

Open Source
MANO

OSM 5th Hackfest – Introduction to NFV and OSM

José Miguel Guzmán (Whitestack)

- Background for moving to a NFV world
- Quick review of NFV
- Introduction to the latest OSM Release
- Contributing to the Community



Open Source
MANO

Background for moving to a NFV world



Credit: Brue? Da

How was this originated?

A white paper was written in 2012 by the world's leading telecom network operators.

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

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 Koliass.
Telecom Italia:	Ivano Guardini, Elena Demaria, Roberto Minerva, Antonio Manzalini.
Telefonica:	Diego López, Francisco Javier Ramón Salguero.
Telstra:	Frank Ruhl.
Verizon:	Prodip Sen.

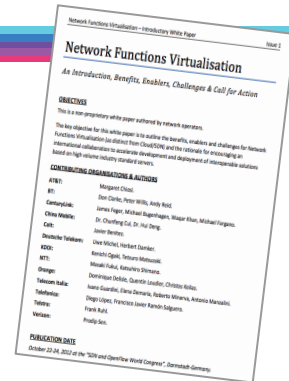
PUBLICATION DATE

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

How was this originated?

The issues

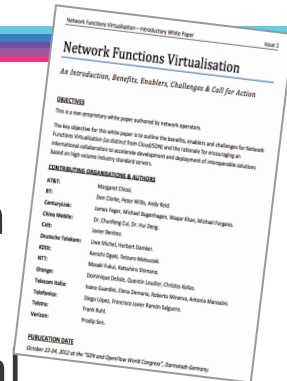
- Network Operators' networks are populated with a **large and increasing variety of proprietary** hardware appliances.
- To launch a new network service often requires **finding the space and power to accommodate these boxes**, what is becoming increasingly difficult;
- **Increasing costs** of energy, capital investment challenges and the **rarity of skills necessary** to design, integrate and operate increasingly complex hardware-based appliances.
- Moreover, hardware-based appliances rapidly reach end of life, **requiring much of the procure-design-integrate-deploy cycle to be repeated with little or no revenue benefit.**



How was this originated?

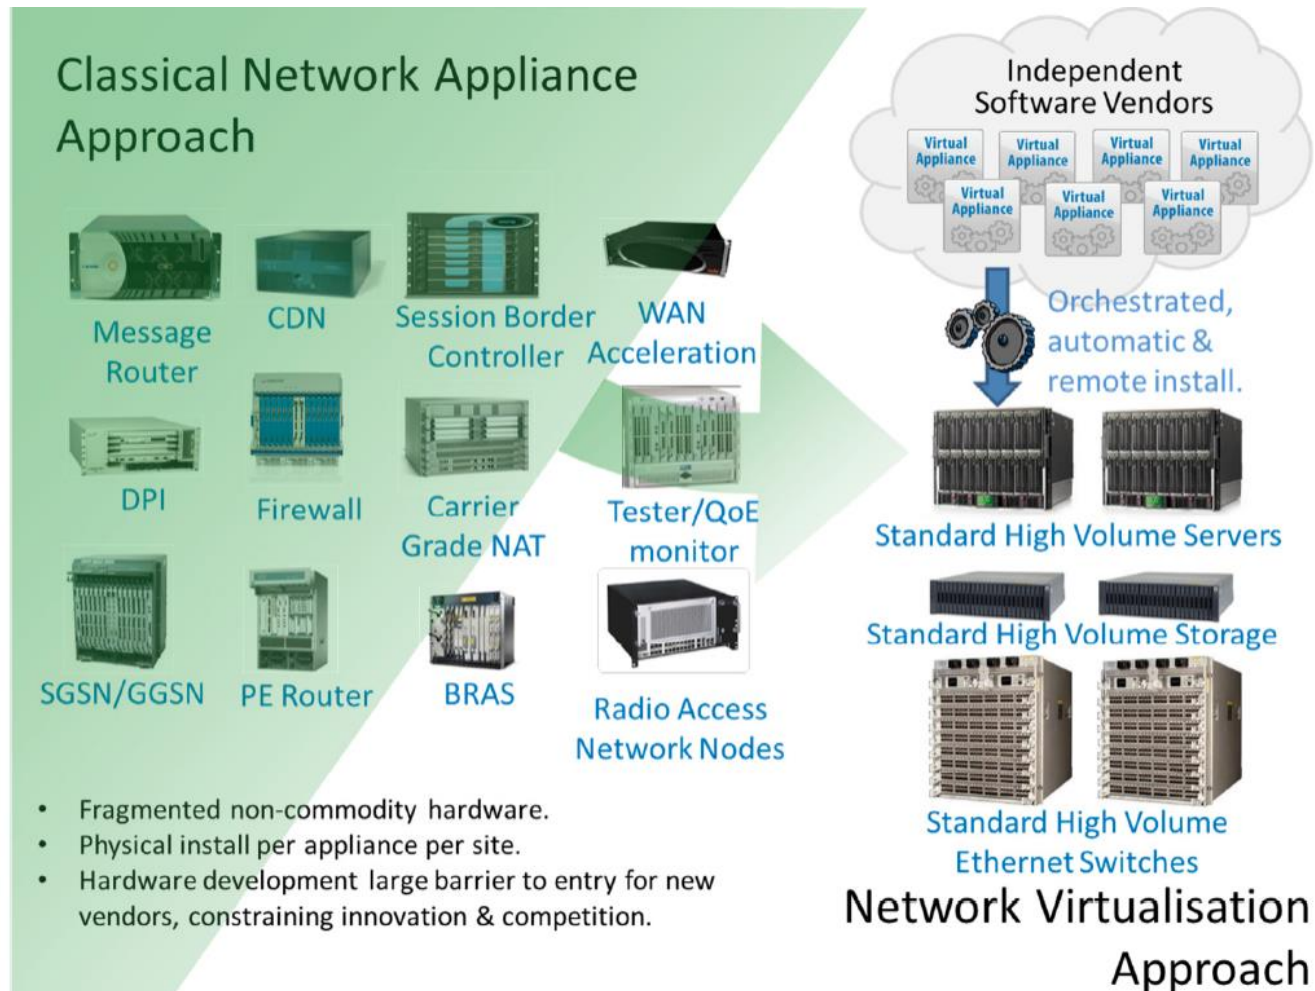
The Benefits

- **Reduced equipment costs** and reduced power consumption through consolidating equipment (scale of the IT industry)
- Increased speed of **Time to Market** by minimising the typical network operator cycle of innovation.
- Availability of network appliance multi-version and multi-tenancy, which allows use of a **single platform for different applications, users and tenants**.
- **Targeted service introduction** based on geography or customer sets is possible. **Services can be rapidly scaled up/down** as required.
- Enables a wide **variety of eco-systems** and encourages **openness**.

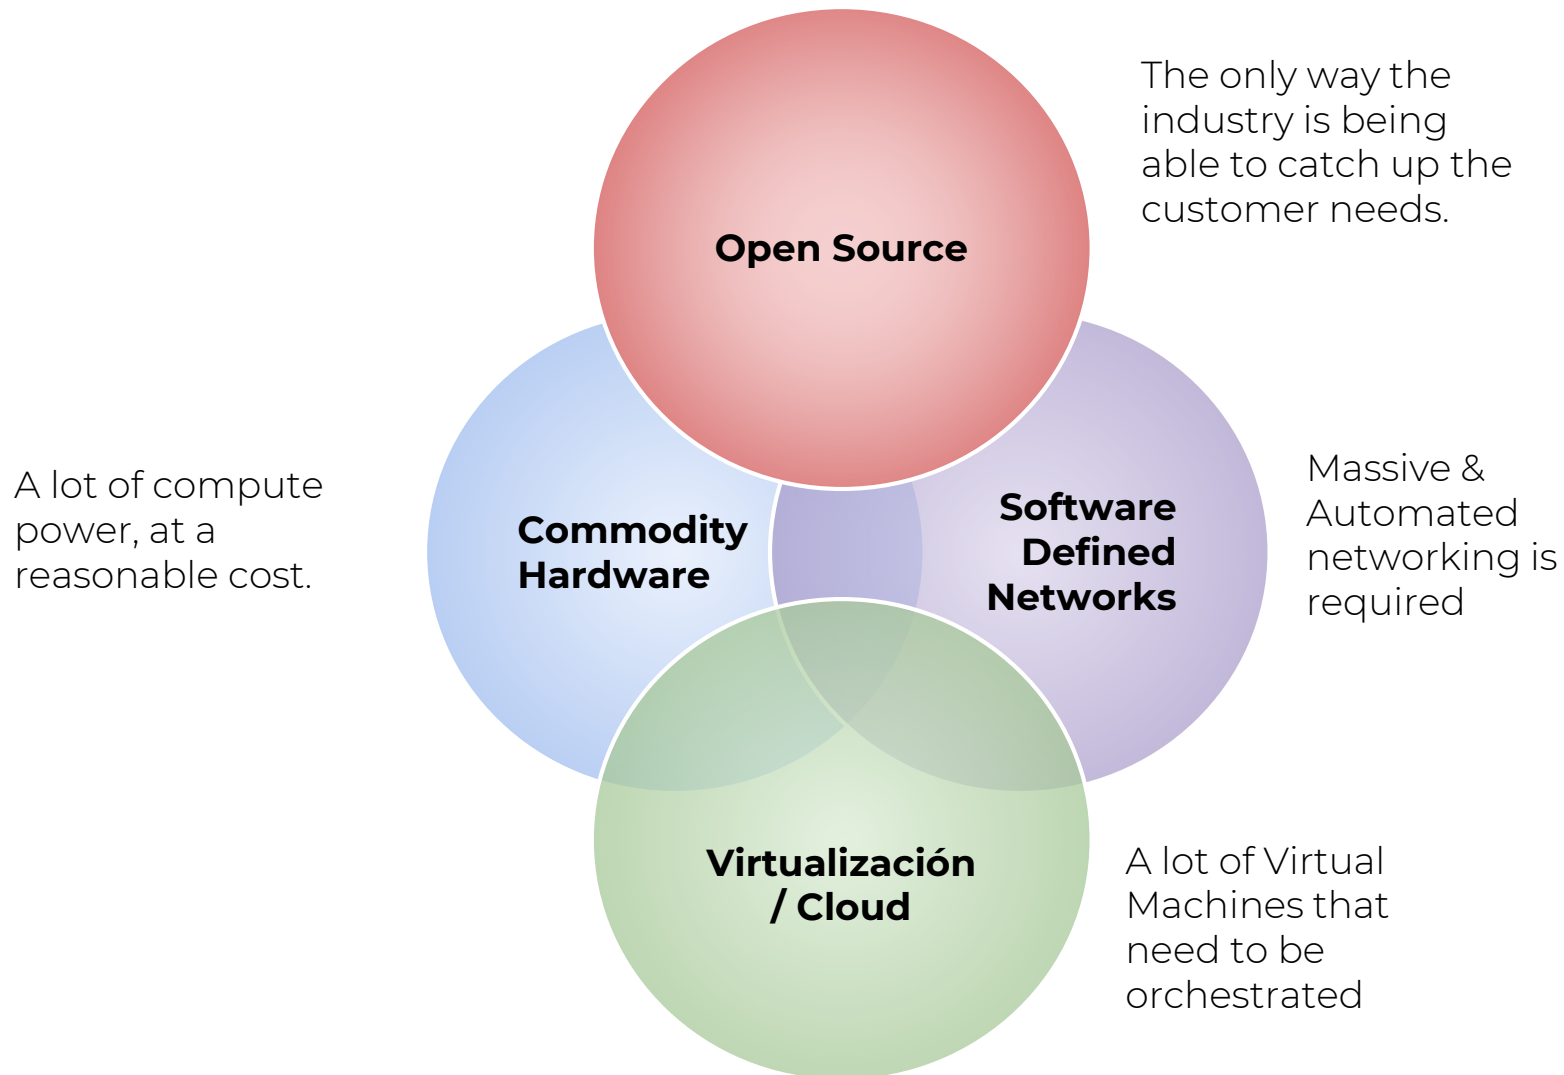


What is NFV trying to address?

