

SPRING Working Group
Internet-Draft
Intended Status: Informational

Madhukar Anand
Sanjoy Bardhan
Ramesh Subrahmaniam
Infinera Corporation

Jeff Tantsura
Individual
July 6, 2016

Expires: January 7, 2017

Packet-Optical Integration in Segment Routing
draft-anand-spring-poi-sr-01

Abstract

This document illustrates a way to integrate a new class of nodes and links in segment routing to represent transport networks in an opaque way into the segment routing domain. An instance of this class would be optical networks that are typically transport centric. In the IP centric network, this will help in defining a common control protocol for packet optical integration that will include optical paths as 'transport segments' or sub-paths as an augmentation to the defined extensions of segment routing. The transport segment option also defines a general mechanism to allow for future extensibility of segment routing into non-packet domains.

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 RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2016 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	3
2.	Reference Taxonomy	3
3.	Use case - Packet Optical Integration	3
4.	Mechanism overview	5
5.	PCEP-LS extensions for supporting the transport segment	6
6.	OSPF extensions for supporting the transport segment	7
7.	OSPFv3 extensions for supporting the transport segment	9
8.	IS-IS extensions for supporting the transport segment	10
9.	BGP-LS extensions for supporting the transport segment	12
10.	Summary	14
11.	Security Considerations	14
12	IANA Considerations	15
12.1	PCEP	15
12.2	OSPF	15
12.3	OSPFv3	15
12.4	IS-IS	15
12.5	BGP-LS	16
13	References	16
13.1	Normative References	16
13.2	Informative References	17
	Authors' Addresses	17

1 Introduction

Packet and optical transport networks have evolved independently with different control plane mechanisms that have to be provisioned and maintained separately. Consequently, coordinating packet and optical networks for delivering services such as end-to-end traffic engineering or failure response has proved challenging. To address this challenge, a unified control and management paradigm that provides an incremental path to complete packet-optical integration while leveraging existing signaling and routing protocols in either domains is needed. This document introduces such a paradigm based on Segment Routing (SR) [I-D.ietf-spring-segment-routing].

This document introduces a new type of segment, Transport segment. Transport segment can be used to model abstracted paths through the optical transport domain and integrate it with the packet network for delivering end-to-end services. In addition, this also introduces a notion of a Packet optical gateway (POG). These are nodes in the network that map packet services to the optical domain that originate and terminate these transport segments. Given a transport segment, a POG will expand it to a path in the optical transport network.

2. Reference Taxonomy

POG - Packet optical gateway Device

SR Edge Router - The Edge Router which is the ingress device

CE - Customer Edge Device that is outside of the SR domain

PCE - Path Computation Engine

Controller - A network controller

3. Use case - Packet Optical Integration

Many operators build and operate their networks that are both multi-layer and multi-domain. Services are built around these layers and domains to provide end-to-end services. Due to the nature of the different domains, such as packet and optical, the management and service creation has always been problematic and time consuming. With segment routing, enabling a head-end node to select a path and embed the information in the packet is a powerful construct that would be used in the Packet Optical Gateways (POG). The path is usually

constructed for each domain that may be manually derived or through a stateful PCE which is run specifically in that domain.

P1-----O1-----P2-----O2-----P3-----O3-----P4

Figure 1: Representation of a packet-optical path

In Figure 1 above, the nodes represent a packet optical network. P1, P2, P3 and P4 are packet optical devices that are connected via optical paths O1, O2 and O3. Nodes P1 and P4 are edge devices that have customer facing devices (denoted as Border POGs) and P2 and P3 are core nodes (denoted as Transit POGs) in the network. A packet service is established by specifying a path between P1 and P4. Note that in defining this path, we will need to specify both the nodes and the links that make up this service. POGs advertise themselves along with their adjacencies and the domains they belong to. To leverage segment routing to define the above service, the ingress node P1 would append all outgoing packets in a SR header consisting of the SIDs that constitute the path. In the packet domain this would mean P1 would send its packets towards P4 using the segment list {P2, P4}. The operator would need to use a different mechanism in the optical domain to set up the optical paths denoted by O1, O2 and O3. Each POG would announce the active optical path as a transport segment - for example, in the case of P1, the optical path O1 would represent an optical path that includes the optical nodes Om and On as shown on Figure 2. This path is not known to the packet SR domain and is only relevant to the optical domain D between P1 and P2. A PCE that is run in Domain D would be responsible for calculating path O1.

```

          |-----Om-----On-----|
P1----|               (D)               |-----P2
          |-----Ox-----Oy-----|

