

TEAS Working Group  
Internet Draft  
Intended status: Informational  
Expires: April 20, 2019

Y. Lee  
Q. Wu  
I. Busi  
Huawei

D. Cekarreli  
Ericsson

J. Tantsura  
Apstra

October 19, 2018

## Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to VPN with the Integration of Packet and Optical Networks

[draft-lee-teas-actn-vpn-poi-00](#)

### Abstract

This document outlines the applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to VPN with the integration of Packet and Optical Networks (POI). It also identifies a number of scenarios where the integration of packet and optical networks is necessary to support VPN service requirements. The role of optical underlay tunnels in the POI is to support certain applications that require a hard isolation with strict deterministic latency and guaranteed constant bandwidth.

### 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."

---

Internet-Draft      ACTN Applicability to VPN with POI

October 2018

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

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

This Internet-Draft will expire on April 20, 2019.

## Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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. Introduction.....</a>	<a href="#">3</a>
<a href="#">1.1. Requirements Language.....</a>	<a href="#">3</a>
<a href="#">2. Background and Scope.....</a>	<a href="#">4</a>
<a href="#">3. POI with L2/L3VPN Service Under Single Network Operator Control .....</a>	<a href="#">5</a>
<a href="#">3.1. POI with single packet and single optical domain.....</a>	<a href="#">5</a>
3.2. POI with multiple packet domains and single optical domain	8
3.3. POI with multiple packet domains and multiple optical domains.....	<a href="#">10</a>
<a href="#">3.4. Transport of Tunnel and VPN information.....</a>	<a href="#">12</a>
<a href="#">3.5. Virtual Switching Instance (VSI) Provisioning for L2VPN..</a>	<a href="#">13</a>
<a href="#">3.6. Inter-domain Links Update.....</a>	<a href="#">13</a>
<a href="#">3.7. End-to-end Tunnel Management.....</a>	<a href="#">13</a>
<a href="#">4. POI with VN Recursion Under Multiple Network Operators Control .....</a>	<a href="#">14</a>
<a href="#">4.1. Service Request Process between Multiple Operators.....</a>	<a href="#">15</a>
<a href="#">4.2. Service/Network Orchestration of Operator 2.....</a>	<a href="#">16</a>
<a href="#">5. Security Considerations.....</a>	<a href="#">16</a>
<a href="#">6. IANA Considerations.....</a>	<a href="#">17</a>

[7.1. Normative References.....17](#)  
[7.2. Informative References.....17](#)  
[8. Contributors.....18](#)  
[Authors' Addresses.....18](#)

## [1. Introduction](#)

Abstraction and Control of Traffic Engineered Networks (ACTN) describes a set of management and control functions used to operate one or more TE networks to construct virtual networks that can be represented to customers and that are built from abstractions of the underlying TE networks so that, for example, a link in the customer's network is constructed from a path or collection of paths in the underlying networks [[RFC8453](#)].

This document outlines the applicability of ACTN to VPN with the integration of packet and optical networks which is known as the Packet and Optical Integration (POI).

It also identifies a number of scenarios where the integration of packet and optical networks is necessary to support VPN service requirements. The role of optical underlay tunnels in the POI is to support certain applications that require a hard isolation with strict deterministic latency and guaranteed constant bandwidth.

Note that there may be other transport technologies that can support the aforementioned service requirements such as TSN or Detnet to name a few. In this particular document, we are focusing on the currently available network settings where packet networks are a client layer to optical transport networks as a server layer. The principle discussed in this document can be applied to other transport technologies when they are available.

As ACTN [[RFC8453](#)] introduces the role of controllers that facilitate network operations, the scope of this document is how controllers can facilitate L2/3VPN service provisioning in the packet and optical transport networks.

### [1.1. Requirements Language](#)

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED",

"MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

## [2](#). Background and Scope

One of the important functions the MDSC performs is to identify which TE Tunnels should carry the L3VPN traffic and to relay this information to the domain-level controllers to ensure proper Virtual routing and forwarding (VRF) table be populated according to the TE binding requirement for the L3VPN. This function is referred to as TE & service mapping function. The YANG model to provide TE & service mapping function is provided in [[TSM](#)]. The role of the TE-service Mapping model [[TSM](#)] is to expose the mapping relationship between service models and TE models so that VN/VPN service instantiations provided by the underlying TE networks can be viewed outside of the MDSC.

The TE-Service Mapping model also provides service-TE binding information for each service instance so that proper TE tunnel should be created.

