

Multi-layer Transport Network Slicing with Hard and Soft Isolation

A. Alcalá¹, S. Barguil¹, V. López², L.M. Contreras², C. Manso³, P. Alemany³,
R. Casellas³, R. Martínez³, D. González-Pérez⁴, X. Liu⁴, J.M. Pulido⁴,
J.P. Fernández-Palacios², R. Muñoz³, R. Vilalta³

¹ Universidad Autónoma de Madrid, Madrid, Spain

² Telefónica I+D, Madrid, Spain

³ Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels (Barcelona), Spain

⁴ Volta Networks, Barcelona, Spain

e-mail: ricard.vilalta@cttc.es

Abstract: We validate the deployment of isolated transport network slices in IP over DWDM network. To this end, an isolated transport network slice is deployed using multi-layer isolation mechanisms based on OpenConfig and ONF Transport API. ©2021 The Authors.

1. Introduction

The provisioning of connectivity services that guarantee a specific set of Service Level Objectives (SLO) regarding network resources is expected to benefit many use cases, such as beyond 5G networks or NFV and data center interconnects. Transport network slices provide connectivity coupled with a set of specific network resources commitment between several endpoints over a shared network infrastructure [1].

Each slice is associated with a tenant. Each tenant can control and manage all its slices. As underlying resource multi-tenancy is supported, multiple isolation options are provided, such as soft slicing and hard slicing. Soft network slicing focuses on QoS mechanisms that provide a dynamic allocation of available network resources to different traffic classes. An example of a soft network slice is a VLAN assigned to a customer to carry voice/data traffic, usually transported by LxVPNs. On the other hand, hard network slicing provides component virtualization and replication. An example of a hard network slice is routers (physical or virtual) under the same administrative domain, where services are configured between routers.

The virtualization of the transport network can be exploited at the time of provisioning and orchestrating the virtualized service functions of the slice [2]. The idea leverages on the concept of Wide-area Infrastructure Manager (WIM) as defined by the ETSI NFV architecture as the element devoted to manage the virtualization capabilities in the WAN. Thus, it is assumed that a control entity will be in charge of handling the WAN transport connectivity for interconnecting given communication end-points. When referring to slices, such an entity could be associated to a Network Slice Controller (NSC), as defined in [1] either complementing or being part of an overarching SDN transport network control environment. A single transport network slice request can be decomposed into multiple control and management steps across all the network components involved. In order to provide the necessary transport network slice, NSC will determine the necessary resources allocation depending on the characteristics of the requested network slice (i.e., soft or hard). Once resources are allocated, they are configured on the network using multiple southbound interfaces (SBI).

an example interface between the NSC and the underlying optical controller is the Open Networking Foundation (ONF) Transport API (T-API) [3], which allows the underlying optical SDN controller topology export and connectivity service provisioning. ONF T-API allows the provisioning of connectivity services with specific connectivity constraints (including QoS requirements), which can later be mapped to specific network slice isolation levels.

Another well demonstrated data model is OpenConfig [4], which provides vendor-neutral YANG data models for various network elements, from routers to optical switches. The IP SDN Domain Controller will be responsible for allocating, instantiating and configuring the multiple (virtual) routers and interfaces depending on the required slice isolation level. OpenConfig data models usage is combined with several protocols such as gRPC or NETCONF. The gRPC protocols provide high performance and scalability to the proposed architecture.

This paper presents an end-to-end architecture to provide transport network slices deployed over multi-layer IP over DWDM networks. Several degrees of isolation (from hard to soft) might be required and implemented in the requested transport network slice. This is the first paper to explore transport network slice isolation using an IP over DWDM network. In order to validate this proposed architecture, we present a proof-of-concept in Telefonica and CTTC Laboratories.




Fig. 1. a) Overall proposed architecture; and b) Sequence diagram for multiple network slice deployment with isolation characteristics.

2. Proposed architecture

Figure 1.a shows a network scheme of the approached concept. It consists of three main domains: SDN domain, IP domain, and optical domain.

The Network Slice Controller (NSC) realizes a transport network slice in the underlying transport infrastructure, maintains and monitors the state of its resources. The NSC will delegate on SDN Domain controllers [6] to configure the network resources. The NSC receives a transport network slice request from the Operation Support System and Business Support System (OSS/BSS). The request is modeled using the YANG data model defined in [5] by means of the RESTCONF protocol [7]. Internal workflow for transport network slice life-cycle management is prepared on top of L2/L3 service management (SM) workflows. In order to interact with underlying IP and Optical Domain controllers via a RESTCONF client.

The OSS/BSS may request that the deployed network slice is isolated from any other network slices or different services delivered to the same customers. Naturally, other network slices or services must not negatively impact the requested transport network slice's delivery. There are several possibilities to provide this isolation, which can be provided at several degrees, such as dedicated allocation of resources for a specific slice or sharing some network resources. Figure 2.a shows the multiple isolation options that range from a hard slice to a soft slice: a) no-isolation, meaning that slices are not separated; b) physical-isolation, where slices are completely physically separated, for example, in different locations; c) logical-isolation, where slices are logically separated, only a certain degree of isolation is performed through QoS mechanisms; d) process-isolation, where slices include process and threads isolation; e) physical-network-isolation, where slices contain physically separated links; f) virtual-resource-isolation, where slices have dedicated virtual resources; g) network-functions-isolation, where Network Function (NF) are dedicated to a single network slice; h) service-isolation, where virtual resources and NFs are shared.

Figure 1.b shows the proposed workflow to deploy hard and soft transport network slices. In the workflow, two isolated network slices are deployed. The first one allocates a connectivity service to interconnect both IP layer domains (see for reference Figure 1.a). This triggers the necessary optical configuration mechanism to each of the underlying ROADMs (e.g., using OpenROADM protocol).

