



Research trends in abstraction of networks and orchestration of network services

Prasad Gorja | Principal Staff Systems & Architecture Engineer

prasad.gorja@freescale.com

JAN . 31 . 2014



External Use

Freescale, the Freescale logo, AltVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. Airfast, BeeKit, BeeStack, CoreNet, Flexis, Layerscape, MagnIV, MXC, Platform in a Package, QorIQ Converge, QUICC Engine, Ready Play, SMARTMOS, Tower, TurboLink, UMEMS, Hybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2014 Freescale Semiconductor, Inc.



Agenda

Programmable Networks

- Limitations of ossified networks
- Need for abstraction of Networks/protocols
- Software Defined Networking (SDN)
- OpenFlow
- Open vSwitch
- Network Function Virtualization(NFV)

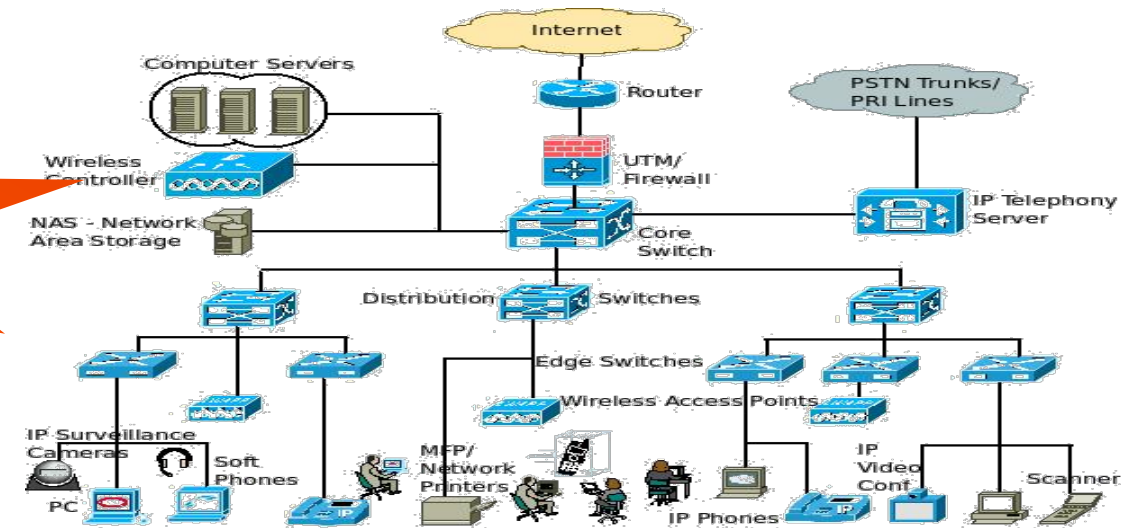
Orchestration of Network Services

- Cloud orchestration landscape
- OpenStack Architecture
- Neutron Plug-ins
- Interactions of OpenStack for NFV & SDN

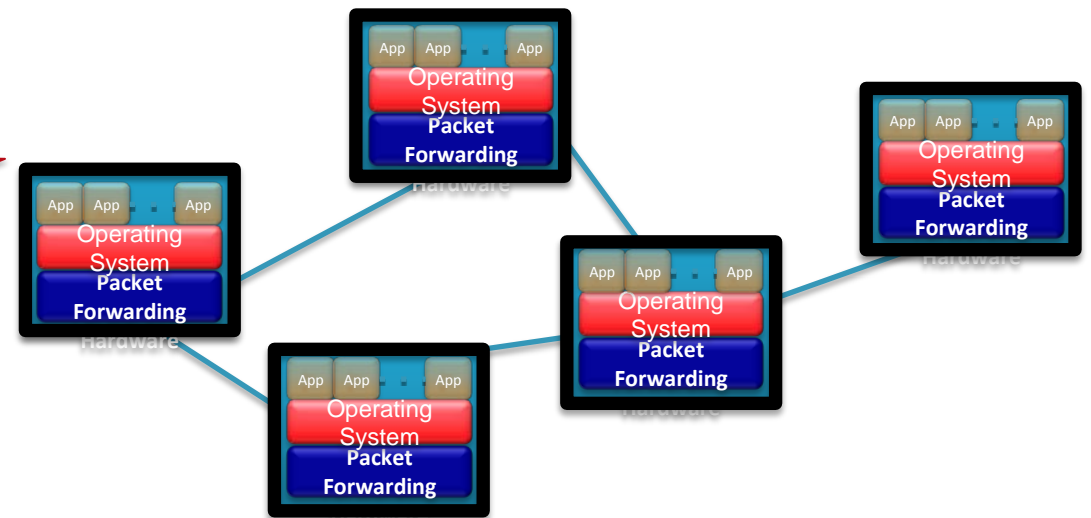
Future work

Limitations of Current Networking Technologies

- Lack of Abstraction
- Static Networks
- Inconsistent policies

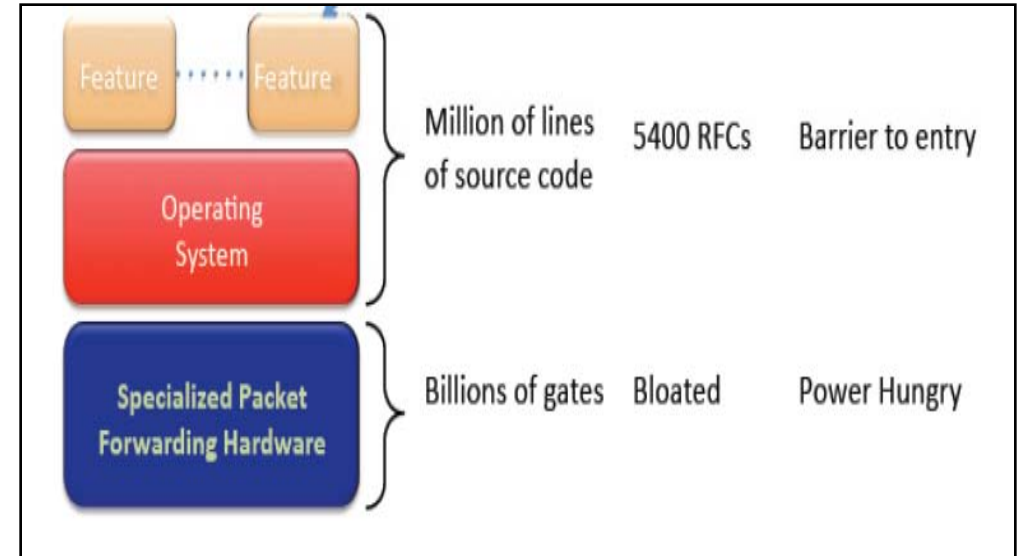


- Proprietary Implementation
- Collaboration among vendors?
- Lack of open & standard interface
- Limited innovation

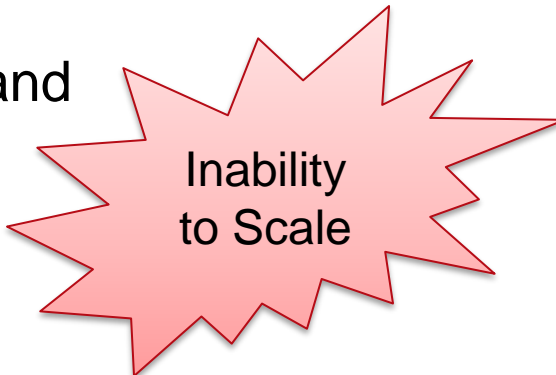


Limitations of Current Networking Technologies


- Complex functions baked into infrastructure
- Routing, traffic monitoring, firewall, NAT, IPSec, IPS/IDS ...



- Exploding Traffic Patterns and data-set exchanges
- High performance and low cost connectivity



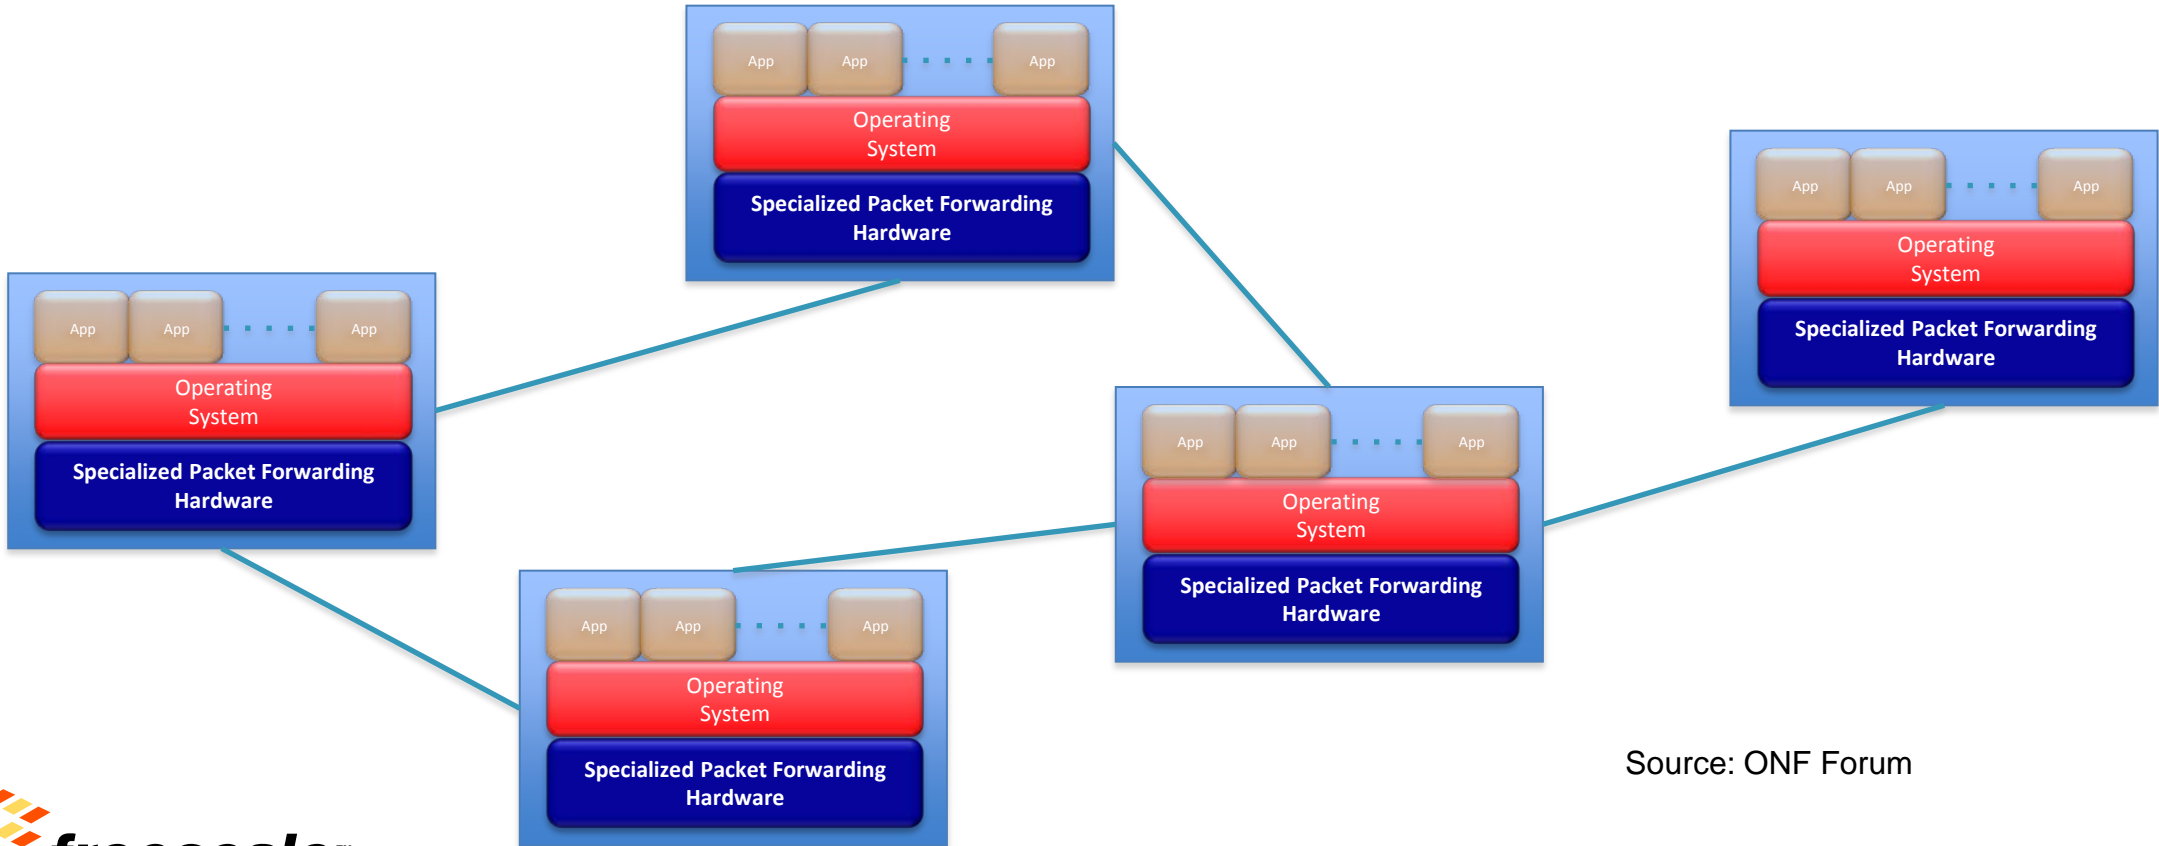
Multi-tenancy Options



An OS for Networks

Control Programs

Network Operating System



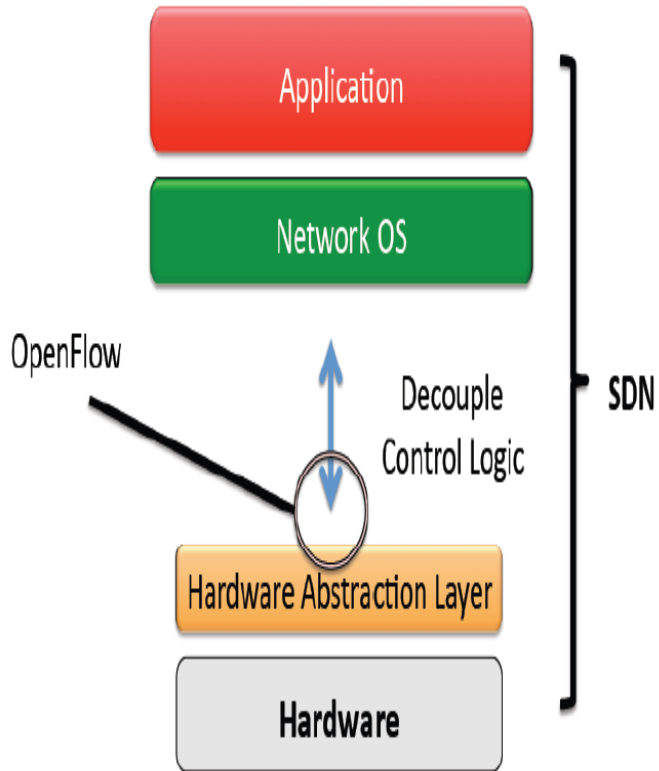
Source: ONF Forum

Need for New Architecture

- Program vs Configure a Network
- Speed-to-Market and faster upgrade for new enhancements
- Ease of customization and integration
- Open Interfaces and facilitate Innovation in Network
- More flexibility with dynamic networking
- Layered Architecture and Innovation at each layer



