

Workgroup: PCE Working Group

Internet-Draft:

draft-sidor-pce-circuit-style-pcep-extensions-05

Published: 1 December 2023

Intended Status: Standards Track

Expires: 3 June 2024

Authors: S. Sidor P. Maheshwari A. Stone
 Cisco Systems, Inc. Airtel India Nokia
 L. Jalil S. Peng
 Verizon Huawei Technologies

PCEP extensions for Circuit Style Policies

Abstract

This document proposes a set of extensions for Path Computation Element Communication Protocol (PCEP) for Circuit Style Policies - Segment-Routing Policy designed to satisfy requirements for connection-oriented transport services. New TLV is introduced to control path recomputation and new flag to add ability to request path with strict hops only.

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.

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 3 June 2024.

Copyright Notice

Copyright (c) 2023 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. Introduction](#)
- [2. Terminology](#)
- [3. Overview of Extensions to PCEP](#)
 - [3.1. LSP-EXTENDED-FLAG TLV](#)
 - [3.2. PATH-RECOMPUTATION TLV](#)
- [4. Operation](#)
 - [4.1. Strict path enforcement](#)
 - [4.2. Path recomputation](#)
- [5. Manageability Considerations](#)
 - [5.1. Control of Function and Policy](#)
 - [5.2. Information and Data Models](#)
 - [5.3. Liveness Detection and Monitoring](#)
 - [5.4. Verify Correct Operations](#)
 - [5.5. Requirements On Other Protocols](#)
 - [5.6. Impact On Network Operations](#)
- [6. Implementation Status](#)
 - [6.1. Cisco](#)
- [7. Security Considerations](#)
- [8. IANA Considerations](#)
 - [8.1. LSP-EXTENDED-FLAG TLV](#)
 - [8.2. PATH-RECOMPUTATION TLV](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Contributors](#)
- [Authors' Addresses](#)

1. Introduction

Usage of Segment-routing and PCEP in connection-oriented transport services require path persistancy and hop-by-hop behavior for PCE computed paths.

Circuit-Style Policy introduced in [[I-D.ietf-spring-cs-sr-policy](#)] requires PCEP extensions, which are covered in this document.

This document:

- *Introduces possibility to request strict path from the PCE by extending LSP-EXTENDED-FLAG TLV

- *Adding new TLV to encode information about disabling path recomputation for specific path to the PCE, to be carried inside the LSPA object, which is defined in [[RFC5440](#)].

- *Clarifies usage of existing O-flag from RP object in Segment-routing

PCEP extensions described in this document are applicable to all Path Setup Types.

2. Terminology

The following terminologies are used in this document:

ERO: Explicit Route Object

IGP: Interior Gateway Protocol

LSP: Label Switched Path.

LSPA: Label Switched Path Attributes.

OTN: Optical Transport Network.

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Protocol.

SDH: Synchronous Digital Hierarchy

SID: Segment Identifier

SONET: Synchronous Optical Network

SR: Segment Routing.

SR-TE: Segment Routing Traffic Engineering.

3. Overview of Extensions to PCEP

3.1. LSP-EXTENDED-FLAG TLV

0-flag is proposed in the LSP-EXTENDED-FLAG TLV, which was introduced in 3.1 of [[RFC9357](#)] and extended with E-flag in [[I-D.peng-pce-entropy-label-position](#)]. TLV format will be added after assigning 0 flag bit position by IANA.


Type (16 bits): the value is TBD1 by IANA.

Length (16 bits): multiple of 4 octets.

0 (Strict-Path): If set to 1, this indicates to the PCE that a path exclusively made of strict hops is required. Strict hop definition is described in Section 4.1

3.2. PATH-RECOMPUTATION TLV

This document defines new TLV for the LSPA Object for encoding information whether path recomputation is allowed for delegated LSP. The TLV is optional. If the TLV is included in LSPA object, the PCE MUST NOT recompute path in cases specified by flags in the TLV.



Type (16 bits): the value is TBD2 by IANA.

