

**RED HAT  
SUMMIT**

**BOSTON, MA  
JUNE 23-26, 2015**



# **CLOUD ARCHITECTURE & PERFORMANCE WORKLOADS**



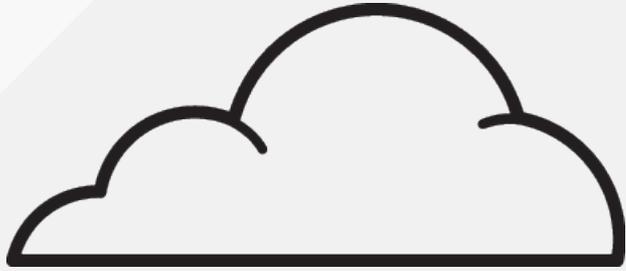
**Field Activities**

**Matt Smith**  
Senior Solution Architect  
Red Hat, Inc  
**@rhmjs**



**Jeremy Eder**  
Principal Performance Engineer  
Red Hat, Inc  
**@jeremyeder**

# CLOUD ARCHITECTURE

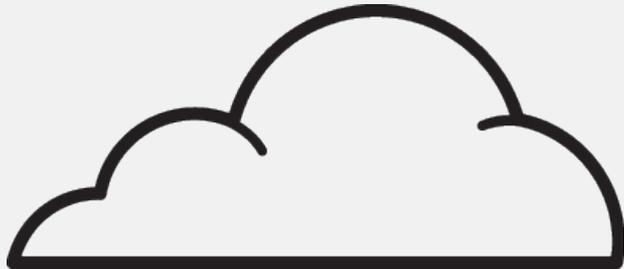
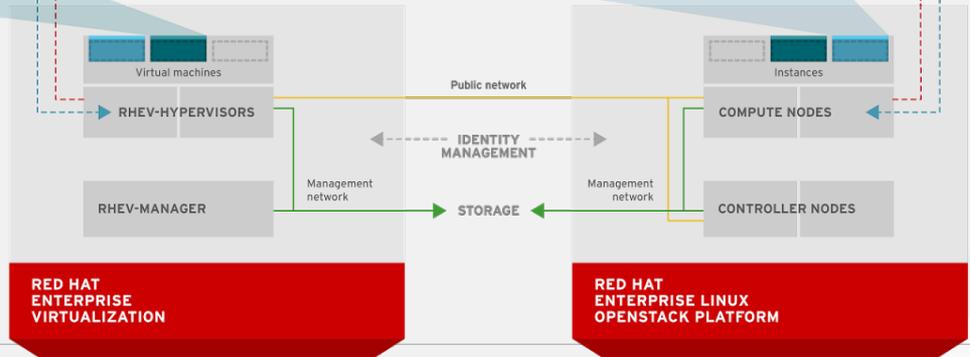
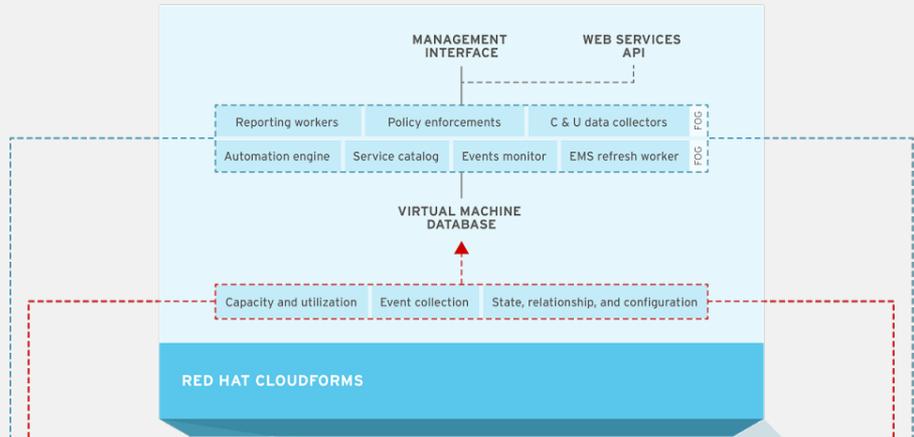
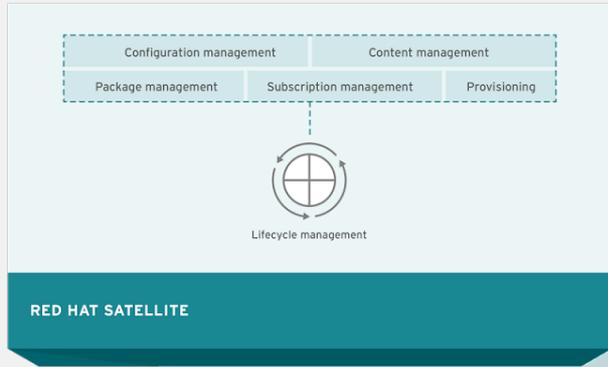


# “Cloud” - A Buzzword Review

- On-demand self-service
  - Accelerate Time-to-Value
  - Multi-tenancy
- Broad network access
  - Heterogeneous, Standardized Interfaces
- Resource pooling
  - Shared Compute, Network, Storage
- Rapid elasticity
  - Grow, Shrink, Repurpose
- Measured service
  - Metering and Reporting



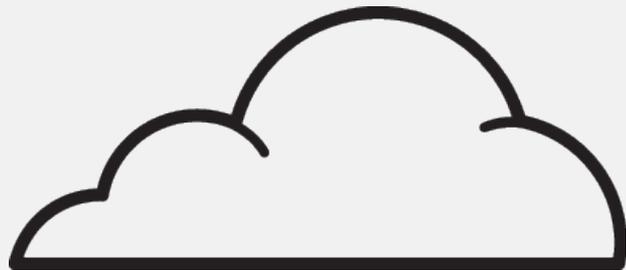
# Red Hat Cloud Infrastructure



# CLOUD ADOPTION

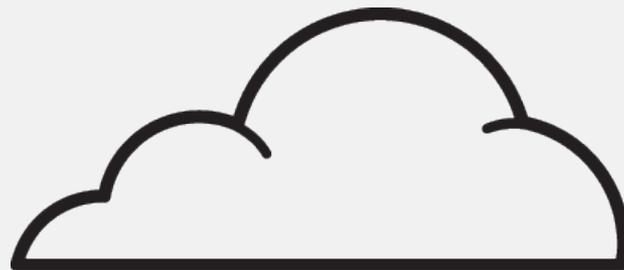
# Common Cloud Workloads

- Web Applications
  - Horizontal scalability/resiliency
  - Rebuild, don't repair
- Development Environments
  - On-demand self-service
  - Templates matching production
  - Automated testing



