Virtio 1 - why do it?
And - are we there yet?

2015

Michael S. Tsirkin
Red Hat

Uses material from
https://lwn.net/Kernel/LDD3/
Gcompris, tuxpaint
Distributed under the
Creative commons license.

Lots of work ...

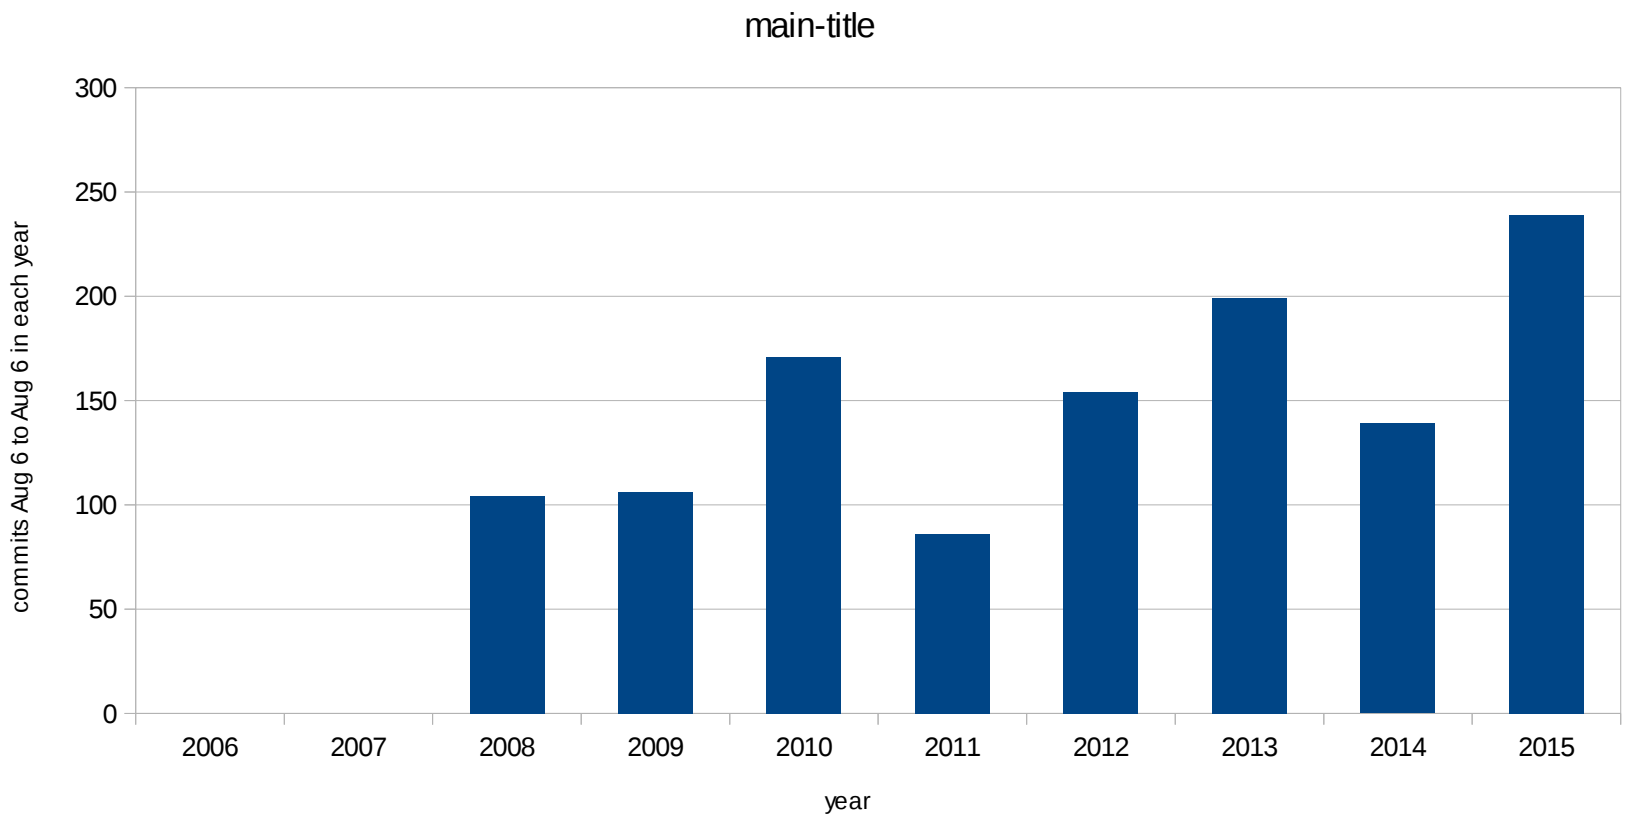
main-title
commits Aug 6 to Aug 6 in each year
300
250
200
150
100
50
0
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
year

Virtio 1: update

- Documented assumptions
- More Robust
- More Extendable

Conformance statements

Virtio 0.9	Virtio 1.0
- DRIVER_OK status bit is set. - The device can now be used. drv→probe(dev); netif_carrier_on(dev)  add_status(dev, DRIVER_OK);	The driver **MUST NOT** notify the device before setting DRIVER_OK. drv→probe(dev); add_status(dev, DRIVER_OK); netif_carrier_on(dev) 

Virtio 0.9: inflate

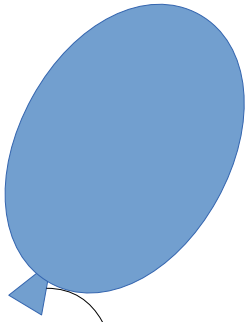

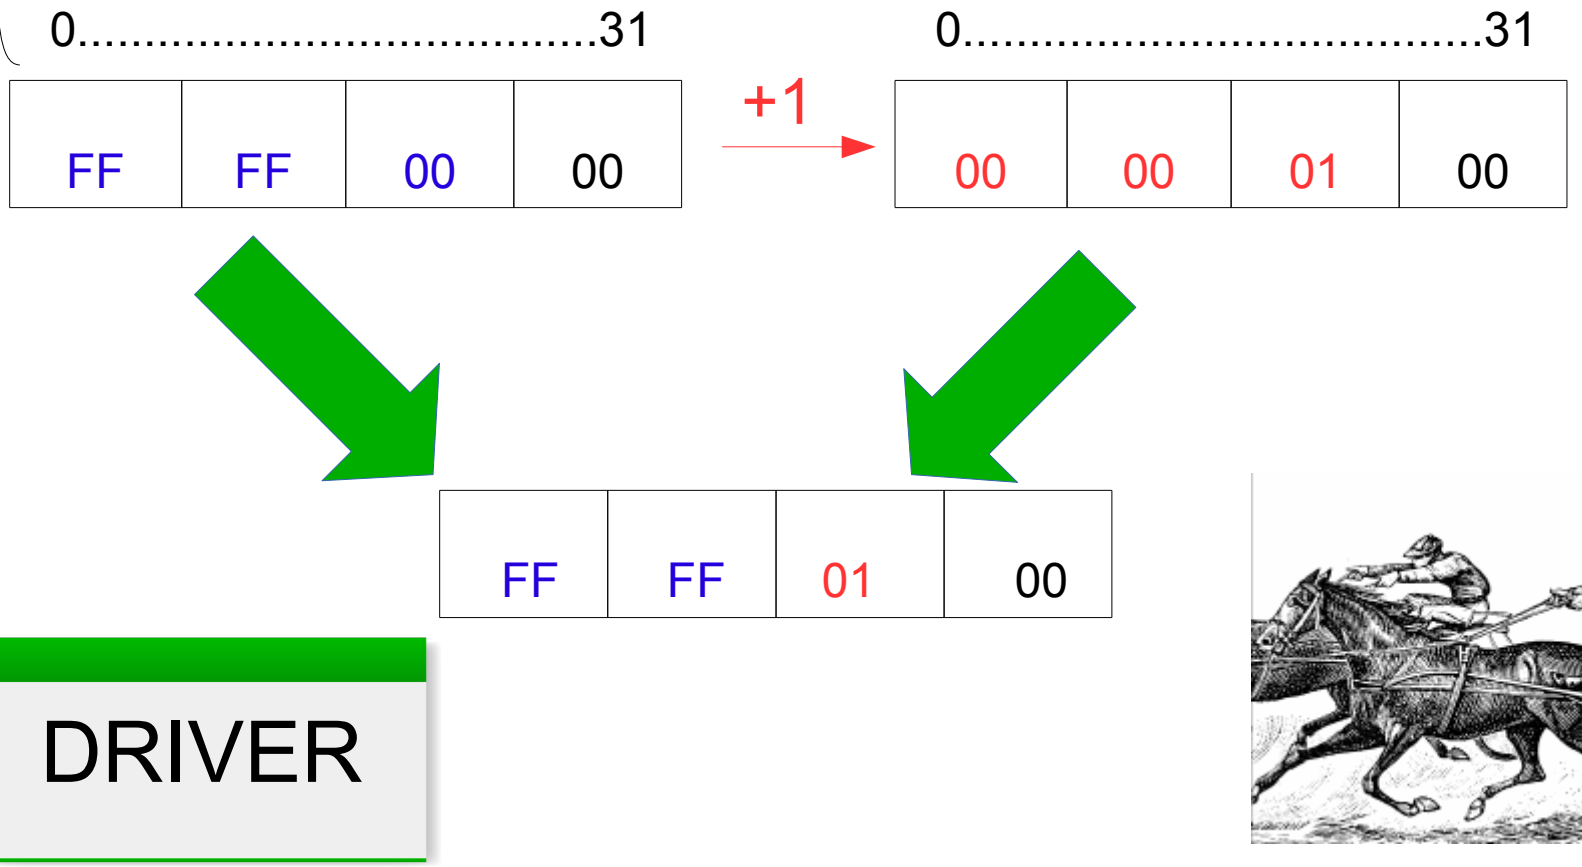

Virtio 1.0: inflate

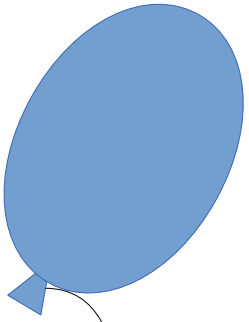

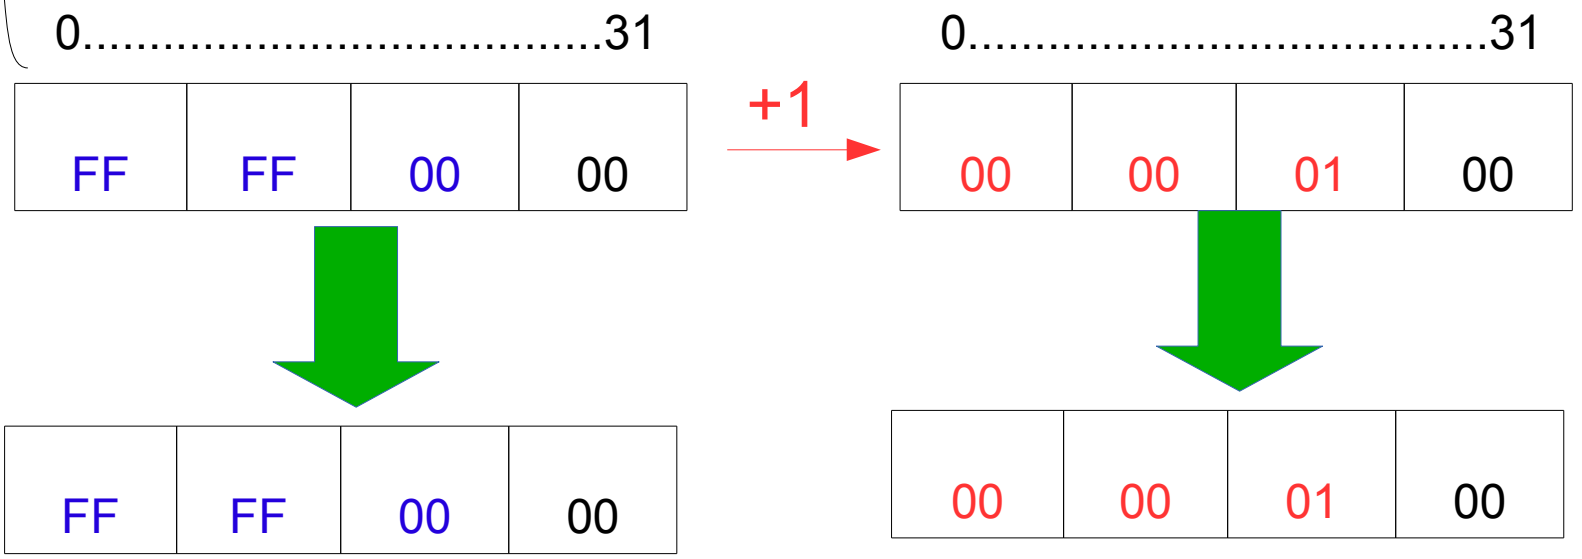

DRIVER

Generation counter

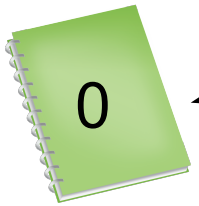

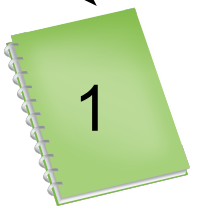

0.....63

FFFFFFFF	00000000

FFFFFFFF	00000001

+1

0.....63

00000000	00000001

00000000	00000001

Memory map

0.9

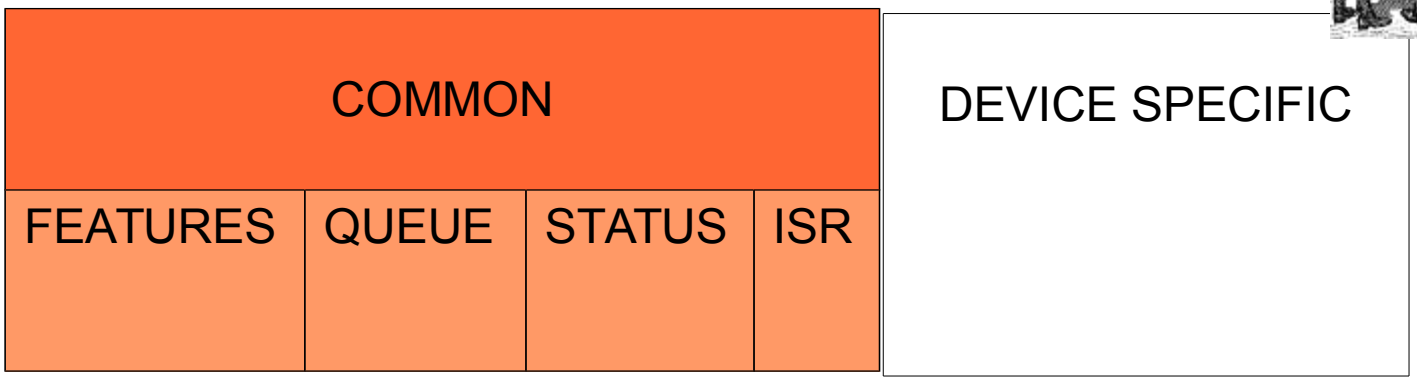

1.0

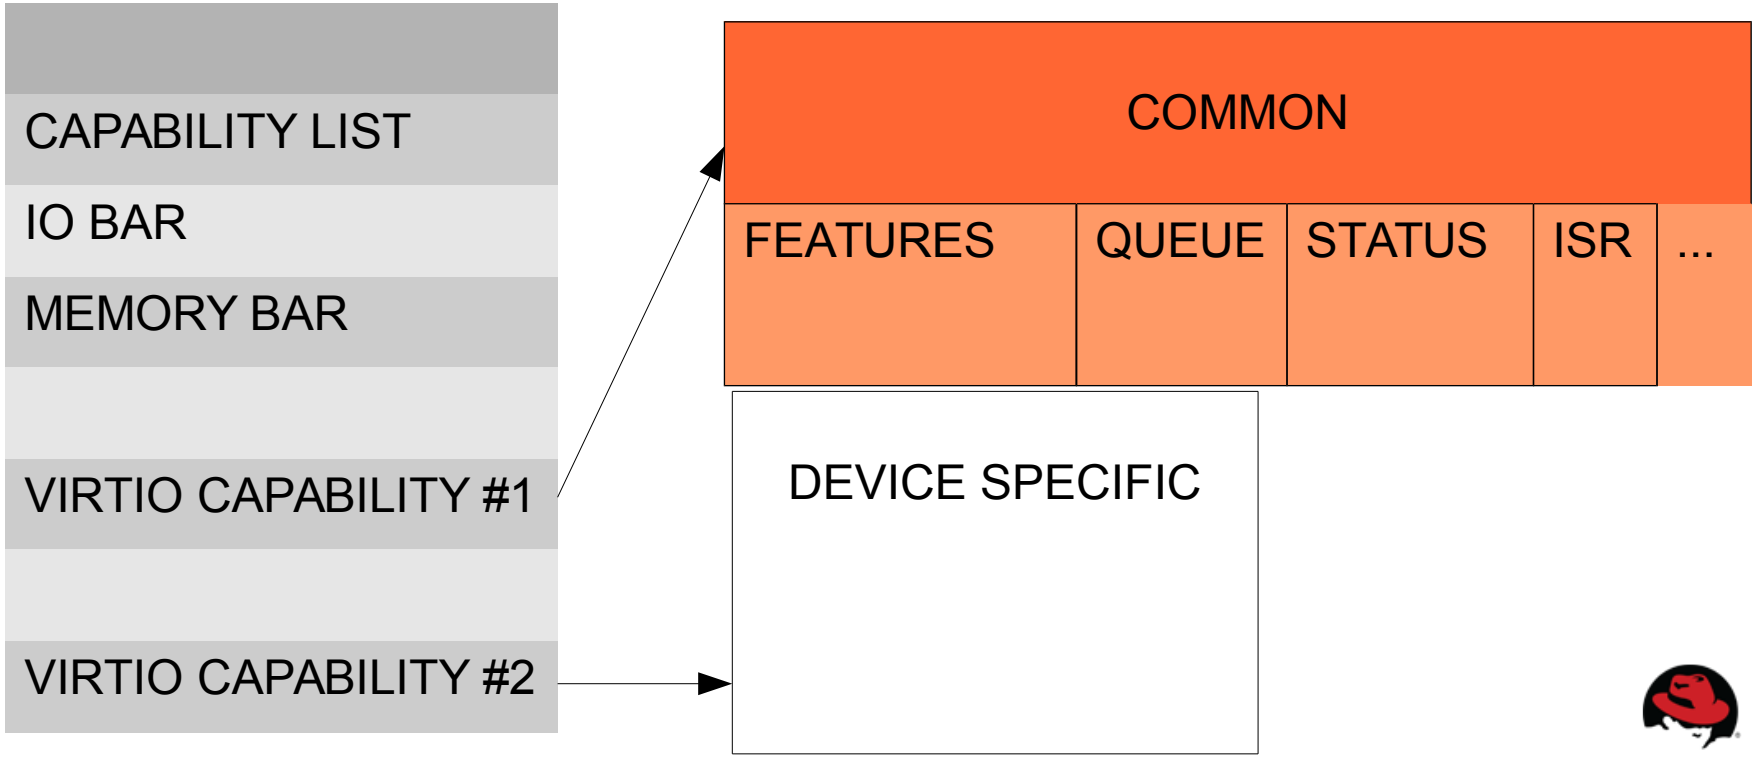

Virtio 0.9: Port IO vs Memory

	Port IO	MM IO
x86 decode: address	✓	✓
x86 decode: data	✓	✗
Fast on x86	✓	✗
32/64 bit	✗	✓
Page tables	✗	✓
Required by PCI Express	✗	✓

Fast MMIO
avoid need to decode data

0.9

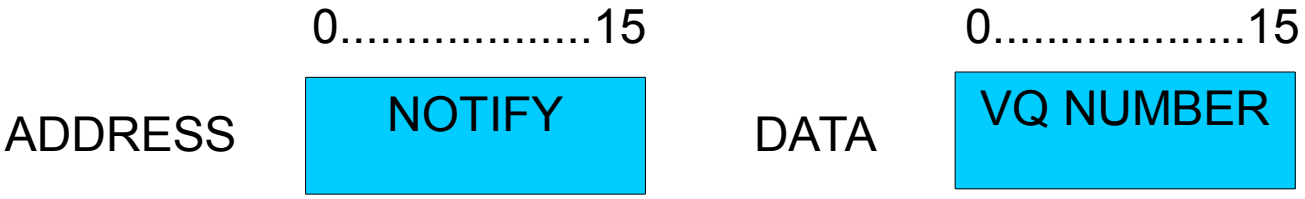

1.0

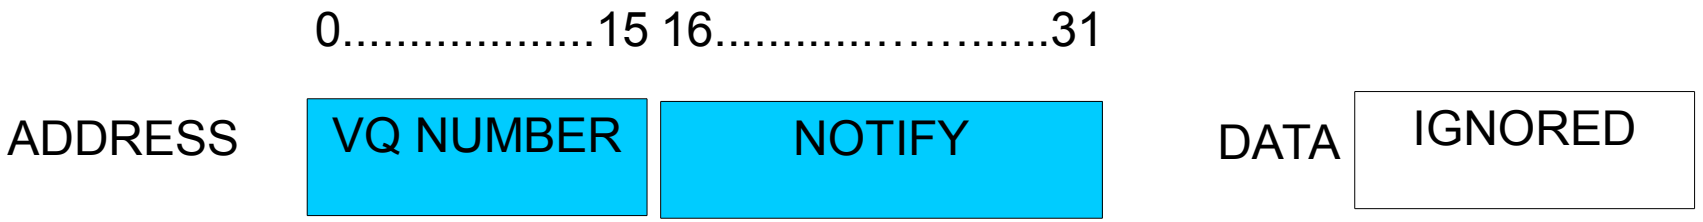

Virtio 1: Access times on KVM x86: Cycles per access (lower is better)

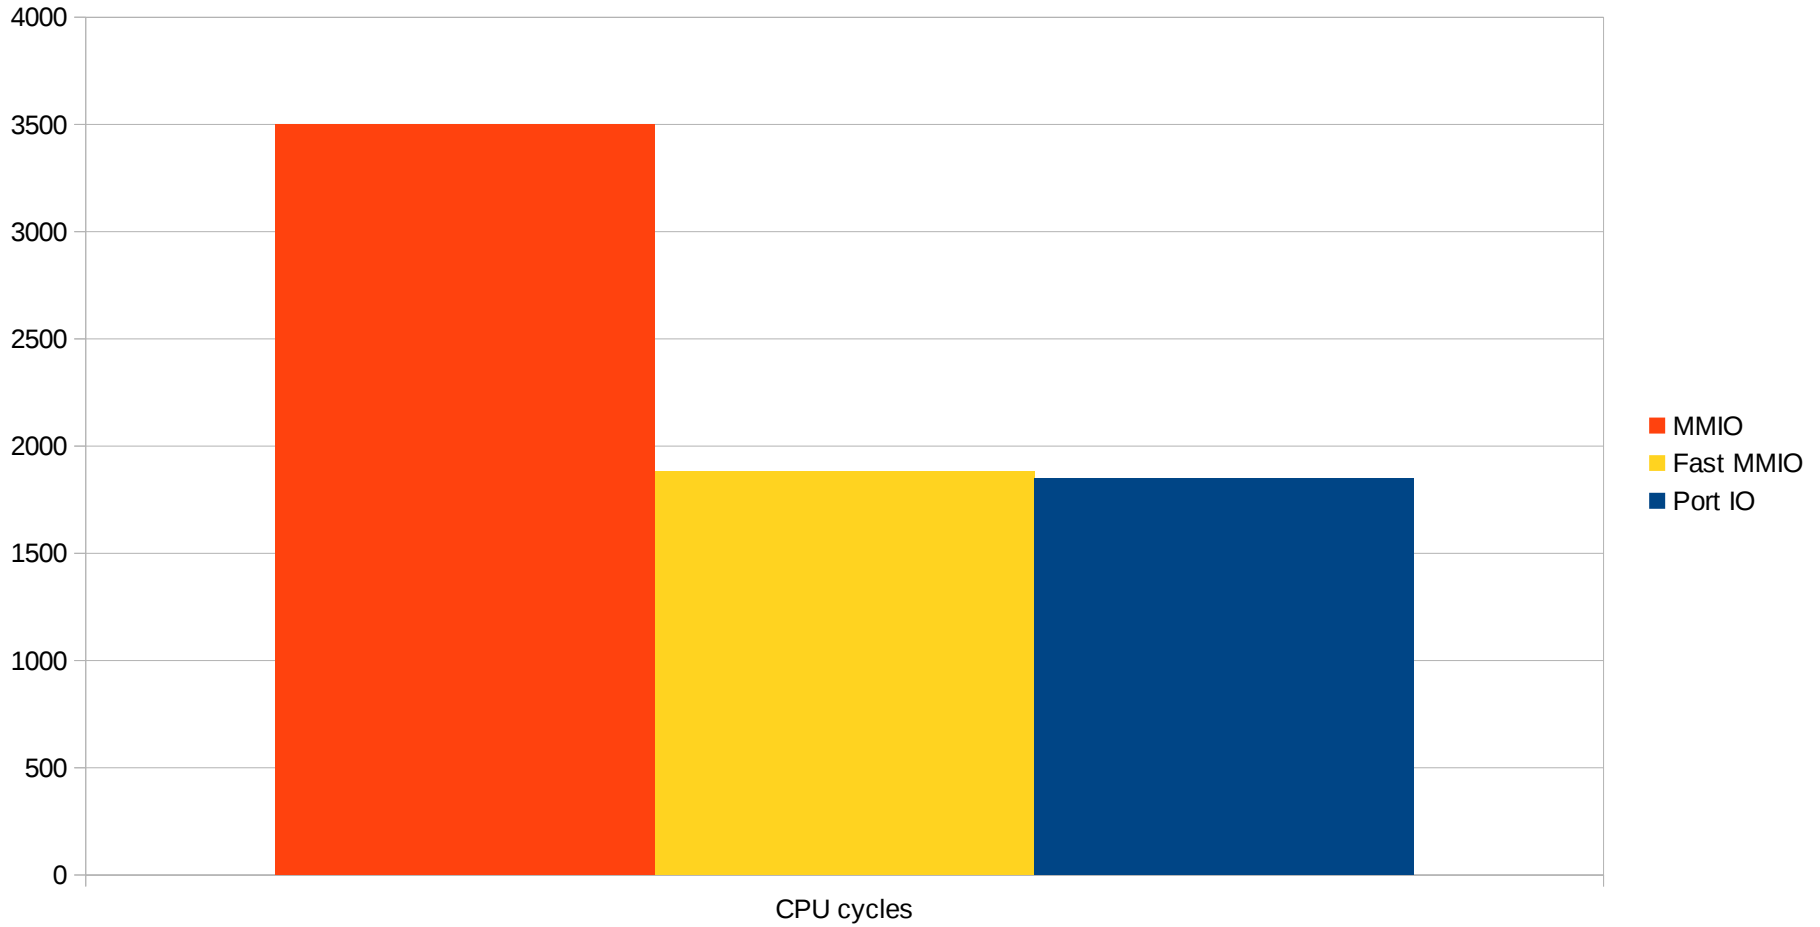

Virtio 1: Port IO vs Memory

	Port IO	MM IO
x86 decode: address	✓	✓
Fast on x86	✓	✓
32/64 bit	✗	✓
Page tables	✗	✓
Required by PCI Express	✗	✓