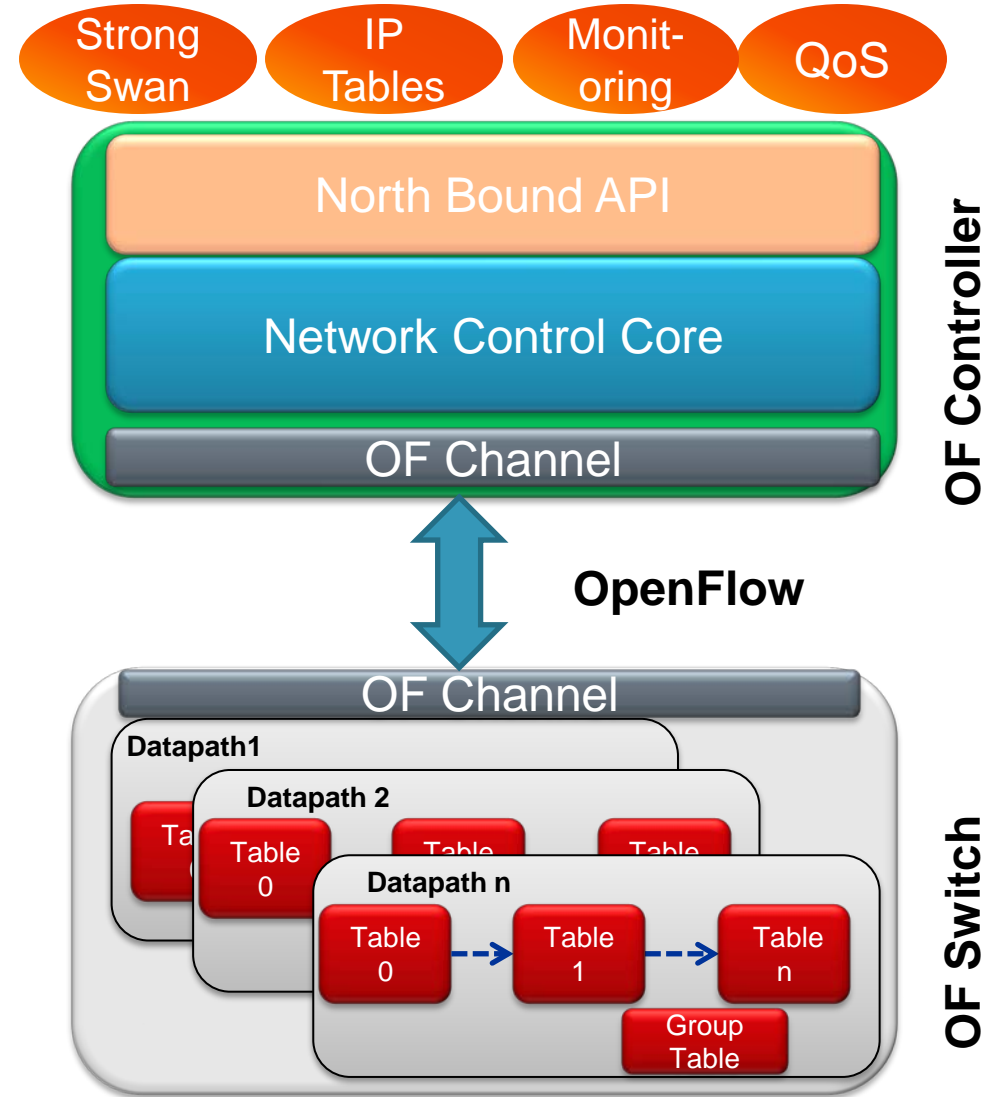
Software Defined Networking



- OpenFlow Controller
 - OpenFlow Channel
 - Network control Core
 - North bound API

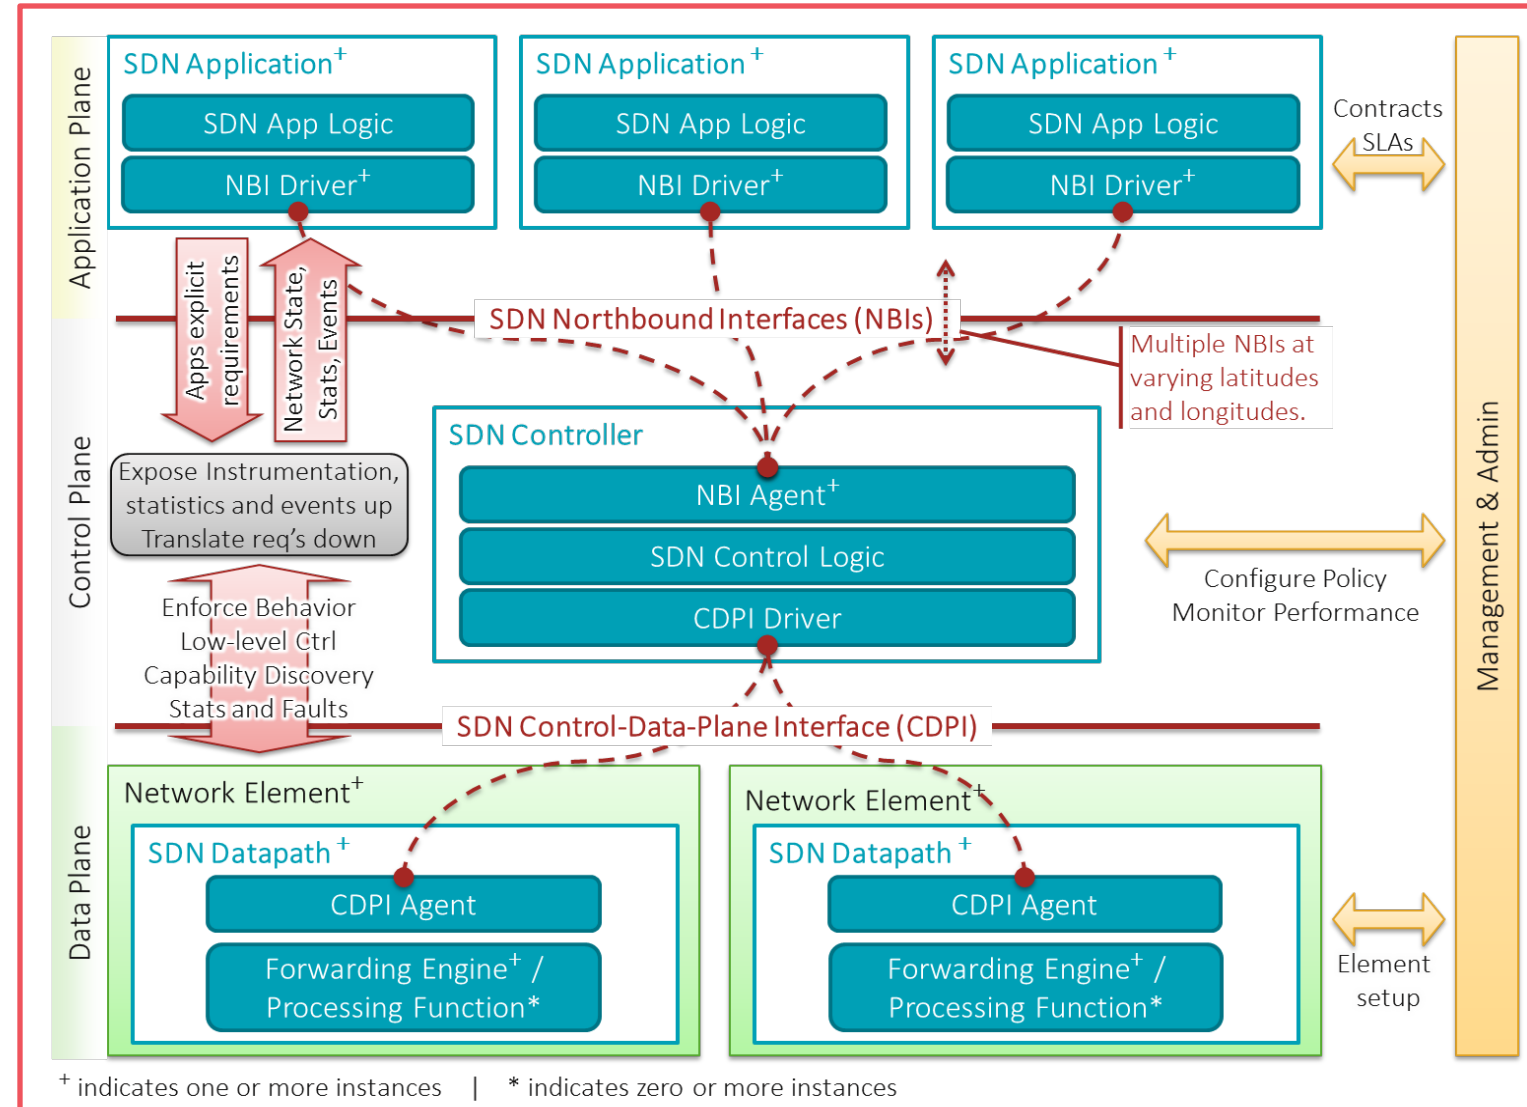
- OpenFlow Switch
 - OpenFlow Channel
 - One or more Flow tables
 - A group table

Source: ONF Forum



Software Defined Networking

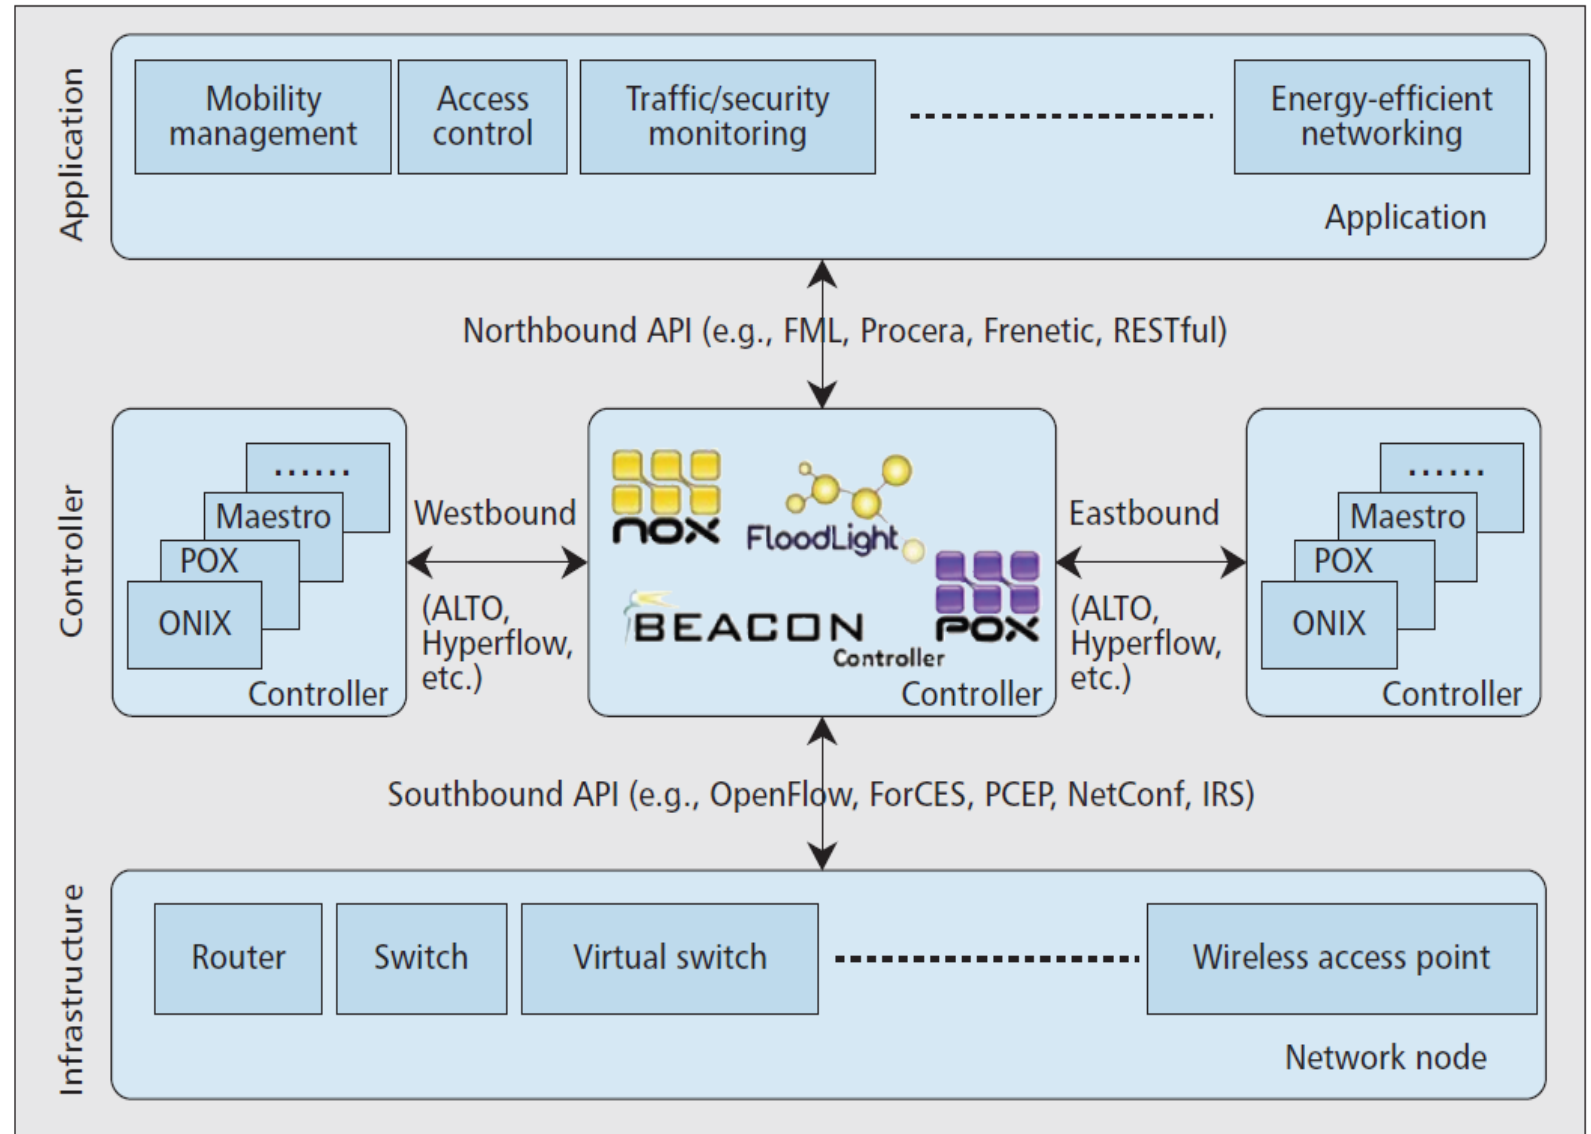
- Centralized control of multi-vendor environment
 - Control any vendor SDN Switch
 - Facilitate on-fly deploy and configure devices across network
- Reduced complexity of automation
 - Reduced operational overhead and Service disruption
 - IT-as-a-service and self-provisioning models
- Open for higher rate of Innovation
- More granular network control
 - Levels: Application, device, user and session



