FM10000

OVS Fall Conference 2015

Dan Daly, Intel

Legal Disclaimer

General Disclaimer:

Technology Disclaimer:

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com].

Performance Disclaimers (include only the relevant ones):

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

What Has…

- 9 100G Ports
- 9 PCIe Interfaces
- FlexPipe® Frame Processor Inside
- Open Standards Software Compatibility for SDN & NFV

"Red Rock Canyon"

Ethernet Multi-Host Controller

Goals

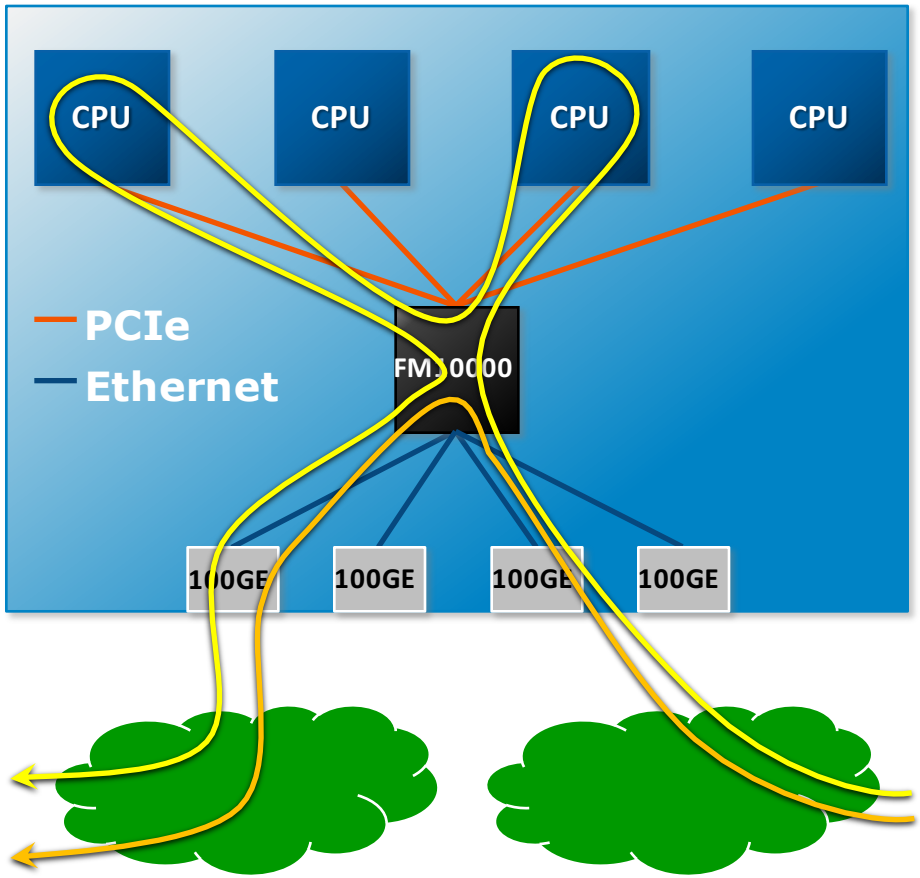

1. Augment packet processing and accelerate virtual switching on Intel® Xeon servers

Net Result: Reduce infrastructure overhead & latency
Enabling more virtual functions
Chained together using SDN

