

**PCEP Extensions for service segment used in Segment Routing  
draft-dw-pce-service-segment-routing-01**

Abstract

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.

This document specifies extensions to the Path Computation Element Protocol (PCEP) that allow a stateful PCE to compute and instantiate SR-TE paths that also have a local service segments involved.

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 7, 2015.

Copyright Notice

Copyright (c) 2015 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>	<a href="#">Introduction . . . . .</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Conventions used in this document . . . . .</a>	<a href="#">3</a>
<a href="#">2.1.</a>	<a href="#">Terminology . . . . .</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Overview of PCEP Operation in SR Networks for Service Chaining . . . . .</a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Object Formats . . . . .</a>	<a href="#">4</a>
<a href="#">4.1.</a>	<a href="#">The SR-ERO Subobject extension for service segment support . . . . .</a>	<a href="#">5</a>
<a href="#">4.2.</a>	<a href="#">Service Segment SR-ERO Processing . . . . .</a>	<a href="#">6</a>
<a href="#">5.</a>	<a href="#">Backward Compatibility . . . . .</a>	<a href="#">6</a>
<a href="#">6.</a>	<a href="#">Management Considerations . . . . .</a>	<a href="#">6</a>
<a href="#">7.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">6</a>
<a href="#">8.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">9.</a>	<a href="#">References . . . . .</a>	<a href="#">7</a>
<a href="#">9.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">7</a>
<a href="#">9.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">7</a>
<a href="#">Appendix A.</a>	<a href="#">Examples . . . . .</a>	<a href="#">7</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">8</a>

## [1.](#) Introduction

Segment Routing (SR) enables Traffic Engineering (TE) without relying on a hop-by-hop signaling. It depends only on "segments" that are advertised by Interior Gateway Protocols (IGPs). These segments made by -

- o Node Segment
- o Adjacency Segment
- o Anycast Segment
- o IGP-Prefix Segment

Further to this list, a segment may also be identify a particular value added service or service function (SF).

[\[I-D.filsfils-spring-segment-routing-use-cases\]](#) describes using local Service-Segment to stand for a BGP-VPN service in an example. A service-segment may also be used to represent specific treatment offered by SR enabled node(s) in the path, a combination of node-segment and service-segment can be used. The service segment is local to the SR enabled node.



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 [\[I-D.ietf-pce-segment-routing\]](#).

An SR-TE path is defined as a path that consists of one or more SID(s) where each SID is associated with the identifier that represents the node or adjacency corresponding to the SID. This document extends the SR-TE path to use Service-SID(s) in the path as well.

The means by which the PCE learns about the Service-SID (e.g., learnt over a management interface or through a variety of other mechanisms) is beyond scope of this document.

## **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](#)].

### **2.1. Terminology**

The following terminologies are used in this document:

ERO: Explicit Route Object

IGP: Interior Gateway Protocol

LSR: Label Switching Router

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Protocol

SF: Service Function

SFC: Service Function Chaining

SID: Segment Identifier

SR: Segment Routing

SR-ERO: Segment Routed Explicit Route Object



SR Path: Segment Routed Path

SR-TE: Segment Routed Traffic Engineering

### **3. Overview of PCEP Operation in SR Networks for Service Chaining**

In SR networks, an ingress node of an SR path appends all outgoing packets with an SR header consisting of a list of Segment IDs (SIDs). The header has all necessary information to guide the packets from the ingress node to the egress node of the path, and hence there is no need for any signaling protocol. In the [\[I-D.ietf-pce-segment-routing\]](#), an SID represents either a nodal segment representing a path to a node or adjacency segment representing path over a specific adjacency. In this document, we allow SID also can represent a service segment representing a specific treatment or SF.

In a PCEP session, path information is carried in the Explicit Route Object (ERO), which consists of a sequence of subobjects. In this document, a PCE needs to specify EROs containing SID of service segments (or service-SID), and a PCC needs to be capable of processing such ERO sub-objects.

The SR-ERO Subobject defined in the [\[I-D.ietf-pce-segment-routing\]](#) can be used to carry SID of service segment. An SR-ERO containing SID of service segment can be included in the same PCEP messages specified in the [\[I-D.ietf-pce-segment-routing\]](#).

When a PCEP session between a PCC and a PCE is established, the corresponding PCEP operation is same as defined in the [\[I-D.ietf-pce-segment-routing\]](#).

