FPGA Augmented Network Interfaces

SOLARFLARE®

Steve Pope – Solarflare CTO

Application Onload Engine (SFA6902)

- Out of the box, it is a Solarflare high performance network adapter
  - Drivers included
  - Works with kernel network stack and kernel bypass (OpenOnload)
- FPGA acceleration by Incremental upgrade
  - Pass-thru by default
  - Accelerate a subset of traffic
  - No new switches, cabling, slot
- Solarflare & 3rd party applications
  - Solve common problems
  - Shrink wrapped
- FDK (developer kit)
  - Reusable IP blocks to minimise effort for FPGA developers

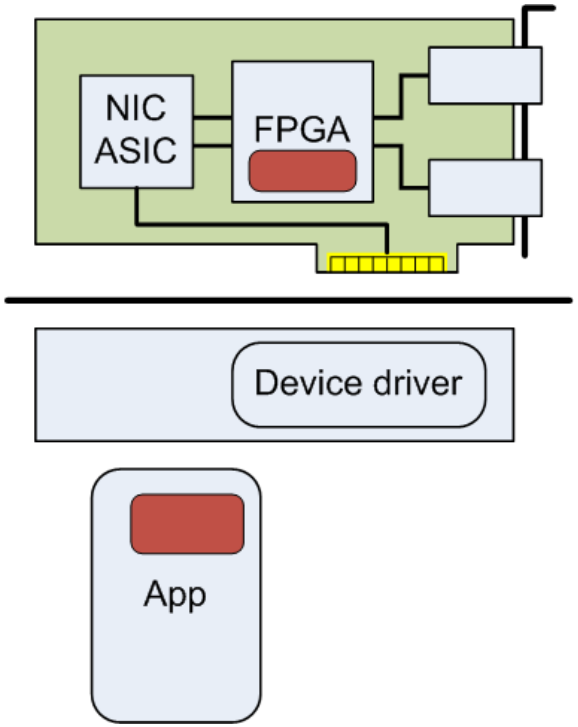

FDK Concept

- Solarflare instantiates and qualifies all external interfaces and presents a Layer2 streaming interface to application logic

