Demonstrating NFV is possible
1st Open Multi-Vendor NFV Showcase

enabled by ETSI Open Source MANO openstack intel whitestack

Gianpietro Lavado
glavado@whitestack.com
Whitestack

Jose Miguel Guzmán
jmguzman@whitestack.com
Whitestack

NFV

OPEN INFRASTRUCTURE SUMMIT Denver, CO | April 29–May 1, 2019

Agenda

- Current models of NFV Adoption
- Introducing Multi Vendor NFV Showcase
- Describe Tests and Results
- Conclusions & Invitation

Quick review about
Current models of NFV Adoption

The foundational paper

A white paper was written in 2012 by the world's leading telecom network operators (US, Europe & Asia).

- Introduction
- Benefits
- Enablers
- Challenges
- Call for Action

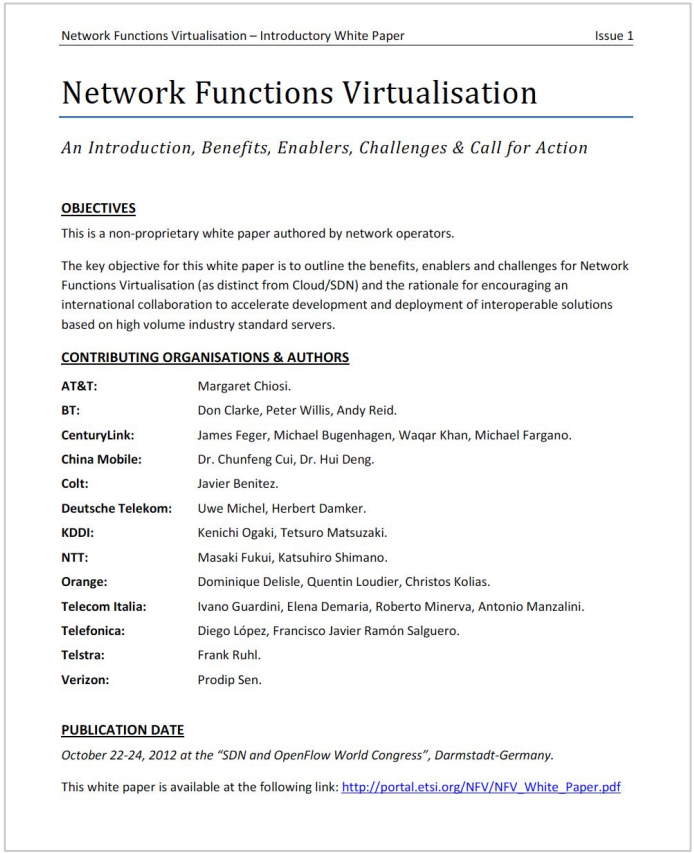

Network Functions Virtualisation – Introductory White Paper Issue 1

Network Functions Virtualisation

An Introduction, Benefits, Enablers, Challenges & Call for Action

OBJECTIVES

This is a non-proprietary white paper authored by network operators.

The key objective for this white paper is to outline the benefits, enablers and challenges for Network Functions Virtualisation (as distinct from Cloud/SDN) and the rationale for encouraging an international collaboration to accelerate development and deployment of interoperable solutions based on high volume industry standard servers.

CONTRIBUTING ORGANISATIONS & AUTHORS

AT&T:	Margaret Chiosi.
BT:	Don Clarke, Peter Willis, Andy Reid.
CenturyLink:	James Feger, Michael Bugenhagen, Waqar Khan, Michael Fargano.
China Mobile:	Dr. Chunfeng Cui, Dr. Hui Deng.
Colt:	Javier Benitez.
Deutsche Telekom:	Uwe Michel, Herbert Damker.
KDDI:	Kenichi Ogaki, Tetsuro Matsuzaki.
NTT:	Masaki Fukui, Katsuhiro Shimano.
Orange:	Dominique Delisle, Quentin Loudier, Christos Kolias.
Telecom Italia:	Ivano Guardini, Elena Demaria, Roberto Minerva, Antonio Manzalini.
Telefonica:	Diego López, Francisco Javier Ramón Salguero.
Telstra:	Frank Ruhl.
Verizon:	Prodip Sen.

PUBLICATION DATE

October 22-24, 2012 at the "SDN and OpenFlow World Congress", Darmstadt-Germany.

This white paper is available at the following link: http://portal.etsi.org/NFV/NFV_White_Paper.pdf

https://portal.etsi.org/nfv/nfv_white_paper.pdf

ETSI NFV Recommendations

- Based on member's feedback, field experiences and proof of concepts, standard documents have evolved.
- 60+ publications exist today, including the following three main documents:

- NFV Architectural Framework
  http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf
- NFV Infrastructure Overview
  http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/001/01.01.01_60/gs_NFV-INF001v010101p.pdf
- NFV Management and Orchestration
  http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf

http://www.etsi.org/standards-search

The NFV framework

ETSI defined an architecture to allow interoperability between VNFs and the infrastructure blocks, in an orchestrated way.

