ONOS-based Data Plane Acceleration Support for 5G

Dec 4, 2018

SKTelecom

Partner for New Possibilities

A Brief Introduction to 5G Network

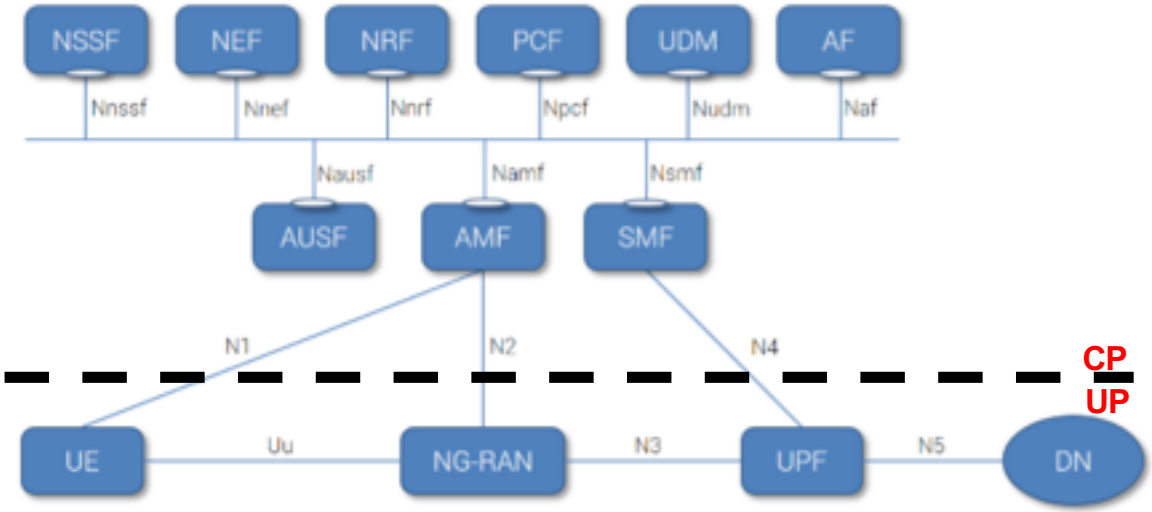

5G Core Service-based Architecture

- **Key Functions**
  - UE Authentication
  - UE Mobility Mgmt.
  - UE Session mgmt.
  - QoS mgmt.
- **Key Changes v.s. 4G**
  - Cloud-native design
  - UP/CP Seperation
  - Telco API exposure
  - Network Slice
  - NF repository support

Why DPA needed in 5G?

- According to 5G *"state of the art"*, each 5G Network Function should be deployed in a virtualized or containerized manner
- Virtualization/containerization inevitably gives network performance degradation
- Still needs line-rate performance for I/O intensive NF's (e.g., RAN, UPF, IMS nodes)

Legacy Technologies

- **As incumbent solutions, PCI-Passthrough and SR-IOV are widely being used**

Functionaliy	No DPA	PCI-PT	SR-IOV
Easy to configure	Very easy	Easy (Flavor, PCI whitelist,, alias)	Difficult (NIC Speific configuration, agent setup, …)
Easy to manage	Easy	Difficult (Cannot monitor this)	Normal
SDN-based management	Easy	Impossible	Impossible
Performance	7~80% of line rate	Line rate	Line rate if VF/PF = 1:1
Live Migration	Support	X	X

New Candidates

Functionaliy	OVS-DPDK	Smart NIC
Versatility	DPDK-enabled NIC required No limitation on Host OS	Specific driver required Host Kernel Ver. Limitation, OpenStack Pike~
Easy to Install	Easy (Well-made wiki, so many test results exist in Google)	Difficult (Private wiki, no Google results available)
Easy to configure	Difficult (So many tuning options)	Easy (same CLI as open vswitch)
CPU Usage	10Gbps per core (What will happen with upcoming 40G I/f era?)	NO CPU Usage in case of Full Offloading
Packet Loss	Increases as throughput increases	0
Live Migration	X	X

What if we use containerized NF?