```

Figure 2: POG with multiple optical paths through an optical domain

Similarly, the transit POGs P2 and P3 in Figure 1 would announce transport segments O2 and O3. The border POG would include the optical paths O1, O2 and O3 to the segment list for P1 to P4. The expanded segment list would read as {O1, P2, O2, P3, O3, P4}.

There are potentially two locations for Borders POGs - one that has last-mile access nodes and the other being Data Center Interconnect nodes. The POGs that are in the core of the network which connect with long haul optical networks are usually Transit POGs.

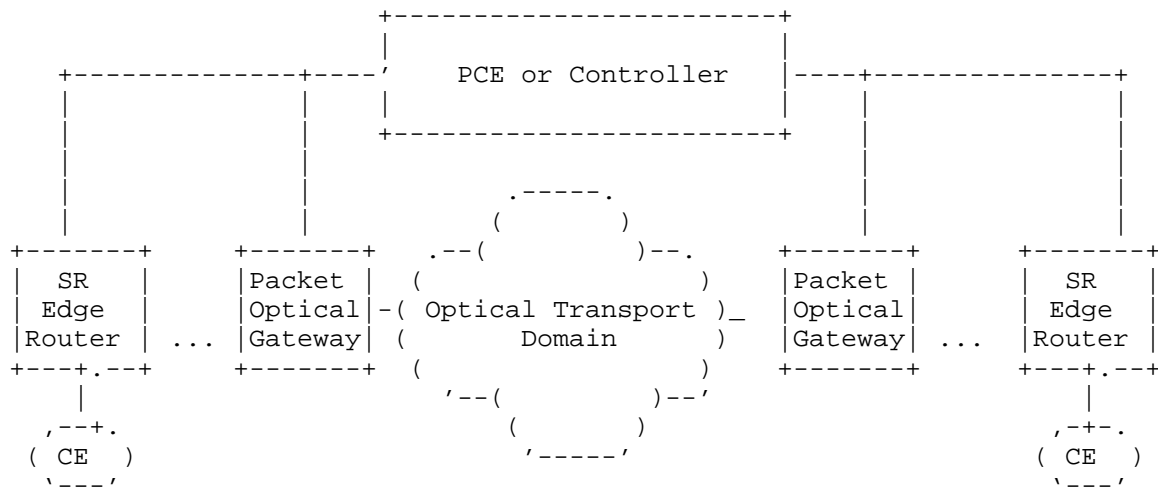


Figure 3. Reference Topology for Transport Segment

4. Mechanism overview

The current proposal assumes that the SR domains run standard IGP protocols to discover the topology and distribute labels without any modification. There are also no modifications to the control plane mechanisms in the Optical transport domains. The mechanism for supporting the transport segment is as follows.

1. Firstly, the Packet Optical Gateway (POG) devices announce themselves in the SR domain. This is indicated by advertising a new SR node capability flag. The exact extensions to support this capability are described in the subsequent sections of this document.

2. Then, the POG devices announce paths to other POGs through the optical transport domain as a transport segment (transport segment binding SID) in the SR domain. The paths are announced with an appropriate optical transport domain ID, and a label (Packet-Optical Label) to be used to bind to the transport segment. The appropriate

IGP segment routing extensions to carry this information is described in the subsequent sections of this document.

