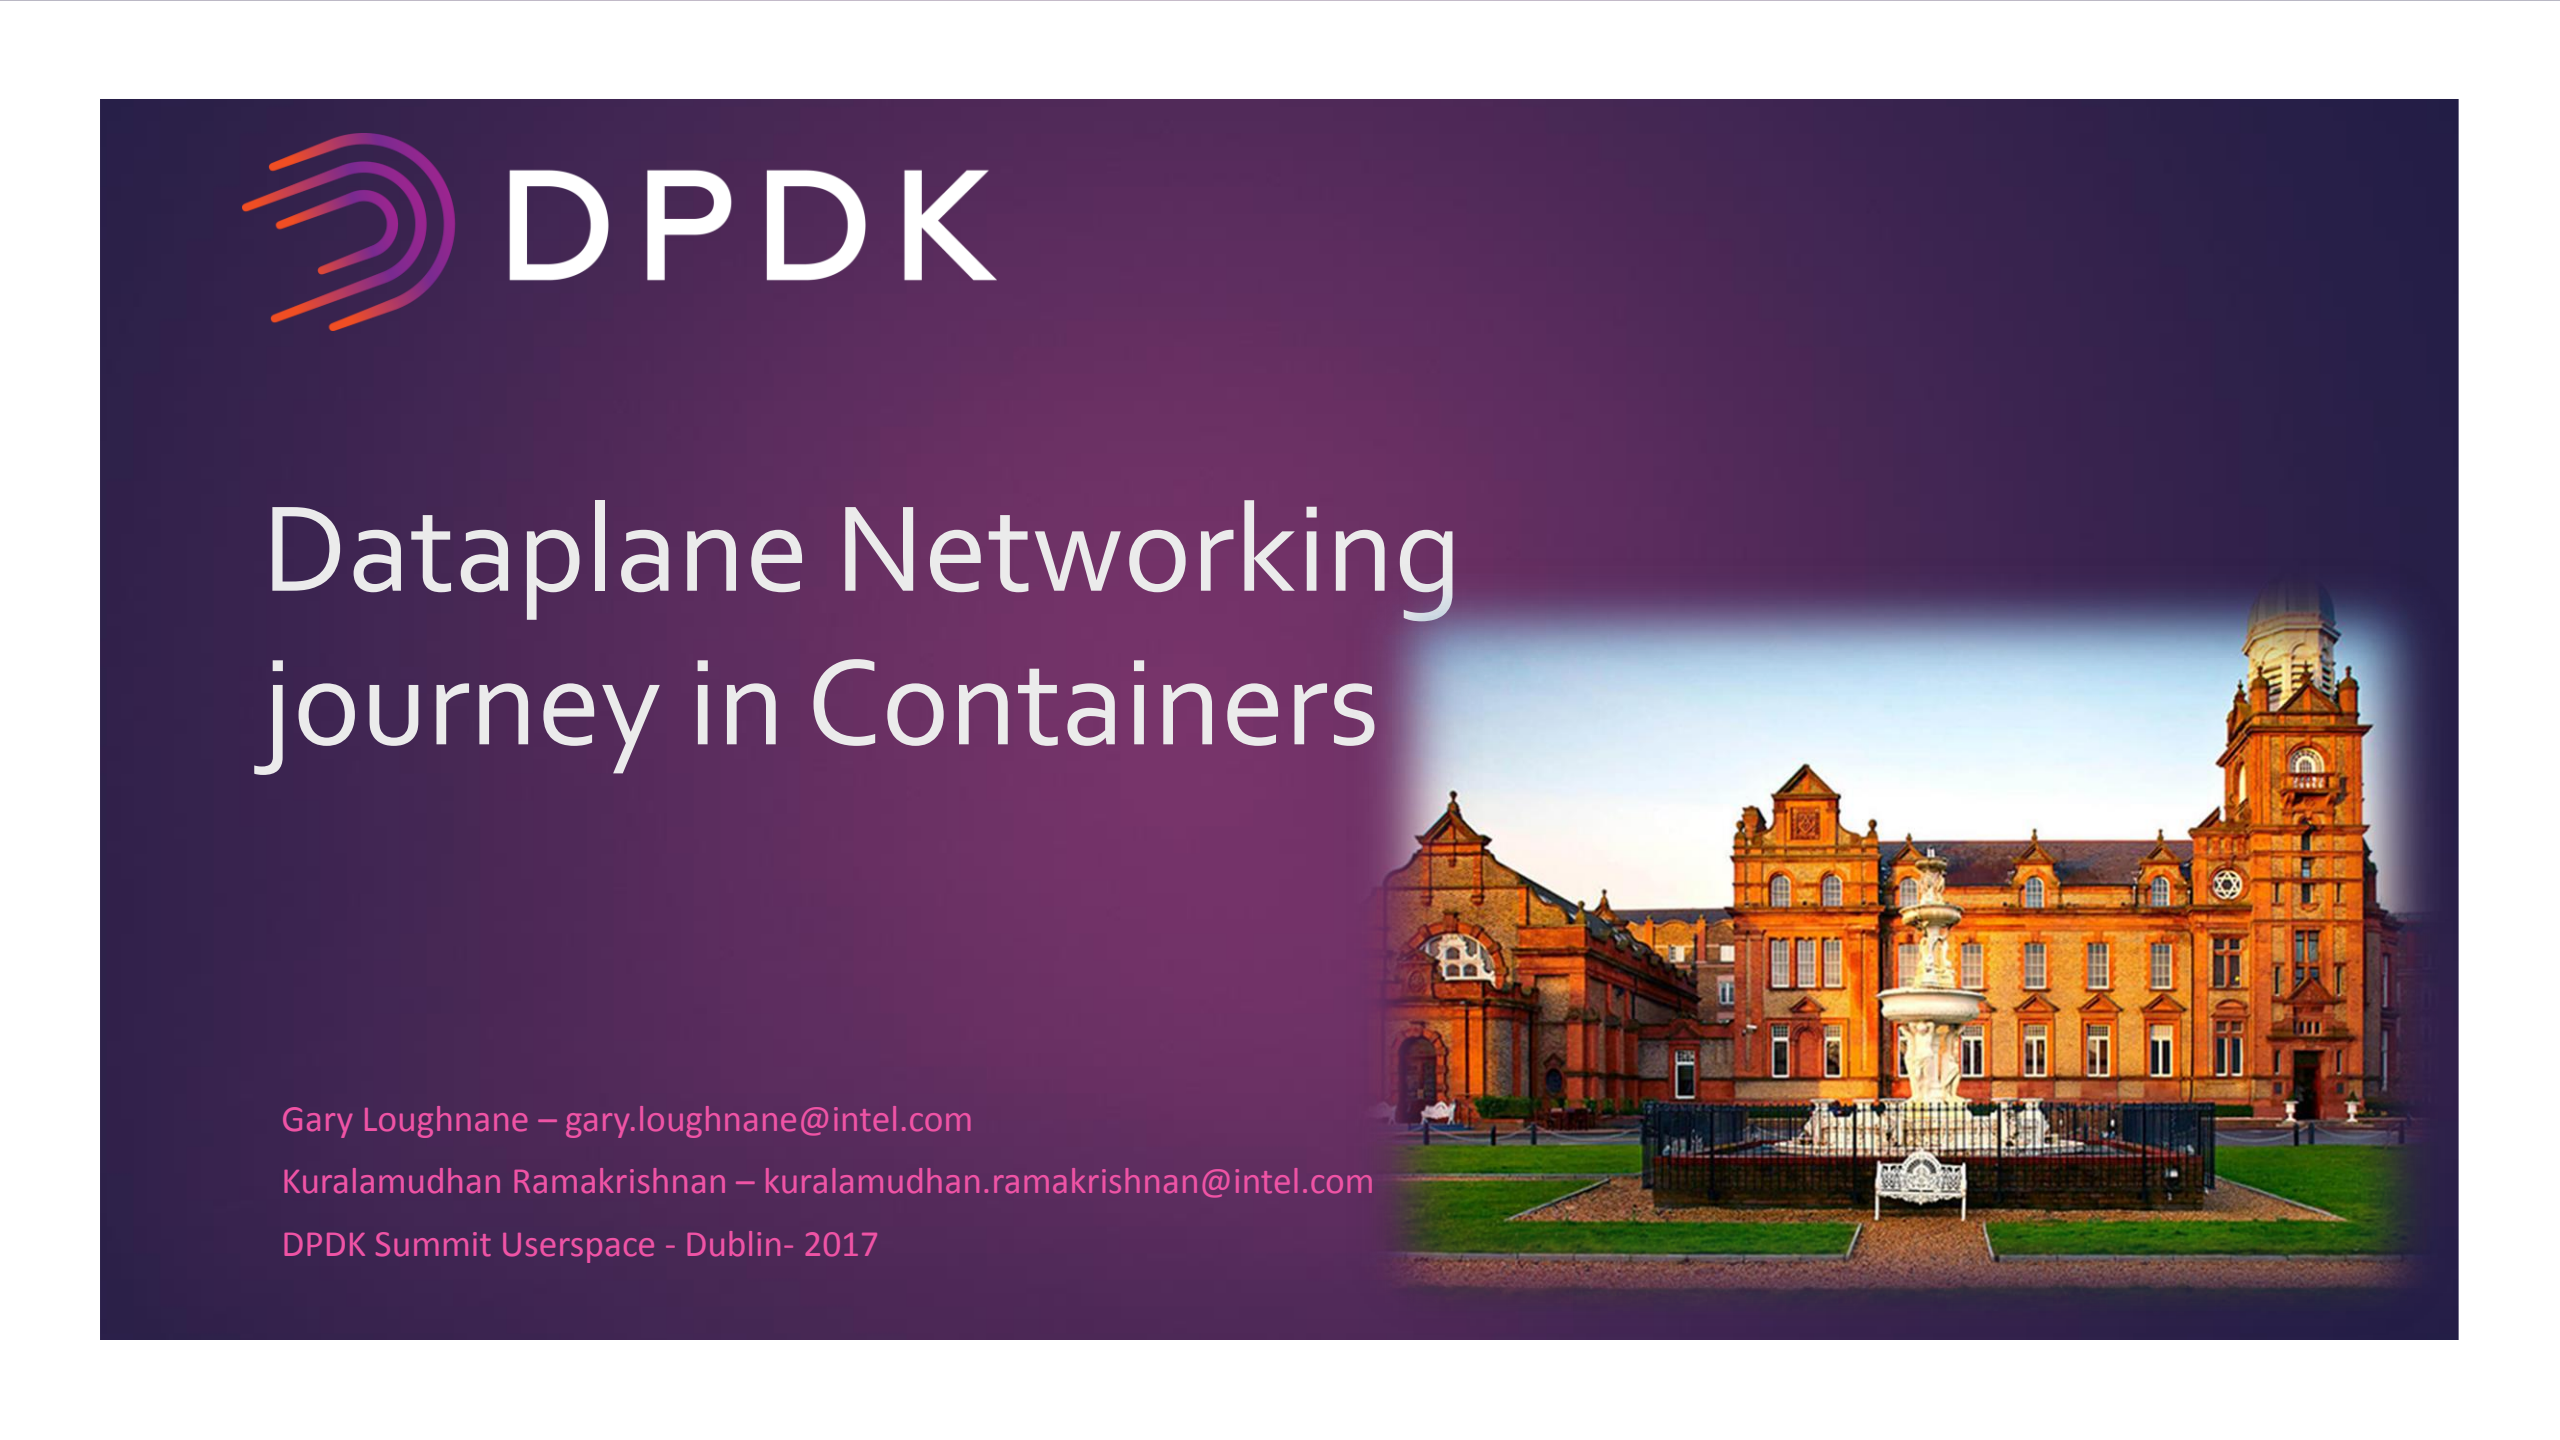

Dataplane Networking journey in Containers

Gary Loughnane – gary.loughnane@intel.com

Kuralamudhan Ramakrishnan – kuralamudhan.ramakrishnan@intel.com

DPDK Summit Userspace - Dublin- 2017

Discussion topics

- Container Deployment Models
- Container Bare Metal Reference Architecture
- Container Unified Infrastructure Reference Architecture

Network Cloudification
Multiple Deployment Models

DPDK

CLOUD NATIVE COMPUTING FOUNDATION

VNFs

vEPC	vNAT	vIMS	vRouter	vGGSN	vFirewall
vCPE	vRNC	vHLR	vSGSN	vMME	vIDS

NFVi- Network

OvS Open vSwitch | fd.io | DPDK | SR-IOV | CNI

NFV Orchestration

