

PCE Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: March 18, 2015

Q. Wu  
D. Dhody  
Huawei  
M. Boucadair  
C. Jacquenet  
France Telecom  
J. Tantsura  
Ericsson  
September 14, 2014

**PCEP Extensions for traffic steering support in Service Function  
Chaining  
draft-wu-pce-traffic-steering-sfc-05**

**Abstract**

This document provides an overview of the usage of Path Computation Element (PCE) with Service Function Chaining (SFC); which is described as the definition and instantiation of an ordered set of such service functions (such as firewalls, load balancers), and the subsequent "steering" of traffic flows through those service functions.

This document specifies extensions to the Path Computation Element Protocol (PCEP) that allow a stateful PCE to compute and instantiate Service Function Paths (SFP).

**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 March 18, 2015.

## Copyright Notice

Copyright (c) 2014 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Conventions used in this document . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Service Function Paths and PCE . . . . .	<a href="#">3</a>
<a href="#">4.</a>	Overview of PCEP Operation in SFC-enabled Networks . . . . .	<a href="#">5</a>
<a href="#">4.1.</a>	SFP Instantiation . . . . .	<a href="#">5</a>
<a href="#">4.2.</a>	SFP Withdrawal . . . . .	<a href="#">5</a>
<a href="#">4.3.</a>	SFP Delegation and Cleanup . . . . .	<a href="#">5</a>
<a href="#">4.4.</a>	SFP State Synchronization . . . . .	<a href="#">5</a>
<a href="#">4.5.</a>	SFP Update and Report . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Object Formats . . . . .	<a href="#">6</a>
<a href="#">5.1.</a>	The OPEN Object . . . . .	<a href="#">6</a>
<a href="#">5.2.</a>	The LSP Object . . . . .	<a href="#">6</a>
<a href="#">5.2.1.</a>	SFP Identifiers TLV . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Backward Compatibility . . . . .	<a href="#">7</a>
<a href="#">7.</a>	Relationship to SR . . . . .	<a href="#">7</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">8</a>
<a href="#">9.</a>	IANA Considerations . . . . .	<a href="#">8</a>
<a href="#">10.</a>	References . . . . .	<a href="#">8</a>
<a href="#">10.1.</a>	Normative References . . . . .	<a href="#">8</a>
<a href="#">10.2.</a>	Informative References . . . . .	<a href="#">8</a>

## [1.](#) Introduction

Service chaining enables the creation of composite services that consist of an ordered set of Service Functions (SF) that must be applied to packets and/or frames selected as a result of classification as described in [[I-D.ietf-sfc-architecture](#)] and referred to as Service Function Chain (SFC). A Service Function Path (SFP) is the instantiation of a SFC in the network. Packets follow a Service Function Path from a classifier through the requisite Service Functions (SF).



[RFC5440] describes the Path Computation Element Protocol (PCEP) as 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).

[I-D.ietf-pce-stateful-pce] specifies extensions to PCEP to enable stateful control of MPLS TE LSPs. [I-D.ietf-pce-pce-initiated-lsp] provides the fundamental extensions needed for stateful PCE-initiated LSP instantiation.

This document specifies extensions to the PCEP that allow a stateful PCE to compute and instantiate Service Function Paths (SFP).

## **2. Conventions used in this document**

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](#) [RFC2119].

The following terminologies are used in this document:

PCC: Path Computation Client.

PCE: Path Computation Element.

PCEP: Path Computation Element Protocol.

PDP: Policy Decision Point.

SF: Service Function.

SFC: Service Function Chain.

SFP: Service Function Path.

SFF: Service Forwarder Function.

UNI: User-Network Interface.

## **3. Service Function Paths and PCE**

Services are constructed as a sequence of SFs that represent an SFC, where a SF can be a virtual instance or be embedded in a physical network element, and one or more SFs may be supported by the same physical network element. A SFC creates an abstracted view of a service and specifies the set of required SFs as well as the order in which they must be executed.



When an SFC is instantiated into the network it is necessary to select the specific instances of SFs that will be used, and to create the service topology for that SFC using SF network locators. Thus, instantiation of the SFC results in the creation of a Service Function Path (SFP) and is used for forwarding packets through the SFC. In other words, an SFP is the instantiation of the defined SFC as described in [[I-D.ietf-sfc-architecture](#)].

The selection of SFP can be based on a set of policy attributes (forwarding and routing, QoS, security, etc., or a combination thereof), ranging from simple to more elaborate selection criteria and the use of stateful PCE with extensions to PCEP are one such way to achieve this.

Stateful pce [[I-D.ietf-pce-stateful-pce](#)] specifies a set of extensions to PCEP to enable stateful control of TE LSPs. [[I-D.ietf-pce-pce-initiated-lsp](#)] provides the fundamental motivations and extensions needed for stateful PCE-initiated LSP instantiation. This document specifies extensions that allow a stateful PCE to compute and instantiate Service Function Paths (SFP) via PCEP.

