Pce Working Group   X. Xu
Internet-Draft   J. You
Intended status: Standards Track   Huawei
Expires: April 30, 2015   S. Sivabalan
                                                           Cisco Systems
                                                                 H. Shah
                                                                   Ciena
                                                            L. Contreras
                                                         Telefonica I+D
                                                        October 27, 2014

PCEP Extensions for SFC in SPRING Networks
draft-xu-pce-sr-sfc-02

Abstract

[I-D.xu-spring-pce-based-sfc-arch] describes a PCE-based SFC architecture in which the PCE is used to compute service function paths in SPRING networks. Based on the above architecture, this document describes extensions to the Path Computation Element Protocol (PCEP) that allow a PCE to compute and instantiate service function paths in SPRING networks. The extensions specified in this document are applicable to both the stateless PCE model and 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 30, 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

1. Introduction

Service Function Chaining [I-D.ietf-sfc-architecture] provides a flexible way to construct services. When applying a particular Service Function Chain (SFC) to the traffic classified by the Classifier, the traffic needs to be steered through an ordered set of Service Functions (SF) in the network. This ordered set of SFs in the network, referred to as a Service Function Path (SFP), is an instantiation of the SFC in the network. For example, as shown in Figure 1, an SFP corresponding to the SFC of {SF1, SF3} can be expressed as {SFF1, SF1, SFF2, SF3}.

```
            +-------+
         +--+  PCE  |
         |  +-------+
         |
         |
         |
         |   +-------------------------------------------------+
         |   |                                     SR Netowrks |
         |   |         +-----+   +-----+                       |
         |   |         | SF1 |   | SF2 |                       |
         |   |         +--+--+   +--+--+                       |
         |   |            |         |                          |
         |   |         ^  |         |                          |
         |   |      (2)|  +---+ +---+                          |
         |   |         +--+   | |                              |
        ++---------+      |   | |          +--------------+    |
        |    +----+|      V   | |          |   +-----+    |    |
        |    |PCC || (1)  +---+-+----+ (3) |   | SF3 |    |    |
    --> |SFC +----+|----> |   SFF1   |---->|   +-----+    |---->
    ----+Classifier+------+          +-----+    SFF2      +--------
        +----------+      +----------+     +--------------+    |
             |                                                 |
             +-------------------------------------------------+
```

Figure 1: PCE-based Service Function Chaining in SR Network

[I-D.xu-spring-pce-based-sfc-arch] describes a PCE-based SFC architecture in which the PCE is used to compute an SFP (i.e., instantiate a service function chain) in SPRING networks (a.k.a., Segment Routing networks or SR networks in short). This document describes extensions to the PCEP on basis of that architecture. The extensions specified in this document are applicable to both the stateless PCE model defined in [RFC5440] and the stateful PCE model defined in [I-D.ietf-pce-stateful-pce].

2. Terminology

This section contains definitions for terms used frequently throughout this document. However, many additional definitions can be found in [RFC5440], [I-D.sivabalan-pce-segment-routing] and [I-D.xu-spring-pce-based-sfc-arch].

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Protocol

ERO: Explicit Route Object

Service Function Chain (SFC): A service function chain defines a set of abstract service functions and ordering constraints that must be applied to packets and/or frames selected as a result of classification.

SF Identifier (SF ID): A unique identifier that represents a service function within an SFC-enabled domain.

Service Function Forwarder (SFF): A service function forwarder is responsible for delivering traffic received from the network to one or more connected service functions according to information carried in the SFC encapsulation, as well as handling traffic coming back from the SF.

Service Function Path (SFP): The SFP provides a level of indirection between the fully abstract notion of service chain as a sequence of abstract service functions to be delivered, and the fully specified notion of exactly which SFF/SFs the packet will visit when it actually traverses the network. Specifically, it is an ordered list of SFFs and SF IDs.

Compact SFP: An ordered list of SFFs.

SID: Segment Identifier

Service Function SID : A locally unique SID indicating a particular service function on an SFF.

SR: Segment Routing

SR-specific SFP: An ordered list of node SIDs (representing SFFs) and Service Function SIDs.

Compact SR-specific SFP: An ordered list of node SIDs (representing SFFs).

3. Overview of PCEP Extensions for SFC in SR Networks

As discussed in [I-D.xu-spring-pce-based-sfc-arch], the PCC provides an ordered list of SF IDs to the PCE and indicates to the PCE that what type SFs and paths are requested (e.g., an SFP, or a compact SFP, or an SR-specific SFP, or a compact SR-specific SFP) through the path computation request message, and then the PCE responds with a corresponding path through the path computation response message. This specification is applicable to both stateful and stateless PCEs.

4. PCEP Message Extensions for SR-based SFC

4.1. PCReq Message

This document does not specify any changes to the PCReq message format. This document requires the PATH-SETUP-TYPE TLV [I-D.sivabalan-pce-lsp-setup-type] to be carried in the RP Object in order for a PCC to request a particular type of path. Four new Path Setup Types need to be defined for SR-based SFC, or SR-SFC in short (Section 5.2). This document also requires the Include Route Object (IRO) to be carried in the PCReq message in order for a PCC to specify SFC. A new IRO sub-object type needs to be defined for SF (Section 5.3).

