

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2017

E. Crabbe
Individual Contributor
I. Minei
Google, Inc.
S. Sivabalan
Cisco Systems, Inc.
R. Varga
Pantheon Technologies SR0
March 4, 2017

**PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model
draft-ietf-pce-pce-initiated-lsp-08**

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 September 5, 2017.

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	4
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 Message	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-IDENTITY-ID TLV	11
5.4.	LSP Deletion	12
6.	LSP delegation and Cleanup	12
7.	Implementation Status	13
8.	IANA Considerations	14
8.1.	PCEP Messages	14
8.2.	LSP Object	14
8.3.	SRP object	14
8.4.	STATEFUL-PCE-CAPABILITY TLV	14
8.5.	PCEP-Error Object	15
9.	Security Considerations	15
9.1.	Malicious PCE	15
9.2.	Malicious PCC	16
10.	Acknowledgements	16
11.	References	16
11.1.	Normative References	16
11.2.	Informative References	17

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 Control Element (PCE), or between PCE and PCE, enabling computation of Multiprotocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP) characteristics.

Stateful pce [[I-D.ietf-pce-stateful-pce](#)] 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, and PCE control of timing and sequence of path computations within and across PCEP sessions and 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 [[I-D.ietf-pce-stateful-pce](#)]: 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 Format (RBNF) encoding as specified in [[RFC5511](#)].

3. Architectural Overview

3.1. Motivation

[I-D.ietf-pce-stateful-pce] 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 make possible agile software-driven network operation, and 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-driven 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

A PCC or PCE indicates its ability to support PCE-provisioned dynamic LSPs during the PCEP Initialization Phase via a new flag in the STATEFUL-PCE-CAPABILITY TLV (see details in [Section 4.1](#)).

The decision when to instantiate or delete a PCE-initiated LSP is out of the scope of this document. To instantiate or delete an LSP, the PCE sends a new message, the Path Computation LSP Initiate Request (PCInitiate) message to the PCC. The LSP Initiate Request MUST include the SRP and LSP objects ([[I-D.ietf-pce-stateful-pce](#)]), and the LSP object MUST include the Symbolic Path Name TLV and MUST have a PLSP-ID ([[I-D.ietf-pce-stateful-pce](#)]) of 0.

For an instantiation operation, the PCE MUST include the ERO and END-POINTS object and may include various attributes as per [[RFC5440](#)]. The PCC creates the LSP using the attributes communicated by the PCE, and local values for the unspecified parameters. It assigns a unique PLSP-ID for the LSP and automatically delegates the LSP to the PCE. It MUST also generate an LSP State Report (PCRpt) for the LSP, carrying the newly assigned PLSP-ID and indicating the delegation via the Delegate flag in the LSP object. In addition to the Delegate flag, the PCC MUST also set the Create flag in the LSP object (see [Section 5.3.1](#)), to indicate that the LSP was created as a result of a PCInitiate message and SHOULD include the optional SPEAKER-IDENTITY-ID TLV defined in [[I-D.ietf-pce-stateful-sync-optimizations](#)] identifying the PCE that requested the LSP creation. This PCRpt message MUST include the SRP object, with the SRP-id-number used in the SRP object of the PCInitiate message. The PCE MAY update the attributes of the LSP via subsequent PCUpd messages. Subsequent LSP State Report and LSP Update Request for the LSP will carry the PCC-assigned PLSP-ID, which uniquely identifies the LSP. See details in [Section 5.3](#).

Once instantiated, the delegation procedures for PCE-initiated LSPs are the same as for PCC-initiated LSPs as described in [[I-D.ietf-pce-stateful-pce](#)], with the exception that the PCC cannot revoke a delegation for a PCE-initiated LSP. This applies to the case of a PCE failure as well. In order to allow for network cleanup without manual intervention, the PCC SHOULD support removal of PCE-initiated LSPs as one of the behaviors applied on expiration of the State Timeout Interval [[I-D.ietf-pce-stateful-pce](#)]. The behavior SHOULD be picked based on local policy, and can result either in LSP removal, or into reverting to operator-defined default parameters. See details in [Section 6](#). A PCE MAY return a delegation to the PCC in order to facilitate re-delegation of its LSPs to an alternate PCE.

To indicate a delete operation, the PCE MUST use the new R flag in the SRP object in a PCInitiate message as described in [Section 5.2](#). As a result of the deletion request, the PCC MUST remove all state related to the LSP, and send a PCRpt with the R flag set in the LSP object for the removed state. See details in [Section 5.4](#).

I (LSP-INANTIATION-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 may attempt to instantiate LSPs. The LSP-INSTANTIATION-CAPABILITY flag must be set by both PCC and PCE in order to support PCE-initiated LSP instantiation.

Unassigned bits are considered reserved. They MUST be set to 0 on transmission and MUST be ignored on receipt.

