

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 24, 2021

A. Farrel, Ed.
Old Dog Consulting
E. Gray
Independent
J. Drake
Juniper Networks
R. Rokui
Nokia
S. Homma
NTT
K. Makhijani
Futurewei
LM. Contreras
Telefonica
J. Tantsura
Juniper Networks
May 23, 2021

Framework for IETF Network Slices
draft-ietf-teas-ietf-network-slices-03

Abstract

This document describes network slicing in the context of networks built from IETF technologies. It defines the term "IETF Network Slice" and establishes the general principles of network slicing in the IETF context.

The document discusses the general framework for requesting and operating IETF Network Slices, the characteristics of an IETF Network Slice, the necessary system components and interfaces, and how abstract requests can be mapped to more specific technologies. The document also discusses related considerations with monitoring and security.

This document also provides definitions of related terms to enable consistent usage in other IETF documents that describe or use aspects of IETF Network Slices.

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 November 24, 2021.

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
1.1. Background	4
2. Terms and Abbreviations	5
2.1. Core Terminology	6
3. IETF Network Slice Objectives	6
3.1. Definition and Scope of IETF Network Slice	6
4. IETF Network Slice System Characteristics	7
4.1. Objectives for IETF Network Slices	7
4.1.1. Service Level Objectives	8
4.1.2. Service Level Expectations	10
4.2. IETF Network Slice Endpoints	12
4.2.1. IETF Network Slice Connectivity Types	14
4.3. IETF Network Slice Decomposition	14
5. Framework	15
5.1. IETF Network Slice Stakeholders	15
5.2. Expressing Connectivity Intents	15
5.3. IETF Network Slice Controller (NSC)	17
5.3.1. IETF Network Slice Controller Interfaces	19
5.3.2. Northbound Interface (NBI)	20
5.4. IETF Network Slice Structure	21
6. Realizing IETF Network Slices	22

6.1. Procedures to Realize IETF Network Slices	22
6.2. Applicability of ACTN to IETF Network Slices	23
6.3. Applicability of Enhanced VPNs to IETF Network Slices . .	23
6.4. Network Slicing and Slice Aggregation in IP/MPLS Networks	24
7. Isolation in IETF Network Slices	24
7.1. Isolation as a Service Requirement	24
7.2. Isolation in IETF Network Slice Realization	24
8. Management Considerations	25
9. Security Considerations	25
10. Privacy Considerations	26
11. IANA Considerations	26
12. Informative References	26
Acknowledgments	30
Contributors	31
Authors' Addresses	31

1. Introduction

A number of use cases benefit from network connections that along with the connectivity provide assurance of meeting a specific set of objectives with respect to network resources use. This connectivity and resource commitment is referred to as a network slice. Since the term network slice is rather generic, the qualifying term "IETF" is used in this document to limit the scope of network slice to network technologies described and standardized by the IETF. This document defines the concept of IETF Network Slices that provide connectivity coupled with a set of specific commitments of network resources between a number of endpoints over a shared network infrastructure. Services that might benefit from IETF Network Slices include, but are not limited to:

- o 5G services (e.g. eMBB, URLLC, mMTC) (See [TS23501])
- o Network wholesale services
- o Network infrastructure sharing among operators
- o NFV connectivity and Data Center Interconnect

IETF Network Slices are created and managed within the scope of one or more network technologies (e.g., IP, MPLS, optical). They are intended to enable a diverse set of applications that have different requirements to coexist on the shared network infrastructure. A request for an IETF Network Slice is technology-agnostic so as to allow a customer to describe their network connectivity objectives in a common format, independent of the underlying technologies used.

This document also provides a framework for discussing IETF Network Slices. This framework is intended as a structure for discussing interfaces and technologies. It is not intended to specify a new set of concrete interfaces or technologies. Rather, the idea is that existing or under-development IETF technologies (plural) can be used to realize the concepts expressed herein.

For example, virtual private networks (VPNs) have served the industry well as a means of providing different groups of users with logically isolated access to a common network. The common or base network that is used to support the VPNs is often referred to as an underlay network, and the VPN is often called an overlay network. An overlay network may, in turn, serve as an underlay network to support another overlay network.

Note that it is conceivable that extensions to these IETF technologies are needed in order to fully support all the ideas that can be implemented with slices. Evaluation of existing technologies, proposed extensions to existing protocols and interfaces, and the creation of new protocols or interfaces is outside the scope of this document.

1.1. Background

Driven largely by needs surfacing from 5G, the concept of network slicing has gained traction ([NGMN-NS-Concept], [TS23501], [TS28530], and [BBF-SD406]). In [TS23501], a Network Slice is defined as "a logical network that provides specific network capabilities and network characteristics", and a Network Slice Instance is defined as "A set of Network Function instances and the required resources (e.g. compute, storage and networking resources) which form a deployed Network Slice." According to [TS28530], an end-to-end network slice consists of three major types of network segments: Radio Access Network (RAN), Transport Network (TN) and Core Network (CN). An IETF Network Slice provides the required connectivity between different entities in RAN and CN segments of an end-to-end network slice, with a specific performance commitment. For each end-to-end network slice, the topology and performance requirement on a customer's use of IETF Network Slice can be very different, which requires the underlay network to have the capability of supporting multiple different IETF Network Slices.

While network slices are commonly discussed in the context of 5G, it is important to note that IETF Network Slices are a narrower concept, and focus primarily on particular network connectivity aspects. Other systems, including 5G deployments, may use IETF Network Slices as a component to create entire systems and concatenated constructs that match their needs, including end-to-end connectivity.

