PCEP Working Group
Internet-Draft
Intended status: Standards Track
Expires: 2 September 2026

V. P. Beeram
HPE
K. Kompella
Juniper Networks
A. Stone
Nokia

1 March 2026

# Path Computation Element Communication Protocol (PCEP) Extensions for Multipath Traffic Engineered Directed Acyclic Graph (MPTED) Tunnels
draft-beeram-pce-pcep-mpted-01

## Abstract

A Multipath Traffic Engineered Directed Acyclic Graph (MPTED) tunnel is a Traffic Engineering (TE) construct that facilitates weighted load balancing of unicast traffic across a constrained set of paths optimized for a specific objective.

This document describes the provisioning of an MPTED Tunnel in a TE network using Path Computation Element Communication Protocol (PCEP) in a stateful PCE model.

## 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 2 September 2026.

## Copyright Notice

Table of Contents



## 1. Introduction

The notion of a Multipath Traffic Engineering Directed Acyclic Graph (MPTED) tunnel is introduced in [I-D.draft-kompella-teas-mpte]. An MPTED tunnel is a Traffic Engineering (TE) construct that contains a constrained set of paths representing an optimized Directed Acyclic Graph (DAG) from one or more ingresses to one or more egresses. The paths that make up an MPTED tunnel traverse a set of junction nodes. An MPTED junction refers to the construct associated with the MPTED tunnel at each junction node and constitutes a set of previous-hops and a set of next-hops over which traffic is load-balanced in a weighted fashion. Provisioning an MPTED tunnel in a TE network involves provisioning the control and forwarding plane state

associated with the MPTED junction at each junction node.

[RFC8231] specifies extensions to the Path Computation Element Protocol (PCEP) that enable the deployment of a stateful Path Computation Element (PCE) model. These extensions allow a Path Computation Client (PCC) to delegate control of the Label Switched Paths (LSPs) associated with its TE Tunnels to a stateful PCE. [RFC8281] specifies extensions that allow a PCE to instantiate and manage PCE-initiated LSPs on a PCC under the stateful PCE model. This document describes the extensions to PCEP that would enable (a) a PCC to delegate control of the MPTED TE Tunnel to a stateful PCE and (b) a PCE to instantiate and manage PCE-initiated MPTED Tunnels on a PCC under the stateful PCE model.

Each MPTED computation request that a PCE caters to carries a set of ingress nodes, a set of egress nodes, a set of constraints, and an optimization objective. The computation result for the MPTED contains a set of unordered elements called JUNCTIONs. Each ingress, transit, and egress node on the DAG is a junction and has a JUNCTION element associated with it. A JUNCTION element contains the information necessary to provision a specific junction node in the computed DAG. The onus is on the MPTED tunnel signaling source to signal and provision the MPTED junction on each junction node. Depending on the type of the MPTED tunnel and the signaling protocol used, the signaling source function may reside either on the PCC or the PCE. This document also describes the extensions to PCEP that enable the PCE to act as an MPTED tunnel signaling source and manage the MPTED junctions on a junction node.

As discussed in [I-D.draft-kompella-teas-mpte], an MPTED tunnel may be realized over a Multiprotocol Label Switching (MPLS) forwarding plane or a native Internet Protocol (IP) v4/v6 forwarding plane using an appropriate tunnel type. The focus of this version of the document is on discussing how the PCEP protocol is extended to facilitate distributed and centralized provisioning of MPTED Tunnels over an MPLS forwarding plane in a TE network.

## 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.

### 1.2. Terminology

The reader is expected to be familiar with the terminology used in [I-D.draft-kompella-teas-mpte].

## 2. Modes of Operation

### 2.1. PCC Initiated and PCC Signaled

In this mode of operation, the PCC is the MPTED tunnel originator and the signaling source. The PCC delegates the control of the MPTED tunnel to the PCE. The PCE computes the MPTED, produces a set of JUNCTIONs, and puts the onus on the PCC to signal and provision the JUNCTION on each junction node. . After the signaling setup sequence is complete, the PCC notifies the PCE of the status of each junction in the DAG. The MPTED tunnel setup is deemed complete on the PCE when all junction notifications are received from the PCC. An example of this mode of operation is an MPLS deployment where RSVP MPTED tunnels that use signaled label switching originate on an ingress node, and the DAG computation is offloaded to the PCE.

