SD-WAN + Business Forum
May 4, 2017

Understanding Forwarding Performance Characteristics of NFV-nodes

Ryota Kawashima

Nagoya Institute of Technology, Japan

Tsunemasa Hayashi & Hiroki Nakayama

BOSCO Technologies, Inc., Japan

What is the NFV-node ?

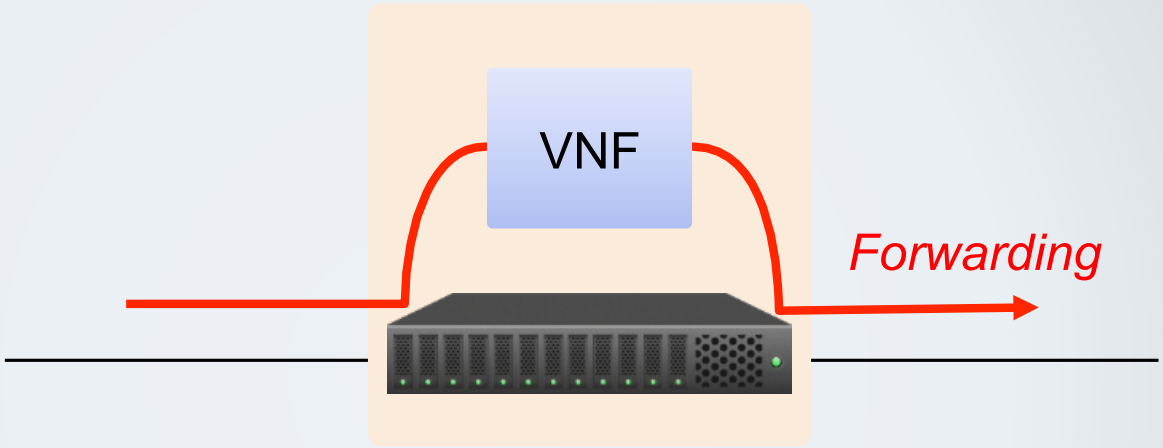

How much the node efficiently forwards packets ?

