



# DPDK

DATA PLANE DEVELOPMENT KIT

## Reflections on Mirroring with DPDK

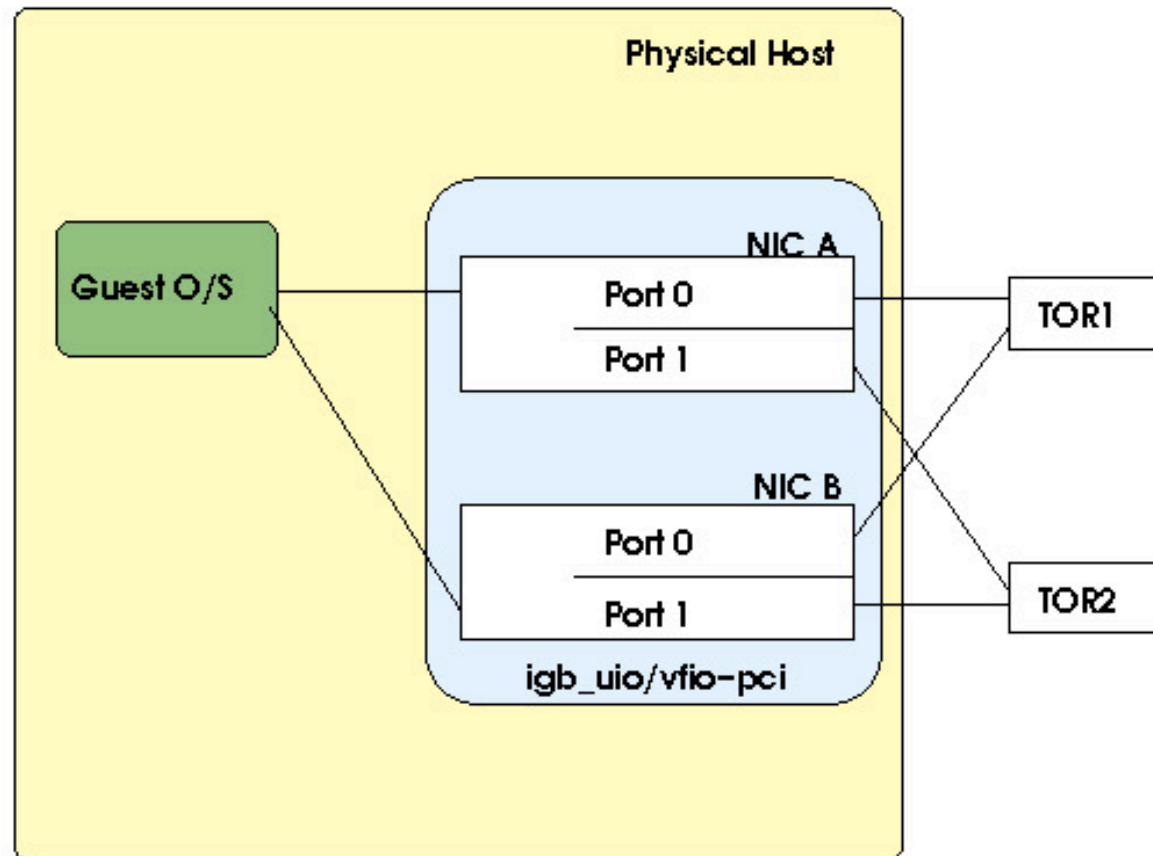
E. Scott Daniels,  
Kaustubh Joshi,  
Alex Zelezniak  
AT&T Labs -- Research

DPDK Summit - San Jose – 2017





- Debugging network problems for an SR-IOV guest is hard
- Attaching `igb_uio` or `vfio-pci` to the PF eliminates use of useful tools (e.g. `tcpdump`)
- Further complicated by network redundancy measures

