

PCE Working Group
Internet-Draft
Updates: [5440](#) (if approved)
Intended status: Standards Track
Expires: 8 September 2022

H.P. Pouyllau
Alcatel-Lucent
R.T. Theillaud
Marben Products
J.M. Meuric
Orange
H. Zheng (Editor)
X. Zhang
Huawei Technologies
7 March 2022

Extensions to the Path Computation Element Communication Protocol for
Enhanced Errors and Notifications
draft-ietf-pce-enhanced-errors-11

Abstract

This document defines new error and notification TLVs for the PCE Communication Protocol (PCEP) specified in [RFC5440](#), and will update it. It identifies the possible PCEP behaviors in case of error or notification. Thus, this draft defines types of errors and how they are disclosed to other PCEs in order to support 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 <https://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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Terminology	3
2.	Conventions used in this document	3
3.	Introduction	3
3.1.	Examples	4
3.1.1.	Error use-case	4
3.1.2.	Notification use-case	4
4.	PCEP Behaviors	4
4.1.	PCEP Behaviors in Case of Error	5
4.2.	PCEP Behaviors in Case of Notification	6
4.3.	PCE Peer Identification	6
5.	PCEP Extensions for Error and Notification Handling	7
5.1.	Propagation TLV	7
5.2.	Error-criticality TLV	7
5.3.	Behaviors and TLV combinations	8
5.4.	Propagation Restrictions TLVs	9
5.4.1.	Time-To-Live (TTL) TLV	9
5.4.2.	DIFFUSION-LIST TLV	9
5.4.3.	Rules Applied to Existing Errors and Notifications	11
6.	Error Handling Guidelines for Future PCEP Extension	15
7.	Backward Compatibility Consideration	15
8.	Implementation Status	15
9.	Security Considerations	16
10.	IANA Considerations	16
10.1.	PCEP TLV Type Indicators	16
10.2.	New DIFFUSION-LIST TLV	17
11.	References	17
11.1.	Normative References	17
11.2.	Informational References	19
Appendix A.	Error and Notification Scenarios	20
A.1.	Error Behavior Type 1	20
A.2.	Error Behavior Type 2	21
A.3.	Error Behavior Type 4	21

A.4. Error Behavior Type 5	22
Authors' Addresses	22

[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, for a given path computation query, initiating the first PCEP request, which may then 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 [[RFC7470](#)]. 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. This document updates [[RFC5440](#)] by introducing mechanism to propagate the error message, with specifying error and notification TLVs.

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. In order to generalize PCEP behaviors in the case of heterogeneous PCE systems, new objects have to be defined. Dealing with heterogeneity is a major challenge considering PCE applicability, particularly in multi-layer, multi-domain and H-PCE contexts [[RFC8751](#)]. Thus, extending such error codes and PCEP behaviors accordingly would improve interoperability among different PCEP implementations and would solve some of these 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 identify and specify new optional TLVs and objects 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 existing error or notification types.

[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.

4. PCEP Behaviors

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 two 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.

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](#)]. [[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. Detailed behavior is mostly let free to any specific implementation. The following sections identifies the PCEP behaviors in case of error or notification and

also introduce the requirement of PCEP peer identification in both cases.

[4.1.](#) 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 PCEP peers;
- * "Criticality level": in different RFCs, error-types affects the state of the PCEP request or session in different manners; 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 PCEP 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 sake of completeness, they have been included in this document.

[4.2.](#) 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 types of behaviors can be identified:

- * "Local": the notification does not imply any forward or backward propagation of the message;
- * "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;
- * "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.

[4.3.](#) PCE Peer Identification

The propagation of errors and notifications affects the state of the PCEP peers along the chain. In some cases, for instance a notification that a PCE is overloaded, the identification of the PCEP peer – or that the sender PCE is not the direct neighbor – might be an important information for the PCEP peers receiving the message. The ID of sender PCE is not carried in the error TLVs, but can be achieved via the speaker entity ID TLV during state synchronization. An example can be found in [[RFC8232](#)].

[5.](#) PCEP Extensions for Error and Notification Handling

This section describes extensions to support error and notification with respect to the PCEP behavior description defined in [Section 4](#). This document does not intend to modify errors and notification types previously defined in existing documents (e.g. [[RFC5440](#)], [[RFC5441](#)], etc.). Error related TLVs have been specified in this section, while the notification functionality can be achieved via using PCNtf

message with RP object with no need to extend further notification type.

