OVS Offload with Intel® FM10000

DPDK Summit, Aug. ‘16

Eyal Cohen

Silicom Ltd.
Connectivity Solutions

OvS
Open vSwitch

Agenda

- Intel® FM10000 NIC Is Unique
- A Walk on the Data Path
- DPDK OVS Benefits
- Benchmark Results
- Roadmap

Agenda

- Intel® FM10000 NIC Is Unique
- A Walk on the Data Path
- DPDK OVS Benefits
- Benchmark Results
- Roadmap

Intel® FM10000 NIC Is Unique

- Connectivity – 1G, 10G / 40G, 25G / 100G
- Switch – TCAM, FlexPipe™ (tunneling, VXLAN, filtering)
- Integrated PCIe MAC – SR-IOV
- All in one die

Agenda

- Intel® FM10000 NIC Is Unique
- A Walk on the Data Path
- DPDK OVS Benefits
- Benchmark Results
- Roadmap

A Walk on the Data Path

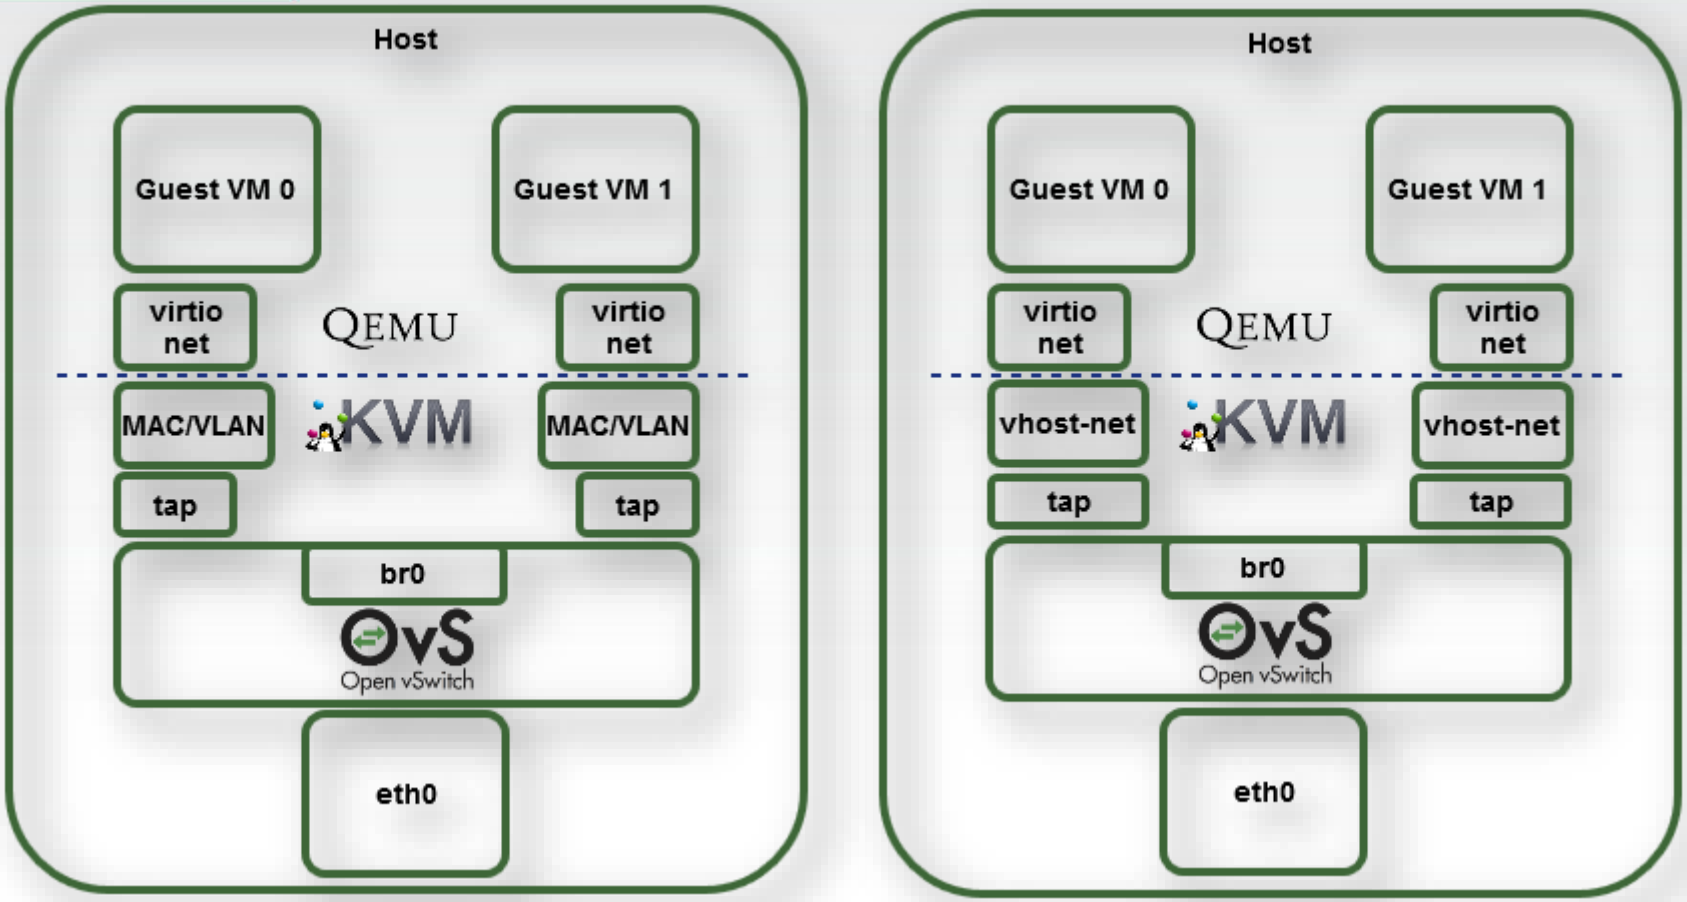

A Walk on the Data Path (cont.)

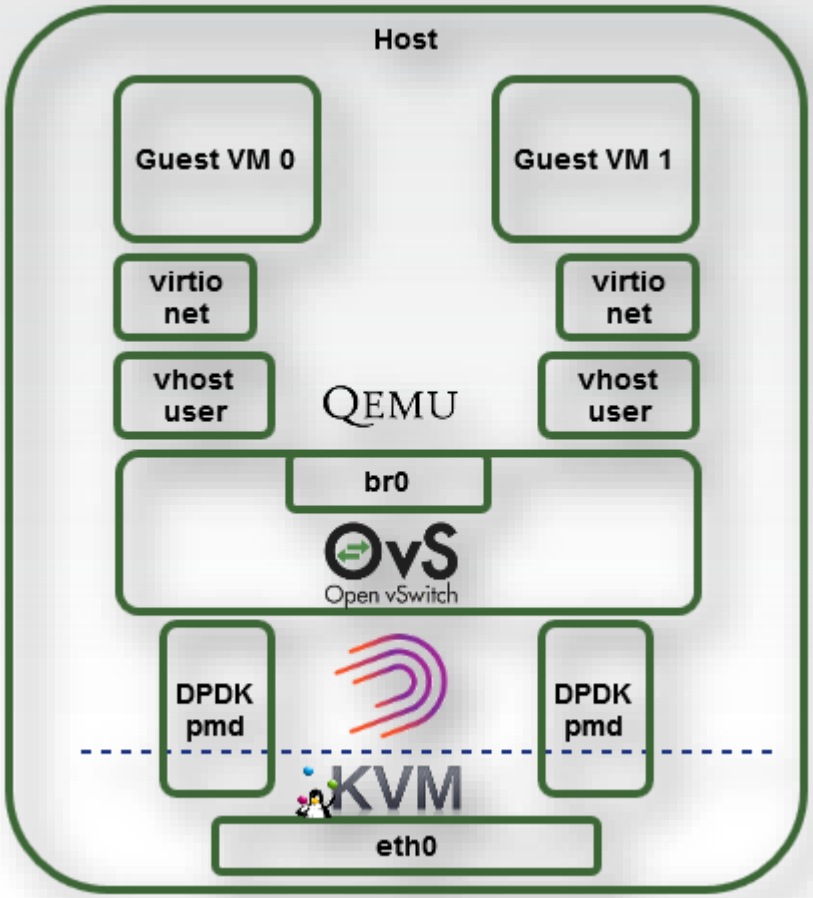

A Walk on the Data Path (cont.)

