

CCAMP Working Group
Internet-Draft
Intended status: Standards Track
Expires: 10 July 2026

A. Guo
Futurewei Technologies
L.M. Contreras
Telefonica
S. Belotti
Nokia
R. Rokui
Ciena
Y. Xu
CAICT
Y. Zhao
China Mobile
X. Liu
Alef Edge
6 January 2026

Framework and Data Model for OTN Network Slicing
draft-ietf-ccamp-yang-otn-slicing-10

Abstract

The requirement of slicing network resources with desired quality of service is emerging at every network technology, including the Optical Transport Networks (OTN). As a part of the transport network, OTN can provide hard pipes with guaranteed data isolation and deterministic low latency, which are highly demanded in the Service Level Agreement (SLA).

This document describes a framework for OTN network slicing and defines YANG data models with OTN technology-specific augments deployed at both the north and south bound of the OTN network slice controller. Additional YANG data model augmentations will be defined in a future version of this draft.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 10 July 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Prefixes in Data Node Names	3
1.3. Definition of OTN Slice	4
2. Use Cases for OTN Network Slicing	5
2.1. Leased Line Services with OTN	6
2.2. Co-construction and Sharing	6
2.3. Wholesale of optical resources	6
2.4. Vertical dedicated network with OTN	7
2.5. End-to-end network slicing	7
3. Framework for OTN slicing	8
4. Realizing OTN Slices	11
5. YANG Data Model for OTN Slicing Configuration	14
5.1. OTN Slicing YANG Model for MPI	14
5.1.1. MPI YANG Model Overview	14
5.1.2. MPI YANG Model Tree	14
5.1.3. MPI YANG Code	15
5.2. OTN Slicing YANG Model for OTN-SC NBI	18
5.2.1. NBI YANG Model Overview	18
5.2.2. NBI YANG Model Tree for OTN slice	19
5.2.3. NBI YANG Code for OTN Slice	21
6. Manageability Considerations	28
7. Security Considerations	29
8. IANA Considerations	29
9. References	30

9.1. Normative References	30
9.2. Informative References	33
Acknowledgments	33
Contributors	33
Authors' Addresses	34

1. Introduction

The requirement of slicing network resources with desired quality of service is emerging at every network technology, including the Optical Transport Networks (OTN). As a part of the transport network, OTN can provide hard pipes with guaranteed data isolation and deterministic low latency, which are highly demanded in the Service Level Agreement (SLA). This document describes a framework for OTN network slicing and defines YANG data models with OTN technology-specific augments deployed at both the north and south bound of the OTN network slice controller. Additional YANG data model augmentations will be defined in a future version of this draft.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The terminology for describing YANG data models is found in [RFC7950].

1.2. Prefixes in Data Node Names

In this document, names of data nodes and other data model objects are prefixed using the standard prefix associated with the corresponding YANG imported modules, as shown in Table 1.

Prefix	YANG Module	Reference
yang	ietf-yang-types	[RFC6991]
inet	ietf-inet-types	[RFC6991]
nt	ietf-network-topology	[RFC8345]
nw	ietf-network-topology	[RFC8345]
tet	ietf-te-topology	[RFC8795]
ietf-nss	ietf-network-slice-service	[RFCVVVV]
ns-topo	ietf-ns-topo	[RFCWWWW]
otnt	ietf-otn-topology	[RFCYYYY]
ll-types	ietf-layer1-types	[RFCZZZZ]
otns	ietf-otn-slice	[RFCXXXX]
otns-mpi	ietf-otn-slice-mpi	[RFCXXXX]

Table 1: Prefixes and Corresponding YANG Modules

RFC Editor Note: Please replace VVVV with the RFC number assigned to [I-D.ietf-teas-ietf-network-slice-nbi-yang]. Please replace WWWW with the RFC number assigned to [I-D.ietf-teas-network-slice-topology-yang]. Please replace XXXX with the RFC number assigned to this document. Please replace YYYY with the RFC number assigned to [I-D.ietf-ccamp-otn-topo-yang]. Please replace ZZZZ with the RFC number assigned to [I-D.ietf-ccamp-layer1-types]. Please remove this note.

1.3. Definition of OTN Slice

An OTN slice is an OTN virtual network topology connecting a number of OTN endpoints using a set of shared or dedicated OTN network resources to satisfy specific service level objectives (SLOs).

An OTN slice is a technology-specific realization of the RFC 9543 network slice service [RFC9543] in the OTN domain, with the capability of configuring slice resources in the term of OTN technologies. Therefore, all the terms and definitions concerning network slicing as defined in [RFC9543] apply to OTN slicing.

An end-to-end OTN slice may be composed of multiple OTN segment slices in a hierarchical or sequential (or stitched) combination.

```
<-----End-to-end OTN Slice----->
<- OTN Segment Slice 1 --->  <-- OTN Segment Slice 2 -->
```

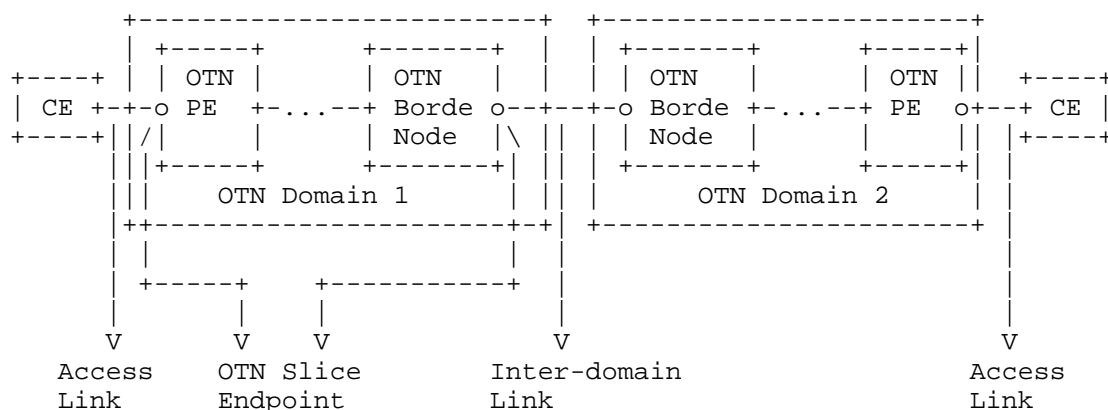


Figure 1: OTN Slice

OTN slices may be pre-configured by the management plane and presented to the customer via the northbound interface (NBI), or be dynamically provisioned by a higher layer slice controller, e.g., an RFC 9543 network slice controller (NSC) through the NBI. The OTN slice is provided by a service provider to a customer to be used as though it was part of the customer's own networks.

