

CCAMP Working Group
Internet Draft
Intended status: Informational

I. Busi
Huawei
D. King
Old Dog Consulting
H. Zheng
Huawei
Y. Xu
CAICT

Expires: September 2019

March 11, 2019

Transport Northbound Interface Applicability Statement
draft-ietf-ccamp-transport-nbi-app-statement-05

Status of this Memo

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

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

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

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

This Internet-Draft will expire on September 11, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

Transport network domains, including Optical Transport Network (OTN) and Wavelength Division Multiplexing (WDM) networks, are typically deployed based on a single vendor or technology platforms. They are often managed using proprietary interfaces to dedicated Element Management Systems (EMS), Network Management Systems (NMS) and increasingly Software Defined Network (SDN) controllers.

A well-defined open interface to each domain management system or controller is required for network operators to facilitate control automation and orchestrate end-to-end services across multi-domain networks. These functions may be enabled using standardized data models (e.g. YANG), and appropriate protocol (e.g., RESTCONF).

This document analyses the applicability of the YANG models being defined by the IETF (Traffic Engineering Architecture and Signaling (TEAS) and Common Control and Measurement Plane (CCAMP) WGs in particular) to support OTN single and multi-domain scenarios.

Table of Contents

1. Introduction.....	4
1.1. The scope of this document.....	4
1.2. Assumptions.....	5
2. Terminology.....	6
3. Conventions used in this document.....	7
3.1. Topology and traffic flow processing.....	7
3.2. JSON code.....	8

4. Scenarios Description.....	9
4.1. Reference Network.....	9
4.2. Topology Abstractions.....	13
4.3. Service Configuration.....	15
4.3.1. ODU Transit.....	16
4.3.2. EPL over ODU.....	17
4.3.3. Other OTN Clients Services.....	18
4.3.4. EVPL over ODU.....	19
4.4. Multi-function Access Links.....	20
4.5. Protection and Restoration Configuration.....	21
4.5.1. Linear Protection (end-to-end).....	21
4.5.2. Segmented Protection.....	23
4.6. Notification.....	23
4.7. Path Computation with Constraint.....	24
5. YANG Model Analysis.....	24
5.1. YANG Models for Topology Abstraction.....	25
5.1.1. Domain 1 Black Topology Abstraction.....	26
5.1.2. Domain 2 Black Topology Abstraction.....	30
5.1.3. Domain 3 White Topology Abstraction.....	31
5.1.4. Multi-domain Topology Merging.....	31
5.2. YANG Models for Service Configuration.....	33
5.2.1. ODU Transit Service.....	36
5.2.1.1. Single Domain Example.....	38
5.2.2. EPL over ODU Service.....	38
5.2.2.1. Single Domain Example.....	40
5.2.3. Other OTN Client Services.....	41
5.2.3.1. Single Domain Example.....	42
5.2.4. EVPL over ODU Service.....	42
5.3. YANG Models for Protection Configuration.....	44
5.3.1. Linear Protection (end-to-end).....	44
5.4. Notifications.....	44
5.5. Path Computation with Constraints.....	44
6. Security Considerations.....	44
7. IANA Considerations.....	45
8. References.....	45
8.1. Normative References.....	45
8.2. Informative References.....	46
9. Acknowledgments.....	47
Appendix A. Validating a JSON fragment against a YANG Model...	49
A.1. Manipulation of JSON fragments.....	49
A.2. Comments in JSON fragments.....	50
A.3. Validation of JSON fragments: DSDL-based approach.....	50
A.4. Validation of JSON fragments: why not using a XSD-based approach.....	51

Appendix B.	Detailed JSON Examples.....	52
B.1.	JSON Examples for Topology Abstractions.....	52
B.1.1.	JSON Code: mpil-otn-topology.json.....	52
B.2.	JSON Examples for Service Configuration.....	68
B.2.1.	JSON Code: mpil-odu2-service-config.json.....	68
B.2.2.	JSON Code: mpil-odu2-tunnel-config.json.....	72
B.2.3.	JSON Code: mpil-epl-service-config.json.....	72

1. Introduction

Transport of packet services are critical for a wide-range of applications and services, including data center and LAN interconnects, Internet service backhauling mobile backhaul and enterprise Carrier Ethernet services. These services are typically setup using stovepipe NMS and EMS platforms, often requiring propriety management platforms and legacy management interfaces. A clear goal of operators is to automate the setup of transport services across multiple transport technology domains.

A common open interface (API) to each domain controller and or management system is pre-requisite for network operators to control multi-vendor and multi-domain networks and also enable service provisioning coordination/automation. This can be achieved by using standardized YANG models, used together with an appropriate protocol (e.g., RESTCONF [RFC8040]).

This document analyses the applicability of the YANG models being defined by IETF (Traffic Engineering Architecture and Signaling (TEAS) and Common Control and Measurement Plane (CCAMP) WGs in particular) to support Optical Transport Networks (OTN) single and multi-domain scenarios.

1.1. The scope of this document

This document assumes a reference architecture, including interfaces, based on the Abstraction and Control of Traffic-Engineered Networks (ACTN), defined in [RFC8453].

The focus of this document is on the MPI (interface between the Multi Domain Service Coordinator (MDSC) and a Physical Network Controller (PNC), controlling a transport network domain).

It is worth noting that the same MPI analyzed in this document could be used between hierarchical MDSC controllers, as shown in Figure 4 of [RFC8453].

Detailed analysis of the CMI (interface between the Customer Network Controller (CNC) and the MDSC) as well as of the interface between service and network orchestrators are outside the scope of this document. However, some considerations and assumptions about the information are described when needed.

The relationship between the current IETF YANG models and the type of ACTN interfaces can be found in [ACTN-YANG]. Therefore, it considers the TE Topology YANG model defined in [TE-TOPO], with the OTN Topology augmentation defined in [OTN-TOPO] and the TE Tunnel YANG model defined in [TE-TUNNEL], with the OTN Tunnel augmentation defined in [OTN-TUNNEL]. It also considers the Ethernet Client Topology augmentation defined in [CLIENT-TOPO] as well as the Client Signal YANG models defined in [CLIENT-SIGNAL] for both Transparent and Ethernet clients.

The ONF Technical Recommendations for Functional Requirements for the transport API in [ONF TR-527] and the ONF transport API multi-domain examples in [ONF GitHub] have been considered as input for defining the reference scenarios analyzed in this document.

1.2. Assumptions

This document is making the following assumptions:

1. The MDSC can request, at the MPI, a PNC to setup a Transit Tunnel Segment using the TE Tunnel YANG model: in this case, since the endpoints of the E2E Tunnel are outside the domain controlled by that PNC, the MDSC would not specify any source or destination TE Tunnel Termination Point (TTP), i.e., it would leave empty the source, destination, src-tp-id and dst-tp-id attributes of the TE tunnel instance, and it would use the explicit-route-object (ERO) or route-object-include-exclude list to specify the ingress and egress links for each path of the Transit Tunnel Segment.
2. Each PNC provides to the MDSC, at the MPI, the list of available timeslots on the inter-domain links using the TE Topology YANG model and OTN Topology augmentation. The TE Topology YANG model in [TE-TOPO] is being updated to report the label set information.

See section 1.7 of [TE-TUTORIAL] for more details.

This document has made the following assumptions:

1. The topology information for the Ethernet Client access links is modelled using the YANG model defined in [CLIENT-TOPO];
2. The topology information for the OTN and Transparent Client access links is modelled using the YANG model defined in [OTN-TOPO];
3. The mapping information for Ethernet and Transparent Client signals are modelled using the YANG model defined in [CLIENT-SIGNAL].

Finally, the Network Elements (NEs) described in the scenarios used in the document are using ODU switching. It is assumed that the ODU links are pre-configured and using mechanisms such as WDM wavelength, which are outside the scope of this document.

2. Terminology

Domain: A domain as defined by [RFC4655] is "any collection of network elements within a common sphere of address management or path computation responsibility". Specifically, within this document we mean a part of an operator's network that is under common management (i.e., under shared operational management using the same instances of a tool and the same policies). Network elements will often be grouped into domains based on technology types, vendor profiles, and geographic proximity.

E-LINE: Ethernet Line

EPL: Ethernet Private Line

EVPL: Ethernet Virtual Private Line

OTN: Optical Transport Network

Service: A service in the context of this document can be considered as some form of connectivity between customer sites across the network operator's network [RFC8309].

Service Model: As described in [RFC8309] it describes a service and the parameters of the service in a portable way that can be used uniformly and independent of the equipment and operating environment.

UNI: User Network Interface

MDSC: Multi-Domain Service Coordinator

CNC: Customer Network Controller

PNC: Provisioning Network Controller

3. Conventions used in this document

3.1. Topology and traffic flow processing

The traffic flow between different nodes is specified as an ordered list of nodes, separated with commas, indicating within the brackets the processing within each node:

<node> [<processing>]{, <node> [<processing>]}

The order represents the order of traffic flow being forwarded through the network.

The processing can be just switching at a given layer "[switching]" or also having an adaptation of a client layer into a server layer "[client -> server]" or [client -> (server)], depending on whether the node is switching in the client or in the server layer.

For example, the following traffic flow:

R1 [(PKT) -> ODU2], S3 [(ODU2)], S5 [(ODU2)], S6 [(ODU2)],
R3 [ODU2 -> (PKT)]

Node R1 is switching at the packet (PKT) layer and mapping packets into an ODU2 before transmission to node S3. Nodes S3, S5 and S6 are switching at the ODU2 layer: S3 sends the ODU2 traffic to S5 which then sends it to S6 which finally sends to R3. Node R3 terminates the ODU2 from S6 before switching at the packet (PKT) layer.

The paths of working and protection transport entities are specified as an ordered list of nodes, separated with commas:

```
<node> {, <node>}
```

The order represents the order of traffic flow being forwarded through the network in the forward direction. In the case of bidirectional paths, the forward and backward directions are selected arbitrarily, but the convention is consistent between working/protection path pairs, as well as across multiple domains.

3.2. JSON code

This document provides some detailed JSON code examples to describe how the YANG models being developed by the IETF (TEAS and CCAMP WG in particular) may be used. The scenario examples are provided using JSON to facilitate readability.

Different objects need to have an identifier. The convention used to create mnemonic identifiers is to use the object name (e.g., S3 for node S3), followed by its type (e.g., NODE), separated by an "-", followed by "-ID". For example, the mnemonic identifier for node S3 would be S3-NODE-ID.

The JSON language does not support the insertion of comments that have been instead found to be useful when writing the examples. This document will insert comments into the JSON code as JSON name/value pair with the JSON name string starting with the "//" characters. For example, when describing the example of a TE Topology instance representing the ODU Abstract Topology exposed by the Transport PNC, the following comment has been added to the JSON code:

```
"// comment": "ODU Abstract Topology @ MPI",
```

The JSON code examples provided in this document have been validated against the YANG models following the validation process described in Appendix A, which would not consider the comments.

In order to have successful validation of the examples, some numbering scheme has been defined to assign identifiers to the different entities which would pass the syntax checks. In that case, to simplify the reading, another JSON name/value pair formatted as a comment and using the mnemonic identifiers is also provided. For example, the identifier of node S3 (S3-NODE-ID) has been assumed to

be "10.0.0.3" and would be shown in the JSON code example using the two JSON name/value pair:

```
"// te-node-id": "S3-NODE-ID",  
  
"te-node-id": "10.0.0.3",
```

The first JSON name/value pair will be automatically removed in the first step of the validation process while the second JSON name/value pair will be validated against the YANG model definitions.

4. Scenarios Description

4.1. Reference Network

The physical topology of the reference network is shown in Figure 1. It represents an OTN network composed of three transport network domains which provide transport connectivity services to an IP customer network through eight access links:




Figure 1 - Reference network

This document assumes that all the transport network switching nodes S_i are capable of switching in the electrical domain (ODU switching) and that all the S_i - S_j OTN links within the transport network (intra-domain or inter-domain) are 100G links while the access R_i - S_j links are 10G links.

It is also assumed that, within the transport network, the physical/optical interconnections supporting the S_i - S_j OTN links (up to the OTU4 trail), are pre-configured using mechanisms which are

outside the scope of this document and are not exposed at the MPIs to the MDSC.

Different technologies can be used on the access links (e.g., Ethernet, STM-N and OTU). Section 4.3 provides more details about the different assumptions on the access links for different services types and section 4.4 describes the control of access links which can support different technology configuration (e.g., STM-64, 10GE or OTU2) depending on the type of service being configured (multi-function access links).

The transport domain control architecture, shown in Figure 2, follows the ACTN architecture and framework document [RFC8453], and functional components:




Figure 2 - Controlling Hierarchies

The ACTN framework facilitates the detachment of the network and service control from the underlying technology. It helps the customer define the network as desired by business needs. Therefore, care must be taken to keep a minimal level of dependency on the CMI (or no dependency at all) with respect to the network domain technologies. The MPI instead requires some specialization according to the domain technology.

The control interfaces within the scope of this document are the three MPIs shown in Figure 2.

It is worth noting that the split of functionality at the MPI in the ACTN architecture between the MDSC and the PNCs is equivalent/analogous to the split of functionality which is assumed for the ONF T-API interface when used between a multi-domain controller and domain controllers, as described in the ONF T-API multi-domain use cases [ONF TR-527], as well as at the MEF PRESTO interface between the Service Orchestration Functionality (SOF) and the Infrastructure Control and Management (ICM) in the MEF LSO Architecture [MEF 55].

This document does not make any assumption about the control architecture of the customer IP network: in line with [RFC8453], the CNC is just a functional component within the customer control architecture which is capable of requesting, at the CMI, transport connectivity between IP routers, when needed.

