

SPRING Working Group
Internet-Draft
Intended Status: Standard Track

Madhukar Anand
Ciena Corporation

Sanjoy Bardhan
Infinera Corporation

Ramesh Subrahmaniam
Individual

Jeff Tantsura
Apstra

Utpal Mukhopadhyaya
Equinix Inc

Clarence Filsfils
Cisco Systems, Inc.

Expires: January 30, 2020

July 29, 2019

Packet-Optical Integration in Segment Routing
draft-anand-spring-poi-sr-08

Abstract

This document illustrates a way to integrate a new class of nodes and links in segment routing to represent transport/optical networks in an opaque way into the segment routing domain. An instance of this class would be optical networks that are typically transport centric by having very few devices with the capability to process packets. 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 packet paths. 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) 2019 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	4
2.	Reference Taxonomy	4
3.	Use case - Packet Optical Integration	5
4.	Mechanism overview	8
5.	Transport Segments as SR Policy	9
6.	PCEP extensions for supporting the transport segment	10
7.	BGP-LS extensions for supporting the transport segment	11
7.1	Node Attributes TLV	11
7.2	SR-Optical-Node-Capability Sub-TLV	11
7.3	Prefix Attribute TLVs	12
7.3.1	Transport Segment SID Sub-TLV	12

8. Note about Transport Segments and Scalability	13
9. Summary	14
10. Security Considerations	14
11. IANA Considerations	14
11.1 PCEP	14
11.2 BGP-LS	15
12. Acknowledgements	15
13. References	15
13.1 Normative References	15
13.2 Informative References	16
Authors' Addresses	16

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) [RFC8402].

This document introduces a new type of segment, Transport segment, as a special case of SR traffic engineering (SR-TE) policy (Type 1, Sec 5. [I-D.draft-ietf-spring-segment-routing-policy]). Specifically, the structure of SR-TE policy and constraints associated in the transport/optical network are different from those outlined for the packet networks. Transport segment can be used to model abstracted paths through the transport/optical 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. A POG can be viewed as SR traffic engineering policy headend.

The concept of POG introduced here allows for multiple instantiations of the concept. In one case, the packet device is distinct from the transport/optical device, and the POG is a logical entity that spans these two devices. In this case, the POG functionality is achieved with the help of external coordination between the packet and optical devices. In another case, the packet and optical components are integrated into one physical device, and the co-ordination required for functioning of the POG is performed by this integrated device. It must be noted that in either case, it is the packet/optical data plane that is either disaggregated or integrated. Control of the devices can be logically centralized or distributed in either scenario. The focus of this document is to define the logical functions of a POG without going into the exact instantiations of the concept.

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.




Figure 1: Representation of a packet-optical path

In Figure 1 above, the nodes represent a packet optical network. P1,...,P5 are packet devices. Nodes P2 and P3 are connected via optical network (e.g., DWDM) comprising of nodes O1,...,O6. Nodes P2 and P3 are POGs that communicate with other packet devices and also with the devices in the transport/optical domain. POGs P2 and P3 are connected to optical nodes O2/O3 and O3/O4 respectively via multiple links that are visible to the packet network. In defining a path between nodes P2 and P3, we will need to specify the nodes and the links in both the packet and transport/optical domains.

To leverage segment routing to define a service between P1 and P4, 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 a segment list {P2, P3, P4} or {P2, P5, P3, P4} as the case may be. The operator would need to use a different mechanism in the optical domain to set up the paths between the two POGs P2 and P3. For instance, if the packet is forwarded on the link from P2 towards O1 with the expectation that it would come out on the link O4-P3, it could be routed in the optical network using either path {O1, O2, O3, O4} or {O1, O6, O5, O4}. Currently, this decision is made in the optical domain, and there are no mechanisms in the packet network to influence that. The transport segment mechanism proposed in this draft has been designed with an explicit goal of providing better control of optical path selection to the packet network and applications running on them.

Under the proposed scheme, each POG would announce active optical paths to the other POG as a transport segment - for example, the optical path from P2 to P3 comprising {O1, O2, O3, O4} could be represented as a transport segment label Om and the optical path from P2 to P3 comprising devices {O1, O5, O6, O4} could be represented as a transport segment label On. Both Om and On will be advertised by POG P2 as two optical paths between P2 and P3 with specific properties. The specifics of the optical paths, including specific intermediate devices, need not be exposed to the packet SR domain and are only relevant to the optical domain between P2 and P3. A PCE that is run in the optical domain would be responsible for calculating paths corresponding to label Om and On. The expanded segment list would read as {P2, Om, P3, P4} or {P2, On, P3, P4}. Multiple optical paths between P2 and P3 corresponding to different properties can be exposed as transport segments in the packet domain. For example, some optical paths can be low operational cost paths, some could be low-latency, and some others can be high-bandwidth paths. Transport segments for all these candidate viable alternative paths may be generated statically or dynamically. They may be pre-computed or may be generated on the fly when a customer at node P1 requests a service towards node P4. A discussion on transport segments and scalability can be found in Section 8.