Bringing “cloud efficiencies” to the Telecom Industry



Conjunction of 4 mature technologies



1) Commodity Hardware

- The Telecom industry has transitioned from proprietary (tailor-made) hardware platforms, to commodity (COTS) x86 hardware



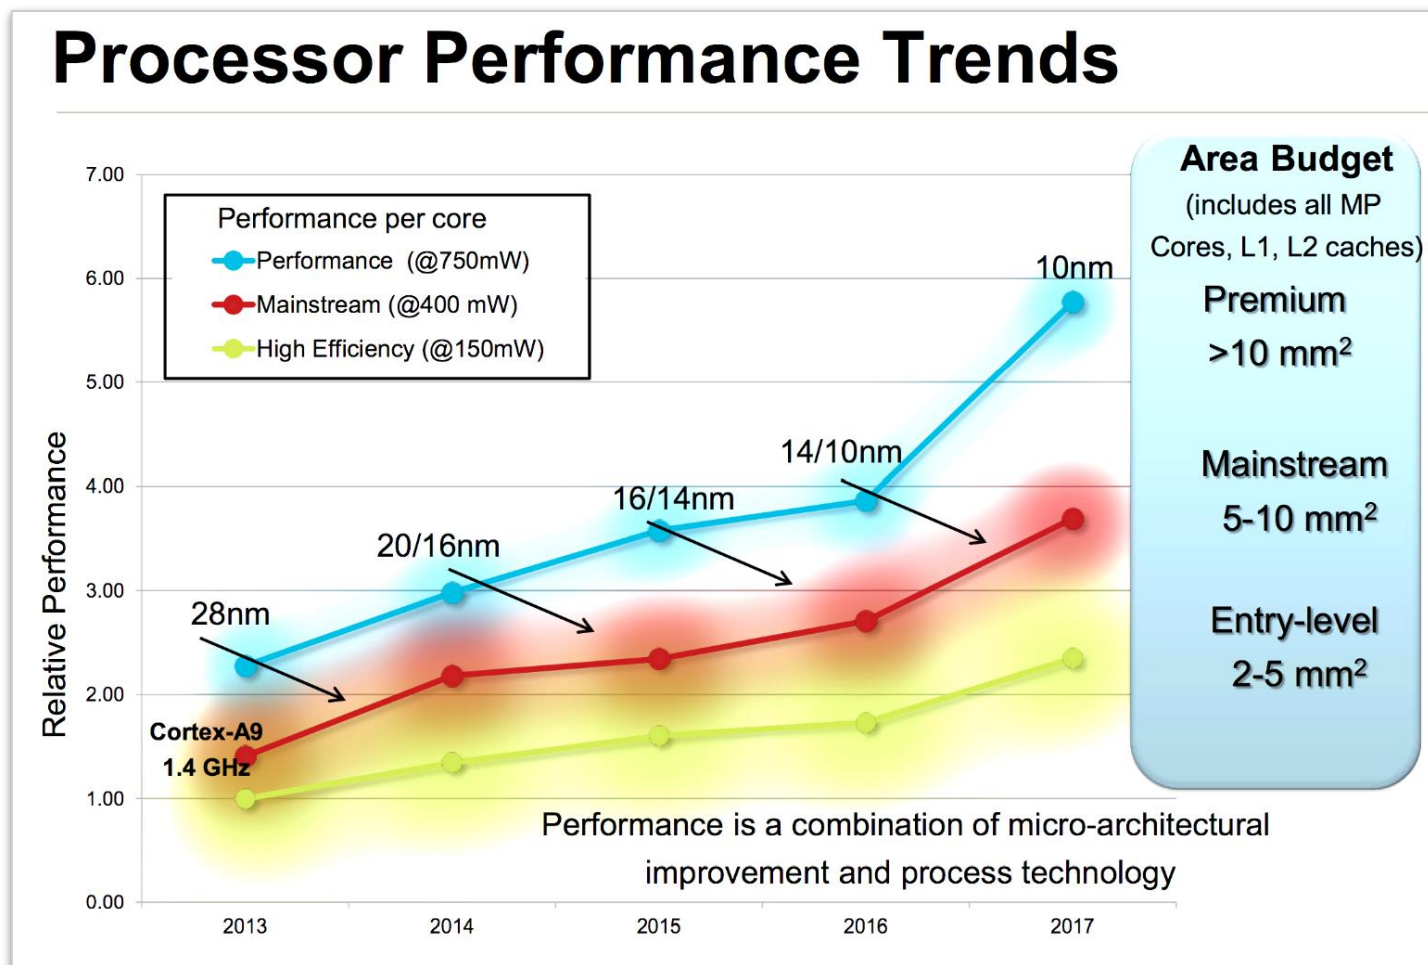
Arquitectura Hardware
Advanced TCA



Arquitectura Hardware
Commodity “off-the-shelf” con
procesador x86

1) Commodity Hardware

- In recent years, the performance of Microprocessors has more than duplicated, specially in the servers segment



1) Commodity Hardware

- New form factors, might achieve the highest compute & storage density levels



Core



96x



= 6 TB of RAM

24x



= 96 TB of SSD

4x



= 16 x 10G Ports

2 RU



1) Commodity Hardware

- A super computer.. in a rack


5K Physical Cores

60 TB of RAM

3840 TB of Solid State Storage


320 x 10G Ports

/Rack
@ ~ 20K Watts



1) Commodity Hardware

- Non-Volatile RAM, improves the performance of traditional SSD storage



1) Commodity Hardware

- New SSD form-factors, improve the SSD density

Built in Serviceability

Programmable LEDs to quickly locate failed drives, offline drives, and unpopulated slots



Carrier-less design with integrated pull tab removes need for drive carriers

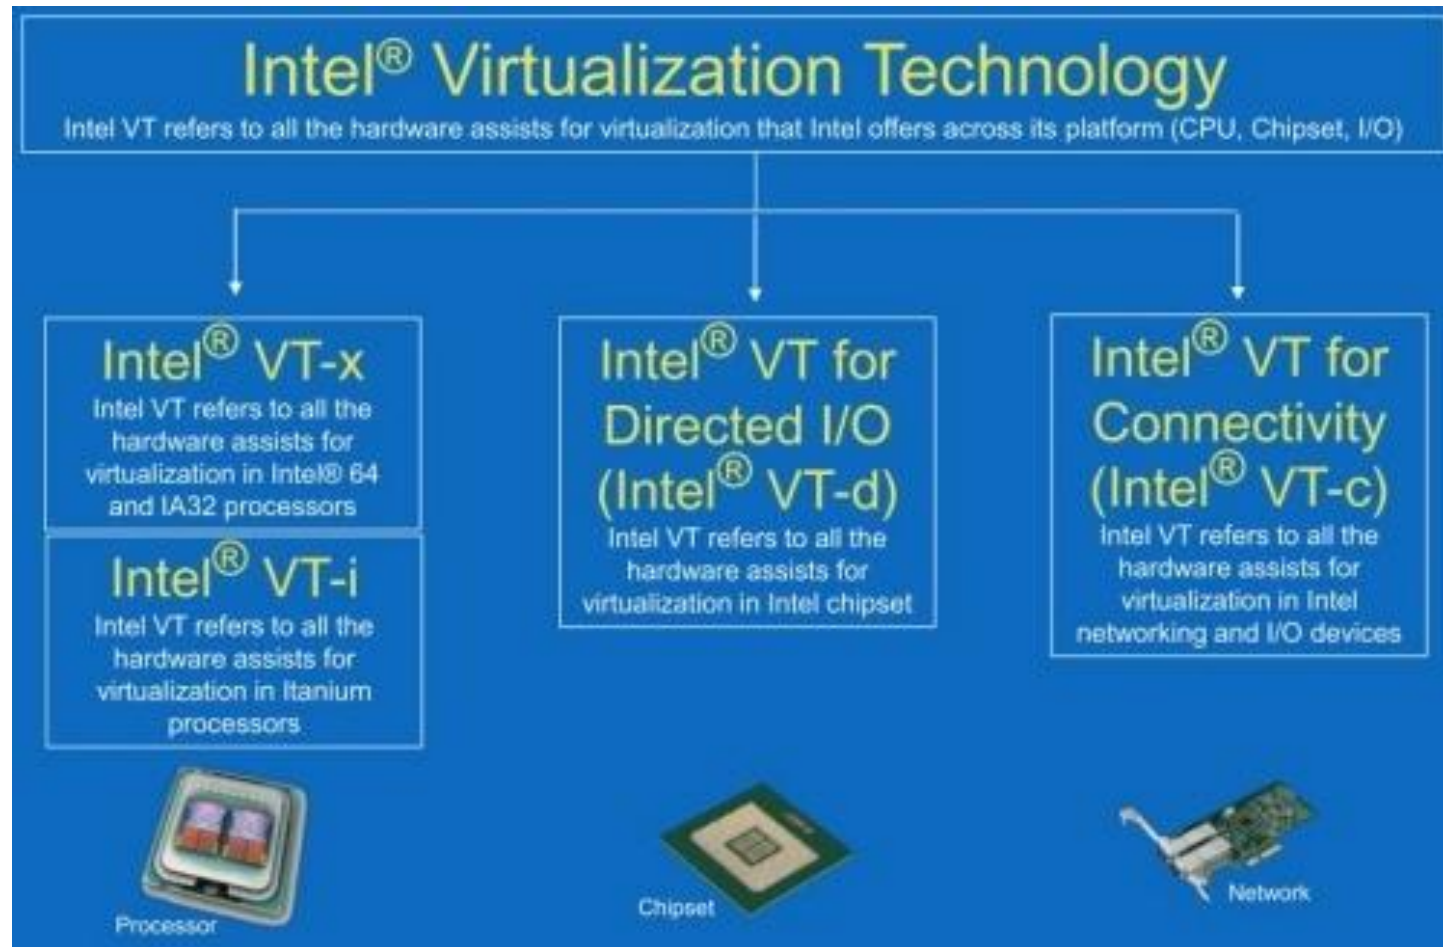
Enclosure Management with slot level power control enables single drive isolation or system level power loss



STORAGE CAPACITY
1PB IN 1U
INTEL® 3D NAND SSD, 32TB RULER IN 2018

2) HW Assisted Virtualization

- Intel/AMD have developed a number of hardware features, to accelerate virtualization