The CNC can request transport connectivity services between IP routers which can be attached to different transport domains (e.g., between R1 and R8 in Figure 1) or to the same transport domain (e.g., between R1 and R3 in Figure 1). Since the CNC is not aware of the transport network controlling hierarchy, the mechanisms used by the CNC to request, at the CMI, transport connectivity services are independent on whether the service request is single-domain or multi-domain.

It is assumed that the CMI allows the CNC to provide all the information that is required by the MDSC to understand the connectivity service request and to decide the network configurations to be requested, at the MPIs, to its underlying PNCs to support the requested connectivity service.

When a single-domain service is requested by the CNC at the CMI (e.g., between R1 and R3 in Figure 1), the MDSC can follow the same procedures, described above for the multi-domain service, and decide the network configuration to request only at the MPI of the PNC controlling that domain (e.g., MPI1 of PNC1 in Figure 2).

4.2. Topology Abstractions

Abstraction provides a selective method for representing connectivity information within a domain. There are multiple methods

to abstract a network topology. This document assumes the abstraction method defined in [RFC7926]:

"Abstraction is the process of applying the policy to the available TE information within a domain, to produce selective information that represents the potential ability to connect across the domain. Thus, abstraction does not necessarily offer all possible connectivity options, but presents a general view of potential connectivity according to the policies that determine how the domain's administrator wants to allow the domain resources to be used."

[RFC8453] Provides the context of topology abstraction in the ACTN architecture and discusses a few alternatives for the abstraction methods for both packet and optical networks. This is an important consideration since the choice of the abstraction method impacts protocol design and the information it carries. According to [RFC8453], there are three types of topology:

- o White topology: This is a case where the PNC provides the actual network topology to the MDSC without any hiding or filtering. In this case, the MDSC has the full knowledge of the underlying network topology;
- o Black topology: The entire domain network is abstracted as a single virtual node with the access/egress links without disclosing any node internal connectivity information;
- o Grey topology: This abstraction level is between black topology and white topology from a granularity point of view. This is an abstraction of TE tunnels for all pairs of border nodes. We may further differentiate from a perspective of how to abstract internal TE resources between the pairs of border nodes:
 - Grey topology type A: border nodes with TE links between them in a full mesh fashion;
 - Grey topology type B: border nodes with some internal abstracted nodes and abstracted links.

Each PNC should provide the MDSC with a network topology abstractions hiding the internal details of the physical domain network topology controlled by the PNC and this abstraction is independent of the abstractions provided by other PNCs. Therefore it

is possible that different PNCs provide different types of topology abstractions and each MPI operates on the abstract topology regardless of, and independently from, the type of abstraction provided by the PNC.

To analyze how the MDSC can operate on abstract topologies independently from the topology abstraction provided by each PNC and, therefore, that different PNCs can provide different topology abstractions, that the following examples are assumed:

- o PNC1 and PNC2 provide black topology abstractions which expose at MPI1, and MPI2 respectively, a single virtual node (representing the whole network domain 1, and domain 2 respectively).
- o PNC3 provides a white topology abstraction which exposes at MPI3 all the physical nodes and links within network domain 3.

The MDSC should be capable of stitching together the abstracted topologies provided by each PNC to build its own view of the multi-domain network topology. The process may require suitable oversight, including administrative configuration and trust models, a method of how to achieve this is out of scope for this document.

The MDSC can also provide topology abstraction of its own view of the multi-domain network topology at its CMIs depending on the customers' needs: it can provide different types of topology abstractions at different CMIs. Analyzing the topology abstractions provided by the MDSC to its CMIs is outside the scope of this document.

4.3. Service Configuration

In the following scenarios, it is assumed that the CNC is capable of requesting service connectivity from the MDSC to support IP routers connectivity.

The type of services could depend on the type of physical links (e.g. OTN link, ETH link or SDH link) between the routers and transport network.

The control of different adaptations inside IP routers, R_i (PKT \rightarrow foo) and R_j (foo \rightarrow PKT), are assumed to be performed by means that are not under the control of, and not visible to, the MDSC nor to

the PNCs. Therefore, these mechanisms are outside the scope of this document.

4.3.1. ODU Transit

The physical links interconnecting the IP routers and the transport network can be 10G OTN links.

In this case, it is assumed that the physical/optical interconnections below the ODU layer (up to the OTU2 trail) are pre-configured using mechanisms which are outside the scope of this document and not exposed at the MPIS between the PNCs and the MDSC.

For simplicity of the description, it is also assumed that these interfaces are not channelized (i.e., they can only support one ODU2).

To setup a 10Gb IP link between R1 and R8, an ODU2 end-to-end connection needs to be created, passing through transport nodes S3, S1, S2, S31, S33, S34, S15 and S18 which belong to different PNC domains (multi-domain service request):

```
R1 [(PKT) -> ODU2], S3 [(ODU2)], S1 [(ODU2)], S2 [(ODU2)],  
S31 [(ODU2)], S33 [(ODU2)], S34 [(ODU2)],  
S15 [(ODU2)], S18 [(ODU2)], R8 [ODU2 -> (PKT)]
```

The MDSC understands that it needs to establish an ODU2 transit service between the access links on S3 and S18, which belong to different PNC domains (multi-domain service request). It also decides the network configurations to request, at the MPIS, to its underlying PNCs, to coordinate the setup of a multi-domain ODU2 segment connection between the access links on S3 and S18.

To setup of a 10Gb IP link between R1 and R3, an ODU2 end-to-end connection needs to be created, passing through transport nodes S3, S5 and S6 which belong to the same PNC domain (single-domain service request):

```
R1 [(PKT) -> ODU2], S3 [(ODU2)], S5 [(ODU2)], S6 [(ODU2)],  
R3 [ODU2 -> (PKT)]
```

As described in section 4.1, the mechanisms used by the CNC at the CMI are independent on whether the service request is single-domain service or multi-domain.

The MDSC can understand that it needs to setup an ODU2 transit service between the access links on S3 and S6, which belong to the same PNC domain (single-domain service request) and it decides the proper network configuration to request PNC1.

4.3.2. EPL over ODU

The physical links interconnecting the IP routers and the transport network can be 10G Ethernet physical links (10GE).

In this case, it is assumed that the Ethernet physical interfaces (up to the MAC layer) are pre-configured using mechanisms which are outside the scope of this document and not exposed at the MPis between the PNCs and the MDSC.

To setup a 10Gb IP link between R1 and R8, an EPL service needs to be created, supported by an ODU2 end-to-end connection, between transport nodes S3 and S18, passing through transport nodes S1, S2, S31, S33, S34 and S15, which belong to different PNC domains (multi-domain service request):

```
R1 [(PKT) -> ETH], S3 [ETH -> (ODU2)], S1 [(ODU2)],  
S2 [(ODU2)], S31 [(ODU2)], S33 [(ODU2)], S34 [(ODU2)],  
S15 [(ODU2)], S18 [(ODU2) -> ETH], R8 [ETH -> (PKT)]
```

The MDSC understands that it needs to setup an EPL service between the access links on S3 and S18, which belong to different PNC domains (multi-domain service request). It also decides the network configurations to request, at the MPis, to its underlying PNCs, to coordinate the setup of an end-to-end ODU2 connection between the nodes S3 and S8, including the configuration of the adaptation functions inside these edge nodes, such as S3 [ETH -> (ODU2)] and S18 [(ODU2) -> ETH].

To setup a 10Gb IP link between R1 and R2, an EPL service needs to be created, supported by an ODU2 end-to-end connection between transport nodes S3 and S6, passing through the transport node S5, which belong to the same PNC domain (single-domain service request):

```
R1 [(PKT) -> ETH], S3 [ETH -> (PKT)] S5 [(ODU2)],  
S6 [(ODU2) -> ETH], R2 [ETH -> (PKT)]
```

As described in section 4.1, the mechanisms used by the CNC at the CMI are independent on whether the service request is single-domain service or multi-domain.

Based on the assumption above, in this case, the MDSC can understand that it needs to setup an EPL service between the access links on S3 and S6, which belong to the same PNC domain (single-domain service request) and it decides the proper network configuration to request PNC1.

4.3.3. Other OTN Clients Services

[ITU-T G.709] defines mappings of different Transparent Client layers into ODU. Most of them are used to provide Private Line services over an OTN transport network supporting a variety of types of physical access links (e.g., Ethernet, SDH STM-N, Fibre Channel, InfiniBand, etc.) interconnecting the IP routers and the transport network.

In order to setup a 10Gb IP link between R1 and R8 using, with for example SDH physical links between the IP routers and the transport network, an STM-64 Private Line service needs to be created, supported by an ODU2 end-to-end connection, between transport nodes S3 and S18, passing through transport nodes S1, S2, S31, S33, S34 and S15, which belong to different PNC domains (multi-domain service request):

```
R1 [(PKT) -> STM-64], S3 [STM-64 -> (ODU2)], S1 [(ODU2)],  
S2 [(ODU2)], S31 [(ODU2)], S33 [(ODU2)], S34[(ODU2)],  
S15 [(ODU2)], S18 [(ODU2) -> STM-64], R8 [STM-64 -> (PKT)]
```

As already described (section 4.1) CNC provides the essential information to permit the MDSC to understand which type of service is needed, in this case, an STM-64 Private Line service between the access links on S3 and S8, and it also decides the network configurations, including the configuration of the adaptation functions inside these edge nodes, such as S3 [STM-64 -> (ODU2)] and S18 [(ODU2) -> STM-64].

To setup a 10Gb IP link between R1 and R3), an STM-64 Private Line service needs to be created between R1 and R3 (single-domain service request):

```
R1 [(PKT) -> STM-64], S3[STM-64 -> (ODU2)], S5 [(ODU2)],  
S6 [(ODU2) -> STM-64], R3[STM-64 -> (PKT)]
```

As described in section 4.1, the mechanisms used by the CNC at the CMI are independent on whether the service request is single-domain service or multi-domain.

Based on the assumption above, in this case, the MDSC can understand that it needs to setup an STM-64 Private Line service between the access links on S3 and S6, which belong to the same PNC domain (single-domain service request) and it decides the proper network configuration to request PNC1.

4.3.4. EVPL over ODU

When the physical links interconnecting the IP routers and the transport network are Ethernet physical links, it is also possible that different Ethernet services (e.g., EVPL) can share the same physical access link using different VLANs.

As described in section 4.3.2, it is assumed that the Ethernet physical interfaces (up to the MAC layer) are pre-configured.

To setup two 1Gb IP links between R1 to R2 and between R1 and R8, two EVPL services need to be created, supported by two ODU0 end-to-end connections:

```
R1 [(PKT) -> VLAN], S3 [VLAN -> (ODU0)], S5 [(ODU0)],  
S6 [(ODU0) -> VLAN], R2 [VLAN -> (PKT)]
```

```
R1 [(PKT) -> VLAN], S3[VLAN -> (ODU0)], S1 [(ODU0)],  
S2 [(ODU0)], S31 [(ODU0)], S33 [(ODU0)], S34 [(ODU0)],  
S15 [(ODU0)], S18 [(ODU0) -> VLAN], R8[VLAN -> (PKT)]
```

It is worth noting that the first EVPL service is required between access links which belong to the same PNC domain (single-domain service request) while the second EVPL service is required between access links which belong to different PNC domains (multi-domain service request).

Since the two EVPL services are sharing the same Ethernet physical link between R1 and S3, different VLAN IDs are associated with different EVPL services: for example, VLAN IDs 10 and 20 respectively.

Based on the assumptions described in section 4.3.2, the CNC requests at the CMI the MDSC to setup these EVPL services and the MDSC understands the network configurations to request to the underlying PNCs, as described in section 4.3.2.

4.4. Multi-function Access Links

Some physical links interconnecting the IP routers and the transport network can be configured in different modes, e.g., as OTU2 trail or STM-64 or 10GE physical links.

This configuration can be done a-priori by means which are outside the scope of this document. In this case, these links will appear at the MPI as links supporting only one mode (depending on the a-priori configuration) and will be controlled at the MPI as discussed in section 4.3: for example, a 10G multi-function access link can be pre-configured as an OTU2 trail (section 4.3.1), a 10GE physical link (section 4.3.2) or an STM-64 physical link (section 4.3.3).

It is also possible not to configure these links a-priori and let the MDSC (or, in case of a single-domain service request, the PNC) decide how to configure these links, based on the service configuration.

For example, if the physical link between R1 and S3 is a multi-functional access link while the physical links between R4 and S6 and between R8 and S18 are STM-64 and 10GE physical links respectively, it is possible to configure either an STM-64 Private Line service between R1 and R4 or an EPL service between R1 and R8.

The traffic flow between R1 and R4 can be summarized as:

```
R1 [(PKT) -> STM-64], S3 [STM-64 -> (ODU2)], S5 [(ODU2)],  
S6 [(ODU2) -> STM-64], R4 [STM-64 -> (PKT)]
```

The traffic flow between R1 and R8 can be summarized as:

```
R1 [(PKT) -> ETH], S3 [ETH -> (ODU2)], S1 [(ODU2)],  
S2 [(ODU2)], S31 [(ODU2)], S33 [(ODU2)], S34 [(ODU2)],  
S15 [(ODU2)], S18 [(ODU2) -> ETH], R8 [ETH -> (PKT)]
```

The CNC is capable to request at the CMI the setup either an STM-64 Private Line service, between R1 and R4, or an EPL service, between R1 and R8.

