



April 26-28, 2022

DoubleTree by Hilton San Jose

SmartNICsSummit.com

Platform Choices for FPGA-Based In-Network Compute Acceleration

Endric Schubert, Ph.D.

CTO

Missing Link Electronics

Key Contributors

Alex Forencich, Ph.D. - UC San Diego

Ulrich Langenbach, Dir. Eng. Missing Link Electronics



Dr. David Boggs
1950 - 2022

Co-Inventor of
Ethernet



Backgrounder MLE

Mission: “If It Is Packets, We Make It Go Faster!”

High-Performance (Embedded) Compute & Connected Systems-of-Systems

- PCIe (CXL, NVMe)
- Ethernet (TCP/IP, TSN)
- Audio/Video (HDMI, SDI)



San Jose, CA April 26-28, 2022


MLE Technology & Manufacturing Partnerships

The Fraunhofer-Gesellschaft undertakes applied research of direct utility to private and public enterprise and of wide benefit to society.



25,327 staff

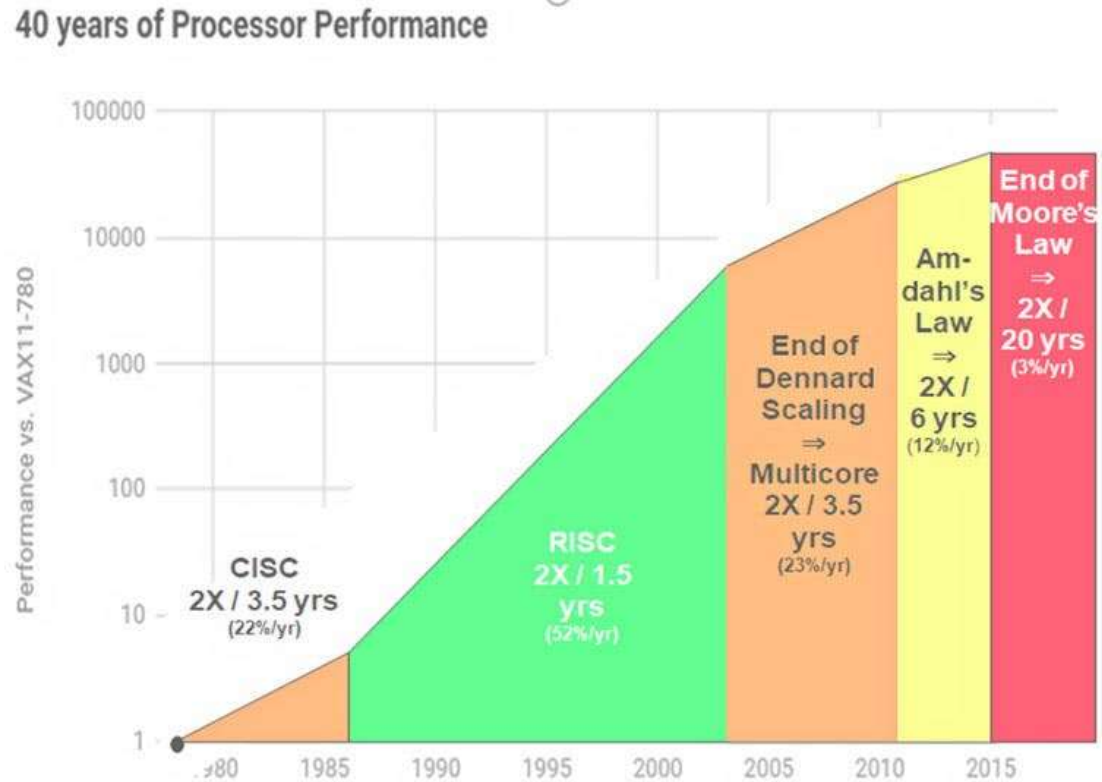
72 institutes and research units



FPGAs Great for Data-in-Motion Processing



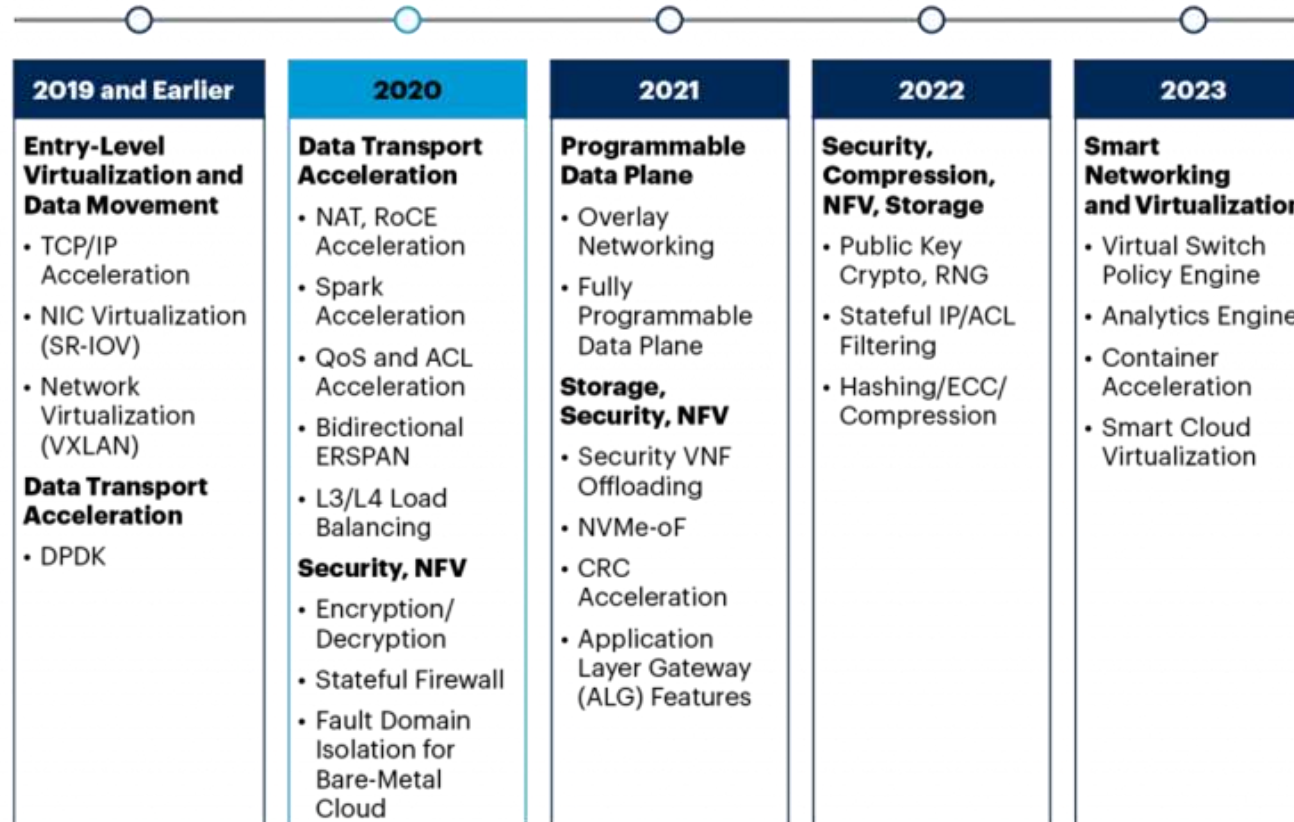
The Need for Domain-Specific Architectures



Network Port Speeds Outstrip CPU Performance



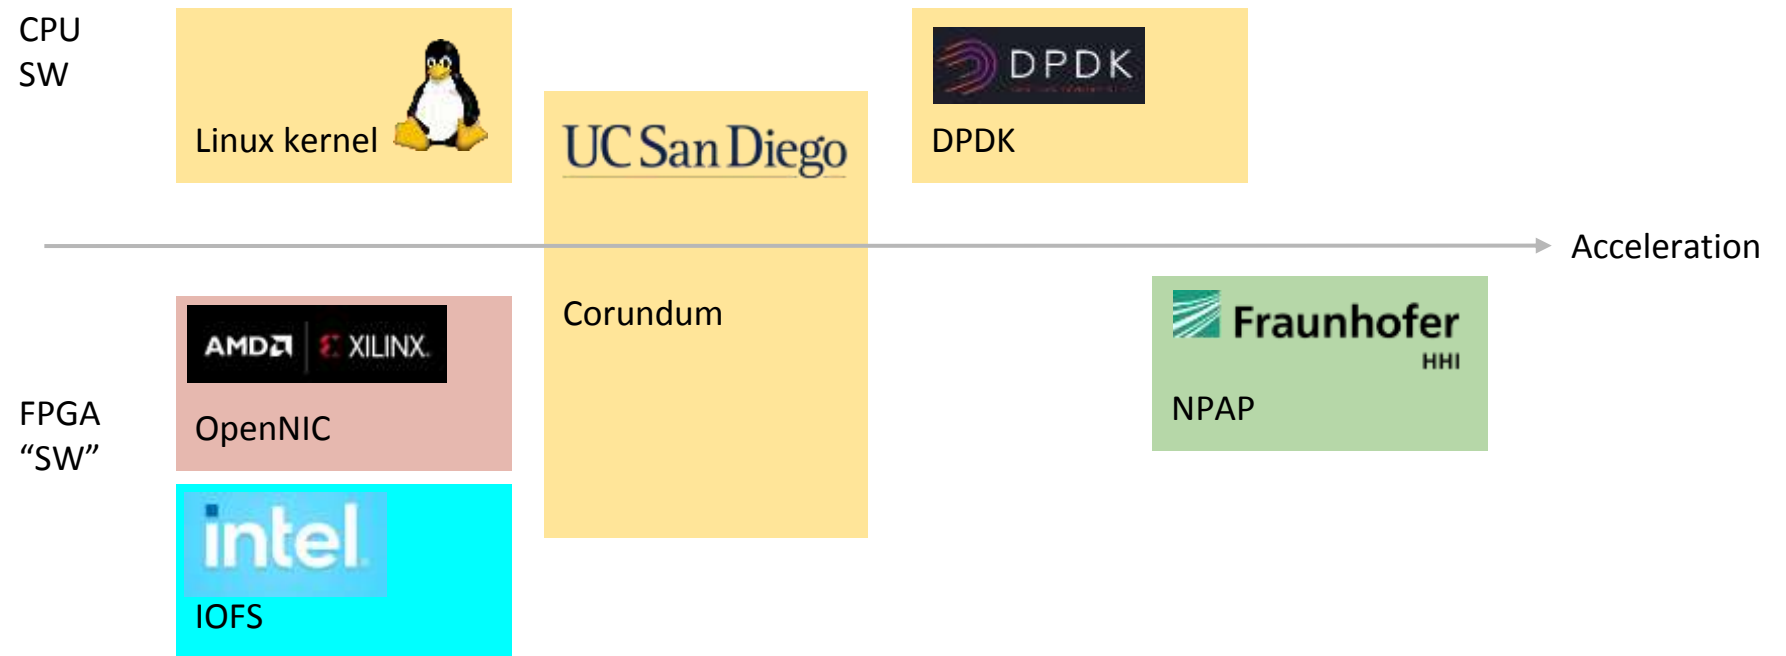
Evolution of Function Accelerators / SmartNICs



