criticality

In this example, a PCC attempts to establish a second PCEP session with the same PCE for another request. Consequently the PCE sent an error of error-type 9. This error stay local and does not affect the former session. The second session is ignored.

	+--++		+--++
	PCC		PCE
	+--++		+--++
1) Path computation event			
2) PCE Selection			
3) Path computation request X sent to the selected PCE		---- PCReq message--->	
			4) Path computation request queued
5) Path computation event			
6) PCE Selection			
7) Path computation request X' sent to the selected PCE		---- PCReq message--->	
			8) Session already opened
		<--- PCErr message----	Error-type=9

6.2. Propagated Error with a low level of criticality

In this example, a PCC sends a path computation requests with no P flag set whereas (e.g. END-POINT object with P-flag cleared). This is detected by another PCE in the sequence. The path computation request can thus be treated but the P-Flag will be ignored. Hence, this error is not critical but the source PCC should be informed of this fact. So, a PCErr message with error-type 10 ("Reception of an

invalid object"). The PCEP-ERROR object of the message contains a TLV "Propagation" at value 1 and a TLV "Error-criticality" at value 0. It is hence propagated backwardly to the source PCC.

+--++	+--+--+	+--++
PCC	PCE PCC	PCE
+--++	+--+--+	+--++
---- PCReq message-->		
	---- PCReq message--->	
		1) Parameter is
		not supported
	<--- PCErr message----	Error-type=10
<--- PCErr message---		

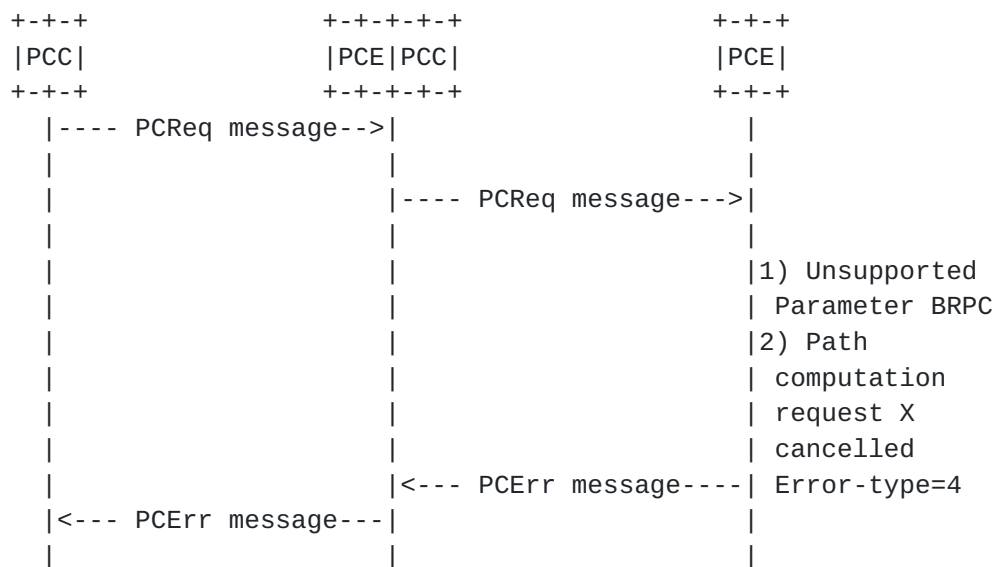
6.3. Local Error with a medium level of criticality

In this example, the PCC sends a DiffServ-aware path computation request. The PCE receiving the request does not support the indicated class-type and thus sends back a PCErr message with error-type=12, error-value=1, a TLV "Propagation" at value 0 and a TLV "Error-criticality" at value 1. Consequently, the request(s) is (are) cancelled.

