none
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

L. Qiang, Ed.
Huawei
P. Martinez-Julia
NICT
L. Geng
China Mobile
J. Dong
K. Makhijani
Huawei
A. Galis
University College London
S. Hares
Hickory Hill Consulting
S. Kuklinski
Orange
July 3, 2017

Gap Analysis for Transport Network Slicing
draft-qiang-netslices-gap-analysis-01

Abstract

This document presents network slicing differentiation from the non-partition network or from simply partition of connectivity resources. It lists 7 standardization gaps related to 4 key requirements for network slicing in transport network. It also presents an analysis of existing related work and other potential solutions on network slicing.

This gap analysis document aims to provide a basis for future works in transport network slicing.

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 http://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 4, 2018.

Copyright Notice

Table of Contents

1. Introduction

Network slicing is an approach to enable flexible isolation of network resources and functions for dedicated services, providing a certain level of customization and quality guarantee. It establishes customized dedicated network upon a common infrastructure for vertical industries with flexible design of functions, different performance requirements, system isolation and OAM tools.

Several SDOs have investigated network slicing. To list a few: NGMN initiated a study of network slicing in the context of 5G from the mobile network point of view [NGMN-2016]. Around the same time ITU-T IMT 2020 and ITU-T SG13 studied network softwarization that also included network slicing concept. ITU-T has issued a number of recommendations, such as: Gap Analysis [IMT2020-2015], Network Softwarization [IMT2020-2016], Terms & Definitions [IMT2020-2016bis]. Open Network Foundation (ONF) has developed a recommendation on applying SDN architecture to Network Slicing [ONF-2016]. Finally, 3GPP standards development for 5G includes network slicing in radio access and core networks. 3GPP issued TS 23.501 [TS23-501] about the system architecture for 5G in 2017. BBF started the project SD-406 focusing on the end-to-end architecture enhancement and requirements gathering for transport networks. Although these SDOs have done a lot of work, potential requirements especially in the transport network and end-to-end enabling need to be investigated in order to elicit and identify the technical gaps in IETF for transport network slicing.

In order to establish a network slice that meets various customer's demands, an infrastructure owner needs to understand how these demands map with the available network resources and accessible capabilities. This also requires end-to-end coverage and inter-domain coordination. Meanwhile, the slice provider provides customized OAM to the tenants under provisioning. Slicing OAM

approach is a fundamental capability to guarantee stable, effective and reliable services for the vertical industries. It is also expected to be capable of operations with customized granularity levels that provides robust management flexibilities.

This document presents the identified key requirements and investigates potential technical gaps accordingly. To assist understanding of this document, Section 2 outlines the terminology. Section 3 introduces overall requirements of network slicing. Sections 4~7 illustrates resource specification, end-to-end consideration, performance guarantee and OAM concerns respectively. Section 8 summarizes the identified gaps.

2. Terminology and Abbreviation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

All of the network slicing related words used in this document are to interpreted as described in [NS-Framework]

3. Overall Requirements in Network Slicing

This section introduces 4 key requirements of network slicing derived from [NS-UseCase] as shown in Table 1. These 4 requirements are organized according to a general network slice working process as shown in Figure 1:

1: describe network slicing resource/functions and capture requirements (Req. 1)

2: network slicing cross-domain coordination (Req. 2)

3: construct a performance guaranteed and isolated end-to-end network slice (Req. 3)

4: provide necessary Operation & Maintenance & Administration (OAM) (Req. 4)