A IETF Network Slice could span multiple technologies and multiple administrative domains. Depending on the IETF Network Slice customer's requirements, an IETF Network Slice could be isolated from other, often concurrent IETF Network Slices in terms of data, control and management planes.

The customer expresses requirements for a particular IETF Network Slice by specifying what is required rather than how the requirement is to be fulfilled. That is, the IETF Network Slice customer's view of an IETF Network Slice is an abstract one.

Thus, there is a need to create logical network structures with required characteristics. The customer of such a logical network can require a degree of isolation and performance that previously might not have been satisfied by traditional overlay VPNs. Additionally, the IETF Network Slice customer might ask for some level of control of their virtual networks, e.g., to customize the service paths in a network slice.

This document specifies definitions and a framework for the provision of an IETF Network Slice service. Section 6 briefly indicates some candidate technologies for realizing IETF Network Slices.

2. Terms and Abbreviations

The following abbreviations are used in this document.

- o NBI: NorthBound Interface
- o NSC: Network Slice Controller
- o NSE: Network Slice Endpoint
- o SBI: SouthBound Interface
- o SLA: Service Level Agreement
- o SLI: Service Level Indicator
- o SLO: Service Level Objective

The meaning of these abbreviations is defined in greater details in the remainder of this document.

2.1. Core Terminology

The following terms are presented here to give context. Other terminology is defined in the remainder of this document.

Customer: A customer is the requester of an IETF Network Slice service. Customers may request monitoring of SLOs. A customer may be an entity such as an enterprise network or a network operator, an individual working at such an entity, a private individual contracting for a service, or an application or software component. A customer may be an external party (classically a paying customer) or a division of a network operator that uses the service provided by another division of the same operator. Other terms that have been applied to the customer role are "client" and "consumer".

Provider: A provider is the organization that delivers an IETF Network Slice service. A provider is the network operator that controls the network resources used to construct the network slice (that is, the network that is sliced). The provider's network maybe a physical network or may be a virtual network supplied by another service provider.

3. IETF Network Slice Objectives

It is intended that IETF Network Slices can be created to meet specific requirements, typically expressed as bandwidth, latency, latency variation, and other desired or required characteristics. Creation is initiated by a management system or other application used to specify network-related conditions for particular traffic flows.

It is also intended that, once created, these slices can be monitored, modified, deleted, and otherwise managed.

It is also intended that applications and components will be able to use these IETF Network Slices to move packets between the specified end-points in accordance with specified characteristics.

3.1. Definition and Scope of IETF Network Slice

The definition of a network slice in IETF context is as follows:

An IETF Network Slice is a logical network topology connecting a number of endpoints using a set of shared or dedicated network resources that are used to satisfy specific Service Level Objectives (SLOs).

An IETF Network Slice combines the connectivity resource requirements and associated network behaviors such as bandwidth, latency, jitter, and network functions with other resource behaviors such as compute and storage availability. IETF Network Slices are independent of the underlying infrastructure connectivity and technologies used. This is to allow an IETF Network Slice service customer to describe their network connectivity and relevant objectives in a common format, independent of the underlying technologies used.

IETF Network Slices may be combined hierarchically, so that a network slice may itself be sliced. They may also be combined sequentially so that various different networks can each be sliced and the network slices placed into a sequence to provide an end-to-end service. This form of sequential combination is utilized in some services such as in 3GPP's 5G network [TS23501].

An IETF Network Slice is technology-agnostic, and the means for IETF Network 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 IETF Network Slices are described in the following sections.

Term "Slice" refers to a set of characteristics and behaviours that separate one type of user-traffic from another. IETF Network Slice assumes that an underlying network is capable of changing the configurations of the network devices on demand, through in-band signaling or via controller(s) and fulfilling all or some of SLOs to all of the traffic in the slice or to specific flows.

4. IETF Network Slice System Characteristics

The following subsections describe the characteristics of IETF Network Slices.

4.1. Objectives for IETF Network Slices

An IETF Network Slice service is defined in terms of quantifiable characteristics known as Service Level Objectives (SLOs) and unquantifiable characteristics known as Service Level Expectations (SLEs). SLOs are expressed in terms Service Level Indicators (SLIs), and together with the SLEs form the contractual agreement between service customer and service provider known as a Service Level Agreement (SLA).

The terms are defined as follows:

- o A Service Level Indicator (SLI) is a quantifiable measure of an aspect of the performance of a network. For example, it may be a measure of throughput in bits per second, or it may be a measure of latency in milliseconds.
- o A Service Level Objective (SLO) is a target value or range for the measurements returned by observation of an SLI. For example, an SLO may be expressed as "SLI <= target", or "lower bound <= SLI <= upper bound". A customer can determine whether the provider is meeting the SLOs by performing measurements on the traffic.
- o A Service Level Expectation (SLE) is an expression of an unmeasurable service-related request that a customer of an IETF network slice makes of the provider. An SLE is distinct from an SLO because the customer may have little or no way of determining whether the SLE is being met, but they still contract with the provider for a service that meets the expectation.
- o A Service Level Agreement (SLA) is an explicit or implicit contract between the customer of an IETF Network Slice and the provider of the slice. The SLA is expressed in terms of a set of SLOs and SLEs that are to be applied to the connections between the service endpoints, and may include commercial terms as well as the consequences of missing/violating the SLOs they contain.

