Experimental Demonstration of End-to-end NFV Orchestration on Top of the ADRENALINE Testbed

Lluis Gifre, Carlos Manso, Ramon Casellas, Ricardo Martínez, Ricard Vilalta, Raul Muñoz
Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels, Spain
E-mail: lluis.gifre@cttc.es

***Abstract*—This demonstration will showcase the end-to-end orchestration of virtual network functions in the full-fledged ADRENALINE Testbed Cloud Platform expanding from the edge to the cloud. The Management and Orchestration (MANO) software ETSI OpenSource MANO (OSM) is used to deploy and handle a multi-site network service involving both edge and core Data Centers (DCs). Besides, the inter- and intra-DC connectivity is directly managed by a novel OSM WAN Infrastructure Manager (WIM) connector using the Transport API (TAPI) interface, thus completely abstracting the details of the underlying SDN controllers handling the programmability of the WAN network interconnecting the DCs.**

***Index Terms*—NFV orchestration, SDN controllers, control interfaces, edge-cloud computing, beyond 5G infrastructures.**

I. INTRODUCTION

The Software-Defined Networks (SDNs) technology is one of the key enablers for Beyond 5G (B5G) networks. It promises to enhance the network dynamicity and programmability with the aim of increasing the efficiency in the configuration and leveraging higher performance and monitoring capacities. However, the reality is that network operators are slowly adopting the SDN-ready solutions, in part, due to the complexity arising on the control of the underlying networking infrastructure. Open Networking Foundation (ONF) Transport Application Programming Interface (API) (TAPI) [1] aims at leveraging a well-defined and complete-enough abstract interface to enable a vendor-agnostic operation of the transport networks. It is an interface for the control of layers 0-2 and is the most-backed SDN control interface by the industry for optical domains. It proposes different data models for the different layers including support for generic Digital Signal Rate (DSR) (TAPI-DSR), for the optical layers (TAPI-PHOTONIC-MEDIA), or for modeling Ethernet networks (TAPI-ETH).

Network Function Virtualization (NFV) aims to move network functions, such as content caches, routers and firewalls, away from proprietary hardware appliances and deploy them on top of virtualized infrastructures as decoupled and softwarized components. The resulting Virtual Network Functions (VNFs) can be conveniently deployed along the cloud network locations where compute capacities are available, and grouped to form network services.

The Management and Orchestration (MANO) of the virtual infrastructures in a Data Center (DC) are usually addressed by means of Virtual Infrastructure Managers (VIMs), such as OpenStack [2]. ETSI Open Source Management and Orchestration (MANO) (OSM) [3], one of the most widely used NFV orchestrators, can manage both single- and multi-site network services involving either a single or multiple VIMs, respectively. It is worth noting that multi-site network services require of appropriate mechanisms to establish the connectivity between the different DCs. That is the mission of the Wide Area Network (WAN) Infrastructure Manager (WIM) connectors. Different WIM connectors, e.g., IETF L2VPN, DynPaC, ONOS OpenFlow, etc, have been released for OSM. However, none of them is able to abstract the details of the underlying WAN domains independently of the technology used.

In our previous work [4], we implemented a similar functionality for the SONATA NFV orchestrator. However, to the best of our knowledge, this is the first experiment integrating TAPI directly into OSM. Our aim is to open-source a WIM connector for OSM able to deal with any TAPI-enabled WAN network, and leverage OSM from dealing with the particularities of the different underlying technologies.

The paper is organized as follows: in section II we present the overview of the demonstration. Then, section III summarizes the key innovations we will showcase. Next, in section IV we present the experimental setup, configurations, and a description of the workflow that will be demonstrated. Besides, it emphasizes the relevance of this demonstration for the NetSoft conference. Finally, section V draws the conclusions of this work and enumerates our future research plans.

