

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2013

E. Crabbe
Google, Inc.
I. Minei
Juniper Networks, Inc.
S. Sivabalan
Cisco Systems, Inc.
R. Varga
Pantheon Technologies SR0
October 9, 2012

PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model
draft-crabbe-pce-pce-initiated-lsp-00

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 described in [[I-D.ietf-pce-stateful-pce](#)] provide stateful 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 12, 2013.

Copyright Notice

Copyright (c) 2012 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	4
2.	Terminology	4
3.	Architectural Overview	4
3.1.	Motivation	5
3.2.	Operation overview	5
4.	Support of PCE-initiated LSPs	6
4.1.	Stateful PCE Capability TLV	6
5.	PCE-initiated LSP creation	7
5.1.	The LSP Create Message	7
6.	LSP delegation and cleanup	9
6.1.	LSP delegation procedures	9
6.2.	LSP cleanup procedures	9
6.2.1.	LSP-CLEANUP TLV	10
7.	IANA considerations	10
7.1.	PCEP-Error Object	11
7.2.	PCEP TLV Type Indicators	11
8.	Security Considerations	11
8.1.	Malicious PCE	11
9.	Acknowledgements	12
10.	References	12
10.1.	Normative References	12
10.2.	Informative References	13
	Authors' Addresses	13

1. Introduction

[RFC5440] describes the Path Computation Element 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 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 [[I-D.ietf-pce-stateful-pce](#)]: Stateful PCE, Delegation, Delegation 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.

LSP cleanup timer: PCE-defined timer for cleanup of PCE-initiated LSPs that are no longer delegated to a 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 stateful control over LSPs that are locally configured on the PCC. This model relies on the 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 one of 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 create and signal it. When the path is no longer required by the application, the PCE can request its teardown.

Another use case is that of 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 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 create a PCE-initiated LSP is out of the scope of this document. To instantiate an LSP, the PCE sends a new message, the LSP Create Request (PCCreate) message to the PCC. The LSP Create Request MUST include the END-POINTS and LSPA objects, and the LSPA object MUST include the SYMBOLIC-PATH-NAME TLV. The PCC creates the LSP using the attributes communicated by the PCE, and local values for the unspecified parameters. It assigns a unique LSP-ID for the LSP and automatically delegates the LSP to the PCE. It then generates an LSP State Report (PCRpt) for the LSP, carrying the LSP-ID and the delegation bit. 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 LSP-ID, which uniquely identifies the LSP.

The LSPA Object included in these messages MUST carry the SYMBOLIC-PATH-NAME TLV which will be used to correlate between the PCC-assigned LSP-ID and the LSP. See details in [Section 5](#).

Removal of PCE-initiated LSPs is done by the PCE by setting the R flag in the LSP Object in the PCUpd message. Upon receiving the PCUpd message with the R Flag set, the PCC deletes the LSP. See details in [Section 5](#).

Once instantiated, a PCRpt is generated for the LSP, with the delegation bit set. After this, the delegation procedures for PCE-initiated LSPs are the same as for PCC initiated LSPs. Upon session failure, PCE-initiated LSPs are not immediately removed, in order to avoid LSP flap and service interruption. However, to allow for network cleanup without manual intervention, such "orphan" PCE-initiated LSPs must be either adopted by a different PCE or cleaned up within a time interval. This time is negotiated between PCE and PCC at session initialization time. See details in [Section 6](#).

4. Support of PCE-initiated LSPs

A PCC indicates its ability to support PCE provisioned dynamic LSPs during the PCEP Initialization Phase. The Open Object in the Open message contains the "Stateful PCE Capability" TLV, defined in [\[I-D.ietf-pce-stateful-pce\]](#).

A new flag, the I (LSP-INSTANTIATION-CAPABILITY) flag is introduced to indicate support for instantiation of PCE-initiated LSPs. A PCE wishing to initiate LSPs, can do so only for PCCs that advertised this capability and a PCC will follow the procedures described in this document only on sessions where the PCE advertised the I flag. A PCE or PCC that advertise support of LSP initiation MUST also advertise a cleanup time for the removal of such LSPs. The cleanup time is advertised via a new TLV in the Open Object, the LSP-CLEANUP TLV, discussed in [Section 6](#), and the value is negotiated to the lower one advertised on a session.

