FPGA Acceleration and Virtualization Technology in DPDK

ROSEN XU
TIANFEI ZHANG

Agenda

- FPGA in Networking Acceleration
- Partial Reconfiguration (PR)
- FPGA Acceleration on DPDK
  - DPDK High Level Design
  - Scan and Probe Work Flow
  - Intel FPGA Acceleration Environment
  - Intel FPGA Acceleration Stack OPAE Intro
  - Intel FPGA Acceleration on DPDK
  - Port Representor and Virtualization Scenario
- Status & WIP
- Acknowledgement

FPGA in Networking Acceleration

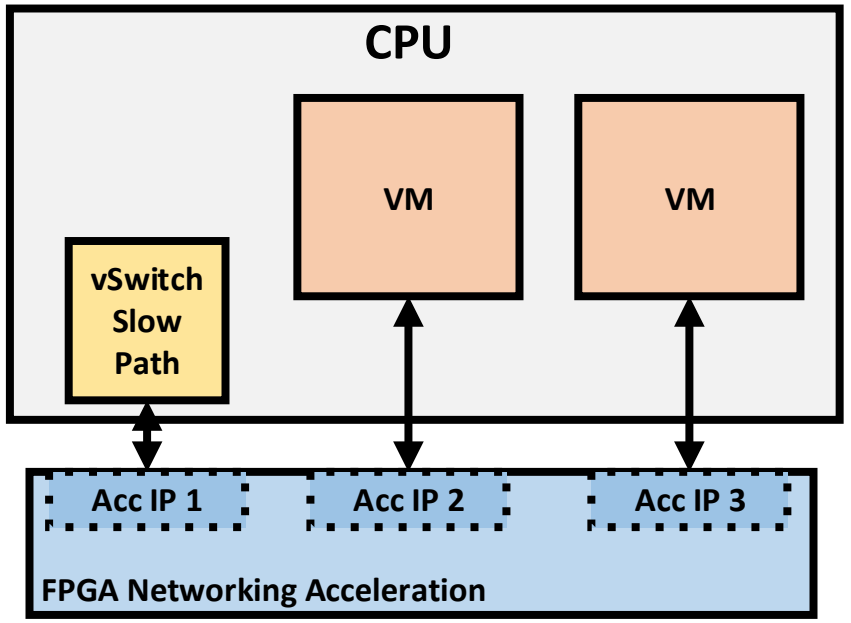

- **Opportunities**
  - Enhancing Performance: Provide NIC ASIC liked performance
  - Changing dynamically: Flexible enough for adding new feature by replace Bit Stream
- **Problems**
  - Longer design cycle than software: Compilation, Analysis & Synthesis, Fitter(Place & Router), Assembler, Timing
  - Update Bit Stream affecting business: PCIe rescan and driver re-probe
  - How to share One FPGA resource with many users

Partial Reconfiguration (PR)

