

TEAS Working Group
Internet Draft
Intended status: Informational
Intended status: Informational
Expires: April 2017

K. Lam
E. Varma
Nokia
P. Doolan
Coriant
N. Davis
Ciena
B. Zeuner
Deutsche Telekom
M. Betts
ZTE
I. Busi
Huawei
S. Mansfield
Ericsson
R. Vilalta
CTTC
V. Lopez
Telefonica
October 27, 2016

Usage of IM for network topology to support TE Topology YANG Module
Development
draft-lam-teas-usage-info-model-net-topology-04.txt

Status of this Memo

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

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may

not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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 April 28, 2017.

Copyright Notice

Copyright (c) 2016 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

The benefits of using a common Information Model (IM) as a foundation for deriving purpose and protocol specific interfaces, particularly

for complex networking domains, has been described in draft-betts-netmod-framework-data-schema-uml. This draft describes existing information model relevant to Network Topology and illustrates how it can be used to help ensure the consistency and completeness of the YANG data modeling for TE topologies solutions work in TEAS.

Table of Contents

1. Introduction.....	4
2. Background and Motivation.....	4
3. The Common Information Model.....	6
3.1. Core Model.....	7
3.1.1. Core Network Model.....	7
3.1.2. Core Foundation Model.....	9
3.1.3. Core Physical Model.....	12
3.1.4. Core Specification Model.....	13
3.2. Other Models.....	15
4. High Level Description of the Topology Subset of the CNM.....	15
4.1. Object Classes of the CNM Topology Subset.....	16
4.1.1. LogicalTerminationPoint (LTP) and LayerProtocol (LP).....	16
4.1.2. ForwardingDomain (FD).....	16
4.1.3. Link and Link Port.....	17
4.1.4. Network Element (NE).....	18
4.2. Relationships between Object Classes of the Topology Subset.....	18
4.2.1. ForwardingDomain Recursive Aggregation (HigherLevelFdEncompassesLowerLevelFds Aggregation).....	18
4.2.2. Network Elements encompassing ForwardingDomains (NeEncompassesFds Aggregation).....	19
4.2.3. ForwardingDomain association with LTPs (FdAggregatesLtpps Composition).....	21
4.2.4. ForwardingDomain aggregating Links (FdEncompassesLinks)	21
4.2.5. ForwardingDomain aggregating NEs.....	21
5. Detailed Description of the Topology Subset.....	21
5.1. Forwarding Entity.....	24
5.2. Characteristics of Topological Entity.....	25
5.2.1. Risk (RiskParameter_Pac).....	26
5.2.2. TransferCost_Pac.....	27
5.2.3. TransferTiming_Pac.....	28
5.2.4. TransferIntegrity_Pac.....	29

5.2.5. TransferCapacity_Pac.....	30
5.2.6. Validation_Pac.....	31
5.2.7. LayerProtocolTransition_Pac.....	31
6. Purpose Specific IM Example - Transport API Topology Service..	32
6.1. T-API IM Constructs.....	32
6.2. T-API Topology Service IM.....	34
7. Usage of the IM Topology Subset regarding TE Topology DM.....	35
8. Security Considerations.....	36
9. IANA Considerations.....	36
10. Conclusions.....	36
11. References.....	36
11.1. Normative References.....	36
11.2. Informative References.....	36
12. Contributors.....	38
13. Acknowledgments.....	38

1. Introduction

This draft describes existing information modeling (IM) relevant to Network Topology [ONF TR-512] [OSSDN SNOWMASS] and illustrates how it can be used to help ensure the consistency and completeness of the YANG data model (DM) for TE topologies solutions development work in TEAS.

2. Background and Motivation

Information Models (IM) and Data Models (DM) are related but different. An IM provides an abstract, conceptual view of the system being modeled in terms of its constituent parts (objects), independent of any specific implementations or protocols used to transport the data; it hides all protocol and implementation details (RFC 3444, TM Forum/NGCOR, ITU-T SG 15). A DM is a concrete specification in a particular language of an interface to, in this case, a controlled/managed system. The intention of the distinction between IMs and DMs has been to separate the modeling of problem space semantics from the modeling of the implementation of those semantics (though the dividing line has not always been clearly articulated).

A DM may be derived from an IM though it is often created without (explicit or obviously implicit) reference to one. When a DM is

derived from an IM, the DM and the components of the system it provides control/management access to are traceable to the definitions provided in the IM. There is no ambiguity between designer, developer, user or operator regarding the name, function, and information elements that are associated with a particular managed object.

As described in [I-D.betts], when DMs are created "in isolation" solely for the purpose of encoding specific interfaces, they may do that job adequately for any particular interface but in complex domains may create opportunities for confusion, duplication of effort, lack of interoperability, and lack of extensibility. In the past, ad-hoc development of DMs has caused significant operational and implementation inefficiencies in our industry.

Since March 2014, upon IESG recommendation that SNMP no longer be used for new work re configuration and that NETCONF/YANG be used instead, there has been an explosion of YANG DM development in IETF. It has consequently been recognized as essential to assure proper coordination of YANG DM development (including reaching out to different SDOs/consortia), as well as to assure that the YANG modules themselves provide a good representation of what is being modeled, to meet expectations of functionality, quality, and interoperability. In order to facilitate this objective, guidance from available pertinent IMs can be valuable.

This draft first describes an existing information model relevant to Network Topology [ONF TR-512], which is part of the Common Information Model (ONF-CIM) of network resources (as described in [I-D.betts]), that can be leveraged to assess the consistency and completeness of related YANG modules under development. It also describes an transport application-specific IM [OSSDN SNOWMASS], derived from CIM pruning and refactoring as explained in [I-D.betts], that is intended to enable further clarity in understanding the modeling. Being part of a Common Information Model, it will not lead to development of incompatible/uncoordinated models that can be difficult to maintain as other purpose-specific interfaces are developed.

3. The Common Information Model

This section provides a high level introduction to the ONF Common Information Model (ONF-CIM), and in particular its Core Model (see [ONF TR-512]), to provide an overall context for the topology relevant subset. The ONF-CIM has been developed through collaboration among several SDOs, including ITU-T, TM Forum, and ONF, and also published as ITU-T Recommendation G.7711 [G.7711].

An information model describes the things in a domain in terms of objects, their properties (represented as attributes), and their relationships.

