

Sukhdev Kapur, Juniper Networks

Vikram Siwach, MobileEdgeX

April, 2019

AI/VR Applications at Edge

 MobileEdgeX

 JUNIPER
NETWORKS

 *tungstenfabric*

 AKRAINO
EDGE STACK

Blueprint Proposal: AI/VR Applications at Edge

Case Attributes	Description	Informational
Type	New Blueprint for enabling AI/ML and low latency AR/VR capabilities at the Edge	
Blueprint Family - Proposed Name	Integrated Edge Cloud	
Use Case	Programmability on switches and I/O Accelerations on programmable NICs & embedded FPGAs to deliver AI/ML workload placement and low latency demands of AR/VR applications onboarding edge stack	
Blueprint proposed Name	AI/ML and AR/VR applications at Edge	
Initial POD Cost (capex)	Leverage white boxes, standard NICs: The cost of POD will depend upon the hardware profiles and peripherals desired	
Scale & Type	Detailed in Resource Requirements slide	
Applications	AI/ML streaming workloads and AR/VR applications	
Power Restrictions	Less than 10Kw	
Infrastructure orchestration	Openstack Queens or above Docker 1.13.1 or above Container Orchestration –K8s 1.10.2 or above OS - Ubuntu 16.x, CentOS	
SDN	Tungsten Fabric, kernel/DPDK/SmartNIC offload vRouter	
Workload Type	Containers over VM or baremetal	
Additional Details	<ul style="list-style-type: none"> ● Runs on Commodity HW ● x86, ARM, SoC, ... ● Multiple options for partial or full NIC offloads (Intel, Netronome, Mellanox) ● Future Supports: ● eBPF/XDP offload 	See next slide for additional details



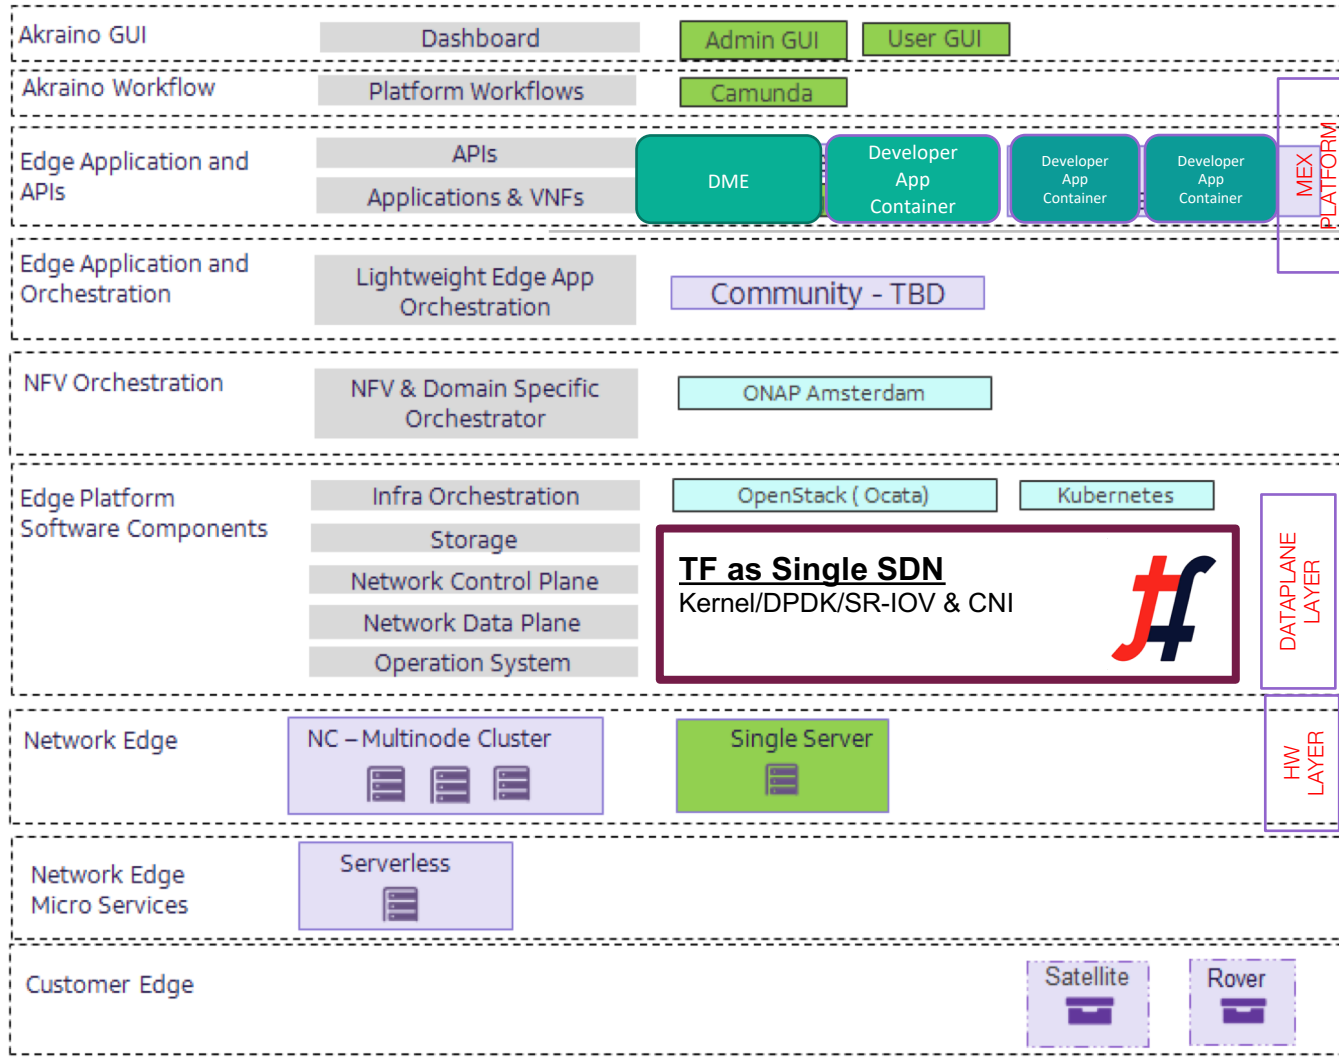
Motivations

- Programmable SW Defined Fabric to control the life of packets going to application containers.
- Low cost per byte for compute cycles, I/O and storage.
- Enable high performance on commodity hardware
- Slice Traffic to offer certain SLAs
- Embedded Security

