

Internet Engineering Task Force (IETF)
Request for Comments: 9168
Category: Standards Track
ISSN: 2070-1721

D. Dhody
Huawei Technologies
A. Farrel
Old Dog Consulting
Z. Li
Huawei Technologies
January 2022

Path Computation Element Communication Protocol (PCEP) Extension for Flow Specification

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.

This document specifies a set of extensions to PCEP to support 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 include the creation, update, and withdrawal of Flow Specifications via PCEP and can be applied to tunnels initiated by the PCE or to tunnels where control is delegated to the PCE by the Path Computation Client (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.

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/rfc9168>.

Copyright Notice

Copyright (c) 2022 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.	Terminology
3.	Procedures for PCE Use of Flow Specifications
3.1.	Context for PCE Use of Flow Specifications
3.2.	Elements of the Procedure
3.2.1.	Capability Advertisement
3.2.1.1.	PCEP Open Message
3.2.1.2.	IGP PCE Capabilities Advertisement
3.2.2.	Dissemination Procedures
3.2.3.	Flow Specification Synchronization
4.	PCE FlowSpec Capability TLV
5.	PCEP FLOWSPEC Object
6.	Flow Filter TLV
7.	Flow Specification TLVs
8.	Detailed Procedures
8.1.	Default Behavior and Backward Compatibility
8.2.	Composite Flow Specifications
8.3.	Modifying Flow Specifications
8.4.	Multiple Flow Specifications
8.5.	Adding and Removing Flow Specifications
8.6.	VPN Identifiers
8.7.	Priorities and Overlapping Flow Specifications
9.	PCEP Messages
10.	IANA Considerations
10.1.	PCEP Objects
10.1.1.	PCEP FLOWSPEC Object Flag Field
10.2.	PCEP TLV Type Indicators
10.3.	Flow Specification TLV Type Indicators
10.4.	PCEP Error Codes
10.5.	PCE Capability Flag
11.	Security Considerations
12.	Manageability Considerations
12.1.	Management of Multiple Flow Specifications
12.2.	Control of Function through Configuration and Policy
12.3.	Information and Data Models
12.4.	Liveness Detection and Monitoring
12.5.	Verifying Correct Operation
12.6.	Requirements for Other Protocols and Functional Components
12.7.	Impact on Network Operation
13.	References
13.1.	Normative References
13.2.	Informative References
	Acknowledgements
	Contributors
	Authors' Addresses

1. Introduction

[RFC4655] defines the Path Computation Element (PCE), 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 state about 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.

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.

This document specifies a set of extensions to PCEP to support dissemination of Flow Specification components. We term the description of a traffic flow using Flow Specification components as a "Flow Specification". This term is conceptually the same as the term used in [RFC8955]; however, no mechanism is provided to distribute an action associated with the Flow Specification because there is only one action that is applicable in the PCEP context (that is, directing the matching traffic to the identified LSP).

The extensions defined in this document include the creation, update, and withdrawal of Flow Specifications via PCEP and 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.

Flow Specifications are carried in TLVs within a new object called the FLOWSPEC object defined in this document. The flow filtering rules indicated by the Flow Specifications are mainly defined by BGP Flow Specifications.

Note that PCEP-installed Flow Specifications are intended to be installed only at the head end of the LSP to which they direct traffic. It is acceptable (and potentially desirable) that other routers in the network have Flow Specifications installed that match the same traffic but direct it onto different routes or to different LSPs. Those other Flow Specifications may be installed using the PCEP extensions defined in this document, distributed using BGP per [RFC8955], or configured using manual operations. Since this

document is about PCEP-installed Flow Specifications, those other Flow Specifications at other routers are out of scope. In this context, however, it is worth noting that changes to the wider routing system (such as the distribution and installation of BGP Flow Specifications, or fluctuations in the IGP link state database) might mean that traffic matching the PCEP Flow Specification never reaches the head end of the LSP at which the PCEP Flow Specification has been installed. This may or may not be desirable according to the operator's traffic engineering and routing policies and is particularly applicable at LSPs that do not have their head ends at the ingress edge of the network, but it is not an effect that this document seeks to address.

2. Terminology

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.

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

[RFC8955] also states that "[a] given Flow Specification may be associated with a set of attributes" and that "...attributes can be used to encode a set of predetermined actions." However, in the context of this document, no action is explicitly specified as associated with the Flow Specification since the action of forwarding all matching traffic onto the associated path is implicit.

How an implementation decides to filter traffic that matches a Flow Specification does not form part of this specification, but a flag is provided to indicate whether the sender of a PCEP message that includes a Flow Specification intends it to be installed as a Longest Prefix Match (LPM) route or as a Flow Specification policy.

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

3. Procedures for PCE Use of Flow Specifications

3.1. Context for PCE Use of Flow Specifications

In the PCE architecture, there are five steps in the setup and use of LSPs:

1. Decide which LSPs to set up. The decision may be made by a user, by a PCC, or by the PCE. There can be a number of triggers for this, including user intervention and dynamic response to changes in traffic demands.
2. Decide what properties to assign to an LSP. This can include bandwidth reservations, priorities, and the Differentiated Services Code Point (DSCP) (i.e., MPLS Traffic Class field). This function is also determined by user configuration or in response to predicted or observed traffic demands.

3. Decide what traffic to put on the LSP. This is effectively determining which traffic flows to assign to which LSPs; practically, this is closely linked to the first two decisions listed above.
4. Cause the LSP to be set up and modified to have the right characteristics. This will usually involve the PCE advising or instructing the PCC at the head end of the LSP, and the PCC will then signal the LSP across the network.
5. Tell the head end of the LSP what traffic to put on the LSP. This may happen after or at the same time as the LSP is set up. This step is the subject of this document.

