Intent Driven Networking with OVS and Vendor Neutral Hardware Offload

Thomas Graf
Noiro Networks, Cisco
OVS Conference 2014

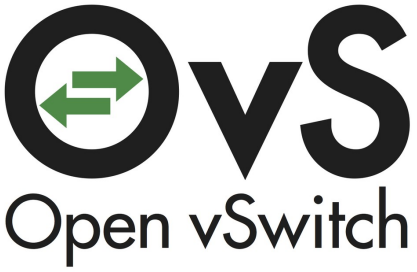

Agenda

Context

- User intent based policy
- Policy driven networking

OVS hardware offload

- Reaching ludicrous+1 speed
- SDN + SR-IOV is hard
- A model that could work

Legacy Forwarding Elements

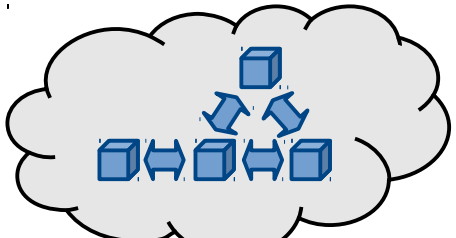

Developer

Service Ticket

NetOps

Hours to Weeks

Vendor UI

SDN solved everything, right?

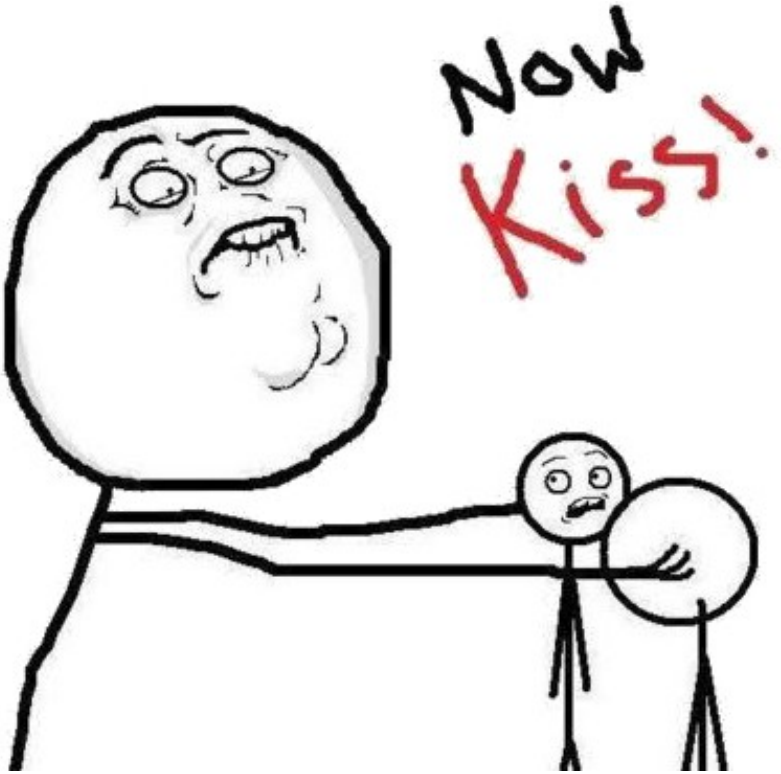

Application developers don't care about 802.1Q, VXLAN, switching or routing.

Group Based Policy

Tell me *"what you want"* instead of *"how you want it implemented"*.

