

TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: January 13, 2022

LM. Contreras
Telefonica
S. Homma
NTT
J. Ordonez-Lucena
Telefonica
J. Tantsura
Microsoft
K. Szarkowicz
Juniper Networks
July 12, 2021

IETF Network Slice Use Cases and Attributes for Northbound Interface of
IETF Network Slice Controllers
draft-contreras-teas-slice-nbi-05

Abstract

This document analyses the needs of potential customers of network slices realized with IETF techniques in several use cases, identifies the functionalities for the North Bound Interface (NBI) of an IETF Network Slice Controller to satisfy such requests.

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 January 13, 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions used in this document and terminology	3
3. Northbound Interface for IETF Network Slices	4
4. IETF Network Slice Use Cases	5
4.1. 5G Services	5
4.1.1. 3GPP network slice	6
4.1.1.1. Topology of the TN-NSS	6
4.1.1.2. Traffic segregation and mapping to S-NSSAI list	7
4.1.1.3. Reachability information	10
4.1.1.4. QoS profiling	10
4.1.2. Generic network Slice Template	10
4.1.3. Categorization of GST attributes	11
4.1.3.1. Attributes with direct impact on the IETF network slice definition	12
4.1.3.2. Attributes with indirect impact on the IETF network slice definition	12
4.1.3.3. Attributes with no impact on the IETF network slice definition	13
4.1.4. Provisioning procedures	14
4.2. NFV-based services	14
4.2.1. Connectivity attributes	15
4.2.2. Provisioning procedures	15
4.3. Network sharing	16
4.3.1. Connectivity attributes	17
4.3.2. Provisioning procedures	17
4.4. SD-WAN	17
4.4.1. SD-WAN Structure	18
4.4.2. Connectivity Attributes	19
4.4.3. SD-WAN Endpoint Attributes	21
4.4.4. SD-WAN UNI Attributes	21
4.5. Radio functional splits	22
4.5.1. Attributes and procedures	23
4.6. Additional use cases	23
5. Security Considerations	23
6. IANA Considerations	23
7. References	23
7.1. Normative References	23
7.2. Informative References	23

Authors' Addresses	24
--------------------	----

1. Introduction

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.ietf-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 and terminology

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].

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.ietf-teas-ietf-network-slice-definition] and [I-D.nsdtd-teas-ns-framework] is primarily used here.

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.

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 Figure 1.




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 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 mMTC, 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.

4.1.1.1. 3GPP network slice

A 3GPP network slice represents a logical network that provides specific capabilities and network characteristics, supporting the service requirements of one or more network slice customers. The service requirements of each network slice customer are captured into a separate "ServiceProfile" artifact within the network slice class (see Network Slicing NRM fragment in TS 28.541).

A 3GPP network slice spans from 5G NR access nodes to the UPF that terminates the PDU session, i.e. PSA UPF. In this in-slice data path, there are TN segments (e.g. backhaul) that are out of scope of 3GPP management domain. For the provisioning and operation of these TN segments, usually referred to as transport Network Slice Subnets (TN-NSS), the 3GPP management system relies on an external TN management system, which hosts (among other components) the IETF NSC. To proceed with this delegation, the 3GPP management system needs to make available to the TN management system the information described in the following sub-sections.

