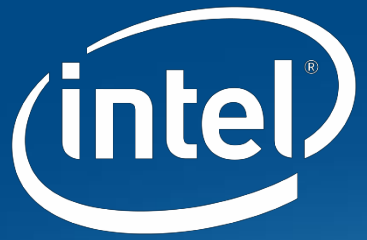

ACCELERATION STACK FOR INTEL® XEON® CPU WITH FPGAS

Marc Gaucheron

Intel Programmable Solutions Group

Multi-Function FPGA: Algorithm + Networking + Data Access Acceleration

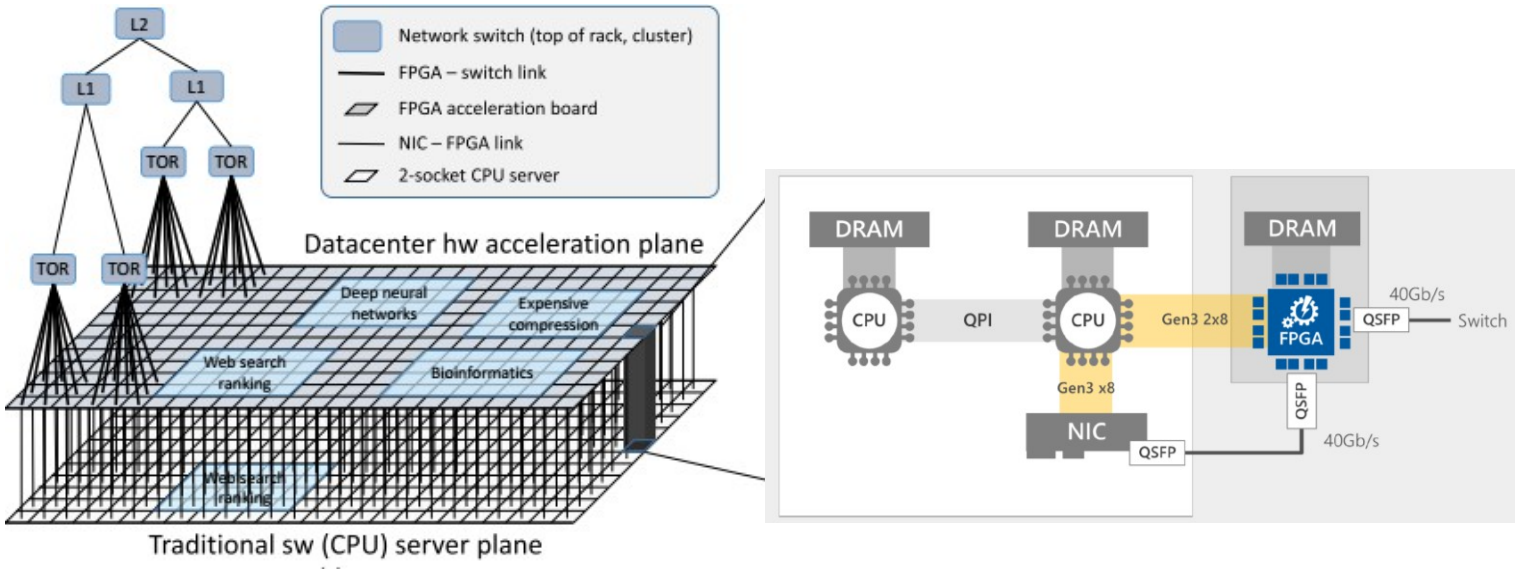

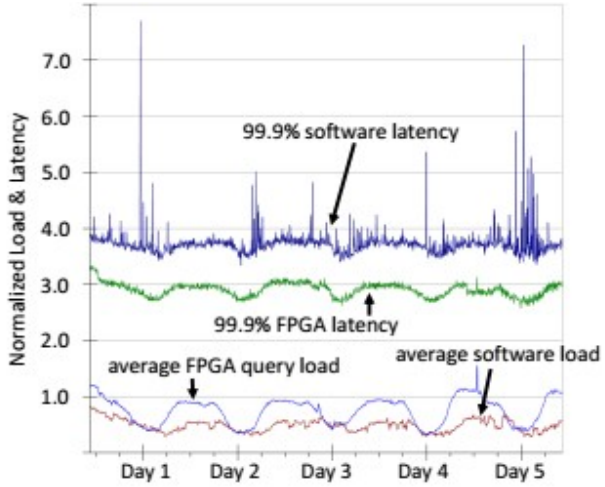

Fig. 7. Five day query throughput and latency of ranking service queries running in production, with and without FPGAs enabled.

Microsoft Scale Out FPGA Multi-Function Accelerator

- "Diversity of cloud workloads and … rapid … change" (weekly or monthly)
  - Search, SmartNIC, machine learning, encrypt, compress, and big data analytics
- Lower and predictable FPGA latency for ranking vs. software
- Excellent FPGA inferencing with reduced precision floating point (FP8)
  - Intel® Stratix® 10 10SG280 FPGA has 90 Teraops (MSFT Hot Chips presentation (Oct 2017))

Source: Microsoft

Intel FPGAs Offer Unique Value

High Throughput Deterministic Low Latency

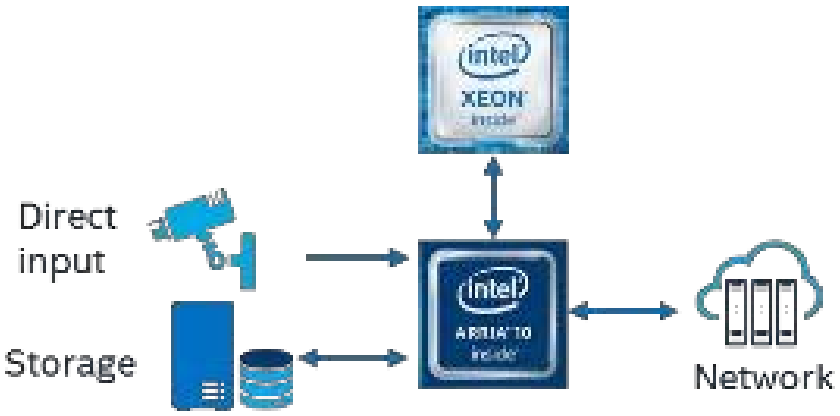

Excellent Power Efficiency

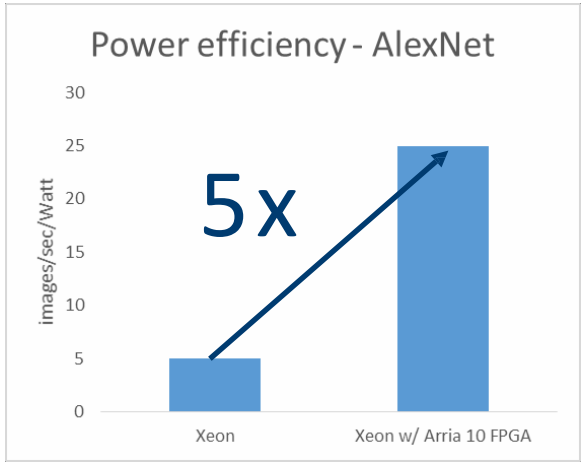

Future Proof

- ✓ Current and future neural network topology
- ✓ Arbitrary precision data types (FloatP32 => FixedP2, sparsity, weight sharing)
- ✓ Inline and offload parallel processing; IO expansion
- ✓ More than 25 year silicon lifespans

