Network Working Group
Internet-Draft
Intended status: Informational
Expires: 25 April 2024

S. Barguil, Ed.
Nokia
L.M. Contreras, Ed.
Telefonica
V. Lopez
Nokia
O. Gonzalez de Dios
Telefonica
M. Boucadair
Orange
R. Rokui
Ciena

23 October 2023

Applicability of IETF-Defined Service and Network Data Models for RFC XXXX Network Slice Service Management
draft-barguil-teas-network-slices-instantation-08

Abstract

This document exemplifies how the various data modeles that are produced in the IETF can be combined in the context of RFC XXXX Network Slice Services delivery.

Specifically, this document describes the relationship between the RFC XXXX Network Slice Service models for requesting Network Slice Services and both Service (e.g., the Layer-3 Service Model, the Layer-2 Service Model) and Network (e.g., the Layer-3 Network Model, the Layer-2 Network Model) models used during their realizations. In addition, this document describes the communication between an RFC XXXX Network Slice Controller (NSC) and the network controllers for the realization of RFC XXXX Network Slices.

The RFC XXXX Network Slice Service YANG model provides a customer-oriented view of the intended Network slice Service. Thus, once an NSC receives a request for a Slice Service request, the NSC has to map it to accomplish the specific objectives expected by the network controllers. Existing YANG network models are analyzed against the RFC XXXX Network Slice requirements, and the gaps in existing models are identified.

Note to the RFC Editor: Please replace "RFC XXXX" with the RFC number assigned to I-D.ietf-teas-ietf-network-slices.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The IETF has produced several YANG data models that are instrumental for automating the provisioning and delivery of connectivity services. An overview of these data models and a framework that describes how these various modules can be glued together are described in [RFC8969].

This document adopts the rationale of [RFC8969], but with a focus on the Network Slice Service [I-D.ietf-teas-ietf-network-slices].

For example, the RFC XXXX Network Slice Service YANG service model provides a customer-oriented view of the Network Slice Service. Once an RFC XXXX Network Slice controller (NSC) receives a Slice Service request, it needs to map it into the underlying network capabilities to accomplish the intended service in a way understandable by the network controller.

Several service models and network models, including the Layer-3 Service Model (L3SM) [RFC8049], the Layer-2 Service Model (L2SM) [RFC8466], and network models (e.g., the Layer-3 Network Model (L3NM) [RFC9182], the Layer-2 Network Model (L2NM) [RFC9291])) which may be utilized for the realization of RFC XXXX Network Slice Services, are analyzed whether they can satisfy the RFC XXXX Network Slice requirements.

The document also identifies some gaps on existing models.

The document outlines an architecture and communication process between an NSC and other network controllers for managing RFC XXXX Network Slice Services, including creation and modification.

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.

This document assumes that the reader is familiar with the contents of [RFC6241], [RFC7950], [RFC8309], and [I-D.ietf-teas-ietf-network-slices] as it uses terms from those RFCs.

This document uses the term "network model" as defined in Section 2.1 of [RFC8969].

2. A Reference Architecture and Main Components

As described in [I-D.ietf-teas-ietf-network-slices], the RFC XXXX Network Slice Controller (NSC) is a functional entity for the control and management of RFC XXXX Network Slices Services. As shown in Figure 1, an NSC exposes set of APIs for higher level systems to request an RFC XXXX Network Slice Service. These APIs can be used to manage other connectivity services, such as managing the underlying delivery setup that that is required for the delivery of an RFC XXXX Network Slice Service. Such setup can be managed prior or during the process of a Network Service Slice. Concretely, the setup can be the management of bearers and attachment circuits that connect Service Demarcation Points (SDPs) to customer premises.

The NSC customer-facing interface is invoked by a customer for managing an RFC XXXX Network Slice Service (i.e., creation, modification, or deletion). Upon receiving a request via a customer-facing interface, an NSC assesses whether it can satisfy the request and then identifies the resources that are needed for realization of the RFC XXXX Network Slice Service. The network-facing interface is used to interact with one or more Network Controllers for the realization of the requested RFC XXXX Network Slice Service.

An NSC exposes a set of RFC XXXX service data models: mainly, the RFC XXXX Network Slice Service Interface [I-D.ietf-teas-ietf-network-slice-nbi-yang] or the Attachment Circuit-as-a-Service Interface [I-D.boro-opsawg-teas-attachment-circuit].

Network Controllers exposes to NSCs a set of network data models, such as the L3NM, the L2NM, or the Service Attachment Points (SAPs) [RFC9408]. Typically, by setting the service type to "network-slice", an NSC can retrieve via the SAPs where Slice Services can be delivered to customers. Likewise, SAPs can also be used to retrieve where such services are delivered to customers through the network configuration described in the L3NM [RFC9182] or the L2NM [RFC9291]. These checks can be used as part of request feasibility checks.