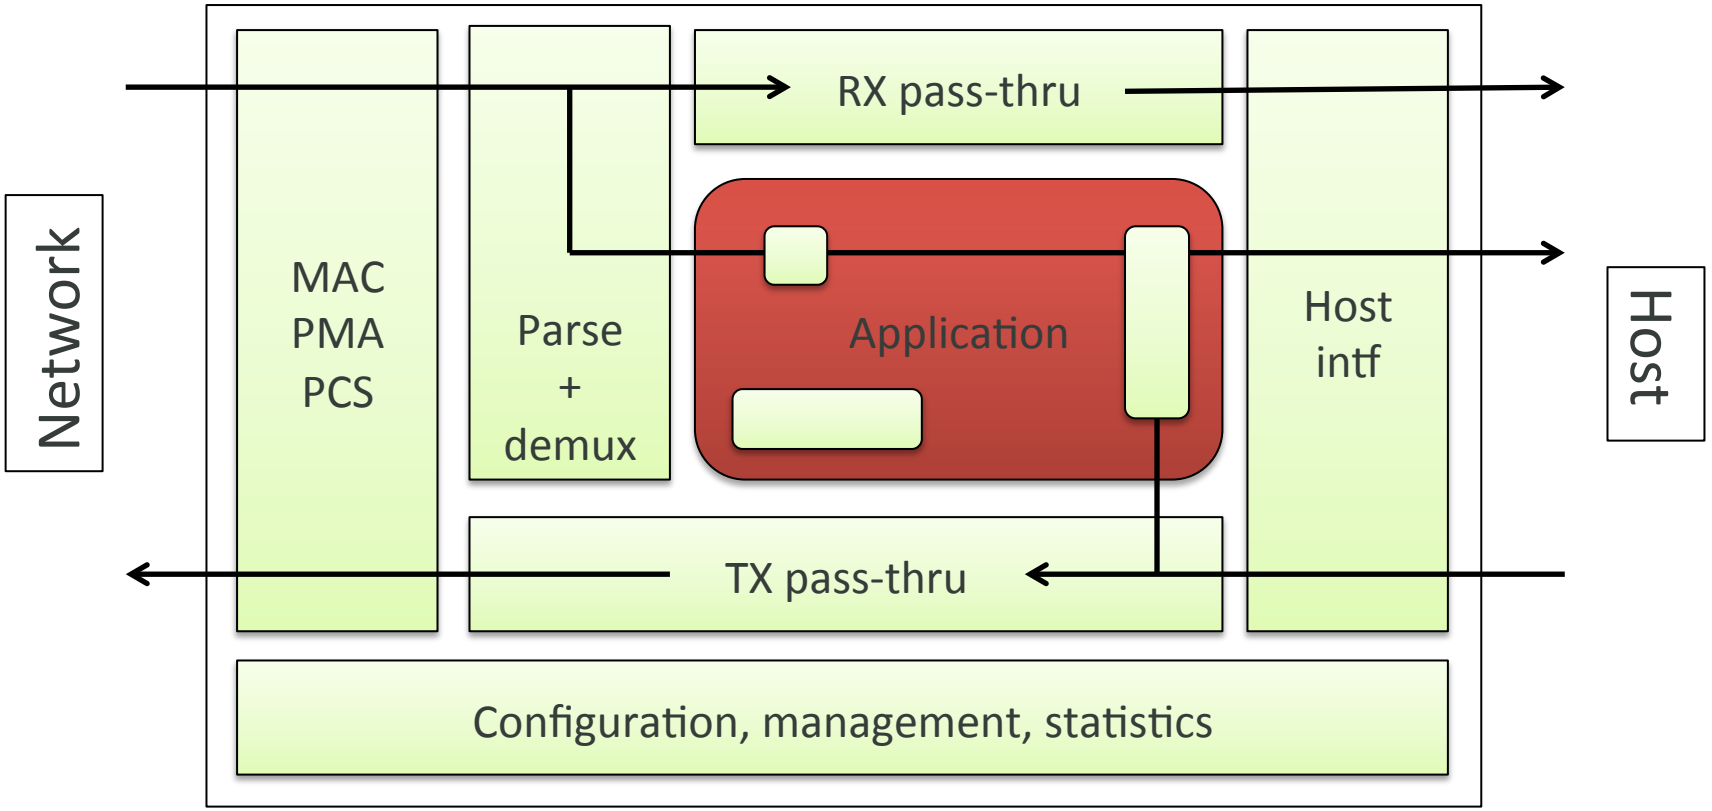

Qsys Application Integration

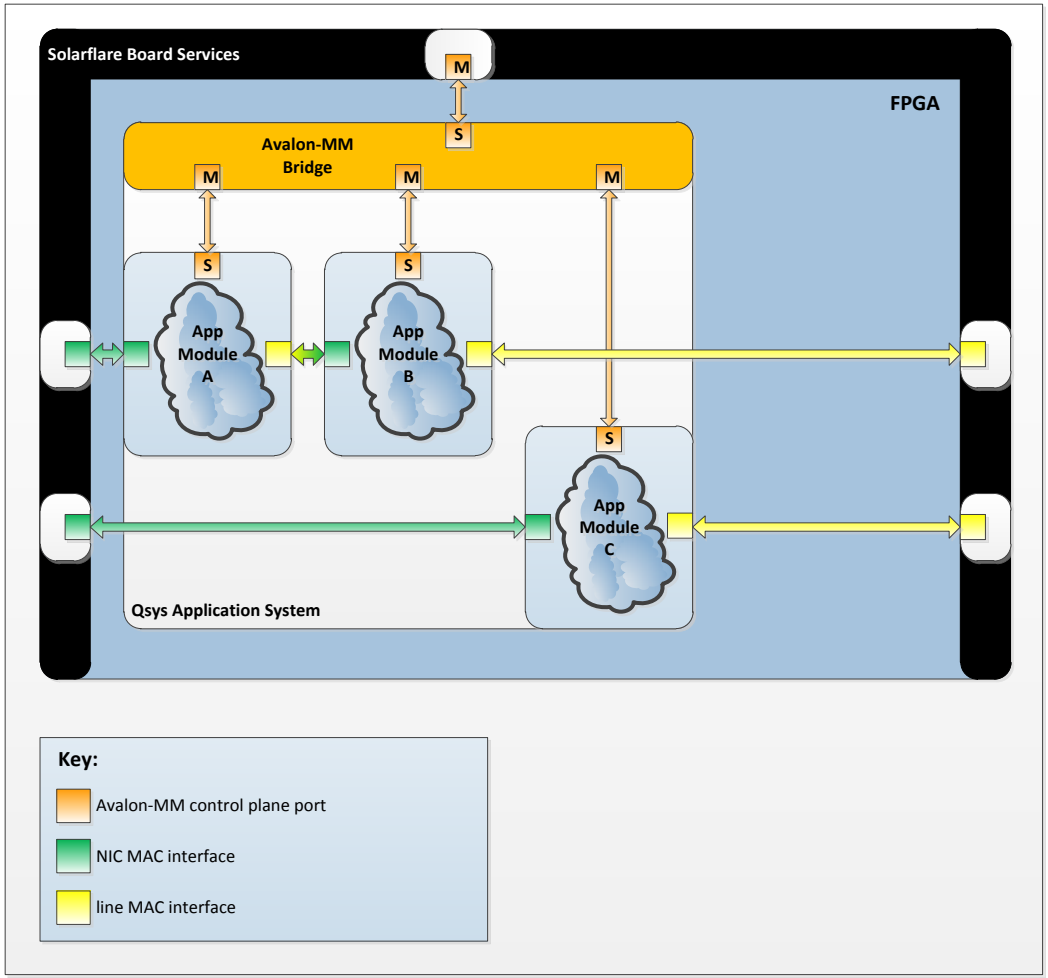

- Board Support Services
  - PMA link tuning
  - MAC statistics driver accessible
  - Avalon-ST interfaces to the MAC from FPGA fabric
  - Avalon-MM interface for side-band control (1GB registers mapped to host user-space)
- All Interfaces are Qsys compatible
  - Users can rely on Qsys for inter-module binding
- Follow the Qsys design flow!

FDK Worked Example:

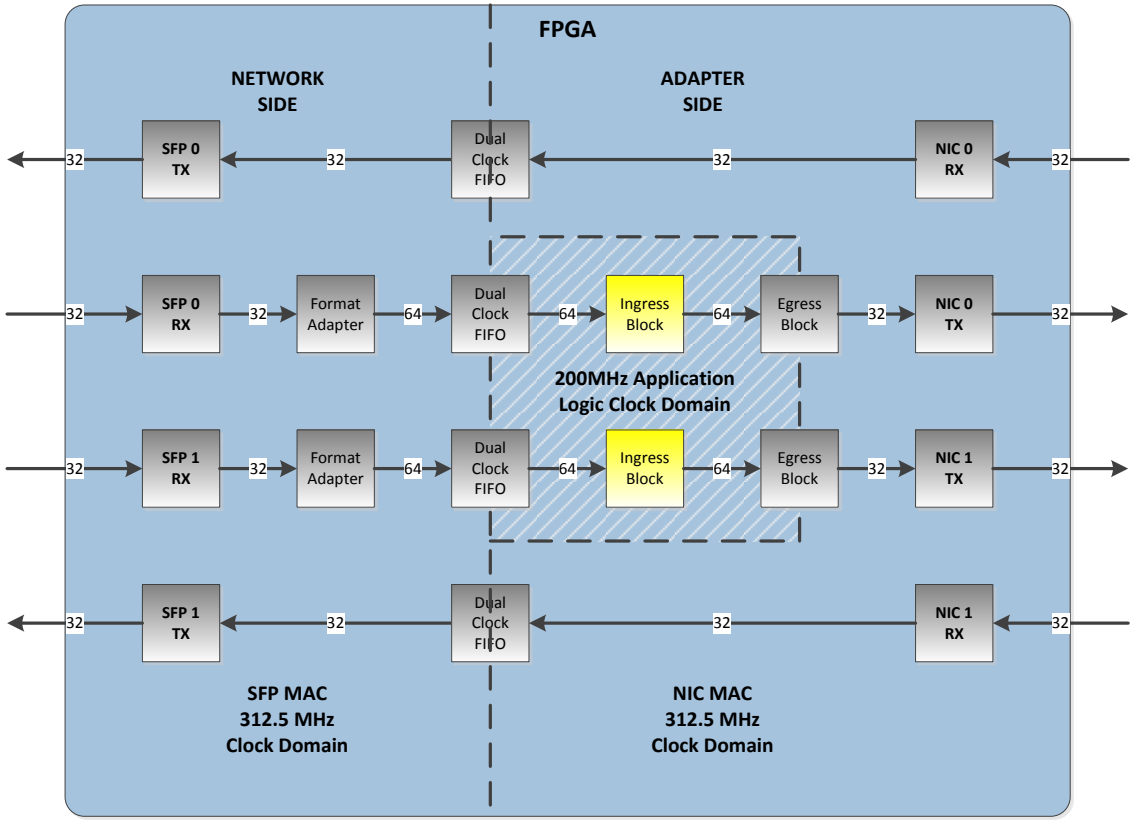

- Data conversion to 64bit/200Mhz
  - Including metadata / ECC
- Ingress parser
  - IPv4, GRE, UDP
- Field extraction
- Context lookup
  - Appended to meta-data
- Egress
  - Strip metadata / ECC for streaming to host
- Host interface
  - POSIX socket