FPGAs Control the Datapath

Compute Acceleration/Offload

- Workload agnostic compute
- FPGAaaS
- Virtualization

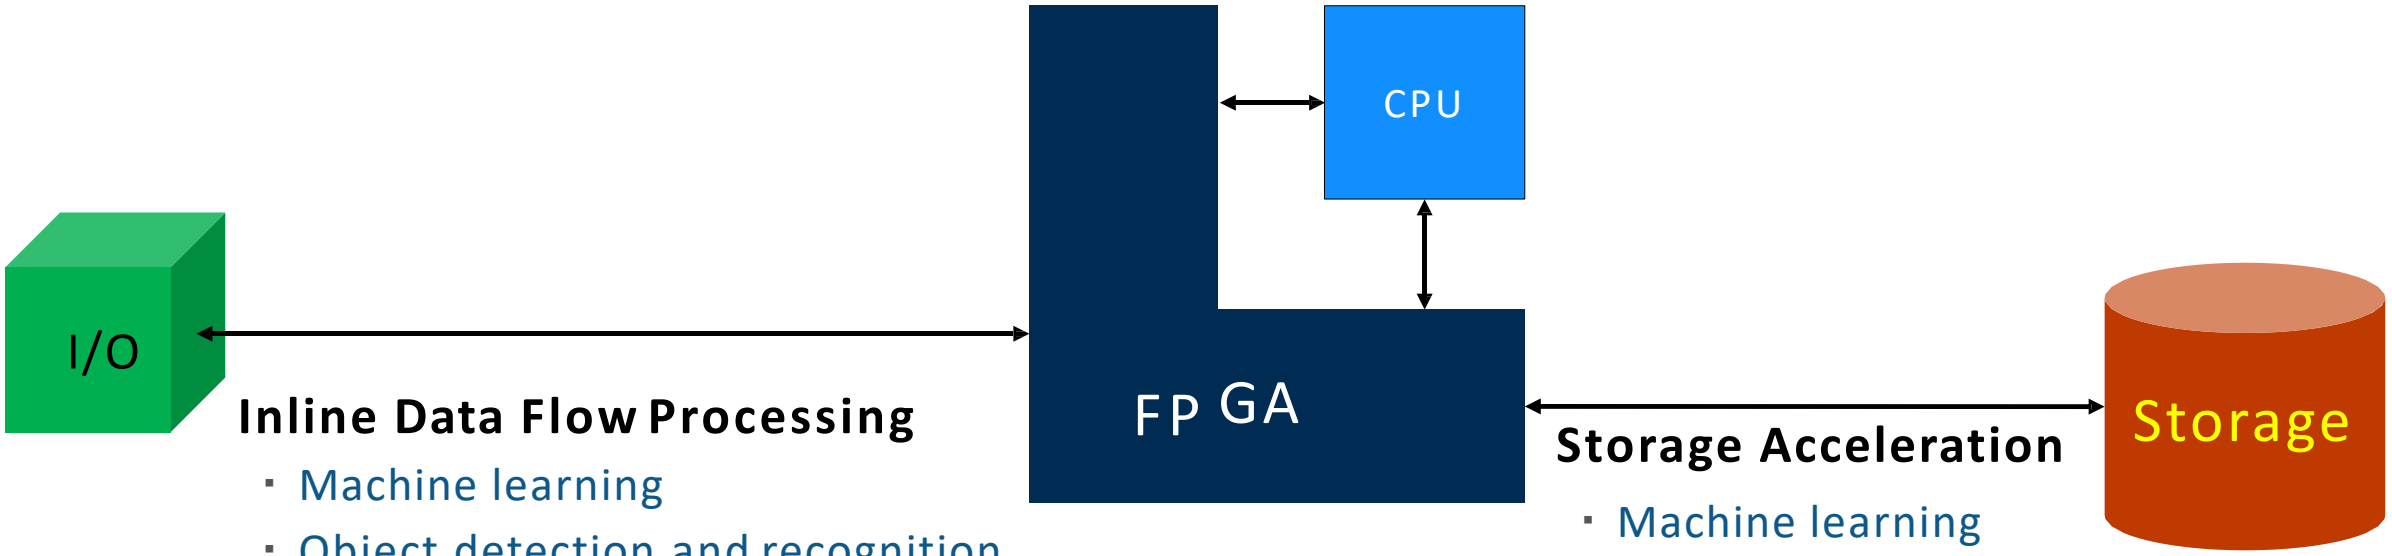

Inline Data Flow Processing

- Machine learning
- Object detection and recognition
- Advanced driver assistance system (ADAS)
- Gesture recognition
- Face detection

Storage Acceleration

- Machine learning
- Cryptography
- Compression
- Indexing

Integrated and Discrete Acceleration

Intel® Xeon® + Integrated FPGA Strengths

- Intel® Xeon® socket-compatible
- High Intel Xeon-to-FPGA bandwidth
- FPGA coherent with Intel Xeon memory
- Very low Intel Xeon-to-FPGA latency

Seamless Application Migration

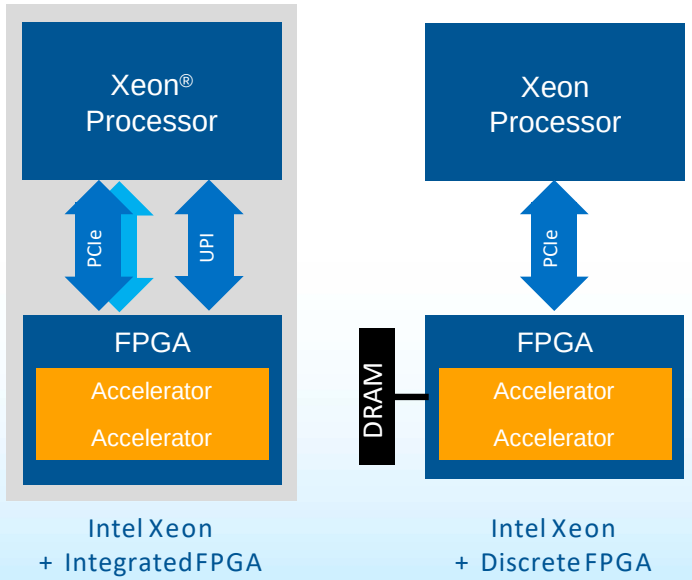

Intel Xeon + Discrete FPGA Strengths

- FPGA size choice
- Provides processor-pairing choice
- Compatible with standard motherboard
- Compatible with 1RU rack server
- Dedicated memory (Arria® 10 board)

Same virtualization framework, drivers, libraries, and development tools

Portfolio choice, consistent development experience and investment protection

Core Cache Interface: CCI-P

CCI-P is the hardware signaling interface between the Blue & Green Bitstreams

Virtual Channels	Physical links are presented to the AFU as channels. The AFU can select the virtual channel for each memory request.
VL0	Low latency virtual channel. (Mapped to UPI)
VH0	High latency virtual channel. (Mapped to PCIe0). Protocol efficiency is better for larger data payloads.
VH1	High latency virtual channel. (Mapped to PCIe1). Protocol efficiency is better for larger data payloads.
VA	Virtual Auto: FIU auto selects the link based on link utilization, request caching hint, and payload size. Latency: expect to see high variance BW: expect to see high steady state BW

