

Virtualização de Redes MC833


Nelson L. S. da Fonseca

<http://www.ic.unicamp.br/~nfonseca/comsoc-school/2017/>


Breve Revisão de Conceitos de Redes

Addressing: routing to another LAN

- ❖ A creates IP datagram with IP source A, destination B
- ❖ A creates link-layer frame with R's MAC address as dest, frame contains A-to-B IP datagram

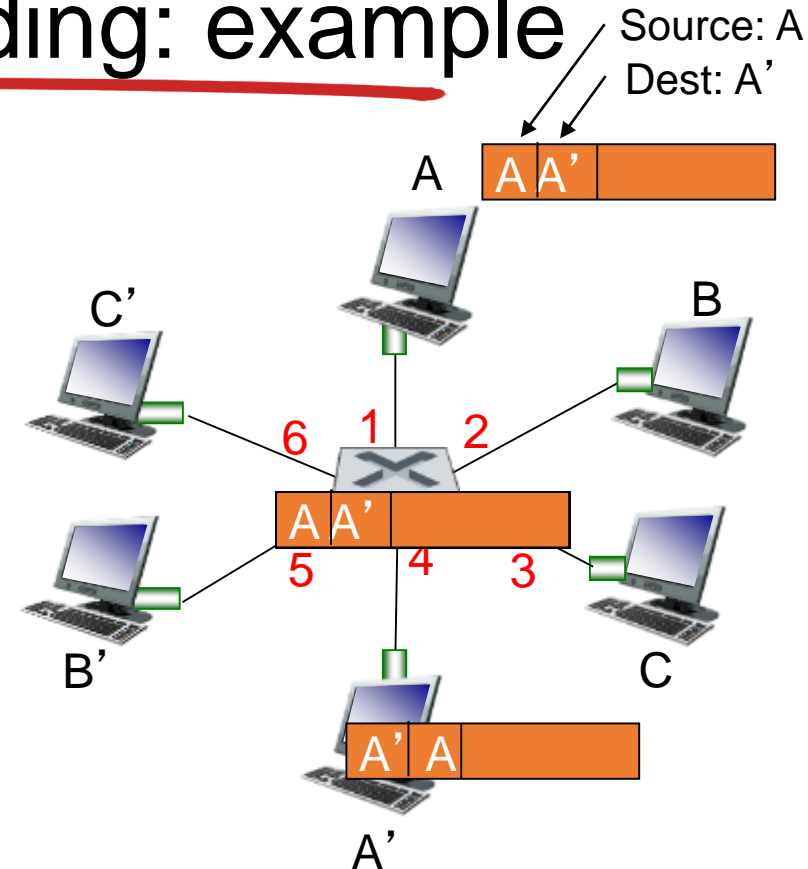


Switch Layer 2



Self-learning, forwarding: example

- frame destination, A', location unknown: *flood*
- destination A location known: *selectively send*
on just one link




MAC addr	interface	TTL
A	1	60
A'	4	60

*switch table
(initially empty)*

Interconnecting switches

self-learning switches can be connected together:



Q: sending from A to G - how does S₁ know to forward frame destined to G via S₄ and S₃?

- A: self learning! (works exactly the same as in single-switch case!)

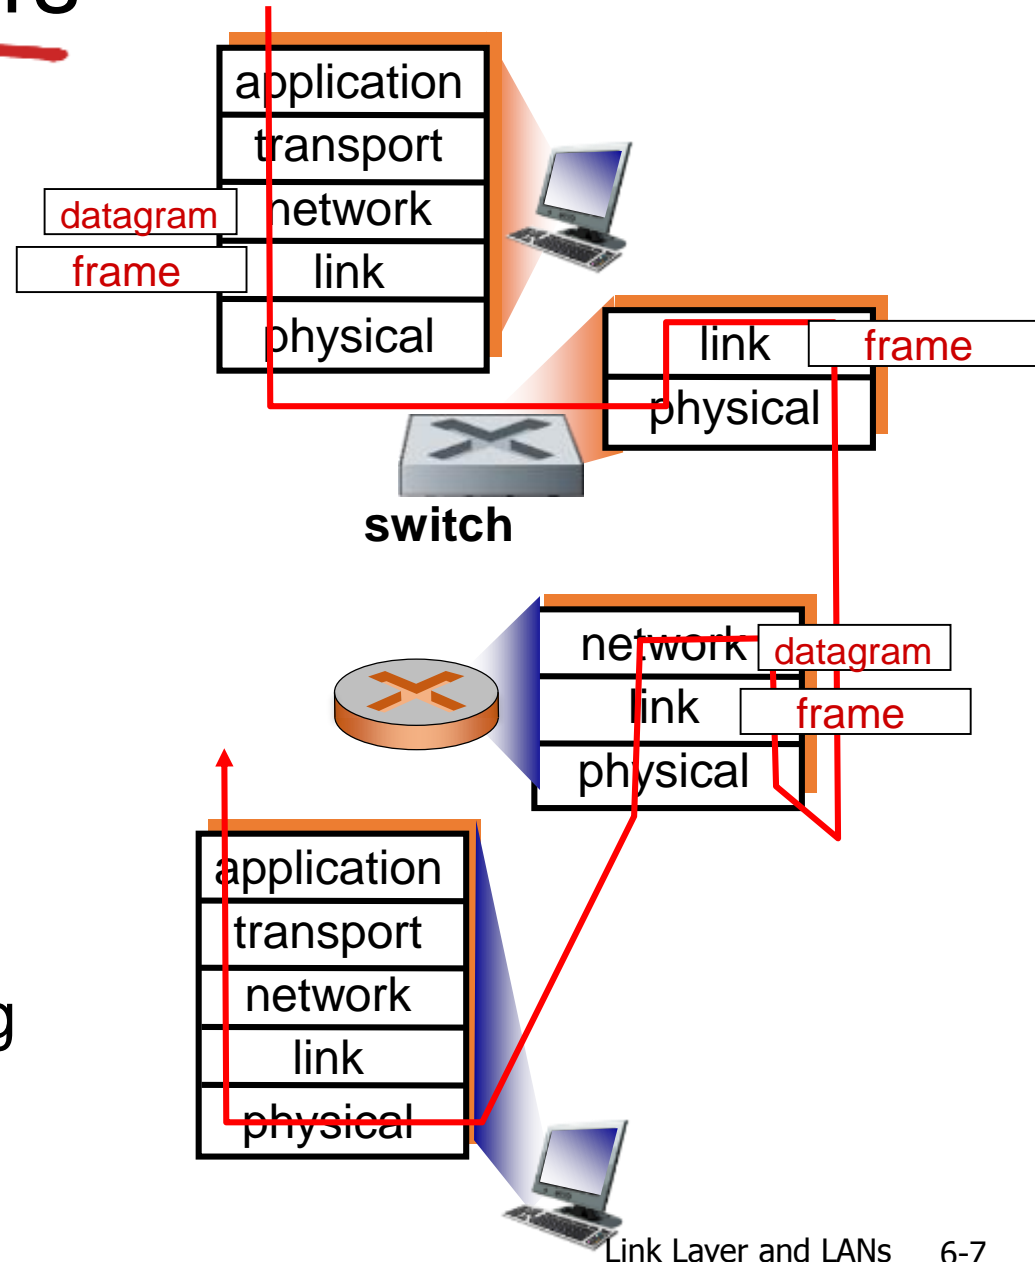
Switches vs. routers

both are store-and-forward:

- **routers:** network-layer devices (examine network-layer headers)
- **switches:** link-layer devices (examine link-layer headers)

both have forwarding tables:

- **routers:** compute tables using routing algorithms, IP addresses
- **switches:** learn forwarding table using flooding, learning, MAC addresses



Data center networks

- 10's to 100's of thousands of hosts, often closely coupled, in close proximity:
 - e-business (e.g. Amazon)
 - content-servers (e.g., YouTube, Akamai, Apple, Microsoft)
 - search engines, data mining (e.g., Google)

- challenges:
 - multiple applications, each serving massive numbers of clients
 - managing/balancing load, avoiding processing, networking, data bottlenecks

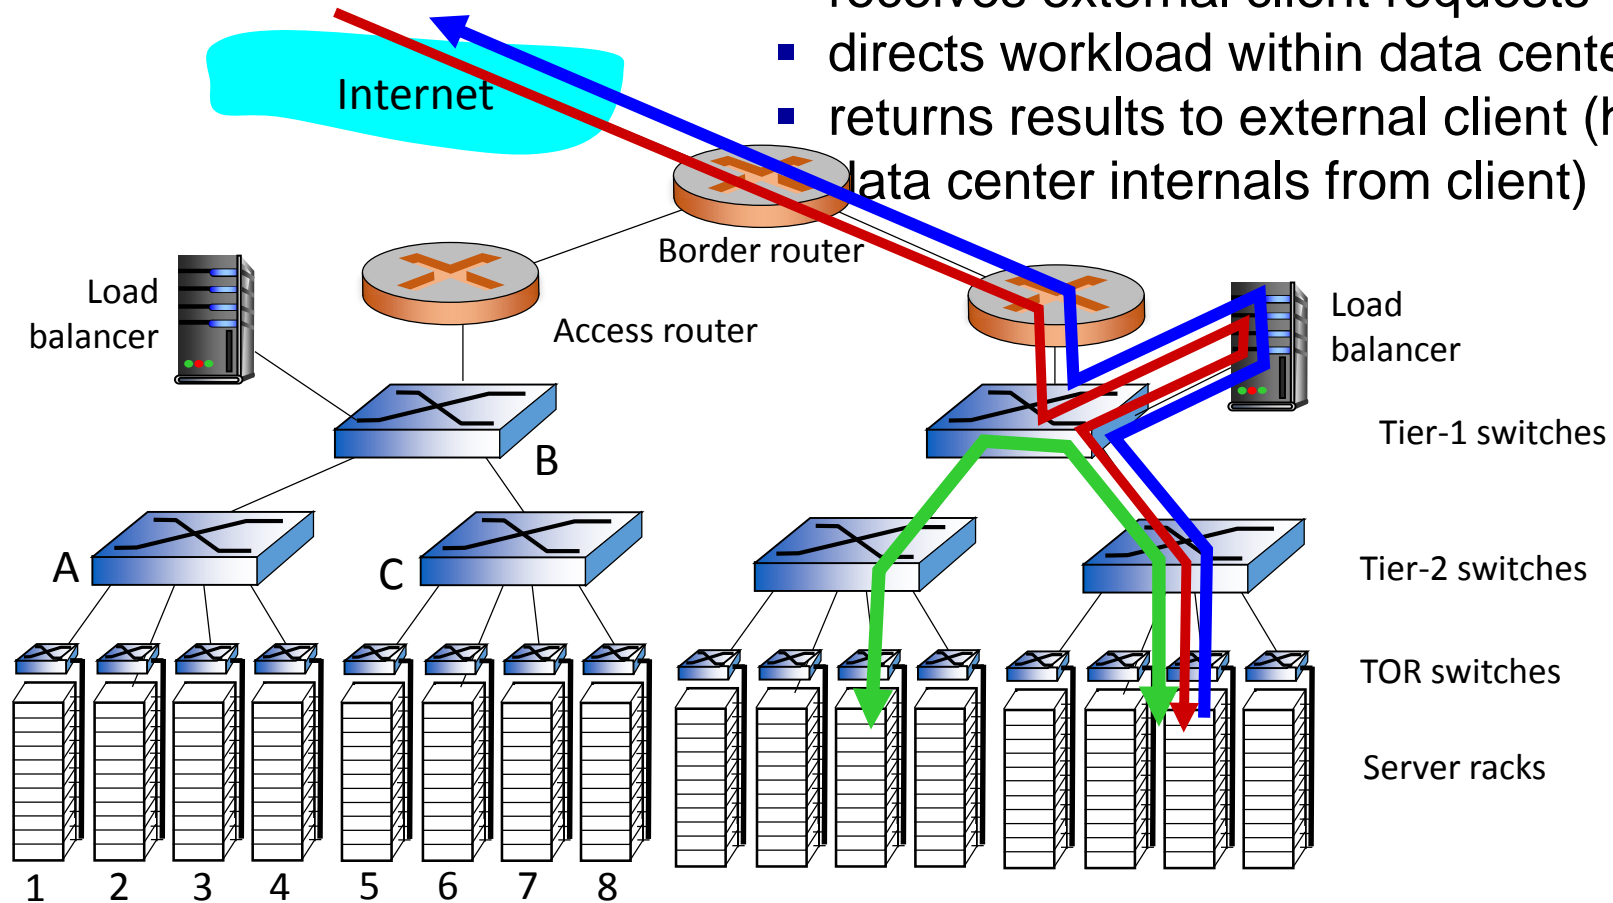


Inside a 40-ft Microsoft container,
Chicago data center

Data center networks


load balancer: application-layer routing

- receives external client requests
- directs workload within data center
- returns results to external client (hiding data center internals from client)



Data center networks

- rich interconnection among switches, racks:
 - increased throughput between racks (multiple routing paths possible)
 - increased reliability via redundancy

