OPSAWG
Internet-Draft
Intended status: Standards Track
Expires: 14 September 2023

S. B. Giraldo, Ed.
Nokia
L. M. Contreras, Ed.
Telefonica
V. Lopez
Nokia
R. Rokui
Ciena
O. G. D. Dios
Telefonica
D. King
Old Dog Consulting
13 March 2023

Instantiation of IETF Network Slices in Service Providers Network
draft-barguil-teas-network-slices-instantation-06

Abstract

Network Slicing (NS) is an integral part of Service Provider networks.

The IETF has produced several YANG data models to support the Software-Defined Networking and network slice architecture and YANG-based service models for network slice (NS) instantiation.

This document describes the relationship between IETF Network Slice models for requesting the IETF Network Slices and (e.g., Layer-3 Service Model, Layer-2 Service Model) and Network Models (e.g., Layer-3 Network Model, Layer-2 Network Model) used during their realizations.

In addition, this document describes the communication between the IETF Network Slice Controller and the network controllers for the realization of IETF network slices.

The IETF Network Slice YANG model provides the customer-oriented view of the network slice. Thus, once the IETF Network Slice controller (NSC) receives a request, it needs to map it to accomplish the specific parameters expected by the network controllers. The network models are analyzed to satisfy the IETF Network Slice requirements, and the gaps in existing models are reported.

The document also provides operational and security considerations when deploying network slices in Service Provider networks.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

1.1. Scope and Intended Use

The IETF has produced several YANG data models to support the Software-Defined Networking and network slice architecture.

The IETF Network Slice YANG service model provides the customer-oriented view of the network slice. Once the IETF Network Slice Controller (NSC) receives a request, it needs to map it to accomplish the specific parameters expected by the network controller.

Several Service Models and Network Models may be utilized for realizing an IETF Network Slice service. Those models are analyzed in this documet to understand to what extent they can satisfy the IETF Network Slice requirements. In addition, identified gaps on existing models are reported.

This document also describes the architecture and communication process between the IETF Network Slice Controller and underneath Network Controllers for IETF network slice creation.

1.2. Terminology

The keywords *MUST*, *MUST NOT*, *REQUIRED*, *SHALL*, *SHALL NOT*, *SHOULD*, *SHOULD NOT*, *RECOMMENDED*, *MAY*, and *OPTIONAL*, when they appear in this document, are to be interpreted as described in [RFC2119].

Same terminology as used in [I-D.ietf-teas-ietf-network-slice-definition] and [I-D.draft-nsdt-teas-ns-framework] is primarily used here.

2. Reference Architecture and Components

As described in [I-D.ietf-teas-ietf-network-slice-definition], the IETF Network Slice Controller (NSC) is a functional entity for control and management of IETF network slices. As shown in Figure 1, the NSC exposes an IETF Network Slice Service Interface that allow a higher level system to request an IETF network slice. The NSC IETF Network Slice Service Interface supports the request for enablement of an IETF Network Slice (i.e., creation, modification or deletion). Upon receiving a request from its IETF Network Slice Service Interface, the NSC finds the resources needed for realization of the IETF Network Slice and in turn interfaces with one or more Network Controllers for the realization of the requested IETF Network Slice, through the Network Configuration Interface.

This document focuses on how IETF NSC can be implemented in operator's network.

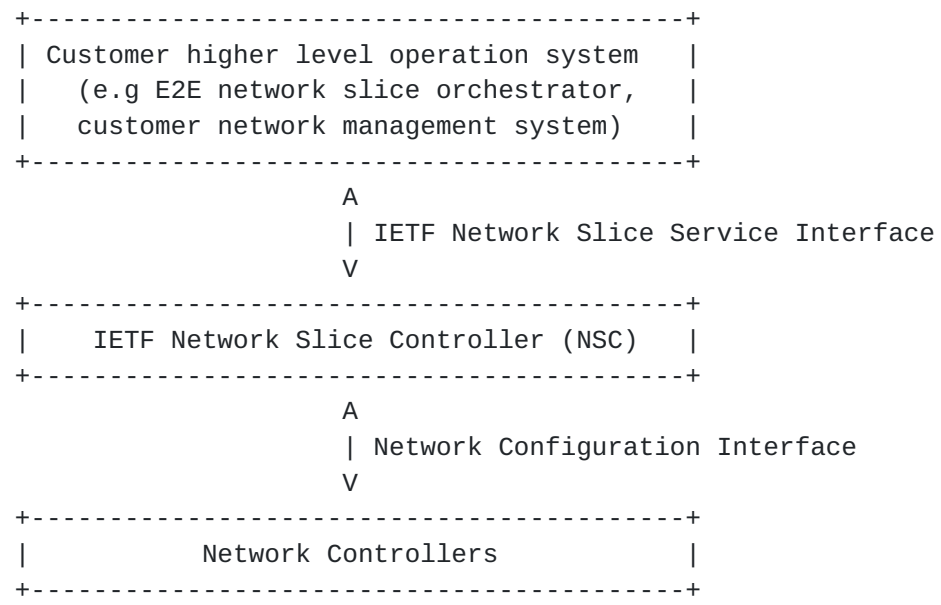

Figure 1: Network Slice Controller as a module of the Hierarchical SDN controller

2.1. Possible architectural options for IETF Network Slice Controller

Several architectural definitions have arisen on the IETF to support SDN and network slicing deployments. The proposal in [I-D.ietf-teas-ietf-network-slices], presented in Figure 1, defines an initial architecture.

Additional approaches are briefly described next.

2.1.1. IETF Network Slice Controller as a module of the Hierarchical SDN controller

The IETF Network Slice Controller function might be part of the Hierarchical network controller (e.g., as the MDSC in the ACTN context, as in [RFC8453]) being a modular function. Below the NSC, a number of network controllers can exist, e.g. each of them handling multiple or single underlay technologies. This approach is represented in Figure 2.

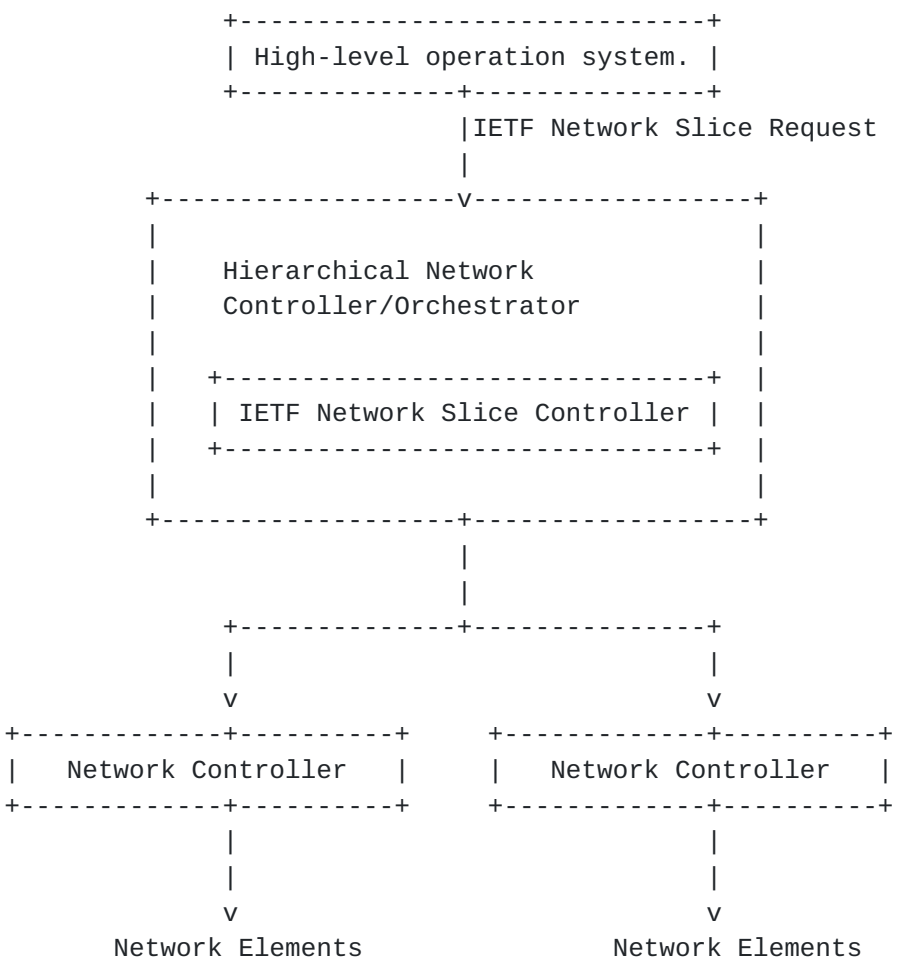

Figure 2: IETF Network Slice Controller as a module of the Hierarchical SDN controller

2.1.2. IETF Network Slice Controller as a stand-alone entity

An alternative implementations can be the one considering the IETF Network Slice Controller as an a stand-alone element, directly interacting with an underlaying network controller, as depicted in Figure 3. In this scenario, the IETF Network Slice Service request can follow a data-enrichment path, where each entity can add more information to the service request.

