TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: 12 April 2026

F. Peruzzini
FiberCop
J.-F. Bouquier
Vodafone
I. Busi
Huawei
D. King
Old Dog Consulting
D. Ceccarelli
Cisco

9 October 2025

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

draft-ietf-teas-actn-poi-applicability-16

Abstract

This document explores the applicability of the Abstraction and Control of TE Networks (ACTN) architecture to Packet Optical Integration (POI) within the context of IP/MPLS and optical internetworking. It examines the YANG data models defined by the IETF that enable an ACTN-based deployment architecture and highlights specific scenarios pertinent to Service Providers.

Existing IETF protocols and data models are identified for each multi-technology scenario (packet over optical), particularly emphasising the Multi-Domain Service Coordinator to Provisioning Network Controller Interface (MPI) within the ACTN architecture

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://IETF-TEAS-WG.github.io/actn-poi/draft-ietf-teas-actn-poi-applicability.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ietf-teas-actn-poi-applicability/.

Discussion of this document takes place on the Traffic Engineering Architecture and Signaling Working Group mailing list (mailto:teas@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/teas/. Subscribe at https://www.ietf.org/mailman/listinfo/teas/.

Source for this draft and an issue tracker can be found at https://github.com/IETF-TEAS-WG/actn-poi.

Status of This Memo

This Internet-Draft will expire on 12 April 2026.

Copyright Notice

Table of Contents

1. Introduction

The full automation of management and control for Service Providers' transport networks, spanning IP/MPLS, optical, and microwave technologies, is crucial to addressing customer demands for high-bandwidth applications, such as ultra-fast mobile broadband for 5G and fiber connectivity services. The Abstraction and Control of TE Networks (ACTN) architecture and interfaces enable the automation and efficient operation of complex optical and IP/MPLS networks using standardized interfaces and data models. This approach supports a broad spectrum of network services that can be requested by upper-layer applications, meeting diverse service-level requirements from a network perspective, such as physical diversity, latency, bandwidth, and topology.

Packet Optical Integration (POI) represents an advanced application of traffic engineering. In wide-area networks, packet networks based on the Internet Protocol (IP), often augmented with Multiprotocol Label Switching (MPLS) or Segment Routing (SR), are typically implemented over an optical transport network utilizing Dense Wavelength Division Multiplexing (DWDM), occasionally with an optional Optical Transport Network (OTN) layer.

There are significant technical differences between the packet and optical technologies (e.g., routers versus optical switches) and their associated network engineering and planning approaches (e.g., inter-domain peering optimization in IP networks versus managing physical impairments in DWDM systems or operating on vastly different time scales). Additionally, customer requirements often differ between packet and optical networks, and it is common for Service Providers to use different vendors for each domain. As a result, the operation of these complex packet and optical networks is often siloed, as each technology domain requires specialized skill sets.

As a consequence, in many existing network deployments, packet and optical networks are engineered and operated independently.

This separation is inefficient for several reasons. Firstly, integrating packet and optical networks can significantly reduce capital expenditures (CAPEX) and operational expenditures (OPEX). Secondly, multi-technology topology insights can optimize troubleshooting (e.g., alarm correlation) and enhance network operation (e.g., coordination of maintenance events). Additionally, detailed inventory and planning information can also improve service assurance quality, such as detecting constraint violations or lack of resource diversity. Thirdly, multi-technology traffic engineering enables more efficient use of available network capacity (e.g., coordination of restoration). Furthermore, provisioning workflows can be simplified or automated across layers, facilitating capabilities such as bandwidth-on-demand and streamlined maintenance activities.

The ACTN framework facilitates seamless integration of packet and optical networks across multiple technologies and vendors. This is achieved through separated Provisioning Network Controllers (PNCs) for both packet and optical domains, which hide the complexities of the technical differences between the packet and optical technologies while providing sufficient abstract information that allows the Multi-Domain Service Coordinator (MDSC) to provide multi-layer coordination between packet and optical networks.

This document uses packet-based Traffic Engineered (TE) service examples. These are described as "TE-path" in this document. Unless otherwise stated, these TE services may be instantiated using Resource Reservation Protocol (RSVP) Traffic Engineering (TE)-based or SR -TE-based, forwarding plane mechanisms.

This document outlines key scenarios for Packet Optical Integration (POI) from the perspective of the packet service layer and highlights the necessary coordination between packet and optical layers to enhance POI deployment and operation. These scenarios emphasize multi-domain packet networks functioning as clients of optical networks.

This document analyzes the scenario in which packet networks support multi-domain TE paths. The optical networks may consist of a DWDM network, an OTN network (without a DWDM layer), or a multi-layer OTN/DWDM network. Additionally, DWDM networks can be either fixed-grid or flexible-grid.

Multi-technology and multi-domain scenarios, based on the reference network described in Section 2 and very relevant for Service Providers, are described in Section 4 and Section 5.

For each scenario, existing IETF protocols and data models, identified in Section 3.1 and Section 3.2, are analyzed with a particular focus on the MPI in the ACTN architecture.

For each multi-technology scenario, the document analyzes how to use the interfaces and data models of the ACTN architecture.

A summary of the gaps identified in this analysis is provided in Section 6.

Understanding the degree of standardization and identifying potential gaps are crucial for evaluating the feasibility of integrating packet and optical DWDM domains (with an optional OTN layer) from an end-to-end, multi-vendor service provisioning perspective.

1.1. Terminology

This document uses the ACTN terminology defined in [RFC8453].

In addition, this document uses the following terminology.

Customer service: The end-to-end service from Customer Edge (CE) to CE.

Network service: The Provider Edge (PE) to PE configuration, including both the network service layer (VRFs, RT import/export policies configuration) and the network transport layer (e.g. RSVP-TE Label Switched Paths (LSPs). This includes the configuration (on the PE side) of the interface towards the CE (e.g. VLAN, IP address, routing protocol, etc.).

Technology domain: short for "switching technology domain", defined as "region" in [RFC5212], where the term "region" is applied to (GMPLS) control domains.

PNC Domain: part of the network under the control of a single PNC instance. It is subject to the capabilities of the PNC which technology is controlled.

Port: The physical entity that transmits and receives physical signals.

Interface: A physical or logical entity that transmits and receives traffic.

Link: An association between two interfaces that can exchange traffic directly.

Intra-domain link: a link between two adjacent nodes that belong to the same PNC domain.

Inter-domain link: a link between two adjacent nodes that belong to different PNC domains.

Ethernet link: A link between two Ethernet interfaces.

Single-technology Ethernet link: An Ethernet link between two Ethernet interfaces on physically adjacent IP routers.

Multi-technology Ethernet link: An Ethernet link between two Ethernet interfaces on logically adjacent IP routers, supported by two cross-technology Ethernet links interconnected through an optical tunnel.

Cross-technology Ethernet link: An Ethernet link connecting an Ethernet interface on an IP router to an Ethernet interface on a physically adjacent optical node.

Inter-domain Ethernet link: An Ethernet link between two Ethernet interfaces on physically adjacent IP routers that belong to different P-PNC domains.

Single-technology intra-domain Ethernet link: An Ethernet link between two Ethernet interfaces on physically adjacent IP routers that belong to the same P-PNC domain.

Multi-technology intra-domain Ethernet link: An Ethernet link between two Ethernet interfaces on logically adjacent IP routers within the same P-PNC domain, supported by two cross-technology Ethernet links interconnected through an optical tunnel.

IP link: A link between two IP interfaces.

Inter-domain IP link: An IP link supported by an inter-domain Ethernet link.

Single-technology intra-domain IP link: An IP link supported by a single-technology intra-domain Ethernet link.

Multi-technology intra-domain IP link: An IP link supported by a multi-technology intra-domain Ethernet link.

2. Reference Network Architecture

This document examines various deployment scenarios for Packet and Optical Integration (POI), where the ACTN hierarchy is implemented to manage a multi-technology, multi-domain network comprising two optical domains and two packet domains, as illustrated in Figure 1:

```
                               +----------+
                               |   MDSC   |
                               +-----+----+
                                     |
                   +-----------+-----+------+-----------+
                   |           |            |           |
              +----+----+ +----+----+  +----+----+ +----+----+
              | P-PNC 1 | | O-PNC 1 |  | O-PNC 2 | | P-PNC 2 |
              +----+----+ +----+----+  +----+----+ +----+----+
                   |           |            |           |
                   |           \            /           |
         +-------------------+  \          /  +-------------------+
    CE1 / PE1             BR1 \  |        /  / BR2             PE2 \ CE2
    o--/---o               o---\-|-------|--/---o               o---\--o
       \   :               :   / |       |  \   :               :   /
        \  : PKT domain 1  :  /  |       |   \  : PKT domain 2  :  /
         +-:---------------:-+   |       |    +-:---------------:--+
           :               :     |       |      :               :
           :               :     |       |      :               :
         +-:---------------:------+     +-------:---------------:--+
        /  :               :       \   /        :               :   \
       /   o...............o        \ /         o...............o    \
       \     optical domain 1       / \       optical domain 2       /
        \                          /   \                            /
         +------------------------+     +--------------------------+
```

Figure 1: Reference Network

The ACTN architecture, as defined in [RFC8453], is utilized to manage this multi-technology, multi-domain network. In this topology, each Packet PNC (P-PNC) is responsible for controlling its respective packet domain, while each Optical PNC (O-PNC) is tasked with managing its optical domain. The packet domains controlled by the P-PNCs can represent Autonomous Systems (ASes), as defined in [RFC1930], or Interior Gateway Protocol (IGP) areas within the same operator network.

The IP routers between the packet domains can be either AS Boundary Routers (ASBR) or Area Border Router (ABR): in this document, the generic term Border Router (BR) is used to represent either an ASBR or an ABR.

The Multi-Domain Service Coordinator (MDSC) is responsible for orchestrating the entire multi-domain, multi-technology network, encompassing both packet and optical domains. A standardized interface, the Multi-Domain Service Coordinator to Provisioning Network Controller Interface (MPI), enables the MDSC to interact with various Provisioning Network Controllers (O-PNCs and P-PNCs).

The MPI interface provides the MDSC with an abstracted topology, concealing technology-specific details of the network and selectively hiding topology information based on the chosen abstraction policy. The level of abstraction is determined by the configuration parameters of the P-PNC and O-PNC, such as offering potential connectivity information between any Provider Edge (PE) and Border Router (BR) within a packet network.

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

* The domain boundaries of the packet and optical domains are congruent. In other words, each optical domain exclusively supports connectivity between IP routers within a single packet domain.

* There are no physical links directly connecting optical domains. Inter-domain physical links exist only under the following conditions:

    - between packet domains (i.e., between BRs belonging to different packet domains): these links are called inter-domain Ethernet or IP links within this document;

    - between packet and optical domains (i.e., between routers and optical nodes): these links are called cross-technology Ethernet links within this document;

    - between customer sites and the packet network (i.e., between CE devices and PE routers): these links are called access links within this document.

* All the physical interfaces at inter-domain links are Ethernet physical interfaces.

Scenarios using coherent optical interfaces on the IP routers are outside the scope of this document.

This document analyzes scenarios in which all multi-technology IP links supported by the optical network are intra-domain (intra-AS/intra-area), such as PE-BR, PE-P, BR-P, and P-P IP links. Consequently, inter-domain IP links are always single-technology connections, supported by single-technology Ethernet links between physically adjacent IP routers.

As described in [RFC7424], in order to increase the bandwidth between two adjacent routers, multiple Ethernet links can be setup between adjacent routers using either Link Aggregation Groups (LAGs) [IEEE_802.1AX] or Equal Cost Multi-Path (ECMP) [RFC2991].

Therefore, if inter-domain links between optical domains exist, they would be utilized to support multi-domain optical services, which fall outside the scope of this document.

The optical nodes within the optical domains can be either:

* WDM nodes, as defined in [I-D.ietf-ccamp-optical-impairment-topology-yang], with an integrated ROADM function with or without integrated optical transponders;

* OTN nodes, with integrated an OTN cross-connect function and with or without integrated ROADM functions or optical transponders.

2.1. Multi-domain Service Coordinator (MDSC) functions

The MDSC in Figure 1 is responsible for coordinating multiple packet and optical domains in a multi-domain, multi-technology environment. It facilitates multi-layer and multi-domain L2/L3 VPN network services as requested by the OSS/Orchestration layer.

From an implementation perspective, the functions associated with MDSC described in [RFC8453] may be grouped differently.

1. The service-related and network-related functions are combined into a single, monolithic implementation. This implementation manages end-customer service requests received through the Customer MDSC Interface (CMI) and adapts the corresponding network models. An example of this architecture is illustrated in Figure 2 of [RFC8453].

2. An implementation may opt to separate the service-related and network-related functions into distinct functional entities, as outlined in [RFC8309] and Section 4.2 of [RFC8453]. In this approach, the MDSC is decomposed into a top-level Service Orchestrator, which interfaces with the customer through the Customer MDSC Interface (CMI), and a Network Orchestrator, which interfaces southbound with the PNCs. The interface between the Service Orchestrator and the Network Orchestrator is not specified in [RFC8453].

3. Another implementation may choose to split the MDSC functions into a "higher-level MDSC" (MDSC-H) and "lower-level MDSCs" (MDSC-Ls). The MDSC-H is responsible for multi-technology coordination across packet and optical domains, while the MDSC-Ls handle domain-specific coordination. Specifically, an Optical MDSC-L manages multi-domain coordination between the O-PNCs, and a Packet MDSC-L manages multi-domain coordination between the P-PNCs. This approach is illustrated, for example, in Figure 9 of [RFC8453].

4. An alternative implementation may choose to integrate the MDSC and P-PNC functions in a single entity.

In current service provider network deployments, the MDSC’s Northbound Interface (NBI) typically connects to an OSS/Orchestration layer rather than a CNC. In this scenario, the MDSC is limited to performing Network Orchestration functions, as described in [RFC8309] (point 2 above). Consequently, the MDSC handles network service requests received from the OSS and/or Orchestration.

The functionality of the OSS and/or Orchestration layer, as well as its interface with the MDSC, is typically operator-specific and falls outside the scope of this draft. Therefore, this document assumes that the OSS and/or Orchestration layer requests the MDSC to provision L2/L3 VPN network services through mechanisms not covered in this document.

There are two prominent workflow cases when the MDSC multi-technology coordination is initiated:

* Initiated by request from the OSS and/or Orchestration layer to setup L2/L3 VPN network services that require multi-layer/multi-domain coordination;

* The MDSC initiates these workflows to perform multi-layer and multi-domain optimizations and/or maintenance activities (e.g., rerouting LSPs and their associated services when a resource, such as a fiber, is placed in maintenance mode during a maintenance window). Unlike service fulfilment, these workflows are not triggered by a network service provisioning request from the OSS or Orchestration layer.

