

PCE
Internet-Draft
Intended status: Standards Track
Expires: 29 November 2025

Q. Xiong
S. Peng
ZTE Corporation
V. Beeram
T. Saad
Juniper Networks
M. Koldychev
Ciena Corporation
28 May 2025

Path Computation Element Communication Protocol (PCEP) Extensions for
Topology Filter
draft-xpbs-pce-topology-filter-04

Abstract

This document proposes a set of extensions for Path Computation Element Communication Protocol (PCEP) to support the topology filter during path computation.

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 29 November 2025.

Copyright Notice

Copyright (c) 2025 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
1.1. Terminology	3
1.2. Requirements Language	3
2. Topology Filter with PCE	3
3. PCEP Extensions	4
3.1. TOPOLOGY-FILTER Object	4
3.1.1. IGP Domain Identifier	4
3.1.1.1. Protocol ID TLV	5
3.1.1.2. Multi-topology ID TLV	5
3.1.1.3. Algorithm ID TLV	6
3.1.1.4. Domain ID TLV	6
3.1.2. TE Topology Identifier	6
3.1.2.1. Provider ID TLV	7
3.1.2.2. Client ID TLV	7
3.1.2.3. Topology ID TLV	7
3.1.3. Filtering Rules TLV	8
3.1.3.1. Link ID sub-TLV	9
3.1.3.2. Admin Group sub-TLV	9
3.1.3.3. Source Protocol sub-TLV	10
3.2. Procedures	10
4. IANA Considerations	10
4.1. TOPOLOGY-FILTER Object	10
5. Acknowledgements	11
6. Security Considerations	11
7. References	11
7.1. Normative References	12
Authors' Addresses	14

1. Introduction

[RFC5440] describes the Path Computation Element Computation Protocol (PCEP) which is used between a Path Computation Element (PCE) and a Path Computation Client (PCC) (or other PCE) to enable computation of Multi-protocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP). PCEP Extensions for the Stateful PCE Model [RFC8231] describes a set of extensions to PCEP to enable active control of MPLS-TE and Generalized MPLS (GMPLS) tunnels. As depicted in [RFC4655], a PCE MUST be able to compute the path of a TE LSP by operating on the TED and considering bandwidth and other constraints applicable to the TE LSP service request.

A PCE may perform path computation based on the network topology information collected through BGP-LS [RFC9552]. BGP-LS can get multiple link-state data from multiple IGP instance, or multiple virtual topologies from a single IGP instance. In other cases, as per [I-D.ietf-teas-yang-topology-filter], a path may be computed within a network topology such as a specified topology, a topology associated with a specific IGP domain, a topology learnt from a specific TE information source, a topology defined by the application of one or more topology filters, a topology associated with an Network Resource Partition (NRP) and so on. The PCE MUST take the topology related constraints into consideration during the path computation.

As defined in [I-D.ietf-teas-yang-topology-filter], a topology filter is a data construct that is used to filter network topologies. This document proposes a set of extensions for PCEP to support the topology filter during path computation.

1.1. Terminology

The terminology is defined as [RFC5440], [RFC9552] and [RFC8795].

1.2. Requirements Language

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].

2. Topology Filter with PCE

As defined in [I-D.ietf-teas-yang-topology-filter], a topology filter specifies the topology reference or a set of filtering rules. The topology filters carry a list of topology filters and a topology filter-set constitutes a list of topology filter references.

The topology reference indicates a predefined TE topology or a specific IGP domain. A TE topology can be identified from a global scope such as a provider ID, a client ID or a topology ID. And a specific IGP domain can be identified by protocol ID, instance ID, division ID, algo ID and MT ID. The PCE should consider these identifiers as topology constraints during path computation.

The filtering rules specify a set of constraints on the topology including include-any, include-all and exclude. A set of attributes that can be used as rules to filter the topology such as link affinity, link name, node prefix, AS number and TE information source. The filtering rules of these attributes can be used to compute path at PCE.