The ONF-CIM is expressed in a formal language called UML (Unified Modeling Language). UML has a number of basic model elements, called UML artifacts. In order to assure a consistent and harmonized modeling approach, only a selected subset of these UML artifacts were used in the development of the ONF-CIM according to guidelines for creating an information model expressed in UML (see the UML Guidelines document in the ONF TR-514 [ONF TR-514]).

The ONF-CIM has been developed using the Papyrus open source UML Tool, for which a detailed guidelines document is available (see the Papyrus Guidelines document in the ONF TR-515 [ONF TR-515]). This guidelines document also describes how the modelers constructing the ONF-CIM can cooperate in the GitHub environment to allow for separate and still coordinated development of the ONF-CIM fragments.

The ONF-CIM includes all of the artifacts (objects, attributes, associations, etc.) that are necessary to describe the domain for the applications being developed.

It will be necessary to continually expand and refine the ONF-CIM over time as, for example to add, new applications, capabilities or forwarding technologies, or to refine the ONF-CIM as new insights are gained. To allow these extensions to be made in a seamless manner, the ONF-CIM is structured into a number of sub-models. This modeling approach enables application specific and forwarding technology specific extensions to be developed by domain experts with appropriate independence. This approach is further articulated in ONF TR-513 [ONF TR-513] and [I-D.betts].

3.1. Core Model

The Core Model of the ONF-CIM consists of model artifacts that are intended for use by multiple applications and/or forwarding technologies.

For navigability, the Core Model is further sub-structured into sub-models. Currently, these consist of the Core Network Model (CNM), Core Foundation Model, Core Physical Model, and the Core Specification Model. The following sub-sections provide an overview of these sub-models. A detailed description is contained in ONF TR-512 [ONF TR-512].

3.1.1. Core Network Model

The Core Network Model (CNM) consists of artifacts that model the essential network aspects that are neutral to the forwarding technology of the network. The CNM currently encompasses Topology, Termination, and Forwarding aspects (subsets of the CNM) as described below:

- Topology Subset of CNM

The Topology subset of the CNM supports the modeling of network topology information, which can be used to build the topology database and depict the topology. Object classes representing topological entities include:

- o Forwarding Domain (FD): Offers the potential to enable forwarding of information.
- o Link (L): Models the adjacency between two or more FDs. A Link has LinkPorts.
- o Logical Termination Point (LTP): Models the ports of a link. It encapsulates the termination, adaptation, and OAM functions of one or more transport layers.
- o Network Element (NE): While not actually part of topology, a NE brings meaning to the FD and the LTP contexts (and hence the links). A NE represents physical equipment "bundling" to

provide a view of management scope, management access, and session.

The Topology subset of the CNM supports network topology abstraction and virtualization. FD abstraction is supported via recursive aggregation and virtualization via partitioning of resources according to the resource dedication criterion.

- Forwarding Subset of CNM

The Forwarding subset of the CNM (not covered in detail in this draft) supports configuration of forwarding entities, including their setup, modification, and tear down. Artifacts representing the forwarding construct include:

- o ForwardingConstruct (FC): Also known as SNC. In conjunction with the FcPort, FC models the enabled forwarding between two FcPorts across a FD.
- o FcPort: Models the access to the FC, and associates the FC to the LTP. When the FC supports protection, the FcPort also indicates its role in the protection scheme, i.e., whether it is a working or protection FcPort.
- o FcRoute: Also known as SncRoute. It models the individual routes of an FC.
- o FcSwitch: Also known as SncSwitch. It models the switched forwarding of traffic (traffic flow) between EPs and is present where there is protection functionality in the FD.

- Termination Subset of CNM

The Termination subset of the CNM (not covered in detail in this draft) supports modeling of the processing of transport characteristic information, such as termination, adaptation, OAM, etc. Artifacts representing the termination and adaptation and OAM construct include:

- o Logical Termination Point (LTP): See the LTP description in the Topology Subset

- o Layer Protocol (LP): This identifies the type of signal and is the anchor for transport layer protocol specific definitions, which are modeled in, e.g., [G.874.1] for OTN, [G.8052] for transport Ethernet, and [G.8152] for MPLS-TP.

- Resilience Subset of CNM

The Resilience subset provides a view of the model for resilience (including protection and restoration) and encompasses:

- o The basic resilience model structure
- o The key attributes relevant to resilience
- o The application of the resilience model to various cases

3.1.2. Core Foundation Model

To communicate about an entity, it is important to have some way of referring to that entity, i.e., to have some way of referencing it. The Core Foundation model defines the artifacts for referencing entities; i.e.:

- Global Unique ID (GUID):

An identifier that is globally unique where an identifier is a property of an entity/role with a value that is unique within an identifier space, where the identifier space is itself unique, and immutable. The identifier therefore represents the identity of the entity/role. An identifier carries no semantics with respect to the purpose of the entity.)

- Local ID:

An identifier that is unique in the context of some scope that is less than the global scope (where an identifier is as defined in GUID above).

- Name:

A property of an entity with a value that is unique in some namespace but may change during the life of the entity. A name carries no semantics with respect to the purpose of the entity.

- Label:

A property of an entity with a value that is not expected to be unique and is allowed to change. A label carries no semantics with respect to the purpose of the entity and has no effect on the entity behavior or state.

The Core Foundation model also provides the opportunity to extend any entity using the Extension structure.

The model also defines two foundation object classes:

- GlobalClass:

Super class of object classes for which their instances can exist on their own right, e.g. NE, LTP, FD, Link, and FC. Global classes shall have one and only one globally unique identifier (GUID) and may have zero or more local identifiers, zero or more names, zero or more labels, zero or more extensions.

- LocalClass:

Super class of object classes for which the existence of their instances depends on instances of global classes; e.g., LP (of LTP), EP (of FC), and LE (of Link). Local classes shall have at least one local identifier, may have zero or more names, zero or more labels, zero or more extensions.

```
-----
I                                     I
I                                     I
I                                     I
I                                     I
I          Artifacts for Referencing of Entities      I
I                                                     I
```

I		I
I		I
I		I
I	(only in PDF version)	I
I		I
I		I
I		I
I		I
I		I
I		I
I		I

Figure 3-1 Artifacts for Referencing of Entities

The Core Foundation model also defines a State_Pac artifact, which is a package of state attributes. The State_Pac is inherited by GlobalClass and LocalClass object classes. The State_Pac consists of the following state-related attributes:

- Operational State:
 - Read-only with values: DISABLED, ENABLED
- Administrative State:
 - Read-write with values: LOCKED, UNLOCKED
- Lifecycle State:
 - Read-write with values: PLANNED, POTENTIAL, INSTALLED, PENDING_REMOVAL

I		I
I		I

I		I
I	States of Objects	I
I		I
I		I
I	(only in PDF version)	I
I		I
I		I
I		I
I		I

Figure 3-2 States of Objects

3.1.1.3. Core Physical Model

The Physical model provides a view of the model for physical entities (including equipment, holders and connectors). This model also specifies the relationship between the connector and the LTP, and the relationship between physical and functional views.

