SUPERFLUIDITY

A super-fluid, cloud-native, converged edge system

La ricerca sui nuovi modelli di rete

Stefano Salsano, *Pier Luigi Ventre*
Univ. of Rome Tor Vergata, Italy / CNIT, Italy

Workshop GARR 2017 - Netvolution
Roma, 6 Aprile 2017

Outline

- Netvolution : Softwarization, SDN, NFV…

- A “softwarization” path for GÉANT – the Vision
  (NB: it’s our vision coming from our activities in GN4…)

- SUPERFLUIDITY project – goals and approach
  - Highly dynamic and granular Network Function Virtualization

- Back to GÉANT – the reality of implementation work

Network Softwarization

We're in the middle (or at the beginning…) of a revolution.

A set of technologies and trends is changing the way networks are designed, provisioned, operated.

Technologies and trends

- Software Defined Networking (SDN)
- Network Function Virtualization (NFV)
- Cloud computing / edge computing
- Open innovation, Open Source

Software Defined Networking (SDN)

Separation of control plane and forwarding plane.

2008: first paper on OpenFlow, 2011 OpenFlow 1.1 released by ONF

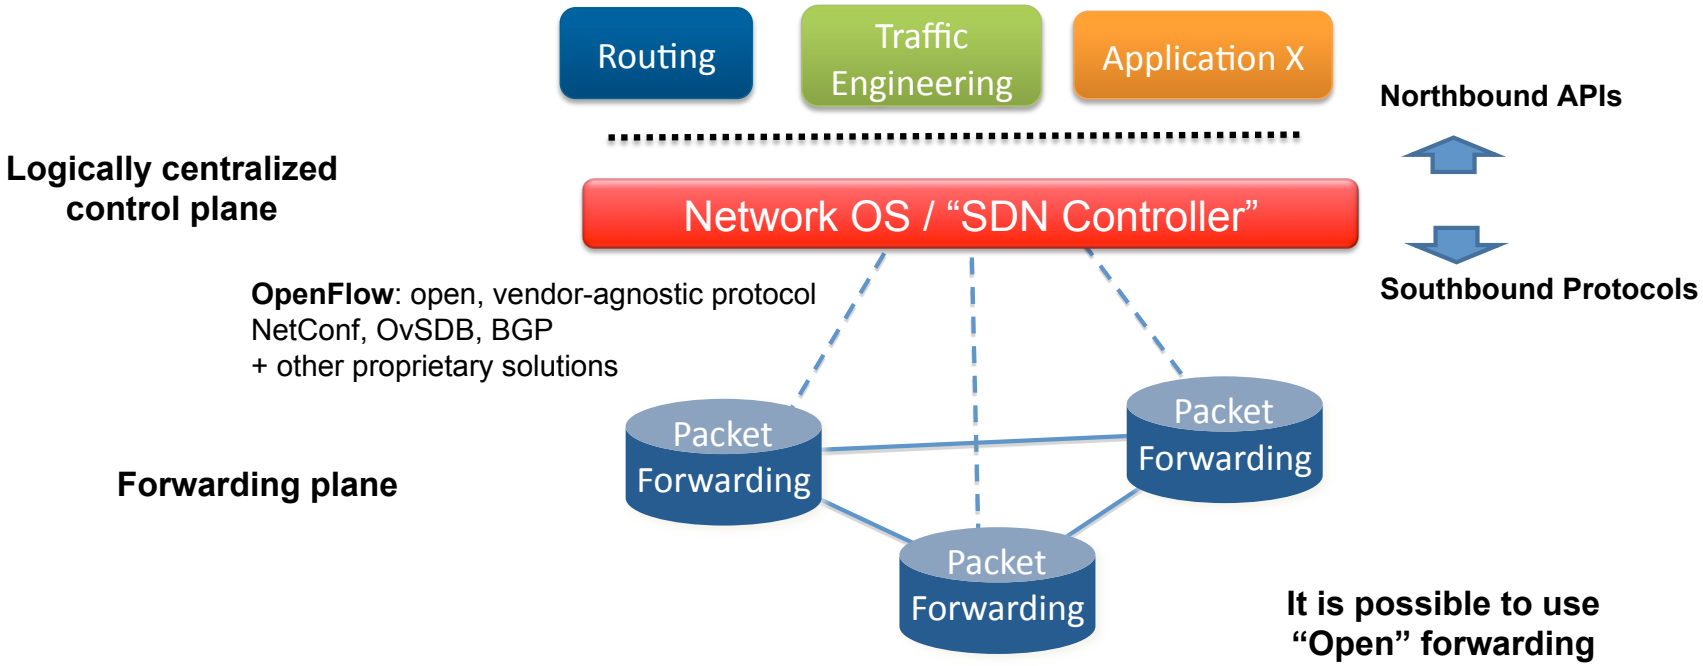

Network Function Virtualization (NFV)

- Cloud computing principles applied to networking equipment:
  - virtualize the Network Functions
  - execute them as “VNFs” into generic purpose servers

- NFV: a Network-operator-driven specification group within ETSI, started in November 2012
- Membership of ETSI NFV ISG has grown to over 290 individual companies including 38 of the world's major service providers (as of May 2016)