Length (16 bits): 4 octets

Reserved: MUST be set to zero by the sender and MUST be ignored by the receiver.

Flags: This document defines the following flag bits. The other bits MUST be set to zero by the sender and MUST be ignored by the receiver.

*P (Permanent): If set to 1, the PCE MUST NOT recompute path even if current path is not satisfying path computation constraints. If this flag is cleared, then the PCE SHOULD recompute path if original path is invalidated.

*F (Force): If set to 1, the PCE MUST NOT update path. If flag is cleared, the PCE MAY update path based on explicit request from operator.

4. Operation

4.1. Strict path enforcement

PCC MAY set the O flag in LSP-EXTENDED-FLAG TLV in PCRpt message to the PCE to indicate that a path exclusively made of strict hops is required.

O flag cleared or LSP-EXTENDED-FLAG TLV not included indicates that a loose path is acceptable.

In PCUpdate or PCInitiate messages, PCE MAY set O bit if strict path is provided.

The flag is applicable only for stateful messages. Existing O flag in RP object MAY be used to indicate similar behavior in PCReq and PCRep messages as described in as described in Section 7.4.1 of [\[RFC5440\]](#).

If O flag is set to 1 for both stateful and stateless messages for SR paths introduced in [\[RFC8664\]](#), PCE MUST use only SIDs, which will use explicitly specified adjacencies for packet forwarding. For example Adjacency SIDs MAY be used, but Prefix SIDs MUST NOT be used (even if there is only one adjacency).

4.2. Path recomputation

PCC MAY set flags in PATH-RECOMPUTATION TLV to control path computation behavior on PCE side. If TLV is not included, then the PCE MAY use local policy to trigger path-computation or LSP path update.

The presence of the TLV is blocking path recomputation based on various triggers like topology update, any periodic update or changed state of other LSPs in the network. LSP path MAY be modified if forwarded packets will still use same path - for example if same path can be encoded using Adjacency and Prefix SIDs, then PCE MAY switch between various representations of same path.

If P flag is cleared, the PCE MAY recompute if current path is not considered valid, for example after topology update resulting in path not satisfying LSP's path constraints, but it MUST NOT recompute path if current path is not optimal.

If P flag is set, the PCE MUST NOT recompute path during LSP lifetime even if path is invalidated. Only exception is explicit request from operator to recompute path

If F flag is cleared, path update triggered manually by operator or any northbound interface of PCE MAY be done. If flag is set the PCE CAN update path only to tear down LSP by sending PCUpdate message with empty ERO.

TLV MAY be included in PCInitiate and PCUpdate messages to indicate, which triggers will be disabled on the PCE. PCC should reflect flag values in PCRpt messages to forward requirement to other PCEs in the network.

5. Manageability Considerations

All manageability requirements and considerations listed in [[RFC5440](#)], [[RFC8231](#)] and [[RFC8281](#)] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

5.1. Control of Function and Policy

A PCE or PCC implementation MAY allow the capability of supporting PCEP extensions introduced in this document to be enabled/disabled as part of the global configuration.

5.2. Information and Data Models

An implementation SHOULD allow the operator to view the capability defined in this document. Section 4.1 and 4.1.1 of [[I-D.ietf-pce-pcep-yang](#)] should be extended to include that capability for PCEP peer.

Section 4.2 of [[I-D.ietf-pce-pcep-yang](#)] module should be extended to add notification for blocked recomputation satisfying specified constraints if recomputation is blocked using PATH-RECOMPUTATION TLV.

5.3. Liveness Detection and Monitoring

Circuit-Style Policy draft [[I-D.ietf-spring-cs-sr-policy](#)] is already describing connectivity verification and path validity considerations for Circuit Style Policies.

5.4. Verify Correct Operations

A PCE implementation SHOULD notify operator in case blocked recomputation for path, which is no longer satisfying specified

constraints and it SHOULD allow the operator to view LSPs on PCE, which are not satisfying specified constraints.

