Survey of ETSI NFV standardization documents

BY
ABHISHEK GUPTA
FRIDAY GROUP MEETING
FEBRUARY 26, 2016

UCDAVIS

VNFaaS (Virtual Network Function as a Service)

- **In our present work, we consider the VNFaaS use-case**

- **Here, the service provider is responsible for deploying, configuring, updating and managing the operation of the VNF instance to provide the expected service level (SLA) for subscribers**

- **Enterprise network functions that can be virtualized : AR(Enterprise Access Router/Enterprise CPE), PE(Provider Edge Router), FW (Enterprise Firewall), NG-FW(Enterprise NG-FW), WOC(Enterprise WAN Optimization Controller), DPI(Deep Packet Inspection - Appliance/Function), IPS(Intrusion Prevention System), Network Performance Monitoring.**

[1] ETSI NFV; Use-Cases

Use cases for NFV

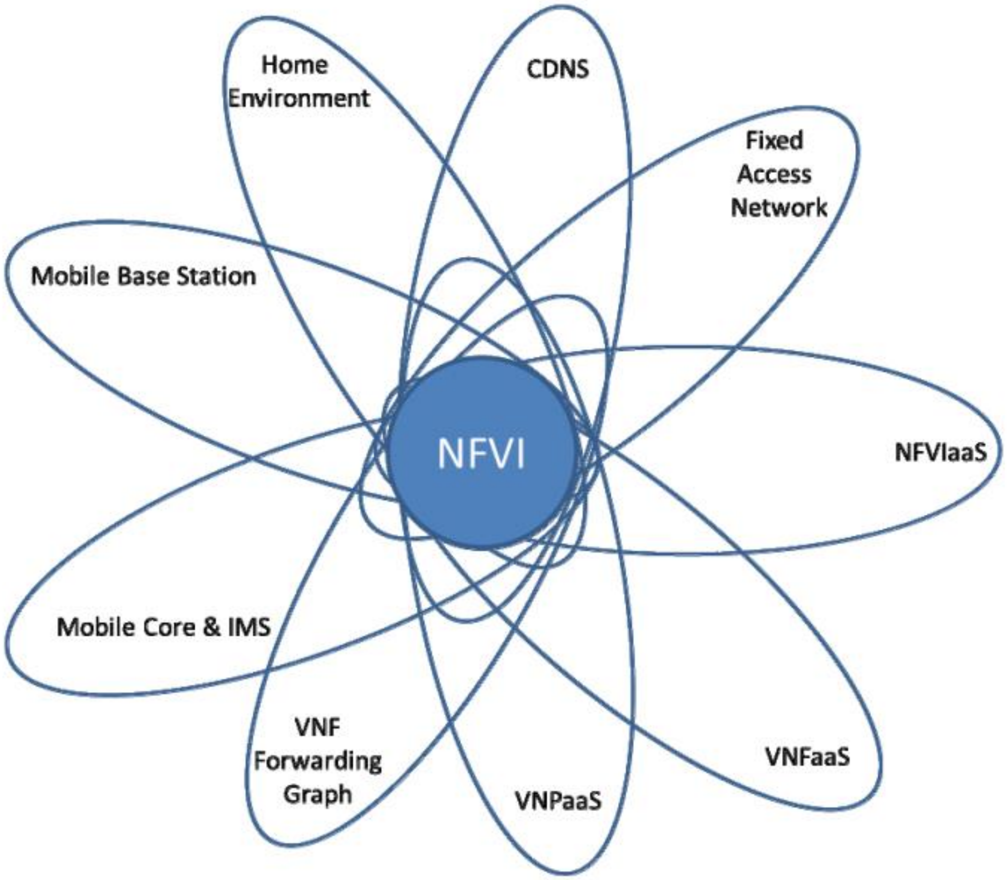

[2] ETSI NFV; Infrastructure Overview

Network-enabled Cloud/ Integrated Cloud (AT&T)

