ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Performance of Network Virtualization in Cloud Computing Infrastructures: The OpenStack Case.

Franco Callegati, Walter Cerroni,
Chiara Contoli, Giuliano Santandrea

Dept. of Electrical, Electronic and Information Engineering
University of Bologna – Italy
chiara.contoli@unibo.it

Context

- Middlebox virtualization will enhance flexibility and configurability in the networks, with the help of Software defined networking (SDN) and Network function virtualization (NFV) paradigms
- Evolution of future internet architecture will take advantage of **cloud computing**
- The cloud computing **infrastructure** will determine the performance of the networking environment
- Do the virtualized network functions pose limitations?
- We want to:
    - investigate performance of a cloud virtual network infrastructure (VNI) and identify the critical bottlenecks
    - isolate the performance of the single components of the VNI
    - investigate performance of a multi-tenant NFV scenario

OpenStack

Open-source cloud management platform

- manages a cluster of hosting servers executing different Hypervisors (Vmware, KVM, Xen)

Cloud customers can instantiate computing, storage and network resources

- a virtual topology can be created, composed of virtual servers and virtual network appliances

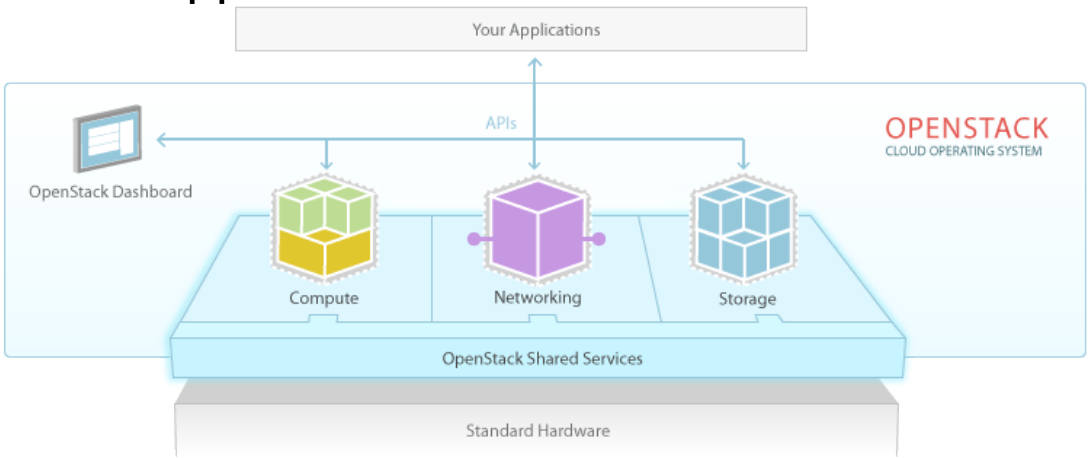

http://www.openstack.org

Our Test-bed

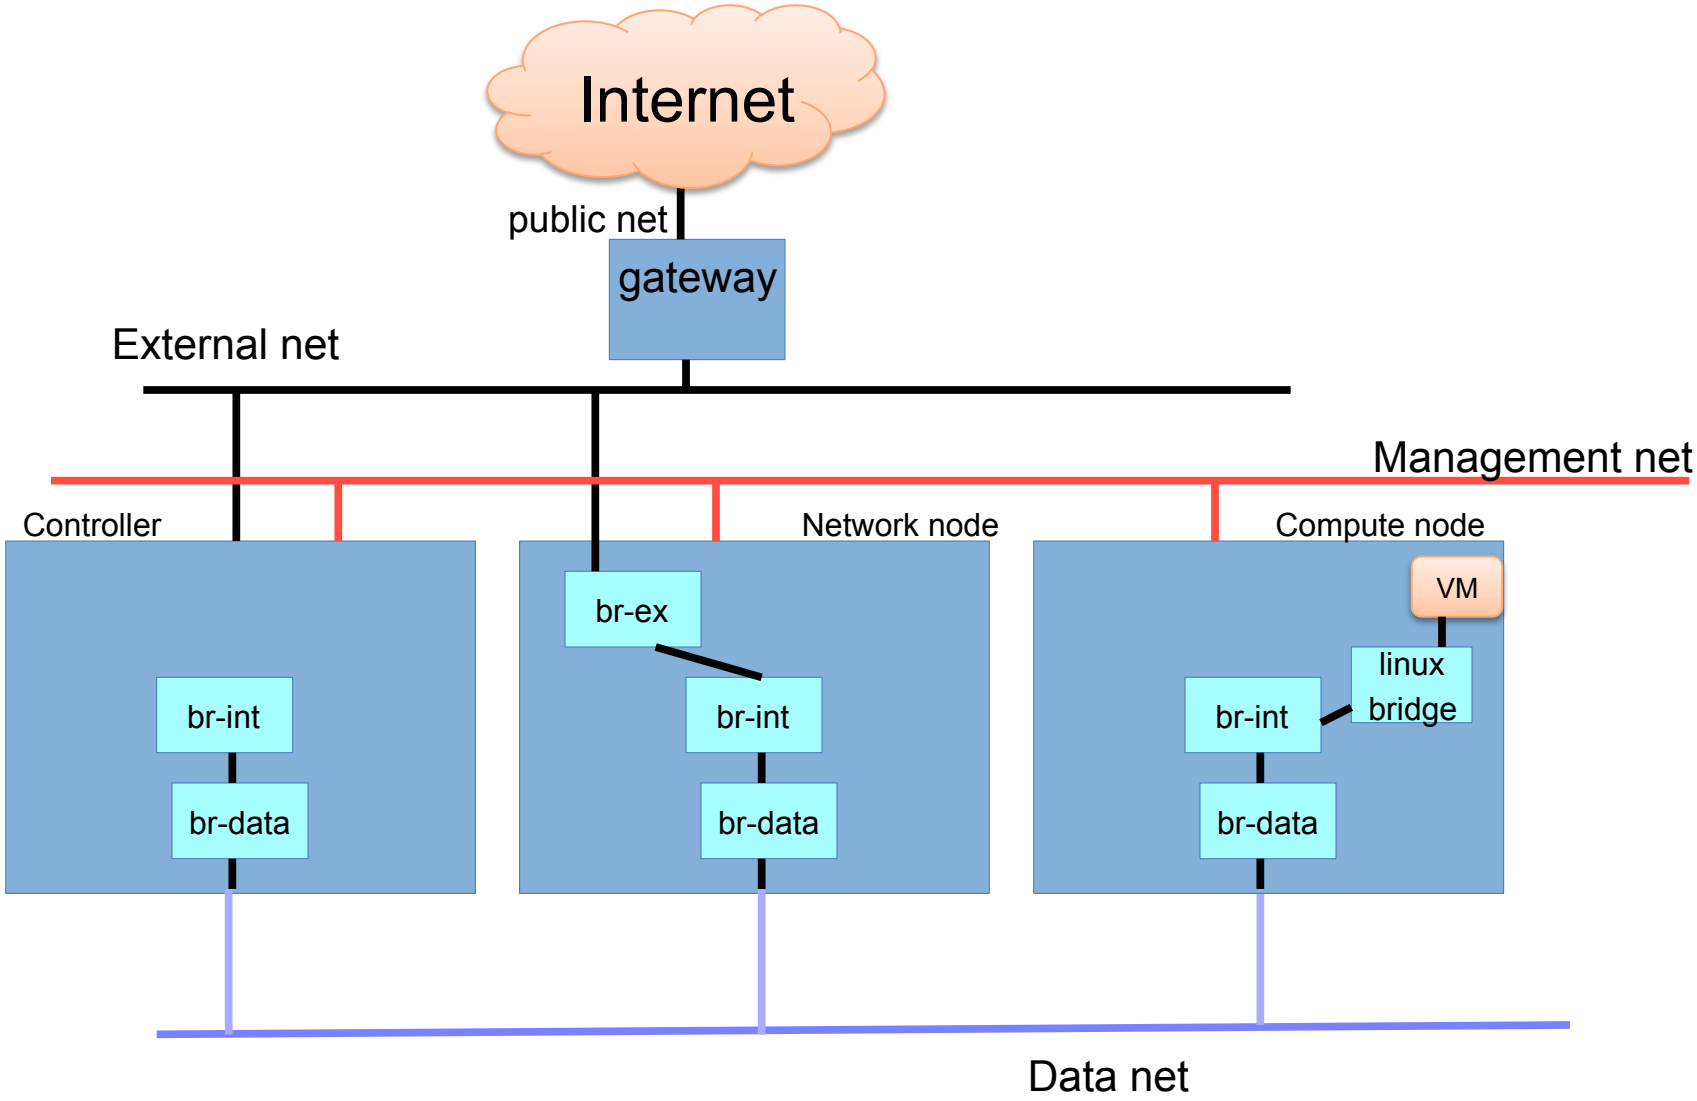
Internet
public net
gateway
External net
Management net
Controller
Network node
Compute node
br-ex
br-int
br-data
br-int
br-data
VM
linux
bridge
br-int
br-data
Data net

