

# Enhancing K8s Networking with SmartNICs

Dave Cremins, Intel

## Agenda



- Edge, Bare Metal & Kubernetes
- Kubernetes Networking & Trends
- SmartNICs
- Disaggregation
- More offloads

#### What is the Edge?

KubeCon CloudNativeCon Virtual

The move to 5G is driving changes in edge computing and adoption of 5G solutions with the promise of:

- Lower latency
- Higher capacity
- Increased bandwidth



## Edge & Bare Metal



#### Why bare metal deployment model for the Edge?

- Greenfield applications are built as microservices, packaged as containers and orchestrated
- Even the Telco industry are adopting these patterns
  - Leveraging the edge to take advantage of the 5G promise
- Edge deployments require small footprints to be successful:
  - Does not have the space of a modern datacenter i.e. resource constrained
  - Little tolerance for extra server capacity for software overheads like virtualization i.e. avoid the tax
  - Lead time for new hardware is too slow i.e. get it right the first time
- Some Edge architectures will consist of various deployment platforms, bare metal, VMs, containers, etc.
- Bare Metal platforms allow full access to hardware:
  - Reduction in the amount of resources needed (note scale up vs scale out)
  - Options to leverage accelerators e.g. FPGAs, QAT, etc.
  - Higher throughput, lower latency and superior performance (note aligns with 5G promise & edge deployment reqs)
  - Agility
    - Dependencies minimized
    - More aligned with dynamic aspects Swappable components
    - Facilitates adapting to infrastructural requirement changes with little side-effects on applications
- Even with a small footprint, the bare metal edge can deliver on the speed & performance required by today's connected world





#### Why Kubernetes for the Edge?

- Forecasters predicting increases in edge computing
  - Sheer quantity of instances compared to centralized cloud centers
  - Edge workload opportunities continue to grow and will become more demanding
  - Infrastructure and platform resources will need careful management to honor these workloads
  - K8s will play a very important role in this space
- Inherent to Kubernetes is the ability to abstract the infrastructural capabilities while still being robust and scalable
  - Single control plane with standardized deployments => Allows us to easily target the Edge
  - Same programming paradigm applies to Edge as it does to Cloud
  - Code can be developed and tested in the Cloud, then seamlessly deployed at the Edge
  - Developers can take advantage of the vast array of DevOps tools
  - Control the pace of innovation
- However, we need to start thinking differently about how we leverage edge platforms?
  - Boundaries & Separation of concerns Disaggregation & Distribution
  - Optimal resource management Get more out of my platform
  - Observability Collect, process, learn, optimize, etc.

## **Kubernetes Networking**



- The Kubernetes networking model has some simple requirements:
  - All Pods/Nodes can communicate with all Pods without NAT
  - The IP that a Pod sees itself as is the same IP that others see it as
- Advanced networking models with complex properties are deployed in K8s clusters all the time:
  - Tunnels & Overlays
  - Side cars & Service Meshes
  - IPSec & Zero Trust
  - Data plane technologies e.g. DPDK, SRIOV, etc.
  - Multiple network interfaces
- These properties require platform resources to deliver on their respective claims
- Let's take an example: ovn-kubernetes

#### **Kubernetes Networking**



**CloudNativeCon** 

North America 2020

KubeCon

nal

Traffic flow I - IV Kubernetes pod creation flow -----> Web API call 1 - 4 Network settings generation flow A - C Network settings application flow Non-web action ----- Route

- Host visible interface
  - —• Virtual interface

#### **Kubernetes Networking**



- Change is inevitable, networking architectures will continue to evolve and become even more complex
  - It is already happening e.g. Network Observability, AI, Service Assurance, etc.
- We need to be careful regarding disproportionate components



- Why utilise your platform resources for infrastructure "boilerplate" if you don't have to?
  - Embrace disaggregation & leverage hardware offloads

## **SmartNICs**



- SmartNIC
  - Network attached acceleration platform
  - A processing environment
  - A Network Pipeline
  - Programmable data plane
  - A location to run infrastructure management components
  - A real option to guarantee integrity of the network – Hardware root of trust
- Delicate balance between flexibility and efficiency
  - Both are required
  - Performance, security & offloads



# **Disaggregation at the Edge**





- Bin pack as many applications as possible
- Provide more predictable models for deterministic performance
- How? Leverage the side car pattern and apply it to the infrastructure
- What can we yield if we migrate these concerns?
  - Network flows can be programmed and offloaded OVS TC & RTE\_FLOW
  - Traffic can be forwarded between PFs & VFs without going through software
  - Inline processing stays in the NIC and lookaside can directly transmit to the target No data movement back/forth with host
  - Programmable security policies completely managed by the SmartNIC ACLs, Network Policies
  - eBPF offloads Filtering, LB, Monitoring, etc.
  - Observability pipelines deployed and managed on the SmartNIC
- Important This shift is not new
  - Follows the same patterns that have been applied across Data Centres & Clouds for the last few years
  - Facilitate advancements in networking without the cost of extra CPU cycles Perfect fit for edge scenarios

## **More Disaggregation**



- SmartNICs provide programmable solutions
  - Can enhance our overall networking model while Kubernetes continues to orchestrate business value
- What other offloads can we leverage?
  - FPGA
  - QAT
  - GPU
  - VPU