3. PCEP Extensions

3.1. TOPOLOGY-FILTER Object

This document defines a new TOPOLOGY-FILTER object to carry the topology filter. The TOPOLOGY-FILTER object is optional and specifies the specific topology to be taken into account by the PCE during path computation. The TOPOLOGY-FILTER object can be carried within a PCReq message, or a PCRep message in case of unsuccessful path computation.

TOPOLOGY-FILTER Object-Class is TBD1.

TOPOLOGY-FILTER Object-Type is TBD2.

The format of the TOPOLOGY-FILTER object body is:




Figure 1: TOPOLOGY-FILTER Body Object Format

Reserved (24 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

Flags (8 bits): No flags are currently defined. Unassigned flags MUST be set to zero on transmission and MUST be ignored on receipt.

The format of optional TLVs is defined in [RFC5440] and may be used to carry topology filter information as defined in section. The existing topology constraints TLVs could also be reused such as Algorithm ID TLV and Domain ID TLV.

3.1.1. IGP Domain Identifier

As defined in [RFC9552], the IGP domain has a unique IGP representation by using the combination of Area-ID, Router-ID, Protocol-ID, Multi-Topology Identifier (MT-ID), and BGP-LS Instance-ID. This document defines some TLVs for topology filter to identify a IGP domain within a referenced topology. The Protocol ID TLV is mandatory to identify a IGP domain and others are optional to carry

the additional information to further filter permitted resources within the domain. These TLVs can be carried but each type can only be presented once. And it MUST be applied to filter the resources that match all presenting TLVs at the same time.

3.1.1.1. Protocol ID TLV

The Protocol ID TLV is mandatory to identify a IGP domain and the format is as following shown:

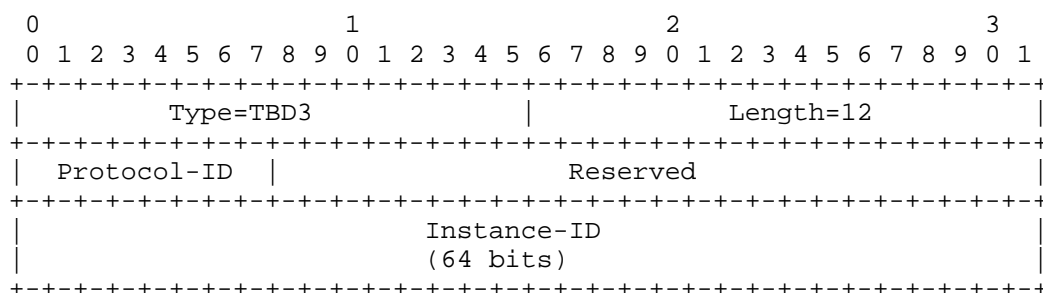


Figure 2: Protocol ID TLV

The code point for the TLV type is TBD3. The TLV length is 12 octets.

Protocol-ID (8 bits): defined in [RFC9552] section 5.2.

Instance-ID (64 bits): defined in [RFC9552] section 5.2.

Reserved: This fields MUST be set to zero on transmission and MUST be ignored on receipt.

3.1.1.2. Multi-topology ID TLV

The Multi-topology ID TLV is optional and is defined to carry the multi-topology-ID.

The format of the Multi-topology ID TLV is :




Figure 3: Multi-topology TLV

The code point for the TLV type is TBD4. The TLV length is 4 octets.

Multi-Topology-ID (12 bits): Semantics of the IS-IS MT-ID are defined in Section 7.2 of [RFC5120]. Semantics of the OSPF MT-ID are defined in Section 3.7 of [RFC4915]. As defined in section 3.2.1.5 of [RFC7752], if the value is derived from OSPF, then the upper 9 bits MUST be set to 0. Bits R are reserved and SHOULD be set to 0 when originated and ignored on receipt.

Reserved (16 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.