ONAP | Open Source MANO | kubernetes | openstack | OPEN DAYLIGHT | OVN (Open Virtual Network)

VM

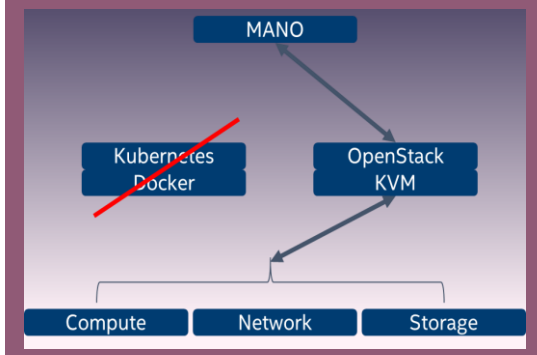

Containers

Bare Metal

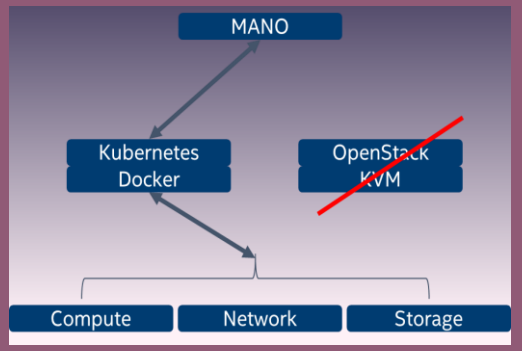

VM
Containers

VM | **Containers**

Hybrid

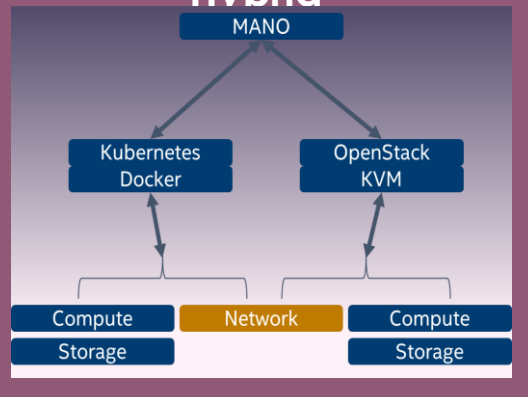

Unified

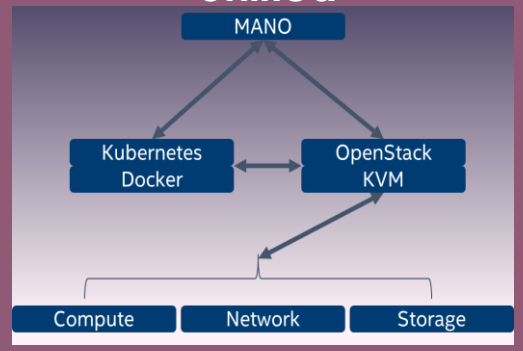

Network Cloudification
Multiple Deployment Models – Today Discussion Focus

DPDK

CLOUD NATIVE COMPUTING FOUNDATION

VNFs

vEPC	vNAT	vIMS	vRouter	vGGSN	vFirewall
vCPE	vRNC	vHLR	vSGSN	vMME	vIDS

NFVi- Network

OvS Open vSwitch | fd.io | DPDK | SR-IOV | CNI

NFV Orchestration

ONAP | Open Source MANO | kubernetes | openstack | OPEN DAYLIGHT | OVN (Open Virtual Network)