Contract

Dir	Condition	Action
In	Port 8080	Service Chain = { LB }
Out	Related	Accept

consumes

provides

WAN Group

IP 0.0.0.0/0

App Group

VM 10.0.0.1

VM 10.0.0.2

LXC 10.0.0.3

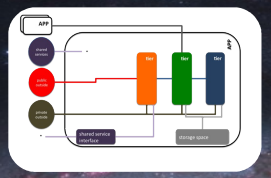

Intent

Governance

Operations Constraints

Capabilities and State

Intent Driven Networking

- How (L2/L3) a packet gets from A to B doesn't matter ...
  - ... but it should only get there if it's allowed to
  - ... and it should get there fast

Intent Based Edge Intelligence

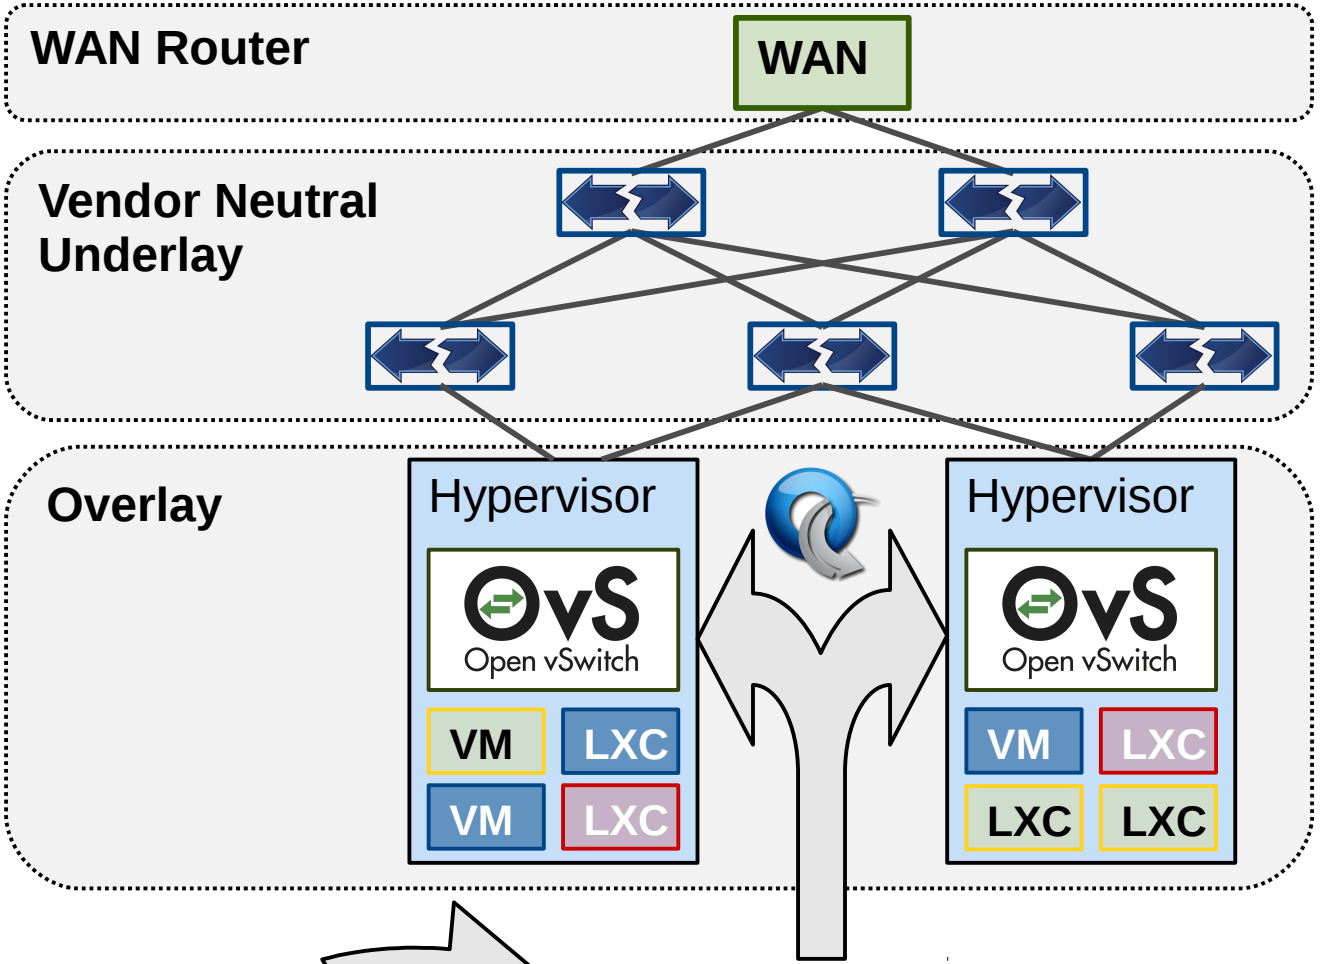

