

TEAS Working Group
Internet Draft
Intended status: Informational

Fabio Peruzzini
TIM
Jean-Francois Bouquier
Vodafone
Italo Busi
Huawei
Daniel King
Old Dog Consulting
Daniele Ceccarelli
Ericsson

Expires: December 2020

June 29, 2020

Applicability of Abstraction and Control of Traffic Engineered
Networks (ACTN) to Packet Optical Integration (POI)

[draft-peru-teas-actn-poi-applicability-05](#)

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), 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/ietf/1id-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 December 29, 2020.

Internet-Draft

ACTN POI

June 2020

Copyright Notice

Copyright (c) 2020 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.

Abstract

This document considers the applicability of the IETF Abstraction and Control of Traffic Engineered Networks (ACTN) to Packet Optical Integration (POI), and IP and Optical DWDM domain internetworking.

In this document, we highlight the IETF protocols and YANG data models that may be used for the ACTN and control of POI networks, with particular focus on the interfaces between the MDSC (Multi-Domain Service Coordinator) and the underlying Packet and Optical Domain Controllers (P-PNC and O-PNC) to support POI use cases.

Table of Contents

1.	Introduction.....	3
2.	Reference Scenario.....	5
2.1.	Generic Assumptions.....	6
3.	Multi-Layer Topology Coordination.....	7
3.1.	Discovery of existing Och, ODU, IP links, IP tunnels and IP services.....	7
3.1.1.	Common YANG Models used at the MPI.....	8
3.1.1.1.	YANG models used at the Optical MPIs.....	8
3.1.1.2.	Required YANG models at the Packet MPIs.....	9
3.1.2.	Inter-domain link Discovery.....	9
3.2.	Provisioning of an IP Link/LAG over DWDM.....	10
3.2.1.	YANG models used at the MPIs.....	10
3.2.1.1.	YANG models used at the Optical MPIs.....	10

3.2.1.2. Required YANG models at the Packet MPIs.....	11
3.2.2. IP Link Setup Procedure.....	12

3.3. Provisioning of an IP link/LAG over DWDM with path constraints.....	12
3.3.1. YANG models used at the MPIs.....	13
3.4. Provisioning Link Members to an existing LAG.....	13
3.4.1. YANG Models used at the MPIs.....	13
4. Multi-Layer Recovery Coordination.....	13
4.1. Ensuring Network Resiliency during Maintenance Events....	13
4.2. Router Port Failure.....	13
5. Service Coordination for Multi-Layer network.....	14
5.1. L2/L3VPN/VN Service Request by the Customer.....	17
5.2. Service and Network Orchestration.....	19
5.3. IP/MPLS Domain Controller and NE Functions.....	23
5.3.1. Scenario A: Shared Tunnel Selection.....	23
5.3.1.1. Domain Tunnel Selection.....	24
5.3.1.2. VPN/VRF Provisioning for L3VPN.....	25
5.3.1.3. VSI Provisioning for L2VPN.....	26
5.3.1.4. Inter-domain Links Update.....	26
5.3.1.5. End-to-end Tunnel Management.....	26
5.3.2. Scenario B: Isolated VN/Tunnel Establishment.....	27
5.4. Optical Domain Controller and NE Functions.....	27
5.5. Orchestrator-Controllers-NEs Communication Protocol Flows	29
6. Security Considerations.....	31
7. Operational Considerations.....	31
8. IANA Considerations.....	31
9. References.....	31
9.1. Normative References.....	31
9.2. Informative References.....	32
Acknowledgments.....	34
Contributors.....	34
Authors' Addresses.....	35

[1. Introduction](#)

Packet Optical Integration (POI) is an advanced use case of traffic engineering. In wide-area networks, a packet network based on the Internet Protocol (IP) and possibly Multiprotocol Label Switching (MPLS) is typically deployed on top of an optical transport network that uses Dense Wavelength Division Multiplexing (DWDM). In many

existing network deployments, the packet and the optical networks are engineered and operated independently of each other. There are technical differences between the technologies (e.g., routers versus optical switches) and the corresponding network engineering and planning methods (e.g., inter-domain peering optimization in IP vs. dealing with physical impairments in DWDM, or very different time scales). In addition, customers and customer needs vary between a packet and an optical network, and it is not uncommon to use

different vendors in both domains. Last but not least, state-of-the-art packet and optical networks use sophisticated but complex technologies, and for a network engineer, it may not be trivial to be a full expert in both areas. As a result, packet and optical networks are often managed by different technical and organizational silos.

This separation is inefficient for many reasons. Both capital expenditure (CAPEX) and operational expenditure (OPEX) could be significantly reduced by better integrating the packet and the optical network. Multi-layer online topology insight can speed up troubleshooting (e.g., alarm correlation) and network operation (e.g., coordination of maintenance events), multi-layer offline topology inventory can improve service quality (e.g., detection of diversity constraint violations) and multi-layer traffic engineering can use the available network capacity more efficiently (e.g., coordination of restoration). In addition, provisioning workflows can be simplified or automated as needed across layers (e.g., to achieve bandwidth on demand, or to perform maintenance events).

Fully leveraging these benefits requires integration between the management and control of the packet and the optical network. The Abstraction and Control of TE Networks (ACTN) framework outlines the functional components and interfaces between a Multi-Domain Service Coordinator (MDSC) and Provisioning Network Controllers (PNCs) that can be used for coordinating the packet and optical layers.

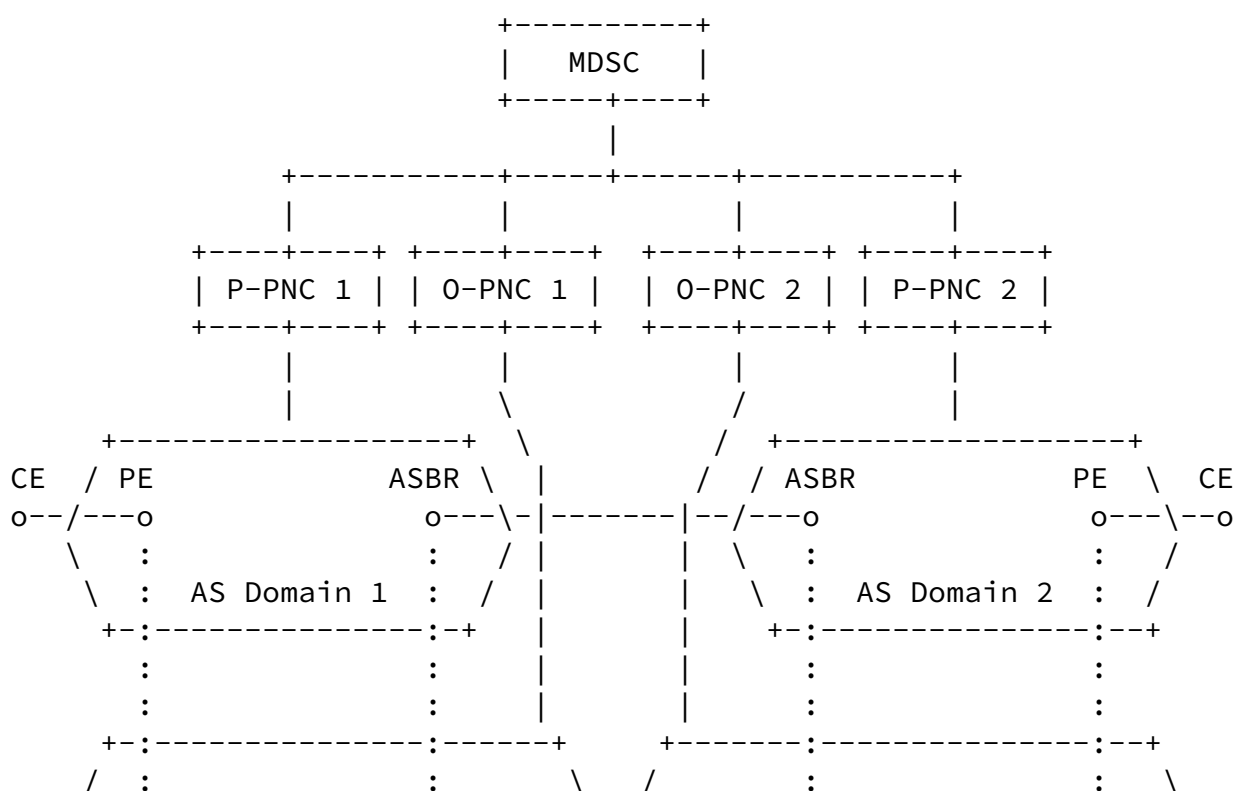
In this document, critical use cases for Packet Optical Integration (POI) are described. We outline how and what is required for the packet and the optical layer to interact to set up and operate services. The IP networks are operated as a client of optical networks. The use cases are ordered by increasing the level of integration and complexity. For each multi-layer use case, the

document analyzes how to use the interfaces and data models of the ACTN architecture.

The document also captures the current issues with ACTN and POI deployment. By understanding the level of standardization and potential gaps, it will help to better assess the feasibility of integration between IP and optical DWDM domain, in an end-to-end multi-vendor network.

2. Reference Scenario

This document uses "Reference Scenario 1" with multiple Optical domains and multiple Packet domains. The following Figure 1 shows this scenario in case of two Optical domains and two Packet domains:






Figure 1 - Reference Scenario 1

The ACTN architecture, defined in [\[RFC8453\]](#), is used to control this multi-domain network where each Packet PNC (P-PNC) is responsible for controlling its IP domain (AS), and each Optical PNC (O-PNC) is responsible for controlling its Optical Domain.