3.2. Elements of the Procedure

There are three elements in the procedure:

1. A PCE and a PCC must be able to indicate whether or not they support the use of Flow Specifications.
2. A PCE or PCC must be able to include Flow Specifications in PCEP messages with a clear understanding of the applicability of those Flow Specifications in each case. This includes whether the use of such information is mandatory, constrained, or optional and how overlapping Flow Specifications will be resolved.
3. Flow Specification information/state must be synchronized between PCEP peers so that, on recovery, the peers have the same understanding of which Flow Specifications apply just as is required in the case of stateful PCE and LSP delegation (see Section 5.6 of [RFC8231]).

The following subsections describe these points.

3.2.1. Capability Advertisement

As with most PCEP capability advertisements, the ability to support Flow Specifications can be indicated in the PCEP Open message or in IGP PCE capability advertisements.

3.2.1.1. PCEP Open Message

During PCEP session establishment, a PCC or PCE that supports the procedures described in this document announces this fact by including the PCE FlowSpec Capability TLV (described in Section 4) in the OPEN object carried in the PCEP Open message.

The presence of the PCE FlowSpec Capability TLV in the OPEN object in a PCE's Open message indicates that the PCE can distribute FlowSpecs to PCCs and can receive FlowSpecs in messages from PCCs.

The presence of the PCE FlowSpec Capability TLV in the OPEN object in a PCC's Open message indicates that the PCC supports the FlowSpec functionality described in this document.

If either one of a pair of PCEP peers does not include the PCE FlowSpec Capability TLV in the OPEN object in its Open message, then the other peer MUST NOT include a FLOWSPEC object in any PCEP message sent to the peer. If a FLOWSPEC object is received when support has not been indicated, the receiver will respond with a PCErr message reporting the objects containing the FlowSpec as described in [RFC5440]: that is, it will use "Unknown Object" if it does not support this specification and "Not supported object" if it supports this specification but has not chosen to support FLOWSPEC objects on this PCEP session.

3.2.1.2. IGP PCE Capabilities Advertisement

The ability to advertise support for PCEP and PCE features in IGP advertisements is provided for OSPF in [RFC5088] and for IS-IS in [RFC5089]. The mechanism uses the PCE Discovery TLV, which has a PCE-CAP-FLAGS sub-TLV containing bit flags, each of which indicates support for a different feature.

This document defines a new PCE-CAP-FLAGS sub-TLV bit, the FlowSpec Capable flag (bit number 16). Setting the bit indicates that an advertising PCE supports the procedures defined in this document.

Note that while PCE FlowSpec capability may be advertised during discovery, PCEP speakers that wish to use Flow Specification in PCEP MUST negotiate PCE FlowSpec capability during PCEP session setup, as specified in Section 3.2.1.1. A PCC MAY initiate PCE FlowSpec capability negotiation at PCEP session setup even if it did not receive any IGP PCE capability advertisement, and a PCEP peer that advertised support for FlowSpec in the IGP is not obliged to support these procedures on any given PCEP session.

3.2.2. Dissemination Procedures

This section describes the procedures to support Flow Specifications in PCEP messages.

The primary purpose of distributing Flow Specification information is to allow a PCE to indicate to a PCC what traffic it should place on a path (such as an LSP or an SR path). This means that the Flow Specification may be included in:

- * PCInitiate messages so that an active PCE can indicate the traffic to place on a path at the time that the PCE instantiates the path.
- * PCUpd messages so that an active PCE can indicate or change the traffic to place on a path that has already been set up.
- * PCRpt messages so that a PCC can report the traffic that the PCC will place on the path.
- * PCReq messages so that a PCC can indicate what traffic it plans to place on a path when it requests that the PCE perform a computation in case that information aids the PCE in its work.
- * PCRep messages so that a PCE that has been asked to compute a path can suggest which traffic could be placed on a path that a PCC may be about to set up.
- * PCErr messages so that issues related to paths and the traffic they carry can be reported to the PCE by the PCC and problems with other PCEP messages that carry Flow Specifications can be reported.

To carry Flow Specifications in PCEP messages, this document defines a new PCEP object called the "PCEP FLOWSPEC object". The object is OPTIONAL in the messages described above and MAY appear more than once in each message.

To describe a traffic flow, the PCEP FLOWSPEC object carries a Flow Filter TLV.

The inclusion of multiple PCEP FLOWSPEC objects allows multiple traffic flows to be placed on a single path.

Once a PCE and PCC have established that they can both support the

use of Flow Specifications in PCEP messages, such information may be exchanged at any time for new or existing paths.

The application and prioritization of Flow Specifications are described in Section 8.7.

As per [RFC8231], any attributes of the path received from a PCE are subject to the PCC's local policy. This holds true for the Flow Specifications as well.

3.2.3. Flow Specification Synchronization

The Flow Specifications are carried along with the LSP state information as per [RFC8231], making the Flow Specifications part of the LSP database (LSP-DB). Thus, the synchronization of the Flow Specification information is done as part of LSP-DB synchronization. This may be achieved using normal state synchronization procedures as described in [RFC8231] or enhanced state synchronization procedures as defined in [RFC8232].

The approach selected will be implementation and deployment specific and will depend on issues such as how the databases are constructed and what level of synchronization support is needed.

4. PCE FlowSpec Capability TLV

The PCE-FLOWSPEC-CAPABILITY TLV is an optional TLV that can be carried in the OPEN object [RFC5440] to exchange the PCE FlowSpec capabilities of the PCEP speakers.