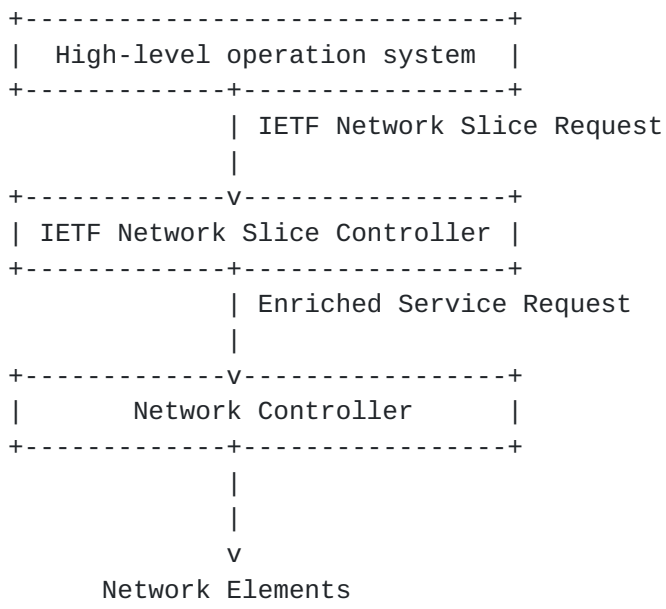

Figure 3: The IETF Network Slice Controller as a stand-alone entity

2.1.3. IETF Network Slice Controller as a module of the Network controller

As another possible implementation, the IETF Network Slice Controller can be an integral part of a Network Controller, directly realizing the network slice service using device data models to configure the network devices. That is, a conventional customer service requests is configured in the form of an IETF Network Slice.

This architecture is depicted in Figure 4.

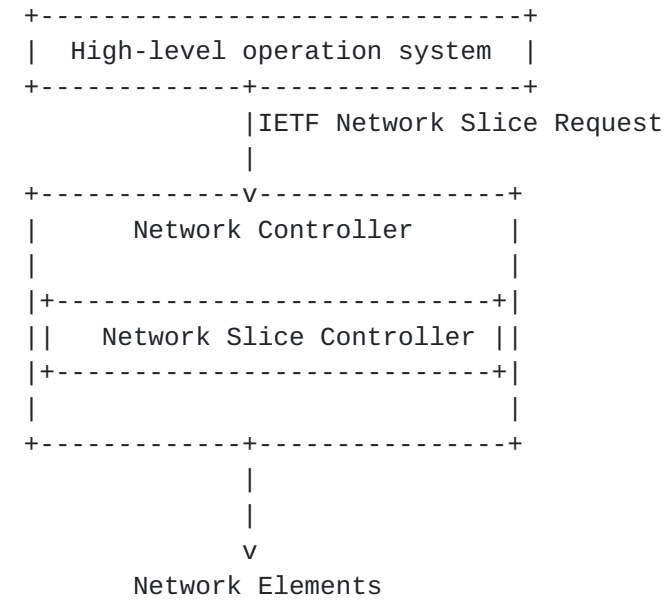

Figure 4: IETF Network Slice Controller as a module of the Network controller

2.2. Possible relationship of IETF Network Slice service model with other models

An IETF Network Slice Service is expected to serve as input from where deriving some other models in the network. According to the architectural options before, different relationships could be considered. Figure 5 reflects such options.

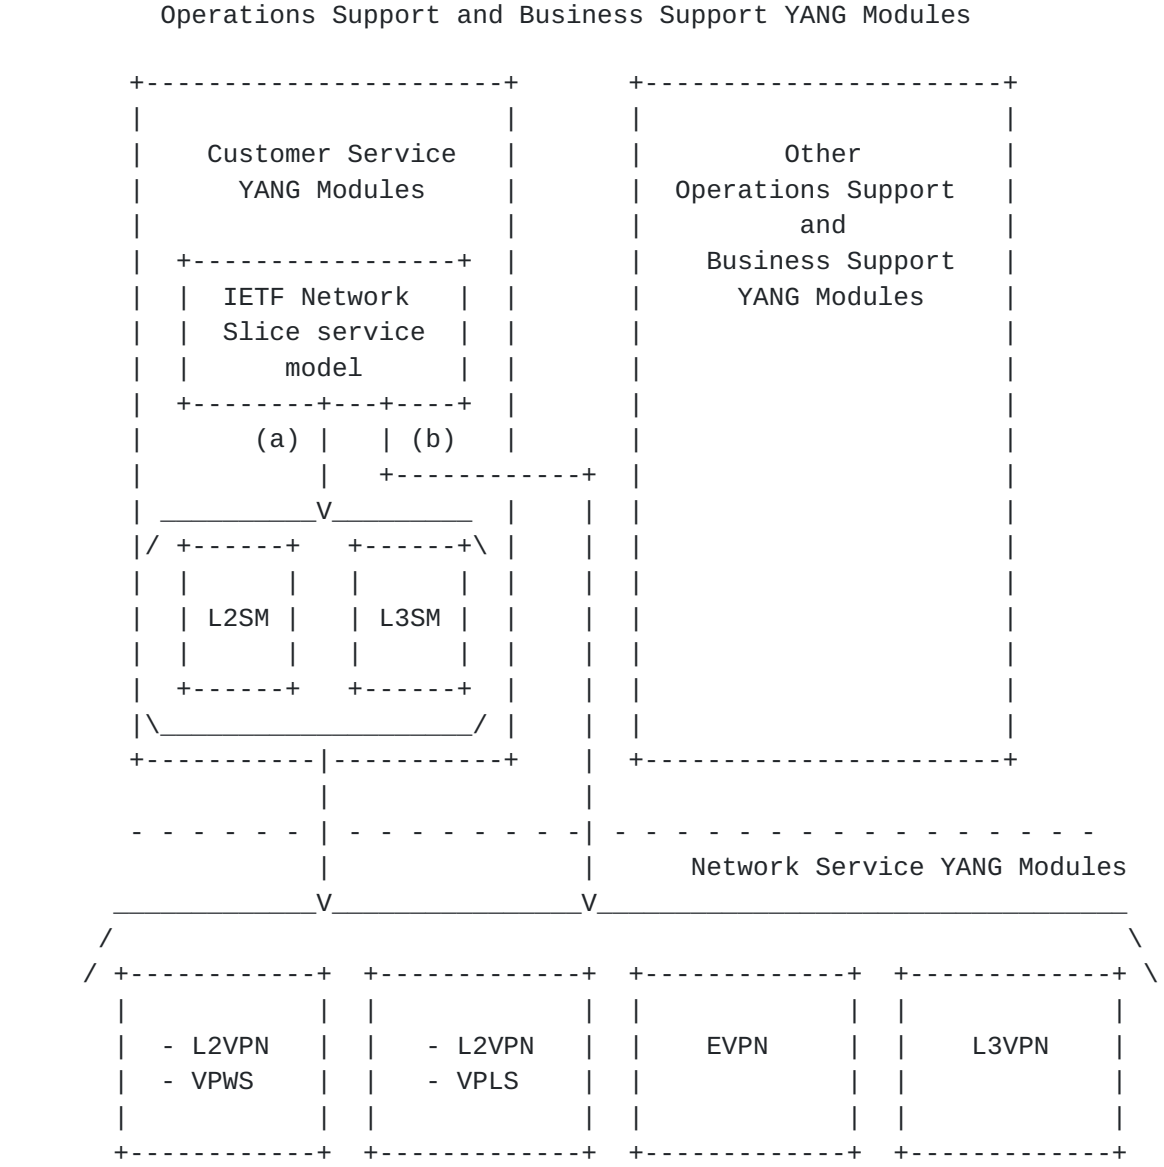

Figure 5: Possible relationships between models

Thus, the IETF Network Slice model (e.g., as defined in [I-D.ietf-teas-ietf-network-slice-nbi-yang] could feed existing service models, such as L2SM or L3SM (case (a) in Figure 5), or could feed existing network models, e.g., EVPN, L3VPN, etc (case (b) in Figure 5). Existing models both for service or network level could require some extensions themselves, or their application in conjunction with some other complementary models (e.g., TE model) to accomplish the service objectives and expectations as declared in the IETF Network Slice model.

3. IETF Network Slice Requirements and Data Models

The main set of requirements for the IETF Slice, based on the high-level slice requirements from multiple organizations and use cases, are compiled in [I-D.ietf-teas-ietf-network-slice-use-cases]. To accomplish those requirements, a set of YANG data models have been proposed.

* [I-D.ietf-teas-ietf-network-slice-nbi-yang]: A Yang Data Model for IETF Network Slice NBI.

* [I-D.liu-teas-transport-network-slice-yang]: Transport Network Slice YANG Data Model.

Those Yang models could be used by an IETF Network Slice Controller to manage CRUD operations on the IETF Network Slice. That is, these models aim capturing the requirements from the consumer of the slice point of view and avoid entering into the detail of how the slice is actually created.

4. Operational Considerations

This section outlines the compliance and operational aspects of Network Controller models with IETF Network slice requirements. [I-D.ietf-teas-ietf-network-slice-use-cases] presents the requirements of the IETF Network slice. In this subsection it is analyzed how available YANG models that can be used by a Network Controller can satisfy those requirements and identify gaps.