[5.1.](#) Propagation TLV

To support the propagation behavior mentioned in [Section 4.1](#) and [Section 4.2](#), a new optional TLV is defined, which can be carried in PCEP-ERROR and NOTIFICATION objects, to indicate whether a 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 10](#).

The description is "Propagation", the length value is 2 bytes and the value field is 1 byte. The value field is set to 0 meaning that the message MUST NOT be propagated. If the value field is set to 1, the message MUST be propagated. [Section 5.4](#) specifies the destination and to limit the number of messages.

[5.2.](#) Error-criticality TLV

To support the shutdown behavior mentioned in [Section 4.1](#), we extend the PCEP-ERROR object by creating a new optional TLV to indicate whether an 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 10](#).

The description is "Error-criticality", the length value is 2 bytes and the value field is 1 byte. The value field is set to 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 identifiers appear in the same message 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 for this PCEP session is closed by the sender PCE peer.

[5.3.](#) Behaviors and TLV combinations

The propagation behavior MAY 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.

For a PCErr message, all the possible behaviors described in [Section 4.1](#) can be covered with TLVs included in a PCEP-ERROR object. The following table captures all combinations of error behaviors:

Error criticality value	\Propagation Value	0 (No Propagation)	1 (Propagation Required)
0 (low)		Type 1	Type 4
1 (medium)		Type 2	Type 5
2 (high)		Type 3	Type 6

- * "Error Behavior Type 1" : Local Error with a low level of criticality;
- * "Error Behavior Type 2": Local Error with a medium level of criticality;
- * "Error Behavior Type 3": Local Error with a high level of criticality;
- * "Error Behavior Type 4": Propagated Error with a low level of criticality;
- * "Error Behavior Type 5": Propagated Error with a medium level of criticality;
- * "Error Behavior Type 6": Propagated Error with a high level of criticality;

[5.4.](#) Propagation Restrictions TLVs

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

A Time-To-Live(TTL) TLV: to limit the number of PCEP peers that will recursively receive the message;

A DIFFUSION-LIST TLV: to specify the PCEP peer addresses or domains of PCEP peers the message must be propagate to;

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

Such mechanisms SHOULD be used jointly or independently depending the error or notification behaviors they are associated to. The conditions of use for the TTL and DIFFUSION-LIST TLVs are described in sections below.

[5.4.1.](#) Time-To-Live (TTL) TLV

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

If the message to be propagated is request-specific and there is no TTL or DIFFUSION-LIST TLVs included, the message MUST reach the source PCC (or alternatively the target PCE).

[5.4.2.](#) DIFFUSION-LIST TLV

The DIFFUSION-LIST TLV can be carried within either the error object of a PCErr message, or the notification object of a PCNtf message. It can either be used in a message sent by a PCC to a PCE or vice versa. The DIFFUSION-LIST MAY be used with propagated errors (TLV "Propagation" at value 1 in PCEP-ERROR object).

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

PCEP peer receiving a PCErr or PCNTf message containing a PCEP-ERROR or a NOTIFICATION object with a TLV "Propagation" at value 1 and where a DIFFUSION-LIST appears, MUST remove the addresses of the PCEP peers from the DIFFUSION-LIST, before sending the message to any other PCEP peers. This is performed by adding the PCEP peer addresses to the Explicit eXclusion Route Sub-object of the DIFFUSION-LIST. If a DIFFUSION-LIST value is empty, the PCEP peer MUST NOT propagate the message to any peer.

Note that, a Diffusion-List 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 its own address to the Explicit eXclusion Route Sub-object (EXRS) of the Diffusion-List for any forwarded messages. Hence, a PCE SHOULD avoid forwarding the same message repeated 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).