The MDSC is responsible for coordinating the whole multi-domain multi-layer (Packet and Optical) network. A specific standard interface (MPI) permits MDSC to interact with the different Provisioning Network Controller (O/P-PNCs).

The MPI interface presents an abstracted topology to MDSC hiding technology-specific aspects of the network and hiding topology details depending on the policy chosen regarding the level of abstraction supported. The level of abstraction can be obtained based on P-PNC and O-PNC configuration parameters (e.g. provide the potential connectivity between any PE and any ABR in an MPLS-TE network).

The MDSC in Figure 1 is responsible for multi-domain and multi-layer coordination across multiple Packet and Optical domains, as well as to provide IP services to different CNCs at its CMIs using YANG-based service models (e.g., using L2SM [\[RFC8466\]](#), L3SM [\[RFC8299\]](#)).

The multi-domain coordination mechanisms for the IP tunnels supporting these IP services are described in [section 5](#). In some cases, the MDSC could also rely on the multi-layer POI mechanisms, described in this draft, to support multi-layer optimizations for these IP services and tunnels.

In the network scenario of Figure 1, it is assumed that:

- o The domain boundaries between the IP and Optical domains are congruent. In other words, one Optical domain supports

connectivity between Routers in one and only one Packet Domain;

- o Inter-domain links exist only between Packet domains (i.e., between ASBR routers) and between Packet and Optical domains (i.e., between routers and ROADMs). In other words, there are no inter-domain links between Optical domains;
- o The interfaces between the routers and the ROADM's are "Ethernet" physical interfaces;
- o The interfaces between the ASBR routers are "Ethernet" physical interfaces.

[2.1](#). Generic Assumptions

This section describes general assumptions which are applicable at all the MPI interfaces, between each PNC (optical or packet) and the MDSC, and also to all the scenarios discussed in this document.

The data models used on these interfaces are assumed to use the YANG 1.1 Data Modeling Language, as defined in [[RFC7950](#)].

The RESTCONF protocol, as defined in [[RFC8040](#)], using the JSON representation, defined in [[RFC7951](#)], is assumed to be used at these interfaces.

As required in [[RFC8040](#)], the "ietf-yang-library" YANG module defined in [[RFC8525](#)] is used to allow the MDSC to discover the set of YANG modules supported by each PNC at its MPI.

[3](#). Multi-Layer Topology Coordination

In this scenario, the MDSC needs to discover the network topology, at both WDM and IP layers, in terms of nodes (NEs) and links, including inter-AS domain links as well as cross-layer links.

Each PNC provides to the MDSC an abstract topology view of the WDM or of the IP topology of the domain it controls. This topology is abstracted in the sense that some detailed NE information is hidden at the MPI, and all or some of the NEs and related physical links

are exposed as abstract nodes and logical (virtual) links, depending on the level of abstraction the user requires. This detailed information is vital to understand both the inter-AS domain links (seen by each controller as UNI interfaces but as I-NNI interfaces by the MDSC) as well as the cross-layer mapping between IP and WDM layer.

The MDSC also maintains an up-to-date network inventory of both IP and WDM layers through the use of IETF notifications through MPI with the PNCs.

For the cross-layer links, the MDSC needs to be capable of automatically correlating physical ports information from the routers (single link or bundle links for link aggregation groups - LAG) to client ports in the ROADM.

[3.1.](#) Discovery of existing Och, ODU, IP links, IP tunnels and IP services

Typically, an MDSC must be able to automatically discover network topology of both WDM and IP layers (links and NE, links between two domains), this assumes the following:

- o An abstract view of the WDM and IP topology must be available;

- o MDSC must keep an up-to-date network inventory of both IP and WDM layers, and it should be possible to correlate such information (e.g., which port, lambda/OTSi, the direction it is used by a specific IP service on the WDM equipment);
- o It should be possible at MDSC level to easily correlate WDM and IP layers alarms to speed-up troubleshooting.

[3.1.1.](#) Common YANG Models used at the MPI

Both optical and packet PNCs use the following common topology YANG models at the MPI to report their abstract topologies:

- o The Base Network Model, defined in the "ietf-network" YANG module of [[RFC8345](#)];
- o The Base Network Topology Model, defined in the "ietf-network-topology" YANG module of [[RFC8345](#)], which augments the Base Network Model;
- o The TE Topology Model, defined in the "ietf-te-topology" YANG module of [[TE-TOP](#)], which augments the Base Network Topology Model.

These IETF YANG models are generic and augmented by technology-specific YANG modules as described in the following sections.

[3.1.1.1](#). YANG models used at the Optical MPIs

The optical PNC also uses at least the following technology-specific topology YANG models, providing WDM and Ethernet technology-specific augmentations of the generic TE Topology Model:

- o The WSON Topology Model, defined in the "ietf-wson-topology" YANG modules of [[WSON-TOP](#)], or the Flexi-grid Topology Model, defined in the "ietf-flexi-grid-topology" YANG module of [[Flexi-TOP](#)].
- o The Ethernet Topology Model, defined in the "ietf-eth-te-topology" YANG module of [[CLIENT-TOP](#)]

The WSON Topology Model or, alternatively, the Flexi-grid Topology model is used to report the fixed-grid or, respectively, the flexible-grid DWDM network topology (e.g., ROADMs and OMS links).

The Ethernet Topology Model is used to report the Ethernet access links on the edge ROADMs.

[3.1.1.2](#). Required YANG models at the Packet MPIs

The Packet PNC also uses at least the following technology-specific topology YANG models, providing IP and Ethernet technology-specific augmentations of the generic Topology Models:

- o The L3 Topology Model, defined in the "ietf-l3-unicast-topology" YANG modules of [[RFC8346](#)], which augments the Base Network

Topology Model

- o The Ethernet Topology Model, defined in the "ietf-eth-te-topology" YANG module of [[CLIENT-TOPO](#)], which augments the TE Topology Model
- o The L3-TE Topology Model, defined in the "ietf-l3-te-topology" YANG modules of [[L3-TE-TOPO](#)], which augments the L3 Topology Model

The Ethernet Topology Model is used to report the Ethernet links between the IP routers and the edge ROADMs as well as the inter-domain links between ASBRs, while the L3 Topology Model is used to report the IP network topology (e.g., IP routers and IP links).

The L3-TE Topology Model reports the relationship between the IP routers and LTPs provided by the L3 Topology Model and the underlying Ethernet nodes and LTPs provided by the Ethernet Topology Model.

[3.1.2](#). Inter-domain link Discovery

In the reference network of Figure 1, there are two types of inter-domain links:

- o Links between two IP domains/ASBRs (ASes)
- o Links between an IP router and a ROADM

Both types of links are Ethernet physical links.

The inter-domain link information is reported to the MDSC by the two adjacent PNCs, controlling the two ends of the inter-domain link, using the Ethernet Topology Model defined in [[CLIENT-TOPO](#)].

The MDSC can understand how to merge these inter-domain Ethernet links together using the plug-id attribute defined in the TE

A more detailed description of how the plug-id can be used to discover inter-domain link is also provided in section 5.1.4 of [\[TNBI\]](#).

Both types of inter-domain Ethernet links are discovered using the plug-id attributes reported in the Ethernet Topologies exposed by the two adjacent PNCs.

The MDSC, when discovering an Ethernet inter-domain link between two Ethernet LTPs which are associated with two IP LTPs, reported in the IP Topologies exposed by the two adjacent P-PNCs, can also discover an inter-domain IP link/adjacency between these two IP LTPs.

Two options are possible to discover these inter-domain Ethernet links:

1. Static configuration
2. LLDP [IEEE 802.1AB] automatic discovery

Since the static configuration requires an administrative burden to configure network-wide unique identifiers, the automatic discovery solution based on LLDP is preferable when LLDP is supported.

As outlined in [\[TNBI\]](#), the encoding of the plug-id namespace as well as of the LLDP information within the plug-id value is implementation specific and needs to be consistent across all the PNCs.

[3.2](#). Provisioning of an IP Link/LAG over DWDM

In this scenario, the MDSC needs to coordinate the creation of an IP link, or a LAG, between two routers through a DWDM network.

It is assumed that the MDSC has already discovered the whole network topology as described in [section 3.1](#).

[3.2.1](#). YANG models used at the MPIs

[3.2.1.1](#). YANG models used at the Optical MPIs

The optical PNC uses at least the following YANG models:

- o The TE Tunnel Model, defined in the "ietf-te" YANG module of [[TE-TUNNEL](#)]
- o The WSON Tunnel Model, defined in the "ietf-wson-tunnel" YANG modules of [[WSON-TUNNEL](#)], or the Flexi-grid Media Channel Model, defined in the "ietf-flexi-grid-media-channel" YANG module of [[Flexi-MC](#)]
- o The Ethernet Client Signal Model, defined in the "ietf-eth-tran-service" YANG module of [[CLIENT-SIGNAL](#)]

The TE Tunnel model is generic and augmented by technology-specific models such as the WSON Tunnel Model and the Flexi-grid Media Channel Model.

The WSON Tunnel Model or, alternatively, the Flexi-grid Media Channel Model are used to setup connectivity within the DWDM network depending on whether the DWDM optical network is based on fixed grid or flexible-grid.

