

# Introducing the DPU Operator

Vendor-Agnostic DPU Integration in K8s for Performance and Security

Vijay Ram Inavolu  
Senior Principal Engineer



William Zhao  
Principal Software Engineer



# Agenda

- Shifting Gears: Next Offloads
- Limitations in DPU Aware Orchestration
- Enterprise use cases to be Unlocked
- Solution Goals & Overview
- Context: Openshift & Microshift
- DPU Operator: A Detailed Overview
- An Example Use Cases: IPSec Offload, Load Balancing
- Roadmap & Availability

# Infra Offload → Network Functions Offload

- A substantial portion of host computing resources are used for Infra.
- Next DPU Opportunity: Accelerating Specialized Network Functions
- Many NFs are already available as ready-to-use VNFs and CNFs.
- Most of these functions can significantly benefit from hardware acceleration.
- They can also be deployed as shared resources along the data path.

# DPU-Aware Orchestration: Gaps and Limitations

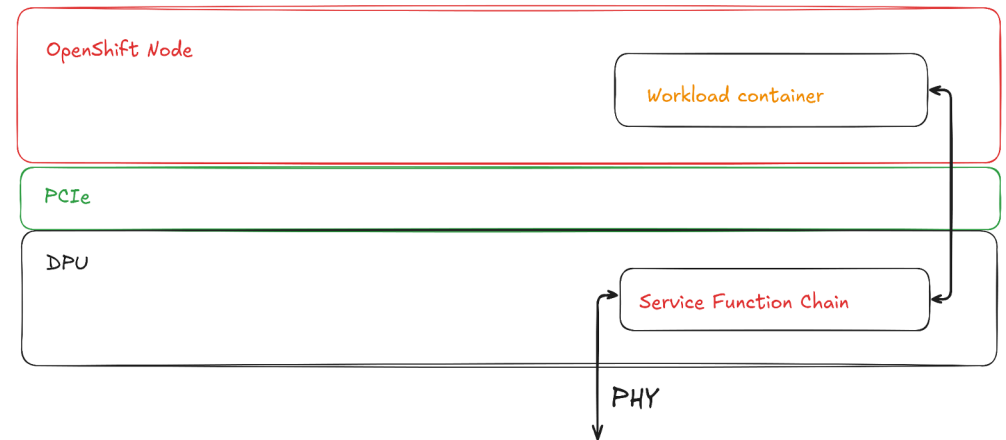
- Lack of DPU awareness, at best it's a node
- There is no standardised method to deploy NFs on DPUs.
- Solutions for traffic steering to NFs are not available.
- Proprietary solutions, vendor lock-in's can't work at scale.
- No concept of NF's plug and play

# Enterprise Use Cases to be Unlocked

Domain	Acceleratable NFs	Candidate Implementations
Load Balancing	Traffic Management, TLS termination	F5 Big-IP, NGINX
Security	Firewalls, VPNs, IDS/IPS	Palo Alto VM-Series, Suricata
Telecom	vEPC, vRAN, Edge Compute	All most all these are deployed today as VNFs or CNFs
Network Services	CDN, WAN Optimization	Cloudflare, Riverbed, Kong

# Solution Goals and Overview

- A Secondary Network
  - Fully programmable
  - Managed independently
- CNFs
  - Package NFs as pods ( CNFs)
  - Deploy using standard mechanism
- Connectivity
  - Service Function Chaining
- Resource Management
  - Discovery
  - Life Cycle Management
- Acceleration
  - DPU vendor neutral



# Red Hat Openshift

- OpenShift is an Open Container platform built on Kubernetes.
- Vibrant ecosystem of partner solutions and partner hardware
- OpenShift on bare metal can scale with the number of baremetal nodes
- Allows for our customers to deploy their containerized or virtualized workload within this platform
- OpenShift will manage the lifecycle of the workloads taking the load off customer's devops engineers
- DPUs to offload workloads from OpenShift nodes

