PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 1, 2016

Q. Wu
D. Dhody
Huawei
M. Boucadair
C. Jacquenet
France Telecom
J. Tantsura
Ericsson

June 30, 2015

PCEP Extensions for traffic steering support in Service Function Chaining
draft-wu-pce-traffic-steering-sfc-07

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 January 1, 2016.

Table of Contents

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) and Service Function Forwarders (SFF).

[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 function path for that SFC using SF network locators. Thus, the instantiation of a SFC results in the establishment of a Service Function Path, either a la hop-by-hop through the ordered sequence of SF functions, or in a pre-computed, traffic-engineered fashion. 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.

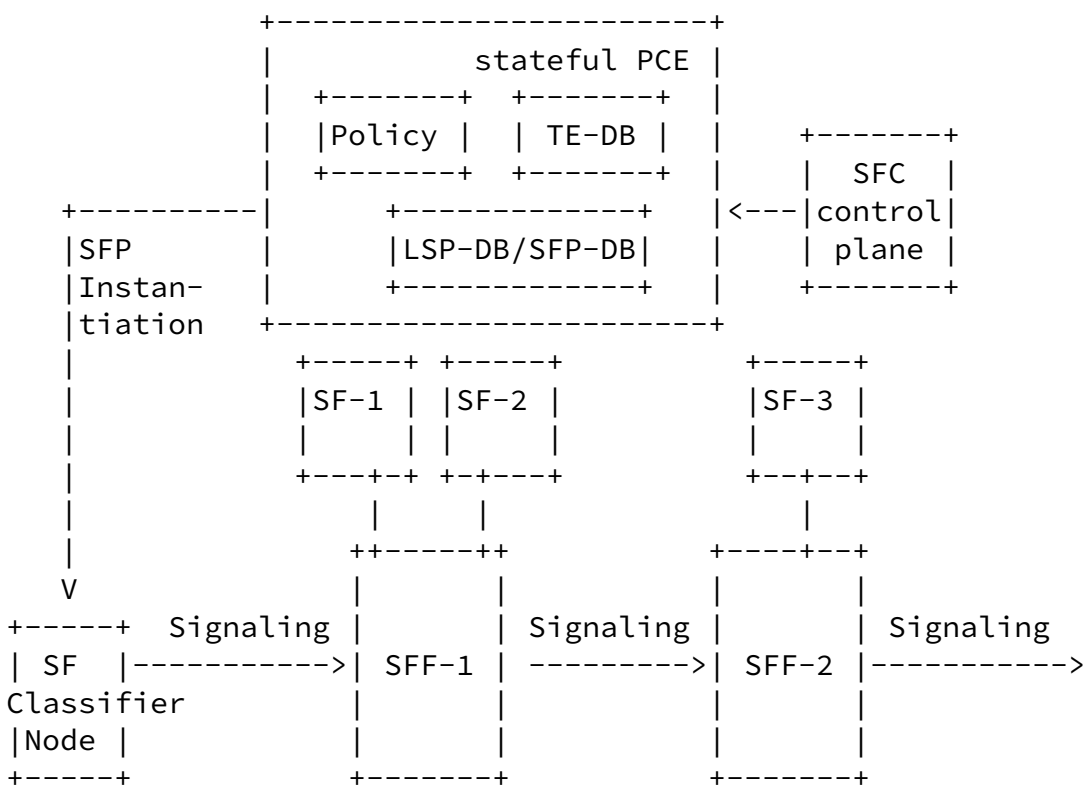
stateful PCE
Policy
TE-DB
LSP-DB/SFP-DB
SFC control plane
SFP Instantiation
SF-1
SF-2
SF-3
SF Classifier Node
Signaling
SFF-1
SFF-2

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 operates 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. The Explicit Route Object (ERO) can be used to encode either a sequence of SF functions or a combination of SFs and SFFs to establish a SFP. If the said SFFs and

SFs can be identified with an IP address, the IP sub-object can be used as a SF/SFF identification means. This document makes no change to the PCInitiate message format but extends LSP objects described in Section 5.2.

Editor-Note: In case a PCE-Initiated Signaling mechanism is used to setup the service function path, then does the classifier / PCE-Initiated signaling protocol needs to understand if the IP address is for SFF or SF or the signaling protocol is only used to signal IP address for SFs?