The Host Interface

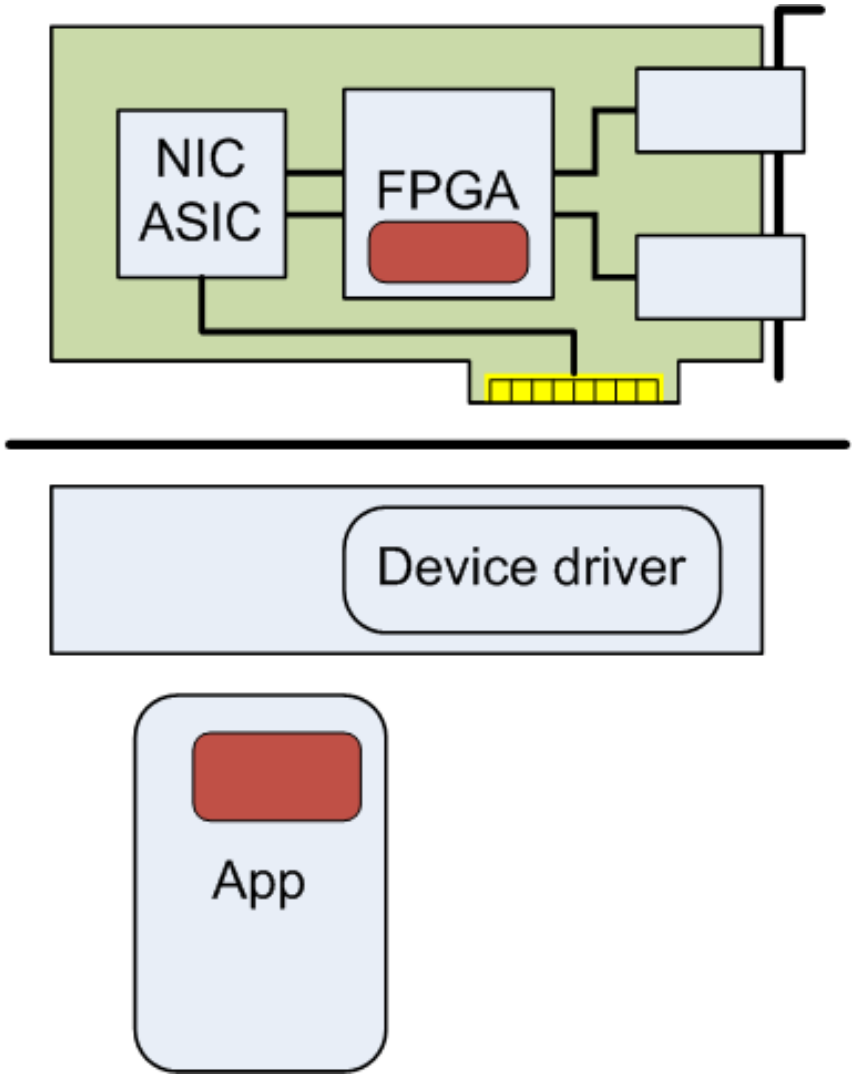

Kernel mode Host Interface

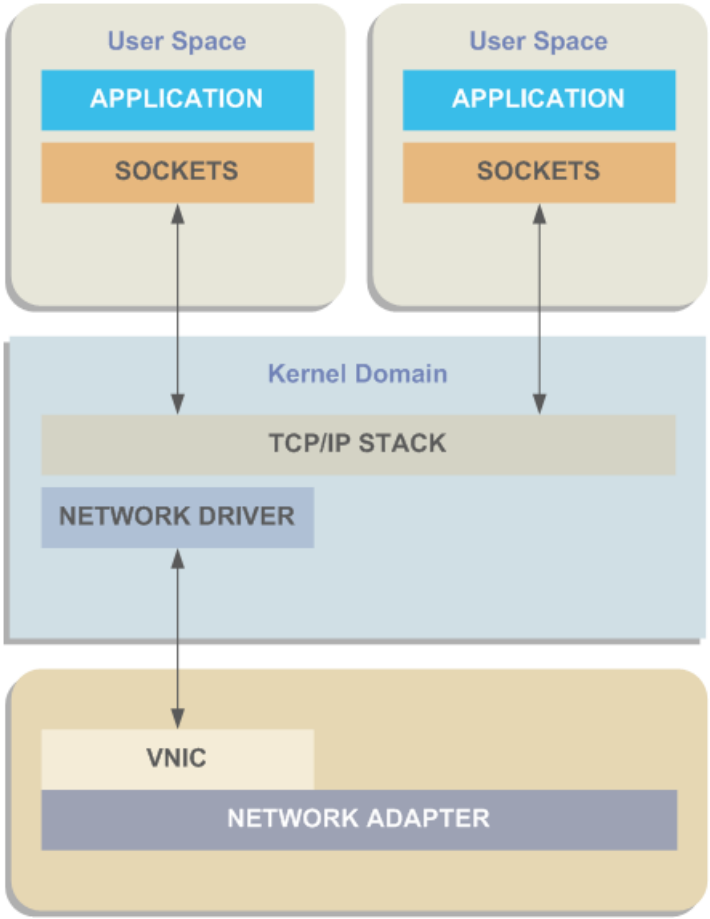

- Standard network model:
  - Good performance
  - < ~5us latency
  - ~1,000,000 msg/sec/core
- Drivers for Linux / ESX / Windows ..

