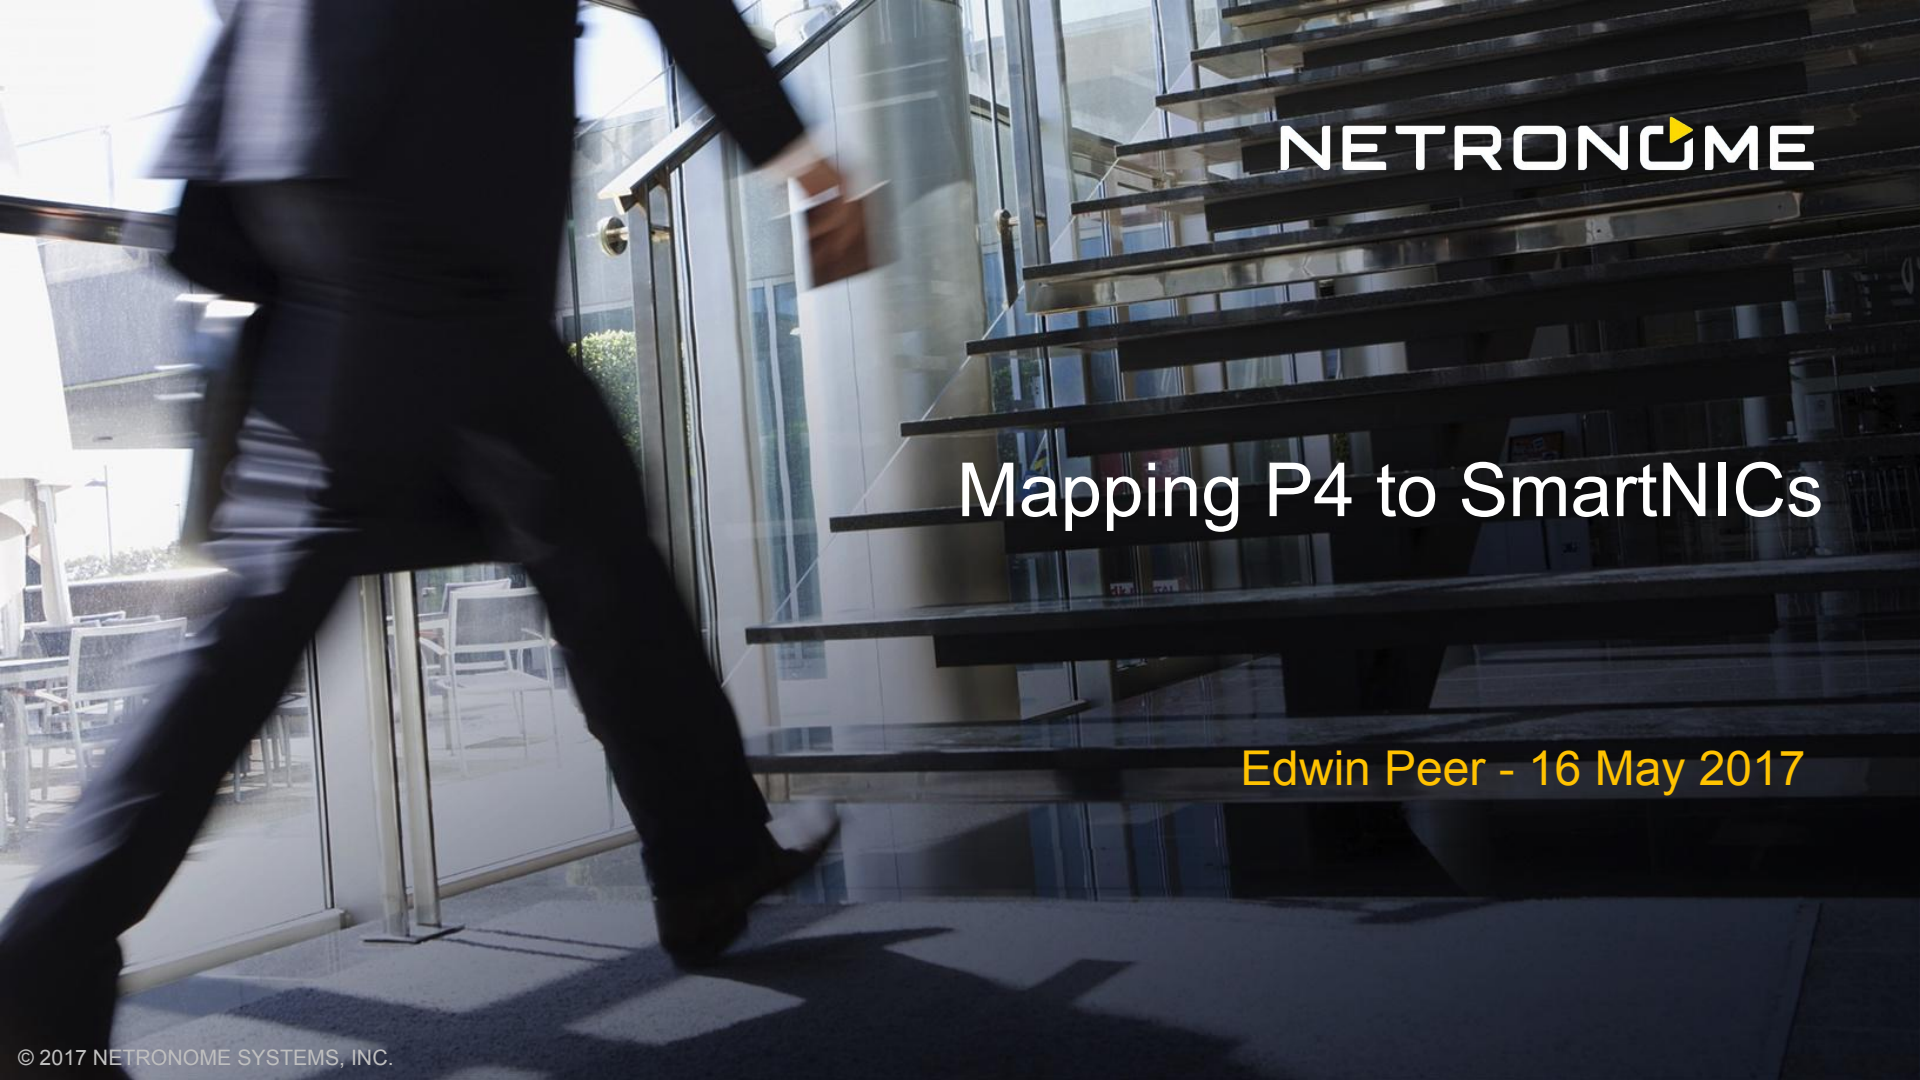

Mapping P4 to SmartNICs

Edwin Peer - 16 May 2017

Agenda

NETRONOME

- SmartNIC Hardware
- Silicon Architecture
- Programming Model
- Mapping P4
- Results & Experience
- Further Reading / Q&A

Agilio™ CX SmartNIC Family

NETRONOME

- Optimized for standard server based cloud data centers
- Low profile half length PCIe form factor, power < 25W
- Based on Netronome's Network Flow Processor 4xxx silicon (72 cores x 8 threads each)
- 2GB DRAM for lookup tables / state tables (millions of entries)

