

Internet Engineering Task Force (IETF)
Request for Comments: 8408
Category: Standards Track
ISSN: 2070-1721

S. Sivabalan
Cisco Systems, Inc.
J. Tantsura
Nuage Networks
I. Minei
Google, Inc.
R. Varga
Pantheon Technologies SRO
J. Hardwick
Metaswitch Networks
July 2018

Conveying Path Setup Type
in PCE Communication Protocol (PCEP) Messages

Abstract

A Path Computation Element (PCE) can compute Traffic Engineering (TE) paths through a network; these paths are subject to various constraints. Currently, TE paths are Label Switched Paths (LSPs) that are set up using the RSVP-TE signaling protocol. However, other TE path setup methods are possible within the PCE architecture. This document proposes an extension to the PCE Communication Protocol (PCEP) to allow support for different path setup methods over a given PCEP session.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8408>.

Copyright Notice

Copyright (c) 2018 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	3
1.1. Requirements Language	4
2. Terminology	4
3. Path Setup Type Capability TLV	4
4. Path Setup Type TLV	6
5. Operation	7
6. Manageability Considerations	8
7. Security Considerations	9
8. IANA Considerations	9
8.1. Additions to PCEP TLV Type Indicators Registry	9
8.2. New PCEP Path Setup Types Registry	9
8.3. Additions to PCEP-ERROR Object Error Types and Values Registry	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Acknowledgements	11
Contributors	11
Authors' Addresses	12

1. Introduction

[RFC5440] describes the PCE Communication Protocol (PCEP) for communication between a Path Computation Client (PCC) and a Path Computation Element (PCE) or between a PCE and a PCE. A PCC requests, from a PCE, a path subject to various constraints and optimization criteria. The PCE responds to the PCC with a hop-by-hop path in an Explicit Route Object (ERO). The PCC uses the ERO to set up the path in the network.

[RFC8231] specifies extensions to PCEP that allow a PCC to delegate its LSPs to a PCE. The PCE can then update the state of LSPs delegated to it. In particular, the PCE may modify the path of an LSP by sending a new ERO. The PCC uses this ERO to reroute the LSP in a make-before-break fashion. [RFC8281] specifies a mechanism that allows a PCE to dynamically instantiate an LSP on a PCC by sending the ERO and the characteristics of the LSP. The PCC creates the LSP using the ERO and other attributes sent by the PCE.

So far, PCEP and its extensions have assumed that the TE paths are label switched and are established via the RSVP-TE signaling protocol. However, other methods of LSP setup are possible in the PCE architecture (see [RFC4655] and [RFC4657]). This document generalizes PCEP to allow other LSP setup methods to be used. It defines two new TLVs and specifies the base procedures to facilitate this:

- o The PATH-SETUP-TYPE-CAPABILITY TLV allows a PCEP speaker to announce which LSP setup methods it supports when the PCEP session is established.
- o The PATH-SETUP-TYPE TLV allows a PCEP speaker to specify which setup method should be used for a given LSP. When multiple path setup types are deployed in a network, a given PCEP session may have to simultaneously support more than one path setup type. A PCEP speaker uses the PATH-SETUP-TYPE TLV to explicitly indicate the intended path setup type in the appropriate PCEP messages, unless the path setup type is RSVP-TE (which is assumed to be the path setup type if no other setup type is indicated). This is so that both the PCC and the PCE can take the necessary steps to set up the path.

This document defines a path setup type code for RSVP-TE. When a new path setup type (other than RSVP-TE) is introduced for setting up a path, a path setup type code and, optionally, a sub-TLV pertaining to the new path setup type will be defined by the document that specifies the new path setup type.

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. Terminology

The following terminology is used in this document:

- ERO: Explicit Route Object
- PCC: Path Computation Client
- PCE: Path Computation Element
- PCEP: PCE Communication Protocol
- PST: Path Setup Type
- TLV: Type, Length, and Value

3. Path Setup Type Capability TLV

A PCEP speaker indicates which PSTs it supports during the PCEP initialization phase using the following process. When the PCEP session is created, it sends an Open message with an OPEN object containing the PATH-SETUP-TYPE-CAPABILITY TLV. The format of this TLV is as follows.

