Internet Engineering Task Force (IETF)
Request for Comments: 8281
Category: Standards Track
ISSN: 2070-1721

E. Crabbe
Individual Contributor
I. Minei
Google, Inc.
S. Sivabalan
Cisco Systems, Inc.
R. Varga
Pantheon Technologies SRO

December 2017

Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model

Abstract

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Client (PCC) requests.

The extensions for stateful PCE provide active control of Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSPs) via PCEP, for a model where the PCC delegates control over one or more locally configured LSPs to the PCE. This document describes the creation and deletion of PCE-initiated LSPs under the stateful PCE model.

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/rfc8281.

Table of Contents

1. Introduction

[RFC5440] describes the Path Computation Element Communication Protocol (PCEP). PCEP defines the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between PCE and PCE, enabling computation of Multiprotocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP) characteristics.

[RFC8231] specifies a set of extensions to PCEP to enable stateful control of TE LSPs between and across PCEP sessions in compliance with [RFC4657]. It includes:

- mechanisms to effect LSP State Synchronization between PCCs and PCEs
- delegation of control of LSPs to PCEs
- PCE control of timing and sequence of path computations within and across PCEP sessions

It focuses on a model where LSPs are configured on the PCC, and control over them is delegated to the PCE.

This document describes the setup, maintenance, and teardown of PCE-initiated LSPs under the stateful PCE model, without the need for local configuration on the PCC, thus allowing for a dynamic network that is centrally controlled and deployed.

2. Terminology

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

This document uses the following terms defined in [RFC8051]: Stateful PCE and Delegation.

This document uses the following terms defined in [RFC8231]: Redelegation Timeout Interval, State Timeout Interval, LSP State Report, and LSP Update Request.

The following terms are defined in this document:

PCE-initiated LSP: LSP that is instantiated as a result of a request from the PCE.

The message formats in this document are specified using Routing Backus-Naur Form (RBNF) encoding as specified in [RFC5511].

2.1. Requirements Language

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

3. Architectural Overview

3.1. Motivation

[RFC8231] provides active control over LSPs that are locally configured on the PCC. This model relies on the Label Edge Router (LER) taking an active role in delegating locally configured LSPs to the PCE and is well suited in environments where the LSP placement is fairly static. However, in environments where the LSP placement needs to change in response to application demands, it is useful to support dynamic creation and teardown of LSPs. The ability for a PCE to trigger the creation of LSPs on demand can be seamlessly integrated into a controller-based network architecture, where intelligence in the controller can determine when and where to set up paths.

A possible use case is a software-defined network, where applications request network resources and paths from the network infrastructure. For example, an application can request a path with certain constraints between two Label Switching Routers (LSRs) by contacting the PCE. The PCE can compute a path satisfying the constraints, and instruct the head end LSR to instantiate and signal it. When the path is no longer required by the application, the PCE can request its teardown.

Another use case is dynamically adjusting aggregate bandwidth between two points in the network using multiple LSPs. This functionality is very similar to auto-bandwidth, but it allows for providing the desired capacity through multiple LSPs. This approach overcomes two of the limitations auto-bandwidth can experience: 1) growing the capacity between the endpoints beyond the capacity of individual links in the path and 2) achieving good bin packing through use of several small LSPs instead of a single large one. The number of LSPs varies based on the demand, and LSPs are created and deleted dynamically to satisfy the bandwidth requirements.

Another use case is demand engineering, where a PCE with visibility into both the network state and the demand matrix can anticipate and optimize how traffic is distributed across the infrastructure. Such optimizations may require creating new paths across the infrastructure.

3.2. Operation Overview

This document defines the new I flag in the STATEFUL-PCE-CAPABILITY TLV to indicate that the sender supports PCE-initiated LSPs (see details in Section 4.1). A PCC or PCE sets this flag in the Open message during the PCEP initialization phase to indicate that it supports the procedures of this document.

This document defines a new PCEP message, the LSP Initiate Request (PCInitiate) message, which a PCE can send to a PCC to request the initiation or deletion of an LSP. The decision when to instantiate or delete a PCE-initiated LSP is out of the scope of this document.