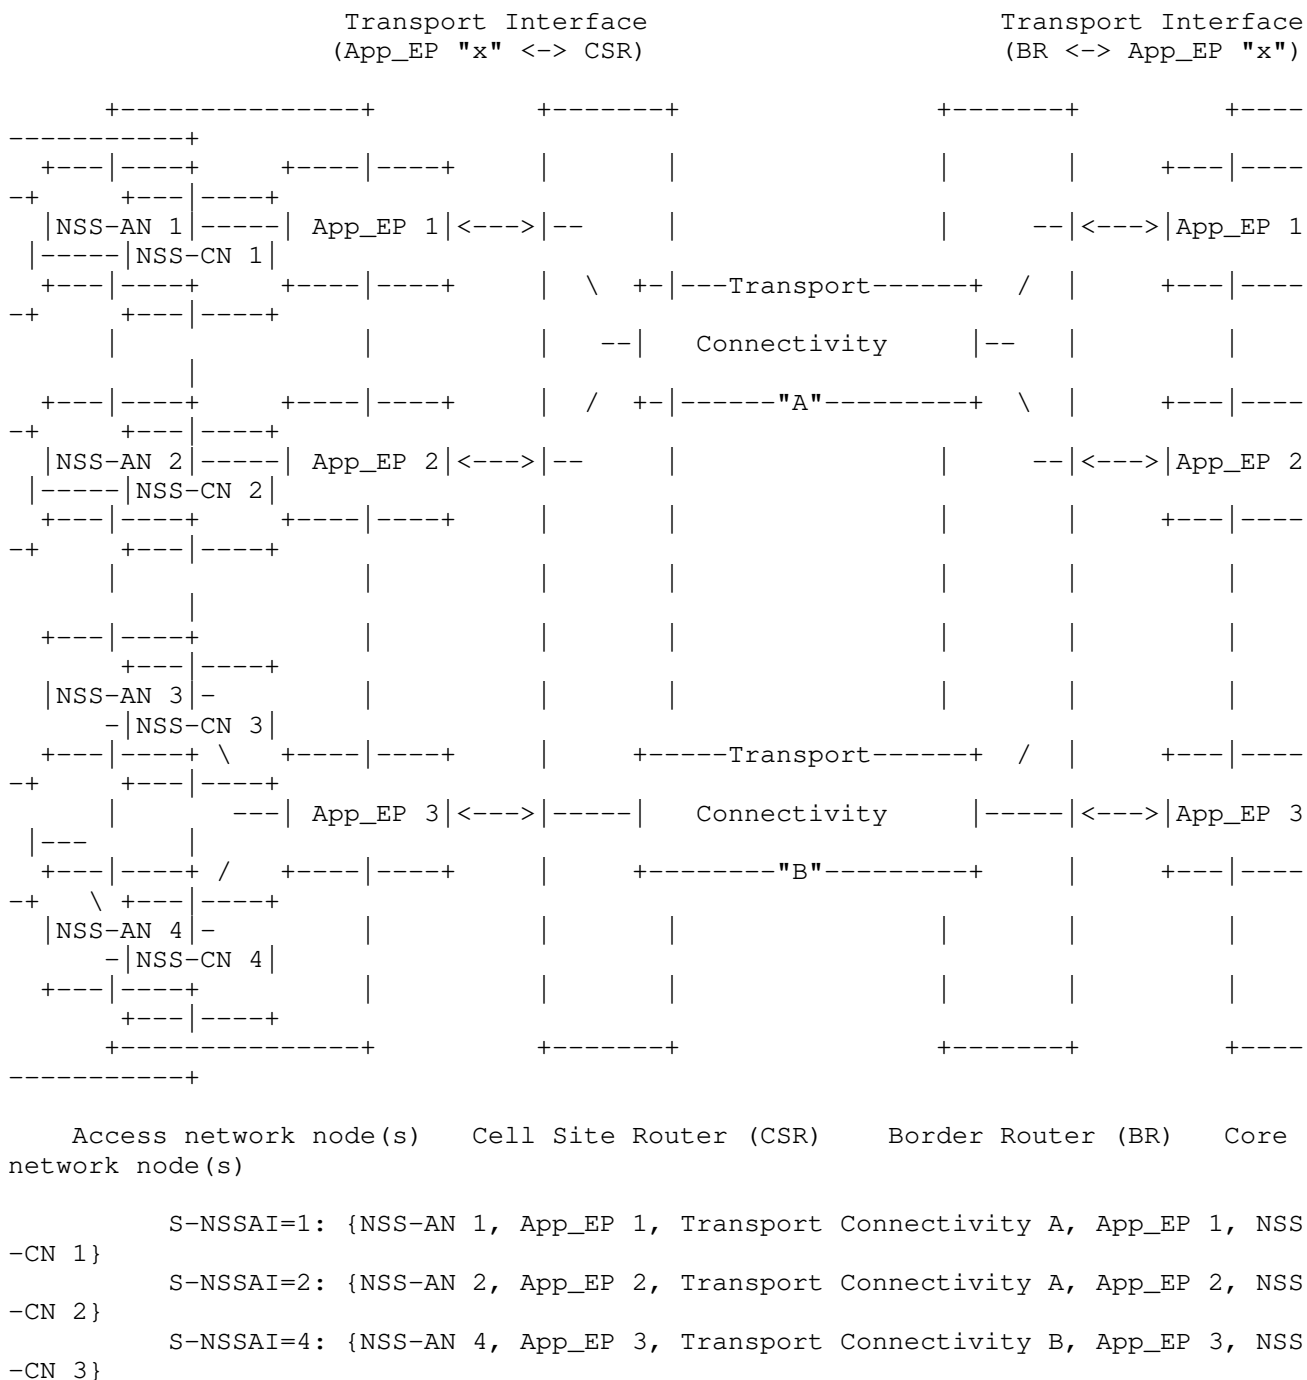
4.1.1.1.1. Topology of the TN-NSS

The TN management system needs to know the transport termination/end points to determine the transport resources, either physical or virtual nodes. 3GPP management system systems need to provide the transport endpoints of 3GPP managed functions that are part of the RAN-NSS (e.g., gNB-CU-UP, gNB-CU-CP) and CN-NSS (e.g., UPF, AMF), and if applicable further information such as the next-hop router IP address configured in a RAN-NSS or CN-NSS. The TN management system should be able to correlate this with the transport network topology and derive the site or border routers connecting to 3GPP managed functions.

4.1.1.2. Traffic segregation and mapping to S-NSSAI list

As network functions can be shared by many network slices, it will be necessary to segregate the traffic belonging to specific slices on transport interfaces.

One option for traffic segregation is to assign application endpoints to specific sets of S-NSSAI values. The transport network can map packets to connectivity services based on local remote or remote endpoints, provided that the allocation of S-NSSAI to endpoints is known and exposed, and provided that the application endpoints are visible on the transport layer. The application endpoints visible in a RAN-NSS and CN-NSS are already mapped to a specific set of S-NSSAI. Figure 2 illustrates an example of this solution, whereby a 3GPP network slice with S-NSSAI=1 is mapped to specific application endpoints (e.g., N3 tunnel endpoint 1) by the access network node. In this example, the TN management system decides to map application endpoints 1 and 2 to the same transport connectivity service A. This mapping is implemented by the site router connecting to the access network node. On the core network slice, a similar mapping is done by the border router. Demultiplexing the packet streams belonging to different transport interfaces is based on regular routing and reachability of endpoint IP addresses.