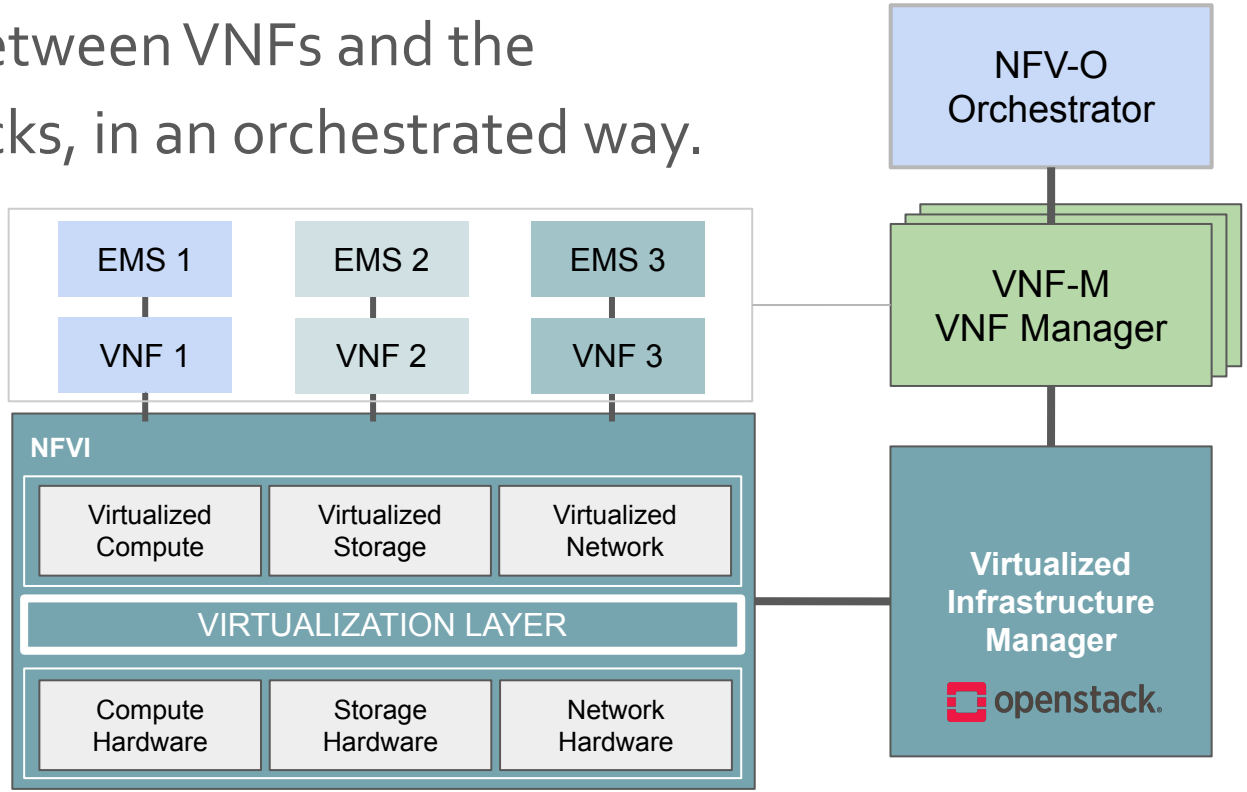

The expected scenario
“Horizontal Virtualization”

An architecture where

- Operators make the NVFI, VIM and MANO stacks available for multiple vendors
- VNF vendors provide “VNF Packages” that can be deployed on top of a compatible NFV infrastructure.

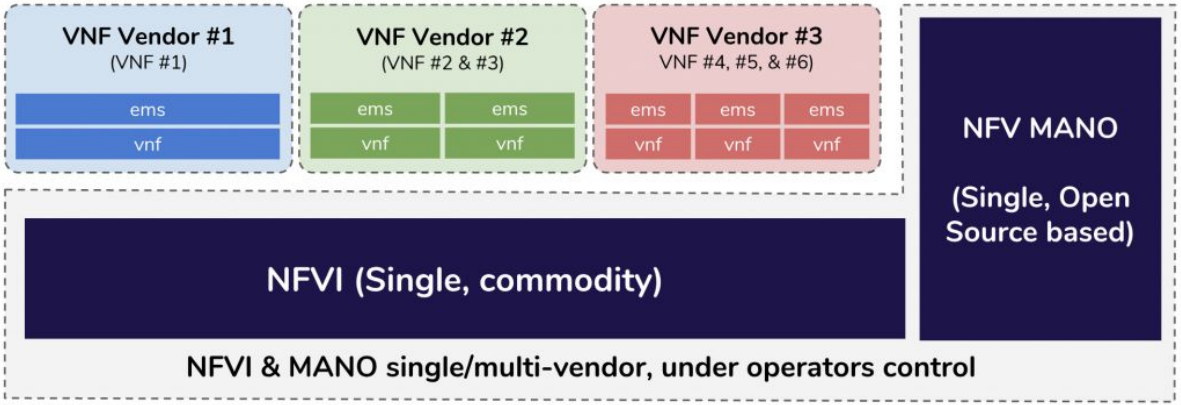

What is really happening
"Vertical Virtualization"

An suboptimal deployment, where

- Each vendor deploy an complete NFV stack and its VNFs
- There are different NFVIs, VNF-Ms and NFV-Os
- Vendors remain the ultimate responsible for the solution.

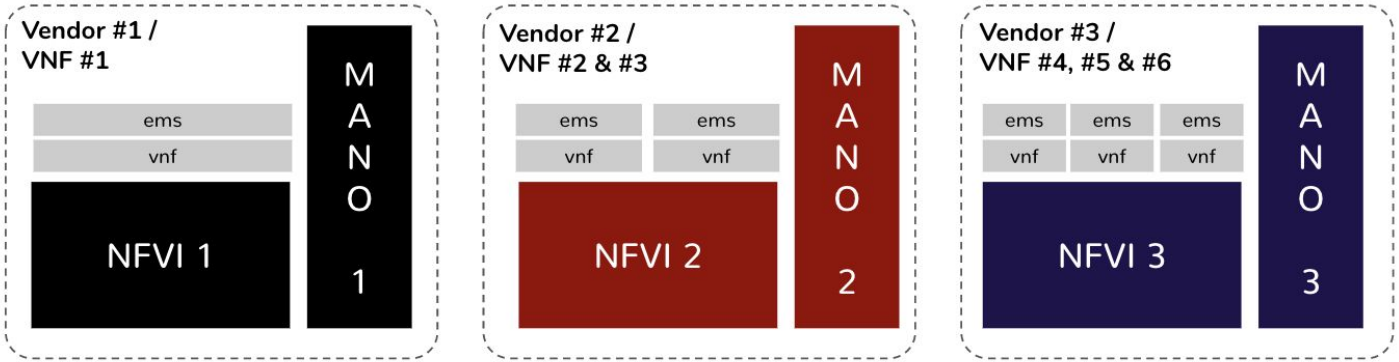

What is delaying Adoption

whitestack

Some operators are still concerned about:

- Running software on other's vendor hardware is too complex. ***Who do I blame?***
- Most required software components are not mature enough.
- If we save CAPEX by going NFV, I will end spending more OPEX.
- Virtualization on commodity hardware is not able to handle much traffic
- VNF Onboarding is very difficult (day-0, day-1 and day-2)
- Composing (multi-vendor) Network Services, is not possible (yet)