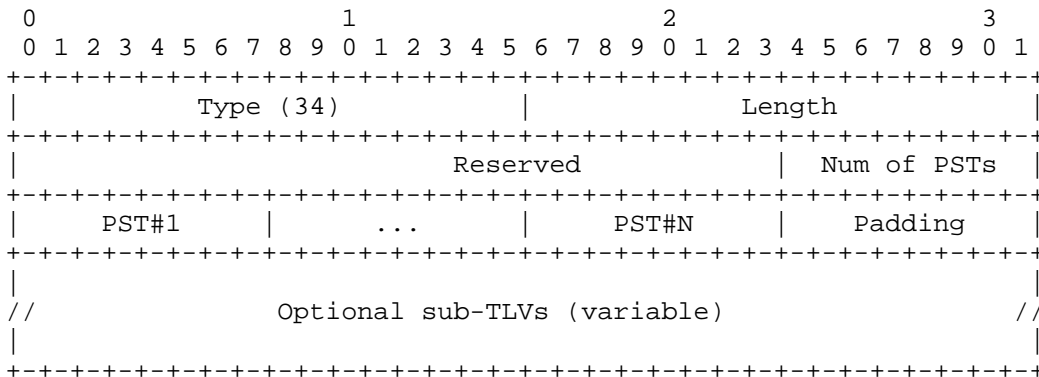


Figure 1: PATH-SETUP-TYPE-CAPABILITY TLV

The TLV Type is 34. Its Reserved field MUST be set to zero by the sender and MUST be ignored by the receiver. The other fields in the TLV are as follows.

Length: The total length in bytes of the remainder of the TLV, that is, excluding the Type and Length fields.

Num of PSTs: The number of PSTs in the following list, excluding padding.

List of PSTs: A list of the PSTs that the PCEP speaker supports. Each PST is a single byte in length. Duplicate entries in this list MUST be ignored. The PCEP speaker MUST pad the list with zeros so that it is a multiple of four bytes in length. This document defines the following PST value:

* PST = 0: Path is set up using the RSVP-TE signaling protocol

Optional sub-TLVs: A list of sub-TLVs associated with the supported PSTs. Each PST has zero or one sub-TLVs associated with it, and each sub-TLV is associated with exactly one PST. Each sub-TLV MUST obey the rules for TLV formatting defined in [RFC5440]. That is, each sub-TLV is padded to a four-byte alignment, and the Length field of each sub-TLV does not include the padding bytes. This document does not define any sub-TLVs; an example sub-TLV can be found in [PCEP-EXTENSIONS].

A PCEP speaker MUST check that this TLV is correctly formatted, as follows.

- o If there are no sub-TLVs, then the TLV Length field MUST be equal to four bytes plus the size of the PST list, excluding any padding bytes.
- o If there are sub-TLVs, then the TLV Length field MUST be equal to four bytes plus the size of the PST list (rounded up to the nearest multiple of four) plus the size of the appended sub-TLVs, excluding any padding bytes in the final sub-TLV.
- o The Num of PSTs field MUST be greater than zero.

If a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV that violates these rules, then the PCEP speaker MUST send a PCERR message with Error-Type = 10 (Reception of an invalid object) and Error-value = 11 (Malformed object) and MUST close the PCEP session. The PCEP speaker MAY include the malformed OPEN object in the PCERR message as well.

If a PCEP speaker receives an OPEN object with more than one PATH-SETUP-TYPE-CAPABILITY TLV, then it MUST ignore all but the first instance of this TLV.

The absence of the PATH-SETUP-TYPE-CAPABILITY TLV from the OPEN object is equivalent to a PATH-SETUP-TYPE-CAPABILITY TLV containing a single PST value of 0 (Path is set up using the RSVP-TE signaling protocol) and no sub-TLVs. A PCEP speaker MAY omit the PATH-SETUP-TYPE-CAPABILITY TLV if the only PST it supports is RSVP-TE. If a PCEP speaker supports other PSTs besides RSVP-TE, then it SHOULD include the PATH-SETUP-TYPE-CAPABILITY TLV in its OPEN object.

If a PCEP speaker does not recognize the PATH-SETUP-TYPE-CAPABILITY TLV, it will ignore the TLV in accordance with [RFC5440].