Memory Region Aliases

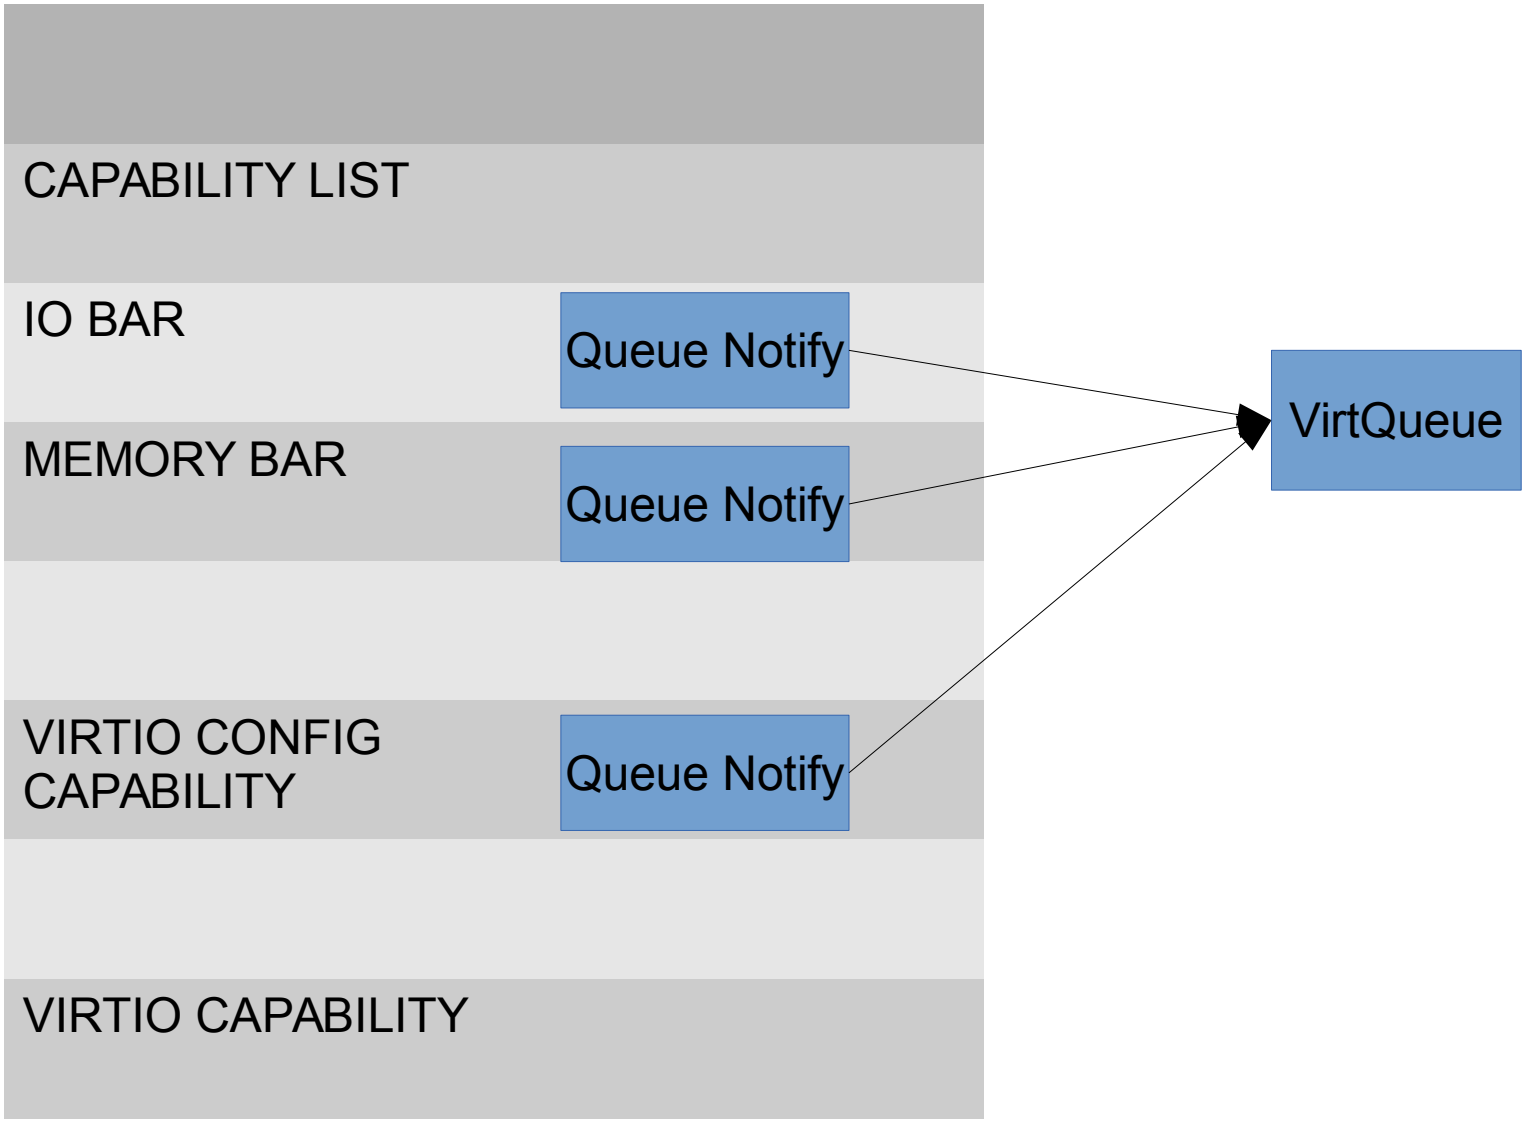

soft mac

Ethernet MAC

52	54	00	12	34	56

DRIVER

0.9

1.0

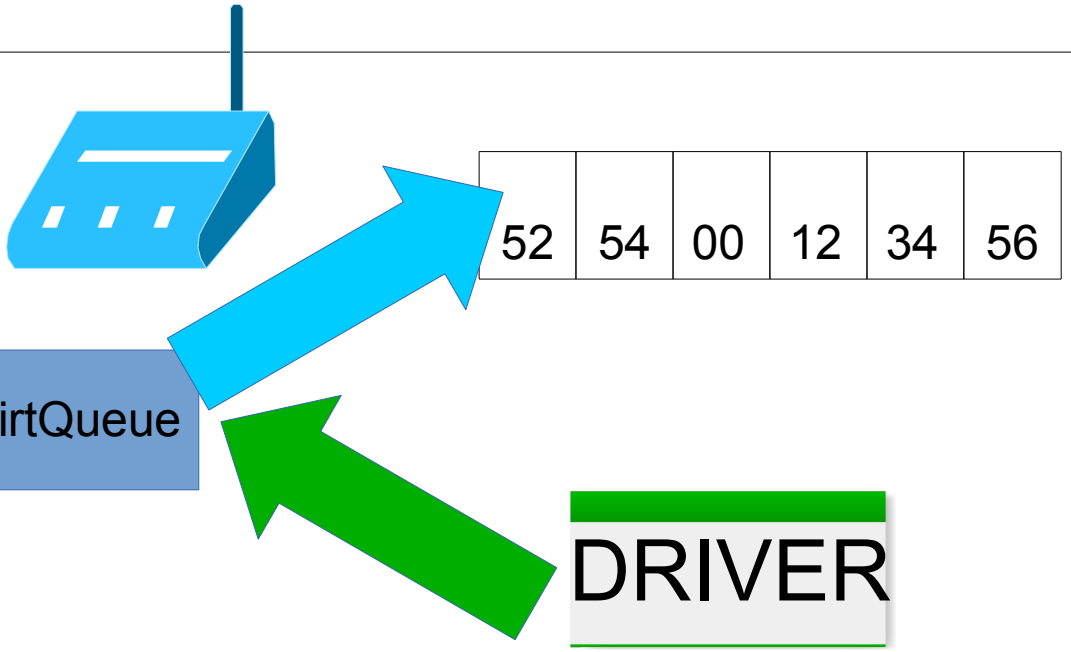

Virtio feature negotiation

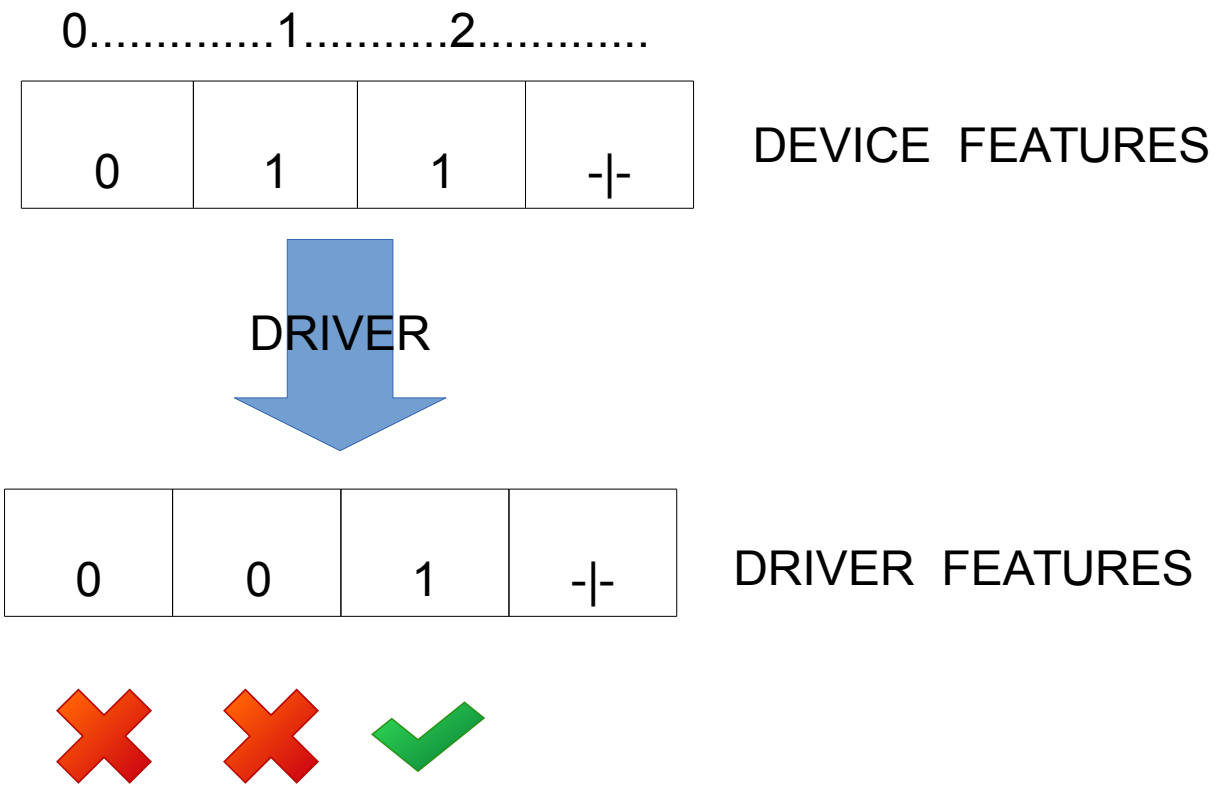

Defaults must be maintained forever!

Virtio 1: Error handling

- DRIVER: set features
- DRIVER: set FEATURES_OK bit

- DEVICE: check features
- DEVICE: clear FEATURES_OK on error

- DRIVER: check FEATURES_OK bit
- DRIVER: fail gracefully if not set

Error handling: Virtio 0.9

- Can't recover from device errors
- Not very useful?
- Just stop guest.

Vhost-user

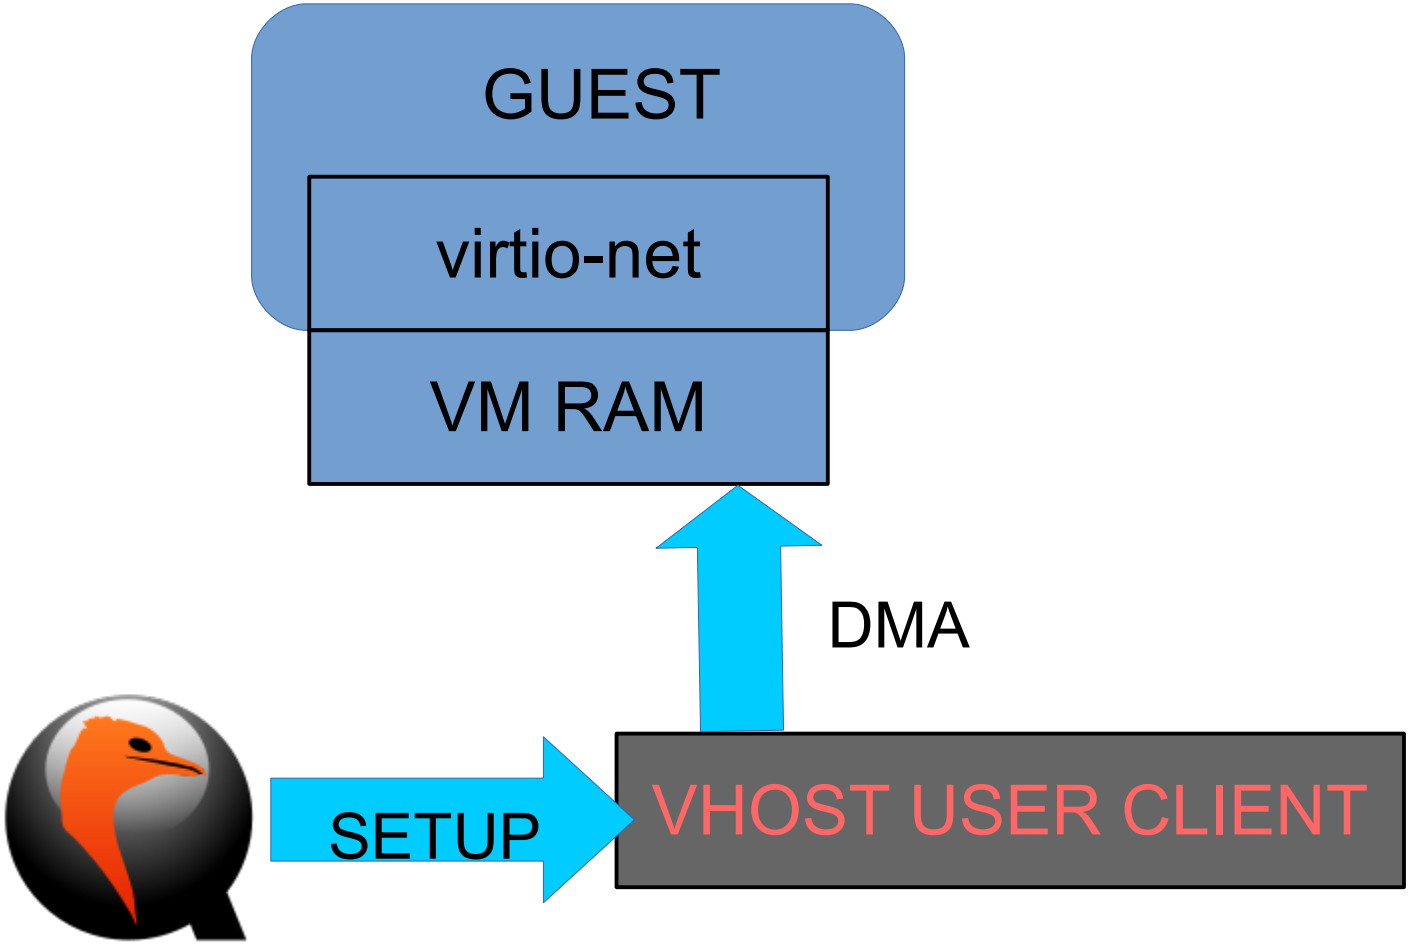

Client crash or restart need not cause guest crash!

DEVICE_NEEDS_RESET

Read STATUS;

Detect:
NEEDS_RESET set

Write
STATUS=0
Will reset device

Reconfigure device.
Write
STATUS=DRIVER_OK
Restart operation.

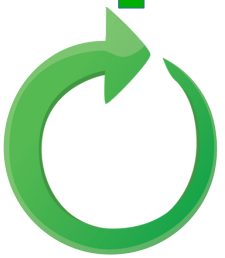

DRIVER

Compatibility

Transitional Device & Driver

Legacy Modern

Legacy Modern

DRIVER

Legacy Driver

Legacy Modern

Legacy

DRIVER

Legacy Device

Legacy

Legacy Modern

DRIVER

Are we there yet?

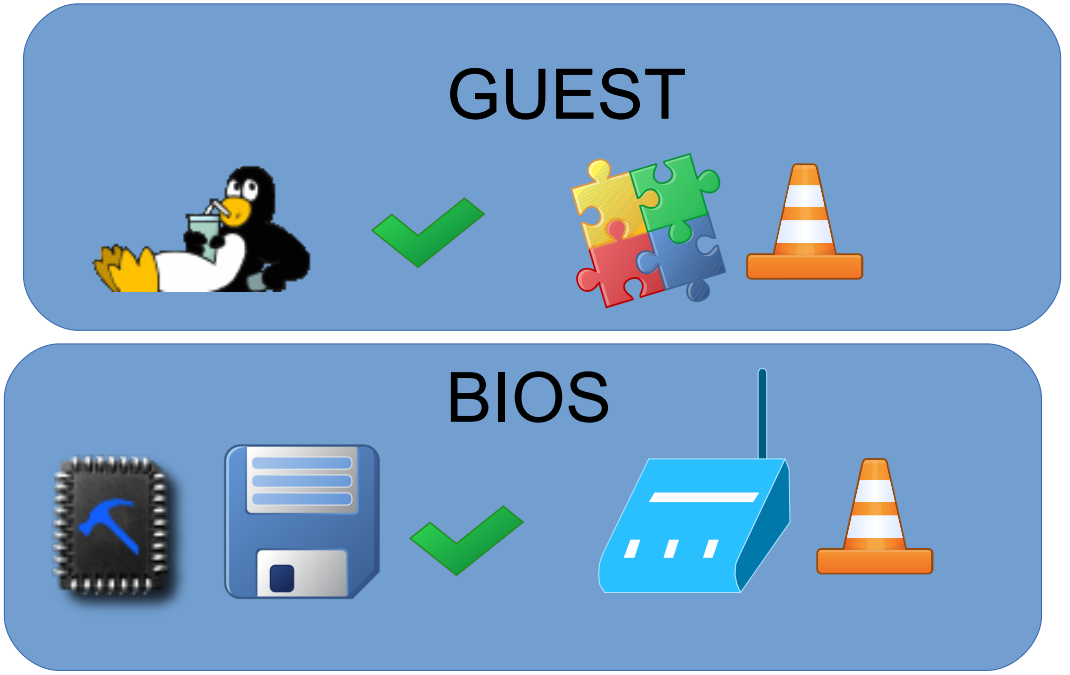

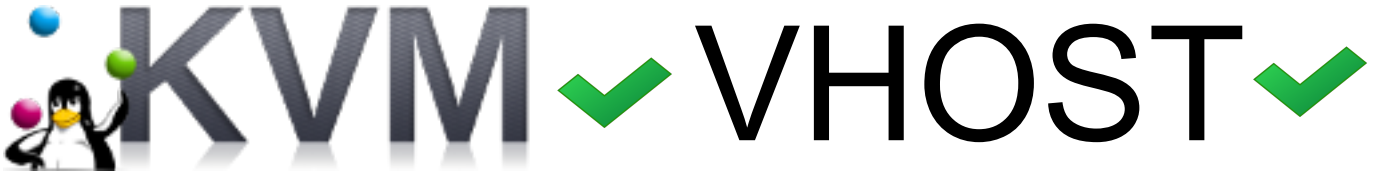

What to expect?

- Current: Virtio-v1.0-cs03
- Next bugfix: Virtio-v1.0-cs04
  - Virtio-blk: writeback / writethrough control
  - More update guidance
- Next feature: Virtio-v1.1-cs01
  - Virtio-input
  - Virtio-gpu
  - Virtio-vsock

