

teas
Internet-Draft
Intended status: Informational
Expires: January 13, 2021

R. Rokui
Nokia
S. Homma
NTT
K. Makhijani
Futurewei
LM. Contreras
Telefonica
J. Tantsura
Apstra, Inc.
July 12, 2020

IETF Definition of Transport Slice
draft-nsdt-teas-transport-slice-definition-03

Abstract

This document describes the definition of a slice in the transport networks and its characteristics. The purpose here is to bring clarity and a common understanding of the transport slice concept and describe related terms and their meaning. It explains how transport slices can be used in combination with end to end network slices, or independently.

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, 2021.

Copyright Notice

Copyright (c) 2020 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	2
1.1.	Rationale	3
2.	Terms and Abbreviations	3
3.	Definition and Scope of Transport Slice	4
4.	Transport Slice System Characteristics	5
4.1.	Service Level Objectives for Transport Slices	5
4.1.1.	Minimal Set of SLOs	5
4.1.2.	Other Objectives	7
4.2.	Transport Slice Endpoints	8
4.2.1.	Transport Slice Connectivity Types	9
4.3.	Vertical Composition of Transport Slice	9
4.4.	Horizontal Composition of Transport Slice	11
5.	Transport Slice Structure	11
5.1.	Stakeholders	13
5.2.	Transport Slice Controller Interfaces	14
5.3.	Transport slice Realization	15
6.	Relationship with End-to-End Network Slicing	15
7.	Security Considerations	17
8.	IANA Considerations	17
9.	Acknowledgment	17
10.	Informative References	17
Appendix A.	Discussions	19
A.1.	On Isolation Requirements In a Transport Slice	19
	Authors' Addresses	20

1. Introduction

A number of use cases benefit from establishing network connectivity providing transport and assurance of a specific set of network resources. In this document, as detailed in the subsequent sections, we refer to this connectivity and resource commitment as the transport slice. Services that might benefit from the transport slices include but not limited to:

- o 5G services (e.g. eMBB, URLLC, mMTC)(See [[TS.23.501-3GPP](#)])

- o Network wholesale services
- o Network infrastructure sharing among operators
- o NFV connectivity and Data Center Interconnect

This document defines the concept of transport slices that provide connectivity with a specific commitment of network resources between a number of end points over a shared network infrastructure.

1.1. Rationale

Transport slices are created and managed within the scope of one or more underlying network technologies (e.g., IP, MPLS, optical). Transport slices are expected to enable a diverse set of applications that have different requirements to coexist on the same network infrastructure.

Transport slice is described as a construct that specifies connectivity requirements, emphasizing on assurance of those requirements. Transport slice is unaware of the underlying infrastructure connectivity (hence, the term "transport"). The types of underlying networking technologies can be based on any combination of IP, Ethernet, MPLS, and optical technologies. Transport slices also include specification of resources related to network functions required by customer applications.

Traditionally, VPNs have focussed on segmentation, i.e., creation and management of the private networks. They are bound to a specific traffic type and are technology specific. In contrast, transport slices concern with the assurance of resources required from the network and provide a common user interface for describing those resources. A service provider can use many aspects of the VPNs to build the transport slices.

Transport slices relate to a more general topic of network slicing. It is not the goal of this document to define this broader concept, but in general, it is to identify the methodology to describe the logical (or abstract) partitioning of network resources associated with a service or an application.

2. Terms and Abbreviations

The terms and abbreviations used in this document are listed below.

- o E2E NS: End to End Network Slice
- o TS: Transport Slice

- o TSC: Transport Slice Controller
- o EP: Endpoint
- o EU: End User
- o NBI: NorthBound Interface
- o SBI: SouthBound Interface
- o SLI: Service Level Indicator A well defined quantitative measure of some aspect of the level of service that is provided.
- o SLO: Service Level Objective A target value or range of values for a service level that is measured by an SLI. A natural structure for SLOs is thus $SLI \leq \text{target}$, or lower bound $\leq SLI \leq$ upper bound.
- o SLA: Service Level Agreement An explicit or implicit contract with the end users that includes consequences of meeting (or missing) the SLOs they contain.

The above terminology is described in greater detail in the remainder of this document.

