



OVS

Open vSwitch

December 10-11, 2019 | Westford, MA

# Utilizing DPDK Virtual Devices in OVS

Yinan Wang / Lei Yao  
Intel

# DPDK Virtual Devices Introduction

## Virtual Devices for Exceptional Path

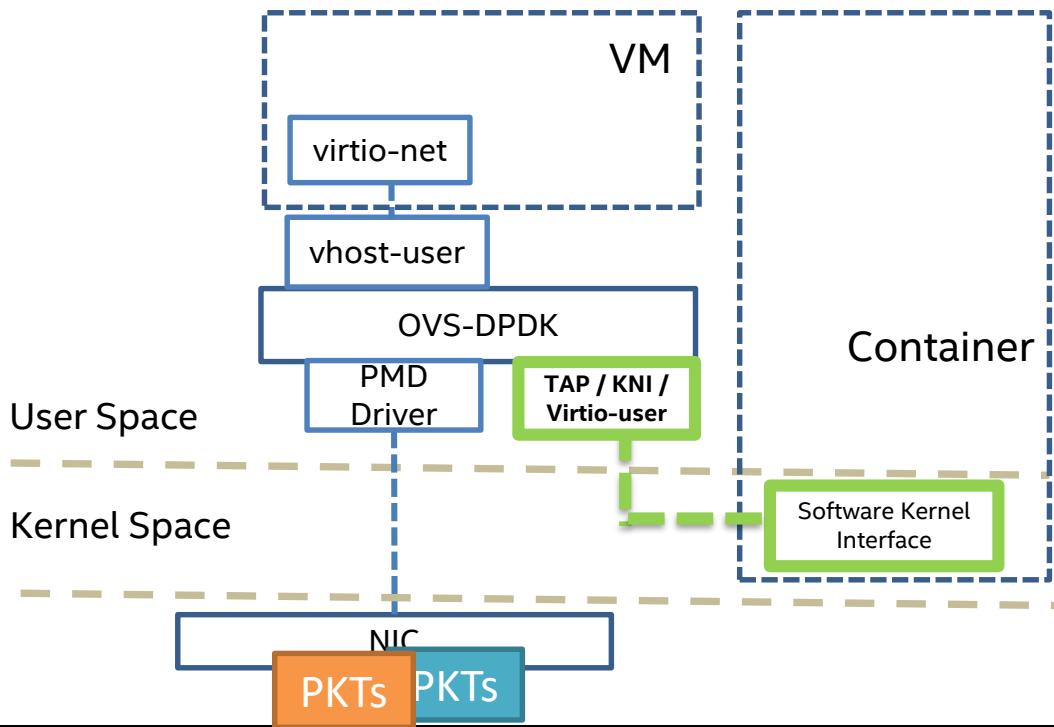
- TAP
- KNI
- Virtio-user

## Virtual Devices Based on NIC in Kernel

- PCAP
- AF\_PACKET
- AF\_XDP

# Virtual Devices for Exceptional Path

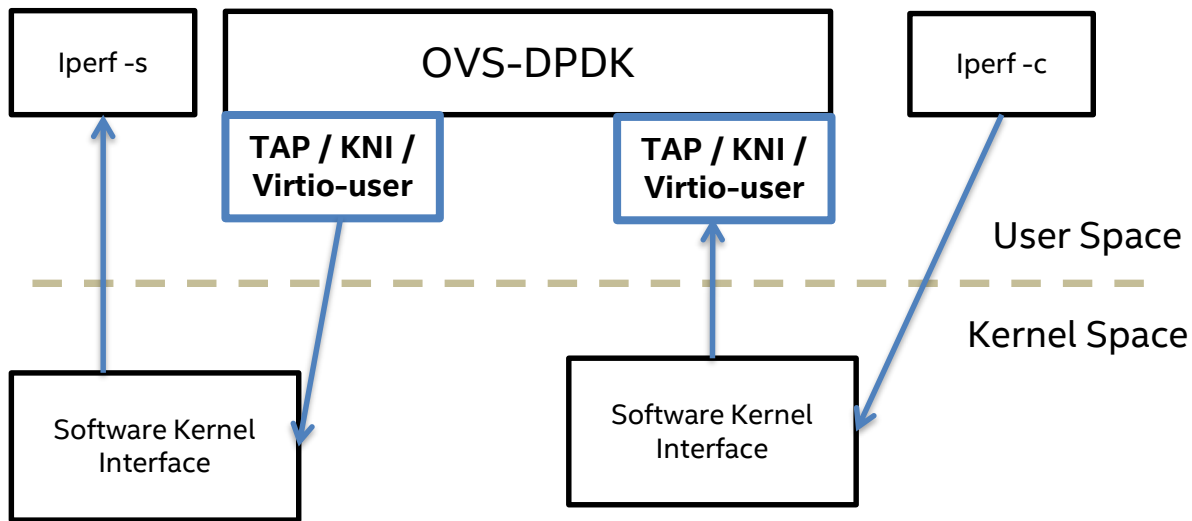
## ❑ TAP / KNI / Virtio-user



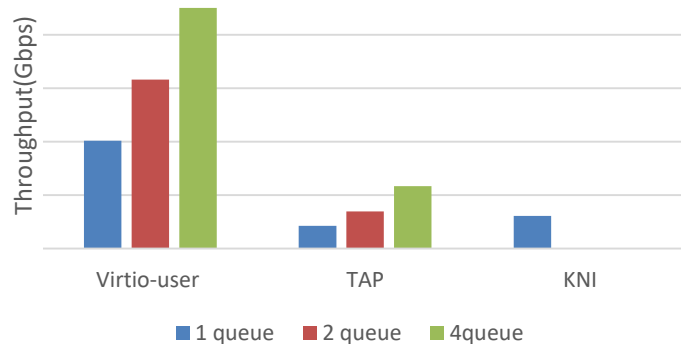
Potential usage in OVS-DPDK scenario:  
Container network.