The latter workflow cases are outside the scope of this document.

This document examines use cases in which multi-layer coordination is initiated by a network service request from the OSS and/or Orchestration layer.

2.1.1. Multi-domain L2/L3 VPN Network Services

Figure 2 and Figure 3 provide an example of a hub & spoke multi-domain L2/L3 VPN with three PEs where the hub PE (PE13) and one spoke PE (PE14) are within the same packet domain, and the other spoke PE (PE23) is within a different packet domain.

```
        ------
       | CE13 |    Packet Domain 1              Packet Domain 2
        ------ ____________________            __________________
        ( |                         )         (                  )
       (  | PE13     P15       BR11  )       (  BR21       P24     )
      (   |____         ___       ____ )      ( ____      ___       )
     (    /    \ _ _ _ /   \ _ _ /    \________/    \    /   \     )
    (     \____/       \___/     \___ /        \____/    \_ _/     )
   (   PE14  :\_ _               /      )  (    /  :      : \__     )
   (    ____  :   \__ P16    ___/      )  (  __/_             _\__  )
    (  /    \  :  /   \- - -/    \__________/    \ :_ _ _ :_ /    \  )
    (  \____/     \___/     \____/     )  ( \____/           \____/ )
      (  / :   :    :         :  BR12  )   (   :    :     :     |  )
       (/                              )   ( BR22           PE23|   )
    ------ :   :    :         :       )      ( :     :    :     |  )
   | CE14 | (__ ____ _________ _____)           (_____ ___ _ ------
    ------ :   :    :         :                :      :   : | CE23 |
                                                             ------
           :   :    :         :                :      :   :
          _ ___ ____ _________ ________         ______ ___ _______
         ( :   :    :         :        )       :      :   :       )
        (      ____  :      ____        )     (      ____  .. ..   )
       (   :  /    \_ _ _ _/    \ NE12   )   ( :    /    \ _    :   )
      (  NE11 \____/ :     \____/         )  ( NE21 \____/   \     )
      (    :  /    \    _ _ /  \          )  ( :     /        \ :   )
      (   ___/      \:_|        \____    )  (   .___/         _\__  )
      (  /    \_ _ /    \ _ _ _ /    \   )  (   /    \ _ _ _ /    \  )
       ( \____/    \____/       \____/  )    (  \____/       \____/  )
        ( NE13      NE14         NE15   )     (  NE22         NE23  )
         (_____________________________)       (___________________)

                Optical Domain 1                  Optical Domain 2


          _____  = Inter-domain links
          .. ..  = Cross-layer links
          _ _ _  = Intra-domain links
```

Figure 2: Multi-domain VPN topology example

```
        ------
       | CE13 |    Packet Domain 1              Packet Domain 2
        ------ ____________________            _________________
        ( |                         )         (                 )
       (  | PE13     P15       BR11  )       (  BR21       P24    )
      (   |____         ___       ____ )      ( ____     ___       )
     (    / H  \       /   \     /    \________/..  \   / ..\ ..  )
    (     \____/.....  \___/     \___ / .. .. ..___:/   \___/   : )
   (   PE14  :      :              .. .. )  (             :        )
   (    ____  :    _:_ P16   ____ :     )  ( ____  :          __:_ )
    (  / S  \  :  / ..\     /   ..__________/    \        :  /  S \ )
    (  \____/     \__:/     \____/     )  ( \____/ :         \____/ )
      (  / :   :     :          :BR12  )   (              :     |  )
       (/  :         :                 )   ( BR22  :        PE23|   )
    ------ :   :     :          :     )      (            :     | )
   | CE14 |:(__ _____:__________ ___)           (__:______ __ ------
    ------ :   :      :         :                         :  | CE23 |
           :           :                           :          ------
           :   :       :        :                         :
          _:___________:________ ______         ___:______ _______
         ( :   :       :        :      )       (          : .. .. )
        (  :   ____    :    ____        )     (     :____          )
       (   :  / .. \.. : ../ .. \ NE12   )   (      /..  \      :   )
      (  NE11 \____/   :   \____/         )  ( NE21 \__:_/          )
      (    :           :                  )  (                  :  )
      (   _:__      ___:         ____    )  (    ____  : ..  ____  )
      (  / :..\..../...:\       /    \   )  (   /    \      /.. :\  )
       ( \____/    \____/       \____/  )    (  \____/      \____/  )
        ( NE13      NE14         NE15   )     (  NE22        NE23  )
         (_____________________________)       (__________________)

                Optical Domain 1                  Optical Domain 2


           H / S = Hub VRF / Spoke VRF

          .....  = Intra-domain TE Path 1 {PE13, P16, NE14, NE13, PE14}
          .. ..  = Inter-domain TE Path 2 {PE13, NE11, NE12, BR12,
                   BR11, BR21, NE21, NE23, P24, PE23}
```

Figure 3: Multi-domain VPN TE paths example

There are many options to implement multi-domain L2/L3 VPNs, including:

1. BGP-Labeled Unicast (BGP-LU) ([RFC8277])

2. Inter-domain RSVP-TE

3. Inter-domain SR-TE

This document explores inter-domain TE options where the TE tunnel model, as defined in [I-D.ietf-teas-yang-te], applies at the MPI for both intra-domain and inter-domain TE configurations. The assessment of alternative options is beyond the scope of this draft.

It is also assumed that:

* the bandwidth of each intra-domain TE path is managed by its respective P-PNC;

* technology-specific mechanisms are employed for inter-domain TE path stitching. In the case of inter-domain SR-TE, a Segment Identifier (SID) is used in Segment Routing (SR) to define a segment (a portion of the path) within a network. A binding SID, a special type of SID, acts as a reference to a precomputed SR policy or path.

* each packet domain in Figure 2 employs technology-specific local protection mechanisms, such as Fast Reroute (FRR) for MPLS-TE or Topology Independent Loop-Free Alternate (TI-LFA) for SR-TE. These mechanisms operate with an awareness of the multi-technology TE path properties, such as the Shared Risk Link Group (SRLG) path properties defined in [RFC8001].

For inter-domain TE paths, it is assumed that each packet domain in Figure 2 and Figure 3 employs the same TE technology. The stitching between two domains is achieved using inter-domain TE mechanisms.

In this scenario, a key function of the MDSC is to identify the multi-domain and multi-layer TE paths for carrying L2/L3 VPN traffic between PEs in different packet domains. The MDSC then relays this information to the P-PNCs to ensure that the forwarding tables of the PEs (e.g., VRF) are correctly configured, allowing the L2/L3 VPN traffic to be routed over the designated multi-domain and multi-layer TE paths.

The selection of the TE path should consider both the TE requirements and the binding requirements of the L2/L3 VPN network service.

In general, the binding requirements for a network service (e.g., L2/L3 VPN) can be categorized into three main cases:

1. The customer is asking for VPN isolation to dynamically create and bind tunnels to the service so that they are not shared by other services (e.g. VPN).

The level of isolation can be different:

a. Hard isolation with deterministic latency implies that the L2/L3 VPN requires a set of dedicated TE tunnels. These tunnels neither share resources with other services, nor compete for bandwidth with other tunnels, ensuring deterministic latency performance.

b. Hard isolation but without deterministic characteristics

c. Soft isolation means the tunnels associated with L2/L3 VPN are dedicated to that but can compete for bandwidth with other tunnels.

2. The customer does not require isolation and may request a VPN service where the associated tunnels are shared across multiple VPNs.

For each TE path required to support the L2/L3 VPN network service, it is possible that:

1. A TE path that meets the TE and binding requirements already exists in the network.

2. An existing TE path could be modified (e.g., through bandwidth increase) to meet the TE and binding requirements:

   a. The TE path characteristics can be modified only in the packet layer.

   b. One or more new underlay optical tunnels need to be setup to support the requested changes of the overlay TE paths (multi-layer coordination is required).

3. A new TE path needs to be setup to meet the TE and binding requirements:

   a. The new TE path reuses existing underlay optical tunnels;

   b. One or more new underlay optical tunnels need to be setup to support the setup of the new TE path (multi-layer coordination is required).

This document examines scenarios in which a single TE path is used to carry VPN traffic between PEs. Scenarios involving multiple parallel TE paths for load-balancing VPN traffic between PEs are possible but are beyond the scope of this document.

2.1.2. Multi-domain and Multi-layer Path Computation

When establishing a new TE path, the MDSC is responsible for coordinating the path computation across multiple layers and domains.

Based on the MDSC’s knowledge of the underlying network topology and configuration, there are three possible approaches for multi-layer and multi-domain path computation:

1. Full Summarization: In this approach, the MDSC maintains an abstracted TE topology view of all its packet and optical underlying domains.

   In this case, the MDSC lacks sufficient TE topology information to perform multi-layer/multi-domain path computation. It delegates the P-PNCs and O-PNCs to compute local paths within their respective domains, then uses the returned information to compute the optimal multi-domain/multi-layer path.

   This approach presents an issue for the P-PNC, as it lacks the ability to perform single-domain/multi-layer path computation. It cannot retrieve topology information from the O-PNCs or delegate optical path computation to the O-PNCs. A possible solution is to include a CNC function within the P-PNC to request the MDSC for multi-domain optical path computation, as shown in Figure 10 of [RFC8453].

   Another solution could involve relying on the MDSC recursive hierarchy, as defined in Section 4.1 of [RFC8453], where each IP and optical domain pair has a "lower-level MDSC" (MDSC-L) for multi-layer correlation, and a "higher-level MDSC" (MDSC-H) for multi-domain coordination.

   In this case, the MDSC-H obtains an abstract view of the underlying multi-layer domain topologies from its MDSC-L. Each MDSC-L gets the full IP domain topology from the P-PNC and an abstracted view of the optical domain topology from its O-PNC. Topology abstraction occurs at the MPIs between MDSC-L and O-PNC, as well as between MDSC-L and MDSC-H.

2. Partial summarization: In this approach, the MDSC has complete visibility of the TE topology of the packet network domains and an abstracted view of the TE topology of the optical network domains.

   The MDSC can then only perform multi-domain/single-layer path computation for the packet layer, where the path can be computed optimally for the two packet domains.

The MDSC must still need to delegate the O-PNCs to perform local path computation within their domains. It uses the information from the O-PNCs and its TE topology view of the multi-domain packet layer to perform multi-layer/multi-domain path computation.

3. Full knowledge: In this approach, the MDSC has a complete and enough detailed view of the TE topology of all the network domains (both optical and packet).

   In such a case, the MDSC has all the information needed to perform multi-domain/multi-layer path computation without relying on PNCs.

   This approach, however, may present scalability issues. As discussed in Section 2.2 of [I-D.ietf-teas-yang-path-computation], performing path computation for optical networks in the MDSC is particularly challenging, as optimal paths also depend on vendor-specific optical attributes, which may vary across domains if provided by different vendors.

This document examines scenarios where the MDSC adopts the partial summarization approach to enable multi-domain and multi-layer path computation.

Typically, O-PNCs are responsible for optical path computation within their respective domains. When setting up a network service, they must consider connection requirements such as bandwidth, amplification, wavelength continuity, and non-linear impairments that may impact the network service path.

The methods and types of path requirements and impairments, such as those detailed in [I-D.ietf-ccamp-optical-impairment-topology-yang], used by the O-PNC for optical path computation, are not exposed at the MPI and therefore are out of scope for this document.

2.2. IP/MPLS Domain Controller and IP router Functions

Each packet domain in Figure 1, corresponding to either an IGP area or an Autonomous System (AS) within the same operator network, is controlled by a packet domain controller (P-PNC).

P-PNCs are responsible for establishing TE paths between any two PEs or BRs within their controlled domains, as requested by the MDSC. They also provide topology information to the MDSC to enable efficient network coordination.

For example, in inter-domain SR-TE, setting up a bidirectional SR-TE path from PE13 in Domain 1 to PE23 in Domain 2, as shown in Figure 3, requires the MDSC to coordinate the following actions:

* P-PNC1: Push a SID list to PE13, including the Binding SID associated with the SR-TE path in Domain 2, with PE23 as the target destination (forward direction).

* P-PNC2: Push a SID list to PE23, including the Binding SID associated with the SR-TE path in Domain 1, with PE13 as the target destination (reverse direction).

With reference to Figure 4, P-PNCs are responsible for the following:

1. To expose to MDSC their respective detailed TE topology

2. To perform single-layer, single-domain local TE path computation, when requested by the MDSC, between two PEs (for single-domain end-to-end TE path) or between PEs and BRs for an inter-domain TE path selected by the MDSC.

3. To configure the routers in their respective domain to setup a TE path;

4. To configure the VRF and PE-CE interfaces (Service access points) for the intra-domain and inter-domain network services requested by the MDSC.

```
        +------------------+            +------------------+
        |                  |            |                  |
        |      P-PNC1      |            |      P-PNC2      |
        |                  |            |                  |
        +--|-----------|---+            +--|-----------|---+
           | 1.TE      | 2.VPN           | 1.TE      | 2.VPN
           | Path      | Provisioning    | Path      | Provisioning
           | Config    |                 | Config    |
           V           V                 V           V
         +---------------------+         +---------------------+
    CE  / PE     TE path 1    BR\       / BR     TE path 2   PE \  CE
    o--/---o..................o--\-----/--o..................o---\--o
       \                         /     \                         /
        \        Domain 1      /       \       Domain 2      /
         +---------------------+         +---------------------+

                            End-to-end TE path
           <------------------------------------------------->
```

Figure 4: Domain Controller & node Functions

When requesting a new TE path setup, the MDSC provides the P-PNCs with the explicit path to be created or modified. In other words, the MDSC communicates the complete list of nodes involved in the path (strict mode). The P-PNC is then responsible for setting up the explicit TE path. For example:

* with SR-TE, the P-PNC pushes to headend PE or BR the list of SIDs to create the explicit SR-TE path, provided by the MDSC;

* with RSVP-TE, the P-PNC requests the headend PE or BR to start signaling the explicit RSVP-TE path, provided by the MDSC.

To scale in large SR-TE packet domains, the MDSC can provide the P-PNC with a loose path and per-domain TE constraints. The P-PNC can then select the complete path within its domain.

In this case, the P-PNC must signal back to the MDSC which path it has chosen, allowing the MDSC to track relevant resource utilization.

From the Figure 3 example, the TE path requested by the MDSC touches PE13 - P16 - BR12 - BR21 - PE23. P-PNC2 is aware of two paths with the same topology metric, e.g., BR21 - P24 - PE23 and BR21 - BR22 - PE23, but with different loads. It may prefer to steer traffic on the latter as it is less loaded.