4. Path Setup Type TLV

When a PCEP session is used to set up TE paths using different methods, the corresponding PCE and PCC must be aware of the path setup method used. This means that a PCE must be able to specify paths in the correct format, and a PCC must be able to take control-plane and forwarding-plane actions appropriate to the PST.




Figure 2: PATH-SETUP-TYPE TLV

The PATH-SETUP-TYPE TLV is an optional TLV associated with the Request Parameters (RP) [RFC5440] and the Stateful PCE Request Parameters (SRP) [RFC8231] objects. Its format is shown in Figure 2. The TLV type is 28. Its Reserved field MUST be set to zero. The one-byte PST field contains the PST as defined for the PATH-SETUP-TYPE-CAPABILITY TLV.

The absence of the PATH-SETUP-TYPE TLV is equivalent to a PATH-SETUP-TYPE TLV with a PST value of 0 (Path is set up using the RSVP-TE signaling protocol). A PCEP speaker MAY omit the TLV if the PST is RSVP-TE. If the RP or SRP object contains more than one PATH-SETUP-TYPE TLV, only the first TLV MUST be processed, and the rest MUST be ignored.

If a PCEP speaker does not recognize the PATH-SETUP-TYPE TLV, it will ignore the TLV in accordance with [RFC5440] and use RSVP-TE to set up the path.

5. Operation

During the PCEP initialization phase, if a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV from its peer, it MUST assume that the peer supports only the PSTs listed in the TLV. If the PCEP speaker and its peer have no PSTs in common, then the PCEP speaker MUST send a PCErr message with Error-Type = 21 (Invalid traffic engineering path setup type) and Error-value = 2 (Mismatched path setup type) and close the PCEP session.

If the peer has sent no PATH-SETUP-TYPE-CAPABILITY TLV, then the PCEP speaker MUST infer that the peer supports path setup using at least RSVP-TE. The PCEP speaker MAY also infer that the peer supports other path setup types, but the means of inference are outside the scope of this document.

When a PCC sends a PCReq message to a PCE [RFC5440], it MUST include the PATH-SETUP-TYPE TLV in the RP object, unless the intended PST is RSVP-TE (in which case it MAY omit the PATH-SETUP-TYPE TLV). If the PCE is capable of expressing the path in a format appropriate to the intended PST, it MUST use the appropriate ERO format in the PCRep message.

When a PCE sends a PCRep message to a PCC [RFC5440], it MUST include the PATH-SETUP-TYPE TLV in the RP object, unless the PST is RSVP-TE (in which case it MAY omit the PATH-SETUP-TYPE TLV). If the PCE does not support the intended PST, it MUST send a PCErr message with Error-Type = 21 (Invalid traffic engineering path setup type) and Error-value = 1 (Unsupported path setup type) and close the PCEP session. If the PSTs corresponding to the PCReq and PCRep messages do not match, the PCC MUST send a PCErr message with Error-Type = 21 (Invalid traffic engineering path setup type) and Error-value = 2 (Mismatched path setup type) and close the PCEP session.

When a stateful PCE sends a PCUpd message [RFC8231] or a PCInitiate message [RFC8281] to a PCC, it MUST include the PATH-SETUP-TYPE TLV in the SRP object, unless the intended PST is RSVP-TE (in which case it MAY omit the PATH-SETUP-TYPE TLV). If the PCC does not support the PST associated with the PCUpd or PCInitiate message, it MUST send a PCErr message with Error-Type = 21 (Invalid traffic engineering path setup type) and Error-value = 1 (Unsupported path setup type) and close the PCEP session.

When a PCC sends a PCRpt message to a stateful PCE [RFC8231], it MUST include the PATH-SETUP-TYPE TLV in the SRP object, unless the PST is RSVP-TE (in which case it MAY omit the PATH-SETUP-TYPE TLV). The PCC MUST include the SRP object in the PCRpt message if the PST is not RSVP-TE, even when the SRP-ID-number is the reserved value of 0x00000000. If the PCRpt message is triggered by a PCUpd or PCInitiate message, then the PST that the PCC indicates in the PCRpt message MUST match the PST that the stateful PCE intended in the PCUpd or PCInitiate message. If it does not match, then the PCE MUST send a PCErr message with Error-Type = 21 (Invalid traffic engineering path setup type) and Error-value = 2 (Mismatched path setup type) and close the PCEP session.