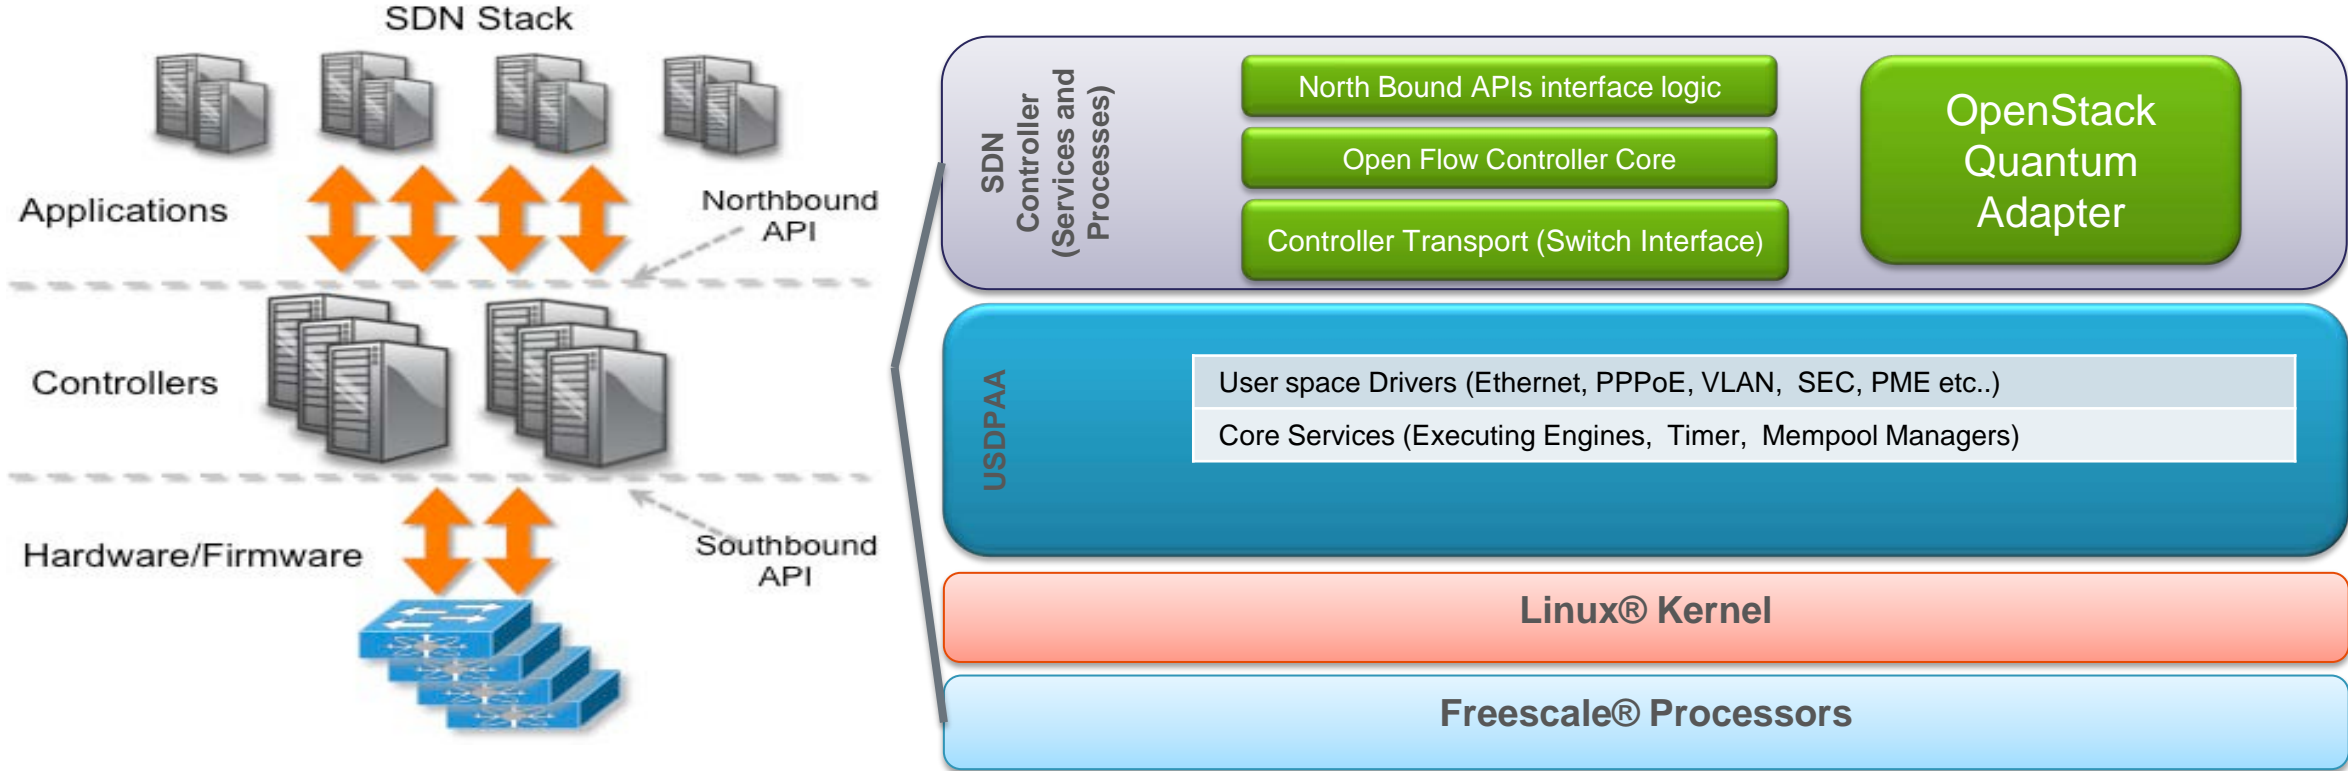
Source: ONF

SDN Interfaces and Enabling Technologies

- Northbound API
 - Applications with controller
- Southbound API
 - Datapath with controller
- Eastbound API
 - Among controllers
- Westbound API
 - Among controllers



OpenFlow Controller

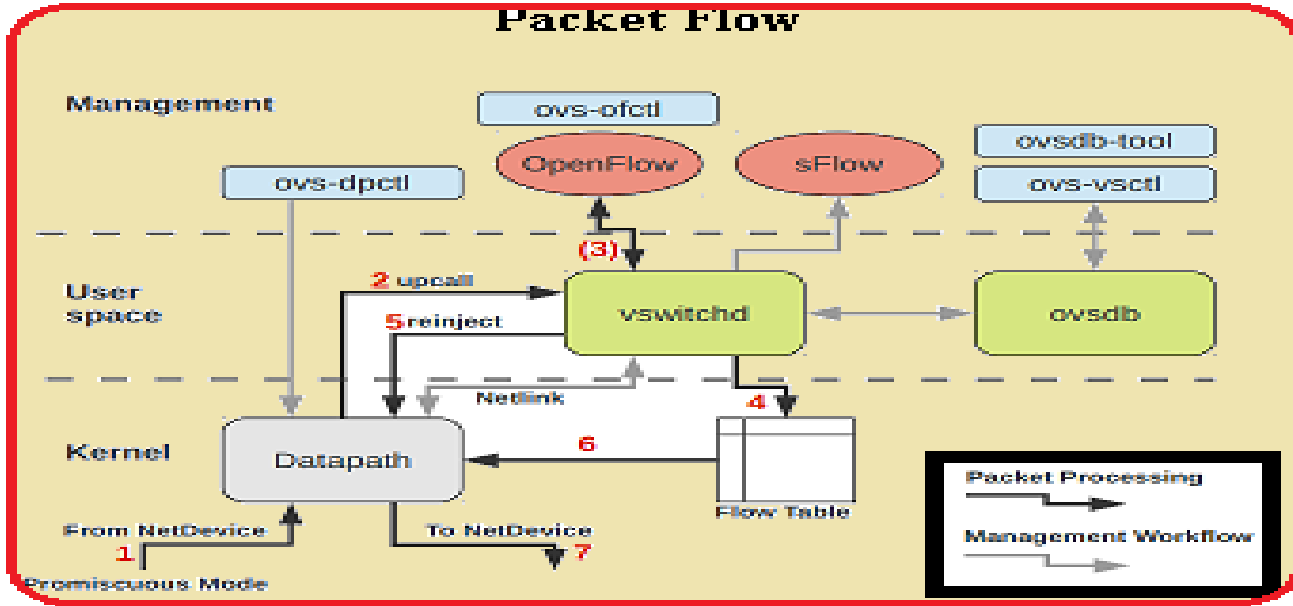
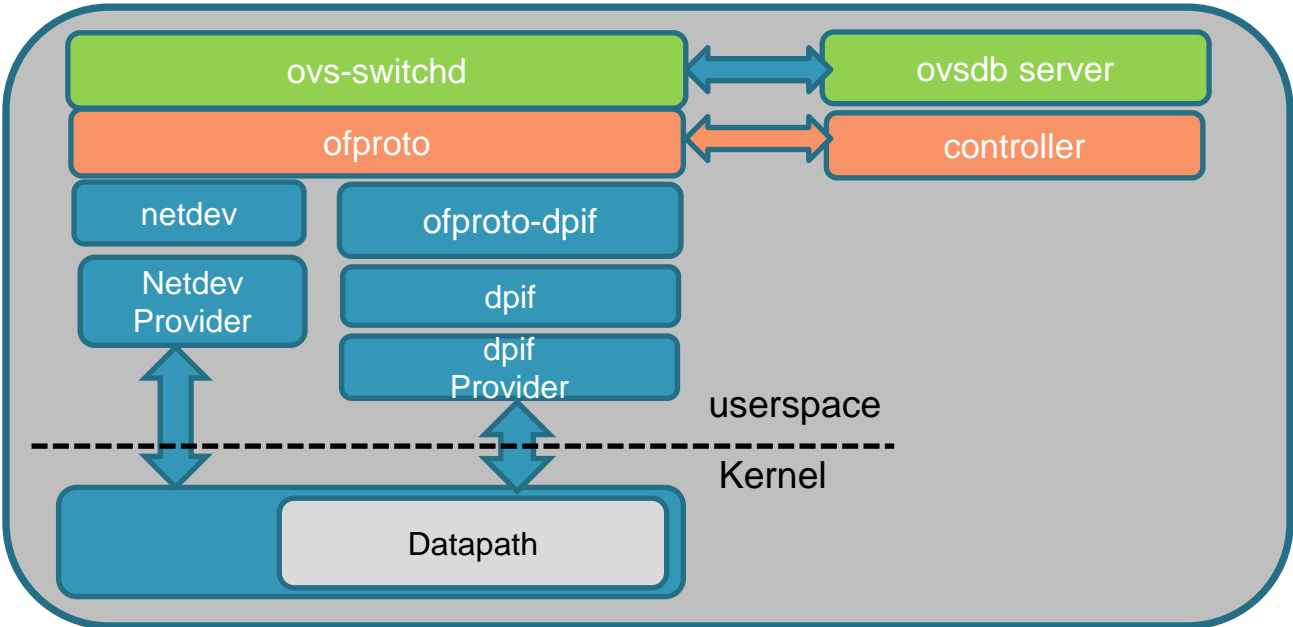


Freescale OpenFlow controller is available at <http://git.freescale.com/git/cgit.cgi/nfv/of-controller.git/>

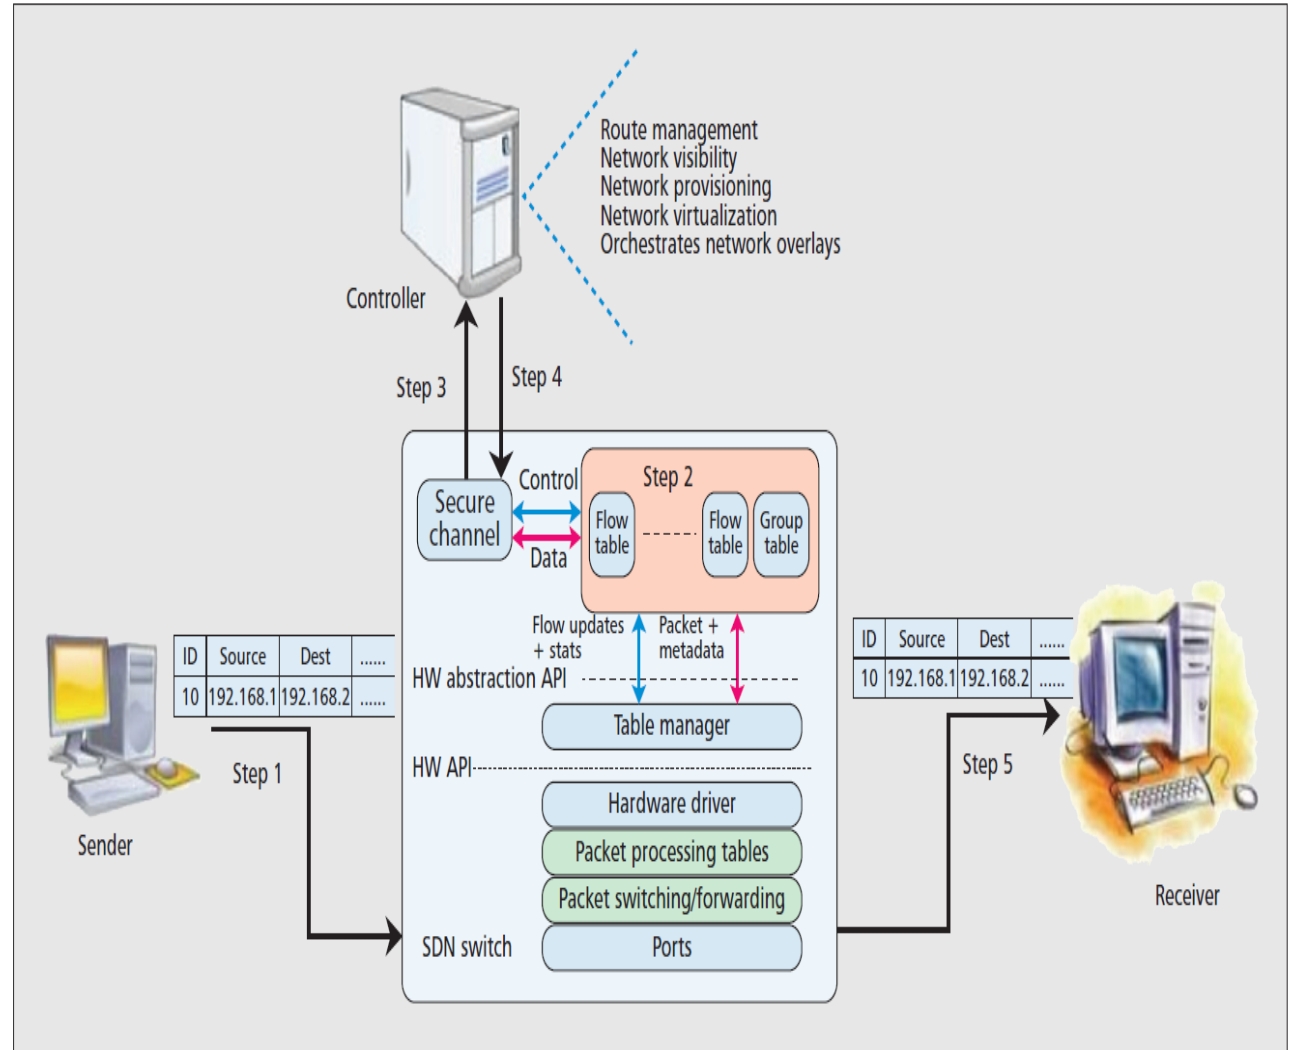
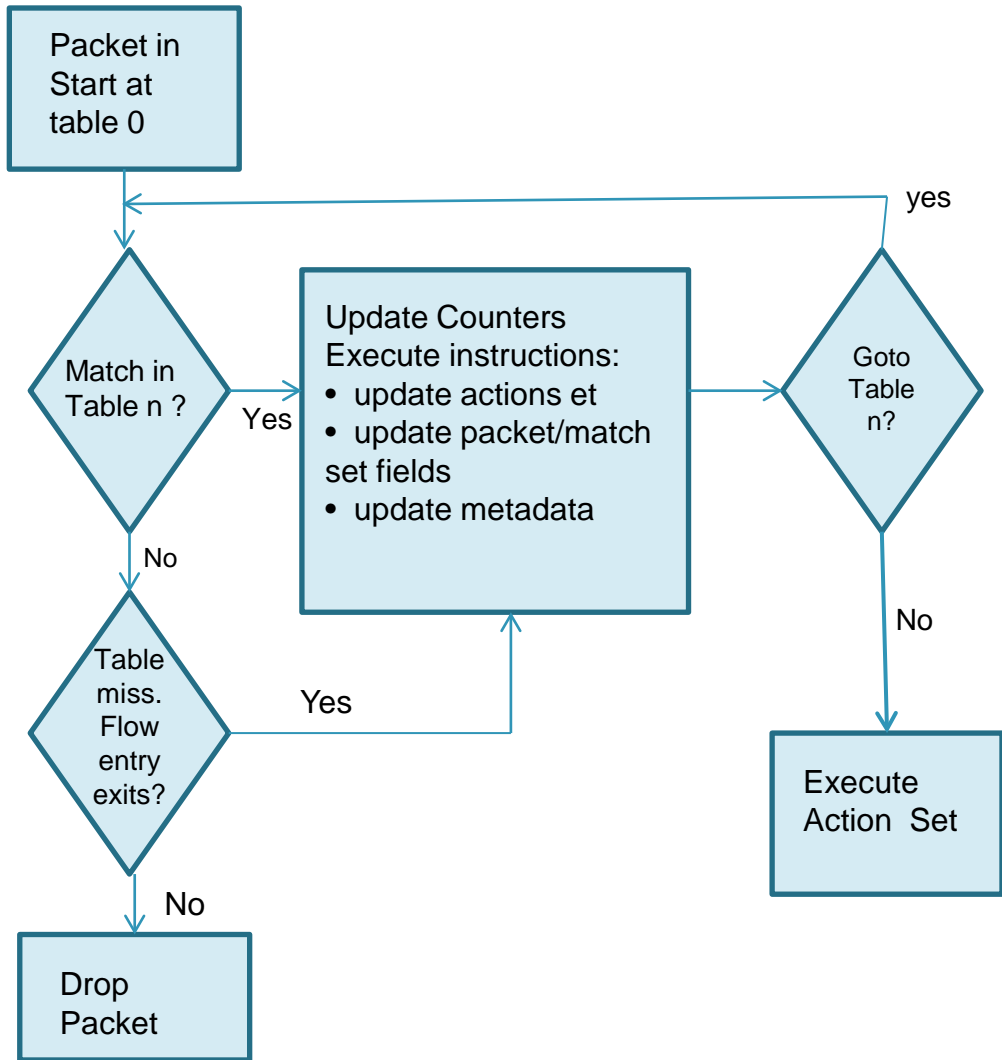
Open vSwitch(OVS)