Editor's note: the requirements here below represent a sub-set of the overall requirements in [I-D.ietf-teas-ietf-network-slice-use-cases]. Further versions of this document will address other requirements not present in this version.

4.1. Availability

As per [I-D.draft-ietf-teas-te-service-mapping-yang], Availability is a probabilistic measure of the length of time that a VPN/VN instance functions without a network failure. As per [RFC8330], The parameter "availability", as described in G.827, F.1703, and P.530, is often used to describe the link capacity. The availability is a time scale, representing a proportion of the operating time that the requested bandwidth is ensured".

The calculation of the availability is not trivial and would need to be clearly scoped to avoid misunderstandings.

The set of Yang models proposed today allow to request tunnels/paths with different resiliency requirements in terms of protection and restoration. However, none of them include the possibility of requesting a specific availability (e.g. 99.9999%).

4.2. Downlink throughput / Uplink throughput.

The LxVPN Models [RFC9182] and [RFC9291] allow to specify the bandwdidth at the interface level between the slice and the customer. In addition, the Service Mapping model [I-D.draft-ietf-teas-te-service-mapping-yang] allows to bind a VPN to a given LSP, which have its bandwidth requirements. Additionally, TE models can enforce a given bandwidth allocation in the connection between Provider Edges.

Previous comment applies to the incoming and outgoing bandwidth parameters required for the NFV-based services use case in [I-D.ietf-teas-ietf-network-slice-use-cases]. The Network sharing use case has Maximum and Guaranteed Bit Rate parameters. These parameters can be mapped to the TE tunnel models when setting up LSPs [I-D.draft-ietf-teas-yang-te].

4.3. Protection scheme

Protection schemes are mechanisms to define how to setup resources for a given connection. TE tunnel models [I-D.draft-ietf-teas-yang-te] includes protection and restoration as two main attributes. The parameters included in the containers for protection and restoration cover the requirements of the IETF NS related with protection schemes. Similarly, TE models cover the parameter 'recovery time' for the network sharing use case.

4.4. Delay

Delay is a critical parameter for several IETF NS types. Every use-case defined in [I-D.ietf-teas-ietf-network-slice-use-cases] contains delay constraints. 5G use cases require 'delay tolerance', NFV-based services have the delay information within 'QoS metrics' and 'Bounded latency' in the network sharing use case.

During the realization of the IETF Network Slice, these parameters are part of the requirements of a TE tunnel configuration [I-D.draft-ietf-teas-yang-te].
They can be included within the 'path-metric-bounds' parameter, so the created LSP fulfils the given metrics bounds like 'path-metric-delay-average' or 'path-metric-delay- minimum'.

4.5. Packet loss rate

The packet loss rate indicates the maximum rate for lost packets that the service tolerates in the link. During the realization of the IETF Network Slice, this attribute will influence the tunnel selection and the value is included in the [I-D.draft-ietf-teas-yang-te] document as the 'path-metric-loss".
The 'path-metric-loss' is a metric type, which measures the percentage of packet loss of all links traversed by a P2P path. This parameter is required for 5G services and network sharing use-case, while it is part of the 'QoS metrics' for the NFV-based services.

5. Relationship between IETF NBI model parameters anf those in service and network models

5.1. Relationship between IETF NBI model parameters and L3SM and L2SM model parameters

This section presents an initial analysis of the relationship between IETF NBI model parameters and L3SM and L2SM service model parameters.

The L3SM service parameters are defined in section 6.2 of [RFC8299].

The following parameters are considered, so far:

* Bandwidth. This parameter indicates the bandwidth requirement between each CE and PE participating in the service, then referrign essentially to the required WAN link bandwidth. It is expressed in terms of bits per second and individually specified for both input and output. Despite it is not stated in RFC 8299, this parameter can be interpreted as the CIR/PIR expected for the CE - PE connection.

* MTU. This parameter indicates the maximum PDU size expected for the layer-3 service. It is relevant since packets could be discarded in case the customer sends packets with longer MTU than the one expressed by this parameter.

* QoS. Regarding QoS, two different kind of parameters are detailed.

  - QoS classification policy. This policy is used to classify the traffic received from the customer, and it is expressed as a set of ordered rules. It is used for marking the input traffic (from CE to PE) when the customer flows match any of the rules in the list, setting the appropriate target class of service (target-class-id).

  - QoS profile. This profile defines the traffic-scheduling to be applied to the flows for either Site-to-WAN, WAN-to-Site, or both directions. It contains the following information per class of service: rate-limit, latency, jitter and guaranteed bandwidth.

* Multicast. This parameter identifies if the service is multicast, and if so, what is the role of the site in the customer multicast service topology (i.e., source, receiver, or both). It also defines the kind of multicast relationship with the customer (i.e., as a router requiring PIM, host requiring either IGMP or MLD, or both), as well as the support of IPv4, IPv6 or both.

Similarly L2SM model parameters are described in section 5.9 and 5.10 of [RFC8466].

* Bandwidth. This parameter is related to the bandwidth between both CE and PE and can be expressed as CIR/EIR/PIR, in the ingress or egress direction, taking the CE as the point of reference.

* MTU. This parameter refers to the maximum layer-2 PDU frame size.

* QoS. The specification of the QoS follows a similar structure to the one described in the case of L3SM. Some differences apply, for instance, at the time of QoS classification, which is performed on top of layer-2 parameters (e.g., MAC addresses).

* BUM traffic. This parameter allows to determine if a site acts as source, receiver, or both.

* Availability. This parameter in the L2SM model relates to the capability of supporting multi-homing.

On the other hand, the IETF NS NBI YANG model supports a number of SLOs and SLEs in the form of network slice service policy attributes. Such policy can apply to per-network slice, per-connection group or per-connection indivudually (over-writting of attributes is allowed as more granular information is provided). The following SLO attributes are detailed:

* One-way / Two-way bandwidth, indicating the guaranteed minimum bandwidth between any two NSEs (unidirectional / bidirectional).

* One-way / Two-way latency, indicating the guaranteed minimum latency between any two NSEs (unidirectional / bidirectional).

* One-way / Two-way delay variation, indicating the maximum permissible delay variation of the slice (unidirectional / bidirectional).

* One-way / Two-way packet loss, indicating the maximum permissible packet loss rate between endpoints (unidirectional / bidirectional).

Additionally, the following SLEs are defined:

* MTU, referring to the the maximum PDU size that the customer may use.

* Security, indicating if encryption or other security measures are required between two endpoints.

* Isolation, as a way of indicating the isolation level expected by the customer in the allocation of network resources.

* Maximum occupancy level, to express the amount of flows to be admitted (and optionally a maximum number of countable resource units such as IP or MAC addresses).

Thus, an initial mapping between L3SM, L2SM and IETF NS NBI model can be performed as indicated in the follwoing table.