4.1.1. Service Level Objectives

SLOs define a set of network attributes and characteristics that describe an IETF Network Slice. SLOs do not describe how the IETF Network Slices are implemented or realized in the underlying network layers. Instead, they are defined in terms of dimensions of operation (time, capacity, etc.), availability, and other attributes. An IETF Network Slice can have one or more SLOs associated with it. The SLOs are combined in an SLA. The SLOs are defined for sets of two or more endpoints and apply to specific directions of traffic flow. That is, they apply to specific source endpoints and specific connections between endpoints within the set of endpoints and connections in the IETF Network Slice.

SLOs define a set of measurable network attributes and characteristics that describe an IETF Network Slice service. SLOs do not describe how the IETF network slices are implemented or realized in the underlying network layers. Instead, they are defined in terms of dimensions of operation (time, capacity, etc.), availability, and other attributes. An IETF Network Slice service can have one or more SLOs associated with it. The SLOs are combined with Service Level Expectations in an SLA.

An IETF network slice service may include multiple connection constructs that associate sets of endpoints. SLOs apply to sets of two or more endpoints and apply to specific directions of traffic flow. That is, they apply to a specific source endpoint and the connection to specific destination endpoints.

4.1.1.1. Some Common SLOs

SLOs can be described as 'Directly Measurable Objectives': they are always measurable. See Section 4.1.2 for the description of Service Level Expectations which are unmeasurable service-related requests sometimes known as 'Indirectly Measurable Objectives'.

Objectives such as guaranteed minimum bandwidth, guaranteed maximum latency, maximum permissible delay variation, maximum permissible packet loss rate, and availability are 'Directly Measurable Objectives'. Future specifications (such as IETF Network Slice service YANG models) may precisely define these SLOs, and other SLOs may be introduced as described in Section 4.1.1.2.

The definition of these objectives are as follows:

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.

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.

Maximum Permissible Delay Variation

Packet delay variation (PDV) as defined by [RFC3393], is the difference in the one-way delay between sequential packets in a flow. This SLO sets a maximum value PDV for packets between two endpoints.

Maximum Permissible Packet Loss Rate

The ratio of packets dropped to packets transmitted between two endpoints over a period of time. See [RFC7680].

Availability

The ratio of uptime to the sum of uptime and downtime, where uptime is the time the IETF Network Slice is available in accordance with the SLOs associated with it.

4.1.1.2. Other Service Level Objectives

Additional SLOs may be defined to provide additional description of the IETF Network Slice service that a customer requests. These would be specified in further documents.

If the IETF network slice service is traffic aware, other traffic specific characteristics may be valuable including 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).

4.1.2. Service Level Expectations

SLEs define a set of network attributes and characteristics that describe an IETF Network Slice service, but which are not directly measurable by the customer. Even though the delivery of an SLE cannot usually be determined by the customer, the SLEs form an important part of the contract between customer and provider.

Quite often, an SLE will imply some details of how an IETF Network Slice service is realized by the provider, although most aspects of the implementation in the underlying network layers remain a free choice for the provider.

SLEs may be seen as aspirational on the part of the customer, and they are expressed as behaviors that the provider is expected to apply to the network resources used to deliver the IETF Network Slice service. An IETF network slice service can have one or more SLEs associated with it. The SLEs are combined with SLOs in an SLA.

An IETF Network Slice service may include multiple connection constructs that associate sets of endpoints. SLEs apply to sets of two or more endpoints and apply to specific directions of traffic flow. That is, they apply to a specific source endpoint and the connection to specific destination endpoints. However, being more general in nature, SLEs may commonly be applied to all connection constructs in an IETF Network Slice service.

4.1.2.1. Some Common SLEs

SLEs can be described as 'Indirectly Measurable Objectives': they are not generally directly measurable by the customer.

Security, geographic restrictions, maximum occupancy level, and isolation are example SLEs as follows.

Security

A customer may request that the provider applies encryption or other security techniques to traffic flowing between endpoints of an IETF Network Slice service. For example, the customer could request that only network links that have MACsec [MACsec] enabled are used to realize the IETF Network Slice service.

This SLE may include the request for encryption (e.g., [RFC4303]) between the two endpoints explicitly to meet architecture recommendations as in [TS33.210] or for compliance with [HIPAA] or [PCI].

Whether or not the provider has met this SLE is generally not directly observable by the customer and cannot be measured as a quantifiable metric.

Please see further discussion on security in Section 9.

Geographic Restrictions

A customer may request that certain geographic limits are applied to how the provider routes traffic for the IETF Network Slice service. For example, the customer may have a preference that its traffic does not pass through a particular country for political or security reasons.

Whether or not the provider has met this SLE is generally not directly observable by the customer and cannot be measured as a quantifiable metric.

Maximal Occupancy Level

The maximal occupancy level specifies the number of flows to be admitted and optionally a maximum number of countable resource units (e.g., IP or MAC addresses) an IETF network slice service can consume. Since an IETF Network Slice service may include multiple connection constructs, this SLE should also say whether it applies for the entire IETF Network Service slice, for group of connections, or on a per connection basis.

Again, a customer may not be able to fully determine whether this SLE is being met by the provider.

Isolation

As described in Section 7, a customer may request that its traffic within its IETF Network Slice service is isolated from the effects of other network services supported by the same provider. That is, if another service exceeds capacity or has a burst of traffic, the customer's IETF Network Slice service should remain unaffected and there should be no noticeable change to the quality of traffic delivered.

In general, a customer cannot tell whether a service provider is meeting this SLE. They cannot tell whether the variation of an SLI is because of changes in the underlying network or because of interference from other services carried by the network. And if the service varies within the allowed bounds of the SLOs, there may be no noticeable indication that this SLE has been violated.

Diversity

A customer may request that traffic on the connection between one set of endpoints should use different network resources from the traffic between another set of endpoints. This might be done to enhance the availability of the IETF Network Slice service.

While availability is a measurable objective (see Section 4.1.1.1) this SLE requests a finer grade of control and is not directly measurable (although the customer might become suspicious if two connections fail at the same time).

4.2. IETF Network Slice Endpoints


As noted in Section 3.1, an IETF Network Slice describes connectivity between multiple endpoints across the underlying network. These connectivity types are: point-to-point, point-to-multipoint, multipoint-to-point, or multipoint-to-multipoint.

Figure 1 shows an IETF Network Slice along with its Network Slice Endpoints (NSEs).

The characteristics of IETF NSEs are as follows:

- o The IETF NSEs are conceptual points of connection to IETF network slice. As such, they serve as the IETF Network Slice ingress/egress points.
- o Each endpoint could map to a device, application or a network function. A non-exhaustive list of devices, applications or network functions might include but not limited to: 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 An NSE should be identified by a unique ID in the context of an IETF Network Slice customer.
- o In addition to an identifier, each NSE should contain a subset of attributes such as IPv4/IPv6 addresses, encapsulation type (i.e., VLAN tag, MPLS Label etc.), interface/port numbers, node ID etc.
- o A combination of NSE unique ID and NSE attributes defines an NSE in the context of the IETF Network Slice Controller (NSC).
- o During the realization of the IETF Network Slice, in addition to SLOs, all or subset of IETF NSE attributes will be utilized by the IETF NSC to find the optimal realization in the IETF network.
- o Similarly to IETF Network Slices, the IETF Network Slice Endpoints are logical entities that are mapped to services/tunnels/paths endpoints in IETF Network Slice during its initialization and realization.

Note that there are various IETF TE terms such as access points (AP) defined in [RFC8453], Termination Point (TP) defined in [RFC8345], and Link Termination Point (LTP) defined in [RFC8795] which are tightly coupled with TE network type and various realization techniques. At the time of realization of the IETF Network Slice, the NSE could be mapped to one or more of these based on the network slice realization technique in use.



Legend:

NSE: IETF Network Slice Endpoint

O: Represents IETF Network Slice Endpoints

Figure 1: An IETF Network Slice Endpoints (NSE)

4.2.1. IETF Network Slice Connectivity Types

The IETF Network Slice connection types can be point to point (P2P), point-to-multipoint (P2MP), multipoint-to-point (MP2P), or multipoint-to-multipoint (MP2MP). They will be requested by the higher level operation system.

4.3. IETF Network Slice Decomposition

Operationally, an IETF Network Slice may be decomposed into two or more IETF Network Slices as specified below. Decomposed network slices are then independently realized and managed.

- o Hierarchical (i.e., recursive) composition: An IETF Network Slice can be further sliced into other network slices. Recursive composition allows an IETF Network Slice at one layer to be used by the other layers. This type of multi-layer vertical IETF Network Slice associates resources at different layers.
- o Sequential composition: Different IETF Network Slices can be placed into a sequence to provide an end-to-end service. In sequential composition, each IETF Network Slice would potentially support different dataplanes that need to be stitched together.

5. Framework

A number of IETF Network Slice services will typically be provided over a shared underlying network infrastructure. Each IETF Network Slice consists of both the overlay connectivity and a specific set of dedicated network resources and/or functions allocated in a shared underlay network to satisfy the needs of the IETF Network Slice customer. In at least some examples of underlying network technologies, the integration between the overlay and various underlay resources is needed to ensure the guaranteed performance requested for different IETF Network Slices.

5.1. IETF Network Slice Stakeholders

An IETF Network Slice and its realization involves the following stakeholders and it is relevant to define them for consistent terminology. The IETF Network Slice Customer and IETF network Slice provider (see Section 2.1) are also stakeholders.

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

IETF Network Slice Controller (NSC): It realizes an IETF Network Slice in the underlying 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 IETF NSCs and different types of orchestrators. An IETF Network Slice operator (or slice operator for short) manages one or more IETF Network Slices using the IETF NSCs.

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

5.2. Expressing Connectivity Intents

The NSC northbound interface (NBI) can be used to communicate between IETF Network Slice customers and the NSC.

An IETF Network Slice customer may be a network operator who, in turn, provides the IETF Network Slice to another IETF Network Slice customer.

Using the NBI, a customer expresses requirements for a particular slice by specifying what is required rather than how that is to be achieved. That is, the customer's view of a slice is an abstract one. Customers normally have limited (or no) visibility into the provider network's actual topology and resource availability information.

This should be true even if both the customer and provider are associated with a single administrative domain, in order to reduce the potential for adverse interactions between IETF Network Slice customers and other users of the underlay network infrastructure.

The benefits of this model can include:

- o Security: because the underlay network (or network operator) does not need to expose network details (topology, capacity, etc.) to IETF Network Slice customers the underlay network components are less exposed to attack;
- o Layered Implementation: the underlay network comprises network elements that belong to a different layer network than customer applications, and network information (advertisements, protocols, etc.) that a customer cannot interpret or respond to (note - a customer should not use network information not exposed via the NSC NBI, even if that information is available);
- o Scalability: customers do not need to know any information beyond that which is exposed via the NBI.

The general issues of abstraction in a TE network is described more fully in [RFC7926].

This framework document does not assume any particular layer at which IETF Network Slices operate as a number of layers (including virtual L2, Ethernet or IP connectivity) could be employed.

Data models and interfaces are of course needed to set up IETF Network Slices, and specific interfaces may have capabilities that allow creation of specific layers.

Layered virtual connections are comprehensively discussed in IETF documents and are widely supported. See, for instance, GMPLS-based networks [RFC5212] and [RFC4397], or Abstraction and Control of TE Networks (ACTN) [RFC8453] and [RFC8454]. The principles and mechanisms associated with layered networking are applicable to IETF Network Slices.

There are several IETF-defined mechanisms for expressing the need for a desired logical network. The NBI carries data either in a protocol-defined format, or in a formalism associated with a modeling language.

For instance:

- o Path Computation Element (PCE) Communication Protocol (PCEP) [RFC5440] and GMPLS User-Network Interface (UNI) using RSVP-TE [RFC4208] use a TLV-based binary encoding to transmit data.
- o Network Configuration Protocol (NETCONF) [RFC6241] and RESTCONF Protocol [RFC8040] use XML and JSON encoding.
- o gRPC/GNMI [I-D.openconfig-rtgwg-gnmi-spec] uses a binary encoded programmable interface;
- o For data modeling, YANG ([RFC6020] and [RFC7950]) may be used to model configuration and other data for NETCONF, RESTCONF, and GNMI - among others; ProtoBufs can be used to model gRPC and GNMI data.

While several generic formats and data models for specific purposes exist, it is expected that IETF Network Slice management may require enhancement or augmentation of existing data models.

5.3. IETF Network Slice Controller (NSC)

The IETF NSC takes abstract requests for IETF Network Slices and implements them using a suitable underlying technology. An IETF NSC is the key building block for control and management of the IETF Network Slice. It provides the creation/modification/deletion, monitoring and optimization of IETF Network Slices in a multi-domain, a multi-technology and multi-vendor environment.

The main task of the IETF NSC is to map abstract IETF Network Slice requirements to concrete technologies and establish required connectivity, and ensuring that required resources are allocated to the IETF Network Slice.

An NSC northbound interface (NBI) is needed for communicating details of a IETF Network Slice (configuration, selected policies, operational state, etc.), as well as providing information to a slice requester/customer about IETF Network Slice status and performance. The details for this NBI are not in scope for this document.

The controller provides the following functions:

- o Provides a technology-agnostic NBI for creation/modification/deletion of the IETF Network Slices. The API exposed by this NBI communicates the endpoints of the IETF network slice, IETF Network Slice SLO parameters (and possibly monitoring thresholds), applicable input selection (filtering) and various policies, and provides a way to monitor the slice.
- o Determines an abstract topology connecting the endpoints of the IETF Network Slice that meets criteria specified via the NBI. The NSC also retains information about the mapping of this abstract topology to underlying components of the IETF network slice as necessary to monitor IETF Network Slice status and performance.
- o Provides "Mapping Functions" for the realization of IETF Network Slices. In other words, it will use the mapping functions that:
 - * map technology-agnostic NBI request to technology-specific SBIs
 - * map filtering/selection information as necessary to entities in the underlay network.
- o Via an SBI, the controller collects telemetry data (e.g., OAM results, statistics, states, etc.) for all elements in the abstract topology used to realize the IETF Network Slice.
- o Using the telemetry data from the underlying realization of a IETF Network Slice (i.e., services/paths/tunnels), evaluates the current performance against IETF Network Slice SLO parameters and exposes them to the IETF Network Slice customer via the NBI. The NSC NBI may also include a capability to provide notification in case the IETF Network Slice performance reaches threshold values defined by the IETF Network Slice customer.

An IETF Network Slice customer is served by the IETF Network Slice Controller (NSC), as follows:

- o The NSC takes requests from a management system or other application, which are then communicated via an NBI. This interface carries data objects the IETF Network Slice customer provides, describing the needed IETF Network Slices in terms of topology, applicable service level objectives (SLO), and any monitoring and reporting requirements that may apply. Note that - in this context - "topology" means what the IETF Network Slice connectivity is meant to look like from the customer's perspective; it may be as simple as a list of mutually (and symmetrically) connected end points, or it may be complicated by details of connection asymmetry, per-connection SLO requirements, etc.

- o These requests are assumed to be translated by one or more underlying systems, which are used to establish specific IETF Network Slice instances on top of an underlying network infrastructure.
- o The NSC maintains a record of the mapping from customer requests to slice instantiations, as needed to allow for subsequent control functions (such as modification or deletion of the requested slices), and as needed for any requested monitoring and reporting functions.

5.3.1. IETF Network Slice Controller Interfaces

The interworking and interoperability among the different stakeholders to provide common means of provisioning, operating and monitoring the IETF Network Slices is enabled by the following communication interfaces (see Figure 2).

NSC Northbound Interface (NBI): The NSC Northbound Interface is an interface between a customer's higher level operation system (e.g., a network slice orchestrator) and the NSC. It is a technology agnostic interface. The customer can use this interface to communicate the requested characteristics and other requirements (i.e., the SLOs) for the IETF Network Slice, and the NSC can use the interface to report the operational state of an IETF Network Slice to the customer.

NSC Southbound Interface (SBI): The NSC Southbound Interface is an interface between the NSC and network controllers. It is technology-specific and may be built around the many network models defined within the IETF.




Figure 2: Interface of IETF Network Slice Controller

5.3.2. Northbound Interface (NBI)

The IETF Network Slice Controller provides a Northbound Interface (NBI) that allows customers of network slices to request and monitor IETF Network Slices. Customers operate on abstract IETF Network Slices, with details related to their realization hidden.

The NBI complements various IETF services, tunnels, path models by providing an abstract layer on top of these models.

The NBI is independent of type of network functions or services that need to be connected, i.e., it is independent of any specific storage, software, protocol, or platform used to realize physical or virtual network connectivity or functions in support of IETF Network Slices.

