Introduction Infrastructure Programming

Presented by // DAN DALY

OCTOBER, 2021

IPDK Overview

IPDK is a **community-driven target agnostic** framework for **infrastructure programming** that runs on a CPU, IPU, DPU, or switch.

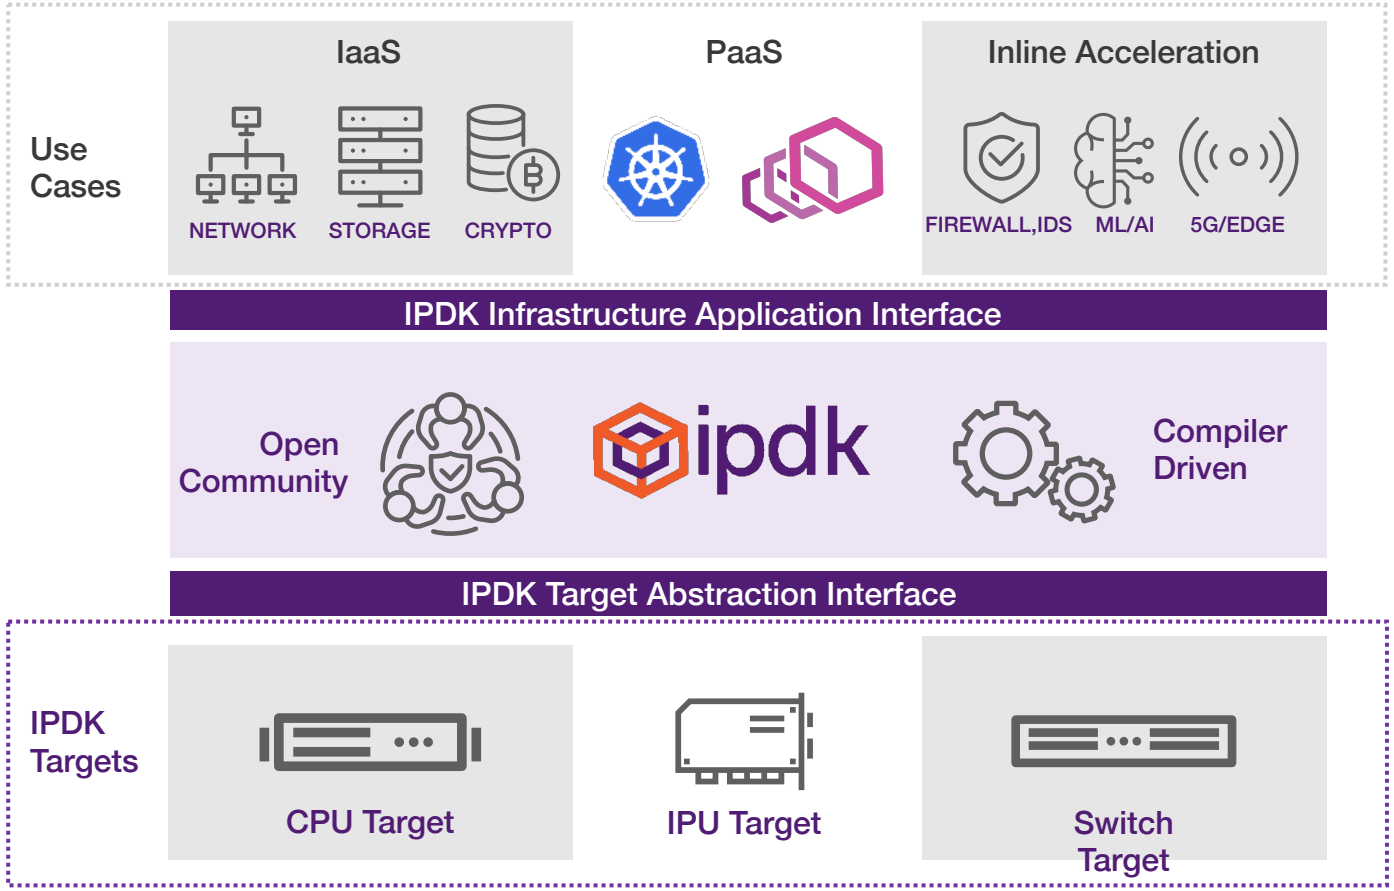

IPDK.io: Infrastructure Programmer Development Kit
Collaborate with the community on Github & Slack

ipdk

Agenda

- Charter
- Approach
- Use cases
- Open-source development
  - Example: Infrastructure-as-a-Service
- Roadmap
- Call to action

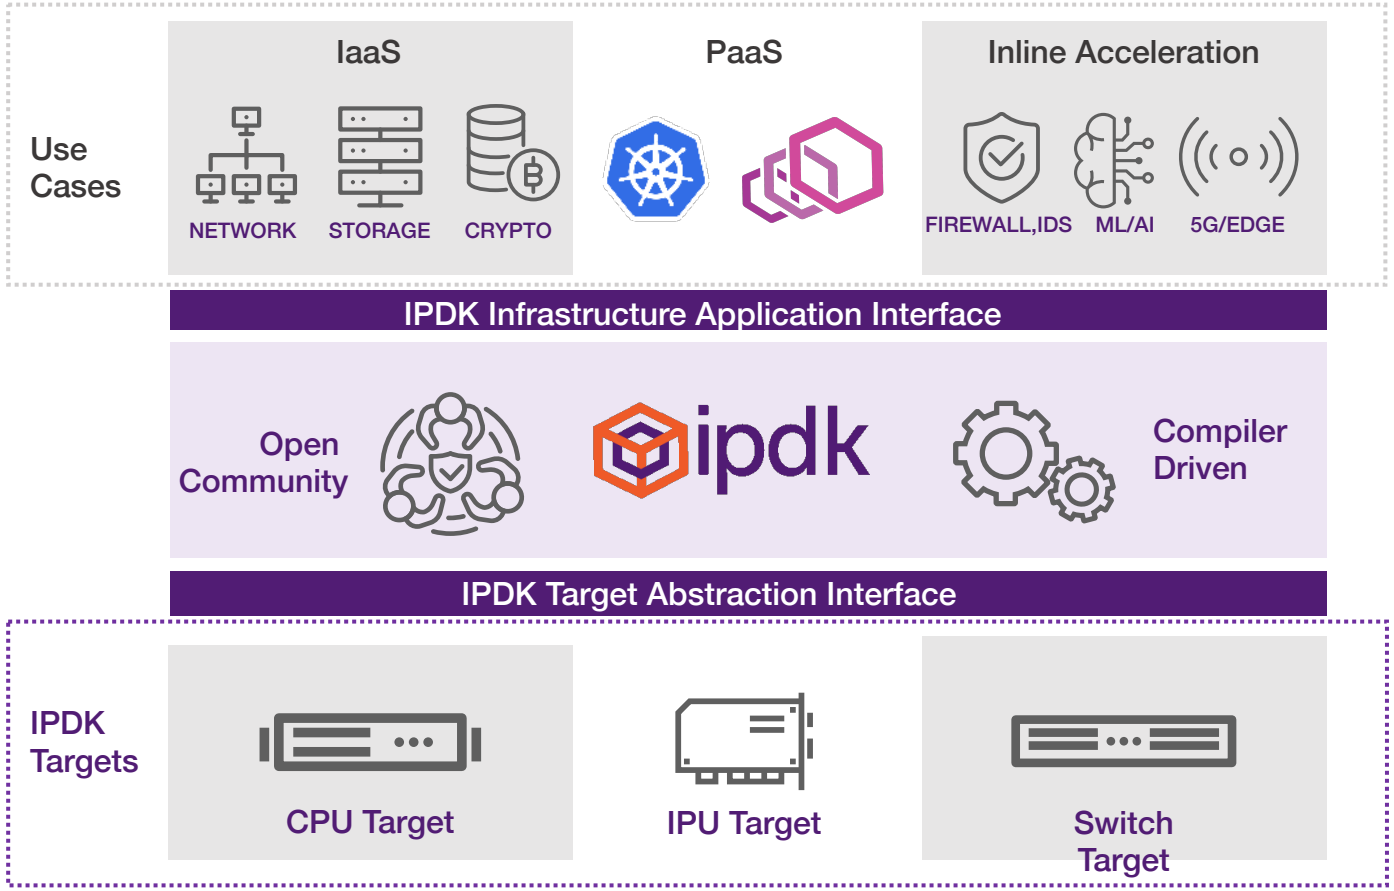

IPDK Charter

Open Community:

Open-source project with standard interfaces

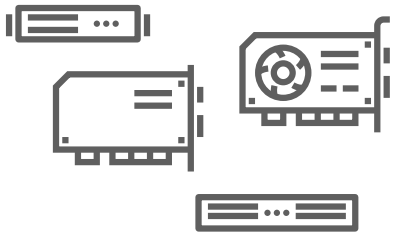

Target Agnostic:

Any CPU, IPU, DPU, or switch

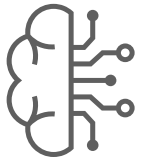

Use-Case Driven:

Use programmability for adaptability

Approach

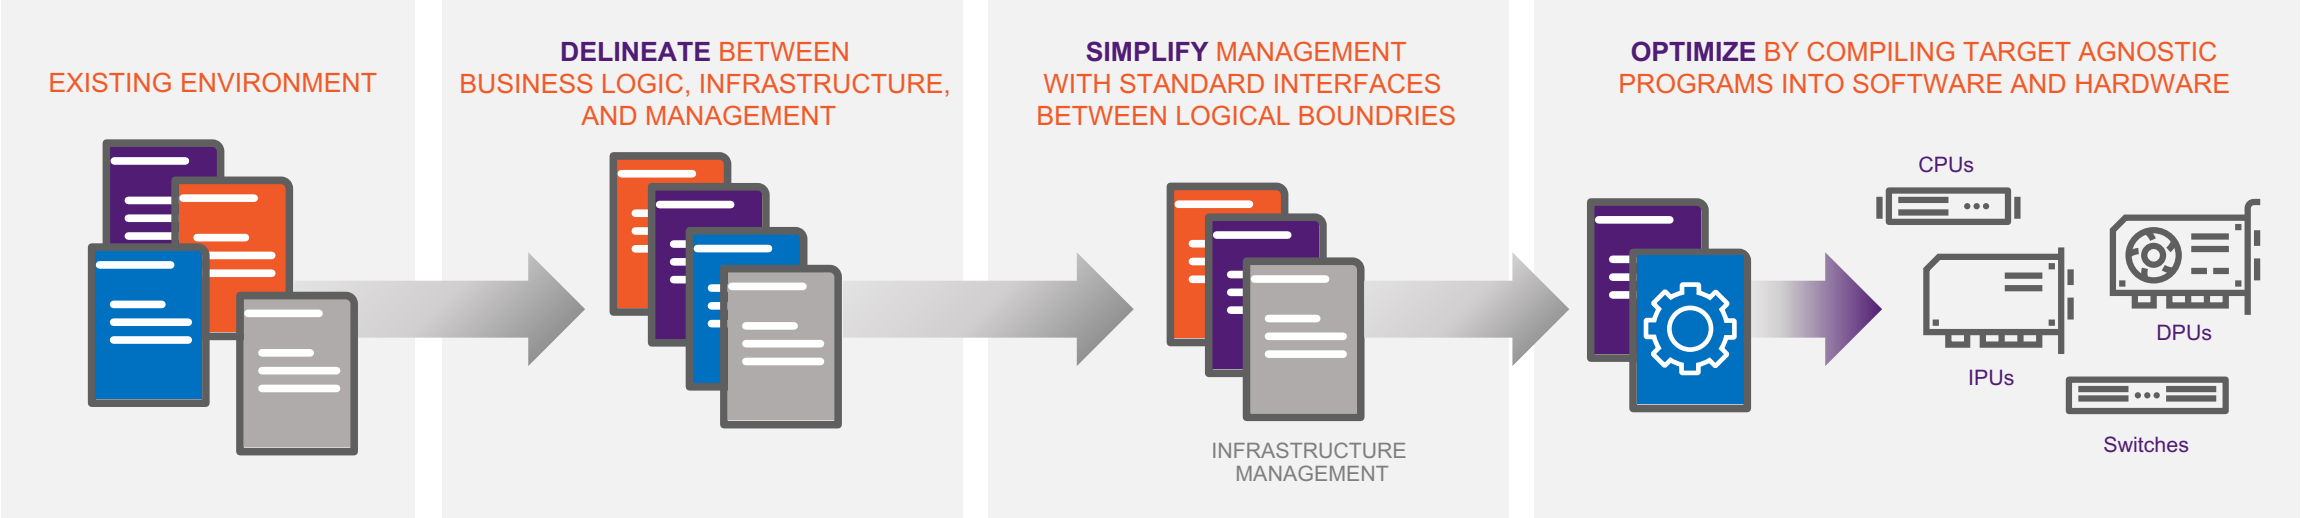

1. **Delineate** Business Logic vs. Infrastructure
2. **Simplify** Infrastructure Management
3. **Optimize** using a Compiler-Driven Target Abstraction

Use-Case Driven

1. **Infrastructure-as-a-Service**
   Virtual networking, storage and Crypto across VMs, containers and bare metal
2. **Platform-as-a-Service**
   Container Networking (Kubernetes) Sidecars (Envoy, MongoDB)
3. **Inline Acceleration**
   Firewall, IDS, Network Telemetry 5G/Wireless Infrastructure, AI/ML

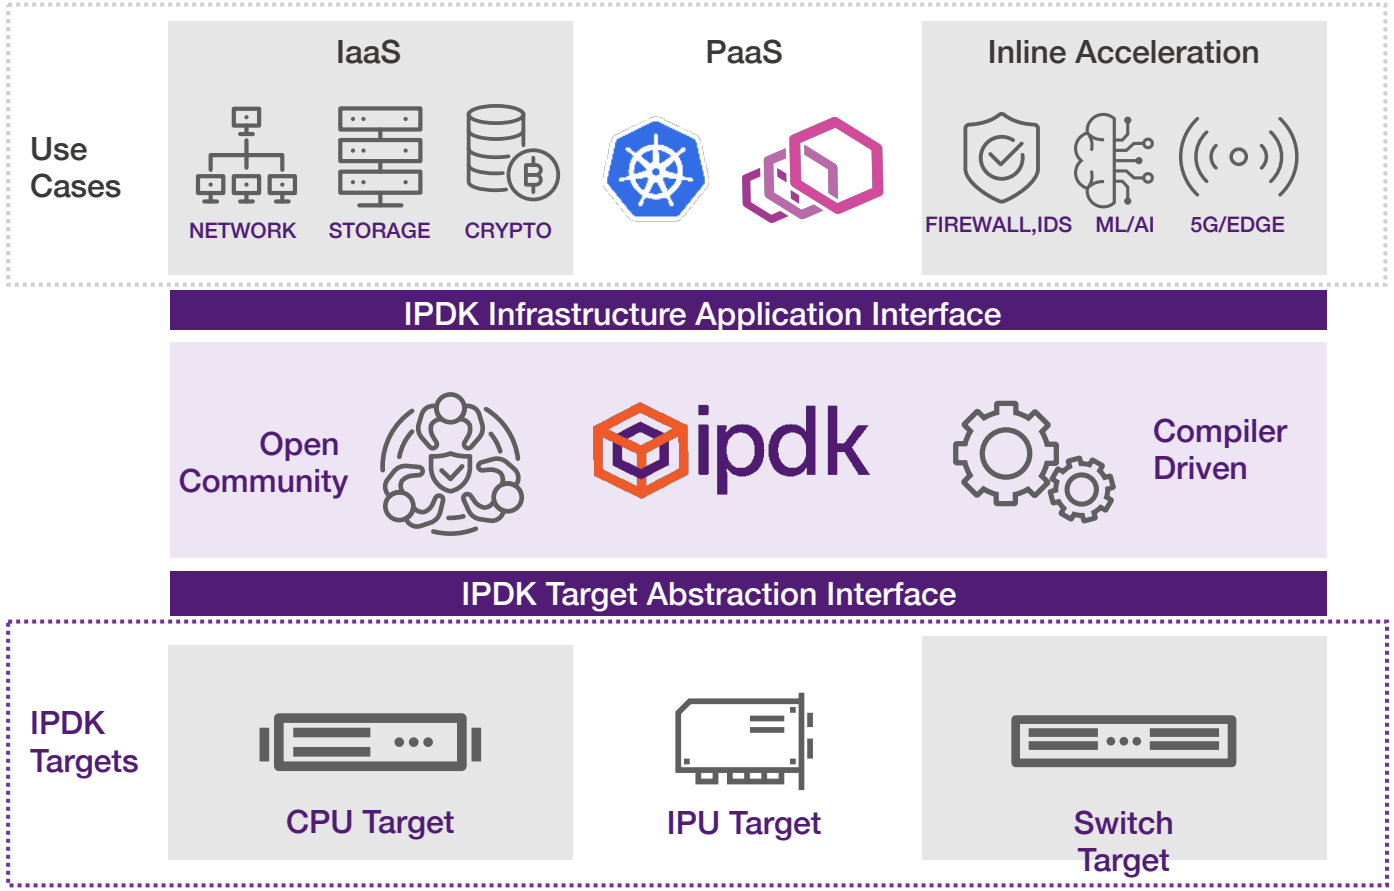

IPDK.io: Infrastructure Programmer Development Kit
Collaborate with the community on Github & Slack

ipdk

Open-Source Development

- **Recipes**
  Delineate, simplify, and optimize for each major use case and functional area
- **Open-Source Development & Governance**
  New patches, agents, and interface support
  Open-Source definitions of interfaces
  Dockerfiles and pre-built containers
- **Development has started, come join!**
  Collaborate on Slack , Github & IPDK.io

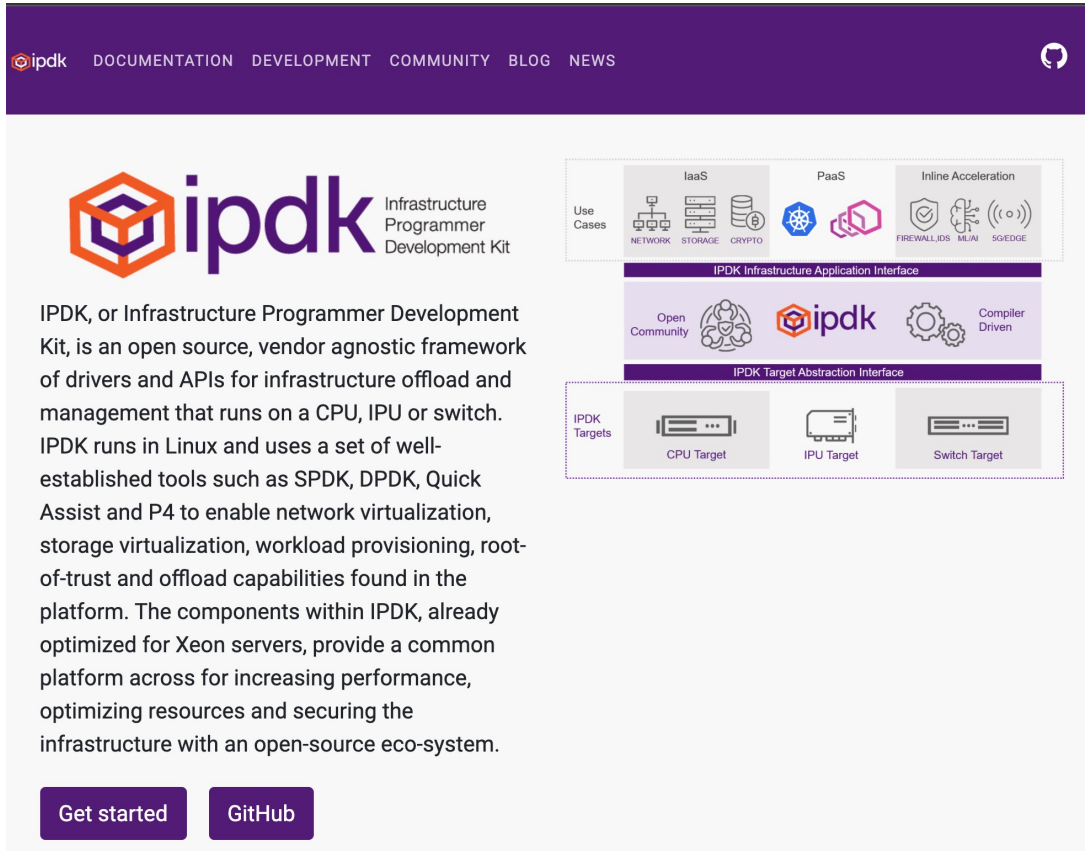

IPDK Standard Interfaces

- **Infrastructure Application Interface**
  Devices & Services for Workload Apps
  Platform Capabilities Offered to Mgmt
- **Target Abstraction Interface (TAI)**
  Target Capabilities Offered to IPDK

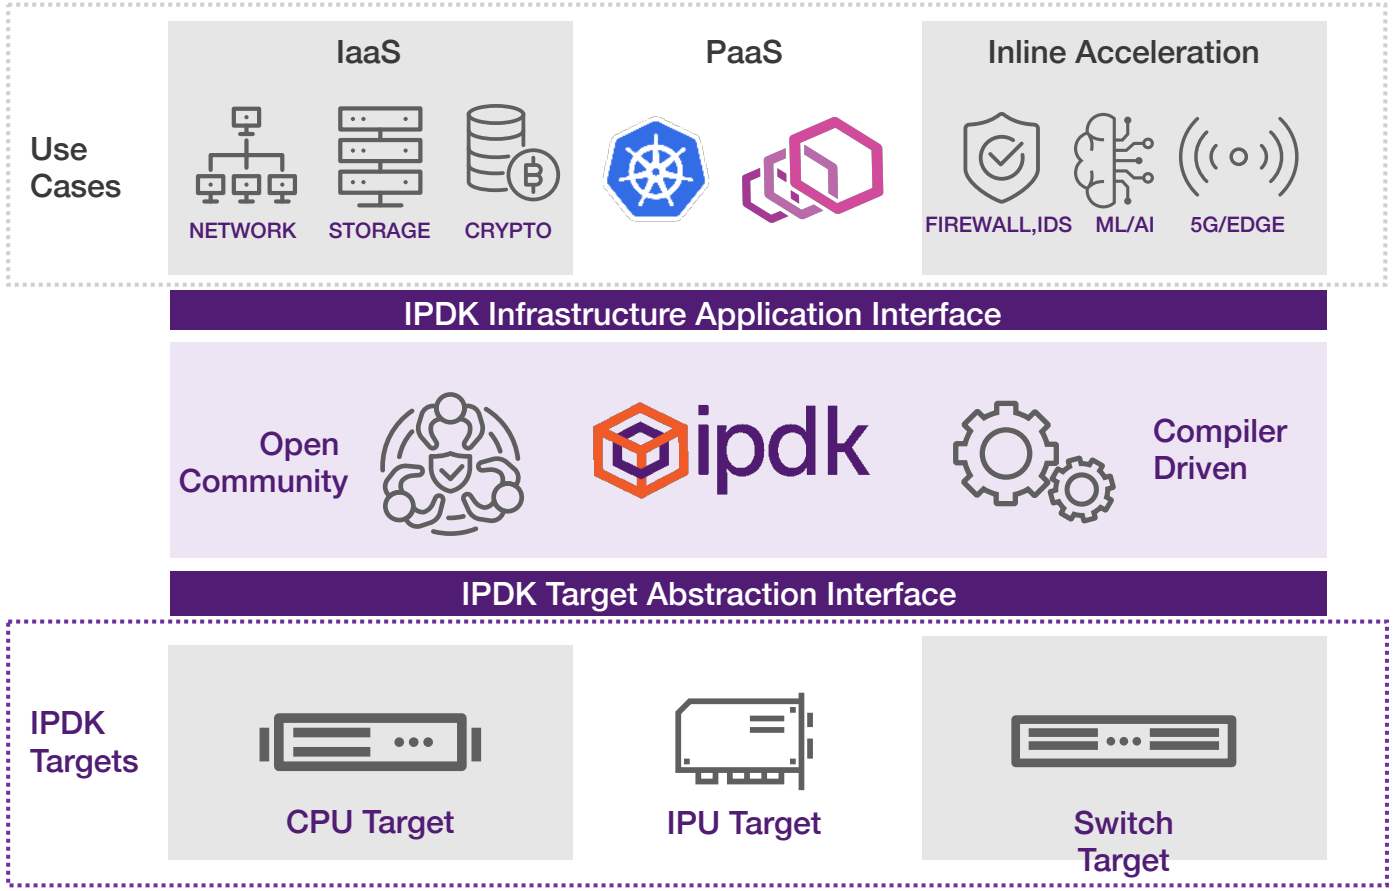

Take-Home Demo

Download Containers

- IPDK Container
- Traffic Gen Container
- Storage Target Container

Create Virtual Devices

- Create virtual port
- Create virtual disk

Generate network traffic and storage read/writes

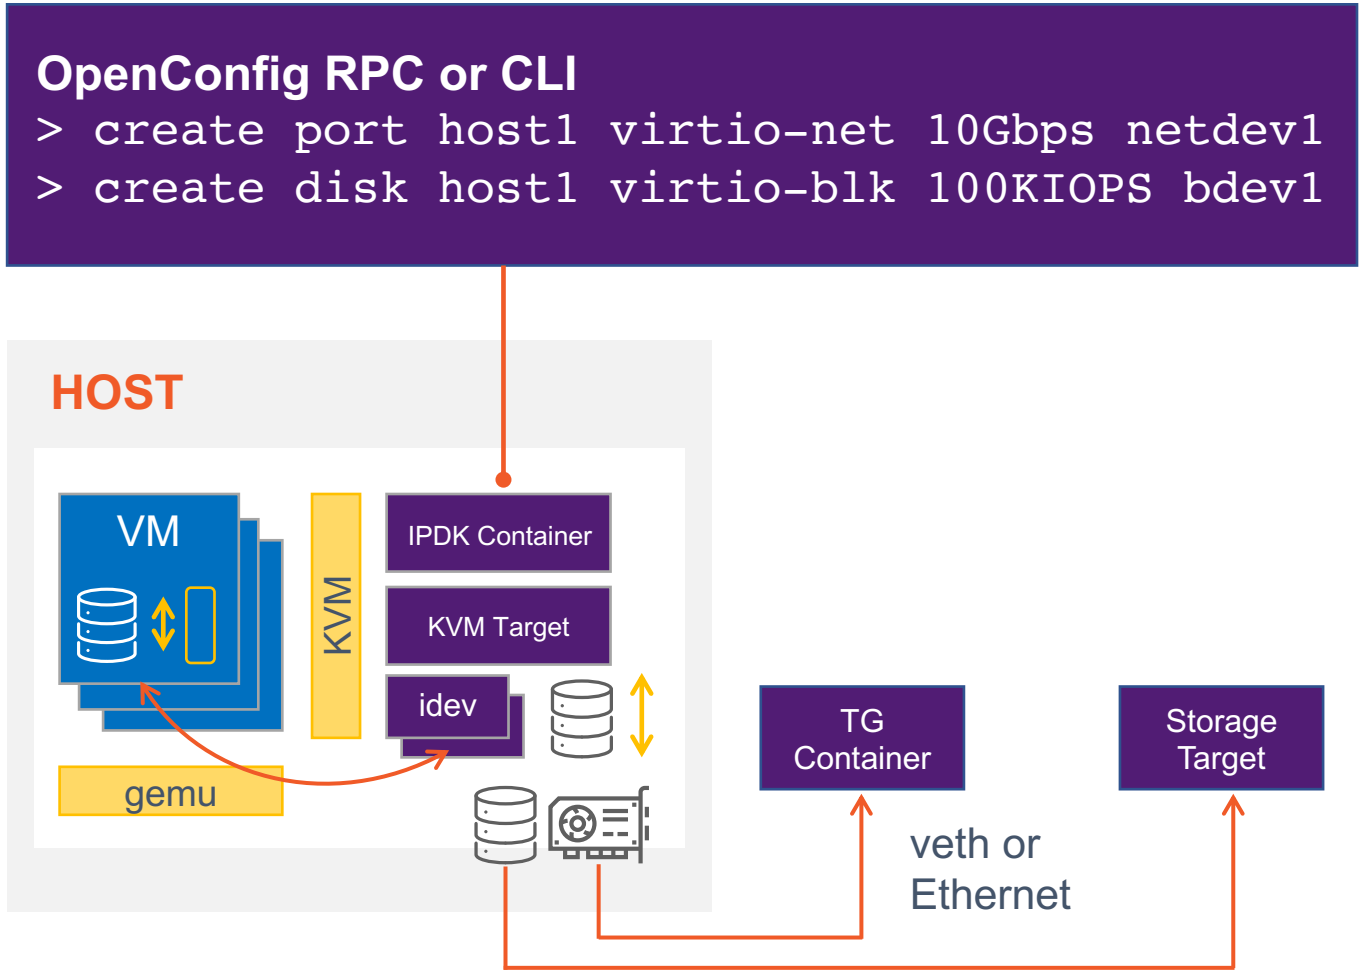

Where Can IPDK Run?

KVM HOST

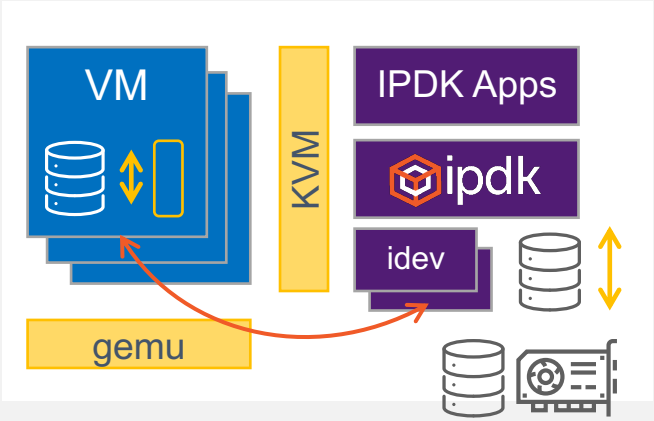

- Insert devices into **VMs**
- Sockets in/out of **VM** apps/containers
- **Direct attached** devices (drives, NICs, Accel)

Software Hypervisor
Software Virtual Switch

VIRTUAL INSTANCE

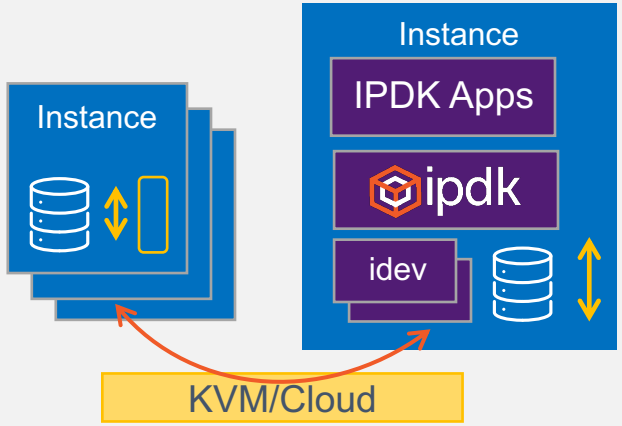

- Insert devices into **instances** (VM, bare metal)
- Sockets in/out of **instance** apps/containers
- **Virtual** devices (drives, NICs, Accel)

Cloud Hypervisor
Cloud Virtual Switch

KVM can also run in each instance

PHYSICAL DEVICE

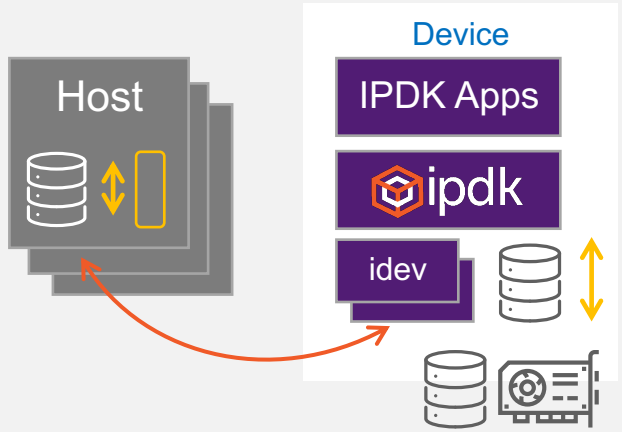

- Insert devices into **hosts** (VM, bare metal)
- Sockets in/out of **host or VM** apps/containers
- **Direct and virtual** devices (drives, NICs, Accel)