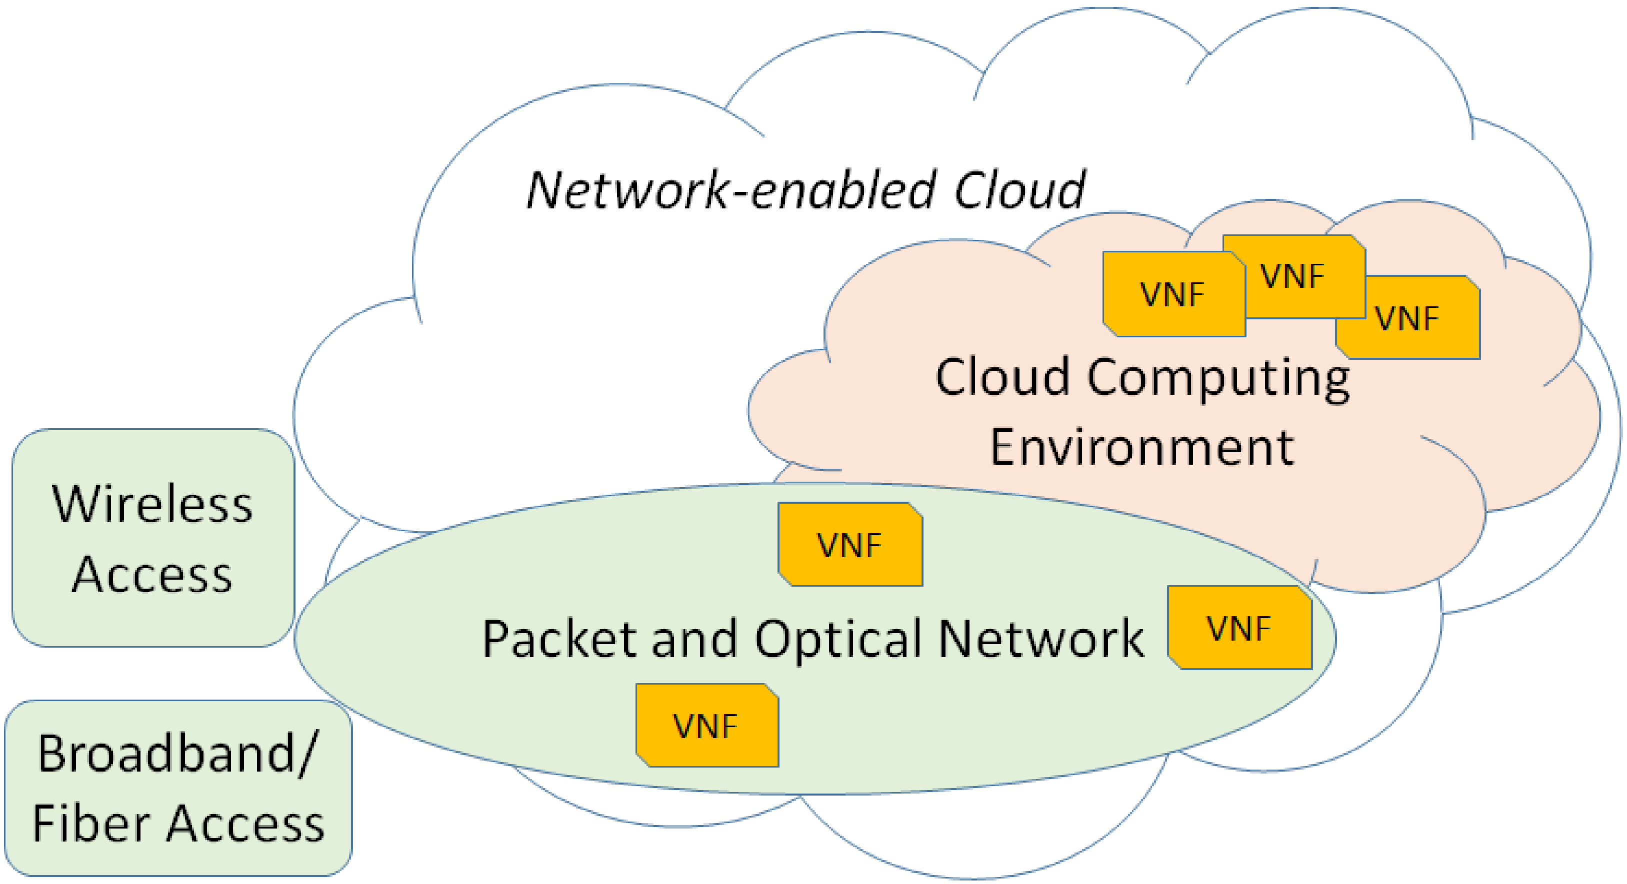

Simulation Topology

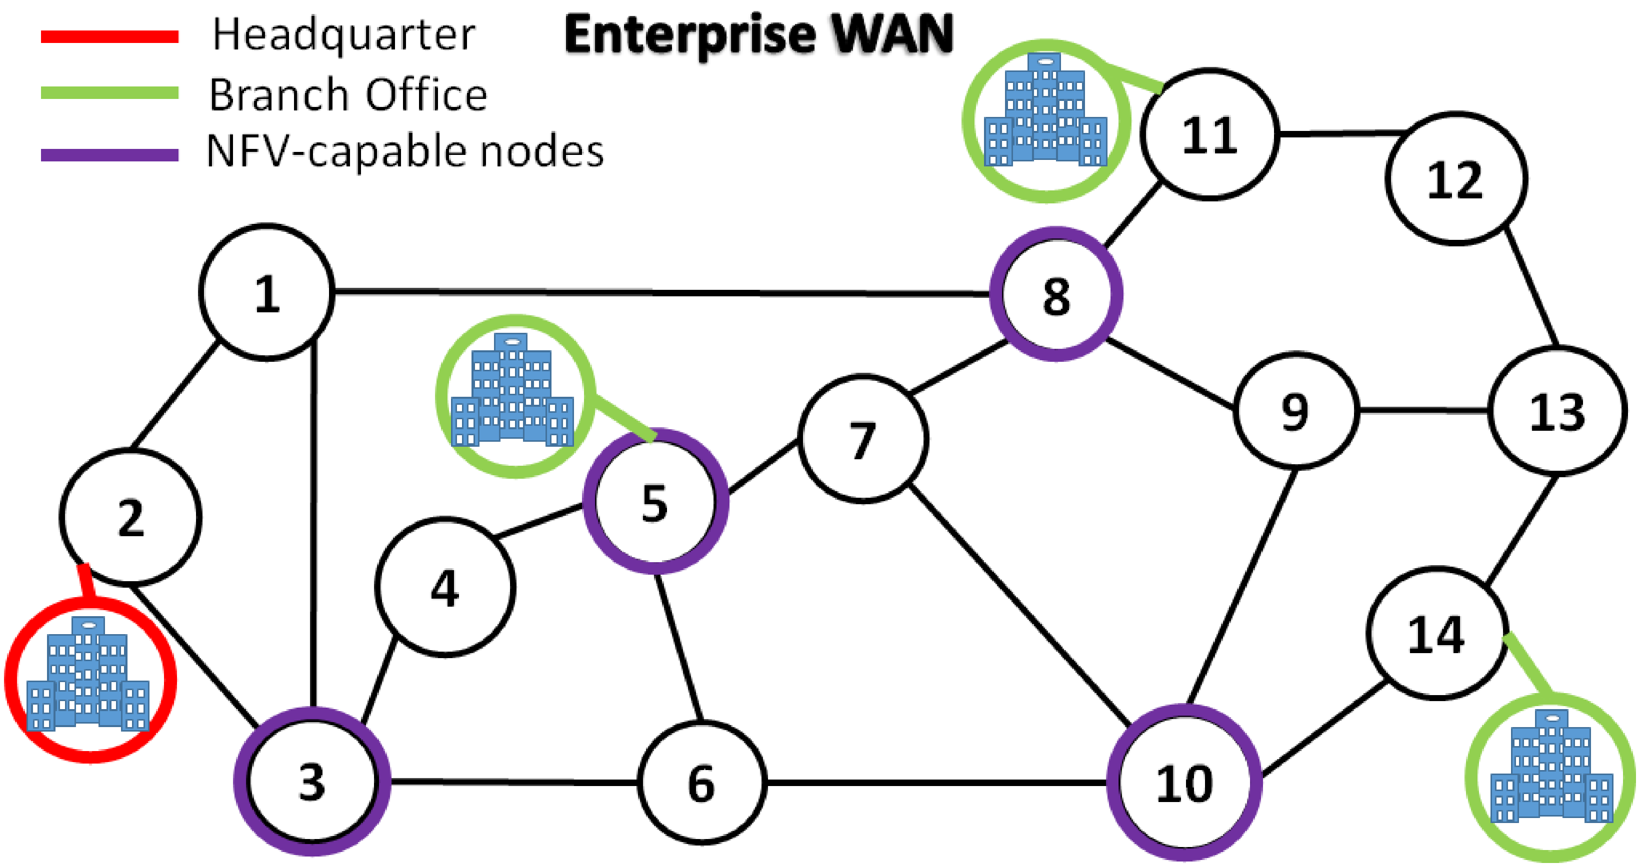

Results

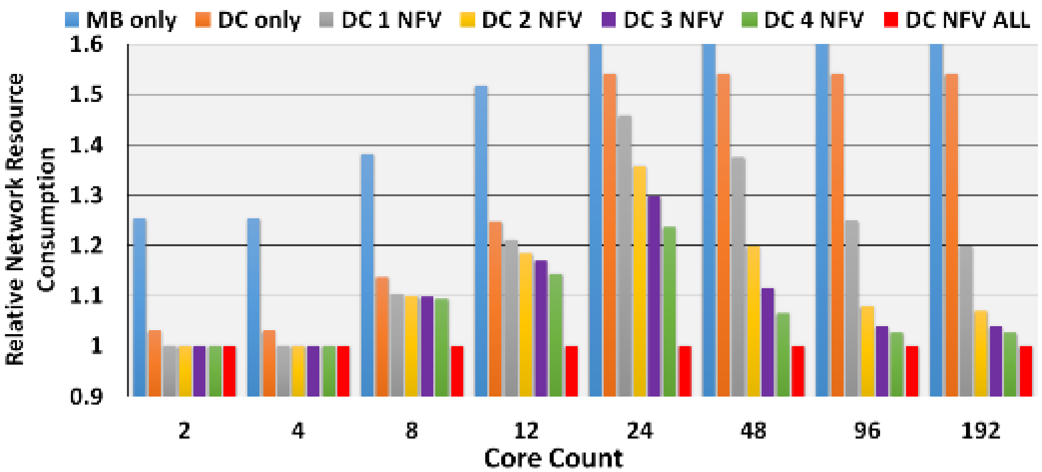

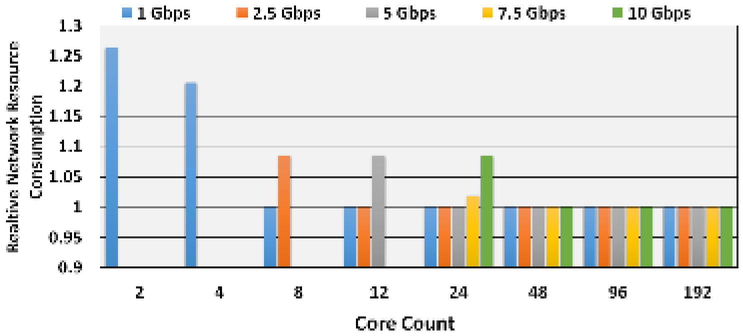

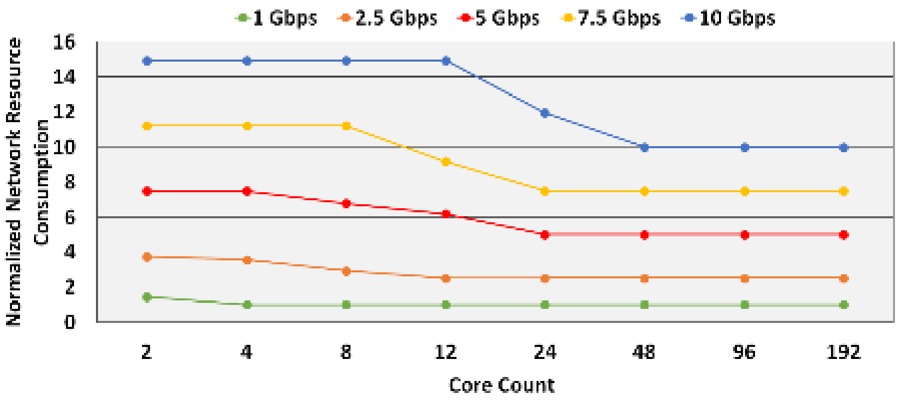

UCDAVIS
UNIVERSITY OF CALIFORNIA

Assumptions

- Only a single type of service was provisioned
  - There may be several services to be provisioned
- Service chain is a single linear chain
  - When multiple services are to be deployed, we will have a VNF-Forwarding Graph (VNF-FG)
  - For services that do not specify the service chain order, we will have a VNF-Set (VNF-S)
- Single datacenter in the scenario
  - Multiple datacenters scenario needs to be investigated
- We consider a clean-slate scenario (no service have been deployed yet)
  - The scenario where services have already been deployed is also to be considered.
- VNF scaling is considered as up/down (parallelizable)
  - VNF scaling can in fact in/out (non-parallelizable)

UCDAVIS
UNIVERSITY OF CALIFORNIA

Continued…

- VNF State is not considered
  - A VNF maybe state-full or stateless
- No overhead of resource allocation
  - Alllocation issues at multiple levels - compute node resources(CPU, I/O, memory, storage, hypervisor), control plane complexity etc.
- No virtual path setup cost
  - Variety of tunneling protocols used.
  - Allocation and de-allocation of these paths must have paths associated with them.
- No SLA & QoS requirements
  - SLA and QoS requirements of the various services to be deployed need to be considered.