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 followingFigure 2 :

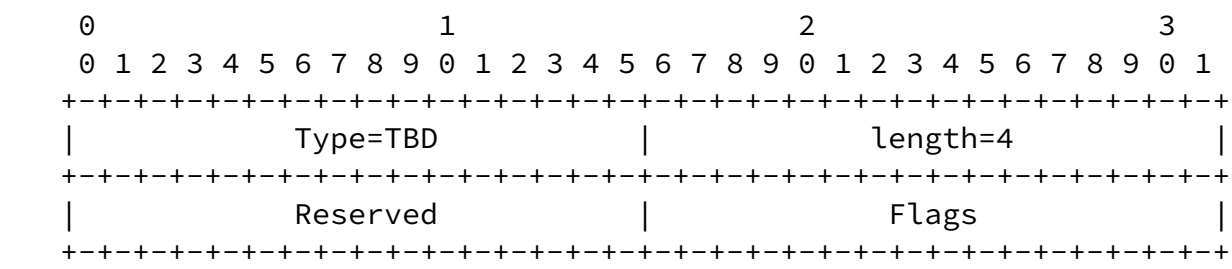

SFC-PCE-CAPABILITY TLV Format

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

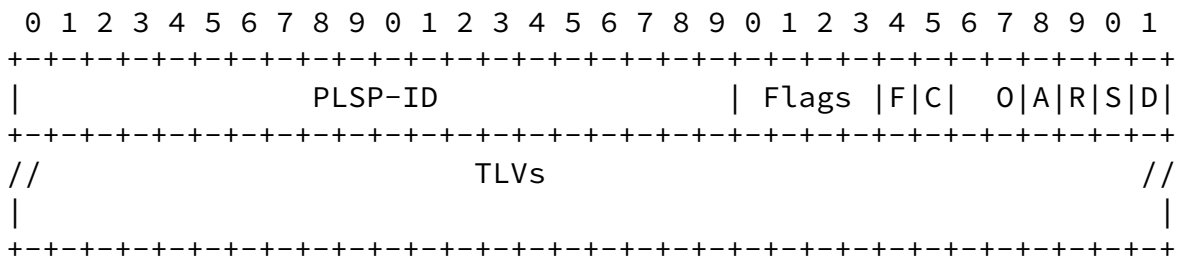

LSP Object Format

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 SFP Identifier TLV is used by the classifier to enable SFP selection for the traffic,i.e.,direct traffic to specific SFP[I-D.ietf-sfc-architecture]. The SFP Identifier carried in the SFC encapsulation can be further used by SFF to select service functions and next SFF,e.g., enable a packet that repeatedly arrives at the same SFF to get the correct services provided each time it arrives, and to go to the correct next SFF each time it arrives.

The format of SFP Identifier TLV is shown in the following figure.

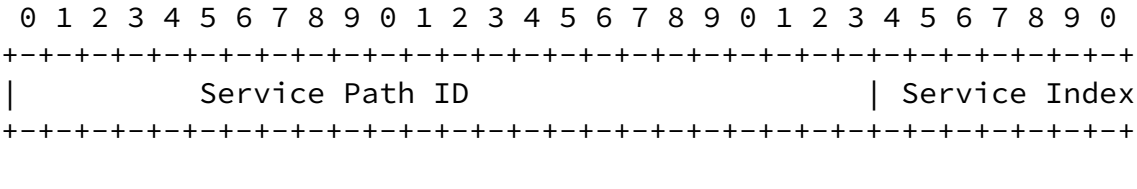

Service path ID (SPI): 24 bits
Service index (SI): 8 bits

SPI: identifies a service path. The same ID is used by the participating nodes for path setup/selection. An administrator can use the SPI for reporting and troubleshooting packets along a specific path. SPI along with PLSP-ID is used in PCEP to identify the Service Path.

SI: provides location within the service path.

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 PCErr 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. SFP signaling and forwarding consideration

The SFP instantiation mechanism described in this document is not tightly coupled to any SFP signaling mechanism. For example,SR-based approach [I-D.ietf-pce-segment-routing] can utilize the mechanism described here and does not need any other specific protocol extensions. Generic SFC Encapsulation [I-D.quinn-sfc-nsh] can also be used together with the mechanism described here to enable SFP forwarding.

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. Acknowledgements

Many thanks to Ron Parker, Hao Wang,Dave Dolson,Jing Huang,Joel M. Halpern for the discussion in formulating the content for the draft.

11. References

11.1. 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-11 (work in progress), April 2015.

[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-04 (work in progress), April 2015.

11.2. 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-09 (work in progress), June 2015.

[I-D.ww-sfc-control-plane]
Li, H., Wu, Q., Boucadair, M., Jacquenet, C., Haeffner, W., Lee, S., Parker, R., Dunbar, L., Malis, A., Halpern, J., Reddy, T., and P. Patil, "Service Function Chaining (SFC) Control Plane Components & Requirements", draft-ww-sfc-control-plane-06 (work in progress), June 2015.

[I-D.ietf-pce-segment-routing]
Sivabalan, S., Medved, J., Filsfils, C., Crabbe, E., Lopez, V., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", draft-ietf-pce-segment-routing-05 (work in progress), May 2015.

[I-D.quinn-sfc-nsh]

Quinn, P., Guichard, J., Surendra, S., Smith, M., Henderickx, W., Nadeau, T., Agarwal, P., Manur, R., Chauhan, A., Halpern, J., Majee, S., Elzur, U., Melman, D., Garg, P., McConnell, B., Wright, C., and K. Kevin, "Network Service Header", draft-quinn-sfc-nsh-07 (work in progress), February 2015.

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