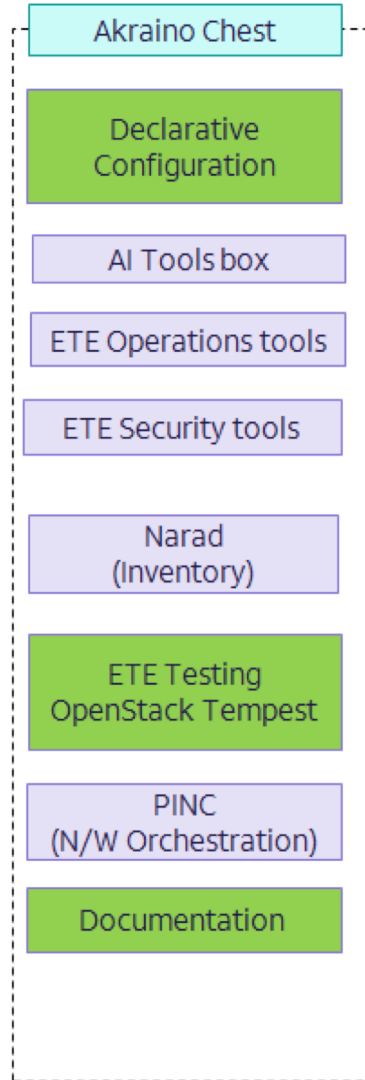
Blueprint Value

- Programmability: Programmable model for fabric | host | VM | container
 - Network slicing programming TOR for SLA contracts per tenant
- Multiple Data Plane Options: SRIOV, Hardware offload – vrouter


CICD (Community)