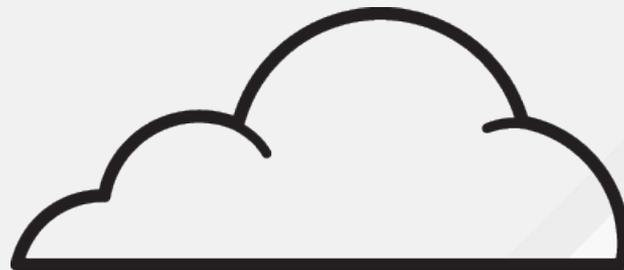
# Growing Cloud Workloads

- Distributed / Grid Compute
  - Resource Reallocation
- Big Data / Data Analytics
  - Horizontal Scalability/Resiliency
  - On-demand self-service “DAaaS”
  - Resource Reallocation



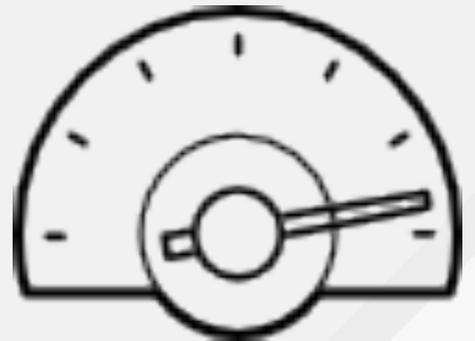
# Emerging Cloud Workloads

- Telco
  - Network Functions Virtualization (NFV)
  - Packet Processing/Analysis
- Retail / Industrial
  - “Internet of Things” (IoT)
- Financial Services Industry
  - Regulatory Compliance
  - Transactional, Low-Latency