VNI: OpenStack Compute Node w/ 2 VMs

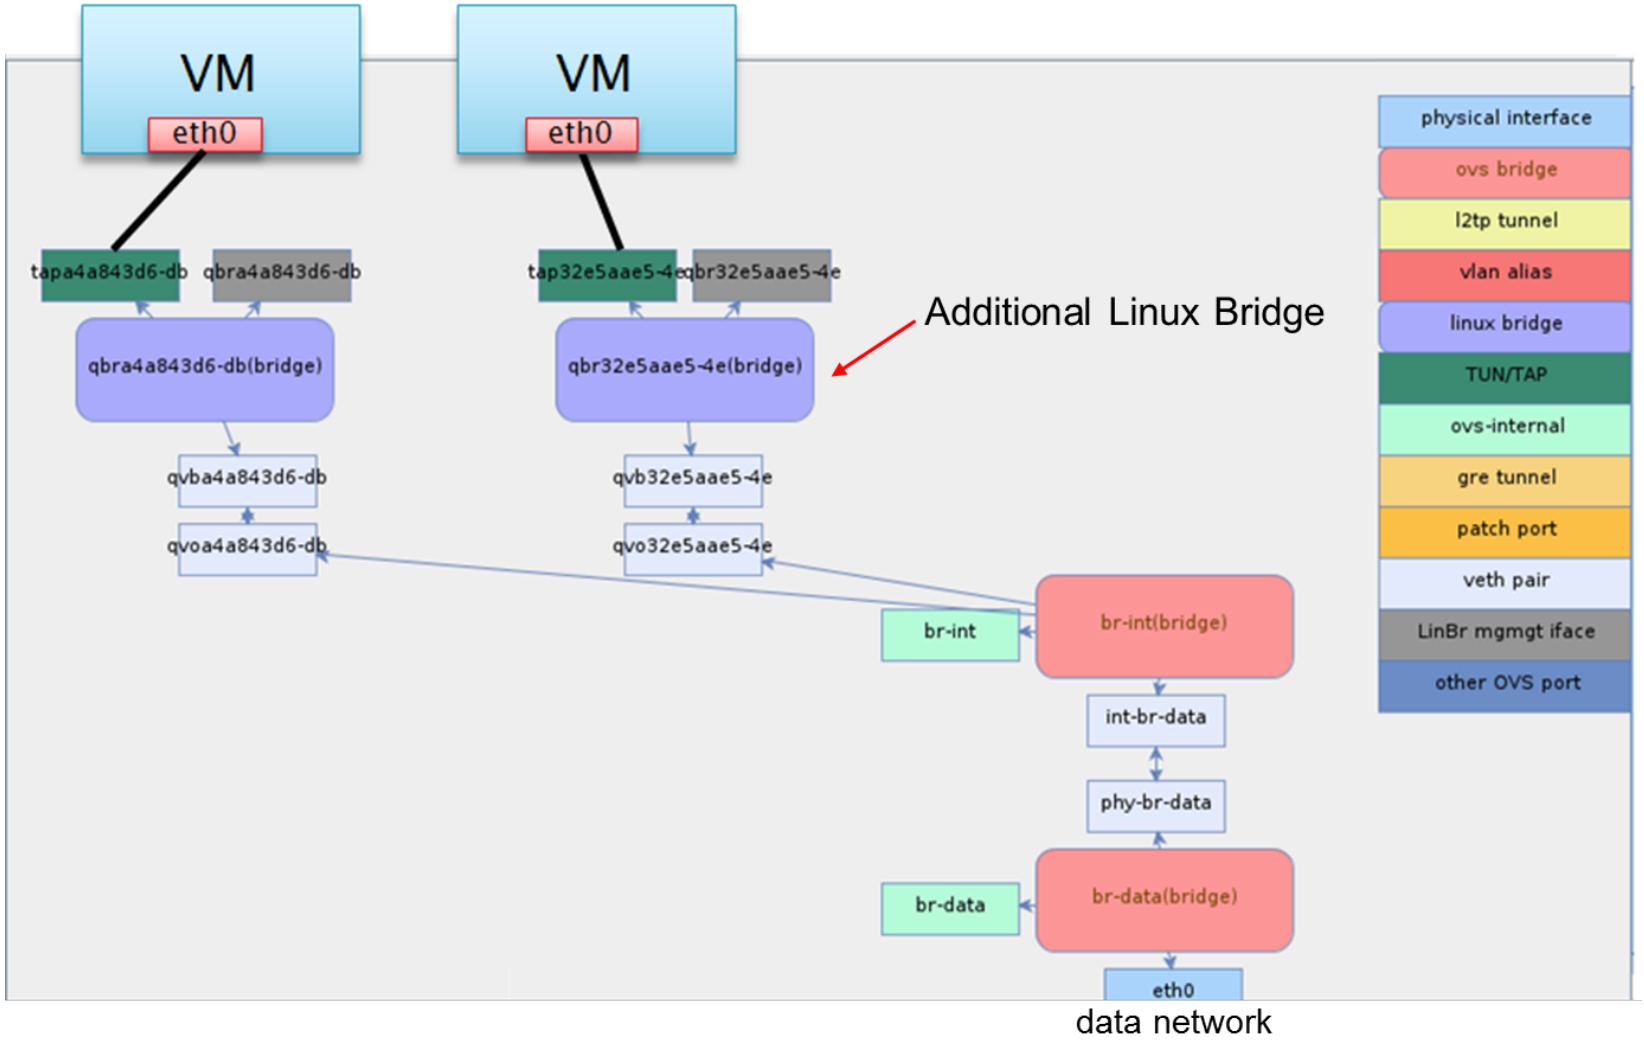

Screenshot taken from ShowMyNetworkState. Available: https://sites.google.com/site/showmynetworkstate

VNI: OpenStack Network Node

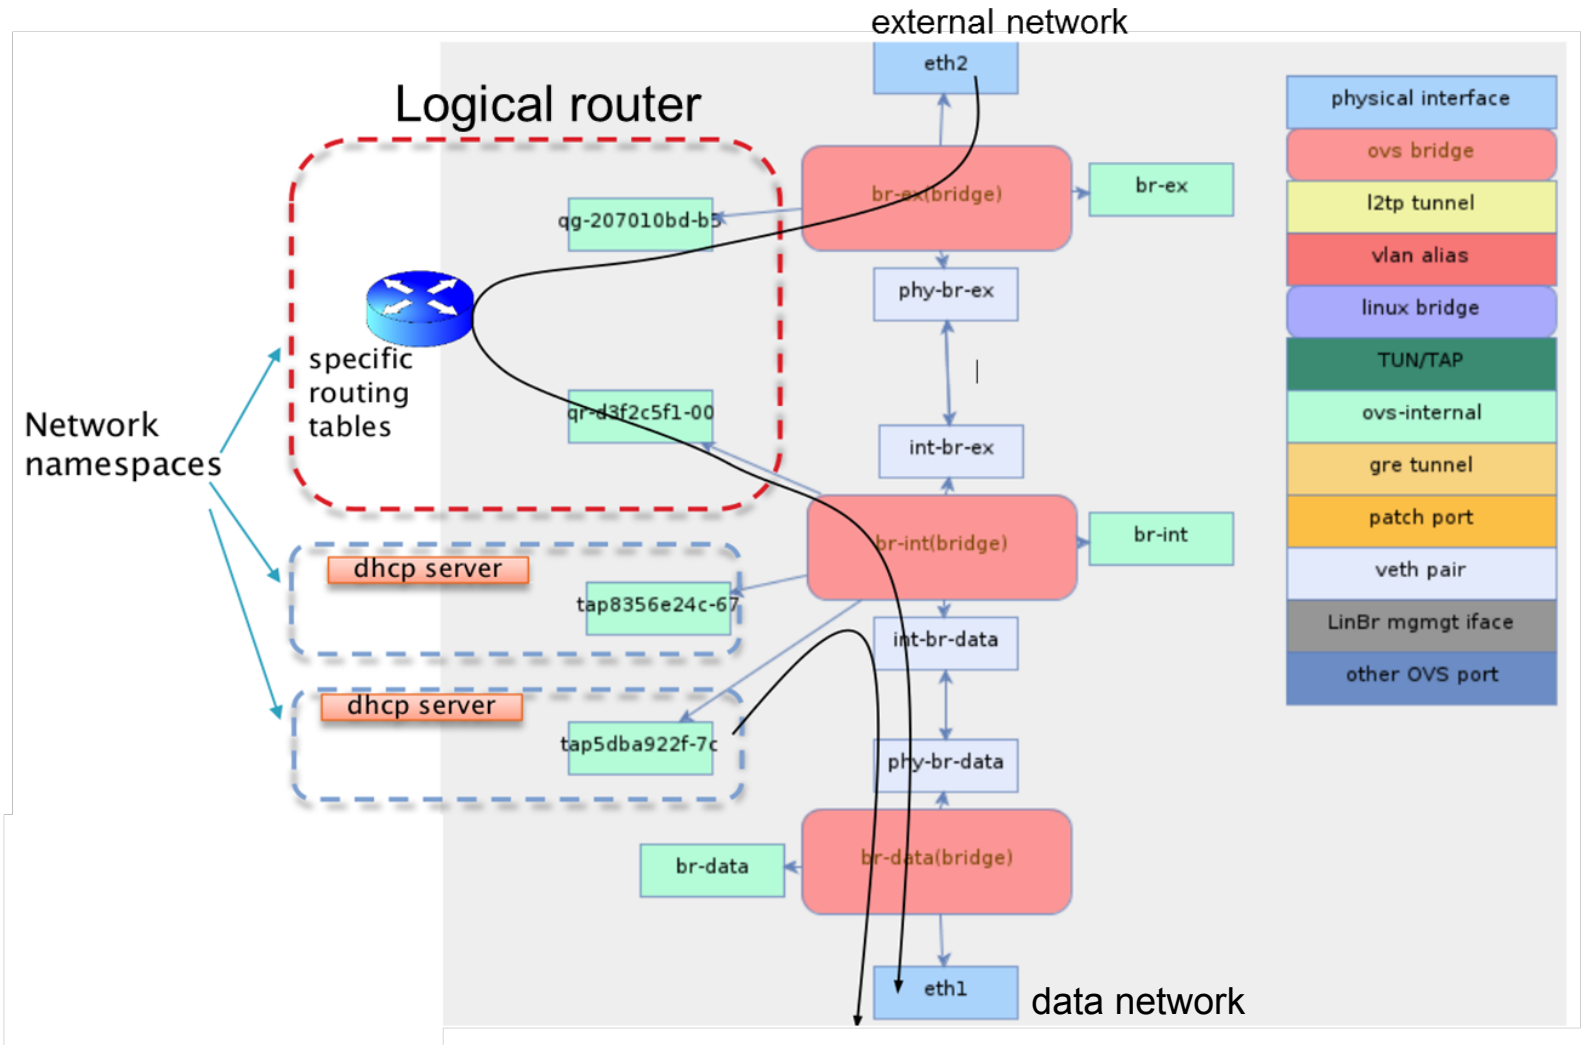

Screenshot taken from ShowMyNetworkState. Available: https://sites.google.com/site/showmynetworkstate

