

ACHIEVE LOW LATENCY NFV WITH OPENSTACK*

Yunhong Jiang Yunhong.Jiang@intel.com



*Other names and brands may be claimed as the property of others.

Agenda

NFV and network latency

Why network latency on NFV

How to achieve low latency NFV

Setups and configurations



NFV AND NETWORK LATENCY

Low network latency is important

Voice quality for VOIP service - 150ms

User experience for interactive game – 100ms

Quality of service for RAN (Radio Access Networks) – 1ms

Profit for algorithmic trading – microseconds



NFV and network latency

What is NFV

- Running VNF (Virtualized Network Function) on high volume servers

Benefit of NFV

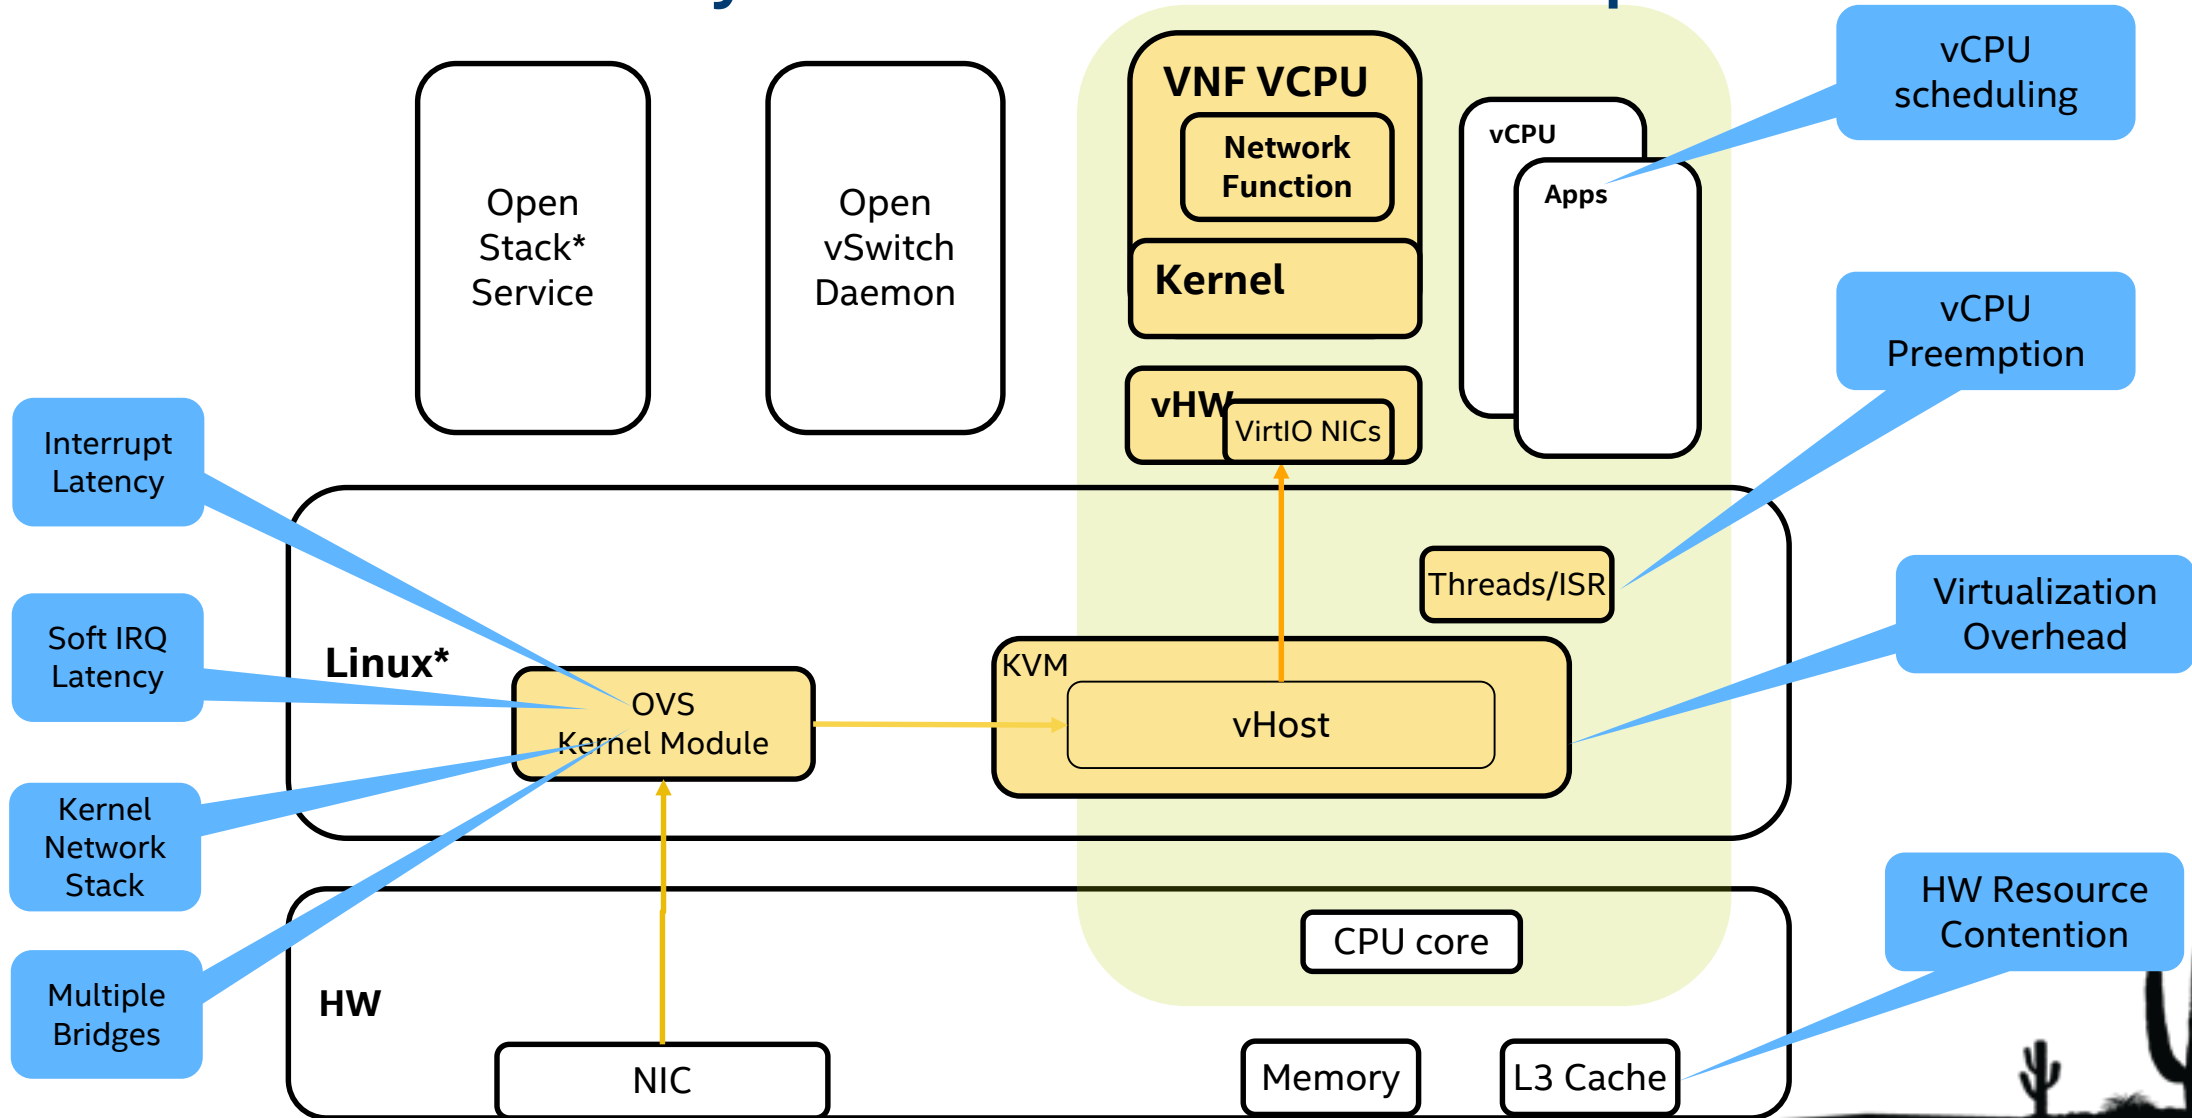
- Flexibility
- Manageability
- CAPEX and OPEX
- Time-to-market

BUT NFV MAY INCREASE LATENCY!