5.4.3. Rules Applied to Existing Errors and Notifications

Many existing normative references states on error definitions (see for instance [[RFC5440](#)], [[RFC5441](#)], [[RFC5455](#)], [[RFC5521](#)], [[RFC5557](#)], [[RFC5886](#)], [[RFC8231](#)], [[RFC8232](#)], [[RFC8253](#)], [[RFC8281](#)], [[RFC8306](#)], [[RFC8408](#)], [[RFC8697](#)]). This section provides processing rules for existing error types handling, as a recommendation. 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 and not propagated. Hence, if a "Propagation" TLV is added to the error message it is recommended to be set to value 0. Error-values 1,2,6,7 have a high level of criticality. Hence, if the "Error-criticality" TLV is included within a PCErr message of type 1 and value 1,2,6 or 7, it is recommended to 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 MAY be propagated using the TLV "Propagation". Their level of criticality is defined as leading to cancel the path computation request [[RFC5440](#)]. Hence, if the "Error-criticality" TLV is included, it usually have a value of 1. The error-value 4 of error-type 4 ("Unsupported parameter") associated to the BRPC procedure [[RFC5441](#)] is suggested to contain the "Propagation" TLV with a DIFFUSION-LIST 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 [[RFC8306](#)]. In [[RFC5440](#)], it is specified that the path computation request MUST be cancelled when an error of type 5 occurs. Hence, if the "Error-criticality" TLV is included, it usually have a value of 1. As such errors might be conveyed to several PCEs, the "Propagation" TLV MAY be used.

Error-type 6 described as "Mandatory object missing" in [[RFC5440](#)], leads to the cancellation of the path computation request. Hence, if the "Error-criticality" TLV is included, it usually have a value of 1. The "Propagation" TLV MAY be used with such errors. The error-value of 4 for Monitoring object missing defined in [[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 reffered synchronized path computation request MUST be cancelled when an error of type 5 occurs. Hence, if the "Error-criticality" TLV is included, it usually have a value of 1. The "Propagation" TLV MAY be used with such errors.

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

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

type 9, it is usually set to a value of 0.

Error-type 10 which refers to the reception of an invalid object as described in [[RFC5440](#)] no indication is provided on the cancellation of the path computation request. Hence, if the "Error-criticality" TLV is included, it usually have a value of 0. The "Propagation" TLV MAY 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 "Error-criticality" TLV is included, it usually have a value of 0. The "Propagation" TLV MAY 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 "Propagation" TLV and a DIFFUSION-LIST 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 "Error-criticality" TLV MAY 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 "Error-criticality" TLV is included, it usually have a value of 1. The "Propagation" TLV MAY be used with such errors.

Error-type 16 relates to "P2MP Capability Error" defined in [\[RFC8306\]](#). Such errors lead to the cancellation of the path computation request. Hence, if the "Error-criticality" TLV is included, it usually have a value of 1. The "Propagation" TLV MAY be used with such errors.

Error-type 17, titled "P2MP END-POINTS Error" is defined [\[RFC8306\]](#). Such errors are thrown when a PCE tries to add or prune nodes to or from a P2MP Tree. [\[RFC8306\]](#) does not specify if such errors lead to cancel the path computation request. Hence, the "Error-criticality" and "Propagation" TLVs MAY be used with this type of error with any value depending on the expected behavior.

Error-type 18 of "P2MP Fragmentation Error" is described [\[RFC8306\]](#) 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 "Error-criticality" TLV MAY 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 "Propagation" TLV MAY also be used with such errors.

Error-type 19 of "Invalid Operation" is described in [\[RFC8231\]](#) and [\[RFC8281\]](#), which implies a wrong capability description for PCEP session. In this case, the PCERR message MUST be returned to PCC, and this message usually contain a "Propagation" TLV and a

DIFFUSION-LIST with a Target-Type of 0 or 2. The "Error-criticality" TLV is recommended be set to 2 in order to guarantee the termination of PCEP session.

Error-type 20 of "LSP State Synchronization Error" is described in [\[RFC8231\]](#) and [\[RFC8232\]](#), which cannot successfully sync up the LSP states. In this case, the "Error-criticality" TLV should be set to 2 in order to guarantee the termination of PCEP session. The "Propagation" TLV MAY also be used with such errors.

Error-type 21 of "Invalid traffic engineering path setup type" is

described in [[RFC8408](#)] . Such errors failed to find a matched path setup type and the PCEP sessions MUST be closed. In this case, the "Error-criticality" TLV is usually set to 2 in order to guarantee the termination of PCEP session. The "Propagation" TLV MAY also be used with such errors.

Error-type 23 of "Bad parameter value" is described in [[RFC8281](#)] . Such errors occur when there is a conflict in path name of C flag not set for PCE initiation. In this case, the "Error-criticality" TLV may be set to either 0 or 1 to indicate whether the request is still valid, with the PCEP session open. The "Propagation" TLV MAY also be used with such errors.