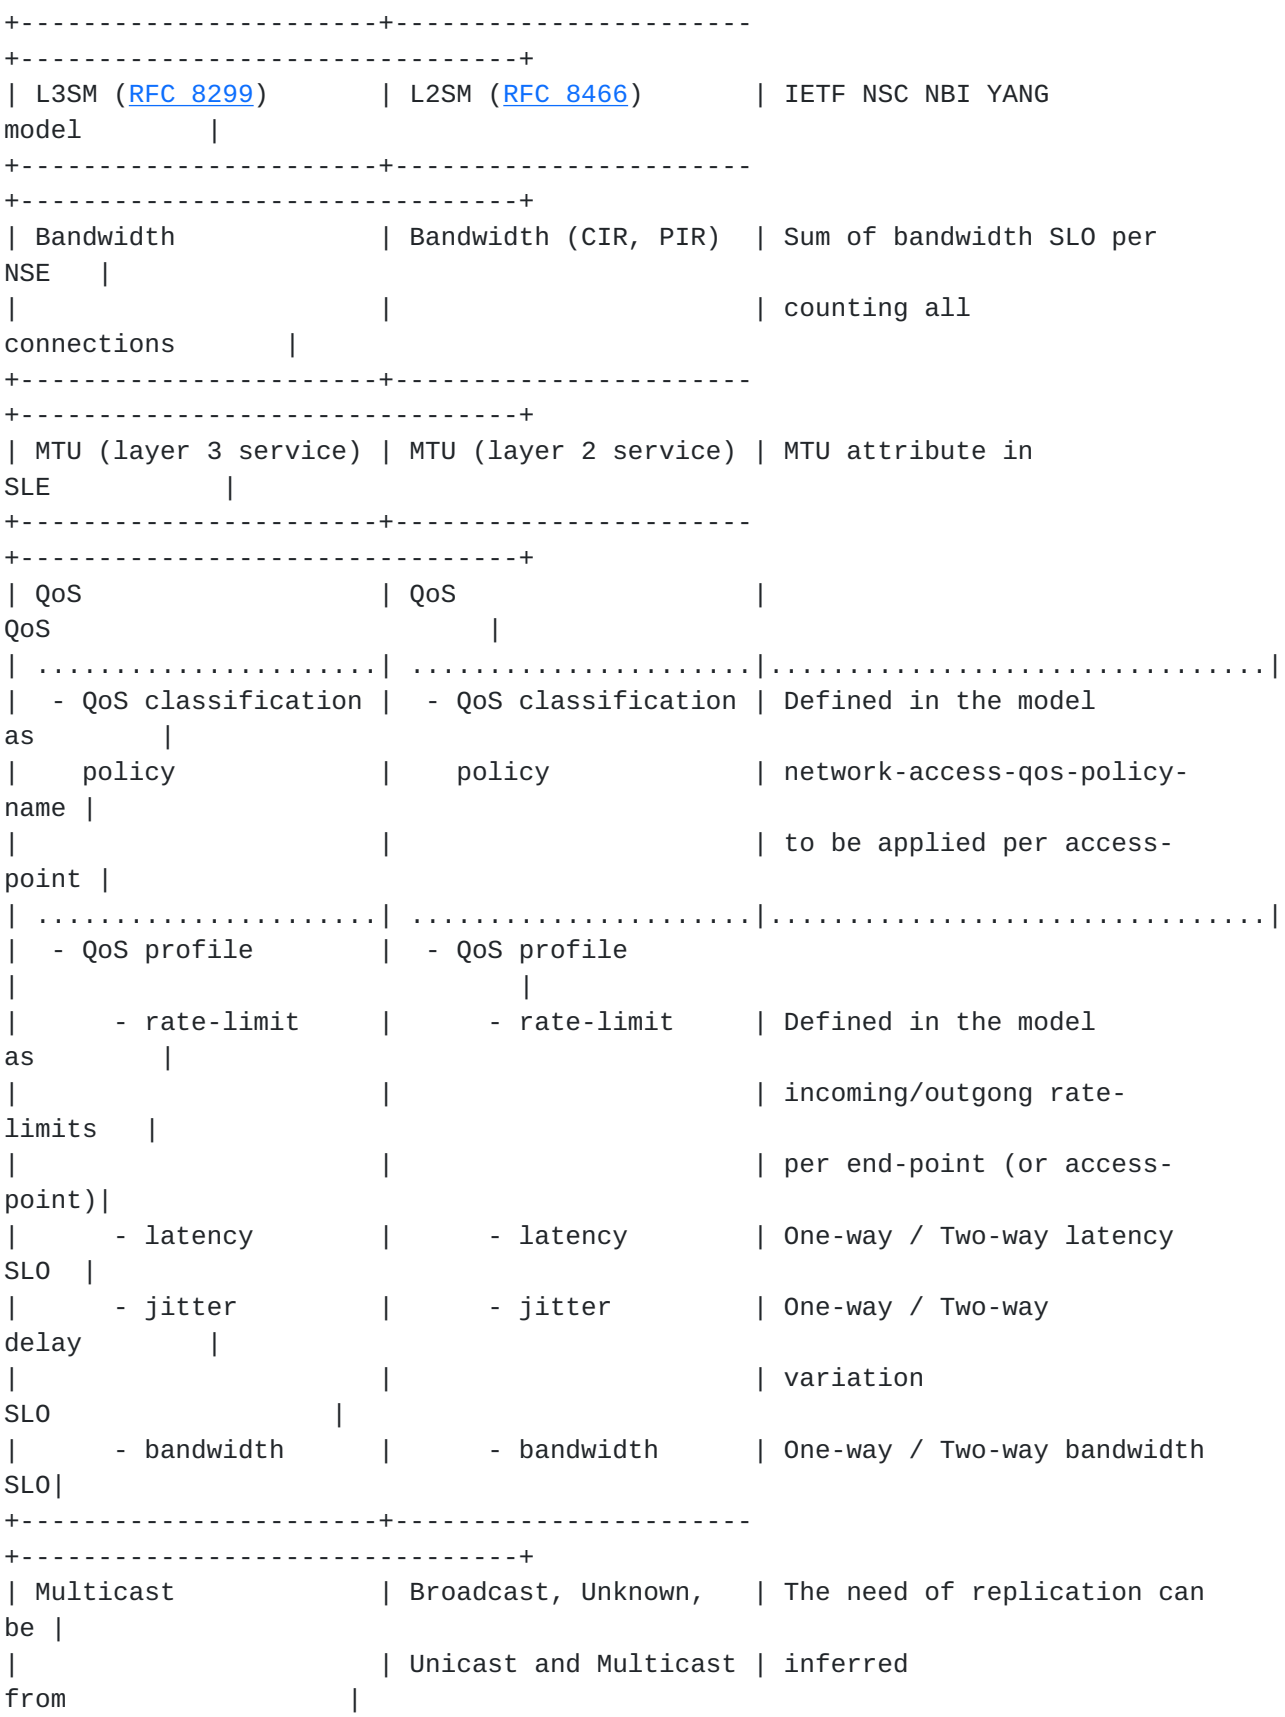
L3SM (RFC 8299)
L2SM (RFC 8466)
IETF NSC NBI YANG model
Bandwidth
Bandwidth (CIR, PIR)
Sum of bandwidth SLO per NSE counting all connections
MTU (layer 3 service)
MTU (layer 2 service)
MTU attribute in SLE
QoS
QoS
QoS
- QoS classification
policy
- QoS classification
policy
Defined in the model as network-access-qos-policy-name to be applied per access-point
- QoS profile
- QoS profile
- rate-limit
- rate-limit
Defined in the model as incoming/outgong rate-limits per end-point (or access-point)
- latency
- latency
One-way / Two-way latency SLO
- jitter
- jitter
One-way / Two-way delay variation SLO
- bandwidth
- bandwidth
One-way / Two-way bandwidth SLO
Multicast
Broadcast, Unknown, Unicast and Multicast
The need of replication can be inferred from

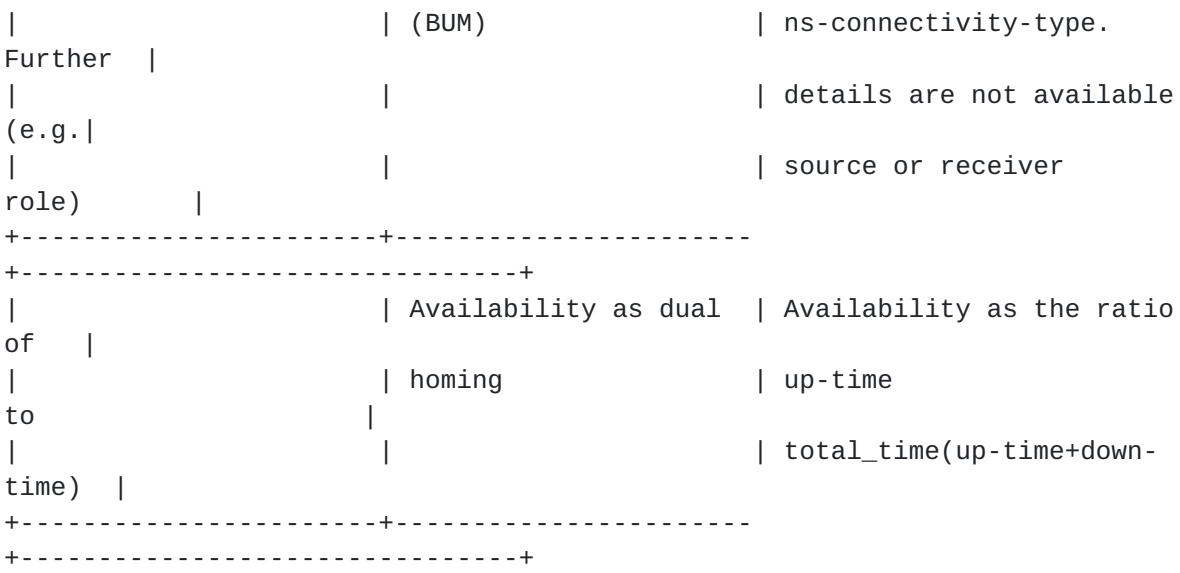

Figure 6: Mapping of IETF NS NBI and LxSM service attribute

The following consideration can be made.

* While the QoS profile in L3SM and L2SM applies per service class, the parameters in IETF NS NBI apply per connection. So if per-class granularity is required in an IETF network slice, then different connections have to be defined between the same end-points, one per service class.

* A number of attributes are not defined in L3SM nor L2SM such as packet loss, isolation or security. Then L3SM and L2SM could not be sufficient to realize IETF network slices with such specific needs, unless those other objectives and expectations are provided by other means (e.g., realizing the L3SM thorugh technologies guaranteing dedicated resource allocation such as OTN).

5.2. Relationship between IETF NBI model parameters and L3NM and L2NM model parameters

This section presents an initial analysis of the relationship between IETF NBI model parameters and L3NM and L2NM network model parameters.

The L3NM service parameters are defined in section 7.6.6 of [RFC9182].

As made in the previous section, some basic parameters are considered:

* Bandwidth: The L3NM defines bandwidth in terms of the 'pe-to-ce-bandwidth' & 'ce-to-pe-bandwidth'. Both values are defined in absolute value in bps per interface. The model supports the usage of QoS policies to include inbound and outbound Rate limits.

* MTU: L3NM only supports the definition at vpn-network-access level.

* QoS: The quality of service is differentiated in three-levels:

  - QoS Profile: Allows the reference of an existing profile. The profile creation is out-scope of the model.