I		I
I		I
I		I
I		I
I	Basic Equipment Pattern	I
I		I
I		I
I		I
I		I
I	(only in PDF version)	I
I		I

I I
I I
I I
I I

Figure 3-3 Basic Equipment Pattern

3.1.1.4. Core Specification Model

There are several related needs that have given rise to the Specification model:

- Provide machine readable form of specific localized behavior:
 - o Representing rules related to restrictions of specific cases of use of the model
 - o Representing capabilities of specific cases of use
- Enable the introduction of run time schema where the essential structure of the model is known up front (at compile time) but the details are not
- Reduce the clutter in a representation where a set of details take the same values for all instances that are related to a specific case
- Allow leverage of existing standards definitions (e.g., technology/application specific) in a machine readable language

The combination of the above resulted in a separation in the model of definitions of structure and content such that an instance of a class from one model fragment could have an association instance to another model fragment to enable the provision of a fragment of definition of the class and of subordinates.

The aim of all specification definitions is that they be rigorous definitions of specific cases of usage and enable machine

interpretation where traditional interface designs would only allow human interpretation.

The following dedicated spec structures have been considered:

- FC spec: Main focus to provide a representation of the effective internal structure of a ForwardingConstruct (FC)
- LTP and LP spec: Main focus to provide a representation of Layer Protocol (LP) specific parameters for the Logical Termination Point (LTP)
- FD and Link spec: Main focus on capacity and forwarding enablement restrictions
- Equipment spec: Main focus to provide a representation of equipping constraints

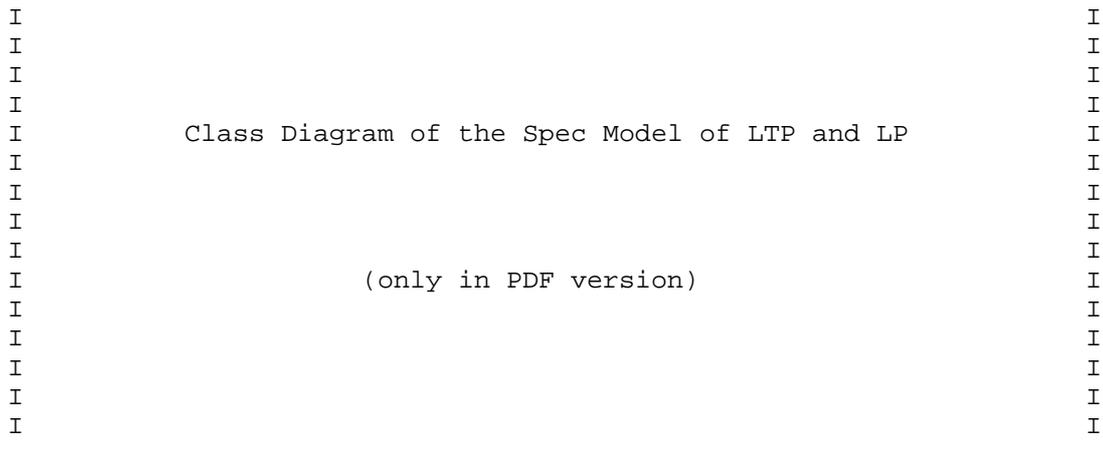


Figure 3-4 Class Diagram of the Spec Model of LTP and LP

3.2. Other Models

In addition to the Core Model, the ONF-CIM includes forwarding technology and application specific models. The forwarding technology models of the ONF-CIM (see [ONF TR-512]) encompasses transport technology layers 0, 1, and 2.

4. High Level Description of the Topology Subset of the CNM

This section provides a high-level overview of the Topology Subset of the CNM. Figure 4-1 below is a skeleton class diagram illustrating the key object classes. To avoid cluttering the figure, not all associations have been shown and all of the attributes were omitted.

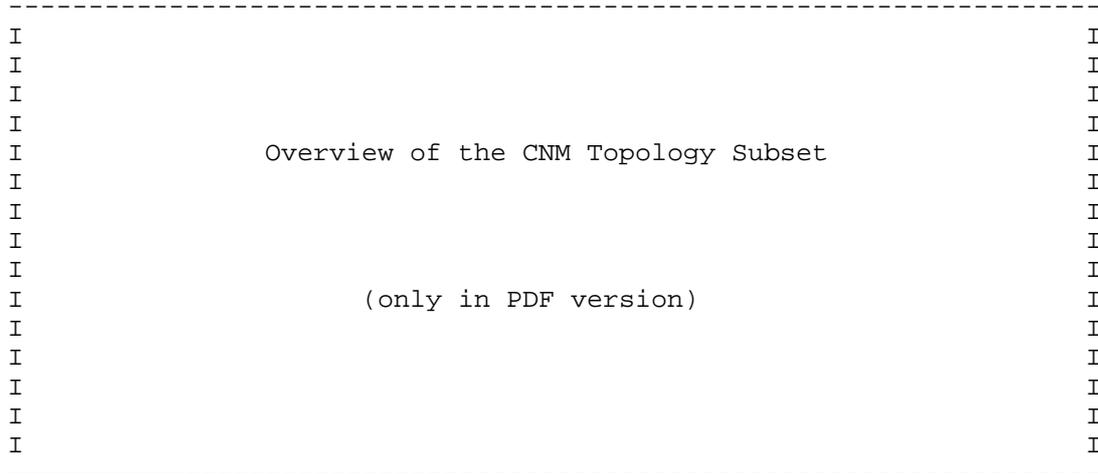


Figure 4-1 Overview of the CNM Topology Subset

4.1. Object Classes of the CNM Topology Subset

This section describes the object classes of the Topology Subset of the CNM. Relationships between these classes are described in section 4.2 below

4.1.1. LogicalTerminationPoint (LTP) and LayerProtocol (LP)

The LogicalTerminationPoint (LTP) object class encapsulates the termination, adaptation and OAM functions of one or more transport protocol layers. The structure of the LTP supports all transport protocols including circuit and packet forms. Each transport layer is represented by a LayerProtocol (LP) instance. The LayerProtocol instances of the LTP can be used for controlling the termination and OAM functionality of that layer. It can also be used for controlling the adaptation (i.e. encapsulation and/or multiplexing of client signal). Where the client/server relationship is fixed 1:1 and immutable, the different layers can be encapsulated in a single LTP instance. Where there is a n:1 relationship between client and server, the layers must be split over separate instances of LTP.