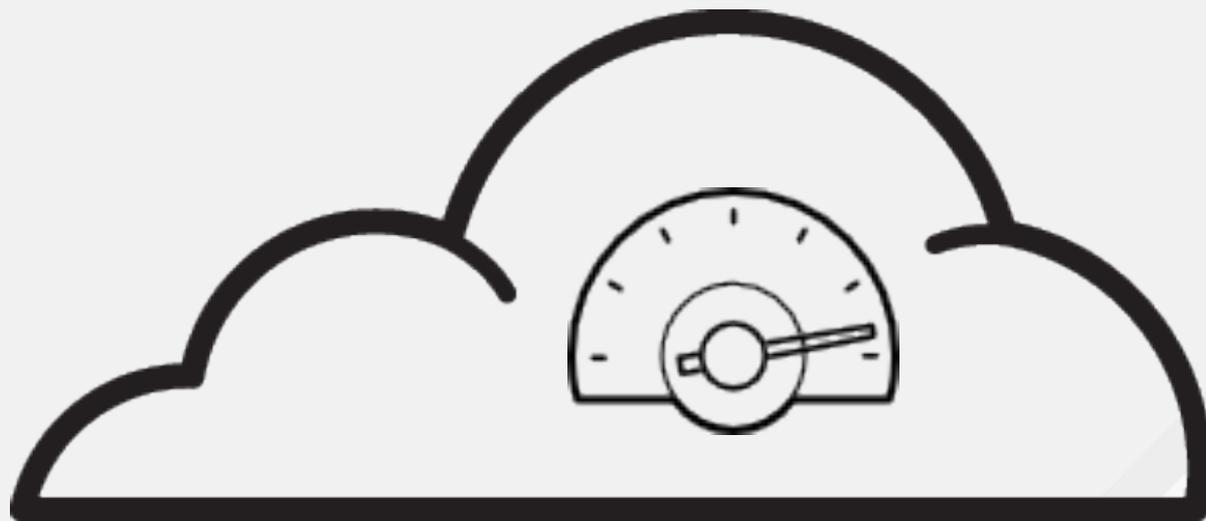


# Performance Requirements

- High Compute Throughput
- High Network Throughput
- Low Network Latency

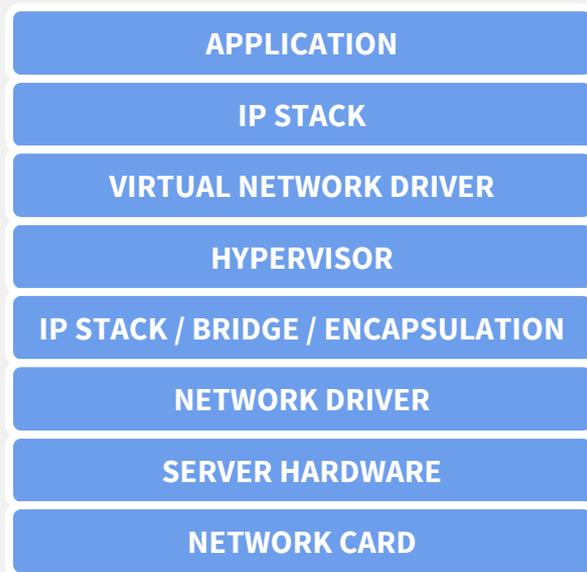


# Performance & Cloud ?

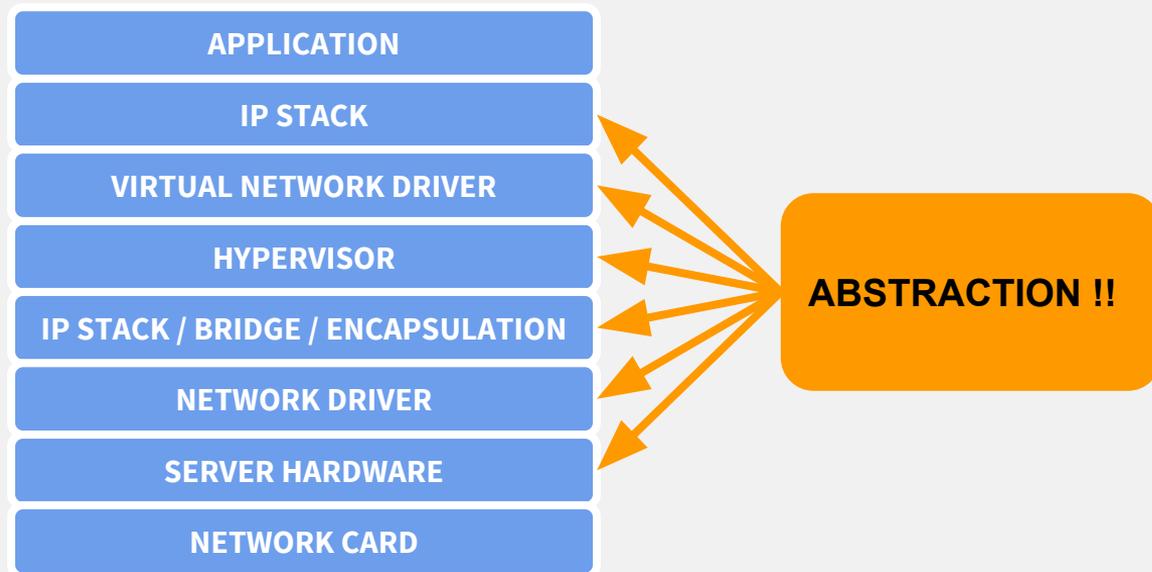


# TECHNICAL APPROACH

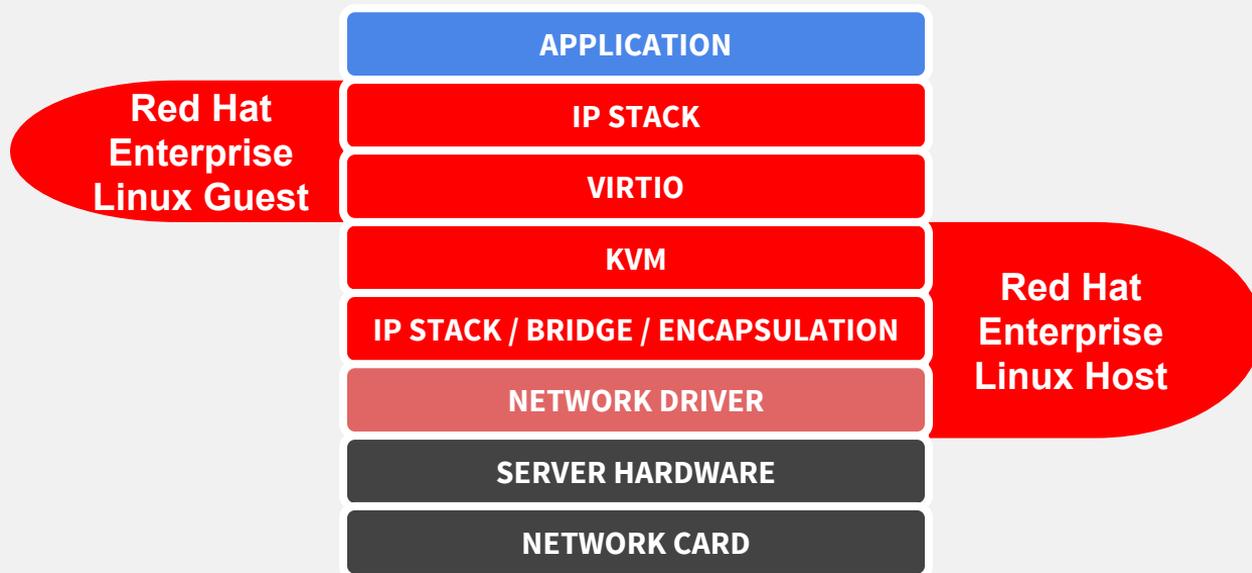
# Typical Virtualization Stack



# Typical Virtualization Stack



# Typical KVM Virtualization Stack

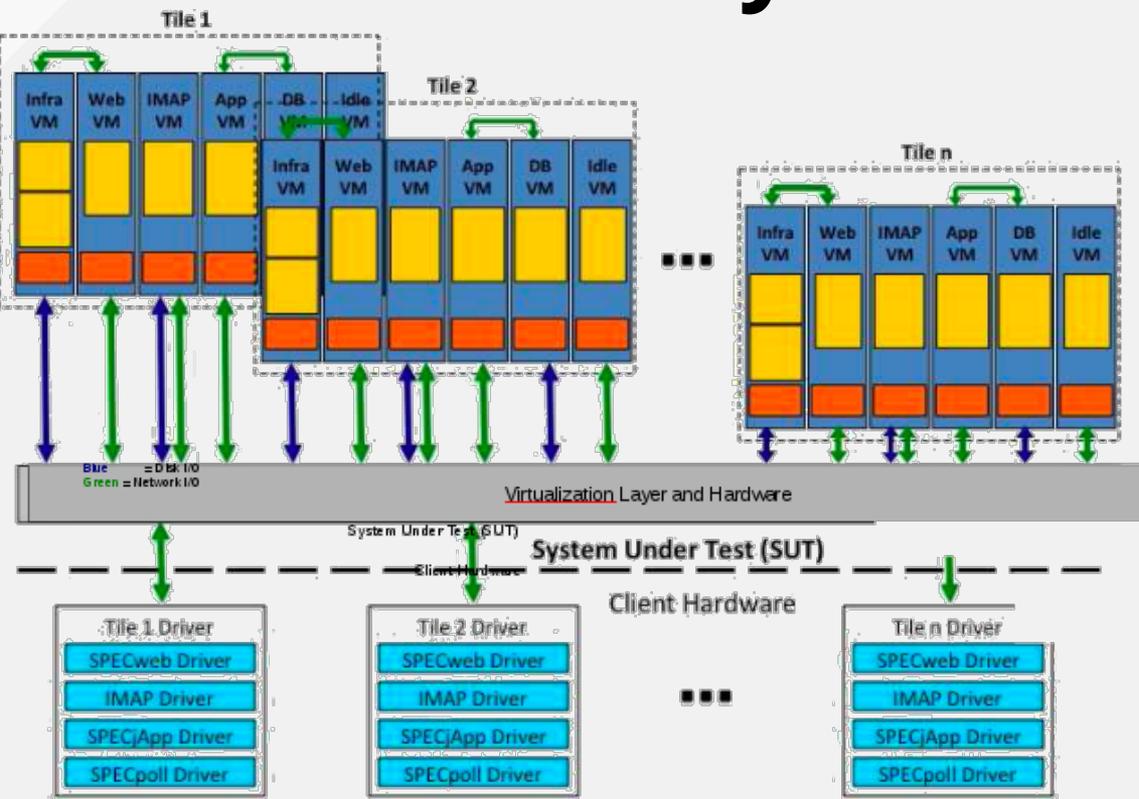


# Red Hat Enterprise Linux

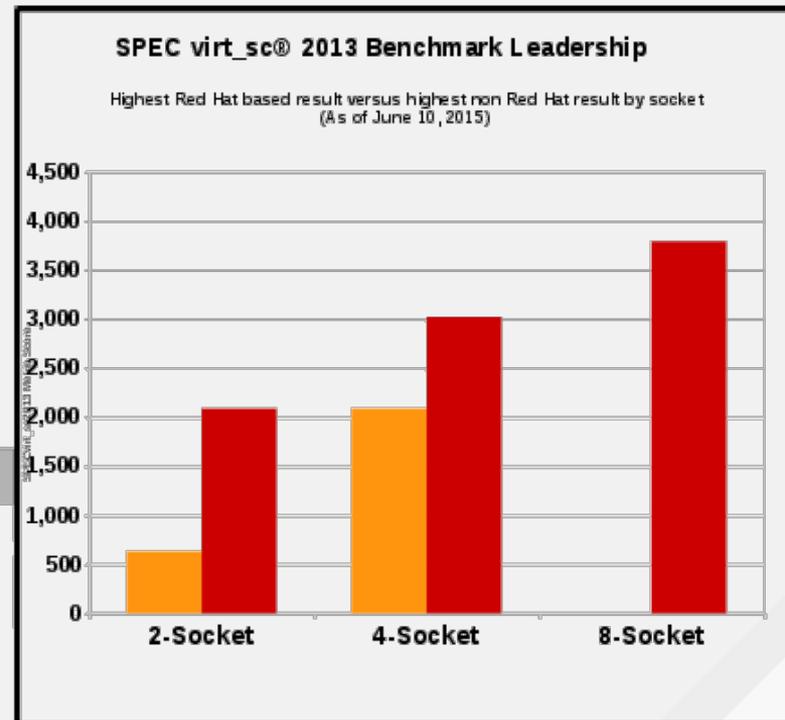
```
# tuned-adm profile virtual-host
```

```
# tuned-adm profile virtual-guest
```

