

PCE Working Group
Internet-Draft
Updates: [8231](#) (if approved)
Intended status: Standards Track
Expires: November 6, 2021

C. Li
H. Zheng
Huawei Technologies
S. Litkowski
Cisco
May 5, 2021

Extension for Stateful PCE to allow Optional Processing of PCEP Objects [draft-dhody-pce-stateful-pce-optional-08](#)

Abstract

This document introduces a mechanism to mark some of the Path Computation Element (PCE) Communication Protocol (PCEP) objects as optional during PCEP messages exchange for the Stateful PCE model to allow relaxing some constraints. This document introduces this relaxation and updates [RFC 8231](#).

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 November 6, 2021.

Copyright Notice

Copyright (c) 2021 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 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.	Introduction	2
1.1.	Requirements Language	3
2.	Overview	3
2.1.	Usage Example	4
3.	PCEP Extension	4
3.1.	STATEFUL-PCE-CAPABILITY TLV	4
3.2.	Handling of P flag	5
3.2.1.	The PCRpt Message	5
3.2.2.	The PCUpd Message and the PCInitiate Message	5
3.3.	Handling of I flag	5
3.3.1.	The PCUpd Message	5
3.3.2.	The PCRpt Message	6
3.3.3.	The PCInitiate Message	6
3.4.	Delegation	6
3.5.	Unknown Object Handling	6
4.	Security Considerations	7
5.	IANA Considerations	7
5.1.	STATEFUL-PCE-CAPABILITY TLV	7
6.	Manageability Considerations	7
6.1.	Control of Function and Policy	7
6.2.	Information and Data Models	7
6.3.	Liveness Detection and Monitoring	8
6.4.	Verify Correct Operations	8
6.5.	Requirements On Other Protocols	8
6.6.	Impact On Network Operations	8
7.	Acknowledgments	8
8.	References	8
8.1.	Normative References	8
8.2.	Informative References	9
Appendix A.	Contributors	11
	Authors' Addresses	11

[1.](#) Introduction

[RFC5440] describes the Path Computation Element Communication Protocol (PCEP) which enables the communication between a Path Computation Client (PCC) and a Path Control Element (PCE), or between two PCEs based on the PCE architecture [RFC4655].

PCEP Extensions for Stateful PCE Model [RFC8231] describes a set of extensions to PCEP to enable active control of Multiprotocol Label Switching Traffic Engineering (MPLS-TE) and Generalized MPLS (GMPLS) tunnels. [RFC8281] describes the setup and teardown of PCE-initiated

LSPs under the active stateful PCE model, without the need for local configuration on the PCC, thus allowing for dynamic control.

[RFC5440] defined the P flag (Processing-Rule) in the Common Object Header to allow a PCC to specify in a Path Computation Request (PCReq) message (sent to a PCE) whether the object must be taken into account by the PCE during path computation or is optional. The I flag (Ignore) is used by the PCE in a Path Computation Reply (PCRep) message to indicate to a PCC whether or not an optional object was considered by the PCE during path computation. Stateful PCE [RFC8231] specified that the P and I flags of the PCEP objects defined in [RFC8231] is to be set to zero on transmission and ignored on receipt, since they are exclusively related to path computation requests. The behavior for P and I flag in other messages defined in [RFC5440] and other extension was not specified. This document clarifies how the P and I flag could be used in the stateful PCE model to identify optional objects in the Path Computation State Report (PCRpt) [RFC8231], the Path Computation Update Request (PCUpd) [RFC8231], and the LSP Initiate Request (PCInitiate) [RFC8281] message.

This document updates [RFC8231] with respect to usage of the P and I flag as well as the handling of unknown objects in the stateful PCEP message exchange.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Overview

[RFC5440] describes the handling of unknown objects as per the setting of the P flag for the PCReq message. Further, [RFC8231] defined the usage of the LSP Error Code TLV in the PCRpt message in response to failed LSP Update Request via the PCUpd message (for example, due to an unsupported object/TLV).

This document clarifies the procedure of marking some objects as 'optional to be processed' by the PCEP peer in the stateful PCEP messages. Furthermore, this document updates the procedure for handling unknown objects in the stateful PCEP messages based on the P flag.

2.1. Usage Example

The PCRpt message is used to report the current state of an LSP. As part of the message both the <intended-attribute-list> and <actual-attribute-list> is encoded (see [RFC8231]). For example, the <intended-attribute-list> could include the METRIC object to indicate a limiting constraint (B flag set) for the Path Delay Variation metric [RFC8233]. In some scenarios, it would be useful to state that this limiting constraint can be relaxed by the PCE in case it cannot find a path. Similarly in the case of an association group [RFC8697] such as Disjoint Association [RFC8800], the PCE may need to completely relax the disjointness constraint in order to provide a path to all the LSPs that are part of the association. In these case it would be useful to mark the objects as 'optional' and it could be ignored by the PCEP peer. Also, it would be useful for the PCEP speaker to learn if the PCEP peer has relaxed the constraint and ignored the processing of the PCEP object.