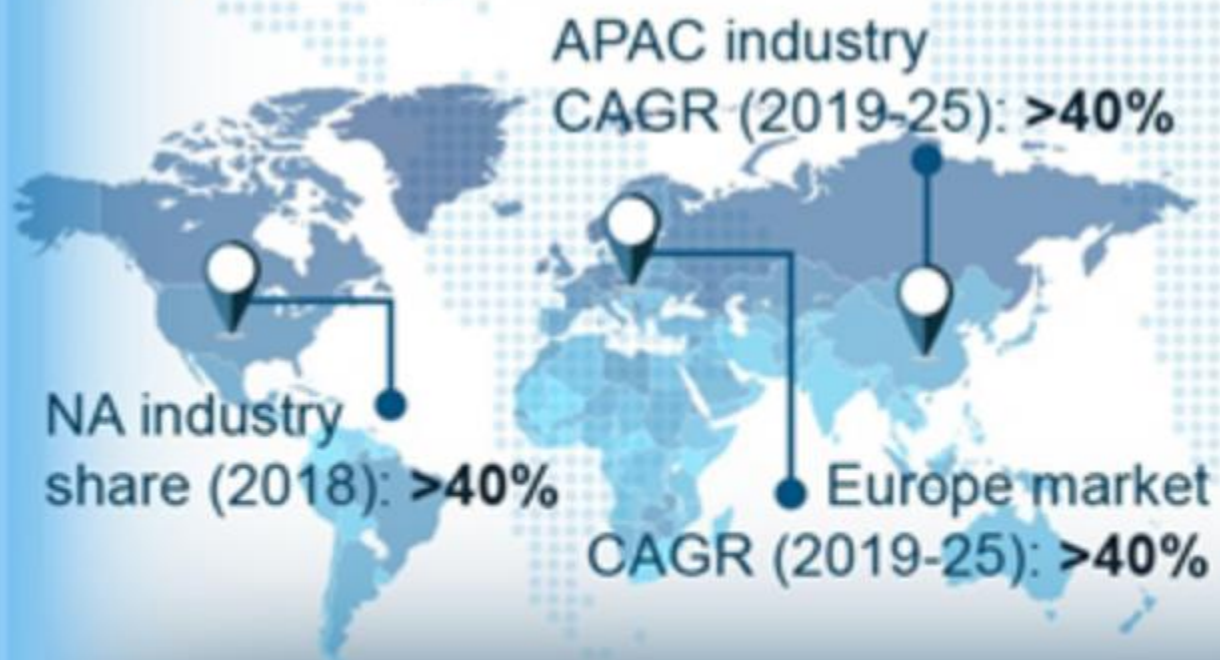


Virtualização
Motivação

SOFTWARE DEFINED NETWORKING (SDN) MARKET

Global Market Insights

Insights to Innovation



END-USE LANDSCAPE



Enterprises end-use sector
share (2018): >45%

Telecom service providers
segment CAGR (2019-25):
>41%

IT-enabled service sector share (2018):
>30%



Software Defined Networking Arc

Find Jobs

Advanced Job Search

job title, keywords or company

Software Defined Networking Architect jobs

[Upload your resume](#) - Let employers find you

Jobs 1 to 10 of 569

Sort by: **relevance** - date

Salary Estimate

- \$90,000+ (517)
- \$100,000+ (411)
- \$110,000+ (308)
- \$115,000+ (243)
- \$125,000+ (135)

Job Type

- Full-time (526)
- Contract (47)
- Part-time (4)
- Temporary (4)
- Commission (3)
- Internship (2)

Location

- San Jose, CA (33)
- Seattle, WA (32)
- Santa Clara, CA (21)
- New York, NY (18)
- Atlanta, GA (14)
- [more »](#)

Company

- Oracle (42)
- Tribridge (35)
- Cisco Systems, Inc. (22)
- Silver Peak Systems (17)

Sr Mgr Systems Engineering

NeuStar - ★★★★★ 23 reviews - Sterling, VA

Deep understanding of Software Defined Services. Networking, Storage, Security. Lead a small group of very talented Engineers and Architects....

16 days ago - [save job](#) - [email](#) - [more...](#)

Dir, Network Engineering

CDK Global - ★★★★★ 70 reviews - Hoffman Estates, IL 60169

Experience in leading technical teams and/or transformation programs related to networking technologies. Provide technical, hands-on management to drive complex...

27 days ago - [save job](#) - [email](#) - [more...](#)

Senior Staff Engineer-Software Defined Mobile Networks

Huawei - ★★★★★ 221 reviews - Santa Clara, CA

Good knowledge of network function virtualization (NFV), software defined networking (SDN) controllers and applications, OpenFlow, and virtualization....

3 days ago - [save job](#) - [email](#) - [more...](#)

Network Engineer - Cloud Network Services

GE Digital - ★★★★★ 7 reviews - Ohio

Keep up to date on developments with cloud-hosted infrastructure and platform services, software defined networking (SDN), and related connectivity solutions....

GE Careers - 23 days ago - [save job](#) - [email](#) - [more...](#)

IT Infrastructure Solutions Architect (Experienced)

Sandia National Laboratories - ★★★★★ 86 reviews - Livermore, CA

Get new jobs for this search by email

My email:

Also get an email with jobs recommended just for me

Activate

You can cancel email alerts at any time.



Analista De Redes Multicamadas Pleno Sãnior

Padtec

Minas Gerais

via IT Jobs UK

 há 10 dias  Tempo integral



Desenvolvedor(a) de Software Sênior - SDN e NFV

globo.com

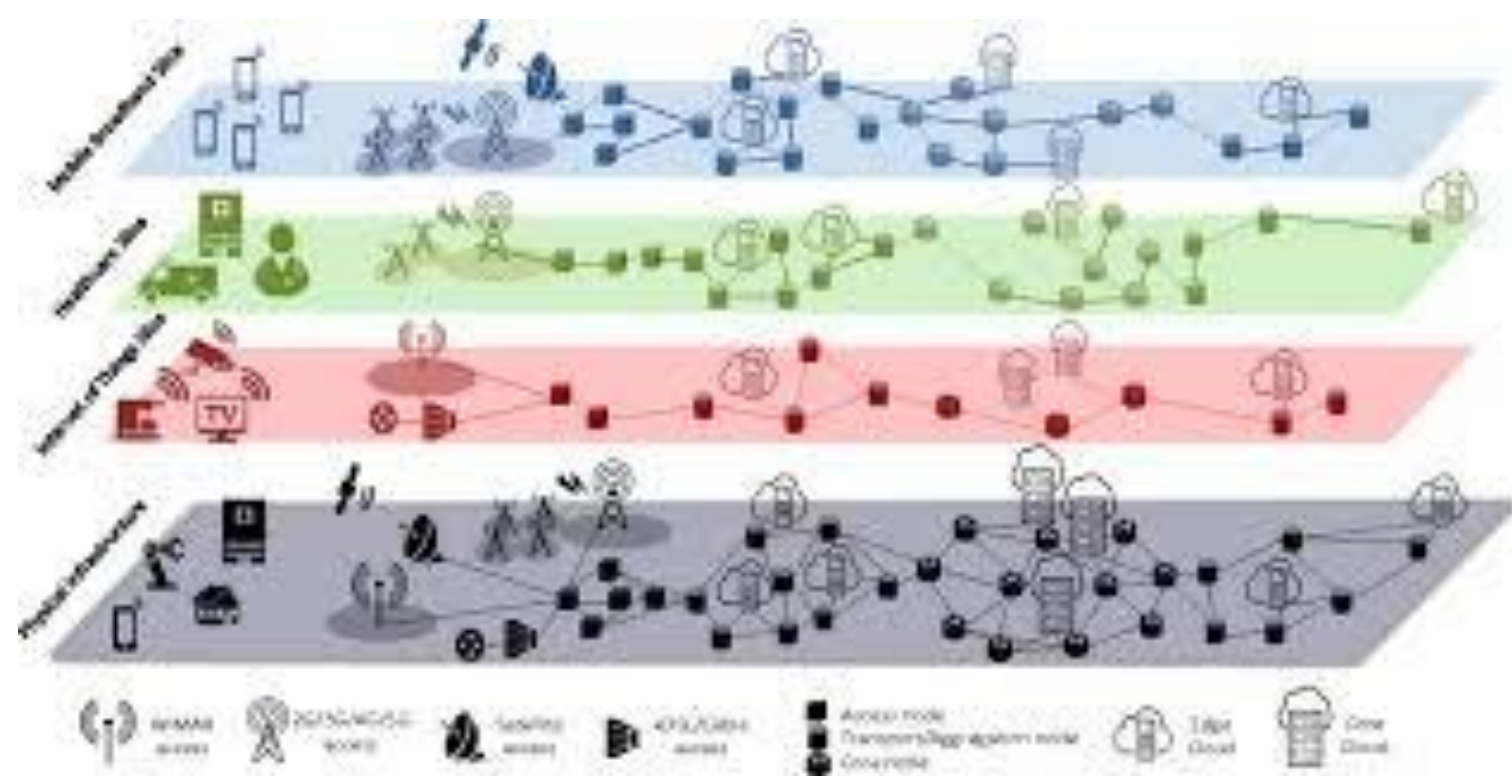
Rio de Janeiro, RJ

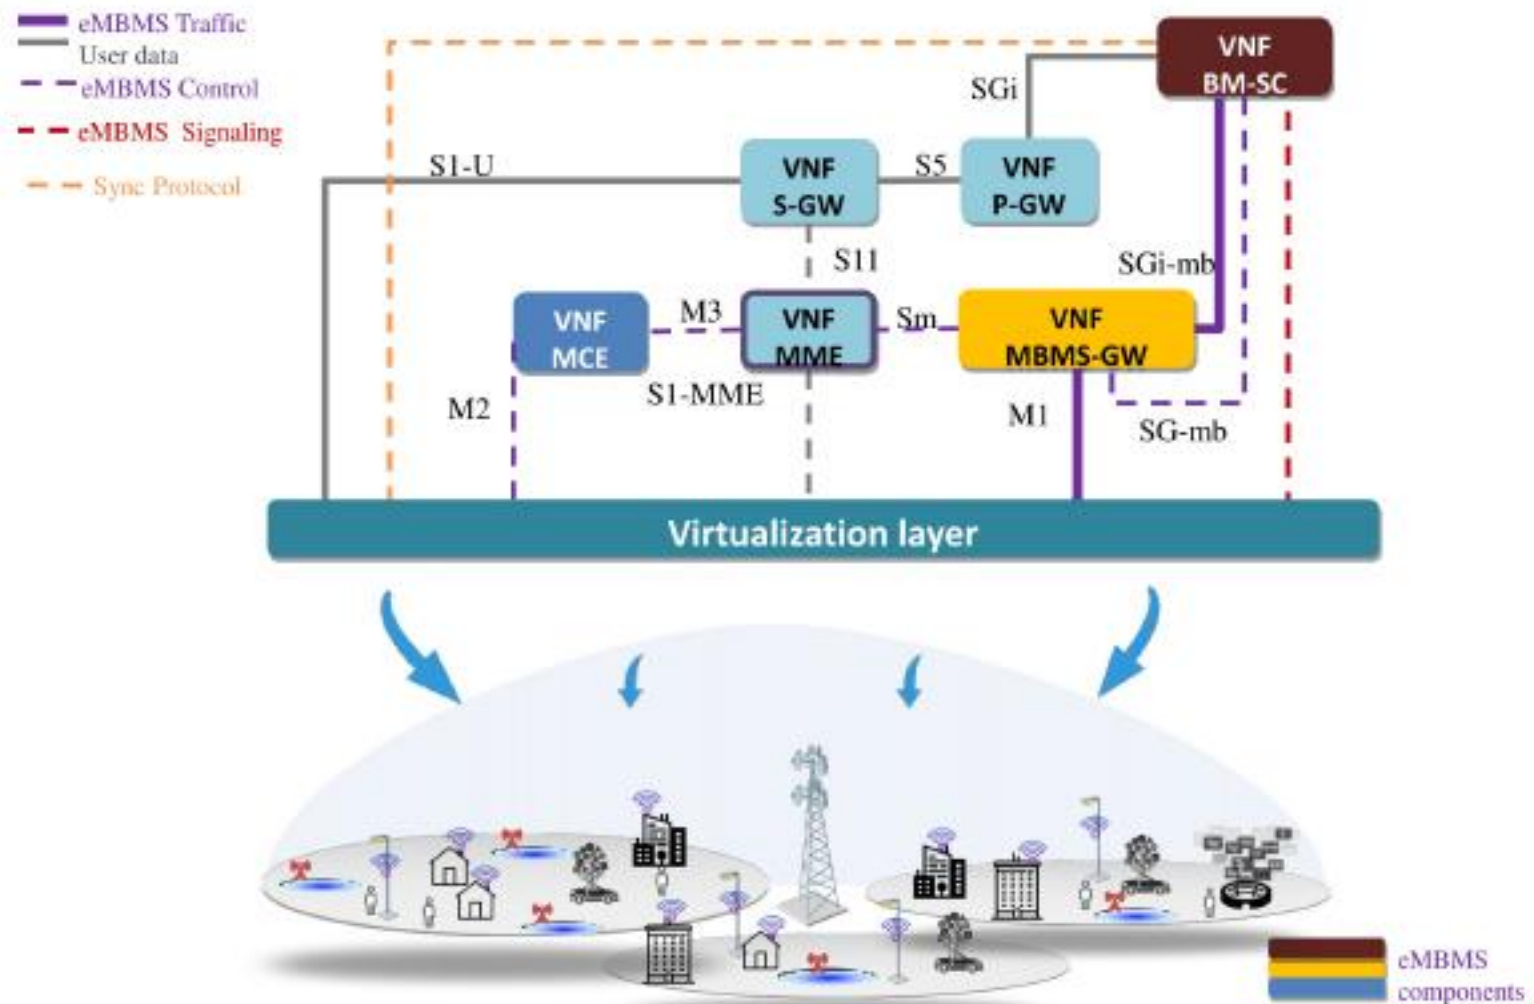
via Recruit.net

 há 4 dias  Tempo integral



[→ Ver vagas](#)