1x 40GbE

2x 10GbE

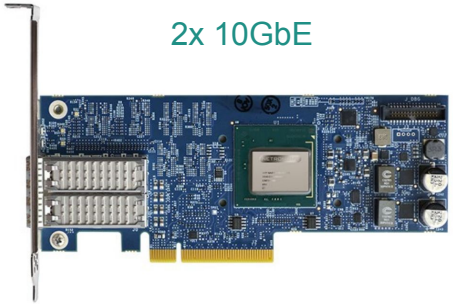

2x 40GbE

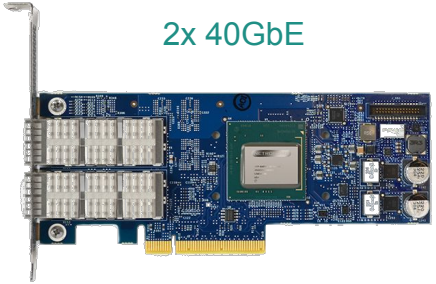

2x 25GbE

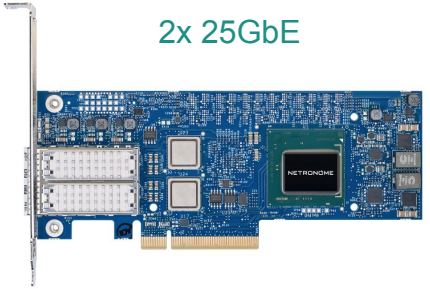

Agilio™ LX SmartNIC Family

NETRONOME

- Optimized for higher throughput requirements - middlebox, gateway, appliance, service node…
- Full height half length PCIe form factor
- Based on Netronome's Network Flow Processor 6xxx silicon (120 cores x 8 threads)
- Memory: 8GB of DDR3 DRAM @ 1866Mhz w/ECC
- Dual PCIe Gen3x8

2x40GbE (QSFP)

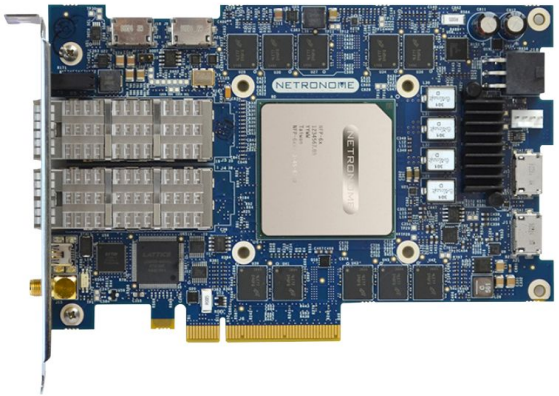

1x100GbE (CXP)

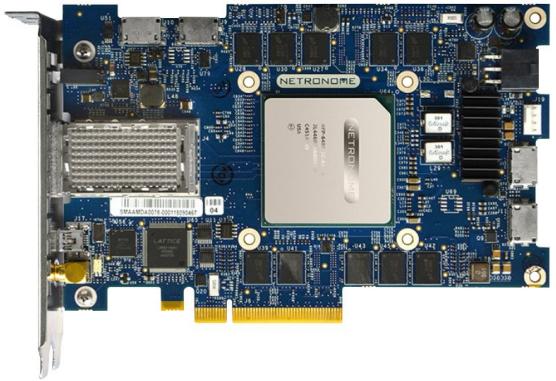

P4 Programmable SmartNIC in Context

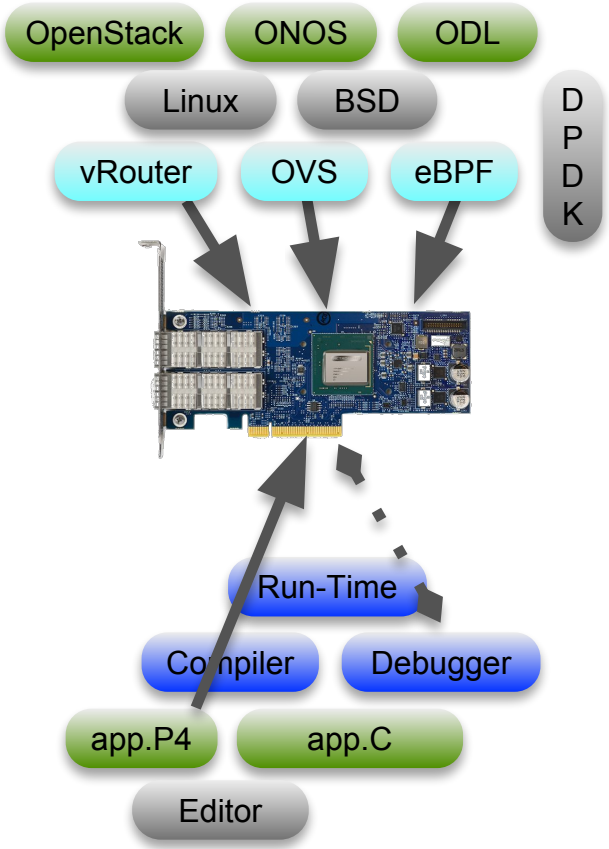

- Transparent acceleration of OVS / Contrail / eBPF
- SmartNIC with dynamic firmware
- Custom datapath in P4 and/or C

Inside the Netronome Flow Processor (NFP6xxx)