### **4. Object Formats**

In the [\[I-D.ietf-pce-segment-routing\]](#), an SR-TE path is defined as a path that consists of one or more SID(s) where each SID is associated with the identifier that represents the node or adjacency corresponding to the SID. In this document, we allow the SR-TE path to include one or more SID of service segments (called service-SID) that are inserted along with node segments in SR-TE path. A service-segment may also be used to represent a set of SF instances. The service-SID is local to the node where the service resides, thus a combination of node-segment and service-segment are used together.



#### 4.1. The SR-ERO Subobject extension for service segment support

The SR-ERO Subobject is defined in section 5.3.1 of [I-D.ietf-pce-segment-routing] and as an mandatory subobject used to advertise SID and NAI ('Node or Adjacency Identifier') associated with SID. In this document, we extend the existing SR-ERO Subobject as specified in section 5.3.1 of [I-D.ietf-pce-segment-routing] to represent service-SID of the service segment.

The SR-ERO Subobject as described in [I-D.ietf-pce-segment-routing]:

```

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   |  ST  |   Flags   |F|S|C|M|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SID                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                                     NAI (variable)                                     //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

L

The L bit SHOULD NOT be set, so that the subobject represents a strict hop in the explicit route in case of service-segment.

Type

The Type is as per [I-D.ietf-pce-segment-routing].

Length

The Length is as per [I-D.ietf-pce-segment-routing].

ST

The ST (SID Type) field is set to specify service-SID. A new SID-Type values is to be assigned.

Flags

All flags (M, C, S, F bit) are as per [I-D.ietf-pce-segment-routing].



## SID

The SID value represents an service segment as described in [[I-D.filsfils-spring-segment-routing-use-cases](#)].

## NAI

The NAI for service-segment may be defined in future based on the service.

### **4.2. Service Segment SR-ERO Processing**

When the SID represents a service segment (as per the SID Type - ST field), its value is local to node segment offering the service. Thus Service-SID MUST be associated with a node-SID preceding it in the SR-ERO. Note that multiple services may be offered by the same node, and in this case node-SID maybe followed by multiple Service-SID. NAI value for service-SID may be defined in future.

If a service segment (or service-SID) cannot be associated with a node segment (or node-SID), PCEP speaker MUST send a PCE error with Error-Type = "Reception of an invalid object" and Error-Value = "Segment List Order Error".

The rest of the processing rules are as per [[I-D.ietf-pce-segment-routing](#)].

### **5. Backward Compatibility**

Backward Compatibility consideration described in section 8 of [[I-D.ietf-pce-segment-routing](#)] can be applied for service segment support as well.

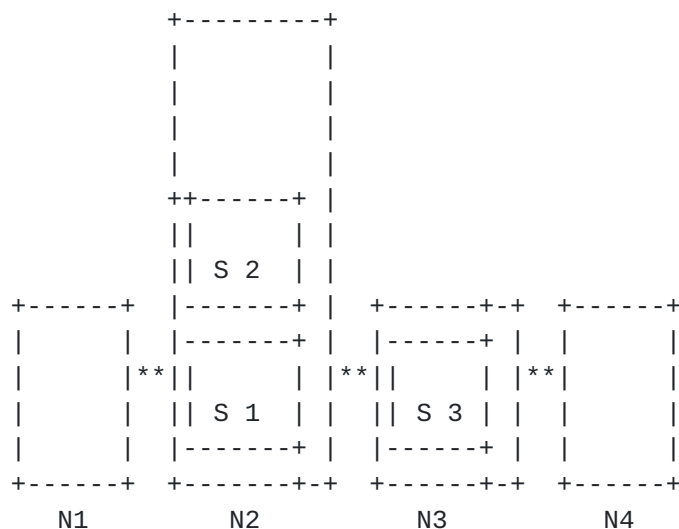
### **6. Management Considerations**

Management consideration described in section 9 of [[I-D.ietf-pce-segment-routing](#)] can be applied to service segment support as well.

### **7. Security Considerations**

The security considerations described in [[RFC5440](#)] and [[I-D.ietf-pce-segment-routing](#)] apply.







- o N1 is Ingress;
- o N4 is Egress;
- o N2 has two services hosted identified as S1 and S2;
- o N3 has one service hosted identified as S3.

The SR-ERO for the SR-TE path including the service segment would be  
-

[{SID\_N2, SID\_S1, SID\_S2}, {SID\_N3, SID\_S3}, {SID\_N4}]

#### Authors' Addresses

Dhruv Dhody  
Huawei  
Divyashree Techno Park, Whitefield  
Bangalore, Karnataka 560037  
India

Email: dhruv.ietf@gmail.com

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

Email: bill.wu@huawei.com