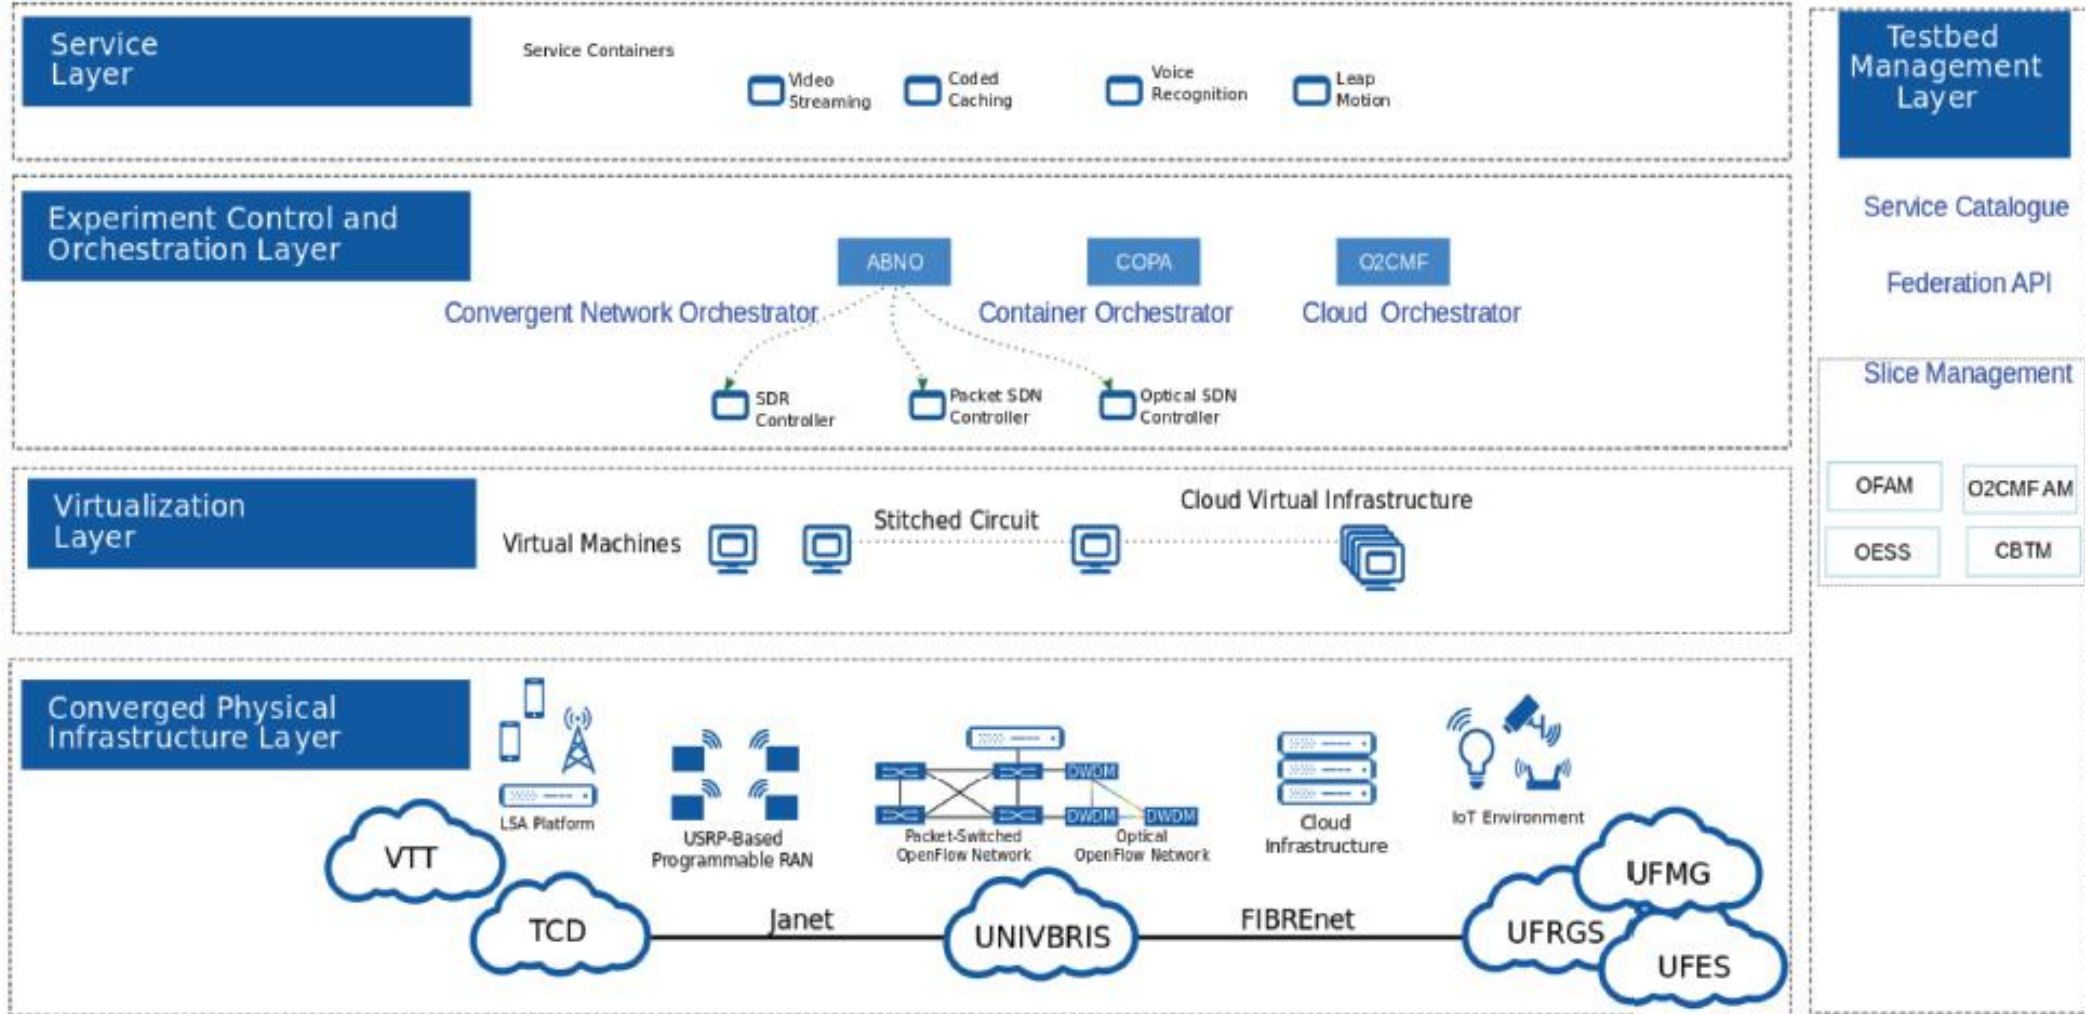




Um exemplo no IC
IoT Virtualizada



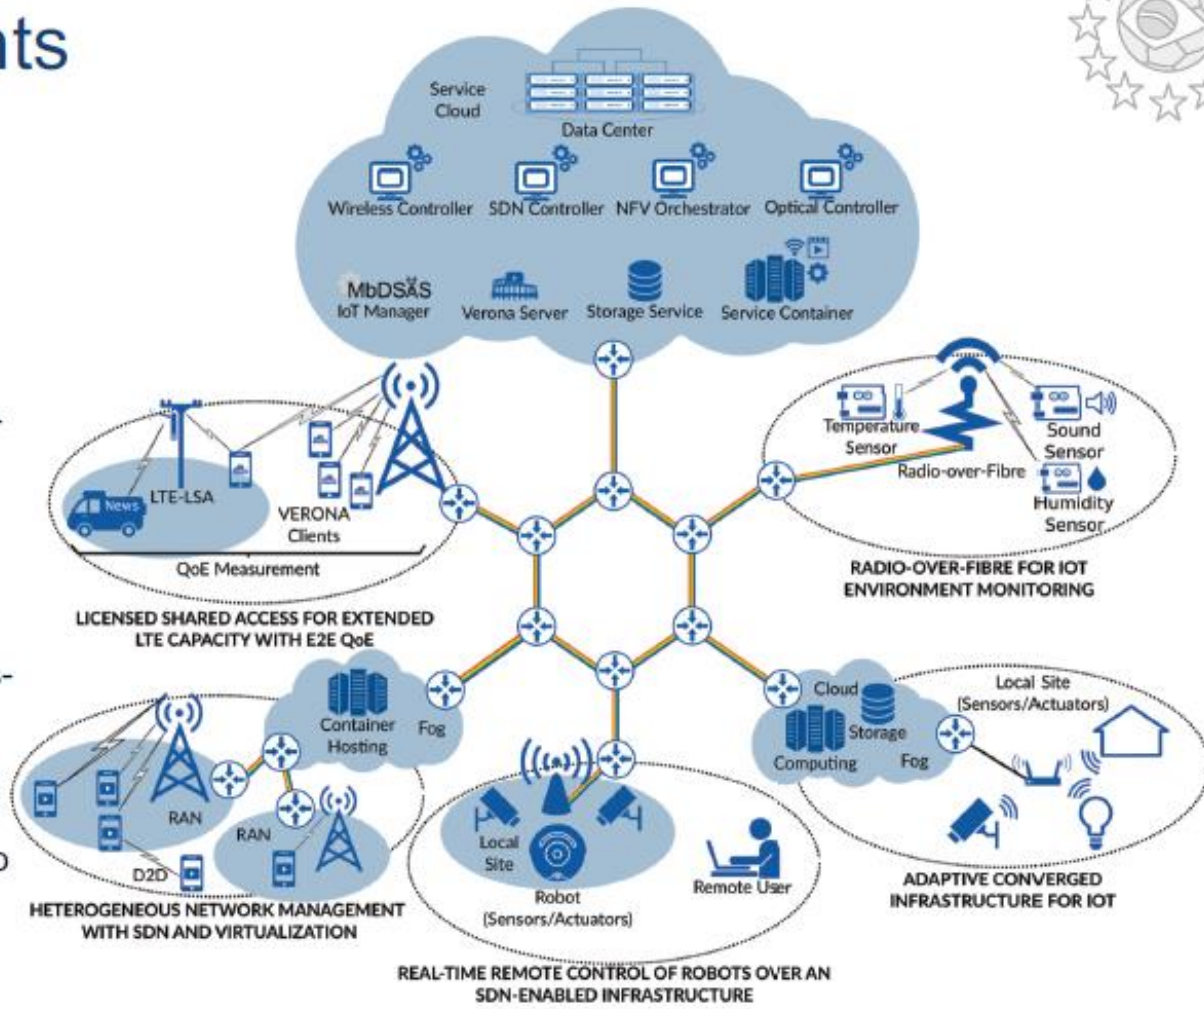
Architectural View of Futebol





Use cases and experiments

- UC1: The impact of broadband wireless on optical backhauling
 - Licensed Shared Access for extended LTE capacity with shared optical backhauling and end-to-end QoE
- UC2: The design of SDN infrastructure for wireless-optical integration
 - Heterogeneous wireless-optical network management with SDN and virtualization
 - Real-time remote control of robots over a wireless-optical SDN infrastructure
- UC3: The interplay between wireless and optical networks for IoT
 - Adaptive cloud/fog computing for IoT, according to network capacity and service latency requirements
 - Radio-over-fiber for IoT environment monitoring





FUTEBOL

Federated Union of Telecommunications
Research

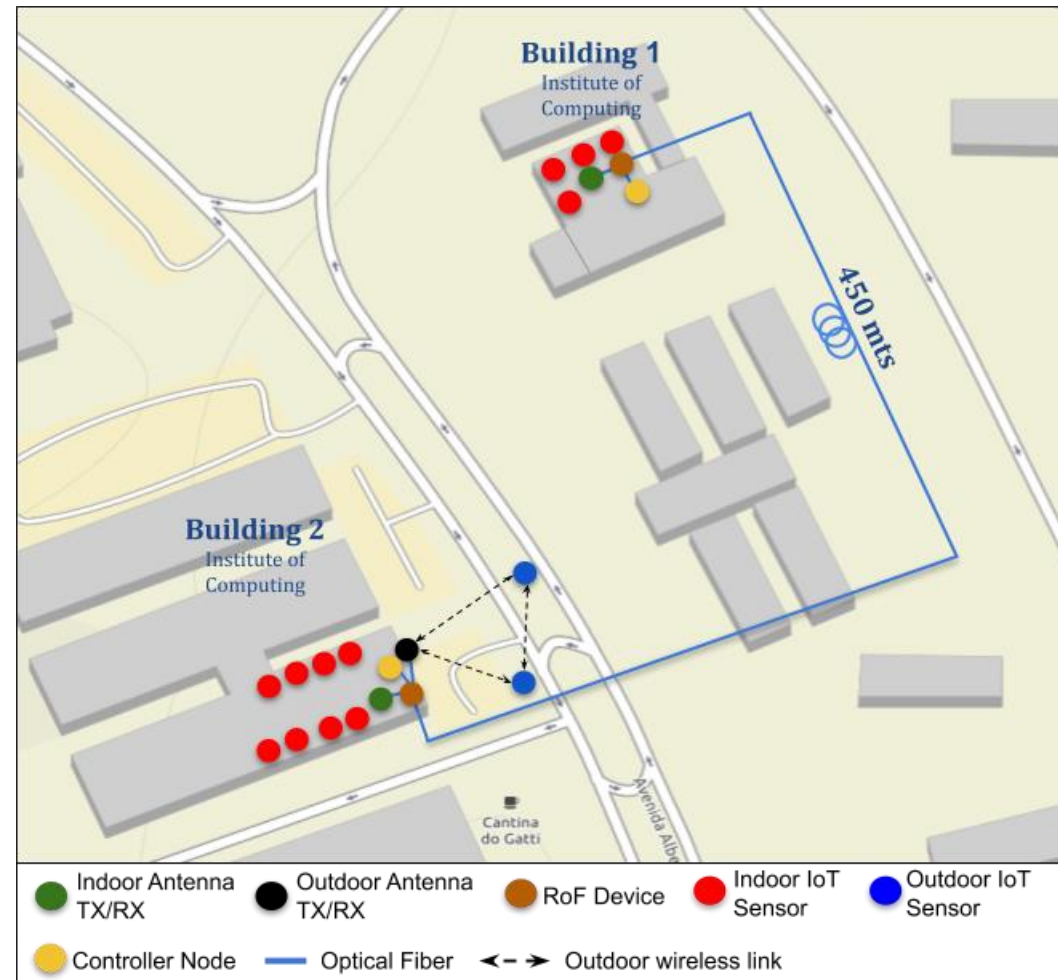
Facilities for an EU-Brazil Open Laboratory