The Ethernet Client Signal Model is used to configure the steering of the Ethernet client traffic between Ethernet access links and TE Tunnels, which in this case could be either WSON Tunnels or Flexi-Grid Media Channels. This model is generic and applies to any technology-specific TE Tunnel: technology-specific attributes are provided by the technology-specific models which augment the generic TE-Tunnel Model.

[3.2.1.2](#). Required YANG models at the Packet MPIs

The Packet PNC uses at least the following topology YANG models:

- o The Base Network Model, defined in the "ietf-network" YANG module of [[RFC8345](#)] (see [section 3.1.1](#))
- o The Base Network Topology Model, defined in the "ietf-network-topology" YANG module of [[RFC8345](#)] (see [section 3.1.1](#))
- o The L3 Topology Model, defined in the "ietf-l3-unicast-topology" YANG modules of [[RFC8346](#)] (see [section 3.1.1.1](#))

If, as discussed in [section 3.2.2](#), IP Links created over DWDM can be automatically discovered by the P-PNC, the IP Topology is needed only to report these IP Links after being discovered by the P-PNC.

Internet-Draft

ACTN POI

June 2020

The IP Topology can also be used to configure the IP Links created over DWDM.

[3.2.2](#). IP Link Setup Procedure

The MDSC requires the O-PNC to setup a WDM Tunnel (either a WSON Tunnel or a Flexi-grid Tunnel) within the DWDM network between the two Optical Transponders (OTs) associated with the two access links.

The Optical Transponders are reported by the O-PNC as Trail Termination Points (TTPs), defined in [\[TE-TOPO\]](#), within the WDM Topology. The association between the Ethernet access link and the WDM TTP is reported by the Inter-Layer Lock (ILL) identifiers, defined in [\[TE-TOPO\]](#), reported by the O-PNC within the Ethernet Topology and WDM Topology.

The MDSC also requires the O-PNC to steer the Ethernet client traffic between the two access Ethernet Links over the WDM Tunnel.

After the WDM Tunnel has been setup and the client traffic steering configured, the two IP routers can exchange Ethernet packets between themselves, including LLDP messages.

If LLDP [IEEE 802.1AB] is used between the two routers, the P-PNC can automatically discover the IP Link being set up by the MDSC. The IP LTPs terminating this IP Link are supported by the ETH LTPs terminating the two access links.

Otherwise, the MDSC needs to require the P-PNC to configure an IP Link between the two routers: the MDSC also configures the two ETH LTPs which support the two IP LTPs terminating this IP Link.

[3.3](#). Provisioning of an IP link/LAG over DWDM with path constraints

MDSC must be able to provision an IP link with a fixed maximum latency constraint, or with the minimum latency available constraint within each domain but as well inter-domain when required (e.g. by monitoring traffic KPIs trends for this IP link). Through the O-PNC fixed latency path/minimum latency path is chosen between PE and ASBR in each optical domain. Then MDSC needs to select the inter-AS domain with less latency (in case we have several interconnection links) to have the right low latency constraint fulfilled end-to-end across domains.

Internet-Draft

ACTN POI

June 2020

MDSC must be able to automatically create two IP links between two routers, over DWDM network, with physical path diversity (avoiding SRLGs communicated by O-PNCs to the MDSC).

MDSC must be responsible for routing each of this IP links through different inter-AS domain links so that end-to-end IP links are fully disjoint.

Optical connectivity must be set up accordingly by MDSC through O-PNCs.

[3.3.1](#). YANG models used at the MPIs

This section is for further study

[3.4](#). Provisioning Link Members to an existing LAG

When adding a new link member to a LAG between two routers with or without path latency/diversity constraint, the MDSC must be able to force the additional optical connection to use the same physical path in the optical domain where the LAG capacity increase is required.

[3.4.1](#). YANG Models used at the MPIs

This is for further study

[4](#). Multi-Layer Recovery Coordination

[4.1](#). Ensuring Network Resiliency during Maintenance Events

Before planned maintenance operation on DWDM network takes place, IP traffic should be moved hitless to another link.

MDSC must reroute IP traffic before the events takes place. It should be possible to lock IP traffic to the protection route until the maintenance event is finished, unless a fault occurs on such path.

[4.2.](#) Router Port Failure

The focus is on client-side protection scheme between IP router and reconfigurable ROADMs. Scenario here is to define only one port in the routers and in the ROADM muxponder board at both ends as back-up ports to recover any other port failure on client-side of the ROADM (either on router port side or on muxponder side or on the link

between them). When client-side port failure occurs, alarms are raised to MDSC by IP-PNC and O-PNC (port status down, LOS etc.). MDSC checks with OP-PNC(s) that there is no optical failure in the optical layer.

There can be two cases here:

- a) LAG was defined between the two end routers. MDSC, after checking that optical layer is fine between the two end ROADMs, triggers the ROADM configuration so that the router back-up port with its associated muxponder port can reuse the OCh that was already in use previously by the failed router port and adds the new link to the LAG on the failure side.

While the ROADM reconfiguration takes place, IP/MPLS traffic is using the reduced bandwidth of the IP link bundle, discarding lower priority traffic if required. Once backup port has been reconfigured to reuse the existing OCh and new link has been added to the LAG then original Bandwidth is recovered between the end routers.

Note: in this LAG scenario let assume that BFD is running at LAG level so that there is nothing triggered at MPLS level when one of the link member of the LAG fails.

- b) If there is no LAG then the scenario is not clear since a router port failure would automatically trigger (through BFD failure) first a sub-50ms protection at MPLS level :FRR (MPLS RSVP-TE case) or TI-LFA (MPLS based SR-TE case) through a protection port. At the same time MDSC, after checking that optical network connection is still fine, would trigger the reconfiguration of the back-up port of the router and of the ROADM muxponder to reuse the same OCh as the one used originally for the failed router

port. Once everything has been correctly configured, MDSC Global PCE could suggest to the operator to trigger a possible re-optimisation of the back-up MPLS path to go back to the MPLS primary path through the back-up port of the router and the original OCh if overall cost, latency etc. is improved. However, in this scenario, there is a need for protection port PLUS back-up port in the router which does not lead to clear port savings.

5. Service Coordination for Multi-Layer network

[Editors' Note] This text has been taken from section 2 of [draft-lee-teas-actn-poi-applicability-00](#) and need to be reconciled with the other sections (the introduction in particular) of this document

This section provides a number of deployment scenarios for packet and optical integration (POI). Specifically, this section provides a deployment scenario in which ACTN hierarchy is deployed to control a multi-layer and multi-domain network via two IP/MPLS PNCs and two Optical PNCs with coordination with L-MDSC. This scenario is in the context of an upper layer service configuration (e.g. L3VPN) across two AS domains which are transported by two transport underlay domains (e.g. OTN).

The provisioning of the L3VPN service is outside ACTN scope but it is worth showing how the L3VPN service provisioning is integrated for the end-to-end service fulfilment in ACTN context. An example of service configuration function in the Service/Network Orchestrator is discussed in [[BGP-L3VPN](#)].

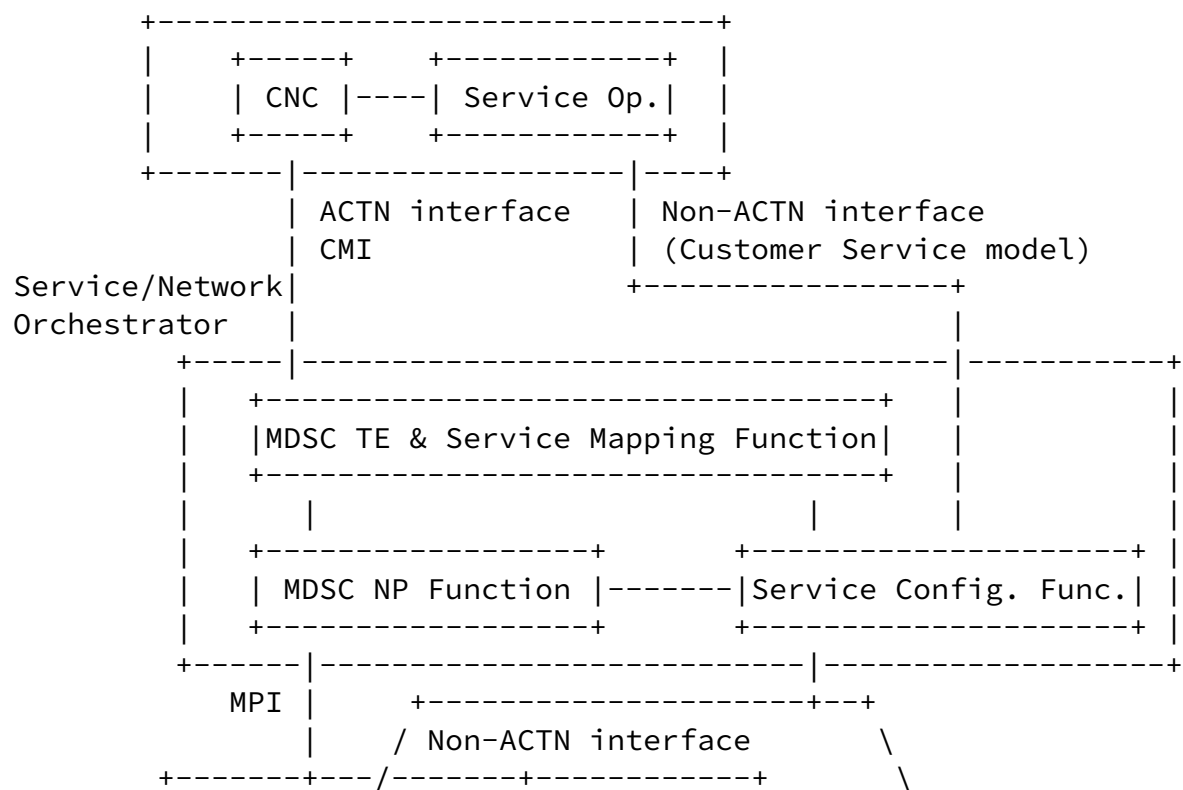
Figure 2 shows an ACTN POI Reference Architecture where it shows ACTN components as well as non-ACTN components that are necessary for the end-to-end service fulfilment. Both IP/MPLS and Optical Networks are multi-domain. Each IP/MPLS domain network is controlled by its' domain controller and all the optical domains are controlled by a hierarchy of optical domain controllers. The L-MDSC function of the optical domain controllers provides an abstract view of the whole optical network to the Service/Network Orchestrator. It is assumed that all these components of the network belong to one single network operator domain under the control of the service/network orchestrator.

