

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 22, 2020

S. Litkowski
S. Sivabalan
Cisco Systems, Inc.
J. Tantsura
Apstra, Inc.
J. Hardwick
Metaswitch Networks
M. Negi
C. Li
Huawei Technologies
April 20, 2020

**Path Computation Element communication Protocol (PCEP) extension for
associating Policies and Label Switched Paths (LSPs)
draft-ietf-pce-association-policy-09**

Abstract

This document introduces a simple mechanism to associate policies to a group of Label Switched Paths (LSPs) via an extension to the Path Computation Element (PCE) Communication Protocol (PCEP).

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 October 22, 2020.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	Terminology	3
3.	Motivation	4
3.1.	Policy based Constraints	5
4.	Overview	6
5.	Policy Association Group	7
5.1.	Policy Parameters TLV	7
6.	Implementation Status	8
6.1.	Cisco's Implementation	9
7.	Security Considerations	9
8.	IANA Considerations	10
8.1.	Association object Type Indicators	10
8.2.	PCEP TLV Type Indicators	10
9.	Manageability Considerations	10
9.1.	Control of Function and Policy	10
9.2.	Information and Data Models	10
9.3.	Liveness Detection and Monitoring	11
9.4.	Verify Correct Operations	11
9.5.	Requirements on Other Protocols	11
9.6.	Impact on Network Operations	11
10.	Acknowledgments	11
11.	References	11
11.1.	Normative References	11
11.2.	Informative References	12
Appendix A.	Contributor Addresses	14
	Authors' Addresses	14

[1.](#) Introduction

[RFC5440] describes the Path Computation Element communication Protocol (PCEP) which enables the communication between a Path Computation Client (PCC) and a Path Control Element (PCE), or between two PCEs based on the PCE architecture [RFC4655]. [RFC5394] provides additional details on policy within the PCE architecture and also provides context for the support of PCE Policy.

PCEP Extensions for Stateful PCE Model [RFC8231] describes a set of extensions to PCEP to enable active control of Multiprotocol Label

Switching Traffic Engineering (MPLS-TE) and Generalized MPLS (GMPLS) tunnels. [RFC8281] describes the set-up and teardown of PCE-initiated LSPs under the active stateful PCE model, without the need for local configuration on the PCC, thus allowing for a dynamic network. Currently, the LSPs can either be signalled via Resource Reservation Protocol Traffic Engineering (RSVP-TE) or can be segment routed as specified in [RFC8664].

[RFC8697] introduces a generic mechanism to create a grouping of LSPs which can then be used to define associations between a set of LSPs and a set of attributes (such as configuration parameters or behaviours) and is equally applicable to stateful PCE (active and passive modes) and stateless PCE.

This document specifies a PCEP extension to associate one or more LSPs with policies using the generic association mechanism.

A PCEP speaker may want to influence the PCEP peer with respect to path selection and other policies. This document describes a PCEP extension to associate policies by creating Policy Association Group (PAG) and encoding this association in PCEP messages. The specification is applicable to both stateful and stateless PCEP sessions.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

The following terminology is used in this document.

Association parameters: As described in [RFC8697], the combination of the mandatory fields Association type, Association ID and Association Source in the ASSOCIATION object uniquely identify the association group. If the optional TLVs - Global Association Source or Extended Association ID are included, then they are included in combination with mandatory fields to uniquely identify the association group.

Association information: As described in [RFC8697], the ASSOCIATION object could include other optional TLVs based on the association types, that provides 'information' related to the association.

LSR: Label Switch Router.

MPLS: Multiprotocol Label Switching.

PAG: Policy Association Group.

PAT: Policy Association Type.

PCC: Path Computation Client. Any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP: Path Computation Element Communication Protocol.

3. Motivation

Paths computed using PCE can be subjected to various policies on both PCE and PCC. For example, in a centralized traffic engineering scenario, network operators may instantiate LSPs and specifies policies for traffic steering, path monitoring, etc., for some LSPs via the Stateful PCE. Similarly, a PCC could request a user- or service-specific policy to be applied at the PCE, such as constraints relaxation to meet optimal QoS and resiliency.

