

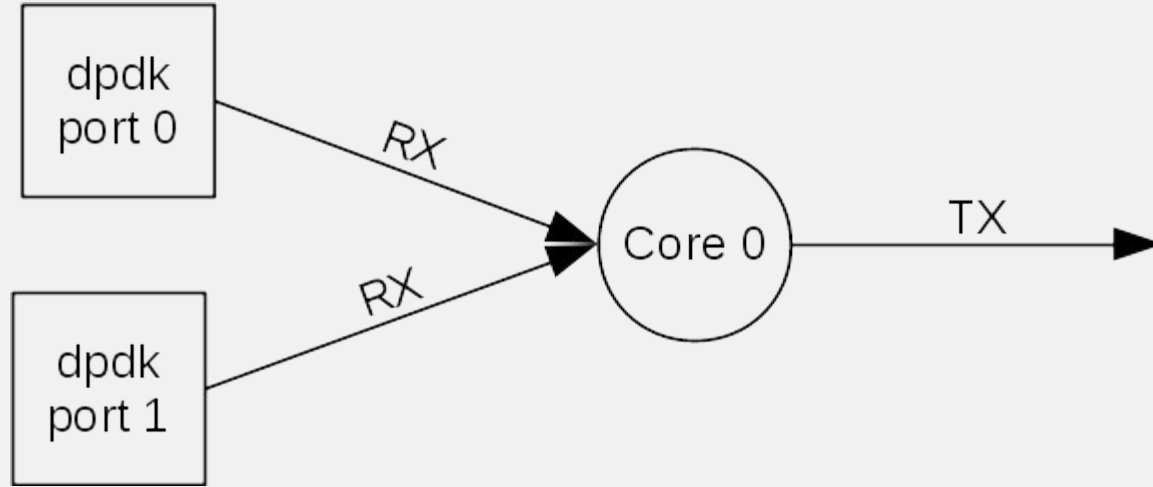


OVS-DPDK: Every cycle counts


OVS-DPDK Datapath workload distribution

Kevin Traynor
ktraynor@redhat.com
16th November 2017

First came OVS-DPDK




Then came multiple PMD's (cores)




```
# ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=0x3
```

Then came multiple Receive queues

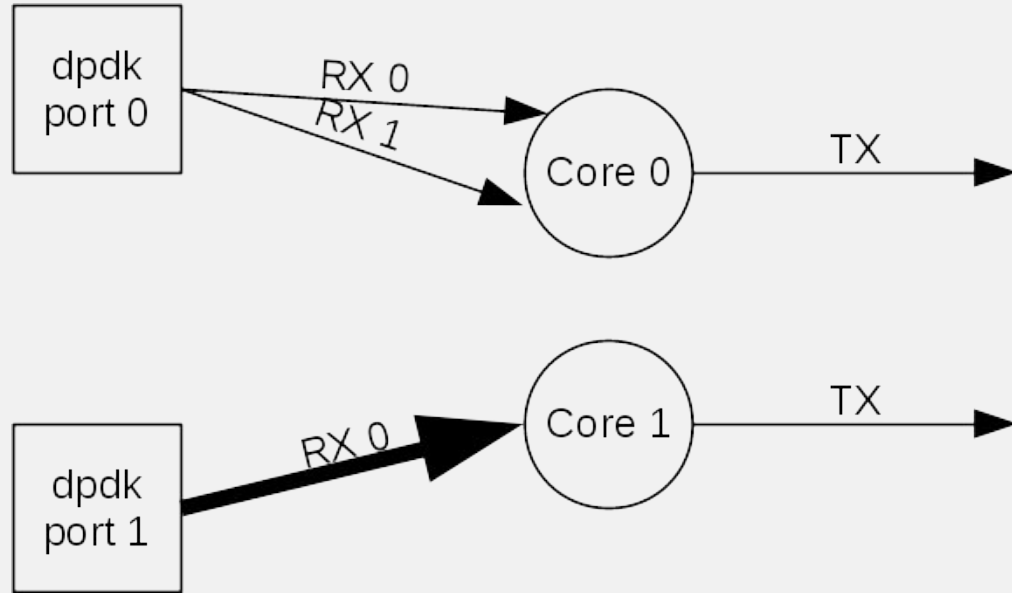


```
# ovs-vsctl set Interface dpdkport0 options:n_rxq=2  
# ovs-vsctl set Interface dpdkport1 options:n_rxq=2
```

But what if this happens...