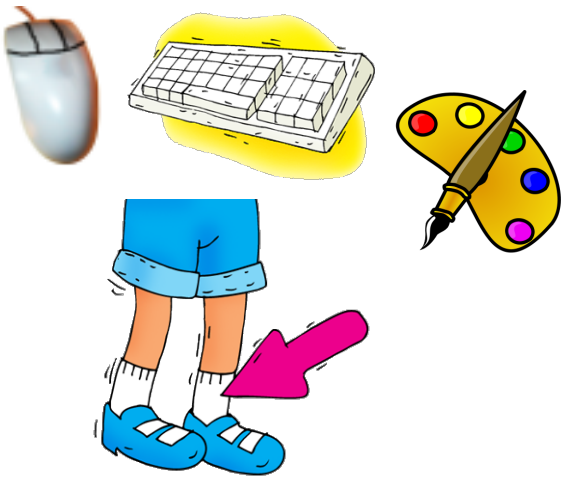

TX: Interrupt avoidance

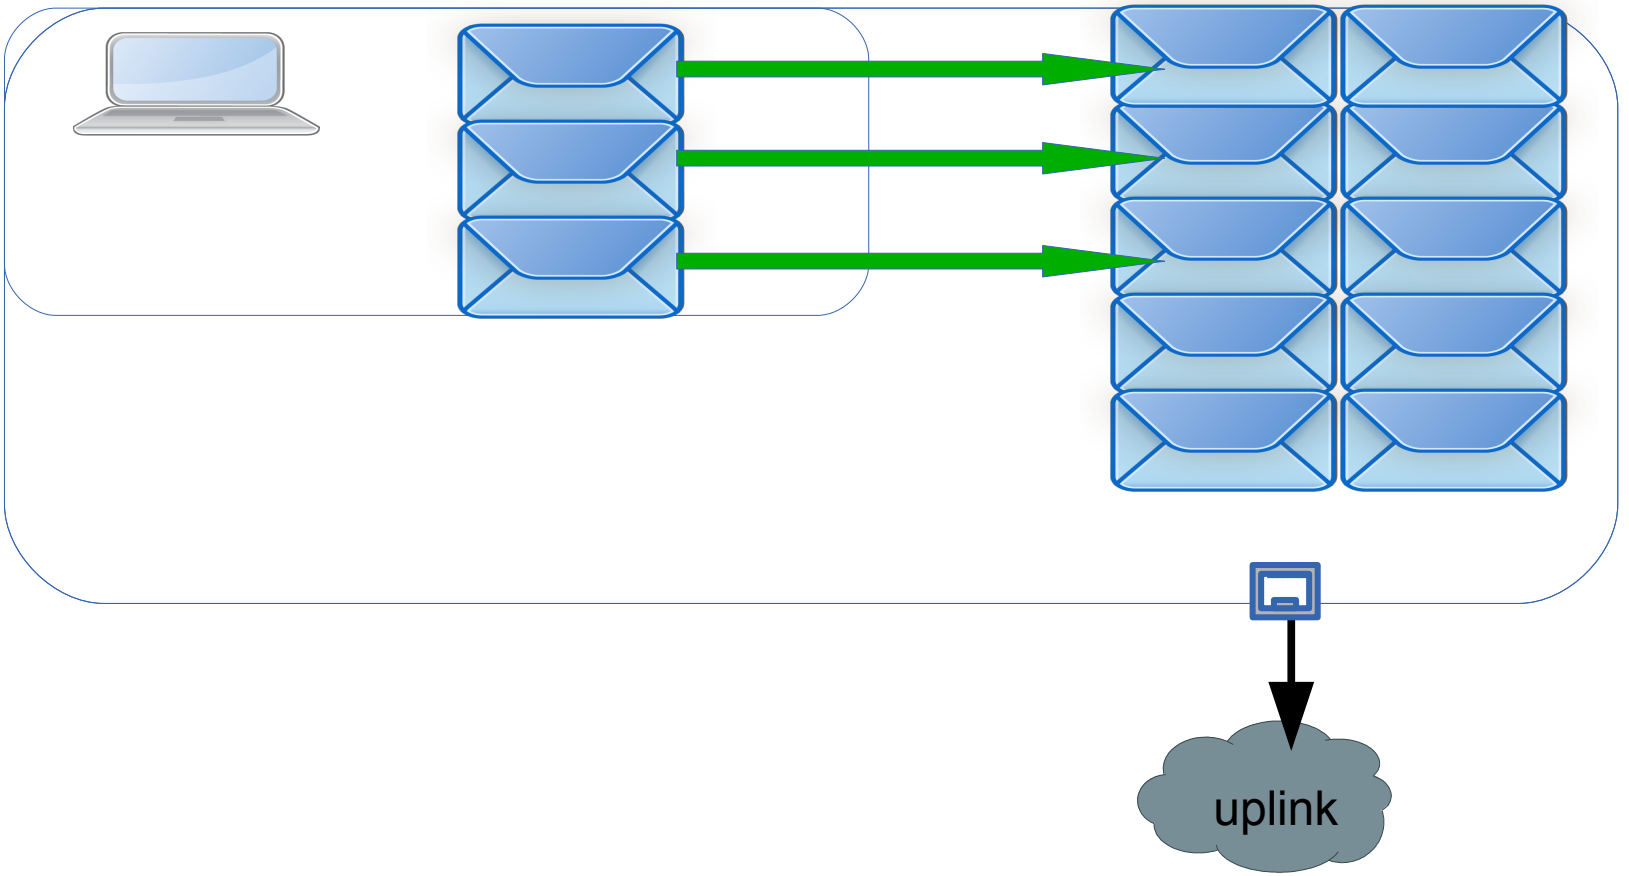

TX: Interrupt coalescing

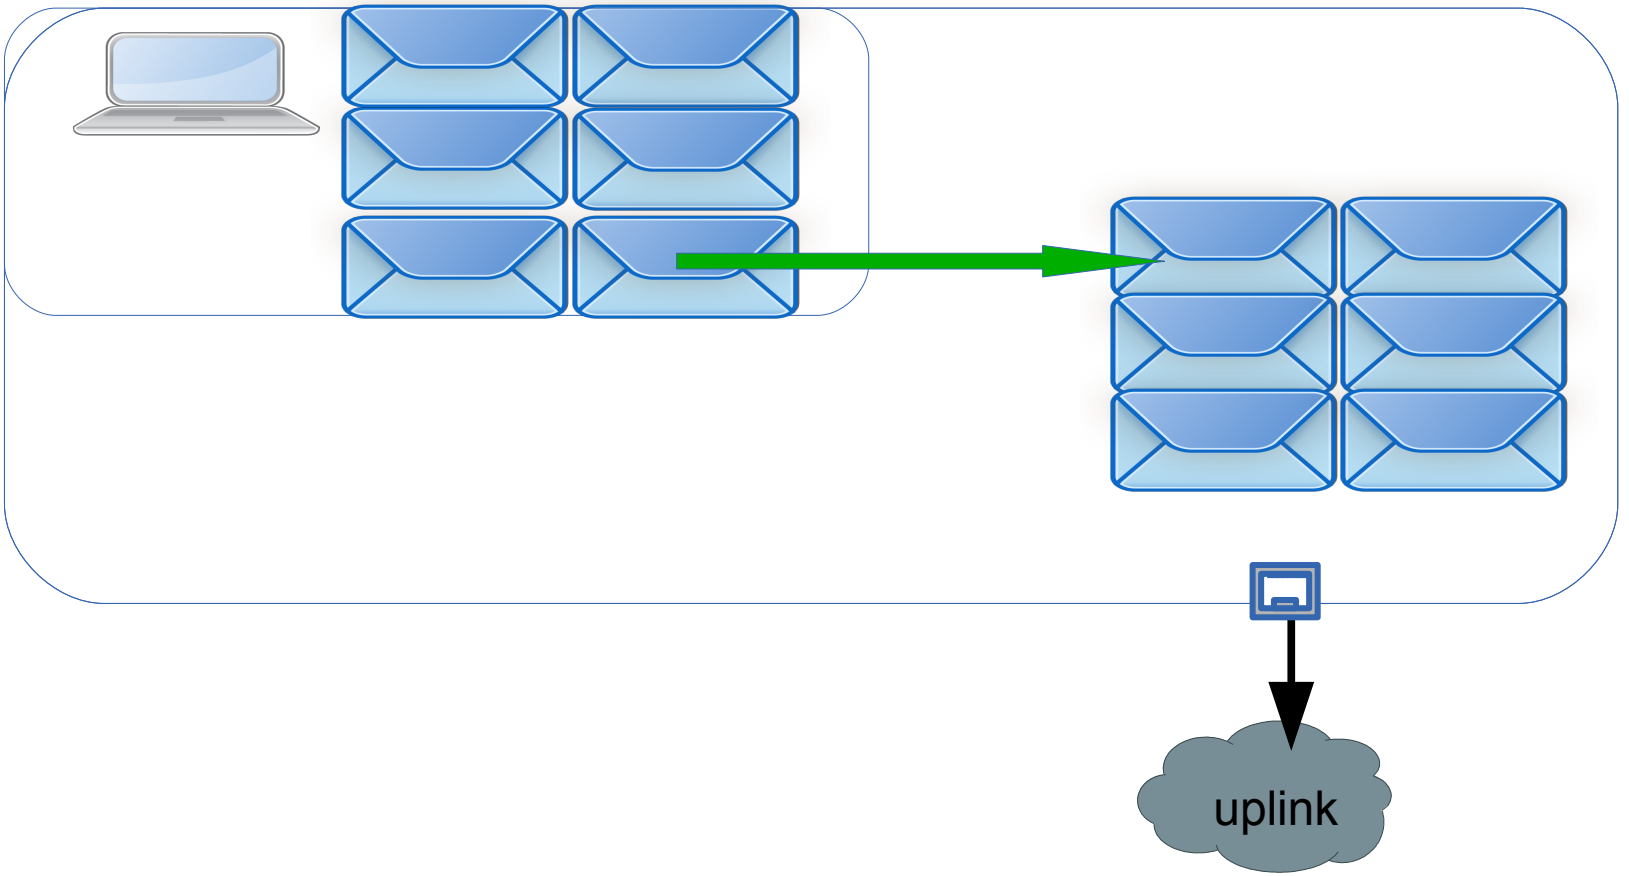

Pass-through for nested virt

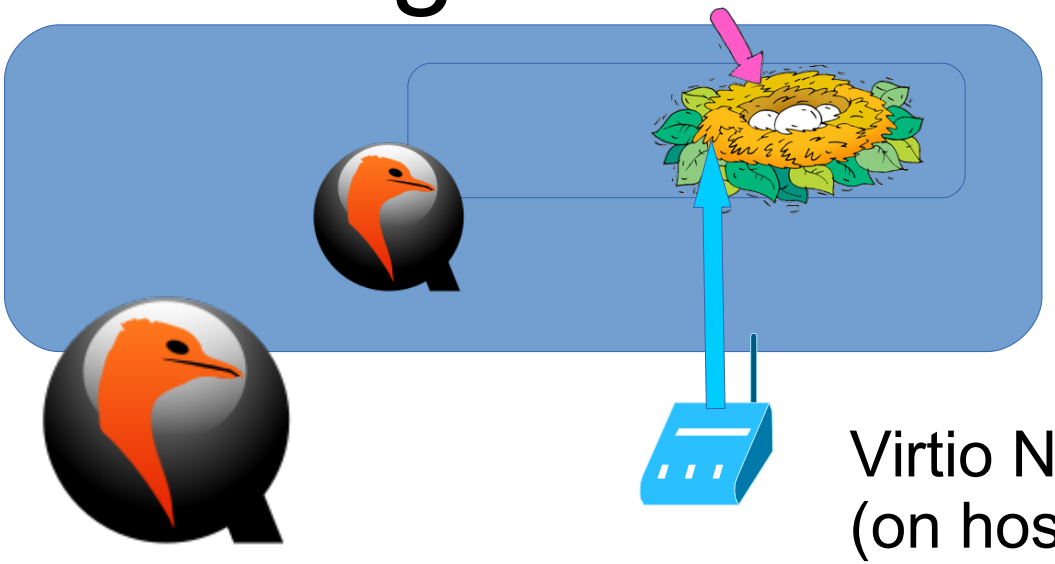

- Memory mapped: use page tables

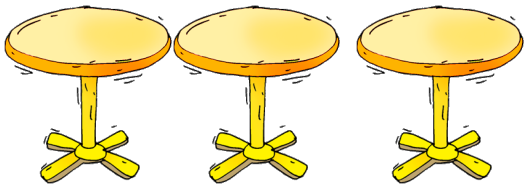

- IOMMU: translate and protect guest memory

Virtio as PCI Express device

- Uses memory mapped IO support
- Multi-root for NUMA
- Native hotplug
- Advanced Error Reporting

Summary

- Why do it?
  - Improved robustness for virtual devices
- Are we there yet?
  - Yes!
  - And there's more to come.

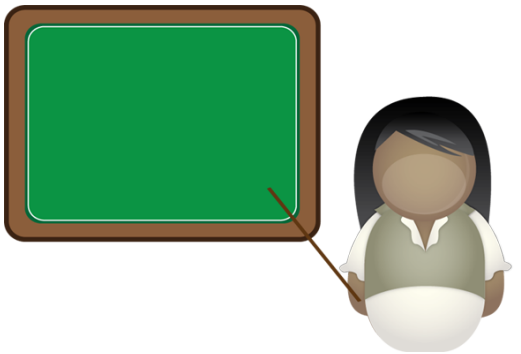

Thank you!

Virtio 0.9: Port IO versus memory on KVM x86: cycles per access (lower is better)

