NFV with OpenStack

Vincent Jardin
CTO
6Wind
@VincentJardin

Nick Barcet
VP Products
eNovance
@nijaba

Hello? The technology behind

Hello? The technology behind

Hello? The technology behind

Proprietary hardware

Hello? The technology behind

Full re-deployment for each new generation network

What if this went went virtual?

What if this went went virtual?

What if this went scalable?

What if this went went virtual?

What if this went scalable?

What if this was just like cloud?

Creation of ETSI NFV

- In November 2012 seven of the world's leading telecoms network operators selected ETSI to be the home of the Industry Specification Group for NFV.

- 2 years later: 230 individual companies including 37 of the world's major service providers as well as representatives from both telecoms and IT vendors.

NFV business opportunities

- Latest major disruption in the telco marketplace since IP introduction
- Create opportunities for service providers:
  - To accelerate development of new services
  - To implement Network and IT convergence
- Force NEP
  - To change their business model (transform themselves as Software Providers)
  - To redesign their software

OpenStack NFV subteam

- NFV support for OpenStack aims to provide the best possible infrastructure for such workloads to be deployed in, while respecting the design principles of a IaaS cloud.
- In order for VNF to perform correctly in a cloud world, the underlying infrastructure needs to provide a certain number of functionalities which range from scheduling to networking and from orchestration to monitoring capacities.
- This means that to correctly support NFV use cases in OpenStack, implementations may be required across most, if not all, main OpenStack projects, starting with Neutron and Nova.

https://wiki.openstack.org/wiki/Teams/NFV

3 main specific problems for NFVi

- North/South connectivity requirements

3 main specific problems for NFVi

- North <-> South connectivity requirements

- High Service Level Agreement requirements

3 main specific problems for NFVi

- North <-> South connectivity requirements

- High Service Level Agreement requirements

- East <-> West connectivity requirements

NFV Simplified Architecture

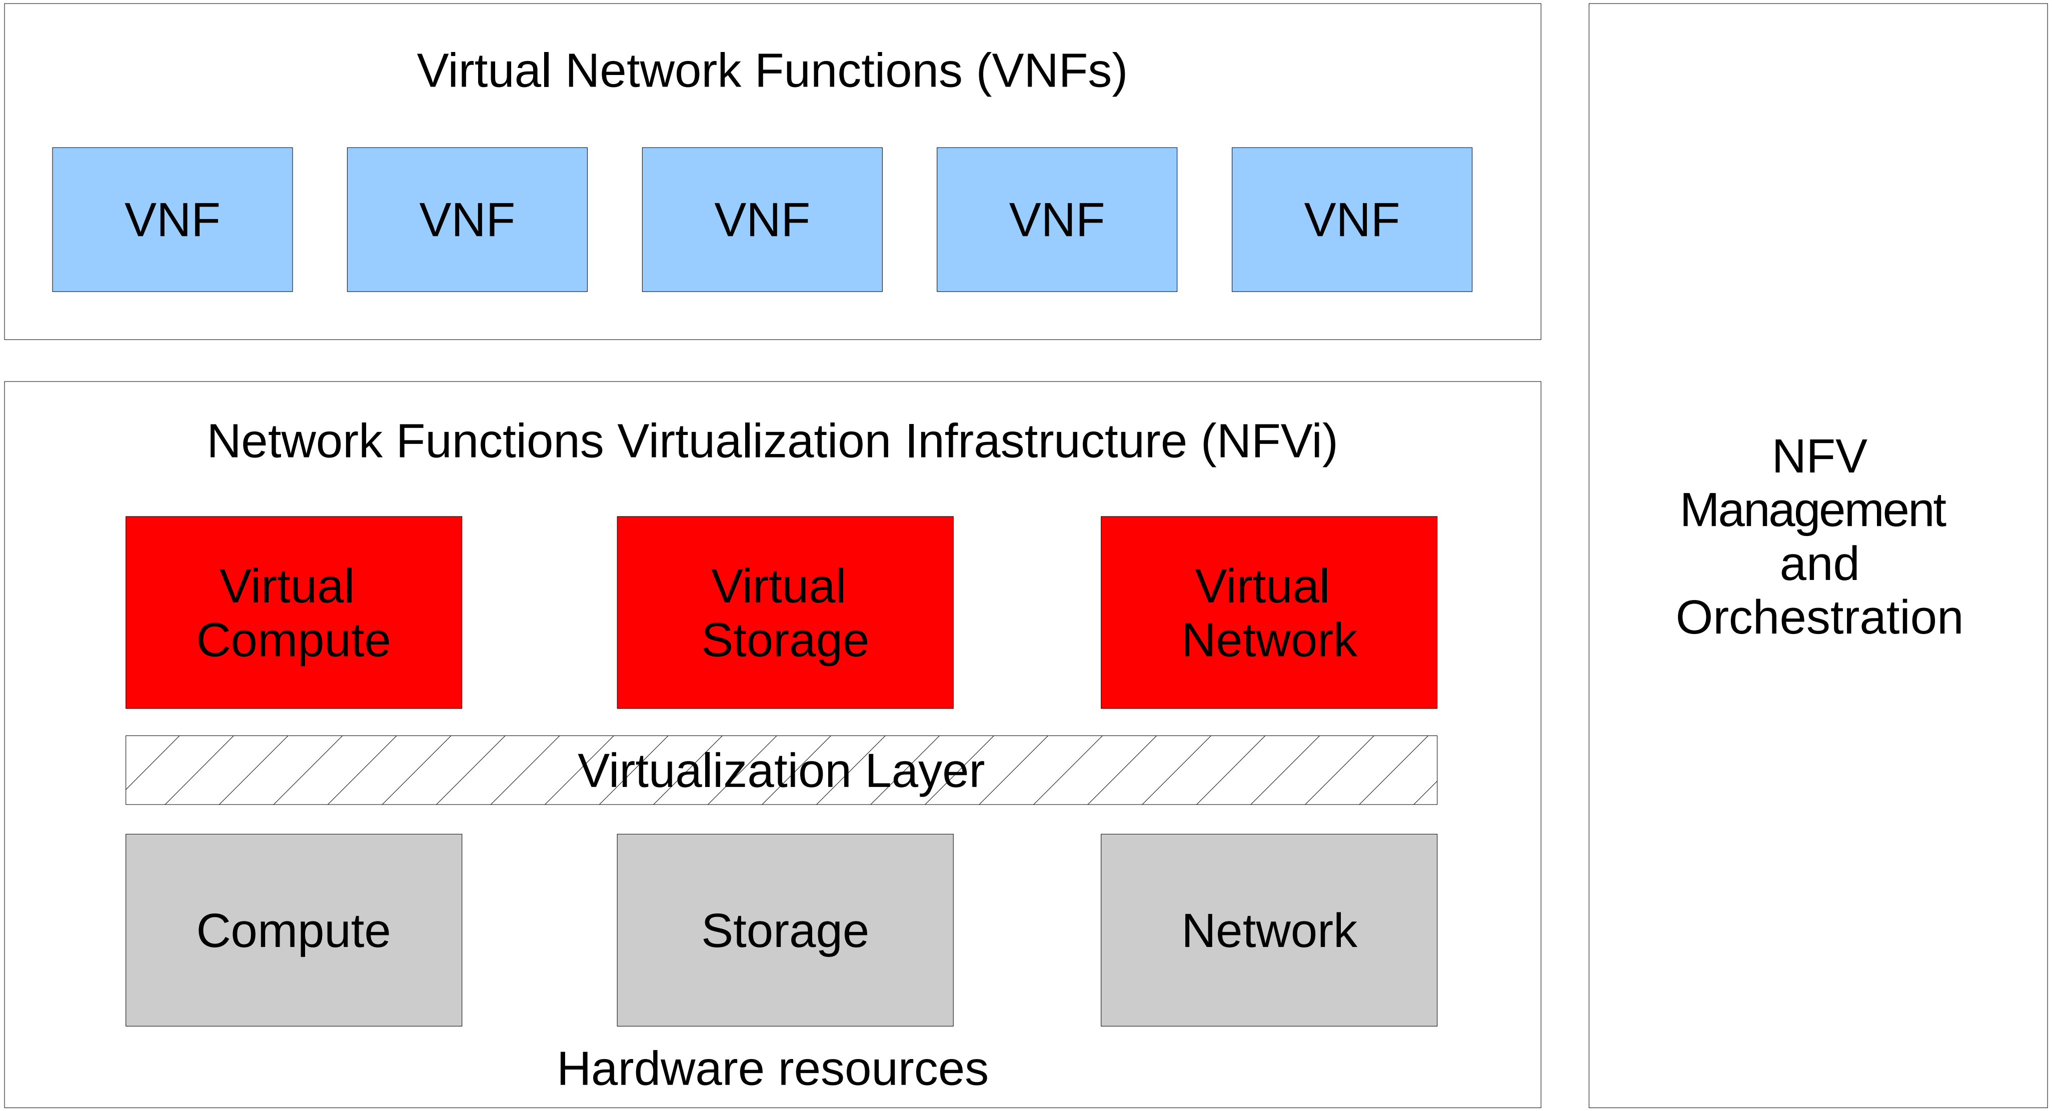

Use Cases

Workload Type	Description	Characteristics	Examples
Data plane	Tasks related to packet handing in an end-to-end communication between edge applications.	Intensive I/O requirements - potentially millions of small VoIP packets per second per core Intensive memory R/W requirements	CDN cache node Router IPSec tunneller Session Border Controller - media relay function
Control plane	Any other communication between network functions that is not directly related to the end-to-end data communication between edge applications.	Less intensive I/O and R/W requirements than data plane, due to lower packets per second More complicated transactions resulting in (potentially) higher CPU load per packet.	PPP session management Border Gateway Protocol (BGP) routing Remote Authentication Dial In User Service (RADIUS) authentication in a Broadband Remote Access Server (BRAS) network function Session Border Controller - SIP signaling function IMS core functions (S-CSCF / I-CSCF / BGCF)
Signal processing	All network function tasks related to digital processing	Very sensitive to CPU processing capacity. Delay sensitive.	Fast Fourier Transform (FFT) decoding Encoding in a Cloud-Radio Access Network (C-RAN) Base Band Unit (BBU) Audio transcoding in a Session Border Controller
Storage	All tasks related to disk storage.	Varying disk, SAN, or NAS, I/O requirements based on applications, ranging from low to extremely high intensity.	Logger Network probe