The TE binding requirement types defined in [[TSM](#)] are:

- a) New VN/Tunnel Binding - A customer could request a VPN service based on VN/Tunnels that are not shared with other existing or future services. This might be to meet VPN isolation requirements.

Under this mode, the following sub-categories can be supported:

- i. Hard Isolation with deterministic characteristics: A customer could request a VPN service using a set of TE Tunnels with deterministic characteristics requirements (e.g., no latency variation) and where that set of TE Tunnels must not be shared with other VPN services and must not compete for bandwidth or other network resources with other TE Tunnels.
- ii. Hard Isolation: This is similar to the above case but without the deterministic characteristics requirements.

- iii. Soft Isolation: The customer requests a VPN service using a set of TE tunnels which can be shared with other VPN services.
- b) VN/Tunnel Sharing - A customer could request a VPN service where new tunnels (or a VN) do not need to be created for each VPN and can be shared across multiple VPNs.

- c) VN/Tunnel Modify - This mode allows the modification of the properties of the existing VN/tunnel (e.g., bandwidth).

This document addresses cases a)-i (hard isolation with deterministic latency) and a)-ii (hard isolation with non-deterministic latency). Both cases warrant consideration of optical undelay bypass tunnels to meet the service requirement.

The optical bypass tunnel could be setup via RSVP-TE signaling and thus tunnel label allocation could be done during signaling. It is also possible that PNC and MDSC coordinates to exchange the TE tunnel label information to setup this optical bypass tunnel. This document focuses on the latter case.

The multi-hop e-BGP session between ingress and egress for multi-domain case would be setup to exchange VPN routes. The rest of the forwarding action is as per the usual BGP L3VPN handling including the use of TE tunnel.

### [3.](#) POI with L2/L3VPN Service Under Single Network Operator Control

This section provides a set of specific deployment scenarios for POI under single network operator control. Specifically, the following deployment scenarios are discussed in this section:

- One optical transport domain overarched by one packet domain (see [Section 3.1](#));
- One optical transport domain overarched by multiple packet domains (see [Section 3.2](#));
- multiple optical transport domains overarched by multiple packet domains (see [Section 3.3](#)).

All scenarios are taking place in the context of an upper layer service configuration (e.g., L3VPN) in the packet and optical transport network.

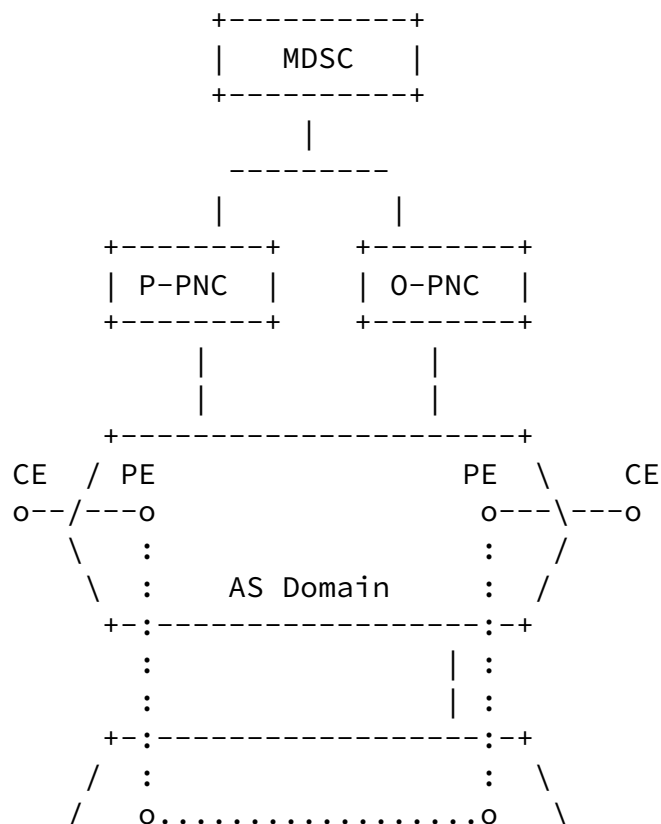
Since this document only addresses the procedure for creating optical underlay bypass tunnels, it does not affect MP-BGP MPLS operations for inter-AS scenarios as specified in [RFC4364].

### 3.1. POI with single packet and single optical domain

This section provides a specific deployment scenario for POI. Specifically, it provides a deployment scenario in which hierarchical controllers (an MDSC and two PNCs, one for packet and

one for optical) facilitate optical bypass tunnel across the packet domain and the optical domain.

Figure 1 shows this scenario.