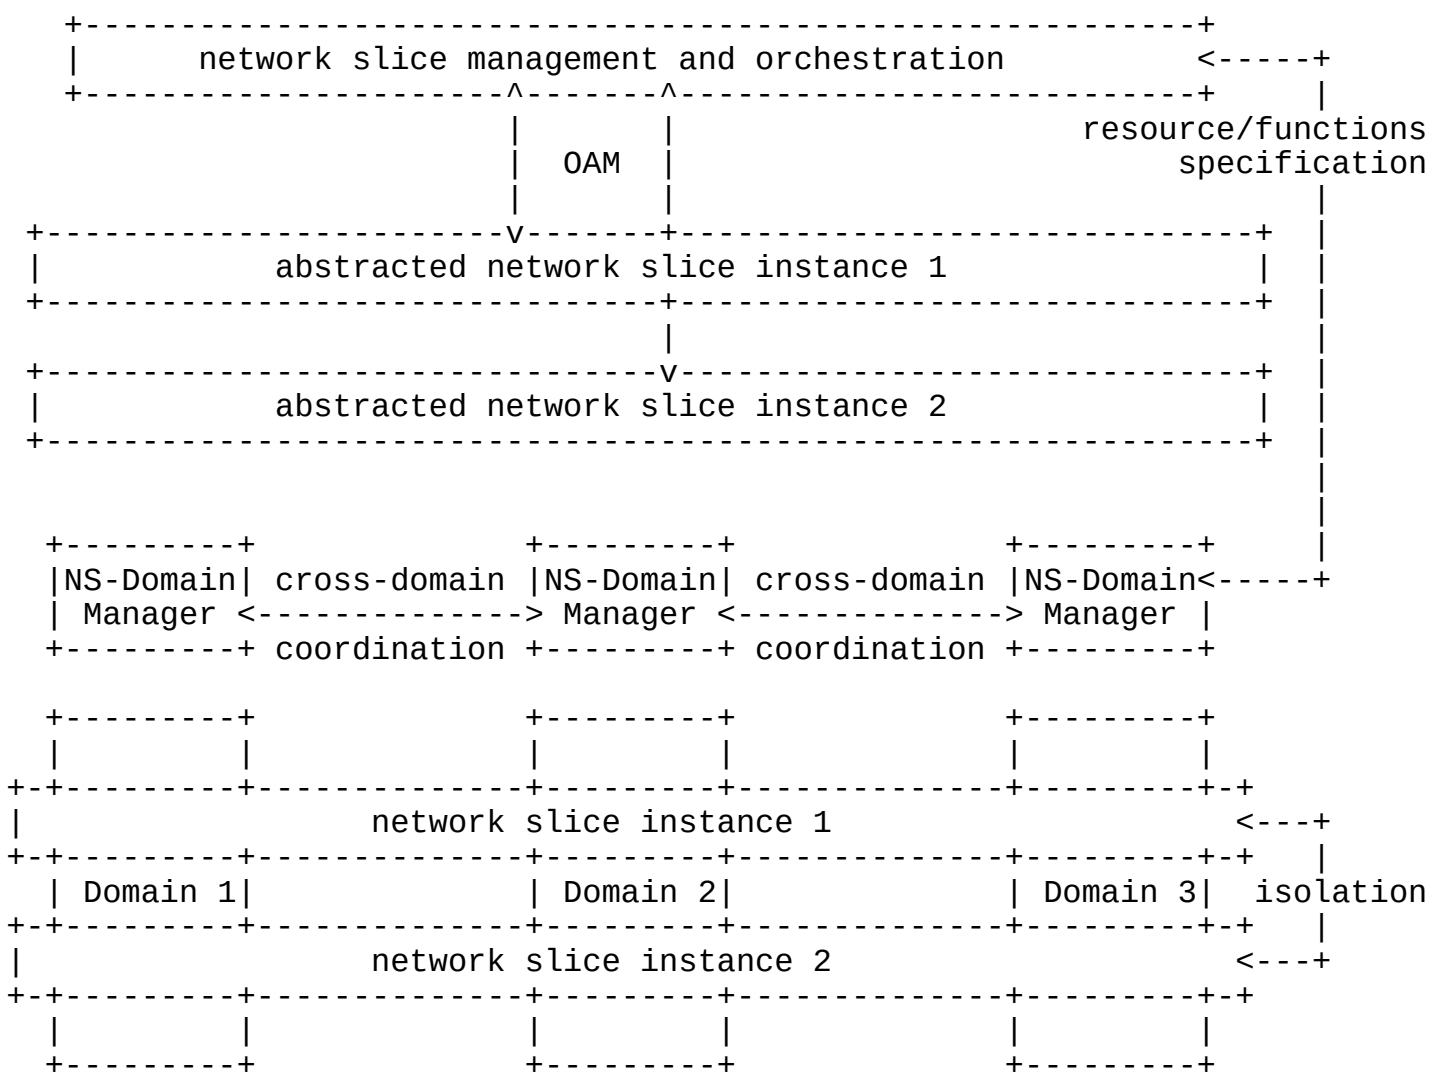

Figure 1: Illustration of Key Requirements

Table 1: Requirement Association

Requirements Illustrated in NS UseCase	Extracted KEY Requirements
1) Resource Reservation 2) Abstraction 3) Multi-Access Knowledge 4) Multi-Dimensional Service Vertical 5) Agile Resource Adjustment	Req 1. Network Slicing Specification
6) Multi-Domain Coordination 7) Resource Assurance	Req 2. Network Slicing Cross-Domain Coordination
8) Performance/Operation Isolation	Req 3. Network Slicing Performance Guarantee and Isolation
9) Independent Slice Management Plane Reliability	Req 4. Network Slicing OAM

Table 1: Requirements Association

- Req 1. Network Slicing Specification (NSS) - The management systems of both network slice providers and operators need to know what and how much resources/network functions they have, so that they can accurately and abstractedly describe the available resources/network functions to tenants or peers. The objective of NSS is to deliver the network slicing requests without incurring any over-utilization of resources. In order to cooperate and provide consistent network slicing service, the way that resources/network functions are described should be homogeneous and compatible among all of the involved technology-specific domains, provides, and slicing platforms.

- Req 2. Network Slicing Cross-Domain Coordination (NS-CDC) - From terminal to server (or other terminal), an end-to-end network slice will involve different infrastructural domains (e.g., AN, TN, CN, etc.) that may be owned by different providers/operators. Each infrastructural domain may be further divided into different administrative domains. That is an end-to-end slice is a logical entity composed by multiple separated components, and the cross-domain coordination is a way to integrate these components together.

- Req 3. Network Slicing Performance Guarantee and Isolation (NS-PGI) - In order to enable the safe, secure, privacy-preservation service for multi-tenancy on a common physical network, the

isolation among network slices in each of the Data/Control/Management/Service planes are needed. Furthermore, network slices that provide differentiated services usually require different resources. The resources allocated to a network slice must be able to guarantee the service performance requirement.