WHY NETWORK LATENCY ON NFV

Network latency on NFV – An example



Root causes of Latency on NFV

Resource contention

- CPU/Cache/TLB/Memory/IO
- vCPU/apps/services/kernel service

Virtualization overhead

- CPU/IO virtualization
- Network virtualization



HOW TO ACHIEVE LOW LATENCY NFV

How to achieve low latency NFV

Resource contention

- Reduce the contention cost
- Priority VNF vCPU for the contention
- Exclusively resource assignment

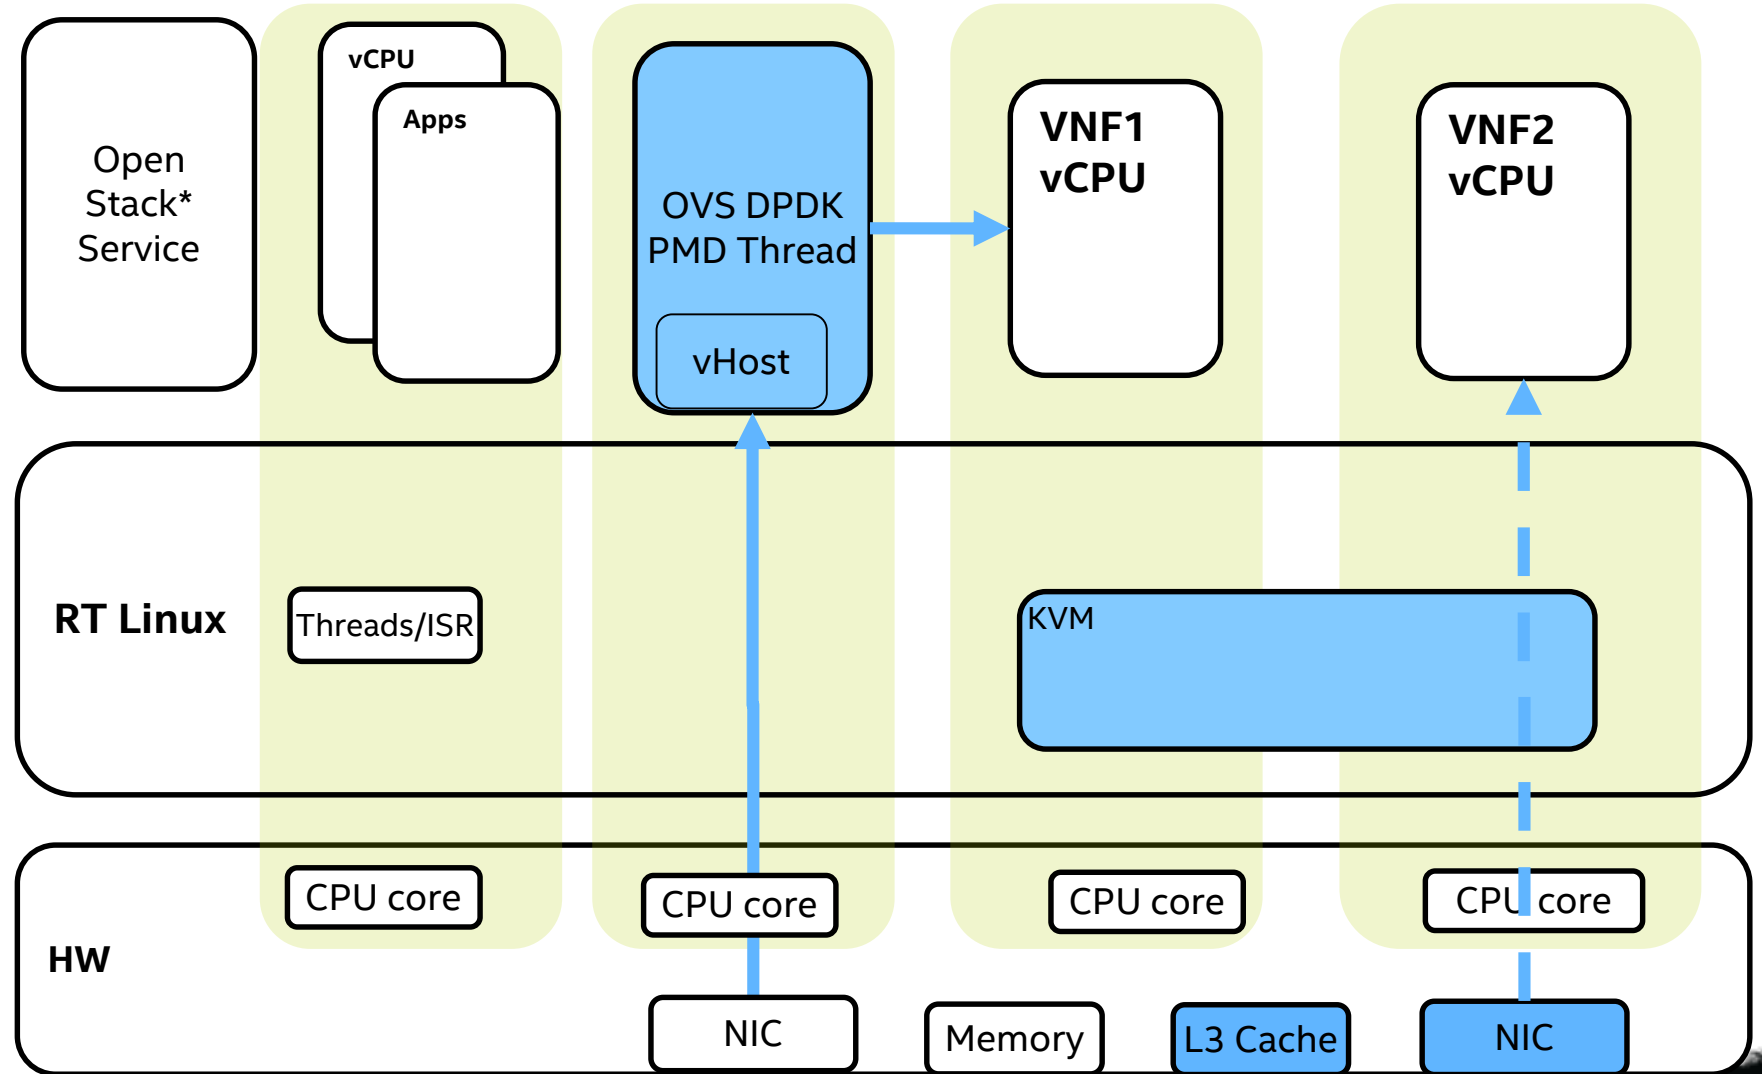
Virtualization overhead

- Advanced virtualization features and hypervisor
- DPDK OVS for network virtualization or SR-IOV




A systematic solution for low latency NFV

- Advanced hardware platform
- Real-time Hypervisor
- DPDK OVS
- OpenStack*



Advanced hardware platform



VT-d and SR-IOV

- I/O device reservation

PI (Posted Interrupt)


- Reducing interrupt virtualization overhead

CAT (Cache Allocation Technology)

- Cache reservation

.....

Real-time hypervisor/Linux*



Real-time scheduler

Full preemption support

All activities (ISR, soft IRQ etc) on thread context

CPU isolation support

.....

KVM4NFV Project


An OPNFV project <https://wiki.opnfv.org/nfv-kvm>

- Based on upstream real-time Linux*
- Included in Brahmaputra release
- Will be integrated into OPNFV test framework in C release

Active contributors from OPNFV community

- Intel, Nokia, Wind River ...