- Default Switch(Bridge) for Linux-based Hypervisors
- OpenFlow supported
- Forwarding path is designed to be amenable for offloading to classic hardware switch or an end-host NIC.
- GRE tunneling and VxLAN support for datacenter segments
- To respond Network dynamics in virtual environments
 - VMs come and go which results in changes to logical network environment
 - Mobility of the soft state while VM migration

<http://openvswitch.org/>

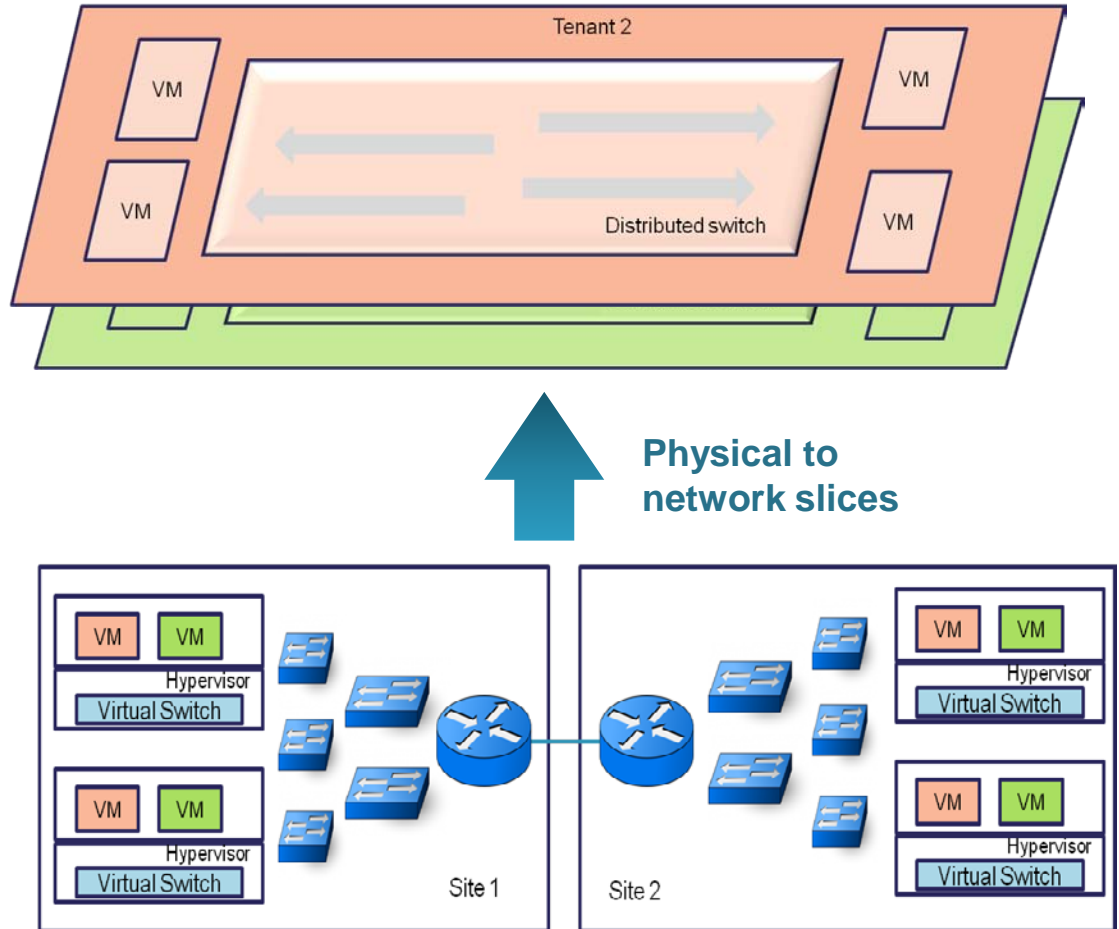


Operation of SDN OF(Controller- Switch)




Network Virtualization

- Need—Similar to Compute and Storage Virtualization
 - On demand provisioning and elasticity
 - Faster provisioning
- Current Technology
 - VLAN (every tenant that signs up to CSP is assigned with set of VLANs)
 - Limited to 4K
- Trend: Overlay Technologies
 - Compute VM traffic (L2 traffic) sent over any network, including L3 by encapsulation
 - VxLAN™(Virtual Extensible LAN) *
 - NVGRE™ (Network Virtualization using GRE)
 - STT (Stateless Transport Tunneling)



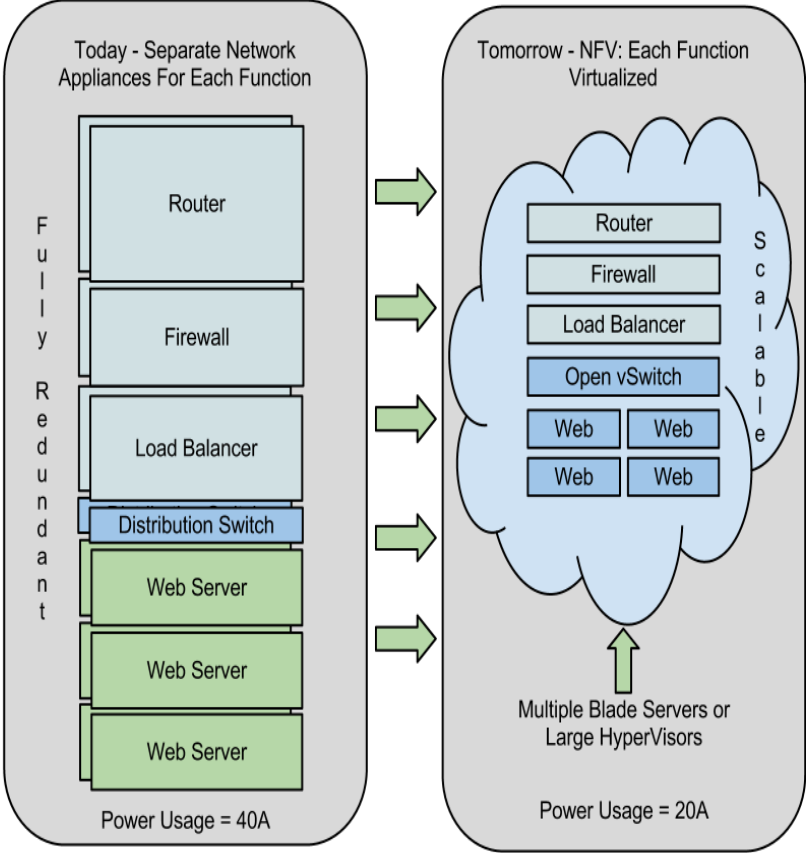
Challenges – Research Trends in SDN

- Standards are evolving and availability of ecosystem
- Distributed multi controllers and co-ordination among controllers
- Northbound API standardization
- Performance of the OpenFlow Switch as table look ups are expensive in table driven datapath
- New extensions for L4-L7 application
- Widely used Open vSwitch yet to support higher level protocols by using OpenFlow
- Configuration of Switch profile and OF config standard is evolving
- Interoperability and Testing



Network Function Virtualization(NFV)™

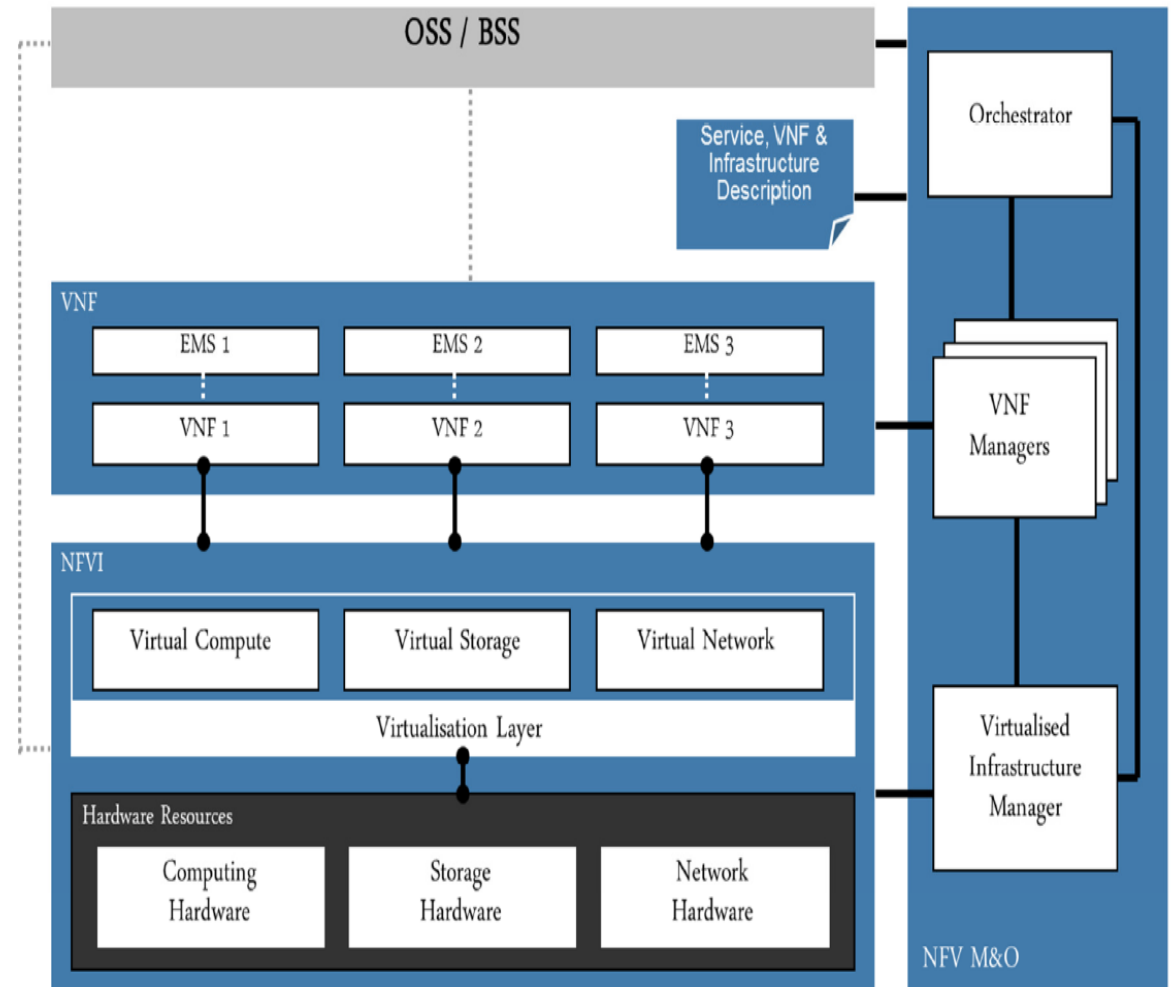
- Each Network function as virtual appliance
- Very Flexible
- Reduced CapEx and OpEx
- On-demand service scale up/down
- Single platform for different apps
- Open for ecosystems and partners