2. Use Cases for OTN Network Slicing

2.1. Leased Line Services with OTN

For end business customers (like OTT or enterprises), leased lines have the advantage of providing high-speed connections with low costs. On the other hand, the traffic control of leased lines is very challenging due to rapid changes in service demands. Carriers are recommended to provide network-level slicing capabilities to meet this demand. Based on such capabilities, private network users have full control over the sliced resources which have been allocated to them and which could be used to support their leased lines, when needed. Users may formulate policies based on the demand for services and time to schedule the resources from the entire network's perspective flexibly. For example, the bandwidth between any two points may be established or released based on the time or monitored traffic characteristics. The routing and bandwidth may be adjusted at a specific time interval to maximize network resource utilization efficiency.

2.2. Co-construction and Sharing

Co-construction and sharing of a network are becoming a popular means among service providers to reduce networking building CAPEX. For Co-construction and sharing case, there are typically multiple co-founders for the same network. For example, one founder may provide optical fibres and another founder may provide OTN equipment, while each occupies a certain percentage of the usage rights of the network resources. In this scenario, the network O&M is performed by a certain founder in each region, where the same founder usually deploys an independent management and control system. The other founders of the network use each other's management and control system to provision services remotely. In this scenario, different founders' network resources need to be automatically (associated) divided, isolated, and visualized. All founders may share or have independent O&M capabilities, and should be able to perform service-level provisioning in their respective slices.

2.3. Wholesale of optical resources

In the optical resource wholesale market, smaller, local carriers and wireless carriers may rent resources from larger carriers, or infrastructure carriers instead of building their networks. Likewise, international carriers may rent resources from respective local carriers and local carriers may lease their owned networks to each other to achieve better network utilization efficiency. From the perspective of a resource provider, it is crucial that a network slice is timely configured to meet traffic matrix requirements requested by its tenants. The support for multi-tenancy within the resource provider's network demands that the network slices are

qualitatively isolated from each other to meet the requirements for transparency, non-interference, and security. Typically, a resource purchaser expects to use the leased network resources flexibly, just like they are self-constructed. Therefore, the purchaser is not only provided with a network slice, but also the full set of functionalities for operating and maintaining the network slice. The purchaser also expects to, flexibly and independently, schedule and maintain physical resources to support their own end-to-end automation using both leased and self-constructed network resources.

2.4. Vertical dedicated network with OTN

Vertical industry slicing is an emerging category of network slicing due to the high demand for private high-speed network interconnects for industrial applications. In this scenario, the biggest challenge is to implement differentiated optical network slices based on the requirements from different industries. For example, in the financial industry, to support high-frequency transactions, the slice must ensure to provide the minimum latency along with the mechanism for latency management. For the healthcare industry, online diagnosis network and software capabilities to ensure the delivery of HD video without frame loss. For bulk data migration in data centers, the network needs to support on-demand, large-bandwidth allocation. In each of the aforementioned vertical industry scenarios, the bandwidth shall be adjusted as required to ensure flexible and efficient network resource usage.

2.5. End-to-end network slicing

In an end-to-end network slicing scenario such as 5G network slicing [TS.28.530-3GPP], an RFC 9543 network slice [RFC9543] provides the required connectivity between other different segments of an end-to-end network slice, such as the Radio Access Network (RAN) and the Core Network (CN) segments, with a specific performance commitment. An RFC 9543 network slice could be composed of network slices from multiple technological and administrative domains. An RFC 9543 network slice can be realized by using or combining multiple underlying OTN slices with OTN resources, e.g., ODU time slots or ODU containers, to achieve end-to-end slicing across the transport domain.

3. Framework for OTN slicing

OTN slices may be abstracted differently depending on the requirement contained in the configuration provided by the slice customer. Whereas the customer requests an OTN slice to provide connectivity between specified endpoints, an OTN slice can be abstracted as a set of endpoint-to-endpoint links, with each link formed by an end-to-end tunnel across the underlying OTN networks. The resources associated with each link of the slice is reserved and commissioned in the underlying physical network upon the completion of configuring the OTN slice and all the links are active.

An OTN slice can also be abstracted as an abstract topology when the customer requests the slice to share resources between multiple endpoints and to use the resources on demand. The abstract topology may consist of virtual nodes and virtual links, and their associated resources are reserved but not commissioned across the underlying OTN networks. The customer can later commission resources within the slice dynamically using the NBI provided by the service provider. An OTN slice could use abstract topology to connect endpoints with shared resources to optimize the resource utilization, and connections can be activated within the slice as needed.

It is worth noting that those means to abstract an OTN slice are similar to the Virtual Network (VN) abstraction defined for higher-level interfaces in [RFC8453], in which context a connectivity-based slice corresponds to Type 1 VN and a resource-based slice corresponds to Type 2 VN, respectively.

A particular resource in an OTN network, such as a port or link, may be sliced with one of the two granularity levels:

- * Link-based slicing, in which a link and its associated link termination points (LTPs) are dedicatedly allocated to a particular OTN slice.
- * Tributary-slot based slicing, in which multiple OTN slices share the same link by allocating different OTN tributary slots in different granularities.

Furthermore, an OTN switch is typically fully non-blocking switching at the lowest ODU container granularity, it is desirable to specify just the total number of ODU containers in the lowest granularity (e.g. ODU0), when configuring tributary-slot based slicing on links and ports internal to an OTN network. In multi-domain OTN network scenarios where separate OTN slices are created on each of the OTN networks and are stitched at inter-domain OTN links, it is necessary to specify matching tributary slots at the endpoints of the inter-

domain links. In some real network scenarios, OTN network resources including tributary slots are managed explicitly by network operators for network maintenance considerations. Therefore, an OTN slice controller shall support configuring an OTN slice with both options.

An OTN slice controller (OTN-SC) is a logical function responsible for the life-cycle management of OTN slices instantiated within the corresponding OTN network domains. The OTN-SC provides technology-specific interfaces at its northbound (OTN-SC NBI) to allow a higher-layer slice controller, such as an RFC 9543 network slice controller (NSC) or an orchestrator, to request OTN slices with OTN-specific requirements. The OTN-SC interfaces at the southbound using the MDSC-to-PNC interface (MPI) with a Physical Network Controller (PNC) or Multi-Domain Service Orchestrator (MDSC), as defined in the ACTN control framework [RFC8453]. The logical function within the OTN-SC is responsible for translating the OTN slice requests into concrete slice realization which can be understood and provisioned at the southbound by the PNC or MDSC.

The presence of OTN-SC provides multiple options for a high-level slice controller or an orchestrator to configure and realize slicing in OTN networks, depending on whether a customer's slice request is technology agnostic or technology specific:

Option 1[opt.1]: An IETF NSC receives a technology-agnostic slice request from the IETF NSC NBI and realizes full or part of the slice in OTN networks directly through MPI provided by the PNC or MDSC. The IETF NSC is responsible for mapping a technology-agnostic slicing request into an OTN technology-specific realization. In this option, the OTN-SC is not used.