6. Manageability Considerations

This document generalizes PCEP to allow path setup methods other than RSVP-TE to be used by the network (but does not define any new path setup types besides RSVP-TE). It is possible that, in a given network, multiple path setup methods will be used. It is also possible that not all devices will support the same set of path setup methods. Managing networks that combine multiple path setup methods may therefore raise some challenges from a configuration and observability point of view.

Each document that defines a new path setup type in the "PCEP Path Setup Types" registry (Section 8.2) must include a Manageability Considerations section. The Manageability Considerations section must explain how operators can manage PCEP with the new path setup type. It must address the following questions, which are generally applicable when working with multiple path setup types in PCEP.

- o What are the criteria for when devices will use the new path setup type in PCEP, and how can the operator control this?
- o How can the network be migrated to the new path setup type, and are there any backwards-compatibility issues that operators need to be aware of?
- o Are paths set up using the new path setup type intended to coexist with other paths over the long term, and if so, how is this situation managed with PCEP?
- o How can operators verify the correct operation of PCEP in the network with respect to the new path setup type? Which fault conditions must be reported to the operators?

- o Are there any existing management interfaces (such as YANG models) that must be extended to model the operation of PCEP in the network with respect to the new path setup type?

See [RFC5706] for further guidance on how to write Manageability Considerations sections in Standards Track documents.

7. Security Considerations

The security considerations described in [RFC5440] and [RFC8281] are applicable to this specification. No additional security measure is required.

Note that if the security mechanisms of [RFC5440] and [RFC8281] are not used, then the protocol described in this document could be attacked in the following new way. An attacker, using a TCP man-in-the-middle attack, could inject error messages into the PCEP session when a particular PST is (or is not) used. Doing this could potentially force the use of a specific PST, which may allow the attacker to subsequently attack a weakness in that PST.

8. IANA Considerations

8.1. Additions to PCEP TLV Type Indicators Registry

IANA has allocated the following code points in the "PCEP TLV Type Indicators" registry.

Value	Description	Reference
28	PATH-SETUP-TYPE	RFC 8408
34	PATH-SETUP-TYPE-CAPABILITY	RFC 8408

8.2. New PCEP Path Setup Types Registry

IANA has created a new sub-registry within the "Path Computation Element Protocol (PCEP) Numbers" registry called "PCEP Path Setup Types". The allocation policy for this new registry is IETF Review [RFC8126]. This new registry contains the following value:

Value	Description	Reference
0	Path is set up using the RSVP-TE signaling protocol	RFC 8408

8.3. Additions to PCEP-ERROR Object Error Types and Values Registry

IANA has allocated the following code points in the "PCEP-ERROR Object Error Types and Values" registry.

Error-Type	Meaning	Reference
10	Reception of an invalid object	RFC 5440
	Error-value = 11: Malformed object	RFC 8408
21	Invalid traffic engineering path setup type	RFC 8408
	Error-value = 0: Unassigned	RFC 8408
	Error-value = 1: Unsupported path setup type	RFC 8408
	Error-value = 2: Mismatched path setup type	RFC 8408

9. References

9.1. 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>>.
- [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>>.

9.2. Informative References

- [PCEP-EXTENSIONS] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", Work in Progress, draft-ietf-pce-segment-routing-12, June 2018.
- [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>>.
- [RFC4657] Ash, J., Ed. and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <<https://www.rfc-editor.org/info/rfc4657>>.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", RFC 5706, DOI 10.17487/RFC5706, November 2009, <<https://www.rfc-editor.org/info/rfc5706>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

Acknowledgements

We would like to thank Marek Zavodsky for valuable comments.

Contributors

The following people contributed to this document:

- Jan Medved
- Edward Crabbe

Authors' Addresses

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

Email: msiva@cisco.com

Jeff Tantsura
Nuage Networks
755 Ravendale Drive
Mountain View, CA 94043
United States of America

Email: jefftant.ietf@gmail.com

Ina Minei
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States of America

Email: inaminei@google.com

Robert Varga
Pantheon Technologies SRO
Mlynske Nivy 56
Bratislava, 821 05
Slovakia

Email: nite@hq.sk

Jon Hardwick
Metaswitch Networks
100 Church Street
Enfield, Middlesex
United Kingdom

Email: jonathan.hardwick@metaswitch.com