Error-type 24 of "LSP instantiation error" is described in [[RFC8281](#)] . Such errors occur when PCC detects problems when establishing the path, the message MUST relay to the PCE, therefore the "Propagation" TLV is usually contained. The "Error-criticality" TLV may be set to either 0 or 1 to indicate whether the request is still valid, with the PCEP session open.

Error-type 25 of "PCEP StartTLS failure" is described in [[RFC8253](#)]. Such errors indicate the security issue in transport layer. In this case, the "Error-criticality" TLV is usually set to 2 in order to close the PCEP session. The "Propagation" TLV MAY also be used with such errors, depending on the detailed security conditions.

Error-type 26 of "Association Error " is described in [[RFC8697](#)] . Such errors occur when there is problem for LSP association. In this case, the "Error-criticality" TLV should be set to either 0 or 1 to indicate whether the request is still valid, with the PCEP session open. The "Propagation" TLV MAY also be used with such errors.

[6.](#) Error Handling Guidelines for Future PCEP Extension

Error and Notification handling in this document should be considered in PCE documents that include new errors and notifications. A

requirement for the authors of these drafts is to evaluate the applicability of the procedure in this document and provide details about the "Error-criticality" TLV and "Propagation" TLV for errors and notifications defined in the draft. Example text is provided as follow.

Error-type XX (fill in value of the Error-type) of " XXXX " (fill in name of the Error-type) is described in [RFCYYYY] (fill in the document reference of the Error-type). Such errors occur when ZZZZ (fill in typical scenario). In this case, the "Error-criticality" TLV should be set to X (fill in the recommended value) to indicate whether the request is still valid, with the PCEP session open. The error messages SHOULD/MAY (select the mandatory level) contain a "Propagation" TLV and a DIFFUSION-LIST with a Target-Type of A(fill in the recommended value).

7. Backward Compatibility Consideration

There would be backward compatibility issue if there are multiple PCEs with different level understanding of error message. In a scenario that PCE(i) propagate the error message to PCE (i+1), it is possible that PCE (i+1) is not capable to extract the message correctly, then such error message would be ignored and not be further propagated.

There can be potential approach to avoid these problem, such as recognizing the incapable PCE and avoiding propagation. However, these approach is not in the scope of this document.

8. Implementation Status

[NOTE TO RFC EDITOR : This whole section and the reference to [\[RFC7942\]](#) is to be removed before publication as an RFC]

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [\[RFC7942\]](#). The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [\[RFC7942\]](#), "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

At the time of posting the -08 version of this document, there are no known implementations of this mechanism. It is believed that two vendors are considering prototype implementations, but these plans are too vague to make any further assertions.

[9.](#) Security Considerations

Within the introduced set of TLVs, the "Propagation" TLV 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.

[10.](#) 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.

[10.1.](#) PCEP TLV Type Indicators

As described in [Section 5.4](#) 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.

Internet-Draft Extensions to PCEP for Enhanced errors March 2022

Value	Description	Reference
TBD	Propagation	this document
TBD	Error-criticality	this document

[10.2.](#) New DIFFUSION-LIST TLV

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: 4-byte AS number	this document
6 OSPF area ID	this document
7 IS-IS Area ID	this document
32: Autonomous system number	this document
33: Explicit Exclusion Route subobject (EXRS)	this document

[11.](#) References

[11.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#),

- [RFC5441] Vasseur, JP., Ed., 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](#), DOI 10.17487/RFC5441, April 2009, <<https://www.rfc-editor.org/info/rfc5441>>.
- [RFC5455] Sivabalan, S., Ed., Parker, J., Boutros, S., and K. Kumaki, "Diffserv-Aware Class-Type Object for the Path Computation Element Communication Protocol", [RFC 5455](#), DOI 10.17487/RFC5455, March 2009, <<https://www.rfc-editor.org/info/rfc5455>>.
- [RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", [RFC 5521](#), DOI 10.17487/RFC5521, April 2009, <<https://www.rfc-editor.org/info/rfc5521>>.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", [RFC 5541](#), DOI 10.17487/RFC5541, June 2009, <<https://www.rfc-editor.org/info/rfc5541>>.
- [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](#), DOI 10.17487/RFC5557, July 2009, <<https://www.rfc-editor.org/info/rfc5557>>.
- [RFC5886] Vasseur, JP., Ed., Le Roux, JL., and Y. Ikejiri, "A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture", [RFC 5886](#), DOI 10.17487/RFC5886, June 2010,

<<https://www.rfc-editor.org/info/rfc5886>>.