For the purposes of this document, it is assumed that the MDSC always provides the explicit list of all hops to the P-PNCs to set up or modify the TE path.

2.3. Optical Domain Controller and NE Functions

The optical network provides underlay connectivity services to IP/MPLS networks. The packet and optical multi-layer coordination is handled by the MDSC, as shown in Figure 1.

The O-PNC is responsible to:

* Provide the MDSC with an abstract TE topology view of its underlying optical network resources;

* perform single-domain local path computation when requested by the MDSC;

* Perform optical tunnel set up when requested by the MDSC.

The mechanisms used by the O-PNC to perform intra-domain topology discovery and path setup are typically vendor-specific and outside the scope of this document.

Depending on the optical network type, TE topology abstraction, path computation, and path setup can be single-layer (either OTN or DWDM) or multi-layer OTN/DWDM. In the latter case, multi-layer coordination between the OTN and DWDM layers is handled by the O-PNC.

3. Interface Protocols and YANG Data Models for the MPIs

This section describes general assumptions applicable to all MPI interfaces between each PNC (Optical or Packet) and the MDSC, to support the scenarios discussed in this document.

3.1. RESTCONF Protocol at the MPIs

The RESTCONF protocol, as defined in [RFC8040], using the JSON representation from [RFC7951], is assumed to be used at these interfaces. Additionally, extensions to RESTCONF, as defined in [RFC8527], to comply with the Network Management Datastore Architecture (NMDA) from [RFC8342], are assumed to be used at these MPI and MDSC NBI interfaces.

3.2. YANG Data Models at the MPIs

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

This section describes the YANG data models applicable to the Packet and Optical MPIs. Some of these YANG data models may be optional, depending on the specific network configuration detailed in Section 4 and Section 5.

3.2.1. Common YANG Data Models at the MPIs

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.

Both Optical and Packet PNCs can use the following common topology YANG data models at the MPI:

* The Base Network Model, defined in the "ietf-network" YANG module of [RFC8345];

* The Base Network Topology Model, defined in the "ietf-network-topology" YANG module of [RFC8345], which augments the Base Network Model;

* The TE Topology Model, defined in the "ietf-te-topology" YANG module of [RFC8795], which augments the Base Network Topology Model.

Optical and Packet PNCs can use the common TE Tunnel Model, defined in the "ietf-te" YANG module of [I-D.ietf-teas-yang-te], at the MPI.

All common YANG data models are generic and augmented by technology-specific YANG modules, as described in the following sections.

Both Optical and Packet PNCs can also use the Ethernet Topology Model, defined in the "ietf-eth-te-topology" YANG module of [I-D.ietf-ccamp-eth-client-te-topo-yang], which augments the TE Topology Model with Ethernet technology-specific information.

Both Optical and Packet PNCs can use the following common notifications YANG data models at the MPI:

* Dynamic Subscription to YANG Events and Datastores over RESTCONF as defined in [RFC8650];

* Subscription to YANG Notifications for Datastores updates as defined in [RFC8641].

PNCs and MDSCs comply with subscription requirements as stated in [RFC7923].

3.2.2. YANG models at the Optical MPIs

The Optical PNC can use the following technology-specific topology YANG data models, which augment the generic TE Topology Model:

* The WSON Topology Model, defined in the "ietf-wson-topology" YANG module of [RFC9094];

* the Flexi-grid Topology Model, defined in the "ietf-flexi-grid-topology" YANG module of [I-D.ietf-ccamp-flexigrid-yang];

* the OTN Topology Model, as defined in the "ietf-otn-topology" YANG module of [I-D.ietf-ccamp-otn-topo-yang].

The optical PNC can use the following technology-specific tunnel YANG data models, which augments the generic TE Tunnel Model:

* The WDM Tunnel Model, defined in the "ietf-wdm-tunnel" YANG module of [I-D.ietf-ccamp-wdm-tunnel-yang];

* the OTN Tunnel Model, defined in the "ietf-otn-tunnel" YANG module of [I-D.ietf-ccamp-otn-tunnel-model].

The optical PNC can use the generic Path Computation YANG RPC, defined in the "ietf-te-path-computation" YANG module of [I-D.ietf-teas-yang-path-computation].

Note that technology-specific augmentations of the generic path computation RPC for WSON, Flexi-grid, and OTN path computation RPCs have been identified as a gap.

The optical PNC can use the following client signal YANG data models:

* the CBR Client Signal Model, defined in the "ietf-trans-client-service" YANG module of [I-D.ietf-ccamp-client-signal-yang];

* the Ethernet Client Signal Model, defined in the "ietf-eth-tran-service" YANG module of [I-D.ietf-ccamp-client-signal-yang].

3.2.3. YANG data models at the Packet MPIs

The Packet PNC can use the following technology-specific topology YANG data models:

* The L3 Topology Model, defined in the "ietf-l3-unicast-topology" YANG module of [RFC8346], which augments the Base Network Topology Model;

* the Packet TE Topology Mode, defined in the "ietf-te-topology-packet" YANG module of [I-D.ietf-teas-yang-l3-te-topo], which augments the generic TE Topology Model;

* The MPLS-TE Topology Model, defined in the "ietf-te-mpls-topology" YANG module of [I-D.ietf-teas-yang-te-mpls-topology], which augments the TE Packet Topology Model with or without the L3 TE Topology Model, defined in "ietf-l3-te-topology" YANG module of [I-D.ietf-teas-yang-l3-te-topo];

* the SR Topology Model, defined in the "ietf-sr-mpls-topology" YANG module of [I-D.ietf-teas-yang-sr-te-topo].

The Packet PNC can use the following technology-specific tunnel YANG data models, which augments the generic TE Tunnel Model:

* The MPLS-TE Tunnel Model, defined in the "ietf-te-mpls" YANG modules of [I-D.ietf-teas-yang-te-mpls];

* the SR-TE Tunnel Model which is to be defined as described in Section 6.

The packet PNC can use the following network service YANG data models:

* L3VPN Network Model (L3NM), defined in the "ietf-l3vpn-ntw" YANG module of [RFC9182];

* L3NM TE Service Mapping, defined in the "ietf-l3nm-te-service-mapping" YANG module of [I-D.ietf-teas-te-service-mapping-yang];

* L2VPN Network Model (L2NM), defined in the "ietf-l2vpn-ntw" YANG module of [RFC9291];

* L2NM TE Service Mapping, defined in the "ietf-l2nm-te-service-mapping" YANG module of [I-D.ietf-teas-te-service-mapping-yang].

3.3. Path Computation Element Protocol (PCEP)

[RFC8637] examines the applicability of a Path Computation Element (PCE) [RFC5440] and PCE Communication Protocol (PCEP) to the ACTN framework. It further describes how the PCE architecture applies to ACTN and lists the PCEP extensions needed to use PCEP as an ACTN interface. The stateful PCE [RFC8231], PCE-Initiation [RFC8281], stateful Hierarchical PCE (H-PCE) [RFC8751], and PCE as a central controller (PCECC) [RFC8283] are key extensions enabling the use of PCE/PCEP for ACTN.

Since PCEP supports path computation in both packet and optical networks, it is well-suited for inter-layer path computation. [RFC5623] describes a framework for applying the PCE-based architecture to interlayer (G)MPLS traffic engineering. Furthermore, section 6.1 of [RFC8751] outlines H-PCE applicability for inter-layer or POI.

[RFC8637] lists various PCEP extensions that apply to ACTN. It also lists the PCEP extension for the optical network and POI.

Note that PCEP can be used in conjunction with the YANG data models described in the rest of this document. Depending on whether ACTN is deployed in a greenfield or brownfield, two options are possible:

1. The MDSC uses a single RESTCONF/YANG interface to each PNC to discover all TE information and request TE tunnels. It may perform full multi-layer path computation or delegate path computation to the underlying PNCs.

This approach is desirable for operators from a multi-vendor integration perspective as it is simple. We need only one type of interface (RESTCONF) and use the relevant YANG data models depending on the operator use case considered. The benefits of having only one protocol for the MPI between MDSC and PNC have already been highlighted in [I-D.ietf-teas-yang-path-computation].

2. The MDSC uses the RESTCONF/YANG interface towards each PNC to discover all the TE information and requests the creation of TE tunnels. However, it uses PCEP for hierarchical path computation.

   As mentioned in Option 1, from an operator perspective, this option can add integration complexity to have two protocols instead of one unless the RESTCONF/YANG interface is added to an existing PCEP deployment (brownfield scenario).

Section 4 and Section 5 of this draft analyze the case where a single RESTCONF/YANG interface is deployed at the MPI (i.e., option 1 above).

4. Inventory, Service and Network Topology Discovery

In this scenario, the MSDC needs to discover the underlying PNCs:

* the network topology, at both optical and IP layers, in terms of nodes and links, including the access links, inter-domain IP links as well as cross-technology Ethernet links;

* the optical tunnels supporting multi-technology intra-domain IP links;

* both intra-domain and inter-domain L2/L3 VPN network services deployed within the network;

* the TE paths supporting those L2/L3 VPN network services;

* the hardware inventory information of IP and optical equipment.

The O-PNC and P-PNC could discover and report the hardware network inventory information of the equipment used by the different management layers. In the context of POI, the inventory information of IP and optical equipment can complement the topology views and facilitate the packet/optical multi-layer view, e.g., by providing a mapping between the lowest-level link termination points (LTPs) in the topology view and corresponding ports in the network inventory view.

The MDSC could also discover the entire network inventory information of both IP and optical equipment and correlate this information with the links reported in the network topology.

Reporting the entire inventory and detailed topology information of packet and optical networks to the MDSC may present scalability issues as a potential drawback. The analysis of the scalability of this approach and mechanisms to address potential issues is outside the scope of this document.

Each PNC provides the MDSC the topology view of the domain it controls, as described in Section 4.1 and Section 4.3. The MDSC uses this information to discover the complete topology view of the multi-layer multi-domain networks it controls.

The MDSC should also maintain up-to-date inventory, service and network topology databases of IP and optical layers through IETF notifications through MPI with the PNCs when any network inventory/topology/service change occurs.

It should also be possible to correlate information from IP and optical layers (e.g., which port, lambda/OTSi, and direction are used by a specific IP service on the WDM node).

In particular, for the cross-technology Ethernet links, it is key for MDSC to automatically correlate the information from the PNC network databases about the physical ports from the routers (single link or bundle links for LAG) to client ports in the ROADM.

The analysis of multi-layer fault management is outside the scope of this document. However, the discovered information should be sufficient for the MDSC to correlate optical and IP layer alarms to speed-up troubleshooting easily.

Alarms and event notifications are required between MDSC and PNCs so that any network changes are reported almost in real-time to the MDSC (e.g., node or link failure). As specified in [RFC7923], MDSC must subscribe to specific objects from PNC YANG datastores for notifications.

4.1. Optical Topology Discovery

The WSON Topology Model and the Flexi-grid Topology model can be used to report the DWDM network topology (e.g., WDM nodes and OMS links), depending on whether the DWDM optical network is based on fixed-grid or flexible-grid or a mix of fixed-grid and flexible-grid.

It is worth noting that, as described in Appendix I of [ITU-T_G.694.1], a fixed-grid can also be described as a flexible grid with constraints: for example, a 50GHz fixed-grid can be described as a flexible-grid which supports only m=4 and values of n which are only multiplier of 8.

As a consequence:

* A flexible-grid DWDM network topology can only be reported using the Flexi-grid Topology model;

* A fixed-grid DWDM network topology, can be reported using either the WSON Topology model or the Flexi-grid Topology model;

* A mixed fixed and flexible grid DWDM network topology can be reported using either the Flexi-grid Topology model or both WSON and Flexi-grid topology models.

Clarifying how both WSON and Flexi-grid topology models could be used together (e.g., through multi-inheritance as described in [I-D.ietf-teas-te-topology-profiles]) has been identified as a gap.

The OTN Topology Model is used to report the OTN network topology (e.g., OTN switching nodes and links), when the OTN switching layer is deployed within the optical domain.

To allow the MDSC to discover the complete multi-layer and multi-domain network topology and to correlate it with the hardware inventory information, the O-PNCs report an abstract optical network topology where:

* one TE node is reported for each optical node deployed within the optical network domain; and

* one TE link is reported for each OMS link and, optionally, for each OTN link.

Since the MDSC delegates optical path computation to its underlay O-PNCs, the following information can be abstracted and not reported at the MPI:

* the optical parameters required for optical path computation, such as those detailed in [I-D.ietf-ccamp-optical-impairment-topology-yang];

* the underlay OTS links and ILAs of OMS links;

* the physical connectivity between the optical transponders and the ROADMs.

The OTN Topology Model also reports the CBR client LTPs that terminates the cross-technology Ethernet links: once CBR client LTP is reported for each CBR or multi-function client interface on the optical nodes (see sections 4.4 and 5.1 of [I-D.ietf-ccamp-transport-nbi-app-statement] for the description of multi-function client interfaces).

The Ethernet Topology Model reports the Ethernet client LTPs that terminate the cross-technology Ethernet links: one Ethernet client LTP is reported for each Ethernet or multi-function client interface on the optical nodes.

The optical transponders and, optionally, the OTN access cards, are abstracted at MPI by the O-PNC as Trail Termination Points (TTPs), defined in [RFC8795], within the optical network topology. This abstraction is valid independently of the fact that optical transponders are physically integrated within the same WDM node or are physically located on a device external to the WDM node since it both cases the optical transponders and the WDM node are under the control of the same O-PNC and abstracted as a single WDM TE Node at the O-MPI.

The association between the Ethernet or CBR client LTPs terminating the Ethernet cross-technology Ethernet links and the optical TTPs is reported using the Inter Layer Lock-id (ILL) identifiers, defined in [RFC8795].

For example, with a reference to Figure 5, the ILL values X and Y are used to associate the client LTPs (7-0) in NE11 and (8-0) in NE12 with the corresponding optical TTPs (7) in NE11 and (8) in NE12, respectively.

```
         +----------------------------------------------------------+
        /                                                          /
       /            <X>                      <Y>                  /
      /    +------O------+                +------O------+        /
     /     |    (7-0)    |                |    (8-0)    |       /
    /      |             |                |             |      /
   /       |    NE11     |                |     NE12    |     /
  /        +-------------+                +-------------+    /
 /                Ethernet or OTN Topology (O-PNC 1)        /
+-----------------------------------------------------------+


         +----------------------------------------------------------+
        /    <X> (7)                            (8) <Y>            /
       /         ---                            ---               /
      /    +-----\ /-----+                +-----\ /-----+        /
     /     |      V      |                |      V      |       /
    /      |             |                |             |      /
   /       |    NE11     |                |    NE12     |     /
  /        +-------------+                +-------------+    /
 /                   Optical Topology (O-PNC 1)             /
+----------------------------------------------------------+

Legenda:
========
  O   LTP
 ---
 \ /  TTP
  V
<   > Inter-Layer Lock-id reported by the PNC
```

Figure 5: Multi-layer optical topology discovery

The intra-domain optical links are discovered by O-PNCs, using mechanisms which are outside the scope of this document, and reported at the MPIs within the optical network topology.

In the case of a multi-layer DWDM/OTN network domain, multi-layer intra-domain OTN links are supported by underlay WDM tunnels: this relationship is reported by the mechanisms described in Section 4.2.

4.2. Optical Path Discovery

The WDM Tunnel Model is used to report all the WDM tunnels established within the optical network.

When the OTN switching layer is deployed within the optical domain, the OTN Tunnel Model is used to report all the OTN tunnels established within the optical network.

The Ethernet client signal model and the Transparent CBR client signal model are used to report all the connectivity services provided by the underlay optical tunnels between Ethernet or CBR client LTPs, depending on whether the connectivity service is frame-based or transparent. The underlay optical tunnels can be either WDM tunnels or, when the optional OTN switching layer is deployed, OTN tunnels.

The WDM tunnels can be used to support either Ethernet or CBR client signals or multi-layer intra-domain OTN links. In the latter case, the hierarchical-link container, defined in [I-D.ietf-teas-yang-te], associates the underlay WDM tunnel with the supported multi-layer intra-domain OTN link, and it allows discovery of the multi-layer path supporting all the connectivity services provided by the optical network.

The O-PNCs report in their operational datastores all the Ethernet and CBR client connectivities and all the optical tunnels deployed within their optical domain regardless of the mechanisms being used to set them up, such as the mechanisms described in Section 5.2, as well as other mechanism (e.g., static configuration), which are outside the scope of this document.

4.3. Packet Topology Discovery

The L3 Topology Model is used to report the IP network topology.

The L3 Topology Model, SR Topology Model, TE Topology Model and the TE Packet Topology Model are used together to report the SR-TE network topology, as described in Figure 2 of [I-D.ietf-teas-yang-sr-te-topo].

The TE Topology Model, TE Packet Topology Model and MPLS-TE Topology Model are used together to report the MPLS-TE network topology, as described in [I-D.ietf-teas-yang-te-mpls-topology].

As described in [I-D.ietf-teas-yang-l3-te-topo], the relationship between the IP network topology and the MPLS-TE network topology depend on whether the two network topologies are congruent or not: in the latter case, the L3 TE Topology Model is used, together with the L3 Topology Model to provide the association between the two network topologies.

To allow the MDSC to discover the complete multi-layer and multi-domain network topology and to correlate it with the hardware inventory information as well as to perform multi-domain TE path computation, the P-PNCs report the full packet network, including all the information that the MDSC requires to perform TE path computation. In particular, one TE node is reported for each IP router and one TE link is reported for each intra-domain IP link. The packet topology also reports the IP LTPs terminating the inter-domain IP links.

The Ethernet Topology Model is used to report the intra-domain Ethernet links supporting the intra-domain IP links as well as the Ethernet LTPs that might terminate cross-technology Ethernet links, inter-domain Ethernet links or access links, as described in detail in Section 4.5 and in Section 4.6.

All the intra-domain Ethernet and IP links are discovered by the P-PNCs, using mechanisms, such as Link Layer Discover Protocol LLDP [IEEE_802.1AB], which are outside the scope of this document, and reported at the MPIs within the Ethernet or the packet network topology.

4.4. TE Path Discovery

We assume that the discovery of existing TE paths, including their bandwidth, at the MPI is done using the generic TE tunnel YANG data model, defined in [I-D.ietf-teas-yang-te], with packet technology-specific (e.g., MPLS-TE or SR-TE) augmentations.

Note that technology-specific augmentations of the generic path TE tunnel model for SR-TE path setup and discovery is outlined in section 1 of [I-D.ietf-teas-yang-te] but are currently identified as a gap in Section 6.

To enable MDSC to discover the full end-to-end TE path configuration, the technology-specific augmentation of the [I-D.ietf-teas-yang-te] should allow the P-PNC to report the TE path within its domain (e.g., the SID list assigned to an SR-TE path).

For example, considering the L3VPN in Figure 2, the TE path 1 in one direction (PE13-P16-PE14) and the TE path in the reverse direction (between PE14 and PE13) should be reported by the P-PNC1 to the MDSC as TE primary and primary-reverse paths of the same TE tunnel instance. The bandwidth of these TE paths represents the bandwidth allocated by P-PNC1 to the two TE paths, which can be symmetric or asymmetric in the two directions.

The P-PNCs use the TE tunnel model to report, at the MPI, all the TE paths established within their packet domain regardless of the mechanism being used to set them up; i.e., independently on whether the mechanisms described in Section 5.3 or other means, such as static configuration, which are outside the scope of this document, are used.

4.5. Inter-domain Link Discovery

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

* Inter-domain Ethernet links supporting inter-domain IP links between two adjacent IP domains;

* Cross-technology Ethernet links between an an IP domain and an adjacent optical domain;

* Access links between a CE device and a PE router.

All the three types of links are Ethernet links.

It is worth noting that the P-PNC may not be aware whether an Ethernet interface terminates a cross-technology Ethernet link, an inter-domain Ethernet link or an access link. The TE Topology Model supports the discovery for all these types of links with no need for the P-PNC to know the type of inter-domain link.

There are two possible models to report the access links between CEs and PEs: the TE Topology Model, defined in [RFC8795], or the Service Attachment Points (SAP) Model, defined in [RFC9408].

Although the discovery of access links is outside the scope of this document, clarifying the relationship between these two models has been identified as a gap.

The inter-domain Ethernet links and cross-technology Ethernet links are discovered by the MDSC using the plug-id attribute, as described in section 4.3 of [RFC8795].

A more detailed description of how the plug-id can be used to discover inter-domain links is also provided in section 5.1.4 of [I-D.ietf-ccamp-transport-nbi-app-statement].

The plug-id attribute can also be used to discover the access-links, but the analysis of the access-link discovery is outside the scope of this document.

This document considers the following two options for discovering inter-domain links:

1. Static configuration

2. LLDP [IEEE_802.1AB] automatic discovery

Other link discovery options are possible but not described in this document.

As outlined in [I-D.ietf-ccamp-transport-nbi-app-statement], the encoding of the plug-id namespace and the specific LLDP information reported within the plug-id value, such as the Chassis ID and Port ID mandatory TLVs, is implementation-specific and needs to be consistent across all PNCs within the network.

The static configuration requires an administrative burden to configure network-wide unique identifiers, making it more viable for inter-domain Ethernet links. For cross-technology Ethernet links, the automatic discovery solution based on LLDP snooping is preferable when possible.

The routers exchange standard LLDP packets as defined in [IEEE_802.1AB], and the optical nodes snoop the LLDP packets received from the local Ethernet interface and report the extracted information, such as the Chassis ID, Port ID, and System Name TLVs, to the O-PNCs.

Note that the optical nodes do not actively participate in the LLDP packet exchange and do not send any LLDP packets.

4.5.1. Cross-technology Ethernet link Discovery

The MDSC can discover a cross-technology Ethernet link by matching the plug-id values of the two LTPs reported by adjacent O-PNC and P-PNCs. In case LLDP snooping is used, the P-PNC reports the LLDP information sent by the corresponding Ethernet interface on the IP router, while the O-PNC reports the LLDP information received by the corresponding Ethernet interface on the optical node, e.g., between LTP 5-0 on PE13 and LTP 7-0 on NE11, as shown in Figure 6.

```
           +-----------------------------------------------------------+
          /             Ethernet Topology (P-PNC)                     /
         /    +-------------+                +-------------+         /
        /     |    PE13     |                |    BR11     |        /
       /      |             |                |             |       /
      /       |    (5-0)    |                |    (6-0)    |      /
     /        +------O------+                +------O------+     /
    /       {PE13,5} ^                              ^ {BR11,6}  /
   +-----------------:------------------------------:----------+
                     :                              :
                     :                              :
                     :                              :
                     :                              :
            +--------:------------------------------:------------------+
           /         :                              :                 /
          / {PE13,5} v                              v {BR11,6}       /
         /    +------O------+                +------O------+        /
        /     |    (7-0)    |                |    (8-0)    |       /
       /      |             |                |             |      /
      /       |    NE11     |                |     NE12    |     /
     /        +-------------+                +-------------+    /
    /                Ethernet or OTN Topology (O-PNC)          /
   +----------------------------------------------------------+

   Legenda:
   ========
     O   LTP
   <...> Link discovered by the MDSC
   {   } LTP Plug-id reported by the PNC
```

Figure 6: Cross-technology Ethernet link discovery

As described in Section 4.1, the LTP terminating a cross-technology Ethernet link is reported by an O-PNC in the Ethernet topology, the OTN topology model, or both, depending on the type of corresponding physical port on the optical node.

It is worth noting that the discovery of cross-technology Ethernet links is based solely on the LLDP information sent by the Ethernet interfaces of the routers and snooped by the Ethernet interfaces of the optical nodes. Therefore, the MDSC can discover these links even before optical paths, supporting overlay multi-technology IP links, are set up.

4.5.2. Inter-domain IP Link Discovery

The MDSC can discover an inter-domain Ethernet link supporting an inter-domain IP link by matching the plug-id values of the two Ethernet LTPs reported by adjacent P-PNCs. The P-PNCs report the LLDP information being sent and received from the corresponding Ethernet interfaces, e.g., between Ethernet LTP 3-1 on BR11 and Ethernet LTP 4-1 on BR21, as shown in Figure 7.

```
           +--------------------------+     +-------------------------+
          /  IP Topology (P-PNC 1)   /     /  IP Topology (P-PNC 2)  /
         /   +-------------+        /     /   +-------------+       /
        /    |    BR11     |       /     /    |    BR21     |      /
       /     |        (3-2)O<................>O(4-2)        |     /
      /      |             |\    /     /     /|             |    /
     /       +-------------+|   /     /      |+-------------+   /
    /                       |  /     /       |                 /
   +------------------------|-+     +-------------------------+
                            |                |
             Supporting LTP |                | Supporting LTP
                            |                |
                            |                |
             +--------------|----------+    +|------------------------+
            /               V         /    / V                       /
           / +-------------+/        /    /  \+-------------+       /
          /  |     {1}(3-1)O<................>O(4-1){1}     |      /
         /   |             |\      /    /    /|             |     /
        /    |    BR11     |V(*)  /    /  (*)V|     BR21    |    /
       /     |             |/    /    /      \|             |   /
      /      |     {2}(3-0)O<~~~~~~~~~~~~~~~~>O(4-0){3}     |  /
     /       +-------------+   /    /         +-------------+ /
    / Eth. Topology (P-PNC 1) /    / Eth. Topology (P-PNC 2) /
   +-------------------------+    +-------------------------+

   Notes:
   =====
   (*) Supporting LTP
   {1} {BR11,3,BR21,4}
   {2} {BR11,3}
   {3} {BR21,4}

   Legenda:
   ========
     O   LTP
   ----> Supporting LTP
   <...> Link discovered by the MDSC
   <~~~> Link inferred by the MDSC
   {   } LTP Plug-id reported by the PNC
```

Figure 7: Inter-domain Ethernet and IP link discovery

Different information is required to be encoded by the P-PNC within the plug-id attribute of the Ethernet LTPs to discover cross-technology Ethernet links and inter-domain Ethernet links.

If the P-PNC does not know a priori whether an Ethernet interface on an IP router terminates a cross-technology Ethernet link or an inter-domain Ethernet link, it must report at the MPI two Ethernet LTPs representing the same Ethernet interface, e.g., both Ethernet LTP 3-0 and Ethernet LTP 3-1, supported by LTP 3-0, as shown in Figure 7.

* The physical Ethernet LTP (e.g., LTP 3-0 in BR11, as shown in Figure 7) represents the physical adjacency between the Ethernet interface on an IP router and the Ethernet interface on its physically adjacent node. This node can be either an IP router (in the case of a single-technology Ethernet link) or an optical node (in the case of a cross-technology Ethernet link). Therefore, as described in Section 4.5.1, the P-PNC reports, within the plug-id attribute of this LTP, the LLDP information sent by the corresponding Ethernet interface on the IP router, such as {BR11,3} and {BR21,4} plug-id values reported by the Ethernet LTP 3-0 on BR11 and the Ethernet LTP 4-0 on BR21, as shown in Figure 7.

* The logical Ethernet LTP (e.g., LTP 3-1 in BR11, as shown in Figure 7), supported by a physical Ethernet LTP (e.g., LTP 3-0 in BR11, as shown in Figure 7), is used to discover the logical adjacency between Ethernet interfaces on IP routers, which can be either single-technology or multi-technology. Therefore, the P-PNC reports, within the plug-id attribute of this LTP, the LLDP information sent and received by the corresponding Ethernet interface on the IP router, such as the {BR11,3,BR21,4} plug-id values reported by the Ethernet LTP 3-1 on BR11 and the Ethernet LTP 4-1 on BR21, as shown in Figure 7.

It is worth noting that in the case of inter-domain Ethernet links, the MDSC cannot discover, using the LLDP information reported in the plug-id attributes, the physical adjacency between two Ethernet interfaces on physically adjacent IP routers, because these plug-id values do not match, such as the {BR11,3} and {BR21,4} plug-id values shown in Figure 7. However, the MDSC may infer the physical intra-domain Ethernet links if it knows a priori, using mechanisms outside the scope of this document, that the Ethernet interfaces on the IP routers either terminate a cross-technology or single-technology (intra-domain or inter-domain) Ethernet link, e.g., as shown in Figure 7.

The P-PNC can omit to report the physical Ethernet LTPs when it knows, through mechanisms outside the scope of this document, that the corresponding Ethernet interfaces terminate inter-domain Ethernet links.

The MDSC can then discover an inter-domain IP link between the two IP LTPs supported by the two Ethernet LTPs terminating an inter-domain Ethernet link, discovered as described in Section 4.5.2, e.g., between IP LTP 3-2 on BR21 and IP LTP 4-2 on BR22, supported respectively by Ethernet LTP 3-1 on BR11 and Ethernet LTP 4-1 on BR21, as shown in Figure 7.