High Performance Host Interface: Kernel bypass

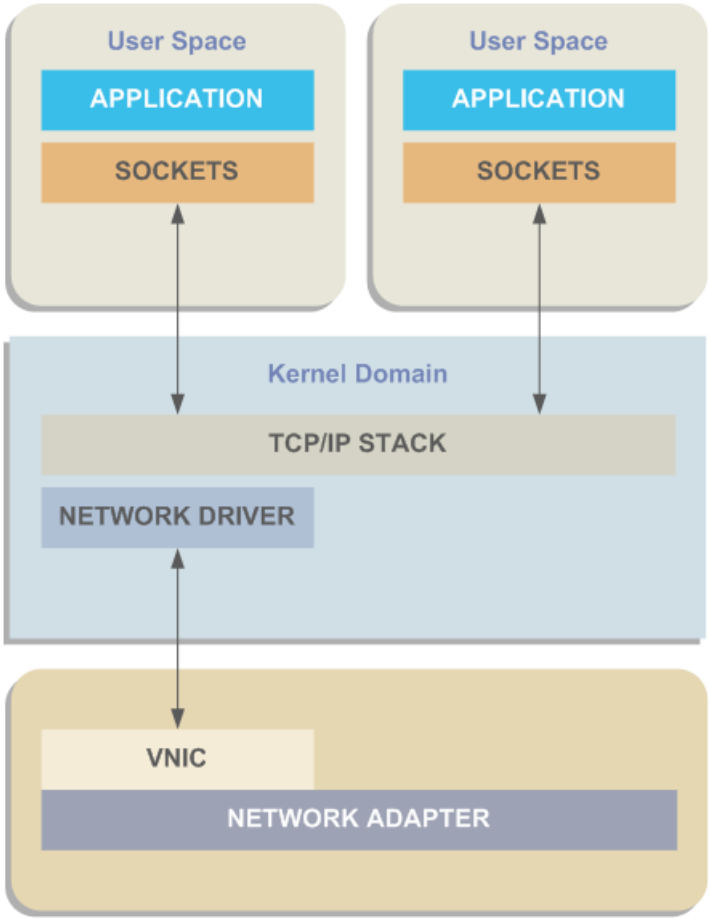

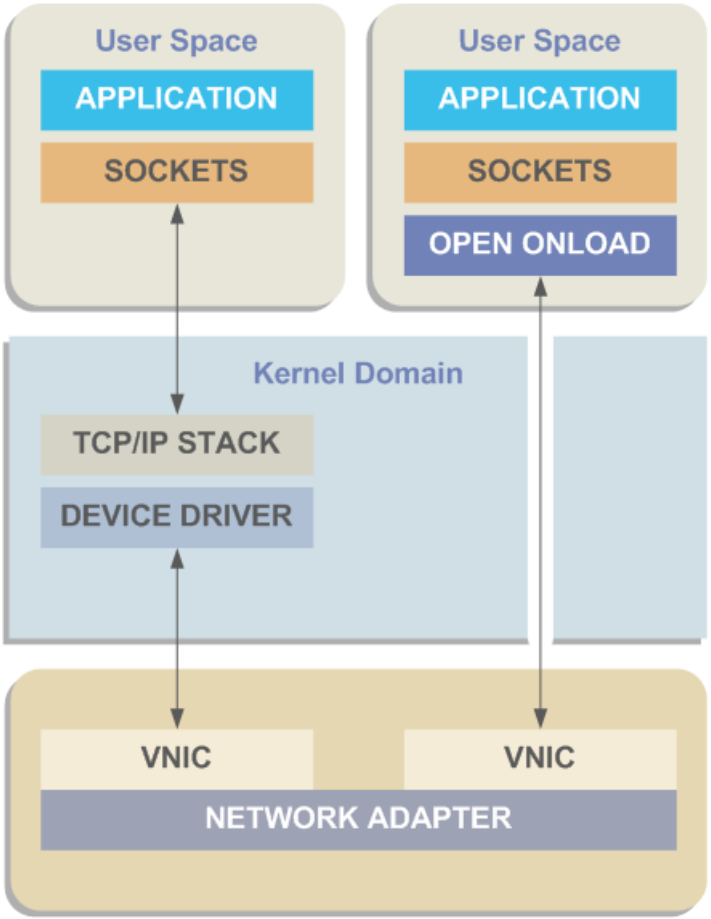

High Performance Host Interface: Kernel bypass

- Performance like Infiniband, but with:
  - Standard BSD sockets API
  - Standard protocols (TCP, UDP)
  - regular Ethernet
- Performance
  - < 2us latency
  - > 4,000,000 msg/sec/core
- Very easy to use
  - No changes needed to applications

