
DPDK
DATA PLANE DEVELOPMENT KIT
SafetyOrange
DPDK Summit - San Jose – 2017

An experiment on a backpack-size DPDK development and test environment

andras.kovacs@ericsson.com
laszlo.vadkeri@ericsson.com

Objectives

- Demonstrate that our infrastructure can eliminate the packet processing impacts of virtualization
- PoC of ‘virtual router on a pendrive’ for portability and demo purposes
- Free yet powerful packet generator (Pktgen-DPDK)
- Efficient and affordable standalone DPDK environment (personal PoD)
- PCIe bifurcation for our use case (increase port density)

The hardware

- Mini-ITX motherboard (6.7” x 6.7”)
- Intel low-power Xeon CPU (14c/28t)
- 32Gbyte memory (4 channels)
- PCIe riser splitter for PCIe bifurcation
- 2 x X710 NICs (both in 4x10G mode)
- NFC S4 Mini case (4.3 liter / 1.14 gallon)
  http://nfc-systems.com/s4-mini/

PCIe bifurcation

"The ability to support multiple PCIe devices using a single PCIe slot"

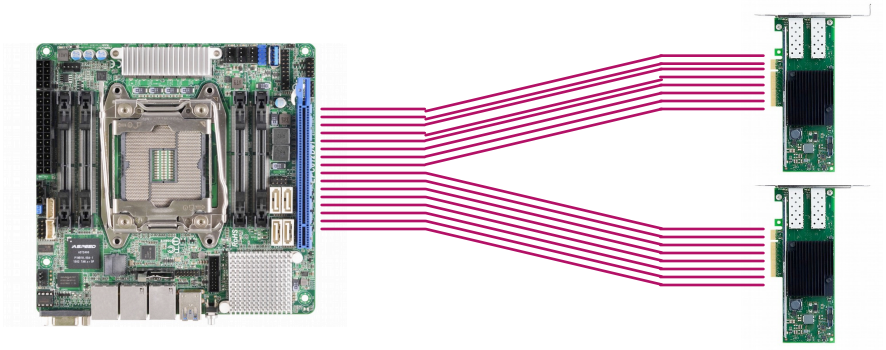

- E.g. Intel Niantic (X520) and Fortville (X710) cards use only x8 PCI-E lanes
- If the motherboard supports it then a x16 PCIe can be split into x8x8 slots (or even further… x4x4x4x4)
- Port density can be increased (even to oversubscription)