+-----+
|NSS-AN 1|
+-----+
|App_EP 1|

+-----+
|NSS-CN 1|
+-----+
|App_EP 1|

+-----+
|NSS-AN 2|
+-----+
|App_EP 2|

+-----+
|NSS-CN 2|
+-----+
|App_EP 2|

+-----+
|NSS-AN 3|
+-----+
|App_EP 3|

+-----+
|NSS-CN 3|
+-----+
|App_EP 3|

+-----+
|NSS-AN 4|
+-----+
|App_EP 3|

+-----+
|NSS-CN 4|
+-----+
|App_EP 3|

Access network node(s)
network node(s)

Cell Site Router (CSR)

Border Router (BR)

Core
network node(s)

S-NSSAI=1:

{NSS-AN 1, App_EP 1, Transport Connectivity A, App_EP 1, NSS-CN 1}

S-NSSAI=2:

{NSS-AN 2, App_EP 2, Transport Connectivity A, App_EP 2, NSS-CN 2}

S-NSSAI=4:


{NSS-AN 4, App_EP 3, Transport Connectivity B, App_EP 3, NSS-CN 3}

Figure 2: Mapping of S-NSSAI to specific application endpoints

Despite the simplicity of the above-referred approach, notice that it is not a universal solution as the application endpoint addresses are not always visible to the TN, for example when they are encrypted by IPSec tunnels. In such a case, the application endpoints are not visible to the site router, and thus cannot be used for transport connectivity mappings. To deal with these situations, an alternative solution is to use the concept of logical transport interfaces. A logical transport interface is a virtual interface separate from

application endpoints; it can be for example a specific IP address / VLAN combination that corresponds to an IPSec termination point, an identifier (e.g., MPLS label, segment ID) that the TN recognizes, or it can be just a logical interface defined on top of top a physical transport interface. As long as the interface identity can derived from packet headers, the TN nodes can perform the mapping to transport connectivity services. In this regard, it is useful to indicate to the TN which traffic types are carried over an interface (e.g., N3 user plane packets, N2 control plane packets, etc.).