Radio over Fiber for IoT Environment Monitoring

UNICAMP and UFRGS

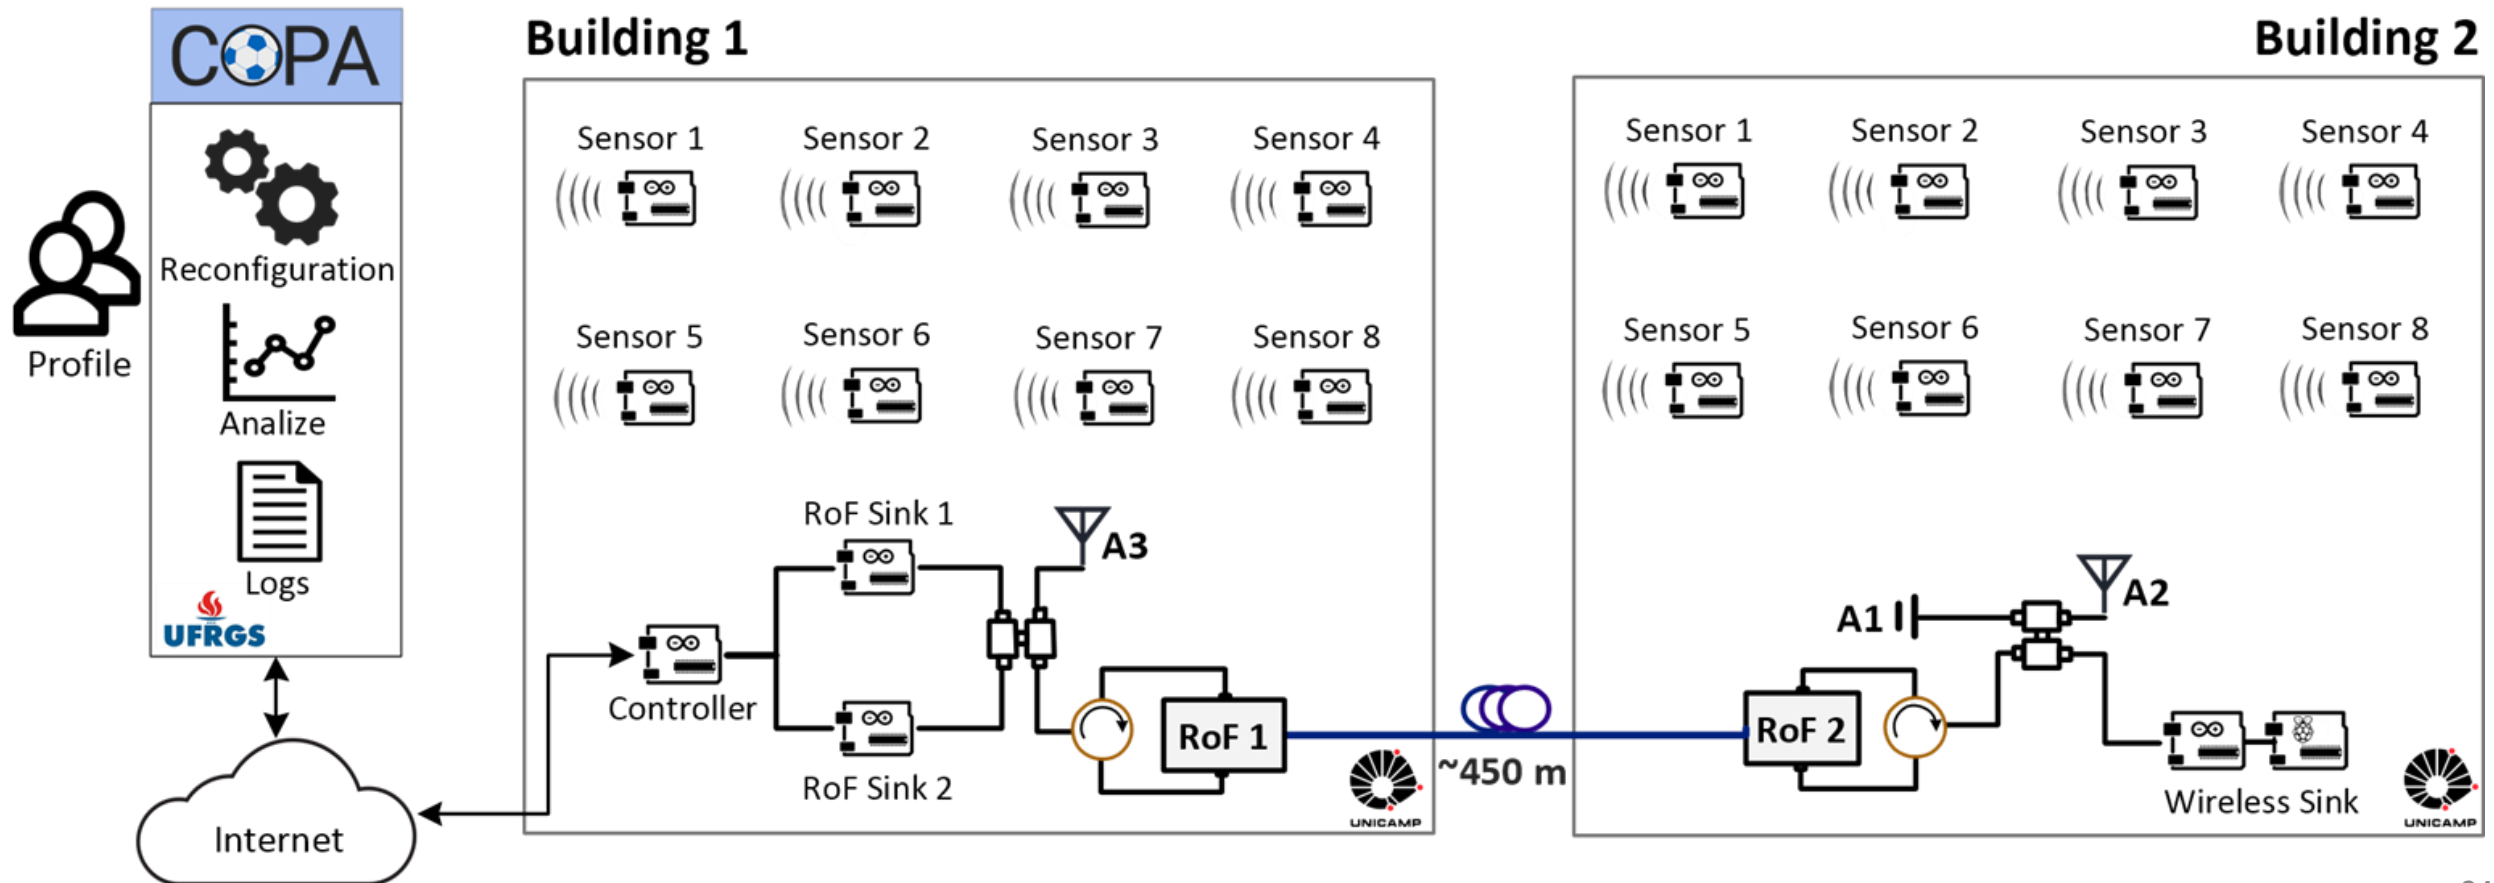
DEMO Setup

- Geographical distribution



DEMO Setup

- Specific experimental setup for extended-coverage architecture



Virtualização




Virtualization

“Virtualization means to create a virtual version of a device or resource, such as a server, storage device, network or even an operating system where the framework divides the resource into one or more execution environments. Devices, applications and human users are able to interact with the virtual resource as if it were a real single logical resource.”

<http://www.webopedia.com/TERM/V/virtualization.html>

Virtualization - Features

- Sharing of resources
- Isolation
- Agregation




Virtualization - advantages

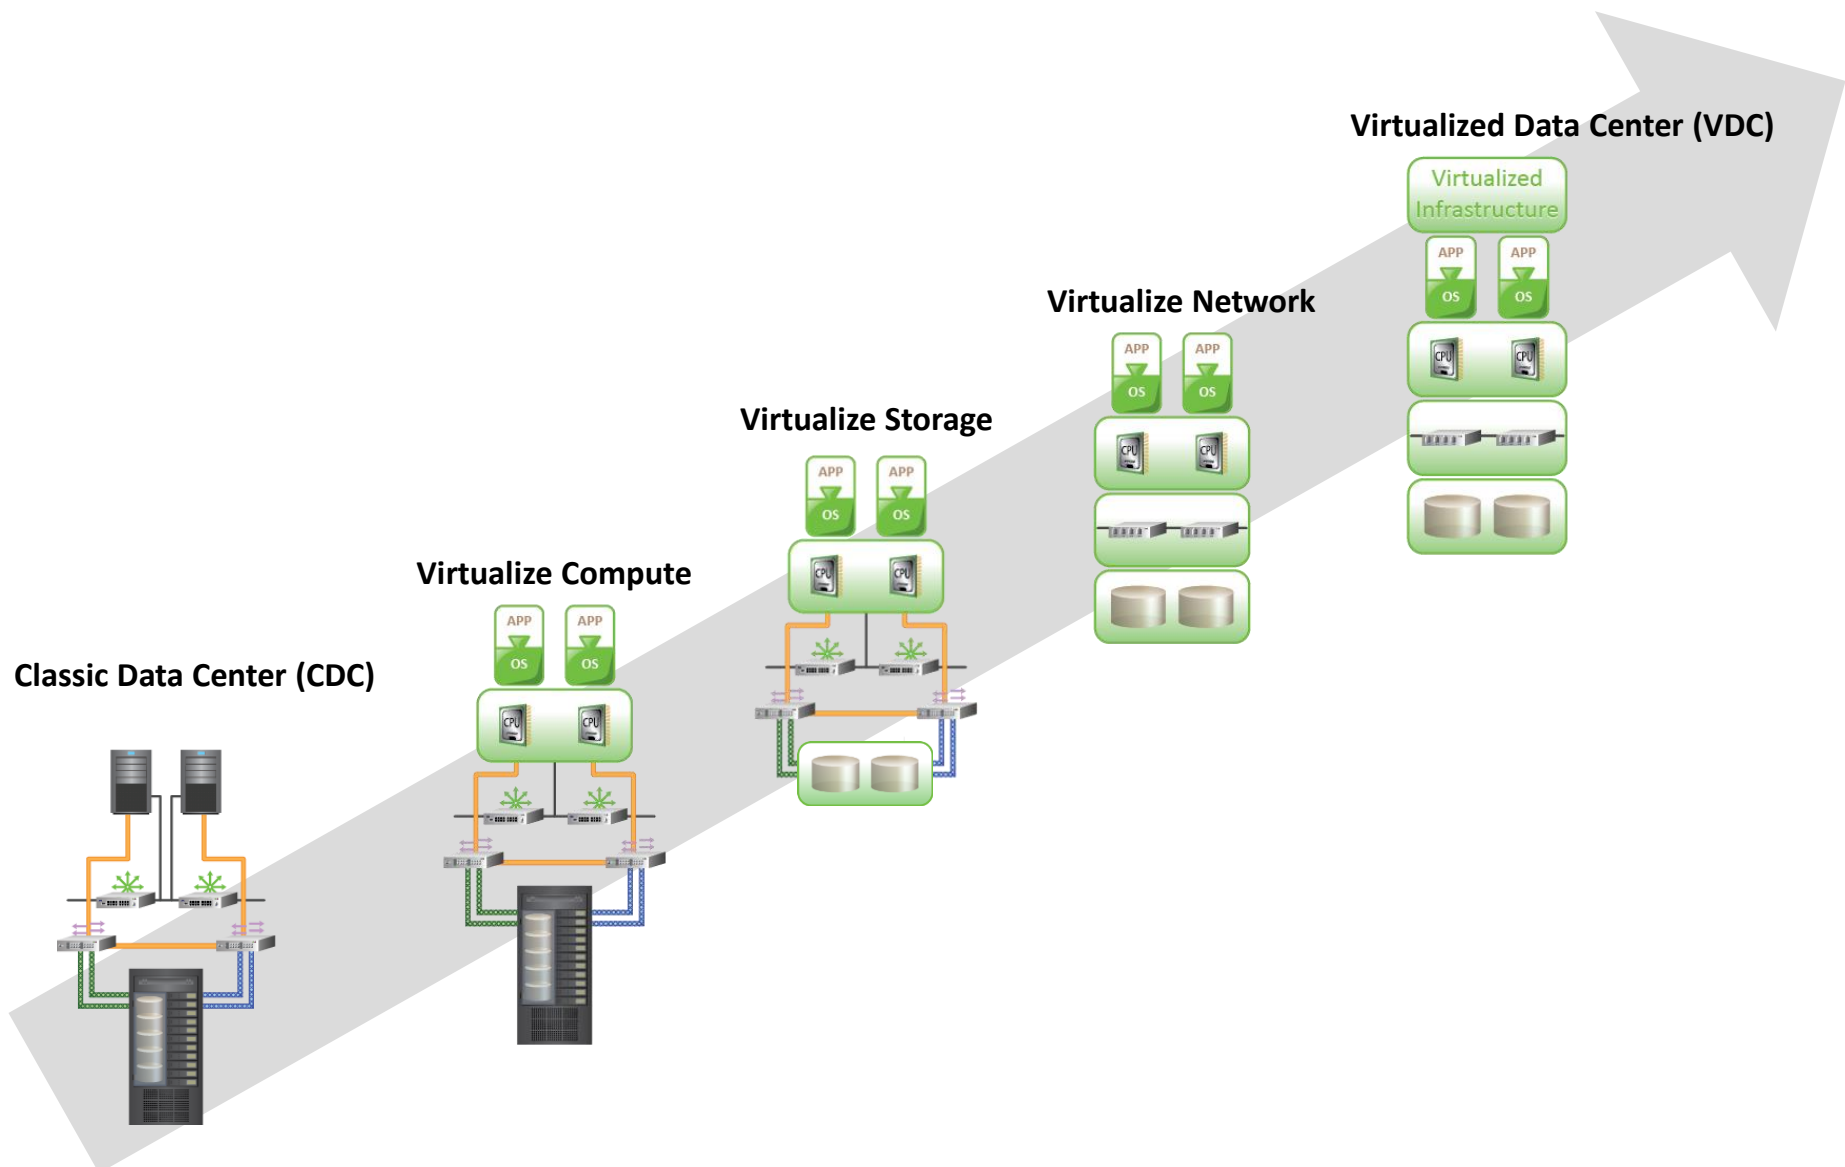
- Fast deployment
- Scalability
- Load consolidation
- Flexibility
- Mobility
- Green



Agenda:

IBM Mainframe Technology Evolution
Over 40 Years of Relentless Innovation and Refinement





Classic Data Center (CDC)

Virtualize Compute

Virtualize Storage

Virtualize Network

Virtualized Data Center (VDC)