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

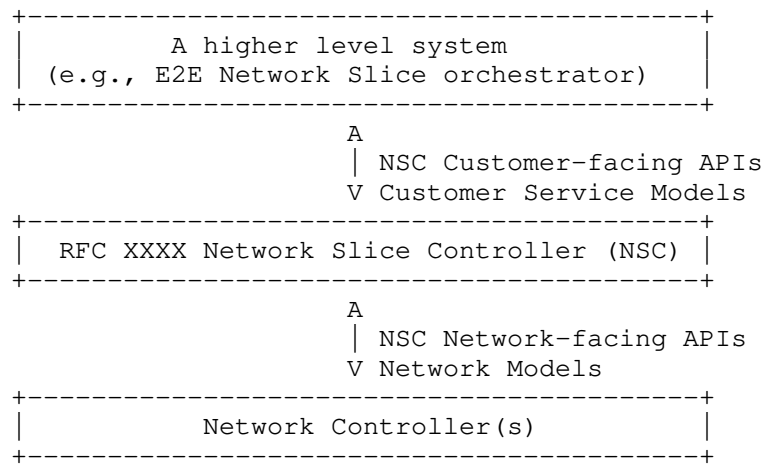

Figure 1: Network Slice Controllers as a Module of a Hierarchical SDN Controller.

Several architectural definitions have arisen on the IETF to support SDN and network slicing deployments. The architecture defined in [I-D.ietf-teas-ietf-network-slices] includes a three-level hierarchy.

depicts a possible architecture using similar concepts. It starts from a consumer or high-level operational systems. Then, the NSC function might be part of a hierarchical network controller (e.g., as the MDSC in the ACTN context [RFC8453]) as a modular function. As an alternative, in the Figure 2 at the bottom, multiple network controllers can be orechestraded from the NSC. Each Network controller can handle multiple or single underlay technologies.

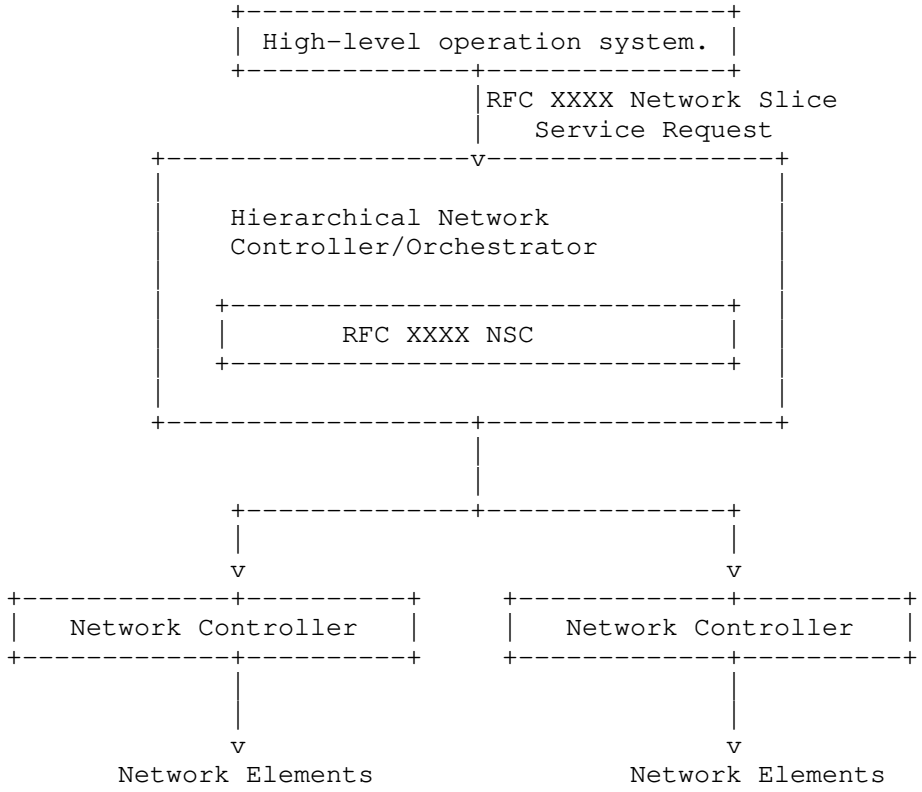

Figure 2: RFC XXXX Network Slice Controller as a Module of a Hierarchical SDN Controller.

In other implementations, an NSC can be a standalone component that directly interact with a network controller, as depicted in Figure 3. In this scenario, a service request follows a "data-enrichment" path, where each entity adds more information to the service request.