3. Definition and Scope of Transport Slice

The definition of a transport slice is as follows:

"A transport slice is a logical network topology connecting a number of endpoints with a set of shared or dedicated network resources, that are used to satisfy specific Service Level Objectives (SLOs)".

The text below describes transport slices in more details.

Transport slice specification is technology-agnostic, and the means for transport slice realization can be chosen depending on several factors such as: service requirements, specifications or capabilities of underlying infrastructure. The structure and different characteristics of transport slices are described in the following sections.

The term "transport" in transport slice is derived from the definition of Transport Network in the [section 1.3.1 of \[RFC5921\]](#) : A Transport Network provides transparent transmission of user traffic between attached client devices by establishing and maintaining point-to-point or point-to-multipoint connections between such devices. "Slice" refers to a set of characteristics that separate

one type of user-traffic from other types. Transport slice assumes that an underlying transport network is capable of changing the configurations of the network devices on demand, through in-band signaling or via controller(s) and to provide transport transmissions with fulfilling all or some of SLOs to all of the traffic in the slice or to specific flows.

4. Transport Slice System Characteristics

The following subsections describe the characteristics needed for support of transport slices.

4.1. Service Level Objectives for Transport Slices

A transport slice is defined in terms of several quantifiable characteristics or service level objectives (SLOs). These objectives define a set of network resource parameters or values necessary to provide a service as requested for a given transport slice. SLOs do not describe 'how' the transport slices will be implemented or realized in the underlying network layers. Instead, they are defined in terms of dimensions of operations (time, capacity, etc.), availability and other attributes. A transport slice can have one or more SLOs associated with it, all SLO's combined to form an SLA. The SLO values are defined unidirectionally and for specific subsets of two or more endpoints (i.e. for a subset of connections in transport slice).

The SLOs and values associated with them that are exposed to the end user, are in the form of Service Level Indicators (SLIs). If for example the range of latencies a network can provide is 50ms-100ms, then this would be the range of values the end user should be able to request, it would be as low as 50ms or as high as 100ms or anything in between. The values of requested SLOs should always be in the range of values supported. The underlying networks must provide means to monitor and measure the performance of transport slices against the SLOs requested and verify that they are being met. Some SLOs can be measured directly through a collection of metrics and statistics from the network (commonly known as 'telemetry'), while others are deduced from measurable objectives and may require additional tools or mechanisms to measure their target values.

4.1.1. Minimal Set of SLOs

This document defines a minimal set of SLOs and later systems or standards could extend this set and define more SLOs. For example, we included Guaranteed bandwidth which is the minimum requested bandwidth for the transport slice. The later standard might define other SLOs related to bandwidth if needed.

Accordingly, SLOs can be categorized in to 'Directly Measurable Objectives' or 'Indirectly Measurable Objectives' as follows:

Some of the 'Directly Measurable Objectives' are:

- o Guaranteed Minimum Bandwidth
- o Guaranteed Maximum Latency
- o Maximum permissible delay variation
- o Maximum permissible packet loss rate
- o Availability
- o Other objectives could be specified

Some of the 'Indirectly Measurable Objectives' are:

- o Security
- o others objectives such as geographical restrictions, maximum occupancy level, etc. could be specified

The definition of these objectives are as follows:

- o Guaranteed Minimum Bandwidth: Minimum guaranteed bandwidth between two endpoints at any time. The bandwidth is measured in data rate units of bits per second and is measured unidirectionally.
- o Guaranteed Maximum Latency: Upper bound of network latency when transmitting between two endpoints. The latency is measured in terms of network characteristics (excluding application-level latency). [[RFC2681](#)] and [[RFC7679](#)] discuss round trip times and one-way metrics, respectively.
- o Maximum permissible delay variation: Packet delay variation (PDV) as defined by [[RFC3393](#)], is measured by the difference in the one-way delay between sequential packets in a flow. Minimizing variations in the delay is important for real-time applications.
- o Maximum permissible packet loss rate: is defined by the ratio of packets dropped to packets transmitted between two endpoints. See [[RFC7680](#)]
- o Availability: is defined as the ratio of uptime to $\text{total_time(uptime+downtime)}$, where uptime is the time the

transport slice is available in accordance with the SLOs associated with it.