The
Multi Vendor NFV Showcase

Objective

whitestack

Demonstrate that **NFV Orchestrated Network Services**, integrating VNF from multiple vendors, on top of commoditized hardware, are possible (*with not too much effort*)

Architecture

"**Multi-vendor NFV Showcase**" with the support of leading NFV-enablers, putting together a number of leading VNF vendors, on top of commoditized **x86** infrastructure, managed by well-known, production-ready, open-source components like **OpenStack** and **Open Source MANO.**

Goal: to demonstrate publicly that multi-vendor networks are possible

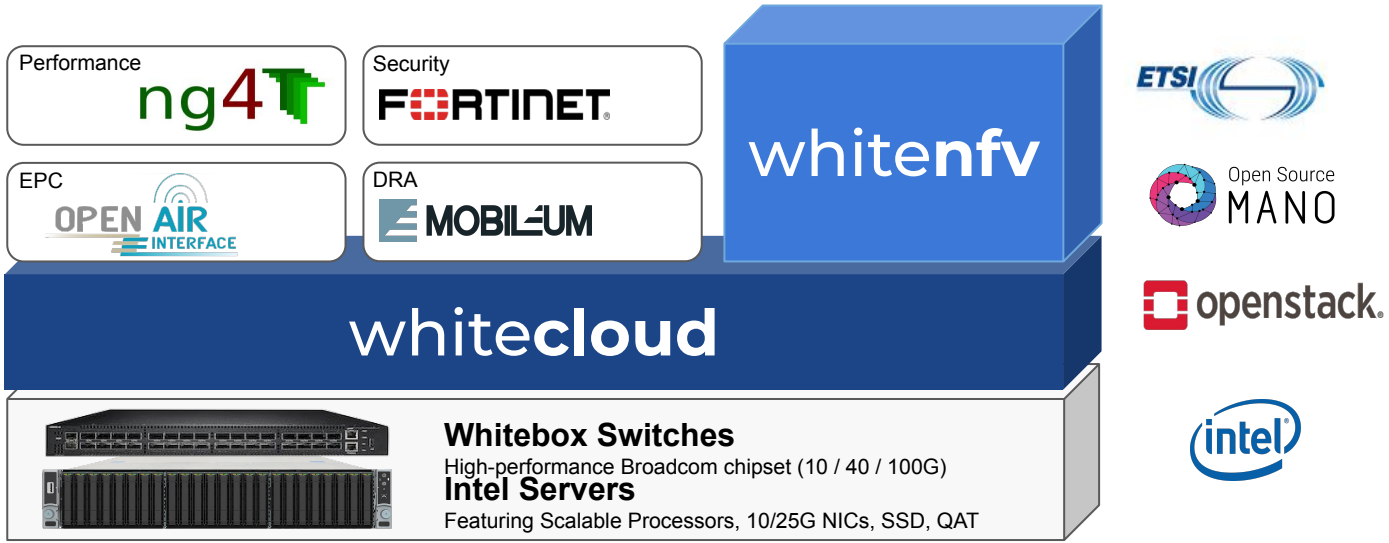

Network Service

- **ng4t VRAN**: Emulates the vRAN
- **OpenAir Interface**: Implement the vEPC (MME, SGW, PGW)
- **Fortinet**: implement security
- **Mobileum**: implement DRA and NTR (Roaming Steering)

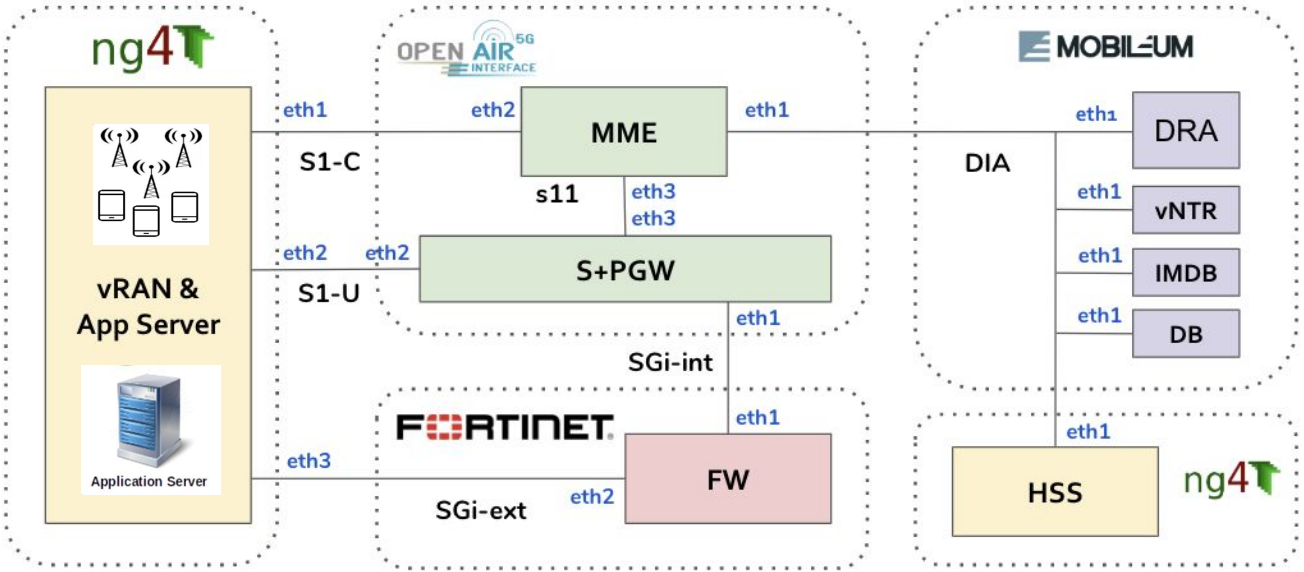

Enablers

- **ETSI**: Provide the **NFV Plugtests Programme** as a framework for the initiative, including access to its HIVE infrastructure (Hub for Interoperability and Validation at ETSI) for facilitating remote testing
- **OpenStack Foundation**: would provide general endorsement to the initiative, highlighting its support to using open technologies in NFV environments.
- **Intel**: Provides the the main hardware conforming the NFVI.
- **Whitestack**: provides the testbed to be used on this event
  - **NFVI**: Linux / Hypervisors
  - **VIM**: Whitecloud (Openstack distro)
  - **MANO**: WhiteNFV (Open Source Mano distro)