D a w n *



*To be contributed to Airship



MobileEdgeX Architecture

DEVICES

NETWORKS

CLOUD

Native Device Support

SDKs

APIs

Edge Services

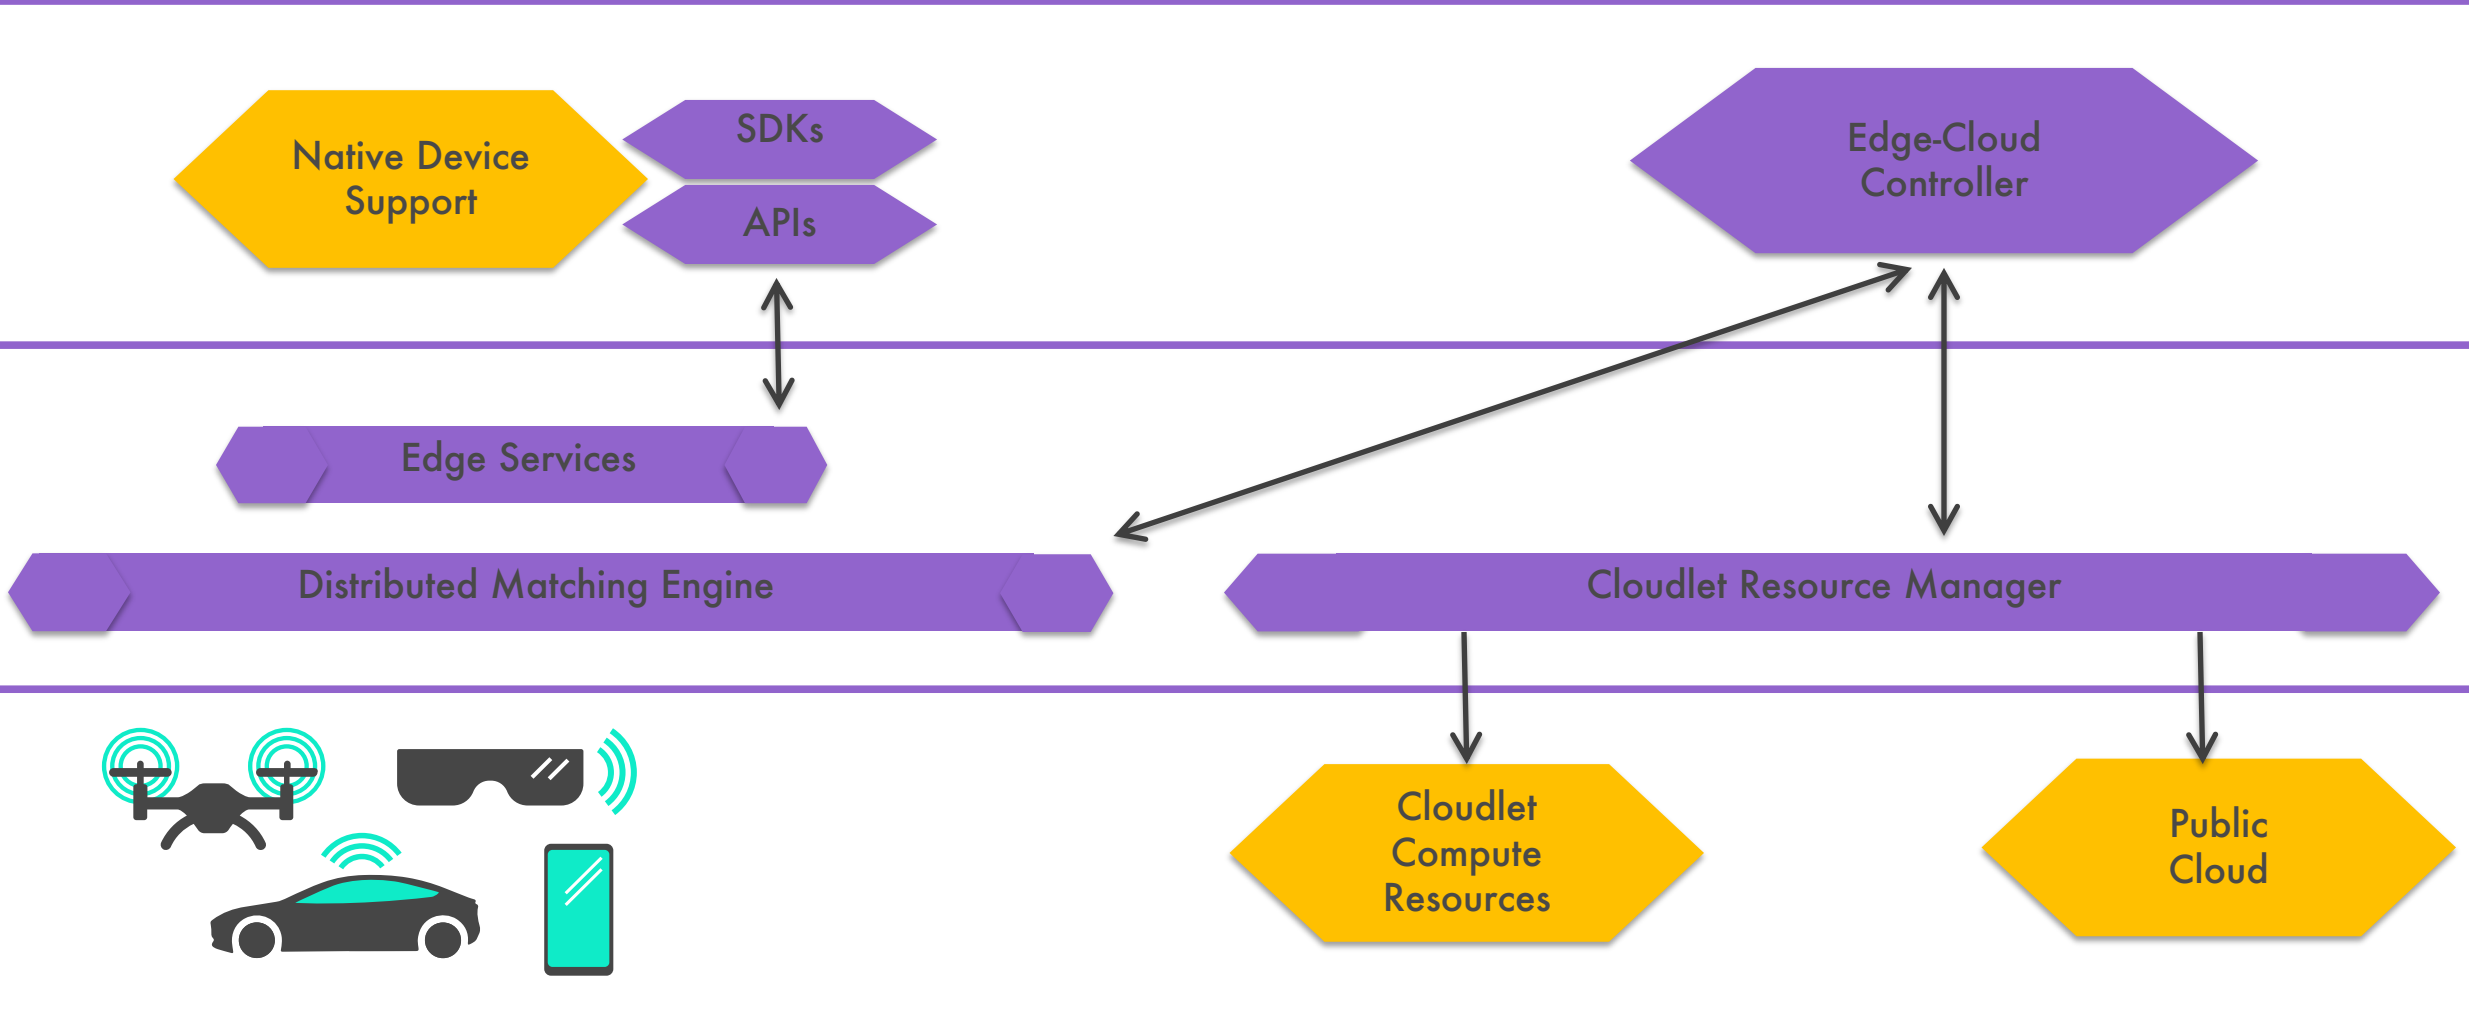

Distributed Matching Engine

Edge-Cloud Controller

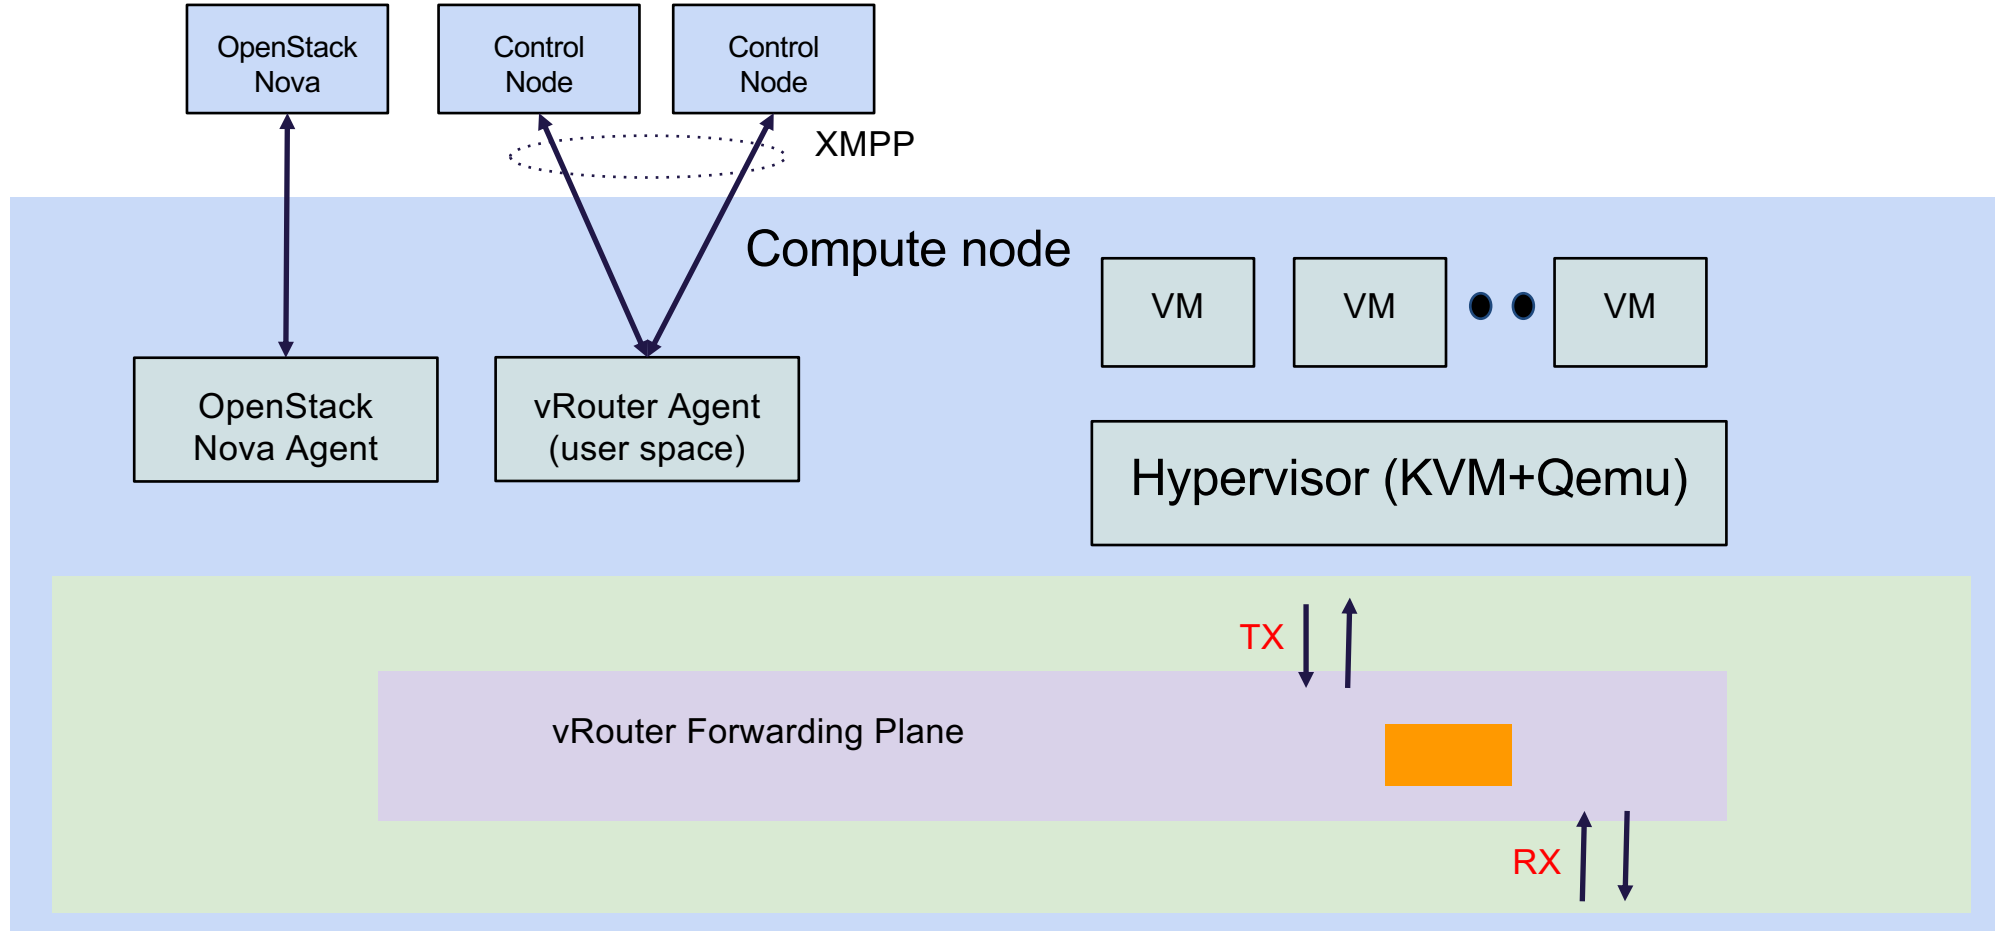
Cloudlet Resource Manager

Cloudlet Compute Resources

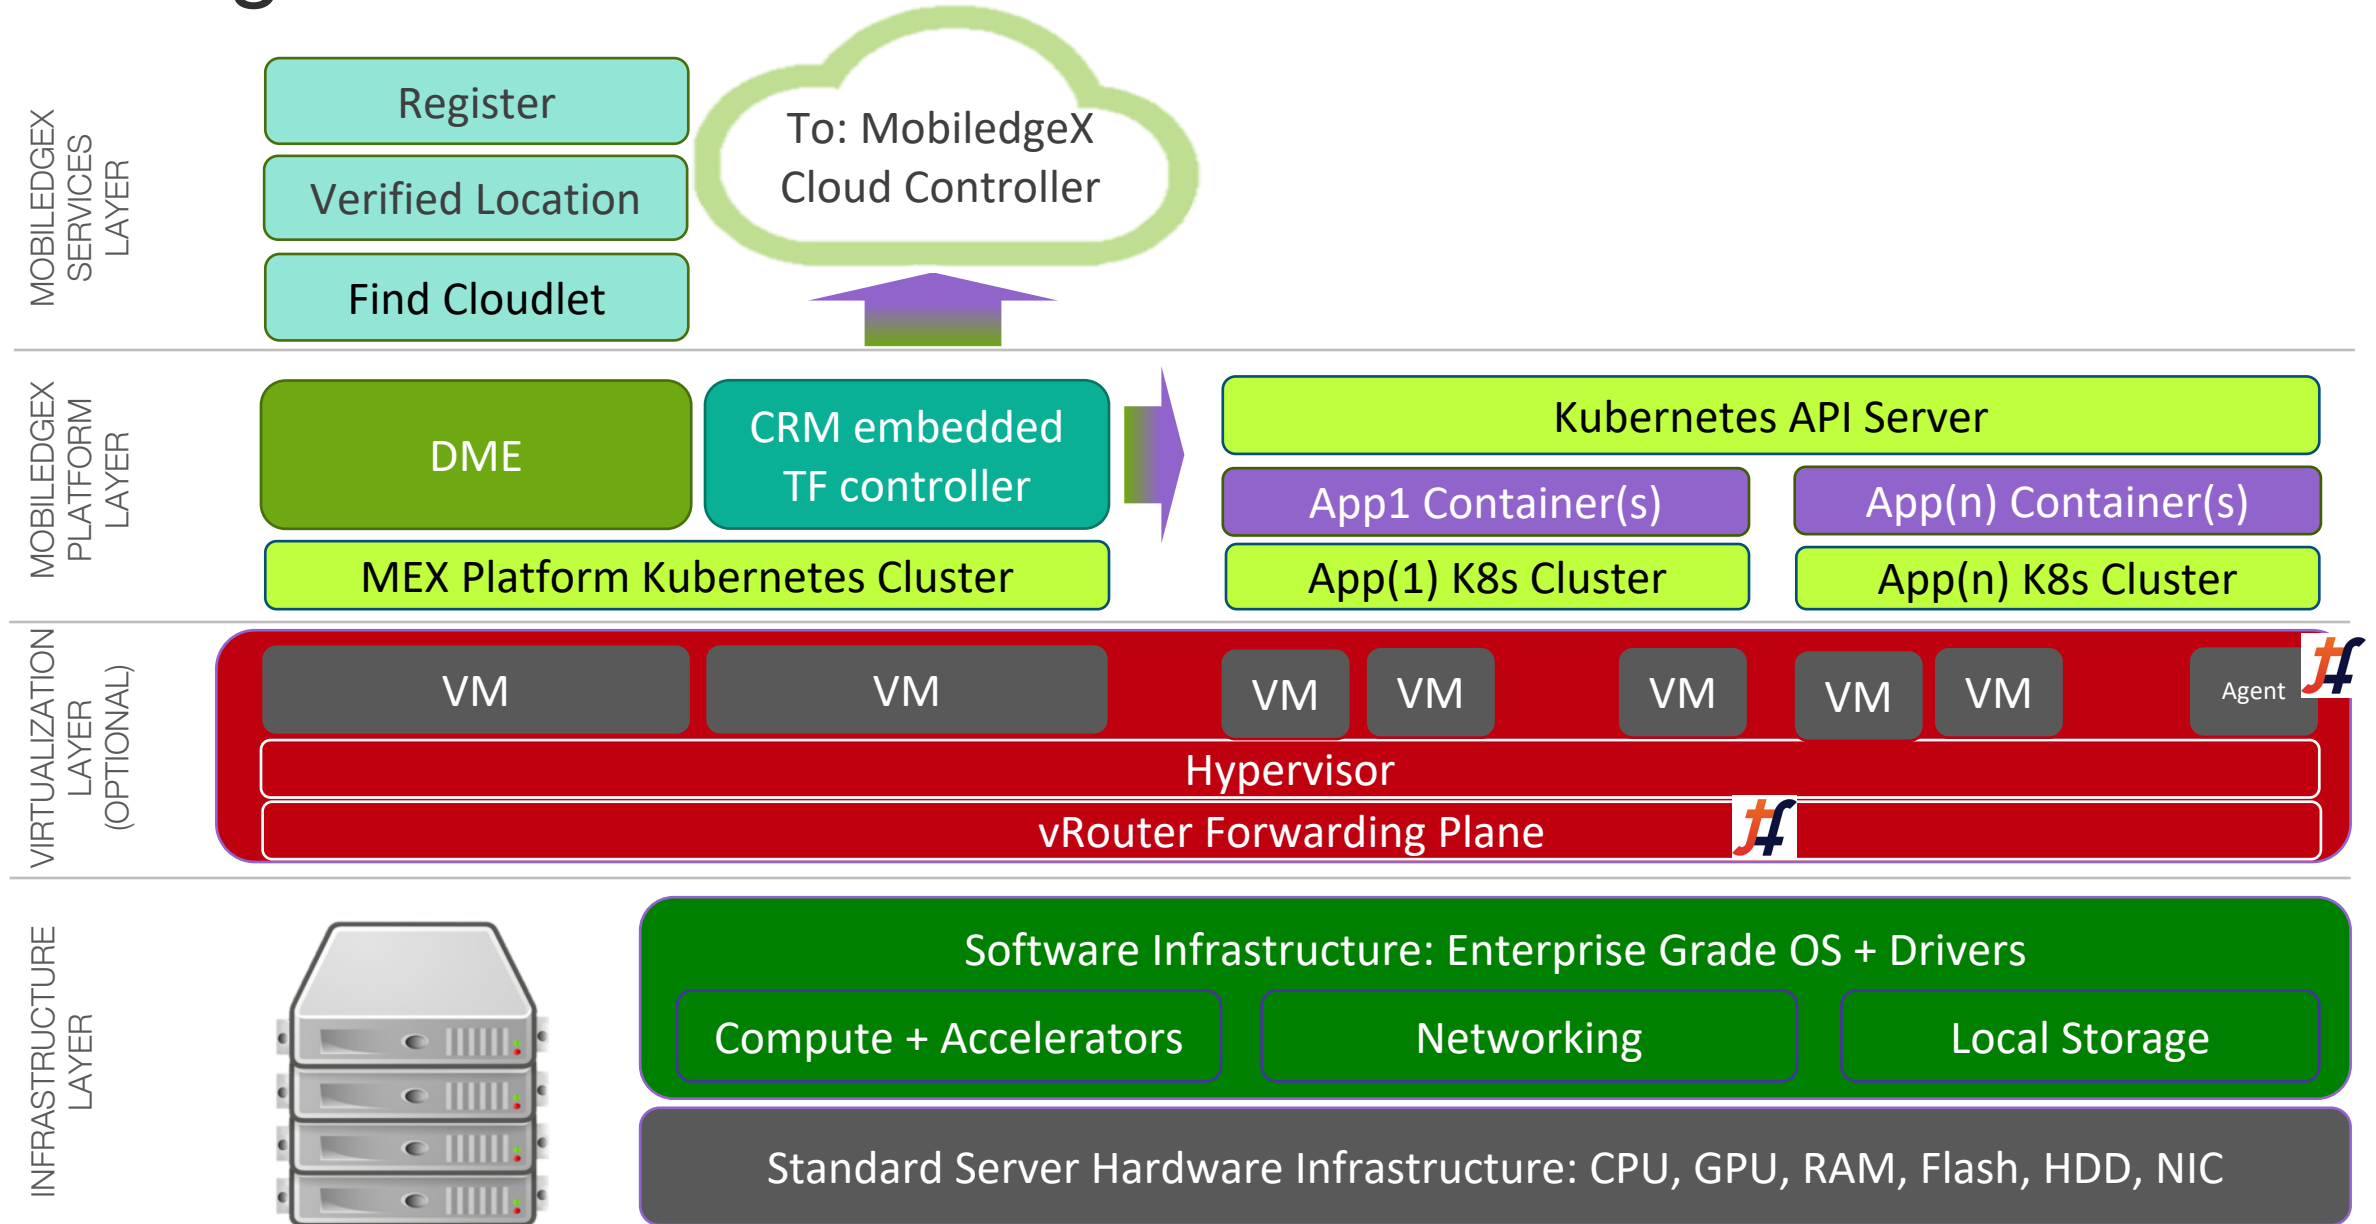
Public Cloud



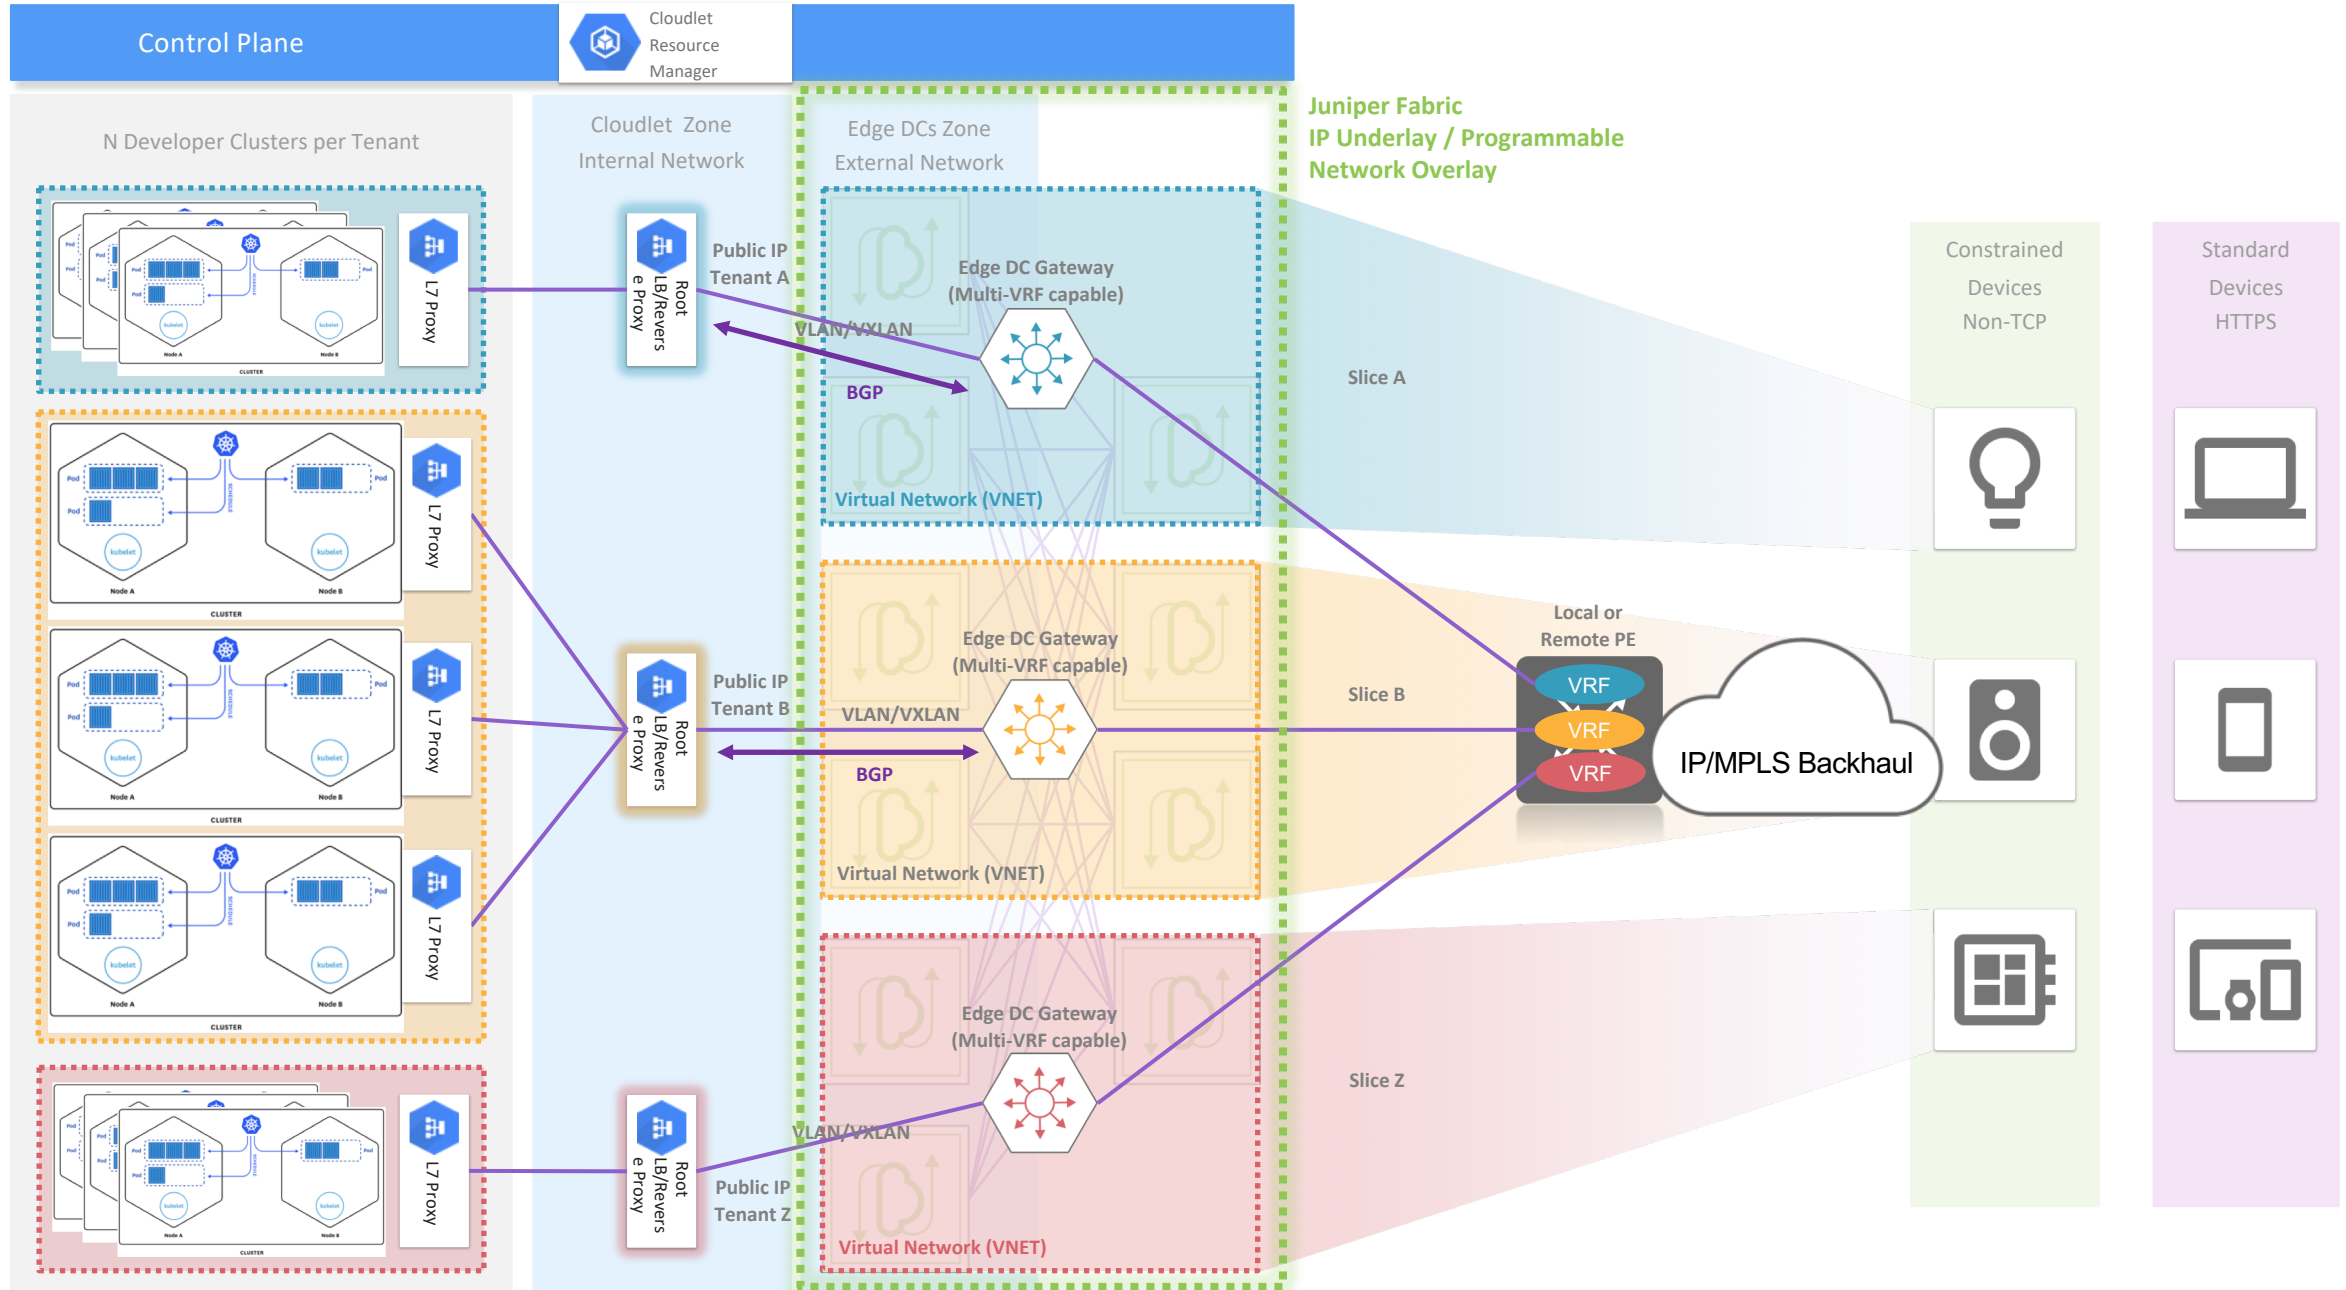
TF Architecture



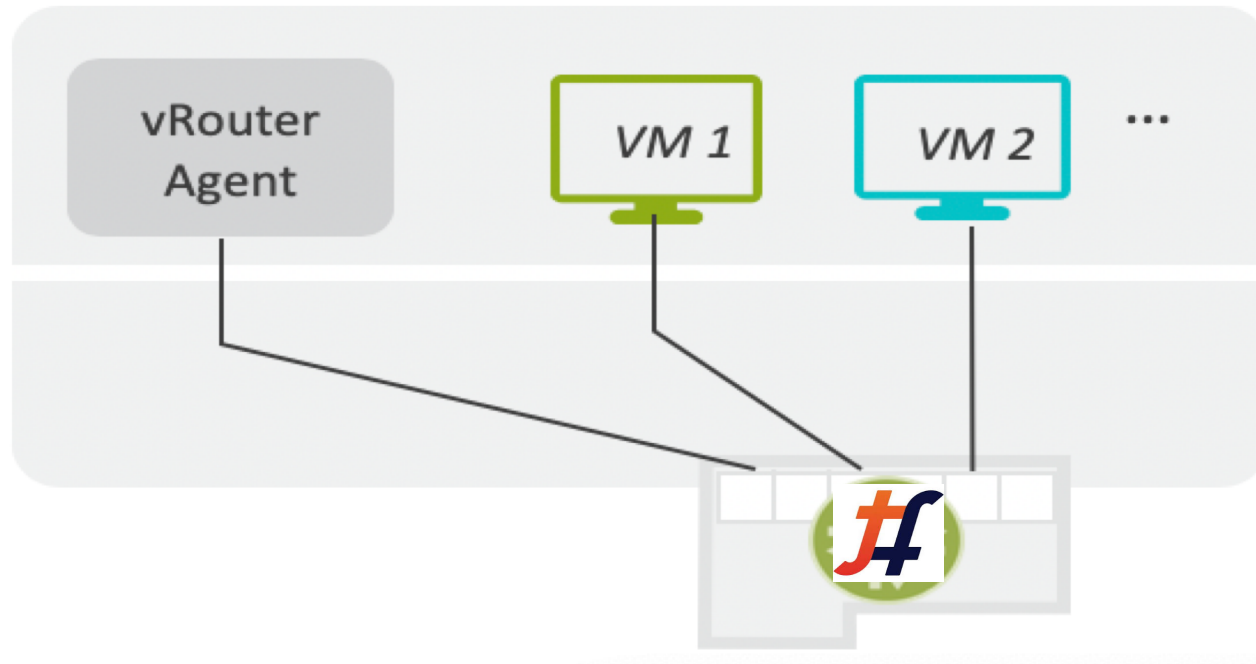
Integrated Cloudlet's Software Stack



Use case Programmability : Tenant Network Slicing – Hardware and Software



