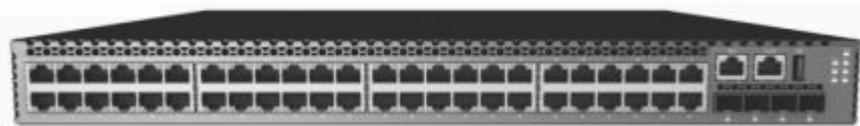


Rocker

Scott Feldman
netdev 0.1
2015 Ottawa





Network OS

Mgmt
CLI

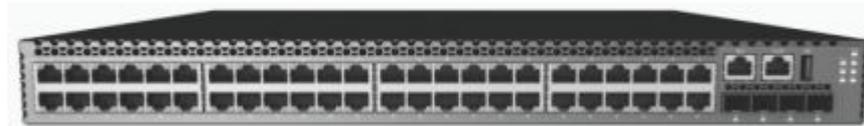
Control
Protocols

Monitoring

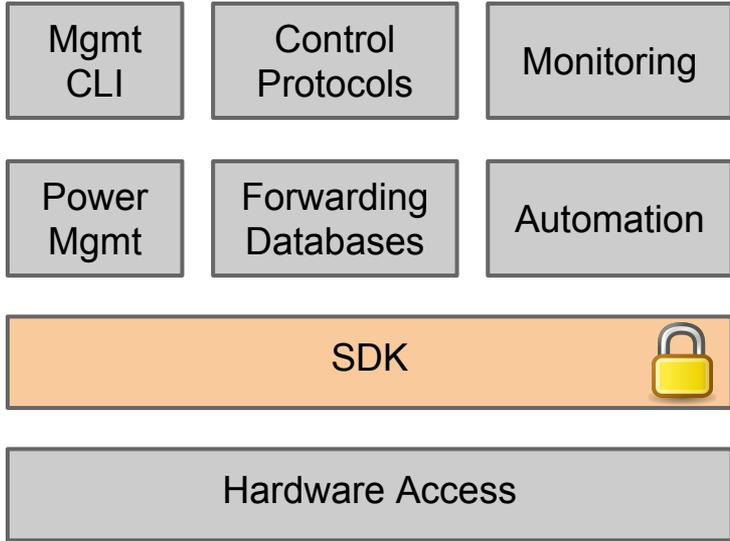
Power
Mgmt

Forwarding
Databases

Automation

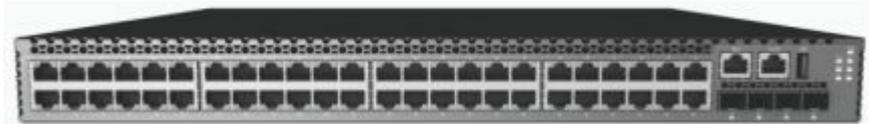


Network OS



NDA

\$\$\$





Linux

Mgmt
CLI

Control
Protocols

Monitoring

Power