Option 2[opt.2]: An IETF NSC receives a technology-agnostic slice request from the IETF NSC NBI and delegates the request to the OTN-SC through the OTN-SC NBI, which is OTN technology specific. The OTN-SC in turn realizes the slice in single or multi domain OTN networks by working with the underlying PNC or MDSC. In this option, the OTN-SC is considered as a realization of IETF NSC, i.e., an NS realizer as per [I-D.ietf-teas-ns-controller-models], when the underlying network is OTN. The OTN-SC is also a subordinate slice controller of the RFC 9543 NSC, which is consistent with the hierarchical control of slices specified by [RFC9543].

Option 3[opt.3]: An OTN-aware orchestrator may request an OTN technology-specific slice with OTN-specific SLOs through the OTN-SC NBI to the OTN-SC. The OTN-SC in turn realizes the slice in single or multi domain OTN networks by working with the underlying PNC or MDSC

An OTN slice may be realized by using standard MPI interfaces, control plane, network management system (NMS) or any other proprietary interfaces as needed. Examples of such interfaces include the abstract TE topology [RFC8795], TE tunnel [I-D.ietf-teas-yang-te], L1VPN [RFC4847], or Netconf/YANG based interfaces such as OpenConfig. Some of these interfaces, such as the TE tunnel model, are suitable for creating connectivity-based OTN slices which represent a slice as a set of TE tunnels, while other interfaces such as the TE topology model are more suitable for creating resource-based OTN slices which represent a slice as a topology.

The OTN-SC NBI is a technology-specific interface that augments the IETF NSC NBI, which is technology-agnostic.

Figure 2 illustrates the OTN slicing control hierarchy, the positioning of the OTN slicing interfaces as well as the options for OTN slice configuration.

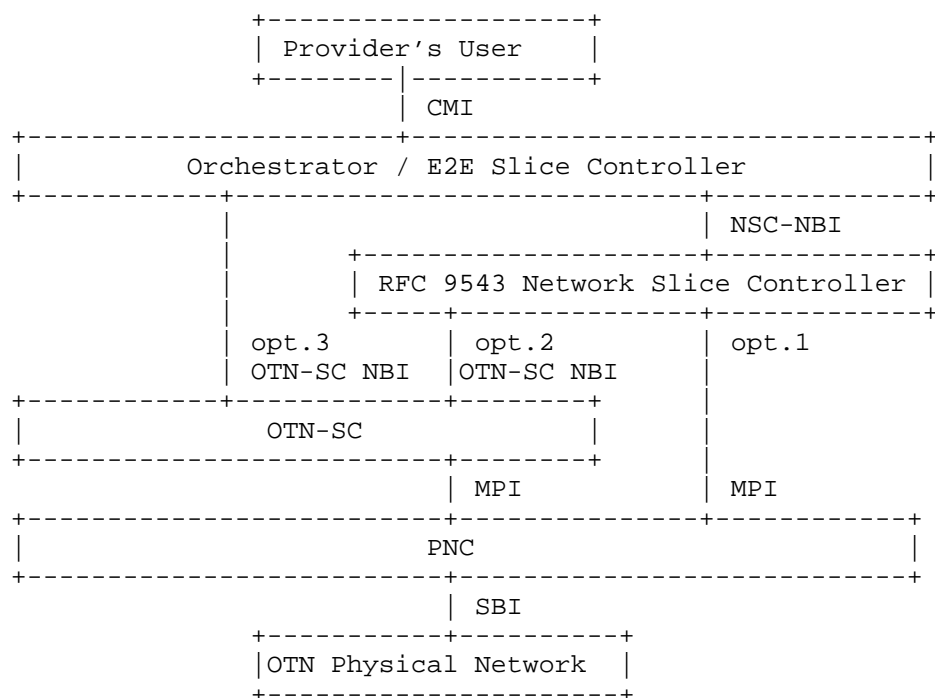


Figure 2: Positioning of OTN Slicing Interfaces

OTN-SC functionalities may be recursive such that a higher-level OTN-SC may designate the creation of OTN slices to a lower-level OTN-SC in a recursive manner. This scenario may apply to the creation of OTN slices in multi-domain OTN networks, where multiple domain-wide OTN slices provisioned by lower-layer OTN-SCs are stitched to support a multi-domain OTN slice provisioned by the higher-level OTN-SC. Alternatively, the OTN-SC may interface with an MDSC, which in turn interfaces with multiple PNCs through the MPI to realize OTN slices in multi-domain OTN networks without OTN-SC recursion. Figure 3 illustrates both options for OTN slicing in multi-domain.

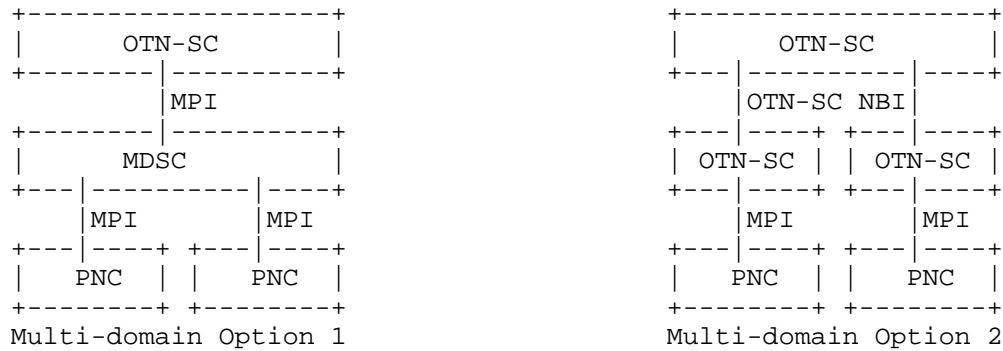


Figure 3: OTN-SC for multi-domain

OTN-SC functionalities are logically independent and may be deployed in different combinations to cater to the realization needs. In reference to the ACTN control framework [RFC8453], an OTN-SC may be deployed

- * as an independent network function;
- * together with a Physical Network Controller (PNC) for single-domain or with a Multi-Domain Service Orchestrator (MDSC) for multi-domain;
- * together with a higher-level network slice controller to support end-to-end network slicing;

4. Realizing OTN Slices

[RFC9543] introduces a mechanism for an RFC 9543 network slice controller to realize network slices by constructing Network Resource Partitions (NRP). A NRP is a collection of resources identified in the underlay network to facilitate the mapping of network slices onto available network resources. An NRP is a scope view of a topology and may be considered as a topology in its own right. Thus, in

traffic-engineered (TE) networks including OTN, an NRP may be simply represented as an abstract TE topology defined by [RFC8795]. For OTN networks, An NRP may be represented as an abstract OTN topology defined by [I-D.ietf-ccamp-otn-topo-yang].

