TEAS Working Group LM. Contreras
Internet-Draft Telefonica
Intended status: Informational S. Homma
Expires: May 3, 2021 NTT
J. Ordonez-Lucena
Telefonica
October 30, 2020

IETF Network Slice use cases and attributes for Northbound Interface of controller

draft-contreras-teas-slice-nbi-03

Abstract

The transport network is an essential component in the end-to-end delivery of services and, consequently, with the advent of network slicing it is necessary to understand what could be the way in which the transport network is consumed as a slice. This document analyses the needs of potential IETF network slice customers (i.e., use cases) in order to identify the functionality required on the North Bound Interface (NBI) of a IETF network slice controller for satisfying such IETF network slice requests.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Editor's Note: the terminology in this draft will be aligned in forthcoming versions with the final terminology selected for describing the notion of IETF network slice when applied to IETF technologies, which is currently under discussion. By now same terminology as used in [I-D.nsdt-teas-ietf-network-slice-definition] and [I-D.nsdt-teas-ns-framework] is primarily used here.

Editor's Note: the term "transport network" in the context of this draft refers in broad sense to WAN, MBH, IP backbone and other network segments implemented by IETF technologies.

A number of new technologies, such as 5G, NFV and SDN are not only evolving the network from a pure technological perspective but also are changing the concept in which new services are offered to the customers [I-D.homma-slice-provision-models] by introducing the concept of network slicing.

The transport network is an essential component in the end-to-end delivery of services and, consequently, it is necessary to understand what could be the way in which the transport network is consumed as a slice. For a definition of IETF network slice refer to [I-D.nsdt-teas-ietf-network-slice-definition].

In this document it is assumed that there exists a (logically) centralized component in the transport network, namely IETF Network Slice Controller (NSC) with the responsibilities on the control and management of the IETF network slices invoked for a given service, as requested by IETF network slice customers.

This document analyses different use cases deriving the needs of potential IETF network slice customers in order to identify the functionality required on the North Bound Interface (NBI) of the NSC to be exposed towards such IETF network slice customers. Solutions to construct the requested IETF network slices are out of scope of this document.

This document addresses some of the discussions of the TEAS Slice Design Team. However, it is not at this stage an official outcome of the Design Team.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [RFC2119].

3. Northbound interface for IETF network slices

In a general manner, the transport network supports different kinds of services. These services consume capabilities provided by the transport network for deploying end-to-end services, interconnecting network functions or applications spread across the network and providing connectivity toward the final users of these services.

Under the slicing approach, a IETF network slice customer requests to a IETF network slice controller a slice with certain characteristics and parametrization. Such request it is assumed here to be done through a NBI exposed by the NSC to the customer, as reflected in Fig. 1.

```
                         +--------------------+
                         |                    |
                         |    IETF Network    |
                         |   Slice Customer   |
                         |                    |
                         +--------------------+
                                   A
                                   |
                                   | IETF Network
                                   | Slice Controller
                                   | NBI
                                   |
                                   V
                         +--------------------+
                         |                    |
                         |    IETF Network    |
                         |  Slice Controller  |
                         |                    |
                         +--------------------+
```

Figure 1: IETF network slice NBI concept

The functionality supported by the NBI depends on the requirements that the slice customer has to satisfy. It is then important to understand the needs of the slice customers as well as the way of expressing them.

4. IETF network slice use cases

Different use cases for slice customers can be identified, as described in the following sections.

4.1. 5G Services

5G services natively rely on the concept of network slicing. 5G is expected to allow vertical customers to request slices in such a manner that the allocated resources and capabilities in the network appear as dedicated for them.

In network slicing scenarios, a vertical customer requests a network operator to allocate a network slice instance (NSI) satisfying a

particular set of service requirements. The content/format of these requirements are highly dependent on the networking expertise and use cases of the customer under consideration. To deal with this heterogeneity, it is fundamental for the network operator to define a a unified ability to interpret service requirements from different vertical customers, and to represent them in a common language, with the purposes of facilitating their translation/mapping into specific slicing-aware network configuration actions. In this regard, model-based network slice descriptors built on the principles of reproducibility, reusability and customizability can be defined for this end.