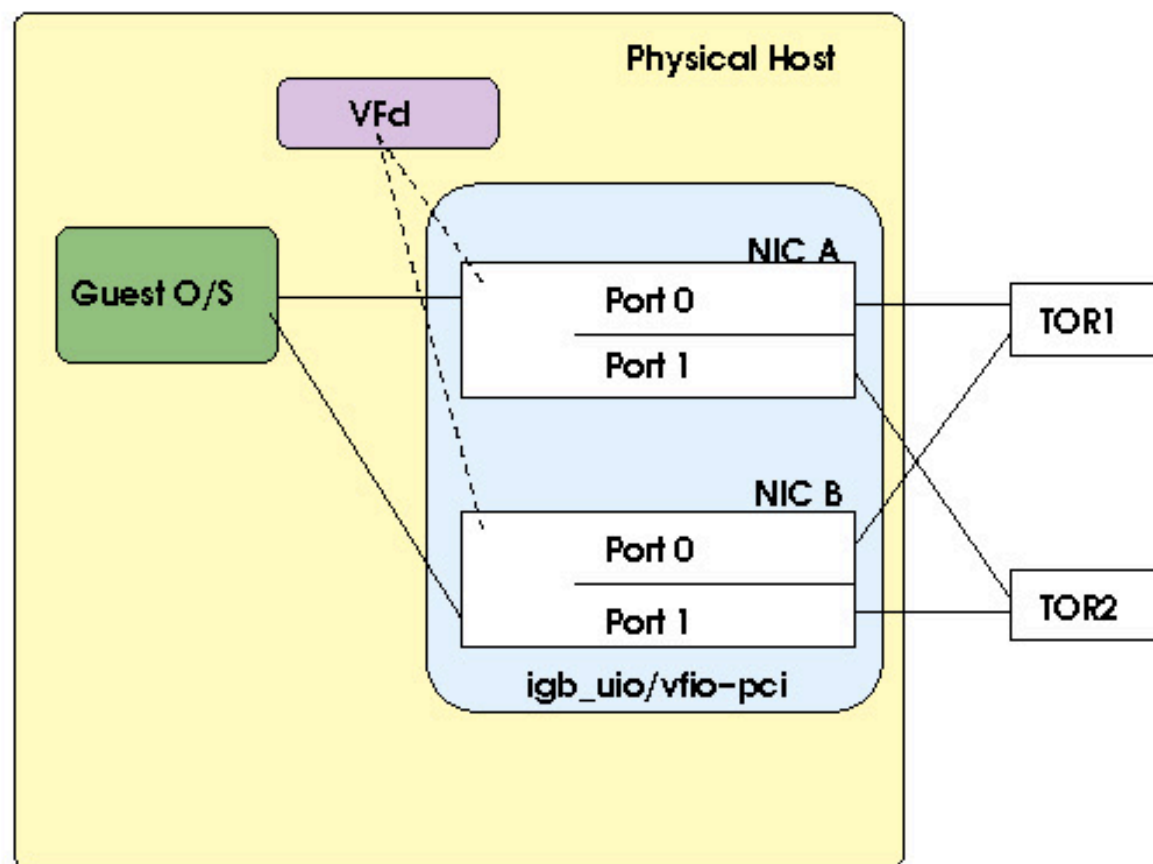


## Why Bind DPDK Driver on Physical Host?



DPDK

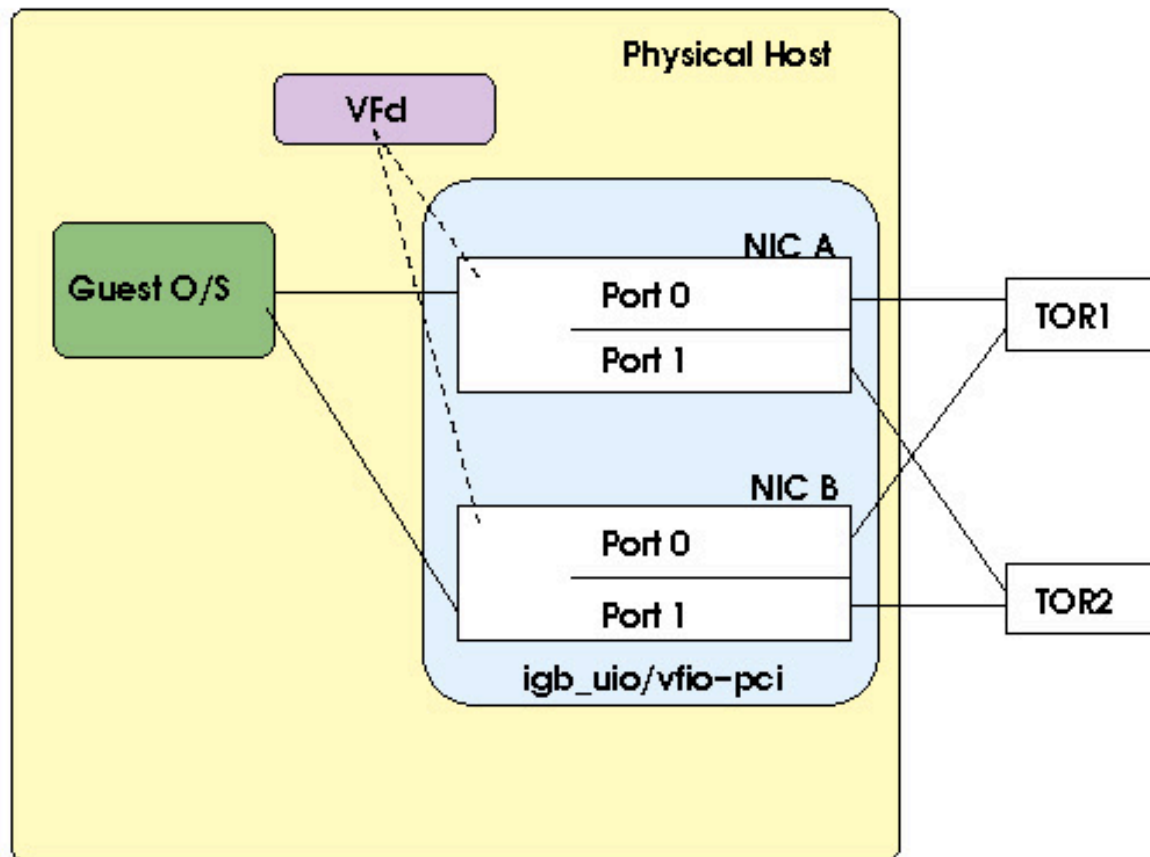
- Want to use **VFd** to provide configuration & implement policy
- VFd:
  - A *Network Hypervisor*
  - A DPDK application
  - VF configuration not available via kernel drivers (e.g. allow VF to change MAC)