- o Security: This objective may request for encryption [[RFC4303](#)] between two end-points explicitly to meet architecture recommendations as in [[TS33.210](#)] or for compliance with [[HIPAA](#)] [[PCI](#)]. Other security requests may be made as specified in [[draft-ietf-i2nsf-capability](#)].
- * Note: Security violations are not directly observable and cannot be measured as quantifiable metrics. Still, the user of the transport slice should be able to request certain criteria for compliance and identify exceptions and unexpected traffic. For this purpose [[i2nsf-nsf-monitoring-data-model](#)] can be leveraged.

[4.1.2.](#) Other Objectives

Additional objectives, such as certain geographical restrictions or well defined domains that a slice may transit may be necessary.

Optionally, when the customer is traffic aware, other traffic specific characteristics may be provided. These include for example, MTU, traffic-type (e.g., IPv4, IPv6, Ethernet or unstructured), or a higher-level behavior to process traffic according to user-application (which may be realized using network functions).

Maximal occupancy for a transport slice should be provided. Since it carries traffic for multiple flows between the two endpoints, the objectives should also say if they are for the entire connection, group of flows or on per flow basis. Maximal occupancy should specify the scale of the flows (i.e. maximum number of accommodatable flows) and optionally a maximum number of countable resource units, e.g IP or MAC addresses a slice might consume.

With these objectives incorporated, a customer sees transport slice as a dedicated network for its exclusive use. Achieving this may require different types of isolation techniques in provider networks as described in [Appendix A.1](#).

Additional description of slice attributes is covered in a broader context of 'Generic Network Slice Template' in [[I-D.contreras-teas-slice-nbi](#)].

4.2. Transport Slice Endpoints

The transport slice endpoints are the conceptual entities that perform any required conversion, or adaptation, and forwarding of the user traffic. The characteristics of the transport slice endpoints (TSE) are:

- o They are conceptual points of connection of a network function, device or application to the transport slice
- o They are identified in a request provided by the customer of transport slice (i.e. higher level operation systems) during the creation of the transport slice
- o They are associated with a device, application and/or network function nodes. A non-exhaustive list of such nodes are routers, switches, firewalls, WAN, 4G/5G RAN nodes, 4G/5G Core nodes, application acceleration, Deep Packet Inspection (DPI), server load balancers, NAT44 [[RFC3022](#)], NAT64 [[RFC6146](#)], HTTP header enrichment functions, and TCP optimizers
- o A TSE is identified by its associated node (its IP address, name , ID, etc.), a unique identifier and/or a unique name and other data. A non-exhaustive list of other data includes IP address (v4 or v6), VLAN, port, connectivity type (P2P, P2MP, MP2MP). TBD for more


Note that the TSE is different from access points (AP) defined in [[RFC8453](#)] as an AP is a logical identifier to identify the shared link between the customer and the operator where as TSE is an identifier of an endpoint. Also TSE is different from TE Link Termination Point (LTP) defined in [[I-D.ietf-teas-yang-te-topo](#)] as it is a conceptual point of connection of a TE node to one of the TE links on a TE node.

The TSE is similar to the Termination Point (TP) defined in [[RFC8345](#)] and can contain more attributes. TSE could be modeled by augmenting the TP model.

There is another type of the endpoints called "Transport Slice Realization endpoints (TSREs)". These endpoints are allocated and assigned by the network controller during the realization of a transport slice and are technology-specific, i.e. they depend on the network technology used during the transport slice realization. They are identified by a node and some associated data. A non-exhaustive list of nodes containing TSREs are routers, switches, PON nodes, Wireless nodes and Optical devices.

Note that there will be a mapping between TSE and TSRE on Transport Slice Controller (TSC). When TSC receives a request via its NBI to create a transport slice between multiple TSEs, it will send the request via its SBI to realize the transport slice. The TSRE will be notified by network controller during TS realization to enable mapping between TSREs and the TSEs.

Figure 1 shows an example of a transport slice and its realization between multiple TSEs and TSREs.



Legend:

DAN: Device, application and/or network function

Figure 1: A transport slice between TSEs and its realization between TSREs

[4.2.1. Transport Slice Connectivity Types](#)


The transport slices connection types can be point to point (P2P), point to multipoint (P2MP), multi-point to point (MP2P), or multi-point to multi-point (MP2MP). The transport slice connection type will requested by the higher level operation system.