# Virtual Devices for Exceptional Path

## Single Host Iperf Performance in OVS-DPDK



OVS-DPDK Iperf Performance



	Multi-queues
TAP	Y
KNI	N
Virtio-user	Y

# Virtual Devices for Exceptional Path

## BKMs on supporting three VDEVs in OVS-DPDK:

### TAP:

Normally, OVS-DPDK calculate Tx queue number according total port numbers.  
TAP device require (`rxq number = txq number`).

### KNI:

Launch OVS-DPDK with KNI, KNI's MTU is **2034** by default .

KNI\_MTU= (`mbuf_size`) - `RTE_ETHER_HDR_LEN`

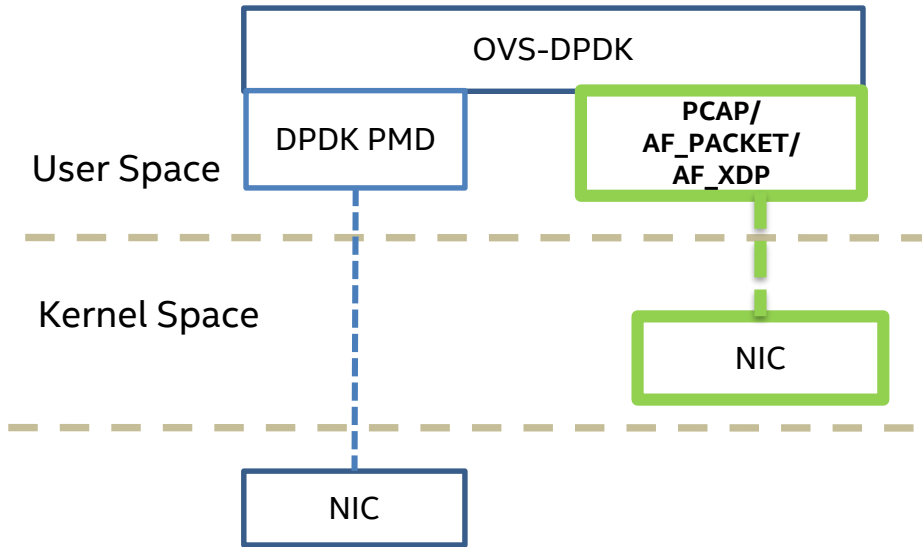
Need re-config the MTU for KNI interface in kernel side for better usage.

