

Internet Engineering Task Force (IETF)
Request for Comments: 9005
Category: Standards Track
ISSN: 2070-1721

S. Litkowski
Cisco Systems, Inc.
S. Sivabalan
Ciena
J. Tantsura
Juniper Networks
J. Hardwick
Metaswitch Networks
李呈 (C. Li)
技术有限公司 (Huawei Technologies)
March 2021

Path Computation Element Communication Protocol (PCEP) Extension for Associating Policies and Label Switched Paths (LSPs)

Abstract

This document introduces a simple mechanism to associate policies with a group of Label Switched Paths (LSPs) via an extension to the Path Computation Element Communication Protocol (PCEP). The extension allows a PCEP speaker to advertise to a PCEP peer that a particular LSP belongs to a particular Policy Association Group (PAG).

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/rfc9005>.

Copyright Notice

Copyright (c) 2021 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
 - 1.1. Requirements Language
2. Terminology
3. Motivation
 - 3.1. Policy-Based Constraints
4. Overview

- 5. Policy Association Group
 - 5.1. POLICY-PARAMETERS-TLV
- 6. Security Considerations
- 7. IANA Considerations
 - 7.1. ASSOCIATION Object Type Indicators
 - 7.2. PCEP TLV Type Indicators
 - 7.3. PCEP Errors
- 8. Manageability Considerations
 - 8.1. Control of Function and Policy
 - 8.2. Information and Data Models
 - 8.3. Liveness Detection and Monitoring
 - 8.4. Verifying Correct Operations
 - 8.5. Requirements on Other Protocols
 - 8.6. Impact on Network Operations
- 9. References
 - 9.1. Normative References
 - 9.2. Informative References
- Appendix A. Example of Policy Parameters
- Acknowledgments
- Contributors
- Authors' Addresses

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.

"Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE" ([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 setup 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 signaled via Resource Reservation Protocol Traffic Engineering (RSVP-TE) or segment routed as specified in [RFC8664].

[RFC8697] introduces a generic mechanism to create a grouping of LSPs that can then be used to define associations between a set of LSPs and a set of attributes (such as configuration parameters or behaviors) 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 a Policy Association Group (PAG) and encoding this association in PCEP messages. The specification is applicable to both stateful and stateless PCEP sessions.

Note that the actual policy definition and the associated parameters are out of scope of this document.

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 identifies 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 provide "information" related to the association.

LSR: Label Switching 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.

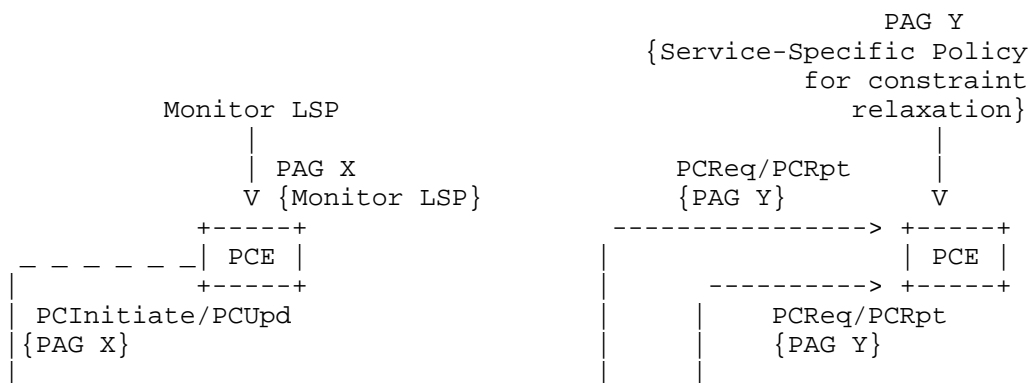
PCE: Path Computation Element; an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP: Path Computation Element Communication Protocol

3. Motivation

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

PCEP speakers can use the generic mechanism of [RFC8697] to associate a set of LSPs with a policy, without the need to know the details of such a policy. This simplifies network operations, avoids frequent software upgrades, and provides the ability to introduce new policies more quickly.



