Performant Deployment of a Virtualised Network Functions in a Data Center Environment using Resource Aware Scheduling

Michael J. McGrath, Vincenzo Riccobene, Guiseppe Petralia
Intel Labs Europe,
Leixlip, Co. Kildare, Ireland

Georgios Xilouris, Michail-Alexandros Kourtis
Institute of Informatics and Telecommunications,
NCSR "Demokritos",
Athens, Greece

***Abstract*—The EU funded FP7 project T-NOVA, with the specific goal of accelerating the evolution of NFV, proposes an open architecture to provide Virtual Network Functions as a Service (VNFaaS), together with a dynamic, and flexible platform for the management of Network Services (NSs) composed by those Virtual Network Functions (VNFs). The proposed architecture allows operators to deploy distinct virtualized network functions, not only for their internal operational needs, but also to offer them to their customers, as value-added services. Virtual network appliances (e.g. gateways, proxies or even traffic analyzers) can be provided on-demand, eliminating the need to acquire, install, and maintain specialized hardware at customer premises. This demo illustrates early work carried out on the deployment of a VNF on a Network Function Virtualization Infrastructure (NFVI) using resource aware scheduling methods to ensure optimal use of resources and performance.**

Keywords—NFV, VNF, virtualisation, NFaaS, T-NOVA;

I. Introduction

Network Functions Virtualization (NFV) has received significant interest an as approach that can address many of the key challenges being experienced by service providers. These challenges are being driven by an exponential growth in data volumes with a corresponding fall in revenues per megabit. Many service providers are currently investigating and in some cases deploying VNFs to replace traditional high cost fixed appliance based architectures particularly for edge of network applications.

Virtualization is the key enabling technology that allows traditional physical network functions to be decoupled from fixed appliances by leveraging standard IT virtualization technologies to consolidate various network equipment types onto industry standard high volume servers, switches and storage located in DCs. This approach will allow service providers to innovate through the rapid deployment of new services, increased customization and flexibility to meet diverse customer needs, increased utilization of capital resources, etc.

The EU funded FP7 T-NOVA project [1-2] is focused on realizing the concept of Network Functions as a Service (NFaaS) by designing and implementing an integrated management architecture, for the automated provision, management, monitoring and optimization of VNFs over Network/IT infrastructures. As a technology, NFV encompasses a wide variety of network functions, which have a diversity of resource requirements. As a result T-NOVA is carrying out research to develop an understanding of workload types and their affinities for certain platform features and technologies.

II. Motivation

While virtualisation brings many benefits to Enterprise IT and more recently to the Telecom domain it also brings many challenges particularly in achieving the same level of performance in comparison to the traditional fixed appliance approach. The composition, configuration and optimization of the virtualised resources are critical in achieving the required levels of performance. Additionally given the origins of many virtualisation technologies such as cloud computing environments there are capability gaps that need to be addressed in order to adequately support VNF/NS type workloads.

Currently, within cloud environments resources are highly abstracted, which again causes issues for the performant deployment of VNFs and Network Services (NSs). For example, it is important to expose the specific platform features, such as unique CPU instructions and attached devices, such as acceleration cards, co-processors or Network Interface Cards (NICs) with advanced capabilities. Additionally many hardware devices, such as NICs, have additional dependencies, such as the availability of supporting software libraries - e.g. DPDK - in order for a VNF to function in an optimal manner.

The demo presents initial work by the T-NOVA project to address some of the limitations of cloud compute environments with respect to scheduling of virtualised resources to host VNFs in a performance manner.

III. Demo Description

This demo illustrates how an NS with a single VNF (in this case a network traffic classification VNF) can exploit the platform requirements specification contained in an ETSI ISG NFV VNF descriptor (VNFD) accompanying it, in order to deploy and appropriately configure the virtualized infrastructure resources in a performant manner as proposed by the project. The Orchestrator parses the VNFD into corresponding metadata are, which is then used to dynamically construct a Heat template. The template is then used in an OpenStack Cloud environment to automatically instantiate the required virtualized resources with specific platforms features that cannot be predictably provisioned using OpenStack's current scheduling mechanism. Finally a side-to-side comparison of the performance of the virtual Traffic Classification VNF's running on 'standard Virtual Machines' (VMs) versus VMs with features to improve packet processing performance is demonstrated using real-time instrumentation.