The MDSC, based on the service being request, decides the network configurations to request, at the MPis, to its underlying PNCs, to coordinate the setup of an end-to-end ODU2 connection, either between nodes S3 and S6, or between nodes S3 and S18, including the configuration of the adaptation functions on these edge nodes, and in particular whether the multi-function access link between R1 and S3 should operate as an STM-64 or as a 10GE physical link.

4.5. Protection and Restoration Configuration

Protection switching provides a pre-allocated survivability mechanism, typically provided via linear protection methods and would be configured to operate as 1+1 unidirectional (the most common OTN protection method), 1+1 bidirectional or 1:n bidirectional. This ensures fast and simple service survivability.

Restoration methods would provide the capability to reroute and restore connectivity traffic around network faults, without the network penalty imposed with dedicated 1+1 protection schemes.

This section describes only services which are protected with linear protection. The description of services using dynamic restoration is outside the scope of this document.

The MDSC needs to be capable of deciding the network configuration to request different PNCs to coordinate the protection switching configuration to support protected connectivity services described in section 4.3.

Since in these service examples, switching within the transport network domain is performed only in the OTN ODU layer. It may also be assumed that protection switching within the transport network domain is provided at the OTN ODU layer.

4.5.1. Linear Protection (end-to-end)

In order to protect the connectivity services described in section 4.3 from failures within the OTN multi-domain transport network, the MDSC can decide to request its underlying PNCs to configure ODU2 linear protection between the access nodes (e.g., nodes S3 and S18 for the services setup between R1 and R8).

It is assumed that the OTN linear protection is configured to with 1+1 unidirectional protection switching type, as defined in [ITU-T G.808.1] and [ITU-T G.873.1], as well as in [RFC4427].

In these scenarios, a working transport entity and a protection transport entity, as defined in [ITU-T G.808.1], (or a working LSP and a protection LSP, as defined in [RFC4427]) should be configured in the data plane.

Two cases can be considered:

- o In one case, the working and protection transport entities pass through the same PNC domains:

Working transport entity: S3, S1, S2,
S31, S33, S34,
S15, S18

Protection transport entity: S3, S4, S8,
S32,
S12, S17, S18

- o In another case, the working and protection transport entities can pass through different PNC domains:

Working transport entity: S3, S5, S7,
S11, S12, S17, S18

Protection transport entity: S3, S1, S2,
S31, S33, S34,
S15, S18

The PNCs should be capable to report to the MDSC which is the active transport entity, as defined in [ITU-T G.808.1], in the data plane.

Given the fast dynamic of protection switching operations in the data plane (50ms recovery time), this reporting is not expected to be in real-time.

It is also worth noting that with unidirectional protection switching, e.g., 1+1 unidirectional protection switching, the active transport entity may be different in the two directions.

4.5.2. Segmented Protection

To protect the connectivity services defined in section 4.3 from failures within the OTN multi-domain transport network, the MDSC can decide to request its underlying PNCs to configure ODU2 linear protection between the edge nodes of each domain.

For example, MDSC can request PNC1 to configure linear protection between its edge nodes S3 and S2:

Working transport entity: S3, S1, S2

Protection transport entity: S3, S4, S8, S2

MDSC can also request PNC2 to configure linear protection between its edge nodes S15 and S18:

Working transport entity: S15, S18

Protection transport entity: S15, S12, S17, S18

MDSC can also request PNC3 to configure linear protection between its edge nodes S31 and S34:

Working transport entity: S31, S33, S34

Protection transport entity: S31, S32, S34

4.6. Notification

To realize the topology update, service update and restoration function, following notification type should be supported:

1. Object create
2. Object delete
3. Object state change
4. Alarm

There are three types of topology abstraction type defined in section 4.2, the notification should also be abstracted. The PNC and MDSC should coordinate together to determine the notification

policy, such as when an intra-domain alarm occurred, the PNC may not report the alarm but the service state change notification to the MDSC.

4.7. Path Computation with Constraint

It is possible to define constraints to be taken into account during path computation procedures (e.g., IRO/XRO).

For example, the CNC can request, at the CMI, an ODU transit service, as described in section 4.3.1, between R1 and R8 with the constraint to pass through the link from S2 to S31 (IRO), such that a qualified path could be:

```
R1 [(PKT) -> ODU2], S3 [(ODU2)], S1 [(ODU2)], S2 [(ODU2)],  
S31 [(ODU2)], S33 [(ODU2)], S34 [(ODU2)],  
S15 [(ODU2)], S18 [(ODU2)], R8 [ODU2 -> (PKT)]
```

If the CNC instead requested to pass through the link from S8 to S12, then the above path would not be qualified, while the following would be:

```
R1 [(PKT) -> ODU2], S3[(ODU2)], S1 [(ODU2)], S2[(ODU2)],  
S8 [(ODU2)], S12[(ODU2)], S15 [(ODU2)], S18[(ODU2)], R8 [ODU2 ->  
(PKT)]
```

The mechanisms used by the CNC to provide path constraints at the CMI are outside the scope of this document. It is assumed that the MDSC can understand these constraints and take them into account in its path computation procedures (which would decide at least which domains and inter-domain links) and in the path constraints to provide to its underlying PNCs, to be taken into account in the path computation procedures implemented by the PNCs (with a more detailed view of topology).

5. YANG Model Analysis

This section analyses how the IETF YANG models can be used at the MPIs, between the MDSC and the PNCs, to support the scenarios described in section 4.

The YANG models described in [ACTN-YANG] are assumed to be used at the MPI.

Section 5.1 describes the different topology abstractions provided to the MDSC by each PNC via its own MPI.

Section 5.2 describes how the MDSC can request different PNCs, via their own MPIs, the network configuration needed to setup the different services described in section 4.3.

Section 5.3 describes how the protection scenarios can be deployed, including end-to-end protection and segment protection, for both intra-domain and inter-domain scenario.

5.1. YANG Models for Topology Abstraction

Each PNC reports its respective OTN abstract topology to the MDSC, as described in section 4.2, using the Topology YANG models, defined in [RFC8345], with the TE Topology YANG augmentations, provided in [TE-TOPO], and the OTN technology-specific YANG augmentations, defined in [OTN-TOPO].

The [OTN-TOPO] model allows reporting within the OTN abstract topology also the access links which are capable of supporting the transparent client layers, defined in section 4.3.3 and in [CLIENT-SIGNAL]. These links can also be multi-function access links that can be configured as a transparent client physical links (e.g., STM-64 physical link) as an OTUk trail.

In order to support the EPL and EVPL services, described in sections 4.3.2 and 4.3.4, the access links, which are capable to be configured as Ethernet physical links, are reported by each PNC within its respective Ethernet abstract topology, using the Topology YANG models, defined in [RFC8345], with the TE Topology YANG augmentations, defined in [TE-TOPO], and the Ethernet client technology-specific YANG augmentations, defined in [CLIENT-TOPO]. These links can also be multi-function access links that can be configured as an Ethernet physical link or as an OTUk trail and/or as transparent client physical links (e.g., STM-64 physical link). In this case, these physical access links are represented in both the OTN and Ethernet abstract topologies.

It is worth noting that in the network scenarios analyzed in this document (where switching is performed only in the ODU layer), the Ethernet abstract topologies reported by the PNCs describes only the Ethernet client access links: no Ethernet TE switching capabilities are reported in these Ethernet abstract topologies.

5.1.1.1. Domain 1 Black Topology Abstraction

PNC1 provides the required black topology abstraction, as described in section 4.2. It exposes at MPI1 to the MDSC, two TE Topology instances with only one TE node each.

The first TE Topology instance reports the domain 1 OTN abstract topology view (MPI1 OTN Topology), using the OTN technology-specific augmentations [OTN-TOPO], with only one abstract TE node (i.e., AN1) and only inter-domain and access abstract TE links (which represent the inter-domain physical links and the access physical links which can support ODU and/or transparent client layers), as shown in Figure 3 below.




Figure 3 - OTN Abstract Topology exposed at MPI1 (MPI1 OTN Topology)

The second TE Topology instance reports the domain 1 Ethernet abstract topology view (MPI1 ETH Topology), using the Ethernet technology-specific augmentations [CLIENT-TOPO], with only one abstract TE node (i.e., AN1) and only access abstract TE links (which represent the access physical links which can support Ethernet client layers), as shown in Figure 4 below.




Figure 4 - ETH Abstract Topology exposed at MPI1 (MPI1 ETH Topology)

As described in section 4.1, it is assumed that the OTU4 trails on the inter-domain physical links (e.g., the link between S2 and S31) are pre-configured and exposed as external TE Links, within the MPI1 OTN topology (e.g., the external TE Link terminating on AN1-7 TE Link Termination Point (LTP) abstracting the OTU4 trail between S2 and S31).

In order to analyze the service scenarios of sections 4.3 and 4.4:

- o the access link between S3 and R1 is assumed to be a multi-function access link, which can be configured as an OTU2 trail or as an STM-64 or a 10GE physical link;
- o the access link between S6 and R2 is assumed to be pre-configured as a 10GE physical link, up to the MAC layer;
- o the access link between S6 and R3 is assumed to be a multi-function access link which can be configured as an OTU2 trail or as an STM-64 physical link;
- o the access link connected to router R4 is assumed to be pre-configured as an STM-64 physical link.

Therefore PNC1 exports at MPI1 the following external TE Links, within the MPI1 OTN topology:

- o two abstract TE Links, terminating on LTP AN1-1 and AN1-2 respectively, abstracting the physical access link between S3 and R1 and the access link between S6 and R3 respectively, reporting that they can support STM-64 client signals as well as ODU2 connections;
- o one abstract TE Link, terminating on LTP AN1-3, abstracting the physical access link between S6 and R4, reporting that it can support STM-64 client signals but no ODU2 connections.

The information about the 10GE access link between S6 and R2 as well as the fact that the access link between S3 and R1 can also be configured as a 10GE link cannot be exposed by PNC1 within the MPI1 OTN topology.

Therefore, PNC1 exports at MPI1, within the MPI1 ETH topology, two abstract TE Links, terminating on LTP AN1-1 and AN1-8 respectively, abstracting the physical access link between S3 and R1 and the access link between S6 and R2 respectively, reporting that they can support Ethernet client signal with port-based and VLAN-based classifications. The PNC1 native topology would represent the physical network topology of the domain controlled by the PNC, as shown in Figure 5.

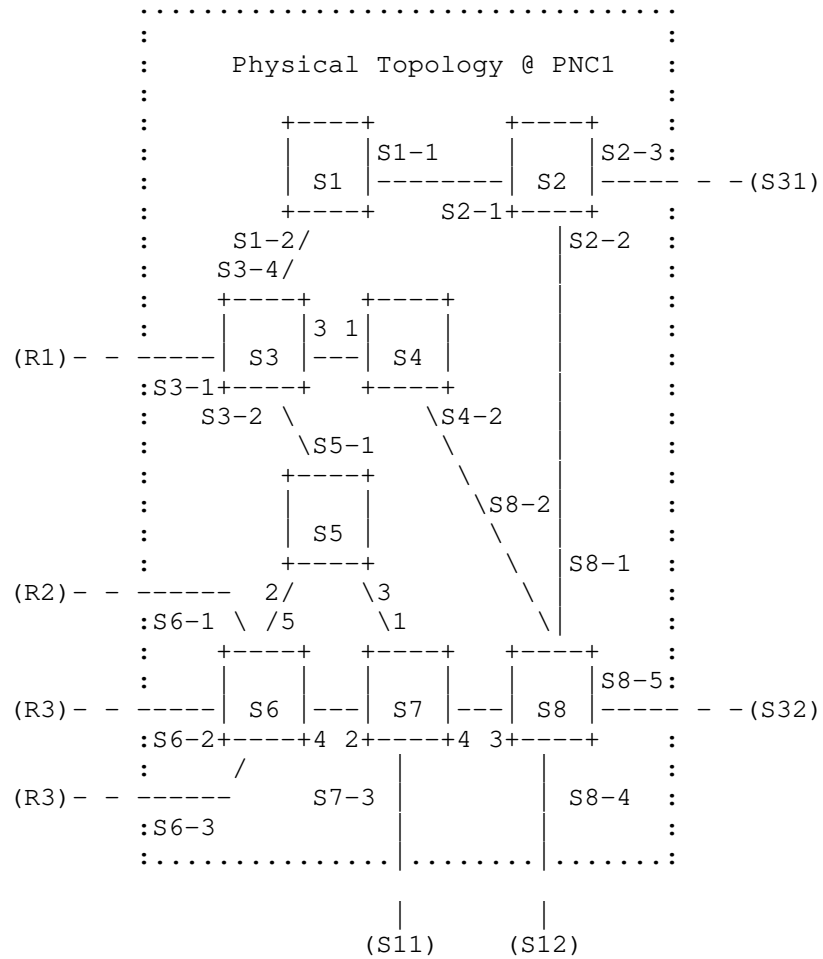


Figure 5 - Physical Topology controlled by PNC1

The PNC1 native topology is not exposed and therefore it under PNC responsibility to abstract the whole domain physical topology as a single TE node and to maintain a mapping between the LTPs exposed at MPI abstract topologies and the associated physical interfaces controlled by the PNC:

Physical Interface (Figure 5)	OTN Topology LTP (Figure 3)	ETH Topology LTP (Figure 4)
S2-3	AN1-7	
S3-1	AN1-1	AN1-1
S6-1		AN1-8
S6-2	AN1-2	
S6-3	AN1-3	
S7-3	AN1-4	
S8-4	AN1-5	
S8-5	AN1-6	

Appendix B.1.1 provides the detailed JSON code example ("mpil-otn-topology.json") describing how the MPI1 ODU Topology is reported by the PNC1, using the [RFC8345], [TE-TOPO] and [OTN-TOPO] YANG models, at MPI1.