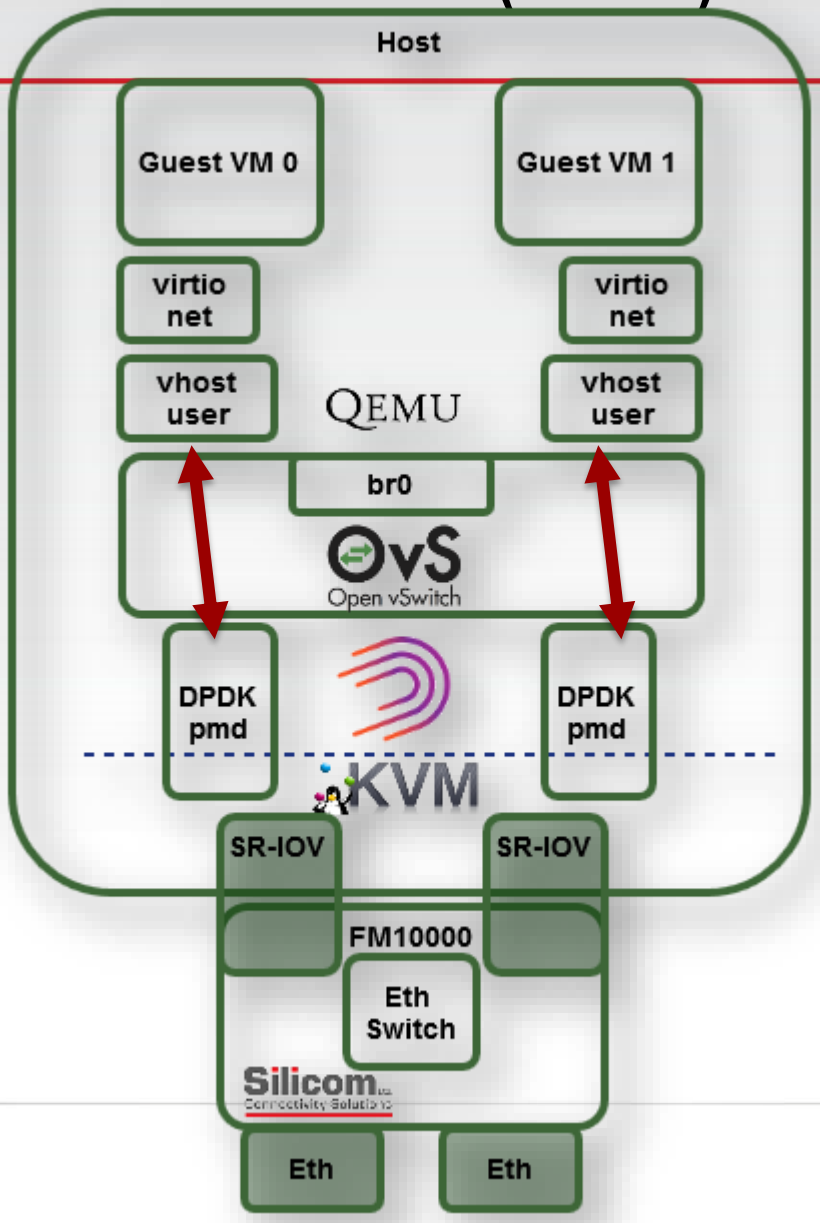

Agenda

- Intel® FM10000 NIC Is Unique
- A Walk on the Data Path
- DPDK OVS Benefits
- Benchmark Results
- Roadmap

DPDK OVS Benefits

- Userspace – Match Interface
  - Predictable core count use
- vhost-user
  - Multiple CPU cores utilization
- SR-IOV and virt-io magical coupling
  - Live migration, etc.

