

TEAS
Internet-Draft
Intended status: Informational
Expires: 8 January 2026

L.M. Contreras, Ed.
Telefonica
R. Rokui
Ciena
J. Tantsura
Nvidia
B. Wu
Huawei
X. Liu
Alef Edge
July 2025

IETF Network Slice Controller and its Associated Data Models
draft-ietf-teas-ns-controller-models-05

Abstract

This document describes an approach for structuring the IETF Network Slice Controller as well as how to use different data models being defined for IETF Network Slice Service provision (and how they are related). It is not the purpose of this document to standardize or constrain the implementation the IETF Network Slice Controller.

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 2 January 2026.

Copyright Notice

Copyright (c) 2025 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	2
2. IETF Network Slice Data Models	3
3. Structure of the IETF Network Slice Controller (NSC)	4
3.1. NS Mapper	7
3.2. NS Realizer	8
4. Model Types in IETF Network Slice Controller Interfaces	8
5. Security Considerations	9
6. IANA Considerations	9
7. References	9
Acknowledgments	11
Contributors	11
Authors' Addresses	11

1. Introduction

A main promise of network slicing is to provide tailored end-to-end network capabilities to customers in the way that they could be perceived as a dedicated infrastructure, despite that it makes use of shared physical infrastructure facilities.

Particularly, the connectivity within and between different segments of a network slice with specific performance characteristics are key in characterizing a slice. Thus, the IETF Network Slice, realized by any of the IETF technologies, emerges as complementary but essential part of an end-to-end network slice.

In order to facilitate the service exposure, service order handling, realization, and lifecycle control and management of a transport slice, a dedicated element called IETF Network Slice Controller (NSC) is proposed in [RFC9543].

The NSC from its customer-facing interface, i.e., the IETF Network Slice Service interface, exposes a set of APIs that allow a third party system to request a transport slice. The NSC receives slice service requests from customers to manage an IETF Network Slice (i.e., creation, modification or deletion). Upon receipt of a request to create a slice, the NSC assess and then identifies the

resources needed for realization of the IETF Network Slice. To that end, the NSC may interact with one or more Network Controllers for the realization of the requested IETF Network Slice request and the management of its lifecycle. Figure 1 presents a high-level view of the IETF NSC [RFC9543].




Figure 1: Interface of Transport Slice Controller

This document describes the characteristics of the NSC as well as a detailed structure of the NSC and its major components. In addition, it describes the characteristics of the data models to identify an IETF Network Slice and its realization. Then the referred data models are mapped to the interfaces among components.

This document describes a potential way of structuring the IETF Network Slice Controller as well as how to use different data models being defined for IETF Network Slice Service provision (and how they are related). It is not the purpose of this document to standardize or constrain the implementation the IETF Network Slice Controllers.

2. IETF Network Slice Data Models

At the time of provisioning and operating IETF Network Slices different views can be identified as necessary:

- * The customer' s view. It is focused on the individual IETF Network Slice request process, reflecting the needs of each particular customer, including SLOs and other characteristics of the slice relevant for it. This view is technology-agnostic and describes

the characteristics of the IETF Network Slice from a customer' s point of view. It can include the customer slice topology intent, performance parameters, endpoints of the slice, traffic characteristics of the slice, and the KPIs to monitor the slice.

- * The provider' s view. In addition to the view that is exposed to customers, the provider maintains an more network-centric view that focuses on the provisioning and operation of IETF Network Slices in the underlay network, considering how a particular IETF Network Slice interplays with other IETF Network Slices maintained by the provider on a shared infrastructure. In other words, the provider' s view shows how an IETF Network Slice is implemented in the operator' s network along with all the resources used during the its realization. This view is not exposed to the customers.

Both views are complementary as they are invoked in different stages of service provisioning and delivery lifecycles. For the sake of automated procedures, some consistency should be ensured between these views to ease the service mapping as per [RFC8969].

It should be noted that for the realization of an IETF Network Slice, the NSC interacts with one or more Network Controllers underneath. Whether one or more NSCs/Network Controllers are used is deployment specific. The data models to be used are specific for each Network Controller (e.g., technology-dependent), as well as the mapping function from its customer-facing interface (i.e., IETF Network Slice Service interface) to network-facing interfaces (i.e., Network Configuration Interface) and the details of this mapping function are both out of the scope of this document.

3. Structure of the IETF Network Slice Controller (NSC)

The NSC should support both service and network data models. The NSC exposes service models to customers. Customers use those models for their slice service request placements. The NSC then process customers requests taking into account local policies and guidelines (e.g., mapping strategy 1:1/1:M/N:M), the overall view of the network resources (e.g., service functions) and the IETF Network Slices already instantiated. Finally, the NSC normalizes the slice instantiation across different technologies, and maps such slice to the provider view.