Test setups

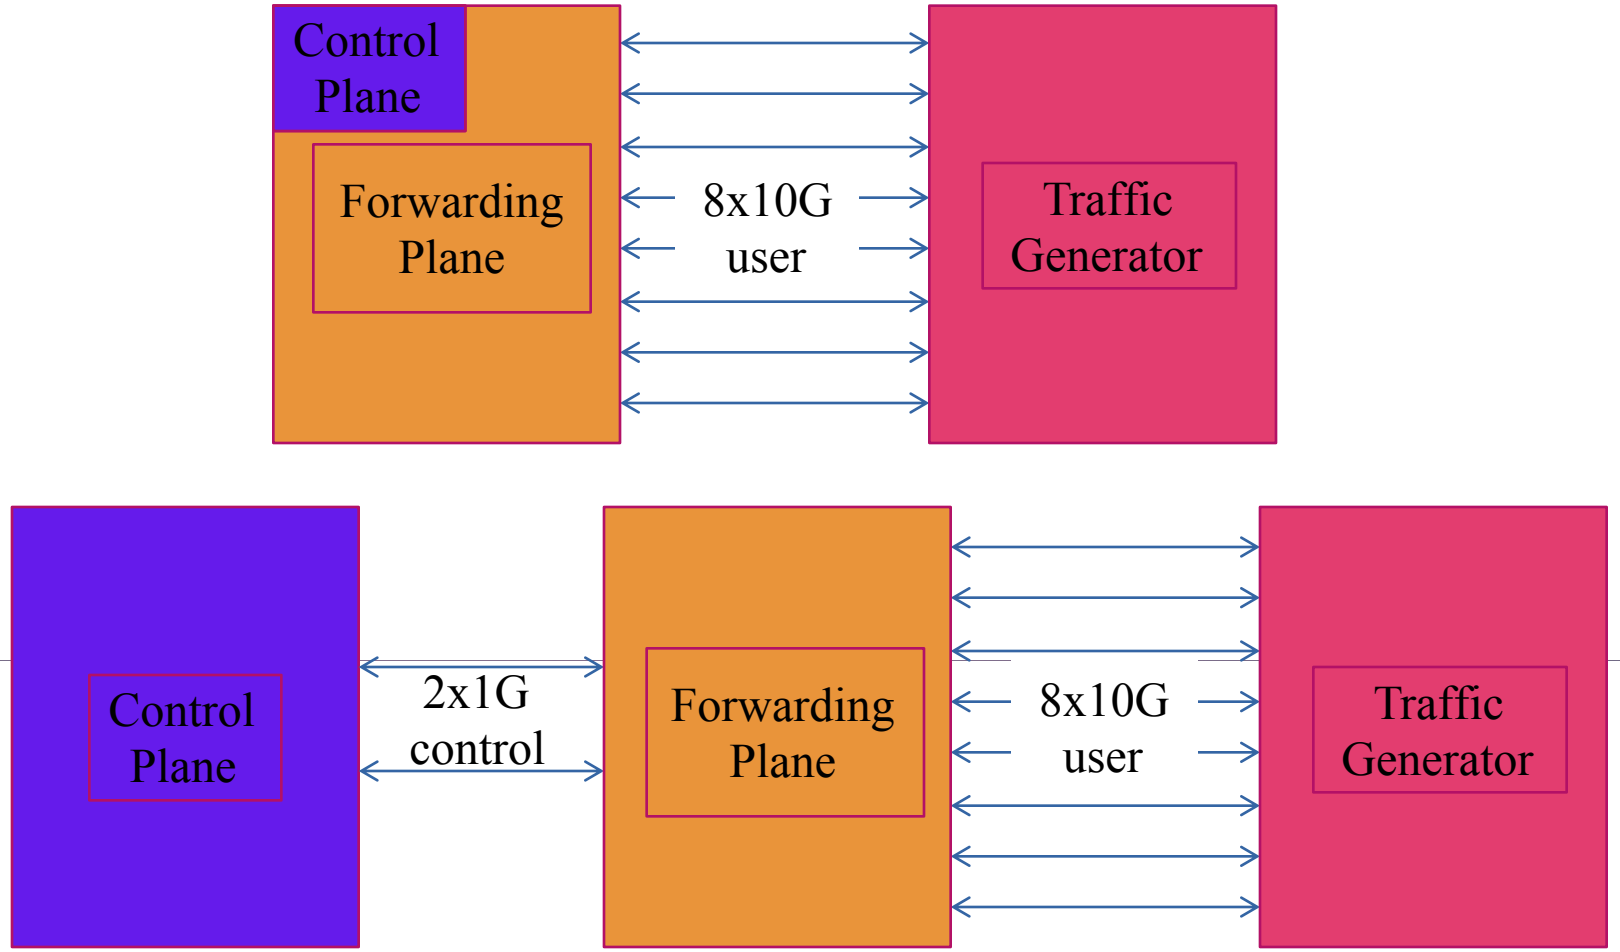
Control Plane
Forwarding Plane
8x10G user
Traffic Generator
Control Plane
2x1G control
Forwarding Plane
8x10G user
Traffic Generator

Previous Work

DPDK

DPDK Summit 2014

DPDK Userspace 2015

The layer that makes it all possible

- Abstraction layer for environment independence
- Memory management forged to make best use of x86 architecture
- Predictable high performance and low latency
- High Availability enhancements
- Runs both native and virtualized mode

Finetuning

- Ability to remap all CPU resources. Allows easy adaptation to environment specific HT sibling mappings. Can also be utilized to change the CPU resources dedicated to IO / forwarding / application packet processing etc.