Mgmt

Forwarding
Databases

Automation

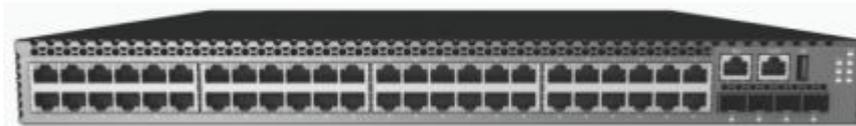
NDA

SDK



\$\$\$

Hardware Access





Linux

Mgmt
CLI

Control
Protocols

Monitoring

Power
Mgmt

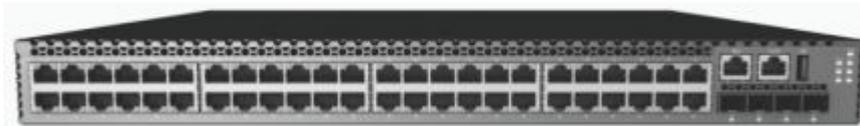
Forwarding
Databases

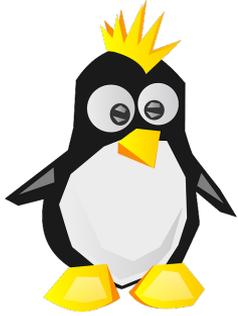
Automation

Switchdev

driver

Hardware Access





Linux

Mgmt
CLI

Control
Protocols

Monitoring

Power
Mgmt