Virtualized Infrastructure


APP OS APP OS

CPU CPU



Network Switch

Storage

Virtual Machine Hardware



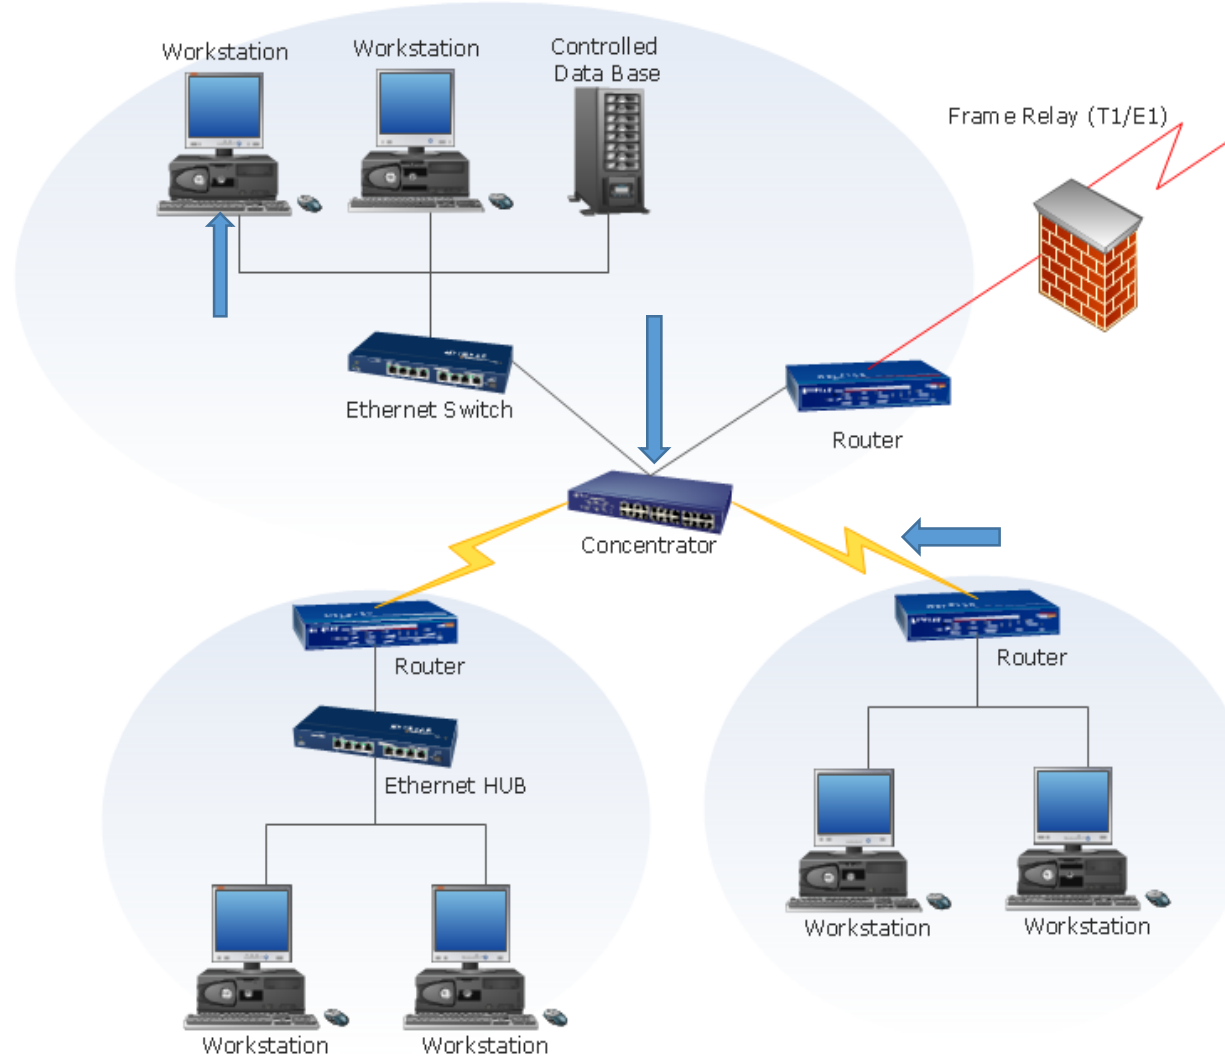
Virtual Machines x Containers



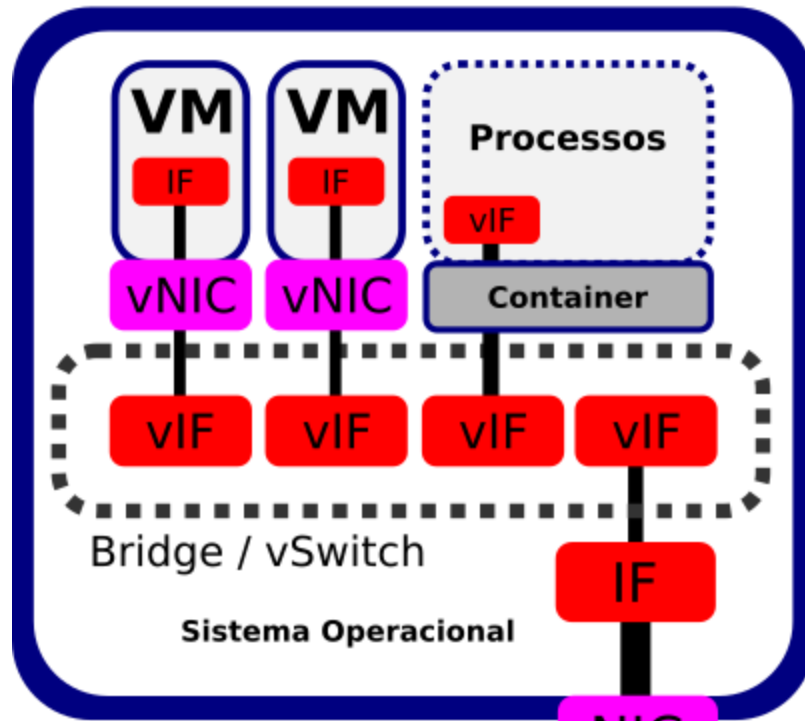
Network Virtualization

Nelson L. S. da Fonseca
IEEE ComSoc Summer School
Trento, July 9th, 2015

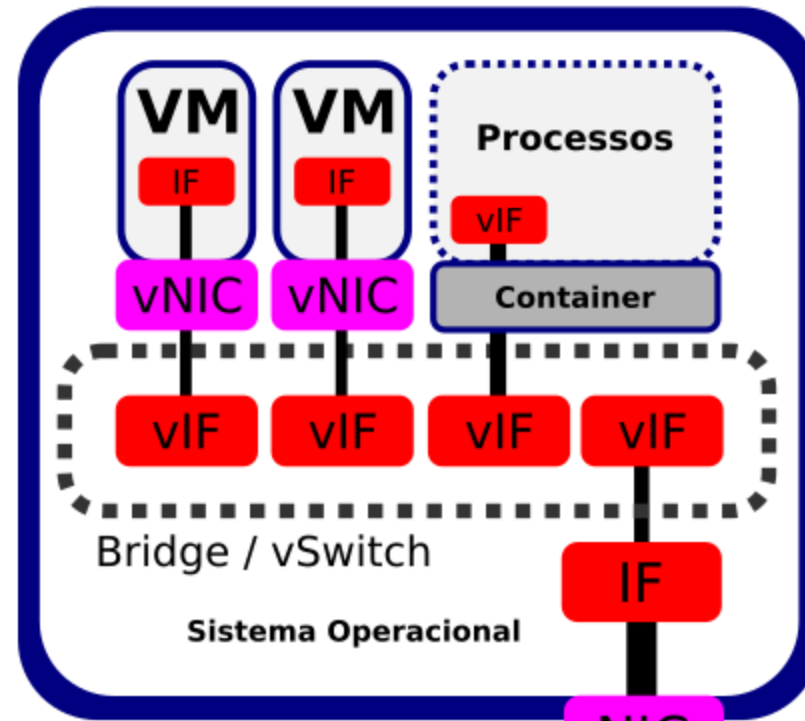
Network Virtualization



Servidor



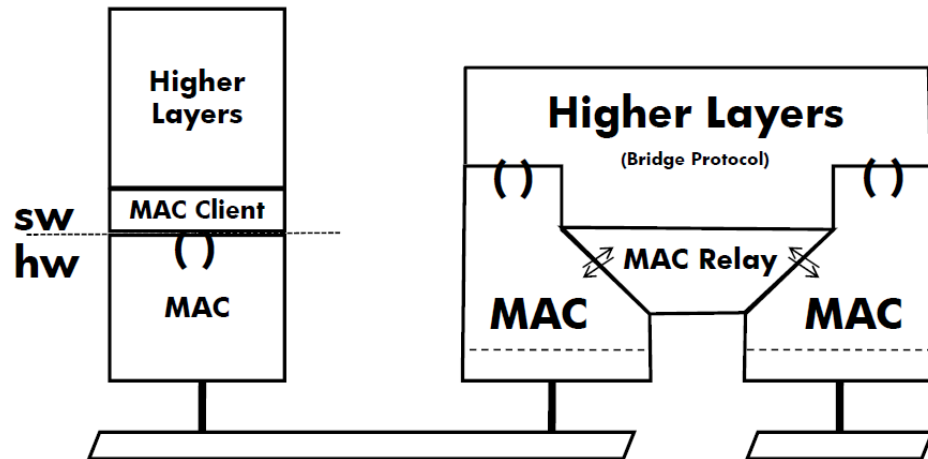
Servidor



Networking


IEEE 802

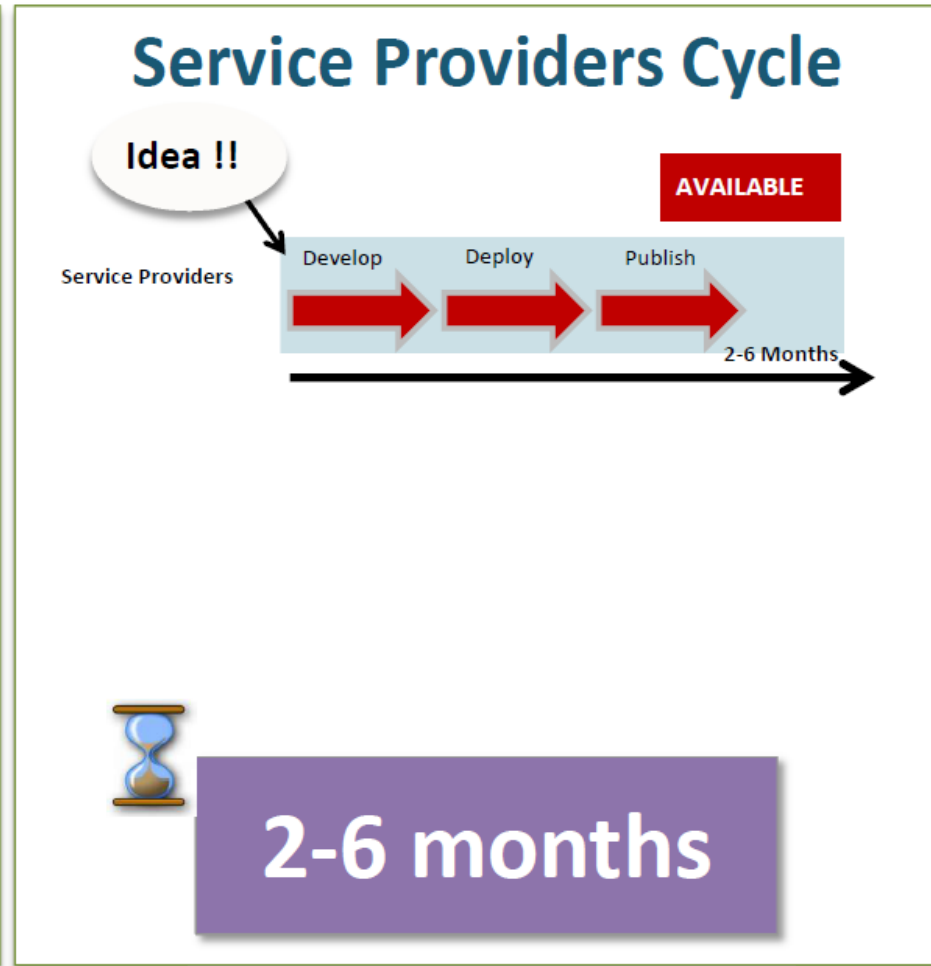
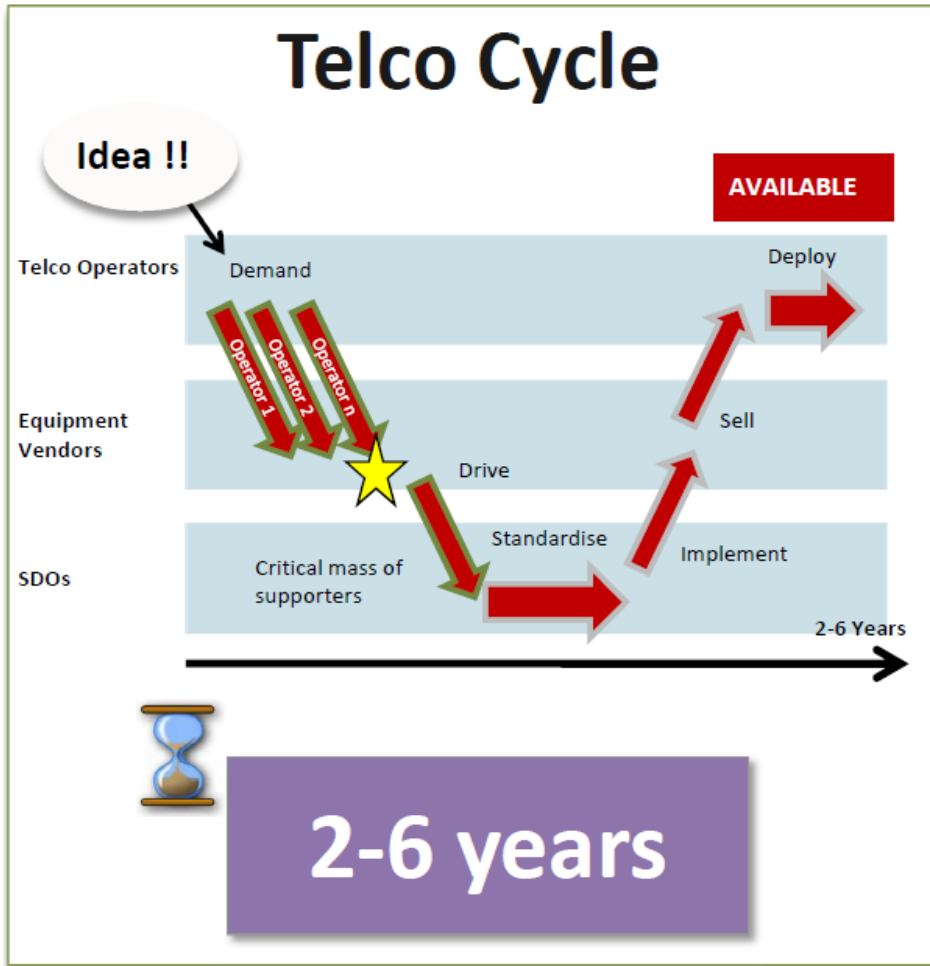
Traditional Networking The end-station and bridge



IEEE 802

Modern Networking The end-station and bridge





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

Multitenancy


Multitenancy is the fundamental technology that clouds use to share IT resources cost-efficiently and securely. Just like in an apartment building in which many tenants cost-efficiently share the common infrastructure of the building but have walls and doors that give them privacy from other tenants - a cloud uses multitenancy technology to share IT resources securely among multiple applications and tenants (businesses, organizations) that use the cloud.