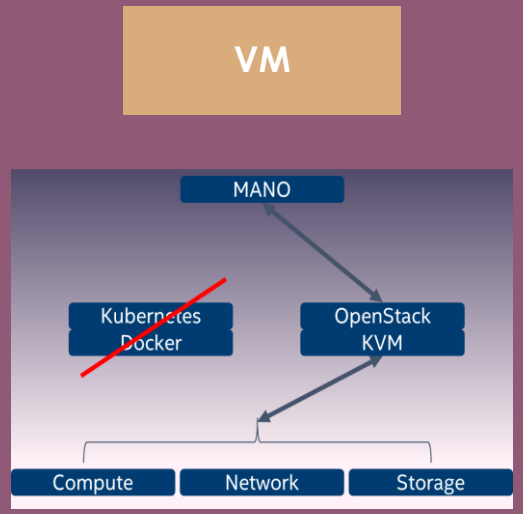

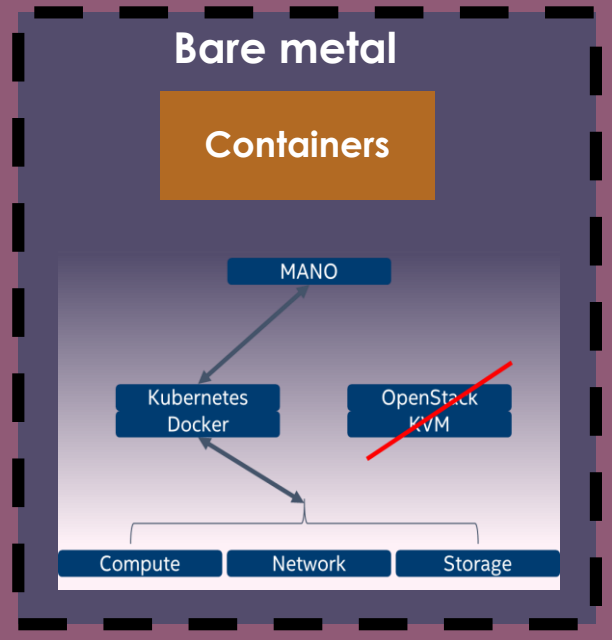

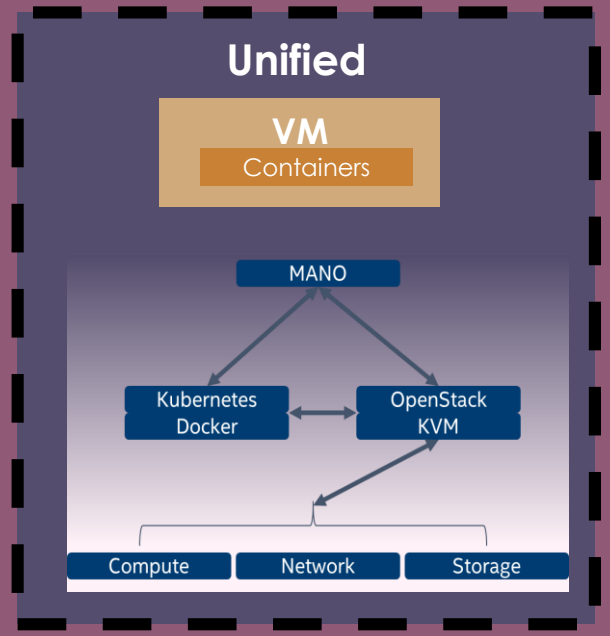

Address Container Networking Industry Gaps
Intel Containers NFV Reference Architectures (Experience Kits)

DPDK

Identified Gaps

Resolve, Integrate

Communicate

Containers Compute & Networking

Existing Solution

MULTUS
Multi-network Support

SR-IOV CNI PLUGIN
Resource Isolation

DPDK CNI PLUGIN
Data Plane Support

NODE FEATURE DISCOVERY
Data Center Heterogeneity

CPU MANAGER for K8s
CPU Core Pinning

RESOURCE MANAGMENT
NUMA Awareness
Huge pages
QAT

SW Contributions Demo Integration

Open Source Projects

VNF

Kubernetes v1.5.0
Docker v1.12

flannel SRIOV CNI MULTUS
Flannel v0.5.0
Multus V1.0
SRIOV v0.2-alpha

DPDK DATA PLANE DEVELOPMENT KIT
DPDK v16.11

CentOS
CentOS Linux 7 (Core)
Linux 3.10.0-327.36.3.el7.x86_64

intel
82599ES 10-Gigabit SFI/SFP+ Network Connection
X710 for 10GbE SFP+

Container BM Ref. Architecture Rel. 1 Oct. 2017

Container UI Ref. Architecture Rel. 1 Feb 2017

Best Practice Guidelines
Experience kits (Examples)

MULTUS

Data Plane Scale

EPA - CMK

Kuryr

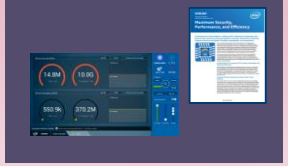

Network Cloudification
Multiple Deployment Models – Today Discussion Focus

DPDK

VNFs

vEPC	vNAT	vIMS	vRouter	vGGSN	vFirewall
vCPE	vRNC	vHLR	vSGSN	vMME	vIDS

NFVi- Network

OvS Open vSwitch | fd.io | DPDK | SR-IOV | CNI

NFV Orchestration

ONAP | Open Source MANO | kubernetes | openstack | OPEN DAYLIGHT | OVN (Open Virtual Network)