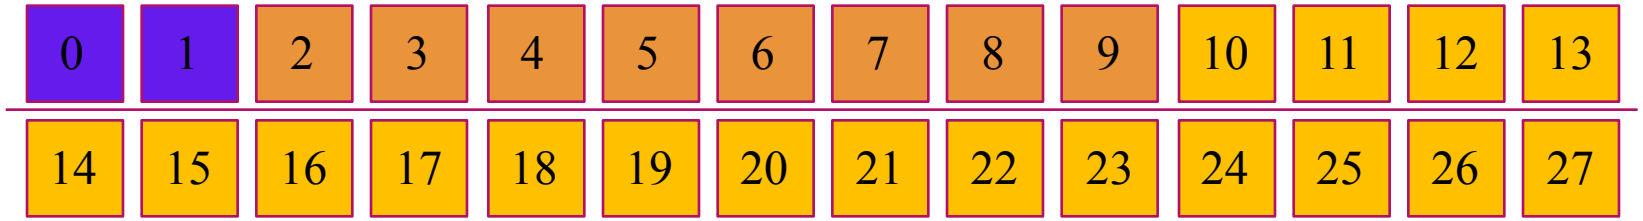

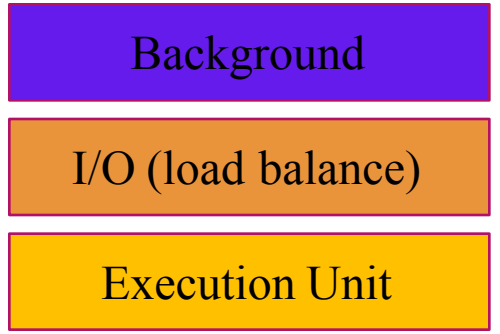

- Tuned CPU allocation after initial measurements and bottlenecks identified (simple reconfiguration)

