

This is a postprint version of the following published document:

Zanzi, L., Salvat, J.X., Sciancalepore, V., Saavedra, A.G. y Costa-Perez, X. (2018). Overbooking Network Slices End-to-End: Implementation and Demonstration. In *Proceedings of the ACM SIGCOMM 2018 Conference on Posters and Demos*, pp. 144-146.

DOI: <https://doi.org/10.1145/3234200.3234230>

Overbooking Network Slices End-to-End: Implementation and Demonstration

Lanfranco Zanzi, Josep Xavier Salvat,
Vincenzo Sciancalepore, Andres Garcia Saavedra, Xavier Costa-Perez
NEC Laboratories Europe GmbH
{name.surname}@neclab.eu

ABSTRACT

The novel *network slicing* paradigm allows service providers to open their infrastructure to new business players such as vertical industries. In this demo, we showcase the benefits of our proposed end-to-end network slicing orchestration solution that blends together *i*) an admission control engine able to handle heterogeneous network slice requests, *ii*) a resource allocation solution across multiple network domains: radio access, edge, transport and core networks and *iii*) a monitoring, forecasting and dynamic configuration solution that maximizes the statistical multiplexing of network slices resources. Our orchestration solution is operated through a dashboard that allows requesting network slices on-demand, monitors their performance once deployed and displays the achieved multiplexing gain through overbooking.

CCS CONCEPTS

- Networks → Network management;

ACM Reference Format:

Lanfranco Zanzi, Josep Xavier Salvat, Vincenzo Sciancalepore, Andres Garcia Saavedra, Xavier Costa-Perez. 2018. Overbooking Network Slices End-to-End: Implementation and Demonstration. In *SIGCOMM Posters and Demos '18: ACM SIGCOMM 2018 Conference Posters and Demos, August 20–25, 2018, Budapest, Hungary*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3234200.3234230>

1 INTRODUCTION

Fifth generation mobile networks (5G) will incorporate novel technologies such as *network programmability* and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCOMM Posters and Demos '18, August 20–25, 2018, Budapest, Hungary

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5915-3/18/08...\$15.00

<https://doi.org/10.1145/3234200.3234230>

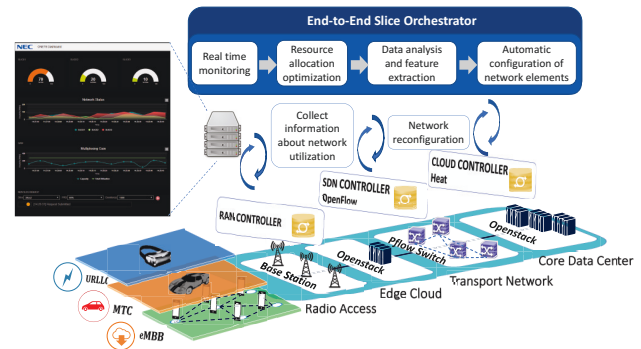


Figure 1: End-to-end Network Slicing Orchestrator

virtualization. In order to do so, the software-defined networking (SDN) and network function virtualization (NFV) paradigms have attracted major interest from academic and industrial stakeholders to enable service providers to open their ecosystem to new business players.

Network Slicing is the main driver of this novel business model where operators may “slice” their infrastructure and offer independent, isolated and self-contained sets of network functions and resources tailored to specific services requirements. End-to-end (E2E) network slices encompass *i*) time/spectrum resources in the radio access network (RAN), *ii*) transport resources on the fronthauling/backhauling links, and *iii*) computing/storage resources at core and edge data centers. Vertical industries—such as automotive, e-health—are considering network slicing as a cost-effective solution for their digital transformation.

While *Network Slicing* has the potential to increase the revenue sources of service providers, it involves a number of technical challenges that must be carefully addressed. An end-to-end network slicing orchestration solution shall *i*) admit network slice requests such that the overall system revenues are maximized [3] and *ii*) provide the required resources across different network domains to fulfill the service level agreements (SLAs) [1].

Fig. 1 illustrates the end-to-end network slicing orchestration concept in a scenario that comprises a centralized (core) data center, e.g., where the traditional evolved packet core (EPC) is deployed, edge data centers, a novel architectural

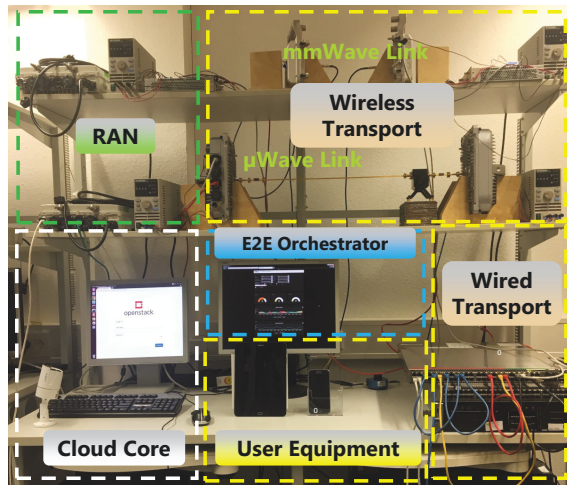


Figure 2: E2E Network Slicing Orchestration Testbed

component to support (mobile) edge computing services, the transport network to connect the data centers to the access network and the radio access points.

