

POLITECNICO DI TORINO
Repository ISTITUZIONALE

Demo: network slices for virtual Content Delivery Networks in 5G infrastructures

Original

Demo: network slices for virtual Content Delivery Networks in 5G infrastructures / Landi, Giada; Giardina, Pietro; Capitani, Marco; Kondepu, Koteswararao; Valcarengi, Luca; Avino, Giuseppe. - STAMPA. - (2019). (Intervento presentato al convegno ACM MOBIHOC 2019 tenutosi a Catania nel 2-5 July).

Availability:

This version is available at: 11583/2731150 since: 2019-04-17T14:12:22Z

Publisher:

ACM

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

ACM postprint/Author's Accepted Manuscript, con Copyr. autore

(Article begins on next page)

Demo: provisioning and automated scaling of network slices for virtual Content Delivery Networks in 5G infrastructures

Giada Landi, Pietro Giardina,
Marco Capitani
Nextworks
Pisa, Italy
{g.landi, p.giardina,
m.capitani}@nextworks.it

Koteswararao Kondepu,
Luca Valcarenghi
Scuola Superiore Sant'Anna
Pisa, Italy
{k.kondepu,
luca.valcarenghi@santannapisa.it}

Giuseppe Avino
Dept. of Electronics and
Telecommunications
Politecnico di Torino
Torino, Italy
giuseppe.avino@polito.it

ABSTRACT

The concept of network slicing in 5G infrastructures allows to deliver multiple virtual services over shared environments and fully customized based on vertical-driven requirements and target performance indicators. A key feature of 5G networks is the capability to dynamically re-optimize the allocation of computing and network resources, from the core up to the edge, to instantiate and manage different types of concurrent services over a shared infrastructure. In this demonstration, we present a network slicing and orchestration solution for vertical services in the media sector, where enhanced Mobile Broadband (eMBB) network slices are instantiated interconnecting physical and virtual functions, provisioned and configured on-demand using the concept of NFV Network Services. The eMBB slices are automatically dimensioned to match the requirements of the virtual Content Delivery Network (vCDN) service, e.g. in terms of number of target users, video quality and geographical coverage area. Arbitration and resource allocation schemas optimize the sharing of mobile communication services and virtual resources among concurrent media service instances, in compliance with the Service Level Agreement between network operators and vertical service providers.

KEYWORDS

5G, network slicing, Content Delivery Network, NFV, edge computing

1 System architecture and demo workflow

The proposed architecture enhances the management and orchestration (MANO) framework, defined in the ETSI Network Function Virtualization (NFV) group, to manage network slices in support of vertical services. In particular, we introduce a novel functional element, called Vertical Slicer (VS) [1], to dynamically provision, operate and monitor network slices on the basis of service requirements specified by vertical service providers.

The VS works on top of the NFV Orchestrator (NFVO) providing an access point for verticals to request the instantiation of end-to-end services. Services can be selected from a catalogue of blueprints and further customized according to the requirements of the vertical. This customization is expressed in terms of service-

level parameters (e.g., number of users to support), without any detail in terms of infrastructure resources, internal structure of the service or dimensioning of its components. The VS translates such requirements into a slice able to implement the given service and coordinates the provisioning process to optimize the sharing of slices' functions and resources among concurrent services.

In our architecture, network slices are mapped over composite NFV Network Services, which are composed of multiple virtual or physical (network) functions, interconnected through service graphs and organized in hierarchical structures of nested Network Services. Nested services typically provide atomic functionalities and represent network slice subnets, i.e. functional sub-components of the end-to-end network slice. A key feature of nested services is their capability of being shared among multiple, coexisting instances of composite services that require the same kind of functionality. This characteristic allows to re-use virtual resources and virtual functions in concurrent network slices. The VS implements the logic to map vertical services into the most efficient set of network slices, coordinating the sharing of their sub-components and arbitrating the allocation of NFV network services among different vertical service instances, according to policies and SLAs. In particular, the arbitration decisions regulate the dimensioning of network slices and network slice subnets, as well as their composition and automated scaling, based on service priorities and characteristics.

In this demonstration, we show the VS capability to deploy network slices for a number of virtual Content Delivery Network (vCDN) services [2] with different characteristics and requirements. As depicted in Figure 1, the multiple end-to-end network slices for the vCDN will share a common network slice subnet that provides a mobile communication service [3]. This slice subnet is deployed through the instantiation of an NFV network service composed by a VNF running a virtual Evolved Packet Core (vEPC), interconnected to Central Unit (CU) and Distributed Unit (DU) functions, both deployed as Physical Network Functions (PNFs). Depending on the number of users to be served, the VS will adjust the vCDN size requesting a variable number of caches. The NFVO will then be in charge of the actual instantiation of the composite network services, selecting the target resources to meet the requirements at the virtual infrastructure level. For example, the

VNFs running the vEPC and vCDN edge caches will be placed into computing nodes closer to the network edge, where DU and CU are deployed, to guarantee the compliance with the latency constraints.