Blueprints implemented in Juno

Description	Project(s)	Status
Support two interfaces from one VM attached to the same network	Nova	Design Approved / Implemented
SR-IOV Networking Support	Nova	Design Approved / Needs Code Review
Virt driver guest vCPU topology configuration	Nova	Design Approved / Implemented
Evacuate instance to scheduled host	Nova	Approved / Implemented (juno-2)

What's brewing for Kilo? (Highest Priority)

Description	Project(s)	Status
• VLAN trunking networks for NFV • VLAN tagged traffic transmissible over a tenant network • decomposition of VLAN trunks to virtual networks • VLAN tagged traffic to a physical appliance • management of VLANs on ports as sub-ports	Neutron	New
Permit unaddressed interfaces for NFV use cases	Neutron	New

What's brewing for Kilo (the rest)

Description	Project(s)	Status	Description	Project(s)	Status
Discless VM	Nova	Under discussion	Port mirroring	Neutron	Under discussion
Framework for Advanced Services in Virtual Machines	Neutron	Under Discussion	Schedule vms per flavour cpu overcommit	Nova	New
I/O (PCIe) Based NUMA Scheduling	Nova	Design Approved / Needs Code Review	Snabb NFV mechanism driver	Neutron	Approved
Network QoS API	Neutron	Under discussion	Solver Scheduler - complex constraints scheduler with NFV use cases	Nova	Design review in progress
Neutron Services Insertion, Chaining, and Steering	Neutron	Design Approved / Needs Code Review	Support userspace vhost in ovs vif bindings	Nova	Design review in progress
NIC state aware scheduling	Nova	Rejected	Traffic Steering Abstraction	Neutron	Design review in progress
Open vSwitch to use patch ports in place of veth pairs for vlan n/w	Neutron	Superseded / Unknown	VIF_VHOSTUSER (qemu vhost-user) support	Nova	Approved
Open vSwitch-based Security Groups: Open vSwitch Implementation of FirewallDriver	Neutron	Design review in progress	Virt driver guest NUMA node placement & topology	Nova	Design Approved / Needs Code Review
OVF Meta-Data Import via Glance	Glance	New	Virt driver large page allocation for guest RAM *	Nova	Design Approved / Needs Code Review
Persist scheduler hints	Nova	Design review in progress	Virt driver pinning guest vCPUs to host pCPUs	Nova	Design Approved / Needs Code Review

Openstack related challenges

- Cultural Changes:
  - Introduction of devops, Continuous Integration, …
  - Work with opensource communities
  - No more standards...
  - 80% is good enough...
- Technical, R&D:
  - Openstack is evolving very fast (major release every 6 months)
  - Limited availability of skills
  - Redesign of application to become cloud application
- Operation:
  - Need to educate/train people

North ↔ South
East ↔ West

Virtual Network Performances
With 6WIND

Performance first

High Performance East-West Communications

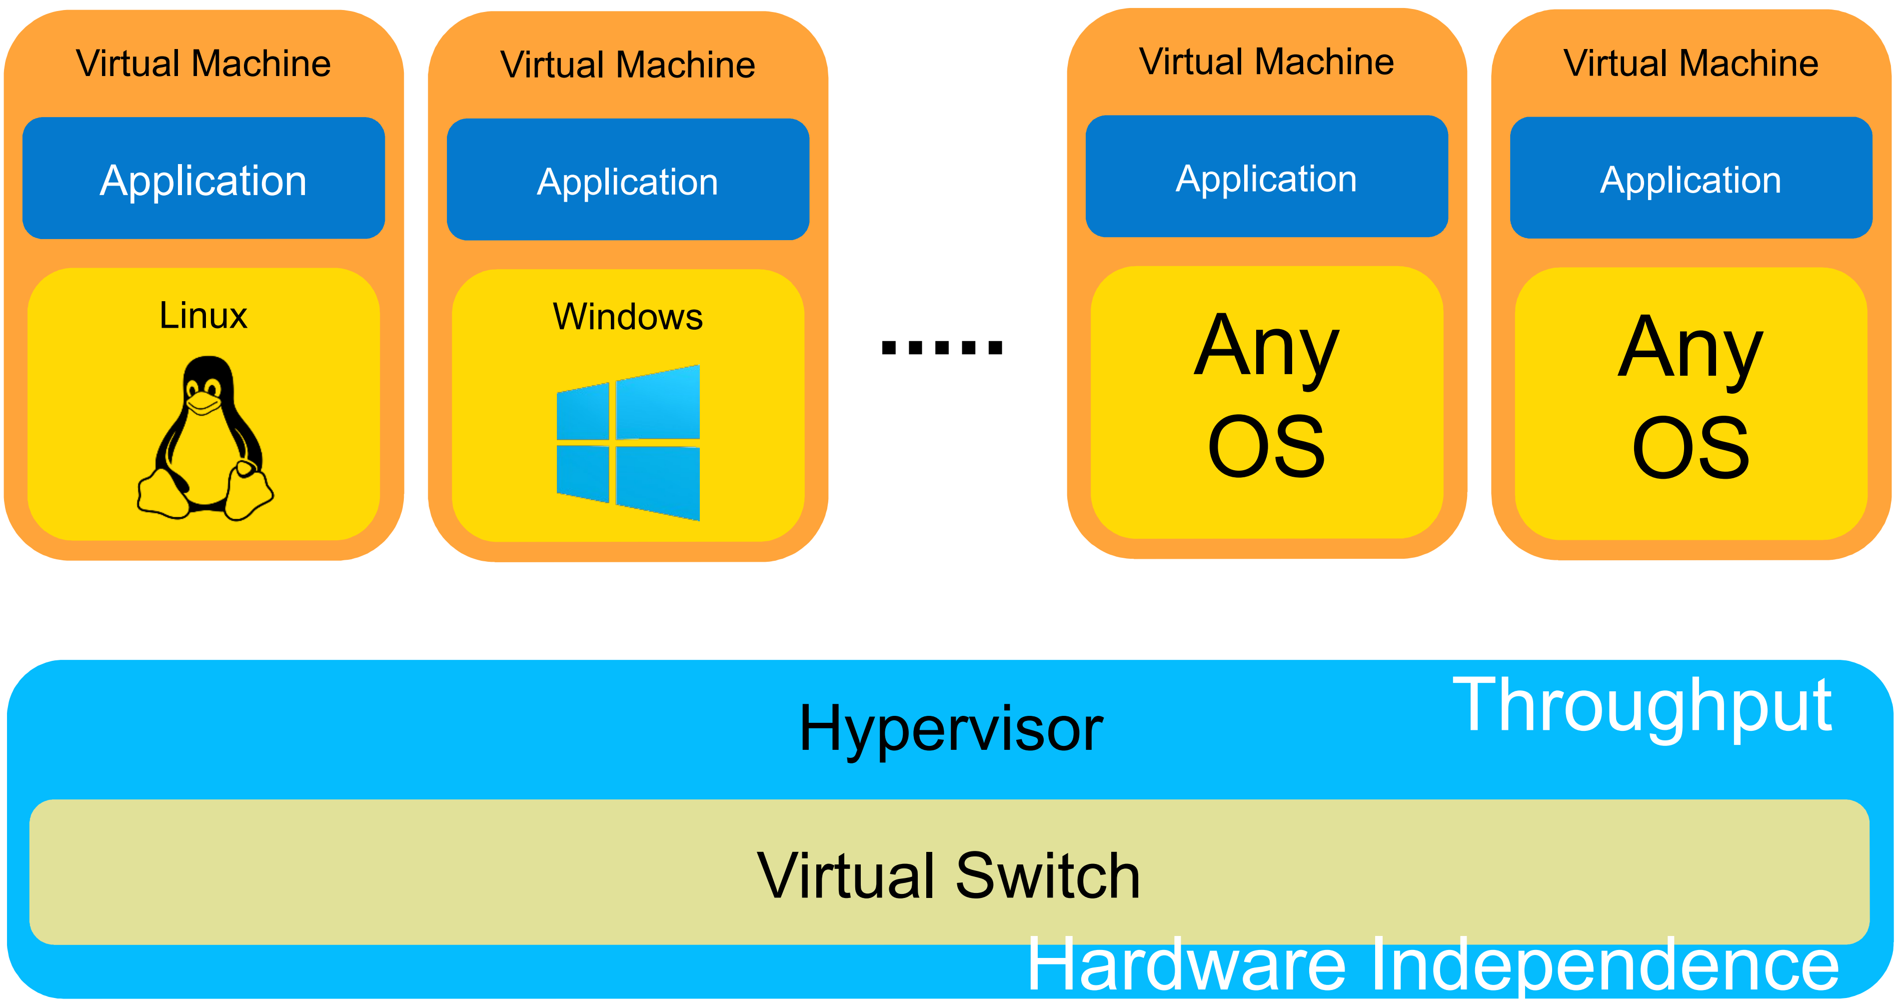

What if SRIOV?
or what if XYZ PCI passthru technologies?