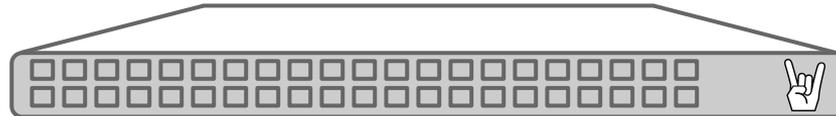
Forwarding
Databases

Automation

Switchdev

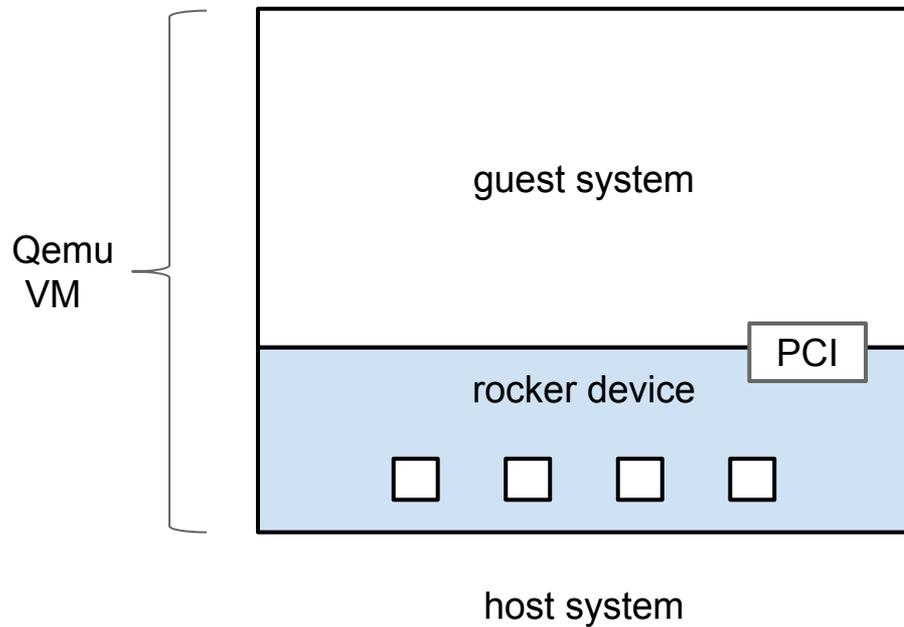
rocker driver

Hardware Access

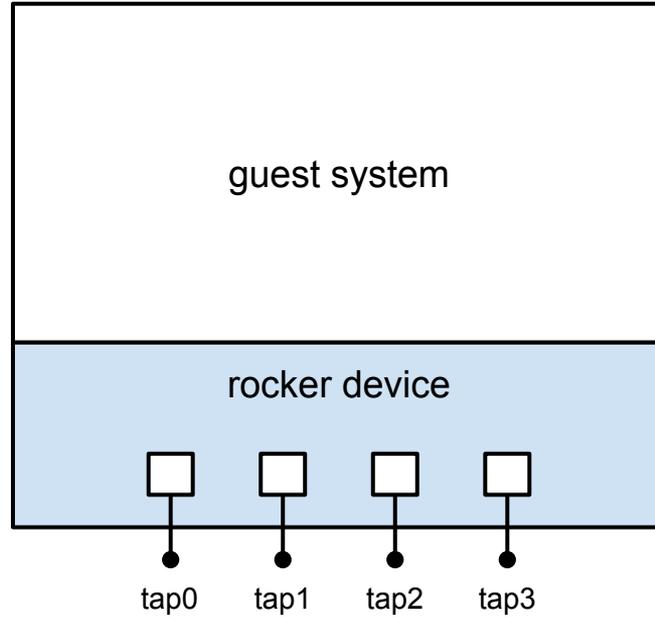




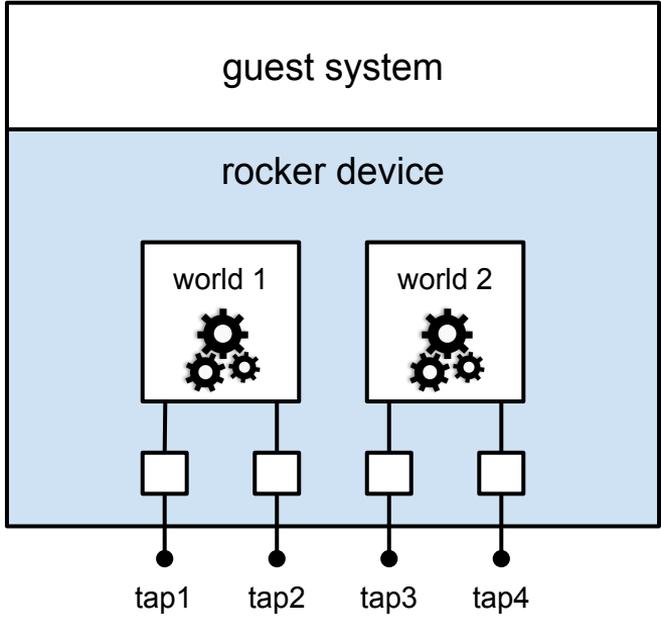




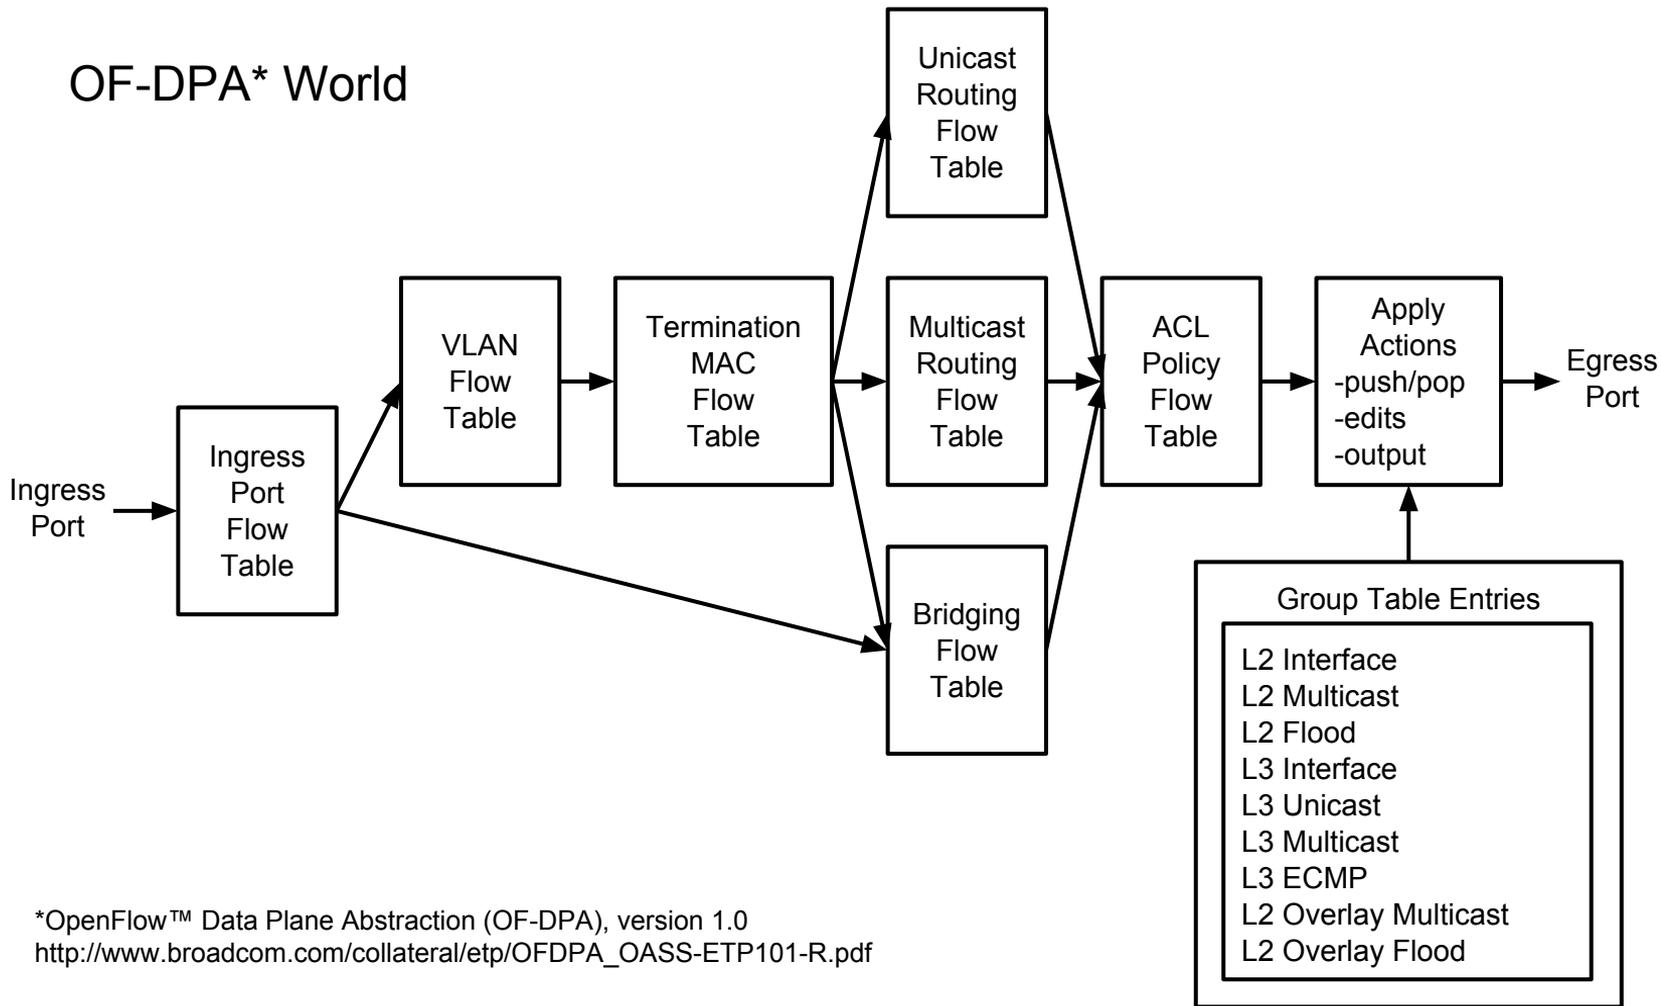
```
-device rocker,name=sw1,len-ports=4,ports[0]=dev1,ports[1]=dev2, \  
ports[2]=dev3,ports[3]=dev4
```



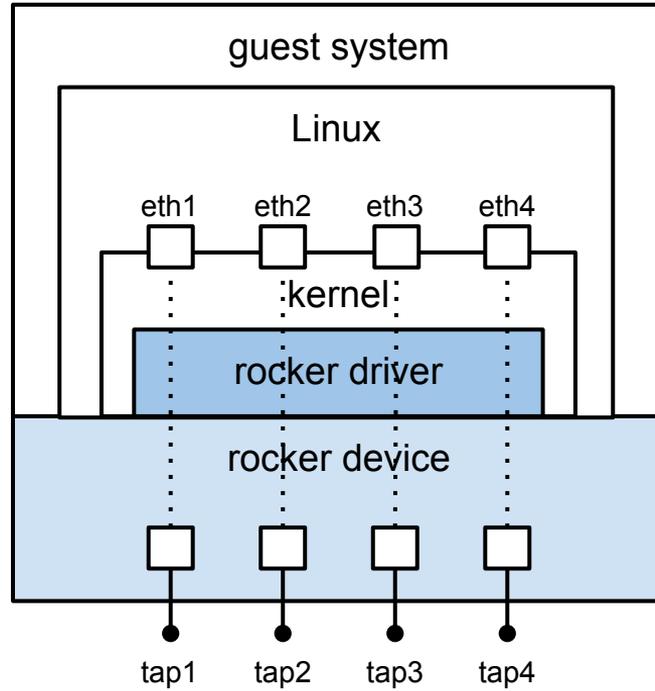
```
-netdev tap,ifname=tap1,dev=dev1 -netdev tap,ifname=tap2,dev=dev2 -netdev \ tap,  
ifname=tap3,dev=dev3 -netdev tap,ifname=tap4,dev=dev4
```