3. The transport segment can also optionally be announced with a set of attributes that characterizes the path in the optical transport domain between the two POG devices. For instance, those attributes could define the OTN mapping used (e.g., ODU4, ODU3, ODU3e1...ODU1), timeslots (1-8 or 4,6,7 or 1-2,5), or optical path protection schemes.

4. The POG device is also responsible for programming its forwarding table to map every transport segment label entry into an appropriate forwarding action relevant in the optical domain, such as mapping it to a label-switched path.

5. The transport segment is communicated to the PCE or Controller using extensions to BGP-LS or PCEP-LS as described in subsequent sections of this document.


6. Finally, the PCE or Controller then uses the transport segment label to influence the path leaving the SR domain into the optical domain, thereby defining the end-to-end path for a given service.

5. PCEP-LS extensions for supporting the transport segment

To communicate the Packet-Optical Gateway capability of the device, we introduce a new PCEP capabilities TLV is defined as follows (extensions to [I-D.draft-sivabalan-pce-segment-routing]):

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

A new type of TLV to accommodate a transport segment is defined by extending Binding SIDs [I-D.draft-sivabalan-pce-binding-label-sid-01]



where:

Type: TBD, suggested value 32

Length: variable.

Binding Type: 0 or 1 as defined in
[I-D.draft-sivabalan-pce-binding-label-sid-01]

Domain ID: An identifier for the transport domain

Binding Value: is the transport segment label

Transport Segment Sub TLVs: TBD

IANA will be requested to allocate a new TLV type (recommended value is 32) for TRANSPORT-SEGMENT-BINDING-TLV as specified in this document:

- 1 Transport Segment Label (This document)

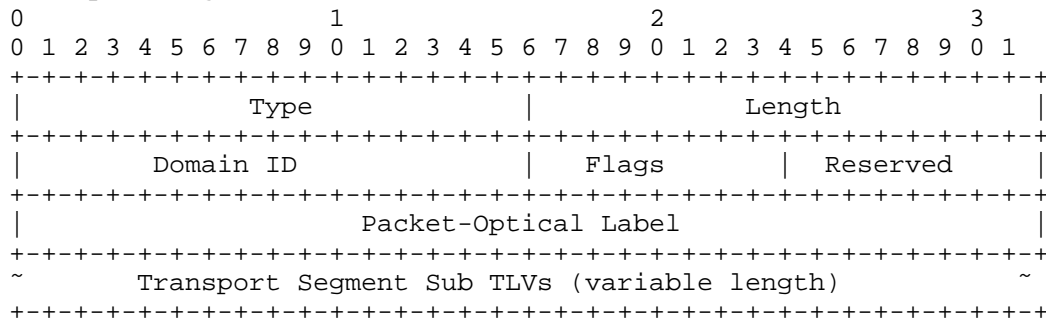
6. OSPF extensions for supporting the transport segment

To communicate the Packet-Optical Gateway capability of the device, we introduce an new optical informational capability bit in the Router Information capabilities TLV (as defined in [RFC4970]).

Bit-24 - Optical - If set, then the router is capable of performing Packet Optical Gateway function.

Further, a new OSPF sub-TLV (similar to the ERO SubTLV) of SID/Label Binding Sub-TLV (TRANSPORT-SEGMENT-BINDING-SUBTLV) to carry the

transport segment label is defined as follows.



where:


Type : TBD, Suggested Value 9

Length: variable.

Domain ID: An identifier for the transport domain

Flags: 1 octet field of following flags:

- V - Value flag. If set, then the optical label carries a value. By default the flag is SET.
- L - Local. Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.



Packet-Optical Label : according to the V and L flags, it contains either:

- * A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- * A 4 octet index defining the offset in the label space advertised by this router. In this case V and L flags MUST be unset.