Goals

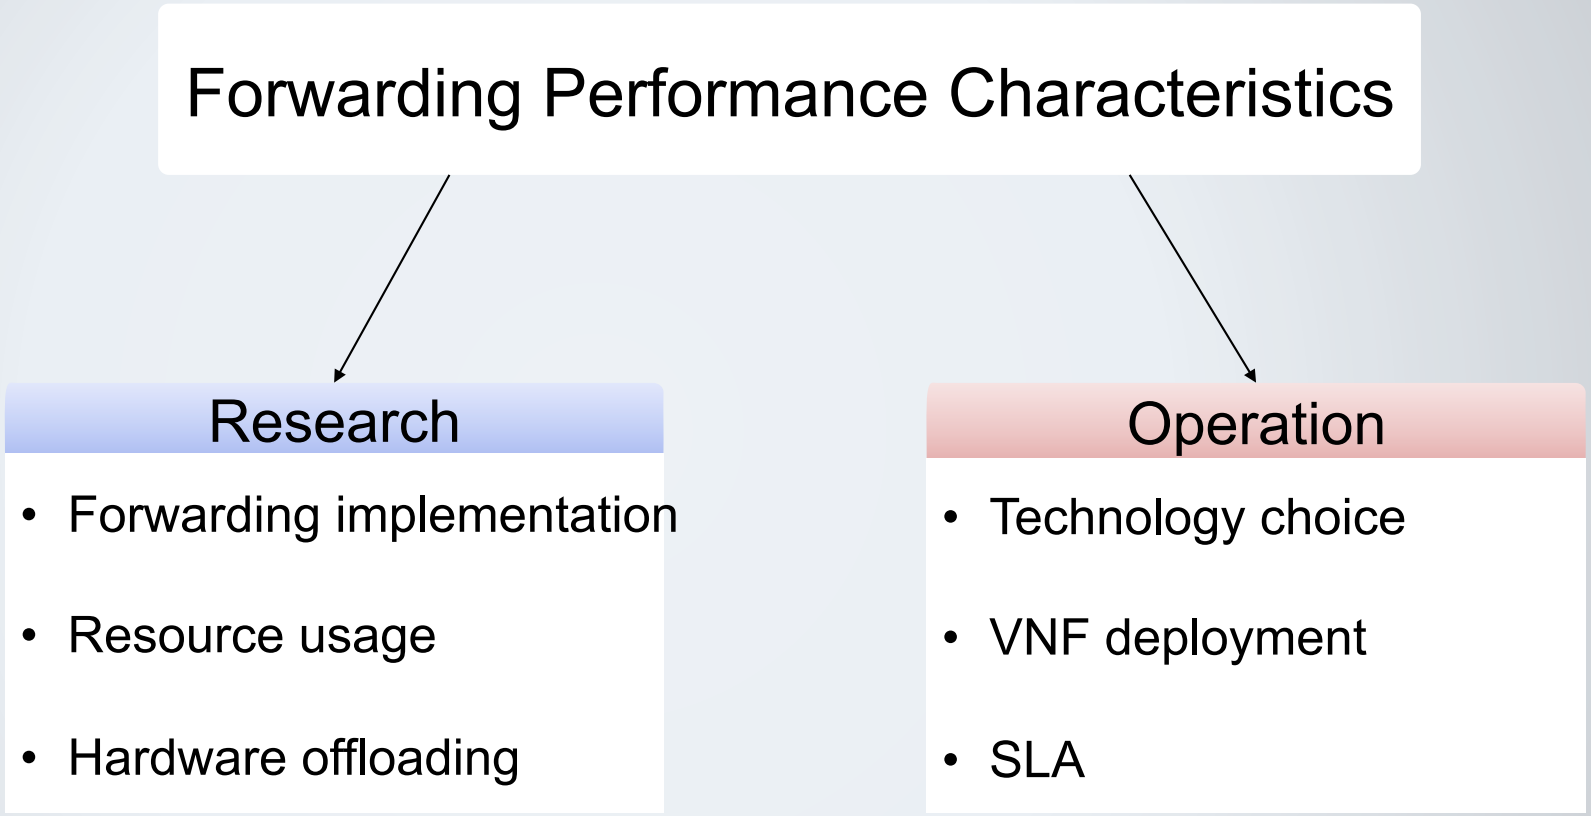
Forwarding Performance Characteristics
Research
• Forwarding implementation
• Resource usage
• Hardware offloading
Operation
• Technology choice
• VNF deployment
• SLA

Outlines

- **Backgrounds**
- **Evaluation**
  - Methods
  - Targets
  - Results
- **Guidelines**
  - Hardware
  - Software
- **Conclusion**

Network Softwarization

Everything is "*Software-based*"