Agenda

- Intel® FM10000 NIC Is Unique
- A Walk on the Data Path
- DPDK OVS Benefits
- Benchmark Results
- Roadmap

Benchmark Results

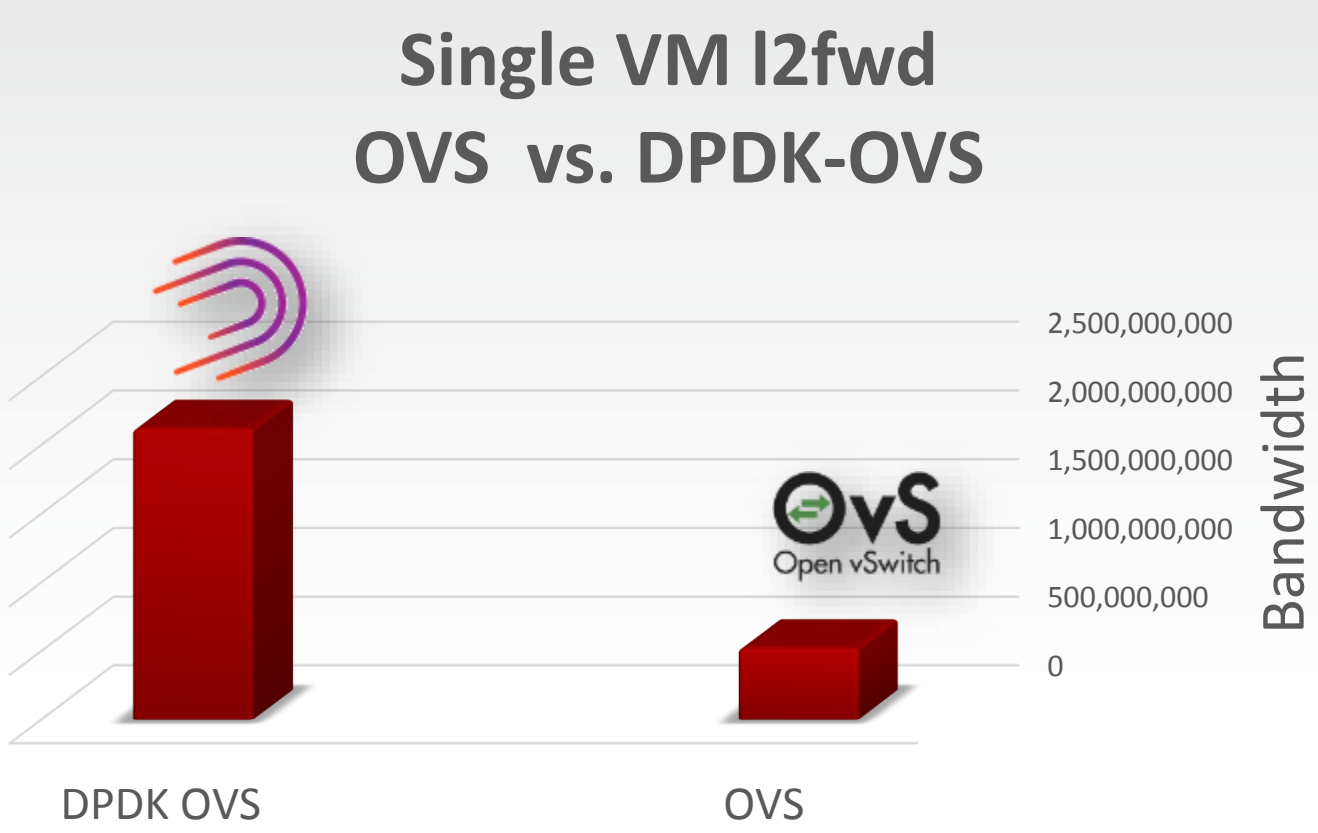

Benchmark Results (cont.)