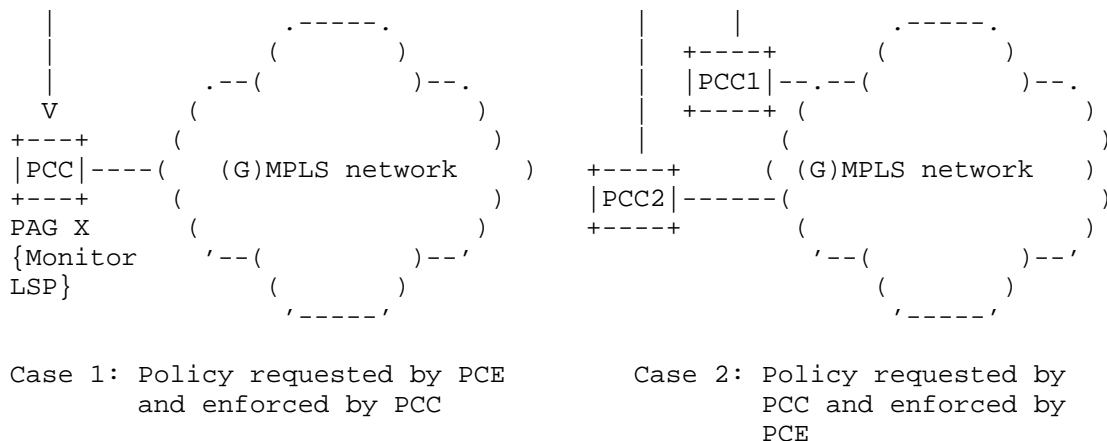


Figure 1: Sample Use Cases for Carrying Policies over PCEP

3.1. Policy-Based Constraints

In the context of a policy-enabled path computation framework [RFC5394], path computation policies may be applied at a PCC, a PCE, or both. A Label Switching Router (LSR) with a policy-enabled PCC can receive:

- * A service request via signaling, including over a Network-Network Interface (NNI) or User-Network Interface (UNI) reference point.
- * A configuration request over a management interface to establish a service.

The PCC may apply user-specific or service-specific policies to decide how the path selection process should be constrained -- that is, which constraints, diversities, optimization criteria, and constraint-relaxation strategies should be applied to increase the likelihood that the service LSP(s) will be successfully established and will provide the necessary QoS and resilience against network failures. The user-specific or service-specific policies are applied to the PCC and are then passed to the PCE along with the path computation request in the form of constraints [RFC5394].

The PCEP speaker can use the generic mechanism as per [RFC8697] to associate a set of LSPs with user-specific or service-specific policies. 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 the association between LSPs and the policies associated with them. As described in [RFC8697], the association group is uniquely identified by the combination of the following fields in the ASSOCIATION object: Association Type, Association ID, Association Source, and (if present) Global Association Source or Extended Association ID. This document defines a new Association Type called "Policy Association" with value 3 based on the generic ASSOCIATION object. This new Association Type is called "Policy Association Type" (PAT).

[RFC8697] specifies the mechanism for the capability advertisement of the Association Types supported by a PCEP speaker by defining an ASSOC-Type-List TLV to be carried within an OPEN object. This capability exchange for the PAT MUST be done before using the Policy Association. Thus, the PCEP speaker MUST include the PAT in the ASSOC-Type-List TLV and MUST receive the same from the PCEP peer

before using the PAG in PCEP messages.

The Policy Association Type (3) is operator configured (as specified in [RFC8697]), i.e., the association is created by the operator manually on the PCEP peers, and an LSP belonging to this association is conveyed via PCEP messages to the PCEP peer. There is no need to convey an explicit Operator-configured Association Range, which could only serve to artificially limit the available Association IDs. Thus, for the Policy Association Type, the Operator-configured Association Range MUST NOT be set and MUST be ignored if received.

A PAG can have one or more LSPs. The association parameters including Association Identifier, Policy Association Type (PAT), as well as the Association Source IP address are manually configured by the operator and are 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 PCEP error (PCErr) message with Error-Type 26 "Association Error" and Error-value 1 "Association type is not supported". 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 a 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".

Associating a particular LSP with multiple policy groups is allowed from a protocol perspective; however, there is no assurance that the PCEP speaker will be able to apply multiple policies. If a PCEP speaker does not support handling of multiple policies for an LSP, it MUST NOT add the LSP into the association group and MUST return a PCErr with Error-Type 26 "Association Error" and Error-value 7 "Cannot join the association group".

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 adds a new Association Type, Policy Association Type (PAT).

PAG may carry optional TLVs including but not limited to:

POLICY-PARAMETERS-TLV:

Used to communicate opaque information useful to applying the policy, described in Section 5.1.

VENDOR-INFORMATION-TLV:

Used to communicate arbitrary vendor-specific behavioral information, described in [RFC7470].