4.1. Stateful PCE Capability TLV

The format of the STATEFUL-PCE-CAPABILITY TLV is shown in the following figure:




Figure 1: STATEFUL-PCE-CAPABILITY TLV format

The type of the TLV is defined in [[I-D.ietf-pce-stateful-pce](#)] and it has a fixed length of 4 octets.

The value comprises a single field - Flags (32 bits). The U and S bits are defined in [[I-D.ietf-pce-stateful-pce](#)].

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 will 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 creation

To create a PCE-initiated LSP, a PCE sends a PCCreate message to a PCC, which include a set of objects and TLVs describing the LSP to be instantiated. The message format, the objects and TLVs are discussed separately below.

5.1. The LSP Create Message

A Path Computation LSP Create message (also referred to as PCCreate message) is a PCEP message sent by a PCE to a PCC to trigger an LSP instantiation. The Message-Type field of the PCEP common header for the PCCreate message is set to [TBD].

The PCCreate message MUST include the END-POINTS and the LSPA objects. In the LSPA object, it MUST include the SYMBOLIC-PATH-NAME TLV for the LSP. The PCCreate message MAY include other attributes for the LSP. If specified, the PCC MUST use them for the LSP instantiation, otherwise it MUST use its locally configured values. The error messages will be specified in a future version of this document.

The format of a PCCreate message is as follows:

```
<PCCreate Message> ::= <Common Header>
                        <lsp-instantiation-list>
```

Where:

```
<lsp-instantiation-list> ::= <lsp-instantiation-request>[<lsp-instantiation-
list>]
```

```
<lsp-instantiation-request> ::= <END-POINTS>
                                <LSPA>
                                [<ERO>]
                                [<BANDWIDTH>]
                                [<metric-list>]
```

Where:

```
<metric-list> ::= <METRIC>[<metric-list>]
```

The END-POINTS Object contains the source and destination addresses for provisioning the PCE-initiated LSP. If the END-POINTS Object is missing, the PCC MUST send a PCErr message with Error-type=6 (Mandatory Object missing) and Error-value=3 (END-POINTS Object missing).

The LSPA Object MUST include the SYMBOLIC-PATH-NAME TLV, which will be used to correlate between the PCC-assigned LSP-ID and the LSP. 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=TBD (Invalid Parameter) and Error-value=TBD (Bad Symbolic Path Name). The only exception to this rule is for LSPs for which the LSP-cleanup timer is running (see [Section 6](#)).

PCE-initiated removal of a PCE-initiated LSP is done by setting the R (remove) flag in the LSP Object in the PCUpd request from the PCE. The definition of the R bit is updated as follows:

R (Remove - 1 bit): On PCRpt messages the R Flag indicates that the LSP has been removed from the PCC. Upon receiving a PCRpt message with the R Flag set to 1, the PCE SHOULD remove all state related to the LSP from its database. In PCUpd messages the R flag indicates that the PCE wishes to disable the LSP. Upon receiving the PCUpd message with the R Flag set for a PCE-initiated LSP, the PCC tears

down the LSP and removes its state.

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 TBD "PCE-initiated limit reached" and is free to drop any incoming PCUpd messages without additional processing.

A PCC SHOULD relay to the PCE errors it encounters in the setup of PCE-initiated LSP. The error codes and error processing will be detailed in a future version of this document.

6. LSP delegation and cleanup

6.1. LSP delegation procedures

PCE-initiated LSPs are automatically delegated by the PCC to the PCE upon instantiation. The PCC MUST delegate the LSP to the PCE by setting the delegation bit to 1 in the PCRpt that includes the assigned LSP-Id. All subsequent messages from the PCC 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 results in a PCErr message of type 19 (Invalid Operation) and value TBD "Delegation for PCE-initiated LSP cannot be revoked".

A PCE MAY return a delegation to the PCC, to allow for LSP transfer between PCEs. Doing so MUST trigger the LSP cleanup timer described in [Section 6.2](#).

Control over PCE-initiated LSPs reverts to the PCC at the expiration of the delegation timeout. To obtain control of a PCE-initiated LSP, a PCE (either the original or one of its backups) sends a PCCreate message specifying the endpoints and symbolic name (the same process used when initiating an LSP from the PCE). See more in the next section.