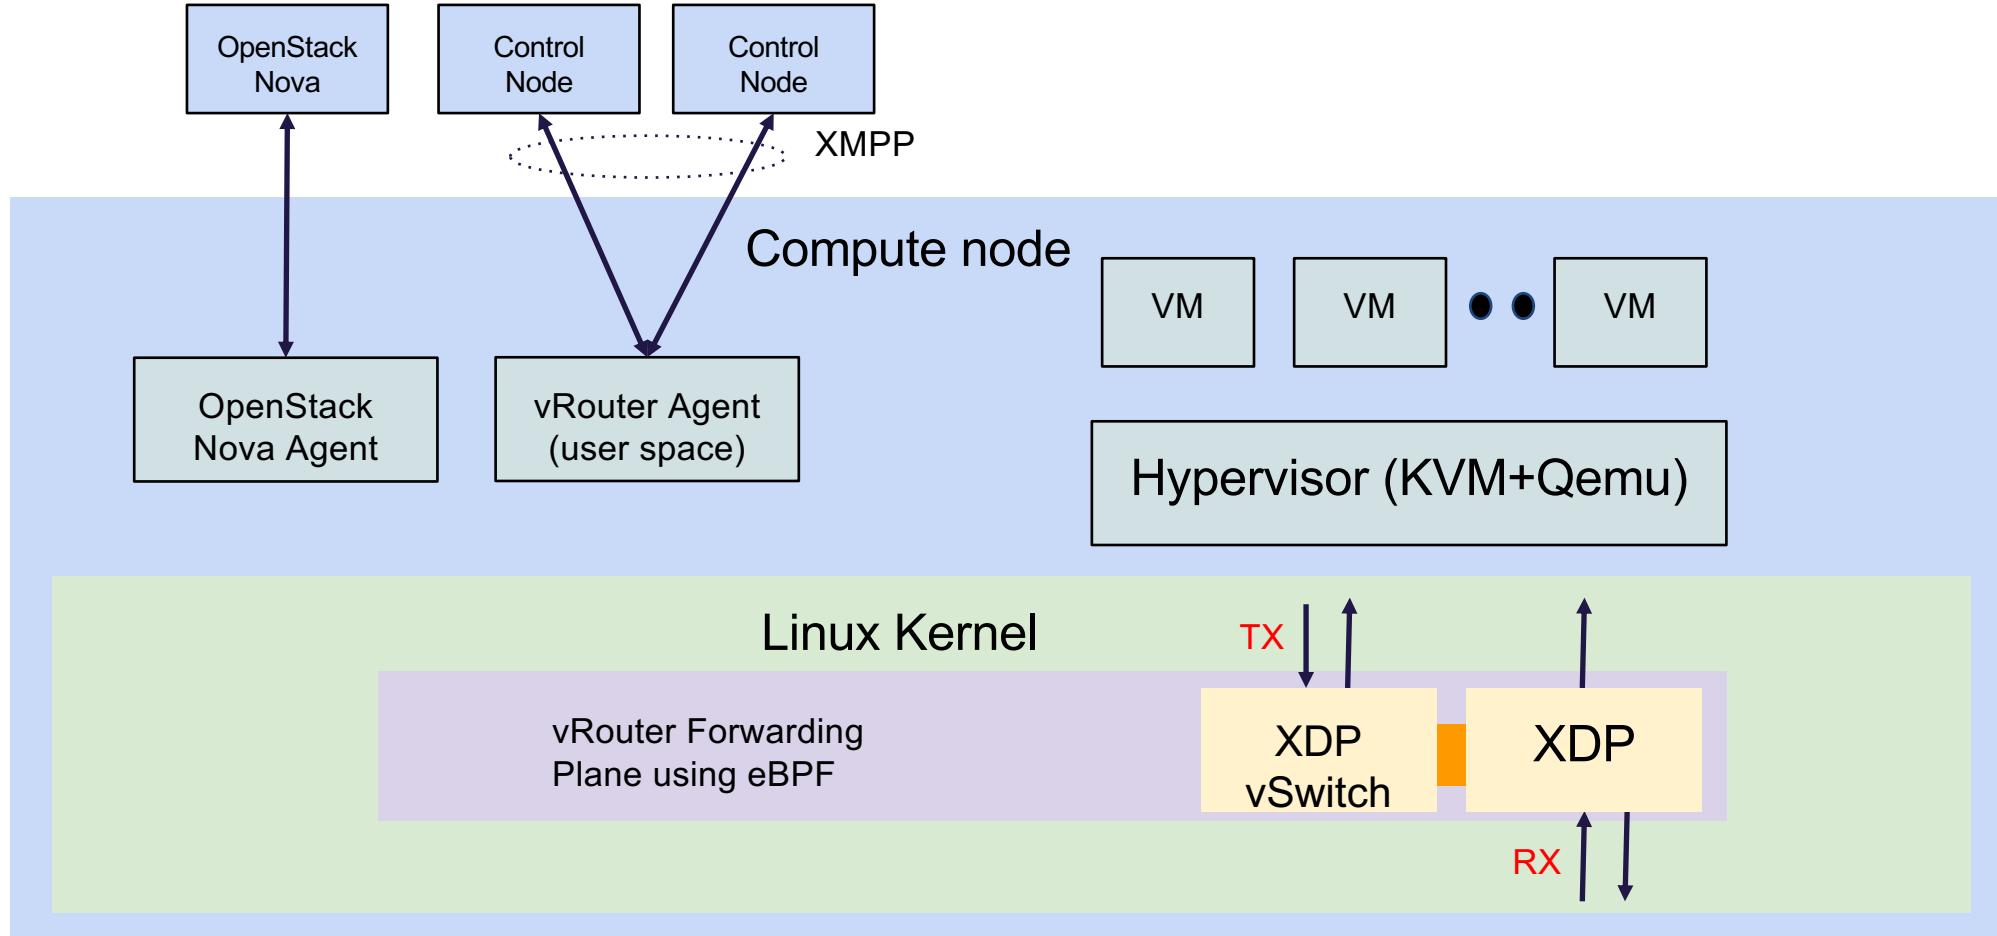
Use case : Host Acceleration using vRouter Offload



Reference:

<https://tungstenfabric.github.io/website/Tungsten-Fabric-Architecture.html>

Future support: ebpf based XDP offload



Lifecycle Management for MEX applications

- Configuration and deployment: Cloudlet resource manager is single point of configuration and is capable to end to end Cloudlet provisioning from apps to connectivity to external world. It is the orchestrator for Kubernetes cluster and containers based on developers needs. Tungsten Fabric controller is embedded within CRM