NETRONOME

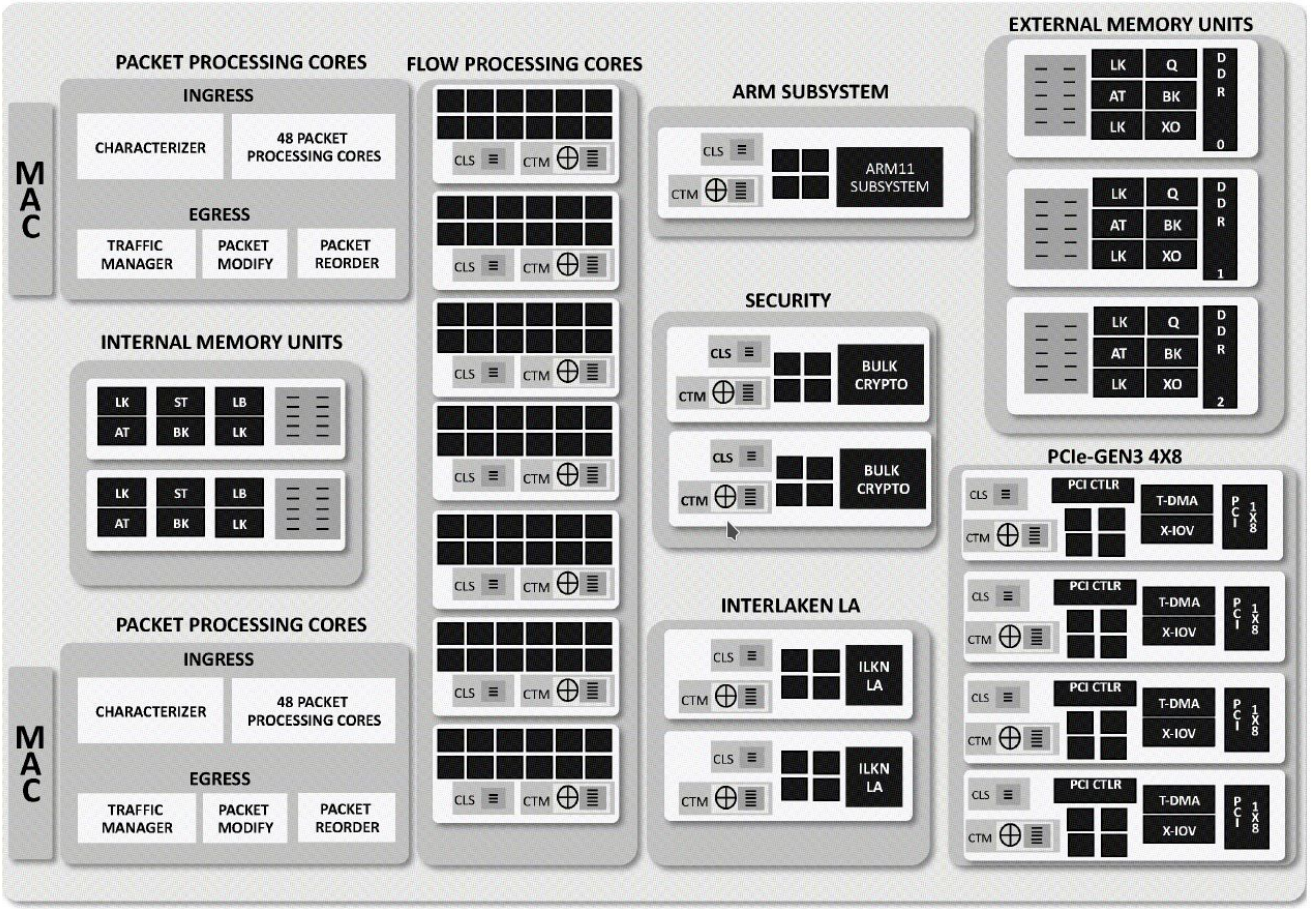

Hierarchical Memory Architecture

NETRONOME

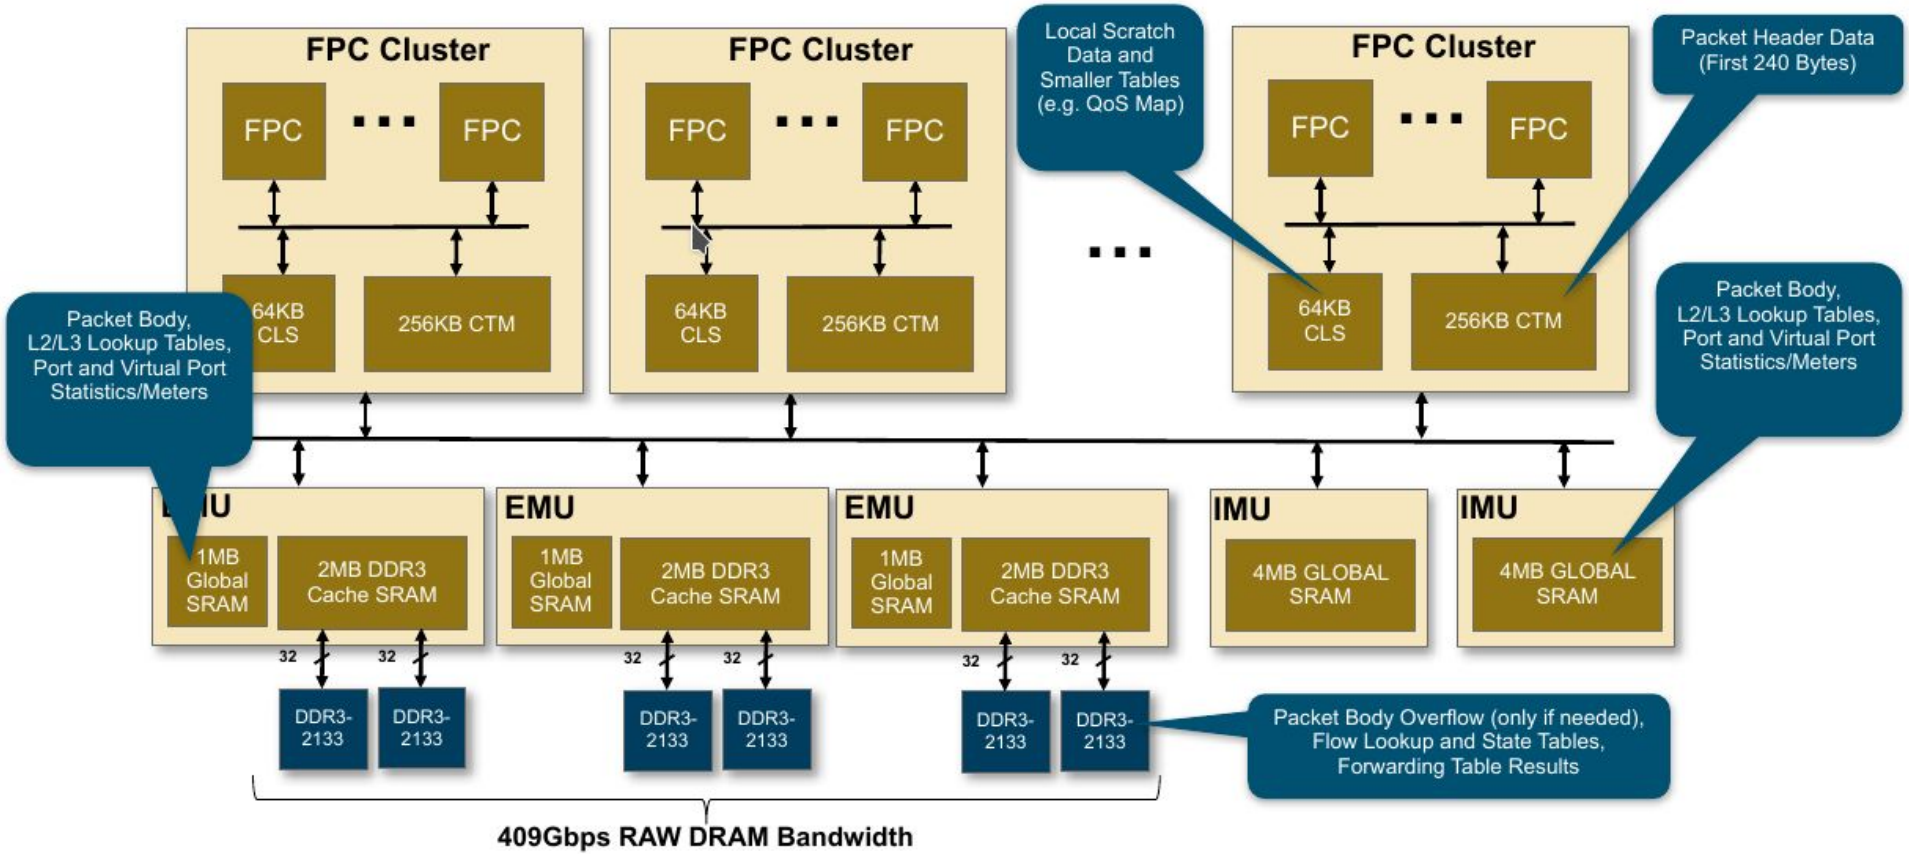

Programming Model

NETRONOME

Network Flow Processor 4xxx / 6xxx

- Highly parallel multithreaded architecture (8 threads / core) for high throughput
- Purpose built Microengines / Flow Processing Cores (72 / 120) maximize flexibility
- H/W accelerators further maximize efficiency (throughput/watt)

Fully software defined feature set - examples:

- Flexible tunneling support (e.g. VXLAN, GRE, VLAN, MPLS, NSH)
- Flexible Match/Action processing - many packet fields / protocols
- Highly scalable and fine grained security policies
- Network and PCIe SR-IOV / VirtIO RX/TX with stateless offloads
- Packet generation / reception with advanced statistics e.g. jitter
- Traffic directing (tapping / mirroring / steering / load balancing)