The LP object class is defined with generic attributes "layerProtocolName" for indicating the supported transport layer protocol.

Transport layer specific properties (such as layer-specific termination and adaptation properties) are modeled as attributes of conditional packages (called "_Pacs" in the UML notation of the ONF-CIM) associated with the LP object class.

4.1.2. ForwardingDomain (FD)

The ForwardingDomain (FD) object class models the switching and routing capabilities (see "subnetwork" topological component in [G.852.2] and [TMF612]), which is used to effect forwarding of transport characteristic information and offers the potential to enable forwarding. It represents the resource that supports flows across the FD. The FD object can hold zero or more instances of ForwardingConstruct (FC) (representing constrained forwarding, not discussed further in this document, covering connections, VLANs etc) of one or more layer networks; e.g., OCh, ODU, ETH, and MPLS-TP. The

FD object provides the context for operations that create/modify/delete FCs.

The FD object class supports a recursive aggregation relationship such that the internal construction of an FD can be exposed as multiple lower level FDs and associated Links (partitioning) (see section 4.2.1.)

At the lowest level of recursion, a FD (within a network element) could represent a switch matrix (i.e., a fabric).

Note that an NE can encompass multiple switch matrices (FDs), as described in section 4.2.2. An instance of FD is associated with zero or more LTP objects, as described in section 4.2.3.

4.1.3. Link and Link Port

The Link object class models the adjacency between two or more ForwardingDomains (FDs).

In its basic form (i.e., point-to-point Link) it associates a set of LTP clients on one FD with an equivalent set of LTP clients on another FD. Like the FC, the Link has endpoints (LinkPort) which take roles in the context of the function of the Link. A point-to-point Link can be a TE Link and support parameters such as capacity, delay etc. These parameters depend on the type of technology that supports the link.

A Link can be terminated on two or more FDs. This provides support for technologies such as PON and Layer 2 MAC in MAC configurations.

The LinkPort further details the relationship between FD and Link for asymmetric cases.

A FD may aggregate Links (see section 4.2.5).

The Link can support multiple transport layers via the associated LTP object. An instance of Link can be formed with the necessary properties according to the degree of virtualization. For implementation optimization, multiple layer-specific links can be merged and represented as a single Link instance.

4.1.4. Network Element (NE)

The NetworkElement (NE) object class represents a network element (traditional NE) in the data plane or a virtual network element visible in an interface where virtualization is used.

In the direct interface from a SDN controller to a network element in the data plane, the NE object defines the scope of control for the resources within the network element, e.g., internal transfer of user information between the external terminations (ports), encapsulation, multiplexing/demultiplexing, and OAM functions, etc. The NE provides the scope of the naming space for identifying objects representing the resources within the network element.

Where virtualization is employed, the NE object represents a virtual NE (VNE). The mapping of the VNE to the NEs is the internal matter of the SDN controller that offers the view of the VNE. Via the interface between hierarchical SDN controllers, NE instances can be created (or deleted) for providing (or removing) virtual views of the combination of slices of network elements in the data plane.

4.2. Relationships between Object Classes of the Topology Subset

4.2.1. ForwardingDomain Recursive Aggregation (HigherLevelFdEncompassesLowerLevelFds Aggregation)

Figure 4-2 below provides a pictorial example of ForwardingDomain (FD) recursion with Links.



I	ForwardingDomain recursion with Links	I
I		I
I		I
I		I
I		I
I	(only in PDF version)	I
I		I
I		I
I		I
I		I
I		I

Figure 4-2 ForwardingDomain recursion with Links

Figure 4-2 shows a UML fragment including the Link and ForwardingDomain (FD). For simplicity it is assumed here that the Links and FDs are for a single LayerProtocol (LP) although it can be seen from the detailed figure earlier in this section that both a FD and link can support a list of LPs.

The pictorial form shows a number of instances of FD interconnected by Links and shows nesting of FDs. The recursive aggregation "HigherLevelFdEncompassesLowerLevelFds" relationship (represented by an open diamond) supports the FD nesting but it should be noted that this is intentionally showing no lifecycle dependency between the lower FDs and the higher ones that nest them (to do this composition, a black diamond would have been used instead of the open diamond). This is to allow for rearrangements of the FD hierarchy (e.g. when regions of a network are split or merged). This emphasizes that the nesting is an abstraction rather than decomposition. The underlying network still operates regardless of how it is perceived in terms of aggregating FDs. The model allows for only one hierarchy.

4.2.2. Network Elements encompassing ForwardingDomains (NeEncompassesFds Aggregation)

Figure 4-3 below provides a pictorial example of ForwardingDomain (FD) recursion with Links and NEs.

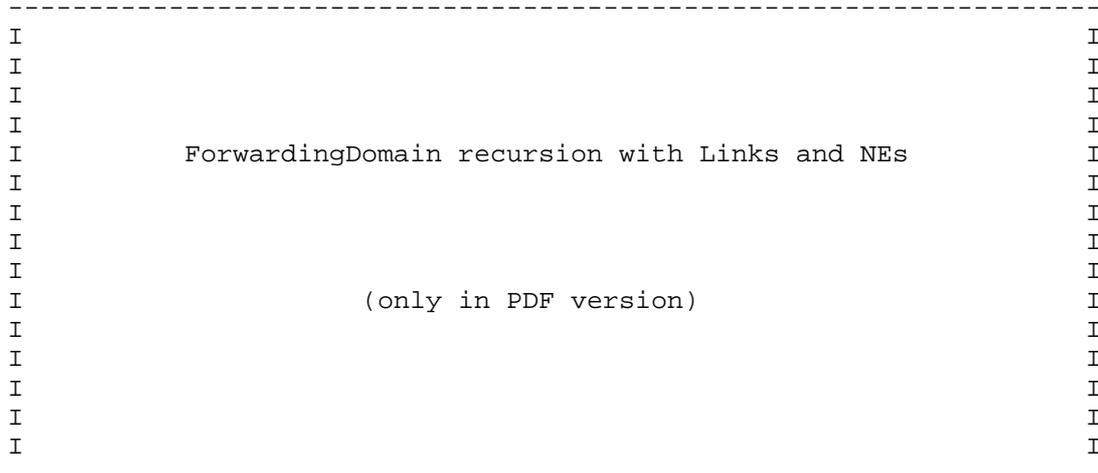


Figure 4-3 ForwardingDomain recursion with Links and NEs

Figure 4-3 above shows an overlay of NetworkElement (NE) on the ForwardingDomains and a corresponding fragment of UML showing only the ForwardingDomain and NetworkElement classes.