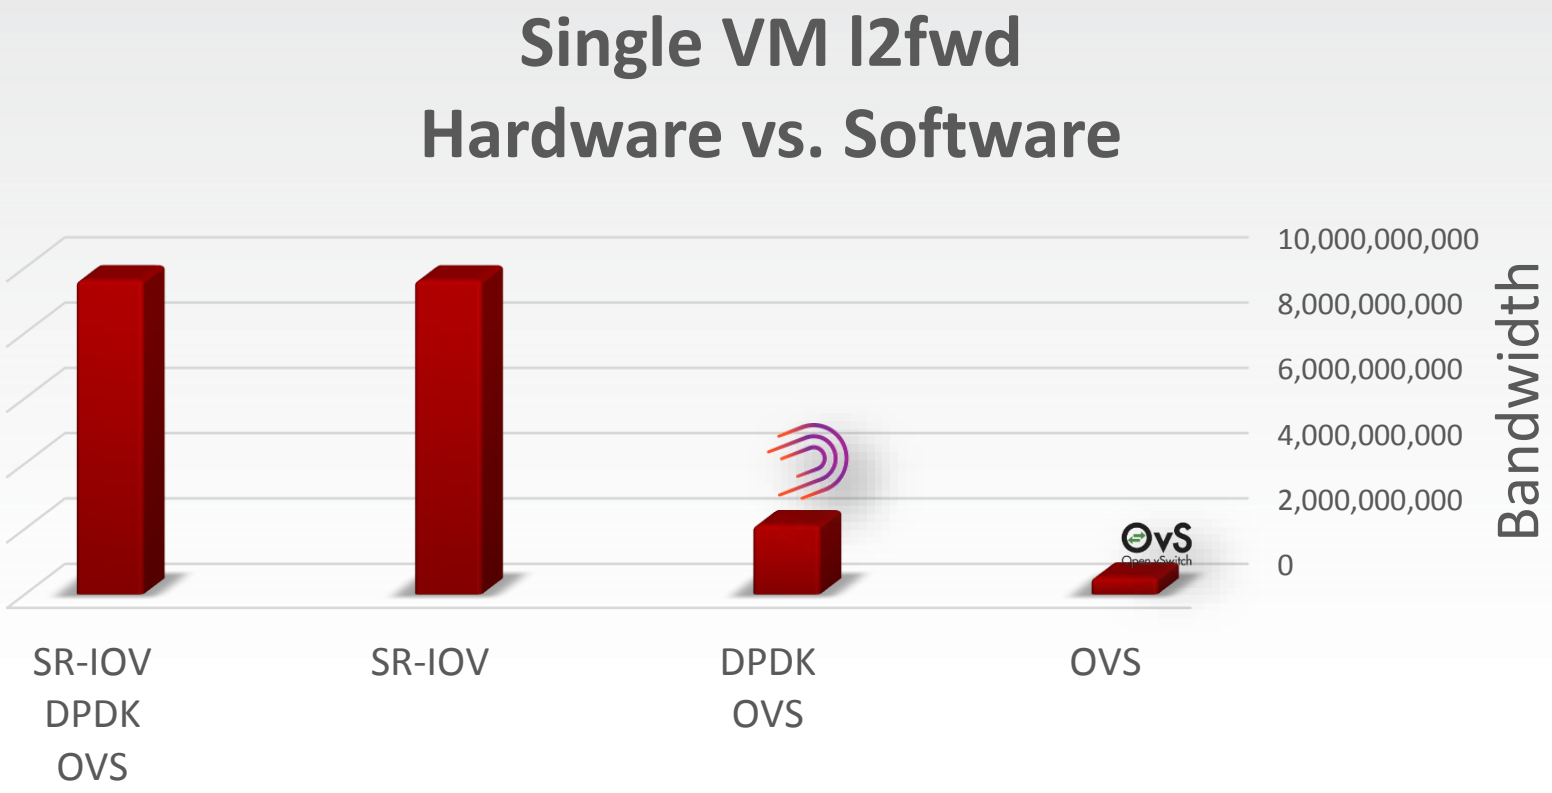

Agenda

- Intel® FM10000 NIC Is Unique
- A Walk on the Data Path
- DPDK OVS Benefits
- Benchmark Results
- Roadmap

Roadmap

- OVS benefits
  - VM L2/L3 networking state (ACL, QoS, SPAN, sFlow) – identifiable and migratable
  - VXLAN, NVGRE
  - Multi host device
- Why FM10K
  - Because it can assist with OVS's whys
- DPIF, Netdev, ofproto