Tests and results

Test

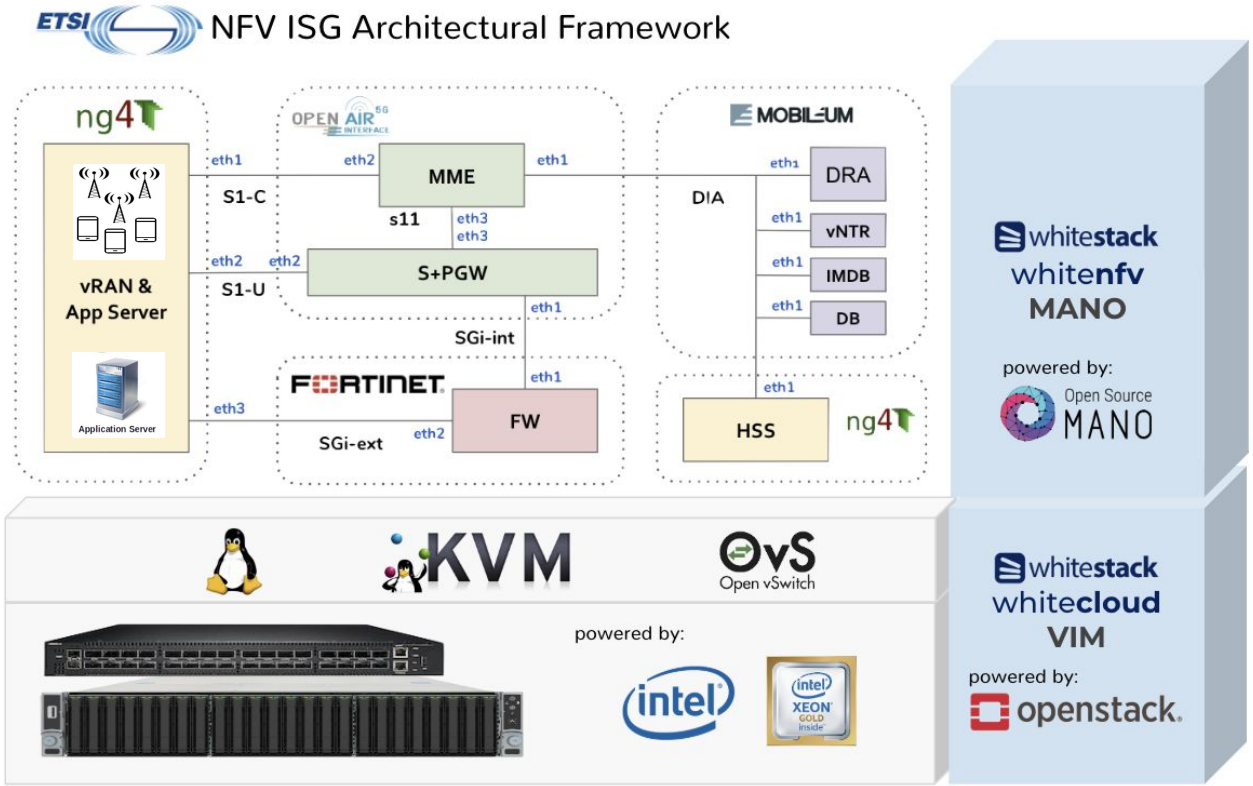

- This setup builds a completely **virtualized LTE Packet Core** with security features at the packet data network, and advanced roaming/analytics functions at the control plane.
- It is built using commodity hardware at NFVI, and open-source based MANO/VIM software.
- Multi-vendor, horizontal NFV is effectively achieved by leveraging ETSI NFV ISG standards.

Hardware for NFVI

Attribute	Description
CPU	Intel(R) Xeon(R) Gold 6139 CPU @ 2.30GHz
Memory	384GB RDIMM RAM
Disk	2 x 512GB SSD (OS) 4 x 1TB nvme (VNFs)

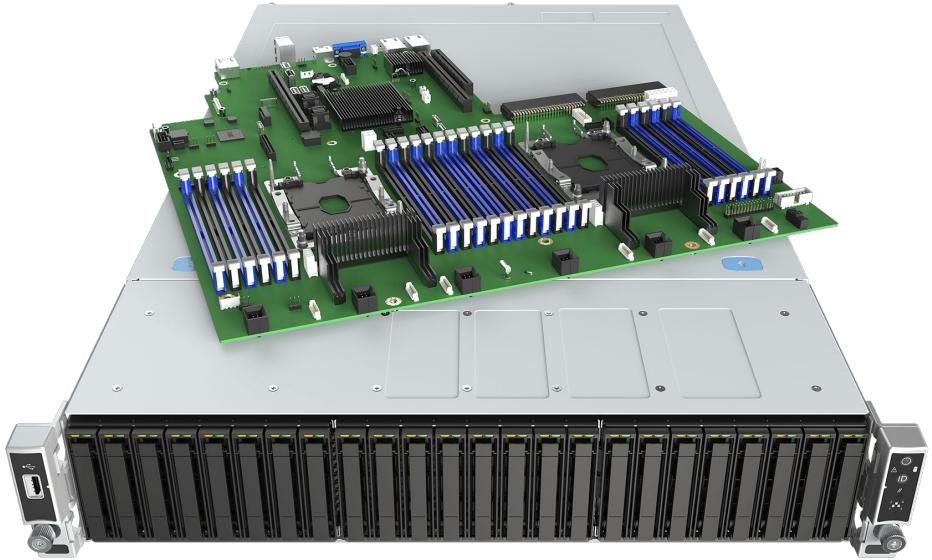

Software for NFVI & MANO

whitestack

Attribute	Description
Component	Open Source MANO
Release	5.0.5
Distribution	WhiteNFV (by Whitestack)
Version	alcobendas-1rc3

Performance

Security

Packet Core

DRA

whitenfv

whitecloud

Attribute	Description
Operating System	Ubuntu Server
Release	18.0.4 LTS
Kernel	4.15.0-48
KVM	4.15.0-48
QEMU	2.11

Attribute	Description
Component	Openstack
Release	Rocky
Distribution	WhiteCloud (by Whitestack)
Version	rio-1

Network Service

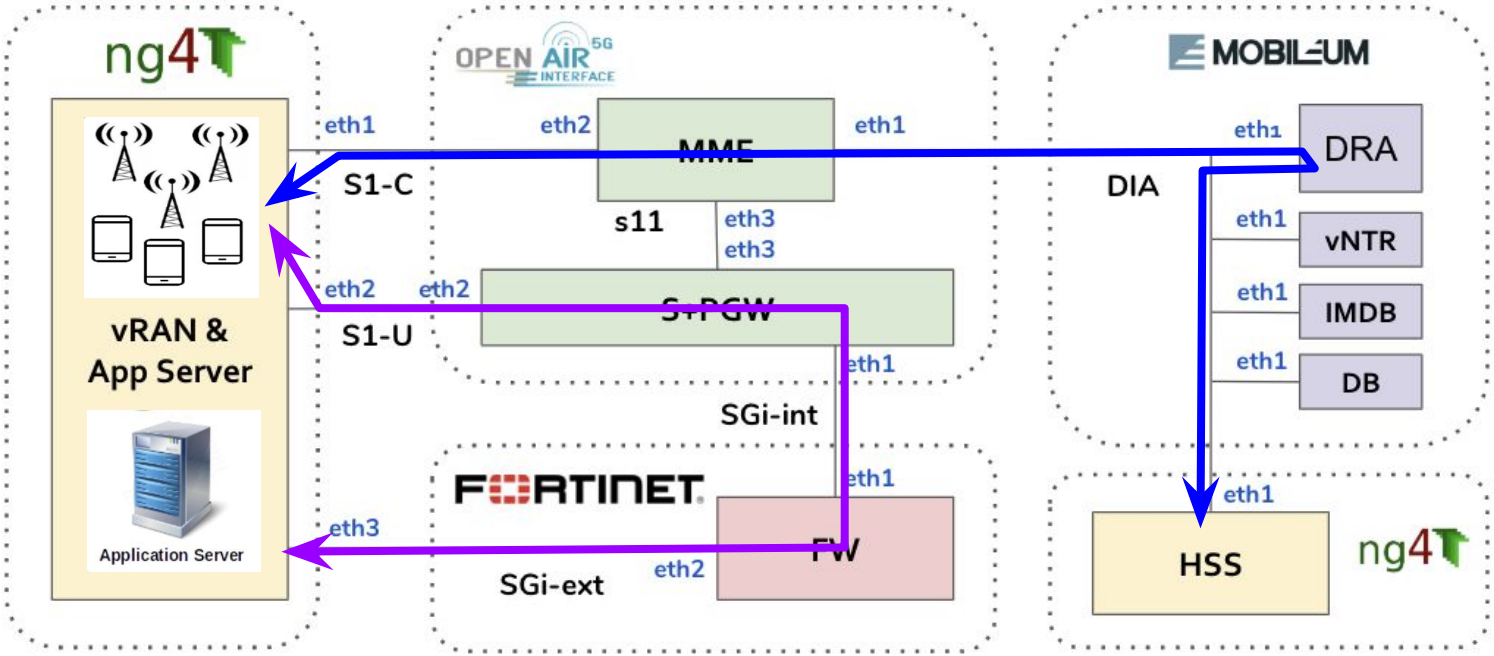

Mobile Data Service
Control Plane

Mobile Data Service
Data Plane

The VNF Onboarding Challenge

What's the real objective?

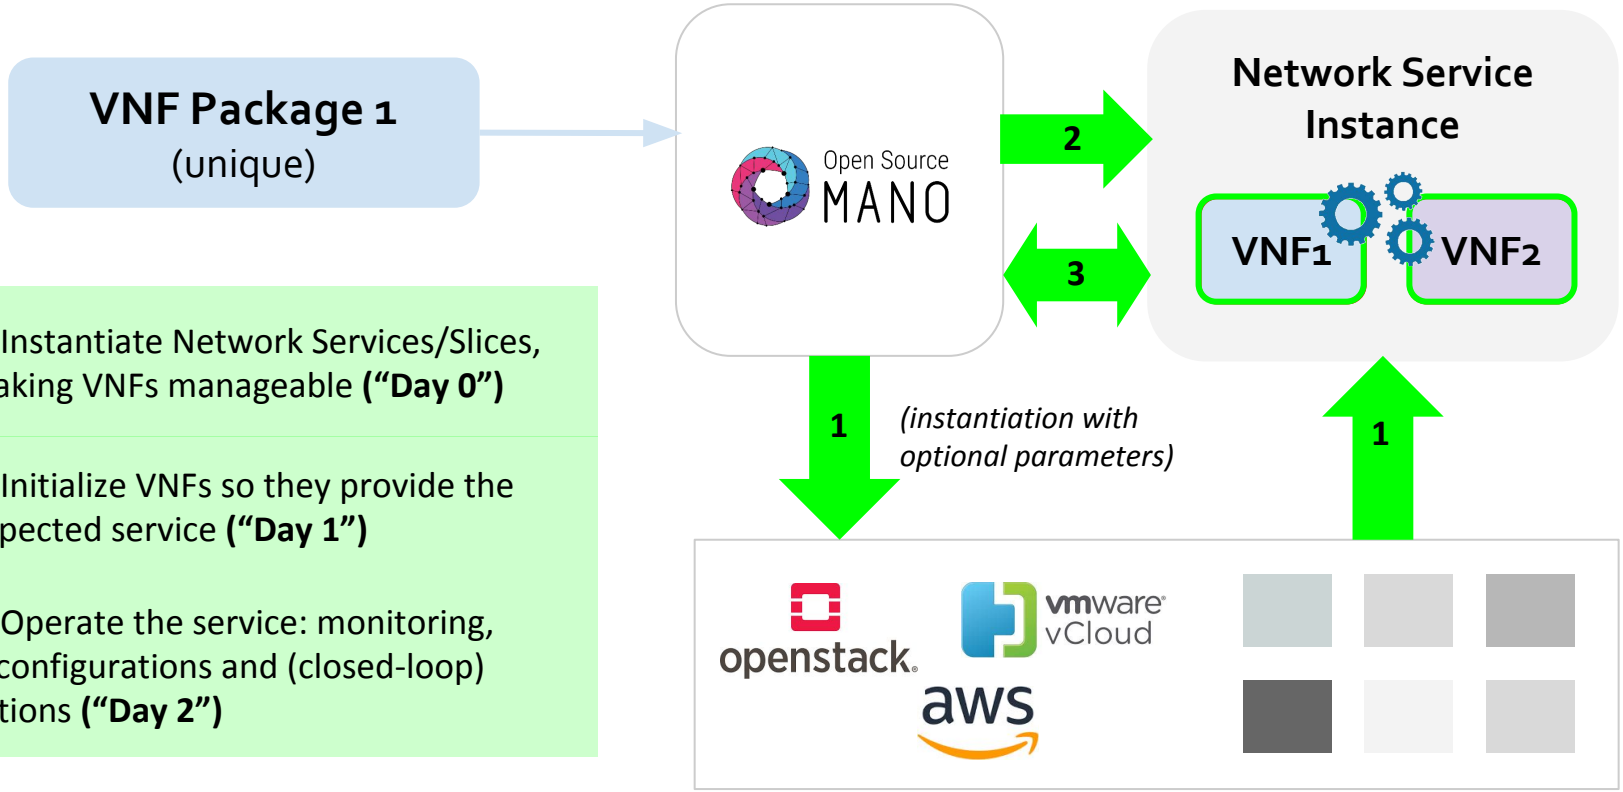

1. Instantiate Network Services/Slices, making VNFs manageable **("Day 0")**

2. Initialize VNFs so they provide the expected service **("Day 1")**

3. Operate the service: monitoring, reconfigurations and (closed-loop) actions **("Day 2")**

The VNF Onboarding Challenge

How do we achieve it?

whitestack

Open Source MANO

Day 0: Instantiate

cloud-init

Isolated VNFs

FORTINET

OPEN AIR INTERFACE

MOBILEUM

ng4T

Open Source MANO

Day 1: Build Service (automated)

juju

Evolved Packet Core

FORTINET

OPEN AIR INTERFACE

MOBILEUM

ng4T

Open Source MANO

Day 2: Operate (on demand)

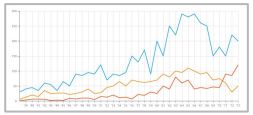

juju

Evolved Packet Core

FORTINET

OPEN AIR INTERFACE

MOBILEUM

ng4T

The VNF Onboarding Challenge

How do we achieve it?

The VNF Onboarding process concludes when we have a package that models the VNF Day-0 to Day-2 requirements

```
nsd:nsd-catalog:
  nsd:
  - constituent-vnfd:
    - member-vnf-index: 1
      vnfd-id-ref: vran-vnfd
    - member-vnf-index: 2
      vnfd-id-ref: c
    - member-vnf-inc
      vnfd-id-ref: h
    - member-vnf-inc
      vnfd-id-ref: m
    - member-vnf-inc
      vnfd-id-ref: m
    - member-vnf-inc
      vnfd-id-ref: m
    - member-vnf-inc
      vnfd-id-ref: m
    - member-vnf-inc
      vnfd-id-ref: f
    - member-vnf-inc
      vnfd-id-ref: c
    id: ng4t-oai-for
    ip-profiles:
    -   name: s1-c
        description:
        ip-profile-p
```

```
vnfd-catalog:
  vnfd:
  - connection-point:
    - name: eth0
      type: VPORT
    - name: eth1
      port-security-enabled: 'false'
      type: VPORT
    - name: eth2
      port-security-enabled: 'false'
      type: VPORT
    - name: eth3
      port-security-enabled: 'false'
      type: VPORT
    description: oai-psgw
    id: oai-psgw-vnfd
    mgmt-interface:
      cp: eth0
    name: oai-psgw-vnfd
    service-function-chain: UNAWARE
    short-name: oai-psgw-vnfd
    vdu:
    - cloud-init-file: psgw-cloud-init
      supplemental-boot-data:
        boot-data-drive: true
      count: '1'
```

```
- hosts: all
  connection: local
  vars:
   username: "admin"
   password: "password"
  tasks:
  - name: Configure interfaces.
    fortios_system_interface:
      host: "{{ fgt_mgmt_ip }}"
      username: "{{ username }}"
      password: "{{ password }}"
      https: "False"
```

```
@when('actions.add-routes')
def add_routes():
    # Adding LTE routes
    ue_network = action_get('ue_network')
    sgiint_fw_ip = action_get('sgiint_fw_ip')
    app_network = action_get('app_network')
    cmd = "sudo ip r add default via " + sgiint_fw_ip + " dev ens4 table lte && \
          sudo ip rule add from " + ue_network + " table lte && \
          sudo route add -net " + app_network + " gw " + sgiint_fw_ip
    charms.sshproxy._run(cmd)
    remove_flag('actions.add-routes')

@when('actions.start-service')
def start_service():
    err = ''
    # Start SPGW
    cmd = "/home/ubuntu/OAI/start_spgw"
    charms.sshproxy._run(cmd)
    remove_flag('actions.start-service')
```

The VNF Onboarding Challenge

What do we get?

1. We are able to instantiate the VNFs successfully.

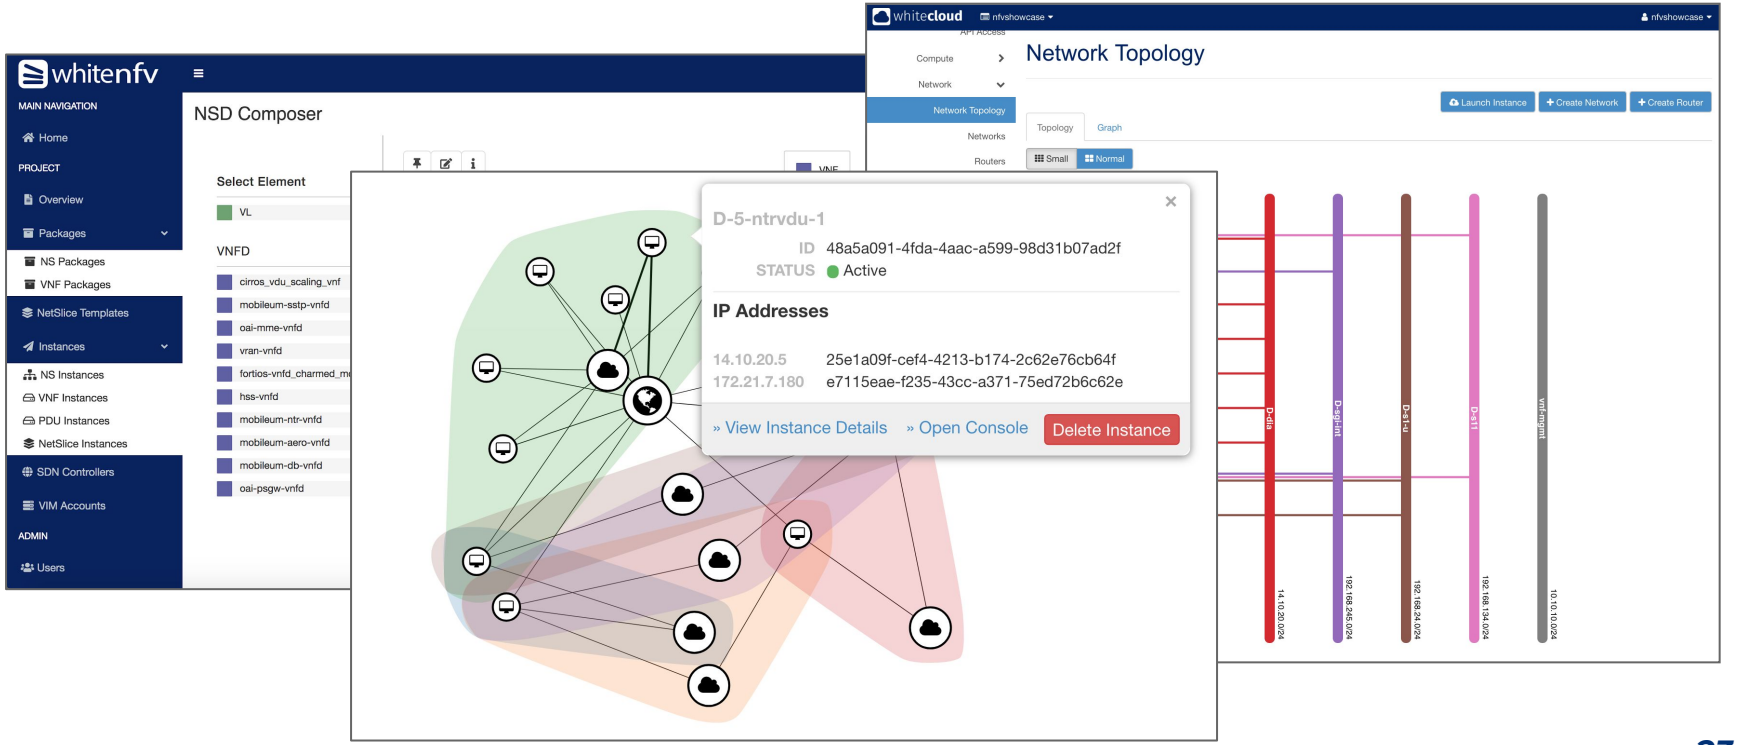

The VNF Onboarding Challenge

What do we get?

2. We get a <u>fully functional Network Service without manual intervention</u>

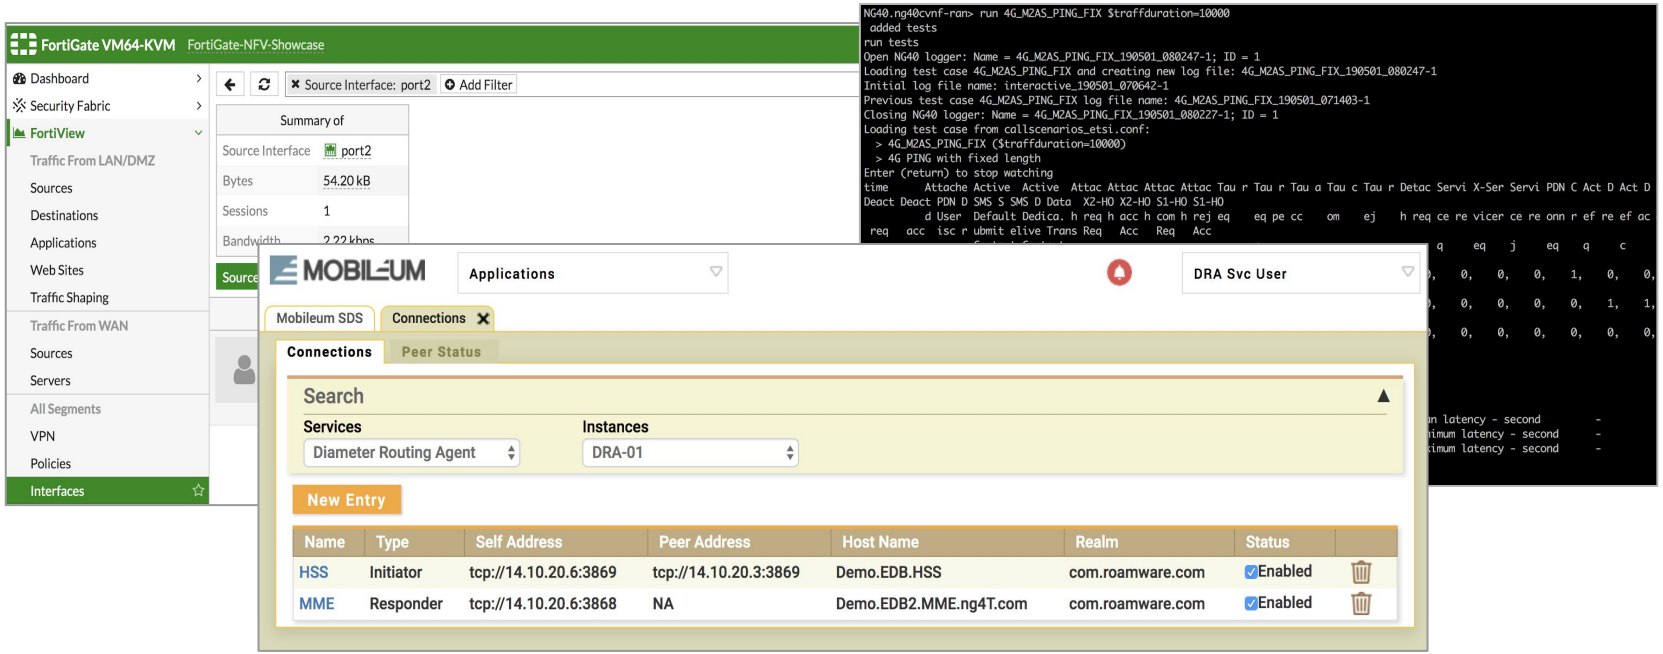

The VNF Onboarding Challenge

What do we get?

3. We are able to operate the service through simple interfaces through the Orchestrator

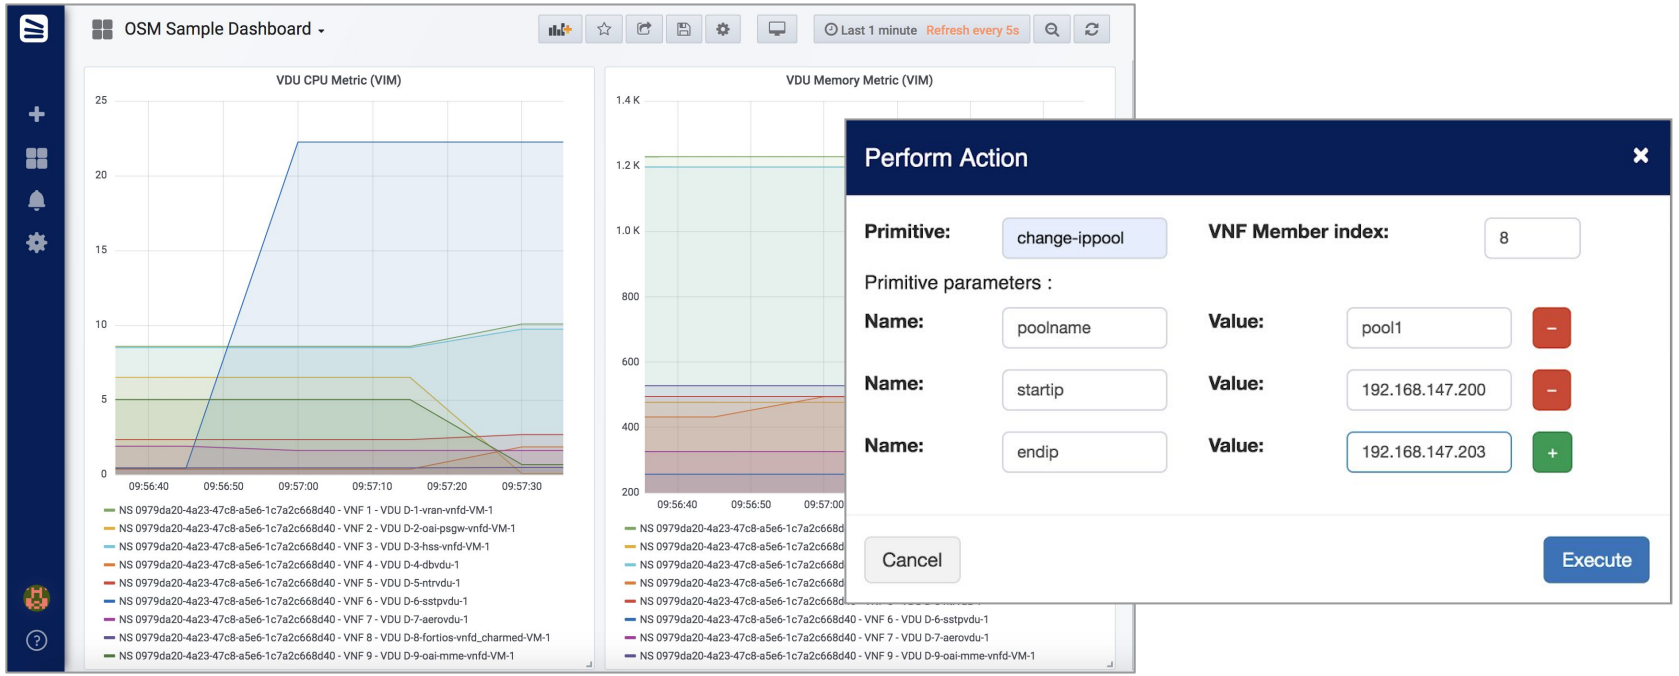

Conclusions & Invitation

Conclusions

In only 4 weeks:

1. VNFs from different vendors were **described** (VNFD) and **onboarded** in the Open Source MANO catalog.
2. **Day-0, Day-1 and Day-2 operations were automated** for each VNF, by using **Open Source MANO VCA** (VNF Configuration and Abstraction), on top of OpenStack, following ETSI NFV Standards.
3. A multi vendor **virtualized Evolved Packet Core was modeled**, by using the ETSI OSM Information Model, including multiple Virtual Links for implementing the service topology.
4. And we deployed it on OpenStack with OSM.

Conclusions

whitestack

Most of the challenges are not related to technology
but to the use of good practices

We are not anymore in the discussion if traditional software would work virtualized on x86,
but how to scale by using the cloud

Dynamic Configuration - Scalability Models - Manageability

(that are also limitations in the baremetal model)

Final Report

The final results, including configurations used for deploying this vEPC, are published, following the guidelines from ETSI Plugtests Programme.

https://www.whitestack.com/posts/results-multivendor-nfv-showcase/

Invitation

whitestack

The Cloud was the result of a **industry effort to reach more efficiencies**, aligning manufacturers, software developers, systems integrators and service providers.

The Telco industry is **struggling to scale**, and needs to move to the Cloud. Virtualizations is not enough!.

Invitation

whitestack

Vendors need to **get on board**, by

- Adopting new **tools and technologies** (cloud-centric)
- Using more recent **best practices**
  (dynamic config, horizontal scaling)
- Creating cloud-native solution,
  and demonstrate them in action.

We encourage vendors to follow that path, and participate in our

2nd Multi Vendor NFV Showcase

(Oct/2019)

whitestack

Thanks!

Download report from:

Gianpietro Lavado
glavado@whitestack.com
Whitestack

Jose Miguel Guzmán
jmguzman@whitestack.com
Whitestack