DPDK OVS for low latency NFV



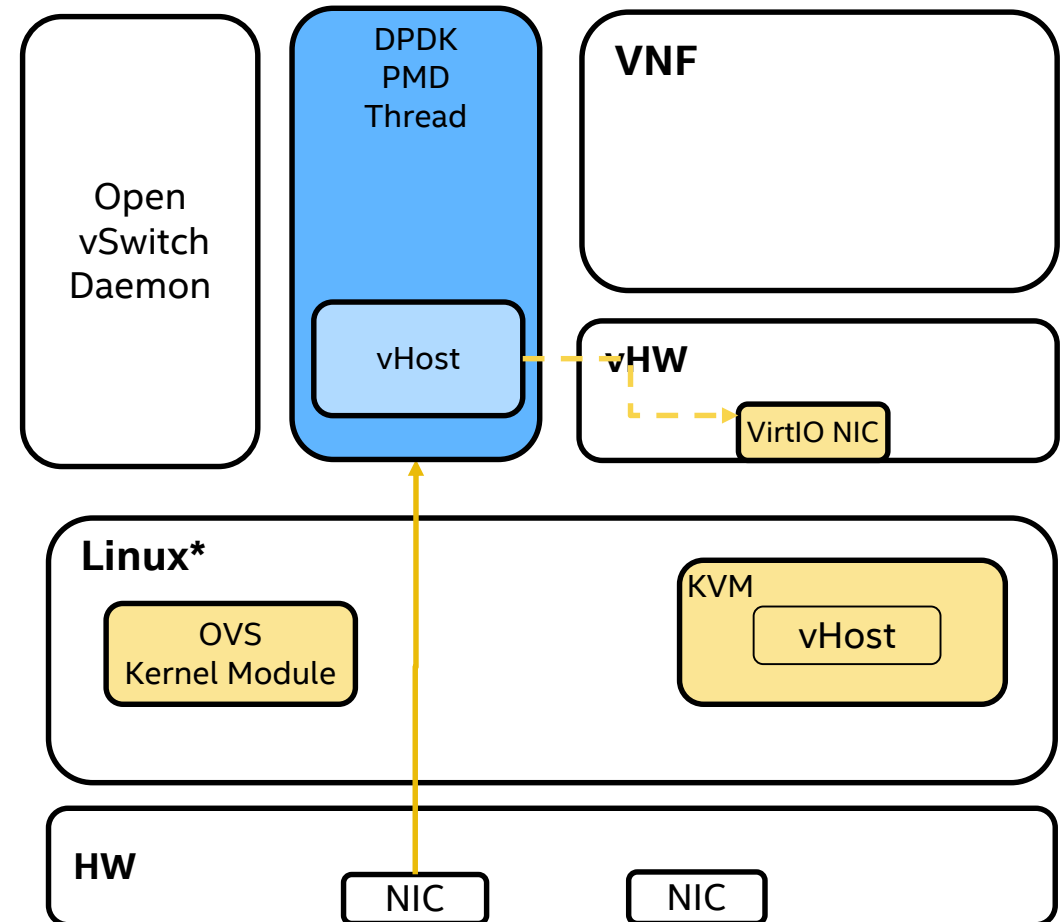
Data Plane Development Kit (DPDK)

- Improves packet processing on Intel® architecture
- Bypasses Linux* network stack and maps hardware registers to user space.
- PMD (Polling Model Driver) thread
- Queue & buffer management, packet flow classification, memory management


DPDK OVS for low latency NFV

Accelerate OVS data plane with DPDK

- PMD thread avoids the interrupt/soft IRQ latency
- Queue & buffer management
- User space vHost



OpenStack*- Nova



Resource tracking on compute node

- Avoid VM resource contention

VM creation

- Reserved resources allocation
- vCPU thread scheduler priority
- Guest NUMA support
- Guest huge page support



OpenStack*- Nuutron


vNIC port with SR-IOV NIC

vNIC port with DPDK OVS and vHost-user

SETUPS AND CONFIGURATIONS



Host configuration



RT Linux* setup

- Static kernel boot parameter (isolcpus/rcu_nocbs/nohz_full)
- Runtime configuration (Interrupt affinity/watchdog/RCU CB

DPDK OVS

- Upgrade to DPDK OVS service
- DPDK PMD thread affinity



Configure OpenStack*

Compute node resource reservation

- vcpu_pin_set/cpu_allocation_ratio/ram_allocation_ratio
- pci_passthrough_whitelist

VM flavor

- hw:cpu_realtime_mask/hw:cpu_thread_policy/
hw:cpu_realtime_mask/hw:cpu_realtime/hw:mem_page_size/pci:alias

Configure OpenStack*

Network node for DPDK OVS *

- datapath_type/ vhostuser_socket_dir

Network node for SR-IOV

- supported_pci_vendor_devs/ physical_device_mappings

CALL TO ACTION

- Document
- Automatic host deployment
- Automatic latency test



Configuration Reference

<https://wiki.opnfv.org/display/kvm/Nfv-kvm-tuning>

<https://github.com/openstack/networking-ovs-dpdk>

<https://github.com/openvswitch/ovs/blob/master/INSTALL.DPDK.md#performance-tuning>

http://docs.openstack.org/liberty/networking-guide/adv_config_sriov.html

Legal notices and disclaimers



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>.

Intel, the Intel logo and Intel architecture are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© 2016 Intel Corporation.



experience
what's inside™