The format of the PCE-FLOWSPEC-CAPABILITY TLV follows the format of all PCEP TLVs as defined in [RFC5440] and is shown in Figure 1.




Figure 1: PCE-FLOWSPEC-CAPABILITY TLV Format

The type of the PCE-FLOWSPEC-CAPABILITY TLV is 51, and it has a fixed length of 2 octets. The Value field MUST be set to 0 and MUST be ignored on receipt. The two bytes of padding MUST be set to zero and ignored on receipt.

The inclusion of this TLV in an OPEN object indicates that the sender can perform FlowSpec handling as defined in this document.

5. PCEP FLOWSPEC Object

The PCEP FLOWSPEC object defined in this document is compliant with the PCEP object format defined in [RFC5440]. It is OPTIONAL in the PCReq, PCRep, PCErr, PCInitiate, PCRpt, and PCUpd messages and MAY be present zero, one, or more times. Each instance of the object specifies a separate traffic flow.

The PCEP FLOWSPEC object MAY carry a FlowSpec filter rule encoded in a Flow Filter TLV as defined in Section 6.

The FLOWSPEC Object-Class is 43 (to be assigned by IANA).

The FLOWSPEC Object-Type is 1.

The format of the body of the PCEP FLOWSPEC object is shown in Figure 2.




Figure 2: PCEP FLOWSPEC Object Body Format

FS-ID (32 bits): A PCEP-specific identifier for the FlowSpec information. A PCE or PCC creates an FS-ID for each FlowSpec that it originates, and the value is unique within the scope of that PCE or PCC and is constant for the lifetime of a PCEP session. All subsequent PCEP messages can identify the FlowSpec using the FS-ID. The values 0 and 0xFFFFFFFF are reserved and MUST NOT be used. Note that [NUMERIC-IDS-SEC] gives advice on assigning transient numeric identifiers such as the FS-ID so as to minimize security risks.

AFI (16 bits): Address Family Identifier as used in BGP [RFC4760]
(AFI=1 for IPv4 or VPNv4, AFI=2 for IPv6 and VPNv6 as per
[RFC8956]).

Reserved (8 bits): MUST be set to zero on transmission and ignored on receipt.

Flags (8 bits): Two flags are currently assigned:

R bit: The Remove bit is set when a PCEP FLOWSPEC object is included in a PCEP message to indicate removal of the Flow Specification from the associated tunnel. If the bit is clear, the Flow Specification is being added or modified.

L bit: The Longest Prefix Match (LPM) bit is set to indicate that the Flow Specification is to be installed as a route subject to LPM forwarding. If the bit is clear, the Flow Specification described by the Flow Filter TLV (see Section 6) is to be installed as a Flow Specification. If the bit is set, only Flow Specifications that describe IPv4 or IPv6 destinations are meaningful in the Flow Filter TLV, and others are ignored. If the L is set and the receiver does not support the use of Flow Specifications that are present in the Flow Filter TLV for the installation of a route subject to LPM forwarding, then the PCEP peer MUST respond with a PCErr message with Error-Type 30 (FlowSpec Error) and Error-value 5 (Unsupported LPM Route).

Unassigned bits MUST be set to zero on transmission and ignored on receipt.

If the PCEP speaker receives a message with the R bit set in the FLOWSPEC object and the Flow Specification identified with an FS-ID does not exist, it **MUST** generate a PCErr with Error-Type 30 (FlowSpec Error) and Error-value 4 (Unknown FlowSpec).

If the PCEP speaker does not understand or support the AFI in the FLOWSPEC message, the PCEP peer MUST respond with a PCErr message

with Error-Type 30 (FlowSpec Error) and Error-value 2 (Malformed FlowSpec).

The following TLVs can be used in the FLOWSPEC object:

Speaker Entity Identifier TLV: As specified in [RFC8232], the SPEAKER-ENTITY-ID TLV encodes a unique identifier for the node that does not change during the lifetime of the PCEP speaker. This is used to uniquely identify the FlowSpec originator and thus is used in conjunction with the FS-ID to uniquely identify the FlowSpec information. This TLV MUST be included. If the TLV is missing, the PCEP peer MUST respond with a PCerr message with Error-Type 30 (FlowSpec Error) and Error-value 2 (Malformed FlowSpec). If more than one instance of this TLV is present, the first MUST be processed, and subsequent instances MUST be ignored.

Flow Filter TLV (variable): One TLV MAY be included. The Flow Filter TLV is OPTIONAL when the R bit is set.

The Flow Filter TLV MUST be present when the R bit is clear. If the TLV is 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).

Filtering based on the L2 fields is out of scope of this document.

6. Flow Filter TLV

One new PCEP TLV is defined to convey Flow Specification filtering rules that specify what traffic is carried on a path. The TLV follows the format of all PCEP TLVs as defined in [RFC5440]. The Type field values come from the code point space for PCEP TLVs and has the value 52 for Flow Filter TLV.

The Value field of the TLV contains one or more sub-TLVs (the Flow Specification TLVs) as defined in Section 7, and they represent the complete definition of a Flow Specification for traffic to be placed on the tunnel. This tunnel is indicated by the PCEP message in which the PCEP FLOWSPEC object is carried. The set of Flow Specification TLVs in a single instance of a Flow Filter TLV is combined to indicate the specific Flow Specification. Note that the PCEP FLOWSPEC object can include just one Flow Filter TLV.

Further Flow Specifications can be included in a PCEP message by including additional FLOWSPEC objects.