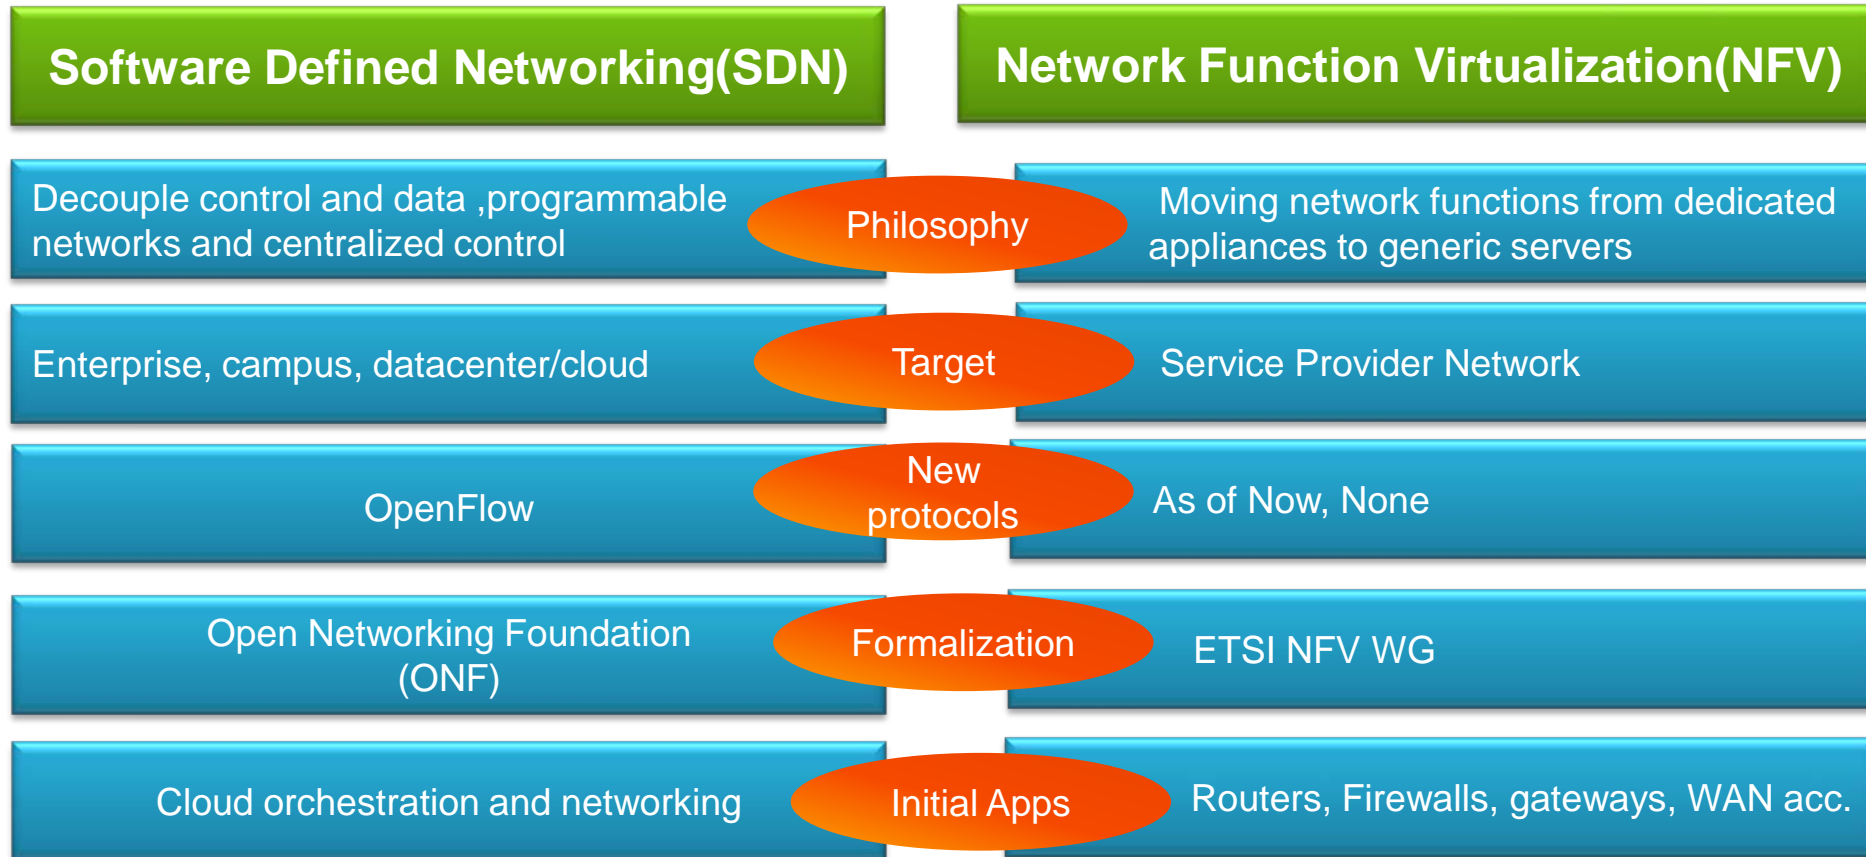
Source: ETSI Network Function Virtualization

NFV Architecture

- Network Function Virtualization Infrastructure (NFVI)
 - ✓ Resources to support the execution of VNF
- Virtual Network Functions(VNF)
- NFV M& O (Management & Orchestration)



SDN vs. NFV



Source: ETSI Network Function Virtualization

NFV Challenges

- Standards are evolving and availability of ecosystem
- Performance
- SDN with NFV
- Industry Acceptance






Orchestration of Network Services

OpenStack™ Orchestration

- Cloud Computing
 - ✓ Deliver applications and services over Internet
 - ✓ make use of **shared** resources and provides **elastic** resources for tenants
- Orchestration provides automatic configuration, **control** and **provisioning** of resources
 - ✓ Computes
 - ✓ Storage
 - ✓ Network



EUCALYPTUS

OpenNebula.org
The Open Source Toolkit for Cloud Computing

openstack™
CLOUD SOFTWARE

cloudstack
open source cloud computing

freescale™



OpenStack™ Architecture

Dash Board (Horizon)

- ✓ Web based portal

Compute(Nova)

- ✓ Spawning, Scheduling, decomposition of machines
- ✓ Manages life cycle of compute instance

Network (Neutron)

- ✓ Network connectivity and define networks

Object Storage (Swift)

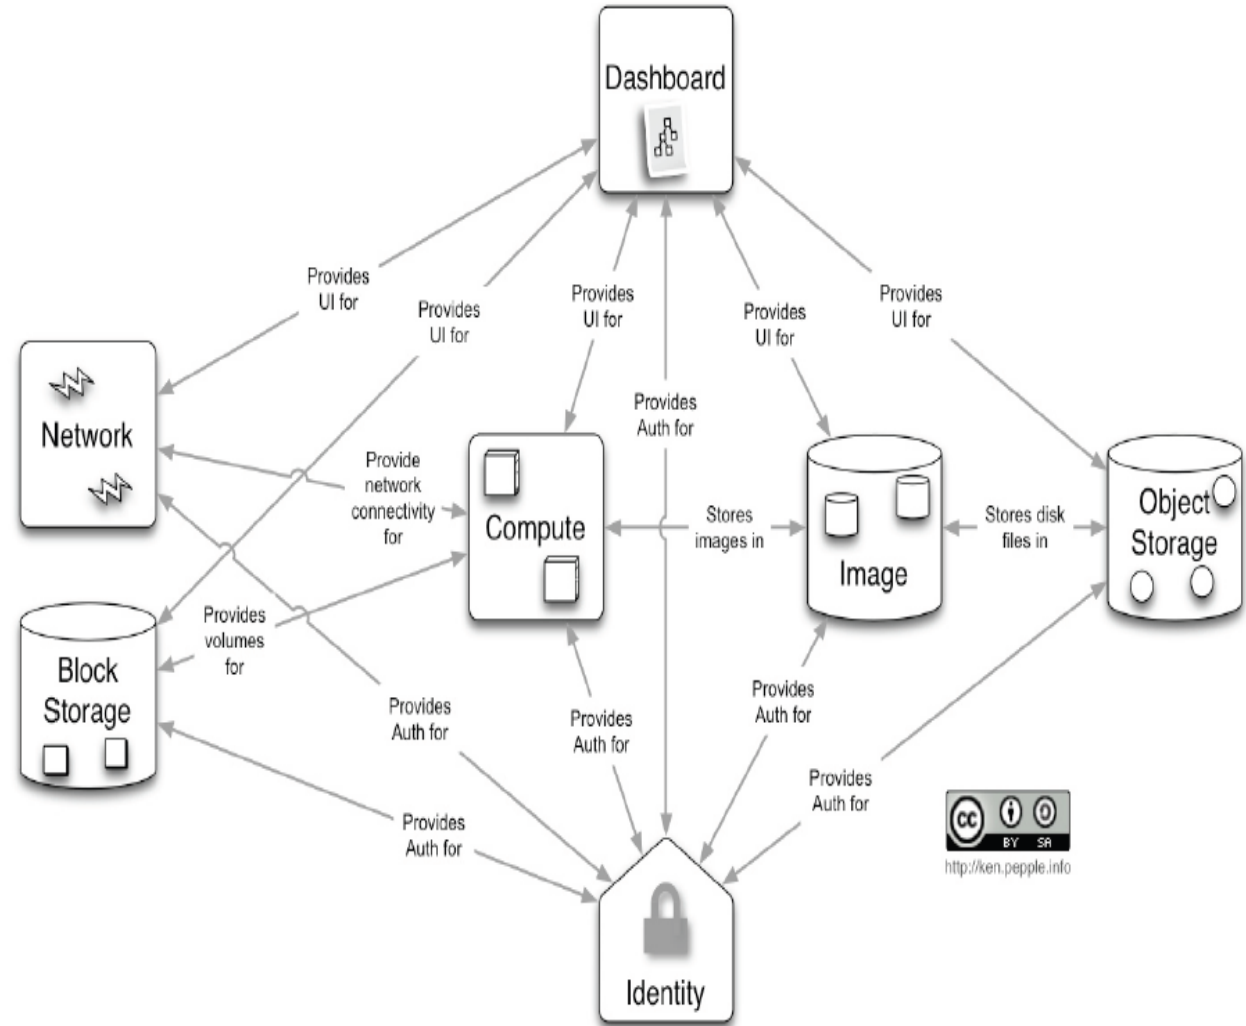
- ✓ Stores and retrieves arbitrary unstructured data

Block Storage (Cinder)

- ✓ Persistent block storage for running instance

Identity (Keystone)- Authentication and authorization

Image service(Glance) – Stores and retrieves virtual machine disk images



OpenStack™ Nova Architecture

nova-api

- accepts and responds to end user compute API calls
- Supports openstack Compute API and Amazon's EC2 API

•network-compute

- Creates and terminates VM instances via hypervisor's API(libvirt for KVM or QEMU, XenAPI for Xen Server, VMware API for VMware etc.)

•network-volume

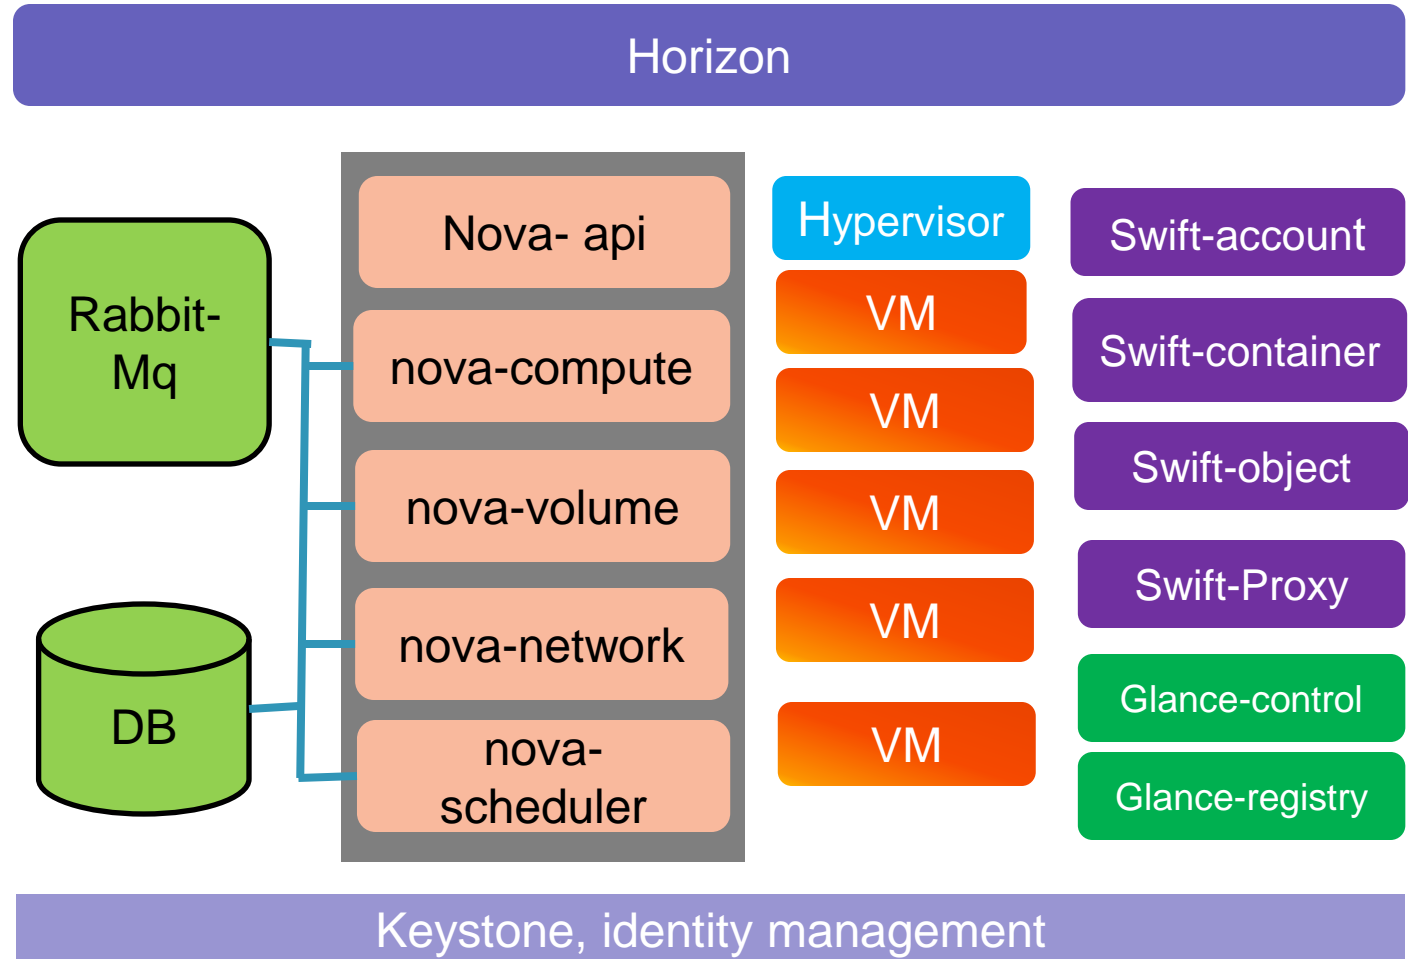
- Creation, attaching, and detaching of volumes to compute instances

•network-network

- Worker daemon accepts networking tasks from queue and performs tasks

•network-schedule

- Accepts VM instance from queue and determines where it should run



Keystone

- Identity Management
 - User management
 - Service catalogue
- Tenant – customer/ organization
- User – person/system uses Service
- Authentication
- Token – Keystone use for authentication
- Service – compute, object store, image service etc
- Endpoint - A network-accessible address, usually described by a URL, from where you access a service

Instances & Volumes

Logged in as: admin [Settings](#) [Sign Out](#)

Success: Instance "test-www.demo.com" launched.