Use-case examples of transport segments.

1. Consider the scenario where there are multiple fibers between two packet end points. The network operator may choose to route packet traffic on the first fiber, and reserve the second fiber only for maintenance or low priority traffic.
2. As a second use-case, consider the case where the packet end points are connected by transport/optical network provided by two different service providers. The packet operator wants to preferentially route traffic over one of the providers and use the second provider as a backup.
3. Finally, let the packet end points be connected by optical paths that may span multiple optical domains i.e. different administrative control. For instance, one transport/optical path may lie completely in one country while the other transport/optical path transits another country. Weather, tariffs, security considerations and other factors may determine how the packet operator wants to route different types of traffic on this network.

All of the above use-cases can be supported by first mapping distinct transport/optical paths to different transport segments and then, depending on the need, affixing appropriate transport segment identifier to the specific packet to route it appropriately through the transport domain.




Figure 3. Reference Topology for Transport Segment Mechanism

4. Mechanism overview

The current proposal assumes that the SR domains run standard protocols without any modification to discover the topology and distribute labels. There are also no modifications necessary in the control plane mechanisms in the transport/optical domains. The only requirement of a transport segment is that the optical path be setup before they are announced to the packet network. For example, the optical paths may be setup using a domain-specific controller or a PCE based on requirements from the packet domain (such as bandwidth, QoS, latency and cost) taking into consideration the constraints in the optical network.

The mechanism for supporting the transport segment is as follows.

1. Firstly, the Packet Optical Gateway (POG) devices are announced in the packet 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 candidate transport/optical paths between that POG (Source POG) and other POGs (Destination POG) via appropriate mechanisms in the packet domain. The paths are announced with an appropriate transport/optical domain ID and a Binding SID representing the transport segment from a source POG to a destination POG. The appropriate protocol-specific extensions to carry path characteristics and Binding SID corresponding to a optical path are described in the subsequent sections of this document.

3. The transport SR policy can also optionally be announced with a set of attributes that characterizes the path in the transport/optical domain between the two POG devices. For instance, those could define the path attributes such as path identifier, latency, bandwidth, quality, directionality, or optical path protection schemes. These attributes can be used to determine the "color" of the SR-TE policy in the tuple <Source POG, Destination POG, color> used to prioritize different candidate paths between the POGs.

4. The POG device is also responsible for programming its forwarding table to map every transport segment Binding SID entry

into an appropriate forwarding action relevant in the optical domain, such as mapping it to a optical label-switched path.

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

6. Finally, the PCE or Controller in the packet domain then uses the transport segment binding SID in the overall SR policy to influence the path traversed by the packet in the optical domain, thereby defining the end-to-end path for a given service.

In the next few sections, we outline a few representative protocol specific extensions to carry the transport segment.

5. Transport Segments as SR Policy

The Segment Routing Traffic Engineering (SRTE) [ietf-spring-segment-routing-policy] process installs the transport segment SR policy in the forwarding plane of the POG. The Transport SR policy is identified by using a transport segment Binding SID. Corresponding to each transport segment Binding SID, the SRTE process MAY learn about multiple candidate paths. The SRTE-DB includes information about the candidate paths including optical domain, topology and path characteristics. All of the information can be learned from different sources including but not limited to: Netconf/Restconf, PCEP and BGP-LS.

The information model for Transport SR policy is as follows:

```
Transport SR Policy F01
  Candidate-paths
    path preference 200 (selected)
      BSID1
    path preference 100
      BSID2
    path preference 100
      BSID3
    path preference 50
      BSID4
```

A transport SR policy is identified through the tuple <Source POG, Destination POG, color>. Each TSR policy is associated with one or more candidate paths, each of them associated with a (locally) unique Binding SID and a path preference. For each transport SR policy, the candidate path with the highest path preference (at most one) is selected and used for forwarding traffic that is being steered onto that policy. When candidate paths change (or a new candidate path is set up), the path

selection process can be re-executed. The validity of each path is to be verified by the POG before announcement in the packet domain. If there are no valid paths, then the transport SR policy is deemed invalid.


The allocation of BSID to a path can include dynamic, explicit or generic allocation strategies as discussed in [ietf-spring-segment-routing-policy]. We discuss PCEP and BGP-LS specific extensions in the subsequent section.