Transport Segment Sub TLVs: TBD

Multiple TRANSPORT-SEGMENT-BINDING-SUBTLV MAY be associated with a pair

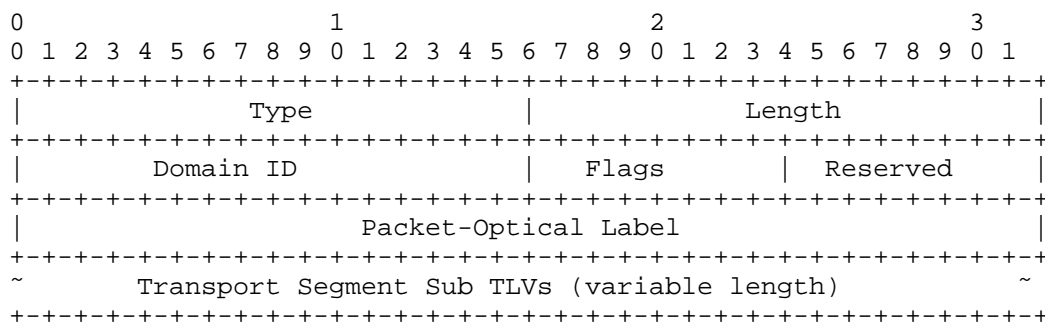
of POG devices to represent multiple paths within the optical domain

7. OSPFv3 extensions for supporting the transport segment

To communicate the Packet-Optical Gateway capability of the device, we introduce an new optical informational capability bit in the Router Information capabilities TLV (as defined in [RFC4970]).

Bit-24 - Optical - If set, then the router is capable of performing Packet Optical Gateway function.

Further, a new OSPFv3 sub-TLV similar to the ERO SubTLV) of SID/Label Binding Sub-TLV (TRANSPORT-SEGMENT-BINDING-SUBTLV) to carry the transport segment label is defined as follows.



where:

Type : TBD, Suggested Value 12


Length: variable.

Domain ID: An identifier for the transport domain

Flags: 1 octet field of following flags:

V - Value flag. If set, then the optical label carries a value.
By default the flag is SET.

L - Local. Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.



Packet-Optical Label : according to the V and L flags, it contains either:

- * A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- * A 4 octet index defining the offset in the label space advertised by this router. In this case V and L flags MUST be unset.

Transport Segment Sub TLVs: TBD

Multiple TRANSPORT-SEGMENT-BINDING-SUBTLV MAY be associated with a pair of POG devices to represent multiple paths within the optical domain

8. IS-IS extensions for supporting the transport segment

To communicate the Packet-Optical Gateway capability of the device, we introduce a new flag O in the SR Node Capabilities sub-TLV:

```

 0 1 2 3 4 5 6 7
+-----+
|I|V|H|O|   |
+-----+
```

I, V, H flags are defined in [I-D.ietf-isis-segment-routing-extensions]

O-Flag: If set, then the router is capable of performing Packet Optical Gateway function.

Further, a new IS-IS sub-TLV (similar to the ERO SubTLV) of SID/Label Binding Sub-TLV (TRANSPORT-SEGMENT-BINDING-SUBTLV) to carry the transport segment label is defined as follows.

First, we define the O flag in the SID/Label Binding TLV

```

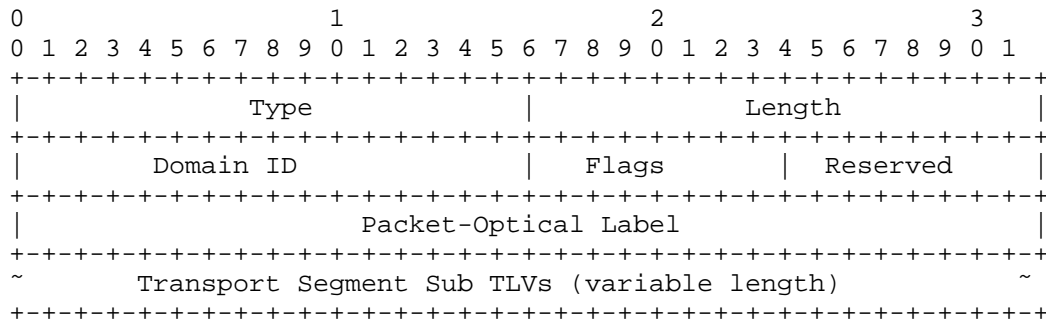
 0 1 2 3 4 5 6 7
+-----+
|F|M|S|D|A|O|   |
+-----+
```