The figure emphasizes that one level of abstraction of ForwardingDomain is bounded by an NE. This is represented in the UML fragment by the composition association (black diamond) that explains that there is a lifecycle dependency in that the ForwardingDomain at this level that cannot exist without the NE. The figure also shows that a ForwardingDomain need not be bounded by an NE (as explained in the UML fragment by the 0..1 composition) and that a ForwardingDomain may have smaller scope than the whole NE (even when considering only a single LayerProtocol as described below).

In one of the cases depicted (e.g., the right hand side NE encompassing two FDs), the two ForwardingDomains in the NE are completely independent. In the other cases depicted (e.g., the left

hand side NE encompassing three FDs) the subordinate ForwardingDomains are themselves joined by Links emphasizing that the NE does not necessarily represent the lowest level of relevant network decomposition.

The figure also emphasizes that just because one ForwardingDomain at a particular level of decomposition of the network happens to be the one bounded by an NE does not mean that all ForwardingDomains at that level are also bounded by NEs.

4.2.3. ForwardingDomain association with LTPs (FdAggregatesLtps Composition)

An instance of FD is associated with zero or more LTP objects via the "FdAggregatesLtps" composition.

4.2.4. ForwardingDomain aggregating Links (FdEncompassesLinks)

A ForwardingDomain can aggregate links. An example of ForwardingDomain Recursive Aggregation with Links is shown in section 4.2.1 above.

However, the FdAggregatesLink association is not modeled because this association can be inferred from the higherLevelFdContainsLowerLevelFd association together with the linkHasAssociatedFds association.

4.2.5. ForwardingDomain aggregating NEs

A ForwardingDomain can aggregate Network Elements. An example of ForwardingDomain Recursive Aggregation with Links and NEs is shown in section 4.2.2 above.

However, the FdAggregatesNe association is not modeled because this association can be inferred from higherLevelFdContainsLowerLevelFd association and together with the NeEncompassesFd association.

5. Detailed Description of the Topology Subset

The two key classes related to Topology are the ForwardingDomain (FD) and the Link. For simple cases the FD represents the switching

capability in the network and the Link represents adjacency. These are depicted in the context of other model classes in Figure 5-1.

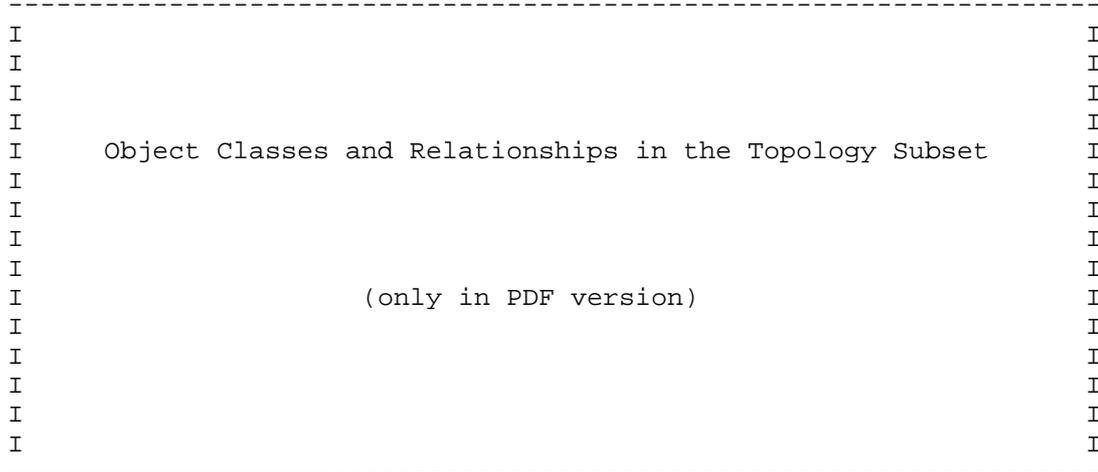


Figure 5-1 Object Classes and Relationships in the Topology Subset

Figure 5-1 shows a lightweight view of the model omitting the attributes (where appropriate these will be described later in this section).

The FD and Link will be described in detail later in the document. Figure 5-1 focuses on interrelationships and these will be the focus of this section. The figure shows that:

- An FD may be a subordinate part of a NetworkElement (NE) or may be larger than, and independent of, any NE.
- An FD may encompass lower level FDs. This may be such that:
 - o A FD directly contained in an NE is divided into smaller parts

- o A FD not encompassed by an NE is divided into smaller parts some of which may be encompassed by NES
- o The FD represents the whole network
- An FD encompasses Links that interconnect any FDs encompassed by the FD
- A Link may aggregate Links in several ways
 - o In parallel where several links are considered as one
 - o In series where Links chain to form a Link of a greater span
 - . Note that this case requires further development in the model
- A Link has associated FDs that it interconnects
 - o A Link may interconnect 2 or more FDs
 - . Note that it is usual for a Link to interconnect 2 FDs but there are cases where many FDs may be interconnected by a Link
- A Link has LinkPorts that represent the ports of the Link itself
 - o LinkPorts are especially relevant for multi-ended asymmetric Link
- A LinkPort aggregates LogicalTerminationPoints (LTPs) that bound the Link. The LTP represent a stack LayerProtocol terminations where the details of each is held in the LayerProtocol (LP). The LTP may be:
 - o Part of an NE
 - o Conceptually independent from any NE

- A LinkPort references LTPs on which the Link associated to the LE terminates

Both the Link and FD are subclasses of ForwardingEntity (an abstract class, i.e. a class that will never be instantiated) and hence they can acquire contents from the conditional packages (`_Pacs`). The conditional packages provide all key topology properties.

5.1. Forwarding Entity

As noted in the previous section the two key topology classes are Forwarding Domain (FD) and Link (L).

The FD topological component is used to show the potential to enable forwarding. At the lowest level of recursion, an FD (within a network element (NE)) represents a switch matrix (e.g., a fabric). Note that an NE can encompass multiple switch matrices (FDs).

