Network Virtualization Controller for Abstraction and Control of OpenFlow-enabled Multi-tenant Multi-technology Transport Networks

R. Vilalta[1], R. Muñoz[1], R. Casellas[1], R. Martínez[1], F. Francois[2], S. Peng[2], R. Nejabati[2], D. Simeonidou[2], N. Yoshikane[3], T. Tsuritani[3], I. Morita[3], V. López[4], T. Szyrkowiec[5], A. Autenrieth[5]

[1] *Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Castellefels (Barcelona), Spain.* ricard.vilalta@cttc.es
[2] *High-Performance Networks Group, University of Bristol, UK*
[3] *KDDI R&D Laboratories Inc., 2-1-15 Ohara, Fujimino-shi, Saitama 356-8502, Japan*
[4] *Telefónica I+D, Madrid, Spain*
[5] *ADVA Optical Networking SE, Martinsried, Germany*

Abstract: A network hypervisor is introduced to dynamically deploy multi-tenant virtual networks on top of multi-technology optical networks. It provides an abstract view of each virtual network and enables its control through an independent SDN controller.
OCIS codes: (060.4250) Networks; (060.4258) Networks, network topology

1. Introduction

Transport networks include different optical technologies (e.g., flexi-grid, OPS) and heterogeneous control technologies (e.g., OpenFlow -OF-, GMPLS/PCE), which enable flexible and dynamic networking but also create a multi-domain multi-technology network scenario in order to support end-to-end service provisioning. Such a transport network scenario raises challenges for efficient network control and operations. An SDN-enabled virtualization platform that allows abstracting and controlling a multi-domain multi-technology transport network is able to support rapid deployment of applications and network services and improve overall network operations. The creation and operation of multi-tenant virtual networks also allows better network utilization.

Optical Network Virtualization has been respectively studied in a single GMPLS-controlled domain [1] and an OF domain [2]. In addition, the provision of a multi-domain Virtual Optical Network (VON) has also been proposed with the usage of per-domain virtualization [1]. A Network Orchestrator (NO) was introduced in order to provision end-to-end connectivity within a multi-domain VON.

In order to achieve multi-domain virtualization, we propose the Network Virtualization Controller (NVC) as a network hypervisor to deploy multi-tenant VONs and allow their own customized control. The NVC runs over a NO, as shown in Fig. 1.a. NO is the responsible for providing end-to-end connectivity on top of multi-technology transport networks. The NVC is in-line with ACTN framework (IETF) [3] and SDN architecture (ONF) [4], where hierarchical SDN controllers have been considered to provide different layers of abstraction.

The NVC is responsible for: a) the interaction with a Network Orchestrator (NO) [5] 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 paper, we present an experimental assessment of the NVC, which allows the deployment of OF-enabled multi-tenant VONs across a multi-domain multi-technology transport infrastructure with heterogeneous control plane technologies. Once a requested VON is deployed, it is controlled with a customer SDN controller.

2. Network Virtualization Controller System Architecture

Fig. 1.a shows the proposed NVC system architecture. Three hierarchical control levels are identified: Customer Controller (CC), NVC&NO and Physical Controller (PC). A CC is a SDN controller run by a VON customer for controlling its deployed VON. The NVC 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). 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.

The NVC is responsible for receiving VON requests, processing them and allocating physical resources. Moreover, the NVC is responsible for the mapping between the allocated physical resources and the abstracted resources that are offered to the CCs, and the control of such abstract networks, acting as a proxy for the OF protocol between a CC and the underlying Physical Controllers. The partitioning of the resources is performed by the NVC, and to this