DCP uses VH0 only

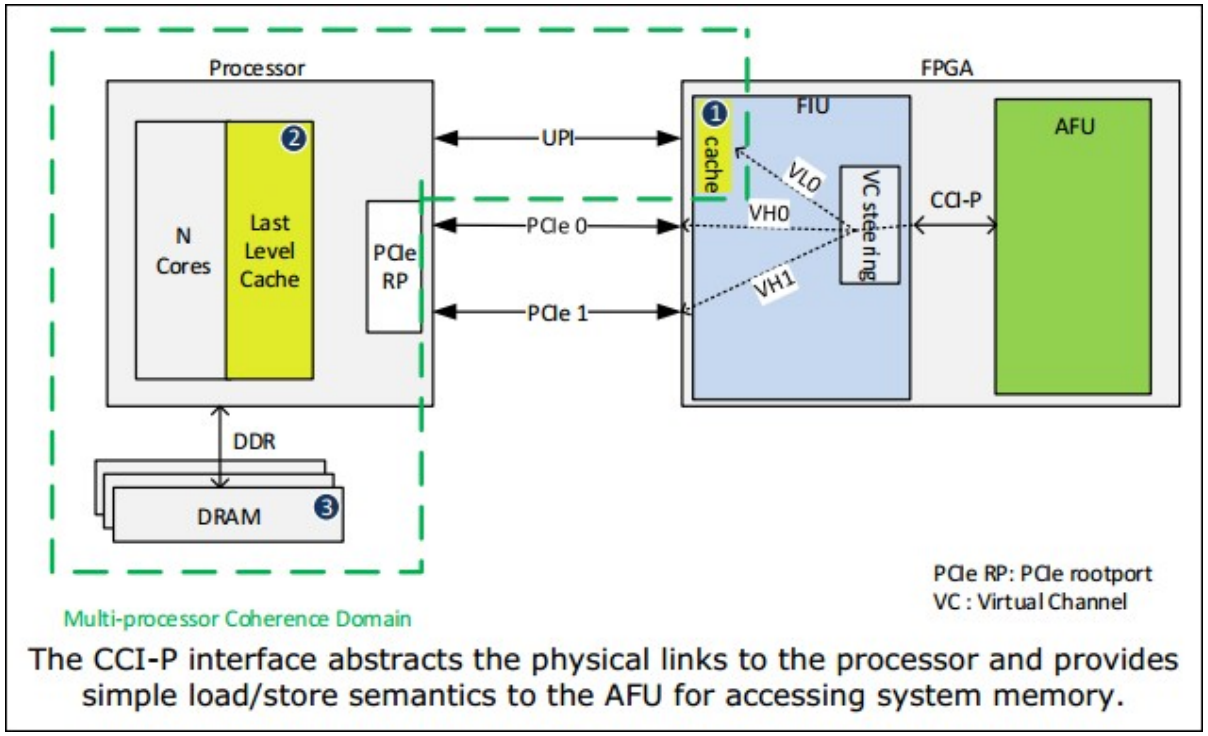

(AFU) Accelerator Function Unit or Green Bitstream
(FIU) FPGA Interface Unit or Blue Bitstream

Acceleration Stack for Intel® Xeon® CPU with FPGAs

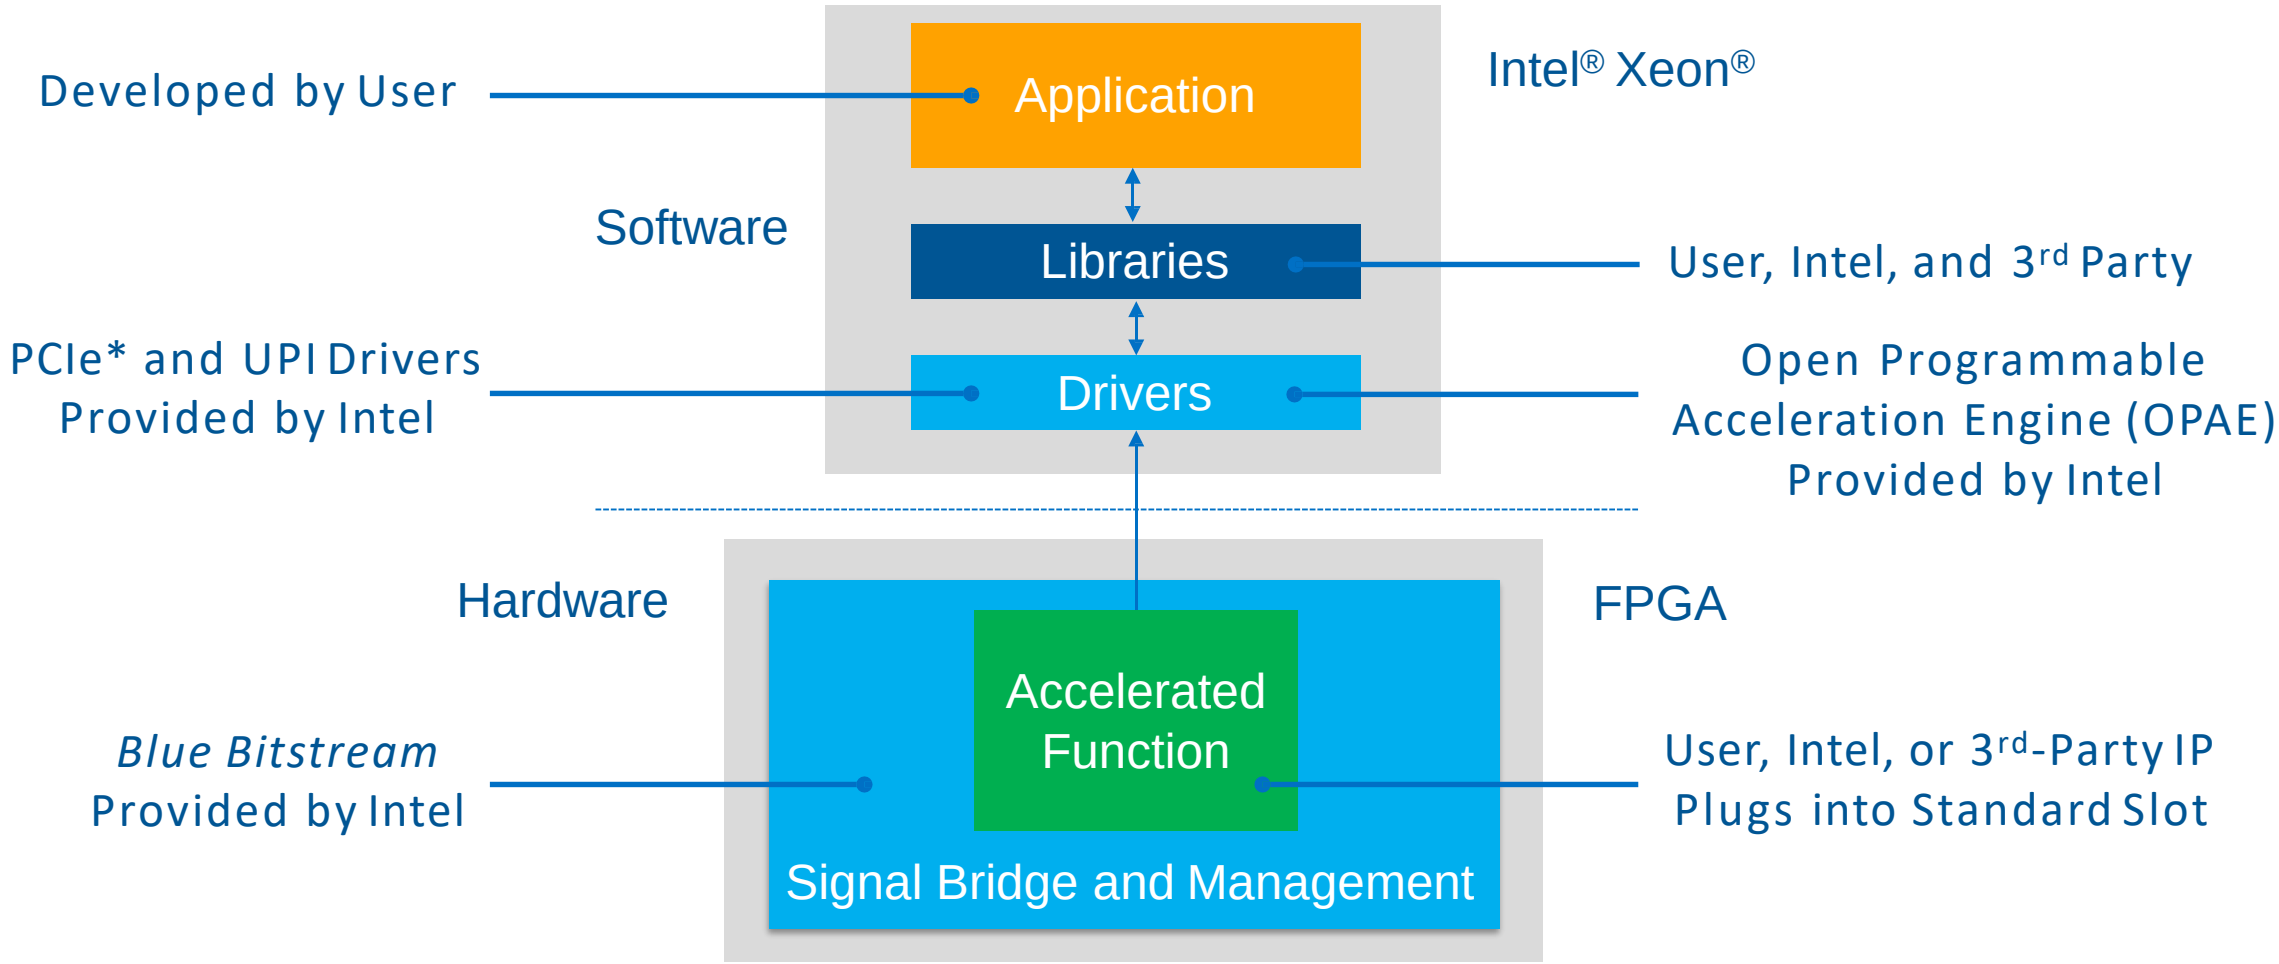

Simplifies the use of FPGAs in virtualized cloud environments

Interfacing with the Software Stack

Intel-developed
virtualization framework

Qualified with ecosystem partners

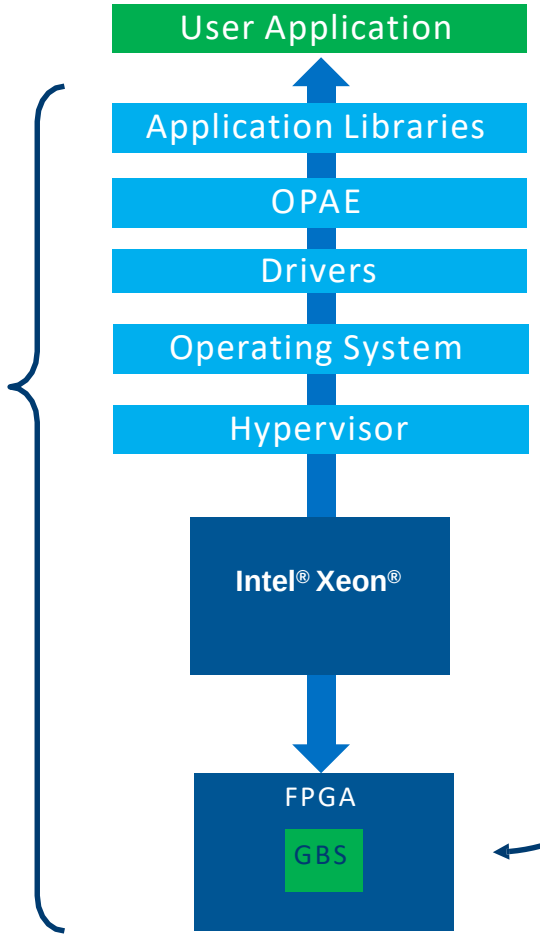

Application developer focuses on
accelerating core workloads

Develop own accelerator (GBS) or
leverage Intel or 3rd party accelerators

Open Programmable Acceleration Engine (OPAE)

Consistent API across product generations and platforms

Abstraction for hardware specific FPGA resource details

Designed for minimal software overhead and latency

- Lightweight user-space library (libfpga)

Open ecosystem for industry and developer community

- License: FPGA API (BSD), FPGA driver (GPLv2)