PCEP speaker can use the generic mechanism as per [[RFC8697](#)] to associate a set of LSPs with a policy, without the need to know the details of such a policy, which simplifies network operations, avoids frequent software upgrades, as well as provides an ability to introduce new policy faster.

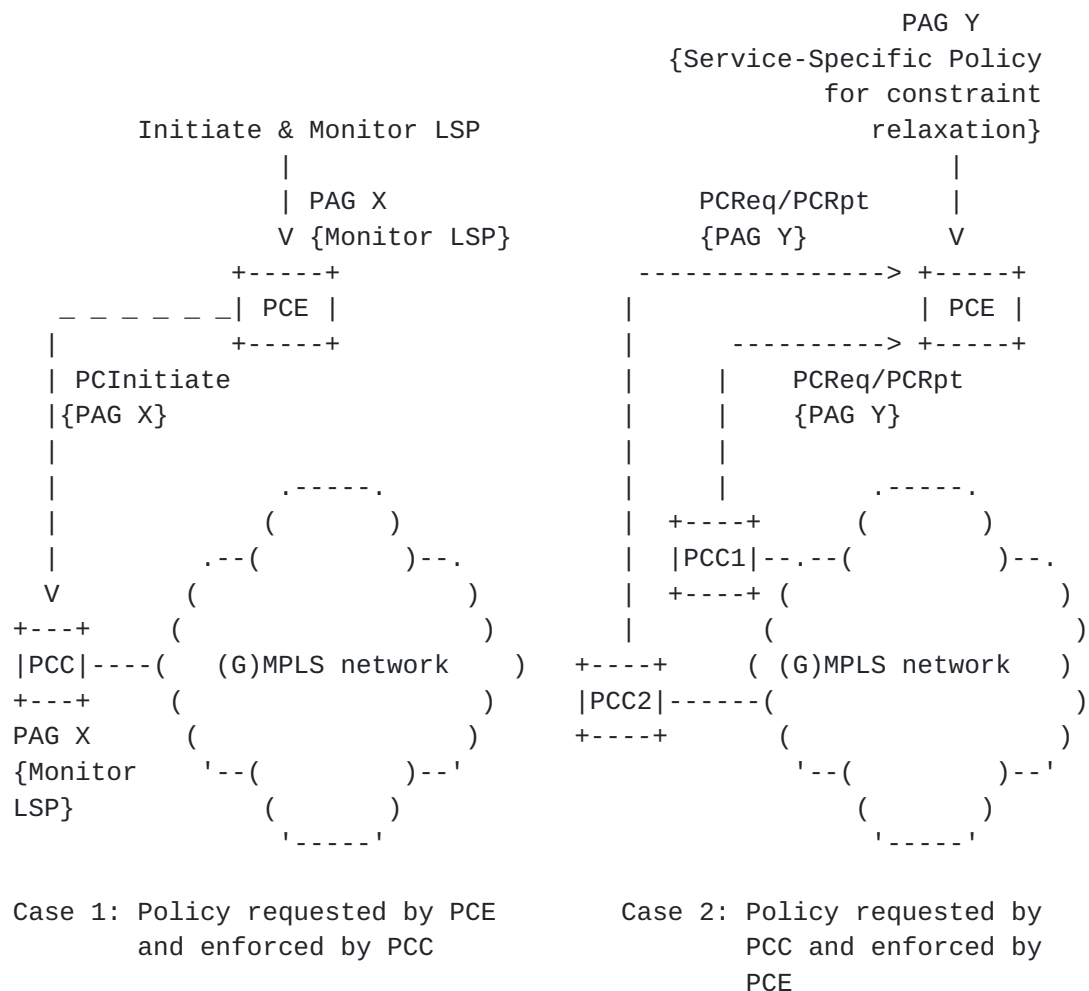


Figure 1: Sample use-cases for carrying policies over PCEP session

3.1. Policy based Constraints

In the context of policy-enabled path computation [[RFC5394](#)], path computation policies may be applied at both a PCC and a PCE. Consider an Label Switch Router (LSR) with a policy enabled PCC, it receives a service request via signalling, including over a Network-Network Interface (NNI) or User Network Interface (UNI) reference point, or receives a configuration request over a management interface to establish a service. The PCC may also apply user- or service-specific policies to decide how the path selection process should be constrained, that is, which constraints, diversities, optimization criterion, and constraint relaxation strategies should be applied in order for the service LSP(s) to have a likelihood to be successfully established and provide necessary QoS and resilience against network failures. The user- or service-specific policies

applied to PCC and are then passed to the PCE along with the Path computation request, in the form of constraints [[RFC5394](#)].

PCEP speaker can use the generic mechanism as per [[RFC8697](#)] to associate a set of LSPs with policy and its resulting path computation constraints. This would simplify the path computation message exchanges in PCEP.

4. Overview

As per [[RFC8697](#)], LSPs are associated with other LSPs with which they interact by adding them to a common association group. Grouping can also be used to define association between LSPs and policies associated to them. One new Association type is defined in this document, based on the generic Association object -

- o Association type = TBD1 ("Policy Association Type (PAT)") for Policy Association Group (PAG).

[[RFC8697](#)] specify the mechanism for the capability advertisement of the Association types supported by a PCEP speaker by defining a ASSOC-Type-List TLV to be carried within an OPEN object. This capability exchange for the association type described in this document (i.e. PAT) MUST be done before using the policy association. Thus the PCEP speaker MUST include the PAT (TBD1) in the ASSOC-Type-List TLV before using the PAG in the PCEP messages.