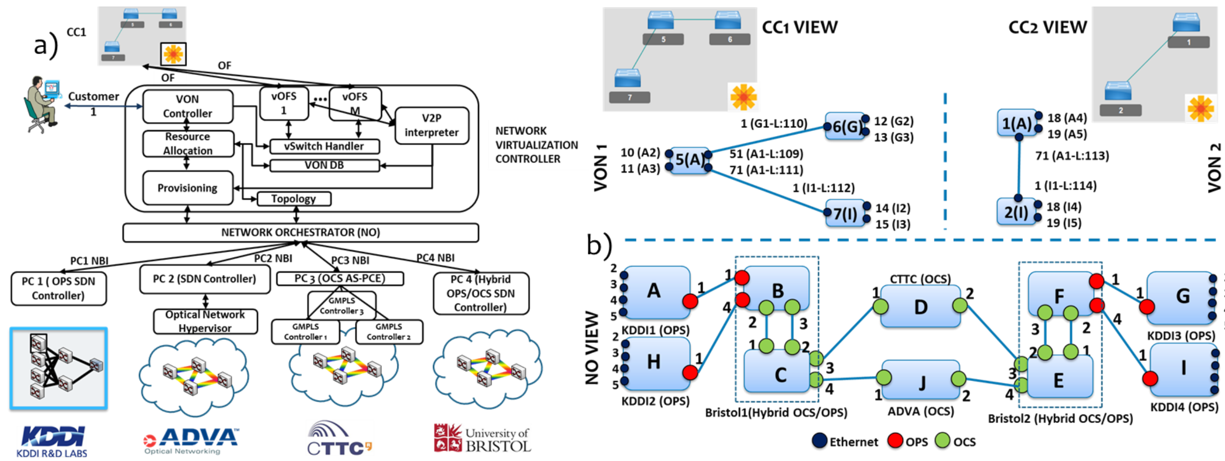

Fig.1. a) NVC system architecture b) Network Orchestrator view, VON 1 and 2 and their relationships with the physical topology

end, the proposed system architecture relies on the NO, which provides a generic network abstraction mechanism for the different transport infrastructure resources (e.g., OPS, flexi-grid).

Each tenant is able to request a VON. Once the VON has been correctly setup, the CC acts as a standard SDN controller where the controlled VON is an abstracted slice of the different allocated physical resources, which are managed by its corresponding PC.

The NVC architecture is as follows. The VON Controller is the component that is responsible for providing the NVC interface to request virtual switches and virtual links to deploy a VON. To do so, also the IP address of the CC is necessary, so that the Virtual Switch Handler is able to provide an abstract network view of the allocated VON to the CC. A virtual switch request includes the related physical domains (abstracted as nodes by the NO) and a number of virtual Ethernet ports. A virtual link request includes the source and destination virtual switches. The Resource Allocation (RA) component is responsible for the allocation of the physical ports of the physical domains to the virtual switches and to request to the NO (through the provisioning component) the necessary multi-domain connections to interconnect the requested virtual switches, which are related to physical domains. Once the connections have been established, the RA allocates the virtual port identifiers, to which the connections are related.

For each VON, the Virtual Switch Handler establishes the necessary OF datapaths with the provided IP address of the corresponding CC. Each OF datapath is provided by an emulated OF virtual switch. The different emulated OF virtual switches are interconnected with virtual links, so when the CC triggers the LLDP to the emulated virtual switches, it is able to recover the VON topology. The emulated virtual OF switches are connected to the Virtual to Physical (V2P) Interpreter, which is the responsible to translate the received OF command (e.g., FLOW_MOD) from the CC using the abstract VON topological view, to the allocated physical resources. To this end, it consults the VON Database for the allocated physical ports and the established LSPs. The processed requests are sent to the provisioning module, which is the responsible to request the provisioning of the physical resources to the NO.

3. Experimental assessment

To experimentally validate the proposed NVC architecture, Fig.1.b shows the NO view of the experimental scenario. The NO topology consists of 4 OPS domains (KDDI R&D Labs), 2 OPS/OCS domains (University of Bristol), and single OCS domains (CTTC and ADVA). Each physical domain has been abstracted as a node, following the previous work on the NO [5]. Two different VONs are requested. Fig. 1.b shows the different CC1 and CC2 views. The relationship between the virtual Ethernet switches and links and the physical domains is provided, including OPS Labels (e.g., A1-L:109 indicates Domain A Port 1 OPS label out: 109) of the established connections.

The workflow to deploy and control VON1 is detailed in Fig. 2.a. Customer 1 triggers the deployment of VON1, by issuing a request to the NVC. Then NVC requests the necessary multi-domain connectivity for deploying VON1 from NO. In the presented scenario, two bidirectional multi-domain connections are requested: A-G, A-I. The connections are established by NO, which provides the port interfaces and the assigned labels (A:1-G:1/L:109-L:110 and A:1-I:1/L:111-L:112). The procedure followed by the NO has been previously detailed in [5]. Once the necessary LSPs have been established and the physical network resources have been allocated and mapped to abstract resources, the NVC creates the OF-enabled VON1and notifies Customer 1.