- [RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", [RFC 8231](#), DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.
- [RFC8232] Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", [RFC 8232](#), DOI 10.17487/RFC8232, September 2017, <<https://www.rfc-editor.org/info/rfc8232>>.

Pouyllau, et al.

Expires 8 September 2022

[Page 18]

Internet-Draft Extensions to PCEP for Enhanced errors

March 2022

- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", [RFC 8253](#), DOI 10.17487/RFC8253, October 2017, <<https://www.rfc-editor.org/info/rfc8253>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", [RFC 8281](#), DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.
- [RFC8306] Zhao, Q., Dhody, D., Ed., Palleti, R., and D. King, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", [RFC 8306](#), DOI 10.17487/RFC8306, November 2017, <<https://www.rfc-editor.org/info/rfc8306>>.
- [RFC8408] Sivabalan, S., Tantsura, J., Minei, I., Varga, R., and J. Hardwick, "Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages", [RFC 8408](#), DOI 10.17487/RFC8408, July 2018, <<https://www.rfc-editor.org/info/rfc8408>>.
- [RFC8697] Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "Path Computation Element

Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)", [RFC 8697](#), DOI 10.17487/RFC8697, January 2020, <<https://www.rfc-editor.org/info/rfc8697>>.

11.2. Informational References

- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.
- [RFC7470] Zhang, F. and A. Farrel, "Conveying Vendor-Specific Constraints in the Path Computation Element Communication Protocol", [RFC 7470](#), DOI 10.17487/RFC7470, March 2015, <<https://www.rfc-editor.org/info/rfc7470>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [BCP 205](#), [RFC 7942](#), DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.

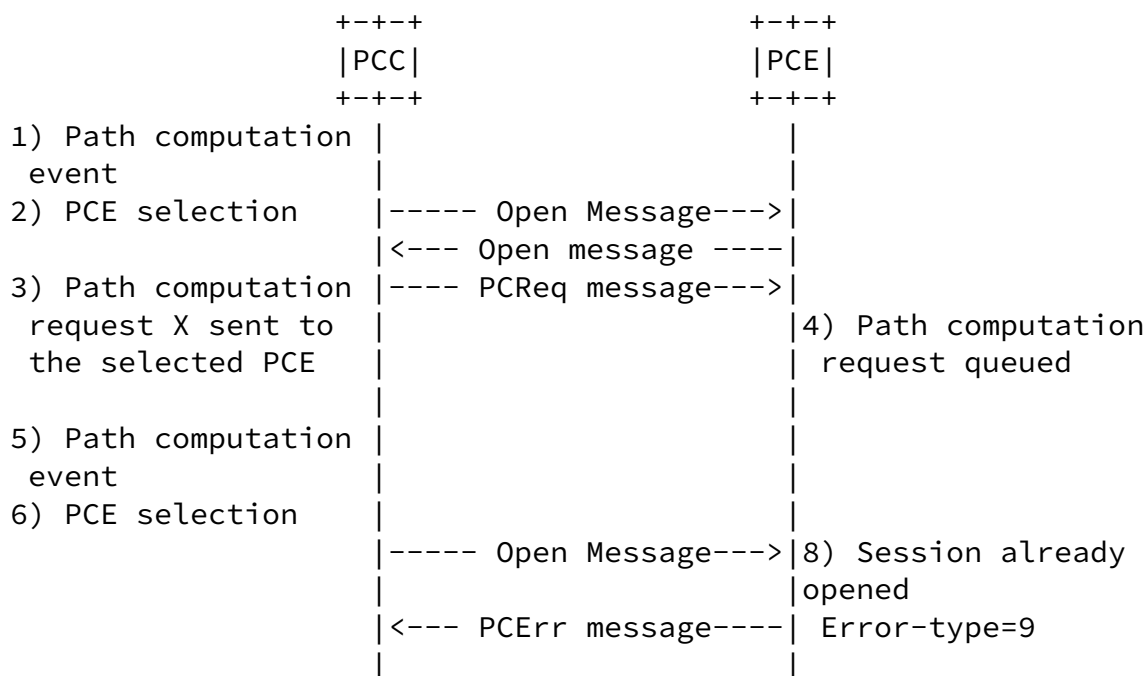
- [RFC8751] Dhody, D., Lee, Y., Ceccarelli, D., Shin, J., and D. King, "Hierarchical Stateful Path Computation Element (PCE)", [RFC 8751](#), DOI 10.17487/RFC8751, March 2020, <<https://www.rfc-editor.org/info/rfc8751>>.