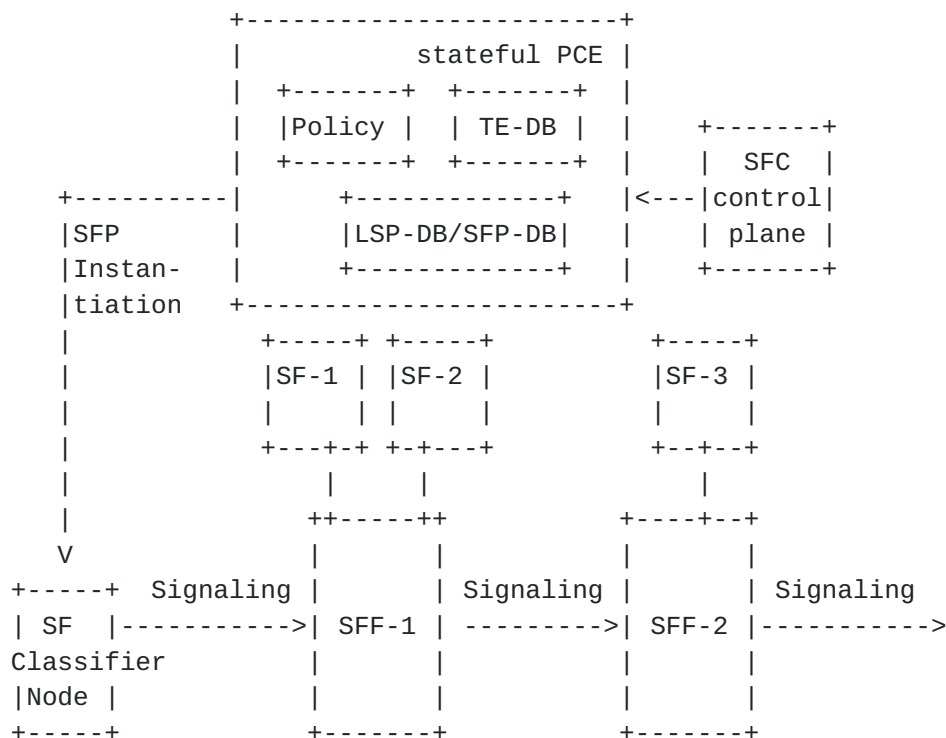


Figure 1: PCE based SFP instantiation

SFC Control plane components are responsible for maintaining SFC Policy Tables and enforcing appropriate policies in SF Classifier and SFF Nodes as described in



[[I-D.ietf-sfc-architecture](#)][I-D.ww-sfc-control-plane]. The SFC Control plane component can be seen as a policy Decision point (PDP, [[RFC5394](#)]). Such PDP can then operate a stateful PCE and its instantiation mechanism to compute and instantiate Service Function Paths (SFP). The PCE maybe co-located with the SFC Control plane component or an external entity.

#### **4. Overview of PCEP Operation in SFC-enabled Networks**

A PCEP speaker indicates its ability to support PCE provisioned dynamic SFP paths during the PCEP Initialization phase via a mechanism described in [Section 5.1](#). A PCE can initiate SFPs 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 this capability. .

As per section 5.1 of [[I-D.ietf-pce-pce-initiated-lsp](#)], the PCE sends a Path Computation LSP Initiate Request (PCInitiate) message to the PCC to instantiate or delete a LSP. This document makes no change to the PCInitiate message format but extends LSP objects described in [Section 5.2](#).

##### **[4.1.](#) SFP Instantiation**

The Instantiation operation of a SFP is the same as defined in [section 5.3](#)[[I-D.ietf-pce-pce-initiated-lsp](#)]. Rules of processing and error codes remain unchanged.

##### **[4.2.](#) SFP Withdrawal**

The withdrawal operation of a SFP is the same as defined in [section 5.4](#) of [[I-D.ietf-pce-pce-initiated-lsp](#)] : the PCE sends an LSP Initiate Message with an LSP object carrying the PLSP-ID of the SFP to be removed and an SRP object with the R flag set (LSP-REMOVE as per section 5.2 of [[I-D.ietf-pce-pce-initiated-lsp](#)]). Rules of processing and error codes remain unchanged.

##### **[4.3.](#) SFP Delegation and Cleanup**

SFP delegation and cleanup operations are similar to those defined in section 6 of [[I-D.ietf-pce-pce-initiated-lsp](#)]. Rules of processing and error codes remains unchanged.

##### **[4.4.](#) SFP State Synchronization**

State Synchronization operations described in Section 5.4 of [[I-D.ietf-pce-stateful-pce](#)] can be applied for SFP state maintenance as well.





#### 4.5. SFP Update and Report


A PCE can send an SFP Update request to a PCC to update one or more attributes of an SFP and to re-signal the SFP with the updated attributes. A PCC can send an SFP state report to a PCE, and which contains the SFP State information. The mechanism is described in [I-D.ietf-pce-stateful-pce] and can be applied for SFPs as well.