Performance Concerns

Lower throughput

Higher latency

Instability

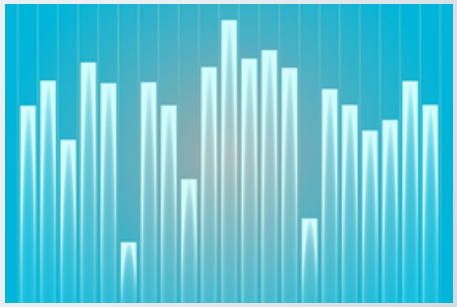

Promising Technologies

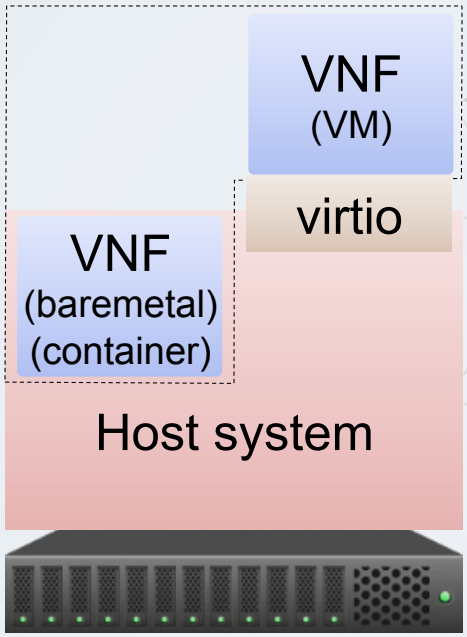

OvS Open vSwitch
BESS
snabb
VOSYSwitch
xDPd
VALE
lagopus

Forwarding

vhost-user
NetVM
SR-IOV
ptnetmap

Virtual I/O

DPDK DATA PLANE DEVELOPMENT KIT
Netmap
PacketShader
Open Onload

Packet I/O

Which technologies should we use?

Outlines

- **Backgrounds**
- **Evaluation**
  - Methods
  - Targets
  - Results
- **Guidelines**
  - Hardware
  - Software
- **Conclusion**