4.6. Multi-technology IP Link Discovery

A multi-technology intra-domain IP link and its supporting multi-technology intra-domain Ethernet link are discovered by the P-PNC like any other intra-domain IP and Ethernet links, as described in Section 4.3, and reported at the MPI within the packet and the Ethernet network topologies, e.g., as shown in Figure 8.

```
           +-----------------------------------------------------------+
          /                    IP Topology (P-PNC 1)                  /
         /    +---------+                        +---------+         /
        /     |  PE13   |                        |   BR11  |        /
       /      |    (5-2)O<======================>O(6-2)    |       /
      /       |         |              |         |         |      /
     /        +---------+              |         +---------+     /
    /                                  |                        /
   +-----------------------------------|-----------------------+
                                       |
                                       | Supporting Link
                                       |
           +---------------------------|-------------------------------+
          / Ethernet Topology (P-PNC 1)|                              /
         /    +-------------+          |    +-------------+          /
        /     |    PE13     |          V    |    BR11     |         /
       /      |        (5-1)O<==============>O(6-1)       |        /
      /       |    (5-0)    |\             /|    (6-0)    |       /
     /        +------O------+|(*)       (*)|+------O------+      /
    /                ^ \<----+             +----->/^            /
   +-----------------:-------------------------------:----------+
                     :                               :
                     :                               :
                     :                               :
           +---------:-------------------------------:-----------------+
          /          :   Ethernet or OTN Topology    :                /
         /           V           (O-PNC 1)           V               /
        /     +------O------+    ETH/CBR     +------O------+        /
```

```
       /      |    (7-0)    |  client sig.   |    (8-0)    |       /
      /       |      X----------+-------------------X      |      /
     /        |    NE11     |   |            |     NE12    |     /
    /         +-------------+   |            +-------------+    /
   +----------------------------|------------------------------+
                                | Underlay
                                | tunnel
                                |
            +----------------------------------------------------------+
           /        (7)         |                  (8)                /
          /         ---         |                  ---               /
         /    +-----\ /-----+   v            +-----\ /-----+        /
        /     |      V      |                |      V      |       /
       /      |      X======|================|======X      |      /
      /       |    NE11     |  Opt. Tunnel   |    NE12     |     /
     /        +-------------+                +-------------+    /
    /                   Optical Topology (O-PNC 1)             /
   +----------------------------------------------------------+

   Notes:
   =====
   (*) Supporting LTP

   Legenda:
   ========
     O   LTP
    ---
    \ /  TTP
     V
   ----> Supporting LTP or Supporting Link or Underlay tunnel
   <===> Link discovered by the PNC and reported at the MPI
   <...> Link discovered by the MDSC
   x---x Ethernet/CBR client signal
   X===X Optical tunnel
```

Figure 8: Multi-technology intra-domain Ethernet and IP link discovery

The Ethernet interface 5 on the P13 router is terminating two Ethernet abstract links: - The multi-technology intra-domain Ethernet link between logical Ethernet LTP 5-1 on PE13 and the logical Ethernet LTP 6-1 on BR11; - The cross-technology Ethernet link, which is supporting that multi-technology intra-domain Ethernet link, between the physical Ethernet LTPs 5-0 on PE13 and the physical Ethernet LTP 7-0 on the optical NE11.

The P-PNC does not report any plug-id information on the logical Ethernet LTPs terminating intra-domain Ethernet links, such as the LTP 5-1 on PE13 and LTP 6-1 in BR11 shown in Figure 8, since these links are discovered by the PNC.

In addition, the P-PNC also reports the physical Ethernet LTPs that terminate the cross-technology Ethernet links supporting the multi-technology intra-domain Ethernet links, e.g., the Ethernet LTP 5-0 on PE13 and the Ethernet LTP 6-0 on BR11, shown in Figure 8.

The MDSC discovers, using the mechanisms described in Section 4.5, which cross-technology Ethernet links support the multi-technology intra-domain Ethernet links, e.g., the link between LTP 5-0 on PE13 and LTP 7-0 on NE11, shown in Figure 8.

The MDSC also discovers, from the information provided by the O-PNC and described in Section 4.2, which optical tunnels support the multi-technology intra-domain IP links and therefore the path within the optical network that supports a multi-technology intra-domain IP link, e.g., as shown in Figure 8.

4.6.1. Intra-domain single-technology IP Links

It is worth noting that the P-PNC may not be aware of whether an Ethernet interface on the IP router terminates a multi-technology or a single-technology intra-domain Ethernet link.

In this case, the P-PNC, always reports two Ethernet LTPs for each Ethernet interface on the IP router, e.g., the Ethernet LTP 1-0 and 1-1 on PE13, shown in Figure 9.

```
           +-----------------------------------------------------------+
          /                    IP Topology (P-PNC 1)                  /
         /    +---------+                       +---------+          /
        /     |  PE13   |                       |    P16  |         /
       /      |    (1-2)O<======================>O(2-2)    |        /
      /       |         |            |          |         |       /
     /        +---------+            |          +---------+      /
    /                                |                          /
   +---------------------------------|-------------------------+
                                     |
                                     | Supporting Link
                                     |
                                     |
             +-----------------------|---------------------------------+
            /                        |                                /
           /  +---------+            v          +---------+          /
          /   |    (1-1)O<======================>O(2-1)    |         /
         /    |         |\                     /|         |        /
        /     |  PE13   |V(*)               (*)V|    P16  |       /
       /      |         |/                     \|         |      /
      /       | {1}(1-0)O<~~~~~~~~~~~~~~~~~~~~~~>O(2-0){2} |     /
     /        +---------+                       +---------+    /
    /                   Ethernet Topology (P-PNC 1)           /
   +---------------------------------------------------------+

   Notes:
   =====
   (*) Supporting LTP
   {1} {PE13,1}
   {2} {P16,2}

   Legenda:
   ========
     O   LTP
   ----> Supporting LTP
   <===> Link discovered by the PNC and reported at the MPI
   <~~~> Link inferred by the MDSC
   {   } LTP Plug-id reported by the PNC
```

Figure 9: Single-technology intra-domain Ethernet and IP link discovery

It is worth noting that in the case of intra-domain single-technology Ethernet links, the MDSC cannot discover, using the LLDP information reported in the plug-id attributes, the physical adjacency between two Ethernet interfaces on physically adjacent IP routers, because the plug-id values do not match, such as {PE13,1} and {P16,2}, as shown in Figure 9. However, the MDSC may infer the physical intra-domain Ethernet links, e.g., between LTP 1-0 on PE13 and LTP 2-0 on P16, as shown in Figure 9, if it knows a priori, using mechanisms outside the scope of this document, that all Ethernet interfaces on the IP routers terminate either a cross-technology or single-technology (intra-domain or inter-domain) Ethernet link, e.g., as shown in Figure 9.

The P-PNC can omit reporting the physical Ethernet LTP if it knows, through mechanisms outside the scope of this document, that the intra-domain Ethernet link is single-technology.

4.7. LAG Discovery

The P-PNCs can discover the configuration of LAG groups within its domain and report each intra-domain LAG as an Ethernet bundle link within the Ethernet topology exposed at the MPI.

This is done by bundling multiple single-domain Ethernet links, as shown in Figure 10. For example, the Ethernet bundled link between Ethernet LTP 5-1 on BR21 and Ethernet LTP 6-1 on P24 is built from the Ethernet links set up respectively:

* between the Ethernet LTP 1-1 on BR21 and the Ethernet LTP 2-1 on P24; and

* between the Ethernet LTP 3-1 on BR21 and the Ethernet LTP 4-1 on P24.

```
           +-----------------------------------------------------------+
          /                    IP Topology (P-PNC 2)                  /
         /    +---------+                        +---------+         /
        /     |  BR21   |                        |    P24  |        /
       /      |    (5-2)O<======================>O(6-2)    |       /
      /       |         |            |           |         |      /
     /        +---------+            |           +---------+     /
    /                                |                          /
   +---------------------------------|-------------------------+
                                     |
                                     | Supporting Link
                                     |
                                     |
             +-----------------------|---------------------------------+
            /                        |                                /
           /  +---------+            v           +---------+         /
          /   |    (5-1)O<======================>O(6-1)    |        /
         /    |  BR21   |  Bundled Link          |    P24  |       /
        /     |         |                        |         |      /
       /      |    (3-1)O<======================>O(4-1)    |     /
      /       |    (1-1)O<======================>O(2-1)    |    /
     /        +---------+                        +---------+   /
    /                   Ethernet Topology (P-PNC 2)           /
   +---------------------------------------------------------+

   Legenda:
   ========
     O   LTP
   <===> Link discovered by the PNC and reported at the MPI
```

Figure 10: LAG

The mechanisms used by the MDSC to discover single-technology and multi-technology intra-domain LAG links are the same, with the only difference being whether the bundled links are single-technology or multi-technology.

However, the mechanisms used by the MDSC to discover single-technology inter-domain LAG links between two BRs are different and outside the scope of this document, as they do not imply cross-technology coordination between packet and optical domains.

As described in Section 4.3, the mechanisms used by the P-PNC to discover the configuration of LAG groups within its domain, such as LLDP [IEEE_802.1AB], are outside the scope of this document.

It is worth noting that according to [IEEE_802.1AB], LLDP can be configured on a LAG group (Aggregated Port) and/or on any number of its LAG members (Aggregation Ports).

If LLDP is enabled on both LAG members and groups, two types of LLDP packets are transmitted by the routers and received by the optical nodes on some cross-technology Ethernet links: one sent for the LLDP session configured at LAG member (Aggregation Port)level and another one for the LLDP session configured at LAG group (Aggregated Port)level. This could cause some issues when LLDP snooping is used to discover the cross-technology Ethernet links, as defined in Section 4.5.1.

The cross-technology Ethernet link discovery is based only on the LLDP session configured on the LAG members (Aggregation Ports) to allow discovery of these links independently from the configuration of the underlay optical tunnel or from the LAG group.

To avoid any ambiguity on how the optical nodes can identify which LLDP packets belong to which LLDP session, the P-PNC can disable the LLDP sessions on the LAG groups configured by the MDSC (e.g., the multi-technology single-domain LAG groups configured using the mechanisms described in Section 5.2.1), keeping the LLDP sessions on the LAG members enabled.

Another option is to rely on other mechanisms (e.g., the Port type field in the Link Aggregation TLV defined in Annex F of [IEEE_802.1AX]) that allow the optical node to identify which LLDP packets belong to which LLDP session: the O-PNC can then use only the LLDP information from the LLDP sessions configured on the LAG members to support the cross-technology Ethernet link discovery mechanisms defined in Section 4.5.1.

4.8. L2/L3 VPN Network Services Discovery

The P-PNC reports the L2/L3 VPN services configured within its domain, using the L2NM and L3NM network service models, and which packet TE tunnels (e.g., MPLS-TE or SR-TE) are used by each L2/L3 VPN service, using the L2NM and L3NM TE service mapping models.

The MDSC can use the information mentioned above together with the packet TE path, packet topology, multi-technology IP links, optical topology and optical path information discovered as described in the previous sections, to discover the multi-technology path used to carry the traffic for each L2/L3 VPN service.

4.9. Inventory Discovery

The are no YANG data models in IETF that could be used to report at the MPI the whole inventory information discovered by a PNC.

[RFC8345] had foreseen some work for inventory as an augmentation of the network model, but no YANG data model has been developed so far.

There are also no YANG data models in IETF that could be used to correlate topology information, e.g., a link termination point (LTP), with inventory information, e.g., the physical port supporting an LTP, if any.

Inventory information through MPI and correlation with topology information is identified as a gap requiring further work and outside of the scope of this draft.

5. Establishment of L2/L3 VPN Services with TE Requirements

In this scenario the MDSC needs to setup a multi-domain L2VPN or a multi-domain L3VPN with some SLA requirements.

The MDSC receives the request to setup a L2/L3 VPN network service from the OSS/Orchestration layer (see Appendix A).

The MDSC translates the L2/L3 VPN SLA requirements into TE requirements (e.g., bandwidth, TE metric bounds, SRLG disjointness, nodes/links/domains inclusion/exclusion) and find the TE paths that meet these TE requirements (see Section 2.1.1).

For example, considering the L3VPN in Figure 2 and Figure 3, the MDSC finds that:

* PE13-P16-PE14 TE path already exists but have not enough bandwidth to support the new L3VPN, as described in Section 4.4;, and that:

  - the IP link(s) between PE13 and P16 has not enough bandwidth to support increasing the bandwidth of that TE path, as described in Section 4.3;

  - a new underlay optical tunnel could be setup to increase the bandwidth of the IP link(s) between PE13 and P16 to support increasing the bandwidth of that overlay TE path, as described in Section 5.1. The dimensioning of the underlay optical tunnel is decided by the MDSC based on the TE requirements (e.g., the bandwidth) requested by the TE path and on its multi-layer optimization policy, which is an internal MDSC implementation issue;

- a new multi-domain TE path needs to be setup between PE13 and PE23, e.g., either because existing TE paths between PE13 and PE23 are not able to meet the TE and binding requirements of the L2/L3 VPN service or because there is no TE path between PE13 and PE23.

As described in Section 2.1.2, with partial summarization, the MDSC will use the TE topology information provided by the P-PNCs and the results of the path computation requests sent to the O-PNCs, as described in Section 5.1, to compute the multi-layer/multi-domain path between PE13 and PE23.

For example, the multi-layer/multi-domain performed by the MDSC could require the setup of:

* a new underlay optical tunnel between PE13 and BR11, supporting a new IP link, as described in Section 5.2;

* a new underlay optical tunnel between BR21 and P24 to increase the bandwidth of the IP link(s) between BR21 and P24, as described in Section 5.2.

When setting up the L2/L3 VPN network service requires multi-domain and multi-layer coordination, the MDSC is also responsible for coordinating the network configuration needed to realize the requested network service across the appropriate optical and packet domains.

The MDSC would therefore request:

* the O-PNC1 to setup a new optical tunnel between the ROADMs connected to PE13 and P16, as described in Section 5.2;

* the P-PNC1 to update the configuration of the existing IP link, in case of LAG, or configure a new IP link, in case of ECMP, between PE13 and P16, as described in Section 5.2;

* the P-PNC1 to update the bandwidth of the selected TE path between PE13 and PE14, as described in Section 5.3.

After that, the MDSC requests P-PNC2 to set up a TE path between BR21 and PE23, with an explicit path (BR21, P24, PE23) to constrain the new TE path to use the underlay optical tunnel setup between BR21 and P24, as described in Section 5.3. The P-PNC2 properly configures the routers within its domain to set up the requested path and returns the information needed for multi-domain TE path stitching. For example, in inter-domain SR-TE, the P-PNC2, knowing the node and adjacency SIDs assigned within its domain, can install the proper SR policy or hierarchical policies within BR21 and return to the MDSC the binding SID assigned to this policy in BR21.