3.1.1.3. Algorithm ID TLV

The Algorithm ID TLV is optional and is defined to carry the Algorithm-ID.

The Algorithm ID TLV MAY be inserted so as to provide the Flex-algo plane information for the computed path. The format of the TLV is defined in [I-D.ietf-pce-sid-algo] section 3.4.

3.1.1.4. Domain ID TLV

The Domain ID TLV is optional and is defined to carry the Domain-ID.

The Domain ID TLV MAY be inserted so as to identify the domains served by the PCE. The format of the TLV is defined in [RFC8685] section 3.2.2.

3.1.2. TE Topology Identifier

This document defines some TE Topology Identifier TLVs for topology filter to identify a predefined TE topology within a referenced topology.


The Provider ID TLV is optional and the format is as following shown:



The code point for the TLV type is TBD5. The TLV length is 4 octets.

Provider-ID (32 bits): an identifier to uniquely identify a provider as defined in [RFC8776].

The Client ID TLV is optional and the format is as following shown:



The code point for the TLV type is TBD6. The TLV length is 4 octets.

Client-ID (32 bits): an identifier to uniquely identify a client as defined in [RFC8776].

The Topology ID TLV is optional and the format is as following shown:




Figure 6: Topology ID TLV

The code point for the TLV type is TBD7. The TLV length is 4 octets.

Topology-ID (32 bits): an identifier for a topology as defined in [RFC8776].

3.1.3. Filtering Rules TLV

This document defines a new Filtering Rules TLV for topology filter to carry a set of constrains on the topology by include-any, include-all and exclude.

The Filtering Rules TLV is optional and the format is as following shown :

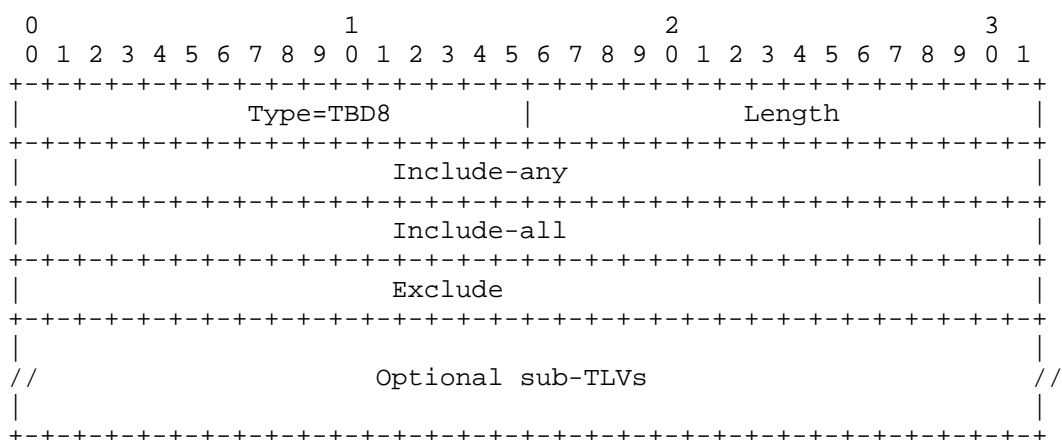


Figure 7: Filtering Rules TLV

The code point for the TLV type is TBD8. The TLV length is variable.

The sub-TLVs carry the attributes that can be used as rules to filter the topology.

3.1.3.1. Link ID sub-TLV

The Link ID is used to identify the link that is used during the path calculation.

The Link ID sub-TLV is defined:




Figure 8: Link ID sub-TLV

The code point for the TLV type is TBD9. The TLV length is 6 octets.

Link-ID (32bits): defined in IS-IS [RFC5307] and OSPF [RFC3630].

3.1.3.2. Admin Group sub-TLV

The Admin Group is used to include the links that is used during the path calculation.