External DRAM accommodates millions of flows / rules

Convenient programmability using P4 and/or C

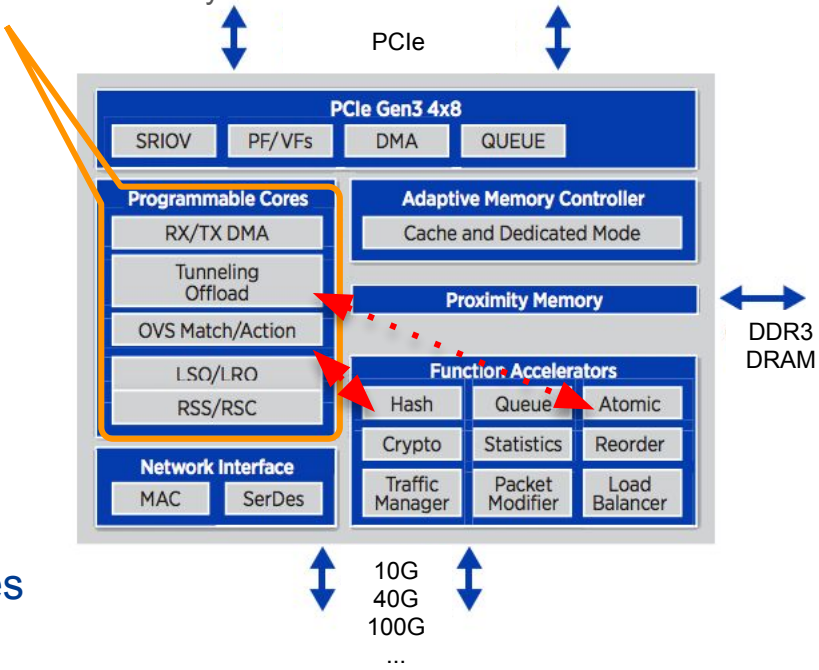

Programming Model - Network to Network

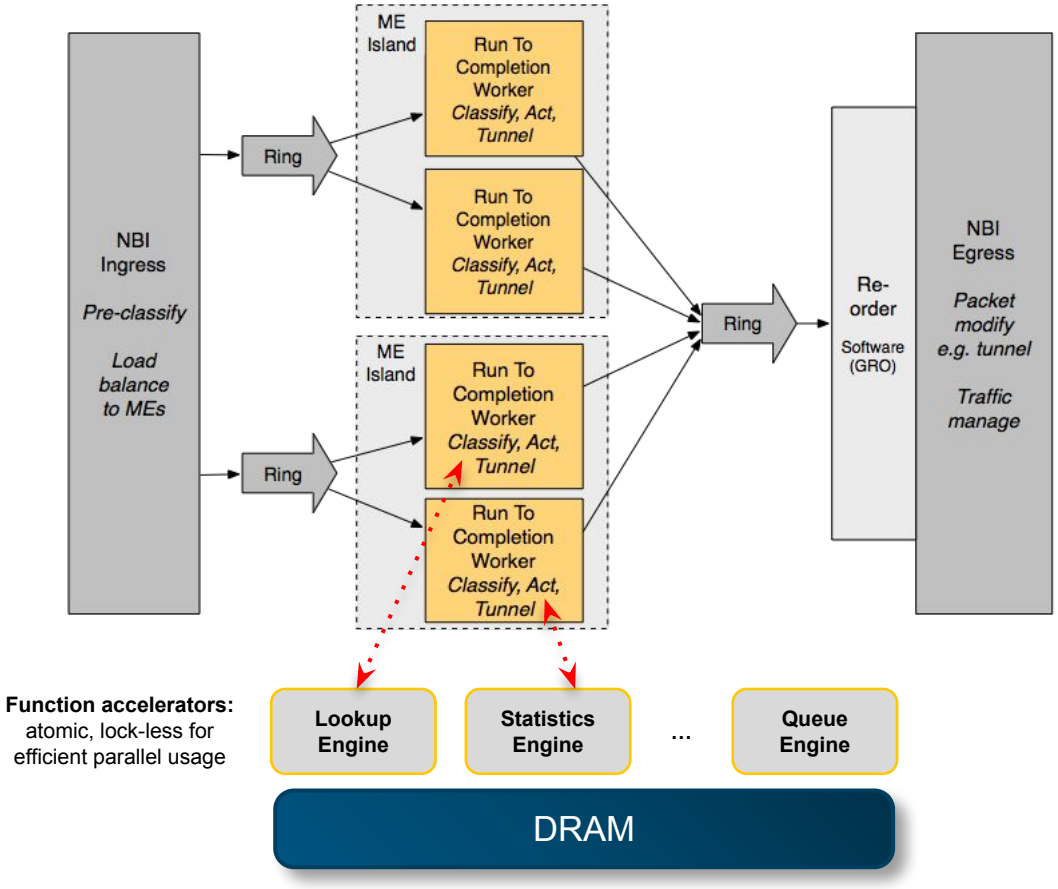

Programming Model - Details Incl. PCIe

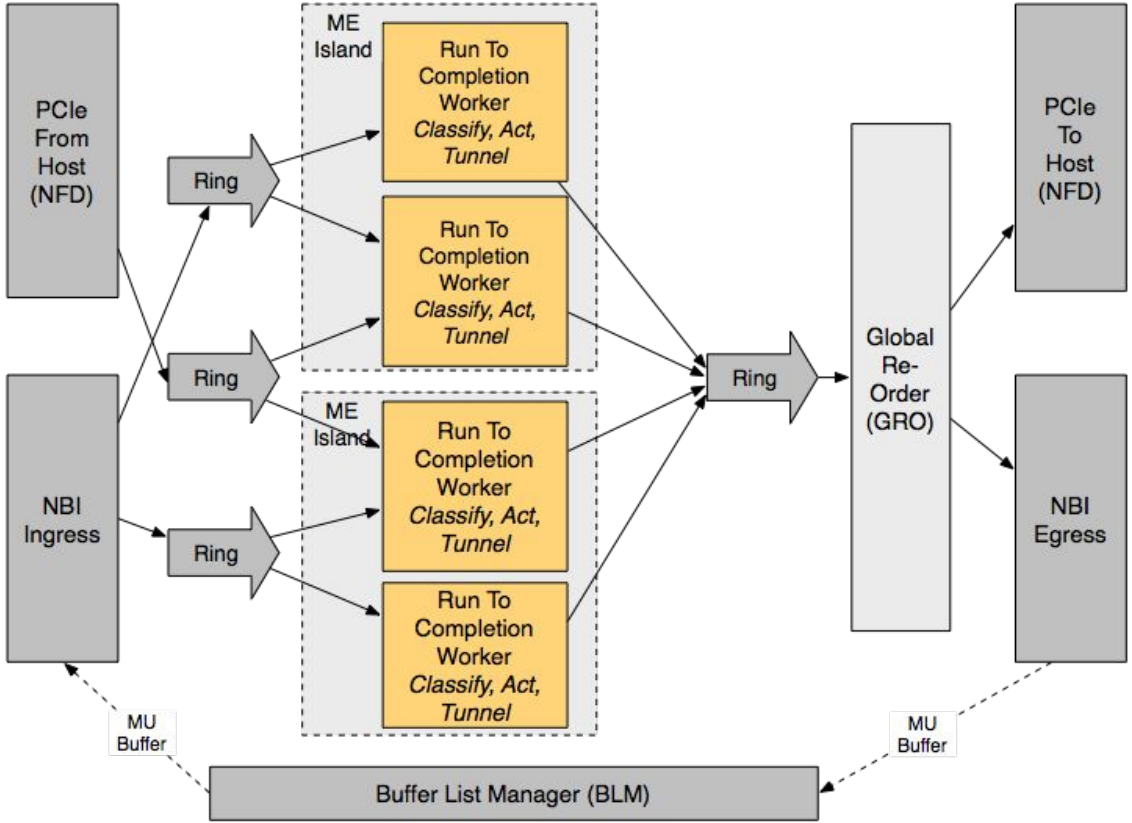

P4 Datapath

NETRONOME

- Load balancer distributes each packet to next available thread for optimum throughput
- Hardware assisted reordering ensures packet order is maintained
- Matching performed using various algorithms, e.g. DRAM-backed “algorithmic TCAM”
- Actions efficiently performed in on-chip memory
- Flow tracker statefully learns / tracks sessions, speeds up classification