Source: Gartner
717089_C

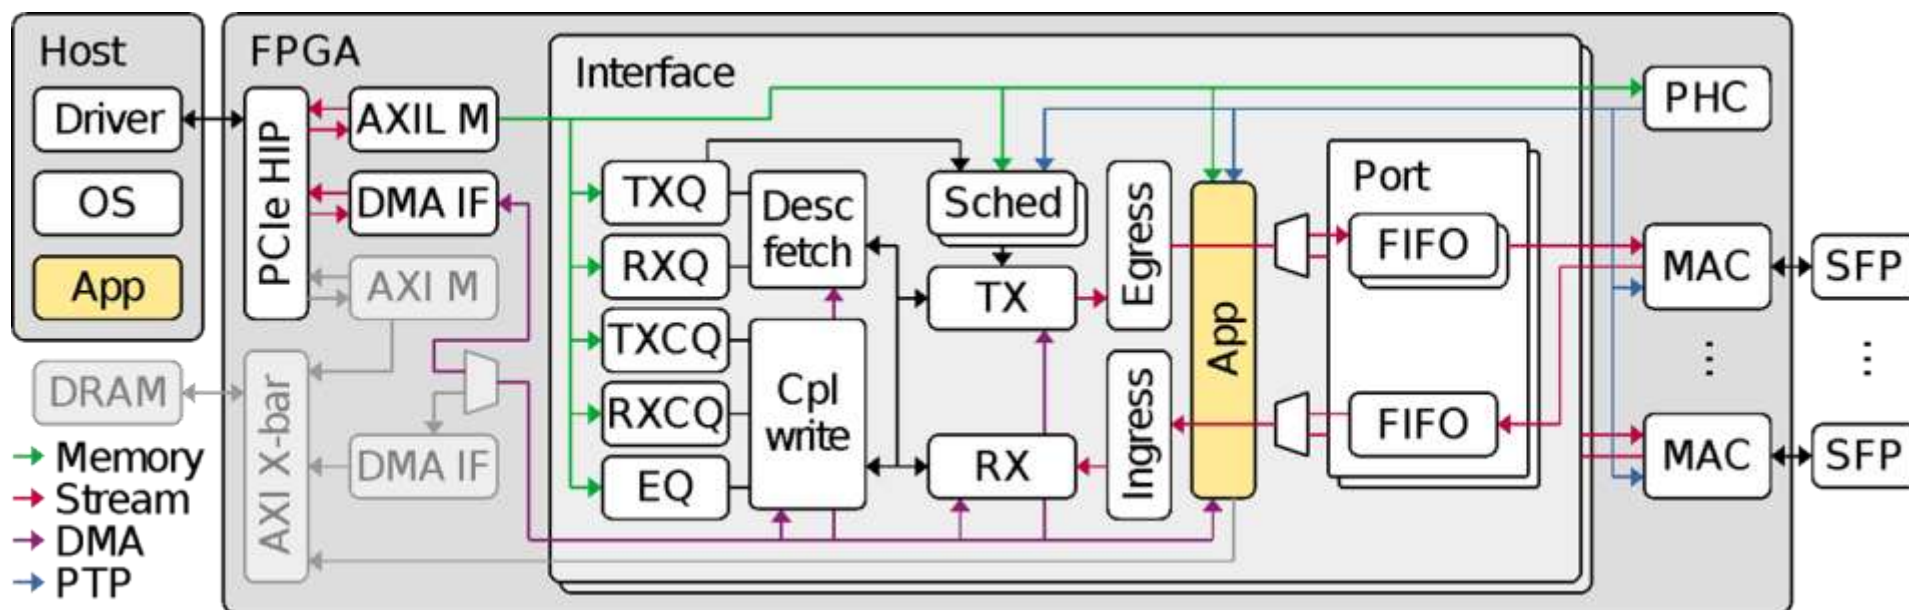
Platforms to Reduce Complexity & to De-Risk

FPGA programming requires special expertise
need for high levels of optimization which makes “App Store” approach difficult.
⇒ Platforms enable small expert teams to deliver solutions more rapidly!