Network Virtualization techniques

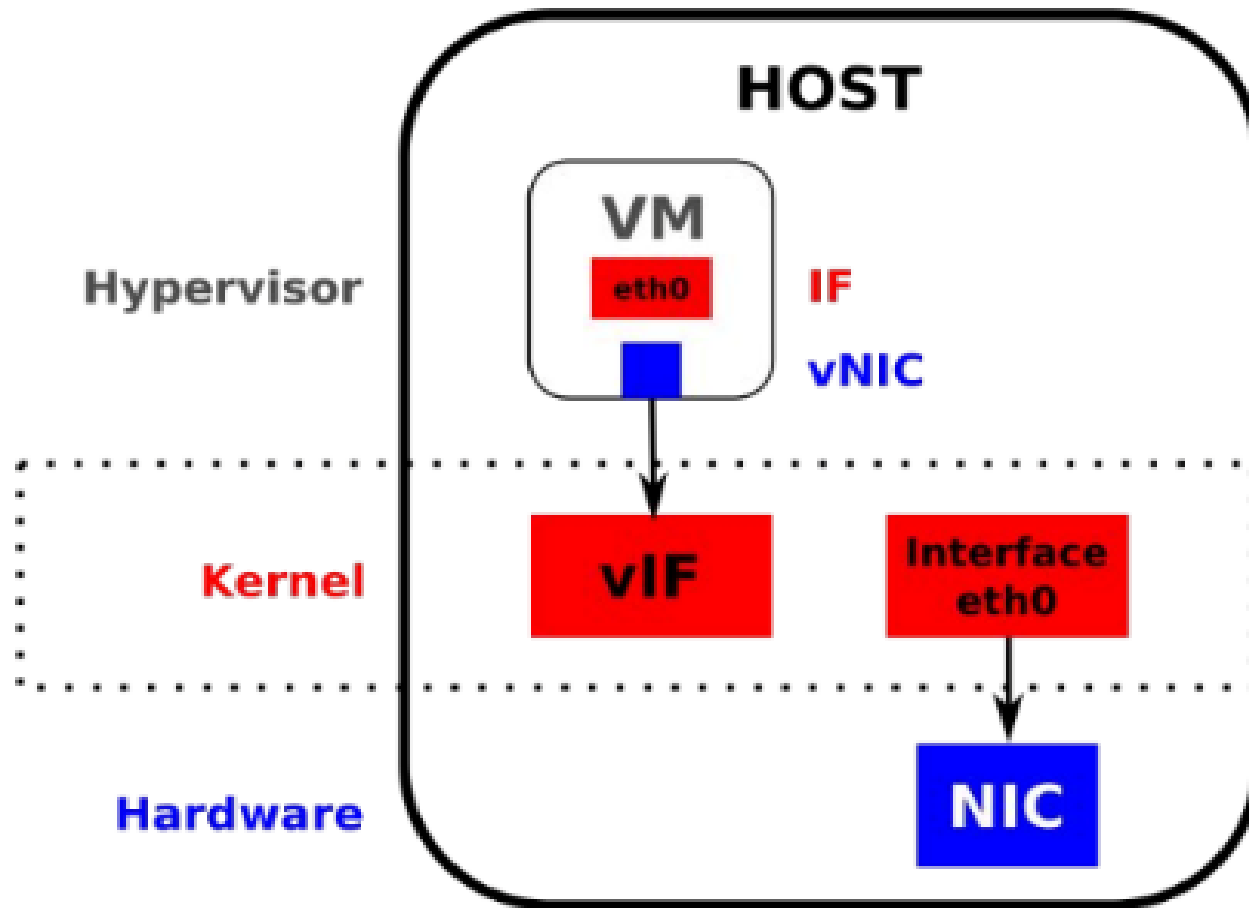
	Technique
NIC	SR-IOV, MR-IOV
Switch	VEB, VEPA, VSS, VBE, DVS, FEX
L2 Link	VLAN
L2 network using L2	VLAN
L2 network using L3	NVO3, VXLAN, NVGRE, STT, TRILL, LISP
Router	VRF, VRRP
L3 network using L3	MPLS, GRE, IPSec

NIC Virtualization

SR-IOV




- *Single Root IOV*
- SR-IOV is a specification that allows a PCIe device to appear to be multiple separate physical PCIe devices.
- With SR-IOV, a card that's SR-IOV-capable has the intelligence to manage the virtual connections so the hypervisor doesn't have to, which means you get a few cycles back in your CPU for other things because it's now offloaded to the card.



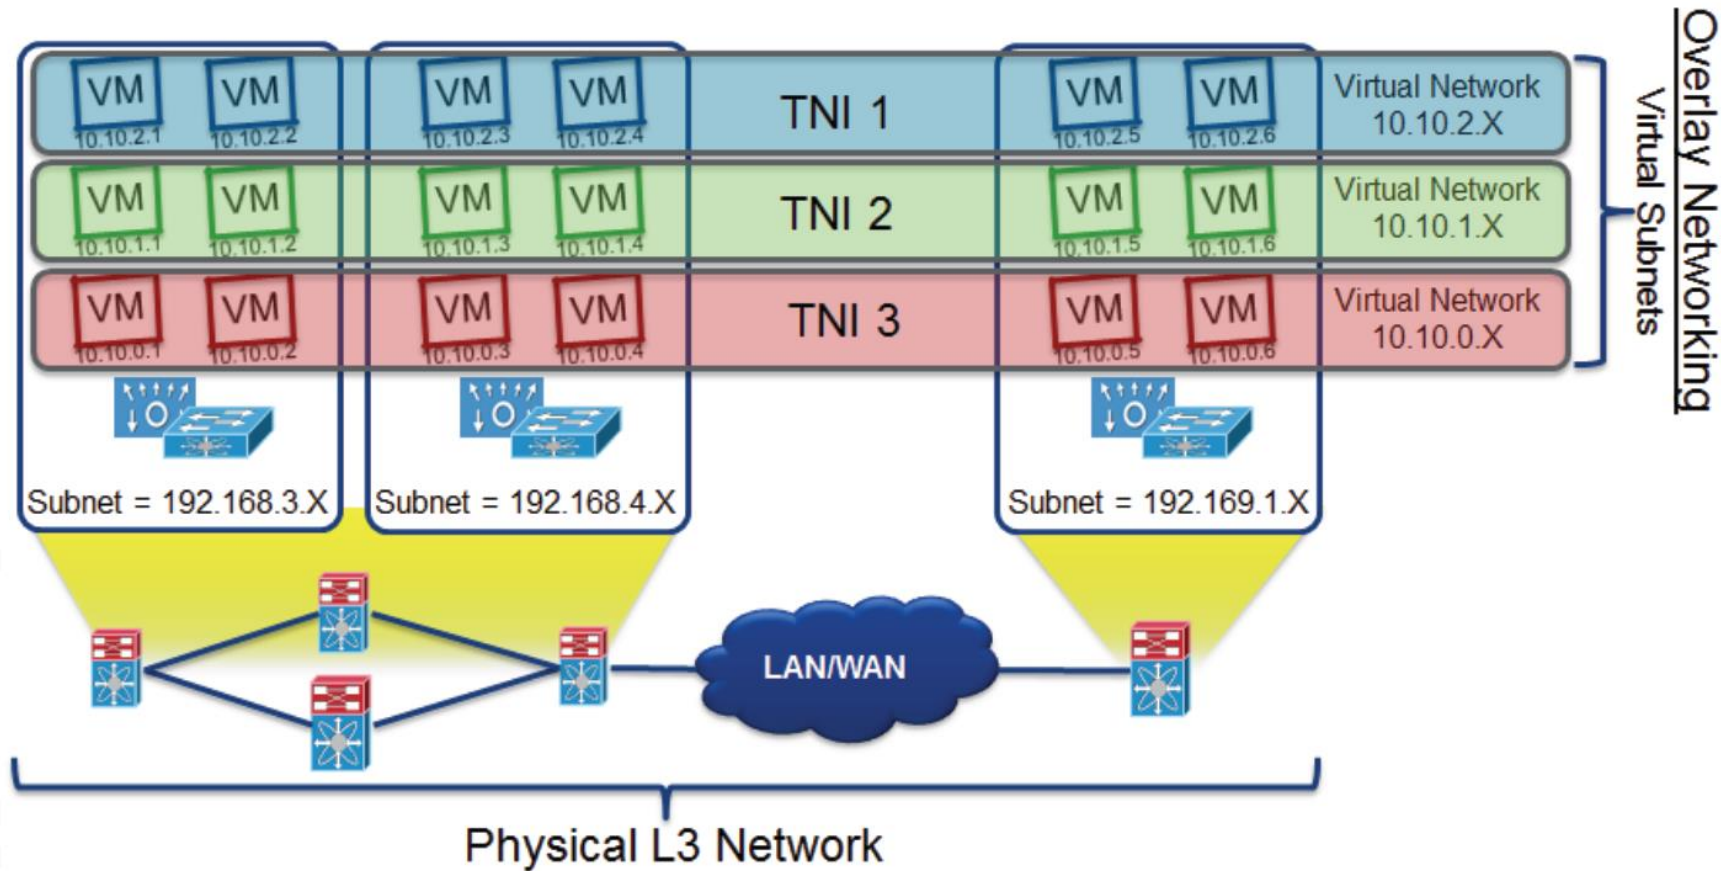
Link Virtualization

Link Aggregation Control Protocol




- IEEE 802.3ad
- Link Aggregation Control Protocol (LACP) provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP)

Network Virtualization using Generic Routing Encapsulation (NVGRE)




Switch Virtualization

vSwitch



- Allows multiple virtual machine to be connected to a physical NIC.
- The vNICs of VMs are connected to a vSwitch
- Hypervisor creates multiplex vNICs, pNIC is controlled by the Hypervisor

Open vSwitch

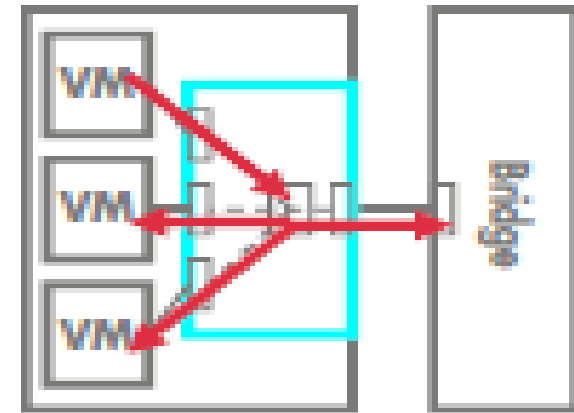


- "Open vSwitch is a production quality, multilayer virtual switch licensed under the open source [Apache 2.0](https://www.apache.org/licenses/LICENSE-2.0) license. It is designed to enable massive network automation through programmatic extension, while still supporting standard management interfaces and protocols (e.g. NetFlow, sFlow, IPFIX, RSPAN, CLI, LACP, 802.1ag). In addition, it is designed to support distribution across multiple physical servers."

<http://openvswitch.org/>

Virtual Ethernet Bridge (VEB)

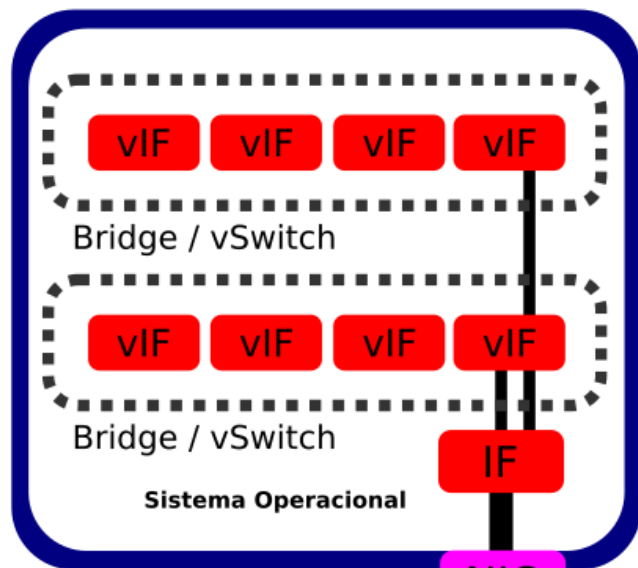
- IEEE 802.1Qbg-2012 standard for vSwitch
- Emulates 802.1 bridges,
- switch internally
- Either in hypervisor or NIC
- Works with all bridges
- Limited bridge visibility
- No changes, legacy solution



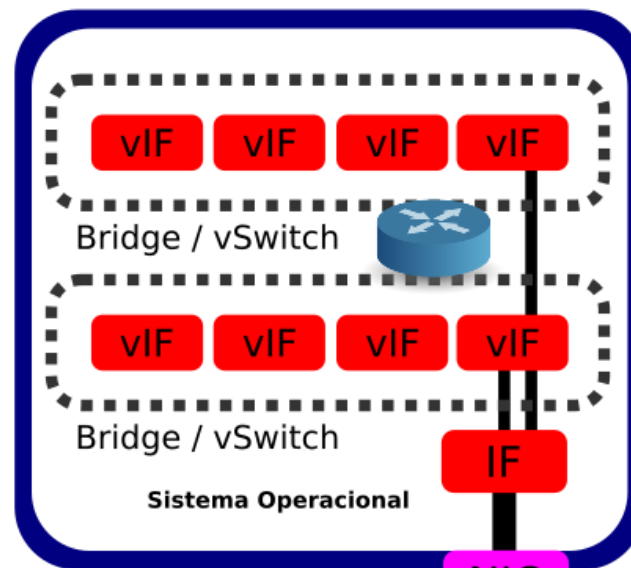
Virtual Ethernet Bridge
(VEB)