MOTD

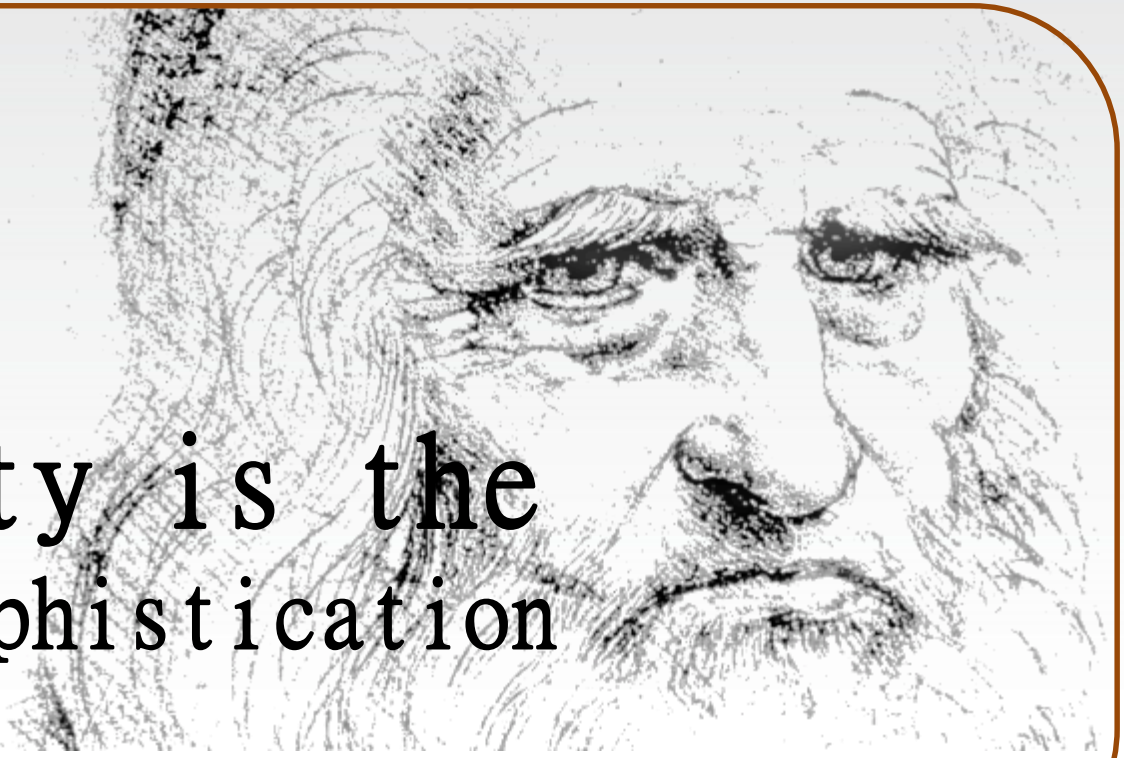

Simplicity is the
ultimate sophistication

Leonardo da Vinci

Summary

- The SR-IOV catch resolved
- No bottlenecks up to wire speed
- Why not Linux bridge(!)
- Gratitude: Maksim Mihailovich

DPDK and QAT Use Cases

DPDK Summit, Aug. '16

Eyal Cohen

Silicom Ltd.
Connectivity Solutions

Agenda

- An Introduction to QAT
- Use Cases for DPDK and QAT
- Monitoring Use Case In-Depth

Agenda

- An Introduction to QAT
- Use Cases for DPDK and QAT
- Monitoring Use Case In-Depth

