

PCE Working Group
Internet-Draft
Updates: 9168 (if approved)
Intended status: Standards Track
Expires: 3 April 2026

D. Dhody
Huawei
A. Farrel
Old Dog Consulting
Z. Li
Huawei Technologies
30 September 2025

PCEP Extension for Layer 2 (L2) Flow Specification
draft-ietf-pce-pcep-l2-flowspec-09

Abstract

The Path Computation Element (PCE) is a functional component capable of selecting paths through a traffic engineering (TE) network. These paths may be supplied in response to requests for computation or may be unsolicited requests issued by the PCE to network elements. Both approaches use the PCE Communication Protocol (PCEP) to convey the details of the computed path.

Traffic flows may be categorized and described using "Flow Specifications". RFC 8955 defines the Flow Specification and describes how Flow Specification components are used to describe traffic flows. RFC 8955 also defines how Flow Specifications may be distributed in BGP to allow specific traffic flows to be associated with routes.

RFC 9168 specifies a set of extensions to PCEP to support the dissemination of Flow Specifications. This allows a PCE to indicate what traffic should be placed on each path that it is aware of. This document updates RFC 9168 by updating the assignment policies for a range of Flow Specification TLV Type Indicators.

The extensions defined in this document extend the support for Ethernet Layer 2 (L2) and Layer 2 Virtual Private Network (L2VPN) traffic filtering rules either by themselves or in conjunction with Layer 3 (L3) flowspecs.

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 April 2026.

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	3
2. Terminology	4
3. L2 Flow Specifications	5
3.1. L2 Flow Specification TLVs	6
4. BGP Flow Specification Version 2	7
5. Update to RFC 9168	7
6. IANA Considerations	7
6.1. PCEP TLV Type Indicators	8
6.2. L2 Flow Specification TLV Type Indicators	8
6.3. Flow Specification TLV Type Indicators	8
7. Implementation Status	8
8. Security Considerations	9
9. Manageability Considerations	9
9.1. Control of Function and Policy	10
9.2. Information and Data Models	10
9.3. Liveness Detection and Monitoring	10
9.4. Verify Correct Operations	10
9.5. Requirements On Other Protocols	10
9.6. Impact On Network Operations	10
10. Acknowledgements	10
11. References	10
11.1. Normative References	10
11.2. Informative References	12
Appendix A. Contributors	13
Authors' Addresses	14

1. Introduction

[RFC4655] defines the Path Computation Element (PCE), as a functional component capable of computing paths for use in traffic engineering networks. PCE was originally conceived for use in Multiprotocol Label Switching (MPLS) for traffic engineering (TE) networks to derive the routes of Label Switched Paths (LSPs). However, the scope of PCE was quickly extended to make it applicable to networks controlled by Generalized MPLS (GMPLS), and more recent work has brought other traffic engineering technologies and planning applications into scope (for example, Segment Routing (SR) [RFC8664]).

[RFC5440] describes the PCE Communication Protocol (PCEP). PCEP defines the communication between a Path Computation Client (PCC) and a PCE, or between PCE and PCE, enabling computation of the path for MPLS-TE LSPs.

Stateful PCE [RFC8231] specifies a set of extensions to PCEP to enable control of TE-LSPs by a PCE that retains the state of the LSPs provisioned in the network (a stateful PCE). [RFC8281] describes the setup, maintenance, and teardown of LSPs initiated by a stateful PCE without the need for local configuration on the PCC, thus allowing for a dynamic network that is centrally controlled. [RFC8283] introduces the architecture for PCE as a central controller and describes how PCE can be viewed as a component that performs computation to place "flows" within the network and decide how these flows are routed.

The description of traffic flows by the combination of multiple Flow Specification components and their dissemination as traffic flow specifications (Flow Specifications) is described for BGP in [RFC8955]. In BGP, a Flow Specification is comprised of traffic filtering rules and is associated with actions to perform on the packets that match the Flow Specification. The BGP routers that receive a Flow Specification can classify received packets according to the traffic filtering rules and can direct packets based on the associated actions. [I-D.ietf-idr-flowspec-v2] specifies version 2 of the BGP flow specification protocol that resolves some of the issues with version 1.