2. Flexible Ports and Interfaces to Server Platforms

Net Result: New Form Factors & Levels of Integration

Form Factors & Integration

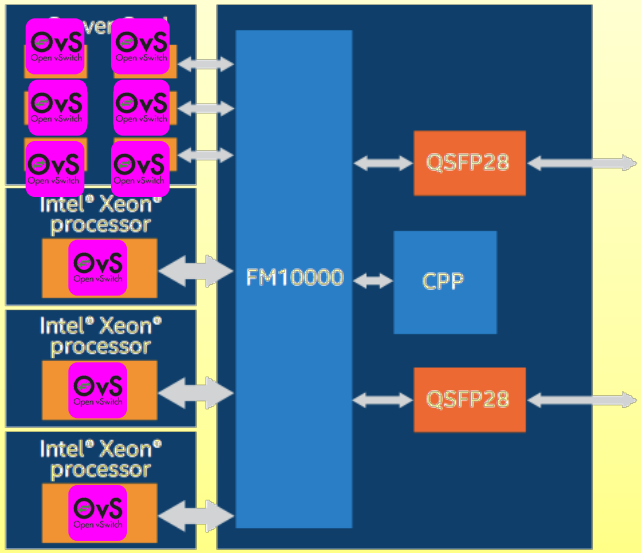

Multi-Host Modular Server Platform

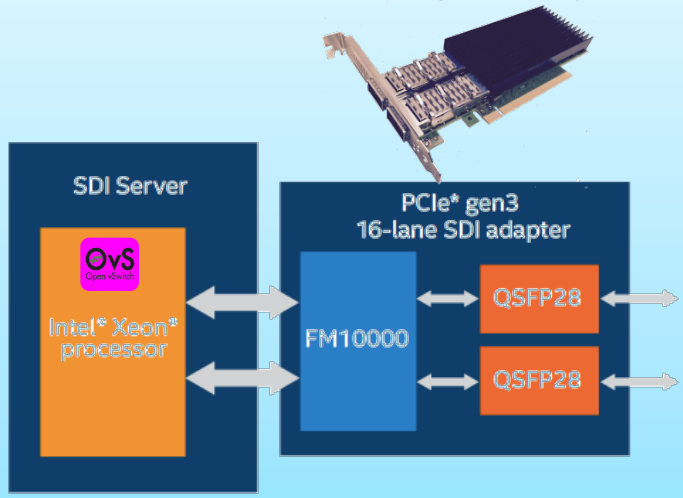

25Gbps / 100Gbps SDI Adapter

Open vSwitch Enabled & Accelerated in All Form Factors

Form Factor & Integration Advantages

1. 25G Serdes for multiple 25G & 100G Ports
   - 2.5x line rate improvement
   - Copper & Optical Cabling Supported
2. Multi-Socket Support
   - Avoid latencies transferring over QPI
   - Balance traffic across sockets
3. Multi-Host Support
   - Integrate multiple hosts to enable sharing of resources & higher density

