

CCAMP Working Group
Internet Draft
Intended status: Informational

I. Busi (Ed.)
Huawei
D. King
Lancaster University

Expires: March 2018

September 20, 2017

**Transport Northbound Interface Applicability Statement and Use Cases
draft-tnbidt-ccamp-transport-nbi-use-cases-03**

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

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

This Internet-Draft will expire on March 20, 2018.

Copyright Notice

Copyright (c) 2017 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.

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 describes the key use cases and requirements for transport network control and management. It reviews proposed and existing IETF transport network data models, their applicability, and highlights gaps and requirements.

Table of Contents

| | |
|--|--------------------|
| 1. Introduction | 3 |
| 1.1. Scope of this document | 4 |
| 2. Terminology | 4 |
| 3. Conventions used in this document | 4 |
| 3.1. Topology and traffic flow processing | 4 |
| 4. Use Case 1: Single-domain with single-layer | 5 |
| 4.1. Reference Network | 5 |
| 4.1.1. Single Transport Domain - OTN Network | 5 |
| 4.2. Topology Abstractions | 8 |
| 4.3. Service Configuration | 9 |
| 4.3.1. ODU Transit | 9 |
| 4.3.2. EPL over ODU | 10 |
| 4.3.3. Other OTN Client Services | 10 |
| 4.3.4. EVPL over ODU | 11 |
| 4.3.5. EVPLAN and EVPTree Services | 12 |
| 4.4. Multi-functional Access Links | 13 |
| 4.5. Protection Requirements | 14 |
| 4.5.1. Linear Protection | 15 |
| 5. Use Case 2: Single-domain with multi-layer | 15 |
| 5.1. Reference Network | 15 |
| 5.2. Topology Abstractions | 16 |
| 5.3. Service Configuration | 16 |
| 6. Use Case 3: Multi-domain with single-layer | 16 |
| 6.1. Reference Network | 16 |
| 6.2. Topology Abstractions | 19 |

| | |
|---|----|
| 6.3. Service Configuration | 19 |
| 6.3.1. ODU Transit | 20 |
| 6.3.2. EPL over ODU | 20 |
| 6.3.3. Other OTN Client Services | 21 |
| 6.3.4. EVPL over ODU | 21 |
| 6.3.5. EVPLAN and EVPTree Services | 21 |
| 6.4. Multi-functional Access Links | 22 |
| 6.5. Protection Scenarios | 22 |
| 6.5.1. Linear Protection (end-to-end) | 23 |
| 6.5.2. Segmented Protection | 23 |
| 7. Use Case 4: Multi-domain and multi-layer | 24 |
| 7.1. Reference Network | 24 |
| 7.2. Topology Abstractions | 25 |
| 7.3. Service Configuration | 25 |
| 8. Security Considerations | 25 |
| 9. IANA Considerations | 26 |
| 10. References | 26 |
| 10.1. Normative References | 26 |
| 10.2. Informative References | 26 |
| 11. Acknowledgments | 27 |

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 will be to automate 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 enable also service provisioning coordination/automation. This can be achieved by using standardized YANG models, used together with an appropriate protocol (e.g., [[RESTCONF](#)]).

This document describes key use cases for analyzing the applicability of the existing models defined by the IETF for transport networks. The intention of this document is to become an applicability statement that provides detailed descriptions of how IETF transport models are applied to solve the described use cases and requirements.

1.1.1. 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 [[ACTN-Frame](#)]

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).

The relationship between the current IETF YANG models and the type of ACTN interfaces can be found in [[ACTN-YANG](#)].

The ONF Technical Recommendations for Functional Requirements for the transport API in [ONF TR-527] and the ONF transport API multi-layer examples in [ONF GitHub] have been considered as an input for this work.

Considerations about the CMI (interface between the Customer Network Controller (CNC) and the MDSC) are outside the scope of this document.

2. Terminology

E-LINE: Ethernet Line

EPL: Ethernet Private Line

EVPL: Ethernet Virtual Private Line

OTH: Optical Network Hierarchy

OTN: Optical Transport Network

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 either an adaptation of a client layer into a server layer "(client -> server)" or switching at a given layer "([switching])". Multi-layer switching is indicated by two layer switching with client/server adaptation: "([client] -> [server])".