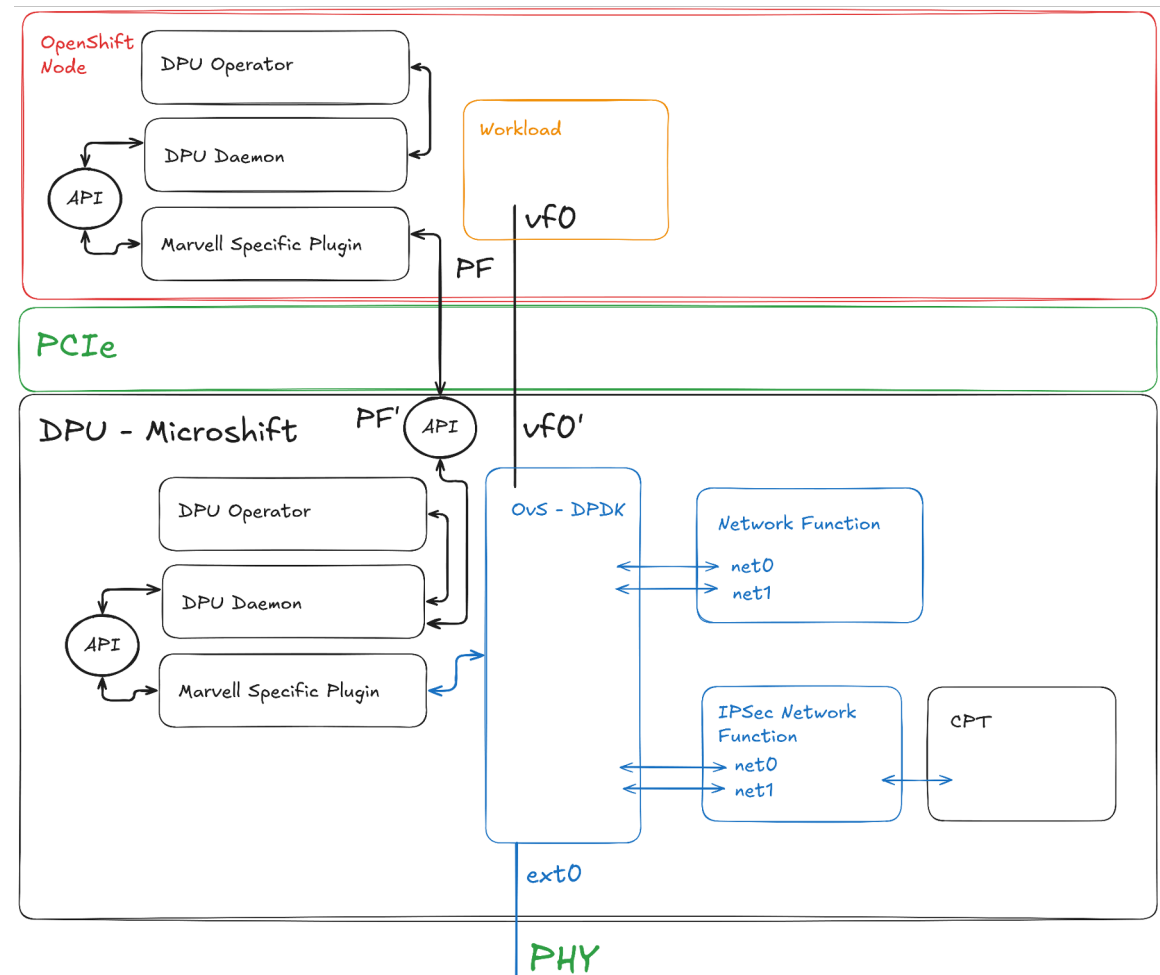
# Red Hat Microshift

- Provides essential container orchestration runtime capability that is built for Internet-of-things (IoT) and Edge computing scenarios that are both CPU and memory constrained
- Container orchestration runtime is binary compatible with OpenShift Container Platform, but it is not 100% API resource compatible.
- Kubernetes conformant
- More control on configuring everything you need manually as an operating system.