# KVM: Industry Leading Results



[http://www.spec.org/virt\\_sc2013/results/](http://www.spec.org/virt_sc2013/results/)



# OpenStack & NUMA

- Non-Uniform Memory Access
- The NUMA topology defines the locality of CPU and Memory
- Processes can be bound to NUMA nodes for optimal performance

# OpenStack & NUMA

- KVM awareness of NUMA aligns the guests virtual NUMA topology with the physical topology of the host
- RHEL-OSP 6 supports specifying NUMA requirements in the “flavor”

# OpenStack NUMA Configuration

Instance with 8 vCPUs and 4GB RAM

```
hw:numa_nodes=2
```

```
hw:numa_cpus.0=0,1,2,3,4,5
```

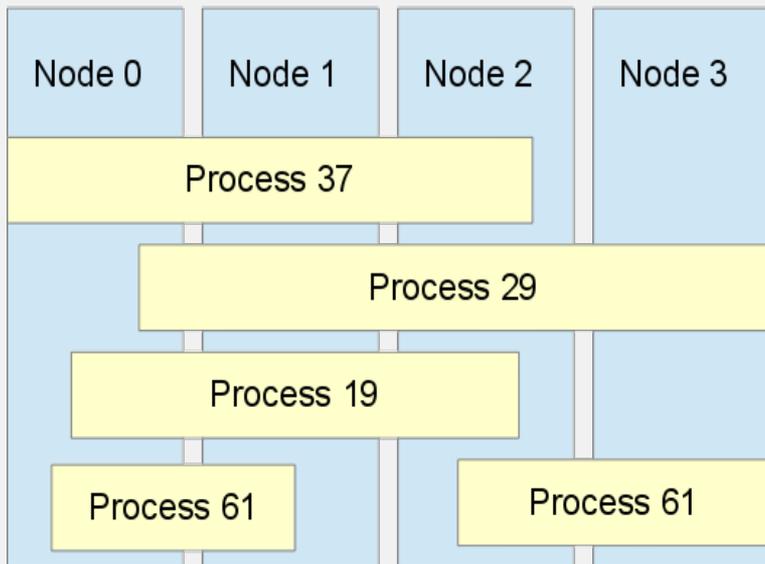
```
hw:numa_cpus.1=6,7
```

```
hw:numa_mem.0=3
```