For example, the following traffic flow:

```
C-R1 (|PKT| -> ODU2), S3 (|ODU2|), S5 (|ODU2|), S6 (|ODU2|),  
C-R3 (ODU2 -> |PKT|)
```

Node C-R1 is switching at the packet (PKT) layer and mapping packets into a 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 C-R3. Node C-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 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.

4. Use Case 1: Single-domain with single-layer

4.1. Reference Network

The current considerations discussed in this document are based on the following reference networks:

- single transport domain: OTN network

4.1.1. Single Transport Domain - OTN Network

As shown in Figure 1 the network physical topology composed of a single-domain transport network providing transport services to an IP network through five access links.




Figure 1 Reference network for Use Case 1

The IP and transport (OTN) domains are respectively composed by five routers C-R1 to C-R5 and by eight ODU switches S1 to S8. The transport domain acts as a transit network providing connectivity for IP layer services.

The behavior of the transport domain is the same whether the ingress or egress service nodes in the IP domain are only attached to the transport domain, or if there are other routers in between the ingress or egress nodes of the IP domain not also attached to the transport domain. In other words, the behavior of the transport network does not depend on whether C-R1, C-R2, ..., C-R5 are PE or P routers for the IP services.

The transport domain control plane architecture follows the ACTN architecture and framework document [[ACTN-Frame](#)], and functional components:

- o Customer Network Controller (CNC) act as a client with respect to the Multi-Domain Service Coordinator (MDSC) via the CNC-MDSC Interface (CMI);

- o MDSC is connected to a plurality of Physical Network Controllers (PNCs), one for each domain, via a MDSC-PNC Interface (MPI). Each PNC is responsible only for the control of its domain and the MDSC is the only entity capable of multi-domain functionalities as well as of managing the inter-domain links;

The ACTN framework facilitates the detachment of the network and service control from the underlying technology and help the customer express the network as desired by business needs. Therefore, care must be taken to keep minimal 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.




Figure 2 Controlling Hierarchy for Use Case 1

Once the service request is processed by the MDSC the mapping of the client IP traffic between the routers (across the transport network) is made in the IP routers only and is not controlled by the transport PNC, and therefore transparent to the transport nodes.

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 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."

[TE-Topo] describes a YANG base model for TE topology without any technology specific parameters. Moreover, it defines how to abstract for TE-network topologies.

[ACTN-Frame] 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 [[ACTN-Frame](#)], there are three types of topology:

- o White topology: This is a case where the Physical Network Controller (PNC) provides the actual network topology to the multi-domain Service Coordinator (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 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 a TE links between them in a full mesh fashion;

- Grey topology type B: border nodes with some internal abstracted nodes and abstracted links.

For single-domain with single-layer use-case, the white topology may be disseminated from the PNC to the MDSC in most cases. There may be some exception to this in the case where the underlay network may have complex optical parameters, which do not warrant the distribution of such details to the MDSC. In such case, the topology disseminated from the PNC to the MDSC may not have the entire TE information but a streamlined TE information. This case would incur another action from the MDSC's standpoint when provisioning a path. The MDSC may make a path compute request to the PNC to verify the feasibility of the estimated path before making the final provisioning request to the PNC, as outlined in [[Path-Compute](#)].

Topology abstraction for the CMI is for further study (to be addressed in future revisions of this document).

[4.3. Service Configuration](#)

In the following use cases, the Multi Domain Service Coordinator (MDSC) needs to be capable to request service connectivity from the transport Physical Network Controller (PNC) to support IP routers connectivity. The type of services could depend of the type of physical links (e.g. OTN link, ETH link or SDH link) between the routers and transport network.

As described in [section 4.1.1](#), the control of different adaptations inside IP routers, C-Ri (PKT -> foo) and C-Rj (foo -> PKT), are assumed to be performed by means that are not under the control of, and not visible to, transport PNC. Therefore, these mechanisms are outside the scope of this document.

[4.3.1. ODU Transit](#)

This use case assumes that the physical links interconnecting the IP routers and the transport network are OTN links. The physical/optical interconnection below the ODU layer is supposed to be pre-configured and not exposed at the MPI to the MDSC.

To setup a 10Gb IP link between C-R1 to C-R3, an ODU2 end-to-end data plane connection needs to be created between C-R1 and C-R3, crossing transport nodes S3, S5, and S6.