Hardware Hypervisor
Hardware Virtual Switch

KVM can also run in each instance

Call to Action

Come Join the Community

- Develop IPDK-compatible Applications
- Contribute an IPDK Target
  - Software, Switch, IPU, DPU, etc.
- Develop New Recipes
- Contribute to Existing Recipes
- Contribute to the Interfaces

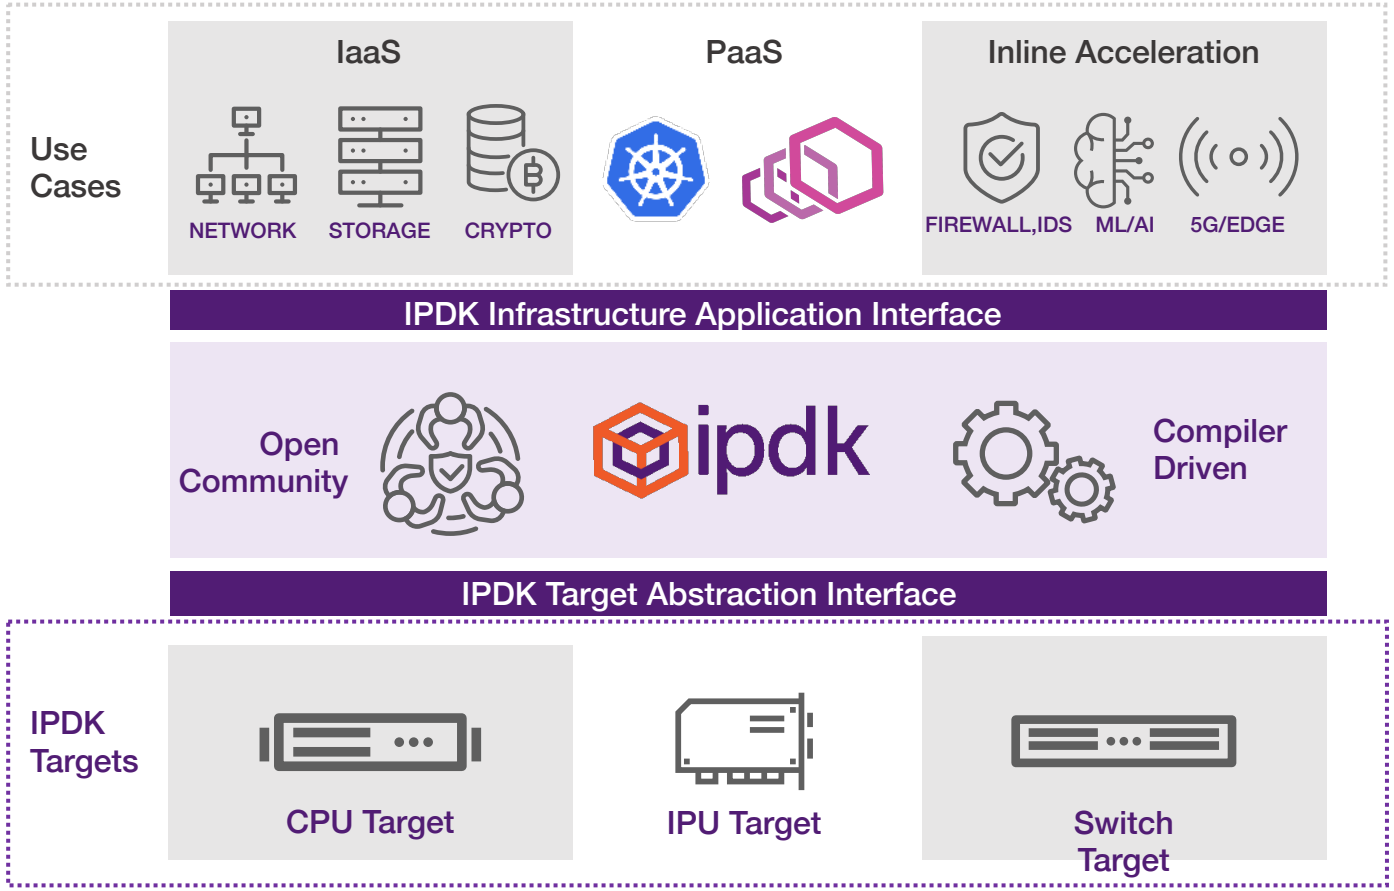

IPDK.io: Infrastructure Programmer Development Kit
Collaborate with the community on Github & Slack

ipdk

IPDK Container

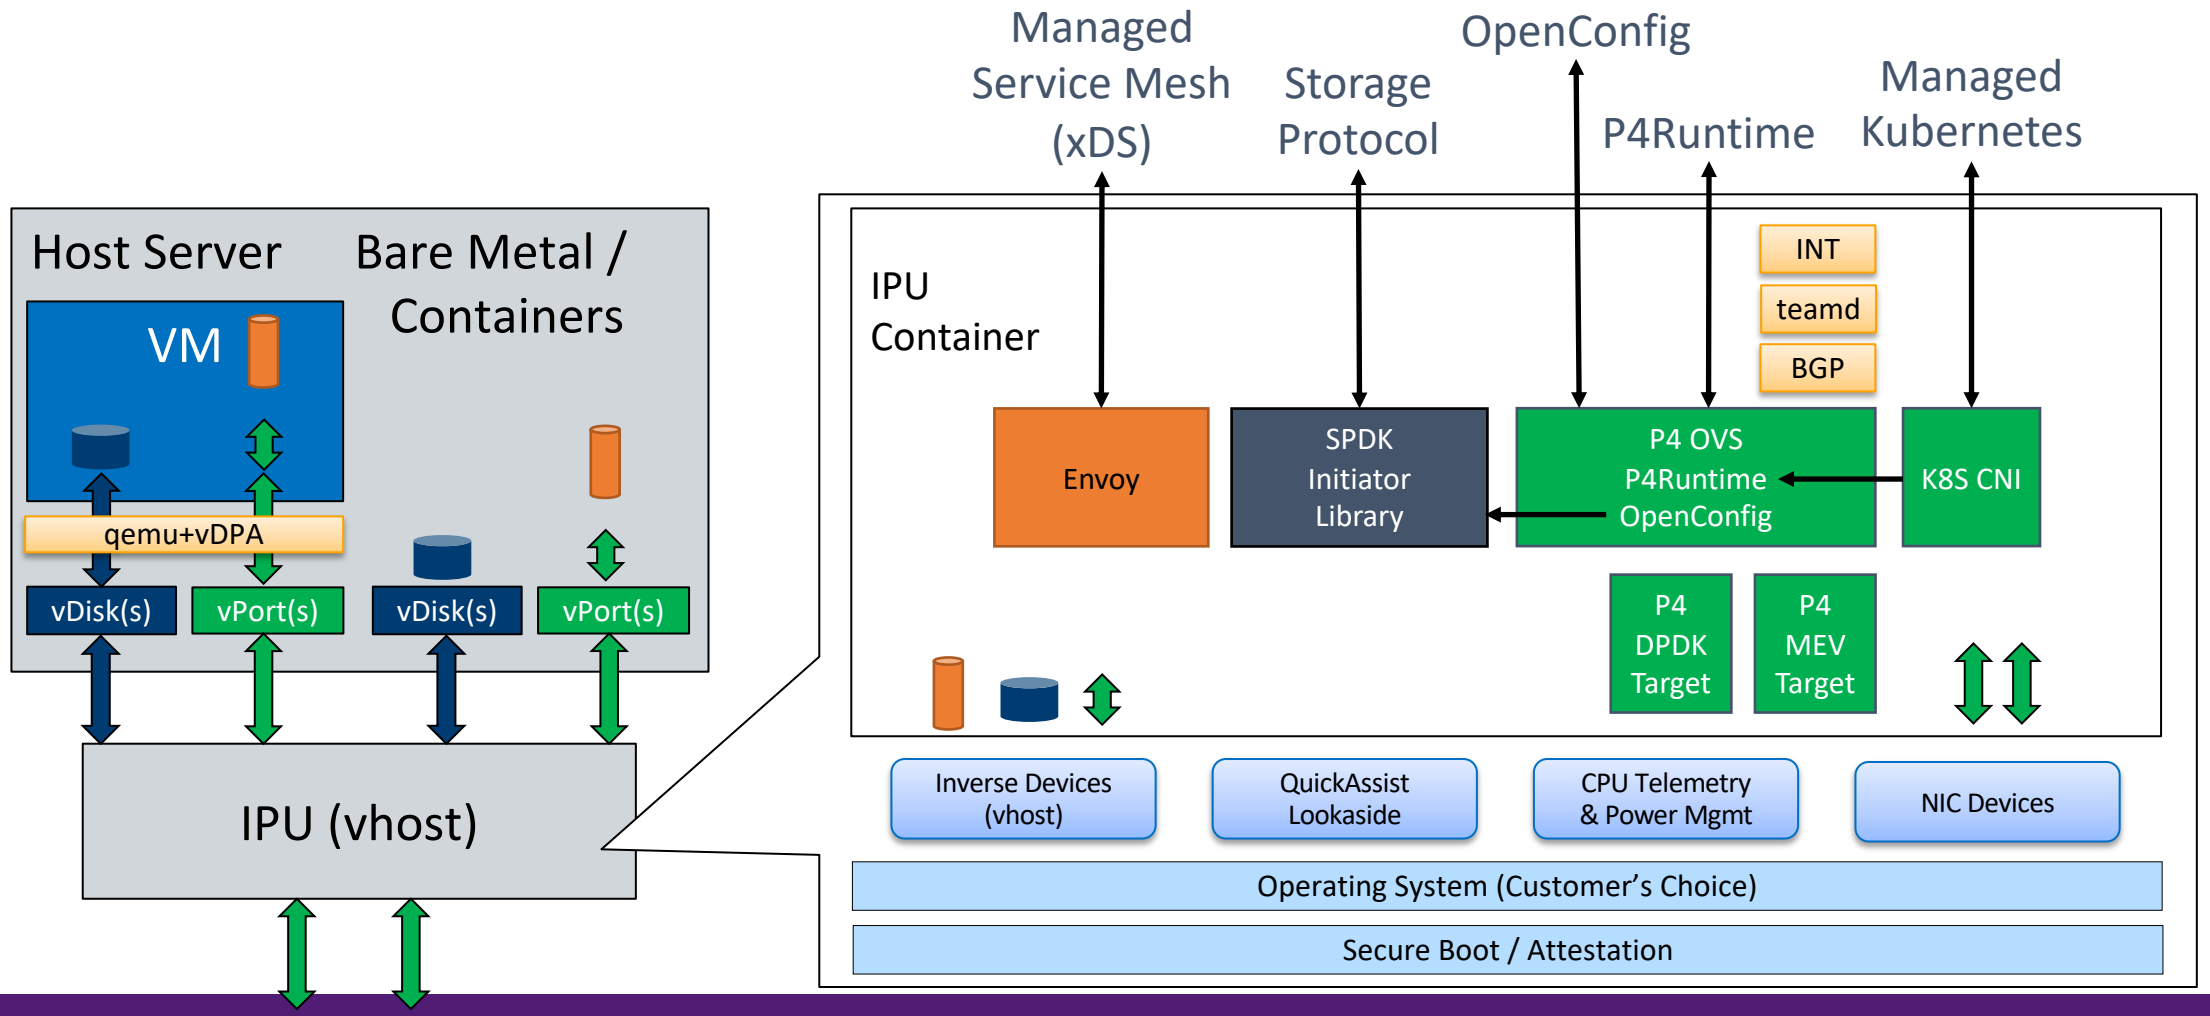

IaaS – Network Virtualization

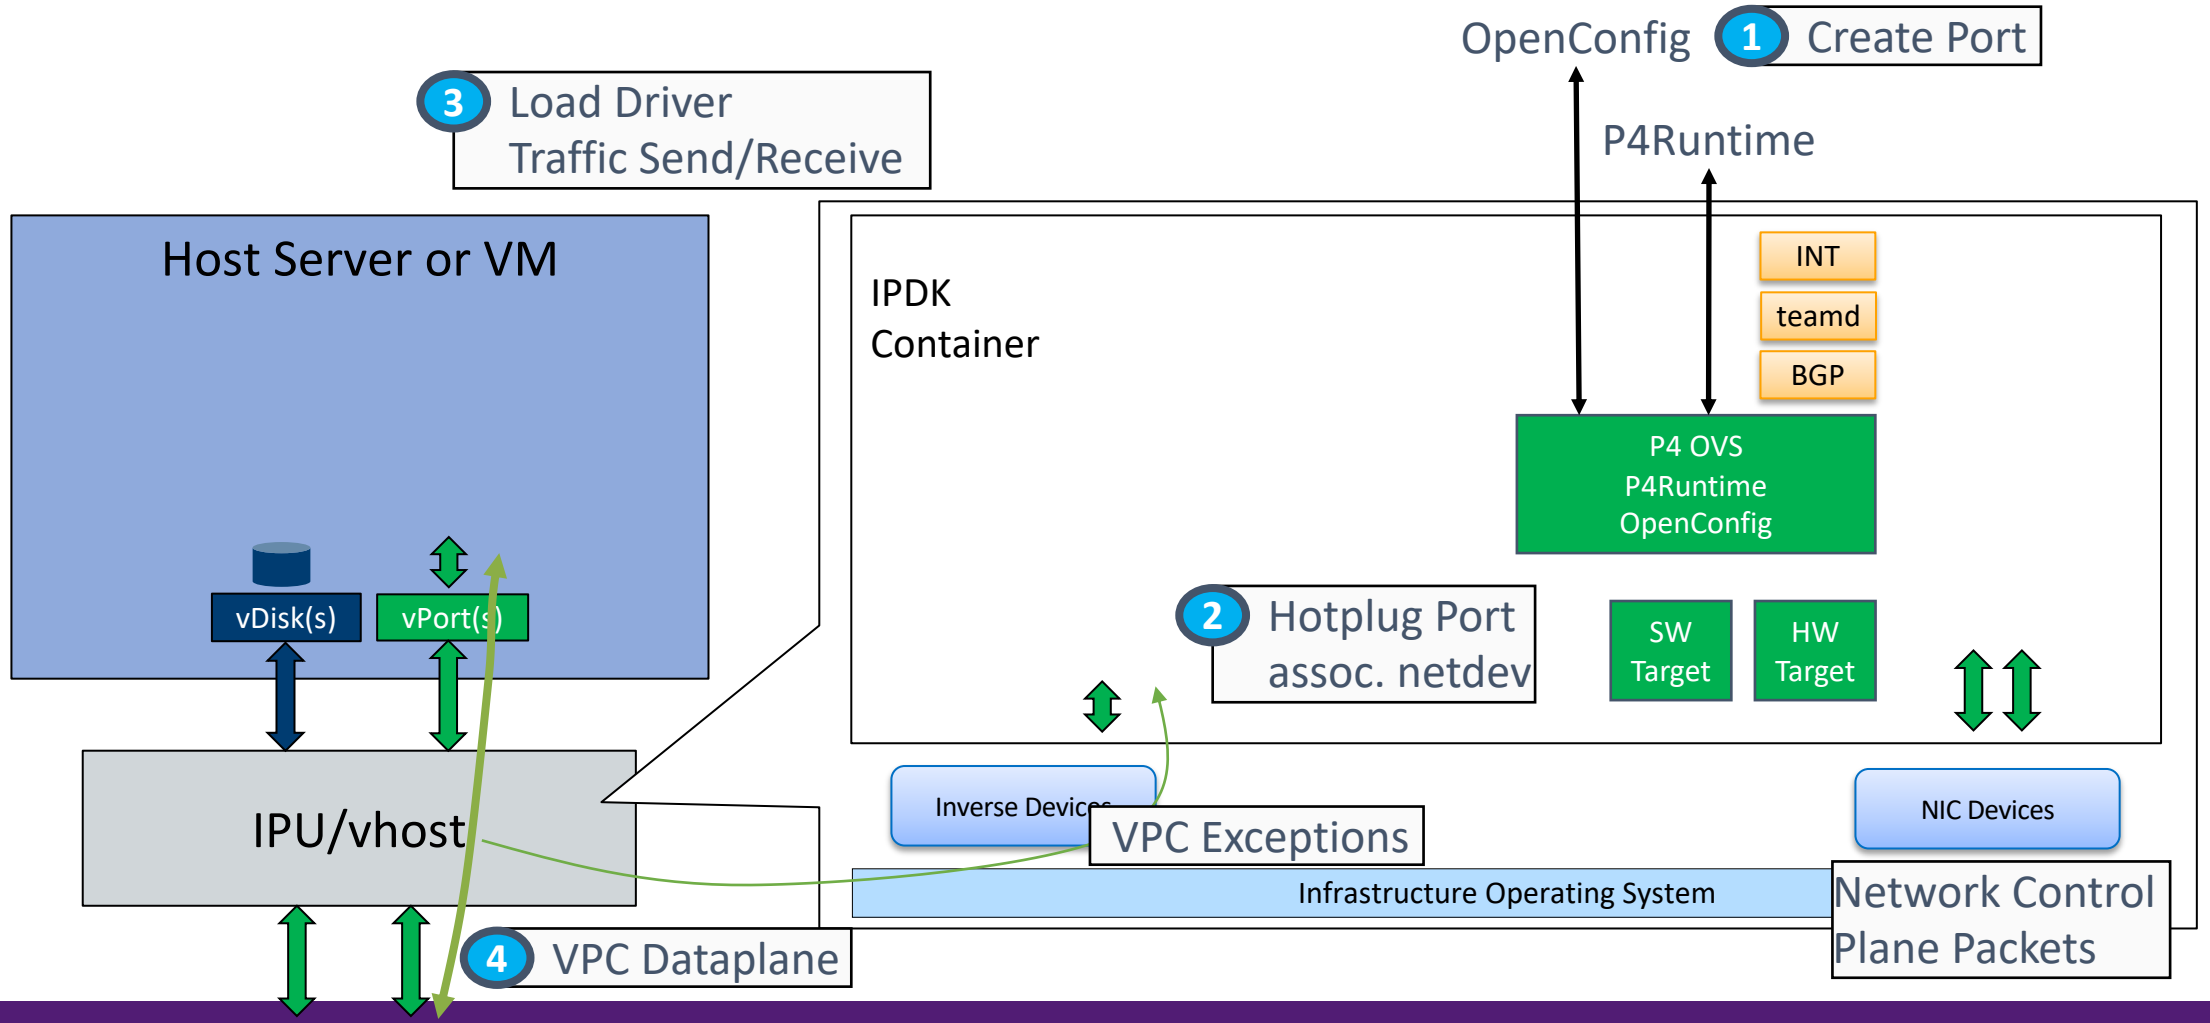

IaaS – Storage Virtualization

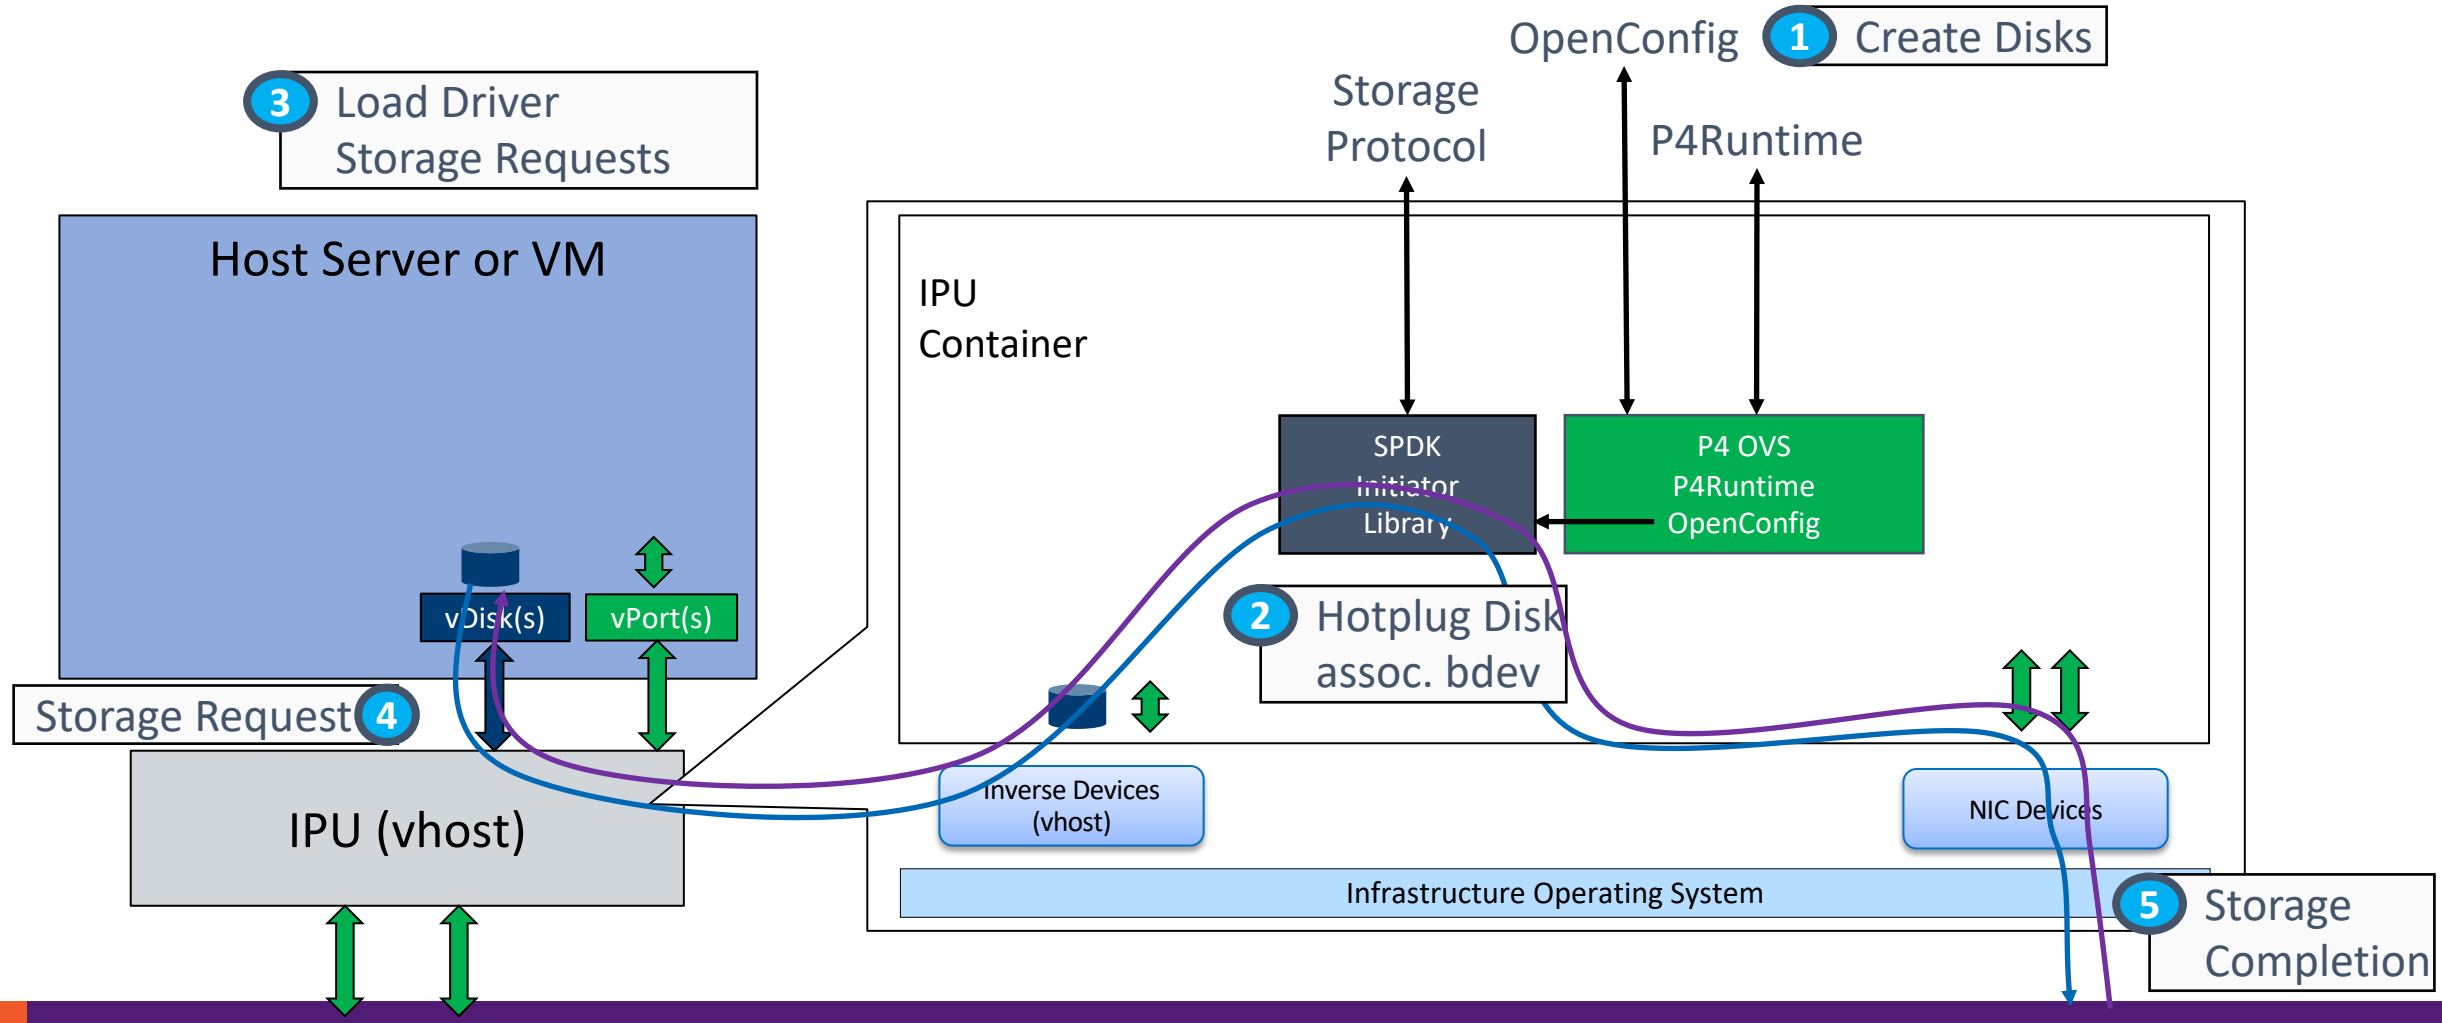

IaaS – Storage Virtualization - Interfaces

1. OpenConfig Operation
   ’add virtio-blk disk1 bdev5 10TB’
2. TDI (Table Driven Interface)
   ‘addEntry(vdev, “disk1”, config1)
3. Backend Device Operation
   - KVM: qemu hotplug
   - Canyon: pci hotplug
   - MEV: simulated hotplug
4. SPDK Operation
   associate device w/ ‘bdev5’