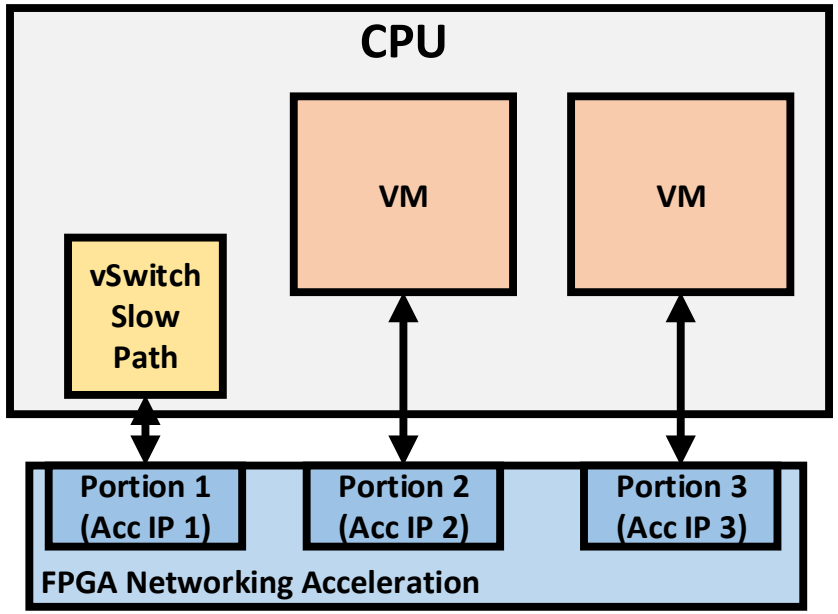

- **With Partial Reconfigure(PR) a portion of the FPGA dynamically, FPGA not only provides one kinds of accelerator but also provides many types of accelerators at the same time**
  - All FPGA vendor support PR function
- **How DPDK fully support FPGA?**
  - Which type of DPDK Device can provide FPGA PR?
  - How can we bind DPDK Driver to FPGA Partial Bit Stream?

DPDK High Level Design

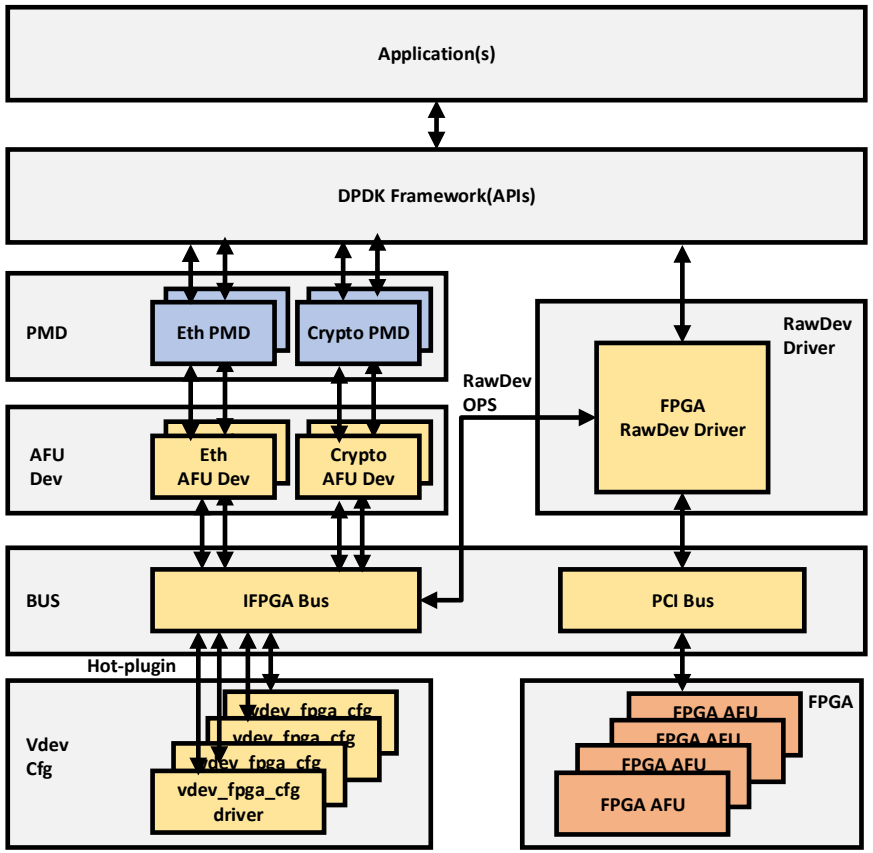

- FPGA is divided to many portions, each portion has its own Partial Bit Stream
- Vdev Cfg takes configuration for each AFU
- IFPGA Bus is a new bus for AFU devices scan and drivers probe
- New added AFU Device for each portion's Acceleration function
- All AFUs' PMD are based on AFU Acceleration

Note:

- AFU(Accelerated Function Unit): Partial-Bitstream for a portion of the FPGA

Scan and Probe Work Flow

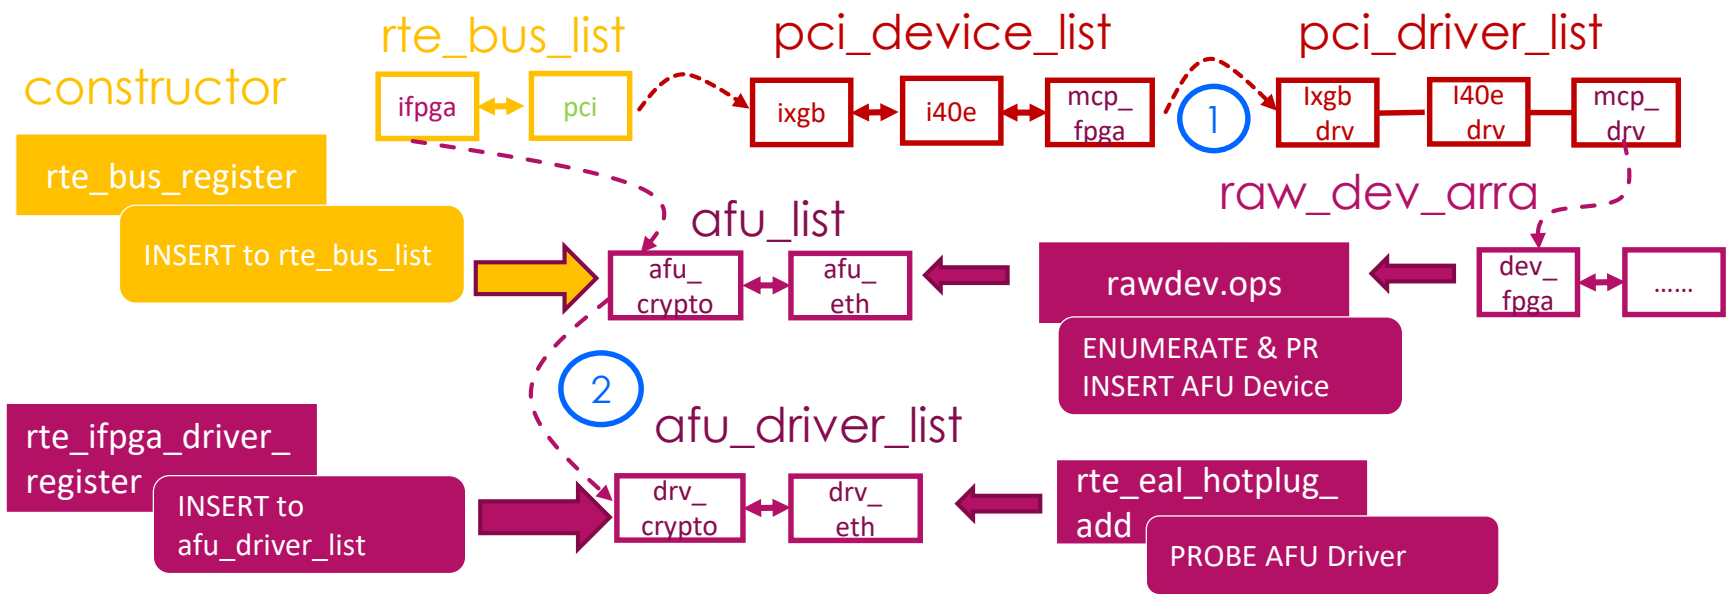

- **Rawdev probed as PCI Driver takes FPGA Configuration(Download/PR)**
- **2 scans: FPGA PCI Device Scan(1st Scan) and AFU Scan(2nd Scan)**
- **OPAE Provides Common lib and API for low level FPGA management & accelerator access**