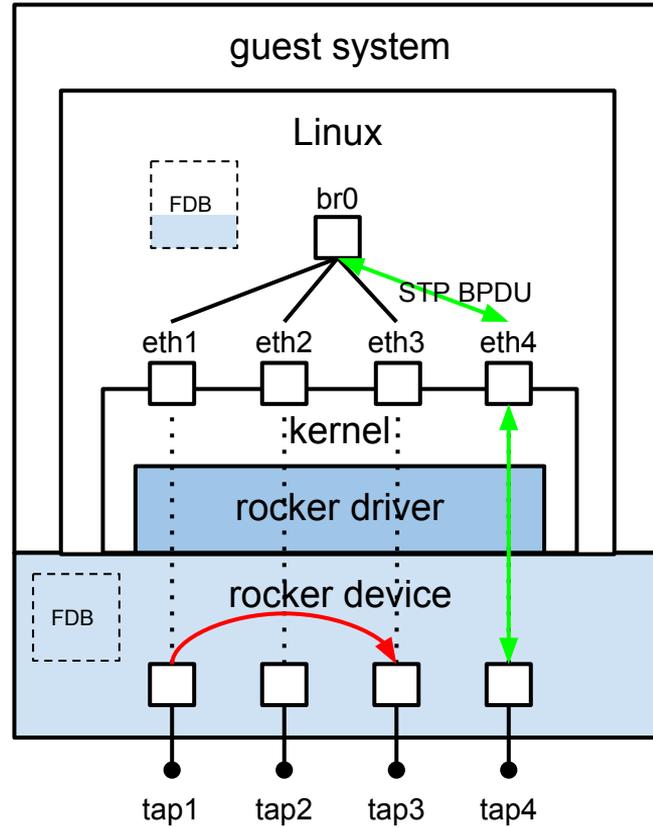
OF-DPA* World



*OpenFlow™ Data Plane Abstraction (OF-DPA), version 1.0
http://www.broadcom.com/collateral/etp/OFDPA_OASS-ETP101-R.pdf



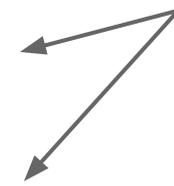
L2 Offload



- offloaded forwarding data path
- control path

```
$ /sbin/bridge fdb
33:33:00:00:00:01 dev eth0 self permanent
01:00:5e:00:00:01 dev eth0 self permanent
33:33:ff:00:00:00 dev eth0 self permanent
01:80:c2:00:00:0e dev eth0 self permanent
52:54:00:12:35:01 dev swp1 master br0 permanent
00:02:00:00:02:00 dev swp1 master br0 external
00:02:00:00:02:00 dev swp1 self
52:54:00:12:35:02 dev swp2 master br0 permanent
00:02:00:00:03:00 dev swp2 master br0 external
00:02:00:00:03:00 dev swp2 self
33:33:00:00:00:01 dev br0 self permanent
01:00:5e:00:00:01 dev br0 self permanent
33:33:ff:12:35:01 dev br0 self permanent
```