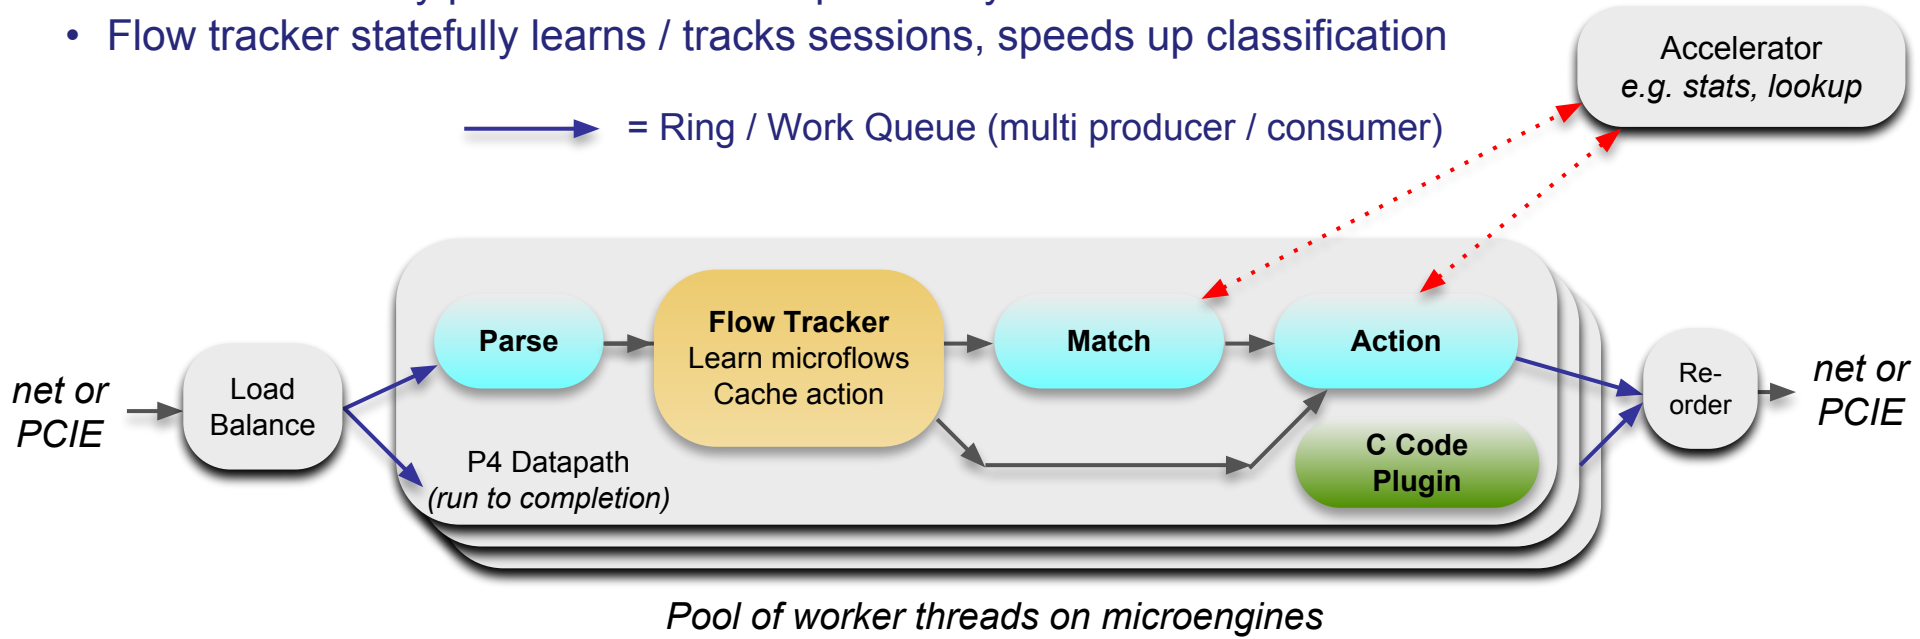

P4 Datapath

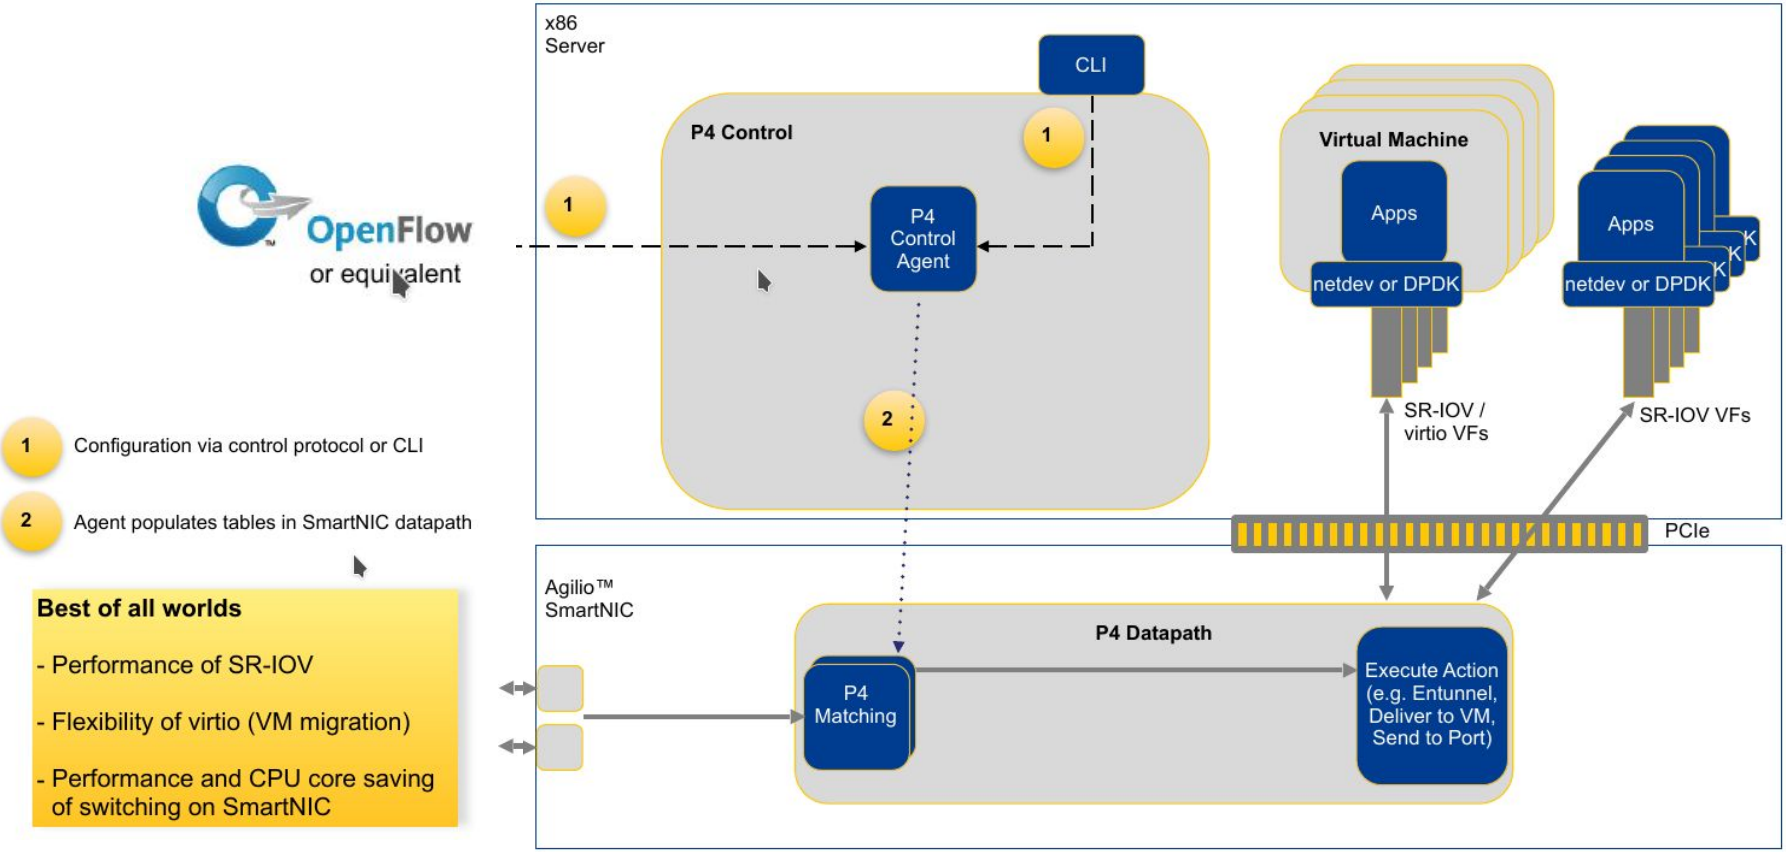

OpenFlow
or equivalent

1 Configuration via control protocol or CLI

2 Agent populates tables in SmartNIC datapath

Best of all worlds

- Performance of SR-IOV
- Flexibility of virtio (VM migration)