Internet-Draft

ACTN POI

June 2020

Customer



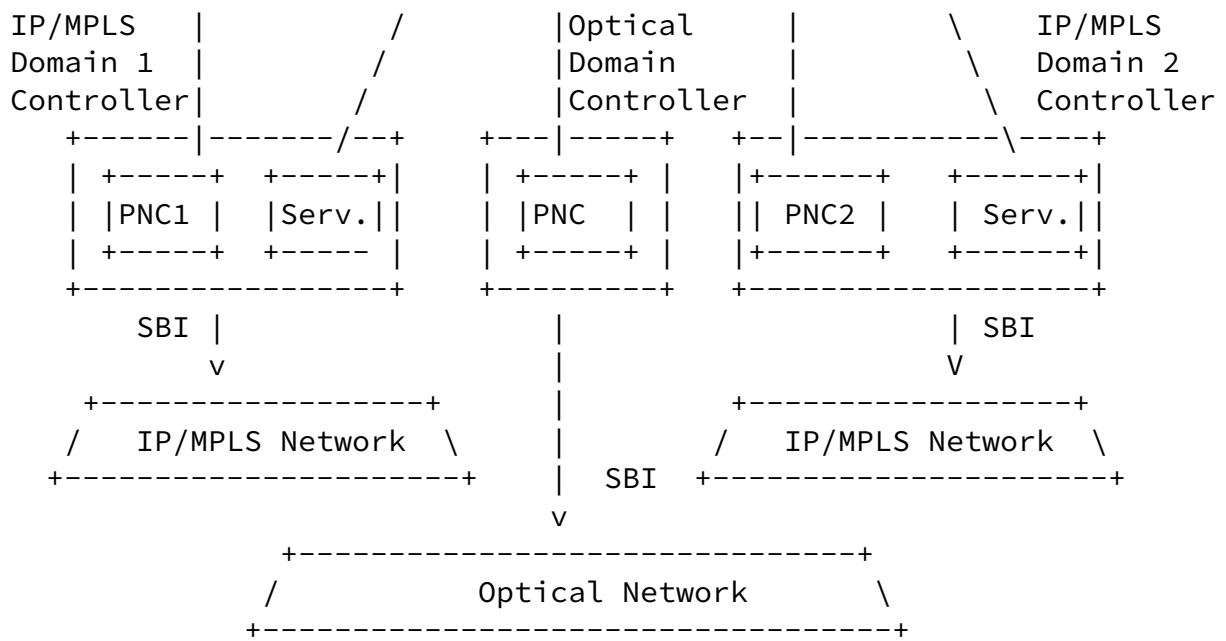


Figure 2 ACTN POI Reference Architecture

Figure 2 shows ACTN POI Reference Architecture where it depicts:

- o CMI (CNC-MDSC Interface) interfacing CNC with MDSC function in the Service/Network Orchestrator. This is where TE & Service Mapping [TSM] and either ACTN VN [ACTN-VN] or TE-topology [TE-TOPO] model is exchanged over CMI.
- o Customer Service Model Interface: Non-ACTN interface in the Customer Portal interfacing Service/Network Orchestrator's Service Configuration Function. This is the interface where L3SM information is exchanged.
- o MPI (MDSC-PNC Interface) interfacing IP/MPLS Domain Controllers and Optical Domain Controllers.
- o Service Configuration Interface: Non-ACTN interface in Service/Network Orchestrator interfacing with the IP/MPLS Domain Controllers to coordinate L2/L3VPN multi-domain service configuration. This is where service specific information such as VPN, VPN binding policy (e.g., new underlay tunnel creation for

isolation), etc. are conveyed.

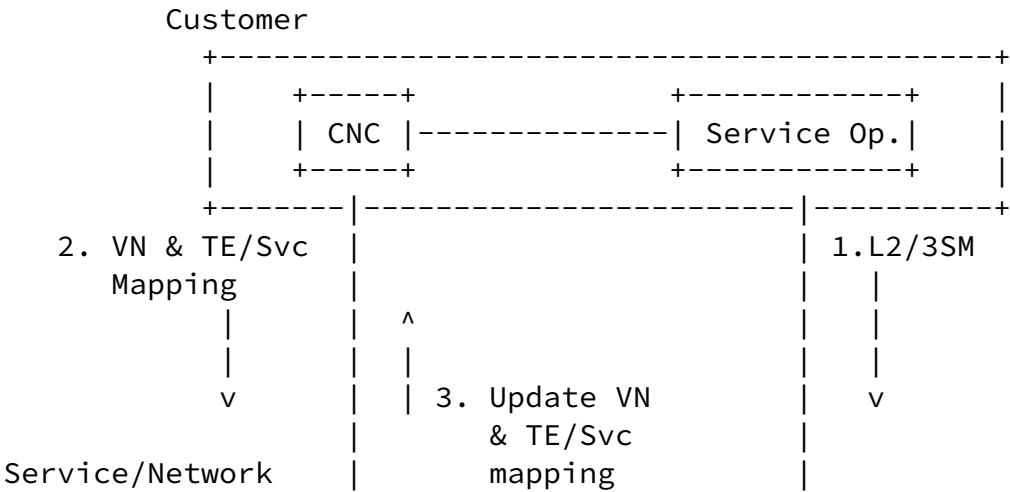
- o SBI (South Bound Interface): Non-ACTN interface in the domain controller interfacing network elements in the domain.

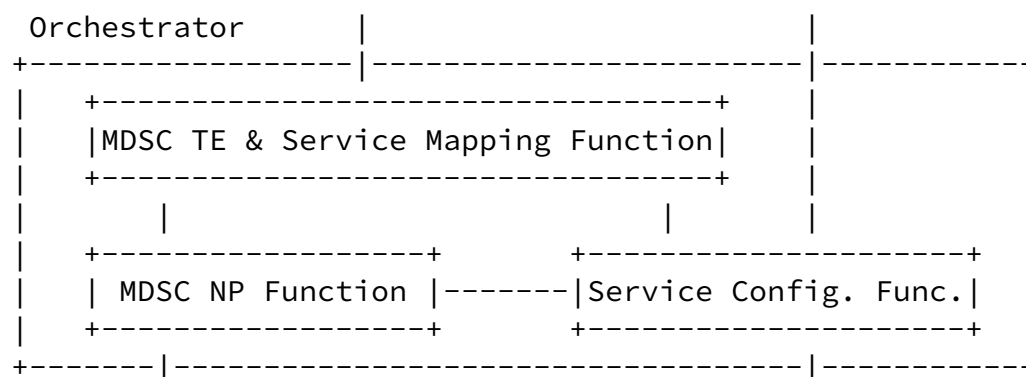
Please note that MPI and Service Configuration Interface can be implemented as the same interface with the two different capabilities. The split is just functional but doesn't have to be also logical.

The following sections are provided to describe key functions that are necessary for the vertical as well as horizontal end-to-end service fulfilment of POI.

5.1. L2/L3VPN/VN Service Request by the Customer

A customer can request L3VPN services with TE requirements using ACTN CMI models (i.e., ACTN VN YANG, TE & Service Mapping YANG) and non-ACTN customer service models such as L2SM/L3SM YANG together. Figure 3 shows detailed control flow between customer and service/network orchestrator to instantiate L2/L3VPN/VN service request.





NP: Network Provisioning

Figure 3 Service Request Process

- o ACTN VN YANG provides VN Service configuration, as specified in [\[ACTN-VN\]](#).
- o It provides the profile of VN in terms of VN members, each of which corresponds to an edge-to-edge link between customer end-points (VNAPs). It also provides the mappings between the VNAPs with the LTPs and between the connectivity matrix with the VN member from which the associated traffic matrix (e.g., bandwidth, latency, protection level, etc.) of VN member is expressed (i.e., via the TE-topology's connectivity matrix).
- o The model also provides VN-level preference information (e.g., VN member diversity) and VN-level admin-status and operational-status.

- o L2SM YANG [\[RFC8466\]](#) provides all L2VPN service configuration and site information from a customer/service point of view.
- o L3SM YANG [\[RFC8299\]](#) provides all L3VPN service configuration and site information from a customer/service point of view.
- o The TE & Service Mapping YANG model [\[TSM\]](#) provides TE-service mapping as well as site mapping.