High Performance East-West Communications

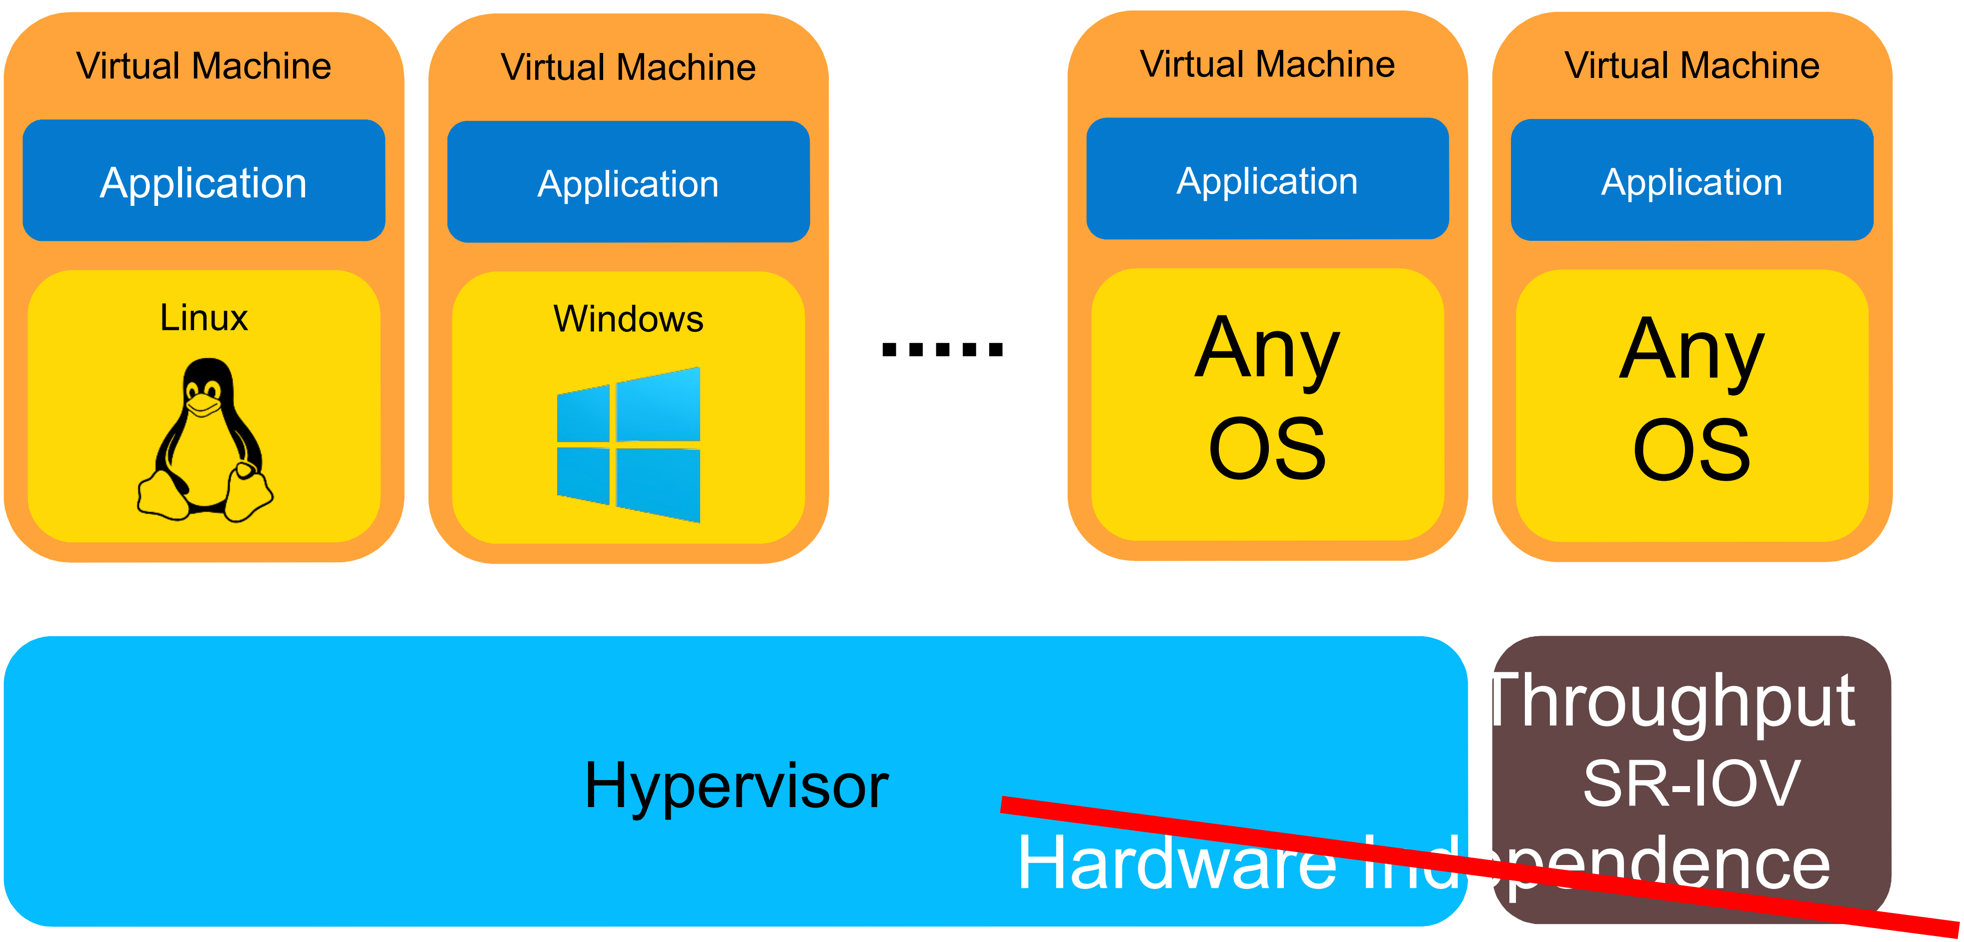

Typical NFV Performance Bottlenecks

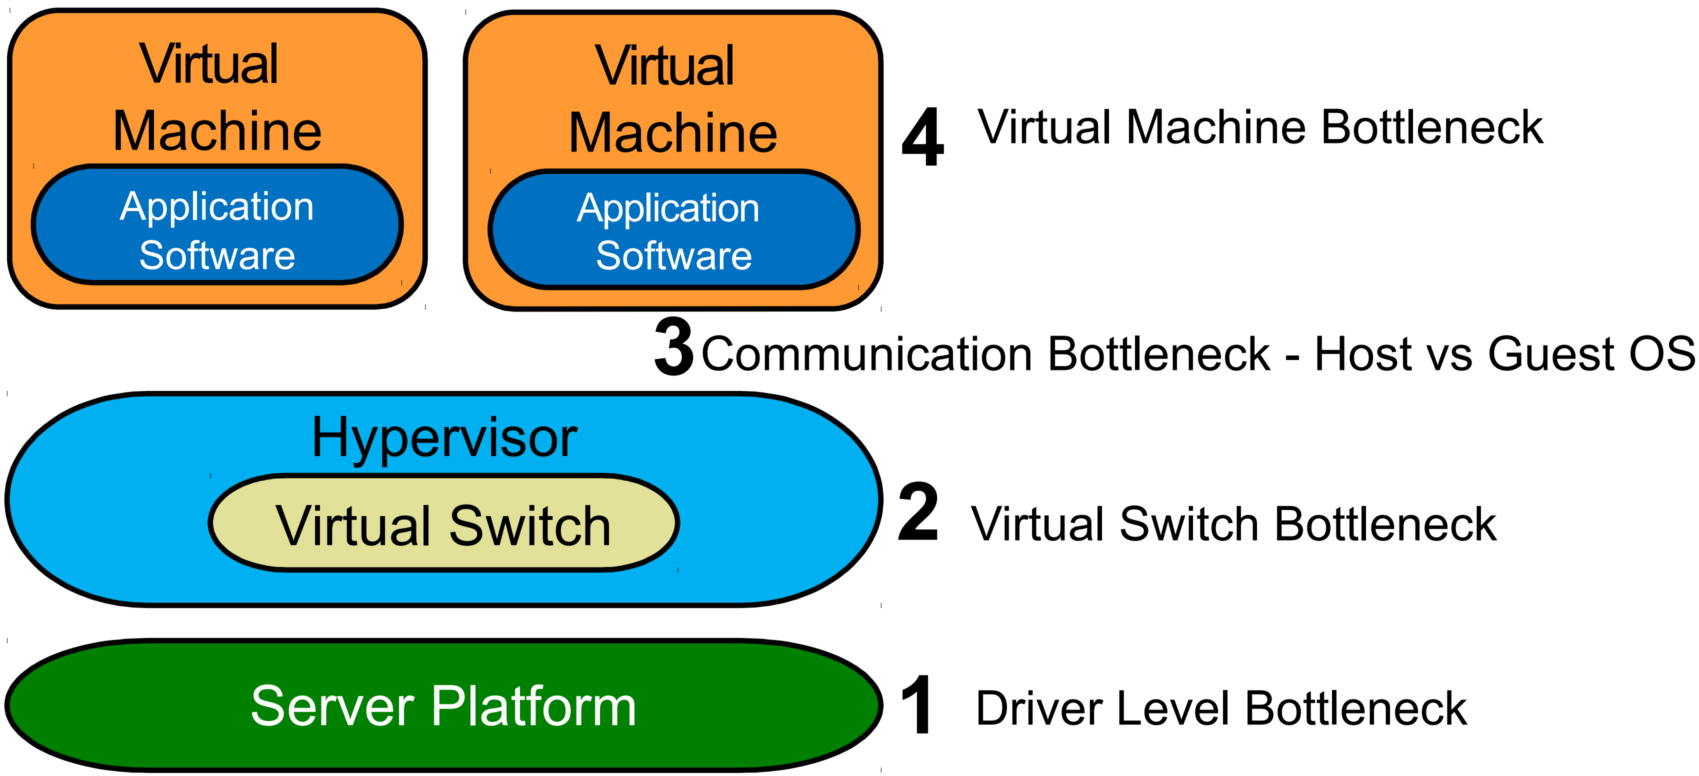

6WINDGate for Industry-Leading Processor Platforms

Architecture-independent "Fast Path Modules"

- **Generic, processor-independent source code**
- **Cycle-level and pipeline-level optimizations**

Architecture-specific "Fast Path Networking SDK"

- **Zero-overhead API for fast path modules**
- **Support for processor-specific features and resources**
- **Leverages processor suppliers' SDKs**