UCDAVIS UNIVERSITY OF CALIFORNIA

General domain architecture and associated Interfaces

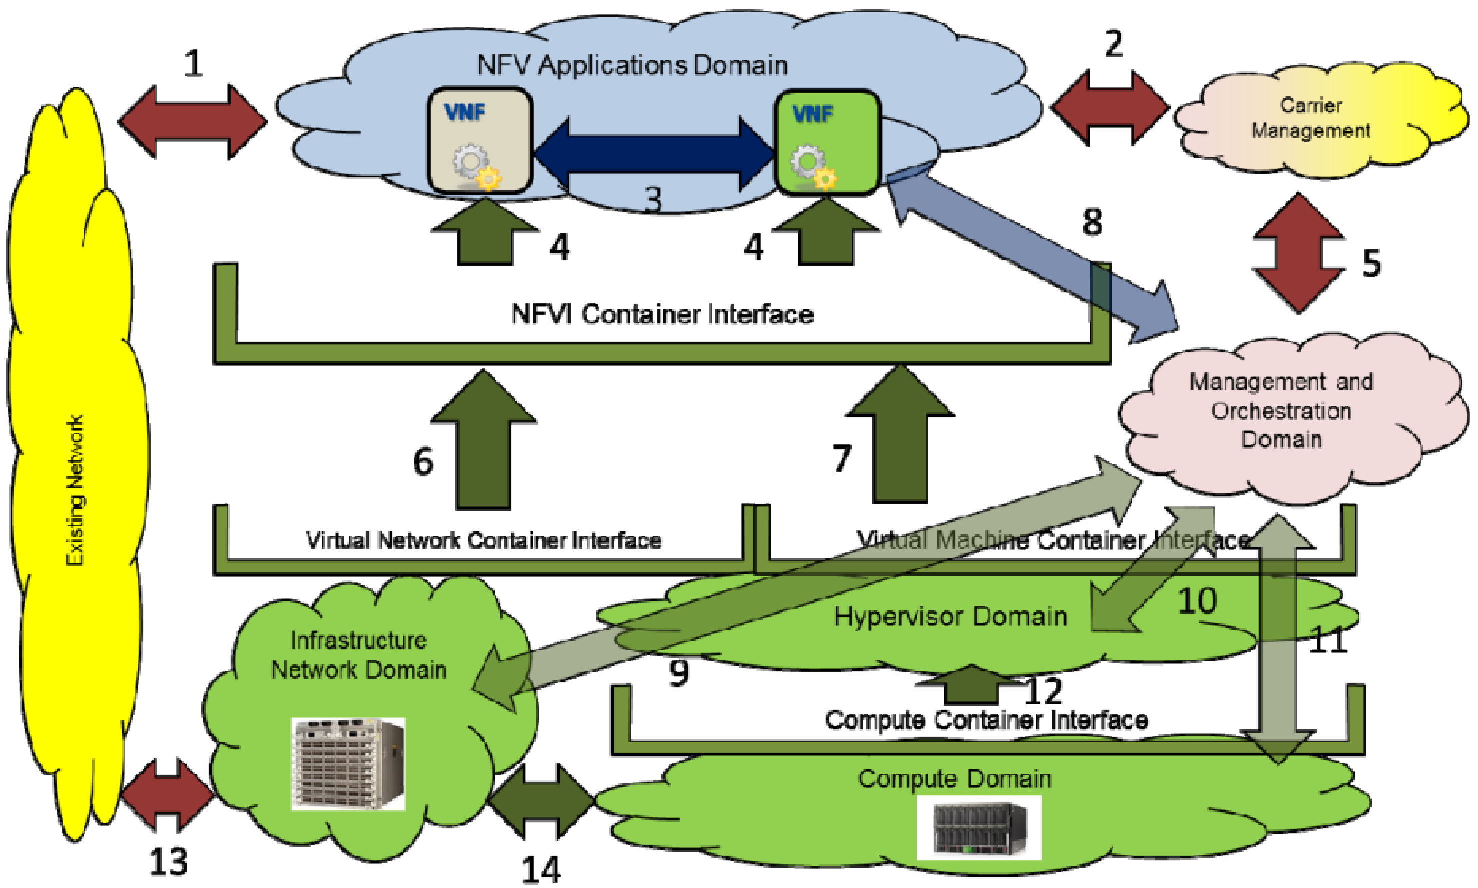

[2] ETSI NFV; Infrastructure Overview

Continued…

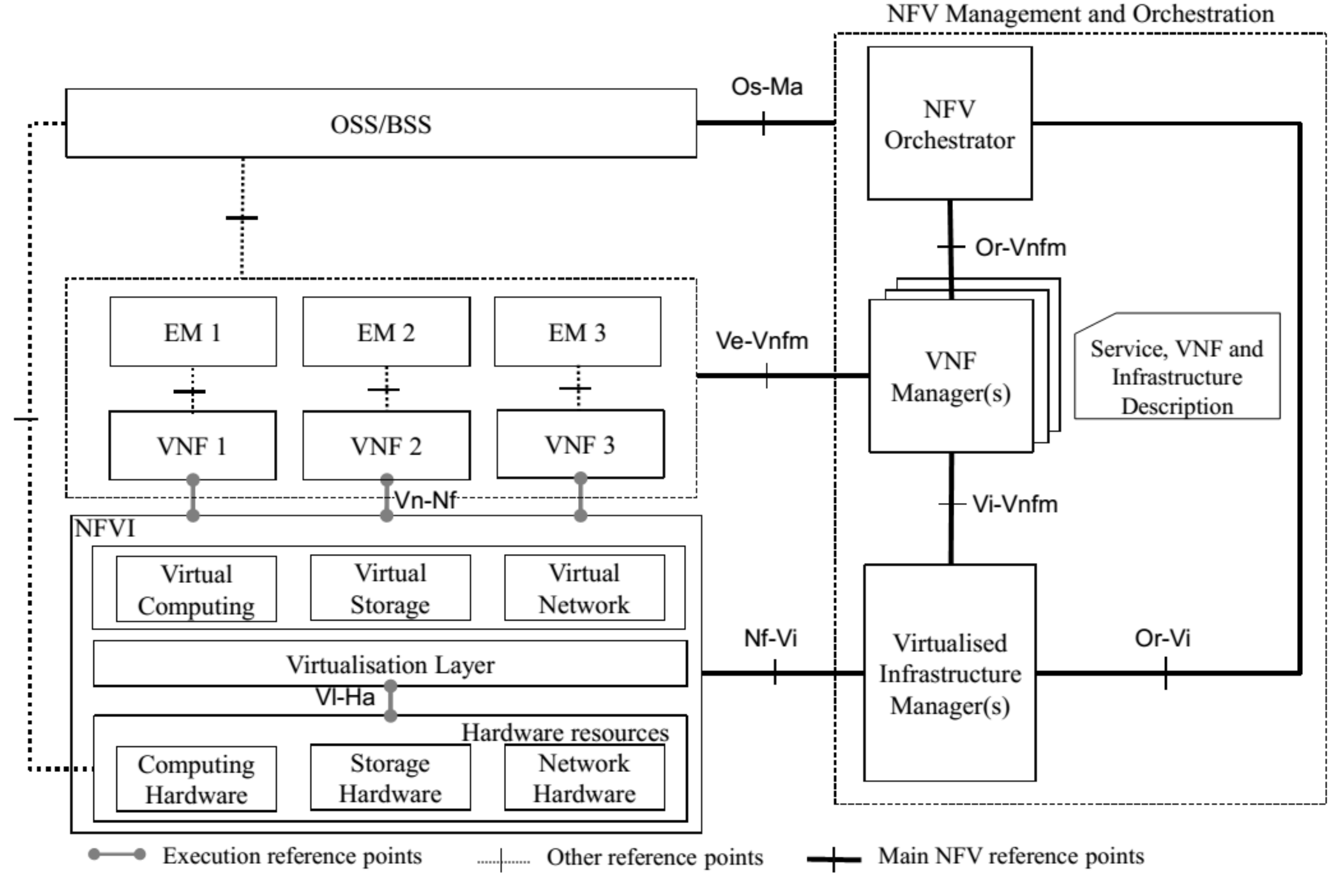

[3] ETSI NFV; Architectural Framework

Various domains of architecture

Hypervisor Domain

Compute Domain

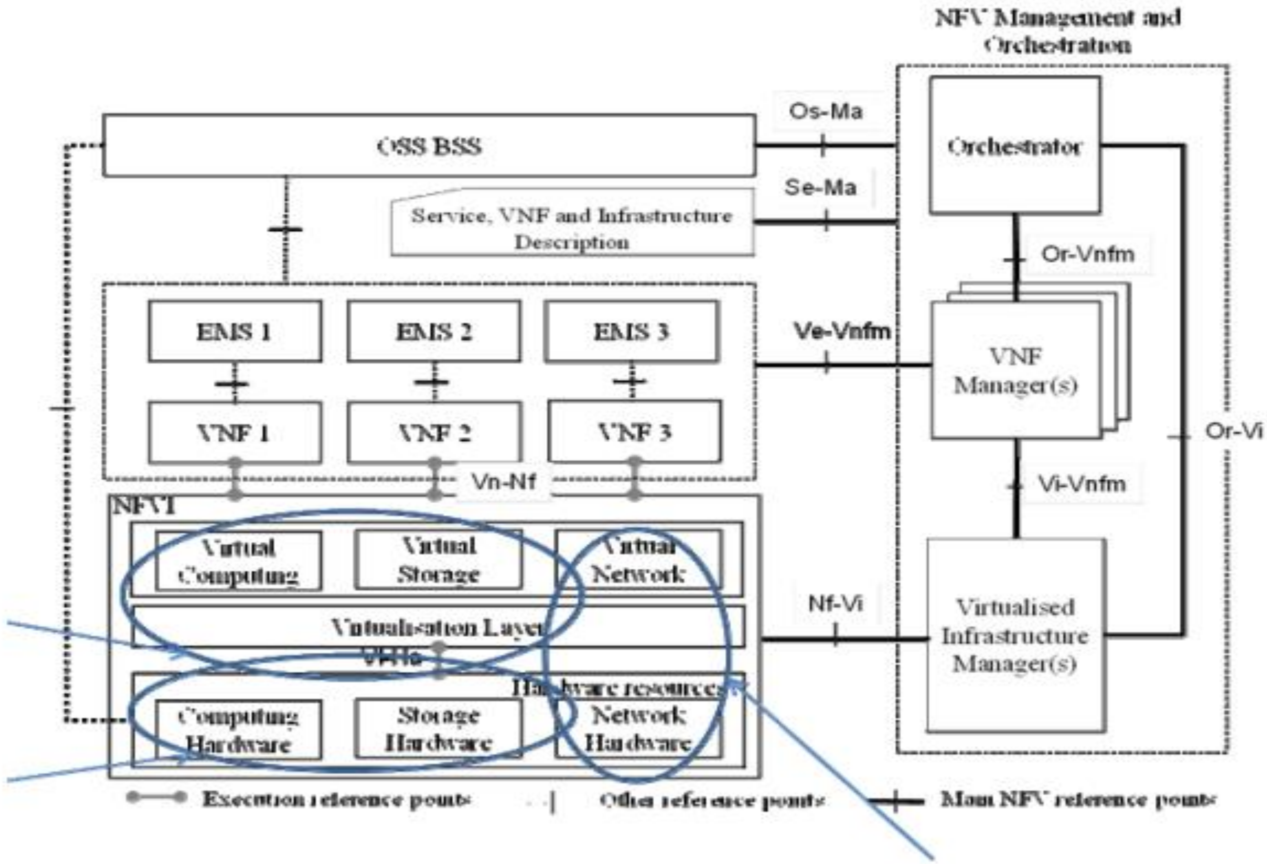

Infrastructure Network Domain

Hypervisor domain

General Architecture of a Cloud Hypervisor

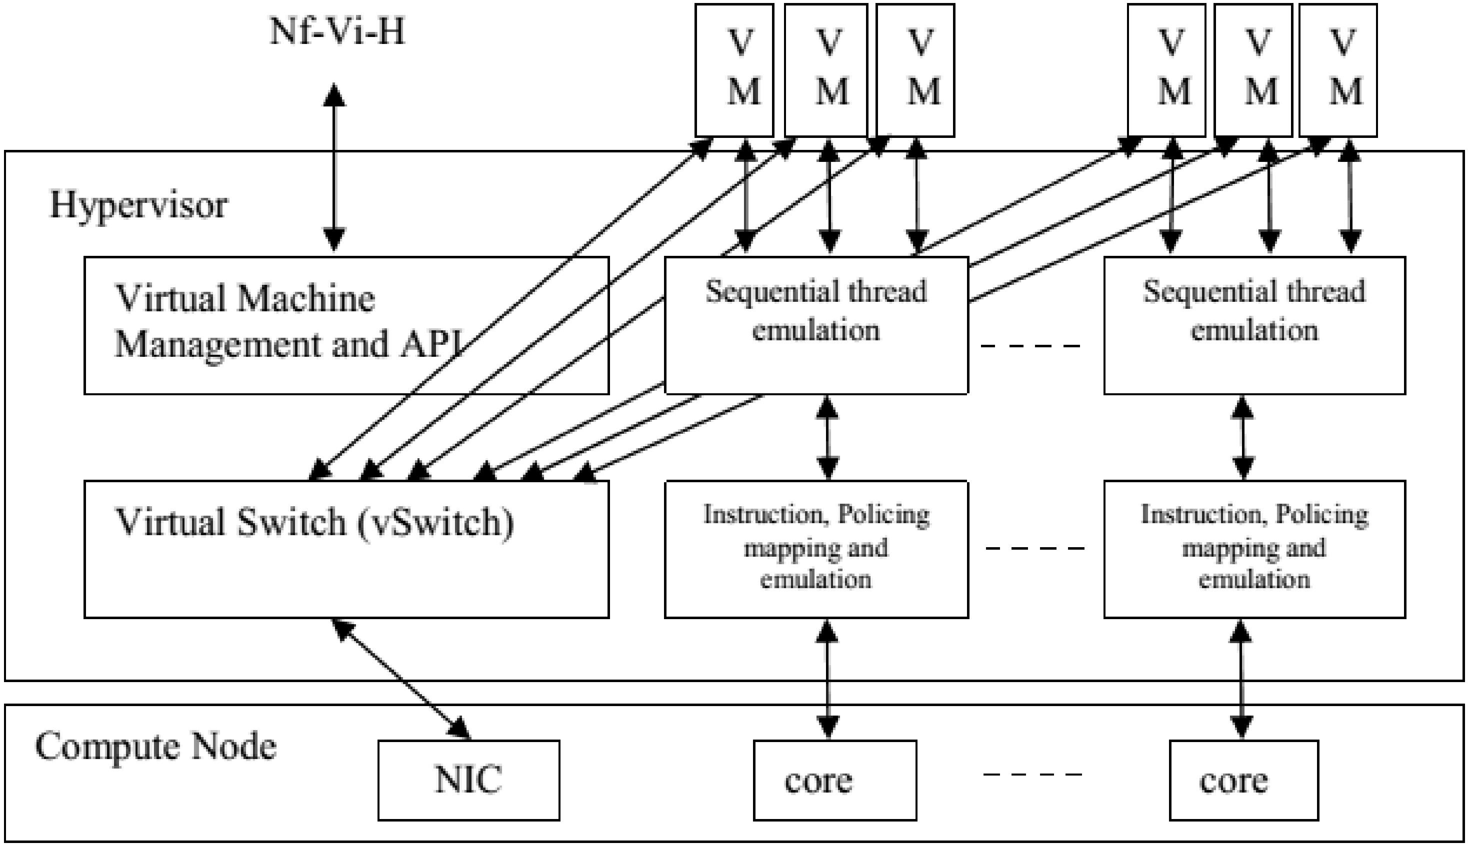

[4] ETSI NFV; Infrastructure; Hypervisor Domain

Hardware support for VM performance improvement

- multicore processors supporting multiple independent parallel threads of execution

- specific CPU enhancements/instructions to control memory allocation and direct access on I/O devices to VM memory allocations

- PCI-e bus enhancements, notably Single Root I/O virtualization (SR-IOV)

[4] ETSI NFV; Infrastructure; Hypervisor Domain

Hypervisor for high performance NFV VMs

- **Exclusive allocation of whole CPU cores to VMs**
- **Direct memory mapped polled drivers for VMs to directly access the physical NICs using user mode instructions requiring no 'context switching‘**
- **Direct memory mapped polled drivers for interVM communications again using user mode instructions requiring no 'context switching‘**
- **vSwitch implementation as a high performance VM again using direct memory mapping and user mode instructions requiring no 'context switching'.**