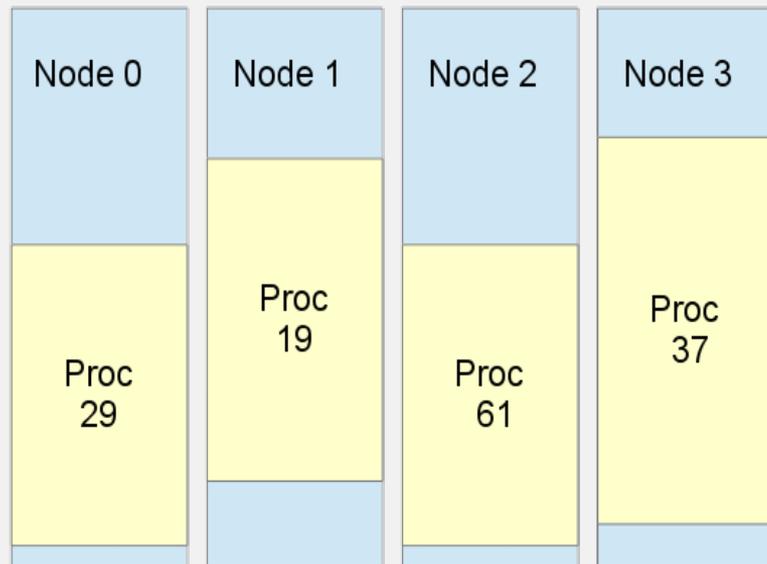
```
hw:numa_mem.1=1
```

# NUMA: Process == KVM thread

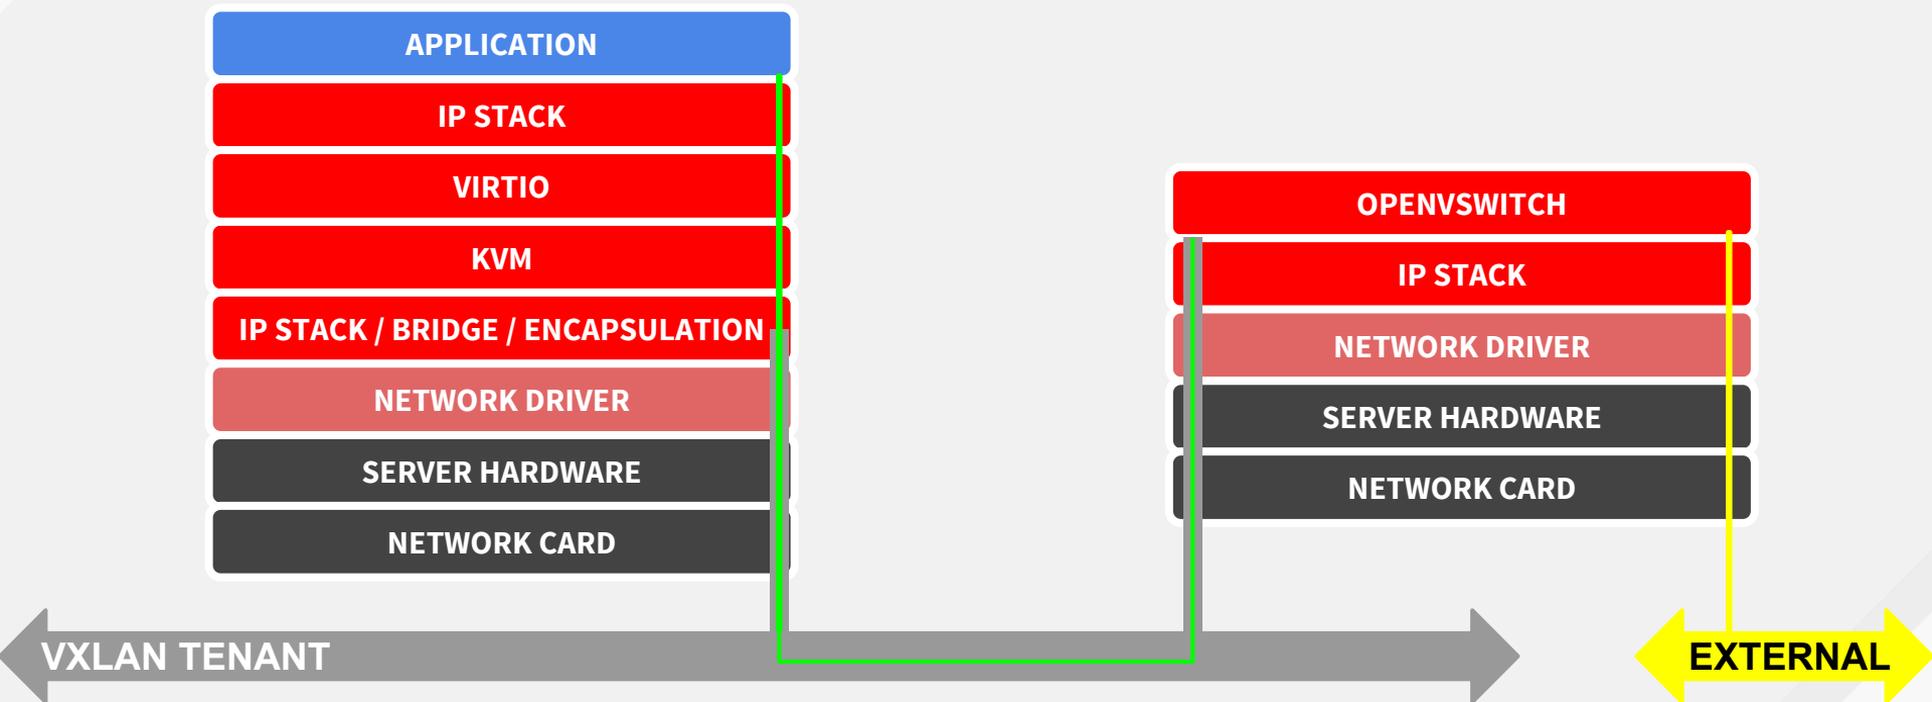
No NUMA scheduling



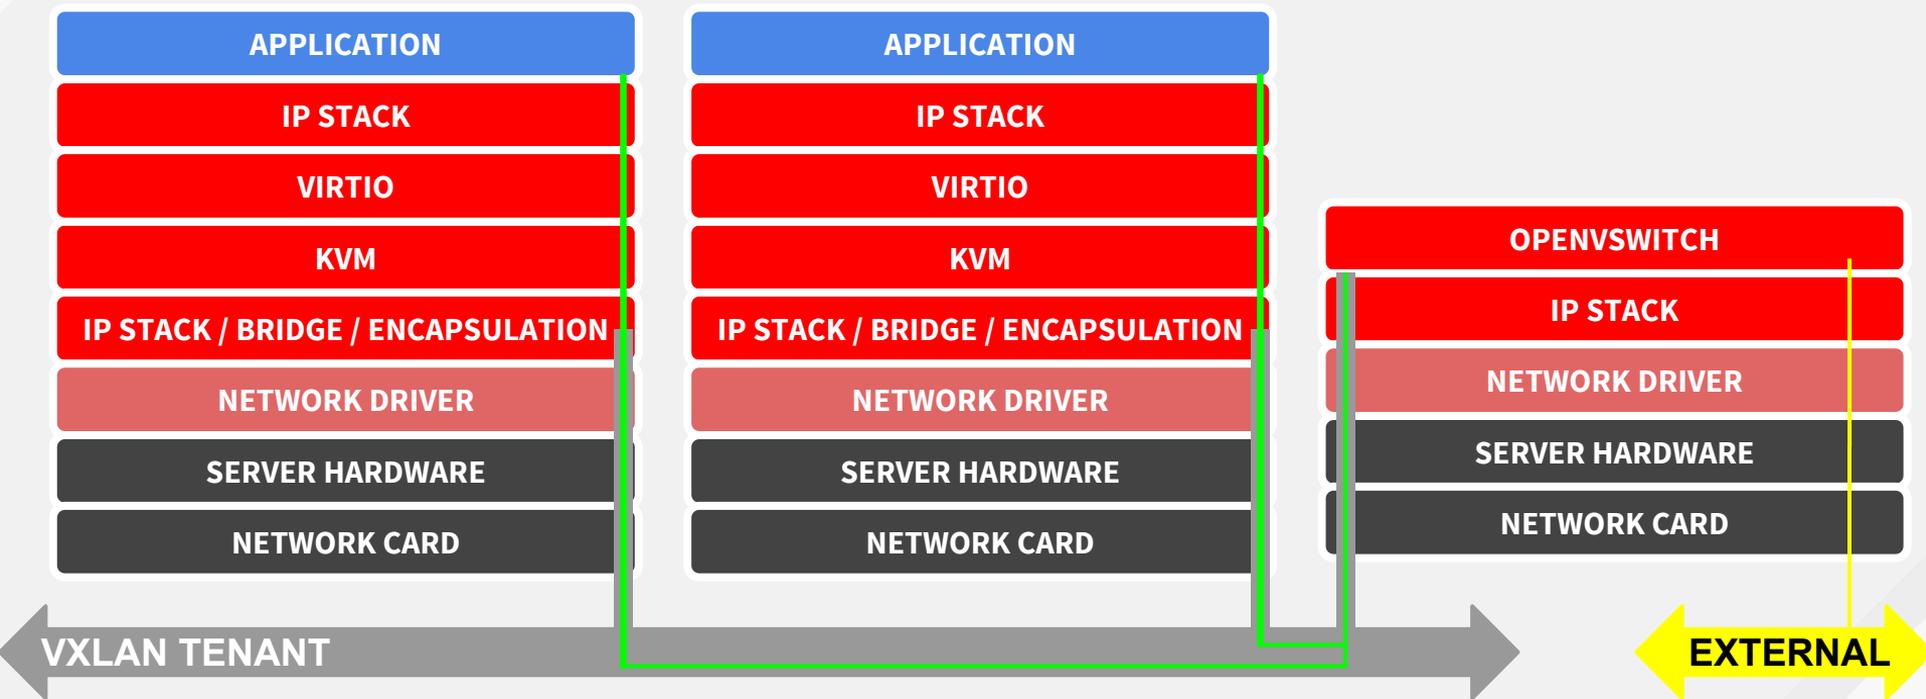
With NUMA Scheduling