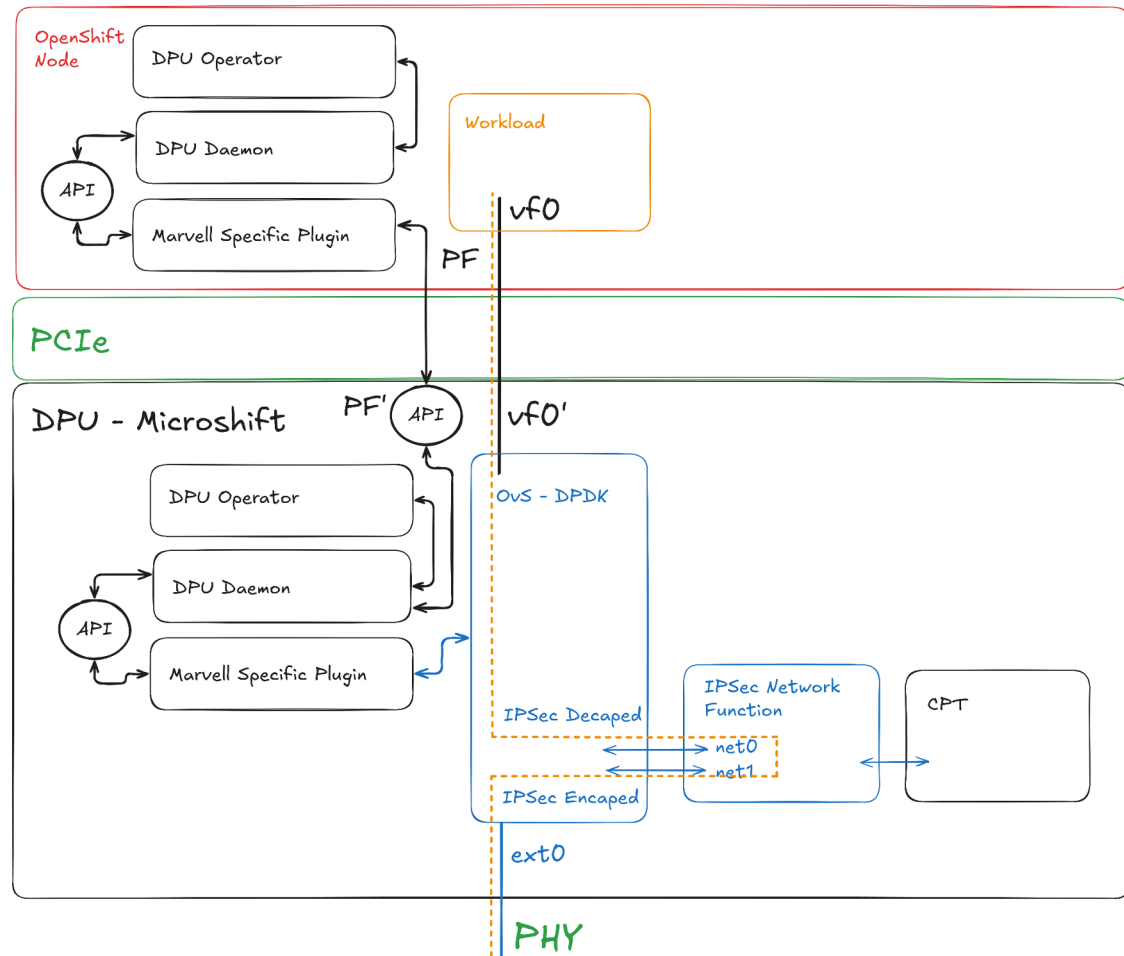
# Detailed overview of the DPU Operator

- Network Functions are part of the Service Function Chain (SFC)
- Marvell Octeon 10 design utilizes OvS DPDK with hardware accelerators to achieve performance and security.
- gRPC APIs across components makes the solution vendor agnostic with the goal of OPI compatibility