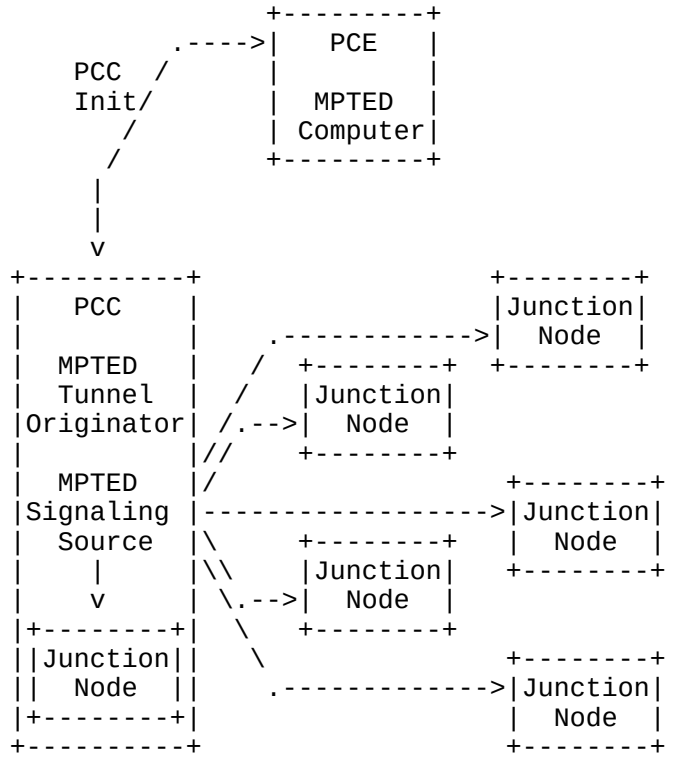

Figure 1: PCC Initiated and PCC Signaled

## 2.2. PCE Initiated and PCC Signaled

In this mode of operation, the PCE serves as the MPTED tunnel originator, and the PCC acts as the signaling source. The PCE computes the MPTED and initiates the setup process by providing the PCC a list of JUNCTIONs. The PCC signals and provisions the JUNCTION on each junction node. After the signaling setup sequence is complete, the PCC notifies the PCE of the status of each junction in the DAG. The MPTED tunnel setup is deemed complete on the PCE when all junction notifications are received from the PCC. An example of this mode of operation is an MPLS deployment where RSVP MPTED tunnels, which use signaled label switching, are originated and computed by the PCE.