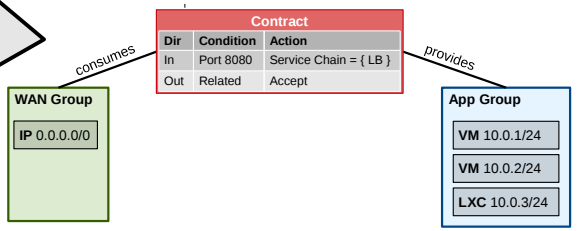

Awesome!

What about throughput and latency?

$(LUDICROUS_SPEED + 1)

Can we leverage SR-IOV?

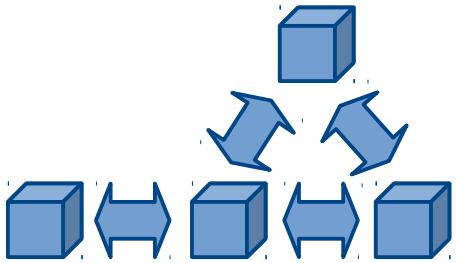

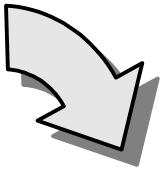

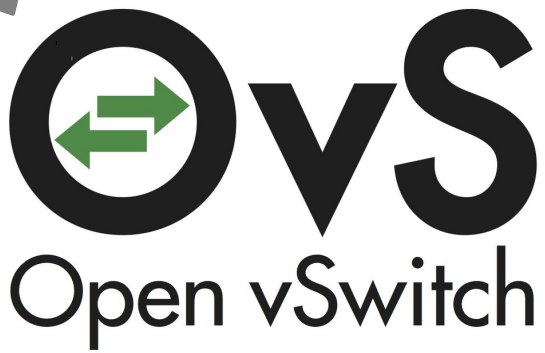

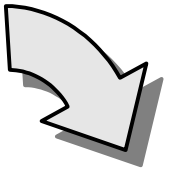

SR-IOV

SDN + SR-IOV is hard

- Almost anything is possible with software
  - That's why OVS rocks
- Constant feature gap
  - Software release cycle < Hardware release cycle
- SR-IOV resources are not unlimited
  - # VFs, # queues, limited flow setup rate

