Experimental Demonstration of Virtual Network Controller for Abstraction and Control of Multi-tenant Multi-technology Transport Networks

Ricard Vilalta, Arturo Mayoral, Raul Muñoz, Ramon Casellas, Ricardo Martínez
Optical Networks and Systems Department
Centre Tecnològic de Telecomunicacions de Catalunya, CTTC
Av. Carl Friedrich Gauss 7, 08860 Castelldefels, Spain
Email: ricard.vilalta@cttc.es

***Abstract*—A network hypervisor is demonstrated. The Virtual Network Controller dynamically deploys multi-tenant virtual networks on top of multi-technology transport networks. It provides an abstract view of each virtual network and enables its control through an independent SDN controller.**

I. Introduction

The Software Defined Networking (SDN) paradigm proposes the decoupling of network data and network control. In SDN, the network control is abstracted from Network Equipment(NE), and the network intelligence is centralized in a SDN Controller. The SDN Controller is responsible for configuring the NE with a network protocol, such as OpenFlow (OF) [1].

SDN Controllers are currently commercialized by Network Product Manufacturers (vendors) and they do not enable network interconnection. It is in this context where the Virtual Network Controller (VNC) is proposed. It is a network hypervisor (a network virtualization enabler), such as happens in computers and virtual machines.

The VNC Network Hypervisor can run on top of heterogeneous network control systems (e.g., Network Management Systems, MPLS control, GMPLS/PCE control or SDN control). This property brings a multi-domain multi-vendor solution, which is highly appreciated by network operators and also can be reused by Data Center (DC) operators. Moreover, the VNC offers virtual network slices which can be controlled by independent instances of SDN control plane. This feature enables multi-tenancy of the underlying network, which is known as network slicing and network virtualization. With the VNC network and DC operators can provide to each customer a dedicated Virtual Network (VN). The VNC has been first proposed by the authors at [2].

The VNC is responsible for: a) the interaction with a Network Orchestrator (NO) [3] for the deployment of the necessary multi-domain connectivity, b) the provision of a virtual network by enabling the abstracted network view to its customer (i.e., tenant) SDN controller, and c) the translation of the issued OF commands from the customer SDN controller to the necessary actions in the multi-domain network with interaction with the NO.

In this demonstration, we present the VNC, which allows the deployment of OF-enabled multi-tenant VNs across a multi-domain multi-technology transport infrastructure with heterogeneous control plane technologies. Once a requested VN is deployed, it is controlled with a customer SDN controller.

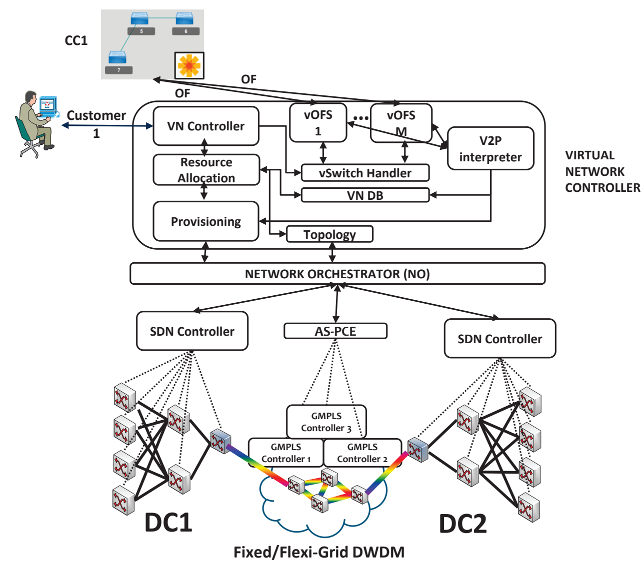

Fig. 1: Experimental demonstration scenario and VNC architecture

II. Experimental demonstration

Fig. 1 shows the proposed VNC system architecture. Three hierarchical control levels are identified: Customer Controller (CC), VNC&NO and Physical Controller (PC). A CC is a SDN controller run by a VN customer for controlling its deployed VN. The VNC and NO are the central components of the virtualization architecture. Finally, a PC is the centralized instance of control in charge of a physical infrastructure (i.e., SDN controller or Active Stateful Path Computation Element - AS-PCE). The PC's northbound interfaces (NBI) are typically technology and vendor dependent, so the NO shall implement different PC plugins for each of the PC's NBI. It is assumed that the PC's NBI are able to provide network topology information and flow programming functionalities.