  - Enforces local policy (e.g. validate VF supplied VLAN)



## Typical Guest Problem



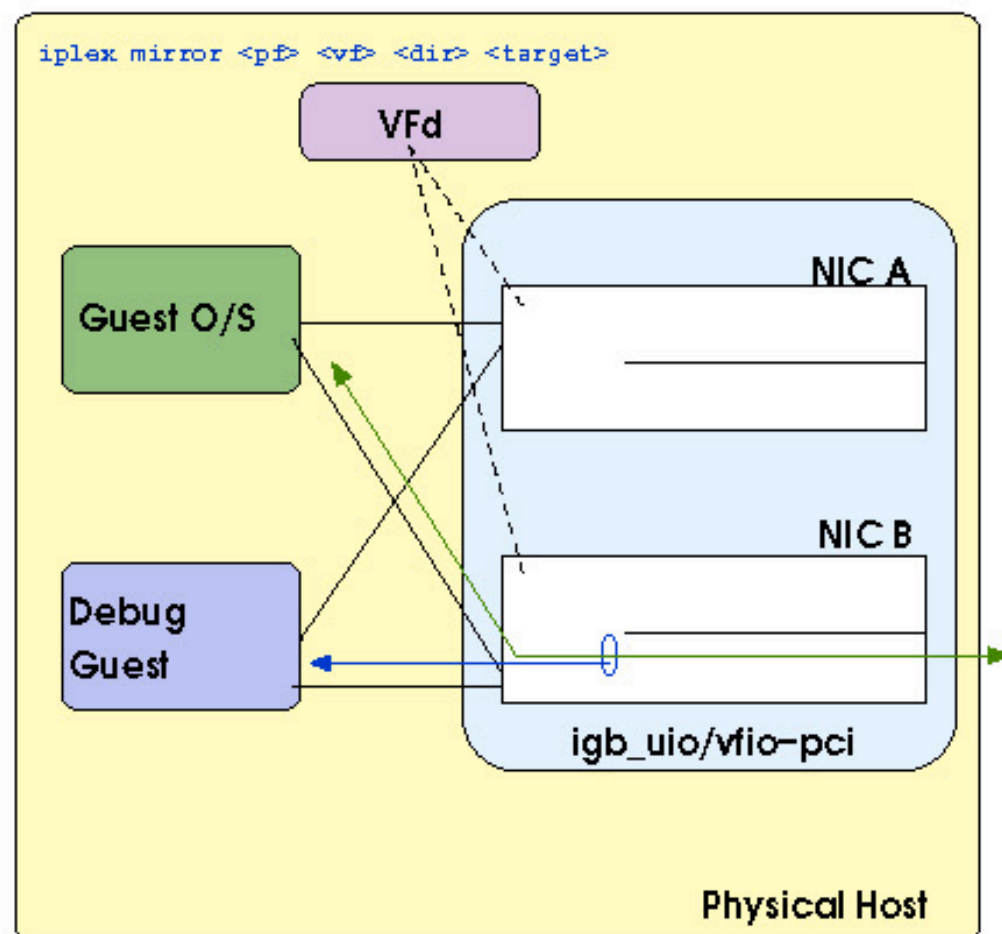
- Sporadic traffic arrives at the guest
- Possible causes:
  - Incorrect VLAN tags on some packets
  - Incorrectly configured VF
  - Incorrectly configured TOR
  - No traffic arriving at all
- Just flip on tcpdump, right?