### 5. Object Formats

#### 5.1. The OPEN Object

This document defines a new optional TLV for use in the OPEN Object to indicate the PCEP speaker's Service function Chaining capability.

The SFC-PCE-CAPABILITY TLV is an optional TLV for use in the OPEN Object to advertise the SFC capability during the PCEP session. The format of the SFC-PCE-CAPABILITY TLV is shown in the following Figure 2 :



SFC-PCE-CAPABILITY TLV Format

The code point for the TLV type is to be defined by IANA. The TLV length is 4 octets.

The value is TBD.

As per [I-D.ietf-pce-stateful-pce], a PCEP speaker advertises the capability of instantiating PCE-initiated LSPs via the Stateful PCE Capability TLV (LSP-INSTANTIATION-CAPABILITY bit) conveyed in an Open message. The inclusion of the SFC-PCE-CAPABILITY TLV in an OPEN object indicates that the sender is SFC-capable. Both mechanisms indicate the SFP instantiation capability of the PCEP speaker.

#### 5.2. The LSP Object

The LSP object is defined in [I-D.ietf-pce-pce-initiated-lsp] and included here for reference (Figure 3).



```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               | Flags |F|C| 0|A|R|S|D|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                               TLVs                               //
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

### LSP Object Format

A new flag, called the SFC (F) flag, is introduced. The F Flag set to 1 indicates that this LSP is actually an SFP. The C flag will also be set to indicate it was created via a PCInitiate message.

#### 5.2.1. SFP Identifiers TLV

The SFP Identifiers TLV MUST be included in the LSP object for Service Function Paths (SFP).

The format and operations are TBD.

### 6. Backward Compatibility

The SFP instantiation capability PCEP protocol extensions described in this document MUST NOT be used if PCCs or the PCE did not advertise its SFP instantiation stateful capability, as per [Section 5.1](#). If this is not the case and stateful operations on SFPs are attempted, then a PCError with error-type 19 (Invalid Operation) and error-value TBD needs to be generated.

[Editor Note: more information on exact error value is needed]

### 7. Relationship to SR

Segment Routing (SR) technology leverages the source routing and tunneling paradigms where a source node can choose a path without relying on hop-by-hop signaling. A stateful PCE can be used for computing one or more SR-TE paths taking into account various constraints and objective functions. Once a path is chosen, the stateful PCE can instantiate an SR-TE path on a PCC using PCEP extensions specified in [[I-D.sivabalan-pce-segment-routing](#)].

The SFP instantiation mechanism described in this document is not tightly coupled to any SFP signaling mechanism. Thus, SR-based SFP can also utilize the mechanism described here and does not need any other specific protocol extensions.



## **8. Security Considerations**

The security considerations described in [[RFC5440](#)] and [[I-D.ietf-pce-pce-initiated-lsp](#)] are applicable to this specification. No additional security measure is required.

## **9. IANA Considerations**

TBD

## **10. References**

### **10.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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-09](#) (work in progress), June 2014.
- [I-D.ietf-pce-pce-initiated-lsp]  
Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", [draft-ietf-pce-pce-initiated-lsp-01](#) (work in progress), June 2014.

### **10.2. Informative References**

- [RFC2753] Yavatkar, R., Pendarakis, D., and R. Guerin, "A Framework for Policy-based Admission Control", [RFC 2753](#), January 2000.
- [RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", [RFC 5394](#), December 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [I-D.ietf-sfc-architecture]  
Halpern, J. and C. Pignataro, "Service Function Chaining (SFC) Architecture", [draft-ietf-sfc-architecture-01](#) (work in progress), September 2014.



[I-D.ww-sfc-control-plane]

Li, H., Wu, Q., Huang, O., Boucadair, M., Jacquenet, C.,  
and W. Haeffner, "Service Function Chain Control Plane  
Framework", [draft-ww-sfc-control-plane-02](#) (work in  
progress), July 2014.

[I-D.sivabalan-pce-segment-routing]

Sivabalan, S., Medved, J., Filsfils, C., Crabbe, E.,  
Raszuk, R., Lopez, V., and J. Tantsura, "PCEP Extensions  
for Segment Routing", [draft-sivabalan-pce-segment-  
routing-03](#) (work in progress), July 2014.

Authors' Addresses

Qin Wu  
Huawei  
101 Software Avenue, Yuhua District  
Nanjing, Jiangsu 210012  
China

EMail: sunseawq@huawei.com

Dhruv Dhody  
Huawei  
Leela Palace  
Bangalore, Karnataka 560008  
INDIA

EMail: dhruv.ietf@gmail.com

Mohamed Boucadair  
France Telecom  
Rennes 35000  
France

EMail: mohamed.boucadair@orange.com

Christian Jacquenet  
France Telecom  
Rennes 35000  
France

EMail: christian.jacquenet@orange.com





Jeff Tantsura  
Ericsson  
300 Holger Way  
San Jose, CA 95134  
US

EMail: [Jeff.Tantsura@ericsson.com](mailto:Jeff.Tantsura@ericsson.com)