The NRP may be used to address the scalability issues where there may be considerable numbers of control and data plane states required to be stored and programmed if network slices are mapped directly to the underlay topology. NRP is internal to a network slice controller, and use of NRPs is optional yet could benefit a network slice realization in large-scale networks, including OTN networks.

For connectivity-based OTN slices, a connection within an OTN slice is typically realized by an OTN tunnel in the underlay topology and resources are reserved by the tunnel, thus use of NRP is optional in this case.

For resource-based OTN slices, the OTN-SC may map an OTN slice directly onto the underlay TE topology presented by the subtended network controller (MDSC or PNC) without creating NRP topologies. Due to the need for reserving resources, the OTN-SC needs to color corresponding link resources of the underlay topology with a slice identifier and maintain the coloring to keep track of the mapping of OTN slices. The OTN-SC may push the colored topology to the subtended MDSC or PNC using the MPI model defined in this draft.

Alternatively, an OTN slice may be mapped to a NRP as an overlay abstract OTN TE topology on top of the underlay topology. The corresponding link resources allocated to the slice is encapsulated in and tracked by the abstract topology, and a given link or port in the NRP topology represents resources that are reserved in the underlay topology. Multiple OTN slices may be mapped to the same NRP, and a single connectivity construct of the slice may be mapped to only one NRP, as per [RFC9543]. The resources of an NRP topology are reserved and shared by all the OTN slices mapped to this NRP, and the NRP topology may be pushed directly to the subtended MDSC or PNC, thus eliminating the need for link coloring if using the underlay topology.

Figure 4 illustrates the relationship between OTN slices and NRP.

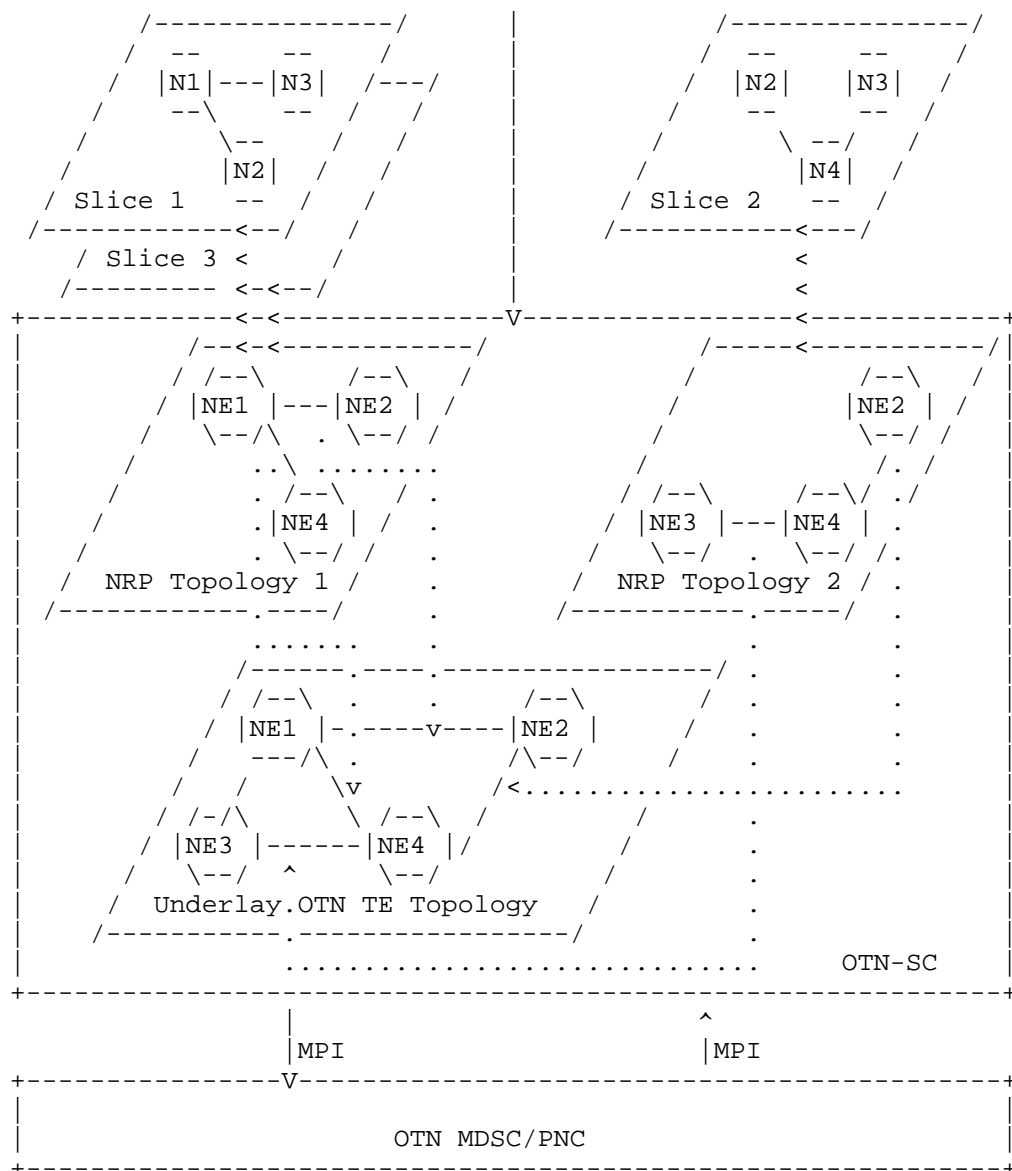


Figure 4: Mapping OTN Slices to NRP

5. YANG Data Model for OTN Slicing Configuration

5.1. OTN Slicing YANG Model for MPI

5.1.1. MPI YANG Model Overview

For the configuration of connectivity-based OTN slices, existing models such as the TE tunnel interface [I-D.ietf-teas-yang-te] may be used and no addition is needed. This model is addressing the case for configuring resource-based OTN slices, where the model permits to reserve resources exploiting the common knowledge of an underlying virtual topology between the OTN-SC and the subtended network controller (MDSC or PNC). The slice is configured by marking corresponding link resources on the TE topology received from the underlying MDSC or PNC with a slice identifier and OTN-specific resource requirements, e.g. the number of ODU time slots or the type/number of ODU containers. The MDSC or PNC, based on the marked resources by the OTN-SC, will update the underlying TE topology with new TE link for each of the colored links to keep booked the reserved OTN resources e.g. time slots or ODU containers.