Methods

- VNF forms

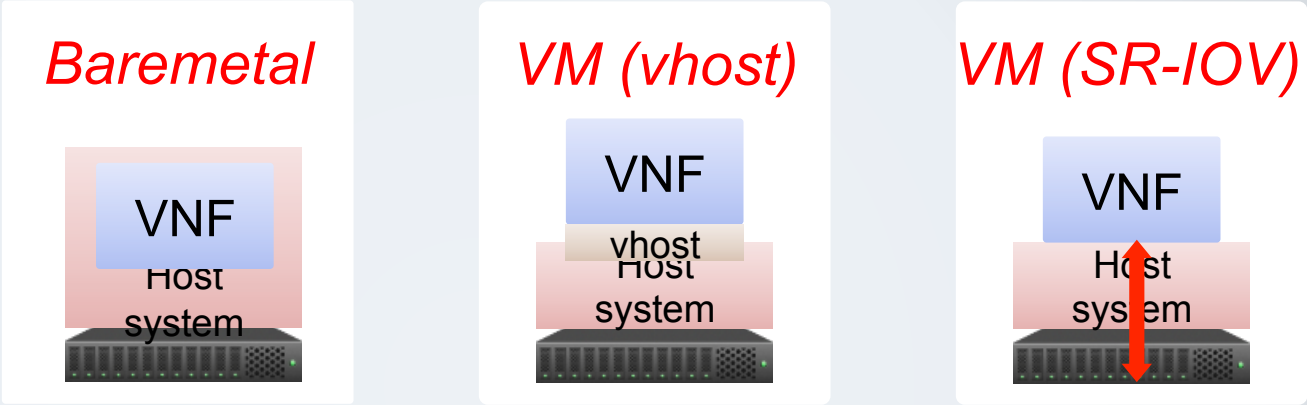

- Combinations

Packet I/O x *Virtual I/O* x *Forwarding*

Environment

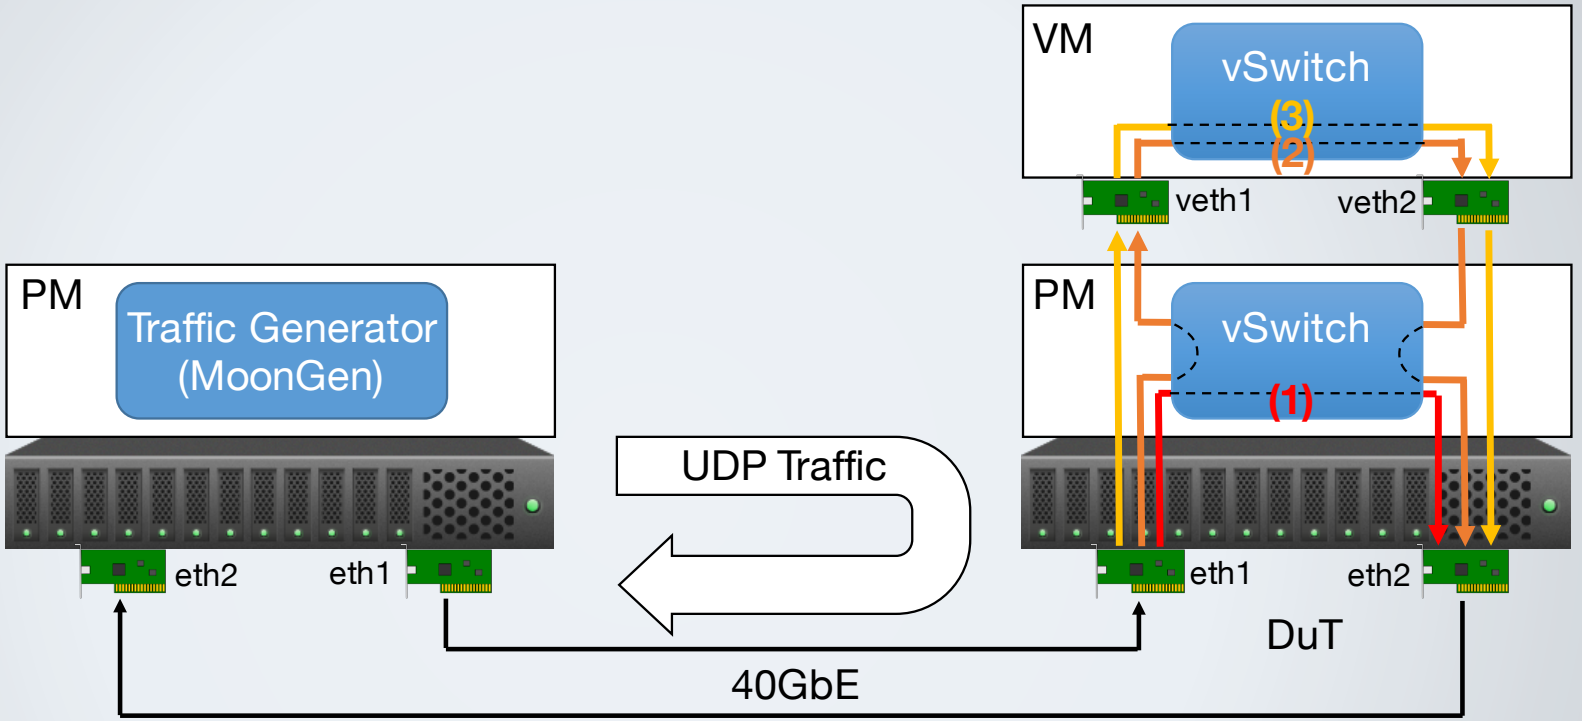

- Single UDP flow (unidirectional)

Machine specs

Physical	Machine A	Machine B
OS	CentOS 7.3 (3.10.0-514.6.1.x86_64)	
CPU	Core i7-6700K 4.00 GHz (4 cores with HT)	Xeon E5-2630 v4 2.20 GHz (2x10 cores w/o HT)
Memory	32 GB (DDR4-2133)	128 GB (DDR4-2133)
VMM	KVM	
NIC	Intel XL710 (i40e) PCI Express 3.0 (x8)	

Virtual	VM (on Machine A)	VM (on Machine B)
OS	CentOS 7.3 (3.10.0-514.6.1.x86_64)	
vCPU	3 cores	4 cores
Memory	4 GB	
vNIC	virtio-net (for virtio) / i40evf (for SR-IOV)	