Once the connectivity service has been established, the NSC is responsible for requesting to IP SDN domain controller the necessary virtual routers (in the proposed scenario, two site-network-access are configured, one network instance on each site). Link Aggregation Control Protocol (LACP) is configured to each network instance. Then interfaces are aggregated and properly configured using dedicated VLAN or MPLS-TE mechanisms.

When a new isolated slice is requested, NSC can request a dedicated and isolated connectivity service to the underlying optical SDN controller. Figure 2.b details the available connectivity constraint to provide disjoint path selection using ONF Transport API. It consists of including a diversity exclusion constraint with the previous connectivity service identifier. Later, at IP layer novel virtual routers are deployed to provide the requested degree of isolation.

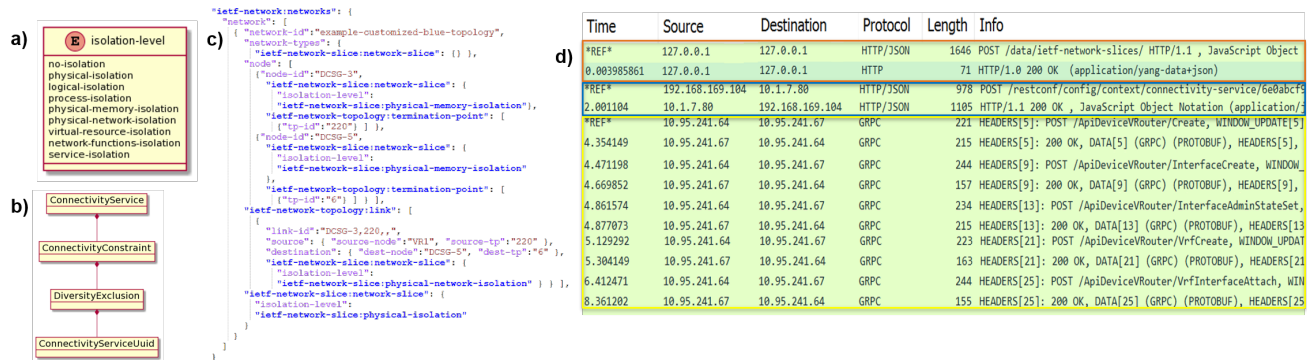


Fig. 2. a) Network Slice isolation levels; b) Disjoint path selection using ONF Transport API; c) Example of transport slice request; and d) Wireshark captures for deployment of a hard slice.

3. Results

The experimental testbed has been set-up at Telefónica (NSC deployment and Volta Elastic Virtual Routing Engine stacked in 7316 Edgecore hardware) and CTTC laboratories (including a T-API SDN controller over a flexi-grid 4-nodes DWDM network), as depicted in Figure 1.a. In order to properly set-up transport, we have introduced the transport network slice YANG data model in the RESTCONF server of NSC. An example of requested transport network slice is depicted in Figure 2.c. It can be observed that several levels of isolation can be requested per node, link and slice.

In order to validate the depicted sequence diagram in Figure 1.b, we provide the captured Wireshark traces among the complete system to deploy a complete hard slice request (the figure provides the selected significant traces). In this figure the different involved protocols can be observed. Firstly it shows the request for the transport network slice. It is quickly stored and answered, as it is processed after response (later an status update might be requested).

The transport network slice requests physical network isolation for the solicited link, thus a T-API connectivity service is requested including as connectivity constraints the diversity exclusion option. For that, we provide the identifiers from the previous established connectivity services and we obtain a disjoint path for the new requested slice, resulting in the requested physical network isolation (Figure 2.b). It can be appreciated that the connectivity service setup time is 2.001s. Consider that the ADRENALINE testbed already has the paths equalized.

Finally, the virtual routers (vRouter) are created and properly setup. The Wireshark capture shows the internal gRPC protocol from Volta Networks, which is equivalent to OpenConfig calls, for the vRouter deployment and configuration. It includes a call to create a new vRouter on top of Edgecore hardware. Then, interfaces are incorporated to the vRouter and configured. Later, Virtual routing and forwarding (VRF) table is setup and finally interfaces are attached to the VRF. This process set-up delay is of 8.36s.

4. Conclusions

We have presented and validated an end-to-end architecture that allows the deployment of transport network slices with several degrees of isolation. The results indicate the feasibility of deployment of multi-layer IP over DWDM transport network slices based on virtual routers and optical disjoint paths to provide hard isolation. Soft isolation is instead reached through connectivity constraints and L2VPN service router configuration.

Acknowledgements

Work supported by the EC H2020 TeraFlow (101015857) and Spanish AURORAS (RTI2018-099178-I00) projects.

References

1. R. Rokui, et al., Definition of IETF Network Slices, IETF draft draft-ietf-teas-ietf-network-slice-definition-00, 2021.
2. R. Casellas, et al., Virtualization of disaggregated optical networks with open data models in support of network slicing. JOCN Feb 1;12(2):A144-54, 2020.
3. R. Vilalta, et al., Transport API extensions for the interconnection of multiple NFV infrastructure points of presence. OFC, 2019.
4. A. Shaikh, et al., Vendor-neutral network representations for transport SDN. OFC, 2016.
5. X. Liu, et al., IETF Network Slice YANG Data Model, IETF draft-liu-teas-transport-network-slice-yang-02, 2020.
6. L. Contreras, et al., FUSION: Standards-based SDN Architecture for Carrier Transport Network. CSCN, 2018.
7. A. Bierman, M. Bjorklund, K. Watsen, R. Fernando, RESTCONF protocol. IETF RFC 8040, 2017.