An introduction to QAT

- Symmetric and asymmetric crypto primitives
- HMAC
- Stream compression

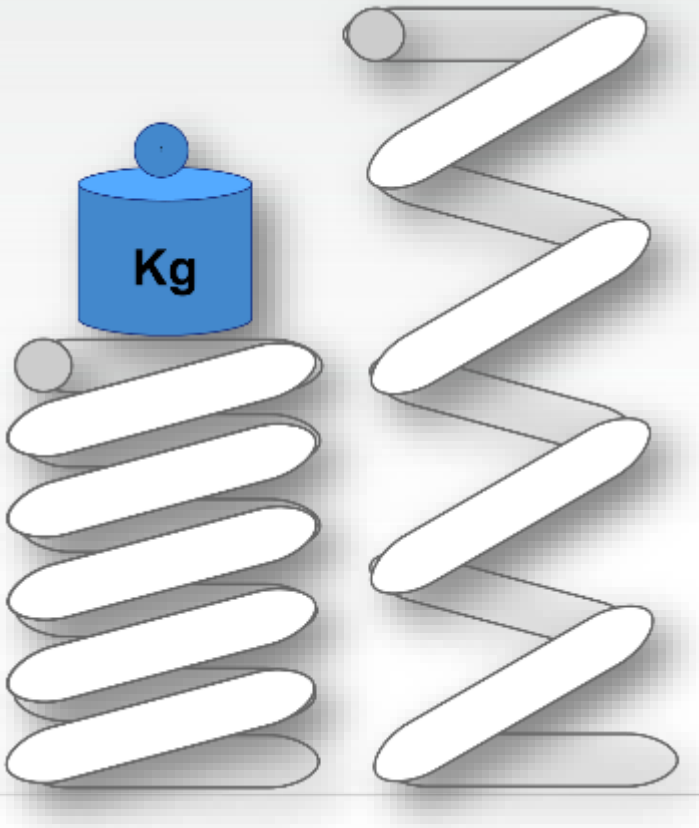

Agenda

- An Introduction to QAT
- Use Cases for DPDK and QAT
- Monitoring Use Case In-Depth

Use Cases for DPDK and QAT

- IPSec
- Monitoring (decryption, no TCP termination)
- Storage (compression and decompression)

Agenda

- An Introduction to QAT
- Use Cases for DPDK and QAT
- Monitoring Use Case In-Depth

Monitoring Use Case In Depth

- SSL web traffic renders monitoring systems blind
- MITM is required
- Benefits of DPDK
- Model of integration with QAT

Monitoring Use Case In Depth (cont.)