Limitations of Container

	VM/Openstack	Container/k8s
PCI-PT	O	X
SR-IOV	O	X
Multiple NIC's	O	X
Huge Page, DPDK	O	X
NUMA aware deployment	O	X
Smart NIC	O	X

Offloading Model

- **Half Offloading Model**
  - Specific to OVS 2.6.1 for Agilio provided
  - Use same CLI as OVS
    - E.g., ovs-vsctl, ovs-ofctl, ovs-appctl
  - ODL Plugin provided

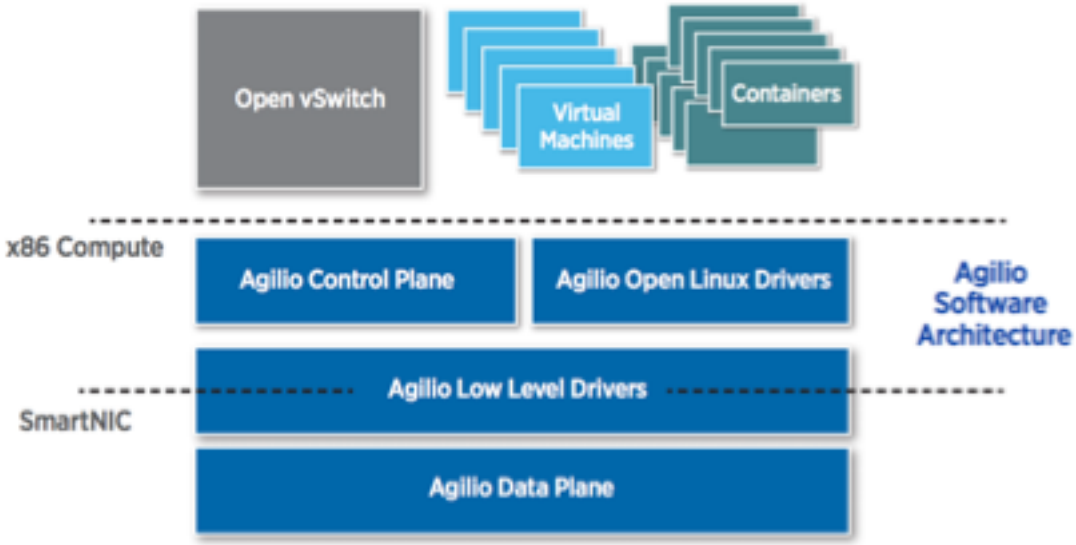

Netronome Agilio

Source: https://www.netronome.com/media/documents/PB_Agilio_OVS_SW.pdf

Offloading Model

- **Full Offloading Model**
  - Specific to OVS 2.8.1 for Cavium provided
  - Use same CLI as OVS
    - E.g., ovs-vsctl, ovs-ofctl, ovs-appctl
  - ODL Plugin provided
  - ONOS plugin provided (SKT's efforts)
  - **Zero CPU usage**