The Software Fallback

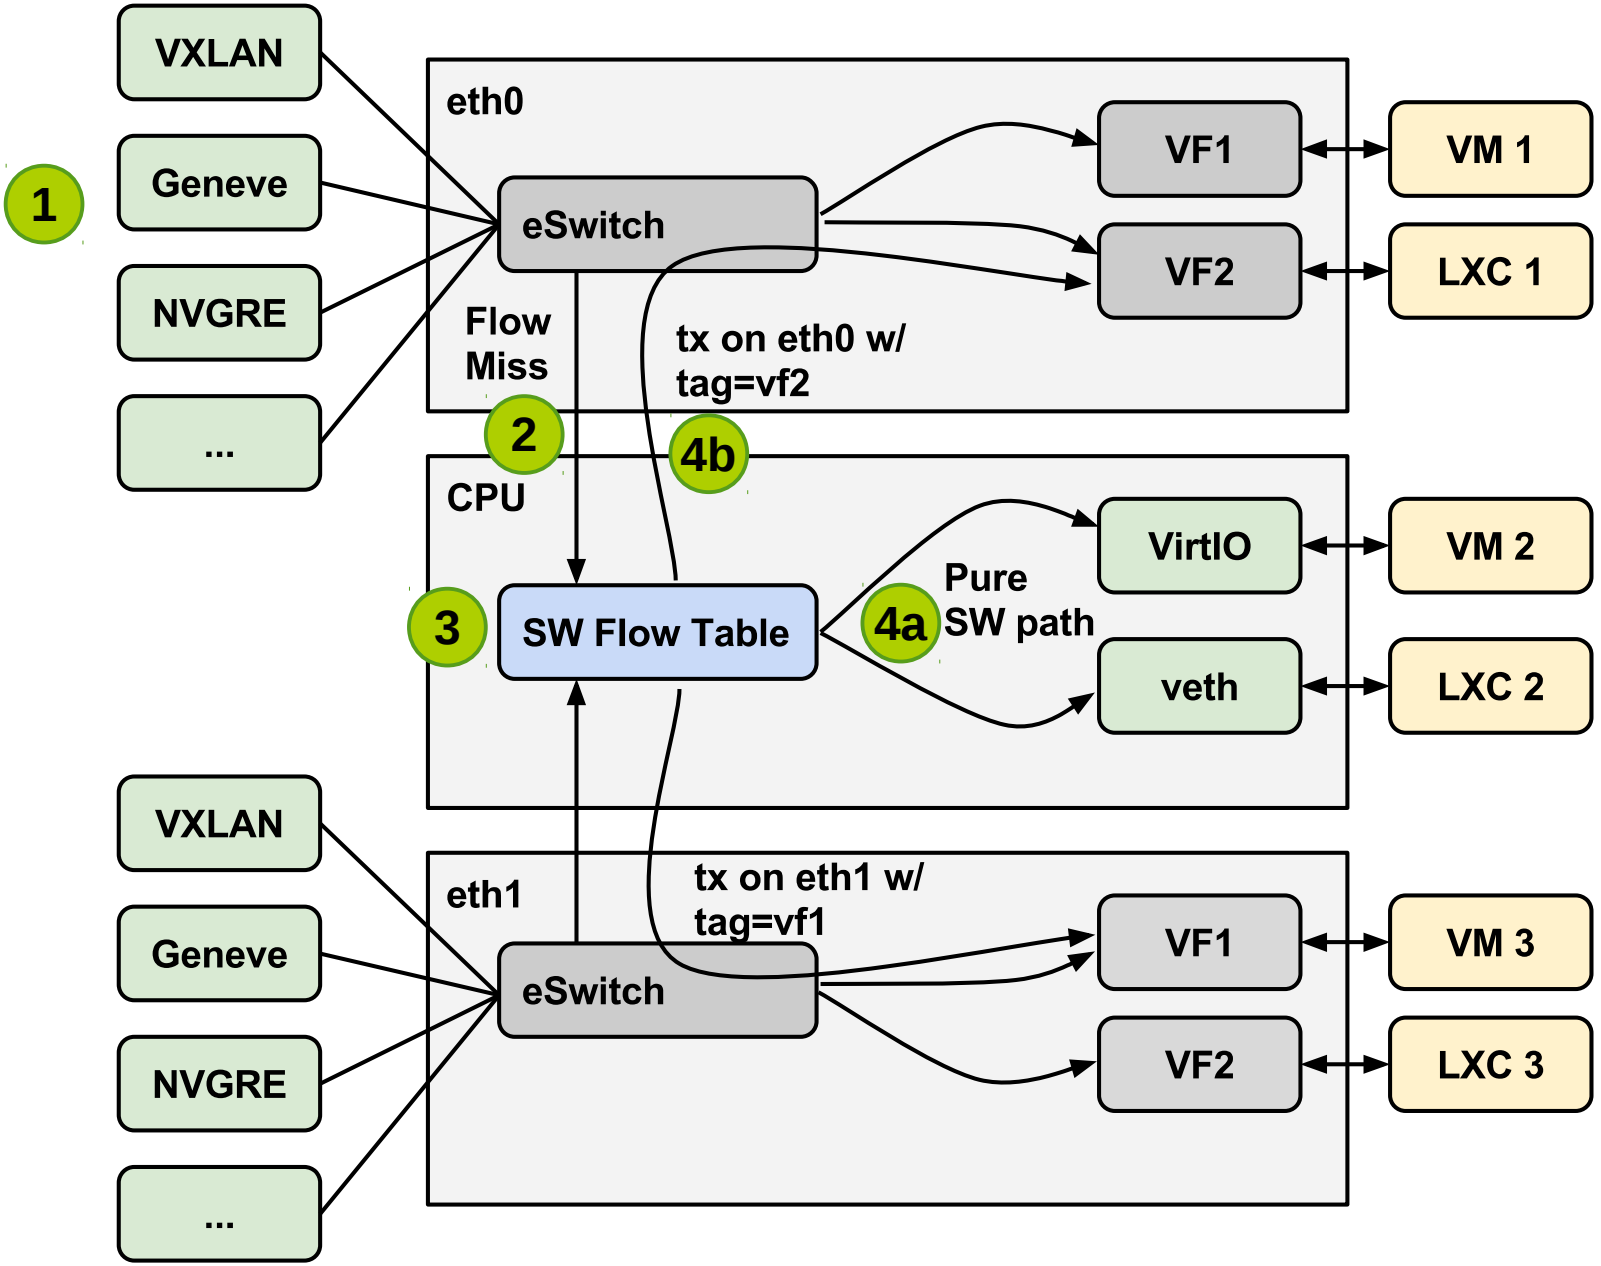

Taking the networking approach

Adding layers will solve any problem.