Then the MDSC requests P-PNC1 to set up a TE path between PE13 and BR11, with an explicit path (PE13, BR11) to constrain the new TE path to use the underlay optical tunnel setup between PE13 and BR11, specifying which inter-domain link should be used to send traffic to BR21 and the information for multi-domain TE path stitching, as described in Section 4.4 (e.g., in inter-domain SR-TE, the binding SID assigned by P-PNC2 to the corresponding SR policy in BR21). The P-PNC1 properly configures the routers within its domain to set up the requested path and the multi-domain TE path stitching. For example, in inter-domain SR-TE, the P-PNC1, knowing the node and adjacency SIDs assigned within its domain and the PE SID assigned by P-PNC1 to the inter-domain link between BR11 and BR21, along with the binding SID assigned by P-PNC2, installs the proper policy or policies within PE13.

Once the TE paths have been selected and, if needed, set up or modified, the MDSC can request both P-PNCs to configure the L3VPN and its binding with the selected TE paths, as described in Section 5.4.

5.1. Optical Path Computation

As described in Section 2.1.2, optical path computation is usually performed by the O-PNCs.

When performing multi-layer/multi-domain path computation, the MDSC can delegate single-domain optical path computation to the O-PNC.

As described in Section 4.1, Section 4.5, and Section 4.6, there is a one-to-one relationship between a multi-layer intra-domain IP link and its underlay optical tunnel. Therefore, the properties of an optical path between two optical TTPs, as computed by the O-PNC, can be used by the MDSC to infer the properties of the associated multi-layer single-domain IP link.

As discussed in [I-D.ietf-teas-yang-path-computation], there are two options to request an O-PNC to perform optical path computation: either via a "compute-only" TE tunnel path, using the generic TE tunnel YANG data model defined in [I-D.ietf-teas-yang-te], or via the path computation RPC defined in [I-D.ietf-teas-yang-path-computation].

This draft assumes that the path computation RPC is used.

There are no YANG data models in IETF that could be used to augment the generic path computation RPC with technology-specific attributes.

Optical technology-specific augmentation for the path computation RPC is identified as a gap requiring further work outside of this draft’s scope.

5.2. Multi-technology IP Link Setup

As described in Section 5.1, there is a one-to-one relationship between a multi-technology intra-domain IP link and its underlay optical tunnel.

Therefore, to set up a new multi-technology intra-domain IP link, the MDSC requires the O-PNC to set up the optical tunnel (using either the WDM Tunnel model or the OTN Tunnel model, if optional OTN switching is supported) within the optical network and steer client traffic between the two cross-technology Ethernet links over that optical tunnel, using either the Ethernet Client Signal Model (for frame-based transport) or the Transparent CBR Client Signal Model (for transparent transport).

For example, with reference to Figure 11, the MDSC can request O-PNC1 to set up an optical tunnel between optical TTPs (7) on NE11 and (8) on NE12 and steer client traffic over this tunnel between LTP (7-0) on NE11 and LTP (8-0) on NE12.

```
           +-----------------------------------------------------------+
          /                    IP Topology (P-PNC 1)                  /
         /    +---------+                        +---------+         /
        /     |  PE13   |                        |   BR11  |        /
       /      |    (5-2)O<======================>O(6-2)    |       /
      /       |         |              |         |         |      /
     /        +---------+              |         +---------+     /
    /                                  |                        /
   +-----------------------------------|-----------------------+
                                       |
                                       | Supporting Link
                                       |
           +---------------------------|-------------------------------+
          / Ethernet Topology (P-PNC 1)|                              /
         /    +-------------+          |     +-------------+         /
```

```
        /     |    PE13     |          V     |    BR11     |        /
       /      |        (5-1)O<==============>O(6-1)        |       /
      /       |    (5-0)    |\              /|    (6-0)    |      /
     /        +------O------+|(*)        (*)|+------O------+     /
    /                ^ \<----+              +----->/^           /
   +-----------------:------------------------------:----------+
                     :                              :
                     :                              :
                     :                              :
           +---------:------------------------------:------------------+
          /          :   Ethernet or OTN Topology   :                 /
         /           V          (O-PNC 1)           V                /
        /     +------O------+    ETH/CBR     +------O------+        /
       /      |    (7-0)    |  client sig.   |    (8-0)    |       /
      /       |      X----------+-------------------X      |      /
     /        |    NE11     |   |            |     NE12    |     /
    /         +-------------+   |            +-------------+    /
   +----------------------------|------------------------------+
                                | Underlay
                                | tunnel
                                |
            +----------------------------------------------------------+
           /        (7)         |                  (8)                /
          /         ---         |                  ---               /
         /    +-----\ /-----+   v            +-----\ /-----+        /
        /     |      V      |                |      V      |       /
       /      |      X======|================|======X      |      /
      /       |    NE11     |  Opt. Tunnel   |    NE12     |     /
     /        +-------------+                +-------------+    /
    /                   Optical Topology (O-PNC 1)             /
   +----------------------------------------------------------+

   Notes:
   =====
   (*) Supporting LTP

   Legenda:
   ========
     O   LTP
    ---
    \ /  TTP
     V
   ----> Supporting LTP or Supporting Link or Underlay tunnel
   <===> Link discovered by the PNC and reported at the MPI
   <...> Link discovered by the MDSC
   x---x Ethernet/CBR client signal
   X===X Optical tunnel
```

Figure 11: Multi-technology IP link setup

Note: Figure 11 is an exact copy of Figure 8.

After the optical tunnel has been set up and the client traffic steering configured, the two IP routers can exchange Ethernet frames between themselves, including LLDP messages.

If LLDP [IEEE_802.1AB] or any other discovery mechanisms, outside the scope of this document, are used between the adjacency of the two IP routers' ports, the P-PNC can automatically discover the underlay multi-technology single-domain Ethernet link set up by the MDSC and report it to the P-PNC, as described in Section 4.6.

Otherwise, if no automatic discovery mechanisms are used, the MDSC can configure this multi-technology single-domain Ethernet link at the MPI of the P-PNC.

The two Ethernet LTPs terminating this multi-technology single-domain Ethernet link are supported by the two underlay Ethernet LTPs terminating the two cross-technology Ethernet links, e.g., LTP 5-1 on PE13 and 6-1 on BR11, as shown in Figure 11.

After the multi-technology single-domain Ethernet link has been configured by the MDSC or discovered by the P-PNC, the corresponding multi-technology single-domain IP link can also be configured either by the MDSC or the P-PNC.

This document assumes that the IP link is configured by the P-PNC.

It is worth noting that if LAG is not supported within the domain controlled by the P-PNC, the P-PNC can configure the multi-technology single-domain IP link as soon as the underlay multi-technology single-domain Ethernet link is either discovered by the P-PNC or configured by the MDSC at the MPI. However, if LAG is supported, the P-PNC lacks enough information to determine whether the discovered or configured multi-technology single-domain Ethernet link would be:

1. Used to support a multi-technology single-domain IP link;

2. Used to create a new LAG group;

3. Added to an existing LAG group.

The MDSC can request the P-PNC to configure a new multi-technology single-domain IP link, supported by the just discovered or configured multi-technology single-domain Ethernet link, by creating an IP link within the running datastore of the P-PNC MPI. Only the IP link, IP

LTPs, and the reference to the supporting multi-technology single-domain Ethernet link are configured by the MDSC. All other configuration is provided by the P-PNC.

For example, with reference to Figure 11, the MDSC can request P-PNC1 to set up a multi-technology single-domain IP link between IP LTP 5-2 on PE13 and IP LTP 6-2 on BR11, supported by the multi-technology single-domain Ethernet link between ETH LTP 5-1 on PE13 and ETH LTP 6-1 on BR11.

The P-PNC configures the requested multi-technology single-domain IP link and, once finished, reports it to the MDSC within the IP topology exposed at its MPI.

5.2.1. Multi-technology LAG Setup

The P-PNC configures a new LAG group between two routers when the MDSC creates a new Ethernet bundled link at the MPI (using the bundled-link container defined in [RFC8795]), bundling the multi-technology single-domain Ethernet link(s) being created, as described above.

When a new LAG link is created, it is recommended to configure the minimum number of active member links required to consider the LAG link as up. For example, a LAG link with three members can be considered up when only one member link fails and down when at least two member links fail.

The attribute required to configure the minimum number of active member links is missing in [I-D.ietf-ccamp-eth-client-te-topo-yang] and is identified as a gap in Section 6.

It is worth noting that a new LAG group can be created to bundle one or more multi-technology single-domain Ethernet link(s).

For example, with reference to Figure 10, the MDSC can request P-PNC2 to set up an Ethernet bundled link between Ethernet LTP 5-1 on BR21 and Ethernet LTP 6-1 on P24, bundling the multi-technology single-domain Ethernet link between Ethernet LTP 1-1 on BR21 and Ethernet LTP 2-1 on P24.

It is also worth noting that the MDSC needs to create the Ethernet LTPs terminating the Ethernet bundled link.

The MDSC can request the P-PNC to configure a new multi-technology single-domain IP link, supported by the just configured Ethernet bundled link, following the same procedure described in Section 5.2 above.

For example, with a reference to Figure 10, the MDSC can request the P-PNC2 to setup a multi-technology single-domain IP Link between IP LTP 5-2 on BR21 and IP LTP 6-2 on P24 supported by the Ethernet bundle link between ETH LTP 5-1 on BR21 and the Ethernet LTP 6-1 on P24.

5.2.2. Multi-technology LAG Update

The P-PNC adds new member(s) to an existing LAG group when the MDSC updates the configuration of an existing Ethernet bundled link at the MPI, adding the multi-technology single-domain Ethernet link(s) being created, as described above.

When member links are added or removed from a LAG link, the minimum number of active member links required to consider the LAG link as up may also need to be updated.

For example, with reference to Figure 10, the MDSC can request P-PNC2 to add the multi-technology single-domain Ethernet link set up between Ethernet LTP 3-1 on BR21 and Ethernet LTP 4-1 on P24 to the existing Ethernet bundle link set up between Ethernet LTP 5-1 on node BR21 and Ethernet LTP 6-1 on node P24.

After the LAG configuration has been updated, the P-PNC can also update the bandwidth information of the multi-technology single-domain IP link supported by the updated Ethernet bundled link.

5.2.3. Multi-technology TE path properties Configuration

The MDSC can discover the TE path properties (e.g., the list of SRLGs, the delay) of a multi-technology IP link from the TE properties of:

* the IP LTPs terminating the multi-technology IP link (e.g., the list of SRLGs reported by the P-PNC using the packet TE topology model);

* the optical path (e.g., the list of SRLGs reported by the O-PNC using the WDM or OTN tunnel model); and

* the cross-domain links (e.g., the list of SRLGs reported by the O-PNC and P-PNC respectively, using the WSON and/or flexi-grid, the OTN and the packet TE topology models).

The MDSC can also report this information to the P-PNC by properly configuring the multi-technology IP link properties using the packet TE topology model at the packet PNC MPI.

This information is used by the P-PNC at least when computing the local protection path, as described in Section 5.3, e.g., to ensure that the local protection path is SRLG disjoint with the primary path.

It is worth noting that the list of SRLGs for a multi-technology IP link can be quite long. Implementation-specific mechanisms can be implemented by the MDSC or by the O-PNC to summarize the SRLGs of an optical tunnel. These mechanisms are implementation-specific and have no impact on the YANG models nor on the interoperability at the MPI, but cares have to be taken to avoid missing information.

5.3. TE Path Setup and Update

This document assumes that TE path setup and update at the MPI could be done using the generic TE tunnel YANG data model, defined in [I-D.ietf-teas-yang-te], with packet technology-specific augmentations, described in Section 3.2.3.

When a new TE path needs to be setup, the MDSC can use the [I-D.ietf-teas-yang-te] model to request the P-PNC to set it up, properly specifying the path constraints, such as the explicit path, to ensure the P-PNC sets up a TE path that meets the end-to-end TE and binding constraints and uses the optical tunnels set up by the MDSC to support this new TE path.

The [I-D.ietf-teas-yang-te] model supports requesting the setup of both end-to-end as well as segment TE tunnels (within one domain).

In the latter case, the technology-specific augmentations should allow the configuration of the information needed for multi-domain TE path stitching.

For example, the SR-TE specific augmentations of the [I-D.ietf-teas-yang-te] model should be defined to allow the MDSC to configure the binding SIDs to be used for the multi-domain SR-TE path stitching and to allow the P-PNC to report the binding SID assigned to the segment TE paths. Note that the assigned binding SID should be persistent in case IP router or P-PNC rebooting.

The MDSC can also use the [I-D.ietf-teas-yang-te] model to request the P-PNC to increase the bandwidth allocated to an existing TE path, and, if needed, also on its reverse TE path. The [I-D.ietf-teas-yang-te] model supports both symmetric and asymmetric bandwidth configuration in the two directions.

[Editor’s Note:] Add some text about the protection options (to further discuss whether to put this text here or in Section 5.2).

The MDSC also request the P-PNC to configure local protection mechanisms. For example, the FRR local protection, as defined in [RFC4090] in case of MPLS-TE domain or the TI-LFA local protection, as defined in [I-D.ietf-rtgwg-segment-routing-ti-lfa] in case of SR-TE domain. The mechanisms to request the configuration TI-LFA local protection for SR-TE paths using the [I-D.ietf-teas-yang-te] are a gap in the current YANG models.

The requested local protection mechanisms within the P-PNC domain are configured by the P-PNC through implementation specific mechanisms which are outside the scope of this document.

The P-PNC takes into account the multi-layer TE path properties (e.g., SRLG information), configured by the MDSC as described in Section 5.2.3, when computing the protection configuration (e.g., in case of SR-TE domains, the TI-LFA post-convergence path or, in case of MPLS-TE domain, the FRR backup tunnel) for multi-technology single-domain IP links.

SR-TE path setup and update (e.g., bandwidth increase) through MPI is identified as a gap requiring further work, which is outside of the scope of this draft.

5.4. L2/L3 VPN Network Service Setup

The MDSC can use the L2NM and L3NM network service models to request the P-PNCs to setup L2/L3 VPN services, and the L2NM and L3NM TE service mapping models to request the P-PNCs to configure the PE routers to steer the L2/L3 VPN traffic to the selected TE tunnels (e.g., MPLS-TE or SR-TE).

It is worth noting that the L2NM and L3NM TE service mapping models, defined in [I-D.ietf-teas-te-service-mapping-yang], provide a list of TE tunnel(s) that should be used to forward L2/L3 VPN traffic between the two PEs terminating the listed TE tunnel(s). If the list contains more than one TE tunnel for the same pair of PEs, these TE tunnels are used to load balance the associated L2/L3 VPN traffic between the same set of two PEs.