The traffic flow between C-R1 and C-R3 can be summarized as:


```
C-R1 (|PKT| -> ODU2), S3 (|ODU2|), S5 (|ODU2|), S6 (|ODU2|),  
C-R3 (ODU2 -> |PKT|)
```

The MDSC should be capable via the MPI to request the setup of an ODU2 transit service with enough information that enable the transport PNC to instantiate and control the ODU2 data plane connection segment through nodes S3, S5, S6.

4.3.2. EPL over ODU

This use case assumes that the physical links interconnecting the IP routers and the transport network are Ethernet links.

In order to setup a 10Gb IP link between C-R1 to C-R3, an EPL service needs to be created between C-R1 and C-R3, supported by an ODU2 end-to-end connection between S3 and S6, crossing transport node S5.

The traffic flow between C-R1 and C-R3 can be summarized as:

```
C-R1 (|PKT| -> ETH), S3 (ETH -> |ODU2|), S5 (|ODU2|),  
S6 (|ODU2| -> ETH), C-R3 (ETH-> |PKT|)
```

The MDSC should be capable via the MPI to request the setup of an EPL service with enough information that can permit the transport PNC to instantiate and control the ODU2 end-to-end data plane connection through nodes S3, S5, S6, as well as the adaptation functions inside S3 and S6: S3&S6 (ETH -> ODU2) and S5&S6 (ODU2 -> ETH).

4.3.3. Other OTN Client Services

[ITU-T G.709-2016] defines mappings of different 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.).

This use case assumes that the physical links interconnecting the IP routers and the transport network are any one of these possible options.

In order to setup a 10Gb IP link between C-R1 to C-R3 using, for example STM-64 physical links between the IP routers and the transport network, an STM-64 Private Line service needs to be created between C-R1 and C-R3, supported by an ODU2 end-to-end data plane connection between S3 and S6, crossing transport node S5.

The traffic flow between C-R1 and C-R3 can be summarized as:

```
C-R1 (|PKT| -> STM-64), S3 (STM-64 -> |ODU2|), S5 (|ODU2|),  
S6 (|ODU2| -> STM-64), C-R3 (STM-64 -> |PKT|)
```

The MDSC should be capable via the MPI to request the setup of an STM-64 Private Line service with enough information that can permit the transport PNC to instantiate and control the ODU2 end-to-end connection through nodes S3, S5, S6, as well as the adaptation functions inside S3 and S6: S3&S6 (STM-64 -> ODU2) and S9&S3 (STM-64 -> PKT).

4.3.4. EVPL over ODU

This use case assumes that the physical links interconnecting the IP routers and the transport network are Ethernet links and that different Ethernet services (e.g, EVPL) can share the same physical link using different VLANs.

In order to setup two 1Gb IP links between C-R1 to C-R3 and between C-R1 and C-R4, two EVPL services need to be created, supported by two ODU0 end-to-end connections respectively between S3 and S6, crossing transport node S5, and between S3 and S2, crossing transport node S1.

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

The traffic flow between C-R1 and C-R3 can be summarized as:

```
C-R1 (|PKT| -> VLAN), S3 (VLAN -> |ODU0|), S5 (|ODU0|),  
S6 (|ODU0| -> VLAN), C-R3 (VLAN -> |PKT|)
```

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

```
C-R1 (|PKT| -> VLAN), S3 (VLAN -> |ODU0|), S1 (|ODU0|),  
S2 (|ODU0| -> VLAN), C-R4 (VLAN -> |PKT|)
```

The MDSC should be capable via the MPI to request the setup of these EVPL services with enough information that can permit the transport PNC to instantiate and control the ODU0 end-to-end data plane connections as well as the adaptation functions on the boundary nodes: S3&S2&S6 (VLAN -> ODU0) and S3&S2&S6 (ODU0 -> VLAN).

4.3.5. EVPLAN and EVPTree Services

This use case assumes that the physical links interconnecting the IP routers and the transport network are Ethernet links and that different Ethernet services (e.g, EVPL, EVPLAN and EVPTree) can share the same physical link using different VLANs.

Note - it is assumed that EPLAN and EPTree services can be supported by configuring EVPLAN and EVPTree with port mapping.