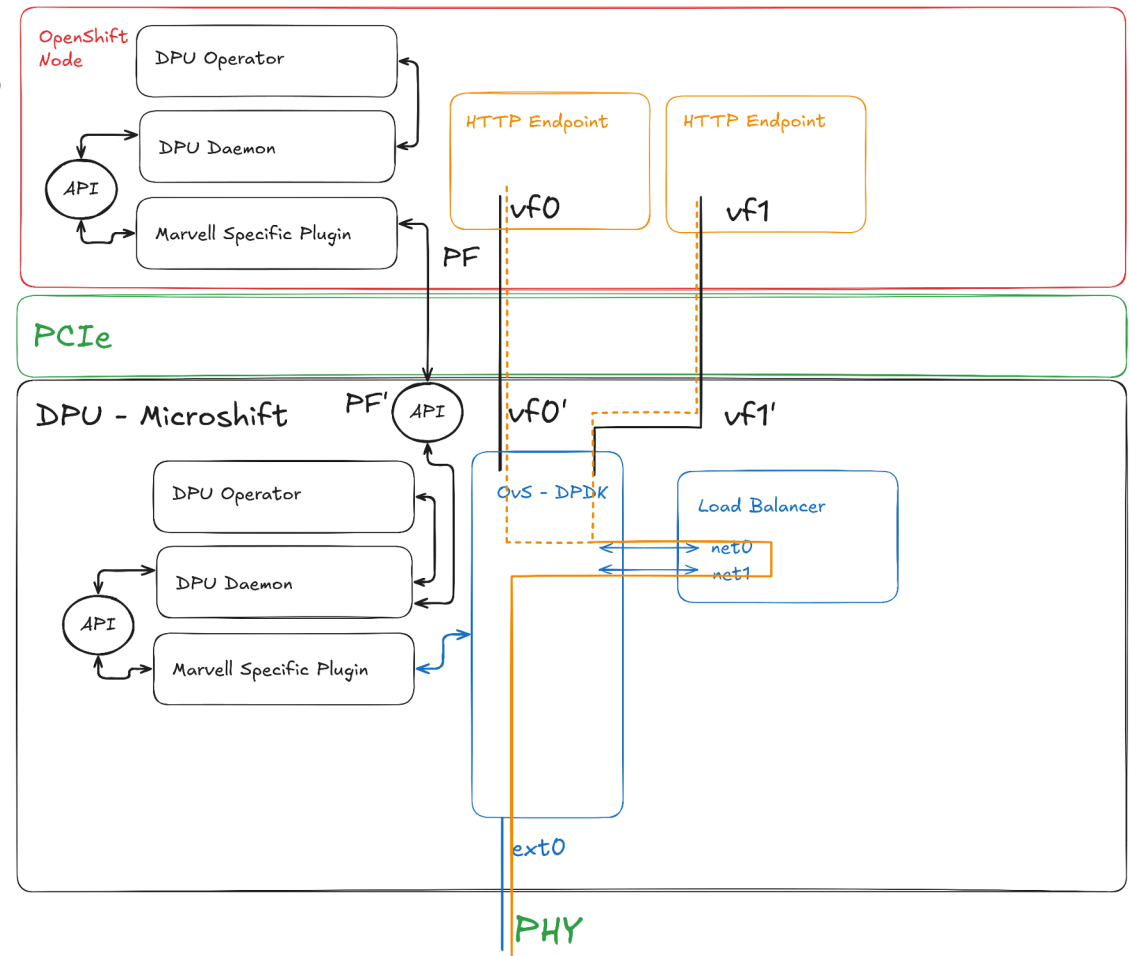
# Example Use case: IPSec Offloading

- Line rate IPSec is achieved with the Marvell Octeon 10
- Offload the work of IPSec encapsulation and decapsulation to the DPU.
- The Host receives and transmits packets without knowledge that the packets will be secured with IPSec



# Example Use case: Load Balancing

- Load Balancer can balance HTTP requests towards multiple HTTP endpoint pods.
- Taking over the role of an CPE (Customer Premise Equipment)



# Roadmap & Availability

OpenShift and Microshift planned GA of the DPU-Operator

- 4.19 for both OpenShift and Microshift
- Underlying RHEL 9.6

The DPU Operator is Open Source and companies like Marvell & Intel are active contributors:

- <https://github.com/openshift/dpu-operator>
- The goal is to build an API for OPI (open programmable infrastructure)

Q n A

Thank You