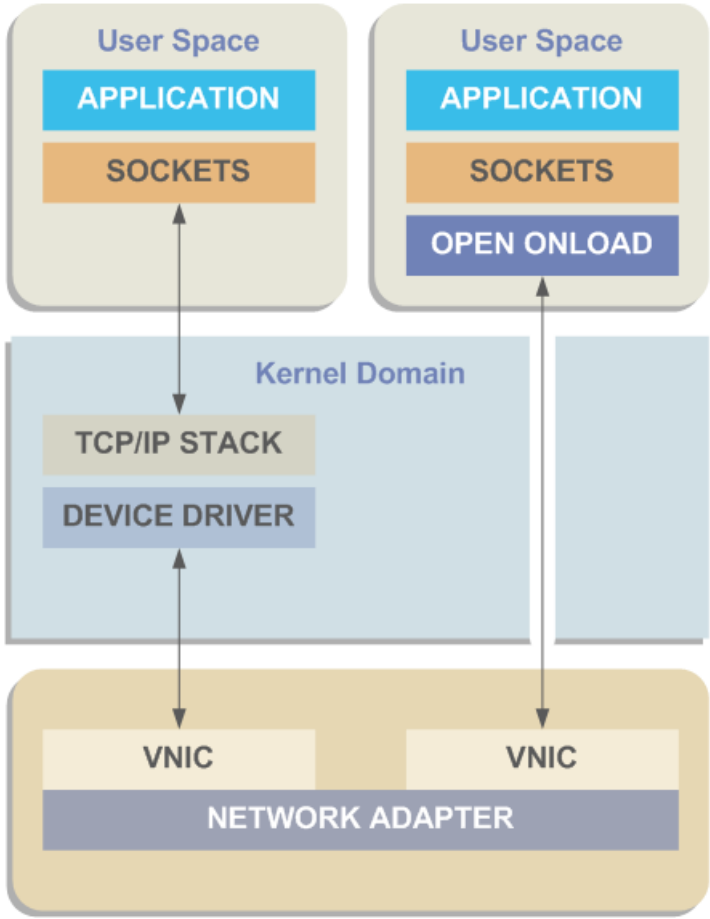

Kernel bypass makes a really big difference…

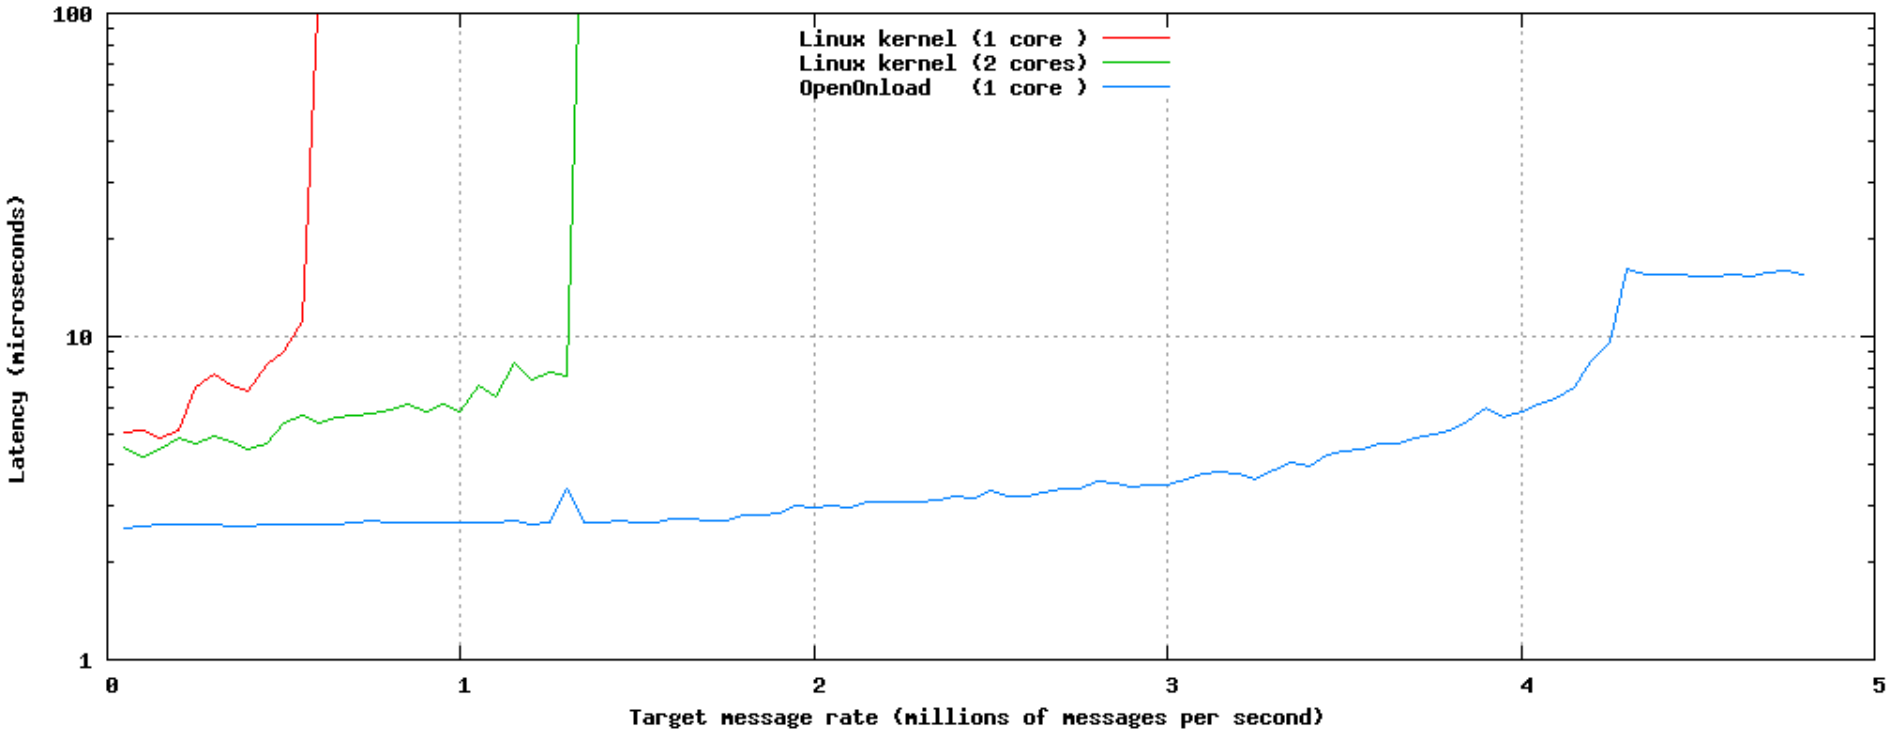

Simultaneous hypervisor and kernel bypass

- SR-IOV allows PCI functions to be passed through to a virtualixed guest operating system
- Ultra-low latency and extreme message rates in a virtual machine
- Enables isolation and fast provisioning of accelerated hardware applications
- Used in trading applications over a Cloud (service) based delivery model
- Under investigation for telecom application virtualisation architectures