Corundum Architectures

Open-source, FPGA-based NIC and platform for in-network compute
Full System Stack implementing a Data Stream Oriented Architecture



Corundum Features

- Open-source, high-performance, FPGA-based NIC
 - PCIe Gen3 x16, multiple 10G/25G/100G Ethernet ports
 - Fully custom, high-performance DMA engine; Linux driver
- Application block for custom logic
 - Access to network traffic, DMA engine, on-card RAM, PTP time
- Fine-grained traffic control
 - 10,000+ hardware queues, customizable schedulers
- PTP timestamping and time synchronization
- Management features (FW update, etc.)
- Wide device support (AMd/Xilinx and Intel)
- Source code: <https://github.com/corundum/corundum>

Corundum Hardware Support & Services

- Alpha Data ADM-PCIE-9V3 (Xilinx Virtex UltraScale+ XCVU3P)
- Exablaze ExaNIC X10/Cisco Nexus K35-S (Xilinx Kintex UltraScale XCKU035)
- Exablaze ExaNIC X25/Cisco Nexus K3P-S (Xilinx Kintex UltraScale+ XCKU3P)
- Silicom fb2CG@KU15P (Xilinx Kintex UltraScale+ XCKU15P)
- NetFPGA SUME (Xilinx Virtex 7 XC7V690T)
- Intel Stratix 10 MX dev kit (Intel Stratix 10 MX 1SM21CHU1F53E1VG)
- Xilinx Alveo U50 (Xilinx Virtex UltraScale+ XCU50)
- Xilinx Alveo U200 (Xilinx Virtex UltraScale+ XCU200)
- Xilinx Alveo U250 (Xilinx Virtex UltraScale+ XCU250)
- Xilinx Alveo U280 (Xilinx Virtex UltraScale+ XCU280)
- Xilinx VCU108 (Xilinx Virtex UltraScale XCVU095)
- Xilinx VCU118 (Xilinx Virtex UltraScale+ XCVU9P)
- Xilinx VCU1525 (Xilinx Virtex UltraScale+ XCVU9P)
- Xilinx ZCU106 (Xilinx Zynq UltraScale+ XCZU7EV)

Growing list pre-compiled and tested systems stacks for COTS FPGA Cards



[Home](#) [User Guides](#) [App Notes](#) [Remote Evals](#) [LXR](#)

Corundum FPGA NIC Project

Corundum is an open-source FPGA-based NIC which features a high-performance datapath between multiple 10/25/50/100 Gigabit Ethernet ports and the PCIe link to the host CPU. Corundum has several unique architectural features: For example, transmit, receive, completion, and event queue states are stored efficiently in block RAM or ultra RAM, enabling support for thousands of individually-controllable queues.

MLE has made various contributions to the open source code basis of the Corundum FPGA NIC which all can be found at the Corundum project [hosted on GitHub](#).

MLE also provides binary releases of the Corundum NIC for the following FPGA boards:

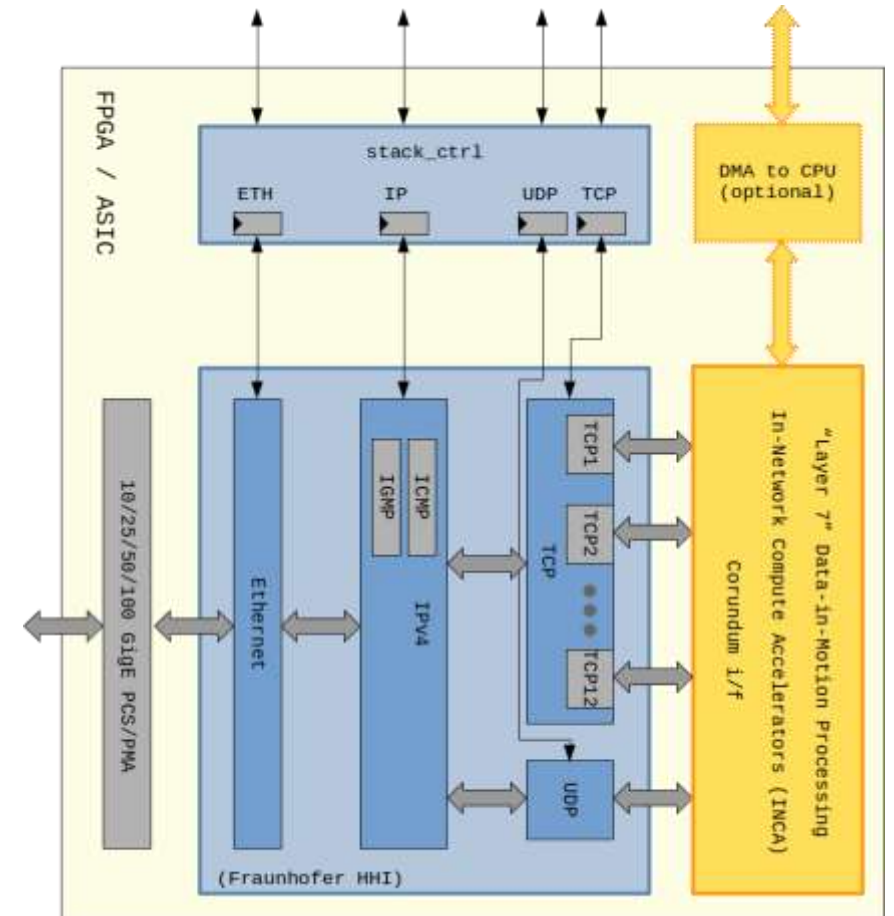
- Alveo U280 from AMD/Xilinx
- NPAC-Ketch from MLE (510GX400)
- Sidewinder-100 from Fidus Systems (ZU19EG)