This document describes how existing service models and network models interact to deliver a Network Slice Service in a service provider environment.

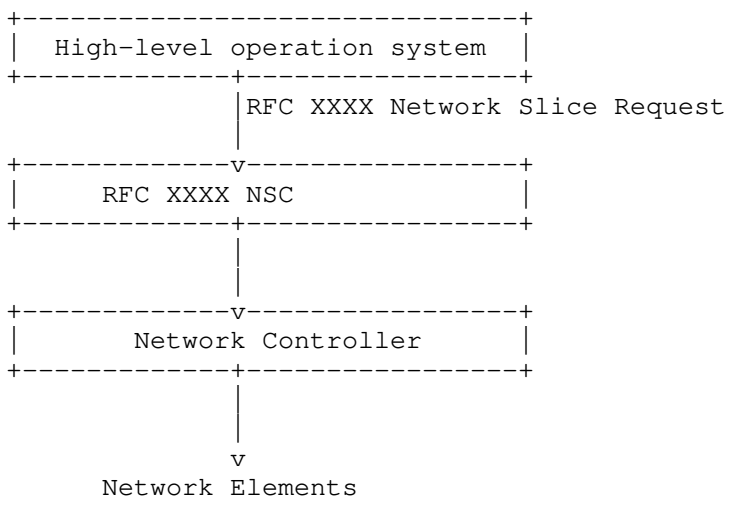

Figure 3: RFC XXXX Network Slice Controller as a Standalone Component

Alternatively, an NSC can be integrated with a network controller and directly realizes the Network Slice Services using device data models to configure the network devices. A sample architecture is depicted in Figure 4.

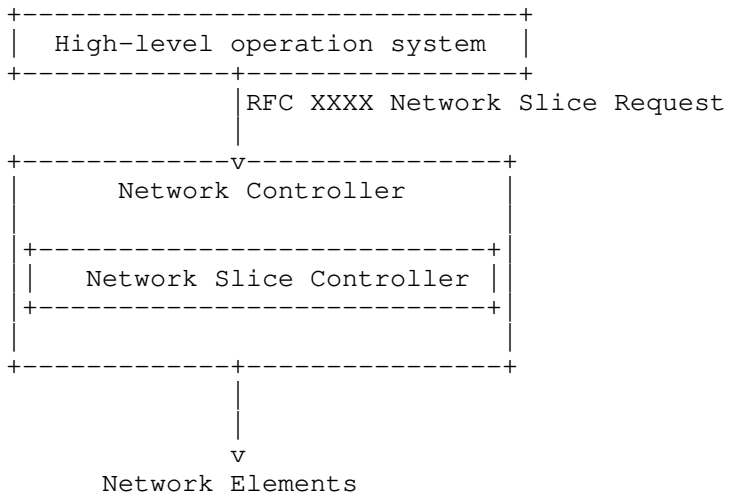

Figure 4: RFC XXXX Network Slice Controller as a Module of a Network Controller.

3. RFC XXXX Network Slice Requirements and Data Models

The main requirements for an RFC XXXX Network Slice Service, 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 RFC XXXX Network Slice Service.

* [RFC9181]: specifies a set of reusable types and groupings to manage VPN services; VPN is used to realize slices.

* [I-D.boro-opsawg-teas-common-ac]: specifies a set of reusable types and groupings to manage Attachment Circuits (ACs).

* [I-D.boro-opsawg-teas-attachment-circuit]: specifies YANG data models for managing 'Attachment Circuits'-as-a-Service (ACaaS) and also bearers. These ACs and bearers are used to identify where to deliver a Network Slice Service.

* [RFC9408]: defines a YANG data model for representing an abstract view of the provider network topology that contains the points from which its services can be attached (e.g., Network Slices). A SAP network topology can be used for one or multiple service types ('service-type'). Setting this data node to 'network-slice' allows a controller to expose where RFC XXXX Network Slices services are being delivered. It can also be used to check where RFC XXXX Network Slice services can be delivered.

* [I-D.boro-opsawg-ntw-attachment-circuit] augments the SAP model with more details for managing ACs at the network level.

* [I-D.dhody-teas-ietf-network-slice-mapping] specifies an RFC XXXX Network Slice Service mapping YANG model.

