WHITE PAPER

Network Functions Virtualization

Mobile World Congress 2016

End-to-End Service Instantiation Using Open-Source Management and Orchestration Components

An open-source orchestration solution based on published information and data models will unleash an OpenNFV ecosystem that addresses multiple use cases for carrier-grade network functions virtualization (NFV).

Executive Summary

With the objective of furthering their goal of fostering an open software ecosystem and cost-effective, rapid service deployment related to network functions virtualization (NFV), lead operators Telefónica, British Telecom, Telenor, and Telecom Austria Group have joined with key NFV cloud solution providers Intel, Canonical, RIFT.io, and Mirantis as founding members to form a new ETSI-hosted open-source project called Open Source MANO (OSM).

The project will focus on evolving an ETSI NFV ISG-compliant information model (IM) to an implementable data model, which will in turn drive an OSM code base to address multi-use-case NFV. Such an open-source implementation begins to realize the promise of an open ecosystem of virtualized network function (VNF) vendors, by working toward a uniform deployment model and providing the path to true multi-use-case network services (NS), which can in turn be deployed by the operator community at the lowest possible integration cost.

OSM´s scope (see Figure 1) encompasses both resource and service orchestration, and it will enable the automated deployment and interconnection of all components for the lead NFV network scenarios as well as network service lifecycle management.

Mapping to ETSI

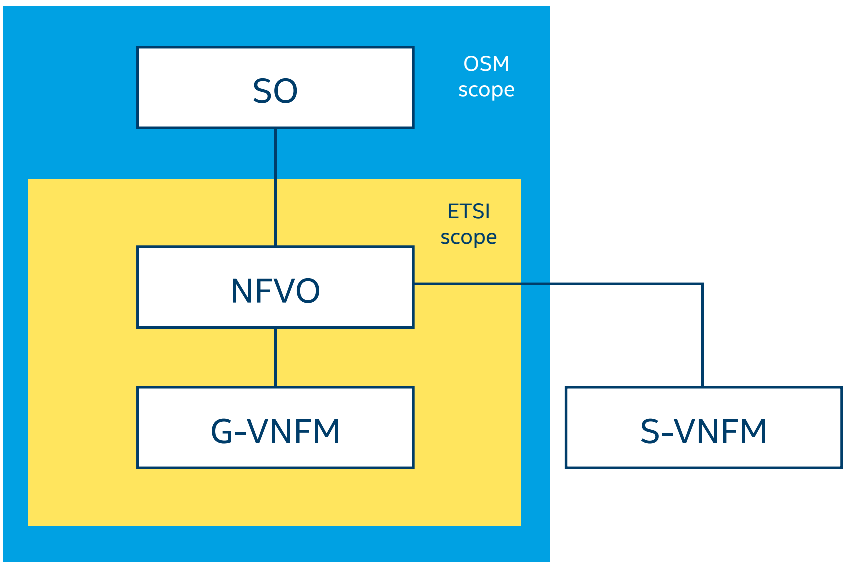

Figure 1. Service orchestration.

The OSM project demonstration described in this paper has been demonstrated at Mobile World Congress 2016 and implements a significant and complex operator use-case deployment. The use case demonstrates that through the use of proper design principles, even the initial implementation of OSM can be adopted to deploy a high-performance inter-site router infrastructure. This in turn enables value-added virtual services such as corporate VoIP (instantiated as an IMS deployment). This entire infrastructure is deployed and configured in minutes using the OSM stack, relying exclusively on development driven by VNF and NS descriptor models. This avoids the need for expensive and time-consuming customization on a per-use-case basis, at the network functions virtualization infrastructure (NFVI), network functions virtualization orchestration (NFVO), or OSS levels.

The implementation showcases the power in the key guiding architectural principles of OSM—layering, abstraction, modularity, and simplicity—which enabled five separate vendors to come together and integrate a working Management and Orchestration (MANO) stack in just a few months. The power of the information-model-driven architecture enabled a community of diverse contributors (Telefónica, Canonical, Intel, RIFT.io, Metaswitch, and 6WIND) to deliver complete functionality for the following features:

- **End-to-end service fulfilment**, allowing full automation of the whole setup via standard open modeling.
- **Enhanced platform awareness (EPA) extensions**, which assure high and predictable performance in the virtual nodes where it is required.
- **SDN underlay control**, to guarantee link bandwidth wherever it's needed.
- **Multi-site capability**, with more than one site involved in the E2E deployment
- **Multi-cloud virtual infrastructure manager (VIM) capability**, embracing the fact that several VIMs are expected to coexist in industry for the near and medium term.
- **Deployment and configuration** of multi-tenant and single-tenant VNFs.
- **Physical network function (PNF) connections** to external physical lines with the existing physical elements.