In order to setup an IP subnet between C-R1, C-R2, C-R3 and C-R4, an EVPLAN/EVPTree service needs to be created, supported by two ODUflex end-to-end connections respectively between S3 and S6, crossing transport node S5, and between S3 and S2, crossing transport node S1.

In order to support this EVPLAN/EVPTree service, some Ethernet Bridging capabilities are required on some nodes at the edge of the transport network: for example Ethernet Bridging capabilities can be configured in nodes S3 and S6 but not in node S2.

Since this EVPLAN/EVPTree service can share the same Ethernet physical links between IP routers and transport nodes (e.g., with the EVPL services described in [section 4.3.4](#)), a different VLAN ID (e.g., 30) can be associated with this EVPLAN/EVPTree service.

In order to support an EVPTree service instead of an EVPLAN, additional configuration of the Ethernet Bridging capabilities on the nodes at the edge of the transport network is required.

The MAC bridging function in node S3 is needed to select, based on the MAC Destination Address, whether the Ethernet frames from C-R1 should be sent to the ODUflex terminating on node S6 or to the other ODUflex terminating on node S2.

The MAC bridging function in node S6 is needed to select, based on the MAC Destination Address, whether the Ethernet frames received from the ODUflex should be set to C-R2 or C-R3, as well as whether the Ethernet frames received from C-R2 (or C-R3) should be sent to C-R3 (or C-R2) or to the ODUflex.

For example, the traffic flow between C-R1 and C-R3 can be summarized as:

```
C-R1 (|PKT| -> VLAN), S3 (VLAN -> |MAC| -> |ODUflex|),  
S5 (|ODUflex|), S6 (|ODUflex| -> |MAC| -> VLAN),  
C-R3 (VLAN -> |PKT|)
```


The MAC bridging function in node S3 is also needed to select, based on the MAC Destination Address, whether the Ethernet frames one ODUflex should be sent to C-R1 or to the other ODUflex.

For example, the traffic flow between C-R3 and C-R4 can be summarized as:

```
C-R3 (|PKT| -> VLAN), S6 (VLAN -> |MAC| -> |ODUflex|),  
S5 (|ODUflex|), S3 (|ODUflex| -> |MAC| -> |ODUflex|),  
S1 (|ODUflex|), S2 (|ODUflex| -> VLAN), C-R4 (VLAN -> |PKT|)
```

In node S2 there is no need for any MAC bridging function since all the Ethernet frames received from C-R4 should be sent to the ODUflex toward S3 and viceversa.

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

```
C-R1 (|PKT| -> VLAN), S3 (VLAN -> |MAC| -> |ODUflex|),  
S1 (|ODUflex|), S2 (|ODUflex| -> VLAN), C-R4 (VLAN -> |PKT|)
```

The MDSC should be capable via the MPI to request the setup of this EVPLAN/EVPTree services with enough information that can permit the transport PNC to instantiate and control the ODUflex end-to-end data plane connections as well as the Ethernet Bridging and adaptation functions on the boundary nodes: S3&S6 (VLAN -> MAC -> ODU2), S3&S6 (ODU2 -> ETH -> VLAN), S2 (VLAN -> ODU2) and S2 (ODU2 -> VLAN).

4.4. Multi-functional Access Links

This use case assumes that some physical links interconnecting the IP routers and the transport network can be configured in different modes, e.g., as OTU2 or STM-64 or 10GE.

This configuration can be done a-priori by means outside the scope of this document. In this case, these links will appear at the MPI either as an ODU Link or as an STM-64 Link or as a 10GE Link (depending on the a-priori configuration) and will be controlled at the MPI as discussed in [section 4.3](#).

It is also possible not to configure these links a-priori and give the control to the MPI to decide, based on the service configuration, how to configure it.

For example, if the physical link between C-R1 and S3 is a multi-functional access link while the physical links between C-R3 and S6 and between C-R4 and S2 are STM-64 and 10GE physical links respectively, it is possible at the MPI to configure either an STM-

64 Private Line service between C-R1 and C-R3 or an EPL service between C-R1 and C-R4.

The traffic flow between C-R1 and C-R3 can be summarized as:

C-R1 (|PKT| -> STM-64), S3 (STM-64 -> |ODU2|), S5 (|ODU2|),
S6 (|ODU2| -> STM-64), C-R3 (STM-64 -> |PKT|)

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

C-R1 (|PKT| -> ETH), S3 (ETH -> |ODU2|), S1 (|ODU2|),
S2 (|ODU2| -> ETH), C-R4 (ETH-> |PKT|)

The MDSC should be capable via the MPI to request the setup of either service with enough information that can permit the transport PNC to instantiate and control the ODU2 end-to-end data plane connection as well as the adaptation functions inside S3 and S2 or S6.