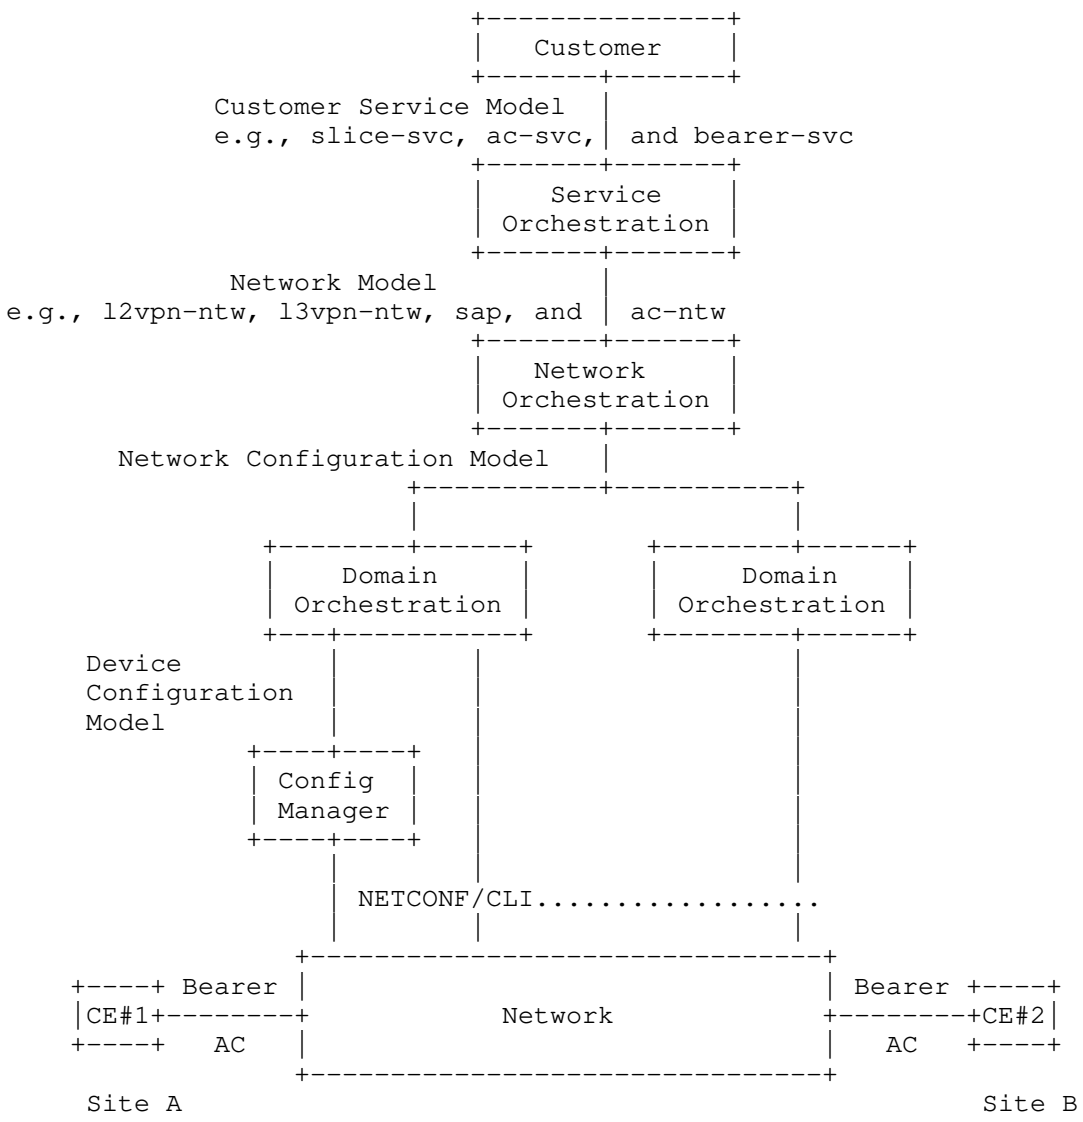

Figure 5: Overview of Data Models used for Network Slicing

4. Operational Considerations

This section outlines the compliance and operational aspects of Network Controller models with RFC XXXX Network Slice requirements. Section presented the requirements of the RFC XXXX 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.

4.1. Availability

As per [I-D.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 RFC 8330, 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 LxNMs ([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.ietf-teas-te-service-mapping-yang] allows to bind a VPN to a given LSP, which have its bandwidth requirements. Additionally, TE models can force a give bandwidth 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.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.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 RFC XXXX 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 RFC XXXX 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 RFC XXXX Network Slice, these parameters are part of the requirements of a TE tunnel configuration [I-D.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 RFC XXXX Network Slice, this attribute will influence the tunnel selection and the value is included in the [I-D.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 RFC XXXX Network Slice Service YANG Model Parameters and those in Lx Service and Network Models

5.1. Relationship Between RFC XXXX Network Slice Service Model Parameters and The L3SM and L2SM Parameters

This section presents an initial analysis of the relationship between the RFC XXXX Network Slice Service model parameters and the 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 referring 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, the L2SMs parameters are described in Sections 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).

* Broadcast, unknown-unicast and multicast (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 RFC XXXX Network Slice Service YANG module 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 individually (overriding 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 maximum Protocol data unit (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 the L3SM, L2SM, and RFC XXXX Network Slice Service model can be performed as indicated in the following table.

L3SM (RFC 8299)	L2SM (RFC 8466)	RFC XXXX NSS 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 (BUM)	The need of replication can be inferred from ns-connectivity-type. Further details are not available (e.g.

		source or receiver role)
	Availability as dual homing	Availability as the ratio of up-time to total_time(up-time+down-time)

{: #Table1 title='Mapping of RFC XXXX Network Slice Service and The LxSM Service Attribute' artwork-align="center"}

The following considerations can be made:

* While the QoS profile in the L3SM and the L2SM applies per service class, the parameters in the RFC XXXX Network Slice Service Interface apply per connection. So if per- class granularity is required in an RFC XXXX 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 the L3SM nor the L2SM, such as packet loss, isolation, or security. Then, the L3SM and L2SM could not be sufficient to realize RFC XXXX Network Slice Services with such specific needs, unless those other objectives and expectations are provided by other means (e.g., realizing the L3SM through technologies guaranteeing dedicated resource allocation such as OTN).

5.2. Relationship Between RFC XXXX Network Slice Service Model Parameters and the L3NM and L2NM Parameters

This section presents an initial analysis of the relationship between RFC XXXX Network Slice Service model parameters and the L3NM and the L2NM 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' and '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: The L3NM only supports the definition at the '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, the L2NM parameters are described in Section 7.6.6 of [RFC9291]:

* Bandwidth: The L2NM considers the same parameters 'pe-to-ce-bandwidth' and '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.

* MTU: The 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 the L3NM, L2NM, and RFC XXXX Network Slice Service model can be performed as indicated in the following table:

L3NM (RFC 9182)	L2NM (RFC 9291)	RFC XXXX NSC Service 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 classification policy (based on layer 3 and 4 info) - QoS profile (not defined)	QoS - QoS classification policy (based on layer 2 info) - QoS profile (not defined)	QoS Defined in the model as network-access-qos-policy-name to be applied per access-point 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
Multicast	Broadcast, Unknown, Unicast and Multicast (BUM)	The need of replication can be inferred from ns-connectivity-type. Further

		details are not available (e.g.
		source or receiver role)
N/A	N/A	Availability as the ratio of up-time to total_time(up-time+down-time)

{: #Table2 title='Mapping of RFC XXXX Network Slice Service and The LxNM Service Attribute' artwork-align="center"}

6. RFC XXXX Network Slice Procedure

6.1. RFC XXXX Network Slice Provisioning Workflow

An RFC XXXX Network Slice may use several underlying technologies. A new RFC XXXX Network Slice may be initiated following these steps:

1. A higher level system requests services with specific characteristics via the customer-facing APIs
2. This request is processed by an NSC which specifies a mapping between service request to any services, tunnels, and paths models.
3. A series of requests for creation of services, tunnels and paths will be sent to the network to realize the transport slice.

Variations of this flow can be considered: - The customer requests bearers and attachment circuits, independent of any service that will be delivered over them. - The customer place a service-specific request with references to ACes. - The curstomer may update the bearers/AC/service delivery points during the lifetime of a service.

As a functional entity responsible for managing a network domain, a network controller can expose a set of YANG models to an NSC. An NSC can invokes these models during the realization of an RFC XXXX Network Slice Service. The following network models can be used for realization of RFC XXXX 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, its exposure in the Network Slice Service 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 Network 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 so called 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

6.3. Traffic Engineering Models

TEAS 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.

6.4. 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 Models Usage in Alternative RFC XXXX 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 RFC XXXX Network Slice Service.

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

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

* RFC XXXX 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 RFC XXXX Network Slice and triggers its realization.

Note that this division in functional components of an RFC XXXX NSC is provided as an implementation option, not constraining any other implementation of functional structure.

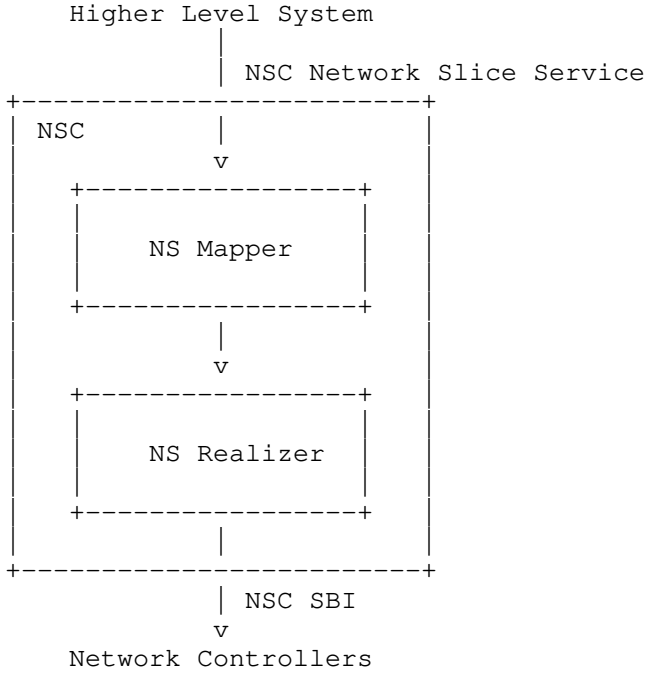

Figure 8: RFC XXXX Network Slice Controller Structure

The details of RFC XXXX Network Slice mapper and realize are provided below for various implementation of NCS.

7.1. RFC XXXX Network Slice Service Requested to Hierarchical Network Controller

Referring to Figure 1 in an integrated architecture, an NCS is part of a Hierarchical SDN controller module, the NSC's and the Hierarchical Network Controller should share the same internal data and the same Network Slice Service interface. Thus, the H-SDN module must be able to:

* Map: The NSC should process the customer request received through [I-D.ietf-teas-ietf-network-slice-nbi-yang]. The mapping process takes the network-slice SLOs selected by the customer selecting available Routing Policies and Forwarding policies for accomplishing those SLOs.

* 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 NSC 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 RFC XXXX Network Slice ID ns-id used of the [I-D.ietf-teas-ietf-network-slice-nbi-yang] must be directly mapped to the transport-instance-id at the VPN-Node level.

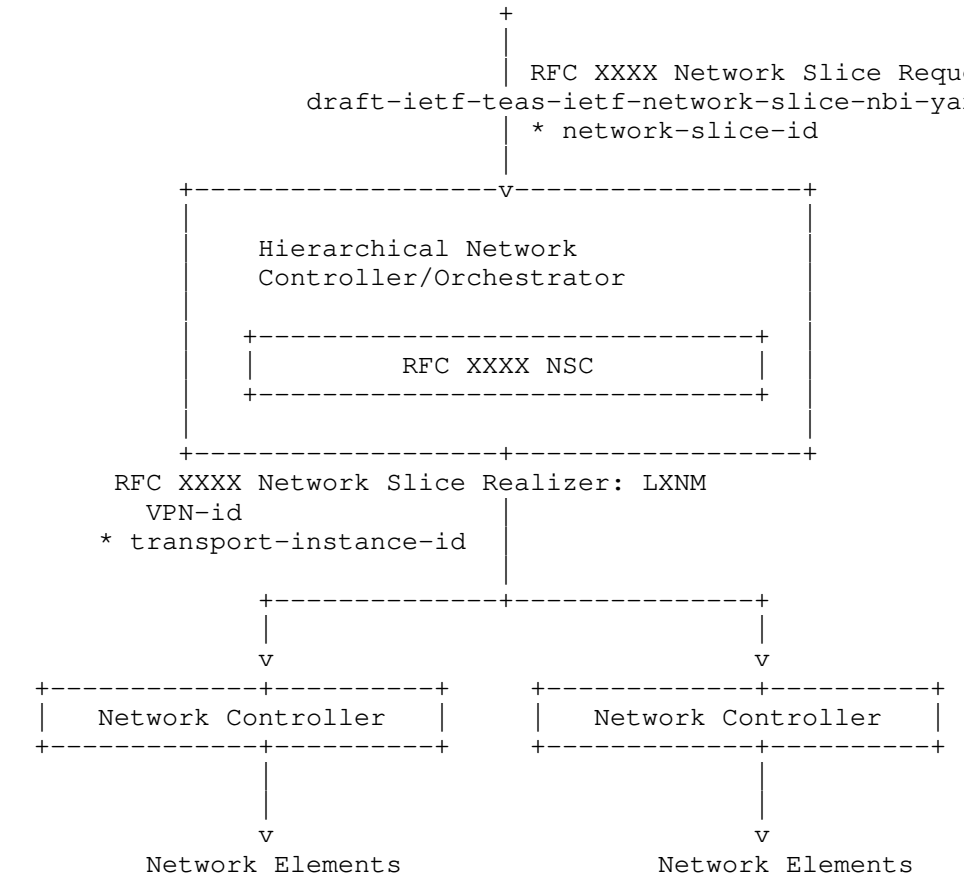

Figure 9 Workflow for the Slice Request in an Integrated Architecture.