(c) Karen Sagovac

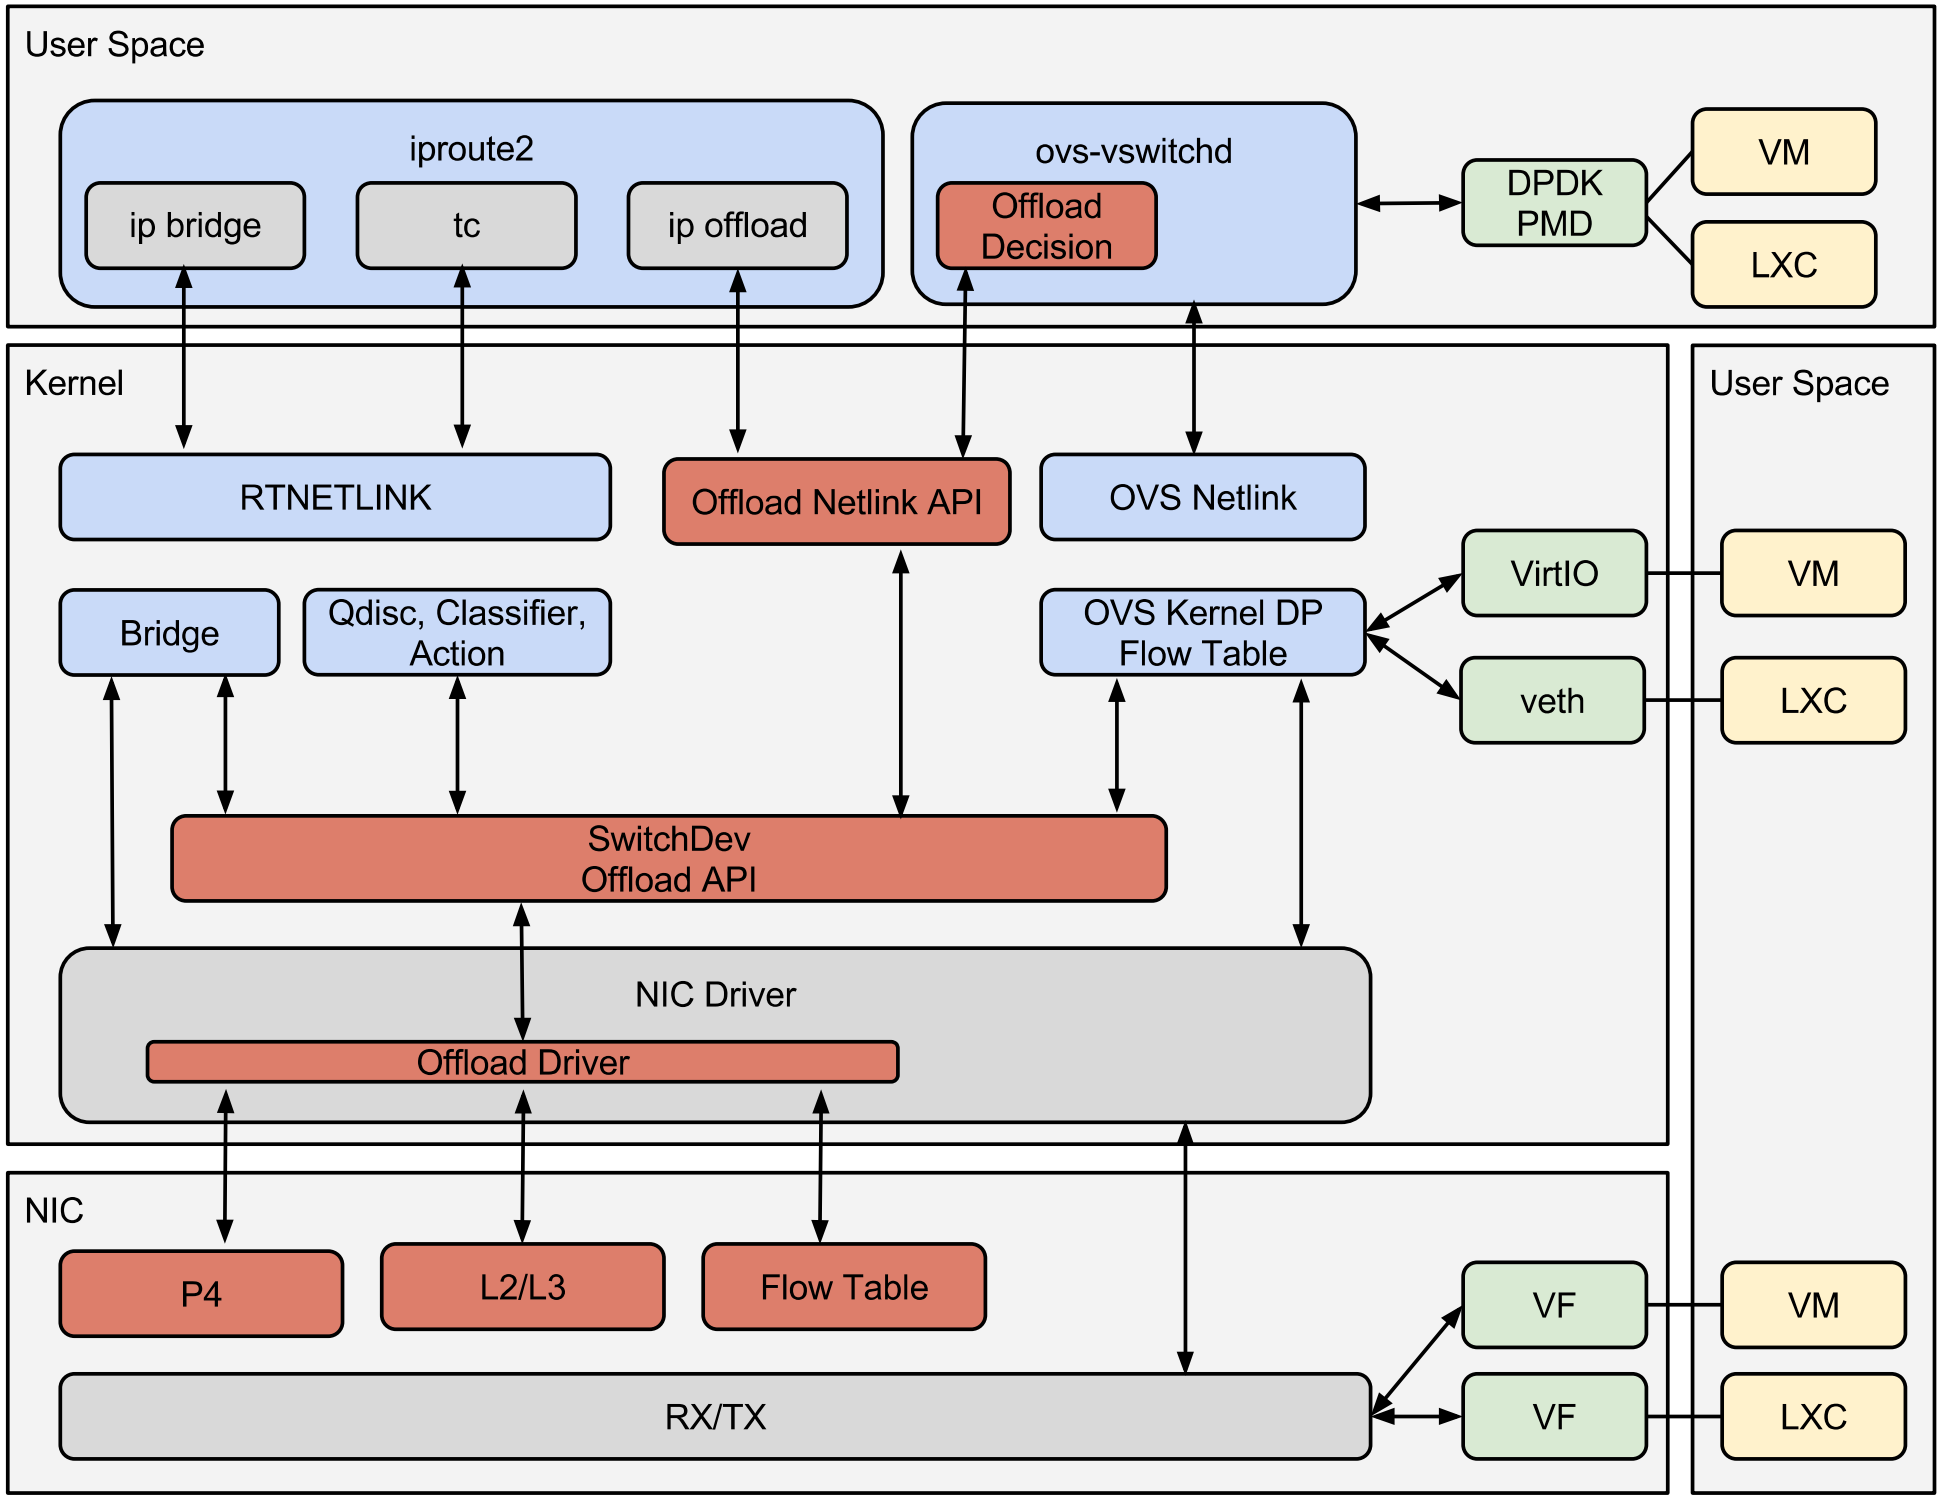
User Space
iproute2
ip bridge
tc
ip offload
ovs-vswitchd
Offload Decision
DPDK PMD
VM
LXC
Kernel
RTNETLINK
Offload Netlink API
OVS Netlink
Bridge
Qdisc, Classifier, Action
OVS Kernel DP Flow Table
SwitchDev Offload API
NIC Driver
Offload Driver
User Space
VirtIO
VM
veth
LXC
NIC
P4
L2/L3
Flow Table
RX/TX
VF
VM
VF
LXC

Q&A

Interested in helping out?

Join the community fun on the netdev and ovs-dev mailing lists

Contact

tgraf@noironetworks.com or @tgraf__