The PCE sends a PCInitiate message to the PCC to request the initiation of an LSP. The PCC creates the LSP using the attributes communicated by the PCE and local values for any unspecified parameters. The PCC generates a Path Computation State Report (PCRpt) for the LSP, carrying a newly assigned PLSP-ID for the LSP and delegating the LSP to the PCE via the Delegate flag in the LSP object.

The PCE can update the attributes of the LSP by sending subsequent Path Computation Update Request (PCUpd) messages. Subsequent PCRpt and PCUpd messages that the PCC and PCE, respectively, send for the LSP will carry the PCC-assigned PLSP-ID, which uniquely identifies the LSP. See details in Section 5.3.

The PCE sends a PCInitiate message to the PCC to request the deletion of an LSP. To indicate a delete operation, this document defines the new R flag in the Stateful PCE Request Parameter (SRP) object in the PCInitiate message, as described in Section 5.2. As a result of the deletion request, the PCC removes the LSP and sends a PCRpt for the removed state. See details in Section 5.4.

Figure 1 illustrates these message exchanges.

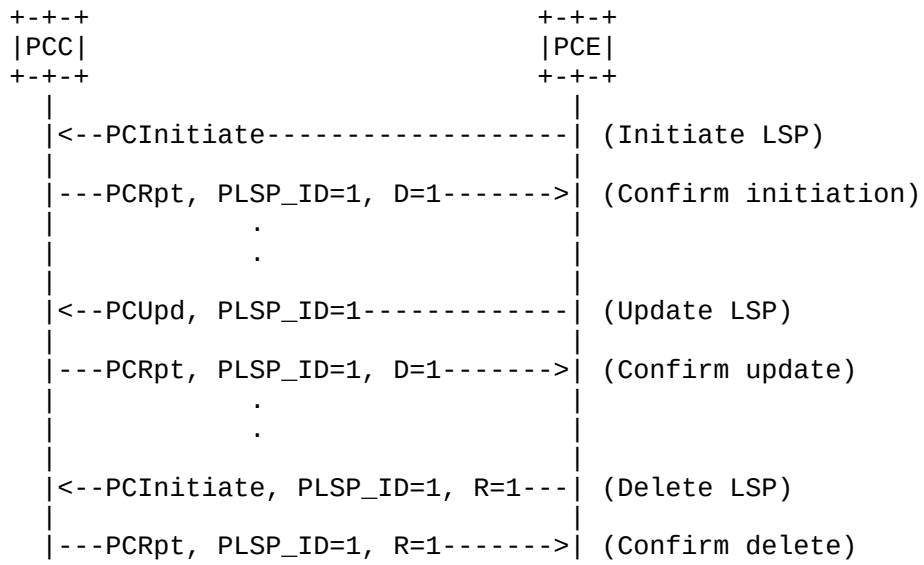

Figure 1: PCE-Initiated LSP Life Cycle

4. Support of PCE-Initiated LSPs

A PCEP speaker indicates its ability to support PCE-initiated LSPs during the PCEP initialization phase, as follows. When the PCEP session is created, it sends an Open message with an OPEN object that contains the STATEFUL-PCE-CAPABILITY TLV, as defined in [RFC8231]. A new flag, the I (LSP-INSTANTIATION-CAPABILITY) flag, is introduced to this TLV to indicate support for instantiation of PCE-initiated LSPs. A PCE can initiate LSPs only for PCCs that advertised this capability. A PCC will follow the procedures described in this document only on sessions where the PCE advertised the I flag.

4.1. STATEFUL-PCE-CAPABILITY TLV

The format of the STATEFUL-PCE-CAPABILITY TLV is defined in [RFC8231] and included here for easy reference with the addition of the new I flag.