7.2. RFC XXXX Network Slice Service Requested to Network Slice Controller

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

* Map: Process the customer request. The customer request can be sent using [I-D.ietf-teas-ietf-network-slice-nbi-yang]. The customer can also perform the Network Slice request using customized topologies.

* 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.

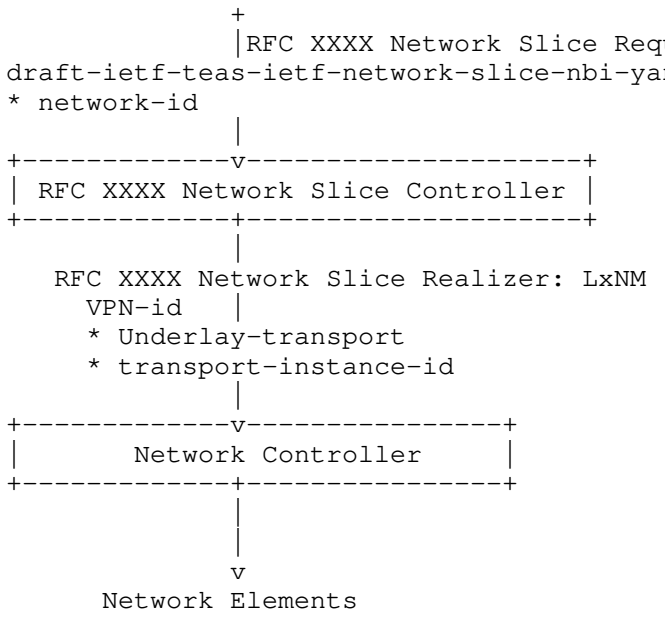

Figure 10 Workflow for the slice request in an stand-alone architecture.

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 customer-facing and network-facing interfaces of the network controller. The direct translation would reduce the service logic implemented at the network controller level, grouping the mapping and translation into a single task:

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

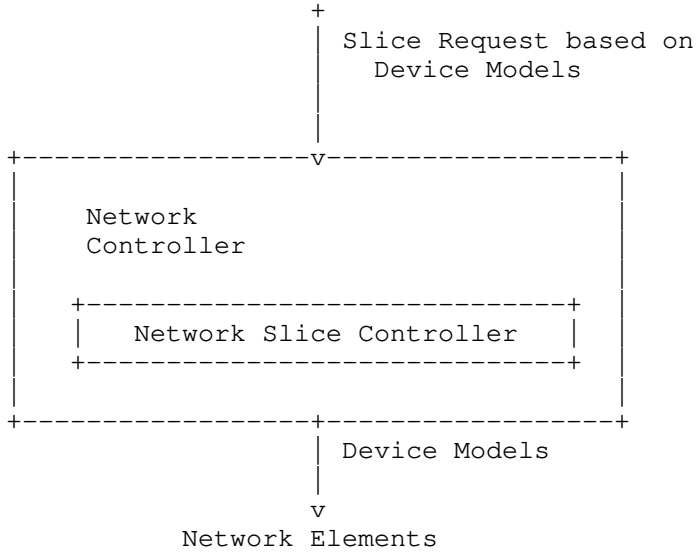

Figure 11 Workflow for the slice request in an stand-alone architecture.

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.

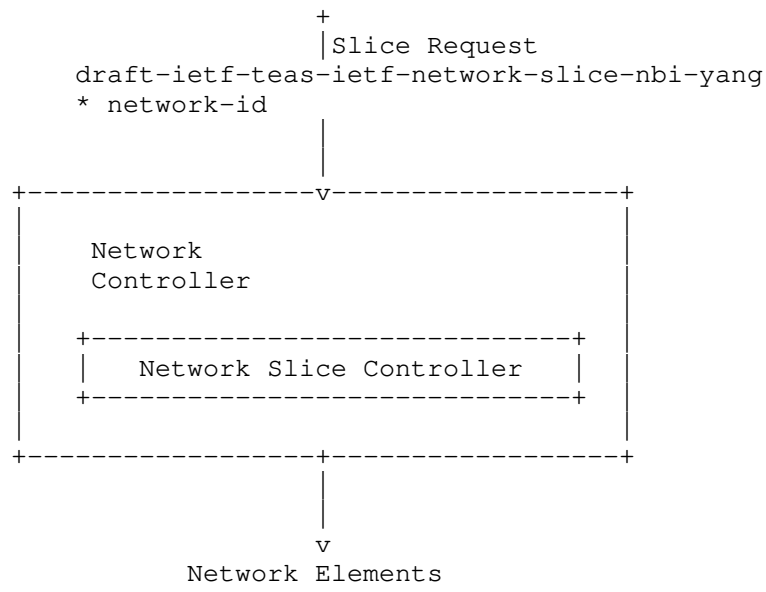

Figure 12 Workflow for the slice request in an stand-alone architecture.