### Virtio-user:

1G hugepage	<i>ok</i>
2M hugepage	<i>ovs-vsctl --no-wait set Open_vSwitch . other_config:dpdk- extra="--single-file-segments"</i>
4K page	<i>System request: VT-d on; NIC bind to vfio-pci ; ovs-vsctl --no-wait set Open_vSwitch . other_config:dpdk- extra="--no-huge"</i>

# Virtual Devices Based on NIC in Kernel

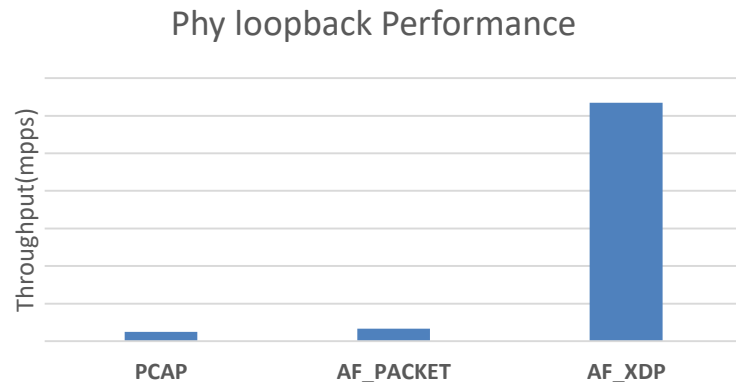
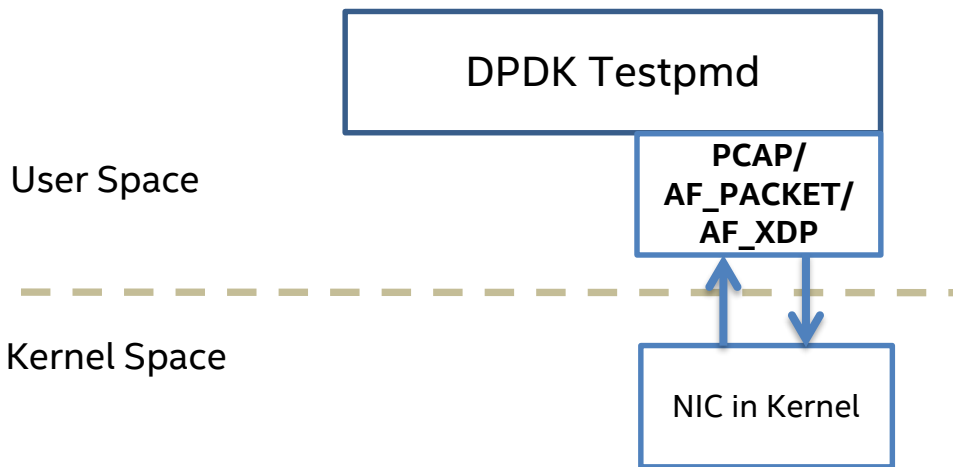
## ❑ PCAP/ AF\_PACKET/ AF\_XDP



Potential usage of these VDEVs in OVS-DPDK scenario:  
Receive packets from kernel NIC to user space w/o nic-pmd.

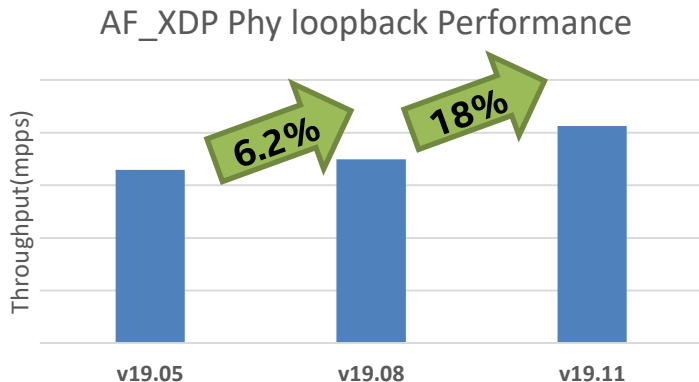
# Virtual Devices Based on NIC in Kernel

## ❑ PCAP/ AF\_PACKET/ AF\_XDP Phy loopback Performance



# Virtual Devices Based on NIC in Kernel

## □ AF\_XDP performance gains in DPDK



Two tips for using AF\_XDP in OVS-DPDK:

- AF\_XDP has dependencies on kernel version, recommend **v5.4**, includes some bug fixes
- `ovs-vsctl add-port br1 afxdp -- set Interface afxdp type=dpdk options:dpdk-devargs=net_af_xdp0,iface=[interface name], start_queue=[x],queue_count=[x]`





Thank You!