Intel FPGA Acceleration Environment

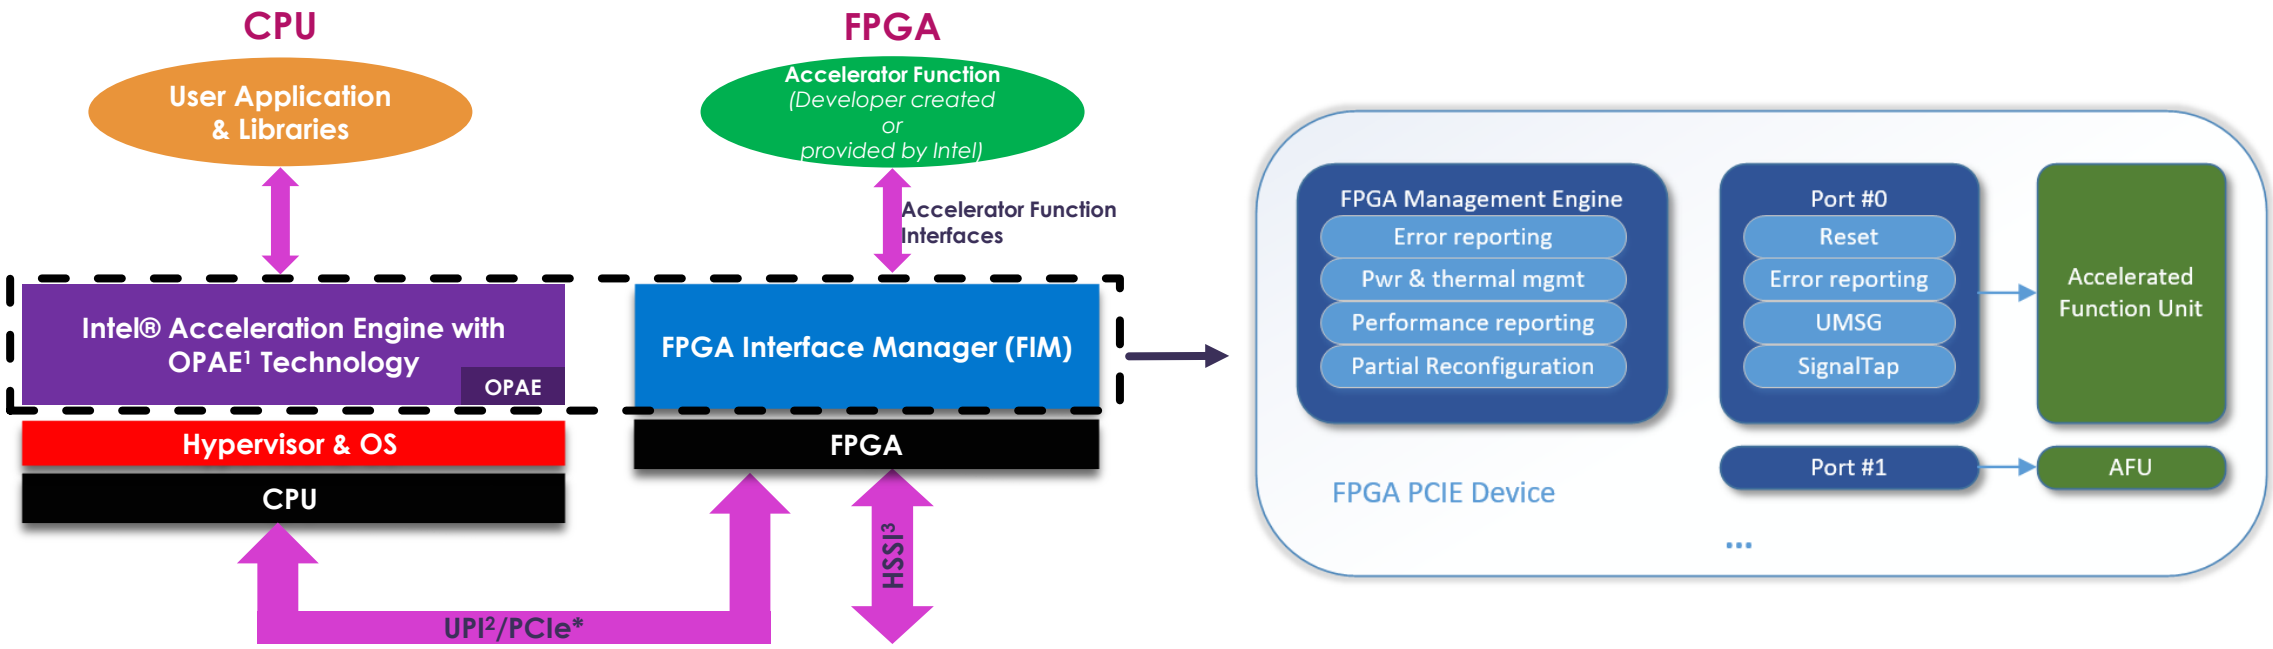

- **FPGA Management Engine(FME)**
  - Provides: power and thermal management, error reporting, partial reconfiguration, performance reporting, and other infrastructure functions.
  - Each FPGA has one FME, accessible through the physical function.
- **Accelerated Function Unit(AFU)**
  - Implements: one Acceleration, can be partial reconfiguration.
  - Each FPGA can support Multiple AFUs.