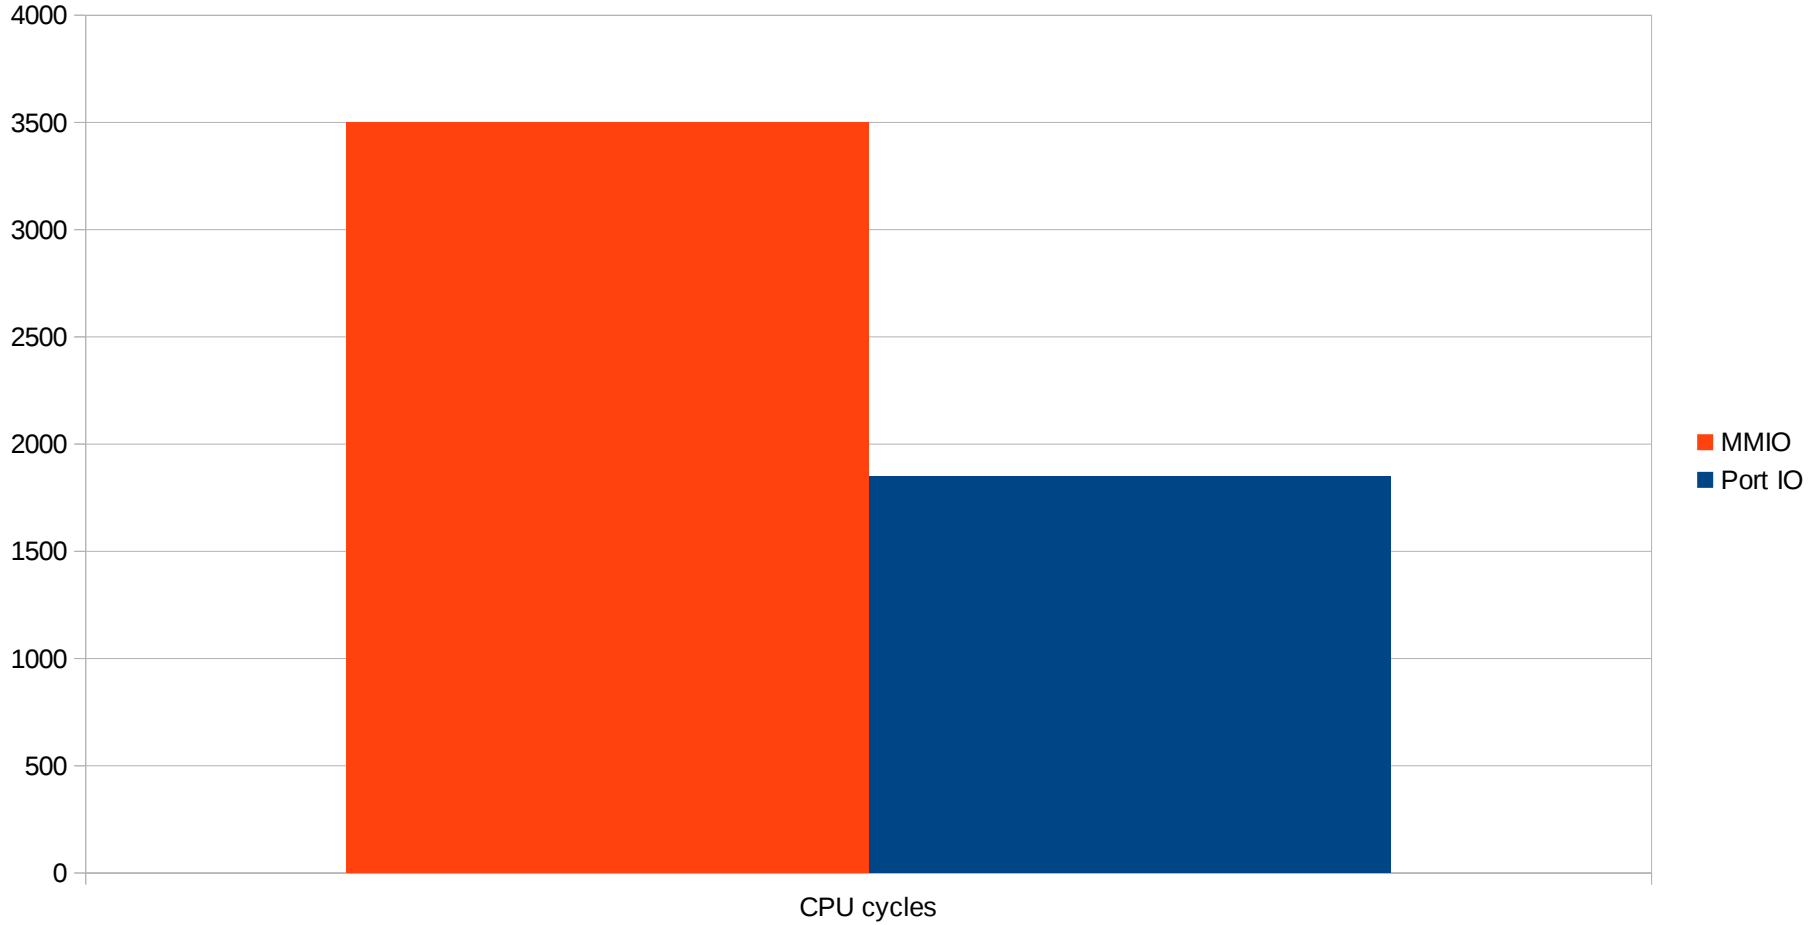

OASIS Virtio TC

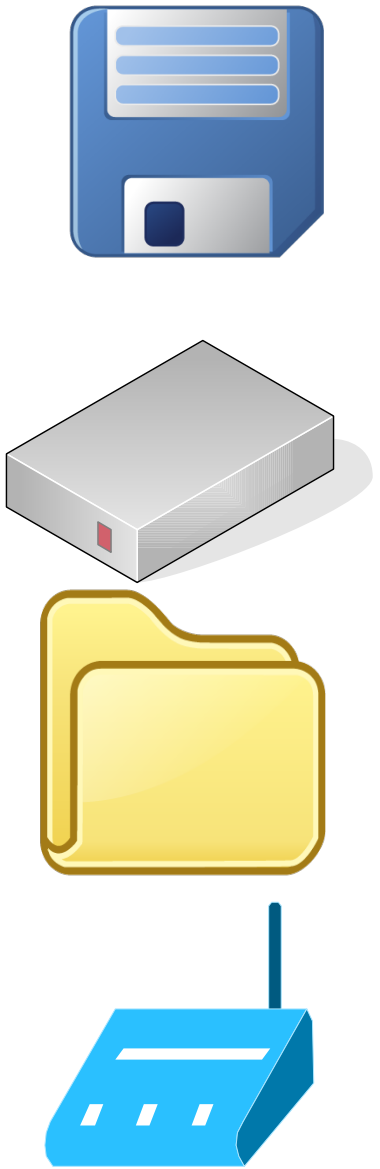

PCI

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 IO: outl

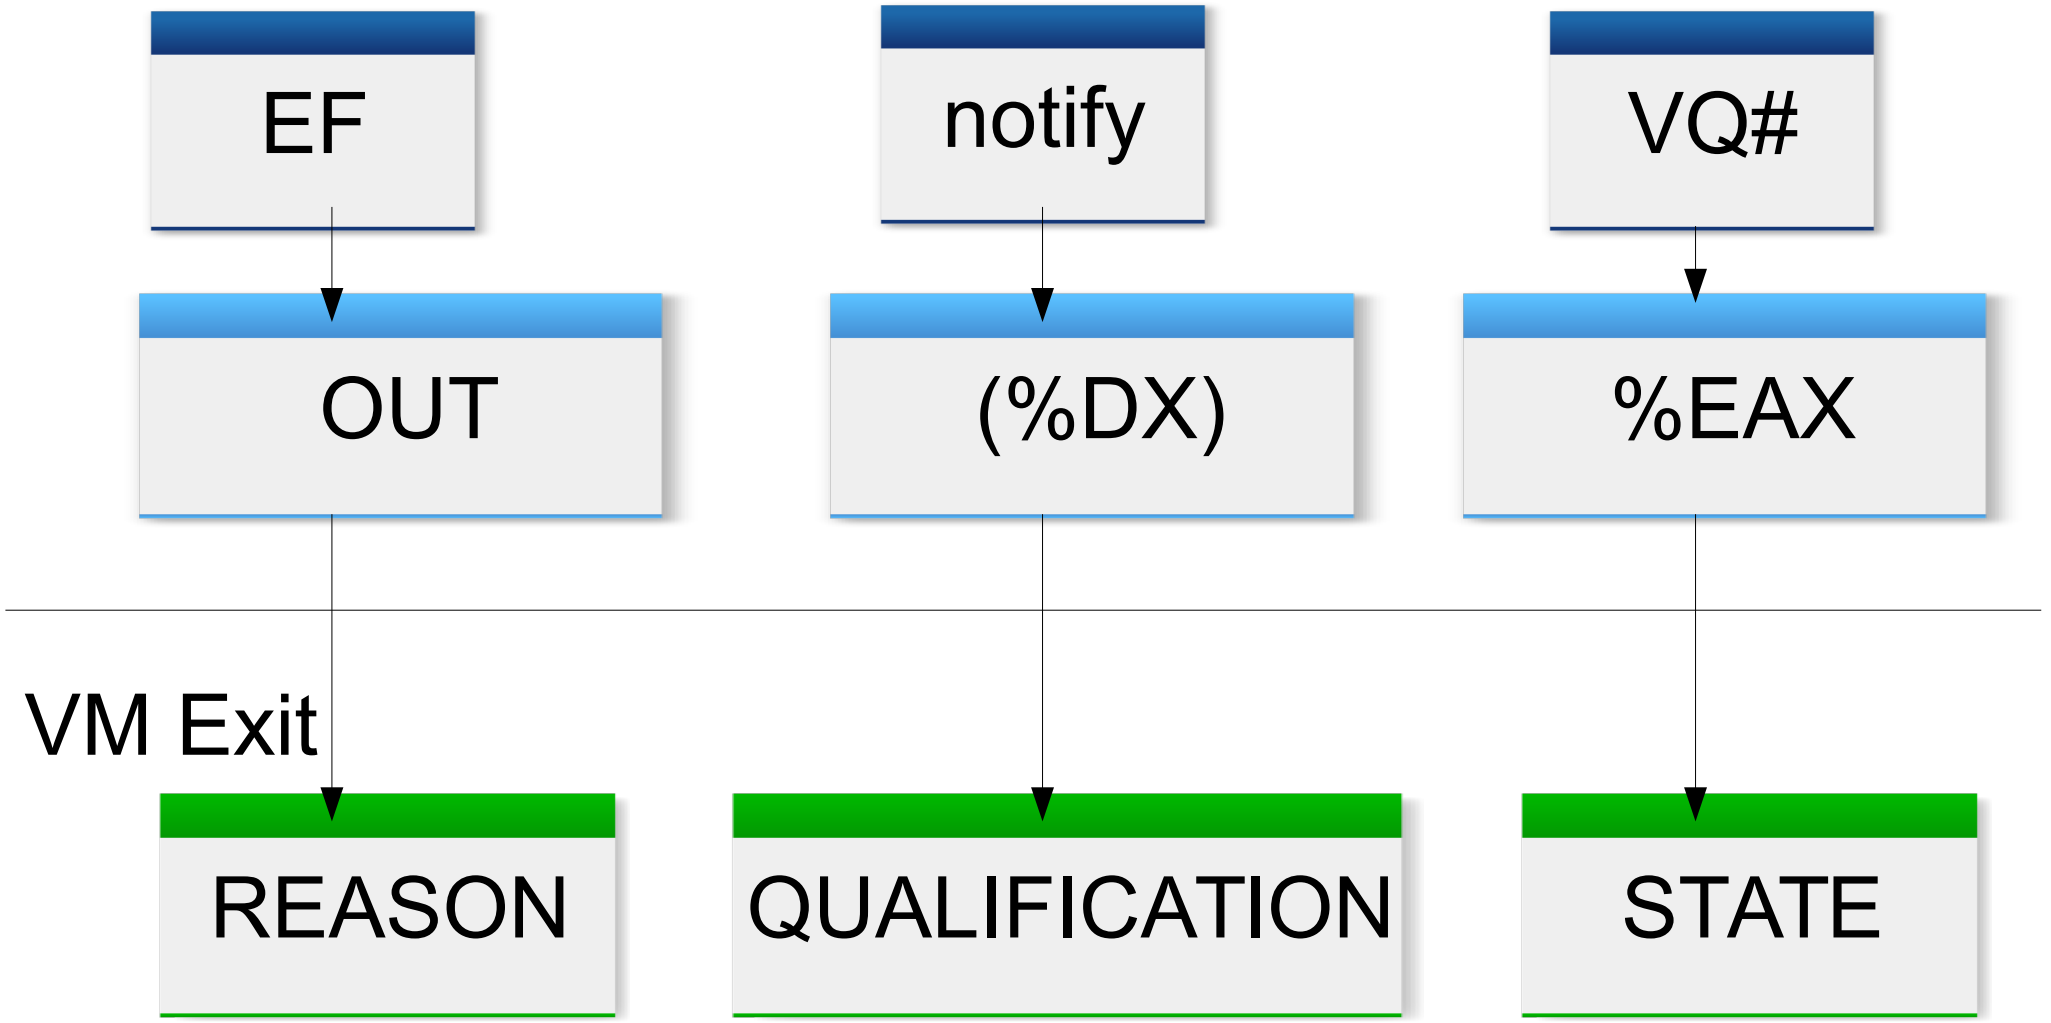

Memory mapped IO: writel

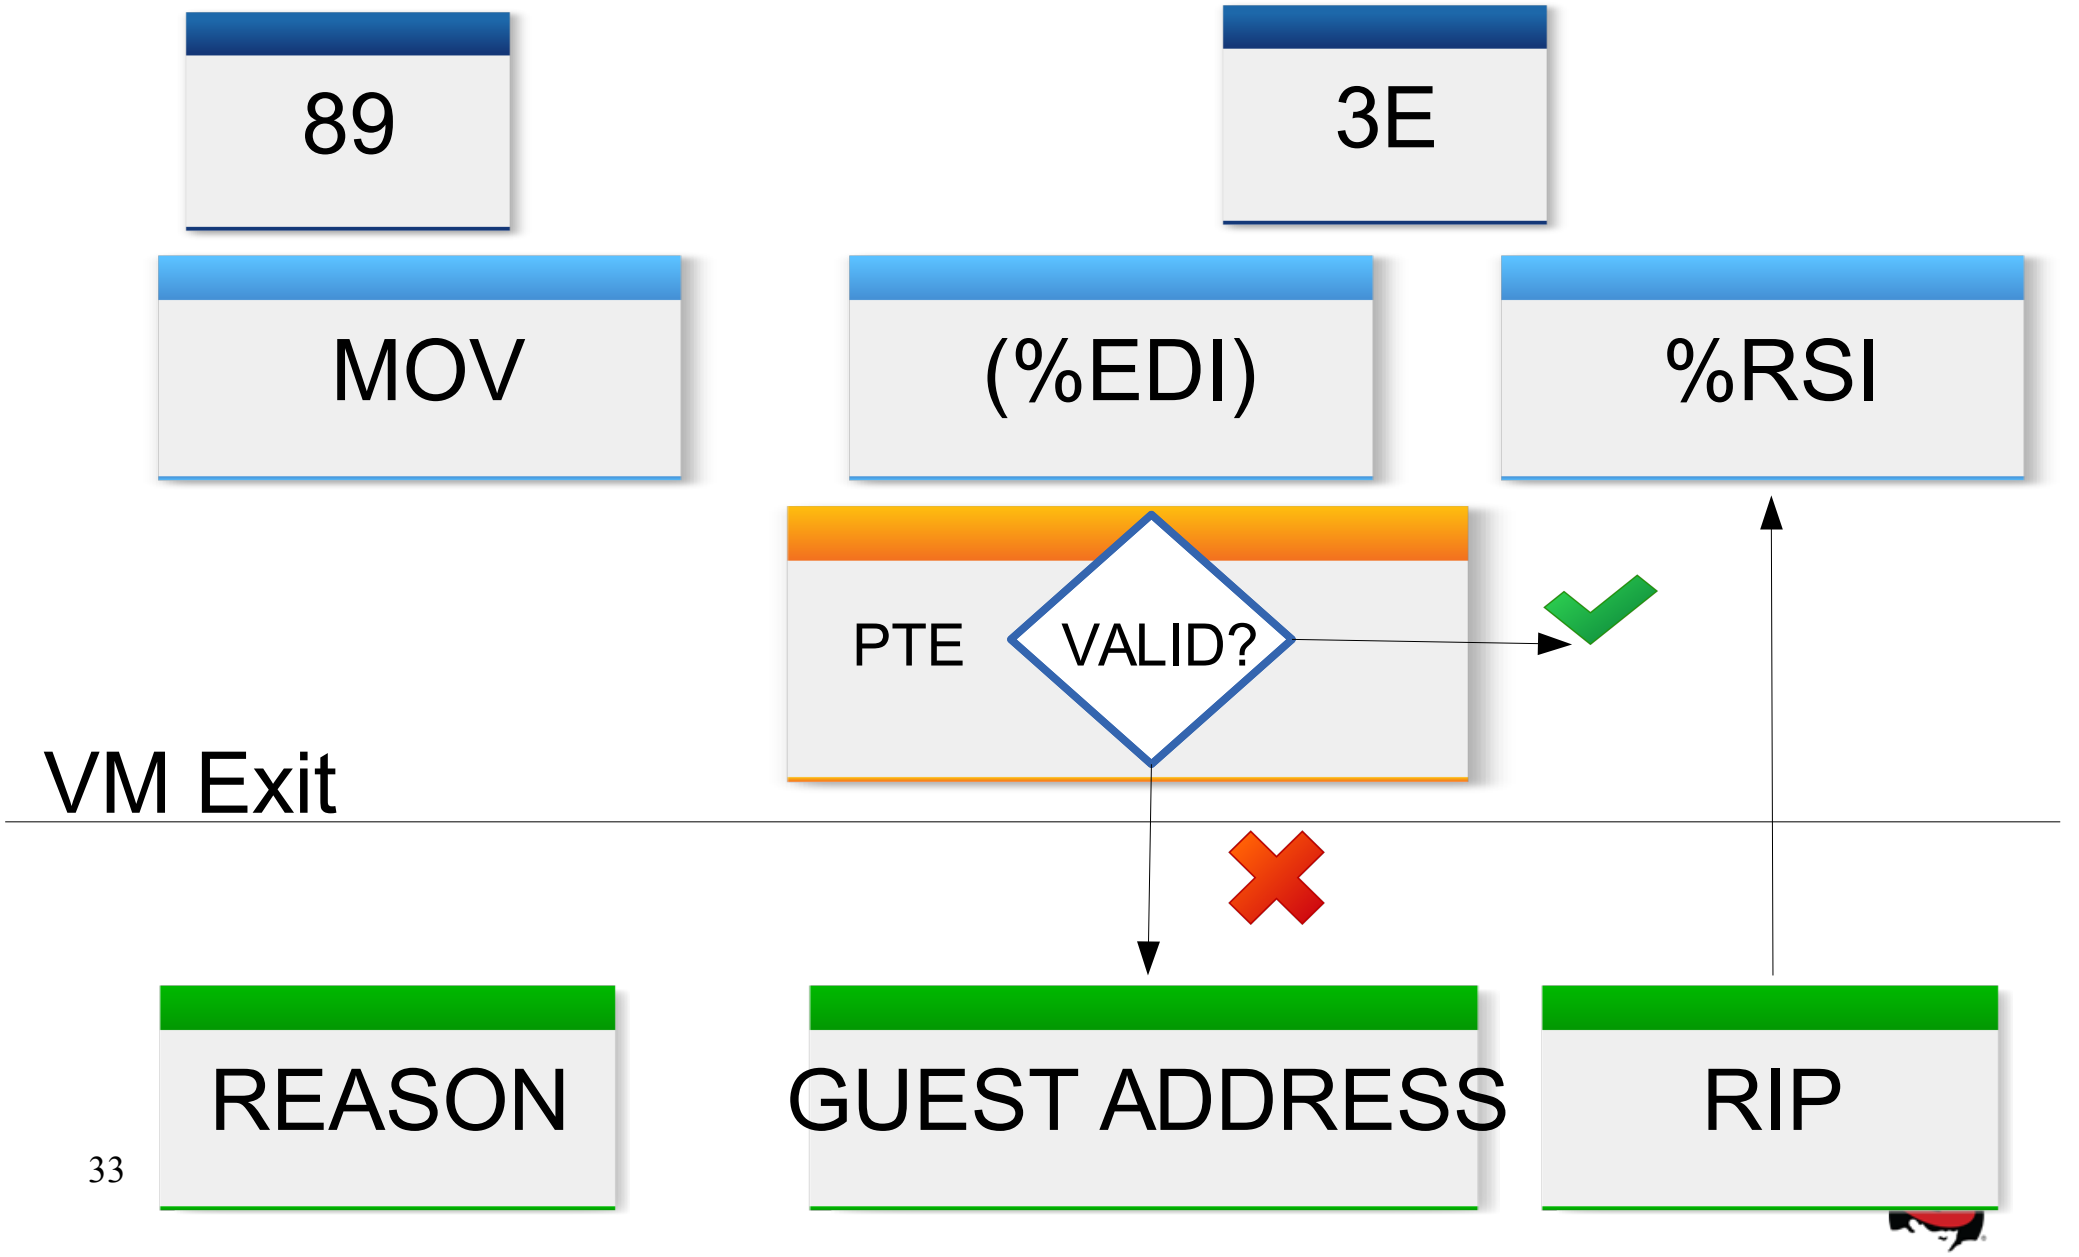

Fast MMIO

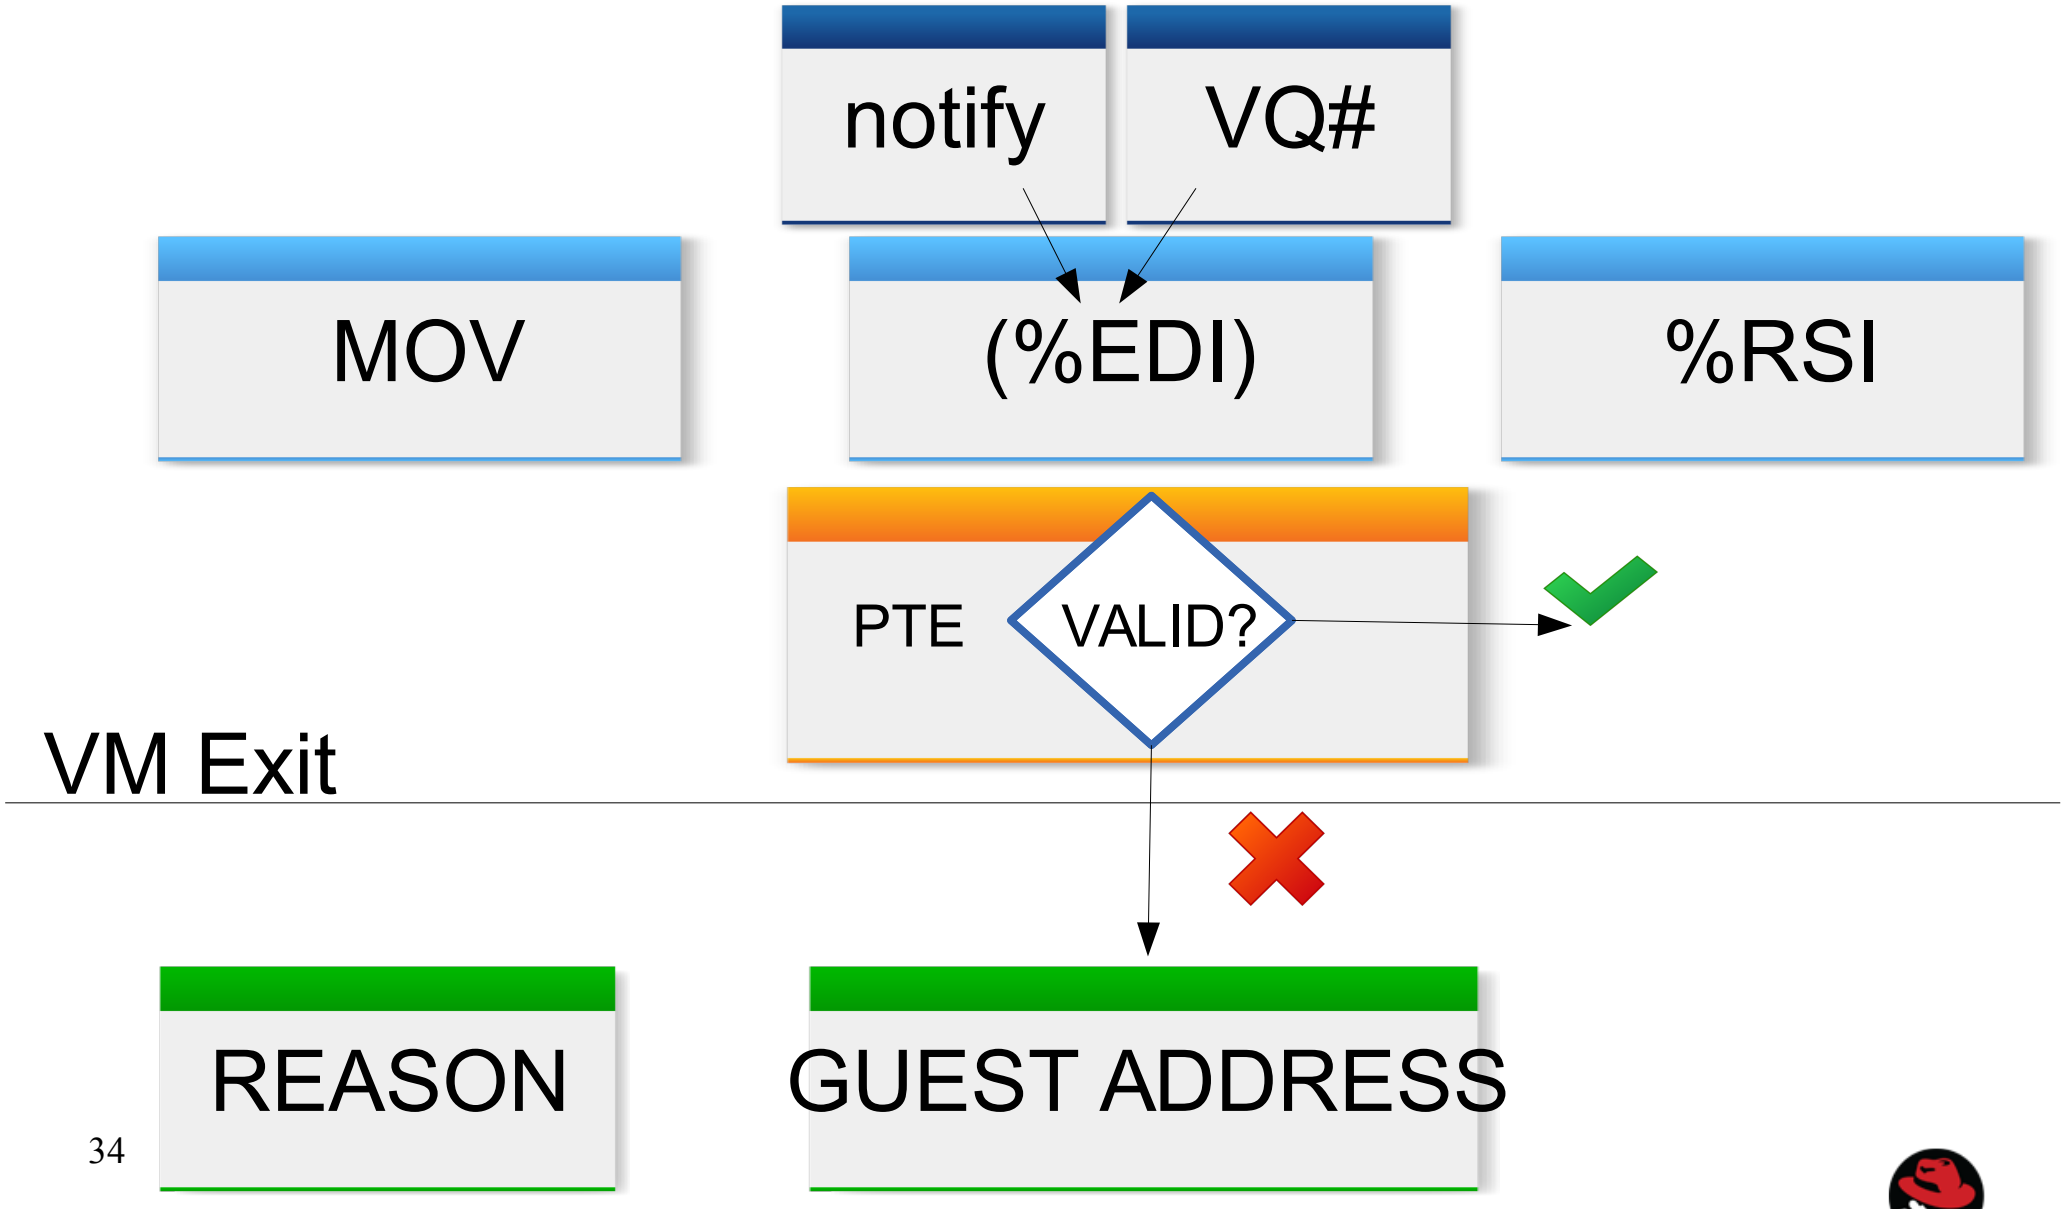

Multiple interfaces

Memory requirements

0.9

1.0

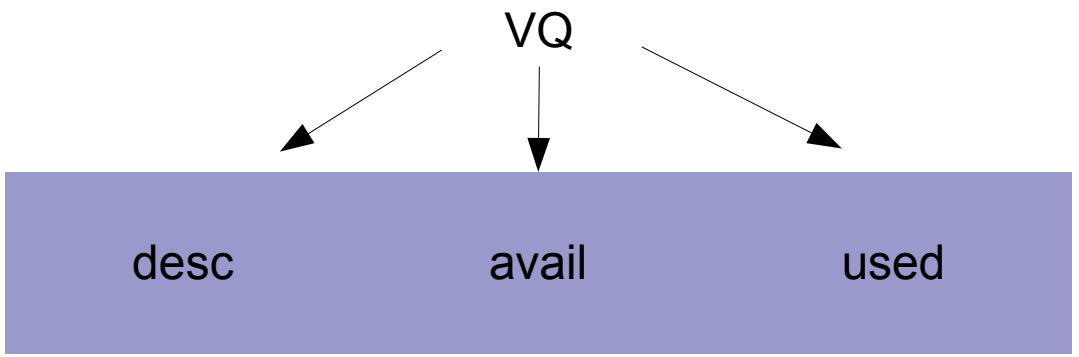

0.9

features

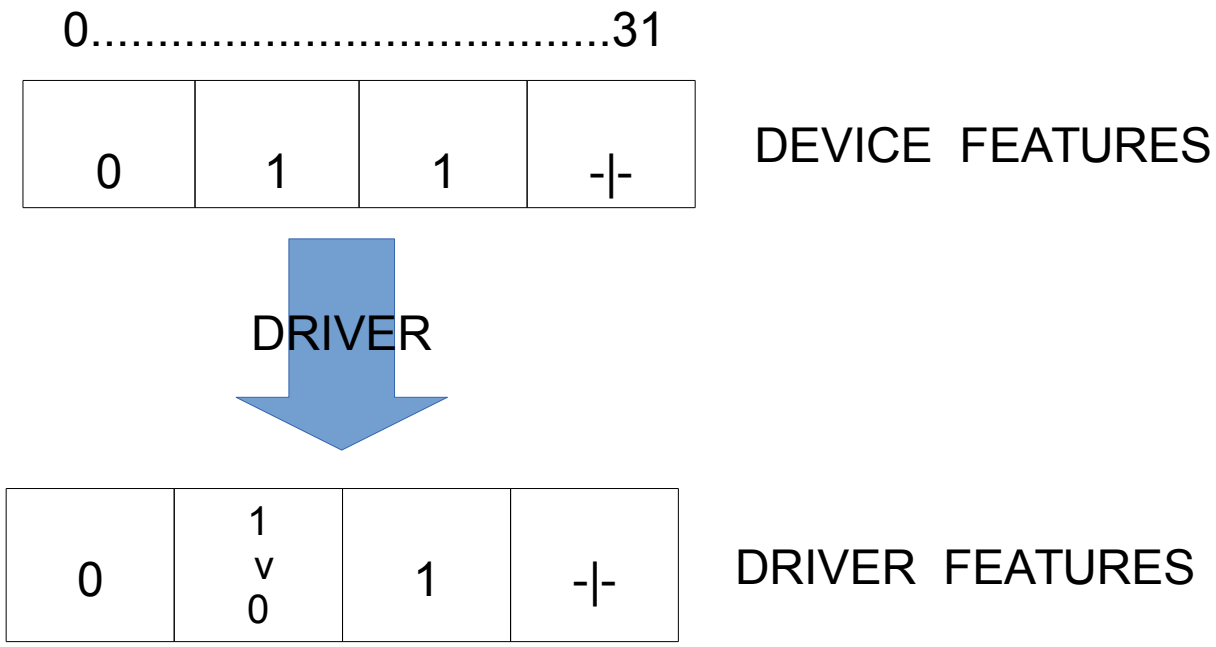

1.0

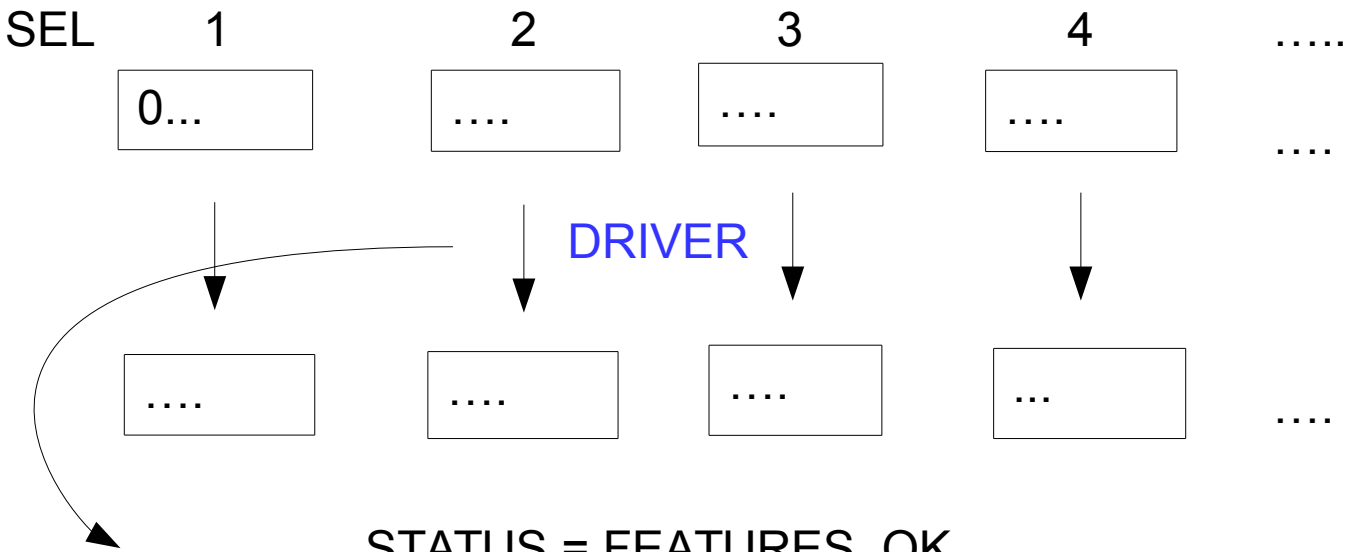

Endianness

Virtio 0.9

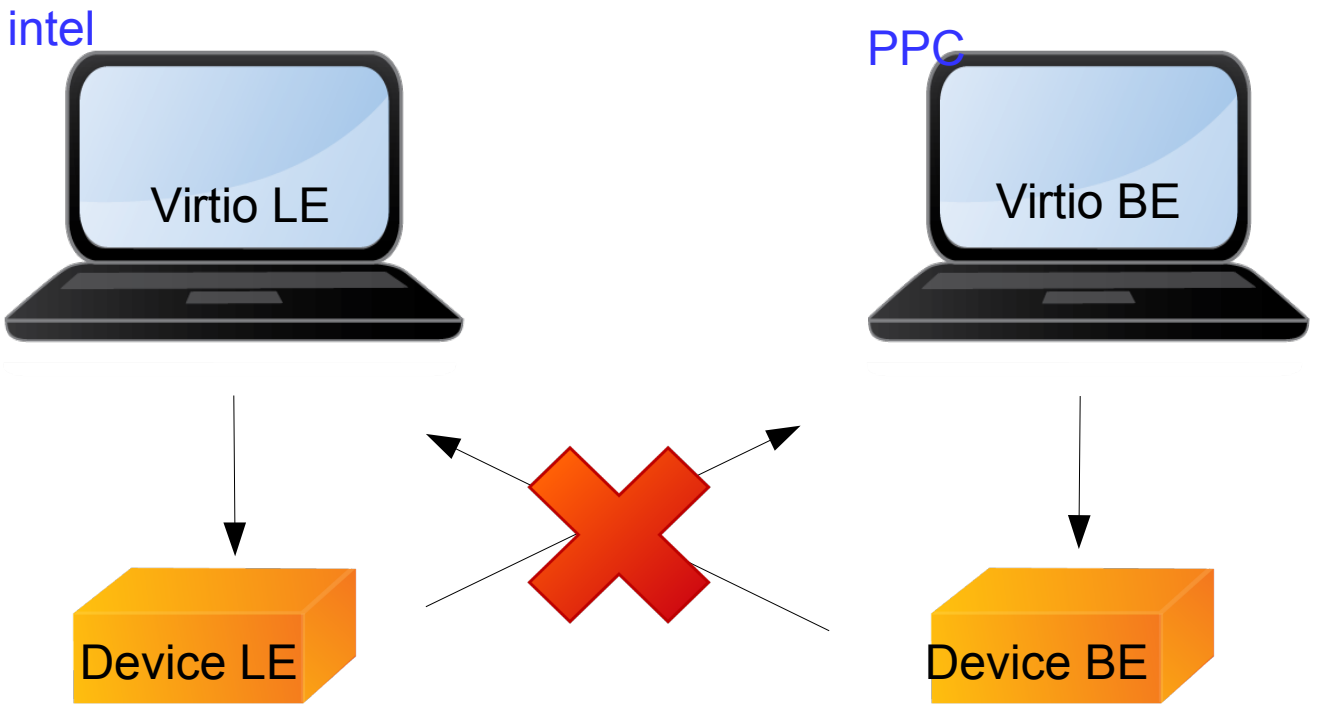

Virtio 1.0

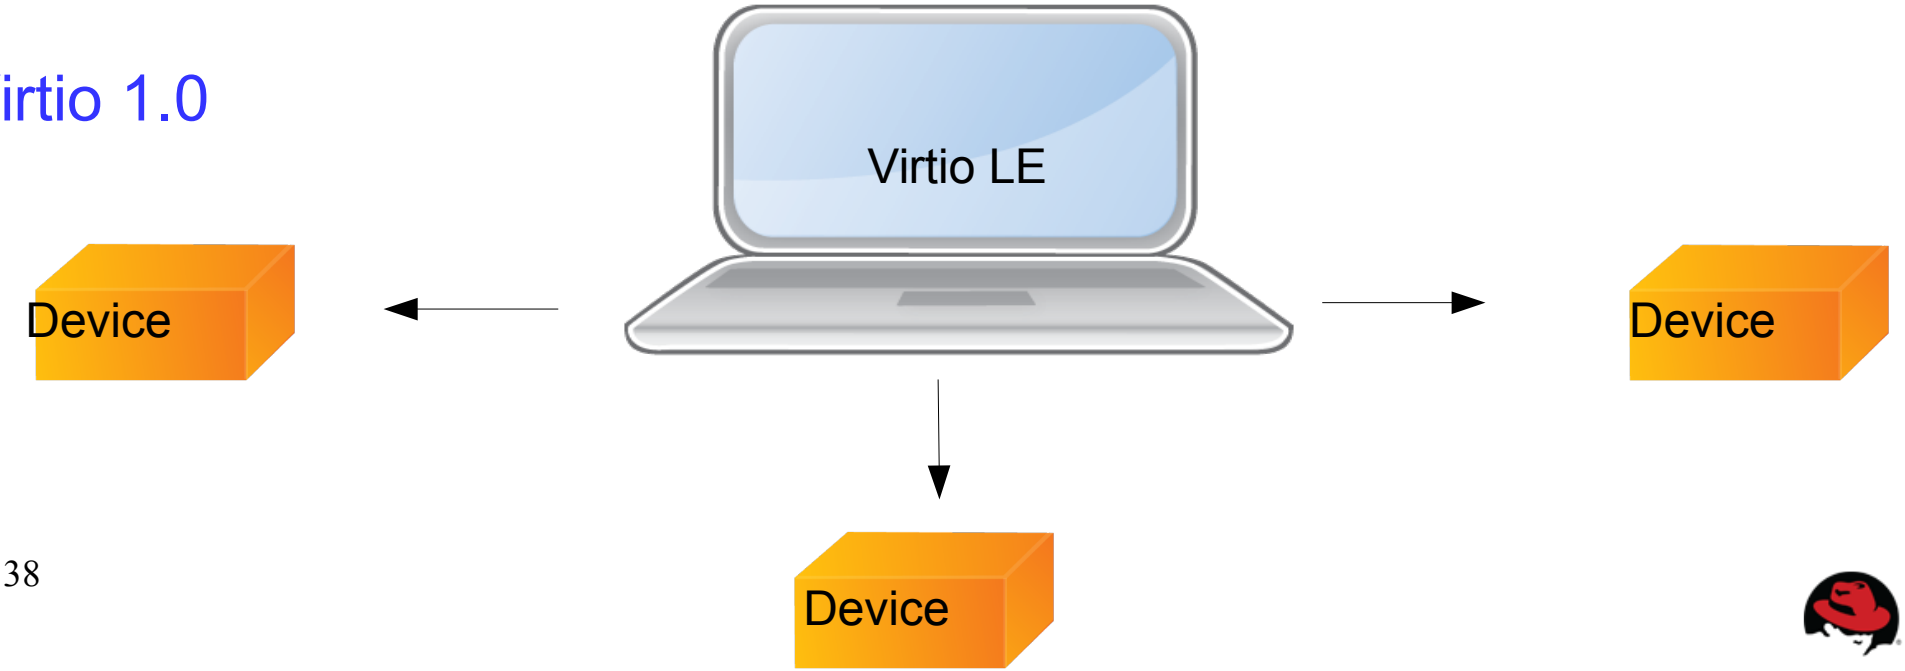