It is worth noting that this JSON code example does not provide all the attributes defined in the relevant YANG models, including:

- o YANG attributes which are outside the scope of this document are not shown;
- o The attributes describing the set of label values that are available at the inter-domain links (label-restriction container) are also not shown to simplify the JSON code example;
- o The comments describing the rationale for not including some attributes in this JSON code example even if in the scope of this document are identified with the prefix "// __COMMENT" and included only in the first object instance (e.g., in the Access Link from the AN1-1 description or in the AN1-1 LTP description).

5.1.2. Domain 2 Black Topology Abstraction

PNC2 provides the required black topology abstraction, as described in section 4.2, to expose to the MDSC, at MPI2, two TE Topology instances with only one TE node each:

- o the first instance reports the domain 2 OTN abstract topology view (MPI2 OTN Topology), with only one abstract node (i.e., AN2) and only inter-domain and access abstract TE links (which represent the inter-domain physical links and the access physical links which can support ODU and/or transparent client layers);

- o the instance reports the domain 2 Ethernet abstract topology view (MPI2 ETH Topology), with only one abstract TE node (i.e., AN2) and only access abstract TE links (which represent the access physical links which can support Ethernet client layers).

5.1.3. Domain 3 White Topology Abstraction

PNC3 provides the required white topology abstraction, as described in section 4.2, to expose to the MDSC, at MPI3, two TE Topology instances with multiple TE nodes, one for each physical node:

- o the first instance reports the domain 3 OTN topology view (MPI3 OTN Topology), with four TE nodes, which represent the four physical nodes (i.e. S31, S32, S33 and S34), and abstract TE links, which represent the inter-domain and internal physical links;
- o the second instance reports the domain 3 Ethernet abstract topology view (MPI3 ETH Topology), with only two TE nodes, which represent the two edge physical nodes (i.e., S31 and S33) and only the two access TE links which represent the access physical links.

5.1.4. Multi-domain Topology Merging

As assumed at the beginning of this section, MDSC does not have any knowledge of the topologies of each domain until each PNC reports its own abstract topologies, so the MDSC needs to merge together these abstract topologies, provided by different PNCs, to build its own topology view of the multi-domain network (MDSC multi-domain native topology), as described in section 4.3 of [TE-TOPO].

Given the topologies reported from multiple PNCs, the MDSC need to merge them into its multi-domain native topology. The topology of each domain may be in an abstracted shape (refer to section 5.2 of [RFC8453] for a different level of abstraction), while the inter-domain link information must be complete and fully configured by the MDSC.

The inter-domain link information is reported to the MDSC by the two PNCs, controlling the two ends of the inter-domain link.

The MDSC needs to understand how to merge together these inter-domain links.

One possibility is to use the plug-id information, defined in [TE-TOPO]: two inter-domain TE links, within two different MPI abstract topologies, terminating on two LTPs reporting the same plug-id value can be merged as a single intra-domain link, within any MDSC native topology.

The value of the reported plug-id information can be either assigned by a central network authority, and configured within the two PNC domains. Alternatively, it may be discovered using an automatic discovery mechanisms (e.g., LMP-based, as defined in [RFC6898]).

In case the plug-id values are assigned by a central authority, it is under the central authority responsibility to assign unique values.

In case the plug-id values are automatically discovered, the information discovered by the automatic discovery mechanisms needs to be encoded as a bit string within the plug-id value. This encoding is implementation specific, but the encoding rules need to be consistent across all the PNCs.

In case of co-existence within the same network of multiple sources for the plug-id (e.g., central authority and automatic discovery or even different automatic discovery mechanisms), it is needed that the plug-id namespace is partitioned to avoid that different sources assign the same plug-id value to different inter-domain links. Also, the encoding of the plug-id namespace within the plug-id value is implementation specific and will need to be consistent across all the PNCs.

This document assumes that the plug-id is assigned by a central authority, with the first octet set to 0x00 to represent the central authority namespace. The configuration method used, within each PNC domain, are outside the scope of this document.

Based on the plug-id values, the MDSC can merge together the abstract topologies exposed by the underlying PNCs, as described in sections 5.1.1, 5.1.2 and 5.1.3 above, into its multi-domain native TE topology as shown in Figure 6.

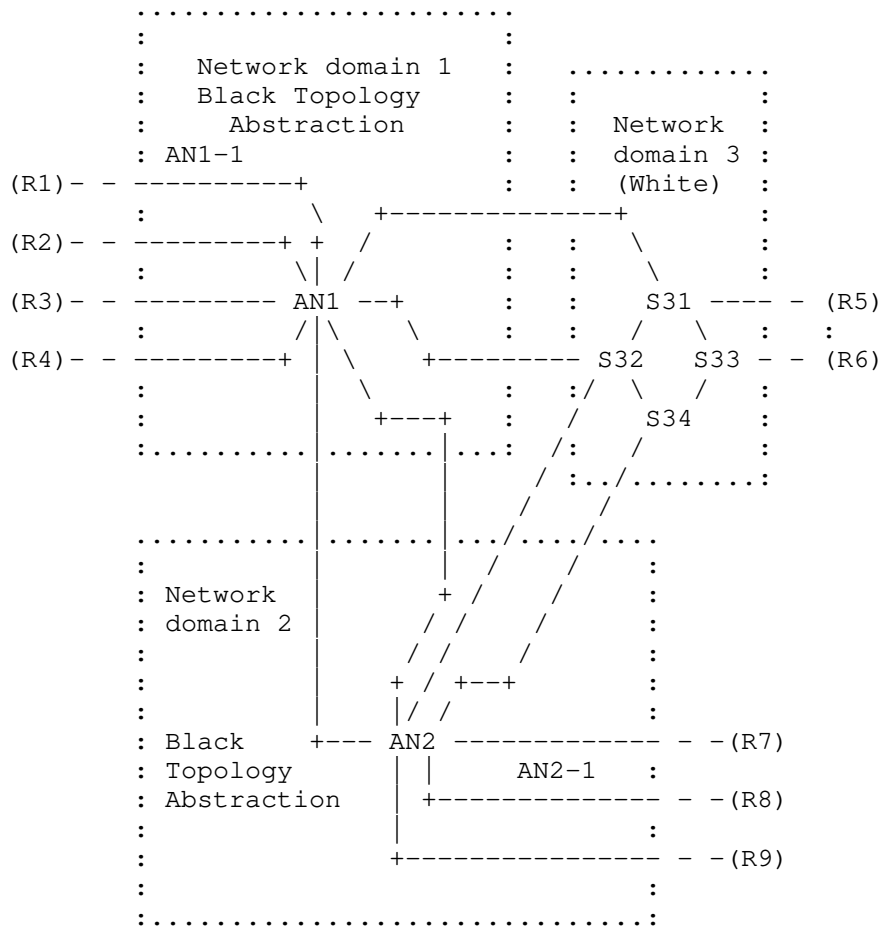


Figure 6 - Multi-domain Abstract Topology controlled by an MDSC

5.2. YANG Models for Service Configuration

The service configuration procedure is assumed to be initiated (step 1 in Figure 7) at the CMI from CNC to MDSC. Analysis of the CMI models is (e.g., L1CSM, L2SM, VN) is outside the scope of this document.

As described in section 4.3, it is assumed that the CMI YANG models provide all the information that allows the MDSC to understand that it needs to coordinate the setup of a multi-domain ODU data plane connection (which can be either an end-to-end connection or a segment connection) and, when needed, also the configuration of the adaptation functions in the edge nodes belonging to different domains.

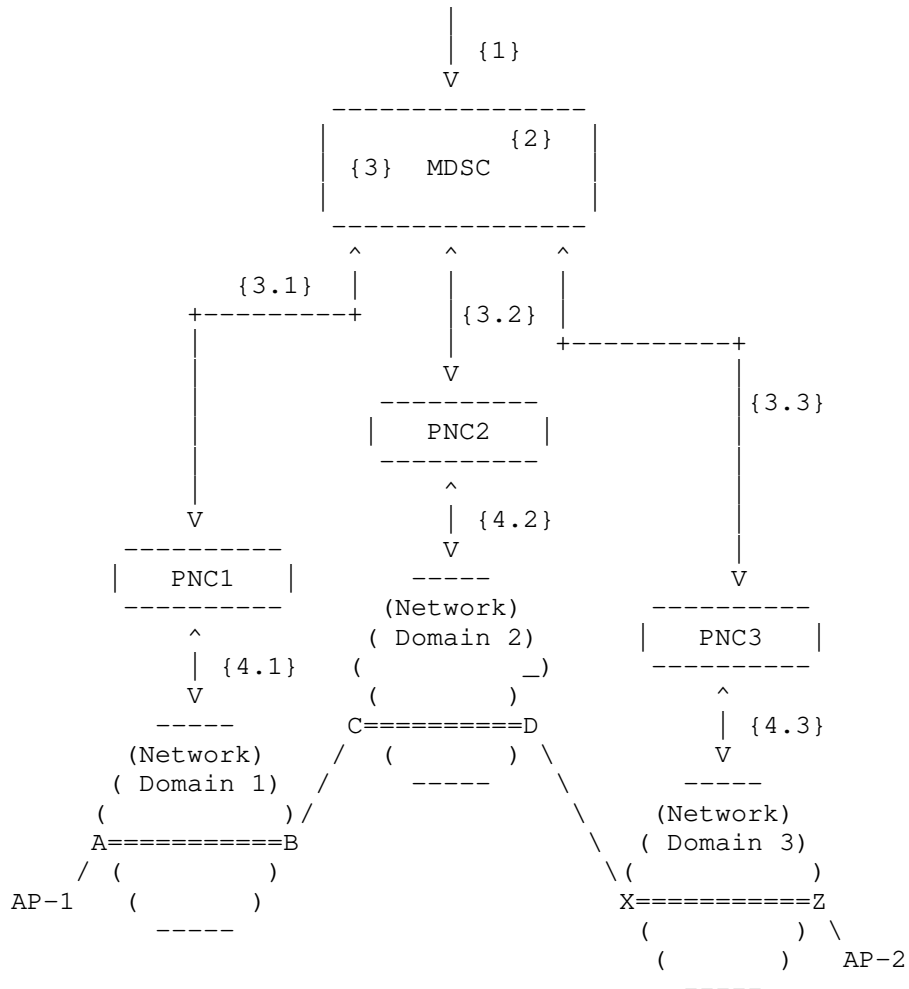


Figure 7 - Multi-domain Service Setup

As an example, the objective in this section is to configure a transport service between R1 and R8, such as one of the services described in section 4.3. The inter-domain path is assumed to be R1 <-> S3 <-> S1 <-> S2 <-> S31 <-> S33 <-> S34 <-> S15 <-> S18 <-> R8 (see the physical topology in Figure 1).

According to the different client signal types, different adaptations can be required to be configured at the edge nodes (i.e., S3 and S18).

After receiving such request, MDSC determines the domain sequence, i.e., domain 1 <-> domain 3 <-> domain 2, with corresponding PNCs and the inter-domain links (step 2 in Figure 7).

As described in [PATH-COMPUTE], the domain sequence can be determined by running the MDSC own path computation on the MDSC native topology, defined in section 5.1.4, if and only if the MDSC has enough topology information. Otherwise, the MDSC can send path computation requests to the different PNCs (steps 2.1, 2.2 and 2.3 in Figure 7) and use this information to determine the optimal path on its internal topology and therefore the domain sequence.

The MDSC will then decompose the tunnel request into a few tunnel segments via tunnel models (both technology agnostic TE tunnel model and OTN tunnel model), and request different PNCs to setup each intra-domain tunnel segment (steps 3, 3.1, 3.2 and 3.3 in Figure 7).

The MDSC will take care of the configuration of both the intra-domain tunnel segment and inter-domain tunnel via corresponding MPI (via TE tunnel model and OTN tunnel model) through all the PNCs controlling the domains selected during path computation. More specifically, for the inter-domain tunnel hand-off, taking into account that the inter-domain links are all OTN links, the list of timeslots and the TPN value assigned to that ODUk connection at the inter-domain link needs to be configured by the MDSC.

In any case, the access link configuration is done only on the PNCs that control the access links (e.g., PNC-1 and PNC-3) and not on the PNCs of transit domain(s) (e.g., PNC-2). An access link will be configured by MDSC after the OTN tunnel is set up.

Access configuration will vary and will be dependent on each type of service. Further discussion and examples are provided in the following sub-sections.

5.2.1. ODU Transit Service

In this scenario, described in section 4.3.1, the physical access links are configured as 10G OTN links and, as described in section 5.1, reported by each PNC as TE Links within the OTN abstract topologies they expose to the MDSC.

To setup this IP link, between R1 and R8, the CNC requests, at the CMI, the MDSC to setup an ODU transit service.

From its native topology, shown in Figure 6, the MDSC understands, by means which are outside the scope of this document, that R1 is attached to the access link terminating on AN1-1 LTP in the MPI1 OTN Abstract Topology (Figure 3), exposed by PNC1, and that R8 is attached to the access link terminating on AN2-1 LTP in the MPI2 Abstract Topology, exposed by PNC2.

MDSC then performs multi-domain path computation (step 2 in Figure 7) and requests PNC1, PNC2 and PNC3, at MPI1, MPI2 and MPI3 respectively, to setup ODU2 (Transit Segment) Tunnels within the OTN Abstract Topologies they expose (MPI1 OTN Abstract Topology, MPI2 OTN Abstract Topology and MPI3 OTN Abstract Topology, respectively).