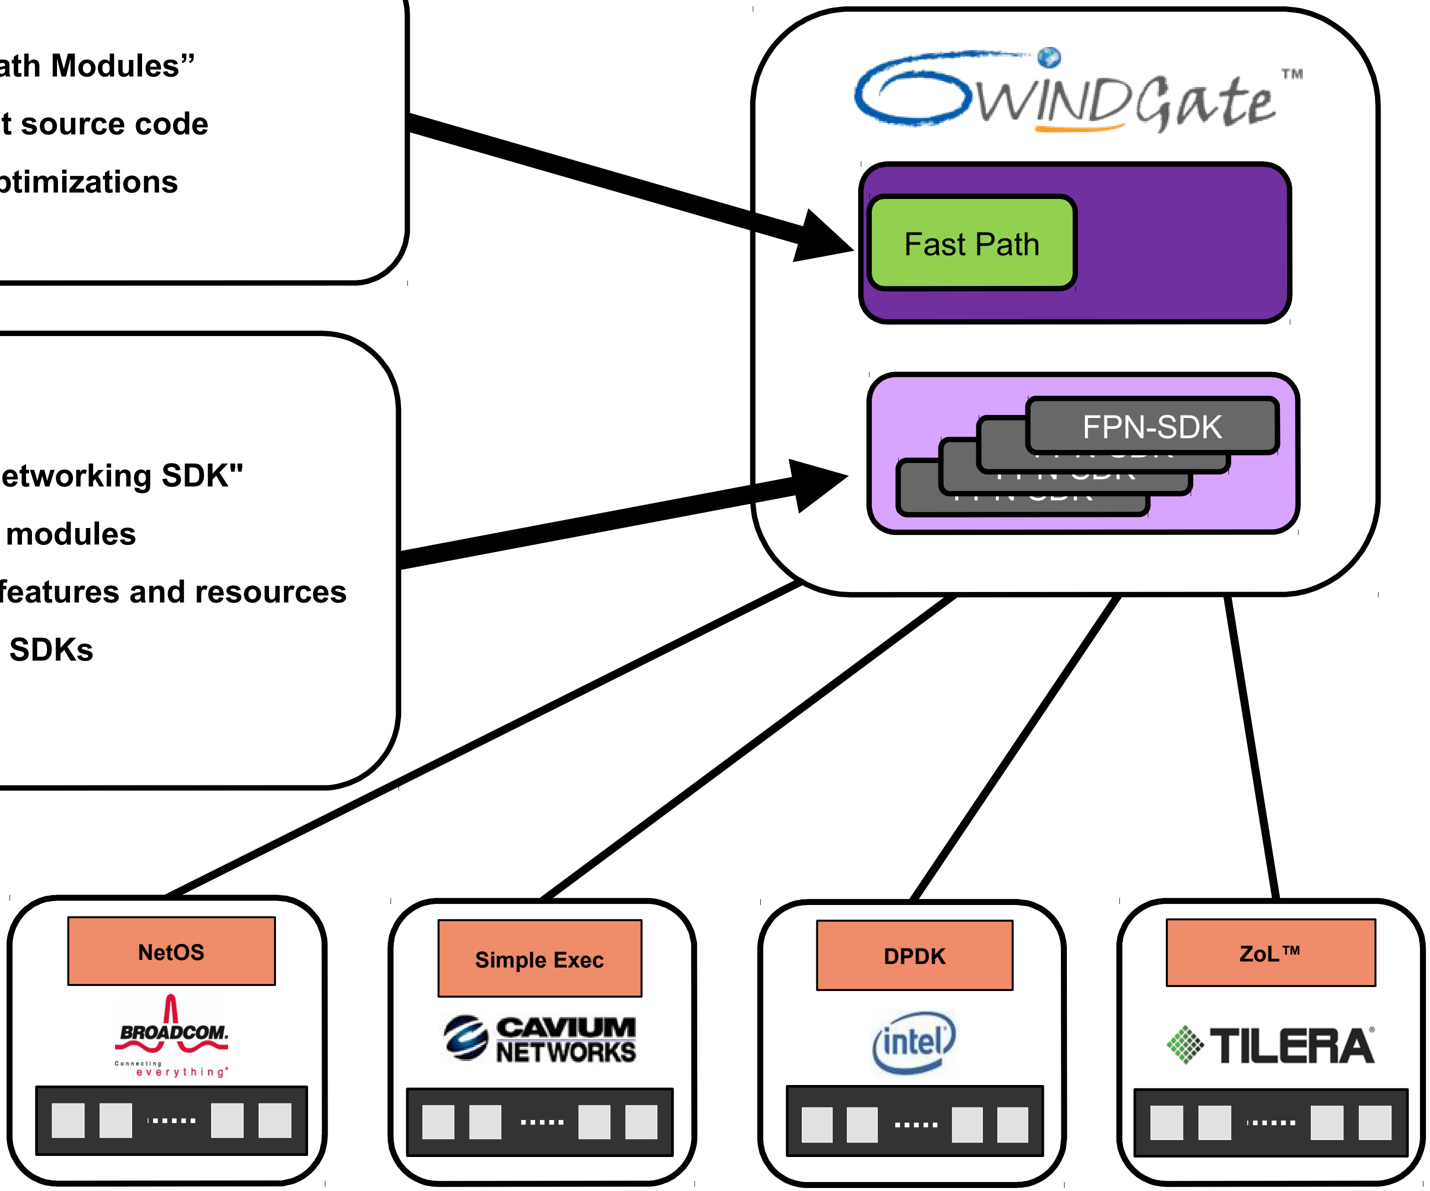

Linux Compatibility is Critical
Linux Acceleration via 6WINDGate

- **Standard Linux functions are accelerated by 6WINDGate**

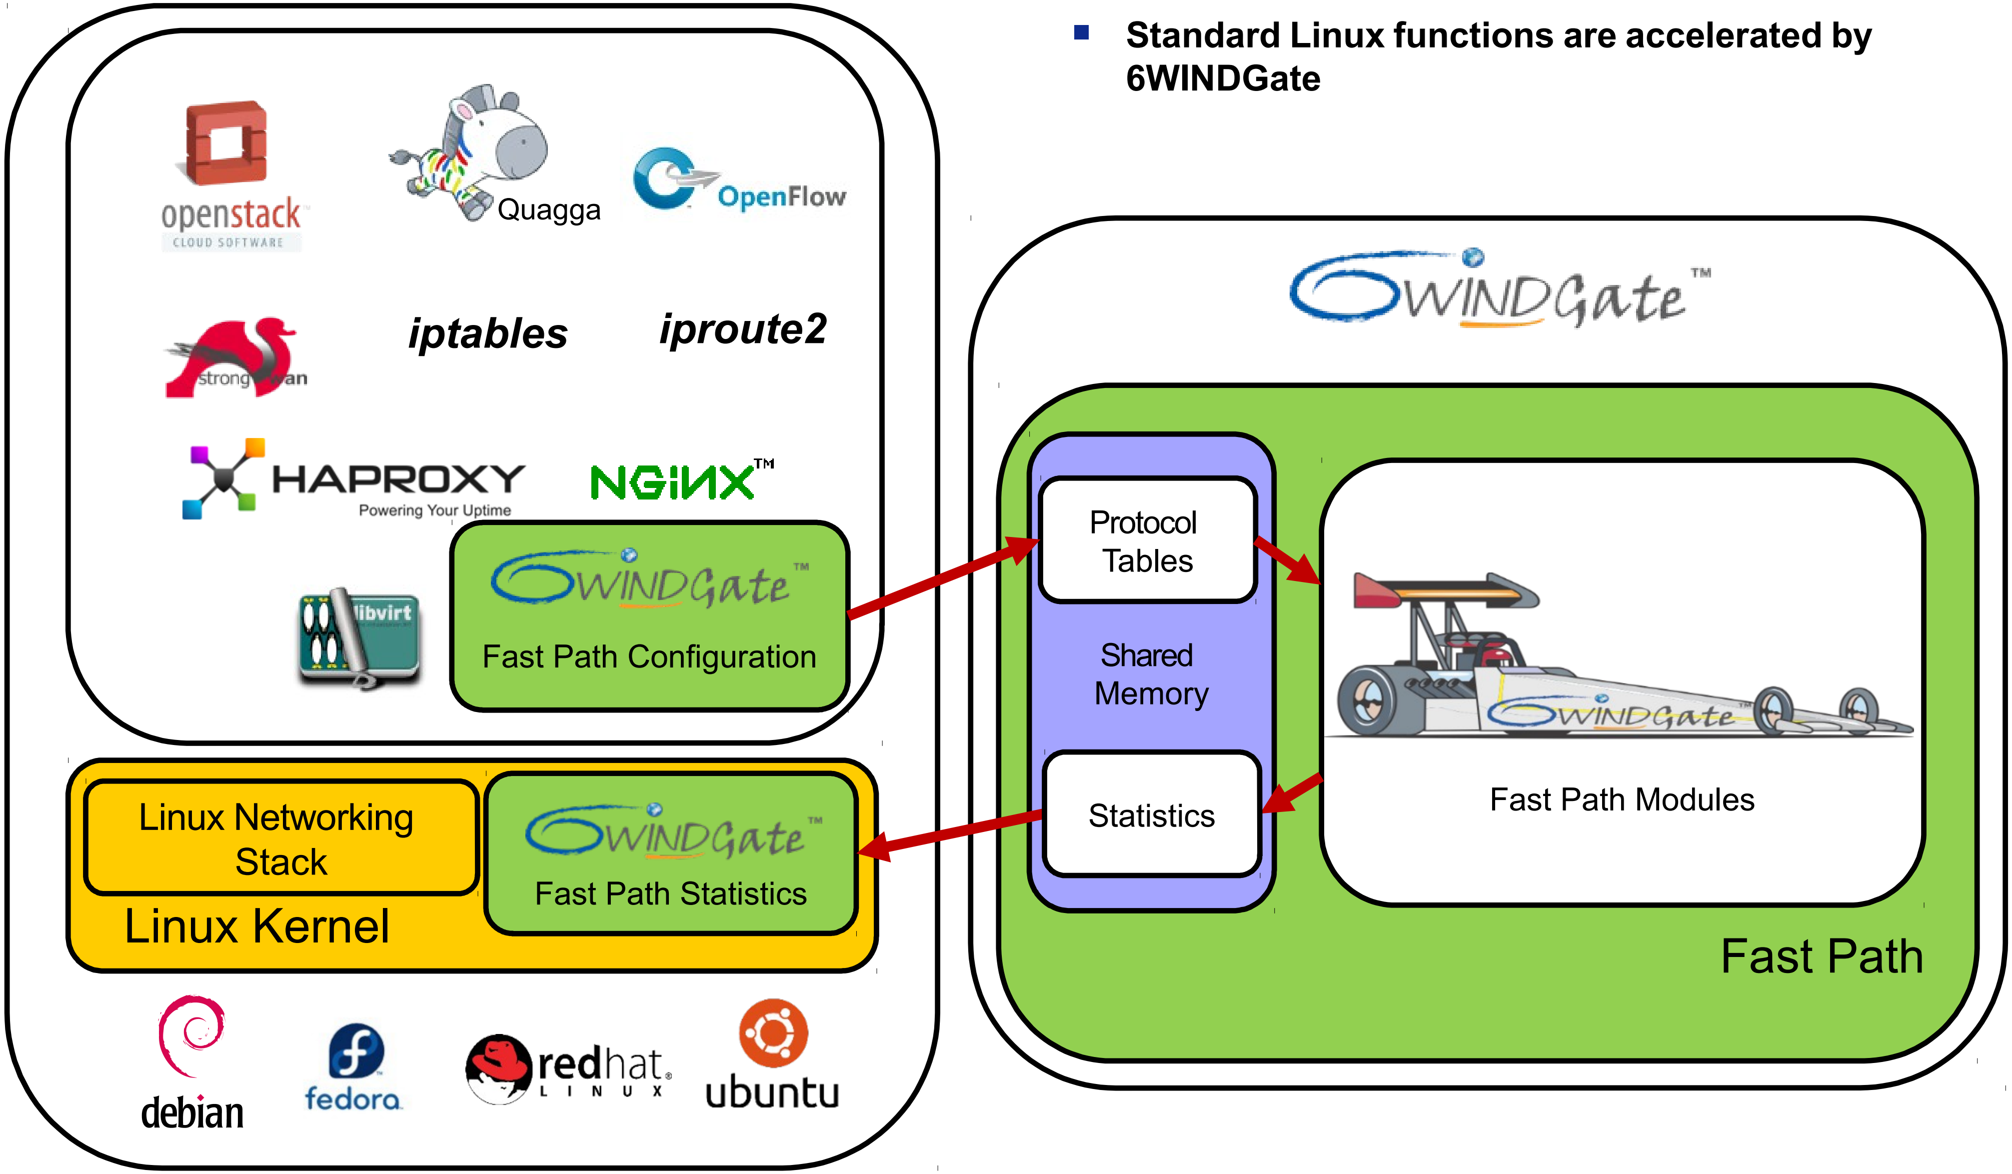

Neutron’s protocols – strong needs from a fast path

Openstack

openstack™
CLOUD SOFTWARE

Nova

Neutron

Say no to proprietary plugins

Say no to SRIOV to be SDN ready.

Accelerate Neutron

6WINDGate FPN-SDK

intel

BROADCOM.
Connecting everything®

CAVIUM NETWORKS

TILERA®

6WINDGate fast path

VLAN

Link Aggregation

GRE

MPLS/VPLS Encapsulation

Ethernet Bridging

IPv4/IPv6 Forwarding

IPv4/IPv6 Multi-cast

Tunneling (IPinIP)

NAT

Filtering IPv4/IPv6

IPsec IPv4/IPv6

QoS

Flow Inspection