Figure 1. One Packet Domain and One Optical Domain

The following control sequence describes the scenario depicted in Figure 1.

- a) The MDSC translates the service instance and its requirement (hard isolation with deterministic latency).
- b) The MDSC computes the path if there is any feasible path to meet the requirement based on the abstracted topology at hand. Note that there would not be any tunnel in the packet domain to meet this requirement (hard isolation with deterministic latency).
- c) The MDSC finds a feasible path in the optical domain.
- d) The MDSC asks the optical PNC to create a tunnel for this VPN instance whose endpoints are the ingress PE and the egress PE

of the packet domain, respectively. The MDSC and Optical PNC need to maintain an instance ID for this VPN instance.

- e) The MDSC asks the Packet PNC to bind a TE-tunnel label (to be allocated by the egress PE to identify the underlay optical tunnel) with the VPN ID and the Ingress and Egress interfaces of the underlay optical tunnel.
- f) The PNC in turn asks the Egress PE to allocate a TE-tunnel label. The Egress PE allocates a TE-tunnel label, populates the VRF for this VPN instance, and updates the Packet PNC with the allocated TE-tunnel label. Please refer to the note below on the details of this procedure in regard to VPN binding.

Note: There are two cases for binding network instance with the TE tunnel label:

1. VRF instance does not exist.
2. VRF instance has already been created.

For case 1, the Egress PE needs to bind the TE-tunnel label and the VPN information (e.g., VPN instance name, VPN label, RD, RT, Destination IP address, etc.) and inform this binding information to the packet PNC.

- g) The packet PNC informs the MDSC the allocated TE-tunnel label for the VPN instance.
- h) The MDSC informs the optical PNC to bind the TE-tunnel label with the VPN instance, which has been created previously in step d).
- i) The optical PNC informs this binding information (i.e., ingress/egress interfaces from packet domain and the TE-tunnel label) to the optical ingress switch.
- j) The packet PNC informs the ingress PE to use the TE-label for this VPN instance. The Ingress PE populates the VRF for the VPN with the TE-label. (Note that the TE-label would need to be PUSHed over the VPN traffic).
- k) When the packet arrives at the ingress PE, it recognizes the VPN instance and PUSHes the VPN label and the TE-tunnel label and forward the traffic to optical ingress switch.
- l) The optical ingress switch recognizes the TE-tunnel label and encapsulate the whole data packet including TE-tunnel label into the OTN payload.
- m) The optical egress switch POPs the ODU label and forwards the data packet to the packet egress PE.
- n) The packet egress PE POPs the TE-tunnel label and forwards the VPN packets to the destination CE.

Note: in steps k) - l), the assumption made was that the packet ingress PE is not OTN-capable router. If the packet ingress PE support channelized OTN interfaces, the data plane behavior in steps k) and l) would change as the following:

k') When the packets arrives at the ingress PE, it recognizes the VPN instance and PUSHes the VPN label and the TE-tunnel label and the ODU label and forward the traffic to optical ingress switch.

l') The optical ingress switch recognizes the incoming ODU label and swap it to outgoing ODU label.

### [3.2.](#) POI with multiple packet domains and single optical domain

This section provides a specific deployment scenario for POI. Specifically, it provides a deployment scenario in which hierarchical controllers (an MDSC and two packet PNCs and one



optical PNC) facilitate optical bypass tunnel across the two packet domains and the optical domain.

Figure 2 shows this scenario.