VM

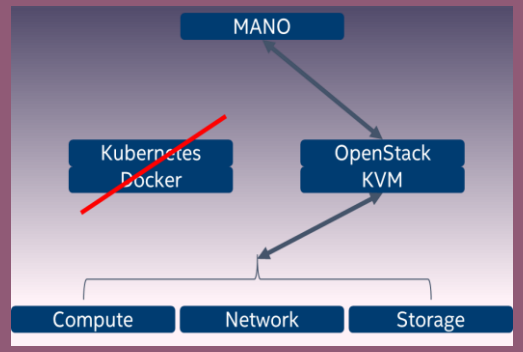

Bare metal

Containers

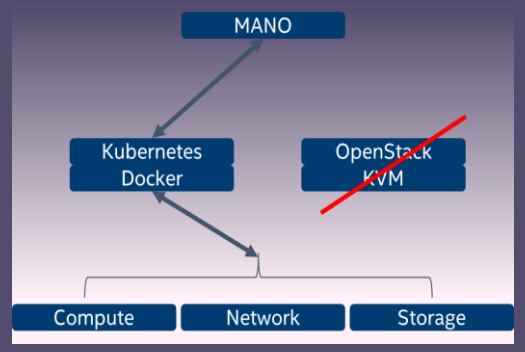

Unified

VM
Containers

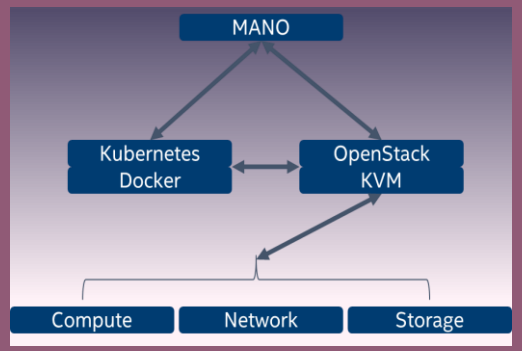

Industry challenges in containers
Bare Metal

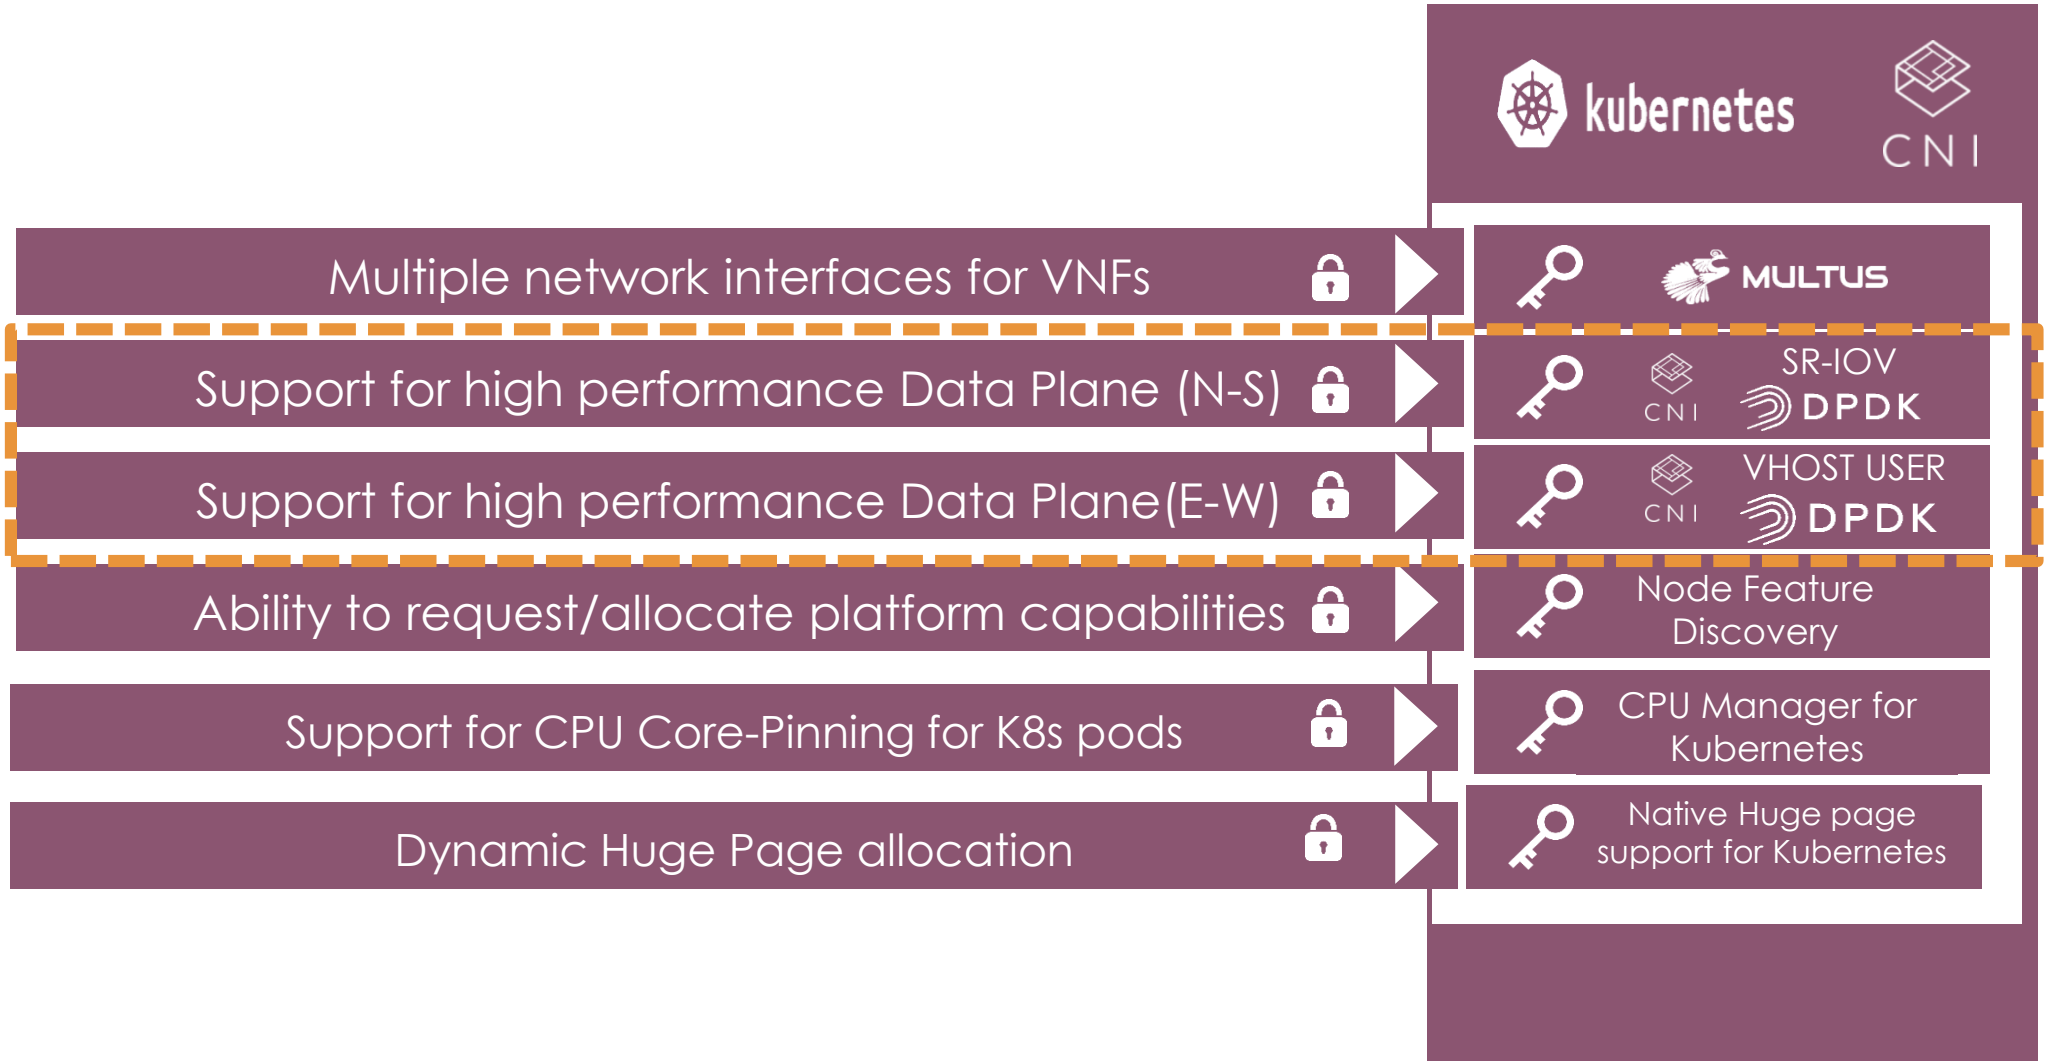

Kubernetes networks via Container Network Interface (CNI)