The demo visitors will interact with the VS web GUI to request the desired vCDN service and verify the network slices instantiated by the system. They will also visualize the video streaming from a mobile phone connected to the DU, verifying the changes in the monitoring data of the network slice due to the increasing traffic.




Figure 1: Composition of network services for deployment of end-to-end networks slices running vCDN services.

2 Setup of the demonstration environment

The demo environment is presented in Figure 2. All the components of the extended MANO framework run as Virtual Machines (VMs) installed in a Lenovo ThinkPad E480 laptop

(Intel® Core™ i7-8550U 1.80GHz with 32GB RAM). More specifically, the MANO deployment includes the prototypes of the VS and the NFVO, also integrating VNF and PNF Manager, the monitoring platform based on Prometheus and Grafana software and the controller of the OpenStack Virtual Infrastructure Manager.

The NFV Infrastructure (NFVI) is a distributed OpenStack deployment (Queens version) that integrates multi-technology computing nodes. In particular, the deployment includes two Intel NUCs (Intel® Core™ i7-8705G, CPU 3.10 GHz, 32GB RAM and 480 GB SSD) to run standard VMs and three Accelerated Processing Units (APUs) [4], more suitable as edge resources for running light virtual functions, e.g. based on unikernels. The OpenStack platform is connected to two external networks, interconnecting to Internet and to the RAN domain respectively. VNFs requiring an external access are configured with floating IP addresses assigned on these networks, where also the PNFs (running in two other NUCs) are attached to. In particular the origin server storing all the CDN contents is attached to the external network towards Internet, while DU and CU are connected to the external network towards the RAN. In our deployment, the CU implements the base-band processing and the DU implements the RF processing following the intra-PHY option 7-1 functional split. The DU is connected to a Universal Software Radio Peripherals (USRPs) Ettus B210, providing the radio access for mobile UEs. All the demo deployment will be installed at the conference premises and it will require 12 power sockets and a table.

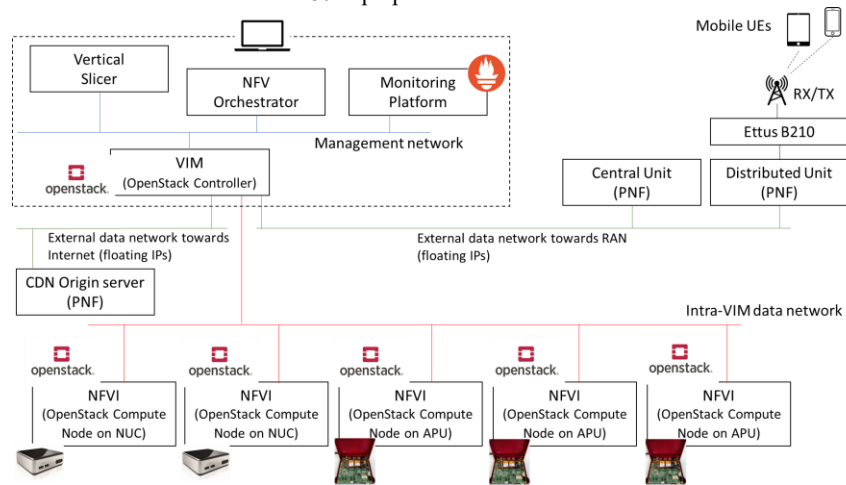


Figure 2: Setup of demonstration environment.

ACKNOWLEDGMENTS

This work has been partially supported by EU H2020 5G-PPP 5G-TRANSFORMER project (Grant 761536) and EU H2020 5G-PPP blueSPACE project (Grant 762055).

REFERENCES

[1] C. Casetti, et al. 2018. "The Vertical Slicer: Verticals' Entry Point to 5G Networks", EuCNC 2018, Ljubljana, Slovenia, June 2018

[2] G. Carrozzo, et al. 2018. "Virtual CDNs over 5G networks: scenarios and requirements for ultra-high definition media distribution", IEEE International Symposium on Broadband Multimedia Systems and Broadcasting 2018, Valencia, Spain, June 2018

[3] M. Capitani et al. 2018. "Experimental Demonstration of a 5G Network Slice Deployment through the 5G-Transformer Architecture", 44th European Conference on Optical Communication (ECOC 2018), Roma, Italy, September 2018

[4] PC Engines, APU2 platform specification, <https://pcengines.ch/apu2.htm>