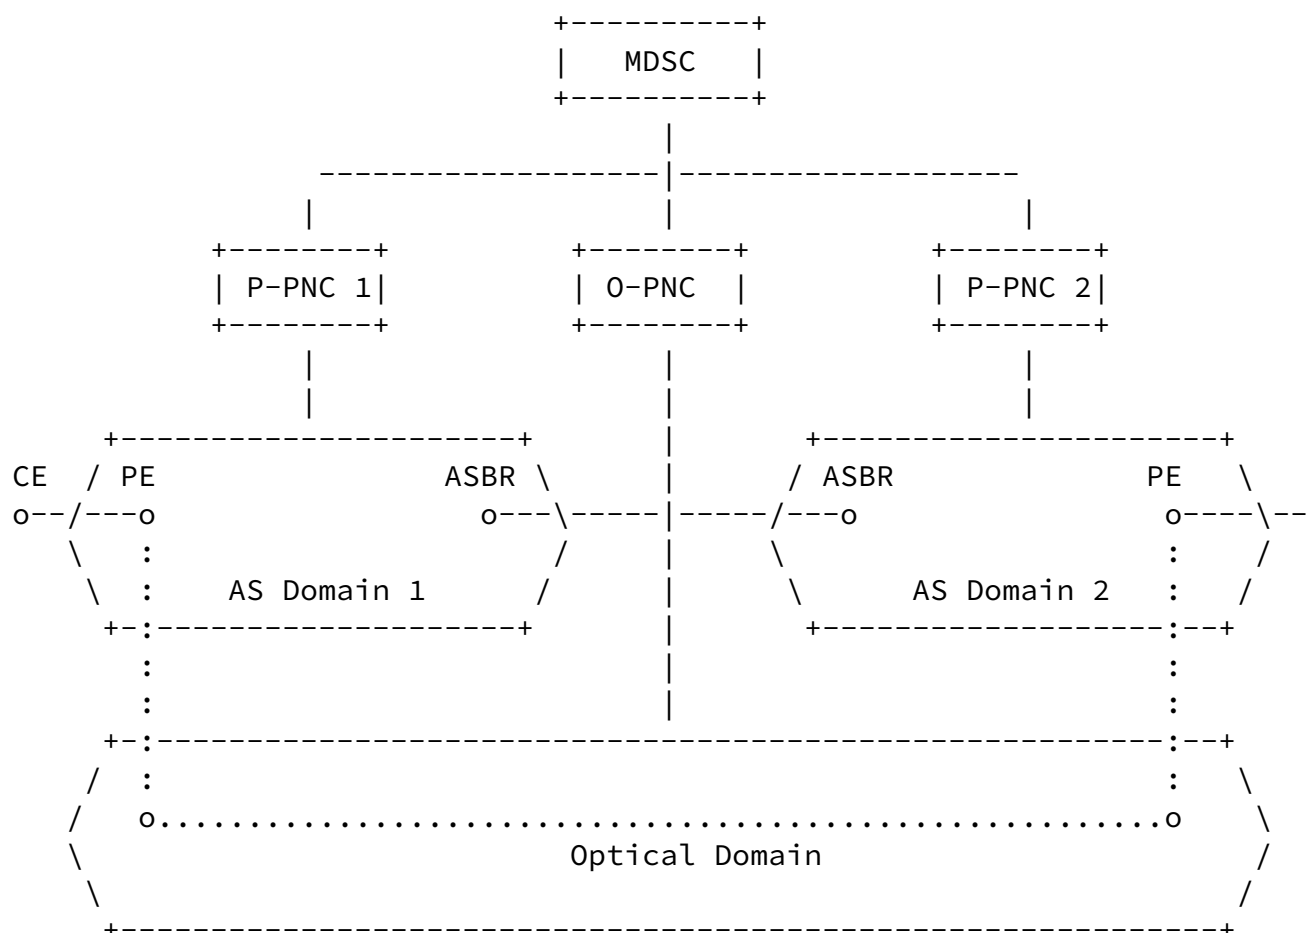


Figure 2. Two Packet Domains and One Optical Domain

The control sequence depicted in Figure 2 is same as the control sequence a)-d) in [Section 3.1](#) with the following differences:

- e) The MDSC asks the Packet PNC 2 to bind a TE-tunnel label (to be allocated by the egress PE to identify the underlay optical tunnel) with the VPN ID and the Ingress and Egress interfaces of the underlay optical tunnel.
- f) The packet PNC 2 in turn asks the Egress PE to allocate a TE-tunnel label. The Egress PE allocates a TE-tunnel label, populates the VRF for this VPN instance, and updates the

packet PNC 2 with the allocated TE-tunnel label. Please refer to the note below on the details of this procedure in regard to VPN binding.

Note: There are two cases for binding network instance with the TE tunnel label:

1. VRF instance does not exist.
2. VRF instance has already been created.

For case 1, the Egress PE needs to bind the TE-tunnel label and the VPN information (e.g., VPN instance name, VPN label, RD, RT, Destination IP address, etc.) and inform this binding information to the packet PNC 2.

- g) The packet PNC 2 informs the MDSC the allocated TE-tunnel label for the VPN instance.
- h) The MDSC informs the packet PNC 1 the allocated TE-tunnel label for the VPN instance.
- i) The MDSC informs the optical PNC to bind the TE-tunnel label with the VPN instance, which has been created previously in step d).
- j) The optical PNC informs this binding information (i.e., ingress/egress interfaces from packet domain and the TE-tunnel label) to the optical ingress switch.
- k) The packet PNC 1 informs the ingress PE in Domain 1 to use the TE-tunnel label for this VPN instance. The Ingress PE in Domain 2 populates the VRF for the VPN and bind with the TE-tunnel label. (Note that the TE-tunnel label would need to be PUSHed over the VPN traffic).
- l) When the packets arrives at the ingress PE in Domain 1, it recognizes the VPN instance and PUSHes the VPN label and the TE-tunnel label and forward the traffic to optical ingress switch.