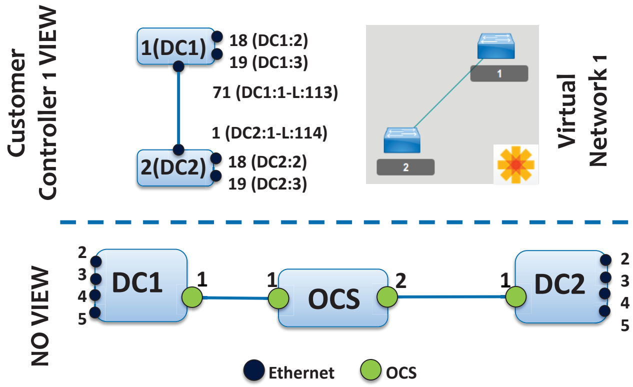

Fig. 2: Virtual Network resources allocation

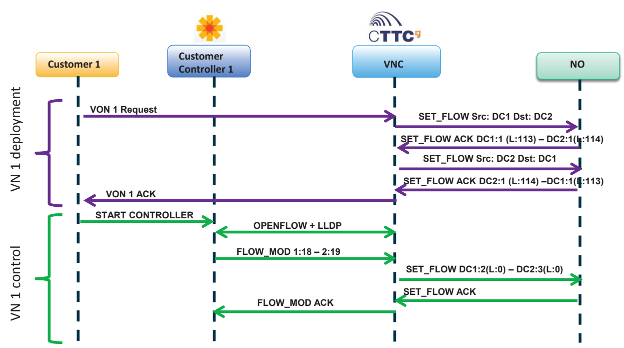

Fig. 3: Experimental workflow

To demonstrate the proposed VNC architecture, Fig. 1 shows the topology of the experimental scenario. Using mininet we have set up two intra-DC Ethernet networks in flat-tree. A GMPLS-controlled Fixed/Flexi-Grid DWDM network is also emulated using CTTC GMPLS control plane. An AS-PCE is located on top of GMPLS-controlled Fixed/Flexi-Grid DWDM as an SDN enabler.

The NO abstracted topology consists of two emulated DC (packet domain) interconnected with a single GMPLS-controlled Fixed/Flexi-Grid DWDM domain. Each domain has been abstracted as a node, following the previous work on the NO [4]. A VN is requested. Fig. 2 shows the CC1 view. The relationship between the virtual Ethernet switches and links and the physical domains is provided of the established connections.

The workflow to deploy and control VN1 is detailed in Fig. 3 Customer 1 triggers the deployment of VN1, by issuing a request to the VNC. Then VNC requests the necessary multi-domain connectivity for deploying VN1 from NO. In the presented scenario, two unidirectional multi-domain connections are requested: DC1-DC2, DC2-DC1. The connections are established by NO, which provides the port interfaces and the assigned labels. The procedure followed by the NO has been previously detailed in [4]. Once the necessary LSPs have been established and the physical network resources have been allocated and mapped to abstract resources, the VNC creates the OF-enabled VN1 and notifies Customer 1.

Once VN1 has been deployed, Customer 1 can connect CC1 (e.g., using ODL) to VNC in order to control it. CC1 and VNC start the different virtual OF switches datapaths for VN1. CC1 can also trigger a topology discovery protocol (i.e., LLDP) to discover VN1 topology. If CC1 triggers an OF command to a virtual OF switch, the VNC receives the OF command and translates it to the necessary actions to be performed by the NO.

Fig. 3 shows the captured messages for VN1 deployment and control. Firstly, we observe the different requested LSPs from VNC to NO. Once the VN1 network resources have been allocated, the communication between the CC1 and VNC is established by means of OF protocol. Secondly, when VN1 has been deployed and CC1 issues an OF command, we observe how this command is processed by the VNC, an action is requested to NO and finally the OF command is acknowledged to CC1.

III. Conclusion

We have demonstrated a Virtual Network Controller that can dynamically deploy and control multi-tenant VNs on top of a multi-domain multi-technology transport networks. Customer control for its own deployed VN is enabled and demonstrated.

Acknowledgment

This work was partially funded by EU FP7 STRAUSS (FP7-ICT-2013-EU-Japan 608528), EU FP7 PACE (619712), and Spanish MINECO project FARO (TEC2012-38119).

References

[1] N. McKeown, et al., "OpenFlow: enabling innovation in campus networks", *ACM SIGCOMM Computer Communication Review*, Vol. 38, Num. 2, 2008.

[2] R. Vilalta, et al., "Network Virtualization Controller for Abstraction and Control of OpenFlow-enabled Multi-tenant Multi-technology Transport Networks", accepted at *OFC Conference*, 2015.

[3] D. King, and A. Farrel, "A PCE-based Architecture for Application-based Network Operations", IETF draft *draft-farrkingel-pce-abno-architecture-13*, 2014.

[4] R. Muñoz, et al., "Experimental Assessment of ABNO-based Network Orchestration of end-to-end Multi-layer (OPS/OCS) Provisioning across SDN/OpenFlow and GMPLS/PCE Control Domains", in *ECOC* 2014.