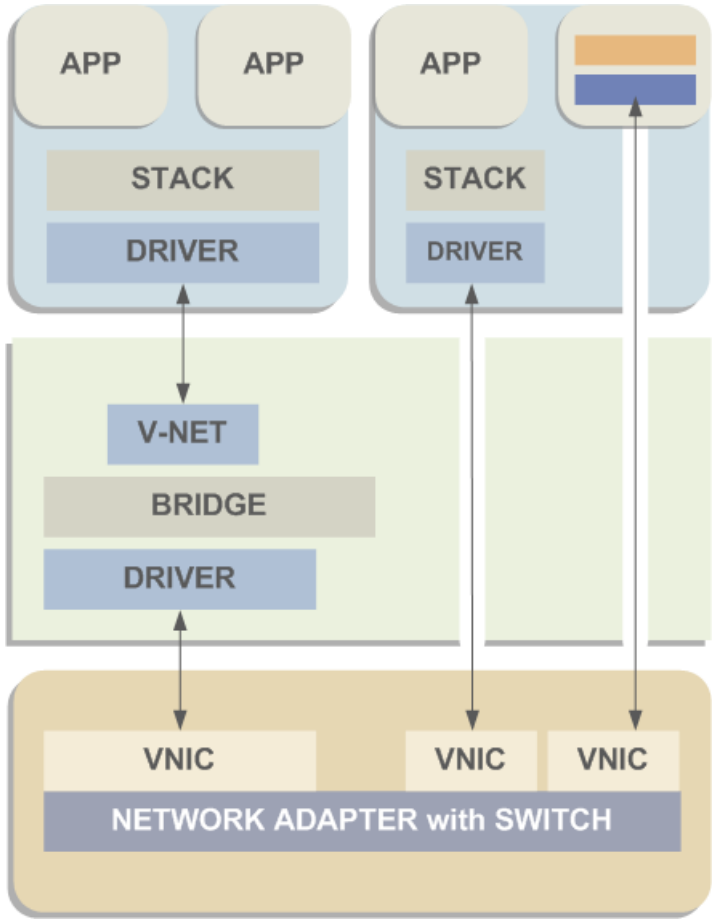

Solarflare University Program

- Overview
  - Program fosters collaborative development, deployment and promotion of AOE applications
  - Open to faculty, researchers and students at university and education institutions worldwide
  - Provides FPGA-based products for classroom instruction in computer science and computer engineering helping to pioneer Custom Compute
- Program benefits
  - Access to the Solarflare AOE platform and firmware development kits at a substantial educational discount
  - Eligibility to enter the Custom Compute Innovation Awards to encourage and reward research, innovation, and application development
  - Connection to Solarflare AOE Partner Network
- Complements Altera University Program
- 50 universities worldwide

Solarflare University Members

APAC

EMEA

Americas

Clemson University

- SFAOENIDS
  - Research project for accelerating network intrusion detection with the AOE
  - Capable of parsing SNORT rules and implementing reconfigurable pattern matching on the AOE
  - Enables user to reprogram the hardware as the ruleset changes without having to resynthesize the hardware
  - String matching unit only requires a small amount of memory to save/resume its state
  - Runs at 10Gbps
  - Open source – being posted on github

Columbia University

- Embedded Systems Course
  - 5 financial projects on the AOE
  - 3 completed last semester
    - Market data processing/ticker plants for NASDAQ ITCH Decoder and Book Builder and CME FIX/FAST Decoder and Book Builder
    - Inline Risk System ("Kill Switch")
    - Projects posted on Columbia and Solarflare websites
  - 2 projects still ongoing
    - FIX Engine (Order Entry) and TCP offload engine
- Future research under consideration includes:
  - Linux (Nios II-based) running inside the AOE
  - Optical memory / pure optical routing research
  - Ultra low latency FPGA-based order matching engine ("Exchange on a Chip")