Experiments

- evaluate the network performance under critical traffic conditions and assessed
    - **maximum sustainable packet rate, maximum bandwidth**
- A traffic source (RUDE/CRUDE) sends a UDP flow to a destination:
    - packet rate = 1K-100K [pps], Ethernet payload = 64B / 1500B
- The destination measures the received packet-rate/throughput
- Because of R/C implementation, granularity is lower at higher packet rates
- Fixed rate background traffic: Iperf3
- Benchmark: back-to-back connection between physical hosts
- Hardware:
    - Tower PCs: HP Compaq 6200 Pro (2 CPU cores, 4 GB RAM)
    - OpenStack compute node (8 logical cores, 8GB RAM)
    - Gigabit Ethernet cards

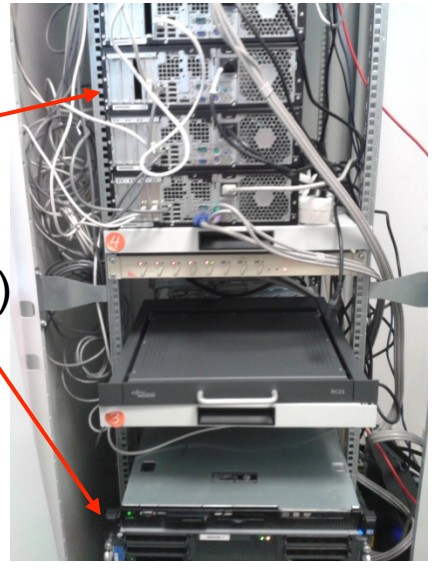

Scenarios

Benchmark

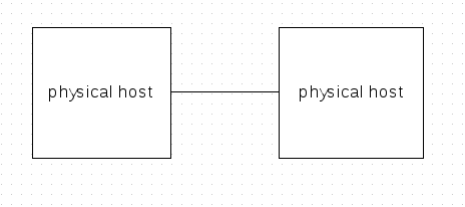

Openstack scenarios

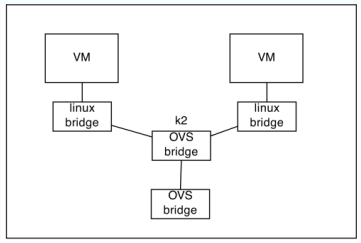

2 VMs in 1 OpenStack compute node

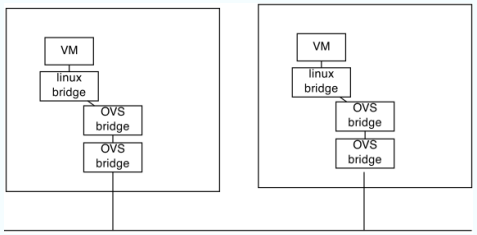

2 VMs in 2 OpenStack compute nodes

Scenarios

Non-Openstack scenarios

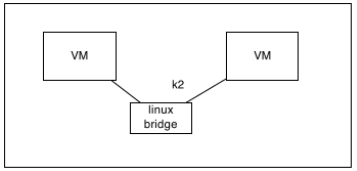

2 colocated VMs
Linux Bridge

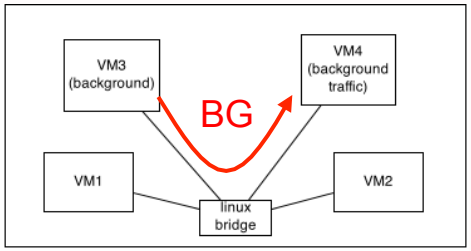

4 colocated VMs
background traffic
Linux bridge