2) HW Assisted Virtualization

- Two leading solutions for virtualization

Commercial



Microsoft
Hyper-V



VirtualBox

CITRIX
XenServer

Open Source

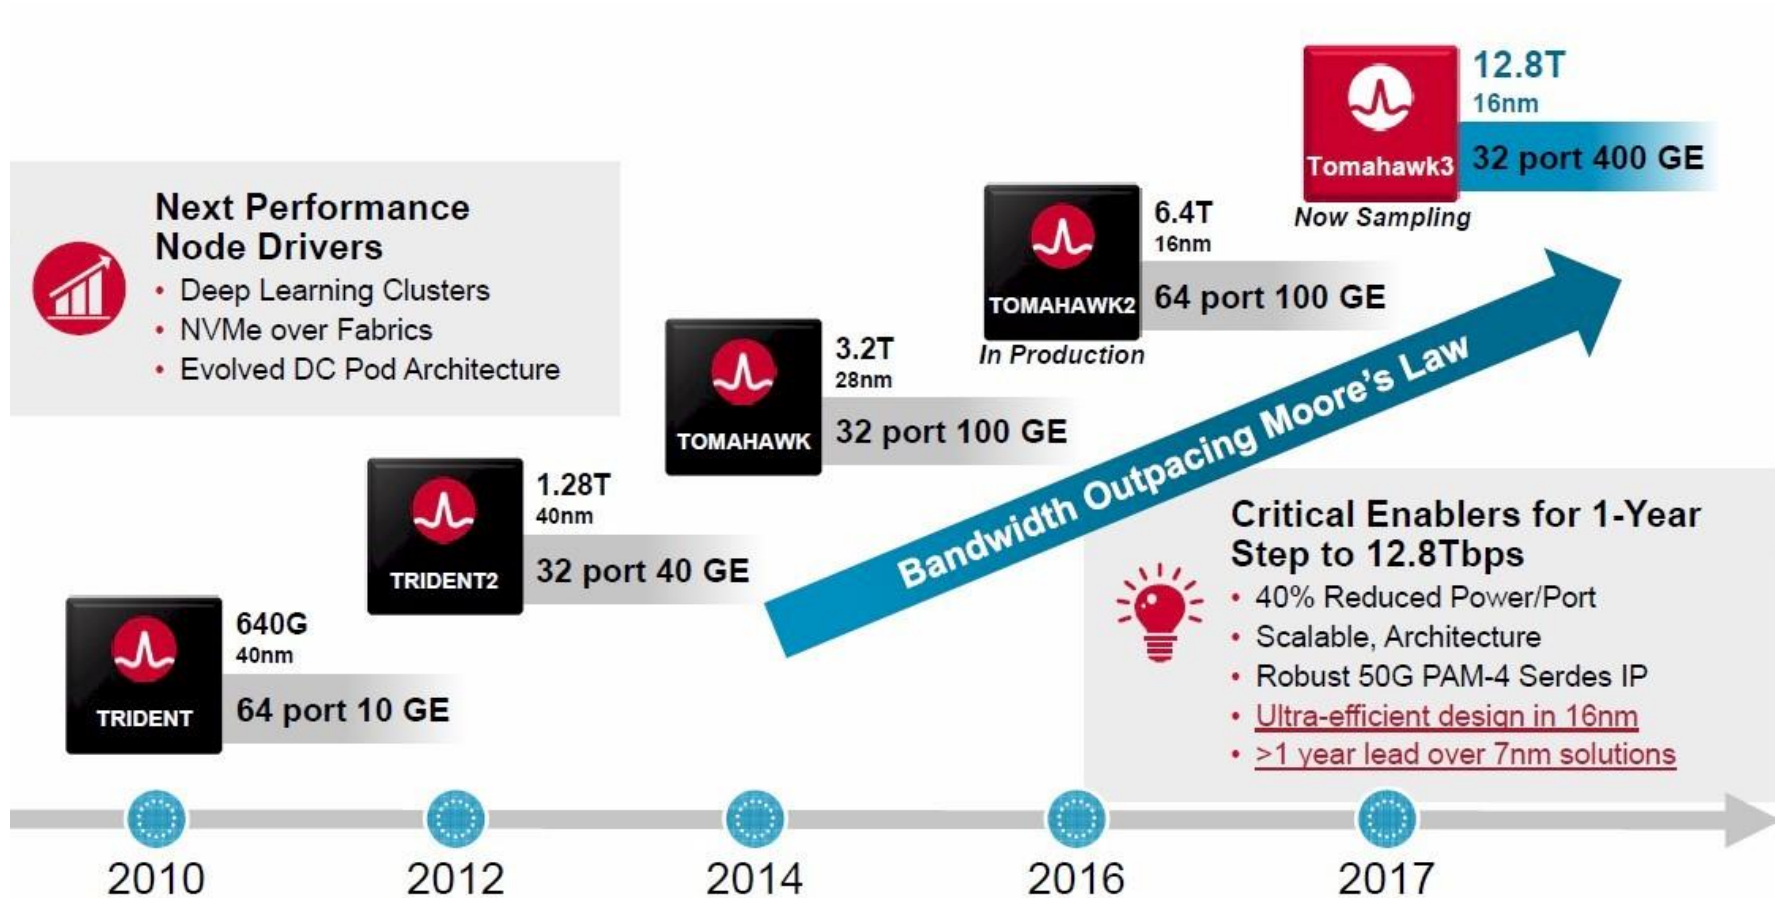
KVM



Xen
Server™

3) SDN / Open Networking

High Capacity switching chipsets



3) SDN / Open Networking

Most vendors adopting “Merchant Silicon”

Juniper QFX3500



IBM BNT RackSwitch G8264



Alcatel-Lucent OmniSwitch 6900



Cisco Nexus® 3064



HP 5900AF 48XG



Dell Force10 S4810

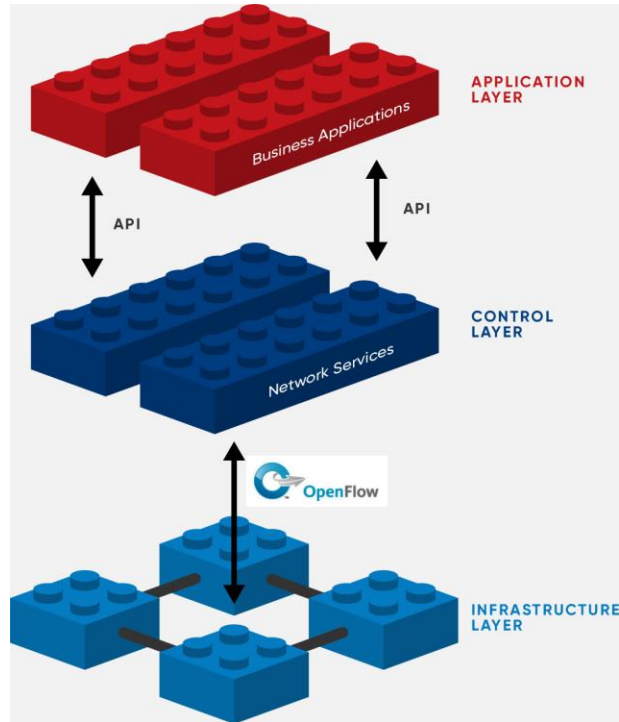


Arista 7050S-64



3) SDN / Open Networking

SDN is not a requirement, but increases the potential of NFV



Easily portable to general purpose **CPUs**, therefore highly virtualizable

Achievable by using **ASICs or CPUs** (dataplane acceleration)




- NFV and SDN aligns closely with the objectives to use commodity servers and switches.
- Control plane can be easily ported to NFV
But, dataplane too (lots of improvements here)

4) Open Source initiatives



Current landscape provides several open source initiatives, that cover most of the industry challenges:





Open Source
MANO


Quick review of NFV

ETSI



Home of NFV

What is NFV trying to address?



- Network Function Virtualization (NFV) proposes to virtualize network functions that typically run in dedicated appliances
- The main goal is to support virtualized functions over COTS servers.
- Virtual Network Functions (VNFs) acquire all the advantages of Cloud Applications!

The original idea.... triggered an industry movement

- Initial white paper was written in 2012 by the world's leading telecom network operators (Europe, América & Asia).
- This group evolved to the ETSI NFV ISG (Industry Specification Group), formed today by 300+ companies.
- Their main motivation had to do with reducing TCO of building a network by using open solutions.

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.

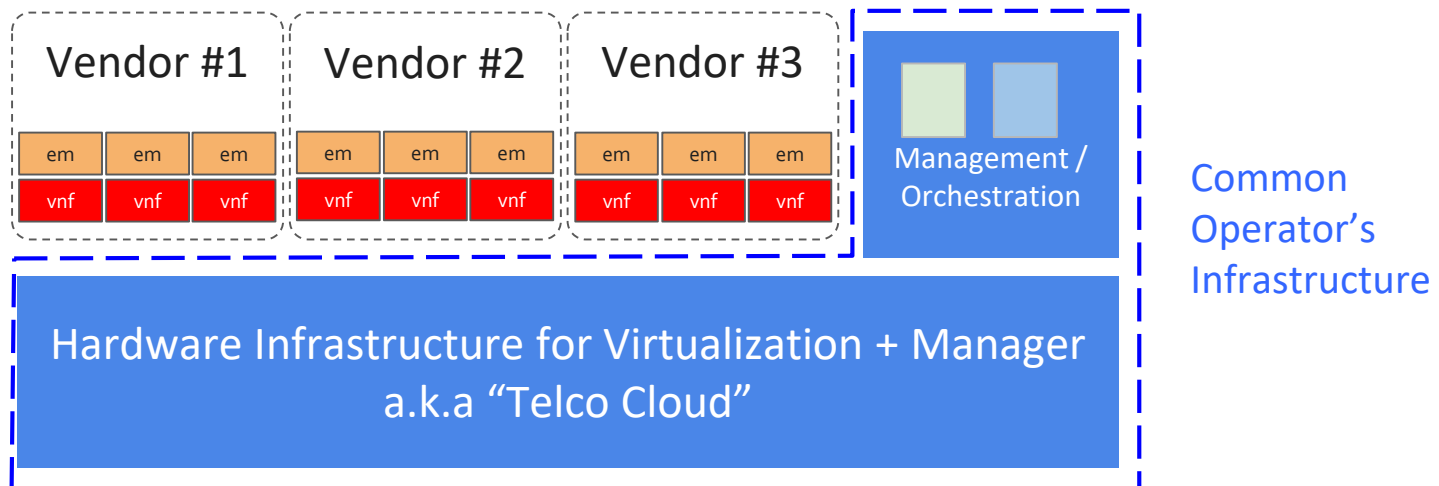
ETSI Publications

- 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



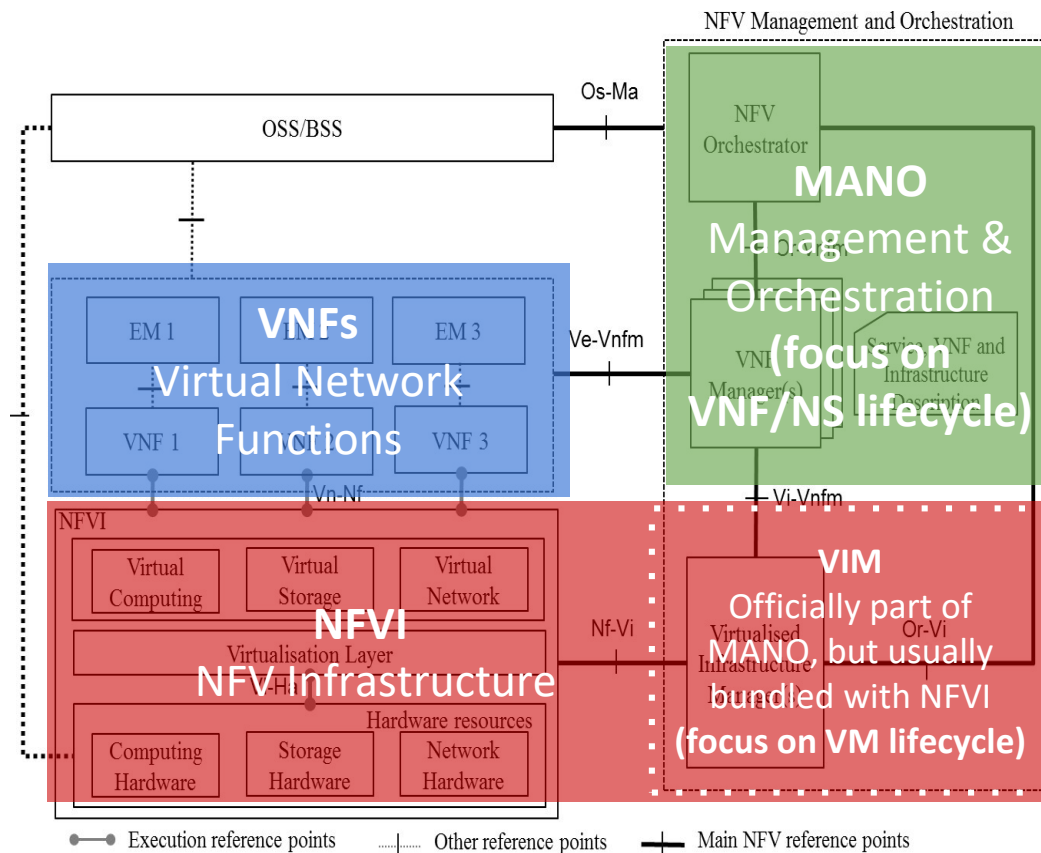
Benefits of a standard NFV architecture

We are looking for a **unified and generic virtualization infrastructure**, compatible with any vendor's Virtual Networking Function (VNF), **so standardization is a must.**




The ETSI NFV Architecture

The standard architecture can be better understood in three blocks:



NFVI: NFV Infrastructure

NFVI goal is to provide a virtualization environment for VNFs, including virtual compute, storage and networking resources.