  - QoS Classification: Customize policy creation rules, including quote name and upper and lower limits.

  - QoS Action: Allows the filtering of incoming and outcoming rate limits.

* Multicast: mVPN is supported at vpn-node and vpn-network-access; Each level includes Rendezvous Point (RP), IGMP, PIM and MLD definitions.

Similarly L2NM model parameters are described in section 7.6.6 of [RFC9291].

* Bandwidth: The L2NM considers the same parameters 'pe-to-ce-bandwidth' & 'ce-to-pe-bandwidth'. However, per definition, the L2NM supports the differentiation of CIR, PIR values. It includes the same set of values described for the L2SM model.

* MTU: L2NM differentiates among Service MTU and interface MTU. The MTU mismatch configuration is also supported as part of the vpn-service configuration.

* QoS: The quality of service is differentiated in two-levels:

  - QoS Profile: Reference an existing profile. Creation is out-scope of the model.

  - QoS Classification: Customize policy creation rules, including quote name and limits.

* Multicast: Discard options are available for unknown Broadcast, Unicast or Multicast (BUM).

Thus, an initial mapping between L3NM, L2NM and IETF NS NBI model can be performed as indicated in the follwoing table.

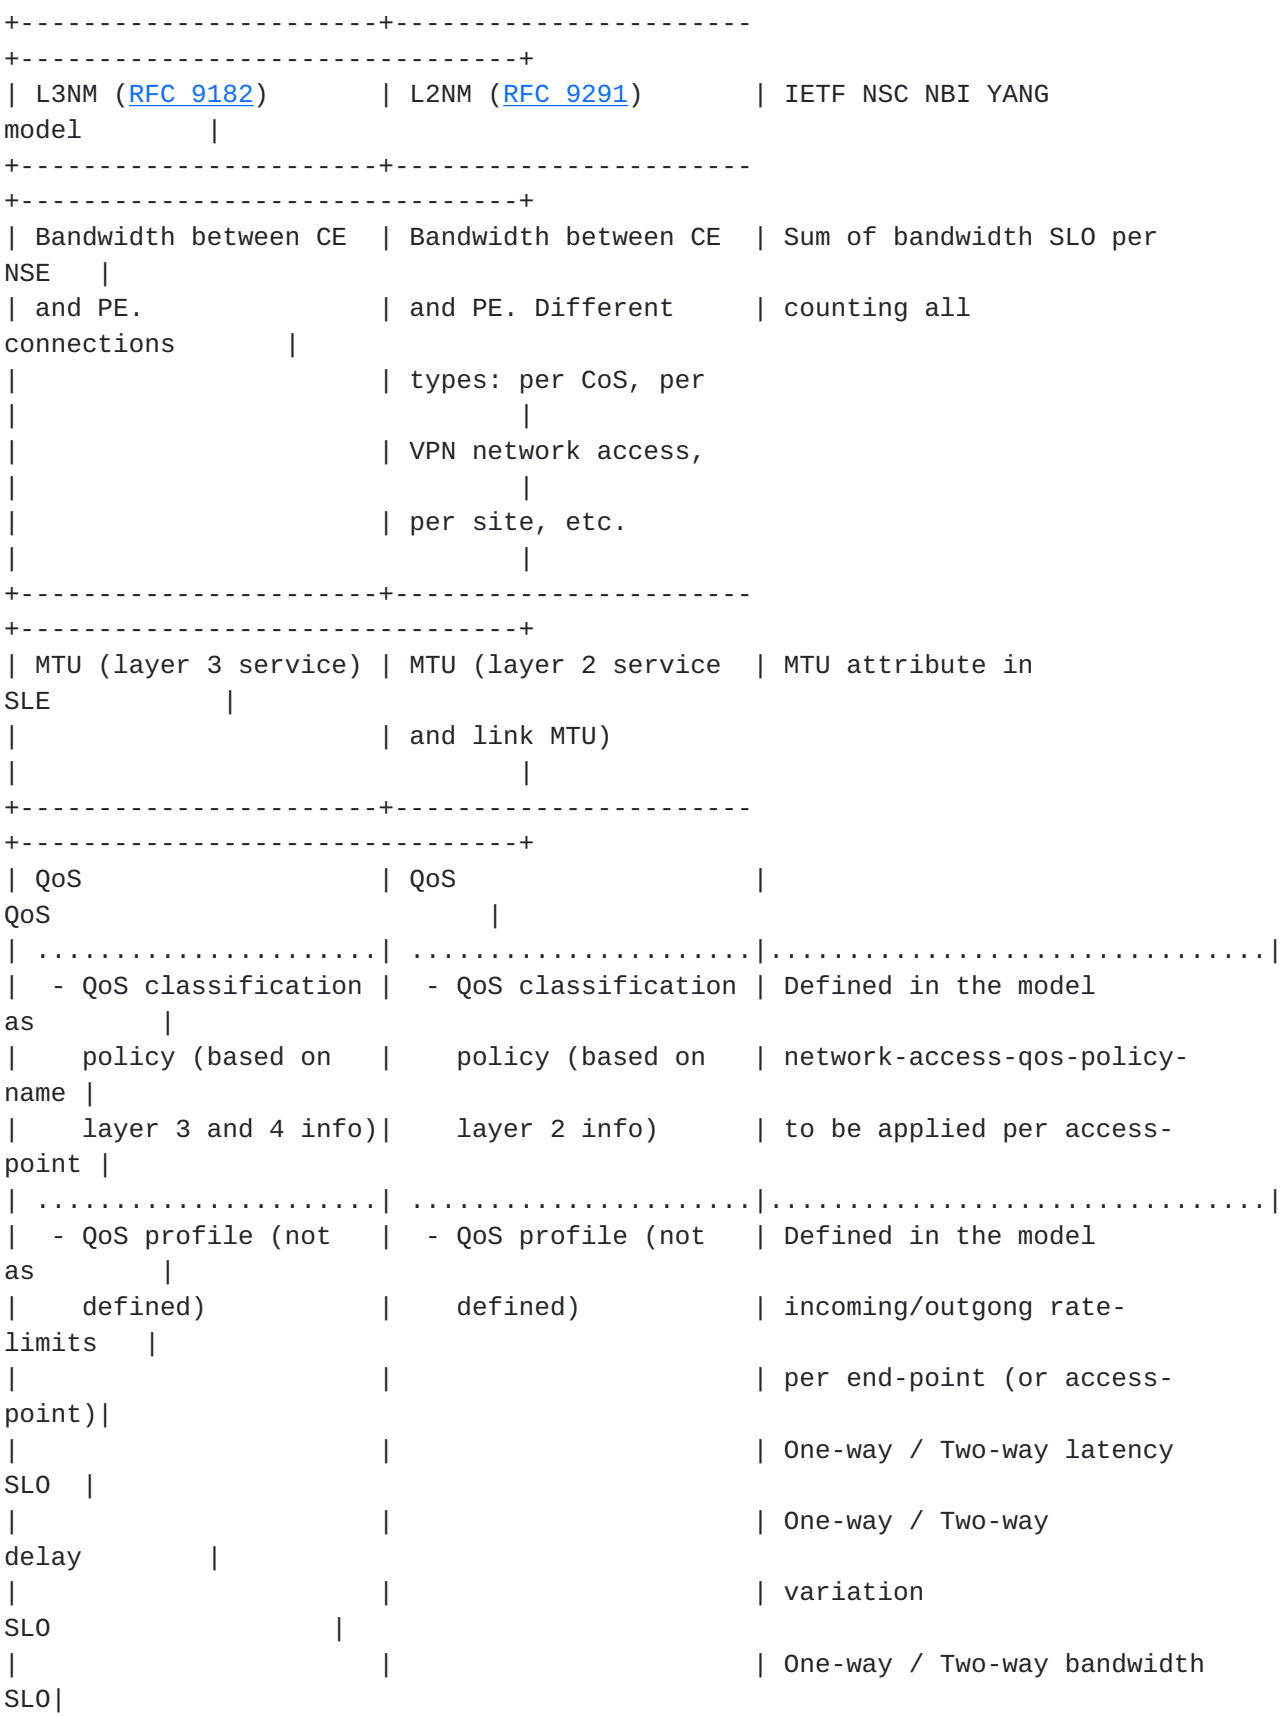
L3NM (RFC 9182)
L2NM (RFC 9291)
IETF NSC NBI YANG model
Bandwidth between CE and PE.
Bandwidth between CE and PE. Different types: per CoS, per VPN network access, per site, etc.
Sum of bandwidth SLO per NSE counting all connections
MTU (layer 3 service)
MTU (layer 2 service and link MTU)
MTU attribute in SLE
QoS
QoS
QoS
- QoS classification as policy (based on layer 3 and 4 info)
- QoS classification policy (based on layer 2 info)
Defined in the model as network-access-qos-policy-name to be applied per access-point
- QoS profile (not defined)
- QoS profile (not defined)
Defined in the model as incoming/outgong rate-limits per end-point (or access-point)
One-way / Two-way latency SLO
One-way / Two-way delay variation SLO
One-way / Two-way bandwidth SLO