- o TE-service mapping provides the mapping of L3VPN instance from [[RFC8299](#)] with the corresponding ACTN VN instance.
- o The TE-service mapping also provides the service mapping requirement type as to how each L2/L3VPN/VN instance is created with respect to the underlay TE tunnels (e.g., whether the L3VPN requires a new and isolated set of TE underlay tunnels or not, etc.). See [Section 5.2](#) for detailed discussion on the mapping requirement types.
- o Site mapping provides the site reference information across L2/L3VPN Site ID, ACTN VN Access Point ID, and the LTP of the access link.

[5.2](#). Service and Network Orchestration

The Service/Network orchestrator shown in Figure 2 interfaces the customer and decouples the ACTN MDSC functions from the customer service configuration functions.

An implementation can choose to split the Service/Network orchestration functions, as described in [[RFC8309](#)] and in [section 4.2 of \[RFC8453\]](#), between a top-level Service Orchestrator interfacing the customer and two low-level Network Orchestrators, one controlling a multi-domain IP/MPLS network and the other controlling the Optical networks.

Another implementation can choose to combine the L-MDSC functions of the Optical hierarchical controller, providing multi-domain coordination of the Optical network together with the MDSC functions in the Service/Network orchestrator.

Without loss of generality, this assumes that the service/network orchestrator as depicted in Figure 2 would include all the required functionalities as in a hierarchical orchestration case.

One of the important service functions the Service/Network orchestrator performs is to identify which TE Tunnels should carry the L3VPN traffic (from TE & Service Mapping Model) and to relay

this information to the IP/MPLS domain controllers, via non-ACTN interface, to ensure proper IP/VRF forwarding table be populated according to the TE binding requirement for the L3VPN.

[Editor's Note] What mechanism would convey on the interface to the IP/MPLS domain controllers as well as on the SBI (between IP/MPLS domain controllers and IP/MPLS PE routers) the TE binding policy dynamically for the L3VPN? Typically, VRF is the function of the device that participate MP-BGP in MPLS VPN. With current MP-BGP implementation in MPLS VPN, the VRF's BGP next hop is the destination PE and the mapping to a tunnel (either an LDP or a BGP tunnel) toward the destination PE is done by automatically without any configuration. It is to be determined the impact on the PE VRF operation when the tunnel is an optical bypass tunnel which does not participate either LDP or BGP.

Figure 4 shows service/network orchestrator interactions with various domain controllers to instantiate tunnel provisioning as well as service configuration.

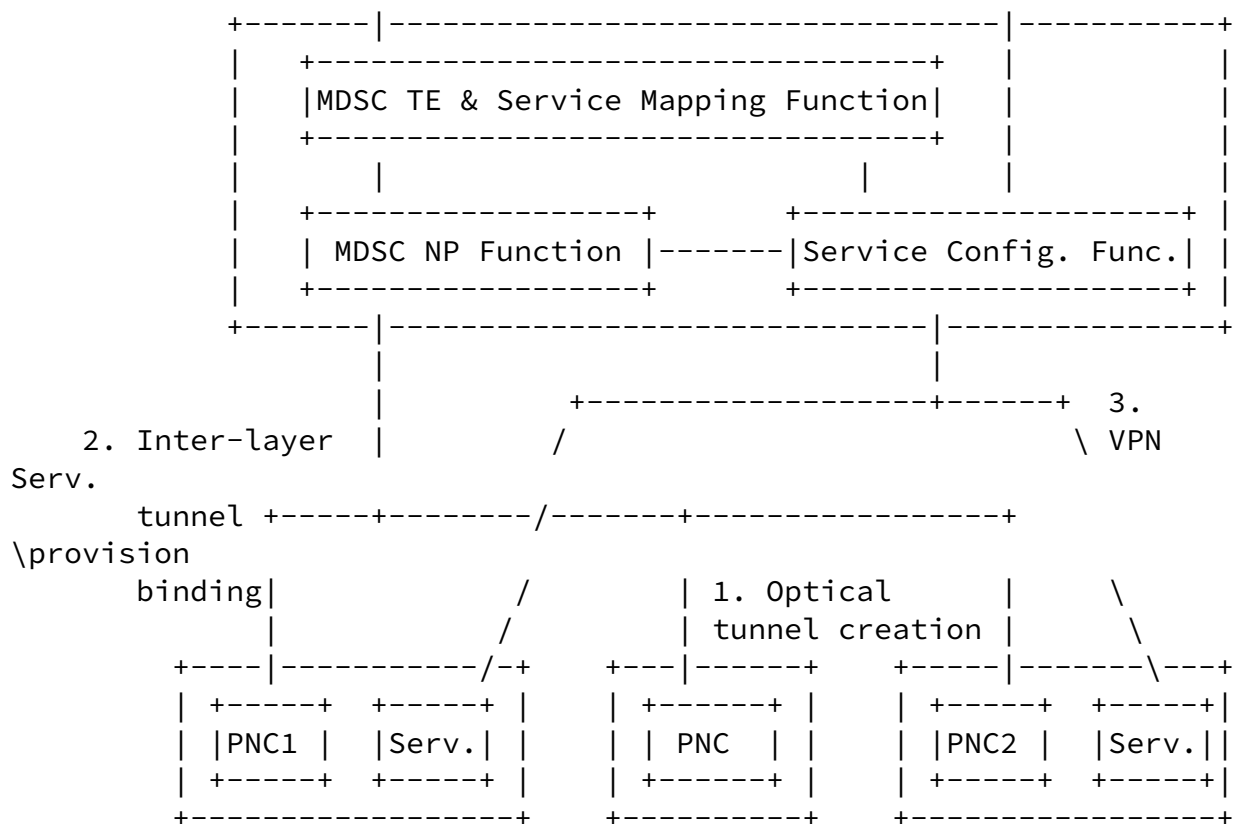


Figure 4 Service and Network Orchestration Process

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

1. Hard Isolation with deterministic latency: Customer would request an L3VPN service [[RFC8299](#)] using a set of TE Tunnels with a deterministic latency requirement and that cannot be not shared with other L3VPN services nor compete for bandwidth with other Tunnels.
2. Hard Isolation: This is similar to the above case without deterministic latency requirements.
3. Soft Isolation: Customer would request an L3VPN service using a set of MPLS-TE tunnel which cannot be shared with other L3VPN services.
4. Sharing: Customer would accept sharing the MPLS-TE Tunnels supporting its L3VPN service with other services.

For the first three types, there could be additional TE binding requirements with respect to different VN members of the same VN

Internet-Draft

ACTN POI

June 2020

associated with an L3VPN service. For the first two cases, VN members can be hard-isolated, soft-isolated, or shared. For the third case, VN members can be soft-isolated or shared.

- o When "Hard Isolation with or w/o deterministic latency" (i.e., the first and the second type) TE binding requirement is applied for a L3VPN, a new optical layer tunnel has to be created (Step 1 in Figure 4). This operation requires the following control level mechanisms as follows:
 - o The MDSC function of the Service/Network Orchestrator identifies only the domains in the IP/MPLS layer in which the VPN needs to be forwarded.
 - o Once the IP/MPLS layer domains are determined, the MDSC function of the Service/Network Orchestrator needs to identify the set of optical ingress and egress points of the underlay optical tunnels providing connectivity between the IP/MPLS layer domains.
 - o Once both IP/MPLS layers and optical layer are determined, the MDSC needs to identify the inter-layer peering points in both IP/MPLS domains as well as the optical domain(s). This implies that the L3VPN traffic will be forwarded to an MPLS-TE tunnel that starts at the ingress PE (in one IP/MPLS domain) and terminates at the egress PE (in another IP/MPLS domain) via a dedicated underlay optical tunnel.
- o The MDSC function of the Service/Network Orchestrator needs to first request the optical L-MDSC to instantiate an optical tunnel for the optical ingress and egress. This is referred to as optical tunnel creation (Step 1 in Figure 4). Note that it is L-MDSC responsibility to perform multi-domain optical coordination with its underlying optical PNCs, for setting up a multi-domain optical tunnel.
- o Once the optical tunnel is established, then the MDSC function of the Service/Network Orchestrator needs to coordinate with the PNC functions of the IP/MPLS Domain Controllers (under which the ingress and egress PEs belong) the setup of a multi-domain MPLS-TE Tunnel, between the ingress and egress PEs. This setup is carried by the created underlay optical tunnel (Step 2 in Figure 4).

Internet-Draft

ACTN POI

June 2020

- o It is the responsibility of the Service Configuration Function of the Service/Network Orchestrator to identify interfaces/labels on both ingress and egress PEs and to convey this information to both the IP/MPLS Domain Controllers (under which the ingress and egress PEs belong) for proper configuration of the L3VPN (BGP and VRF function of the PEs) in their domain networks (Step 3 in Figure 4).

[5.3.](#) IP/MPLS Domain Controller and NE Functions

IP/MPLS networks are assumed to have multiple domains and each domain is controlled by IP/MPLS domain controller in which the ACTN PNC functions and non-ACTN service functions are performed by the IP/MPLS domain controller.

Among the functions of the IP/MPLS domain controller are VPN service aspect provisioning such as VRF control and management for VPN services, etc. It is assumed that BGP is running in the inter-domain IP/MPLS networks for L2/L3VPN and that the IP/MPLS domain controller is also responsible for configuring the BGP speakers within its control domain if necessary.