4.2. PCRep Message

This document defines the format of the PCRep message carrying an SFP. The message is sent by a PCE to a PCC in response to a previously received PCReq message, where the PCC requested an SFP. The format of the SFC-specific PCRep message is as follows:

```
<PCRep Message>::=<Common Header>
                  <response-list>

Where:

 <response-list>::=<response>[<response-list>]

 <response>::=<RP>
             [<NO-PATH>]
             [<path-list>]


 Where:

  <path-list>::=<SR-SFC-ERO>[<path-list>]
```

The RP and NO-PATH Objects are defined in [RFC5440]. The <SR-SFC-ERO> object contains the SFP and is defined in Section 5.4.

4.3. PCUpd Message

This document defines the format of the PCUpd message carrying an SFP update. The message is sent forwardly by a PCE to a PCC to update an previously computed SFP.

The format of the PCUpd message is as follows:

```
<PCUpd Message>::=<Common Header>
                  <udpate-request-list>

Where:

 <udpate-request-list>::=<udpate-request>[<udpate-request-list>]

 <udpate-request>::=<SRP><path-list>


 Where:

  <path-list>::=<SR-SFC-ERO>[<path-list>]
```

4.4. PCRpt Message

PCPRpt message sent from a PCC to PCE as a respond to a PCUpd message or in an unsolicited manner (e.g., during state synchronization).

The format of the PCUpd message is as follows:

```
<PCUpd Message>::=<Common Header>
                  <state-report-list>

Where:

 <state-report-list>::=<state-report>[<state-report-list>]

 <state-report>::=[<SRP>]<path-list>


 Where:

  <path-list>::=<SR-SFC-ERO>[<path-list>]
```

5. Object Formats

5.1. OPEN Object

This document defines a new optional TLV for use in the OPEN Object.

5.1.1. SR-SFC PCE Capability TLV

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

```
  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
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |         Type=TBD              |       Length=4                |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |          Reserved             |  Flags        |      MSD      |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: SR-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 32-bit value is formatted as follows. The "Maximum SID Depth" (1 octet) field (MSD) specifies the maximum number of SIDs that a PCC is capable of imposing on a packet. The "Flags" (1 octet) and "Reserved" (2 octets) fields are currently unused, and MUST be set to zero and ignored on receipt.

5.1.1.1. Negotiating SR-SFC Capability

The SR-SFC capability TLV is contained in the OPEN object. By including the TLV in the OPEN message to a PCE, a PCC indicates its support for SFPs. By including the TLV in the OPEN message to a PCC, a PCE indicates that it is capable of computing SFPs.

5.2. RP/SRP Object

In order to setup an SFP, the RP or SRP object MUST carry a PATH-SETUP-TYPE TLV specified in [I-D.sivabalan-pce-lsp-setup-type]. This document defines four new Path Setup Types (PST) for SR-SFC as follows:

PST = 2: The path is an SFP.

PST = 3: The path is a compact SFP.

PST = 4: The path is an SR-specific SFP.

PST = 5: The path is a compact SR-specific SFP.

5.3. Include Route Object

The IRO (Include Route Object) MUST be carried within PCReq messages to indicate a particular SFC. Furthermore, the IRO MAY be carried in PCRep messages. When carried within a PCRep message with the NO-PATH object, the IRO indicates the set of service functions that cause the PCE to fail to find a path.

This document defines a new sub-object type for the SR-SFC as follows:

Type	Sub-object
5	Service Function ID

5.4. SR-SFC-ERO Object

Generally speaking, an SR-SFC-ERO object consists of one or more ERO subobjects described in the following sub-sections to represent a particular type of service function path. In the ERO subobject, each SID is associated with an identifier that represents either a service node or a service function. This identifier is referred to as the 'Node or Service Identifier' (NSI). As described later, an NSI can be represented in various formats (e.g., IPv4 address, IPv6 address, SF identifier, etc). Specifically, in the SFP case, the NSI of every ERO subobject contained in the SR-SFC-ERO object represents a service

node or a service function while the SID of each ERO subobject is set to null. In the compact SFP case, the NSI of every ERO subobject contained in the SR-SFC-ERO object only represents an SFF meanwhile the SID of every ERO subobject is set to null. In the SR-specific SFP, the NSI of every ERO subobject contained in the SR-SFC-ERO object represents an SFF or a service function while the SID of every ERO subject MUST NOT be null. In the compact SR-specific SFP, the NSI of every ERO subobject contained in the SR-SFC-ERO object represents an SFF meanwhile the SID of every ERO subobject MUST NOT be null.

5.4.1. SR-SFC-ERO Subobject

An SR-SFC-ERO subobject (as shown in Figure 3) consists of a 32-bit header followed by the SID and the NSI associated with the SID. The SID is a 32-bit or 128 bit number. The size of the NSI depends on its respective type, as described in the following sub-sections.