http://www.etsi.org/technologies-clusters/technologies/nfv

ETSI MANagement and Orchestration (MANO) Model

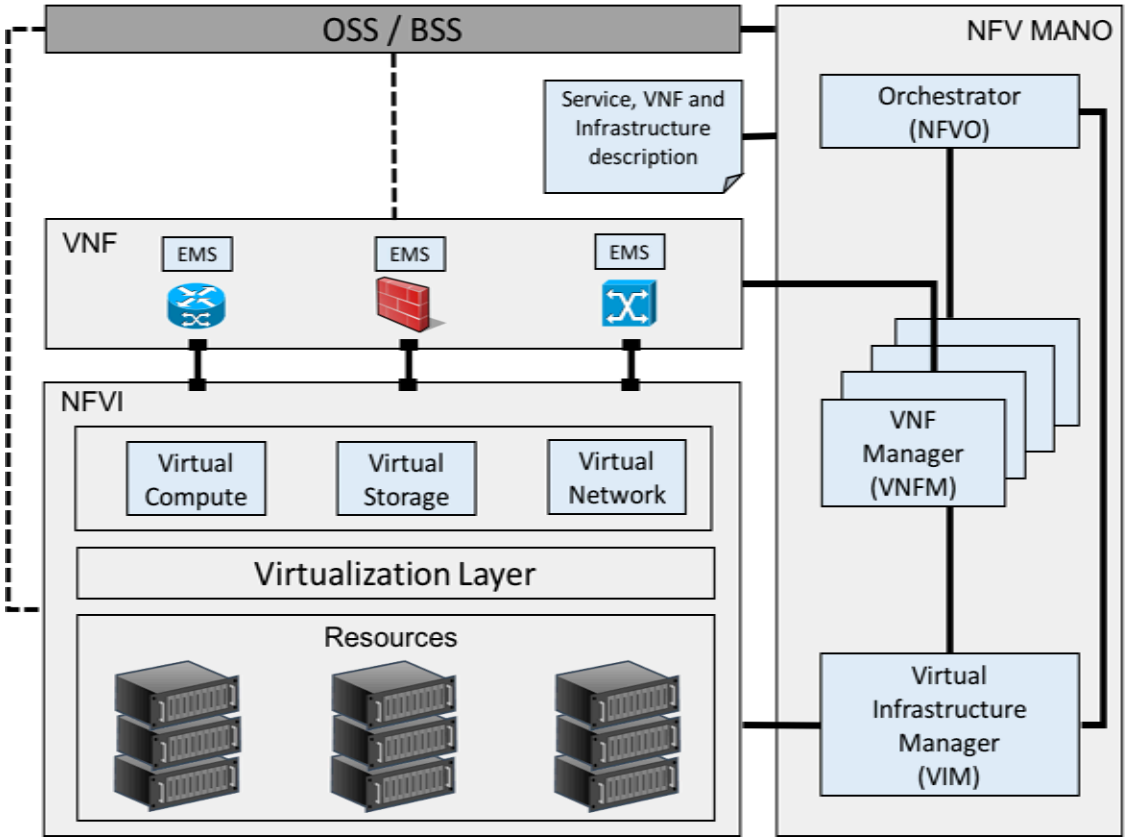

ETSI MANagement and Orchestration (MANO) Model

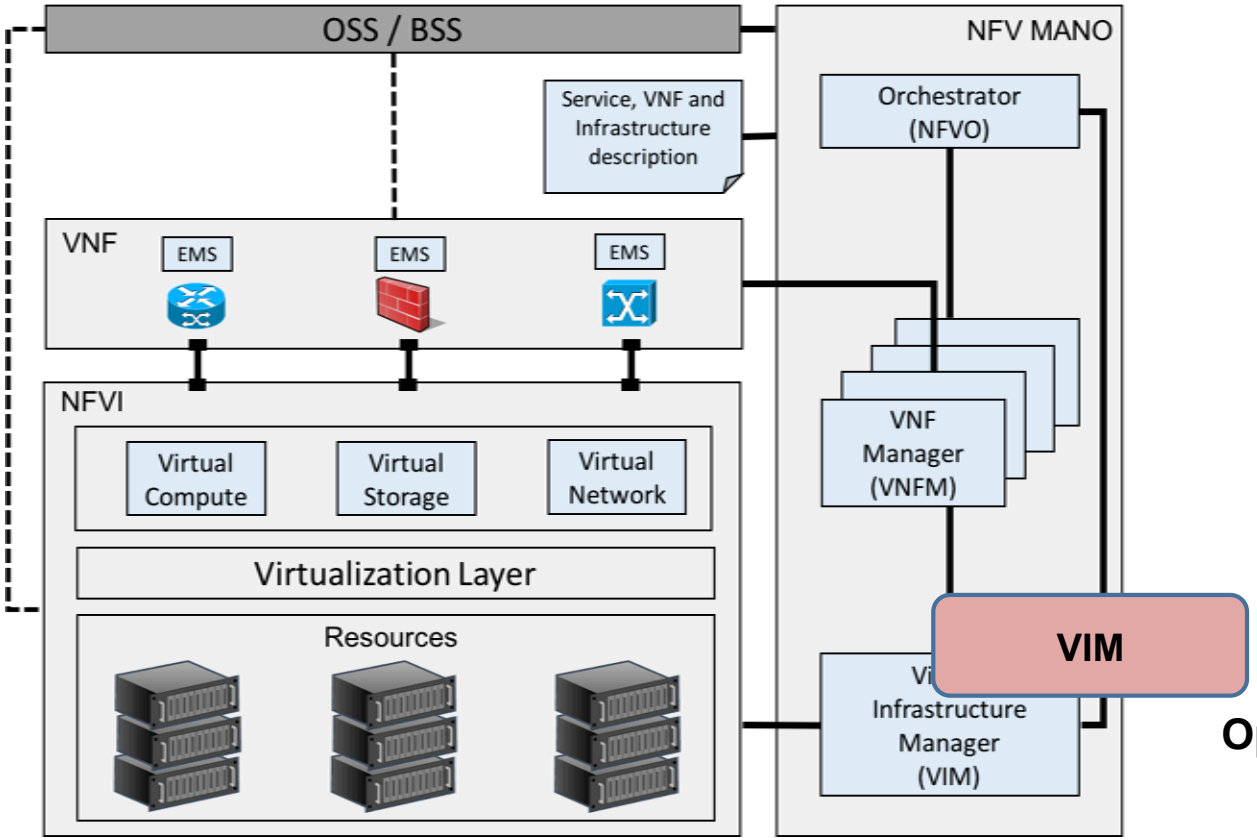

ETSI MANagement and Orchestration (MANO) Model

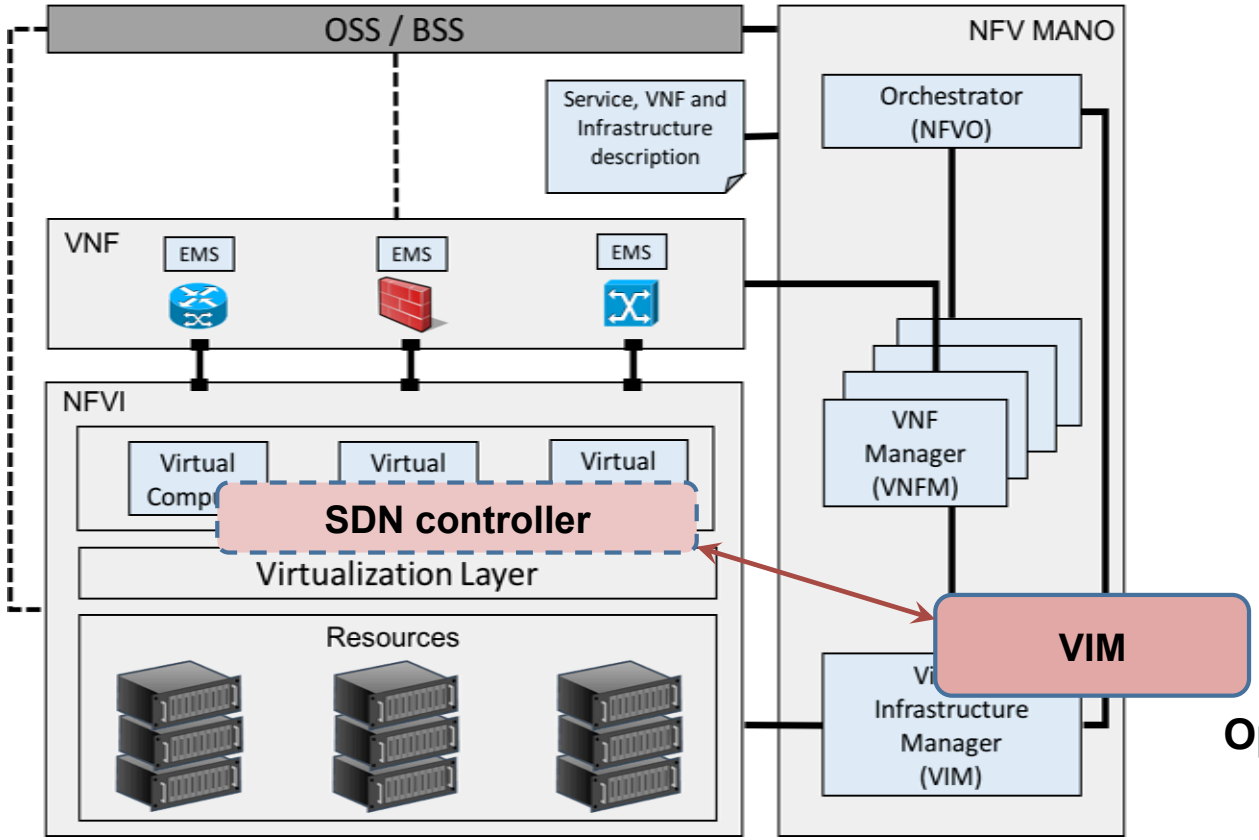

OpenDayLight, ONOS

OpenStack

ETSI MANagement and Orchestration (MANO) Model

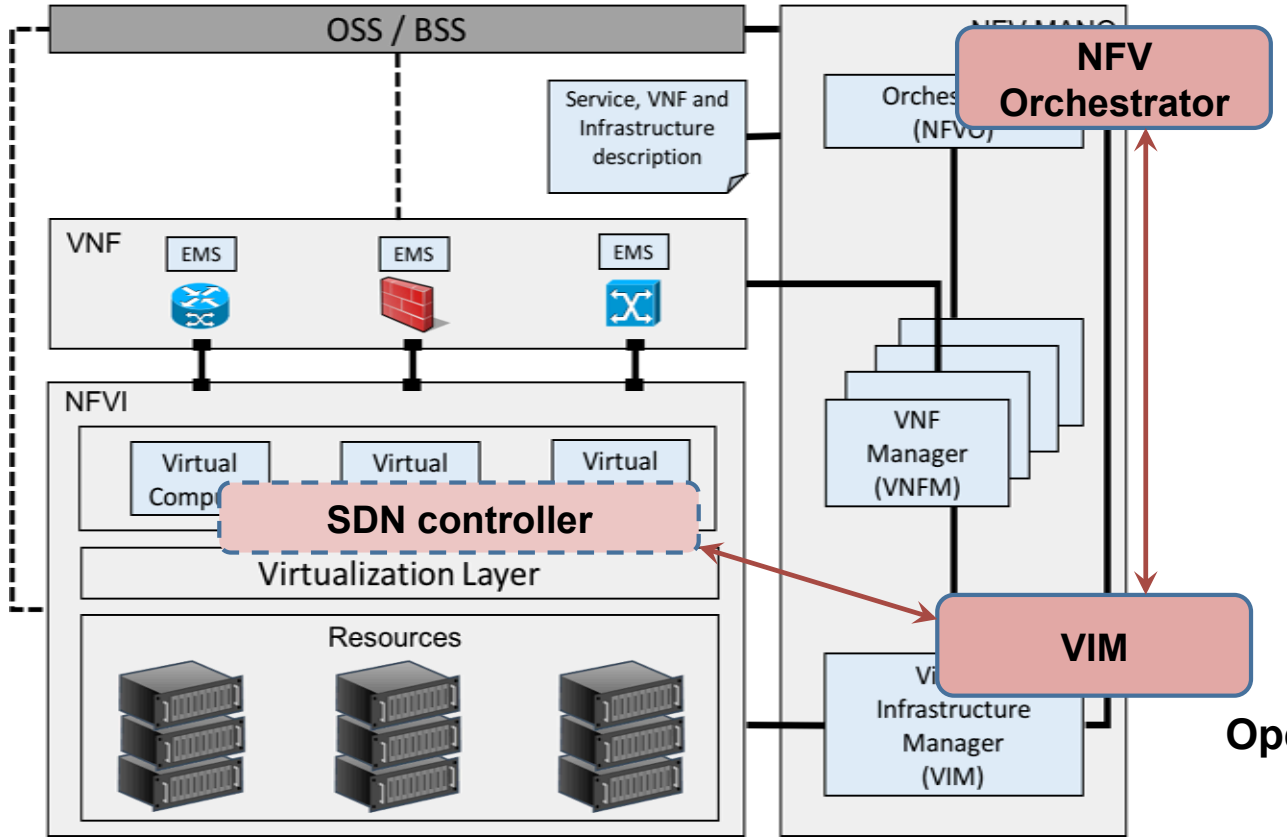

ONAP
OpenO
ECOMP
OpenBaton
OSM
Aria/Cloudify
…

OpenDayLight,
ONOS

OpenStack

Standard Defining Organizations and Open Source Communities

Standards
- ETSI NFV ISF
- IETF (SFC)
- ONF (SDN, OpenFlow)

NFV Orchestrators
- OSM (Open Source MANO)
- ONAP (Open-O+ECOMP)
- OpenBaton, …

Virtual Infrastructure Managers (VIMs)
- OpenStack (30 mainstream projects)

SDN controllers
- OpenDayLight
- ONOS

Data plane forwarding
- Open vSwitch
- FD.IO, …

Overall integration
- OPNFV (50 ongoing projects)

Outline

- Netvolution : Softwarization, SDN, NFV…
- A “softwarization” path for GÉANT – the Vision
  (NB: it’s our vision coming from our activities in GN4…)
- SUPERFLUIDITY project – goals and approach
  - Highly dynamic and granular Network Function Virtualization
- Back to GÉANT – the reality of implementation work

Towards a new era for GÉANT

- *SDN* and *Softwarization* benefits:
  - Provisioning procedures drastically simplified (from days to ?)
  - Cheaper hardware in place of current equipment
  - Open APIs, reduction of running protocols, simplification of the control plane
  - Less proprietary implementations
  - Easier development and introduction of new services

A result from GN4 phase 1

Pier Luigi Ventre, Stefano Salsano, Matteo Gerola, Elio Salvadori, Mian Usman, Sebastiano Buscaglione, Luca Prete, Jonathan Hart, and William Snow

"SDN-Based IP and Layer 2 Services with Open Networking Operating System in the GÉANT Service Provider Network",

IEEE Communication Magazine, "SDN Use Cases for Service Provider Networks", to appear, April 2017