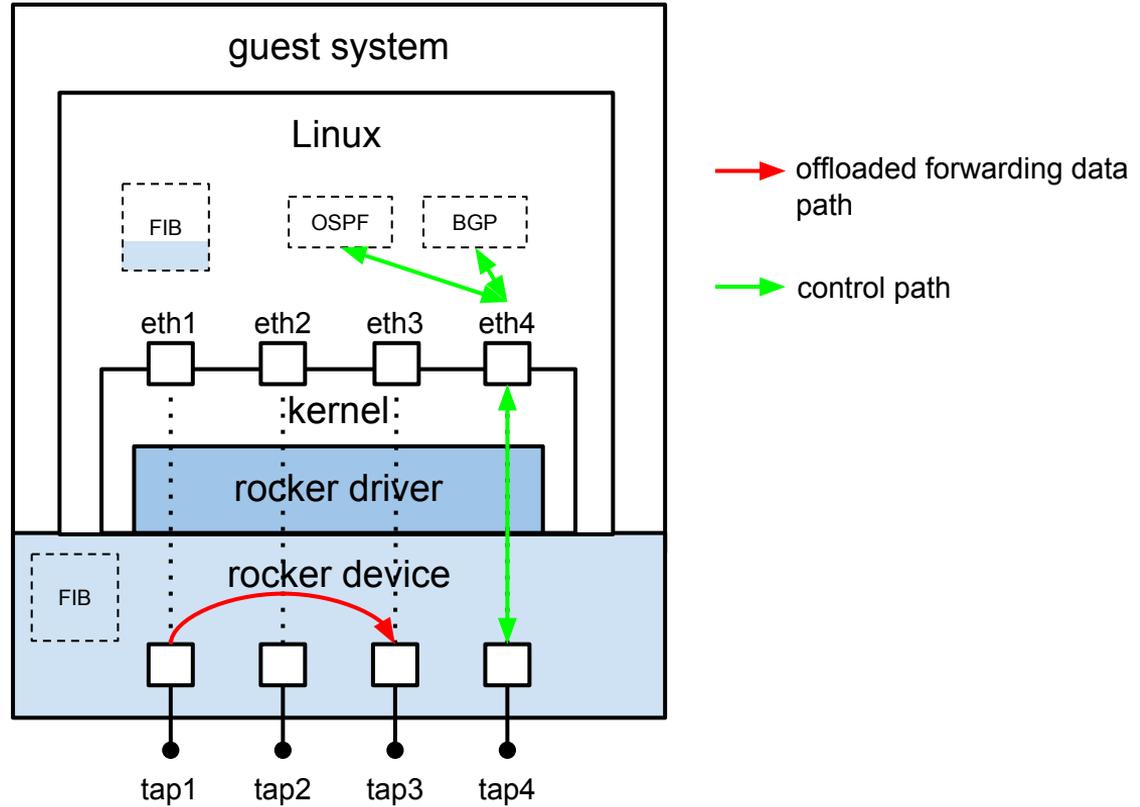
Learned by switch



```
(qemu) info rocker-of-dpa-flows sw1
prio tbl hits key(mask) --> actions
3 50 10 vlan 3840 dst 00:02:00:00:03:00 --> write group 0x0f000002 goto tbl 60
3 50 13 vlan 3840 dst 00:02:00:00:02:00 --> write group 0x0f000001 goto tbl 60
1 30 IP dst 11.0.0.0/32 --> write group 0x0f000000 goto tbl 60
1 30 IP dst 11.0.0.0/24 --> write group 0x0f000000 goto tbl 60
1 30 IP dst 11.0.0.255/32 --> write group 0x0f000000 goto tbl 60
1 30 10 IP dst 11.0.0.3/32 --> write group 0x0f000000 goto tbl 60
0 20 pport 2 vlan 3840 IPv6 dst 52:54:00:12:35:02 --> goto tbl 30
0 20 pport 2 vlan 3840 IP dst 52:54:00:12:35:02 --> goto tbl 30
1 20 8 pport 2 vlan 3840 IPv6 dst 33:33:00:00:00:00(ff:ff:00:00:00:00) --> goto tbl 40
1 20 pport 2 vlan 3840 IP dst 01:00:5e:00:00:00(ff:ff:ff:80:00:00) --> goto tbl 40
3 60 1 pport 2 vlan 3840 dst 01:80:c2:00:00:00(ff:ff:ff:ff:ff:f0) --> write group 0x0f000000
0 20 pport 1 vlan 3840 IPv6 dst 52:54:00:12:35:01 --> goto tbl 30
0 20 10 pport 1 vlan 3840 IP dst 52:54:00:12:35:01 --> goto tbl 30
1 50 6 vlan 3840 --> write group 0x4f000000 goto tbl 60
1 20 8 pport 1 vlan 3840 IPv6 dst 33:33:00:00:00:00(ff:ff:00:00:00:00) --> goto tbl 40
1 20 pport 1 vlan 3840 IP dst 01:00:5e:00:00:00(ff:ff:ff:80:00:00) --> goto tbl 40
3 60 1 pport 1 vlan 3840 dst 01:80:c2:00:00:00(ff:ff:ff:ff:ff:f0) --> write group 0x0f000000
1 10 22 pport 2 vlan 0 --> apply new vlan 3840 goto tbl 20
1 10 33 pport 1 vlan 0 --> apply new vlan 3840 goto tbl 20
1 0 55 pport 0(0xfffff0000) --> goto tbl 10
```

