

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 8, 2018

E. Crabbe
Individual Contributor
I. Minei
Google, Inc.
S. Sivabalan
Cisco Systems, Inc.
R. Varga
Pantheon Technologies SR0
October 5, 2017

PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model
draft-ietf-pce-pce-initiated-lsp-11

Abstract

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

The extensions for stateful PCE provide active control of Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSP) 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.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

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 <http://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 April 8, 2018.

Copyright Notice

Copyright (c) 2017 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 (<http://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	3
2.	Terminology	3
3.	Architectural Overview	4
3.1.	Motivation	4
3.2.	Operation Overview	5
4.	Support of PCE-initiated LSPs	6
4.1.	STATEFUL-PCE-CAPABILITY TLV	6
5.	PCE-initiated LSP Instantiation and Deletion	7
5.1.	The LSP Initiate Request	7
5.2.	The R flag in the SRP Object	8
5.3.	LSP Instantiation	9
5.3.1.	The Create Flag	11
5.3.2.	The SPEAKER-ENTITY-ID TLV	11
5.4.	LSP Deletion	12
6.	LSP Delegation and Cleanup	12
7.	LSP State Synchronization	13
8.	Implementation Status	14
9.	IANA Considerations	14
9.1.	PCEP Messages	14
9.2.	LSP Object	15
9.3.	SRP object	15
9.4.	STATEFUL-PCE-CAPABILITY TLV	15
9.5.	PCEP-Error Object	15
10.	Security Considerations	16
10.1.	Malicious PCE	16
10.2.	Malicious PCC	17
11.	Acknowledgements	17
12.	References	17
12.1.	Normative References	17

12.2. Informative References	18
Authors' Addresses	18

[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

- o mechanisms to effect LSP state synchronization between PCCs and PCEs
- o delegation of control of LSPs to PCEs
- o 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, PCEP Peer.

This document uses the following terms defined in [\[RFC8051\]](#): Stateful PCE, Delegation.

This document uses the following terms defined in [\[RFC8231\]](#): Redelegation Timeout Interval, State Timeout Interval, LSP State Report, 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](#)].

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 tear down 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 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 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 an LSP 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 PCUpd messages. Subsequent LSP State Report (PCRpt) and LSP Update Request (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 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.

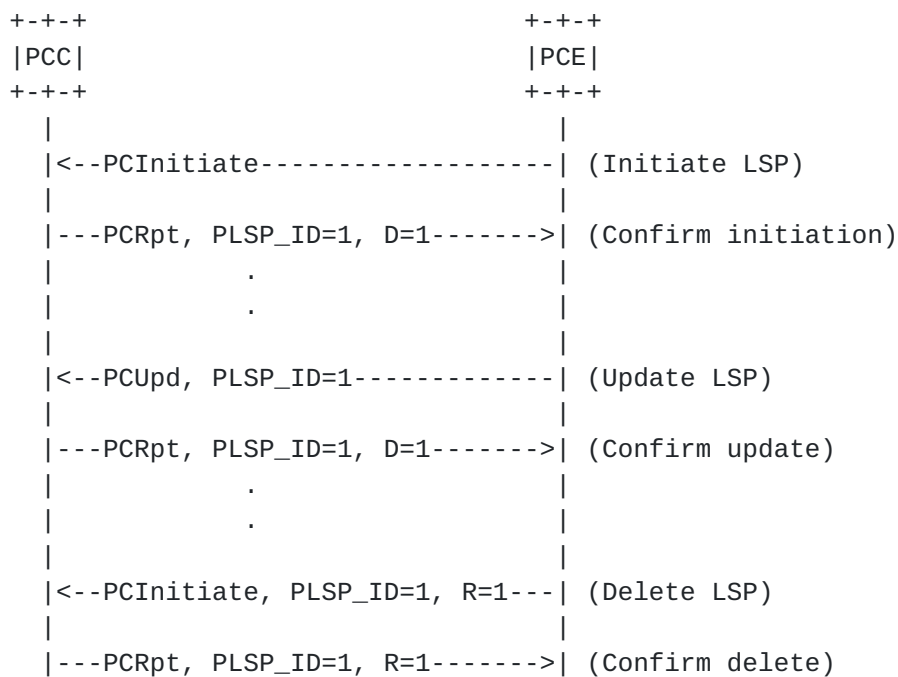


Figure 1: PCE-Initiated LSP lifecycle

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, defined in [RFC8231]. A new flag, the I (LSP-INANTIATION-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

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 PCC and PCE in order to enable PCE-initiated LSP instantiation.