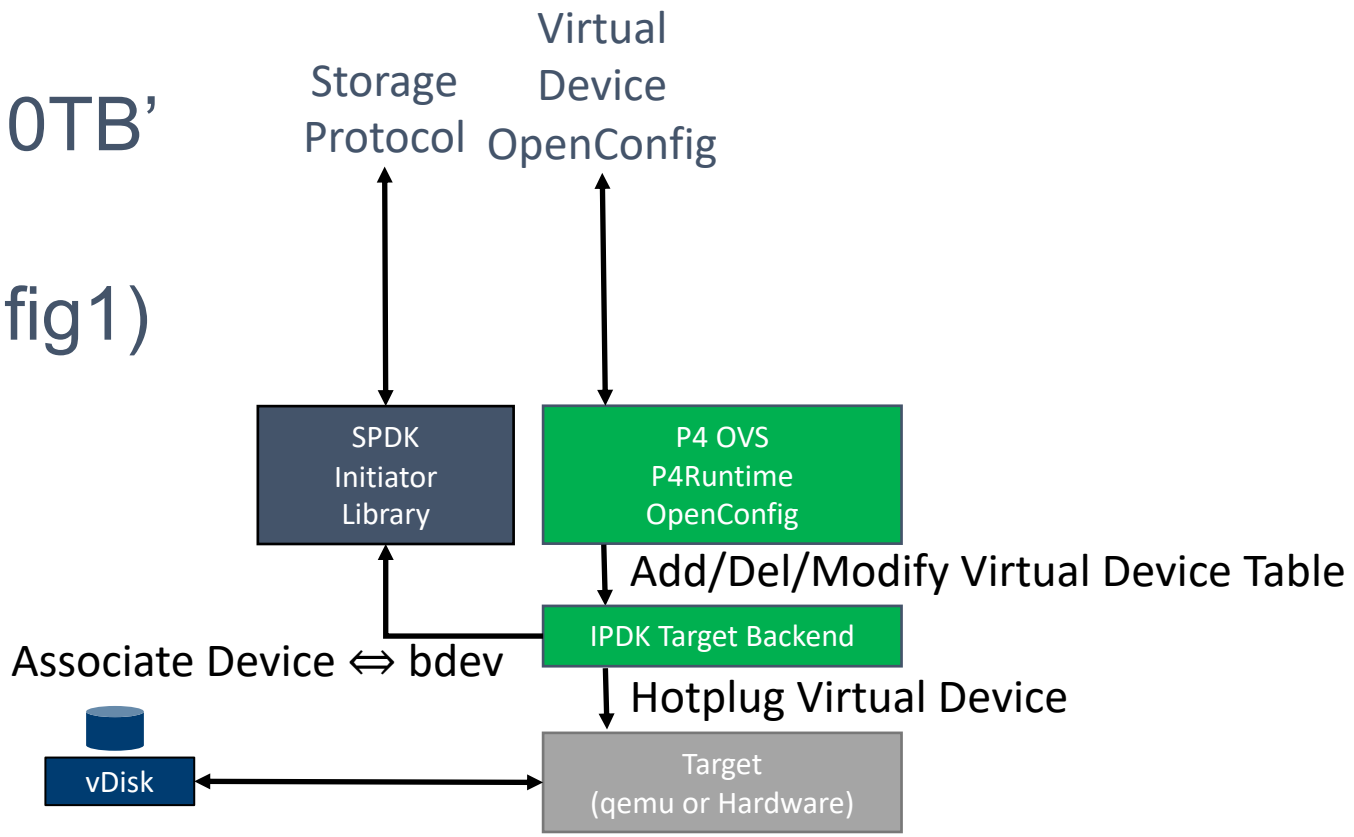

IaaS – Host Virtualization – Network

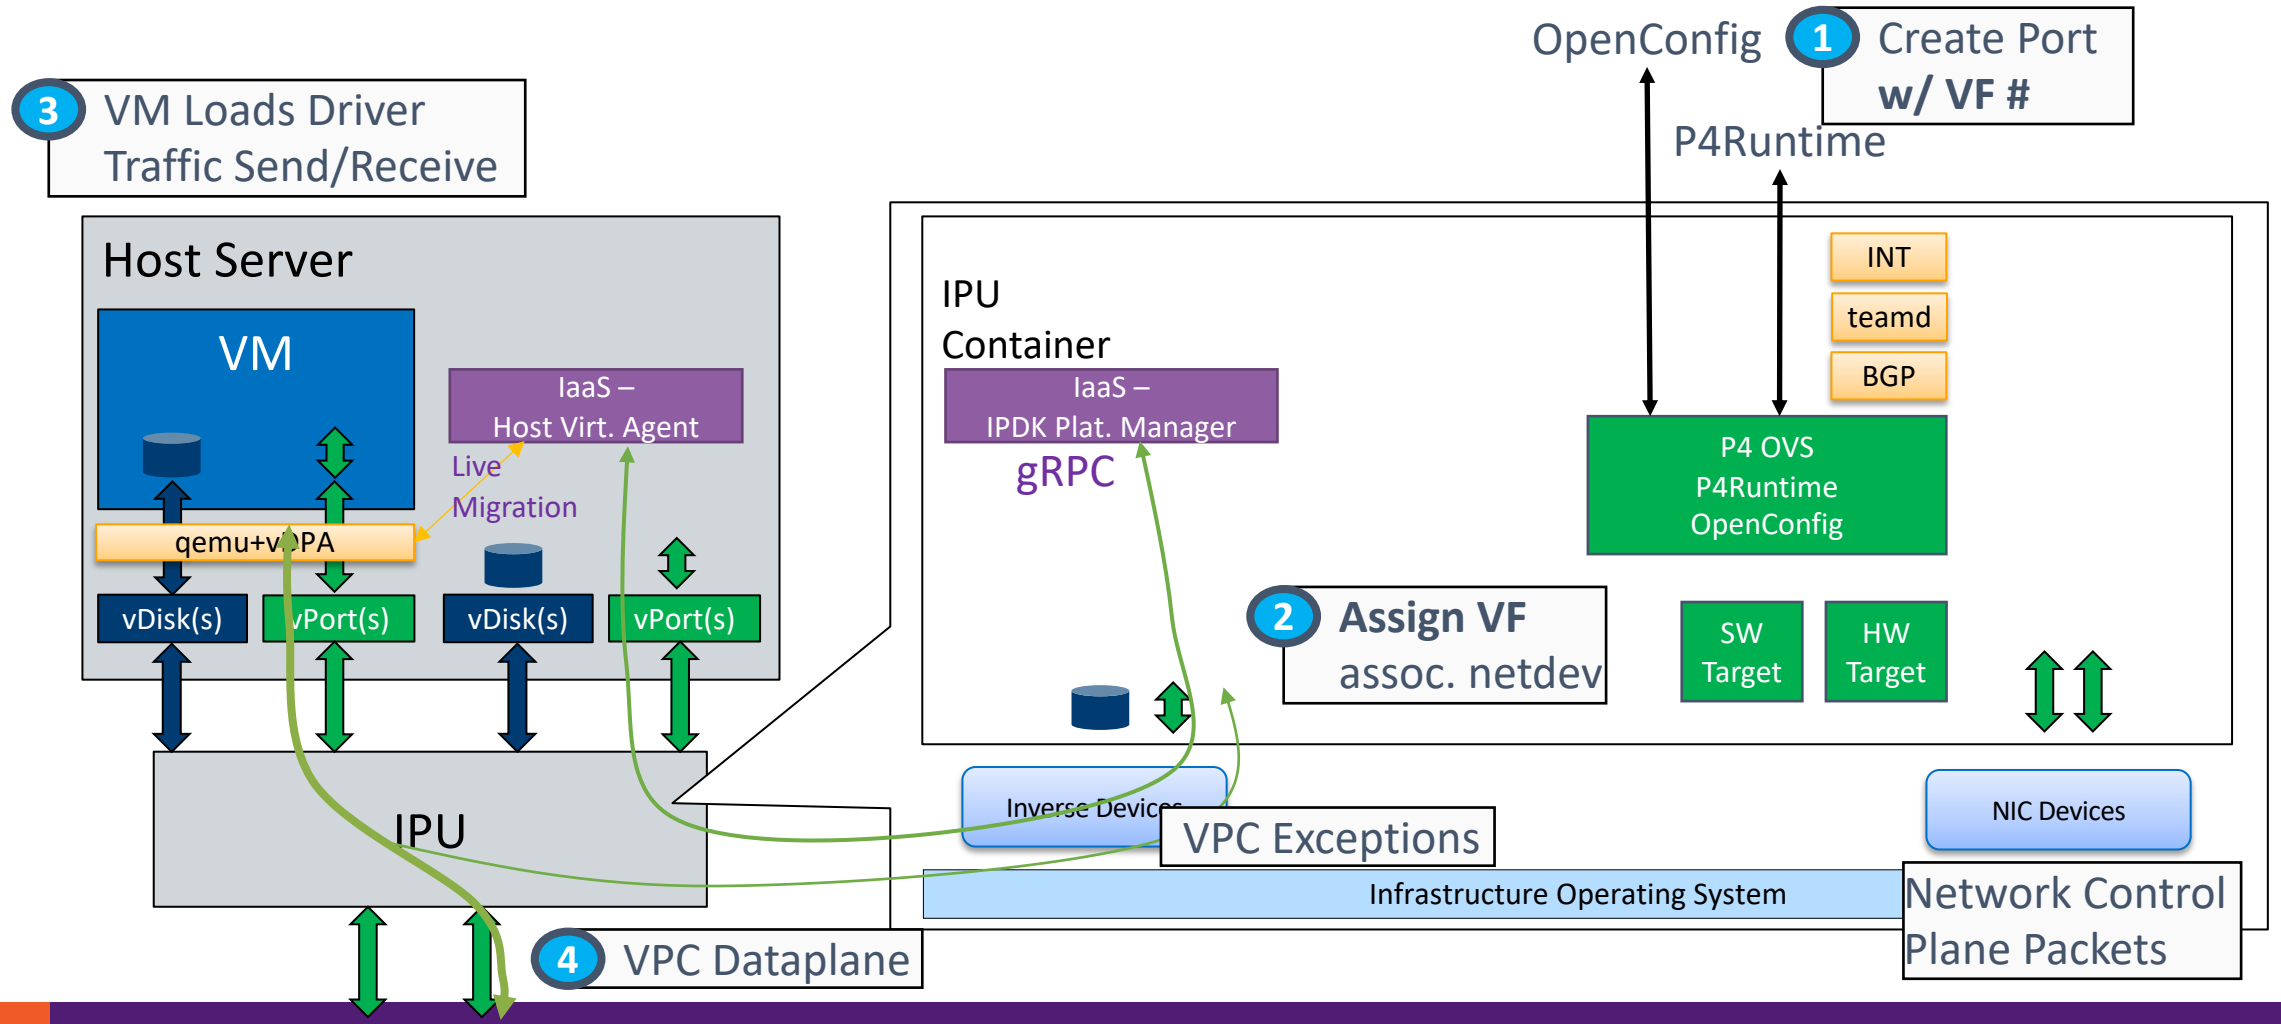

ipdk

IaaS – Host Virtualization – Disk