This Association type is operator-configured association in nature and created by the operator manually on the PCEP peers. An LSP belonging to this association is conveyed via PCEP messages to the PCEP peer. Operator-configured Association Range need not be set for this association-type, and MUST be ignored, so that the full range of association identifier can be utilized.

A PAG can have one or more LSPs and its associated policy. The association parameters including association identifier, Association type (Policy), as well as the association source IP address is manually configured by the operator and is used to identify the PAG as described in [[RFC8697](#)]. The Global Association Source and Extended Association ID MAY also be included.

As per the processing rules specified in [section 6.4 of \[\[RFC8697\]\(#\)\]](#), if a PCEP speaker does not support this Policy Association type, it would return a PCErr message with Error-Type 26 "Association Error" and Error-Value 1 "Association type is not supported". Since the PAG is opaque in nature, the PAG and the policy MUST be configured on the PCEP peers as per the operator-configured association procedures. All further processing is as per [section 6.4 of \[\[RFC8697\]\(#\)\]](#). If a PCE

speaker receives PAG in a PCEP message, and the policy association information is not configured, it MUST return a PCErr message with Error-Type 26 "Association Error" and Error-Value 4 "Association unknown". If some of the association information [[RFC8697](#)] (the TLVs defined in this document) received from the peer does not match the local configured values, the PCEP speaker MUST reject the PCEP message and send a PCErr message with Error-Type 26 "Association Error" and Error-Value 5 "Operator-configured association information mismatch".

Associating a particular LSP to multiple policy groups is authorized from a protocol perspective, however there is no assurance that the PCE will be able to apply multiple policies.

5. Policy Association Group

Association groups and their memberships are defined using the ASSOCIATION object defined in [[RFC8697](#)]. Two object types for IPv4 and IPv6 are defined. The ASSOCIATION object includes "Association type" indicating the type of the association group. This document add a new Association type -

Association type = TBD1 ("Policy Association type") for PAG.

PAG may carry optional TLVs including but not limited to -

- o POLICY-PARAMETERS-TLV: Used to communicate opaque information useful to apply the policy, described in [Section 5.1](#).
- o VENDOR-INFORMATION-TLV: Used to communicate arbitrary vendor specific behavioural information, described in [[RFC7470](#)].