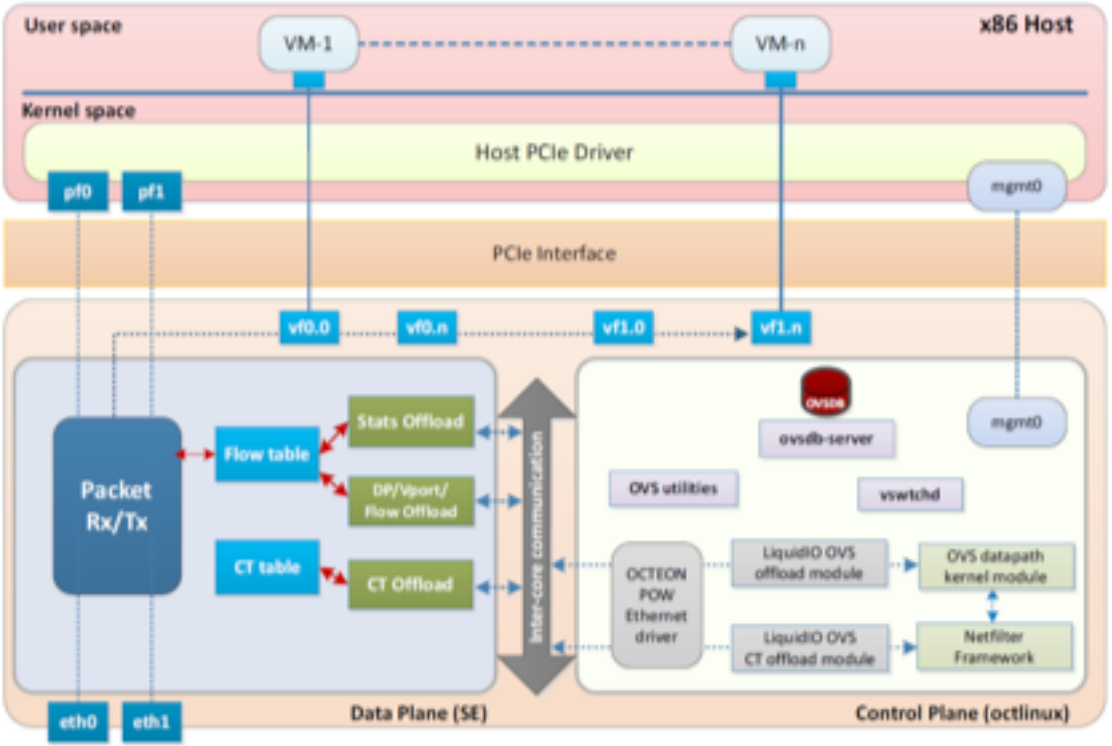

Cavium LiquidIO

Source: https://www.cavium.com/Documents/WhitePapers/Adapters/WP_LiquidIO_OVS_Software_Architecture.pdf?x=2

LiquidIO PoC: Requirements and Features

- **Requirements**
  - CentOS 7.4, Ubuntu 14.04
  - LiquidIO driver installation (upstreamed since CentOS 7.5)
  - OpenStack Pike
- **Main Features**
  - Hardware: 2 x 10/25G SFP+
  - VxLAN/NVGRE/GENEVE encap/decap offloaded
    - In OpenStack point of view, the performance of overlay network is identical to FLAT network
  - SR-IOV mode suggested and 126 VF's supported
    - In OpenStack point of view, all VMs should be bound to VF

LiquidIO PoC: PoC Scenario

- **Experiment Setup**
  - OpenStack version: Pike
  - Network mode
    - VxLAN
  - VM flavor
    - vCPU 4core, 8GB RAM
  - OVS version
    - 2.8.2, 2.6.1
  - Packet generator
    - iperf3
  - Metrics
    - Inter VM throughput