As a starting point for such a definition, GSMA developed the idea of having a universal blueprint that, being offered by network operators, can be used by any vertical customer to order the deployment of an NSI based on a specific set of service requirements. The result of this work has been the definition of a baseline network slice descriptor called Generic network Slice Template (GST). The GST contains multiple attributes that can be used to characterize a network slice. A Network Slice Type (NEST) describes the characteristics of a network slice by means of filling GST attributes with values based on specific service requirements. Basically, a NEST is a filled-in version of a GST. Different NESTs allow describing different types of network slices. For slices based on standardized service types, e.g. eMBB, uRLLC and mIoT, the network operator may have a set of readymade, standardized NESTs (S-NESTs). For slices based on specific industry use cases, the network operator can define additional NESTs.

Service requirements from a given vertical customer are mapped to a NEST, which provides a self-contained description of the network slice to be provisioned for that vertical customer. According to this reasoning, the NEST can be used by the network operator as input to the NSI preparation phase, which is defined in [TS28.530]. 3GPP is working on the translation of the GST/NEST attributes into NSI related requirements, which are defined in the "ServiceProfile" data type from the Network Slice Information Object Class (IOC) in [TS28.541]. These requirements are used by the 3GPP Management System to allocate the NSI across all network domains, including transport network. The IETF network slice defines the part of that NSI that is deployed across the transport network.

Despite the translation is an on-going work in 3GPP it seems convenient to start looking at the GST attributes to understand what kind of parameters could be required for the IETF network slice NBI.

4.1.1. Generic network Slice Template

The structure of the GST is defined in [GSMA]. The template defines a total of 35 attributes. For each of them, the following information is provided:

- o Attribute definition, which provides a formal definition of what the attribute represents.

- o Attribute parameters, including:

  * Value, e.g. integer, float.

  * Measurement unit, e.g. milliseconds, Gbps

  * Example, which provides examples of values the parameter can take in different use cases.

  * Tag, which allow describing the type of parameter, according to its semantics. An attribute can be tagged as a characterization attribute or a scalability attribute. If it is characterization attribute, it can be further tagged as a performance-related attribute, a functionality-related attribute or an operation-related attribute.

  * Exposure, which allow describing how this attribute interact with the slice customer, either as an API or a KPI.

- o Attribute presence, either mandatory, conditional or optional.

Attributes from GST can be used by the network operator (slice controller) and a vertical customer (slice customer) to agree SLA.

GST attributes are generic in the sense that they can be used to characterize different types of network slices. Once those attributes become filled with specific values, it becomes a NEST which can be ordered by slice customers.

4.1.2. Categorization of GST attributes

Not all the GST attributes as defined in [GSMA] have impact in the transport network since some of them are specific to either the radio or the mobile core part.

In the analysis performed in this document, the attributes have been categorized as:

- Directly impactive attributes, which are those that have direct impact on the definition of the IETF network slice, i.e., attributes that can be directly translated into requirements required to be satisfied by a IETF network slice.

- Indirectly impactive attributes, which are thise that impact in an indirect manner on the definition of the IETF network slice, i.e., attributes that indirectly impose some requirements to a IETF network slice.

- Non-impactive attributes, that are those which do not have impact on the IETF network slice at all.

The following sections describe the attributes falling into the three categories.

4.1.2.1. Attributes with direct impact on the IETF network slice definition

The following attributes impose requirements in the IETF network slice

- Availability
- Deterministic communication
- Downlink throughput per network slice
- Energy efficiency
- Group communication support
- Isolation level
- Maximum supported packet size
- Mission critical support
- Performance monitoring
- Slice quality of service parameters
- Support for non-IP traffic
- Uplink throughput per network slice
- User data access (i.e., tunneling mechanisms)

4.1.2.2. Attributes with indirect impact on the IETF network slice definition

The following attributes indirectly impose requirements in the IETF network slice to support the end-to-end service.

- Area of service (i.e., the area where terminals can access a particular network slice)
- Delay tolerance (i.e., if the service can be delivered when the system has sufficient resources)
- Downlink (maximum) throughput per UE
- Network functions owned by Network Slice Customer
- Maximum number of (concurrent) PDU sessions
- Performance prediction (i.e., capability to predict the network and service status)
- Root cause investigation
- Session and Service Continuity support
- Simultaneous use of the network slice
- Supported device velocity
- UE density
- Uplink (maximum) throughput per UE
- User management openness (i.e., capability to manage users' network services and corresponding requirements)
- Latency from (last) UPF to Application Server

4.1.2.3. Attributes with no impact on the IETF network slice definition

The following attributes do not impact the IETF network slice.