II. OVERVIEW

A. Testbed Architecture

The architecture of the ADRENALINE Testbed Cloud Platform [5], as well as the Control, Orchestration and Management (COM) plane are illustrated in Fig. 1. It is composed of 2 packet domains and 1 optical domain. The optical domain is composed of 4 optical cross-connects which are managed by the Open Line System (OLS) [7] controller by means of a REST API. The OLS controller is managed by the child SDN controller using the TAPI interface. The latter also controls the 10Gbit/s XFP transponders using an HTTP interface. This way, the upper-layer systems can provision packet services without the knowledge of the internal architecture of the underlying optical domain. The parent SDN controller manages the child optical-domain SDN controller by means of the TAPI interface

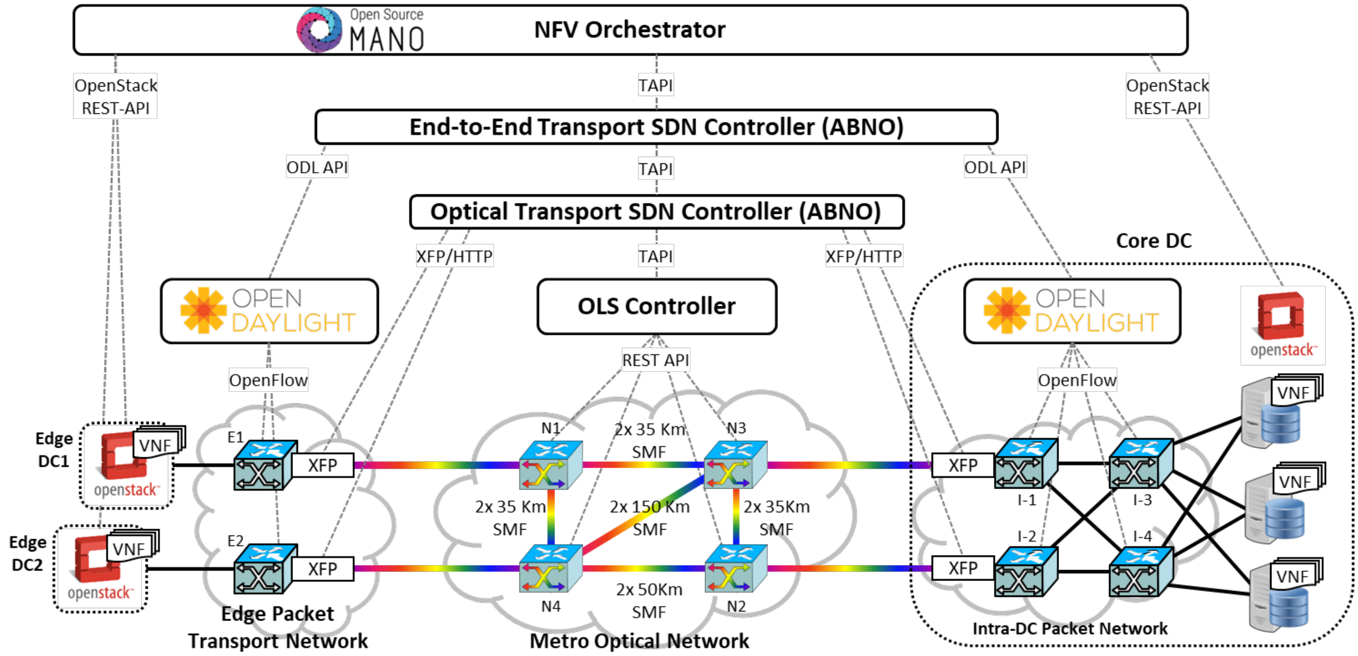

Fig. 1. ADRENALINE Testbed Cloud Platform Architecture

and the OpenDayLight (ODL)-API to control the ODL controllers for the packet domains. Both parent and child SDN controllers follow the architecture described in section II-B. The ODL controllers are in charge of managing the Open vSwitch (OVS) virtual switches on their respective packet domains using the OpenFlow protocol. Each DC (edge and core) is managed by means of an OpenStack VIM.

B. SDN Controller Architecture

The COM plane of the testbed follows a hierarchical architecture with two main SDN controllers in a parent/child relationship. The architecture of both of the controllers is depicted in Fig. 2; it is based on an adaptation of the IETF Application-Based Network Operations (ABNO) [6] architecture, and an implementation carried out by CTTC. As it is shown in the figure, the controller consists of a set of software modules that inter-operate among them. The main module is the Service Orchestrator, which is the module that contains the main workflows and coordinates the other modules. It also acts as the NorthBound Interface (NBI), receiving the requests from the user or from upper layers, using the NBI TAPI plugin. The Virtual Link Manager module is used to create virtual Ethernet links from the DSR connections requested to lower layers. This way, after a virtual link has been created, subsequent Ethernet connections can use the link as part of their routes until there is no more bandwidth available in these virtual links. The Path Computation Element (PCE) is the module that provides the route for the connection, it returns a full route between the two endpoints that are requested. The Transponder Manager is a module used to control the transponders, it features a plugin architecture to support different transponder interfaces.