MDSC requests, at MPI1, PNC1 to setup an ODU2 (Transit Segment) Tunnel with one primary path between AN-1 and AN1-7 LTPs, within the MPI1 OTN Abstract Topology (Figure 4), using the TE Tunnel YANG model, defined in [TE-TUNNEL], with the OTN technology-specific augmentations, defined in [OTN-TUNNEL]:

- o Source and Destination TTPs are not specified (since it is a Transit Tunnel)
- o Ingress and egress points are indicated in the route-object-include-exclude list of the explicit-route-objects of the primary path:
 - o The first element references the access link terminating on AN1-1 LTP
 - o The last two element reference respectively the inter-domain link terminating on AN1-7 LTP and the data plane resources (i.e., the list of timeslots and the TPN) used by the ODU2 connection over that link.

The configuration of the timeslots used by the ODU2 connection on the internal links within a PNC domain (i.e., on the internal links within domain1) is outside the scope of this document since it is a matter of the PNC domain internal implementation.

However, the configuration of the timeslots used by the ODU2 connection at the transport network domain boundaries (e.g., on the inter-domain links) needs to take into account the timeslots available on physical nodes belonging to different PNC domains (e.g., on node S2 within PNC1 domain and on node S31 within PNC3 domain).

The MDSC, when coordinating the setup of a multi-domain ODU connection, also configures the data plane resources (i.e., the list of timeslots and the TPN) to be used on the inter-domain links. The MDSC can know the timeslots which are available on the physical OTN nodes terminating the inter-domain links (e.g., S2 and S31) from the OTN Topology information exposed, at the MPIs, by the PNCs controlling the OTN physical nodes (e.g., PNC1 and PNC3 controlling the physical nodes S2 and S31 respectively).

Appendix B.2.1 provides the detailed JSON code ("mpil-odu2-service-config.json") describing how the setup of this ODU2 (Transit Segment) Tunnel can be requested by the MDSC, using the [TE-TUNNEL] and [OTN-TUNNEL] YANG models at MPI1.

PNC1 knows, as described in the mapping table in Section 5.1.1, that AN-1 and AN1-7 LTPs within the MPI1 OTN Abstract Topology it exposes at MPI1 correspond to the S3-1 and S2-3 LTPs, respectively, within its native topology. Therefore it performs path computation, for an ODU2 connection between these LTPs within its native topology, and sets up the ODU2 cross-connections within the physical nodes S3, S1 and S2, as shown in section 4.3.1.

Since the R1-S3 access link is a multi-function access link, PNC1 also configures the OTU2 trail before setting up the ODU2 cross-connection in node S3.

As part of the OUD2 cross-connection configuration in node S2, PNC1 configures the data plane resources (i.e., the list of timeslots and the TPN), to be used by this ODU2 connection on the S2-S31 inter-domain link, as requested by the MDSC.

Following similar requests from MDSC to setup ODU2 (Transit Segment) Tunnels within the OTN Abstract Topologies they expose, PNC2 then sets up ODU2 cross-connections on nodes S31 and S33 while PNC3 sets up ODU2 cross-connections on nodes S15 and S18, as shown in section 4.3.1. PNC2 also configures the OTU2 trail on the S18-R8 multi-function access link.

5.2.1.1. Single Domain Example

To setup an ODU2 end-to-end connection, supporting an IP link, between R1 and R3, the CNC requests, at the CMI, the MDSC to setup an ODU transit service.

Following the procedures described in section 5.2.1, MDSC requests only PCN1 to setup the ODU2 (Transit Segment) Tunnel between the access links terminating on AN-1 and AN1-2 LTPs within the MPI1 Abstract Topology and PNC1 sets up ODU2 cross-connections on nodes S3, S5 and S6, as shown in section 4.3.1. PNC1 also configures the OTU2 trails on the R1-S3 and R3-S6 multi-function access links.

5.2.2. EPL over ODU Service

In this scenario, described in section 4.3.2, the access links are configured as 10GE Links and, as described in section 5.1, reported by each PNC as TE Links within the ETH abstract topologies they expose to the MDSC.

To setup this IP link, between R1 and R8, the CNC requests, at the CMI, the MDSC to setup an EPL service.

From its native topology, shown in Figure 6, the MDSC understands, by means which are outside the scope of this document, that R1 is attached to the access link terminating on AN1-1 LTP in the MPI1 ETH Abstract Topology, exposed by PNC1, and that R8 is attached to the access link terminating on AN2-1 LTP in the MPI2 ETH Abstract Topology, exposed by PNC2.

The MDSC also understands that it needs to coordinate the setup of a multi-domain ODU2 Tunnel between the TTPs, abstracting nodes S3 and S18 within the OTN Abstract Topologies exposed by PNC1 and PNC2, respectively.

MDSC then performs multi-domain path computation (step 2 in Figure 7) and then requests:

- o PNC1, at MPI1, to setup an ODU2 (Head Segment) Tunnel within the MPI1 OTN Abstract Topology;
- o PNC1, at MPI1, to steer the Ethernet client traffic from/to AN1-1 LTP, within the MPI1 ETH Abstract Topology, through that ODU2 (Head Segment) Tunnel;
- o PNC3, at MPI3, to setup an ODU2 (Transit Segment) Tunnel within the MPI3 OTN Abstract Topology;
- o PNC2, at MPI2, to setup ODU2 (Tail Segment) Tunnel within the MPI2 OTN Abstract Topology;
- o PNC2, at MPI2, to steer the Ethernet client traffic to/from AN2-1 LTP, within the MPI2 ETH Abstract Topology, through that ODU2 (Tail Segment) Tunnel.

MDSC requests, at MPI1, PNC1 to setup an ODU2 (Head Segment) Tunnel with one primary path between the source TTP and AN1-7 LTP, within the MPI1 OTN Abstract Topology (Figure 4), using the TE Tunnel YANG model, defined in [TE-TUNNEL], with the OTN technology-specific augmentations, defined in [OTN-TUNNEL]:

- o Only the Source TTP is specified (since it is a Head Segment Tunnel): therefore the Destination TTP is not specified
- o The egress point is indicated in the route-object-include-exclude list of the explicit-route-objects of the primary path:
 - o The last two element reference respectively the inter-domain link terminating on AN1-7 LTP and the data plane resources (i.e., the list of timeslots and the TPN) used by the ODU2 connection over that link.

Appendix B.2.2 provides the detailed JSON code ("mpil-odu2-tunnel-config.json") describing how the setup of this ODU2 (Head Segment) Tunnel can be requested by the MDSC, using the [TE-TUNNEL] and [OTN-TUNNEL] YANG models at MPI1.

MDSC requests, at MPI1, PNC1 to steer the Ethernet client traffic from/to AN1-2 LTP, within the MPI1 ETH Abstract Topology (Figure 4), through the MPI1 ODU2 (Head Segment) Tunnel, using the Ethernet Client YANG model, defined in [CLIENT-SIGNAL].

Appendix B.2.3 provides the detailed JSON code ("mpil-epl-service-config.json") describing how the setup of this EPL service using the ODU2 Tunnel can be requested by the MDSC, using the [CLIENT-SIGNAL] YANG model at MPI1.

PNC1 knows, as described in the table in section 5.1.1, that the tunnel source TTP and AN1-7 LTP, within the MPI1 OTN Abstract Topology it exposes at MPI1, correspond to node S3 and S2-3 LTP, respectively, within its native topology. Therefore it performs path computation, for an ODU2 connection between node S3 and S2-3 LTP within its native topology, and sets up the ODU2 cross-connections within the physical nodes S3, S1 and S2, as shown in section 4.3.2.

As part of the ODU2 cross-connection configuration in node S2, PNC1 configures the data plane resources (i.e., the list of timeslots and the TPN), to be used by this ODU2 connection on the S2-S31 inter-domain link, as requested by the MDSC.

After the configuration of the ODU2 cross-connection in node S3, PNC1 also configures the [ETH -> (ODU)] and [(ODU2) -> ETH] adaptation functions, within node S3, as shown in section 4.3.2.

Since the R1-S3 access link is a multi-function access link, PNC1 also configures the 10GE link before this step.

Following similar requests from MDSC to setup ODU2 (Segment) Tunnels within the OTN Abstract Topologies they expose as well as the steering of the Ethernet client traffic, PNC3 then sets up ODU2 cross-connections on nodes S31 and S33 while PNC2 sets up ODU2 cross-connections on nodes S15 and S18 as well as the [ETH -> (ODU2)] and [(ODU2) -> ETH] adaptation functions in node S18, as shown in section 4.3.2. PNC2 also configures the 10GE link on the S18-R8 multi-function access link.

5.2.2.1. Single Domain Example

To setup this IP link, between R1 and R2, the CNC requests, at the CMI, the MDSC to setup an EPL service.

Following the procedures described in section 5.2.2, the MDSC requests PCN1 to:

- o setup an ODU2 (end-to-end) Tunnel between the TTPs, abstracting nodes S3 and S6 within MPI1 OTN Abstract Topology exposed by PNC1 at MPI1;
- o steer the Ethernet client traffic between the AN1-1 and AN1-8 LTPs, exposed by PNC1 within MPI1 ETH Abstract Topology, through that ODU2 (end-to-end) Tunnel.

Then PNC1 sets up ODU2 cross-connections on nodes S3, S5 and S6 as well as the [ETH -> (ODU)] and [(ODU2) -> ETH] adaptation functions in nodes S3 and S6, as shown in section 4.3.2. PNC1 also configures the 10GE link on the R1-S3 multi-function access link (the R2-S6 access link has been pre-configured as a 10GE link).

5.2.3. Other OTN Client Services

In this scenario, described in section 4.3.3, the access links are configured as STM-64 links and, as described in section 5.1, reported by each PNC as TE Links within the OTN Abstract Topologies they expose to the MDSC.

The CNC requests, at the CMI, MDSC to setup an STM-64 Private Line service between R1 and R8.

Following similar procedures as described in section 5.2.2, MDSC understands that:

- o R1 is attached to the access link terminating on AN1-1 LTP in the MPI1 OTN Abstract Topology, exposed by PNC1, and that R8 is attached to the access link terminating on AN2-1 LTP in the MPI2 OTN Abstract Topology, exposed by PNC2;
- o it needs to coordinate the setup of a multi-domain ODU2 Tunnel between the TTPs, abstracting nodes S3 and S18 within the OTN Abstract Topologies exposed by PNC1 and PNC2, respectively.

The MDSC then performs multi-domain path computation (step 2 in Figure 7) and then requests:

- o PNC1, at MPI1, to setup an ODU2 (Head Segment) Tunnel within the MPI1 OTN Abstract Topology;

- o PNC1, at MPI1, to steer the STM-64 transparent client traffic from/to AN1-1 LTP, within the MPI1 OTN Abstract Topology, through that ODU2 (Head Segment) Tunnel;
- o PNC3, at MPI3, to setup an ODU2 (Transit Segment) Tunnel within the MPI3 OTN Abstract Topology;
- o PNC2, at MPI2, to setup ODU2 (Tail Segment) Tunnel within the MPI2 OTN Abstract Topology;
- o PNC2, at MPI2, to steer the STM-64 transparent client traffic to/from AN2-1 LTP, within the MPI2 ETH Abstract Topology, through that ODU2 (Tail Segment) Tunnel.

PNC1, PNC2 and PNC3 then sets up the ODU2 cross-connections within the physical nodes S3, S1, S2, S31, S33, S15 and S18 as well as the [STM-64 -> (ODU)] and [(ODU2) -> STM-64] adaptation functions in nodes S3 and S18, as shown in section 4.3.3. PNC1 and PNC2 also configure the STM-64 links on the R1-S3 and R8-S18 multi-function access links, respectively.

5.2.3.1. Single Domain Example

To setup this IP link, between R1 and R3, the CNC requests, at the CMI, the MDSC to setup an STM-64 Private Line service.

The MDSC and PNC1 follows similar procedures as described in section 5.2.2.1 to set up ODU2 cross-connections on nodes S3, S5 and S6 as well as the [STM-64 -> (ODU)] and [(ODU2) -> STM-64] adaptation functions in nodes S3 and S6, as shown in section 4.3.3. PNC1 also configures the STM-64 links on the R1-S3 and R3-S6 multi-function access links.

5.2.4. EVPL over ODU Service

In this scenario, described in section 4.3.4, the access links are configured as 10GE links, as described in section 5.2.2 above.

The CNC requests, at the CMI, the MDSC to setup two EVPL services: one between R1 and R2, and another between R1 and R8.

Following similar procedures as described in section 5.2.2 and 5.2.2.1, MDSC understands that:

- o R1 and R2 are attached to the access links terminating respectively on AN1-1 and AN1-8 LTPs in the MPI1 ETH Abstract Topology, exposed by PNC1, and that R8 is attached to the access link terminating on AN2-1 LTP in the MPI2 ETH Abstract Topology, exposed by PNC2;
- o to setup the first (single-domain) EVPL service, between R1 and R2, it needs to coordinate the setup of a single-domain ODU0 Tunnel between the TTPs, abstracting nodes S3 and S6 within the OTN Abstract Topology exposed by PNC1;
- o to setup the second (multi-domain) EVPL service, between R1 and R8, it needs to coordinate the setup of a multi-domain ODU0 Tunnel between the TTPs, abstracting nodes S3 and S18 within the OTN Abstract Topologies exposed by PNC1 and PNC2, respectively.