5.1. Policy Parameters TLV

The POLICY-PARAMETERS-TLV is an optional TLV that can be carried in ASSOCIATION object (for PAT) to carry opaque information needed to apply the policy at the PCEP peer. In some cases to apply a PCE policy successfully, it is required to also associate some policy parameters that needs to be evaluated, to successfully apply the said policy. This TLV is used to carry those policy parameters. The TLV could include one or more policy related parameter. The encoding format and the order MUST be known to the PCEP peers, this could be done during the configuration of the policy (and its association parameters) for the PAG. The TLV format is as per the format of the PCEP TLVs, as defined in [[RFC5440](#)], and shown in Figure 2. Only one POLICY-PARAMETERS-TLV can be carried and only the first occurrence is processed and any others MUST be ignored.

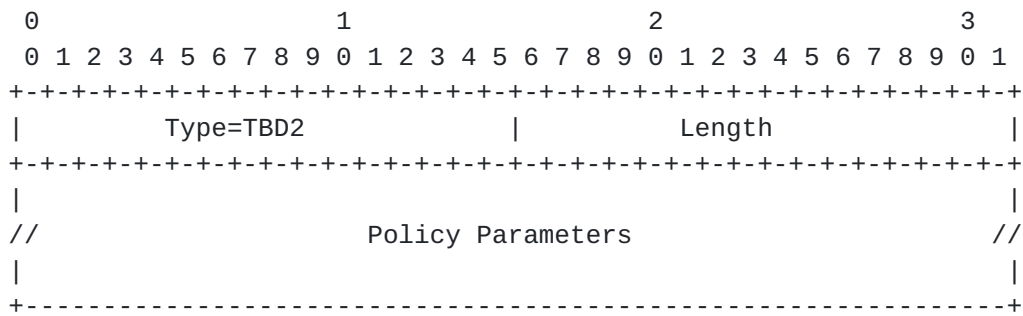


Figure 2: The POLICY-PARAMETERS-TLV format

The type of the POLICY-PARAMETERS-TLV is TBD2 and it has a variable length. The Value field is variable field padded to a 4-bytes alignment; padding is not included in the Length field. The PCEP peer implementation need to be aware of the encoding format, order, and meaning of the 'Policy Parameters' well in advance based on the policy. Note that from the protocol point of view this data is opaque and can be used to carry parameters in any format understood by the PCEP peers and associated to the policy. The exact use of this TLV is beyond the scope of this document.

If the PCEP peer is unaware of the policy parameters associated with the policy and it receives the POLICY-PARAMETERS-TLV, it MUST ignore the TLV and SHOULD log this event. Further, if one or more parameters received in the POLICY-PARAMETERS-TLV received by the PCEP speaker are considered as unacceptable in the context of the associated policy (e.g. out of range value, badly encoded value...), the PCEP speaker MUST NOT apply the received policy and SHOULD log this event.

Note that, the vendor specific behavioural information is encoded in VENDOR-INFORMATION-TLV which can be used along with this TLV.

6. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC 7942](#).]

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 catalogue 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".

6.1. Cisco's Implementation

- o Organization: Cisco Systems, Inc.
- o Implementation: IOS-XR PCE and PCC.
- o Description: The PCEP extension specified in this document is used to convey traffic steering policies.
- o Maturity Level: In shipping product.
- o Coverage: Partial.
- o Contact: msiva@cisco.com.

7. Security Considerations

This document defines one new type for association, which do not add any new security concerns beyond those discussed in [[RFC5440](#)], [[RFC8231](#)] and [[RFC8697](#)] in itself.

Extra care needs to be taken by the implementation with respect to POLICY-PARAMETERS-TLV while decoding, verifying and applying these policy variables. This TLV parsing could be exploited by an attacker.

Some deployments may find policy associations and their implications as extra sensitive and thus securing the PCEP session using Transport Layer Security (TLS) [[RFC8253](#)], as per the recommendations and best current practices in [BCP 195](#) [[RFC7525](#)], is RECOMMENDED.