compatibility

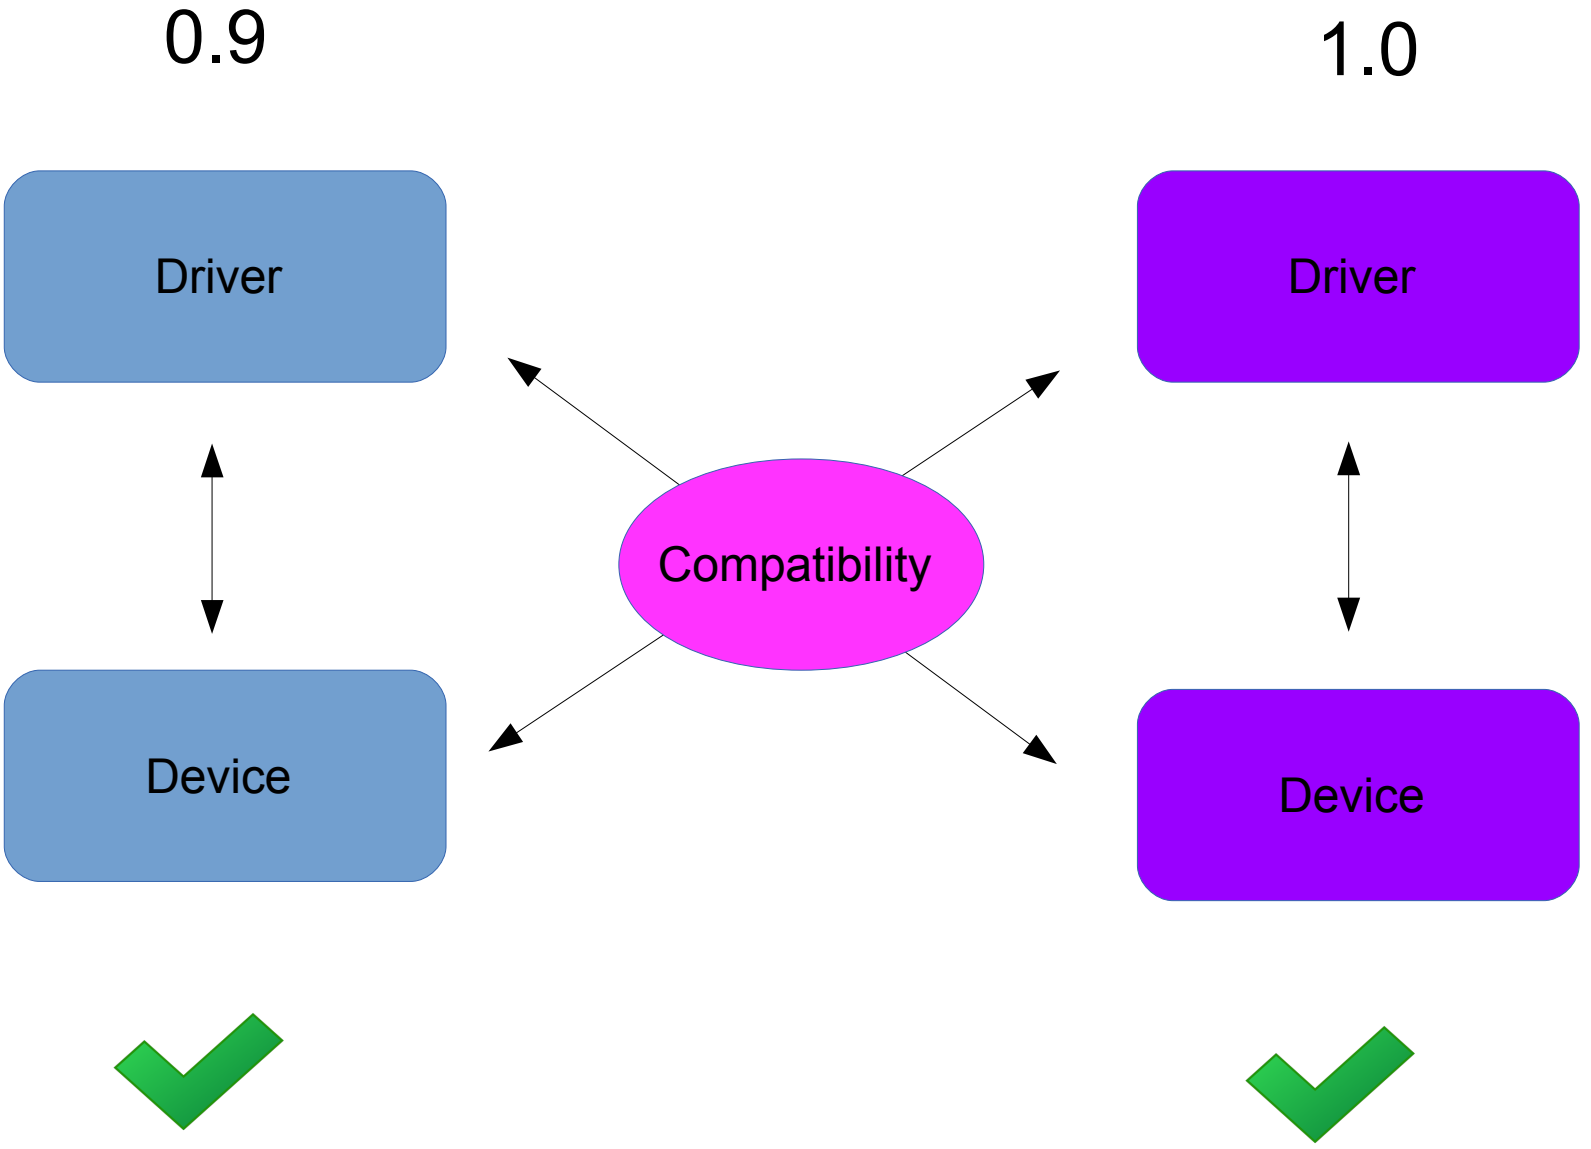

Packet layout

Virtio 0.9

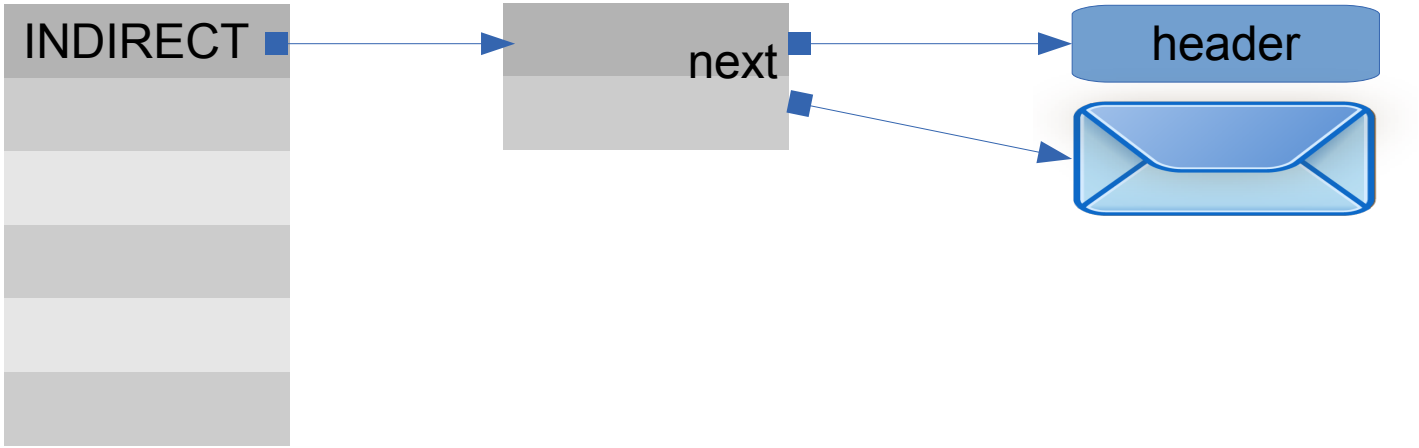

Virtio 1.0

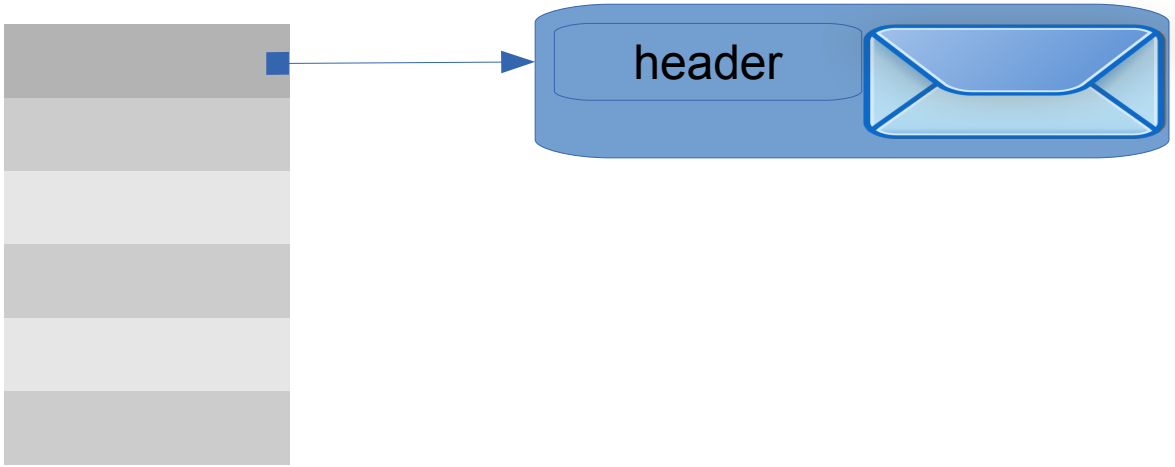

Packet layout: transactions per sec (higher is better)

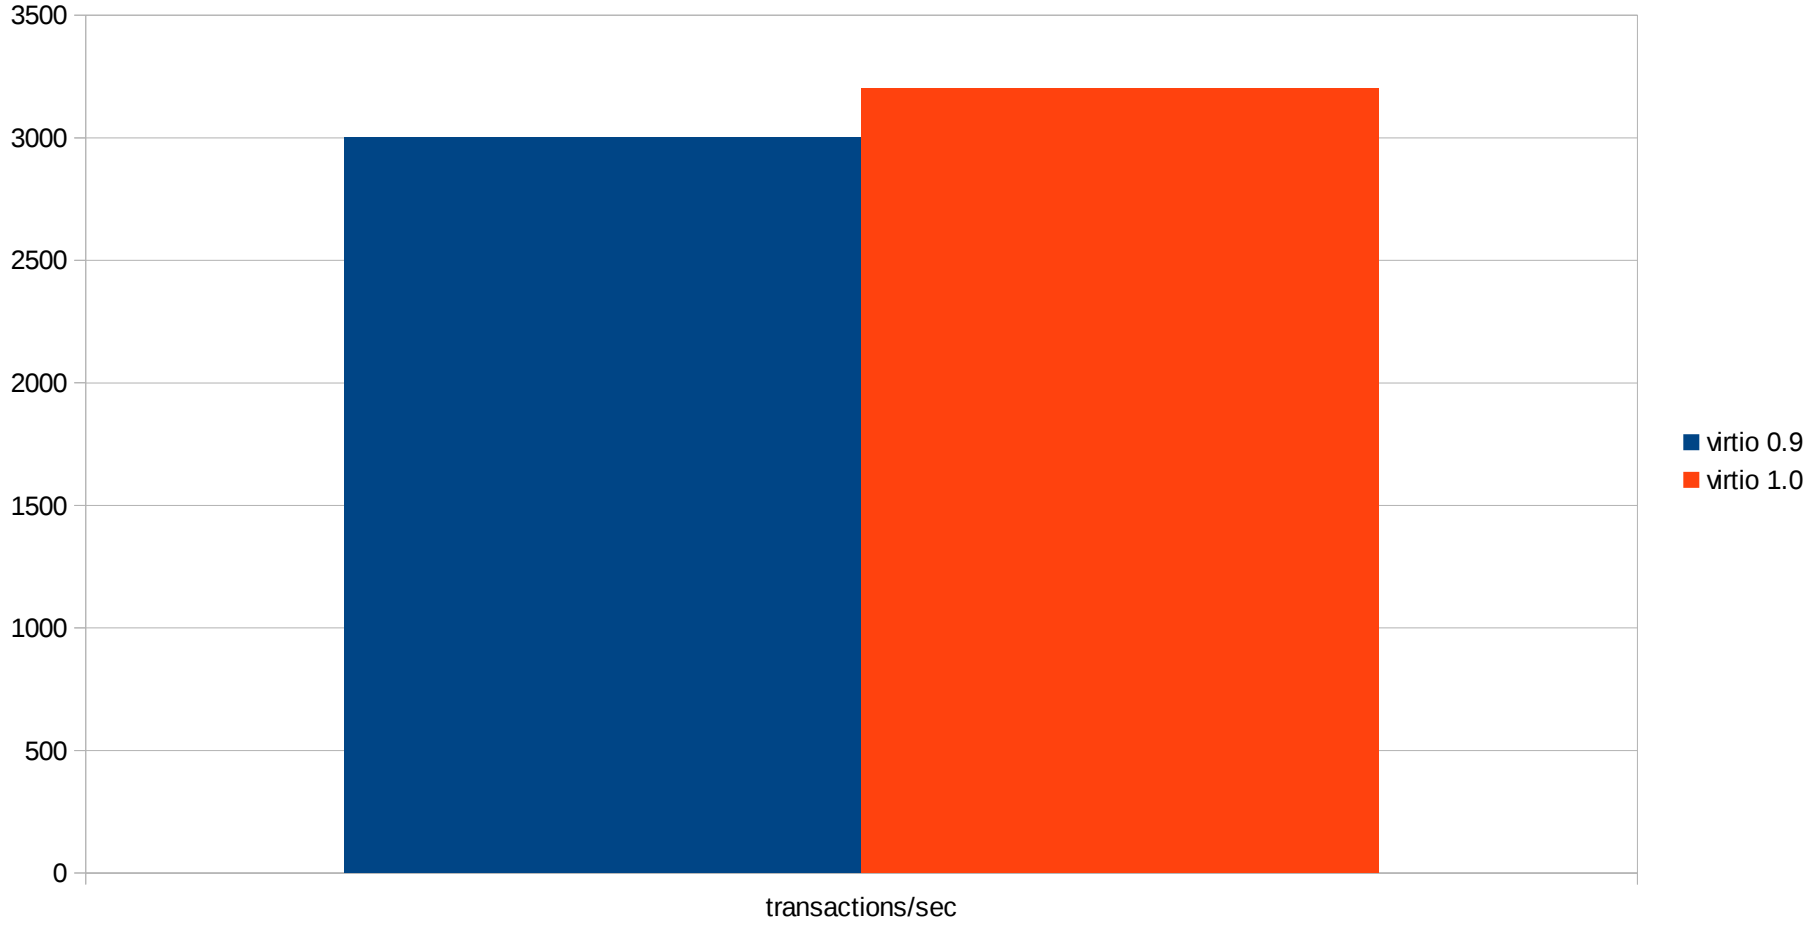

More: virtio 1.0 versus 0.9.5

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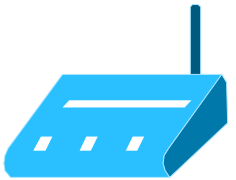
- Virtio-net dynamic offloads
- Already upstream (based on spec draft)

vhost updates

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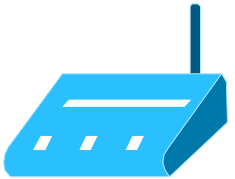
- No need for driver changes