As noted earlier the Link models adjacency between two or more Forwarding Domains (FD).

Both the link and the FD have the potential to handle more than one layerProtocol (both have a `layerProtocolNameList` attribute).

As shown in Figure 5-1 an object class "ForwardingEntity" has been defined to collect topology-related properties (characteristics etc.) that are common for FD and Link.

A ForwardingEntity is an abstract representation of the emergent effect of the combined functioning of an arrangement of components (running hardware, software running on hardware, etc). The effect can be considered as the realization of the potential for apparent communication adjacency for entities that are bound to the terminations at the boundary of the ForwardingEntity.

The ForwardingEntity enables the creation of constrained forwarding to achieve the apparent adjacency. The apparent adjacency has intended performance degraded from perfect adjacency and a statement of that degradation is conveyed via the attributes of the packages associated with this class. In the model both ForwardingDomain and Link are ForwardingEntities.

This abstract class is used as a modeling approach to apply packages of attributes to both Link and ForwardingDomain. Link and ForwardingDomain are the key ForwardingEntities.

5.2. Characteristics of Topological Entity

As noted above the characteristic of a TopologicalEntity are covered by the conditional packages (_PACs).

I I
I I
I I
I I
I Conditional Packages of Topological Entity I
I I
I I
I I
I I
I (only in PDF version) I
I I
I I
I I
I I
I I
I I
I I

Figure 5-2 Conditional Packages of Topological Entity

5.2.1. Risk (RiskParameter_Pac)

The risk characteristics of a ForwardingEntity come directly from the underlying physical realization.

The risk characteristics propagate from the physical realization to the client and from the server layer to the client layer, this propagation may be modified by protection.

A ForwardingEntity may suffer degradation or failure as a result of a problem in a part of the underlying realization.

The realization can be partitioned into segments which have some relevant common failure modes.

There is a risk of failure/degradation of each segment of the underlying realization.

Each segment is a part of a larger physical/geographical unit that behaves as one with respect to failure (i.e. a failure will have a high probability of impacting the whole unit (e.g. all fibers in the same cable)).

Disruptions to that larger physical/geographical unit will impact (cause failure/errors to) all ForwardingEntities that use any part of that larger physical/geographical entity.

Any ForwardingEntity that uses any part of that larger physical/geographical unit will suffer impact and hence each ForwardingEntity shares risk.

The identifier of each physical/geographical unit that is involved in the realization of each segment of a Topological entity can be listed in the RiskParameter_Pac of that ForwardingEntity.

A segment has one or more risk characteristic.

Shared risk between two ForwardingEntities compromises the integrity of any solution that use one of those ForwardingEntity as a backup for the other.

Where two ForwardingEntities have a common risk characteristic they have an elevated probability of failing simultaneously compared to two ForwardingEntities that do not share risk characteristics.

- riskCharacteristicList: A list of risk characteristics (RiskCharacteristic) for consideration in an analysis of shared risk. Each element of the list represents a specific risk consideration.
- RiskCharacteristic: The information for a particular risk characteristic where there is a list of risk identifiers related to that characteristic. It includes:
 - o riskCharacteristicName: The name of the risk characteristic. The characteristic may be related to a specific degree of closeness. For example a particular characteristic may apply to failures that are localized (e.g. to one side of a road) where as another characteristic may relate to failures that have a broader impact (e.g. both sides of a road that crosses a bridge). Depending upon the importance of the traffic being routed different risk characteristics will be evaluated.
 - o riskIdentifierList: A list of the identifiers of each physical/geographic unit (with the specific risk characteristic) that is related to a segment of the ForwardingEntity.

5.2.2. TransferCost_Pac

The cost characteristics of a ForwardingEntity not necessarily correlated to the cost of the underlying physical realization.

They may be quite specific to the individual ForwardingEntity e.g. opportunity cost. Relates to layer capacity

There may be many perspectives from which cost may be considered for a particular ForwardingEntity and hence many specific costs and potentially cost algorithms.

Using an entity will incur a cost.

- costCharacteristicList: The list of costs (CostCharacteristic) where each cost relates to some aspect of the Link
 - o CostCharacteristic: The information for a particular cost characteristic
 - . costName: The cost characteristic will related to some aspect of the ForwardingEntity (e.g. \$ cost, routing weight). This aspect will be conveyed by the costName
 - . costValue: The specific cost.
 - . costAlgorithm: The cost may vary based upon some properties of the ForwardingEntity. The rules for the variation are conveyed by the costAlgorithm.

5.2.3. TransferTiming_Pac

A link will suffer effects from the underlying physical realization related to the timing of the information passed by the link.

- fixedLatencyCharacteristic: A ForwardingEntity suffers delay caused by the realization of the servers (e.g. distance related; FEC encoding etc.) along with some client specific processing. This is the total average latency effect of the ForwardingEntity
- jitterCharacteristic: High frequency deviation from true periodicity of a signal and therefore a small high rate of change of transfer latency. Applies to TDM systems (i.e., not packet based systems).
- wanderCharacteristics: Low frequency deviation from true periodicity of a signal and therefore a small low rate of

change of transfer latency. Applies to TDM systems (i.e., not packet based systems).

- queuingLatencyList: The effect on the latency of a queuing process. This only has significant effect for packet based systems and has a complex characteristic (QueuingLatency).
 - o QueuingLatency: Provides information on latency characteristic for a particular stated trafficProperty.

5.2.4. TransferIntegrity_Pac

Transfer integrity characteristic covers expected (specified) error, loss and duplication signal content as well as any damage of any form to total link and to the client signals.

- errorCharacteristic: describes the degree to which the signal propagated can be errored. Applies to TDM systems as the errored signal will be propagated and not packet as errored packets will be discarded.
- lossCharacteristic: Describes the acceptable characteristic of lost packets where loss may result from discard due to errors or overflow. Applies to packet systems and not TDM (as for TDM errored signals are propagated unless grossly errored and overflow/underflow turns into timing slips).
- repeatDeliveryCharacteristic: Primarily applies to packet systems where a packet may be delivered more than once (in fault recovery for example). It can also apply to TDM where several frames may be received twice due to switching in a system with a large differential propagation delay.
- deliveryOrderCharacteristic: Describes the degree to which packets will be delivered out of sequence. Does not apply to TDM as the TDM protocols maintain strict order.
- unavailableTimeCharacteristic: Describes the duration for which there may be no valid signal propagated.