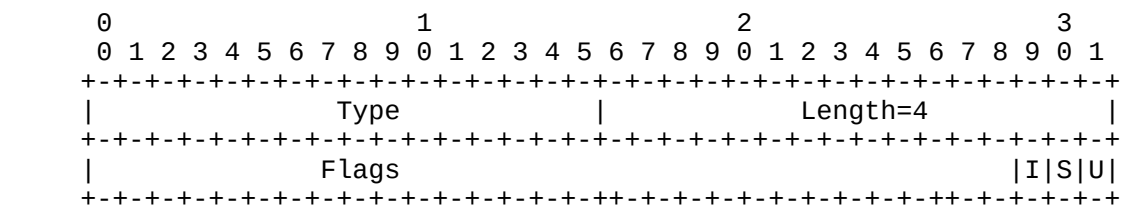

Figure 2: STATEFUL-PCE-CAPABILITY TLV Format

A new flag is defined to indicate the sender's support for LSP instantiation by a PCE:

I (LSP-INSTANTIATION-CAPABILITY -- 1 bit): If set to 1 by a PCC, the I flag indicates that the PCC allows instantiation of an LSP by a PCE. If set to 1 by a PCE, the I flag indicates that the PCE supports instantiating LSPs. The LSP-INSTANTIATION-CAPABILITY flag must be set by both the PCC and PCE in order to enable PCE-initiated LSP instantiation.

5. PCE-Initiated LSP Instantiation and Deletion

To initiate an LSP, a PCE sends a PCInitiate message to a PCC. The message format, objects, and TLVs are discussed separately below for the creation and the deletion cases.

5.1. The LSP Initiate Request

An LSP Initiate Request (PCInitiate) message is a PCEP message sent by a PCE to a PCC to trigger LSP instantiation or deletion. The Message-Type field of the PCEP common header for the PCInitiate message is set to 12. The PCInitiate message MUST include the SRP and the LSP objects and MAY contain other objects, as discussed later in this section.

The format of a PCInitiate message is as follows:

```
  <PCInitiate Message> ::= <Common Header>
                           <PCE-initiated-lsp-list>
Where:
  <Common Header> is defined in RFC 5440

  <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>]

  <PCE-initiated-lsp-deletion> ::= <SRP>
                                   <LSP>
```

Where:
<attribute-list> is defined in RFC 5440 and extended by PCEP extensions.

The LSP object is defined in [RFC8231]. The END-POINTS and Explicit Route Objects (EROs) are defined in [RFC5440].

The SRP object is defined in [RFC8231]. The SRP object contains an SRP-ID-number that is unique within a PCEP session. The PCE increments the last-used SRP-ID-number before it sends each PCInitiate message. The PCC MUST echo the value of the SRP-ID-number in PCEP Error (PCErr) and PCRpt messages that it sends as a result of the PCInitiate; this allows the PCE to correlate them with the corresponding PCInitiate message.

5.2. The R Flag in the SRP Object

The format of the SRP object is defined in [RFC8231] and included here for easy reference with the addition of the new R flag.

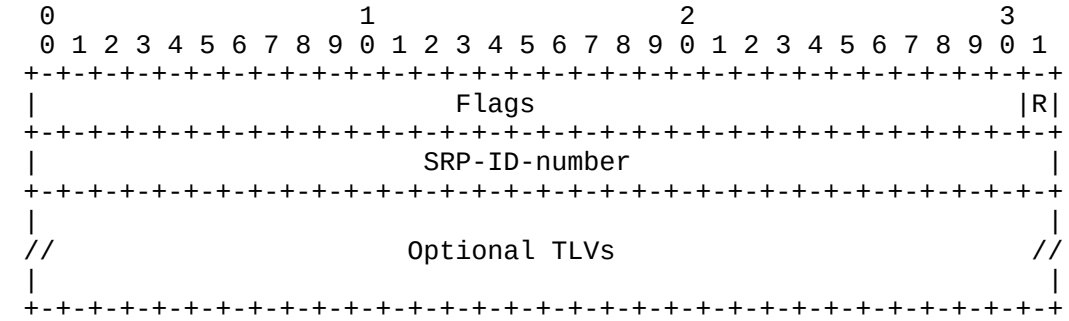

Figure 3: The SRP Object Format

A new flag is defined to indicate a delete operation initiated by the PCE:

R (LSP-REMOVE -- 1 bit): If set to 0, it indicates a request to create an LSP. If set to 1, it indicates a request to remove an LSP.