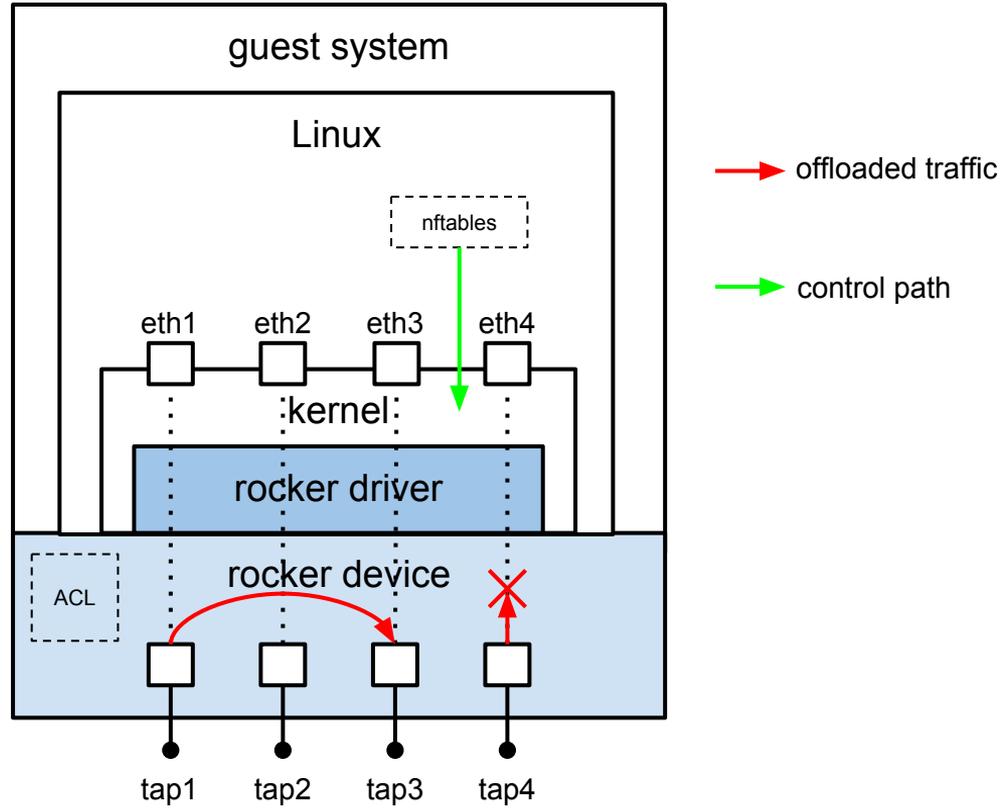
```
(qemu) info rocker-of-dpa-groups sw1
id (decode) --> buckets
0x0f000002 (type L2 interface vlan 3840 pport 2) --> pop vlan out pport 2
0x0f000001 (type L2 interface vlan 3840 pport 1) --> pop vlan out pport 1
0x4f000000 (type L2 flood vlan 3840 index 0) --> groups [0x0f000002,0x0f000001]
0x0f000000 (type L2 interface vlan 3840 pport 0) --> pop vlan out pport 0
```