[4.3. Vertical Composition of Transport Slice](#)

Transport slice may follow a hierarchical relationship to provide a vertical structure to it. This is used for composing multi-layer slices in which each layer provides an abstraction, as well as an independent monitoring, performance, control and management of the

resources. The vertical transport slice characteristic could be used in 2 forms:

- o The Transport slice itself where it represents a hierarchy of abstracted transport slices. In this case, the realization will be done just once with a particular technology. Thus, the lowest transport slice in the hierarchy that can not be decomposed further will be one to one mapping to its instance of the realized transport slice.
- o Each layer (physical, datalink, or IP) has its own set of resources that can be provided to the upper layer as a transport slice. Thus, transport slice at one layer is used by the layer above. This type of multi-layer vertical transport slice associates resources at different layers. For example, an IP transport slice would utilize one or more optical transport slice. In this case, the realization will be done for a particular technology at that particular layer. Thus, the lowest transport slice in this type of hierarchy that can not be decomposed further will be an instance of realized physical layer transport slice.



Legend:

TSnnn: Level 3 vertical transport slice nnn

TSnn: Level 2 vertical transport slice nn

TSn: Level 1 transport Slice n

Figure 2: Transport Slice Vertical and Horizontal Composition

Figure 2 shows the transport slice hierarchy. Slices TS11 and TS12 are composed together to form TS1 that is the top level transport slice definition, TS121 and TS122 collectively define TS12. The SLO for bandwidth guarantee will be shared and latency guarantee will be split into latency in networks 2 and 3. To emphasize the hierarchical structure, consider Network-2 and Network-3 are in the same administrative domain but use different transport technologies respectively. Then instead of presenting 2 transport slices,

Operator-z can expose only one transport slice TS12 abstracting the underlying transport technology details.

Note: The specification to connect TS121 and TS122 are similar to those connecting TS12 and TS11.

4.4. Horizontal Composition of Transport Slice

In contrast, horizontal transport slices enable the composition of multiple realized transport slices. Since transport slices are not necessarily a single encapsulation tunnel and may traverse through different data planes, each realized transport slice will require a stitching, interworking or mapping function. These stitching functions can be viewed as a type of intermediate network function endpoints. For instance in Figure 2, TS11 and TS12 are horizontal transport slices. If we assume that TS11 is an L2 tunnel and TS12 is an SRV6 based path, then a 'Service type EP' (not shown in the figure) is needed for translation.

Author's notes: This service type EP is a new type of transport slice specific service function. We may call it transport slice gateway.

5. Transport Slice Structure

A transport slice is a set of connections among various endpoints to form a logical network that meets the SLOs agreed upon.

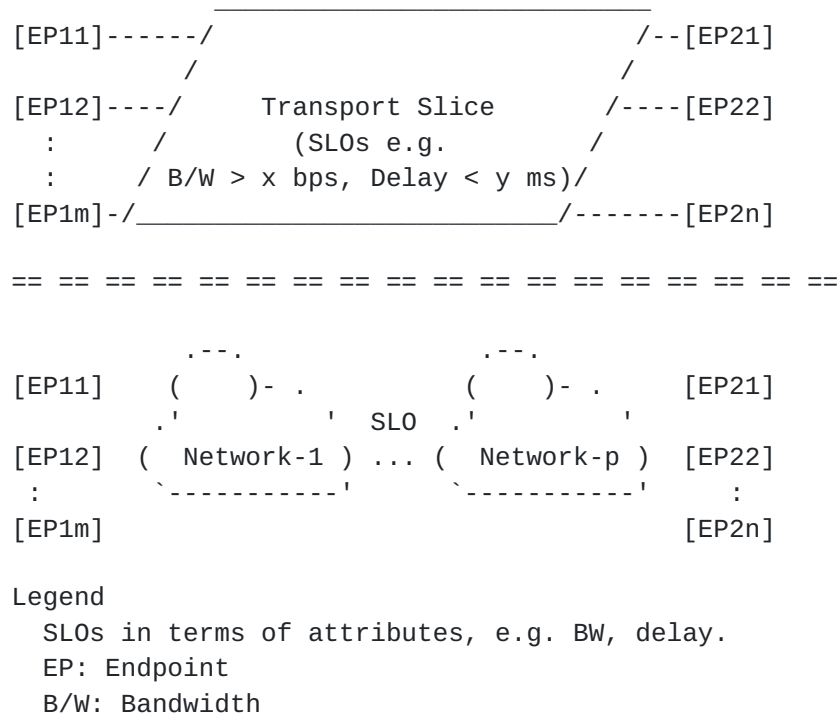


Figure 3: Transport slice

Figure 3 illustrates a case where a transport slice provides connectivity between a set of endpoints pairs with specific characteristics for each SLO (e.g. guaranteed minimum bandwidth of x bps and guaranteed delay of less than y ms). The endpoints may be distributed in the underlay networks, and a transport slice can be deployed across multiple network domains. Also, the endpoints on the same transport slice may belong to the same address space.