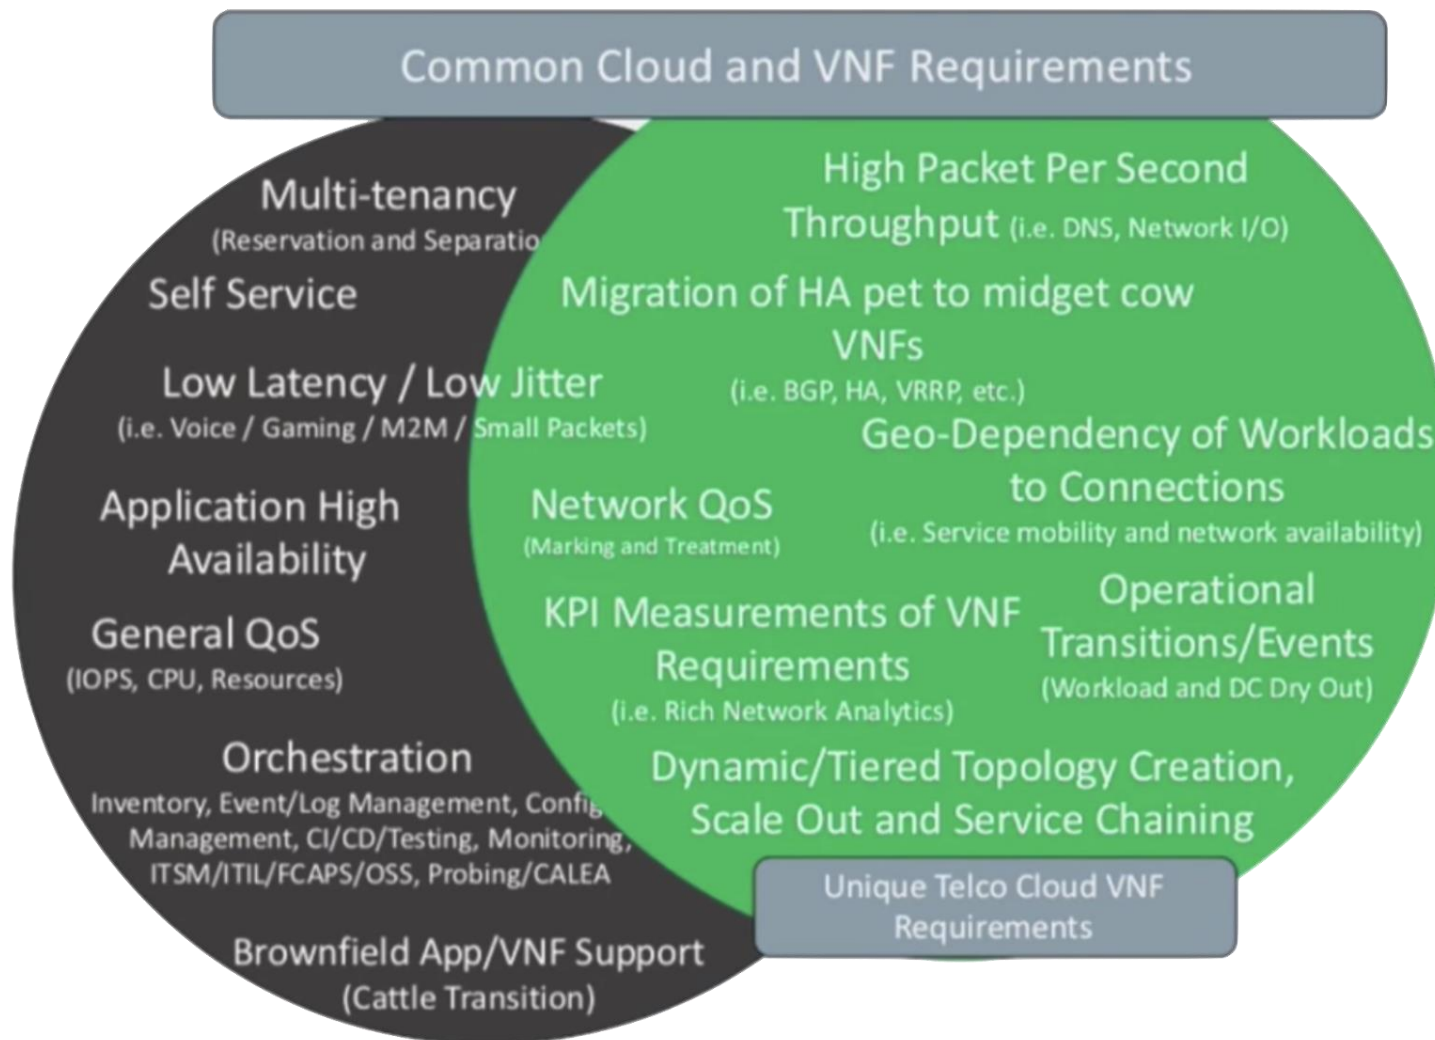


- But! Networking applications may have more strict performance requirements, we will discuss that later.



NFVI: NFV Infrastructure

VNF Special Requirements



NFVI: NFV Infrastructure

VNF Special Requirements



VNFs, especially data-plane ones, usually have additional requirements than common cloud applications, including:

- **Minor latency** (disk I/O & network)
 - Faster hardware (More cores, SSD disks, faster buses)
 - Dataplane acceleration
- **Higher throughput or PPS**
 - Dataplane acceleration
 - EPA: Enhanced Platform Awareness
- **Geographical distribution**
 - multi-site cloud
- **Horizontal auto-scaling**
 - automated operations (orchestration)

NFVI: NFV Infrastructure


VNF Special Requirements

EPA covers the different approaches that can be taken at the NFVI layer to increase performance while maintaining a generic (COTS) infrastructure. VIM and MANO should be able to request them.


Huge Pages




NUMA Topology Awareness



CPU Pinning




Data Plane assignment



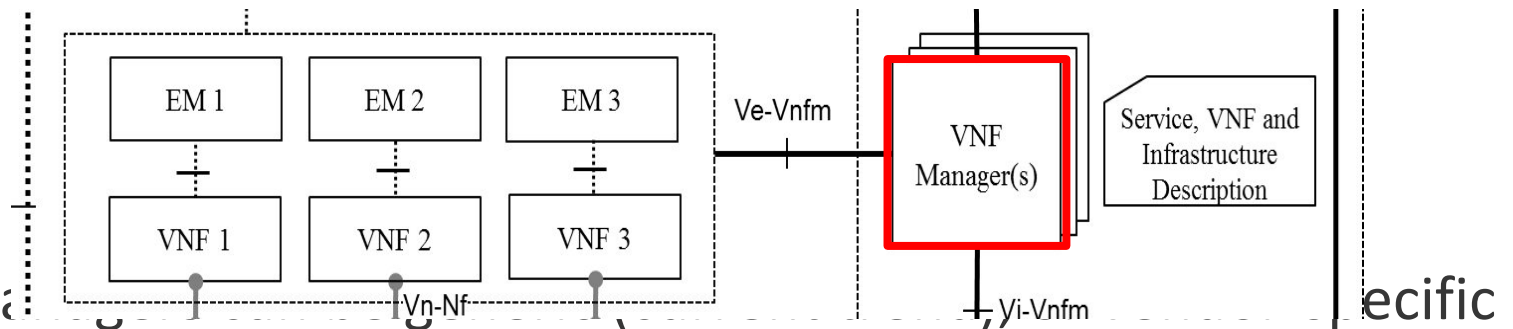
VIM: Virtualized Infrastructure Manager

- The Virtualized Infrastructure Manager is part of the 'MANO Stack' and addresses provides lifecycle management for virtualized resources (VMs, volumes, networking paths and connectivity, etc.)



VNF-M: VNF Manager

- The VNF Manager, also part of the 'MANO Stack', covers **lifecycle management for Virtual Network Functions (VNFs)**, either directly or through their own Element Management System (EMS).



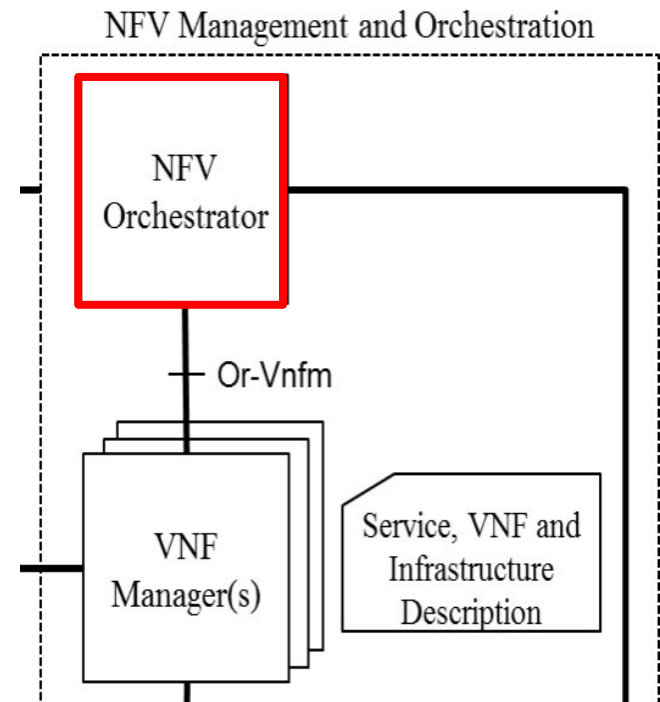
- VNF Manager(s) are specific

CANONICAL



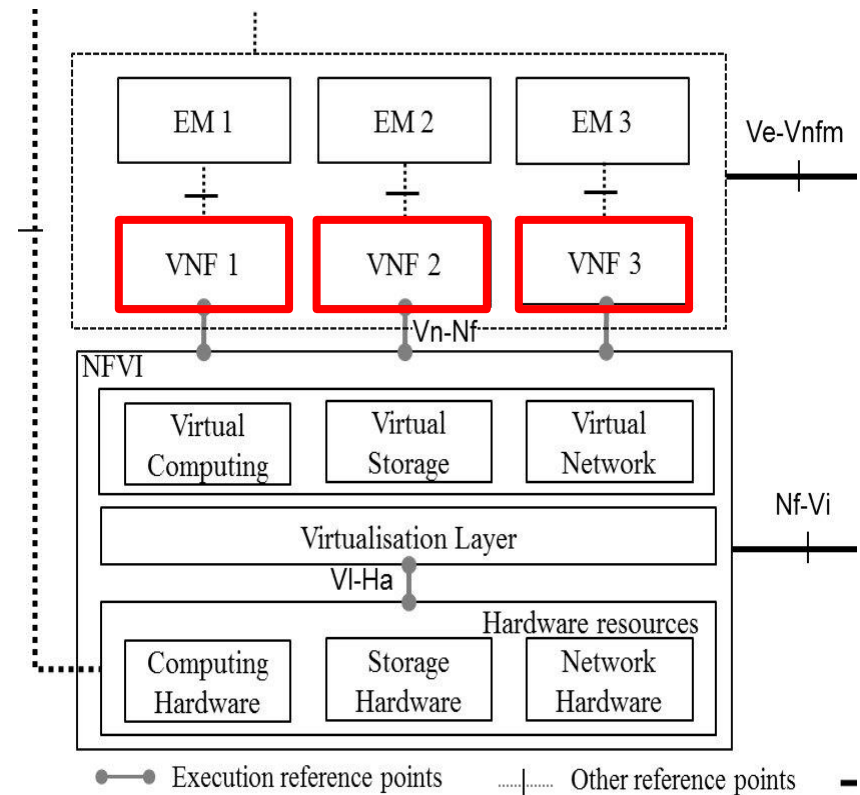
MANO: NFV Orchestrator (NFV-O)

- The NFV Orchestrator, the higher entity in the 'MANO Stack', covers general resource orchestration and services lifecycle, which comprise multiple VNFs and define their roles (traffic paths, scaling decisions, and other service-related requirements)
- It can interact with a generic VNF Manager, or vendor-specific ones.



Virtual Network Functions (VNF)

- Finally, the VNFs, which are supported by the underlying NFVI, and managed by their own EM (internal manager) and the VNF Manager (external, 'context-aware' manager)
- They should be able to provide any networking function and interact with other VNFs.



VNF Descriptor files (VNFD)

One of the most important aspects of achieving a unified VNF catalogue, is having a standard way of describing VNFs.