GÉANT Softwarization path

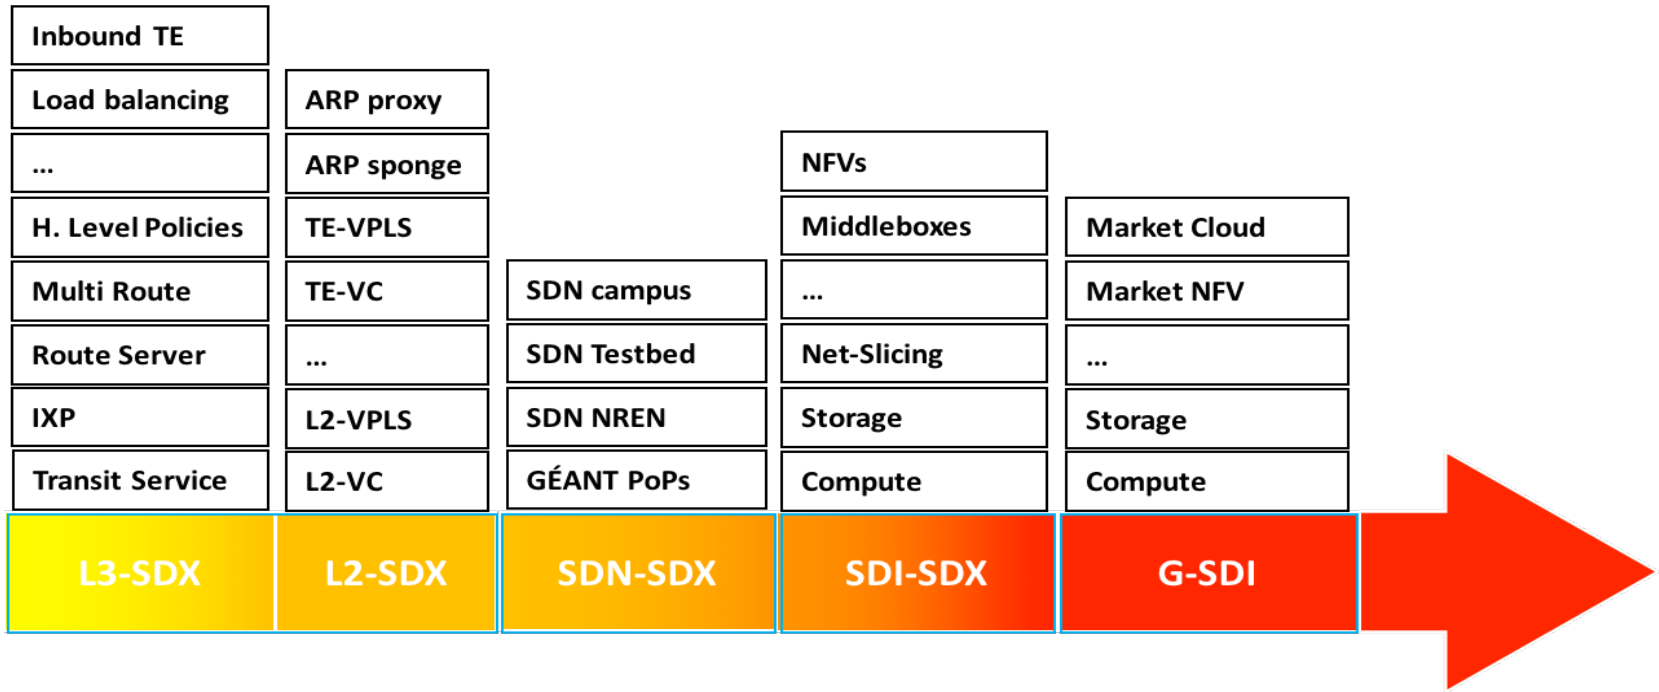

GÉANT Softwarization path

SDNization of the GÉANT layer 2 and layer 3 services

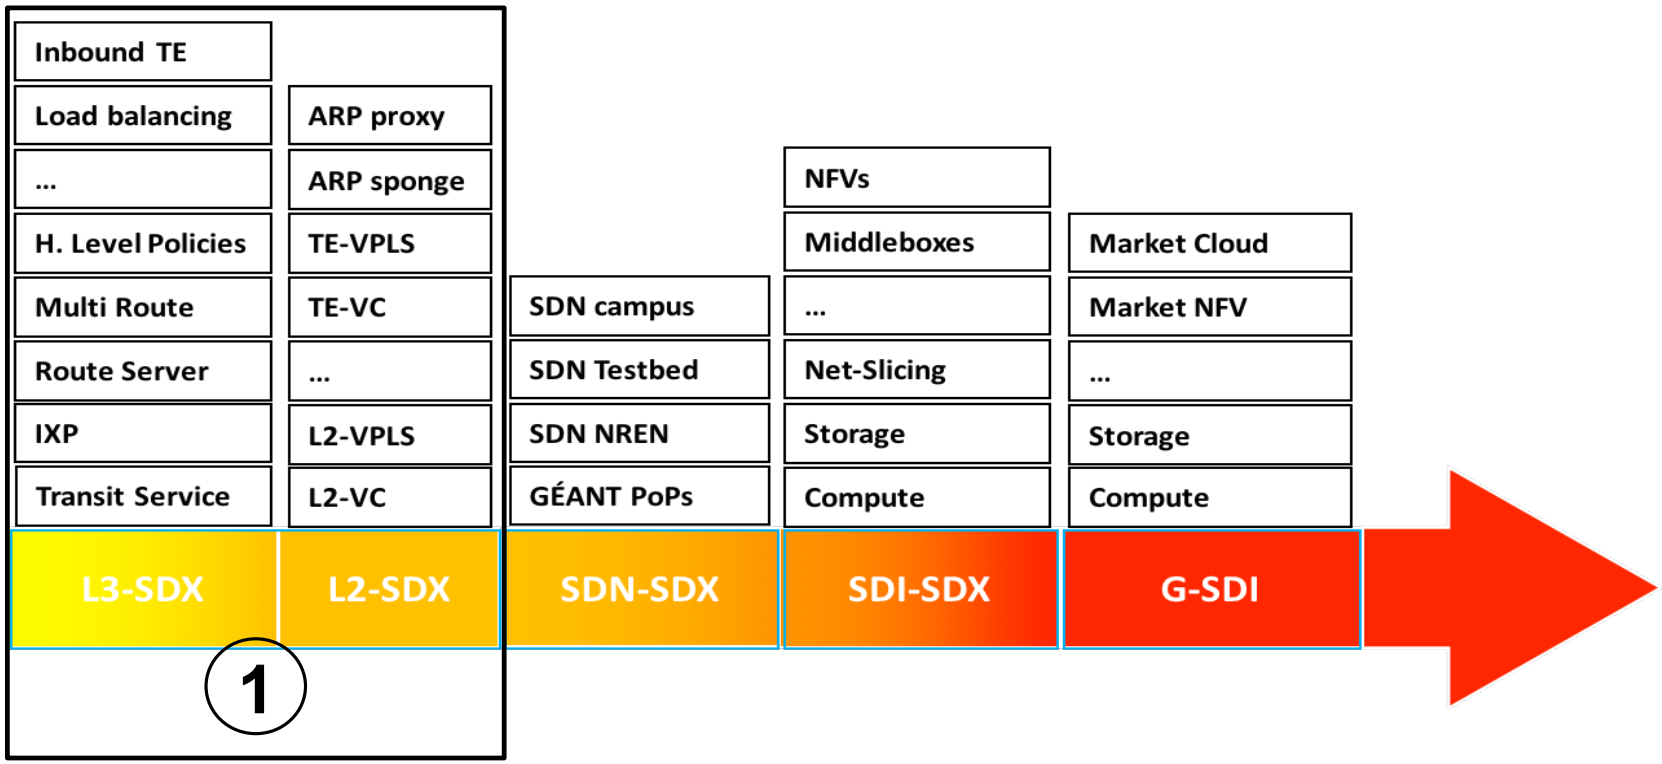

GÉANT Softwarization path

SDNization of the GÉANT layer 2 and layer 3 services

Interconnection and harmonization of the SDN infrastructures between GÉANT and the NRENs

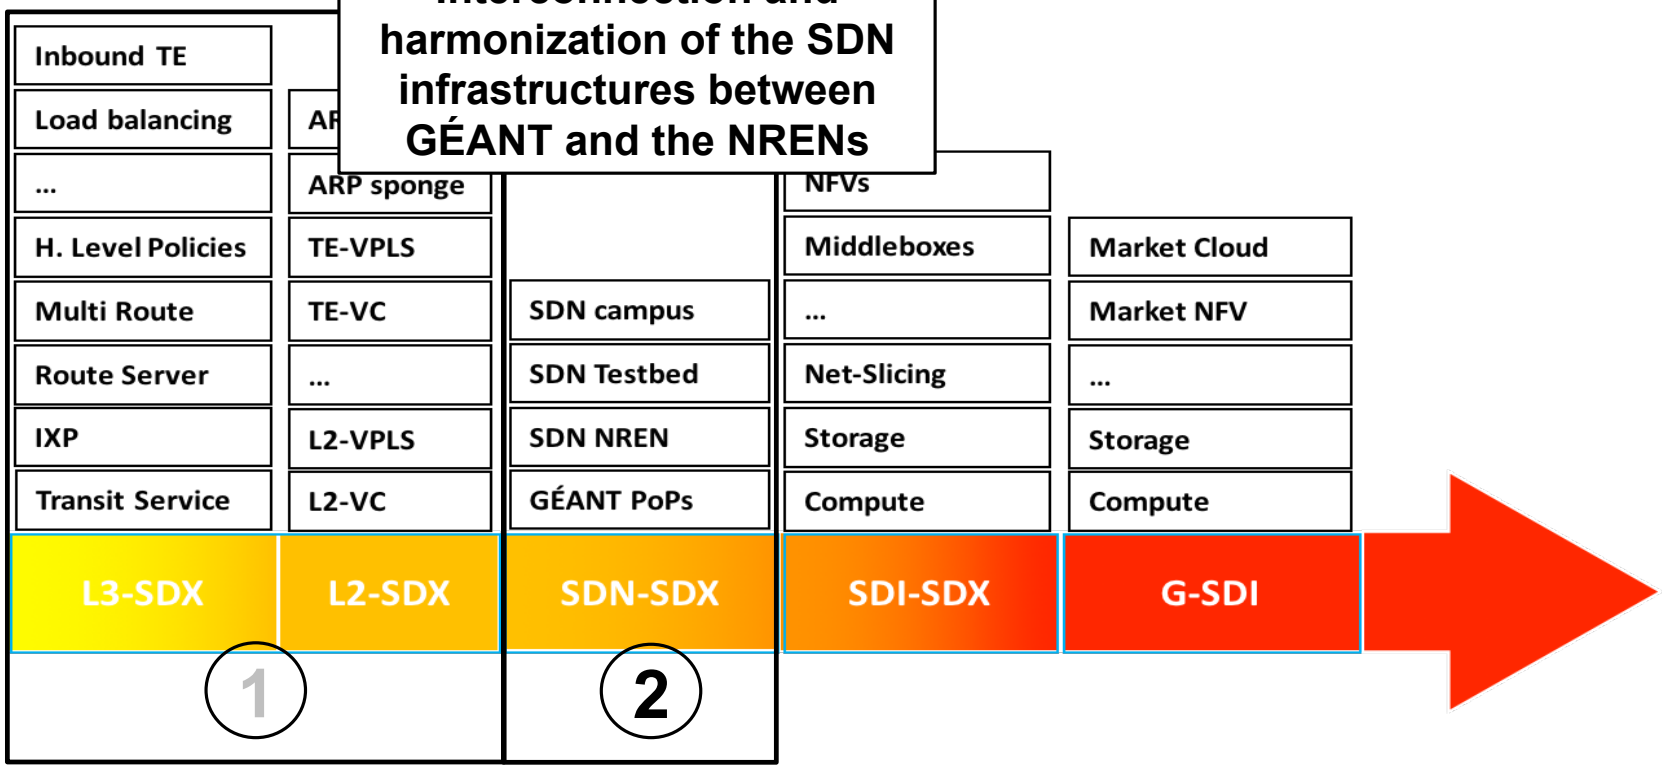

GÉANT Softwarization path

SDNization of the GÉANT layer 2 and layer 3 services

GÉANT also provides computing / storage / NFV

Interconnection and harmonization of the SDN infrastructures between GÉANT and the NRENs

