

PCE
Internet-Draft
Intended status: Standards Track
Expires: 21 June 2026

Y. Liu
ZTE Corporation
18 December 2025

PCEP Extension for Tunnelled Flow Specification
draft-liu-pce-pcep-tunnel-flowspec-00

Abstract

Traffic flows may be categorized and described using "Flow Specifications". RFC8955 defines the Flow Specification and describes how Flow Specification components are used to describe traffic flows. RFC8955 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.

The extensions defined in this document extend the support for tunneled traffic filtering rules.

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 21 June 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 | 2 |
| 2. Terminology | 3 |
| 2.1. Requirements Language | 4 |
| 3. Tunneled Flow Specifications | 4 |
| 3.1. Tunneled Flow Specification TLVs | 5 |
| 4. IANA Considerations | 6 |
| 4.1. PCEP TLV Type Indicators | 6 |
| 4.2. Tunneled Flow Specification TLV Type Indicators | 6 |
| 5. Security Considerations | 7 |
| 6. Manageability Considerations | 7 |
| 6.1. Control of Function and Policy | 7 |
| 6.2. Information and Data Models | 7 |
| 6.3. Liveness Detection and Monitoring | 7 |
| 6.4. Verify Correct Operations | 7 |
| 6.5. Requirements On Other Protocols | 8 |
| 6.6. Impact On Network Operations | 8 |
| 7. References | 8 |
| 7.1. Normative References | 8 |
| 7.2. Informative References | 8 |
| Author's Address | 9 |

1. Introduction

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-pce-pcep-l2-flowspec] further extends the support for Ethernet Layer 2 (L2) and Layer 2 Virtual Private Network (L2VPN) traffic filtering rules in PCEP Flow Specifications.

[I-D.ietf-idr-flowspec-nvo3] defines a BGP flowspec extension to disseminate tunneled traffic filtering rules and flow specification components are specified for certain tunneling header fields.

This document extends the same support for PCEP by defining a new Tunneled 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].

2.1. 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.

3. Tunneled 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 tunneled flow specifications by creating a new PCEP TLV called Tunneled 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 the fields of the tunnel encapsulation, a new Tunneled Flow Filter TLV is introduced by this document. The PCEP FLOWSPEC object could carry no TLV or any combination of its TLVs, i.e., Flow Filter TLV[RFC9168], L2 Flow Filter TLV[I-D.ietf-pcep-l2-flowspec] and Tunneled Flow Filter TLV (this document), and TLV of the same type can only appear at most once in the object.

At most one Tunneled Flow Filter TLV MAY be included in the PCEP FLOWSPEC object. The TLV is OPTIONAL when the R (remove) bit [RFC9168] is set in the object. At least one Flow Filter TLV or one L2 Flow Filter TLV or one Tunneled Flow Filter TLV MUST be present when the R bit is clear. If all the TLVs are missing when the R bit is clear, the PCEP peer MUST respond with a PCerr message with Error-Type 30 (FlowSpec Error) and Error-value 2 (Malformed FlowSpec). When filtering is based on the tunneling header fields and the L3 or L2 fields of the flow, a Flow Filter TLV or an L2 Flow Filter TLV (or both of them) MAY be present together with the Tunneled Flow Filter TLV.

The Tunneled 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 TBA1. The value field of Tunneled Flow Filter TLV contains one or more sub-TLVs (Section 3.1), and they are specified for certain tunneling header fields.

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

3.1. Tunneled Flow Specification TLVs

The Tunneled Flow Filter TLV carries one or more Tunneled Flow Specification TLVs. The Tunneled 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 4.2) rather than from the common PCEP TLV registry.

Type values are chosen so that there can be commonality with Tunneled Flow Specifications defined for use with BGP [I-D.ietf-idr-flowspec-nvo3]. 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-nvo3]. Not to be allocated in this registry. |
| 256 .. 65535 | New PCEP Flow Specifications allocated according to the registry defined in this document. |

Figure 1: Tunneled Flow Specification TLV Type Ranges

[I-D.ietf-idr-flowspec-nvo3] is the reference for the registry "Tunneled 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-nvo3], but may also be inherited from future BGP specifications.

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

An implementation that receives a PCEP message carrying a Tunneled 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 30 (FlowSpec Error) and Error-value 1 (Unsupported FlowSpec) and MUST NOT install the Flow Specification.

All Tunneled 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-nvo3]) 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. 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.

4.1. 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 |
|-------|--------------------------|---------------|
| TBA1 | Tunneled FLOW FILTER TLV | this document |

4.2. Tunneled Flow Specification TLV Type Indicators

IANA is requested to create a new registry called the "PCEP Tunneled 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 Tunnelled FlowSpec registry [I-D.ietf-idr-flowspec-nvo3]. |
| 256 .. 64506 | IETF Review |
| 64507 .. 65531 | First Come First Served |
| 65532 .. 65535 | Experimental |

This document makes no allocations in the newly created registry.

5. Security Considerations

The security considerations in [RFC9168] apply to this document as well.

No new security issues are introduced to the PCEP protocol by this specification.

6. Manageability Considerations

6.1. 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 Tunnelled flowspec defined in this document.

6.2. Information and Data Models

The PCEP YANG module [RFC9826] would need to be augmented to cover tunnelled flowspec.

6.3. 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].

6.4. Verify Correct Operations

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

6.5. Requirements On Other Protocols

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

6.6. Impact On Network Operations

Mechanisms defined in this document do not have any new impact on network operations in addition to those already listed in [RFC9168].

7. References

7.1. Normative References

- [I-D.ietf-idr-flowspec-nvo3]
Eastlake, D. E., Weiguo, H., Zhuang, S., Li, Z., and R. Gu, "BGP Dissemination of Flow Specification Rules for Tunneled Traffic", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-nvo3-23, 5 December 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-nvo3-23>>.
- [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>>.
- [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>>.

7.2. Informative 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>>.

[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-l2-flowspec]

Dhody, D., Farrel, A., and Z. Li, "PCEP Extension for Layer 2 (L2) Flow Specification", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-l2-flowspec-09, 30 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-l2-flowspec-09>>.

[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>>.

[RFC9826] Dhody, D., Ed., Beeram, V., Hardwick, J., and J. Tantsura, "A YANG Data Model for the Path Computation Element Communication Protocol (PCEP)", RFC 9826, DOI 10.17487/RFC9826, September 2025, <<https://www.rfc-editor.org/info/rfc9826>>.

Author's Address

Yao Liu
ZTE Corporation
China
Email: liu.yao71@zte.com.cn