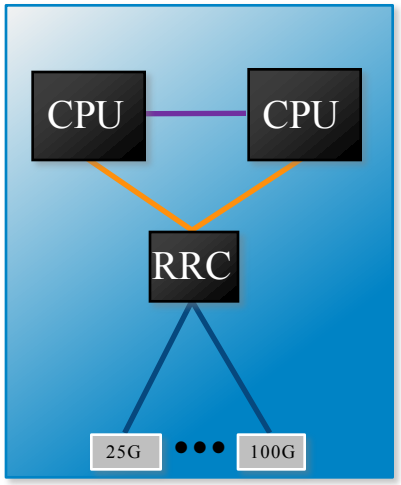

FlexPipe® Forwarding Use Cases

Lowering the cost of network virtualization within the infrastructure

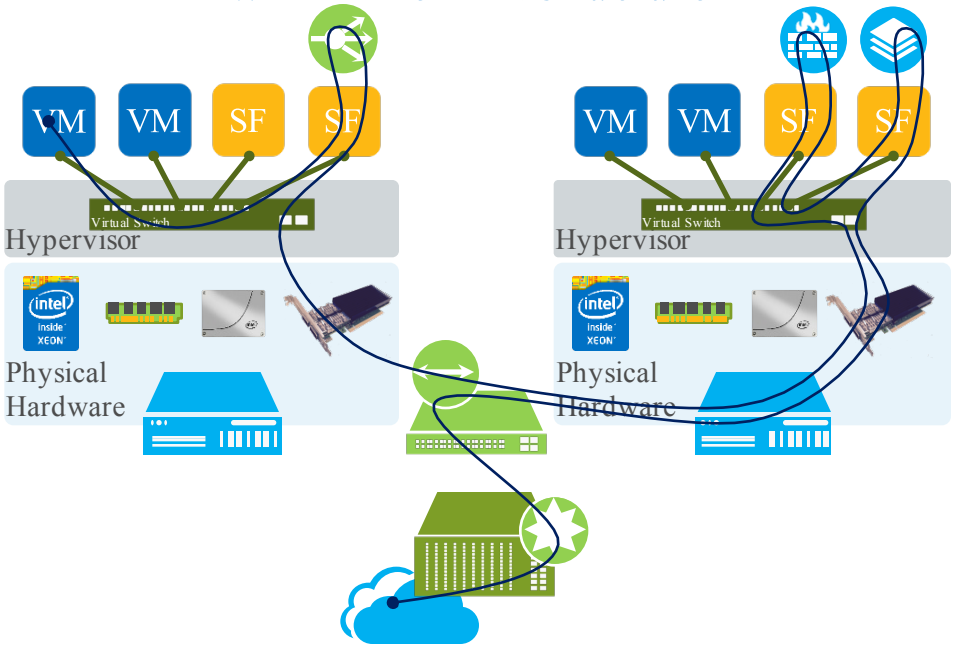

Focus on Virtual Switching & Service Function Forwarding

Scaling Multi-Core VNFs Running DPDK

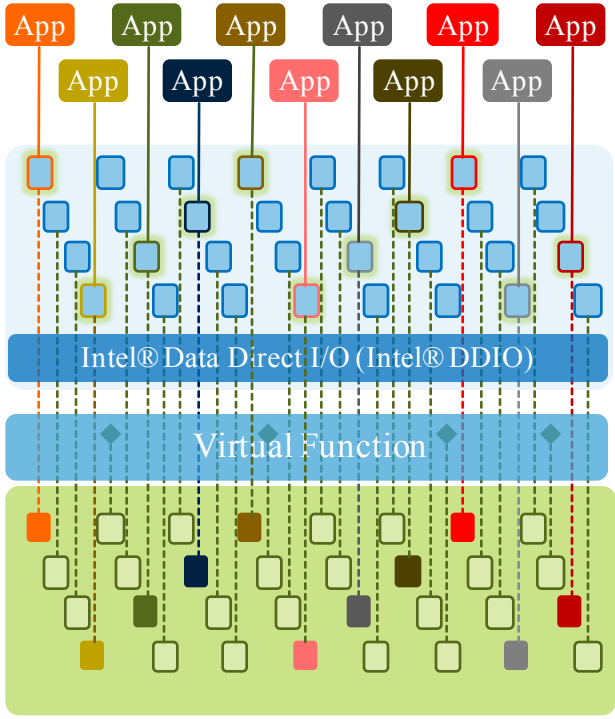

VNF Programmable Packet Classification

Open vSwitch Software Advantages

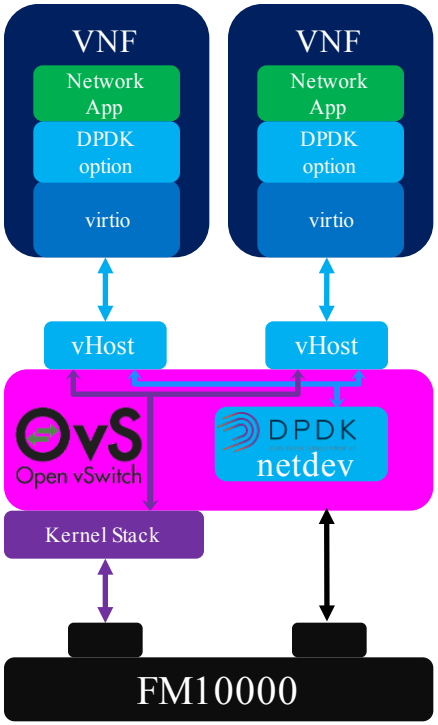

- Support for Kernel & DPDK OvS Data Paths
  - Choose data paths depending on use case
  - Supports simultaneous operation
- DPDK Poll Mode Driver Optimizations
  - Vector PMD Driver (DPDK 2.2)
  - Stateless Offloads (TSO, RSS, checksums)
  - Statelesss Offloads in the presence of tunnels