Figure 3 illustrates an example on the use of this solution. As seen, logical transport need to be exposed from 3GPP management system to TN management system, so that the latter can create transport network topology and determine the TN resources to support the 3GPP slice.



S-NSSAI=1: {NSS-AN 1, Logical Transport Interface 1, Transport Connectivity A, Logical Transport Interface 1, NSS-CN 1}

S-NSSAI=2: {NSS-AN 2, Logical Transport Interface 2, Transport Connectivity A, Logical Transport Interface 2, NSS-CN 2}

S-NSSAI=4: {NSS-AN 4, Logical Transport Interface 3, Transport Connectivity B, Logical Transport Interface 3, NSS-CN 4}

Figure 3: Logical Transport Interfaces

For traffic segregation, though solutions might be valid, 3GPP prefers the second solution: on the use of concept of transport logical interface. The reason is that it does not impose 1:1 mapping between application endpoint and transport interface (allowing for better redundancy) and that it always works, no matter if encryption. To support this solution, the 3GPP has recently extended the Network Slice NRM fragment, including a new Information Object Class called EP_Transport. This class provides a complete characterization of the logical transport interface, including transport level information (i.e., IP address, reachability information, QoS profile) and the set

of application endpoints aggregated to this interface. For further information on reachability information and QoS profile, see next subsections. For further details on fields of EP_Transport, see Network Slice NRM fragment in TS 28.541.

4.1.1.3. Reachability information

Each physical or logical transport interface will carry the traffic associated with some 3GPP application endpoints that may be using IP addresses separate from the transport interface. These IP addresses must be reachable within the TN-NSS, and hence they need to be advertised to populate forwarding tables. A 3GPP network function can advertise such reachability information by running a dynamic routing protocol towards the next hop router. If that is not possible, it can create association between the reachability data with the logical transport interface and expose it towards the 3GPP and TN management system. This information can be derived from the IP addresses available for application and transport endpoints.

4.1.1.4. QoS profiling

Each TN-NSS may be associated a "TNSliceSubnetProfile", which hosts the SLO requirements (e.g., guaranteed throughput, bounded latency, maximum jitter) that the TN-NSS must support. "TNSliceSubnetProfile" is a 3GPP artifact that result from the decomposition of e2e service requirements ("ServiceProfile" artifact) into domain-specific service requirements ("RANSliceSubnetProfile", "CNSliceSubnetProfile" and "TNSliceSubnetProfile") applicable to RAN-NSS, CN-NSS and TN-NSS respectively. Unlike "RANSliceSubnetProfile" and "CNSliceSubnetProfile", there is not agreement yet on the specific parameters to be captured by the "TNSliceSubnetProfile". Further work in this regard in the upcoming 3GPP SA5 meetings.

Upon receiving the "TNSliceSubnetProfile" from the 3GPP management system, the TN management system translates the SLO requirements therein into a QoS profile, which includes applicability and use of DSCPs and other QoS related properties onto the TN-NSS realization. To enable this, each logical interface may have an associated QoS profile. The QoS profile is just a reference to the detailed profile parameters which are logically provisioned on both sides of a logical transport interface.

4.1.2. 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.

4.1.3. 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:

- o Directly impactful 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.
- o Indirectly impactful attributes, which are those 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.

- o 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.

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

The following attributes impose requirements in the IETF network slice

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

4.1.3.2. 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.

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

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

The following attributes do not impact the IETF network slice.

- o Location based message delivery (not related to the geographical spread of the network slice itself but with the localized distribution of information)
- o MMTel support, i.e. support of and Multimedia Telephony Service (MMTel) as well as IP Multimedia Subsystem (IMS) support.
- o NB-IoT Support, i.e., support of NB-IoT in the RAN in the network slice.
- o Maximum number of (simultaneous) UEs

- o Positioning support
- o Radio spectrum
- o Synchronicity (among devices)
- o V2X communication mode
- o Network Slice Specific Authentication and Authorization (NSSAA)

4.1.4. 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:

- o Slice instance allocation: this procedure permits to create a new slice instance (or reuse an existing one).
- o Slice instance de-allocation: this procedure decommissions a previously instantiated slice.
- o Slice instance modification: this procedure permits the change in the characteristics of an existing slice instance.
- o Get slice instance status: this procedure helps to retrieve run-time information on the status of a deployed slice instance.
- o 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 or 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 MSCSSs.
- o Performance Management: this interface assists on the retrieval of performance information (measurement results collection and notifications) related to MSCSSs.