	+--++	+--++
	PCC	PCE
	+--++	+--++
1) Path computation event		
2) PCE Selection		
3) Path computation request X sent to the selected PCE	---- PCReq message--->	
		4) Path computation request queued
		5) DiffServ class-type not supported
		6) Path computation request X cancelled
	<--- PCErr message----	Error-type=12

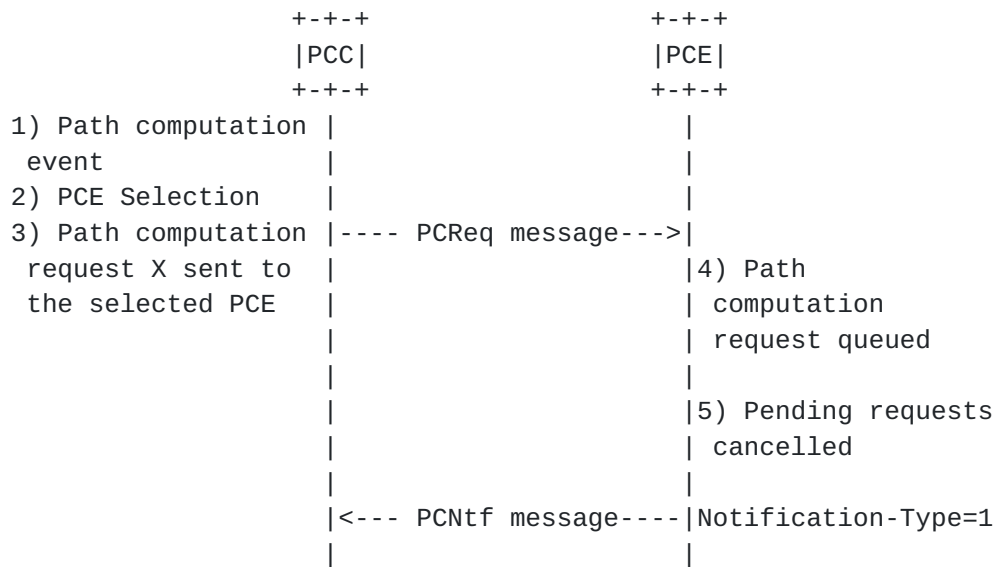
6.4. Propagated Error with a medium level of criticism

In this example, PCEs are using the BRPC procedure to treat a path computation request [RFC5441]. However, one of the PCE does not support a parameter of the request. Hence, a PCErr message with error-type 4 and error-value 4 is sent by this PCE and has to be forwarded to the source PCC. The PCEP-ERROR object includes a TLV "Propagation" at value 1 and TLV "Error-criticality" at value 1 and the message is propagated backwardly to the source PCC. Consequently, the request(s) is (are) cancelled.



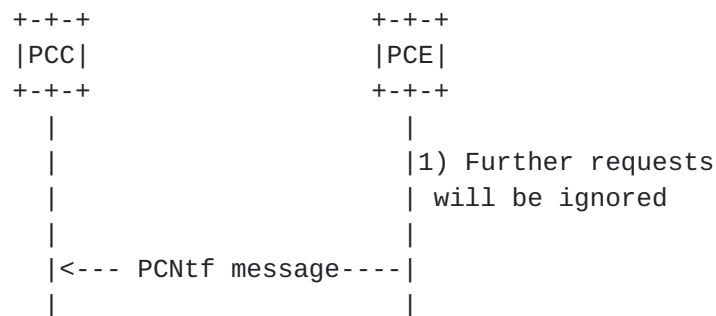
6.5. Local request-specific notification

In this example, a PCE sends a request-specific notification indicating that, a set of pending requests are cancelled (e.g. notification-type=1, notification-value=1 as described in [\[RFC5440\]](#)). Hence, a PCNtf message is sent to the PCC with a NOTIFICATION object including a TLV "Propagation" at value 0 and a TLV "Notification Type" at value 1.



6.6. Local non request-specific notification

In this example, a PCE sends a non request-specific notification indicating that, due to multiple sendings (or for other reason), further requests from this PCC will be ignored. Hence, a PCNtf message is sent to the PCC with a NOTIFICATION object including a TLV "Propagation" at value 0 and a TLV "Notification Type" at value 0.



6.7. Propagated request-specific notification

In this example, a PCE receives a request but it is temporarily congested. However, it can treat the request after few minutes which might cause some time-out in the predecessor PCEs. Hence, a PCNtf message with a NOTIFICATION object containing a TLV "Propagation" at value 1 and a TLV "Notification Type" at value 1 is send to the PCC and propagated backwardly in the PCE sequence. Such a notification could include an OVERLOAD object as described in [[RFC5886](#)].