To setup the first (single-domain) EVPL service between R1 and R2, the MDSC and PNC1 follows similar procedures as described in section 5.2.2.1 to set up ODU0 cross-connections on nodes S3, S5 and S6 as well as the [VLAN -> (ODU0)] and [(ODU0) -> VLAN] adaptation functions, in nodes S3 and S6, as shown in section 4.3.4. PNC1 also configures the 10GE link on the R1-S3 multi-function access link.

As part of the [VLAN -> (ODU0)] and [(ODU0) -> VLAN] adaptation functions configurations in nodes S2 and S6, PNC1 configures also the classification rules required to associated only the Ethernet client traffic received with VLAN ID 10 on the R1-S3 and R2-S6 access links with this EVPL service. The MDSC provides this information to PNC1 using the [CLIENT-SIGNAL] model.

To setup the second (multi-domain) EVPL service between R1 and R8, the MDSC, PNC1, PNC2 and PNC3 follows similar procedures as described in section 5.2.2 to setup the ODU0 cross-connections within the physical nodes S3, S1, S2, S31, S33, S15 and S18 as well as the [VLAN -> (ODU0)] and [(ODU0) -> VLAN] adaptation functions in nodes S3 and S18, as shown in section 4.3.4. PNC2 also configures the 10GE link on the R8-S18 multi-function access link (the R1-S3 10GE link has been already configured when the first EVPL service has been setup).

As part of the [VLAN -> (ODU0)] and [(ODU0) -> VLAN] adaptation functions configurations in nodes S3 and S18, PNC1 and, respectively, PNC2 configure also the classification rules required to associated only the Ethernet client traffic received with VLAN ID

20 on the R1-S3 and R8-S18 access links with this EVPL service. The MDSC provides this information to PNC1 and PNC2 using the [CLIENT-SIGNAL] model.

5.3. YANG Models for Protection Configuration

5.3.1. Linear Protection (end-to-end)

To be discussed in future versions of this document.

5.4. Notifications

Further detailed analysis is outside the scope of this document

5.5. Path Computation with Constraints

The path computation constraints that can be supported at the MPI using the IETF YANG models defined in [TE-TUNNEL] and [PATH-COMPUTE].

When there is a technology specific network (e.g., OTN), the corresponding technology (e.g., OTN) model should also be used to specify the tunnel information on MPI, with the constraint included in TE Tunnel model.

Further detailed analysis is outside the scope of this document

6. Security Considerations

This document analyses the applicability of the YANG models being defined by the IETF to support OTN single and multi-domain scenarios.

Inherently OTN networks ensure privacy and security via hard partitioning of traffic onto dedicated circuits. The separation of network traffic makes it difficult to intercept data transferred between nodes over OTN-channelized links.

In OTN the (General Communication Channel) GCC is used for OAM functions such as performance monitoring, fault detection, and signaling. The GCC control channel should be secured using a suitable mechanism.

There are no additional or new security considerations introduced by this document.

7. IANA Considerations

This document requires no IANA actions.

8. References

8.1. Normative References

- [RFC4427] Mannie, E., Papadimitriou, D., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, March 2006.
- [RFC4655] Farrel, A. et al., "A Path Computation Element (PCE)-Based Architecture", RFC4655, August 2006.
- [RFC7926] Farrel, A. et al., "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, July 2016.
- [RFC8345] Clemm, A., Medved, J. et al., "A Yang Data Model for Network Topologies", RFC8345, March 2018.
- [RFC8453] Ceccarelli, D., Lee, Y. et al., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC8453, August 2018.
- [ITU-T G.709] ITU-T Recommendation G.709 (06/16), "Interfaces for the optical transport network", June 2016.
- [ITU-T G.808.1] ITU-T Recommendation G.808.1 (05/14), "Generic protection switching - Linear trail and subnetwork protection", May 2014.
- [ITU-T G.873.1] ITU-T Recommendation G.873.1 (05/14), "Optical transport network (OTN): Linear protection", May 2014.
- [TE-TOPO] Liu, X. et al., "YANG Data Model for TE Topologies", draft-ietf-teas-yang-te-topo, work in progress.
- [OTN-TOPO] Zheng, H. et al., "A YANG Data Model for Optical Transport Network Topology", draft-ietf-ccamp-otn-topo-yang, work in progress.

[CLIENT-TOPO] Zheng, H. et al., "A YANG Data Model for Client-layer Topology", draft-zheng-ccamp-client-topo-yang, work in progress.

[TE-TUNNEL] Saad, T. et al., "A YANG Data Model for Traffic Engineering Tunnels and Interfaces", draft-ietf-teas-yang-te, work in progress.

[PATH-COMPUTE] Busi, I., Belotti, S. et al, "Yang model for requesting Path Computation", draft-ietf-teas-yang-path-computation, work in progress.

[OTN-TUNNEL] Zheng, H. et al., "OTN Tunnel YANG Model", draft-ietf-ccamp-otn-tunnel-model, work in progress.

[CLIENT-SIGNAL] Zheng, H. et al., "A YANG Data Model for Optical Transport Network Client Signals", draft-zheng-ccamp-otn-client-signal-yang, work in progress.

8.2. Informative References

[RFC5151] Farrel, A. et al., "Inter-Domain MPLS and GMPLS Traffic Engineering --Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 5151, February 2008.

[RFC6898] Li, D. et al., "Link Management Protocol Behavior Negotiation and Configuration Modifications", RFC 6898, March 2013.

[RFC8040] Bierman, A. et al., "RESTCONF Protocol", RFC 8040, January 2017.

[RFC8309] Wu, Q. et al., "Service Models Explained", RFC 8309, January 2018.

[ACTN-YANG] Zhang, X. et al., "Applicability of YANG models for Abstraction and Control of Traffic Engineered Networks", draft-zhang-teas-actn-yang, work in progress.

[RFC-FOLD] Watsen, K. et al., "Handling Long Lines in Artwork in Internet-Drafts and RFCs", draft-ietf-netmod-artwork-folding, work in progress

[TE-TUTORIAL] Bryskin, I. et al., "TE Topology and Tunnel Modeling for Transport Networks", draft-ietf-teas-te-topo-and-tunnel-modeling, work in progress

[ONF TR-527] ONF Technical Recommendation TR-527, "Functional Requirements for Transport API", June 2016.

[ONF GitHub] ONF Open Transport (SNOWMASS),
<<https://github.com/OpenNetworkingFoundation/Snowmass-ONFOpenTransport>>

[MEF 55] Metro Ethernet Forum, "Lifecycle Service Orchestration (LSO): Reference Architecture and Framework", Technical Specification MEF 55, March 2016,
<https://www.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf>

9. Acknowledgments

The authors would like to thank all members of the Transport NBI Design Team involved in the definition of use cases, gap analysis and guidelines for using the IETF YANG models at the Northbound Interface (NBI) of a Transport SDN Controller.

The authors would like to thank Xian Zhang, Anurag Sharma, Sergio Belotti, Tara Cummings, Michael Scharf, Karthik Sethuraman, Oscar Gonzalez de Dios, Hans Bjursrom and Italo Busi for having initiated the work on gap analysis for transport NBI and having provided foundations work for the development of this document.

The authors would like to thank the authors of the TE Topology and Tunnel YANG models [TE-TOPO] and [TE-TUNNEL], in particular, Igor Bryskin, Vishnu Pavan Beeram, Tarek Saad and Xufeng Liu, for their support in addressing any gap identified during the analysis work.

The authors would like to thank Henry Yu and Aihua Guo for their input and review of the URIs structures used within the JSON code examples.

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

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

Appendix A. Validating a JSON fragment against a YANG Model

The objective is to have a tool that allows validating whether a piece of JSON code embedded in an Internet-Draft is compliant with a YANG model without using a client/server.

A.1. Manipulation of JSON fragments

This section describes the various ways JSON fragments are used in the I-D processing and how to manage them.

Let's call "folded-JSON" the JSON embedded in the I-D: it fits the 72 chars width and it is acceptable for it to be invalid JSON.

We then define "unfolded-JSON" a valid JSON fragment having the same contents of the "folded-JSON" without folding, i.e. limits on the text width. The folding/unfolding operation may be done according to [RFC-FOLD]. The "unfolded-JSON" can be edited by the authors using JSON editors with the advantages of syntax validation and pretty-printing.

Both the "folded" and the "unfolded" JSON fragments can include comments having descriptive fields and directives we'll describe later to facilitate the reader and enable some automatic processing.

The presence of comments in the "unfolded-JSON" fragment makes it an invalid JSON encoding of YANG data. Therefore we call "naked JSON" the JSON where the comments have been stripped out: not only it is valid JSON but it is a valid JSON encoding of YANG data.

The following schema resumes these definitions:

unfold_it -->		stripper -->	
Folded-JSON		Unfolded-JSON	Naked JSON
	<-- fold_it		<-- author edits
<=72-chars?	MUST	MAY	MAY
valid JSON?	MAY	MUST	MUST
JSON-encoding	MAY	MAY	MUST

of YANG data

Our validation toolchain has been designed to take a JSON in any of the three formats and validate it automatically against a set of relevant YANG modules using available open-source tools. It can be found at: <https://github.com/GianmarcoBruno/json-yang/>

A.2. Comments in JSON fragments

We found useful to introduce two kinds of comments, both defined as key-value pairs where the key starts with "//":

- free-form descriptive comments, e.g. "// COMMENT" : "refine this" to describe properties of JSON fragments.

- machine-usable directives e.g. "// __REFERENCES__DRAFTS__" : { "ietf-routing-types@2017-12-04": "rfc8294", } which can be used to automatically download from the network the relevant I-Ds or RFCs and extract from them the YANG models of interest. This is particularly useful to keep consistency when the drafting work is rapidly evolving.

A.3. Validation of JSON fragments: DSDL-based approach

The idea is to generate a JSON driver file (JTOX) from YANG, then use it to translate JSON to XML and validate it against the DSDL schemas, as shown in Figure 8.

Useful link: <https://github.com/mbj4668/pyang/wiki/XmlJson>




Figure 8 - DSDL-based approach for JSON code validation

In order to allow the use of comments following the convention defined in section 3, without impacting the validation process, these comments will be automatically removed from the JSON-file that will be validate.

A.4. Validation of JSON fragments: why not using a XSD-based approach

This approach has been analyzed and discarded because no longer supported by pyang.

The idea is to convert YANG to XSD, JSON to XML and validate it against the XSD, as shown in Figure 9:




Figure 9 - XSD-based approach for JSON code validation

The pyang support for the XSD output format was deprecated in 1.5 and removed in 1.7.1. However pyang 1.7.1 is necessary to work with YANG 1.1 so the process shown in Figure 9 will stop just at step (1).

Appendix B. Detailed JSON Examples

The JSON code examples provided in this appendix have been validated using the tools in Appendix A and folded using the tool in [RFC-FOLD].

B.1. JSON Examples for Topology Abstractions

B.1.1. JSON Code: mpil-otn-topology.json

This is the JSON code reporting the OTN Topology @ MPI:

===== NOTE: '\n' line wrapping per BCP XX (RFC XXXX) =====

```
{
  "// __TITLE__": "ODU Black Topology @ MPI1",
  "// __LAST_UPDATE__": "October 18, 2018",
  "// __MISSING_ATTRIBUTES__": true,
  "// __REFERENCE_DRAFTS__": {
    "ietf-routing-types@2017-12-04": "rfc8294",
    "ietf-otn-types@2017-10-30": "draft-ietf-ccamp-otn-tunnel-model-\
01",
    "ietf-network@2018-02-26": "rfc8345",
    "ietf-network-topology@2018-02-26": "rfc8345",
    "ietf-te-types@2018-06-12": "draft-ietf-teas-yang-te-15",
    "ietf-te-topology@2018-06-15": "draft-ietf-teas-yang-te-topo-18",
    "ietf-otn-topology@2017-10-30": "draft-ietf-ccamp-otn-topo-yang-\
02"
  },
  "// __RESTCONF_OPERATION__": {
    "operation": "GET",
    "url": "http://{{PNC1-ADDR}}/restconf/data/ietf-network:networks"
  },
  "ietf-network:networks": {
    "network": [
      {
        "network-id": "providerId/201/clientId/300/topologyId/otn-bl\
ack-topology",
        "network-types": {
          "ietf-te-topology:te-topology": {
            "ietf-otn-topology:otn-topology": {}
          }
        },
        "ietf-te-topology:provider-id": 201,
        "ietf-te-topology:client-id": 300,
        "ietf-te-topology:te-topology-id": "otn-black-topology",
        "// ietf-te-topology:te": "presence container requires: prov\
ider, client and te-topology-id",
        "ietf-te-topology:te": {
          "name": "OTN Black Topology @ MPI1"
        },
        "// ietf-network:node": "Access LTPs to be reviewed in a fut\
ure update",
        "ietf-network:node": [
          {
            "// __NODE__:__DESCRIPTION__": {
```

```

    "name": "AN1",
    "identifier": "10.0.0.1",
    "type": "Abstract Node",
    "physical node(s)": "whole network domain 1"
  },
  "node-id": "10.0.0.1",
  "ietf-te-topology:te-node-id": "10.0.0.1",
  "ietf-te-topology:te": {
    "te-node-attributes": {
      "name": "AN11",
      "admin-status": "up",
      "// __DISCUSS__ is-abstract": "To be discussed with \
\TE Topology authors",
      "// __DISCUSS__ underlay-topology": "To be discussed\
\ with TE Topology authors"
    },
    "oper-status": "up",
    "// __DISCUSS__ tunnel-termination-point": []
  },
  "ietf-network-topology:termination-point": [
    {
      "// __DESCRIPTION__:__LTP__": {
        "name": "AN1-1 LTP",
        "link type(s)": "OTU-2",
        "physical node": "S3",
        "unnumberd/ifIndex": 1,
        "port type": "tributary port",
        "connected to": "R1"
      },
      "tp-id": "1",
      "ietf-te-topology:te-tp-id": 1,
      "ietf-te-topology:te": {
        "name": "AN1-1 LTP",
        "admin-status": "up",
        "// __DISCUSS__ interface-switching-capability": "\
\See Link attributes (teNodeId/10.0.0.1/teLinkId/1)",
        "// __DISCUSS__ inter-domain-plug-id": "Access Lin\
\k",
        "// __COMMENT__ inter-layer-lock-id": "Empty: OTN \
\Links are pre-configured",
        "oper-status": "up",
        "// __DISCUSS__ ietf-otn-topology:supported-payload\
\ types": "List of ODU clients?",
        "// __DISCUSS__ ietf-otn-topology:client-facing": \