F, M, S, D, and A flags: are defined in [I-D.ietf-isis-segment-routing-extensions]

O-Flag: If set, then the F flag, Range, Prefix Length FEC Prefix, must

be ignored in the SID/Label Binding TLV

Secondly, we define the SubTLV of the SID/Label Binding Sub-TLV:



where:


Type: TBD, Suggested Value 151

Length: variable.

Domain ID: An identifier for the transport domain

Flags: 1 octet field of following flags:

- V - Value flag. If set, then the optical label carries a value. By default the flag is SET.
- L - Local. Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.



Packet-Optical Label : according to the V and L flags, it contains either:

- * A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- * A 4 octet index defining the offset in the label space advertised by this router. In this case V and L flags MUST be unset.

Transport Segment Sub TLVs: TBD

Multiple TRANSPORT-SEGMENT-BINDING-SUBTLV MAY be associated with a pair of POG devices to represent multiple paths within the optical domain with perhaps different characteristics.

9. BGP-LS extensions for supporting the transport segment

9.1 Node Attributes TLV

To communicate the Packet-Optical Gateway capability of the device, we introduce an new optical informational capability the following new Node Attribute TLV is defined:

TLV Code Point	Description	Length	Section
1172	SR-Optical-Node-Capability TLV	variable	

Table 1: Node Attribute TLVs

These TLVs can ONLY be added to the Node Attribute associated with the node NLRI that originates the corresponding SR TLV.

9.2 SR-Optical-Node-Capability TLV

The SR Capabilities sub-TLV has following format:

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		Length	
Flags		RESERVED	

where:

Type : TBD, Suggested Value 1157

Length: variable.

Flags: The Flags field currently has only one bit defined. If the bit is set it has the capability of an Packet Optical Gateway.

9.3 Prefix Attribute TLVs

The following Prefix Attribute Binding SID Sub-TLVs have been added:

TLV Code Point	Description	Length	Section
1173	TRANSPORT-SEGMENT-SID	12	

Table 4: Prefix Attribute - Binding SID Sub-TLVs

The Transport segment TLV allows a node to advertise an transport segment within a single IGP domain. The transport segment SID TLV TRANSPORT-SEGMENT-TLV has the following format:

9.3.1 Transport Segment SID Sub-TLV

Further, a new sub-TLV (similar to the IPV4 ERO SubTLV) of Binding SID Sub-TLV (TRANSPORT-SEGMENT-BINDING-SUBTLV) to carry the transport segment label is defined as follows.

```

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                      |          Length                  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Domain ID                 |          Flags                   |          Reserved                |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Packet-Optical Label      |
+-----+-----+-----+-----+-----+-----+-----+-----+
~          Transport Segment Sub TLVs (variable length)          ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

Type : TBD

Length: variable.

Domain ID: An identifier for the transport domain

Flags: 1 octet field of following flags:

- V - Value flag. If set, then the optical label carries a value.
By default the flag is SET.
- L - Local. Local Flag. If set, then the value/index carried by
the Adj-SID has local significance. By default the flag is SET.