Support for Open vSwitch 2.4

FlexPipe™ Used Under a vSwitch

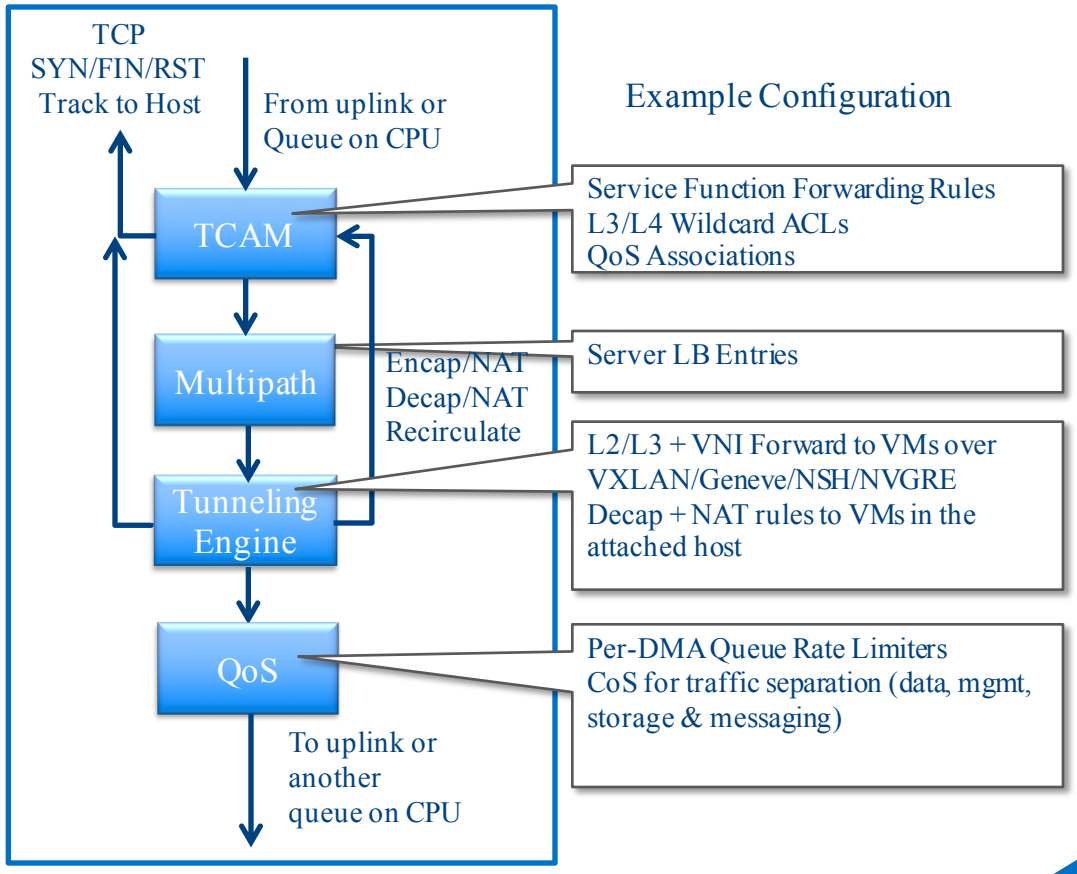

Acceleration Using FlexPipe™

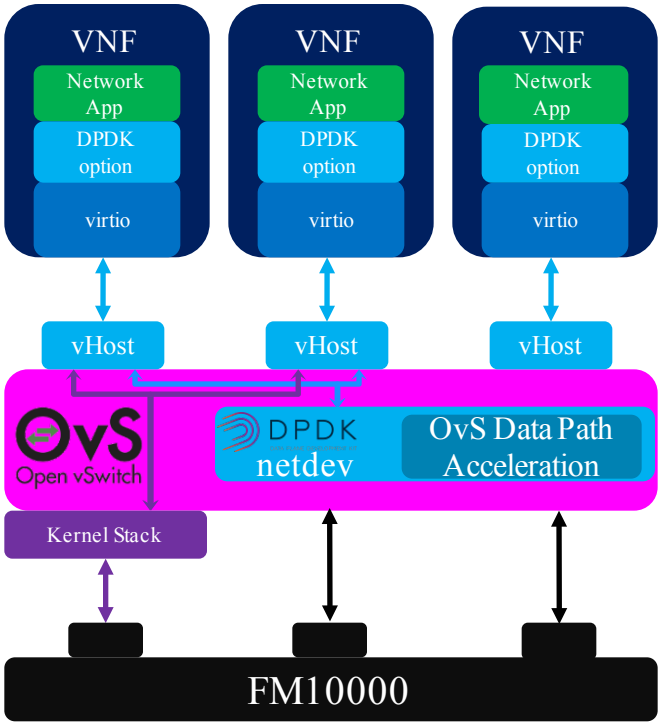

- Accelerate Wildcard Match
  - Tag with Metadata
  - Filter, Count, Mirror, Sample
- Accelerate Tunnel & SFC Encap/Decap
  - Tunneling & service function chaining information put into the DPDK netdev
- Accelerate Multi-Queue Virtio
  - Allow FM10000 to choose virtio queue (RSS, FlowDirector, & filtering)