4.5. Protection Requirements

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.

The MDSC needs to be capable to request the transport PNC to configure protection when requesting the setup of the connectivity services described in [section 4.3](#).

Since in this use case it is assumed that switching within the transport network domain is performed only in one layer, also protection switching within the transport network domain can only be provided at the OTN ODU layer, for all the services defined in [section 4.3](#).

It may be necessary to consider not only protection, but also restoration functions in the future. Restoration methods would provide capability to reroute and restore connectivity traffic around network faults, without the network penalty imposed with dedicated 1+1 protection schemes.

4.5.1. Linear Protection

It is possible to protect any service defined in [section 4.3](#) from failures within the OTN transport domain by configuring OTN linear protection in the data plane between node S3 and node S6.

It is assumed that the OTN linear protection is configured to with 1+1 unidirectional protection switching type, as defined in [ITU-T G.808.1-2014] and [ITU-T G.873.1-2014], as well as in [[RFC4427](#)].

In these scenarios, a working transport entity and a protection transport entity, as defined in [ITU-T G.808.1-2014], (or a working LSP and a protection LSP, as defined in [[RFC4427](#)]) should be configured in the data plane, for example:

Working transport entity: S3, S5, S6

Protection transport entity: S3, S4, S8, S7, S6

The Transport PNC should be capable to report to the MDSC which is the active transport entity, as defined in [ITU-T G.808.1-2014], 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.

5. Use Case 2: Single-domain with multi-layer

5.1. Reference Network

The current considerations discussed in this document are based on the following reference network:

- single transport domain: OTN and OCh multi-layer network

In this use case, the same reference network shown in Figure 1 is considered. The only difference is that all the transport nodes are capable to switch in the ODU as well as in the OCh layer.

All the physical links within the transport network are therefore assumed to be OCh links. Therefore, with the exception of the access

links, no ODU internal link exists before an OCh end-to-end data plane connection is created within the network.

The controlling hierarchy is the same as described in Figure 2.

The interface within the scope of this document is the Transport MPI which should be capable to control both the OTN and OCh layers.

5.2. Topology Abstractions

A grey topology type B abstraction is assumed: abstract nodes and links exposed at the MPI corresponds 1:1 with the physical nodes and links controlled by the PNC but the PNC abstracts/hides at least some optical parameters to be used within the OCh layer.

5.3. Service Configuration

The same service scenarios, as described in [section 4.3](#), are also applicable to these use cases with the only difference that end-to-end OCh data plane connections will need to be setup before ODU data plane connections.

6. Use Case 3: Multi-domain with single-layer

6.1. Reference Network

In this section we focus on a multi-domain reference network with homogeneous technologies:

- multiple transport domains: OTN networks

Figure 3 shows the network physical topology composed of three transport network domains providing transport services to an IP customer network through eight access links:




Figure 4 Controlling Hierarchy for Use Case 3

In this section we address the case where the CNC controls the customer IP network and requests transport connectivity among IP routers, via the CMI, to an MDSC which coordinates, via three MPIs, the control of a multi-domain transport network through three PNCs.

The interfaces within the scope of this document are the three MPIs while the interface between the CNC and the IP routers is out of its scope and considerations about the CMI are outside the scope of this document.

6.2. Topology Abstractions

Each PNC should provide the MDSC a topology abstraction of the domain's network topology.

Each PNC provides topology abstraction of its own domain topology independently from each other and therefore it is possible that different PNCs provide different types of topology abstractions.

As an example, we can assume that:

- o PNC1 provides a white topology abstraction (likewise use case 1 described in [section 4.2](#))
- o PNC2 provides a type A grey topology abstraction
- o PNC3 provides a type B grey topology abstraction, with two abstract nodes (AN31 and AN32). They abstract respectively nodes S31+S33 and nodes S32+S34. At the MPI, only the abstract nodes should be reported: the mapping between the abstract nodes (AN31 and AN32) and the physical nodes (S31, S32, S33 and S34) should be done internally by the PNC.