This initial OSM demonstration code will provide the basis for the seed code framework adopted by the ETSI OSM project. The intention is to place the appropriate components of this stack into the open-source community at http://osm.etsi.org and continue the evolution of the project.

Introduction and Problem Statement

One of the key tenets of the initial NFV vision (see reference 1) was the vision of "the network becoming a data center," enabling a radically new approach to network service delivery by leveraging the commodity price points and capabilities of emerging data center technology. Toward enabling true NFV service delivery, it has become clear that many of these new carrier-grade virtual services will be constructed through a growing and varied ecosystem of VNF providers. The deployment and evolution of this virtual service model is best enabled through the emergence of an open-source management and orchestration stack built upon information and data models that are open and accessible to the general telco community, including independent software vendors, system integrators, equipment manufacturers, and operators, all coming together to unleash the true value of NFV.

The NFV and MANO problem statement was explored and is well understood by the ETSI ISG and the publicised NFV architecture (see reference 1). This initiative did a fine job of focusing on the additional components (to a conventional network) required to enable NFV service. There was a great effort in defining information models, which can in turn drive a lifecycle management function enabling resource orchestration through the NFVO and VIM. The challenge was initially so vast that some aspects were left outside of scope for further study. These included the higher-level issue of end-to-end service orchestration and also the current state of play in the open-source community with respect to cultivating possible management and end to end service orchestration solutions.

NFVO comprises both resource and service orchestration (see Figure 1), which combine to drive the end-to-end service deployment of a specific network service. This includes such tasks as initial Day 0 VNF deployment through the resource-orchestration (RO) layer, assuring the subsequent Day 1 coherent E2E service connectivity and the application of any subsequent VNF Day 2+ configuration required to bring the service into final operation.

Service orchestration is one of the areas where the complexities of deploying a telco network service particularly arise, and it presents many of the additional complexities that typically do not exist in the context of data centers and clouds, as it is the point in the architecture where the legacy world of operator OSS/BSS must adapt and meet the new world of NFV MANO. The OSM solution stack must also take into account the additional work and interaction required at the VIM layer to enable the EPA capabilities required to achieve acceptable carrier-grade VNF service performance.

The telecommunications provider deploys and guarantees the end-to-end service. Since the emergence of NFV in 2011, huge industry investment has enabled high-volume, standards-based servers (see reference 5) to deal effectively with I/O-intensive workloads as required in today's telco edge environment. Thus, most recent Intel® processors working together with suitably optimized hypervisors and the availability of open-source software libraries such as the Intel® Data Plane Development Kit (Intel® DPDK) have enabled vendors to build high-performance VNF solutions. Projects such as OPNFV (see reference 2) and OpenStack* (see reference 3) are enabling vendors and operators to understand, build, and deploy a generic and open NFVI layer onto which these NFV services can be deployed. But as already mentioned, the currently missing piece from the picture is the component required to deploy and manage the coherent end-to-end service. This allows the service providers to connect the world of software (VNFs) onto their open NFVI infrastructure. Open source creates the opportunity of enabling true end-to-end open NFV deployments, while also ensuring that true telco-grade SLAs are achieved.

In this emerging NFV market, the VNFs will be delivered by many different vendors. Unless the community enables and embraces a well-understood and open service instantiation implementation driven by a simple data model and extendable information models (evolved by the community), this new flexibility will become difficult to execute, and the associated system management and integration costs will rise due to the extra proprietary integration work that is required on a case-by-case basis for service implementations. Failure to provide this type of structured evolutionary approach would pave the way for the emergence of proprietary service models and MANO implementations, which would in turn hinder the emergence of an open VNF ecosystem, limiting service innovation while also driving up integration costs and the prevalence of walled-garden implementations.