5.1.2. MPI YANG Model Tree

```

module: ietf-otn-slice-mpi

augment /nw:networks/nw:network/nt:link/tet:te
  /tet:te-link-attributes:
    +--rw (otn-slice-granularity)?
      +--:(link)
        | +--rw slice-id?      uint32
        +--:(link-resource)
          +--rw slices* [slice-id]
            +--rw slice-id          uint32
            +--rw (technology)?
              +--:(otn)
                +--rw (slice-bandwidth)?
                  +--:(containers)
                    +--rw odulist* [odu-type]
                      +--rw odu-type      identityref
                      +--rw number?       uint16
                  +--:(time-slots)
                    +--rw otn-ts-num?    uint32
            +--ro sliced-link-ref?
              -> ../../../../../../nt:link/link-id

```

Figure 5: OTN slicing MPI tree diagram

5.1.3. MPI YANG Code

```
<CODE BEGINS> file "ietf-otn-slice-mpi@2025-07-03.yang"
module ietf-otn-slice-mpi {
  yang-version 1.1;
  namespace "urn:ietf:params:xml:ns:yang:ietf-otn-slice-mpi";
  prefix "otns-mpi";

  import ietf-network {
    prefix "nw";
    reference
      "RFC 8345: A YANG Data Model for Network Topologies";
  }

  import ietf-network-topology {
    prefix "nt";
    reference
      "RFC 8345: A YANG Data Model for Network Topologies";
  }

  import ietf-te-topology {
    prefix "tet";
    reference
      "RFC8795: YANG Data Model for Traffic Engineering
      (TE) Topologies";
  }

  import ietf-otn-topology {
    prefix "otnt";
    reference
      "draft-ietf-ccamp-otn-topo-yang-20:
      RFC YYYY: A YANG Data Model for Optical Transport
      Network Topology";
  }

  import ietf-layer1-types {
    prefix "l1-types";
    reference
      "draft-ietf-ccamp-layer1-types-18:
      RFC ZZZZ: A YANG Data Model for Layer 1 Types";
  }

  organization
    "IETF CCAMP Working Group";
  contact
    "WG Web: <http://tools.ietf.org/wg/ccamp/>
    WG List: <mailto:ccamp@ietf.org>
```

Editor: Haomian Zheng
<mailto:zhenghaomian@huawei.com>

Editor: Italo Busi
<mailto:italo.busi@huawei.com>

Editor: Aihua Guo
<mailto:aihuaguo.ietf@gmail.com>

Editor: Sergio Belotti
<mailto:sergio.belotti@nokia.com>";

description

"This module defines a YANG data model for network slice realization in Optical Transport Networks (OTN).

The model fully conforms to the Network Management Datastore Architecture (NMDA).

Copyright (c) 2022 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Revised BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

```
revision "2025-07-03" {
  description
    "Latest revision of MPI YANG model for OTN slicing.";
  reference
    "draft-ietf-ccamp-yang-otn-slicing-09: Framework and Data
     Model for OTN Network Slicing";
}
```

```
/*
 * Groupings
 */
```

```
grouping otn-link-slice-profile {
  description
    "Profile of an OTN link slice.";
  choice otn-slice-granularity {
```



```
default "link";
description
  "Link slice granularity.";
case link {
  leaf slice-id {
    type uint32;
    description
      "Slice identifier";
  }
}
case link-resource {
  list slices {
    key slice-id;
    description
      "List of slices.";
    leaf slice-id {
      type uint32;
      description
        "Slice identifier";
    }
  }
  choice technology {
    description
      "Data plane technology types.";
    case otn {
      choice slice-bandwidth {
        description
          "Bandwidth specification for OTN slices.";
        case containers {
          uses ll-types:otn-link-bandwidth;
        }
        case time-slots {
          leaf otn-ts-num {
            type uint32;
            description
              "Number of OTN tributary slots allocated
               for the slice.";
          }
        }
      }
    }
  }
  leaf sliced-link-ref {
    type leafref {
      path "../.../.../.../nt:link/nt:link-id";
    }
    config false;
    description
      "Relative reference to virtual links generated from
```

```

        this TE link.";
    }
}
}
}
}
/*
 * Augments
 */
augment "/nw:networks/nw:network/nt:link/tet:te/"
    + "tet:te-link-attributes" {
    when "../../../nw:network-types/tet:te-topology/"
        + "otnt:otn-topology" {
        description
            "Augmentation parameters apply only for networks with
            OTN topology type.";
    }
    description
        "Augment OTN TE link attributes with slicing profile.";
    uses otn-link-slice-profile;
}
}
<CODE ENDS>

```

Figure 6: OTN slicing MPI YANG model

5.2. OTN Slicing YANG Model for OTN-SC NBI

5.2.1. NBI YANG Model Overview

The YANG model for OTN-SC NBI is OTN-technology specific, but shares many common constructs and attributes with the common network slicing YANG model defined in [I-D.ietf-teas-ietf-network-slice-nbi-yang]. Furthermore, the OTN-SC NBI YANG is expected to support both connectivity-based and resource-based slice configuration, which is likely a common requirement for supporting slicing at other transport network layers, e.g. WDM or MPLS(-TP).

The OTN slicing model augments the common network slicing YANG model by extending OTN technology-specific SLO and SLE attributes which can be requested by OTN-aware customers and allows the customer to specify desired OTN signal quality. These attributes include:

- * The performance objective for Optical Data Unit (ODU) containers as defined in [ITU-T-G.8201-Amd.1].
- * Bandwidth specification in the type and number of ODU containers.

5.2.2. NBI YANG Model Tree for OTN slice

```

module: ietf-otn-slice

augment /ietf-nss:network-slice-services/ietf-nss:slo-sle-templates
  /ietf-nss:slo-sle-template/ietf-nss:slo-policy:
  +--rw otn
    +--rw odu-signal-quality
      |   +--rw odu-pm-objective* [duration pm-type]
      |   |   +--rw duration          identityref
      |   |   +--rw pm-type           identityref
      |   |   +--rw pm-threshold?    uint64
      +--rw odulist* [odu-type]
        +--rw odu-type          identityref
        +--rw number?          uint16
augment /ietf-nss:network-slice-services/ietf-nss:slice-service
  /ietf-nss:slo-sle-policy/ietf-nss:custom
  /ietf-nss:service-slo-sle-policy/ietf-nss:slo-policy:
  +--rw otn
    +--rw odu-signal-quality
      |   +--rw odu-pm-objective* [duration pm-type]
      |   |   +--rw duration          identityref
      |   |   +--rw pm-type           identityref
      |   |   +--rw pm-threshold?    uint64
      +--rw odulist* [odu-type]
        +--rw odu-type          identityref
        +--rw number?          uint16
augment /nw:networks/nw:network/ns-topo:slo-sle-policy
  /ns-topo:custom/ns-topo:service-slo-sle-policy
  /ns-topo:slo-policy:
  +--rw otn
    +--rw odu-signal-quality
      |   +--rw odu-pm-objective* [duration pm-type]
      |   |   +--rw duration          identityref
      |   |   +--rw pm-type           identityref
      |   |   +--rw pm-threshold?    uint64
      +--rw odulist* [odu-type]
        +--rw odu-type          identityref
        +--rw number?          uint16
augment /nw:networks/nw:network/nw:node/ns-topo:slo-sle-policy
  /ns-topo:custom/ns-topo:service-slo-sle-policy
  /ns-topo:slo-policy:
  +--rw otn
    +--rw odu-signal-quality
      |   +--rw odu-pm-objective* [duration pm-type]
      |   |   +--rw duration          identityref
      |   |   +--rw pm-type           identityref
      |   |   +--rw pm-threshold?    uint64