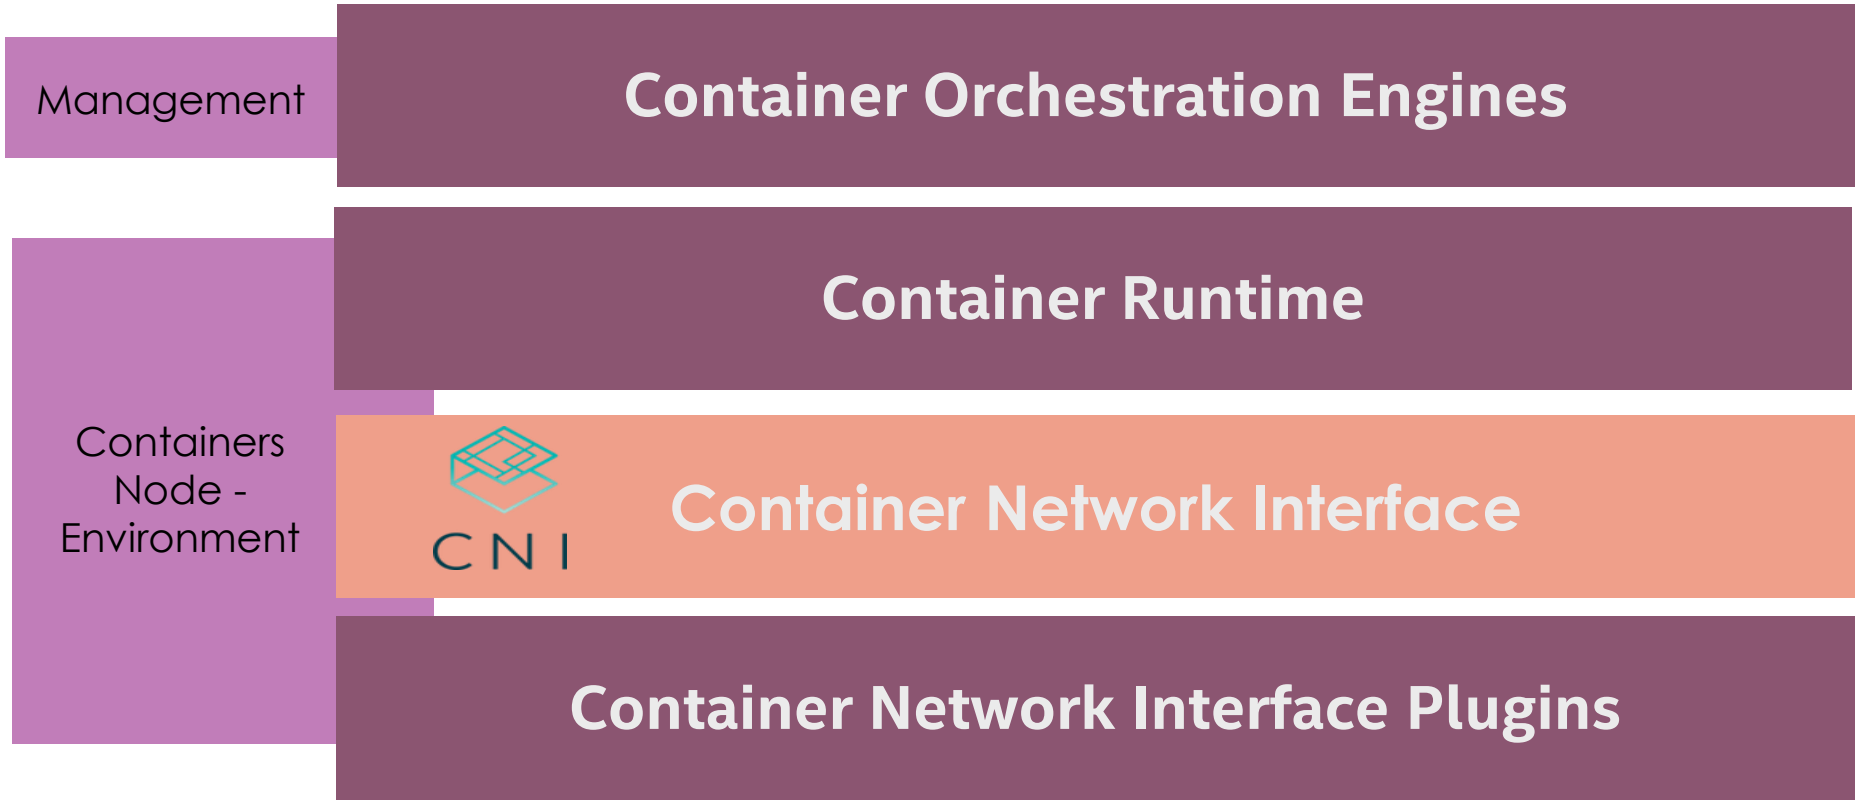

Kubernetes networks via Container Network Interface (CNI)

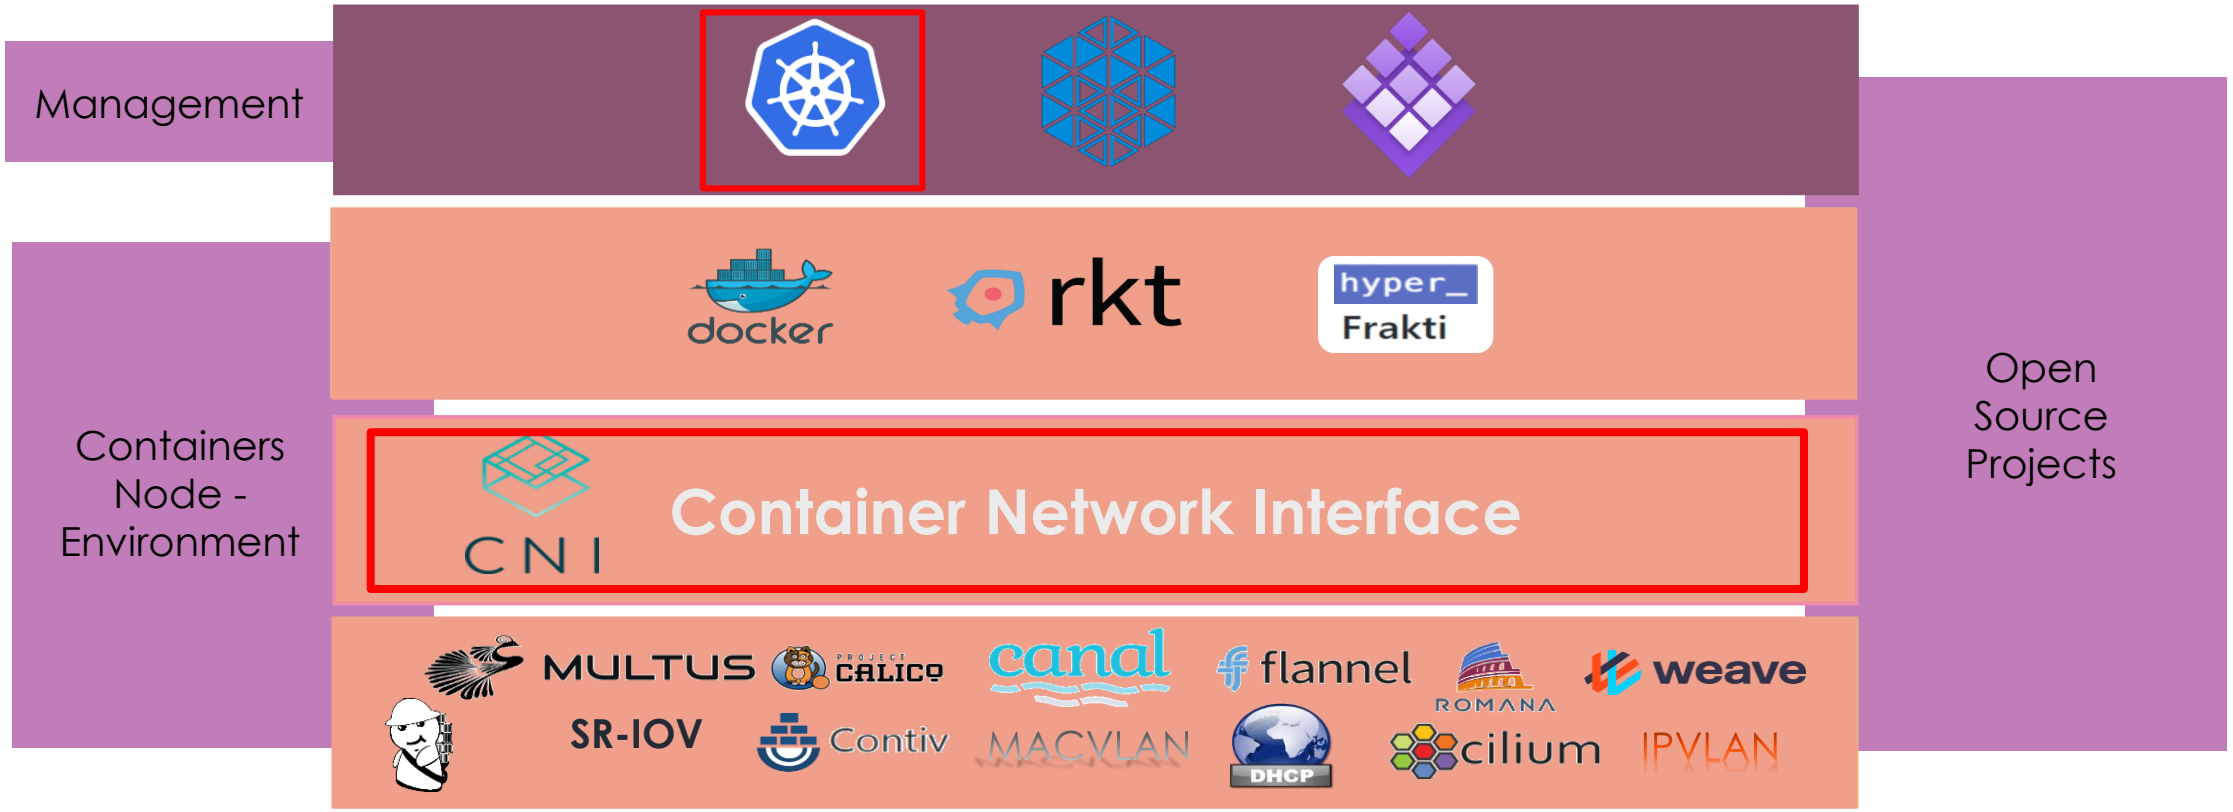

Can "mix and match" with CNI as the API

DPDK - SRIOV CNI Plugin

PROBLEM

Lack of support for physical platform resource isolation

No guaranteed network IO performance

No support for Data Plane Networking

SOLUTION

Allows SRIOV support in Kubernetes via a CNI plugin

Intel contributor and maintainer of SR-IOV CNI plugin

Supports two modes of operation:

SR-IOV : SR-IOV VFs are allocated to pod network namespace

DPDK : SR-IOV VFs are bounded to DPDK drivers in the userspace

REFERENCE

https://github.com/Intel-Corp/sriov-cni

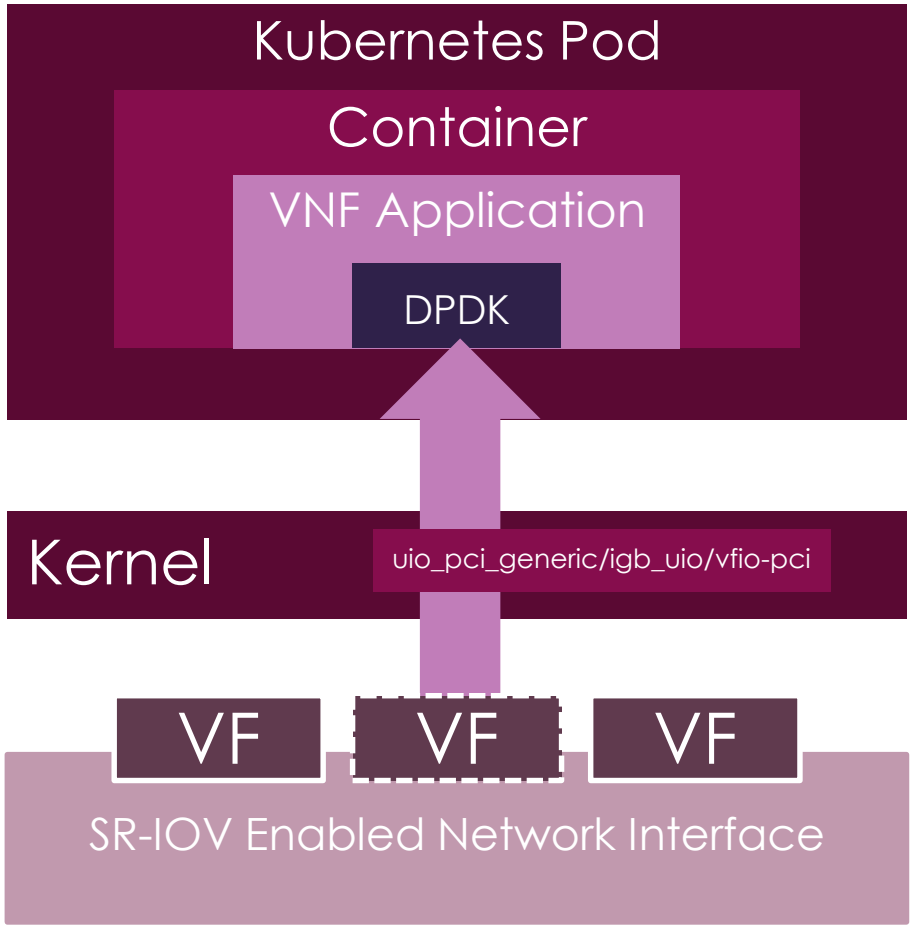

Vhost user CNI Plugin

PROBLEM

No Container Networking with software acceleration for NFV particularly for East – West Traffic

SOLUTION

Virtio_user/ vhost_user gives boosted performance than VETH pairs

Support VPP as well as DPDK OVS

Vhost_user CNI plugin enables K8s to leverage data plane acceleration

REFERENCE

https://github.com/intel/vhost-user-net-plugin

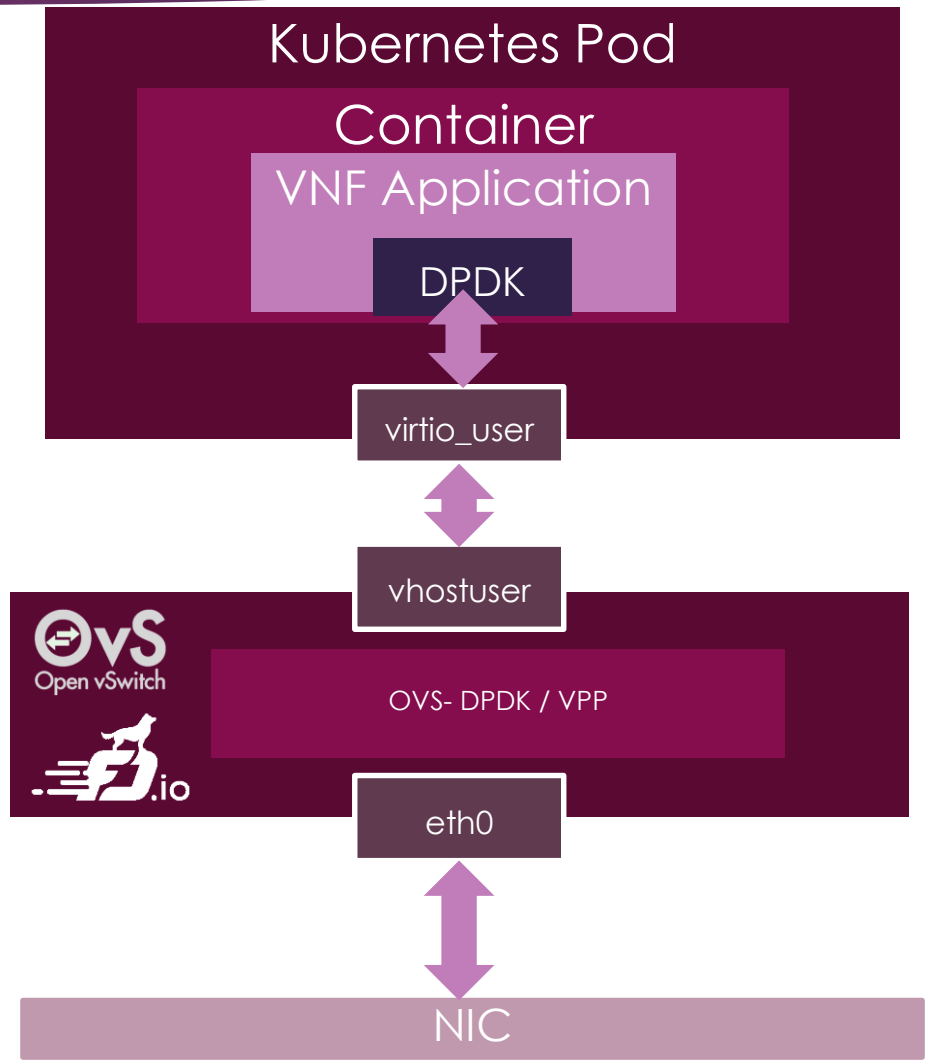

Network Cloudification
Multiple Deployment Models – Today Discussion Focus

DPDK

CLOUD NATIVE COMPUTING FOUNDATION

VNFs

vEPC	vNAT	vIMS	vRouter	vGGSN	vFirewall
vCPE	vRNC	vHLR	vSGSN	vMME	vIDS

NFVi- Network

OvS Open vSwitch | fd.io | DPDK | SR-IOV | CNI

NFV Orchestration

ONAP | Open Source MANO | kubernetes | openstack | OPEN DAYLIGHT | OVN (Open Virtual Network)