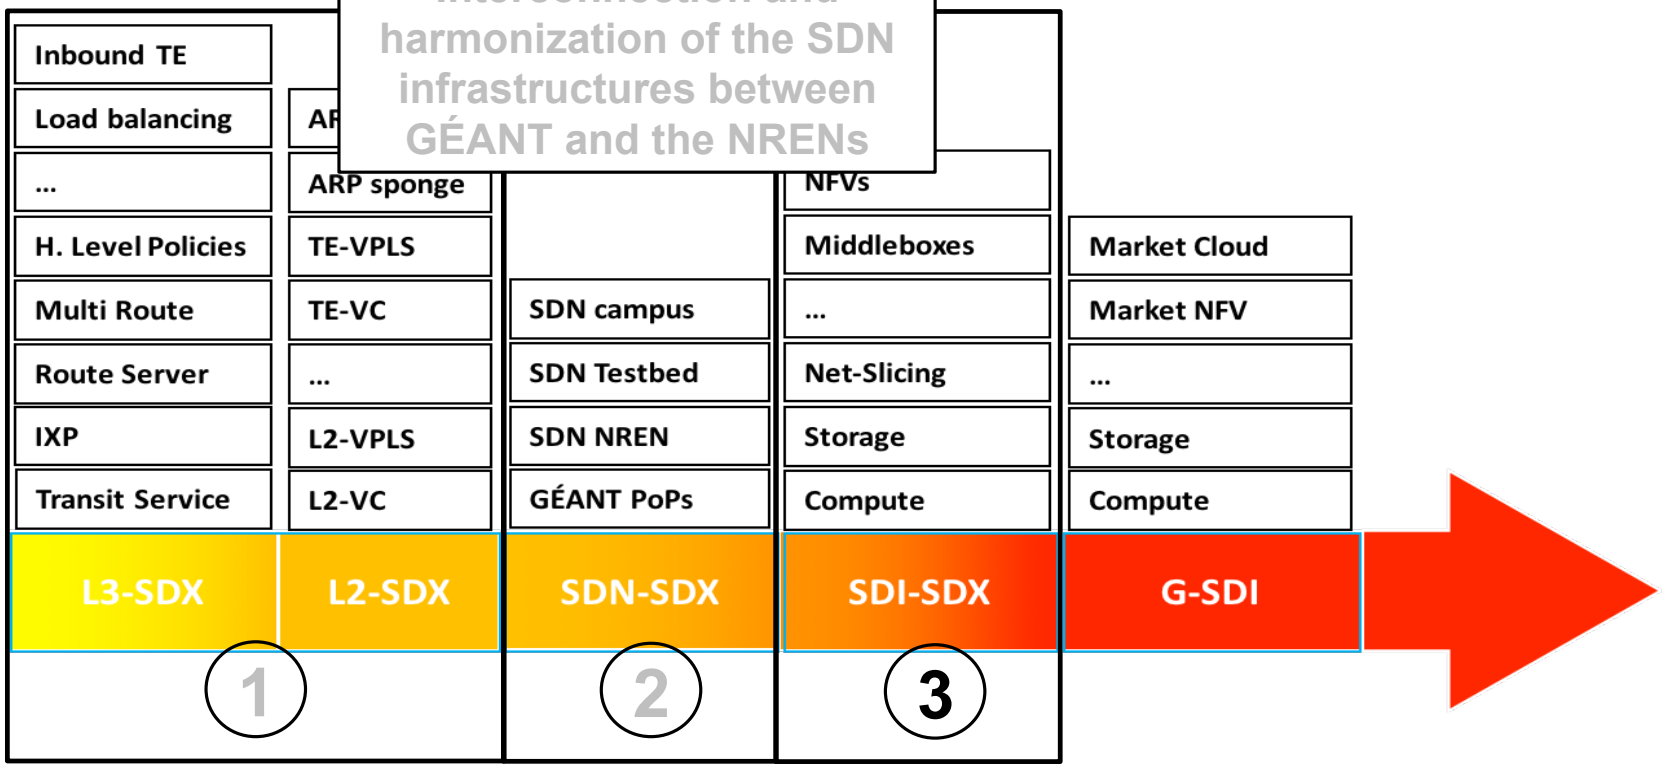

GÉANT Softwarization path

SDNization of the GÉANT layer 2 and layer 3 services

GÉANT also provides computing / storage / NFV

Interconnection and harmonization of the SDN infrastructures between GÉANT and the NRENs

Full Softwarization, Harmonization and integration of computing / storage / NFV between GÉANT and NRENs

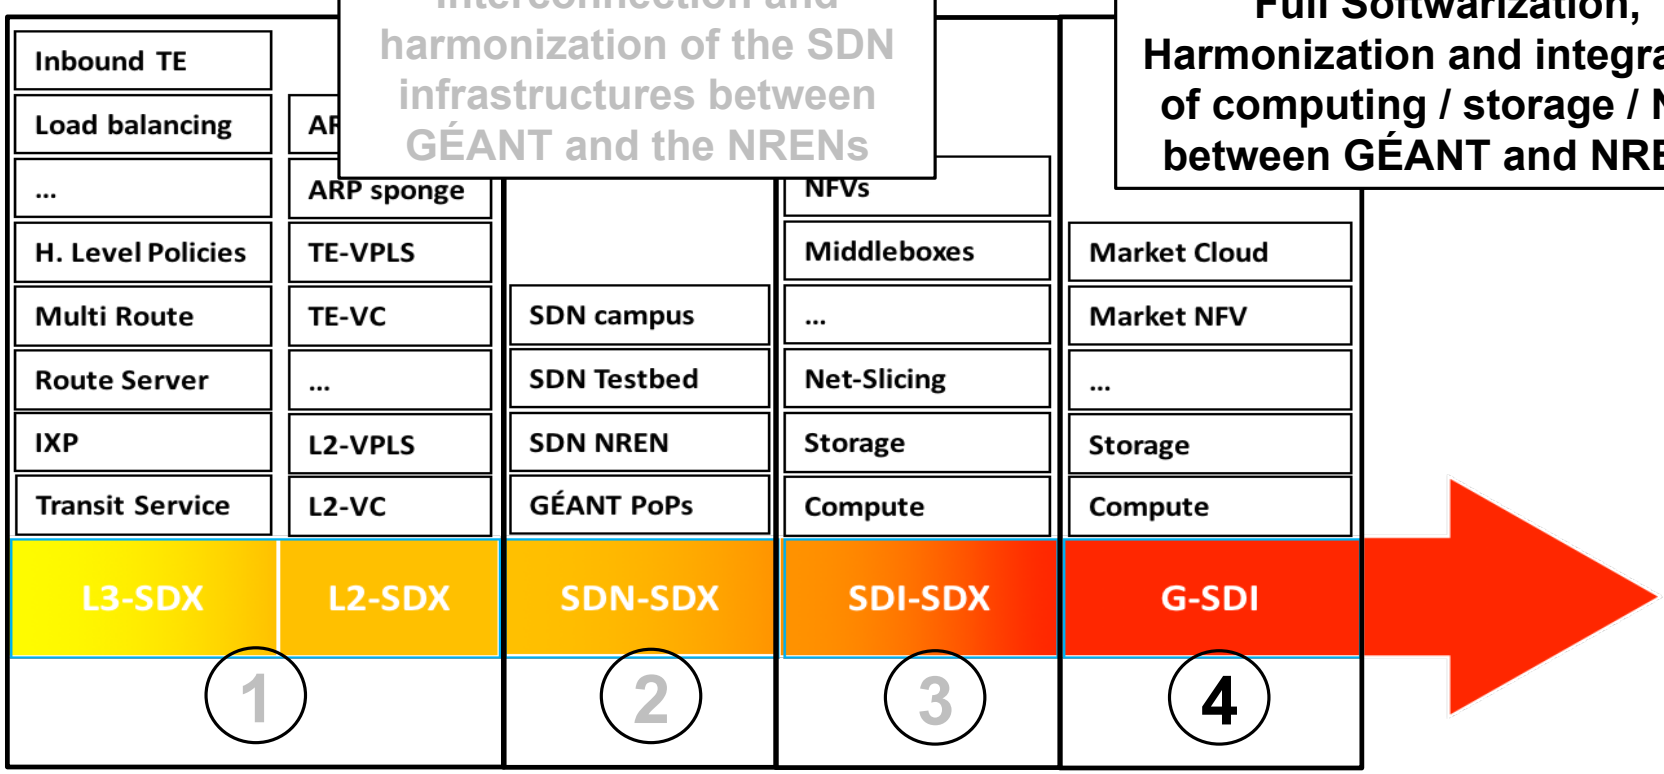

Outline

- Netvolution : Softwarization, SDN, NFV…
- A “softwarization” path for GÉANT – the Vision
  (NB: it’s our vision coming from our activities in GN4…)
- SUPERFLUIDITY project – goals and approach
  - Highly dynamic and granular Network Function Virtualization
  - Toward sub 10 ms service instantiation
- Back to GÉANT – the reality of implementation work

SUPERFLUIDITY project - http://superfluidity.eu/

Goal: a superfluid NFV approach

- Instantiate network functions and services on-the-fly
- Run them anywhere in the network (core, aggregation, edge), across heterogeneous infrastructure environments (computing and networking), taking advantage of specific hardware features, such as high performance accelerators, when available

Approach

- Decomposition of network components and services into elementary and reusable primitives ("Reusable Functional Blocks – RFBs")
- Platform-independent abstractions, permitting reuse of network functions across heterogeneous hardware platforms