Kvm networking

- Openvswitch – if time allows
- Ethernet bridge

Bridge FDB

London	Heathrow
Paris	CDG

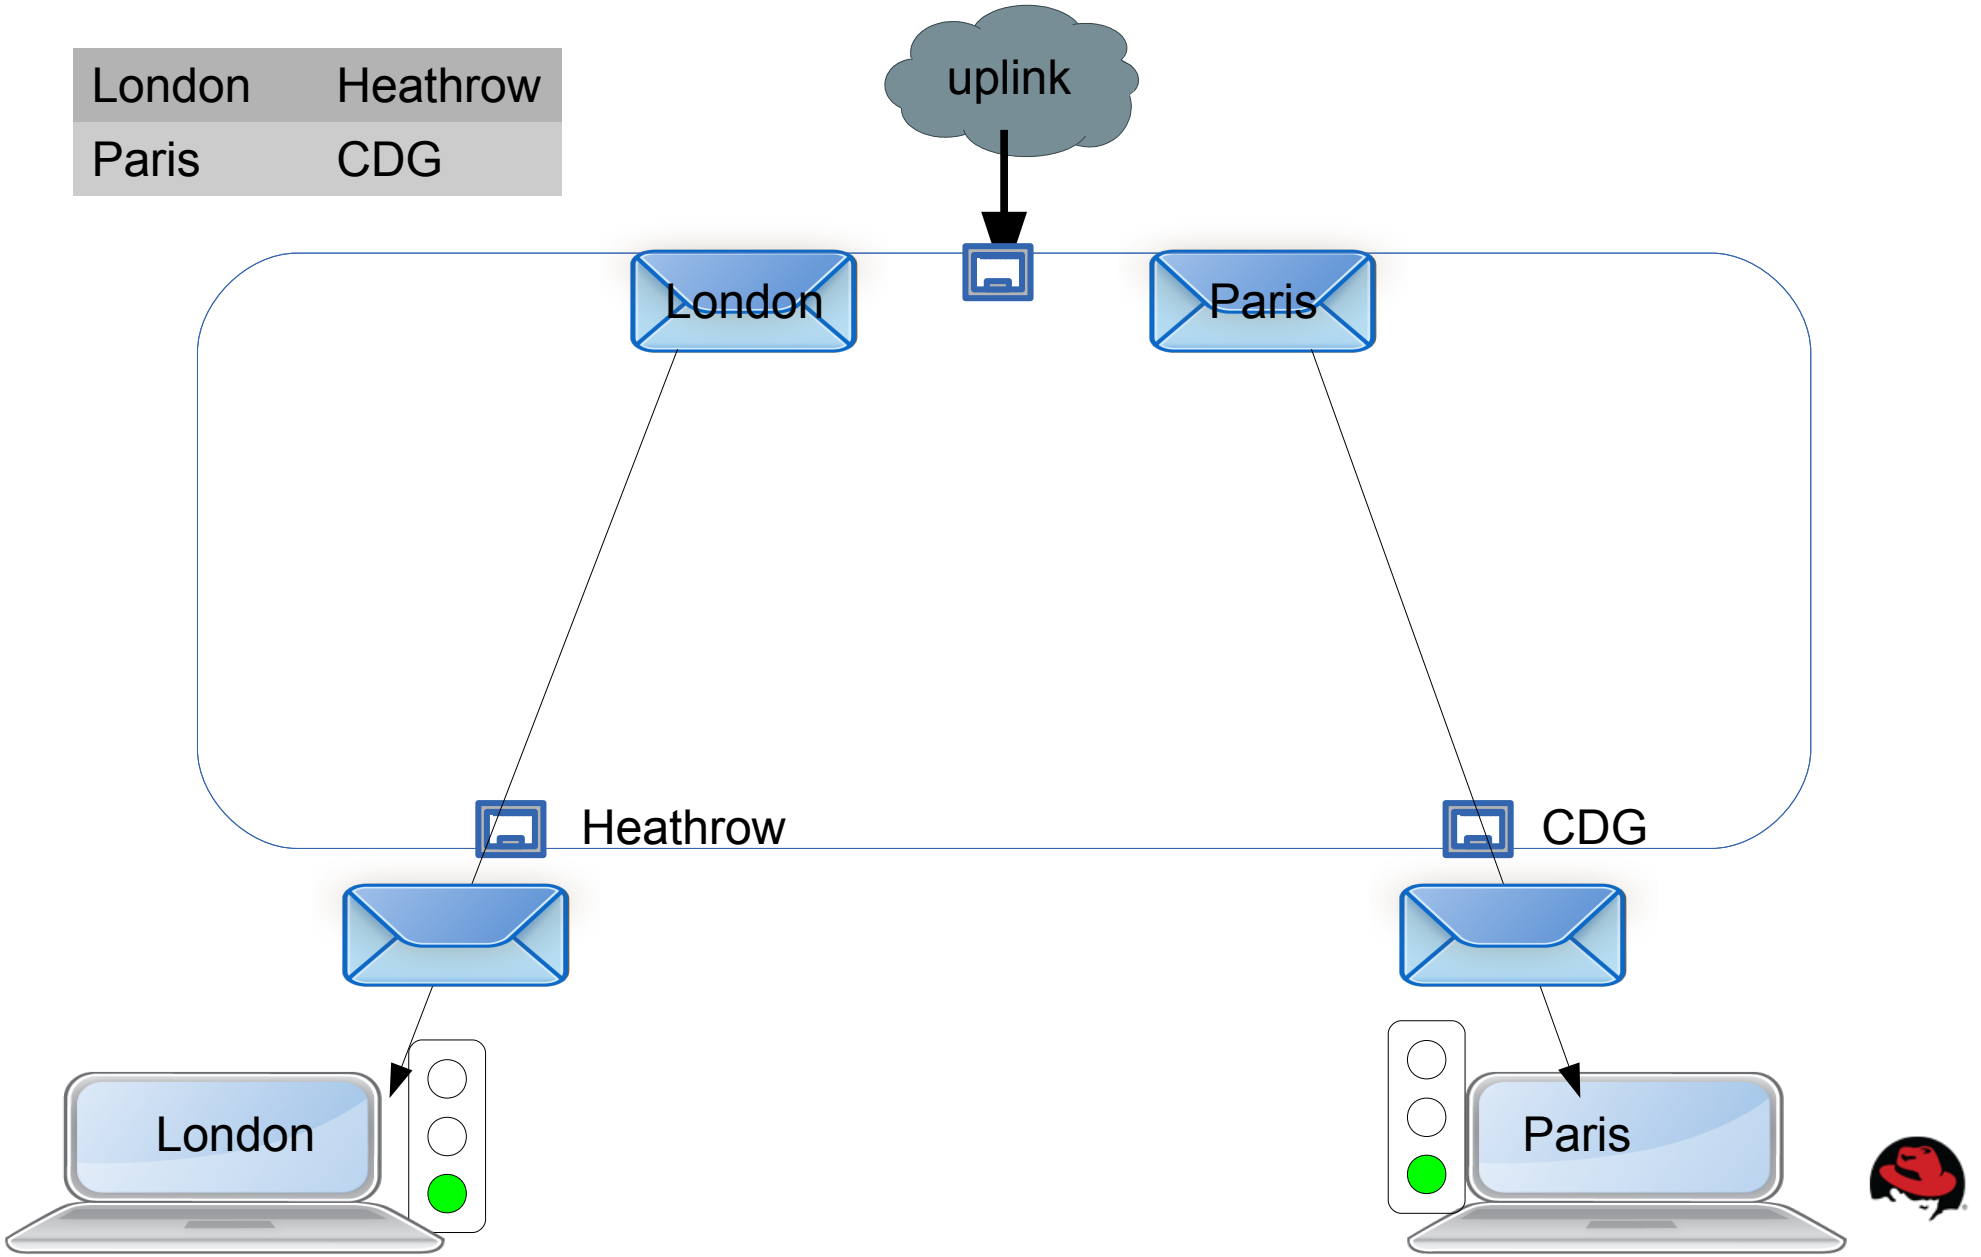

Flood: DOS potential

London	Heathrow
Paris	CDG

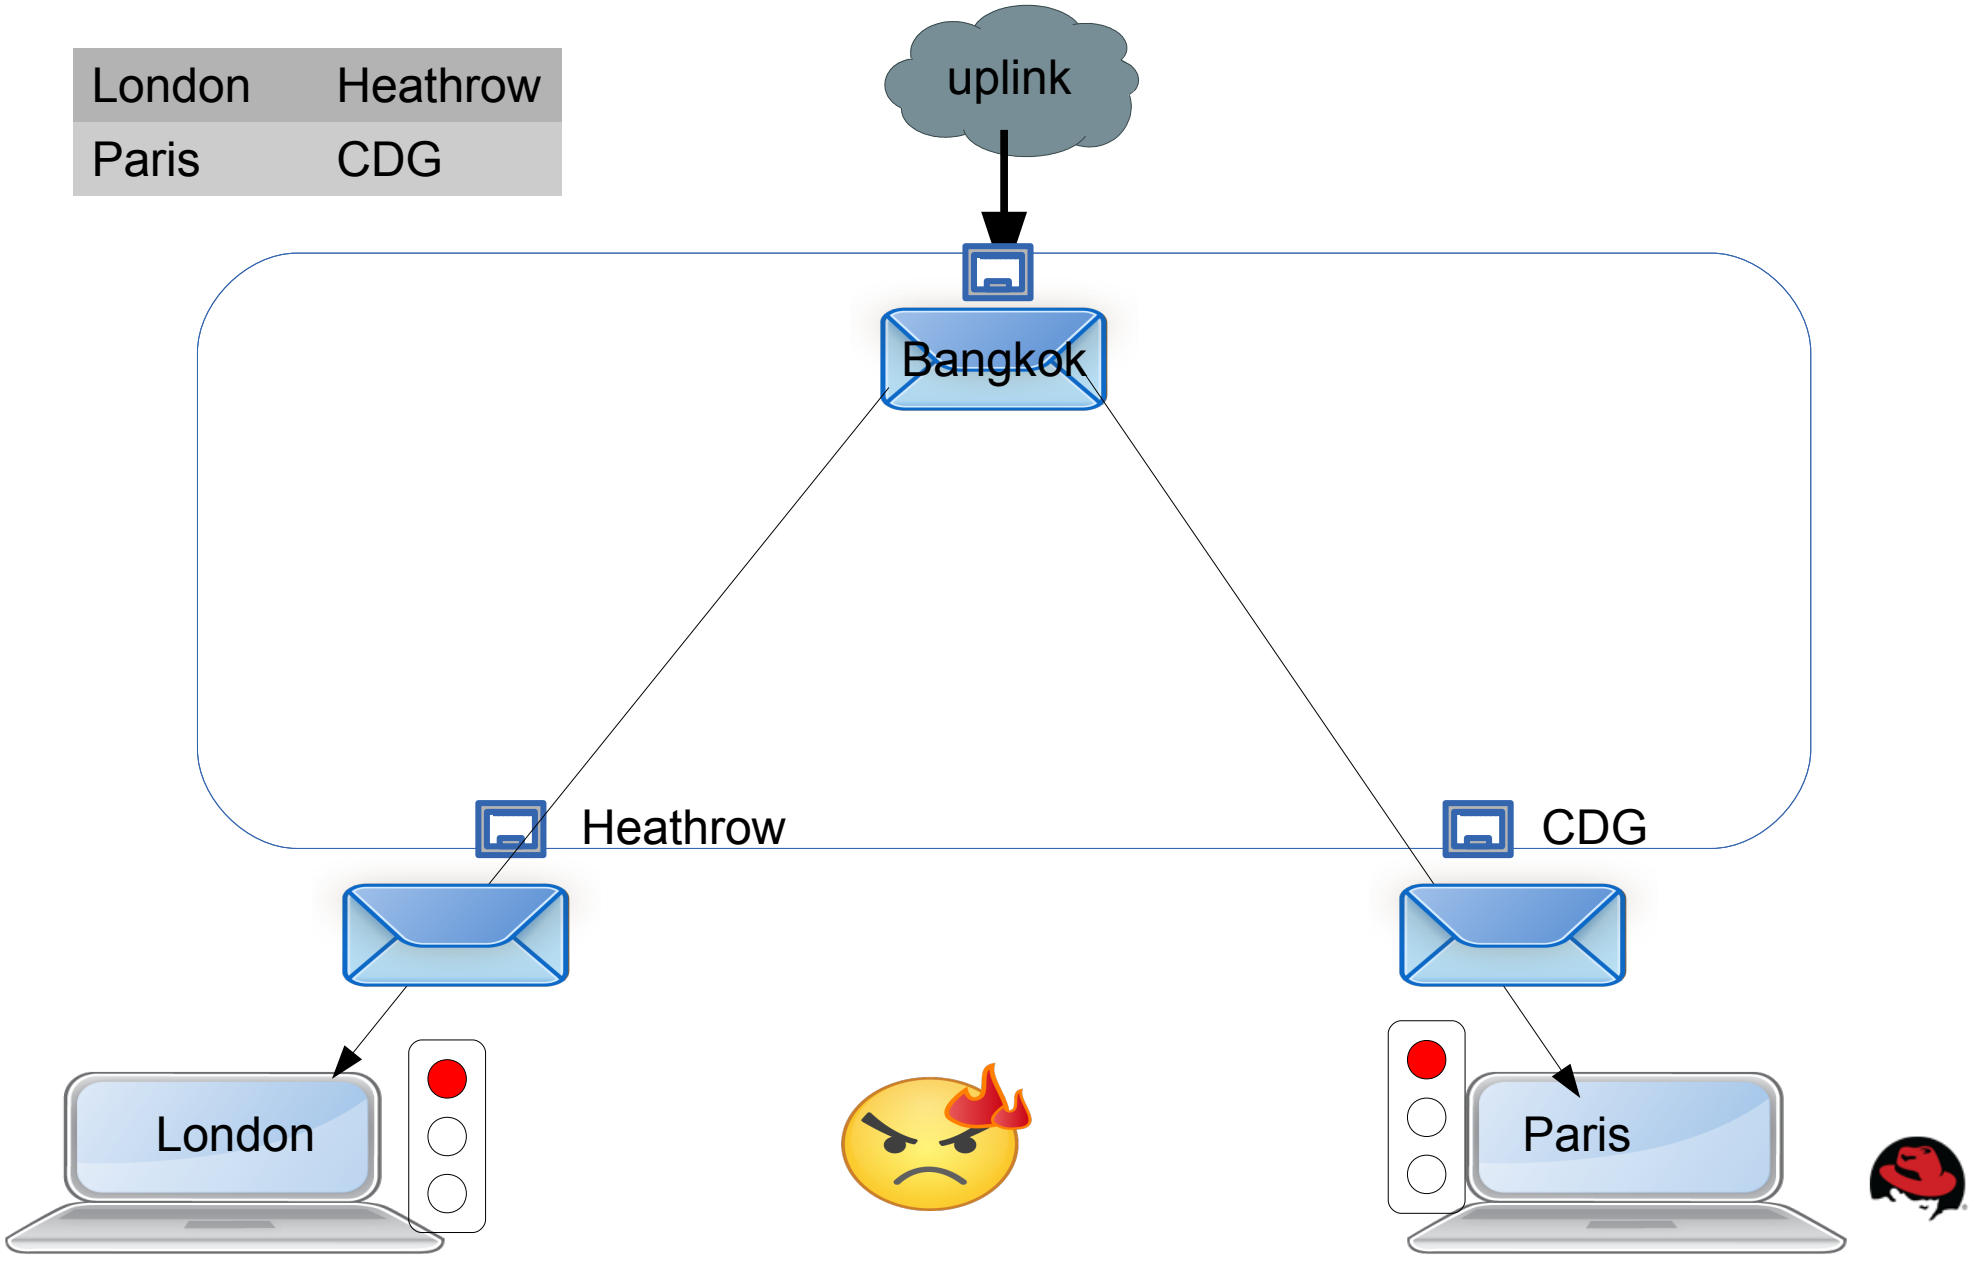

Disable flood

London	Heathrow
Paris	CDG

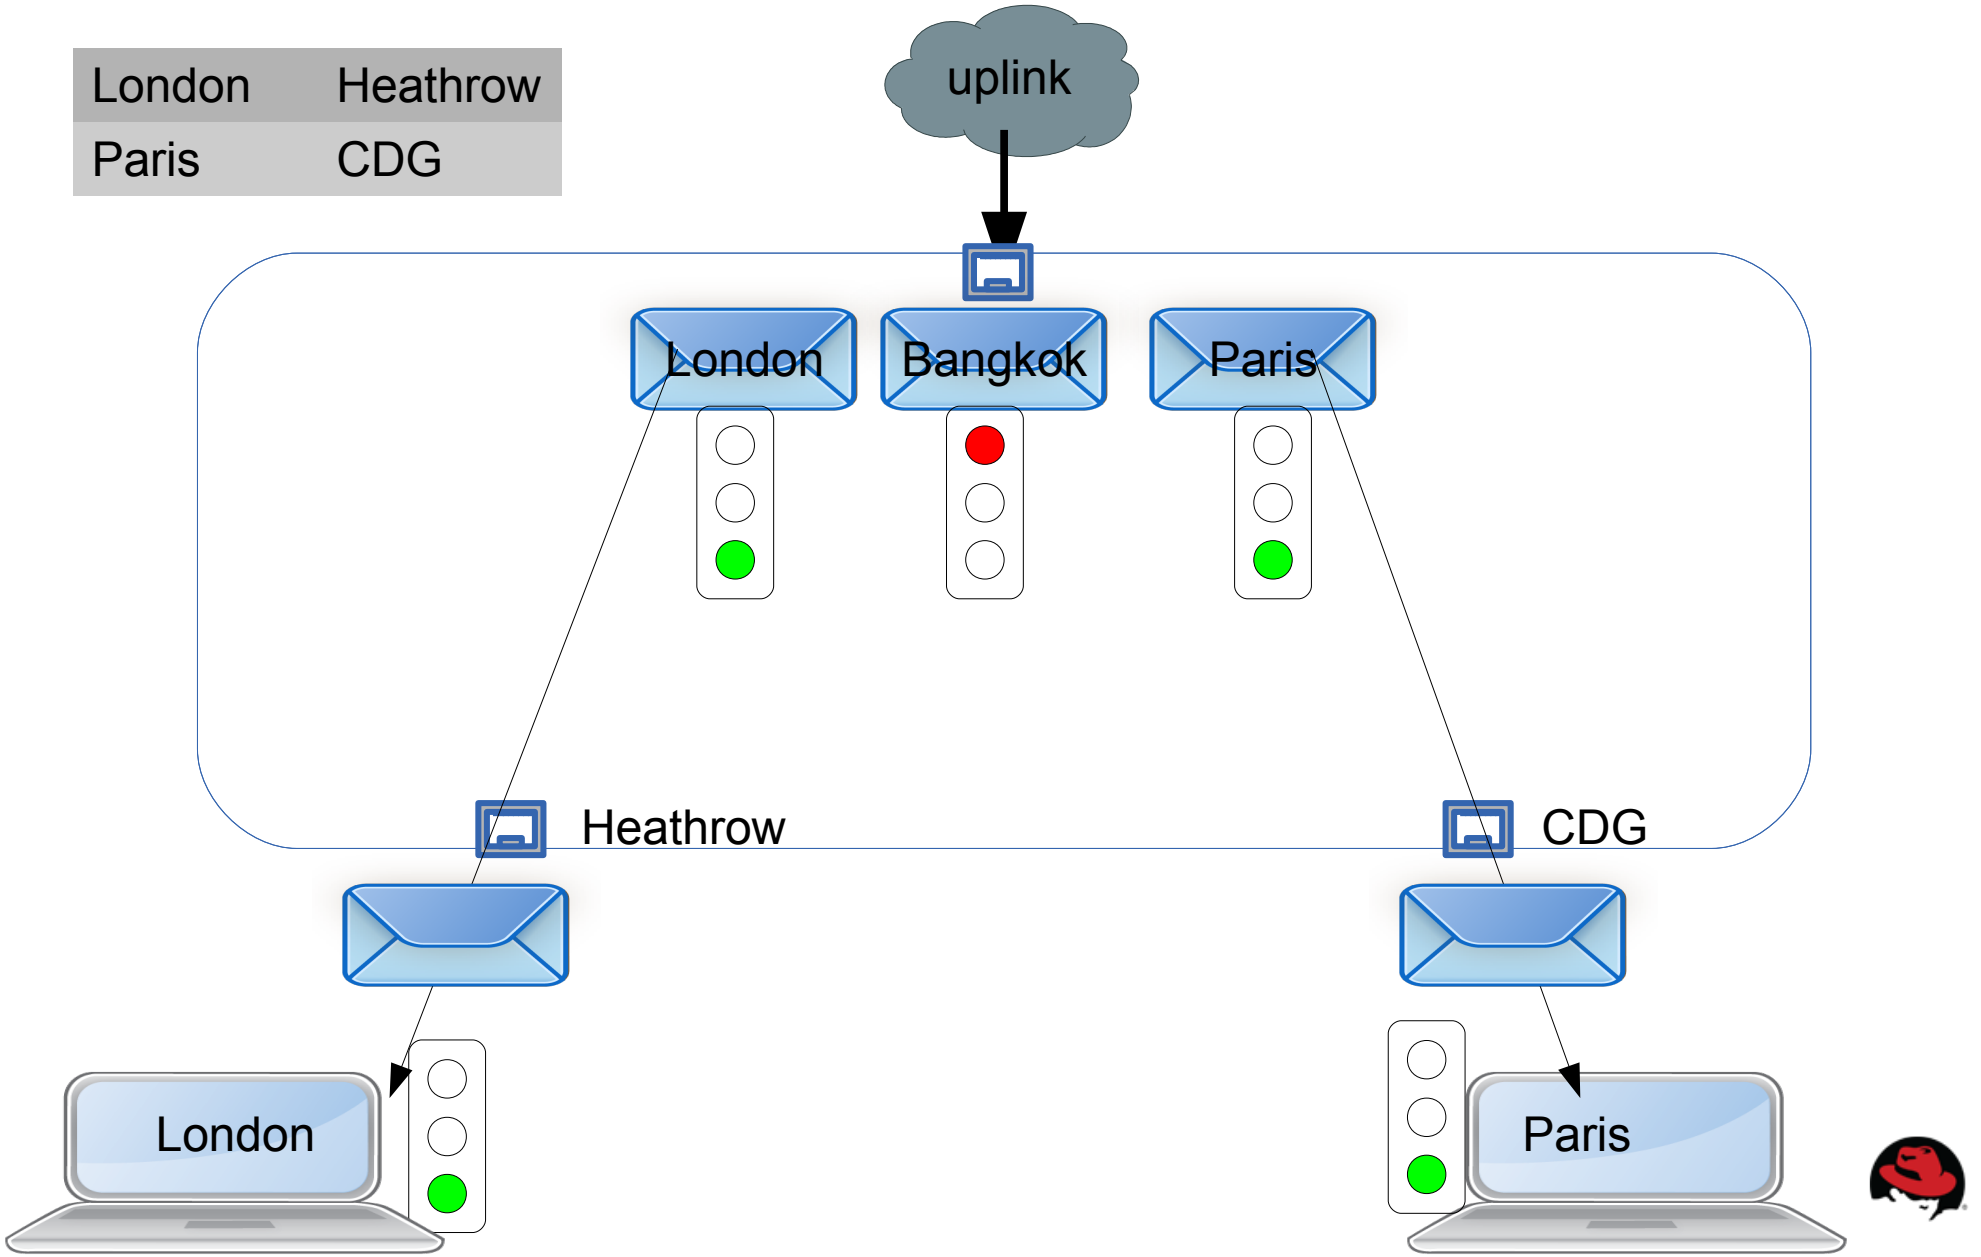

softmac

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

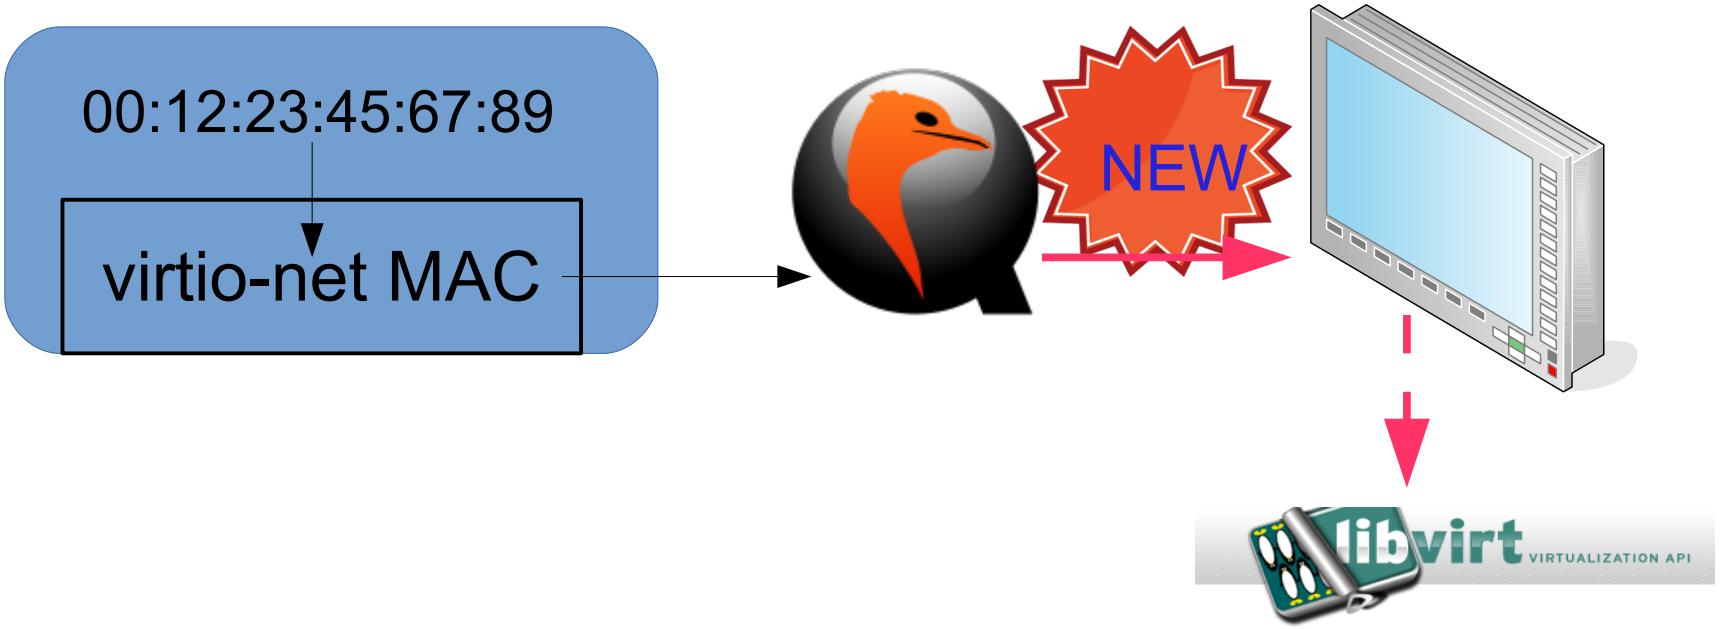

Using softmac/non promiscuous

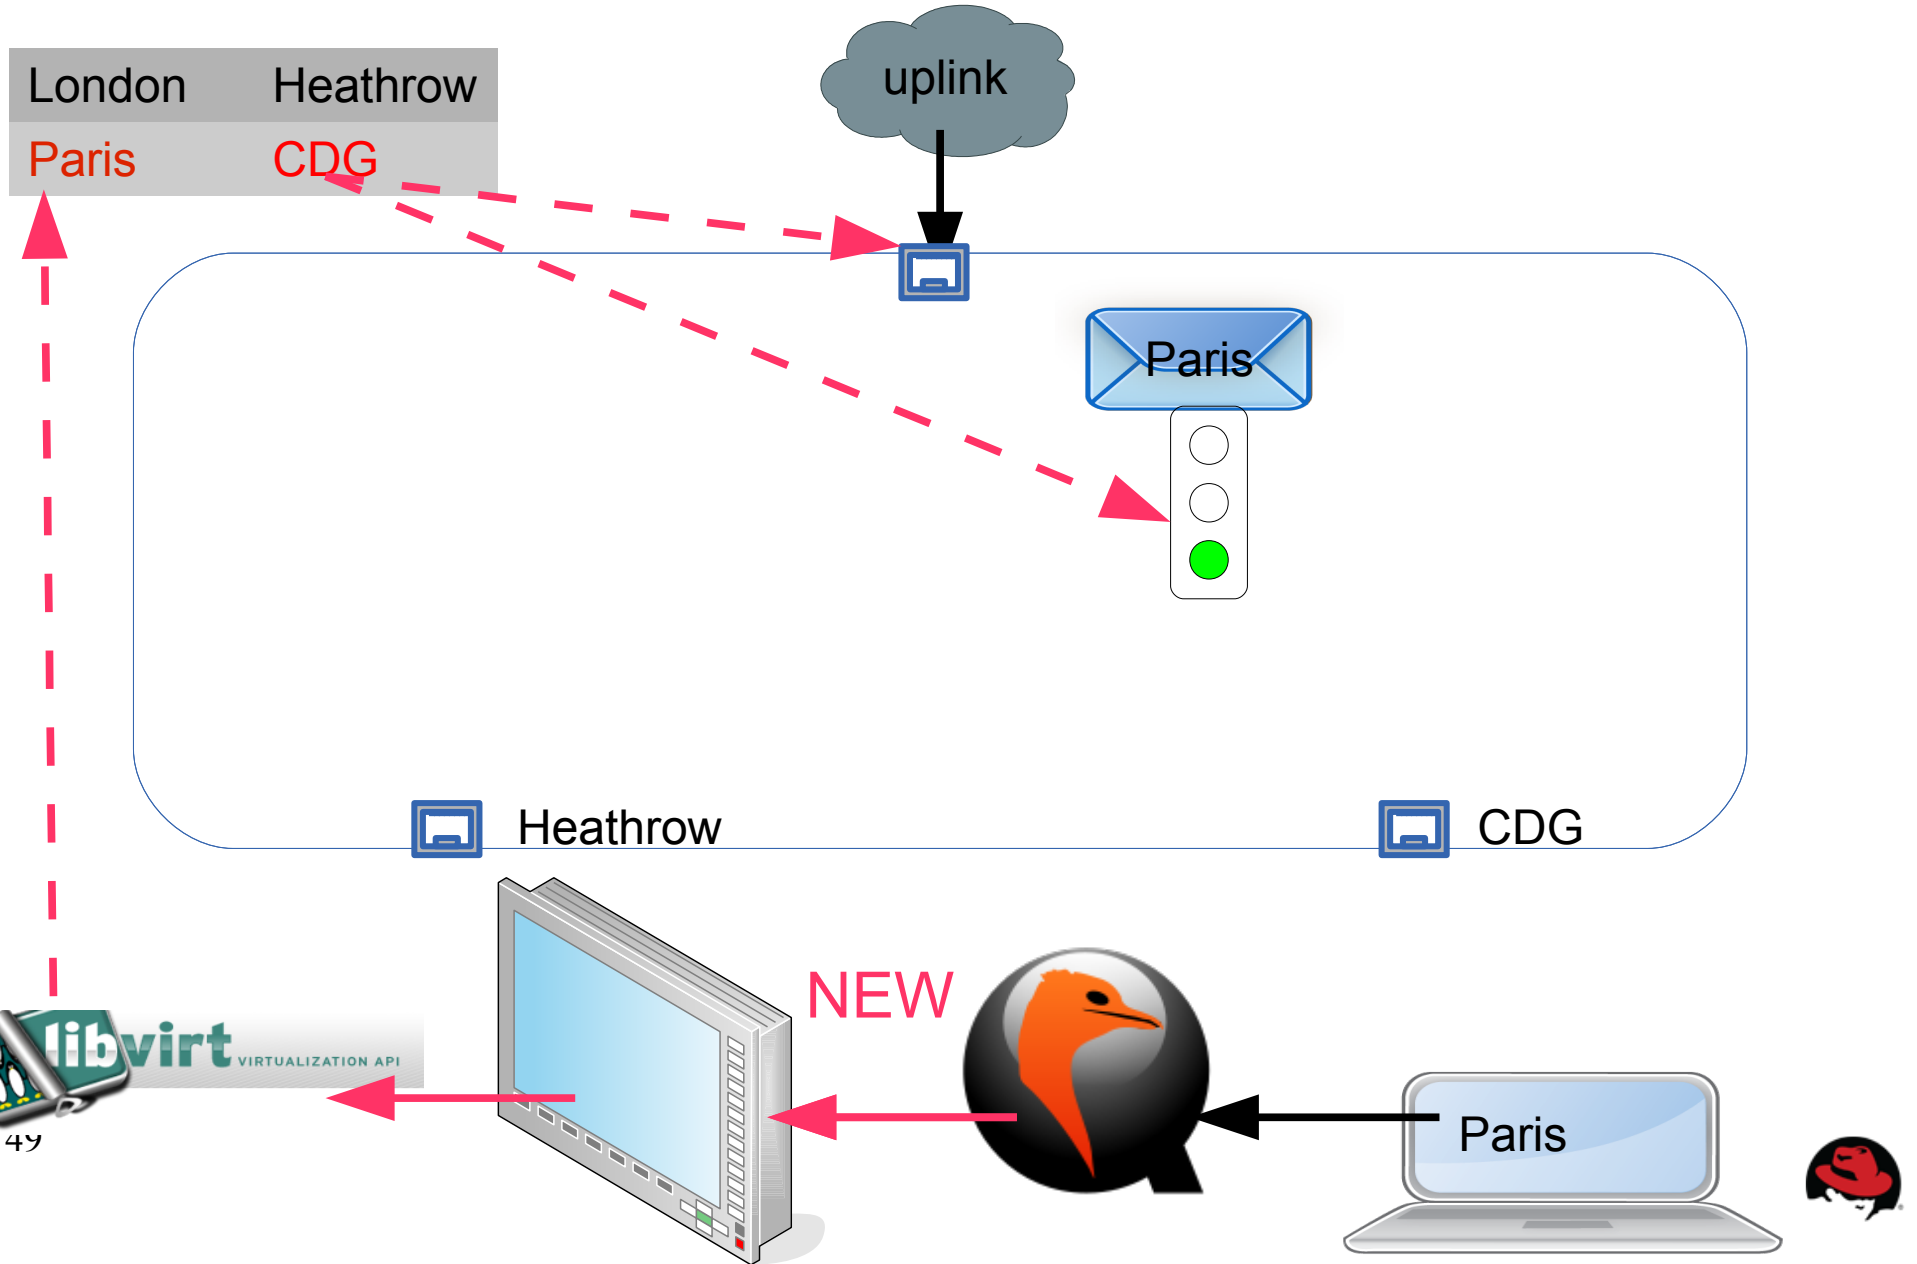