Transport slices involve both customer's and provider's views. A customer 'describes' its requirements in terms of connectivity with specific SLOs. Provider networks address those requirements through 'transport slice realization' (its implementation) using provider network specific technologies.

A transport slice is requested from an entity (such as an orchestrator or a system-wide controller) performing broader service or application specific functions. The interface from such an entity should express the needed connectivity in a technology-agnostic way and donot need to recognize configurations based on the technologies (e.g. being more declarative than imperative). The request to instantiate a transport slice is only represented with some indicators such as SLOs based on which the underlying technologies are selected and managed.

Often, in other SDOs the term sub-slice or slice-subnet comes up. Some of those are mapped to transport network requirements in the form of a transport slice. With in the scope of transport slices (w.r.t. the IP/MPLS based transport networks) there are no definitions for 'sub-slice' or 'slice subnets'. 'Transport slice' term universally represents SLO and connectivity requirements from the transport networks.

Furthermore, the structure of transport slices may be layered vertically or composed horizontally, i.e. operationally, a transport slice maybe decomposed in two or more transport slices which are then independently realized and managed. This is further described in [Section 4.3](#).

[5.1](#). Stakeholders

A transport slice and its realization involves the following stakeholders and it is relevant to define them for consistent terminology.

Customer or User: A customer is a user of a transport slice. Customers may request monitoring of associated resources or specific changes. A user may either directly manage its service by interfacing with the transport slice controller or indirectly through an orchestrator.

Orchestrator: An orchestrator is an entity that composes different services, resource and network requirements. It interfaces with the transport slice controllers.

Transport Slice Controller (TSC): It realizes a transport slice in the network, maintains and monitors the run-time state of resources and topologies associated with it. A well-defined interface is needed between different types of transport slice controllers and different types of orchestrators. A transport slice operator (or slice operator for short) manages one or more transport slices using the Transport Slice Controller(s).

Transport Network Controller: is a form of network infrastructure controller that offers network resources to TSC to realize a particular transport slice. These may be existing network controllers associated with one or more specific technologies that may be adapted to the function of realizing transport slices in a network.

5.2. Transport Slice Controller Interfaces

The interworking and interoperability among the different stakeholders to provide common means of provisioning, operating and monitoring the transport slices is a mandatory requirement. The following communication interfaces are identified (see Figure 4).

TSC Northbound Interface (NBI): The TSC Northbound Interface is an interface between a higher level operation system, e.g. 'E2E network slice orchestrator' and the 'Transport slice controller'. It is a technology agnostic interface. Over this NBI, slice characteristics and other requirements can be communicated to TSC and the operational state of a transport slice may be requested.

TSC Southbound Interface (SBI): The TSC Southbound Interface is an interface between 'Transport slice controller (TSC)' and network controller(s). These interfaces are technology-specific and utilize many of the network models.




Figure 4: Interface of Transport Slice Controller

5.3. Transport slice Realization

Realization of a Transport Slice is a mapping of underlying infrastructure with its definition. It is a technology specific entity that is created and maintained over its southbound interfaces. The Network controller(s) export the connectivity and resource mappings to the TSC. The network controller abstracts the details of underlying resources from the TSC.

The realization can be achieved in the form of either physical or logical connectivity through VPNs, a variety of tunneling technologies such as Segment Routing, SFC, etc. Accordingly, endpoints may be realized as physical or logical service or network functions.