- MANO solutions should give the possibility to describe VNFs through 'descriptor files'
- The industry's goal is a unified and standard descriptor file format across different platforms.
- Both NS (comprised of VNFs) and VNFs should be described in a simple way.

```
vnfd:vnfd-catalog:
  vnfd:vnfd:
  - vnfd:connection-point:
    - vnfd:name: eth0
      vnfd:type: VPORT
    vnfd:description: Generated by OSM pacakage generator
    vnfd:id: ubuntuvmf_vnfd
    vnfd:mgmt-interface:
      vnfd:cp: eth0
    vnfd:name: ubuntuvmf_vnfd
    vnfd:service-function-chain: UNAWARE
    vnfd:short-name: ubuntuvmf_vnfd
    vnfd:vdu:
  - vnfd:cloud-init-file: cloud_init
    vnfd:count: '1'
    vnfd:description: ubuntuvmf_vnfd-VM
    vnfd:guest-epa:
      vnfd:cpu-pinning-policy: ANY
    vnfd:id: ubuntuvmf_vnfd-VM
    vnfd:image: ubuntu_admin
    vnfd:interface:
  - rw-vnfd:floating-ip-needed: 'false'
    vnfd:external-connection-point-ref: eth0
```

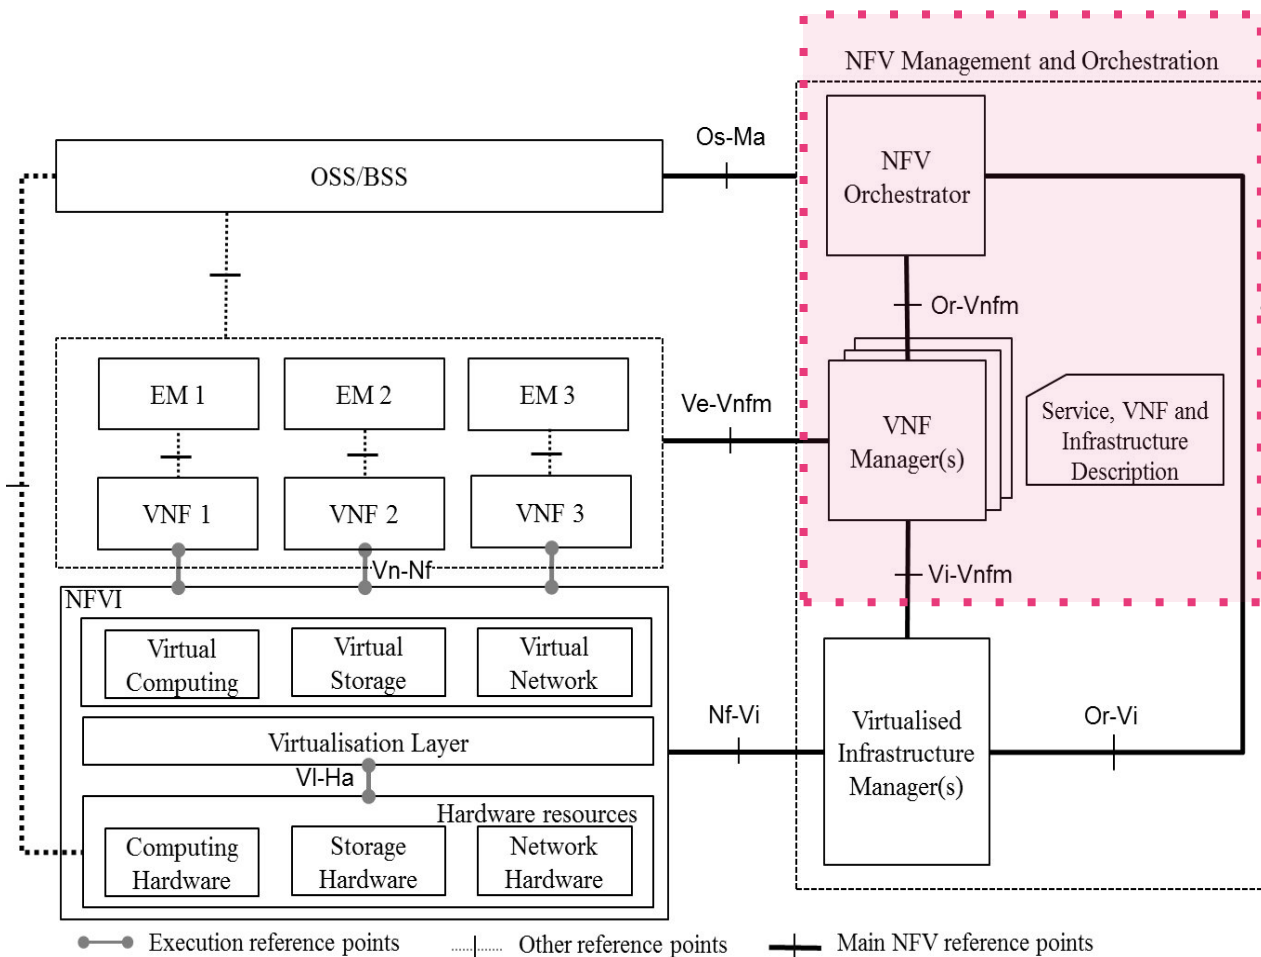


Open Source
MANO

Introduction to OSM Release Five




The Open Source MANO Project



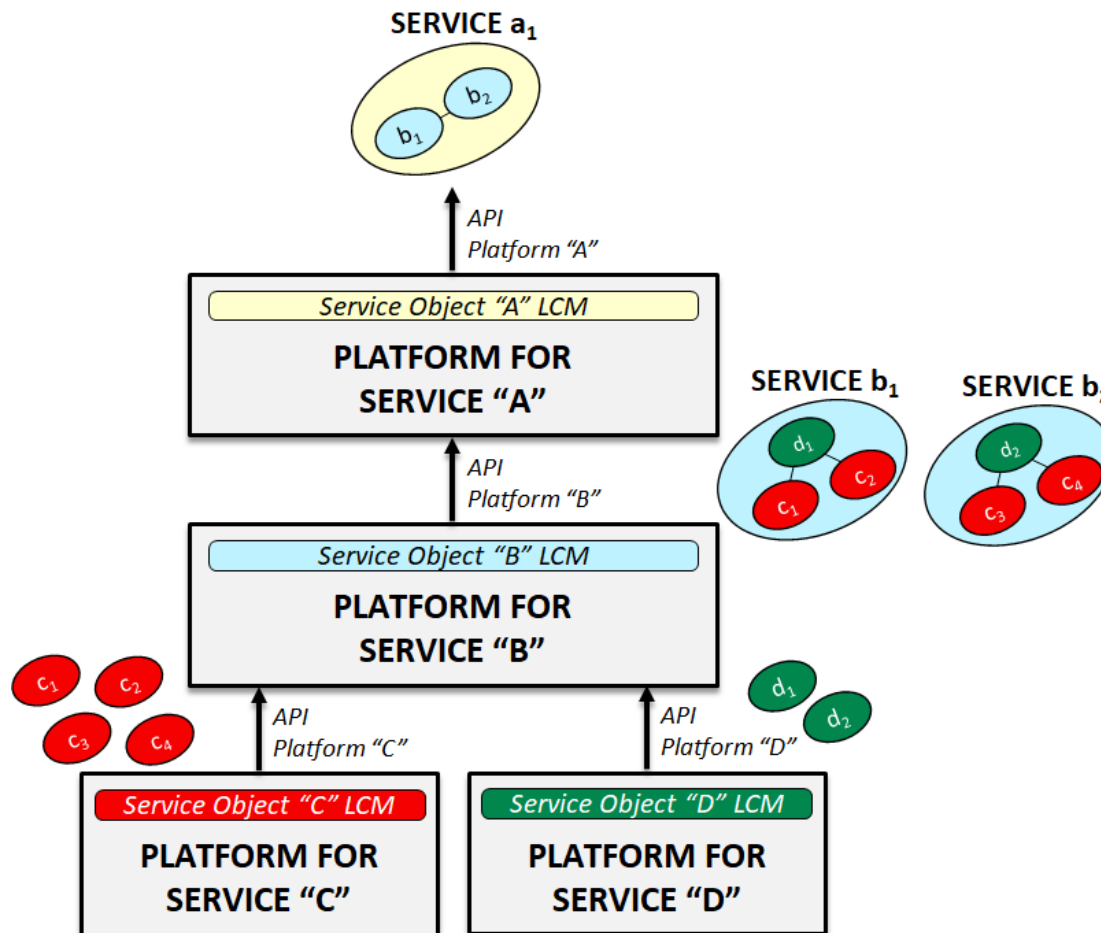
We are here!
Open Source MANO is an ETSI-hosted project to develop an Open Source NFV Management and Orchestration (MANO) software stack aligned with ETSI NFV.

OSM Architectural Principles



Layered Model


A multilayered model, where each layer provides a “service object”, composed by service objects provided by lower layers.



OSM's approach aims to minimize integration efforts

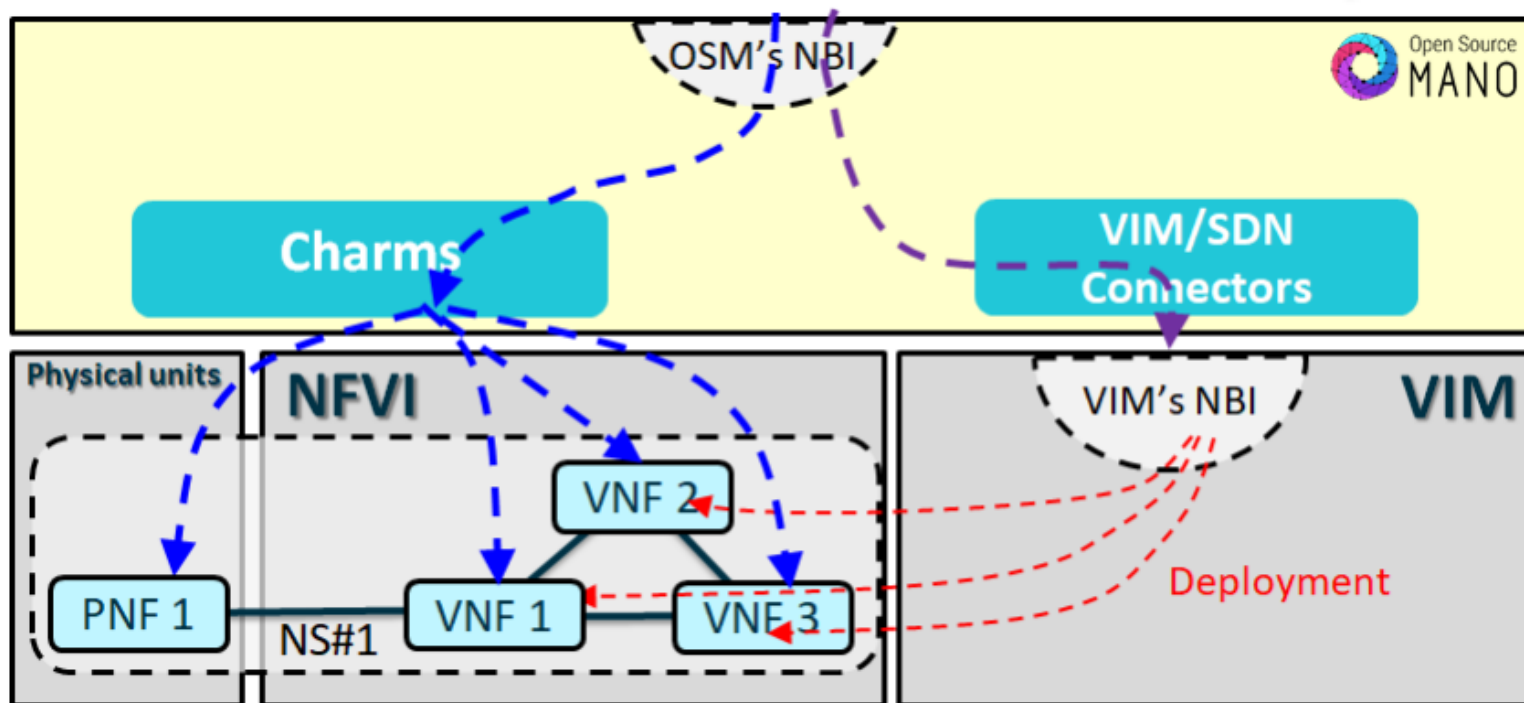
1. A well-known Information Model (IM), aligned with ETSI NFV, that is capable of modelling and automating the full lifecycle of Network Functions:

- VNFD (VNF Descriptor),
VNFR (VNF Record),
- NSD (Network Service Descriptor),
NSR (Network Service Record),
- NST (Network Slice Template),
NSI (Network Slice Instance)



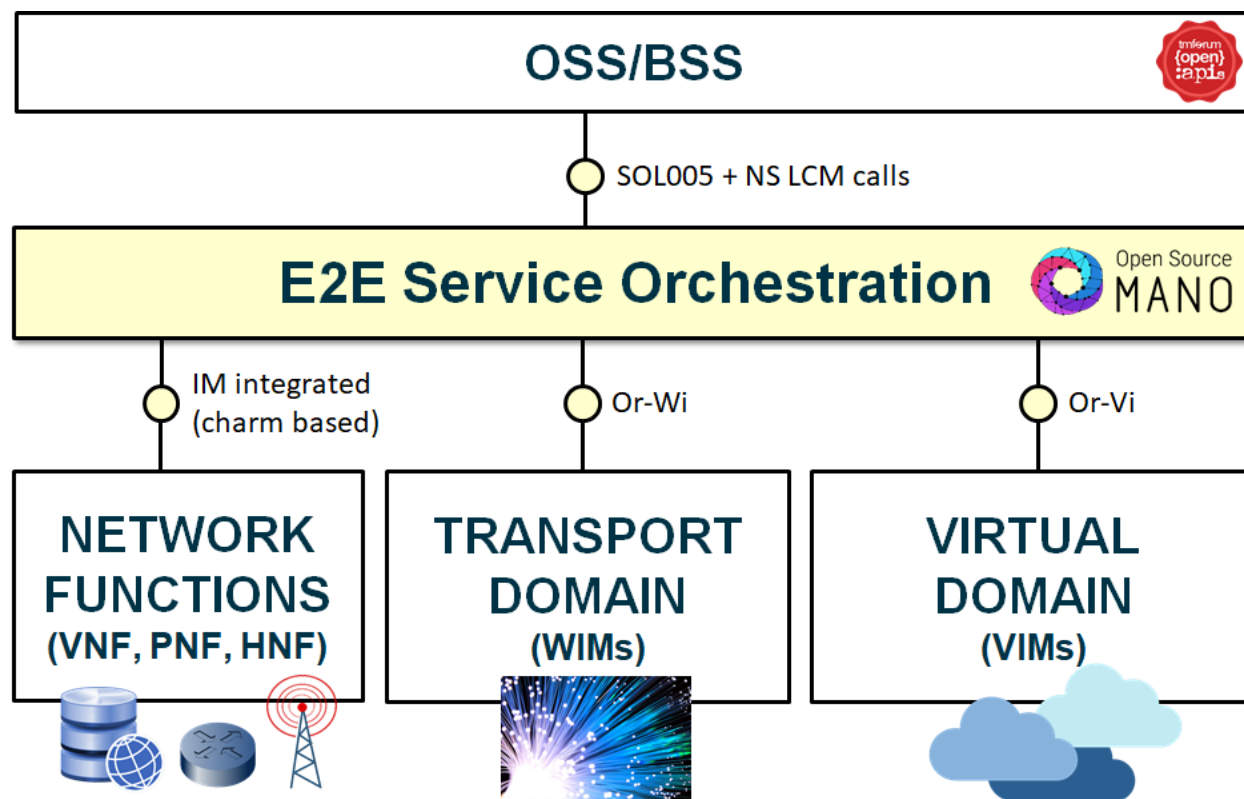
OSM's approach aims to minimize integration efforts