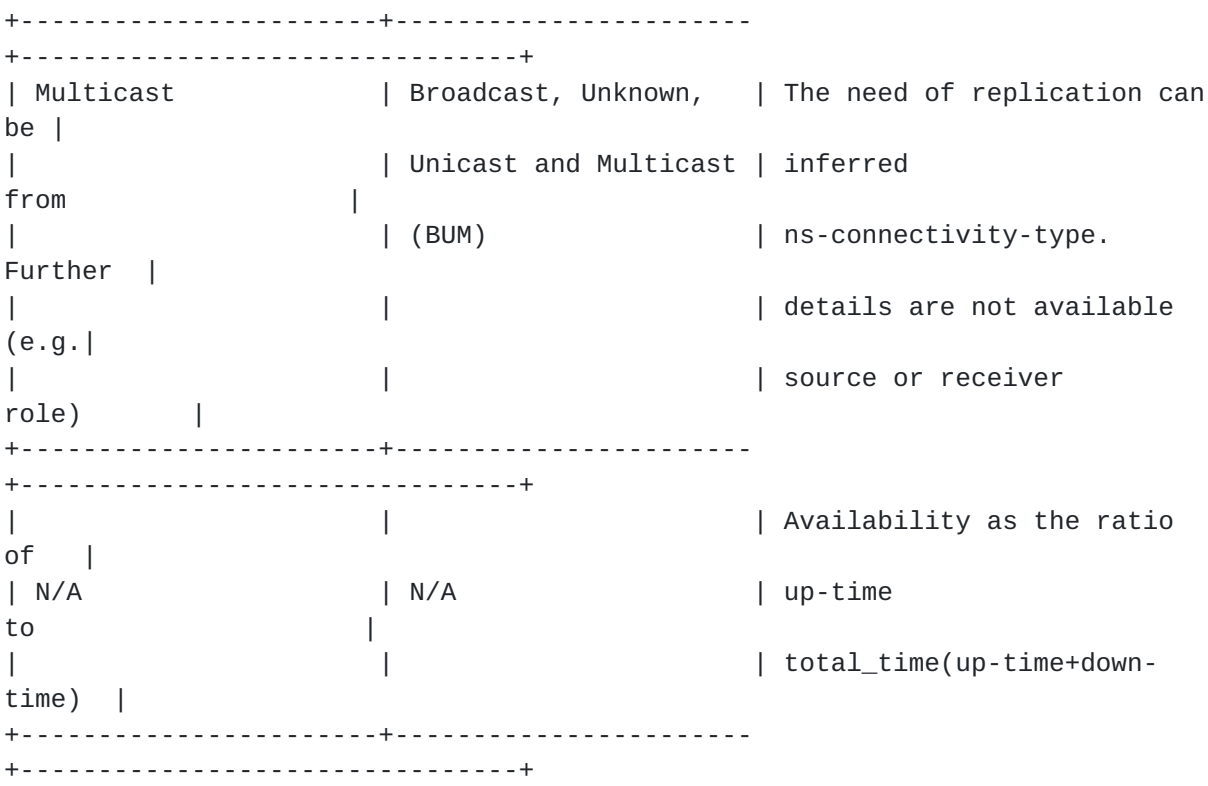

Figure 7: Mapping of IETF NS NBI and LxNM service attribute

6. IETF Network Slice Procedure

6.1. IETF Network Slice provisioning workflow

An IETF Network Slice may use several underlying technologies. The creation of a new IETF Network Slice will be initiated with following three steps:

1. A higher level system requests connections with specific characteristics via the IETF Network Slice Service interface.

2. This request is processed by an IETF NSC which specifies a mapping between the IETF Network Slice Service request to any of the IETF Services, Tunnels, and paths models.

3. A series of requests for creation of services, tunnels and paths is sent out to the network controllers underneath to realize the IETF Network Slice.

4. The final configuration is performed by means of Network Controller operations

As a functional entity responsible for managing a network domain, a network controller can expose its northbound interface based on YANG models. The IETF Network Slice Controller can use the network controller's NBI during the realization of IETF Network Slice. The following network models can be used for realization of IETF Network slices:

* LxVPN Network models:

  - These models describe a VPN service from the network point of view. It supports the creation of Layer 3 and Layer 2 services using several control planes.

* Traffic Engineering models:

  - These models allow to manipulate Traffic Engineering tunnels within the network segment. Technology-specific extensions allow to work with a desired technology (e.g. MPLS RSVP-TE tunnels, Segment Routing paths, OTN tunnels, etc.)

* TE Service Mapping extensions:

  - These extensions allow to specify for LxVPN the details of an underlay based on TE.

* ACLs and routing policies models:

  - Even though ACLs and routing policies are device models, it's exposure in the NBI of a domain controller allows to provide an additional granularity that the network domain controller is not able to infer on its own.

6.2. LxVPN Service Models

The framework defined in [RFC8969] compiles a set of YANG data models for automating network services. The data models can be used during the service and network management life cycle (e.g., service instantiation, service provisioning, service optimization, service monitoring, service diagnosing, and service assurance). The Service models could be a realization of IETF Network slice requests.

The following models are examples of Network models that describe services.

* [RFC8049]: YANG Data Model for L3VPN Service Delivery.

* [RFC8466]: A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery.

6.3. LxVPN Network Models

Similar to the Service Models, the framework defined in [RFC8969] compiles a set of YANG data models for automating network services. The Network models could be reused for the realization of Network slice requests.

The following models are examples of Network models that describe services.

* [RFC9182]: A Layer 3 VPN Network YANG Model

* [RFC9291]: A Layer 2 VPN Network YANG Model

6.4. Traffic Engineering Models

The TEAS WG has defined a collection of models to allow the management of Traffic Engineering tunnels.

* [I-D.ietf-teas-yang-te]: A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces. The model allows to instantiate paths in a TE enabled network. Note that technology augmented models are require to particular per-technology instantiations.

6.5. Traffic Engineering Service Mapping

The IETF has defined a YANG model to set up the procedure to map VPN service/network models to the TE models. This model, known as service mapping, allows the network controller to assign/retrieve transport resources allocated to specific services. At the moment there is just one service mapping model [I-D.ietf-teas-te-service-mapping-yang]. The "Traffic Engineering (TE) and Service Mapping Yang Model" augments the VPN service and network models.

7. Potential usage of models in alternative IETF NSC architectures

This section does not intend to be prescriptive but descriptive about the potential usage of existing and proposed models for the provision of an IETF Network Slice service.

[I-D.draft-contreras-teas-slice-controller-models] shows a potential internal structure of an IETF Network Slice Controller which can be divided into two components:

* IETF Network Slice Mapper: this high-level component processes the customer request, putting it into the context of the overall IETF Network Slices in the network.

* IETF Network Slice Realizer: this high-level component processes the complete view of transport slices including the one requested by the customer, decides the proper technologies for realizing the IETF Network Slice and triggers its realization.

Note that this division in functional components of the IETF NSC is just a potential option, not constraining any other implementation of functional structure.

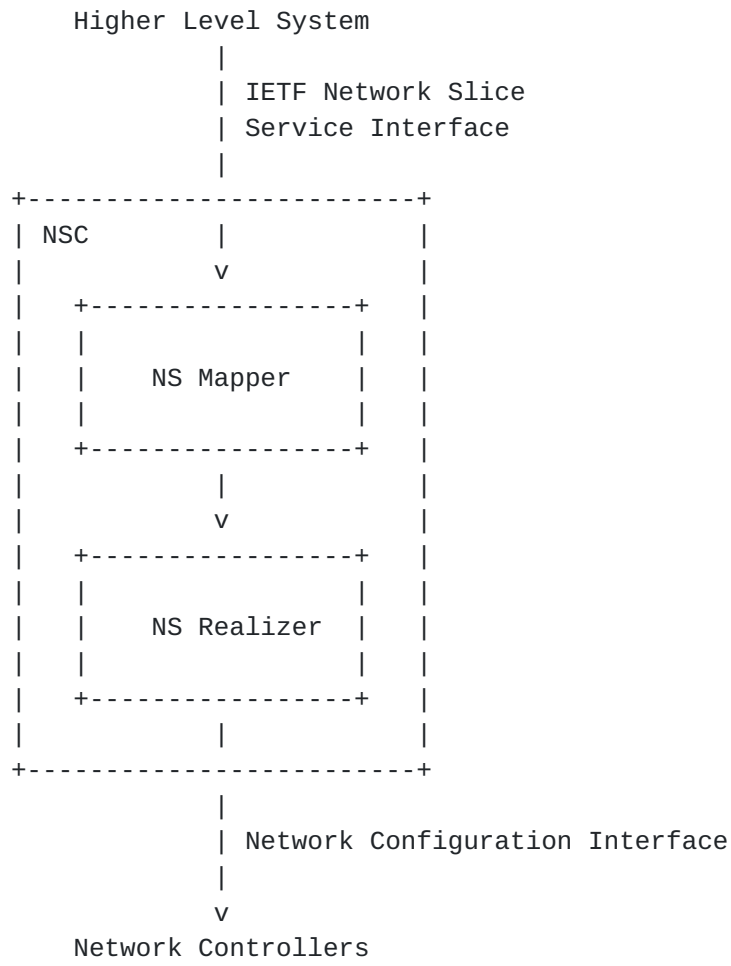

Figure 8: IETF Network Slice Controller Structure

The details of IETF network slice mapper and realizer are provided below for various implementation of NSC.