- Monitoring: MEX dashboard supports native Kubernetes monitoring solution Prometheus which collects key stats and dash board visualizes the metric thus collected.
- Scaling: CRM launches cluster and containers within based on declarative configuration by developer but elasticity to scale up and down based on stress is built in its logic.

Resource Requirements for MEX cluster

Minimal Cloudlet Resource Requirements Scenario

	VMs	vCPU	RAM GB	Disk GB
MEX Platform	2	8	16	160
MEX Agent	4	16	32	320
Cluster Instance Node	3	12	24	240
vRouter		2	8	
Total	9	38	80	720

8-cluster Cloudlet Resource Requirements Scenario

	VMs	vCPU	RAM GB	Disk GB
MEX Platform	2	8	16	160
MEX Agent	4	16	32	320
Cluster Instance Node	24	96	192	1920
vRouter		2	8	
Total	30	122	248	2400

Provisioning Requirements for MEX

- Access: VPN and Openstack
- Resource Allocation for MEX Applications
 - Compute & 3-6 VMs per cloudlet required for MEX
 - MEX Agent VM (Minimum 1, preferred 4)
 - Ingress Load Balancing
 - K8S orchestration
 - MEX Platform VM (2 VMs)
 - Distributed Matching Engine (DME)
 - API traffic from SDK
 - Communication to other services
 - Cloudlet Resource Manager (CRM)
 - OpenStack management
 - Receives commands from Controller node to orchestrate cluster and app instances