- `serverIntegrityProcessCharacteristic`: Describes the effect of any server integrity enhancement process on the characteristics of the `ForwardingEntity`.

5.2.5. `TransferCapacity_Pac`

The `ForwardingEntity` derives capacity from the underlying realization.

A `ForwardingEntity` may be an abstraction and virtualization of a subset of the underlying capability offered in a view or may be directly reflecting the underlying realization.

A `ForwardingEntity` may be directly used in the view or may be assigned to another view for use.

The clients supported by a multi-layer `ForwardingEntity` may interact such that the resources used by one client may impact those available to another. This is derived from the LTP spec details.

A `ForwardingEntity` represents the capacity available to user (client) along with client interaction and usage.

A `ForwardingEntity` may reflect one or more client protocols and one or more members for each profile.

- `totalPotentialCapacity`: A "best case" view of the capacity of the `ForwardingEntity` assuming that any shared capacity is available to be taken.

Note that this area is still under development to cover concepts such as:

- `exclusiveCapacityList`: The capacity allocated to this `ForwardingEntity` for its exclusive use
- `sharedCapacityList`: The capacity allocated to this `ForwardingEntity` that is not exclusively available as it is shared with others.

- `assignedAsExclusiveCapacityList`: The capacity assigned from this `TopologicalEntity` to another `ForwardingEntity` for its exclusive use
- `assignedAsSharedCapacityList`: The capacity assigned to one or more other `ForwardingEntities` for shared use where the interaction follows some stated algorithm.
- Capacity which includes:
 - o `totalSize`
 - o `numberOfUsageInstances`
 - o `maximumUsageSize`
 - o `numberingRange`

5.2.6. `Validation_Pac`

Validation covers the various adjacent discovery and reachability verification protocols. Also may cover Information source and degree of integrity.

- `validationMechanismList`: Provides details of the specific validation mechanism(s) used to confirm the presence of an intended `ForwardingEntity`.

5.2.7. `LayerProtocolTransition_Pac`

Relevant for a Link that is formed by abstracting one or more LTPs (in a stack) to focus on the flow and deemphasize the protocol transformation.

This abstraction is relevant when considering multi-layer routing.

The layer protocols of the LTP and the order of their application to the signal is still relevant and need to be accounted for. This is derived from the LTP spec details.

This Pac provides the relevant abstractions of the LTPs and provides the necessary association to the LTPs involved.

Links that included details in this Pac are often referred to as Transitional Links.

- transitionedLayerProtocolList: Provides the ordered structure of layer protocol transitions encapsulated in the ForwardingEntity. The ordering relates to the LinkEnd role.

6. Purpose Specific IM Example - Transport API Topology Service

In order to provide some further clarity, this section provides a high level introduction to a Purpose Specific IM, the Transport API (T-API) Topology service, which has been derived from the ONF Common Information Model (ONF-CIM) according to the principles in [I-D.betts].

The context of the T-API refers to the scope and control and naming that a particular SDN controller, manager or a client application has with respect to the information it operates on internally or exchanges over an interface. The following sections further describe this purpose specific IM and relationship to the ONF-CIM.

6.1. T-API IM Constructs

The T-API IM uses terminology that is considered to be more familiar to the transport network management community and maps to the constructs defined in the ONF-CIM CNM Topology model. The following table provides a high level summary of the mapping of the constructs relevant to the T-API Topology Service.

Mapping of CIM and T-API IM Constructs			
	ONF-CIM CNM Terminology	T-API IM Terminology	
	NetworkControlDomain	Context	ForwardingDomain
(FD)	Node	Topology	
	TransitionalLink	Link	Link
	nationPoint (LTP) ServiceEndPoint	NodeEdgePoint	LogicalTermi

The following provides a brief description of these T-API IM constructs.

- o Link: A Link is an abstract representation of the effective adjacency between two or more associated Nodes in a Topology. It is terminated by Node-Edge-Points of the associated Nodes.
- o Node: A Node is an abstract representation of the forwarding-capabilities of a particular set of Network Resources. It is described in terms of an aggregation of set of ports (Node-Edge-Point) belonging to those Network Resources and the potential to enable forwarding of information between those edge ports.
- o Node-Edge-Point: A Node-Edge-Point represents the inward network-facing aspects of the edge-port functions that access the forwarding capabilities provided by the Node. Hence it provides an encapsulation of addressing, mapping, termination, adaptation and OAM functions of one or more transport layers (including circuit and packet forms) performed at the entry and exit points of the Node.
- o Topology: A Topology is an abstract representation of the topological-aspects of a particular set of Network Resources. It is described in terms of a network of set of Nodes and Links that enable the forwarding-capabilities of that particular set of Network Resources.
- o Service-End-Point: A Service-End-Point represents the outward customer-facing aspects of the edge-port functions that access the forwarding capabilities provided by the Node. Hence it provides a limited, simplified view of interest to external clients (e.g. shared addressing, capacity, resource availability, etc) that enable the clients to request connectivity without the need to understand the provider network internals.
- o Transitional Link: A topological component that consists of the link port at the edge of one node and a corresponding link port at the edge of another node that operates on different layers or whose layer is the same but with different Layer

Information. A transitional link is supported/implemented by transport processing functions (e.g., adaptation/termination). A transitional link can be partitioned into parallel transitional links, or a concatenation of transitional links. It can also be partitioned into a concatenation of transitional links and zero or more links.

6.2. T-API Topology Service IM

The resultant high-level description for the T-API Topology Service constructs, based upon the pruned and refactored ONF-CIM, and the related Topology Service APIs are provided in Figure 6-1 below.