- m) The optical ingress switch recognizes the TE-tunnel label and encapsulate the whole data packet including TE-tunnel label into the OTN payload.
- n) The optical egress switch POPs the ODU label and forwards the data packet to the packet egress PE.
- o) The packet egress PE in Domain 2 POPs the TE-tunnel label and forwards the VPN packets to the destination CE.

Note: in steps l) - m), the assumption made was that the packet ingress PE is not OTN-capable router. If the packet ingress PE supports channelized OTN interfaces, the data plane behavior in steps l) and m) would change as the following:

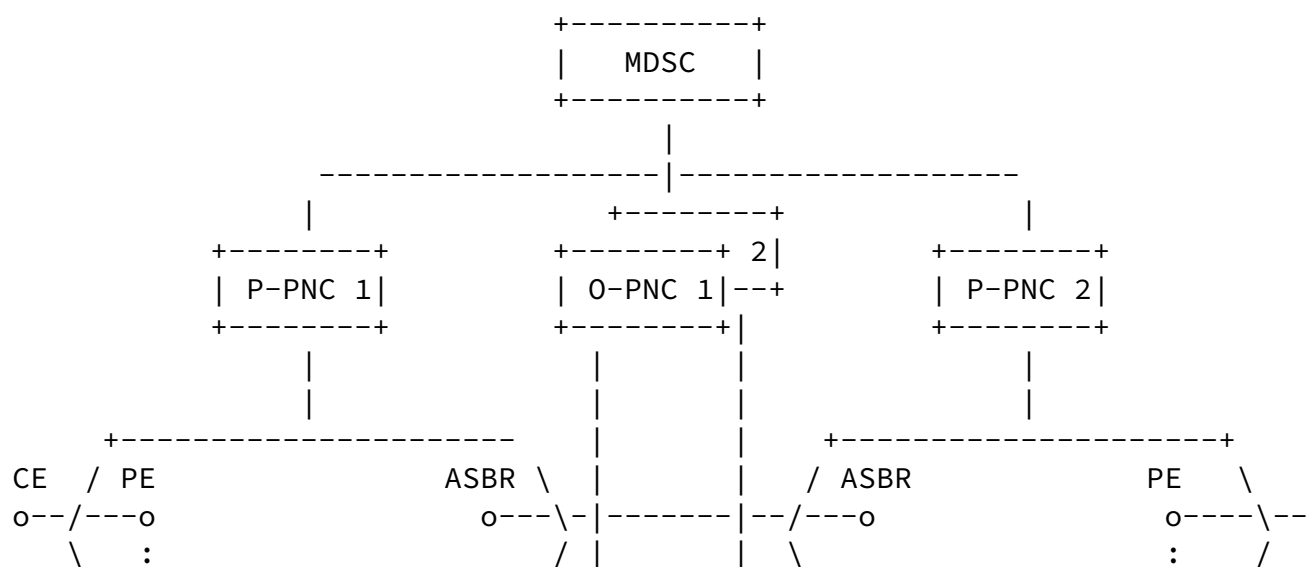
l') When the packets arrives at the ingress PE, it recognizes the VPN instance and PUSHes the VPN label and the TE-tunnel label and the ODU label and forward the traffic to optical ingress switch.

m') The optical ingress switch recognizes the incoming ODU label and swap it to outgoing ODU label.

### 3.3. POI with multiple packet domains and multiple optical domains

This section provides a specific deployment scenario for POI. Specifically, it provides a deployment scenario in which hierarchical controllers (an MDSC and two packet PNCs and two optical PNCs) facilitate optical bypass tunnel across two packet domains and two optical domains.

Figure 3 shows this scenario.