Thus, this document simply clarifies, how the already existing P and I flag in the PCEP common object header could be used during the stateful PCEP message exchange.

3. PCEP Extension

3.1. STATEFUL-PCE-CAPABILITY TLV

A PCEP speaker indicates its ability to support the handling of the P and I flag in the stateful PCEP message exchange during the PCEP initialization phase, as follows. When the PCEP session is established, a PCC sends an Open message with an OPEN object that contains the STATEFUL-PCE-CAPABILITY TLV, as defined in [RFC8231]. A new flag, the R (RELAX) flag, is added in this TLV to indicate the support for relaxing the processing of some objects via the use of the P and I flag in the PCEP common object header.

R (RELAX bit - TBD1): If set to 1 by a PCEP Speaker, the R flag indicates that the PCEP Speaker is willing to send and receive PCEP objects with the P and I flags in the PCEP common object header for the stateful PCE messages. In case the bit is unset, it indicates that the PCEP Speaker would not handle the P and I flags in the PCEP common object header for stateful PCE messages.

The R flag MUST be set by both a PCC and a PCE to indicate support for the handling of the P and I flag in the PCEP common object header to allow relaxing some constraints by marking objects as optional to process. If the PCEP speaker did not set the R flag but receives PCEP objects with P or I bit set, it MUST behave as per the processing rule in [RFC8231] i.e., the bits are simply ignored.

3.2. Handling of P flag

3.2.1. The PCRpt Message

The P flag in the PCRpt message [[RFC8231](#)] allows a PCC to specify to a PCE whether the object must be taken into account by the PCE (during path computation, re-optimization, or state maintenance) or is optional o process. When the P flag is set in the PCRpt message received on a PCEP session on which R bit was set by both peers, the object MUST be taken into account by the PCE. Conversely, when the P flag is cleared, the object is optional and the PCE is free to ignore it. The P flag for the mandatory objects such as the LSP and the ERO object (intended path) MUST be set in the PCRpt message. If a mandatory object is received with the P flag set incorrectly according to the rules stated above, the receiving peer MUST send a PCErr message with Error-Type=10 (Reception of an invalid object) and Error-value=1 (reception of an object with P flag not set). On a PCEP session on which R bit was set by both peers, the PCC SHOULD set the P flag by default, unless a local configuration or local policy indicates that some constraints (corresponding PCEP objects) can be marked as optional and could be ignored by the PCE.

3.2.2. The PCUpd Message and the PCInitiate Message

The P flag in the PCUpd message [[RFC8231](#)] and the PCInitiate message [[RFC8281](#)] allows a PCE to specify to a PCC whether the object must be taken into account by the PCC (during path setup) or is optional to process. When the P flag is set in the PCUpd/PCInitiate message received on a PCEP session on which R bit was set by both peers, the object MUST be taken into account by the PCC. Conversely, when the P flag is cleared, the object is optional and the PCC is free to ignore it. The P flag for the mandatory objects such as the SRP, the LSP and the ERO MUST be set in the PCUpd/PCInitiate message. If a mandatory object is received with the P flag set incorrectly according to the rules stated above, the receiving peer MUST send a PCErr message with Error-Type=10 (Reception of an invalid object) and Error-value=1 (reception of an object with P flag not set). By default, the PCE SHOULD set the P flag, unless a local configuration or local policy indicates that some constraints (corresponding PCEP objects) can be marked as optional and could be ignored by the PCC.

3.3. Handling of I flag

3.3.1. The PCUpd Message

The I flag in the PCUpd message [[RFC8231](#)] allows a PCE to indicate to a PCC whether or not an optional object was processed. The PCE MAY include the ignored optional object in its update request and set the

I flag to indicate that the optional object was ignored. When the I flag is cleared, the PCE indicates that the optional object was processed.

3.3.2. The PCRpt Message

The I flag in the PCRpt message [[RFC8231](#)] allows a PCC to indicate to a PCE whether or not an optional object was processed in response to an LSP Update Request (PCUpd) or LSP Initiate Request (PCInitiate). The PCC MAY include the ignored optional object in its report and set the I flag to indicate that the optional object was ignored at PCC. When the I flag is cleared, the PCC indicates that the optional object was processed. The I flag has no meaning if the PCRpt message is not in response to a PCUpd or PCInitiate message (i.e. without the SRP object in the PCRpt message).

3.3.3. The PCInitiate Message

The I flag has no meaning in the PCinitiate message [[RFC8281](#)] and is ignored.