NPAP - A TCP/IP Full Accelerator

- Interface to 1 / 2.5 / 5 / 10 / 25 / 40 / 50 / 100 Gigabit Ethernet
- Bidirectional datapath width 128 bit each
- Line rate >60 Gbps per individual TCP session in FPGA
- Line rate >100 Gbps per individual TCP session in ASIC
- Low round trip time NPAP-to-NPAP
700 nanoseconds for 100 Bytes RTT

<http://MLEcorp.com/NPAP>




NPAP Bandwidth & Latency

Ongoing Engineering Work to Optimize Bandwidth and Latency:

- Tuning for Intel Hyperflex
Was 40 Gbps (@ 311 MHz)
Now 67 Gbps (@525 MHz)
- Next tune for AMD/Xilinx Versal

Verification & Latency Analysis using Siemens Questa

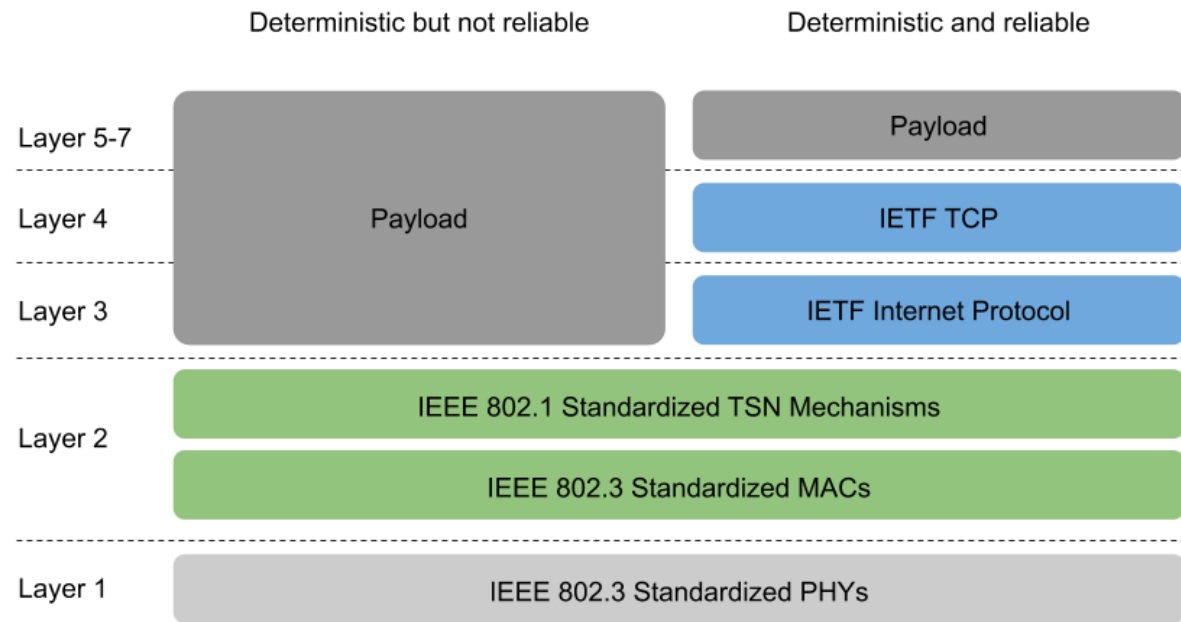


Continuous integration:

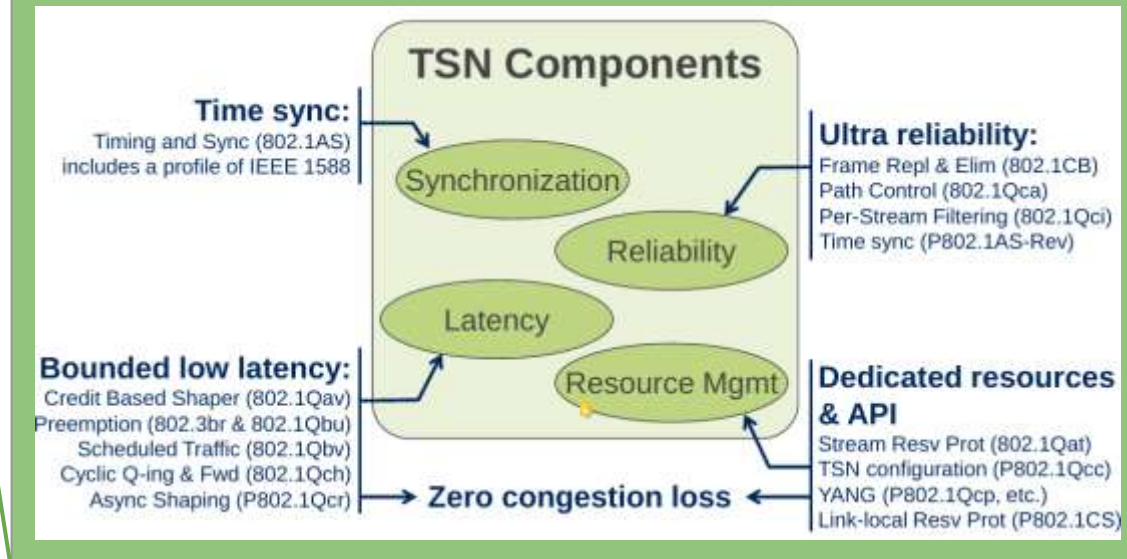
- 10G on in Xilinx ZCU102 with ZU9EG MPSoC
- 10G Xilinx ZC706 with Zynq-7045 SoC
- 10G on Intel Cyclone 10 GX Development Kit
- 10G on Intel Stratix 10 GX Development Kit
- 10G on Microsemi PolarFire MPF300-EVAL-KIT
- 25G on Xilinx ZCU111 with ZU28EG RFSoc
- 25G on Fidus Sidewinder 100 ZU19EG MPSoC
- 25G/100G on Xilinx Versal
- 25G/100G on Intel Agilex

Deterministic Networking With NPAP + TSN

Following ideas from IETF RFC8655 “DetNet”