<pre> +---+ PCC +---+ ---- PCReq message--> <--- PCNtf message--- </pre>	<pre> +---+---+ PCE PCC +---+---+ ---- PCReq message---> <--- PCNtf message---- </pre>	<pre> +---+ PCE +---+ 1)PCE is busy but will answer to X in M minutes (time-out update) Notification-type=2 </pre>
--	--	---

6.8. Propagated non request-specific notification

In this example, a PCE is temporarily congested. A PCNtf message with a NOTIFICATION object containing a TLV "Propagation" at value 1 and a TLV "Notification Type" at value 0 is send to a PCE and propagated to a sequence of PCEs. Here, PCEk is congested and send a PCNtf message to PCEi with the approapriate TLVs, an OVERLOAD object as described in [[RFC5886](#)], and a DIFFUSION-LIST object indicating PCEj as a target of the notification.

<pre> +---+---+ PCEj +---+---+ <--- PCNtf message--- </pre>	<pre> +---+---+ PCEi +---+---+ <--- PCNtf message---- </pre>	<pre> +---+---+ PCEk +---+---+ 1)PCE is busy for M minutes (time-out update) Notification-type=2 </pre>
--	---	--

7. Security Considerations

Within the introduced set of TLVs , the TLV "Propagation" affects PCEP security considerations since it forces propagation behaviors. Thus, a PCEP implementation SHOULD activate stateful mechanism when receiving PCEP-ERROR or NOTIFICATION object including this TLV in order to avoid DoS attacks.

8. IANA Considerations

IANA maintains a registry of PCEP parameters. This includes a sub-registry for PCEP Objects.

IANA is requested to make an allocation from the sub-registry as follows. The values here are suggested for use by IANA.

8.1. PCEP TLV Type Indicators

As described in [Section 5](#) the newly defined TLVs allows a PCE to enforce specific error and notification behaviors within PCEP-ERROR and NOTIFICATION objects. IANA is requested to make the following allocations from the "PCEP TLV Type Indicators" sub-registry.

Value	Description	Reference
7	Propagation	this document
8	Error-criticality	this document
9	Notification type	this document

8.2. New TTL object

TBC

8.3. New DLO object

Object-class Value	Object-Type and Name	Reference
25	1: Diffusion list object	this document

Target-Type Value	Meaning	Reference
0	Any PCEP peers	this document
1	PCEs but excludes PCC-only peers	this document
2	PCEs and PCCs with which a session is still opened	this document

Subobjects	Reference
1: IPv4 prefix	this document
2: IPv6 prefix	this document
4: Unnumbered Interface ID	this document
5: OSPF Area ID	this document
32: Autonomous system number	this document
33: Explicit Exclusion Route subobject (EXRS)	this document

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [RFC5441] Vasseur, JP., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", [RFC 5441](#), April 2009.
- [RFC5455] Sivabalan, S., Parker, J., Boutros, S., and K. Kumaki, "Diffserv-Aware Class-Type Object for the Path Computation Element Communication Protocol", [RFC 5455](#), March 2009.
- [RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", [RFC 5521](#), April 2009.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", [RFC 5541](#), June 2009.
- [RFC5557] Lee, Y., Le Roux, JL., King, D., and E. Oki, "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", [RFC 5557](#), July 2009.
- [RFC5886] Vasseur, JP., Le Roux, JL., and Y. Ikejiri, "A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture", [RFC 5886](#), June 2010.
- [RFC6006] Zhao, Q., King, D., Verhaeghe, F., Takeda, T., Ali, Z., and J. Meuric, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", [RFC 6006](#), September 2010.

9.2. Informational References

- [I-D.ietf-pce-vendor-constraints]
Farrel, A. and G. Bernstein, "Conveying Vendor-Specific Constraints in the Path Computation Element Protocol",

[draft-ietf-pce-vendor-constraints-01](#) (work in progress),
March 2010.

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

Authors' Addresses

Helia Pouyllau
Alcatel-Lucent
Route de Villejust
NOZAY 91620
FRANCE

Phone: + 33 (0)1 30 77 63 11
Email: helia.pouyllau@alcatel-lucent.com

Remi Theillaud
Marben Products
176 rue Jean Jaures
Puteaux 92800
FRANCE

Phone: + 33 (0)1 79 62 10 22
Email: remi.theillaud@marben-products.com

Julien Meuric
France Telecom Orange
2, avenue Pierre Marzin
Lannion 22307
FRANCE

Email: julien.meuric@orange.com