In the future, there may be a desire to add support for L2 Flow Specifications (such as described in [BGP-L2VPN]).

7. Flow Specification TLVs

The Flow Filter TLV carries one or more Flow Specification TLVs. The 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 10.3) rather than from the common PCEP TLV registry.

Type values are chosen so that there can be commonality with Flow Specifications defined for use with BGP [RFC8955] [RFC8956]. This is possible because the BGP Flow Spec encoding uses a single octet to encode the type, whereas PCEP uses 2 octets. Thus, the space of values for the Type field is partitioned as shown in Table 1.

+=====+	
Range	Description
+=====+	

0-255	Per BGP Flow Spec registry defined by [RFC8955] and [RFC8956].
	Not to be allocated in this registry.
256-65535	New PCEP Flow Specifications allocated according to the registry defined in this document.

Table 1: Flow Specification TLV Type Ranges

[RFC8955] is the reference for the "Flow Spec Component Types" registry and defines the allocations it contains. [RFC8956] requested the creation of the "Flow Spec IPv6 Component Types" registry, as well as its initial allocations. If the AFI (in the FLOWSPEC object) is set to IPv4, the range 0..255 is as per "Flow Spec Component Types" [RFC8955]; if the AFI is set to IPv6, the range 0..255 is as per "Flow Spec IPv6 Component Types" [RFC8956].

The content of the Value field in each TLV is specific to the type/AFI 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 [RFC8955] and [RFC8956], but it may also be inherited from future BGP specifications.

When multiple Flow Specification TLVs are present in a single Flow Filter TLV, they are combined to produce a more detailed specification of a flow. For examples and rules about how this is achieved, see [RFC8955]. As described in [RFC8955], where it says "A given component type MAY (exactly once) be present in the Flow Specification", a Flow Filter TLV MUST NOT contain more than one Flow Specification TLV of the same type: an implementation that receives a PCEP message with a Flow Filter TLV that contains more than one Flow Specification TLV of the same type MUST respond with a PCerr message with Error-Type 30 (FlowSpec Error) and Error-value 2 (Malformed FlowSpec) and MUST NOT install the Flow Specification.

An implementation that receives a PCEP message carrying a Flow Specification TLV with a type value that it does not recognize or 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.

When used in other protocols (such as BGP), these Flow Specifications are also associated with actions to indicate how traffic matching the Flow Specification should be treated. In PCEP, however, the only action is to associate the traffic with a tunnel and to forward matching traffic onto that path, so no encoding of an action is needed.

Section 8.7 describes how overlapping Flow Specifications are prioritized and handled.

All Flow Specification TLVs with Types in the range 0 to 255 have values defined for use in BGP (for example, in [RFC8955] and [RFC8956]) 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 the following new types:

Type	Description	Value Defined In
------	-------------	------------------

256	Route Distinguisher	RFC 9168
257	IPv4 Multicast Flow	RFC 9168
258	IPv6 Multicast Flow	RFC 9168

Table 2: Flow Specification TLV Types Defined in this Document

To allow identification of a VPN in PCEP via a Route Distinguisher (RD) [RFC4364], a new TLV, ROUTE-DISTINGUISHER TLV, is defined in this document. A Flow Specification TLV with Type 256 (ROUTE-DISTINGUISHER TLV) carries an RD value, which is used to identify that other flow filter information (for example, an IPv4 destination prefix) is associated with a specific VPN identified by the RD. See Section 8.6 for further discussion of VPN identification.




Figure 3: The Format of the ROUTE-DISTINGUISHER TLV

The format of the RD is as per [RFC4364].

Although it may be possible to describe a multicast Flow Specification from the combination of other Flow Specification TLVs with specific values, it is more convenient to use a dedicated Flow Specification TLV. Flow Specification TLVs with Type values 257 and 258 are used to identify a multicast flow for IPv4 and IPv6, respectively. The Value field is encoded as shown in Figure 4.

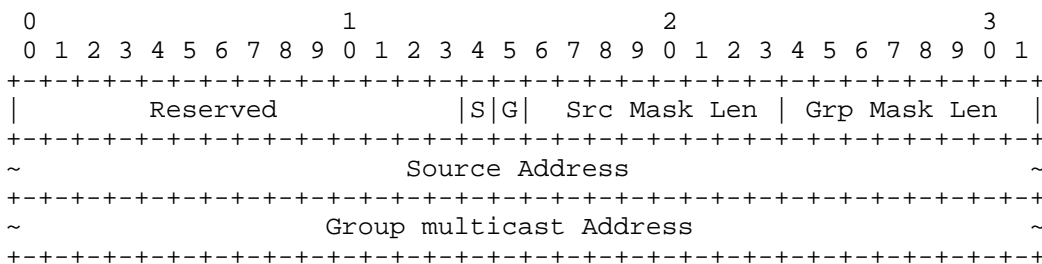


Figure 4: Multicast Flow Specification TLV Encoding

The address fields and address mask lengths of the two Multicast Flow Specification TLVs contain source and group prefixes for matching against packet flows. Note that the two address fields are 32 bits for an IPv4 Multicast Flow and 128 bits for an IPv6 Multicast Flow.

The Reserved field MUST be set to zero and ignored on receipt.

Two bit flags (S and G) are defined to describe the multicast wildcarding in use. If the S bit is set, then source wildcarding is in use, and the values in the Source Mask Length and Source Address fields MUST be ignored. If the G bit is set, then group wildcarding is in use, and the values in the Group Mask Length and Group multicast Address fields MUST be ignored. The G bit MUST NOT be set

unless the S bit is also set: if a Multicast Flow Specification TLV is received with S bit = 0 and G bit = 1, the receiver MUST respond with a PCERR with Error-Type 30 (FlowSpec Error) and Error-value 2 (Malformed FlowSpec).