UC DAVIS
UNIVERSITY OF CALIFORNIA

[4] ETSI NFV; Infrastructure; Hypervisor Domain

NFV Hypervisor Architecture

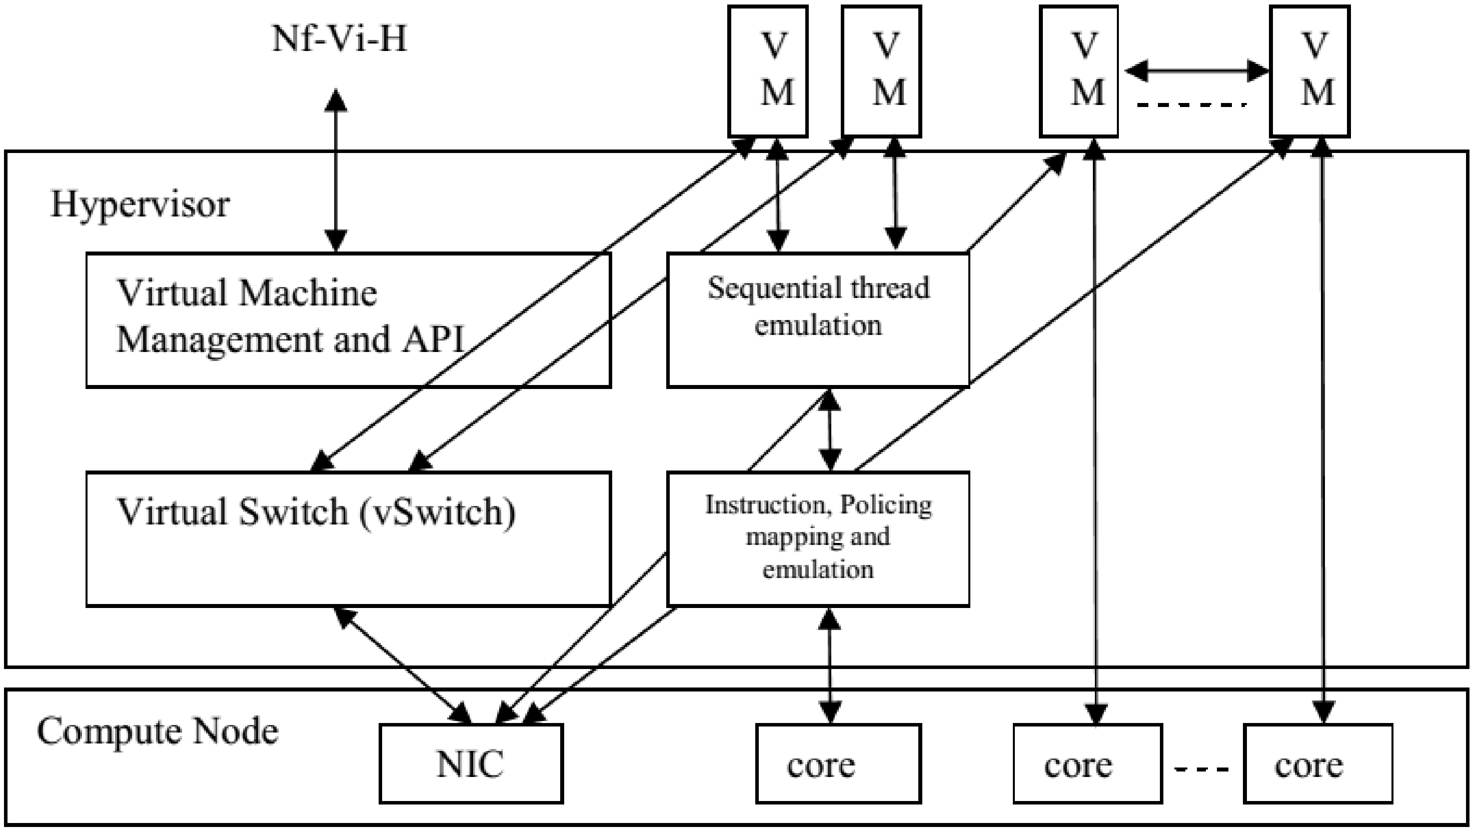

[4] ETSI NFV; Infrastructure; Hypervisor Domain

Network domain

Layering Model

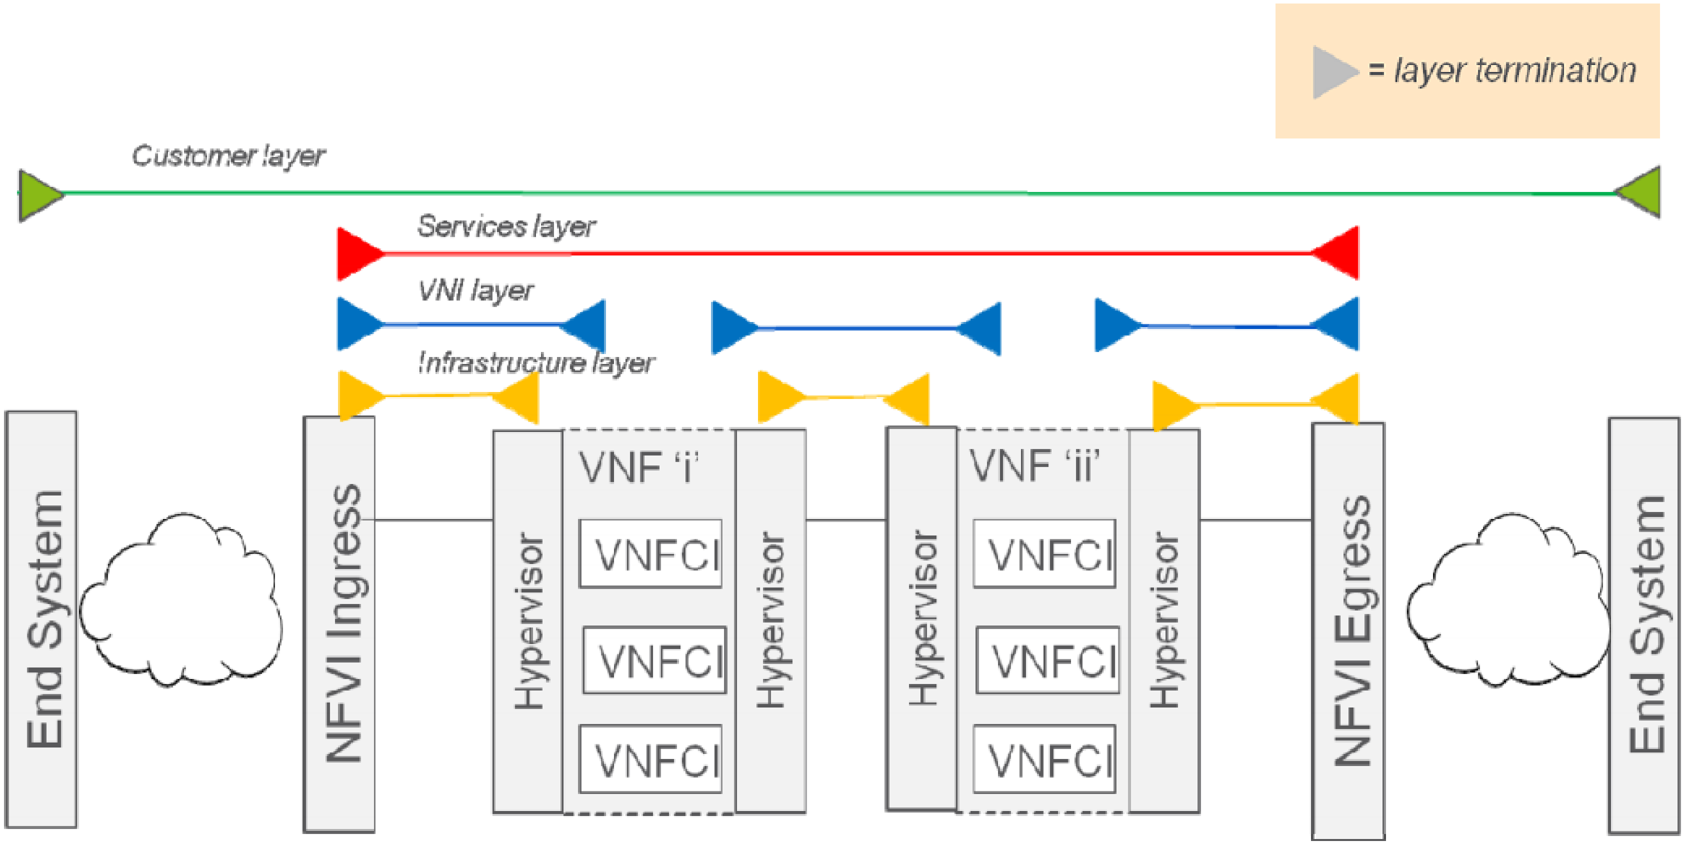

[5] ETSI NFV; Infrastructure; Network Domain

Traffic Isolation

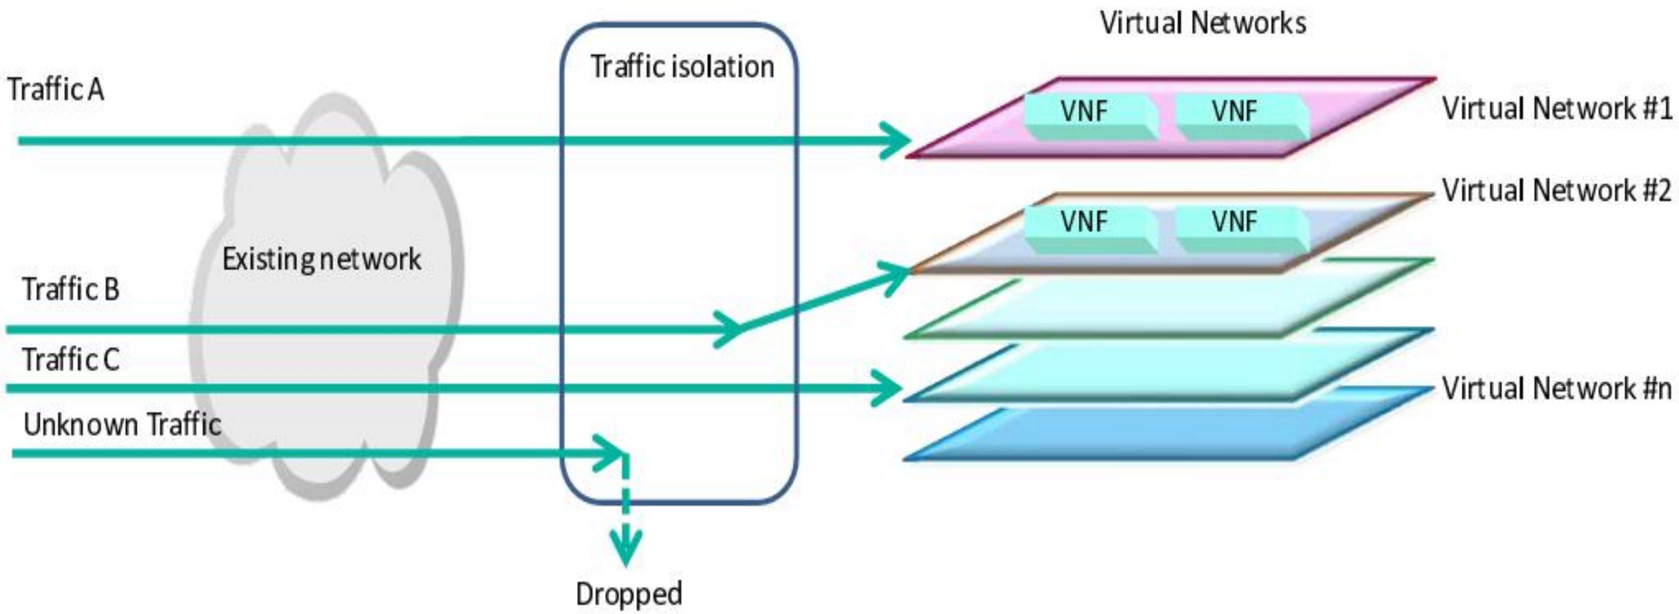

[5] ETSI NFV; Infrastructure; Network Domain

VNF – FG (Virtual Network Function – Forwarding Graph)

Logical View of a VNF-FG

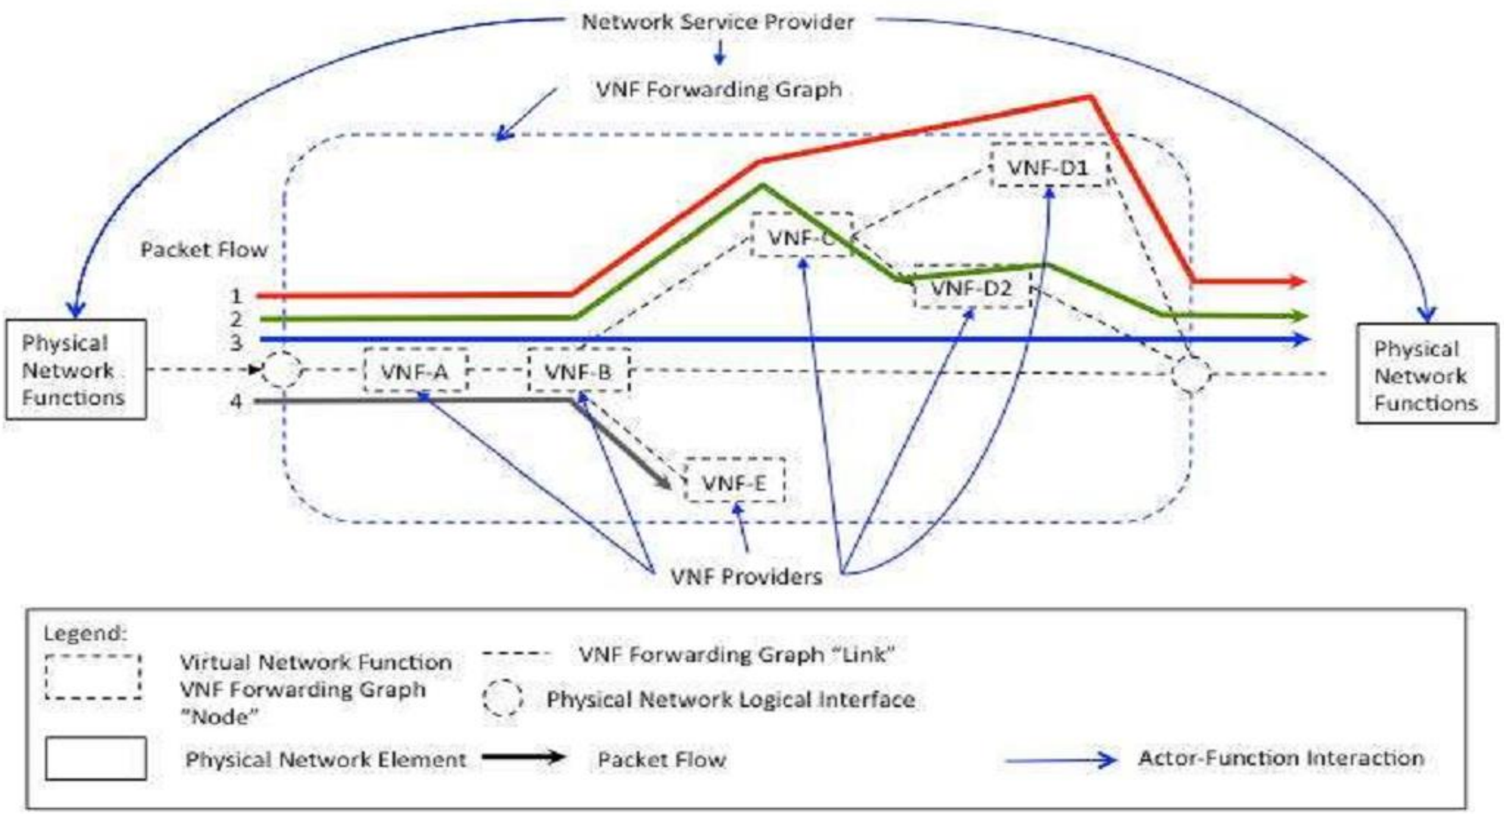

[1] ETSI NFV; Use Cases

Continued….