7.1. IETF Network Slice requested to Hierarchical Network Controller

Referring to Figure 2, in an integrated architecture the IETF Network Slice Controller (NSC) is part of a Hierarchical SDN controller module. The NSC and the Hierarchical Network Controller should share the same internal data and the same IETF Network Slice Service interface. Thus, the H-SDN module must be able to:

* Map: The customer request received using the [I-D.ietf-teas-ietf-network-slice-nbi-yang] must be processed by the NCS. The mapping process takes the network-slice SLAs selected by the customer to available Routing Policies and Forwarding policies.

* Realize: Create necessary network requests. The slice's realization can be translated into one or several LxNM Network requests, depending on the number of underlay controllers. Thus, the NCS must have a complete view of the network to map the orders and distribute them across domains. The realization should include the expansion/selection of Forwarding Policies, Routing Policies, VPN policies, and Underlay transport preference.

To maintain the data coherence between the control layers, the IETF Network Slice ID used in [I-D.ietf-teas-ietf-network-slice-nbi-yang] could be directly mapped to the transport-instance-id at the VPN-Node level.

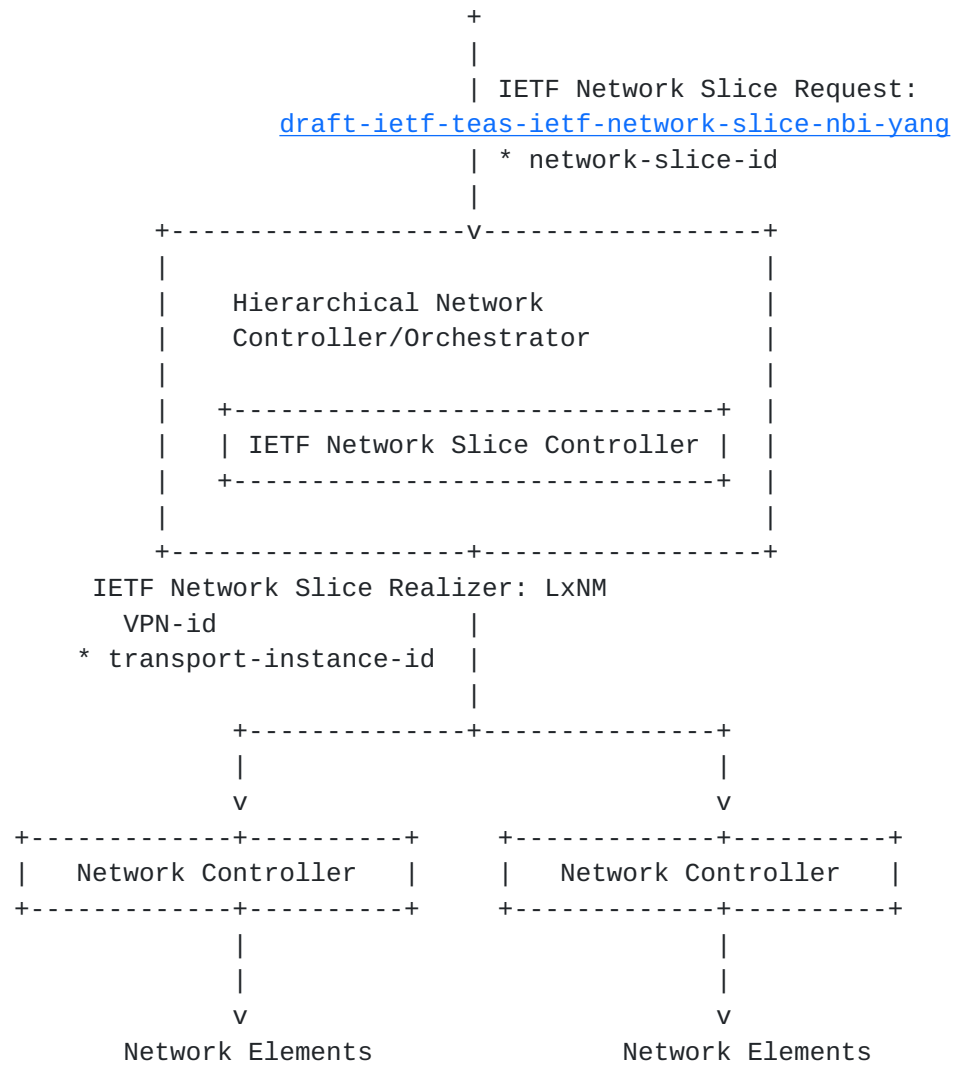

Figure 9: Workflow for the slice request in an integrated architecture

7.2. IETF Network Slice requested to Network Slice Controller

Referring to Figure 2 when the Network Slice Controller is a stand-alone controller module, the NSC should perform the same two tasks described in section 6.1:

* Map: Process the customer request. The customer request can be sent using the [I-D.draft-liu-teas-transport-network-slice-yang]. This draft allows the topology mapping of the Slice request.

* Realize: Create necessary network requests. The slice's realization will be translated into one LXNM Network request. As the NCS has a topological view of the network, the realization can include the customer's traffic engineering transport preferences and policies.

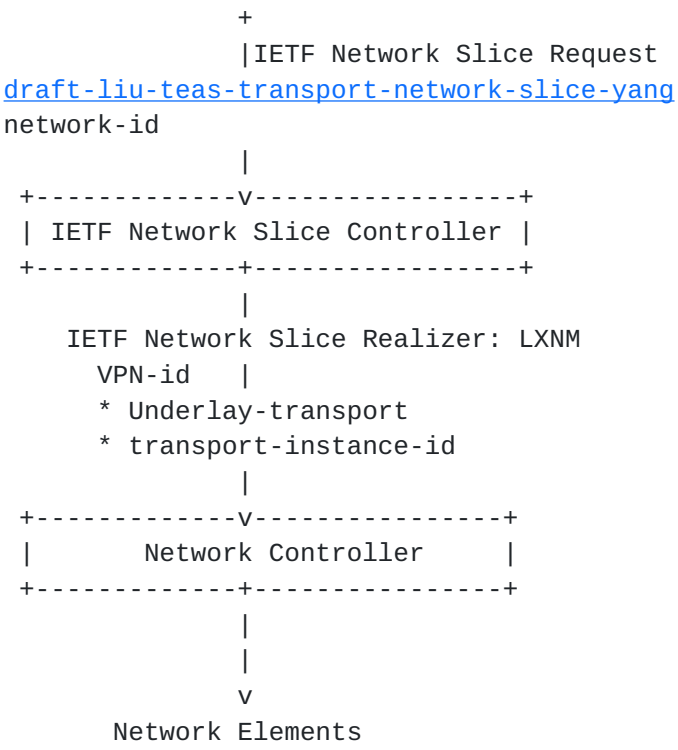

Figure 10: Workflow for the slice request in an stand-alone architecture

7.3. Network Slice Controller as part of the domain controller

The Network Slice Controller can be a module of the Network controller. In that case, two options are available. One is to share the same device data model in the NBI and SBI of the SDN controller. The direct translation would reduce the service logic implemented at the SDN controller level, grouping the mapping and translation into a single task:

* Realize: As the device models are part of the network controller's NBI thus, the realization can be done by the network controller applying a simple service logic to send the Network elements.

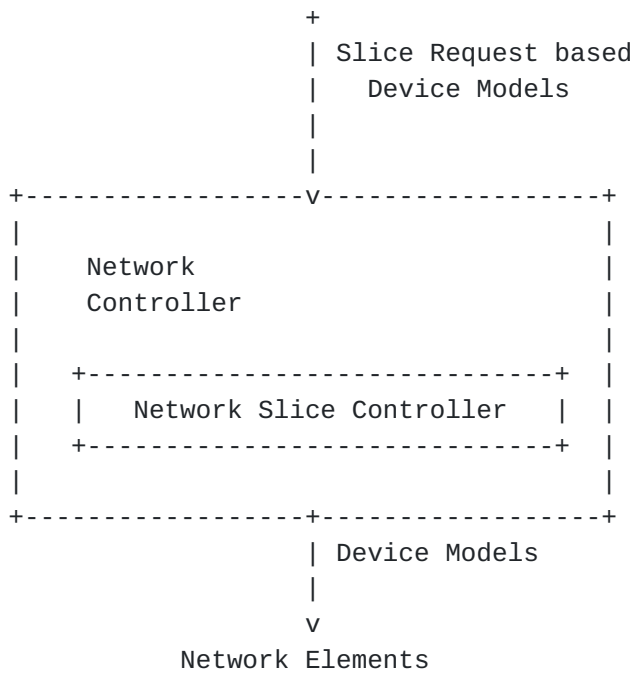

Figure 11: Workflow for the slice request in an stand-alone architecture

A second option introduces a more complex logic in the network controller and creates an abstraction layer to process the transport slices. In that case, the controller should receive network slices creation requests and maintain the whole set of implemented slices:

* Map & Realize: The mapping and realization can be done by the Domain controller applying the service logic to create policies directly on the Network elements.