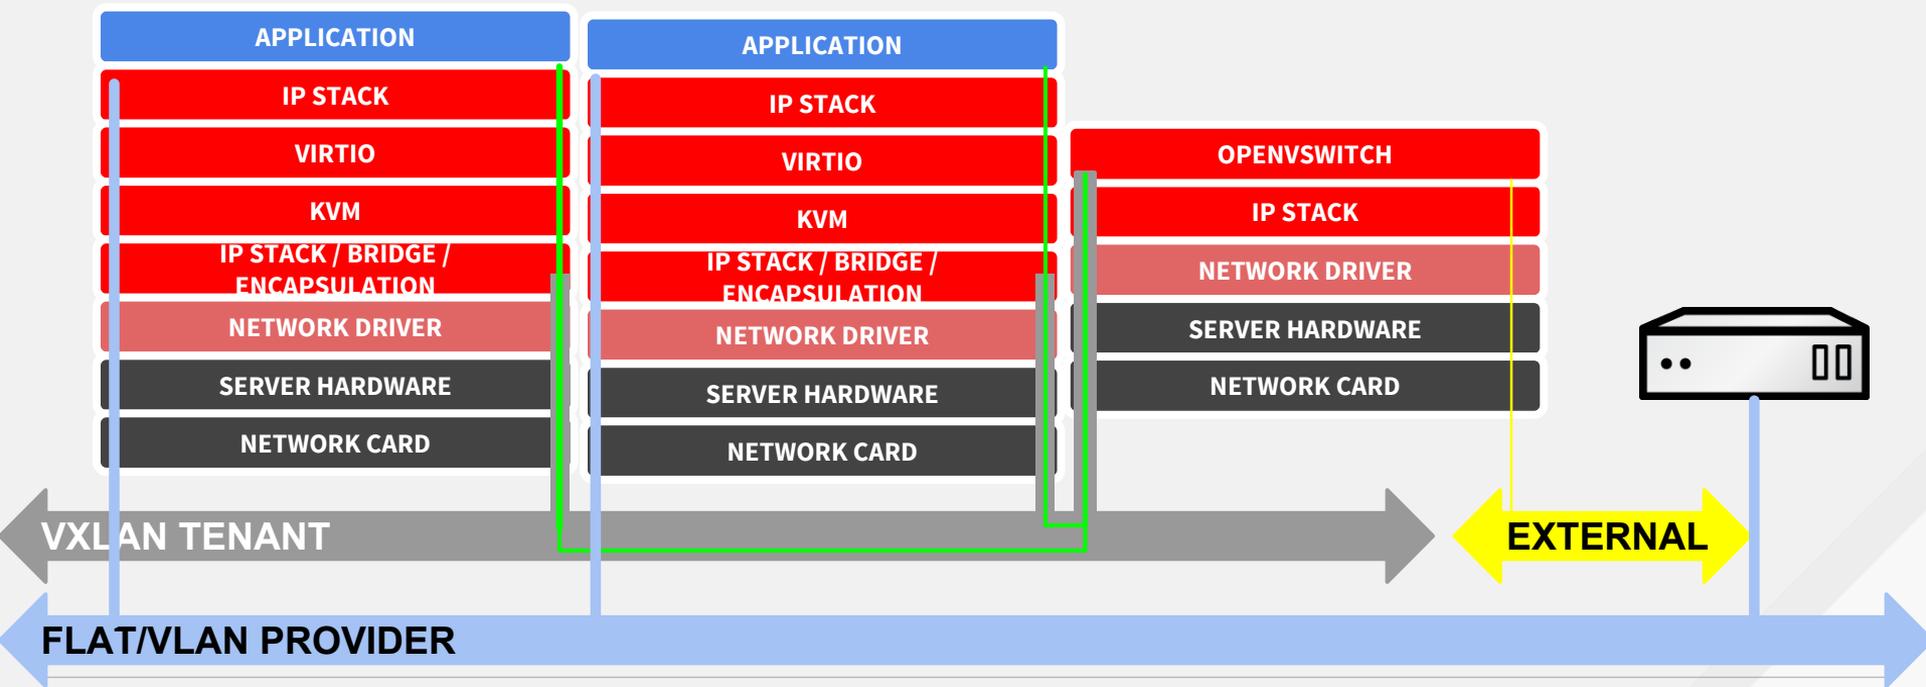
# Simple OpenStack Deployment



# Typical OpenStack Deployment



# OpenStack Provider Network



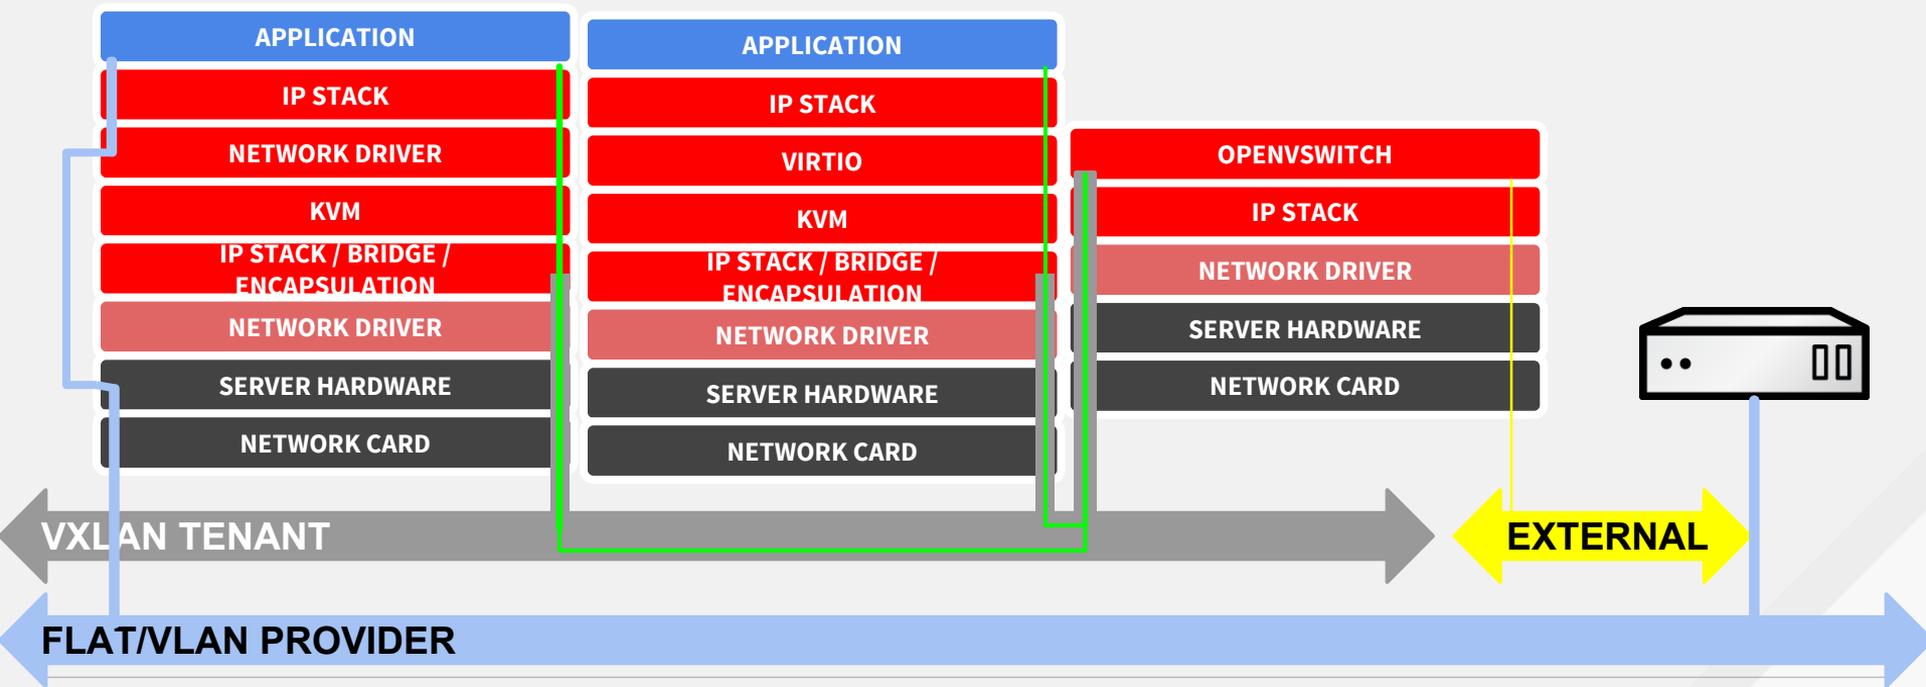
# Single Root I/O Virtualization (SR-IOV)