The proposed approach described in this white paper embraces this challenge through the formation of a new open-source initiative (ETSI OSM) with the stated objective of driving an open-source NFV MANO implementation based on community-available information and data models.

OSM will connect the ecosystem of VNF vendors with an underlying NFVI, in turn enabling a more agile and competitive ecosystem of network functions providers supported through the published open information model approach.

At Mobile World Congress 2016, Telefónica, Canonical, Intel, RIFT.io, 6WIND, and Metaswitch have collaborated toward implementing and demonstrating such an OSM solution. The intent of this collaboration is to quickly prove the concept and architectural principles and to reuse as much of the implementation as possible to form the seed code as the foundational framework for the operator-led open-source MANO community project (ETSI OSM, see http://osm.etsi.org/welcome/).

OSM Proof of Concept and Use-Case Overview

The following overview describes the multi-site, multi-service use case used to demonstrate the key functions and benefits of the OSM project.

The service chosen (see Figure 2) combines the many layers of a typical telco deployment. The setup is composed of a core of PE routers, which provides the high-speed interconnection between different sites. On top of this core, an IMS VoIP service is deployed. Enabled by this virtual VPN infrastructure, the value-added enterprise IMS service is deployed and configured, providing the VoIP service to the users inside the connected organizations. The OSM orchestrator manages the diverse clouds (OpenStack and OpenVIM) and instantaneously deploys, connects, and configures the combined E2E NFV service with all of the primitives and attribute configurations required for each of the VNFs.