FPGA driver being upstreamed into Linux kernel

Supports both virtual machines and bare metal platforms

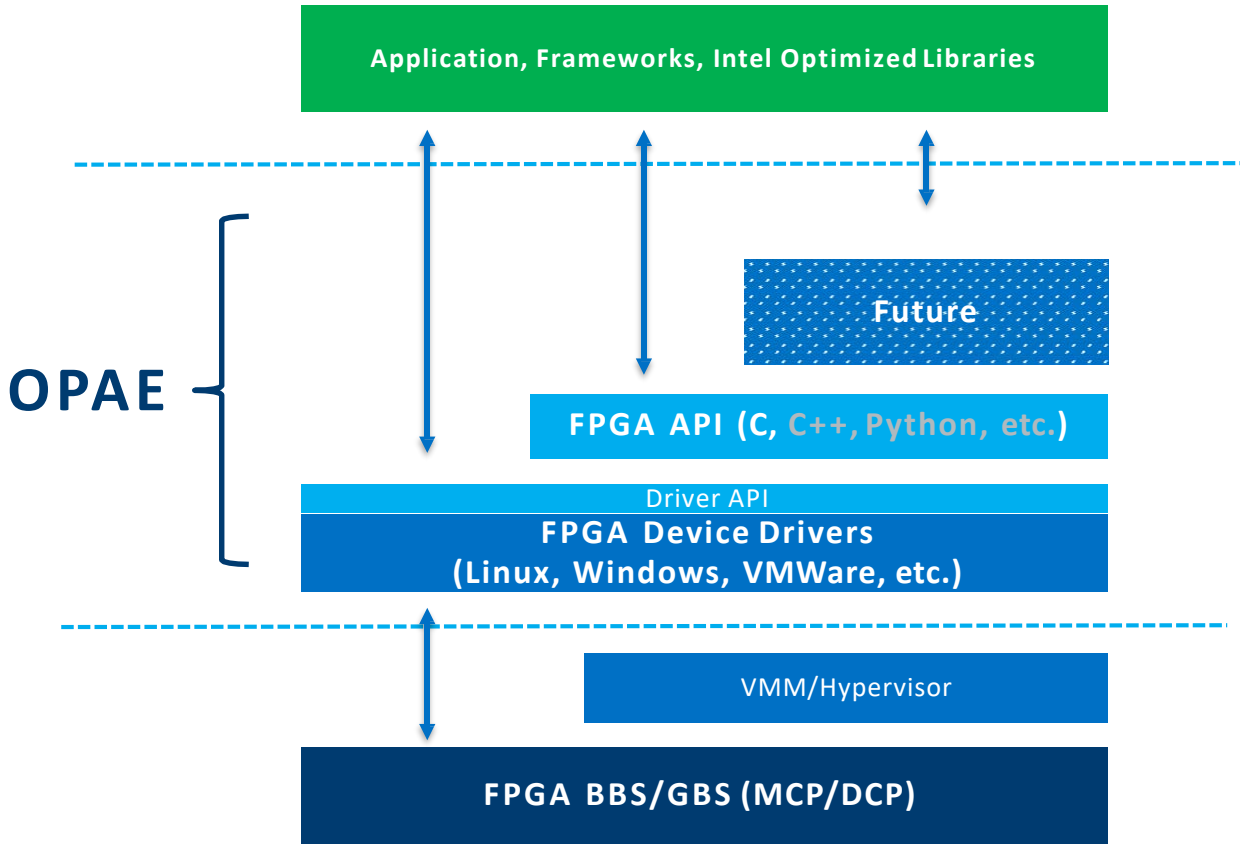

https://opae.github.io
https://github.com/OPAE/opae-sdk

Interfacing with the AFU

FPGA exposed as Standard SR-IOV device

AFU can be accessed by multiple VMs

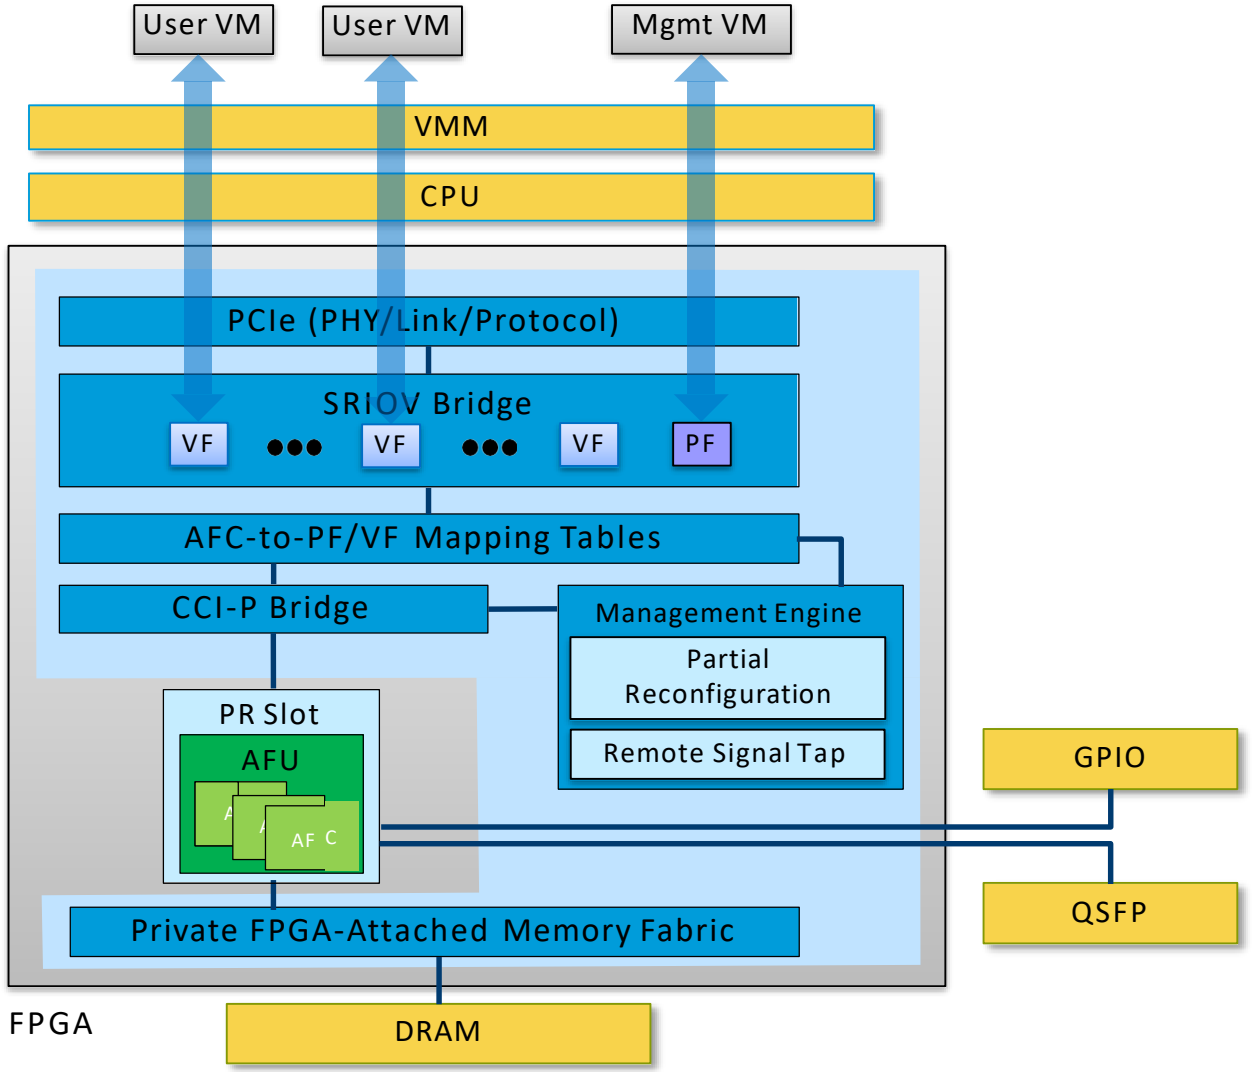

Intel® FPGA Driver Architecture

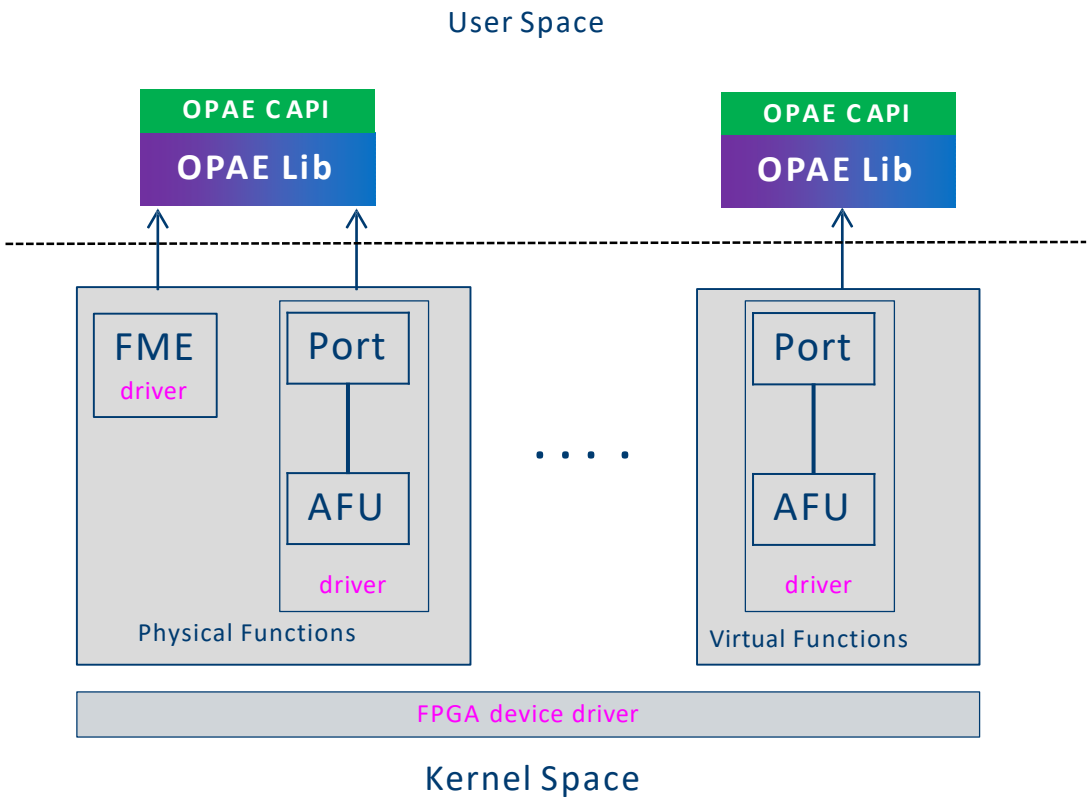

FME: FPGA Management Engine

- Static circuits for power/thermal management, reconfiguration, debugging, error reporting, performance counters, etc.