Appendix A. Error and Notification Scenarios

This section provides some examples depicting how the error 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 (i.e., Error behaviors 3 and 6) since such errors are very specific and until now have been normalized only during the session establishment (error-type of 1).

A.1. Error Behavior Type 1

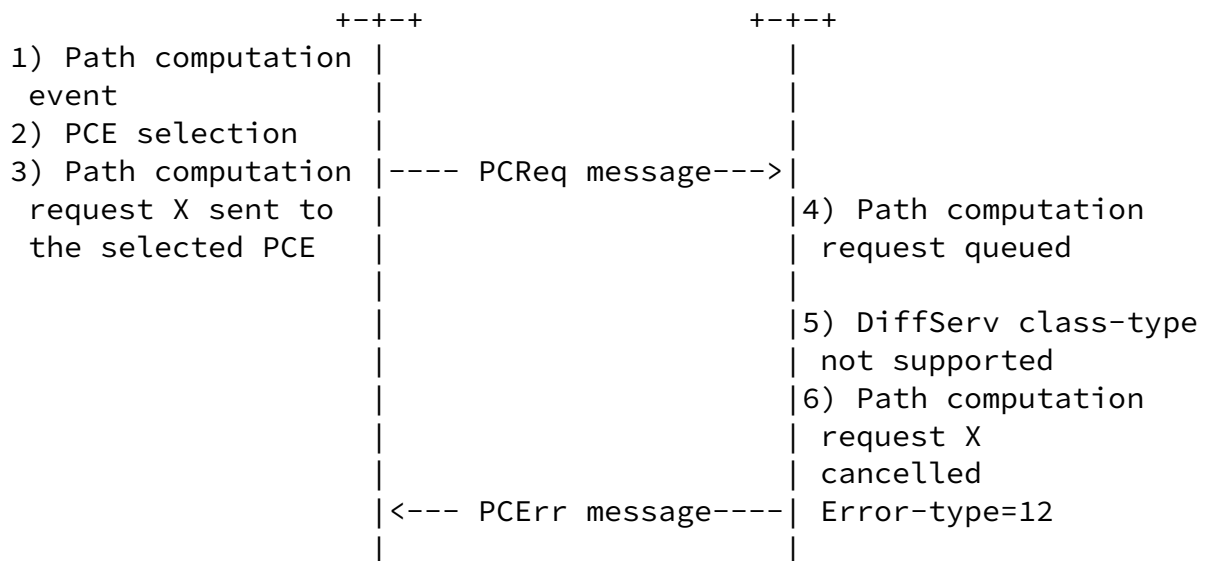
In this example, a PCC attempts to establish a second PCEP session with the same PCE for another request. Consequently the PCE sends back an error message with error-type 9. This error stays local and does not affect the former session. The second session is ignored. If the "Propagation" TLV and "Error-criticality" TLV are used, they should be both set to value 0.



[A.2.](#) Error Behavior Type 2


In this example, the PCC sends a DiffServ-aware path computation request. If the PCE receiving the request does not support the indicated class-type, it thus sends back a PCErr message with error-type=12 and error-value=1. If the "Propagation" TLV and "Error-criticality" TLV are present, they should carry value 0 and value 1 respectively. Consequently, the request is cancelled.