5.1. POLICY-PARAMETERS-TLV

The ASSOCIATION object (for PAT) can carry an optional POLICY-PARAMETERS-TLV with opaque information that is 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 need to be evaluated. This TLV is used to carry those policy parameters. The TLV could include one or more policy-related parameters. 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. Any others MUST be ignored.

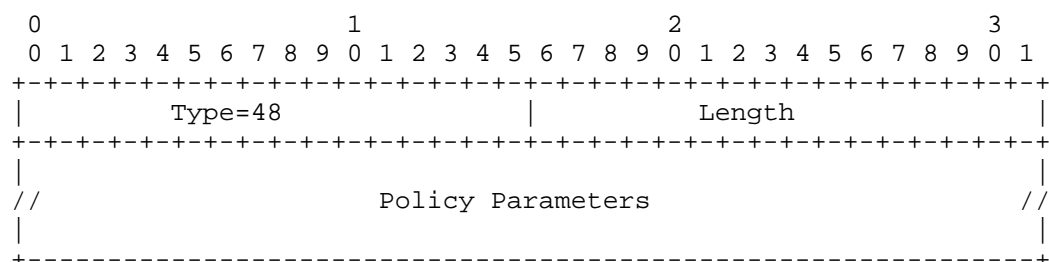


Figure 2: The POLICY-PARAMETERS-TLV Format

The POLICY-PARAMETERS-TLV type is 48, and it has a variable length. The Value field is variable and padded to a 4-byte alignment; padding is not included in the Length field. The PCEP peer implementation needs 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 with the policy. The exact use of this TLV is beyond the scope of this document. Examples are included for illustration purposes in Appendix A.

If the PCEP peer is unaware of the policy parameters associated with the policy and it receives the POLICY-PARAMETERS-TLV, it MUST reject the PCEP message and send a PCerr message with Error-Type 26 "Association Error" and Error-value 12 "Not expecting policy parameters". Further, if at least one parameter in the POLICY-PARAMETERS-TLV received by the PCEP speaker is considered unacceptable in the context of the associated policy (e.g., out of range value, badly encoded value, etc.), the PCEP speaker MUST reject the PCEP message and send a PCerr message with Error-Type 26 "Association Error" and Error-value 13 "Unacceptable policy parameters".

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

6. Security Considerations

The security considerations described in [RFC8697], [RFC8231], [RFC5394], and [RFC5440] apply to the extensions described in this document as well. In particular, a malicious PCEP speaker could be spoofed and used as an attack vector by creating spurious Policy Associations as described in [RFC8697]. Further, as described in [RFC8697], a spurious LSP can have policies that are inconsistent with those of the legitimate LSPs of the group and, thus, cause problems in the handling of the policy for the legitimate LSPs. It should be noted that Policy Association could provide an adversary with the opportunity to eavesdrop on the relationship between the LSPs. [RFC8697] suggests that the implementations and operators use indirect values as a way to hide any sensitive business relationships. 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.

Further, extra care needs to be taken by the implementation with respect to the POLICY-PARAMETERS-TLV while decoding, verifying, and applying these policy variables. This TLV parsing could be exploited by an attacker; thus, extra care must be taken while configuring a Policy Association that uses the POLICY-PARAMETERS-TLV and making

sure that the data is easy to parse and verify before use. Ensuring agreement among all relevant PCEP peers as to the format and layout of the policy parameters information is key for correct operations. Note that the parser for POLICY-PARAMETERS-TLV is particularly sensitive since it is opaque to PCEP and can be used to convey data with many different internal structures/formats. The choice of decoder is dependent on the additional metadata associated with the policy; thus, additional risk of using a wrong decoder and getting garbage results is incurred. Using standard and well-known policy formats could help alleviate those risks.

7. IANA Considerations

7.1. ASSOCIATION Object Type Indicators

This document defines a new Association Type in the subregistry "ASSOCIATION Type Field" of the "Path Computation Element Protocol (PCEP) Numbers" registry that was originally defined in [RFC8697].

Value	Name	Reference
3	Policy Association	RFC 9005

Table 1

7.2. PCEP TLV Type Indicators

The following TLV Type Indicator value has been registered within the "PCEP TLV Type Indicators" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry.