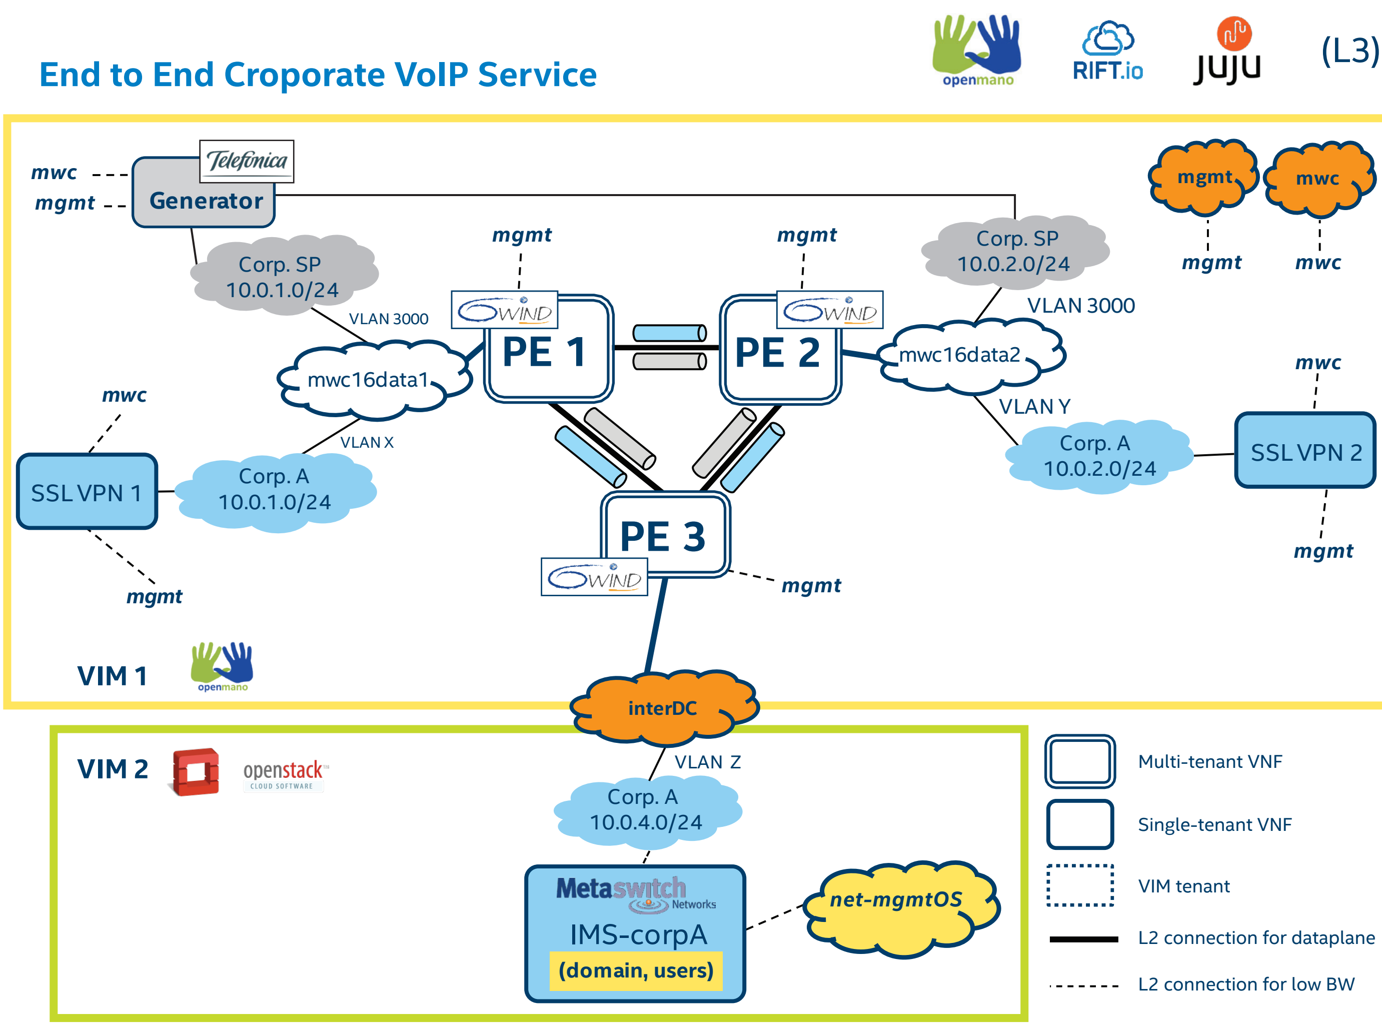

Figure 2. L3 VPN and IMS Service.

The objective of this OSM proof of concept is to show how it is possible for a multi-vendor community to re-use and extend concepts based on open-source components to build and evolve an open-source MANO stack capable of orchestrating services using commercial multi-vendor VNFs, which are entirely described through a well-known published data model.

NOTE: The use case selected is particularly challenging and realistic, as it covers most of the stringent requirements that a real production-like network scenario would require, such as performance, multi-site orchestration, and multiple vendors. It also involves configuration of both brown-field (physical network) and green-field (virtual network) parts of the infrastructure.

The demonstration addresses many of the more complex NFV service orchestration obstacles facing operators that wish to deploy NFV, including the following:

- Multi-use-case, end-to-end service orchestration
- Full support of enhanced platform awareness (EPA) to assure predictable high performance for advanced data-plane workloads required to enable telco-grade services
- Automated configuration and connectivity of the transport connectivity underlay
- Service control of multi (PE) and single-tenant (corporate IMS) VNFs in the same E2E scenario
- Operation and control of a multi-cloud/VIM environment
- Deployment across multiple sites
- Management of overlapping IP addresses, corresponding to different corporations
- Configuration of connectivity to external physical networks
- Combination of VNFs from different vendors in the same end-to-end service scenario

The work is hosted at Telefónica's NFV reference lab in Madrid, and it provides an NFVI environment at each of the sites. Each of them includes two NFVI clouds managed by two different VIMs. The IMS functions are deployed onto OpenStack, and OpenVIM (see reference 9) implements the complete set of EPA features required to host and correctly deploy the high performance PE (6WIND) routing VNF.

The OSM stack ingests the service descriptor, deploys the high-speed Intel DPDK-enabled 6WIND edge routers, and configures the Layer-2 underlay. It then automatically configures the VRF service routes between the sites, creating the multi-site VPN connectivity. The Traffic Gen/Sink VNFs are deployed as traffic generators to test and verify the initial inter-site connectivity and also to stress the PE routers (performance) during the trial to show how 10G inter-site line rate is achieved and service is maintained.

After this central core is deployed, and once the VPNs are provisioned, value-added service VNFs are deployed into the enterprise sites. In this case, Metaswitch' s Project Clearwater Open Source vIMS Core (see reference 4) is deployed for private IMS VoIP service to the corresponding enterprise. At this stage, the OSM implementation performs the IMS configuration and provisioning operations (setting the home domain, adding users, etc.), enabling VoIP SIP calls between any of the sites across the three corporations.