Depending on the TE binding requirement types discussed in [Section 5.2](#), there are two possible deployment scenarios.

[5.3.1.](#) Scenario A: Shared Tunnel Selection

When the L2/L3VPN does not require isolation (either hard or soft), it can select an existing MPLS-TE and Optical tunnel between ingress and egress PE, without creating any new TE tunnels. Figure 5 shows this scenario.

Internet-Draft

ACTN POI

June 2020

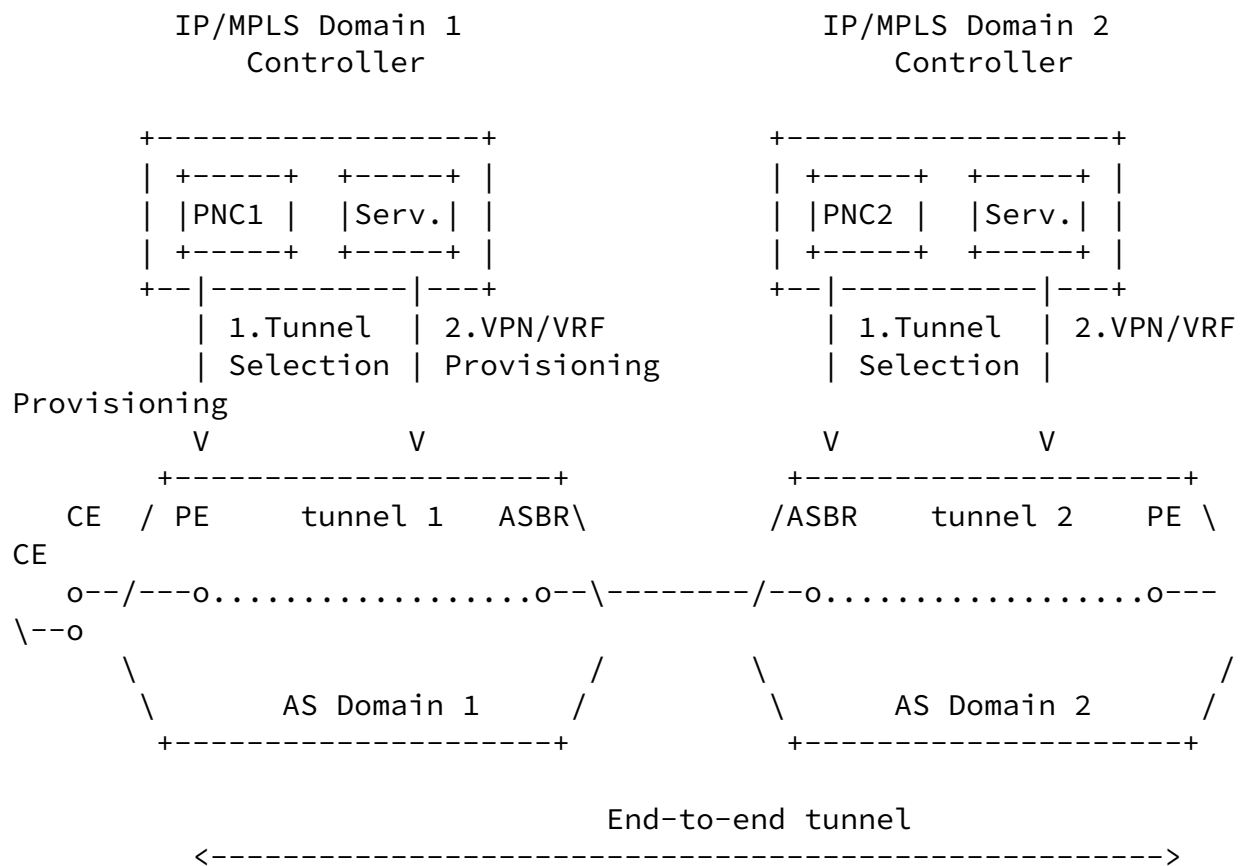


Figure 5 IP/MPLS Domain Controller & NE Functions

How VPN is disseminated across the network is out of the scope of this document. We assume that MP-BGP is running in IP/MPLS networks and VPN is made known to ASBRs and PEs by each IP/MPLS domain controllers. See [[RFC4364](#)] for detailed descriptions on how MP-BGP works.

There are several functions IP/MPLS domain controllers need to provide in order to facilitate tunnel selection for the VPN in both domain level and end-to-end level.

[5.3.1.1](#). Domain Tunnel Selection

Each domain IP/MPLS controller is responsible for selecting its domain level tunnel for the L3VPN. First it needs to determine which existing tunnels would fit for the L2/L3VPN requirements allotted to the domain by the Service/Network Orchestrator (e.g., tunnel binding, bandwidth, latency, etc.). If there are existing tunnels that are feasible to satisfy the L3VPN requirements, the IP/MPLS domain controller selects the optimal tunnel from the candidate

pool. Otherwise, an MPLS tunnel with modified bandwidth or a new MPLS Tunnel needs to be setup. Note that with no isolation requirement for the L3VPN, existing MPLS tunnel can be selected. With soft isolation requirement for the L3VPN, an optical tunnel can be shared with other L2/L3VPN services while with hard isolation requirement for the L2/L3VPN, a dedicated MPLS-TE and a dedicated optical tunnel MUST be provisioned for the L2/L3VPN.

[5.3.1.2](#). VPN/VRF Provisioning for L3VPN

Once the domain level tunnel is selected for a domain, the Service Function of the IP/MPLS domain controller maps the L3VPN to the selected MPLS-TE tunnel and assigns a label (e.g., MPLS label) with the PE. Then the PE creates a new entry for the VPN in the VRF forwarding table so that when the VPN packet arrives to the PE, it will be able to direct to the right interface and PUSH the label assigned for the VPN. When the PE forwards a VPN packet, it will push the VPN label signaled by BGP and, in case of option A and B [[RFC4364](#)], it will also push the LSP label assigned to the configured MPLS-TE Tunnel to reach the ASBR next hop and forwards the packet to the MPLS next-hop of this MPLS-TE Tunnel.

In case of option C [[RFC4364](#)], the PE will push one MPLS LSP label signaled by BGP to reach the destination PE and a second MPLS LSP label assigned to the configured MPLS-TE Tunnel to reach the ASBR next-hop and forward the packet to the MPLS next-hop of this MPLS-TE Tunnel.

With Option C, the ASBR of the first domain interfacing the next domain should keep the VPN label intact to the ASBR of the next domain so that the ASBR in the next domain sees the VPN packets as if they are coming from a CE. With Option B, the VPN label is swapped. With option A, the VPN label is removed.

With Option A and B, the ASBR of the second domain does the same procedure that includes VPN/VRF tunnel mapping and interface/label assignment with the IP/MPLS domain controller. With option A, the ASBR operations are the same as of the PEs. With option B, the ASBR operates with VPN labels so it can see the VPN the traffic belongs to. With option C, the ASBR operates with the end-to-end tunnel labels so it may be not aware of the VPN the traffic belongs to.

This process is repeated in each domain. The PE of the last domain interfacing the destination CE should recognize the VPN label when

the VPN packets arrive and thus POP the VPN label and forward the packets to the CE.

5.3.1.3. VSI Provisioning for L2VPN

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

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

[5.3.1.4](#). 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.

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

It is foreseen that the Service/Network orchestrator should control and manage end-to-end tunnels for VPNs per VPN policy.

As discussed in [[ACTN-PM](#)], the Orchestrator 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.

[5.3.2](#). Scenario B: Isolated VN/Tunnel Establishment

When the L3VPN requires hard-isolated Tunnel establishment, optical layer tunnel binding with IP/MPLS layer is necessary. As such, the following functions are necessary.

- o The IP/MPLS Domain Controller of Domain 1 needs to send the VRF instruction to the PE:
 - o To the Ingress PE of AS Domain 1: Configuration for each L3VPN destination IP address (in this case the remote CE's IP address for the VPN or any customer's IP addresses reachable through a remote CE) of the associated VPN label assigned by the Egress PE and of the MPLS-TE Tunnel to be used to reach the Egress PE: so that the proper VRF table is populated to forward the VPN traffic to the inter-layer optical interface with the VPN label.
- o The Egress PE, upon the discovery of a new IP address, needs to send the mapping information (i.e., VPN to IP address) to its'

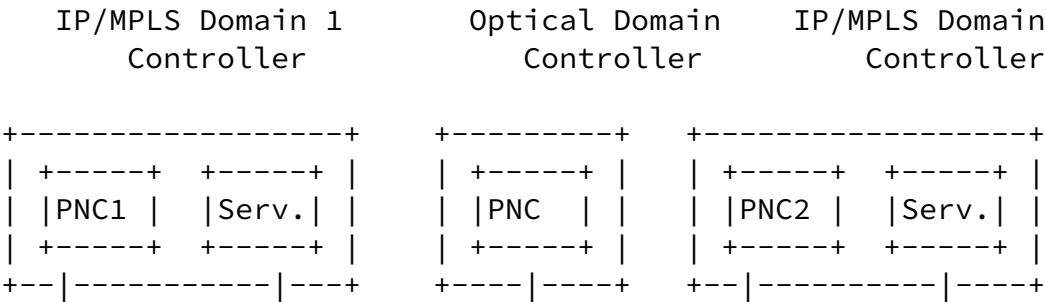
IP/MPLS Domain Controller of Domain 2 which sends, in turn, to the service orchestrator. The service orchestrator would then propagate this mapping information to the IP/MPLS Domain Controller of Domain 1 which sends it, in turn, to the ingress PE so that it may override the VPN/VRF forwarding or VSI forwarding, respectively for L3VPN and L2VPN. As a result, when packets arriving at the ingress PE with that IP destination address, the ingress PE would then forward this packet to the inter-layer optical interface.