OVS Controlled SR-IOV

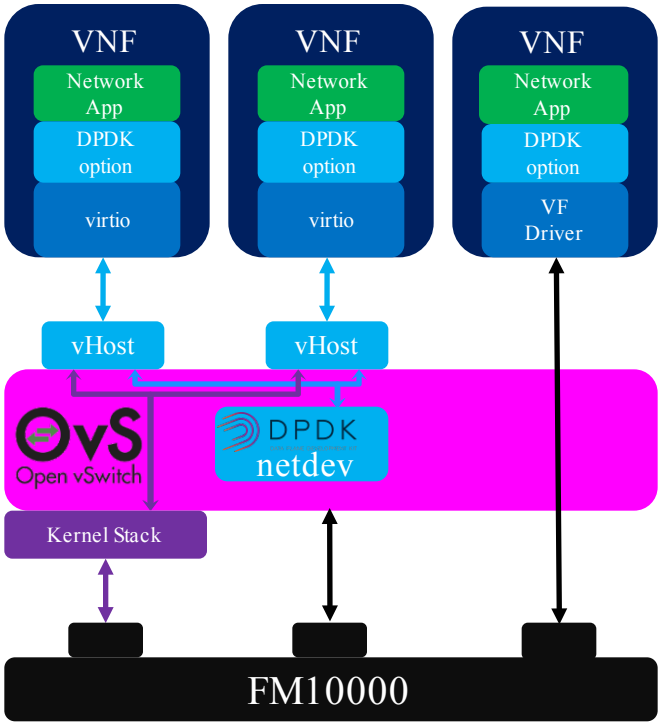

- Preserves OVS Control Point
  - OVS kernel and/or DPDK are the default data paths
- Performance & Latency Sensitive Flows
  - Pushed down into FM10000
  - Directly forwarded in/out VF
- Consistent Performance
  - BW & latency stays the same independent of number of tunnels, ACLs, mirrors, etc.