The three multicast mappings may be achieved as follows:

(S, G) - S bit = 0, G bit = 0, the Source Address and Group multicast Address prefixes are both used to define the multicast flow.

(*, G) - S bit = 1, G bit = 0, the Group multicast Address prefix is used to define the multicast flow, but the Source Address prefix is ignored.

(*, *) - S bit = 1, G bit = 1, the Source Address and Group multicast Address prefixes are both ignored.

8. Detailed Procedures

This section outlines some specific detailed procedures for using the protocol extensions defined in this document.

8.1. Default Behavior and Backward Compatibility

The default behavior is that no Flow Specification is applied to a tunnel. That is, the default is that the FLOWSPEC object is not used, as is the case in all systems before the implementation of this specification.

In this case, it is a local matter (such as through configuration) how tunnel head ends are instructed in terms of what traffic to place on a tunnel.

[RFC5440] describes how receivers respond when they see unknown PCEP objects.

8.2. Composite Flow Specifications

Flow Specifications may be represented by a single Flow Specification TLV or may require a more complex description using multiple Flow Specification TLVs. For example, a flow indicated by a source-destination pair of IPv6 addresses would be described by the combination of Destination IPv6 Prefix and Source IPv6 Prefix Flow Specification TLVs.

8.3. Modifying Flow Specifications

A PCE may want to modify a Flow Specification associated with a tunnel, or a PCC may want to report a change to the Flow Specification it is using with a tunnel.

It is important to identify the specific Flow Specification so it is clear that this is a modification of an existing flow and not the addition of a new flow as described in Section 8.4. The FS-ID field of the PCEP FLOWSPEC object is used to identify a specific Flow Specification in the context of the content of the Speaker Entity Identifier TLV.

When modifying a Flow Specification, all Flow Specification TLVs for the intended specification of the flow MUST be included in the PCEP FLOWSPEC object. The FS-ID MUST be retained from the previous description of the flow, and the same Speaker Entity Identifier TLV MUST be used.

8.4. Multiple Flow Specifications

It is possible that traffic from multiple flows will be placed on a single tunnel. In some cases, it is possible to define these within a single PCEP FLOWSPEC object by widening the scope of a Flow Specification TLV: for example, traffic to two destination IPv4 prefixes might be captured by a single Flow Specification TLV with type "Destination" with a suitably adjusted prefix. However, this is unlikely to be possible in most scenarios, and it must be recalled that it is not permitted to include two Flow Specification TLVs of the same type within one Flow Filter TLV.

The normal procedure, therefore, is to carry each Flow Specification in its own PCEP FLOWSPEC object. Multiple objects may be present on a single PCEP message, or multiple PCEP messages may be used.

8.5. Adding and Removing Flow Specifications

The Remove bit in the PCEP FLOWSPEC object is left clear when a Flow Specification is being added or modified.

To remove a Flow Specification, a PCEP FLOWSPEC object is included with the FS-ID matching the one being removed, and the R bit is set to indicate removal. In this case, it is not necessary to include any Flow Specification TLVs.

If the R bit is set and Flow Specification TLVs are present, an implementation MAY ignore them. If the implementation checks the Flow Specification TLVs against those recorded for the FS-ID and Speaker Entity Identifier of the Flow Specification being removed and finds a mismatch, the Flow Specification matching the FS-ID MUST still be removed, and the implementation SHOULD record a local exception or log.

8.6. VPN Identifiers

VPN instances are identified in BGP using RDs [RFC4364]. These values are not normally considered to have any meaning outside of the network, and they are not encoded in data packets belonging to the VPNs. However, RDs provide a useful way of identifying VPN instances and are often manually or automatically assigned to VPNs as they are provisioned.

Thus, the RD provides a useful way to indicate that traffic for a particular VPN should be placed on a given tunnel. The tunnel head end will need to interpret this Flow Specification not as a filter on the fields of data packets but rather using the other mechanisms that it already uses to identify VPN traffic. These mechanisms could be based on the incoming port (for port-based VPNs) or may leverage knowledge of the VPN Routing and Forwarding (VRF) that is in use for the traffic.

8.7. Priorities and Overlapping Flow Specifications

Flow Specifications can overlap. For example, two different Flow Specifications may be identical except for the length of the prefix in the destination address. In these cases, the PCC must determine how to prioritize the Flow Specifications so as to know which path to assign packets that match both Flow Specifications. That is, the PCC must assign a precedence to the Flow Specifications so that it checks each incoming packet for a match in a predictable order.

The processing of BGP Flow Specifications is described in [RFC8955]. Section 5.1 of that document explains the order of traffic filtering rules to be executed by an implementation of that specification.

PCCs MUST apply the same ordering rules as defined in [RFC8955].

Furthermore, it is possible that Flow Specifications will be distributed by BGP as well as by PCEP as described in this document. In such cases, implementations supporting both approaches MUST apply the prioritization and ordering rules as set out in [RFC8955] regardless of which protocol distributed the Flow Specifications. However, implementations MAY provide a configuration control to allow one protocol to take precedence over the other; this may be particularly useful if the Flow Specifications make identical matches on traffic but have different actions. It is RECOMMENDED that a message be logged for the operator to understand the behavior when two Flow Specifications distributed by different protocols overlap, especially when one acts to replace another.

Section 12.1 of this document covers manageability considerations relevant to the prioritized ordering of Flow Specifications.