[Editor's Note] in case of hard isolated tunnel required for the VPN, we need to create a separate MPLS TE tunnel and encapsulate the MPLS packets of the MPLS Tunnel into the ODU so that the optical NE would route this MPLS Tunnel to a separate optical tunnel from other tunnels.]

5.4. Optical Domain Controller and NE Functions

Optical network provides the underlay connectivity services to IP/MPLS networks. The multi-domain optical network coordination is performed by the L-MDSC function shown in Figure 2 so that the whole multi-domain optical network appears to the service/network orchestrator as one optical network. The coordination of Packet/Optical multi-layer and IP/MPLS multi-domain is done by the service/network orchestrator where it interfaces two IP/MPLS domain controllers and one optical L-MDSC.

Figure 6 shows how the Optical Domain Controllers create a new optical tunnel and the related interaction with IP/MPLS domain controllers and the NEs to bind the optical tunnel with proper forwarding instruction so that the VPN requiring hard isolation can be fulfilled.



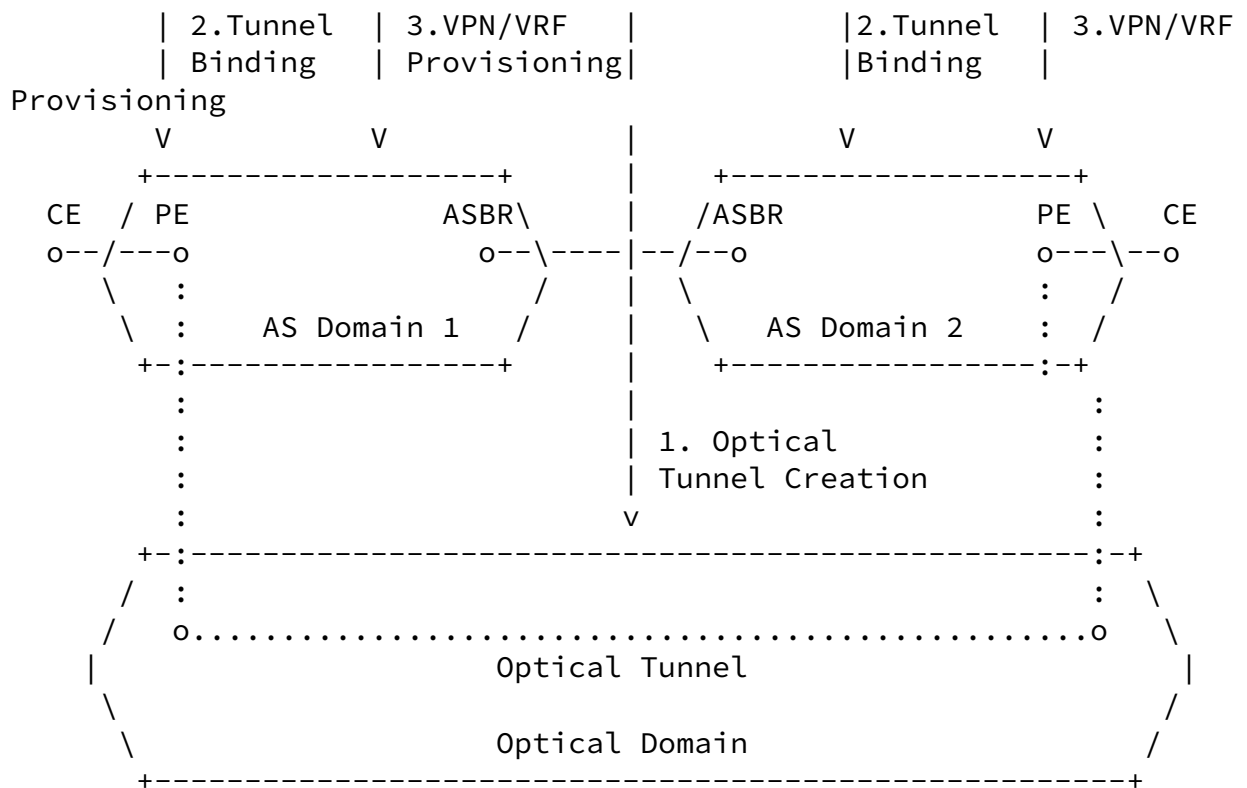


Figure 6 Domain Controller & NE Functions (Isolated Optical Tunnel)

As discussed in 5.2, in case that VPN has requirement for hard-isolated tunnel establishment, the service/network orchestrator will coordinate across IP/MPLS domain controllers and Optical L-MDSC to ensure the creation of a new optical tunnel for the VPN in proper sequence. Figure 6 shows this scenario.

- o The MDSC of the service/network orchestrator requests the L-MDSC to setup and Optical tunnel providing connectivity between the inter-layer interfaces at the ingress and egress PEs and requests the two IP/MPLS domain controllers to setup an inter-domain IP link between these interfaces
- o The MDSC of the service/network orchestrator then should provide the ingress IP/MPLS domain controller with the routing instruction for the VPN so that the ingress IP/MPLS domain controller would help its ingress PE to populate forwarding

table. The packet with the VPN label should be forwarded to the optical interface the MDSC provided.

- o The Ingress Optical Domain PE needs to recognize MPLS-TE label on its ingress interface from IP/MPLS domain PE and encapsulate the MPLS packets of this MPLS-TE Tunnel into the ODU.

[Editor's Note] We assumed that the Optical PE is LSR.]

- o The Egress Optical Domain PE needs to POP the ODU label before sending the packet (with MPLS-TE label kept intact at the top level) to the Egress PE in the IP/MPLS Domain to which the packet is destined.

[Editor's Note] If there are two VPNs having the same destination CE requiring non-shared optical tunnels from each other, we need to explain this case with a need for additional Label to differentiate the VPNs]

5.5. Orchestrator-Controllers-NEs Communication Protocol Flows

This section provides generic communication protocol flows across orchestrator, controllers and NEs in order to facilitate the POI scenarios discussed in [Section 5.3.2](#) for dynamic optical Tunnel establishment. Figure 7 shows the communication flows.

+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
Orchestr.	Optical	Packet	Packet	Ing.PE	Egr.PE
	Ctr.	Ctr-D1	Ctr-D2	D1	D2
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