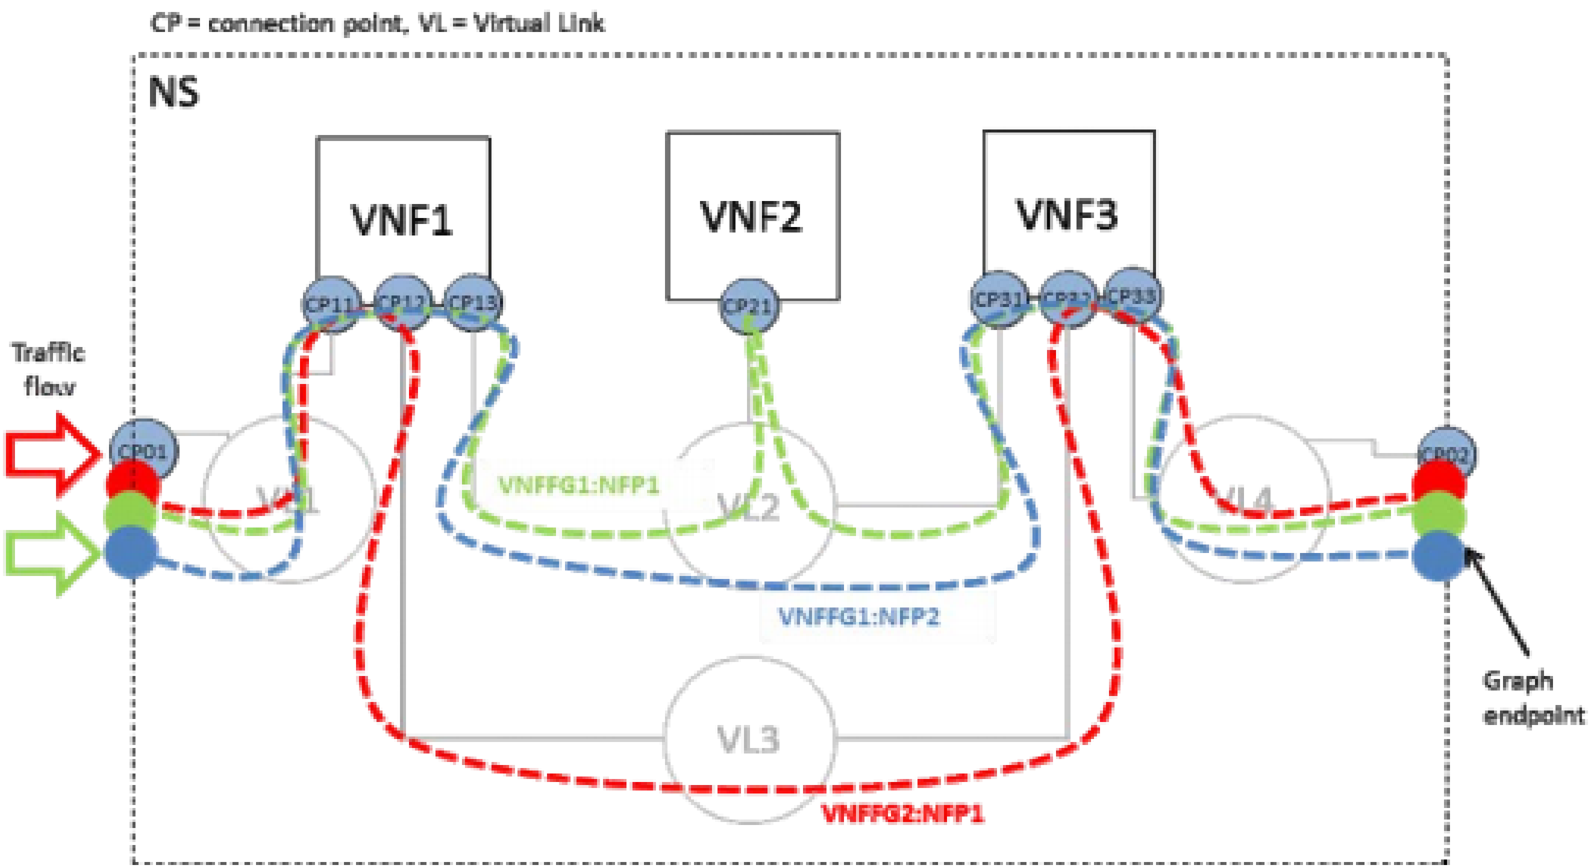

[6] ETSI NFV; Management and Orchestration

Continued…

Physical View of a VNF-FG

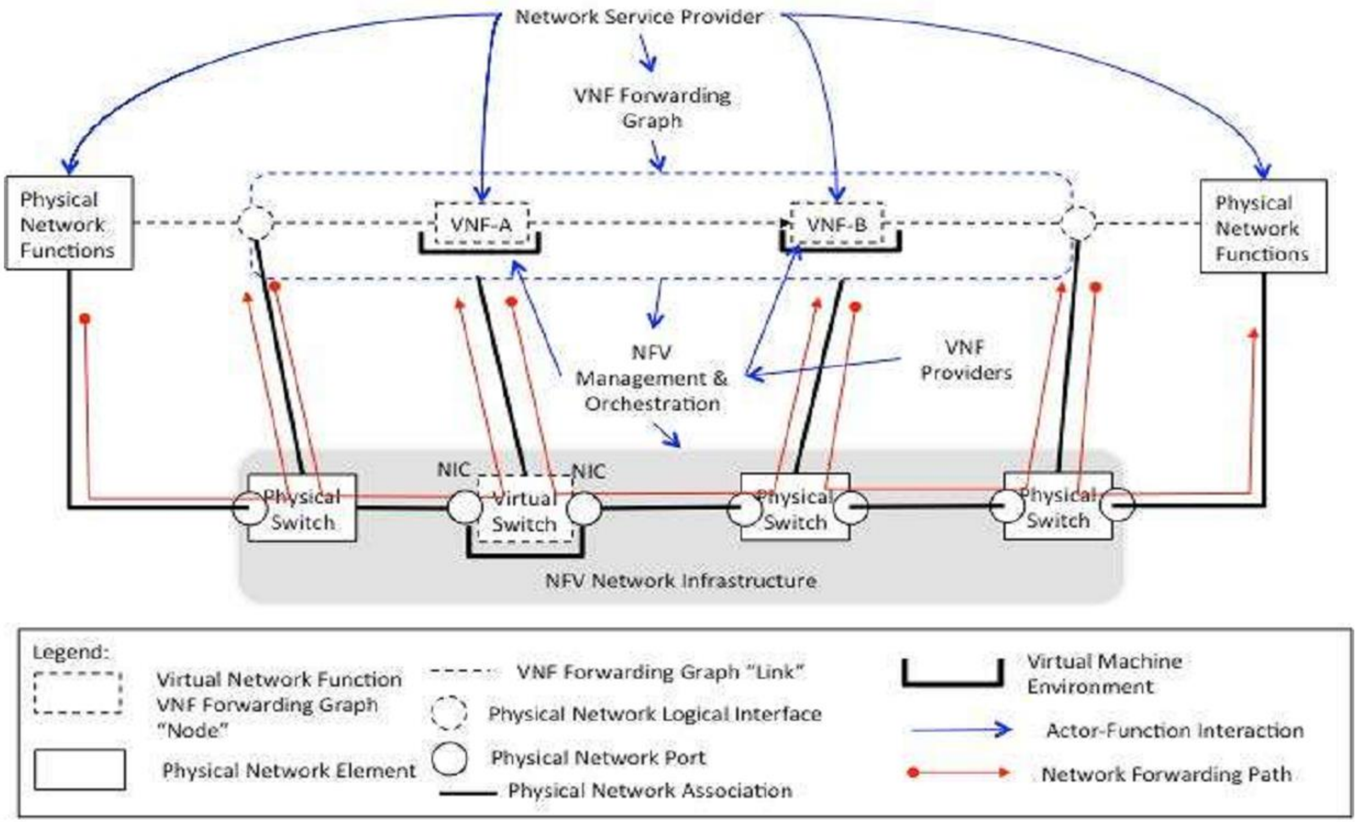

Network Service

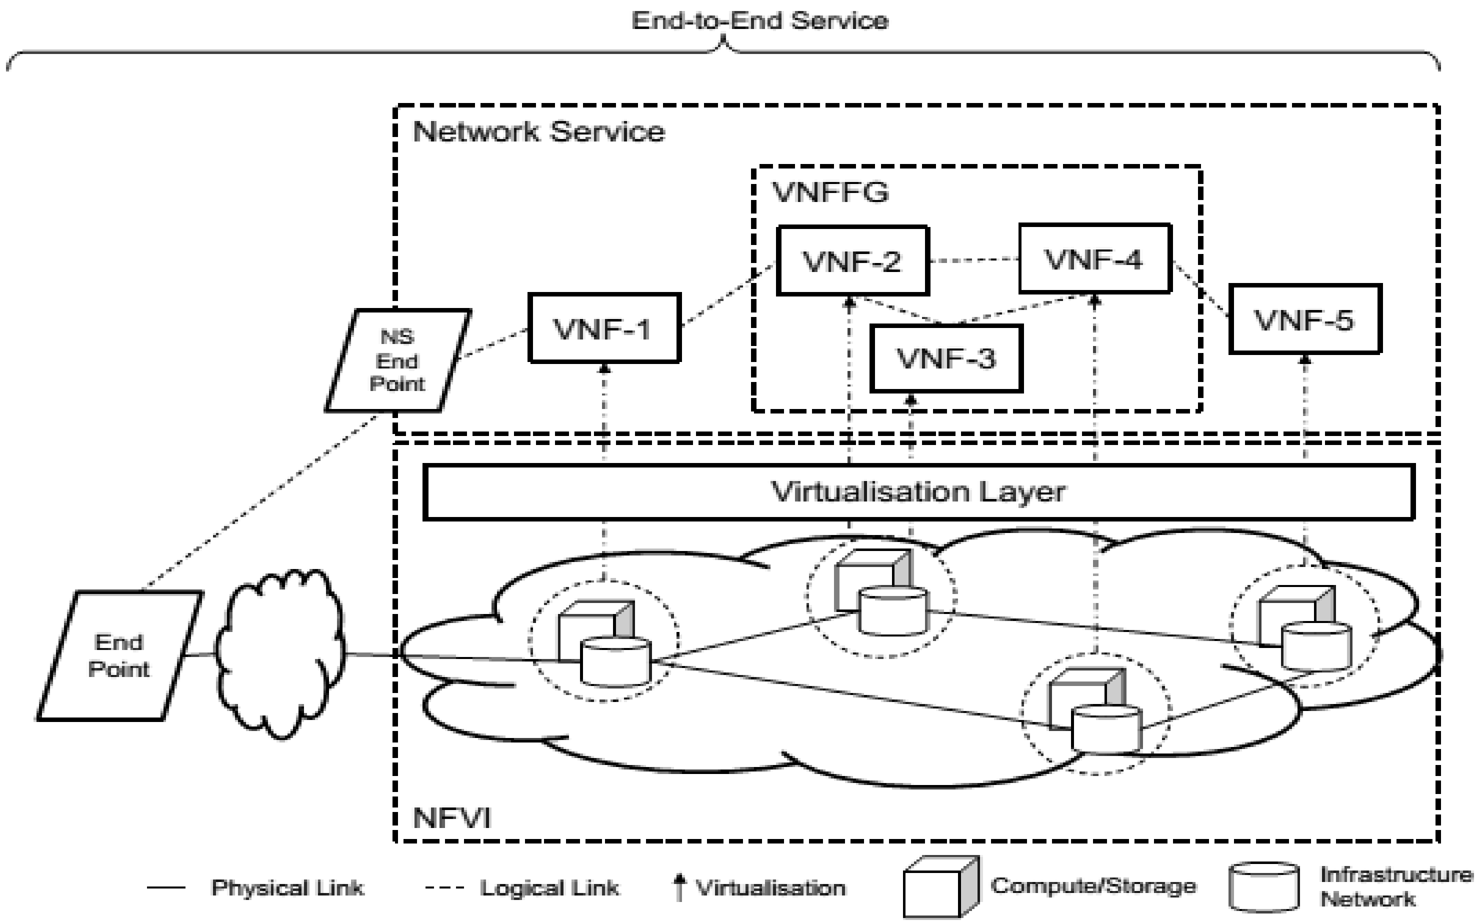

[6] ETSI NFV; Management and Orchestration

Continued…

- A service may be
  - Orchestrated out of existing services i.e. the VNF-FG
  - Deployment of new VNFs i.e. add new nodes to the VNF-FG
  - Or can be a combination of both

- A FG overall may be composed of both physical and virtual network functions

- If the VNFs are not required to be in a strict order then we will have a VNF-Set (VNF-S)

UC DAVIS UNIVERSITY OF CALIFORNIA

[1] ETSI NFV; Use Cases

Challenges in VNF-FG deployment

- **Specifying attributes of the VNF-FG. Especially, how to best define and also determine the dependency between the VNFs, that are to be deployed for a service.**

- **Finding measuring techniques for the above.**

- **Supporting network services that cross administrative domains.**

[1] ETSI NFV; Use Cases

Compute domain

Functional elements of the compute domain

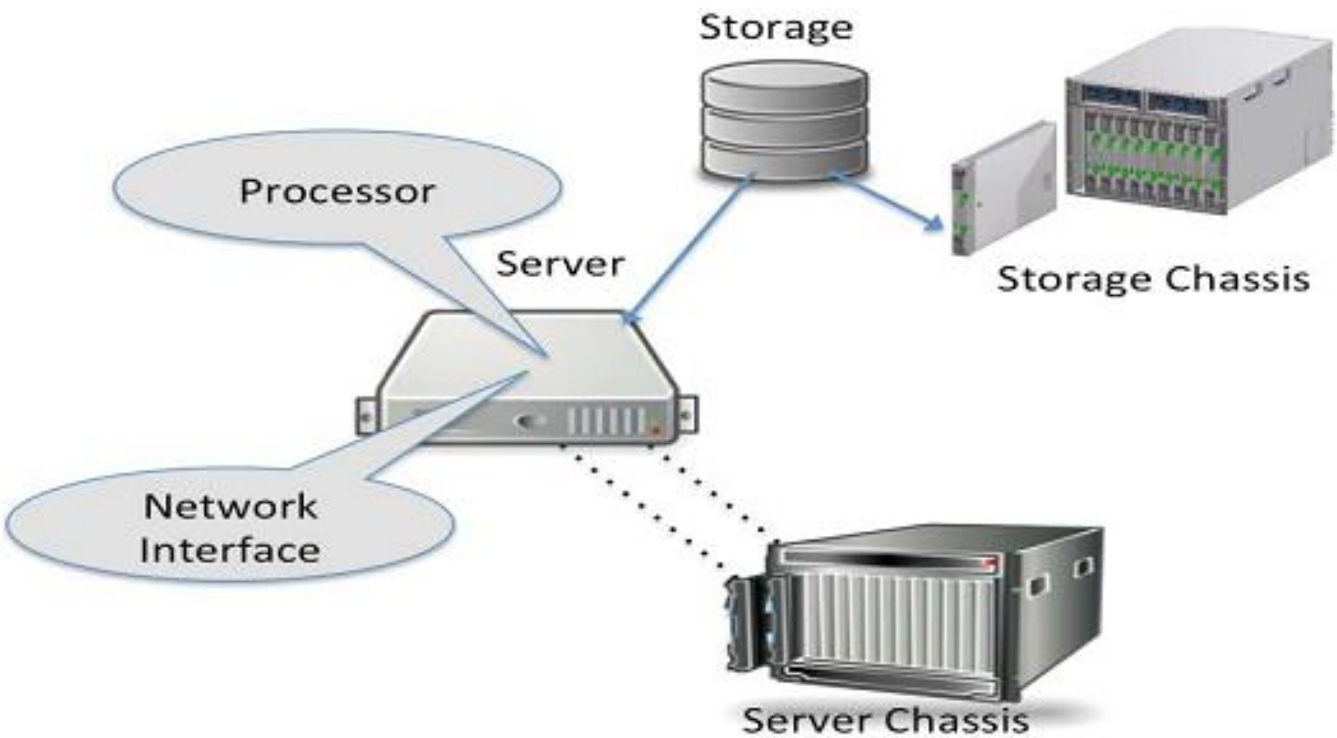

[7] ETSI NFV; Infrastructure; Compute Domain

Disaggregation Model

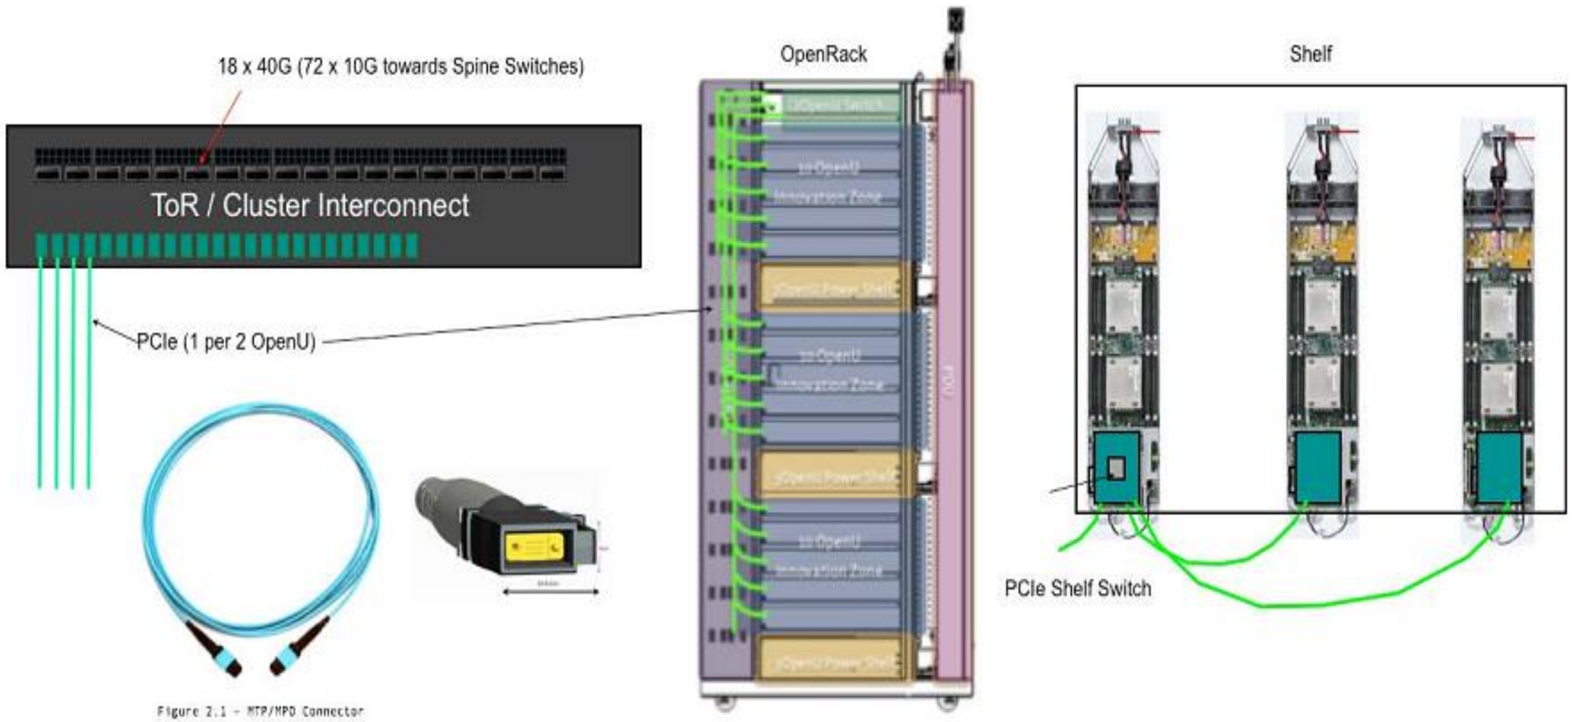

[7] ETSI NFV; Infrastructure; Compute Domain

NFVI deployment scenarios

Deployment Scenario	Building	Host Hard-ware	Hyper-visor	Guest VNF	cf. NIST Cloud Model
Monolithic Operator	N	N	N	N	Private Cloud
Network Operator Hosting Virtual Network Operators	N	N	N	N, N1, N2	Hybrid Cloud
Hosted Network Operator	H	H	H	N	
Hosted Communications Providers	H	H	H	N1, N2, N3	Community Cloud
Hosted Communications and Application Providers	H	H	H	N1, N2, N3, P	Public Cloud
Managed Network Service on Customer Premises	C	N	N	N	
Managed Network Service on Customer Equipment	C	C	N	N	
NOTE: The different letters represent different companies or organisations, and are chosen to represent different roles, e.g. H = hosting provider, N = network operator, P = public, C = customer.					