The Superfluidity vision

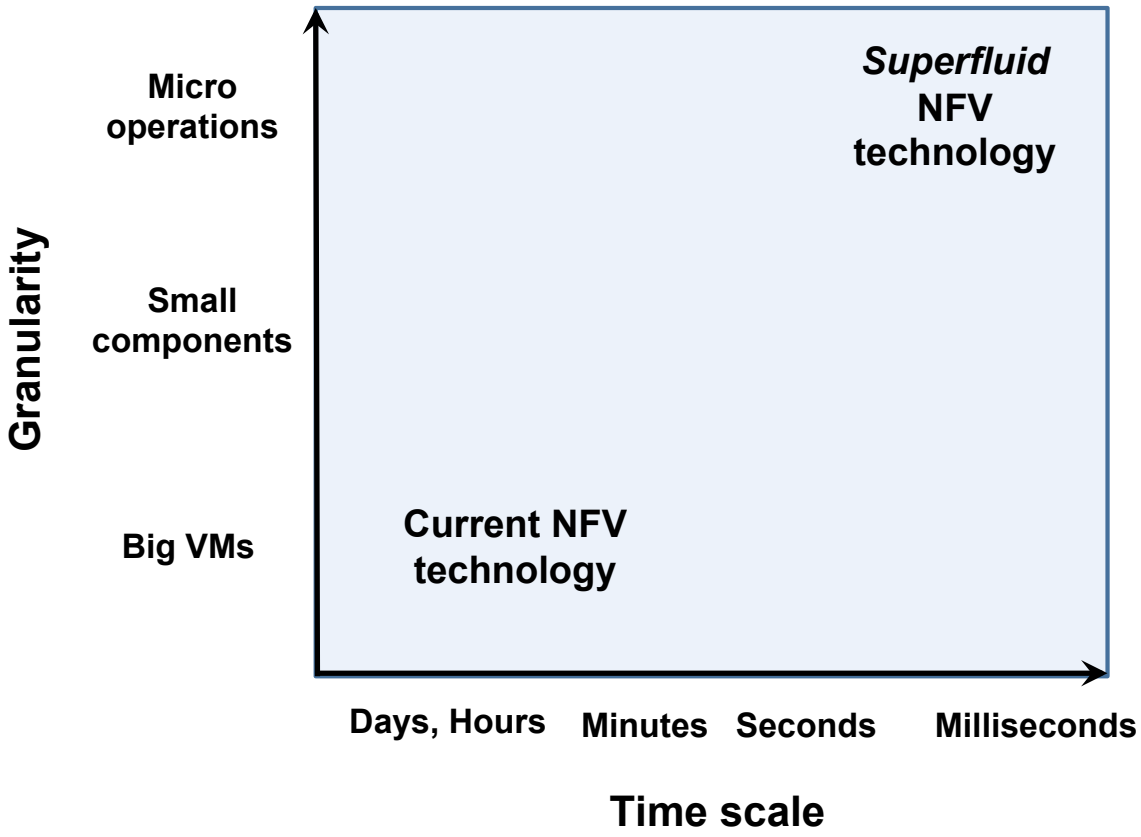

- From VNF
  Virtual Network Functions
  to RFB
  Reusable Functional Blocks

- Heterogeneous RFB execution environments
  - Hypervisors
  - Modular routers
  - Packet processors
  - …

Virtualization technologies

Unikernels

Specialized VMs (e.g. MiniOS, ClickOS…)

- Strong isolation
- Very Lightweight
- Very good security properties

Containers

e.g. Docker

- Lightweight (not enough?)
- Poor isolation

Hypervisors (traditional VMs)

e.g. XEN, KVM, wmware…

- Strong isolation
- Heavyweight

Outline

- Netvolution : Softwarization, SDN, NFV…
- A “softwarization” path for GÉANT – the Vision
  (NB: it’s our vision coming from our activities in GN4…)
- SUPERFLUIDITY project – goals and approach
  - Highly dynamic and granular Network Function Virtualization
- Back to GÉANT – the reality of implementation work

Simple and straightforward ?

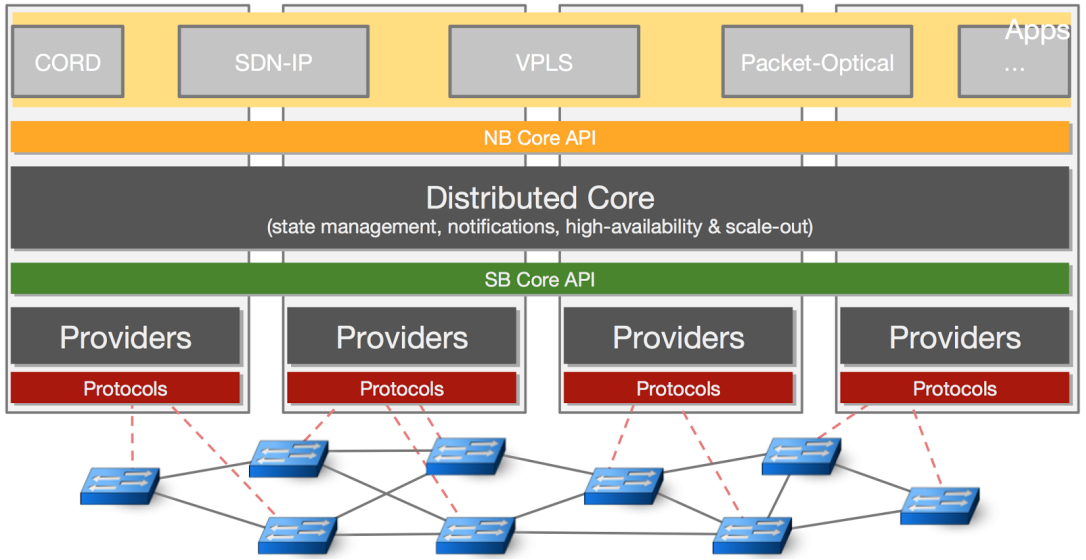

- Challenges: networking, distributed systems and software engineering
- Expertise: not only networking !!
- Vendors: Cisco, Juniper, Corsa, Broadcom, HP, Arista, Ciena, …

Developing in the Softwarization era

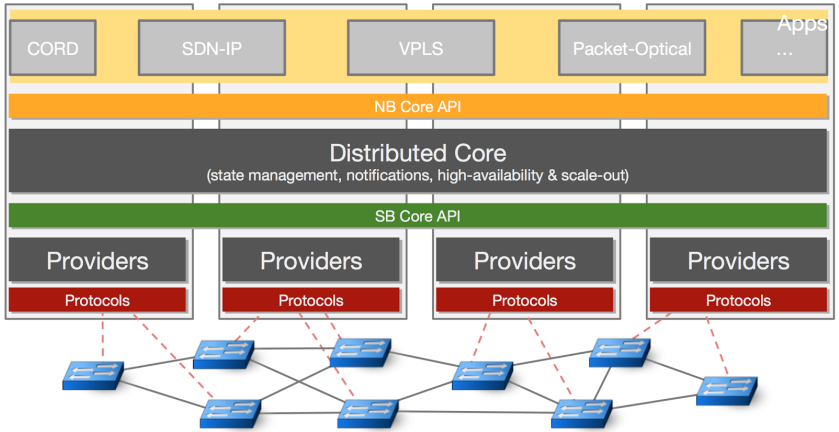

- Service X is re-engineered to run as application on top of the NOS Y
  - A number of MM to develop the service
- Unit tests, Functional Tests, Non-Functional tests, Proof of Concepts
  - A number of MM to spend
- Is that all ?

Where is my driver ?

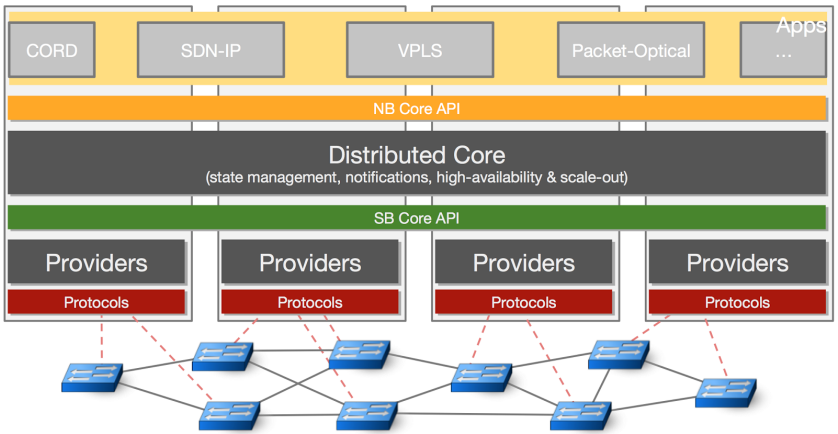

- Assuming ONOS as control plane and CORSA devices for the data plane;
- Should ONF/ON.LAB develop the driver for CORSA device ?
- Should Corsa develop the driver for ONOS ?
- **You have to develop the driver**

Open APIs, simplified control planes… really ?

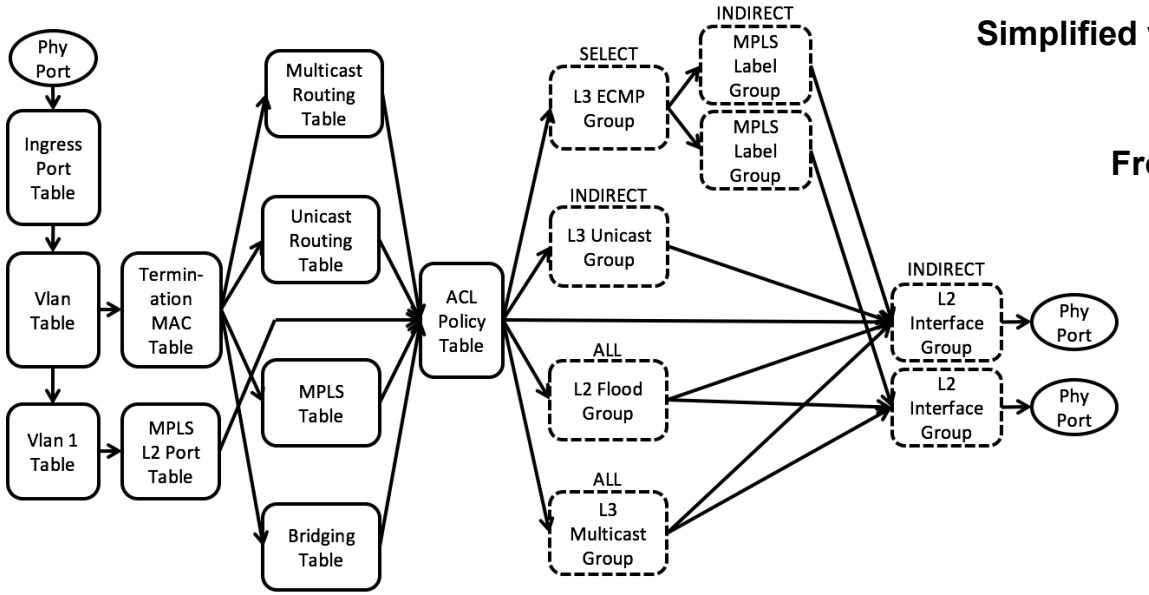

Simplified view of the Broadcomm pipeline

From wiki.opencord.org

- A lot of functionalities: ECMP, Multicast, VPWS and so on
- Mastering a huge complexity. You have to do this and there is no CISCO/Juniper…
- What if you change vendor ?

Déjà vu – is it happening again ?