An implementation that receives a PCEP message carrying a Flow Specification that it cannot resolve against other Flow Specifications already installed (for example, because the new Flow Specification has irresolvable conflicts with other Flow Specifications that are already installed) MUST respond with a PCErr message with Error-Type 30 (FlowSpec Error) and Error-value 3 (Unresolvable Conflict) and MUST NOT install the Flow Specification.

9. PCEP Messages

This section describes the format of messages that contain FLOWSPEC objects. The only difference from previous message formats is the inclusion of that object.

The figures in this section use the notation defined in [RFC5511].

The FLOWSPEC object is OPTIONAL and MAY be carried in the PCEP messages.

The PCInitiate message is defined in [RFC8281] and updated as below:

```
<PCInitiate Message> ::= <Common Header>
                          <PCE-initiated-lsp-list>
```

Where:

```
<PCE-initiated-lsp-list> ::= <PCE-initiated-lsp-request>
                              [<PCE-initiated-lsp-list>]

<PCE-initiated-lsp-request> ::=
    ( <PCE-initiated-lsp-instantiation> |
      <PCE-initiated-lsp-deletion> )

<PCE-initiated-lsp-instantiation> ::= <SRP>
                                       <LSP>
                                       [<END-POINTS>]
                                       <ERO>
                                       [<attribute-list>]
                                       [<flowspec-list>]
```

Where:

```
<flowspec-list> ::= <FLOWSPEC> [<flowspec-list>]
```

The PCUpd message is defined in [RFC8231] and updated as below:

```
<PCUpd Message> ::= <Common Header>
                    <update-request-list>
```

Where:

```
<update-request-list> ::= <update-request>
```

[<update-request-list>]

<update-request> ::= <SRP>
 <LSP>
 <path>
 [<flowspec-list>]

Where:

<path> ::= <intended-path><intended-attribute-list>

<flowspec-list> ::= <FLOWSPEC> [<flowspec-list>]

The PCRpt message is defined in [RFC8231] and updated as below:

<PCRpt Message> ::= <Common Header>
 <state-report-list>

Where:

<state-report-list> ::= <state-report> [<state-report-list>]

<state-report> ::= [<SRP>]
 <LSP>
 <path>
 [<flowspec-list>]

Where:

<path> ::= <intended-path>
 [<actual-attribute-list><actual-path>]
 <intended-attribute-list>

<flowspec-list> ::= <FLOWSPEC> [<flowspec-list>]

The PCReq message is defined in [RFC5440] and updated in [RFC8231];
it is further updated below for a Flow Specification:

<PCReq Message> ::= <Common Header>
 [<svec-list>]
 <request-list>

Where:

<svec-list> ::= <SVEC> [<svec-list>]

<request-list> ::= <request> [<request-list>]

<request> ::= <RP>
 <END-POINTS>
 [<LSP>]
 [<LSPA>]
 [<BANDWIDTH>]
 [<metric-list>]
 [<RRO> [<BANDWIDTH>]]
 [<IRO>]
 [<LOAD-BALANCING>]
 [<flowspec-list>]

Where:

<flowspec-list> ::= <FLOWSPEC> [<flowspec-list>]

The PCRep message is defined in [RFC5440] and updated in [RFC8231];
it is further updated below for a Flow Specification:

<PCRep Message> ::= <Common Header>
 <response-list>

Where:

<response-list> ::= <response> [<response-list>]

```

<response>::=<RP>
    [<LSP>]
    [<NO-PATH>]
    [<attribute-list>]
    [<path-list>]
    [<flowspec-list>]

```

Where:

```

<flowspec-list> ::= <FLOWSPEC> [<flowspec-list>]

```

10. IANA Considerations

This document requests that IANA allocate code points for the protocol elements defined in this document.

10.1. PCEP Objects

IANA maintains a subregistry called "PCEP Objects" within the "Path Computation Element Protocol (PCEP) Numbers" registry. Each PCEP object has an Object-Class and an Object-Type, and IANA has allocated new code points in this subregistry as follows:

Object-Class Value	Name	Object-Type	Reference
43	FLOWSPEC	0: Reserved	RFC 9168
		1: Flow Specification	RFC 9168

Table 3: PCEP Objects Subregistry Additions

10.1.1. PCEP FLOWSPEC Object Flag Field

This document requests that a new subregistry, "FLOWSPEC Object Flag Field", be created within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flag field of the FLOWSPEC object. New values are to be assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit number (counting from bit 0 as the most significant bit)
- * Capability description
- * Defining RFC

The initial population of this registry is as follows:

Bit	Description	Reference
0-5	Unassigned	
6	LPM (L bit)	RFC 9168
7	Remove (R bit)	RFC 9168

Table 4: Initial Contents of the
FLOWSPEC Object Flag Field
Registry

10.2. PCEP TLV Type Indicators

IANA maintains a subregistry called "PCEP TLV Type Indicators" within the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA has made the following allocations in this subregistry:

Value	Description	Reference
51	PCE-FLOWSPEC-CAPABILITY TLV	RFC 9168
52	FLOW FILTER TLV	RFC 9168

Table 5: PCEP TLV Type Indicators Subregistry Additions

10.3. Flow Specification TLV Type Indicators

IANA has created a new subregistry called "PCEP Flow Specification TLV Type Indicators" within the "Path Computation Element Protocol (PCEP) Numbers" registry.