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 Message

A Path Computation LSP Initiate Message is referred to as PCInitiate message. It 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 (suggested value, to be assigned by IANA). The PCInitiate message MUST include the SRP and the LSP objects, and MAY contain other objects, as discussed later in this section. Missing SRP and LSP objects in the PCInitiate message MUST trigger the same PCErr procedures as specified in [[I-D.ietf-pce-stateful-pce](#)]). for PCUpd. LSP instantiation is done by sending a PCInitiate message with an LSP object with the reserved PLSP-ID 0. LSP deletion is done 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](#)).

The format of a PCInitiate message for LSP instantiation is as follows:


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

Where:

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

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

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

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

Where:

<attribute-list> is defined in [[RFC5440](#)] and extended by PCEP extensions.

The SRP object is used to correlate between initiation requests sent by the PCE and the error reports and state reports sent by the PCC. Every request from the PCE receives a new SRP-ID-number. This number is unique per PCEP session and is incremented each time an operation (initiation, update, etc) is requested from the PCE. The value of the SRP-ID-number MUST be echoed back by the PCC in PCErr and PCRpt messages to allow for correlation between requests made by the PCE and errors or state reports generated by the PCC. Details of the SRP object and its use can be found in [[I-D.ietf-pce-stateful-pce](#)].

5.2. The R flag in the SRP Object

The format of the SRP object is shown Figure 2:




Figure 2: The SRP Object format

The type object is defined in [[I-D.ietf-pce-stateful-pce](#)].

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

R (LSP-REMOVE - 1 bit): If set to 1, it indicates a removal request initiated by the PCE.

5.3. LSP Instantiation

LSP instantiation is done by sending an PCInitiate Message with an LSP object with the reserved PLSP-ID 0. The LSP is set up using RSVP-TE, extensions for other setup methods are outside the scope of this draft.

Receipt of a PCInitiate Message with a non-zero PLSP-ID and the R flag in the SRP object set to zero MUST result in a PCErr message of type 19 (Invalid Operation) and value to be assigned by IANA, suggested value 8 (Non-zero PLSP-ID in LSP initiation request).

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 may be either specified or left up to the PCC decision using the 0.0.0.0 value. For LSPs to be setup by other means (e.g., Segment Routing), the END-POINTS object SHOULD be omitted.

The ERO Object is mandatory for an instantiation request. It contains the ERO for the LSP. If the ERO Object is missing, the PCC MUST send a PCErr message with Error-type=6 (Mandatory Object missing) and Error-value=9 (ERO Object missing).

The LSP Object MUST include the SYMBOLIC-PATH-NAME TLV, which will be used to correlate between the PCC-assigned PLSP-ID and the LSP. If the TLV is missing, the PCC MUST send a PCErr message with Error-type=10 (Invalid object) and Error-value to be assigned by IANA, suggested value 8, (SYMBOLIC-PATH-NAME TLV missing). The symbolic name used for provisioning PCE-initiated LSPs must not have conflict with the LSP name of any existing LSP in the PCC. (Existing LSPs may be either statically configured, or initiated by another PCE). If there is conflict with the LSP name, the PCC MUST send a PCErr message with Error-type to be assigned by IANA, suggested value 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 timer is running (see [Section 6](#)).

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 determines that the LSP parameters proposed in the PCInitiate message are unacceptable, it MUST trigger a PCErr with error-type to be assigned by IANA, suggested value 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 trigger a PCErr with error-type to be assigned by IANA, suggested vlaue 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 with error-type to be assigned by IANA, suggeseted value 24 (PCE instantiation error) and error-value=3 (Signaling error). The PCErr MUST echo the SRP-ID-number of the PCInitiate message. The PCEP-ERROR object SHOULD include the RSVP Error Spec TLV (if an ERROR SPEC 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 [\[I-D.ietf-pce-stateful-pce\]](#).

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 of type 19 (Invalid Operation) and value to be assigned by IANA (PCE-initiated 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 LSP initiation requests pending for a particular PCC, or on the time it waits for a response (positive or negative) to a PCInitiate request from a PCC and MAY take further action (such as closing the session or removing all its LSPs) if this limit is reached.

On successful completion of the LSP instantiation, the PCC assigns a PLSP-ID, and immediately delegates the LSP to the PCE by sending a PCRpt with the Delegate flag set.

5.3.1. The Create Flag

The LSP object is defined in [[I-D.ietf-pce-stateful-pce](#)] and included here for easy reference.

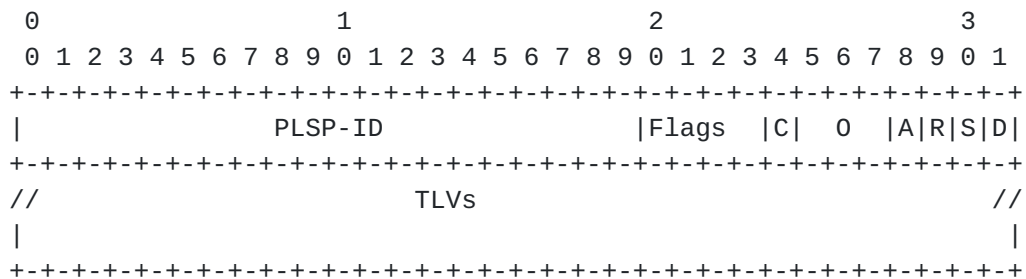


Figure 3: The LSP Object format

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-IDENTITY-ID TLV

The optional SPEAKER-IDENTITY-ID TLV defined in [[I-D.ietf-pce-stateful-sync-optimizations](#)] 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-IDENTITY-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 TLV 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