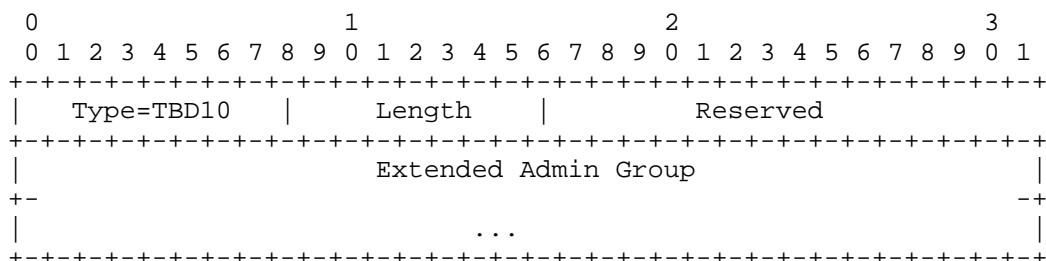


Figure 9: Admin Group sub-TLV

The code point for the sub-TLV type is TBD10. The length is variable.

Extended Administrative Group: Extended Administrative Group as defined in [RFC7308].

3.1.3.3. Source Protocol sub-TLV

The format of the Source Protocol sub-TLV is:

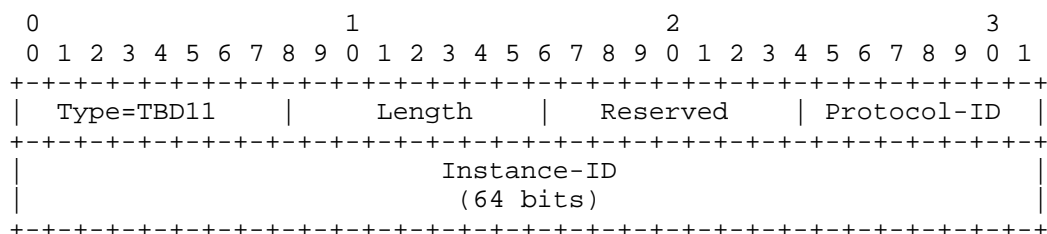


Figure 10: Source Protocol sub-TLV

The code point for the TLV type is TBD11. The TLV length is 10 octets.

Protocol-ID (8 bits): defined in [RFC9552] section 5.2.

Instance-ID (64 bits): defined in [RFC9552] section 5.2.

3.2. Procedures

A PCC MAY insert a TOPOLOGY-FILTER object in PCReq message to indicate the specific topology that MUST be considered by the PCE during path computation. The PCE will compute the path with the constraints with the filtering rules and reply the result to the PCC with a PCRep message.

The PCE could perform path computation based on the topology identified by the topology filter rules that can be applied on either the native topology or a user specified topology. The absence of the IGP Domain Identifier TLV and TE Topology Identifier TLV indicate that the PCE should compute within a native topology and only Filtering Rules TLV is applied as the filtering rules.

4. IANA Considerations

4.1. TOPOLOGY-FILTER Object

IANA is requested to make allocations for Topology Filter Object from the registry, as follows:

TOPOLOGY-FILTER Object-Class is TBD1.

TOPOLOGY-FILTER Object-Type is TBD2.

The TLVs for Topology Filter Object is as follows:

Type	TLV	Reference
TBD3	Protocol ID TLV	[this document]
TBD4	Multi-topology ID TLV	[this document]
TBD5	Provider ID TLV	[this document]
TBD6	Client ID TLV	[this document]
TBD7	Topology ID TLV	[this document]
TBD8	Filtering Rules TLV	[this document]

Table 1: TLVs for Topology Filter Object

IANA is requested to make allocations for sub-TLVs from the Filtering Rules TLV registry, as follows:

Type	sub-TLVs for Filtering Rules TLV	Reference
TBD9	Link ID sub-TLV	[this document]
TBD10	Admin Group sub-TLV	[this document]
TBD11	Source Protocol sub-TLV	[this document]

Table 2: Sub-TLVs

5. Acknowledgements

The authors would like to thank Dhruv Dhody, Andrew Stone for their review, suggestions and comments to this document.