Figure 3. Two Packet Domains and One Optical Domain

The control sequence depicted in Figure 3 is same as the control sequence a)-c) in [Section 3.1](#) with the following differences:

- d) The MDSC asks the optical PNC 1 to create a tunnel for this VPN instance whose endpoints are the ingress PE of the packet domain 1 and the optical inter-domain interface toward optical domain 2; and the optical PNC 2 to create a tunnel for this VPN instance whose endpoints are the optical inter-domain interface from optical domain 1 and the egress PE of the packet domain 2. The MDSC and Optical PNC 1 and PNC 2 need to maintain an instance ID for this VPN instance.
- e) The MDSC asks the Packet PNC 2 to bind a TE-tunnel label with the VPN ID and the Ingress and Egress interfaces of the underlay optical tunnel.
- f) The packet PNC 2 in turn asks the Egress PE to allocate a TE-tunnel label. The Egress PE allocates a TE-tunnel label, populates the VRF for this VPN instance, and updates the packet PNC 2 with the allocated TE-tunnel label. Please refer to the note below on the details of this procedure in regard to VPN binding.

Note: There are two cases for binding network instance with the TE tunnel label:

1. VRF instance does not exist.
2. VRF instance has already been created.

For case 1, the Egress PE needs to bind the TE-tunnel label and the VPN information (e.g., VPN instance name, VPN label, RD, RT, Destination IP address, etc.) and inform this binding information to the packet PNC 2.

- g) The packet PNC 2 informs the MDSC the allocated TE-tunnel label for the VPN instance.

- h) The MDSC informs the packet PNC 1 the allocated TE-tunnel label for the VPN instance.
- i) The MDSC informs the optical PNC 1 and PNC 2 to bind the TE-tunnel label with the instance, which has been created previously in step d).
- j) The optical PNC 1 informs this binding information (i.e., ingress/egress interfaces from packet domain and the TE-tunnel label) to the optical ingress switch in Domain 1. Likewise, the optical PNC 2 to the optical egress switch in Domain 2. (Note we assume that the optical border switches in Domains 1 and 2 would do the normal OTN switching).
- k) The packet PNC 1 informs the ingress PE in Domain 1 to use the TE-tunnel label for this VPN instance. The Ingress PE in Domain 2 populates the VRF for the VPN with the TE-label. (Note that the TE-tunnel label would need to be PUSHed over the VPN traffic).
- l) When the VPN packet arrives at the ingress PE in Domain 1, it recognizes the VPN label and PUSHes the TE-tunnel label and forward the traffic to optical ingress switch in optical domain 1.
- m) The optical ingress switch in optical domain 1 recognizes the TE-tunnel label and encapsulate the whole data packets including TE-tunnel label into the OTN payload.
- n) The optical egress switch in optical domain 2 POPs the OTN label and forwards the data packet to the packet egress PE.
- o) The packet egress PE in Domain 2 POPs the TE-tunnel label and forwards the VPN packet to the destination CE.

Note: in steps l) - m), the assumption made was that the packet ingress PE is not OTN-capable router. If the packet ingress PE supports channelized OTN interfaces, the data plane behavior in steps l) and m) would change as the following:

l') When the packets arrives at the ingress PE, it recognizes the VPN instance and PUSHes the VPN label and the TE-tunnel label and the ODU label and forward the traffic to optical ingress switch in Domain 1.

m') The optical ingress switch in Domain 1 recognizes the incoming ODU label and swap it to outgoing ODU label.

#### 3.4. Transport of Tunnel and VPN information

The discussions in [Section 3](#) as to the transport mechanism of the TE-tunnel label used for the underlay bypass tunnel with the VPN instance information has the undertone of making use of the controllers. Note that other mechanisms may also be possible and

---

Internet-Draft      ACTN Applicability to VPN with POI      October 2018

that such mechanisms are not precluded when solutions are sought out.

### [3.5.](#) Virtual Switching Instance (VSI) Provisioning for L2VPN

The VSI provisioning for L2VPN is similar to the VPN/VRF provision for L3VPN. L2VPN service types include:

- . Point-to-point Virtual Private Wire Services (VPWSs) that use LDP-signaled Pseudowires or L2TP-signaled Pseudowires [[RFC6074](#)];
- . Multipoint Virtual Private LAN Services (VPLSs) that use LDP-signaled Pseudowires or L2TP-signaled Pseudowires [[RFC6074](#)];
- . Multipoint Virtual Private LAN Services (VPLSs) that use a Border Gateway Protocol (BGP) control plane as described in [[RFC4761](#)] and [[RFC6624](#)];
- . IP-Only LAN-Like Services (IPLSs) that are a functional subset of VPLS services [[RFC7436](#)];
- . BGP MPLS-based Ethernet VPN Services as described in [[RFC7432](#)] and [[RFC7209](#)];
- . Ethernet VPN VPWS specified in [[RFC8214](#)] and [[RFC7432](#)].

### [3.6.](#) Inter-domain Links Update

In order to facilitate inter-domain links for the VPN, we assume that the service/network orchestrator would know the inter-domain link status and its resource information (e.g., bandwidth available, protection/restoration policy, etc.) via some mechanisms (which are beyond the scope of this document). We also assume that the inter-domain links are pre-configured prior to service instantiation.

### [3.7.](#) End-to-end Tunnel Management

It is foreseen that the MDSC should control and manage end-to-end tunnels for VPNs per VPN policy.

As discussed in [[ACTN-Telemetry](#)], the MDSC is responsible to


collect domain LSP-level performance monitoring data from domain controllers and to derive and report end-to-end tunnel performance monitoring information to the customer.


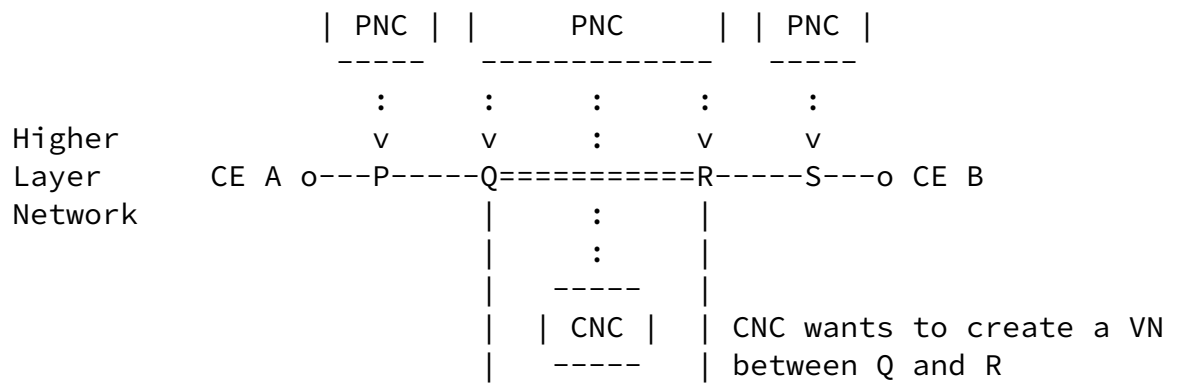
#### 4. POI with VN Recursion Under Multiple Network Operators Control

[RFC8453] briefly introduces a case for the VN supplied to a customer may be built using resources from different technology layers operated by different operators. For example, one operator may run a packet TE network and use optical connectivity provided by another operator.

Figure 4, extracted from [RFC8453], shows the case where a customer asks for end-to-end connectivity between CE A and CE B, a virtual network. The customer's CNC makes a request to Operator 1's MDSC. The MDSC works out which network resources need to be configured and sends instructions to the appropriate PNCs. However, the link between Q and R is a virtual link supplied by Operator 2: Operator 1 is a customer of Operator 2.

To support this, Operator 1 has a CNC that communicates with Operator 2's MDSC. Note that Operator 1's CNC in Figure 10 is a functional component that does not dictate implementation: it may be embedded in a PNC.





Where

--- is a link

=== is a virtual link

Figure 4: VN Recursion with Network Layers

The CMI in Figure 4 interfaces Operator 1's CNC with Operator 2's MDSC. The functions to perform and the information carried over the inter-operator CMI are identical to those of the Customer's CNC and Operator 1's MDSC. In other words, the two CMIs depicted in Figure 4 are recursive in nature.

From a data plane perspective, the interaction between operator 1 and operator 2 is similar to the POI case discussed in [section 3.2](#) (See Figure 2) with an exception that the packet domains belong to operator 1 while optical domain to operator 2.



The control interface depicted in Figure 4 (i.e., the CNC of operator 1 and the MDSC of operator 2) should behave similarly to

#### [4.1.](#) Service Request Process between Multiple Operators

As discussed previously, the reclusiveness principle applies seamlessly over the two CMIs. This implies that Operator 1's MDSC needs to pass all customer service requirements transparently to Operator 2's MDSC so that Operator 2 should provision its underlay network tunnels to meet the service requirements of the original customer. The MDSC of Operator 1 should translate/map the original customer's intent and service requirements and pass down to the corresponding PNC(s) which is(are) responsible for interfacing another operator (in this example, Operator 2) that provides

Lee, et al.

Expires April 2019

[Page 15]

---

Internet-Draft

ACTN Applicability to VPN with POI

October 2018

transport services for the segment of the customer's VN. The PNC in turn performs as a CNC when interfacing its southbound with Operator 2's MDSC.