OSM Architecture Guiding Principles and Partner Contributions

OSM Guiding Architectural Principles

The OSM architecture is founded on the following guiding principles; all have the basis in the Open Information Model concept which is the primary founding principle of the project and its proposed evolution:

- **Layering**. Clear delineation is required between the layers and modules and the information model components each layer consumes/operates on. The layers must be broadly aligned with the ETSI-NFV information models and its recommendations. This allows the plug-in replacement of the various layers (SO, RO, VIM, etc.) with additional or alternative components that may be more applicable to the operator. The key is that the top-level information model must remain constant and coherent.
- **Abstraction**. Moving up or down the layers should offer clear differentiation in the levels of abstraction and detail presented in the information model and its usage within the different OSM layers.
- **Modularity**. Even within layers, clear modularity enabled with a plug-in model is preferred, to facilitate module replacements by later joining members as the OSM community develops and evolves.
- **Simplicity**. The solution must have the minimal complexity necessary to implement the governing information models and no more. The OSM stack will be evolved in a coherent manner by continuous evolution of the top-level information model, which drives the appropriate control (features and scope) into the architecture.

OSM Architectural Component Descriptions

The architecture (see Figure 3) is now described in the context of layers that are a part of the OSM stack necessary to deliver an end-to-end network scenario.

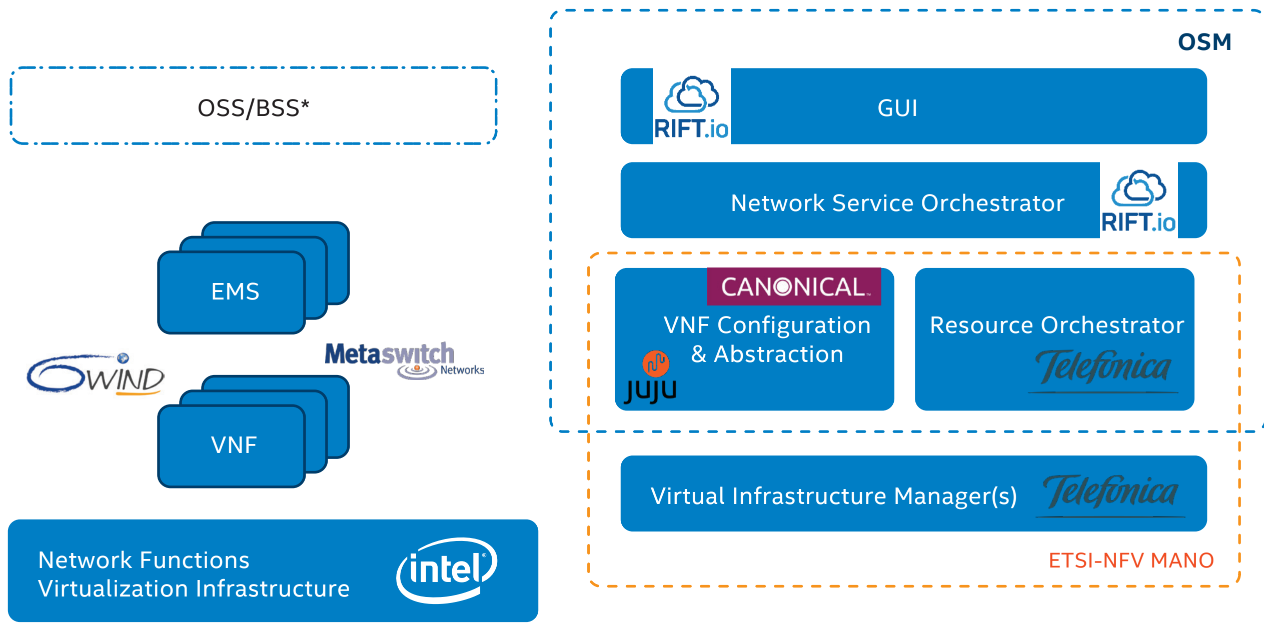

Figure 3. OSM architecture (simplified view).