```

```

\true
    }
  },
  {
    "// __DESCRIPTION__:__LTP__": {
      "name": "AN1-2 LTP",
      "link type(s)": "OTU-4",
      "physical node": "S2",
      "unnumberd/ifIndex": 1,
      "port type": "inter-domain port",
      "connected to": "S31"
    },
    "tp-id": "2",
    "ietf-te-topology:te-tp-id": 2,
    "ietf-te-topology:te": {
      "name": "AN1-2 LTP",
      "admin-status": "up",
      "// __DISCUSS__ interface-switching-capability": "\
See Link attributes (teNodeId/10.0.0.1/teLinkId/2)",
      "// __DISCUSS__ inter-domain-plug-id": "Inter-doma\
in Link",
      "oper-status": "up",
      "// __DISCUSS__ ietf-otn-topology:supported-payload\
types": "Empty? (inter-domain OTN link)",
      "// __DEFAULT__ ietf-otn-topology:client-facing": \
false
    }
  },
  {
    "// __DESCRIPTION__:__LTP__": {
      "name": "AN1-3 LTP",
      "link type(s)": "OTU-2",
      "physical node": "S6",
      "unnumberd/ifIndex": 1,
      "port type": "tributary port",
      "connected to": "R2"
    },
    "tp-id": "3",
    "ietf-te-topology:te-tp-id": 3,
    "ietf-te-topology:te": {
      "name": "AN1-3 LTP",
      "admin-status": "up",
      "// __DISCUSS__ interface-switching-capability": "\
See Link attributes (teNodeId/10.0.0.1/teLinkId/3)",

```

```

        "/// __DISCUSS__ inter-domain-plug-id": "Access Lin\
\k",
        "oper-status": "up",
        "/// __DISCUSS__ ietf-otn-topology:supported-payloa\
\d-types": "List of ODU clients?",
        "/// __DISCUSS__ ietf-otn-topology:client-facing": \
\true
    }
},
{
    "/// __DESCRIPTION__:__LTP__": {
        "name": "AN1-4 LTP",
        "link type(s)": "OTU-4",
        "physical node": "S8",
        "unnumberd/ifIndex": 1,
        "port type": "inter-domain port",
        "connected to": "S32"
    },
    "tp-id": "4",
    "ietf-te-topology:te-tp-id": 4,
    "ietf-te-topology:te": {
        "name": "AN1-4 LTP",
        "admin-status": "up",
        "/// __DISCUSS__ interface-switching-capability": "\
\See Link attributes (teNodeId/10.0.0.1/teLinkId/4)",
        "/// __DISCUSS__ inter-domain-plug-id": "Inter-doma\
\in Link",
        "oper-status": "up",
        "/// __DISCUSS__ ietf-otn-topology:supported-payloa\
\d-types": "Empty? (inter-domain OTN link)",
        "/// __DEFAULT__ ietf-otn-topology:client-facing": \
\false
    }
},
{
    "/// __DESCRIPTION__:__LTP__": {
        "name": "AN1-5 LTP",
        "link type(s)": "OTU-4",
        "physical node": "S8",
        "unnumberd/ifIndex": 5,
        "port type": "inter-domain port",
        "connected to": "S12"
    },
    "tp-id": "5",

```



```

        "ietf-te-topology:te-tp-id": 5,
        "ietf-te-topology:te": {
            "name": "AN1-5 LTP",
            "admin-status": "up",
            "// __DISCUSS__ interface-switching-capability": "\
See Link attributes (teNodeId/10.0.0.1/teLinkId/5)",
            "// __DISCUSS__ inter-domain-plug-id": "Inter-doma\
in Link",
            "oper-status": "up",
            "// __DISCUSS__ ietf-otn-topology:supported-payload\
d-types": "Empty? (inter-domain OTN link)",
            "// __DEFAULT__ ietf-otn-topology:client-facing": \
false
        }
    },
    {
        "// __DESCRIPTION__:__LTP__": {
            "name": "AN1-6 LTP",
            "link type(s)": "OTU-4",
            "physical node": "S7",
            "unnumberd/idxIndex": 4,
            "port type": "inter-domain port",
            "connected to": "S11"
        },
        "tp-id": "6",
        "ietf-te-topology:te-tp-id": 6,
        "ietf-te-topology:te": {
            "name": "AN1-6 LTP",
            "admin-status": "up",
            "// __DISCUSS__ interface-switching-capability": "\
See Link attributes (teNodeId/10.0.0.1/teLinkId/6)",
            "// __DISCUSS__ inter-domain-plug-id": "Inter-doma\
in Link",
            "oper-status": "up",
            "// __DISCUSS__ ietf-otn-topology:supported-payload\
d-types": "Empty? (inter-domain OTN link)",
            "// __DEFAULT__ ietf-otn-topology:client-facing": \
false
        }
    },
    {
        "// __DESCRIPTION__:__LTP__": {
            "name": "AN1-7 LTP",
            "link type(s)": "OTU-2",

```

```

        "physical node": "S6",
        "unnumbered/ifIndex": 2,
        "port type": "tributary port",
        "connected to": "R3"
    },
    "tp-id": "7",
    "ietf-te-topology:te-tp-id": 7,
    "ietf-te-topology:te": {
        "name": "AN1-7 LTP",
        "admin-status": "up",
        "// __DISCUSS__ interface-switching-capability": "\
See Link attributes (teNodeId/10.0.0.1/teLinkId/7)",
        "// __DISCUSS__ inter-domain-plug-id": "Access Lin\
k",
        "oper-status": "up",
        "// __DISCUSS__ ietf-otn-topology:supported-payloa\
d-types": "List of ODU clients?",
        "// __DISCUSS__ ietf-otn-topology:client-facing": \
true
    }
}
]
],
    "// ietf-network-topology:link": "Access links to be reviewe\
d in a future update",
    "ietf-network-topology:link": [
        {
            "// __DESCRIPTION__:__LINK__": {
                "name": "Access Link from AN1-1",
                "type": "access link",
                "physical link": "Link from S3-1 to R1"
            },
            "link-id": "teNodeId/10.0.0.1/teLinkId/1",
            "ietf-te-topology:te": {
                "te-link-attributes": {
                    "name": "Access Link from AN1-1",
                    "// __DISCUSS__ access-type": "Can we assume point-t\
o-point as the default value?",
                    "access-type": "point-to-point",
                    "// __COMMENT__ external-domain": "Empty: the plug-i\
d is used instead of this container",
                    "// __DISCUSS__ is-abstract": "To be discussed with \
TE Topology authors",

```

```

        "// __DISCUSS__ underlay": "To be discussed with TE \
\Topology authors",
        "admin-status": "up",
        "interface-switching-capability": [
            {
                "switching-capability": "ietf-te-types:switching\
\otn",
                "encoding": "ietf-te-types:lsp-encoding-oduk",
                "max-lsp-bandwidth": [
                    {
                        "priority": 0,
                        "// __DISCUSS__ te-bandwidth": "ODU2"
                    }
                ]
            }
        ],
        "// __COMMENT__ label-restrictions": "Not described \
\in this JSON example",
        "// __DISCUSS__ link-protection-type": "Can we assum\
\e unprotected as the default value?",
        "link-protection-type": "unprotected",
        "max-link-bandwidth": {
            "// __DISCUSS__ te-bandwidth": "1xODU2"
        },
        "max-resv-link-bandwidth": {
            "// __DISCUSS__ te-bandwidth": "1xODU2"
        },
        "unreserved-bandwidth": [
            {
                "priority": 0,
                "// __DISCUSS__ te-bandwidth": "1xODU2"
            }
        ],
        "oper-status": "up",
        "// __EMPTY__ is-transitional": "It is not a transitio\
\nal link",
        "// __DISCUSS__ underlay ": "To be discussed with TE T\
\opology authors"
    },
    "source": {
        "source-node": "10.0.0.1",
        "source-tp": 1
    },
},

```

```

    "// __EMPTY__ destination": "access link"
  },
  {
    "// __DESCRIPTION__:__LINK__": {
      "name": "Inter-domain Link from AN1-2",
      "type": "inter-domain link",
      "physical link": "Link from S2-1 to S31"
    },
    "link-id": "teNodeId/10.0.0.1/teLinkId/2",
    "ietf-te-topology:te": {
      "te-link-attributes": {
        "name": "Inter-domain Link from AN1-2",
        "// __DISCUSS__ access-type": "Can we assume point-to-
\o-point as the default value?",
        "access-type": "point-to-point",
        "// __DISCUSS__ is-abstract": "To be discussed with \
\TE Topology authors",
        "// __DISCUSS__ underlay": "To be discussed with TE \
\Topology authors",
        "admin-status": "up",
        "interface-switching-capability": [
          {
            "switching-capability": "ietf-te-types:switching\
\otn",
            "encoding": "ietf-te-types:lsp-encoding-oduk",
            "max-lsp-bandwidth": [
              {
                "priority": 0,
                "// __DISCUSS__ te-bandwidth": "ODU4"
              }
            ],
            "// __DISCUSS__ label-restrictions": "To be adde\
\d?"
          }
        ],
        "// __DISCUSS__ link-protection-type": "Can we assum\
\e unprotected as the default value?",
        "link-protection-type": "unprotected",
        "max-link-bandwidth": {
          "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
        },
        "max-resv-link-bandwidth": {
          "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
        }
      }
    }
  }

```

```

        "unreserved-bandwidth": [
            {
                "priority": 0,
                "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
            }
        ],
        "oper-status": "up",
        "// __EMPTY__ is-transitional": "It is not a transitio\
\nal link",
        "// __DISCUSS__ underlay ": "To be discussed with TE T\
\nology authors",
    },
    "source": {
        "source-node": "10.0.0.1",
        "source-tp": 2
    },
    "// __EMPTY__ destination": "inter-domain link"
},
{
    "// __DESCRIPTION__:__LINK__": {
        "name": "Access Link from AN1-3",
        "type": "access link",
        "physical link": "Link from S6-1 to R2"
    },
    "link-id": "teNodeId/10.0.0.1/teLinkId/3",
    "ietf-te-topology:te": {
        "te-link-attributes": {
            "name": "Access Link from AN1-3",
            "// __DISCUSS__ access-type": "Can we assume point-t\
\no-point as the default value?",
            "access-type": "point-to-point",
            "// __DISCUSS__ is-abstract": "To be discussed with \
\nTE Topology authors",
            "// __DISCUSS__ underlay": "To be discussed with TE \
\nTopology authors",
            "admin-status": "up",
            "interface-switching-capability": [
                {
                    "switching-capability": "ietf-te-types:switching\
\n-otn",
                    "encoding": "ietf-te-types:lsp-encoding-oduk",
                    "max-lsp-bandwidth": [

```

```

        "priority": 0,
        "// __DISCUSS__ te-bandwidth": "ODU2"
    }
],
    "// __DISCUSS__ label-restrictions": "To be adde\
\d?"
    }
],
    "// __DISCUSS__ link-protection-type": "Can we assum\
e unprotected as the default value?",
    "link-protection-type": "unprotected",
    "max-link-bandwidth": {
        "// __DISCUSS__ te-bandwidth": "1xODU2"
    },
    "unreserved-bandwidth": [
        {
            "priority": 0,
            "// __DISCUSS__ te-bandwidth": "1xODU2"
        }
    ],
    "max-resv-link-bandwidth": {
        "// __DISCUSS__ te-bandwidth": "1xODU2"
    }
},
    "oper-status": "up",
    "// __EMPTY__ is-transitional": "It is not a transitio\
nal link",
    "// __DISCUSS__ underlay ": "To be discussed with TE T\
opology authors"
},
    "source": {
        "source-node": "10.0.0.1",
        "source-tp": 3
    },
    "// __EMPTY__ destination": "access link"
},
{
    "// __DESCRIPTION__:__LINK__": {
        "name": "Inter-domain Link from AN1-4",
        "type": "inter-domain link",
        "physical link": "Link from S8-1 to S32"
    },
    "link-id": "teNodeId/10.0.0.1/teLinkId/4",
    "ietf-te-topology:te": {

```

```

      "te-link-attributes": {
        "name": "Inter-domain Link from AN1-4",
        "// __DISCUSS__ access-type": "Can we assume point-t\
\o-point as the default value?",
        "access-type": "point-to-point",
        "// __DISCUSS__ is-abstract": "To be discussed with \
\TE Topology authors",
        "// __DISCUSS__ underlay": "To be discussed with TE \
\Topology authors",
        "admin-status": "up",
        "interface-switching-capability": [
          {
            "switching-capability": "ietf-te-types:switching\
\otn",
            "encoding": "ietf-te-types:lsp-encoding-oduk",
            "max-lsp-bandwidth": [
              {
                "priority": 0,
                "// __DISCUSS__ te-bandwidth": "ODU4"
              }
            ],
            "// __DISCUSS__ label-restrictions": "To be adde\
\d?"
          }
        ],
        "// __DISCUSS__ link-protection-type": "Can we assum\
\e unprotected as the default value?",
        "link-protection-type": "unprotected",
        "max-link-bandwidth": {
          "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
        },
        "unreserved-bandwidth": [
          {
            "priority": 0,
            "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
          }
        ],
        "max-resv-link-bandwidth": {
          "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
        }
      },
      "oper-status": "up",
      "// __EMPTY__ is-transitional": "It is not a transitio\
\nal link",