Work in progress

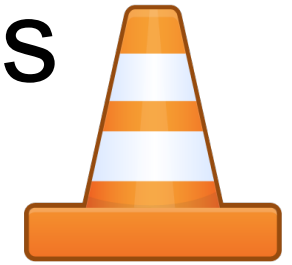

- ELVIS (vhost blk/vhost net)
- Virgl
- Vhost-net performance

RX latency

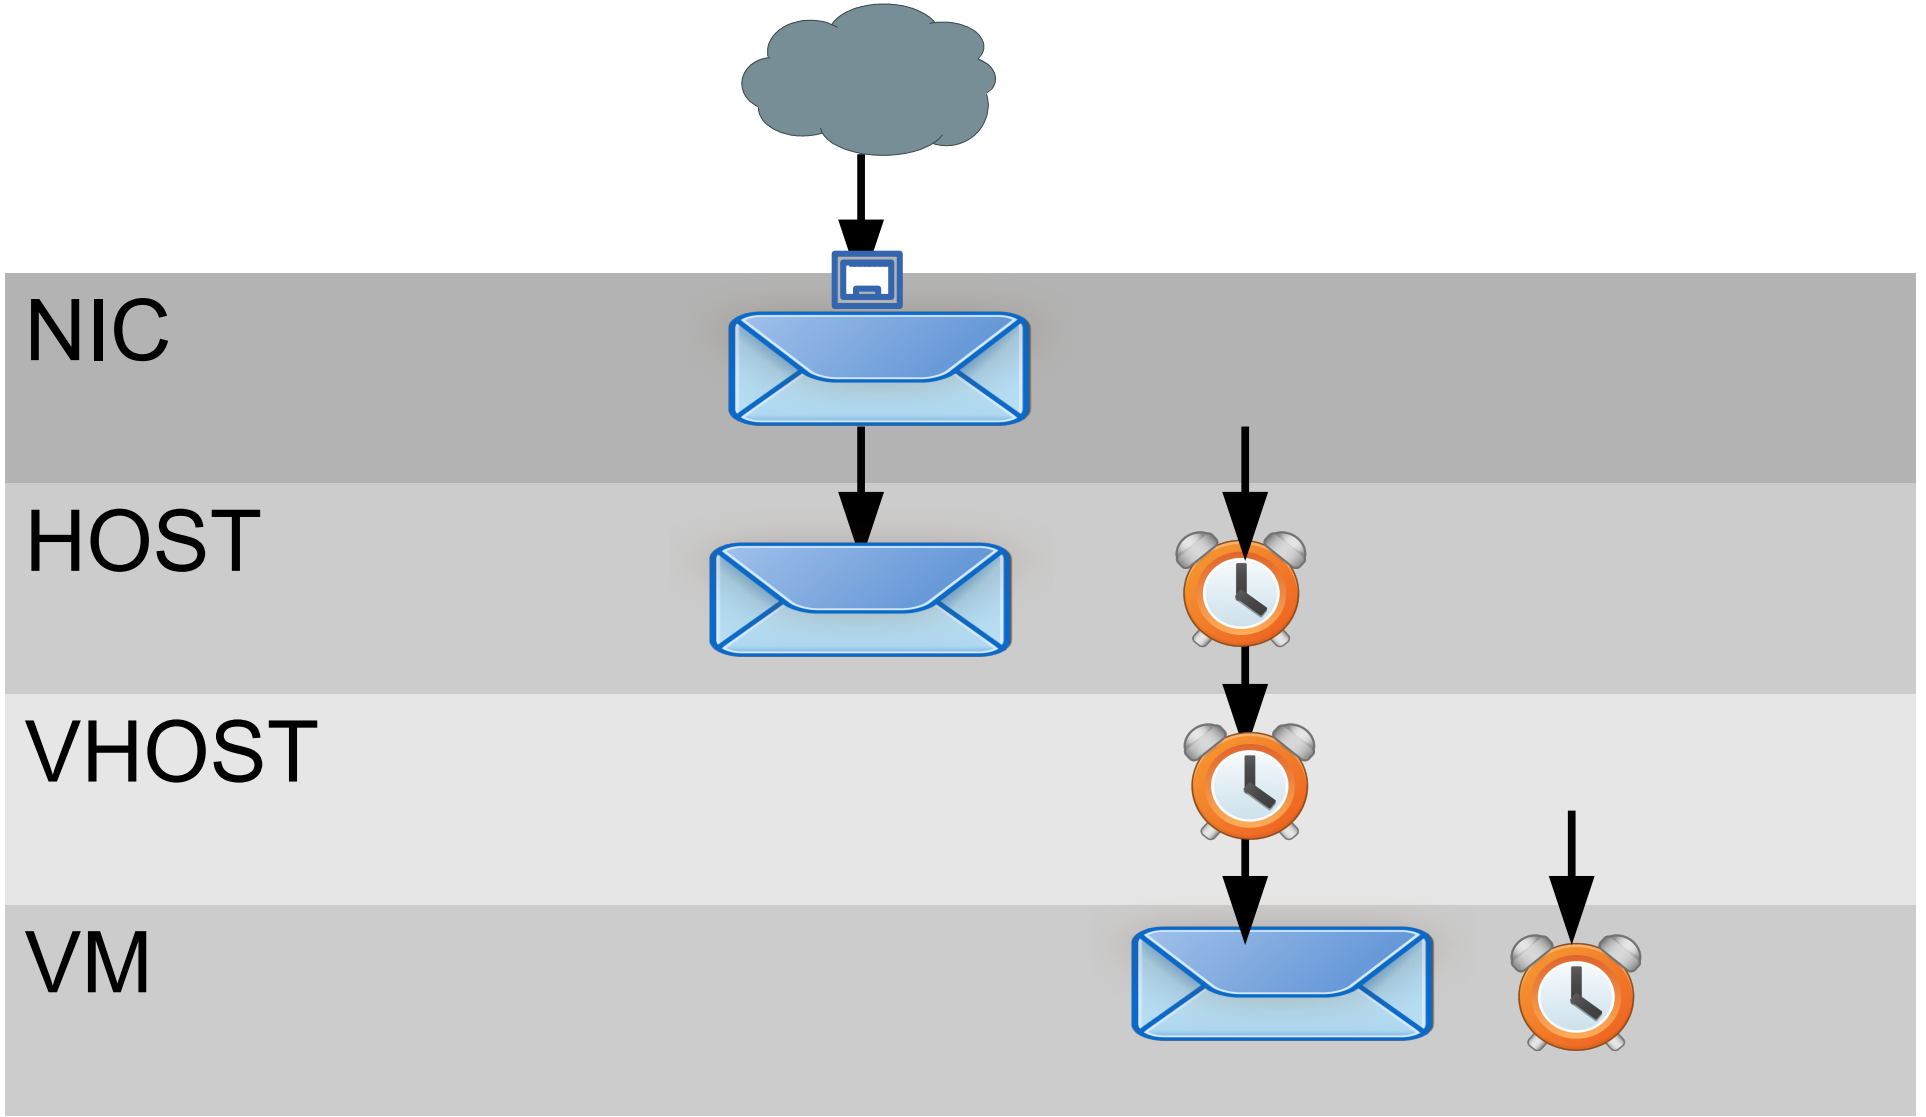

Fast rx

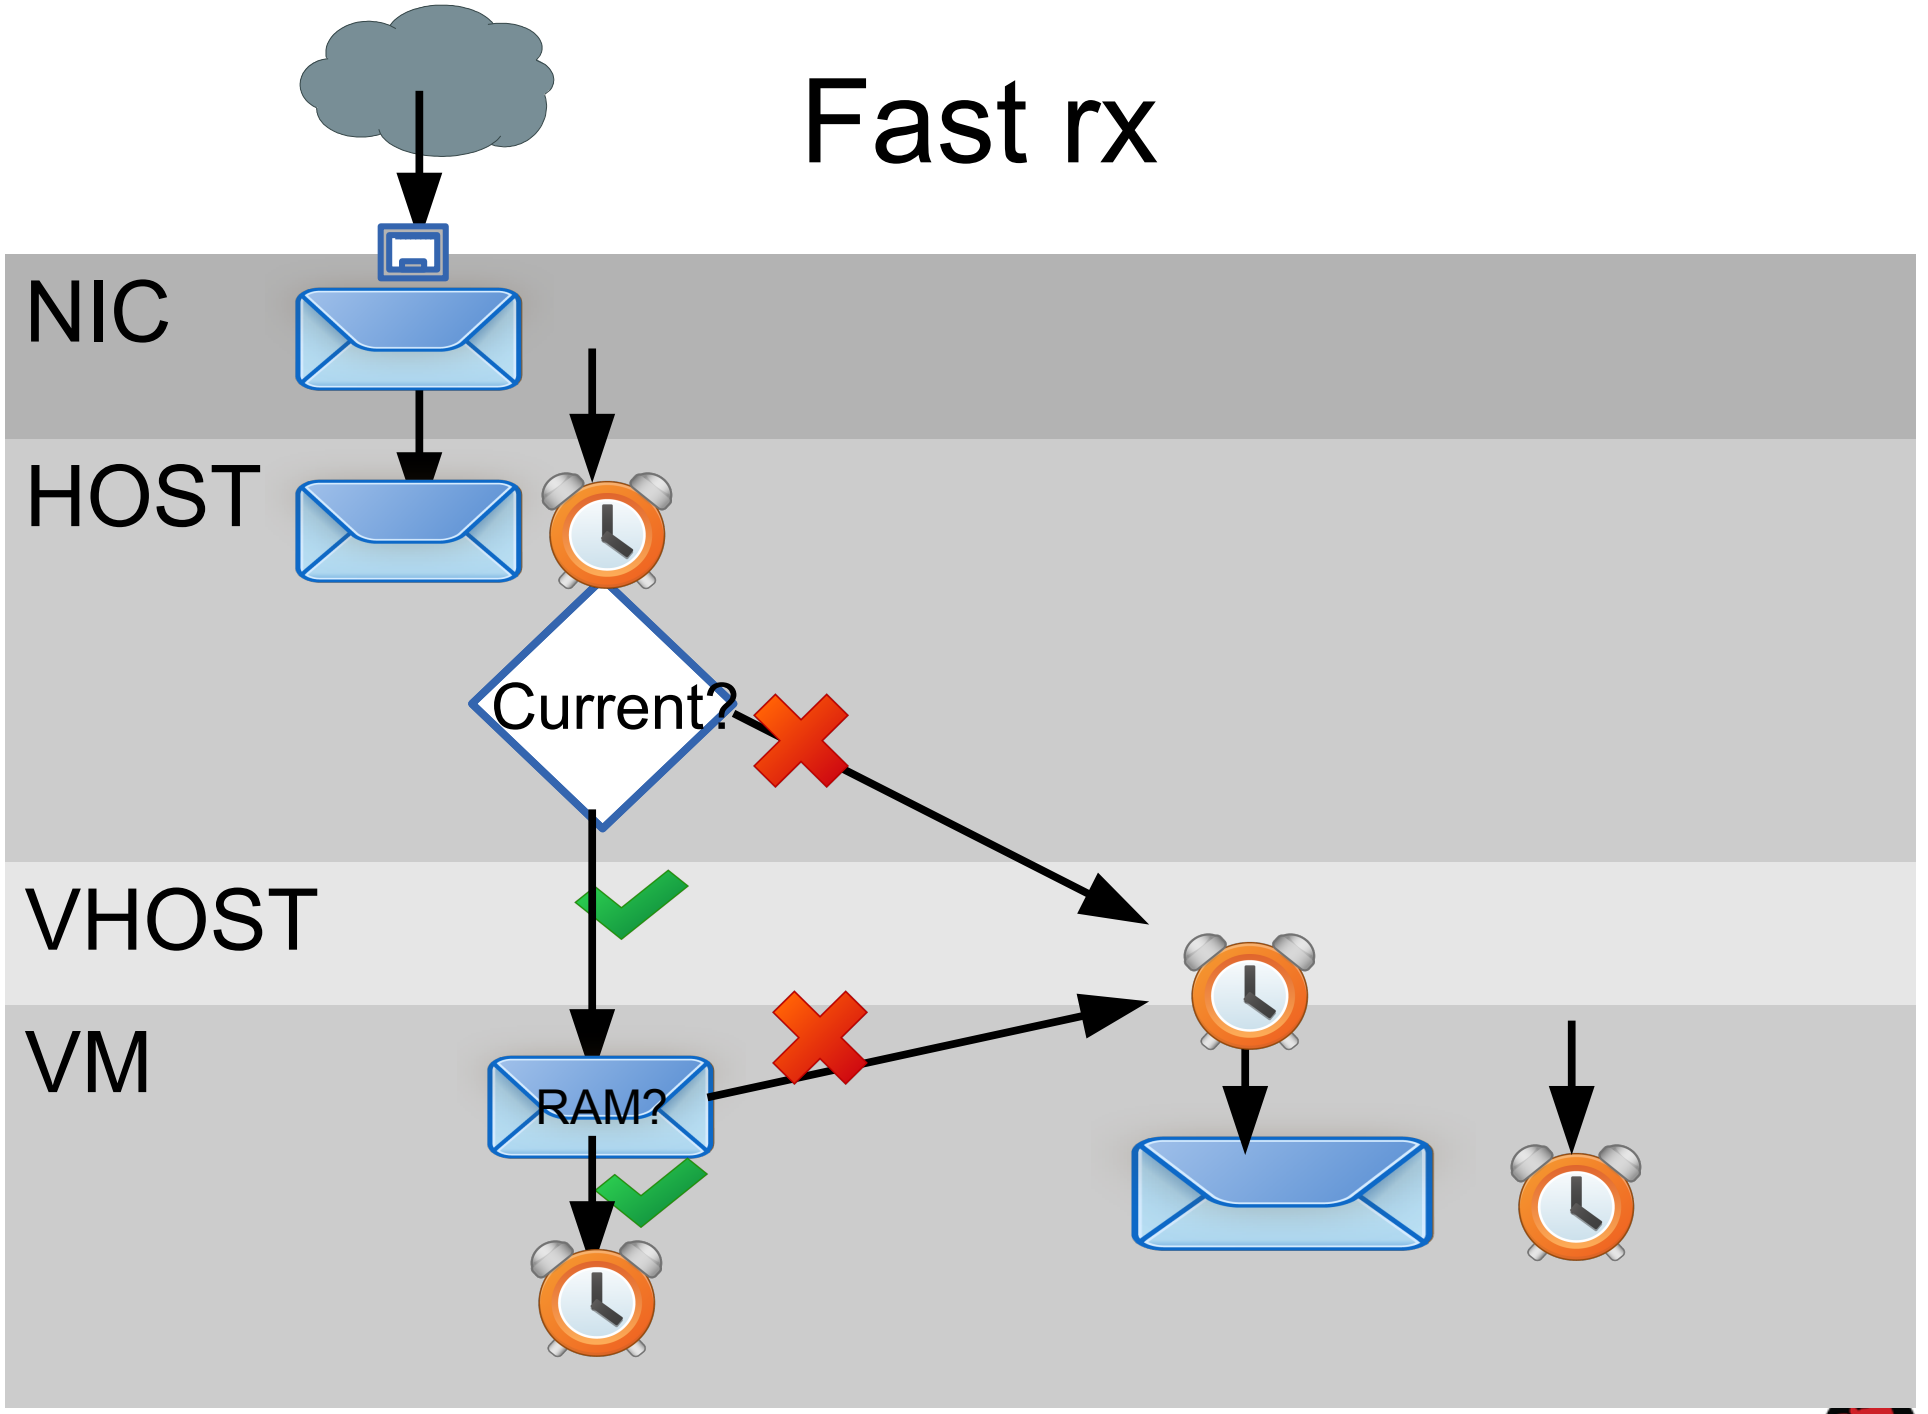

Fast rx: transactions per sec (higher is better)

Hit	331668
Miss	79

Vhost-net threading

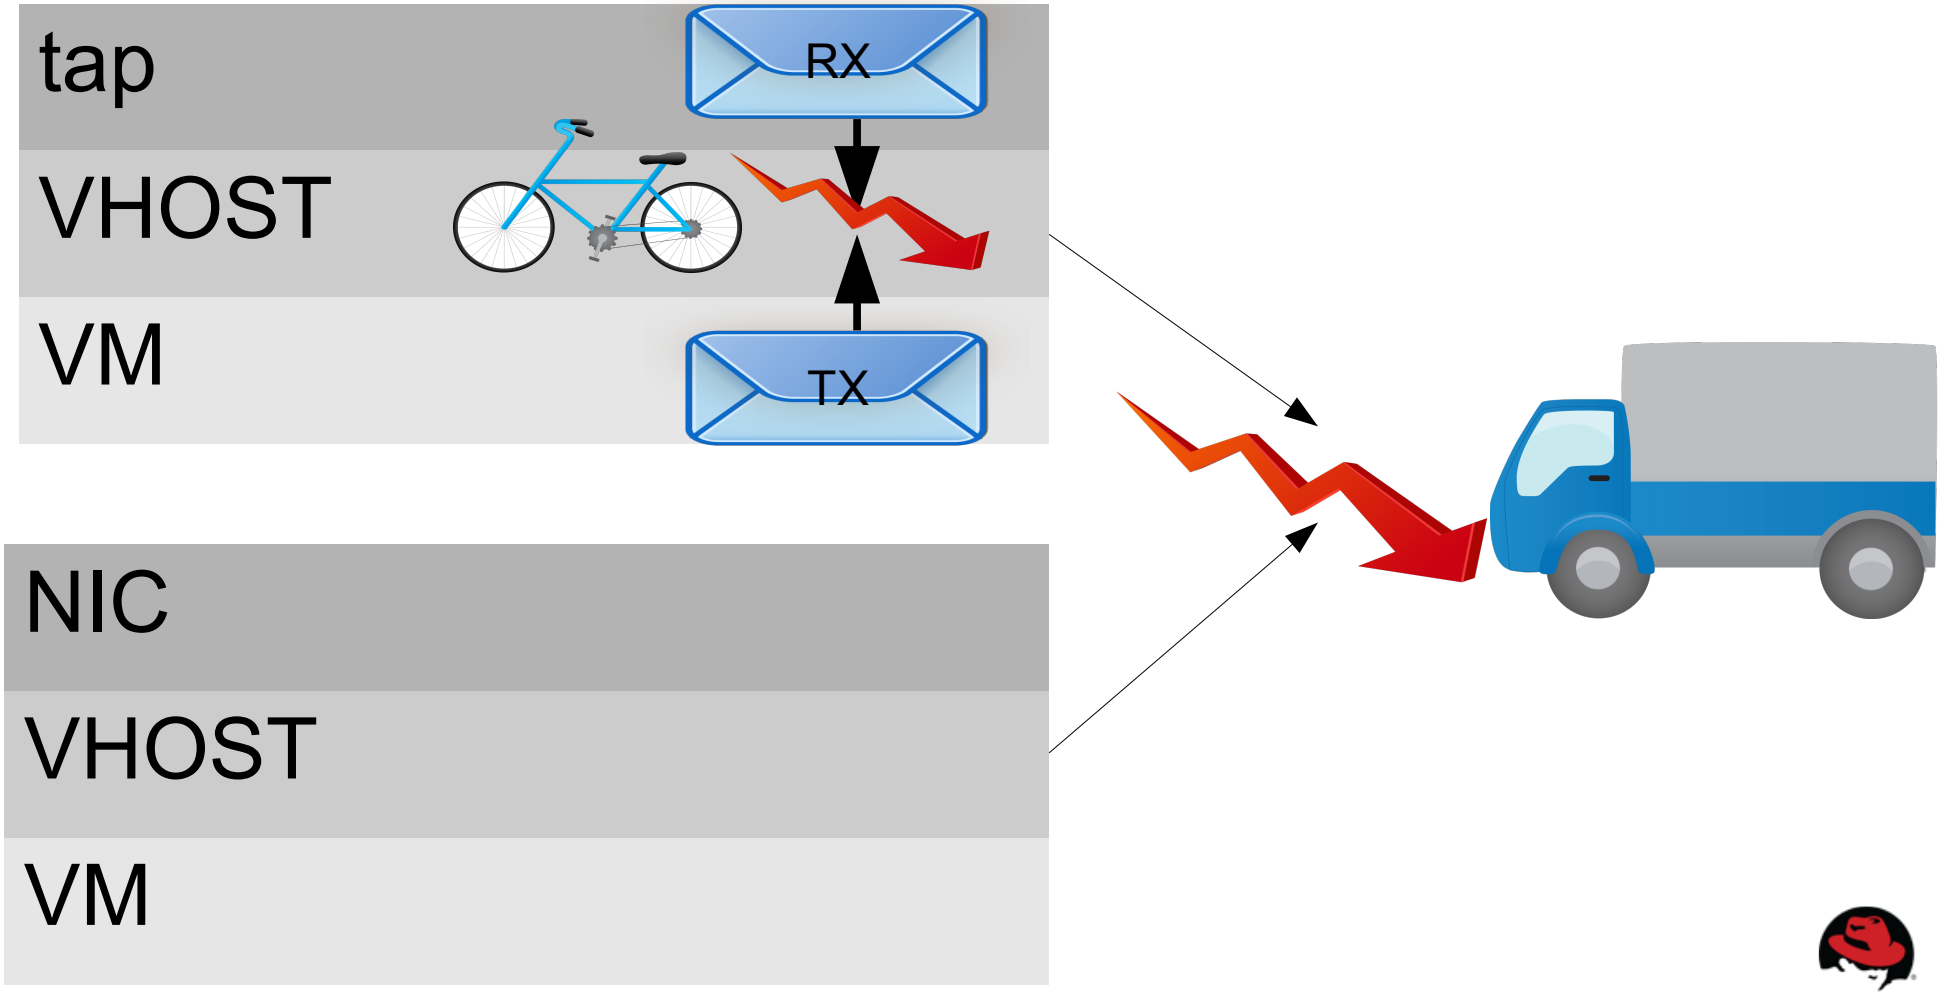

Vhost-net thread pool

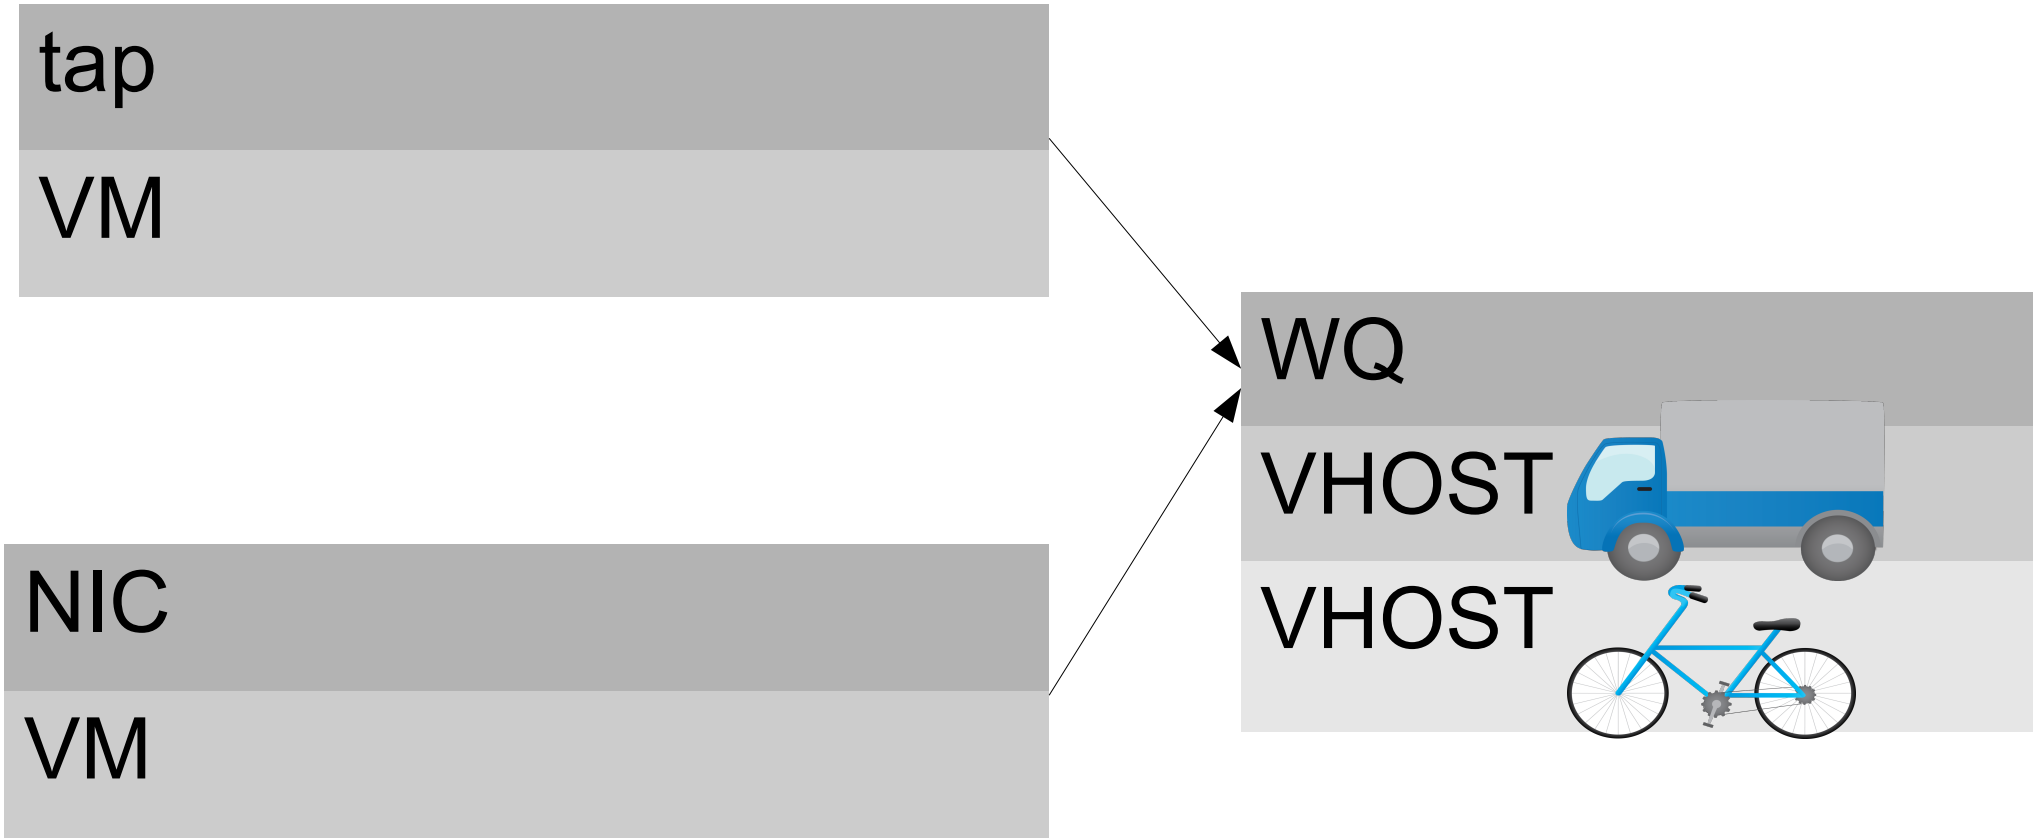

threading: UDP RR transactions/sec (higher is better)