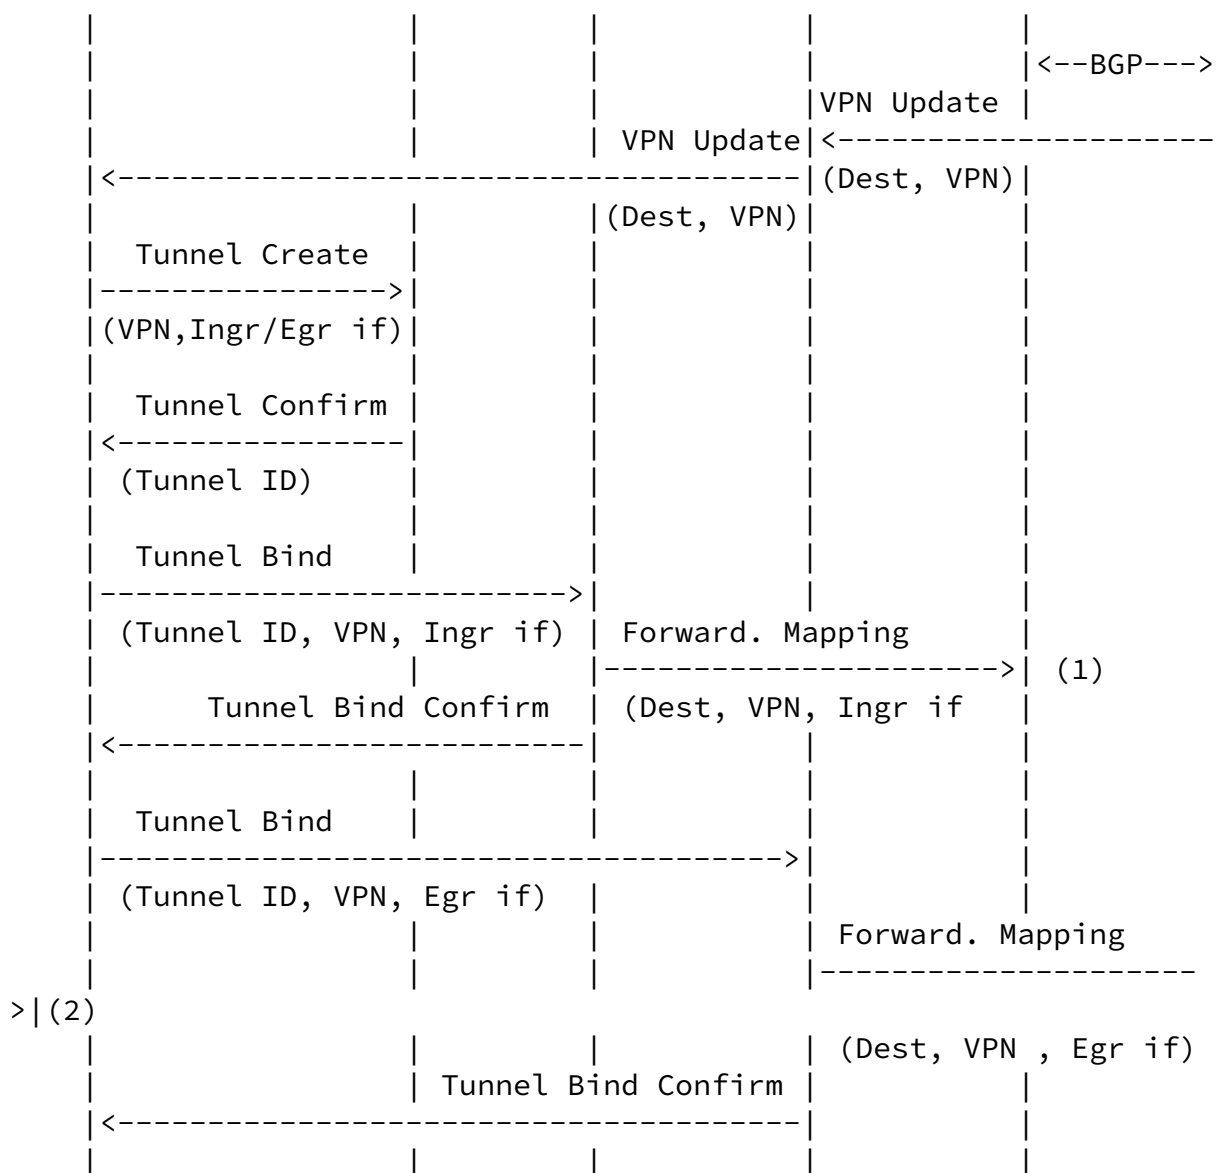


Figure 7 Communication Flows for Optical Tunnel Establishment and binding.

When Domain Packet Controller 1 sends the forwarding mapping information as indicated in (1) in Figure 7, the Ingress PE in Domain 1 will need to provision the VRF forwarding table based on the information it receives. Please see the detailed procedure in [section 5.3.1.2](#). A similar procedure is to be done at the Egress PE in Domain 2.

6. Security Considerations

Several security considerations have been identified and will be discussed in future versions of this document.

7. Operational Considerations

Telemetry data, such as the collection of lower-layer networking health and consideration of network and service performance from POI domain controllers, may be required. These requirements and capabilities will be discussed in future versions of this document.

8. IANA Considerations

This document requires no IANA actions.

9. References

9.1. Normative References

- [RFC7950] Bjorklund, M. et al., "The YANG 1.1 Data Modeling Language", [RFC 7950](#), August 2016.
- [RFC7951] Lhotka, L., "JSON Encoding of Data Modeled with YANG", [RFC 7951](#), August 2016.
- [RFC8040] Bierman, A. et al., "RESTCONF Protocol", [RFC 8040](#), January 2017.
- [RFC8345] Clemm, A., Medved, J. et al., "A Yang Data Model for Network Topologies", [RFC8345](#), March 2018.
- [RFC8346] Clemm, A. et al., "A YANG Data Model for Layer 3 Topologies", [RFC8346](#), March 2018.
- [RFC8453] Ceccarelli, D., Lee, Y. et al., "Framework for Abstraction and Control of TE Networks (ACTN)", [RFC8453](#), August 2018.
- [RFC8525] Bierman, A. et al., "YANG Library", [RFC 8525](#), March 2019.
- [IEEE 802.1AB] IEEE 802.1AB-2016, "IEEE Standard for Local and metropolitan area networks – Station and Media Access Control Connectivity Discovery", March 2016.
- [TE-TOP0] Liu, X. et al., "YANG Data Model for TE Topologies", [draft-ietf-teas-yang-te-topo](#), work in progress.

- [WSON-TOP0] Lee, Y. et al., "A YANG Data Model for WSON (Wavelength Switched Optical Networks)", [draft-ietf-ccamp-wson-yang](#), work in progress.
- [Flexi-TOP0] Lopez de Vergara, J. E. et al., "YANG data model for Flexi-Grid Optical Networks", [draft-ietf-ccamp-flexigrid-yang](#), work in progress.
- [CLIENT-TOP0] Zheng, H. et al., "A YANG Data Model for Client-layer Topology", [draft-zheng-ccamp-client-topo-yang](#), work in progress.
- [L3-TE-TOP0] Liu, X. et al., "YANG Data Model for Layer 3 TE Topologies", [draft-ietf-teas-yang-l3-te-topo](#), work in progress.
- [TE-TUNNEL] Saad, T. et al., "A YANG Data Model for Traffic Engineering Tunnels and Interfaces", [draft-ietf-teas-yang-te](#), work in progress.
- [WSON-TUNNEL] Lee, Y. et al., "A Yang Data Model for WSON Tunnel", [draft-ietf-ccamp-wson-tunnel-model](#), work in progress.
- [Flexi-MC] Lopez de Vergara, J. E. et al., "YANG data model for Flexi-Grid media-channels", [draft-ietf-ccamp-flexigrid-media-channel-yang](#), work in progress.
- [CLIENT-SIGNAL] Zheng, H. et al., "A YANG Data Model for Transport Network Client Signals", [draft-ietf-ccamp-client-signal-yang](#), work in progress.

[9.2](#). Informative References

- [RFC4364] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [RFC4761] K. Kompella, Ed., Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), January 2007.
- [RFC6074] E. Rosen, B. Davie, V. Radoaca, and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", [RFC 6074](#), January 2011.

Internet-Draft

ACTN POI

June 2020

- [RFC6624] K. Kompella, B. Kothari, and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", [RFC 6624](#), May 2012.
- [RFC7209] A. Sajassi, R. Aggarwal, J. Uttaro, N. Bitar, W. Henderickx, and A. Isaac, "Requirements for Ethernet VPN (EVPN)", [RFC 7209](#), May 2014.
- [RFC7432] A. Sajassi, Ed., et al., "BGP MPLS-Based Ethernet VPN", [RFC 7432](#), February 2015.
- [RFC7436] H. Shah, E. Rosen, F. Le Faucheur, and G. Heron, "IP-Only LAN Service (IPLS)", [RFC 7436](#), January 2015.
- [RFC8214] S. Boutros, A. Sajassi, S. Salam, J. Drake, and J. Rabadan, "Virtual Private Wire Service Support in Ethernet VPN", [RFC 8214](#), August 2017.
- [RFC8299] Q. Wu, S. Litkowski, L. Tomotaki, and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", [RFC 8299](#), January 2018.
- [RFC8309] Q. Wu, W. Liu, and A. Farrel, "Service Model Explained", [RFC 8309](#), January 2018.
- [RFC8466] G. Fioccola, ed., "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", [RFC8466](#), October 2018.
- [TNBI] Busi, I., Daniel, K. et al., "Transport Northbound Interface Applicability Statement", [draft-ietf-ccamp-transport-nbi-app-statement](#), work in progress.
- [ACTN-VN] Y. Lee, et al., "A Yang Data Model for ACTN VN Operation", [draft-ietf-teas-actn-vn-yang](#), work in progress.
- [TSM] Y. Lee, et al., "Traffic Engineering and Service Mapping Yang Model", [draft-ietf-teas-te-service-mapping-yang](#), work in progress.
- [ACTN-PM] Y. Lee, et al., "YANG models for VN & TE Performance Monitoring Telemetry and Scaling Intent Autonomics", [draft-lee-teas-actn-pm-telemetry-autonomics](#), work in progress.

progress.

[BGP-L3VPN] D. Jain, et al. "Yang Data Model for BGP/MPLS L3 VPNs",
[draft-ietf-bess-l3vpn-yang](#), work in progress.

Peruzzini et al. Expires December 29, 2020

[Page 33]

Internet-Draft

ACTN POI

June 2020

Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

Some of this analysis work was supported in part by the European Commission funded H2020-ICT-2016-2 METRO-HAUL project (G.A. 761727).

Contributors

Sergio Belotti
Nokia

Email: sergio.belotti@nokia.com

Gabriele Galimberti
Cisco

Email: ggalimbe@cisco.com

Zheng Yanlei
China Unicom

Email: zhengyanlei@chinaunicom.cn

Anton Snitser
Sedona

Email: antons@sedonasys.com

Washington Costa Pereira Correia
TIM Brasil

Email: wcorreia@timbrasil.com.br

Peruzzini et al. Expires December 29, 2020

[Page 34]

Internet-Draft

ACTN POI

June 2020

Michael Scharf
Hochschule Esslingen - University of Applied Sciences

Email: michael.scharf@hs-esslingen.de

Young Lee
Sung Kyun Kwan University

Email: younglee.tx@gmail.com

Jeff Tantsura
Apstra

Email: jefftant.ietf@gmail.com

Authors' Addresses

Fabio Peruzzini
TIM

Email: fabio.peruzzini@telecomitalia.it

Jean-Francois Bouquier
Vodafone

Email: jeff.bouquier@vodafone.com

Italo Busi
Huawei

Email: Italo.busi@huawei.com

Daniel King
Old Dog Consulting

Email: daniel@olddog.co.uk

Peruzzini et al. Expires December 29, 2020

[Page 35]

Internet-Draft

ACTN POI

June 2020

Daniele Ceccarelli
Ericsson

Email: daniele.ceccarelli@ericsson.com