The MDSC should be capable to glue together these different abstract topologies to build its own view of the multi-domain network topology. This might require proper administrative configuration or other mechanisms (to be defined/analysed).

6.3. Service Configuration

In the following use cases, it is assumed that the CNC is capable to request service connectivity from the MDSC to support IP routers connectivity.

The same service scenarios, as described in [section 4.3](#), are also application to this use cases with the only difference that the two IP routers to be interconnected are attached to transport nodes which belong to different PNCs domains and are under the control of the CNC.

Likewise, the service scenarios in [section 4.3](#), the type of services could depend of the type of physical links (e.g. OTN link, ETH link or SDH link) between the customer's routers and the multi-domain transport network and the configuration of the different adaptations inside IP routers is performed by means that are outside the scope of this document because not under control of and not visible to the MDSC nor to the PNCs. It is assumed that the CNC is capable to

request the proper configuration of the different adaptation functions inside the customer's IP routers, by means which are outside the scope of this document.

It is also assumed that the CNC is capable via the CMI to request the MDSC the setup of these services with enough information that enable the MDSC to coordinate the different PNCs to instantiate and control the ODU2 data plane connection through nodes S3, S1, S2, S31, S33, S34, S15 and S18, as well as the adaptation functions inside nodes S3 and S18, when needed.

As described in [section 6.2](#), the MDSC should have its own view of the end-to-end network topology and use it for its own path computation to understand that it needs to coordinate with PNC1, PNC2 and PNC3 the setup and control of a multi-domain ODU2 data plane connection.

6.3.1. ODU Transit

In order to setup a 10Gb IP link between C-R1 and C-R5, an ODU2 end-to-end data plane connection needs be created between C-R1 and C-R5, crossing transport nodes S3, S1, S2, S31, S33, S34, S15 and S18 which belong to different PNC domains.

The traffic flow between C-R1 and C-R5 can be summarized as:

```
C-R1 (|PKT| -> ODU2), S3 (|ODU2|), S1 (|ODU2|), S2 (|ODU2|),  
S31 (|ODU2|), S33 (|ODU2|), S34 (|ODU2|),  
S15 (|ODU2|), S18 (|ODU2|), C-R5 (ODU2 -> |PKT|)
```

6.3.2. EPL over ODU

In order to setup a 10Gb IP link between C-R1 and C-R5, an EPL service needs to be created between C-R1 and C-R5, supported by an ODU2 end-to-end data plane connection between transport nodes S3 and S18, crossing transport nodes S1, S2, S31, S33, S34 and S15 which belong to different PNC domains.

The traffic flow between C-R1 and C-R5 can be summarized as:

```
C-R1 (|PKT| -> ETH), S3 (ETH -> |ODU2|), S1 (|ODU2|),  
S2 (|ODU2|), S31 (|ODU2|), S33 (|ODU2|), S34 (|ODU2|),  
S15 (|ODU2|), S18 (|ODU2| -> ETH), C-R5 (ETH -> |PKT|)
```


6.3.3. Other OTN Client Services

In order to setup a 10Gb IP link between C-R1 and C-R5 using, for example SDH physical links between the IP routers and the transport network, an STM-64 Private Line service needs to be created between C-R1 and C-R5, supported by ODU2 end-to-end data plane connection between transport nodes S3 and S18, crossing transport nodes S1, S2, S31, S33, S34 and S15 which belong to different PNC domains.

The traffic flow between C-R1 and C-R5 can be summarized as:

```
C-R1 (|PKT| -> STM-64), S3 (STM-64 -> |ODU2|), S1 (|ODU2|),  
S2 (|ODU2|), S31 (|ODU2|), S33 (|ODU2|), S34 (|ODU2|),  
S15 (|ODU2|), S18 (|ODU2| -> STM-64), C-R5 (STM-64 -> |PKT|)
```

6.3.4. EVPL over ODU

In order to setup two 1Gb IP links between C-R1 to C-R3 and between C-R1 and C-R5, two EVPL services need to be created, supported by two ODU0 end-to-end connections respectively between S3 and S6, crossing transport node S5, and between S3 and S18, crossing transport nodes S1, S2, S31, S33, S34 and S15 which belong to different PNC domains.