VM

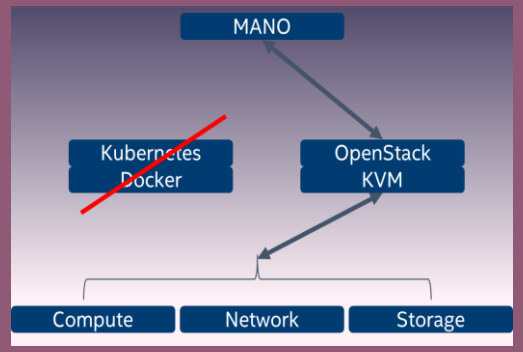

Bare metal

Containers

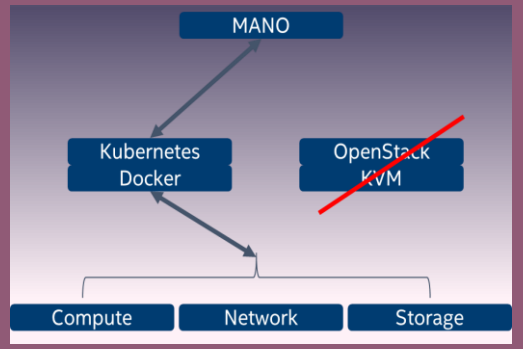

Unified

VM
Containers

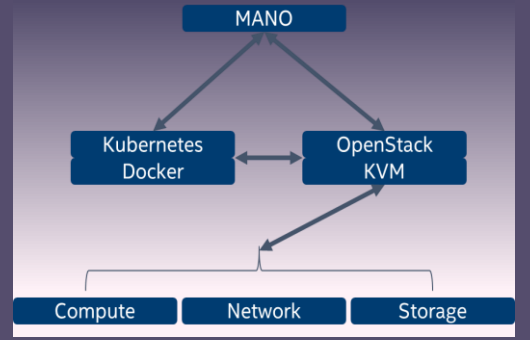

Industry challenges in Nested Containers

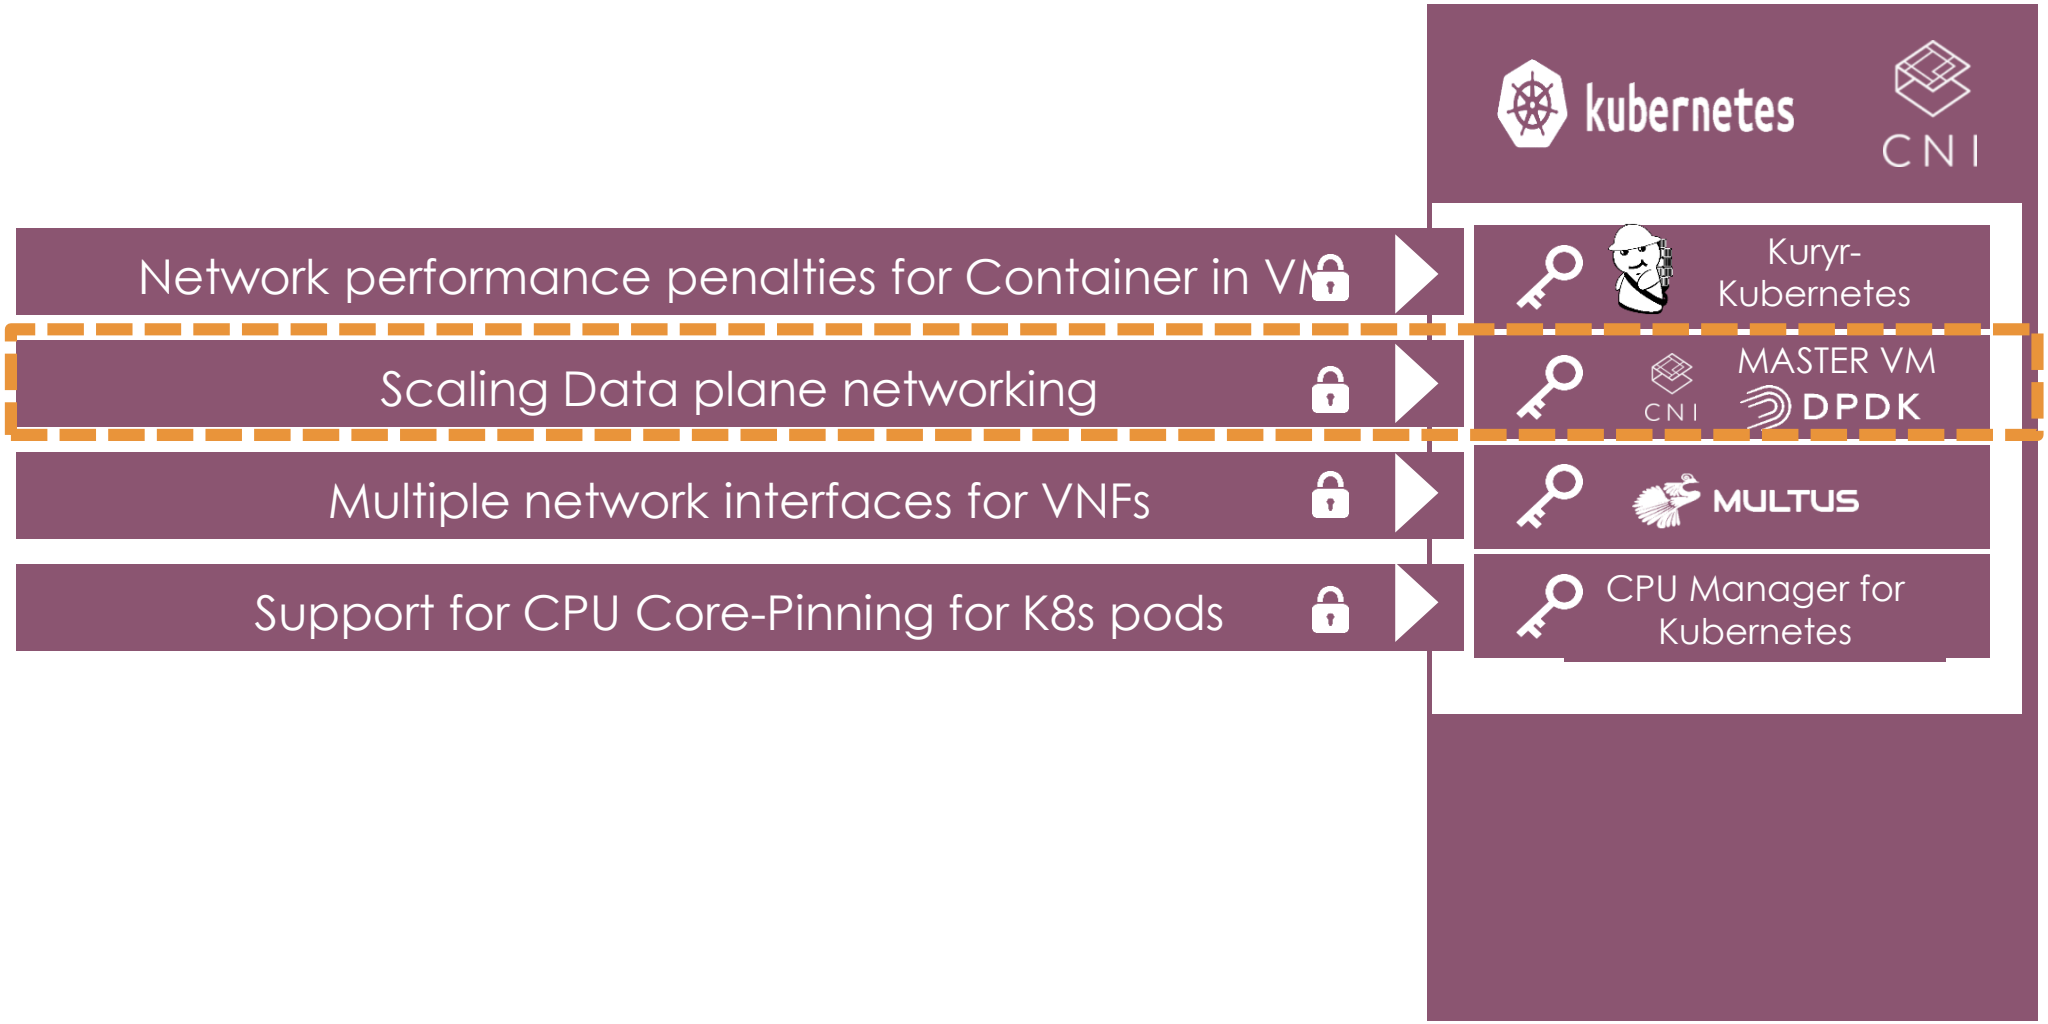

Master VM For Containers
Enabling DPDK in Nested Containers

DPDK

OBJECTIVES

One Virtual Machine to many Containers
Target: 1k Containers per VM
Container Data Plane performance

USE CASE

Elasticity and scalability of containerized VNF application in VM

BENEFITS

- VT-x ring de-privileging to move the VM and Container into userspace, making it accessible to the userspace vSwitch with just a single copy.
- Standard Virtio interface that supports both interrupt and poll modes, VNF and Cloud based applications.
- Standard Vhost shared memory interface between DPDK vSwitch and VNF.