 - One Kubernetes Cluster per Cluster Instance for Applications
 - 3 VMs (vCPU: 4, RAM: 8GB, Disk: 80GB)
 - Applicable to:
 - MEX Agent (1-4 VMs per cloudlet)
 - MEX Platform (1-2 VMs per cloudlet)
 - K8S Node (3 VMs per cluster instance)
 - Server Hardware
 - Minimum Requirement: Dual Socket E5 Xeon Processor, 256GB Memory, 1TB of SSD
 - Prefer NVIDIA Tesla P100 GPUs
- Network
 - External Network
 - Created in advance
 - Needs access to OpenStack API endpoint
 - 1 IP for each MEX Agent – L7 LB
 - 1 IP for MEX Platform VM (DME, CRM)
 - 1 IP for each application requiring L4 access
 - DHCP assigned public IP range on ext network is preferred
 - For short term, a pool of 8 public IPs is adequate
 - MEX K8s Network
 - Created by MobileEdgeX via APIs
 - Shared by all cluster instances
 - Subnets created dynamically by MEX
- Storage
 - Local storage, Block Storage, Object Storage, Shared File System
- Openstack Cloudlet API Requirements:
 - server list, create, delete, set properties
 - image list, save, create, delete
 - network list, create, delete
 - subnet list, create, delete
 - router create, delete, add and delete ports
 - flavor list, show create
 - security group rule list and create
 - show limits