UCDAVIS
UNIVERSITY OF CALIFORNIA

[7] ETSI NFV; Infrastructure; Compute Domain

Potential NFVI scale

# N-PoPs	Example Location Types	Example NE types	NFV Use cases
1 - 10	IT Data Center	firewall	NFVIaaS, VNFaaS, VNPaaS, VNF Forwarding Graph
10 - 100	Major Central Office	router	vIMS
100 - 1 000	Minor Central Office	BRAS	
1 000 - 1 000 000	Curbside Cabinet, Cell Tower	eNodeB, DSLAM	vCDN vMobile Base Station
1 000 000 +	Subscriber Premises	CPE, Mobile Devices, IoT	vCPE, vFixed Access Network

[7] ETSI NFV; Infrastructure; Compute Domain

NFVI hardware elements

- Rack
- Fabric
- TOR switch
- Power grid
- Rack shelf
- Server chassis
- Storage chassis
- Accelerator blades/chassis

[7] ETSI NFV; Infrastructure; Compute Domain

NFV architecture scope within the NFV reference architecture framework

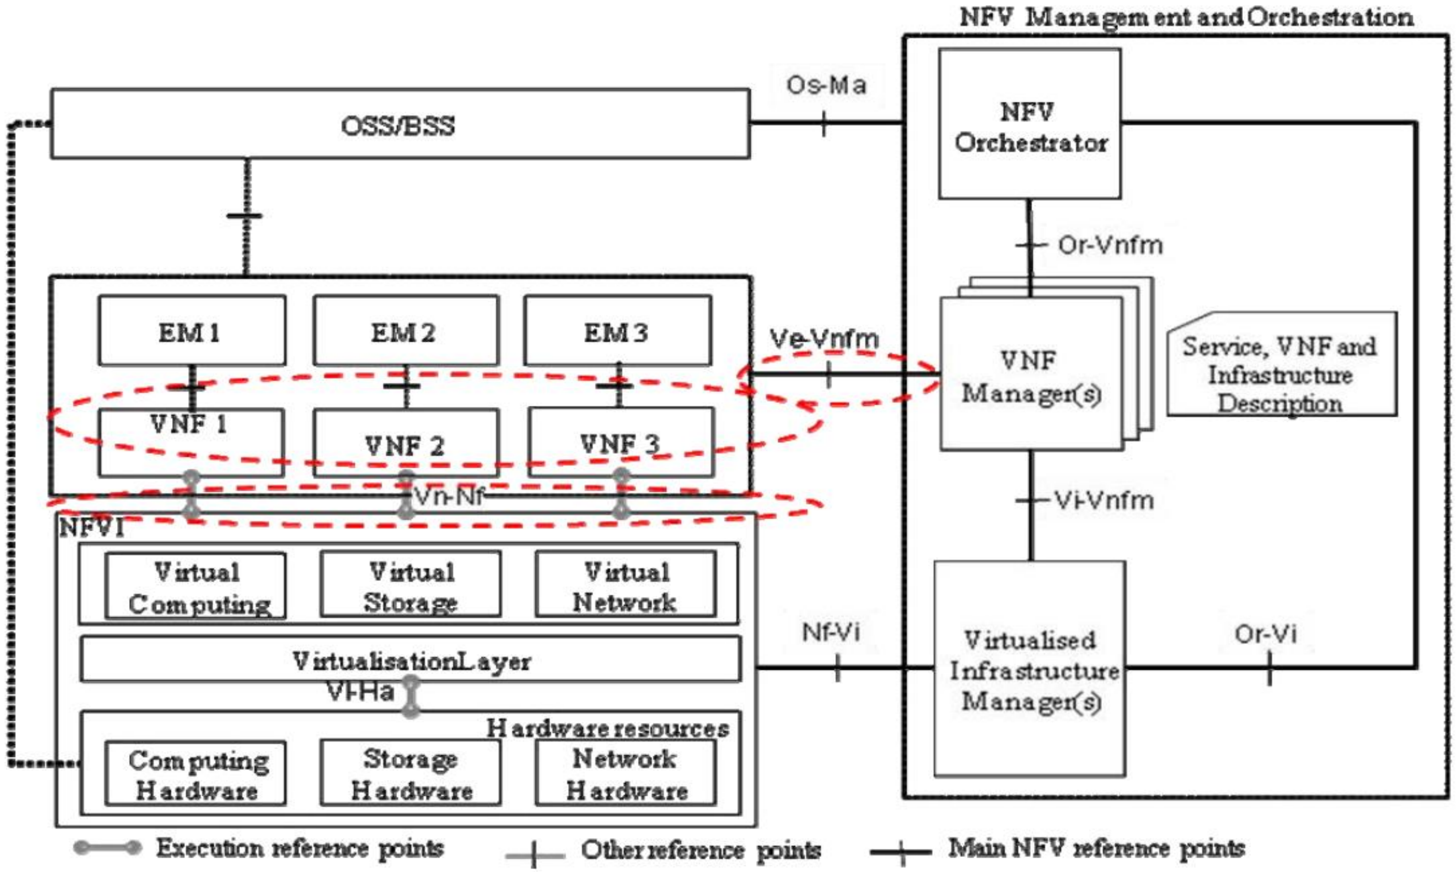

[8] ETSI NFV; Virtual Network Function Architecture

VNF Architecture

- **VNF Descriptor (VNFD) – Gives the resource requirements of a VNF instance. Provided by the VNF vendor**
- **VNFC – Virtual Network Function Component**
- **VNFCI – Virtual Network Function Component Instance**

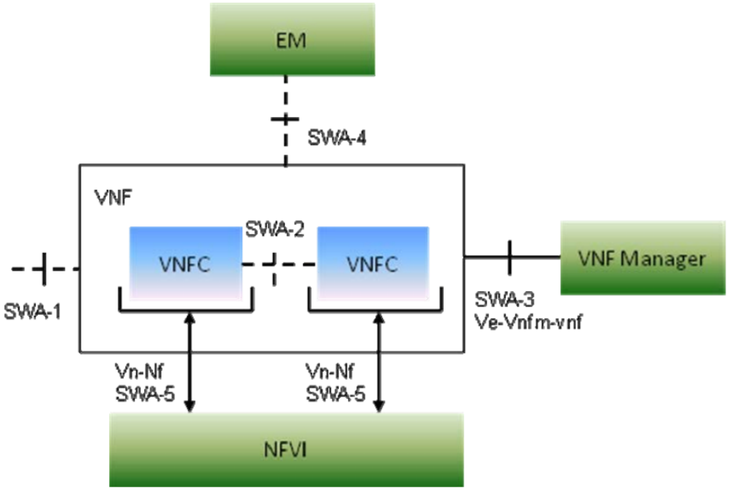

[8] ETSI NFV; Virtual Network Function Architecture

VNF Workload Types

- **Data plane workload e.g. CDN cache node, vRouter**
- **Control plane workload e.g. authentication**
- **Signal processing workloads e.g. FFT decoding and encoding in a C-RAN BBU**
- **Storage workloads**

[9]ETSI NFV; NFV Performance & Portability Best Practices

Workloads and NFV use cases

Use case	VNF	Data plane			Control plane			Signal processing	Storage	
		Edge NF	Intermediate NF	Intermediate NF with encryption	Routing	Authentication	Session management	Signal processing	Non-intensive	R/W intensive
1. NFVI as a Service	N/A									
2. Virtual Network Platform as a Service	N/A									
3. VNF as a Service	E-CPE (Enterprise-CPE)		X		X	X			X	
	PE		X	X	X	X			X	
	FW		X						X	
	DPI		X							X
4. Virtualisation of Mobile Core Network and IMS	MME				X		X		X	
	S-GW			X	X				X	
	P-GW		X		X	X			X	
	PCRF				X		X		X	
	SGSN			X	X		X		X	
	GGSN		X		X	X	X		X	
	P-CSCF		X		X	X	X		X	
	S-CSCF				X	X	X		X	
	I-CSCF				X		X		X	
	MGCF				X		X		X	
	AS	X					X		X	
5. Virtualisation of Mobile Base Station	BBU						X	X		
6. Virtualisation of the Home Environment	RGW		X		X	X			X	
	STB	X							X	
7. Service chains	N/A									
8. Virtualisation of CDNs	CDN Cache Node	X								X
	CDN Controller				X		X		X	
9. Fixed Access Network Functions Virtualisation	OLT		X		X			X	X	
	DSLAM		X		X	X		X	X	
	ONU		X					X	X	
	ONT		X		X			X	X	

[9]ETSI NFV; NFV Performance & Portability Best Practices

VNF States

- Statefulness will create another level of complexity, e.g. a session (transaction) consistency has to be preserved and has to be taken into account in procedures such as load balancing
- The data repository holding the externalized state may in the same VNF
- The data repository holding the externalized state may be an external VNF

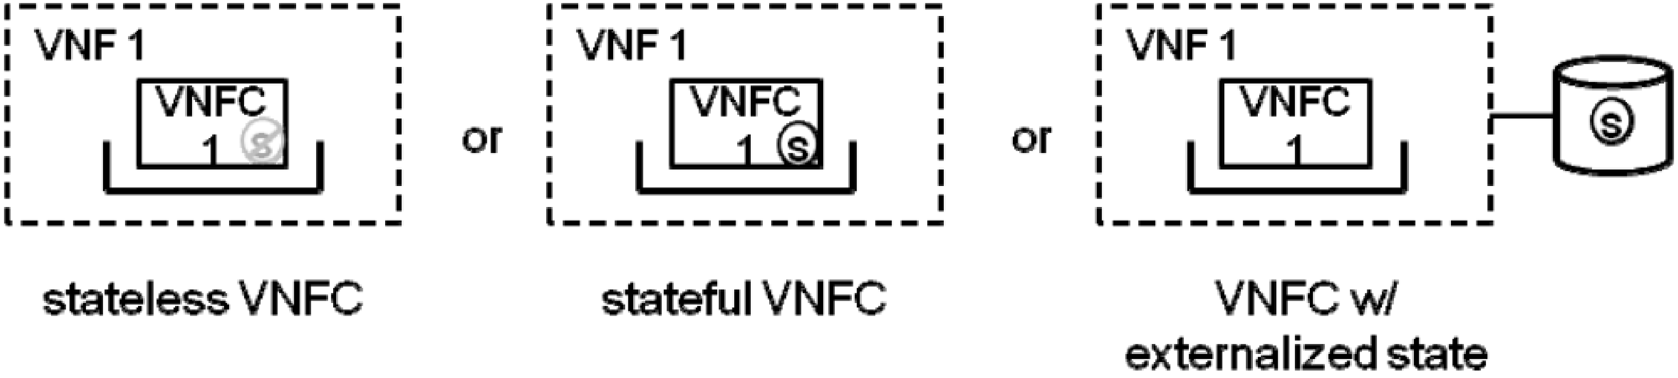

[8] ETSI NFV; Virtual Network Function Architecture

VNF Load Balancing models

- **VNF internal load balancer**

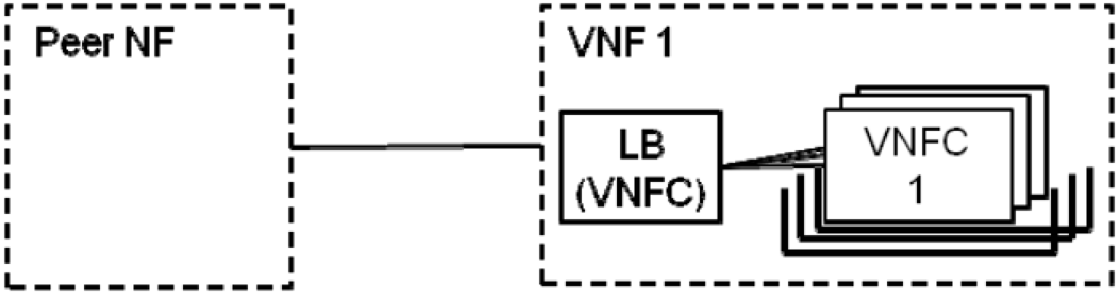

VNF-internal load balancer

- **VNF-external Load Balancer**

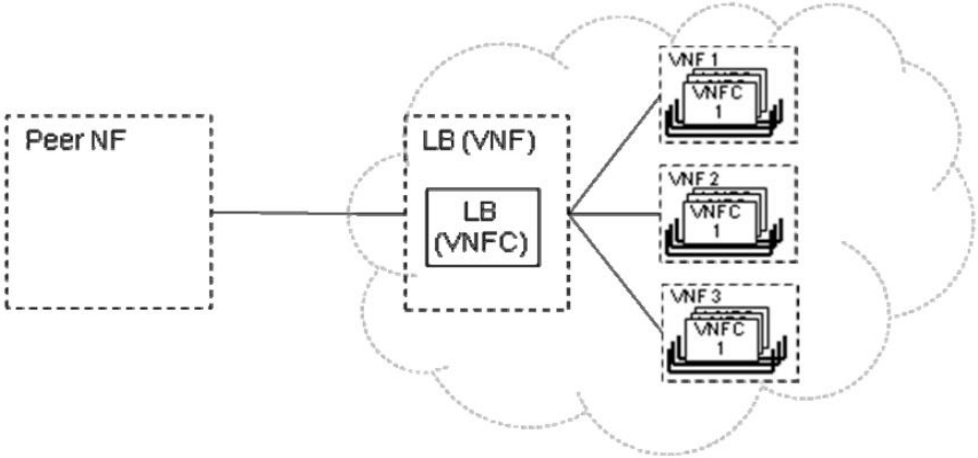

[8] ETSI NFV; Virtual Network Function Architecture

Continued…

- **End-to-End Load Balancing**

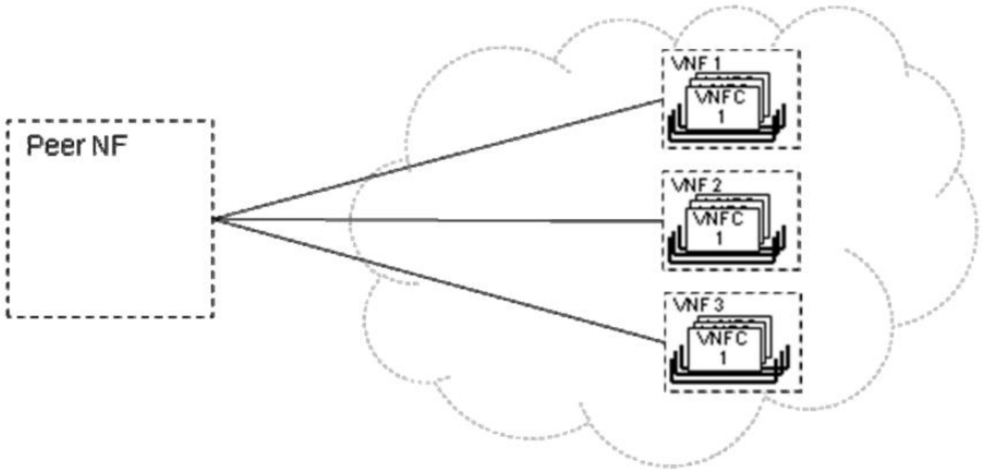

- **Infrastructure Network Load Balancer**

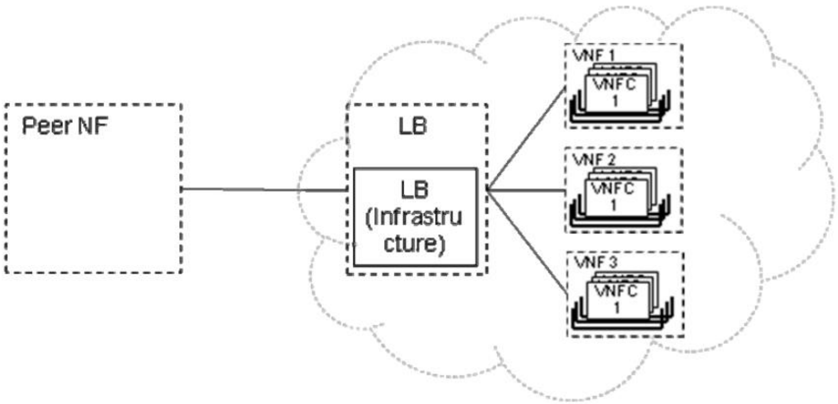

[8] ETSI NFV; Virtual Network Function Architecture

Elasticity of VNFs

- No elasticity
  - The VNF requires a fixed set of resources that cannot be changed.
- Elasticity by scaling up/down only
  - The NFV framework can increase or decrease the size, performance or bandwidth of the virtual resources
- Elasticity by scaling out/in only
  - The VNF is able to perform scaling operations by separately adding/removing instances of specified VNFCs
- Elasticity in either dimensions
  - The VNF has VNFCs that may be scaled out/in, up/down or both

[8] ETSI NFV; Virtual Network Function Architecture

VNF migration capability

- No live migration supported
- Live migration supported
- Partial Migration
- Other schemes – shutdown/restart while moving from one hardware to the other
- There may be further scalability/reliability/redundancy constraints on how many VNFs can be deployed on a single physical host

[8] ETSI NFV; Virtual Network Function Architecture

VNF States and Transitions (from orchestration)

State	Description
Null	A VNF Instance does not exist and is about to be created.
Instantiated Not Configured	VNF Instance does exist but is not configured for service.
Instantiated Configured - Inactive	A VNF Instance is configured for service.
Instantiated Configured - Active	A VNF Instance that participates in service.
Terminated	A VNF Instance has ceased to exist.

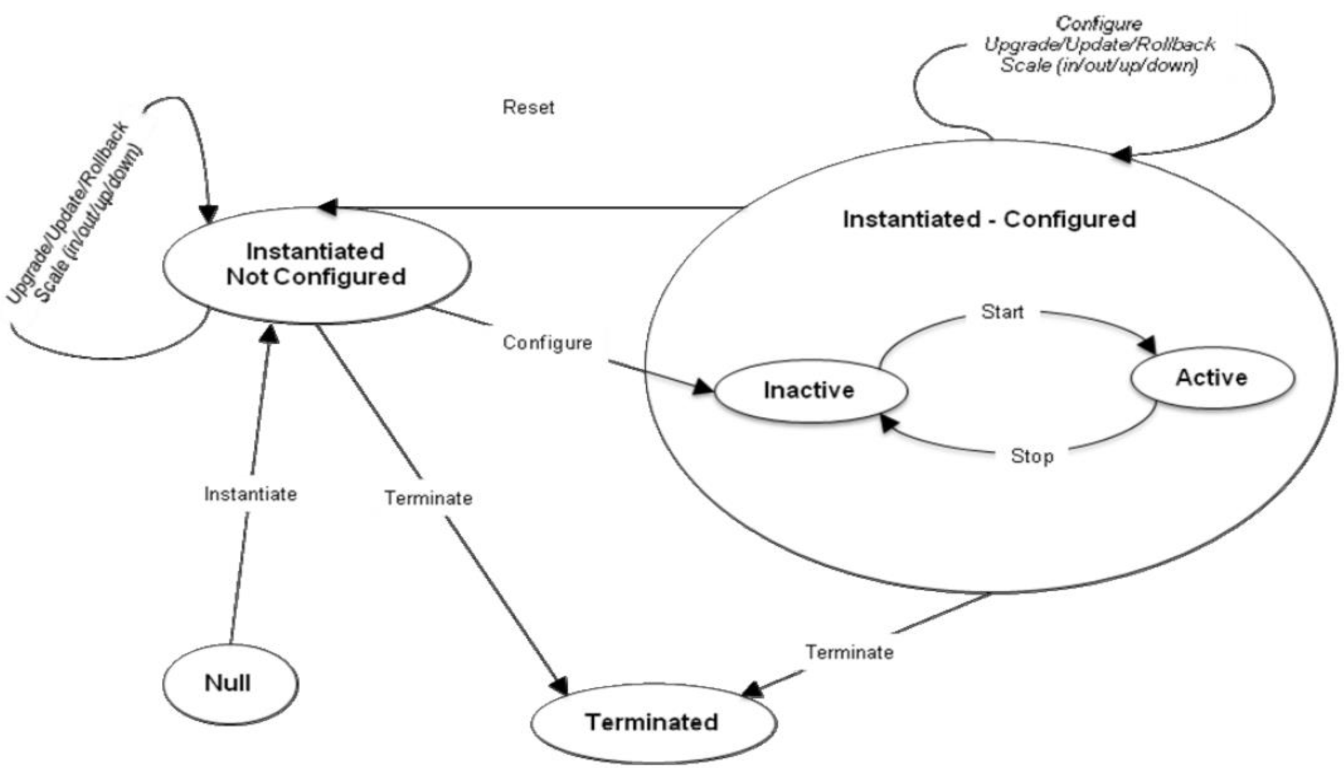

[8] ETSI NFV; Virtual Network Function Architecture

VNF Architecture Design Example

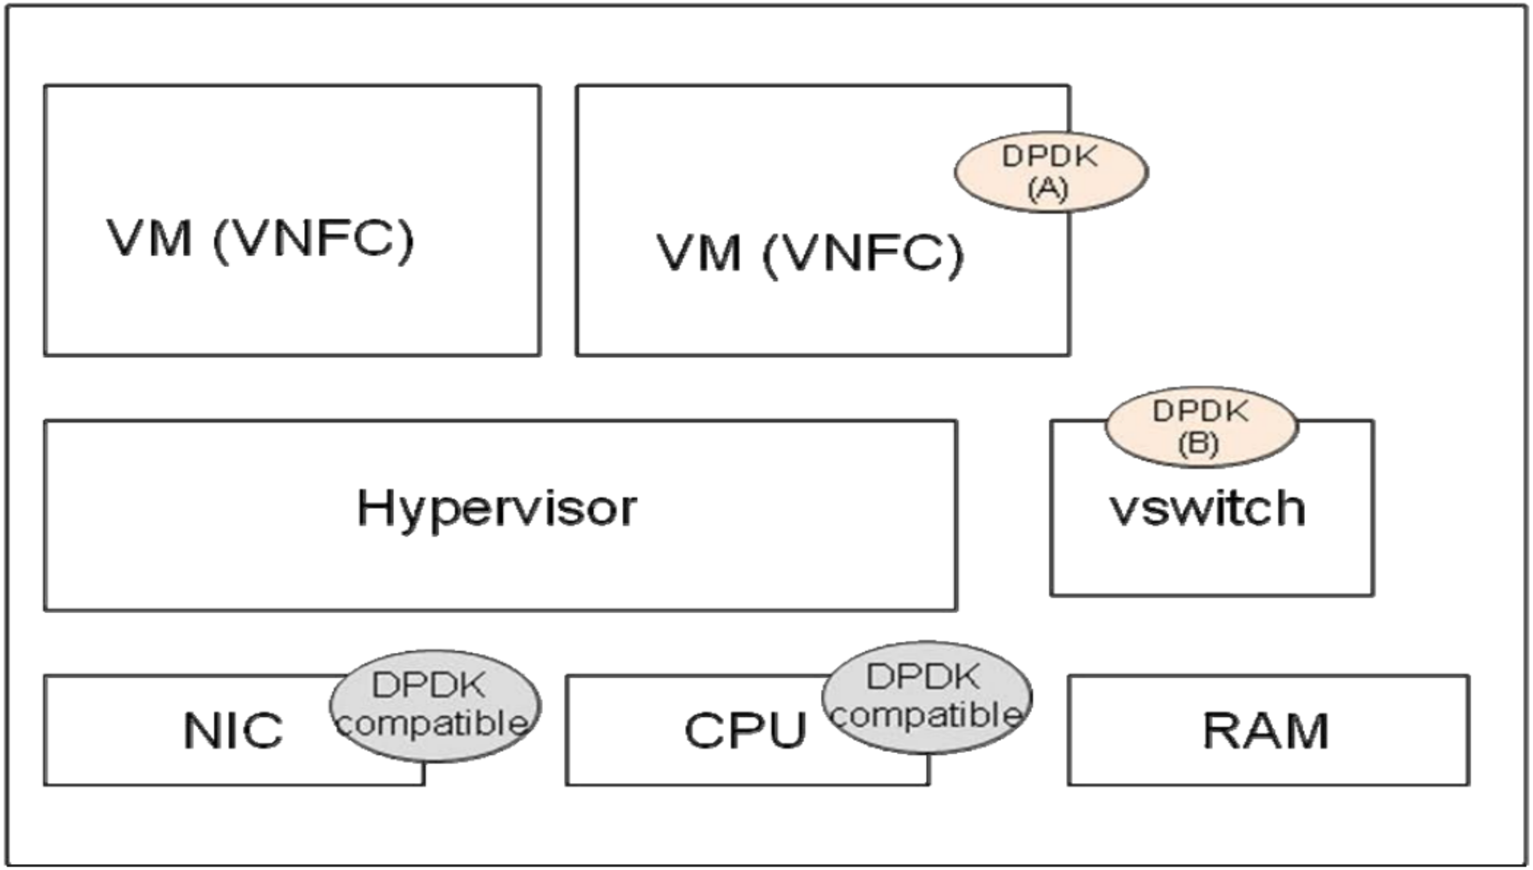

[8] ETSI NFV; Virtual Network Function Architecture

VNFC/VNF to VNFC/VNF communication

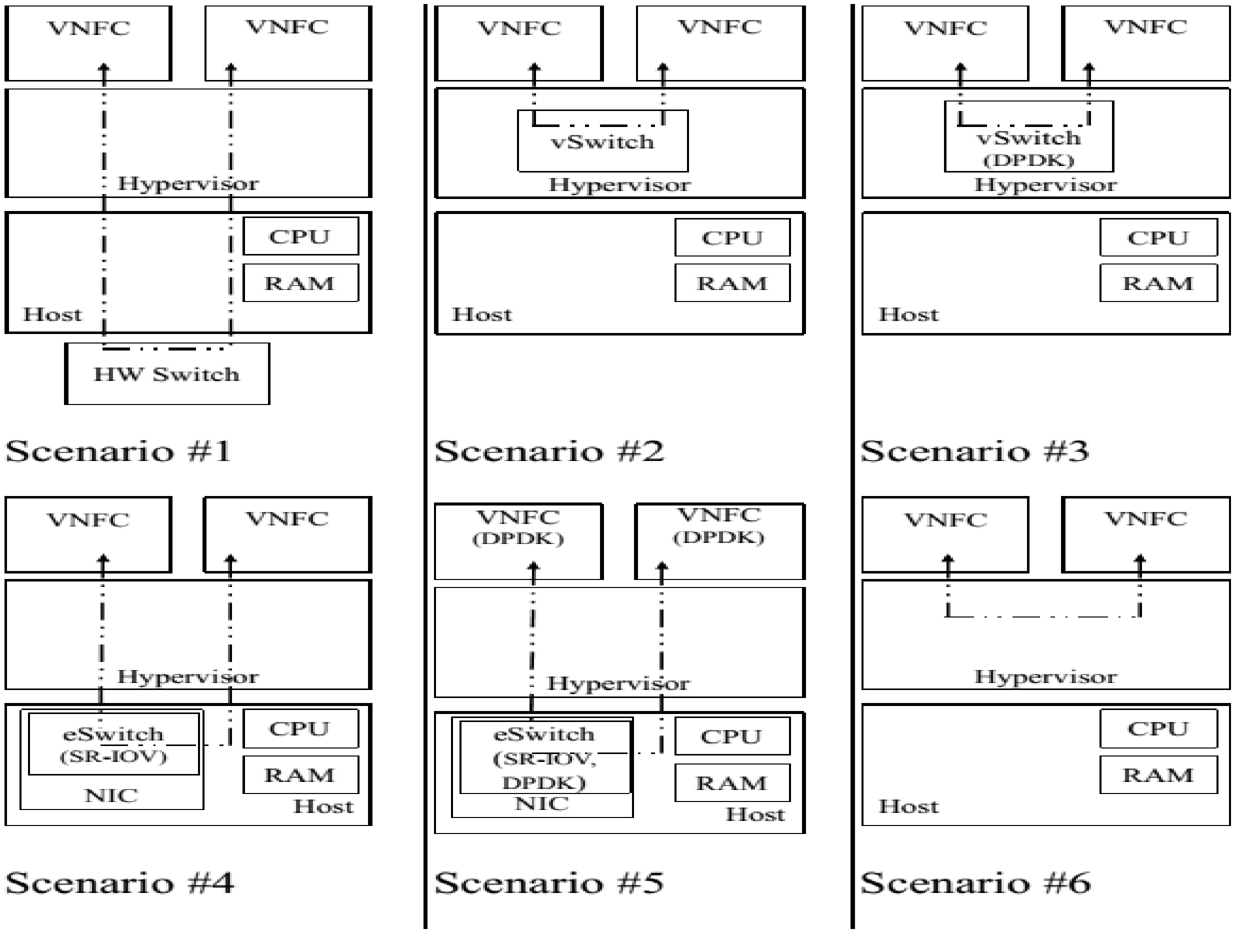

[8] ETSI NFV; Virtual Network Function Architecture

NFV Interconnection options

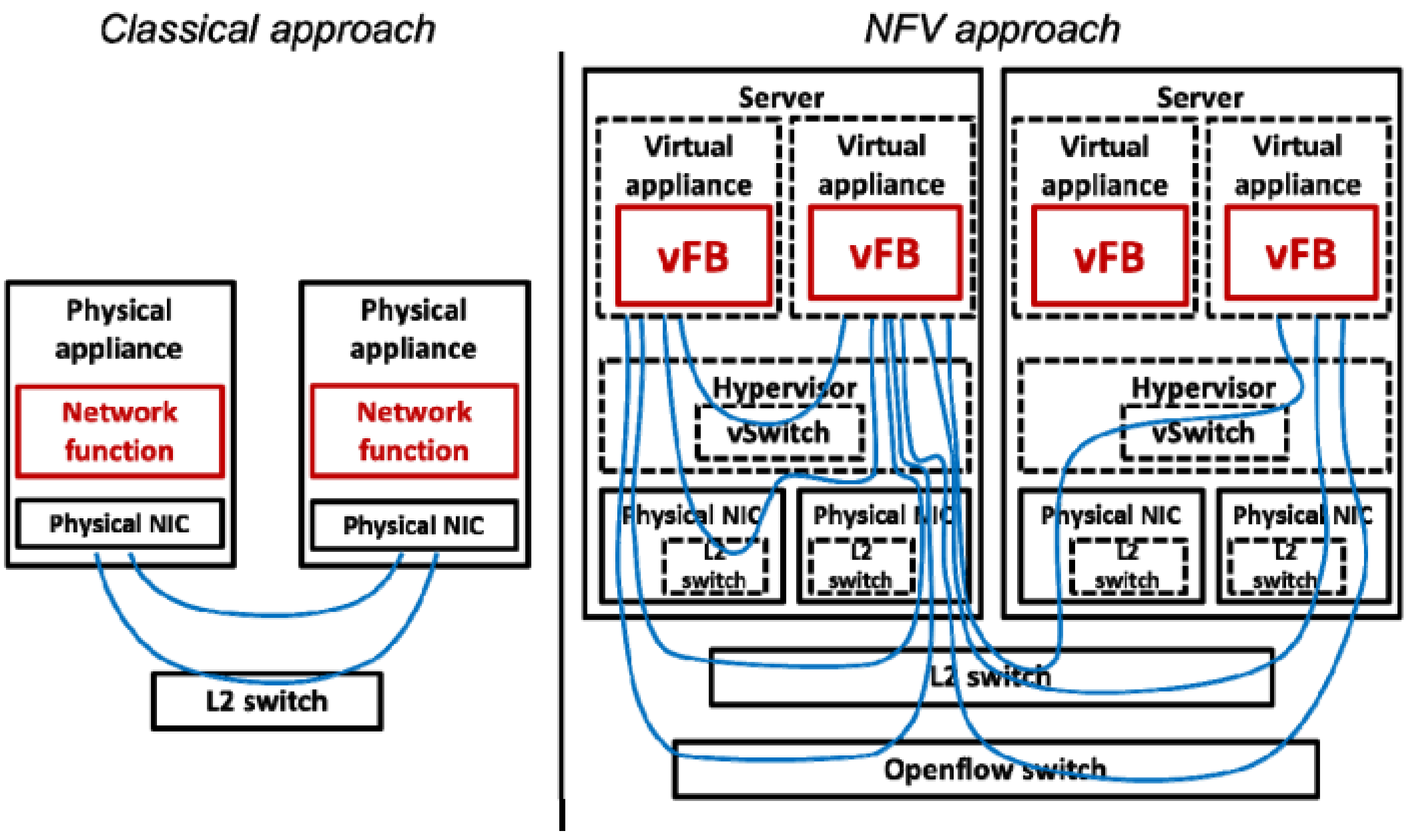

[1]"ETSI NFV; Infrastructure Overview"

Continued…

- **VNFC Memory to VNFC Memory - VNFC must have affinity for shared memory access (for low latency data paths between VNFs)**

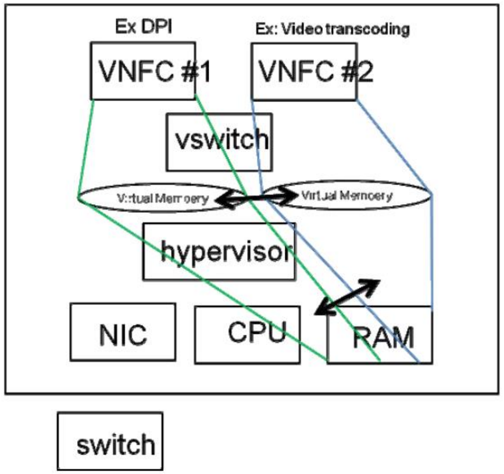

- **Faster Network Access - SR-IOV (single root input/output virtualization) is used ensure latency between the NIC and the VM is optimized. This enables bare metal performance but requires that the server NIC's support SRIOV.**

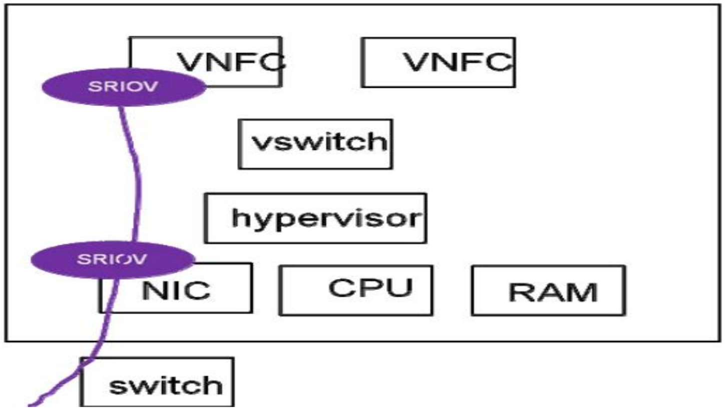

[8] ETSI NFV; Virtual Network Function Architecture

Fast storage access

- RDMA (Remote Direct Memory Access) bypasses the operating system and the CPU
- iSCSI extensions for RDMA (iSER) accelerate hypervisor traffic, includes storage access, VM migration, and data and VM replication.
- RoCE : RDMA over Converged Ethernet.
- ETH (ethernet support from 40 Gbps today to 100 Gbps in the future) and IB (InfiniBand support from 56 Gbps today to 104 Gbps in the future)networks are utilized.

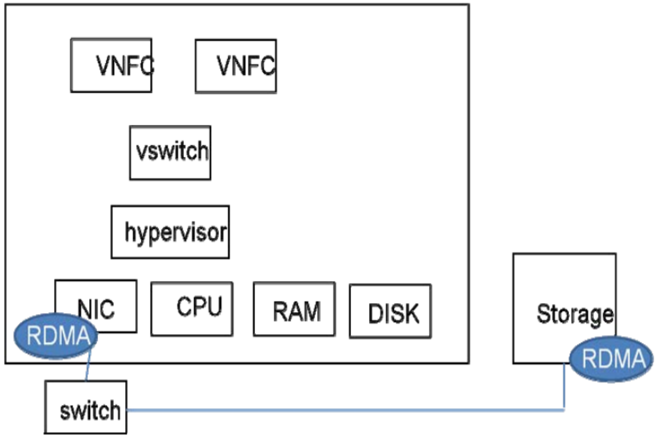

[8] ETSI NFV; Virtual Network Function Architecture

Interfaces within the compute domain

- PCIe
- SR-IOV
- RDMA and RoCE support
- InfiniBand
- DPDK & ODP Support

UC DAVIS
UNIVERSITY OF CALIFORNIA

[7] ETSI NFV; Infrastructure; Compute Domain

Technologies to setup virtual network paths (virtual path setup cost)

- **Virtual Local Area Network (VLAN)**
- **Virtual Private LAN Service (VPLS)**
- **Virtual extensible Local Area Network (VxLAN)**
- **Network Virtualisation using Generic Routing Encapsulation (NVGRE)**

[3] ETSI NFV; Architectural Framework

Performance – HW Implementation

- CPU Architecture and basic performance levels - PCI-E, Infiniband, clock speed, Number of Cores, etc.
- NFVI Compute Node architecture and connectivity - Bus structure and BW, NIC performance, Memory and Storage structure.
- Cache structure and sizes.
- Support for large memory pages and extended TLB caches with large pages.
- Support for SR-IOV (classification of packets in independent, per VM TX/RX queues) and other traffic & data optimization technologies.