Targets

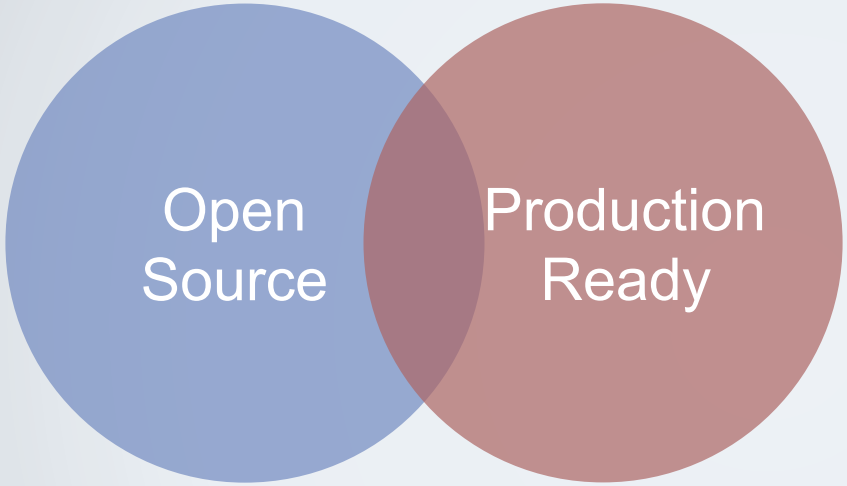

OvS Open vSwitch (Bridge)†

VALE (L2FWD)†

xDPd lagopus

Forwarding

vhost-user

SR-IOV

Virtual I/O

DPDK DATA PLANE DEVELOPMENT KIT

Netmap (NAPI)†

Packet I/O

† benchmarking

Throughput Results (DuT: Machine B)

(1) Baremetal

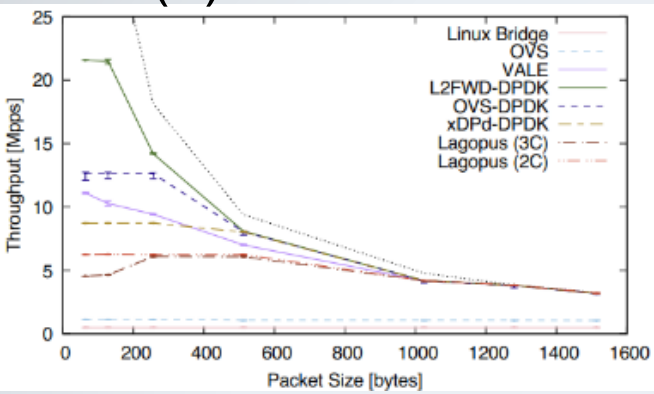

(2) VM (vhost-user)

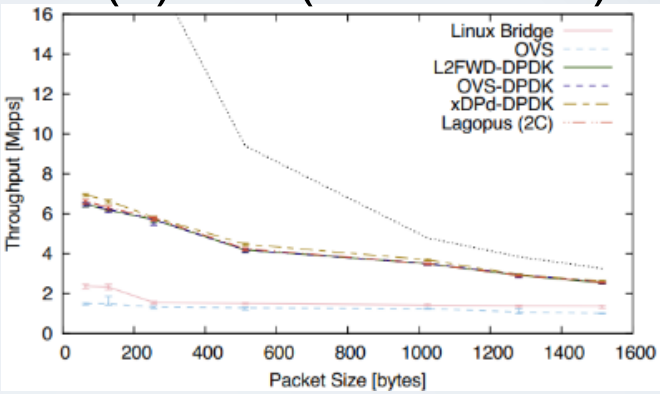

(3) VM (SR-IOV)