OVS Acceleration

VXLAN

6WINDGate DPDK

Mellanox ConnectX®-3 EN Series PMD

Intel® QuickAssist Crypto

Cavium NITROX SDK 5.X Crypto

VIRTIO Guest XEN-KVM PMD

Emulex OCE14000 Series PMD

Intel® Multi-Buffer Crypto

VMXNET3 Guest VMware PMD

Fast vNIC PMD

6WINDGate Extensions to Virtualization

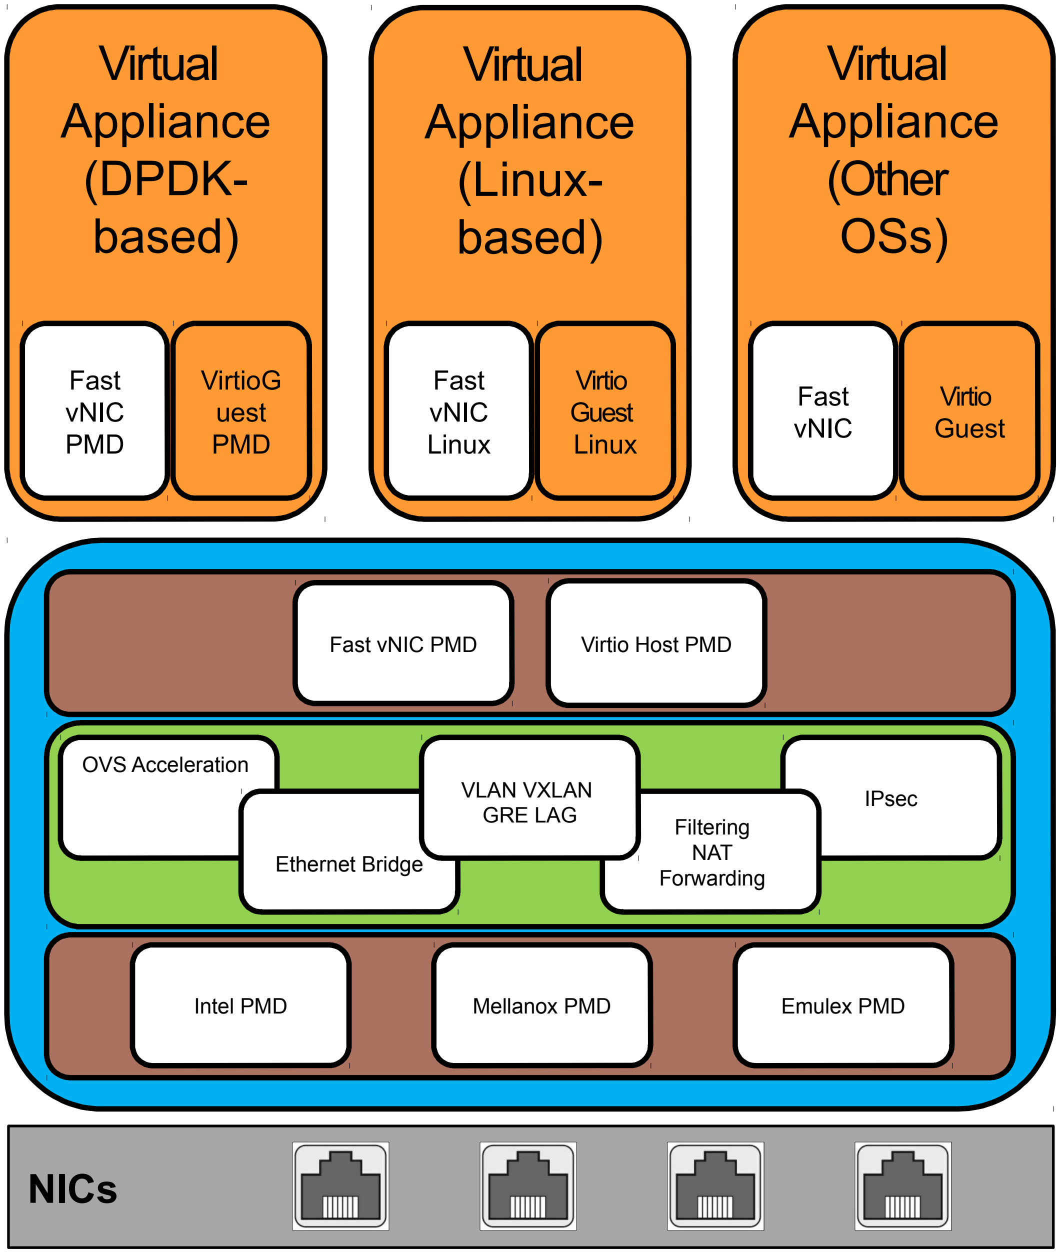

Drivers for Virtual Appliance

- **6WIND drivers for high performance communications**
- **Standard drivers for existing Virtual Appliances**
- **Extensible for all OSs**

Virtual Acceleration

- **6WIND drivers for high performance communications**
- **Accelerated virtual switch and bridging**
- **Extended network services**
- **Dpdk.org with multi-vendor NIC support**