OSM Layers

Resource Orchestrator

The resource orchestrator was provided by Telefónica (OpenMANO, see reference 10). The resource-orchestrator layer is responsible for processing the resource-allocation requirements of the VNF as per parts of the VNF descriptor (VNFD) and driving the VIM to allocate appropriate compute, network, and storage resources for the deployment of VNFs with their interconnection.

The resource orchestrator is responsible for coordinating resource allocation across multiple VIM types and multiple sites. It requires an understanding of EPA capabilities (see reference 5) at the VIM layer and is responsible for ensuring the VIMs are driven in a way that facilitates high performance VNF deployments and more efficient infrastructure utilization. This component is broadly aligned with the NFVO component in the ETSI-NFV MANO specification (see reference 11).

VNF Configuration and Abstraction

The VNF configuration component (Juju) was provided by Canonical. The VNF configurator and abstraction layer functionality provided by Juju offers generic VNF manager functionality and is responsible for processing configuration attributes of the VNFD and handling Day-0 and Day-1 level network service descriptor requests sent from the Network Service Orchestrator.

This is aligned with the VNF Manager component in the ETSI-NFV MANO specification in the deployment model, where NFV-O is responsible for resource orchestration. Juju provides open-source, generic application modeling, which enables OSM to focus on industry-specific orchestration challenges that uses charms to model services, relationships, and scale, while being independent of substrate.

Services are first-class concepts in the Juju model, making it particularly well suited to the role of generic VNF manager as described in Figure 1.

Network Service Orchestrator

RIFT.io provided the network service orchestration (NSO) Layer, which is responsible for end-to-end network-service delivery and drives the coherent service delivery through the RO and VNF configuration components.

NSO provides a simple interface for Network Service (NS) lifecycle management, NS-VNF-PNF on-boarding, translation between data models, catalogue-management for NS-VNF-PNF, configuration management for NS-VNF, and role-based access control.

Graphical User Interface (GUI)

RIFT.io provided the GUI and system-visibility layer. The GUI is the single-pane-of-glass view into configuring and reviewing the state of the orchestration environment. It provides a simple and intuitive way to interact with the NSO, tools for easy on-boarding of VNFs, and a drag-and-drop environment for creating NSs from the VNF catalog.

Instantiating NSs from the dashboard is as easy as browsing the catalog and specifying the required input parameters. Once the service has started, the GUI provides a model-driven view of the monitoring parameters exposed by the services' constituent VNFs, as well as the metrics for the resources the service is consuming. If the service exposes configuration-management endpoints, they can be invoked from the service details page.

The GUI also provides a detailed view of the underlying compute and network topologies that the orchestrator is managing.

Intel® Architecture-Based NFVI

The NFVI is an Intel architecture-based infrastructure. The ETSI-NFV NFVI layer is aligned with the ETSI-NFV002 architectural framework document (see reference 6) and includes all of the software and hardware components that contribute to the infrastructural resources for compute, networking, and storage.

For the Mobile World Congress OSM demo, the NFVI layer was built on an Intel architecture-based white-label server platform (see reference 13). The dual-processor platform is powered by two Intel® Xeon® processors E5-2699 v3 (18 cores) and uses dual-port network adapters based on the Intel® 82599 10 Gigabit Ethernet Controller.

The other NFVI components used were the OpenStack Kilo release, KVM, and Intel DPDK-enabled VNFs; all are integrated and released on the Intel® Open Networking Platform (Intel® ONP, see reference 12).

Virtual Infrastructure Managers

The OpenVIM (EPA Enhanced VIM) was provided by Telefónica. The VIM is the layer in the system that manages the compute, network, and storage resources of the infrastructural domain. The VIM is outside the scope of the OSM project but of course necessary for a demo and any real deployment. Two VIMs were used in the demonstration:

- **OpenVIM** from the OpenMANO project (see reference 10) was used to manage the infrastructure for the PE Router deployments. This included an OpenDaylight SDN (see reference 14) controller plugin to support automated data-plane switch configuration, and layer 2 underlay configuration.
- **OpenStack Kilo** was used to manage the infrastructure for the Metaswitch Clearwater IMS deployments with a Neutron Modular Layer 2 (ML2) plugin to support automated data-plane switch configuration.

NOTE: OSS/BSS layers have been shown in Figure 3 but are not a part of the OSM or MWC demo.