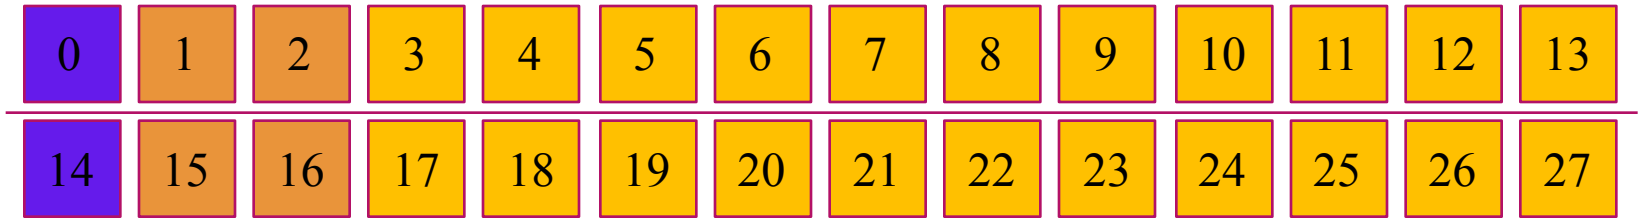

Pktgen-DPDK sample

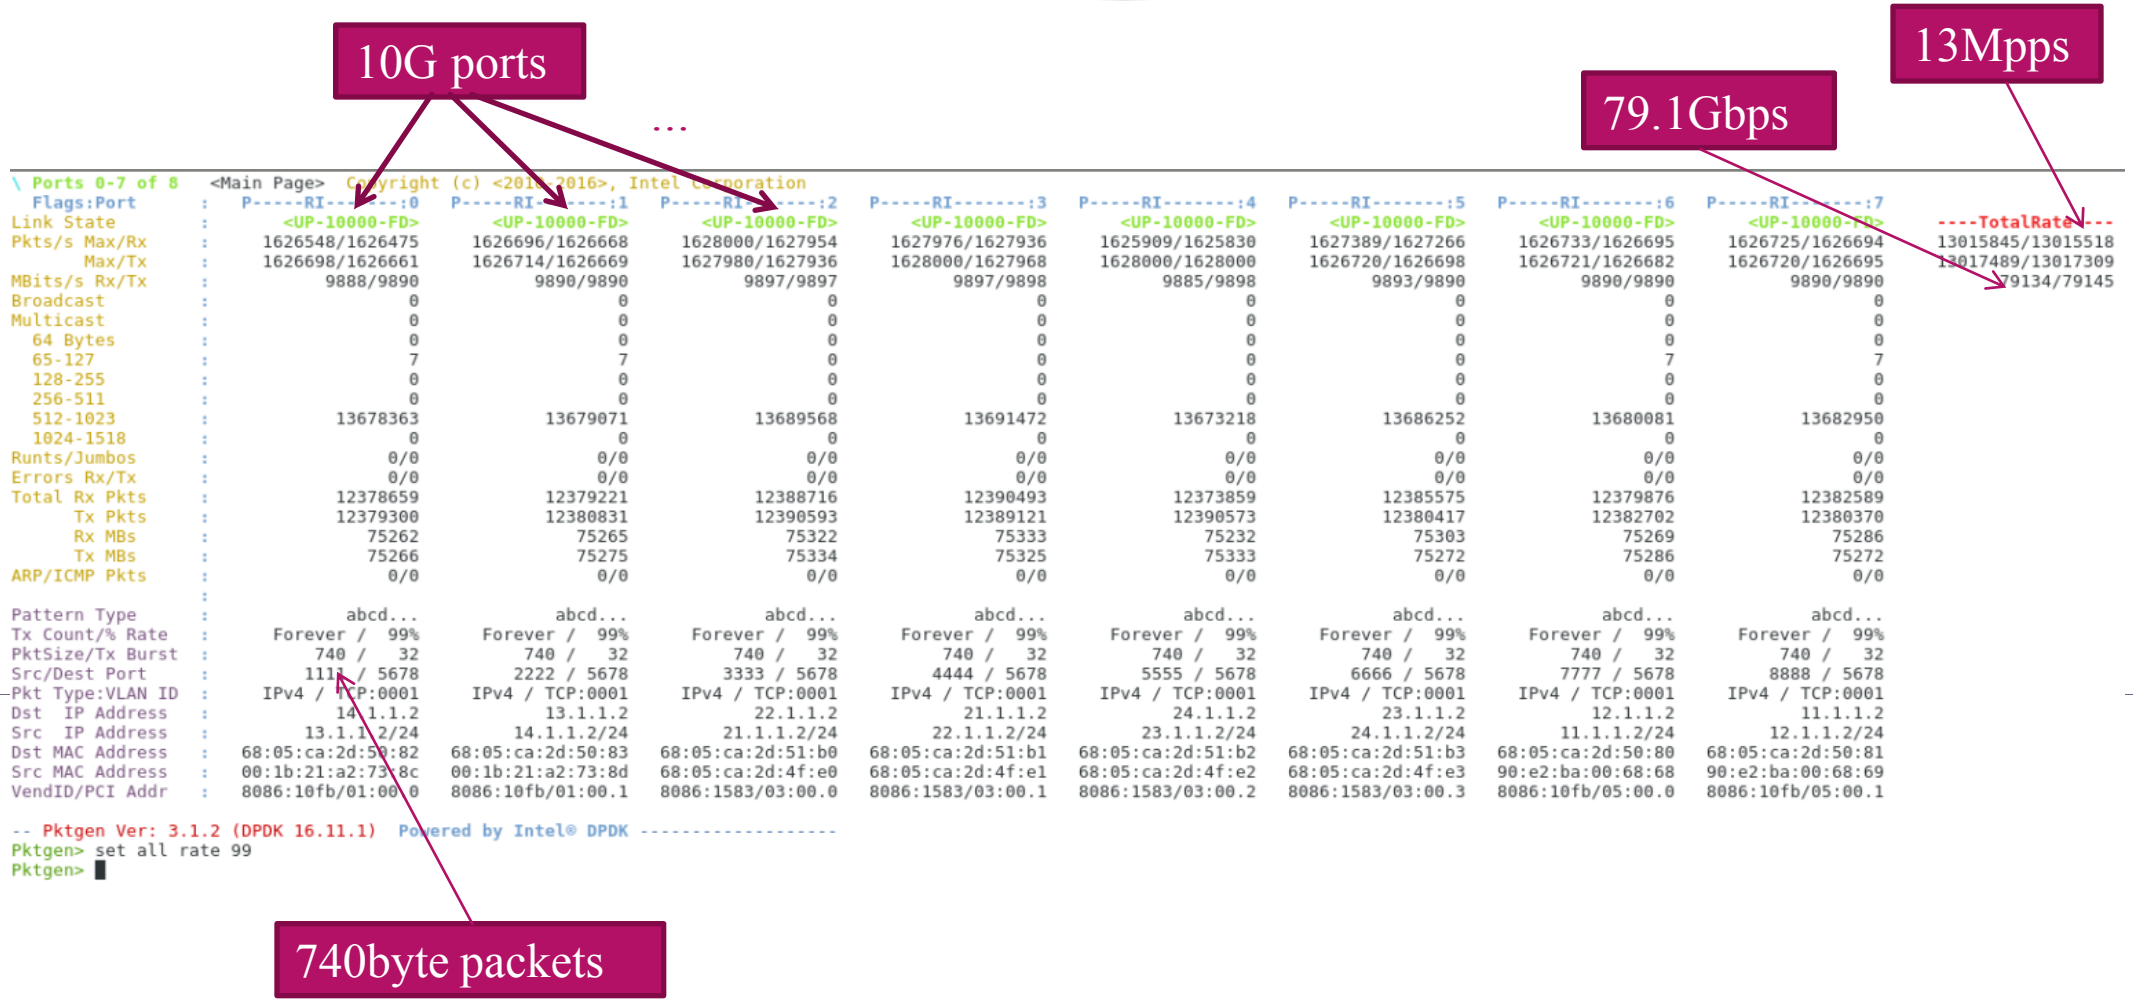


```
\ Ports 0-7 of 8   <Main Page>  Copyright (c) <2010-2016>, Intel Corporation
  Flags:Port      :   P-----RI--------:0   P-----RI--------:1   P-----RI--------:2   P-----RI--------:3   P-----RI--------:4   P-----RI--------:5   P-----RI--------:6   P-----RI--------:7
Link State        :       <UP-10000-FD>        <UP-10000-FD>        <UP-10000-FD>        <UP-10000-FD>        <UP-10000-FD>        <UP-10000-FD>        <UP-10000-FD>        <UP-10000-FD>      ----TotalRate----
Pkts/s Max/Rx     :     1626548/1626475      1626696/1626668      1628000/1627954      1627976/1627936      1625909/1625830      1627389/1627266      1626733/1626695      1626725/1626694      13015845/13015518
      Max/Tx      :     1626698/1626661      1626714/1626669      1627980/1627936      1628000/1627968      1628000/1628000      1626720/1626698      1626721/1626682      1626720/1626695      13017489/13017309
MBits/s Rx/Tx     :          9888/9890            9890/9890            9897/9897            9897/9898            9885/9898            9893/9890            9890/9890            9890/9890            79134/79145
Broadcast         :                  0                    0                    0                    0                    0                    0                    0                    0
Multicast         :                  0                    0                    0                    0                    0                    0                    0                    0
  64 Bytes        :                  0                    0                    0                    0                    0                    0                    0                    0
  65-127          :                  7                    7                    0                    0                    0                    0                    7                    7
  128-255         :                  0                    0                    0                    0                    0                    0                    0                    0