In this demo, we present an end-to-end network slicing orchestration solution that applies admission control policies based on a revenue maximization strategy while, at the same time, assigns resources in terms of expected throughput (on the access network), computational resources (on the edge and core network) and latency constraints (on the transport network). Our proposed approach optimally exploits multiplexing gains through machine-learning techniques. By monitoring past slices traffic behaviors [4], our orchestrator forecasts future traffic demands so as to schedule slice resources while pursuing the overall resource efficiency maximization. All operations are displayed in a control dashboard that shows the installed network slices resource utilization as well as the achieved multiplexing gains.

2 DESIGN

We build our demonstration on top of a fully-fledged LTE-compliant network, see Fig. 2. The radio access network comprises two commercial eNBs¹ that support the Multi Operator Core Network (MOCN) RAN sharing model. This allows to broadcast different Public Land Mobile Networks (PLMNs) while being able to reserve radio resources for each particular network. Since there is currently no commercial network slicing equipment, we map network slices onto dedicated PLMNs dynamically installed in our network.

The transport network is composed of mmWave and μ wave wireless links as well as of an OpenFlow programmable

switch² that enables different transport network topology configurations with predefined capacity and delay characteristics. Our deployment also includes two different data centers configured on top of OpenStack deployments to host mobile edge and core networks. Dynamic configurations of computational resources are performed through Heat, an OpenStack orchestration solution. The EPC is realized with OpenEPC 7 [2] and placed as virtualized instance into the core network. Our end-to-end orchestration solution is hierarchically placed on top of three controllers separately managing the radio, transport and core network domains. The controllers dynamically issue resource assignments as well as implement monitoring activities on the respective resources utilization. The gathered monitoring information is promptly fed to the end-to-end orchestrator through REST APIs. This is used as input to optimally solve an admission control and multi-domain resource reservation problem, accounting for resource availability, ongoing slice reservations and upcoming requests.

3 DEMONSTRATION

We use a control dashboard that provides multiple options for requesting network slices. In particular, we can select the slice time duration, the maximum latency allowed, the expected throughput, the price willing to be paid for the network slice and finally the penalty expected in case of SLA violation. With this information our end-to-end orchestration algorithm checks the infrastructure resources availability in each domain and performs traffic forecasting, considering past and current network slices information. Based on this, allocated network slices might be dynamically re-configured (*overbooked*) to accommodate new slice requests.

If successfully accepted, network slices are installed into our system: radio resources (Physical Resource Blocks - PRBs) are reserved through the RAN controller, dedicated paths are selected to guarantee the required delay and capacity in the transport network and, cloud (or mobile edge) data centers are selected to satisfy the network slice SLAs. Thus, OpenEPC instances are deployed and network links dynamically set up to provide connectivity to the end-users. After few seconds, user devices associated with the PLMN-id of the new slices are allowed to connect to the respective services. Conversely, if the network slice requests are rejected, this is shown in the dashboard. The machine-learning engine implemented into the orchestration algorithm trades off between multiplexing gain and SLA violations: our dashboard shows the current gains vs. penalties when multiple network slices are running.

¹NEC LTE small-cell MB4420. Please refer to <http://www.nec.com/en/global/solutions/nsp/sc2/prod/e-nodeb.html>.

²ProgrammableFlow PF5240 Switch. Please refer to <https://www.nec.com/en/global/prod/pflow/pf5240.html>.

REFERENCES

- [1] P. Caballero, A. Banchs, G. de Veciana, and X. Costa-Perez. 2017. Multi-Tenant Radio Access Network Slicing: Statistical Multiplexing of Spatial Loads. *IEEE/ACM Transactions on Networking* 25 (Oct. 2017), 3044–3058.
- [2] CND. 2017. OpenEPC 7. <http://www.openepc.com/>. (2017).
- [3] Konstantinos Samdanis, Xavier Costa-Pérez, and Vincenzo Sciancalepore. 2016. From network sharing to multi-tenancy: The 5G network slice broker. *IEEE Communications Magazine* 54, 7 (Jul. 2016), 32–39.
- [4] Vincenzo Sciancalepore, Konstantinos Samdanis, Xavier Costa-Pérez, Dario Bega, Marco Gramaglia, and Albert Banchs. 2017. Mobile Traffic Forecasting for Maximizing 5G Network Slicing Resource Utilization. In *IEEE INFOCOM*. 1–9.