Instances [Launch Instance](#) [Terminate Instances](#)

<input type="checkbox"/>	Instance Name	IP Address	Size	Status	Task	Power State	Actions
<input type="checkbox"/>	test-www.demo.com	10.4.128.20	4GB RAM 2 VCPU 10.0GB Disk	Active	None	Running	Edit Instance
<input type="checkbox"/>	test-www.demo.com	10.4.128.19	4GB RAM 2 VCPU 10.0GB Disk	Build	Spawning	No State	Edit Instance
<input type="checkbox"/>	myserve	10.4.128.18	2GB RAM 1 VCPU 10.0GB Disk	Active	None	Running	Edit Instance
<input type="checkbox"/>	myserver	10.4.128.16	2GB RAM 1 VCPU 10.0GB Disk	Active	None	Running	Edit Instance

Displaying 4 items

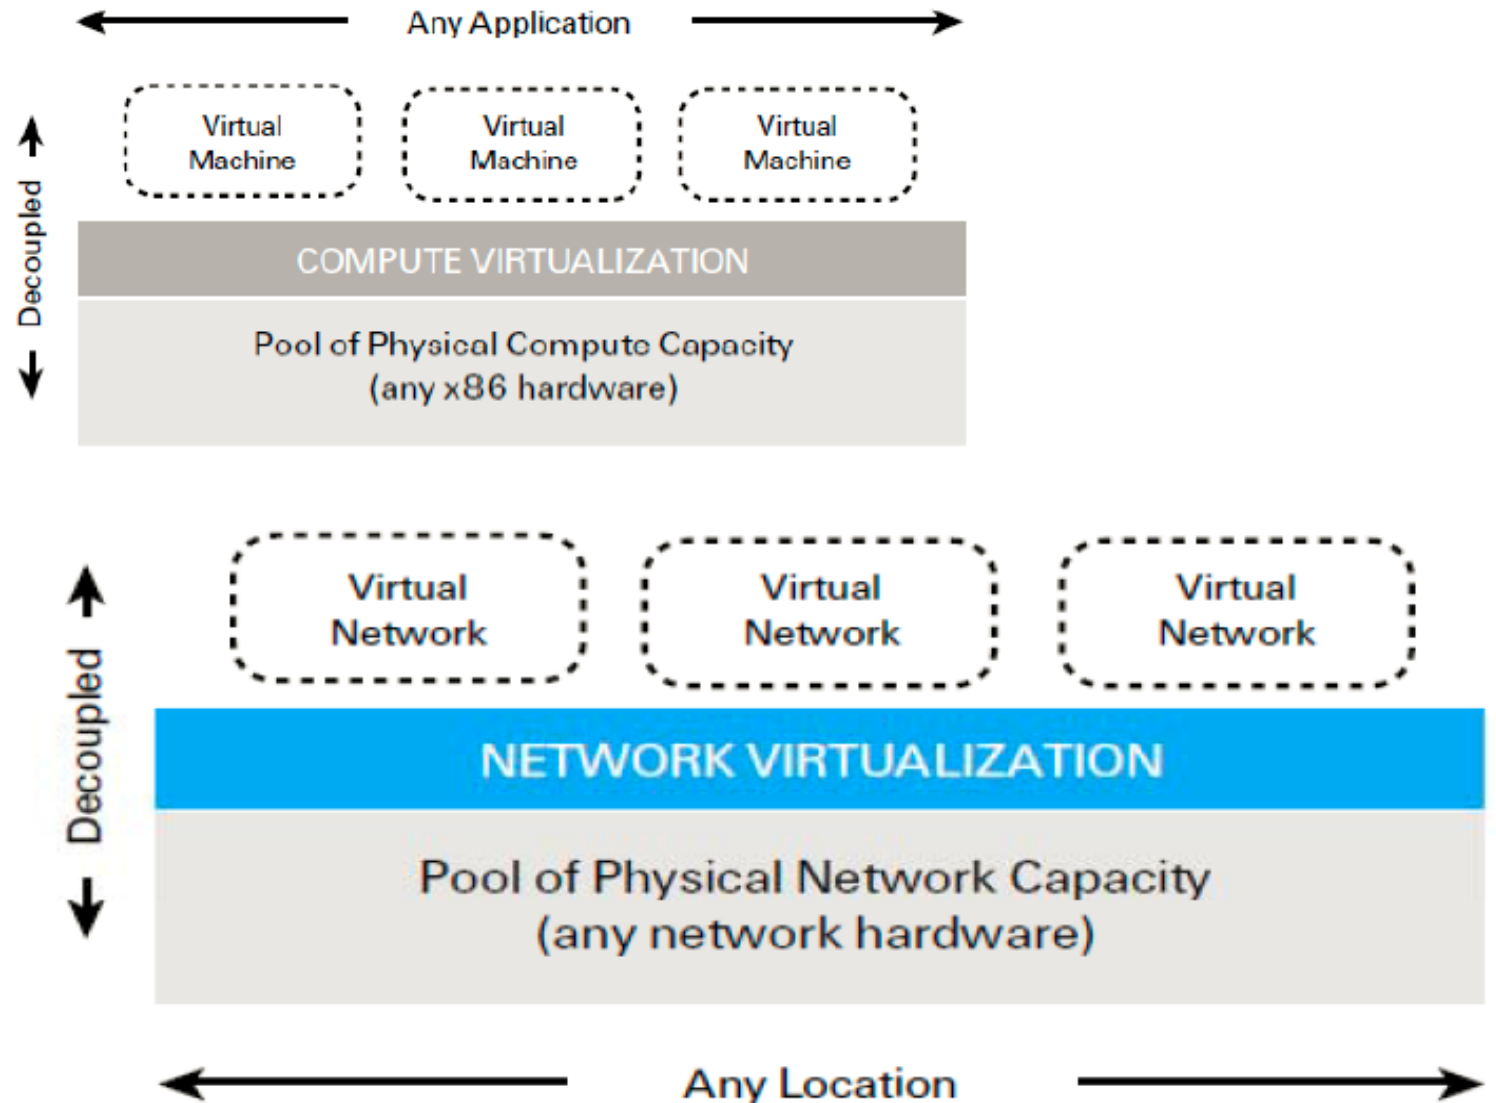
Volumes [Create Volume](#)

<input type="checkbox"/>	Name	Description	Size	Status	Attachments	Actions
--------------------------	------	-------------	------	--------	-------------	---------

Neutron uses network virtualization

- Hypervisor
- nova-scheduler
- nova-compute

- Neutron-server
- Plug-in
- agents



Anatomy of Neutron

- Plug-In Architecture

 - Different network Technologies

- Neutron server

 - Implement neutron API and its extensions
 - Enforce network model
 - Network, subnet, and port
 - IP addressing to each port

- Plug-in agent

 - Run on each compute node on hypervisor performs local vSwitch configuration
 - Connect instances to network port

- DHCP agent

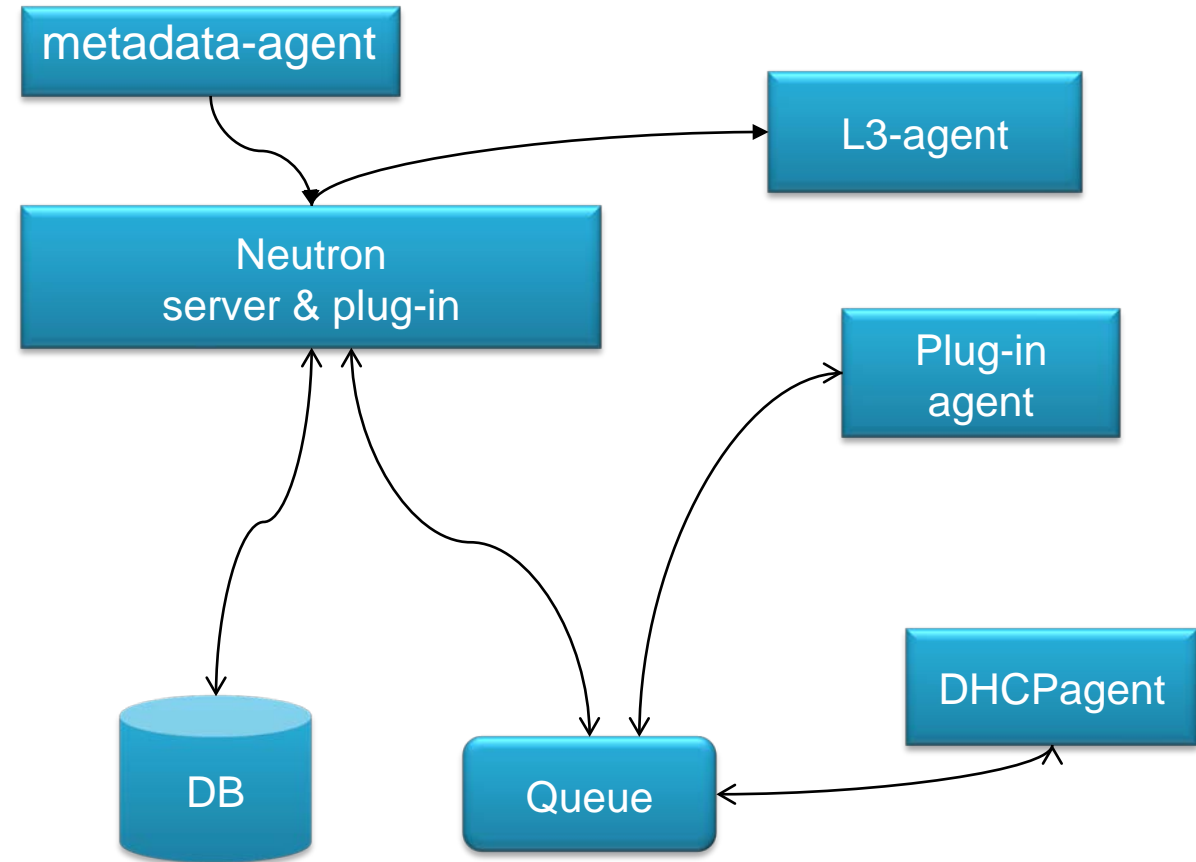
 - Provides DHCP services to tenant networks
 - Start/stop dhcp server maintain dhcp configuration

- L3-agent

 - To implement floating IPs and other L3 features, such as NAT
 - One per network and provides external n/w access to VMs

- Metadata-agent

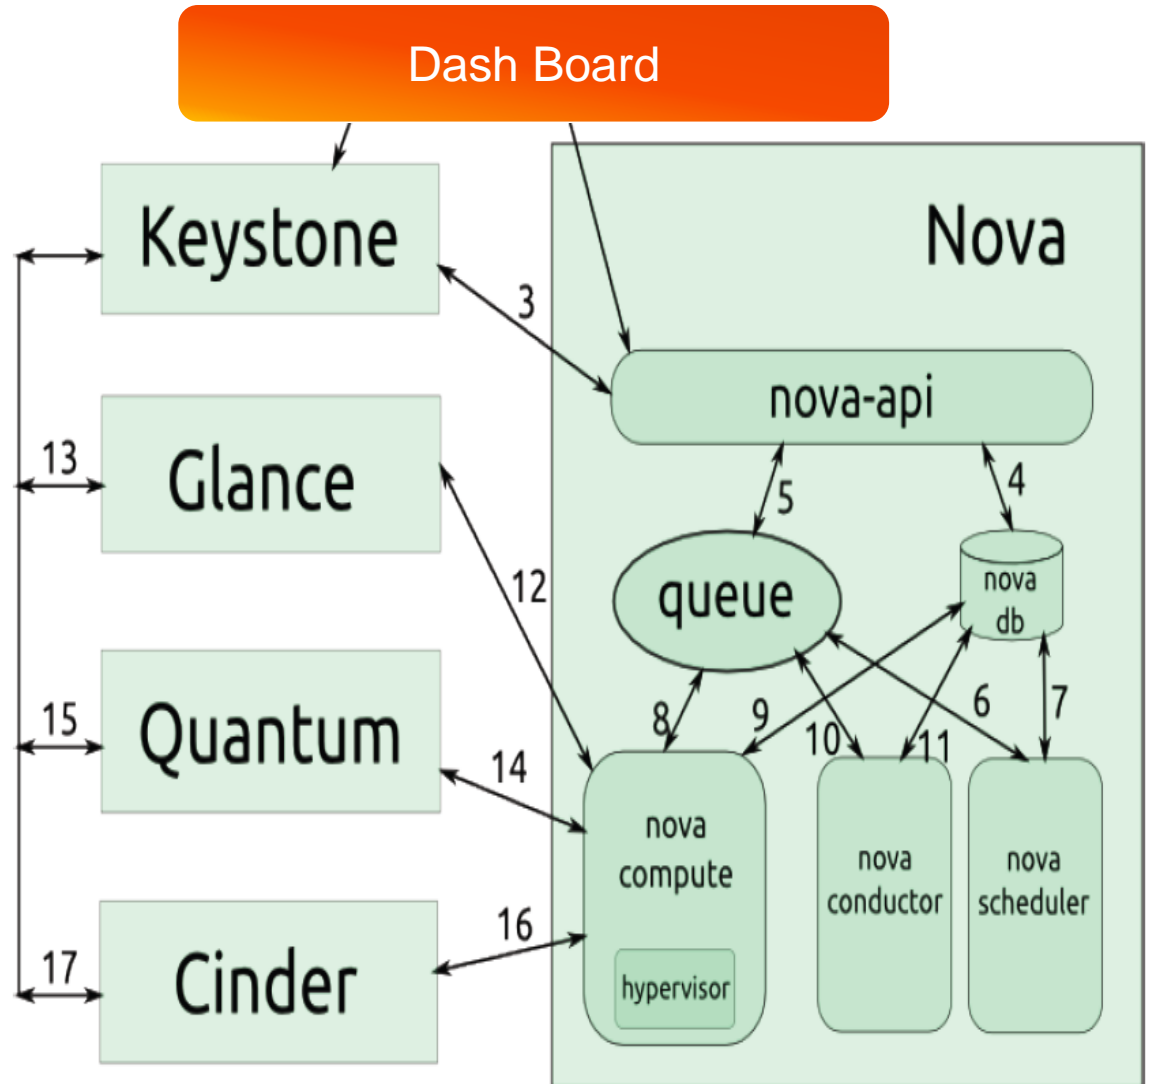
 - Mediate between quantum L3-agent, Dhcp agent with Openstack nova metadata API server



DB service and Queue will be shared with other OpenStack stack services

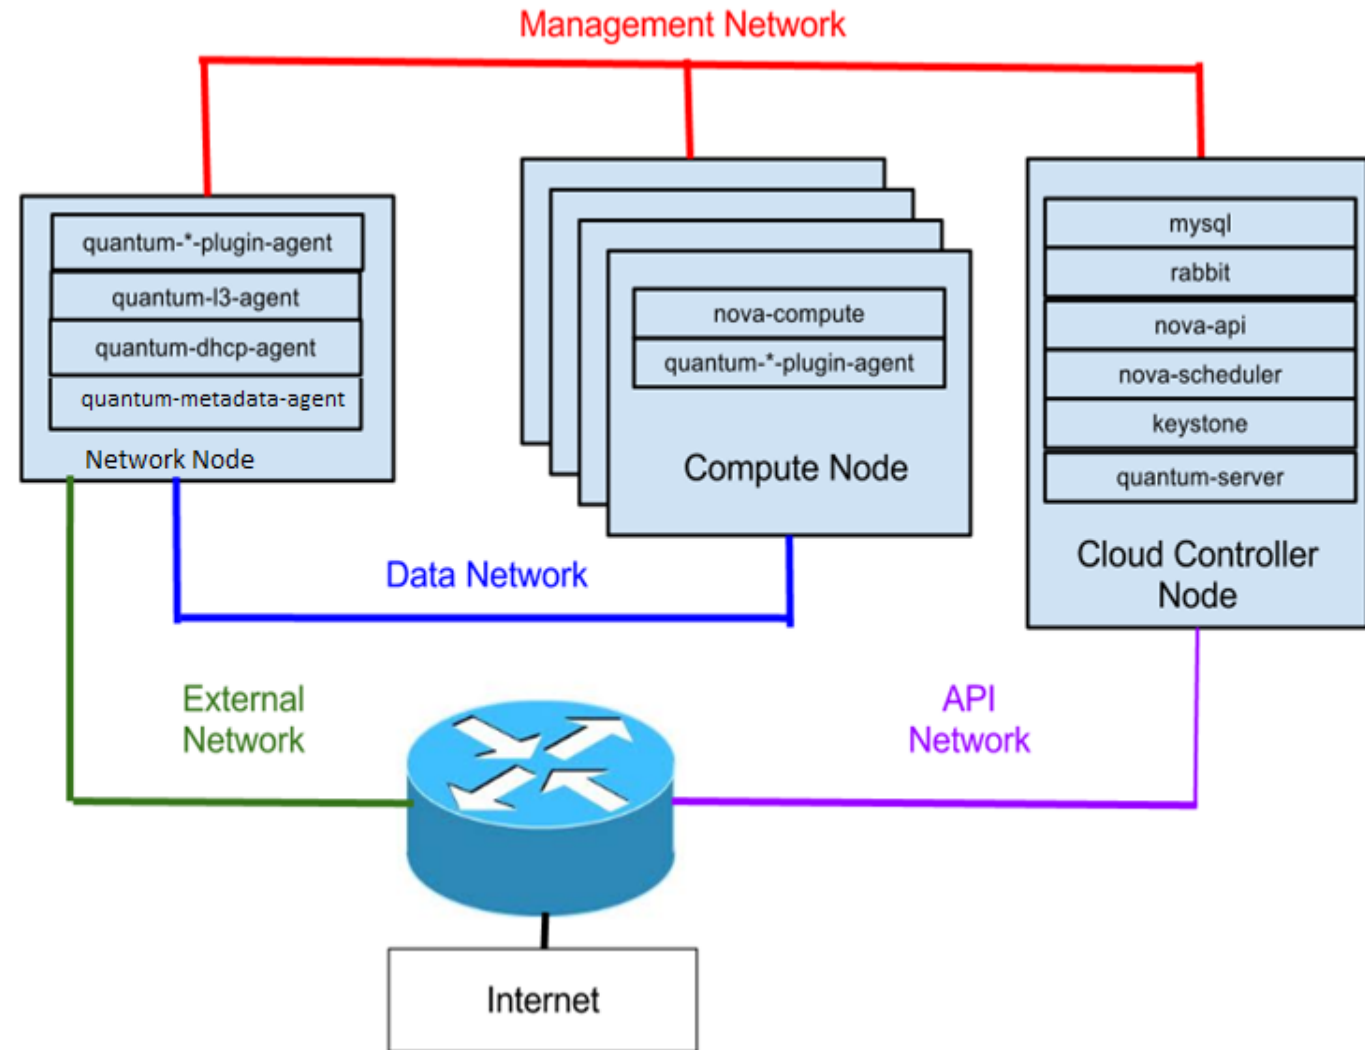
Provisioning of VM

- Dash board gets the credentials authenticates with keystone via REST
- nova-api receives the request and sends the request for validation auth token and access permission to keystone
- nova-api interacts with nova-database
- Creates initial db entry for new instance
- Nova-api sends rpc call to nova scheduler
- Scheduler picks the request from Queue
- Scheduler finds the host ID
- Nova-compute picks the request from the queue
- Nova compute sends rpc call to nova-conductor to fetch the instance information
- Nova-compute picks the request from the queue
- Nova-compute interacts with nova-database and returns instance information