4.3. Network 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 growth of each particular site.
- o Collection of management data, including performance measurements, fault alarms and trace data.

4.4. SD-WAN

SD-WAN is a solution to provide a virtual overlay network for connecting between customer's sites, (virtual) private cloud, or public cloud/Internet. SD-WAN operates over one or more underlay networks, and enables to offer more differentiated service delivery capabilities. SD-WAN can be esteemed as a type of network slices or can be established over underlay networks provided as network slices. The definitions, specification, service attributes, and framework of SD-WAN is defined in Metro Ethernet Forum ([MEF-70]).

SD-WAN forwards traffic based on application flows, and the policies include rules and constraints on the forwarding of the application

flows. In SD-WAN, it may be required from the customer to adjust the behaviors based on its needs in near real time. The service provider is required to monitor the performance of the service and modify the forwarding policies based on the real-time telemetry from the underlying network components.

4.4.1. SD-WAN Structure


SD-WAN has three logical constructs:

- o SD-WAN virtual connection
- o SD-WAN virtual connection endpoint
- o SD-WAN UNI

Several additional components may be visible to the customer. These include:

- o Customer network
- o Service provider network
- o Underlay connectivity
- o Tunnel virtual connection

The following figure shows the overview of SD-WAN structure. In this case, the customer sites are connected with underlay connectivity#1 and they are also connected to remote private cloud with underlay connectivity#2. An SD-WAN endpoint is usually located in each customer network site as a CPE or a customer edge, and it allocates application flow to appropriate underlay connectivity.



* Legend

. . . : Underlay connectivity#1

==== : Underlay Connectivity#2

EP : SD-WAN Endpoint

Figure 4: Overview of SD-WAN Structure

SD-WAN may be provided as a network slice, or it is realized on several network slices provided as underlay connectivities. In the former case, a network slice PE will be mapped to CE in SD-WAN. In the later case, PEs of the provider of underlay connectivities will behave as network slice PEs.

4.4.2. Connectivity Attributes

SD-WAN defined in MEF-70 has several attributes on its connectivity as below:

- o SD-WAN Identifier: the value is a string that is used by the customer and service provider to uniquely identify an SD-WAN connectivity.
- o Endpoint list: the value is a list contains endpoint identifiers and their connected endpoints.
- o Service Uptime Objective: the value is the proportion of time that the connectivity service is working during a given time period.

- o **Reserved Prefixes:** the values are IP prefixes reserved by the service provider for use for SD-WAN within its own network or for distribution to the customer via DHCP or SLAAC.
- o **List for Policies:** the value is a list of policies applied to application flows and application flow groups at endpoints. An SD-WAN policy list contains policy name and list of policy criteria. Support of the criteria listed below would be required:
 - * **Encryption:** indicates whether or not the application flow requires encryption
 - * **Public-Private:** indicates whether the application flow can traverse public or private underlay connectivity services (or both).
 - * **Internet-Breakout:** indicates whether the application flow should be forwarded to an Internet destination.
 - * **Billing-Method:** indicate the application flow can be sent over an underlay connectivity service that has usage-based or flat-rate billing.
 - * **Backup:** indicates whether this application flow can use a TVC designated as aEURbackupaEUR.
 - * **Bandwidth:** specifies a rate limit on the application flow.
- o **List of Application Flow Groups:** the value is a list of application flow groups that application flows can be members of. An application flow group list contains application flow group name and application flow group policy.
- o **List of Application Flows:** the value is a list of the application flows that are recognized by the SD-WAN. An application flow list contains application flow name, list of application flow criteria, and application flow group name. The criteria is listed below:
 - * **Ethertype**
 - * **C-VLAN ID list**
 - * **IPv4 source address**
 - * **IPv4 destination address**
 - * **IPv4 source or destination address**

- * IPv4 protocol list
- * IPv6 source address
- * IPv6 destination address
- * IPv6 source or destination address
- * IPv6 next header list
- * TCP/UDP source port list
- * TCP/UDP destination port list
- * Application identifier
- * any

4.4.3. SD-WAN Endpoint Attributes

SD-WAN contains some endpoints as boundary nodes between underlay connections and customers sites. [MEF-70] defines some attributes for SD-WAN endpoints as below:

- o Endpoint Identifier: the value is for identification of SD-WAN endpoint for management purposes.
- o Endpoint UNI: the value is for identification of the UNI that the endpoint is associated with.
- o Endpoint policy map: the value is for mapping policies to application flows and application flow groups.