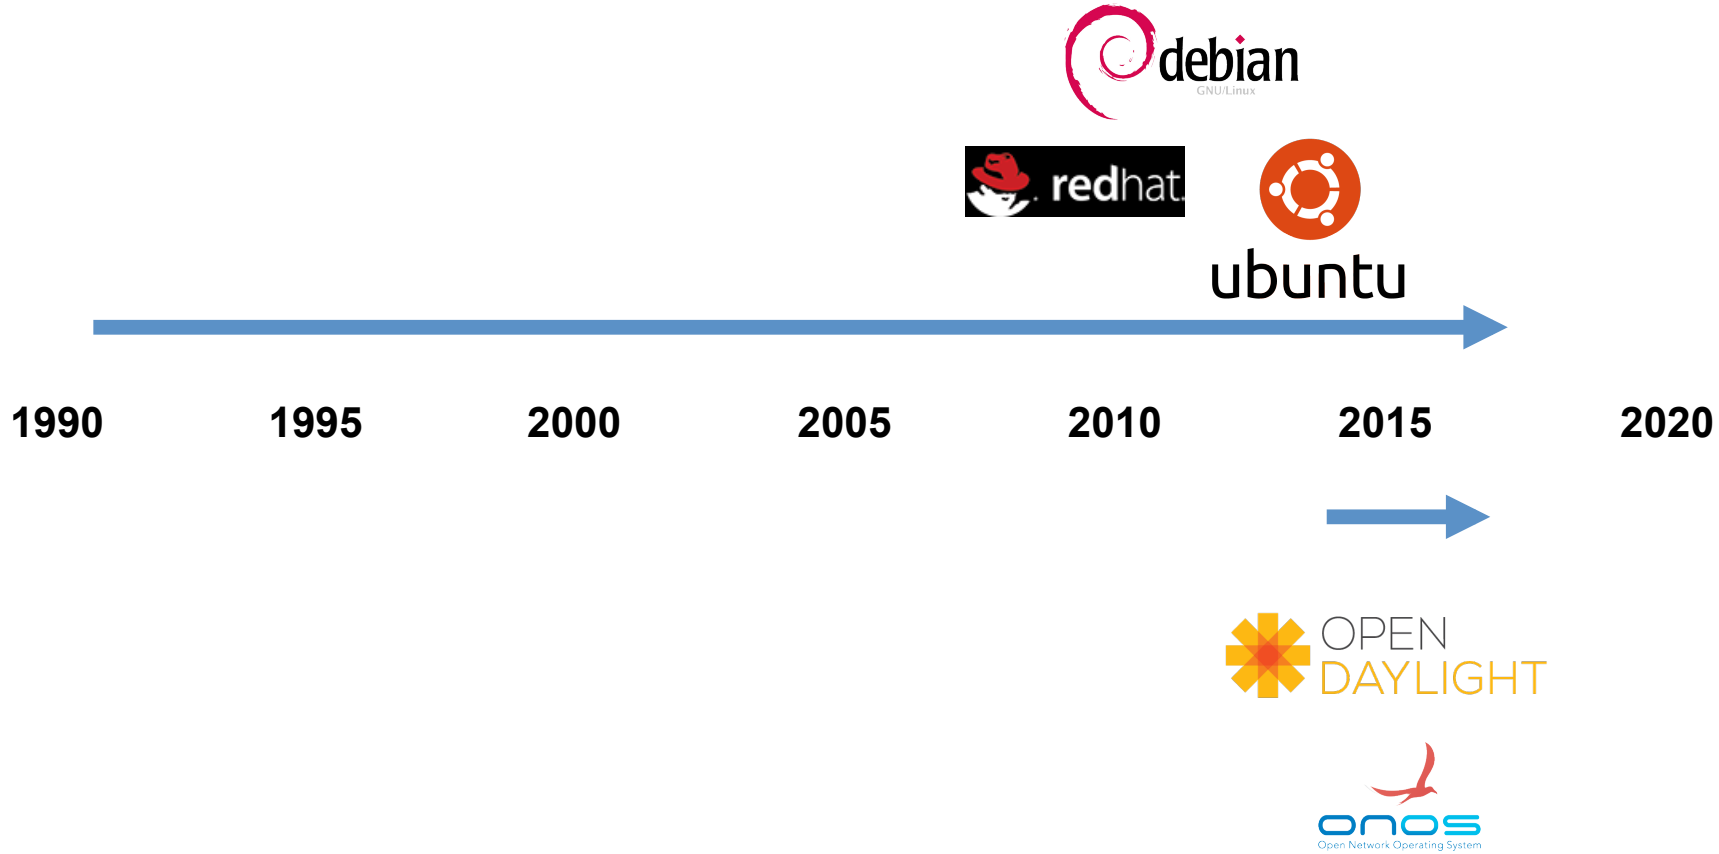

Thank you. Questions?

Contacts

Stefano Salsano

University of Rome Tor Vergata / CNIT

stefano.salsano@uniroma2.it

The SUPERFLUIDITY project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No.671566 (Research and Innovation Action).

GN4-1 and GN4-2 projects received funding from the European Union's Horizon 2020 research and innovation programme respectively under Grant Agreements No. 691567 and No. 731122.

The information given is the author's view and does not necessarily represent the view of the European Commission (EC). No liability is accepted for any use that may be made of the information contained.

Where is my orchestrator ??

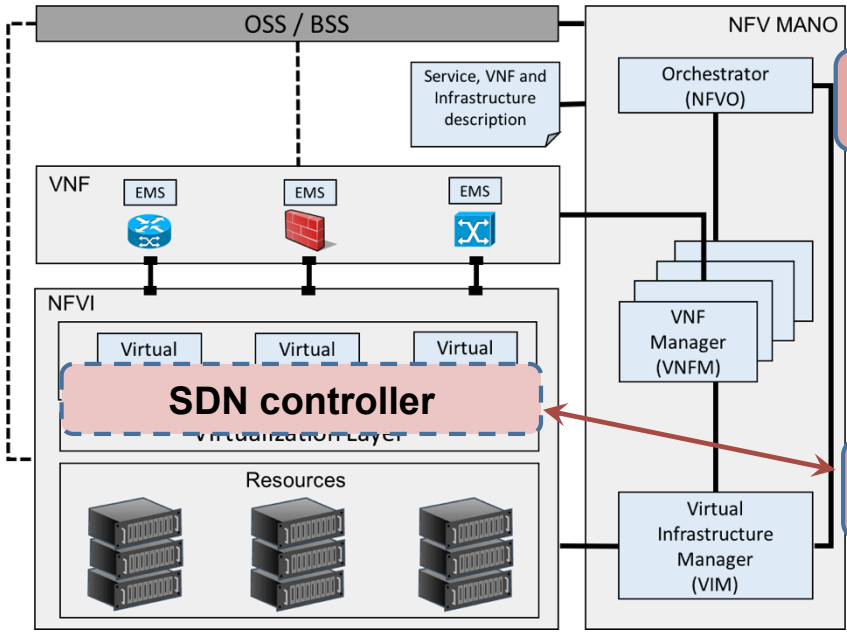

OpenDayLight
ONOS

NFV Orchestrator

OpenO
OpenBaton
OSM
Aria/Cloudify
…

VIM

OpenStack
30 mainstream projects (Heat, Neutron, Magnum, …)