5.3. LSP Instantiation

The LSP is instantiated by sending a PCInitiate message. The LSP is set up using RSVP-TE. Extensions for other setup methods are outside the scope of this document.

The PCInitiate message, when used to instantiate an LSP, MUST contain an LSP object with the reserved PLSP-ID 0. The LSP object MUST include the SYMBOLIC-PATH-NAME TLV, which is used to correlate between the PCC-assigned PLSP-ID and the LSP.

The PCInitiate message, when used to instantiate an LSP, MUST contain an ERO for the LSP.

For an instantiation request of an RSVP-signaled LSP, the destination address may be needed. The PCC MAY determine it from a provided object (e.g., ERO) or a local decision. Alternatively, the END-POINTS object MAY be included to explicitly convey the destination addresses to be used in the RSVP-TE signaling. The source address MUST be either specified or left for the PCC to choose by setting it to "0.0.0.0" (if the destination is an IPv4 address) or "::" (if the destination is an IPv6 address).

The PCE MAY include various attributes as per [RFC5440]. The PCC MUST use these values in the LSP instantiation and local values for unspecified parameters. After the LSP setup, the PCC MUST send a PCRpt to the PCE, reflecting these values. The SRP object in the PCRpt message MUST echo the value of the PCInitiate message that triggered the setup. LSPs that were instantiated as a result of a PCInitiate message MUST have the Create flag (Section 5.3.1) set in the LSP object.

If the PCC receives a PCInitiate message with a non-zero PLSP-ID and the R flag in the SRP object set to zero, then it MUST send a PCErr message with Error-type=19 (Invalid Operation) and Error-value=8 (Non-zero PLSP-ID in the LSP Initiate Request).

If the PCC receives a PCInitiate message without an ERO and the R flag in the SRP object set to zero, then it MUST send a PCErr message with Error-type=6 (Mandatory Object missing) and Error-value=9 (ERO object missing).

If the PCC receives a PCInitiate message without a SYMBOLIC-PATH-NAME TLV, then it MUST send a PCErr message with Error-type=10 (Reception of an invalid object) and Error-value=8 (SYMBOLIC-PATH-NAME TLV missing).

The PCE MUST NOT provide a symbolic path name that conflicts with the symbolic path name of any existing LSP in the PCC. (Existing LSPs may be either statically configured or initiated by another PCE.) If there is a conflict with the symbolic path name of an existing LSP, the PCC MUST send a PCErr message with Error-type=23 (Bad Parameter value) and Error-value=1 (SYMBOLIC-PATH-NAME in use). The only exception to this rule is for LSPs for which the State Timeout Interval timer is running (see Section 6).

If the PCC determines that the LSP parameters proposed in the PCInitiate message are unacceptable, it MUST send a PCErr message with Error-type=24 (PCE instantiation error) and Error-value=1 (Unacceptable instantiation parameters). If the PCC encounters an internal error during the processing of the PCInitiate message, it MUST send a PCErr message with Error-type=24 (PCE instantiation error) and Error-value=2 (Internal error).

A PCC MUST relay errors it encounters in the setup of a PCE-initiated LSP to the PCE by sending a PCErr message with Error-type=24 (PCE instantiation error) and Error-value=3 (Signaling error). The PCErr message MUST echo the SRP-ID-number of the PCInitiate message. The PCEP-ERROR object SHOULD include the RSVP_ERROR_SPEC TLV (if an RSVP ERROR_SPEC object was returned to the PCC by a downstream node).

After the LSP is set up, errors in RSVP signaling are reported in PCRpt messages, as described in [RFC8231].

On successful completion of the LSP instantiation, the PCC MUST send a PCRpt message. The LSP object message MUST contain a non-zero PLSP-ID that uniquely identifies the LSP within this PCC and MUST have the Create flag (Section 5.3.1) and Delegate flag set. The SRP object MUST contain an SRP-ID-number that echoes the value from the PCInitiate message that triggered the setup. The PCRpt MUST include the attributes that the PCC used to instantiate the LSP.