6WINDGate NFVI + VM Performance Comparison

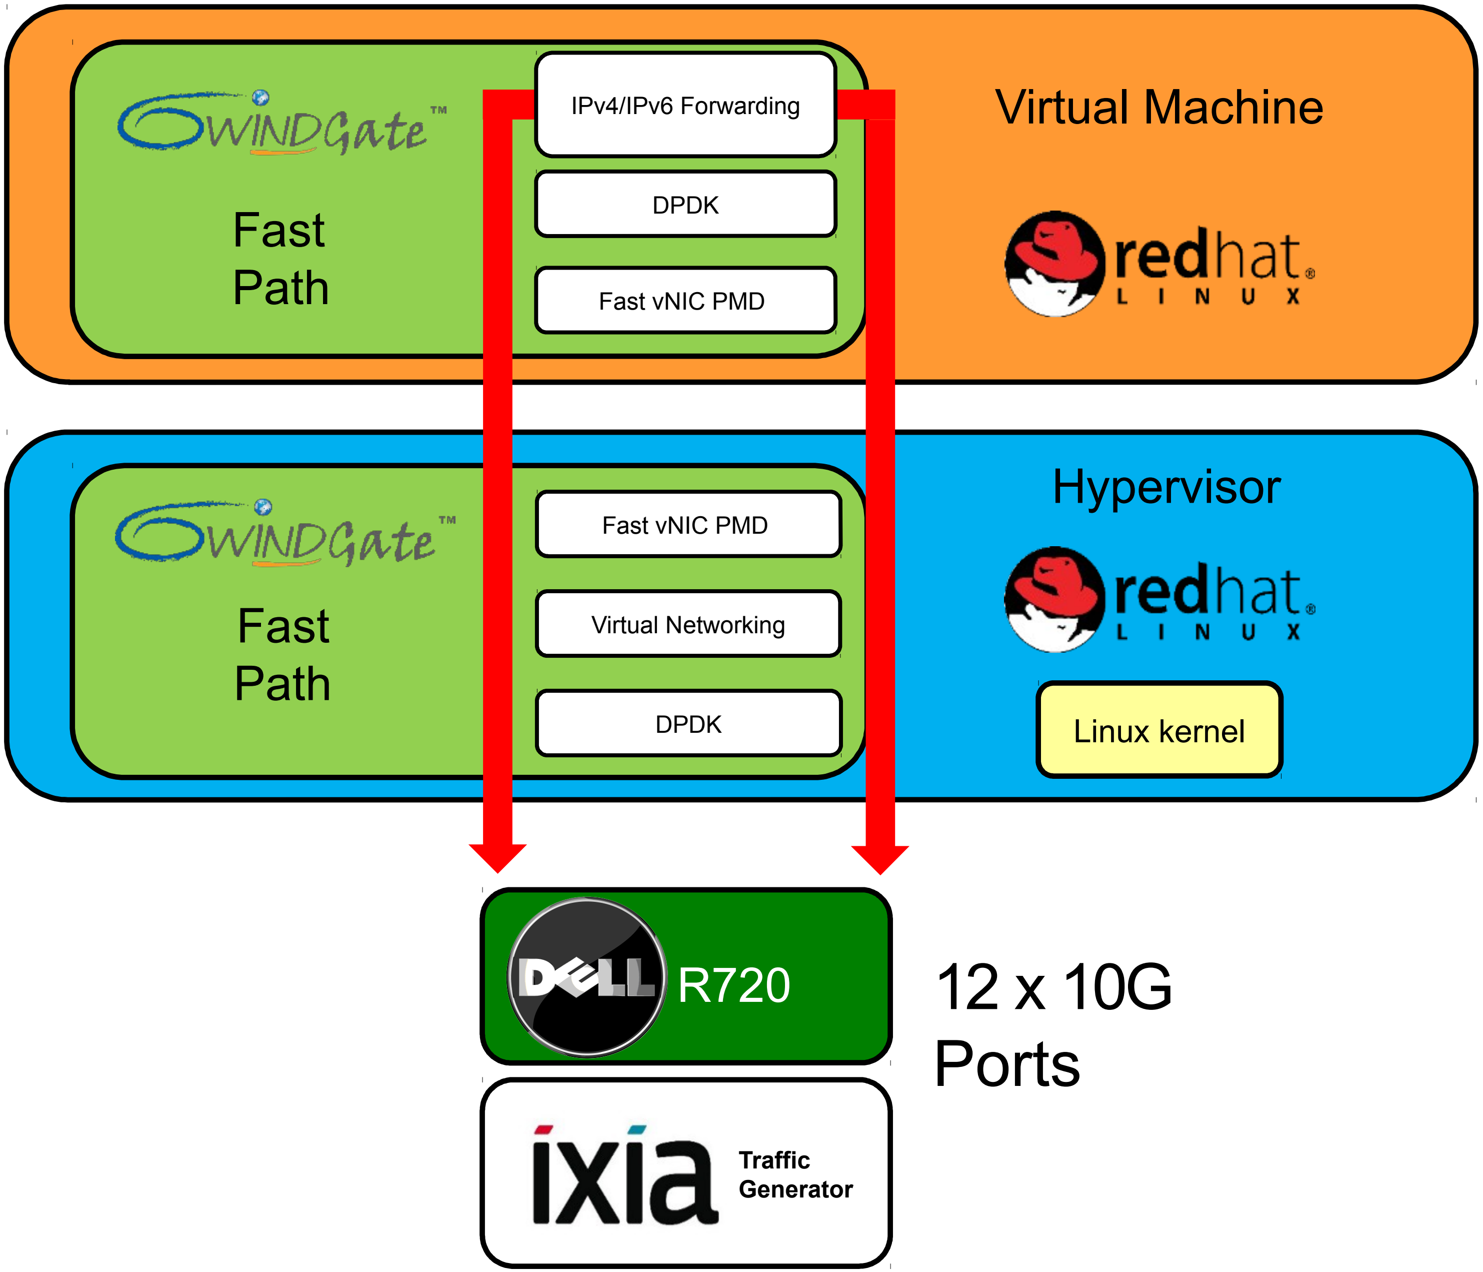

Test 1:
Standard Open vSwitch + Virtio

L2 Throughput

7,2 Gbps

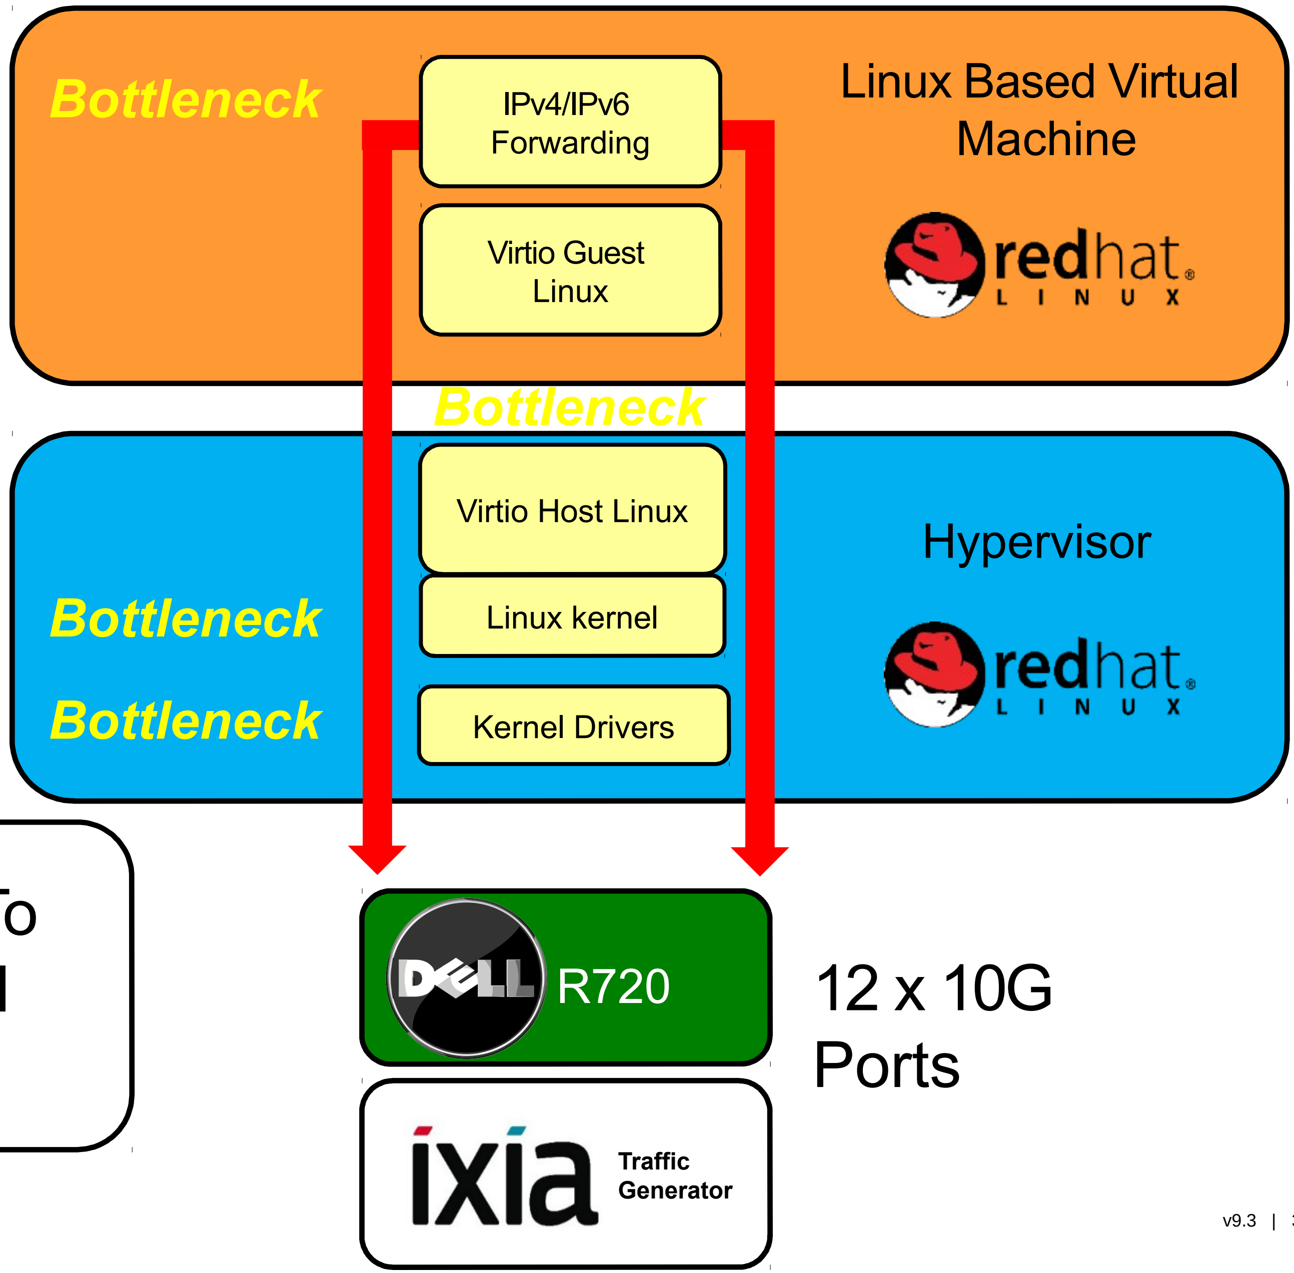

Limited Bandwidth To Linux Based Virtual Machines

Test 2:
6WINDGate OVS Acceleration + Fast vNIC Linux

L2 Throughput

7,2 Gbps
59 Gbps

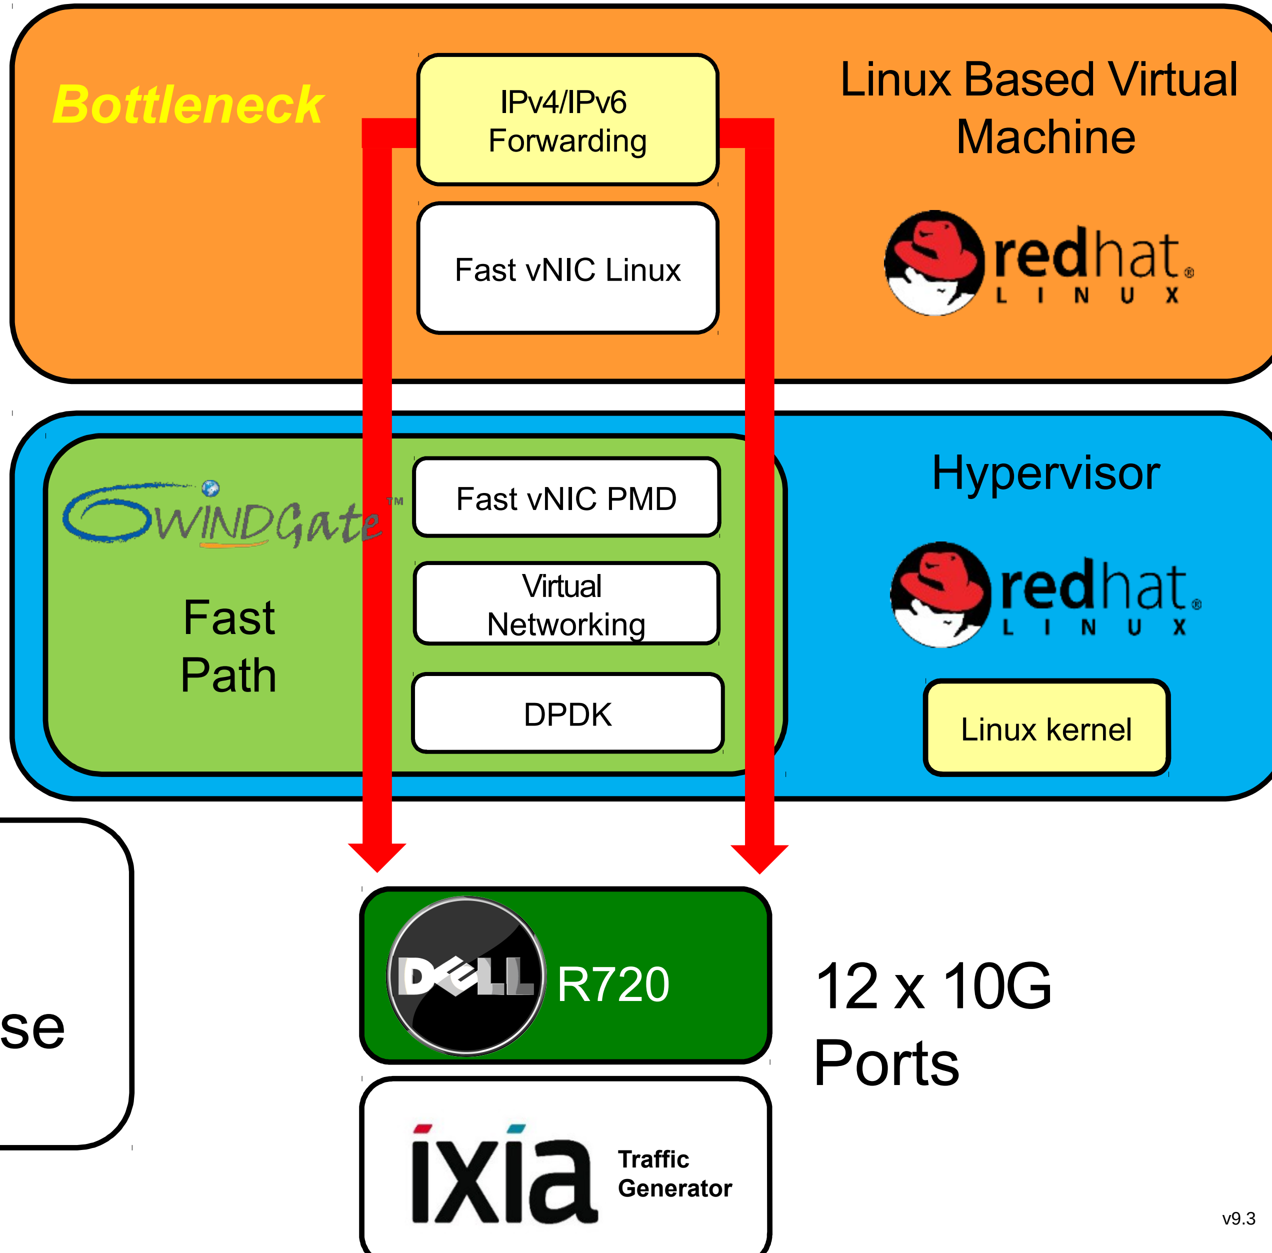

9X Throughput
Performance Increase

Test 3:
6WINDGate OVS Acceleration + Fast vNIC PMD

L2 Throughput

7,2 Gbps
59 Gbps
118 Gbps

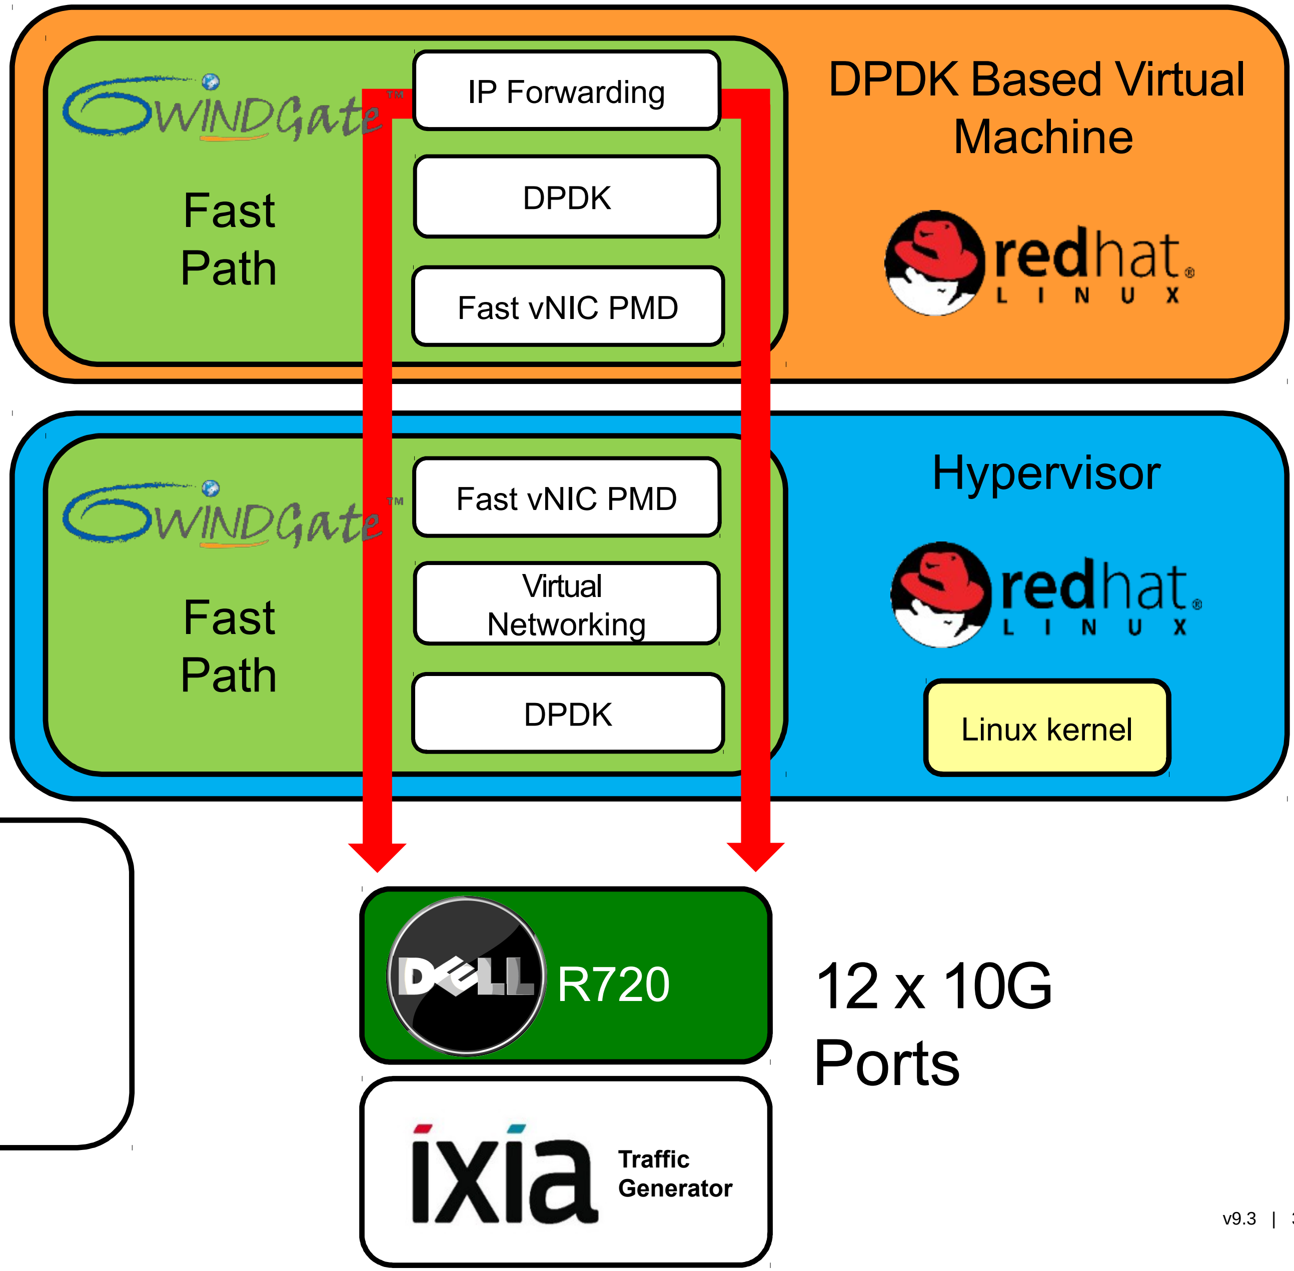

12 x 10G Ports

Wire Speed Performance

Compute nodes with 6WINDGate, Openstack-horizon