2. A unified northbound interface (NBI), based on NFV SOL005



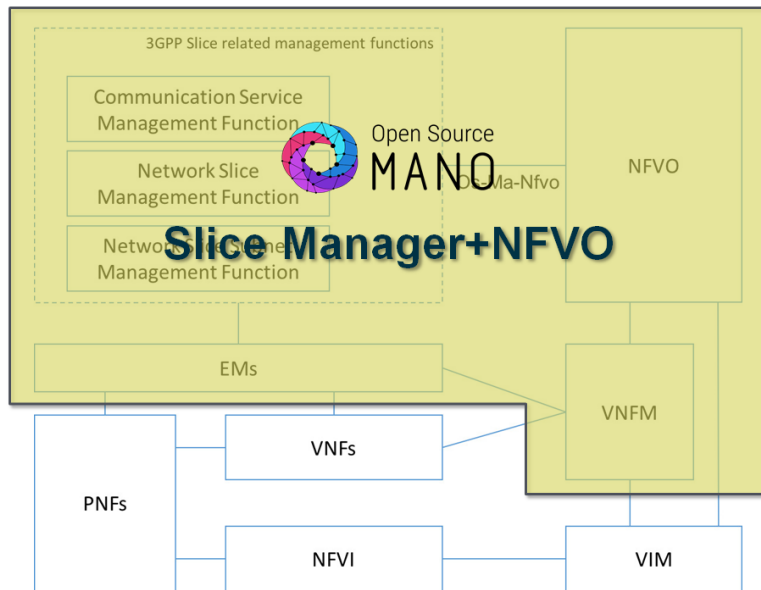
OSM's approach aims to minimize integration efforts

3. The extended concept of “**Network Service**” in OSM, so that an NS can span across the different domains identified and therefore control the full lifecycle of an NS interacting with VNFs, PNFs and HNFs.

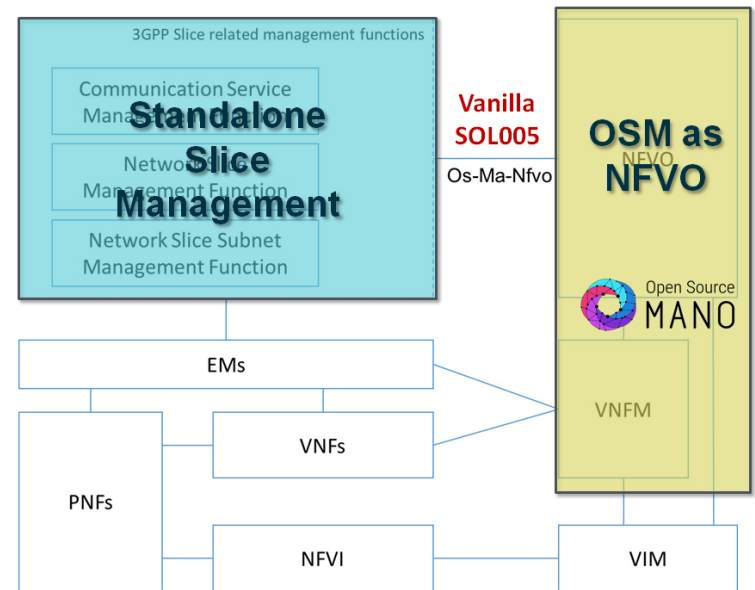


OSM's approach aims to minimize integration efforts

4. In addition, OSM can also manage the lifecycle of **Network Slices**, assuming if required the role of Slice Manager, or integrating with an external Slice Manager



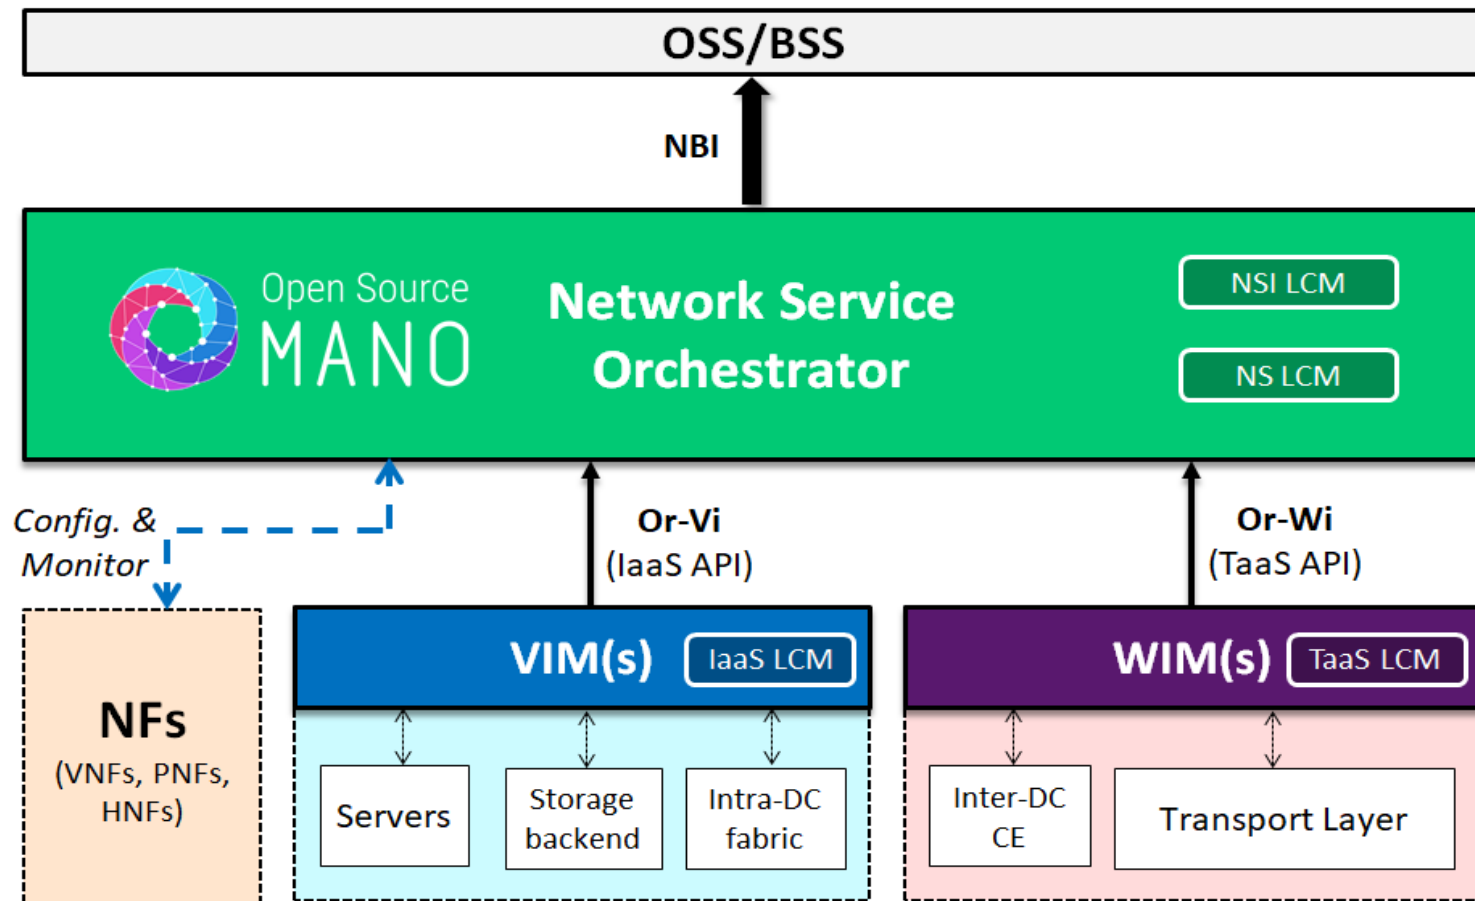
Full E2E Management
(Integrated Modelling)



Standalone Management
(Vanilla NFV/3GPP)

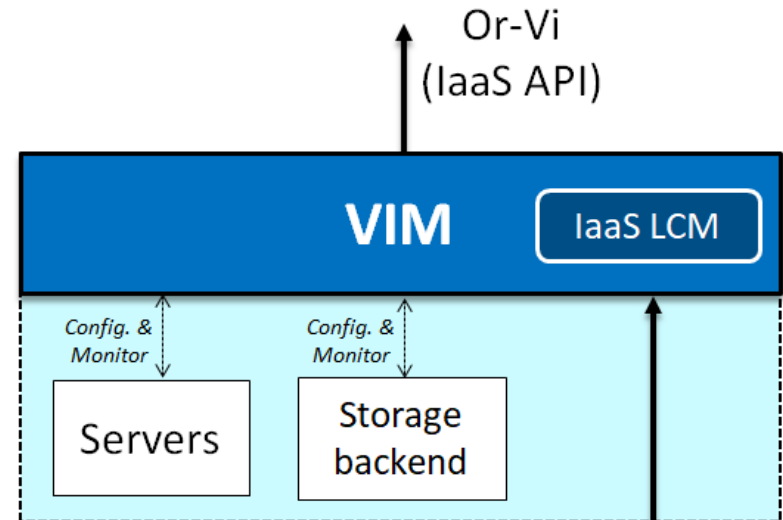
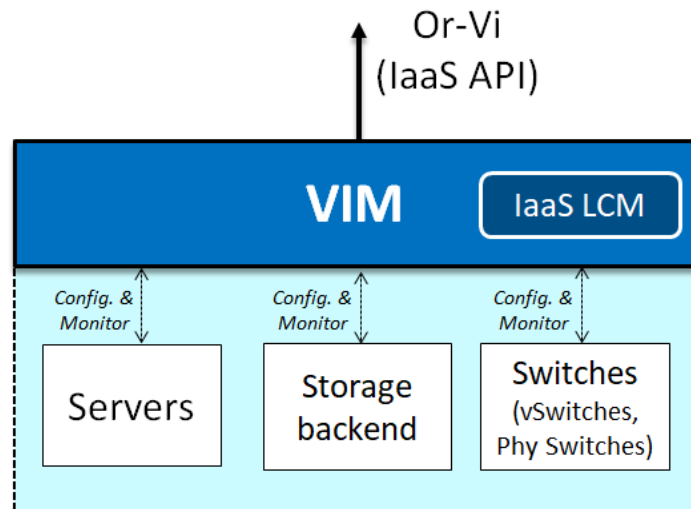
Service Platform view

OSM as a Network Service Orchestrator (NSO)




Service Platform view


VIM manages the virtual network to support VNF's connectivity



In advanced cases, the VIM might (transparently) control an external SDN Controller, to fulfill the connectivity requirements

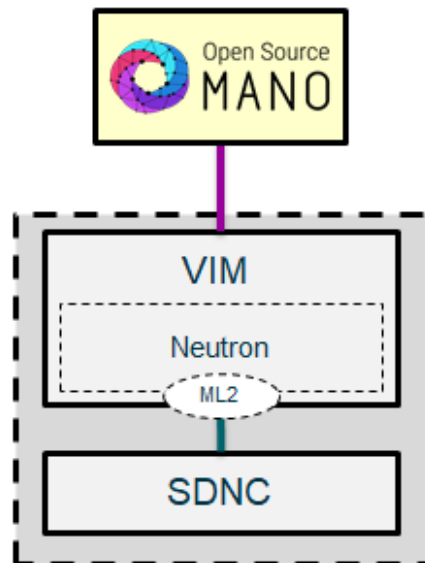


Integration with SDN Controller



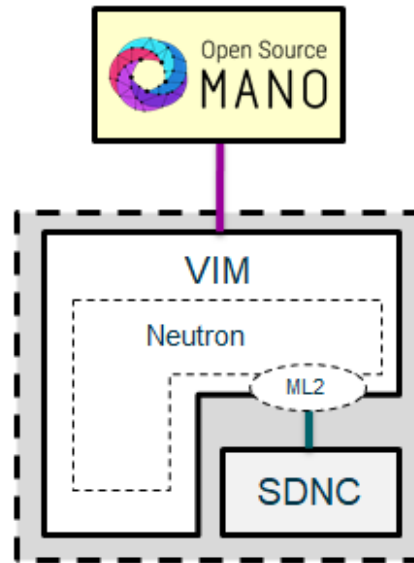
CASE #1: Vanilla

- Overlay: Native
- No underlay



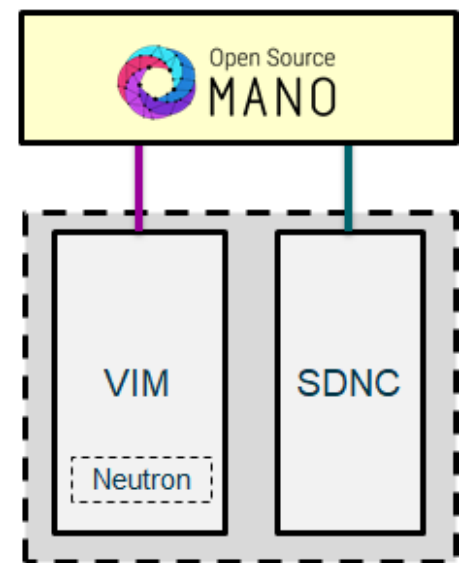
CASE #2: VIM + all SDN

- Overlay: SDNC
- Underlay: if available, via SDNC



CASE #3: VIM with partial SDN

- Overlay: Native
- Underlay: SDNC



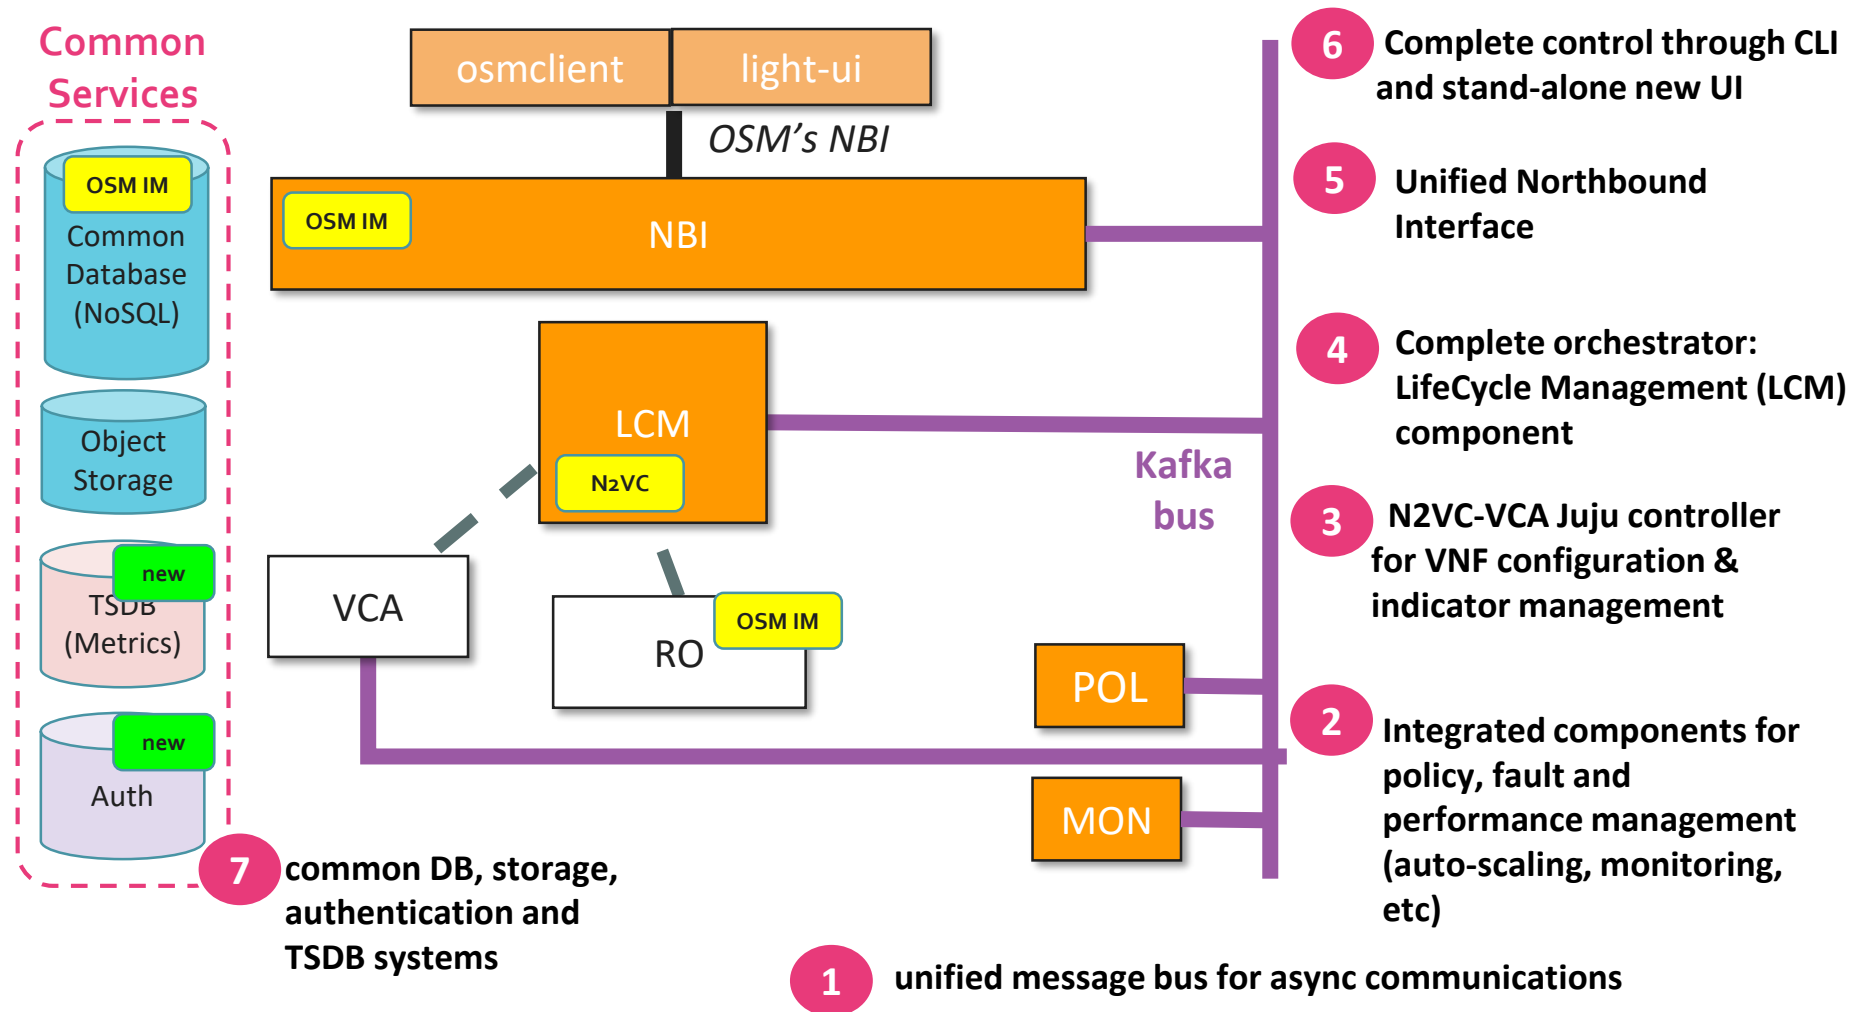
CASE #4: SDN Assist

- Overlay: VIM native
- Underlay: SDNC, via OSM

SDN Assist

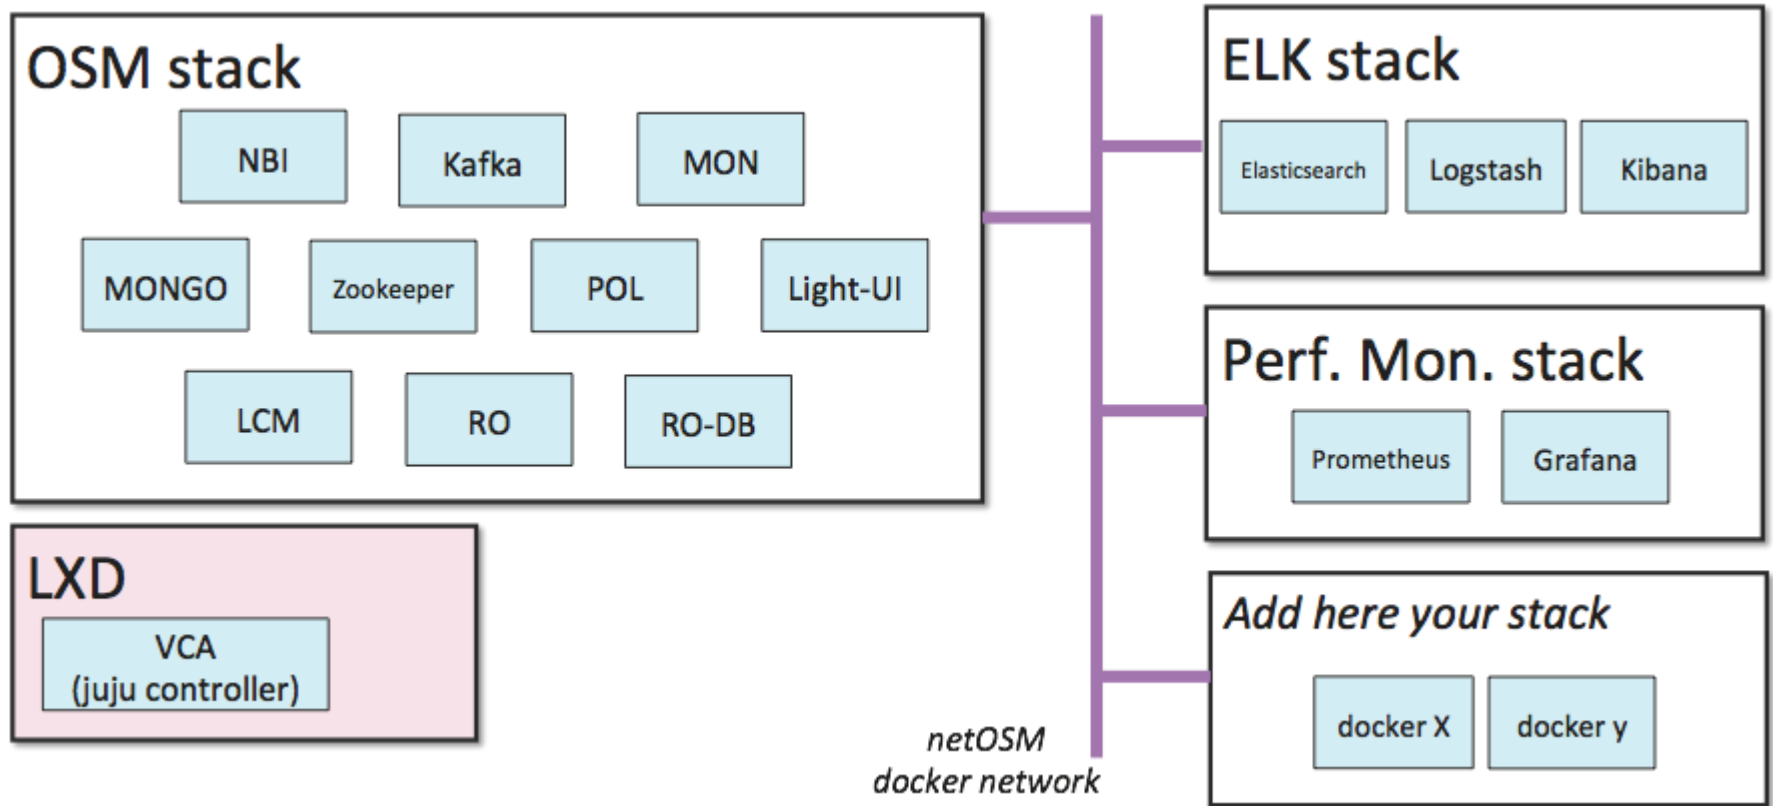
Allows OSM to control SDN connectivity, even when not possible by the VIM (eg: PCI Passthrough, SR-IOV)

Release FOUR+ architectural view



Release FOUR+ architecture

Microservice architecture to enable extensibility



Why is OSM Awesome?



It has a large and diverse community! More than 110 members!

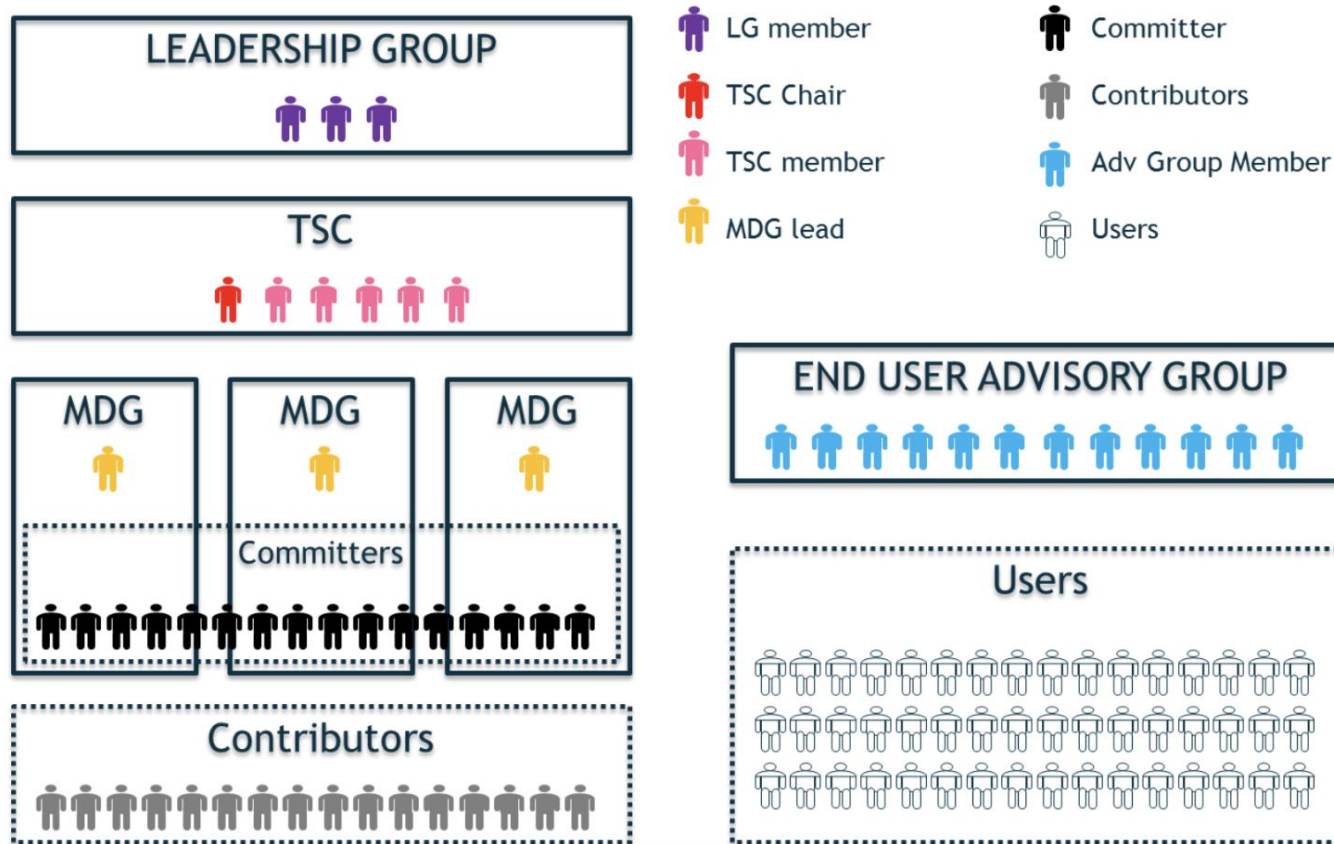
Logos shown include: Sprint, CableLabs, amazon web services, Bell, verizon, BT, telenor, Telefonica, PT, SK telecom, and STC.

- 12 Global Service Providers
- Leading IT/Cloud players
- VNF providers



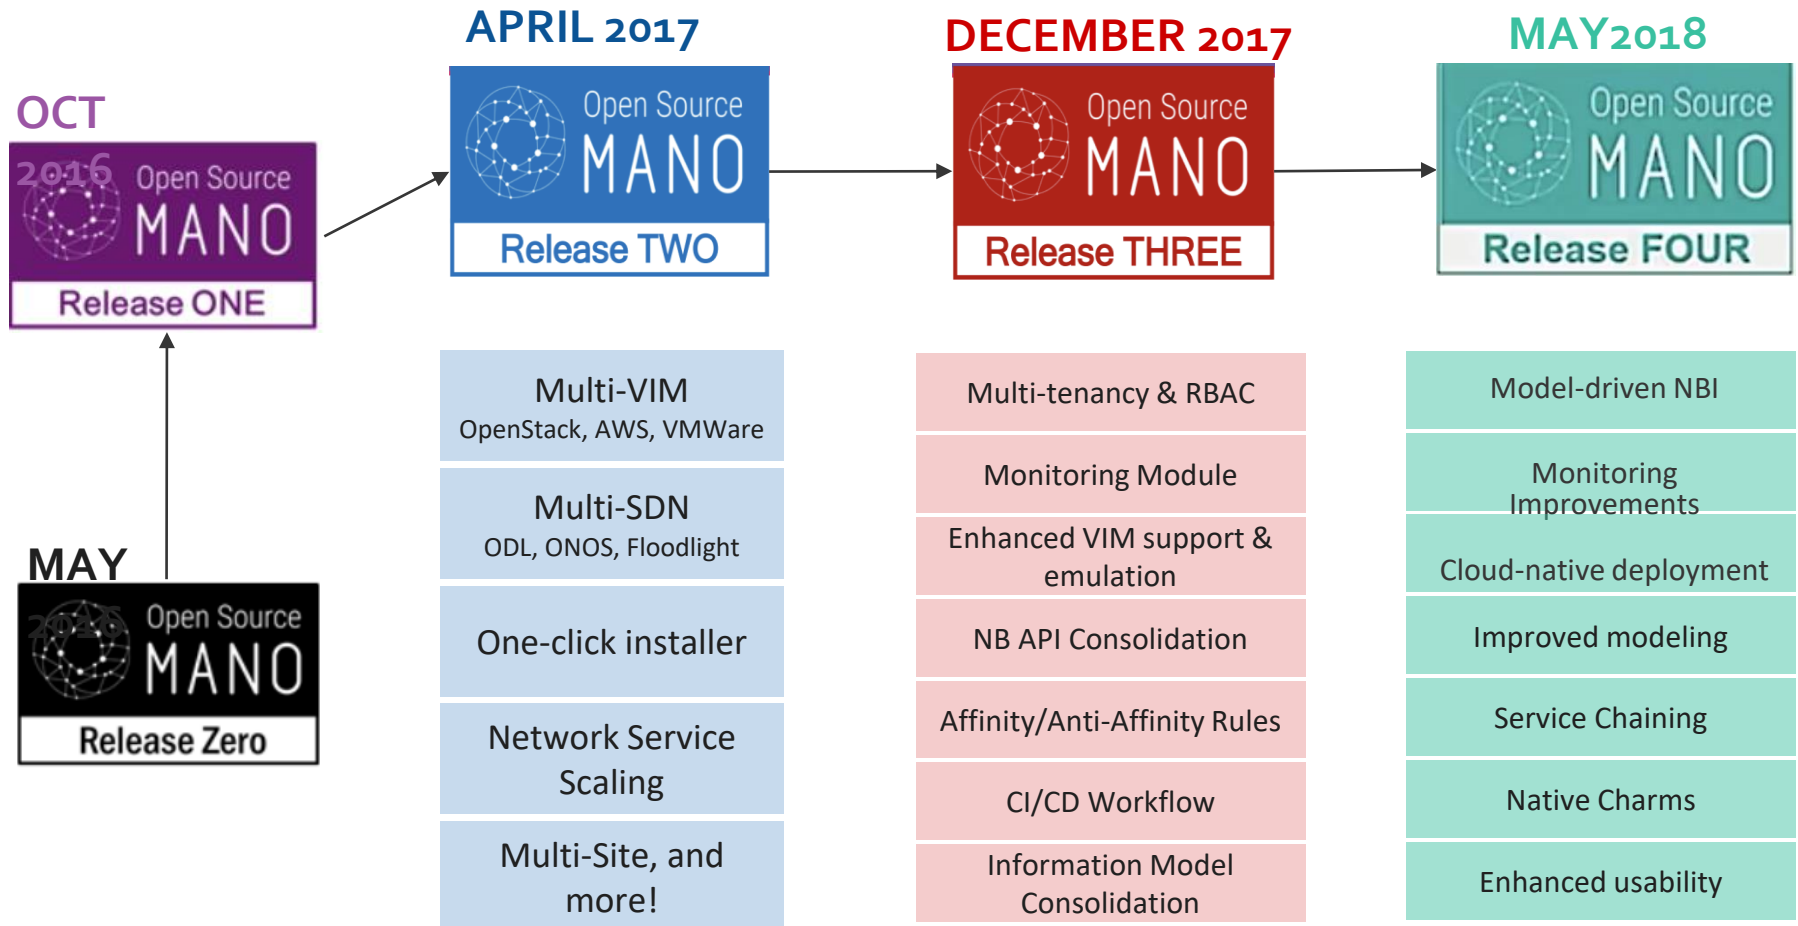
Why is OSM Awesome?

It is well organized for producing production-ready upstream code



Why is OSM Awesome?

It prioritizes features for production readiness...



Why is OSM Awesome?



...and launched Release FIVE in December 2018, with key features for 5G and production environments!


The most ambitious release so far!

- Support for hybrid Network Services! VNFs and PNFs
- Network Slicing for 5G
- Multi-site Inter-DC extensions
- Monitoring & Policy improvements (VNF + VIM metrics, autoscaling)
- Improved modelling of VNF networking (SFC, multi-VDU relations)
- Better user experience (GUI package composer, faster installation, events/logs GUI)



Open Source
MANO

Contributing
to the
Community



- Join [here](https://osm.etsi.org/about/how-to-join) as a company or individual contributor!
<https://osm.etsi.org/about/how-to-join>

HOW TO GET INVOLVED IN OSM

There are two paths to get involved in OSM as an organisation: as an ETSI Member, or as an OSM Participant.

Check first if your organization is already involved by consulting the list of [OSM Members](#) and [Participants](#).

Get involved as an ETSI Member

To take part in the development of OSM and participate to the meetings, ETSI Members need to sign the [OSM Membership Agreement and CCLA](#). In doing this, they agree to the OSM operating rules which in some cases are different from those in ETSI's Technical Working Procedures. [Check if your company is an ETSI Member](#).

Get involved as an OSM Participant

Organizations who are not members of ETSI may also participate in OSM, attend meetings and help to develop OSM by making technical contributions. They are not applicable for leadership (LG) positions and must pay a participation fee to attend OSM meetings. To get involved as a Participant, please sign the [OSM Participant Agreement and the CCLA](#).


Developers and Users

Individual developers and end users are welcome to contribute code and feedback to OSM, they just need to [create an individual contributor or user account](#).

- **Weekly Conference Calls**
 - Technical, leadership, DevOps, and more!
- **Face to Face Meetings**
 - Plenaries and Mid-Release meetings (every 3 months)
 - Next location: Santa Clara, US (May 2019)
- **OSM Hackfest**
 - Sixth edition taking place on May 2019 at Santa Clara, expecting to keep co-locating with OSM Face-to-Face meetings.

Ways to contribute to OSM

- **Try OSM** and give feedback to the community.
- Join as a developer to **make contributions to the code**.
- Join the community to **contribute to design discussions**.
- **Start building your own distribution** of OSM as an integrator.
- **Host an OSM meeting** to contribute to the community's growth and diversity.



Open Source
MANO

Find us at:
osm.etsi.org
osm.etsi.org/wikipub