- Req 4. Network Slicing OAM (NS-OAM) - On one end of the spectrum we have those operators that will require a finalized service that they will simply commercialize. On the other end we have those operators that may want to fine-tune all the low-level aspects of the network resources that form their system or service. Moreover, in the middle there is plenty of room for variations. Therefore, the underlying network layers must offer different levels of granularity for the management of their resources, that the upper layer operators can choose according to their needs and objectives.

4. Network Slicing Specification

4.1. Description

Network Slicing Specification (NSS) is meant to describe the network slicing resources and capture requirements from tenants or peer networks to characterize the service expected to be delivered by a network. These requirements include (non-exhaustive): reachability scope (e.g., limited scope, Internet-wide), direction, bandwidth requirements, performance metrics (e.g., one-way delay [RFC2679], loss [RFC2680], or one-way delay variation [RFC3393]), protection and high-availability guidelines (e.g., uRLLC service restoration in less than 50 ms, 100 ms, or 1 second), traffic isolation constraints, and flow identification. NSS is used by a network provider to decide whether existing network slice instances can be reused or (some of them) even combined, or if another network slice instance is needed for a given service.

Technology-specific actions are then derived from the technology-agnostic requirements depicted in an NSS. Such actions include configuration tasks and operational procedures. A standard definition of NSS is needed to facilitate the dynamic/ automated negotiation procedure of NSS parameters, but also to homogenize the processing of service requirements.

To explain by an example, a network slice may cross multiple domains:

- A cloud deployed, NFV enabled, chain of network functions in a virtualized 5G core.

- A segment routing [I-D.ietf-spring-segment-routing] based IGP network transport/aggregation or slice-specific application functions.

- A PCE [RFC4655] monitored TE-tunnel with ingress and egress points.

- Optical, carrier Ethernet or cellular networks.

The network slice is a combination of the above technologies. It creates a compelling need for a common resource specification interface across these domains.

4.2. Related Work in IETF

4.2.1. YANG Data Models

As rightfully discussed in [I-D.wu-opsawg-service-model-explained], the IETF has already published several YANG data models that are used to model monolithic functions as well as very few services (e.g., L2SM, L3SM, EVPN). These models may be used in the context of network slicing if corresponding technologies are required for a given network slice, but none of them can be used to model an NSRD.

[RFC7297] describes the Connectivity Provisioning Profile (CPP) and proposes a CPP template to capture connectivity requirements to be met within a service delivery context. Such a generic CPP template is meant to

- facilitate the automation of the service negotiation and activation procedures, thus accelerating service provisioning;

- set (traffic) objectives of Traffic Engineering (TE) functions and service management functions;

- improve service and network management systems with ’decision-making’ capabilities based upon negotiated/offered CPPs.

[RFC7297] may be considered as a candidate specification for NSRD. Releasing a RFC7297-bis to take into account specific requirements from network slicing is needed. Since [RFC7297] may not be implemented by all providers, the [SLA-Exchange] may be adopted to implement indirect SLA negotiation and SLA events report. [SLA-Exchange] provides an in-band method to exchange the SLA parameters, and then by the receiving devices to translate SLA in technical specific provisioning languages. However, there still does not exist any standard protocol to translate SLA agreements into technical clauses and configurations.

4.2.2. Building NSS from Protocol Independent Traffic Engineering Models

The NSS requirement for reachability, direction, bandwidth requirements, performance metrics, traffic isolation constraints, and flow identification can be built utilizing protocol which can perform operations (read, write, notification, actions (aka rpcs)) on a yang service layer that supports these traffic engineer and resource definition at the service layers. The network slicing service data model can extend existing work in the TEAS and I2RS working group for protocol-independent topology models. These models support configuration or the dynamic datastores defined in [NMDA] which will be abbreviated as NMDA in this section. This section provides the detail on how the NSS can be built from these models and the RESTCONF protocol.

4.2.2.1. Basic Topology Model

The basic topology model is defined in [I2RS-Yang] to include a service layer. This topology model is protocol independent and can be utilized as a configuration data model or a dynamic datastores model. The configuration data model must abide by the configuration persistence and referential requirements. The dynamic datastores do not need to abide by the same requirements as the configuration datastore. I2RS is defining a dynamic datastores reference model for a data store which ephemeral. The network slices may want to use configuration, ephemeral datastores, or define a third type of dynamic datastores. The I2RS WG provides a place to collaborate this work on the dynamic datastores.

4.2.2.2. TEAS Model Utilization of Basic Topology Model

The TEAS topology model [TE-Yang] provides a general description of a Traffic engineering model that provides:

- abstract topologies with TE constraints (bandwidth, delay metrics, links to lower layers, some traffic isolation constraints, and some link identifiers);
- templates for links or resources;
- functionality to read, write, notification, and rpcs.

Options that need to be consider are:

Augmenting TEAS - The TEAS models provide substantial traffic engineering. It was envisioned in the early topology model that a service resource model would be part of the service layer. This

work was delayed until the maturation of the service requirements from L2VPN, L3VPN, and EVPN plus the maturation of resource requirements from 5G. Network slicing provides a good application use case for this work.