Intel FPGA SW Stack OPAE Intro

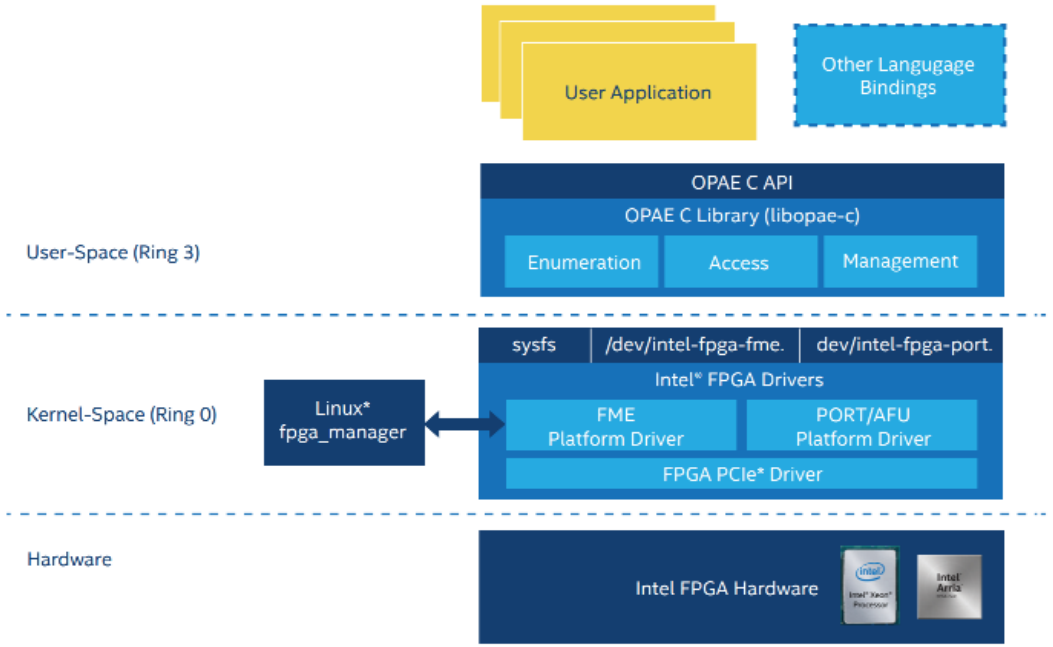

- **Optimized and simplified hardware and software APIs provided by Intel**
- **Consistent cross-platform API**
- **Minimal software overhead and latency**
- **Supports virtual machines and bare metal platforms**
- **Open source code licensing and developer community**
  - Intel FPGA drivers being upstreaming to Linux kernel
  - Intel FPGA user space drivers have merged into DPDK

https://01.org/sites/default/files/downloads/opae/open-programmable-acceleration-engine-paper.pdf