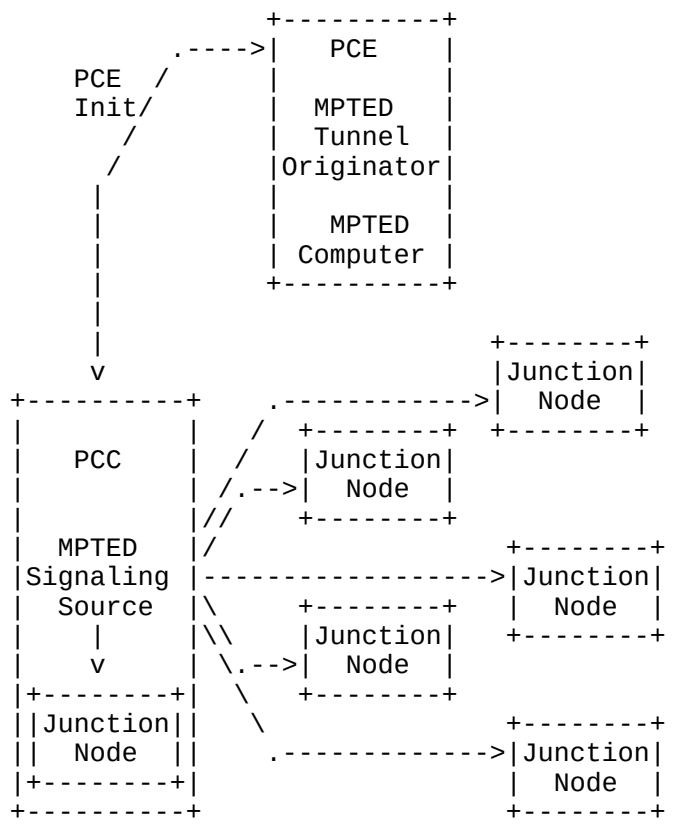


Figure 2: PCE Initiated and PCC Signaled

## 2.3. PCC Initiated and PCE Signaled

In this mode of operation, the PCC serves as the MPTED tunnel originator, and the PCE acts as the signaling source. The PCC delegates the control of the MPTED tunnel to the PCE. The PCE computes the MPTED, produces a set of JUNCTIONs, and uses a signaling protocol to provision the JUNCTION on each junction node. PCEP MAY be used as the signaling protocol on the PCE for junction management. The MPTED tunnel setup is deemed complete on the PCE when junction notifications are received from all junction nodes. An example of this mode of operation with PCEP signaling for junction management is an MPLS deployment where Segment Routing MPTED tunnels, which use static labels, originate on an ingress node and are provisioned in the TE network by the PCE.

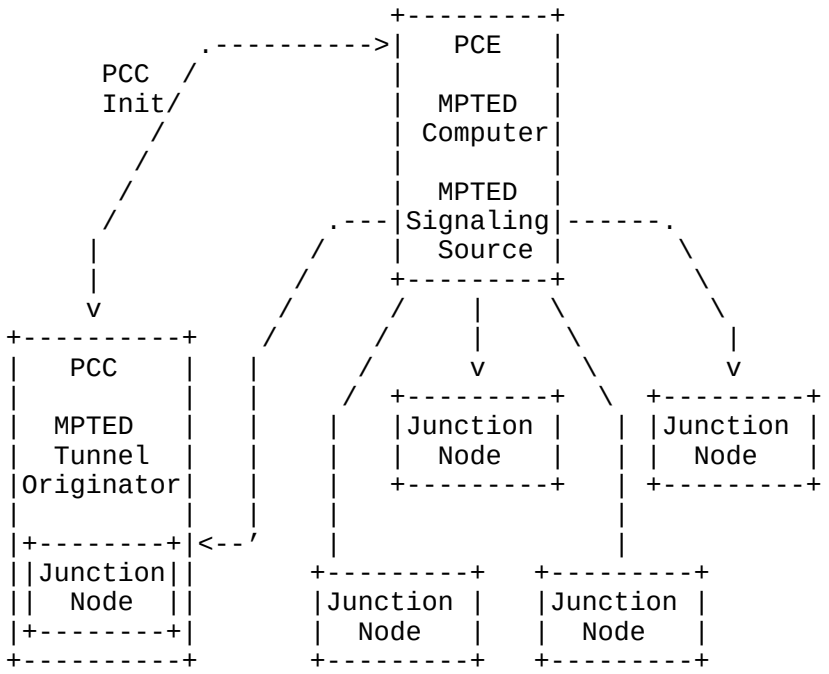


Figure 3: PCC Initiated and PCE Signaled

## 2.4. PCE Initiated and PCE Signaled

In this mode of operation, the PCE serves as the MPTED tunnel originator and the signaling source. The PCE computes the MPTED, produces a set of JUNCTIONs, and uses a signaling protocol to provision the JUNCTION on each junction node. PCEP MAY be used as the signaling protocol on the PCE for junction management. The MPTED tunnel setup is deemed complete on the PCE when junction notifications are received from all junction nodes. An example of this mode of operation with PCEP signaling for junction management is an MPLS deployment where Segment Routing MPTED tunnels, which use static labels, are originated and provisioned in the TE network by the PCE.