The VLAN configuration on the access links is the same as described in [section 4.3.4](#).

The traffic flow between C-R1 and C-R3 is the same as described in [section 4.3.4](#).

The traffic flow between C-R1 and C-R5 can be summarized as:

```
C-R1 (|PKT| -> VLAN), S3 (VLAN -> |ODU2|), S1 (|ODU2|),  
S2 (|ODU2|), S31 (|ODU2|), S33 (|ODU2|), S34 (|ODU2|),  
S15 (|ODU2|), S18 (|ODU2| -> VLAN), C-R5 (VLAN -> |PKT|)
```

6.3.5. EVPLAN and EVPTree Services

In order to setup an IP subnet between C-R1, C-R2, C-R3 and C-R7, an EVPLAN/EVPTree service needs to be created, supported by two ODUflex end-to-end connections respectively between S3 and S6, crossing transport node S5, and between S3 and S18, crossing transport nodes S1, S2, S31, S33, S34 and S15 which belong to different PNC domains.

The VLAN configuration on the access links is the same as described in [section 4.3.5](#).

The configuration of the Ethernet Bridging capabilities on nodes S3 and S6 is the same as described in [section 4.3.5](#) while the configuration on node S18 similar to the configuration of node S2 described in [section 4.3.5](#).

The traffic flow between C-R1 and C-R3 is the same as described in [section 4.3.5](#).

The traffic flow between C-R1 and C-R5 can be summarized as:

```
C-R1 (|PKT| -> VLAN), S3 (VLAN -> |MAC| -> |ODUflex|),  
S1 (|ODUflex|), S2 (|ODUflex|), S31 (|ODUflex|),  
S33 (|ODUflex|), S34 (|ODUflex|),  
S15 (|ODUflex|), S18 (|ODUflex| -> VLAN), C-R5 (VLAN -> |PKT|)
```

6.4. Multi-functional Access Links

The same considerations of [section 4.4](#) apply with the only difference that the ODU data plane connections could be setup across multiple PNC domains.

For example, if the physical link between C-R1 and S3 is a multi-functional access link while the physical links between C-R7 and S31 and between C-R5 and S18 are STM-64 and 10GE physical links respectively, it is possible to configure either an STM-64 Private Line service between C-R1 and C-R7 or an EPL service between C-R1 and C-R5.

The traffic flow between C-R1 and C-R7 can be summarized as:

```
C-R1 (|PKT| -> STM-64), S3 (STM-64 -> |ODU2|), S1 (|ODU2|),  
S2 (|ODU2|), S31 (|ODU2| -> STM-64), C-R3 (STM-64 -> |PKT|)
```

The traffic flow between C-R1 and C-R5 can be summarized as:

```
C-R1 (|PKT| -> ETH), S3 (ETH -> |ODU2|), S1 (|ODU2|),  
S2 (|ODU2|), S31 (|ODU2|), S33 (|ODU2|), S34 (|ODU2|),  
S15 (|ODU2|), S18 (|ODU2| -> ETH), C-R5 (ETH -> |PKT|)
```

6.5. Protection Scenarios

The MDSC needs to be capable to coordinate different PNCs to configure protection switching when requesting the setup of the connectivity services described in [section 6.3](#).

Since in this use case it is assumed that switching within the transport network domain is performed only in one layer, also

protection switching within the transport network domain can only be provided at the OTN ODU layer, for all the services defined in [section 6.3](#).

[6.5.1](#). Linear Protection (end-to-end)

In order to protect any service defined in [section 6.3](#) from failures within the OTN multi-domain transport network, the MDSC should be capable to coordinate different PNCs to configure and control OTN linear protection in the data plane between nodes S3 and node S18.

The considerations in [section 4.5.1](#) are also applicable here with the only difference that MDSC needs to coordinate with different PNCs the setup and control of the OTN linear protection as well as of the working and protection transport entities (working and protection LSPs).

Two cases can be considered.

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

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

[6.5.2](#). Segmented Protection

In order to protect any service defined in [section 6.3](#) from failures within the OTN multi-domain transport network, the MDSC should be capable to request each PNC to configure OTN intra-domain protection when requesting the setup of the ODU2 data plane connection segment.

If linear protection is used within a domain, the considerations in [section 4.5.1](#) are also applicable here only for the PNC controlling the domain where intra-domain linear protection is provided.