8. IANA Considerations

8.1. Association object Type Indicators

This document defines a new Association type. The sub-registry "ASSOCIATION Type Field" of the "Path Computation Element Protocol (PCEP) Numbers" registry was originally defined in [[RFC8697](#)]. IANA is requested to make the following allocation.

Value	Name	Reference
TBD1	Policy Association	[This.I-D]

8.2. PCEP TLV Type Indicators

The following TLV Type Indicator value is requested within the "PCEP TLV Type Indicators" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA is requested to make the following allocation.

Value	Description	Reference
TBD2	POLICY-PARAMETERS-TLV	[This.I-D]

9. Manageability Considerations

9.1. Control of Function and Policy

An operator MUST be allowed to configure the policy associations at PCEP peers and associate it with the LSPs. They MAY also allow configuration to related policy parameters, in which case the an operator MUST also be allowed to set the encoding format and order to parse the associated policy parameters TLV.

9.2. Information and Data Models

[RFC7420] describes the PCEP MIB, there are no new MIB Objects for this document.

The PCEP YANG module is defined in [[I-D.ietf-pce-pcep-yang](#)]. This module supports associations as defined in [[RFC8697](#)] and thus support the Policy Association groups.

An implementation SHOULD allow the operator to view the PAG configured. Further implementation SHOULD allow to view associations reported by each peer, and the current set of LSPs in the PAG.

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](#)], [[RFC8231](#)], and [[RFC8281](#)].

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 [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)].

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

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [[RFC5440](#)], [[RFC8231](#)], and [[RFC8281](#)].

10. Acknowledgments

A special thanks to author of [[RFC8697](#)], this document borrow some of the text from it.

11. References

11.1. 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>>.
- [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>>.
- [RFC8697] Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)", [RFC 8697](#), DOI 10.17487/RFC8697, January 2020, <<https://www.rfc-editor.org/info/rfc8697>>.

11.2. Informative References

- [RFC4655] Farrel, A., Vasseur, J., 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>>.
- [RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", [RFC 5394](#), DOI 10.17487/RFC5394, December 2008, <<https://www.rfc-editor.org/info/rfc5394>>.
- [RFC7420] Koushik, A., Stephan, E., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module", [RFC 7420](#), DOI 10.17487/RFC7420, December 2014, <<https://www.rfc-editor.org/info/rfc7420>>.
- [RFC7470] Zhang, F. and A. Farrel, "Conveying Vendor-Specific Constraints in the Path Computation Element Communication Protocol", [RFC 7470](#), DOI 10.17487/RFC7470, March 2015, <<https://www.rfc-editor.org/info/rfc7470>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", [BCP 195](#), [RFC 7525](#), DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [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>>.

- [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>>.
- [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>>.
- [I-D.ietf-pce-pcep-yang] Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", [draft-ietf-pce-pcep-yang-13](#) (work in progress), October 2019.

Appendix A. Contributor Addresses

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

EMail: dhruv.ietf@gmail.com

Qin Wu
Huawei Technologies
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

EMail: sunseawq@huawei.com

Xian Zhang
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R.China

EMail: zhang.xian@huawei.com

Udayasree Palle

EMail: udayasreereddy@gmail.com

Authors' Addresses

Stephane Litkowski
Cisco Systems, Inc.
11 Rue Camille Desmoulins
Issy-les-Moulineaux 92130
France

EMail: slitkows@cisco.com

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

EMail: msiva@cisco.com

Jeff Tantsura
Apstra, Inc.

EMail: jefftant.ietf@gmail.com

Jonathan Hardwick
Metaswitch Networks
100 Church Street
Enfield, Middlesex
UK

EMail: Jonathan.Hardwick@metaswitch.com

Mahendra Singh Negi
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

EMail: mahend.ietf@gmail.com

Cheng Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

EMail: chengli13@huawei.com