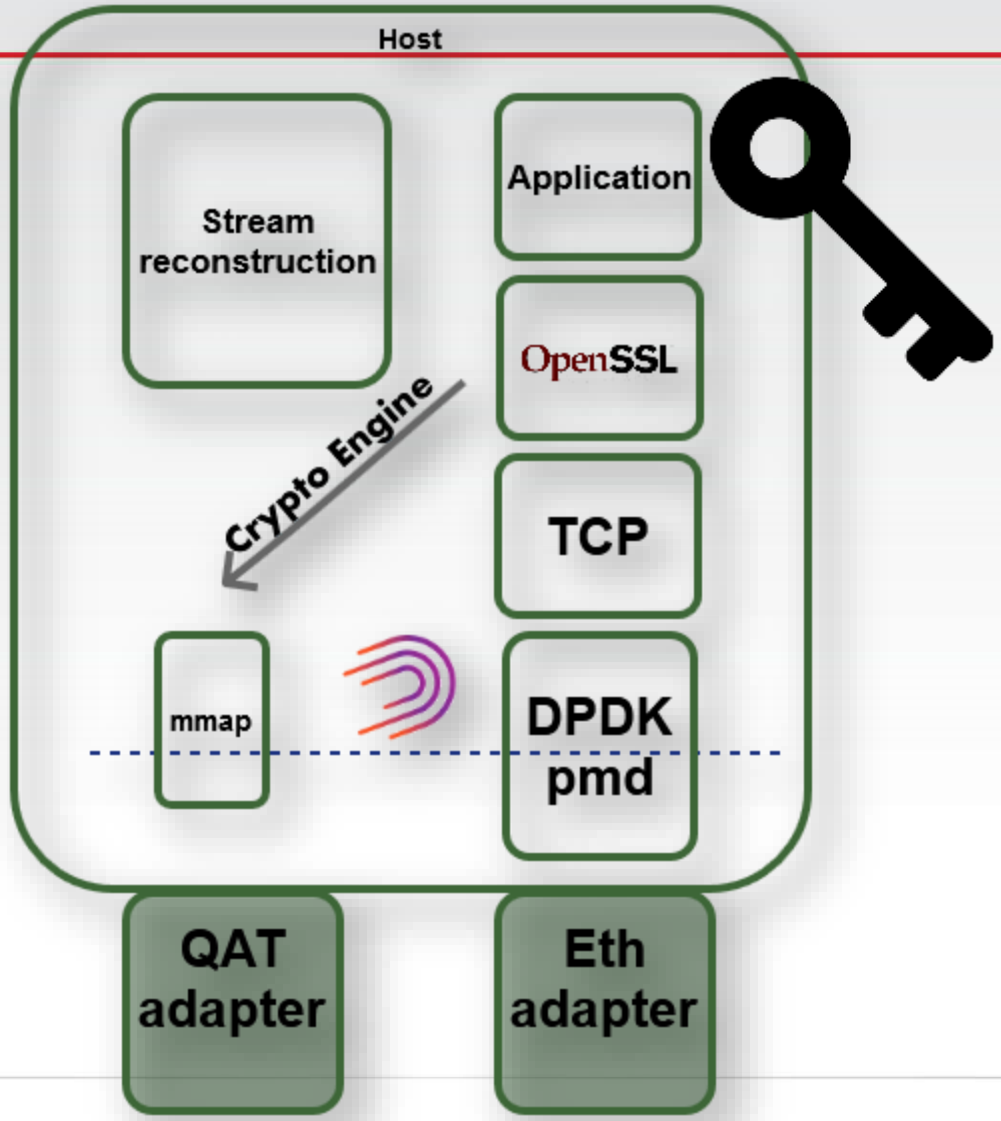

Monitoring Use Case In Depth (cont.)

- MITM on an Intel® Xeon D QuickAssist adapter

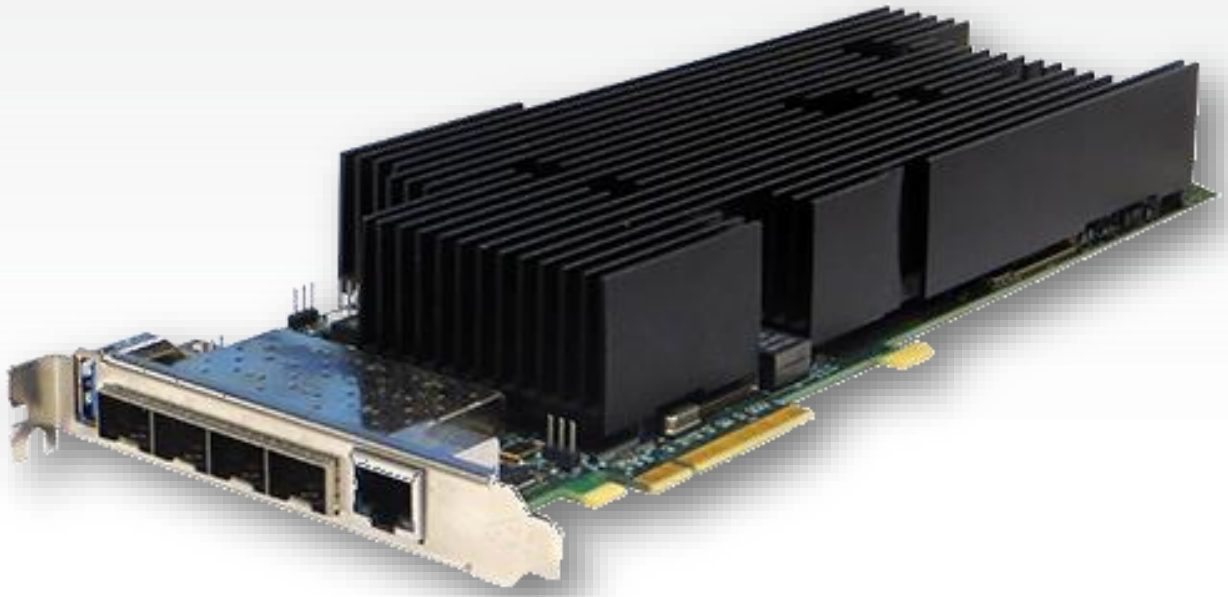

Thank You

Silicom Ltd.
Connectivity Solutions

Backup Data

Silicom Ltd.
Connectivity Solutions

Intel® FM10000

- Programmable Intelligent 100G/40/25/10/ 1 GBE PCI Express Filtering NIC
- Front End Packet Processing offload
- Based on Intel standard controllers and drivers
- Based on Intel Multilayer switches
- Side band management path API to download rules
- The multilayer switch operates at line rate for all packets sizes and conditions
- HW based execution: Drop, Forward, Route, TAP, Police, count, Tag, load balance
- 32K 40b TCAM rules
- Smart NIC / Filtering NIC simplest integration
- Based on standard ASIC, Integrated Switch and NIC
- Cost and power efficient

Thank You

Silicom Ltd.
Connectivity Solutions