If PNC1 provides linear protection, the working and protection transport entities could be:

Working transport entity: S3, S1, S2

Protection transport entity: S3, S4, S8, S2

If PNC2 provides linear protection, the working and protection transport entities could be:

Working transport entity: S15, S18

Protection transport entity: S15, S12, S17, S18

If PNC3 provides linear protection, the working and protection transport entities could be:

Working transport entity: S31, S33, S34

Protection transport entity: S31, S32, S34

[7. Use Case 4: Multi-domain and multi-layer](#)

[7.1. Reference Network](#)

The current considerations discussed in this document are based on the following reference network:

- multiple transport domains: OTN and OCh multi-layer networks

In this use case, the reference network shown in Figure 3 is used. The only difference is that all the transport nodes are capable to switch either in the ODU or in the OCh layer.

All the physical links within each transport network domain are therefore assumed to be OCh links, while the inter-domain links are assumed to be ODU links as described in [section 6.1](#) (multi-domain with single layer - OTN network).

Therefore, with the exception of the access and inter-domain links, no ODU link exists within each domain before an OCh single-domain end-to-end data plane connection is created within the network.

The controlling hierarchy is the same as described in Figure 4.

The interfaces within the scope of this document are the three MPIs which should be capable to control both the OTN and OCh layers within each PNC domain.

7.2. Topology Abstractions

Each PNC should provide the MDSC a topology abstraction of its own network topology as described in [section 5.2](#).

As an example, it is assumed that:

- o PNC1 provides a type A grey topology abstraction (likewise in use case 2 described in [section 5.2](#))
- o PNC2 provides a type B grey topology abstraction (likewise in use case 3 described in [section 6.2](#))
- o PNC3 provides a type B grey topology abstraction with two abstract nodes, likewise in use case 3 described in [section 6.2](#), and hiding at least some optical parameters to be used within the OCh layer, likewise in use case 2 described in [section 5.2](#).

7.3. Service Configuration

The same service scenarios, as described in [section 6.3](#), are also applicable to these use cases with the only difference that single-domain end-to-end OCh data plane connections needs to be setup before ODU data plane connections.

8. Security Considerations

Typically, OTN networks ensure a high level of security and data privacy through hard partitioning of traffic onto isolated circuits.

There may be additional security considerations applied to specific use cases, but common security considerations do exist and these must be considered for controlling underlying infrastructure to deliver transport services:

- o use of RESCONF and the need to reuse security between RESTCONF components;
- o use of authentication and policy to govern which transport services may be requested by the user or application;

- o how secure and isolated connectivity may also be requested as an element of a service and mapped down to the OTN level.

9. IANA Considerations

This document requires no IANA actions.

10. References

10.1. Normative References

- [RFC7926] Farrel, A. et al., "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", [BCP 206](#), [RFC 7926](#), July 2016.
- [RFC4427] Mannie, E., Papadimitriou, D., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4427](#), March 2006.
- [ACTN-Frame] Ceccarelli, D., Lee, Y. et al., "Framework for Abstraction and Control of Transport Networks", [draft-ietf-teas-actn-framework](#), work in progress.
- [ITU-T G.709-2016] ITU-T Recommendation G.709 (06/16), "Interfaces for the optical transport network", June 2016.
- [ITU-T G.808.1-2014] ITU-T Recommendation G.808.1 (05/14), "Generic protection switching - Linear trail and subnetwork protection", May 2014.
- [ITU-T G.873.1-2014] ITU-T Recommendation G.873.1 (05/14), "Optical transport network (OTN): Linear protection", May 2014.

10.2. Informative References

- [TE-Topo] Liu, X. et al., "YANG Data Model for TE Topologies", [draft-ietf-teas-yang-te-topo](#), work in progress.
- [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.
- [Path-Compute] Busi, I., Belotti, S. et al., "Yang model for requesting Path Computation", [draft-busibel-teas-yang-path-computation](#), work in progress.

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

[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>

11. 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.

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

Authors' Addresses

Italo Busi (Editor)
Huawei
Email: italo.busi@huawei.com

Daniel King (Editor)
Lancaster University
Email: d.king@lancaster.ac.uk

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

Haomian Zheng
Huawei
Email: zhenghaomian@huawei.com