Manually pin fat queue...but it doesn't scale

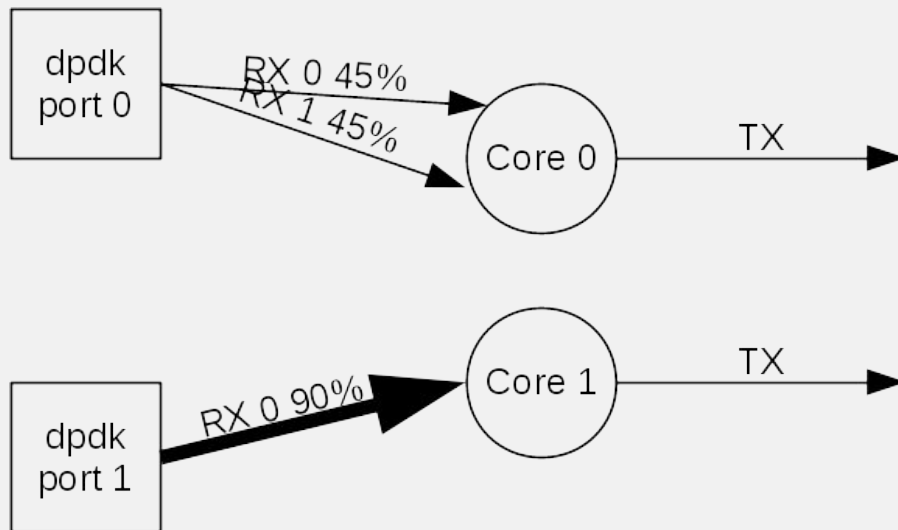


```
# ovs-vsctl set Interface dpdkport1 options:pmd-rxq-affinity=0:1
```

Another way...count cycles/core utilisation



The next time queues are assigned to PMDs



Also, manually with

```
# ovs-appctl dpif-netdev/pmd-rxq-rebalance
```


Demo

```
DUT
Every 1.0s: ovs-appctl dpif-netdev/pmd-rxq-show Thu Oct 26 15:11:36 2017
pmd thread numa_id 0 core_id 4:
  isolated: false
  port: dpdkport0 queue-id: 0 core %: 43
  port: dpdkport0 queue-id: 1 core %: 42
pmd thread numa_id 0 core_id 6:
  isolated: false
  port: dpdkport1 queue-id: 0 core %: 100

DUT
[DUT]# ovs-appctl dpif-netdev/pmd-rxq-rebalance
pmd rxq rebalance requested.
[DUT]# █

Packet Generator
/ Ports 0-1 of 2 <Main Page> Copyright (c) <2010-2017>, Intel Corporation
Flags:Port : P.....R.....:0 P.....:1
Link State : <UP-10000-FD> <UP-10000-FD> ----TotalRate----
Pkts/s Max/Rx : 13375972/13375712 5205286/5205250 18581245/18580962
Max/Tx : 5205279/5205248 13375842/13375729 18581121/18580977
MBytes/s Rx/Tx : 8988/3497 3497/8988 12486/12486
Broadcast : 0 0
Multicast : 0 0
64 Bytes : 68372045 14396995
65-127 : 0 0
128-255 : 0 0
256-511 : 0 0
512-1023 : 0 0
1024-1518 : 0 0
Runts/Jumbos : 0/0 0/0
Errors Rx/Tx : 0/0 0/0
Total Rx Pkts : 279350568 133409984
Tx Pkts : 133409536 342815552
Rx MBs : 187723 89651
Tx MBs : 89651 230372
ARP/ICMP Pkts : 0/0 0/0
Pattern Type : abcd... abcd...
Tx Count/% Rate : Forever /35% Forever /90%
PktSize/Tx Burst : 64 / 32 64 / 32
Src/Dest Port : 1234 / 5678 1234 / 5678
Pkt Type:VLAN ID : IPv4 / TCP:0001 IPv4 / TCP:0001
Dst IP Address : 192.168.1.1 192.168.0.1
Src IP Address : 192.168.0.1/24 192.168.1.1/24
Dst MAC Address : ec:f4:bb:d0:2f:6a ec:f4:bb:d0:2f:68
Src MAC Address : ec:f4:bb:d0:2f:68 ec:f4:bb:d0:2f:6a
VendID/PCI Addr : 8086:10fb/01:00.0 8086:10fb/01:00.1

.. Pktgen Ver: 3.4.0 (DPDK 17.08.0) Powered by DPDK .....

Pktgen:> str
Pktgen:> █
```

<https://youtu.be/gkYLFtoQrui>



THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos