XXXII Simpósio Brasileiro de
Redes de Computadores e Sistemas Distribuídos
Florianópolis, 5 a 9 de Maio de 2014

Minicurso 1

Network Function Virtualization: Perspectivas, Realidades e Desafios.

Autores:
Raphael Rosa (UNICAMP)
Marcos Siqueira (UNICAMP)
Emerson Barea (UFSCar)
Cesar Marcondes (UFSCar)
Christian Rothenberg (UNICAMP)

Network Functions Virtualisation (NFV)

A joint operator initiative and call-for-action to industry

A joint operator push to the IT and Telecom industry,
to provide a new network production environment,
based on modern virtualization technology,
to lower cost, raise efficiency and to increase agility.

We believe Network Functions Virtualisation is applicable to any data plane packet processing and control plane function in fixed and mobile network infrastructures (WP)

at&t BT CenturyLink 中国移动通信 CHINA MOBILE colt smarter / faster / further T KDDI NTT orange TELECOM ITALIA Telefónica verizon

Agenda

- Motivation;
  - Problem Statement, Trends in IT & Telecom challenges
- Network Functions Virtualization
  - Vision; Approach; Benefits & Promises
  - The ETSI NFV ISG; WG; Architecture
- NFV Requirements and Challenges
- Use Cases, Proof-of-Concepts
- Enabling Technologies
- DEMO: Vyatta

Motivation
Problem Statement

- Complex carrier networks
  - with a large variety of proprietary nodes and hardware appliances.
- Launching new services is difficult and takes too long
  - **Space and power to accommodate**
  - requires just another variety of box, which needs to be integrated.
- Operation is expensive
  - **Rapidly reach end of life**
  - due to existing procure-design,-integrate-deploy cycle.

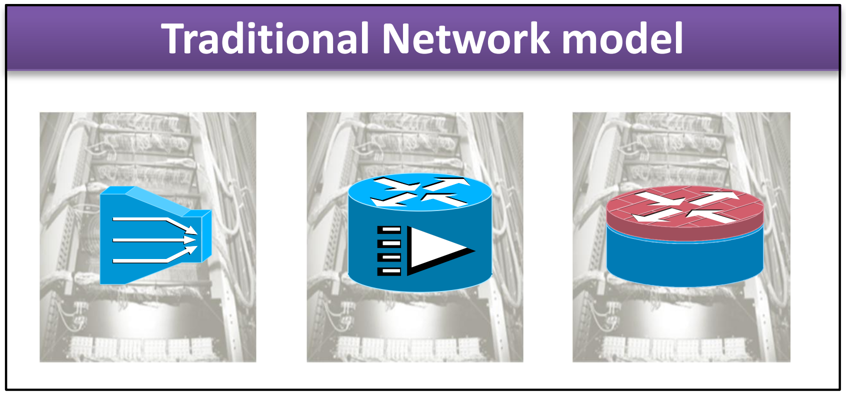

- Network functionalities are **based on specific HW&SW**
- **One physical node per role**

Sisyphus on Different Hills

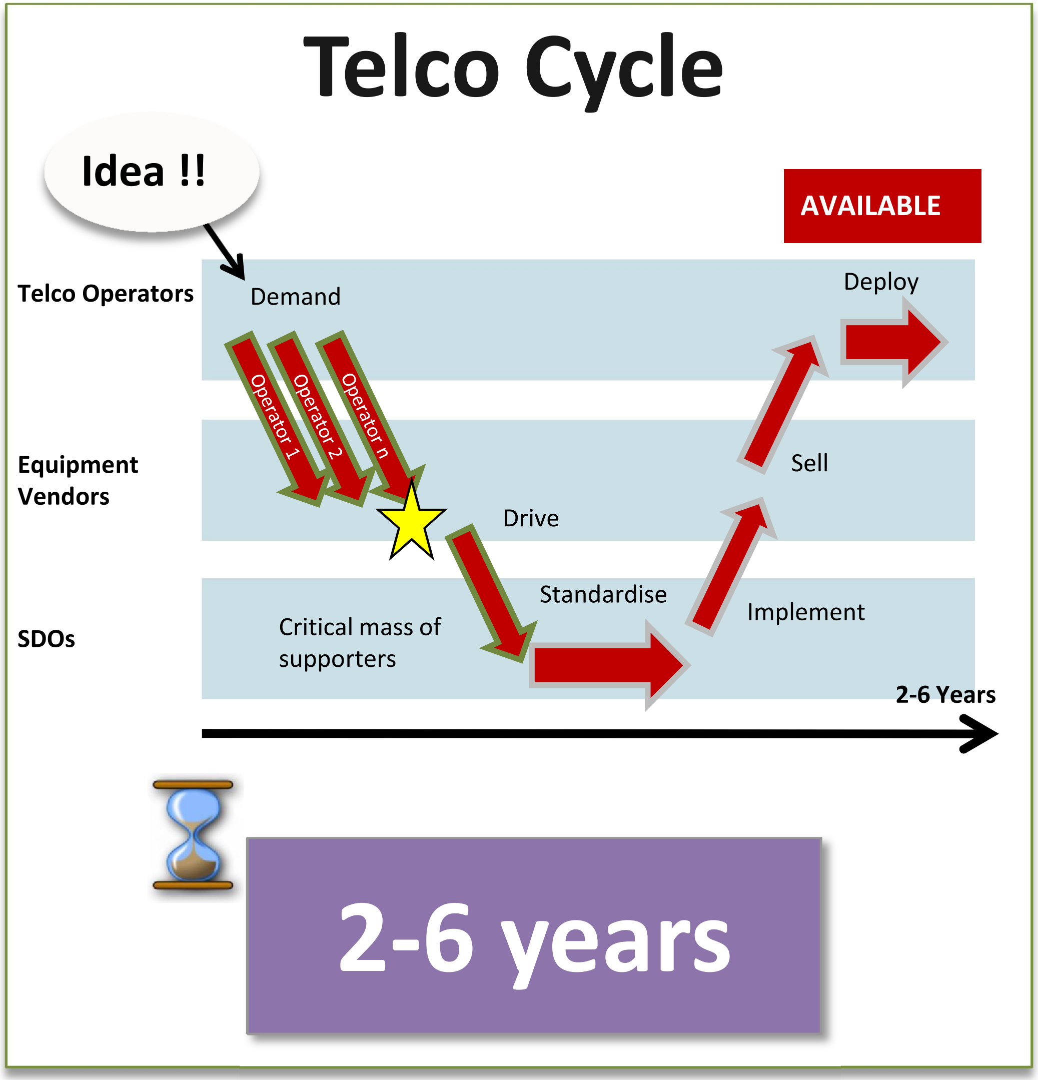

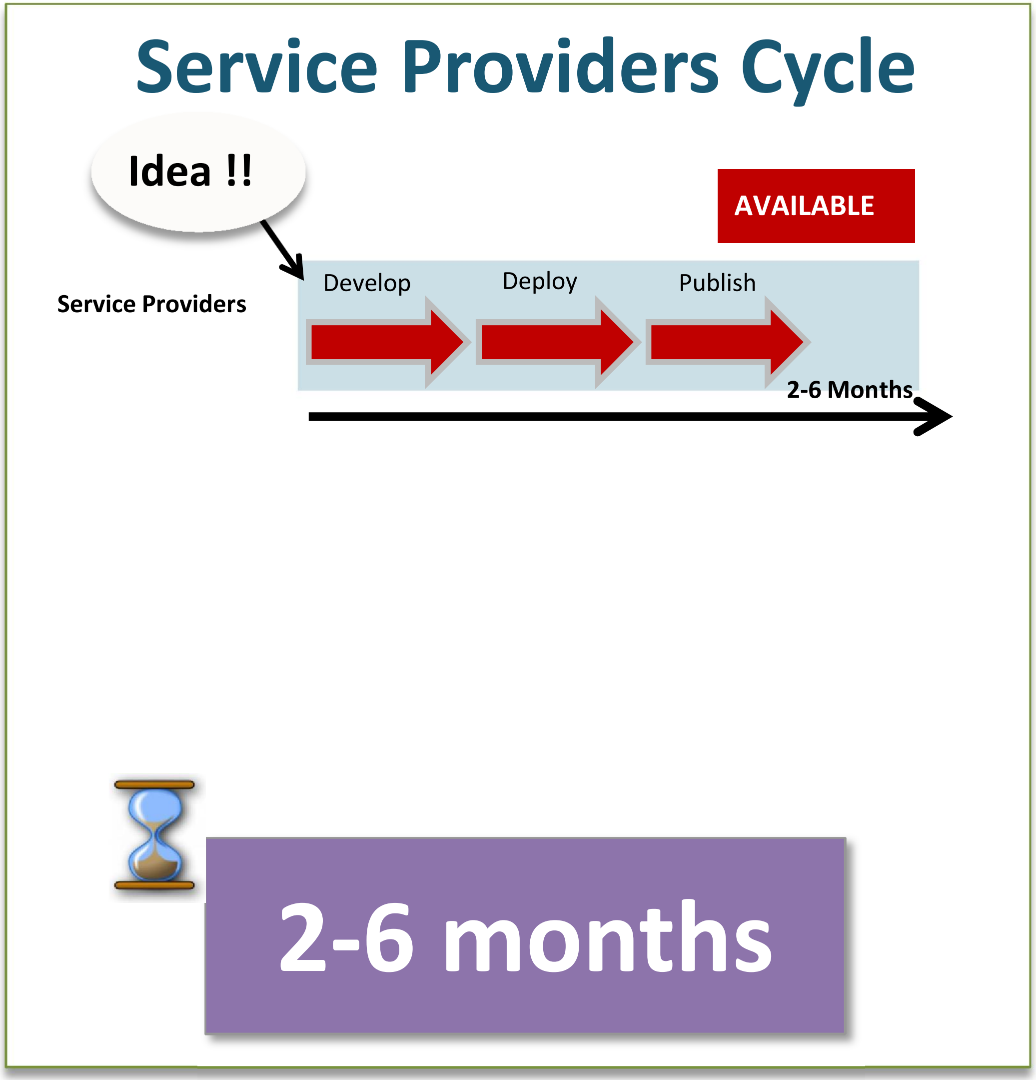

Source: Adapted from D. Lopez Telefonica I+D, NFV

Enter the Software-Defined Era

Traditional telcos

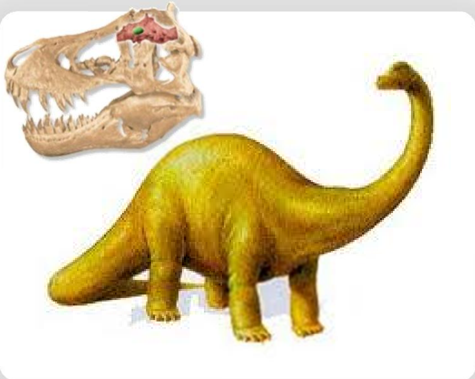

- Very **intensive in hardware**
- Software not at the core

Internet players

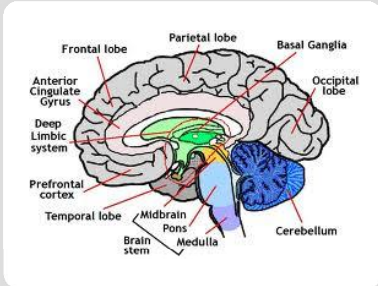

- Very **intensive in software**
- Hardware is a necessary base

HARDWARE

SOFTWARE

AT&T, Telefonica, Telebras

Google, Facebook

Adapt to survive: Telco evolution focus shifting from hardware to software

Source: Adapted from D. Lopez Telefonica I+D, NFV

Scale and Virtualization in the Timeline

Early twentieth century

- Manual Switching
- Very intensive in human resources
- Era **dominated by hardware**

Mid-twentieth century

- Electromechanical Switching
- Less intensive in human resources
- Era **dominated by complex hardware**

Virtualization technologies enables overcoming physical constraints and generating multiplexing gains…

Second half of the twentieth century

- Digital Switching
- Much less intensive in human resources
- Era **dominated by complex and specific hardware**. **Software appears and is important**
- Services defined by telco

Early twenty-first century

- Internet connectivity opens the door to the development of OTT services (without operator)
- **Software becomes a differentiation asset**

Source: Adapted from D. Lopez Telefonica I+D, NFV

Observation

- **Commercial-off-the-shelf IT-platforms**
  - allow to host a large variety of applications.
- **New virtualization technology allows to abstract HW,**
  - enables elasticity, scalability and automation.
- **Network Technology suppliers already use such vTech,**
  - but in a proprietary way.

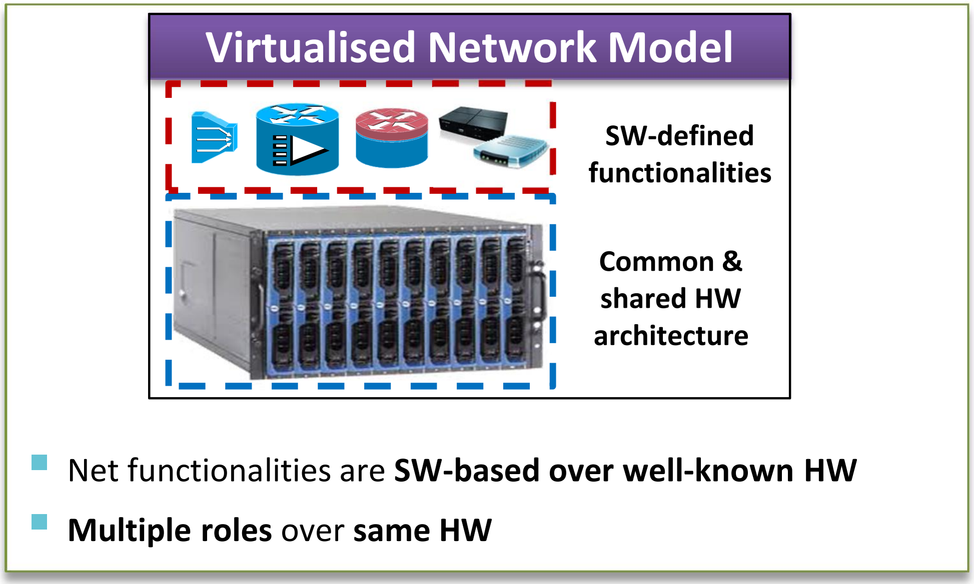

Early adopters offer virtualized versions of their products

Source: NFV

Trends

- High performance industry **standard** servers shipped in very high volume
- **Convergence** of computing, storage and networks
- New **virtualization technologies** that abstract underlying hardware yielding elasticity, scalability and automation
- **Software-defined networking**
- **Cloud** services
- **Mobility**, explosion of devices and traffic

Challenges

- Huge **capital investment** to deal with current trends
- Network operators face an increasing **disparity between costs and revenues**
- **Complexity**: large and increasing variety of proprietary hardware appliances in operator's network
- Reduced **hardware lifecycles**
- **Lack of flexibility and agility**: cannot move network resources where & when needed
- **Launching new services is difficult and takes too long**. Often requires yet another proprietary box which needs to be integrated

Source: Adapted from D. Lopez Telefonica I+D, NFV

The NFV Concept

A means to make the network more flexible and simple by minimising dependence on HW constraints

Traditional Network Model:
APPLIANCE APPROACH

- Network Functions are **based on specific HW&SW**
- **One physical node per role**

Virtualised Network Model:
VIRTUAL APPLIANCE APPROACH

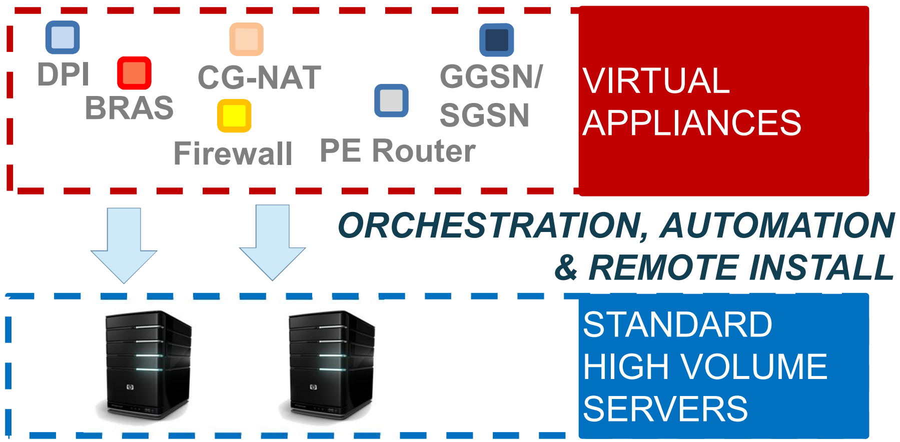

- Network Functions are **SW-based over well-known HW**
- **Multiple roles** over **same HW**

Source: Adapted from D. Lopez Telefonica I+D, NFV

Target

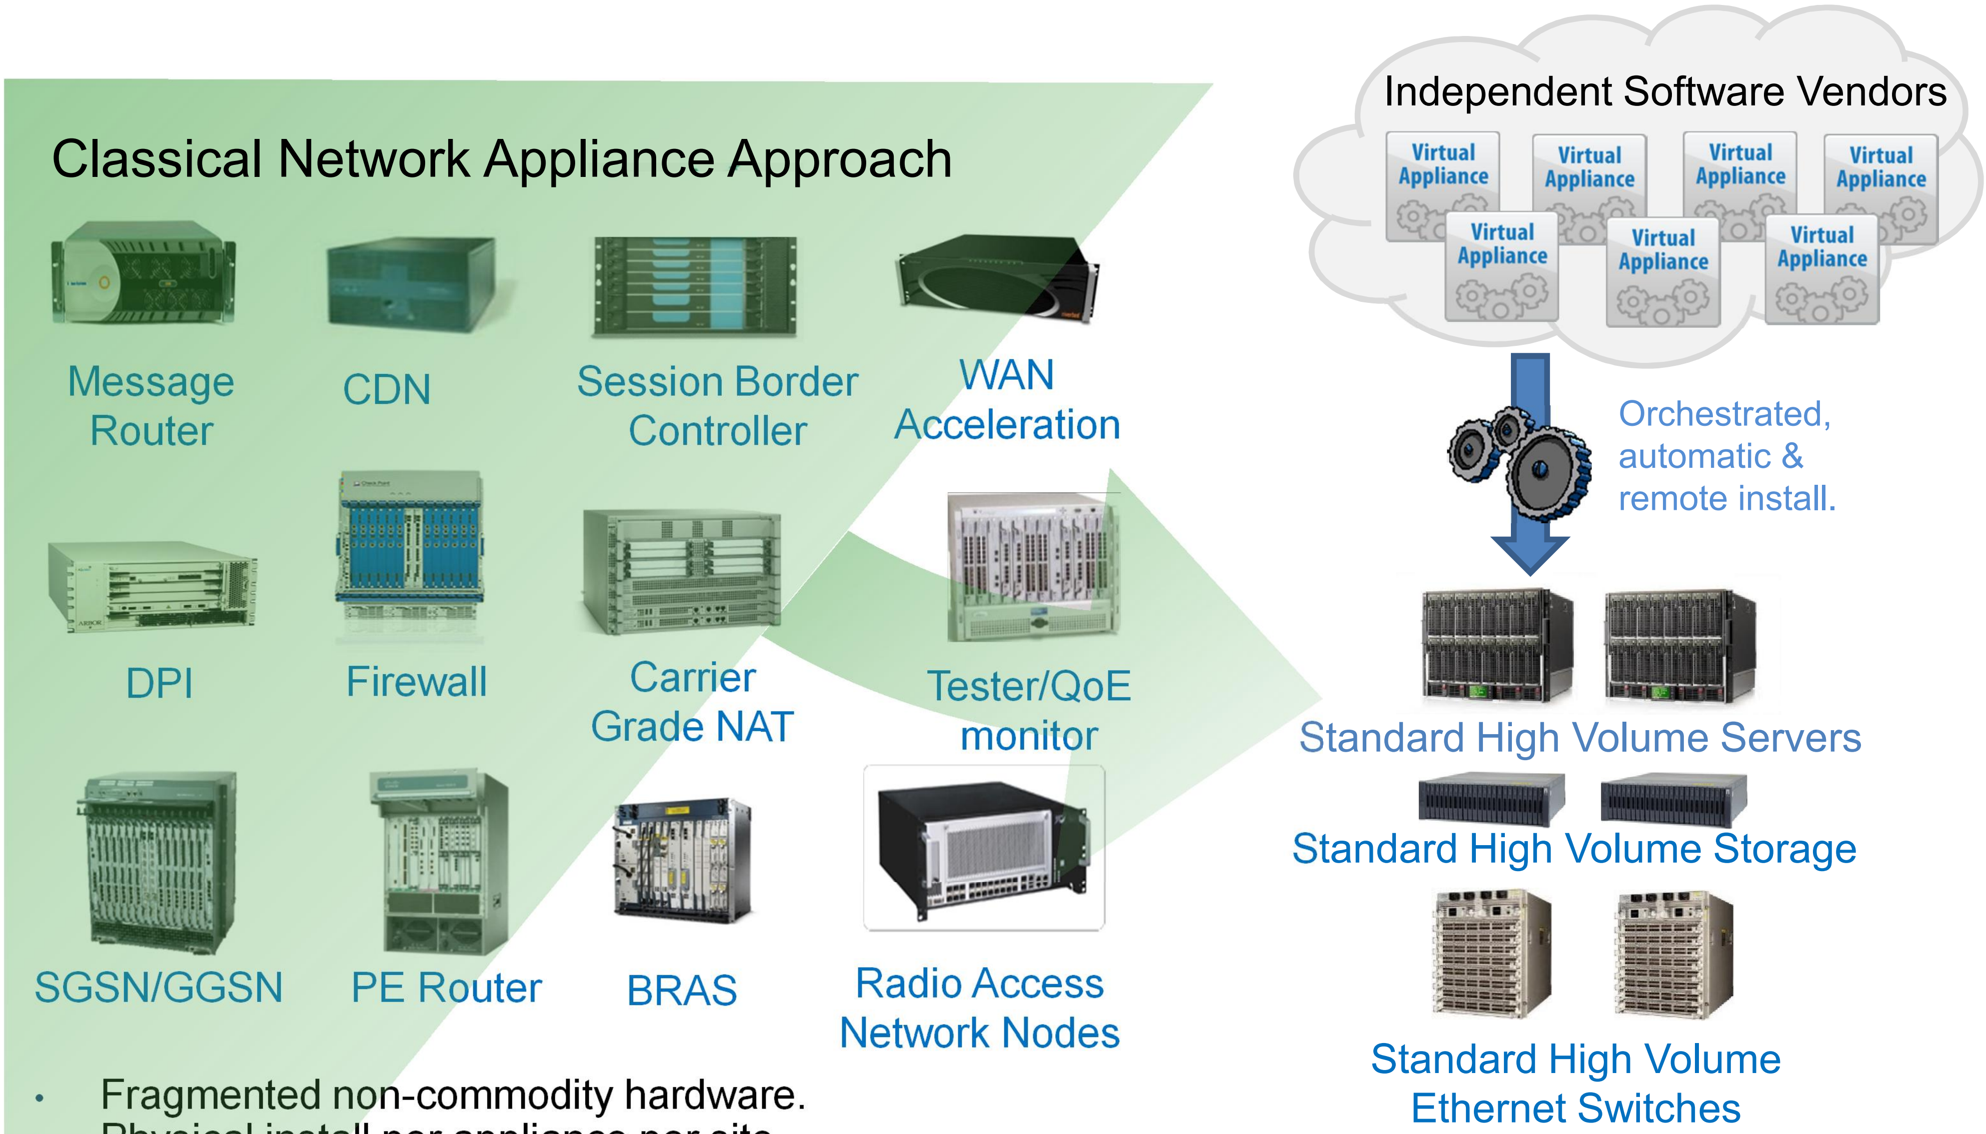

- Fragmented non-commodity hardware.
- Physical install per appliance per site.
- Hardware development large barrier to entry for new vendors, constraining innovation & competition.

Network Virtualisation Approach

Source: NFV

Network Functions Virtualization

- Network Functions Virtualization is about implementing network functions in software - that today run on proprietary hardware - leveraging (high volume) standard servers and IT virtualization
- Supports multi-versioning and multi-tenancy of network functions, which allows use of a single physical platform for different applications, users and tenants
- Enables new ways to implement resilience, service assurance, test and diagnostics and security surveillance
- Provides opportunities for pure software players
- Facilitates innovation towards new network functions and services that are only practical in a pure software network environment
- Applicable to any data plane packet processing and control plane functions, in fixed or mobile networks
- NFV will only scale if management and configuration of functions can be automated
- NFV aims to ultimately transform the way network operators architect and operate their networks, but change can be incremental

Source: Adapted from D. Lopez Telefonica I+D, NFV

Benefits & Promises of NFV

- Reduced equipment costs (CAPEX)
  - through consolidating equipment and economies of scale of 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,
  - allows a single platform for different applications, users and tenants.
- Enables a variety of eco-systems and encourages openness.
- Encouraging innovation to bring new services and generate new revenue streams.

Source: NFV

Benefits & Promises of NFV

- Flexibility to easily, rapidly, dynamically provision and instantiate new services in various locations
- Improved operational efficiency
  - by taking advantage of the higher uniformity of the physical network platform and its homogeneity to other support platforms.
- Software-oriented innovation to rapidly prototype and test new services and generate new revenue streams
- More service differentiation & customization
- Reduced (OPEX) operational costs: reduced power, reduced space, improved network monitoring
- IT-oriented skillset and talent

Source: Adapted from D. Lopez Telefonica I+D, NFV

Setting the Ground

Traditional Network model

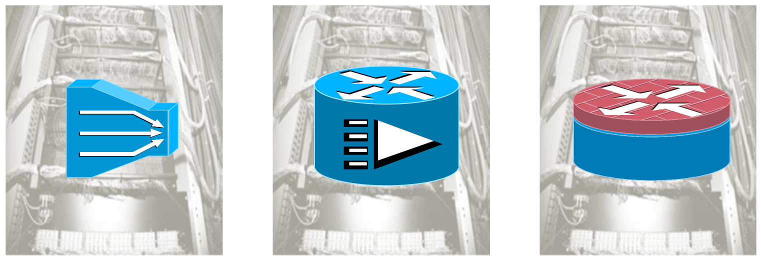

- Network functionalities are **based on specific HW&SW**
- **One physical node per role**

Virtualised Network Model

- Net functionalities are **SW-based over well-known HW**
- **Multiple roles** over **same HW**

Telco industry is optimised to work with the Telco cycle, and this cannot be changed overnight

EVOLUTIONARY APPROACH

- New specific needs for different nodes
- Reuse of equipment still in amortization
- Leverage on new planned elements in architecture

BY RELEVANT USE CASES

- Virtualised PoP
- Traffic analysis
- SW-defined datacenter
- Pervasive security

- Effect of **scale** (better statistical multiplexing)
- **Reduction** of geographical **dispersion** of HW
- **Reuse same HW** for different role when no longer applicable

Source: Telefonica I+D

So, why we need/want NFV(/SDN)?

1. Virtualization: Use network resource without worrying about where it is physically located, how much it is, how it is organized, etc.
2. Orchestration: Manage thousands of devices
3. Programmable: Should be able to change behavior on the fly.
4. Dynamic Scaling: Should be able to change size, quantity
5. Automation
6. Visibility: Monitor resources, connectivity
7. Performance: Optimize network device utilization
8. Multi-tenancy
9. Service Integration
10. Openness: Full choice of modular plug-ins

Note: These are exactly the same reasons why we need/want SDN.

Source: Adapted from Raj Jain

NFV and SDN

- NFV and SDN are highly complementary
- Both topics are mutually beneficial but not dependent on each other

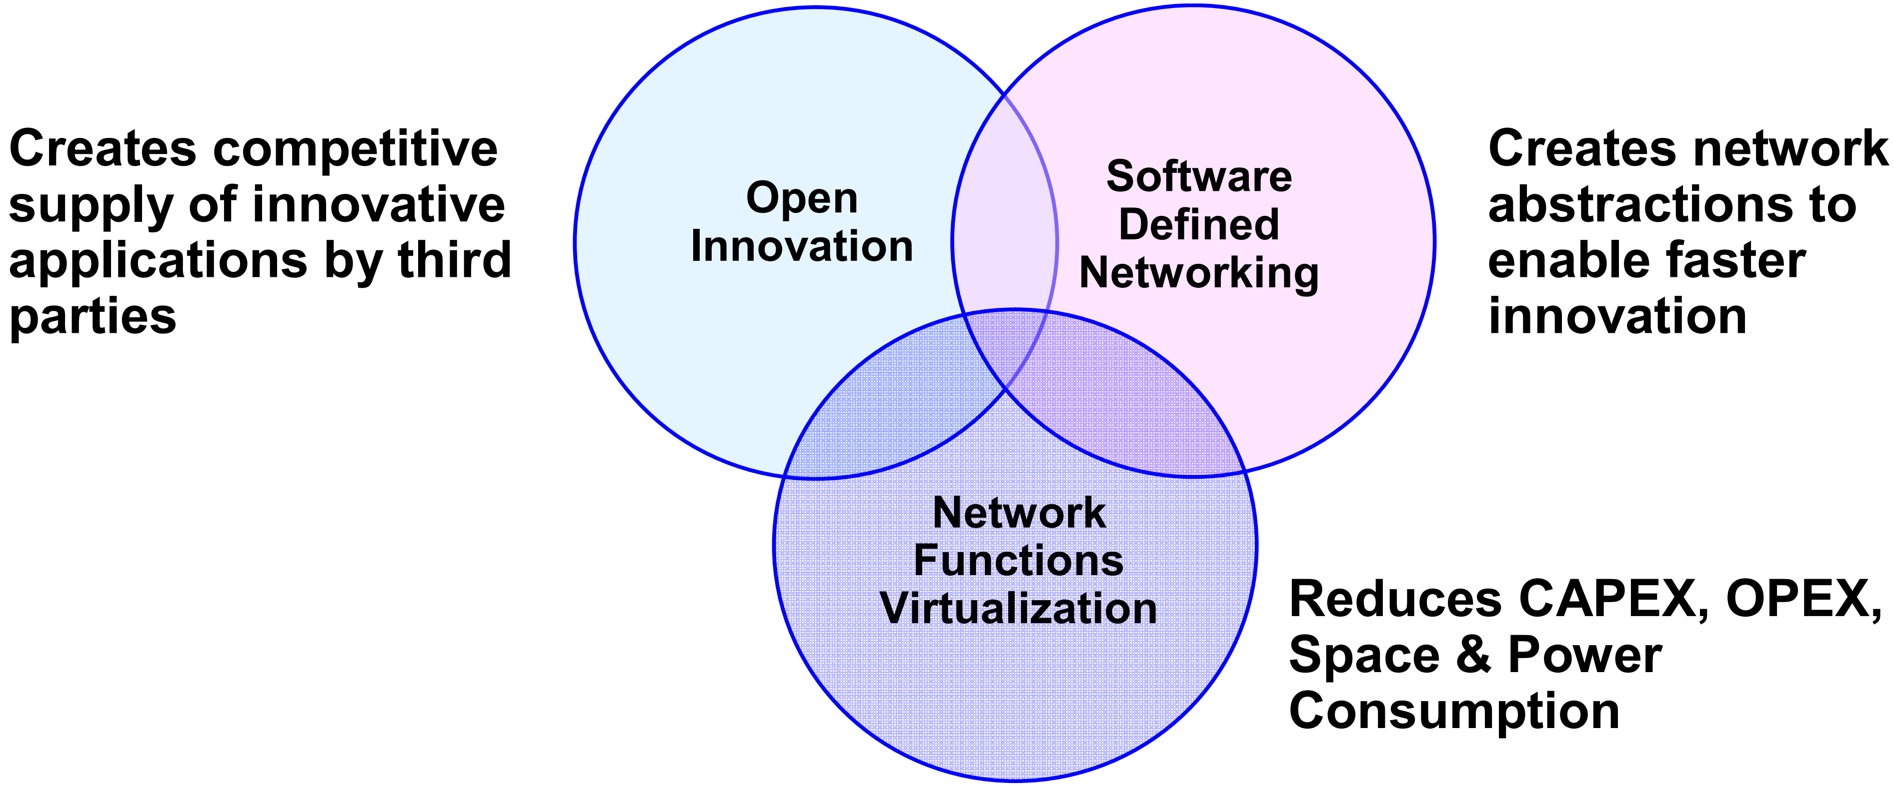

Source: NFV

NFV vs SDN

- NFV: re-definition of **network equipment architecture**
- NFV was born to meet Service Provider (SP) needs:
  - Lower CAPEX by reducing/eliminating proprietary hardware
  - Consolidate multiple network functions onto industry standard platforms
- SDN: re-definition of **network architecture**
- SDN comes from the IT world:
  - Separate the data and control layers, while centralizing the control
  - Deliver the ability to program network behavior using well-defined interfaces

Software Defined Networking

Network equipment as Black boxes

Open interfaces (OpenFlow) for instructing the boxes what to do

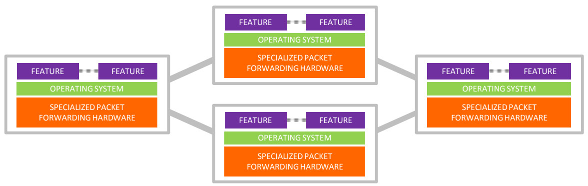

Boxes with autonomous behaviour

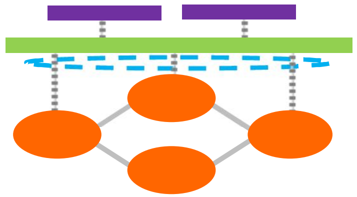

Decisions are taken out of the box

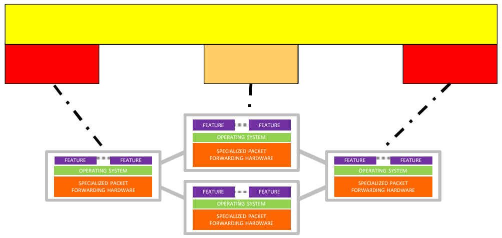

Adapting OSS to manage black boxes

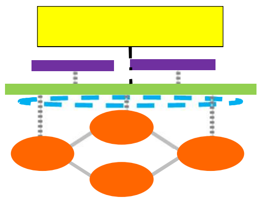

Simpler OSS to manage the SDN controller

Source: Adapted from D. Lopez Telefonica I+D, NFV

Scope of NFV and OpenFlow/SDN

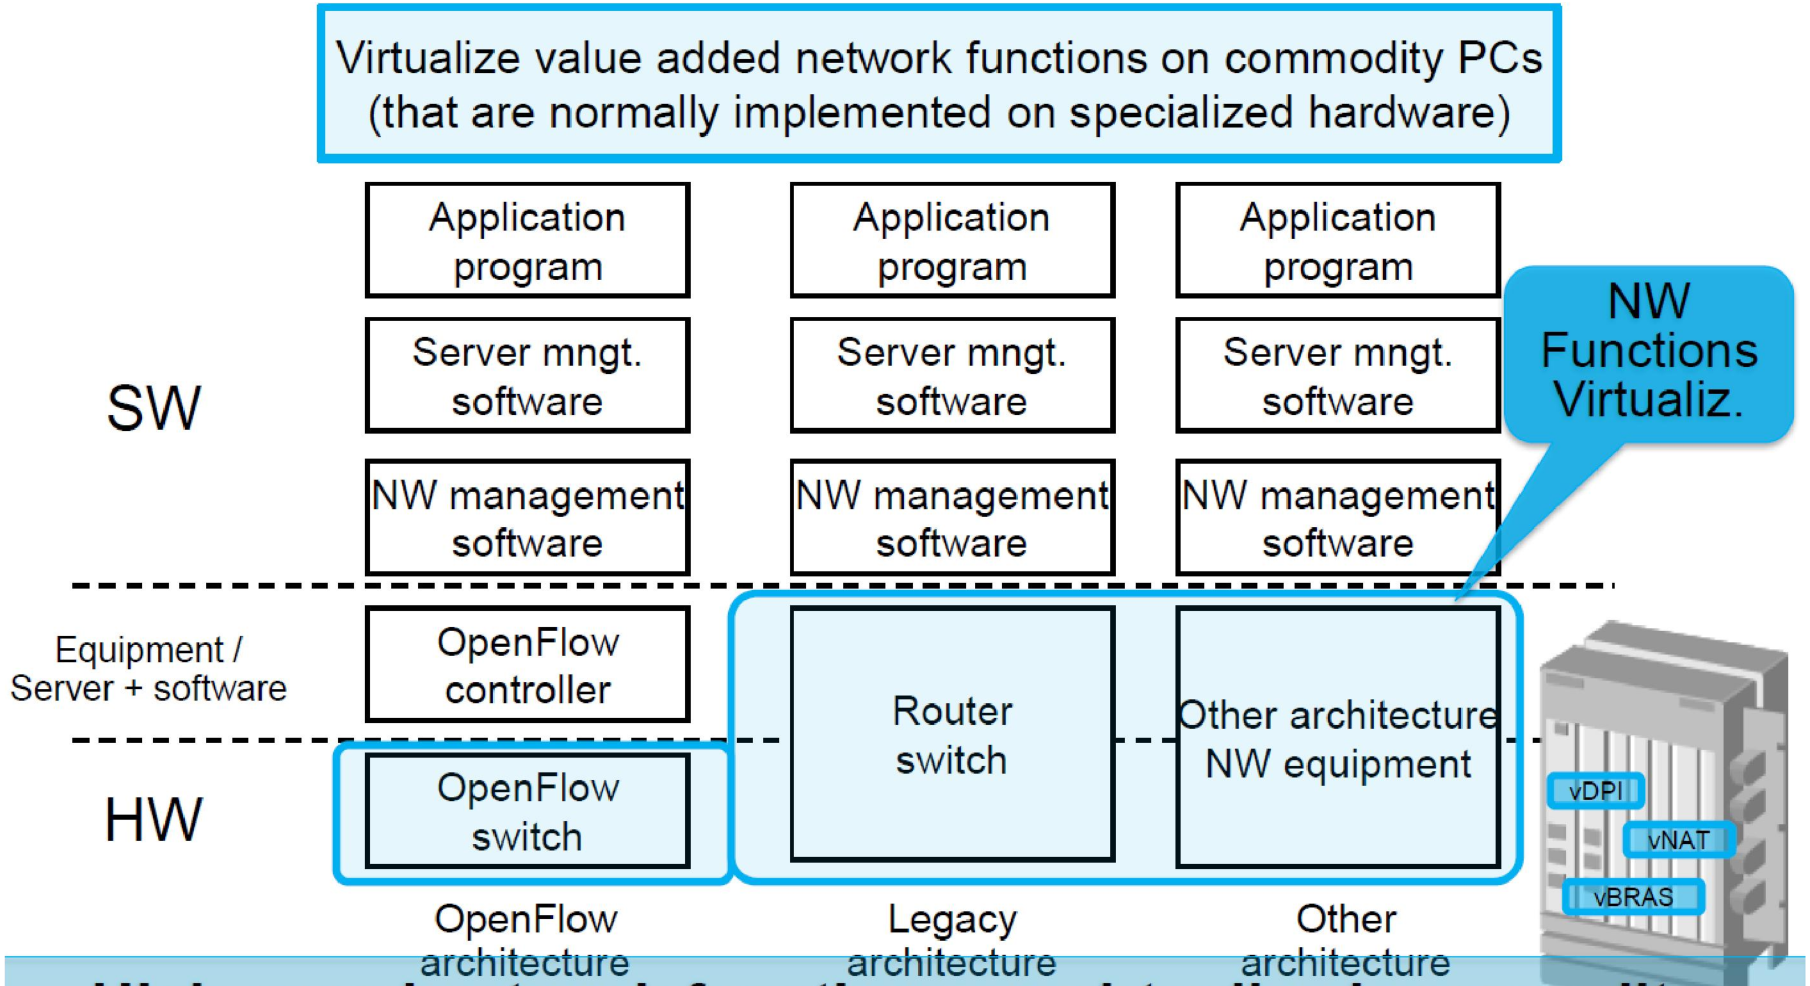

High speed network functions on virtualized commodity hardware: optimized data plane under virtualization

Source: NEC

Networking with SDN & NFV

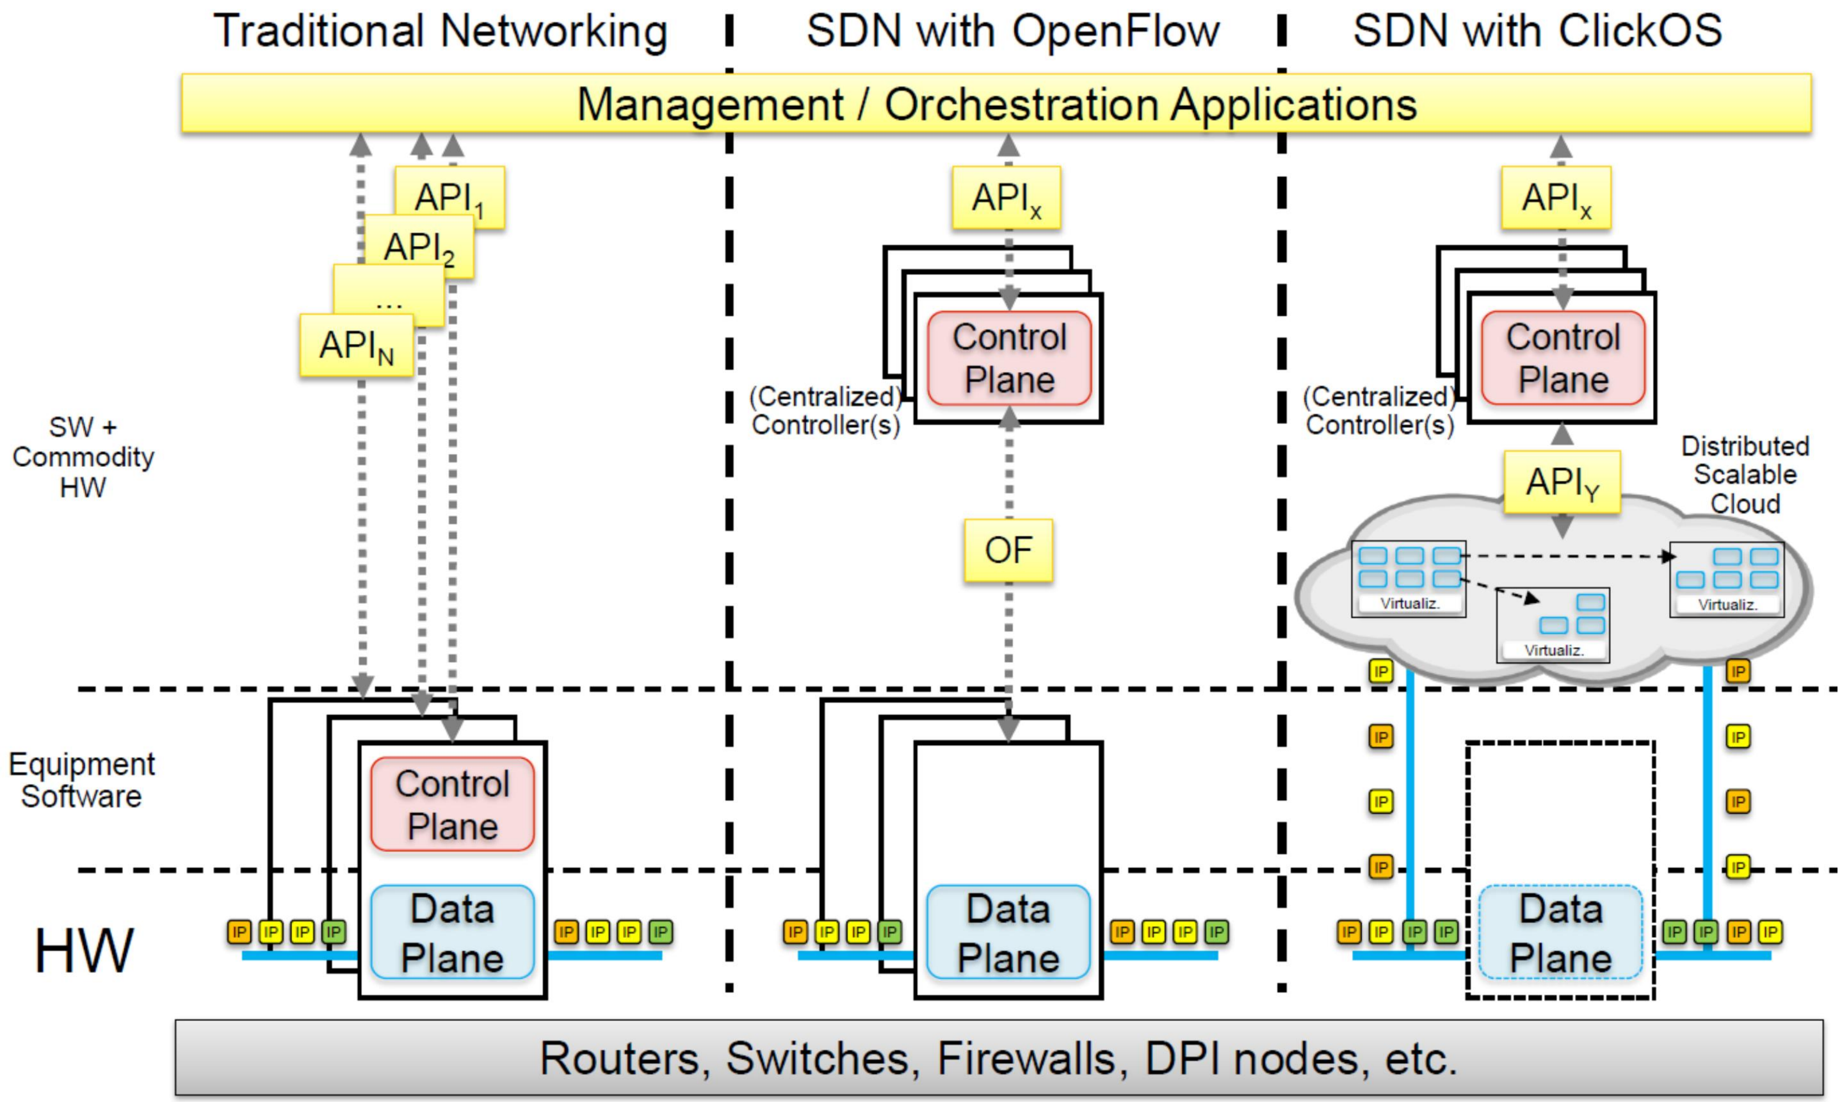

Source: NEC

(Network Virtualization)^2 = SDN + NFV

SDN: Software Defined Networking

NFV: Network Functions Virtualisation

Source: Adapted from D. Lopez Telefonica I+D, NFV

Some Drivers

Virtual CPE

Complex home environment

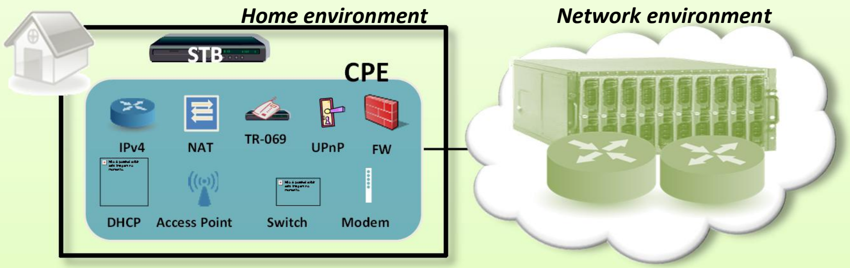

Home simplification

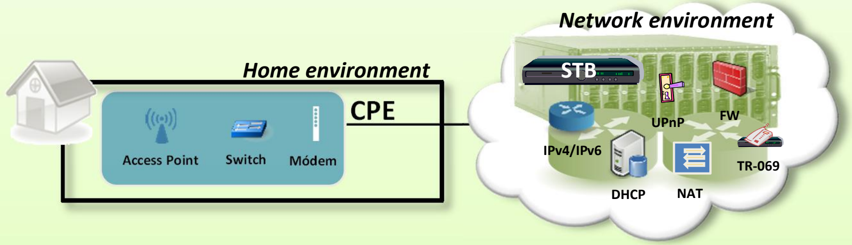

- Simplification or even supression (STB)
- No need for home router replacement as it is updated by configuration
- Fast deployment for new services
- Inexpensive IPv6 migration maintaining legacy home routers

Virtual IP Edge

Multiple IP Edges

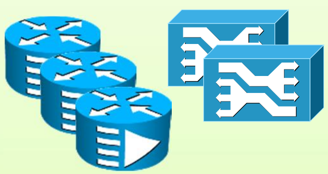

- An IP Edge for each service (voice, video content, Internet)
- Scattered and not well integrated control functions (e.g. DPI, BRAS, PCRF)

A unified software IP Edge

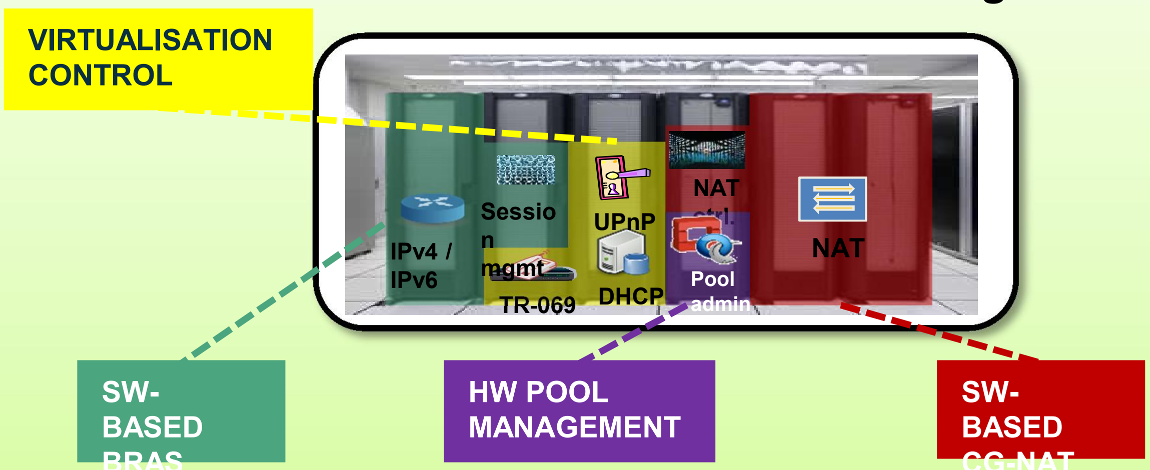

Source: Adapted from D. Lopez Telefonica I+D, NFV

…More Drivers…

Mobile Network Virtualisation

- All the network concentrated in the base station

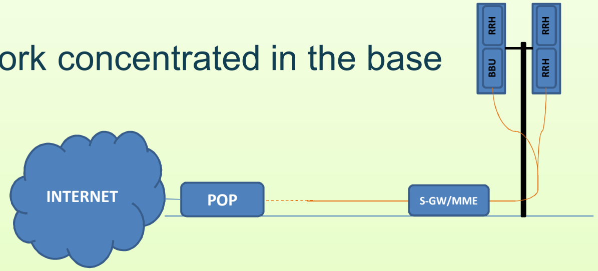

- C-RAN: All the base station functionalities, except for the antennas and power amplifiers, concentrated in a centralized location

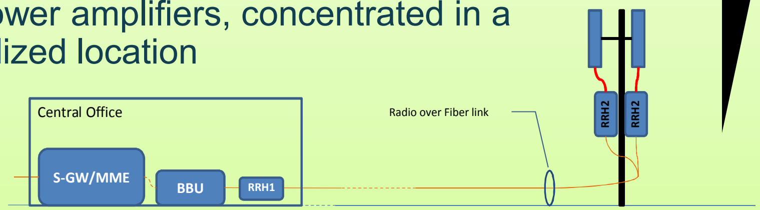

Having the flexibility of **moving functionalities between different locations** may help to network to adopt the best option in each case

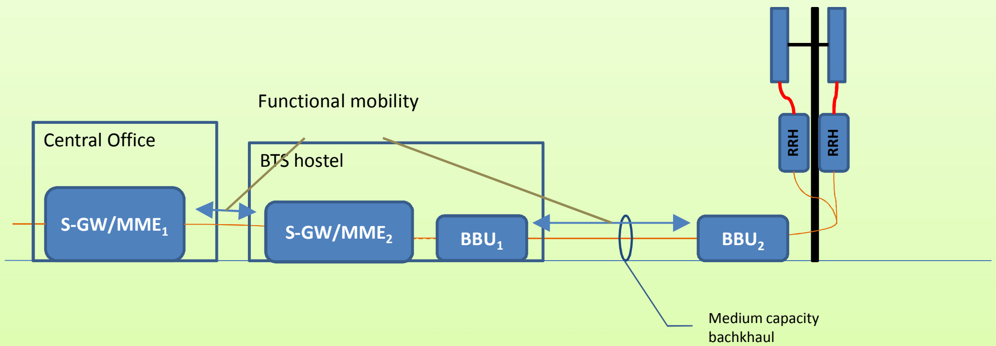

Monitoring/enforcement loop

Current DPI

Everything replicated in 100s of boxes which need to be orchestrated!

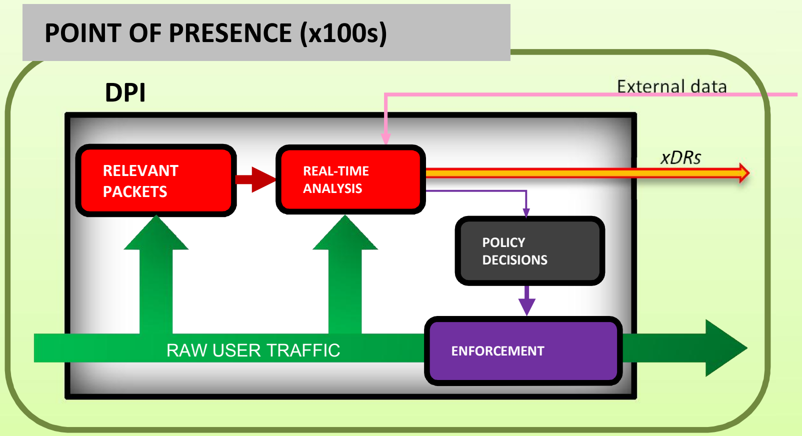

Virtual DPI

Centralised intelligence & orchestration

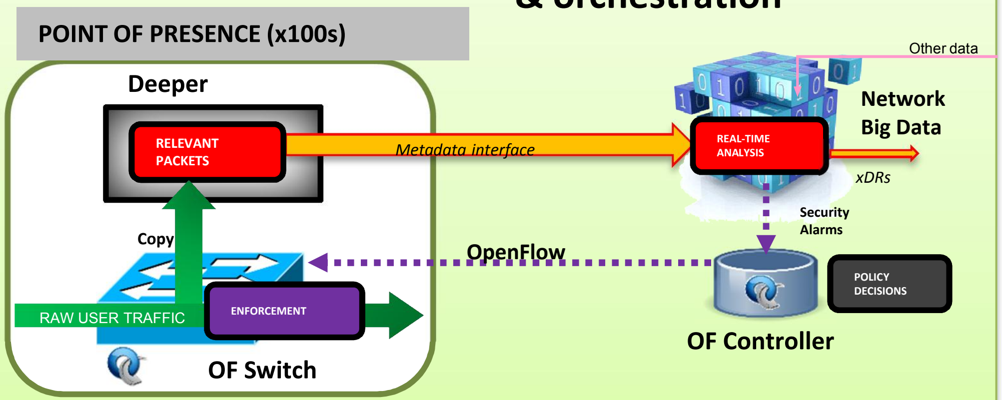

Source: Adapted from D. Lopez Telefonica I+D, NFV

… And a Couple More

Virtualized CGNAT

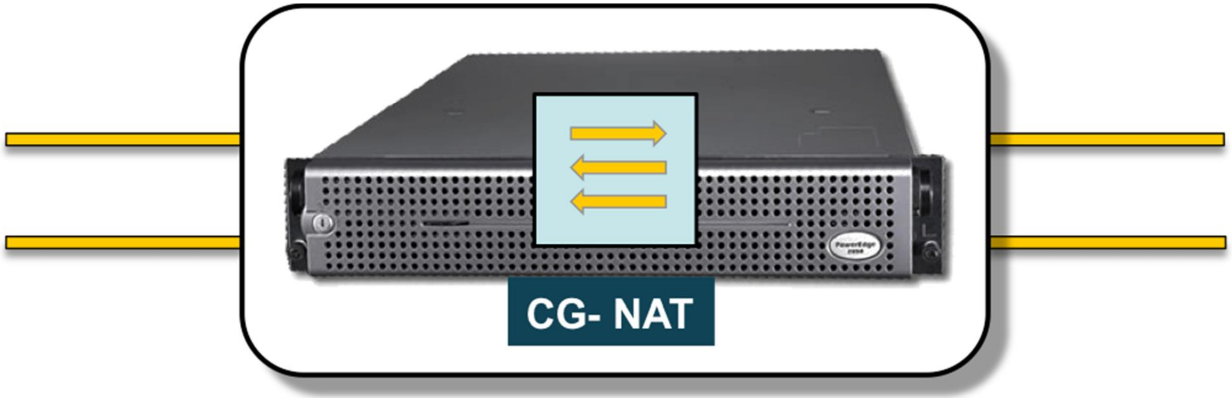

- NAT44 function, extensible to IPv6 transition
- 40 Gbps full-duplex line rate per server
- Support of overlapping addresses and tunnelling
- Auto-provisioning of NAT sessions per access line

Optimized Quagga data plane

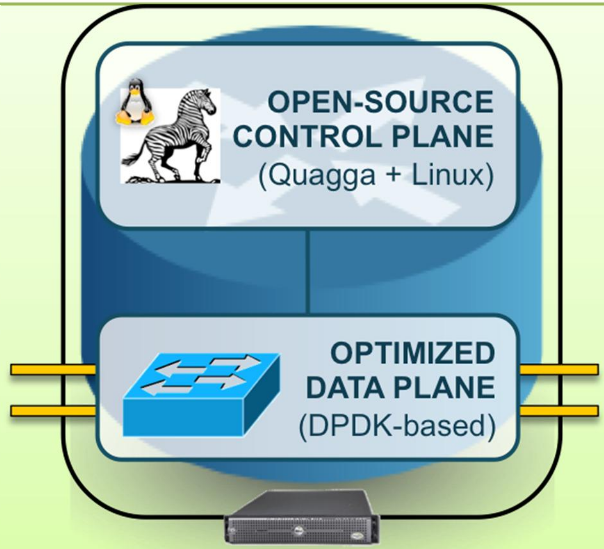

- Leverage on open source routing project as rich and widely tested protocol suite while assuring data plane performance
  - Common routing protocols supported and extended by open source project
  - High-performance line-rate data plane
  - Running in separate process, does not lead to licensing issues

Source: Adapted from D. Lopez Telefonica I+D, NFV

Some Use Case Examples

…not in any particular order

- Switching elements: BNG, CG-NAT, routers.
- Mobile network nodes: HLR/HSS, MME, SGSN, GGSN/PDN-GW.
- Home networks: Functions contained in home routers and set top boxes to create virtualised home environments.
- Tunnelling gateway elements: IPSec/SSL VPN gateways.
- Traffic analysis: DPI, QoE measurement.
- Service Assurance: SLA monitoring, Test and Diagnostics.
- NGN signalling: SBCs, IMS.
- Converged and network-wide functions: AAA servers, policy control and charging platforms.
- Application-level optimisation: CDNs, Cache Servers, Load Balancers, Application Accelerators.
- Security functions: Firewalls, virus scanners, intrusion detection systems, spam protection.

Source: NFV

History of NFV

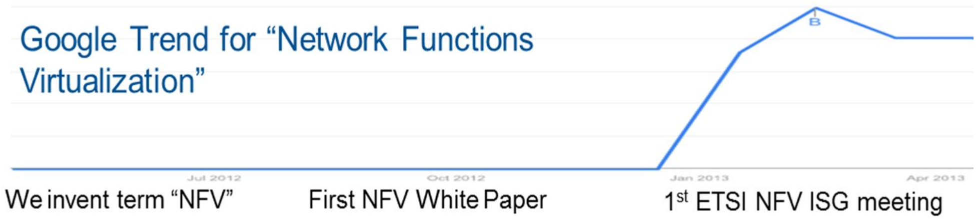

- Network operators had independently discovered that NFV technology now has sufficient performance for real-world network work loads
- Informal discussions on cooperation to encourage industry progress began at ONS in San Francisco in April 2012
- At an operator meeting in Paris in June 2012 we coined the new term "Network Functions Virtualisation (NFV)".
- We decided to convene a new industry forum, and publish a joint white paper to galvanise the industry
- At a meeting in San Francisco in September 2012 we decided to parent the new forum under ETSI
- In October 2012 we published the first joint-operator NFV white paper as a "call to action".
- This paper is widely regarded as the seminal paper heralding this new approach for networks.
- The first NFV ISG plenary session was held in January 2013
- In October 2013 the first NFV ISG documents were released after only 10 months, and a second joint-carrier NFV white paper published to provide our perspectives on progress.

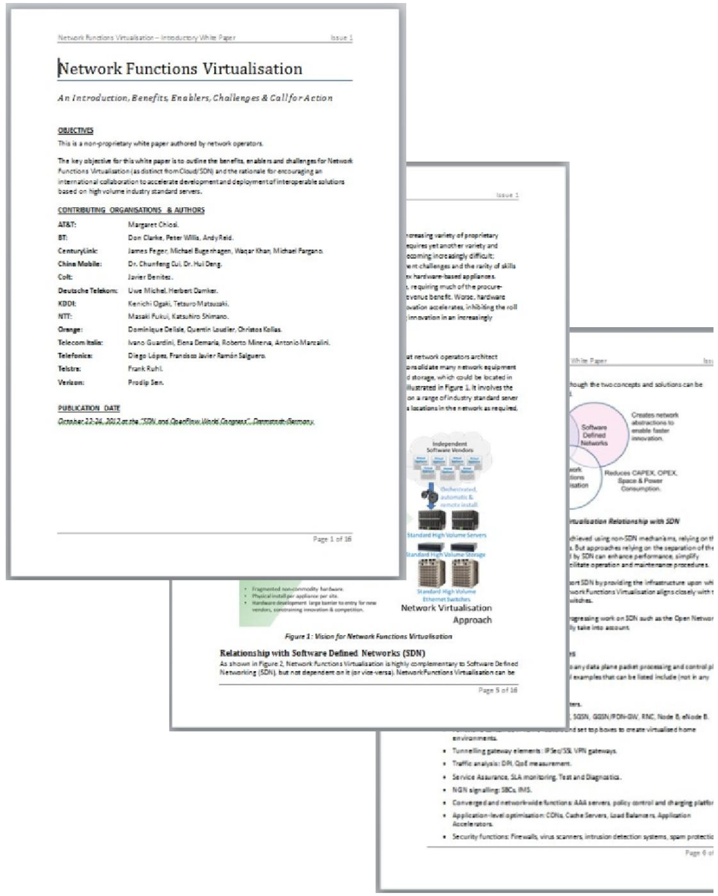

Source: Adapted from D. Lopez Telefonica I+D, NFV

The ETSI NFV ISG

- Global operators-led Industry Specification Group (ISG) under the auspices of ETSI
  - ~150 member organisations
- Open membership
  - ETSI members sign the “Member Agreement”
  - Non-ETSI members sign the “Participant Agreement”
  - Opening up to academia
- Operates by consensus
  - Formal voting only when required
- Deliverables: White papers addressing challenges and operator requirements, as input to SDOs
  - Not a standardisation body by itself

- Currently, four WGs and two EGs
  - Infrastructure
  - Software Architecture
  - Management & Orchestration
  - Reliability & Availability
  - Performance & Portability
  - Security

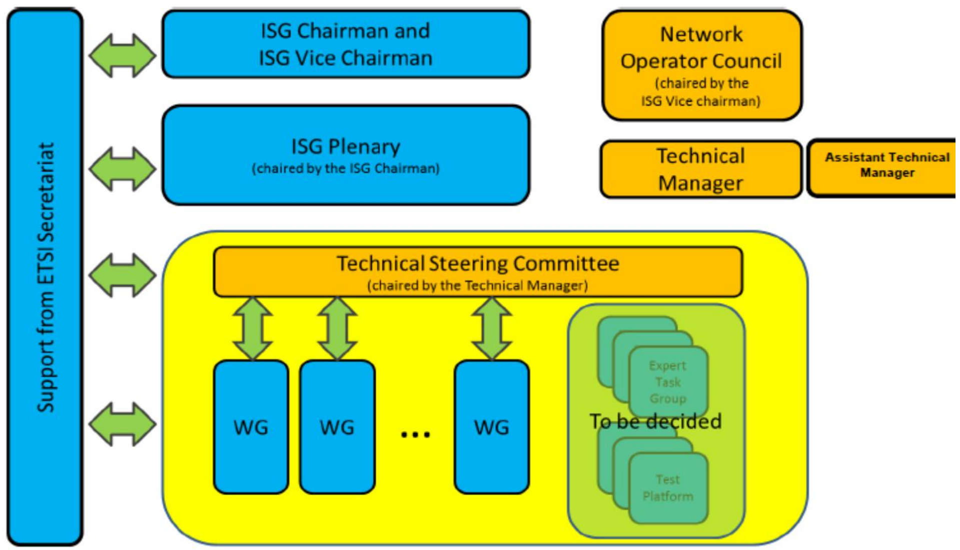

Source: Adapted from D. Lopez Telefonica I+D, NFV

ISG Working Group Structure

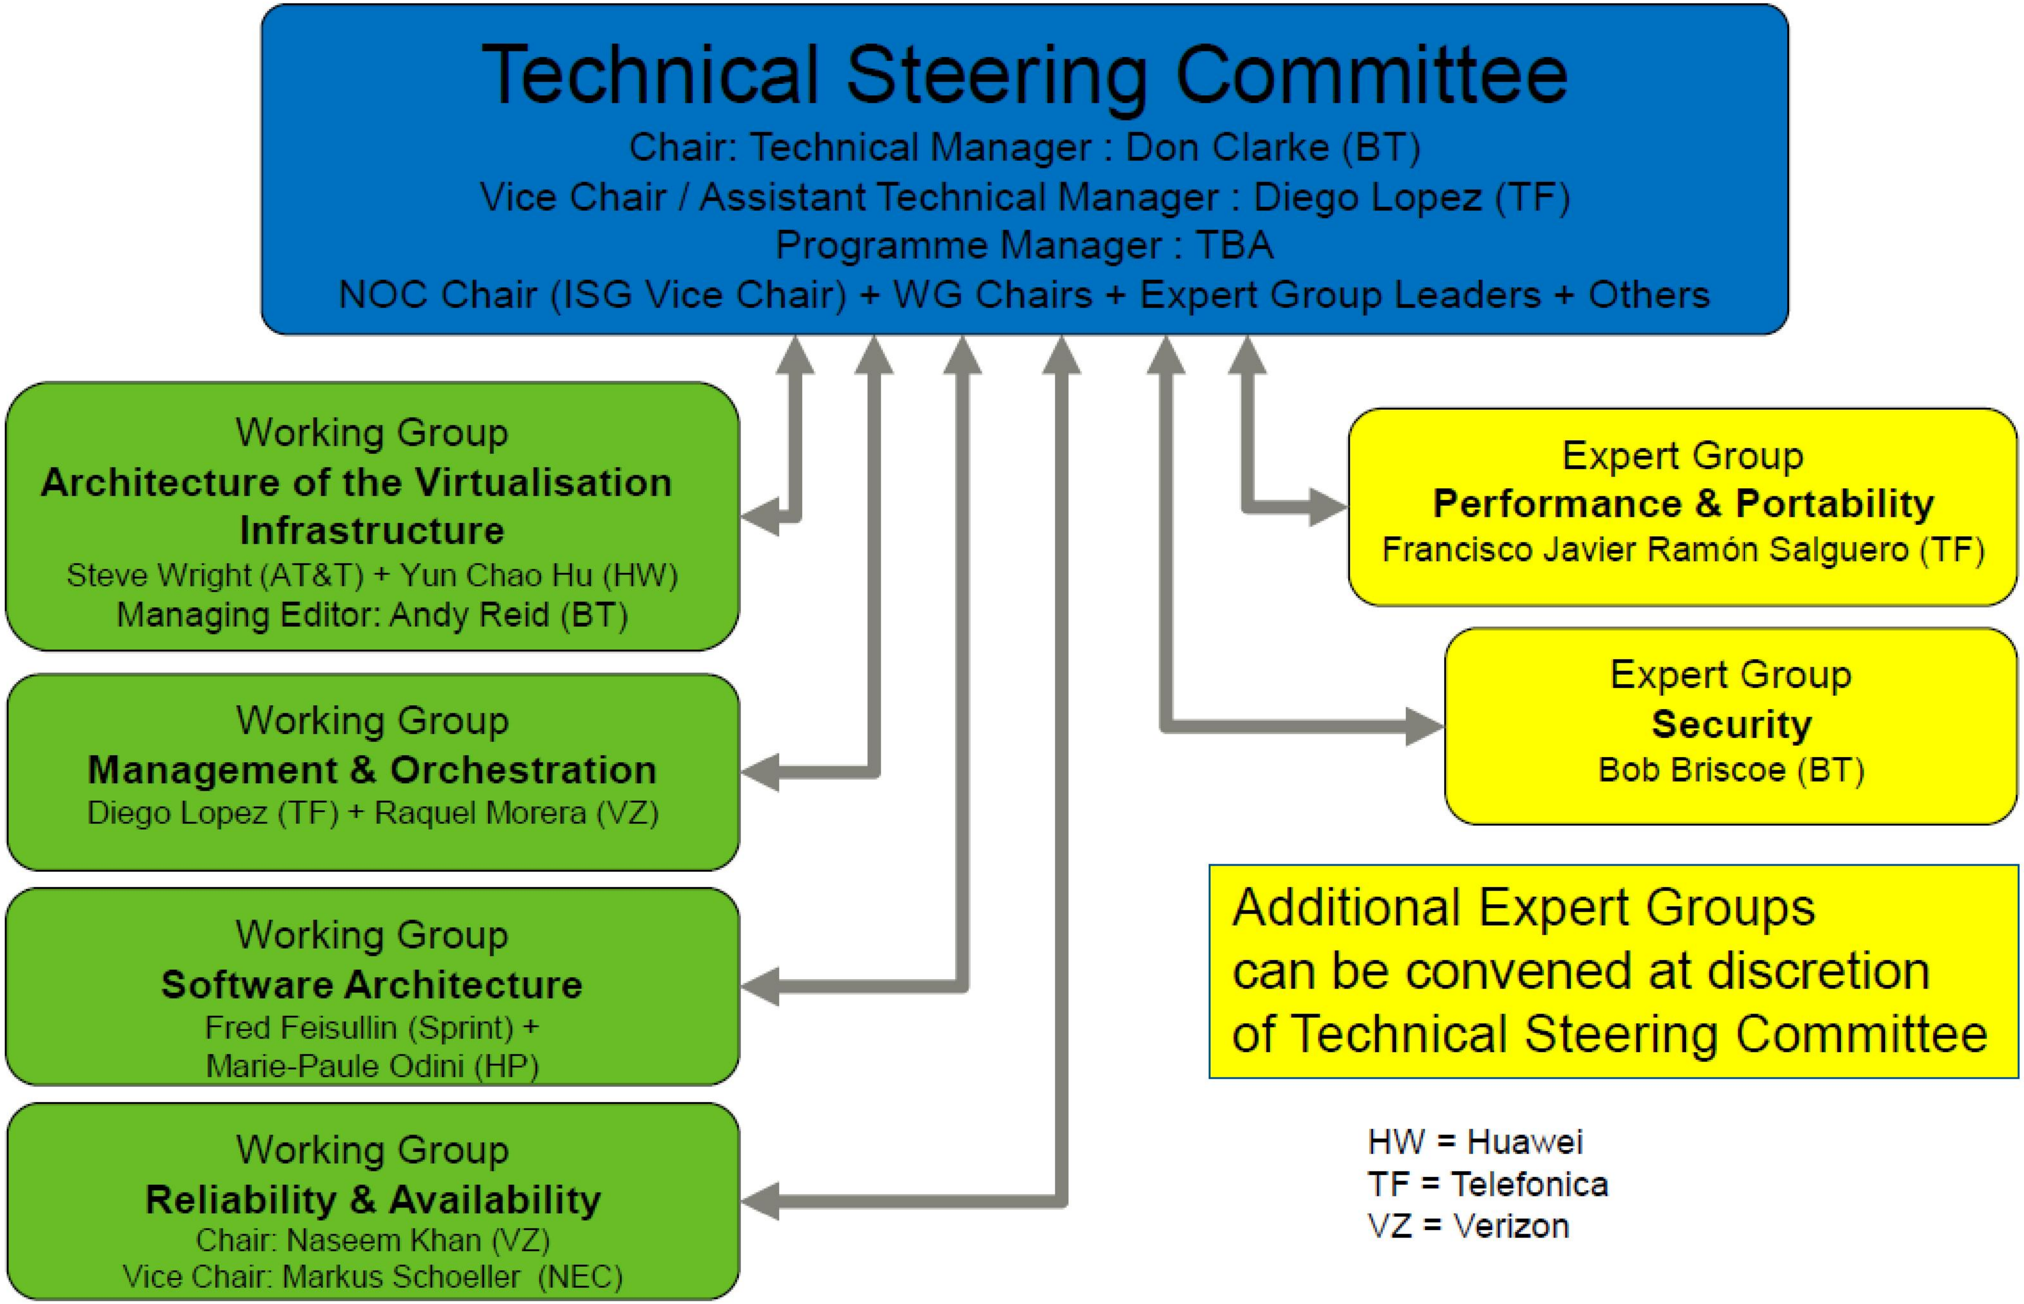

Architectural Working Groups

- Related to functional requirements
- Have a clear location in the NFV architecture
  - Keep consistency with both requirements and architecture

- INF: Supporting infrastructure interfaces and elements
- MANO: External interfaces and behaviour of a VNF
- SWA: Internals of a VNF

- Refining the architecture
- Addressing use cases
- Mostly oriented to produce reference documents

Source: Adapted from D. Lopez Telefonica I+D, NFV

Transversal Working and Expert Groups

- Related to non-functional requirements
- Transversal to the architecture
  - And influencing the architectural groups

- PER: Predictability in the data plane and function portability
- REL: Specify resiliency requirements, mechanisms , and architectures
- SEC: Function by function and infrastructure

- Refining the requirements
- Assessing use cases
- Mostly concerned with recommendations and arch models

Source: Adapted from D. Lopez Telefonica I+D, NFV

ETSI NFV External Consolidation

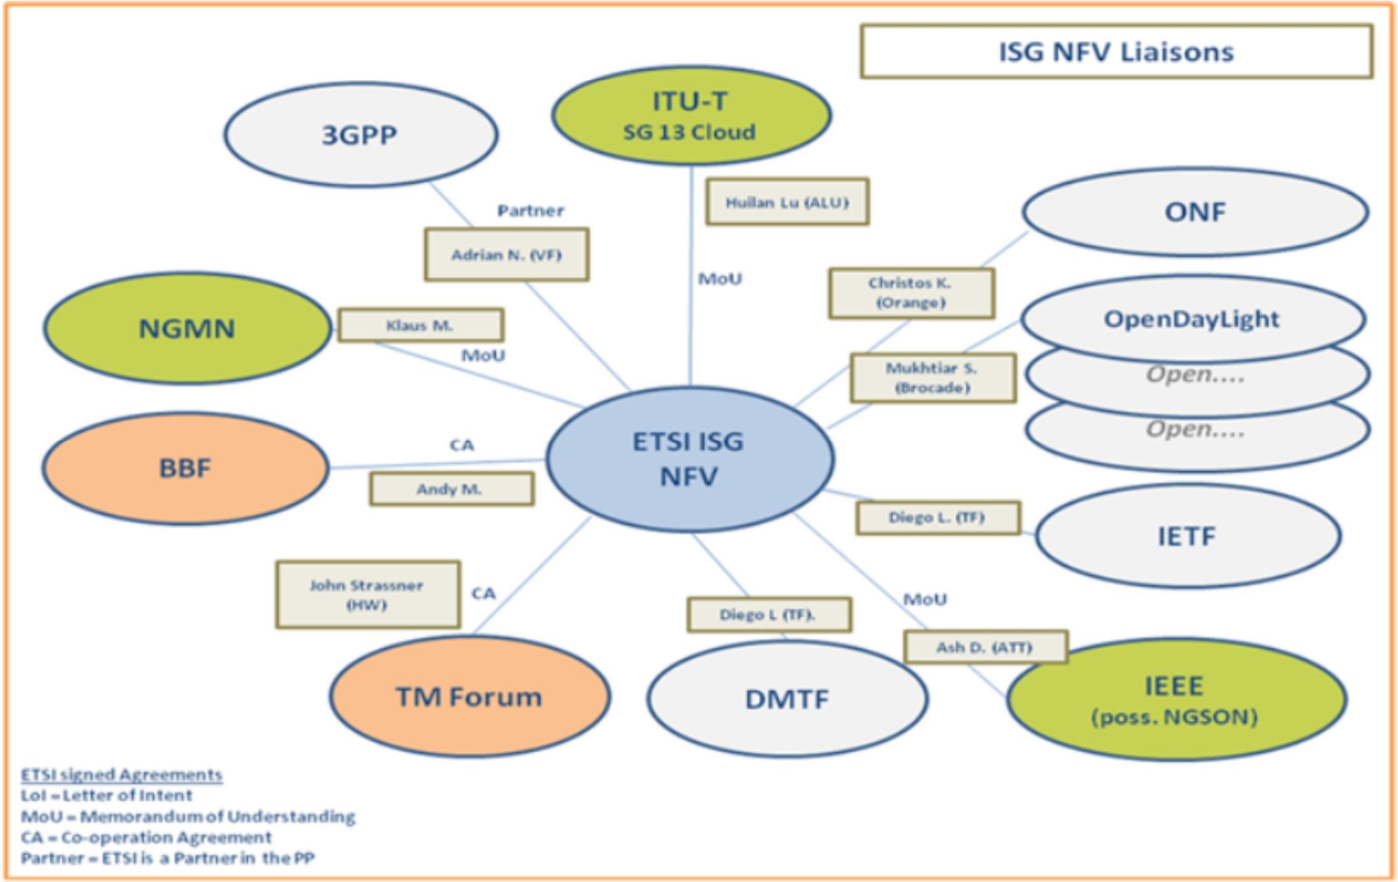

- Most relevant SDOs
- Open Source projects
- Identifying concrete areas of cooperation
- Fruitful results already achieved

- Public documents
- Early access to stable drafts
- Participation in joint events
- Coordinated individual contributions

High-level Architecture

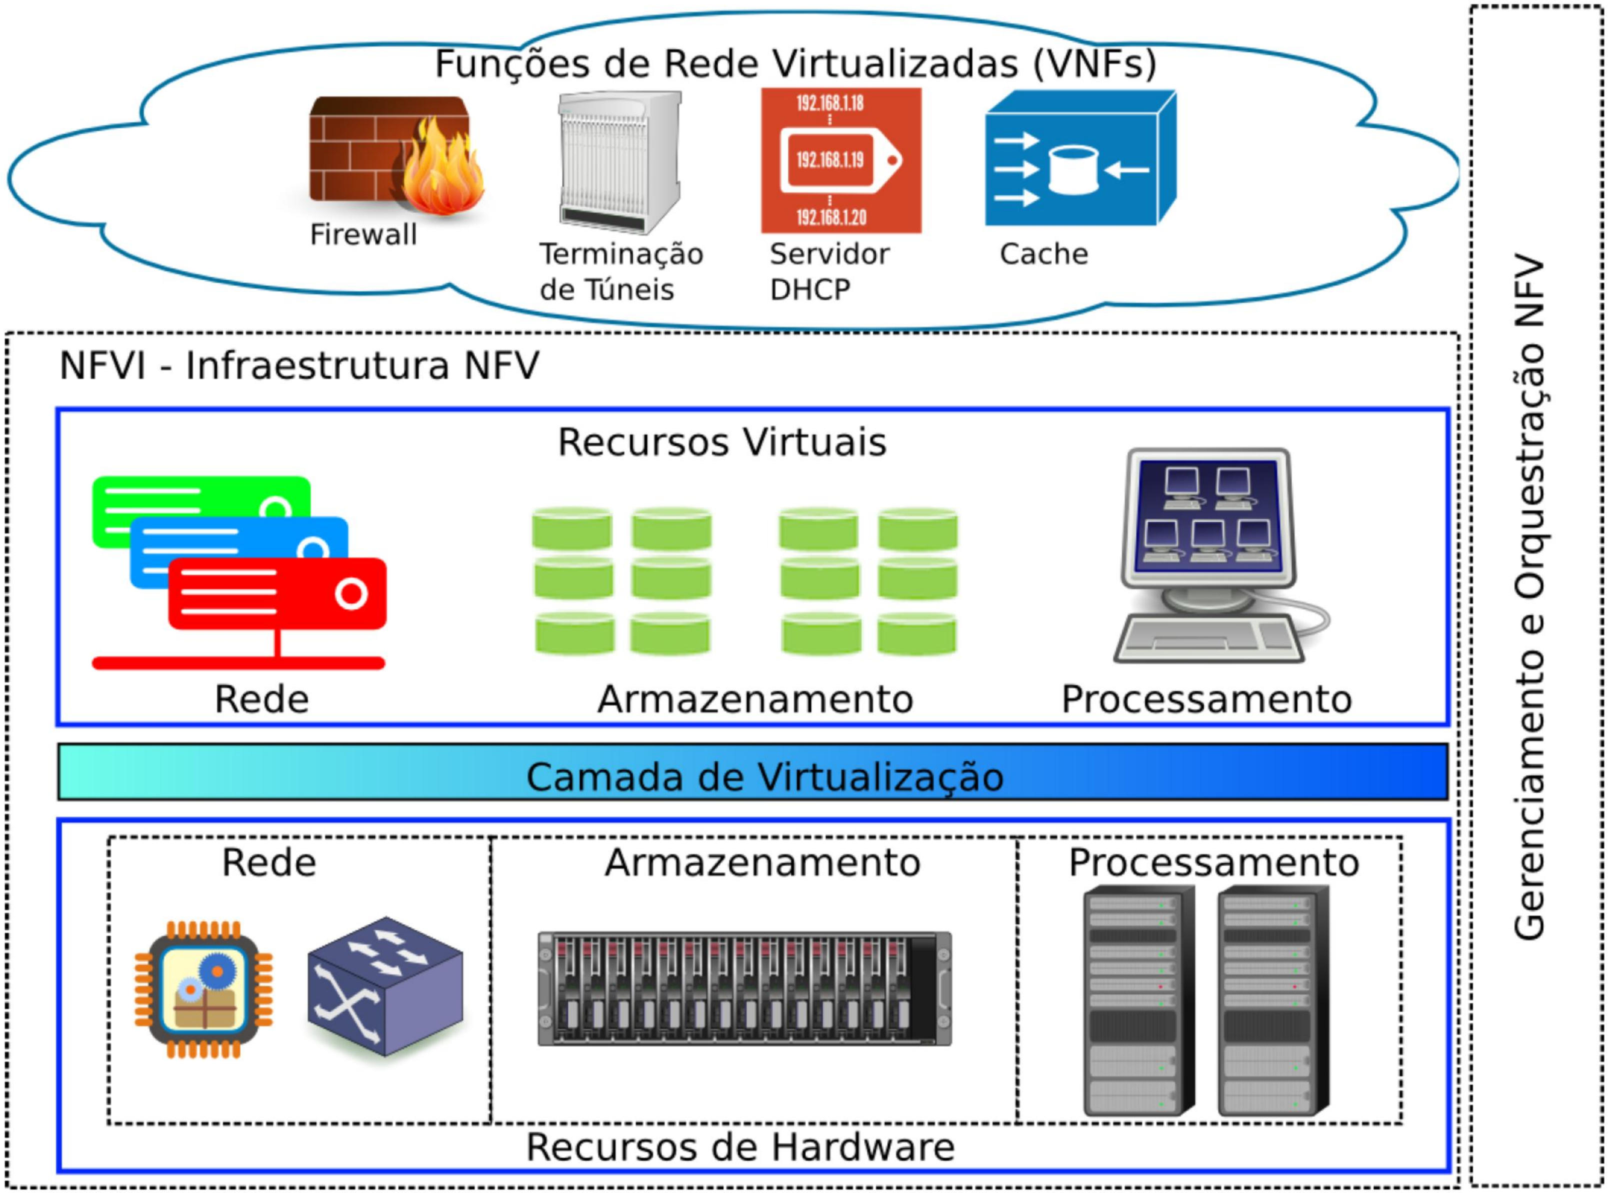

NFV Layers

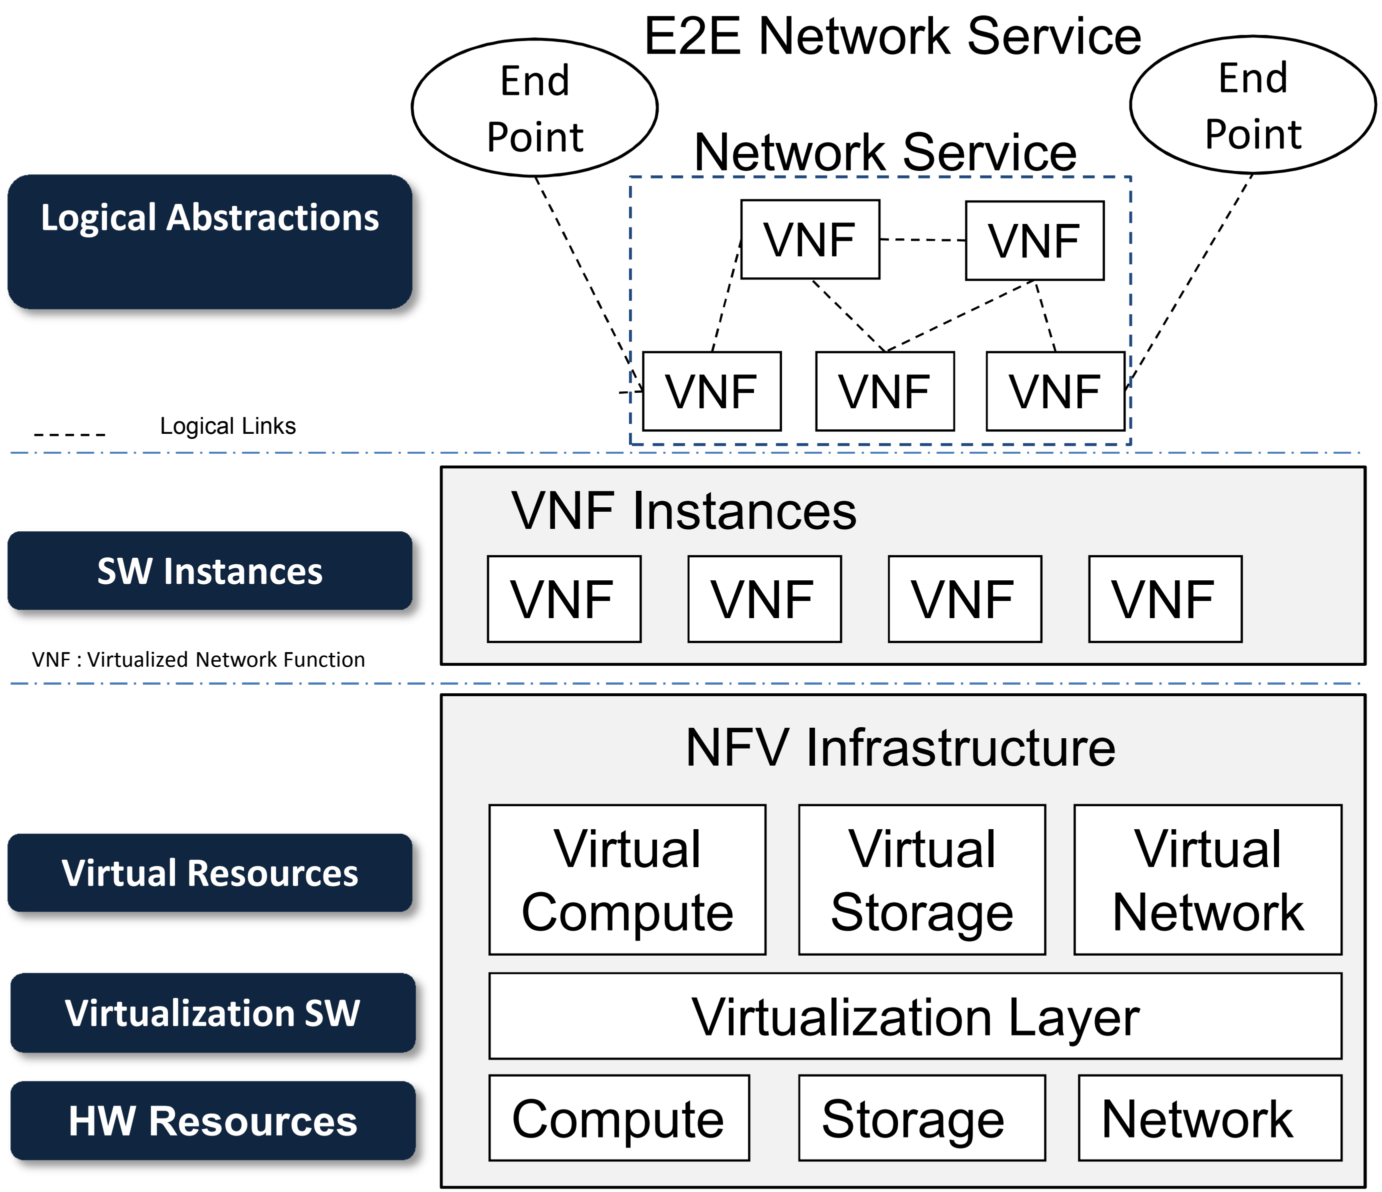

Source: Adapted from D. Lopez Telefonica I+D, NFV

Rethinking relayering

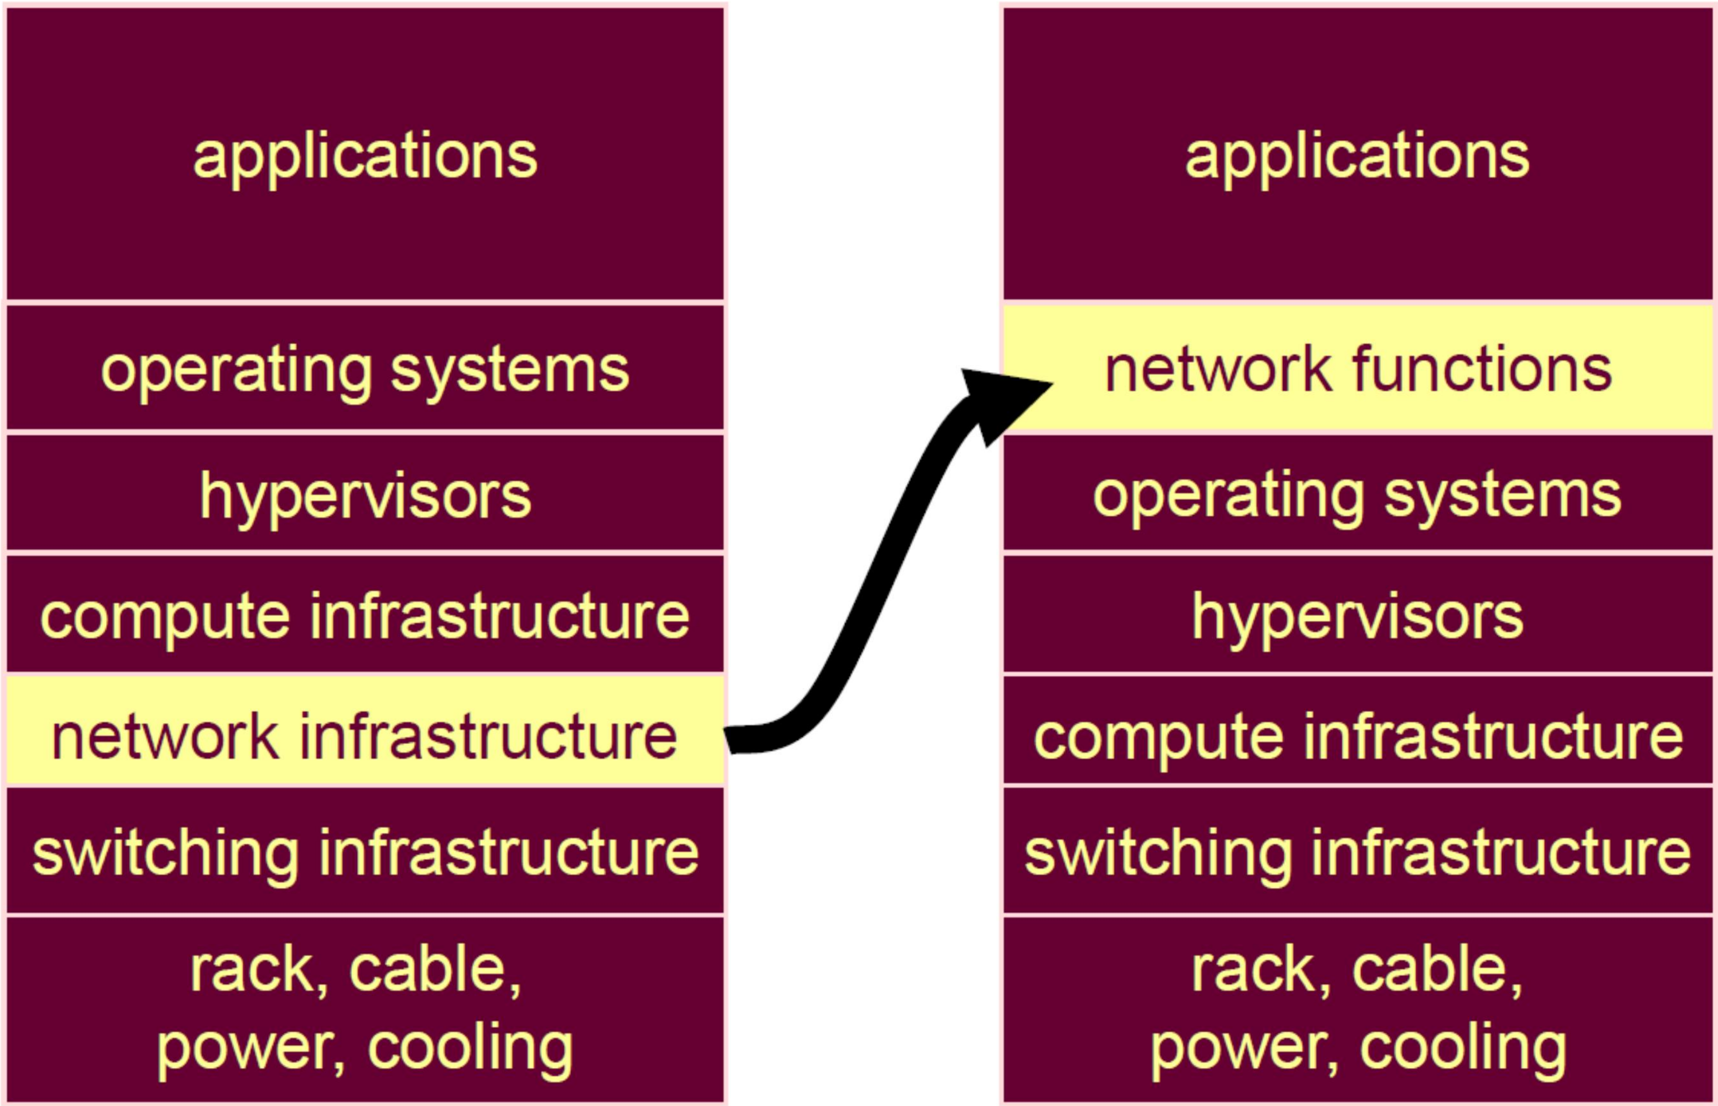

Domains and Implementations

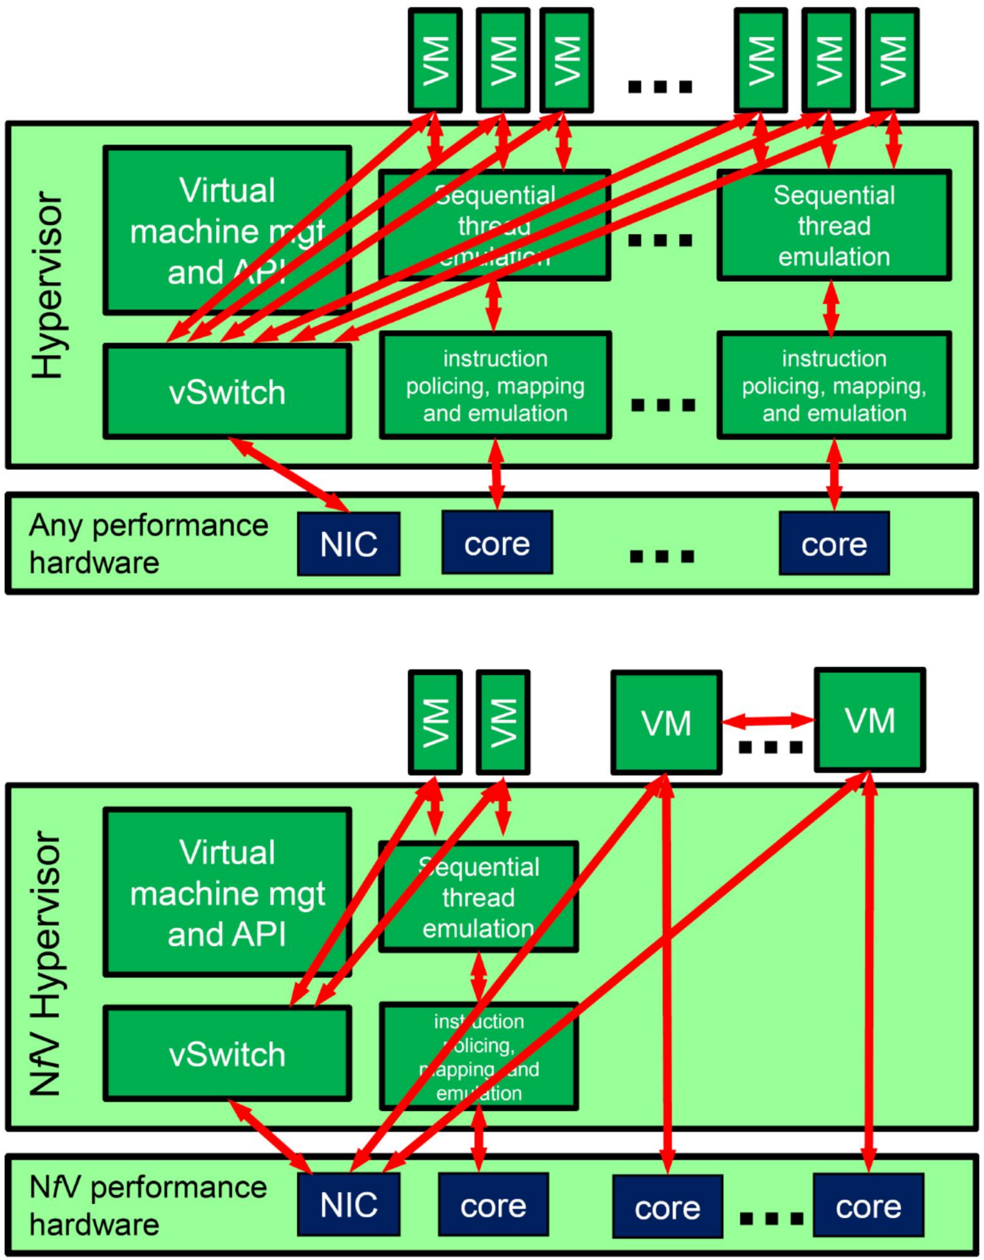

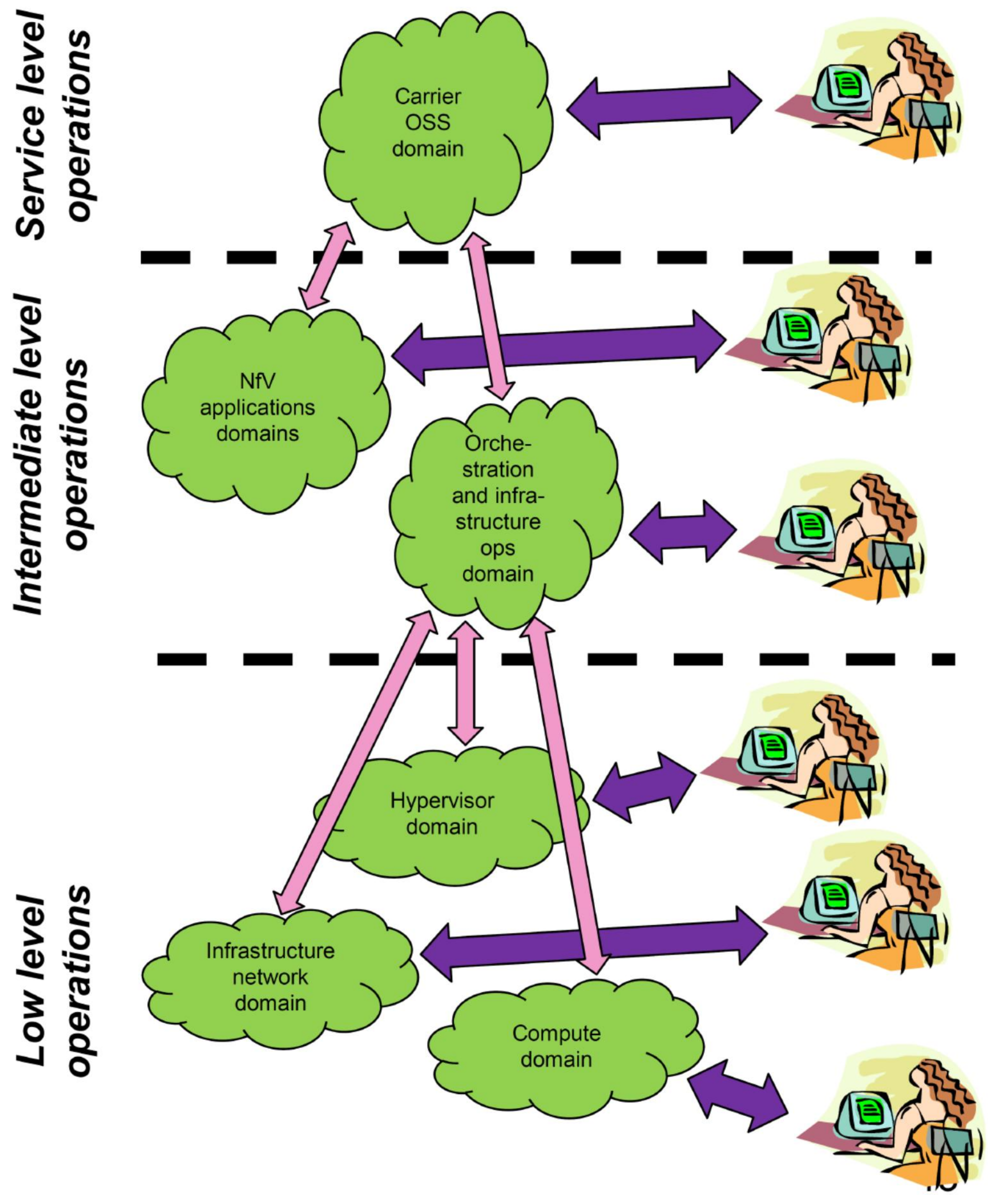

Alternative options to virtualize NFV apps

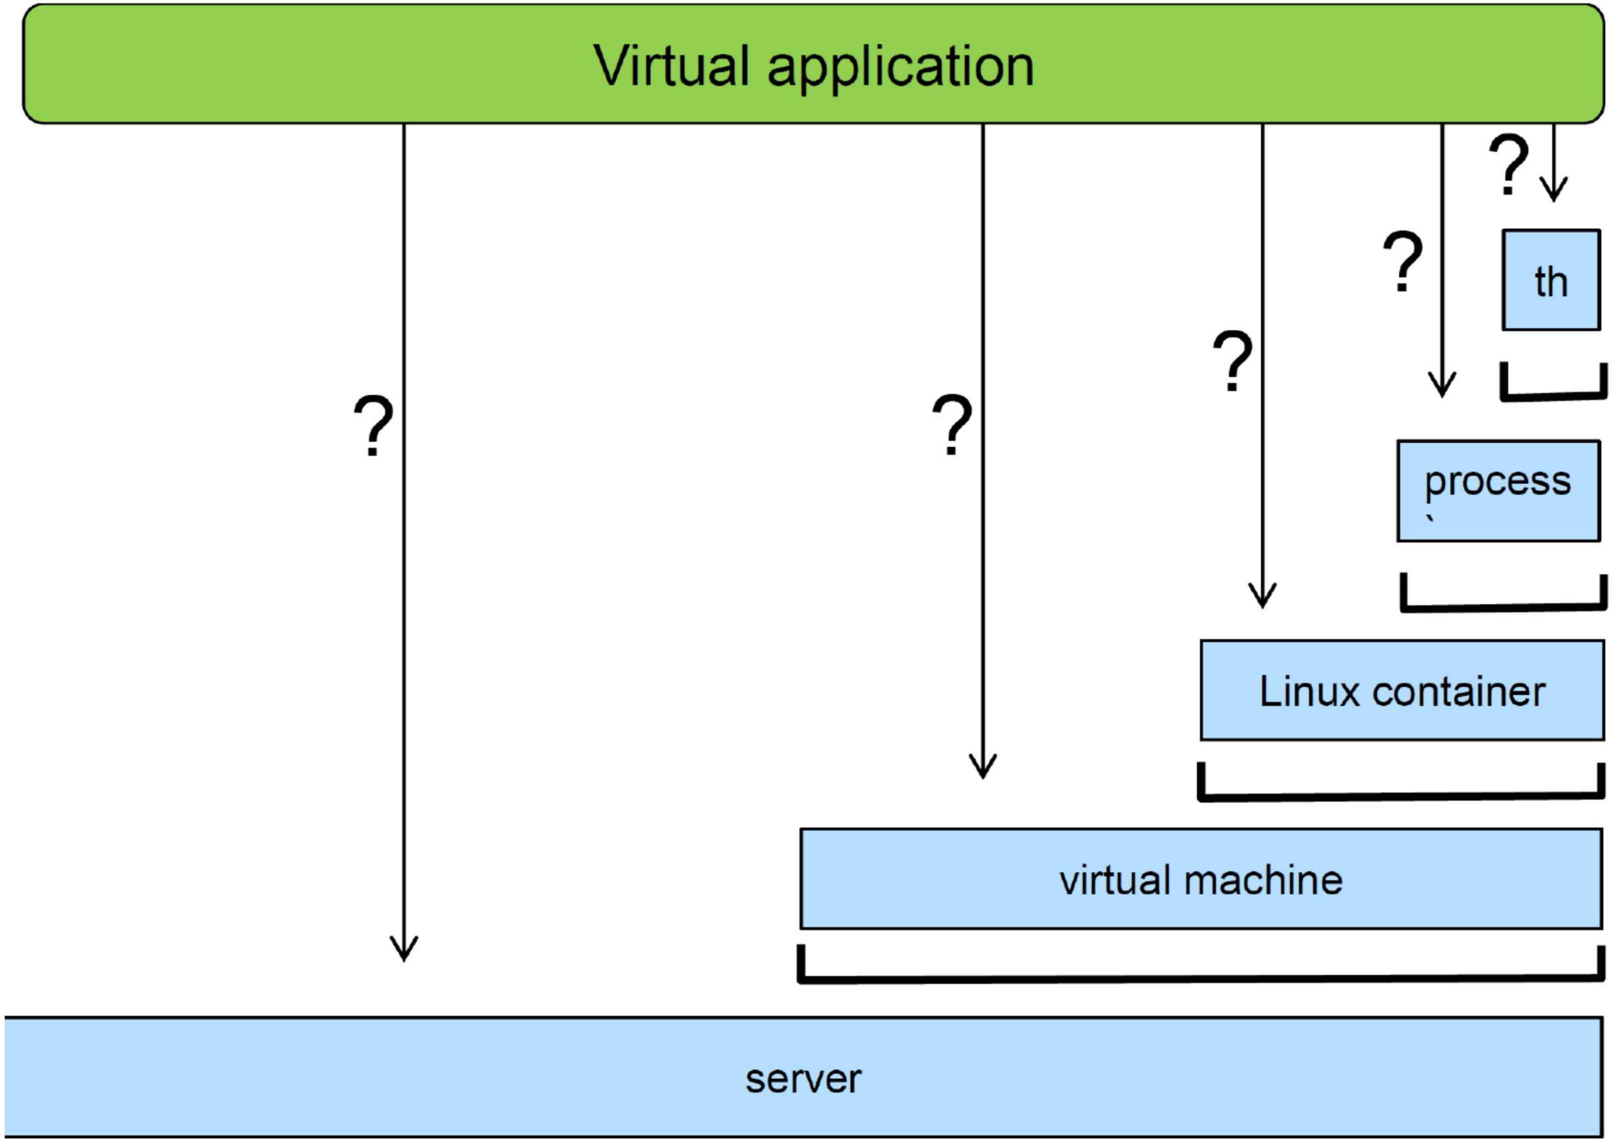

NFV Concepts

- **Network Function (NF):** Functional building block with a well defined interfaces and well defined functional behavior
- **Virtualized Network Function (VNF):** Software implementation of NF that can be deployed in a virtualized infrastructure
- **VNF Set:** Connectivity between VNFs is not specified, e.g., residential gateways
- **VNF Forwarding Graph:** Service chain when network connectivity order is important, e.g., firewall, NAT, load balancer
- **NFV Infrastructure (NFVI):** Hardware and software required to deploy, mange and execute VNFs including computation, networking, and storage.
- **NFV Orchestrator:** Automates the deployment, operation, management, coordination of VNFs and NFVI.

Source: Adapted from Raj Jain

NFV Concepts

- **NFVI Point of Presence (PoP):** Location of NFVI
- **NFVI-PoP Network:** Internal network
- **Transport Network:** Network connecting a PoP to other PoPs or external networks
- **VNF Manager:** VNF lifecycle management e.g., instantiation, update, scaling, query, monitoring, fault diagnosis, healing, termination
- **Virtualized Infrastructure Manager:** Management of computing, storage, network, software resources
- **Network Service:** A composition of network functions and defined by its functional and behavioral specification
- **NFV Service:** A network services using NFs with at least one VNF.

Source: Adapted from Raj Jain

Network Forwarding Graph

- An end-to-end service may include nested forwarding graphs

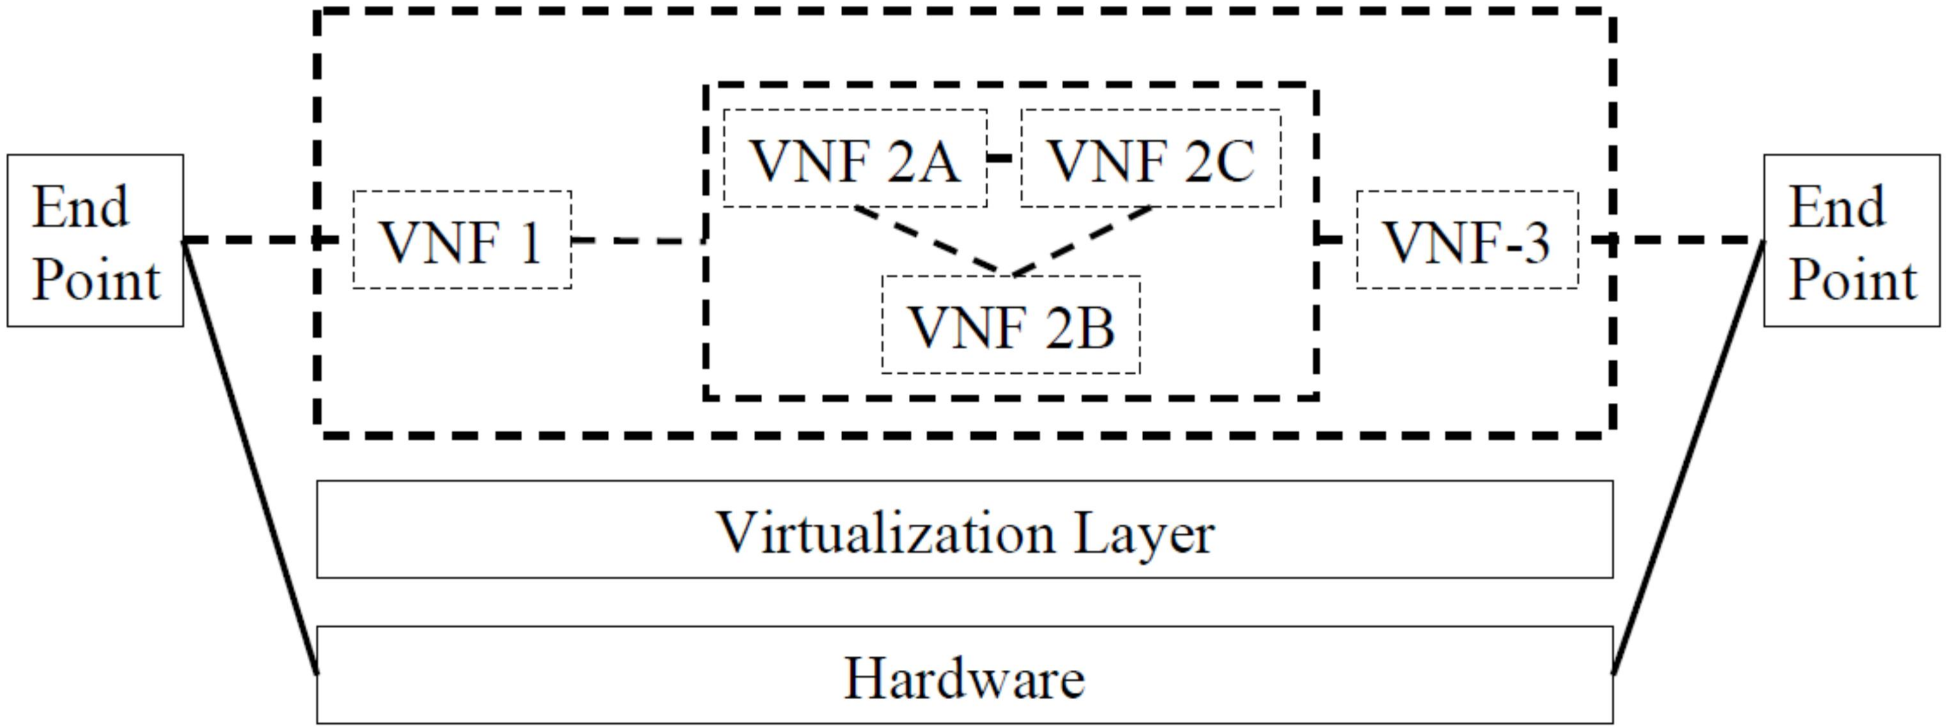

Source: Adapted from Raj Jain

NFV Concepts

- **User Service:** Services offered to end users/customers/subscribers.
- **Deployment Behavior:** NFVI resources that a VNF requires, e.g., Number of VMs, memory, disk, images, bandwidth, latency
- **Operational Behavior:** VNF instance topology and lifecycle operations, e.g., start, stop, pause, migration, …
- **VNF Descriptor:** Deployment behavior + Operational behavior

Source: Adapted from Raj Jain

The NFV Architecture Framework

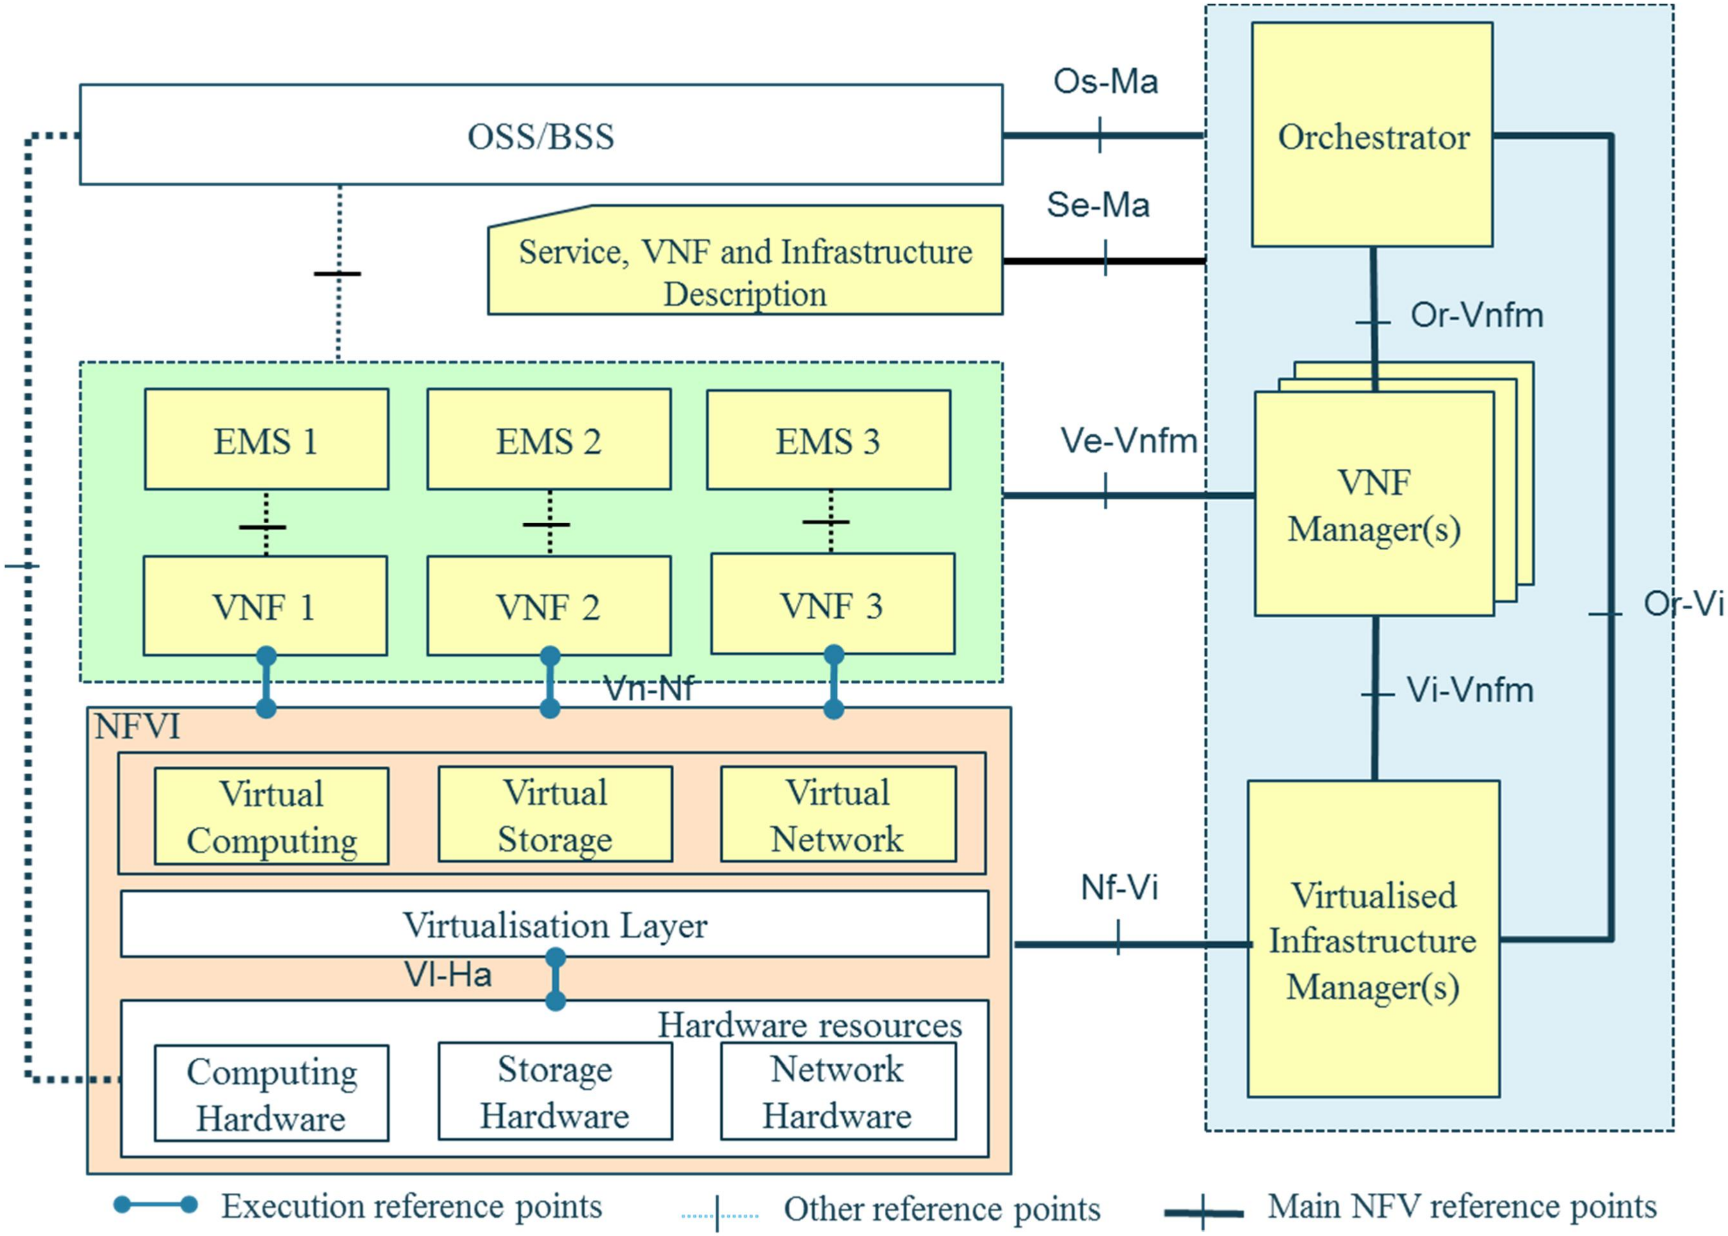

Reference Point:
Points for inter-module specification

- **(Os-Ma)** Operation Support System (OSS)/Business Support Systems (BSS) – NFV Management and Orchestration
- **(Se-Ma)** Service, VNF and Infrastructure Description – NFV Management and Orchestration: VNF Deployment template, VNF Forwarding Graph, service-related information, NFV infrastructure information
- **(Or-Vnfm)** Orchestrator – VNF Manager
- **(Vi-Vnfm)** Virtualized Infrastructure Manager – VNF Manager
- **(Ve-Vnfm)** VNF/ Element Management System (EMS) – VNF Manager
- **(Or-Vi)** Orchestrator – Virtualized Infrastructure Manager
- **(Nf-Vi)** NFVI-Virtualized Infrastructure Manager
- **(Vl-Ha)** Virtualization Layer-Hardware Resources
- **(Vn-Nf)** VNF – NFVI

The NFV Entities

Network Service (NS):

- Described by the NS descriptor, orchestrated by NFVO
- May cover 1 or more VNF Graphs, VNFs and PNFs

VNF Forwarding Graph (VNFFG):

- Described by the VNFFG descriptor, orchestrated by NFVO
- May cover VNFFGs, VNFs and NFs

VNF:

- Described by the VNF descriptor, instantiated by the VNF Manager
- Covers VNF components each mapped to a VM and described as a Virtual Deployment Unit

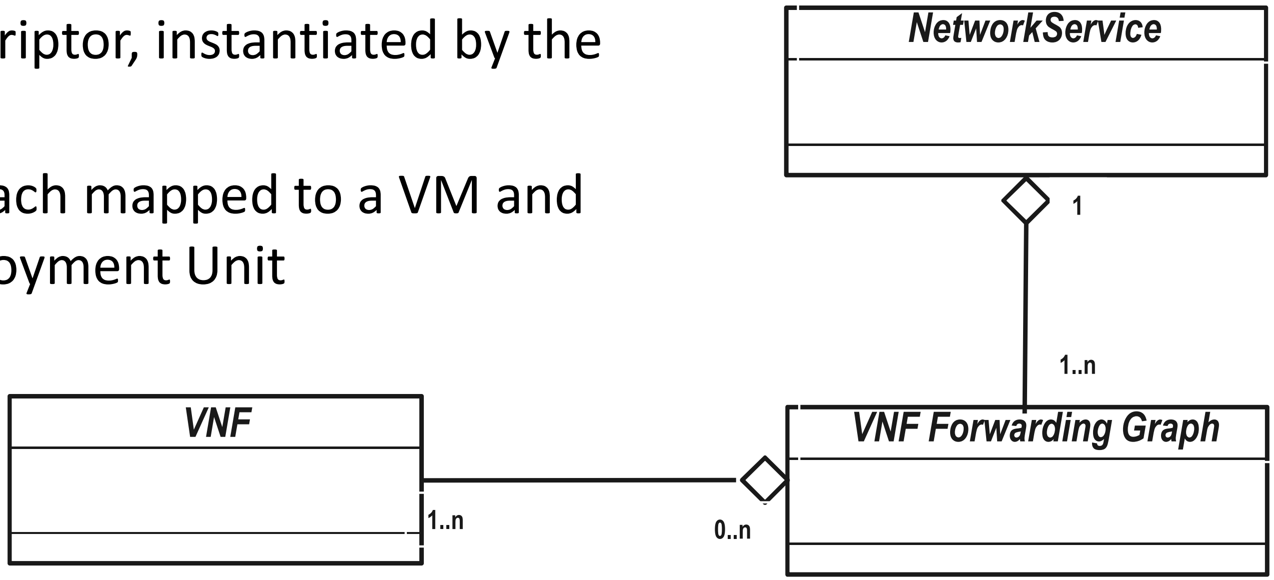

Architectural Use Cases

- Network Functions Virtualisation Infrastructure as a Service
  - Network functions go to the cloud
- Virtual Network Function as a Service
  - Ubiquitous, delocalized network functions
- Virtual Network Platform as a Service
  - Applying multi-tenancy at the VNF level
- VNF Forwarding Graphs
  - Building E2E services by composition

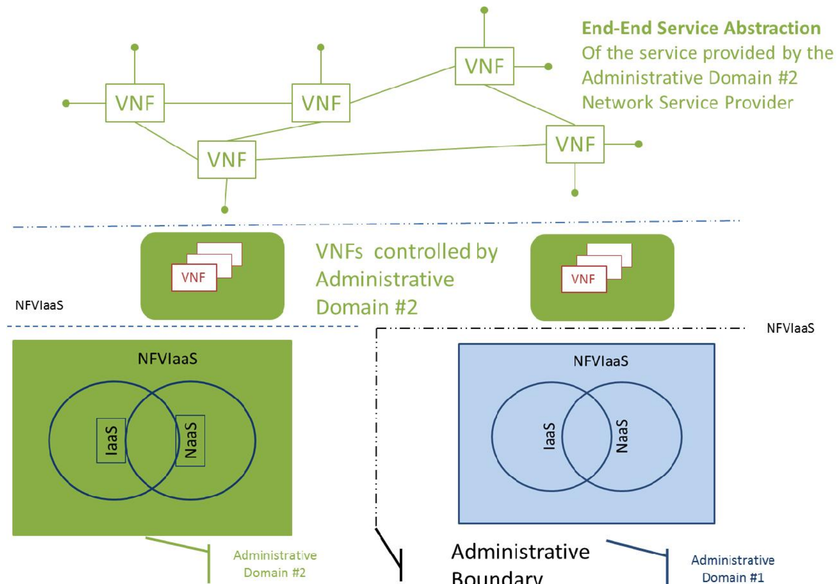

XaaS for Network Services

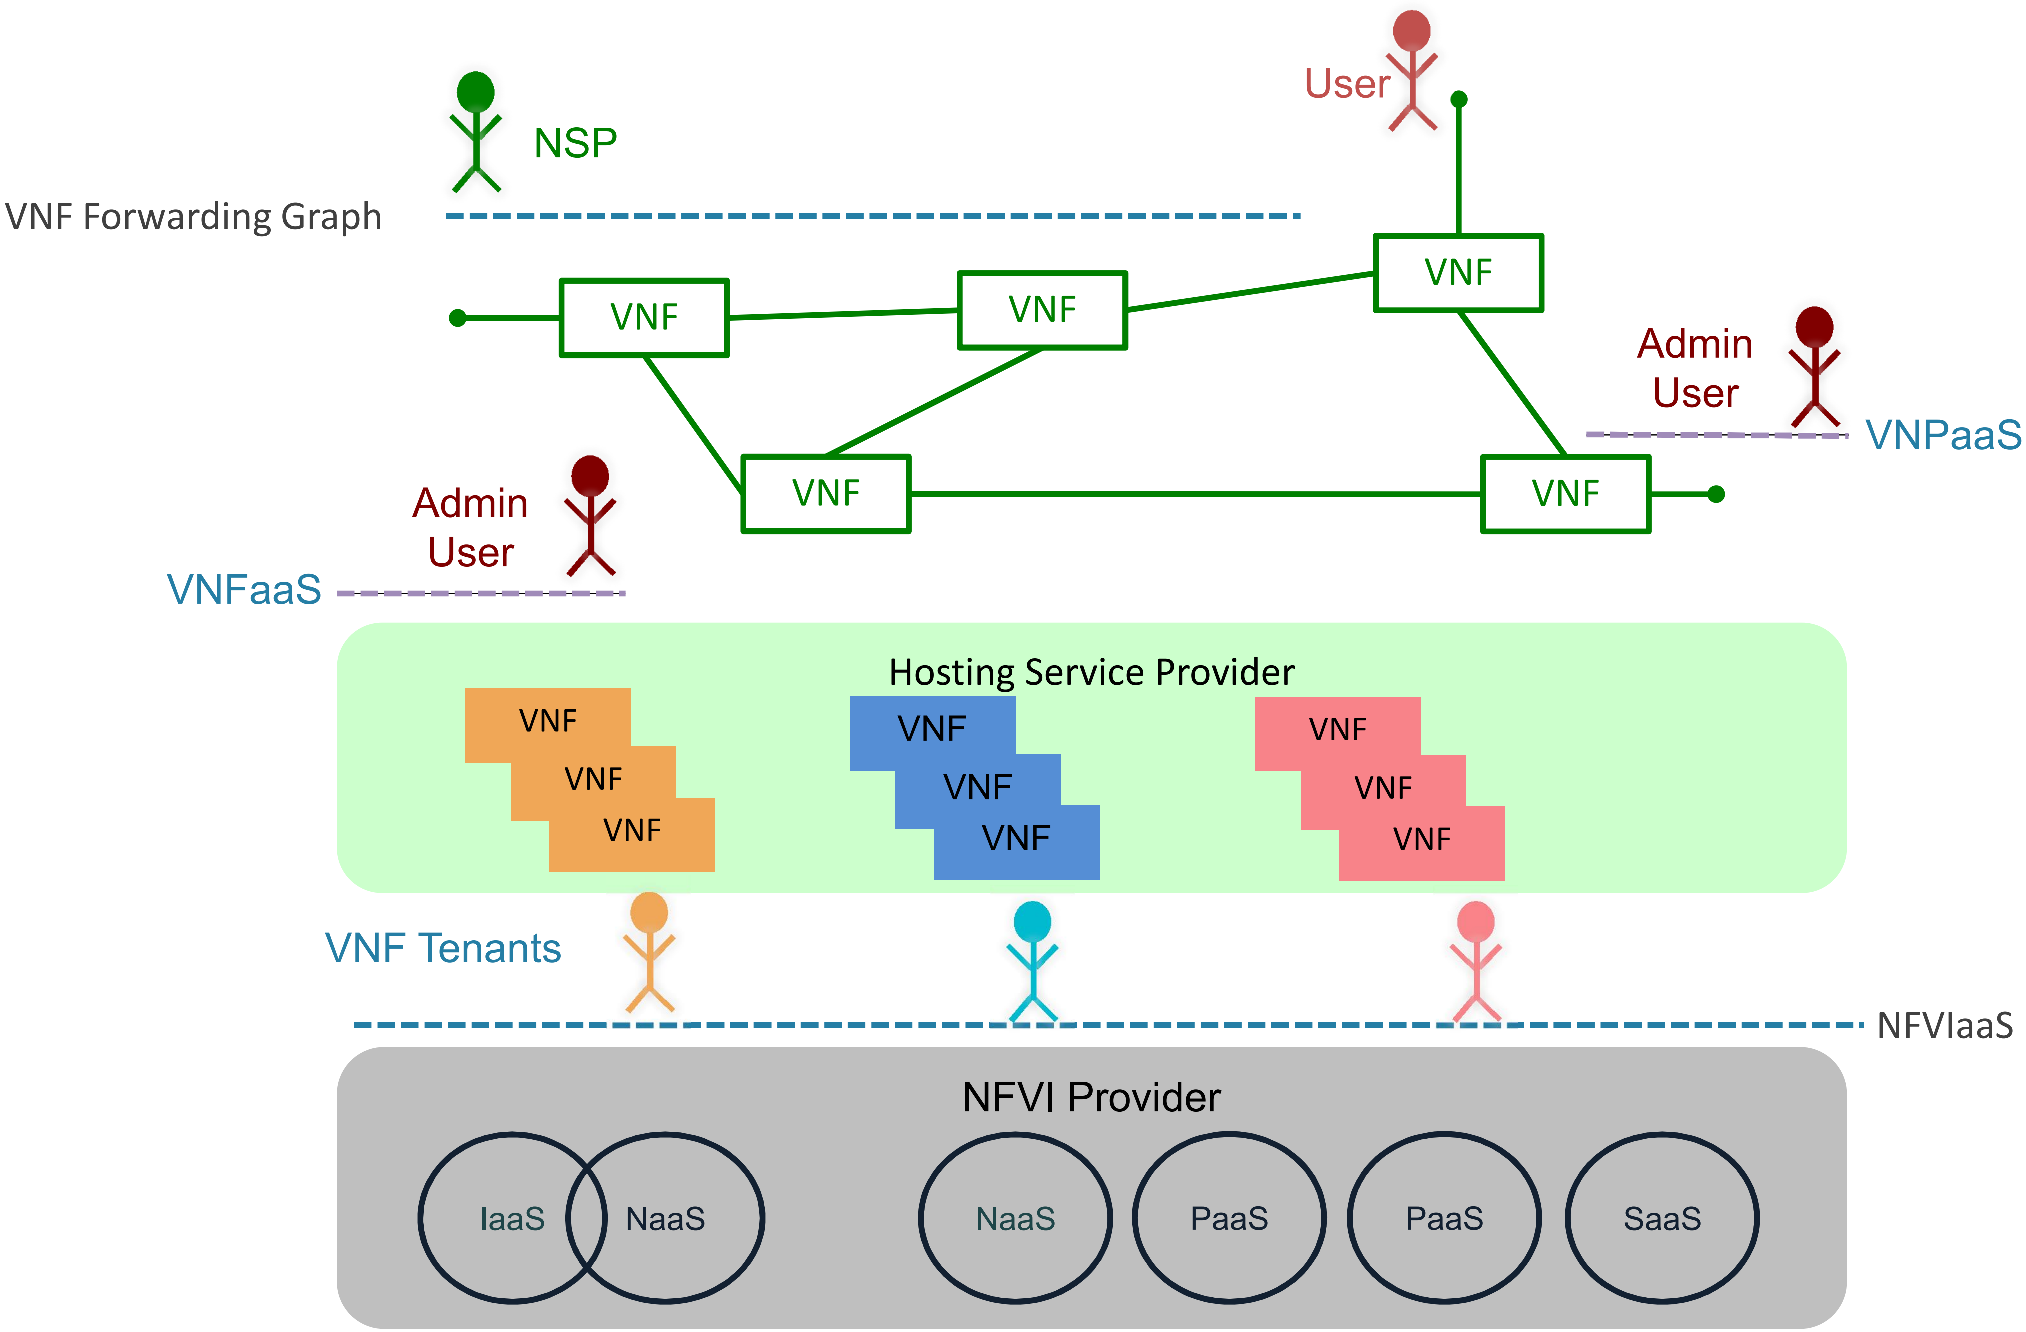

Virtual Infrastructure Management

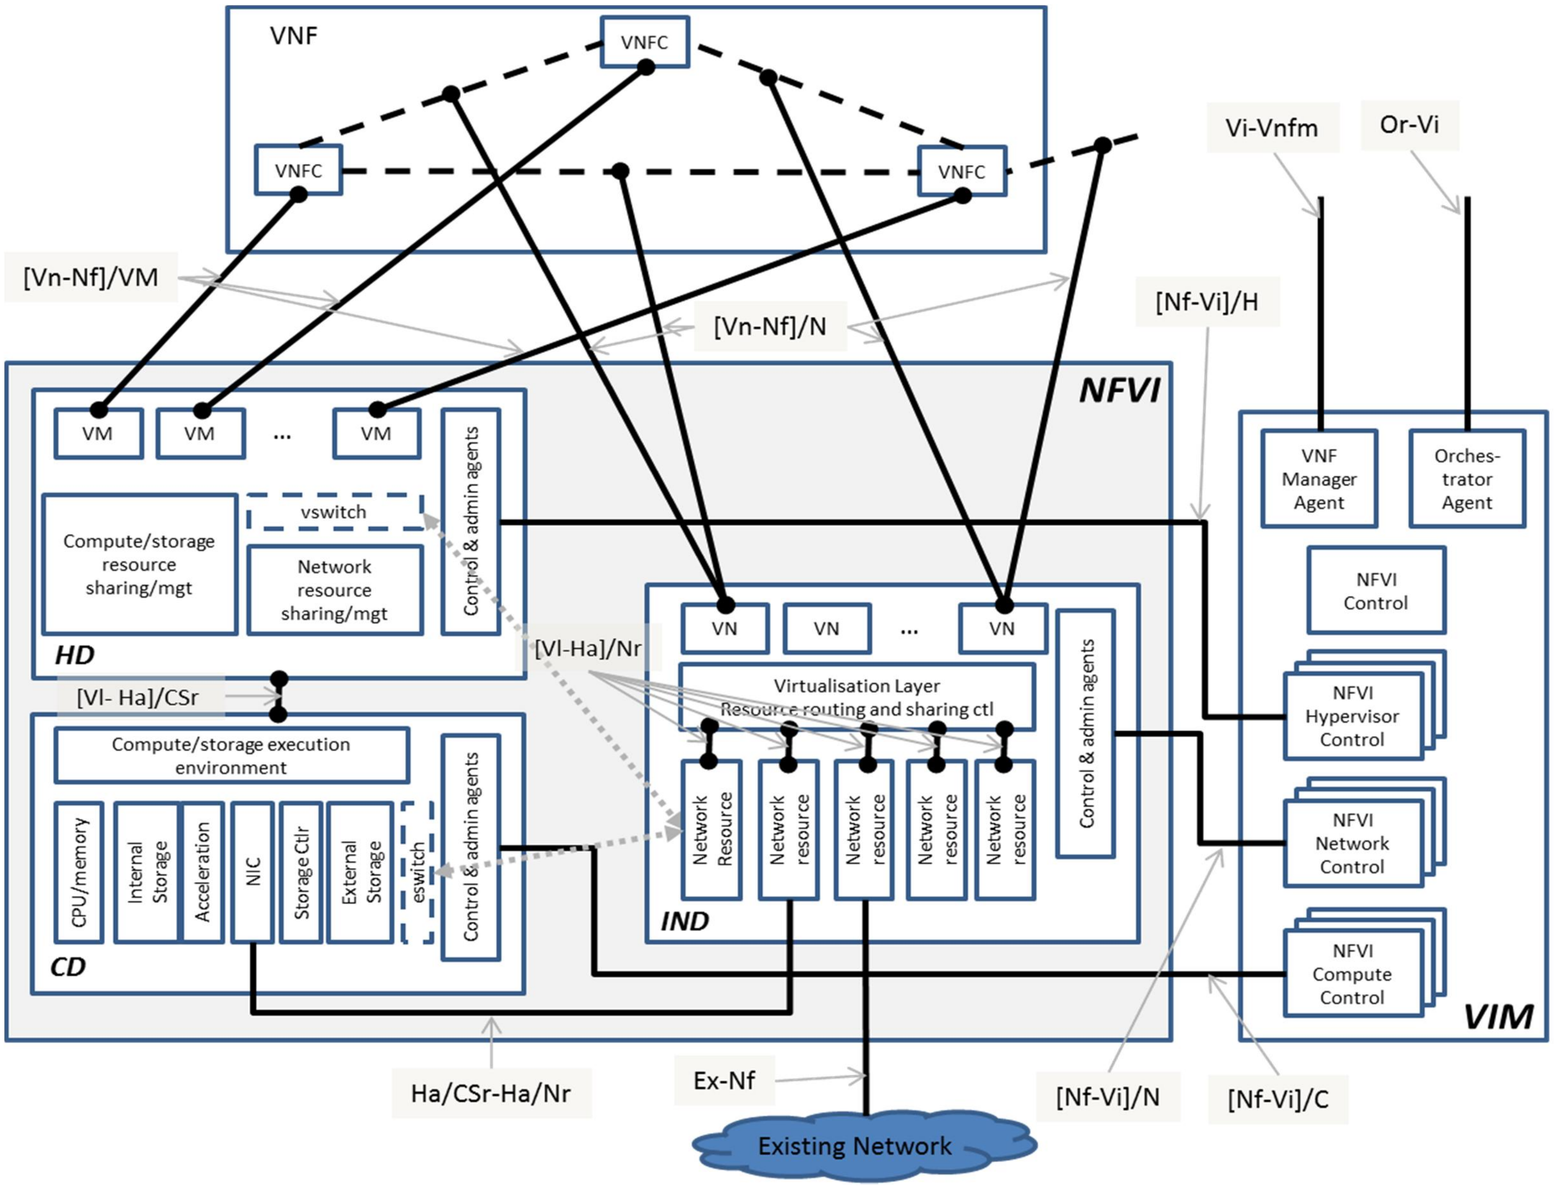

Infrastructure Network Control Hierarchy

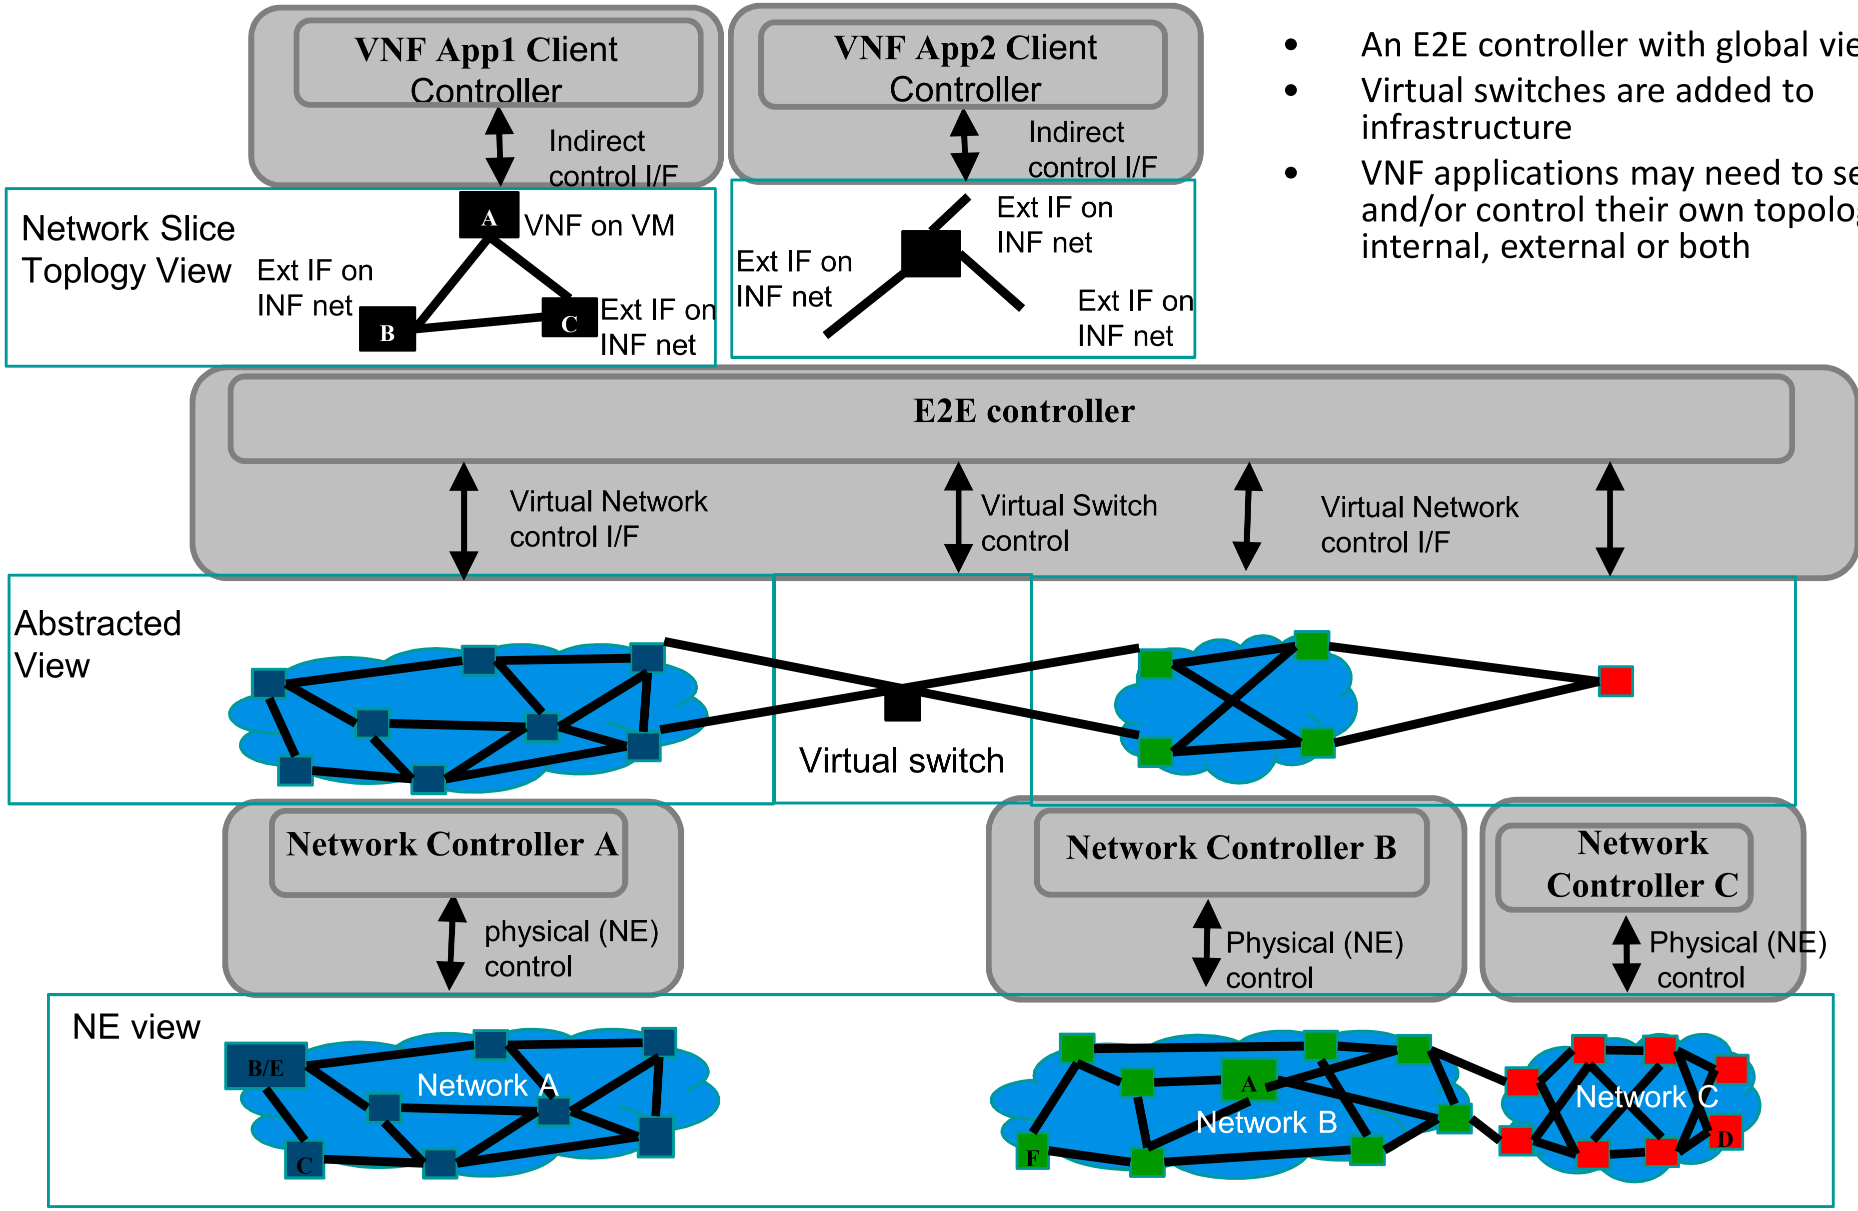

- An E2E controller with global view
- Virtual switches are added to infrastructure
- VNF applications may need to see and/or control their own topology, internal, external or both

CloudNFV Initiative

- An open implementation of the ETSI NFV ISG work

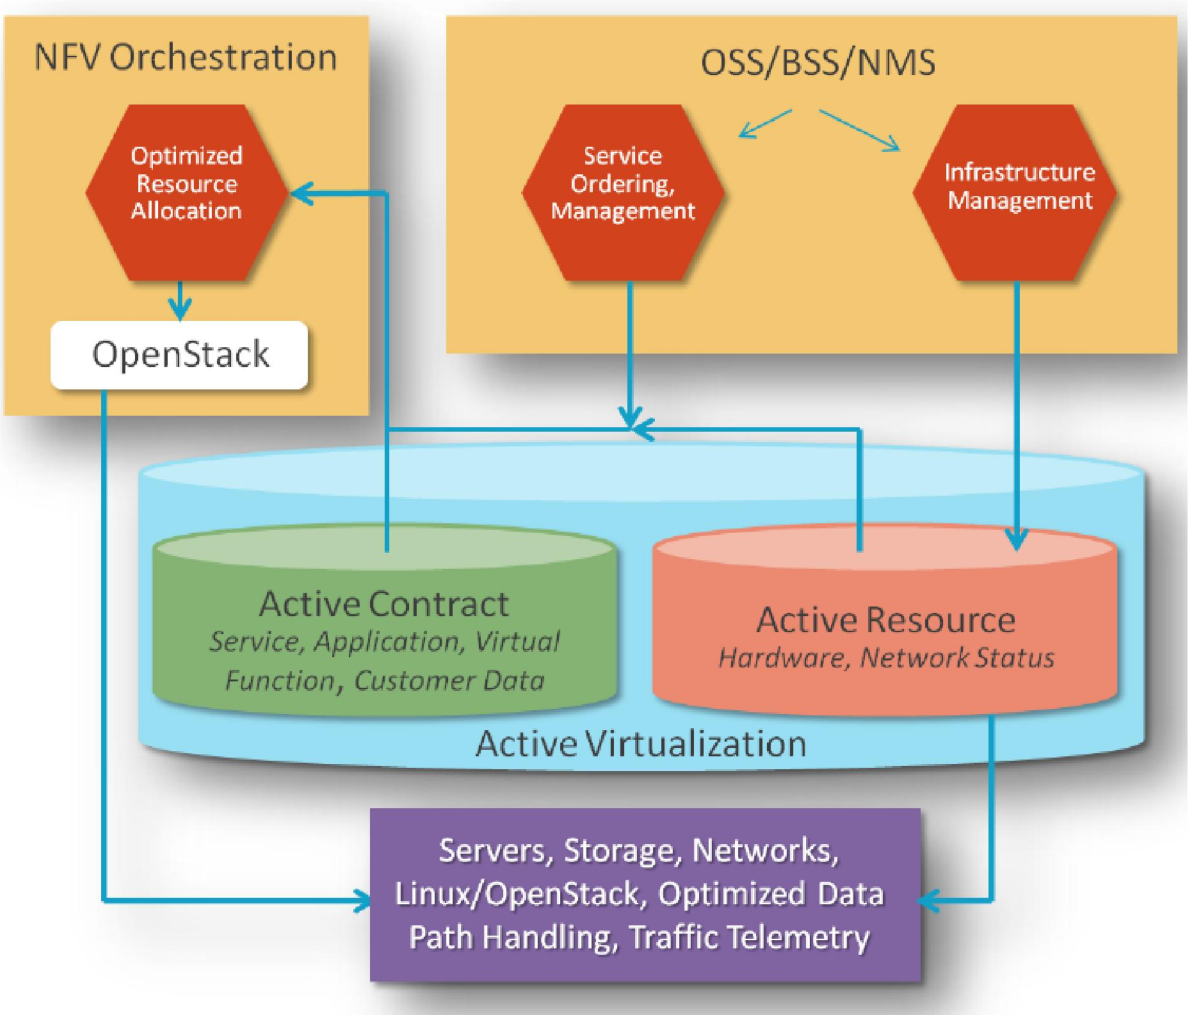

Founding members of the CloudNFV initiative include:

- 6WIND
- CIMI Corporation
- Dell
- EnterpriseWeb
- Overture Networks
- Qosmos

The Road to NFV

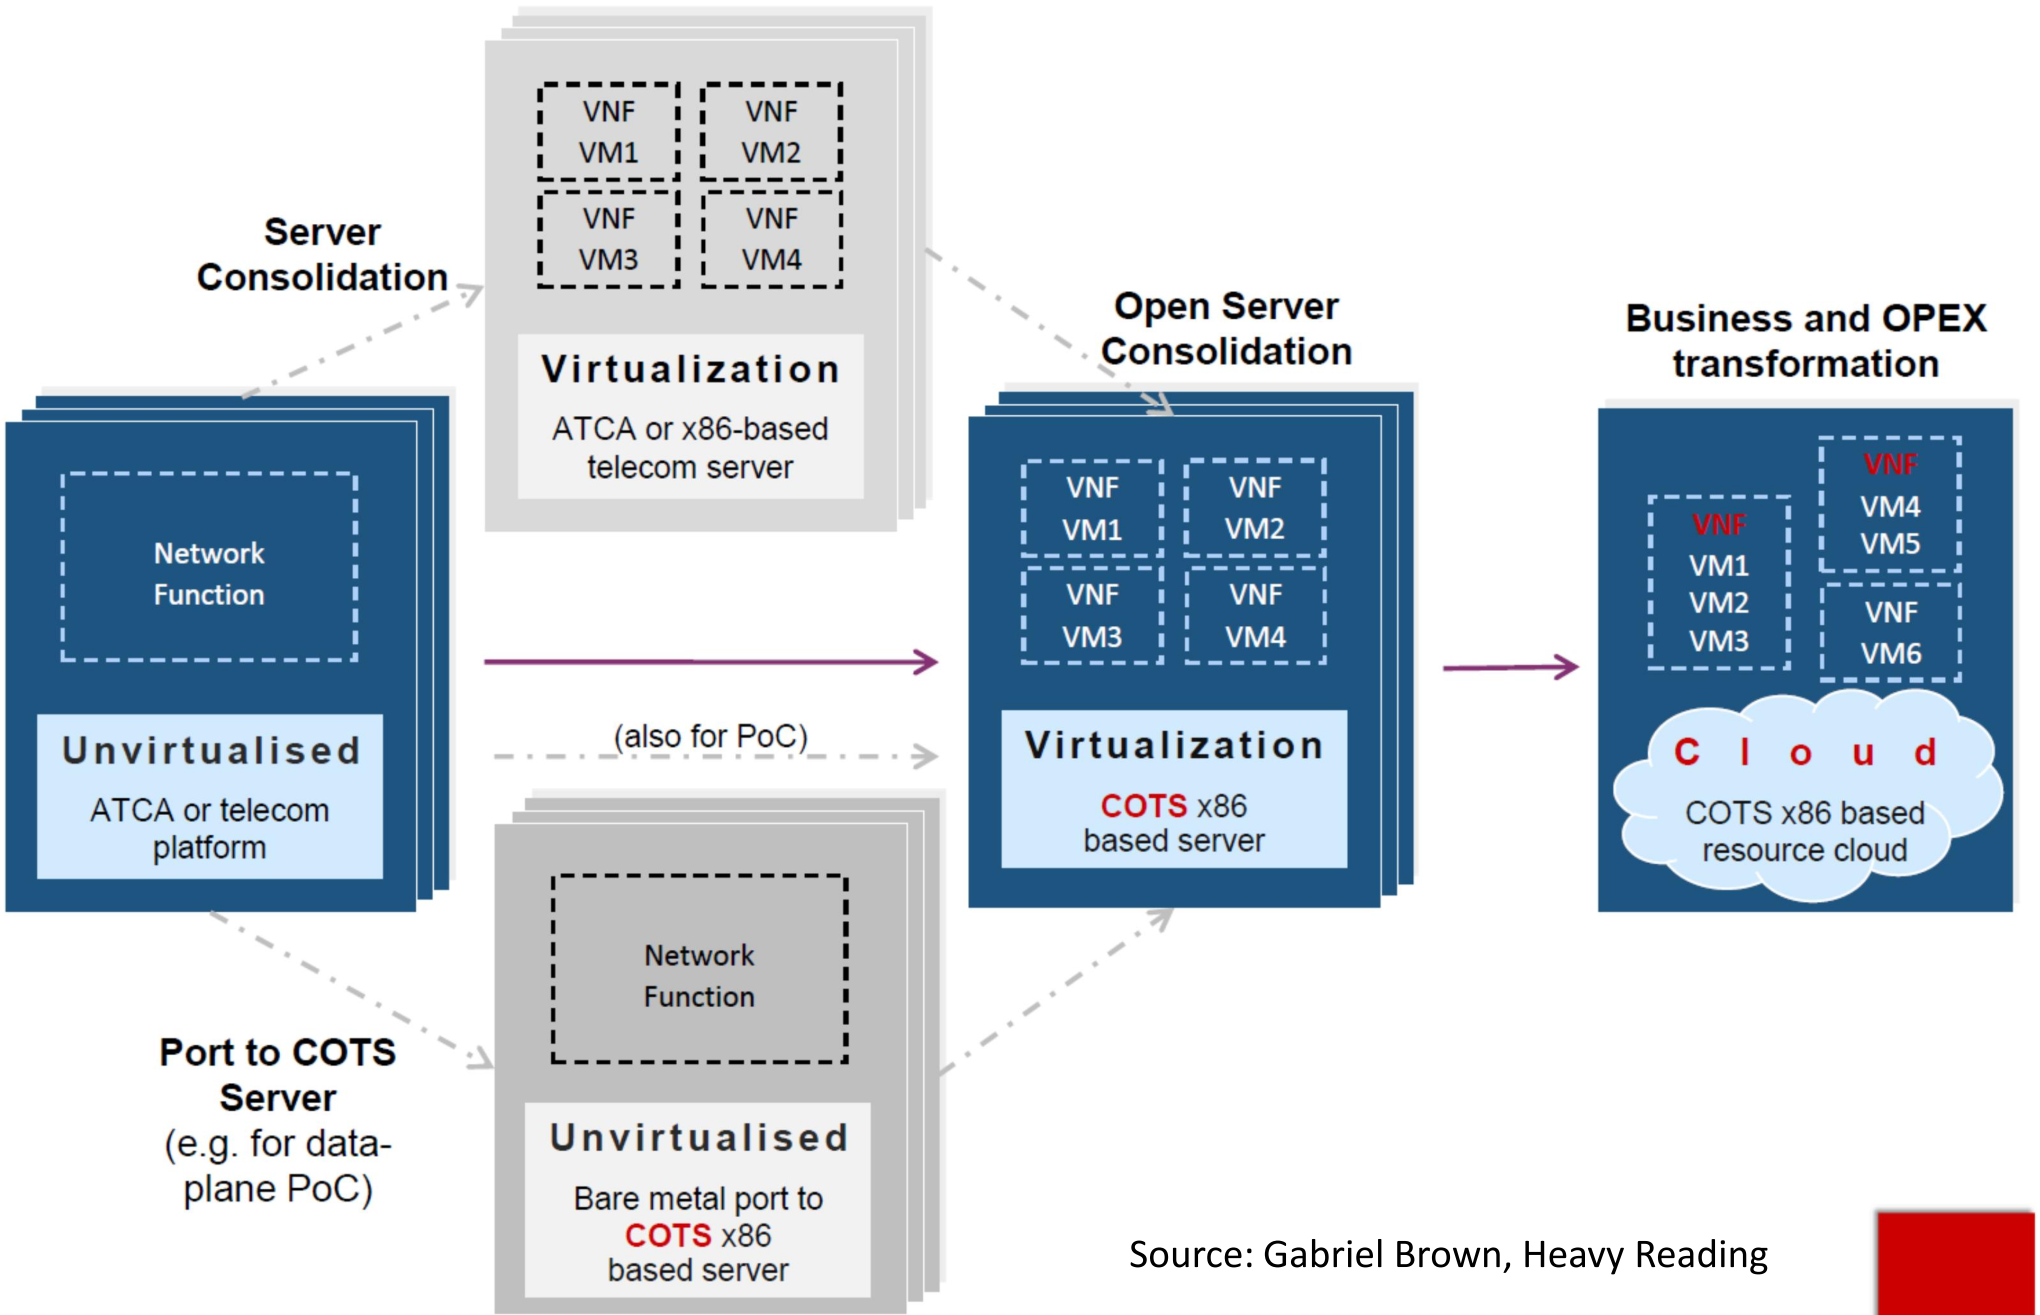

Source: Gabriel Brown, Heavy Reading

Two Approaches to NFV

(to be pursued simultaneously)

Application-driven NFV

- Operator starts with a particular function or domain e.g. IMS
- Increase VNFs over time as technology & opportunity allow
- Faster, less risky; an opportunity to experiment

Platform-driven NFV

- Operator starts to develop a horizontal platform to run VNFs
- Evolve platform to support demanding workloads; add VNFs
- Strategic, disruptive, expensive; long-term

Source: Gabriel Brown, Heavy Reading

Recommendation / Call-for-action

Invitation towards IT and Telecom industries
to combine their complementary expertise and resources
in a joint collaborative effort, to reach broad agreement
on standardised approaches and common architectures,
and which are interoperable and have economies of scale.

- A new carrier-led Industry Specification Group (ISG) is being setup under the auspices of ETSI.
  - Initial face-to-face meeting of the ISG NFV is planned for Jan 2013, and will be announced via the usual ETSI procedures.
- Deliverables: White papers addressing challenges and operator requirements, as input to standardisation bodies.

Source: NFV

NFV Myths

- The ETSI NFV ISG is a standards body.
- NFV equates to “The Cloud.”
- NFV is about CAPEX.
- NFV winds down in January 2015.

Wrapping up: Innovations of NFV

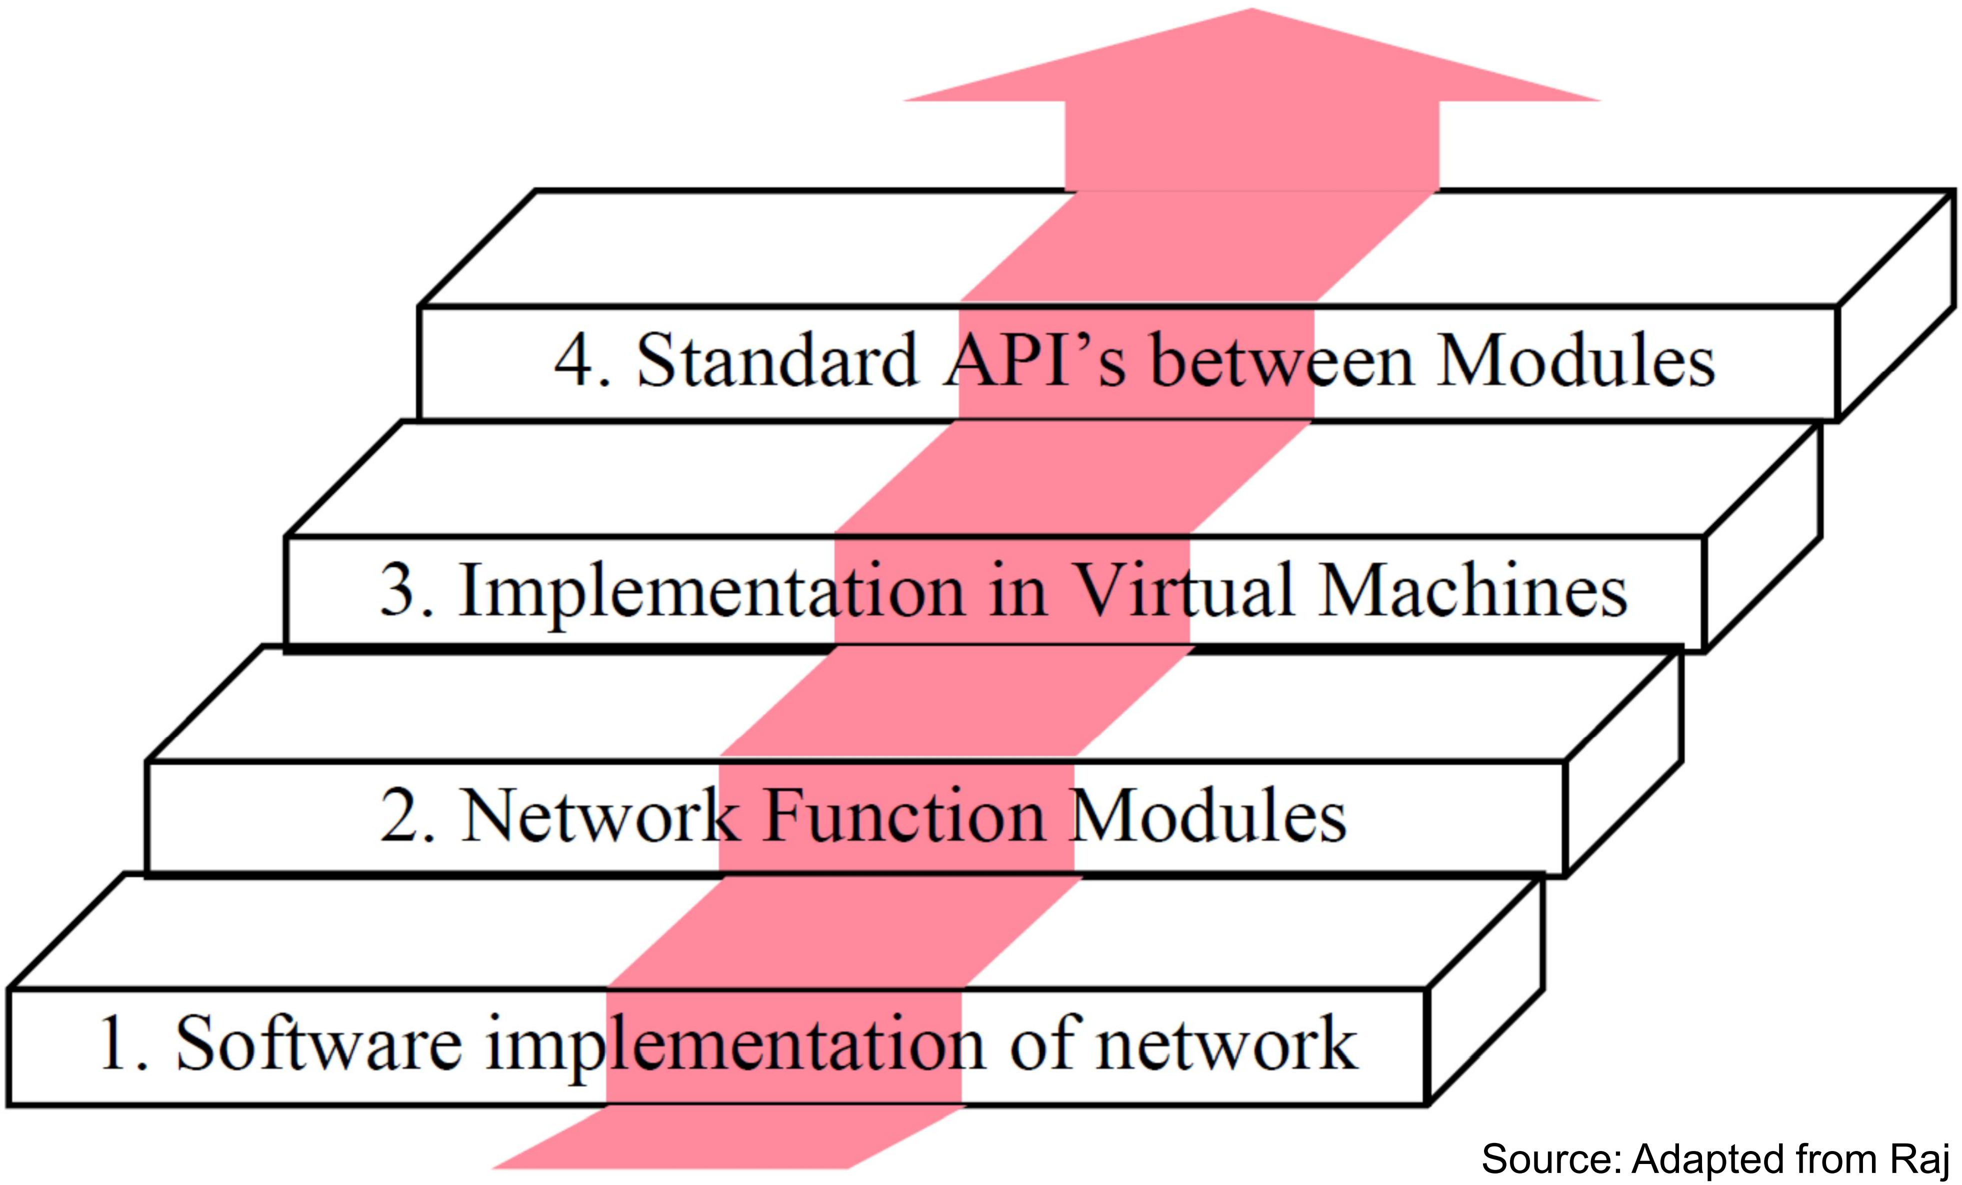

Source: Adapted from Raj Jain

Challenges

- Achieving high performance virtualised network appliances
  - portable between different HW vendors, and with different hypervisors.
- Co-existence with bespoke HW based network platforms
  - enabling efficient migration paths to fully virtualised network platforms.
- Management and orchestration of virtual network appliances
  - ensuring security from attack and misconfiguration.
- NFV will only scale if all of the functions can be automated.
- Appropriate level of resilience to HW and SW failures.
- Integrating multiple virtual appliances from different vendors.
  - Network operators need to be able to "mix & match" HW,
  - hypervisors from different vendors,
  - and virtual appliances from different vendors
  - without incurring significant integration costs and avoiding lock-in.

Source: NFV

Requirements and Challenges

NFV

NFV Requirements

- Too many
  - Based on the environment
    - Specifically bounded

New abstractions,

but new and some already old challenges and requirements

NFV Framework Requirements

1. General: Partial or full Virtualization, Predictable performance
2. Portability: Decoupled from underlying infrastructure
3. Performance: Conforming and proportional to NFs specifications and facilities to monitor
4. Elasticity: Scalable to meet SLAs. Movable to other servers.
5. Resiliency: Be able to recreate after failure.

Specified packet loss rate, calls drops, time to recover, etc.

6. Security: Role-based authorization, authentication
7. Service Continuity: Seamless or non-seamless continuity after failures or migration

NFV Framework Requirements

8. Service Assurance: Time stamp and forward copies of packets for Fault detection

9. Energy Efficiency Requirements: Should be possible to put a subset of VNF in a power conserving sleep state

10. Operational and Management Requirements: Incorporate mechanisms for automation of operational and management functions

11. Transition: Coexistence with Legacy and Interoperability among multi-vendor implementations

12. Service Models: Operators may use NFV infrastructure operated by other operators

Where should we go?

Two Approaches to NFV
(to be pursued simultaneously)

Application-driven NFV

- Operator starts with a particular function or domain e.g. IMS
- Increase VNFs over time as technology & opportunity allow
- Faster, less risky; an opportunity to experiment

Platform-driven NFV

- Operator starts to develop a horizontal platform to run VNFs
- Evolve platform to support demanding workloads; add VNFs
- Strategic, disruptive, expensive; long-term

Arising of challenges

Source: Gabriel Brown, Heavy Reading

First: A Few Challenges

- Achieving high performance virtualised network appliances
  - portable between different HW vendors, and with different hypervisors.
- Co-existence with bespoke HW based network platforms
  - enabling efficient migration paths to fully virtualised network platforms.
- Management and orchestration of virtual network appliances
  - ensuring security from attack and misconfiguration.
- NFV will only scale if all of the functions can be automated.
- Appropriate level of resilience to HW and SW failures.
- Integrating multiple virtual appliances from different vendors.
  - Network operators need to be able to “mix & match” HW,
  - hypervisors and virtual appliances from different vendors,
  - without incurring significant integration costs and avoiding lock-in.
- NFV and SDN

But... Based on what?
Use Cases
Then... More challenges!

Challenging Path upfront:
Not as simple as cloud applied to telco

The network differs from the computing environment in 2 key factors…

1. **Data plane workloads** (which are huge!) → HIGH PRESSURE ON **PERFORMANCE**

2. **Network requires shape** (+ E2E interconnection) → GLOBAL NETWORK VIEW IS REQUIRED FOR **MANAGEMENT**

…which are big challenges for vanilla cloud computing.

AN ADAPTED VIRTUALISATION ENVIRONMENT IS NEEDED
TO OBTAIN **CARRIER-CLASS** BEHAVIOUR

NFV Forwarding Graphs

- Network Service Chaining
  - Networks paths: old stratified vs. dynamic new

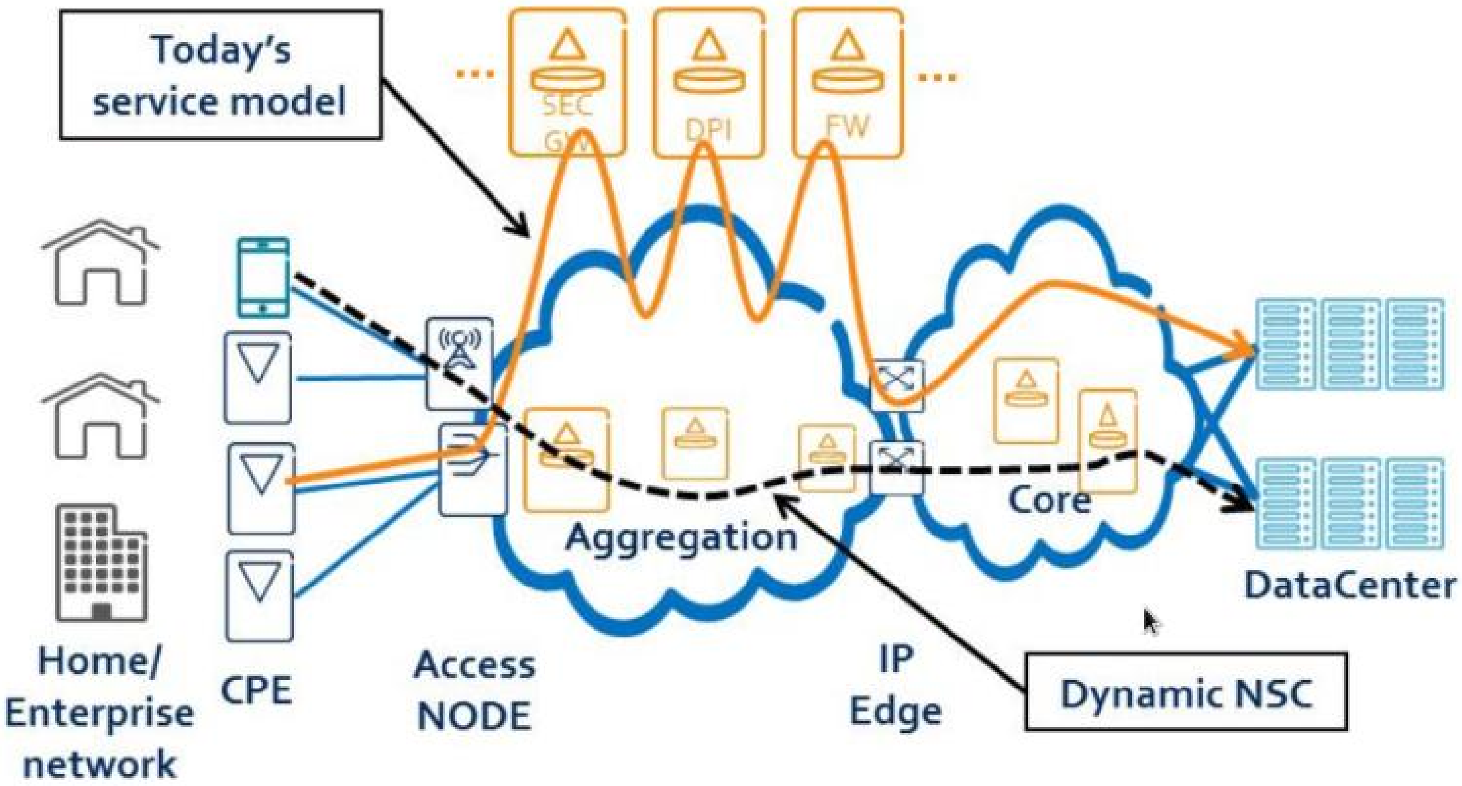

Source: Ericsson, EU UNIFY

NSC and NFV-FG

- Constitution of NSC
  - NF Set to NFV-FG
    - NFs well defined interfaces and behavior
- NFV-FGs topics:
  - Processing semantics
  - Performance guarantees
  - Chargings

Is NFV Technology Good Enough?

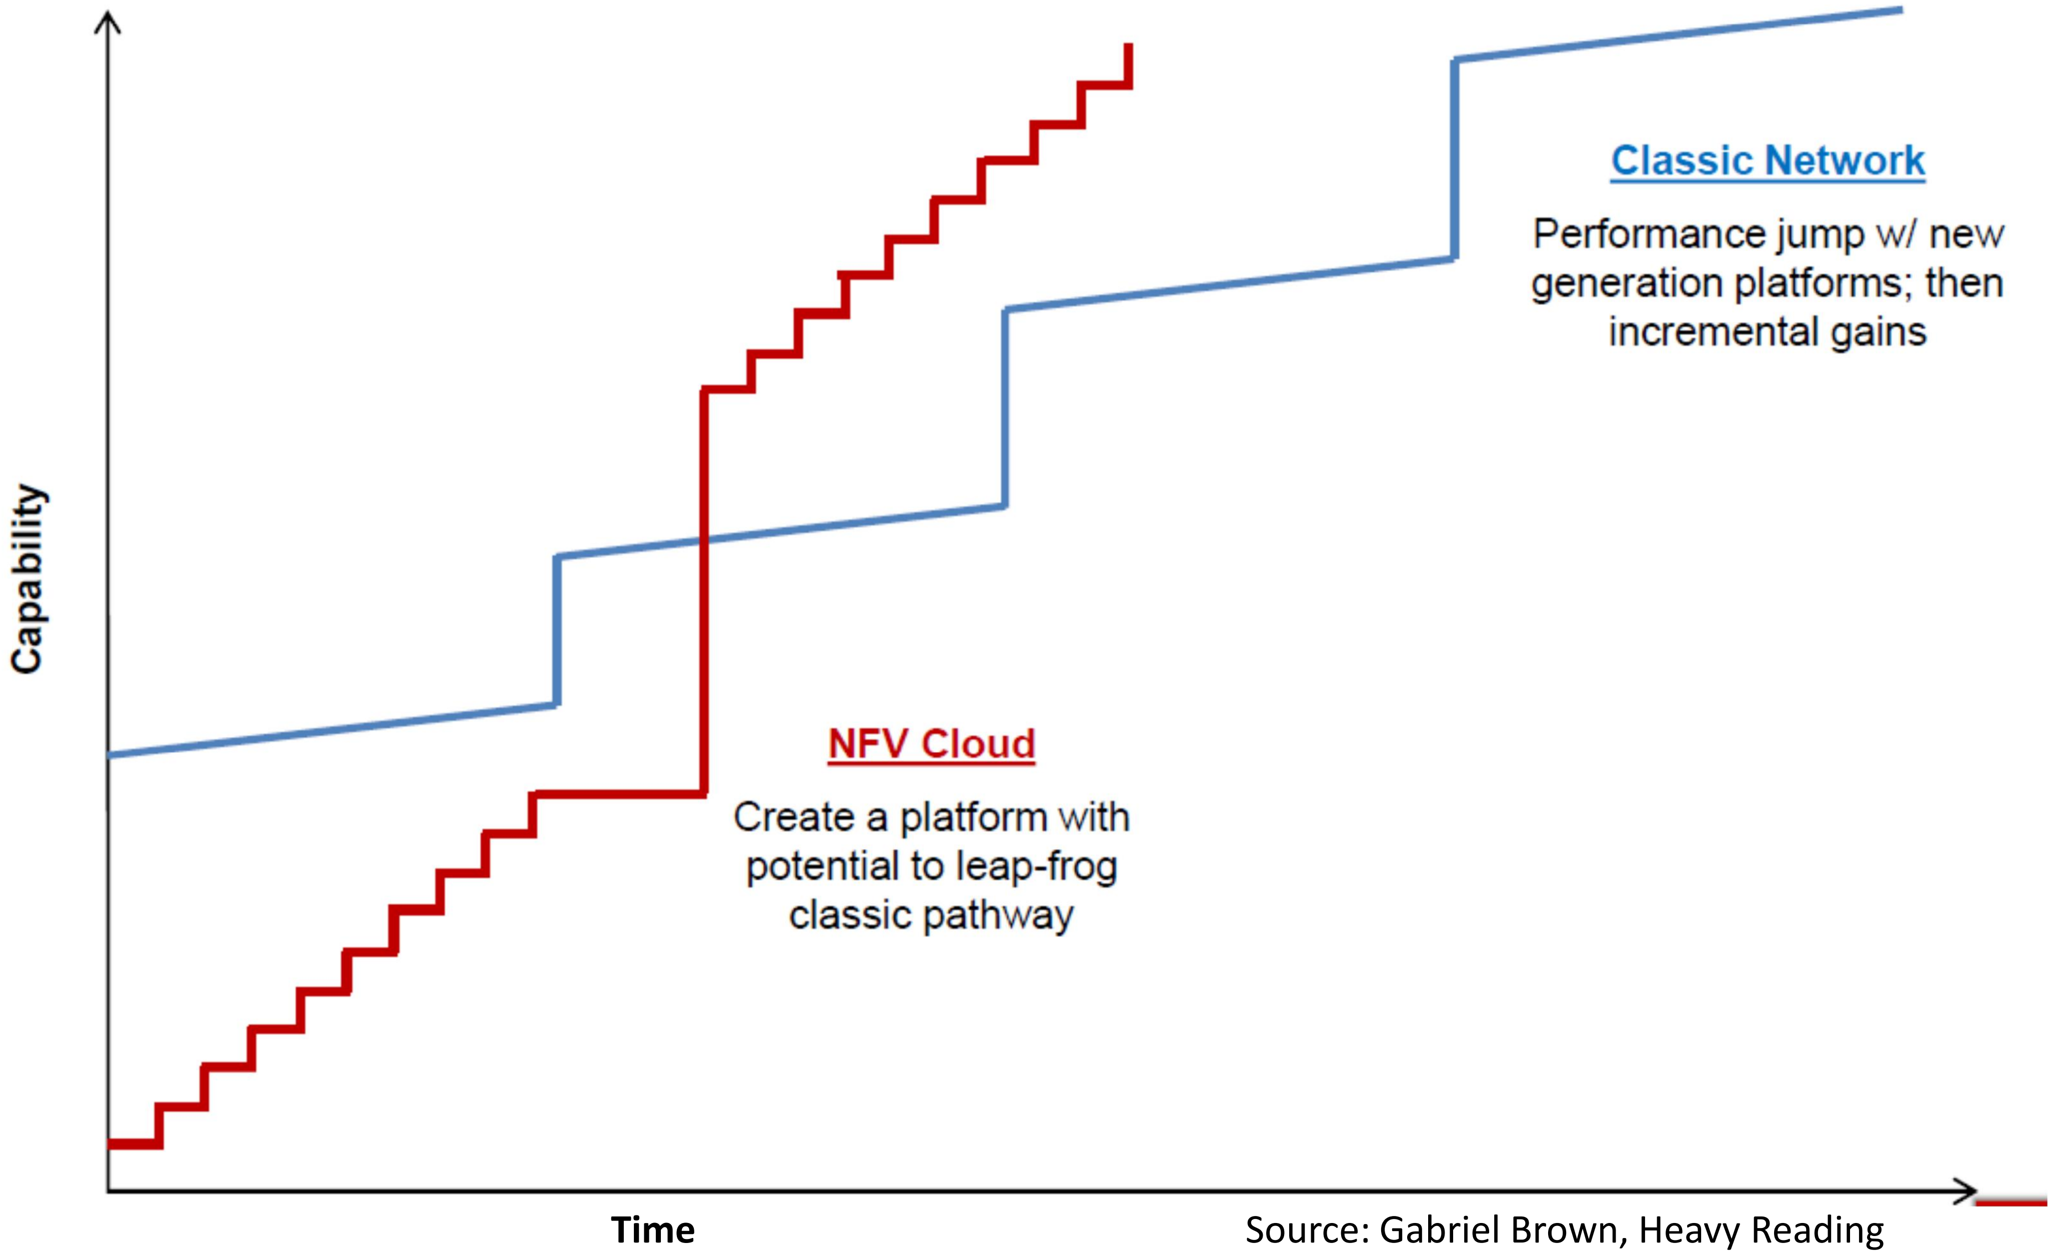

Source: Gabriel Brown, Heavy Reading

Performance

- Different network technologies have a cost...

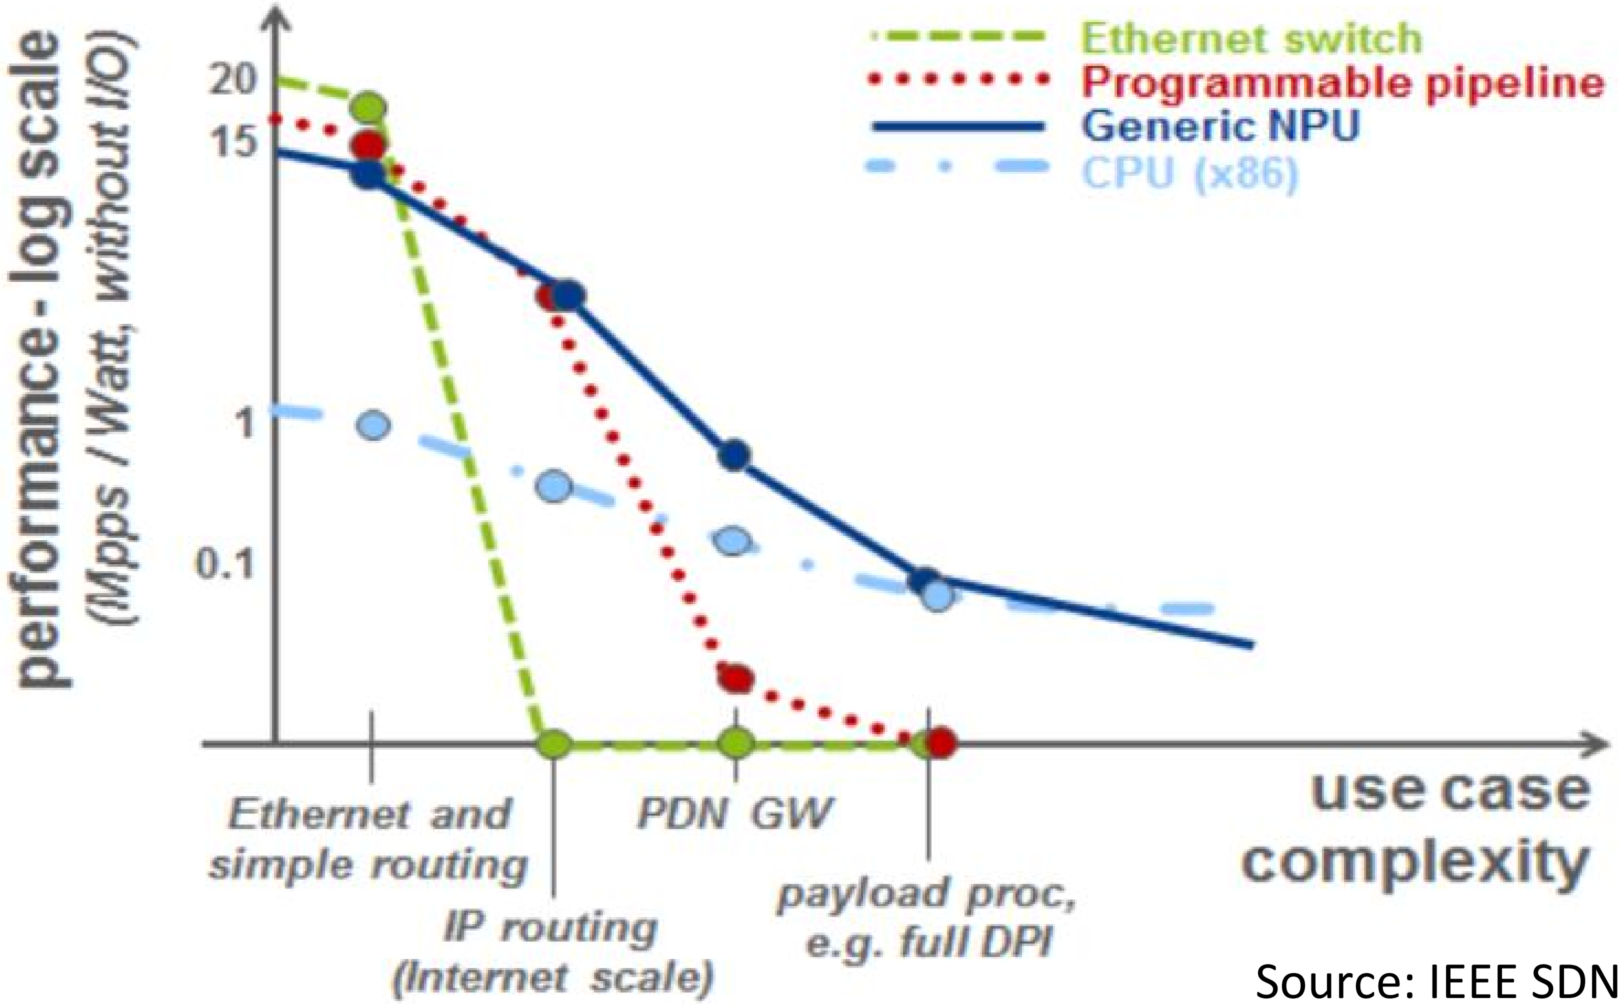

Source: IEEE SDN4NFS, 2013

NFV Performance Challenges

Typical performance

- 3-4 Gbps per CPU core assuming very light per-packet processing
- An order of magnitude less than what the hardware could do (more than 10Gbps per core, 40+ Gbps per x86 server)

Bottlenecks

- TCP stack and Linux kernel in NFV virtual machines
- Hypervisor virtual switch
- NIC TCP offload works only with VLANs

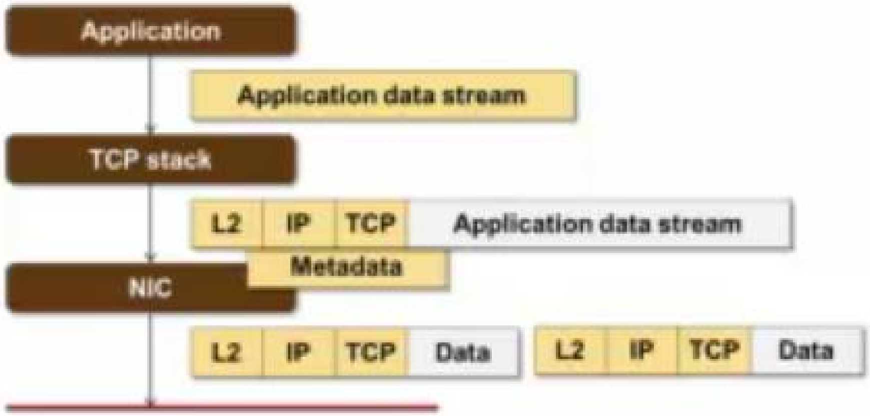

Solutions

- Optimized virtual switches (example: Intel DPDK)
- Dedicated virtual NICs (hypervisor bypass)
- Dedicated packet processing CPU cores
- User-mode packet processing (example: PF_RING)

Source: Ivan Pepelnjak

Scalability

- Real world vs. virtualized perspective
  - Network devices: FIB size, queue length, # of ports
- NFVI existence?
  - Distributed: storage, processing, connecting
  - Distributed NFs
    - Latency and Bandwidth requirements (e.g., BRAS, DPI)

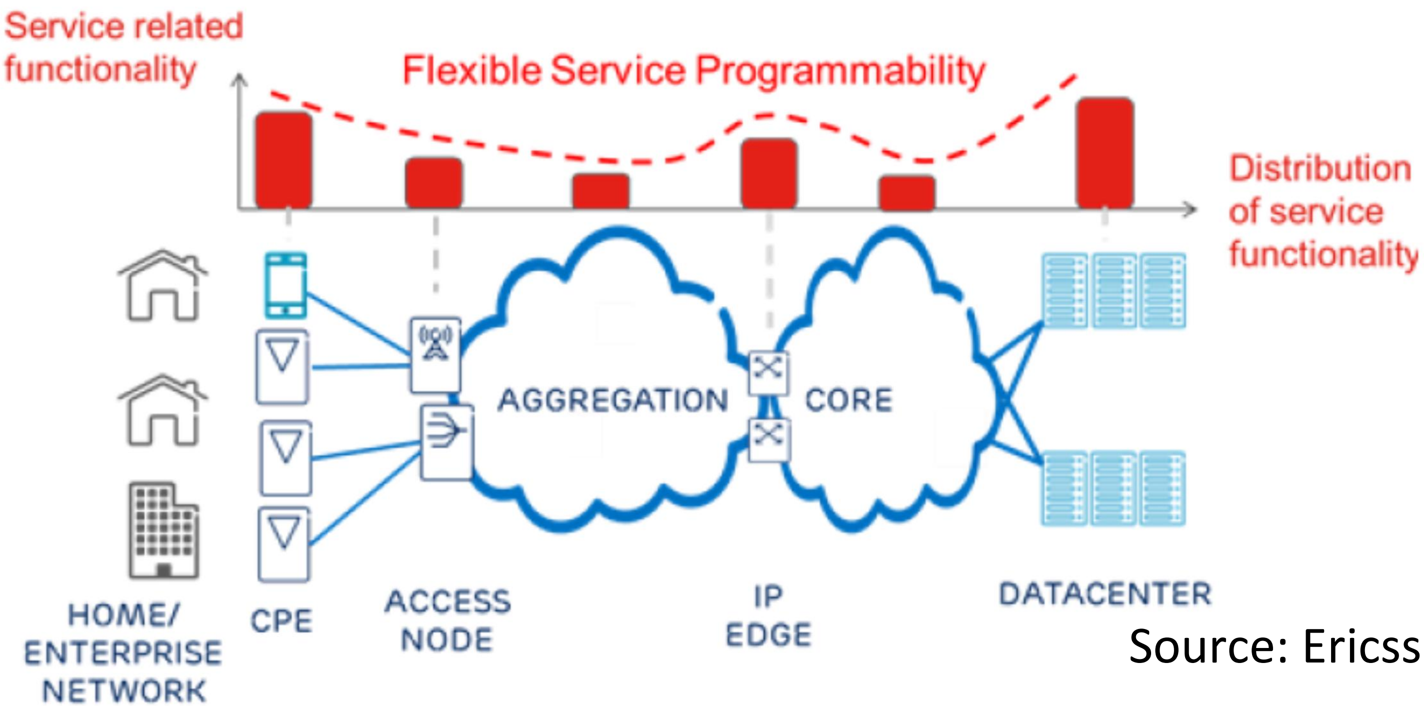

Source: Ericsson, EU UNIFY

Performance and Scalability

- PFs and NFs
  - Lack of performance -> Scalability decreased
- Performance
  - NF vs. NFV-FG
- Proportional performance of NFs and services according to available:
  - Network latency and bandwidth
  - Compute capacity

Performance and Scalability

Networking Workloads x **VMs** x **Traffic >>> Compute**

Networking Workload	Compute Cycles / Packet
L2 Forwarding	70
IP Routing	175
L2-4 Classification	750
TCP Termination	1500
Stateful Firewall	2250
OpenFlow Process	5000
IDS/IPS	5000
NextGen Firewall	8500
IPsec / SSL	9500
Firewall + SSL	18000

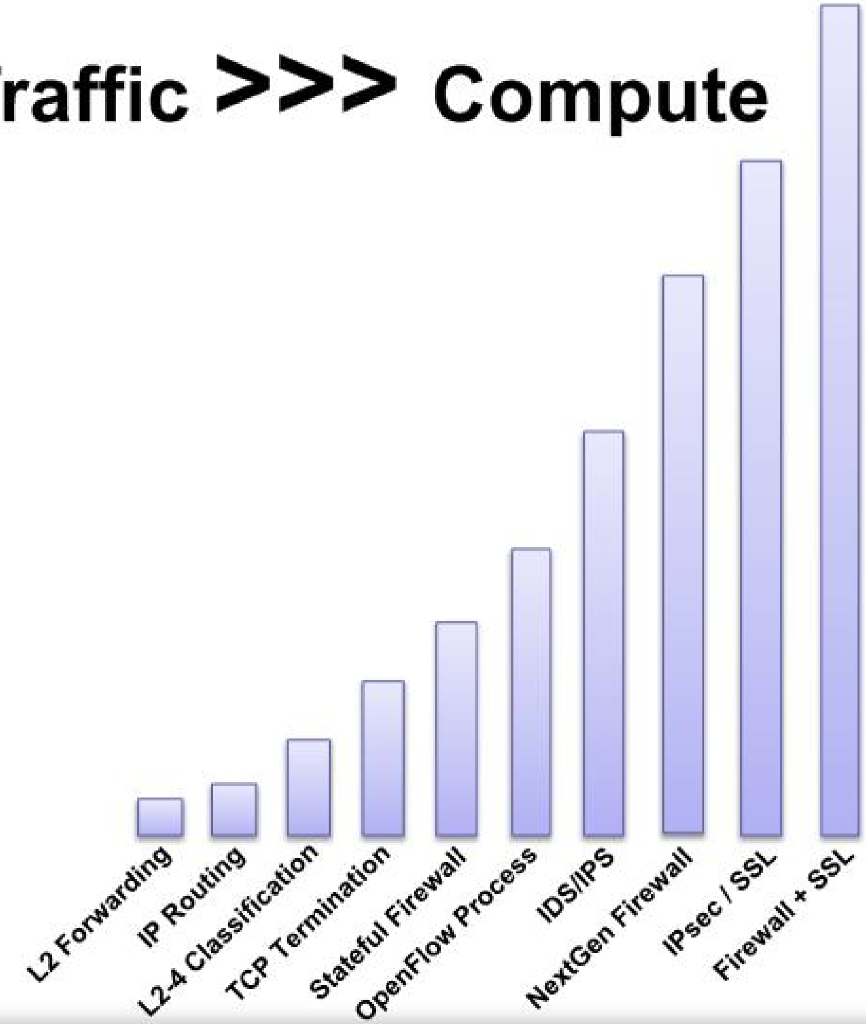

Source: Michael Zimmerman, Tilera

Management and Orchestration

- The key: Elasticity!
  - Pieces at all infrastructure layer
  - Need to go beyond to just fit them together
  - Multi-technology support, and open interfaces

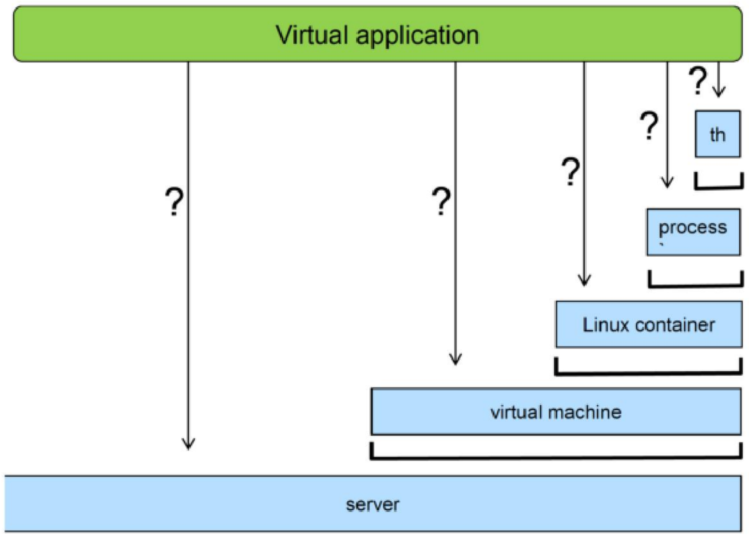

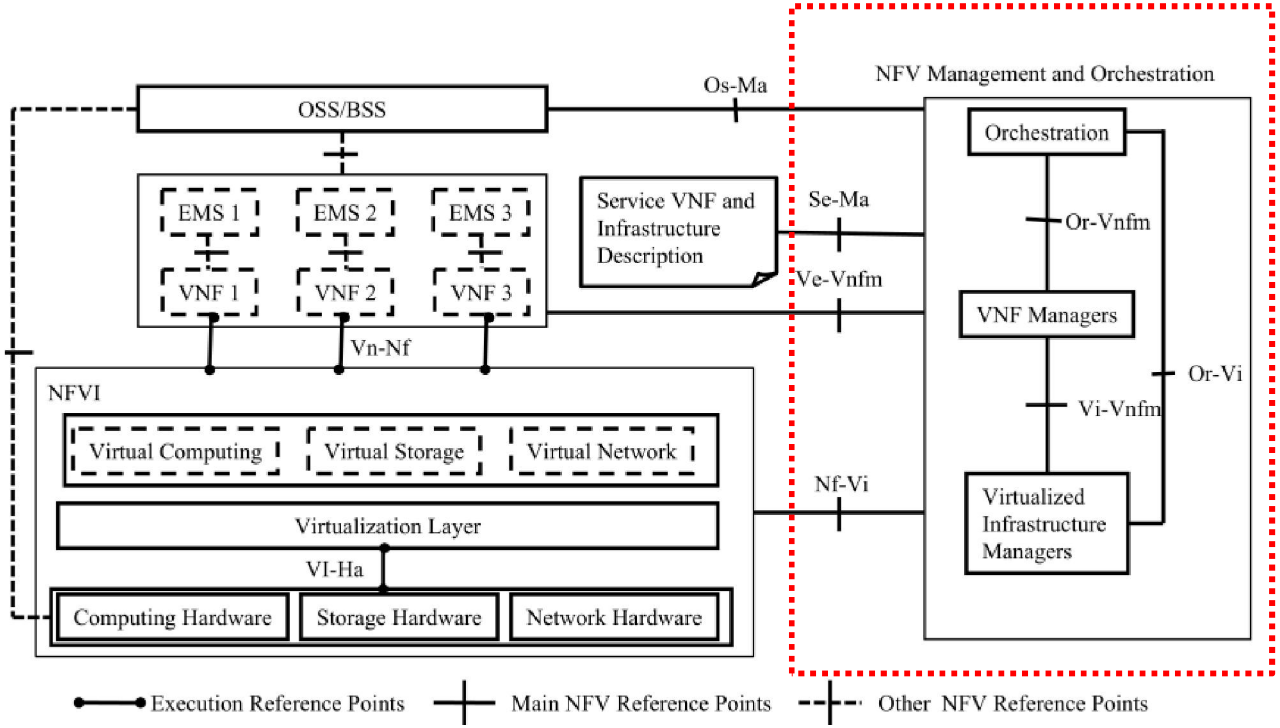

Source: Raj Jain

Orchestration

- An end-to-end perspective
  - May include nested forwarding graphs

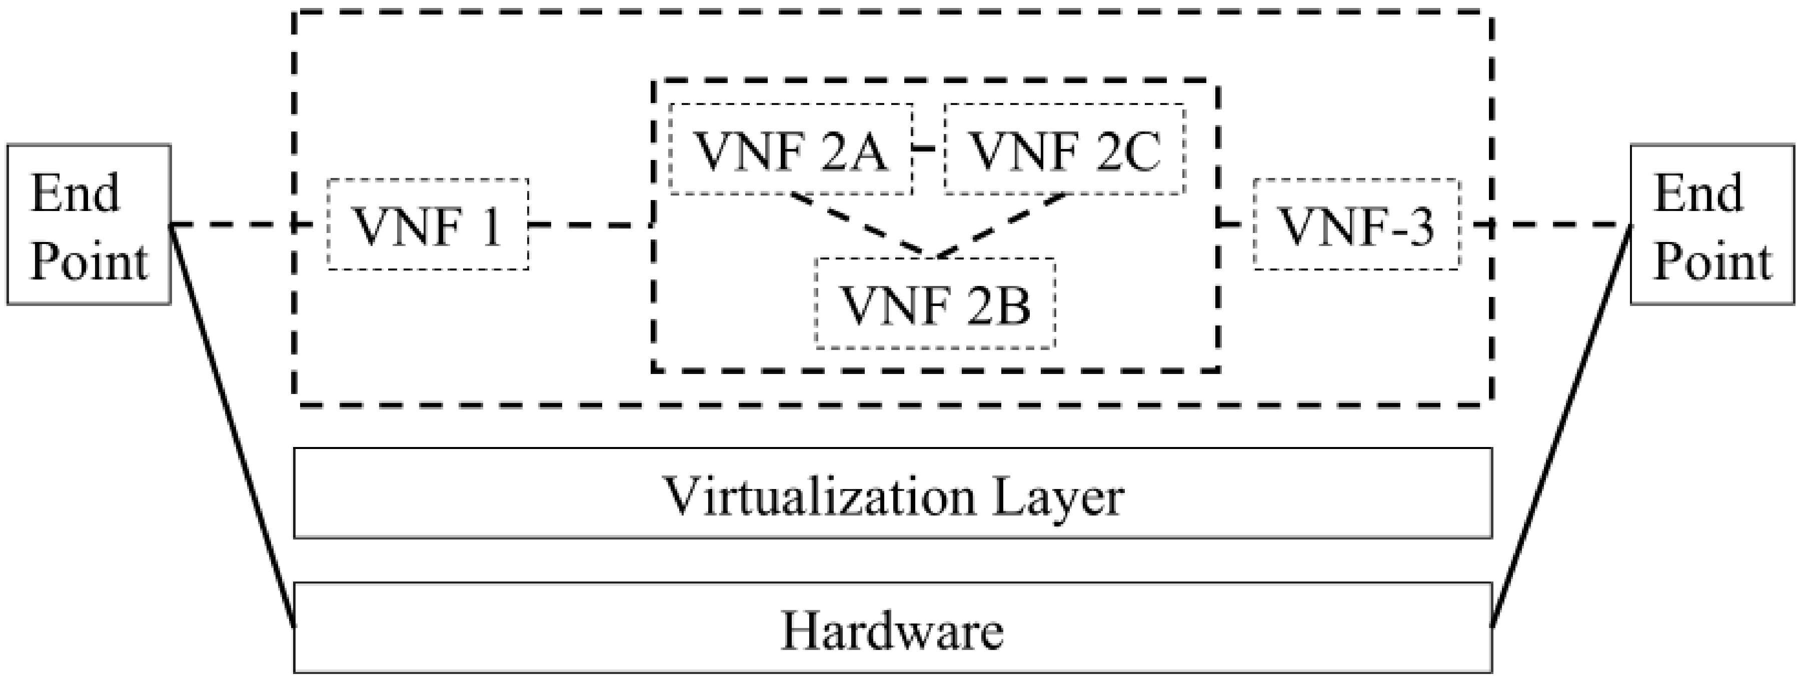

Source: ETSI NFV Framework

Orchestration

- Automated deployment of NFV applications
  - OpenStack, CloudStack...
- NFVI profile for NF
  - Select and start host, VM
- Applications (NFs)
  - Service address
  - Location specific configuration

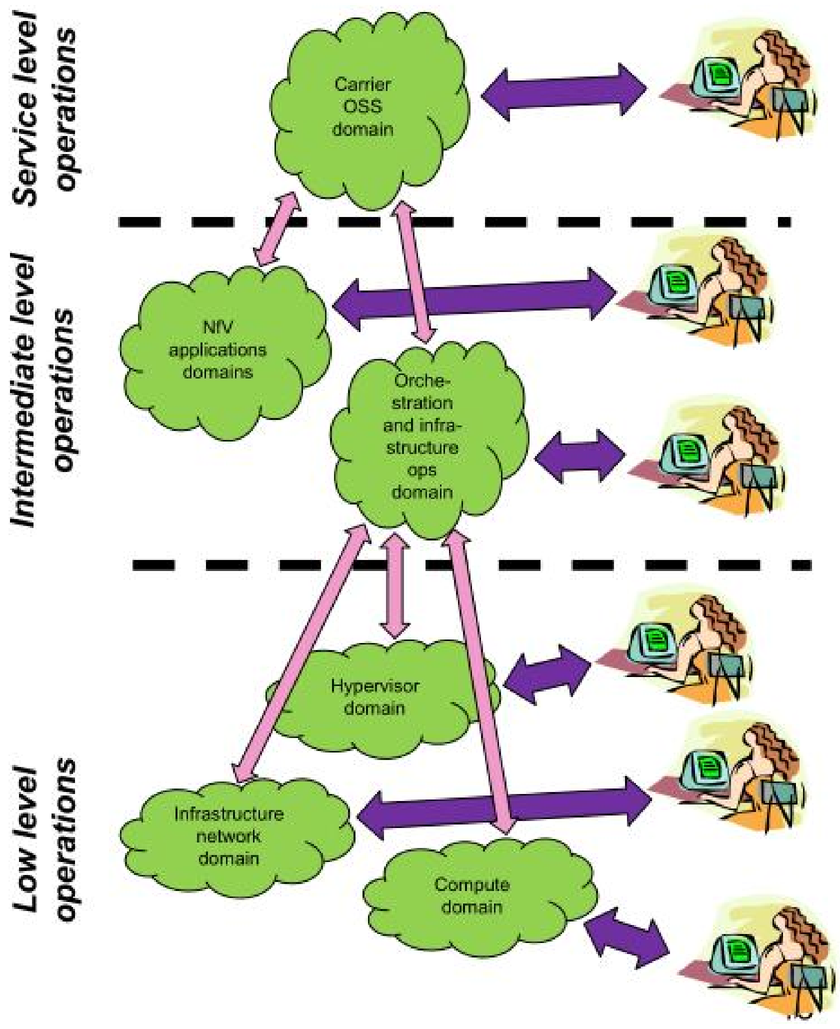

Source: Bob Briscoe, BT

Overall Management & Orchestration

- Control functions and state in all network levels
  - Heterogeneous environments and services

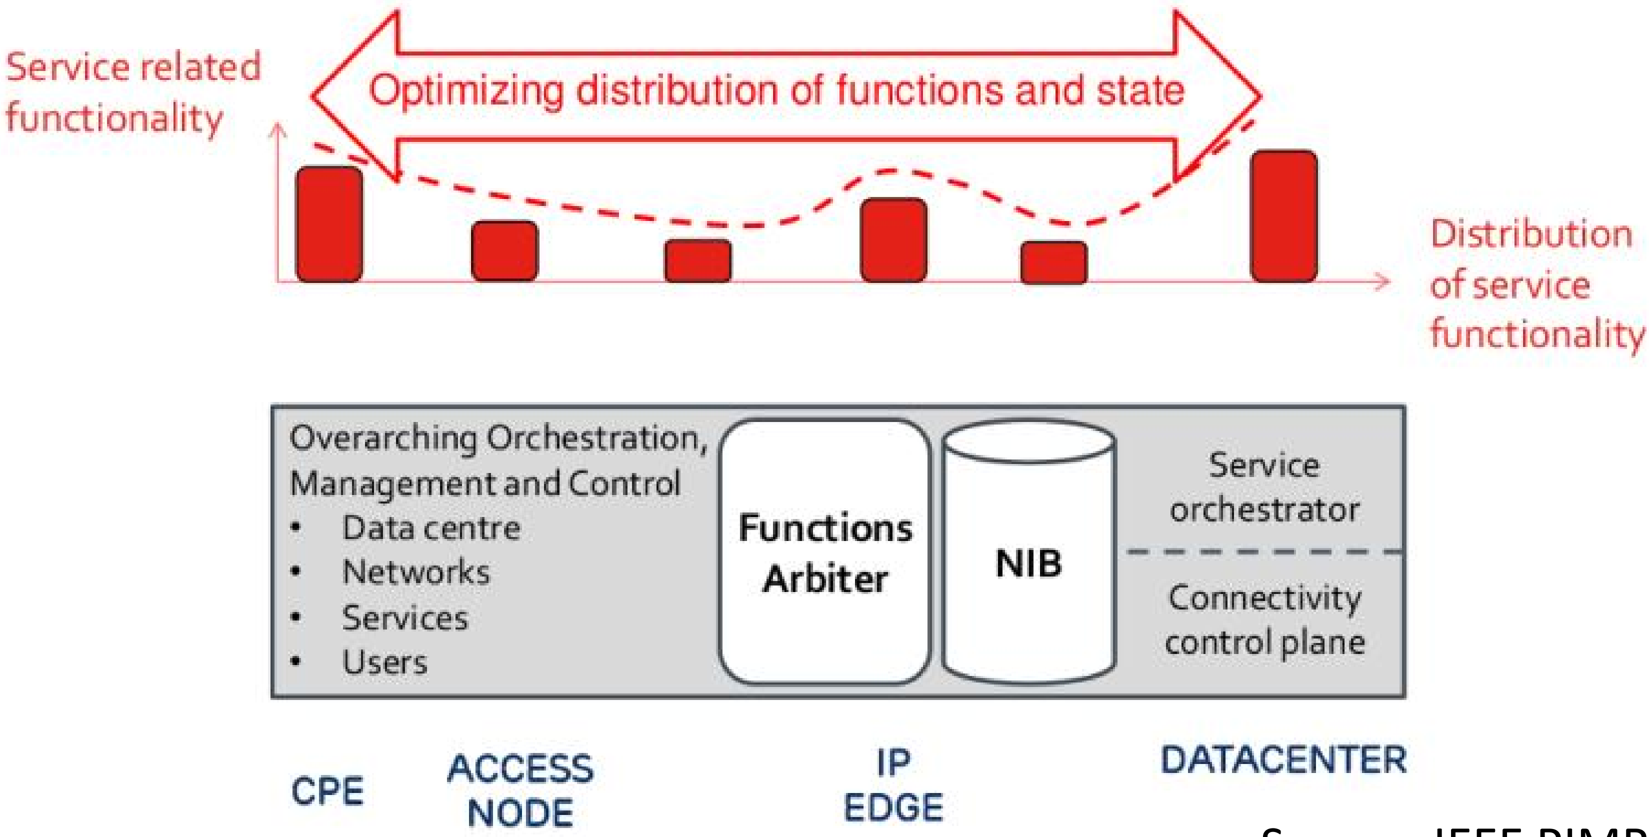

Source: IEEE PIMRC, 2013

Portability

- Move VNF across N-PoPs
- Decoupled NFV framework from NFVI
- Optimize VNF resources:
  - Location
  - Allocation
  - Reservation
- Compatibility
  - Integration/internetworkin g
  - Meeting SLA requirements
- Example: NfV hypervisors

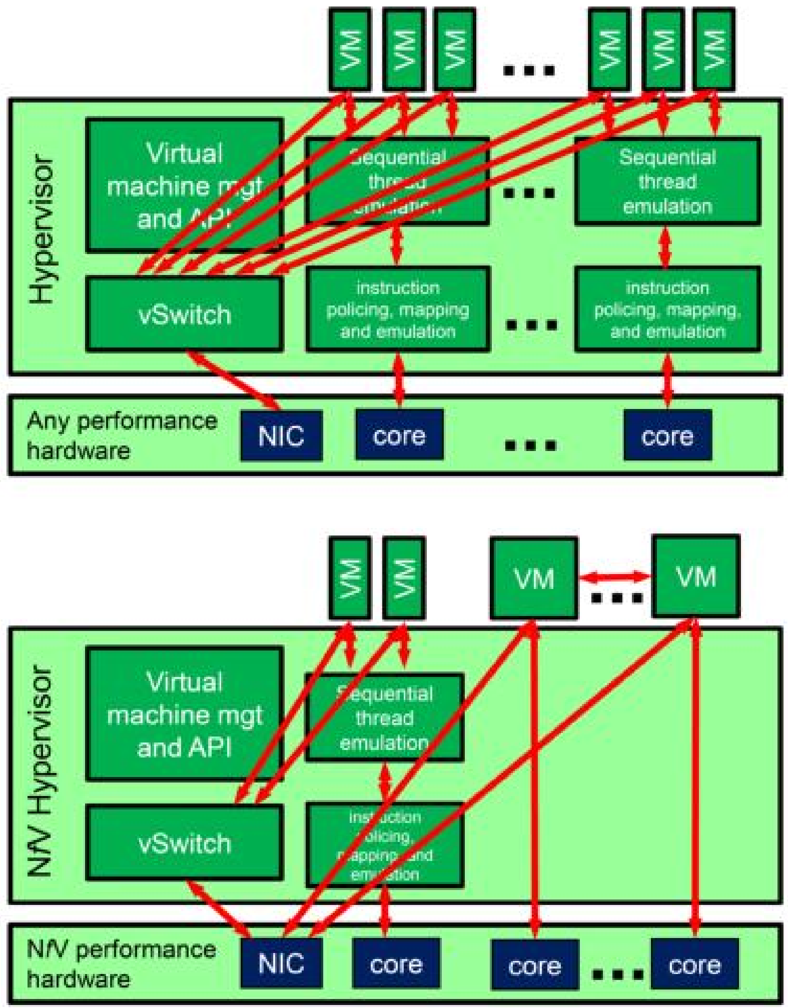

Source: Bob Briscoe, BT

Interoperability and Legacy Networks

- End-to-end network services
  - Transparent management and orchestration
- No place for one-size-fits-all solutions
  - Dynamic and heterogeneous new technologies
- Handle different old and new characteristics
  - Impact on the other requirements:
    - Performance, resilience, security...
- Maintain SLAs
- Avoid disruptions!

Security

- New Threats
  - Virtualization Network Layer
  - Several identity layers and accounting
- Protection of interfaces exposed by NFV architecture principles.
- Secure separation and management of NF entities.
- Heterogeneous network domains
- NFVI shared resources
  - Isolation of VNF sets
  - User privilege resources access (APIs)
- Mechanisms:
  - Control and verify the configuration of soft/hardware

Resilience

- Different Levels
  - PFs, NFs, NFVI, NFV-FG
- Monitoring, synchronisation and trigger mechanisms in the event of failure of NFs
- Correlated failures in NFV-FG
  - Chained resilience plans
- Service Continuity
  - SLA minimum insurance
    - Zero impact vs. Measurable impact
- Orchestration: NOT a single point of failure

SDN

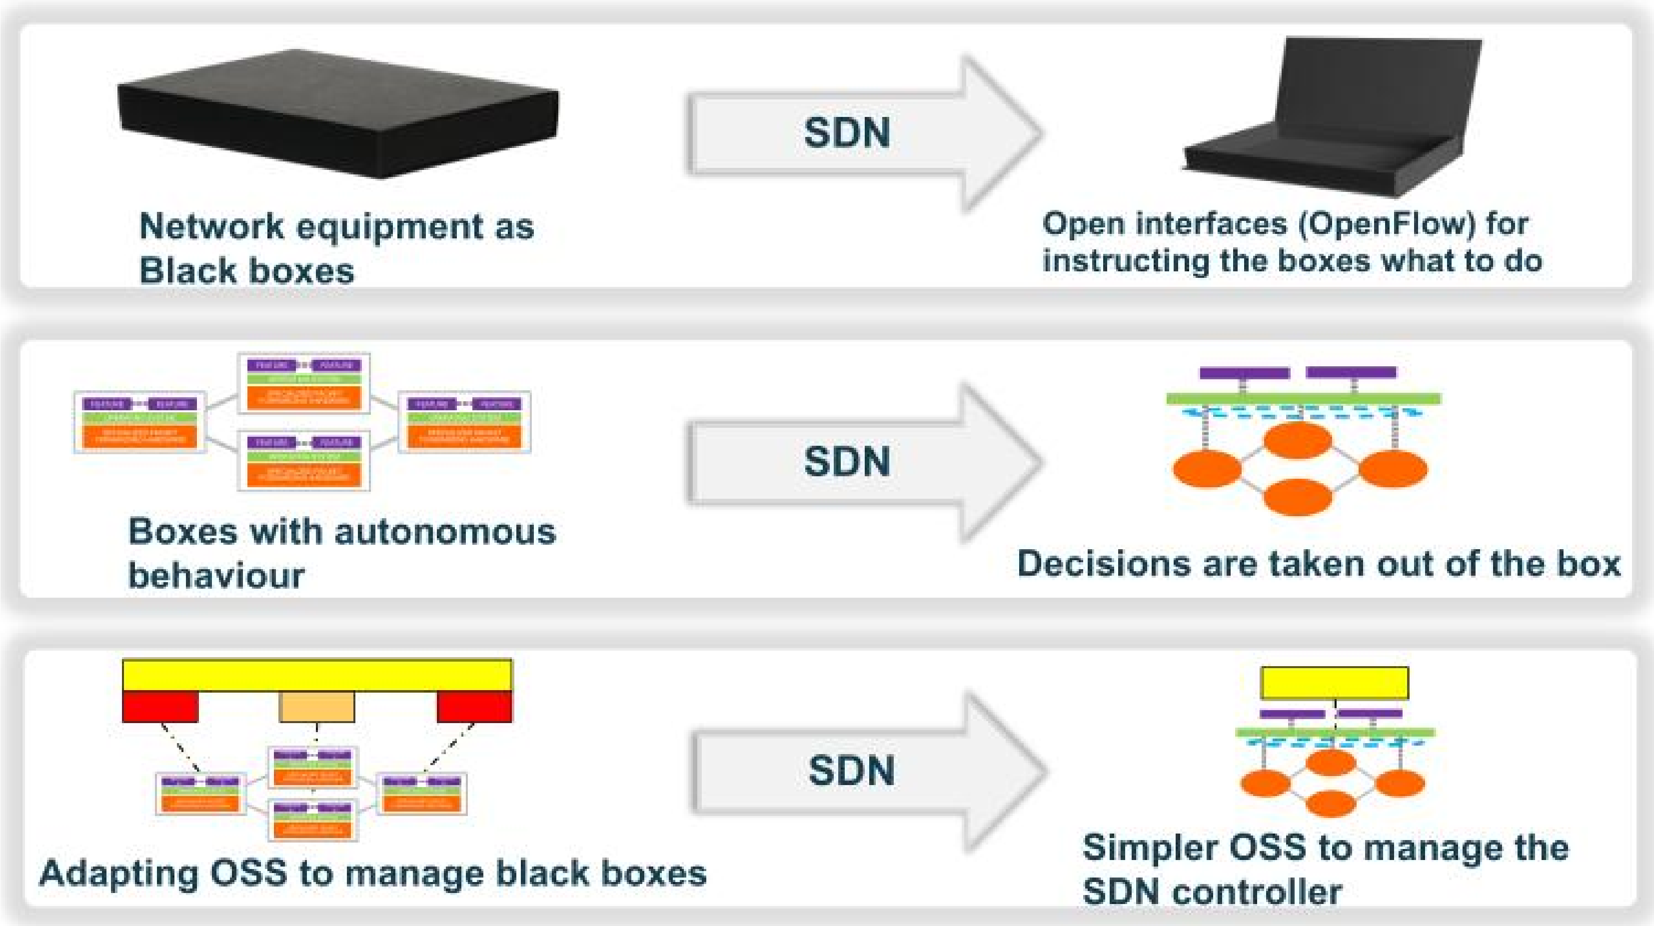

Source: Diego Lopez, Telefonica

SDN and NFV

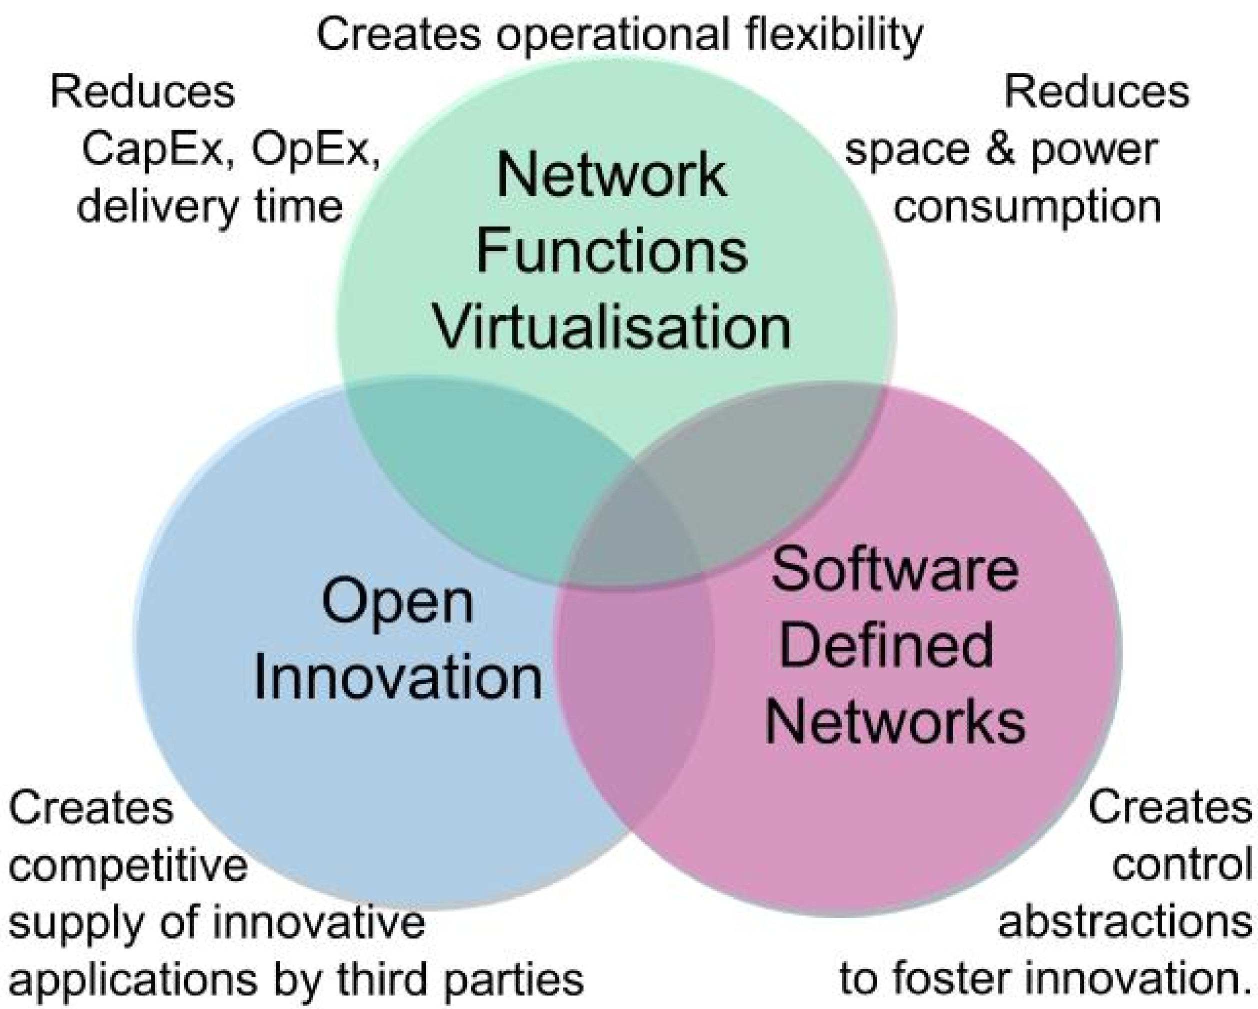

Source: Bob Briscoe, BT

SDN and NFV

- SDN and NFV do NOT depend on each other

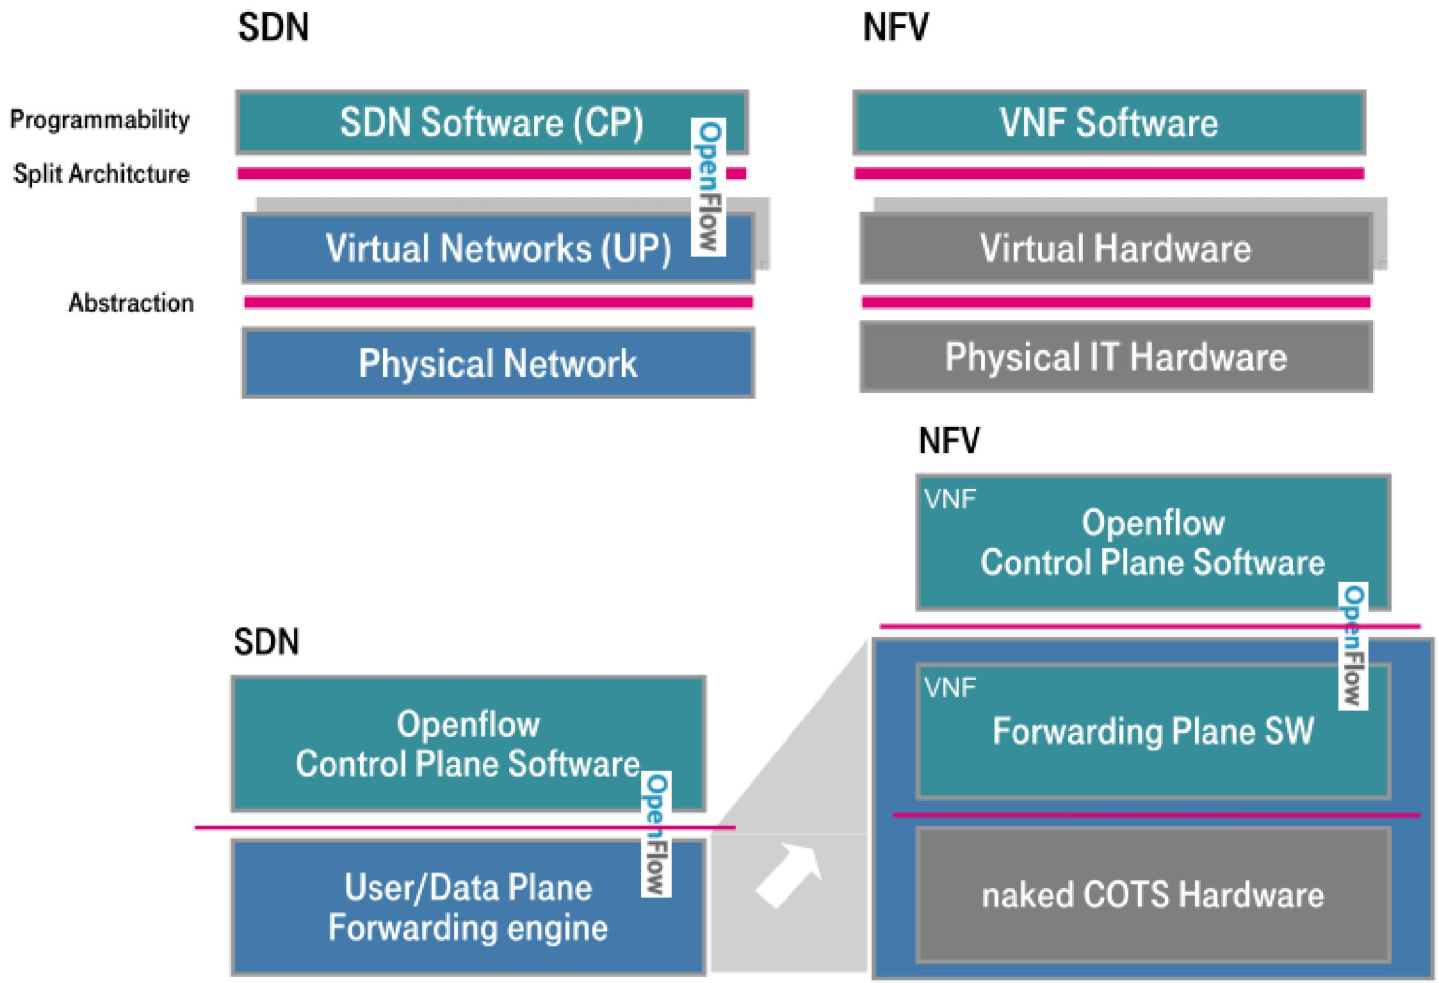

Source: Uwe Michel, T-Systems

SDN and NFV

- SDN poses to NFV:
  - Central point of contact / Orchestrate VNFs (NSC)

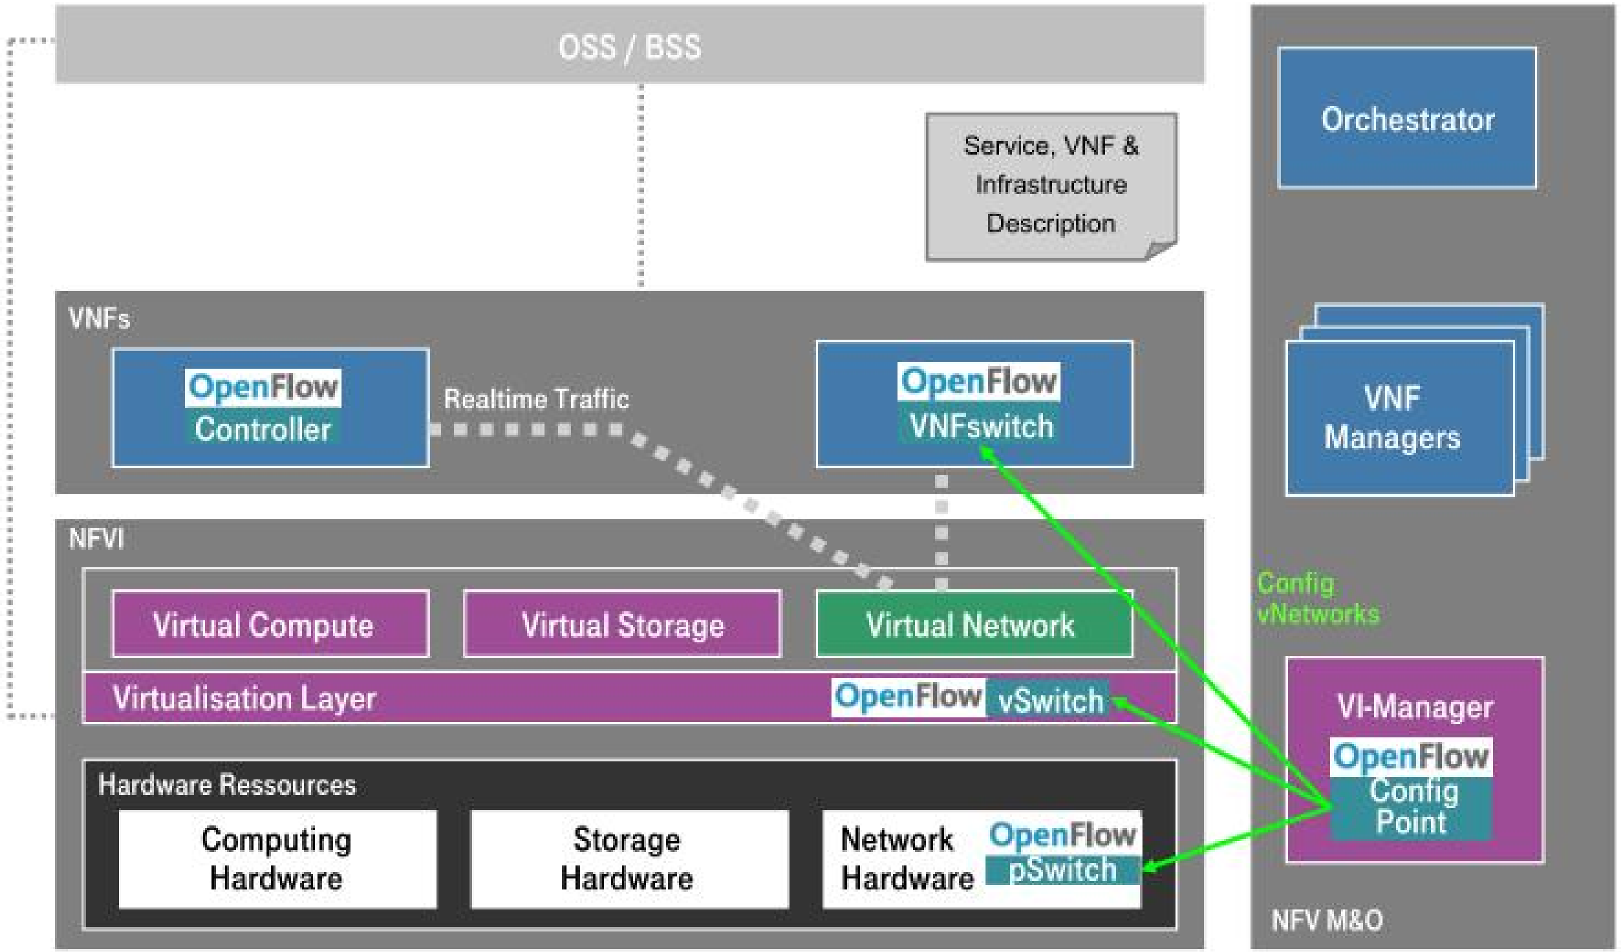

Source:
Uwe Michel
T-Systems

Scope of NFV and OpenFlow/SDN

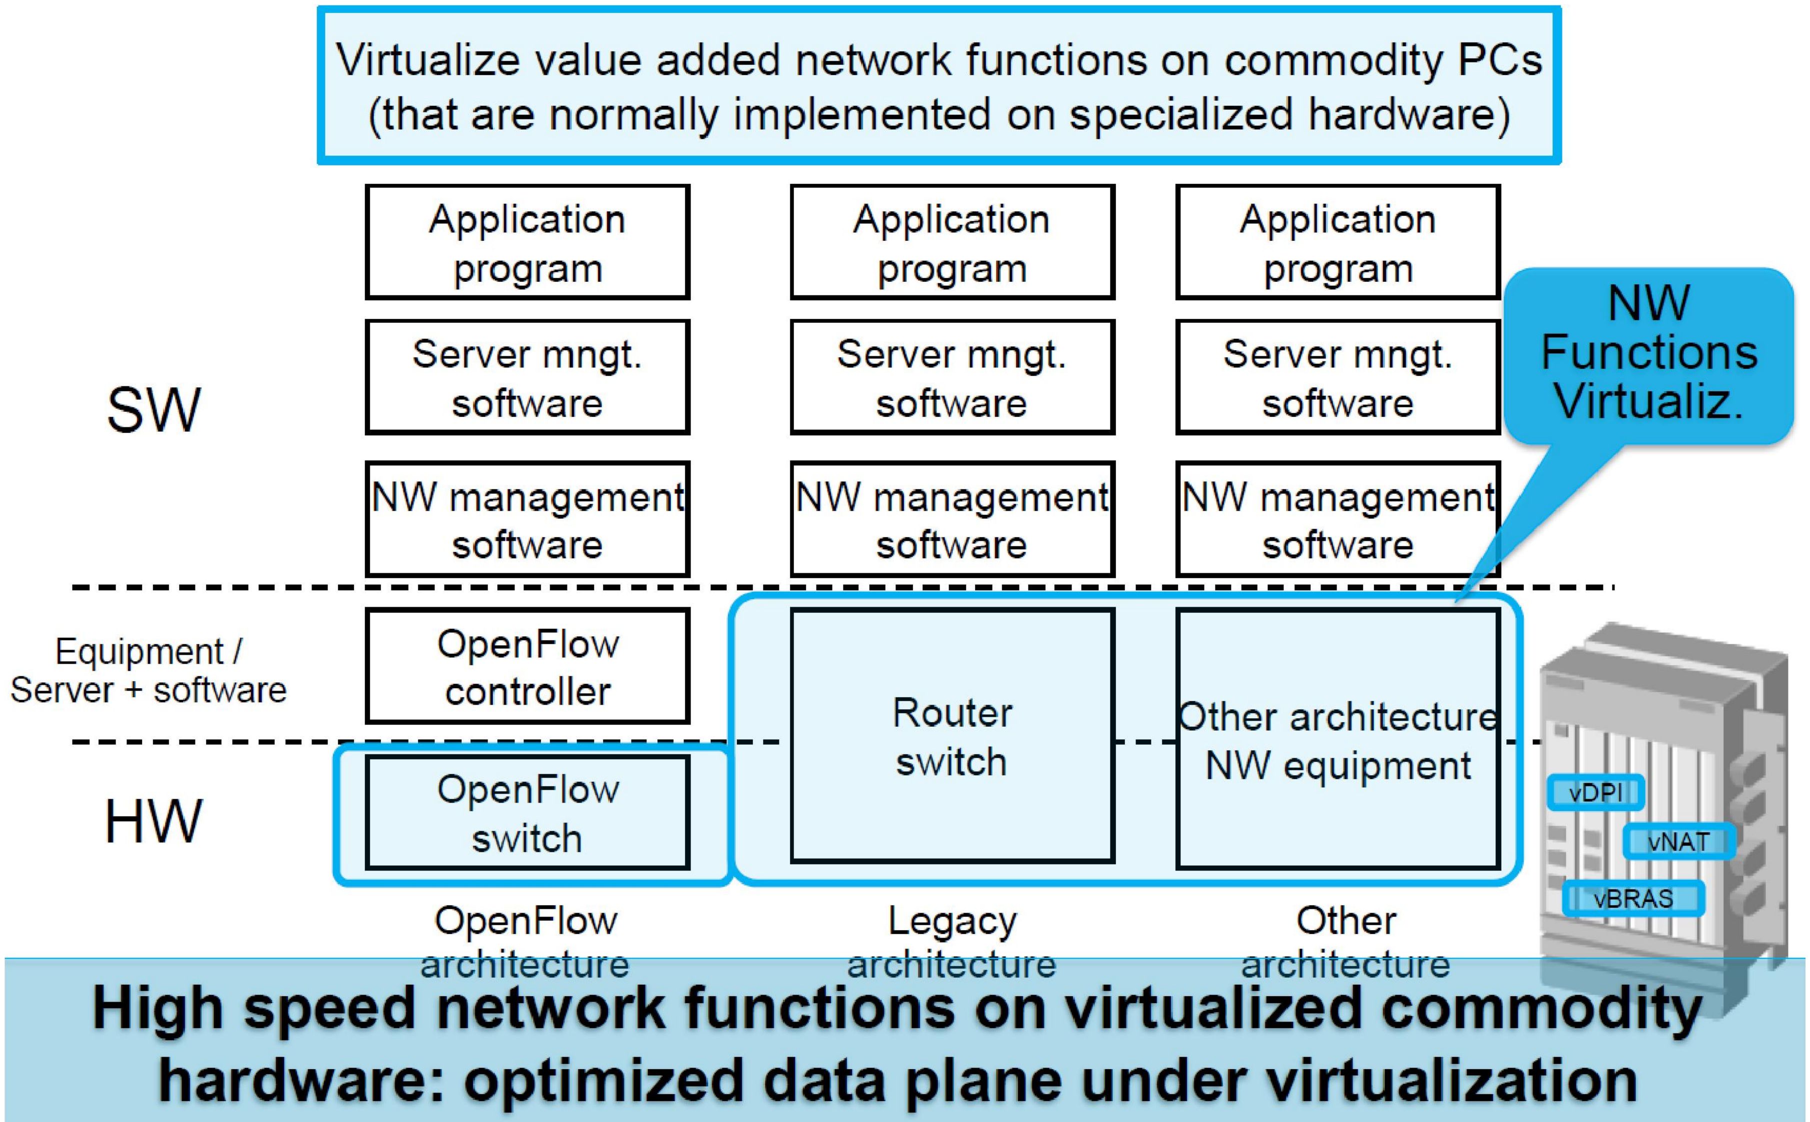

Source: NEC

Networking with SDN & NFV

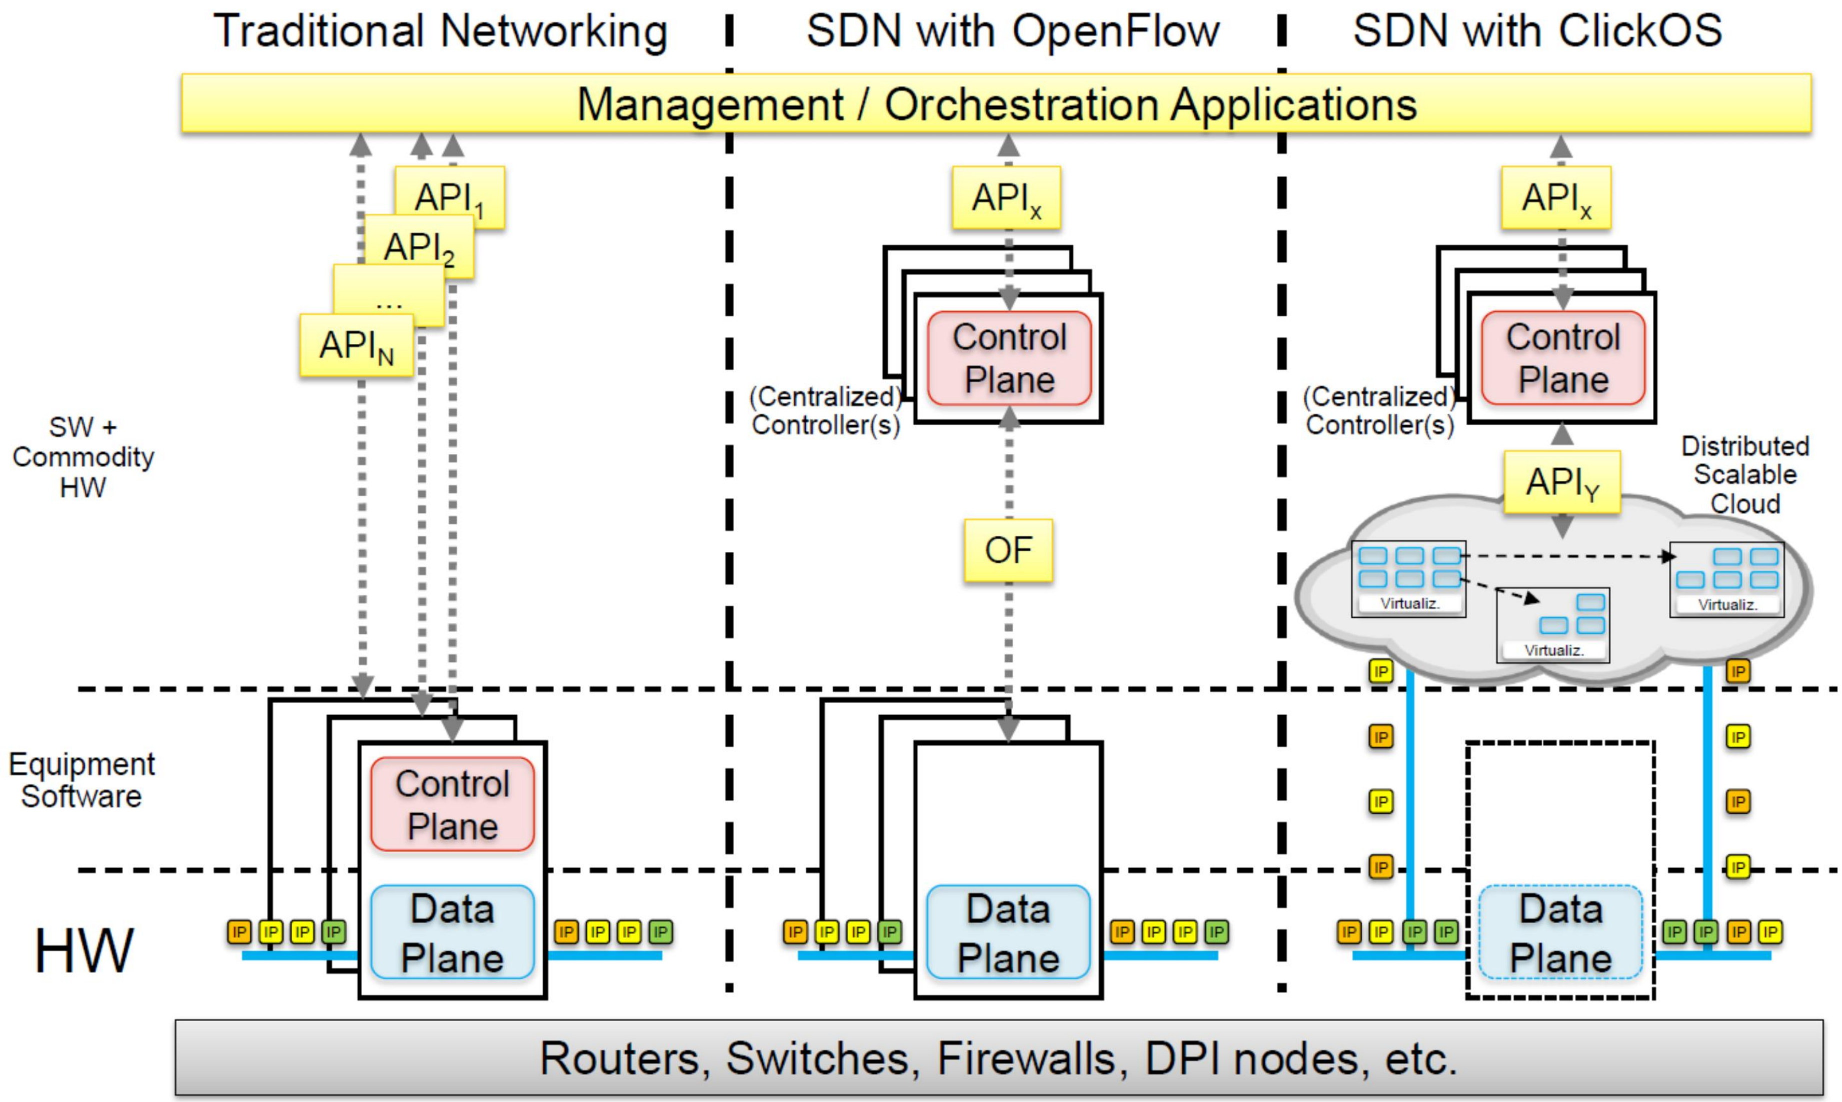

Source: NEC

NFV Challenges for Networking Research

NFV Resiliency

- NFV-based service continuity.
- Coexistence of virtualised and non-virtualised Network Functions (NFs)
- Virtual Network Functions (VNF) Software (VM, Hypervisor) failure or congestion protection.
- Monitoring, synchronisation and trigger mechanisms in the event of failure of NFs.

NFV Control & Orchestration

- Providing automation and elasticity.
- NF Instance instantiation, scaling and migration.
- End-to-end service setup, operation and monitoring.
- Multi-technology support, and open interfaces.

NFV Security

- Securing VNF instances.
- Vulnerabilities introduced in the new virtualisation layer.
- Protection of interfaces exposed by NFV architecture principles.
- Secure separation and management of NF entities.

Source: D. King, Comnnet Workshop

Use Cases

NFV

NFV ISG Use Cases

- First use case proposal: 2010
- Main idea: **contribute to thrive NFV**
  - Real Scenarios
- Fast service inovation based on software and operational end-to-end NFs
  - **Operational eficiency**
  - **Energy consumption reduce (workloads migration)**
  - **Open and standard interfaces**
  - **Flexibility between VNF and hardware**;
  - **Eficient revenues return**

Use Cases Matrix

Cloud Use Cases	Mobile Use Cases	Data Center Use Cases	Access/ Residen-tial
NFV Infrastructure as a Service (IaaS) (NFVIaaS)	Virtualization of the Mobile Core Network and IMS	Virtualization of CDNs	Virtualization of the Home Environ-ment
Virtual Network Functions as a Ser-vice (VNFaaS)	Virtualization of Mobile Base Sta-tion		Fixed Access Net-work Functions Virtualization
Service Chairs (VNF Forwarding Graphs)			
Virtual Network Platform as a Ser-vice (VNPaaS)			

Use Case Matrix – 4 big horizontal themes, and 9 use cases

NFV Infrastructure as a Service (NFVIaaS)

- Cloud Computing Services are typically offered to consumers in one of three service models
  - Infrastructure as a Service (IaaS)
  - Platform as a Service (PaaS)
  - Software as a Service (SaaS)
- **IaaS** is defined as the capability to offer to consumers processing, storage and fundamental computing resources
- Some literature also refers to a capability to offer network connectivity services as **Network as a Service (NaaS).** One application for NaaS appears to be the on demand creation of network connectivity between Cloud Service Provider and Customer

NFV Infrastructure as a Service (NFVIaaS)

NFV Infrastructure :

- provide the capability or functionality of providing an environment in which Virtualized network functions (VNF) can execute

- **NFVIaaS** provides compute capabilities comparable to an **IaaS cloud computing service** as a run time execution environment **as well as support the dynamic network connectivity services** that may be considered as comparable to **NaaS**

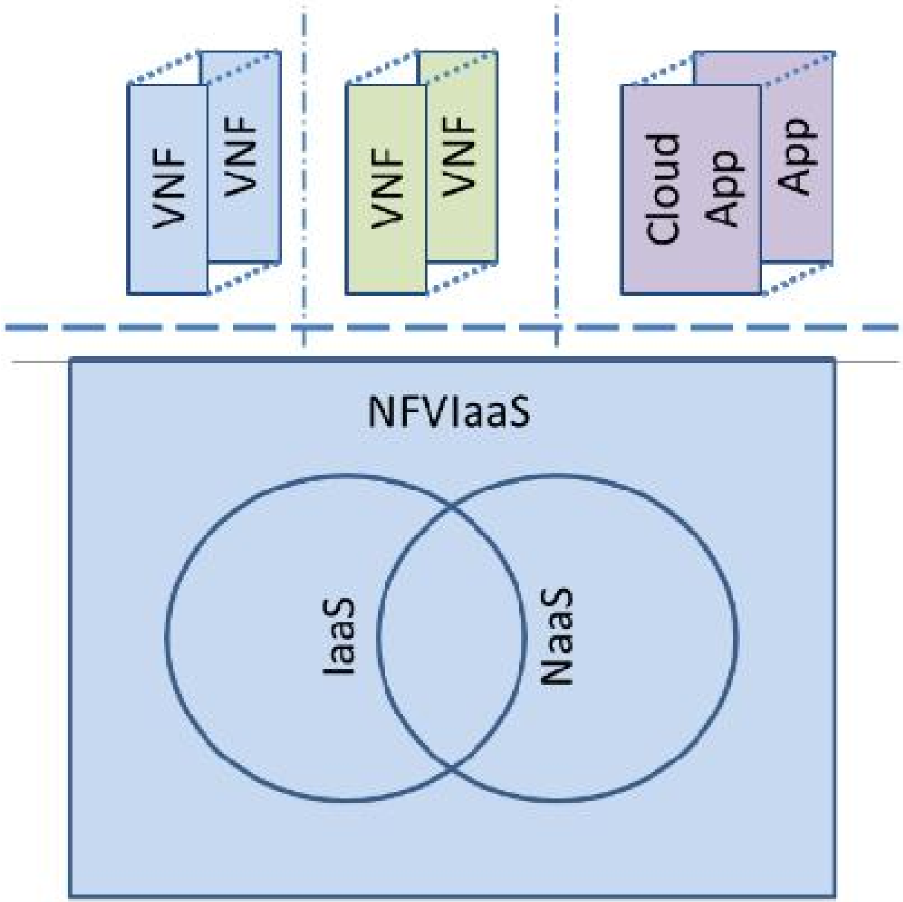

NFVIaaS: Multi-domain Example

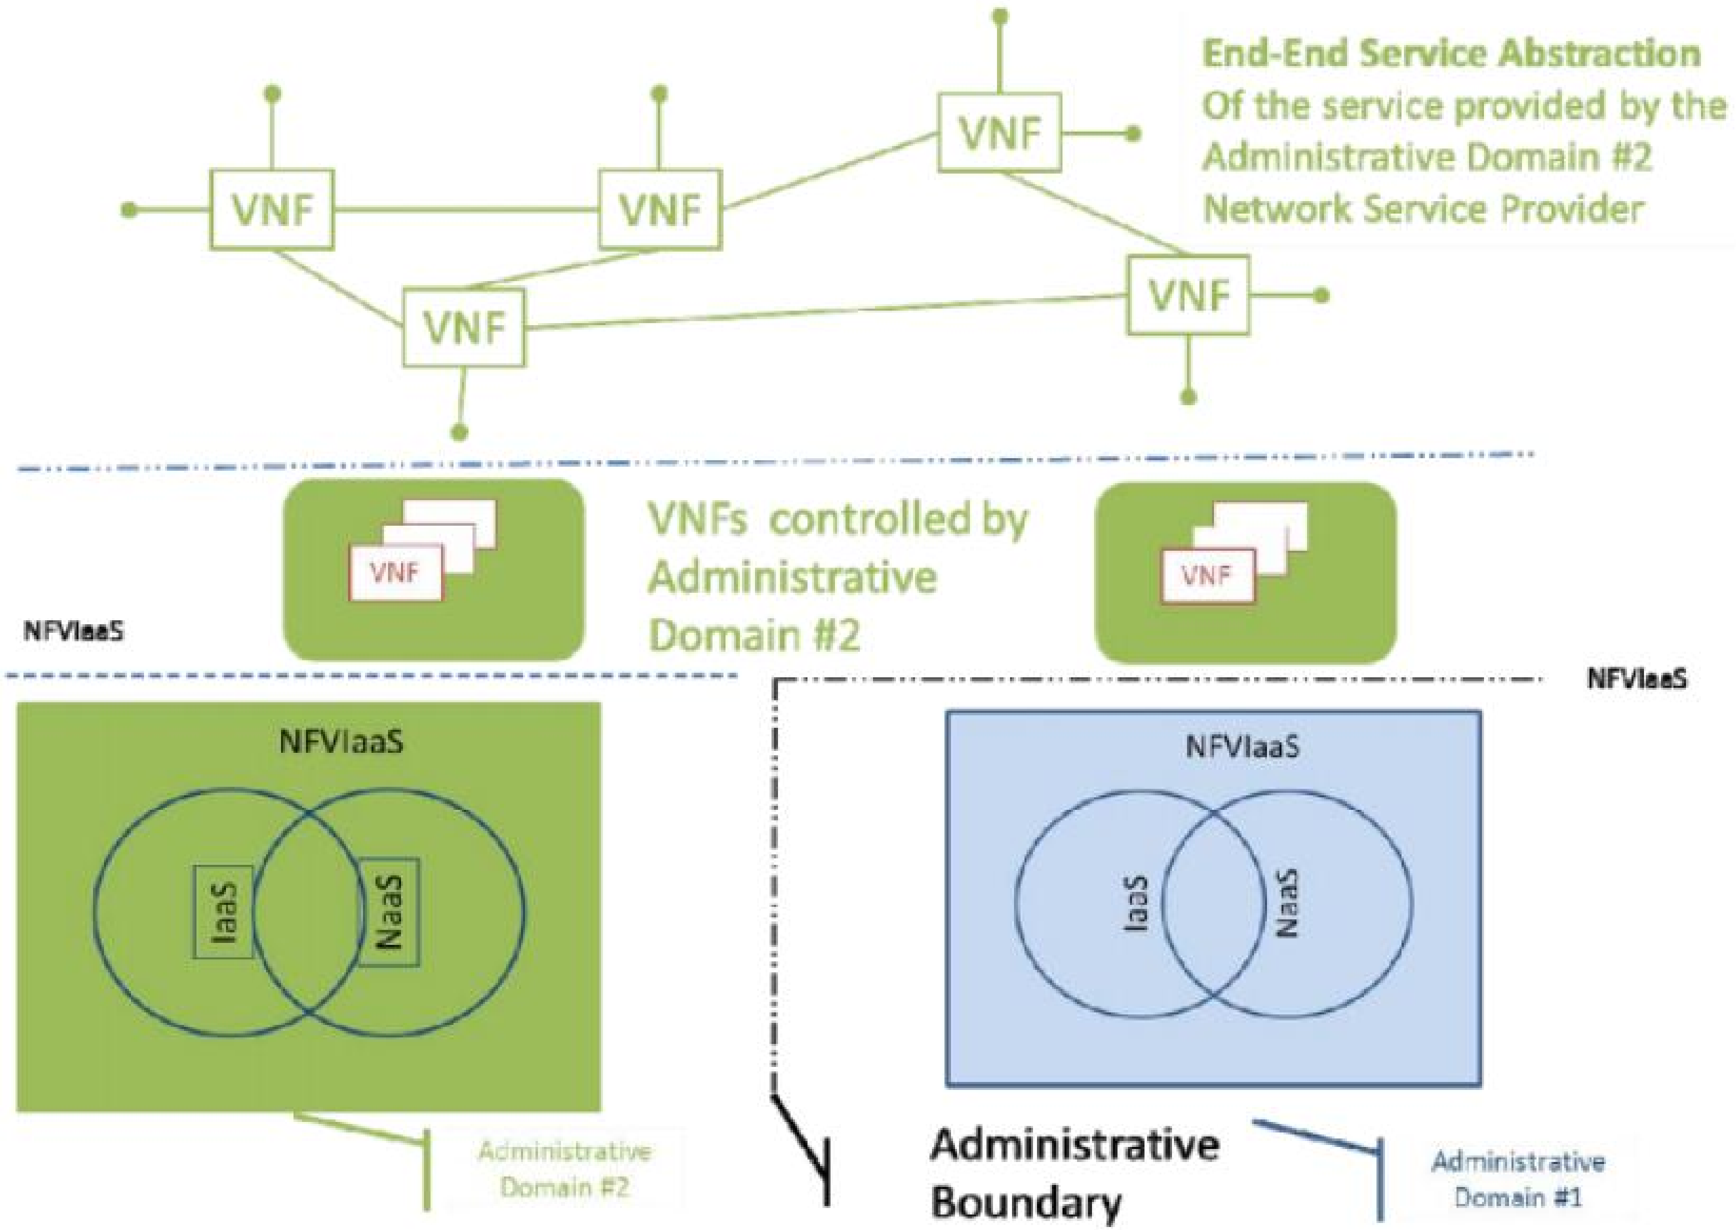

VNFaaS Motivation: CPE e PE

Pre-NFV service provider networks include a Provider Edge (PE) router at the edge of the core, facing the Customer Premises Equipment (CPE) device

Virtual Network Functions as a Service (VNFaaS)

- Substantial saving may be possible by **moving routing functionality from purpose-built routers to equivalent functionality** implemented in COTS hardware environments providing cloud computing capabilities such as the NFVI

- Rather than the Enterprise **investing its own capital in deployment of networking infrastructure**, the service provider may be able to **provide advanced networking features as a measured service**

- The service provider could operate a **VNF instance using its NFVI which provides the functionality required to implement the enterprise CPE** and potentially another VNF instance for the control plane of the PE router improving its scalability

VNFaaS

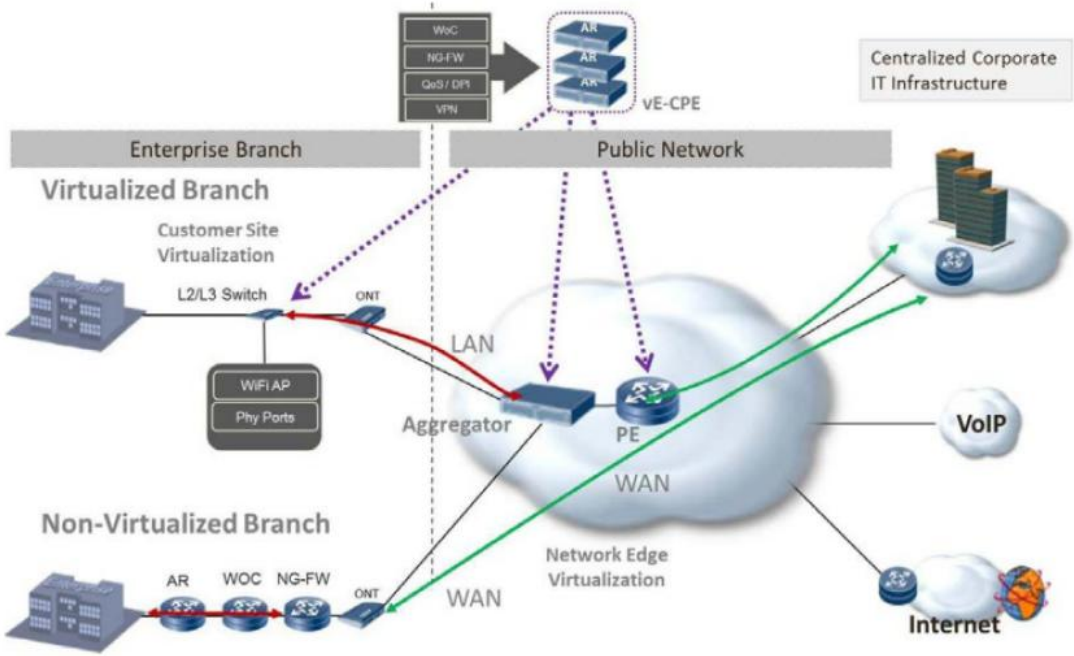

Physical CPE & vE-CPE
(routing, VPN termination, QoS support, DPI, NG-FW and a WOC (WAN Optimization Controller)

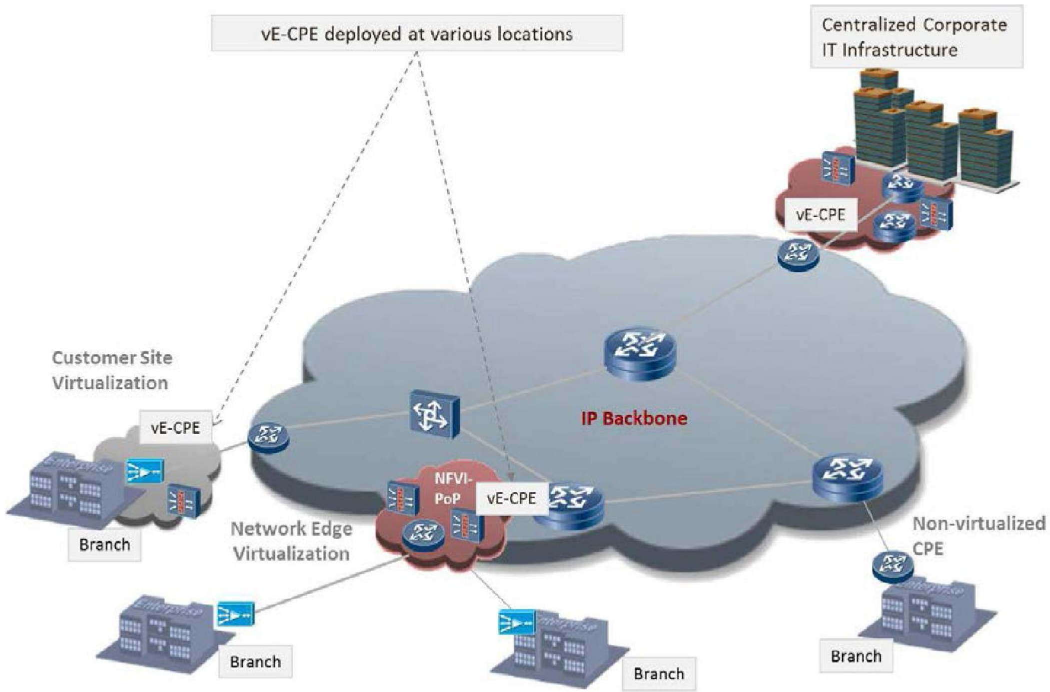

vCPE functionality in many locations

Virtual Network Platform as a Service (VNPaaS)

- Network resources are more and more **often not exclusively used** by the operator

- **Platform as a Service** (PaaS) as the possibility for the **consumer to deploy his own applications** using the computing platform supported by the provider

- Service Provider provides a **toolkit of networking and computing infrastructure as well as potentially some VNFs** as a platform for the creation of virtual network
  i.e. a **Virtual Network Platform as a Service**

VNPaaS

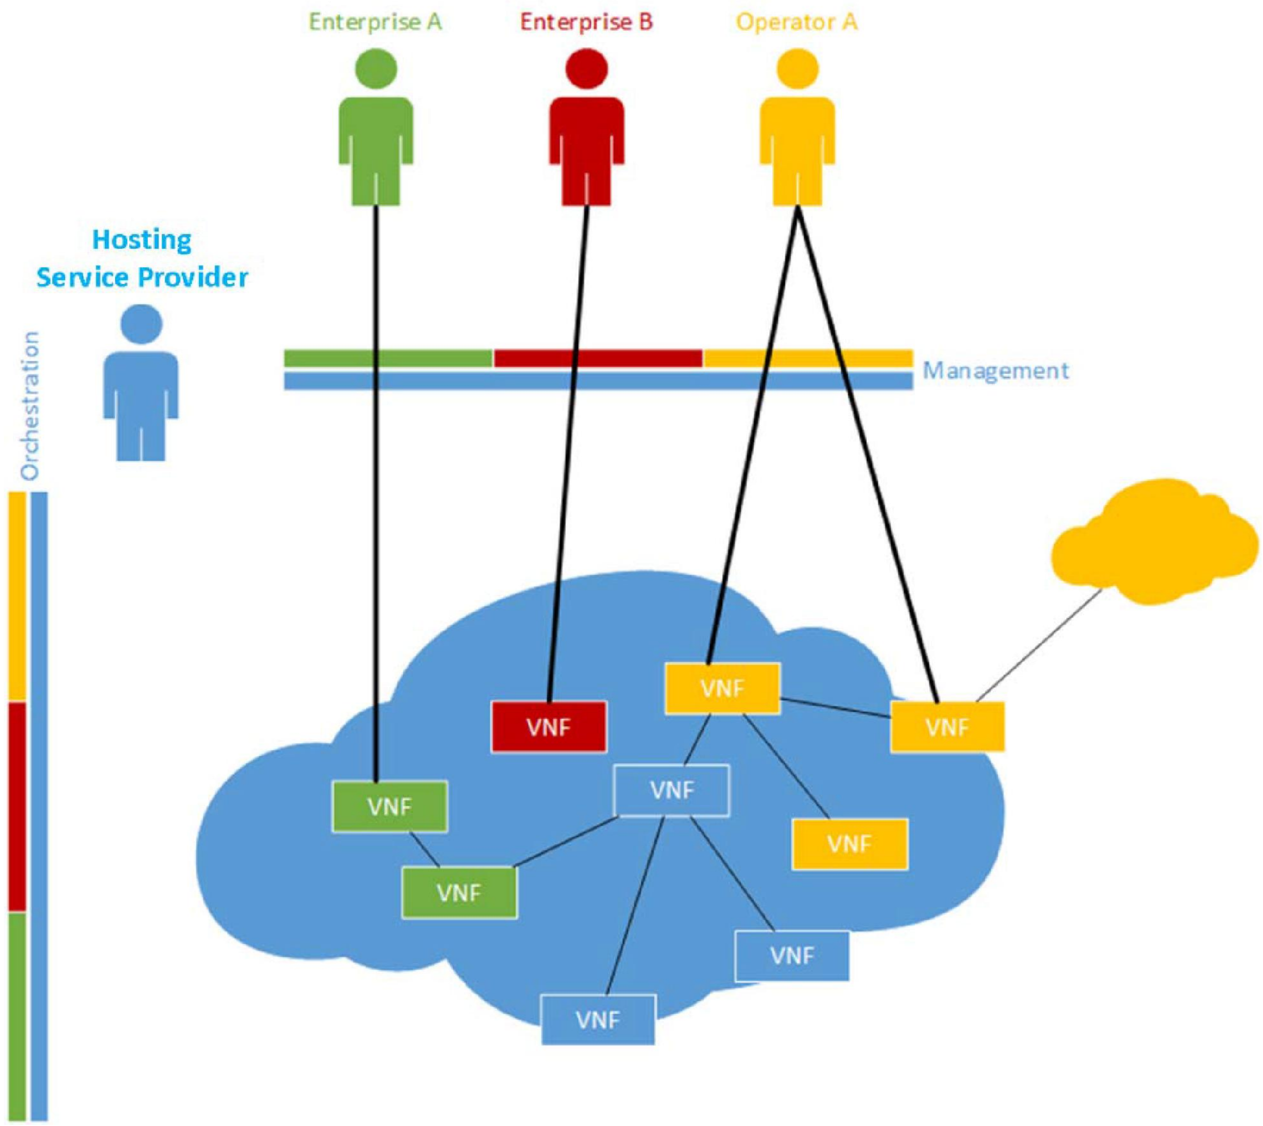

The **VNPaaS is similar to the VNFaaS**, but differs mainly in the scale of the service and programmability

VNPaaS provides a **larger scale service** typically providing a **virtual network** rather than a **single virtual network function**.

Service Chains (VNF Forwarding Graphs)

- VNF FGs are the analogue of connecting existing Physical Appliances via cables as described in the NFV

- Cables are bidirectional and so are most data networking technologies that will be used in Virtualised deployments in the near term (e.g. Ethernet). In other words, a VNF Forwarding Graph provides the logical connectivity between virtual appliances (i.e. VNFs).

VNF FG

VNF FG Logical View

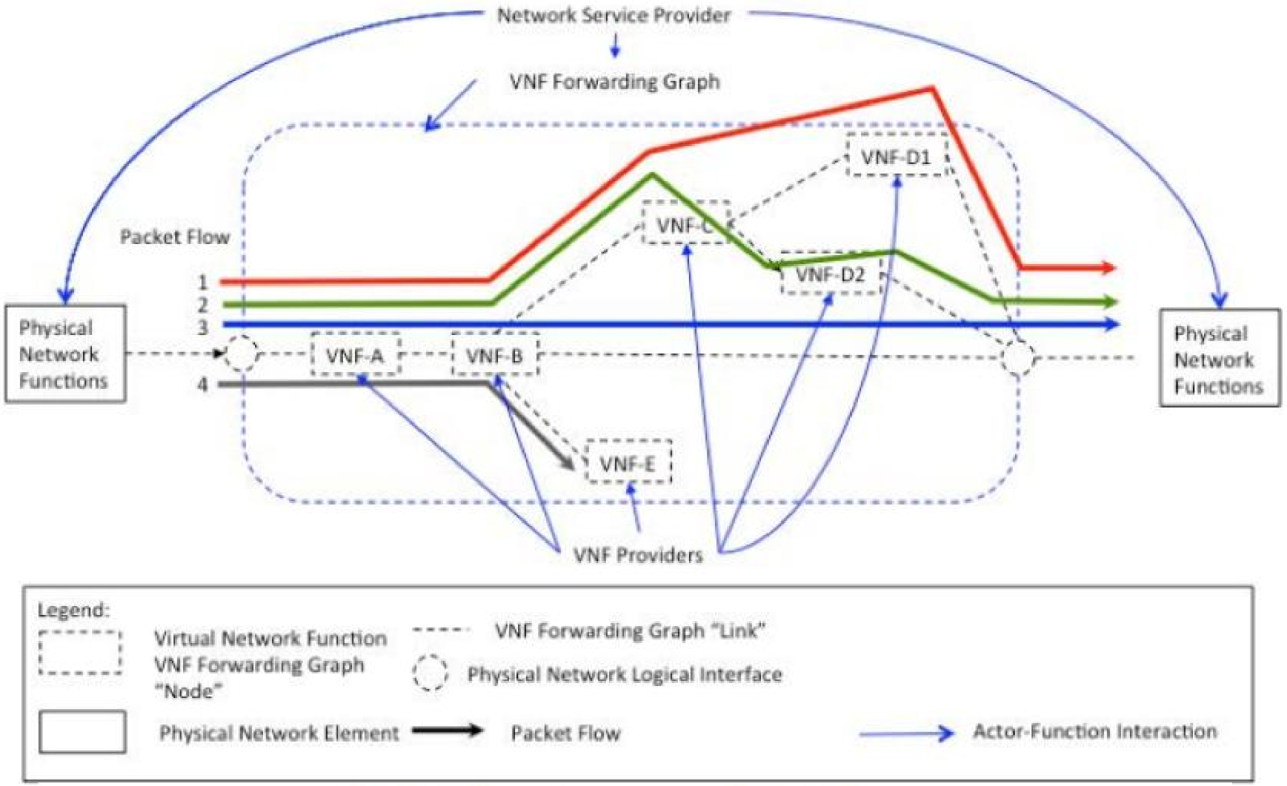

VNF FG Physical View

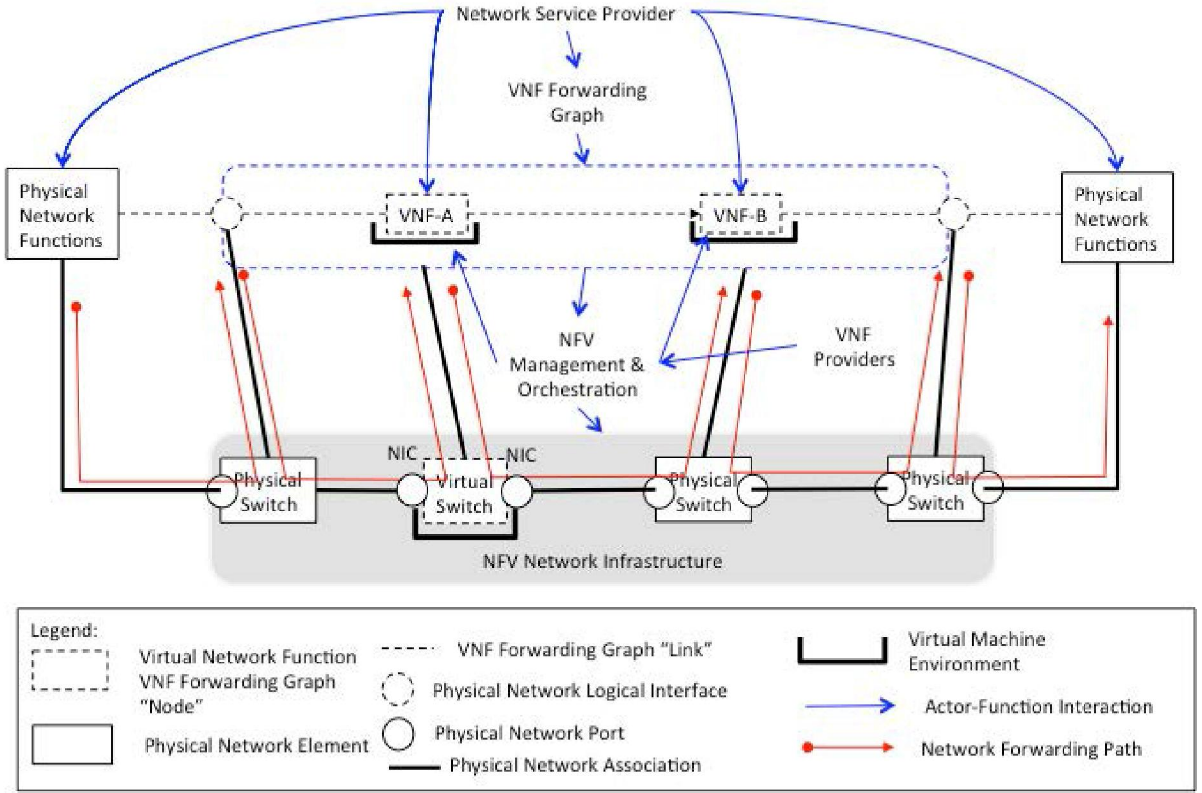

Appliances & VNF-FG

Atributo	Equipamentos Físicos	VNF FG
Eficiência	Funcionalidade e capacidade dimensionada para o pico de carga	Funcionalidade e capacidade dinamicamente dimensionados pela utilização corrente
Resiliência	Backups utilizam hardware específico e redes dedicadas	Backups podem compartilhar hardware e recursos de rede
Flexibilidade	Lentidão em ativações de novas capacidades quando baseado em hardware	Pequeno tempo para atualizações ou ativações de novas funcionalidades por ser baseado em software
Complexidade	Necessário adicionar novas configurações, interfaces físicas e/ou suportar novos sistemas para implementar novos grafos de encaminhamento	Virtualização da função de encaminhamento e/ou configuração de VNF implementam grafos de encaminhamento mais facilmente
Ativação	Ativação em outras operadoras requer equipamentos físicos, interfaces e configurações para conectar usuários finais	Virtualização de funções e comutação torna o processo de ativação mais simples

Mobile Core Network and IMS

- **Mobile networks are populated with a large variety of proprietary hardware** appliances
- **Flexible allocation** of Network Functions on such **hardware resource pool** could highly improve network usage efficiency
- **Accommodate increased demand for particular services** (e.g. voice) without fully relying on the call restriction control mechanisms in a **large-scale natural disaster scenario** such as the Great East Japan Earthquake

V-EPC

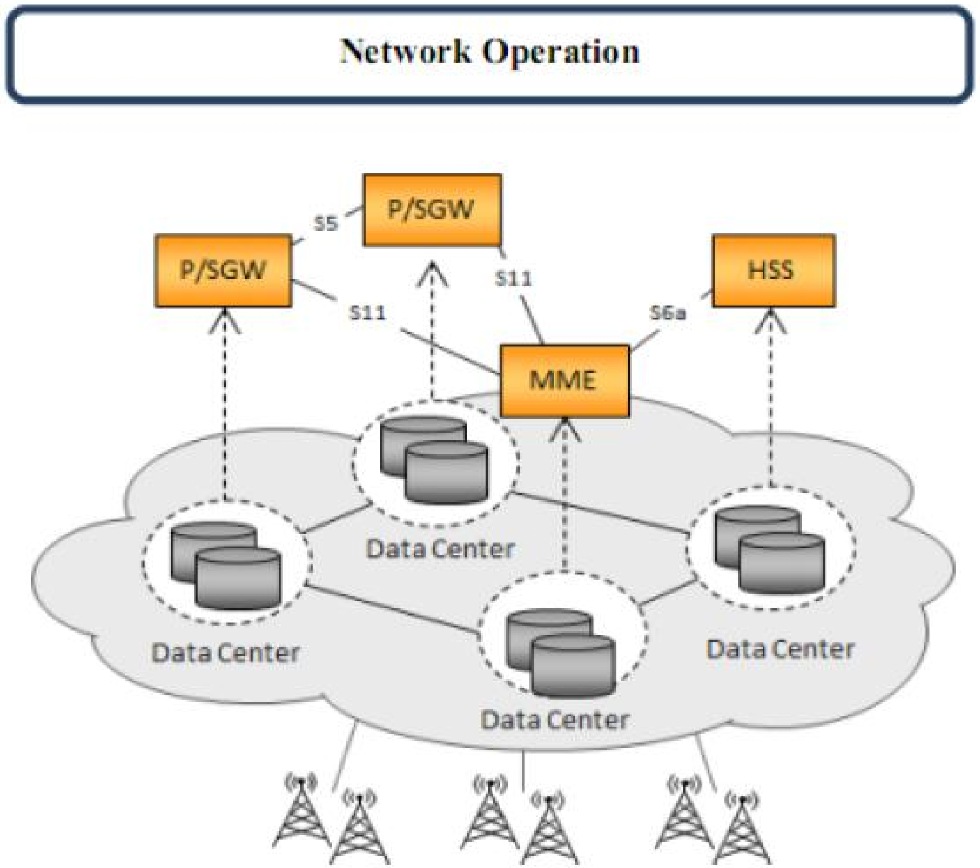

- Examples of Network Functions include MME, S/P-GW, etc
- This use case aims at applying virtualization to the EPC, the IMS, and these other Network Functions mentioned above

Other use cases for v-IMS

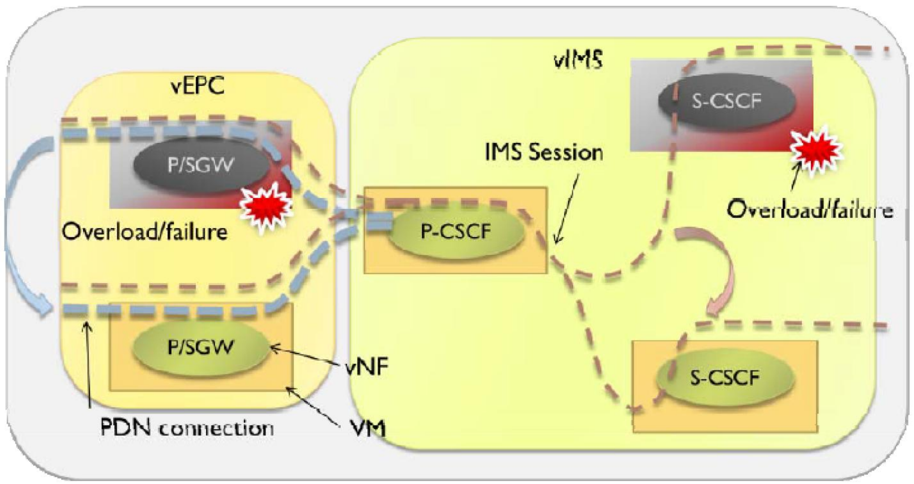

VNF relocation

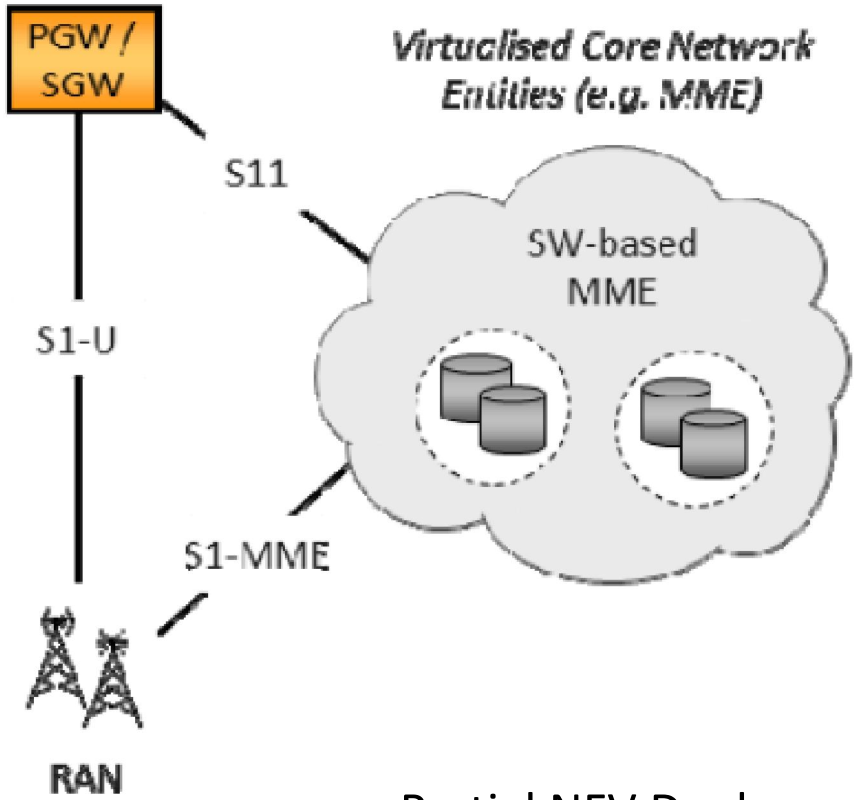

Partial NFV Deployment

Virtualization of Mobile Base Station

- **Mobile network traffic is significantly increasing** by the demand generated by application of mobile devices, while the **ARPU (revenue) is difficult to increase**

- **LTE is also considered as radio access part of EPS (Evolved Packet System)** which is required to fullfil the requirements of **high spectral efficiency, high peak data rates, short round trip time and frequency flexibility** in radio access network (RAN)

- **Virtualisation of mobile base station leverages** IT virtualisation technology to realize at least a part of RAN nodes onto **standard IT servers, storages and switches**

Virtualization of Mobile Base Station

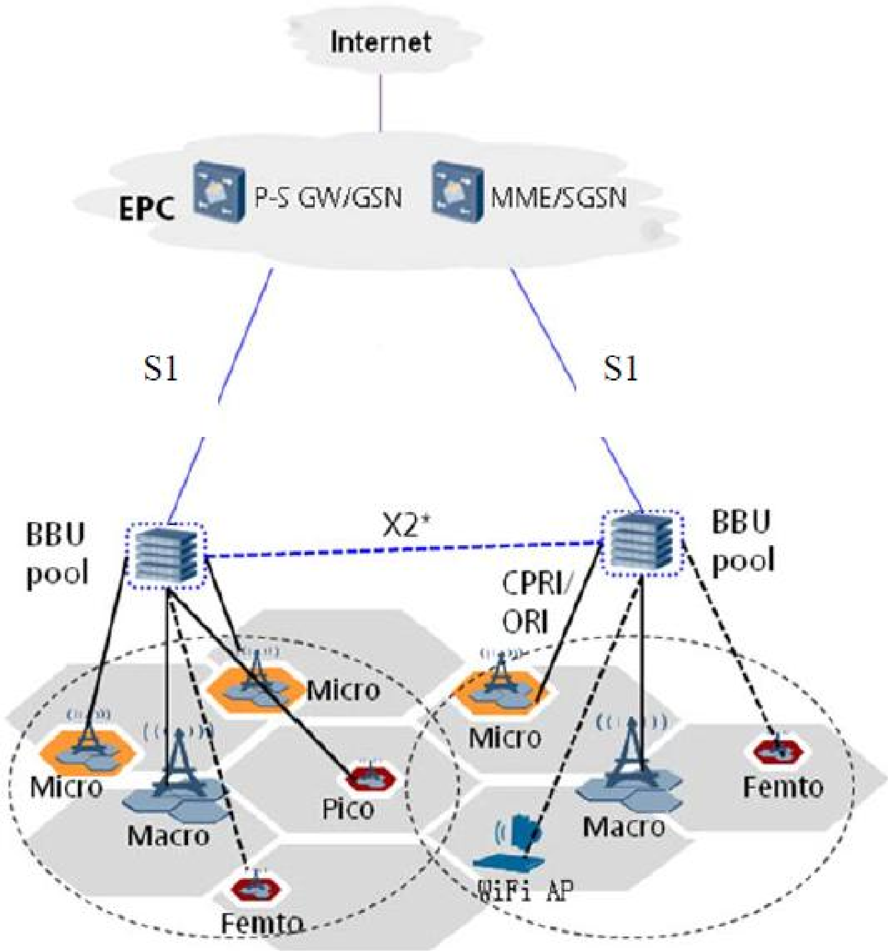

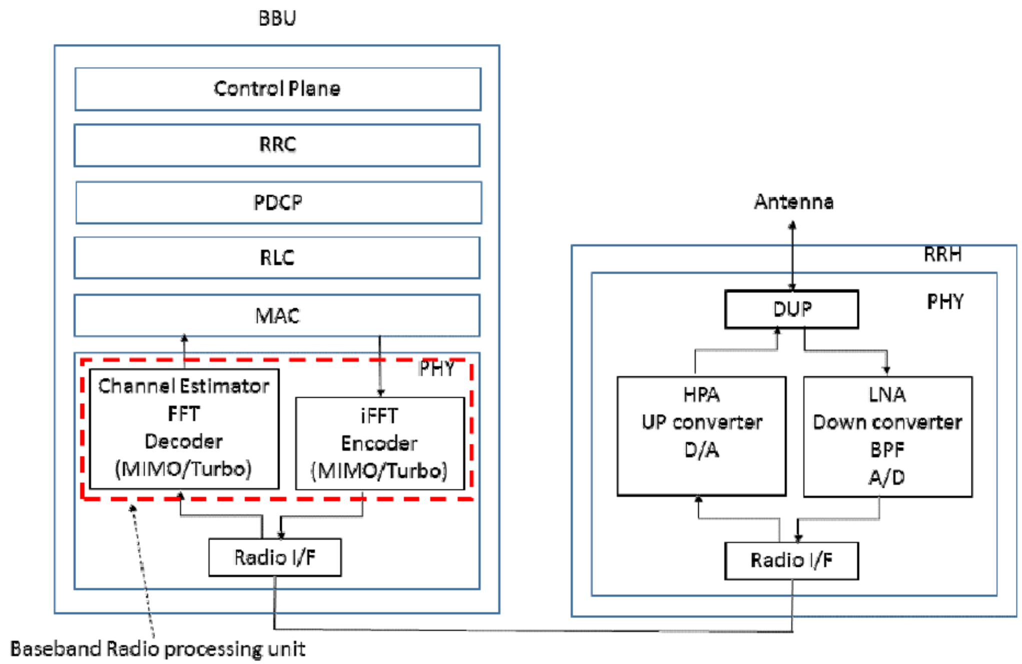

Functional blocks in C-RAN

LTE RAN architecture evolution by centralized BBU pool (Telecom Baseband Unit)

Virtualization of CDNs

- **Delivery of content, especially of video, is one of the major challenges** of all operator networks due to massive growing amount of traffic to be delivered to end customers of the network
- **Integrating nodes of Content Delivery Networks** into operator networks can be an effective and cost-efficient way to answer to the challenges of Video Traffic Delivery
- **CDN providers ask operators to deploy their proprietary cache nodes into the ISP network** (e.g. Netflix OpenConnect program, Akamai Aura CDN). This comes with **benefits for both sides** but also with the challenge that eventually the **operators will host a zoo of different cache devices** side by side in their premises

vCDN

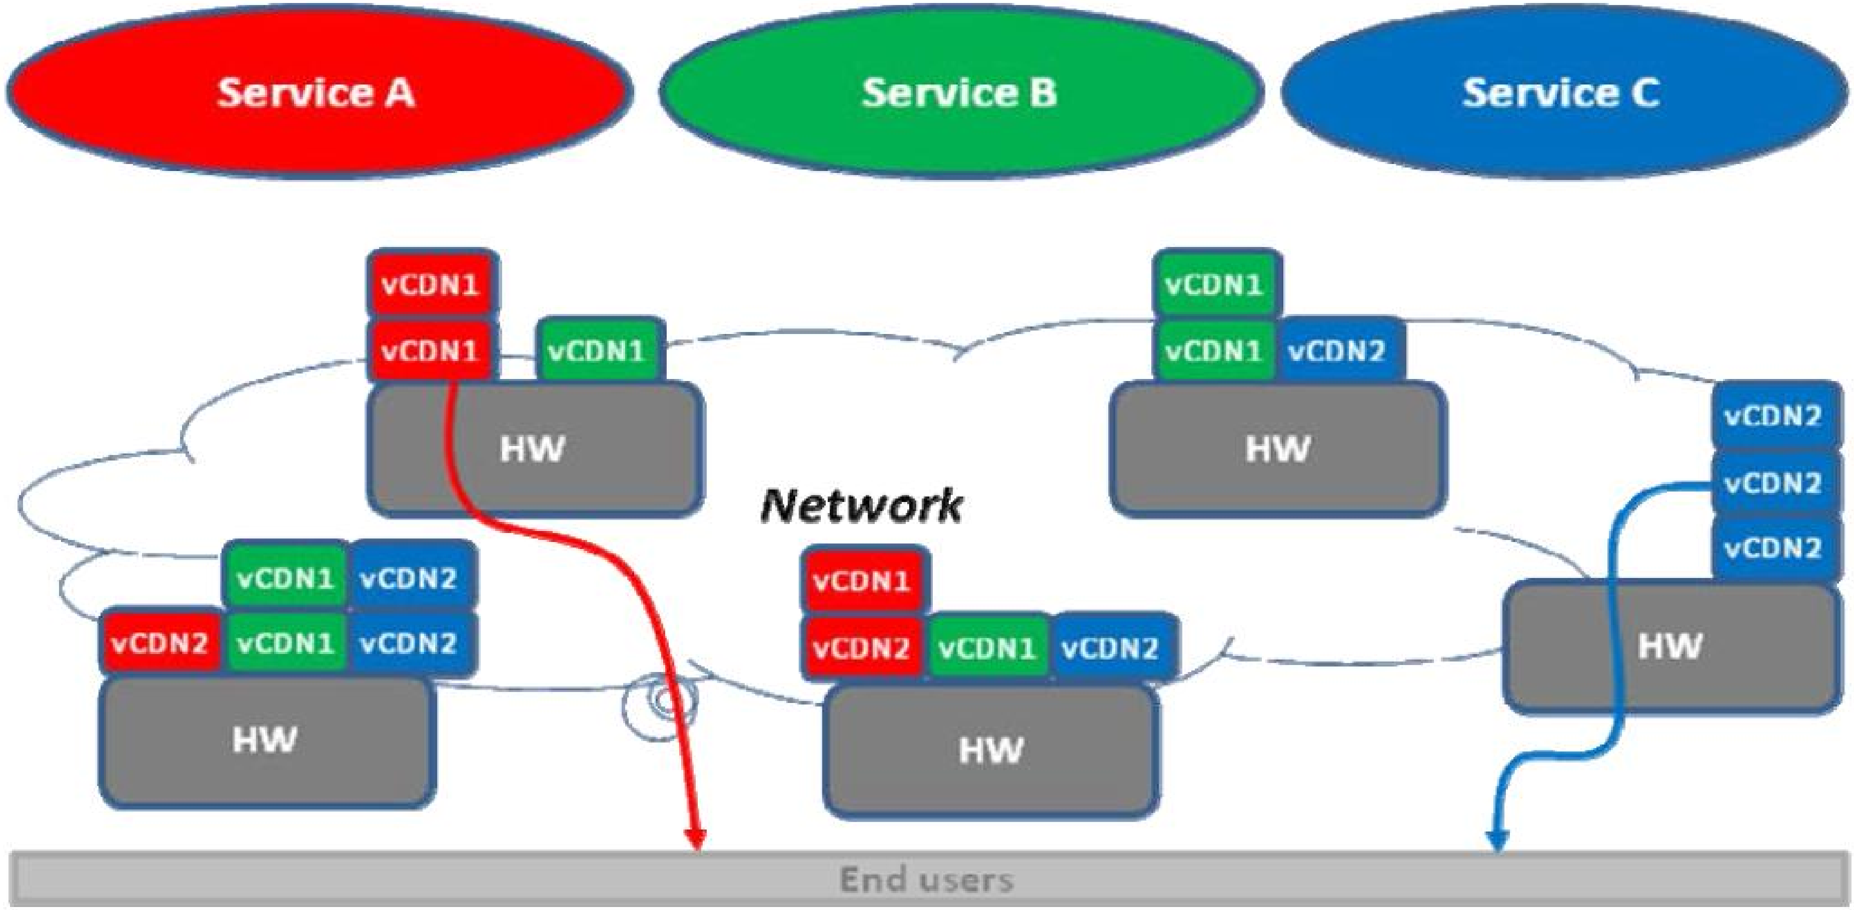

Principle of different vCDN cache nodes deployment in Virtualised environment

Home Environment

- Current **network operator provided home services** are architected using network-located **backend systems and dedicated CPE devices located as part of the home network**.
- These CPE devices mark the operator and/or service provider presence at the customer premises and usually include:
  - Residential Gateway (RGW) for Internet
  - VOIP services, and a
  - Setup Box (STB) for Media services normally supporting local storage for PVR services

Virtualization of the Home Environment

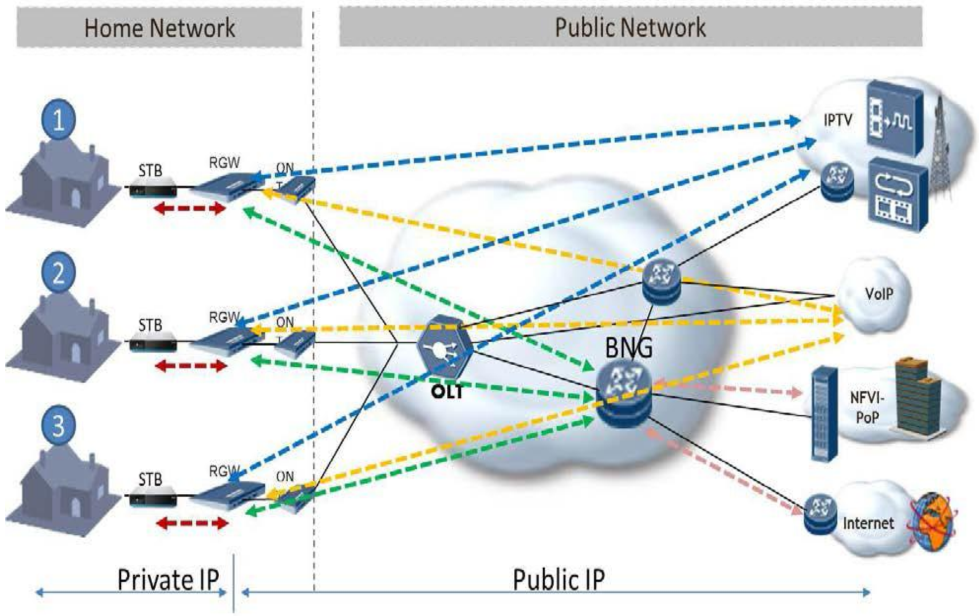

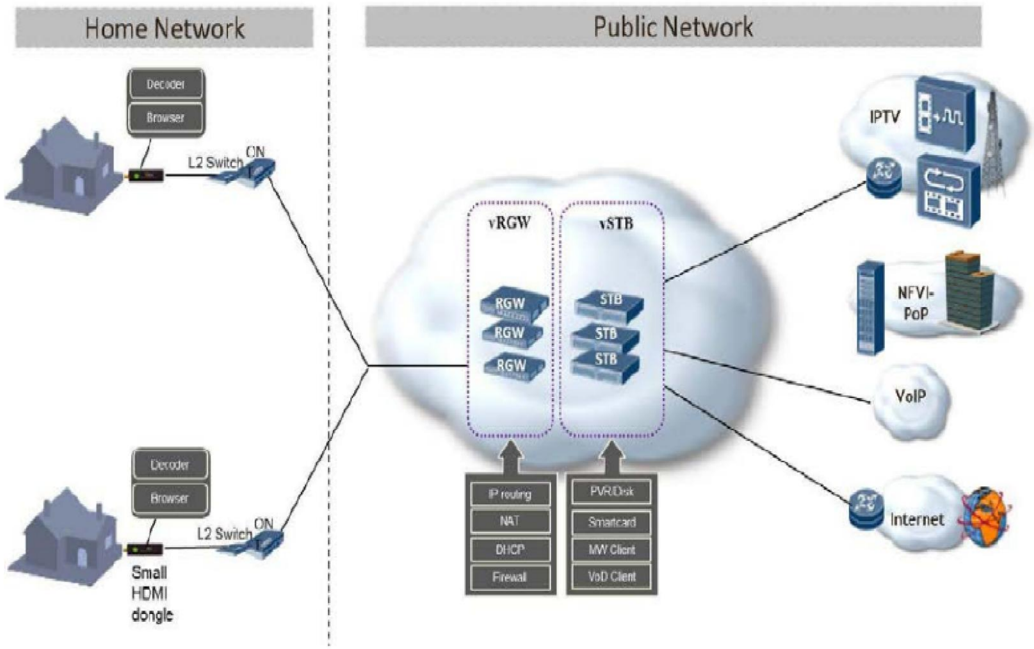

Virtualização das funções do ambiente doméstico. A criação do vRGW e vSTB permitiu a transferência de funções de rede e mídia para o centro da rede

Simplifying Operation and Service Deployment

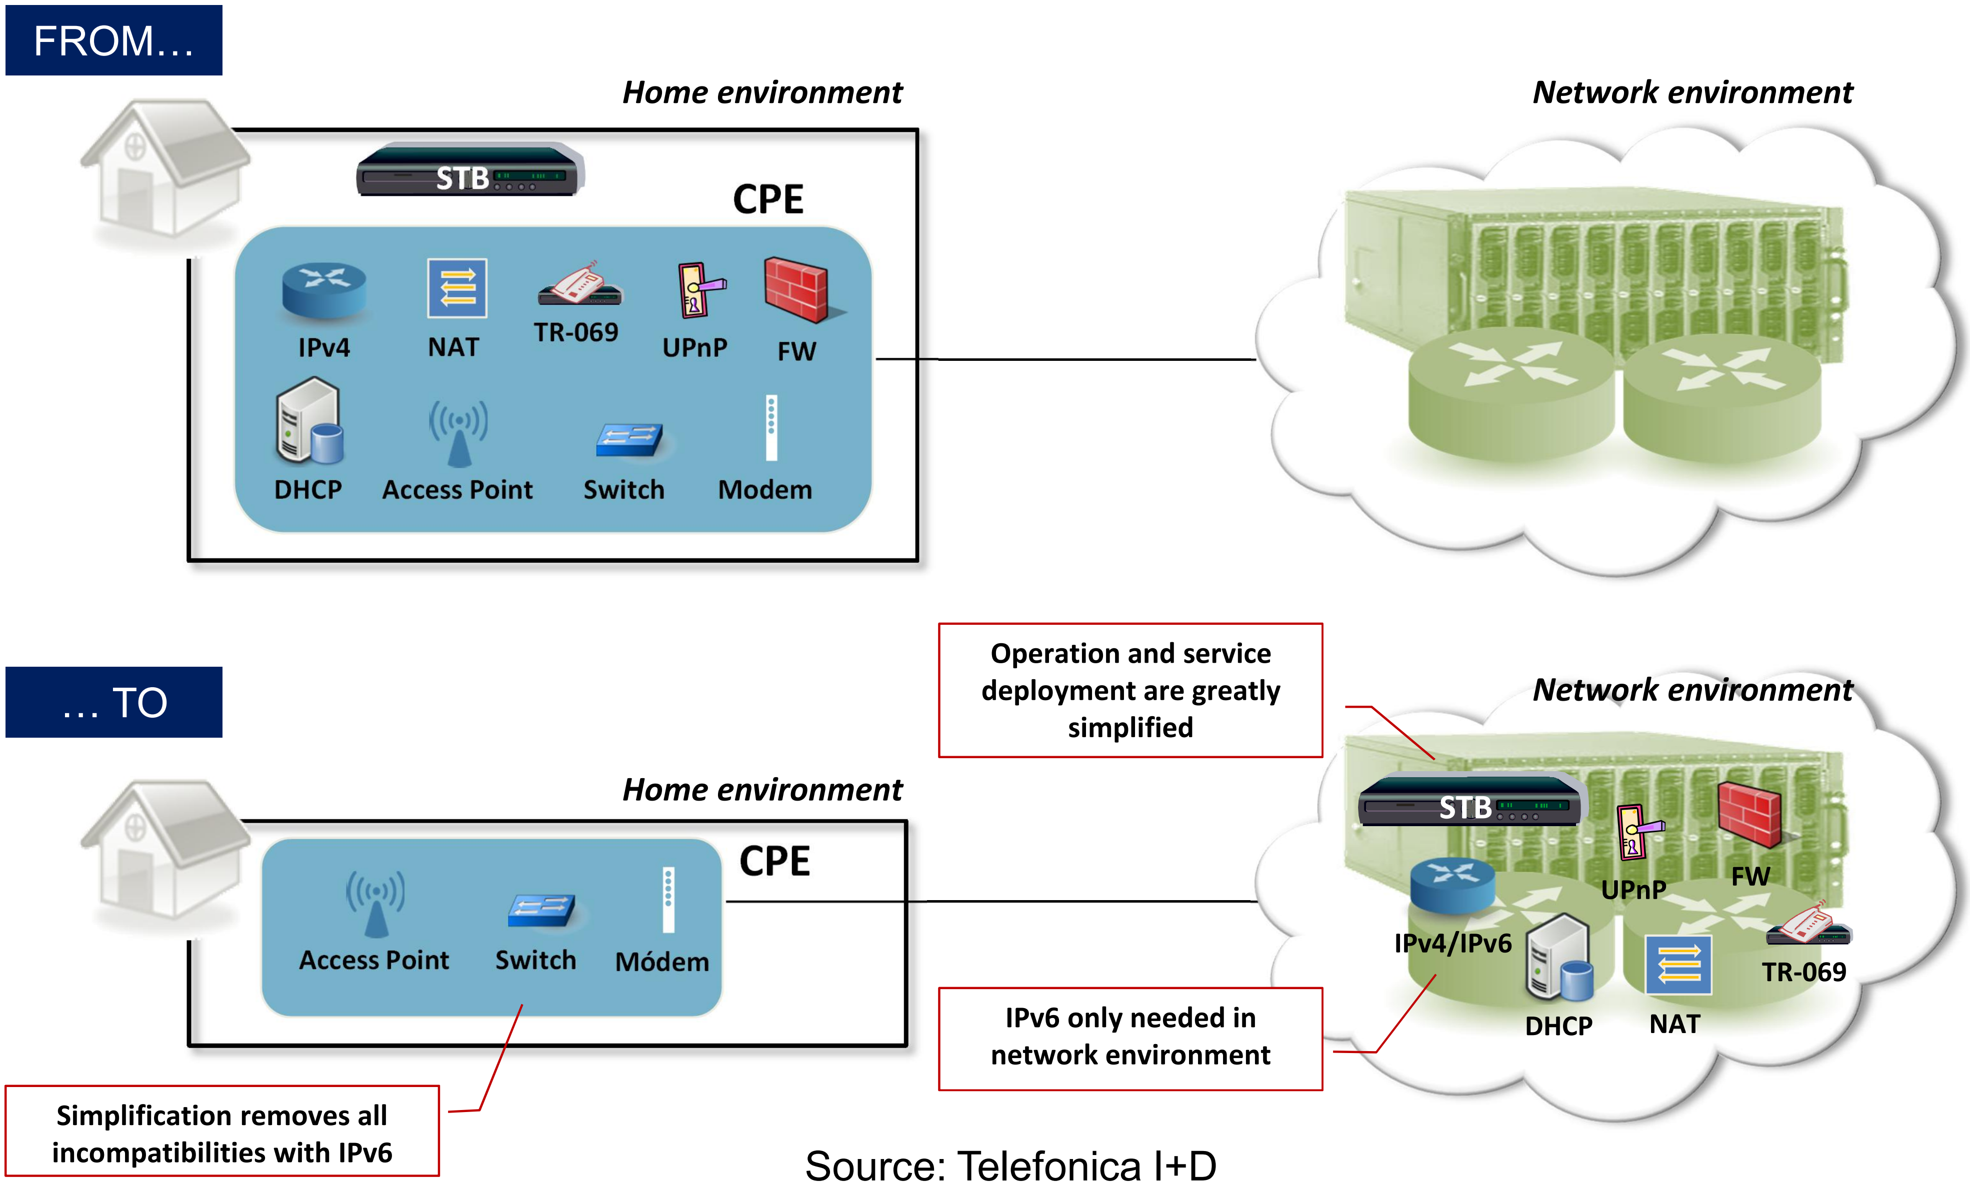

Source: Telefonica I+D

Virtual Residential Gateway

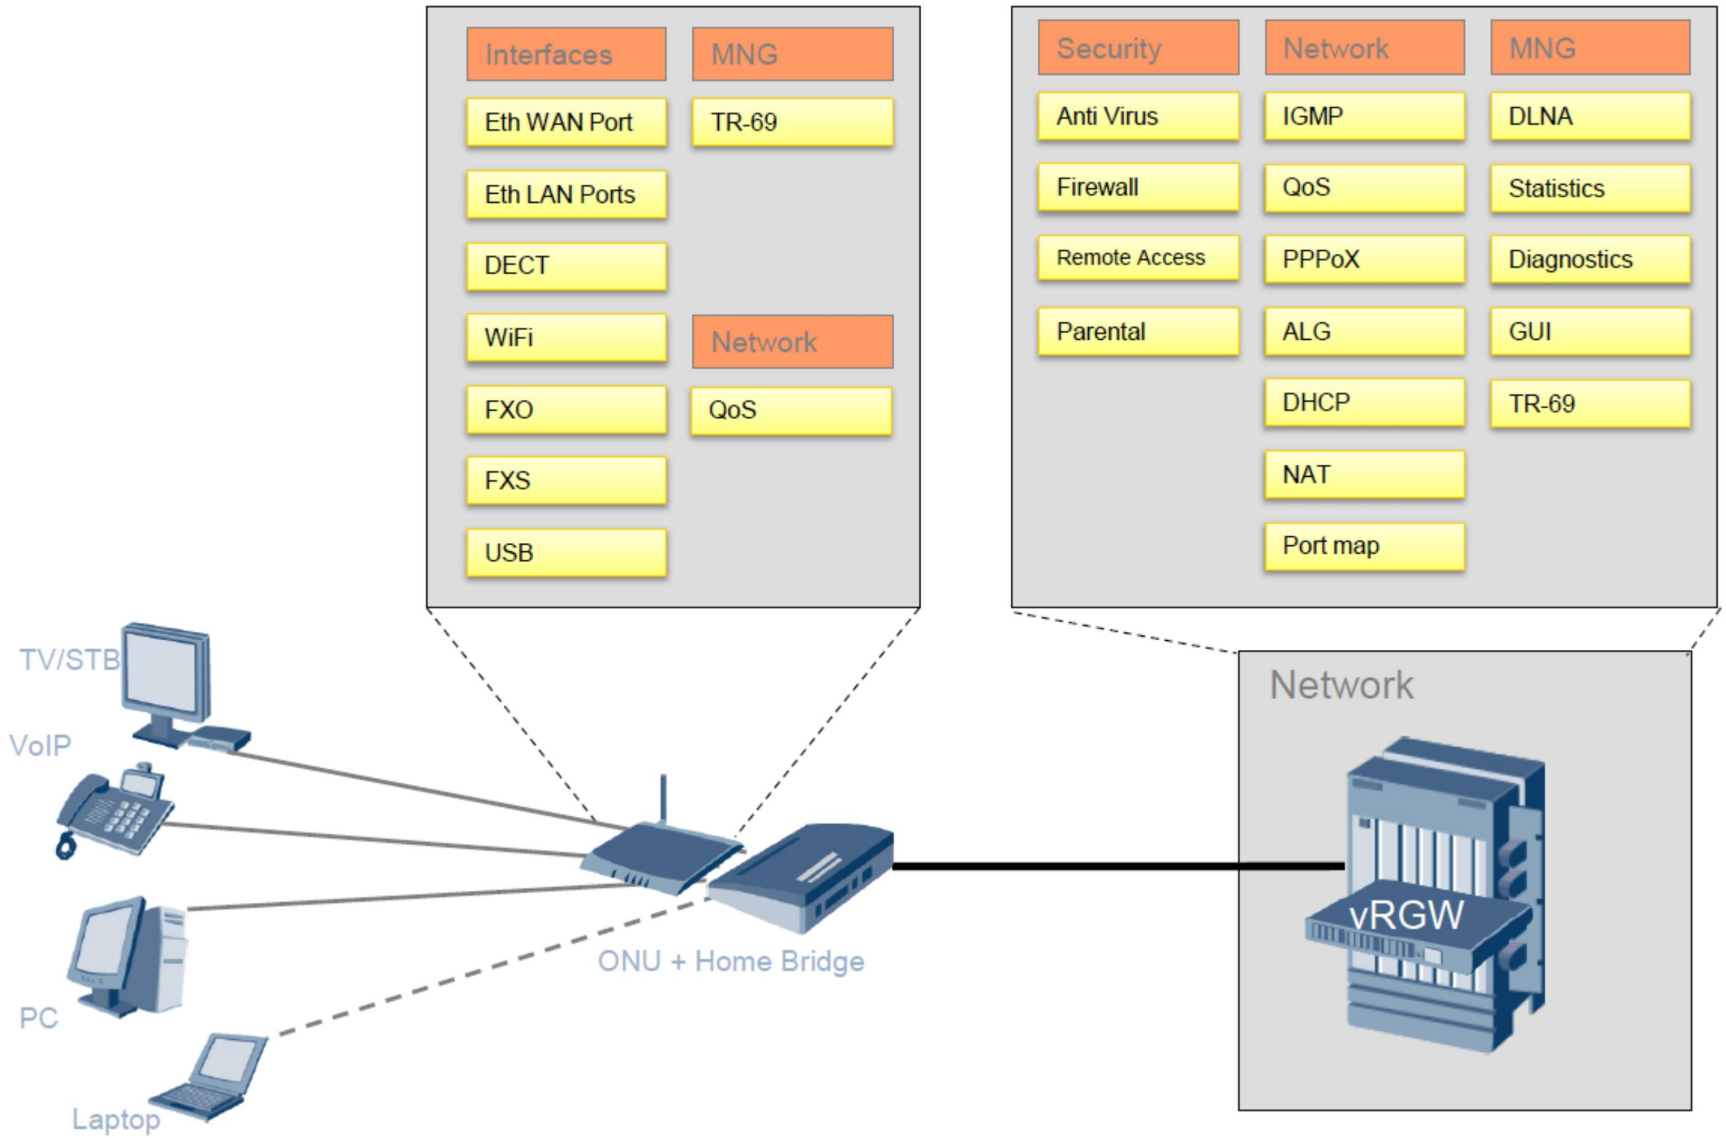

Source: NEC

Fixed Access NFV

- **Main costs and bottlenecks in a network often occur in the access.**
  - For the wireline fixed access network, the most prevalent broadband access technologies today are based on DSL, with the most widely deployed variant being **ADSL2+ which has a maximum downstream bit rate of ~26 Mb/s.**

- The **trend** however is to **replace exchange-based equipment** with equipment based on **VDSL2 in new street cabinets with fibre backhaul (FTTcab)**

Access Networks Virtualisation

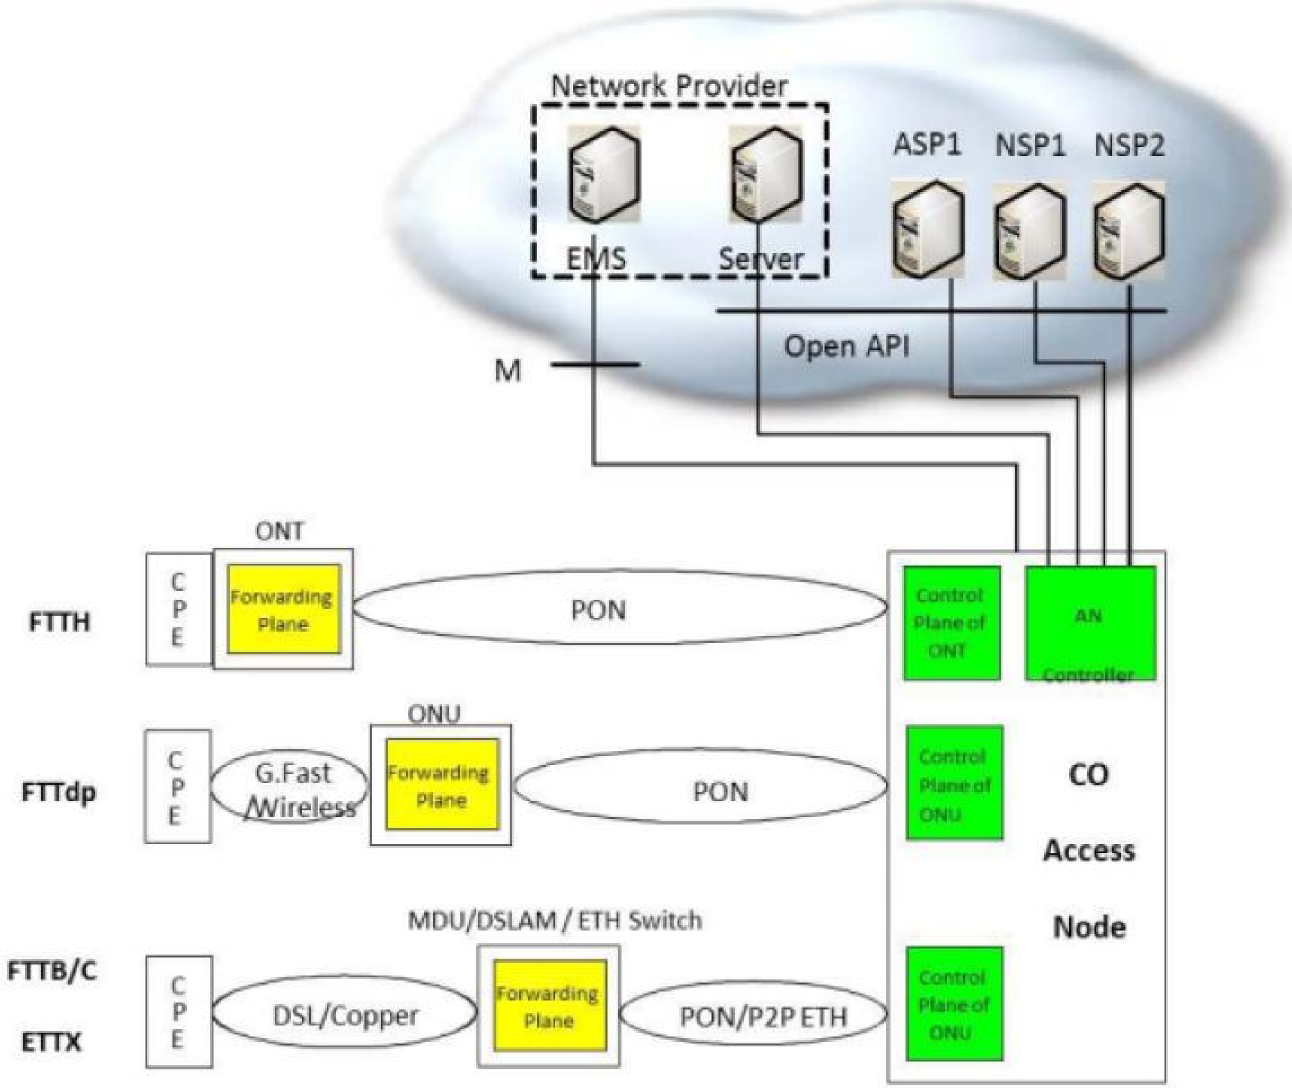

Target Network functions for virtualisation may include control functions from:

OLT
DSLAM
ONU
ONT
MDU
DPU

Access Network Functions Virtualisation will be initially applied to hybrid fibre-DSL nodes such as FTTcab and FTTdp

More Use Cases

Virtual CPE

Complex home environment

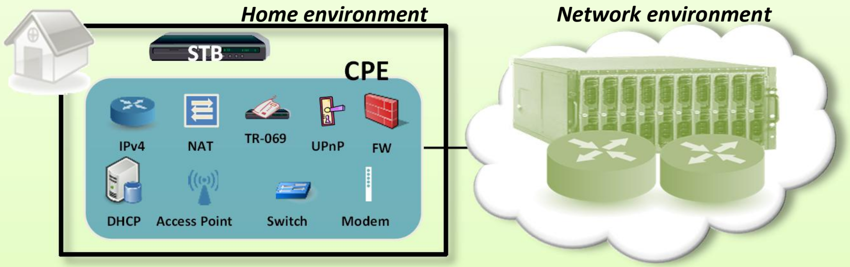

Home simplification

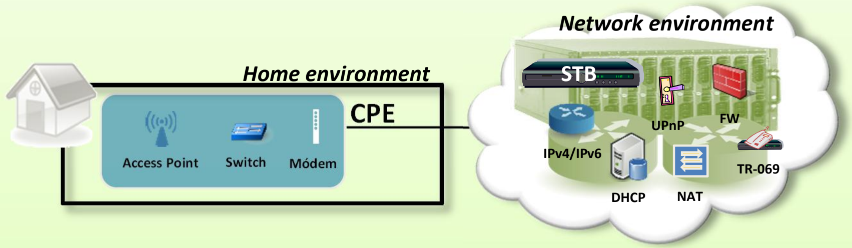

- Simplification or even supression (STB)
- No need for home router replacement as it is updated by configuration
- Fast deployment for new services
- Inexpensive IPv6 migration maintaining legacy home routers

Virtual IP Edge

Multiple IP Edges

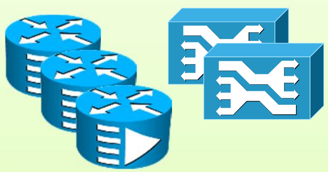

- An IP Edge for each service (voice, video content, Internet)
- Scattered and not well integrated control functions (e.g. DPI, BRAS, PCRF)

A unified software IP Edge

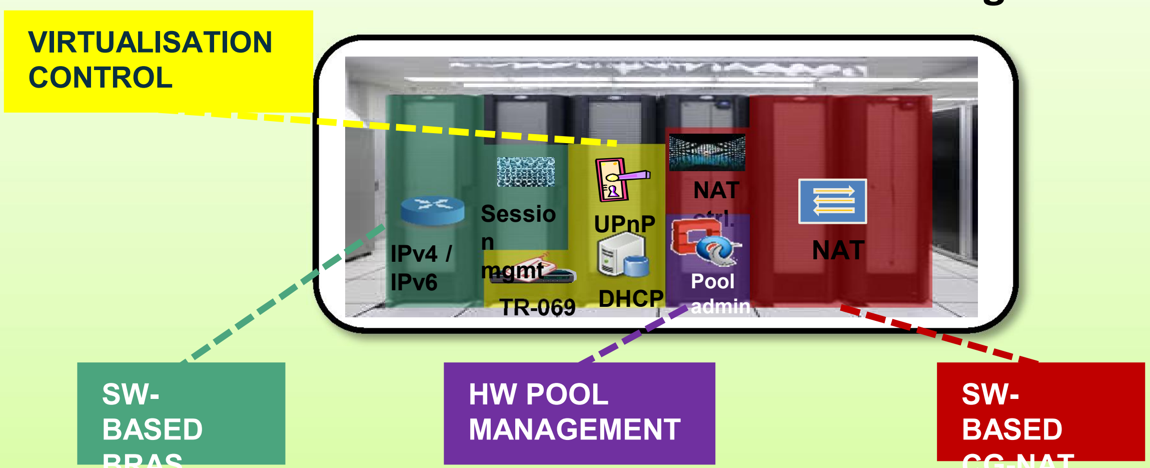

...More ...

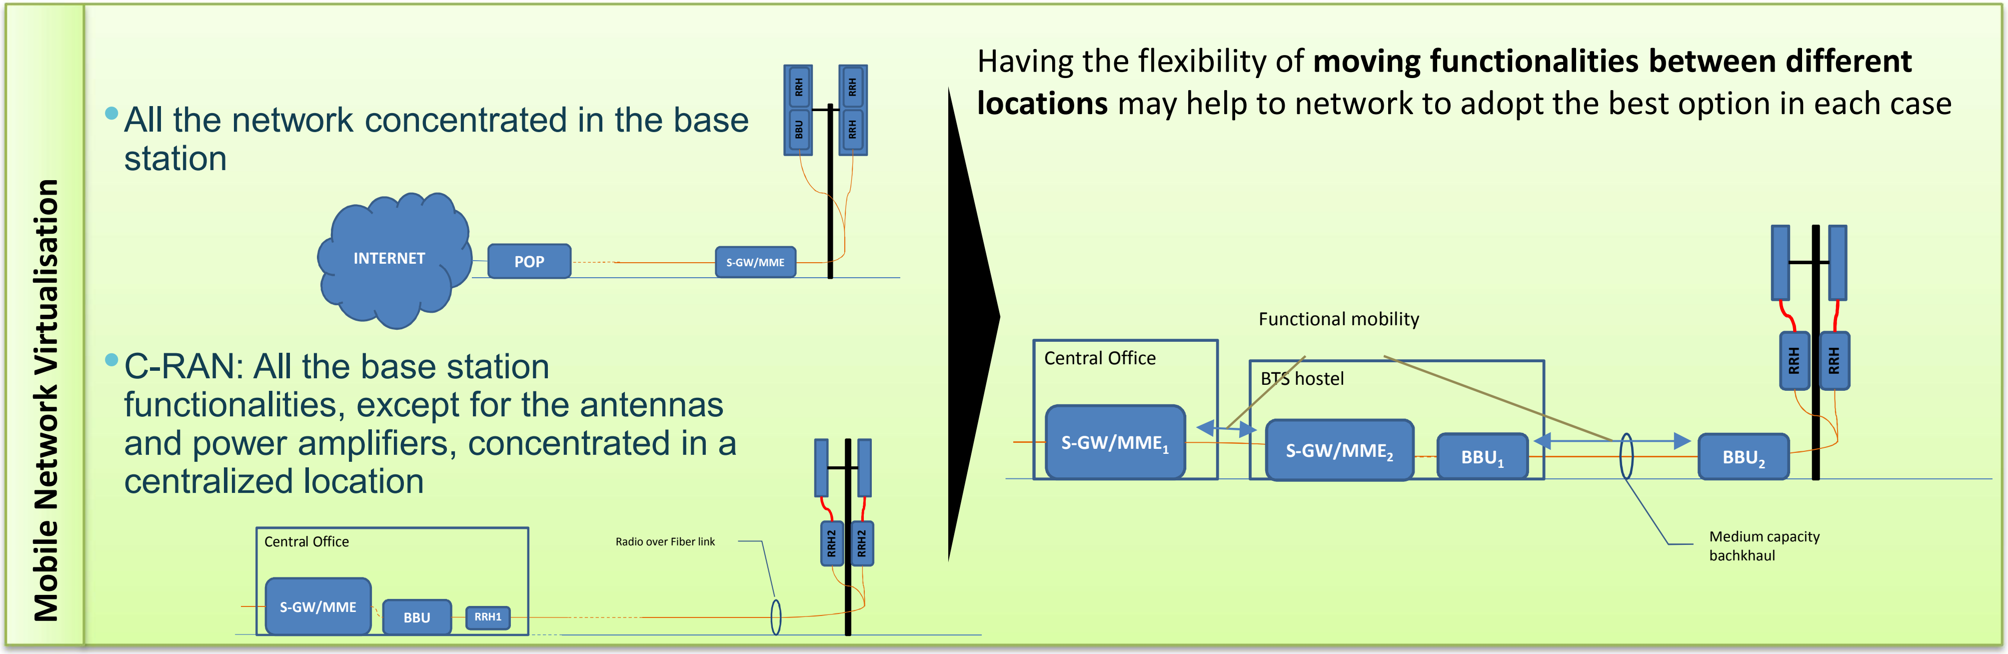
Mobile Network Virtualisation
All the network concentrated in the base station
INTERNET
POP
S-GW/MME
RRH
BBU
C-RAN: All the base station functionalities, except for the antennas and power amplifiers, concentrated in a centralized location
Central Office
S-GW/MME
BBU
RRH1
Radio over Fiber link
RRH2
Having the flexibility of **moving functionalities between different locations** may help to network to adopt the best option in each case
Functional mobility
Central Office
BTS hostel
S-GW/MME1
S-GW/MME2
BBU1
BBU2
RRH
Medium capacity bachkhaul

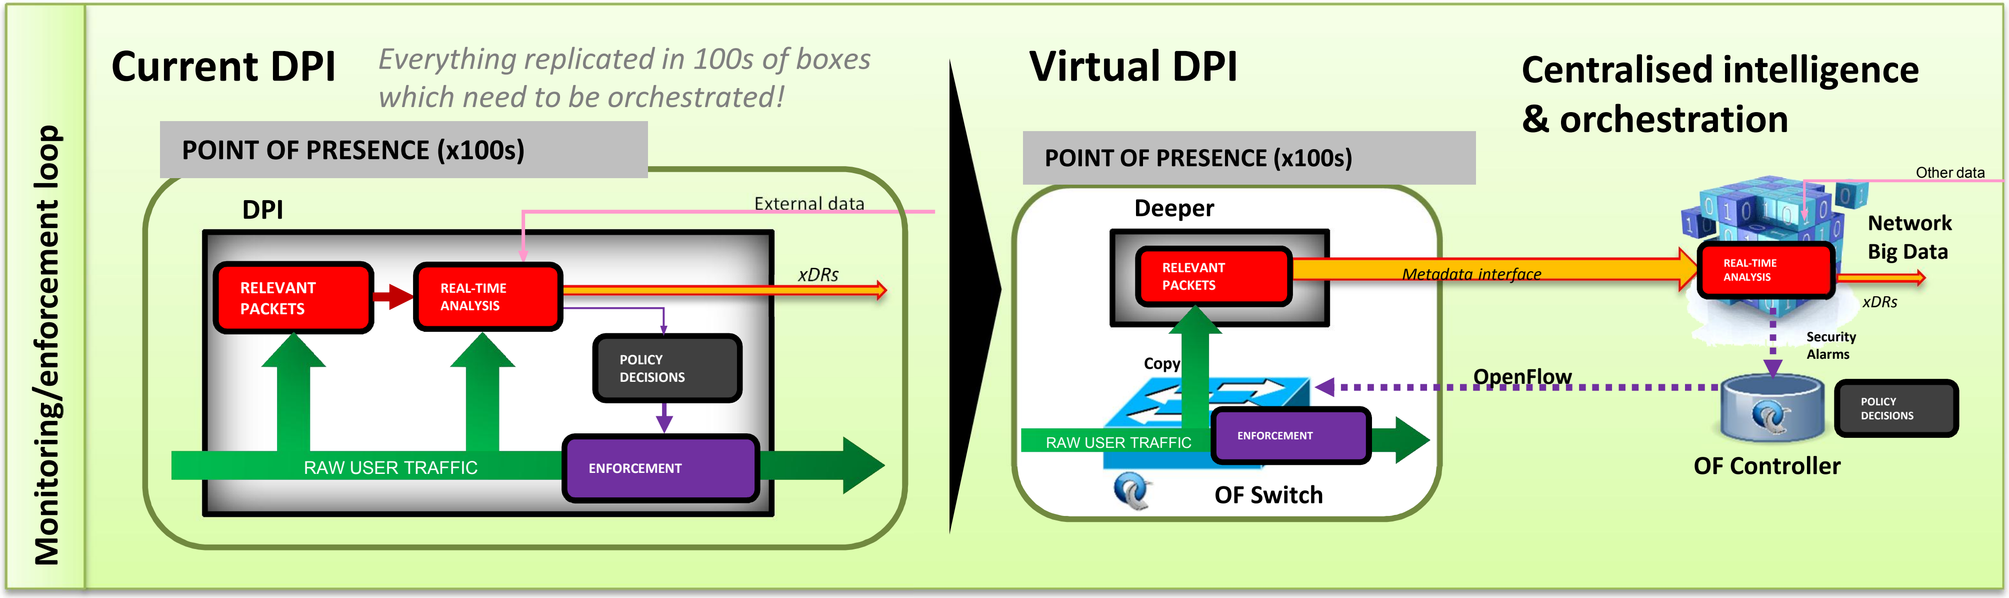
Monitoring/enforcement loop
Current DPI
Everything replicated in 100s of boxes which need to be orchestrated!
POINT OF PRESENCE (x100s)
DPI
External data
RELEVANT PACKETS
REAL-TIME ANALYSIS
xDRs
POLICY DECISIONS
RAW USER TRAFFIC
ENFORCEMENT
Virtual DPI
Centralised intelligence & orchestration
POINT OF PRESENCE (x100s)
Deeper
RELEVANT PACKETS
Metadata interface
Other data
Network Big Data
REAL-TIME ANALYSIS
xDRs
Security Alarms
Copy
OpenFlow
RAW USER TRAFFIC
ENFORCEMENT
OF Switch
POLICY DECISIONS
OF Controller

… And a Couple More

Virtualized CGNAT

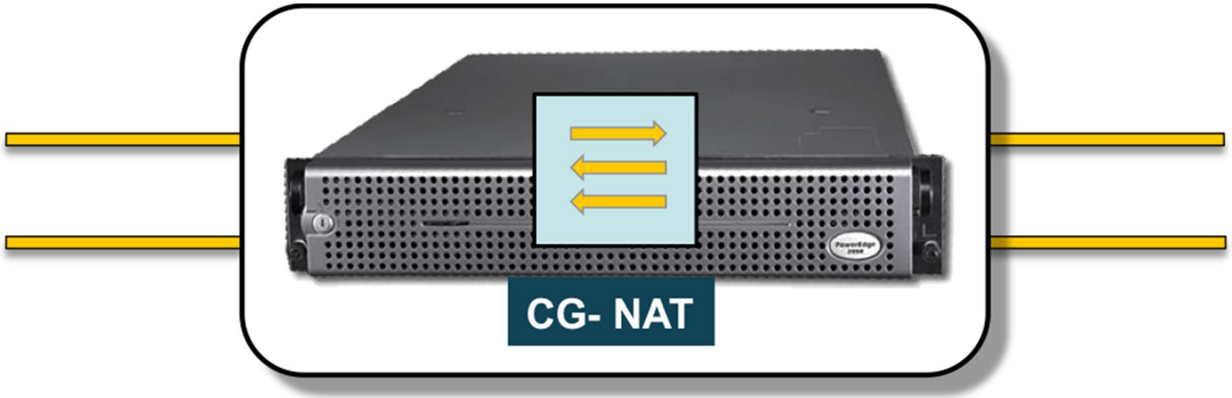

- NAT44 function, extensible to IPv6 transition
- 40 Gbps full-duplex line rate per server
- Support of overlapping addresses and tunnelling
- Auto-provisioning of NAT sessions per access line

Optimized Quagga data plane

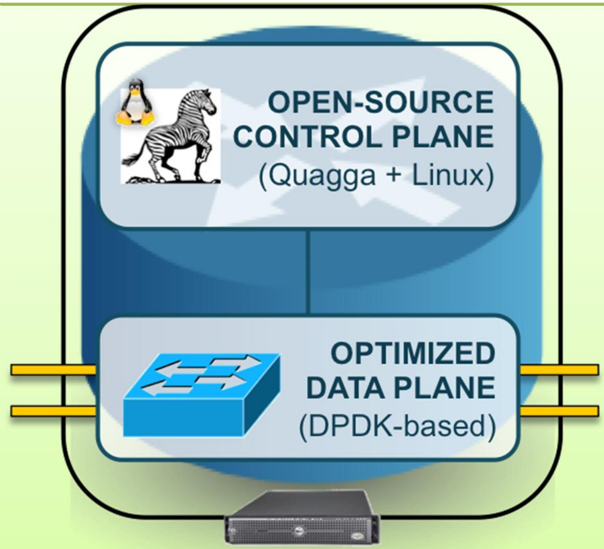

- Leverage on open source routing project as rich and widely tested protocol suite while assuring data plane performance
  - Common routing protocols supported and extended by open source project
  - High-performance line-rate data plane
  - Running in separate process, does not lead to licensing issues

Some Other Use Case Examples

…not in any particular order

- Switching elements: BNG, CG-NAT, routers.
- Mobile network nodes: HLR/HSS, MME, SGSN, GGSN/PDN-GW.
- Home networks: Functions contained in home routers and set top boxes to create virtualised home environments.
- Tunnelling gateway elements: IPSec/SSL VPN gateways.
- Traffic analysis: DPI, QoE measurement.
- Service Assurance: SLA monitoring, Test and Diagnostics.
- NGN signalling: SBCs, IMS.
- Converged and network-wide functions: AAA servers, policy control and charging platforms.
- Application-level optimisation: CDNs, Cache Servers, Load Balancers, Application Accelerators.
- Security functions: Firewalls, virus scanners, intrusion detection systems, spam protection.

Source: NFV

Current Use Cases

CUSTOM DPI

SW-based DPI over general-purpose HW

- *Advanced traffic analysis*
- *Multidimensional reports*

✓ *"Off-the-shelf" hardware working at line rate*

SW-NODE

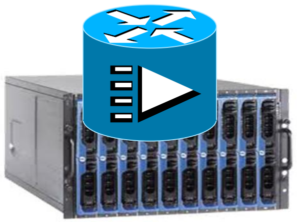

Soft-Node

SW-defined IP nodes (BRAS/CG-NAT) over general-purpose HW

- *Software-defined network over COTS equipment*
- *Effective reuse of HW resources after NAT use declines*

✓ *Escape from telco cycle in IP network*

OPENFLOW

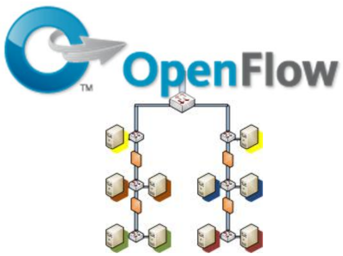

SW-controlled connectivity in datacenter environments

- *Minimise network management complexity at large-scale infrastructures*

✓ *Smooth coordination between network and cloud services*

VIRTUAL EDGE

DS vCPE

Virtualisation of legacy home routers: inexpensive IPv6 migration

- *Cost-effective solution for massive IPv6 migration*

✓ *No need of CPE replacement as it is updated by remote configuration*

Source: Telefonica I+D

Use case: Software BRAS

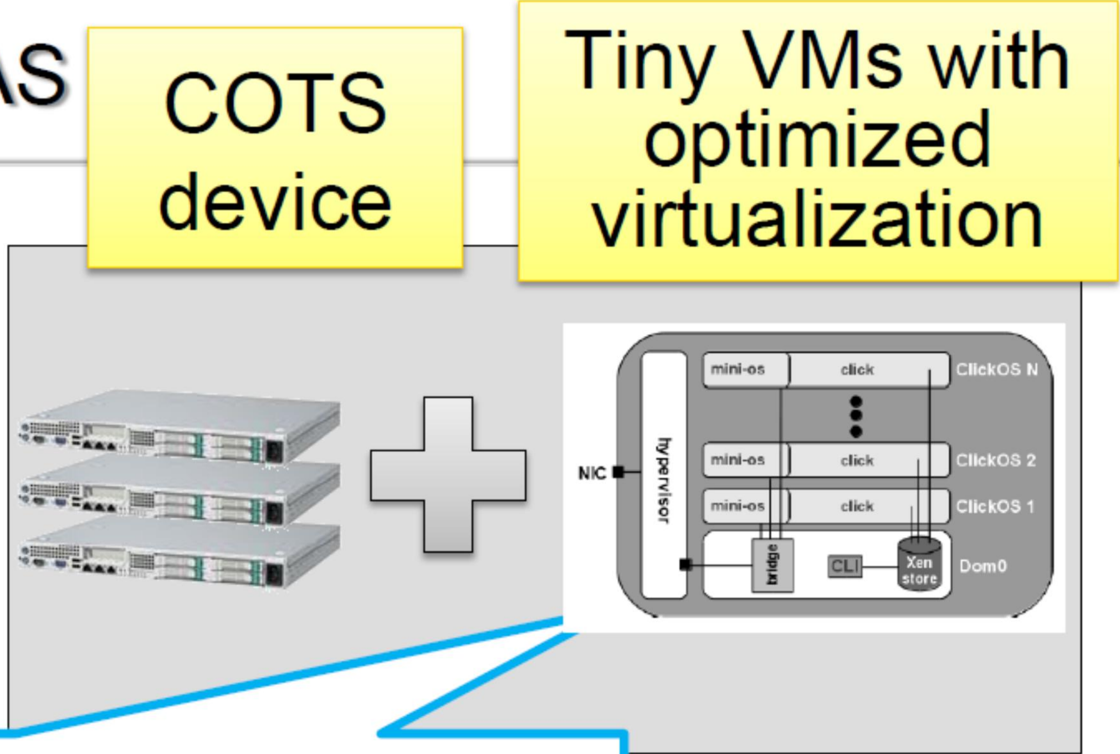

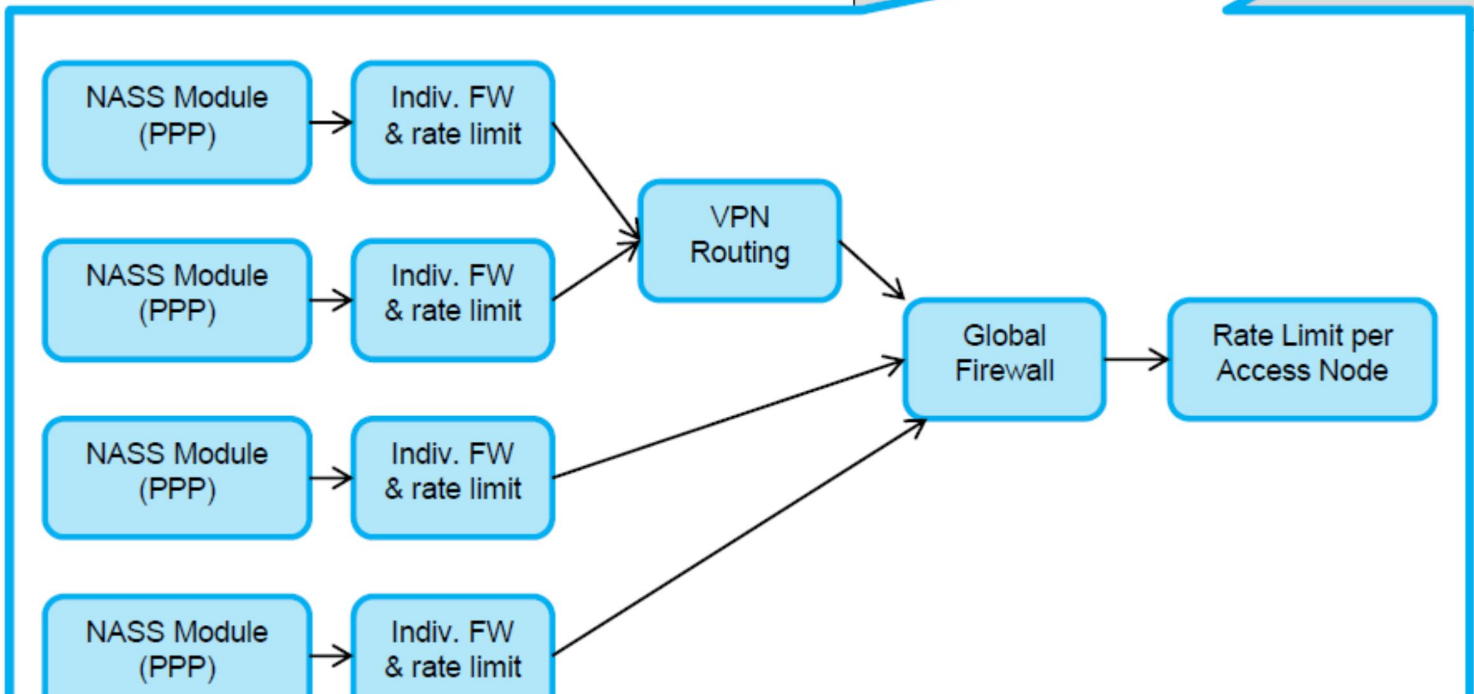

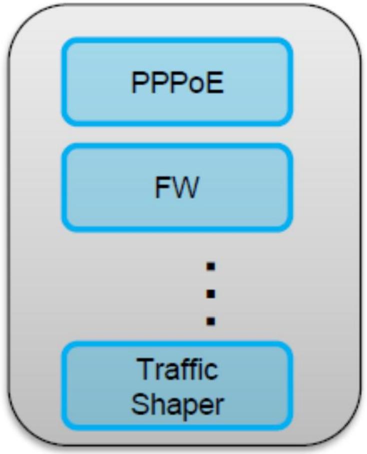

SW-defined NW Functions in ClickOS

Source: NEC

Casos de Uso : Requisitos & Desafios

Casos de Uso	**Requisitos**	**Desafios**
NFVIaaS	Desempenho e interoperabilidade	Gerenciamento, orquestração e automação de arquiteturas e funções de rede
VNFaaS	Elasticidade e desempenho	Segurança e resiliência
VNF FG	Interoperabilidade e co-existência de plataformas legadas	Elasticidade e desempenho
VNPaaS	Desempenho, interoperabilidade e orquestração	Gerenciamento e segurança
Mobile Core e IMS	Desempenho, interoperabilidade, gerenciamento e orquestração	Elasticidade, segurança e resiliência
CDNs	Interoperabilidade, elasticidade	Gerenciamento e desempenho
Home Environment	Portabilidade e interoperabilidade	Elasticidade e desempenho
Mobile BS	Desempenho e interoperabilidade	Segurança, resiliência e automação de arquitetura e funções de rede
Fixed Access	Portabilidade e elasticidade	Segurança e resiliência

Proof-of-Concepts

NFV

Proof of Concepts
ETSI Evaluation Process

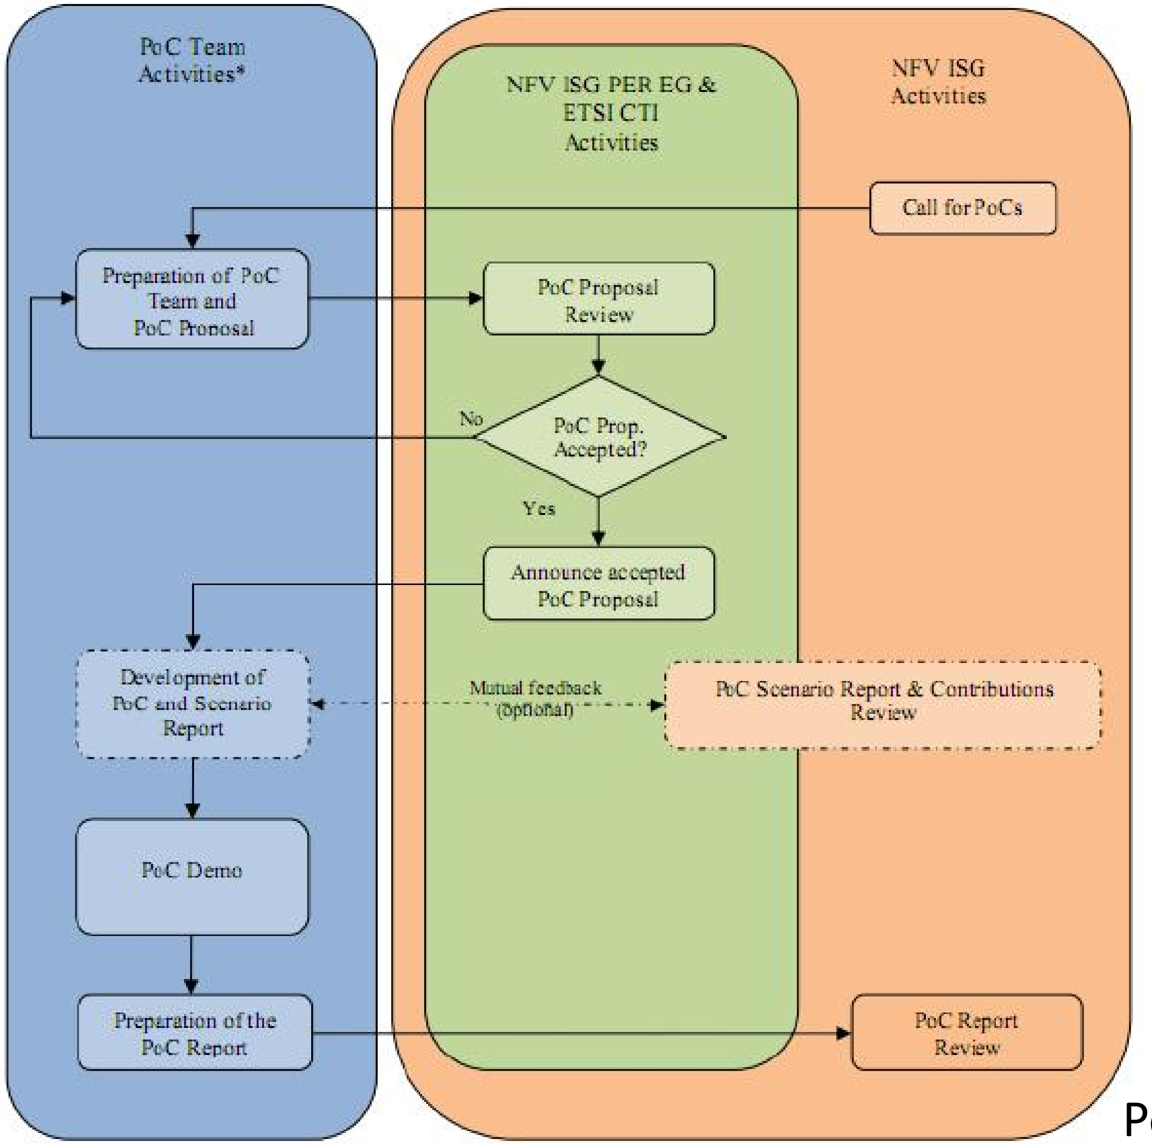

PoCs NFV ISG Diagram

Ongoing Proof of Concepts

- **CloudNFV Open NFV Framework Project**
  - Telefonica - Sprint - 6WIND - Dell - EnterpriseWeb – Mellanox - Metaswitch - Overture Networks - Qosmos - Huawei - Shenick
- **Service Chaining for NW Function Selection in Carrier Networks**
  - NTT - Cisco - HP - Juniper Networks
- **Virtual Function State Migration and Interoperability**
  - AT&T - BT - Broadcom Corporation - Tieto Corporation
- **Multi-vendor Distributed NFV**
  - CenturyLink - Certes - Cyan - Fortinet - RAD
- **E2E vEPC Orchestration in a multi-vendor open NFVI environment**
  - Telefonica - Sprint - Intel - Cyan - Red Hat - Dell - Connectem
- **Virtualised Mobile Network with Integrated DPI**
  - Telefonica - Intel - Tieto - Qosmos - Wind River Systems - Hewlett Packard
- **C-RAN virtualisation with dedicated hardware accelerator**
  - China Mobile - Alcatel-Lucent - Wind River Systems - Intel
- **Automated Network Orchestration**
  - Deutsche Telekom - Ericsson - x-ion GmbH - Deutsche Telekom Innovation Laboratories
- **VNF Router Performance with DDoS Functionality**
  - AT&T - Telefonica - Brocade - Intel - Spirent
- **NFV Ecosystem**
  - Telecom Italia - DigitalWave - SunTec - Svarog Technology Group - Telchemy - EANTC
- **Multi-Vendor on-boarding of vIMS on a cloud management framework**
  - Deutsche Telekom - Huawei Technologies - Alcatel-Lucent
- **Demonstration of multi-location, scalable, stateful Virtual Network Function**
  - NTT - Fujitsu - Alcatel-Lucent

CloudNFV

Dell Lab infrastructure for CloudNFV

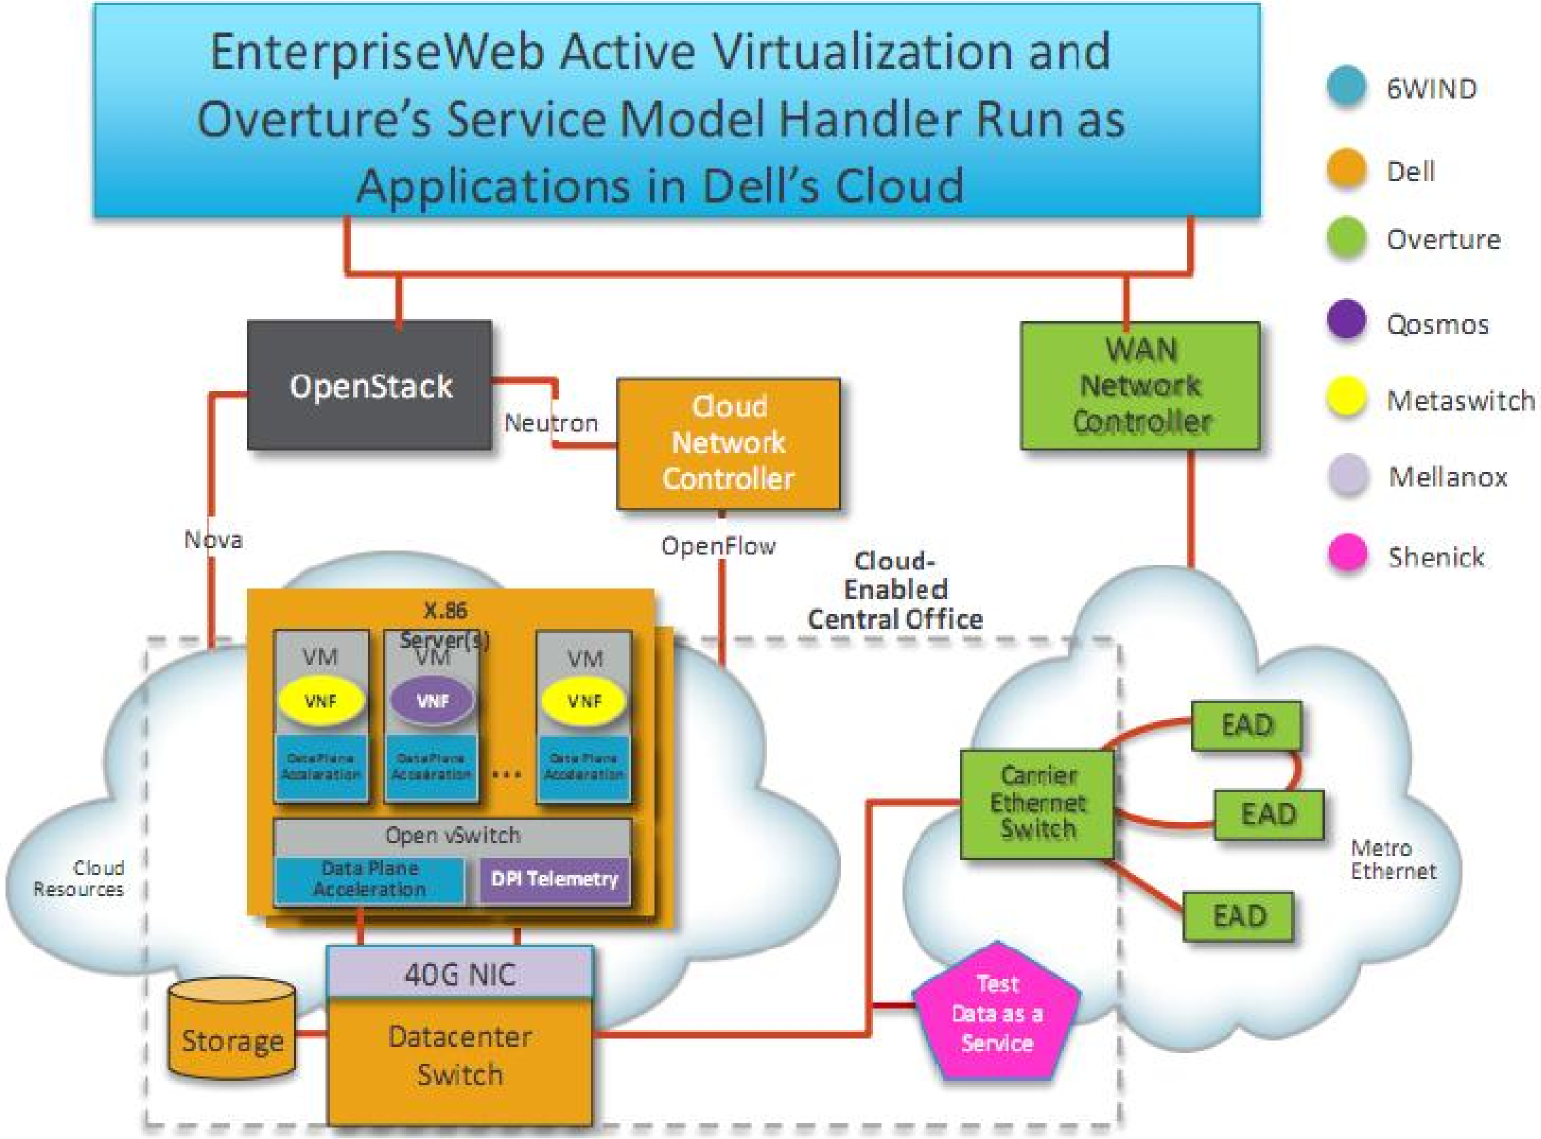

Source: ETSI Ongoing PoC
http://nfvwiki.etsi.org/index.php?title=On-going_PoCs

Service Chaining for NW Function Selection in Carrier Networks

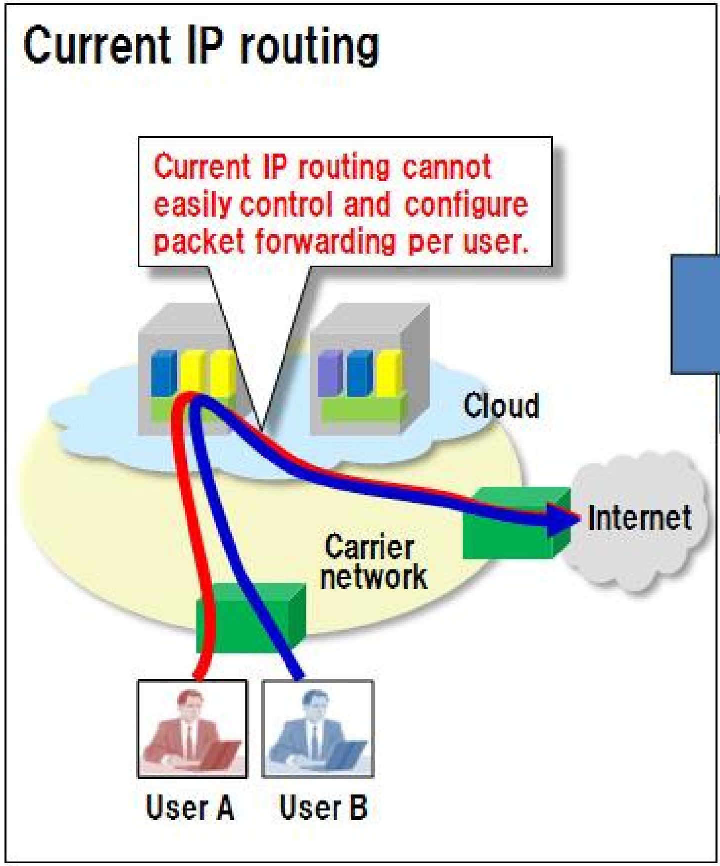

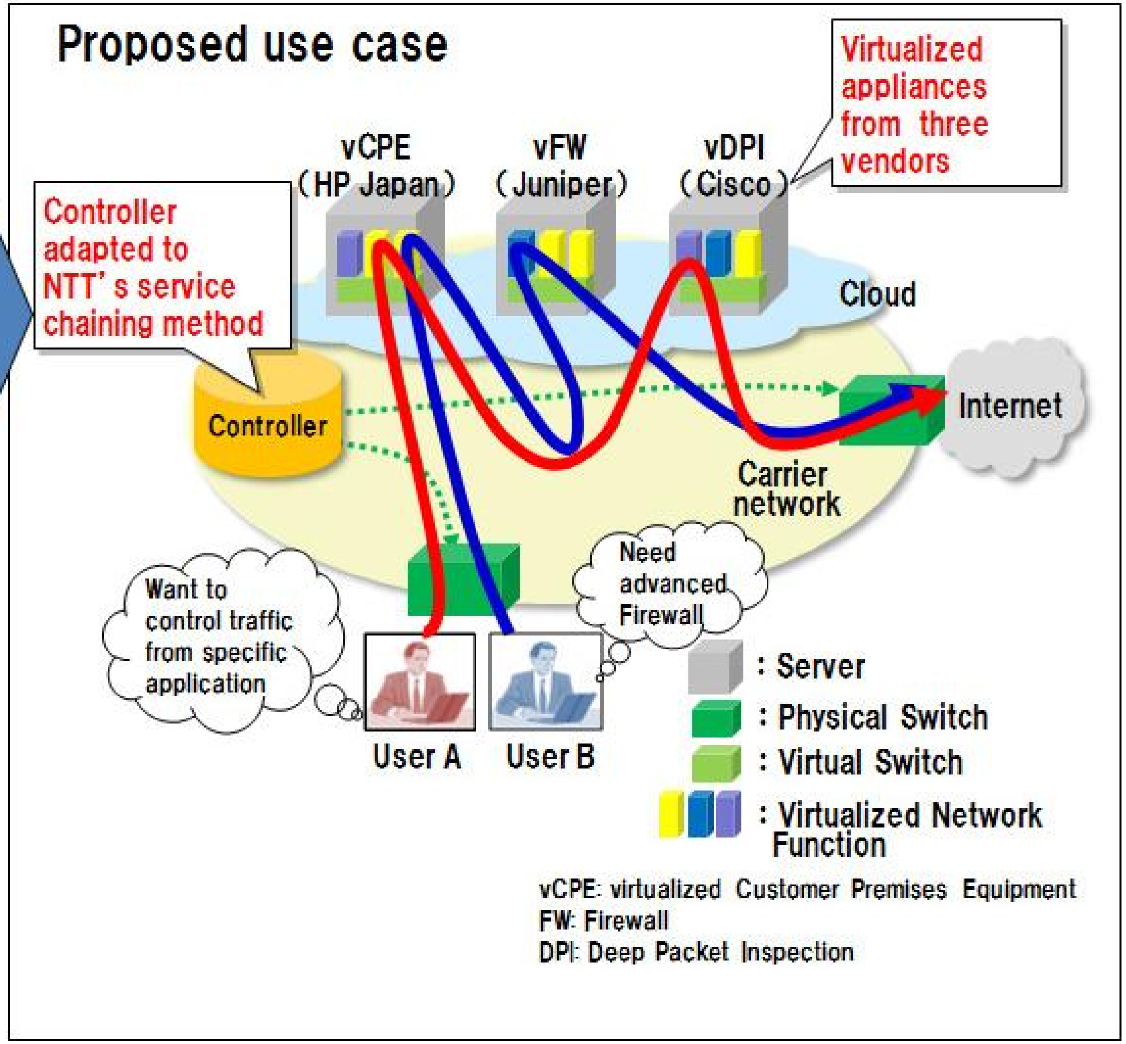

vDPI: CSR 1000v (Cisco Systems)
vCPE: VSR1000 (Hewlett-Packard)
vFW: FireFly (Juniper Networks)
VIM (NW Controller): Service Chaining Function (prototype) + Ryu (NTT)

Source: ETSI Ongoing PoC

Virtual Function State Migration and Interoperability

- Different Hardware BUT Portable Software
- Open Source + Linux + KVM
- Recompiling with GCC or LLVM - Low Level Virtual Machine

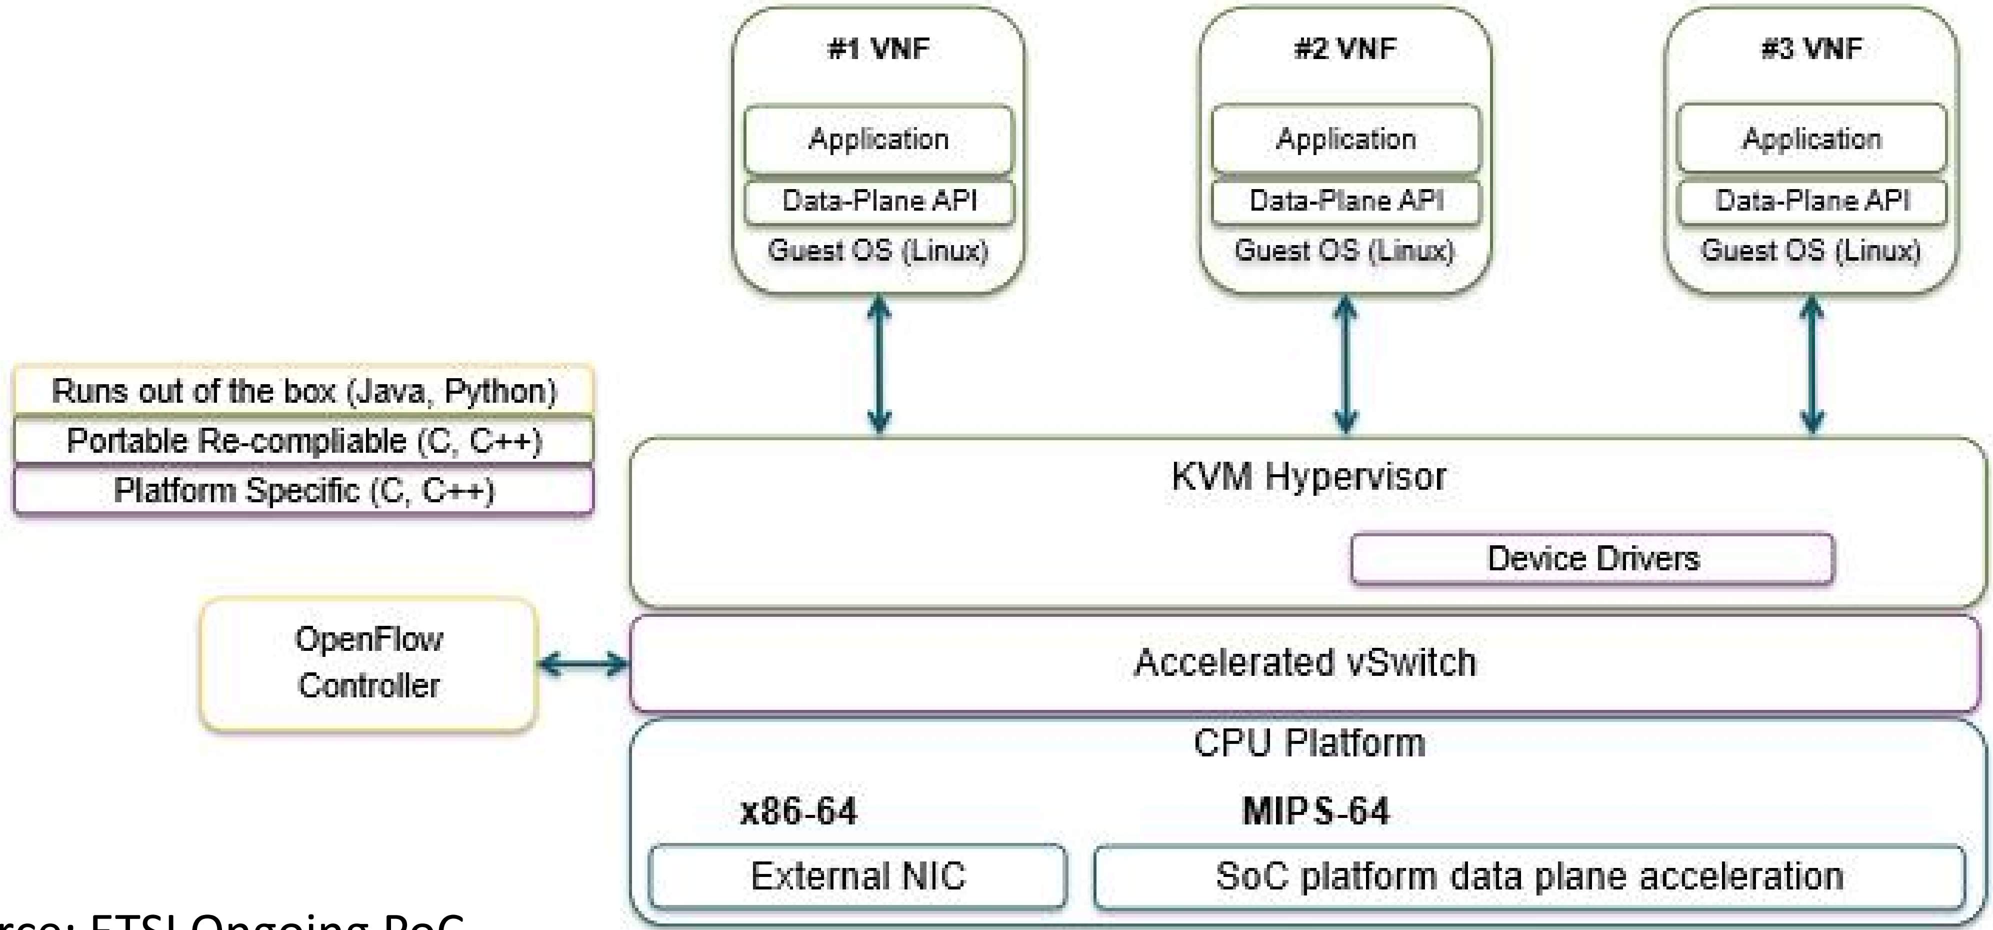

Source: ETSI Ongoing PoC

Virtual Function State Migration and Interoperability: x86 e MIPS-64

VNF state migration from x86 to Broadcom XLP MIPS-64 based blade server

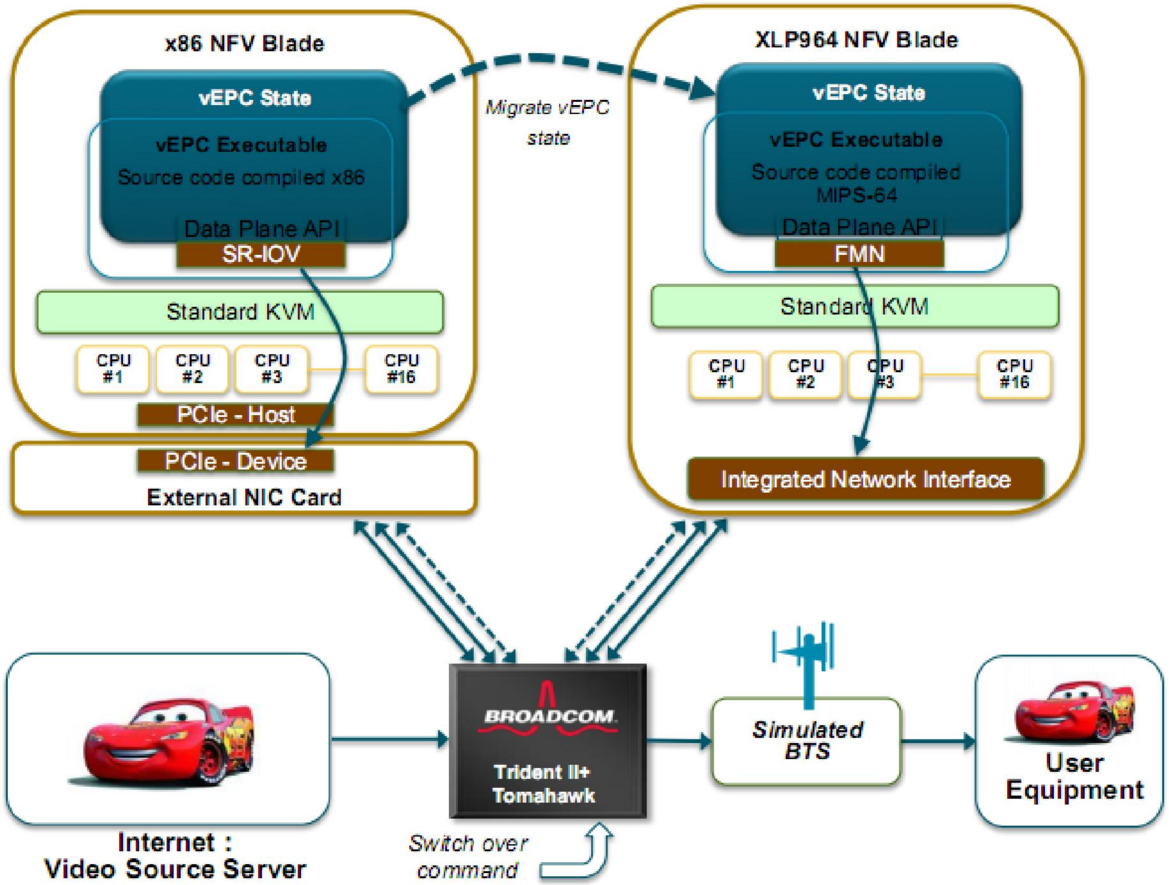

PoC Details:

- 1. Show vEPC with video traffic flowing through to the simulated BTS running on x86 system.
- 2. VNF on the x86 will be placed in a standby state, and the state information will be frozen.
- 3. VNF state will be moved from the x86 blade to the XLP MIPS-64 blade.
- 4. VNF will be restarted on the MIPS-64 blade with the transitioned state.
- 5. Traffic will flow will be rerouted to the vEPC running on the XLP MIPS-64

Source: ETSI Ongoing PoC

Distributed-NFV

PoC are being developed based on **centralized NFVI architectures** and centralized VNF deployment

However, there is **also a need to deploy some functions out at the customer edge**. The ability to support the deployment of virtualized functions at the customer edge **requires a Distributed NFV (D-NFV) architecture**

Omniscient D-NFV orchestrator handles all VNFs and virtual machine (VM) infrastructure, wherever they may be located, and exploits SDN-like mechanisms to achieve optimal VNF placement

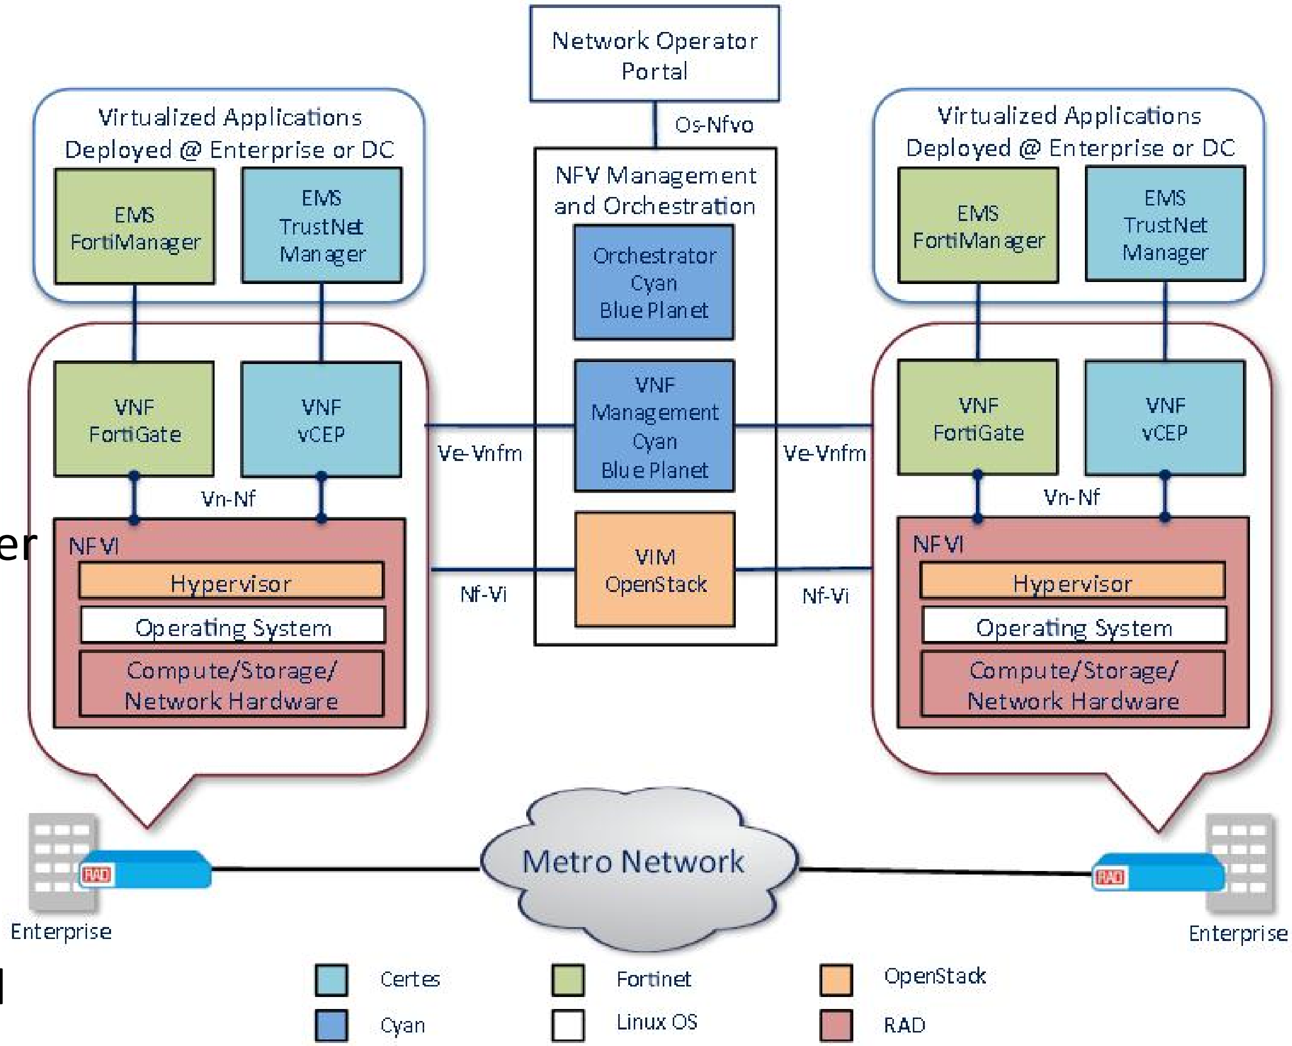

Source: ETSI Ongoing PoC

Multi Vendor on-boarding of vIMS on Cloud Management Frame

Scenario 1 – One-click service deployment.

IMS service is provided by several 3GPP Network Functions, such as CSC, HSS, MMTel, etc. These functions, all from Huawei, are virtualized. **With the pre-defined templates and scripts, all functions can be deployed automatically**, onto the cloud platform provided by DT and ALU.

Scenario 2 – Auto-scaling of VNF

Traffic load generator by a simulator increases and pushes up the workload of the VNF. When the workload exceeds the pre-defined threshold, additional resources (VM) are automatically allocated. In situations of reducing VNF capacity due to decreasing traffic load, similar in reverse direction

Scenario 3 – Automated healing of VNF

When a VM containing a component of a VNF (VNFC) fails, a new VM will be automatically allocated and created with appropriate component instantiated on it. This process heals the VNF with no service interruption.

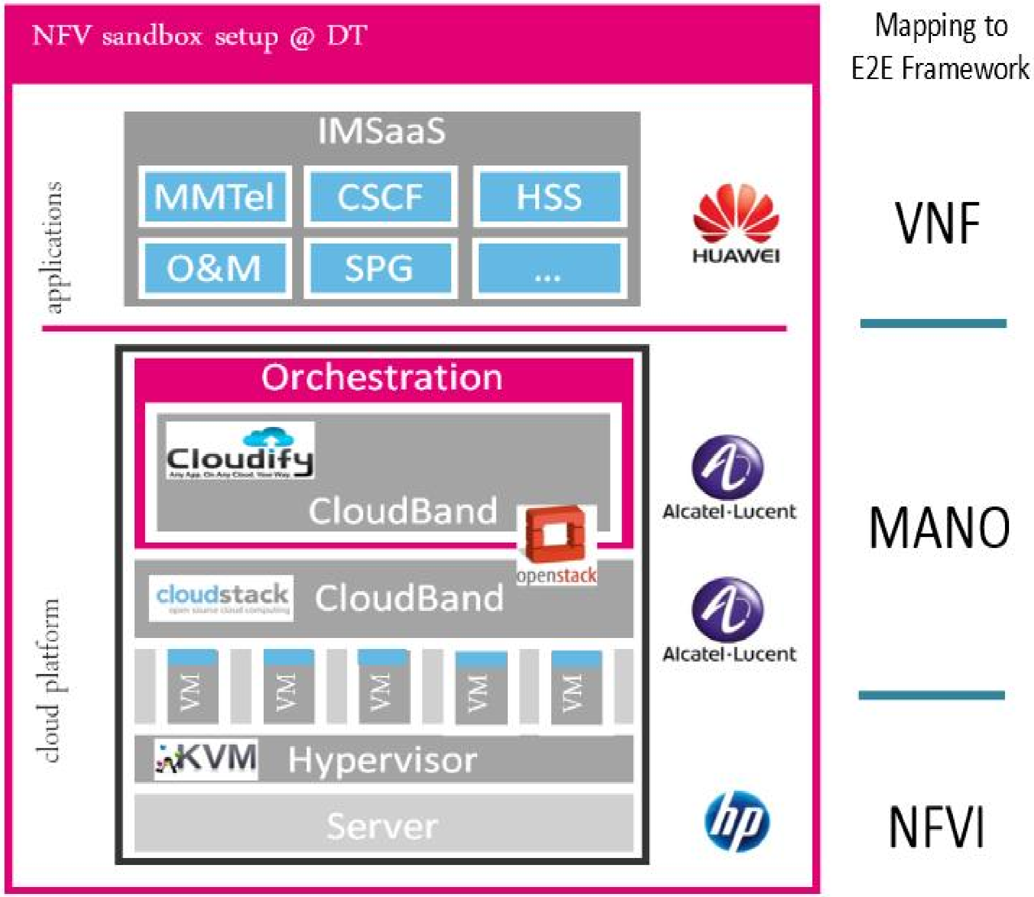

CloudBand is the Alcatel-Lucent Cloud Platform

Source: ETSI Ongoing PoC

INDEPENDENT POCS

OpenNaaS

- OpenNaaS is an open source platform for provisioning network resources.
  - It allows the deployment and automated configuration of dynamic network infrastructures and defines a vendor-independent interface to access services provided by these resources
- OpenNaaS provides support for a variety of resources such as:
  - optical switches, routers, IP networks and Bandwidth on Demand domains,
  - but, more importantly, it is easy to add new resources and their capabilities as an extension
- The core development team is part of Professional Services of the DANA department at i2CAT Foundation (Mantychore FP7)

OpenNaaS Architecture

Intelligence Layer
common web services
connectors for open source
cloud management

Abstract Resource Layer
NaaS resides

Single CLI for Resources
Reusable Building Blocks

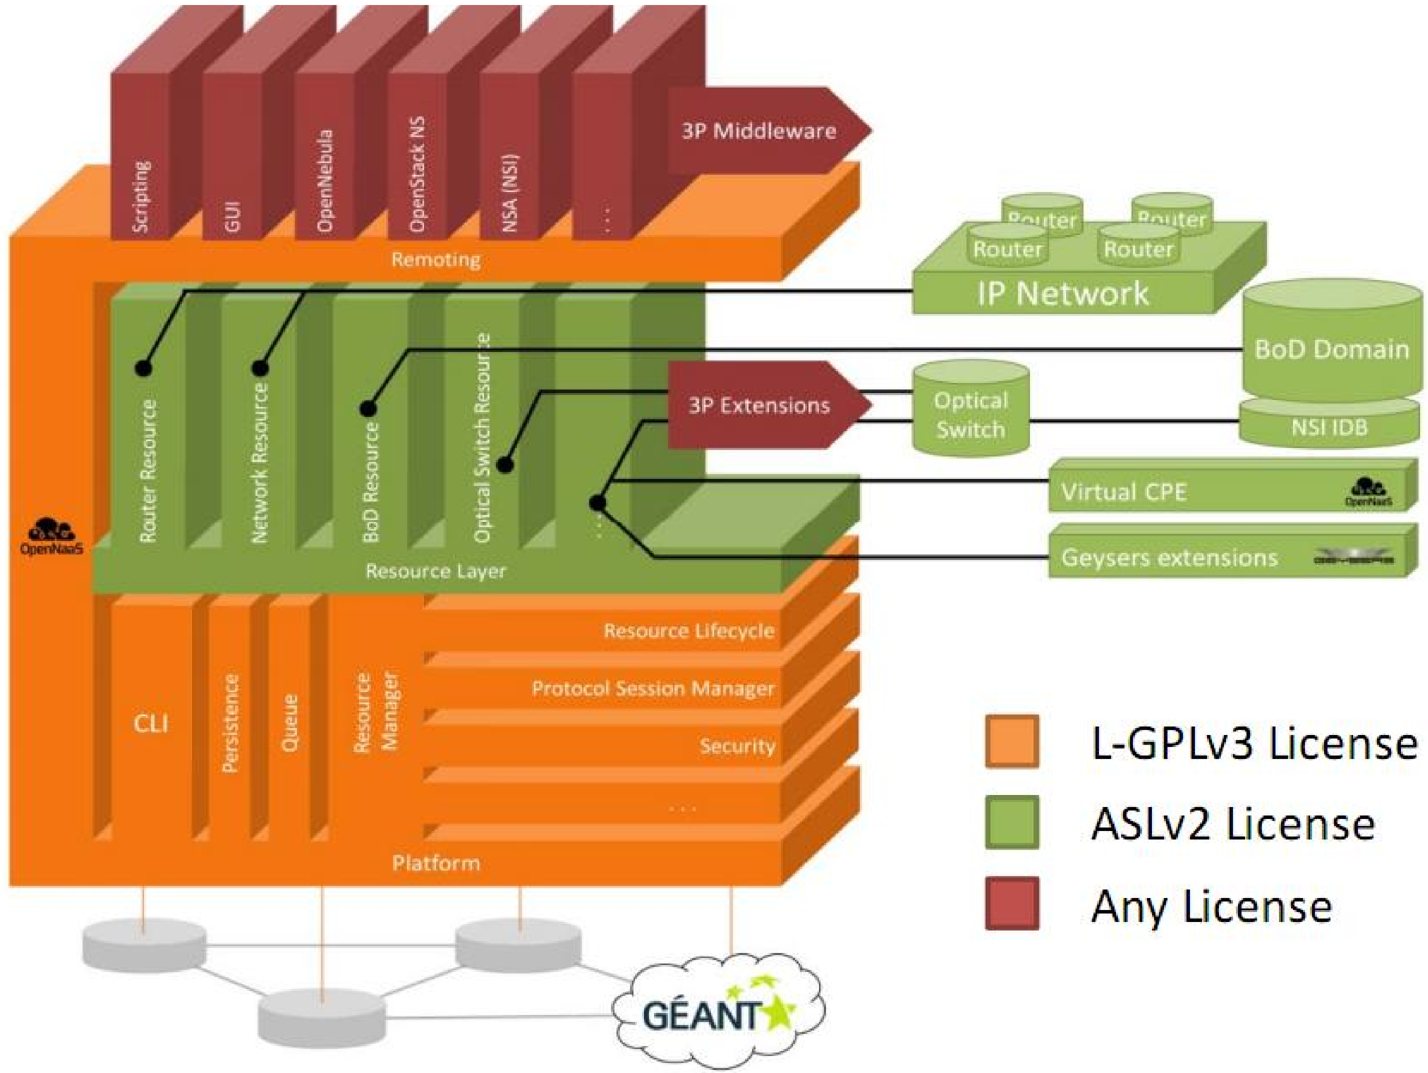

The platform is based on a OSGI (Open Service Gateway initiative) R4 component container

EANTC-NFV Showcase

- European Advanced Networking Test Center (EANTC Berlin, Germany)
  - Vendor independent network quality assurance since 1991
  - Test and certification of network components for manufacturers
  - Network design consultancy and proof of concept testing for service providers

EANTC-NFV Multi-Vendor NFV Showcase

Platform for NFV demonstrations

- Based on ETSI NFV ISG use cases (NFV-009)
- Focused on requirements defined in NFV-012 (Proof of Concept Framework)
- Provides feedback to the ETSI NFV ISG

Target participants

- Open to all Virtual Network Function vendors
- Open to all Virtual Machine/Hypervisor vendors

Real-World Validation

- Verifying advantages provided by Virtual Network Functions
- Monitoring that requirements are met while subscriber traffic is not effected
- Highlighting practical aspects for service providers and carriers

NFV Requirements Verified During the Tests

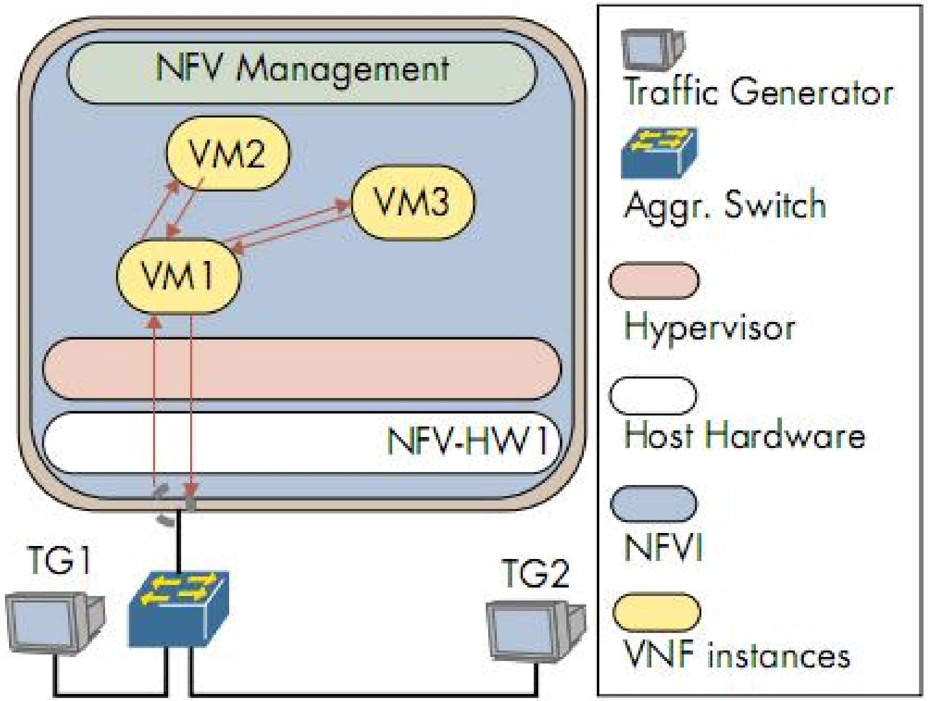

- Instantiation and Provisioning
  - Creation and configuration of virtual network functions
- Portability
  - Moving VNF across hardware
- Elasticity
  - Adjusting resources to the VNF load

EANTC – NFV ShowCases

Huawei VNF Forwarding Graphs and Carrier Grade NAT

- The CG-NAT service intends to provide a **solution for the increasing shortage of IPv4 addresses and transition to IPv6, by implementing nearly any NAT and IPv4-via-IPv6 technique.** The Service Chains make it possible to chain DPI, Parental Controls or other similar functions for flexible services.

Metavswitch Perimeta Session Border Controller

- Metaswitch selected to showcase their Perimeta Session Border Controller (SBC) Virtual Network Function as a Service use case.
- **It uses the concept behind NFV to provide independent distribution and scaling of its signaling (SSC) and media (MSC) components**.

Procera Deep Packet Inspection

- Procera explained that the **Virtualized PacketLogic solution enable network operators to deploy Internet Intelligence pervasively throughout their infrastructure**.
- The solution demonstrated the policy enforcement capabilities of the **PacketLogic solution including application identification, traffic management, and intelligent charging in an NFV** environment.

NFV Showcase

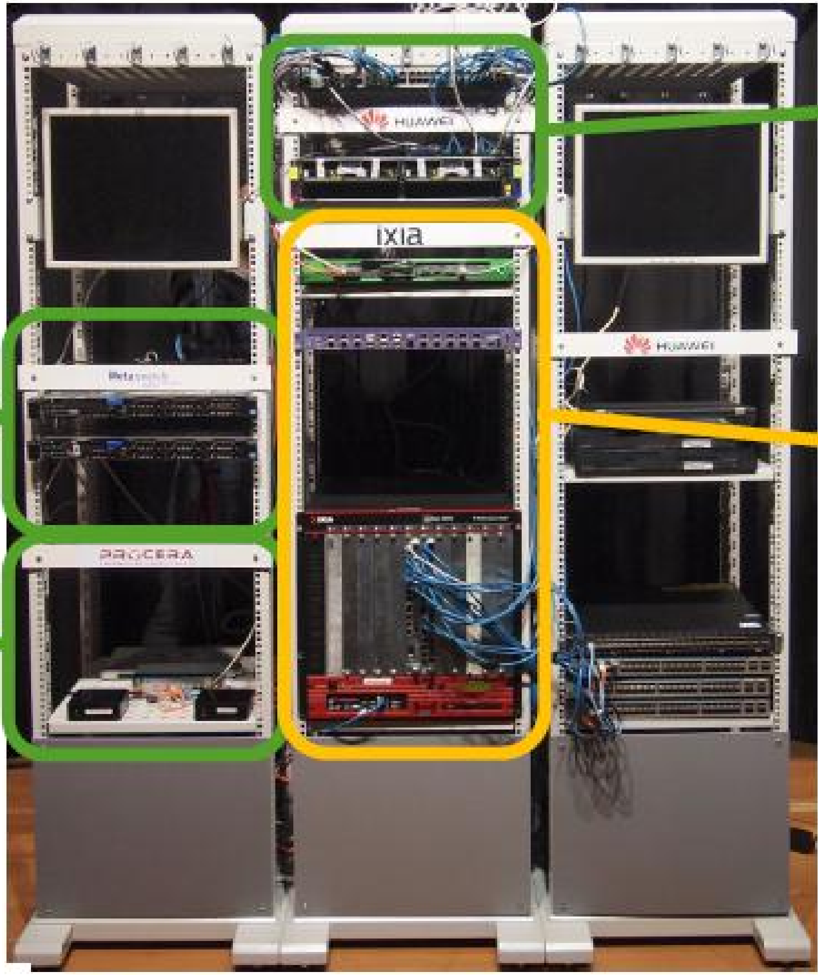

HUAWEI Metaswitch Networks PROCERA EMPOWERING INTELLIGENCE ixia TESTED

DPDK Accelerated OpenvSwitch

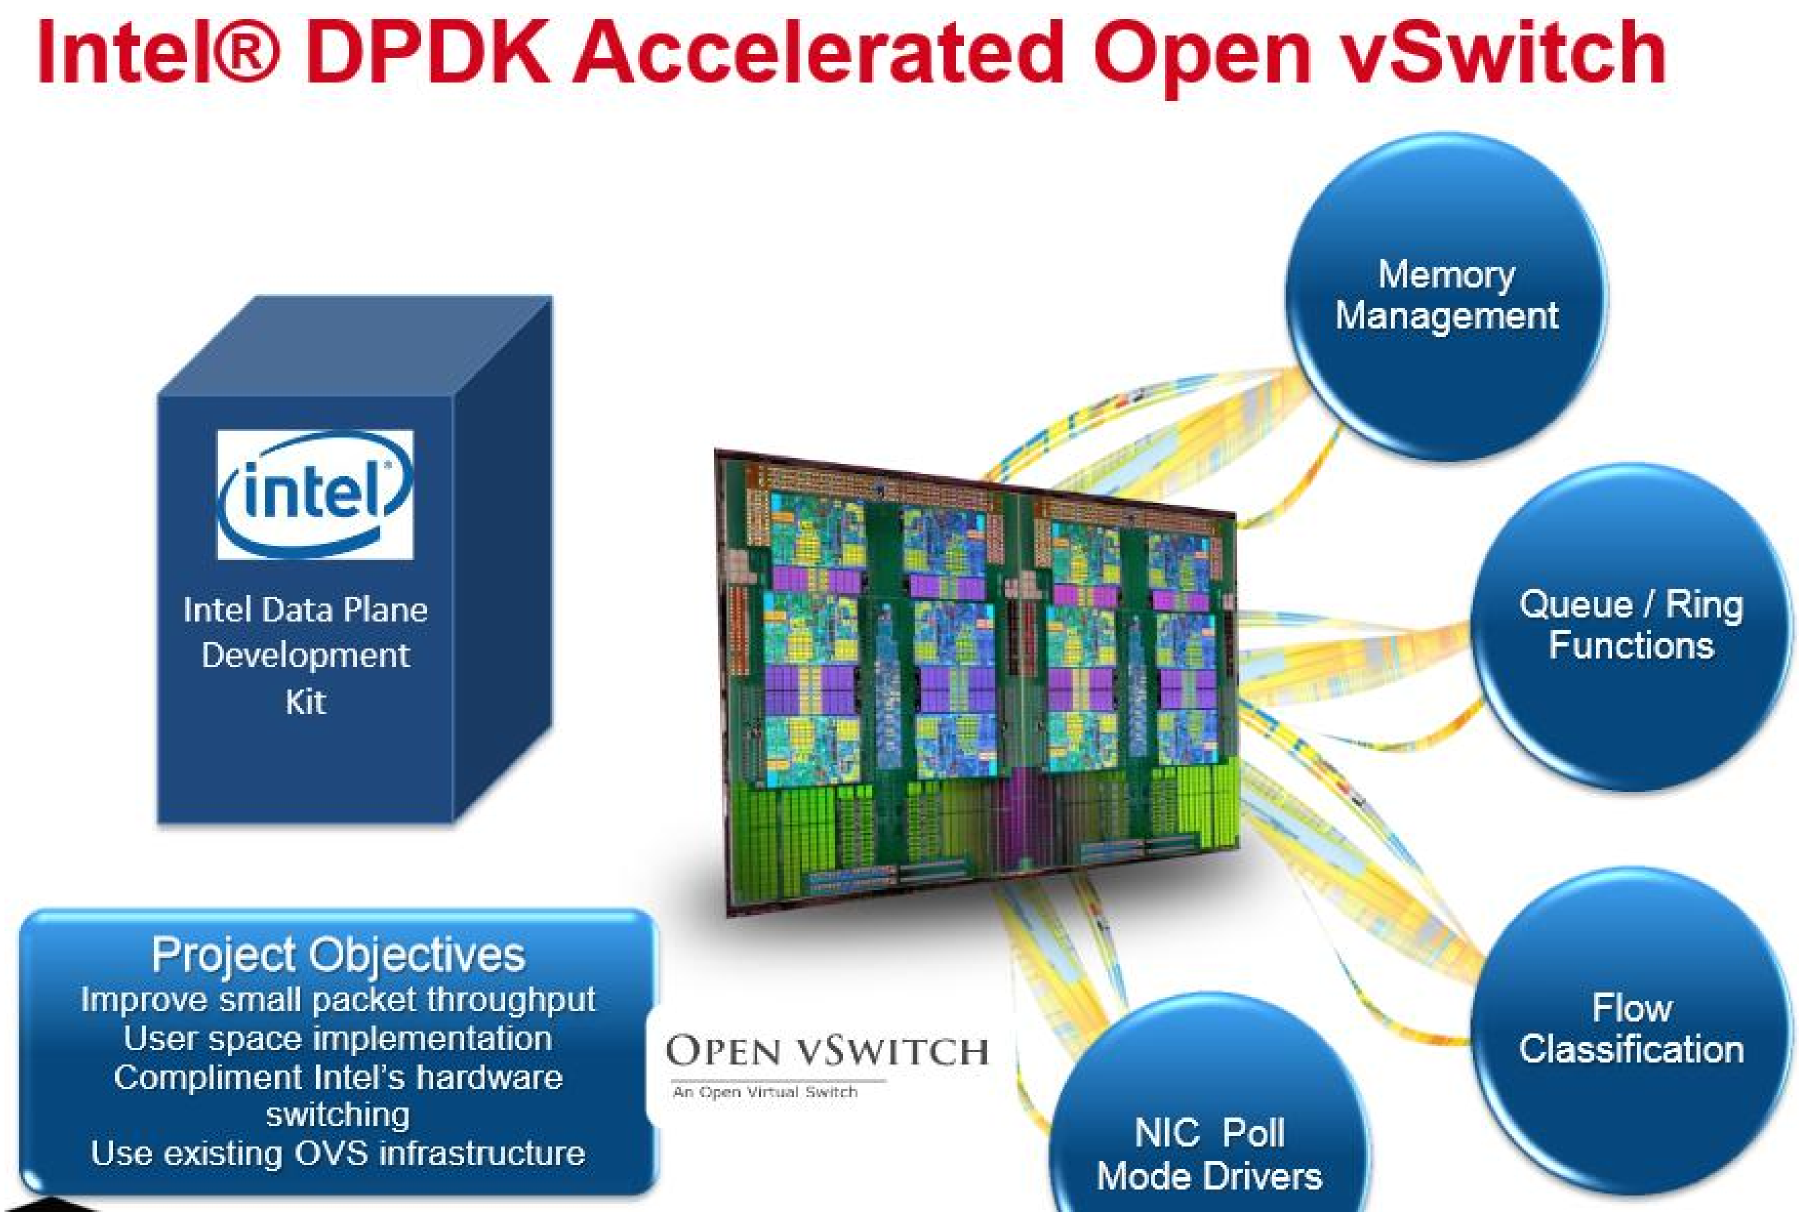

Intel/HP/Wind River Accelerated vSwitch

Figure 4: Intel & Wind River Accelerated Open vSwitch

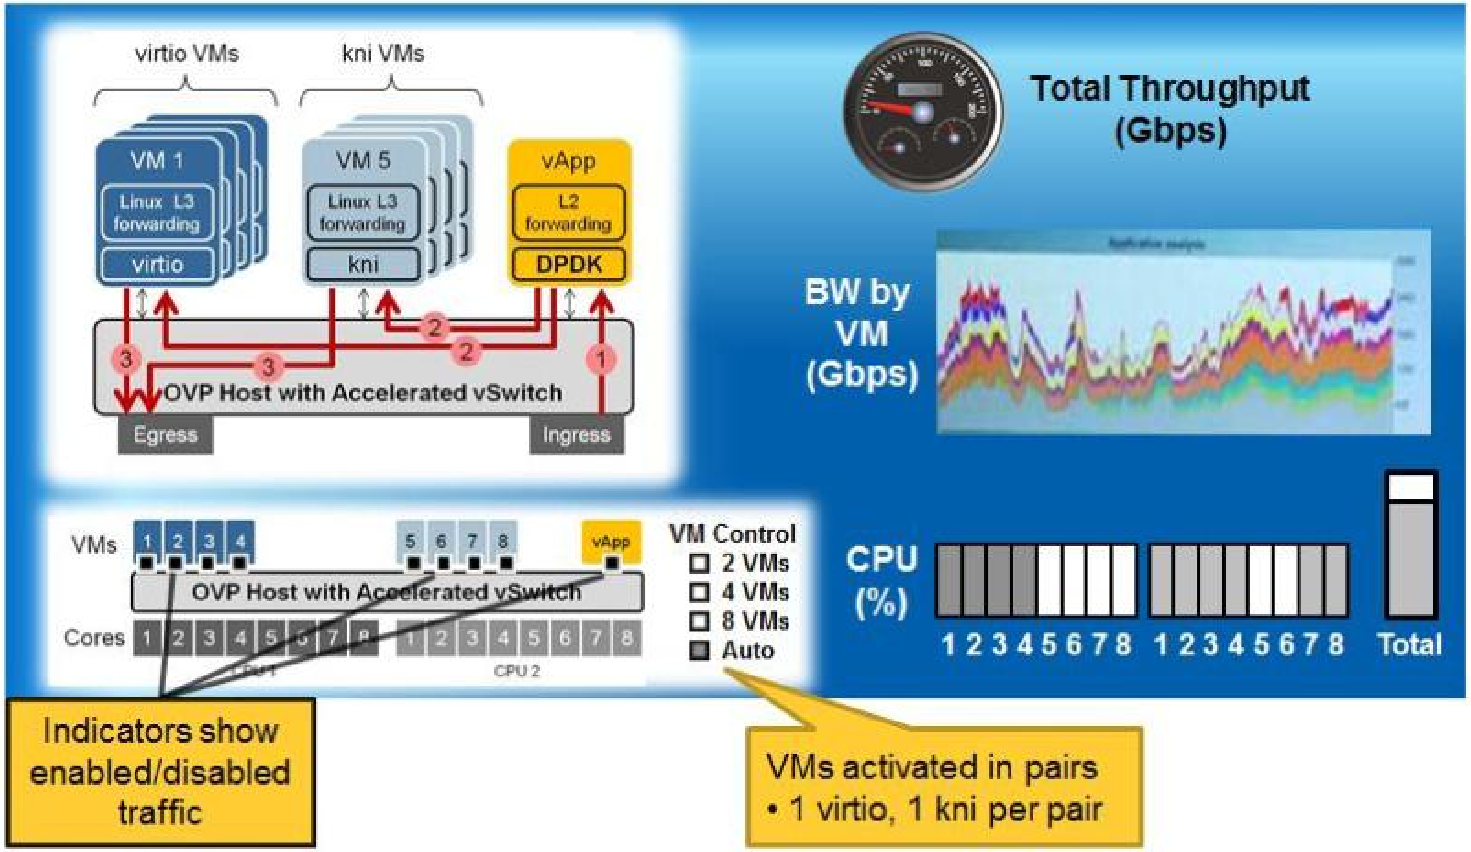

- Combined Intel DPDK, Wind River OVP, and HP hardware
- Reported 10x performance gain in packet switching by bypassing the vSwitch in the Linux kernel
- Provides a "horizontal" platform that can be used across multiple use cases emerging for both SDN and NFV

ENABLING TECHNOLOGIES

Remarkable Enabling Technologies

- Minimalistic OS
  - ClickOS
- Improving Linux i/O
  - Netmap, VALE, Linux NAPI
- Programmable virtual switches / bridges
  - Open vSwitch
- Exploiting x86 for packet processing
  - Intel DPDK
- Some example start-ups
  - LineRate Systems, 6WIND, Midonet, Vyatta (bought by BCD)

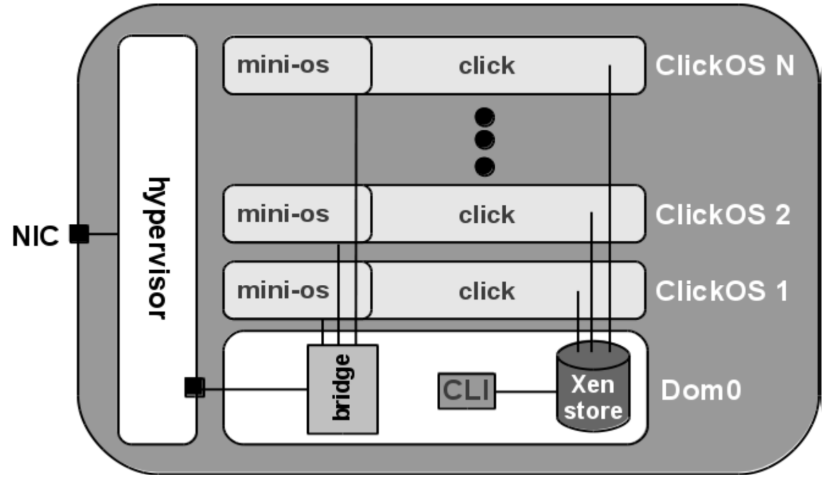

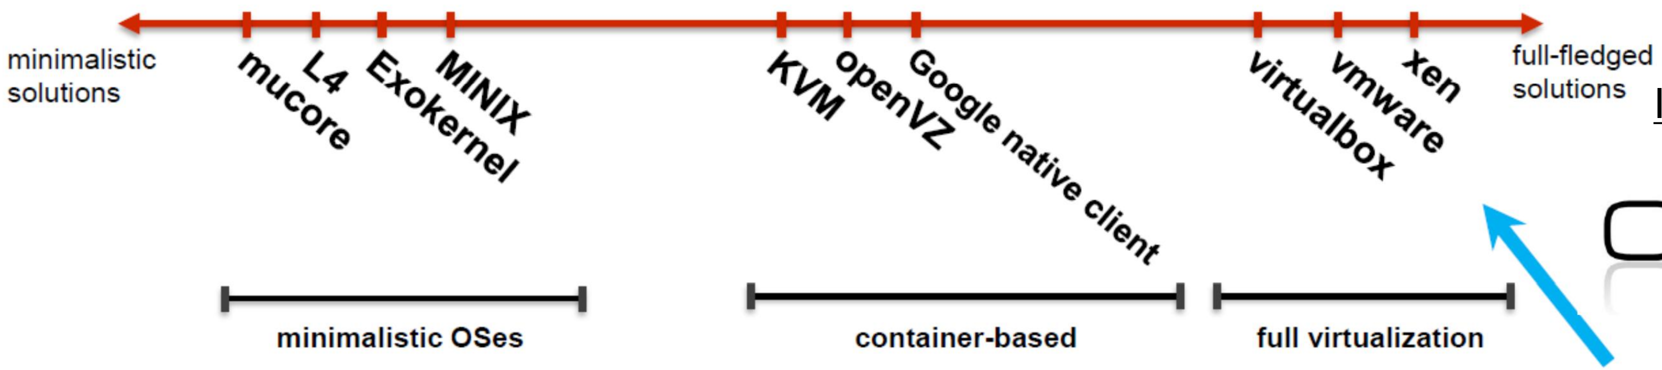

Image source: NEC

Middlebox World

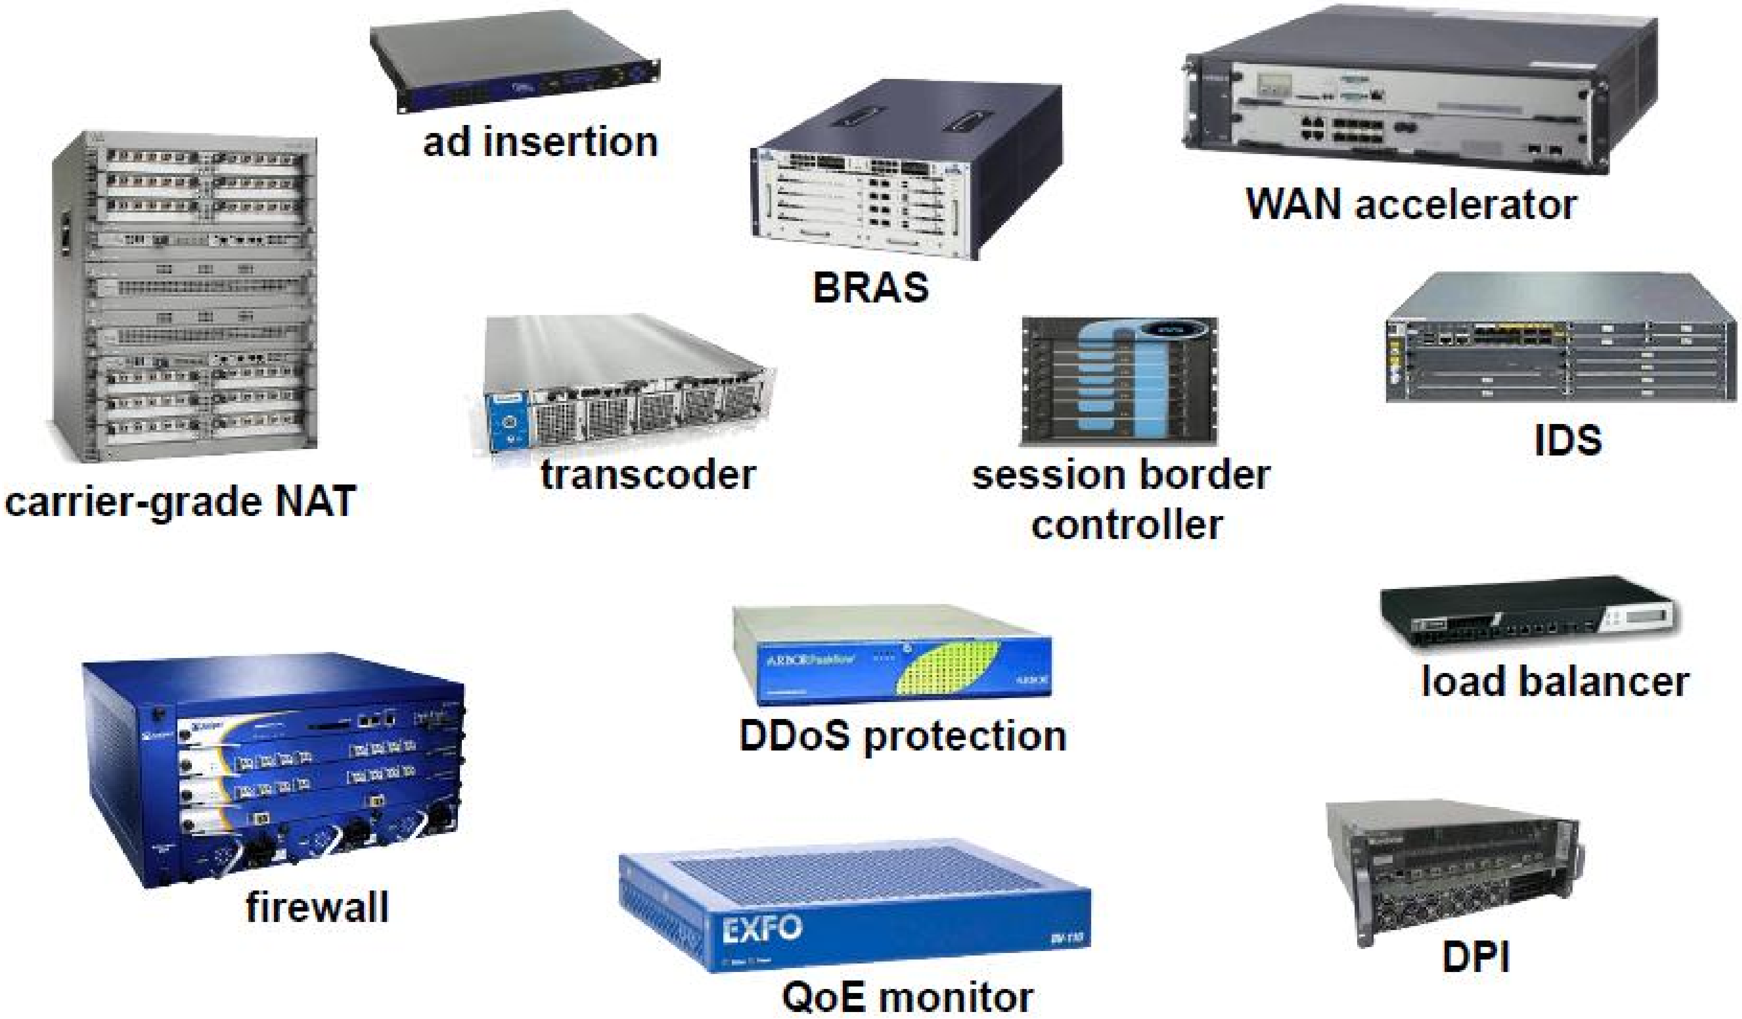

Hardware Middlebox Drawbacks

- Middleboxes are useful, but…
  - Expensive
  - Difficult to add new features, lock-in
  - Difficult to manage
  - Cannot be scaled with demand
  - Cannot share a device among different tenants
  - Hard for new players to enter market
- Clearly shifting middlebox processing to a software-based, multi-tenant platform would address these issues
  - But can it be built using commodity hardware while still achieving high performance?
- ClickOS: tiny Xen-based virtual machine that runs Click

A solution - Click Modular Router

Underlying Idea: Programmable Routing

- Divide -

  Individual router **functionalities like Address-lookup**, Switching Packets, Classification represented as Basic Building blocks
- And Conquer -
  - **Complete Router designed by assembling building blocks and connecting them using a graph**
  - Follows the Bottom-up architecture
- Two new concepts -
  - **Pull – processing**
  - Flow based Router Context

A little bit of Click Modular Router

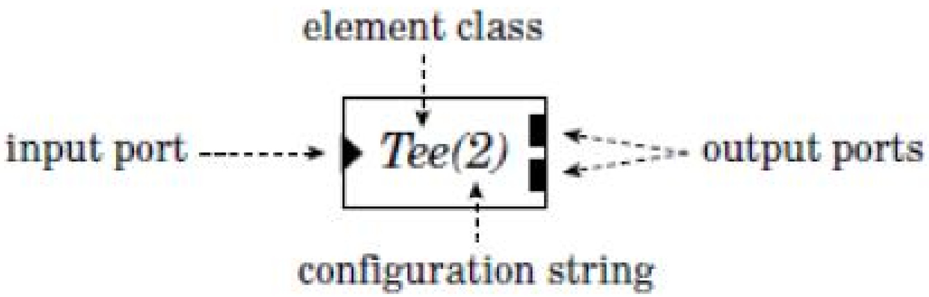

Fig. 1. A sample element. Triangular ports are inputs and rectangular ports are outputs.

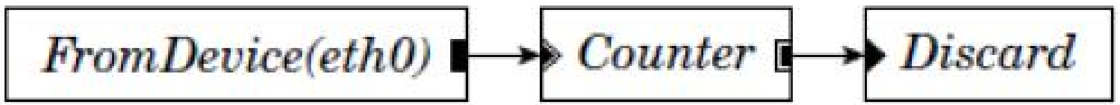

Fig. 2. A router configuration that throws away all packets.

A little bit of Click Modular Router

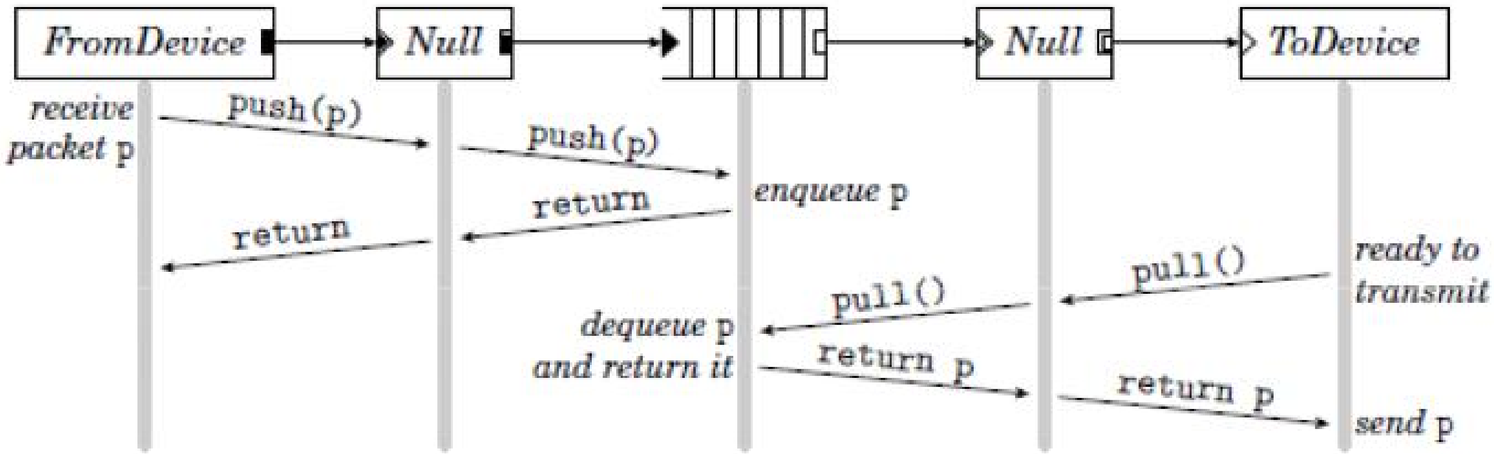

Fig. 3. Push and pull control flow. This diagram shows functions called as a packet moves through a simple router; time moves downwards. The central element is a *Queue*. During the push, control flow moves forward through the element graph starting at the receiving device; during the pull, control flow moves backward through the graph, starting at the transmitting device. The packet p always moves forward.

A little bit of Click Modular Router

Implementation

'Element' is the parent class supporting **all functionalities as virtual functions**

Other classes defined as sub-classes of 'Element'

NullElement derived from Element

NullElement constructor for initializations

Push and Pull in NullElement override parent functions

```
class NullElement : public Element {
 public:
  NullElement()
         { add_input(); add_output(); }
  const char *class_name() const
         { return "Null"; }
  PushOrPull default_processing() const
         { return AGNOSTIC; }
  NullElement *clone() const
         { return new NullElement; }
  void push(int port_number, Packet *p)
         { output(0).push(p); }
  Packet *pull(int port_number)
         { return input(0).pull(); }
};
```

Click Modular Router

Click Language

Simple textual descriptions about declarations and connections

- Configuration string passed as is, as a list separated by commas to the element
- Earlier defined elements used as primitives to define compound elements
- Scope for preprocessing doing pattern matching (namespace)
- Scope for type-systems in place

Declarations: New elements

Connections

Elided form of above specification

```
# a trivial router that drops everything
src :: FromDevice(eth0);
ctr :: Counter;
sink :: Discard;
src -> ctr;
ctr -> sink;

# the same, with anonymous elements
FromDevice(eth0) -> Counter -> Discard;
```

A little bit of Click Modular Router

An example:
IP router

Local information: Elements easy to compose

Annotations

Global information: Chained elements

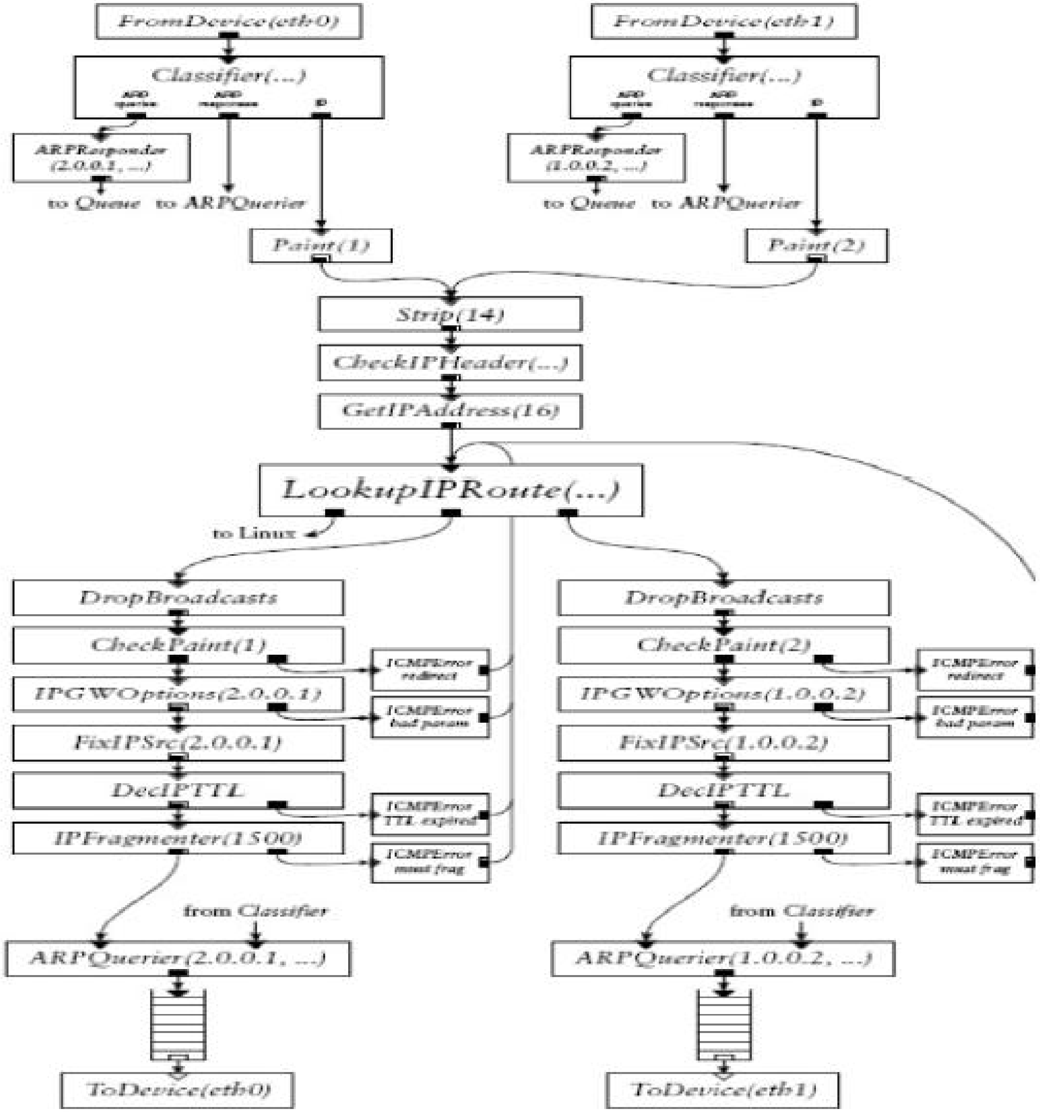

ClickOS Archtecture

Martins, J. et al. Enabling Fast, Dynamic Network Processing with ClickOS. HotSDN 2013.

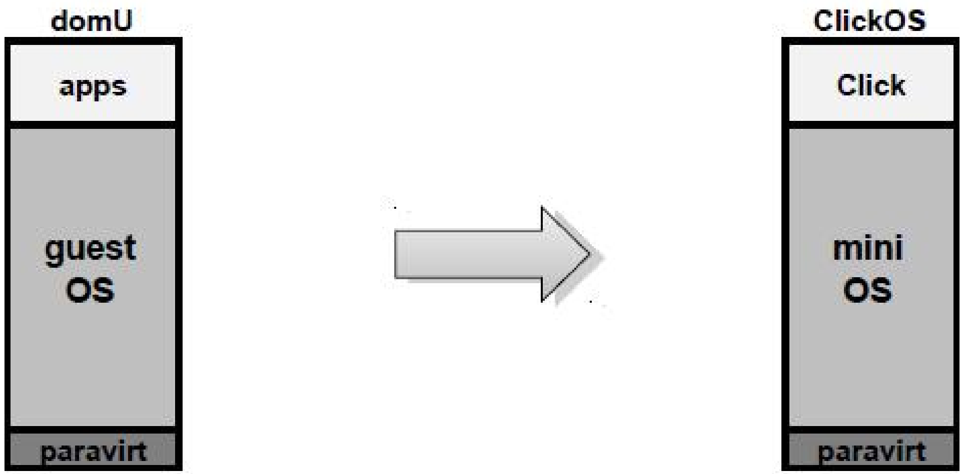

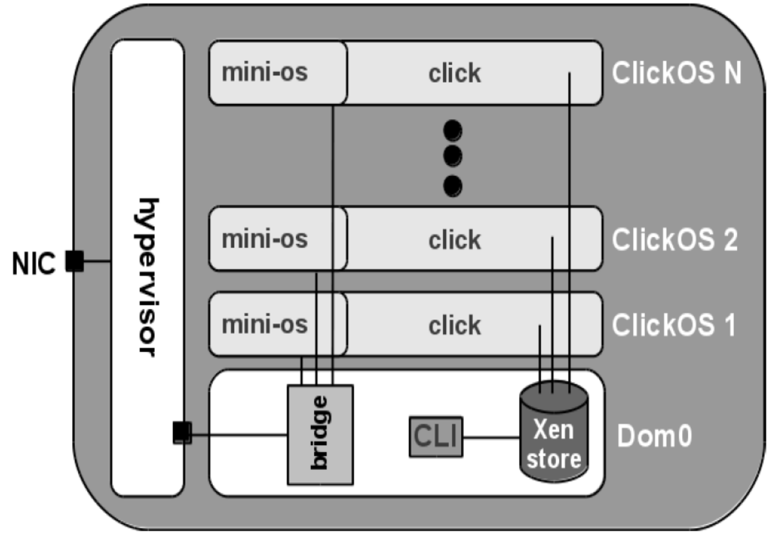

Work consisted of:

- Build system to create ClickOS images (5 MB in size)
- Emulating a Click control plane over MiniOS/Xen
- Reducing boot times (roughly 30 miliseconds)
- Optimizations to the data plane (10 Gb/s for almost all pkt sizes)

Performance Analysis (low performance)

pkt size (bytes)	10Gb rate
64	14.8 Mp/s
128	8.4 Mp/s
256	4.5 Mp/s
512	2.3 Mp/s
1024	1.2 Mp/s
1500	810 Kp/s

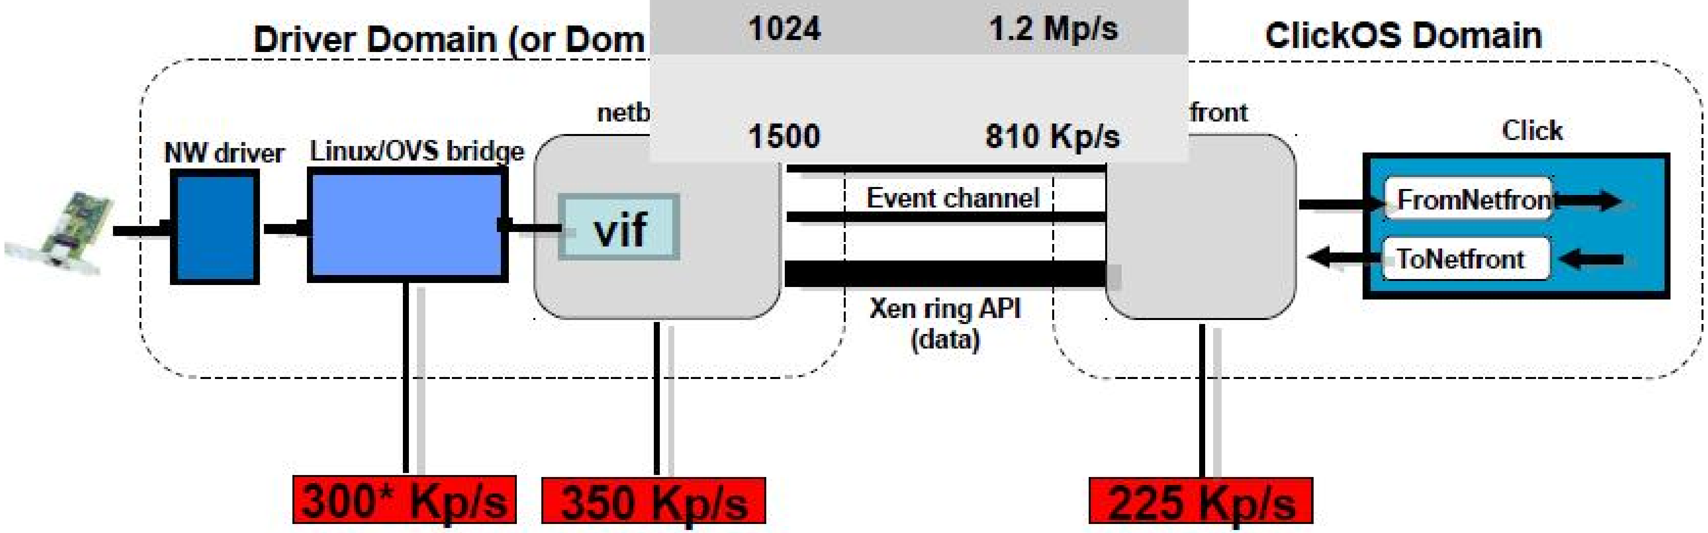

*: maximum-sized packets

Martins, J. et al. Enabling Fast, Dynamic Network Processing with ClickOS. HotSDN 2013.

Background - VALE Software Switch

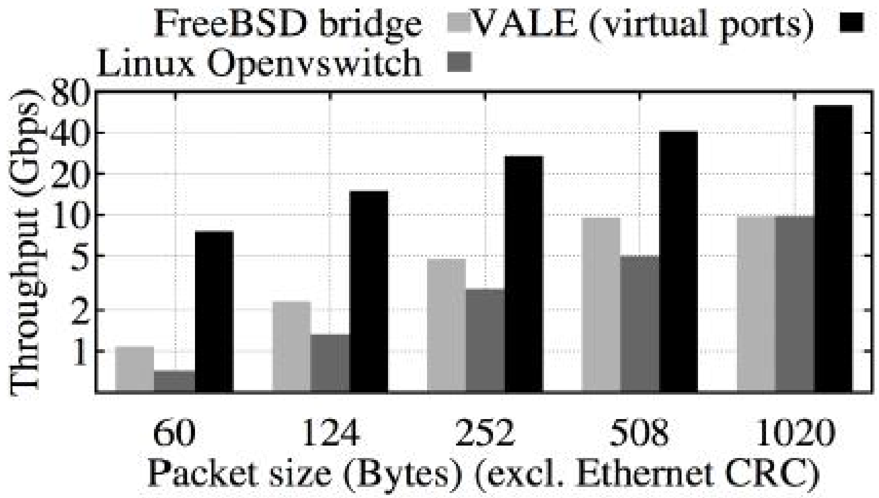

- High performance switch based on netmap API (18 Mpps between virtual ports, one CPU core)
- Packet processing is “modular”
  - Default as learning bridge
  - Modules are independent kernel modules
- Applications use the netmap API

VALE, a Virtual Local Ethernet
http://info.iet.unipi.it/~luigi/vale/
Luigi Rizzo, Giuseppe Lettieri
Universita di Pisa

Optimizing Network I/O – Backend Switch

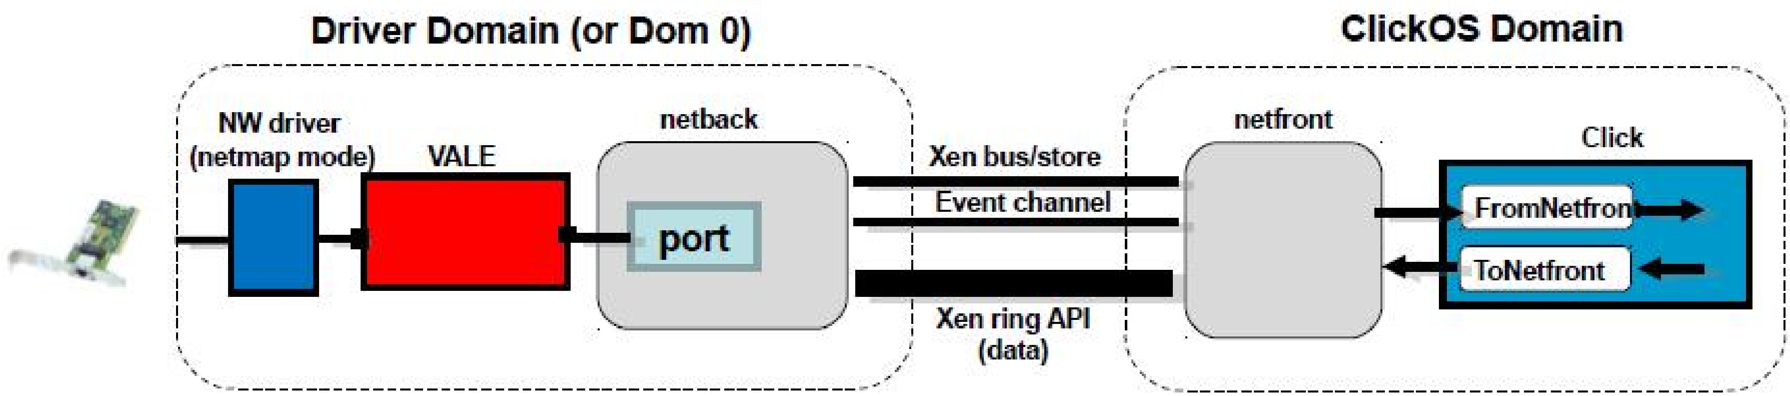

- Introduce VALE as the backend switch
  - NIC switches to netmap-mode
- Slight modifications to the netback driver only
- Batch more I/O requests through multi-page rings
- Removed packet metadata manipulation
- 625 Kpps (1500 size, 2.7x improvement) and 1.2 Mpps (64 size, 4.2x improvement)

Martins, J. et al. Enabling Fast, Dynamic Network Processing with ClickOS. HotSDN 2013.

Background - netmap

- Fast packet I/O framework
  - 14.88 Mpps on 1 core at 900 Mhz
- Available in FreeBSD 9+
  - Also runs on Linux
- Minimal device driver modifications
  - Critical resources (NIC registers, physical buffer addresses, and descriptors) not exposed to the user
  - NIC works in special mode, bypassing the host stack
- Amortize syscalls cost by using large batches
- Preallocated packet buffers, and memory mapped to userspace

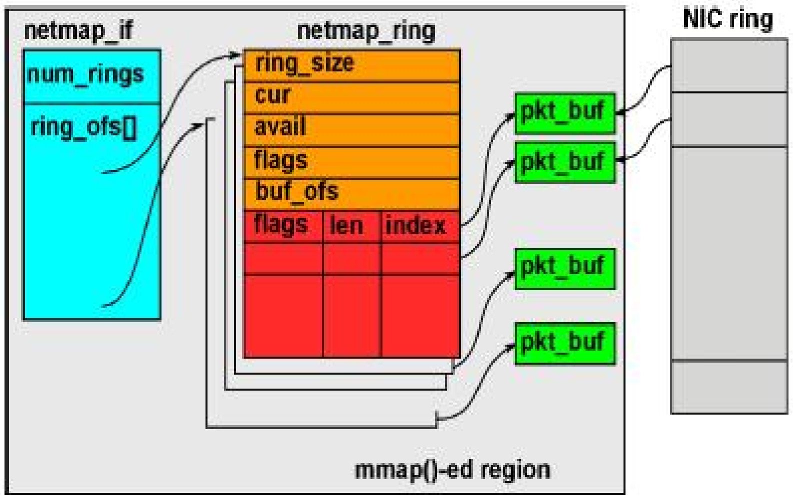

Netmap – a novel framework for fast packet I/O
http://info.iet.unipi.it/~luigi/netmap/
Luigi Rizzo
Universita di Pisa

ClickOS (virtualized) Middlebox Performance

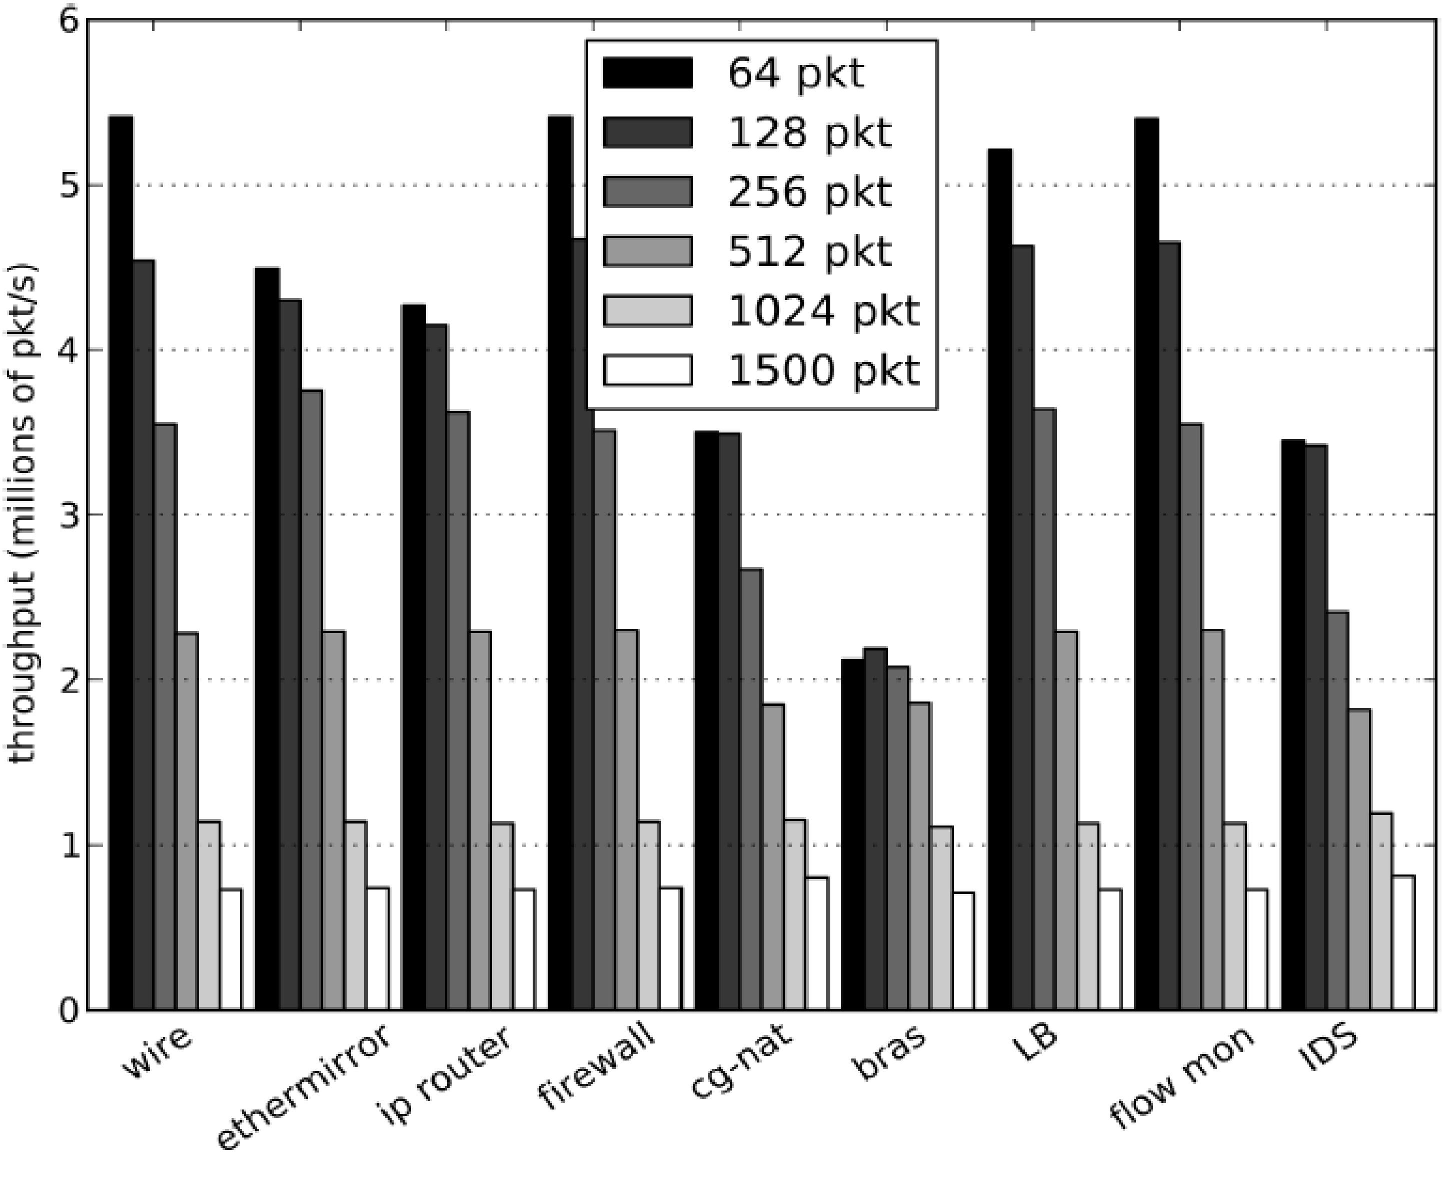

Martins, J. et al. Enabling Fast, Dynamic Network Processing with ClickOS. HotSDN 2013.

ClickOS boot costs and performance

description	function	time
issue create hypercall	libxl__domain_make2	5.244
paravirt. bootloader	libxl_run_bootloader	0.049
prepare domain boot	libxl__build2_pre	0.089
parse, allocate and boot vm image	xc_dom_allocate	0.016
	xc_dom_kernel_path	0.047
	xc_dom_ramdisk	0.001
	xc_dom_boot_xen_init	0.011
	xc_dom_parse_image	0.286
	xc_dom_mem_init	0.007
	xc_dom_boot_mem_init	0.650
	xc_dom_build_image	7.091
	xc_dom_boot_image	0.707
write xen store entries, notify xen store daemon	libxl__build2_post	2.202
init console	init_console_info	0.004
	libxl__need_xenpv_qemu	0.006
	libxl__device_console_add	4.371
	TOTAL	**20.789**

Table 1: Costs of creating a ClickOS virtual machine and booting it up, in milliseconds.

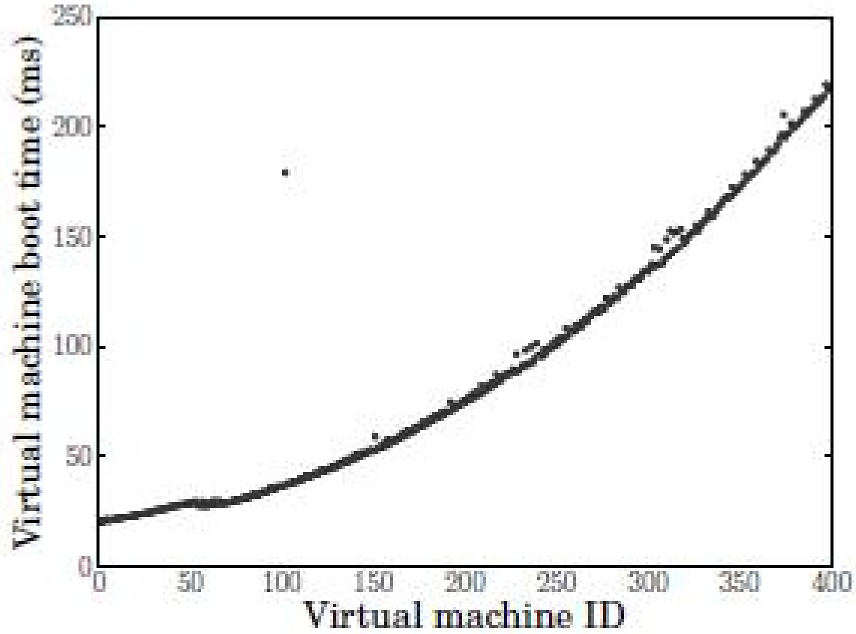

Figure 4: Time to create and boot 400 ClickOS virtual machines on a single server.

Martins, J. et al. Enabling Fast, Dynamic Network Processing with ClickOS. HotSDN 2013.

Intel DPDK

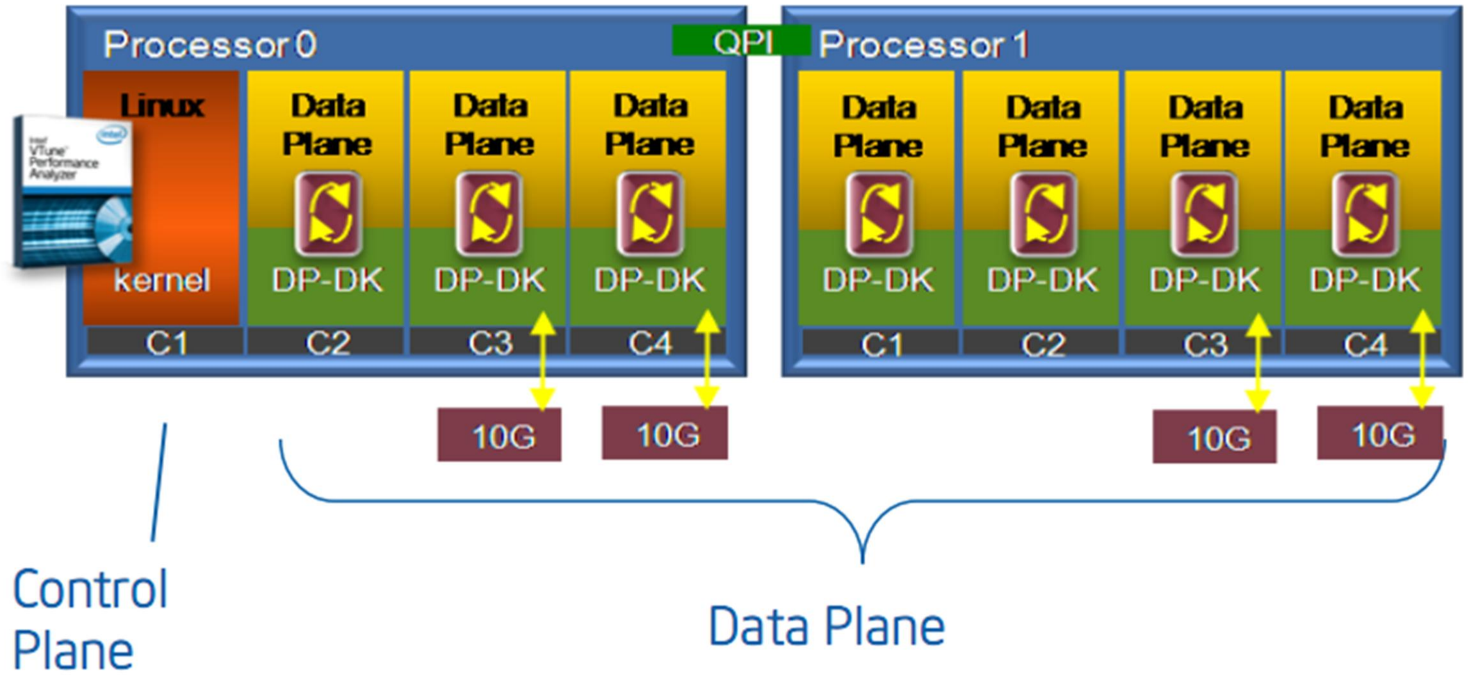

Fonte: Intel Data Plane Development Kit (Intel DPDK) Overview – Packet Processing on Intel Architecture

Intel DPDK

- ⇨ Supported since Intel Atom up to latest Intel Xeon
- ⇨ 32-bit and 64-bit with or without NUMA
- ⇨ No limit on the number of cores or processors
- ⇨ Ideal DRAM allocation for all packets pipelines
- ⇨ Bunch of examples of networking softwares that show the performance improvement
  - Best practices for software architecture
  - Tips on modeling and storing data structures
  - Help compiler to improve the network code
  - Reach levels up to 80Mpps per socket of CPU

Intel DPDK

- ⇨ Optimized NIC Drivers in the user-space
  - Drivers 1/10Gbps
- ⇨ BSD License
- ⇨ Source code available in Intel website (and others)

Intel DPDK

- Buffer and Memory Manager
  - Manage the allocation of objects non-*NUMA* using *hugepages* through *rings*, reducing TLB access, also, perform a pre-allocation of fixed buffer space for each core
- Queue Manager
  - Implements *lockless queues*, allow packets to be processed by different software components with no contention
- Flow Classification
  - Implements hash functions from information tuples, allow packets to be positioned rapidly in their flow paths. Improves *throughput*
- Pool Mode Driver
  - Temporary hold times thus avoiding raise NIC interruptions

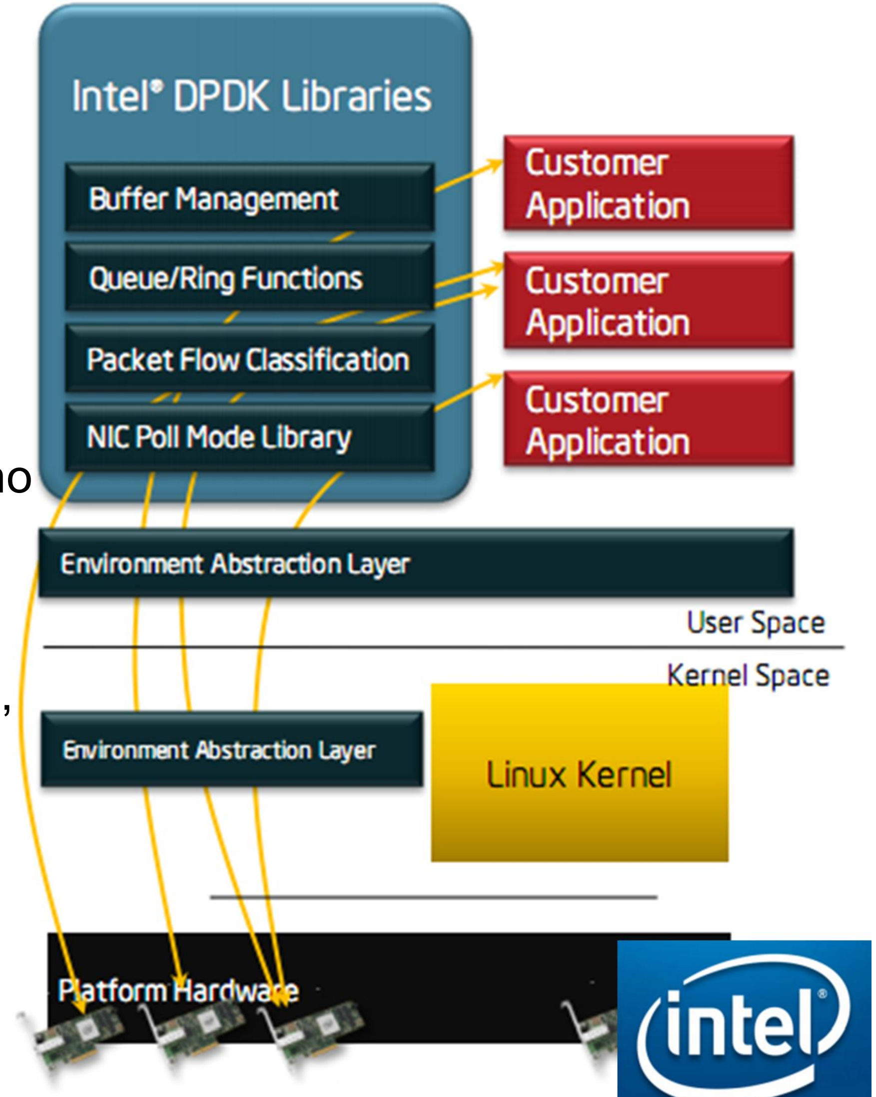

Installing Intel DPDK

- ⇨ Verify CPU
  - `cat /proc/cpuinfo`
- ⇨ Hugepages
  - `cat /proc/meminfo`

```
vi /etc/default/grub
GRUB_CMDLINE_LINUX=... default_hugepagesz=1G hugepagesz=1G
hugepages=4"
```

```
vi /etc/default/grub
GRUB_CMDLINE_LINUX=...hugepages=4096
```

Installing Intel DPDK

- Loading Kernel Module IGB_UIO
  - `modprobe uio`
  - `insmod kmod/igb_uio.ko`
- Free Intel NICs to use IGB_UIO
  - `./tools/pci_unbind.py -status`
  - `./tools/pci_unbind.py --bind=igb_uio <device>`
- Desabling CPU *powersave*
  - `cat /sys/devices/system/cpu/*/cpufreq/scaling_governor`
  - `for gov in /sys/devices/system/cpu/*/cpufreq/scaling_governor ; do echo performance >$gov ; done`

Installing Intel DPDK

- Execute Helloworld
  - `./helloworld -c f -n 4`

  Parameter -c f describes the number of cores and -n the hugepages mem

```
EAL: Requesting 4096 pages of size 2MB from socket 0
EAL: TSC frequency is ~3899765 KHz
EAL: Master core 0 is ready (tid=4e473880)
EAL: Core 1 is ready (tid=4cd72700)
EAL: Core 3 is ready (tid=4a1ce700)
EAL: Core 2 is ready (tid=4a9cf700)
hello from core 1
hello from core 2
hello from core 3
hello from core 0
```

Vyatta Value Proposition

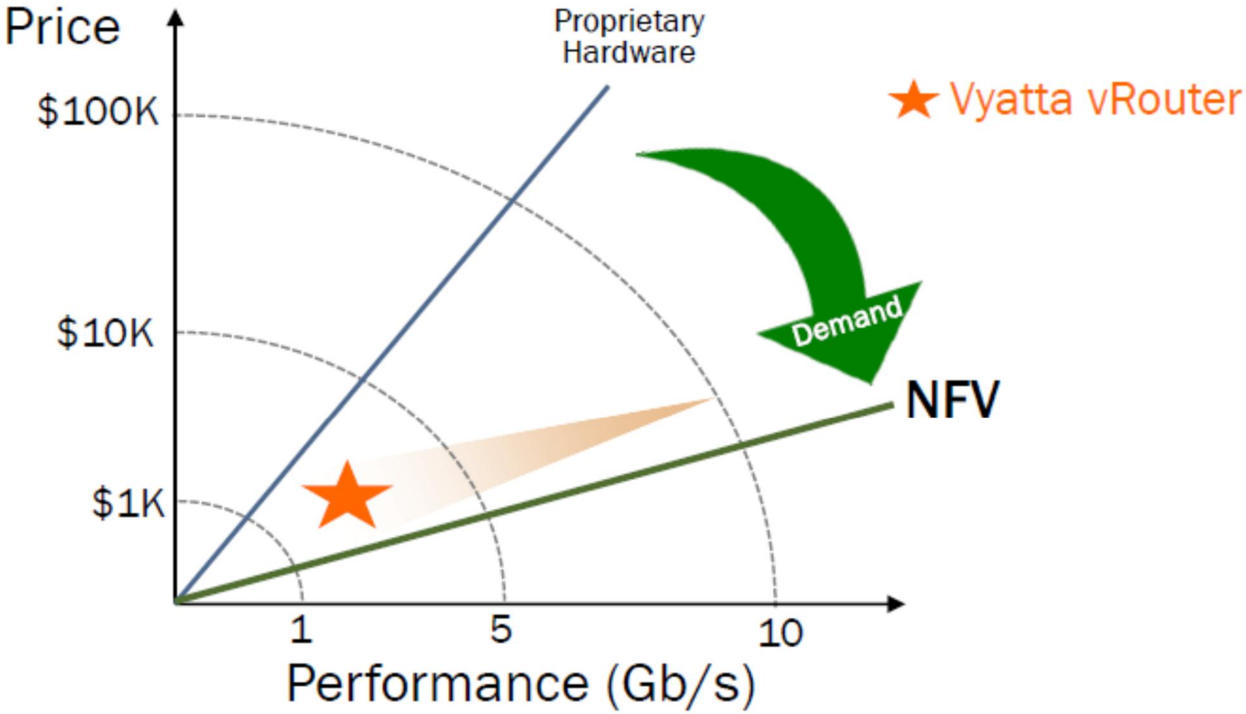

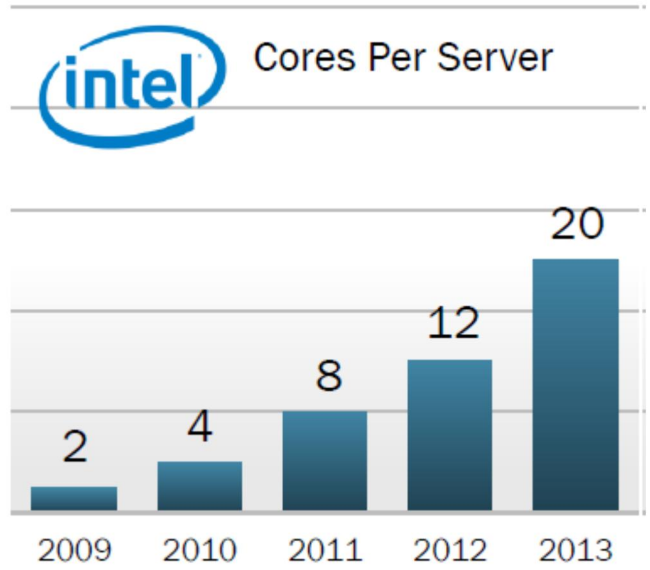

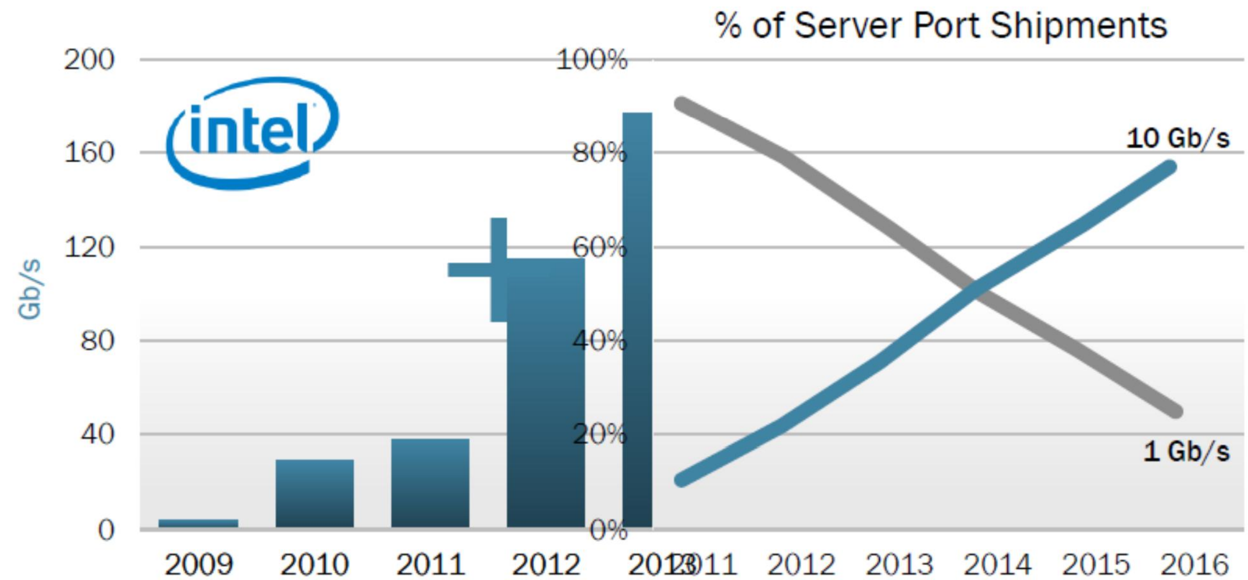

Source: Kelly Herrel (Brocade)

Vyatta vRouter (5400 e 5600)

⇨ Vrouter 5400

⇨Licensing *bare metal*, VM and Amazon

⇨*Features*:

- Network Conectivity
- Firewall
- IPv6
- CLI, GUI and Brocade Vyatta Remote Access API
- Authentication (TACACS+, RADIUS)
- Monitoring and log
- IPSec VPN
- QoS
- High-Availability

BROCADE

Vyatta vRouter (5400 e 5600)

⇨ Vrouter 5600

⇨Licensing *bare metal*, VM and Amazon

⇨*Features*:

- Network Conectivity
- Firewall
- IPv6
- CLI, GUI and Brocade Vyatta Remote Access API
- Authentication (TACACS+, RADIUS)
- Monitoring and log
- IPSec VPN
- QoS
- High-Availability
- vPlane

BROCADE

vPlane

- ⇨ L3 Forwarding Plane is separated from control plane
- ⇨ Based on Intel DPDK
- ⇨ Allow the allocation of the forwarding plane to specific and multiple cores

⇨ ***By isolating forwarding plane to specific x86 cores and cut-through using DPDK (zero copy), allowed vRouter reach 10Gbps per core***

Source: Kelly Herrel (Brocade)

BROCADE

vPlane (high level overview)

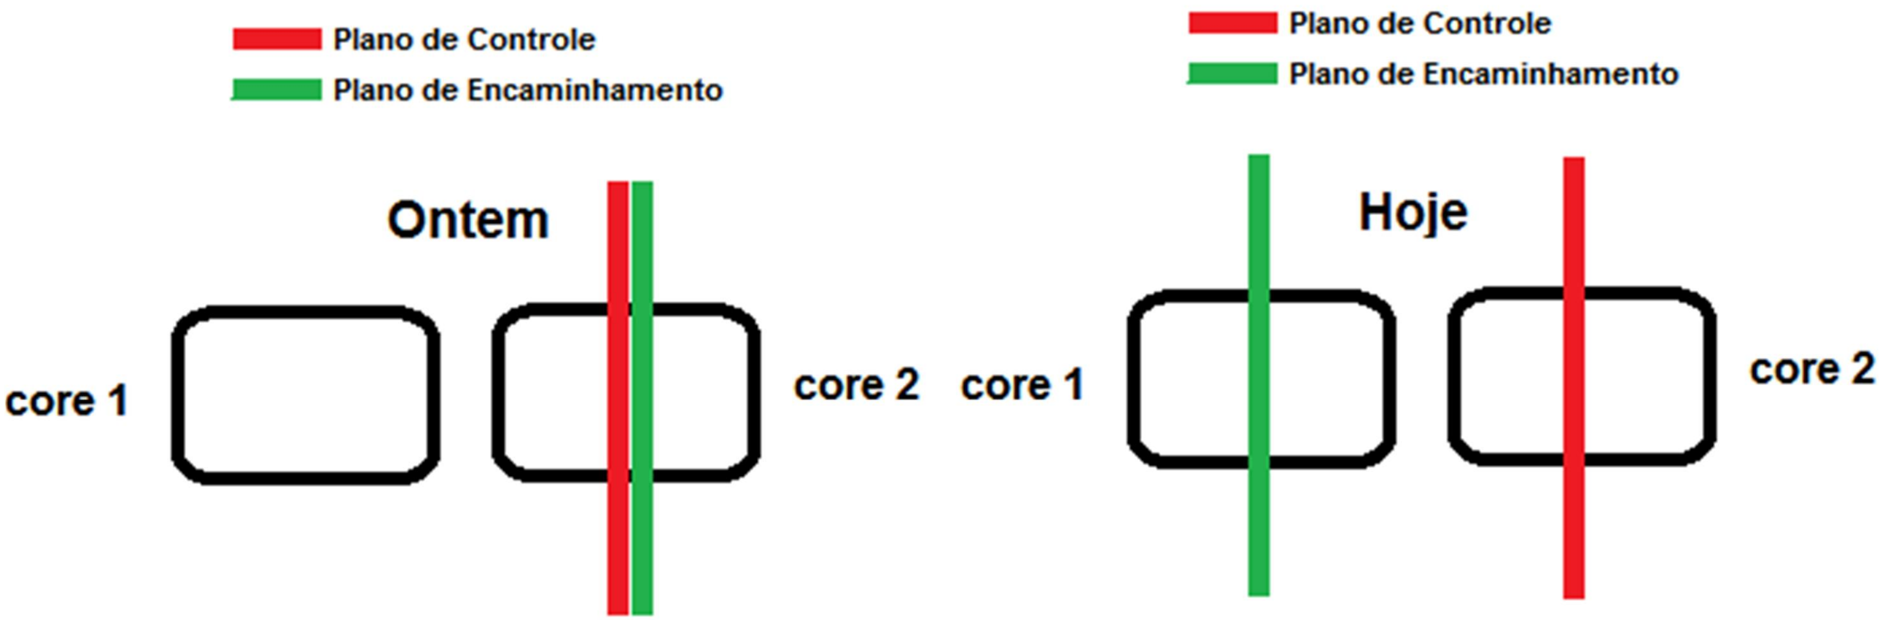

BROCADE

Vyatta: Current Architecture (5400)

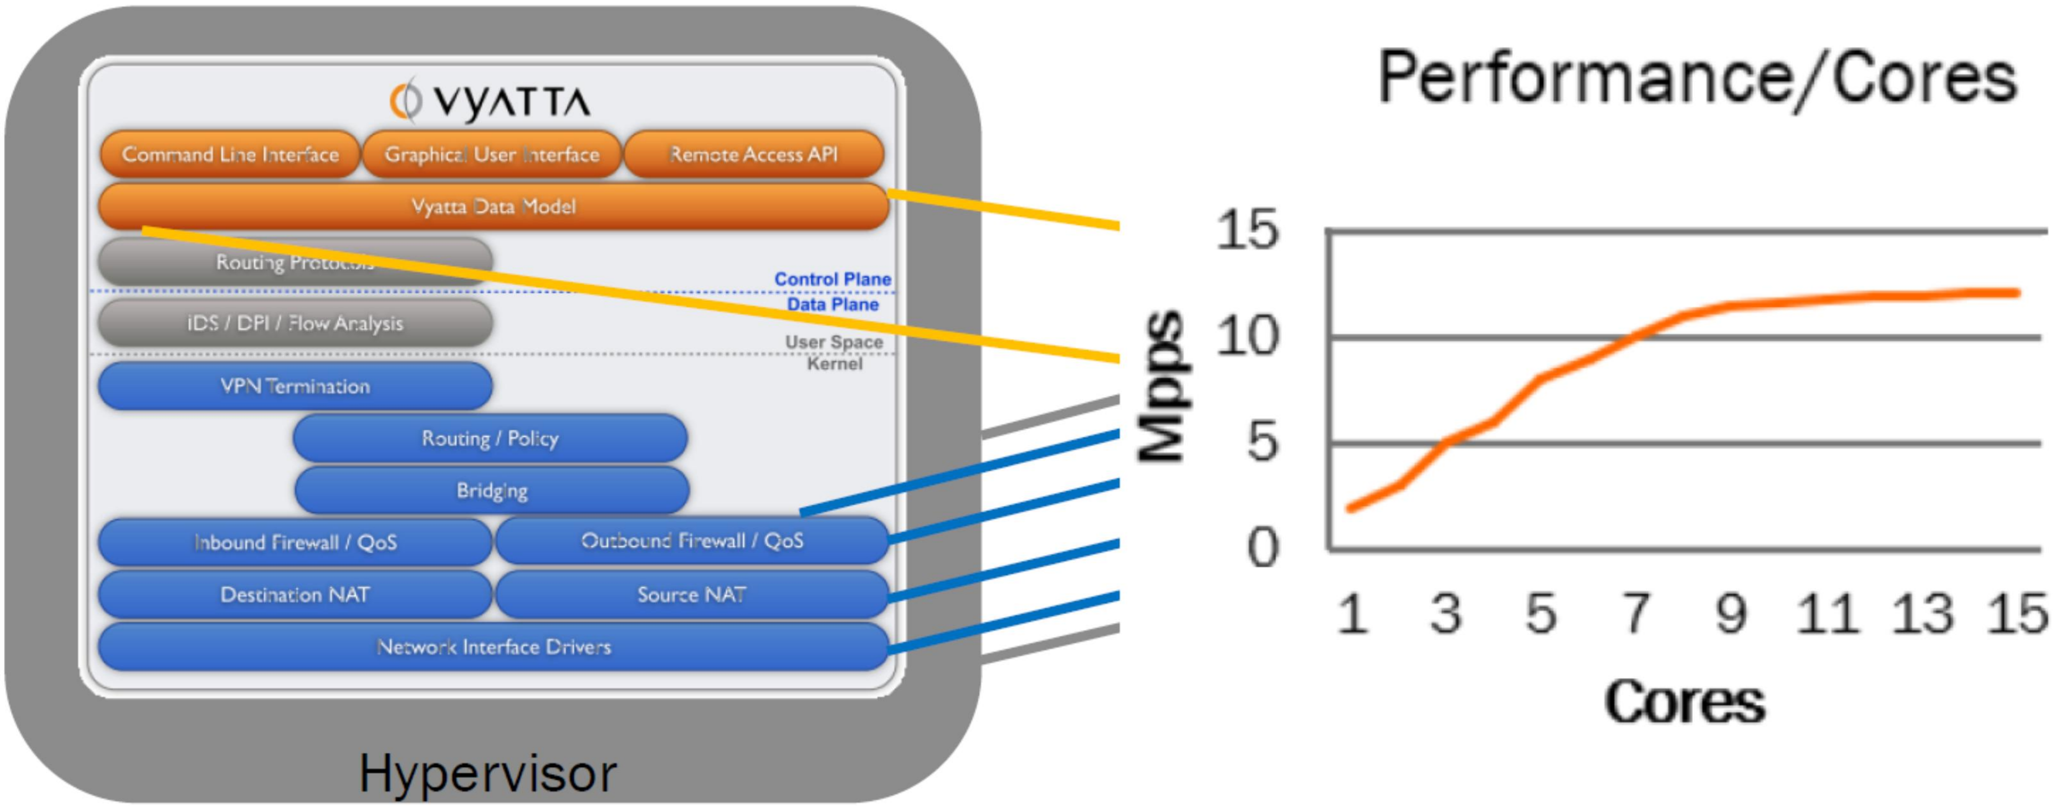

Source: Kelly Herrel (Brocade)

Vyatta: Architecture (5600)

Intel DPDK

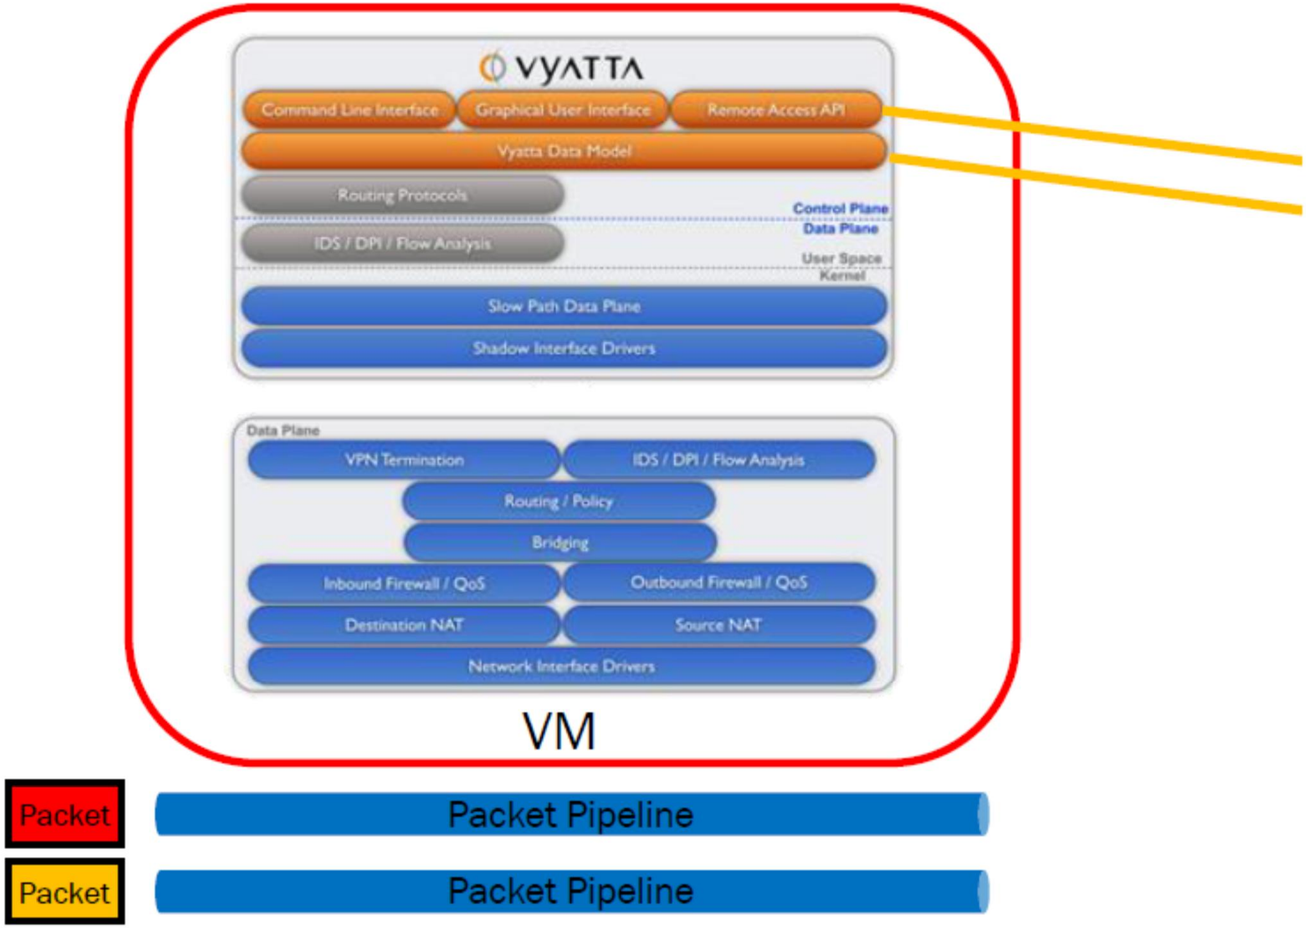

Performance/Cores

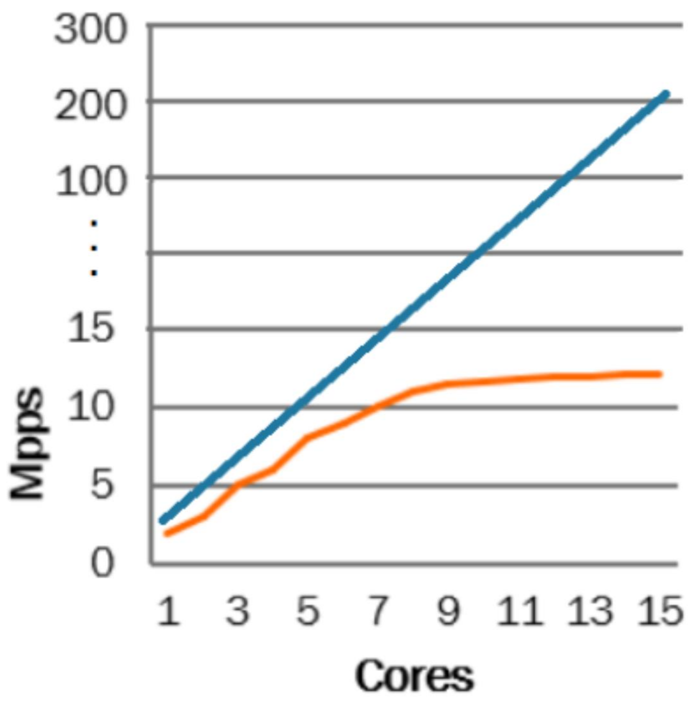

Source: Kelly Herrel (Brocade)

Comparison:
Vyatta vRouter vs Hardware-based

Network virtualization applications.

Requirement	Brocade Vyatta vRouter	Hardware-based Networking
Multifunction Layer 3+ (routing, firewall, VPN, and more)	Yes	Vendor-dependent
Elasticity/scalability	Seamless addition of underlying processor cores	Platform-limited
Multitenancy	Platform-independent Virtual Machine (VM)	Hardware-bound
Hypervisor agnosticism/ awareness	VMware, Hyper-V, Open Xen/ Xen/XenServer, KVM	None
Open management API	Yes	No
On-demand provisioning	Yes	No

BROCADE

Vyatta vRouter 5600

⇨ Using Intel >= 2 cores, NIC Intel (DPDK enabled)

```
[ ok ] Starting configd...done.
[ ok ] Starting vPlane services: huge igb_uio dataplane...
[....] Starting routing daemons: imi nsm mribd pimd msdpd rip
[ ok pd.
[ ok ] Starting vPlane controller...done.
```

⇨ Otherwise

```
[ ok ] Loading cpufreq kernel modules...done (none).
failed (vPlane requires 2 or more logical CPU's).
[ ok ] Starting vPlane controller...done.
```

⇨Without the vPlane, the networking interfaces (even using Intel hardware, thus the need for DPDK) are not recognized by the dataplane, thus vRouter doesn't work.

⇨Obs.: dataplane is a network interface for the vPlane

BROCADE

Vyatta vRouter 5600 (Brocade Company)

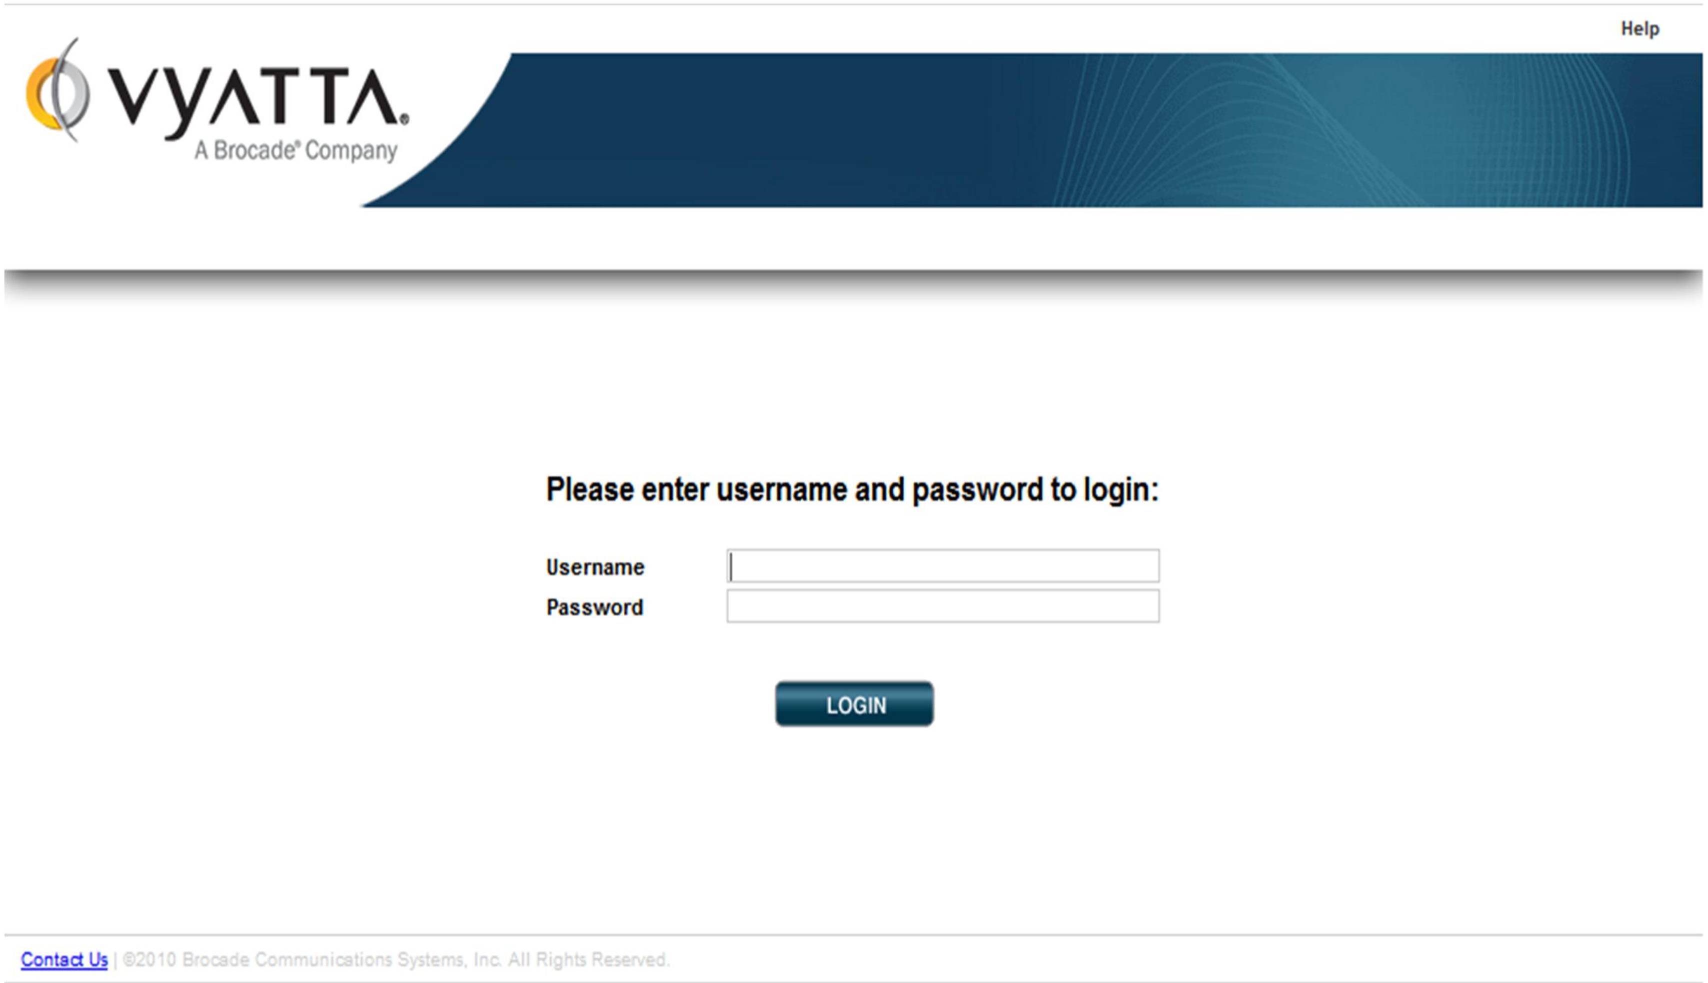

BROCADE

Openstack

OpenStack is a global collaboration of developers and cloud computing technologists producing the ubiquitous open source cloud computing platform for public and private clouds. The project aims to deliver solutions for all types of clouds by being simple to implement, massively scalable, and feature rich. The technology consists of a series of interrelated projects delivering various components for a cloud infrastructure solution.

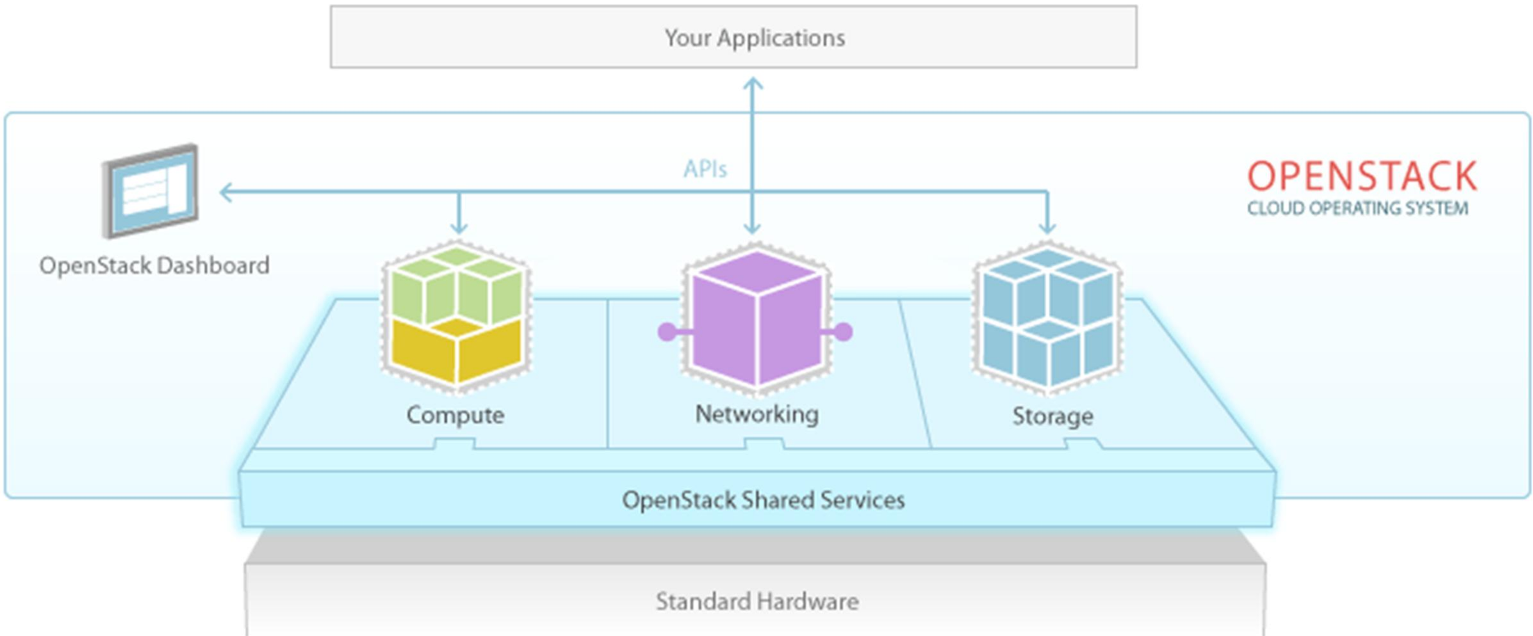

Source: Openstack.org

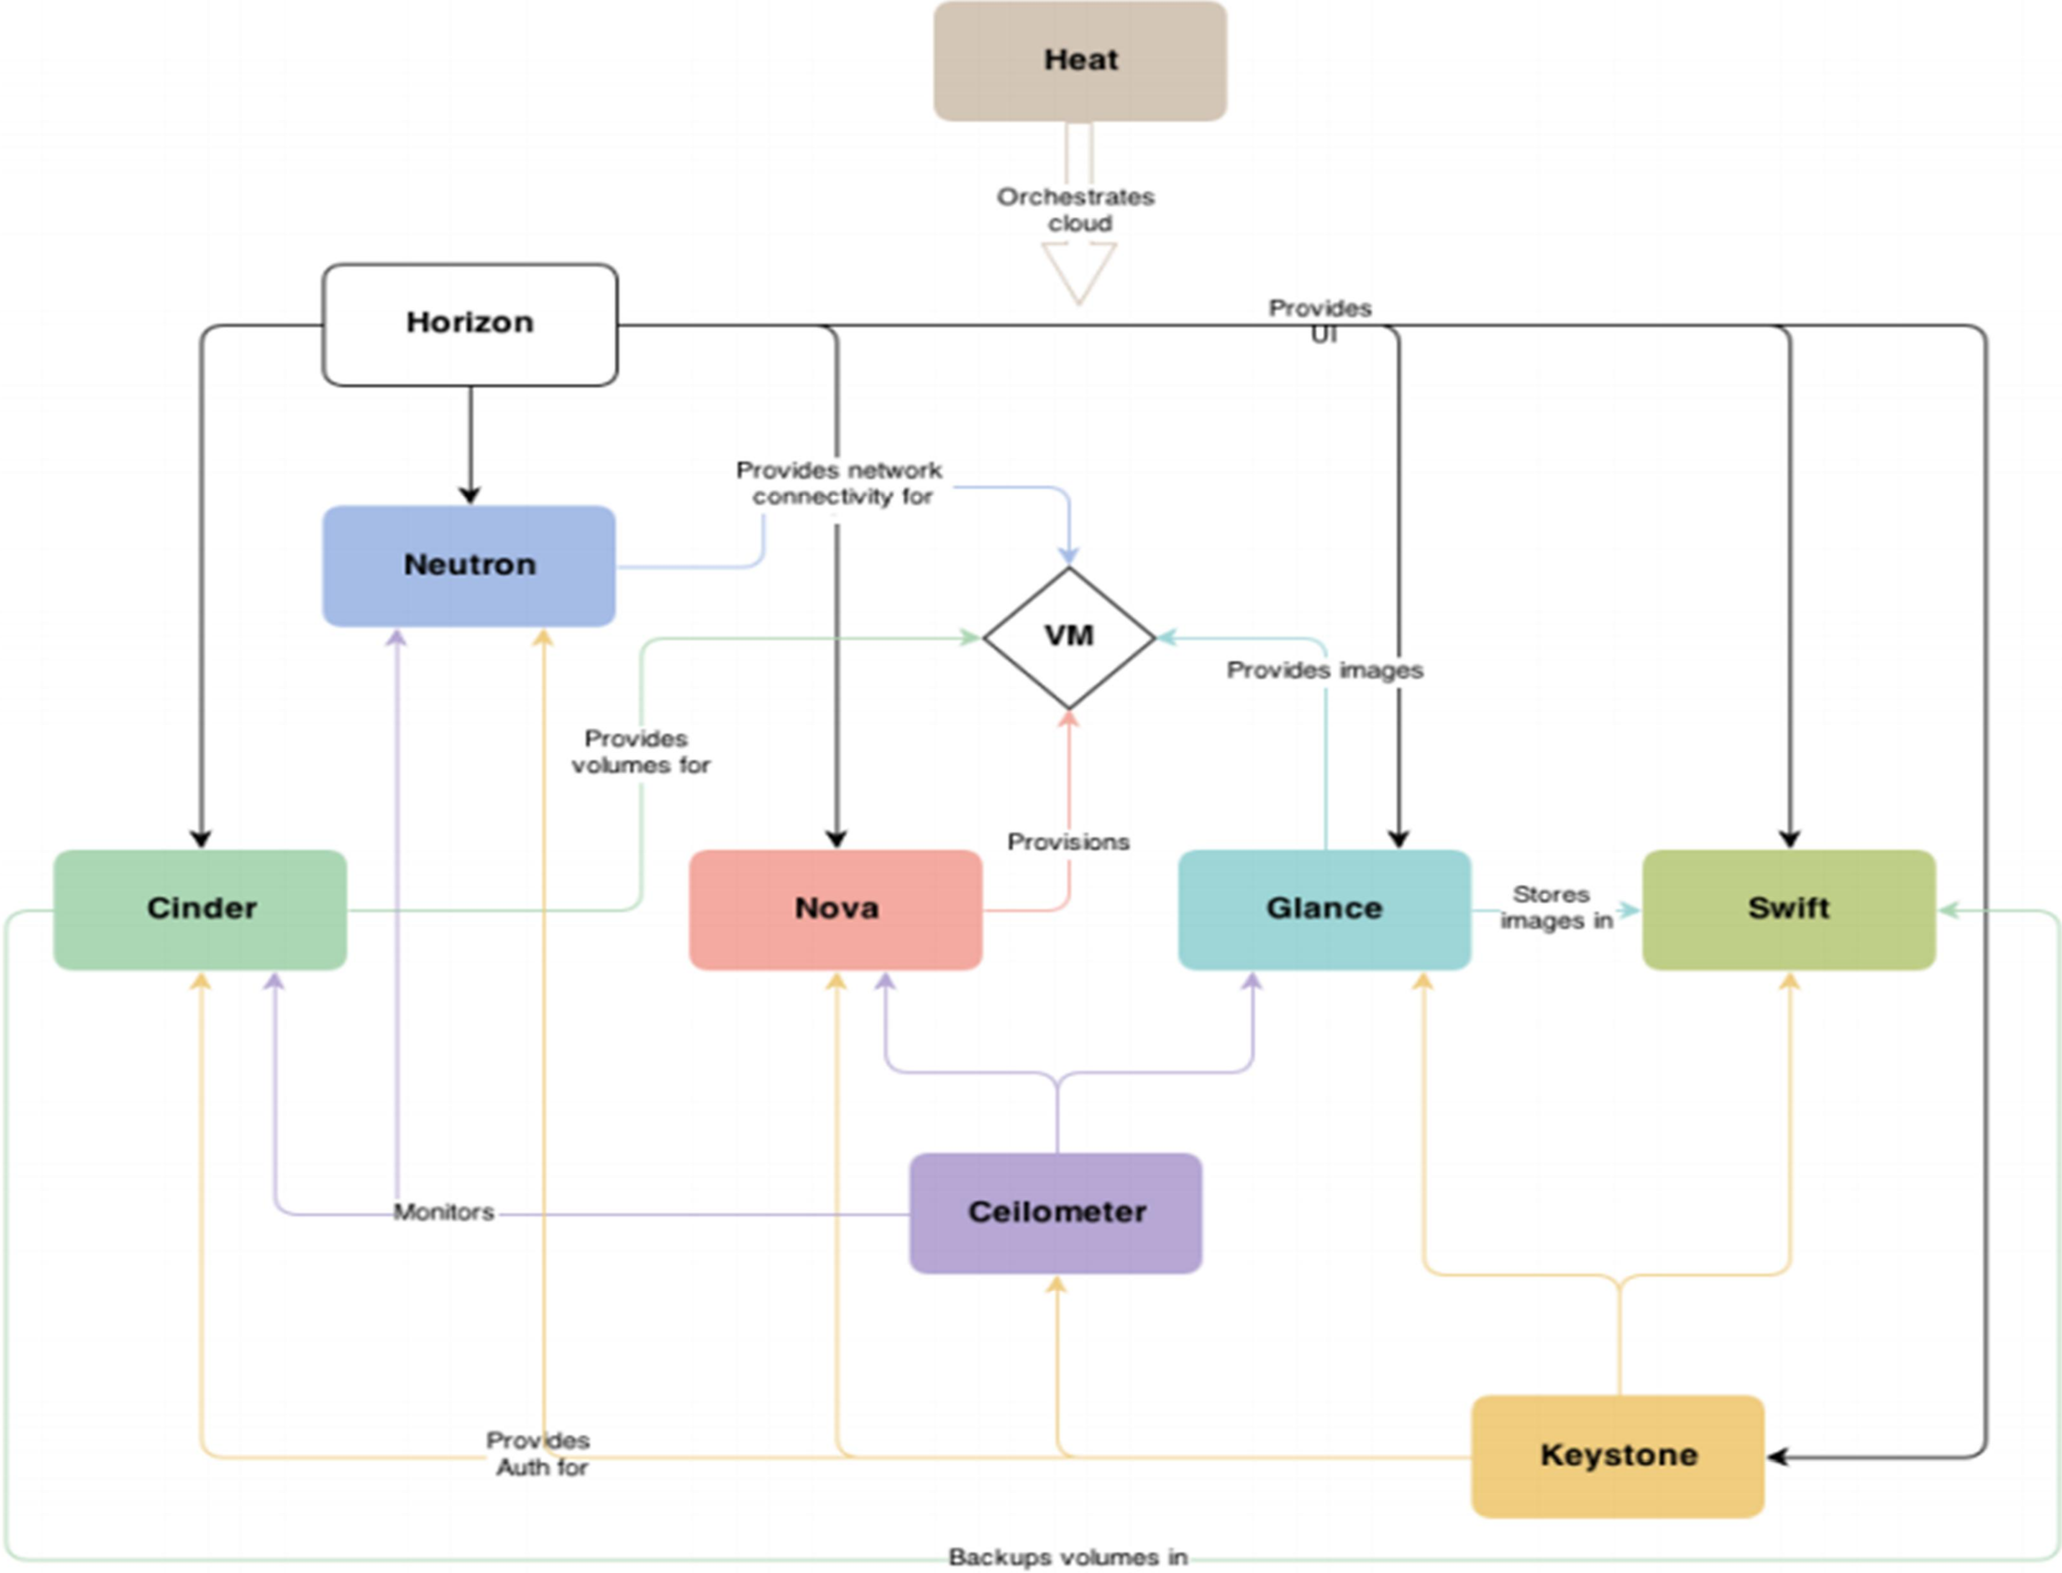
Heat
Orchestrates cloud
Horizon
Provides UI
Neutron
Provides network connectivity for
VM
Provides images
Provides volumes for
Provisions
Cinder
Nova
Glance
Stores images in
Swift
Monitors
Ceilometer
Provides Auth for
Keystone
Backups volumes in

Openstack

- Releases

Series	Status	Date
Juno	Under development	October 2014
Icehouse	Current stable release	Apr 17, 2014
Havana	Security-supported	Apr 03, 2014
Grizzly	EOL	Mar 20, 2014
Folsom	EOL	Apr 11, 2013
Essex	EOL	Oct 12, 2012
Diablo	EOL	Jan 19, 2012
Cactus	Deprecated	Apr 15, 2011
Bexar	Deprecated	Feb 3, 2011
Austin	Deprecated	Oct 21, 2010

OpenStack Distro

- ⇨ www.openstack.org
- ⇨Download, compile, install (module by module) – not that easy

- ⇨ All in One Distros
- ⇨RDO, Fuel, Trystack, DevStack, Rackspace etc

OpenStack Management

⇨ GUI

⇨

⇨ CLI

```
[root@openstack ~(keystone_admin)]# nova image-list
+--------------------------------------+------------+--------+--------+
| ID                                   | Name       | Status | Server |
+--------------------------------------+------------+--------+--------+
| 9eb7bc10-c34b-4c1f-b901-5334b7f3e3f0 | Fedora19   | ACTIVE |        |
| d8cc16d1-0eba-4dfb-9ea2-03ceabadace2 | Vyatta5400 | ACTIVE |        |
| f39f35ec-32c1-4b36-a8d3-b676351a2122 | Vyatta5600 | ACTIVE |        |
| db726128-0d7f-41e2-89f3-783e9906173f | cirros     | ACTIVE |        |
+--------------------------------------+------------+--------+--------+
[root@openstack ~(keystone_admin)]#
```

Image (VM) Management

- ⇨All VM instances has to be associated to a pre-stored image
- ⇨ CLI

```
[root@openstack ~(keystone_admin)]# nova image-list
+--------------------------------------+-------------+--------+--------+
| ID                                   | Name        | Status | Server |
+--------------------------------------+-------------+--------+--------+
| 9eb7bc10-c34b-4c1f-b901-5334b7f3e3f0 | Fedora19    | ACTIVE |        |
| d8cc16d1-0eba-4dfb-9ea2-03ceabadace2 | Vyatta5400  | ACTIVE |        |
| f39f35ec-32c1-4b36-a8d3-b676351a2122 | Vyatta5600  | ACTIVE |        |
| db726128-0d7f-41e2-89f3-783e9906173f | cirros      | ACTIVE |        |
+--------------------------------------+-------------+--------+--------+
[root@openstack ~(keystone_admin)]#
```

VM Instantiation using Openstack

- Via GUI [Vídeo: Instanciando uma máquina simples]

-

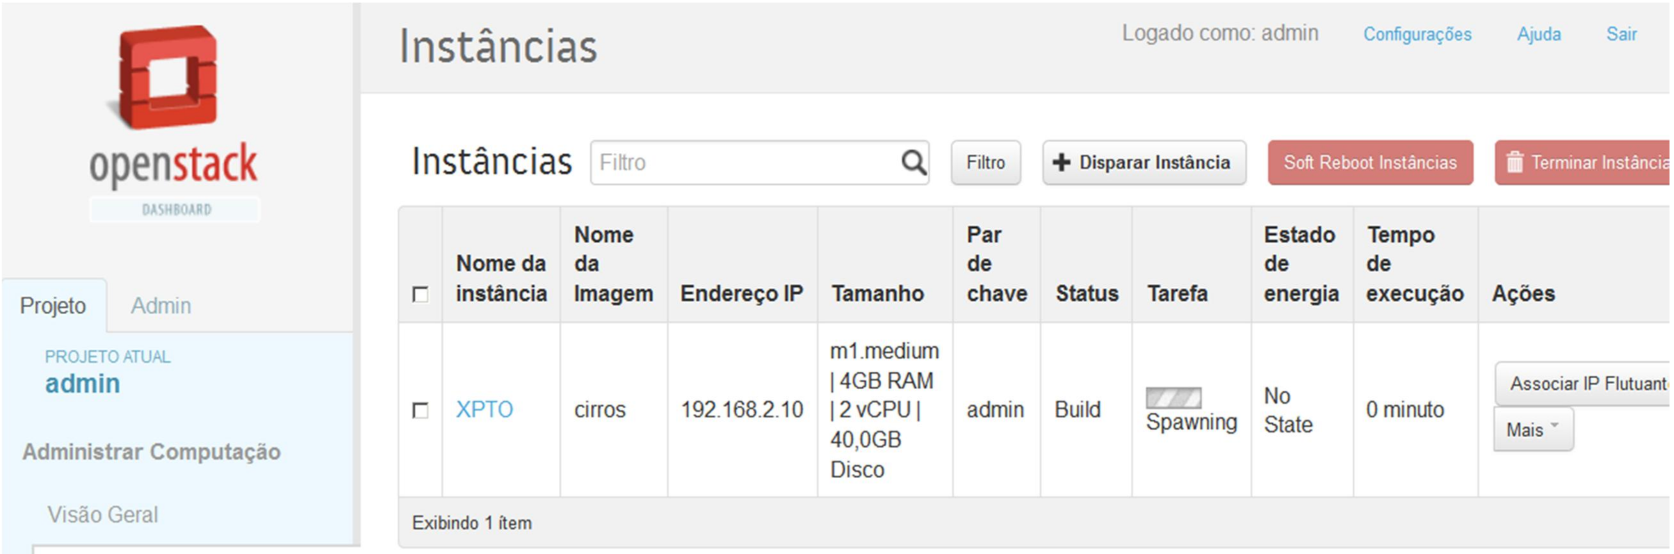

Brocade Vyatta

USE CASE & DEMO

Demo Topology

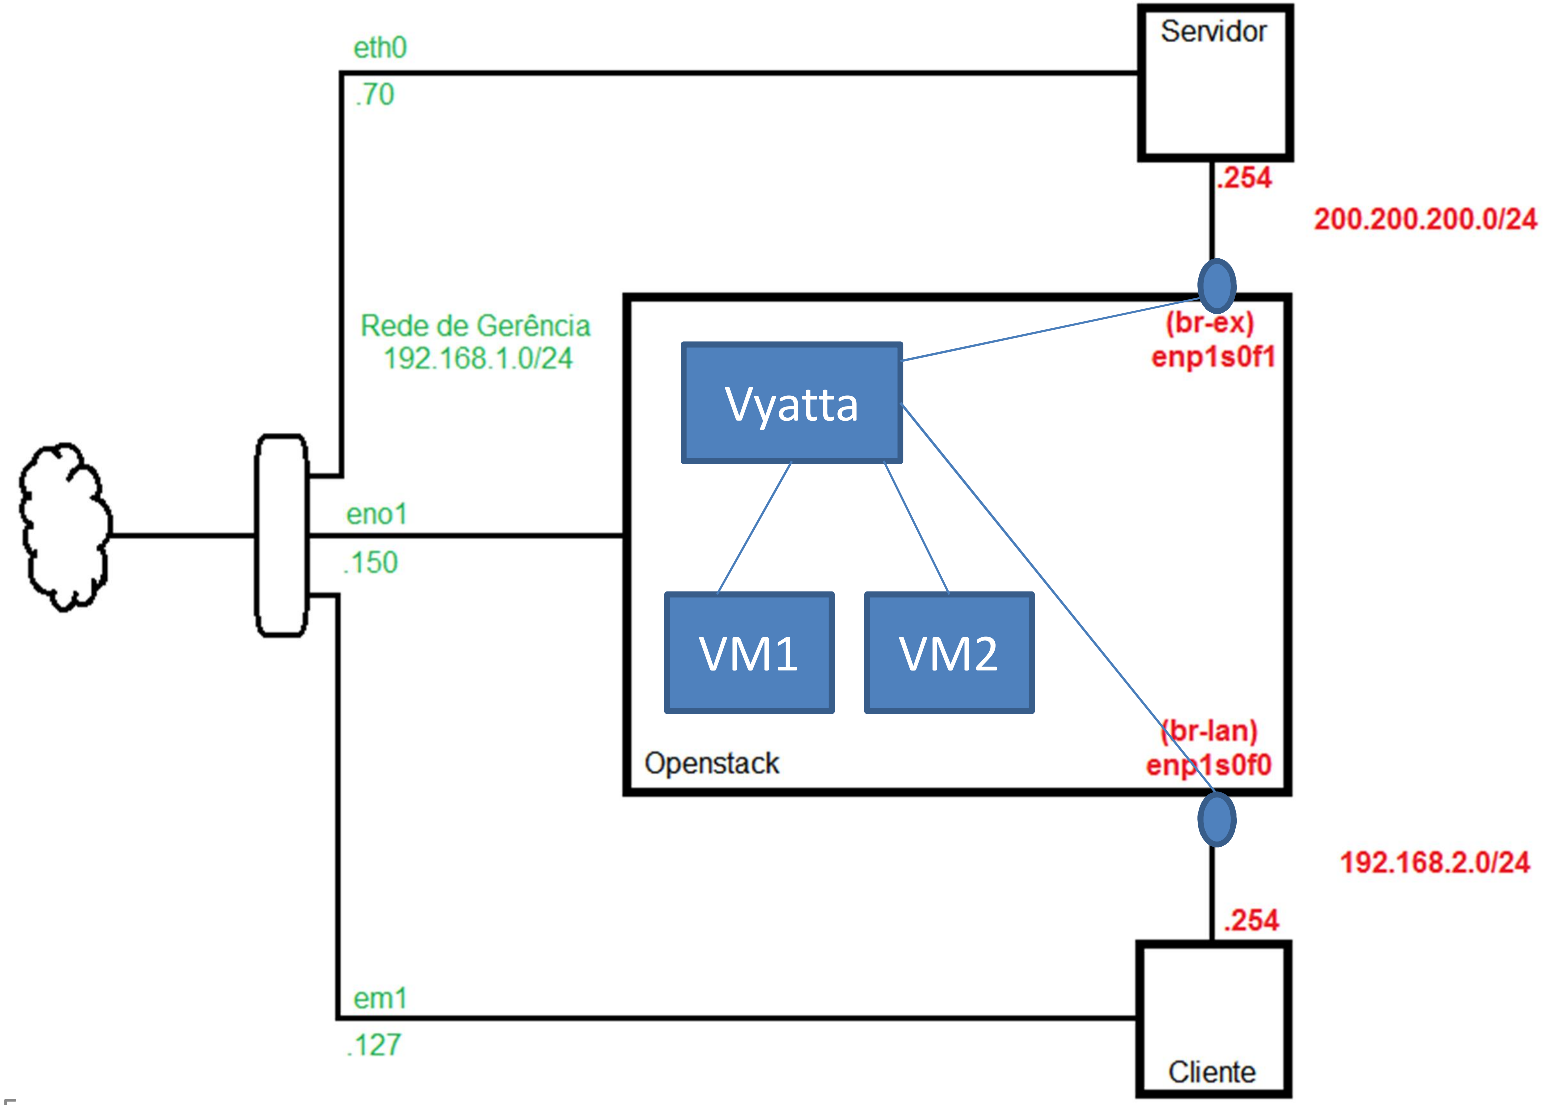

Video Playlist

- Installing OpenStack
- Accessing CLI and GUI
- Setup OVS Bridges
- Download Openstack Images
- Vyatta 5600 Instance
- Iperf Experiments between virtual machines
- Iperf Experiments between physical machines

Ex: Telco Use Case: Virtual BGP Route Reflector

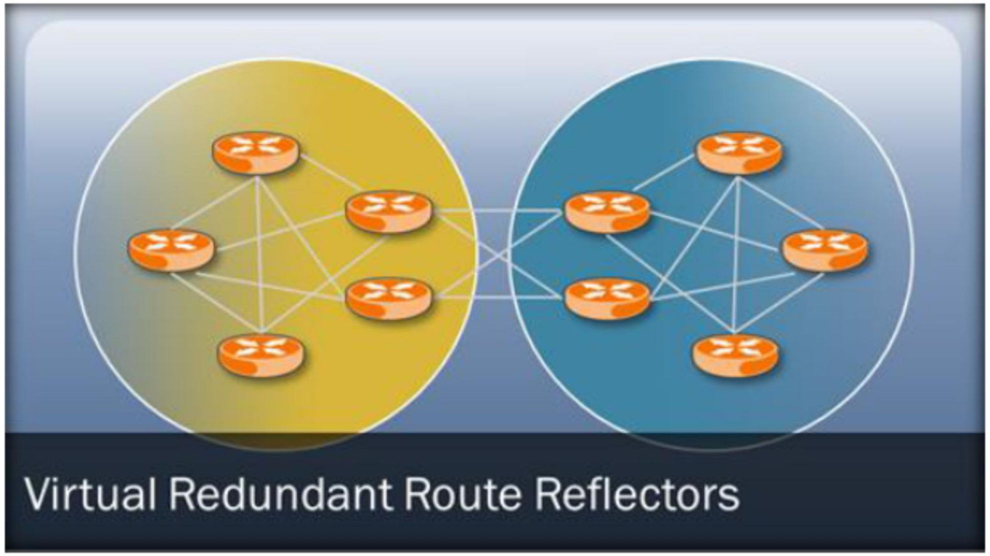

Infrastructure	Routes
Incumbent Hardware	546,488
5400	2,995,770
5600	13,695,151

- BGP Route Reflectors (RR) do not participate in BGP data plane
- BGP RRs today deployed by H/W routers capable of handling data forwarding
- High cost associated in H/W with larger scale – limited memory in custom H/W
- Reduce H/W costs with better scale using S/W based routers

Source: Kelly Herrel (Brocade)

Future Work: Path Ahead

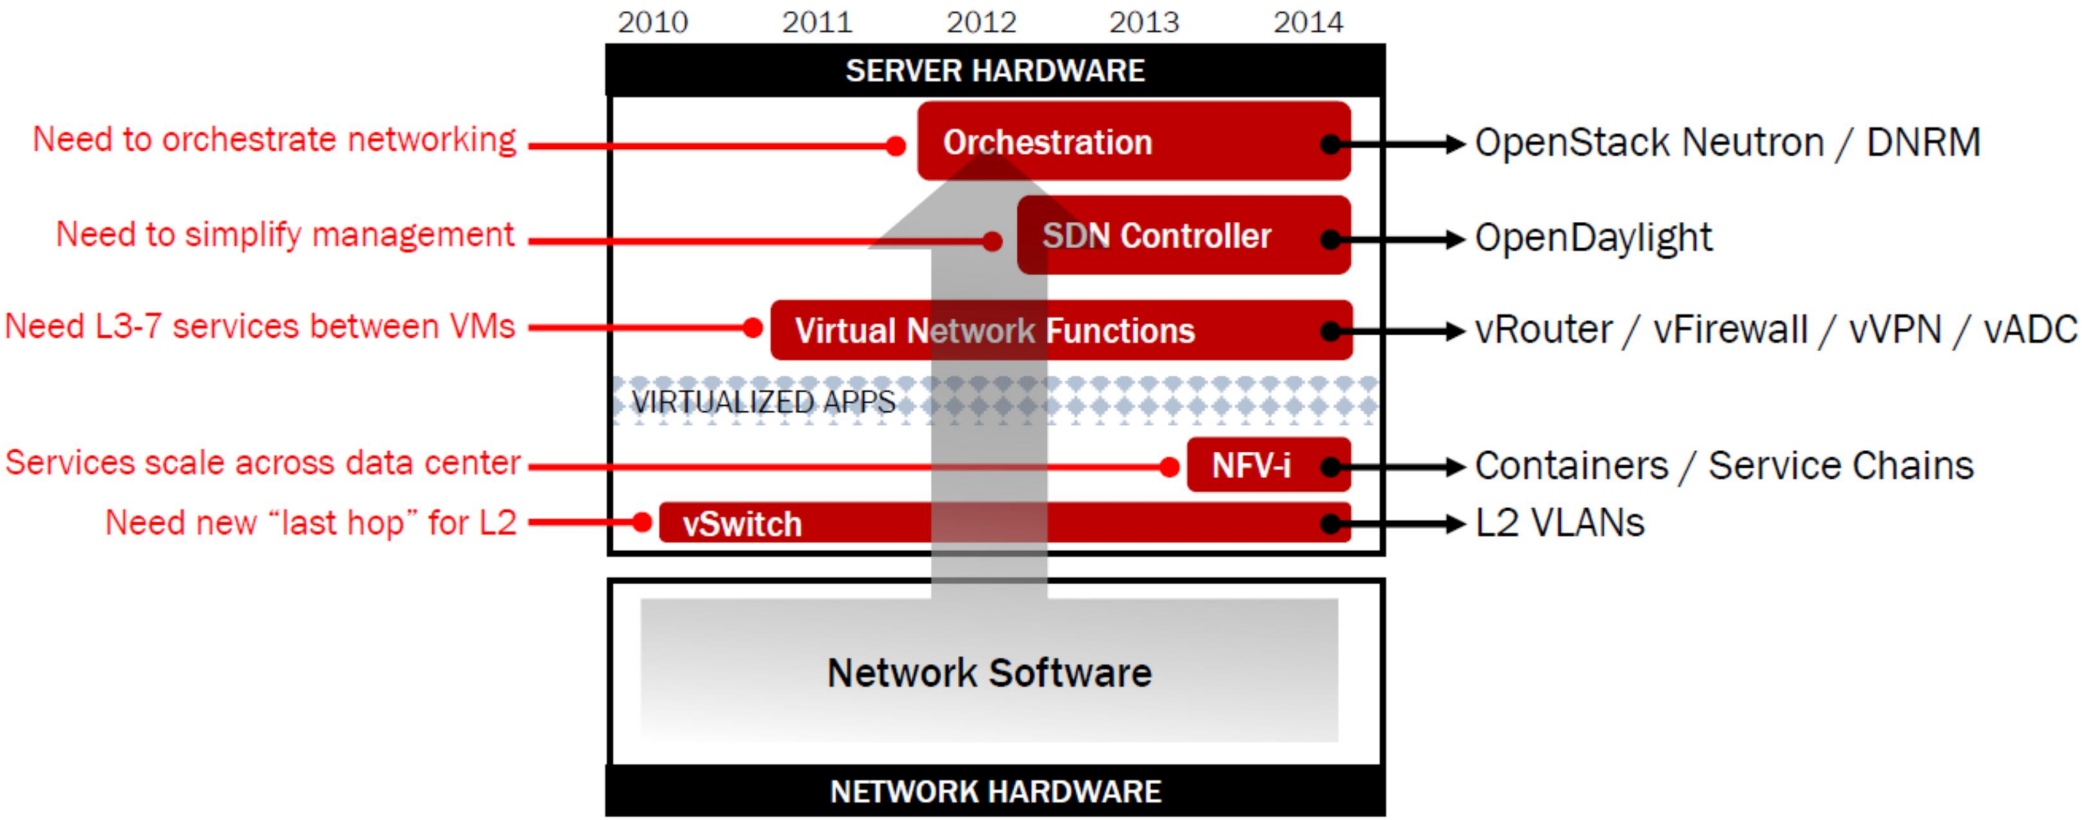

Networking's Path Into The Server

Source: Kelly Herrel (Brocade)

Conclusions

1. NFV aims to reduce OpEx by automation and scalability provided by implementing network functions as virtual appliances
2. NFV allows all benefits of virtualization and cloud computing including orchestration, scaling, automation, hardware independence, pay-per-use, fault-tolerance, …
3. NFV and SDN are independent and complementary. You can do either or both.
4. NFV requires standardization of reference points and interfaces to be able to mix and match VNFs from different sources
5. NFV can be done now. Several of virtual functions have already been demonstrated by carriers.

References / **Acknowledgements**

- **ETSI NFV ISG, http://portal.etsi.org/portal/server.pt/community/NFV/367**
- **Diego R. Lopez, Telefónica I+D, NFV ISG Technical Manager, Network Functions Virtualization - Beyond Carrier-grade Clouds**
- **Raj Jain, Introduction to Network Function Virtualization (NFV), http://www.cse.wustl.edu/~jain/cse570-13/m_17nfv.htm**
- M. Cohn, “NFV Insider’s Perspective, Part 2: There’s a Network in NFV –The Business Case for SDN,” Sep 2013, http://www.sdncentral.com/education/nfv-insiders-perspective-part-2-theres-network-nfv-business-case-sdn/2013/09/
- M. Cohn, “NFV Group Flocks to Proof-of-Concept Demos,” Aug 2013, http://www.sdncentral.com/technology/nfv-group-flocks-to-proof-ofconcept-models/2013/08/
- W. Xu, et al., “Data Models for NFV,” IETF Draft, Sep 2013, http://tools.ietf.org/html/draft-xjz-nfv-model-datamodel-00
- CloudNFV, http://www.cloudnfv.com/page1.html
- Project Clearwater, http://www.projectclearwater.org/
- B. Briscoe, et al., “NFV,” IETF, March 2012, http://www.ietf.org/proceedings/86/slides/slides-86-sdnrg-1.pdf
- Intel, “Open simplified Networking Based on SDN and NFV,” 2013, 7 pp., http://www.intel.com/content/dam/www/public/us/en/documents/whitepapers/sdn-part-1-secured.pdf
- J. DiGiglio, and D. Ricci, “High Performance, Open Standard Virtualization with NFV and SDN,” http://www.windriver.com/whitepapers/ovp/ovp_whitepaper.pdf

Acronyms

- API Application Programming Interface
- BRAS Broadband Remote Access Server
- BSS Business Support Systems
- CapEx Capital Expenditure
- CDN Content Distribution Network
- CGNAT Carrier-Grade Network Address Translator
- CGSN Combined GPRS Support Node
- COTS Commercial-off-the-shelf
- DDIO Data Direct I/O Technology
- DHCP Dynamic Host control Protocol
- DPI Deep Packet Inspection
- EMS Element Management System
- ETSI European Telecom Standards Institute
- GGSN Gateway GPRS Support Node
- GPRS
- HLR Home Location Register
- IaaS Infrastructure as a Service

Acronyms

- IETF Internet Engineering Task Force
- IMS IP Multimedia System
- INF Architecture for the virtualization Infrastructure
- IP Internet Protocol
- ISG Industry Specification Group
- LSP Label Switched Path
- MANO Management and orchestration
- MME Mobility Management Entity
- NAT Network Address Translation
- NF Network Function
- NFV Network Function Virtualization
- NFVI Network Function Virtualization Infrastructure
- NFVIaaS NFVI as a Service
- NIC Network Interface Card
- OpEx Operational Expences
- OS Operating System

Acronyms

- OSS Operation Support System
- PaaS Platform as a Service
- PE Provider Edge
- PGW Packet Data Network Gateway
- PoC Proof-of-Concept
- PoP Point of Presence
- PSTN Public Switched Telephone Network
- QoS Quality of Service
- REL Reliability, Availability, resilience and fault tolerance group
- RGW Residential Gateway
- RNC Radio Network Controller
- SaaS Software as a Service
- SBC Session Border Controller
- SDN Software Defined Networking
- SGSN Serving GPRS Support Node
- SGW Serving Gateway

Acronyms

- SIP Session Initiation Protocol
- SLA Service Level Aggrement
- SWA Software architecture
- TAS Telephony Application Server
- TMF Forum
- vEPC
- VM Virtual Machine
- VNF Virtual Network Function
- VNFaaS VNF as a Service
- vSwitch Virtual Switch
- VT-d Virtualization Technology for Direct IO
- VT-x Virtualization Technology

BACKUP

NFV ISG PoC	NFV Use Case	Operators	Vendors
CloudNFV Open NFV Framework	Use Case #5 Virtualization of the Mobile Core and IMS	Sprint Telefonica	6Wind, Dell Enterprise Web Huawei, Mellanox Overture, Qosmos
Service Chaining for NW Function Selection in Carrier Networks	Use Case #2 Virtual Network Function as a Service (VNFaaS) Use Case #4 Virtual Network Forwarding Graphs	NTT	Cisco, HP Juniper
Virtual Function State Migration and Interoperability	Use Case #1 NFV Infrastructure as a Service (NFVIaaS)	AT&T BT	Broadcom Tieto
Multi-vendor Distributed NFV	Use Case #2 VNFaaS Use Case #4 Virtual Network Forwarding Graphs	CenturyLink	Certes Cyan Fortinet RAD
E2E vEPC Orchestration in a multi-vendor open NFVI environment	Use Case #1 NFVIaaS Use Case #5 Virtualization of the Mobile Core and IMS	Sprint Telefonica	Connectem Cyan Dell Intel
Virtualised Mobile Network with Integrated DPI	Use Case #2 VNFaaS Use Case #5 Virtualization of the Mobile Core and IMS Use Case #6 Virtualisation of Mobile base station	Telefonica	HP Intel Qosmos Tieto Wind River
C-RAN virtualisation with dedicated hardware accelerator	Use Case #6 Virtualisation of Mobile base station	China Mobile	Alcatel-Lucent Intel Wind River
Automated Network Orchestration	Use Case #1 NFVIaaS	Deutsche Telekom	Ericsson x-ion
VNF Router Performance with DDoS Functionality	Use Case #2 VNFaaS	AT&T Telefonica	Brocade Intel

VNF Design Patterns and VNFCs

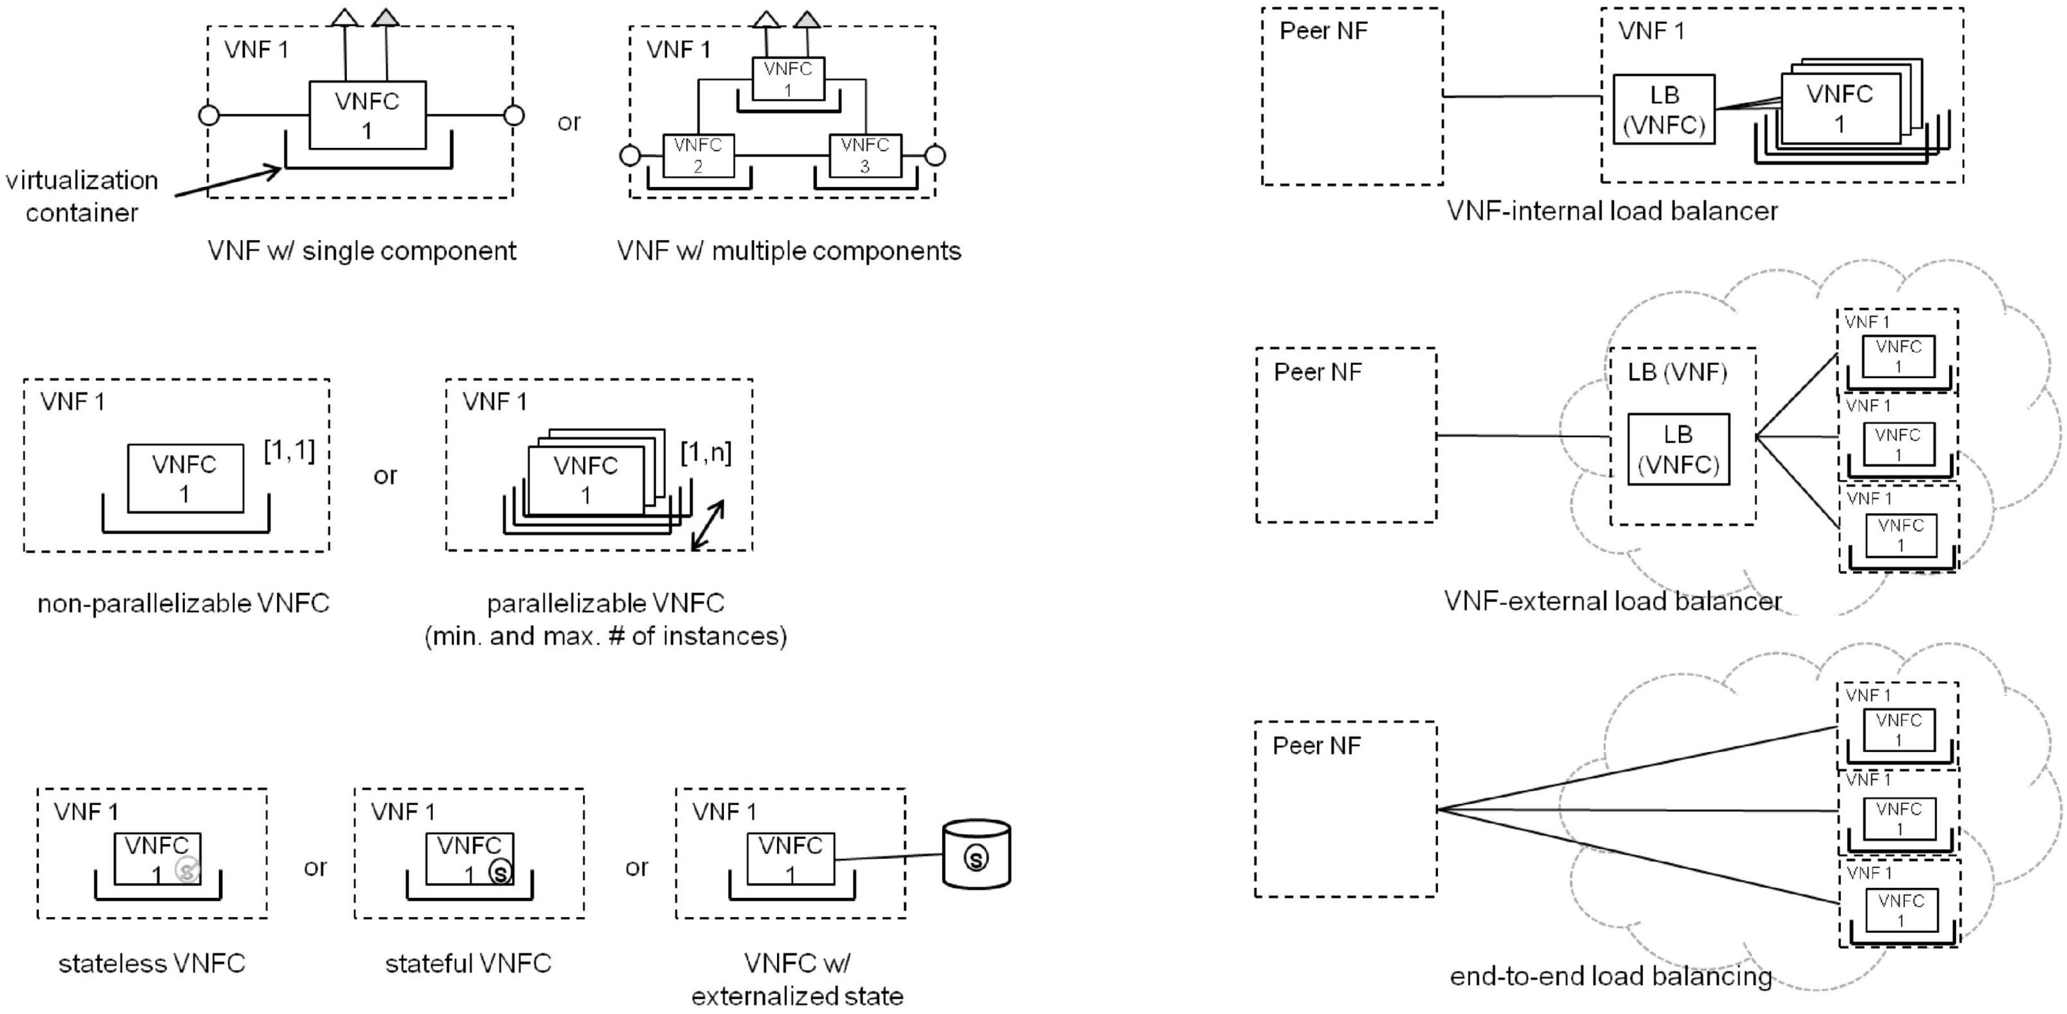

Management and Orchestration Architecture

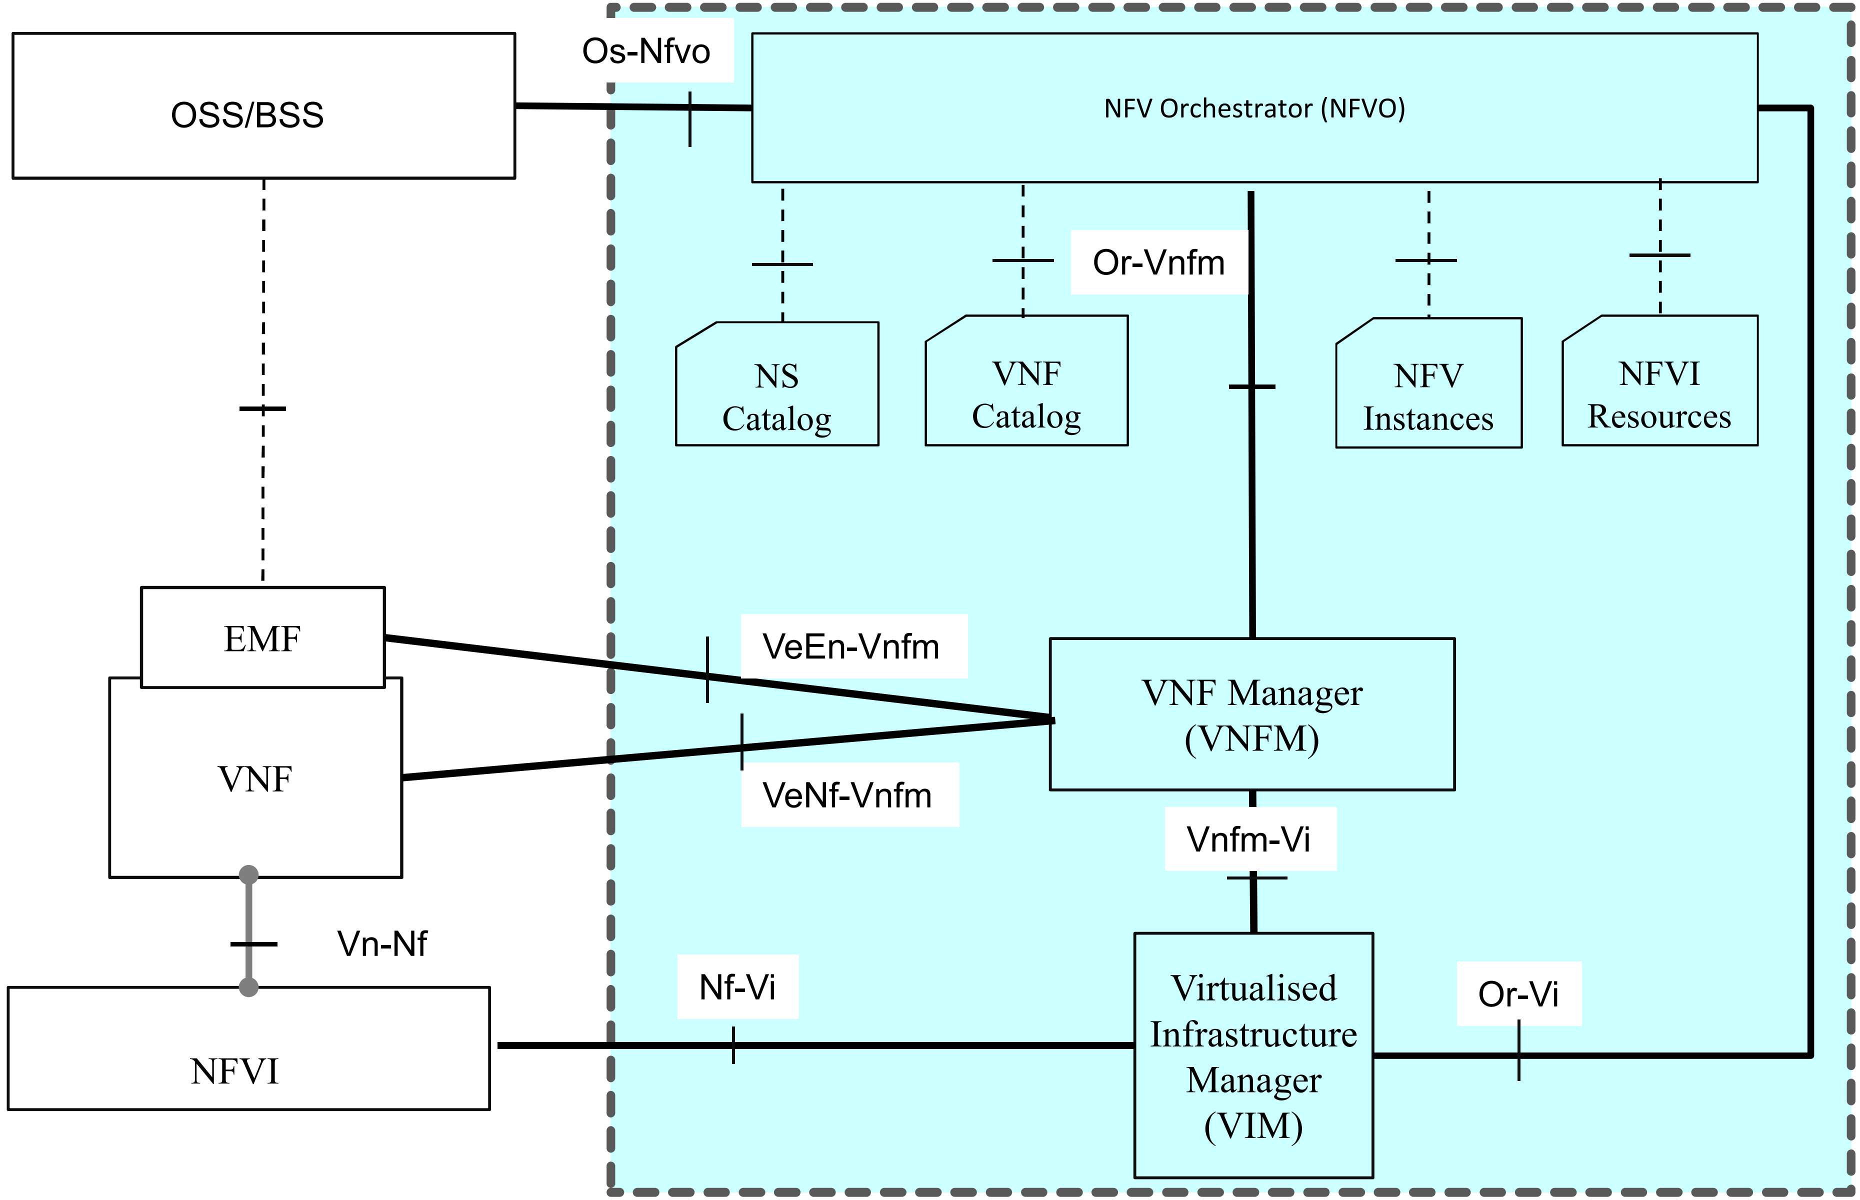

Descriptor Information Model