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

Range	Registration Procedures
0-255	Reserved - must not be allocated. Usage mirrors the BGP Flow Spec registry [RFC8955] [RFC8956].
256-64506	Specification Required
64507-65531	First Come First Served
65532-65535	Experimental Use

Table 6: Registration Procedures for the PCEP Flow Specification TLV Type Indicators Subregistry

IANA has populated this registry with values defined in this document as follows, taking the new values from the range 256 to 64506:

Value	Meaning
256	Route Distinguisher
257	IPv4 Multicast
258	IPv6 Multicast

Table 7: Initial Contents of the PCEP Flow Specification TLV Type Indicators Subregistry

10.4. PCEP Error Codes

IANA maintains a subregistry called "PCEP-ERROR Object Error Types and Values" within the "Path Computation Element Protocol (PCEP) Numbers" registry. Entries in this subregistry are described by

Error-Type and Error-value. IANA has added the following assignment to this subregistry:

Error-Type	Meaning	Error-value	Reference
30	FlowSpec error	0: Unassigned	RFC 9168
		1: Unsupported FlowSpec	RFC 9168
		2: Malformed FlowSpec	RFC 9168
		3: Unresolvable Conflict	RFC 9168
		4: Unknown FlowSpec	RFC 9168
		5: Unsupported LPM Route	RFC 9168
		6-255: Unassigned	RFC 9168

Table 8: PCEP-ERROR Object Error Types and Values Subregistry Additions

10.5. PCE Capability Flag

IANA has registered a new capability bit in the OSPF Parameters "Path Computation Element (PCE) Capability Flags" registry as follows:

Bit	Capability Description	Reference
16	FlowSpec	RFC 9168

Table 9: Path Computation Element (PCE) Capability Flags Registry Additions

11. 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].

The description of Flow Specifications associated with paths set up or controlled by a PCE adds a further detail that could be attacked without tearing down LSPs or SR paths but causes traffic to be misrouted within the network. Therefore, the use of the security mechanisms for PCEP referenced above is important.

Visibility into the information carried in PCEP does not have direct privacy concerns for end users' data; however, knowledge of how data is routed in a network may make that data more vulnerable. Of course, the ability to interfere with the way data is routed also makes the data more vulnerable. Furthermore, knowledge of the connected endpoints (such as multicast receivers or VPN sites) is usually considered private customer information. Therefore,

implementations or deployments concerned with protecting privacy MUST apply the mechanisms described in the documents referenced above, in particular, to secure the PCEP session using IPsec per Sections 10.4 to 10.6 of [RFC5440] or TLS per [RFC8253]. Note that TCP-MD5 security as originally suggested in [RFC5440] does not provide sufficient security or privacy guarantees and SHOULD NOT be relied upon.

Experience with Flow Specifications in BGP systems indicates that they can become complex and that the overlap of Flow Specifications installed in different orders can lead to unexpected results. Although this is not directly a security issue per se, the confusion and unexpected forwarding behavior may be engineered or exploited by an attacker. Furthermore, this complexity might give rise to a situation where the forwarding behaviors might create gaps in the monitoring and inspection of particular traffic or provide a path that avoids expected mitigations. Therefore, implementers and operators SHOULD pay careful attention to the manageability considerations described in Section 12 and familiarize themselves with the careful explanations in [RFC8955].

12. Manageability Considerations

The feature introduced by this document enables operational manageability of networks operated in conjunction with a PCE and using PCEP. In the case of a stateful active PCE or with PCE-initiated services, in the absence of this feature, additional manual configuration is needed to tell the head ends what traffic to place on the network services (LSPs, SR paths, etc.).

This section follows the advice and guidance of [RFC6123].

12.1. Management of Multiple Flow Specifications

Experience with Flow Specification in BGP suggests that there can be a lot of complexity when two or more Flow Specifications overlap. This can arise, for example, with addresses indicated using prefixes and could cause confusion about what traffic should be placed on which path. Unlike the behavior in a distributed routing system, it is not important to the routing stability and consistency of the network that each head-end implementation applies the same rules to disambiguate overlapping Flow Specifications, but it is important that:

- * a network operator can easily find out what traffic is being placed on which path and why. This will facilitate analysis of the network and diagnosis of faults.
- * a PCE be able to correctly predict the effect of instructions it gives to a PCC. This will ensure that traffic is correctly placed on the network without causing congestion or other network inefficiencies and that traffic is correctly delivered.

To that end, a PCC MUST enable an operator to view the Flow Specifications that it has installed, and these MUST be presented in order of precedence such that when two Flow Specifications overlap, the one that will be serviced with higher precedence is presented to the operator first.

A discussion of precedence ordering for Flow Specifications is found in Section 8.7.

12.2. Control of Function through Configuration and Policy

Support for the function described in this document implies that a functional element that is capable of requesting that a PCE compute

and control a path is also able to configure the specification of what traffic should be placed on that path. Where there is a human involved in this action, configuration of the Flow Specification must be available through an interface (such as a graphical user interface or a Command Line Interface). Where a distinct software component (i.e., one not co-implemented with the PCE) is used, a protocol mechanism will be required that could be PCEP itself or a data model, such as extensions to the YANG model for requesting path computation [TEAS-YANG-PATH].

Implementations MAY be constructed with a configurable switch to indicate whether they support the functions defined in this document. Otherwise, such implementations MUST indicate that they support the function as described in Section 4. If an implementation allows configurable support of this function, that support MAY be configurable per peer or once for the whole implementation.

As mentioned in Section 12.1, a PCE implementation SHOULD provide a mechanism to configure variations in the precedence ordering of Flow Specifications per PCC.

12.3. Information and Data Models

The YANG model in [PCE-PCEP-YANG] can be used to model and monitor PCEP states and messages. To make that YANG model useful for the extensions described in this document, it would need to be augmented to cover the new protocol elements.