Virtualized Network Functions (Service VNFs)

Metaswitch vIMS Core

Released to the open-source community (see reference 4) in 2013, the Metaswitch vIMS core initiative, Project Clearwater, was built for NFV. With a micro-services architecture and innovations that enable asynchronous communications services to operate with carrier-grade scaling and reliability in the cloud, this complete IMS was deployed by the OSM stack in a just matter of minutes to deliver standards-based rich multimedia applications.

6WIND Turbo Router

The 6WIND Turbo Router is a high-performance, NFV-ready software appliance providing 10 Gbps per core of IP throughput. 6WIND Turbo Router quickly and easily integrates with OSM deployments. Acting as PE/edge routers, 6WIND Turbo Router demonstrated high-speed L3 forwarding in a multi-tenant environment supporting VRF. VRF allows independent L3 forwarding domains per customer network in a multi-tenant environment.

This project was initiated in late October as an output activity (interested vendors) from the OSM framework meeting held in Dusseldorf at the time of the October SDN-NFV conference there. Its stated intent is to evolve this new integrated code base and approach as the input seed code for the initial ETSI OSM initiative.

"Management and Orchestration (MANO) is, at the same time, one of the key components and most controversial concepts, in network virtualization architecture. Telefónica has long been working from the point of view of innovation in its development. A first result and seed of OSM is OpenMANO, a highly functional framework pioneering the first open-source NFV Orchestration and Management stack and, currently, a key component of Telefónica's NFV Reference Lab. By joining this community, we aim to accelerate the development of MANO while recognizing the value of opensource implementations of NFV and a need to harmonize efforts there."

– Antonio Elizondo, Head of Network Virtualisation Strategy and Technology,
Global CTO Unit, Telefónica

"Canonical's Juju provides open-source, generic application modeling, which enables OpenMANO to focus on industry-specific orchestration challenges, and is central to the delivery of the OSM project stack. The use of an open-source generic VNF manager such as Juju enables industry to collaborate and crowd-source expertise in performance, security and integration. We are committed to NFV and are delighted to participate in OSM."

– Mark Shuttleworth, Founder, Canonical

"Orchestration has a critical role determining how SDN-NFV workloads are provisioned, deployed, and managed to best take advantage of the platform capabilities of underlying standard high volume servers. Intel believes that Communication Service Provider-led ecosystems, such as Open Source MANO (OSM), working in conjunction with other vibrant SDN-NFV communities and Intel® Network Builders will play a critical shaping role to accelerate SDN-NFV deployments due to speed of innovation and the enablement of industry scale."

– John Healy, General Manager, Software Defined Networking Division, Intel

"RIFT.io's RIFT.ware automates end-to-end network service delivery, including network service lifecycle management, VNF and NSD on-boarding and interfaces with VNF configuration elements. RIFT.ware also provides a model-driven graphical user interface that provides a single view of the orchestration environment and multi-layer resource utilization. We are proud to join the OSM community and believe OSM provides a crucial element that will accelerate NFV deployments and expand the universe of deployable virtual network services. "

– Tony Schoener, CEO RIFT.IO

"Data-driven automated deployment and life-cycle management of complex VNFs is essential to realizing the full potential of NFV, but it's a complex and challenging problem to solve. In our view, the best way to approach this problem is to share the combined perspectives and expertise of a wide range of players in the NFV ecosystem, including network operators and vendors in the infrastructure, operations support and VNF spaces. This is what the Open Source MANO project is setting out to do, and as a result it's making a really valuable contribution to the state of the art in the MANO."

– Martin Taylor, CTO, Metaswitch Networks

"Successful Open Source MANO environments require VNFs with scalable performance that are NFV-ready for ease of integration. 6WIND Turbo Router was selected as the PE/edge router for the Open Source MANO demonstration for its quick and easy integration, high speed L3 forwarding, and key features such as VRF, which are required for multi-tenant environments."

– Eric Carmès, Founder and CEO, 6WIND

Recommendations to the OSM Community

During the four-month integration activity, many technical issues were resolved and architectural decisions made. It is currently envisioned that the ETSI OSM project will be officially initiated at a kick-off meeting in early April. At that meeting, the participating vendors expect to elaborate in much greater detail on each of the key output topics described in the remainder of this section.