- Physical Functions (PFs): full PCIe devices that include the SR-IOV capabilities
- Virtual Functions (VFs): simple PCIe functions, derived from PFs, that only process I/O
- KVM Passthrough of VFs to Guests

# OpenStack + SR-IOV

- Ensure that Intel VT-d or AMD IOMMU are enabled in the BIOS and Operating System
- Load the driver with VF configuration
- Enable the **sriovnicswitch** driver in Neutron
- Enable SR-IOV on the Controller and Compute nodes
- Create an SR-IOV port on a Neutron network
- Launch an instance configured to use the port

# OpenStack with SR-IOV



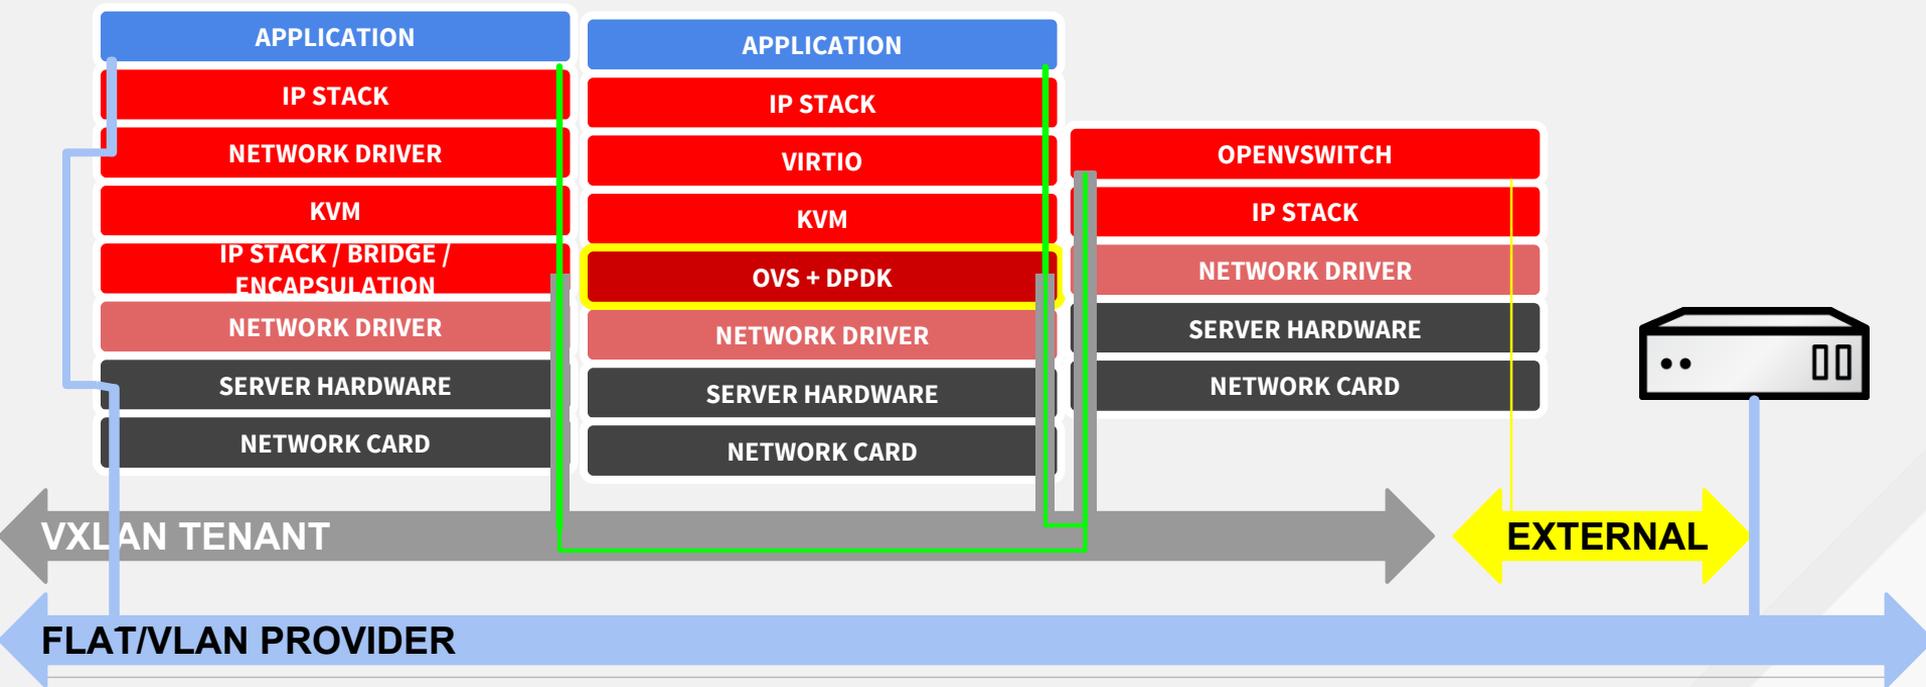
# Data-plane Development Kit (DPDK)

- Primary development by Intel and 6WIND
- Set of libraries and drivers for fast packet processing
- OpenVSwitch with DPDK for accelerated packet processing with a focus on NFV use-cases, under active development