- Performance and CPU core saving of switching on SmartNIC

Example P4 Application

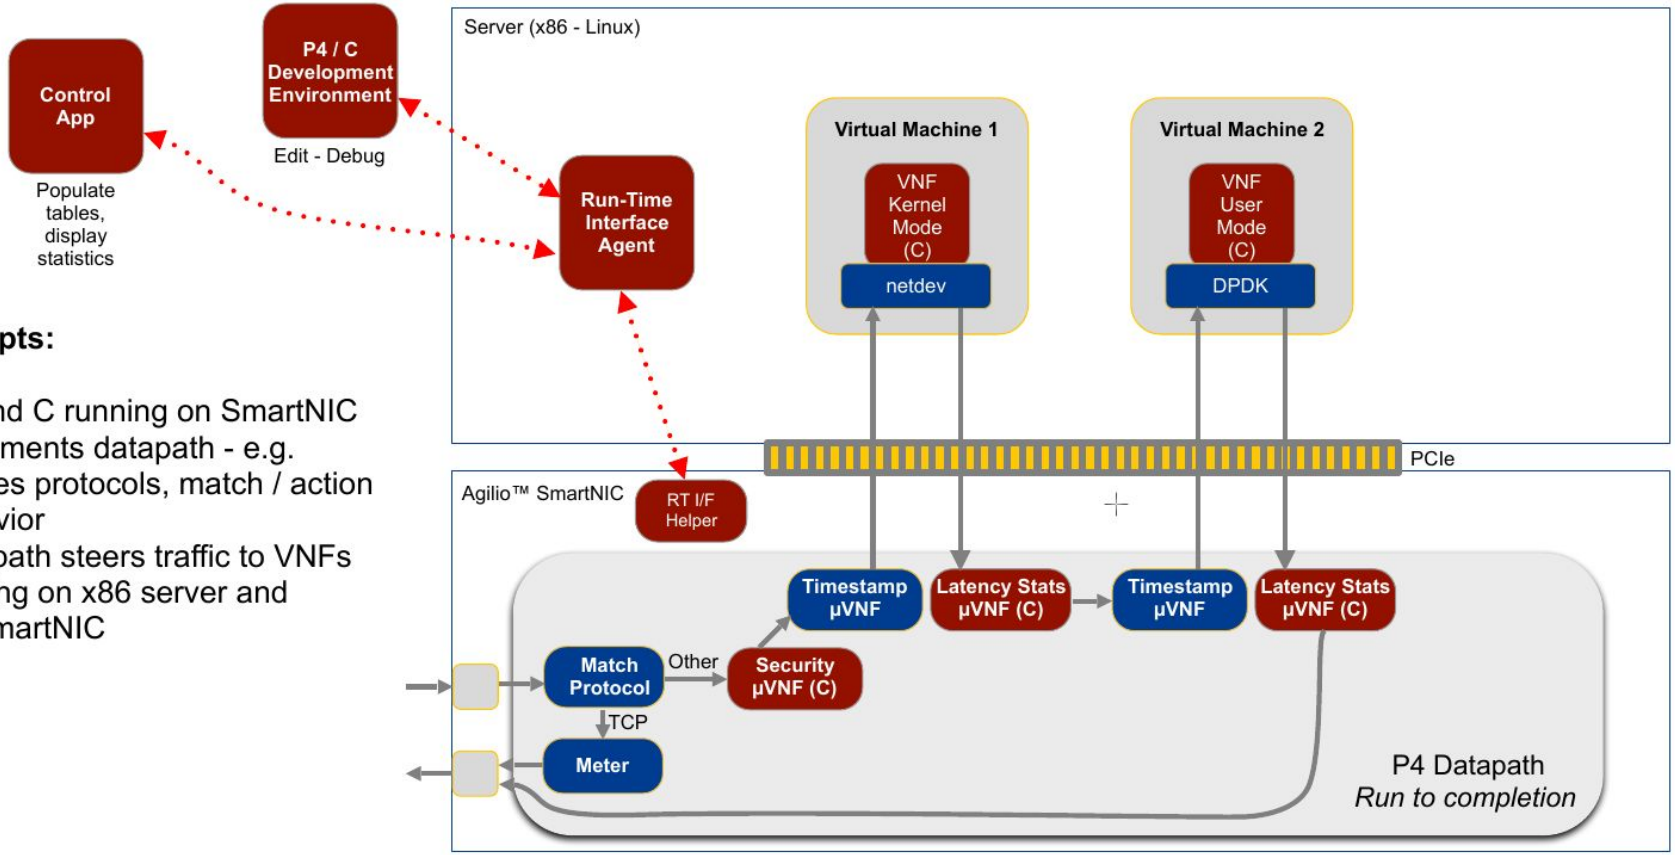

Concepts:

- P4 and C running on SmartNIC implements datapath - e.g. defines protocols, match / action behavior
- Datapath steers traffic to VNFs running on x86 server and on SmartNIC