  256-511         :                  0                    0                    0                    0                    0                    0                    0                    0
  512-1023        :           13678363             13679071             13689568             13691472             13673218             13686252             13680081             13682950
  1024-1518       :                  0                    0                    0                    0                    0                    0                    0                    0
Runts/Jumbos      :                0/0                  0/0                  0/0                  0/0                  0/0                  0/0                  0/0                  0/0
Errors Rx/Tx      :                0/0                  0/0                  0/0                  0/0                  0/0                  0/0                  0/0                  0/0
Total Rx Pkts     :           12378659             12379221             12388716             12390493             12373859             12385575             12379876             12382589
      Tx Pkts     :           12379300             12380831             12390593             12389121             12390573             12380417             12382702             12380370
      Rx MBs      :              75262                75265                75322                75333                75232                75303                75269                75286
      Tx MBs      :              75266                75275                75334                75325                75333                75272                75286                75272
ARP/ICMP Pkts     :                0/0                  0/0                  0/0                  0/0                  0/0                  0/0                  0/0                  0/0
                  :
Pattern Type      :               abcd...              abcd...              abcd...              abcd...              abcd...              abcd...              abcd...              abcd...
Tx Count/% Rate   :        Forever /  99%       Forever /  99%       Forever /  99%       Forever /  99%       Forever /  99%       Forever /  99%       Forever /  99%       Forever /  99%
PktSize/Tx Burst  :           740 /   32           740 /   32           740 /   32           740 /   32           740 /   32           740 /   32           740 /   32           740 /   32
Src/Dest Port     :          1111 / 5678          2222 / 5678          3333 / 5678          4444 / 5678          5555 / 5678          6666 / 5678          7777 / 5678          8888 / 5678
Pkt Type:VLAN ID  :      IPv4 / TCP:0001      IPv4 / TCP:0001      IPv4 / TCP:0001      IPv4 / TCP:0001      IPv4 / TCP:0001      IPv4 / TCP:0001      IPv4 / TCP:0001      IPv4 / TCP:0001
Dst  IP Address   :           14.1.1.2             13.1.1.2             22.1.1.2             21.1.1.2             24.1.1.2             23.1.1.2             12.1.1.2             11.1.1.2
Src  IP Address   :          13.1.1.2/24          14.1.1.2/24          21.1.1.2/24          22.1.1.2/24          23.1.1.2/24          24.1.1.2/24          11.1.1.2/24          12.1.1.2/24
Dst MAC Address   :   68:05:ca:2d:50:82   68:05:ca:2d:50:83   68:05:ca:2d:51:b0   68:05:ca:2d:51:b1   68:05:ca:2d:51:b2   68:05:ca:2d:51:b3   68:05:ca:2d:50:80   68:05:ca:2d:50:81
Src MAC Address   :   00:1b:21:a2:73:8c   00:1b:21:a2:73:8d   68:05:ca:2d:4f:e0   68:05:ca:2d:4f:e1   68:05:ca:2d:4f:e2   68:05:ca:2d:4f:e3   90:e2:ba:00:68:68   90:e2:ba:00:68:69
VendID/PCI Addr   :   8086:10fb/01:00.0   8086:10fb/01:00.1   8086:1583/03:00.0   8086:1583/03:00.1   8086:1583/03:00.2   8086:1583/03:00.3   8086:10fb/05:00.0   8086:10fb/05:00.1

-- Pktgen Ver: 3.1.2 (DPDK 16.11.1)  Powered by Intel® DPDK --------------------
Pktgen> set all rate 99
Pktgen>
```

Power consumption

- Measured with a Kill-A-Watt P4400.01 at the wall
- 35W idle (half of a dual socket blade with 2x E5-2620v3 (6c/12ht))
- 99W with 2BG/4IO/22EU without traffic
- 111W max at highest load (range of a desktop replacement notebook)

Thank you!