6. Relationship with End-to-End Network Slicing

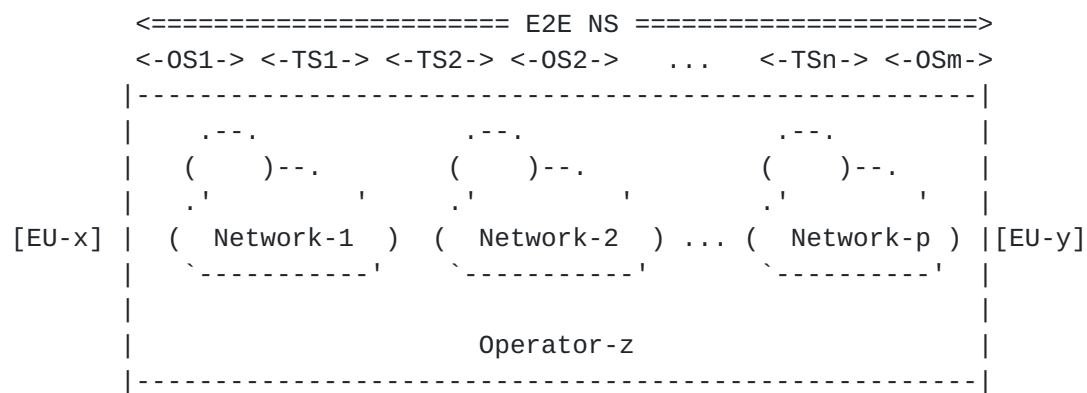
An end-to-end (E2E) network slice is a complete logical network that provides a service in its entirety with a specific assurance to the customer. A transport slice concerns with those assurance aspects only within the transport networks. Consider Figure 5, where a network operator has an E2E network slice that traverses multiple technology-specific networks. Each of these networks might use any number of technologies, including but not limited to IP, MPLS, Fiber-Optics (e.g. WDM, DWDM), Passive Optical Networking (PON), Microwave, etc.

Each of these networks includes multiple (physical or virtual) nodes and may also provide network functions beyond simply carrying of technology-specific protocol data units. The types of nodes used in any of these networks may include:

- o Packet/frame processing nodes (e.g., Routers, Switches)
- o Application servers
- o Service Functions(e.g., Firewall, Loadbalancer)
- o Radio Access Network (RAN) components
- o Mobile Core components
- o Microwave transceivers
- o Optical repeaters
- o etc.

Each network may support different technologies and an E2E network slice is a combination of these networks. As an example:

- o Network 1 might contain multiple 5G RAN nodes connected to a few Cell Site Gateways (CSG) routers.
- o Network 2 might have one or more layer-3 routers and layer-2 switches which may run on top of an optical network.
- o Network 3 might have a number of 5G RAN nodes connected to Passive Optical Network (PON) switches.



Legend:

E2E NS: End-to-end network slice
 TSn: Transport Slice n
 OSm: Other Slice m
 EU-x: End User-x
 EU-y: End User-y

Figure 5: E2E network slice

When operator-z creates a specific E2E network slice, it may create one or more of transport slices and other slices (application logic or other system functions).

An independent E2E logical network (called E2E network slice) is created for a service (e.g. CCTV, autonomous driving, HD map, etc.) with a specific network SLOs, e.g. a secure connection with an E2E latency less than 5ms, from End User-x (EU-x) to End User-y (EU-y). EU-x maybe a 5G user equipment such as an infotainment unit in a car, CCTV, or a car for autonomous driving, etc. and EU-y in 5G is 5G application server, IMS, etc.

In Figure 5, "E2E NS" is that logical network with requested SLO between EU-x to EU-y and is associated with a customer and a specific service type.

7. Security Considerations

Not applicable in this memo.

8. IANA Considerations

This memo includes no request to IANA.

9. Acknowledgment

The entire TEAS NS design team and everyone participating in those discussion has contributed to this draft. Particularly, Eric Gray, Xufeng Liu, Jie Dong, and Jari Arkko for a thorough review among other contributions.

10. Informative References

[HIPAA] HHS, "Health Insurance Portability and Accountability Act - The Security Rule", February 2003, <<https://www.hhs.gov/hipaa/for-professionals/security/index.html>>.

[I-D.contreras-teas-slice-nbi]
Contreras, L., Homma, S., and J. Ordonez-Lucena, "Considerations for defining a Transport Slice NBI", [draft-contreras-teas-slice-nbi-01](#) (work in progress), March 2020.

[I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Services", [draft-ietf-teas-enhanced-vpn-05](#) (work in progress), February 2020.