Courtesy: Janos.Farkas@ericsson.com

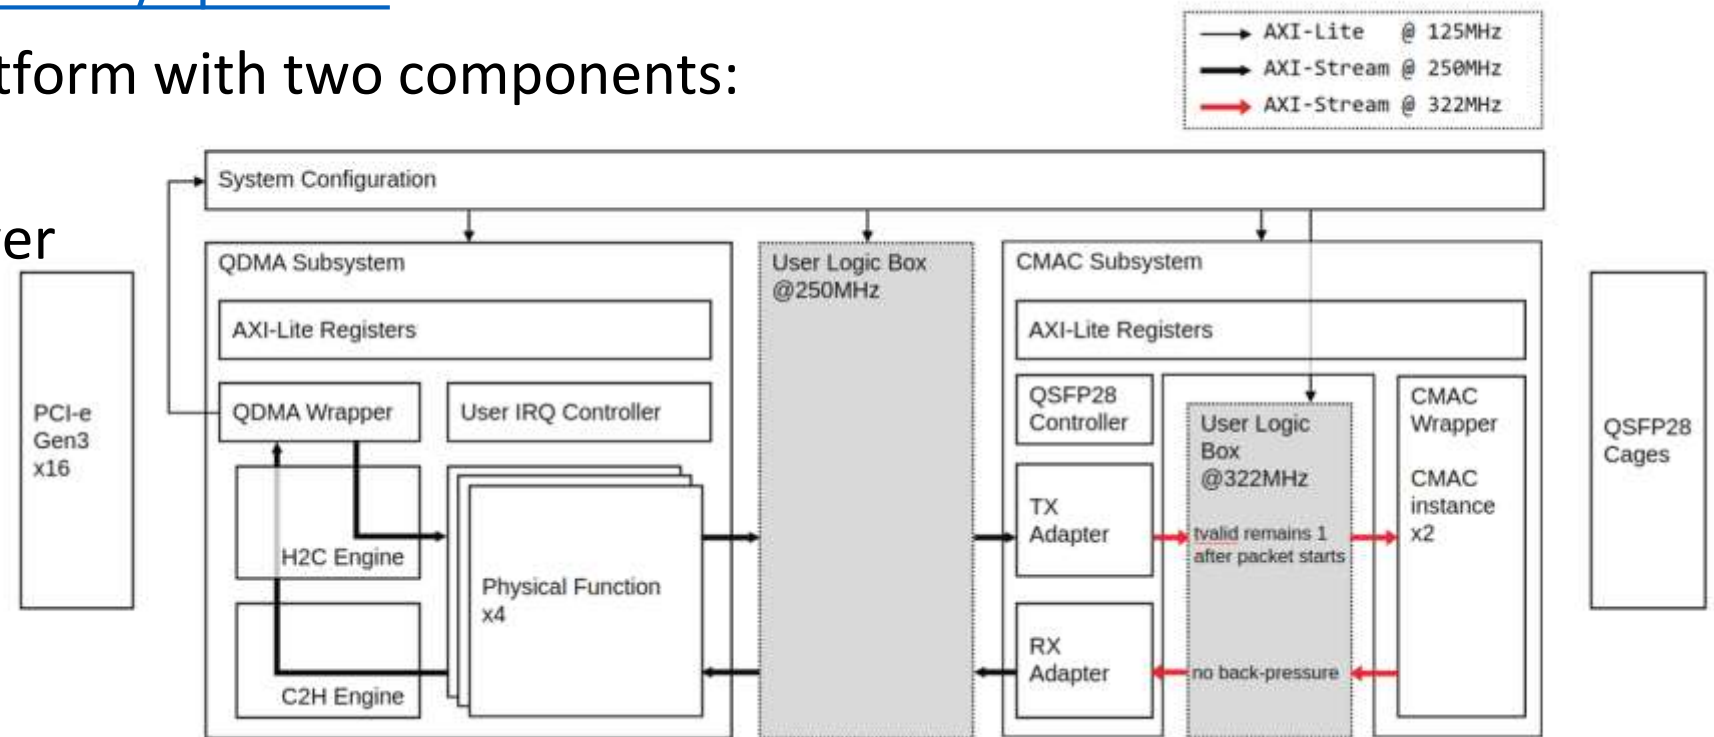


AMD OpenNIC

<https://github.com/Xilinx/open-nic>

FPGA-based NIC platform with two components:

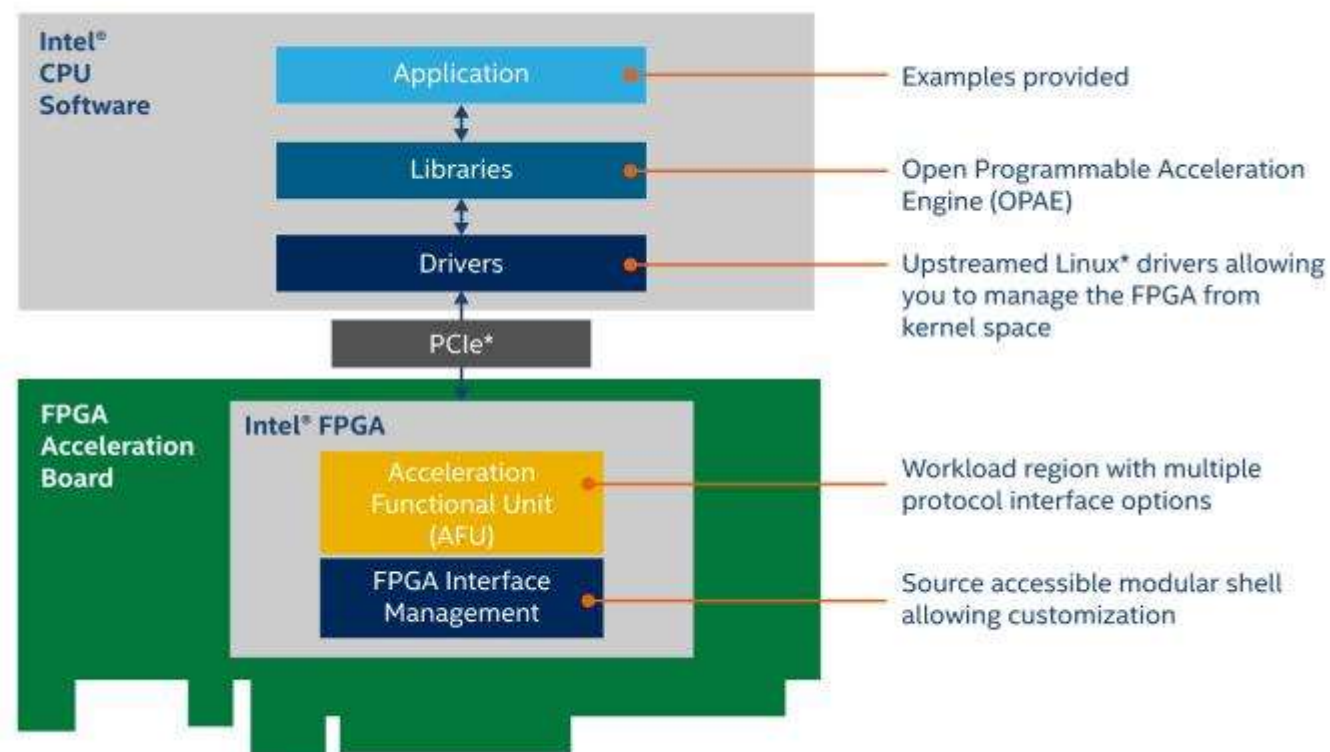
- FPGA shell
- Linux kernel driver



IOFS - Intel Open FPGA Stack

- Scalable
- Open-Source Access

<https://github.com/OPAE>




Putting Together Value-Optimized SmartNICs

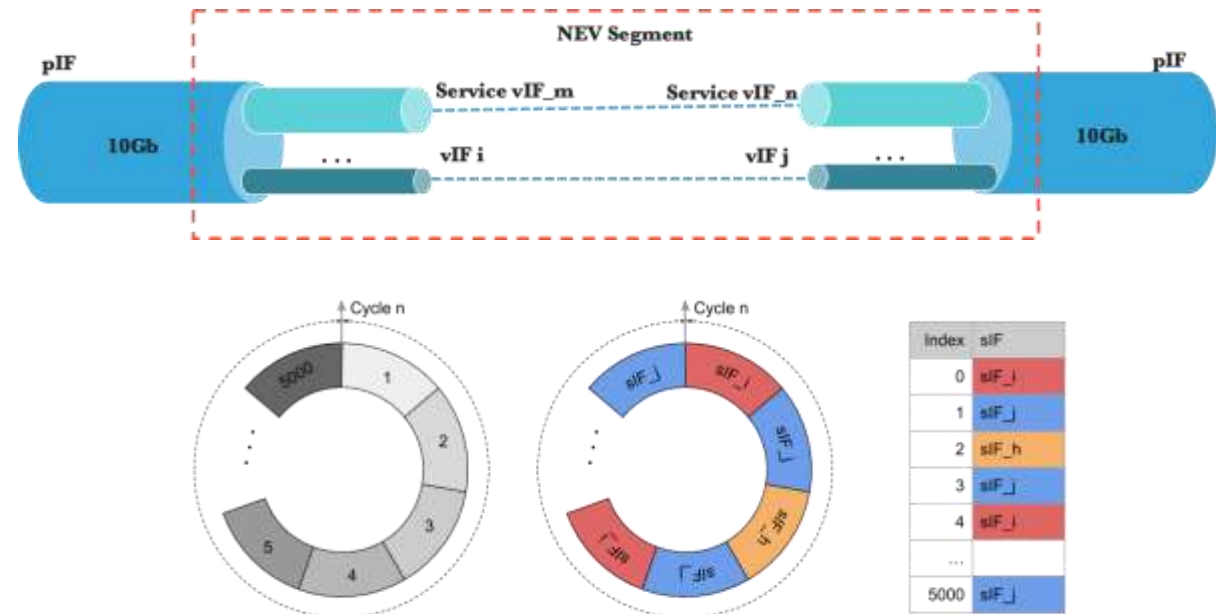
- Use IP-Cores and subsystems as “Lego blocks”
- Make extensive use of High-Level Synthesis
 - Achieves reasonable device independence
 - Vivado HLS
 - Intel Compiler for SystemC <https://github.com/intel/systemc-compiler>

Use case: FPGA SmartNIC for Algotblu, a NaaS provider

4x 10 GigE, PCIe Gen3 x8
Cost Optimized FPGA



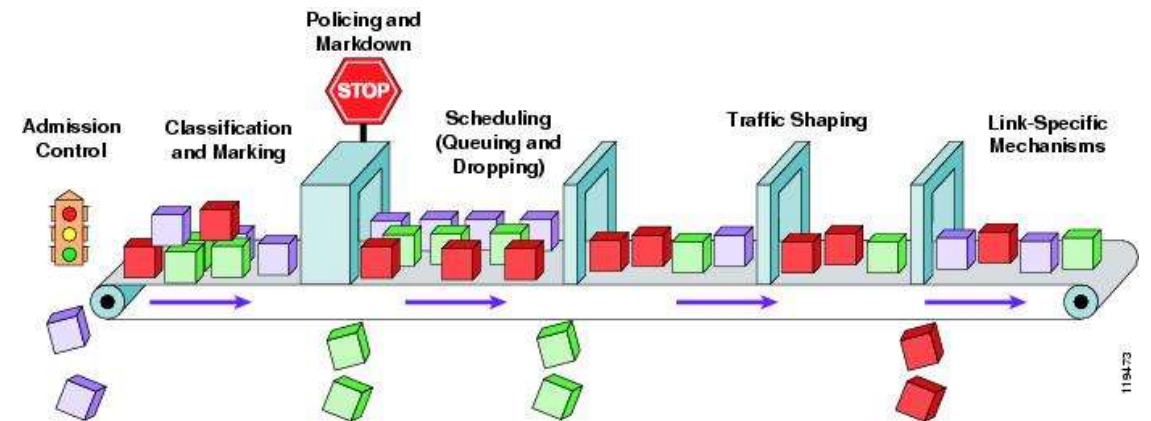
Network Element Virtualization (NEV) implementation targeting for TSN low-latency applications



TSN&NEV implementation on a FPGA



- TDMA based hardware(FPGA) control
- Much more simpler configurations
- Strict SLA guarantee
- TSN low-latency application support



- Software based control
- Complex configurations and CPU intensive
- No strict SLA guarantee

Use case: Algoblu application broadband service for cloud game



- End to end SLA guaranteed from player's home to cloud game platform
- Dedicated bandwidth(30M), ultra low latency(11ms), zero packet loss
- Great gaming experience far beyond the Internet



- Frame rate: 60
- Latency: 11ms
- Packet loss: 0
- Jitter: 0.1ms
- Distance: 650Km

Conclusion

Freedom of Choice:

- Wide range of off-the-shelf FPGA cards available, more on the horizon
- Industry-wide collaboration starts producing useful (open source) platforms

Freedom from Choice (a'la Alberto Sangiovanni-Vincentelli):

- To be useful FPGA platforms must be complete Hardware, “FPGA-ware” and Software
- Wedged between standards:
 - open source SDN software at top level
 - and IEEE Ethernet at bottom level