A PCC SHOULD be able to place a limit on either the number of LSPs or the percentage of resources that are allocated to honor PCE-initiated LSP requests. As soon as that limit is reached, the PCC MUST send a PCErr message with Error-type=19 (Invalid Operation) and Error-value=6 (PCE-initiated LSP limit reached) and is free to drop any incoming PCInitiate messages without additional processing.

Similarly, the PCE SHOULD be able to place a limit on either the number of PCInitiate messages pending for a particular PCC or the time it waits for a response (positive or negative) to a PCInitiate message from a PCC, and it MAY take further action (such as closing the session or removing all its LSPs) if this limit is reached.

5.3.1. The Create Flag

The LSP object is defined in [RFC8231] and included here for easy reference with the addition of the new Create (C) flag.

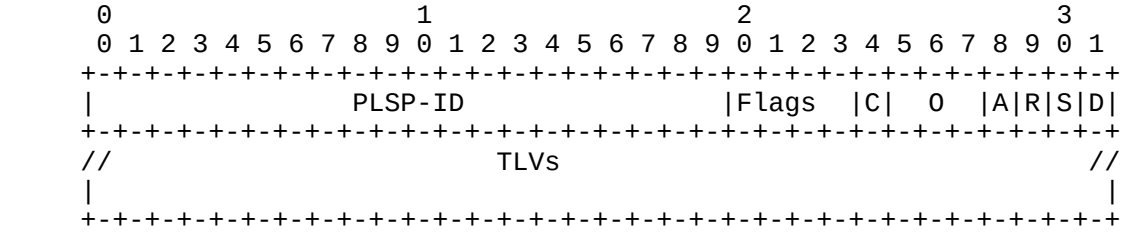

Figure 4: The LSP Object Format

A new flag, the C flag, is introduced. On a PCRpt message, the C flag set to 1 indicates that this LSP was created via a PCInitiate message. The C flag MUST be set to 1 on each PCRpt message for the LSP's duration of existence. The C flag allows PCEs to be aware of which LSPs were PCE initiated (a state that would otherwise only be known by the PCC and the PCE that initiated them).

5.3.2. The SPEAKER-ENTITY-ID TLV

The optional SPEAKER-ENTITY-ID TLV defined in [RFC8232] MAY be included in the LSP object in a PCRpt message as an optional TLV for LSPs for which the C flag is 1. The SPEAKER-ENTITY-ID TLV identifies the PCE that initiated the creation of the LSP on all PCEP sessions, a state that would otherwise only be known by the PCC and the PCE that initiated the LSP. If the TLV appears in a PCRpt for an LSP for which the C flag is 0, the LSP MUST be ignored, and the PCE MUST send a PCErr message with Error-type=23 (Bad parameter value) and Error-value=2 (Speaker identity included for an LSP that is not PCE initiated).

5.4. LSP Deletion

A PCE can initiate the removal of a PCE-initiated LSP by sending a PCInitiate message with an LSP object carrying the PLSP-ID of the LSP to be removed and an SRP object with the R flag set (see Section 5.2). A PLSP-ID of zero removes all LSPs with the C flag set to 1 (in their LSP object) that are delegated to the PCE.

If the PLSP-ID is unknown, the PCC MUST send a PCErr message with Error-type=19 (Invalid Operation) and Error-value=3 (Unknown PLSP-ID) [RFC8231].

If the PLSP-ID specified in the PCInitiate message is not delegated to the PCE, the PCC MUST send a PCErr message with Error-type=19 (Invalid operation) and Error-value=1 (LSP is not delegated) [RFC8231].

If the PLSP-ID specified in the PCInitiate message was not created by a PCE, the PCC MUST send a PCErr message with Error-type=19 (Invalid operation) and Error-value=9 (LSP is not PCE initiated).

Following the removal of the LSP, the PCC MUST send a PCRpt as described in [RFC8231]. The SRP object in the PCRpt MUST include the SRP-ID-number from the PCInitiate message that triggered the removal. The R flag in the SRP object MUST be set.

6. LSP Delegation and Cleanup

The PCC MUST delegate PCE-initiated LSPs to the PCE upon instantiation. The PCC MUST set the delegation bit to 1 in the PCRpt that includes the assigned PLSP-ID.