Why not Augment TEAS - The TEAS models are TE specific, lack of the abstraction for Layer 3+ resources.

Dynamic models to combine TEAS models for network-slicing - The network slicing controller operating across domains may wish to create a multiple-domain data model based on the service layer data models exposed by different providers. These service models would not need to be configured, but only learned as providers exchange data with one another. The rules for combining these models could be defined as part of the dynamic datastore for network-slicing.

Protocol within a domain - The RESTCONF and NETCONF protocol can support read, write, notification and actions (rpcs) within a domain.

Protocol across domains: The RESTCONF protocol currently supports Configuration protocols and 90% of the dynamic datastores. The RESTCONF protocol is being enhanced to support the push of telemetry messages. The RESTCONF protocol could be used to exchange a specific Yang network-slicing service-layer topology (TE and Resources) and for the I2NSF security capabilities between domains.

If a multicast of telemetry data is required between domains, then the push model for telemetry information or the IPFIX protocol may be utilized.

5. Network Slicing Cross-Domain Coordination

5.1. Description

The network slicing cross-domain coordination (NS-CDC) requirement includes the following aspects:

- Network slice resource/functions coordination: for example, a tenant requests for a network slice with at most 10 ms latency from terminal to server. Different infrastructure/administrative domains should coordinate and negotiate to reach an agreement such as RAN provides at most 2 ms service, TN domain I provides at most 4ms service, TN domain II provides at most 2 ms service and CN provides at most 2 ms service;

- Configuration information coordination: for example, for a given TN domain, the configuration information such as VLAN ID, remote IP address, physical port ID, etc. need to be coordinated with other TN domains;

- Other coordination: for example, RAN (or other access network) needs to notify TN about the information of new attachment point when user moves.

From terminal to server, an end-to-end network slice will involve different infrastructure domains (e.g., RAN, TN and CN). An infrastructure domain may be further divided into multiple domains due to geographic isolation, administrative isolation and other reasons. There are two ways to enable an end-to-end network slice: based on a common platform or based on cross-domain coordination.

If all of the involved domains belong to the same operator or the same operator union, the common platform solution may be work. In this case, all of the domain controllers only need to communicate with the common platform, and follow the coordination management of this common platform. Whilst the most common case is that the domains belong to different owners/operators/administrators, making it difficult to realize such a common platform. Consequently, the cross-domain coordination will be essential throughout the whole lifecycle of an end-to-end network slice.

5.2. Related Work in IETF

There are some related works studies the inter-operation/coordination between different entities. Coordination of different components of a slice requires automation. It can be achieved either by

1. Coordination protocols such as ANIMA, CPNP

2. Or through abstraction and corresponding interfaces as in ACTN.

This subsection will briefly review these related work to provide a basis for the gap analysis.

5.2.1. Autonomic Networking Integrated Model and Approach (ANIMA)

Autonomic Networking Integrated Model and Approach (ANIMA) WG provides a series of tools for distributed and automatic management, which includes: Generic Autonomic Signaling Protocol (GRASP), Autonomic Networking Infrastructure (ANI), etc.

GRASP [ANIMA-GRASP] is a protocol for the negotiation between ASAs (Autonomic Service Agent). In GRASP, ASAs could be considered as

"APPs" installed on a device. Different ASAs fulfill different management tasks such as parameter configuration, service delivery, etc. Based on GRASP, the same purpose ASAs that installed on different devices are able to inter-operate and negotiate with each other. Network slicing could make use of GRASP for the coordination among devices in the underlying infrastructure layer, as well as the negotiation among different domain managers. However, the security issue incurred by cross-domain usage should be fixed in GRASP.

ANI [ANI] is a technical packet consisting of BootStrap (for authentication, domain certification distribution, etc.), ACP (a separate control plane), and GRASP (for control message coordination). ANI could be used to construct the management tunnel among devices in underlying infrastructure layer within a single domain. While the network slicing and cross-domain oriented extensions are necessary.

5.2.2. Connectivity Provisioning Negotiation Protocol (CPNP)

[I-D.boucadair-connectivity-provisioning-protocol] defines the Connectivity Provisioning Negotiation Protocol (CPNP) that is meant to dynamically exchange and negotiate connectivity provisioning parameters, and other service-specific parameters, between a Customer and a Provider. CPNP is a tool that introduces automation in service negotiation and activation procedures, thus fostering the overall service provisioning process.

CPNP runs between a Customer and a Provider carrying service orders from the Customer and respective responses from the Provider to the end of reaching a connectivity service provisioning agreement. As the services offered by the Provider are well-described, by means of the CPP template, the negotiation process is essentially a value-settlement process, where an agreement is pursued on the values of the commonly understood information items (service parameters) included in the service description template.

The protocol is transparent to the content that it carries and to the negotiation logic, at Customer and Provider sides, that manipulates the content.

The protocol aims at facilitating the execution of the negotiation logic by providing the required generic communication primitives.

CPNP can be used in the context of network slicing to request for network resources together with a set of requirements that need to be satisfied by the Provider. Such requirements are not restricted to basic IP forwarding capabilities, but may also include a characterization of a set of service functions that may be invoked.