Common and Published Information Models – How and When

A key principle necessary for both the governance and architectural evolution of the OSM project is that its evolution will be driven by an open information and data-model implementation. As in the initial technical sessions, the participants will make proposals about how the current, ESTI-based model can be adopted and evolved by the community using the current integration work as the starting point.

Adoption of MWC OSM Stack

The participants will make recommendations on the suitability of the demonstrated components for adoption by the new OSM community, in terms of architectural relevance and open-source license compatibility. The expectation is that many of these components will be adopted and evolved in the OSM community project and create the basis and seed code for the initial OSM release.

Key Technical Findings

There were some features which, due to timeframe, could not be addressed, and additional workarounds were implemented. Such issues and approaches will be exposed and explained at the first relevant OSM technical meeting. There were also gaps in external technology (e.g., EPA awareness in OpenStack) that drove certain decisions. These issues and proposals to address them will also be exposed to the community participants.

Finally, the best way to stay involved in OSM is to participate. There are different levels of OSM participation roles, including end-user advisors (operators), code contributors, solution testers, and general advisors in the discussions through the new OSM portal (http://osm.etsi.org/index.php).

Additional relevant organizations are encouraged that take the time to understand the project objectives and play an active part in its evolution, ideally through official project membership.

Conclusion

End-to-end service orchestration is the key to realizing carrier NFV, and it will deliver broad multi-use-case benefits to the operator community. The ability to integrate multi-vendor VNFs that enable end-to-end NFV solutions will unleash the true value of NFV service automation to the operator community. This will drastically reduce network operating expense by significantly lowering the cost of integration and helping a new competitive ecosystem of VNF providers to flourish.

To enable a true open VNF provider ecosystem, it is crucial that the industry embrace and evolve an open-source management and orchestration collaborative project that embraces an open information and open data-model approach as its key evolutionary driver.

The community must understand the performance and connectivity challenges unique to the operators and the associated complexities of creating end-to-end connectivity services. This requires exposing the appropriate levels of abstracted EPA intelligence up through the layered NFV delivery stack to maximize application performance and ensure that telco-grade SLA deployments are delivered.

These goals must be achieved while also reducing integration complexity and enabling scalable automation. The highest levels of openness and widespread community access to parties such as VNF vendors, system integrators, operators, and academia will be enabled through widespread adoption of the ETSI OSM Project.

References

1. https://portal.etsi.org/NFV/NFV_White_Paper.pdf
2. https://www.opnfv.org/
3. https://www.openstack.org/
4. http://www.projectclearwater.org/
5. ETSI GS NFV-PER 001 V1.1.2 - "Network Functions Virtualisation (NFV); NFV Performance & Portability Best Practices" https://www.etsi.org/deliver/etsi_gs/NFV-PER/001_099/001/01.01.01_60/gs_nfv-per001v010101p.pdf
6. http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf
7. https://software.intel.com/sites/default/files/managed/72/a6/OpenStack_EPA.pdf
8. http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/open-network-platform-server-paper.pdf
9. https://networkbuilders.intel.com/docs/Intel_Network_Builders_Directory_Sept2014.pdf
10. https://github.com/nfvlabs/openmano
11. http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf
12. http://www.intel.com/content/www/us/en/communications/open-network-platform-server.html
13. http://ark.intel.com/products/codename/75757/Wildcat-Pass
14. https://www.opendaylight.org/

Acronyms

BSS	Business Support System
CPU	Central Processing Unit
DPDK	Data Plane Development Kit
EPA	Enhanced Platform Awareness
ETSI	European Telecommunications and Standards Institute
IMS	IP Multimedia System
MANO	Management and Orchestration
NSD	Network Service Descriptor
NFV	Network Function Virtualization
NFV – O	Network Function Virtualization Orchestrator
OSS	Operations Support System
OPNFV	Open Network Virtualization Program
OSM	Open Source MANO
PE	Provider Edge Router
PNF	Physical Network Function
RO	Resource Orchestration
SDN	Software Defined Networking
SLA	Service Level Agreement
SO	Service Orchestration
TCO	Total Cost of Ownership
VIM	Virtual Infrastructure Manager
VNF	Virtual Network Function
VNFD	Virtual Network Function Descriptor
VRF	Virtual Routing Function

Solution provided by:

0215/JL/PDF 333958-001US