## Typical Guest Problem -- Yes, Use tcpdump!



- Use VFd to enable mirroring
  - One or both guest interfaces
  - Either in, out, or both directions
- Then:  
Start a debugging, guest and run tcpdump in the guest
- Or:  
Attach driver to target PF on host and run natively





### **The good**

- Implementation is straight forward. Completely NIC agnostic -- entirely implemented with DPDK API calls
- Possible expansion to include VLAN sensitive mirroring
- Enable/disable without traffic disruption

### **The bad**

- Can mirror only to same PF
- Must engage NIC *loopback*
- Some inconsistencies between NICs (single direction on some NICs)
- Inbound mirroring applied only after VLAN filtering (Niantic)

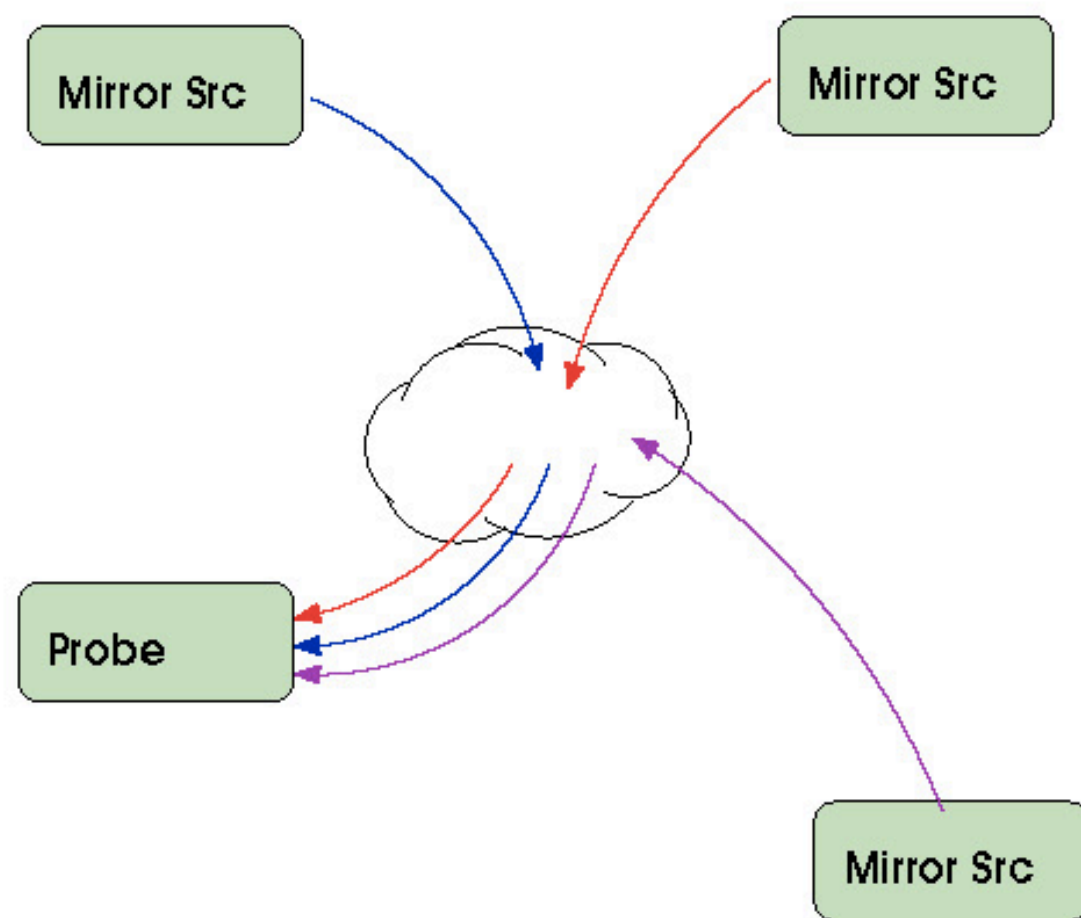


## Possibilities

- Application level protocol analysis
- *Probes* can generate things like step ladder diagram call traces

## But.....

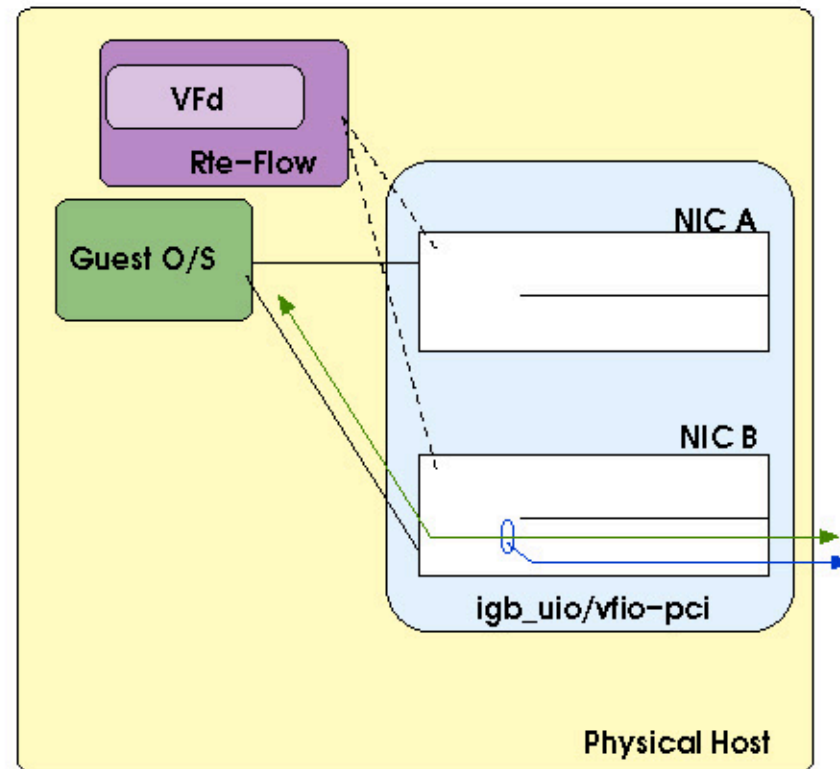
- Probes need mirrored traffic from multiple sources
- Usually requires encapsulation
- Tricky with current techniques



Imagine -- Add mirroring control to rte\_flow



- Semantics well established and cover most mirroring requirements
- Encapsulation and transmission off host
- Mirror only what is needed (5-tuple matching)
- Some NICs support these features today; let's use them.





KDPD

DATA PLANE DEVELOPMENT KIT



"Consider anything, only don't cry!"<sup>\*</sup>

Many thanks to

The VFd Community:

**Contributors**

Ajit Khaparde  
Alex Zelezniak  
Ariel Levkovich  
Dhanunjaya Naidu  
Hamdy Khader  
Kaustubh Joshi  
Munish Mehan  
Scott Daniels  
Stephen Hurd  
Tanvir Kekan

**Encouragers**

Aaron Rowden  
Bernard Iremonger  
Christian Dechamplain  
Ferruh Yigit  
Helin Zhang  
Jim Michelson  
Jingjing Wu  
John Craig  
Lian-min Wang  
Qi Z Zhang  
Rahul Shah  
Wenzhuo Lu

<sup>\*</sup> Lewis Carrol; *Through The Looking-Glass*,