Servidor

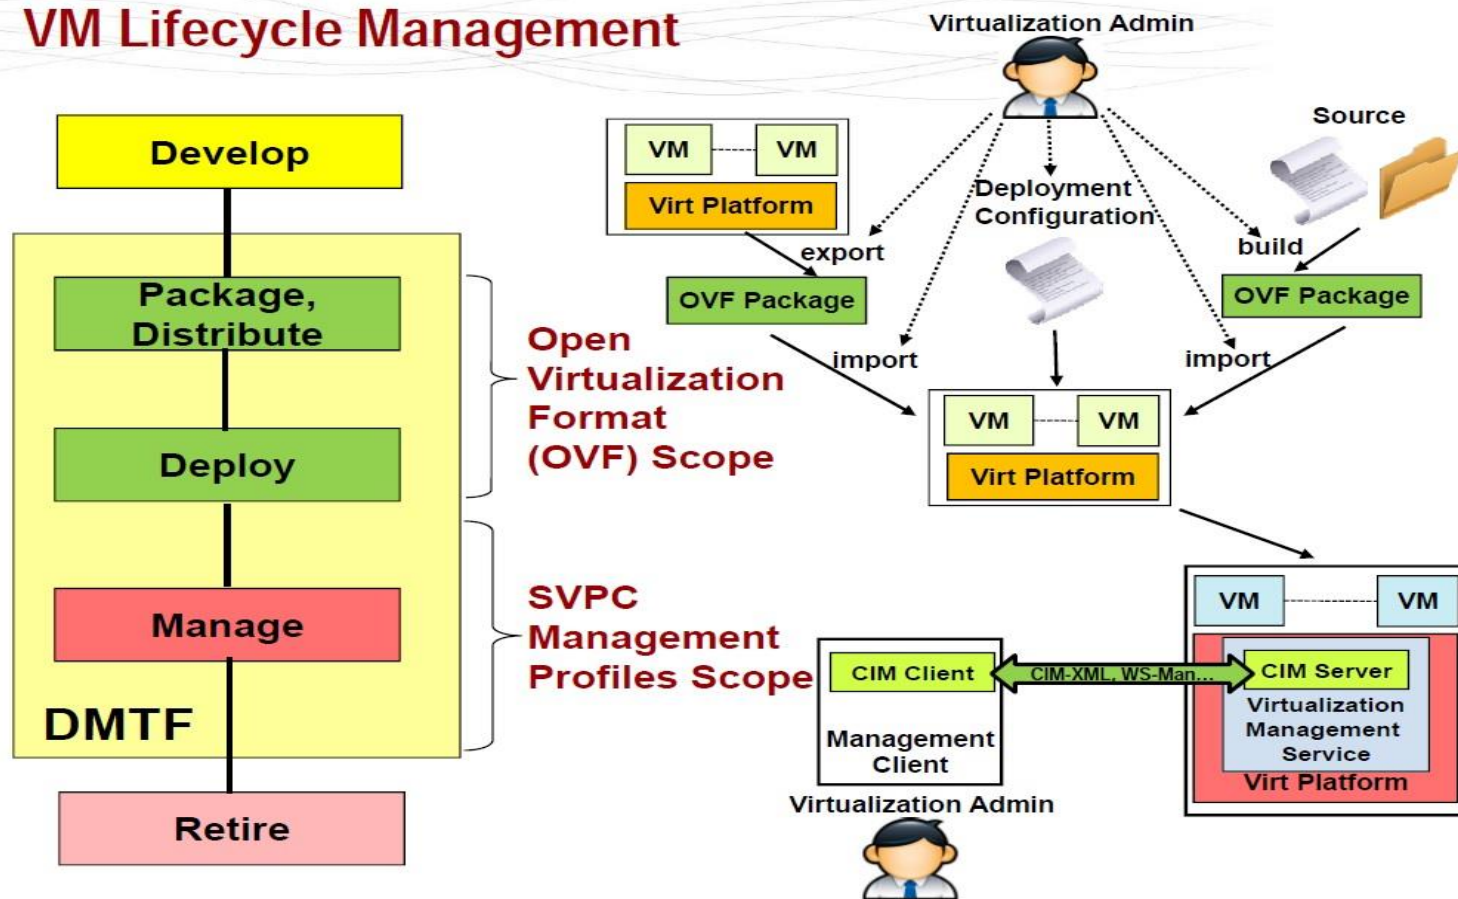


Servidor



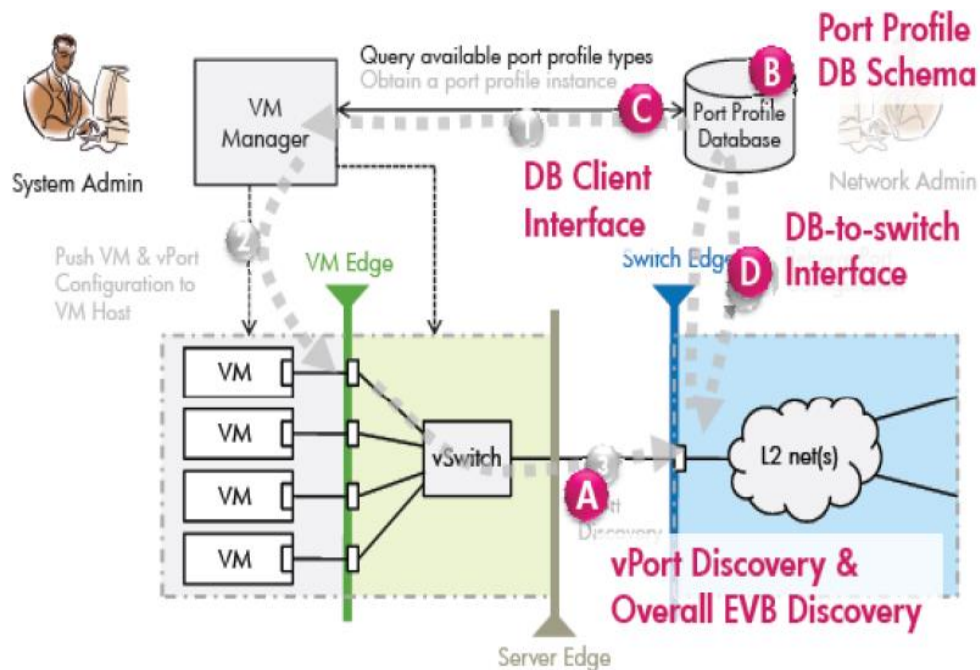
VM Lifecycle

VM Lifecycle Management



Network Port Profile

- Set of attributes that can be applied to one or more virtual machine



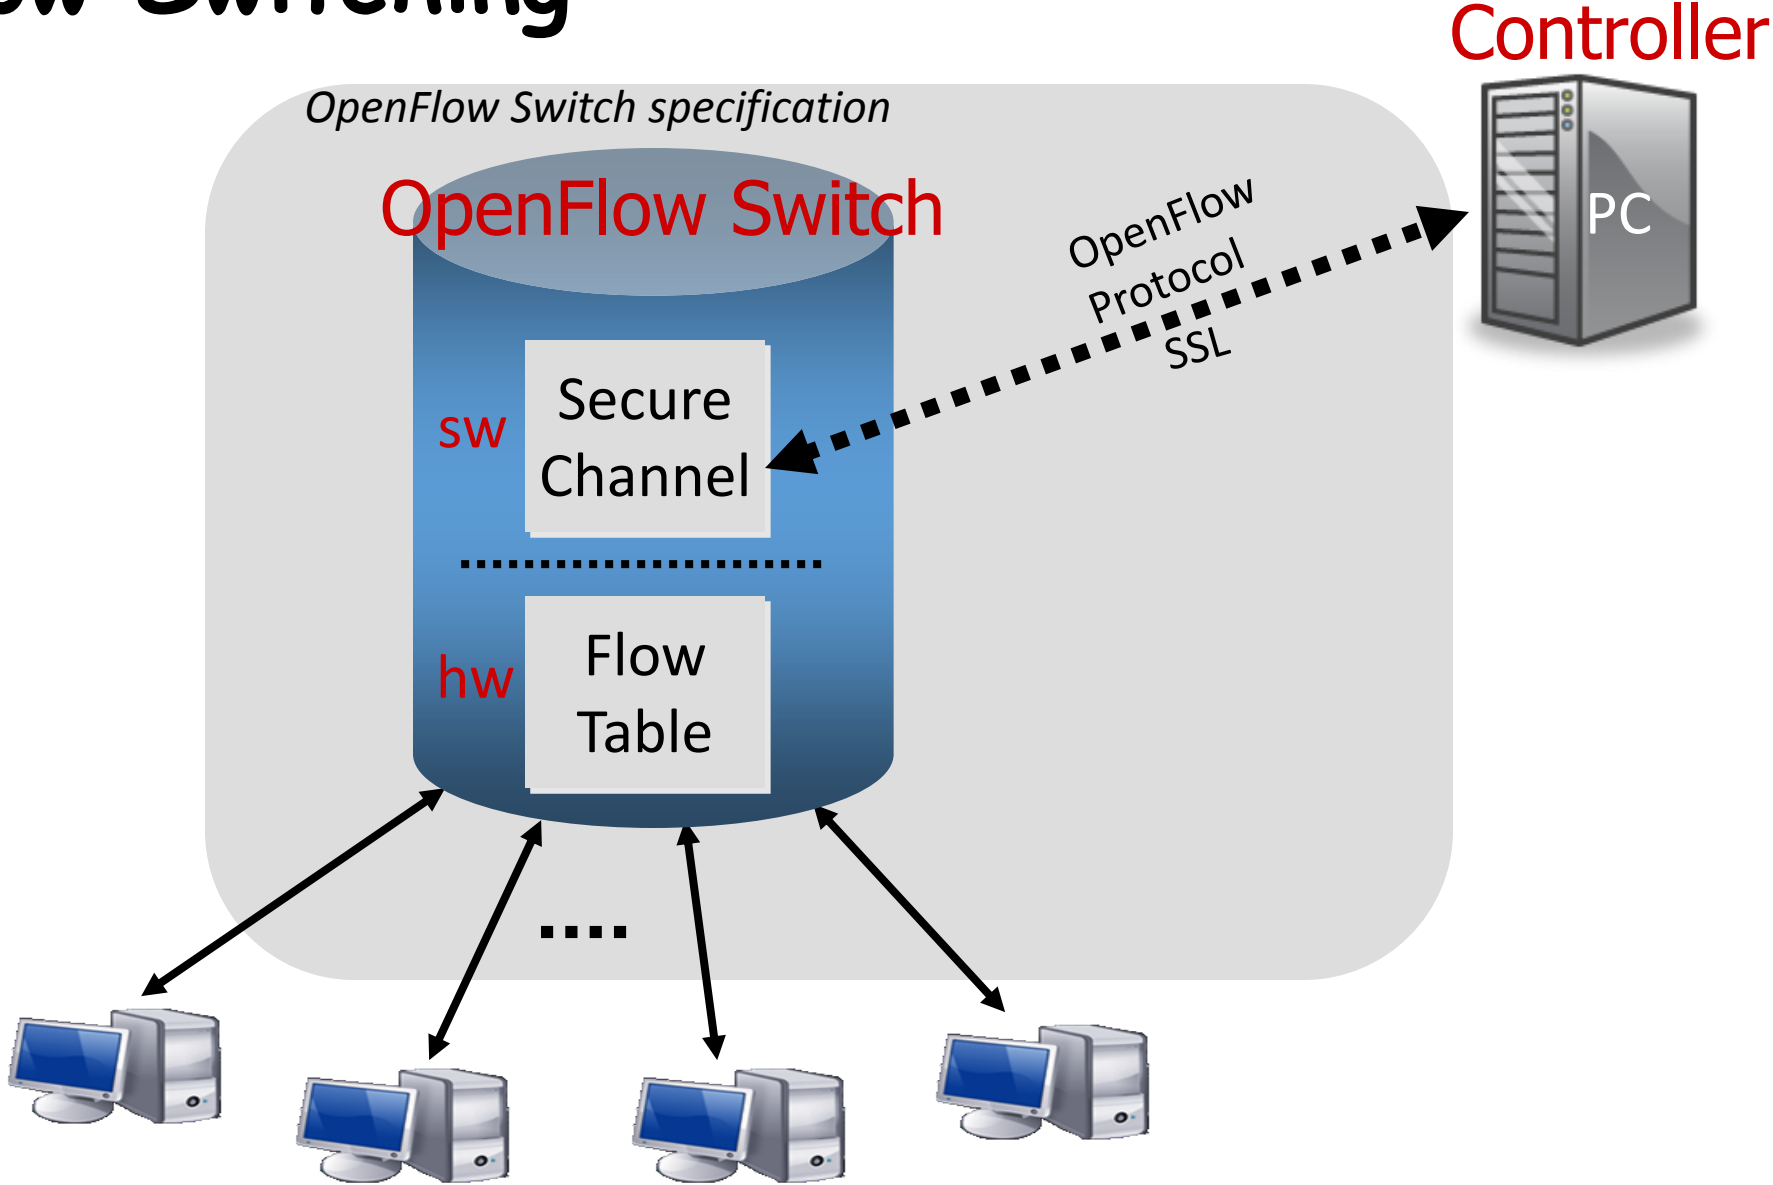
```
<xs:schema xmlns:ppns="http://schemas.dmtf.org/svpc/portprofile/1"
xmlns:xs="http://www.w3.org/2001/XMLSchema"
xmlns:cim="http://schemas.dmtf.org/wbem/wscim/1/common"
xmlns:xm1="http://www.w3.org/XML/1998/namespace"
xmlns:rasd="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/CIM_ResourceAllocationSettingData"
xmlns:epasd="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/CIM_EthernetPortAllocationSettingData"
xmlns:ns1="http://schemas.dmtf.org/svpc/portprofile/1"
targetNamespace="http://schemas.dmtf.org/svpc/portprofile/1" elementFormDefault="qualified" attributeFormDefault="qualified">
<xs:import namespace="http://www.w3.org/XML/1998/namespace" schemaLocation="http://www.w3.org/2001/xml.xsd"/>
<xs:import namespace="http://schemas.dmtf.org/wbem/wscim/1/common"
schemaLocation="http://schemas.dmtf.org/wbem/wscim/1/common.xsd"/>
<xs:import namespace="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/CIM_ResourceAllocationSettingData"
schemaLocation="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2.22.0/CIM_ResourceAllocationSettingData.xsd"/>
<xs:import namespace="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2/CIM_EthernetPortAllocationSettingData"
schemaLocation="http://schemas.dmtf.org/wbem/wscim/1/cim-schema/2.27.0/CIM_EthernetPortAllocationSettingData.xsd"/>
.....
<xs:element name="PortProfile">
<xs:annotation>
<xs:documentation>Root element of Port Profile </xs:documentation>
</xs:annotation>
<xs:complexType>
<xs:sequence>
<xs:element name="Item" type="epasd:CIM_EthernetPortAllocationSettingData_Type" minOccurs="0" maxOccurs="unbounded"/>
<xs:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
</xs:sequence>
</xs:complexType>
</xs:element>
</xs:schema>
```