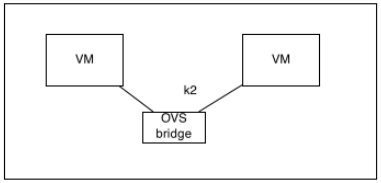

2 colocated VMs
Open vSwitch Bridge

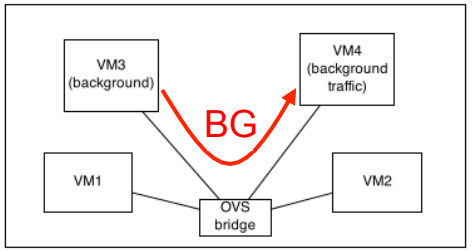

4 colocated VMs
background traffic
Open vSwitch Bridge

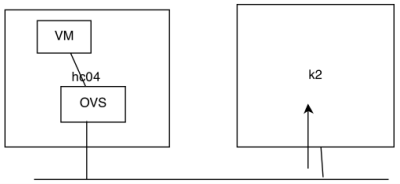

VM to physical host

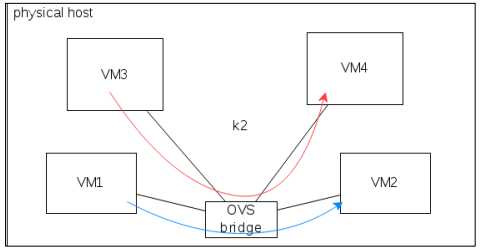

4 colocated VMs
2 flows,
Open vSwitch

Benchmark

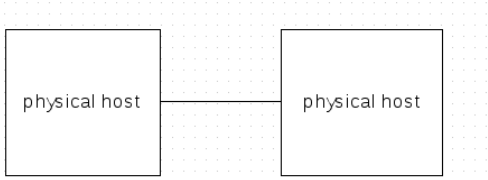

IP throughput =

1 Gbps * (1500/1538) = ~ 975 Mbps OK!

NIC physical bandwidth limit for 1500B packet size

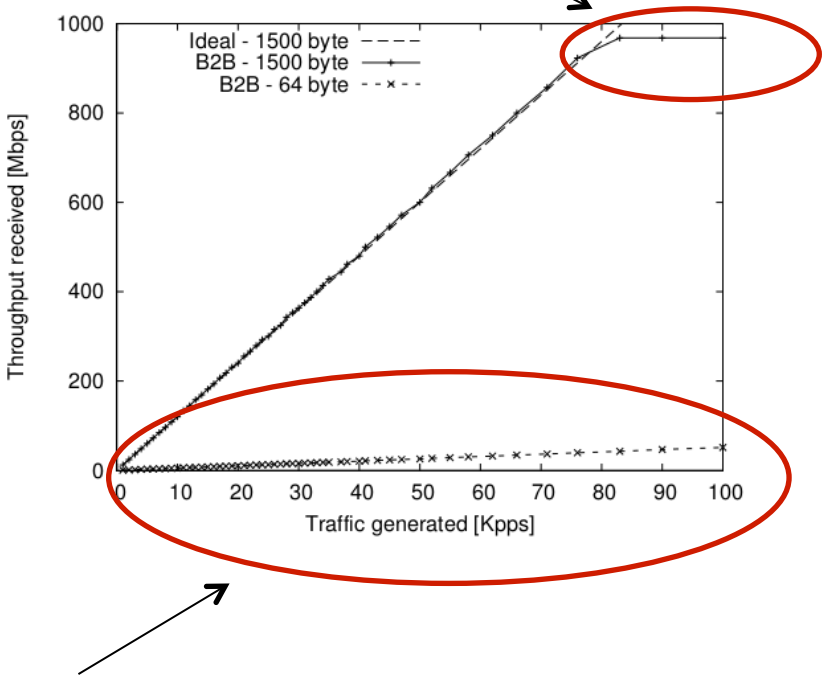

packet processing limit for 64B packet size

OpenStack scenario

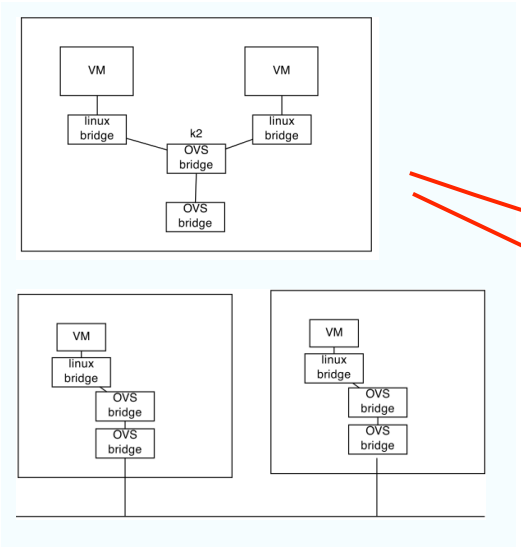

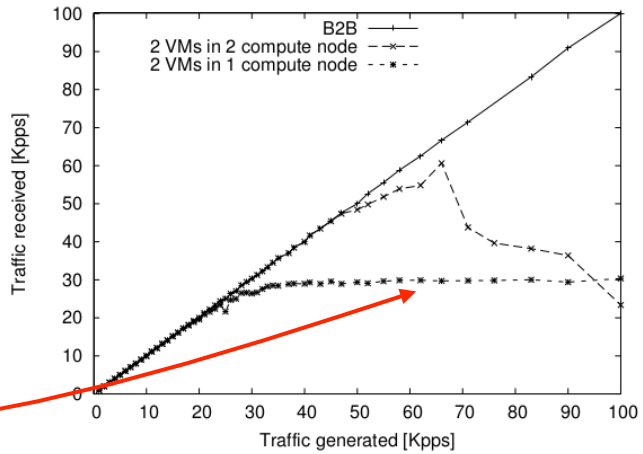

64B

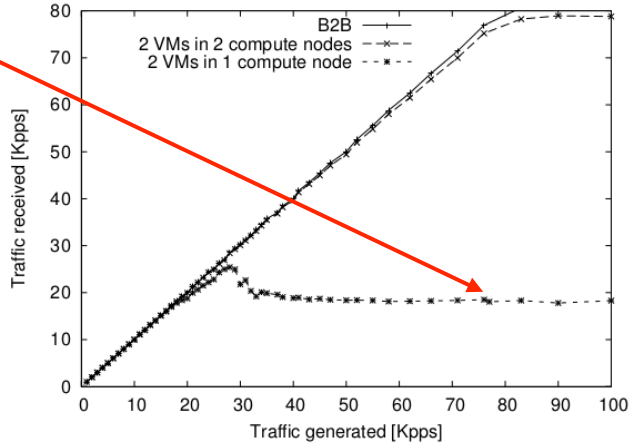

1500B

resorce contention causes the performance bottleneck

OpenStack scenario

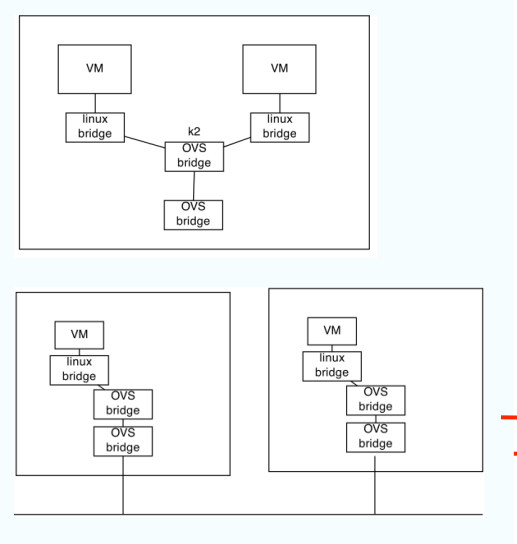

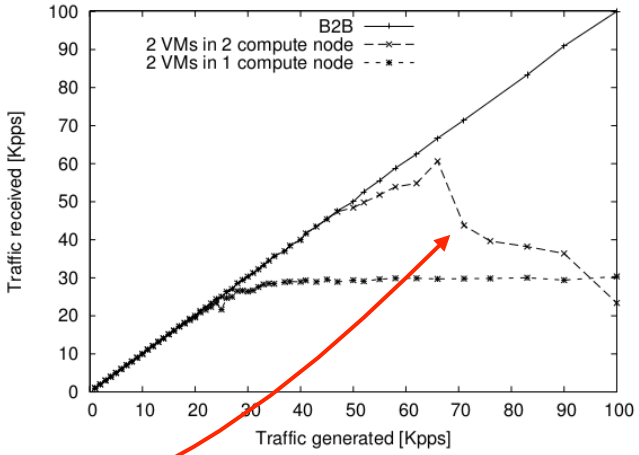

64B

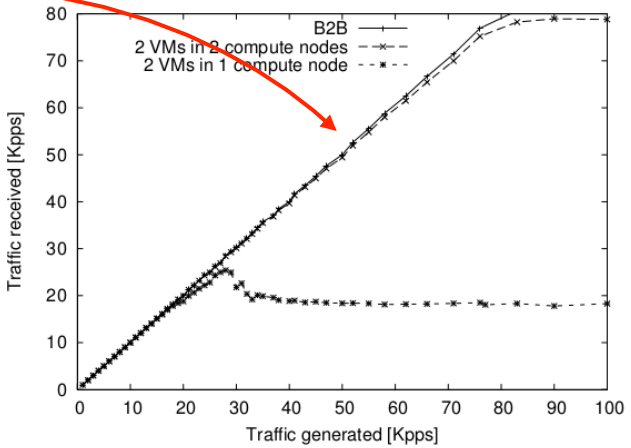

1500B

Workload: shared between the hosts

computational overhead for small packets

Linux bridge

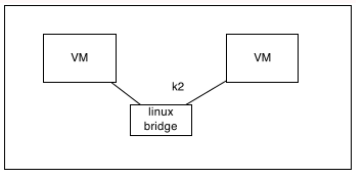

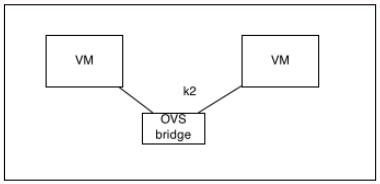

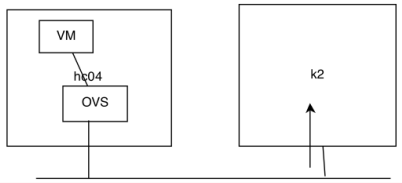

Linux Bridge:

performance degradation!