```

```

    +--rw odulist* [odu-type]
      +--rw odu-type    identityref
      +--rw number?     uint16
augment /nw:networks/nw:network/nw:node/nt:termination-point
      /ns-topo:slo-sle-policy/ns-topo:custom
      /ns-topo:service-slo-sle-policy/ns-topo:slo-policy:
+--rw otn
  +--rw odu-signal-quality
  |   +--rw odu-pm-objective* [duration pm-type]
  |   |   +--rw duration      identityref
  |   |   +--rw pm-type       identityref
  |   |   +--rw pm-threshold? uint64
  |   +--rw odulist* [odu-type]
  |   |   +--rw odu-type      identityref
  |   |   +--rw number?       uint16
augment /nw:networks/nw:network/nt:link/ns-topo:slo-sle-policy
      /ns-topo:custom/ns-topo:service-slo-sle-policy
      /ns-topo:slo-policy:
+--rw otn
  +--rw odu-signal-quality
  |   +--rw odu-pm-objective* [duration pm-type]
  |   |   +--rw duration      identityref
  |   |   +--rw pm-type       identityref
  |   |   +--rw pm-threshold? uint64
  |   +--rw odulist* [odu-type]
  |   |   +--rw odu-type      identityref
  |   |   +--rw number?       uint16
augment /ietf-nss:network-slice-services/ietf-nss:slice-service
      /ietf-nss:connection-groups/ietf-nss:connection-group
      /ietf-nss:slo-sle-policy/ietf-nss:custom
      /ietf-nss:service-slo-sle-policy/ietf-nss:slo-policy:
+--rw otn
  +--rw odu-signal-quality
  |   +--rw odu-pm-objective* [duration pm-type]
  |   |   +--rw duration      identityref
  |   |   +--rw pm-type       identityref
  |   |   +--rw pm-threshold? uint64
  |   +--rw odulist* [odu-type]
  |   |   +--rw odu-type      identityref
  |   |   +--rw number?       uint16
augment /ietf-nss:network-slice-services/ietf-nss:slice-service
      /ietf-nss:connection-groups/ietf-nss:connection-group
      /ietf-nss:connectivity-construct/ietf-nss:slo-sle-policy
      /ietf-nss:custom/ietf-nss:service-slo-sle-policy
      /ietf-nss:slo-policy:
+--rw otn
  +--rw odu-signal-quality
  |   +--rw odu-pm-objective* [duration pm-type]

```

```

|      +--rw duration          identityref
|      +--rw pm-type           identityref
|      +--rw pm-threshold?     uint64
+--rw odulist* [odu-type]
    +--rw odu-type             identityref
    +--rw number?              uint16
augment /ietf-nss:network-slice-services/ietf-nss:slice-service
  /ietf-nss:connection-groups/ietf-nss:connection-group
  /ietf-nss:connectivity-construct/ietf-nss:type
  /ietf-nss:a2a/ietf-nss:a2a-sdp/ietf-nss:slo-sle-policy
  /ietf-nss:custom/ietf-nss:service-slo-sle-policy
  /ietf-nss:slo-policy:
+--rw otn
  +--rw odu-signal-quality
  |   +--rw odu-pm-objective* [duration pm-type]
  |   |   +--rw duration          identityref
  |   |   +--rw pm-type           identityref
  |   |   +--rw pm-threshold?     uint64
  +--rw odulist* [odu-type]
    +--rw odu-type             identityref
    +--rw number?              uint16

```

Figure 7: Tree diagram for OTN slice

5.2.3. NBI YANG Code for OTN Slice

```

<CODE BEGINS> file "ietf-otn-slice@2025-07-03.yang"
module ietf-otn-slice {
  yang-version 1.1;
  namespace
    "urn:ietf:params:xml:ns:yang:ietf-otn-slice";
  prefix "otns";

  import ietf-network {
    prefix "nw";
    reference
      "RFC 8345: A YANG Data Model for Network Topologies";
  }
  import ietf-network-topology {
    prefix "nt";
    reference
      "RFC 8345: A YANG Data Model for Network Topologies";
  }

  import ietf-layer1-types {
    prefix "l1-types";
    reference
      "draft-ietf-ccamp-layer1-types-18:

```

```
    RFC ZZZZ: A YANG Data Model for Layer 1 Types";
}

import ietf-network-slice-service {
    prefix "ietf-nss";
    reference
        "draft-ietf-teas-ietf-network-slice-nbi-yang-25:
        RFC VVVV: A YANG Data Model for the RFC 9543 Network Slice
        Service";
}

import ietf-ns-topo {
    prefix "ns-topo";
    reference
        "draft-ietf-teas-network-slice-topology-yang-01:
        RFC WWWW: IETF Network Slice Topology YANG Data Model";
}

organization
    "IETF CCAMP Working Group";
contact
    "WG Web: <http://tools.ietf.org/wg/ccamp/>
    WG List: <mailto:ccamp@ietf.org>

    Editor: Haomian Zheng
            <mailto:zhenghaomian@huawei.com>

    Editor: Italo Busi
            <mailto:italo.busi@huawei.com>

    Editor: Aihua Guo
            <mailto:aihuaguo.ietf@gmail.com>

    Editor: Sergio Belotti
            <mailto:sergio.belotti@nokia.com>";

description
    "This module defines a YANG data model for configuring
    technology-specific network slices in optical transport
    networks, e.g., Optical Transport Network (OTN).

    The model fully conforms to the Network Management Datastore
    Architecture (NMDA).

    Copyright (c) 2022 IETF Trust and the persons identified as
    authors of the code. All rights reserved.

    Redistribution and use in source and binary forms, with or
```

without modification, is permitted pursuant to, and subject to the license terms contained in, the Revised BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>).