AFU: Accelerator Function Unit

- Reconfigurable circuits for application specific functions.
- Exposes a 256KB region as control registers.
- User process can share memory buffers with AFU.

Port:

- Interface between the static circuits and the reconfigurable region.
- Each port can attach an AFU. There may be multiple ports.
- A port can be assigned to a VM and expose the AFU.

Intel® Programmable Acceleration Card

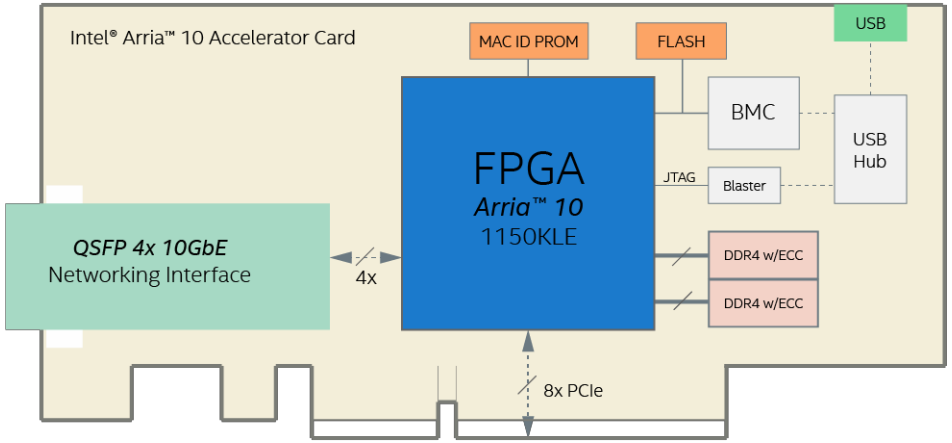

Arria® 10 GX FPGA [10AX115N2F40E2LG]

- High-perf, multi-gigabit SerDes transceivers up to 15 Gbps
- 1150K logic elements available (-2L speed grade)
- 53 Mb of embedded memory
- Passively cooled

On-board Memory

- 8 Gbytes DDR4 Memory Banks with ECC (2 banks)
- 1Gb Mbit (128 MB) Flash

Interfaces

- PCIe x8 Gen3 electrical, x16 mechanical
- USB 2.0 interface for debug and prog FPGA and Flash
- 1x QSFP28 with 4x 10GbE support

Board Management Controller

- Voltage, current, temperature monitoring
- Power sequencing and reset
- Field upgrades
- FPGA configuration and control