Once a new request is processed and tagged as feasible, an NSC triggers its realization by interacting with the relevant Network Controllers underneath and reporting to the higher level controller for accounting/billing purposes. The actual start of the billing process is deployment specific and depends on whether a slice request is a scheduled request or has immediate effect.

In order to accommodate these procedures, an NSC may be structured to embed the following components:

- * IETF Network Slice Service Mapper: this high-level component processes the customer requests, 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.

The IETF Network Slice Mapper and Realizer are considered to be internal modules of the IETF Network Slice Controller. However, anything prevents that these modules could be separated components, communicating through standard protocols. The intention of this document is to figure out how different models interplay in the transition from the technology-agnostic IETF Network Slice request up to the technology-specific IETF Network Slice realization. Whatever implementation guideline is out of scope of this document.

Figure 2 illustrates the components described and the associated models, as follows

- * (a) -> customer' s view, e.g.
[I-D.ietf-teas-ietf-network-slice-nbi-yang], which can be complemented by [I-D.ietf-opsawg-teas-attachment-circuit] and / or [I-D.ietf-teas-network-slice-topology-yang].
- * (b) -> provider' s view, including more detailed but yet technology-agnostic resource view as e.g.
[I-D.ietf-teas-network-slice-topology-yang], and/or alternative technology-specific augmentations as e.g. for OTN [I-D.ietf-ccamp-yang-otn-slicing] or for IP/MPLS NRP [I-D.ietf-teas-nrp-yang]. Note that the provider view could permit network operators to retrieve information about the slices being provided and how they are realized.
- * (c) -> models per network controller, out of scope of this document. An example of applicability of existing models is in [I-D.ietf-teas-ns-models-applicability].




Figure 2: IETF Network Slice Controller Structure and associated Data Models

IETF Network Slices with different level of detail could be requested:

- * The IETF network slice can be abstracted as a set of edge-to-edge links (Type 1).
- * The IETF network slice can be abstracted as a topology of virtual nodes and virtual links (Type 2) which represent the partitioning of underlay network resources for use by network slice connectivity.

The use cases of these two types of networks are further described by [RFC8453].

Regarding IETF Network Slice service requests, it is possible to model the Type 1 service by means of [I-D.ietf-teas-ietf-network-slice-nbi-yang], while it is possible to model the Type 2 service using [I-D.ietf-teas-network-slice-topology-yang]. Moreover, when a

customer intends to request a Type 2 service, [I-D.ietf-teas-network-slice-topology-yang] can also be used at the point (a) in Figure 2 for expressing intent-based topologies for resource reservation or realization intentions within the provider's network. It should be noted that according to [RFC9543], the customer might ask for some level of control of the IETF Network Slice, for instance to customize the service paths in a network slice. The abstract topology defined in [I-D.ietf-teas-network-slice-topology-yang] could serve to enable this capability and optimize the resource utilization for network slice connections activated on top of the abstract topology.

In respect to IETF Network Slice realization, as an example, when ACTN is used to realize an IETF network slice, model mappings are described in more details in [I-D.ietf-teas-actn-yang].

3.1. NS Mapper

The Mapper will receive an IETF Network Slice Service request from the customer. It will process it obtaining an overall view of how this new request complements or fits with the rest of IETF Network Slices, if any, as provisioned in the network. As part of that processing, a single customer IETF Network Slice Service request could result in the need of actually provisioning different IETF Network Slices in the network. The Mapper will maintain the relationship among customer IETF Network Slice request and provisioned IETF Network Slices. The Mapper also will provide performance notifications in relation with the SLOs dictated in the slice request by the customer.

The Mapper performs resource partitions of the filtered topologies provided by the Realizer component, generating specific Network Resource Partitions (NRPs). An NRP represents a collection of resources such as buffers or queues of the links of a filtered topology. The Mapper, when processing the slice request, will map the connectivity constructs to one or more NRPs, e.g., according to specific SLOs.

As part of the performance monitoring of the IETF Network Slice service, the Mapper will aggregate performance information from the distinct NRPs used for mapping the connectivity constructs forming the slice.

3.2. NS Realizer

The Realizer will receive from the Mapper one or more requests for provision of IETF Network Slices, potentially including some technology-specific information (e.g., an indication about the use of Layer 2 or Layer 3 capabilities to put into effect a slice). With that information, the Realizer will determine the realization of each IETF Network Slice Service interacting with technology-specific Network Controllers.

The Realizer will be in charge of generating filtered topologies from the underlying (physical) network information provided by the Network Controllers. The handling of filtered topologies is optional, then if not filtering is applied, the Realizer could expose the physical network. The filtered topologies represent a selection of nodes and links from the underlying network(s), e.g., as result of applying certain policies.