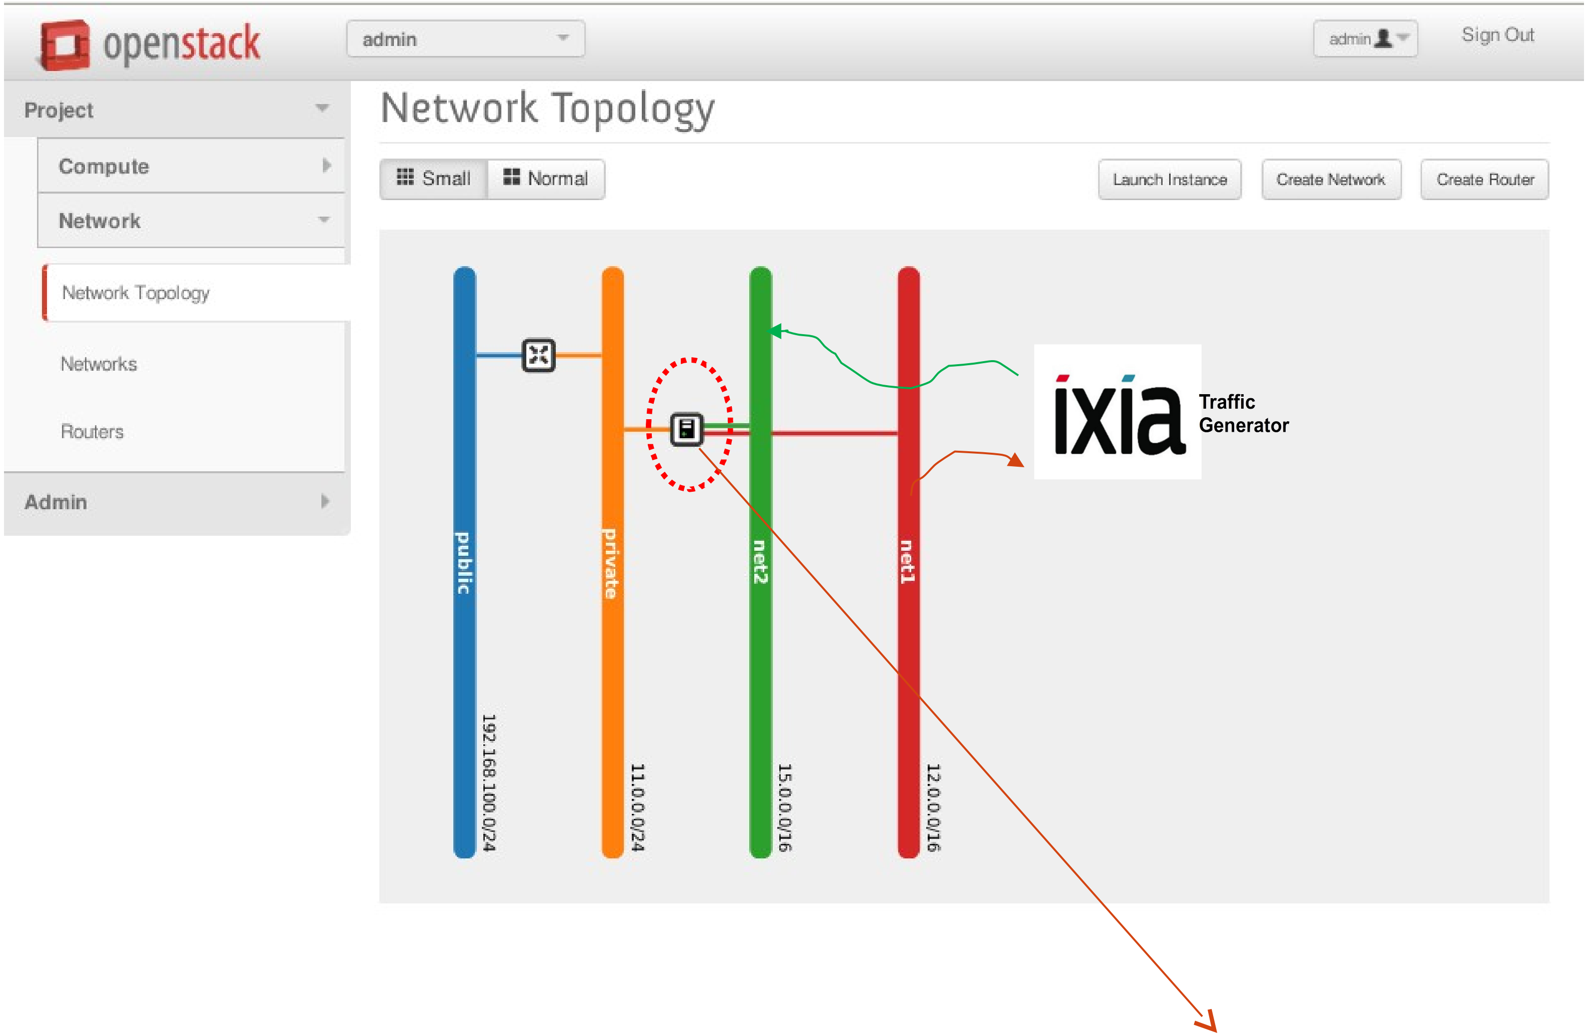

Compute node/host:
yum install 6windgate*.rpm
systemctl enable 6windgate.service

Virtual Switch-Based NFVI
Lowest Latency and Flexible Chaining

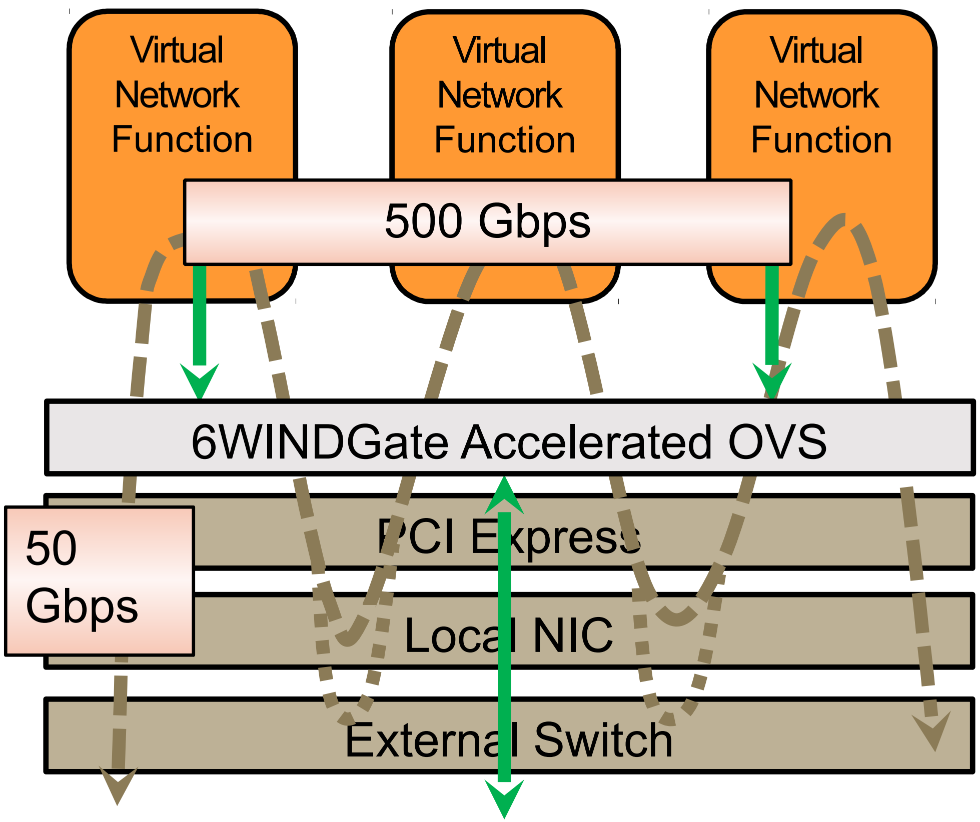

Virtual Switching With 6WINDGate

- Hardware independent virtual switching (NIC driver)
- Aggregate 500 Gbps bandwidth with low latency
- No external limit to number of chained VNFs

Physical Switching Limitations

- Hardware dependent switching (SR-IOV, RDMA, NIC embedded switching)
- Throughput is limited by PCI Express (50 Gbps) and faces PCI Express and DMA additional latencies
- Available PCI slots limit the number of chained VNFs
- At 30 Gbps a **single** VNF is supported per node!


6WIND
SPEED MATTERS
Turbo Boost Linux
The OEM Advantage
6WINDGate™
Packet Processing Software
Up To 10X Network Performance
Increase Data Plane Performance
No Change To Linux Environments
Portable Across All Major Platforms
Support Extensive Set Of Protocols
L2-L4 Acceleration
IPsec VPN Gateways
TCP / UDP Termination
Virtual Switching
DPDK
And More...
fast path
6WIND
Unlock Hidden Performance
Reduce Time-To-Market
Enable Transition To SDN / NFV

DPDK.org meetup – Jeudi prochain – Santa Clara @ 6WIND

THANK YOU!