Brackets

- High and low profile brackets

AFU Development Approaches

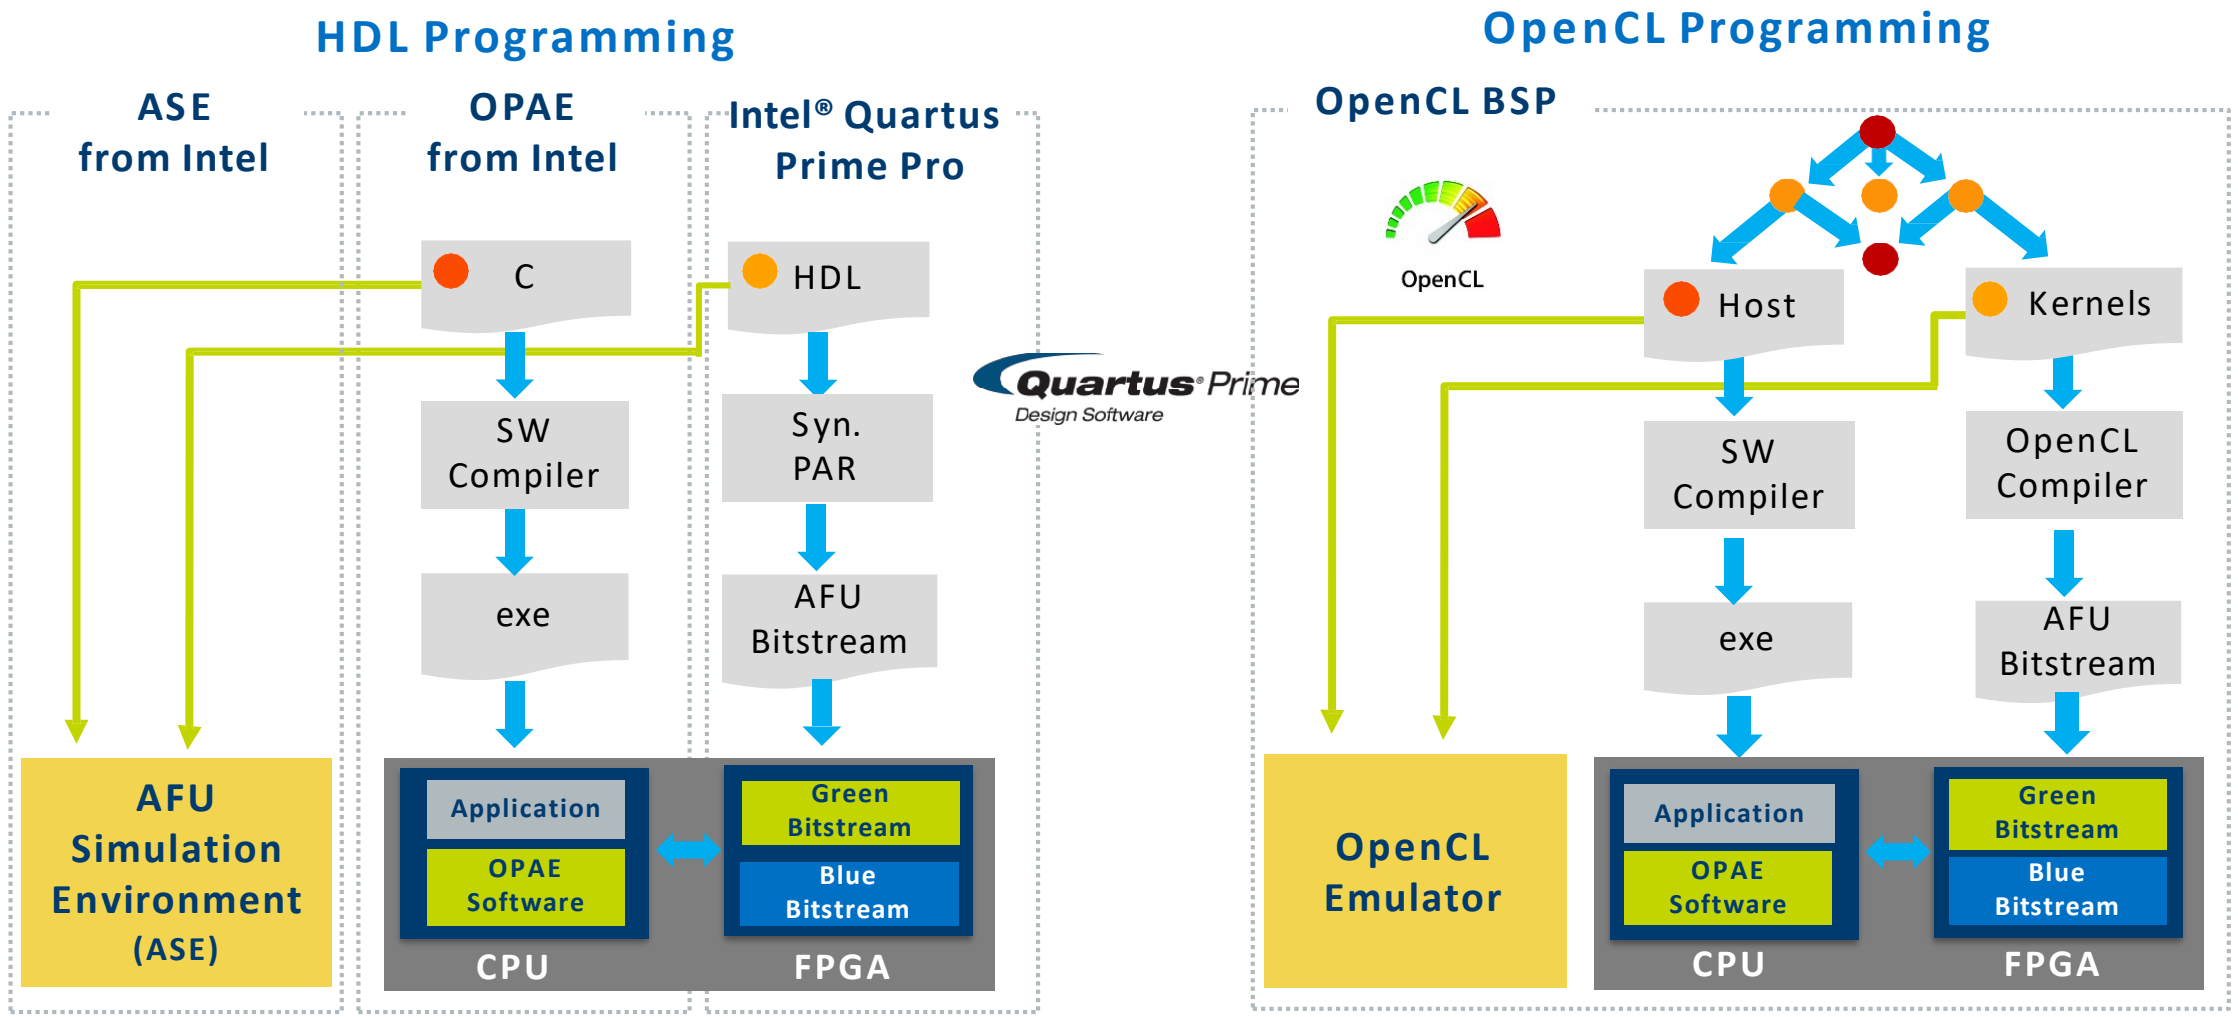

ASE – Transaction Level Simulator

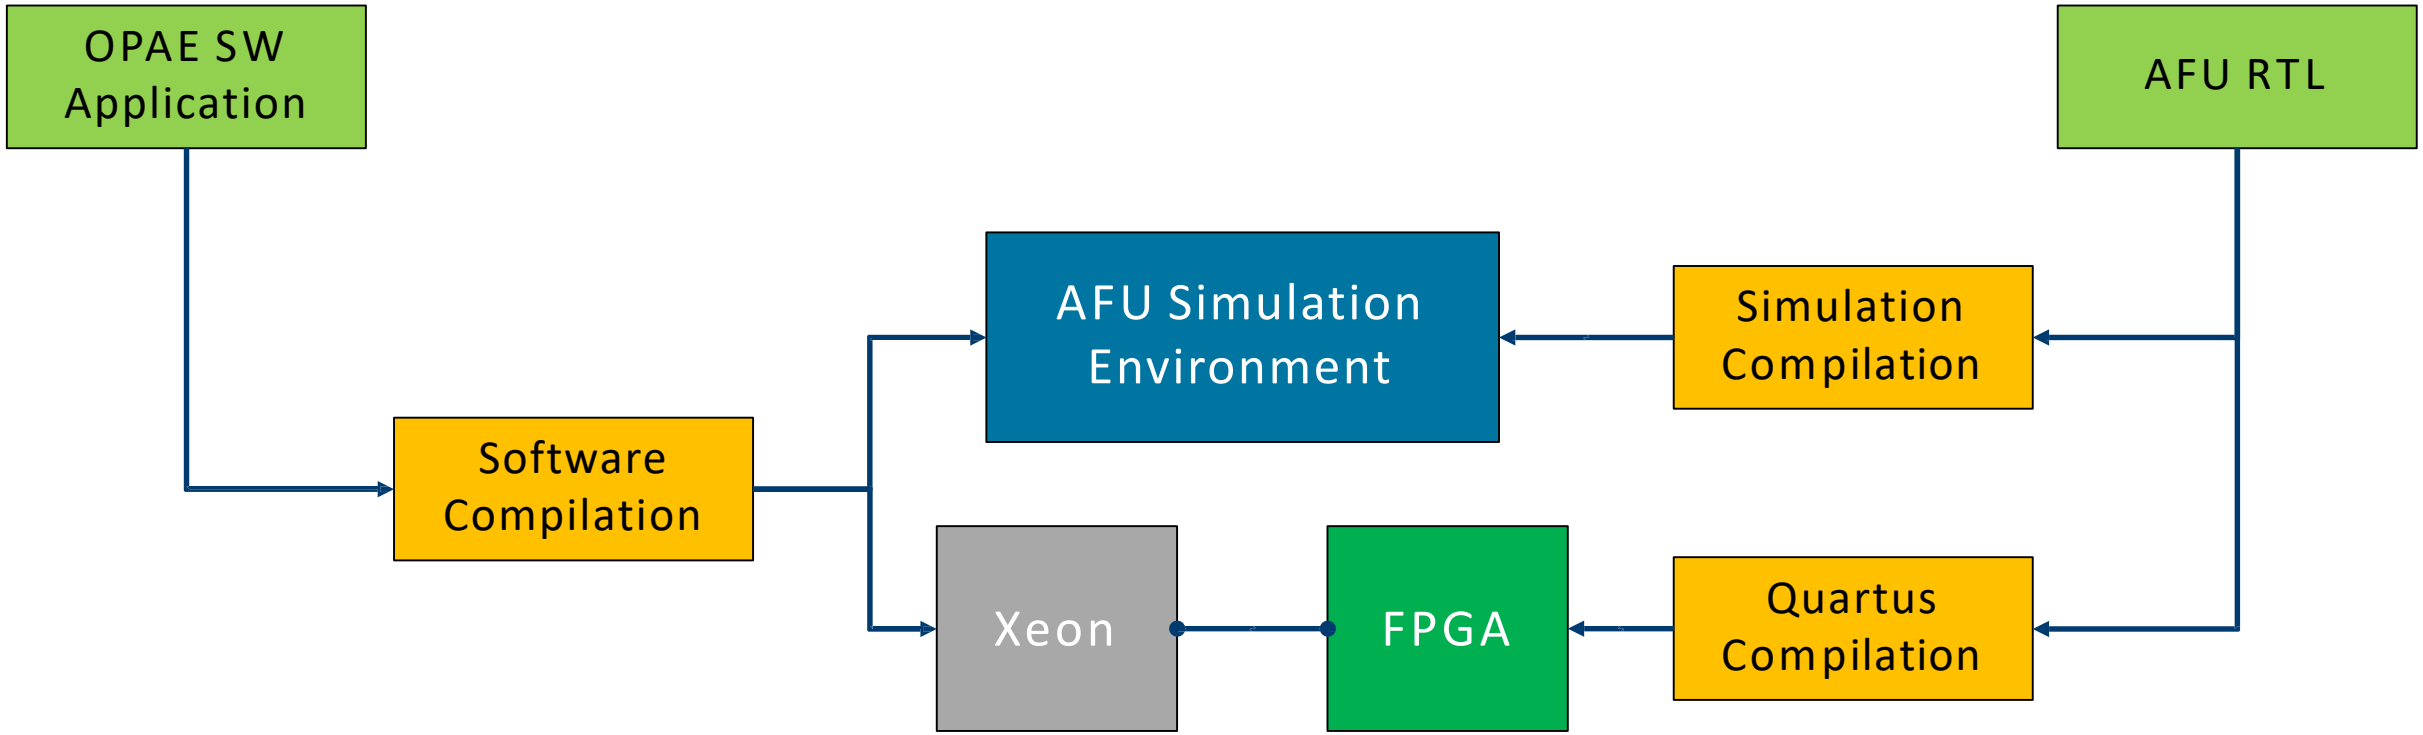

ASE Details

Transaction Simulator

- Models CCI transactions and behaviour
- Models realistic memory map to the memory
- Identifies CCI protocol correctness
- Converts SW API to CCI Transactions and vice-versa
- No administrator privileges needed to run ASE (completely User-level code)

ASE limitations

- Not cycle accurate and cannot be used for performance measurement
- Not a cache simulator
- Does not guarantee synthesizability or timing closure
- Not a latency model

SignalTap and Other Remote Debug Features

Remote debug operates over PCIe

Can connect to the server remotely

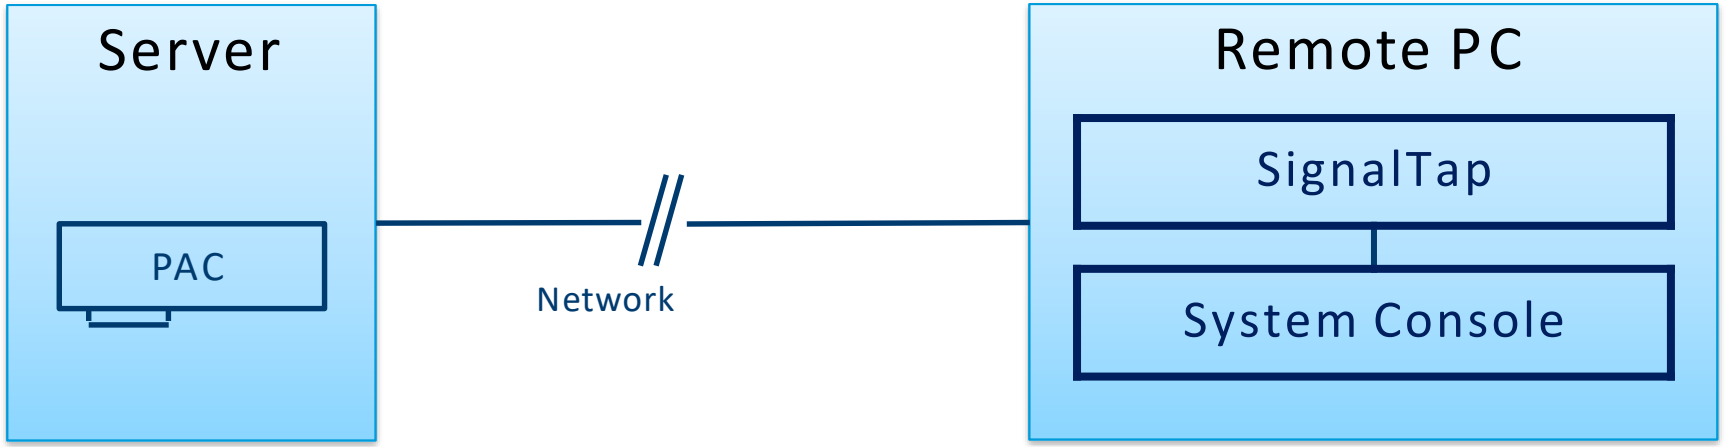

Code Sample

Set Properties

Enumerate

Open

Map MMIO

```
/* Look for AFC with MY_AFC_ID */
res = fpgaGetProperties(NULL, &filter);
ON_ERR_GOTO(res, out_exit, "creating properties object");

res = fpgaPropertiesSetObjectType(filter, FPGA_ACCELERATOR);
ON_ERR_GOTO(res, out_destroy_prop, "setting object type");

res = fpgaPropertiesSetGUID(filter, guid);
ON_ERR_GOTO(res, out_destroy_prop, "setting GUID");
```

```
res = fpgaEnumerate(&filter, 1, &afc_token, 1, &num_matches);