The PCC MUST NOT revoke the delegation for a PCE-initiated LSP on an active PCEP session. Therefore, all PCRpt messages from the PCC to the PCE that owns the delegation MUST have the delegation bit set to 1. If the PCE that owns the delegation receives a PCRpt message with the delegation bit set to 0, then it MUST send a PCErr message with Error-type=19 (Invalid Operation) and Error-value=7 (Delegation for PCE-initiated LSP cannot be revoked). The PCE MAY further react by closing the session.

Control over a PCE-initiated LSP can revert to the PCC in two ways. A PCE MAY return a delegation to the PCC to allow for LSP transfer between PCEs. Alternatively, the PCC gains control of an LSP if the PCEP session that it was delegated on fails and the Redelegation Timeout Interval timer expires. In both cases, the LSP becomes an orphan until the expiration of the State Timeout Interval timer [RFC8231].

The PCC MAY attempt to redelegate an orphaned LSP by following the procedures of [RFC8231]. Alternatively, if the orphaned LSP was PCE-initiated, then a PCE MAY obtain control over it, as follows.

A PCE (either the original or one of its backups) sends a PCInitiate message that includes just the SRP and LSP objects and carries the PLSP-ID of the LSP it wants to take control of. If the PCC receives a PCInitiate message with a PLSP-ID pointing to an orphaned PCE-initiated LSP, then it MUST redelegate that LSP to the PCE. Any other non-zero PLSP-ID MUST result in the generation of a PCErr message using the rules described in Section 5.4. The State Timeout Interval timer for the LSP is stopped upon the redelegation. After obtaining control of the LSP, the PCE may remove it using the procedures described in this document.

The State Timeout Interval timer ensures that a PCE crash does not result in automatic and immediate disruption for the services using PCE-initiated LSPs. PCE-initiated LSPs are not removed immediately upon PCE failure. Instead, they are cleaned up on the expiration of this timer. This allows for network cleanup without manual intervention. The PCC MUST support removal of PCE-initiated LSPs as one of the behaviors applied on expiration of the State Timeout Interval timer. The behavior MUST be picked based on local policy and can result in either LSP removal or reverting to operator-defined default parameters.

7. LSP State Synchronization

LSP State Synchronization procedures are described in Section 5.6 of [RFC8231]. During State Synchronization, a PCC reports the state of its LSPs to the PCE using PCRpt messages, setting the SYNC flag in the LSP object. For PCE-initiated LSPs, the PCC MUST also set the Create flag in the LSP object and MAY include the SPEAKER-ENTITY-ID TLV identifying the PCE that requested the LSP creation. At the end of State Synchronization, the PCE SHOULD send a PCInitiate message to initiate any missing LSPs and/or remove any LSPs that are not wanted. Under some circumstances, depending on the deployment, it might be preferable for a PCE not to send this PCInitiate immediately, or at all. For example, the PCC may be a slow device, or the operator might prefer not to disrupt active flows.

8. IANA Considerations

As detailed below, IANA has allocated code points for the protocol elements defined in this document.

8.1. PCEP Messages

IANA has registered the following message type within the "PCEP Messages" subregistry of the PCEP Numbers registry. (Note that the early allocation for this message type was called "Initiate"; it has been changed as follows.)

Value	Meaning	Reference
12	LSP Initiate Request	RFC 8281

8.2. LSP Object

[RFC8231] defines the LSP object; per that RFC, IANA created a registry to manage the value of the LSP object's Flag field. IANA has allocated a new bit in the "LSP Object Flag Field" subregistry, as follows:

Bit	Description	Reference
4	Create	RFC 8281

8.3. SRP object

IANA has created a new subregistry, named "SRP Object Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry, to manage the Flag field of the SRP object. New values are to be assigned by Standards Action [RFC8126]. Each bit is tracked with the following qualities: bit number (counting from bit 0 as the most significant bit), description, and defining RFC.