Recent Network Virtualization Techniques

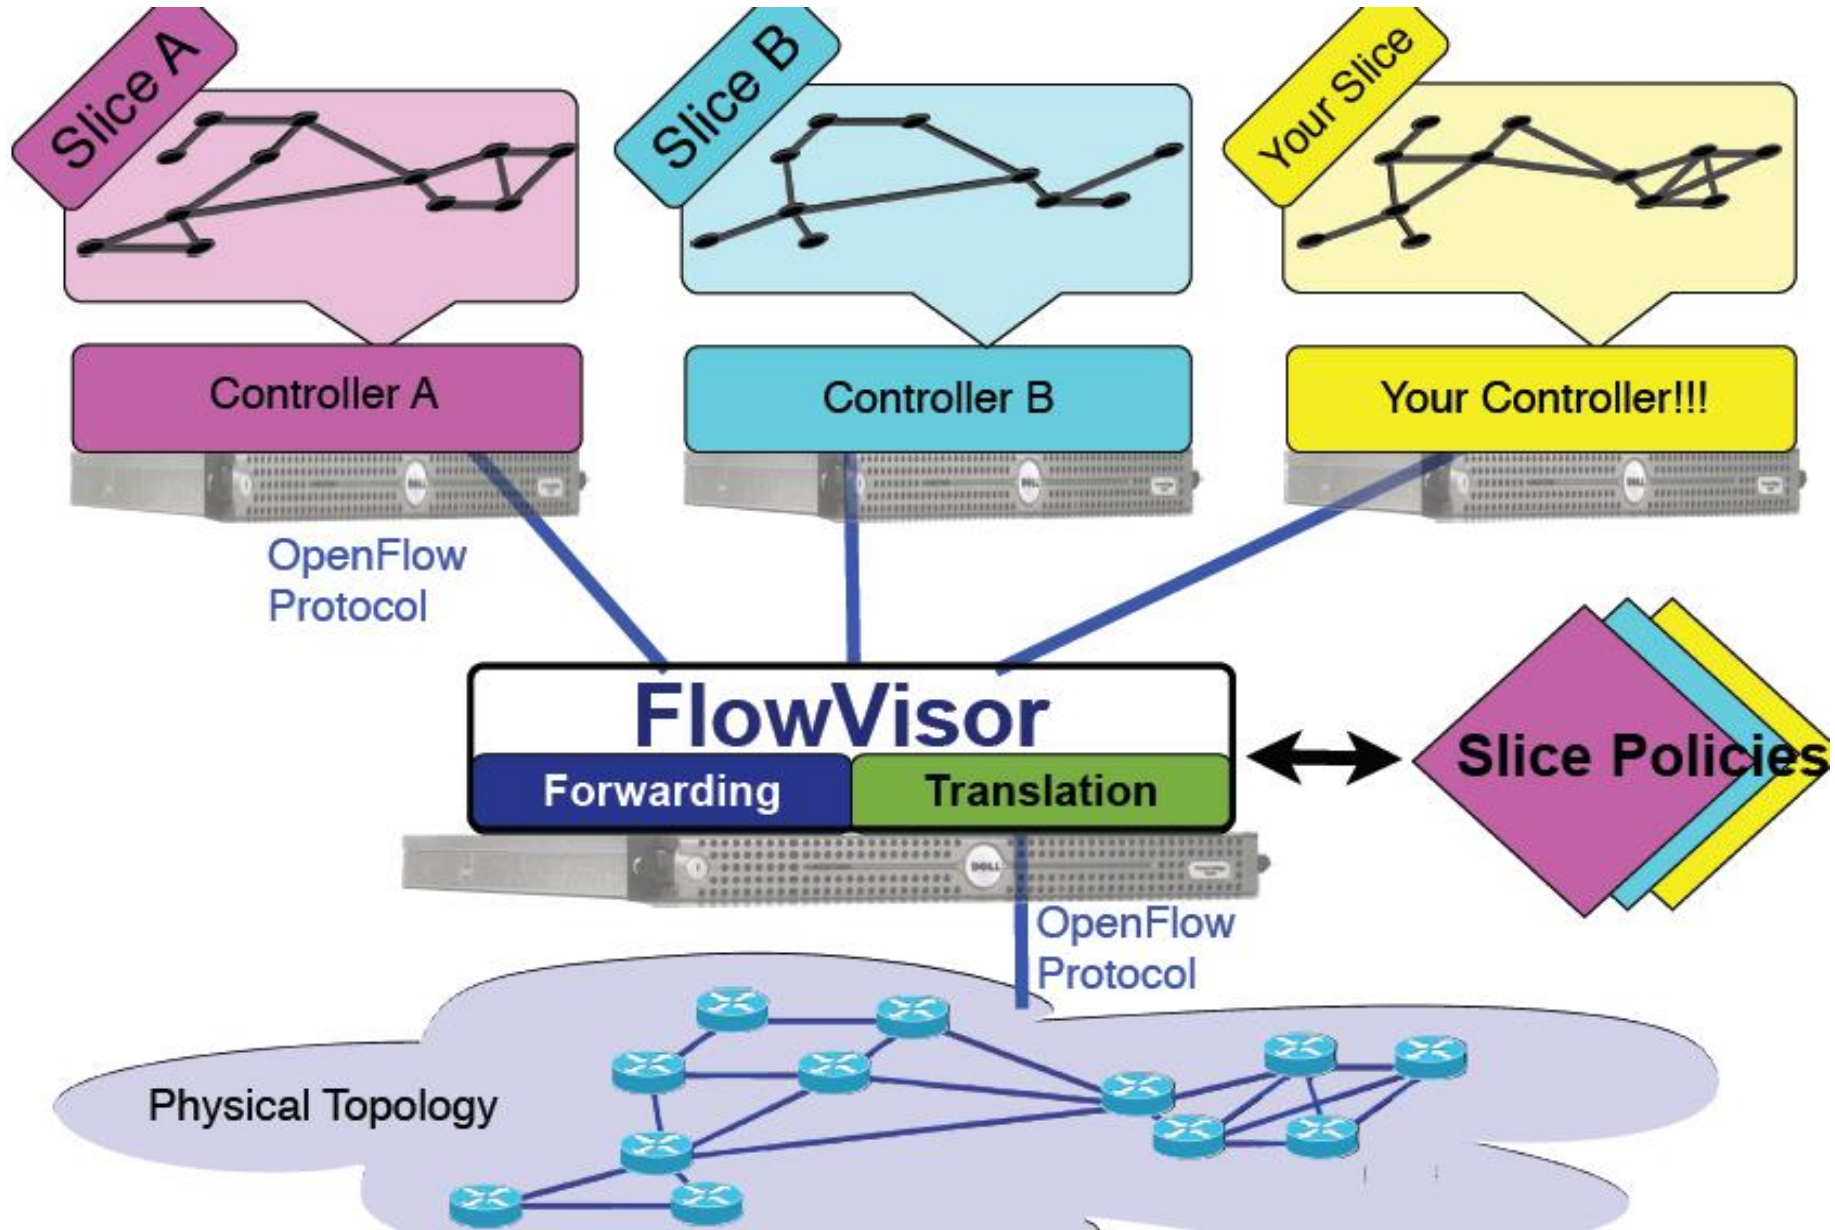
OpenFlow



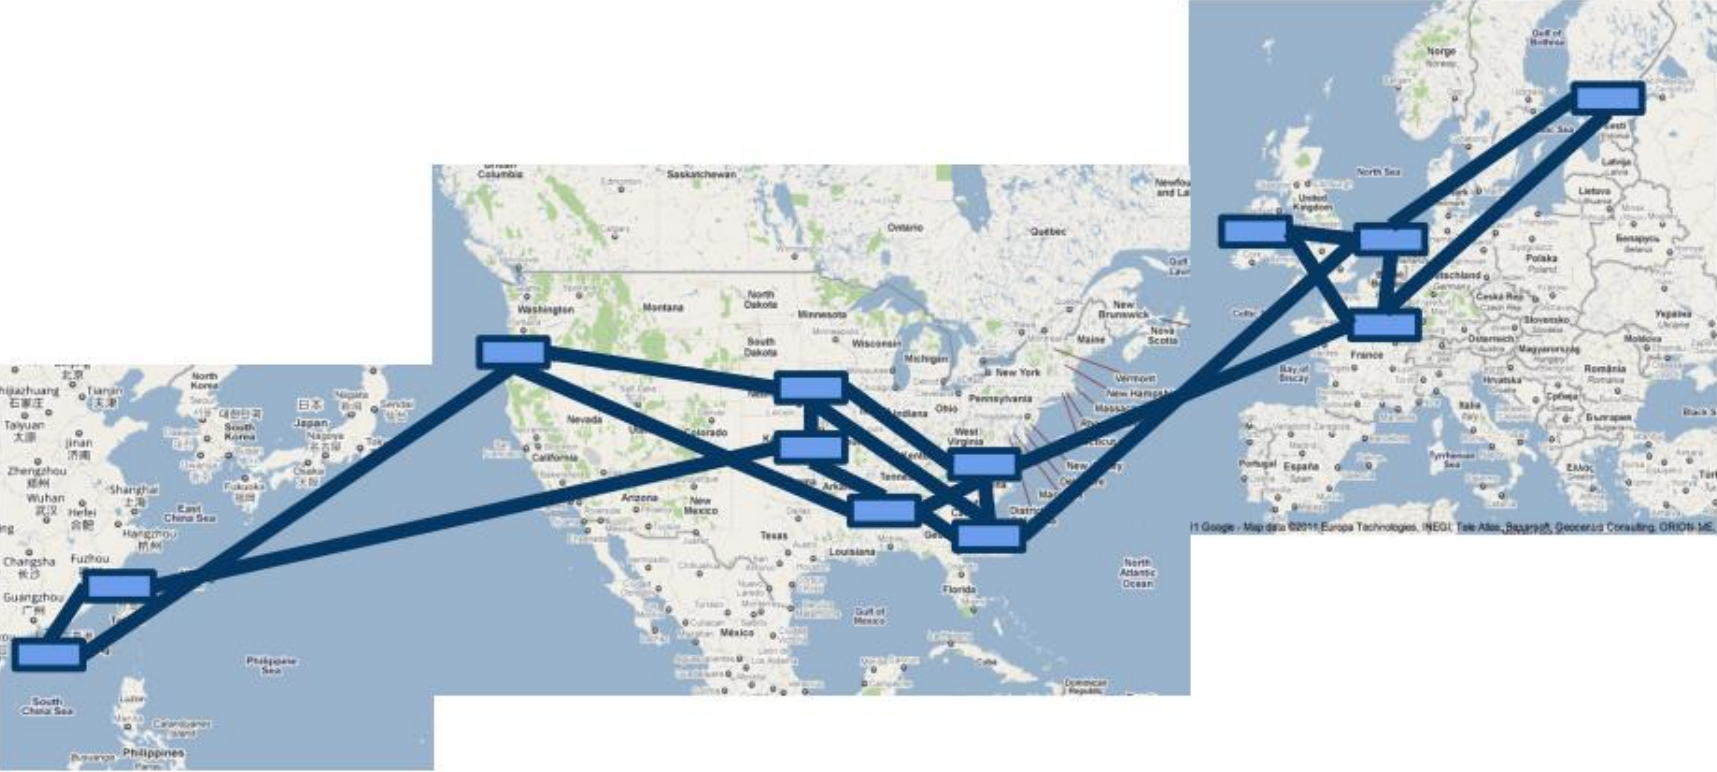
OpenFlow Switching



Flowvisor Virtualization



Google WAN




ElasticTree

Goal: Reduce energy usage in data center networks

Approach:

1. Reroute traffic
2. Shut off links and switches to reduce power



Software Defined Network (SDN)

Software Defined Networking


In the Software Defined Networking architecture, the control and data planes are decoupled, network intelligence and state are logically centralized, and the underlying network infrastructure is abstracted from the applications.

**Software-Defined Networking:
The New Norm for Networks
ONF White Paper
April 13, 2012**

Trend

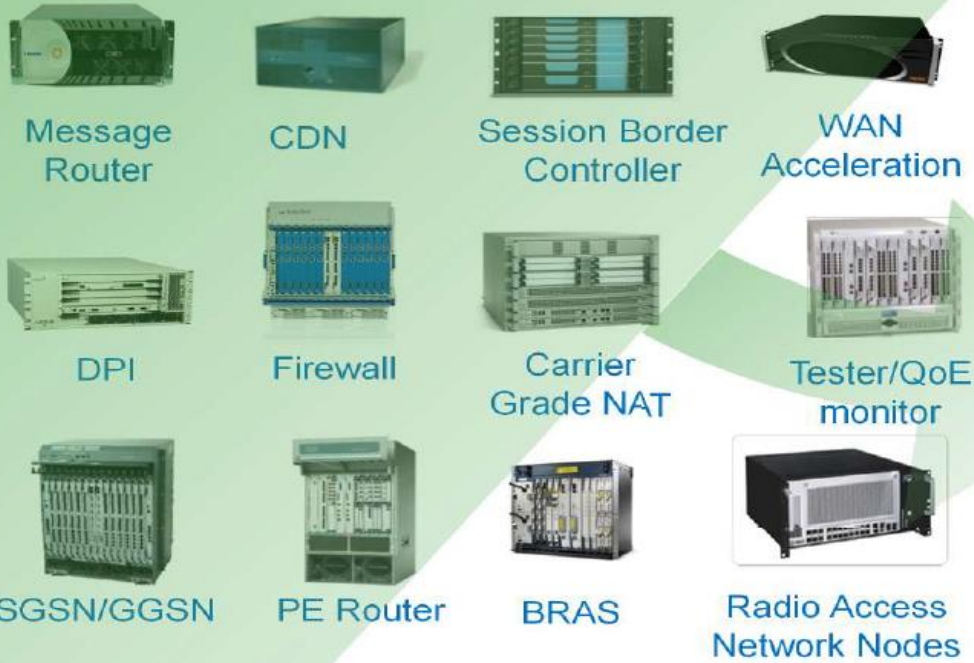


“Mainframe”



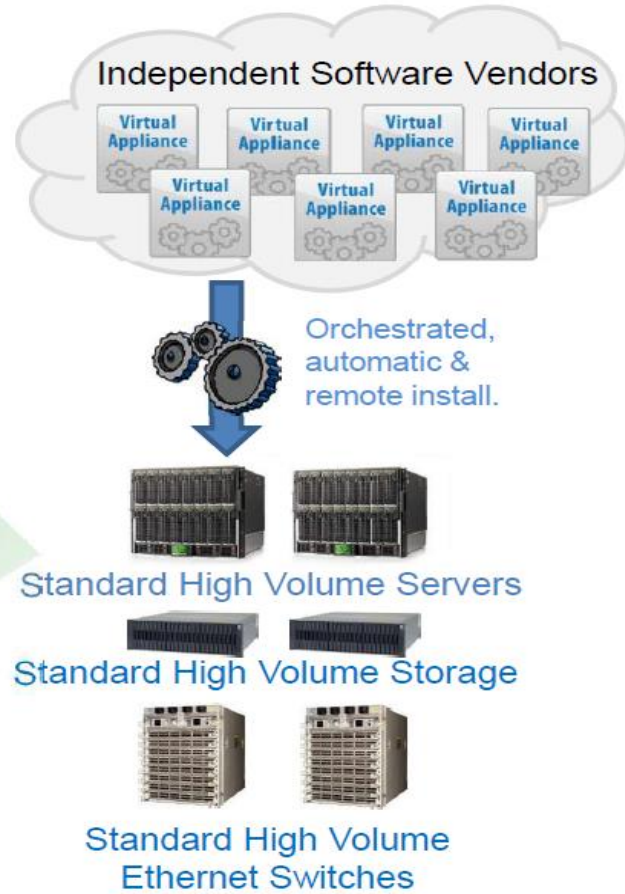
Network Function Virtualization (NFV)

Classical Network Appliance Approach



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

Source: NFV



Network Virtualisation Approach