- Direct I/O access to processor memory & OS, e.g. Architectures that support DMA (Direct Memory Access) and RDMA.
- IOMMU or translation services for I/O.

UC DAVIS UNIVERSITY OF CALIFORNIA

[7] ETSI NFV; Infrastructure; Compute Domain

Performance – Acceleration technologies

- **On Chip HW based acceleration - e.g. AES, CRC, Cryptography, Transcoding.**
- **Compute intensive acceleration - e.g. Heterogeneous Computing/GPU.**
- **Compute acceleration pool.**
- **Network intensive function acceleration - e.g. NP, FPGA, CPU based support for data plane workload acceleration and data traffic optimization - e.g. NAT, ACL, DPI.**
- **Storage acceleration - e.g. Storage Clusters.**
- **NIC based acceleration - e.g. SR-IOV, vSwitch Bypass, Network Intensive processing.**

UC DAVIS
UNIVERSITY OF CALIFORNIA

[7] ETSI NFV; Infrastructure; Compute Domain

Service Quality Metrics

- Virtual machine service quality metrics.

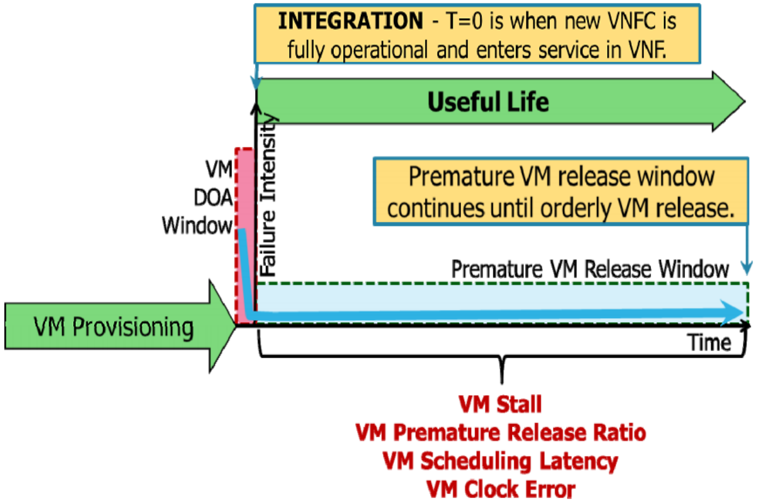

- Virtual network service quality metrics.
- Technology components offered 'as-a-Service' (e.g. Database-as-a-Service) quality metrics.
- Orchestration service quality metrics.

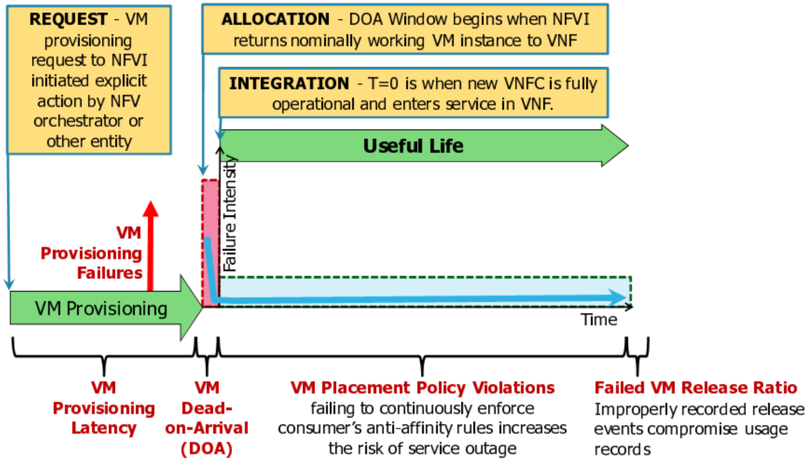

UC DAVIS UNIVERSITY OF CALIFORNIA

[10] ETSI NFV; Service Quality Metric

QoS metrics

- Throughput
- Latency
- Frame loss rate
- Back-to-Back Frame Rate
- Packet delay variation
- Service Disruption Time for Fail-over convergence

[7] ETSI NFV; Infrastructure; Compute Domain

Dynamic Hardware Metrics (RA costs?)

Resource	Metrics	Examples and Units
Physical Server Blade	Initialization Time	• Time from power activation until "in-service", informing all necessary managers and orchestrators that the resources on the server/blade are ready for commissioning (*in the form of an operating Hypervisor ready to instantiate a VM to serve an VNFi.* • Is this wording sufficiently accurate?).
	Failed Initializations	• Count of attempted Initializations that do not result in the server/blade reaching the state of readiness for commissioning.
Storage Unit (i.e. disk)	Activation Time	Time to fully stable, active use in array (such as adding/replacing a member of a RAID array? Better example?).
	Failed Activations	• Count of attempted Activations that do not result in the storage unit reaching the state of readiness.

[4] ETSI NFV; Infrastructure; Hypervisor Domain

CPU-related dynamic hardware activation metrics

	Speed	Accuracy	Reliability
Activation	Initialization time		Failed Initializations,
Operation	Available core count, Per core: temperature, Idle power state residency Operating voltage/frequency point residency, Cache utilization	LLC misses, TLB misses	Time to Failure Error-free Processing Time
De-activation/Deletion/Take-down			

[4] ETSI NFV; Infrastructure; Hypervisor Domain

References

- [1]ETSI NFV; Use-Cases
- [2]ETSI NFV; Infrastructure Overview
- [3]ETSI NFV; Architectural Framework
- [4]ETSI NFV; Infrastructure; Hypervisor Domain
- [5]ETSI NFV; Infrastructure; Network Domain
- [6]ETSI NFV; Management and Orchestration
- [7] ETSI NFV; Infrastructure; Compute Domain
- [8] ETSI NFV; Virtual Network Function Architecture
- [9]ETSI NFV; NFV Performance & Portability Best Practices
- [10] ETSI NFV; Service Quality Metric

UCDAVIS
UNIVERSITY OF CALIFORNIA