Tools

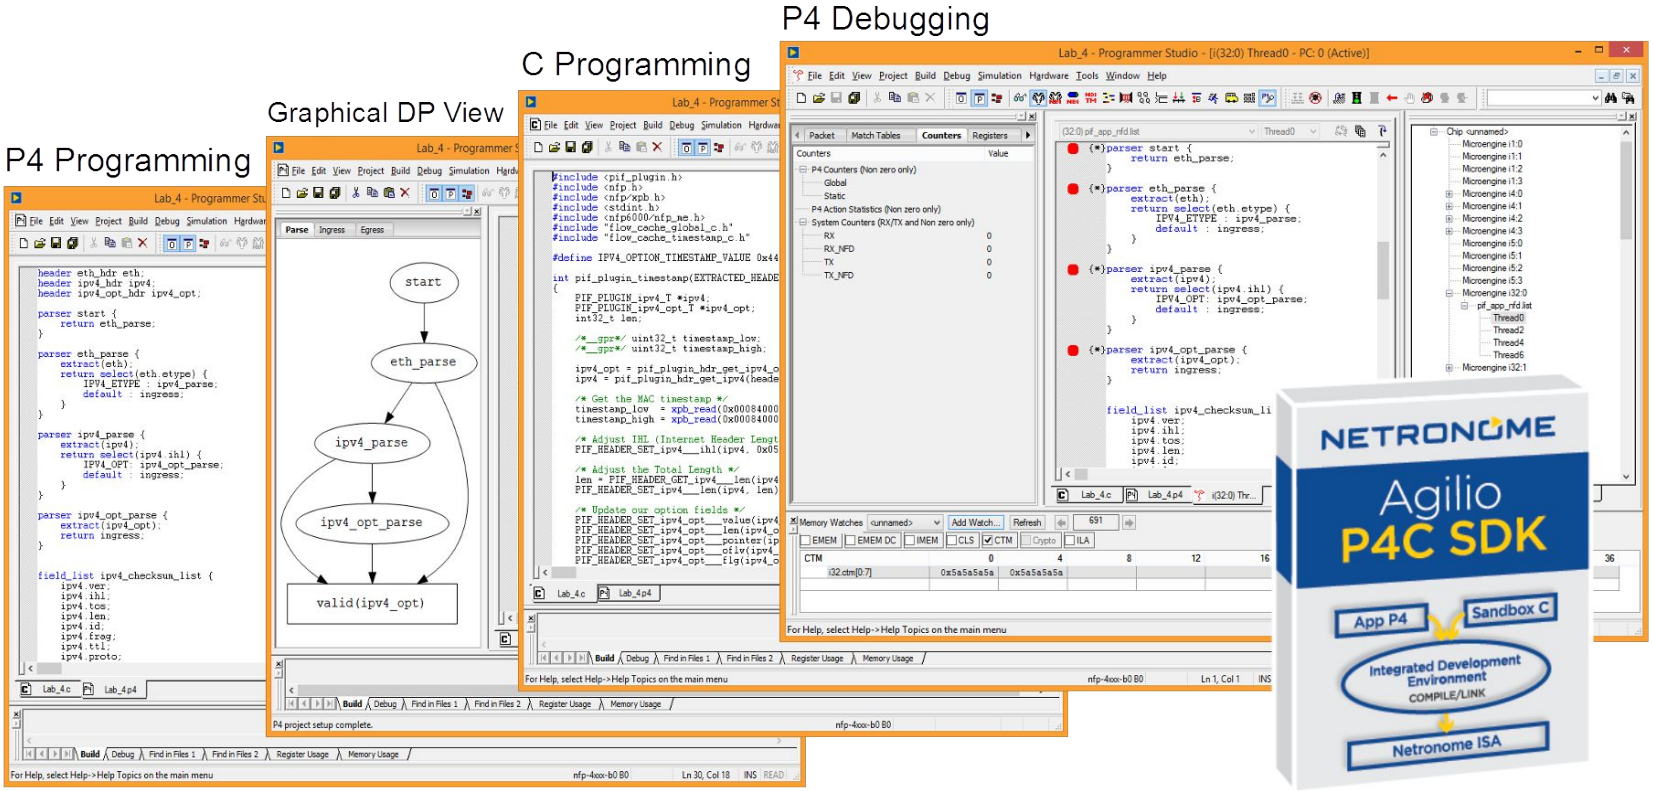

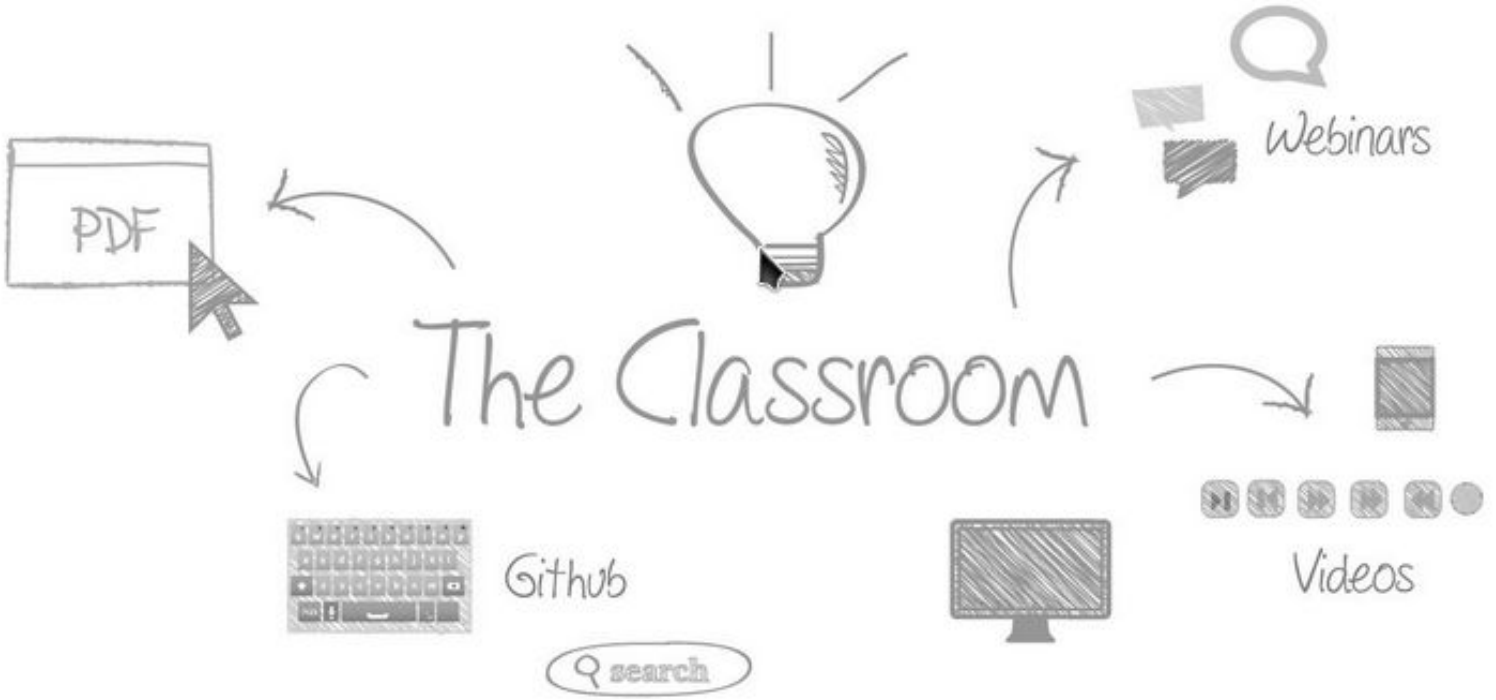

Performance Results

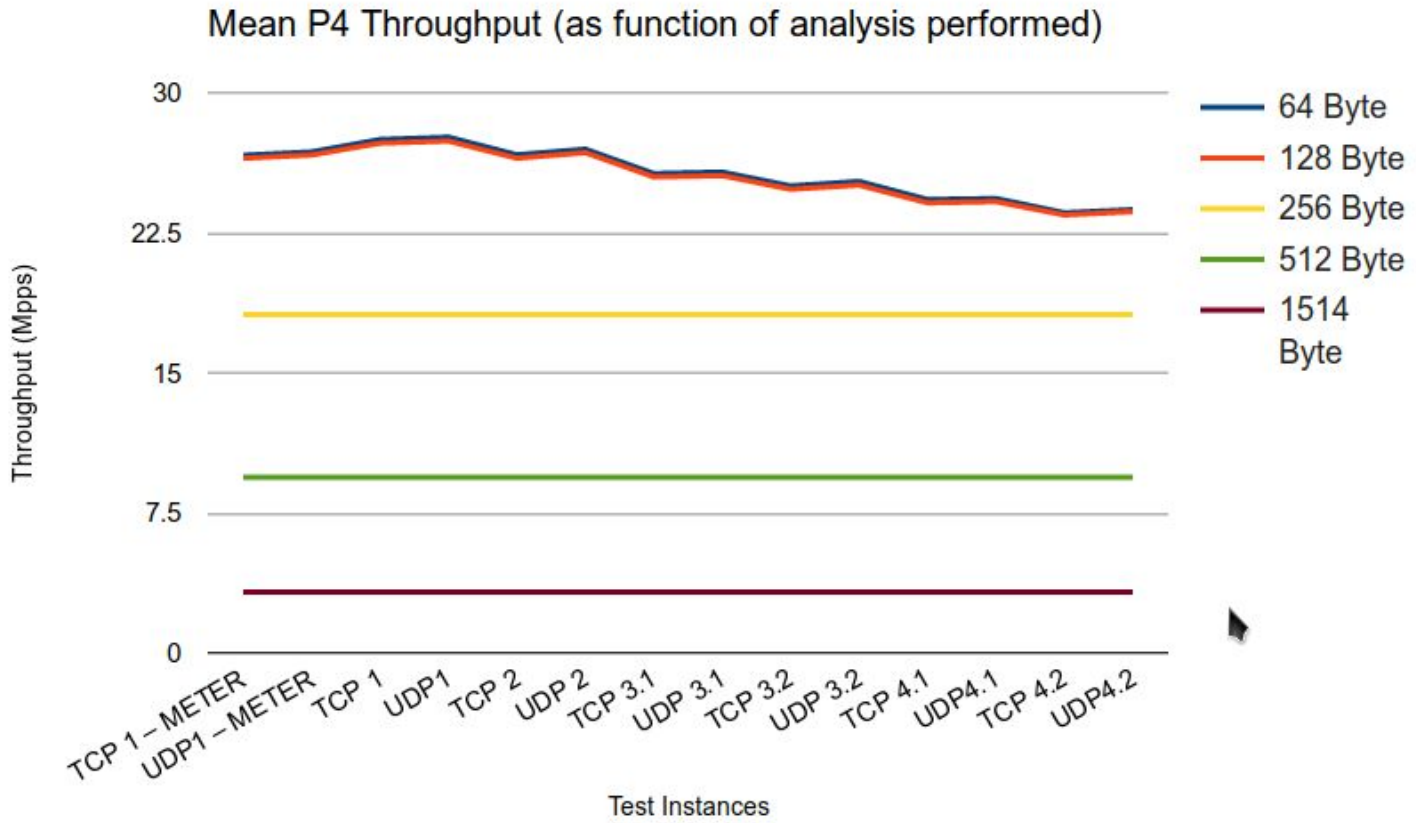

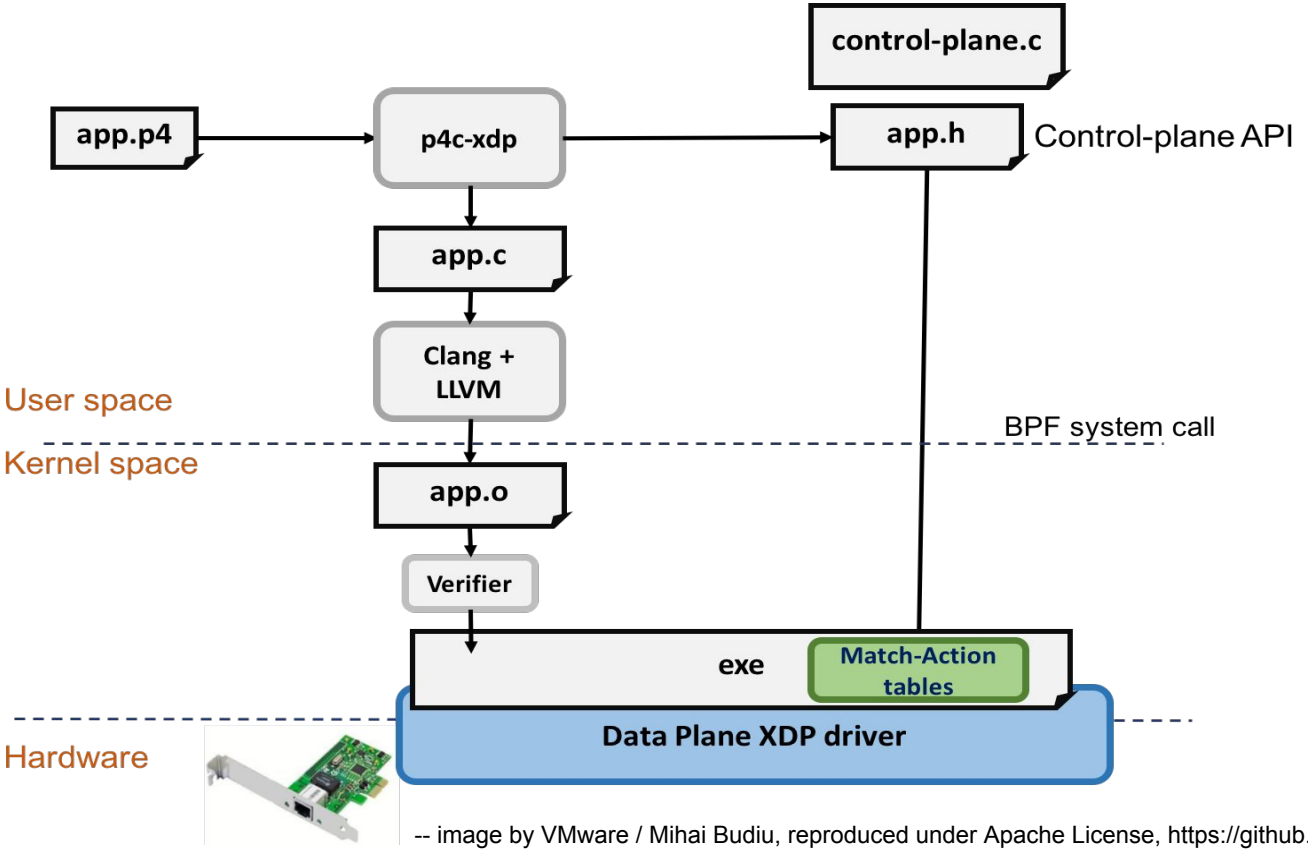

-- image by VMware / Mihai Budiu, reproduced under Apache License, https://github.com/vmware/p4c-xdp

Further Reading / Q&A

NETRONOME

- open-nfp.org
- Netronome white papers
  - www.netronome.com/media/redactor_files/WP_Programming_with_P4_and_C.pdf
  - www.netronome.com/media/redactor_files/WP_P4.pdf
  - www.netronome.com/media/redactor_files/WP_NFP_Programming_Model.pdf
- github.com/vmware/p4c-xdp

Questions?