The NBI uses protocol mechanisms and information passed over those mechanisms to convey desired attributes for IETF Network Slices and their status. The information is expected to be represented as a well-defined data model, and should include at least endpoint and connectivity information, SLO specification, and status information.

To accomplish this, the NBI needs to convey information needed to support communication across the NBI, in terms of identifying the IETF Network Slices, as well providing the above model information.

5.4. IETF Network Slice Structure

An IETF Network Slice is a set of connections among various endpoints to form a logical network that meets the SLOs agreed upon.

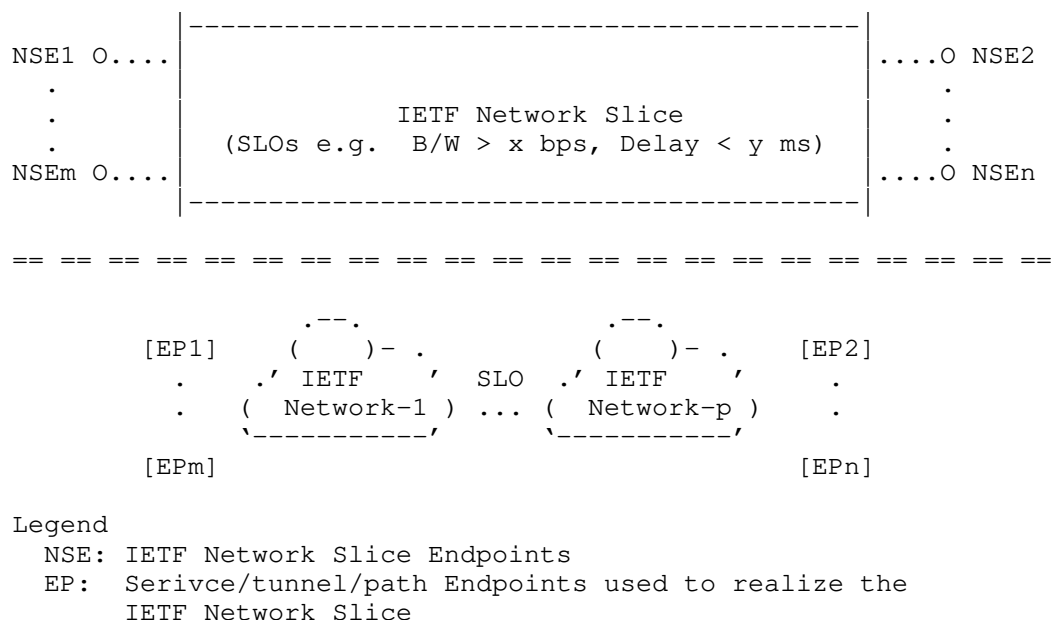


Figure 3: IETF Network Slice

Figure 3 illustrates a case where an IETF Network Slice provides connectivity between a set of IETF network slice endpoints (NSE) pairs with specific SLOs (e.g., guaranteed minimum bandwidth of x bps and guaranteed delay of no more than y ms). The IETF Network Slice endpoints are mapped to the service/tunnel/path Endpoints (EPs) in the underlay network. Also, the IETF NSEs in the same IETF network slice may belong to the same or different address spaces.

IETF Network Slice structure fits into a broader concept of end-to-end network slices. A network operator may be responsible for delivering services over a number of technologies (such as radio networks) and for providing specific and fine-grained services (such as CCTV feed or High definition realtime traffic data). That operator may need to combine slices of various networks to produce an end-to-end network service. Each of these networks may include multiple physical or virtual nodes and may also provide network functions beyond simply carrying of technology-specific protocol data

units. An end-to-end network slice is defined by the 3GPP as a complete logical network that provides a service in its entirety with a specific assurance to the customer [TS23501].

An end-to-end network slice may be composed from other network slices that include IETF Network Slices. This composition may include the hierarchical (or recursive) use of underlying network slices and the sequential (or stitched) combination of slices of different networks.

6. Realizing IETF Network Slices

Realization of IETF Network Slices is out of scope of this document. It is a mapping of the definition of the IETF Network Slice to the underlying infrastructure and is necessarily technology-specific and achieved by the NSC over the SBI.

The realization can be achieved in a form of either physical or logical connectivity using VPNs, virtual networks (VNs), or a variety of tunneling technologies such as Segment Routing, MPLS, etc. Accordingly, endpoints (NSEs) may be realized as physical or logical service or network functions.

6.1. Procedures to Realize IETF Network Slices

There are a number of different technologies that can be used in the underlay, including physical connections, MPLS, time-sensitive networking (TSN), Flex-E, etc.

An IETF Network Slice can be realized in a network, using specific underlying technology or technologies. The creation of a new IETF Network Slice will be initiated with following three steps:

- o Step 1: A higher level system requests connections with specific characteristics via the NBI.
- o Step 2: This request will be processed by an IETF NSC which specifies a mapping between northbound request to any IETF Services, Tunnels, and paths models.
- o Step 3: A series of requests for creation of services, tunnels and paths will be sent to the network to realize the transport slice.

It is very clear that, regardless of how IETF Network Slice is realized in the network (i.e., using tunnels of different types), the definition of the IETF Network Slice does not change at all. The only difference is how the slice is realized. The following sections briefly introduce some existing architectural approaches that can be applied to realize IETF Network Slices.

6.2. Applicability of ACTN to IETF Network Slices

Abstraction and Control of TE Networks (ACTN - [RFC8453]) is a management architecture and toolkit used to create virtual networks (VNs) on top of a traffic engineering (TE) underlay network. The VNs can be presented to customers for them to operate as private networks.