- Location based message delivery (not related to the geographical spread of the network slice itself but with the localized distribution of information)
- MMTel support, i.e. support of and Multimedia Telephony Service (MMTel)as well as IP Multimedia Subsystem (IMS) support.

- NB-IoT Support, i.e., support of NB-IoT in the RAN in the network slice.

- Maximum number of (simultaneous) UEs

- Positioning support

- Radio spectrum

- Synchronicity (among devices)

- V2X communication mode

- Network Slice Specific Authentication and Authorization (NSSAA)

4.1.3. Provisioning procedures

3GPP identifies in [TS28.541] a number of procedures for the provisioning of a network slice in general. It can be assumed that similar procedures may also apply to a transport slice, facilitating a consistent management and control of end-to-end slices.

The envisioned procedures are the following:

- Slice instance allocation: this procedure permits to create a new slice instance (or reuse an existing one).

- Slice instance de-allocation: this procedure decommissions a previously instantiated slice.

- Slice instance modification: this procedure permits the change in the characteristics of an existing slice instance.

- Get slice instance status: this procedure helps to retrieve run-time information on the status of a deployed slice instance.

- Retrieval of slice capabilities: this procedure assists on getting information about the capabilities (e.g. maximum latency supported).

All these procedures fit in the operation of transport network slices.

4.2. NFV-based services

NFV technology allows the flexible and dynamic instantiation of virtualized network functions (and their composition into network services) on top of a distributed, cloud-enabled compute

infrastructure. This infrastructure can span across different points of presence in a carrier network. By leveraging on transport network slicing, connectivity services established across geographically remote points of presence can be enriched by providing additional QoS guarantees with respect present state-of-the-art mechanisms, as conventional L2/L3 VPNs.

4.2.1. Connectivity attributes

The connectivity services are expressed through a number of attributes as listed:

o Incoming and outgoing bandwidth: bandwidth required for the connectivity services (in Mbps).

o Qos metrics: set of metrics (e.g., cost, latency and delay variation) applicable to a specific connectivity service

o Directionality: indication if the traffic is unidirectional or bidirectional.

o MTU: value of the largest PDU to be transmitted in the connectivity service.

o Protection scheme: indication of the kind of protection to be performed (e.g., 1;1, 1+1, etc.)

o Connectivity mode: indication of the service is point-to-point of point-to-multipoint

All those attributes will assist on the characterization of the connectivity slice to be deployed, and thus, are relevant for the definition of a IETF network slice supporting such connectivity.

4.2.2. Provisioning procedures

ETSI NFV defines the role of WAN Infrastructure Manager (WIM) as the component in charge of managing and controlling the connectivity external to the PoPs. In [IFA032] a number of interfaces are identified to be exposed by the WIM for supporting the multi-site connectivity, thus representing the capabilities expected for a transport network slice, as well, in case of satisfying such connectivity needs by means of the slice concept.

The interfaces considered are the following:

o Multi-Site Connectivity Service (MSCS) Management: this interface permits the creation, termination, update and query of MSCSs,

including reservation. It also enables subscription for notifications and information retrieval associated to the connectivity service.

- o Capacity Management: this interface allows querying about the capacity (e.g. bandwidth), topology, and network edge points of the connectivity service, as well as about information of consumed and available capacity on the underlying network resources.

- o Fault Management: this interface serves for the provision of alarms related to the MSCSs.

- o Performance Management: this interface assists on the retrieval of performance information (measurement results collection and notifications) related to MSCSs.

4.3. RAN sharing

Network sharing is one of the means network operators exploit for increasing efficiencies. There are different scenarios of network sharing, being especially popular in the deployment of mobile networks, typically referred to as Radio Access Network (RAN) sharing. From an operational perspective, in RAN sharing we have two roles: master operator, being the actor (e.g. infrastructure provider, network operator) to which the deployment and daily operation of shared RAN elements are entrusted to; and the participant operators, who are the mobile operators who share the RAN facilities provided by the master operator. Note that in this context the master and participant operator can be seen as provider and customer, respectively.