3.4. Delegation

Delegation is an operation to grant a PCE temporary rights to modify a subset of parameters on one or more LSPs by a PCC as described in [[RFC8051](#)]. Note that for the delegated LSPs, the PCE can update and mark some objects as ignored even when the PCC had set the P flag during delegation. Similarly, the PCE can update and mark some object as a must to process even when the PCC had not set the P flag during delegation.

The PCC MUST acknowledge this by sending the PCRpt message with the P flag set as per the PCE expectation for the corresponding object. In case PCC cannot accept this, it would react as per the processing rules of unacceptable update in [[RFC8231](#)].

3.5. Unknown Object Handling

This document updates the handling of unknown objects in the stateful PCEP messages as per the setting of the P flag in the common object header in a similar way as [[RFC5440](#)], i.e. if a PCEP speaker does not understand an object with the P flag set or understands the object but decides to ignore the object, the entire stateful PCEP message MUST be rejected and the PCE MUST send a PCErr message with Error-Type="Unknown Object" or "Not supported Object" [[RFC5440](#)]. In case the P flag is not set, the PCEP speaker is free to ignore the object and continue with message processing as defined.

[RFC8231] defined LSP Error Code TLV to be carried in PCRpt message in the LSP object to convey error information. This document does not change that procedure.

4. Security Considerations

This document clarifies how the already existing P and I flag in PCEP common object header could be used during stateful PCEP exchanges. It updates the unknown object error handling in stateful PCEP message exchange. These changes on their own do not add any new security concerns. The security considerations identified in [RFC5440], [RFC8231], and [RFC8281] continue to apply.

As per [RFC8231], it is RECOMMENDED that these PCEP extensions only be activated on authenticated and encrypted sessions across PCEs and PCCs belonging to the same administrative authority, using Transport Layer Security (TLS) [RFC8253] as per the recommendations and best current practices in [RFC7525] (unless explicitly set aside in [RFC8253]).

5. IANA Considerations

5.1. STATEFUL-PCE-CAPABILITY TLV

[RFC8231] defines the STATEFUL-PCE-CAPABILITY TLV; per that RFC, IANA created a "STATEFUL-PCE-CAPABILITY TLV Flag Field" subregistry to manage the value of the STATEFUL-PCE-CAPABILITY TLV's Flag field. IANA is requested to allocate a new bit in the subregistry, as follows:

Bit	Description	Reference

TBD1	RELAX bit	[This-I.D.]

6. Manageability Considerations

6.1. Control of Function and Policy

An operator MUST be allowed to configure the capability to support relaxation of constraints in the stateful PCEP message exchange. They SHOULD also allow configuration of related LSP constraints (or parameters) that are optional to process.

6.2. Information and Data Models

An implementation SHOULD allow the operator to view the capability defined in this document. To serve this purpose, the PCEP YANG module [I-D.ietf-pce-pcep-yang] could be extended in the future.

6.3. Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [[RFC5440](#)].

6.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [[RFC5440](#)].

6.5. Requirements On Other Protocols

Mechanisms defined in this document do not imply any new requirements on other protocols.

6.6. Impact On Network Operations

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [[RFC5440](#)].

7. Acknowledgments

Thanks to Jonathan Hardwick for discussion and suggestions around this draft.

8. References

8.1. 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](#), DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

- [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>>.

8.2. Informative References

- [I-D.ietf-pce-pcep-yang] Dhody, D., Hardwick, J., Beeram, V. P., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", [draft-ietf-pce-pcep-yang-16](#) (work in progress), February 2021.
- [RFC4655] Farrel, A., Vasseur, J., 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>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", [BCP 195](#), [RFC 7525](#), DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", [RFC 8051](#), DOI 10.17487/RFC8051, January 2017, <<https://www.rfc-editor.org/info/rfc8051>>.
- [RFC8233] Dhody, D., Wu, Q., Manral, V., Ali, Z., and K. Kumaki, "Extensions to the Path Computation Element Communication Protocol (PCEP) to Compute Service-Aware Label Switched Paths (LSPs)", [RFC 8233](#), DOI 10.17487/RFC8233, September 2017, <<https://www.rfc-editor.org/info/rfc8233>>.
- [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>>.

- [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>>.
- [RFC8800] Litkowski, S., Sivabalan, S., Barth, C., and M. Negi, "Path Computation Element Communication Protocol (PCEP) Extension for Label Switched Path (LSP) Diversity Constraint Signaling", [RFC 8800](#), DOI 10.17487/RFC8800, July 2020, <<https://www.rfc-editor.org/info/rfc8800>>.

Appendix A. Contributors

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: dhruv.ietf@gmail.com

Authors' Addresses

Cheng Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

EMail: c.l@huawei.com

Haomian Zheng
Huawei Technologies
H1, Huawei Xiliu Beipo Village, Songshan Lake
Dongguan, Guangdong 523808
China

EMail: zhenghaomian@huawei.com

Stephane Litkowski
Cisco

EMail: slitkows.ietf@gmail.com

