TEAS   LM. Contreras
Internet-Draft   Telefonica
Intended status: Informational   R. Rokui
Expires: 27 April 2023   Ciena
                                                          J.T. Tantsura
                                                             Microsoft
                                                                 B. Wu
                                                                Huawei
                                                                X. Liu
                                                       IBM Corporation
                                                              D. Dhody
                                                                Huawei
                                                            S. Belloti
                                                                 Nokia
                                                          October 2022

IETF Network Slice Controller and its associated data models
draft-contreras-teas-slice-controller-models-04

Abstract

   This document describes the major functional components of an IETF
   Network Slice Controller (NSC) as well as references the data models
   required for supporting the requests of IETF network slice services
   and their realization.

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 4 April 2023.

Table of Contents

1. Introduction

The generic idea of network slicing intends to provide tailored end-
to-end network capabilities to customers in the way that they could
be perceived as a dedicated network, despite the fact that it makes
use of shared physical infrastructure facilities.

Among the capabilities mentioned, connectivity of different parts of
a network slice with particular performance characteristics play a
central role. Thus, the concept of 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 request, realization and lifecycle control
and management of a transport slice, a new element named IETF Network
Slice Controller (NSC) is being proposed in
[I-D.ietf-teas-ietf-network-slices].

The NSC from its North Bound Interface (NBI), i.e. the IETF Network
Slice Service interface, exposes set of APIs that allow a higher
level system to request an end-to-end transport slice. It receives
the request of enablement of an IETF Network Slice by a customer
(i.e. creation, modification or deletion). Upon receiving a request
from its NBI, the NSC finds the resources needed for realization of
the IETF Network Slice and in turn interfaces from its South Bound
Interface (SBI) 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 [I-D.ietf-teas-ietf-network-slices].

```
            +------------------------------------------+
            | Customer higher level operation system   |
            |    (e.g E2E network slice orchestrator,   |
            |       customer network management system) |
            +------------------------------------------+
                            A
                            | IETF Network Slice Service Interface
                            V
            +------------------------------------------+
            |     IETF Network Slice Controller (NSC)   |
            +------------------------------------------+
                            A
                            | Network Configuration Interface
                            V
            +------------------------------------------+
            |              Network Controllers          |
            +------------------------------------------+
```
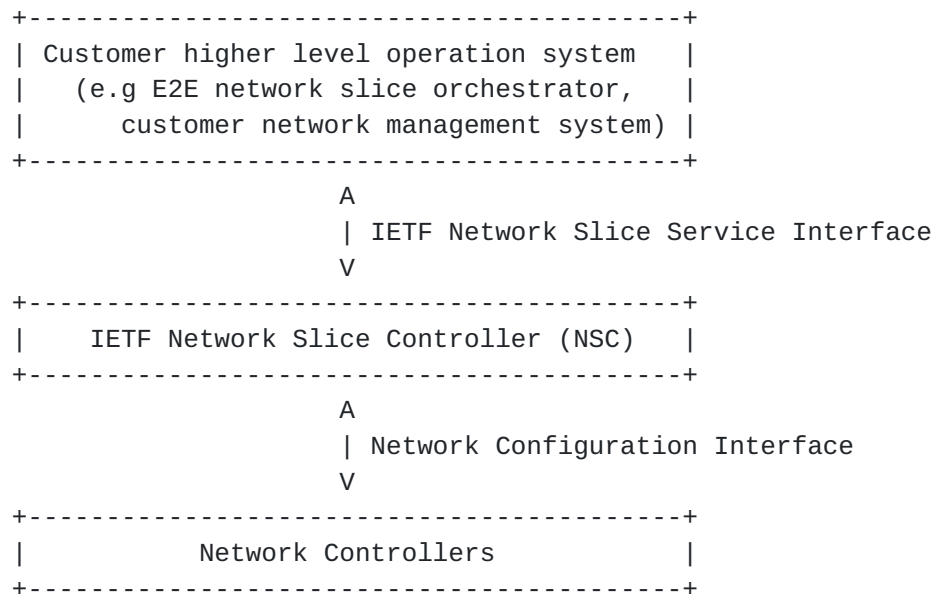
Figure 1: Interface of Transport Slice Controller

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 the
IETF Network Slice and its realization. Then the referred data
models are mapped to the interfaces among components.

2. IETF Network Slice data models

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

* Customer's view, mostly 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 slice topology, performance
   parameters, endpoints of the slice, traffic characteristics of the
   slice, and the KPIs to monitor the slice.

* Provider's view, mostly focused 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, operator's view shows how an IETF Network Slice is realized in operator's network along with all the resources used during the its realization.

Both views are complementary, each of them specialized for a given purpose. In consequence, it should be consistency between both in order to ensure alignment.

Currently there are two different models proposed, one for each of the categories above. The model in [I-D.ietf-teas-ietf-network-slice-nbi-yang] fits into the customer view, while the model defined in [I-D.liu-teas-transport-network-slice-yang] fits in to the provider view.

It should be noted that for the realization of an IETF Network Slice, the NSC interacts with one or more Network Controllers underneath. In that case, the data models to be used are particular for each Network Controller (e.g., technology dependent), as well as the mapping function from its NBI (i.e., IETF Network Slice Service interface) to SBI (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 work with both data models. The NSC takes first the customer's view by analyzing the needs of the customer, processing such requests taking into account the overall view of the network and the IETF Network Slices already instantiated, normalizing its instantiation across different technologies, and finally generates the provider view.

Once the new request is processed and declared as feasible, the NSC triggers its realization by interacting with the Network Controllers underneath and communicates back to the higher level controller to start the billing cycle.