6.2. LSP cleanup procedures

The LSP cleanup 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 LSP cleanup timer is advertised in the session open message via a mandatory TLV for sessions where PCE-initiated

LSPs are supported. The timer is started upon PCEP session failure and is stopped when the LSP is delegated to a PCE. Both PCE and PCC advertise a value for this timer, and the timer value is negotiated to the lower value of the two.

6.2.1. LSP-CLEANUP TLV

The LSP-CLEANUP TLV is advertised in the Open Object and is mandatory when the I flag is set in the STATEFUL-PCE-CAPABILITY TLV. The LSP-CLEANUP TLV contains the time in seconds that the PCC has to wait before cleaning up any PCE-initiated LSPs belonging to a particular PCEP session when a PCEP session terminates. Both PCE and PCC advertise a value for the cleanup time, and the cleanup timer is set to the lower of the two. The timer is triggered on PCEP session failure and reset when the LSP is delegated to a PCE.

Failure to include the mandatory LSP-CLEANUP TLV in the Open Object when the I flag is set MUST trigger PCErr of type 6 (Mandatory Object missing) and value TBD (LSP-CLEANUP TLV missing).

The format of the LSP-CLEANUP TLV is shown in the following figure:




Figure 2: LSP-CLEANUP TLV format

The type of the TLV is TBD and it has a fixed length of 4 octets.

The value comprises a single field, the LSP cleanup timeout value.

The time in seconds to wait before cleaning up PCE-initiated LSPs. Zero means immediate removal. The value 0xFFFFFFFF is reserved.

A PCE may take control of the dynamic LSPs for which the LSP cleanup timer is running by sending an PCCreate request for the LSP. In this case, the "Bad Symbolic Path Name" error MUST NOT be generated, the LSP MUST be delegated and the cleanup timer MUST be stopped.

7. IANA considerations

7.1. PCEP-Error Object

This document defines new Error-Type and Error-Value for the following new error conditions:

Error-Type	Meaning
6	Mandatory Object missing Error-value=8: LSP cleanup TLV missing
19	Invalid operation Error-value=TBD: PCE-initiated LSP limit reached Error-value=TBD: Delegation for PCE-initiated LSP cannot be revoked

7.2. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs:

Value	Meaning	Reference
???	LSP cleanup	This document

8. 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.

8.1. 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 TBD "PCE-initiated LSP limit reached" (XXX TBD add to the IANA section) and is free to drop any incoming PCUpd messages without additional processing.

Rapid flaps triggered by the PCE can also be an attack vector. This will be discussed in a future version of this document.

9. Acknowledgements

We would like to thank Jan Medved, Ambrose Kwong and Raveendra Trovi for their contributions to this document.

10. References

10.1. Normative References

- [I-D.ietf-pce-stateful-pce]
Crabbe, E., Medved, J., Varga, R., and I. Minei, "PCEP Extensions for Stateful PCE",
[draft-ietf-pce-stateful-pce-01](#) (work in progress),
July 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.
- [RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", [RFC 5088](#), January 2008.
- [RFC5089] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", [RFC 5089](#), January 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [RFC5511] Farrell, A., "Routing Backus-Naur Form (RBNF): A Syntax

Used to Form Encoding Rules in Various Routing Protocol Specifications", [RFC 5511](#), April 2009.

10.2. Informative References

- [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", [RFC 2702](#), September 1999.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), January 2001.
- [RFC3346] Boyle, J., Gill, V., Hannan, A., Cooper, D., Awduche, D., Christian, B., and W. Lai, "Applicability Statement for Traffic Engineering with MPLS", [RFC 3346](#), August 2002.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.
- [RFC4657] Ash, J. and J. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", [RFC 4657](#), September 2006.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), October 2008.
- [RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", [RFC 5394](#), December 2008.
- [RFC5557] Lee, Y., Le Roux, J.L., King, D., and E. Oki, "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", [RFC 5557](#), July 2009.

Authors' Addresses

Edward Crabbe
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: edc@google.com

Ina Minei
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: ina@juniper.net

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

Email: msiva@cisco.com

Robert Varga
Pantheon Technologies SR0
Mlynske Nivy 56
Bratislava 821 05
Slovakia

Email: robert.varga@pantheon.sk