The following values are defined in this document:

Bit	Description	Reference
31	LSP-Remove	RFC 8281

8.4. STATEFUL-PCE-CAPABILITY TLV

[RFC8231] defines the STATEFUL-PCE-CAPABILITY TLV; per that RFC, IANA created a registry to manage the value of the STATEFUL-PCE-CAPABILITY TLV's Flag field. IANA has allocated a new bit in the STATEFUL-PCE-CAPABILITY TLV Flag Field registry, as follows:

Bit	Description	Reference
29	LSP-INSTANTIATION-CAPABILITY (I)	RFC 8281

8.5. PCEP-Error Object

IANA has registered the following error types and error values within the "PCEP-ERROR Object Error Types and Values" subregistry of the PCEP Numbers registry.

Error-Type	Meaning
10	Reception of an invalid object
	Error-value=8: SYMBOLIC-PATH-NAME TLV missing
19	Invalid Operation
	Error-value=6: PCE-initiated LSP limit reached
	Error-value=7: Delegation for PCE-initiated LSP cannot be revoked
	Error-value=8: Non-zero PLSP-ID in LSP Initiate Request
	Error-value=9: LSP is not PCE initiated
	Error-value=10: PCE-initiated operation-frequency limit reached
23	Bad parameter value
	Error-value=1: SYMBOLIC-PATH-NAME in use
	Error-value=2: Speaker identity included for an LSP that is not PCE initiated
24	LSP instantiation error
	Error-value=1: Unacceptable instantiation parameters
	Error-value=2: Internal error
	Error-value=3: Signaling error

9. Security Considerations

The security considerations described in [RFC8231] apply to the extensions described in this document. Additional considerations related to a malicious PCE are introduced.

9.1. Malicious PCE

The LSP instantiation mechanism described in this document allows a PCE to generate state on the PCC and throughout the network. As a result, it introduces a new attack vector: an attacker may flood the PCC with LSP instantiation requests and consume network and LSR resources by either spoofing messages or compromising the PCE itself.

A PCC can protect itself from such an attack by imposing a limit on either the number of LSPs or the percentage of resources that are allocated to honor PCE-initiated LSP requests. As soon as that limit is reached, the PCC MUST send a PCErr message with Error-type=19 (Invalid Operation) and Error-value=6 (PCE-initiated LSP limit reached) and is free to drop any incoming PCInitiate messages for LSP initiation without additional processing.

Rapid flaps triggered by the PCE can also be an attack vector. A PCC can protect itself from such an attack by imposing a limit on the number of flaps per unit of time that it allows a PCE to generate. As soon as that limit is reached, a PCC MUST send a PCErr message with Error-type=19 (Invalid Operation) and Error-value=10 (PCE-initiated operation-frequency limit reached) and is free to treat the session as having reached the limit in terms of resources allocated to honor PCE-initiated LSP requests, either permanently or for a locally-defined cool-off period.

9.2. Malicious PCC

The LSP instantiation mechanism described in this document requires the PCE to keep state for LSPs that it instantiates and relies on the PCC responding (with either a state report or an error message) to requests for LSP instantiation. A malicious PCC or one that reached the limit of the number of PCE-initiated LSPs can ignore PCE requests and consume PCE resources. A PCE can protect itself by imposing a limit on the number of requests pending or by setting a timeout, and it MAY take further action such as closing the session or removing all the LSPs it initiated.

10. References

10.1. 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>.

[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>.

10.2. Informative References

[RFC4657] Ash, J., Ed. and J. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <https://www.rfc-editor.org/info/rfc4657>.

[RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", RFC 8051, DOI 10.17487/RFC8051, January 2017, <https://www.rfc-editor.org/info/rfc8051>.

[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>.

Acknowledgments

We would like to thank Jan Medved, Ambrose Kwong, Ramon Casellas, Cyril Margaria, Dhruv Dhody, Raveendra Trovi, and Jon Hardwick for their contributions to this document.

Authors' Addresses

Edward Crabbe
Individual Contributor

Email: edward.crabbe@gmail.com

Ina Minei
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States of America

Email: inaminei@google.com

Siva Sivabalan
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
United States of America

Email: msiva@cisco.com

Robert Varga
Pantheon Technologies SRO
Mlynske Nivy 56
Bratislava 821 05
Slovakia

Email: robert.varga@pantheon.tech