Similarly, as noted in Section 12.2, the YANG model defined in [TEAS-YANG-PATH] could be extended to allow the specification of Flow Specifications.

Finally, as mentioned in Section 12.1, a PCC implementation SHOULD provide a mechanism to allow an operator to read the Flow Specifications from a PCC and to understand in what order they will be executed. This could be achieved using a new YANG model.

12.4. Liveness Detection and Monitoring

The extensions defined in this document do not require any additional liveness detection and monitoring support. See [RFC5440] and [RFC5886] for more information.

12.5. Verifying Correct Operation

The chief element of operation that needs to be verified (in addition to the operation of the protocol elements as described in [RFC5440]) is the installation, precedence, and correct operation of the Flow Specifications at a PCC.

In addition to the YANG model, for reading Flow Specifications described in Section 12.3, tools may be needed to inject Operations and Management (OAM) traffic at the PCC that matches specific criteria so that it can be monitored while traveling along the desired path. Such tools are outside the scope of this document.

12.6. Requirements for Other Protocols and Functional Components

This document places no requirements on other protocols or components.

12.7. Impact on Network Operation

The use of the features described in this document clearly have 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 mechanisms that have to be applied rather than a change to how the traffic is passed through the network.

13. References

13.1. 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>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [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>>.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, DOI 10.17487/RFC5511, April 2009, <<https://www.rfc-editor.org/info/rfc5511>>.
- [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>>.
- [RFC8232] Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", RFC 8232, DOI 10.17487/RFC8232, September 2017, <<https://www.rfc-editor.org/info/rfc8232>>.
- [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>>.

[RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed.,
"Dissemination of Flow Specification Rules for IPv6",
RFC 8956, DOI 10.17487/RFC8956, December 2020,
<<https://www.rfc-editor.org/info/rfc8956>>.

13.2. Informative References

[BGP-L2VPN]

Hao, W., 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-18, 24 October 2021,
<[https://datatracker.ietf.org/doc/html/draft-ietf-idr-
flowspec-l2vpn-18](https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-l2vpn-18)>.

[NUMERIC-IDS-SEC]

Gont, F. and I. Arce, "Security Considerations for
Transient Numeric Identifiers Employed in Network
Protocols", Work in Progress, Internet-Draft, draft-gont-
numeric-ids-sec-considerations-06, 5 December 2020,
<[https://datatracker.ietf.org/doc/html/draft-gont-numeric-
ids-sec-considerations-06](https://datatracker.ietf.org/doc/html/draft-gont-numeric-ids-sec-considerations-06)>.

[PCE-PCEP-YANG]

Dhody, D., Hardwick, J., Beeram, V. P., and J. Tantsura,
"A YANG Data Model for Path Computation Element
Communications Protocol (PCEP)", Work in Progress,
Internet-Draft, draft-ietf-pce-pcep-yang-17, 23 October
2021, <[https://datatracker.ietf.org/doc/html/draft-ietf-
pce-pcep-yang-17](https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-yang-17)>.

[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>>.

[RFC5088] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R.
Zhang, "OSPF Protocol Extensions for Path Computation
Element (PCE) Discovery", RFC 5088, DOI 10.17487/RFC5088,
January 2008, <<https://www.rfc-editor.org/info/rfc5088>>.

[RFC5089] Le Roux, JL., Ed., Vasseur, JP., Ed., Ikejiri, Y., and R.
Zhang, "IS-IS Protocol Extensions for Path Computation
Element (PCE) Discovery", RFC 5089, DOI 10.17487/RFC5089,
January 2008, <<https://www.rfc-editor.org/info/rfc5089>>.

[RFC5886] Vasseur, JP., Ed., Le Roux, JL., and Y. Ikejiri, "A Set of
Monitoring Tools for Path Computation Element (PCE)-Based
Architecture", RFC 5886, DOI 10.17487/RFC5886, June 2010,
<<https://www.rfc-editor.org/info/rfc5886>>.

[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/info/rfc6123>>.

[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>>.

- [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>>.
- [TEAS-YANG-PATH] Busi, I., Belotti, S., Lopez, V., Sharma, A., and Y. Shi, "YANG Data Model for requesting Path Computation", Work in Progress, Internet-Draft, draft-ietf-teas-yang-path-computation-16, 6 September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-path-computation-16>>.

Acknowledgements

Thanks to Julian Lucek, Sudhir Cheruathur, Olivier Dugeon, Jayant Agarwal, Jeffrey Zhang, Acee Lindem, Vishnu Pavan Beeram, Julien Meuric, Deborah Brungard, ric Vyncke, Erik Kline, Benjamin Kaduk, Martin Duke, Roman Danyliw, and Alvaro Retana for useful discussions and comments.

Contributors

Shankara
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India

Email: shankara@huawei.com

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

Email: liangqiandeng@huawei.com

Cyril Margaria
Juniper Networks
200 Somerset Corporate Boulevard, Suite 4001
Bridgewater, NJ 08807
United States of America

Email: cmargaria@juniper.net

Colby Barth
Juniper Networks
200 Somerset Corporate Boulevard, Suite 4001
Bridgewater, NJ 08807
United States of America

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
Huawei Campus, No. 156 Beiqing Rd.
Beijing, 100095
China

Email: c.l@huawei.com

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India

Email: dhruv.ietf@gmail.com

Adrian Farrel
Old Dog Consulting

Email: adrian@olddog.co.uk

Zhenbin Li
Huawei Technologies
Huawei Bldg., No. 156 Beiqing Rd.
Beijing
100095
China

Email: lizhenbin@huawei.com