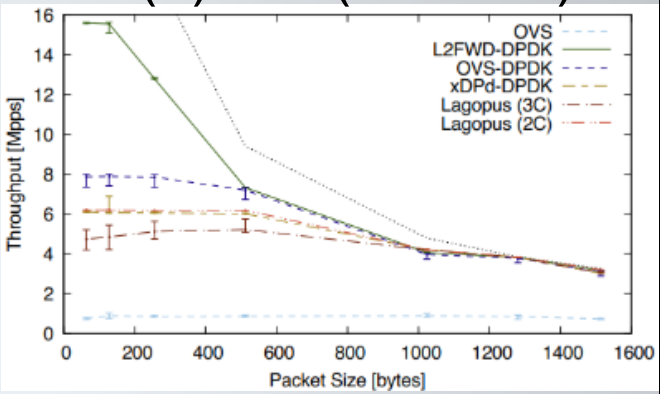

- Low throughput for short packets (<< 59.5 Mpps)
- Throughput variations in DPDK-based switches
- SR-IOV is faster but less stability

Latency/Jitter Results (DuT: Machine B)

(1) Baremetal

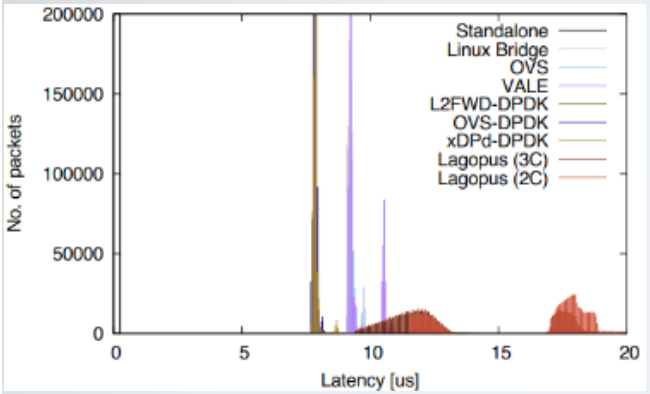

(2) VM (vhost-user)

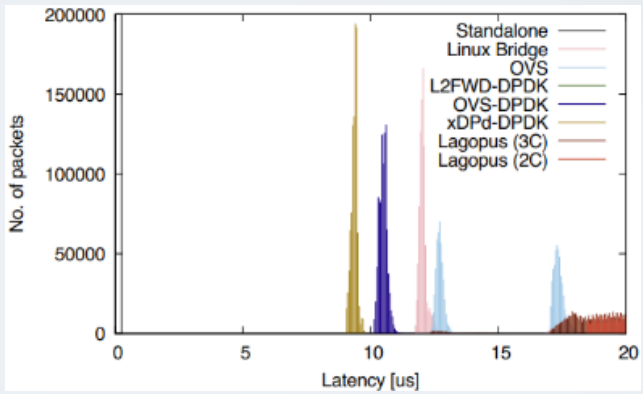

(3) VM (SR-IOV)

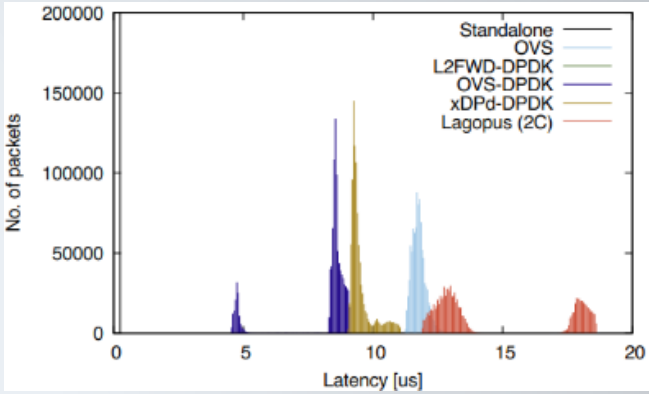

- Under 10 μs latency is possible with minute jitter
- Latency/Jitter variations in DPDK-based switches
- Lagopus and L2FWD showed worse latency/jitter

Effects of CPU Speed Differences

DuT: Machine A (Faster)