In order to accommodate these procedures, the internal structure of the NSC can be divided into:

*   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.

   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].

   * (b) -> provider's view, including more detailed but yet
      technology-agnostic resource view as e.g.
      [I-D.liu-teas-transport-network-slice-yang], and/or alternative
      technology-specific augmentations as e.g.
      [I-D.ietf-ccamp-yang-otn-slicing].

   * (c) -> models per network controller, out of scope of this
      document. An example of applicability of existing models is in
      [I-D.barguil-teas-network-slices-instantation].

```
                         Higher Level System
                                 |
                                 |
                  ---------------------------
                 | NSC          | (a)        |
                 |              v            |
                 |   ------------------      |
                 |   |                | |    |
                 |   |   NS Mapper    | |    |
                 |   |                | |    |
                 |   ------------------      |
                 |            | (b)          |
                 |            v              |
                 |   ------------------      |
                 |   |                | |    |
                 |   |   NS Realizer  | |    |
                 |   |                | |    |
                 |   ------------------      |
                 |            | (c)          |
                  ---------------------------
                                 |
                                 v
                        Network Controllers
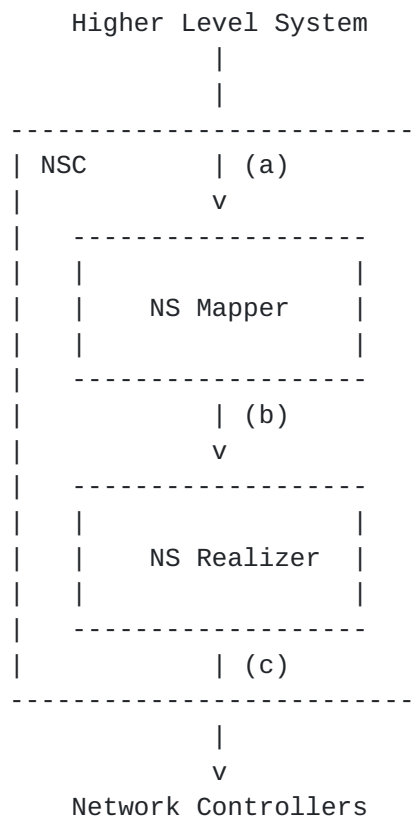```

           Figure 2: IETF Network Slice Controller structure and asspociated
                              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].  [I-D.ietf-teas-ietf-network-slice-nbi-yang] models the
Type 1 service, while [I-D.liu-teas-transport-network-slice-yang]
models the Type 2 service.  When a customer intends to request a Type
2 service, [I-D.liu-teas-transport-network-slice-yang] can also be
used at the point (a) in Figure 2.  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 the IETF Network Slice 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 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, queues, etc, 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.  With that information, the Realizer will determine the realization of each particular IETF Network Slice 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 Slcie Controller interfaces is as follows:

*  NBI of the IETF NSC, i.e. 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 NBI as considered in [I-D.ietf-teas-ietf-network-slices].

*  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.

*  SBI of the IETF NSC, i.e. 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 coincideent with the
        expected behavior of the IETF NSC SBI as considered in in
        [I-D.ietf-teas-ietf-network-slices].

5.  Security Considerations

        This draft 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.barguil-teas-network-slices-instantation]
                    Barguil, S., Contreras, L. M., Lopez, V., Rokui, R., and
                    O. G. D. Dios, "Instantiation of IETF Network Slices in
                    Service Providers Networks", Work in Progress, Internet-
                    Draft, draft-barguil-teas-network-slices-instantation-05,
                    24 October 2022,
                    <https://datatracker.ietf.org/api/v1/doc/document/draft-
                    barguil-teas-network-slices-instantation/>.

        [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-03, 22 October 2022,
                    <https://datatracker.ietf.org/api/v1/doc/document/draft-
                    ietf-ccamp-yang-otn-slicing/>.

[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-10, 4 September
          2022, <https://www.ietf.org/archive/id/draft-ietf-teas-
          actn-yang-10.txt>.

[I-D.ietf-teas-ietf-network-slice-nbi-yang]
          Wu, B., Dhody, D., Rokui, R., Saad, T., and L. Han, "IETF
          Network Slice Service YANG Model", Work in Progress,
          Internet-Draft, draft-ietf-teas-ietf-network-slice-nbi-
          yang-02, 11 July 2022, <https://www.ietf.org/archive/id/
          draft-ietf-teas-ietf-network-slice-nbi-yang-02.txt>.

[I-D.ietf-teas-ietf-network-slices]
          Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani,
          K., Contreras, L. M., and J. Tantsura, "Framework for IETF
          Network Slices", Work in Progress, Internet-Draft, draft-
          ietf-teas-ietf-network-slices-14, 3 August 2022,
          <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-
          network-slices-14.txt>.

[I-D.liu-teas-transport-network-slice-yang]
          Liu, X., Tantsura, J., Bryskin, I., Contreras, L. M., Wu,
          Q., Belotti, S., and R. Rokui, "IETF Network Slice YANG
          Data Model", Work in Progress, Internet-Draft, draft-liu-
          teas-transport-network-slice-yang-05, 6 March 2022,
          <https://www.ietf.org/archive/id/draft-liu-teas-transport-
          network-slice-yang-05.txt>.

[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>.

Authors' Addresses

Luis M. Contreras
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
Microsoft
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 Liu
IBM Corporation
Email: xufeng.liu.ietf@gmail.com


Dhruv Dhody
Huawei Technologies
Divyashree Techno Park
Bangalore, Karnataka 560066
India
Email: dhruv.ietf@gmail.com


Sergio Belloti
Nokia
Email: sergio.belotti@nokia.com