PCE-initiated removal of a PCE-initiated LSP is done by setting the R (remove) flag in the SRP Object in the PCInitiate message from the PCE. The LSP is identified by the PLSP-ID in the LSP object. If the PLSP-ID is unknown, the PCC MUST generate a PCErr with error type 19, error value 3, "Unknown PLSP-ID" ([[I-D.ietf-pce-stateful-pce](#)]). A PLSP-ID of zero removes all LSPs that were initiated by the PCE. If the PLSP-ID specified in the PCInitiate message is not delegated to the PCE, the PCC MUST send a PCErr message indicating "LSP is not delegated" (Error code 19, error value 1 as per [[I-D.ietf-pce-stateful-pce](#)]). If the PLSP-ID specified in the PCInitiate message was not created by a PCE, the PCC MUST send a PCErr message indicating that the LSP is not PCE-initiated, Error code 19, error value to be assigned by IANA, suggested value 9 (LSP is not PCE-initiated). Following the removal of the LSP, the PCC MUST send a PCRpt as described in [[I-D.ietf-pce-stateful-pce](#)]. 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

PCE-initiated LSPs are automatically delegated by the PCC to the PCE upon instantiation. The PCC MUST set the delegation bit to 1 in the PCRpt that includes the assigned PLSP-ID. All subsequent messages from the PCC to the PCE that initiated the LSP MUST have the delegation bit set to 1. The PCC cannot revoke the delegation for PCE-initiated LSPs for an active PCEP session. Sending a PCRpt message with the delegation bit set to 0 to the PCE that initiated the LSP results in a PCErr message of type 19 (Invalid Operation) and error-value value to be assigned by IANA, suggested value 7, (Delegation for PCE-initiated LSP cannot be revoked). The PCE MAY further react by closing the session.

A PCE MAY return a delegation to the PCC, to allow for LSP transfer between PCEs. Doing so MUST trigger the State Timeout Interval timer ([[I-D.ietf-pce-stateful-pce](#)]) for that particular LSP.

In case of PCEP session failure, control over PCE-initiated LSPs reverts to the PCC at the expiration of the redelegation timeout. At this point, the LSP is an "orphan" until the expiration of the State Timeout timer. To obtain control of a PCE-initiated LSP, 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. On receipt of a PCInitiate message with a PLSP-ID pointing to an orphan LSP, the PCC MUST redelegate that LSP to the PCE. Any other non-zero PLSP-ID MUST

result in the generation of a PCErr. The State Timeout 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 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 be 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 SHOULD support removal of PCE-initiated LSPs as one of the behaviors applied on expiration of the State Timeout Interval [[I-D.ietf-pce-stateful-pce](#)]. The behavior SHOULD be picked based on local policy, and can result either in LSP removal, or into reverting to operator-defined default parameters.

7. 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 [[RFC6982](#)]. 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 6982](#), "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.

8. IANA Considerations

8.1. PCEP Messages

IANA is requested to allocate a new message type within the "PCEP Messages" sub-registry of the PCEP Numbers registry, as follows:

Value	Meaning	Reference
12	Initiate	This document

8.2. LSP Object

[I-D.ietf-pce-stateful-pce] 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

8.3. 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 [[RFC5226](#)]. 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

8.4. STATEFUL-PCE-CAPABILITY TLV

[I-D.ietf-pce-stateful-pce] 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

8.5. PCEP-Error Object

IANA is requested to allocate new error types and error values within the "PCEP-ERROR Object Error Types and Values" sub-registry of the PCEP Numbers registry, as follows:

Error-Type	Meaning
10	Invalid Object
19	Error-value=8: SYMBOLIC-PATH-NAME TLV missing Invalid operation
23	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 initiation request Error-value=9: LSP is not PCE-initiated Error-value=10: PCE-initiated operation-frequency limit reached Bad parameter value
24	Error-value=1: SYMBOLIC-PATH-NAME in use Error-value=2: Speaker identity included for an LSP that is not PCE-initiated 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 [[I-D.ietf-pce-stateful-pce](#)] 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, 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 of type 19 (Invalid Operation) and value to be assigned by IANA, suggested 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 of type 19 (Invalid Operation) and value to be assigned by IANA, suggested 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.

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. Acknowledgements

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

11. References

11.1. Normative References

[I-D.ietf-pce-stateful-pce]

Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", [draft-ietf-pce-stateful-pce-18](#) (work in progress), December 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc5511>>.

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

11.2. Informative References

[I-D.ietf-pce-stateful-sync-optimizations]

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", [draft-ietf-pce-stateful-sync-optimizations-09](#) (work in progress), February 2017.

[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, <<http://www.rfc-editor.org/info/rfc4657>>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

[RFC6982] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [RFC 6982](#), DOI 10.17487/RFC6982, July 2013, <<http://www.rfc-editor.org/info/rfc6982>>.

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

