Virtio and kvm networking status update and plans

2013

Michael S. Tsirkin
Red Hat

VIRTIO / VHOST

KVM NETWORKING

OASIS Virtio TC

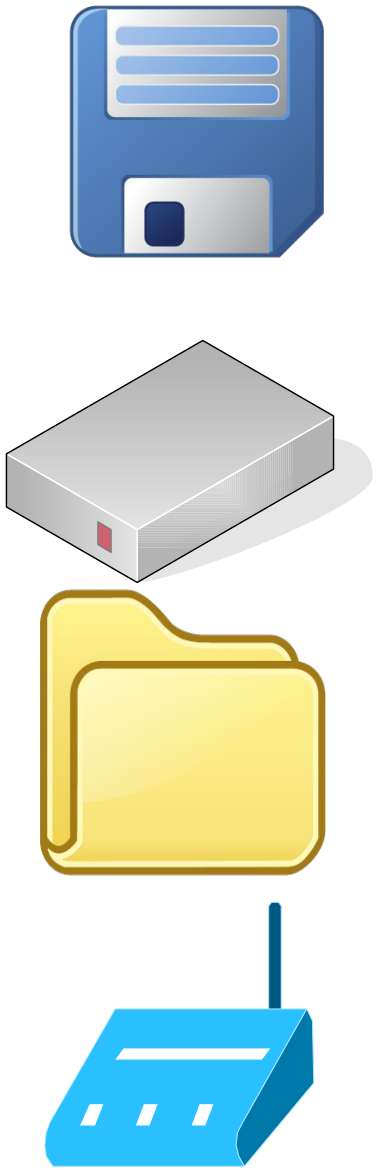

PCI

MMIO (ARM)

CCW (PPC)

Virtio 1.0

- Virtio PCI:
  - Replace Port IO with Memory mapped IO
  - PCI Express (hotplug, AER, multi-root, SRIOV)
  - Infinite features
- Reduced memory requirements
- Fixed endianness
- Compatibility

Port vs Memory mapped IO

	Port IO	MM IO
Mandated for PCI Express	✗	✓
Portable	✗	✓
HW Virtualization	✗	✓
Fast on x86	✓	✗

Port IO versus memory mapped IO on KVM x86: cycles per access (lower is better)

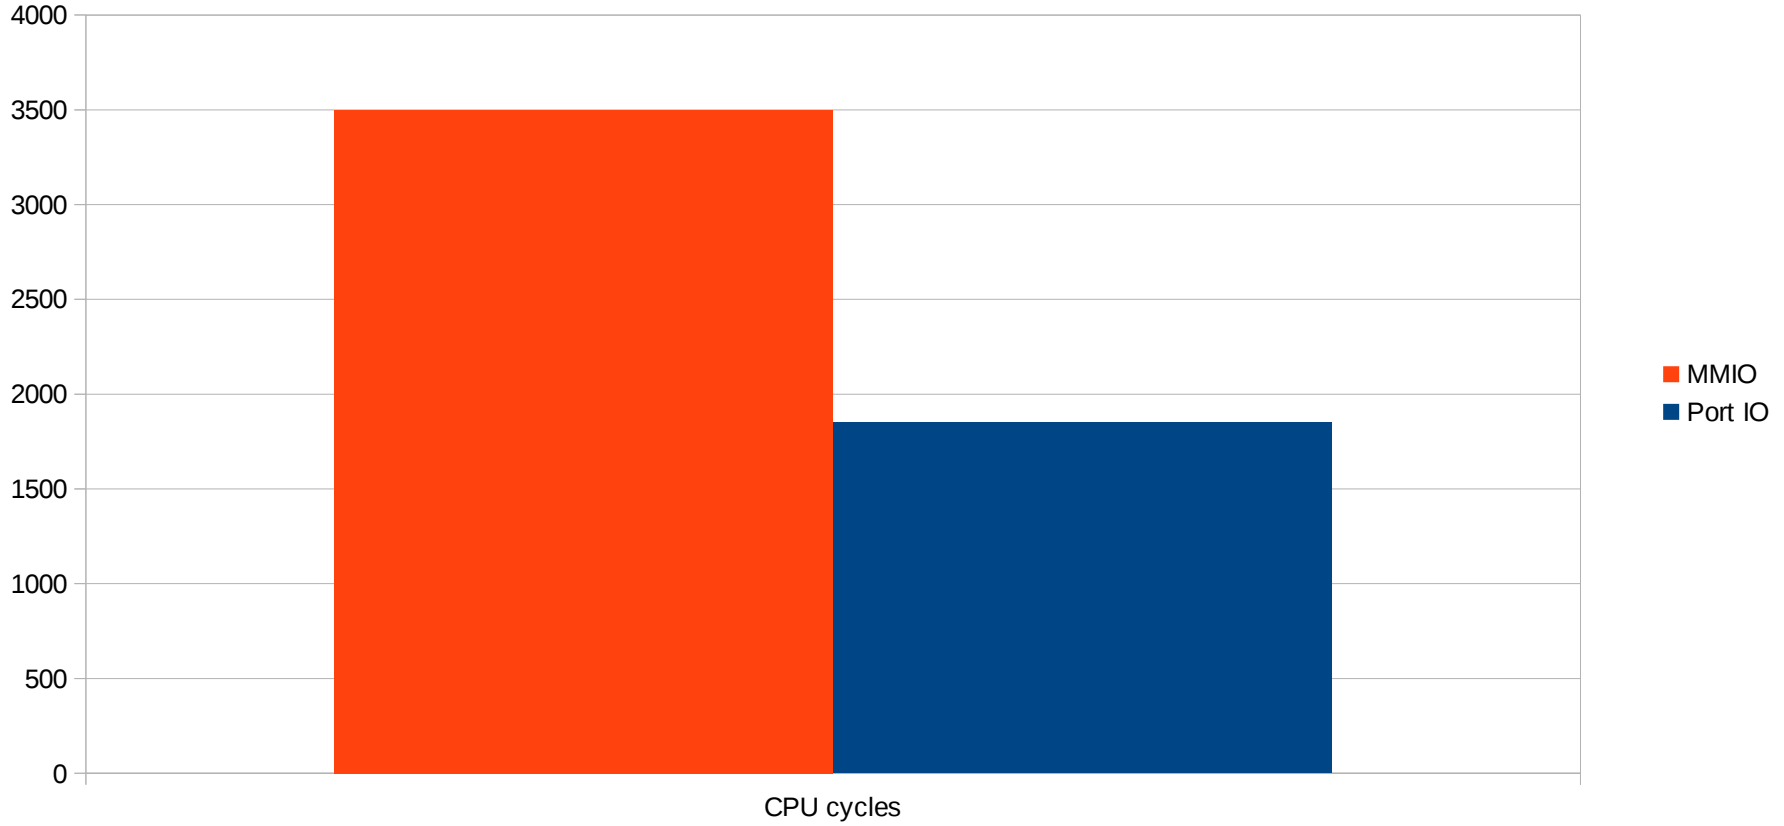

Port IO: outl

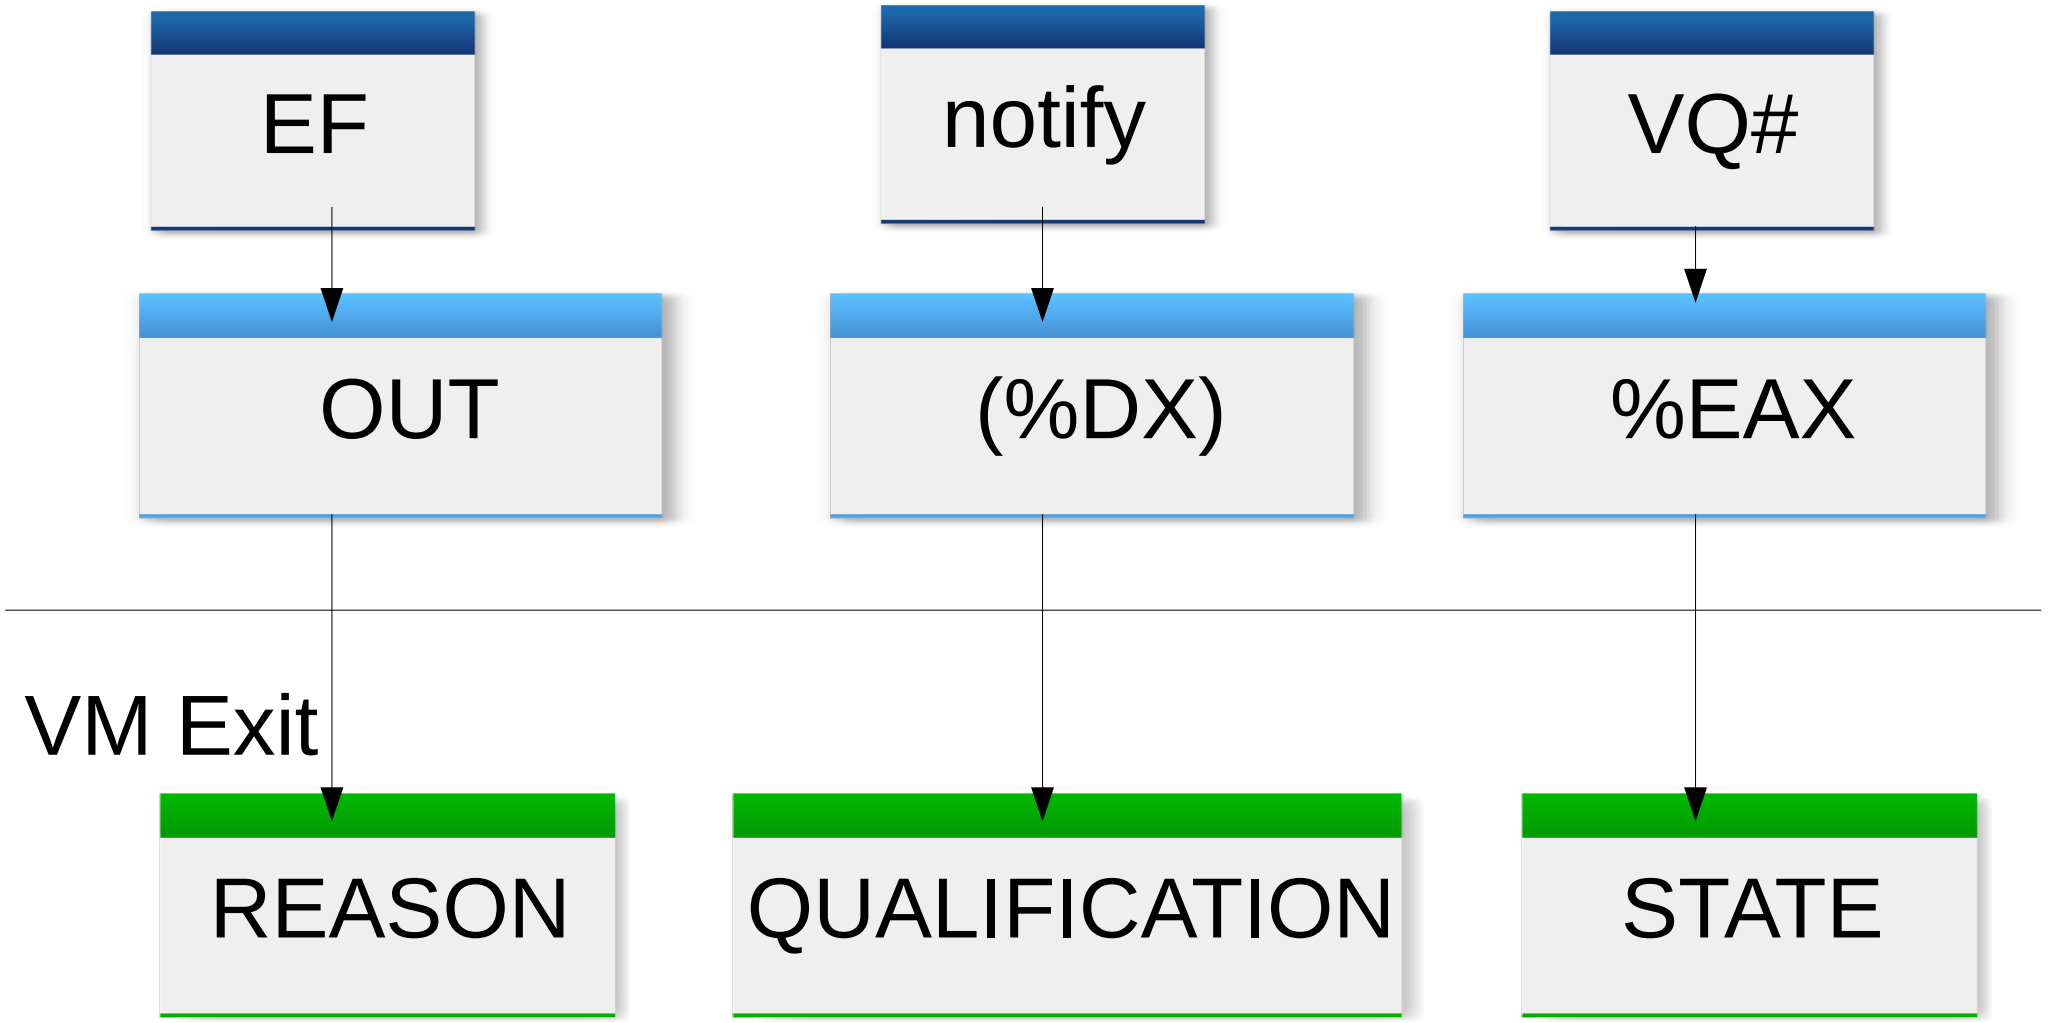

Memory mapped IO: writel

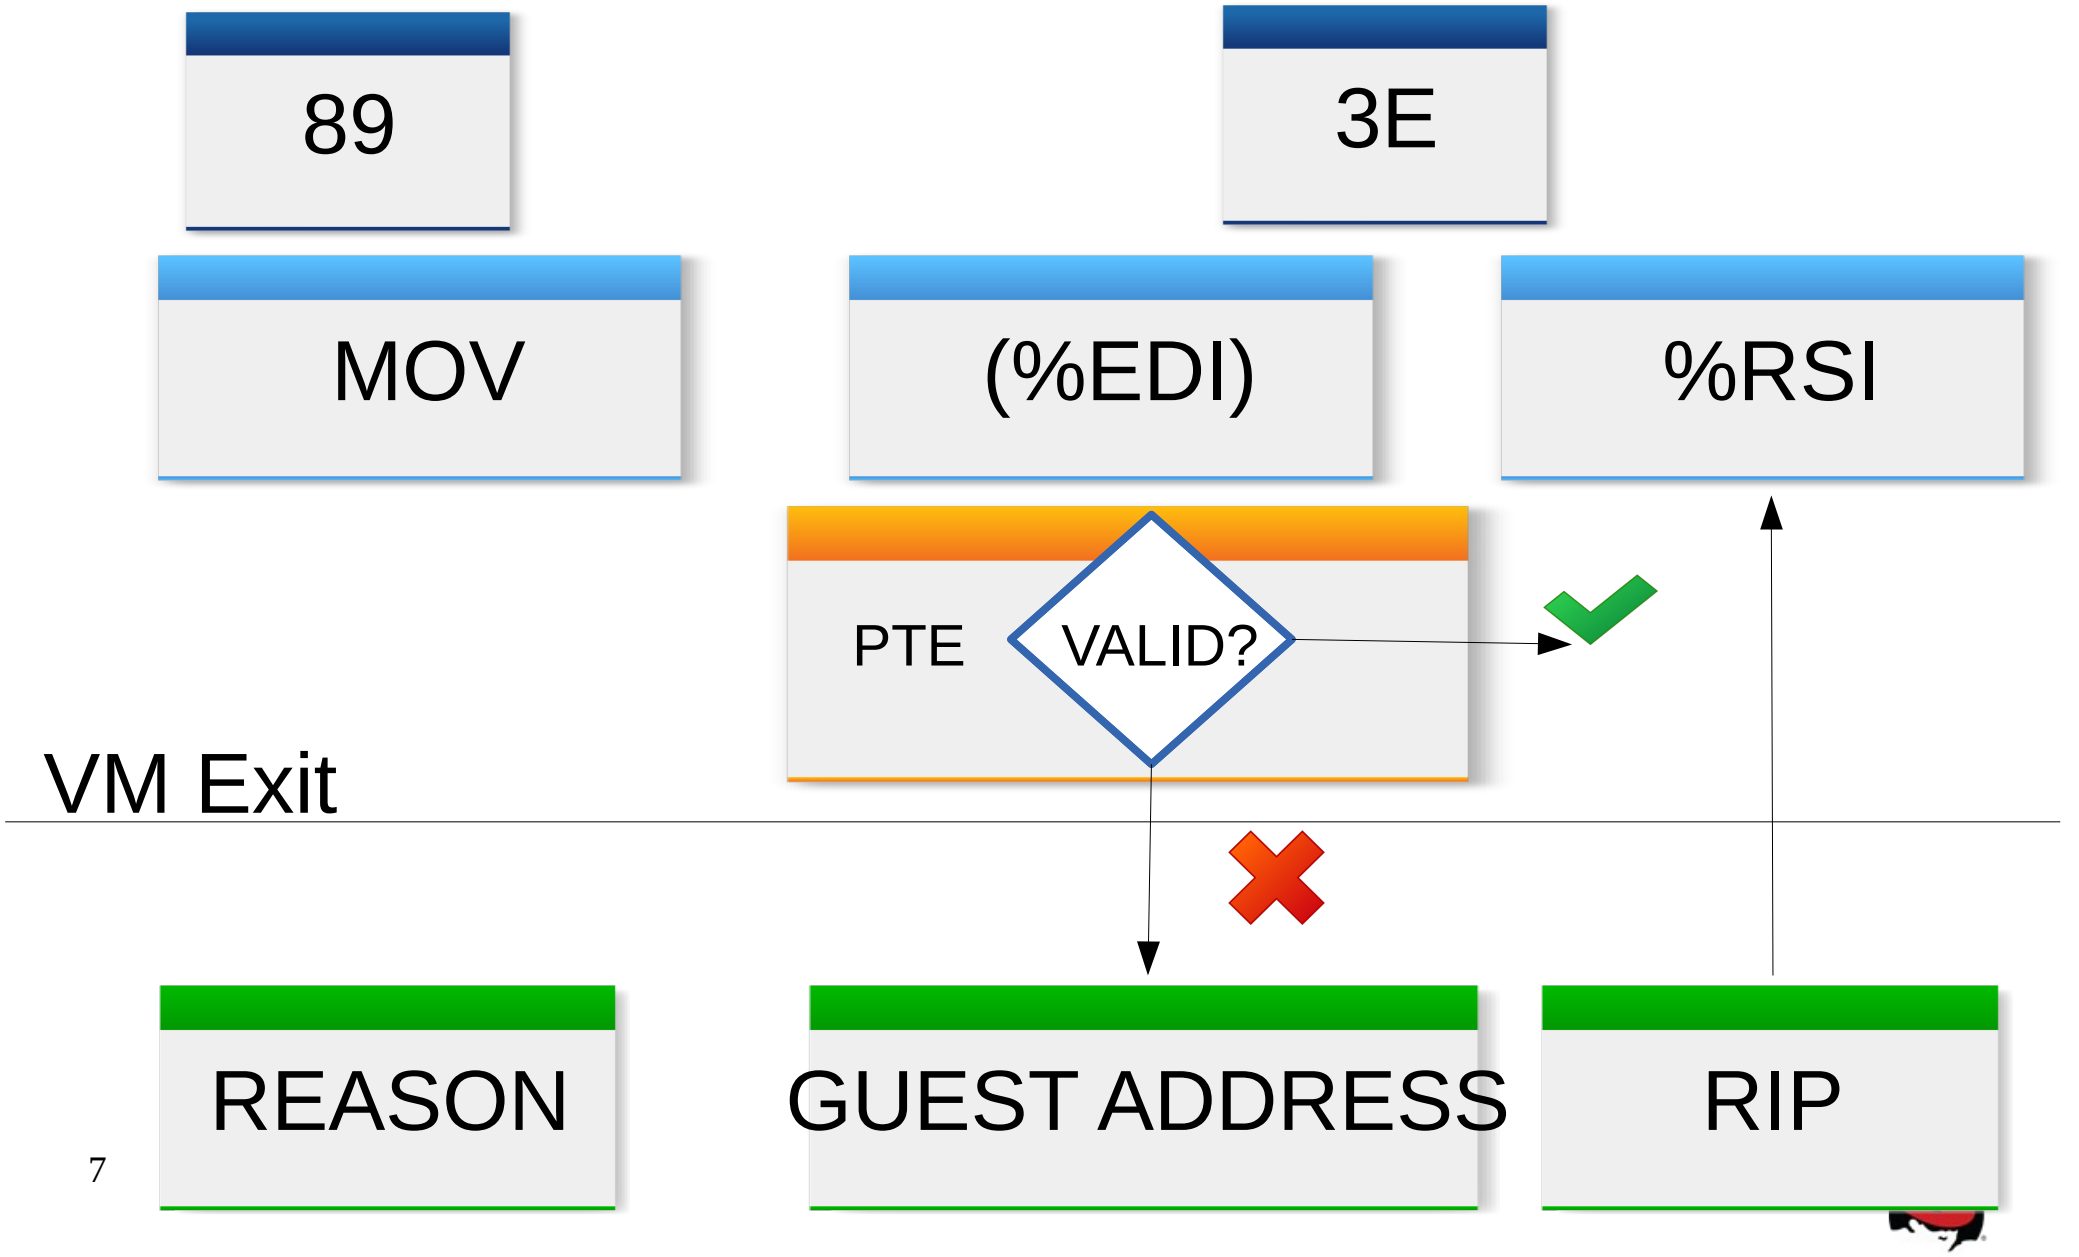

Fast MMIO

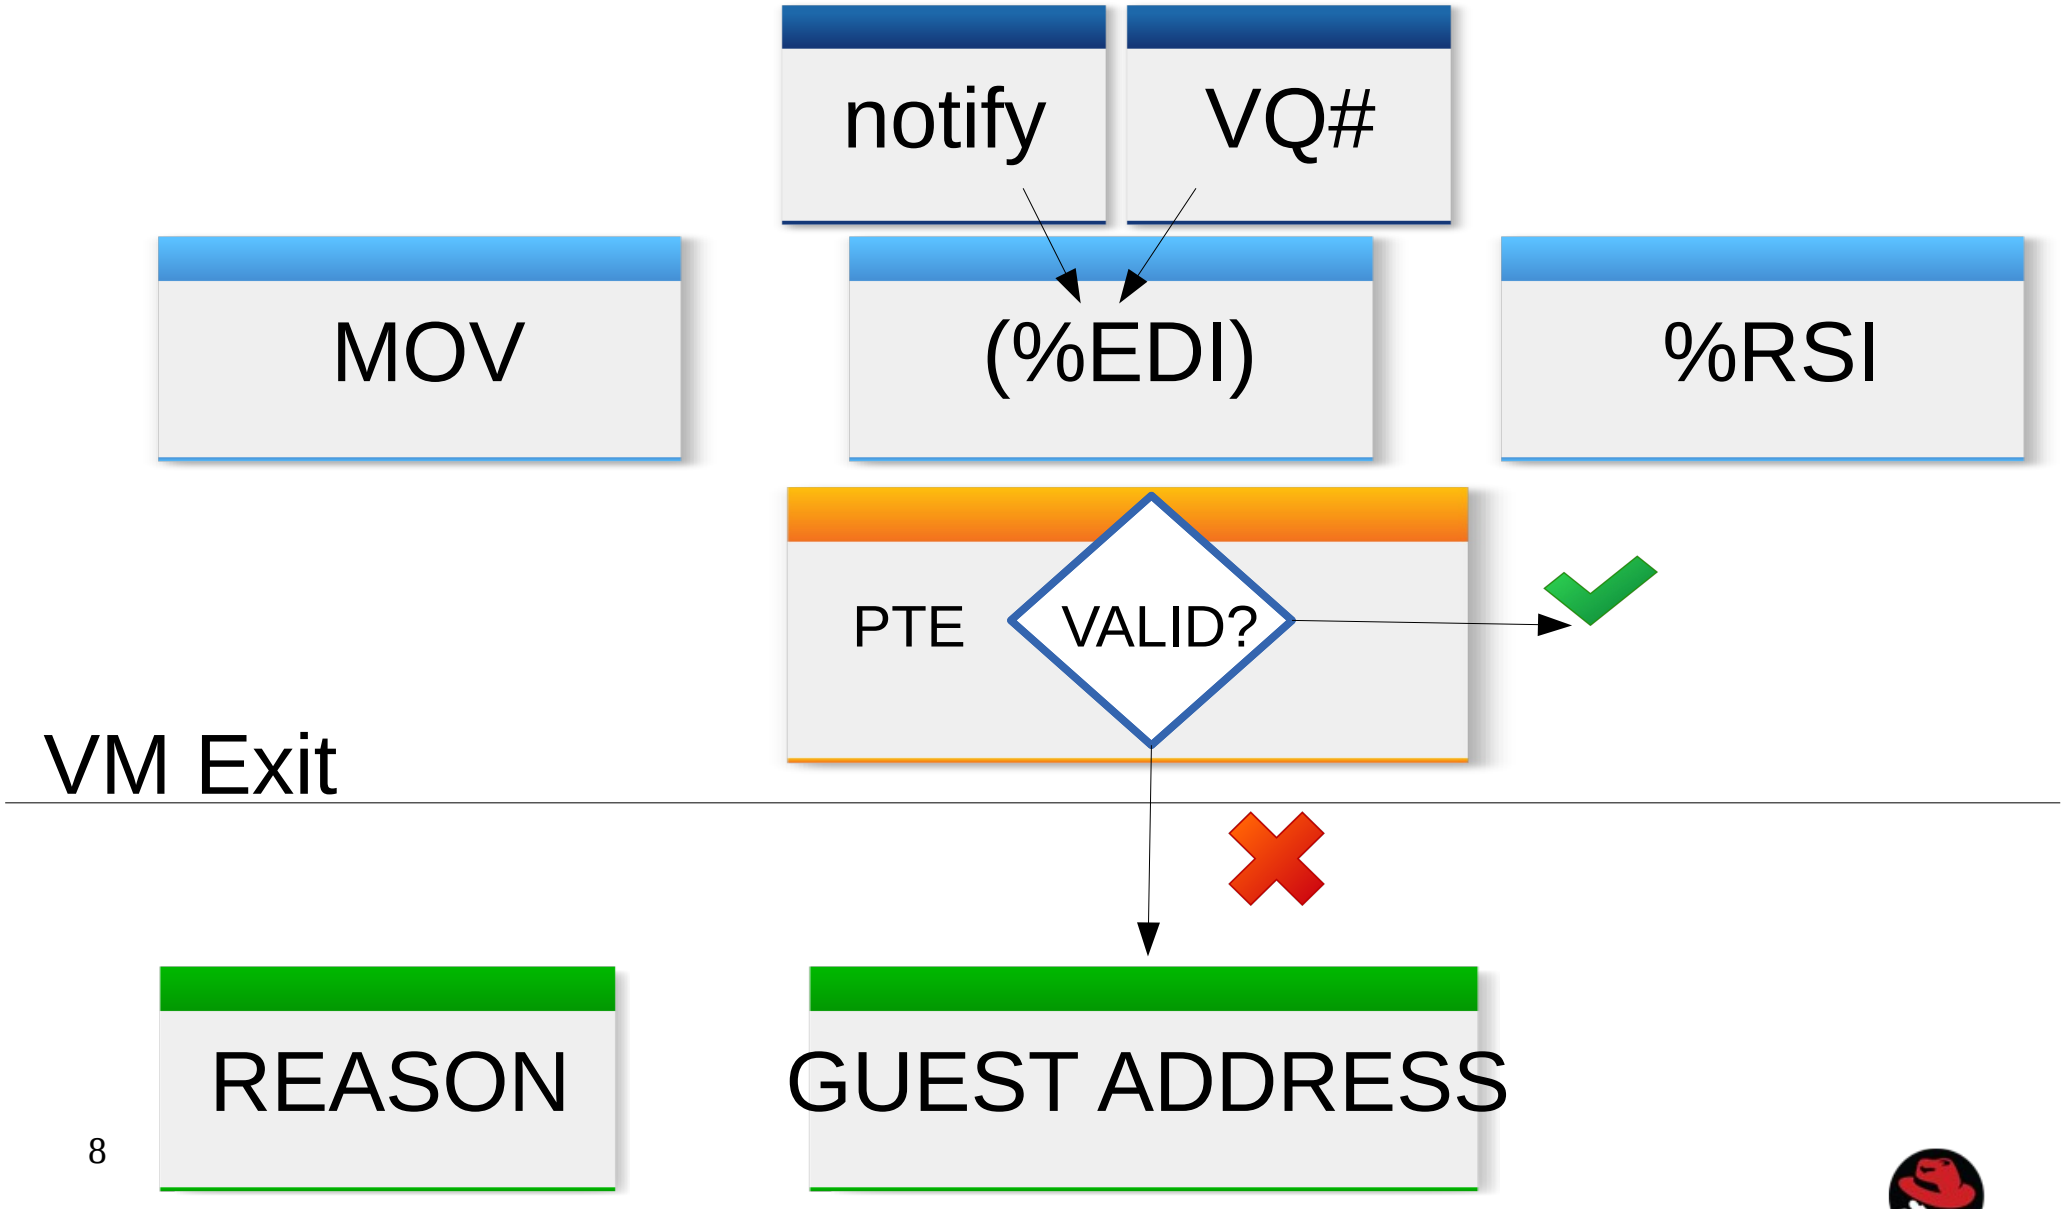

Access times on KVM x86: Cycles per access (lower is better)

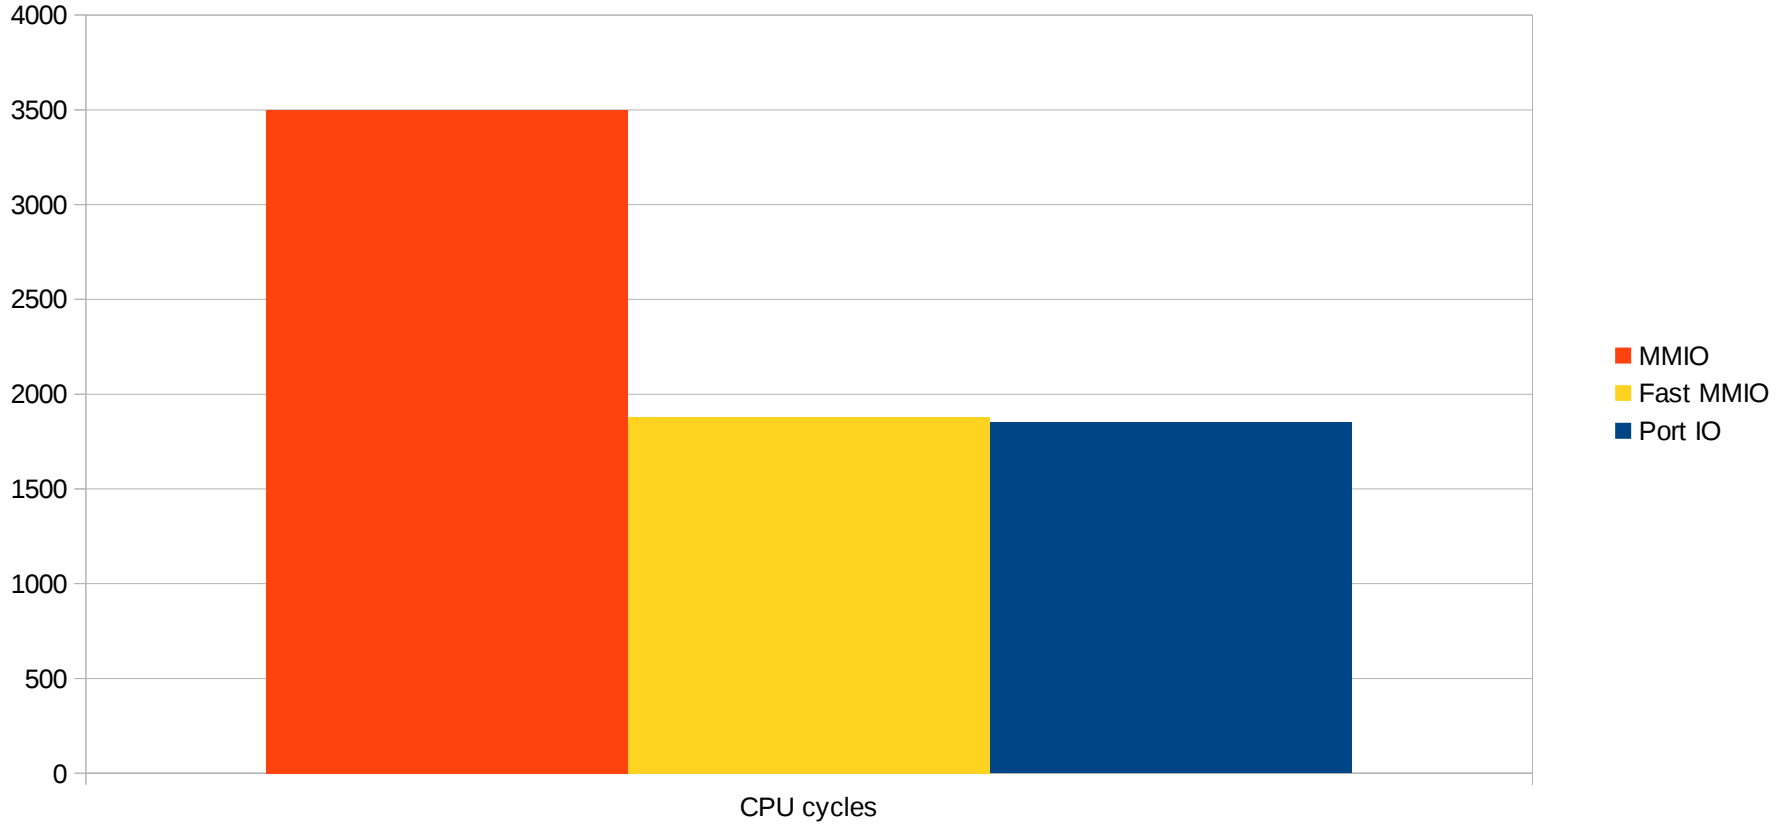

Multiple interfaces

Memory requirements

0.9

1.0

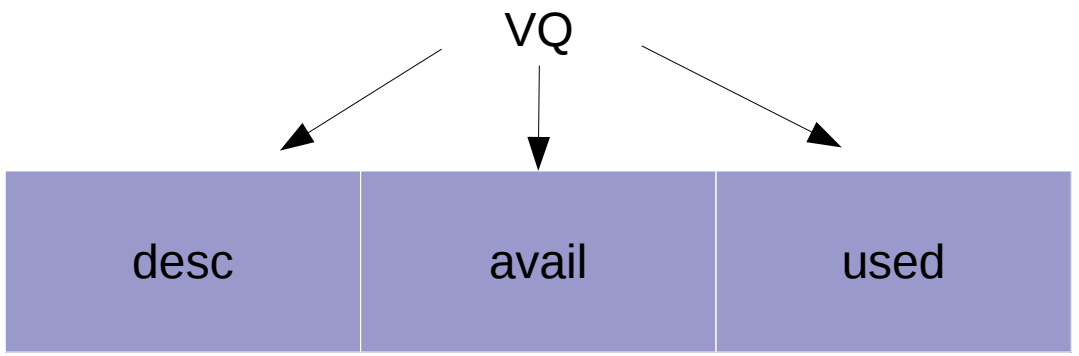

0.9

features

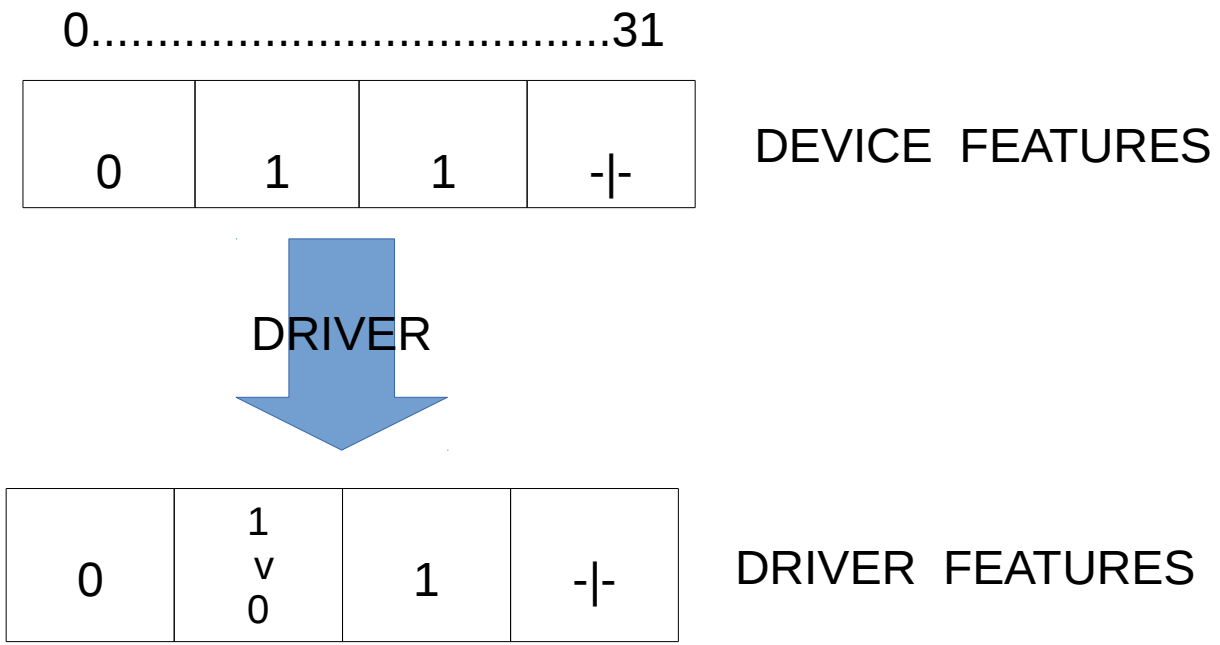

1.0

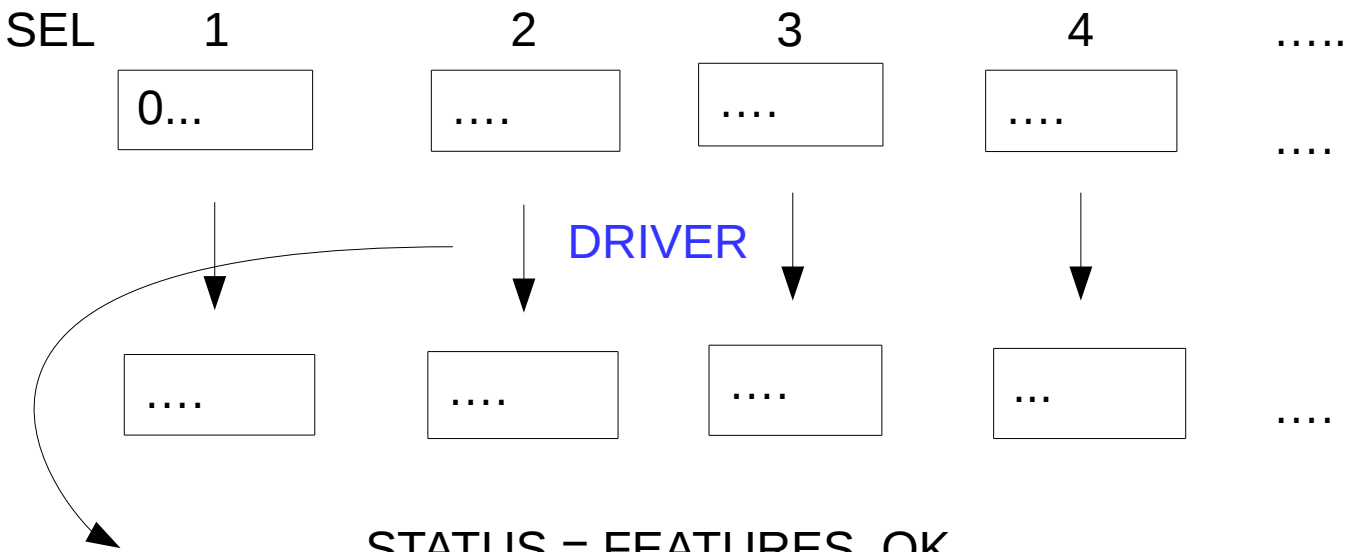

Endianness

Virtio 0.9

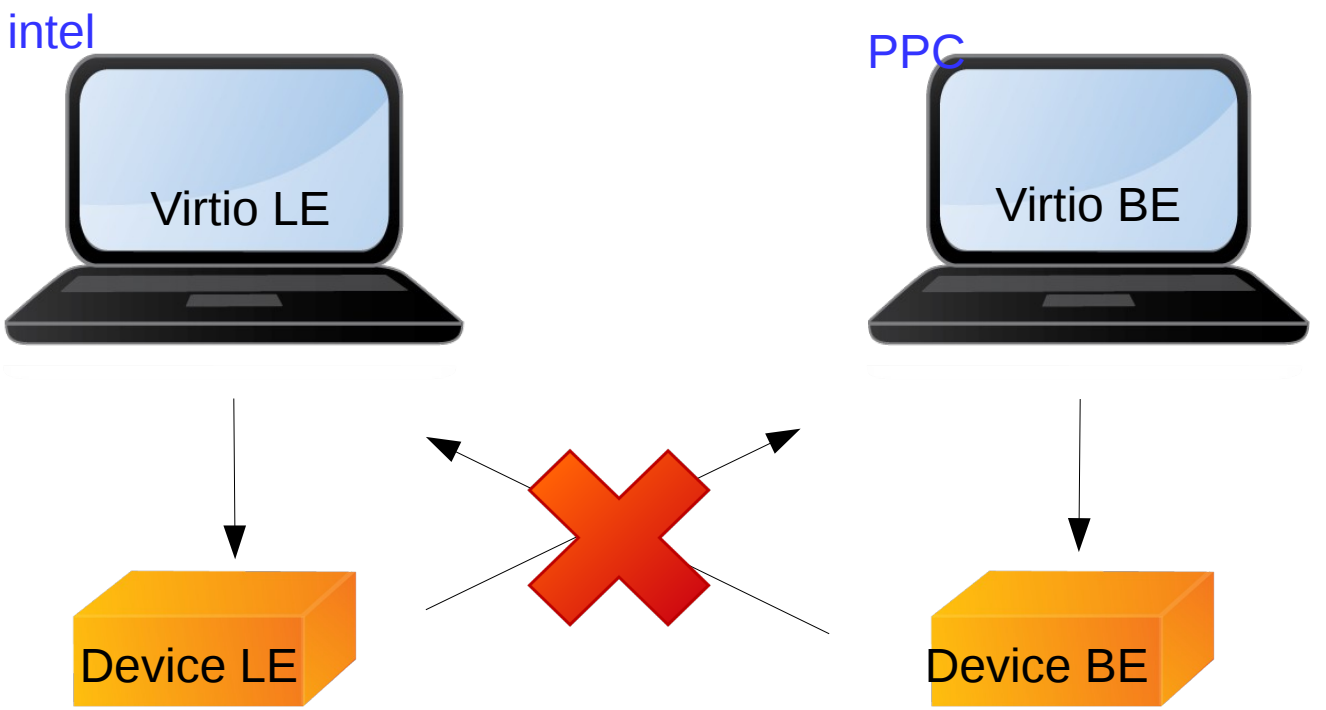

Virtio 1.0

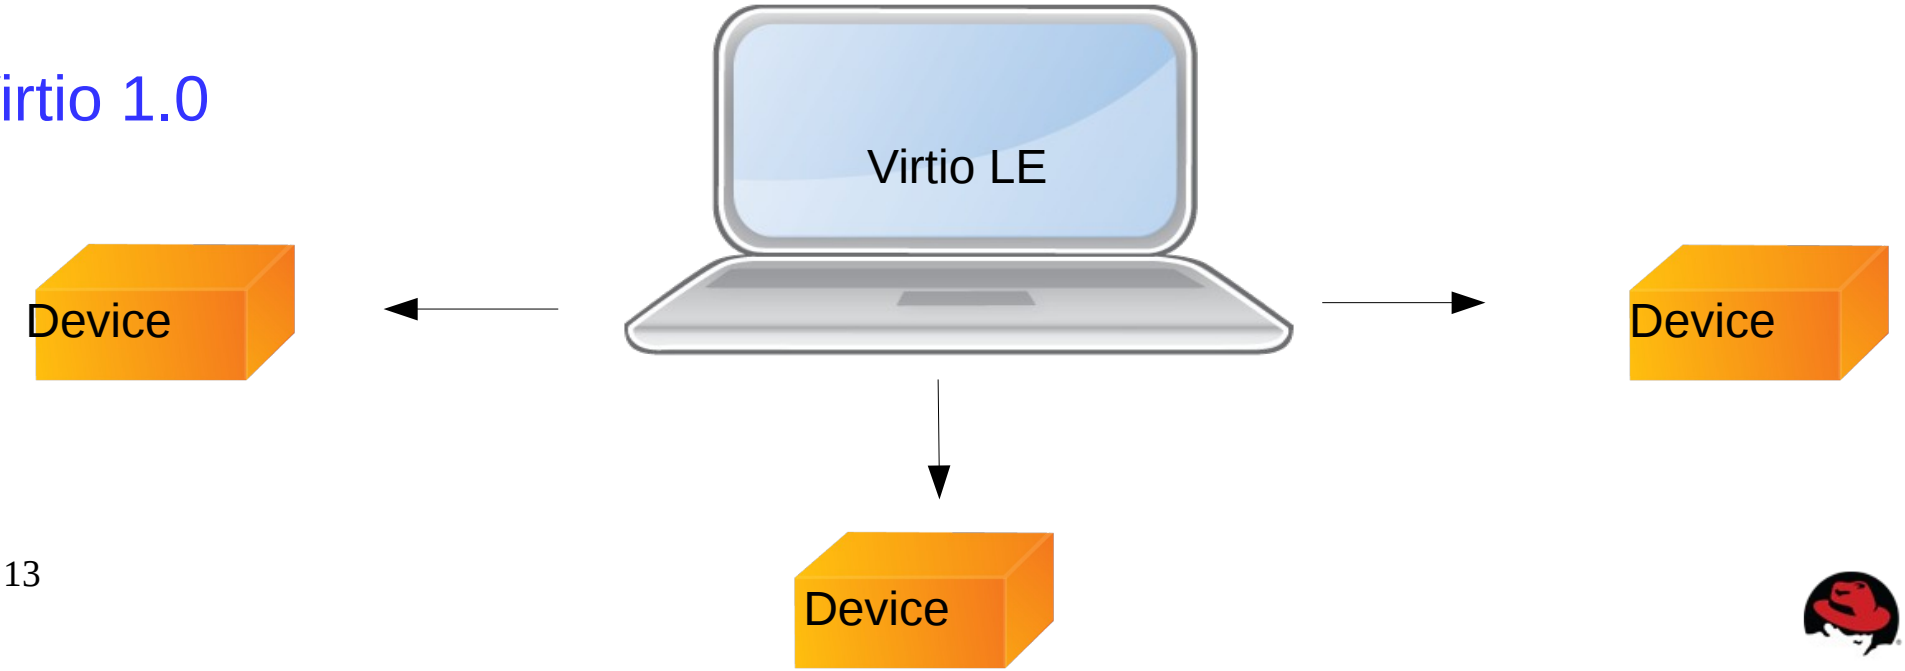

compatibility

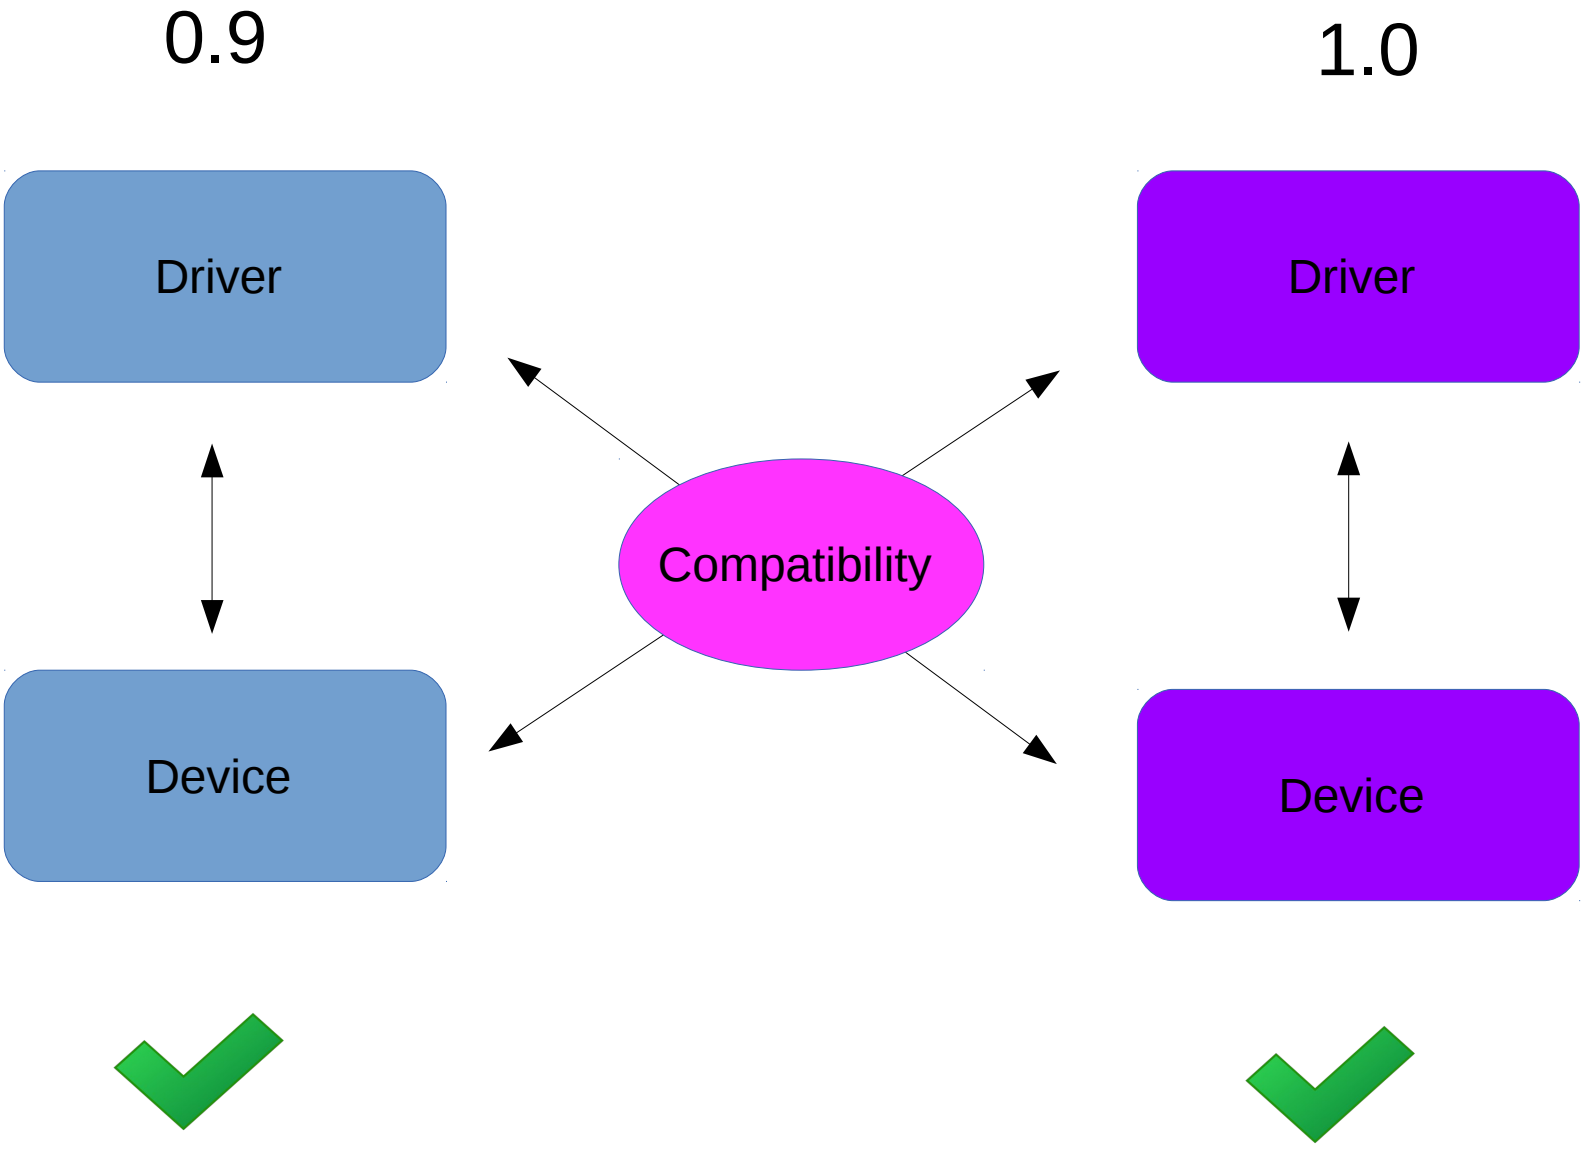

Packet layout

Virtio 0.9

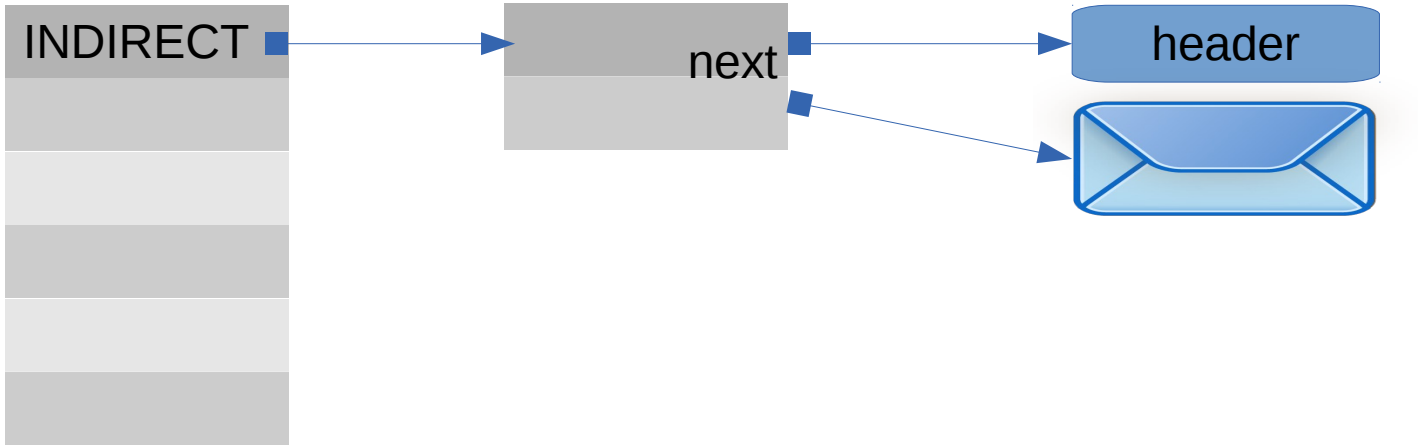

Virtio 1.0

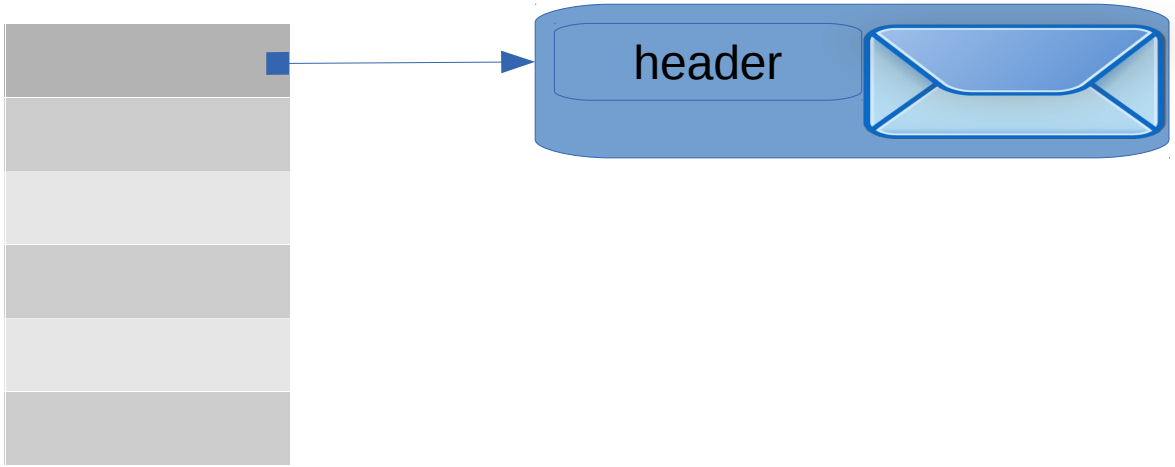

Packet layout: transactions per sec (higher is better)

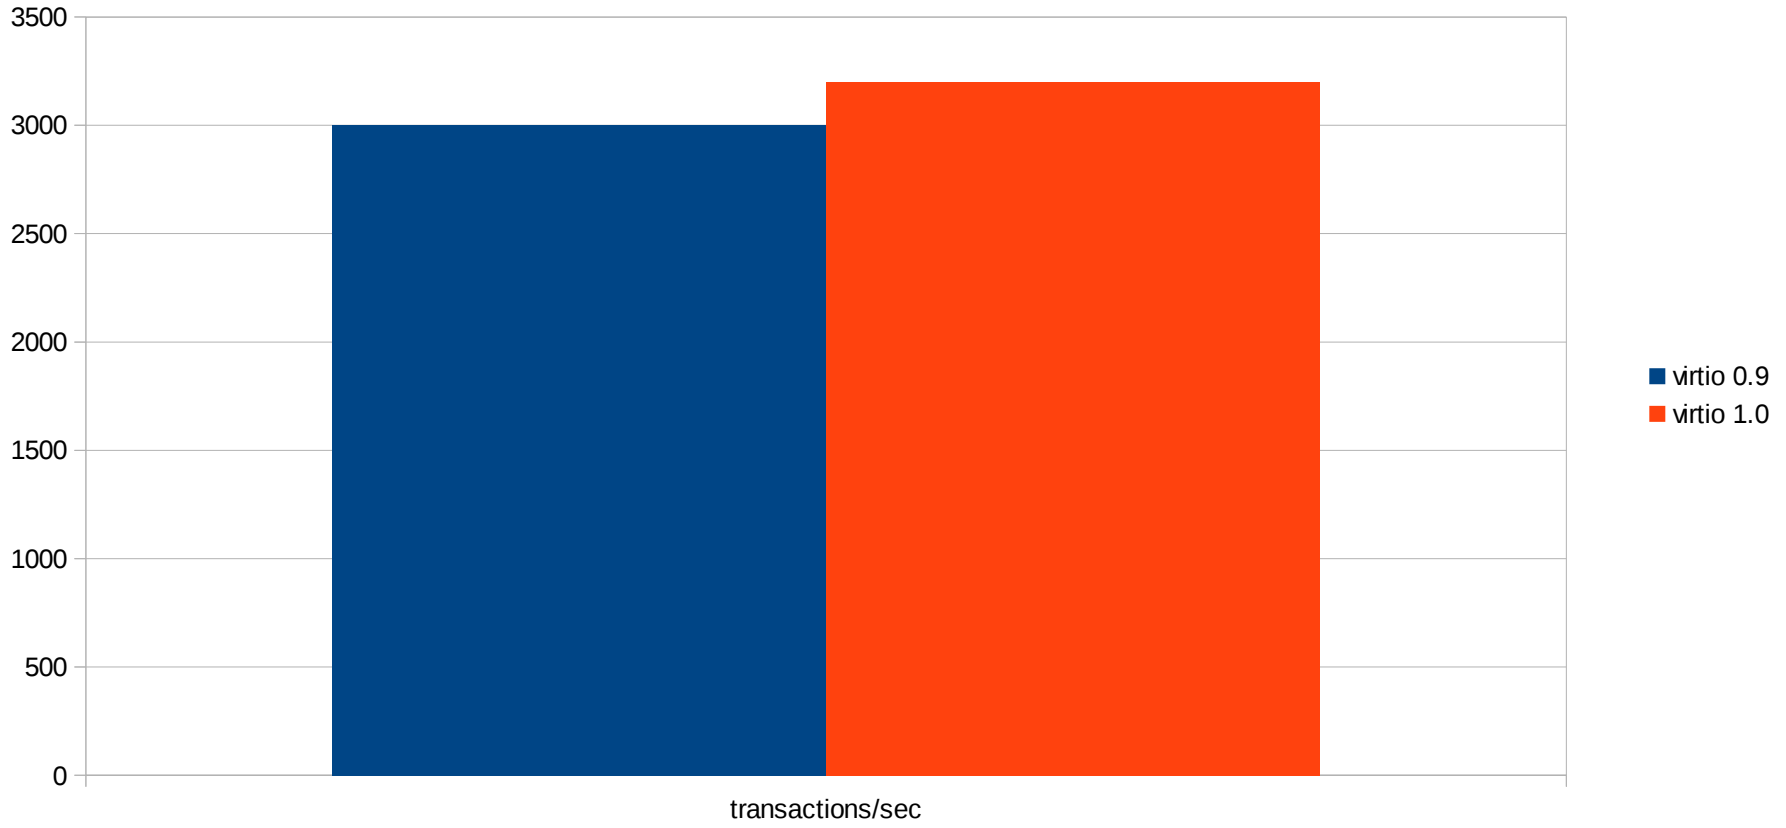

More: virtio 1.0 versus 0.9.5

- Virtio 9p 
- Virtio blk: WCE 
- Virtio-net Multiqueue 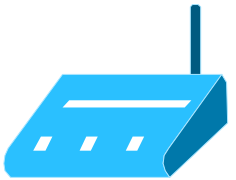
- Virtio-net dynamic offloads
- Already upstream (based on spec draft)

vhost updates

- Vhost scsi  
- Vhost-net zero copy transmit  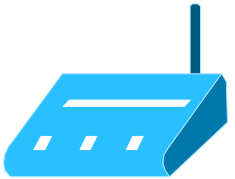
- No need for driver changes

Kvm networking

- Openvswitch – if time allows

- Ethernet bridge

Bridge FDB

London	Heathrow
Paris	CDG

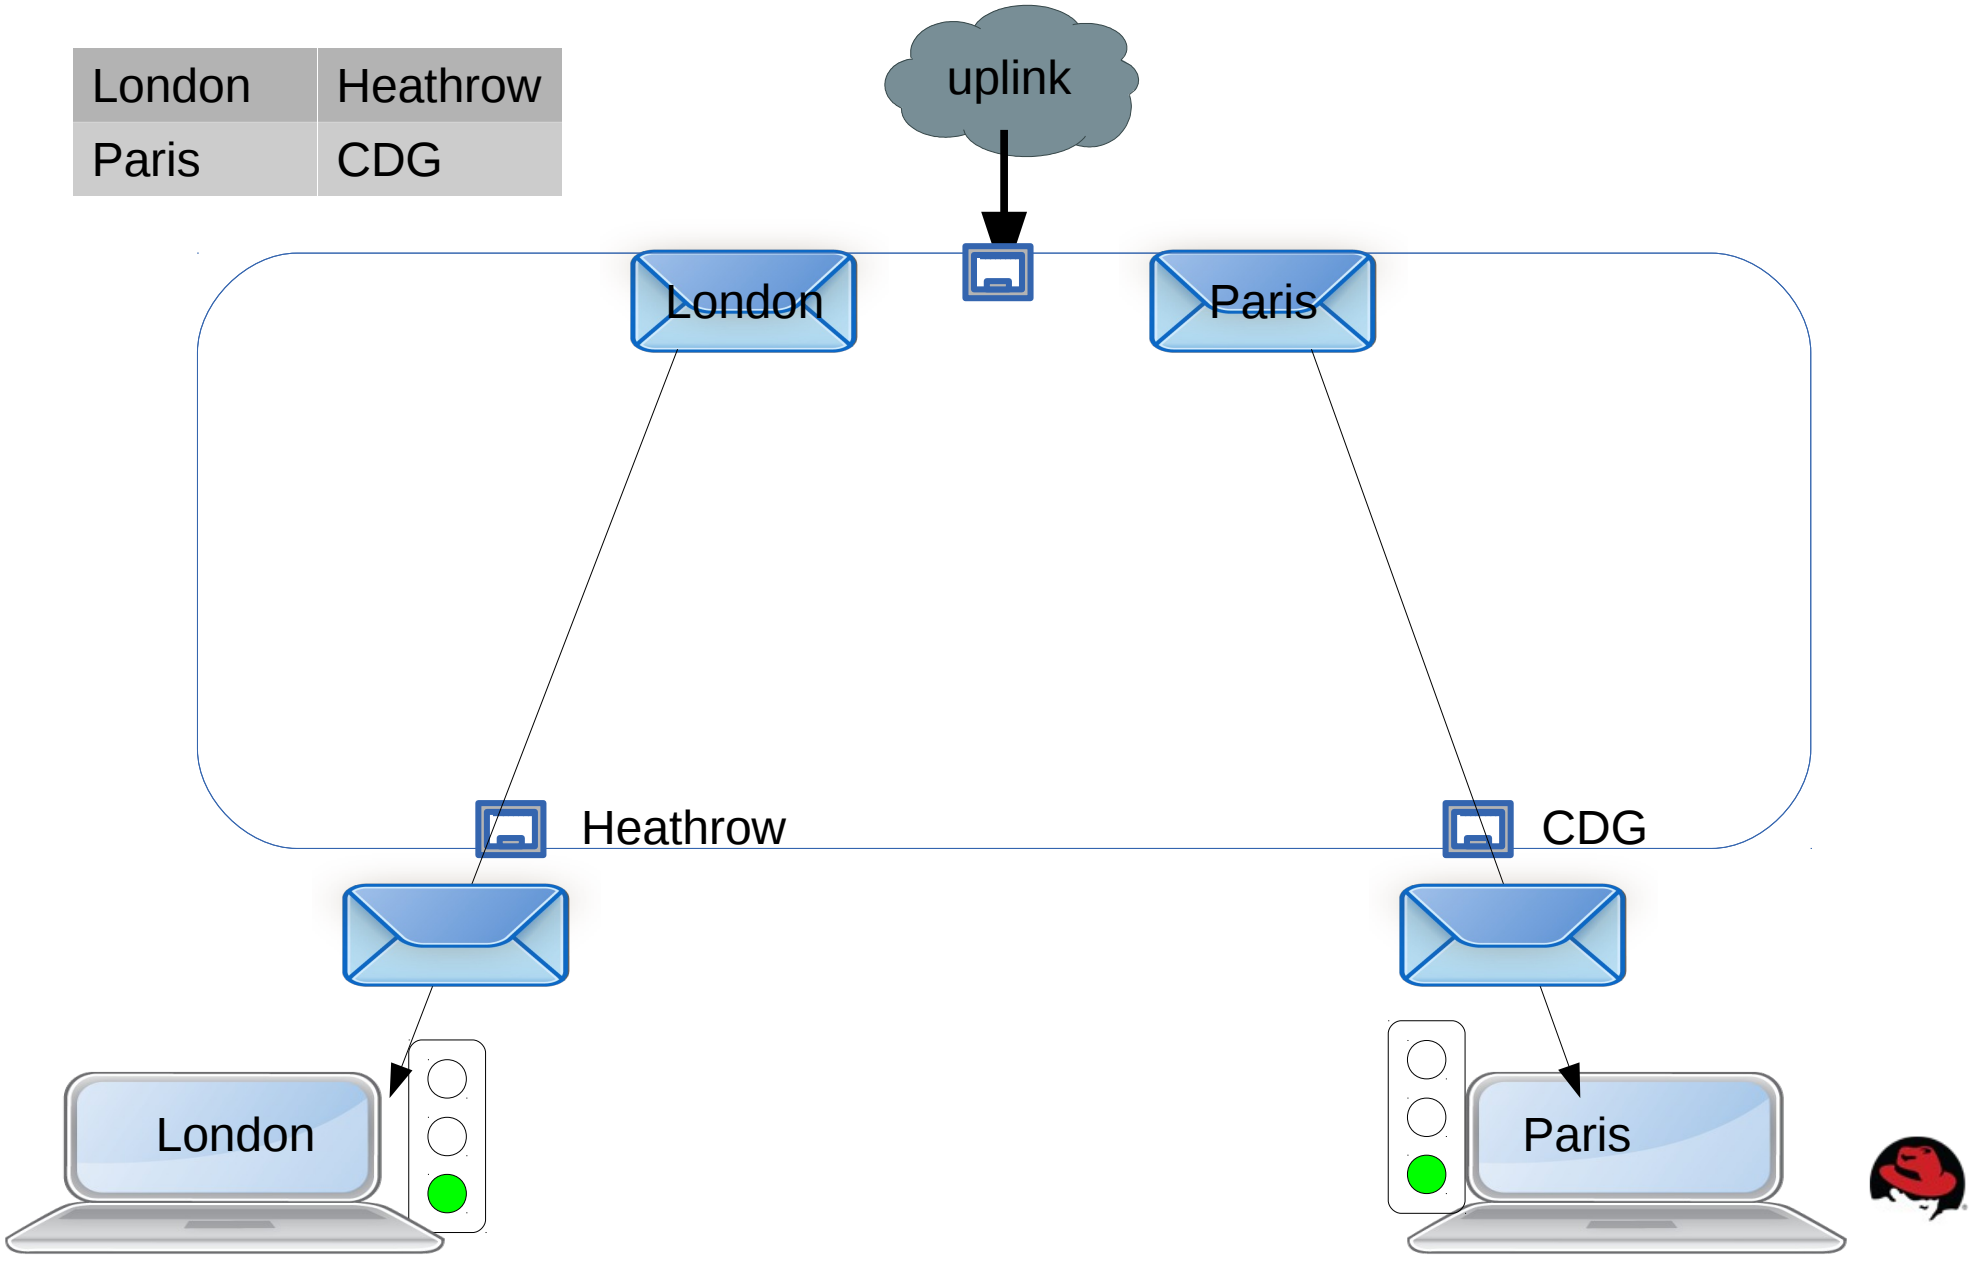

Flood: DOS potential

London	Heathrow
Paris	CDG

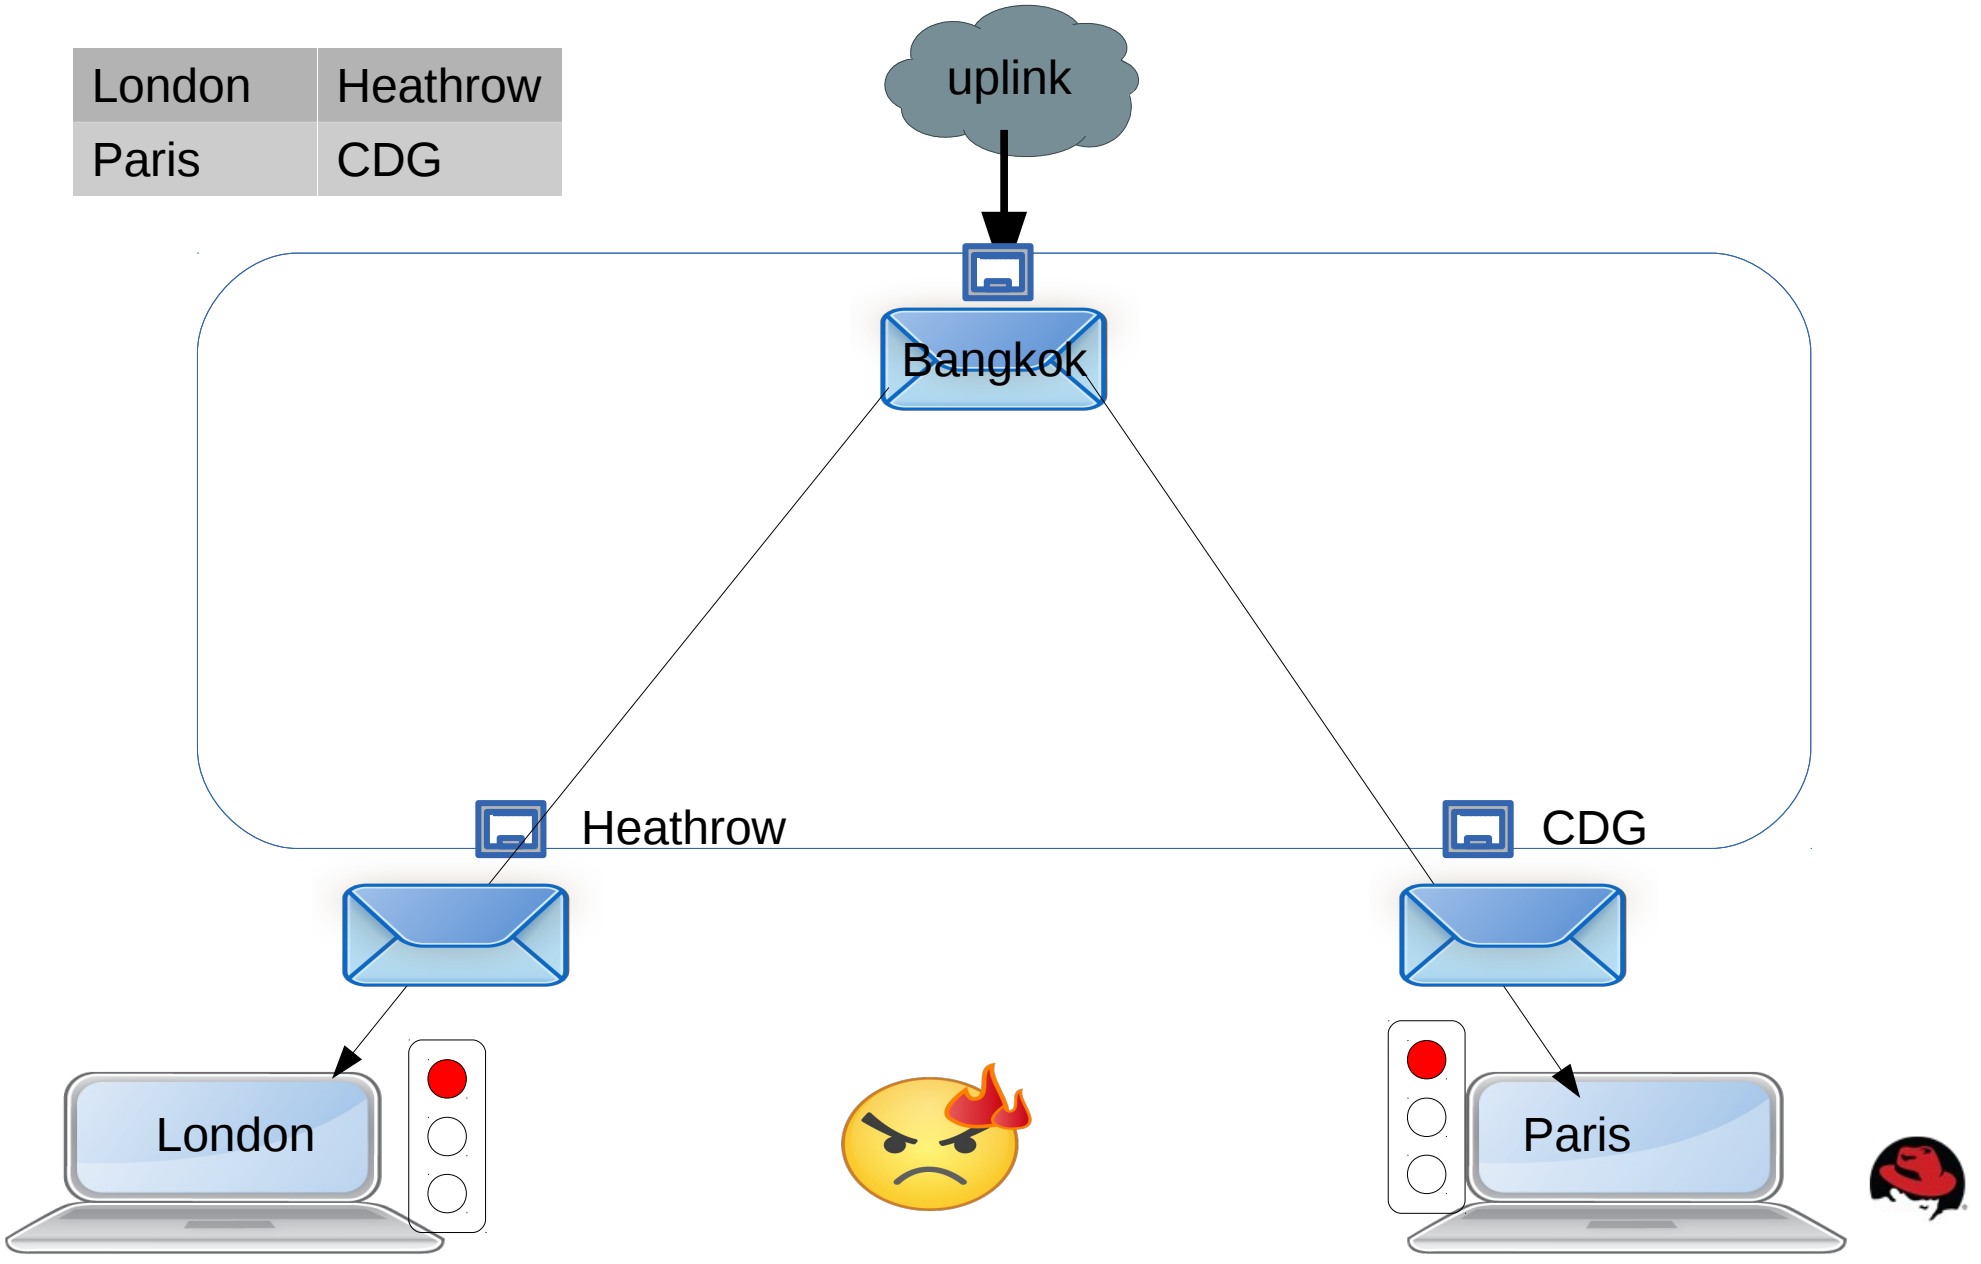

Disable flood

London	Heathrow
Paris	CDG

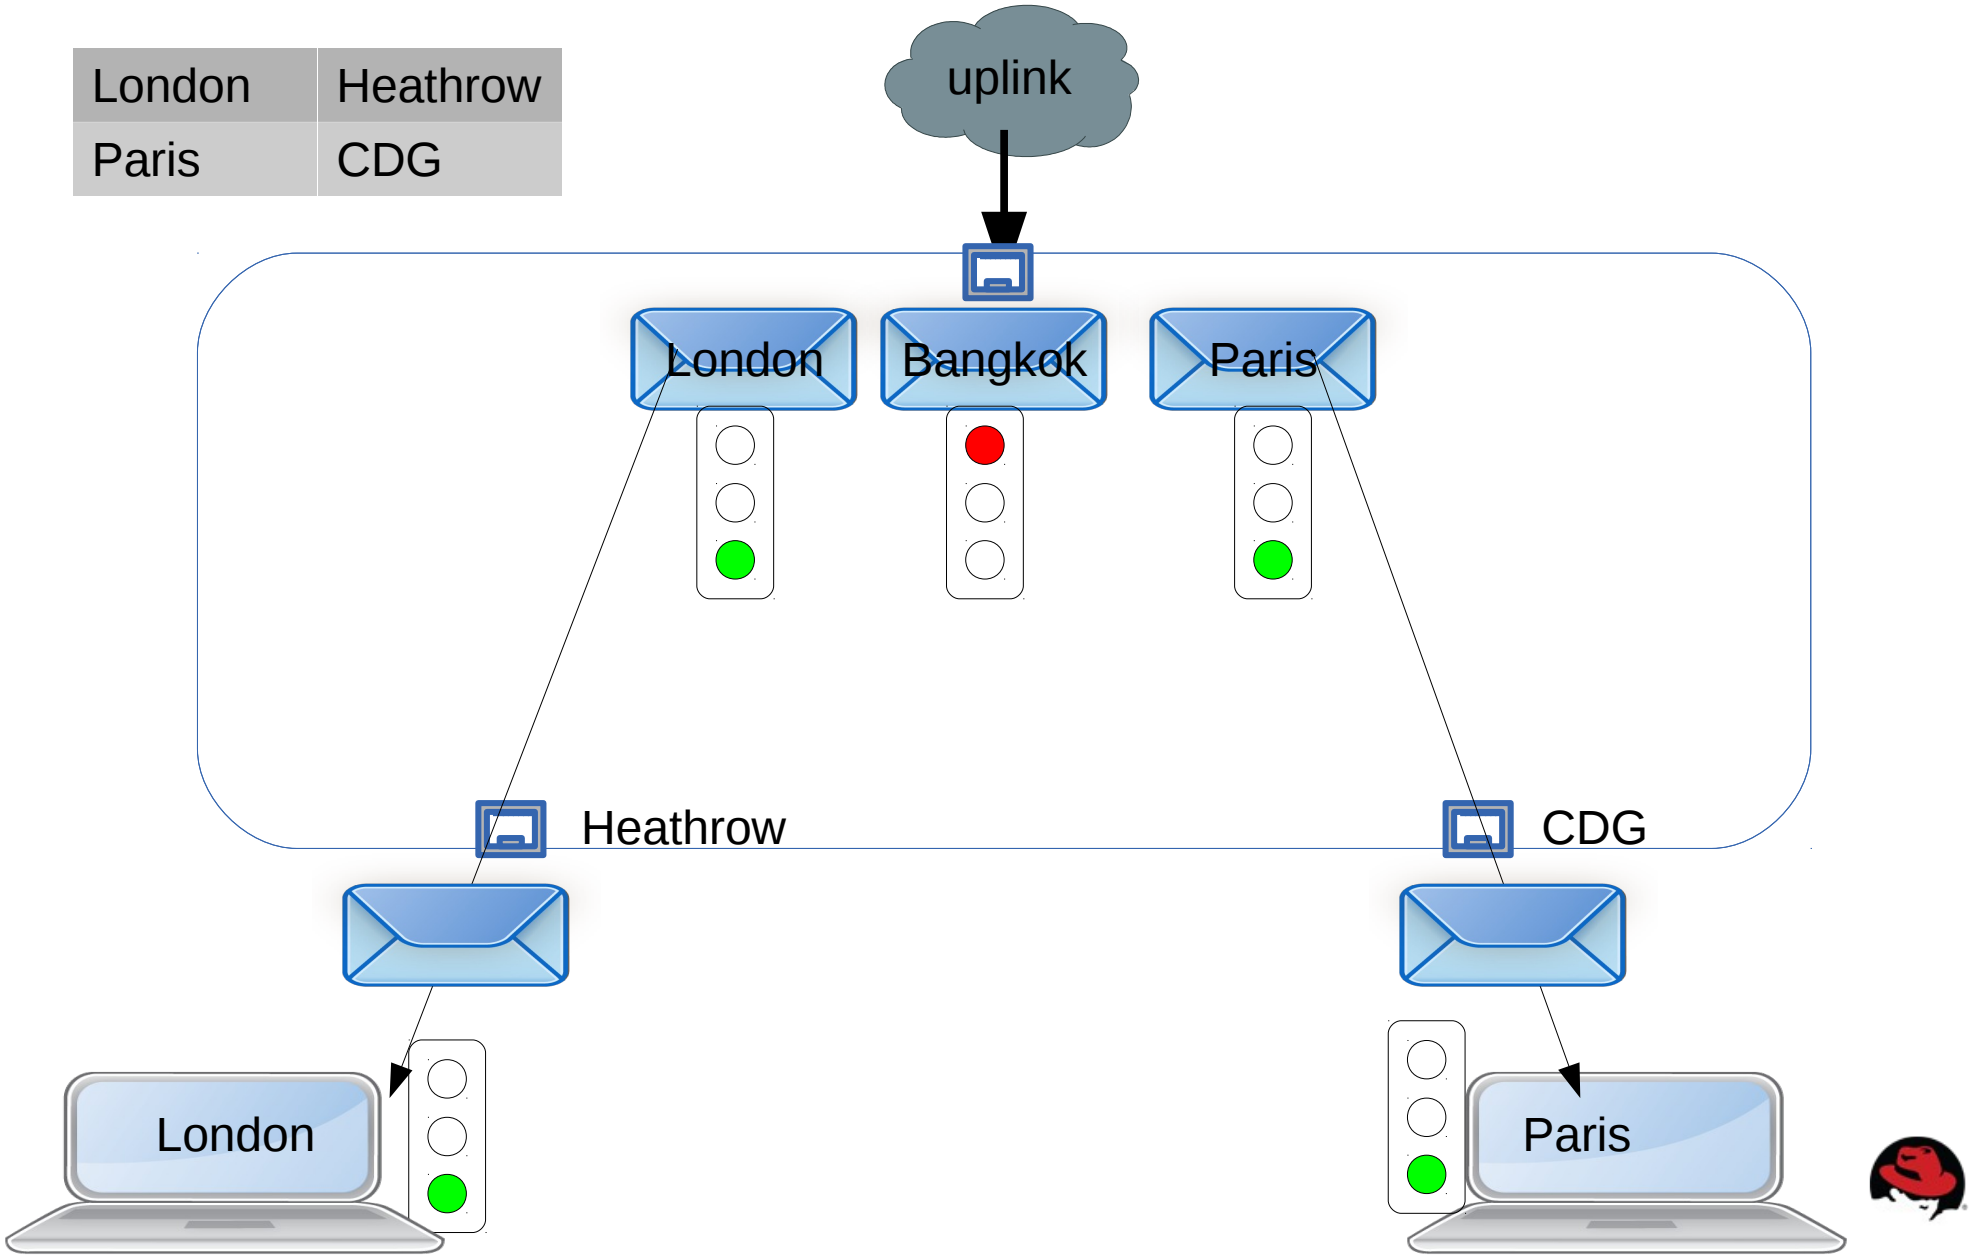

softmac

- Ifconfig eth0 hw ether 00:12:23:45:67:89

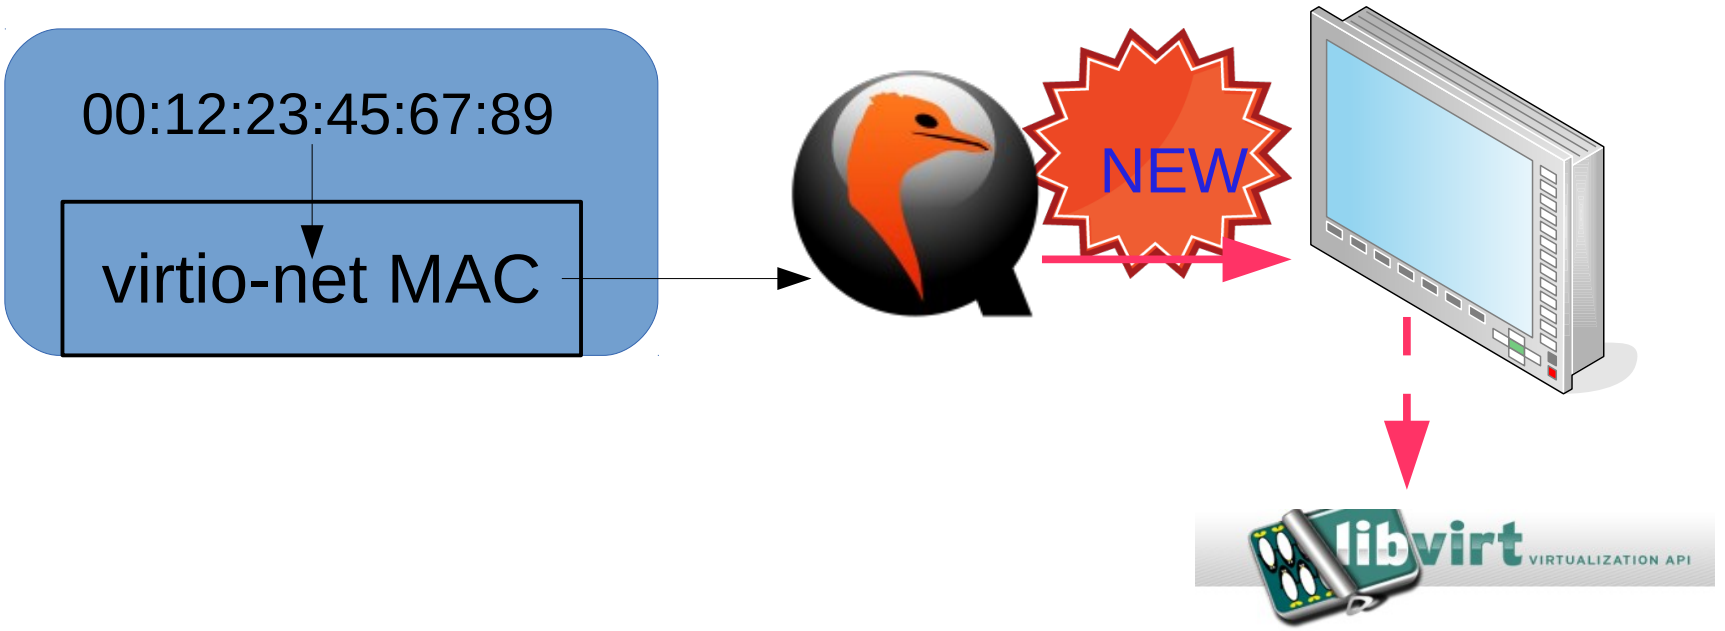

Using softmac/non promiscuous

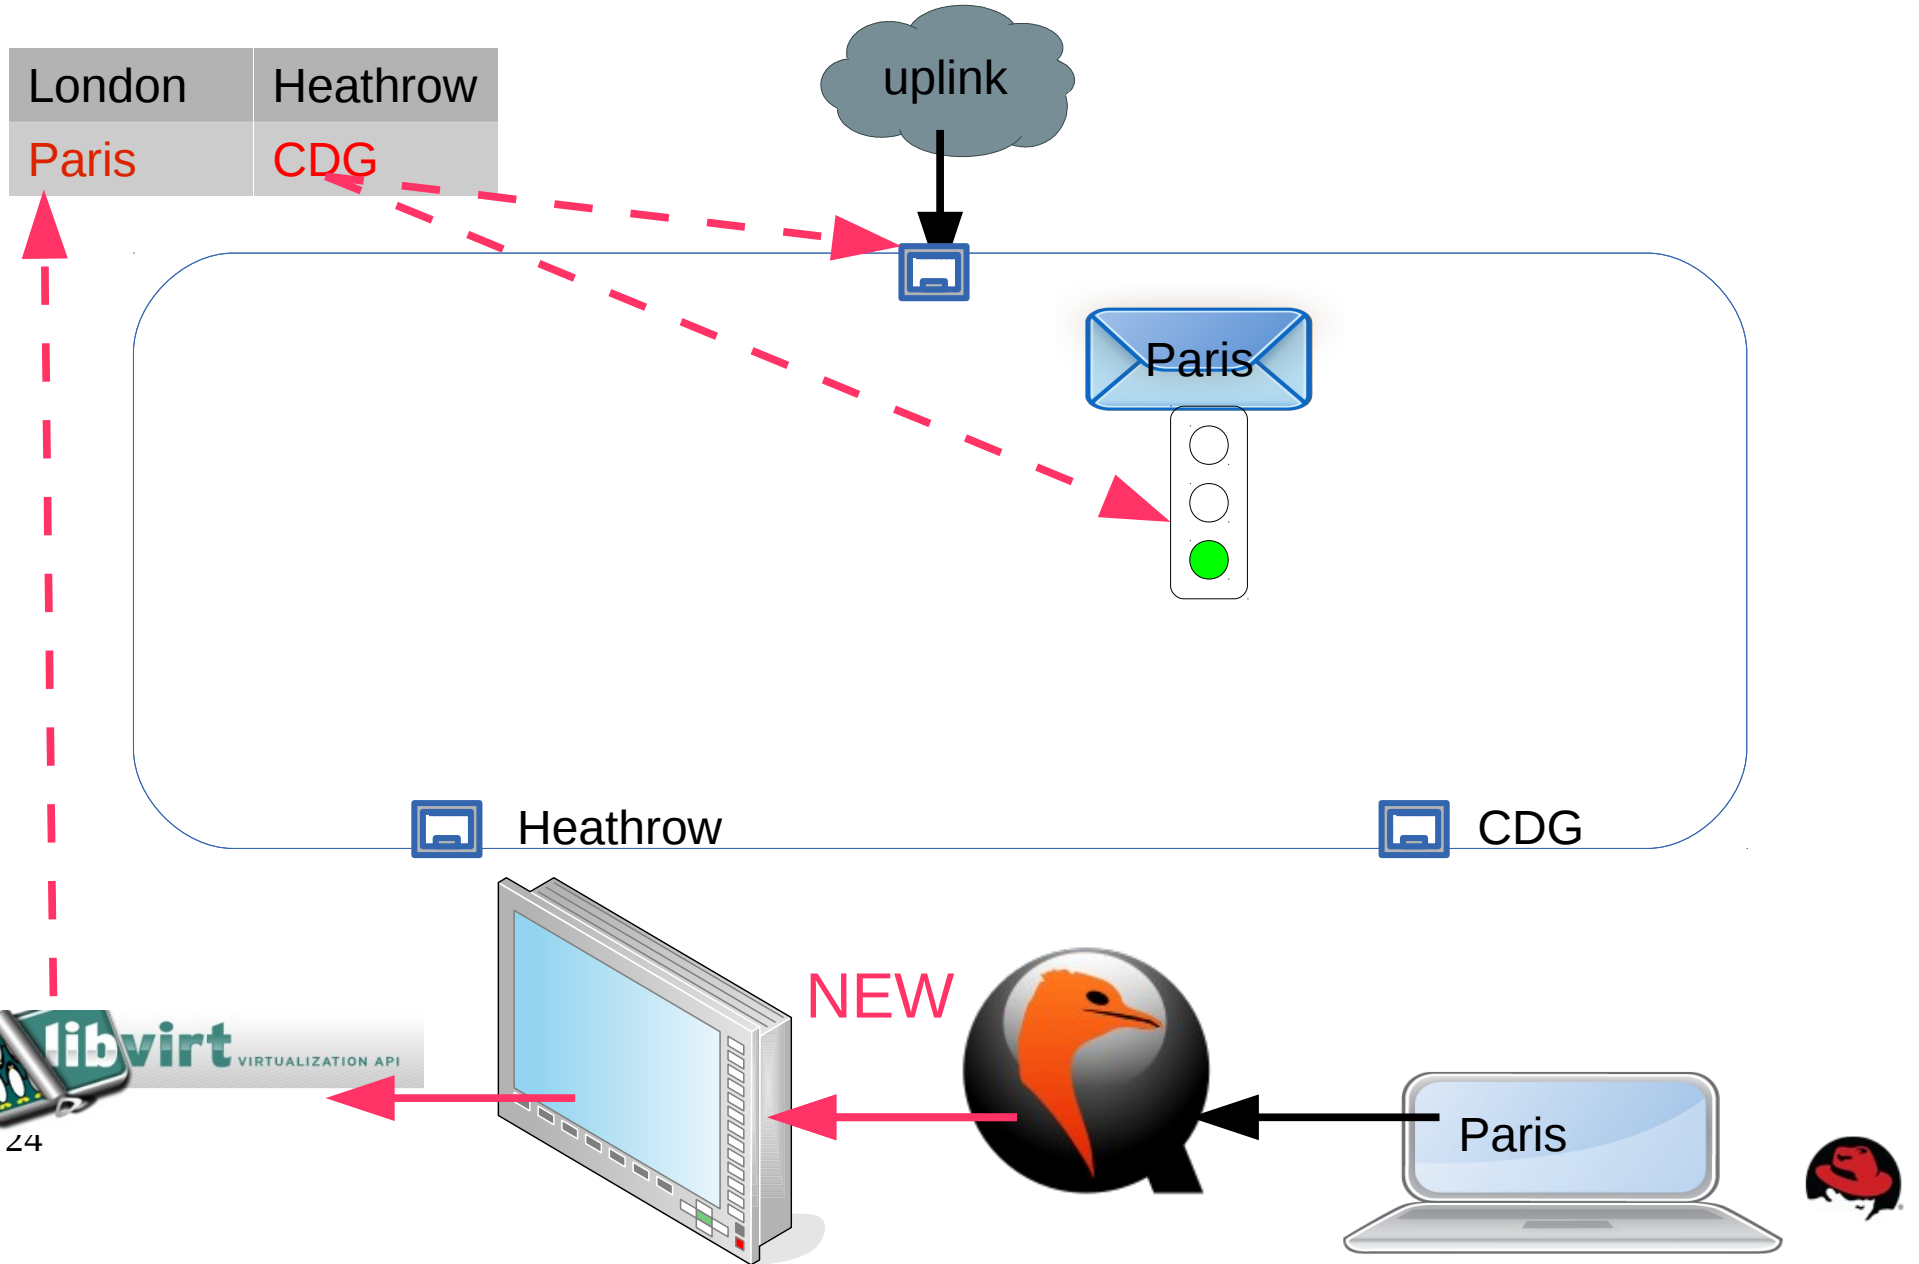

Work in progress

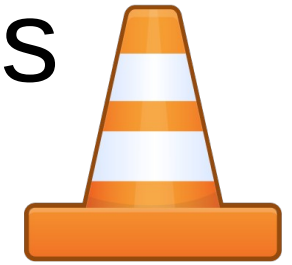

- ELVIS (vhost blk/vhost net)
- Virgl
- Vhost-net performance

RX latency

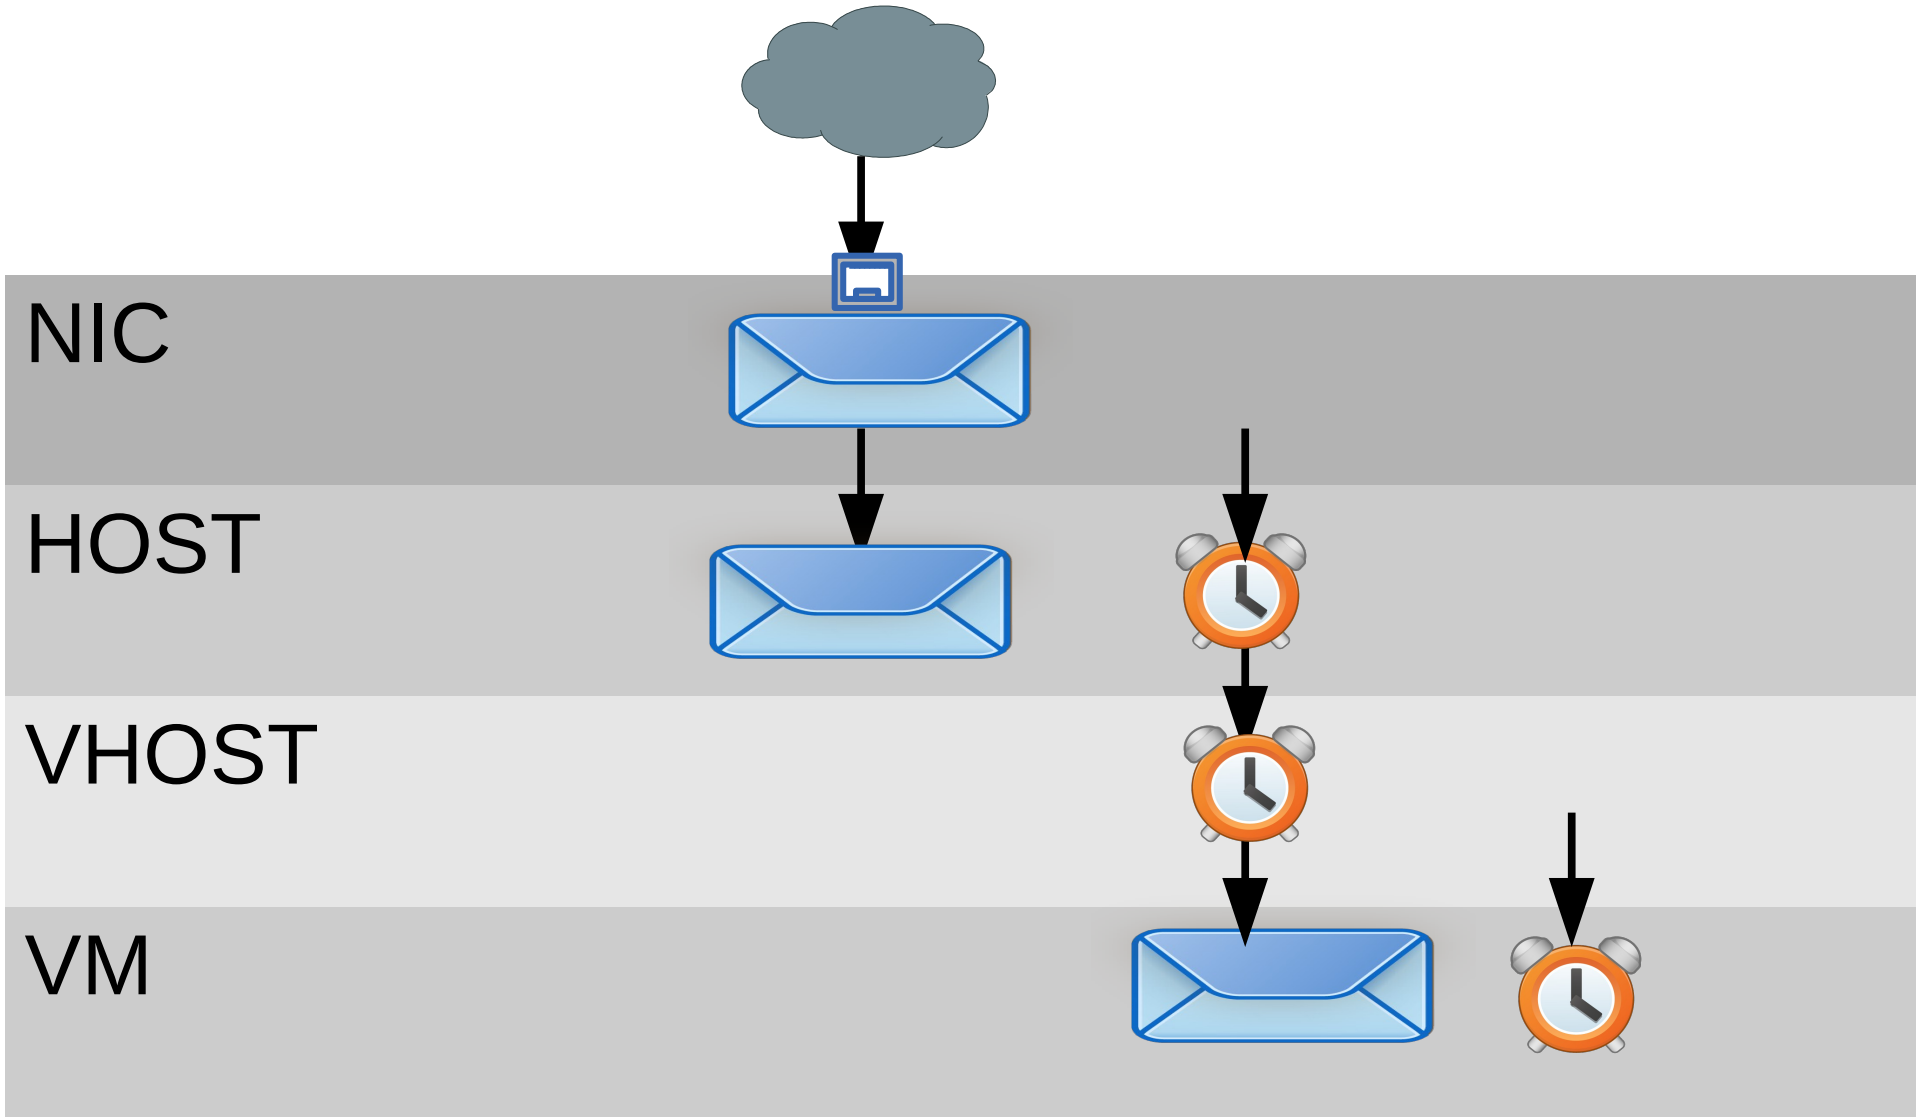

Fast rx

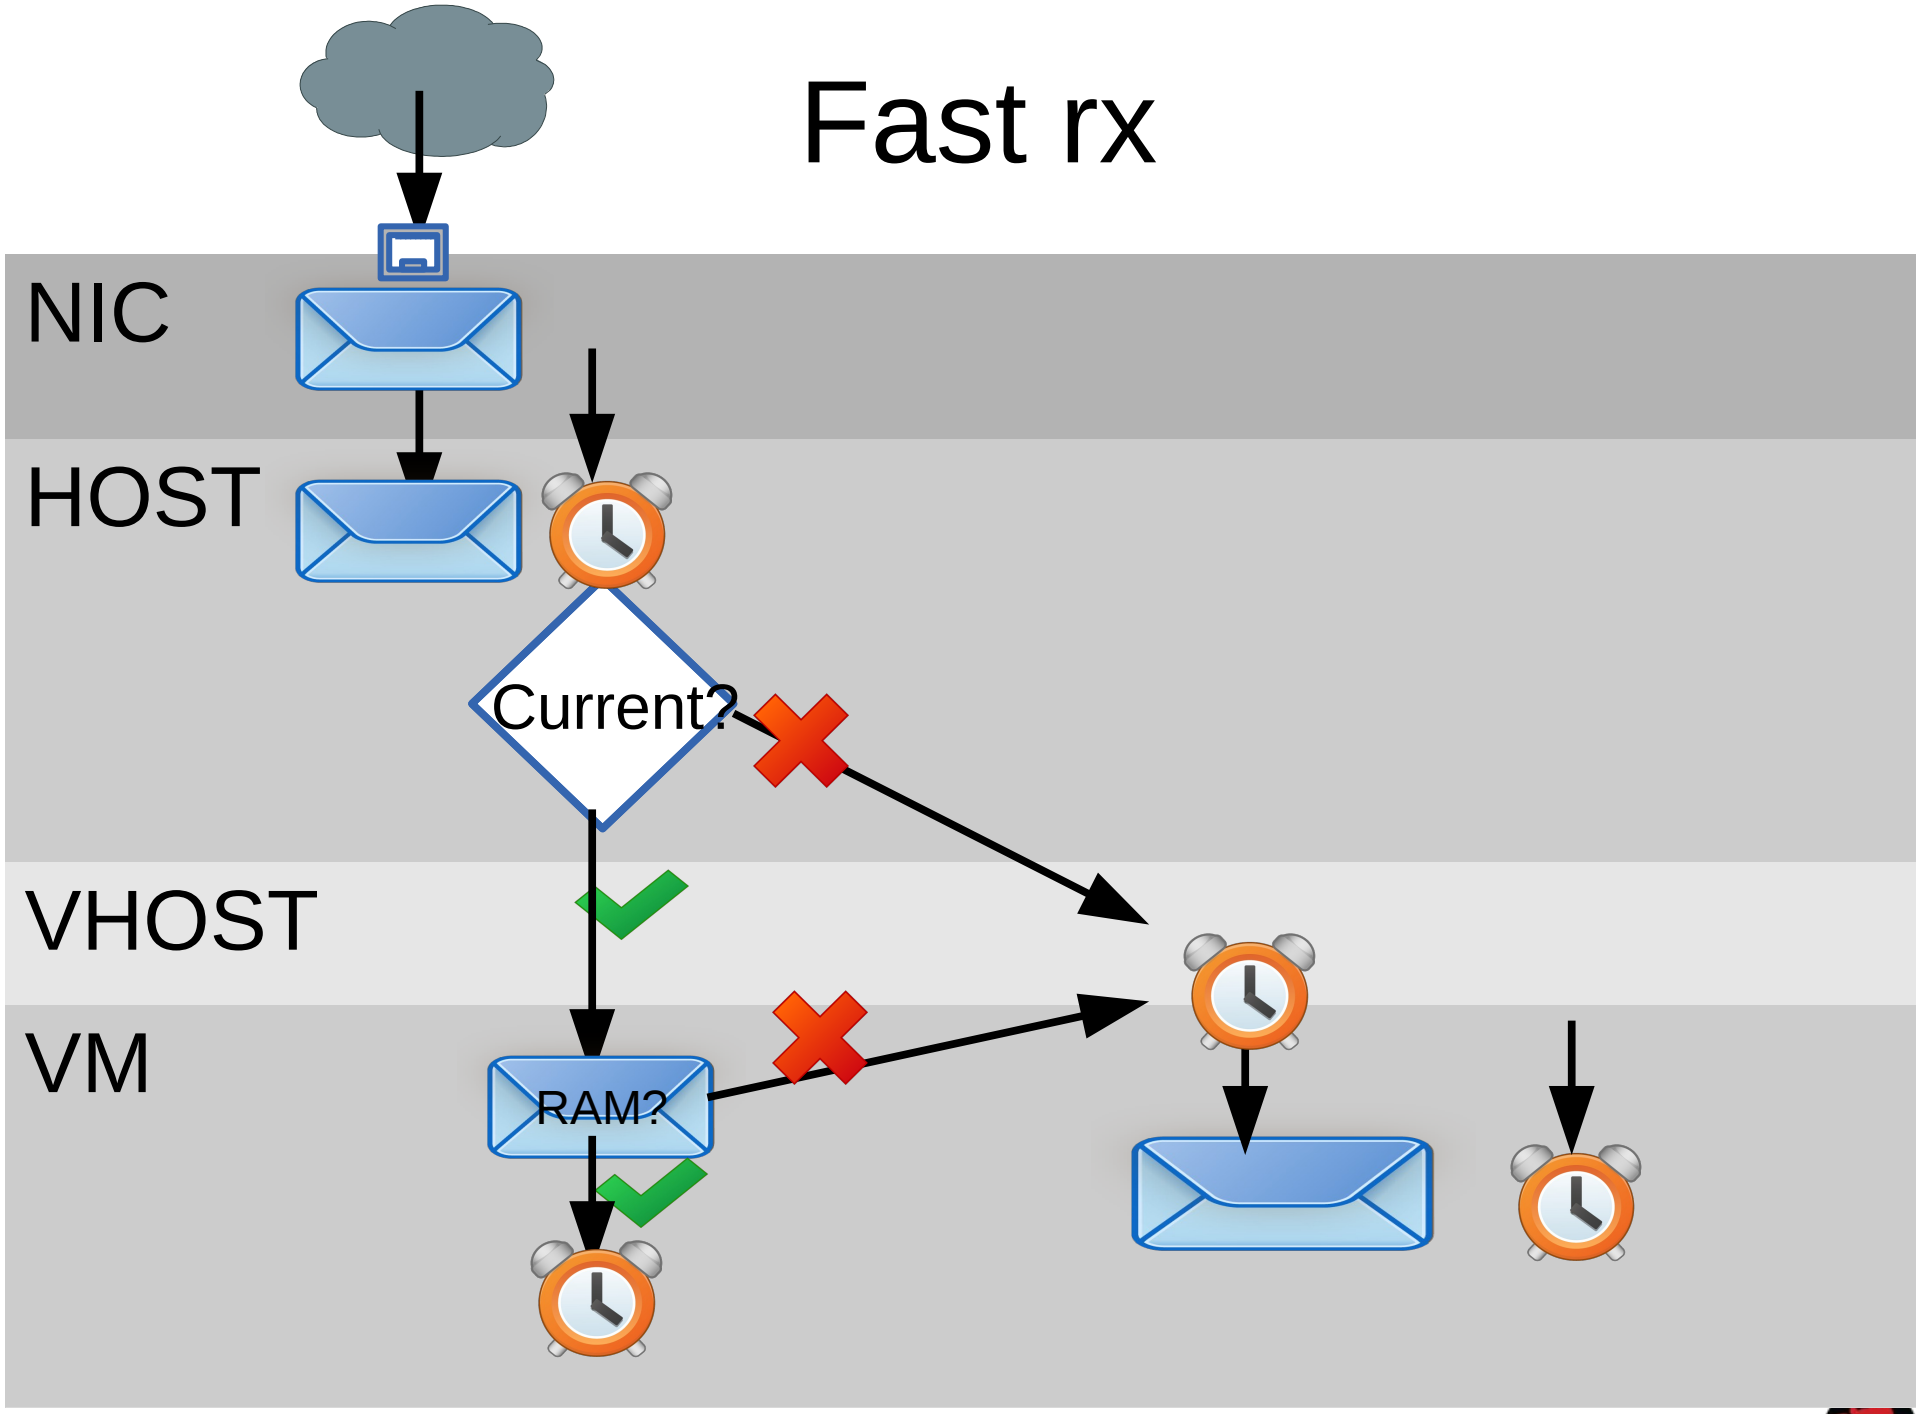

Fast rx: transactions per sec (higher is better)

Hit	331668
Miss	79

Vhost-net threading

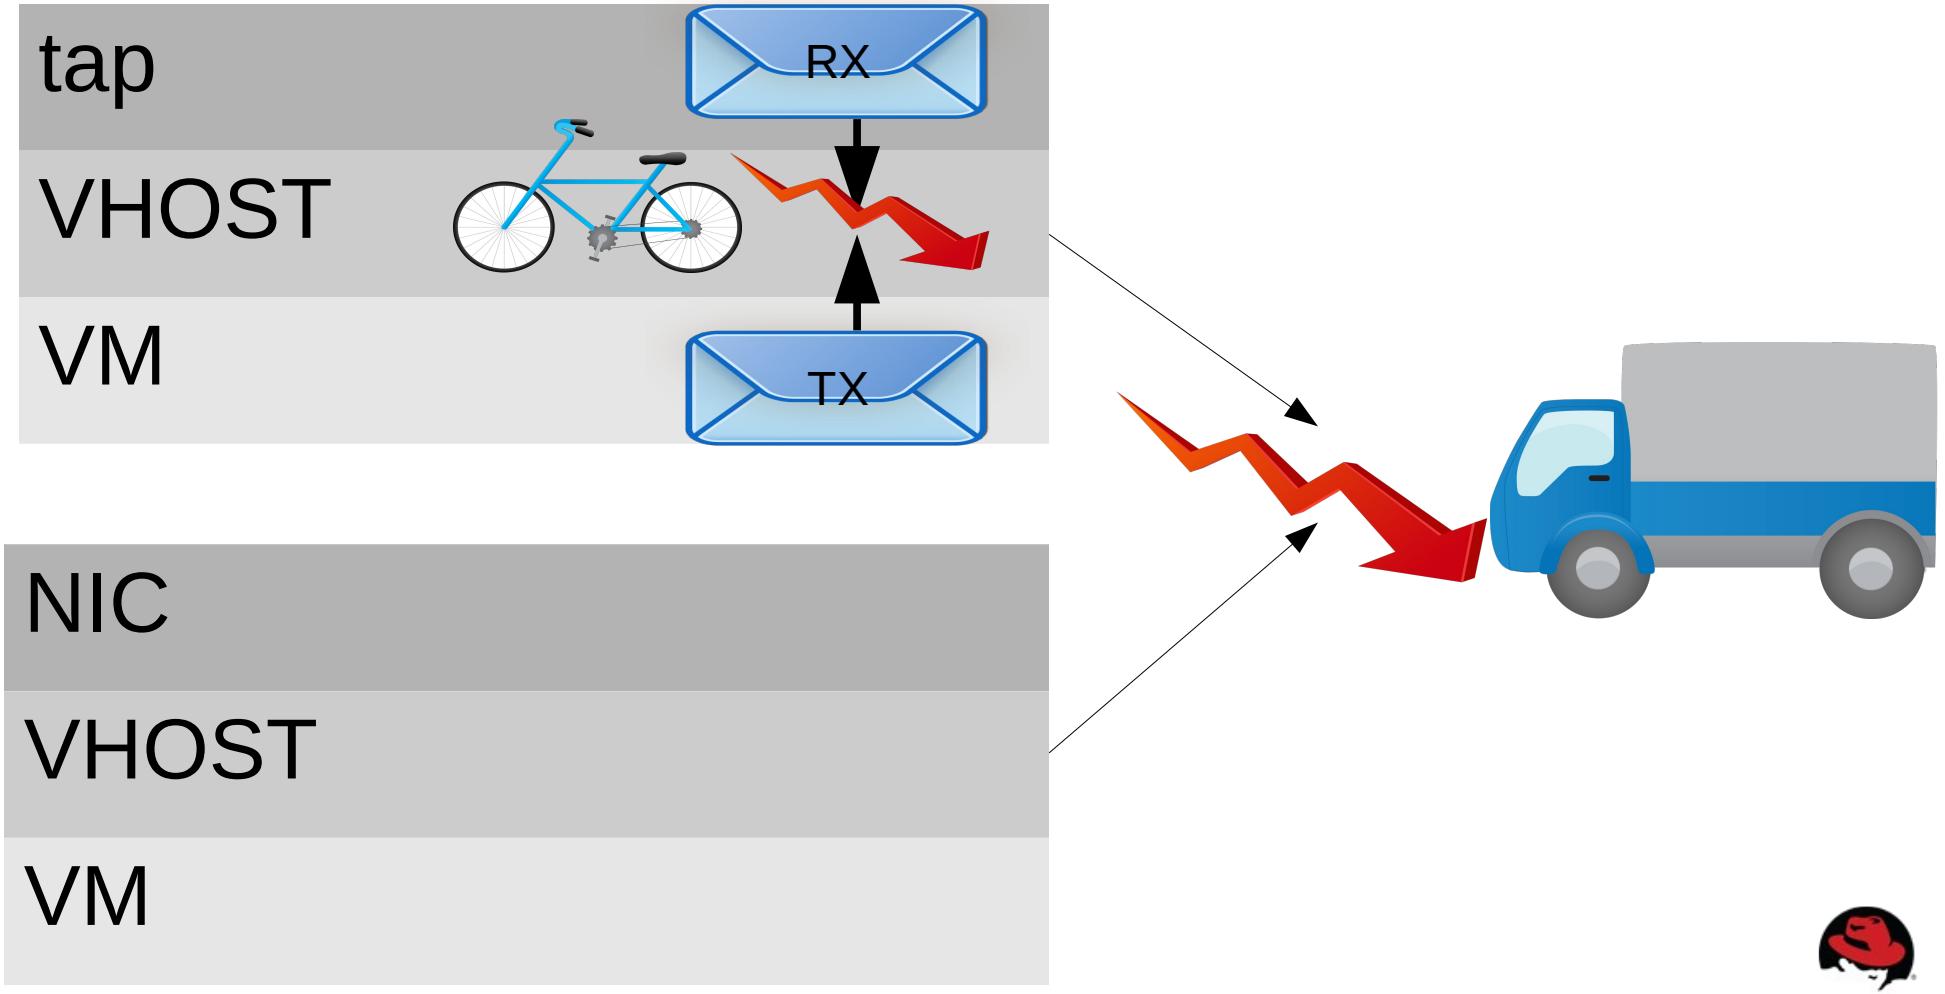

Vhost-net thread pool

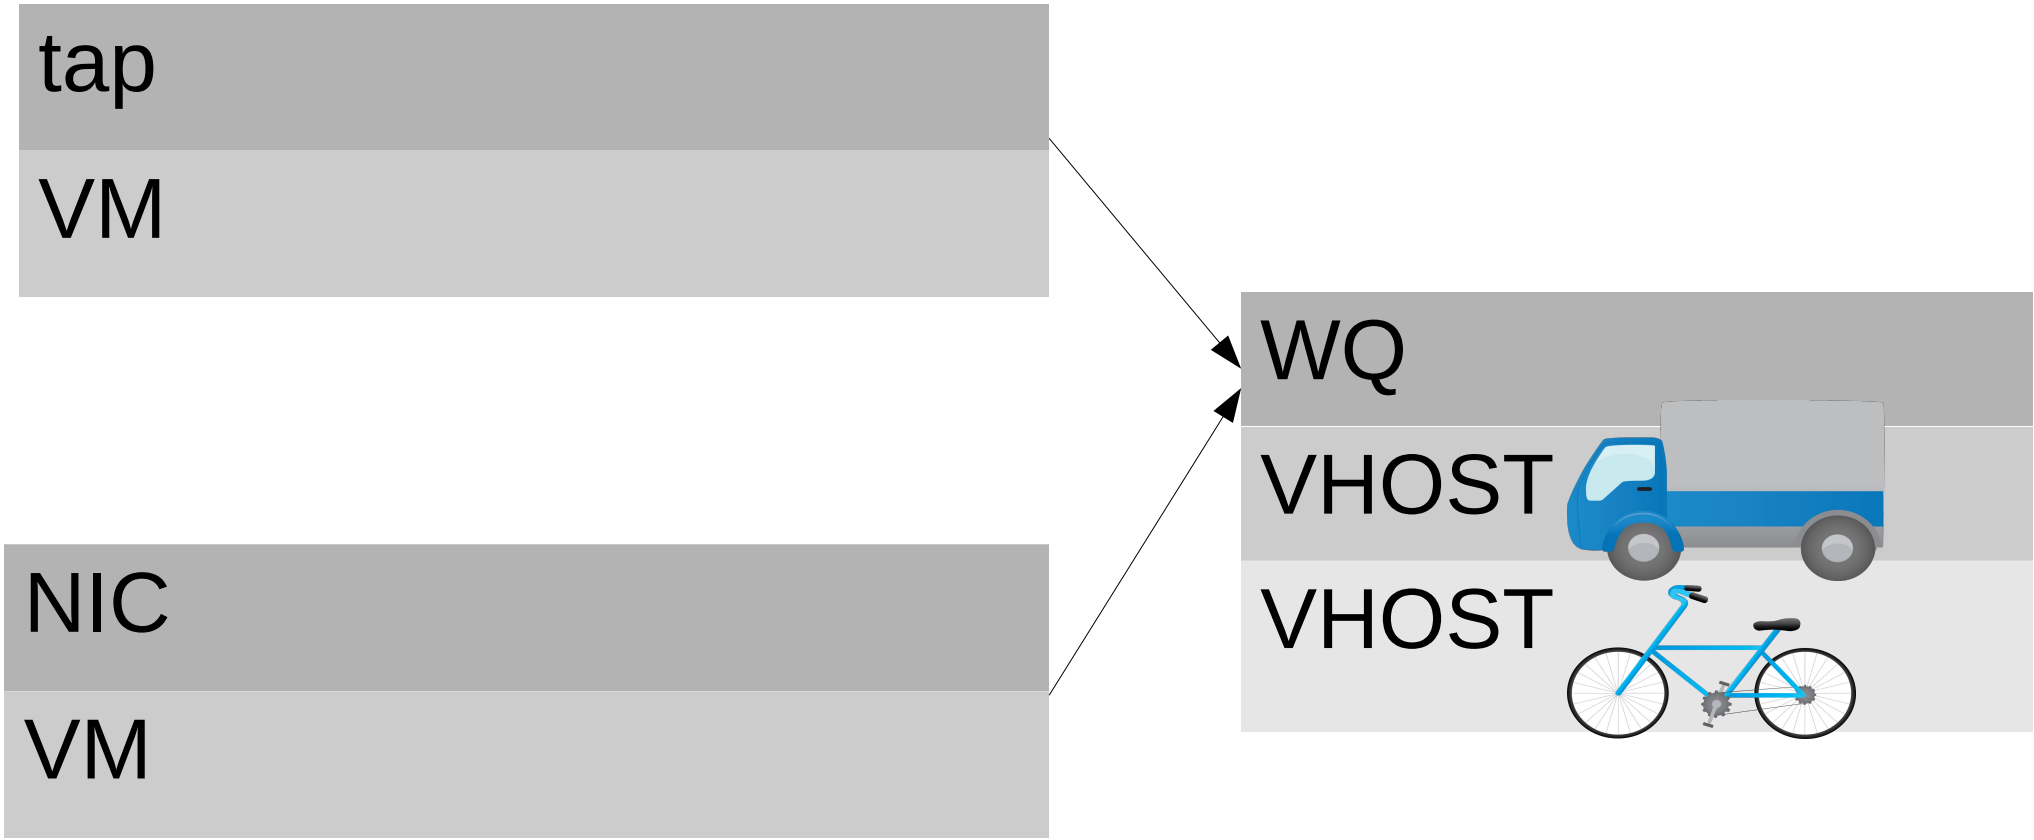

threading: UDP RR transactions/sec (higher is better)

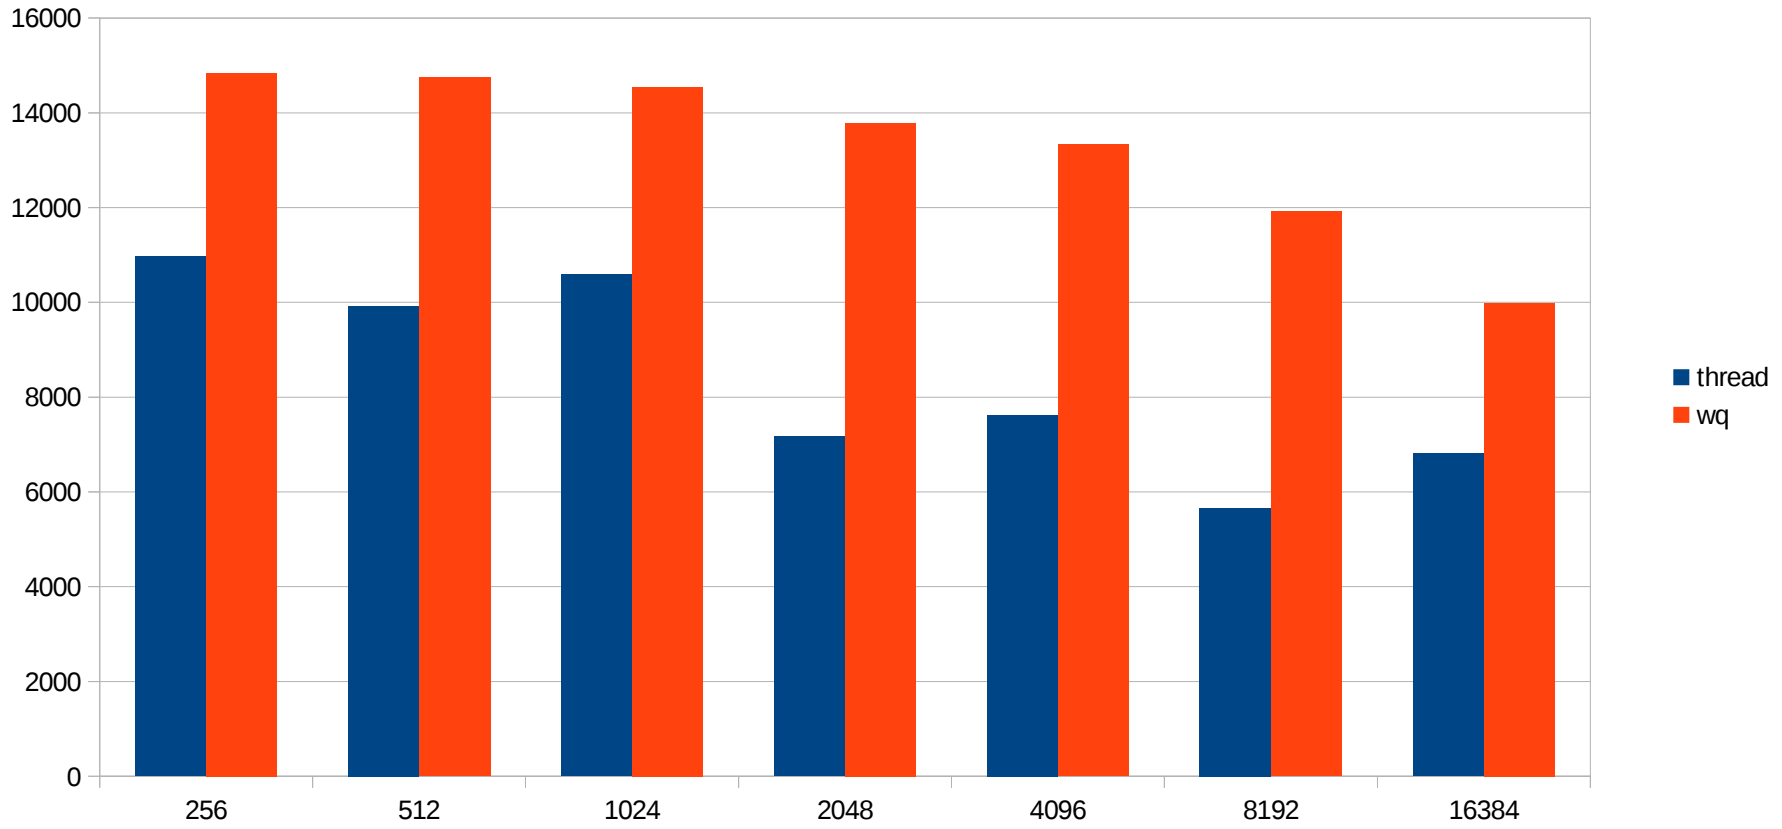

threading: TCP STREAM transactions/sec (higher is better)

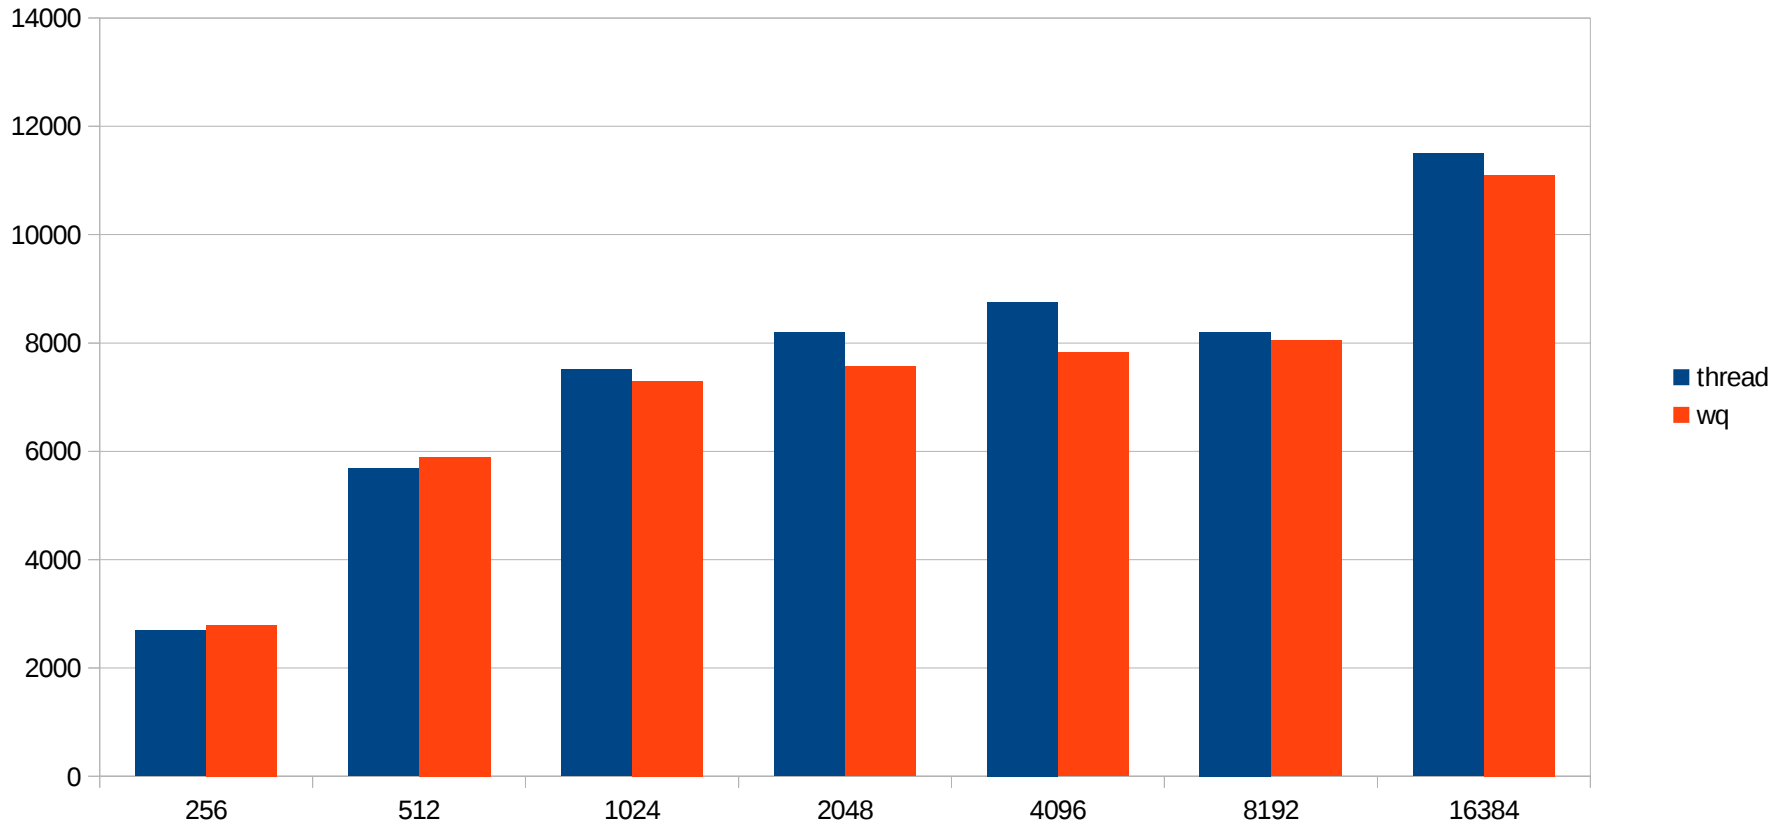

summary

- Performance
- Manageability
- Security

Questions?

OVS: flow match

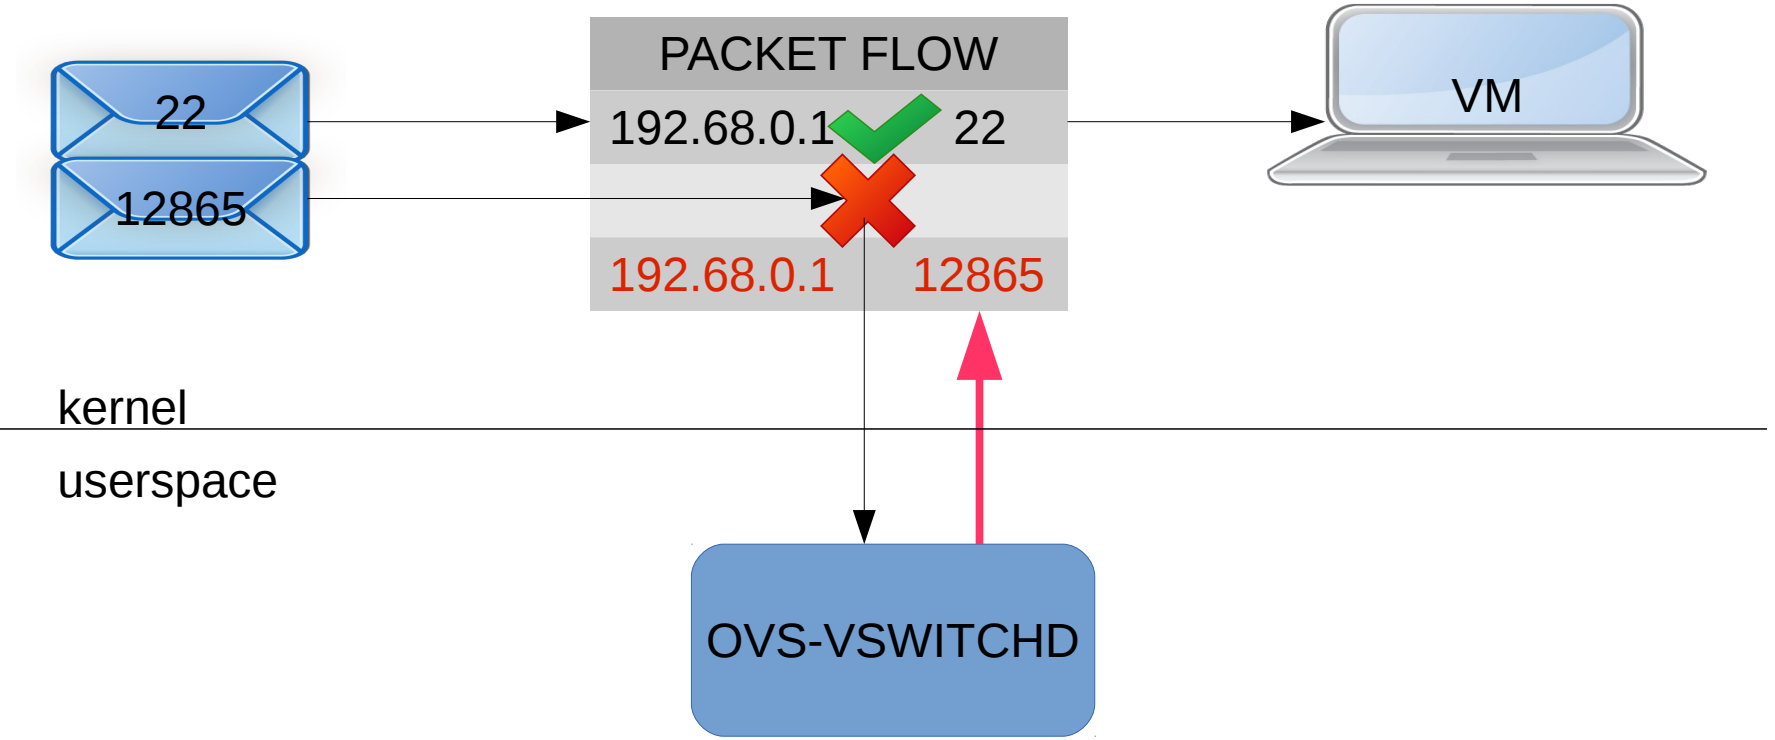

OVS: wildcard match

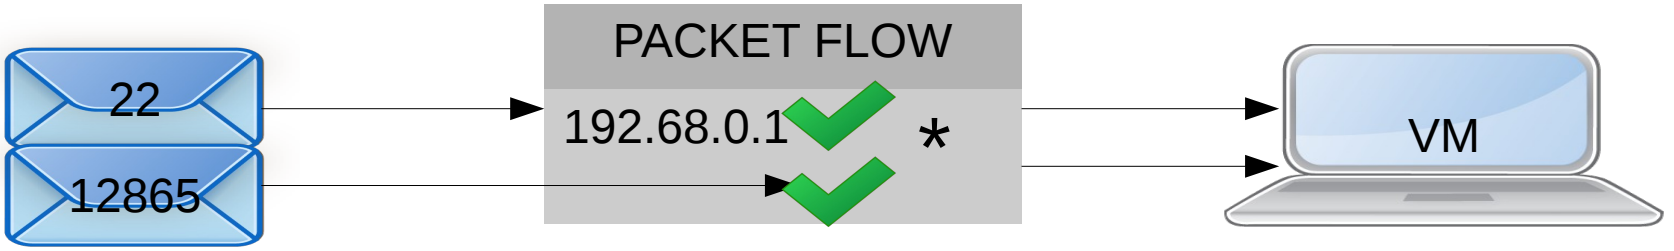

kernel

userspace

OVS-VSWITCHD

Wilcard: netperf CRR (higher is better)

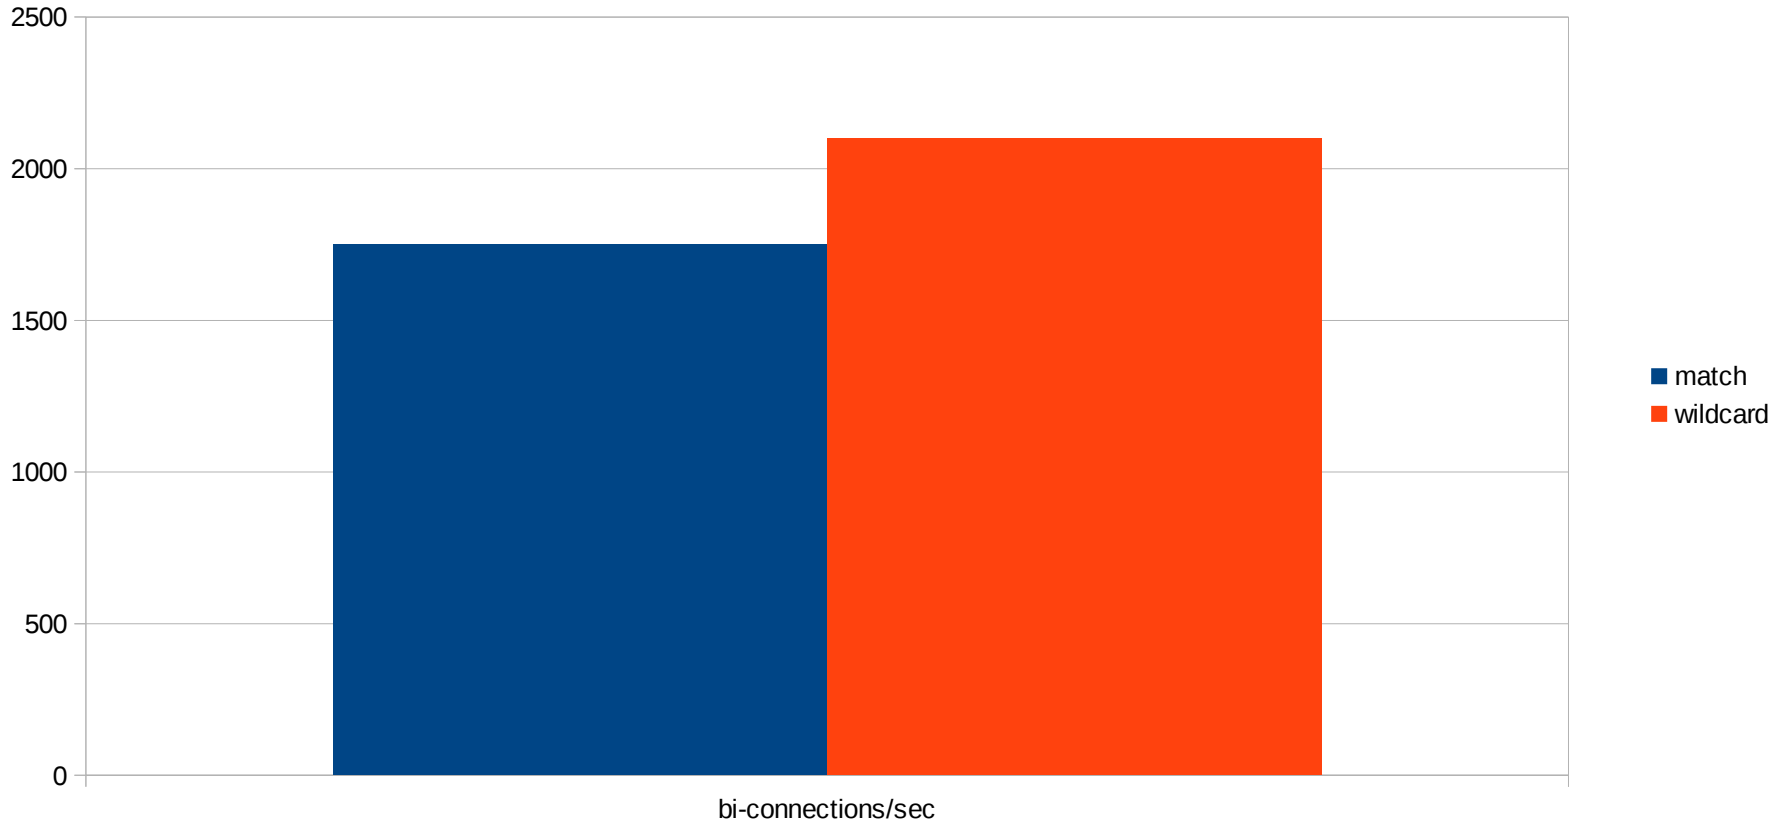


libvirt
VIRTUALIZATION API