6. Security Considerations

TBA

7. References

7.1. Normative References

- [I-D.ietf-pce-sid-algo]
Sidor, S., Rose, Z., Peng, S., Peng, S., and A. Stone,
"Carrying SR-Algorithm in Path Computation Element
Communication Protocol (PCEP).", Work in Progress,
Internet-Draft, draft-ietf-pce-sid-algo-19, 20 May 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-pce-sid-
algo-19](https://datatracker.ietf.org/doc/html/draft-ietf-pce-sid-algo-19)>.
- [I-D.ietf-teas-yang-topology-filter]
Beeram, V. P., Saad, T., Gandhi, R., and X. Liu, "YANG
Data Model for Topology Filter", Work in Progress,
Internet-Draft, draft-ietf-teas-yang-topology-filter-00,
11 October 2024, <[https://datatracker.ietf.org/doc/html/
draft-ietf-teas-yang-topology-filter-00](https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-topology-filter-00)>.
- [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>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering
(TE) Extensions to OSPF Version 2", RFC 3630,
DOI 10.17487/RFC3630, September 2003,
<<https://www.rfc-editor.org/info/rfc3630>>.
- [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>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P.
Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF",
RFC 4915, DOI 10.17487/RFC4915, June 2007,
<<https://www.rfc-editor.org/info/rfc4915>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi
Topology (MT) Routing in Intermediate System to
Intermediate Systems (IS-IS)", RFC 5120,
DOI 10.17487/RFC5120, February 2008,
<<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions
in Support of Generalized Multi-Protocol Label Switching
(GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008,
<<https://www.rfc-editor.org/info/rfc5307>>.

- [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>>.
- [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>>.
- [RFC6549] Lindem, A., Roy, A., and S. Mirtorabi, "OSPFv2 Multi-Instance Extensions", RFC 6549, DOI 10.17487/RFC6549, March 2012, <<https://www.rfc-editor.org/info/rfc6549>>.
- [RFC7308] Osborne, E., "Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)", RFC 7308, DOI 10.17487/RFC7308, July 2014, <<https://www.rfc-editor.org/info/rfc7308>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC8202] Ginsberg, L., Previdi, S., and W. Henderickx, "IS-IS Multi-Instance", RFC 8202, DOI 10.17487/RFC8202, June 2017, <<https://www.rfc-editor.org/info/rfc8202>>.
- [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>>.
- [RFC8685] Zhang, F., Zhao, Q., Gonzalez de Dios, O., Casellas, R., and D. King, "Path Computation Element Communication Protocol (PCEP) Extensions for the Hierarchical Path Computation Element (H-PCE) Architecture", RFC 8685, DOI 10.17487/RFC8685, December 2019, <<https://www.rfc-editor.org/info/rfc8685>>.
- [RFC8776] Saad, T., Gandhi, R., Liu, X., Beeram, V., and I. Bryskin, "Common YANG Data Types for Traffic Engineering", RFC 8776, DOI 10.17487/RFC8776, June 2020, <<https://www.rfc-editor.org/info/rfc8776>>.

- [RFC8795] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Gonzalez de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", RFC 8795, DOI 10.17487/RFC8795, August 2020, <<https://www.rfc-editor.org/info/rfc8795>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

Authors' Addresses

Quan Xiong
ZTE Corporation
China
Email: xiong.quan@zte.com.cn

Shaofu Peng
ZTE Corporation
No.50 Software Avenue
Nanjing
Jiangsu, 210012
China
Email: peng.shaofu@zte.com.cn

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net

Tarek Saad
Juniper Networks
Email: tsaad@juniper.net

Mike Koldychev
Ciena Corporation
385 Terry Fox Dr.
Kanata Ontario K2K 0L1
Canada
Email: mkoldych@proton.me