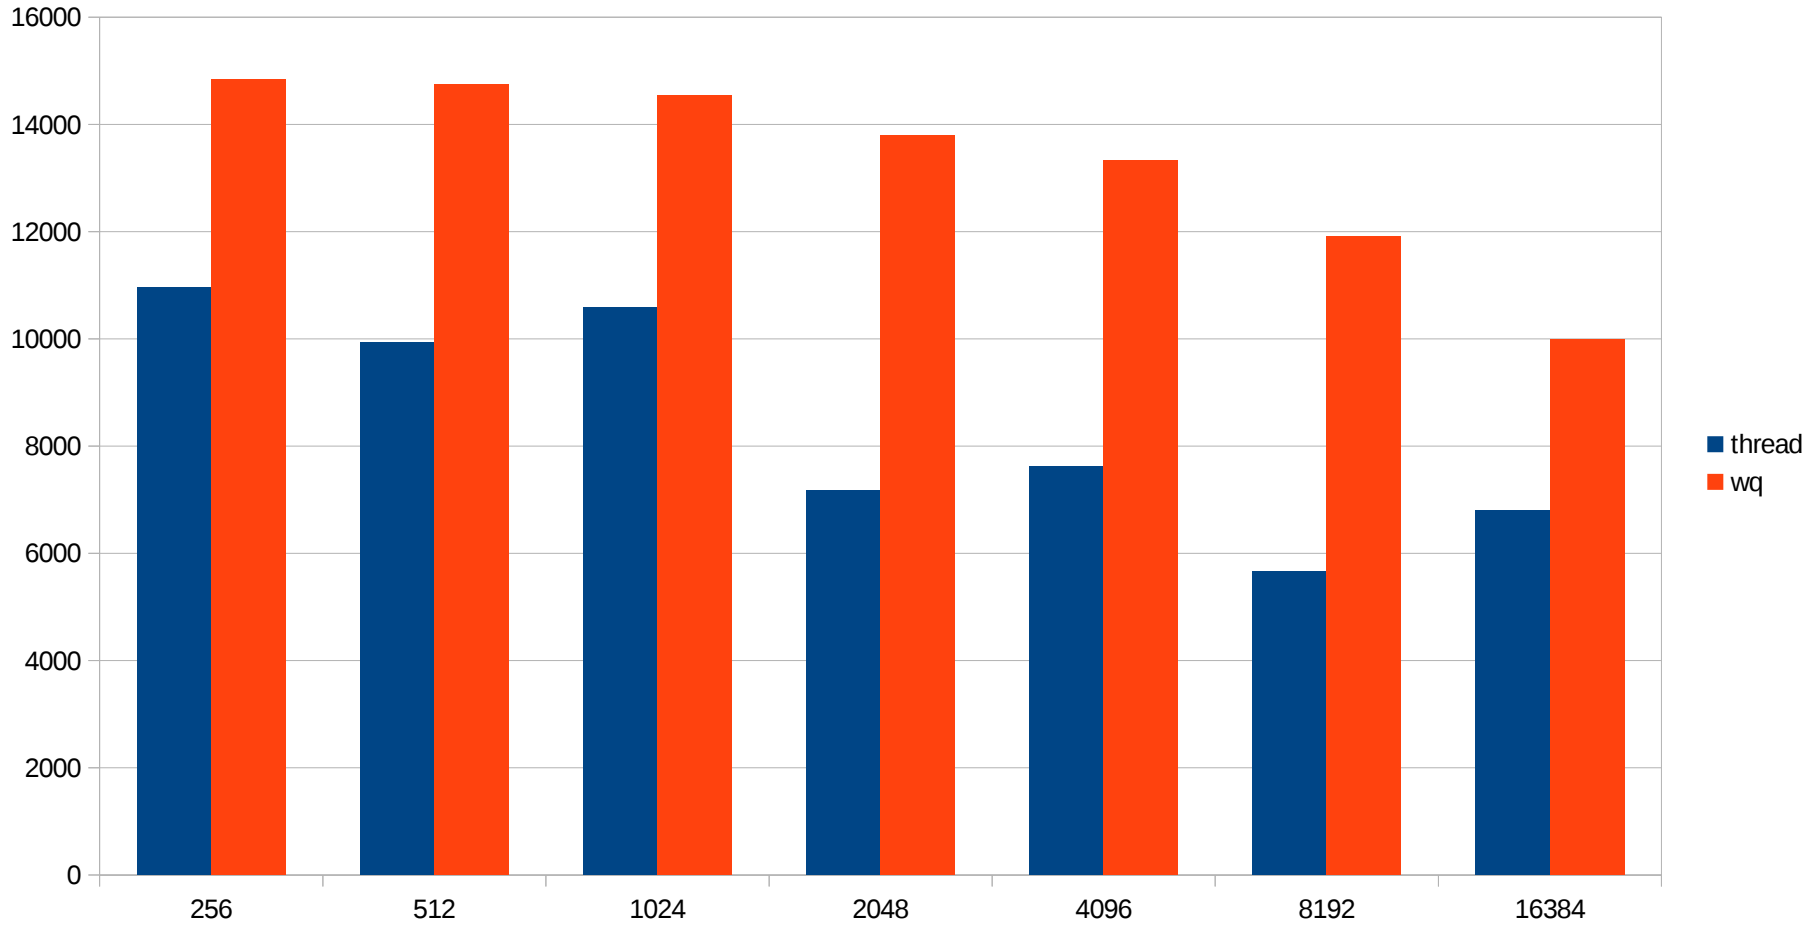

threading: TCP STREAM transactions/sec (higher is better)

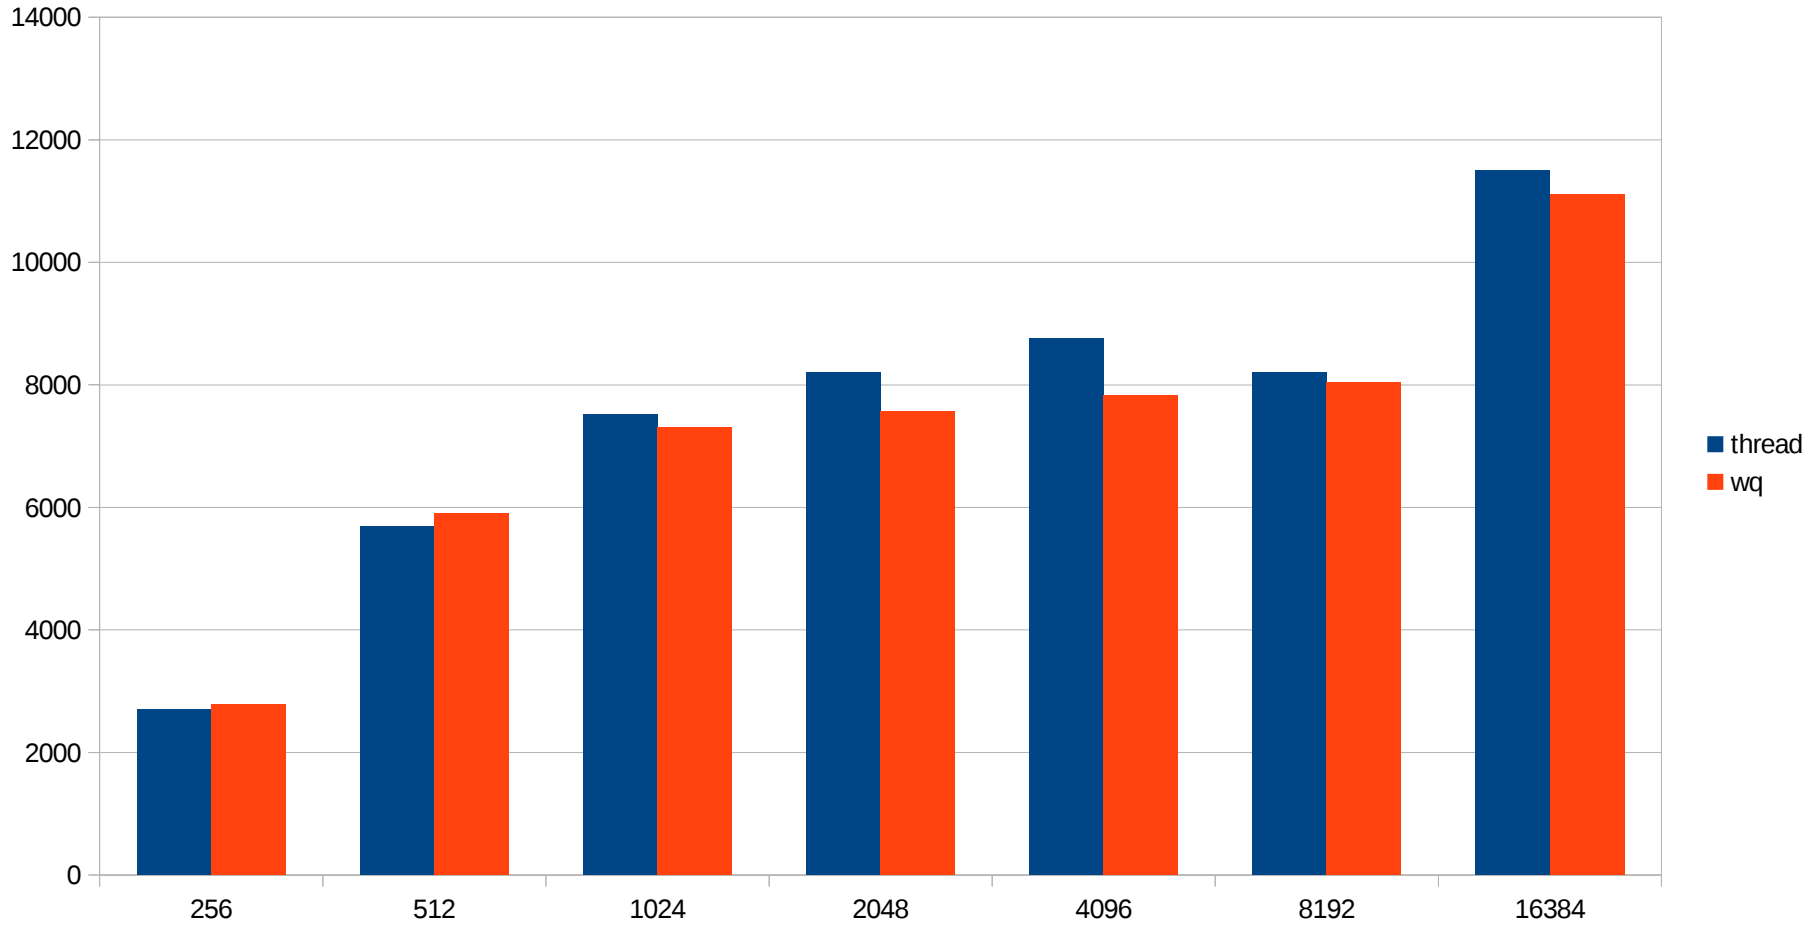

summary

- Performance
- Manageability
- Security

Questions?

OVS: flow match

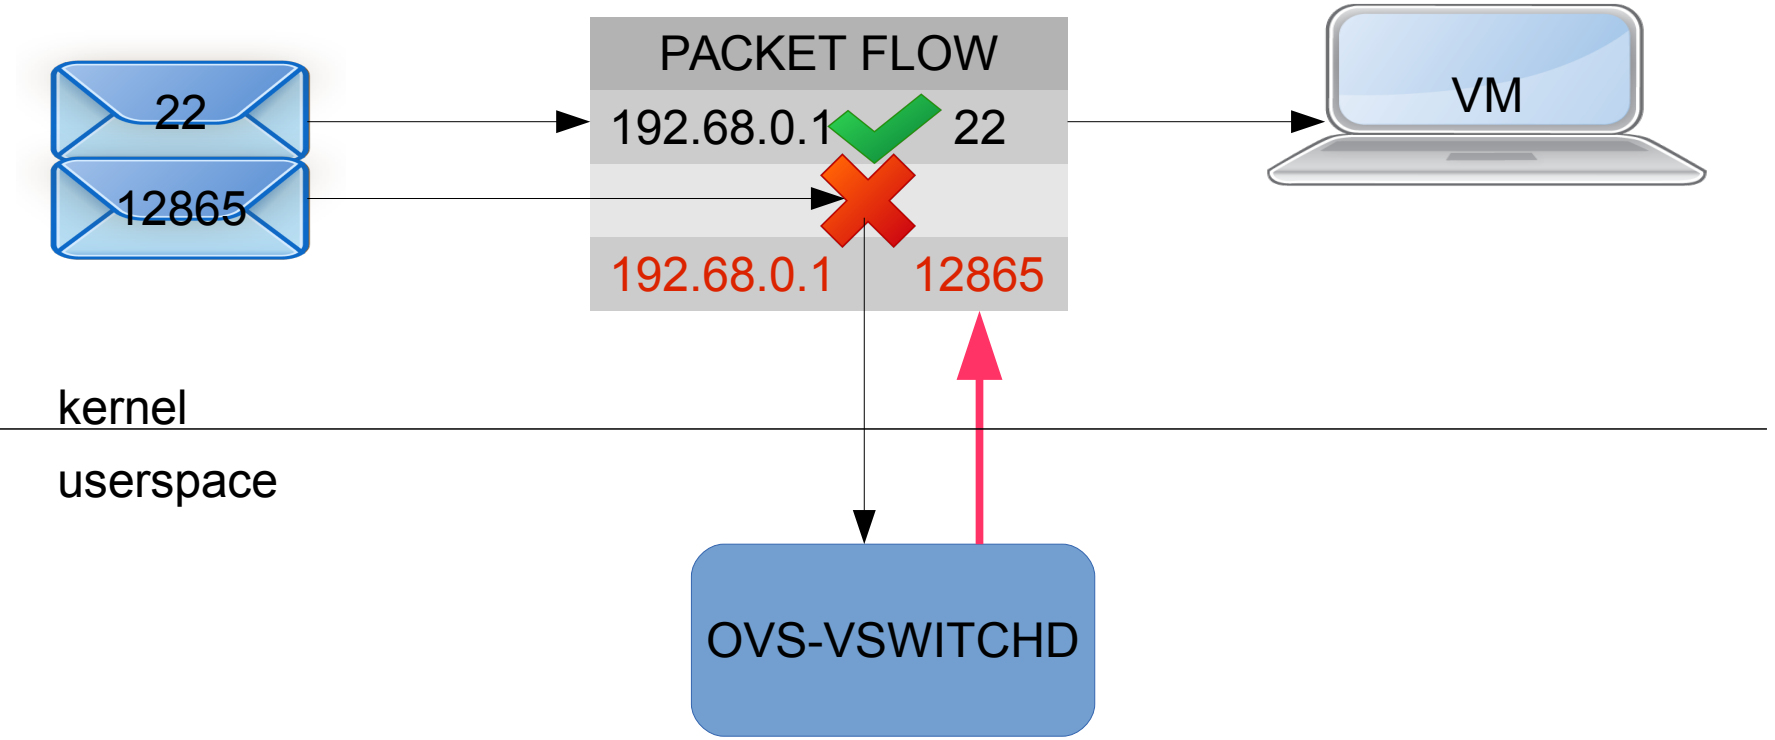

OVS: wildcard match

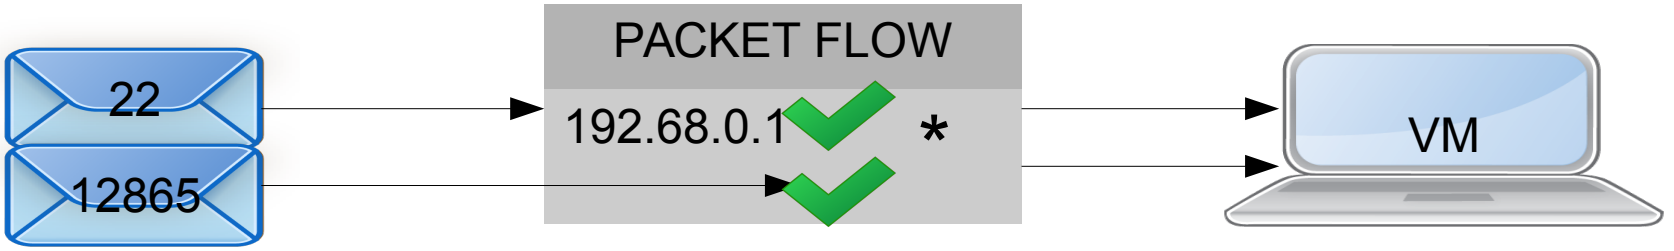

kernel

userspace

OVS-VSWITCHD

Wilcard: netperf CRR (higher is better)

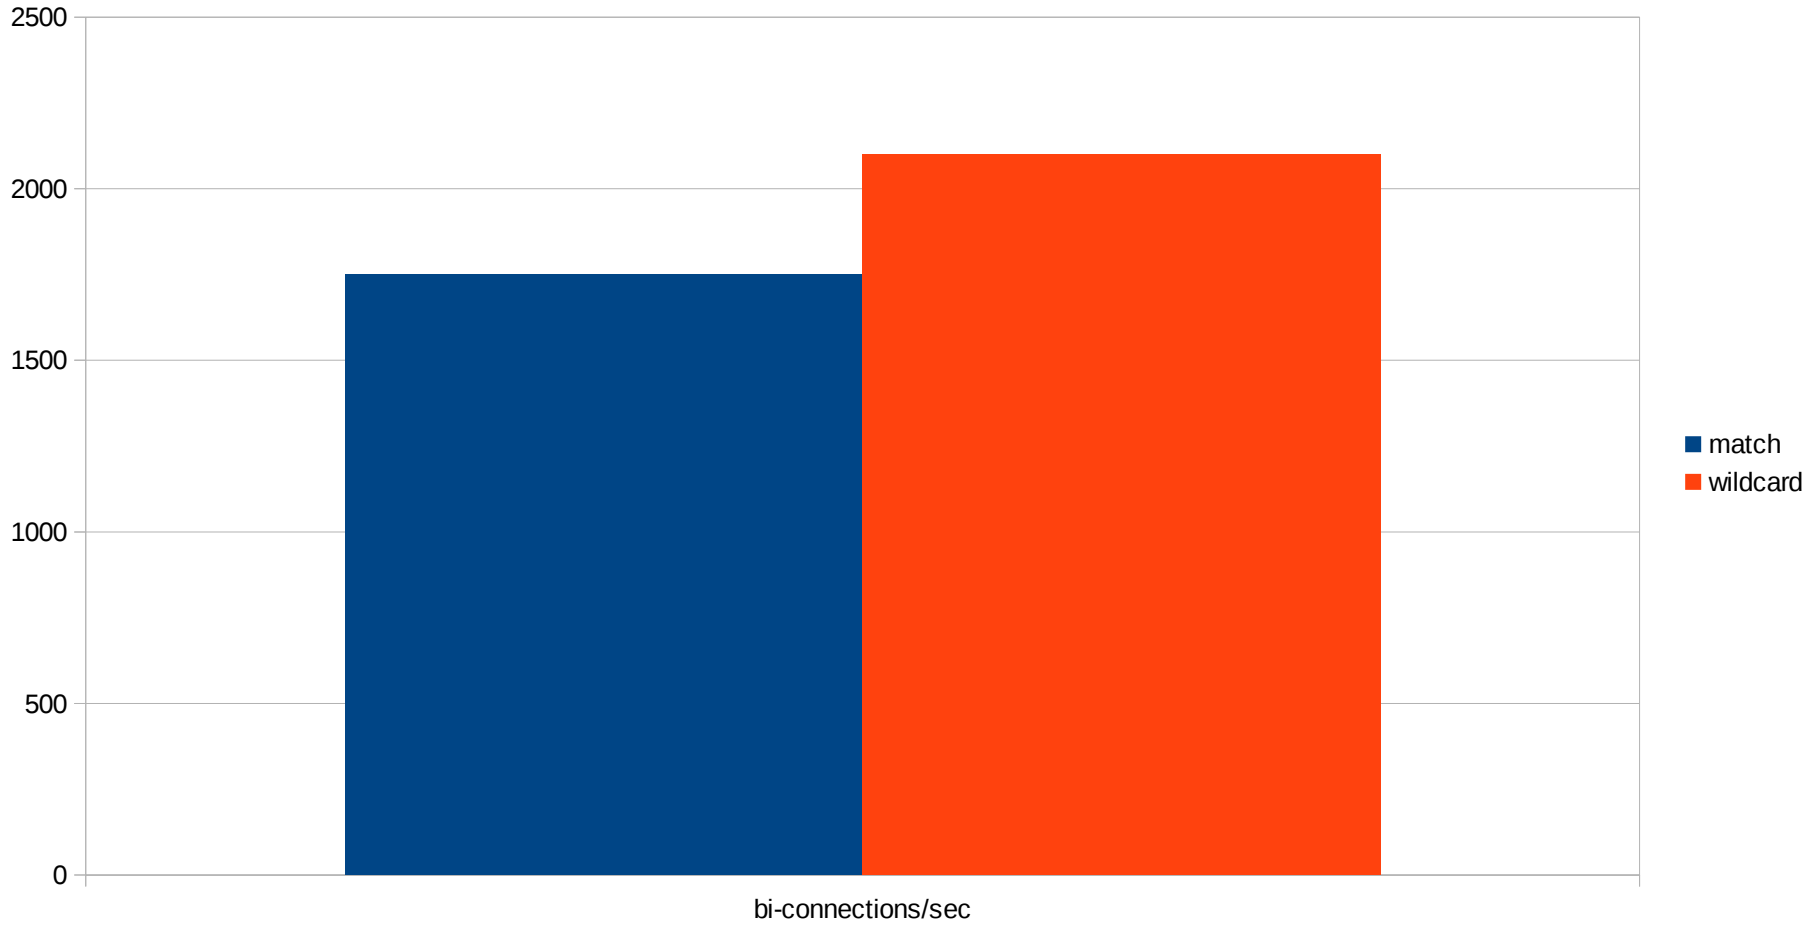


libvirt
VIRTUALIZATION API