L3 Offload

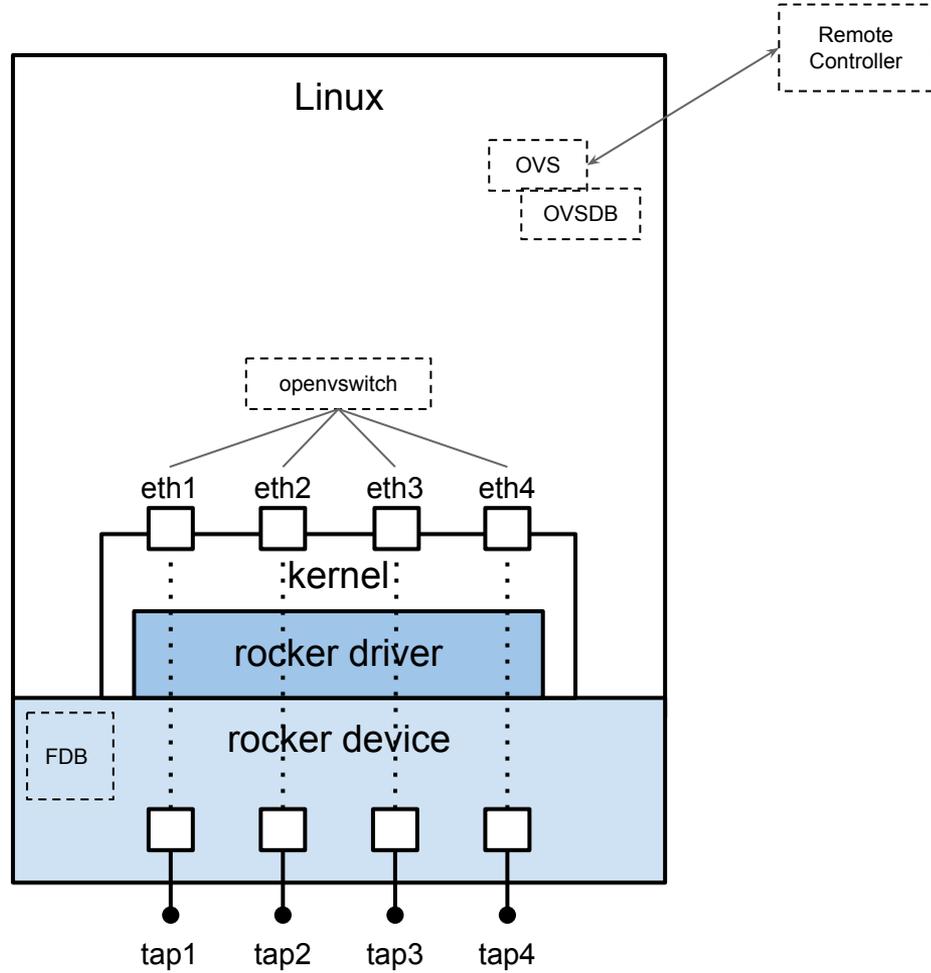


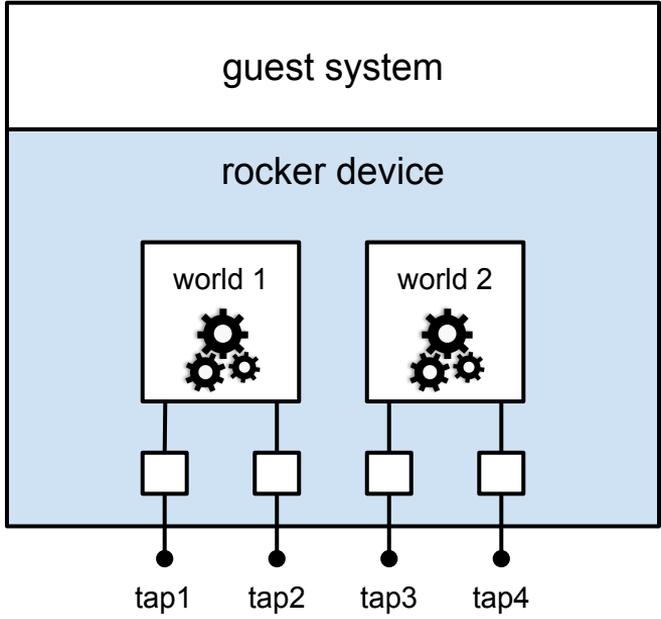
Future

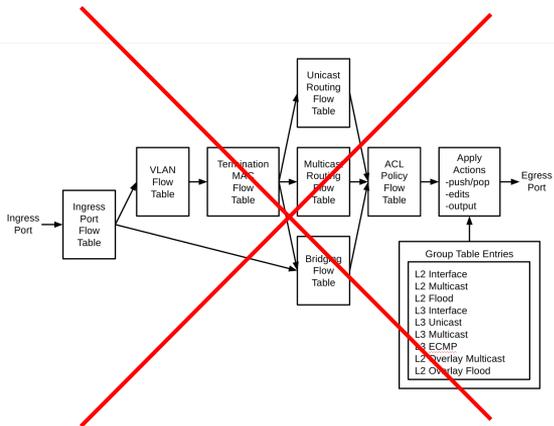
ACL Offload



OVS Offload







Fixed pipeline, c. 2014

P4

```

blah {
  blah {
    blah blah
  }
}

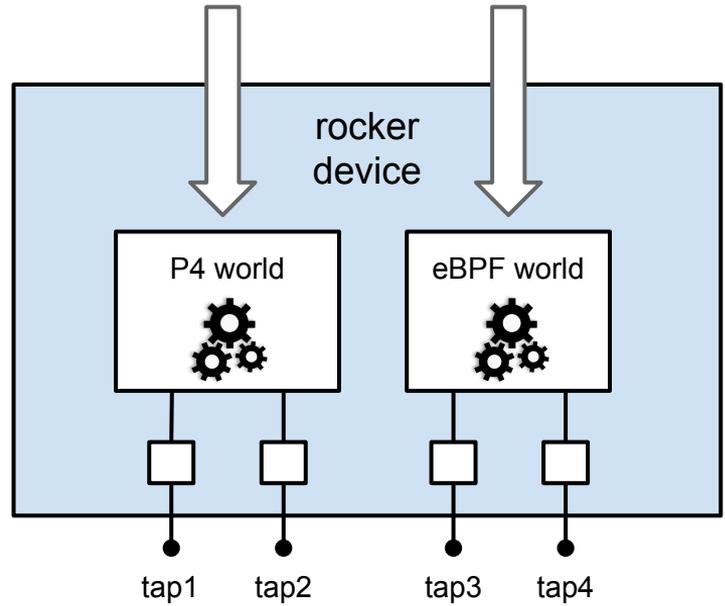
```

eBPF

```

struct
bpf_insn prog
[] = {
  BLAH(blah),
}

```



Questions? Answers?