5.2.3. Abstraction and Control of Traffic Engineered Networks (ACTN)

ACTN [TEAS-ACTN] is an information model proposed by TEAS WG, which enables the multi-domain coordination in Traffic Engineering (TE) network. In order to enable network slicing in transport networks, portion of transport domain will need to be engineered. In particular, building a TE entity and stitching service for this entity is within the scope of ACTN. As an end-to-end network slicing solution, ACTN is able to provide cross-domain coordination. In ACTN, each physical transport network domain is under the control of a Physical Network Controller (PNC) as shown in Figure 2. A Multi-Domain Service Coordinator (MDSC) controls multiple PNCs. Although the MDSCs may form a hierarchical structure, a hierarchical MDSC can still be regarded as a logical common platform. As Section 5.1 discussed, such a common platform solution has a strict presumption that all domains are assumed to follow a common coordination management.

While ACTN does carry out network slicing-related work, some proposed concepts are similar the concepts of today's network slicing: in particular, the virtual network (VN) is similar to a slice instance. ACTN enables VN based on LSP technique, different LSP tunnels correspond to different VNs. However, ACTN focuses on resource abstraction and management on Layer 2 and Layer 1. For transport network slicing, resources abstraction and management on Layer 3+ (e.g., IP routing table, etc.) may also be necessary but have not been addressed by ACTN.

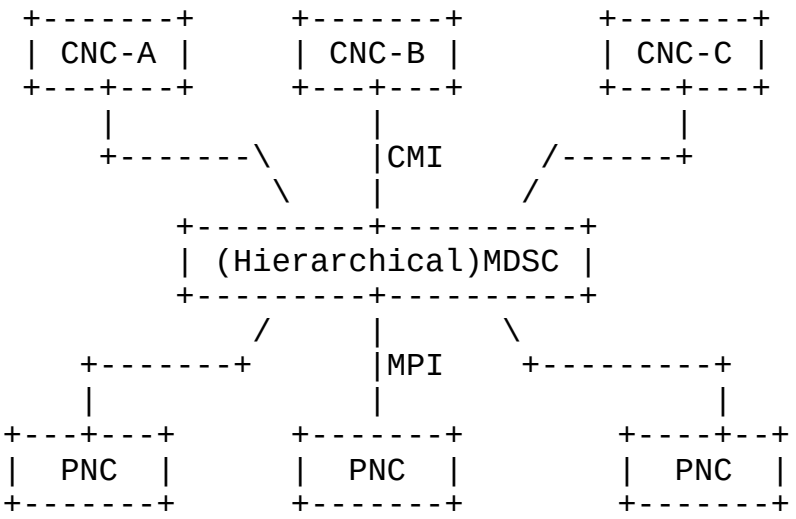

Figure 2: A Three-tier ACTN Control Hierarchy

5.3. Other Potential Solutions

5G Exchange (5GEx) [FGEx] is a 5G-PPP project which aims to enable cross-domain orchestration of services over multiple administrations or over multi-domain single administration networks. The main infrastructure considered in 5GEx is the NFV/SDN compatible software defined infrastructure, which limits the scope of network slicing to SDN based architecture.

6. Network Slicing Performance Guarantee and Isolation

6.1. Description

Network slicing is expected to enable the deployment of various services with diverse requirements, independently on a common physical network. Each network slice is characterized by particular service requirements, which usually are expressed using in the form of several key performance indicators (KPIs) such as bandwidth, latency, jitter, packet loss, etc., and different degrees of isolation. Isolation requirements include performance isolation, which means performance guarantee are maintained regardless of activity in other slices, as well as secure isolation (e.g., including privacy), and management (or OAM) isolation. Additionally, performance isolation in network slicing has to maintain while scaling up or down computing capabilities of a slice (i.e., for elastic scaling). Moreover, since IoT is also a use case for NS, and since some IoT applications are sensitive to data plane or bits on wire overheads, data path encapsulation in the form of labels, VLANs, VxLANs should be optional, or minimized for those cases.

As we will discuss in the detailed sections below, each of these technologies can address some but not all performance and isolation requirements:

- RSVP-TE, Segment Routing (SR), DETNET, FlexE are mostly related to performance guarantee and performance isolation requirements

- Virtual Private Networks (VPN), NVO3 are mostly related to security and management isolation requirements

A Network Slicing solution, to support performance guarantee and isolation requirements, will therefore need to merge in some way characteristics from these two families of technologies, through the combination of (possibly enhanced) existing technologies and/or specifically developed ones. We can also consider the possibility that multiple such technology stacks may be deployed in different domains, and rely on cross-domain coordination, as described inSection 5, to form a single abstracted network slice.

6.2. Related Work in IETF

6.2.1. Virtual Private Networks

VPN technologies such as L3VPN [RFC4364], L2VPN [RFC4664], EVPN [RFC7432], etc. have been widely deployed to provide different virtual networks on common service provider networks. Although VPNs can provide logically separated routing/bridging domains between different VPN customers, essentially it is an overlay network technology with little control of the network resources, so it is challenging for VPN to meet the performance and isolation requirement of some emerging application scenarios such as industrial verticals. VPNs essentially are private networks of enterprises by connecting remote sites. The following two issues illustrate limitations of VPNs for network slicing:

- An end-to-end VPN tunnel competes with other traffic in the network and end-to-end network resource policies cannot be guaranteed.