The Connection Manager is the module that provides the SouthBound Interface (SBI) towards lower controllers. Using different plugins, it can support connections to different kinds of controllers and layers, such as the TAPI plugin for abstracted domains and the ODL plugin for packet domains. The Context Manager is the module that gets the topology from lower controllers, and also provides an interface enable other ABNO modules to retrieve that topological information.

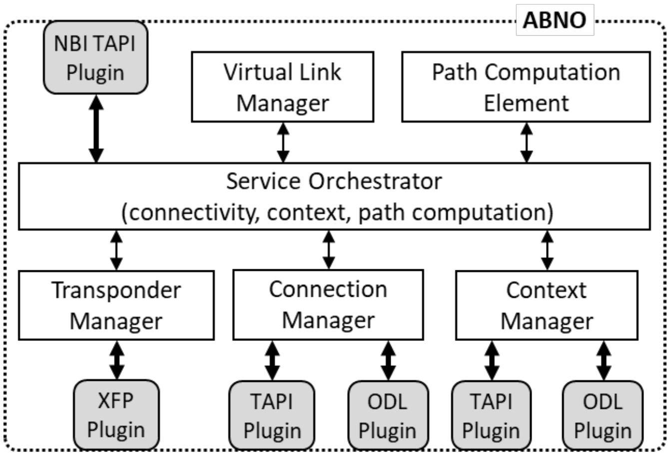

Fig. 2. ABNO-based SDN Controller architecture

III. INNOVATION

This demonstration will showcase, for the first time, how the TAPI driver can be directly exploited by the OSM NFV orchestrator to manage the connectivity between the VIMs involved in a multi-site network service. Besides, TAPI is also used between parent and child SDN controllers, thus enabling network operators to rely on a single interface. Thanks to the use of TAPI, the connectivity establishment request can be abstracted from the underlying networking details simplifying the interconnection of distant VIMs. Nowadays, no other data model brings this level of flexibility to network operators. We think this contribution is of paramount importance for accelerating the adoption of the SDN principles applied to the NFV orchestration.

IV. DEMONSTRATION AND RELEVANCE

The setup, illustrated in Fig. 1, has been deployed on top of the ADRENALINE Testbed Cloud Platform. The edge DC VIMs runs MicroStack Ussuri (R245) on top of Ubuntu 20.04.4 LTS, while the core DC runs OpenStack Xena on top of CentOS Linux 7.

Regarding the packet layer, ODL v0.7.3 is used to control the packet networks, composed of OVSs v2.10.0, through the OpenFlow v1.3 protocol. The 10Gbit/s XFPs transponders are configured through an XFP HTTP interface.

The Transport SDN controllers (end-to-end WIM, and optical layer controller) are implemented in Python, and implement the TAPI interface v2.1, as well as the appropriate plugins for XFP and ODL API. On top of the network control plane and the VIMs, OSM v10.0 has been deployed extended with our contributed change 11733 [8] (under review) enabling automatic selection of a feasible WIM to interconnect VIMs, and our WIM TAPI connector (to be contributed).

We use the multi-site "hackfest_multivdu" network service provided by OSM for this demonstration as an example to illustrate the flexibility of TAPI. This network service, depicted in Fig. 3, consists of two VNFs, each to be deployed in a different DC. While VNF1 is deployed at Edge DC1 (in Fig. 1), VNF2 is deployed at Core DC. Each of the VNFs contains two Virtual Deployment Units (VDUs), implemented as one Virtual Machines (VMs) each, named as "mgmtVM" and "dataVM", representing the management unit and the data processing unit of the VNF. Both of the VMs within the VNF are interconnected by means of an internal virtual network. The VNFs are managed through an external management network pre-configured in the VIMs, named as "mgmt-net". The "data-net" external network provides the data-plane connectivity between the VNFs in the different DCs. The WIM is the responsible for establishing the L2 inter-DC connectivity through the "data-net" network by means of the hierarchy of SDN controllers. This L2 connectivity supports the virtual L3 networks interconnecting the distant VNFs.