Throughput

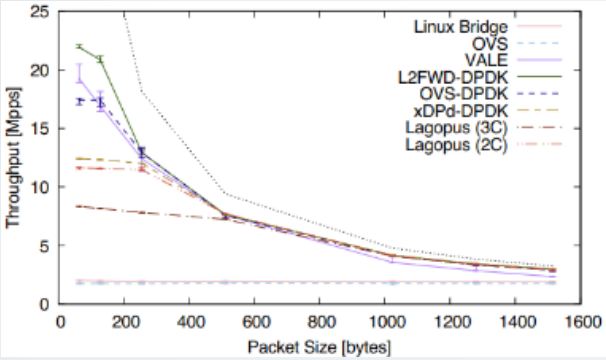

Latency/Jitter (64 bytes)

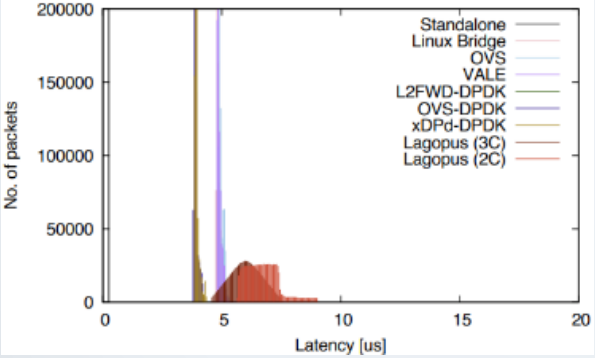

DuT: Machine B (Slower)

Throughput

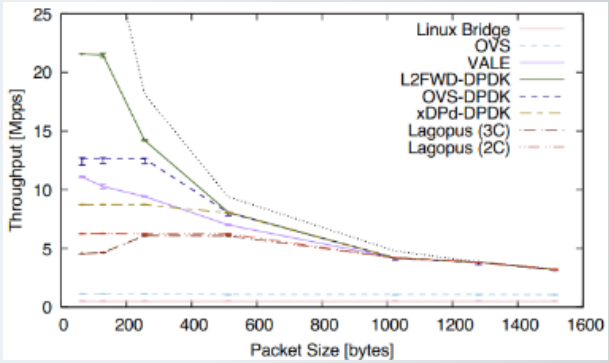

Latency/Jitter (64 bytes)

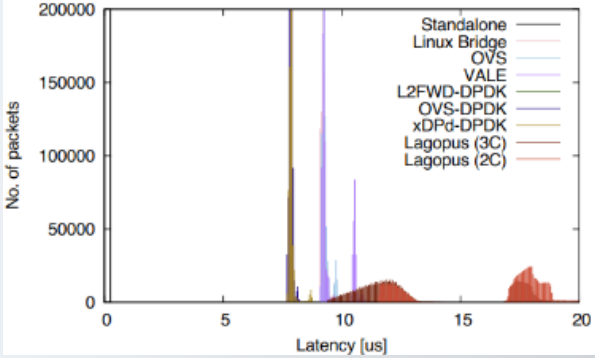

Outlines

- **Backgrounds**
- **Evaluation**
  - Methods
  - Targets
  - Results
- **Guidelines**
  - Hardware
  - Software
- **Conclusion**

Hardware

CPU speed is absolutely critical for network performance

Memory/PCIE† bus speed is not a bottleneck

† PCIE 3.0 (x16) is needed for dual-port 40G NIC (bidirectional)

SR-IOV is not preferable for production use

The performance is better, but …

Software

- How to use DPDK is important !
- OVS-DPDK is well-balanced for throughput and latency/jitter
- VM – Hypervisor comm. is the performance bottleneck

Outlines

- **Backgrounds**
- **Evaluation**
  - Methods
  - Targets
  - Results
- **Guidelines**
  - Hardware
  - Software
- **Conclusion**

Conclusion

Research

- Throughput for short packets
- Novel methods for VM – Hypervisor communications

Operation

- Using faster CPU
- OVS-DPDK/vhost-user is a reasonable approach
- Avoiding the use of SR-IOV