In many ways, the function of ACTN is similar to IETF network slicing. Customer requests for connectivity-based overlay services are mapped to dedicated or shared resources in the underlay network in a way that meets customer guarantees for service level objectives and for separation from other customers' traffic. [RFC8453] the function of ACTN as collecting resources to establish a logically dedicated virtual network over one or more TE networks. Thus, in the case of a TE-enabled underlying network, the ACTN VN can be used as a basis to realize an IETF network slicing.

While the ACTN framework is a generic VN framework that can be used for VN services beyond the IETF network slice, it also a suitable basis for delivering and realizing IETF network slices.

Further discussion of the applicability of ACTN to IETF network slices including a discussion of the relevant YANG models can be found in [I-D.king-teas-applicability-actn-slicing].

6.3. Applicability of Enhanced VPNs to IETF Network Slices

An enhanced VPN (VPN+) is designed to support the needs of new applications, particularly applications that are associated with 5G services, by utilizing an approach that is based on existing VPN and Traffic Engineering (TE) technologies and adds characteristics that specific services require over and above traditional VPNs.

An enhanced VPN can be used to provide enhanced connectivity services between customer sites (a concept similar to an IETF Network Slice) and can be used to create the infrastructure to underpin network slicing.

It is envisaged that enhanced VPNs will be delivered using a combination of existing, modified, and new networking technologies.

[I-D.ietf-teas-enhanced-vpn] describes the framework for Enhanced Virtual Private Network (VPN+) services.

6.4. Network Slicing and Slice Aggregation in IP/MPLS Networks

Network slicing provides the ability to partition a physical network into multiple isolated logical networks of varying sizes, structures, and functions so that each slice can be dedicated to specific services or customers.

Many approaches are currently being worked on to support IETF Network Slices in IP and MPLS networks with or without the use of Segment Routing. Most of these approaches utilize a way of marking packets so that network nodes can apply specific routing and forwarding behaviors to packets that belong to different IETF Network Slices. Different mechanisms for marking packets have been proposed (including using MPLS labels and Segment Routing segment IDs) and those mechanisms are agnostic to the path control technology used within the underlay network.

These approaches are also sensitive to the scaling concerns of supporting a large number of IETF Network Slices within a single IP or MPLS network, and so offer ways to aggregate the slices so that the packet markings indicate an aggregate or grouping of IETF Network Slices where all of the packets are subject to the same routing and forwarding behavior.

At this stage, it is inappropriate to mention any of these proposed solutions that are currently work in progress and not yet adopted as IETF work.

7. Isolation in IETF Network Slices

7.1. Isolation as a Service Requirement

An IETF network slice customer may request that the IETF network slice delivered to them is delivered such that changes to other IETF network slices or services do not have any negative impact on the delivery of the IETF Network Slice. The IETF Network Slice customer may specify the degree to which their IETF Network Slice is unaffected by changes in the provider network or by the behavior of other IETF Network Slice customers. The customer may express this via an SLE it agrees with the provider. This concept is termed 'isolation'

7.2. Isolation in IETF Network Slice Realization

Isolation may be achieved in the underlying network by various forms of resource partitioning ranging from dedicated allocation of resources for a specific IETF Network Slice, to sharing of resources with safeguards. For example, traffic separation between different

IETF Network Slices may be achieved using VPN technologies, such as L3VPN, L2VPN, EVPN, etc. Interference avoidance may be achieved by network capacity planning, allocating dedicated network resources, traffic policing or shaping, prioritizing in using shared network resources, etc. Finally, service continuity may be ensured by reserving backup paths for critical traffic, dedicating specific network resources for a selected number of IETF Network Slices.

8. Management Considerations

IETF Network Slice realization needs to be instrumented in order to track how it is working, and it might be necessary to modify the IETF Network Slice as requirements change. Dynamic reconfiguration might be needed.

9. Security Considerations

This document specifies terminology and has no direct effect on the security of implementations or deployments. In this section, a few of the security aspects are identified.

- o Conformance to security constraints: Specific security requests from customer defined IETF Network Slices will be mapped to their realization in the underlay networks. It will be required by underlay networks to have capabilities to conform to customer's requests as some aspects of security may be expressed in SLEs.
- o IETF NSC authentication: Underlying networks need to be protected against the attacks from an adversary NSC as they can destabilize overall network operations. It is particularly critical since an IETF Network Slice may span across different networks, therefore, IETF NSC should have strong authentication with each those networks. Furthermore, both SBI and NBI need to be secured.
- o Specific isolation criteria: The nature of conformance to isolation requests means that it should not be possible to attack an IETF Network Slice service by varying the traffic on other services or slices carried by the same underlay network. In general, isolation is expected to strengthen the IETF Network Slice security.
- o Data Integrity of an IETF Network Slice: A customer wanting to secure their data and keep it private will be responsible for applying appropriate security measures to their traffic and not depending on the network operator that provides the IETF Network Slice. It is expected that for data integrity, a customer is responsible for end-to-end encryption of its own traffic.

Note: see NGMN document[NGMN_SEC] on 5G network slice security for discussion relevant to this section.

IETF Network Slices might use underlying virtualized networking. All types of virtual networking require special consideration to be given to the separation of traffic between distinct virtual networks, as well as some degree of protection from effects of traffic use of underlying network (and other) resources from other virtual networks sharing those resources.

For example, if a service requires a specific upper bound of latency, then that service can be degraded by added delay in transmission of service packets through the activities of another service or application using the same resources.