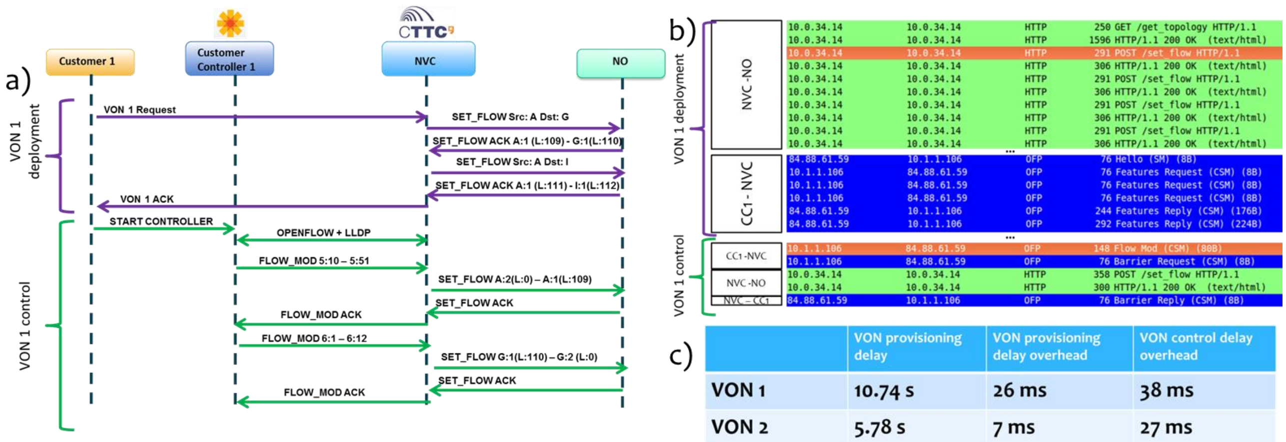

	VON provisioning delay	VON provisioning delay overhead	VON control delay overhead
VON 1	10.74 s	26 ms	38 ms
VON 2	5.78 s	7 ms	27 ms

Fig.2.a) NVC workflow b) Wireshark capture between NVC, NO and CC1 c) VON establishment and control average measurements

Once VON1 has been deployed, Customer 1 can connect CC1 (e.g., using ODL) to NVC in order to control it. CC1 and NVC start the different virtual OF switches datapaths for VON1, following standard OF 1.0 protocol. CC1 can also trigger a topology discovery protocol (i.e., LLDP) to discover VON1 topology. If CC1 triggers an OF command to a virtual OF switch, the NVC receives the OF command and translates it to the necessary actions to be performed by the NO. For example a FLOW_MOD command with match (virtual switch 5, input port 10) action (output port 51), it is translated to an instruction to NO to interconnect the ports 2 and 1 of physical domain A, with labels 0/109. Once this action is processed by the NO, the NVC sends a FLOW_MOD acknowledge to CC1.

We have developed a proof-of-concept prototype of the NVC, which sits on top of the NO of the international testbed of the STRAUSS project as shown in Fig.1.a, in order to experimentally assess the multi-tenant transport network virtualization across multiple heterogeneous SDN/OpenFlow and GMPLS/PCE domains. The experimental assess is performed at the control plane level (i.e., no hardware configuration is performed at the domains). Each domain's PC is connected to the NVC and NO through OpenVPN. Eight domains are involved in the experiment: four OF-controlled OPS domains from KDDI R&D Labs, two OF-controlled OPS/OCS domains from BRISTOL, a GMPLS-controlled OCS domain with active stateful PCE from CTTC and an OF-controlled domain from ADVA, as shown in the network topology in Fig.1.b. The NVC has been implemented using the virtual switch of POX. The different NVC components have been implemented as POX components and a REST API has been provided. The NO has been previously detailed in [5].

Fig. 2.b. shows the captured messages for VON1 deployment and control. Firstly, we observe the different requested LSPs from NVC to NO. Once the VON1 network resources have been allocated, the communication between the CC1 and NVC is established by means of OF protocol. Secondly, when VON1 has been deployed and CC1 issues an OF command, we observe how this command is processed by the NVC, an action is requested to NO and finally the OF command is acknowledged to CC1.

The measurements for the deployment and control of VON1 and VON2 are shown in Fig.2.c. The VON provisioning delay is in average around 10.7s and 5.7s for VON1 and VON2, respectively. We have defined the VON provisioning delay overhead as the required time for the NVC to set up a virtual OF VON, which is in the order of milliseconds. Finally, the VON control delay overhead is the delay introduced by the OF commands translation process at the NVC, which is also in the order of milliseconds.

4. Conclusions

We have proposed and experimentally assessed a Network Virtualization Controller that can dynamically deploy and control multi-tenant VONs on top of a multi-domain multi-technology transport network. Customer control for its own deployed VON is enabled. The proposed architecture has been validated over an international testbed.

5. Acknowledgments

EU FP7 STRAUSS (FP7-ICT-2013-EU-Japan 608528) and Spanish MINECO project FARO (TEC2012-38119).

References

[1] R. Vilalta et al., "Dynamic Multi-domain Virtual Optical Networks Deployment with Heterogeneous Control Domains", M3H.4, OFC14.
[2] A. Autenrieth et al., "Evaluation of Virtualization Models for Optical Connectivity Service Providers", S9_2, ONDM 2014.
[3] D. Ceccarelli et al., "Framework for Abstraction and Control of Transport Networks", IETF draft, 2014.
[4] ONF, "SDN Architecture Issue 1", 2014.
[5] 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", We.3.2.4, ECOC 2014.