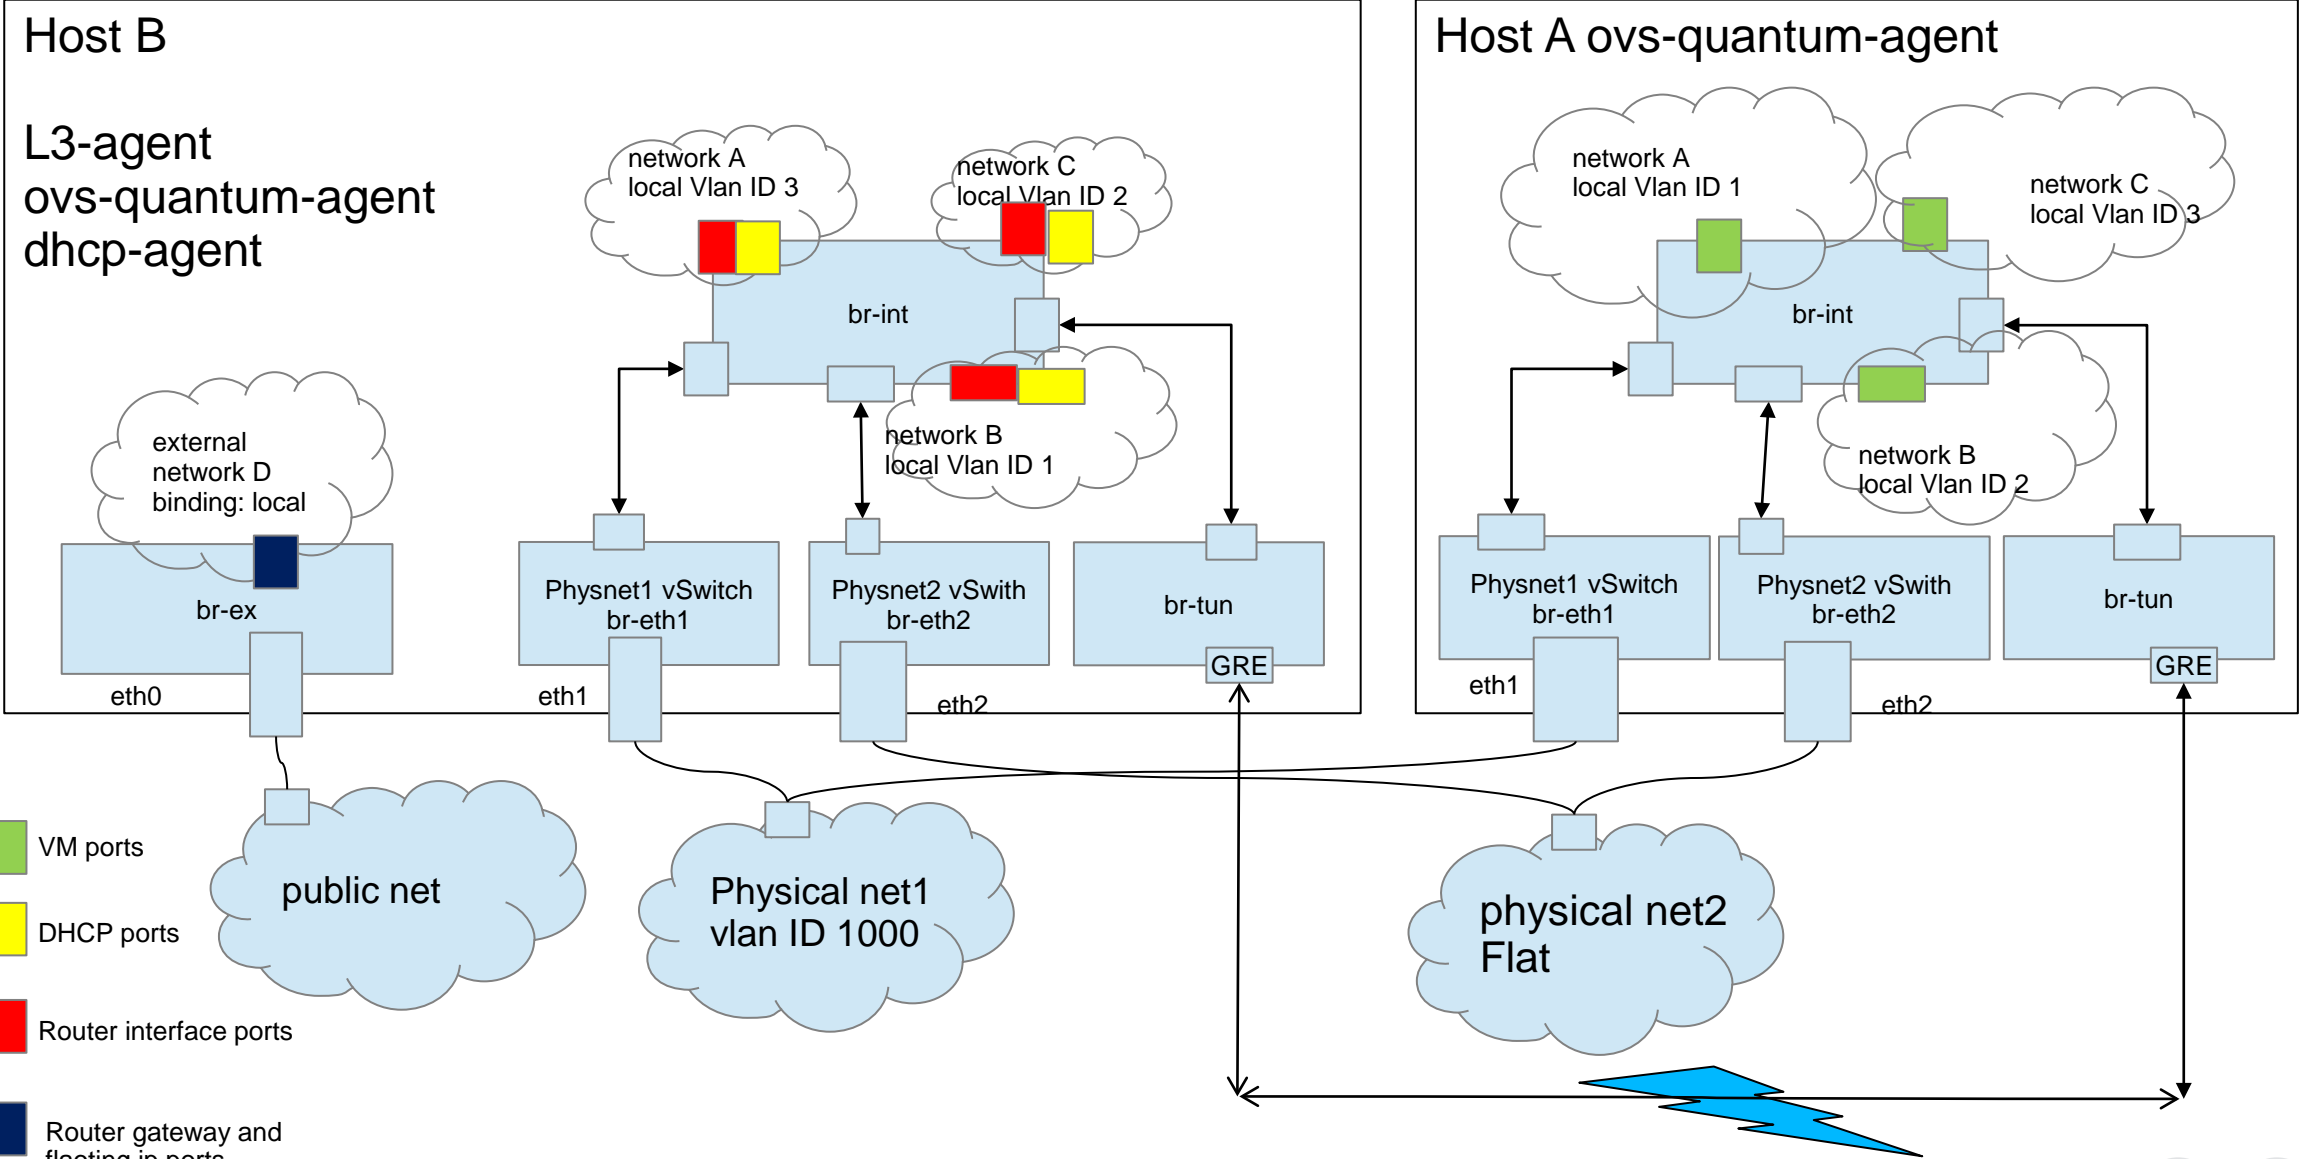


Deployment Model

- **Management network.** Used for internal communication between OpenStack Components. The IP addresses on this network should be reachable only within the data center.
- **Data network.** Used for VM data communication within the cloud deployment. The IP addressing requirements of this network depend on the Quantum plugin in use.
- **External network.** Used to provide VMs with Internet access in some deployment scenarios. The IP addresses on this network should be reachable by anyone on the Internet.
- **API network.** Exposes all OpenStack APIs, including the Quantum API, to tenants. The IP addresses on this network should be reachable by anyone on the Internet. This may be the same network as the external network, as it is possible to create a quantum subnet for the external network that uses IP allocation ranges to use only less than the full range of IP addresses in an IP block.