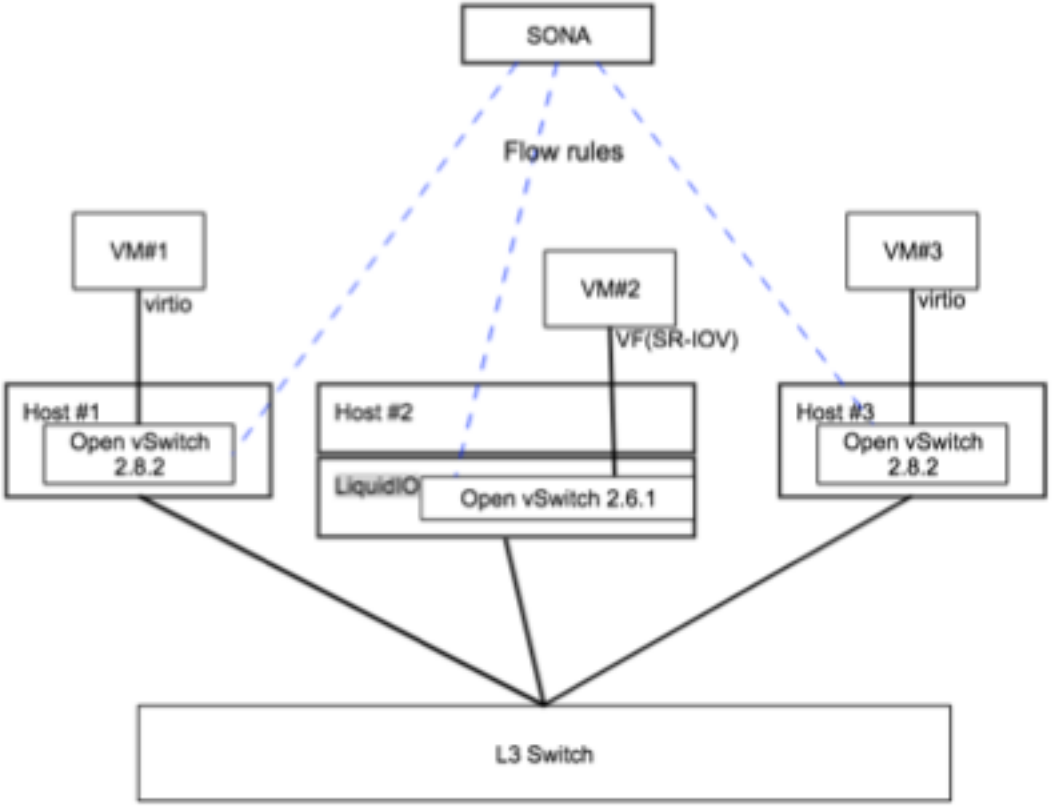

LiquidIO PoC: Test Result

- **Result**
  - Typically VxLAN shown performance degradation from 26% to 54% w.r.t. MTU sizes
  - More than 512byte of MTU, line rate performance guaranteed
  - Network agnostic throughput performance guaranteed

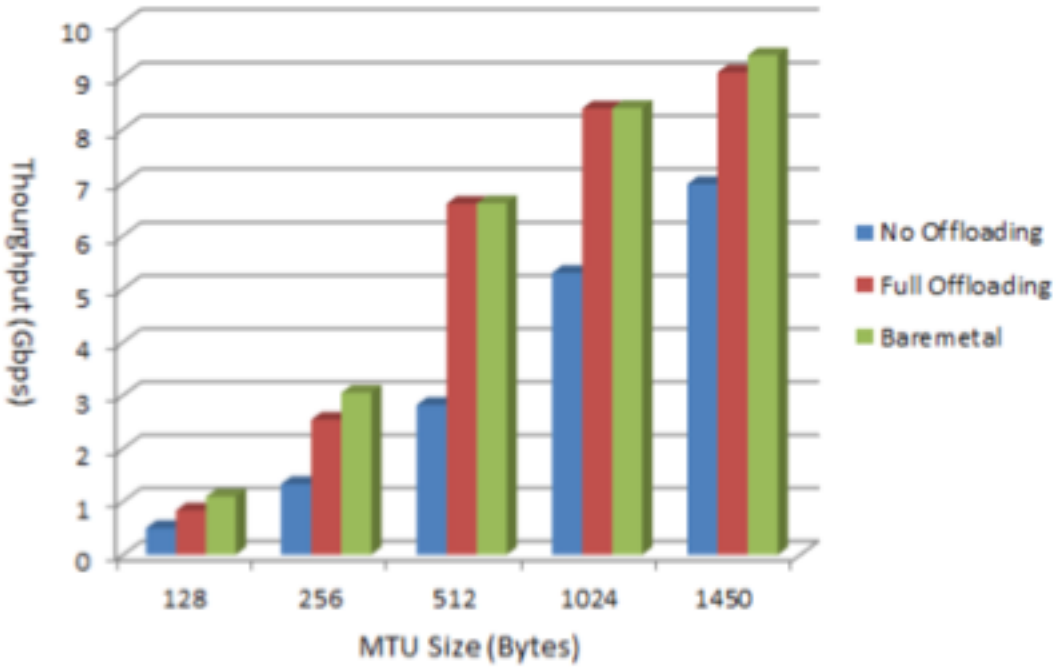

LiquidIO PoC: Lessons Learned

- **Needs more interfaces**
  - Telco VNF's has multiple, isolated physical interfaces per purpose (e.g., 10G x 2, 1G x 4)
- **Needs help from manufacturer**
- **Multiple smart NIC's in a same host should be considered**
  - So far, a host cannot have the integrated view of its multiple smart NIC's
- **Should improve stability**
  - Multiple bugs are reported in this PoC
- **General-purpose smart NIC market is downsizing**
  - Marvel acquires Cavium ☺

Conclusion: What DPA tech. will be used in 5G?

	Comments
PCI-PT	Will be disappear
SR-IOV	Still has a strong position
OVS-DPDK	Configuration should be simple Packet loss problem should be solved(Telco allow 0.001% PER) BW per core should be improved
Smart NIC	Still needs time to be matured Variety of form factor should be available
P4 Smart NIC	Reasonable price should be given

Thank You!