ON_ERR_GOTO(res, out_destroy_prop, "enumerating AFCs");

if (num_matches < 1) {
    fprintf(stderr, "AFC not found.\n");
    res = fpgaDestroyProperties(&filter);
    return FPGA_INVALID_PARAM;
}
```

```
/* Open AFC and map MMIO */
res = fpgaOpen(afc_token, &afc_handle, 0);
ON_ERR_GOTO(res, out_destroy_tok, "opening AFC");

res = fpgaMapMMIO(afc_handle, 0, NULL);
ON_ERR_GOTO(res, out_close, "mapping MMIO space");
```

Key Documentation

dcp-a10-quick-start.pdf

- Installation, environment and basic run through

ase-dcp-quick-start.pdf

- Intro to AFU Simulation Environment

afu-developers-guide.pdf

- Development guidelines including GBS generation and Signal Tap

cci-p-specification.pdf

- Hardware interface specification

opae-programming-guide.pdf

- Software programming

Additional Resources

Intel® FPGA Acceleration Hub

- www.altera.com/solutions/acceleration-hub/solutions.html

OPAE Video

- https://www.brighttalk.com/webcast/10773/275799?utm_source=Intel+-+Data+Center+Group&utm_medium=brighttalk&utm_campaign=275799

Customer training

- Introduction to the Acceleration Stack for Intel® Xeon® CPU with FPGA
- RTL Development and Acceleration with the Acceleration Stack for Intel® Xeon® CPU with FPGA

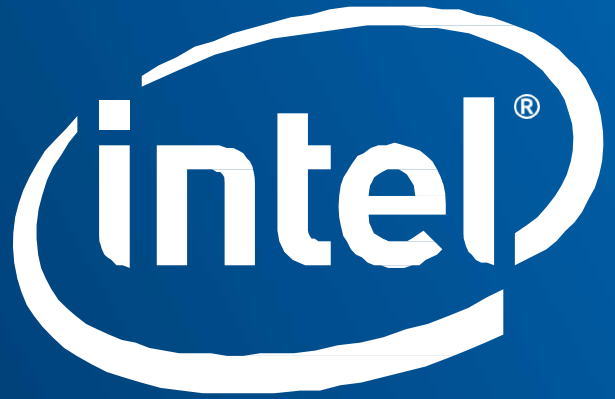
intel