Unikernels (ClickOS) memory footprint and boot time

Memory footprint

- Hello world guest VM : 296 KB
- Ponger (ping responder) guest VM : ~700KB

Boot time, state of the art results

VM configuration: MiniOS, 1 VCPU, 8MB RAM, 1 VIF

- 87.77 msecs

Recent results (from SUPERFLUIDITY), by redesigning the toolstack

- Without libxl: 6.67 msecs
- Without xen store: 1.43 ms

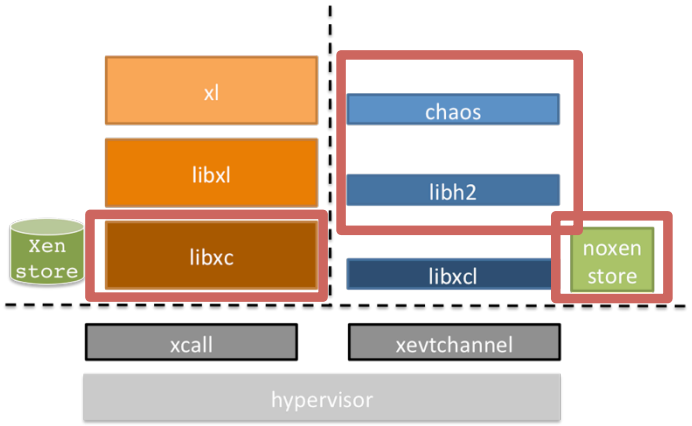

Unikernels (ClickOS) memory footprint and boot time

Memory footprint

- Hello world guest VM : 296 KB
- Ponger (ping responder) guest VM : ~700KB

Boot time, state of the art results

VM configuration: MiniOS, 1 VCPU, 8MB RAM, 1 VIF

- 87.77 msecs

Recent results (from SUPERFLUIDITY), by redesigning the toolstack

- Without libxl: 6.67 msecs
- Without xen store: 1.43 ms

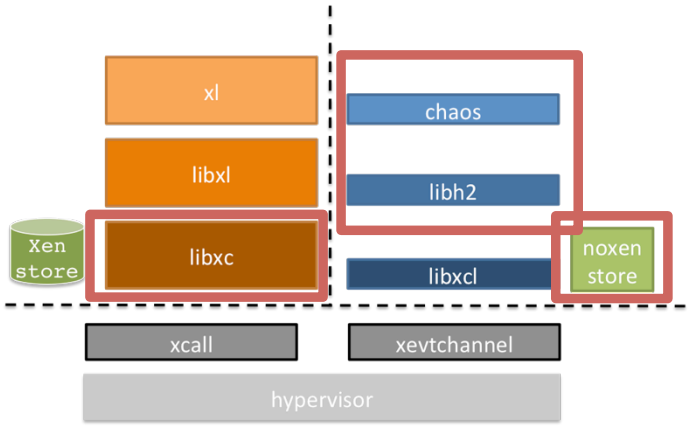

ETSI MANagement and Orchestration (MANO) Model

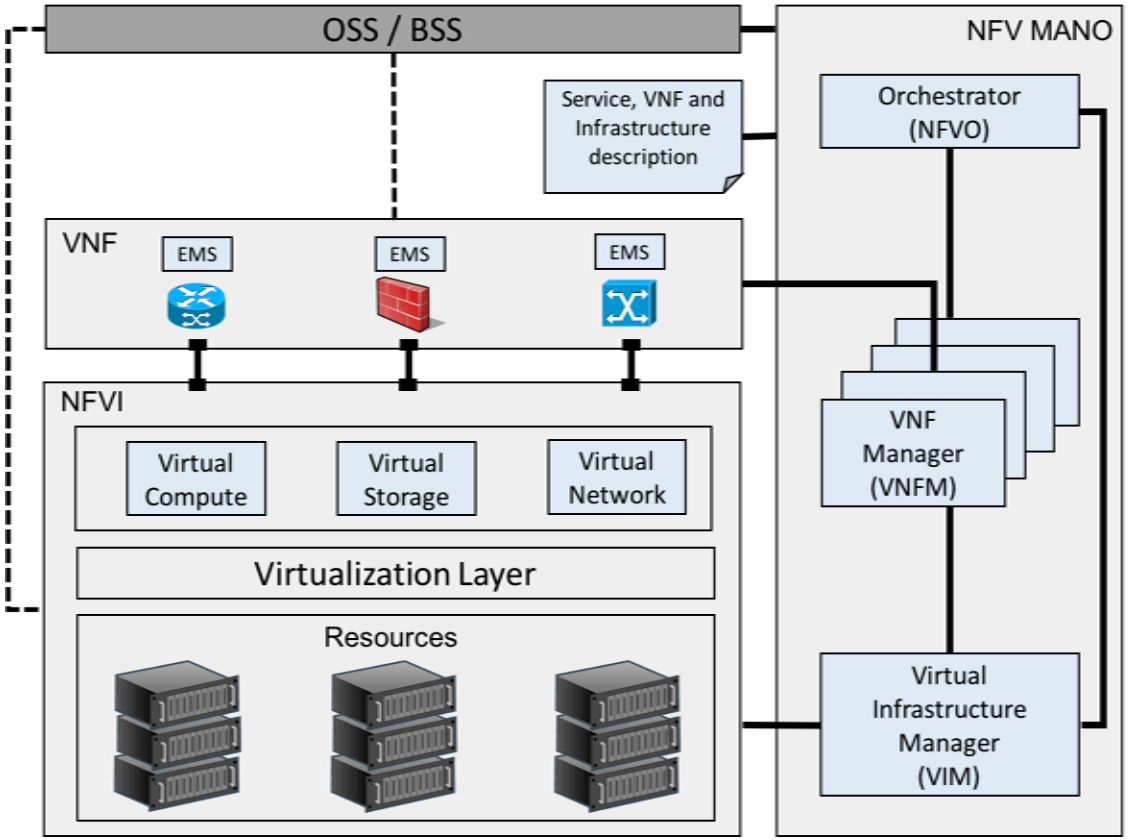

VM instantiation and boot time typical performance (no Unikernels)

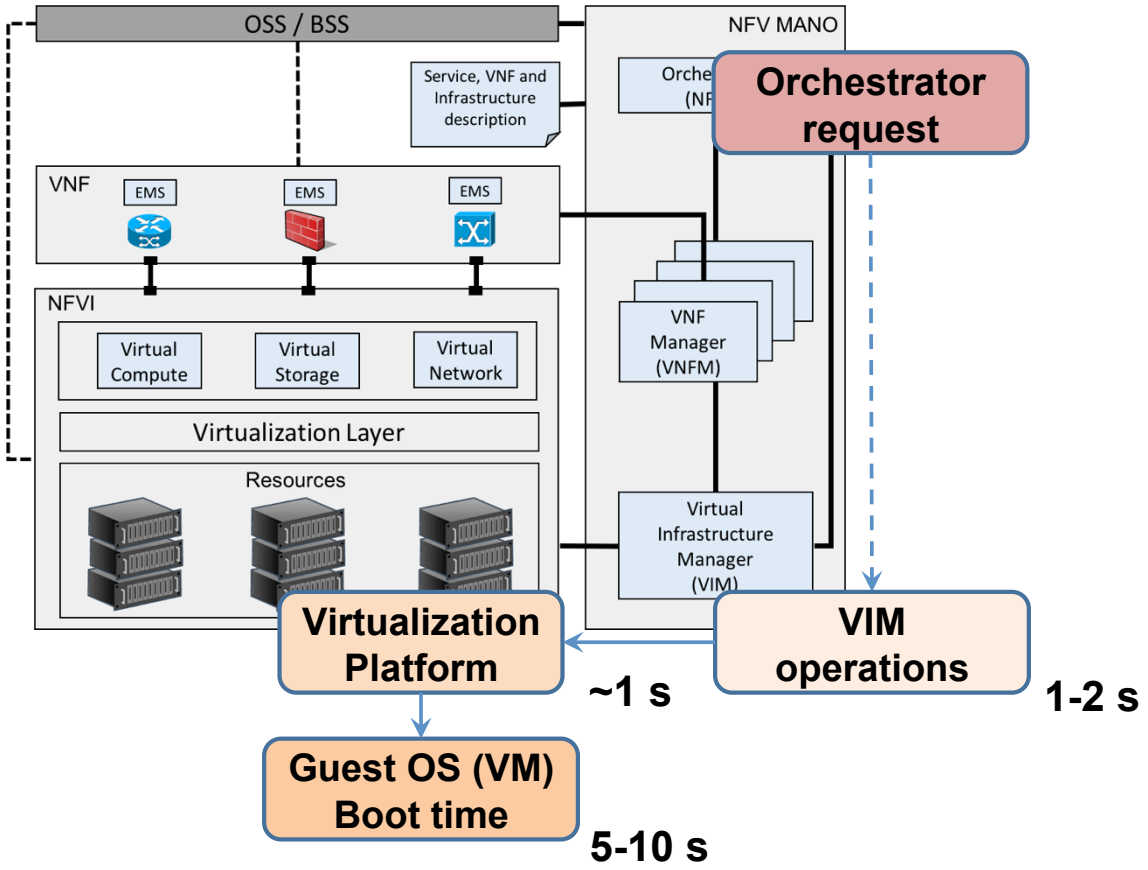

VM instantiation and boot time (using Unikernels)

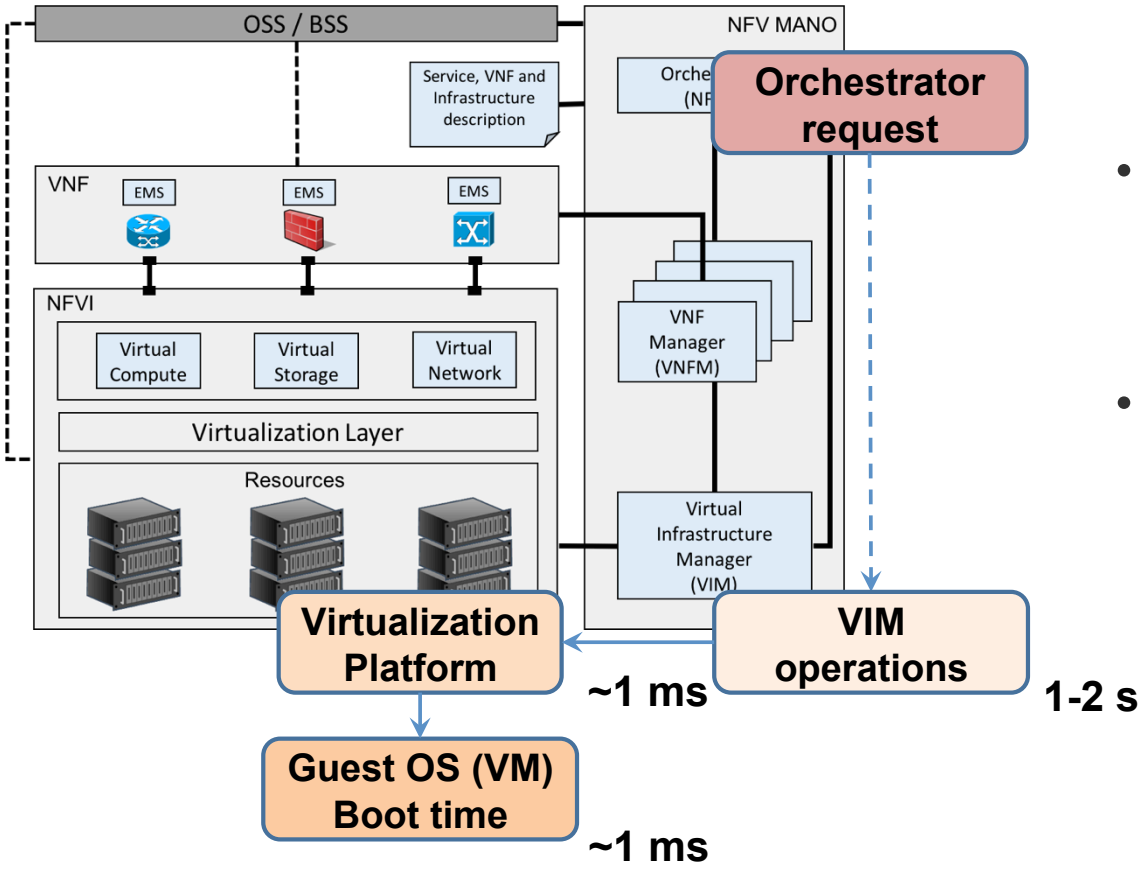

- Unikernels can provide low latency instantiation times for "Micro-VNF"
- What about VIMs (Virtual Infrastructure Managers) ?

Performance analysis and Tuning of Virtual Infrastructure Managers (VIMs) for Unikernel VNFs

- We considered 3 VIMs (Openstack, Nomad, OpenVIM)

- General model of the VNF instantiation process
- Modifications to VIMs to instantiate Micro-VNFs based on ClickOS Unikernel
- Performance Evaluation