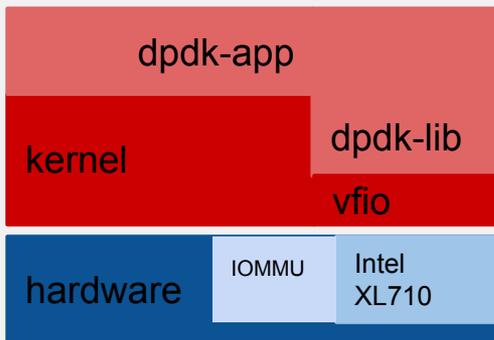
**Leading  
Edge**

# OpenStack with OVS+DPDK

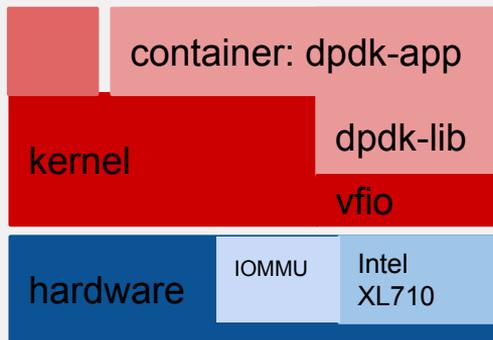


# DPDK config, bare-metal, container, kvm

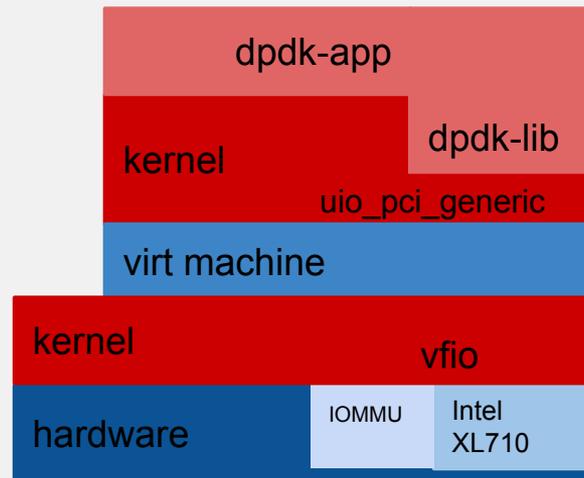
bare-metal



PF assigned to container



PF assigned to VM

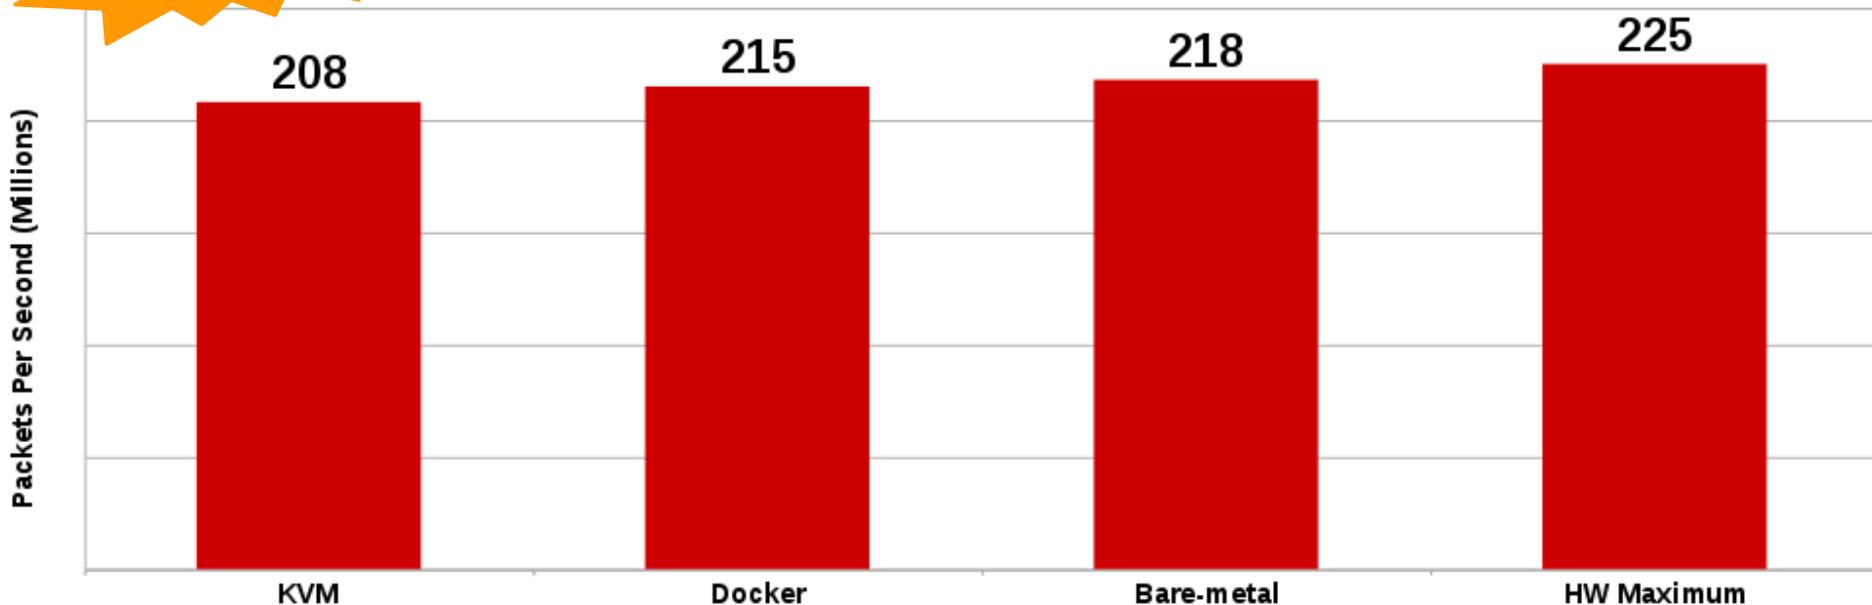


# Network Function Virtualization (NFV) Throughput and Packets/sec (RHEL7.x+DPDK)

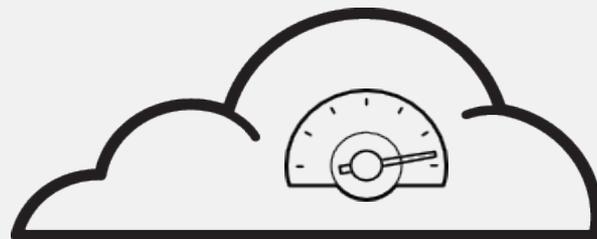
208Mpps+  
INTO KVM  
DPDK

NFV: Millions of Packets Per Second

RHEL7.x, L2 Forwarding, 12 x 40Gb NICs



# Technology Review



- Red Hat Enterprise Linux
- Kernel-based Virtual Machine (KVM)
- Non-Uniform Memory Access (NUMA)
- OpenStack Flat/VXLAN Network Topologies
- Single Root I/O Virtualization (SR-IOV)
- OpenVSwitch (OVS) + Data-plane Development Kit (DPDK)

# Other Cloud/Performance Sessions

- **Performance analysis & tuning of Red Hat Enterprise Linux**
  - Wednesday, June 24 1:20 pm - 3:20 pm
- **Performance of OpenStack Cinder on Ceph**
  - Thursday, June 25 4:50 pm - 5:50 pm
- **Containers versus virtualization**
  - Friday 9:45 am - 10:45 am
- **Open source & network functions virtualization**
  - Friday 11:00 am - 12:00 pm

**Matt Smith**  
**@rhmjs**

[verticalindustriesblog.redhat.com](http://verticalindustriesblog.redhat.com)



**Jeremy Eder**  
**@jeremyeder**

[developerblog.redhat.com](http://developerblog.redhat.com)  
[www.breakage.org](http://www.breakage.org)

**RED HAT**  
**SUMMIT**

**LEARN. NETWORK.  
EXPERIENCE OPEN SOURCE.**