Value	Description	Reference
48	POLICY-PARAMETERS-TLV	RFC 9005

Table 2

7.3. PCEP Errors

This document defines new Error-values for Error-Type 26 "Association Error" defined in [RFC8697]. IANA has allocated new error values within the "PCEP-ERROR Object Error Types and Values" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry as follows:

Error-Type	Meaning	Error-value	Reference
26	Association Error		[RFC8697]
		12: Not expecting policy parameters	RFC 9005
		13: Unacceptable policy parameters	RFC 9005

Table 3

8. Manageability Considerations

8.1. Control of Function and Policy

An operator MUST be allowed to configure the Policy Associations at PCEP peers and associate them with the LSPs. They MAY also allow configuration to related policy parameters and provide information on the encoding format and order to parse the associated POLICY-PARAMETERS-TLV.

8.2. 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 [PCE-PCEP-YANG]. That module supports associations as defined in [RFC8697]; thus, it supports the Policy Association Groups.

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

8.3. Liveness Detection and Monitoring

The mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440] and [RFC8231].

8.4. Verifying Correct Operations

Verifying the correct operation of a policy can be performed by monitoring various parameters as described in [RFC5440] and [RFC8231]. A PCEP implementation SHOULD provide information on failed path computation due to applying policy and log error events, e.g., parsing failure for a POLICY-PARAMETERS-TLV.

8.5. Requirements on Other Protocols

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

8.6. Impact on Network Operations

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

9. References

9.1. 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>>.

[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>>.

[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>>.

9.2. Informative References

[PCE-PCEP-YANG]
Dhody, D., Ed., Hardwick, J., Beeram, V., and J. Tantsura,
"A YANG Data Model for Path Computation Element
Communications Protocol (PCEP)", Work in Progress,
Internet-Draft, draft-ietf-pce-pcep-yang-16, 22 February
2021,
<<https://tools.ietf.org/html/draft-ietf-pce-pcep-yang-16>>.

[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>>.

[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>>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch,
"Network Time Protocol Version 4: Protocol and Algorithms
Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010,
<<https://www.rfc-editor.org/info/rfc5905>>.

[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>>.

[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>>.

Appendix A. Example of Policy Parameters

An example could be a monitoring and telemetry policy, P1, that is dependent on a profile (GOLD/SILVER/BRONZE) as set by the operator. The PCEP peers need to be aware of policy P1 (and its associated characteristics) in advance as well the fact that the policy parameter will encode the profile of a type string in the POLICY-PARAMETERS-TLV. As an example, LSP1 could encode the PAG with the POLICY-PARAMETERS-TLV using the string "GOLD".

The following is another example where the path computation at the PCE could be dependent on when the LSP was configured at the PCC. For such a policy, P2, the timestamp can be encoded in the POLICY-PARAMETERS-TLV, and the exact encoding could be the 64-bit timestamp format as defined in [RFC5905].

While the above example has a single field in the POLICY-PARAMETERS-TLV, it is possible to include multiple fields, but the exact order, encoding format, and meanings need to be known in advance at the PCEP peers.

Acknowledgments

We would like to acknowledge and thank Santiago Alvarez, Zafar Ali, Luis Tomotaki, Victor Lopez, Rob Shakir, and Clarence Filsfils for working on earlier draft versions with similar motivation.

Special thanks to the authors of [RFC8697]. This document borrowed some of its text. The authors would like to thank Aijun Wang, Peng Shuping, and Gyan Mishra for their useful comments.

Thanks to Hariharan Ananthakrishnan for shepherding this document. Thanks to Deborah Brungard for providing comments and being the responsible AD for this document.

Thanks to Nic Leymann for the RTGDIR review.

Thanks to Benjamin Kaduk and Murray Kucherawy for their comments during the IESG review.

Contributors

The following individuals have contributed extensively:

Mahendra Singh Negi
RtBrick Inc
N-17L, 18th Cross Rd, HSR Layout
Bangalore 560102
Karnataka
India

Email: mahend.ietf@gmail.com

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

Email: dhruv.ietf@gmail.com

The following individuals have contributed text that was incorporated:

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
China

Email: zhang.xian@huawei.com

Udayasree Palle

Email: udayasreereddy@gmail.com

Mike Koldychev
Cisco Systems, Inc.
Canada

Email: mkoldych@cisco.com

Authors' Addresses

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

Email: slitkows@cisco.com

Siva Sivabalan
Ciena
385 Terry Fox Drive
Kanata Ontario K2K 0L1
Canada

Email: msiva282@gmail.com

Jeff Tantsura
Juniper Networks

Email: jefftant.ietf@gmail.com

Jonathan Hardwick

Metaswitch Networks
33 Genotin Road
Enfield
United Kingdom

Email: Jonathan.Hardwick@metaswitch.com

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

Email: c.l@huawei.com

Additional contact information:

李呈
中国
100095
北京
北研所
技术有限公司