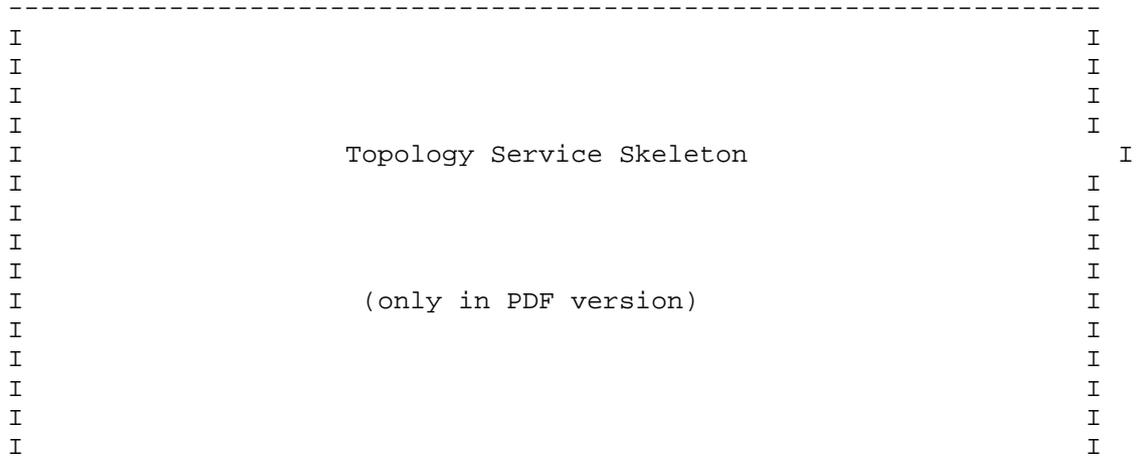


Figure 6-1 Topology Service Skeleton

The T-API Topology Service API enables the API client to, for example, retrieve Topology, Node, Link, and Edge-Point details.

- o Topology details: returns attributes of the Topology identified by the provided input ID. This includes references to lower-level Nodes and Links encompassed by that Topology. A NULL input value is expected to return the top-most Topology that corresponds to the scope of the entire Context including any Off-Network-Links.
- o Node details: Returns attributes of the Node identified by the provided input ID. Includes references to Node-Edge-Points aggregated by the Node, and attributes representing the identification, naming, states and forwarding capabilities of the Node.
- o Link details: Returns attributes of the Link identified by the provided input ID. Includes references to Node-Edge-Points terminating the Link, and references to the Nodes associated by the Link.
- o Node-Edge-Point details: Returns attributes of the Node-Edge-Point identified by the provided input ID, including references to Service-End-Points mapped to this Node-Edge-Point.

The API supports a retrieve-scope filter: LayerProtocol list. If set, the API call will return output that is relevant to the specified Layer only.

7. Usage of the IM Topology Subset regarding TE Topology DM

As discussed earlier, a data model (DM) may be derived from an IM. Examples of YANG DMs derived according to automated translation tools based upon mapping guidelines are provided in [OSSDN SNOMASS] at <https://github.com/OpenNetworkingFoundation/Snowmass-ONFOpenTransport/tree/develop/YANG>. It is possible to leverage the IM

Topology Subset to assess the consistency and completeness of related YANG modules under development.

8. Security Considerations

This informational document is intended only to provide a description of an interface-protocol-neutral information model, and the security concerns are therefore out of the scope of this document.

9. IANA Considerations

This document includes no request to IANA.

10. Conclusions

The information modeling described in this draft, which is relevant to Network Topology [ONF TR-512] [OSSDN SNOWMASS], can be leveraged in assessing the consistency and completeness of related YANG modules under development.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

11.2. Informative References

[I-D.betts] Betts, M., Davis, N., Lam, K., Zeuner, B., Mansfield, S. and P. Doolan, "Framework for Deriving Interface Data Schema from UML Information Models", draft-betts-netmod-framework-data-schema-uml-04 (work in progress), October 2016

[I-D.mansfield] Mansfield, S., Zeuner, B., Davis, N., Yun, X., Tochio, Y., Lam, K. and E. Varma, "Guidelines for Translation of UML Information Model to YANG Data Model", draft-mansfield-netmod-uml-to-yang-03 (work in progress), October 2016

- [ONF TR-512] ONF TR-512 "ONF-CIM Core Model base document 1.2"
([https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-512_CIM_\(CoreModel\)_1.2.zip](https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-512_CIM_(CoreModel)_1.2.zip)),
September 2016
- [ONF TR-513] ONF TR-513 " Common Information Model Overview 1.2"
(https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-513_CIM_Overview_1.2.pdf),
September 2016
- [ONF TR-514] ONF TR-514 "UML Modeling Guidelines 1.2"
(https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-514_UML_Modeling_Guidelines_v1.2.pdf), September 2016
- [ONF TR-515] ONF TR-515 "Papyrus Guidelines 1.2"
(https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-515_Papyrus_Guidelines_v1.2.pdf), September 2016
- [OSSDN SNOWMASS] Open Source SDN SNOWMASS/Open Transport API
Specifications
(<https://github.com/OpenNetworkingFoundation/Snowmass-ONFOpenTransport>)
- [G.7711] Recommendation ITU-T G.7711/Y.17022 "Generic protocol-neutral information model for transport resources",
September 2016
- [G.874.1] Recommendation ITU-T G.874.1 "Optical transport network: Protocol-neutral management information model for the network element view", September 2016
- [G.8052] Recommendation ITU-T G.8052/Y.1346 "Protocol-neutral management information model for the Ethernet Transport capable network element", September 2016
- [G.8152] Recommendation ITU-T G.8152/Y.1375 "Protocol-neutral management information model for the MPLS-TP network element", September 2016

[G.852.2] Recommendation ITU-T G.852.2 "Enterprise viewpoint
description of transport network resource model", March 1999

[TMF612] TM Forum 612 "MTOSI Information Agreement", October 2014

12. Contributors

Karthik Sethuraman
NEC

Email: karthik.sethuraman@necam.com

13. Acknowledgments

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

Authors' Addresses

Kam Lam
Alcatel-Lucent, USA

Phone: +1 732 331 3476
Email: kam.lam@alcatel-lucent.com

Eve Varma
Alcatel-Lucent, USA

Email: eve.varma@alcatel-lucent.com

Paul Doolan
Coriant, Germany

Phone: +1 972 357 5822
Email: paul.doolan@coriant.com

Malcolm Betts
ZTE, China

Phone: +1 678 534 2542
Email: malcolm.betts@zte.com.cn

Nigel Davis
Ciena, UK

Email: ndavis@ciena.com

Bernd Zeuner
Deutsche Telekom, Germany

Phone: +49 6151 58 12086
Email: b.zeuner@telekom.de

Italo Busi
Huawei, China

Email: Italo.Busi@huawei.com

Scott Mansfield
Ericsson, Sweden

Phone: 1 724 931 9316
Email: scott.mansfield@ericsson.com

Yuji Tochio
Fujitsu, Japan

Phone: 81 44 754 8829
Email: tochio@jp.fujitsu.com

Ricard Vilalta
CTTC, Spain

Phone:
Email: ricard.vilalta@cttc.es

Victor Lopez
Telefonica, Spain

Phone:
Email: victor.lopezalvarez@telefonica.com