6. PCEP 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.ietf-pce-segment-routing]):

Value	Meaning	Reference
TBD1	TRANSPORT-SR-PCE-CAPABILITY	This document

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



where:

Type: TBD

Length: variable.

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

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 for
TRANSPORT-SEGMENT-BINDING-TLV as specified in this document:

TBD Transport Segment Label (This document)

7. BGP-LS extensions for supporting the transport segment

7.1 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
TBD	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.

7.2 SR-Optical-Node-Capability Sub-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								
+-----+																																							

```

+---+---+---+---+---+---+---+---+---+---+

```

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.

7.3 Prefix Attribute TLVs

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

TLV Code Point	Description	Length	Section
TBD	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:

7.3.1 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                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          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

8. Note about Transport Segments and Scalability

In most operational scenarios, there would be multiple, distinct paths between the POGs. There is no requirement that every distinct path in the optical domain be advertised as a separate transport segment. Transport segments are designed to be consumed in the packet domain, and the correspondence between transport segments and exact paths in the optical domain are determined by their utility to the packet world. Therefore, the number of transport segments is to be determined by the individual packet-optical use-case. The number of actual paths in the

optical domain between the POG is expected to be large (counting the number of active and passive devices in the optical network), it is likely that multiple actual paths are to be advertised as one transport segment. Of course, in the degenerate case, it is possible that there is a one-to-one correspondence between an optical path and a transport segment. Given this view of network operation, the POG is not expected to handle a large number of transport segments (and identifiers). This framework does leave open the possibility of handling a large number of transport segments in future. For instance, a hierarchical partitioning of the optical domain along with stacking of multiple transport segment identifiers could be explored towards reducing the overall number of transport segment identifiers.

9. Summary

The motivation for introducing a new type of segment - transport segment - is to integrate transport/optical networks with the segment routing domain and expose characteristics of the transport/optical domain into the packet domain. An end-to-end path across packet and transport/optical 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 have been abstracted using binding SIDs. Extensions to various protocols to announce the transport segment have been proposed in this document.

10. Security Considerations

This document does not introduce any new security considerations.

11 IANA Considerations

This documents request allocation for the following TLVs and subTLVs.

11.1 PCEP

Packet-Optical Gateway capability of the device

Value	Meaning	Reference
TBD1	TRANSPORT-SR-PCE-CAPABILITY	This document

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

Value	Description	Reference
-------	-------------	-----------

TBD2 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
TBD3	Transport Segment Label	This document

11.2 BGP-LS

Node Attributes TLV:

Value	Description	Reference
TBD4	TRANSPORT-SR-BGPLS-CAPABILITY	This document

Prefix Attribute Binding SID SubTLV:

Value	Description	Reference
TBD5	TRANSPORT-SR-BGPLS-TLV	This document

12 Acknowledgements

We would like to thank Peter Psenak, Bruno Decraene, Ketan Talaulikar and Radhakrishna Valiveti for their comments and review of this document.

13 References

13.1 Normative References

[RFC8402]

Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, July 2018.

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

Sivabalan, S., Tantsura, J., Filsfils, C., Previdi, S., Hardwick, J., and Dhody, D., "Carrying Binding Label/Segment-ID in PCE-based Networks.", draft-sivabalan-pce-binding-label-sid-07 (work in progress), July 2019.

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

Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W.,

and J. Hardwick, "PCEP Extensions for Segment Routing",
draft-ietf-pce-segment-routing-16 (work in progress),
Mar 2019.

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

Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A.,
and Mattes, P., "Segment Routing Policy Architecture",
draft-ietf-spring-segment-routing-policy-03.txt
(work in progress), May 2019.

13.2 Informative References

- [RFC5513] Farrel, A., "IANA Considerations for Three Letter
Acronyms", RFC 5513, DOI 10.17487/RFC5513, April 1 2009,
<<http://www.rfc-editor.org/info/rfc5513>>.
- [RFC5514] Vyncke, E., "IPv6 over Social Networks",
RFC 5514, DOI 10.17487/RFC5514, April 1 2009,
<<http://www.rfc-editor.org/info/rfc5514>>.
- [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
Ciena Corporation
3939, N 1st Street, San Jose, CA, 95134
Email: madanand@ciena.com

Sanjoy Bardhan
Infinera Corporation
169 W Java Dr, Sunnyvale, CA 94089
Email: sbardhan@infinera.com

Ramesh Subrahmaniam

Email: svr_fremont@yahoo.com

Jeff Tantsura
Apstra
333 Middlefield Road Suite 200
Menlo Park, CA 94025
Email: jefftant.ietf@gmail.com

Utpal Mukhopadhyaya
Equinix Inc
1188 E. Arques, Sunnyvale, CA 94085
Email: umukhopadhyaya@equinix.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
BE
Email: cfilsfil@cisco.com