```

```

        "// __DISCUSS__ underlay ": "To be discussed with TE T\
\opology authors"
    },
    "source": {
        "source-node": "10.0.0.1",
        "source-tp": 4
    },
    "// __EMPTY__ destination": "inter-domain link"
},
{
    "// __DESCRIPTION__:__LINK__": {
        "name": "Inter-domain Link from AN1-5",
        "type": "inter-domain link",
        "physical link": "Link from S8-5 to S12"
    },
    "link-id": "teNodeId/10.0.0.1/teLinkId/5",
    "ietf-te-topology:te": {
        "te-link-attributes": {
            "name": "Inter-domain Link from AN1-5",
            "// __DISCUSS__ access-type": "Can we assume point-t\
\o-point as the default value?",
            "access-type": "point-to-point",
            "// __DISCUSS__ is-abstract": "To be discussed with \
\TE Topology authors",
            "// __DISCUSS__ underlay": "To be discussed with TE \
\Topology authors",
            "admin-status": "up",
            "interface-switching-capability": [
                {
                    "switching-capability": "ietf-te-types:switching\
\otn",
                    "encoding": "ietf-te-types:lsp-encoding-oduk",
                    "max-lsp-bandwidth": [
                        {
                            "priority": 0,
                            "// __DISCUSS__ te-bandwidth": "ODU4"
                        }
                    ],
                    "// __DISCUSS__ label-restrictions": "To be adde\
\d?"
                }
            ],
            "// __DISCUSS__ link-protection-type": "Can we assum\
\e unprotected as the default value?",

```



```

    "link-protection-type": "unprotected",
    "max-link-bandwidth": {
      "/// __DISCUSS__ te-bandwidth": "1xODU4, ..."
    },
    "max-resv-link-bandwidth": {
      "/// __DISCUSS__ te-bandwidth": "1xODU4, ..."
    },
    "unreserved-bandwidth": [
      {
        "priority": 0,
        "/// __DISCUSS__ te-bandwidth": "1xODU4, ..."
      }
    ]
  },
  "oper-status": "up",
  "/// __EMPTY__ is-transitional": "It is not a transitio\
\nal link",
  "/// __DISCUSS__ underlay ": "To be discussed with TE T\
\nology authors"
},
"source": {
  "source-node": "10.0.0.1",
  "source-tp": 5
},
"/// __EMPTY__ destination": "inter-domain link"
},
{
  "/// __DESCRIPTION__:__LINK__": {
    "name": "Inter-domain Link from AN1-6",
    "type": "inter-domain link",
    "physical link": "Link from S7-4 to S11"
  },
  "link-id": "teNodeId/10.0.0.1/teLinkId/6",
  "ietf-te-topology:te": {
    "te-link-attributes": {
      "name": "Inter-domain Link from AN1-6",
      "/// __DISCUSS__ access-type": "Can we assume point-t\
\no-point as the default value?",
      "access-type": "point-to-point",
      "/// __DISCUSS__ is-abstract": "To be discussed with \
\nTE Topology authors",
      "/// __DISCUSS__ underlay": "To be discussed with TE \
\nTopology authors",
      "admin-status": "up",

```

```

        "interface-switching-capability": [
            {
                "switching-capability": "ietf-te-types:switching\
\ -otn",
                "encoding": "ietf-te-types:lsp-encoding-oduk",
                "max-lsp-bandwidth": [
                    {
                        "priority": 0,
                        "// __DISCUSS__ te-bandwidth": "ODU4"
                    }
                ],
                "// __DISCUSS__ label-restrictions": "To be adde\
\ d?"
            }
        ],
        "// __DISCUSS__ link-protection-type": "Can we assum\
\ e unprotected as the default value?",
        "link-protection-type": "unprotected",
        "max-link-bandwidth": {
            "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
        },
        "max-resv-link-bandwidth": {
            "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
        },
        "unreserved-bandwidth": [
            {
                "priority": 0,
                "// __DISCUSS__ te-bandwidth": "1xODU4, ..."
            }
        ]
    },
    "oper-status": "up",
    "// __EMPTY__ is-transitional": "It is not a transitio\
\ nal link",
    "// __DISCUSS__ underlay ": "To be discussed with TE T\
\ opology authors"
},
"source": {
    "source-node": "10.0.0.1",
    "source-tp": 6
},
"// __EMPTY__ destination": "inter-domain link"
},
{

```

```

    "// __DESCRIPTION__:__LINK__": {
      "name": "Access Link from AN1-7",
      "type": "access link",
      "physical link": "Link from S6-2 to R3"
    },
    "link-id": "teNodeId/10.0.0.1teLinkId/7",
    "ietf-te-topology:te": {
      "te-link-attributes": {
        "name": "Access Link from AN1-7",
        "// __DISCUSS__ access-type": "Can we assume point-to-
\o-point as the default value?",
        "access-type": "point-to-point",
        "// __DISCUSS__ is-abstract": "To be discussed with \
\TE Topology authors",
        "// __DISCUSS__ underlay": "To be discussed with TE \
\Topology authors",
        "admin-status": "up",
        "interface-switching-capability": [
          {
            "switching-capability": "ietf-te-types:switching\
\otn",
            "encoding": "ietf-te-types:lsp-encoding-oduk",
            "max-lsp-bandwidth": [
              {
                "priority": 0,
                "// __DISCUSS__ te-bandwidth": "ODU2"
              }
            ],
            "// __DISCUSS__ label-restrictions": "To be adde\
\d?"
          }
        ],
        "// __DISCUSS__ link-protection-type": "Can we assum\
\e unprotected as the default value?",
        "link-protection-type": "unprotected",
        "max-link-bandwidth": {
          "// __DISCUSS__ te-bandwidth": "1xODU2"
        },
        "max-resv-link-bandwidth": {
          "// __DISCUSS__ te-bandwidth": "1xODU2"
        },
        "unreserved-bandwidth": [
          {
            "priority": 0,

```

```

        "/// __DISCUSS__ te-bandwidth": "1xODU2"
      }
    ]
  },
  "oper-status": "up",
  "/// __EMPTY__ is-transitional": "It is not a transitio\
\nal link",
  "/// __DISCUSS__ underlay ": "To be discussed with TE T\
\nopology authors"
},
"source": {
  "source-node": "10.0.0.1",
  "source-tp": 7
},
"/// __EMPTY__ destination": "access link"
}
]
}
]
}
}

```

B.2. JSON Examples for Service Configuration

B.2.1. JSON Code: mpil-odu2-service-config.json

This is the JSON code reporting the ODU2 transit service configuration @ MPI:

===== NOTE: '\n' line wrapping per BCP XX (RFC XXXX) =====

```
{
  "// __TITLE__": "ODU2 Service Configuration @ MPI1",
  "// __LAST_UPDATE__": "October 22, 2018",
  "// __MISSING_ATTRIBUTES__": true,
  "// __REFERENCE_DRAFTS__": {
    "ietf-routing-types@2017-12-04": "rfc8294",
    "ietf-otn-types@2018-06-07": "draft-ietf-ccamp-otn-tunnel-model-
\02",
    "ietf-te-types@2018-07-01": "draft-ietf-teas-yang-te-16",
    "ietf-te@2018-07-01": "draft-ietf-teas-yang-te-16",
    "ietf-otn-tunnel@2018-06-07": "draft-ietf-ccamp-otn-tunnel-model-
\02"
  },
  "// __RESTCONF_OPERATION__": {
    "operation": "PUT",
    "url": "http://{PNC1-ADDR}/restconf/data/ietf-te:te"
  },
  "ietf-te:te": {
    "tunnels": {
      "tunnel": [
        {
          "name": "mpil-odu2-service",
          "// identifier": "ODU2-SERVICE-TUNNEL-ID @ MPI1",
          "identifier": 1,
          "description": "ODU2 Service implemented by ODU2 OTN Tunne
\1 Segment @ MPI1",
          "// encoding and switching-type": "ODU",
          "encoding": "ietf-te-types:lsp-encoding-oduk ",
          "switching-type": "ietf-te-types:switching-otn",
          "// source": "None: transit tunnel segment",
          "// destination": "None: transit tunnel segment",
          "// src-tp-id": "None: transit tunnel segment",
          "// dst-tp-id": "None: transit tunnel segment",
          "// ietf-otn-tunnel:src-client-signal": "None: ODU transit
\ tunnel segment",
          "// ietf-otn-tunnel:dst-client-signal": "None: ODU transit
\ tunnel segment",
          "bidirectional": true,
          "// protection": "No protection",
          "// __ DEFAULT __ protection": {
            "// __ DEFAULT __ enable": false
          },

```

```

    "// restoration": "No restoration",
    "// __ DEFAULT __ restoration": {
        "// __ DEFAULT __ enable": false
    },
    "// te-topology-identifier": "ODU Black Topology @ MPI1",
    "te-topology-identifier": {
        "provider-id": 201,
        "client-id": 300,
        "topology-id": "otn-black-topology"
    },
    "te-bandwidth": {
        "ietf-otn-tunnel:odu-type": "ietf-otn-types:prot-ODU2"
    },
    "// hierarchical-link": "None: transit tunnel segment",
    "p2p-primary-paths": {
        "p2p-primary-path": [
            {
                "name": "mpil-odu2-service-primary-path",
                "path-scope": "ietf-te-types:path-scope-segment",
                "// te-bandwidth": "None: only the tunnel bandwidth \
needs to be specified in transport applications",
                "explicit-route-objects": {
                    "route-object-include-exclude": [
                        {
                            "// comment": "Tunnel hand-off OTU2 ingress in\
terface (S3-1)",
                            "index": 1,
                            "explicit-route-usage": "ietf-te-types:route-i\
nclude-ero",
                            "num-unnum-hop": {
                                "// node-id": "AN1 Node",
                                "node-id": "10.0.0.1",
                                "// link-tp-id": "AN1-1 LTP",
                                "link-tp-id": 1,
                                "hop-type": "STRICT",
                                "direction": "INCOMING"
                            }
                        }
                    ],
                    {
                        "// comment": "Tunnel hand-off ODU2 ingress la\
bel (ODU2 over OTU2) at S3-1",
                        "index": 2,
                        "explicit-route-usage": "ietf-te-types:route-i\
nclude-ero",

```

```

        "label-hop": {
            "te-label": {
                "// __ DISCUSS __ odu-label": "How are HO-\
\ODU (ODUk over OTUk) label represented?",
                "// __ DISCUSS __ direction": "Check with \
\TE Tunnel authors",
                "direction": "FORWARD "
            }
        },
        {
            "// comment": "Tunnel hand-off OTU4 egress int\
\erface (S2-1)",
            "index": 3,
            "explicit-route-usage": "ietf-te-types:route-i\
\nclude-ero",
            "num-unnum-hop": {
                "// node-id": "AN1 Node",
                "node-id": "10.0.0.1",
                "// link-tp-id": "AN1-2 LTP",
                "link-tp-id": 1,
                "hop-type": "STRICT",
                "direction": "OUTGOING"
            }
        },
        {
            "// comment": "Tunnel hand-off ODU2 egress lab\
\el (ODU2 over OTU4) at S2-1",
            "index": 4,
            "explicit-route-usage": "ietf-te-types:route-i\
\nclude-ero",
            "label-hop": {
                "te-label": {
                    "ietf-otn-tunnel:tpn": 1,
                    "ietf-otn-tunnel:tsg": "ietf-otn-types:tsg\
\n1.25G",
                    "ietf-otn-tunnel:ts-list": "1-8",
                    "// __ DISCUSS __ direction": "Check with \
\TE Tunnel authors",
                    "direction": "FORWARD "
                }
            }
        }
    ]

```

```
    }  
  }  
]  
}  
]  
}  
}
```

B.2.2. JSON Code: `mpil-odu2-tunnel-config.json`

The JSON code for this use case will be added in a future version of this document

An incomplete version is located on GitHub at:

<https://github.com/danielkinguk/transport-nbi>

B.2.3. JSON Code: `mpil-epl-service-config.json`

The JSON code for this use case will be added in a future version of this document

An incomplete version is located on GitHub at:

<https://github.com/danielkinguk/transport-nbi>

Authors' Addresses

Italo Busi (Editor)
Huawei

Email: italo.busi@huawei.com

Daniel King (Editor)
Old Dog Consulting

Email: daniel@olddog.co.uk

Haomian Zheng (Editor)
Huawei

Email: zhenghaomian@huawei.com

Yunbin Xu (Editor)
CAICT

Email: xuyunbin@ritr.cn

Yang Zhao
China Mobile

Email: zhaoyangyjy@chinamobile.com

Sergio Belotti
Nokia

Email: sergio.belotti@nokia.com

Gianmarco Bruno
Ericsson

Email: gianmarco.bruno@ericsson.com

Young Lee
Huawei

Email: leeyoung@huawei.com

Victor Lopez
Telefonica

Email: victor.lopezalvarez@telefonica.com

Carlo Perocchio
Ericsson

Email: carlo.perocchio@ericsson.com

Ricard Vilalta
CTTC

Email: ricard.vilalta@cttc.es