While there exist different modes of RAN sharing [TS23.251], including passive RAN sharing (infrastructure site sharing) and active RAN sharing (e.g. Multi-Operator Core Networks or MOCN), most of the cases require the establishment of separated connections in order to separate the traffic per participant operator. Such connections typically extend from the cell site to some pre-defined and agreed interconnection points, from which the traffic is routed and delivered to individual participant operators.

The above-referred connections can have specific attributes. Aspects like guaranteed bandwidth (in line with the expected load from the aggregated cells), redundancy, bounded latency (per kind of traffic), or secure delivery of the information should be considered.

The master operator is the one in charge of provisioning the connections and collecting management data (e.g. performance measurements, telemetry, fault alarms, trace data) for individual

participant operators. The use of network slicing could make the network sharing approach more flexible by allowing the other operators control and manage the established connections [MEF].

The implications of the RAN sharing scenario here described can be extended to either fixed networks or even to mobile networks leveraging on radio functional split (i.e., including fronthaul and midhaul network segments).

4.3.1. Connectivity attributes

The connections for RAN sharing typically consider attributes like:

- o Maximum and Guaranteed Bit Rate (MBR and GBR respectively).

- o Bounded latency (e.g., for user plane, control plane, etc)

- o Packet loss rate.

- o IP addressing (consistent among the operators sharing the infrastructure).

- o L2/L3 reachability.

- o Recovery time (on the event of failures).

- o Secure connection (e.g., encryption support).

4.3.2. Provisioning procedures

The expected provisioning procedures are:

- o Connection provisioning between site and interconnection point. Those connections could evolve in time in terms of capacity depending on the capacity grow of each particular site.

- o Collection of management data, including performance measurements, fault alarms and trace data.

4.4. Additional use cases

This is a placeholder for describing additional use cases (e.g., data center interconnection, etc). To be completed.

5. Security Considerations

This draft does not include any security considerations.

6. IANA Considerations

This draft does not include any IANA considerations

7. References

7.1. Normative References

[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>.

7.2. Informative References

[GSMA] "Generic Network Slice Template, version 3.0", NG.116 , May 2020.

[I-D.homma-slice-provision-models] Homma, S., Nishihara, H., Miyasaka, T., Galis, A., OV, V., Lopez, D., Contreras, L., Ordonez-Lucena, J., Martinez-Julia, P., Qiang, L., Rokui, R., Ciavaglia, L., and X. Foy, "Network Slice Provision Models", draft-homma-slice-provision-models-02 (work in progress), November 2019.

[I-D.nsdt-teas-ietf-network-slice-definition] Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "Definition of IETF Network Slices", draft-nsdt-teas-ietf-network-slice-definition-00 (work in progress), October 2020.

[I-D.nsdt-teas-ns-framework] Gray, E. and J. Drake, "Framework for Transport Network Slices", draft-nsdt-teas-ns-framework-04 (work in progress), July 2020.

[IFA032] "IFA032 Interface and Information Model Specification for Multi-Site Connectivity Services V3.2.1.", ETSI GS NFV-IFA 032 V3.2.1 , April 2019.

[MEF] "Slicing for Shared 5G Fronthaul and Backhaul", MEF White paper , April 2020.

[TS23.251]
"TS 23.251 Network Sharing; Architecture and functional description (Release 16) V16.0.0.", 3GPP TS 23.251 V16.0.0 , July 2020.

[TS28.530]
"TS 28.530 Management and orchestration; Concepts, use cases and requirements (Release 16) V16.0.0.", 3GPP TS 28.530 V16.0.0 , September 2019.

[TS28.541]
"TS 28.541 Management and orchestration; 5G Network Resource Model (NRM); Stage 2 and stage 3 (Release 16) V16.2.0.", 3GPP TS 28.541 V16.2.0 , September 2019.

Authors' Addresses

Luis M. Contreras
Telefonica
Ronda de la Comunicacion, s/n
Sur-3 building, 3rd floor
Madrid 28050
Spain

Email: luismiguel.contrerasmurillo@telefonica.com
URI: http://lmcontreras.com/

Shunsuke Homma
NTT
Japan

Email: shunsuke.homma.ietf@gmail.com

Jose A. Ordonez-Lucena
Telefonica
Ronda de la Comunicacion, s/n
Sur-3 building, 3rd floor
Madrid 28050
Spain

Email: joseantonio.ordonezlucena@telefonica.com