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.

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 [[RFC5440](#)]

```
<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 [[RFC5440](#)] and extended by PCEP extensions.

The LSP object is defined in [[RFC8231](#)]. The END-POINTS and ERO objects are defined in [[RFC5440](#)].

The SRP object is defined in [[RFC8231](#)]. The SRP Object contains an SRP-ID-number which 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 PCErr and PCRpt messages that it sends as a result of the PCInitiate to allow 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 draft.

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 Explicit Route Object (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 either be 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 PCInitiate message).

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 (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 to the PCE errors it encounters in the setup of PCE-initiated LSP 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 on the time it waits for a response (positive or negative) to a PCInitiate message from a PCC and 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 C flag.

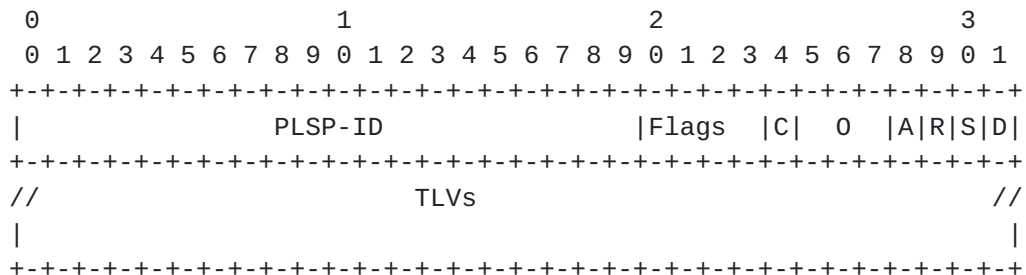


Figure 4: The LSP Object format

A new flag, the Create (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 duration of existence of the LSP. The Create 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 which 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 Redelegating 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, including just the SRP and LSP objects, and carrying 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 either in LSP removal, or in reverting to operator-defined default parameters.

7. LSP State Synchronization

LSP State Synchronization procedures are described in [section 5.4 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. Implementation Status

This section to be removed by the RFC editor.

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 catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC 7942](#), "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".

Two vendors are implementing the extensions described in this draft and have included the functionality in releases that will be shipping in the near future. An additional entity is working on implementing these extensions in the scope of research projects.

9. IANA Considerations

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

9.1. PCEP Messages

IANA is requested to confirm the early allocation of the following new message type within the "PCEP Messages" sub-registry of the PCEP Numbers registry, and to update the reference in the registry to point to this document, when it is an RFC:

Value	Meaning	Reference
12	LSP Initiate Request	This document

Note to IANA: The early allocation was done for a message called "Initiate". This name has changed to "LSP Initiate Request" as above.

9.2. LSP Object

[RFC8231] defines the LSP Object and requests that IANA creates a registry to manage the value of the LSP Object's Flag field. IANA is requested to allocate a new bit in the LSP Object Flag Field registry, as follows:

Bit	Description	Reference
4	Create	This document

9.3. SRP object

This document requests that a new sub-registry, named "SRP Object Flag Field", is created 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 should be 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	This document

9.4. STATEFUL-PCE-CAPABILITY TLV

[RFC8231] defines the STATEFUL-PCE-CAPABILITY TLV and requests that IANA creates a registry to manage the value of the STATEFUL-PCE-CAPABILITY TLV's Flag field. IANA is requested to allocate a new bit in the STATEFUL-PCE-CAPABILITY TLV Flag Field registry, as follows:

Bit	Description	Reference
29	I (LSP-INstantiation-Capability)	This document

9.5. PCEP-Error Object

IANA is requested to confirm the early allocation of the following new error types and error values within the "PCEP-Error Object Error Types and Values" sub-registry of the PCEP Numbers registry, and to update the reference in the registry to point to this document, when it is an RFC:

Error-Type	Meaning
10	Invalid Object
19	Error-value=8: SYMBOLIC-PATH-NAME TLV missing 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 PCInitiate message 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

10. 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.

10.1. 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, either by spoofing messages or by 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 instantiation 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 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.

10.2. 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.

11. Acknowledgements

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.

12. References

12.1. 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>>.

- [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>>.

12.2. 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>>.
- [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>>.
- [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>>.

Authors' Addresses

Edward Crabbe

Individual Contributor

Email: edward.crabbe@gmail.com

Ina Minei
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: inaminei@google.com

Siva Sivabalan
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
US

Email: msiva@cisco.com

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

Email: robert.varga@pantheon.tech