The possibility to request splitting the traffic between multiple TE tunnels for the same PE pair in a way other than load balancing is identified as a gap requiring further work and is outside the scope of this draft.

6. Conclusions

The analysis provided in this document shows that the IETF YANG models described in Section 3.2 provide useful support for Packet Optical Integration (POI) scenarios for resource discovery (network topology, service, tunnels, and network inventory discovery), as well as for supporting multi-layer/multi-domain L2/L3 VPN network services.

The following gaps were identified that may need to be addressed by the relevant IETF Working Groups:

* how both WSON and Flexi-grid topology models could be used together (through multi-inheritance): this gap has been identified in Section 4.1;.

* network inventory model: this gap has been identified in Section 4.9 and the solution in [I-D.ietf-ivy-network-inventory-yang] has been proposed to resolve it;

* technology-specific augmentations of the path computation RPC, defined in [I-D.ietf-teas-yang-path-computation] for optical networks: this gap has been identified in Section 5.1 and the solution in [I-D.ietf-ccamp-optical-path-computation-yang] has been proposed to resolve it;

* relationship between a common discovery mechanisms applicable to access links, inter-domain IP links and cross-technology Ethernet links and the UNI topology discover mechanism defined in [RFC9408]: this gap has been identified in Section 4.3;

* a mechanism applicable to the P-PNC NBI to configure the SR-TE paths. Technology-specific augmentations of TE Tunnel model, defined in [I-D.ietf-teas-yang-te], are foreseen in section 1 of [I-D.ietf-teas-yang-te] but not yet defined: this gap has been identified in Section 5.3;

* an attribute, which is used to configure the minimum number of active member links required to consider the LAG link as being up, is missing from the topology model defined in [I-D.ietf-ccamp-eth-client-te-topo-yang]: this gap has been identified in Section 5.2.1;

* a mechanism to configure splitting the L2/L3 VPN traffic, between multiple TE tunnels for the same PEs pair, in a different way than load balancing: this gap has been identified in Section 5.4;

* a mechanism to report client connectivity constraints imposed by some muxponder design: this gap has been identified in Appendix A.3.

Although not applicable to this document, it has been noted that being able to use WSON and Flexi-grid topology models together (through multi-inheritance) is not only useful for mixed fixed-grid and flexible-grid DWDM network topologies but also the only viable option for a mixed CWDM and DWDM network topology.

Although not applicable to this document, it has been noted that the WDM tunnel model would also support optical tunnel setup in the case of a mixed CWDM and DWDM network topology.

Although not analyzed in this document, it has been noted that the TE Tunnel model, defined in [I-D.ietf-teas-yang-te], needs enhancement to support scenarios where multiple parallel TE paths are used in load-balancing to carry traffic between two end-points (e.g., VPN traffic between two PEs).

7. Security Considerations

This document highlights how the ACTN architecture can deploy packet over optical infrastructure services. It highlights how existing IETF protocols and data models may be used for multi-layer services. It reuses several existing IETF protocols and data models for the MPI interfaces between each PNC (Optical or Packet) and the MDSC, including:

* RESTCONF

* NETCONF

* PCEP

* YANG

Several existing authentication and encryption practices and techniques may be used to help secure these MPI interfaces. These mechanisms include using Transport Layer Security (TLS) to provide secure transport for RESTCONF, NETCONF and PCEP. Furthermore, access control techniques can also provide additional security. NETCONF supports an Access Control Model (NACM), and RESCONF supports Role Based Access Control (RBAC), which should also ensure that MDSC to PNC communication is based on authorised use and granular control of connectivity and resource requests.

7.1. LLDP Snooping Security Considerations

Earlier in the document, LLDP is discussed as a mechanism for the PNCs to discover the intra-domain Ethernet and IP links. While LLDP provides valuable information for network management and troubleshooting, it also presents several security issues:

* Eavesdropping: LLDP transmissions are not encrypted. Potentially, LLDP packets could be captured using a packet sniffer. An attacker can leverage this information to gain insights into the network topology, device types, and configurations, which could be used for further attacks;

* Unauthorized Access: Information disclosed by LLDP can include device types, software versions, and network configuration details. This might help an attacker identify vulnerable devices or configurations that can be exploited to gain unauthorized access or escalate privileges within the network;

* Data Manipulation: If an attacker gains access to a network device, they could manipulate LLDP information to advertise false device information, leading to potential misconfigurations or trust relationships being exploited. This can disrupt network operations or redirect traffic to malicious devices;

* Denial of Service (DoS): By flooding the network with fake LLDP packets, an attacker could overwhelm network devices or management systems, potentially leading to a denial of service where legitimate network traffic is disrupted;

* Spoofing: An attacker could spoof LLDP packets to impersonate other network devices. Potentially, this might lead to incorrect network mappings or trust relationships being established with malicious devices;

* Lack of Authentication: LLDP does not include mechanisms for authenticating the source of LLDP messages, which means that devices accept LLDP information from any source as legitimate.

To mitigate these security issues, network administrators might implement several security measures, including:

* Disabling LLDP on ports where it is not needed, especially those facing untrusted networks;

* Using network segmentation and Access Control Lists (ACLs) to limit who can send and receive LLDP packets;

* Employing network monitoring and anomaly detection systems to identify unusual LLDP traffic patterns that may indicate an attack;

* Regularly updating and patching network devices to address known vulnerabilities that could be exploited through information gathered via LLDP.

8. Operational Considerations

This document has identified the need and enabling components for automating the management and control of multi-layer Service Providers' transport networks, combining the optical and microwave transport layer with the packet (IP/MPLS) layer to create a more efficient and scalable network infrastructure. This approach is particularly beneficial for Service Providers and large enterprises dealing with high bandwidth demands and looking for cost-effective ways to expand their networks. However, integrating these two traditionally separate network layers involves several operational considerations:

* Network Design and Capacity Planning: Deciding the degree of integration between the packet and optical layers is critical. Furthermore, this includes determining whether to pursue a loose integration (keeping layers distinct but coordinated) or a tight integration (combining layers more closely, potentially at the hardware level) coordinated via the MDSC. Accurate forecasting and planning will also be essential to ensure that the integrated ACTN infrastructure can handle future capacity demand without excessive over-provisioning;

* System Interoperability: Networks often comprise equipment from various vendors. Ensuring that packet and optical devices can interoperate seamlessly and the PNCs can manage them is crucial for a successful integration. The Service Provider must also check with the vendors to ensure they support the IETF-based technologies outlined in this document;

* Performance Monitoring: The integrated POI network will require comprehensive monitoring solutions that can provide visibility to the PNCs across both packet and optical layers. Identifying and diagnosing issues may become more complex with integrated layers. Telemetry data may also be required to collect lower-layer networking health and consider network and service performance. This topic is further discussed in [I-D.poidt-teas-actn-poi-assurance];

* Fault Management and Recovery: The POI networks should be resilient, including considerations for automatic protection switching and fast reroute mechanisms that span both layers. Fault isolation and recovery may become more challenging, as issues in one layer can have cascading effects on the other. Effective fault management strategies must be in place to quickly identify and rectify such issues. This topic is further discussed in [I-D.poidt-teas-actn-poi-assurance];

Specific Security Considerations are discussed in Section 7.

9. IANA Considerations

This document requires no IANA actions.

10. References

10.1. Normative References

[I-D.ietf-ccamp-client-signal-yang]
Zheng, H., Guo, A., Busi, I., Snitser, A., and C. Yu, "A YANG Data Model for Transport Network Client Signals", Work in Progress, Internet-Draft, draft-ietf-ccamp-client-signal-yang-16, 28 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-client-signal-yang-16>.

[I-D.ietf-ccamp-eth-client-te-topo-yang]
Yu, C., Zheng, H., Guo, A., Busi, I., Xu, Y., Zhao, Y., and X. Liu, "A YANG Data Model for Ethernet TE Topology", Work in Progress, Internet-Draft, draft-ietf-ccamp-eth-client-te-topo-yang-09, 14 April 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-eth-client-te-topo-yang-09>.

[I-D.ietf-ccamp-flexigrid-yang]
de Madrid, U. A., Burrero, D. P., King, D., Lee, Y., and H. Zheng, "A YANG Data Model for Flexi-Grid Optical Networks", Work in Progress, Internet-Draft, draft-ietf-ccamp-flexigrid-yang-18, 17 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-flexigrid-yang-18>.

[I-D.ietf-ccamp-otn-topo-yang]
Zheng, H., Busi, I., Liu, X., Belotti, S., and O. G. de Dios, "A YANG Data Model for Optical Transport Network Topology", Work in Progress, Internet-Draft, draft-ietf-ccamp-otn-topo-yang-20, 7 November 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-otn-topo-yang-20>.

[I-D.ietf-ccamp-otn-tunnel-model]
Zheng, H., Busi, I., Belotti, S., Lopez, V., and Y. Xu, "A YANG Data Model for Optical Transport Network (OTN) Tunnels and Label Switched Paths", Work in Progress, Internet-Draft, draft-ietf-ccamp-otn-tunnel-model-23, 4 June 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-otn-tunnel-model-23>.

[I-D.ietf-ccamp-wdm-tunnel-yang]
Guo, A., Belotti, S., Galimberti, G., de Madrid, U. A., and D. P. Burrero, "A YANG Data Model for WDM Tunnels", Work in Progress, Internet-Draft, draft-ietf-ccamp-wdm-tunnel-yang-05, 3 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-wdm-tunnel-yang-05>.

[I-D.ietf-teas-yang-l3-te-topo]
Liu, X., Bryskin, I., Beeram, V. P., Saad, T., Shah, H., and O. G. de Dios, "YANG Data Model for Layer 3 TE Topologies", Work in Progress, Internet-Draft, draft-ietf-teas-yang-l3-te-topo-18, 7 July 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-l3-te-topo-18>.

[I-D.ietf-teas-yang-path-computation]
Busi, I., Belotti, S., de Dios, O. G., Sharma, A., and Y. Shi, "A YANG Data Model for requesting path computation", Work in Progress, Internet-Draft, draft-ietf-teas-yang-path-computation-25, 11 August 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-path-computation-25>.

[I-D.ietf-teas-yang-sr-te-topo]
Liu, X., Bryskin, I., Beeram, V. P., Saad, T., Shah, H., and S. Litkowski, "YANG Data Model for SR and SR TE Topologies on MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-teas-yang-sr-te-topo-19, 4 July 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-sr-te-topo-19>.

[I-D.ietf-teas-yang-te]
Saad, T., Gandhi, R., Liu, X., Beeram, V. P., and I. Bryskin, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-38, 29 May 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-38>.

[I-D.ietf-teas-yang-te-mpls]
Saad, T., Gandhi, R., Liu, X., Beeram, V. P., and I. Bryskin, "A YANG Data Model for MPLS Traffic Engineering Tunnels", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-mpls-04, 26 May 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-mpls-04>.

[I-D.ietf-teas-yang-te-mpls-topology]
Busi, I., Guo, A., Liu, X., Saad, T., and R. Gandhi, "A YANG Data Model for MPLS-TE Topology", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-mpls-topology-03, 7 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-mpls-topology-03>.

[IEEE_802.1AB]
Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery", IEEE 802.1AB-2016 , March 2016, <https://ieeexplore.ieee.org/document/7433915>.

[IEEE_802.1AX]
Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and metropolitan area networks - Link Aggregation", IEEE 802.1AX-2014 , December 2014, <https://ieeexplore.ieee.org/document/7055197>.

[ITU-T_G.694.1]
International Telecommunication Union, "Spectral grids for WDM applications: DWDM frequency grid", ITU-T Recommendation G.694.1 , October 2020, <https://www.itu.int/rec/T-REC-G.694.1-202010-I>.

[RFC5212] Shiomoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, DOI 10.17487/RFC5212, July 2008, <https://www.rfc-editor.org/rfc/rfc5212>.

[RFC7923] Voit, E., Clemm, A., and A. Gonzalez Prieto, "Requirements for Subscription to YANG Datastores", RFC 7923, DOI 10.17487/RFC7923, June 2016, <https://www.rfc-editor.org/rfc/rfc7923>.

[RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <https://www.rfc-editor.org/rfc/rfc7950>.

[RFC7951] Lhotka, L., "JSON Encoding of Data Modeled with YANG", RFC 7951, DOI 10.17487/RFC7951, August 2016, <https://www.rfc-editor.org/rfc/rfc7951>.

[RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <https://www.rfc-editor.org/rfc/rfc8040>.

[RFC8342] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "Network Management Datastore Architecture (NMDA)", RFC 8342, DOI 10.17487/RFC8342, March 2018, <https://www.rfc-editor.org/rfc/rfc8342>.

[RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", RFC 8345, DOI 10.17487/RFC8345, March 2018, <https://www.rfc-editor.org/rfc/rfc8345>.

[RFC8346] Clemm, A., Medved, J., Varga, R., Liu, X., Ananthakrishnan, H., and N. Bahadur, "A YANG Data Model for Layer 3 Topologies", RFC 8346, DOI 10.17487/RFC8346, March 2018, <https://www.rfc-editor.org/rfc/rfc8346>.

[RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <https://www.rfc-editor.org/rfc/rfc8453>.

[RFC8525] Bierman, A., Bjorklund, M., Schoenwaelder, J., Watsen, K., and R. Wilton, "YANG Library", RFC 8525, DOI 10.17487/RFC8525, March 2019, <https://www.rfc-editor.org/rfc/rfc8525>.

[RFC8527] Bjorklund, M., Schoenwaelder, J., Shafer, P., Watsen, K., and R. Wilton, "RESTCONF Extensions to Support the Network Management Datastore Architecture", RFC 8527, DOI 10.17487/RFC8527, March 2019, <https://www.rfc-editor.org/rfc/rfc8527>.

[RFC8641] Clemm, A. and E. Voit, "Subscription to YANG Notifications for Datastore Updates", RFC 8641, DOI 10.17487/RFC8641, September 2019, <https://www.rfc-editor.org/rfc/rfc8641>.

[RFC8650] Voit, E., Rahman, R., Nilsen-Nygaard, E., Clemm, A., and A. Bierman, "Dynamic Subscription to YANG Events and Datastores over RESTCONF", RFC 8650, DOI 10.17487/RFC8650, November 2019, <https://www.rfc-editor.org/rfc/rfc8650>.

[RFC8795] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Gonzalez de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", RFC 8795, DOI 10.17487/RFC8795, August 2020, <https://www.rfc-editor.org/rfc/rfc8795>.