[I-D.ietf-teas-sf-aware-topo-model]
Bryskin, I., Liu, X., Lee, Y., Guichard, J., Contreras, L., Ceccarelli, D., and J. Tantsura, "SF Aware TE Topology YANG Model", [draft-ietf-teas-sf-aware-topo-model-05](#) (work in progress), March 2020.

- [I-D.ietf-teas-yang-te-topo]
Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", [draft-ietf-teas-yang-te-topo-22](#) (work in progress), June 2019.
- [I-D.nsdt-teas-ns-framework]
Gray, E. and J. Drake, "Framework for Transport Network Slices", [draft-nsdt-teas-ns-framework-02](#) (work in progress), March 2020.
- [NFVGST] ETSI, "NFVI Compute and Network Metrics Specification", February 2018, <https://www.etsi.org/deliver/etsi_gs/NFV-TST/001_099/008/02.04.01_60/gs_nfv-tst008v020401p.pdf>.
- [PCI] PCI Security Standards Council, "PCI DSS", May 2018, <<https://www.pcisecuritystandards.org>>.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", [RFC 2681](#), DOI 10.17487/RFC2681, September 1999, <<https://www.rfc-editor.org/info/rfc2681>>.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), DOI 10.17487/RFC3022, January 2001, <<https://www.rfc-editor.org/info/rfc3022>>.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", [RFC 3393](#), DOI 10.17487/RFC3393, November 2002, <<https://www.rfc-editor.org/info/rfc3393>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", [RFC 5921](#), DOI 10.17487/RFC5921, July 2010, <<https://www.rfc-editor.org/info/rfc5921>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.

- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, [RFC 7679](#), DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.
- [RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, [RFC 7680](#), DOI 10.17487/RFC7680, January 2016, <<https://www.rfc-editor.org/info/rfc7680>>.
- [RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", [RFC 8345](#), DOI 10.17487/RFC8345, March 2018, <<https://www.rfc-editor.org/info/rfc8345>>.
- [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>>.
- [TS.23.501-3GPP]
3rd Generation Partnership Project (3GPP), "3GPP TS 23.501 (V16.2.0): System Architecture for the 5G System (5GS); Stage 2 (Release 16)", September 2019, <http://www.3gpp.org/ftp//Specs/archive/23_series/23.501/23501-g20.zip>.
- [TS33.210]
3GPP, "3G security; Network Domain Security (NDS); IP network layer security (Release 14).", December 2016, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2279>>.

[Appendix A](#). Discussions

[A.1](#). On Isolation Requirements In a Transport Slice

Transport slices are perceived as if slice was provisioned for the customer as a dedicated network with specific SLOs. These committed SLOs for a given customer should be maintained during the lifetime of the slice, even in the face of potential disruptions. Such disruptions include sudden traffic volume changes either from the customer itself or others, equipment failures in the service provider network, and various misbehaviors or attacks.

The service provider needs to ensure that its network can provide the requested slices with the availability agreed with its customers. Some of the main technical approaches to ensuring guarantees are

about network planning, managing capacity, prioritizing, policing or shaping customer traffic, selecting dedicated resources, and so on.

One term that has commonly been used in this context is "isolation" and is also discussed in the [[I-D.ietf-teas-enhanced-vpn](#)].

A transport slice customer may ask for traffic separation, selection of dedicated resources, or interference avoidance from other traffic. The term "isolation" can refer to any or all of them. For instance, dedicated resources can help assure that traffic in other slices does not affect a given slice. Similarly, VPN technologies can provide traffic separation, and interference avoidance may be provided by mechanisms such as technical approaches mentioned in the previous paragraph (network planning, capacity management, etc). Moreover, these are some of the examples of a particular realization of the requirement for guarantees; other mechanisms may also be used.

Authors' Addresses

Reza Rokui
Nokia
Canada

Email: reza.rokui@nokia.com

Shunsuke Homma
NTT
Japan

Email: shunsuke.homma.ietf@gmail.com

Kiran Makhijani
Futurewei
USA

Email: kiranm@futurewei.com

Luis M. Contreras
Telefonica
Spain

Email: luismiguel.contrerasmurillo@telefonica.com

Jeff Tantsura
Apstra, Inc.

Email: jefftant.ietf@gmail.com