This version of this YANG module is part of RFC XXXX; see the RFC itself for full legal notices.";

```
revision "2025-07-03" {
  description
    "Latest revision of NBI YANG model for OTN slicing.";
  reference
    "draft-ietf-ccamp-yang-otn-slicing-09: Framework and Data
     Model for OTN Network Slicing";
}

/*
 * Identities
 */
identity bit-error-rate {
  base ietf-nss:service-slo-metric-type;
  description
    "ODU bit error rate";
}

identity odu-tca-threshold-type {
  description
    "Base identity for ODU performance counter";
}

identity odu-bbe {
  base odu-tca-threshold-type;
  description
    "ODU Background Block Error (BBE) threshold";
}

identity odu-es {
  base odu-tca-threshold-type;
  description
    "ODU Errored Seconds (ES) threshold";
}

identity odu-ses {
  base odu-tca-threshold-type;
  description
    "ODU Severely Errored Seconds (SES) threshold";
}
```

```
}

identity odu-uas {
  base odu-tca-threshold-type;
  description
    "ODU Unavailable Seconds (UAS) threshold";
}

identity odu-ber {
  base odu-tca-threshold-type;
  description
    "ODU Bit Error Rate (BER) threshold";
}

identity pm-duration {
  description
    "Base identity for ODU performance monitoring interval";
}

identity pm-15m {
  base pm-duration;
  description
    "15 minutes pm duration";
}

identity pm-24h {
  base pm-duration;
  description
    "24 hours pm duration";
}

/*
 * Groupings
 */
grouping odu-signal-quality {
  description
    "Grouping for ODU signal quality.";

  container odu-signal-quality {
    description
      "Container for ODU signal quality attributes.";

    list odu-pm-objective {
      key "duration pm-type";
      description
        "List of ODU performance requirements.";
      leaf duration {
        type identityref {
```



```
        base pm-duration;
    }
    description
        "Time duration.";
}
leaf pm-type {
    type identityref {
        base odu-tca-threshold-type;
    }
    description
        "ODU PM metric type.";
}

leaf pm-threshold {
    type uint64;
    description
        "ODU PM metric threshold.";
}
}
}
}

grouping otn-slice-slo-policy {
    description
        "Policy grouping for OTN network slices.";

    container otn {
        description
            "OTN technology-specific SLO/SLE policy container";

        uses odu-signal-quality;
        uses ll-types:otn-link-bandwidth;
    }
}

/*
 * Augmented data nodes
 */
/* slice template */
augment "/ietf-nss:network-slice-services" +
    "/ietf-nss:slo-sle-templates" +
    "/ietf-nss:slo-sle-template" +
    "/ietf-nss:slo-policy" {
    description
        "Augment IETF network slice service templates with
        OTN technology-specific SLO/SLE policy attributes.";

    uses otn-slice-slo-policy;
}
```

```
}

/* slice augments */
augment "/ietf-nss:network-slice-services" +
    "/ietf-nss:slice-service" +
    "/ietf-nss:slo-sle-policy" +
    "/ietf-nss:custom" +
    "/ietf-nss:service-slo-sle-policy" +
    "/ietf-nss:slo-policy" {
    description
        "Augment IETF network slice services to include technology-
        specific SLO/SLE policy for connectivity-based OTN
        slices.";

    uses otn-slice-slo-policy;
}

/* network topology augments */
augment "/nw:networks/nw:network" +
    "/ns-topo:slo-sle-policy" +
    "/ns-topo:custom" +
    "/ns-topo:service-slo-sle-policy" +
    "/ns-topo:slo-policy" {
    when "../../../nw:network-types/ns-topo:network-slice" {
        description "Augment only for Network Slice topology.";
    }
    description "Augment topology configuration and state.";
    uses otn-slice-slo-policy;
}

/* network node augments */
augment "/nw:networks/nw:network/nw:node" +
    "/ns-topo:slo-sle-policy" +
    "/ns-topo:custom" +
    "/ns-topo:service-slo-sle-policy" +
    "/ns-topo:slo-policy" {
    when "../../../nw:network-types/ns-topo:network-slice" {
        description "Augment only for Network Slice topology.";
    }
    description "Augment node configuration and state.";
    uses otn-slice-slo-policy;
}

/* network node's termination point augments */
augment "/nw:networks/nw:network/nw:node" +
    "/nt:termination-point" +
    "/ns-topo:slo-sle-policy" +
    "/ns-topo:custom" +
```

```
        "/ns-topo:service-slo-sle-policy" +
        "/ns-topo:slo-policy" {
when "../../../nw:network-types/ns-topo:network-slice" {
    description "Augment only for Network Slice topology.";
}
description "Augment node configuration and state.";

    uses otn-slice-slo-policy;
}

/* network link augments */
augment "/nw:networks/nw:network/nt:link" +
    "/ns-topo:slo-sle-policy" +
    "/ns-topo:custom" +
    "/ns-topo:service-slo-sle-policy" +
    "/ns-topo:slo-policy" {
when "../../../nw:network-types/ns-topo:network-slice" {
    description "Augment only for Network Slice topology.";
}
description "Augment link configuration and state.";
    uses otn-slice-slo-policy;
}

/* connectivity construct augments */
augment "/ietf-nss:network-slice-services" +
    "/ietf-nss:slice-service" +
    "/ietf-nss:connection-groups" +
    "/ietf-nss:connection-group" +
    "/ietf-nss:slo-sle-policy" +
    "/ietf-nss:custom" +
    "/ietf-nss:service-slo-sle-policy" +
    "/ietf-nss:slo-policy" {
    description
        "Augment IETF network slice services to include technology-
        specific SLO/SLE policy for connection groups within
        a connectivity-based transport network slice.";
    uses otn-slice-slo-policy;
}

augment "/ietf-nss:network-slice-services" +
    "/ietf-nss:slice-service" +
    "/ietf-nss:connection-groups" +
    "/ietf-nss:connection-group" +
    "/ietf-nss:connectivity-construct" +
    "/ietf-nss:slo-sle-policy" +
    "/ietf-nss:custom" +
    "/ietf-nss:service-slo-sle-policy" +
    "/ietf-nss:slo-policy" {
```

```

    description
      "Augment IETF network slice services to include technology-
       specific SLO/SLE policy for connectivity-constructs within
       a connectivity-based transport network slice.";
    uses otn-slice-slo-policy;
  }

  augment "/ietf-nss:network-slice-services" +
    "/ietf-nss:slice-service" +
    "/ietf-nss:connection-groups" +
    "/ietf-nss:connection-group" +
    "/ietf-nss:connectivity-construct" +
    "/ietf-nss:type" +
    "/ietf-nss:a2a" +
    "/ietf-nss:a2a-sdp" +
    "/ietf-nss:slo-sle-policy" +
    "/ietf-nss:custom" +
    "/ietf-nss:service-slo-sle-policy" +
    "/ietf-nss:slo-policy" {
    description
      "Augment IETF network slice services to include technology-
       specific SLO/SLE policy for a2a connectivity-constructs
       within a connectivity-based transport network slice.";
    uses otn-slice-slo-policy;
  }
}
<CODE ENDS>

```

Figure 8: YANG model for transport network slice

6. Manageability Considerations

To ensure the security and controllability of physical resource isolation, slice-based independent operation and management are required to achieve management isolation. Each optical slice typically requires dedicated accounts, permissions, and resources for independent access and O&M. This mechanism is to guarantee the information isolation among slice tenants and to avoid resource conflicts. The access to slice management functions will only be permitted after successful security checks.