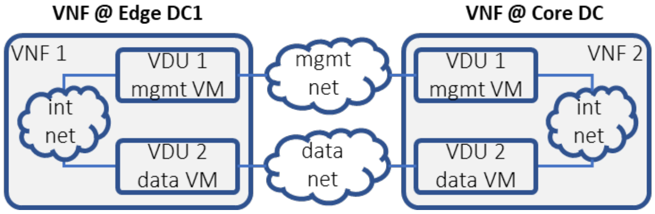

Fig. 3. Hackfest Multi-Site MultiVDU Network Service

This demonstration will illustrate the complete process to deploy the aforementioned multi-site network service including both the operational requests issued to OSM, as well as the traffic captures involving the TAPI requests triggered by OSM. Through the traffic captures we will empirically assess the performance in terms of the contribution of the control plane in the deployment elapsed time. Besides, we will illustrate the configurations implemented in the different virtual infrastructure and transport network elements.

This demonstration is very relevant and adheres very well with the scope of the NetSoft conference. We expect this demonstration can bring light and motivate further discussions in the research and industrial communities in charge of developing the COM planes as well as the NFV orchestrators.

V. CONCLUSION AND FUTURE WORK

This demonstration will showcase an end-to-end orchestration of a multi-site network service on top of the realistic ADRENALINE Testbed Cloud Platform. ETSI OSM will be used to orchestrate the virtual infrastructure resources, while a hierarchical COM plane, mainly based on the ABNO architecture, will be used to establish the WAN connectivity between the involved sites. The management of the WAN infrastructure will be carried out by means of a novel WIM connector for OSM implementing the TAPI interface, thus abstracting the details of the underlying WAN network. This WIM connector will be open-sourced and contributed to the OSM v10, v11 and the master branches for future releases.

Our future work plans include the integration of connectivity service constraints to OSM to bring Quality of Service (QoS) assurance specifying service and traffic requirements such as guaranteed capacity, maximum permitted latency and/or jitter, etc. Thanks to the novel WIM connector based on TAPI, we expect the integration of these QoS constraints within the OSM core will be simplified regardless of the underlying network domains and technologies.

ACKNOWLEDGMENT

This work is partially funded by the EC through the EC/5GPPP TeraFlow PROJECT (101015857) and the "Grant RTI2018-099178-B-I00 funded by MCIN/AEI/ 10.13039/501100011033 and by ERDF A way of making Europe".

REFERENCES

[1] "ONF Transport API (TAPI) GitHub," May 2022. Online: https://github.com/OpenNetworkingFoundation/TAPI
[2] "OpenStack," May 2022. Online: https://www.openstack.org/
[3] "ETSI OpenSource MANO," May 2022. Online: https://osm.etsi.org/
[4] C. Manso, P. Alemany, R. Vilalta, R. Muñoz, R. Casellas, and R. Martinez, "End-to-end SDN/NFV orchestration of multi-domain transport networks and distributed computing infrastructure for beyond-5G services," IEICE Transactions on Communications, Vol. E104-B, No. 3, March 2021.
[5] "ADRENALINE Testbed at CTTC," May 2022. Online: https://www.cttc.cat/adrenaline-testbed/
[6] D. King, A. Farrel, "A PCE-based architecture for application-based network operations," IETF RFC 7491, March 2015.
[7] R. Casellas, F.J. Vilchez, L. Rodriguez, R. Vilalta, J.M. Fabrega, R. Martinez, L. Nadal, M. S. Moreolo, and R. Muñoz, "An OLS controller for hybrid fixed/flexi grid disaggregated networks with open interfaces," 2020 Optical Fiber Communications Conference And Exhibition (OFC), pp. 1-3, March 2020.
[8] "ETSI OSM Gerrit - LCM Change 11733," May 2022. Online: https://osm.etsi.org/gerrit/#/c/osm/LCM/+/11733/