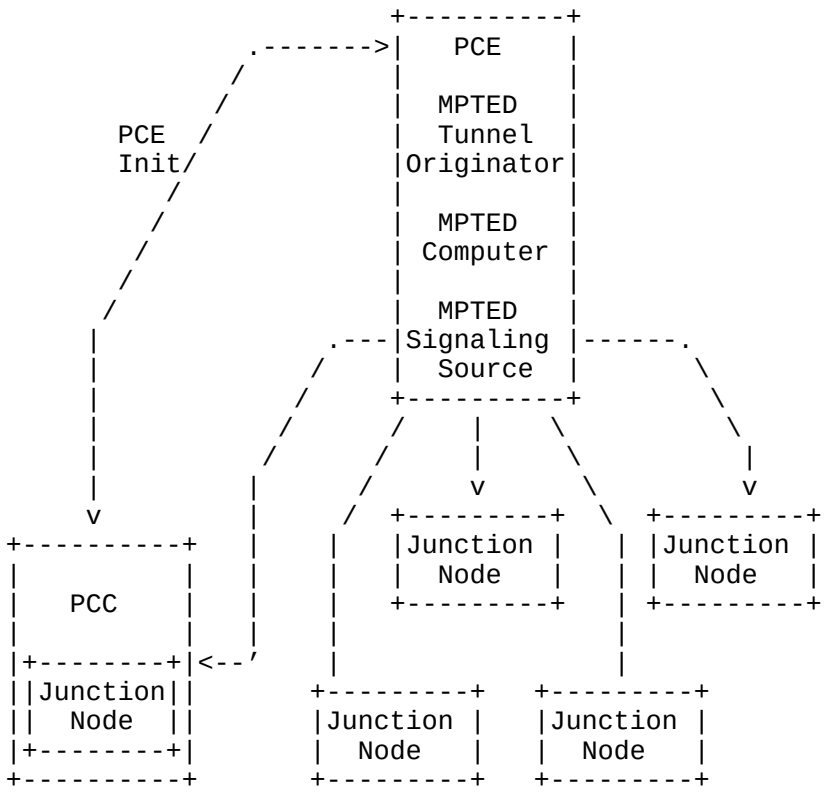


Figure 4: PCE Initiated and PCE Signaled

# 3. Protocol Extensions

### 3.1. Capability Negotiation

Depending on the deployment needs and the signaling protocol used for junction management, one or more entities can perform the roles of (a) MPTED tunnel originator, (b) MPTED computer, and (c) MPTED signaling source. The stakeholders involved in provisioning an MPTED tunnel in a stateful PCE model are the PCC, the PCE, and the junction nodes. The capabilities of each stakeholder must be advertised to enable the assignment of an appropriate role for each stakeholder when originating an MPTED tunnel.

The extensions to the capability negotiation procedures on the PCE and PCC will be discussed in a subsequent version.

### 3.2. PCC Initiated MPTED Tunnels

The protocol procedures to facilitate PCC initiated MPTED tunnels will be discussed in a subsequent version.

### 3.3. PCE Initiated MPTED Tunnels

The protocol procedures to facilitate PCE initiated MPTED tunnels will be discussed in a subsequent version.

### 3.4. Signaling for Junction Management

The procedures to facilitate junction management using PCEP signaling will be discussed in a subsequent version.

### 3.5. PCEP Messages and Objects

The definition of MPTED PCEP messages and objects will be furnished in a subsequent revision.

## 4. Security Considerations

The security considerations described in [RFC5440], [RFC8231], and [RFC8281] are applicable to this specification. An unauthorized PCE may maliciously influence the paths traversed by traffic flows steered onto the MPTED tunnel by manipulating the state on the junction nodes. The procedures described in [RFC8253] and [RFC9325] can be used to protect against this attack.

## 5. Manageability Considerations

This section will be updated in a subsequent version and will follow the advice and guidance of [RFC6123]

## 6. IANA Considerations

This version of the document has no IANA actions. The requests for appropriate IANA actions will be added in a subsequent version.

## 7. References

### 7.1. Normative References

[I-D.draft-kompella-teas-mpte]
Kompella, K., Jalil, L., Khaddam, M., and A. Smith, "Multipath Traffic Engineering", Work in Progress, Internet-Draft, draft-kompella-teas-mpte-01, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-kompella-teas-mpte-01>.

[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/rfc/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/rfc/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/rfc/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/rfc/rfc8231>.

[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/rfc/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/rfc/rfc8281>.

[RFC9325] Sheffer, Y., Saint-Andre, P., and T. Fossati, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 9325, DOI 10.17487/RFC9325, November 2022, <https://www.rfc-editor.org/rfc/rfc9325>.

### 7.2. Informative References

[RFC6123] Farrel, A., "Inclusion of Manageability Sections in Path Computation Element (PCE) Working Group Drafts", RFC 6123, DOI 10.17487/RFC6123, February 2011, <https://www.rfc-editor.org/rfc/rfc6123>.

## Acknowledgments

The authors would like to thank Sudharsana Venkatraman for her input from discussions.

This document was prepared using kramdown.

## Contributors

Chandrasekar Ramachandran
HPE
Email: chandrasekar.ramachandran@hpe.com

## Authors' Addresses

Vishnu Pavan Beeram
HPE
Email: vishnupavan.ietf@gmail.com

Kireeti Kompella
Juniper Networks
Email: kireeti.ietf@gmail.com

Andrew Stone
Nokia
Email: andrew.stone@nokia.com