[RFC9094] Zheng, H., Lee, Y., Guo, A., Lopez, V., and D. King, "A YANG Data Model for Wavelength Switched Optical Networks (WSONs)", RFC 9094, DOI 10.17487/RFC9094, August 2021, <https://www.rfc-editor.org/rfc/rfc9094>.

10.2. Informative References

[I-D.ietf-ccamp-optical-impairment-topology-yang]
Beller, D., Le Rouzic, E., Belotti, S., Galimberti, G., and I. Busi, "A YANG Data Model for Optical Impairment-aware Topology", Work in Progress, Internet-Draft, draft-ietf-ccamp-optical-impairment-topology-yang-19, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-optical-impairment-topology-yang-19>.

[I-D.ietf-ccamp-optical-path-computation-yang]
Busi, I., Guo, A., and S. Belotti, "YANG Data Models for requesting Path Computation in WDM Optical Networks", Work in Progress, Internet-Draft, draft-ietf-ccamp-optical-path-computation-yang-06, 29 August 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-optical-path-computation-yang-06>.

[I-D.ietf-ccamp-transport-nbi-app-statement]
Busi, I., King, D., Zheng, H., and Y. Xu, "Transport Northbound Interface Applicability Statement", Work in Progress, Internet-Draft, draft-ietf-ccamp-transport-nbi-app-statement-17, 10 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-transport-nbi-app-statement-17>.

[I-D.ietf-ivy-network-inventory-yang]
Yu, C., Belotti, S., Bouquier, J., Peruzzini, F., and P. Bedard, "A Base YANG Data Model for Network Inventory", Work in Progress, Internet-Draft, draft-ietf-ivy-network-inventory-yang-08, 22 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-ivy-network-inventory-yang-08>.

[I-D.ietf-rtgwg-segment-routing-ti-lfa]
Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", Work in Progress, Internet-Draft, draft-ietf-rtgwg-segment-routing-ti-lfa-21, 12 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-21>.

[I-D.ietf-teas-actn-vn-yang]
Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B. Y. Yoon, "A YANG Data Model for Virtual Network (VN) Operations", Work in Progress, Internet-Draft, draft-ietf-teas-actn-vn-yang-29, 22 June 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-vn-yang-29>.

[I-D.ietf-teas-te-service-mapping-yang]
Lee, Y., Dhody, D., Fioccola, G., Wu, Q., Ceccarelli, D., and J. Tantsura, "Traffic Engineering (TE) and Service Mapping YANG Data Model", Work in Progress, Internet-Draft, draft-ietf-teas-te-service-mapping-yang-17, 29 January 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-service-mapping-yang-17>.

[I-D.ietf-teas-te-topology-profiles]
Busi, I., Liu, X., Bryskin, I., Saad, T., and O. G. de Dios, "Profiles for Traffic Engineering (TE) Topology Data Model and Applicability to non-TE Use Cases", Work in Progress, Internet-Draft, draft-ietf-teas-te-topology-profiles-03, 14 April 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-topology-profiles-03>.

[I-D.poidt-teas-actn-poi-assurance]
Busi, I., Bouquier, J., Peruzzini, F., Volpato, P., and P. Manna, "Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) for Packet Optical Integration (POI) service assurance", Work in Progress, Internet-Draft, draft-poidt-teas-actn-poi-assurance-05, 26 February 2025, <https://datatracker.ietf.org/doc/html/draft-poidt-teas-actn-poi-assurance-05>.

[RFC1930] Hawkinson, J. and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", BCP 6, RFC 1930, DOI 10.17487/RFC1930, March 1996, <https://www.rfc-editor.org/rfc/rfc1930>.

[RFC2991] Thaler, D. and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, DOI 10.17487/RFC2991, November 2000, <https://www.rfc-editor.org/rfc/rfc2991>.

[RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <https://www.rfc-editor.org/rfc/rfc4090>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/rfc/rfc5440>.

[RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, DOI 10.17487/RFC5623, September 2009, <https://www.rfc-editor.org/rfc/rfc5623>.

[RFC7424] Krishnan, R., Yong, L., Ghanwani, A., So, N., and B. Khasnabish, "Mechanisms for Optimizing Link Aggregation Group (LAG) and Equal-Cost Multipath (ECMP) Component Link Utilization in Networks", RFC 7424, DOI 10.17487/RFC7424, January 2015, <https://www.rfc-editor.org/rfc/rfc7424>.

[RFC8001] Zhang, F., Ed., Gonzalez de Dios, O., Ed., Margaria, C., Hartley, M., and Z. Ali, "RSVP-TE Extensions for Collecting Shared Risk Link Group (SRLG) Information", RFC 8001, DOI 10.17487/RFC8001, January 2017, <https://www.rfc-editor.org/rfc/rfc8001>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/rfc/rfc8231>.

[RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <https://www.rfc-editor.org/rfc/rfc8277>.

[RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <https://www.rfc-editor.org/rfc/rfc8281>.

[RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", RFC 8283, DOI 10.17487/RFC8283, December 2017, <https://www.rfc-editor.org/rfc/rfc8283>.

[RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <https://www.rfc-editor.org/rfc/rfc8309>.

[RFC8637] Dhody, D., Lee, Y., and D. Ceccarelli, "Applicability of the Path Computation Element (PCE) to the Abstraction and Control of TE Networks (ACTN)", RFC 8637, DOI 10.17487/RFC8637, July 2019, <https://www.rfc-editor.org/rfc/rfc8637>.

[RFC8751] Dhody, D., Lee, Y., Ceccarelli, D., Shin, J., and D. King, "Hierarchical Stateful Path Computation Element (PCE)", RFC 8751, DOI 10.17487/RFC8751, March 2020, <https://www.rfc-editor.org/rfc/rfc8751>.

[RFC9182] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <https://www.rfc-editor.org/rfc/rfc9182>.

[RFC9291] Boucadair, M., Ed., Gonzalez de Dios, O., Ed., Barguil, S., and L. Munoz, "A YANG Network Data Model for Layer 2 VPNs", RFC 9291, DOI 10.17487/RFC9291, September 2022, <https://www.rfc-editor.org/rfc/rfc9291>.

[RFC9408] Boucadair, M., Ed., Gonzalez de Dios, O., Barguil, S., Wu, Q., and V. Lopez, "A YANG Network Data Model for Service Attachment Points (SAPs)", RFC 9408, DOI 10.17487/RFC9408, June 2023, <https://www.rfc-editor.org/rfc/rfc9408>.

Appendix A. Additional Scenarios

A.1. OSS/Orchestration Layer

The OSS/Orchestration layer is a vital part of the architecture framework for a service provider:

* to abstract (through MDSC and PNCs) the underlying transport network complexity to the Business Systems Support layer;

* to coordinate NFV, Transport (e.g. IP, optical and microwave networks), Fixed Acess, Core and Radio domains enabling full automation of end-to-end services to the end customers;

* to enable catalogue-driven service provisioning from external applications (e.g. Customer Portal for Enterprise Business services), orchestrating the design and lifecycle management of these end-to-end transport connectivity services, consuming IP and/or optical transport connectivity services upon request.

As discussed in Section 2.1, in this document, the MDSC interfaces with the OSS/Orchestration layer and, therefore, it performs the functions of the Network Orchestrator, defined in [RFC8309].

The OSS/Orchestration layer requests the creation of a network service to the MDSC specifying its end-points (PEs and the interfaces towards the CEs) as well as the network service SLA and then proceeds to configuring accordingly the end-to-end customer service between the CEs in the case of an operator managed service.

A.1.1. MDSC NBI

As explained in Section 2, the OSS/Orchestration layer can request the MDSC to setup L2/L3VPN network services (with or without TE requirements).

Although the OSS/Orchestration layer interface is usually operator-specific, typically it would be using a RESTCONF/YANG interface with a more abstracted version of the MPI YANG data models used for network configuration (e.g. L3NM, L2NM).

Figure 12 shows an example of possible control flow between the OSS/ Orchestration layer and the MDSC to instantiate L2/L3 VPN network services, using the YANG data models under the definition in [I-D.ietf-teas-actn-vn-yang], [RFC9291], [RFC9182] and [I-D.ietf-teas-te-service-mapping-yang].

```
                  +-------------------------------------------+
                  |                                           |
                  |          OSS/Orchestration layer          |
                  |                                           |
                  +-----------------------+-------------------+
                                          |
                    1.VN    2. L2/L3NM &  |            ^
                      |          TSM      |            |
                      |           |       |            |
                      |           |       |            |
                      v           v       |      3. Update VN
                                          |
                  +-----------------------+-------------------+
                  |                                           |
                  |                  MDSC                     |
                  |                                           |
                  +-------------------------------------------+
```

Figure 12: Service Request Process

* The VN YANG data model, defined in [I-D.ietf-teas-actn-vn-yang], whose primary focus is the CMI, can also provide VN Service configuration from an orchestrated network service point of view when the L2/L3 VPN network service has TE requirements. However, this model is not used to setup L2/L3 VPN service with no TE requirements.

  - 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 the connectivity matrix with the VN member. 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).

  - The model also provides VN-level preference information (e.g., VN member diversity) and VN-level admin-status and operational-status.

* The L2NM and L3NM YANG data models, defined in [RFC9291] and [RFC9182], whose primary focus is the MPI, can also be used to provide L2VPN and L3VPN network service configuration from a orchestrated connectivity service point of view.

* The TE & Service Mapping YANG data model [I-D.ietf-teas-te-service-mapping-yang] provides TE-service mapping.

* TE-service mapping provides the mapping between a L2/L3 VPN instance and the corresponding VN instances.

* The TE-service mapping also provides the binding requirements as to how each L2/L3 VPN/VN instance is created concerning the underlay TE tunnels (e.g., whether they require a new and isolated set of TE underlay tunnels or not).

* Site mapping provides the site reference information across L2/L3 VPN Site ID, VN Access Point ID, and the LTP of the access link.

A.2. Multi-layer and Multi-domain Resiliency

A.2.1. Maintenance Window

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.

A.2.2. Router Port Failure

The focus is on client-side protection scheme between IP router and reconfigurable ROADM. 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 the IP router port side or on the 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:

1. LAG was defined between the IP routers at the two ends. MDSC, after checking that optical layer is fine between the two edge WDM nodes, triggers the WDM edge node re-configuration so that the IP router's back-up port with its associated muxponder port can reuse the WDM tunnel that was already in use previously by the failed IP 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 back-up port has been reconfigured to reuse the existing WDM tunnel and the 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.

2. If there is no LAG then the scenario is not clear since a IP 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 IP router and of the muxponder to re-use the same WDM tunnel as the one used originally for the failed IP router port. Once everything has been correctly configured, MDSC Global PCE could suggest to the operator to trigger a possible re-optimization of the back-up MPLS path to go back to the MPLS primary path through the back-up port of the IP router and the original WDM tunnel if overall cost, latency etc. is improved. However, in this scenario, there is a need for protection port PLUS back-up port in the IP router which does not lead to clear port savings.

A.3. Muxponders

The setup of a client connectivity service between two transponders is relatively clear and its implementation simple.

There is a one to one relationship between the tranponder's client and trunk (or DWDM) port. The client port bitrate determines the trunk port bit rate which will also determine the Baud-rate, the modulation format, the FEC etc.

The controller, when asked to set up a client connectivity service, needs to find a WDM tunnel suitable to comply the DWDM port parameters.

The setup of a client connectivity service between two muxponders is different since there is a one to many relationship between the muxponder's trunk (or DWDM) port and client ports. For example, there might be a 100Gb/s trunk port shared by ten 10GE client ports.

The controller, when asked to set a 10GE client connectivity service between two muxponder's client ports, needs first to check whether there is already an existing WDM tunnel between the two muxponders and then take different actions:

1. if the WDM tunnel already exists, the controller needs only to enable the 10GE client ports to establish the 10GE client connectivity service;

2. if the WDM tunnel does not exist, the controller has to first establish the WDM tunnel, finding a proper optical path matching the optical parameters of the two muxponders' trunk ports (e.g., an OTSi carrying an OTU4), and then enable the 10GE client ports to establish the 10GE client connectivity service.

Since multiple client connectivity services are sharing the same WDM tunnel, a multiplexing label shall be assigned to each client connectivity service. The multiplexing label can either be a standard label (e.g., an OTN timeslot) or a vendor-specific label. The multiplexing label can be either configurable (flexible configuration) or assigned by design to each muxponder's client port (fixed configuration). In the former case, any muxponder client port can be connected with any other client port of the peer muxponder (for example client port 1 on one muxponder can be connected with client port 5 on the peer muxponder) while in the latter case only client ports with the same port number can be connected (for example client port 2 on one muxponder can be connected only with client port 2 on the peer muxponder and not with any other client port).

In case of flexible configuration, since the two muxponders are under the control of the same O-PNC, the configuration of the multiplexing label, regardless of whether it is a standard or vendor-specific label, can be done by the O-PNC using mechanisms which are vendor-specific and outside the scope of this document. The MDSC can just request the O-PNC to setup a client connectivity service over a WDM tunnel.

In case of fixed configuration, the multiplexing label is assigned by the muxponder but the O-PNC and MDSC needs to be aware of the connectivity constraints to avoid try and fail.

It is worth noting that the current WSON and Flexi-grid topology models in [RFC9094] and [I-D.ietf-ccamp-flexigrid-yang] do not provide sufficient information to the MDSC about this connectivity constraint and this is identified as a gap.

Acknowledgments

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

The authors would like to thank Young Lee for his valuable input on the initial discusssions which have triggered this work as well as for his contribution to the first drafts of this document.

Previous versions of document were prepared using 2-Word-v2.0.template.dot.

This document was prepared using kramdown.

Contributors

Sergio Belotti
Nokia
Email: sergio.belotti@nokia.com

Gabriele Galimberti
Email: ggalimbe56@gmail.com

Yanlei Zheng
China Unicom
Email: zhengyanlei@chinaunicom.cn

Anton Snitser
Cisco
Email: asnizar@cisco.com

Washington Costa Pereira Correia
TIM Brasil
Email: wcorreia@timbrasil.com.br

Michael Scharf
Hochschule Esslingen - University of Applied Sciences
Email: michael.scharf@hs-esslingen.de

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

Paolo Volpato
Huawei
Email: paolo.volpato@huawei.com

Brent Foster
Cisco
Email: brfoster@cisco.com

Oscar Gonzalez de Dios
Telefonica
Email: oscar.gonzalezdedios@telefonica.com

Authors' Addresses

Fabio Peruzzini
FiberCop
Email: fabio.peruzzini@fibercop.com

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

Daniele Ceccarelli
Cisco
Email: daniele.ietf@gmail.com