OpenVIM model

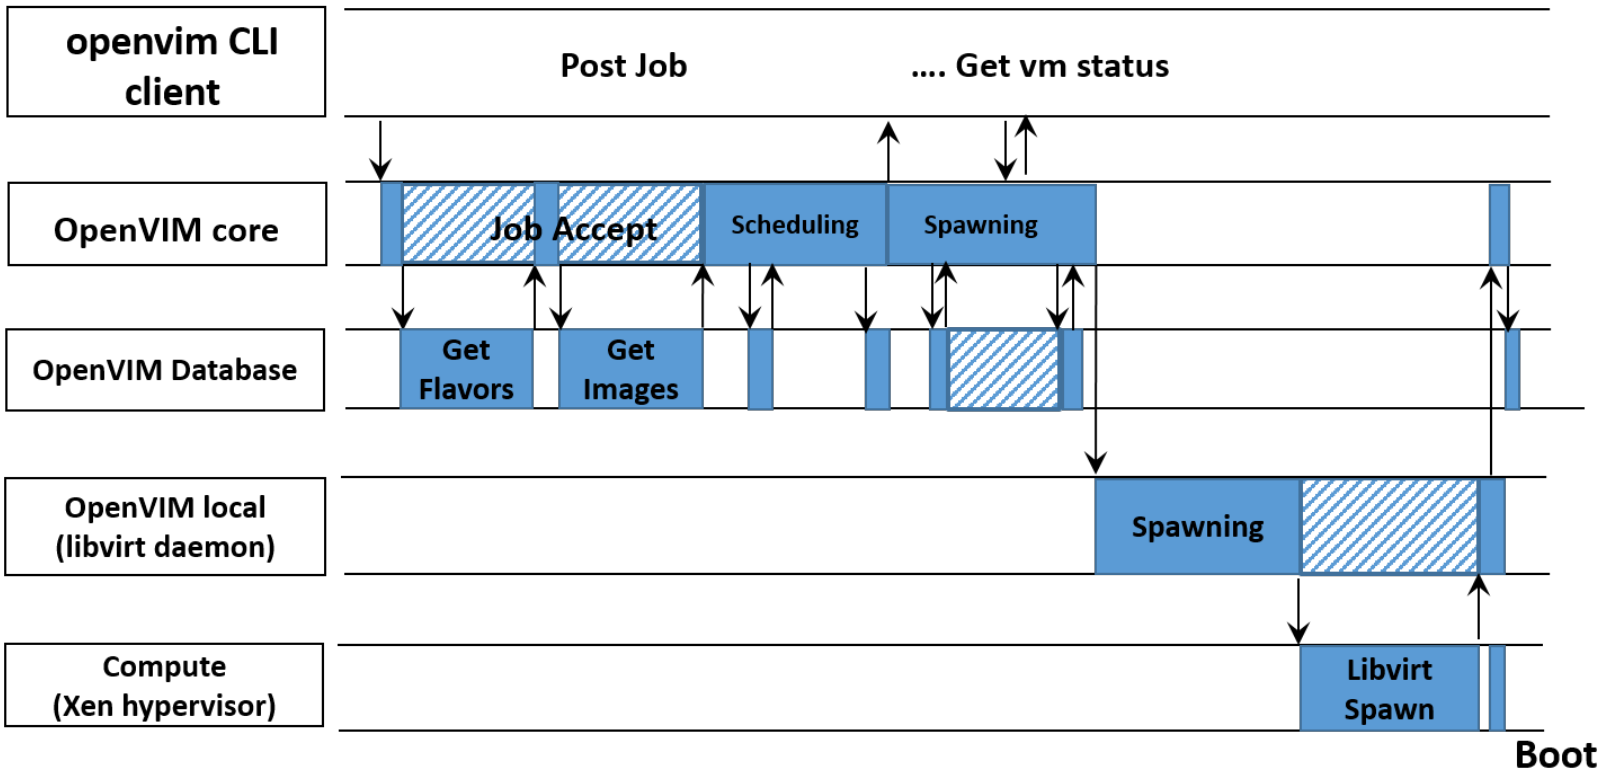

Supporting Unikernel in OpenVIM

OpenVIM does not have a native support to Unikernels

- Added primitive support to Para Virtualized Hypervisor such as Xen
- Changed the code to manage the new images and configurations

Results – ClickOS instantiation times (OpenStack, Nomad)

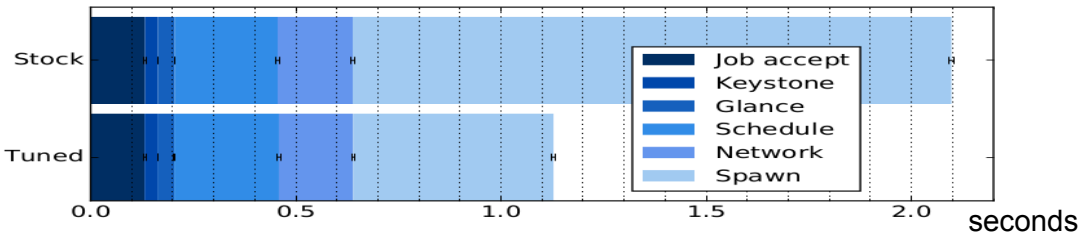

OpenStack Nova

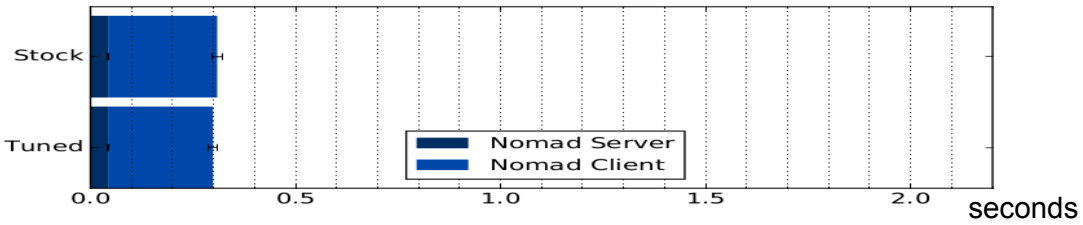

Nomad

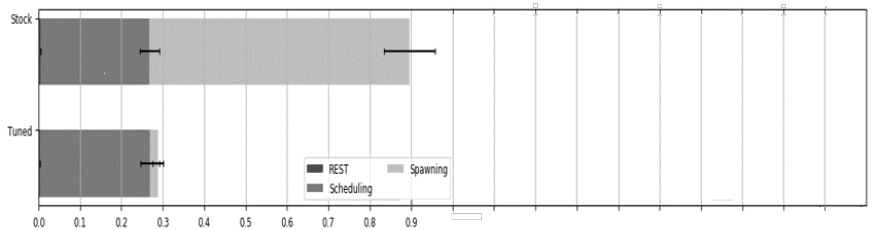

OpenVIM

Chaining of Unikernel VNFs

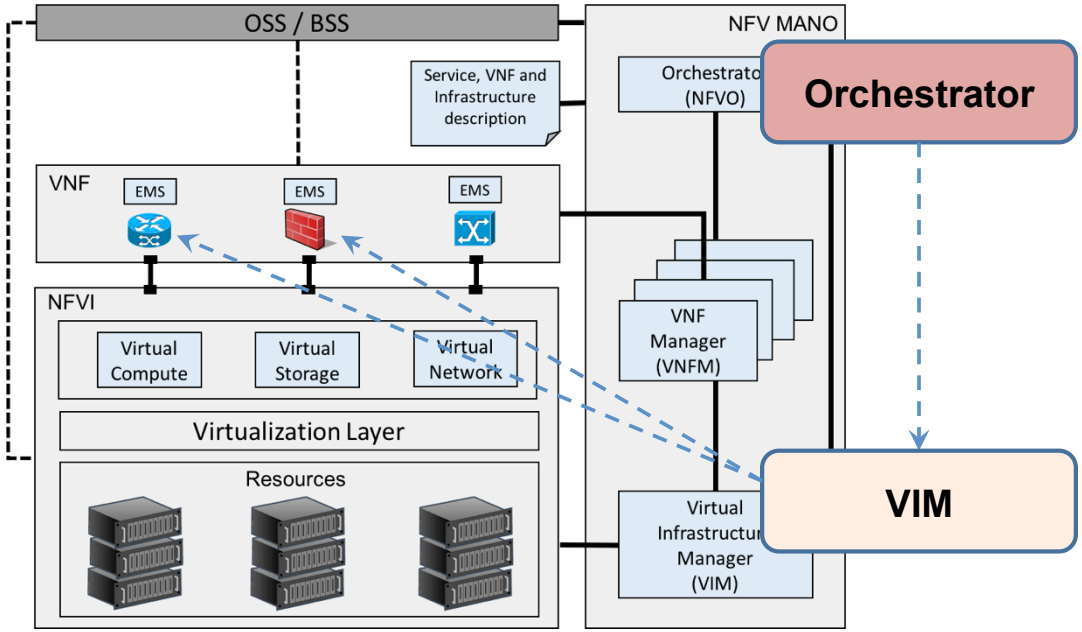

So far, we have discussed how the Unikernel VMs can be instantiated by the Orchestrator/VIM…

Ongoing cooperation with CISCO

- We are cooperating on this topics with the CISCO IPv6 Segment Routing team led by Clarence Filsfils.
- Joint demo planned at MPLS World Congress in Paris (21/24 March 2017)
- Joint paper on the Linux based solution:
  A. AbdelSalam, F. Clad, C. Filsfils, S. Salsano, G. Siracusano, L. Veltri, “Implementation of Virtual Network FunctionChaining through Segment Routingin a Linux-based NFV Infrastructure”
  IEEE Netsoft 2017, Bologna, Italy, https://arxiv.org/abs/1702.05157
- Co-authored IETF draft:
  C. Fisfils, J. Leddy, D. Voyer, D. Bernier, D. Steinberg, R. Raszuk, S. Matsushima, D. Lebrun, B. Decraene, B. Peirens, S. Salsano, G. Naik, H. Elmalky, P. Jonnalagadda, M. Sharif, A. Ayyangar, S. Mynam, A. Bashandy, K. Raza, D. Dukes, F. Clad, P. Camarillo (Ed.), "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-00, IETF draft, March 2017

References & papers for download

- SUPERFLUIDITY project Home Page http://superfluidity.eu/

- G. Bianchi, et al. “Superfluidity: a flexible functional architecture for 5G networks”, Transactions on Emerging Telecommunications Technologies 27, no. 9, Sep 2016

- P. L. Ventre, C. Pisa, S. Salsano, G. Siracusano, F. Schmidt, P. Lungaroni, N. Blefari-Melazzi, “Performance Evaluation and Tuning of Virtual Infrastructure Managers for (Micro) Virtual Network Functions”, IEEE NFV-SDN Conference, Palo Alto, USA, 7-9 November 2016
  http://netgroup.uniroma2.it/Stefano_Salsano/papers/salsano-ieee-nfv-sdn-2016-vim-performance-for-unikernels.pdf

Unikernel virtualization in the SUPERFLUIDITY vision

- We have considered the optimization of Unikernel virtualization and the needed enhancements to Virtual Infrastructure Managers to support Unikernels.
- In the SUPERFLUIDITY vision, Unikernels are interesting as they support the decomposition of network services in “smaller” components that can be deployed on the fly (NB: Unikernels are complementary to other approaches!)
- The NFV Infrastructure should be extended in order to support Unikernel virtualization in addition to traditional VMs. This way it will be possible to design services that exploit the most efficient solutions depending on several factors.

Conclusions

- Unikernel virtualization can provide VM instantiation and boot time in the order of ms, VM memory footprints in the order of few MB

- Work is needed at the level of Virtual Infrastructure Managers
  - e.g. OpenStack (~ 1 s), Nomad (~ 300 ms), OpenVIM (~ 300 ms)

- IPv6 Segment Routing offers a perfect match to the need of dynamic "chaining" of Unikernel VMs
  - Configuring the chain only at the edge speeds up the (re)configuration
  - The amount of state information maintained in the noted is greatly reduced

References – Speed up of Virtualization Platforms / Guests

- J. Martins, M. Ahmed, C. Raiciu, V. Olteanu, M. Honda, R. Bifulco, F. Huici, "ClickOS and the art of network function virtualization", NSDI 2014, *11th USENIX Conference on Networked Systems Design and Implementation,* 2014.

- F. Manco, J. Martins, K. Yasukata, J. Mendes, S. Kuenzer, F. Huici, "The Case for the Superfluid Cloud", 7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 15), 2015

- Recent unpublished results are included in this presentation:
  S. Salsano, F. Huici, "Superfluid NFV: VMs and Virtual Infrastructure Managers speed-up for instantaneous service instantiation", invited talk at EWSDN 2016 workshop, 10 October 2016, The Hague, Netherlands
  http://www.slideshare.net/stefanosalsano/superfluid-nfv-vms-and-virtual-infrastructure-managers-speedup-for-instantaneous-service-instantiation