SOLUTIONS

- Enabling DPDK in containers using VIRTIO
- Using Kuryr–Kurbernetes, orchestrate the Dataplane networking – DPDK workload using the vSwitch

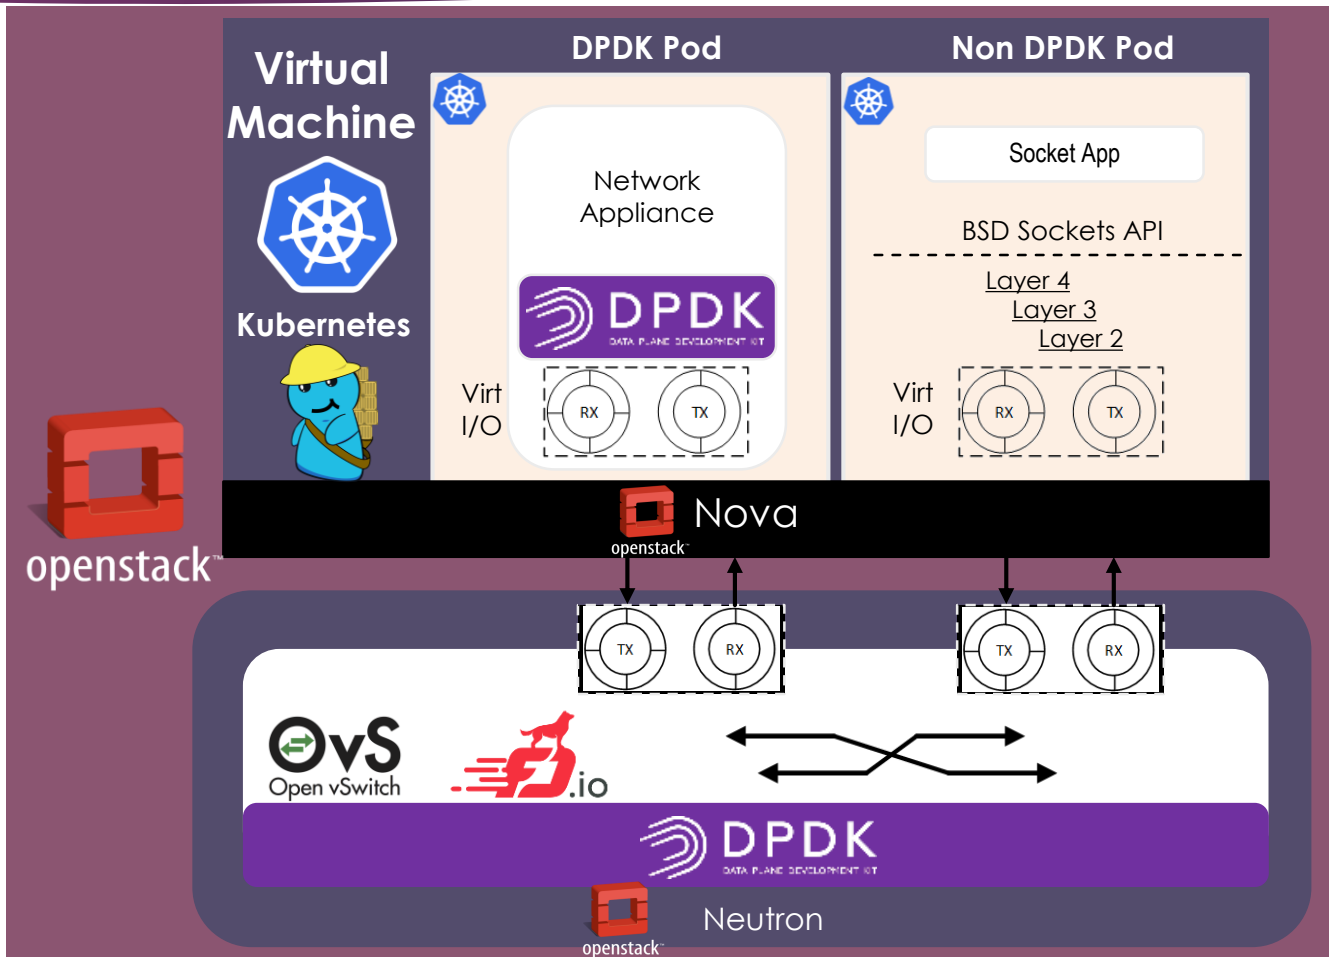

Master VM:
Co-existence of Containers and Virtual Machines
DPDK based vSwitch, independent method to accelerate the Container Data Plane.

Kuryr – Kurbernetes with Dataplane Networking with DPDK based vSwitch

PROBLEM (Current Status)

No support for Data Plane Networking for nested containers in Unified Infrastructure RA

SOLUTION (WIP)

Working on the PoC for the development of DPDK net plugin support in kuryr CNI
DPDK: Virtio are bounded to DPDK drivers in the userspace

REFERENCE

https://blueprints.launchpad.net/kuryr-kubernetes/+spec/nested-dpdk-support

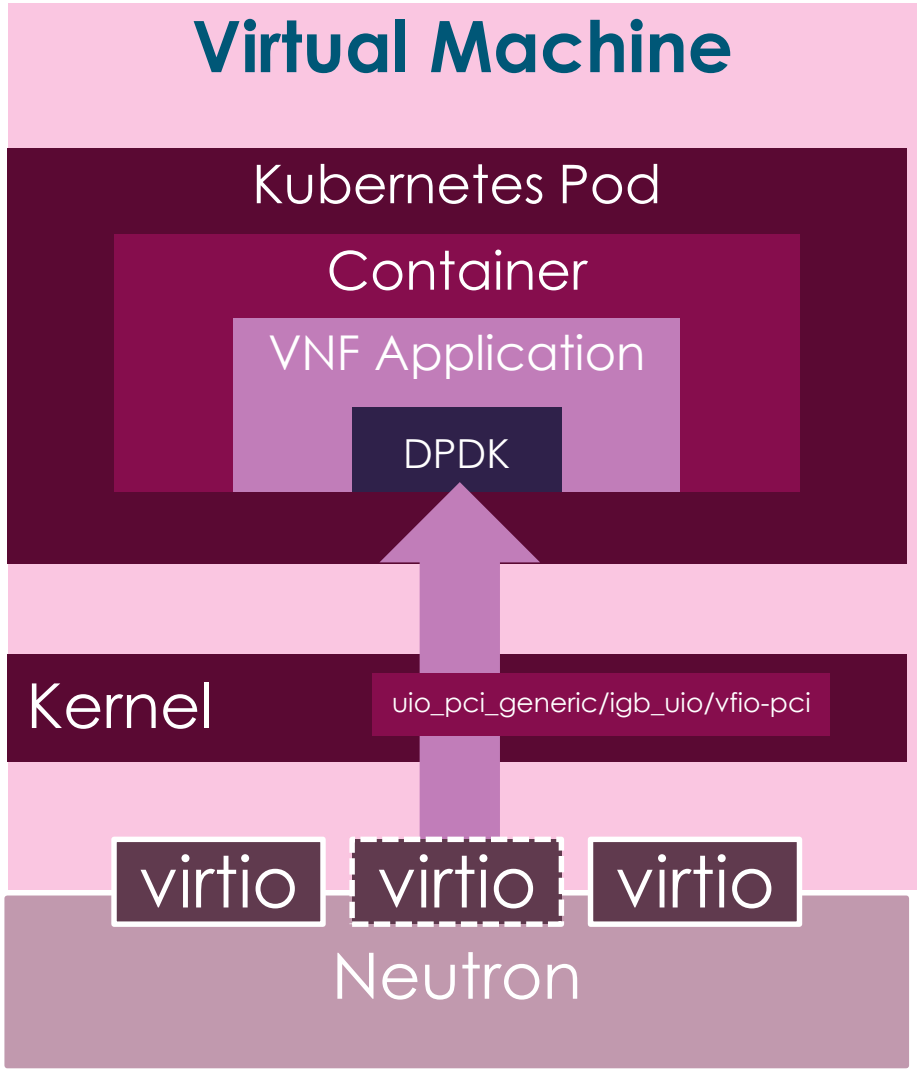

Questions?

Gary Loughnane

gary.loughnane@intel.com

Kuralamudhan Ramakrishnan

kuralamudhan.ramakrishnan@intel.com