Intel FPGA Acceleration on DPDK

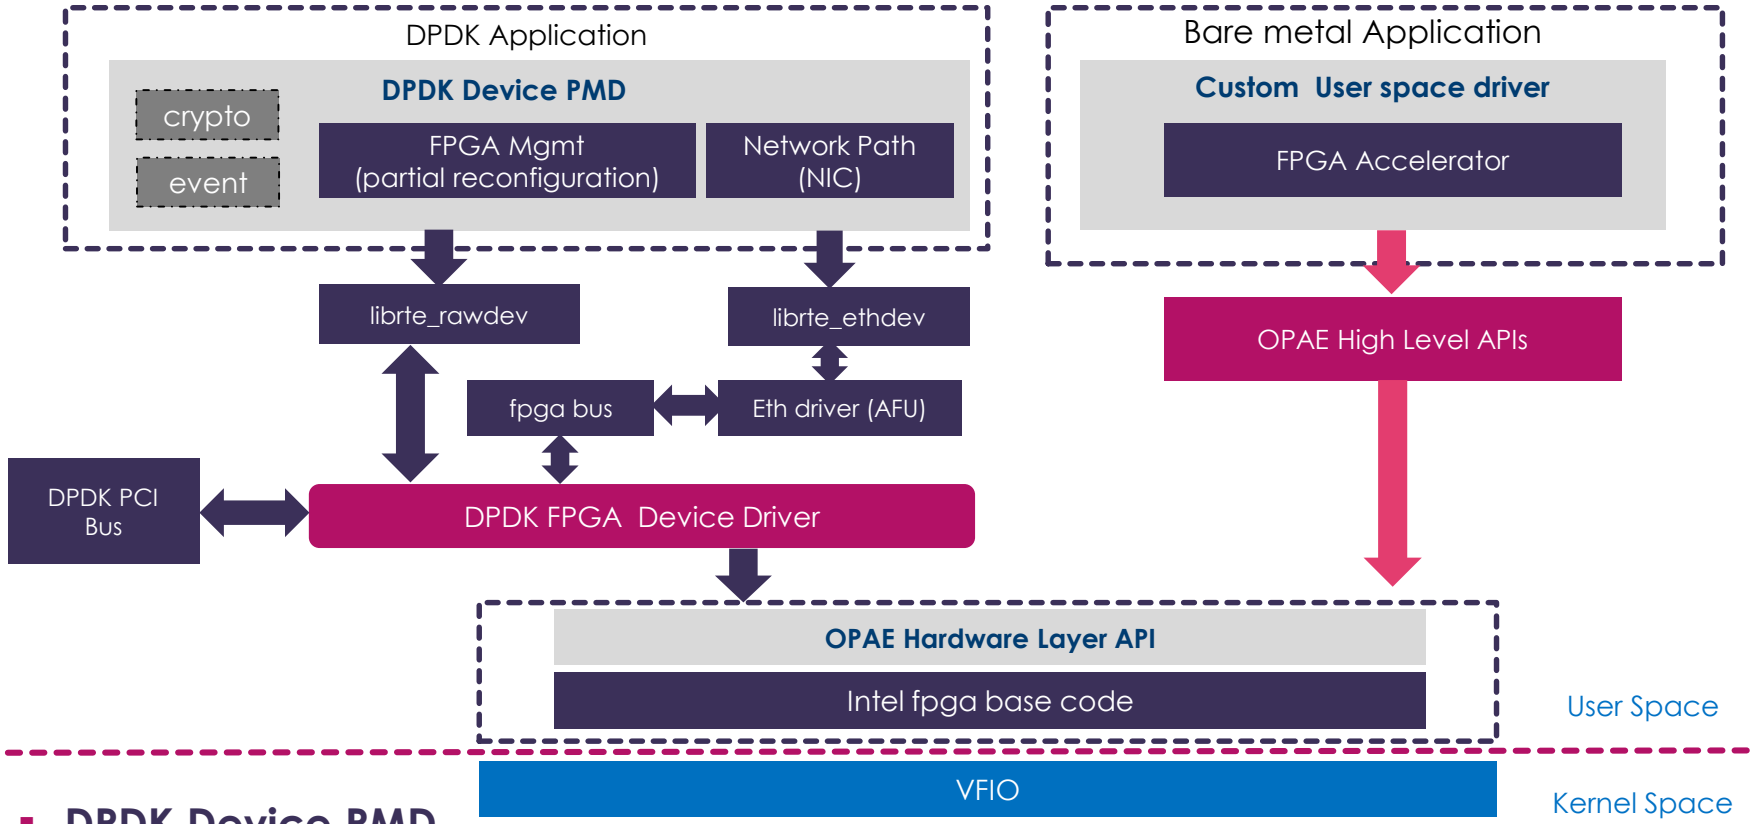

- **DPDK Device PMD**
  - For defined DPDK Device(Ethdev/Cryptodev/Eventdev)
- **Non-DPDK User-Space Driver**
  - For Customized Device, Transparent DPDK
- **Rawdev Support PR**
  - Rawdev is submited in 18.02

Note:

- **DPDK and NON-DPDK mode will not run at the same time**

Port Representor and Virtualization Scenario

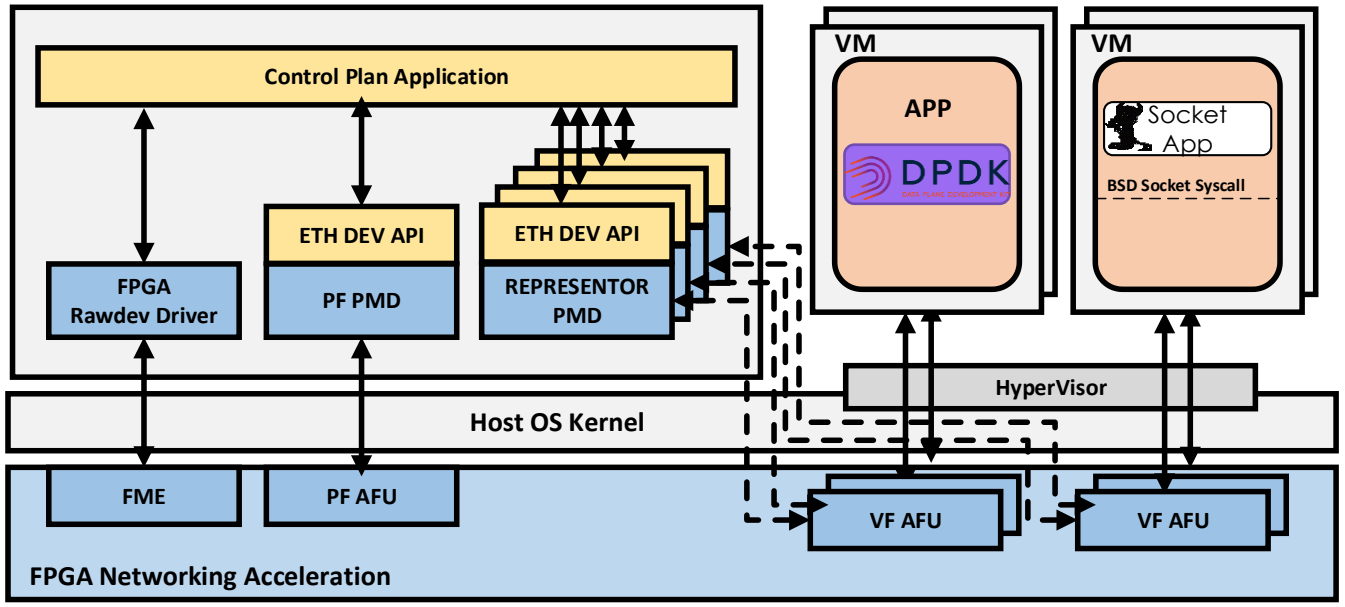

- **Port Representor**
  - Each VF port binging to one representor port
  - Control Plan Application take perception of VF port by its representor port
- **Virtualization**
  - FPGA can support both DPDK APP and Socket APP in VM
  - FPGA Rawdev Driver take FPGA configuration

Status & Working in Progress

- 2017'Q3 OPAE 0.9 Release [**DONE**]
  - Fully support Intel FPGA Acceleration Environment
  - Support FIM 6.3.0
- 2018'Q1 OPAE 0.13 Release, [**DONE**]
  - Support FIM 6.4.0
- 2018'Q2 OPAE 1.0 Release [**DONE**]
- 2018'Q1 DPDK with OPAE User Space Driver PoC [**DONE**]
- 2018'Q1 IFPGA Bus RFC patch [**DONE**]
- 2018'Q2 IFPGA Bus patch set upstream to DPDK 18.05 [**DONE**]

DPDK supports FPGA Acceleration is ready.
Welcome on board!

Acknowledgement

- Song Liu
- Zhang Zhang
- Hao Wu
- Yilun Xu
- Heqing Zhu
- Cunming Liang
- Helin Zhang
- Zhihong Wang
- Yanglong Wu
- Jingjing Wu
- Qi Zhang
- Yuwei Zhang
- Declan Doherty
- Bruce Richardson
- Ferruh Yigit
- Roy Fan Zhang

Thanks !