When a PCE is used to initiate tunnels (such as TE-LSPs or SR paths) using PCEP, it is important that the head end of the tunnels understands what traffic to place on each tunnel. The data flows intended for a tunnel can be described using Flow Specification components. When PCEP is in use for tunnel initiation it makes sense for that same protocol to be used to distribute the Flow Specification components that describe what data is to flow on those tunnels.

[RFC9168] specifies a set of extensions to PCEP to support the dissemination of Flow Specification components. It includes the creation, update, and withdrawal of Flow Specifications via PCEP. It can be applied to tunnels initiated by the PCE or to tunnels where control is delegated to the PCE by the PCC. Furthermore, a PCC requesting a new path can include Flow Specifications in the request to indicate the purpose of the tunnel allowing the PCE to factor this into the path computation.

[I-D.ietf-idr-flowspec-l2vpn] defines a BGP flowspec extension to disseminate Ethernet Layer 2 (L2) and Layer 2 Virtual Private Network (L2VPN) traffic filtering rules either by themselves or in conjunction with L3 flowspecs as per [I-D.ietf-idr-flowspec-v2]. This document extends the same support for PCEP by defining a new L2 Flow Filter TLV to be carried within the FLOWSPEC object. The context and the procedures for the use of Flow Specifications are as per [RFC9168].

2. Terminology

This document uses the following terms defined in [RFC5440]: PCC, PCE, PCEP Peer.

The following term from [RFC8955] is used frequently throughout this document:

A Flow Specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic. A given IP packet is said to match the defined Flow Specification if it matches all the specified criteria.

Its usage in PCEP is further clarified in [RFC9168].

This document uses the terms "stateful PCE" and "active PCE" as advocated in [RFC7399].

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.

3. L2 Flow Specifications

As per [RFC9168], to carry Flow Specifications in PCEP messages, a PCEP object called the PCEP FLOWSPEC object is defined. To describe a traffic flow, a PCEP TLV called the Flow Filter TLV is also defined. This document extends the support for L2 flow specifications by creating a new PCEP TLV called L2 Flow Filter TLV and updating the processing rules.

The PCEP FLOWSPEC object carries a FlowSpec filter rule encoded in a TLV. To describe a traffic flow based on both L3 and L2 fields, a new L2 Flow Filter TLV is introduced by this document. The PCEP FLOWSPEC object could carry one of the following combinations of TLVs:

- * no TLV
- * one Flow Filter TLV
- * one L2 Flow Filter TLV
- * both a Flow Filter TLV and an L2 Flow Filter TLV

At most one L2 Flow Filter TLV MAY be included in the PCEP FLOWSPEC object. The TLV is OPTIONAL when the R (remove) bit is set in the object. At least one Flow Filter TLV or one L2 Flow Filter TLV MUST be present when the R bit is clear. If both TLVs are missing when the R bit is clear, the PCEP peer MUST respond with a PCErr message with error-type TBD1 (FlowSpec Error) and error-value 2 (Malformed FlowSpec). A Flow Filter TLV and an L2 Flow Filter TLV MAY be present when filtering is based on L3 and L2 fields.

The TLV follows the format of all PCEP TLVs as defined in [RFC5440]. The Type field values come from the codepoint space for PCEP TLVs and has the value TBD2. The value field of L2 Flow Filter TLV contains one or more sub-TLVs (Section 3.1, and they represent the complete definition of a Flow Specification for traffic to be placed on the tunnel. The set of Flow Specification TLVs and L2 Flow Filter TLVs in a single instance of a Flow Filter TLV are combined to indicate the specific Flow Specification. Note that the PCEP FLOWSPEC object can include just one Flow Filter TLV, just one L2 Flow Filter TLV, or one of each TLV.

The rest of the procedures are same as [RFC9168].

3.1. L2 Flow Specification TLVs

The L2 Flow Filter TLV carries one or more L2 Flow Specification TLV. The L2 Flow Specification TLV follows the format of all PCEP TLVs as defined in [RFC5440]. However, the Type values are selected from a separate IANA registry (see Section 6.2) rather than from the common PCEP TLV registry.

Type values are chosen so that there can be commonality with L2 Flow Specifications defined for use with BGP [I-D.ietf-idr-flowspec-l2vpn]. This is possible because the BGP Flow Spec encoding uses a single octet to encode the type whereas PCEP uses two octets. Thus the space of values for the Type field is partitioned as shown in Figure 1.

Range	
0 .. 255	Per BGP registry defined by [I-D.ietf-idr-flowspec-l2vpn]. Not to be allocated in this registry.
256 .. 65535	New PCEP Flow Specifications allocated according to the registry defined in this document.

Figure 1: L2 Flow Specification TLV Type Ranges

[I-D.ietf-idr-flowspec-l2vpn] is the reference for the registry "L2 Flow Spec Component Types" and defines the allocations it contains.

The content of the Value field in each TLV is specific to the type and describes the parameters of the Flow Specification. The definition of the format of many of these Value fields is inherited from BGP specifications. Specifically, the inheritance is from [I-D.ietf-idr-flowspec-l2vpn], but may also be inherited from future BGP specifications.

When multiple L2 Flow Specification TLVs are present in a single L2 Flow Filter TLV they are combined to produce a more detailed specification of a flow. Similarly, when both Flow Filter TLV and L2 Flow Filter TLV are present, they are combined to produce a more detailed specification of a flow.

An implementation that receives a PCEP message carrying a L2 Flow Specification TLV with a type value that it does not recognize or does not support MUST respond with a PCErr message with error-type TBD1 (FlowSpec Error), error-value 1 (Unsupported FlowSpec) and MUST NOT install the Flow Specification.

All L2 Flow Specification TLVs with Types in the range 0 to 255 have their Values interpreted as defined for use in BGP (for example, in [I-D.ietf-idr-flowspec-l2vpn]) and are set using the BGP encoding, but without the type octet (the relevant information is in the Type field of the TLV). The Value field is padded with trailing zeros to achieve 4-byte alignment.

This document defines no new types.

4. BGP Flow Specification Version 2

As per [I-D.ietf-idr-flowspec-v2], Flow Specification v2 allows the user to order the flow specification rules and the actions associated with a rule. Each FSv2 rule may have one or more match conditions and one or more associated actions. It further lists the rules and principles to keep filters in a deterministic order between FSv1 and FSv2.

Note that this document relies on the processing rules as per [RFC9168]. A future PCEP specification could consider updating rules to align to [I-D.ietf-idr-flowspec-v2] (FSv2 adds explicit "order" for instance).

5. Update to RFC 9168

[RFC9168] created the "PCEP Flow Specification TLV Type Indicators" registry and set the assignment policies for the range "256-64506" to "Specification Required". This memo changes the policy from Specification Required to IETF Review in alignment with the rest of the registries in the "Path Computation Element Protocol (PCEP) Numbers" registry group.

Note that [RFC9168] did not follow the guidelines for "Specification Required" as per [RFC8126].

6. IANA Considerations

IANA maintains the "Path Computation Element Protocol (PCEP) Numbers" registry. This document requests IANA actions to allocate code points for the protocol elements defined in this document.

6.1. PCEP TLV Type Indicators

IANA maintains a registry called "PCEP TLV Type Indicators" under the "Path Computation Element Protocol (PCEP) Numbers" registry group. IANA is requested to make an assignment from this registry as follows:

Value	Meaning	Reference
TBD2	L2 FLOW FILTER TLV	[This.I-D]

6.2. L2 Flow Specification TLV Type Indicators

IANA is requested to create a new registry called the "PCEP L2 Flow Specification TLV Type Indicators" registry.

Allocations from this registry are to be made according to the following assignment policies [RFC8126]:

Range	Assignment policy
0 .. 255	Reserved - must not be allocated. Usage mirrors the BGP L2 FlowSpec registry [I-D.ietf-idr-flowspec-l2vpn].
256 .. 64506	IETF Review
64507 .. 65531	First Come First Served
65532 .. 65535	Experimental

This document makes no allocations in the newly created registry.

6.3. Flow Specification TLV Type Indicators

[RFC9168] created the "PCEP Flow Specification TLV Type Indicators" registry. IANA is requested to update the assignment policies for the range "256-64506" from "Specification Required" to "IETF Review" [RFC8126].

7. Implementation Status

[NOTE TO RFC EDITOR : This whole section and the reference to RFC 7942 is to be removed before publication as an RFC]

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".

At the time of posting the -05 version of this document, there are no known implementations of this mechanism. It is believed that some vendors are considering prototype implementations, but these plans are too vague to make any further assertions.

8. Security Considerations

We may assume that a system that utilizes a remote PCE is subject to a number of vulnerabilities that could allow spurious LSPs or SR paths to be established or that could result in existing paths being modified or torn down. Such systems, therefore, apply security considerations as described in [RFC5440], Section 2.5 of [RFC6952], [RFC8253], and [RFC8955].

As per [RFC9168], the description of Flow Specifications associated with paths set up or controlled by a PCE add a further detail that could be attacked without tearing down LSPs or SR paths, but causing traffic to be misrouted within the network. Therefore, the use of the security mechanisms for PCEP referenced above is important. It further lists the security considerations with respect to flow specifications which are applicable to L2 flowspec as well.

9. Manageability Considerations

9.1. Control of Function and Policy

[RFC9168] describe the management of multiple flowspecs as well as control via configurations and policies. This is applicable to the L2 flowspec defined in this document.

9.2. Information and Data Models

As per [RFC9168], the PCEP YANG module [I-D.ietf-pce-pcep-yang] would need to be augmented to cover flowspec include L2.

9.3. 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].

9.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC9168].

9.5. Requirements On Other Protocols

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

9.6. Impact On Network Operations

The use of the features described in this document clearly has an important impact on network traffic since they cause traffic to be routed on specific paths in the network. However, in practice, these changes make no direct changes to the network operation because traffic is already placed on those paths using some pre-existing configuration mechanism. Thus, the significant change is the reduction in the mechanisms that have to be applied, rather than a change to how the traffic is passed through the network.

10. Acknowledgements

Thanks to Susan Hares for the discussion related to BGP Flowspec V2.

11. References

11.1. Normative References

- [I-D.ietf-idr-flowspec-l2vpn]
Weiguo, H., Eastlake, D. E., Litkowski, S., and S. Zhuang,
"BGP Dissemination of L2 Flow Specification Rules", Work
in Progress, Internet-Draft, draft-ietf-idr-flowspec-
l2vpn-26, 23 September 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-idr-
flowspec-l2vpn-26](https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-l2vpn-26)>.
- [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>>.
- [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>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M.
Bacher, "Dissemination of Flow Specification Rules",
RFC 8955, DOI 10.17487/RFC8955, December 2020,
<<https://www.rfc-editor.org/info/rfc8955>>.

- [RFC9168] Dhody, D., Farrel, A., and Z. Li, "Path Computation Element Communication Protocol (PCEP) Extension for Flow Specification", RFC 9168, DOI 10.17487/RFC9168, January 2022, <<https://www.rfc-editor.org/info/rfc9168>>.

11.2. Informative References

- [I-D.ietf-idr-flowspec-v2]
Hares, S., Eastlake, D. E., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-v2-04, 28 April 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-v2-04>>.
- [I-D.ietf-pce-pcep-yang]
Dhody, D., Beeram, V. P., Hardwick, J., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-30, 26 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-yang-30>>.
- [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>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <<https://www.rfc-editor.org/info/rfc6952>>.
- [RFC7399] Farrel, A. and D. King, "Unanswered Questions in the Path Computation Element Architecture", RFC 7399, DOI 10.17487/RFC7399, October 2014, <<https://www.rfc-editor.org/info/rfc7399>>.
- [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>>.
- [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>>.

- [RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", RFC 8283, DOI 10.17487/RFC8283, December 2017, <<https://www.rfc-editor.org/info/rfc8283>>.
- [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>>.

Appendix A. Contributors

Shankara
India

Email: shankara.odl@gmail.com

Qiandeng Liang
Huawei Technologies
101 Software Avenue,
Yuhuatai District
Nanjing
210012
China

Email: liangqiandeng@huawei.com

Cyril Margaria
Juniper Networks
200 Somerset Corporate Boulevard, Suite 4001
Bridgewater, NJ
08807
USA

Email: cmargaria@juniper.net

Colby Barth
Juniper Networks
200 Somerset Corporate Boulevard, Suite 4001
Bridgewater, NJ
08807
USA

Email: cbarth@juniper.net

Xia Chen
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China

Email: jescia.chenxia@huawei.com

Shunwan Zhuang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China

Email: zhuangshunwan@huawei.com

Cheng Li
Huawei Technologies

Email: c.l@huawei.com

Authors' Addresses

Dhruv Dhody
Huawei
India
Email: dhruv.ietf@gmail.com

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: lizhenbin@huawei.com