7. Security Considerations

The YANG module specified in this document defines a schema for data that is designed to be accessed via network management protocols such as NETCONF [RFC6241] or RESTCONF [RFC8040]. The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

The NETCONF access control model [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a preconfigured subset of all available NETCONF or RESTCONF protocol operations and content.

There are a number of data nodes defined in this YANG module that are writable/creatable/deletable (i.e., config true, which is the default). These data nodes may be considered sensitive or vulnerable in some network environments. Write operations (e.g., edit-config) to these data nodes without proper protection can have a negative effect on network operations. Considerations in Section 8 of [RFC8795] are also applicable to their subtrees in the module defined in this document.

Some of the readable data nodes in this YANG module may be considered sensitive or vulnerable in some network environments. It is thus important to control read access (e.g., via get, get-config, or notification) to these data nodes. Considerations in Section 8 of [RFC8795] are also applicable to their subtrees in the module defined in this document.

8. IANA Considerations

It is proposed to IANA to assign new URIs from the "IETF XML Registry" [RFC3688] as follows:

URI: urn:ietf:params:xml:ns:yang:ietf-otn-slice
Registrant Contact: The IESG
XML: N/A; the requested URI is an XML namespace.

URI: urn:ietf:params:xml:ns:yang:ietf-otn-slice-mpi
Registrant Contact: The IESG
XML: N/A; the requested URI is an XML namespace.

This document registers two YANG modules in the YANG Module Names registry [RFC6020].

name: ietf-otn-slice
namespace: urn:ietf:params:xml:ns:yang:ietf-otn-slice
prefix: otns
reference: RFC XXXX

name: ietf-otn-slice-mpi
namespace: urn:ietf:params:xml:ns:yang:ietf-otn-slice-mpi
prefix: otns-mpi
reference: RFC XXXX

9. References

9.1. Normative References

[I-D.ietf-ccamp-layer1-types]

Zheng, H. and I. Busi, "Common YANG Data Types for Layer 1 Networks", Work in Progress, Internet-Draft, draft-ietf-ccamp-layer1-types-18, 23 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-layer1-types-18>>.

[I-D.ietf-ccamp-otn-topo-yang]

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

[I-D.ietf-teas-ietf-network-slice-nbi-yang]

Wu, B., Dhody, D., Rokui, R., Saad, T., and J. Mullooly, "A YANG Data Model for the RFC 9543 Network Slice Service", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-nbi-yang-25, 9 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slice-nbi-yang-25>>.

[I-D.ietf-teas-network-slice-topology-yang]

Liu, X., Contreras, L. M., Belotti, S., Guo, A., and I. Busi, "IETF Network Slice Topology YANG Data Model", Work in Progress, Internet-Draft, draft-ietf-teas-network-slice-topology-yang-02, 21 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-network-slice-topology-yang-02>>.

[I-D.ietf-teas-ns-controller-models]

Contreras, L. M., Rokui, R., Tantsura, J., Wu, B., and X. Liu, "IETF Network Slice Controller and its Associated

Data Models", Work in Progress, Internet-Draft, draft-ietf-teas-ns-controller-models-06, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-controller-models-06>>.

[I-D.ietf-teas-yang-te]

Saad, T., Gandhi, R., Liu, X., Beeram, V. P., and I. Bryskin, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths, and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-40, 2 November 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-40>>.

[ITU-T-G.8201-Amd.1]

ITU Telecommunication Standardization Sector (ITU-T), "Error performance parameters and objectives for multi-operator international paths within optical transport networks (Amendment 1)", ITU-T G.8201 (2011) Amd. 1, December 2021, <<https://www.itu.int/rec/T-REC-G.8201>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, DOI 10.17487/RFC3688, January 2004, <<https://www.rfc-editor.org/info/rfc3688>>.

[RFC4847] Takeda, T., Ed., "Framework and Requirements for Layer 1 Virtual Private Networks", RFC 4847, DOI 10.17487/RFC4847, April 2007, <<https://www.rfc-editor.org/info/rfc4847>>.

[RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.

[RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/info/rfc6242>>.

- [RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration Access Control Model", STD 91, RFC 8341, DOI 10.17487/RFC8341, March 2018, <<https://www.rfc-editor.org/info/rfc8341>>.
- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", RFC 8345, DOI 10.17487/RFC8345, March 2018, <<https://www.rfc-editor.org/info/rfc8345>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.
- [RFC8795] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Gonzalez de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", RFC 8795, DOI 10.17487/RFC8795, August 2020, <<https://www.rfc-editor.org/info/rfc8795>>.
- [TS.28.530-3GPP] 3rd Generation Partnership Project (3GPP), "3GPP TS 28.530 V15.1.0 Technical Specification Group Services and System Aspects; Management and orchestration; Concepts, use cases and requirements (Release 15)", 3GPP TS 28.530 , December 2018, <http://ftp.3gpp.org/Specs/archive/28_series/28.530/28530-f10.zip>.

9.2. Informative References

[RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/info/rfc9543>>.

Acknowledgments

This document was prepared using kramdown.

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

The authors would like to thank Adrian Farrel, Danielle Ceccarelli, Igor Bryskin, Bo Wu, Gyan Mishra, Joel M. Halpen, Dhruv Dhoddy and Loa Andersson for providing valuable insights.

Contributors

Haomian Zheng
Huawei Technologies
H1, Xiliu Beipo Village, Songshan Lake
Dongguan
China
Email: zhenghaomian@huawei.com

Italo Busi
Huawei Technologies
Email: italo.busi@huawei.com

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

Victor Lopez
Nokia
Email: victor.lopez@nokia.com

Dieter Beller
Nokia
Email: Dieter.Beller@nokia.com

Henry Yu
Huawei Technologies Canada
Email: henry.yu@huawei.com

Jiang Sun
China Mobile
Email: sunjiang@chinamobile.com

Authors' Addresses

Aihua Guo
Futurewei Technologies
Email: aihuaguo.ietf@gmail.com

Luis M. Contreras
Telefonica
Email: luismiguel.contrerasmurillo@telefonica.com

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

Reza Rokui
Ciena
Email: rrokui@ciena.com

Yunbin Xu
CAICT
Email: xuyunbin@caict.ca.cn

Yang Zhao
China Mobile
Email: zhaoyangyjy@chinamobile.com

Xufeng Liu
Alef Edge
Email: xufeng.liu.ietf@gmail.com