Similarly, in a network with virtual functions, noticeably impeding access to a function used by another IETF Network Slice (for instance, compute resources) can be just as service degrading as delaying physical transmission of associated packet in the network.

While a IETF Network Slice might include encryption and other security features as part of the service, customers might be well advised to take responsibility for their own security needs, possibly by encrypting traffic before hand-off to a service provider.

10. Privacy Considerations

Privacy of IETF Network Slice service customers must be preserved. It should not be possible for one IETF Network Slice customer to discover the presence of other customers, nor should sites that are members of one IETF Network Slice be visible outside the context of that IETF Network Slice.

In this sense, it is of paramount importance that the system use the privacy protection mechanism defined for the specific underlying technologies used, including in particular those mechanisms designed to preclude acquiring identifying information associated with any IETF Network Slice customer.

11. IANA Considerations

This document makes no requests for IANA action.

12. Informative References

- [BBF-SD406] Broadband Forum, "End-to-end network slicing", BBF SD-406, <<https://wiki.broadband-forum.org/display/BBF/SD-406+End-to-End+Network+Slicing>>.
- [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.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Network (VPN+) Services", draft-ietf-teas-enhanced-vpn-07 (work in progress), February 2021.
- [I-D.king-teas-applicability-actn-slicing] King, D., Drake, J., Zheng, H., and A. Farrel, "Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to Network Slicing", draft-king-teas-applicability-actn-slicing-10 (work in progress), March 2021.
- [I-D.openconfig-rtgw-gnmi-spec] Shakir, R., Shaikh, A., Borman, P., Hines, M., Lebsack, C., and C. Morrow, "gRPC Network Management Interface (gNMI)", draft-openconfig-rtgw-gnmi-spec-01 (work in progress), March 2018.
- [MACsec] IEEE, "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Security", 2018, <<https://1.ieee802.org/security/802-1ae>>.
- [NGMN-NS-Concept] NGMN Alliance, "Description of Network Slicing Concept", https://www.ngmn.org/uploads/media/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf , 2016.
- [NGMN_SEC] NGMN Alliance, "NGMN 5G Security - Network Slicing", April 2016, <https://www.ngmn.org/wp-content/uploads/Publication_s/2016/160429_NGMN_5G_Security_Network_Slicing_v1_0.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>>.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, DOI 10.17487/RFC4208, October 2005, <<https://www.rfc-editor.org/info/rfc4208>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC4397] Bryskin, I. and A. Farrel, "A Lexicography for the Interpretation of Generalized Multiprotocol Label Switching (GMPLS) Terminology within the Context of the ITU-T's Automatically Switched Optical Network (ASON) Architecture", RFC 4397, DOI 10.17487/RFC4397, February 2006, <<https://www.rfc-editor.org/info/rfc4397>>.
- [RFC5212] Shiimoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, DOI 10.17487/RFC5212, July 2008, <<https://www.rfc-editor.org/info/rfc5212>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.

- [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>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [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>>.
- [RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, DOI 10.17487/RFC7926, July 2016, <<https://www.rfc-editor.org/info/rfc7926>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [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>>.

- [RFC8454] Lee, Y., Belotti, S., Dhody, D., Ceccarelli, D., and B. Yoon, "Information Model for Abstraction and Control of TE Networks (ACTN)", RFC 8454, DOI 10.17487/RFC8454, September 2018, <<https://www.rfc-editor.org/info/rfc8454>>.
- [RFC8795] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Gonzalez de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", RFC 8795, DOI 10.17487/RFC8795, August 2020, <<https://www.rfc-editor.org/info/rfc8795>>.
- [TS23501] 3GPP, "System architecture for the 5G System (5GS)", 3GPP TS 23.501, 2019.
- [TS28530] 3GPP, "Management and orchestration; Concepts, use cases and requirements", 3GPP TS 28.530, 2019.
- [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>>.

Acknowledgments

The entire TEAS Network Slicing design team and everyone participating in related discussions has contributed to this document. Some text fragments in the document have been copied from the [I-D.ietf-teas-enhanced-vpn], for which we are grateful.

Significant contributions to this document were gratefully received from the contributing authors listed in the "Contributors" section. In addition we would like to also thank those others who have attended one or more of the design team meetings, including the following people not listed elsewhere:

- o Aihua Guo
- o Bo Wu
- o Greg Mirsky
- o Lou Berger
- o Rakesh Gandhi
- o Ran Chen

- o Sergio Belotti
- o Stewart Bryant
- o Tomonobu Niwa
- o Xuesong Geng

Further useful comments were received from Daniele Ceccarelli, Uma Chunduri, Pavan Beeram, Tarek Saad, Med Boucadair, Kenichi Okagi, Oscar Gonzalez de Dios, and Xiaobing Niu.

This work is partially supported by the European Commission under Horizon 2020 grant agreement number 101015857 Secured autonomic traffic management for a Tera of SDN flows (Teraflow).

Contributors

The following authors contributed significantly to this document:

Jari Arkko
Ericsson
Email: jari.arkko@piuha.net

Dhruv Dhody
Huawei, India
Email: dhruv.ietf@gmail.com

Jie Dong
Huawei
Email: jie.dong@huawei.com

Xufeng Liu
Volta Networks
Email: xufeng.liu.ietf@gmail.com

Authors' Addresses

Adrian Farrel (editor)
Old Dog Consulting
UK

Email: adrian@olddog.co.uk

Eric Gray
Independent
USA

Email: ewgray@graiymage.com

John Drake
Juniper Networks
USA

Email: jdrake@juniper.net

Reza Rokui
Nokia

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
Juniper Networks

Email: jefftant.ietf@gmail.com