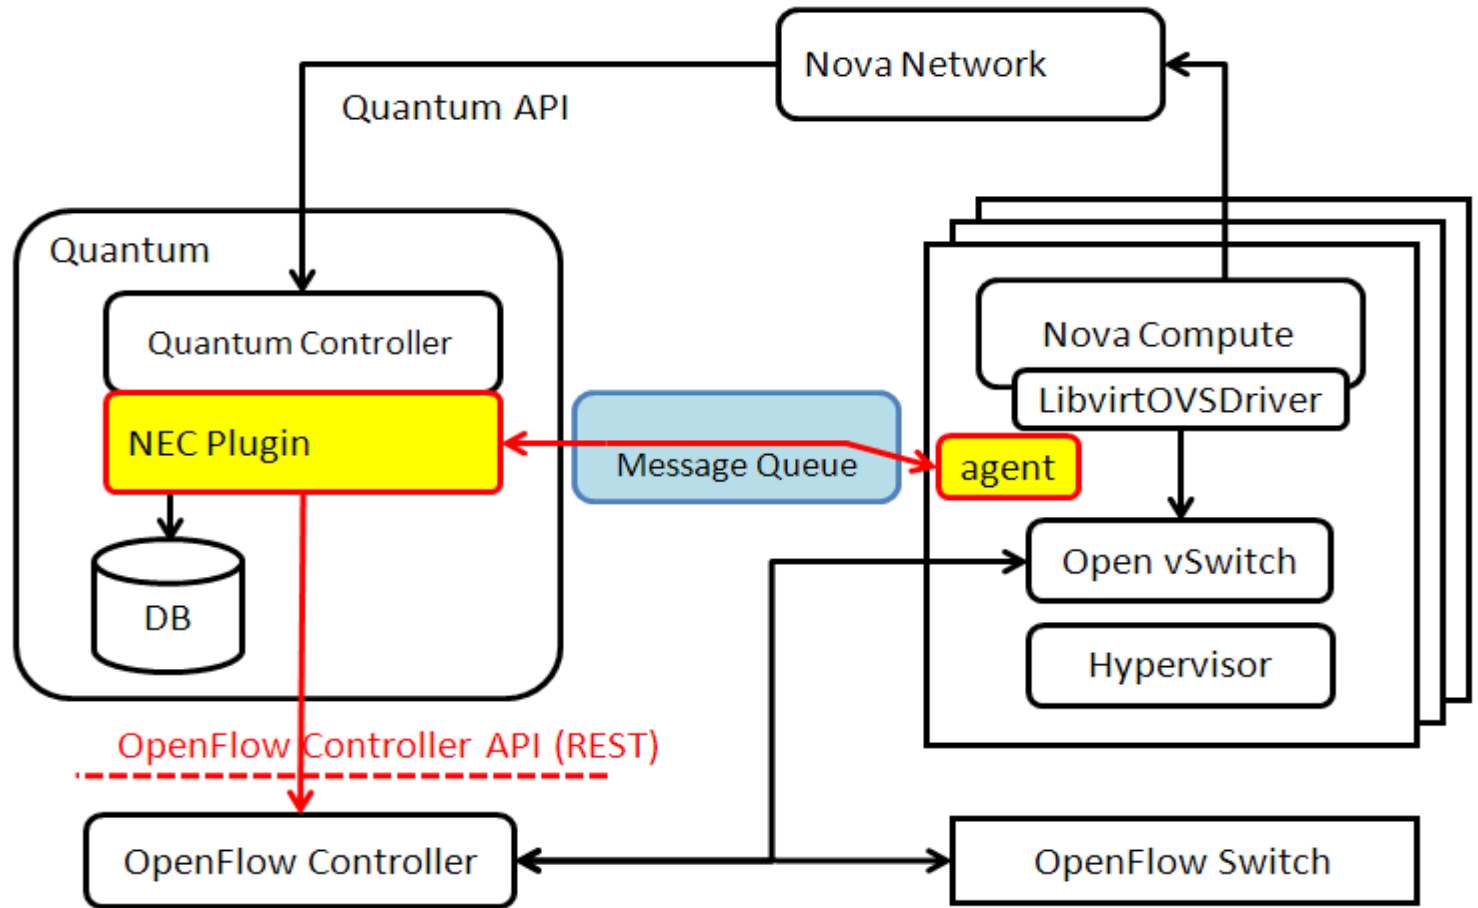


OVS plug-in – Network sample for fixed IPs and routers



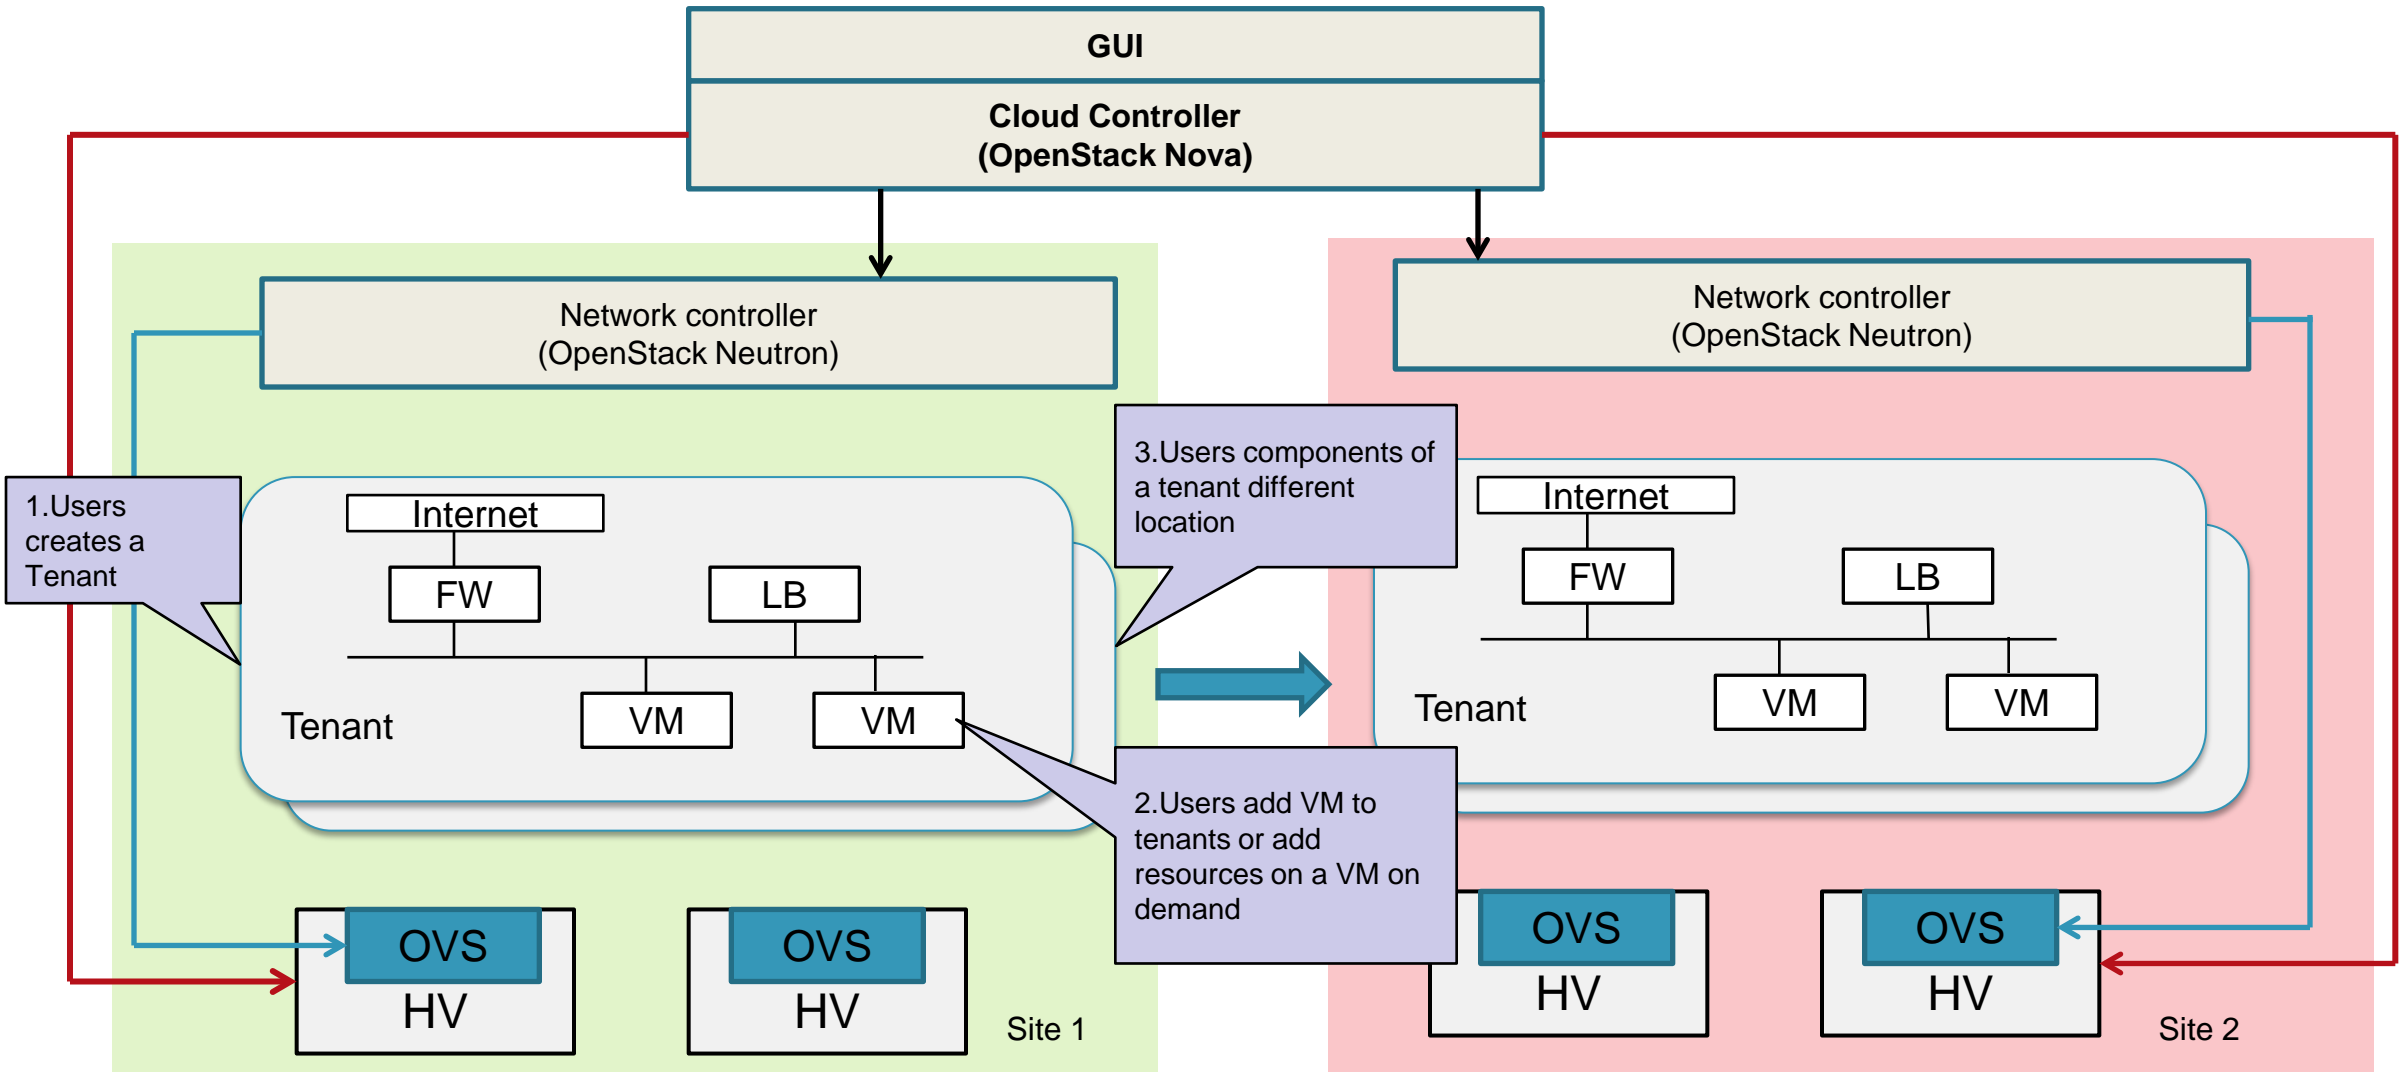
Neutron NEC® Plug-in

- Quantum is earlier version of neutron
- Cloud network dynamics, network resources to OpenFlow controller

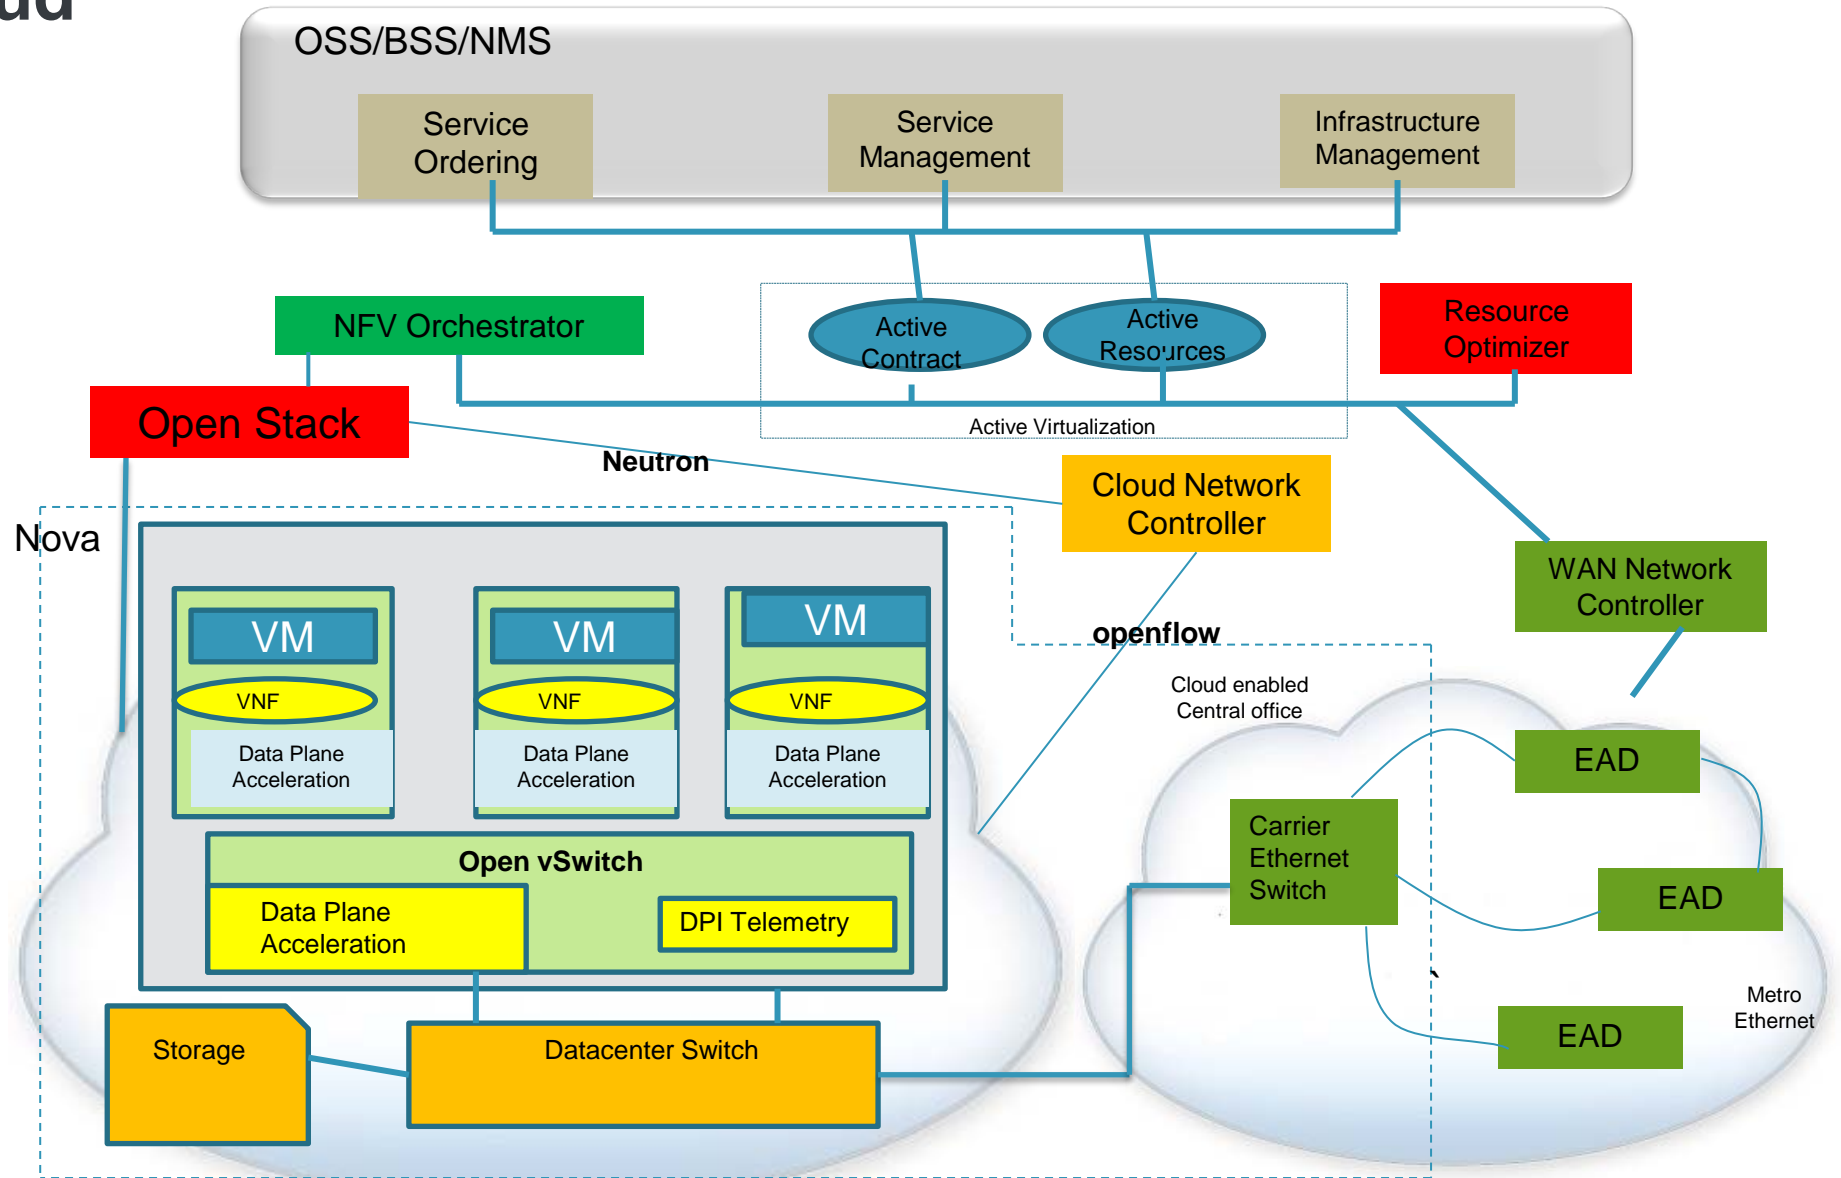


https://wiki.openstack.org/wiki/Neutron/NEC_OpenFlow_Plugin

Openstack and OpenFlow in Virtualized Datacenter



NFV, SDN and Cloud





Future Work

Road map and challenges.....

- Eco-system for Storage Virtualization, Compute virtualisation and Network virtualization
- Performance
- Parallel Innovation
- Connection Rate/setup rate
- Latency
- Fault-tolerance
- High Availability and Security
- Going to cloud is inevitable
- Conceptuality vs. reality vs. hype



References

www.opennetworking.org
<https://wiki.openstack.org/wiki/Neutron>
<http://cleanslate.stanford.edu>
www.openstack.org
www.opennebula.org
www.cloudfv.com
<http://openvswitch.org/>

Q & A



www.Freescale.com