5.5. Requirements On Other Protocols

The PCEP extensions defined in this document do not imply any new requirements on other protocols. Overall concept of Circuit Style policies requires interaction with other protocols, but those requirements are already described in [\[I-D.ietf-spring-cs-sr-policy\]](#).

5.6. Impact On Network Operations

The mechanisms defined in [\[RFC5440\]](#), [\[RFC8231\]](#), and [\[RFC8281\]](#) also apply to the PCEP extensions defined in this document.

6. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to RFC 7942.]

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".

6.1. Cisco

*Organization: Cisco Systems

*Implementation: IOS-XR PCC and PCE.

*Description: PCEP extensions supported using VENDOR-INFORMATION Object.

*Maturity Level: Production.

*Coverage: Partial.

*Contact: ssidor@cisco.com

7. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8253], [RFC8281] and [RFC8664] in itself.

Note that this specification introduces possibility to block path recomputation after various topology events. This creates an additional vulnerability if the security mechanisms of [RFC5440], [RFC8231], and [RFC8281] are not used. If there is no integrity protection on the session, then an attacker could block path updates from PCE potentially resulting in traffic drop.

8. IANA Considerations

8.1. LSP-EXTENDED-FLAG TLV

[RFC9357] defines the LSP-EXTENDED-FLAG TLV. IANA is requested to make the following assignment from the "LSP-EXTENDED-FLAG TLV Flag Field" registry:

Bit	Description	Reference
TBD1	Strict-Path Flag (0)	This document

Table 1

8.2. PATH-RECOMPUTATION TLV

IANA is requested to make the assignment of a new value for the existing "PCEP TLV Type Indicators" registry as follows:

TLV Type	TLV Name	Reference
TBD2	PATH-RECOMPUTATION TLV	This document

Table 2

9. References

9.1. Normative References

[I-D.peng-pce-entropy-label-position] Xiong, Q., Peng, S., and F. Qin, "PCEP Extension for SR-MPLS Entropy Label Position", Work in Progress, Internet-Draft, draft-peng-pce-entropy-label-position-10, 5 March 2023, <<https://www.ietf.org/archive/id/draft-peng-pce-entropy-label-position-09.txt>>.

- [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>>.
- [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>>.
- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [RFC9357] Xiong, Q., "Label Switched Path (LSP) Object Flag Extension for Stateful PCE", RFC 9357, February 2023, <<https://www.rfc-editor.org/info/rfc9357>>.

9.2. 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)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-22, 31 October 2019, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-yang-22>>.
- [I-D.ietf-spring-cs-sr-policy] Schmutzer, C., Ali, Z., Maheshwari, P., Rokui, R., and A. Stone, "Circuit Style Segment Routing Policies", Work in Progress, Internet-Draft, draft-ietf-spring-cs-sr-policy-01, 16 June 2023,

<<https://www.ietf.org/archive/id/draft-ietf-spring-cs-sr-policy-01.txt>>.

[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>>.

[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>>.

[SDH] International Telecommunication Union, "Network node interface for the synchronous digital hierarchy (SDH)", ITU-T Recommendation G.707, October 2020.

[SONET] American National Standards Institute, "Synchronous Optical Network (SONET) Basic Description including Multiplex Structure, Rates, and Formats", ANSI T1.105, January 1995.

Contributors

Daniel Voyer
Bell Canada

Email: daniel.voyer@bell.ca

Reza Rokui
Ciena

Email: rrokui@ciena.com

Tarek Saad
Cisco Systems, Inc.

Email: tsaad.net@gmail.com

Zafar Ali
Cisco Systems, Inc.

Email: zali@cisco.com

Authors' Addresses

Samuel Sidor
Cisco Systems, Inc.
Eurovea Central 3.

811 09 Bratislava
Slovakia

Email: ssidor@cisco.com

Praveen Maheshwari
Airtel India

Email: Praveen.Maheshwari@airtel.com

Andrew Stone
Nokia

Email: andrew.stone@nokia.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com

Shuping Peng
Huawei Technologies

Email: pengshuping@huawei.com