The Realizer will provide the telemetry information from the filtered topologies to the Mapper for further processing in support of the performance assurance of the IETF Network Slices.

4. Model Types in IETF Network Slice Controller Interfaces

Both [RFC8309] and [RFC8969] offer a complete view of customer, service and network model types. In this sense a potential mapping of models to IETF Network Slice Controller interfaces is as follows:

- * IETF Network Slice Service interface (interface (a) in Figure 2) -> Customer service model. According to [RFC8309] "a customer's service request is (or should be) technology agnostic. That is, a customer is unaware of the technology that the network operator has available to deliver the service, so the customer does not make requests specific to the underlying technology but is limited to making requests specific to the service that is to be delivered". This definition matches the expected behavior of the IETF NSC Slice Service interface as considered in [RFC9543].
- * Interface between NS Mapper and NS Realizer (interface (b) in Figure 2) -> Service Delivery model. According to [RFC8309] "a service delivery module is expressed as a core set of parameters that are common across a network type and technology [...] Service delivery modules include technology-specific modules." Furthermore, [RFC8969] (in its Figures 3 and 5) considers L3SM or VN Service models to be later on fed into a controller.

- * Network Configuration interface (interface (c) in Figure 2) -> Network Configuration model. According to [RFC8309] “the orchestrator must map the service request to its view, and this mapping may include a choice of which networks and technologies to use depending on which service features have been requested”. This is coincident with the expected behavior of the IETF NSC network configuration as considered in in [RFC9543].

5. Security Considerations

This document considers both the Mapper and the Realizer component as internal modules of the IETF Network Slice Controller. However, anything prevents that these modules could be separated components, communicating through standard protocols (i.e., not as an internal communication to the IETF NSC).

In that case, some security requirements apply such as:

- * Authentication between Mapper and Realizer, to prevent malicious behaviors.
- * Privacy of the information shared between components.
- * Secure transport between components based on the kind of interface used in the communication (e.g., NETCONF, RESTCONF, etc).

6. IANA Considerations

This draft does not include any IANA considerations

7. References

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

Guo, A., Contreras, L. M., Belotti, S., Rokui, R., Xu, Y., Zhao, Y., and X. Liu, "Framework and Data Model for OTN Network Slicing", Work in Progress, Internet-Draft, draft-ietf-ccamp-yang-otn-slicing-09, 3 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-yang-otn-slicing-09>>.

[I-D.ietf-opsawg-teas-attachment-circuit]

Boucadair, M., Roberts, R., de Dios, O. G., Barguil, S., and B. Wu, "YANG Data Models for Bearers and 'Attachment Circuits'-as-a-Service (ACaaS)", Work in Progress, Internet-Draft, draft-ietf-opsawg-teas-attachment-circuit-20, 23 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-teas-attachment-circuit-20>>.

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

Lee, Y., Zheng, H., Ceccarelli, D., Yoon, B. Y., and S. Belotti, "Applicability of YANG models for Abstraction and Control of Traffic Engineered Networks", Work in Progress, Internet-Draft, draft-ietf-teas-actn-yang-11, 7 March 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-yang-11>>.

[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-01, 3 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-network-slice-topology-yang-01>>.

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

Wu, B., Dhody, D., Beeram, V. P., Saad, T., and S. Peng, "YANG Data Models for Network Resource Partitions (NRPs)", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-yang-03, 2 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-yang-03>>.

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

Barguil, S., Contreras, L. M., Lopez, V., de Dios, O. G., and M. Boucadair, "Applicability of IETF-Defined Service and Network Data Models for Network Slice Service Management", Work in Progress, Internet-Draft, draft-ietf-teas-ns-models-applicability-01, 24 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-models-applicability-01>>.

[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>>.
- [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>>.
- [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

The authors would like to thank (in alphabetical order) to Swamynathan B, Adrian Farrel, Aihua Guo, Joel Halpern and Kiran Makhijani for their valuable comments received.

Contributors

The following people (in alphabetical order) contributed substantially to the content of this document and should be considered coauthors.

Sergio Belotti, Nokia (sergio.belotti@nokia.com)

Med Boucadair, Orange (mohamed.boucadair@orange.com)

Dhruv Dhody, Huawei Technologies (dhruv.ietf@gmail.com)

Authors' Addresses

Luis M. Contreras (editor)
Telefonica
Ronda de la Comunicacion, s/n
Sur-3 building, 3rd floor
28050 Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com
URI: <http://lmcontreras.com/>

Reza Rokui
Ciena
Canada
Email: rrokui@ciena.com

Jeff Tantsura
Nvidia
United States of America
Email: jefftant.ietf@gmail.com

Bo Wu
Huawei Technologies
101 Software Avenue, Yuhua District
Nanjing, Jiangsu
210012
China
Email: lana.wubo@huawei.com

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