- The reachability and resource reservation protocols are not tightly integrated and often solutions require centralized PCE-P like methods.

6.2.2. NVO3

[NVO3-WG] defines several network encapsulations which support the network virtualization and multi-tenancy in the data center networks. Similar to the VPN technologies of service provider networks, NVO3 is also an overlay network technology, which relies on the performance characteristics provided by the IP-based underlay networks. Thus NVO3 may not meet by itself the performance and isolation requirements of network slicing.

6.2.3. RSVP-TE

RSVP-TE [RFC3209] is the signaling protocol to establish end-to-end traffic-engineered Label Switched Paths (LSPs). It can reserve the required link bandwidth along an end-to-end path for specific network flows, which is suitable for services with particular requirement on traffic bandwidth. RSVP-TE LSPs can be used as the underlay tunnels of the VPN service connections. However, the requirement of some emerging services is not only about traffic bandwidth, but also has quite strict requirement on latency, jitter, etc. Such requirements can hardly be met with existing RSVP-TE.

6.2.4. Segment Routing

[I-D.ietf-spring-segment-routing] provides the ability to specify a traffic-engineered path by the source node of data packets. It can provide traffic-engineering features comparable to RSVP-TE with better scalability, by eliminating the per-path state in the transit network nodes. It is therefore a candidate method of creating an NSI, mapping a packet into an NSI and specifying the passage of the packet through the resources dedicated to the NSI. Further study will be required to determine if/how SR as designed today can be used as a core technology for building an NSI. With respect to performance guarantee and isolation, some further investigation may be needed to understand whether SR can provide the same or better performance characteristics as RSVP-TE. In addition, it is not clear whether SR-based LSPs can provide the guaranteed latency and jitter performance required by network slicing.

6.2.5. Deterministic Networking

[DETNET-WG] is working on the deterministic data paths over layer 2 and layer 3 network segments. Such deterministic paths can provide identified flows with extremely low packet loss rates, low packet delay variation (jitter) and assured maximum end-to-end delivery latency. This is accomplished by dedicating network resources such as link bandwidth and buffer space to DetNet flows and/or classes of DetNet flows. DetNet also aims to provide high reliability by replicating packets along multiple paths. It is a characteristic of DetNet that it is concerned solely with worst-case values for the end-to-end latency.

The primary target of DetNet is real-time systems and as such average, mean, or typical latency values are not protected, because they do not affect the ability of a real-time system to perform their tasks. This contrasts with a normal priority-based queuing scheme which will give better average latency to a data flow than DetNet, but, on the other side, the worst-case latency can be essentially unbounded. As such DetNet seems to be a useful technique that may be applied to either a complete NSI, or to part of the traffic within an NSI to address the emerging low latency requirement for real time application.

DetNet can therefore address some of the requirements of NS. It was however not designed with network slicing in mind, which means a mapping between an NSI and a DetNet service may need to be defined.

6.2.6. Flexible Ethernet

[FLEXE-1.0] was initially defined by Optical Internetworking Forum (OIF) as an interface technology which allows the complete decoupling of the Media Access Control layer (MAC) data rates and the standard-based Ethernet Physical layer (PHY) rates. The channelization capability of FlexE can be used to partition a FlexE interface into several independent sub-interfaces, which can be considered as a useful component for the slicing of network interfaces. Currently there is ongoing work in IETF to define the control plane framework for FlexE [FlexE-FWK], which aims to identify the routing and signaling extensions needed for establishing FlexE-based end-to-end LSPs in IP/MPLS networks.

7. Network Slicing OAM with Customized Granularity

7.1. Description

In accordance with [RFC6291], OAM is used to denote the following:

- Operations: refer to activities that are undertaken to keep the network and the services it deliver up and running. It includes monitoring the underlying resources and identifying problems.

- Administration: refer to activities to keep track of resources within the network and how they are used.

- Maintenance: refer to activities to facilitate repairs and upgrades. Maintenance also involves corrective and preventive measures to make the managed network run more effectively, e.g., adjusting configuration and parameters.

As per [RFC6291], network slicing provisioning operations are not considered as part of OAM. Provisioning operations are discussed in other sections.

Maintaining automatically-provisioned slices within a network raises the following requirements:

- Ability to run OAM activities on a provider’s customized granularity level. In other words, ability to run OAM activities at any level of granularity that a service provider see fit. In particular:

  * Per slice OAM: An operator must be able to execute OAM tasks on a per slice basis.

    * Per domain OAM: These tasks can cover the "whole" slice within a domain or a portion of that slice (for troubleshooting purposes, for example).

    * Per service OAM: When a given slice is shared among multiple services/customers, an operator must be able to execute (per-slice) OAM tasks for a particular service or customer.

    * For example, OAM tasks can consist in tracing resources that are bound to a given slice, tracing resources that are invoked when forwarding a given flow bound to a given network slice, assessing whether flow isolation characteristics are in conformance with the NS Resource Specification, or assessing the compliance of the allocated slice resource against flow/ customer requirements.