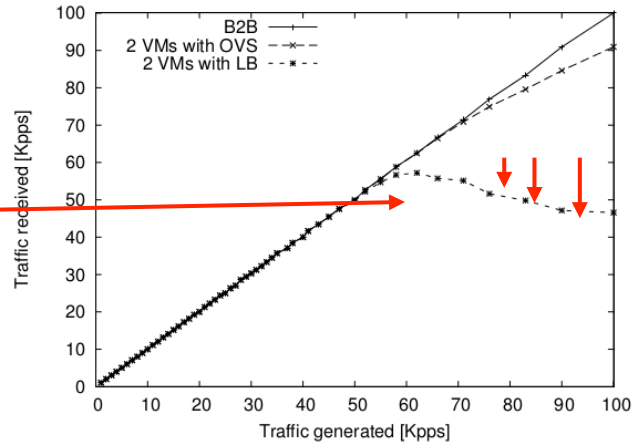

64B

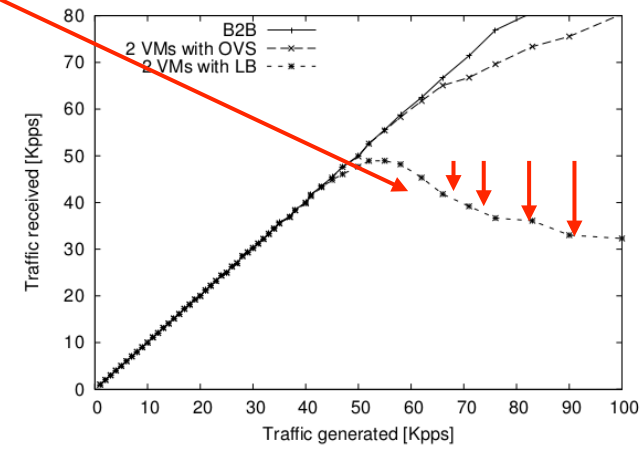

1500B

Open vSwitch

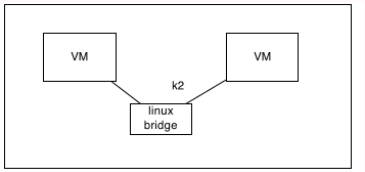

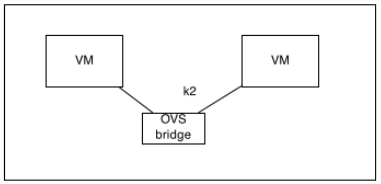

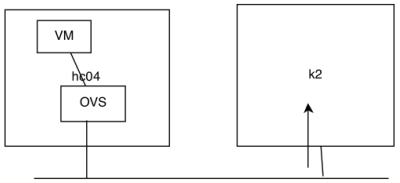

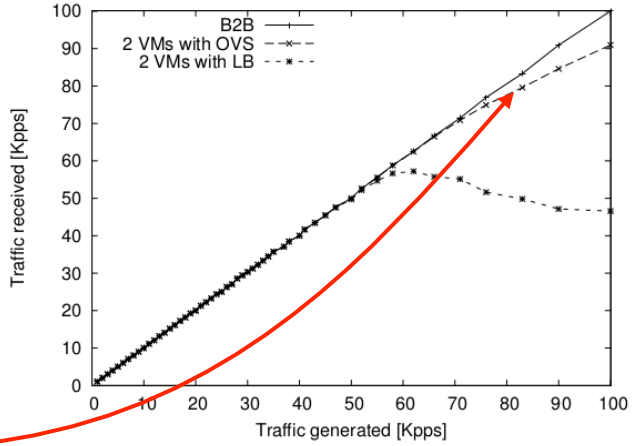

64B

Open vSwitch

no performance degradation

deviates later from the benchamrk

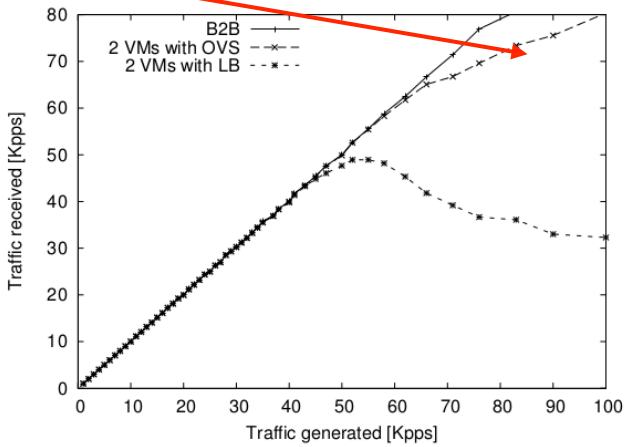

1500B

Physical receiver

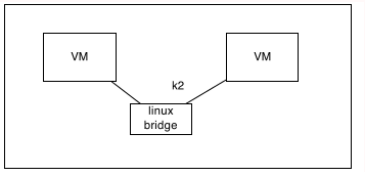

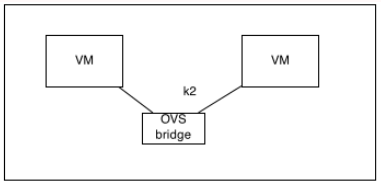

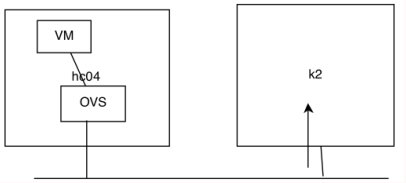

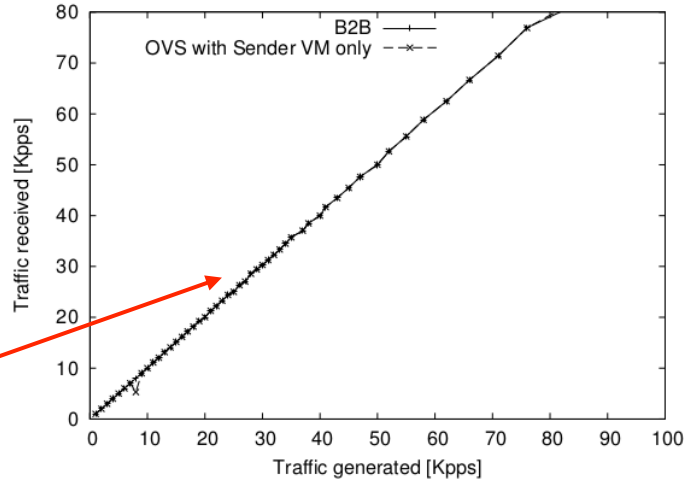

1500B

As good as the benchmark

Linux Bridge with BG traffic

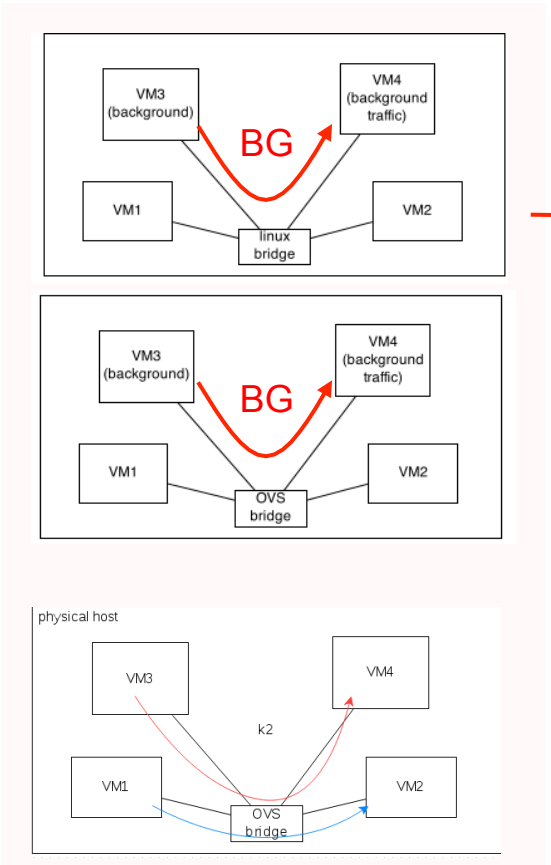

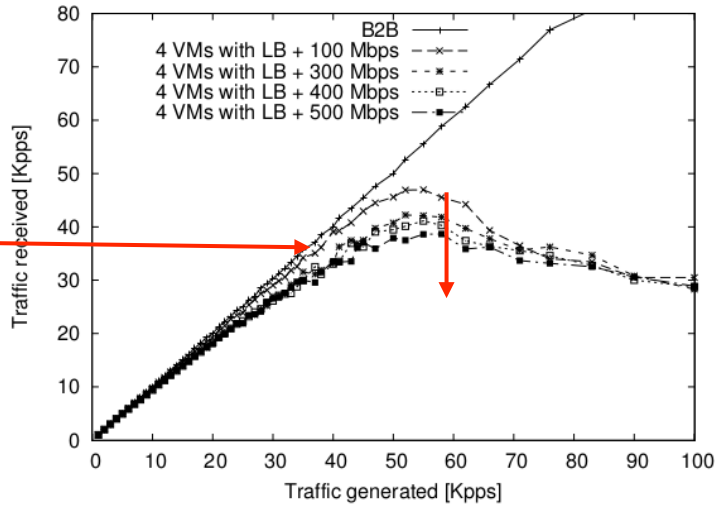

1500B

Additive effect of the background traffic

Background traffic comparison

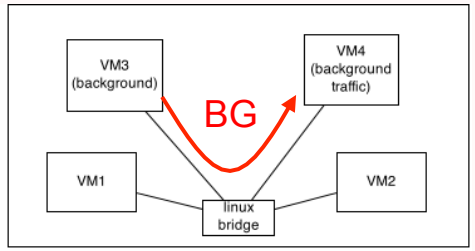

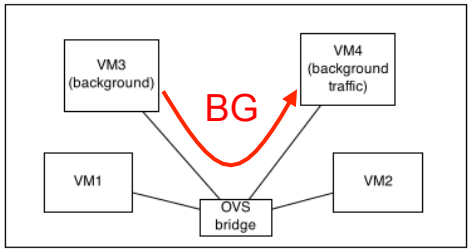

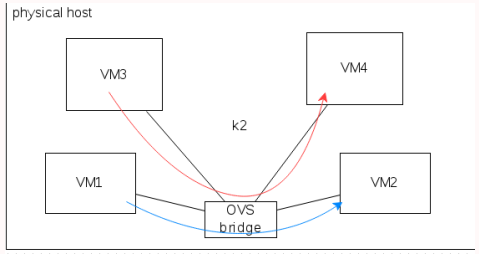

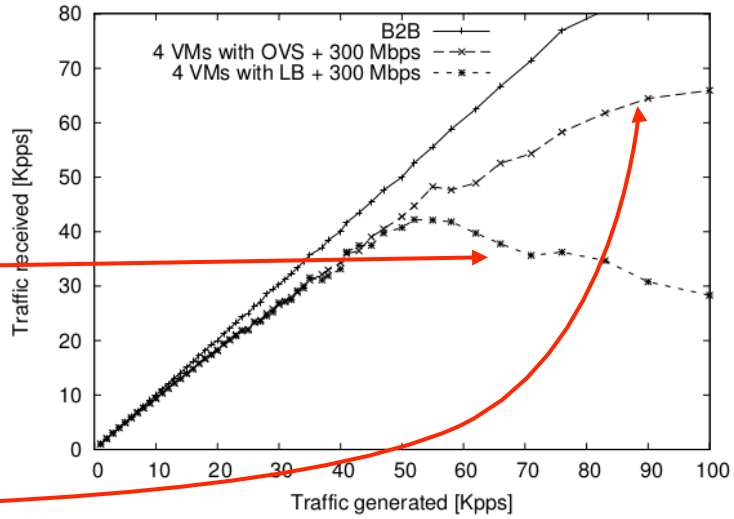

1500B

Comparison:

- effect of the additive traffic on both bridges
- OVS performs better than LB

Two RUDE/CRUDE flows

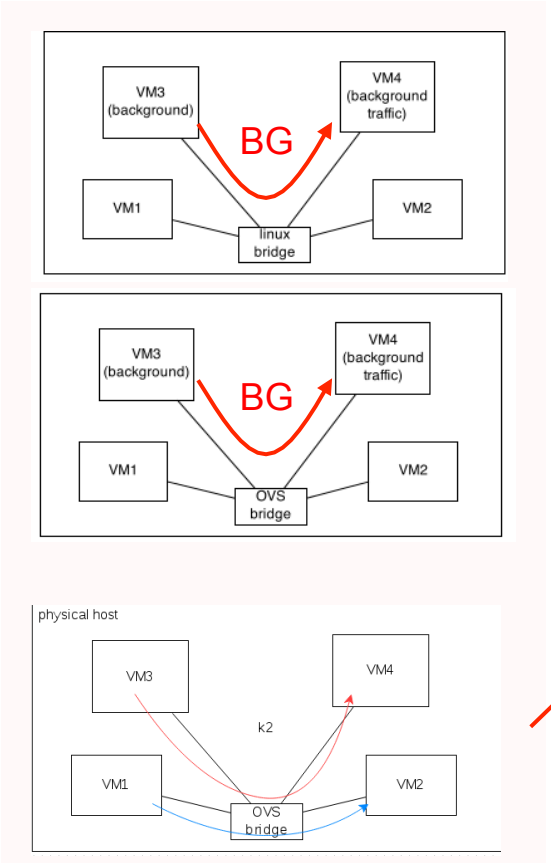

2 flows

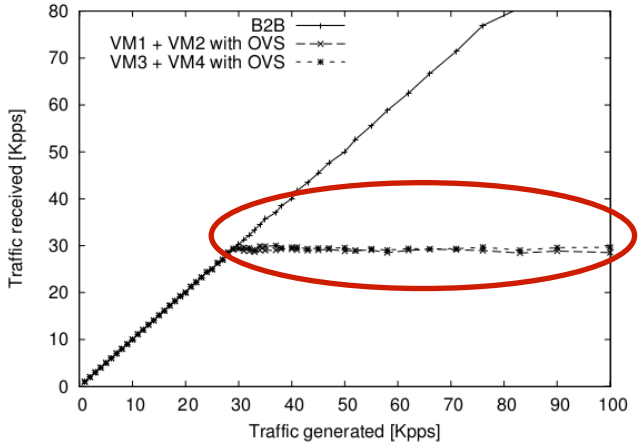

the bandwidths overlap at the saturation limit

30 Kpps + 30 Kpps < 80 Kpps (single-flow case)

Effect of the resource contention

Multi-tenant test scenario: NFV case study

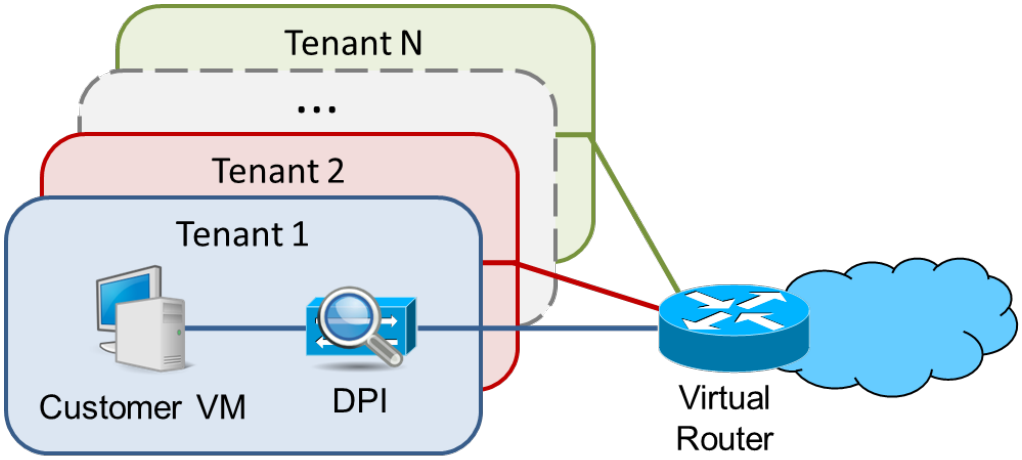

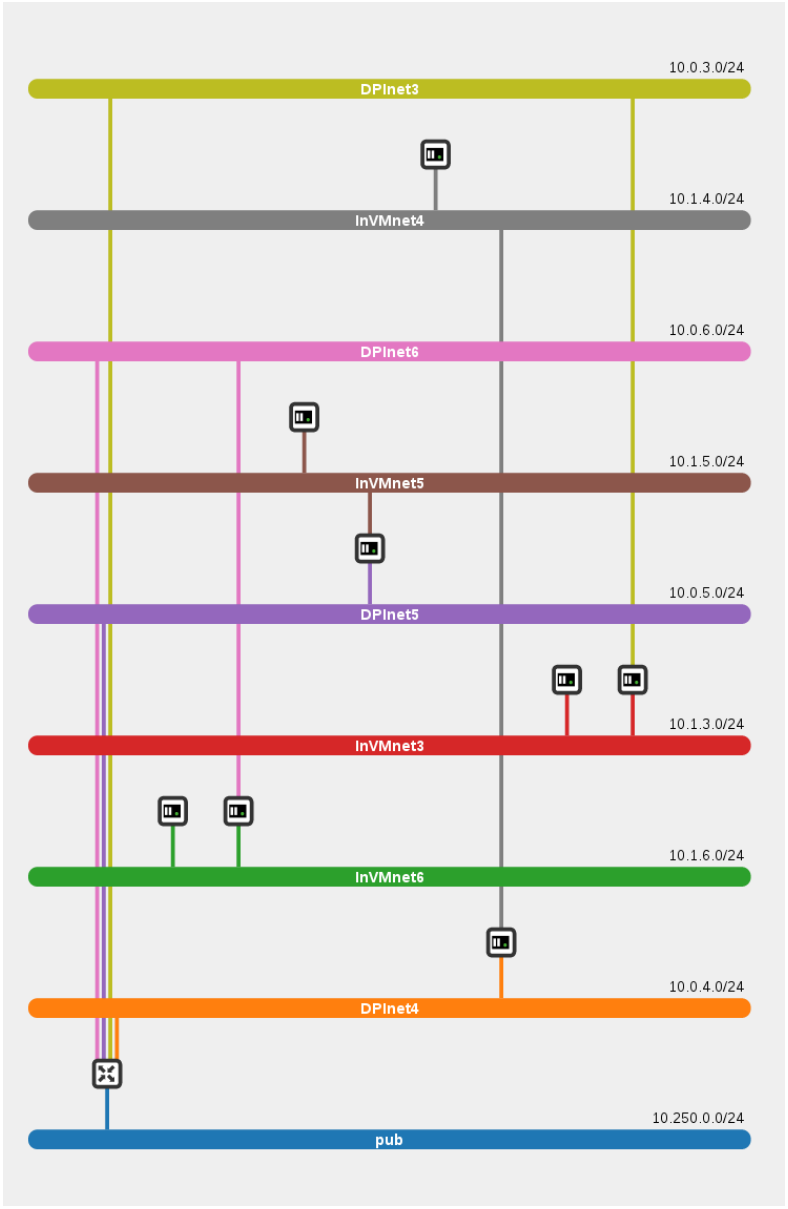

OpenStack dashboard view when N = 4

Performance evaluation

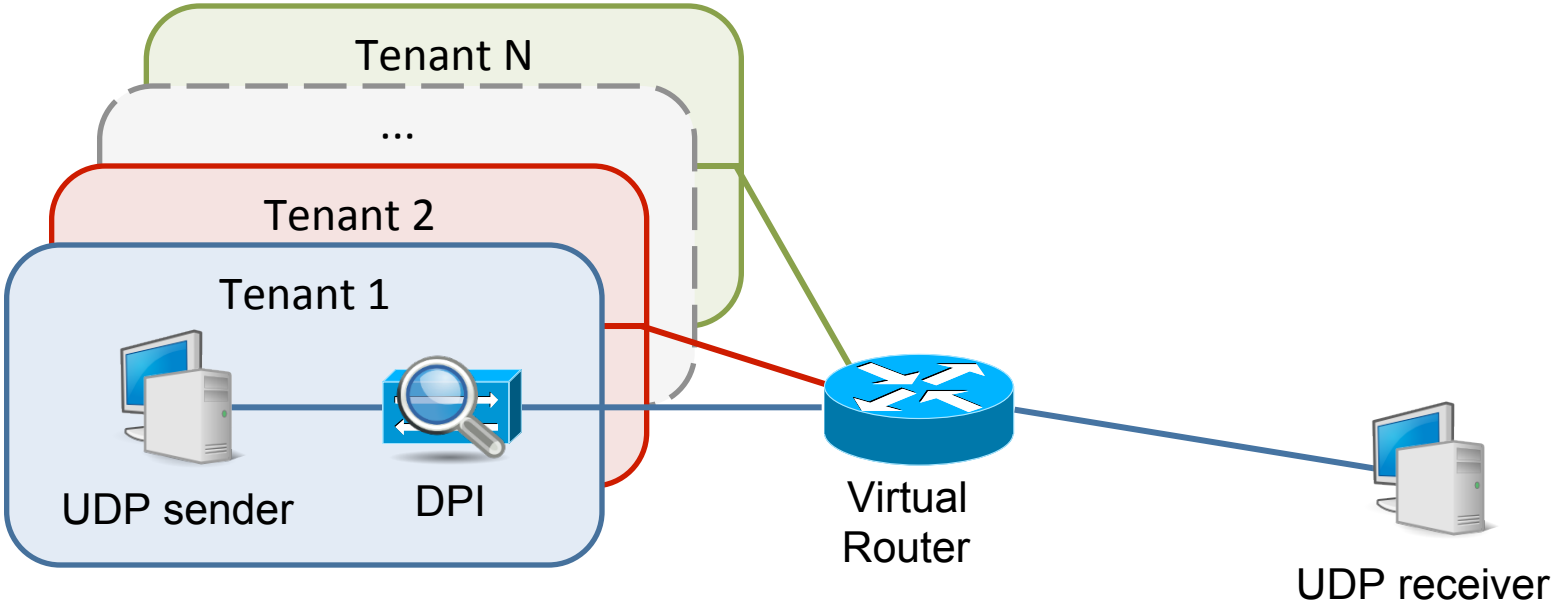

- Different numbers of tenants simultaneously active on the same compute node
- Each sender generates UDP traffic ranging from 10^3 to 10^5 packets per second (pps), for both 64 and 1500-byte IP packet sizes
- RUDE & CRUDE tool is used for traffic generation and detection
- All physical interfaces are Gigabit Ethernet network cards

Packet rate – 1500 byte packets

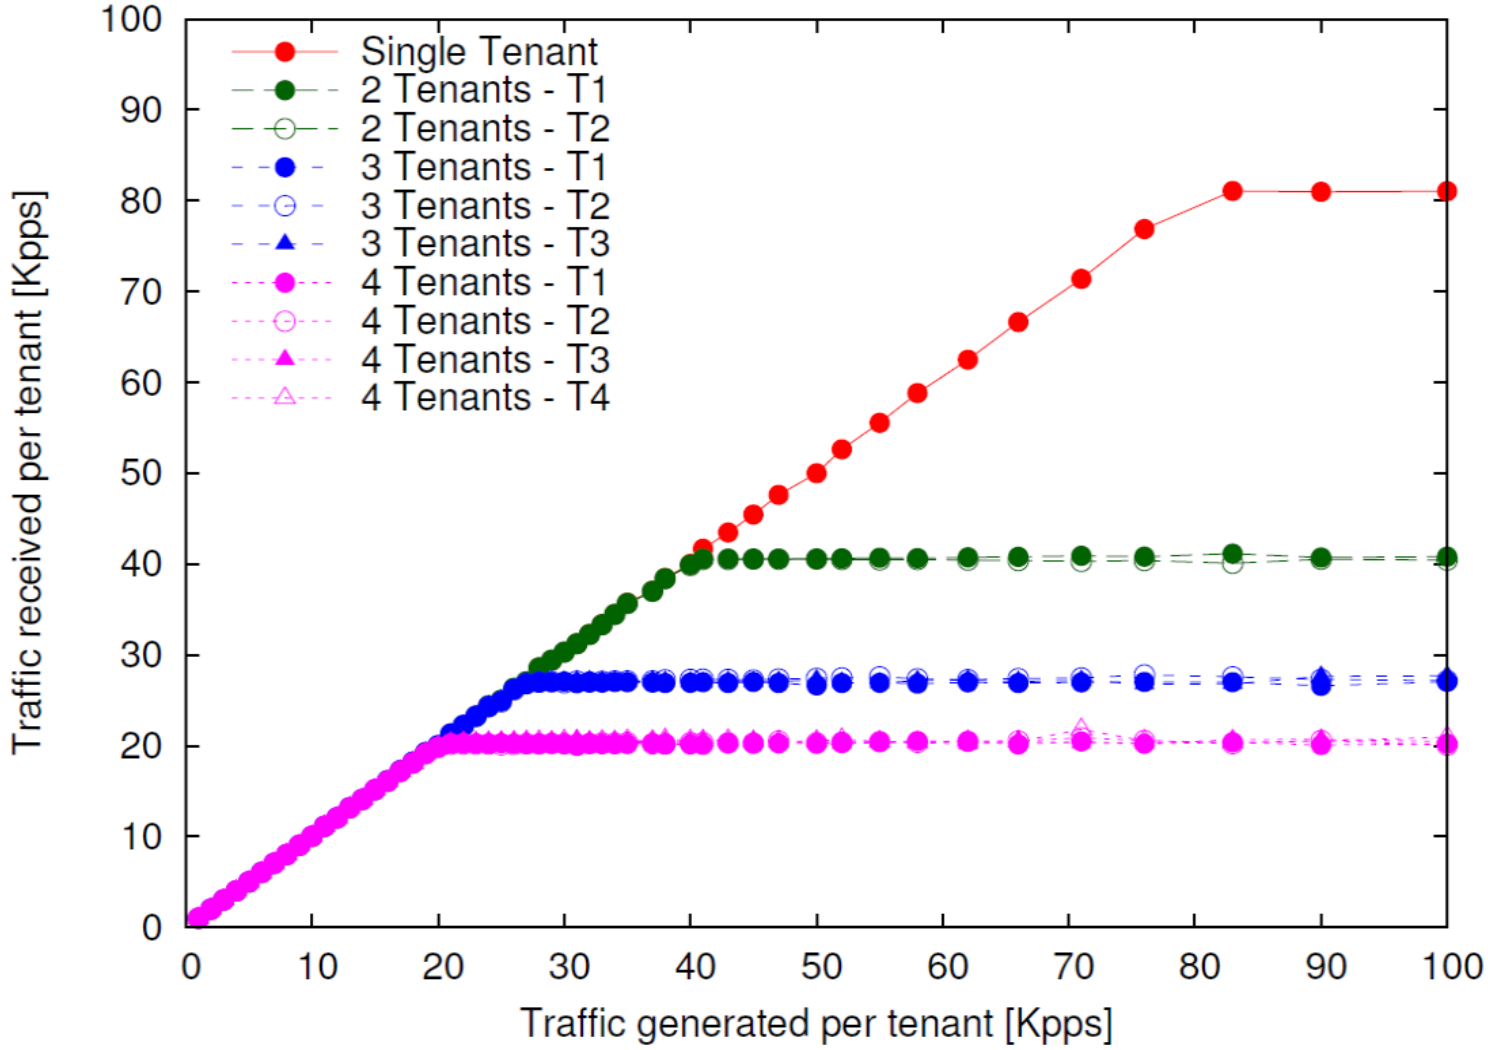

Packet rate – 64 byte packets

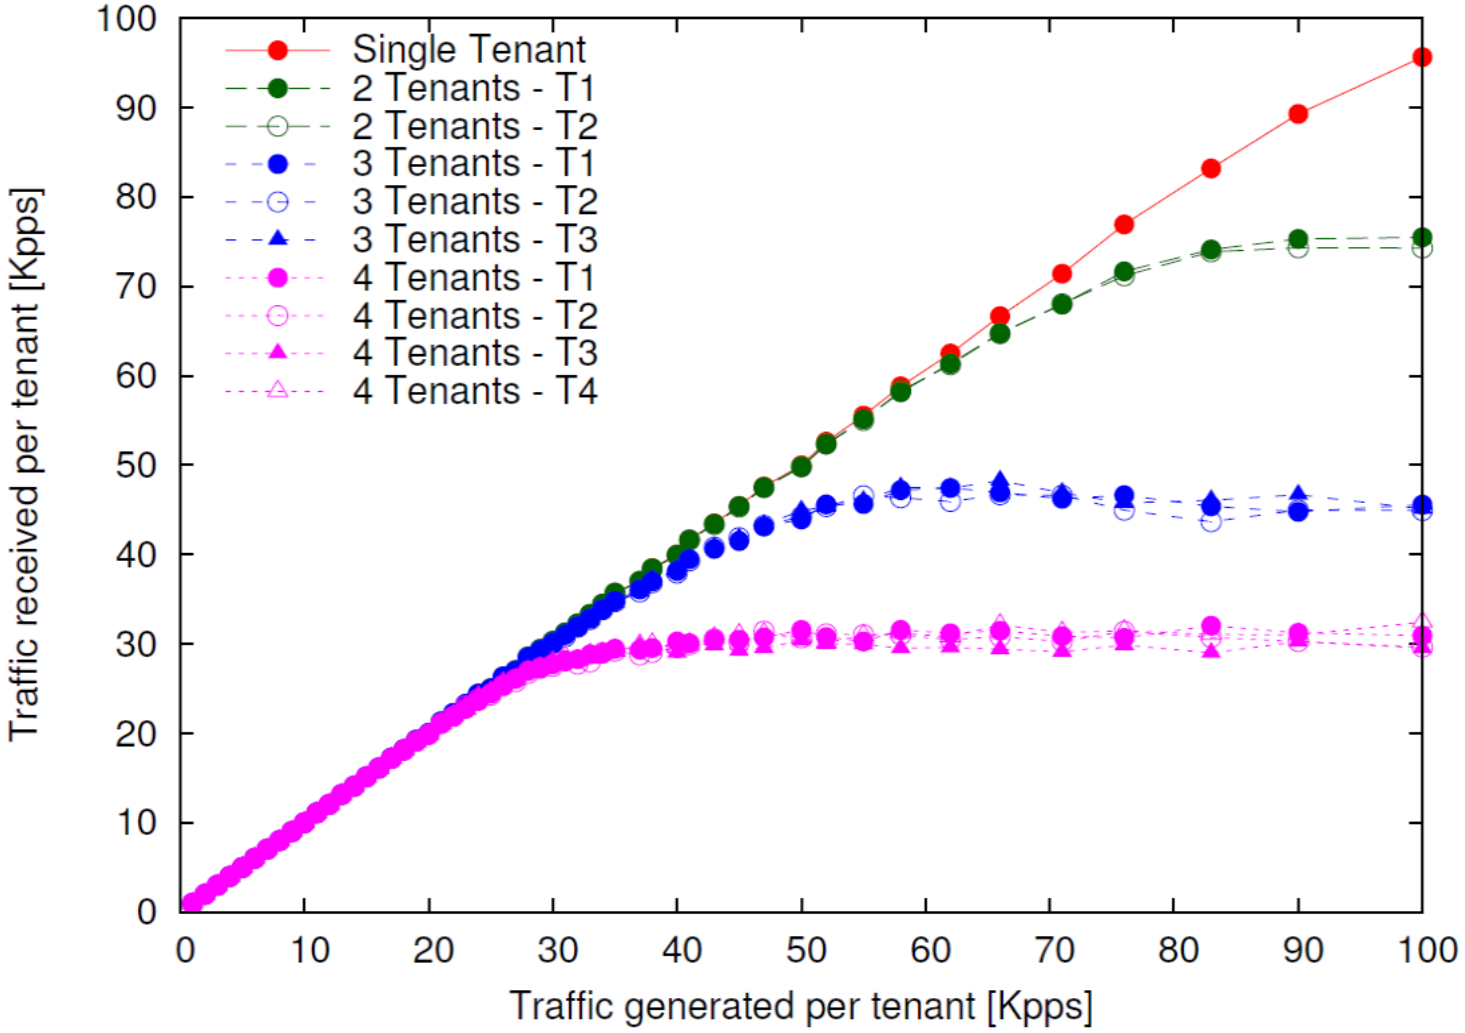

Linux Bridge (LB) bypass on Compute Node

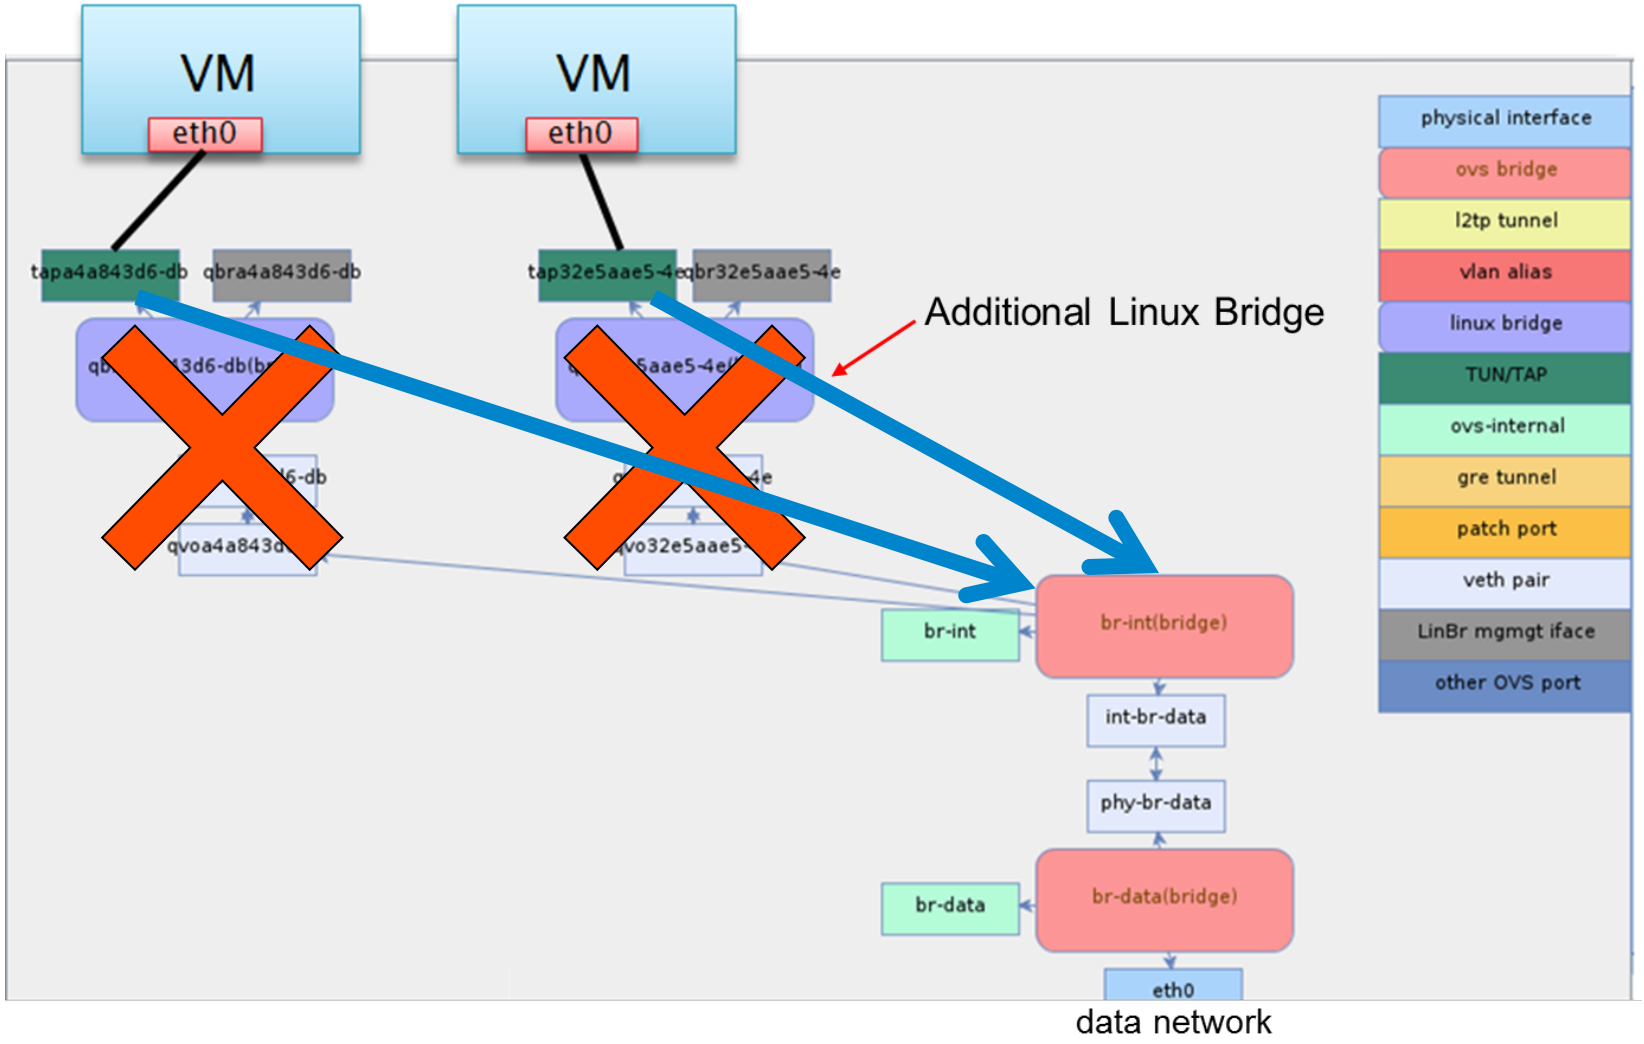

Throughput – 64 byte packets w/ LB bypass

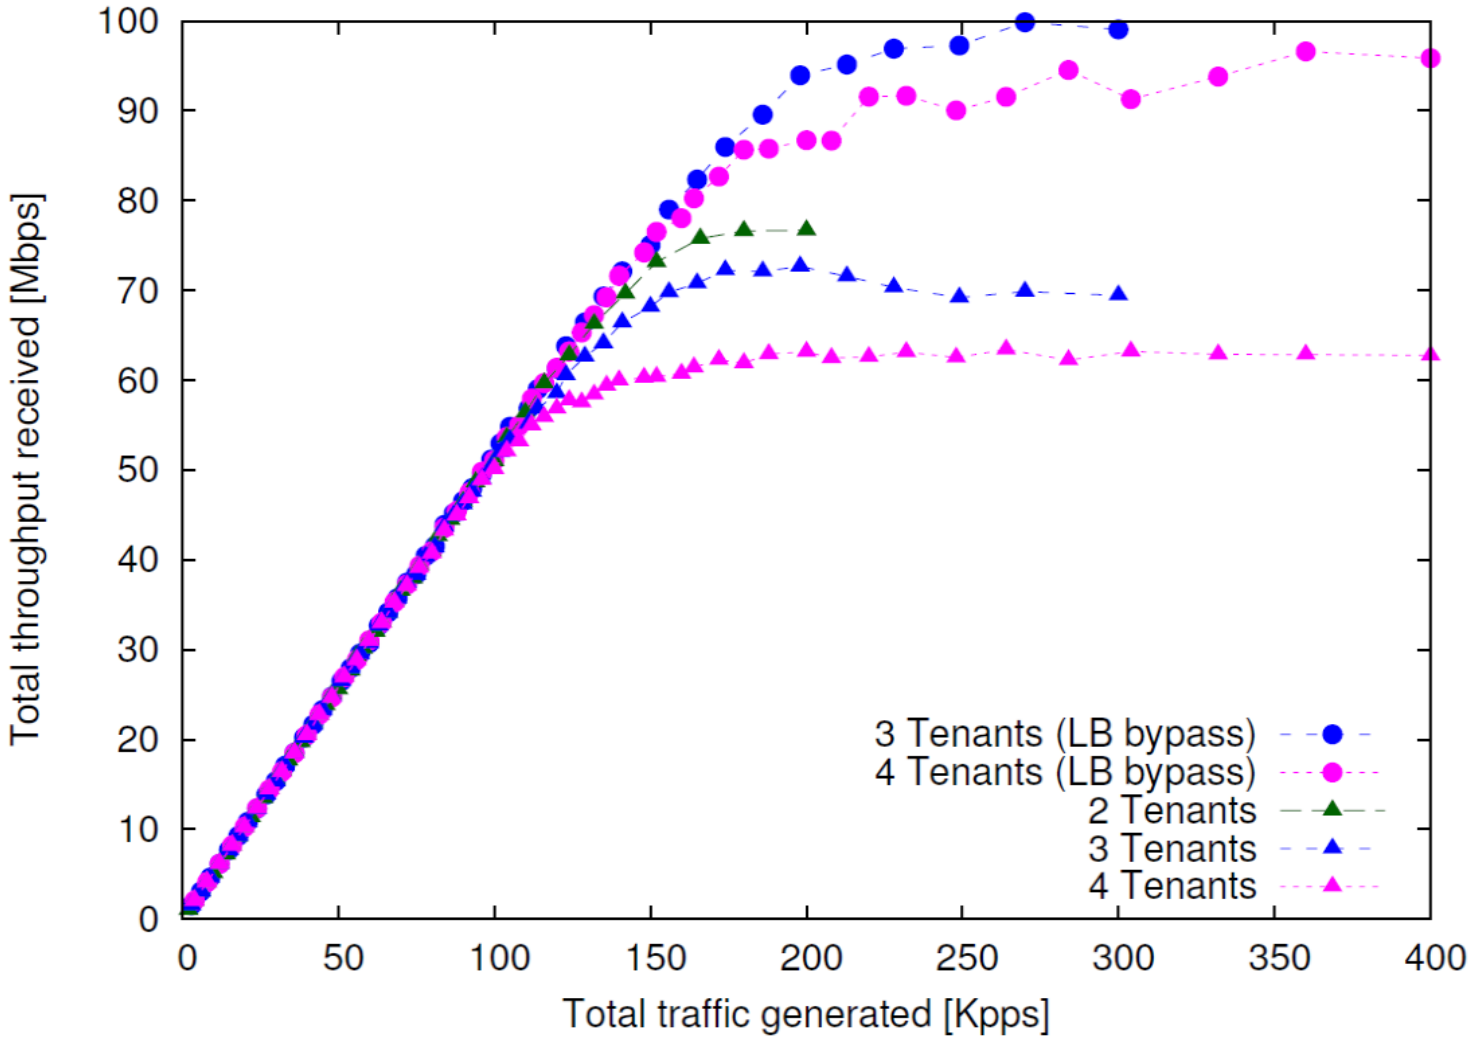

Conclusion

- Performance and complexity assessment of OpenStack virtual network infrastructure under multi-tenant scenario
  - simple NFV chain implementation
  - maximum bandwidth/sustainable packet rate
- Cloud-based architecture poses some limitations to the network performance
  - depends on the hosting hardware maximum capacity
  - but also to the complex OpenStack virtual network architecture
- Some of these limitations can be mitigated with a careful re-design of the virtual network infrastructure
  - remove Linux bridge, used for tenant isolation functions
  - implement filtering using OpenFlow rules in OVS
- Such limits must be taken into account for any engineering activity in the virtual networking arena

THANKS FOR YOUR ATTENTION!

QUESTIONS?