```

0 1 2 3 4 5 6 7
+---+---+---+---+---+
|V|L|
+---+---+---+---+---+

```

Packet-Optical Label : according to the V and L flags, it contains either:

- * A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- * A 4 octet index defining the offset in the label space advertised by this router. In this case V and L flags MUST be unset.

Transport Segment Sub TLVs: TBD

Multiple TRANSPORT-SEGMENT-TLV MAY be associated with a pair of POG devices to represent multiple paths within the optical domain

10. Summary

The motivation for introducing a new type of segment - transport segment - is to integrate transport networks with the segment routing domain and expose characteristics of the transport domain into the packet domain. An end-to-end path across packet and transport domains can then be specified by attaching appropriate SIDs to the packet. An instance of transport segments has been defined here for optical networks, where paths between packet-optical gateway devices has been abstracted using binding SIDs. Extensions to various protocols to announce the transport segment have been proposed in this document.

11. Security Considerations

This document does not introduce any new security considerations.

12 IANA Considerations

This documents request allocation for the following TLVs and subTLVs.

12.1 PCEP

Packet-Optical Gateway capability of the device

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

A new type of TLV to accommodate a transport segment is defined by extending Binding SIDs [I-D.draft-sivabalan-pce-binding-label-sid-01]

Value	Description	Reference
32	TRANSPORT-SR-PCEP-TLV	This document

This document requests that a registry is created to manage the value of the Binding Type field in the TRANSPORT-SR-PCEP TLV.

Value	Description	Reference
1	Transport Segment Label	This document

12.2 OSPF

Transport-Segment SubTLV of OSPF Extended Prefix LSA

Value	Description	Reference
9	TRANSPORT-SR-OSPF-SUBTLV	This document

12.3 OSPFv3

Transport-Segment SubTLV of OSPFv3 Extend-LSA Sub-TLV registry

Value	Description	Reference
12	TRANSPORT-SR-OSPFv3-SUBTLV	This document

12.4 IS-IS

Transport-Segment SubTLV of Segment Identifier / Label Binding TLV

Value	Description	Reference
151	TRANSPORT-SR-ISIS-SUBTLV	This document

12.5 BGP-LS

Node Attributes TLV:

Value	Description	Reference
1172	TRANSPORT-SR-BGPLS-CAPABILITY	This document

Prefix Attribute Binding SID SubTLV:

Value	Description	Reference
1173	TRANSPORT-SR-BGPLS-TLV	This document

13 References

13.1 Normative References

[I-D.ietf-spring-segment-routing]

Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and r. rjs@rob.sh, "Segment Routing Architecture", draft-ietf-spring-segment-routing-04 (work in progress), July 2015.

[I-D.ietf-isis-segment-routing-extensions]

Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-05 (work in progress), June 2015.

[I-D.ietf-ospf-segment-routing-extensions]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", draft-ietf-ospf-segment-routing-extensions-05 (work in progress), June 2015.

[RFC4915] L. Nguyen, P. Psenak, S. Mirtorabi, P. Pillay-Esnault, and A. Roy, "Multi-Topology (MT) Routing in OSPF.", RFC4915, <<http://tools.ietf.org/html/rfc4915>>.

[I-D.ietf-ospf-ospfv3-segment-routing-extensions]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPFv3 Extensions for Segment Routing", draft-ietf-ospf-ospfv3-segment-routing-extensions-03 (work in progress), June 2015.

[I-D.ietf-idr-ls-distribution]

Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-13 (work in progress), October 2015.

[RFC4970] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, DOI 10.17487/RFC4970, July 2007, <<http://www.rfc-editor.org/info/rfc4970>>.

[I-D.sivabalan-pce-binding-label-sid]

Sivabalan, S., Filsfils, C., Previdi, S., Tantsura, J., Hardwick, J., and M. Nanduri, "Carrying Binding Label/Segment-ID in PCE-based Networks.", draft-sivabalan-pce-binding-label-sid-01 (work in progress), March 2016.

[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-07 (work in progress), March 2016.

13.2 Informative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

Authors' Addresses

Madhukar Anand
Infinera Corporation
169 W Java Dr, Sunnyvale, CA 94089

Email: manand@infinera.com

Sanjoy Bardhan
Infinera Corporation
169 W Java Dr, Sunnyvale, CA 94089

Email: sbardhan@infinera.com

Ramesh Subrahmaniam
Infinera Corporation
169 W Java Dr, Sunnyvale, CA 94089

Email: RSubrahmaniam@@infinera.com

Jeff Tantsura

Email: jefftant.ietf@gmail.com