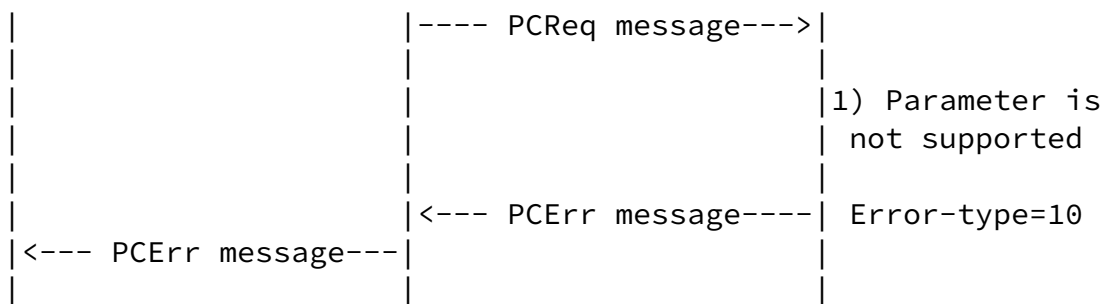




A.3. Error Behavior Type 4

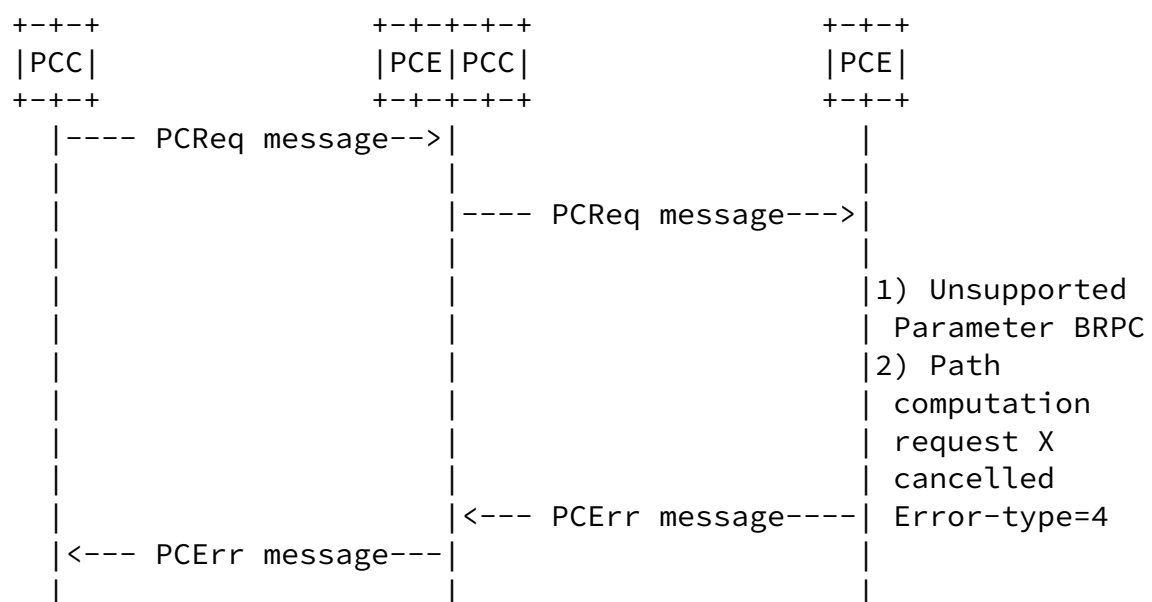
In this example, a PCC sends a path computation requests with no P flag set (e.g. END-POINT object with P-flag cleared). This is detected by another PCE in the sequence. The path computation request can thus be treated but the P-Flag will be ignored. Hence, this error is not critical but the source PCC 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 "Propagation" TLV at value 1 and a "Error-criticality" TLV at value 0. It is hence propagated backwardly to the source PCC.





A.4. Error Behavior Type 5

In this example, PCEs are using the BRPC procedure to treat a path computation request [RFC5441]. However, one of the PCEs 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 "Propagation" TLV at value 1 and "Error-criticality" TLV at value 1 and the message is propagated backwardly to the source PCC. Consequently, the request is cancelled.



Authors' Addresses

Helia Pouyllau
Alcatel-Lucent
Route de Villejust
91620 NOZAY
France
Phone: + 33 (0)1 30 77 63 11
Email: helia.pouyllau@alcatel-lucent.com

Remi Theillaud
Marben Products
176 rue Jean Jaures
92800 Puteaux
France
Phone: + 33 (0)1 79 62 10 22
Email: remi.theillaud@marben-products.com

Julien Meuric
Orange
2, avenue Pierre Marzin
22307 Lannion
France
Email: julien.meuric@orange.com

Haomian Zheng (Editor)
Huawei Technologies
H1, Xiliu Beipo Village, Songshan Lake,
Dongguan
Guangdong, 523808
China
Email: zhenghaomian@huawei.com

Xian Zhang
Huawei Technologies
A10, Huawei Industrial Base, Bantian, Longgang District
Shenzhen
Guangdong, 518129
China
Email: zhang.xian@huawei.com