```
   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
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
  |L|    Type     |     Length    |  NSIT   |  Flags    |P|F|S|C|M|
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
  //         SID (variable:4 or 16 octets)                       //
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
  //              NSI (variable)                                 //
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: SR-SFC-ERO Subobject Format

The fields in the ERO Subobject are as follows:

'L' Flag: indicates whether the subobject represents a loose-hop in the explicit route [RFC3209]. If this flag is unset, a PCC MUST not overwrite the SID value present in the SR-SFC-ERO subobject. Otherwise, a PCC MAY expand or replace one or more SID value(s) in the received SR-SFC-ERO based on its local policy.

Type: is the type of the SR-SFC-ERO Subobject. This document defines the SR-SFC-ERO Subobject type. A new code point will be requested for the SR-SFC-ERO Subobject from IANA.

Length: contains the total length of the subobject in octets, including the L, Type and Length fields. Length MUST be at least 4, and MUST be a multiple of 4.

NSI Type (NSIT): indicates the type of NSI associated with the SID. The NSI-Type values are described later in this document.

Flags: is used to carry any additional information pertaining to SID. Currently, the following flag bits are defined:

M: When this bit is set, the SID value represents an MPLS label stack entry as specified in [RFC5462], where only the label value is specified by the PCE. Other fields (TC, S, and TTL) fields MUST be considered invalid, and PCC MUST set these fields according to its local policy and MPLS forwarding rules.

C: When this bit as well as the M bit are set, then the SID value represents an MPLS label stack entry as specified in [RFC5462], where all the entry's fields (Label, TC, S, and TTL) are specified by the PCE. However, a PCC MAY choose to override TC, S, and TTL values according its local policy and MPLS forwarding rules.

S: When this bit is set, the SID value in the subobject body is null. In this case, the PCC is responsible for choosing the SID value, e.g., by looking up its Traffic Engineering Database (TED) using node/service identifier in the subobject body.

F: When this bit is set, the NSI value in the subobject body is null.

P: When this bit is set, the SID value represents an IPv6 address.

SID: is the 4-octect or 16-octect Segment Identifier

NSI: contains the NSI associated with the SID. Depending on the value of NSIT, the NSI can have different format as described in the following sub-section.

5.4.2. NSI Associated with SID

This document defines the following NSIs:

'IPv4 Node ID': is specified as an IPv4 address. In this case, NSIT and Length are 1 and 12 respectively.

'IPv6 Node ID': is specified as an IPv6 address. In this case, NSIT and Length are 2 and 24 respectively.

'Service Function ID': is specified as an SF ID. In this case, NSIT and Length are TBD.

5.4.3. SR-SFC-ERO Processing

TBD.

6. IANA Considerations

6.1. PCEP Objects

IANA is requested to allocate an ERO subobject type (recommended value= 6) for the SR-SFC-ERO subobject.

6.2. PCEP-Error Object

TBD.

6.3. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs:

Value	Meaning	Reference
27	SR-SFC-PCE-CAPABILITY	This document

6.4. New Path Setup Type

This document defines a new setup type for the PATH-SETUP-TYPE TLV as follows:

Value	Description	Reference
2	The path is an SFP.	This document
3	The path is a compact SFP.	This document
4	The path is an SR-specific SFP.	This document
5	The path is a compact SR-specific SFP.	This document

6.5. New IRO Sub-object Type

This document defines a new IRO sub-object type for the SFC as follows:

Type	Sub-object
5	Service Function ID

7. Security considerations

This document does not introduce any new security considerations.

8. Acknowledgement

TBD.

9. References

9.1. 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-09 (work in progress), June 2014.

[I-D.ietf-sfc-architecture]
Halpern, J. and C. Pignataro, "Service Function Chaining (SFC) Architecture", draft-ietf-sfc-architecture-02 (work in progress), September 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 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.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, February 2009.

9.2. Informative References

[I-D.sivabalan-pce-lsp-setup-type]
Sivabalan, S., Medved, J., Minei, I., Crabbe, E., and R. Varga, "Conveying path setup type in PCEP messages", draft-sivabalan-pce-lsp-setup-type-02 (work in progress), June 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.

[I-D.xu-spring-pce-based-sfc-arch]
Xu, X., You, J., Shah, H., and L. Contreras, "PCE-based SFC Architecture in SR Networks", draft-xu-spring-pce-based-sfc-arch-01 (work in progress), June 2014.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Jianjie You
Huawei
101 Software Avenue, Yuhuatai District
Nanjing, 210012
China

Email: youjianjie@huawei.com

Siva Sivabalan
Cisco Systems
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

Email: msiva@cisco.com

Himanshu Shah
Ciena
3939 North First Street
San Jose, CA 95134
USA

Email: hshah@ciena.com

Luis M. Contreras
Telefonica I+D
Ronda de la Comunicacion, s/n
Sur-3 building, 3rd floor
Madrid, 28050
Spain

Email: luismiguel.contrerasmurillo@telefonica.com
URI: http://people.tid.es/LuisM.Contreras/