Example: 4 DPDK VMs, OVS, VXLAN

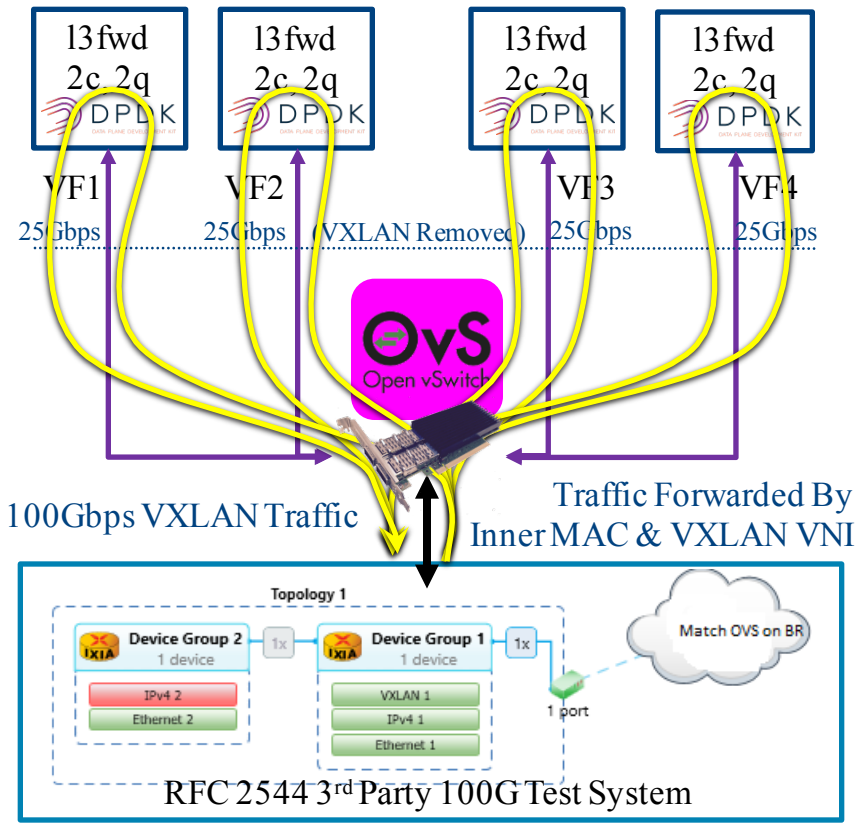

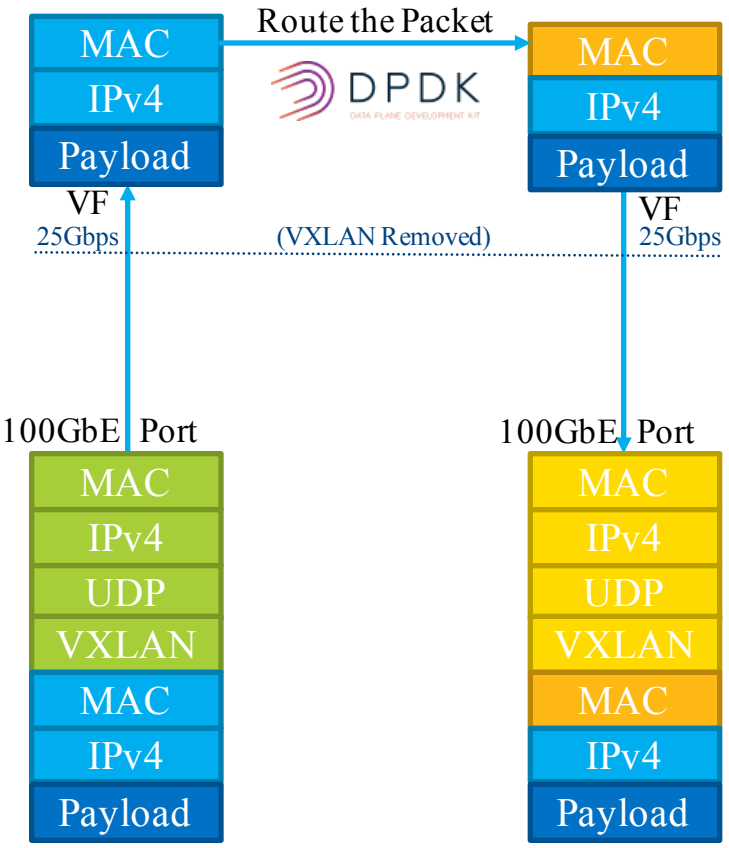

FM10000: Software Flexibility

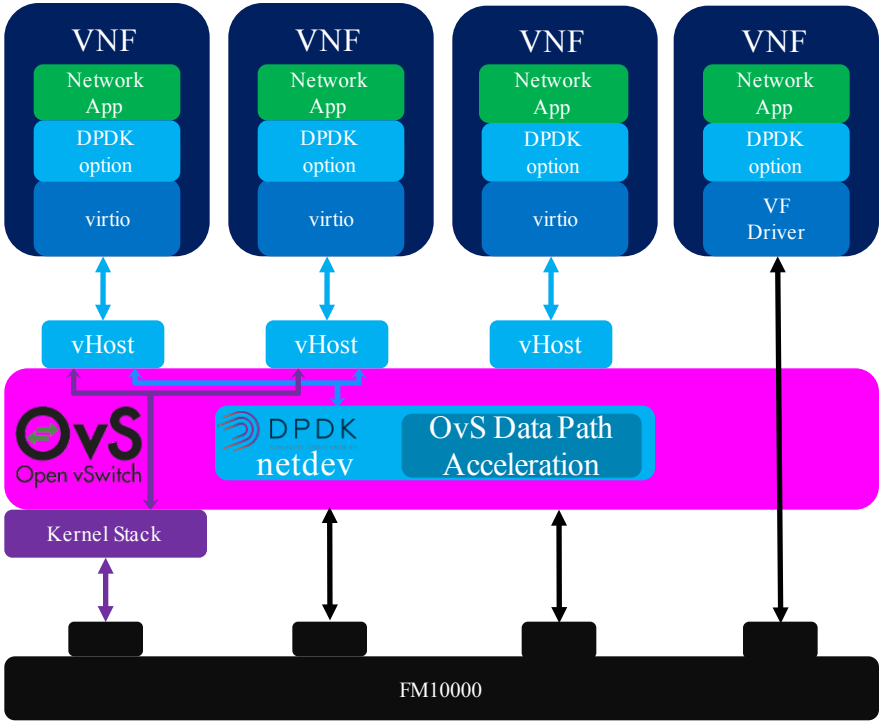

100Gb Line Rate Virtual Switching

Scale Virtual Applications to 100Gb

DPDK

Open Source Eco-System Compatibility

openstack CLOUD SOFTWARE

OvS Open vSwitch

OPEN DAYLIGHT

FM10000 – Ethernet Multi-host Controller

Up to 200Gbps of PCIe bandwidth

- Supports up to eight hosts

Flexible Ethernet ports

- 1/10/25/40/100GbE

Frame processing up to 960Mpps

- Integrated TCAM further accelerates performance

Two integrated tunneling engines (encap/decap)

- Tunneling for network service chaining headers (NSH)
- VXLAN, NVGRE and GENEVE tunneling

DPDK Acceleration Enhancements

- Allows vSwitch accelerations which enable more efficient NFV platforms