    * An operator must be able to enable differentiated failure detect and repair features for a specific/subset of network slices. For example, a given slice may require fast detect and repair mechanisms (e.g., as a function of the nature of the traffic (pattern) forwarded through the NS), while others may not be engineered with such means.

- Ability to automatically discover the underlying service functions and the slices they are involved in or they belong to.

- Ability to dynamically discover the set of network slicing that are enabled within a network. Such dynamic discovery capability facilitates the detection of any mismatch between the view maintained by the control plane and the actual network configuration. When mismatches are detected, corrective actions must be undertaken accordingly.

- Ability to efficiently OAM on shared resources. If multiple network slices share some resources, the same kind of OAM operations from different network slices should be performed only once for efficiency. For example, several network slices share a link. We only need to execute once status query, and directly return the queried result to other status query requests.

7.2. Related Work in IETF

7.2.1. Overview of OAM tools

The reader may refer to [RFC7276] for an overview about available OAM tools. These technology-specific tools can be reused in the context of network slicing. Providers that deploy network slicing capabilities should be able to select whatever OAM technology-

specific feature that would be address their needs. No gap that would legitimate specific requirements has been identified so far.

7.2.2. Overlay OAM

[I-D.ooamdt-rtgwg-ooam-header] specifies a generic OAM header that can be used if overlay technologies are enabled. Obviously, this effort can be reused in the context of network slicing when overlay techniques are in use. Nevertheless, For slice designs that do not assume an overlay technology, OAM packets must be able to fly over the appropriate slice and for a given service/customer. This is possible by reusing some existing tools if and only if no specific fields are required (e.g., carry a slice identifier as Req. 5 stated).

7.2.3. Service Function Chaining

SFC WG [SFCWG] is chartered to describe data plane service encapsulation, control and manageability aspects of service functions. Extensions that will be specified by the SFC WG will be reused in the context of network slicing. Nevertheless, The current charter of the WG does not imply work on the automated discovery of SF instances and their capabilities, nor the automatic discovery of control elements. An additional specification effort is therefore required in this area.

7.2.4. Slice Identification

A network slice data plane, may or may not follow traditional data plane tagging/labeling. However, each network element (router/ switch) still has to classify an incoming packet and associated with the slice instance for proper treatment. Network slice instance identification is essential for network element to make local decisions on forwarding policies, QoS mechanism and etc. The performance requirements of a network slice instance can therefore been met by making the correct decision. Meanwhile, it is also important for OAM so that configuration and provisioning can be delicately performed to particular network slice instances by their identifications.

For flow identification, many existing technologies provide mature solutions. These approaches might be able to be re-used in network slicing by adding an additional layer of mapping to a network slice instance ID. The network slice instance ID further maps to a group of performance requirements and OAM profiles, based on which the network elements within the slice can make local decisions. However, per flow level identification could have adverse impact on the scale of the forwarding entries in the routers.

With traditional IP/MPLS VPNs, the set of Route Targets configured for the VPN can be used as some sort of identifier of the VPN in the control plane, and in the data plane, the VPN service labels can be used to identify the data packets belonging to a particular VPN. NVO3 uses the Virtual Network Identifiers (VNIs) in the header of data packets to identify different overlay network tenants. However, It is not clear if the existing identifiers can meet the requirements of network slicing in terms of making local decisions on forwarding policy, QoS and OAM mechanisms, etc.

8. Summary

The following table is a summary of the identified gaps based on previous analysis in this document.

Requirements	Gaps
Req 1. Network Slicing Specification (NSS)	1) A detailed specification of NSS 2) A companion YANG data model for NSS
Req 2. Network Slicing Cross-Domain Coordination (NS-CDC)	3) A companion data model for NS-CDC
Req 3. Network Slicing Performance Guarantee and Isolation (NS-PGI)	4) Slicing specific extension on existing technologies
Req 4. Network Slicing OAM (NS-OAM)	5) Mechanisms for dynamic discovery of service function instances and their capabilities. Mechanisms for dynamic discovery of instantiated network slices 6) non-overlay OAM solution 7) Mechanisms for customized granularity OAM

Table 2: Summary of Gaps

9. Security Considerations

This document analyzes the standardization work on network slicing in different WGs. As no solution proposed in this document, no security concern raised.

10. IANA Considerations

There is no IANA action required by this document.

11. Acknowledgements

The authors wish to thank Hannu Flinck, Akbar Rahman, Ravi Ravindran, Xavier de Foy, Young Lee and Igor Bryskin for their detailed and constructive reviews. Many thanks to Mohamed Boucadair, Christian Jacquenet and Stewart Bryant for their valuable contributions and comments.

12. References

12.1. Normative References

[NS-Framework]
"NS Framework", <https://datatracker.ietf.org/doc/draft-geng-netslices-architecture/>.

[NS-UseCase]
"NS Use Case", <https://datatracker.ietf.org/doc/draft-netslices-usecases/>.

12.2. Informative References

[ANI] "A Reference Model for Autonomic Networking", <https://datatracker.ietf.org/doc/draft-ietf-anima-reference-model/?include_text=1>.

[ANIMA-GRASP]
"A Generic Autonomic Signaling Protocol (GRASP)", <https://datatracker.ietf.org/doc/draft-ietf-anima-grasp/>.