8. Security Considerations

There are two main aspects to consider. On the one hand, the RFC XXXX 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 RFC XXXX Network Slice Controller 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 RFC XXXX network Slices:

* Malicious clients attempting to delete or modify VPN services that implements an RFC XXXX 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 RFC XXXX Network Slice service.
* Unauthorized clients attempting to read VPN services related information hat implements an RFC XXXX 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 RFC XXXX Network Slices. Some of the RFC XXXX Network 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 RFC XXXX Network Slices.

Contributors

Many thanks to Daniel King for their perspectives on the Series and their ongoing support.

Normative References

[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 Network Slices in Networks Built from IETF Technologies", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-25, 14 September 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-25>.

[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-14, 12 September 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-service-mapping-yang-14>.

[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-34, 1 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-34>.

[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>.

[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>.

[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>.

[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/info/rfc8299>.

[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>.

[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/info/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/info/rfc8969>.

[RFC9181] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., and Q. Wu, "A Common YANG Data Model for Layer 2 and Layer 3 VPNs", RFC 9181, DOI 10.17487/RFC9181, February 2022, <https://www.rfc-editor.org/info/rfc9181>.

[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/info/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/info/rfc9291>.

[RFC9408] Boucadair, M., Ed., Gonzalez de Dios, O., Barguil, S., Wu, Q., and V. Lopez, "A YANG Network Data Model for Service Attachment Points (SAPs)", RFC 9408, DOI 10.17487/RFC9408, June 2023, <https://www.rfc-editor.org/info/rfc9408>.

Informative References

[I-D.boro-opsawg-ntw-attachment-circuit]
Boucadair, M., Roberts, R., de Dios, O. G., Barguil, S., and B. Wu, "A Network YANG Data Model for Attachment Circuits", Work in Progress, Internet-Draft, draft-boro-opsawg-ntw-attachment-circuit-03, 5 September 2023, <https://datatracker.ietf.org/doc/html/draft-boro-opsawg-ntw-attachment-circuit-03>.

[I-D.boro-opsawg-teas-attachment-circuit]
Boucadair, M., Roberts, R., de Dios, O. G., Barguil, S., and B. Wu, "YANG Data Models for 'Attachment Circuits'-as-a-Service (ACaaS)", Work in Progress, Internet-Draft, draft-boro-opsawg-teas-attachment-circuit-07, 10 July 2023, <https://datatracker.ietf.org/doc/html/draft-boro-opsawg-teas-attachment-circuit-07>.

[I-D.boro-opsawg-teas-common-ac]
Boucadair, M., Roberts, R., de Dios, O. G., Barguil, S., and B. Wu, "A Common YANG Data Model for Attachment Circuits", Work in Progress, Internet-Draft, draft-boro-opsawg-teas-common-ac-02, 3 May 2023, <https://datatracker.ietf.org/doc/html/draft-boro-opsawg-teas-common-ac-02>.

[I-D.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-05, 13 March 2023, <https://datatracker.ietf.org/doc/html/draft-contreras-teas-slice-controller-models-05>.

[I-D.dhody-teas-ietf-network-slice-mapping]
Dhody, D. and B. Wu, "IETF Network Slice Service Mapping YANG Model", Work in Progress, Internet-Draft, draft-dhody-teas-ietf-network-slice-mapping-04, 12 September 2023, <https://datatracker.ietf.org/doc/html/draft-dhody-teas-ietf-network-slice-mapping-04>.

[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-07, 20 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slice-nbi-yang-07>.

[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/info/rfc8049>.

[RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <https://www.rfc-editor.org/info/rfc8309>.

[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>.

Authors' Addresses

Samier Barguil (editor)
Nokia
Calle de María Tubau, 9
28050 Madrid
Spain
Email: samier.barguil_giraldo@nokia.com

Luis Miguel Contreras (editor)
Telefonica
Distrito T
28050 Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com

Victor Lopez
Nokia
Calle de María Tubau, 9
28050 Madrid
Spain
Email: victor.lopez@nokia.com

Oscar Gonzalez de Dios
Telefonica
Distrito T
28050 Madrid
Spain
Email: oscar.gonzalezdedios@telefonica.com

Mohamed Boucadair
Orange
Email: mohamed.boucadair@orange.com

Reza Rokui
Ciena
Email: reza.rokui@nokia.com