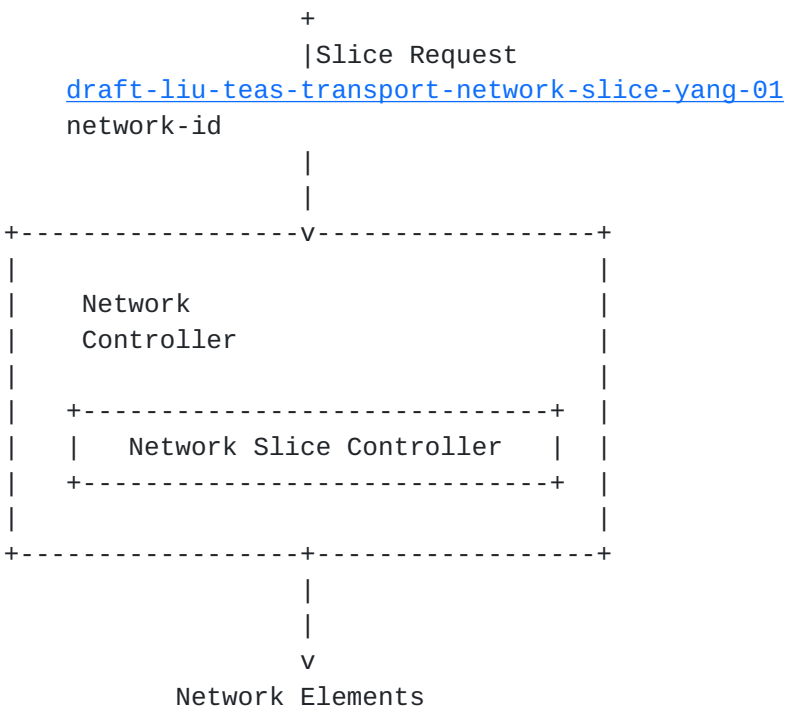

Figure 12: Workflow for the slice request in an stand-alone architecture

8. Security Considerations

There are two main aspects to consider. On the one hand, the IETF Network Slice has a set of security related requirements, such as hard isolation of the slice, or encryption of the communications through the slice. All those requirements need to be analyzed in detailed and clearly mapped to the Network Controller and device interfaces.

On the other hand, the communication between the IETF network slicer and the network controller (or controllers or hierarchy of controllers) need to follow the same security considerations as with the network models.

The network YANG modules defines schemas 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 [RFC8466].

The Network Configuration Access Control Model (NACM) [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.

The following summarizes the foreseen risks of using the Network Models to instantiate IETF network Slices:

* Malicious clients attempting to delete or modify VPN services that implements an IETF network slice. The malicious client could manipulate security related aspects of the network configuration that impact the requirements of the slice, failing to satisfy the customer requirement.

* Unauthorized clients attempting to create/modify/delete a VPN hat implements an IETF network slice service.

* Unauthorized clients attempting to read VPN services related information hat implements an IETF network slice

* Malicious clients attempting to leak traffic of the slice.

9. IANA Considerations

This document is informational and does not require IANA allocations.

10. Conclusions

A wide variety of yang models are currently under definition in IETF that can be used by Network Controllers to instantiate IETF network slices. Some of the IETF slice requirements can be satisfied by multiple means, as there are multiple choices available. However, other requirements are still not covered by the existing models. A more detailed definition of those uncovered requirements would be needed. Finally, a consensus on the set of models to be exposed by Network Controllers would facilitate the deployment of IETF network slices.

Acknowledgments

This work is partially supported by the European Commission under Horizon 2020 grant agreement number 101015857 Secured autonomic traffic management for a Tera of SDN flows (Teraflow).

Informative References

[I-D.draft-contreras-teas-slice-controller-models]
Contreras, L. M., Rokui, R., Tantsura, J., Wu, B., Liu, X., Dhody, D., and S. Belotti, "IETF Network Slice Controller and its associated data models", Work in Progress, Internet-Draft, draft-contreras-teas-slice-controller-models-04, 24 October 2022, <https://datatracker.ietf.org/doc/html/draft-contreras-teas-slice-controller-models-04>.

[I-D.draft-ietf-teas-te-service-mapping-yang]
Lee, Y., Dhody, D., Fioccola, G., Wu, Q., Ceccarelli, D., and J. Tantsura, "Traffic Engineering (TE) and Service Mapping YANG Data Model", Work in Progress, Internet-Draft, draft-ietf-teas-te-service-mapping-yang-13, 11 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-service-mapping-yang-13>.

[I-D.draft-ietf-teas-yang-te]
Saad, T., Gandhi, R., Liu, X., Beeram, V. P., Bryskin, I., and O. G. de Dios, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-32, 12 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-32>.

[I-D.draft-liu-teas-transport-network-slice-yang]
Liu, X., Tantsura, J., Bryskin, I., Contreras, L. M., Wu, Q., Belotti, S., Rokui, R., Guo, A., and I. Busi, "IETF Network Slice Topology YANG Data Model", Work in Progress, Internet-Draft, draft-liu-teas-transport-network-slice-yang-06, 13 March 2023, <https://datatracker.ietf.org/doc/html/draft-liu-teas-transport-network-slice-yang-06>.

[I-D.draft-nsdt-teas-ns-framework]
Gray, E. W. and J. Drake, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-nsdt-teas-ns-framework-05, 2 February 2021, <https://datatracker.ietf.org/doc/html/draft-nsdt-teas-ns-framework-05>.

[I-D.ietf-teas-ietf-network-slice-definition]
Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Definition of IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-definition-01, 22 February 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slice-definition-01>.

[I-D.ietf-teas-ietf-network-slice-nbi-yang]
Wu, B., Dhody, D., Rokui, R., Saad, T., Han, L., and J. Mullooly, "A YANG Data Model for the IETF Network Slice Service", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-nbi-yang-04, 13 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slice-nbi-yang-04>.

[I-D.ietf-teas-ietf-network-slice-use-cases]
Contreras, L. M., Homma, S., Ordonez-Lucena, J. A., Tantsura, J., and H. Nishihara, "IETF Network Slice Use Cases and Attributes for the Slice Service Interface of IETF Network Slice Controllers", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-use-cases-01, 24 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slice-use-cases-01>.

[I-D.ietf-teas-ietf-network-slices]
Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "A Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-19, 21 January 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-19>.

[I-D.ietf-teas-te-service-mapping-yang]
Lee, Y., Dhody, D., Fioccola, G., Wu, Q., Ceccarelli, D., and J. Tantsura, "Traffic Engineering (TE) and Service Mapping YANG Data Model", Work in Progress, Internet-Draft, draft-ietf-teas-te-service-mapping-yang-13, 11 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-service-mapping-yang-13>.

[I-D.ietf-teas-yang-te]
Saad, T., Gandhi, R., Liu, X., Beeram, V. P., Bryskin, I., and O. G. de Dios, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-32, 12 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-32>.

[I-D.liu-teas-transport-network-slice-yang]
Liu, X., Tantsura, J., Bryskin, I., Contreras, L. M., Wu, Q., Belotti, S., Rokui, R., Guo, A., and I. Busi, "IETF Network Slice Topology YANG Data Model", Work in Progress, Internet-Draft, draft-liu-teas-transport-network-slice-yang-06, 13 March 2023, <https://datatracker.ietf.org/doc/html/draft-liu-teas-transport-network-slice-yang-06>.

[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/rfc/rfc2119>.

[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/rfc/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/rfc/rfc6242>.

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

[RFC8049] Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8049, DOI 10.17487/RFC8049, February 2017, <https://www.rfc-editor.org/rfc/rfc8049>.

[RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8299, DOI 10.17487/RFC8299, January 2018, <https://www.rfc-editor.org/rfc/rfc8299>.

[RFC8330] Long, H., Ye, M., Mirsky, G., D'Alessandro, A., and H. Shah, "OSPF Traffic Engineering (OSPF-TE) Link Availability Extension for Links with Variable Discrete Bandwidth", RFC 8330, DOI 10.17487/RFC8330, February 2018, <https://www.rfc-editor.org/rfc/rfc8330>.

[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/rfc/rfc8341>.

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

[RFC8466] Wen, B., Fioccola, G., Ed., Xie, C., and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", RFC 8466, DOI 10.17487/RFC8466, October 2018, <https://www.rfc-editor.org/rfc/rfc8466>.

[RFC8969] Wu, Q., Ed., Boucadair, M., Ed., Lopez, D., Xie, C., and L. Geng, "A Framework for Automating Service and Network Management with YANG", RFC 8969, DOI 10.17487/RFC8969, January 2021, <https://www.rfc-editor.org/rfc/rfc8969>.

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

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

Authors' Addresses

Samier Barguil Giraldo (editor)
Nokia
Email: samier.barguil_giraldo@nokia.com

Luis M. Contreras (editor)
Telefonica
Ronda de la Comunicacion, s/n
28050 Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com
URI: http://lmcontreras.com

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

Reza Rokui
Ciena

Email: reza.rokui@nokia.com

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

Daniel King
Old Dog Consulting
Email: daniel@olddog.co.uk