[DETNET-WG]
"Deterministic Networking", <https://datatracker.ietf.org/wg/detnet/about/ >.

[FGEx] "5G Exchange (5GEx) - Multi-domain Orchestration for Software Defined Infrastructures", <https://www.researchgate.net/publication/296486303_5G_Exchange_5GEx_-_Multi-domain_Orchestration_for_Software_Defined_Infrastructures>.

[FLEXE-1.0]
"Flexible Ethernet 1.0", <http://www.oiforum.com/wp-content/uploads/OIF-FLEXE-01.0.pdf>.

[FlexE-FWK]
"FlexE-FWK", <https://datatracker.ietf.org/doc/draft-izh-ccamp-flexe-fwk/>.

[I-D.boucadair-connectivity-provisioning-protocol]
Boucadair, M., Jacquenet, C., Zhang, D., and P. Georgatsos, "Connectivity Provisioning Negotiation Protocol (CPNP)", draft-boucadair-connectivity-provisioning-protocol-14 (work in progress), May 2017.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-12 (work in progress), June 2017.

[I-D.ooamdt-rtgwg-ooam-header]
Mirsky, G., Kumar, N., Kumar, D., Chen, M., Yizhou, L., and D. Dolson, "OAM Header for use in Overlay Networks", draft-ooamdt-rtgwg-ooam-header-03 (work in progress), March 2017.

[I-D.wu-opsawg-service-model-explained]
Wu, Q., LIU, W., and A. Farrel, "Service Models Explained", draft-wu-opsawg-service-model-explained-06 (work in progress), May 2017.

[I2RS-Yang]
"A Data Model for Network Topologies", <https://datatracker.ietf.org/doc/draft-ietf-i2rs-yang-network-topo/ >.

[IMT2020-2015]
"Report on Gap Analysis", <http://www.itu.int/en/ITU-T/focusgroups/imt-2020/Pages/default.aspx >.

[IMT2020-2016]
"Draft Technical Report Application of network softwarization to IMT-2020 (O-041)", <http://www.itu.int/en/ITU-T/focusgroups/imt-2020/Pages/default.aspx >.

[IMT2020-2016bis]
"Draft Terms and definitions for IMT-2020 in ITU-T (O-040)", <http://www.itu.int/en/ITU-T/focusgroups/imt-2020/Pages/default.aspx >.

[NGMN-2016]
"Description of Network Slicing Concept", <https://www.ngmn.org/uploads/media/160113_Network_Slicing_v1_0.pdf>.

[NMDA] "Network Management Datastore Architecture", <https://datatracker.ietf.org/doc/draft-ietf-netmod-revised-datastores/>.

[NVO3-WG] "Network Virtualization Overlays".

[ONF-2016] TS, "Applying SDN Architecture to 5G Slicing", <https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/Applying_SDN_Architecture_to_5G_Slicing_TR-526.pdf >.

[RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, DOI 10.17487/RFC2679, September 1999, <http://www.rfc-editor.org/info/rfc2679>.

[RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, DOI 10.17487/RFC2680, September 1999, <http://www.rfc-editor.org/info/rfc2680>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <http://www.rfc-editor.org/info/rfc3209>.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <http://www.rfc-editor.org/info/rfc3393>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <http://www.rfc-editor.org/info/rfc4364>.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <http://www.rfc-editor.org/info/rfc4655>.

[RFC4664] Andersson, L., Ed. and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <http://www.rfc-editor.org/info/rfc4664>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <http://www.rfc-editor.org/info/rfc5440>.

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <http://www.rfc-editor.org/info/rfc6291>.

[RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <http://www.rfc-editor.org/info/rfc7276>.

[RFC7297] Boucadair, M., Jacquenet, C., and N. Wang, "IP Connectivity Provisioning Profile (CPP)", RFC 7297, DOI 10.17487/RFC7297, July 2014, <http://www.rfc-editor.org/info/rfc7297>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <http://www.rfc-editor.org/info/rfc7432>.

[SFCWG] "\Service Function Chaining (sfc)", <https://datatracker.ietf.org/wg/sfc/about/>.

[SLA-Exchange] "Inter-domain SLA Exchange Attribute", <https://datatracker.ietf.org/doc/draft-ietf-idr-sla-exchange/>.

[TE-Yang] "YANG Data Model for TE Topologies", <https://datatracker.ietf.org/doc/draft-ietf-teas-yang-te-topo/>.

[TEAS-ACTN] "Information Model for Abstraction and Control of TE Networks (ACTN)", <https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-info-model>.

[TS23-501] "System Architecture for the 5G System", <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

Authors' Addresses

Li Qiang (editor)
Huawei

Email: qiangli3@huawei.com

Pedro Martinez-Julia
NICT

Email: pedro@nict.go.jp

Liang Geng
China Mobile

Email: gengliang@chinamobile.com

Jie Dong
Huawei

Email: jie.dong@huawei.com

Kiran Makhijani
Huawei

Email: Kiran.Makhijani@huawei.com

Alex Galis
University College London

Email: a.galis@ucl.ac.uk

Susan Hares
Hickory Hill Consulting

Email: shares@ndzh.com

Slawomir
Orange

Email: slawomir.kuklinski@orange.com