The network virtualization process starts at the lower levels of the ISO/OSI protocol stack: VMs require vNIC/NICs to provide connections to switches that in turn manage various Virtual Local Area Networks (VLANs). Virtualization of network functions typically involves two approaches; software-assisted and hardware-assisted.

In software-assisted network virtualization, communications are provided by hypervisors through virtual switches (vSwitches) and protocols designed to coordinate and manage virtualized network architectures at the edge of the network. However, despite the flexibility of a software-assisted approach, a significant disadvantage is the potential for the virtual switch to act as a bottleneck with increasing number of guests running on the host and network traffic volumes.

In hardware-assisted network virtualization, physical hardware is directly assigned to the virtual guests in order to both increase the performance and avoid the bottlenecks. Intel's suite of input/output virtualizations technologies, called virtualization Technology for Connectivity (VT-c) is an example of such an approach, which is complementary to Intel's VT-d (Virtualization Technology for Directed I/O). It includes: Data Plane Development Kit (DPDK), Virtual machine Device Queues (VMDq) and Virtual Machine Direct Connect (VMDc). The latter is implemented using the PCI-SIG standard called Single Root I/O Virtualizations (SR-IOV), which enables a single PCI Express (PCIe) network adapter to appear as many special-purpose adapters, called Virtual Functions (VFs) that are available for direct presentation to VMs, through VT-d.

The demo architecture shown in Figure 1 comprises of an OpenStack Juno based cloud environment. The cloud environment comprises of a controller and two compute nodes, and a traffic generator connected on the same network domain through a 10 Gbps switch. The compute node configuration enables the VMs to use both vNICs connected to Open virtual Switch (OvS) in the form of software-assisted solution and physical NICs with PCIe passthrough/SR-IOV functionality in the form of a hardware-assisted solution.

On the same host as the controller, a simulated Orchestrator is running which receives a deployment request (which includes an ETSI compliant VNFD) and converts it into metadata representing the platform deployment requirements of the VNF, which is used to dynamically generate a Heat template. The template orchestrates the setup SR-IOV ports, deployment of Traffic Classification VNF Components (VNFCs) and configuration of the VNFCs The VNF used in this demo comprises of two VNFCs, namely the DPI engine and the Classification and Forwarding functionality. The VNFCs are implemented and contained in the form of two VMs. For the purpose of the demo, two different versions of the VNFD are utilized: the first version contains no platform specific features allowing the Orchestrator request a 'standard' deployment (based on OvS) of the DPI; the second version contains specific platform features i.e. an SR-IOV capable NIC which must be available on the physical sever were the VMs are hosting the vDPI.

Once the two different instances of the Traffic Classification VNF are running, traffic is generated by the packet generator and sent to both the instances. The performance of both deployments is compared using real-time display of parametric data from the VMs captured using embedded instrumentation agents in the VMs to highlight differences in packet processing performance.

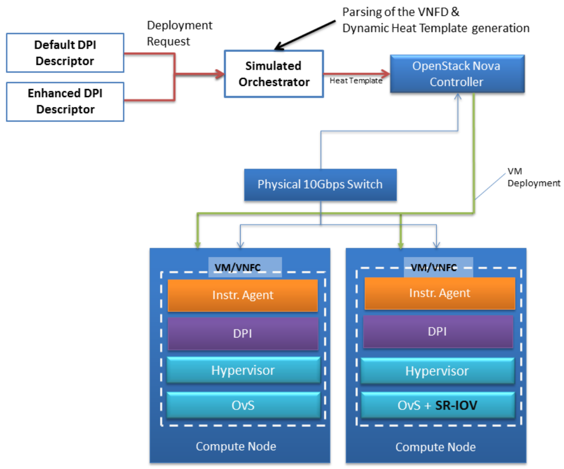

Fig. 1. Architecture of the demo testbed.

ACKNOWLEDGMENT

This work was undertaken under the Information Communication Technologies, EU FP7 T-NOVA project, which is partially funded by the European Commission under the grant 619520.

REFERENCES

[1] T-NOVA Consortium, Deliverable 2.21- Overall system architecture and Interfaces, June 2014, on-line:http://www.t-nova.eu/results

[2] G. Xilouris et. al, T-NOVA: A Marketplace for Virtualized Network Functions, European Conference on Networks and Communications, Bologna, Italy, June 2014.