It is possible that additional recursive relationships may also exist between Operator 2 and other operators.

#### [4.2.](#) Service/Network Orchestration of Operator 2

Operator 2 that provides transport service for Operator 1 may also need to perform service/network orchestration function just as the case for Operator 1.

From a data plane perspective, the interaction between operator 1 and operator 2 is similar to the POI case discussed in [section 3.2](#) (See Figure 2) with an exception that the packet domains belong to operator 1 while optical domain to operator 2.

The control interface depicted in Figure 4 (i.e., the CNC of operator 1 and the MDSC of operator 2) should behave similarly to that of the MDSC and the PNCs discussed in [Section 3](#).

### [5.](#) Security Considerations

This document defines key components and interfaces for managed traffic engineered networks. Securing the request and control of resources, confidentiality of the information, and availability of

function, should all be critical security considerations when deploying and operating ACTN POI platforms.

Several distributed ACTN functional components are required, and implementations should consider encrypting data that flows between components, especially when they are implemented at remote nodes, regardless these data flows are on external or internal network interfaces.

From a security and reliability perspective, ACTN POI may encounter many risks such as malicious attack and rogue elements attempting to connect to various ACTN POI components. Furthermore, some ACTN POI components represent a single point of failure and threat vector, and must also manage policy conflicts, and eavesdropping of communication between different ACTN POI components.

The conclusion is that all protocols used to realize the ACTN POI should have rich security features, and customer, application and network data should be stored in encrypted data stores. Additional

security risks may still exist. Therefore, discussion and applicability of specific security functions and protocols will be better described in documents that are use case and environment specific.

## [6.](#) IANA Considerations

This document has no actions for IANA.

## [7.](#) References

### [7.1.](#) Normative References

[RFC8453] D. Ceccarelli and Y. Lee, "Framework for Abstraction and Control of Transport Networks", [RFC 8453](#), August 2018.

### [7.2.](#) Informative References

[DHODY] D. Dhody, et al., "Packet Optical Integration (POI) Use Cases for Abstraction and Control of Transport Networks (ACTN)", [draft-dhody-actn-poi-use-case](#), work in progress.

[bgp-l3vpn] D. Jain, et al. "Yang Data Model for BGP/MPLS L3 VPNs",

[draft-ietf-bess-l3vpn-yang](#), work in progress.

[RFC4364] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.

[ACTN-VN] Y. Lee, et al., "A Yang Data Model for ACTN VN Operation", [draft-lee-teas-actn-vn-yang](#), work in progress.

[TSM] Y. Lee, et al., "Traffic Engineering and Service Mapping Yang Model", [draft-lee-teas-te-service-mapping-yang](#), work in progress.

[TE-Topo] X. Liu, et al., "YANG Data Model for Traffic Engineering (TE) Topologies", [draft-ietf-teas-yang-te-topo](#), work in progress.

[RFC8309] Q. Wu, W. Liu, and A. Farrel, "Service Models Explained", [RFC 8309](#), January 2018.

[L3SM] Q. Wu, S. Litkowski, L. Tomotaki, and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", [RFC 8299](#), January 2018.

[L2SM] G. Fioccola (Ed), "A YANG Data Model for L2VPN Service Delivery", [draft-ietf-l2sm-l2vpn-service-model](#), work in progress.

Lee, et al.

Expires April 2019

[Page 17]

---

Internet-Draft

ACTN Applicability to VPN with POI

October 2018

[ACTN-Telemetry] Y. Lee, et al., "YANG models for ACTN TE Performance Monitoring Telemetry and Network Autonomics", [draft-lee-teas-actn-pm-telemetry-autonomics](#), work in progress.

## 8. Contributors

Adrian Farrel  
Old Dog Consulting  
Email: [adrian@olddog.co.uk](mailto:adrian@olddog.co.uk)

Dhruv Dhody  
Huawei  
Email: [dhruv.dhody@huawei.com](mailto:dhruv.dhody@huawei.com)

Haomian Zheng

Huawei  
Email: haomianzheng@hauwei.com

#### Authors' Addresses

Young Lee  
Huawei Technologies  
Email: leeyoung@huawei.com

Qin Wu  
Huawei Technologies  
Email: bill.wu@huawei.com

Italo Busi  
Huawei Technologies  
Email: Italo.Busi@huawei.com

Daniele Ceccarelli  
Ericsson  
Email: daniele.ceccarelli@ericsson.com

Jeff Tantsura  
Nuage  
Email: jefftant.ietf@gmail.com