4.4.4. SD-WAN UNI Attributes

SD-WAN UNI is a reference point that represents the demarcation between the responsibility of the customer and the responsibility of the provider. Some attributes for UNI is defined in [MEF-70] as below:

- o SD-WAN UNI Identifier: the value is for identification of the UNI for management purposes.
- o SD-WAN UNI L2 Interface: the value describes the underlay L2 interface for the UNI.
- o SD-WAN UNI Maximum L2 Frame Size: the value specifies the maximum length L2 frame that is accepted by the provider.

- o SD-WAN UNI IPv4 connection addressing: the value describes IPv4 connection address mechanisms (e.g., Static or DHCP).
- o SD-WAN UNI IPv6 connection addressing: the value describes IPv6 connection address mechanisms (e.g., DHCP, SLAAC, Static or Link-Local-only).

4.5. Radio functional splits

The disaggregation of the software stack in radio base stations allows the centralization of some of the radio processing functions. O-RAN is promoting the interoperability of implementations of radio functional splits, defining an architecture where three main entities can be considered: the Radio Unit (RU), with some basic processing, the Distributed Unit (DU) with the rest of real-time processing capabilities, and the Centralized Unit (CU) with the non-real-time processing of the software stack. The network segment between RU and DU is known as fronthaul (FH), while the segment between DU and CU is referred as midhaul (MH). Figure 5 shows this situation.




Figure 5: Logical Transport Interfaces

The fronthaul leverages on eCPRI protocol which can be transported directly on Ethernet frames or encapsulated in IP/UDP (for the user plane). The midhaul can be transported in a similar way as the backhaul.

With current specifications, individual service flows being carried by FH cannot be distinguished, so no possibility of differentiating connectivity slices at that point. Similar thing happens for MH. The only possible differentiation per flow can happen in downstream direction from CU to DU, but this basically can only help for policing traffic at that point (i.e., slice is yet the same).

Advanced scenarios such as RU sharing could allow traffic differentiation per mobile operator based on e.g. vlans, being each of those vlans mapped to a different slice.

4.5.1. Attributes and procedures

The attributes of IETF network slices for the conveniently supported the radio functional split are based on main characteristics of FH/MH: Latency, BW, and packet loss, as specified in [O-RAN]. Geographical location could have an impact due to latency restrictions for FH.

Regarding slice management procedures, it can be assumed a similar lifecycle as in 3GPP slices.

4.6. 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. R., Lopez, D. R., Contreras, L. M., Ordonez-Lucena, J. A., Martinez-Julia, P., Qiang, L., Rokui, R., Ciavaglia, L., and X. D. Foy, "Network Slice Provision Models", draft-homma-slice-provision-models-02 (work in progress), November 2019.

- [I-D.ietf-teas-ietf-network-slice-definition]
Rokui, R., Homma, S., Makhiyani, K., Contreras, L. M., and J. Tantsura, "Definition of IETF Network Slices", draft-ietf-teas-ietf-network-slice-definition-01 (work in progress), February 2021.
- [I-D.nsdt-teas-ns-framework]
Gray, E. and J. Drake, "Framework for IETF Network Slices", draft-nsdt-teas-ns-framework-05 (work in progress), February 2021.
- [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.
- [MEF-70] "SD-WAN Service Attributes and Services", MEF-70 , July 2019.
- [O-RAN] "O-RAN Xhaul Transport Requirements 1.0", O-RAN.WG9.XTRP-REQ-v01.00 , November 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

Jeff Tantsura
Microsoft

Email: jefftant.ietf@gmail.com

Krzysztof Szarkowicz
Juniper Networks

Email: kszarkowicz@juniper.net