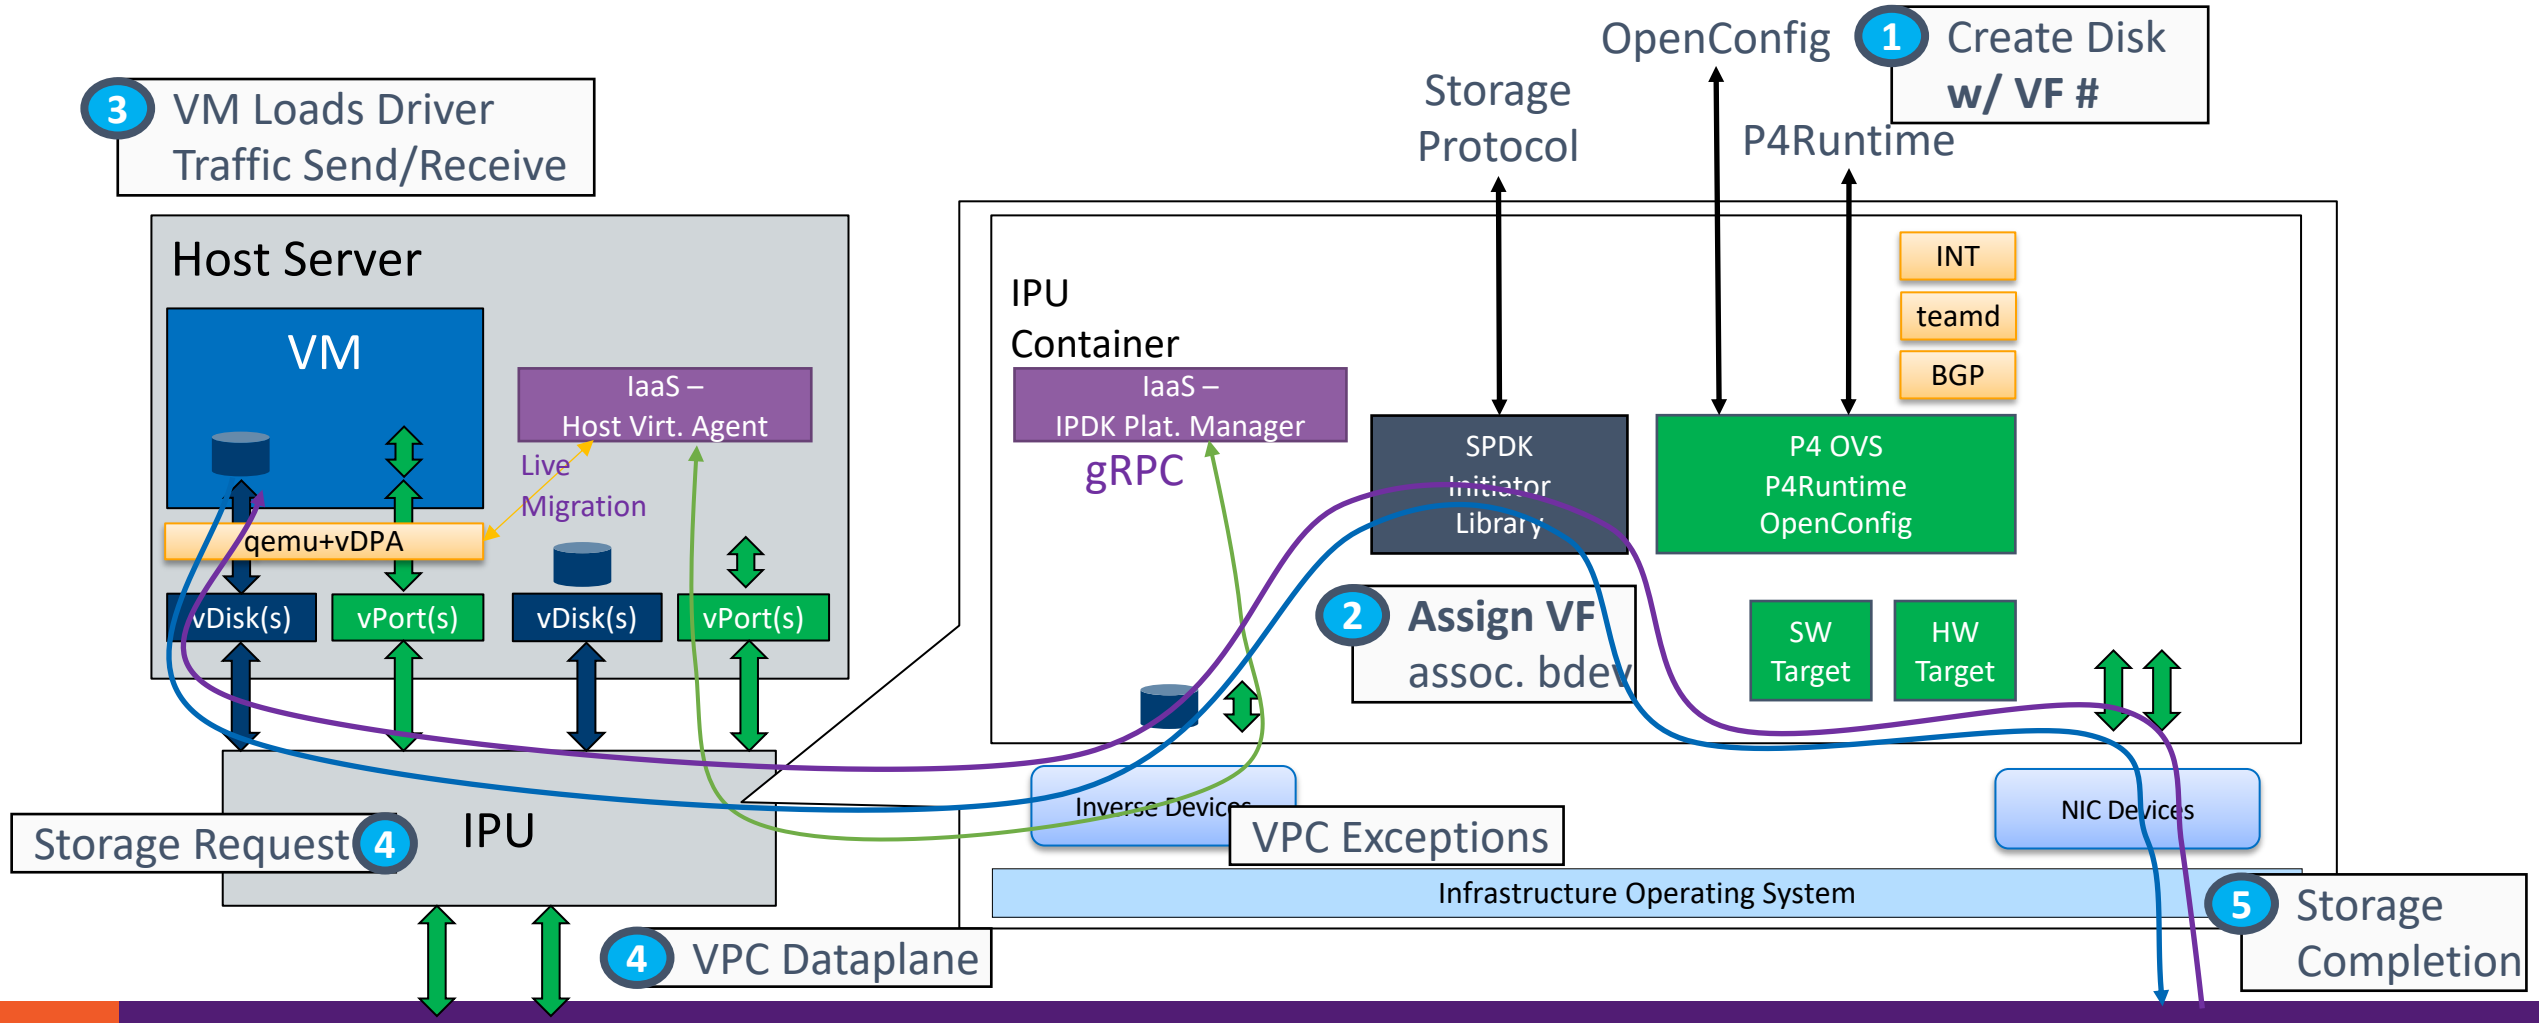

ipdk

IaaS – Host Virtualization – Net+Disk

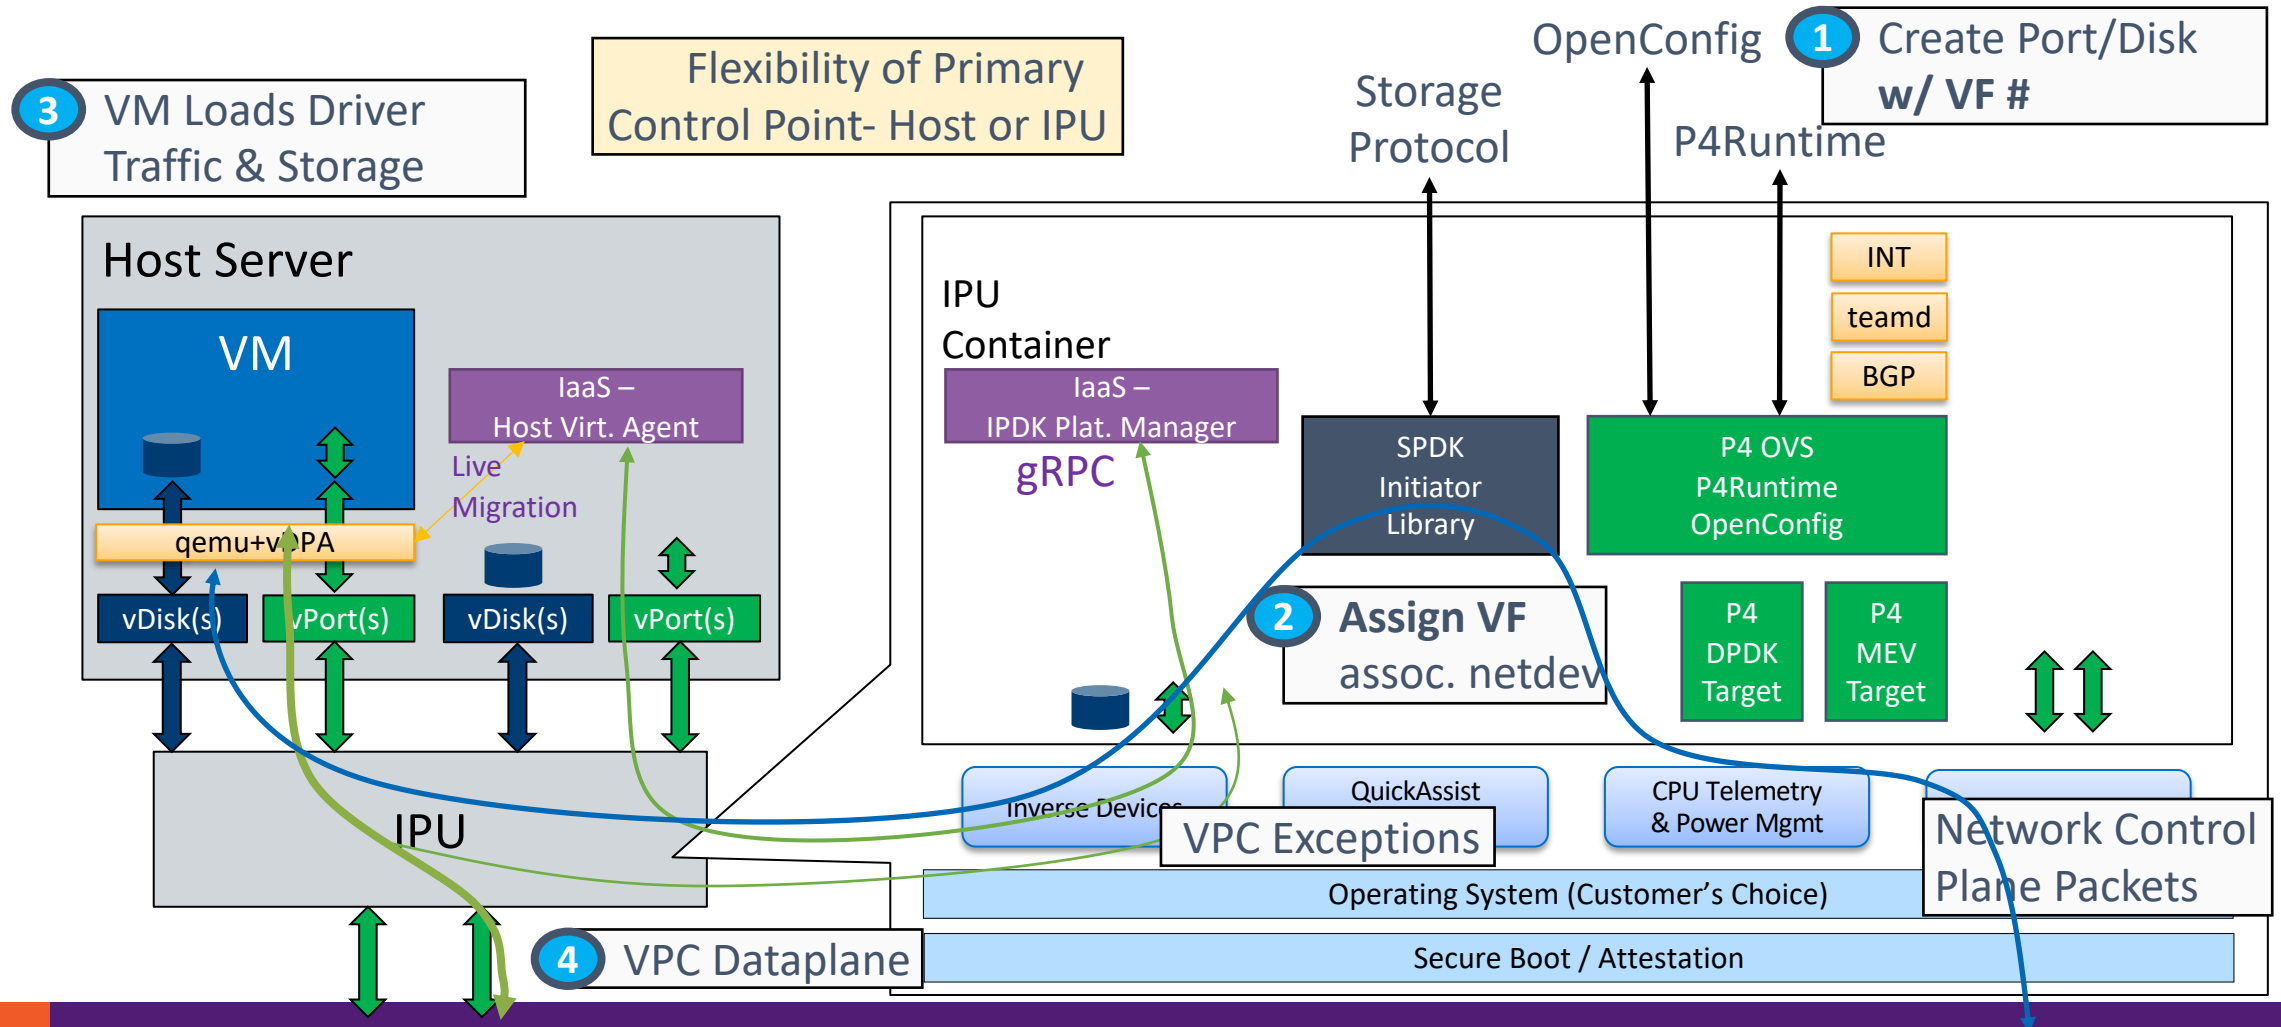

PaaS – CSI & CNI

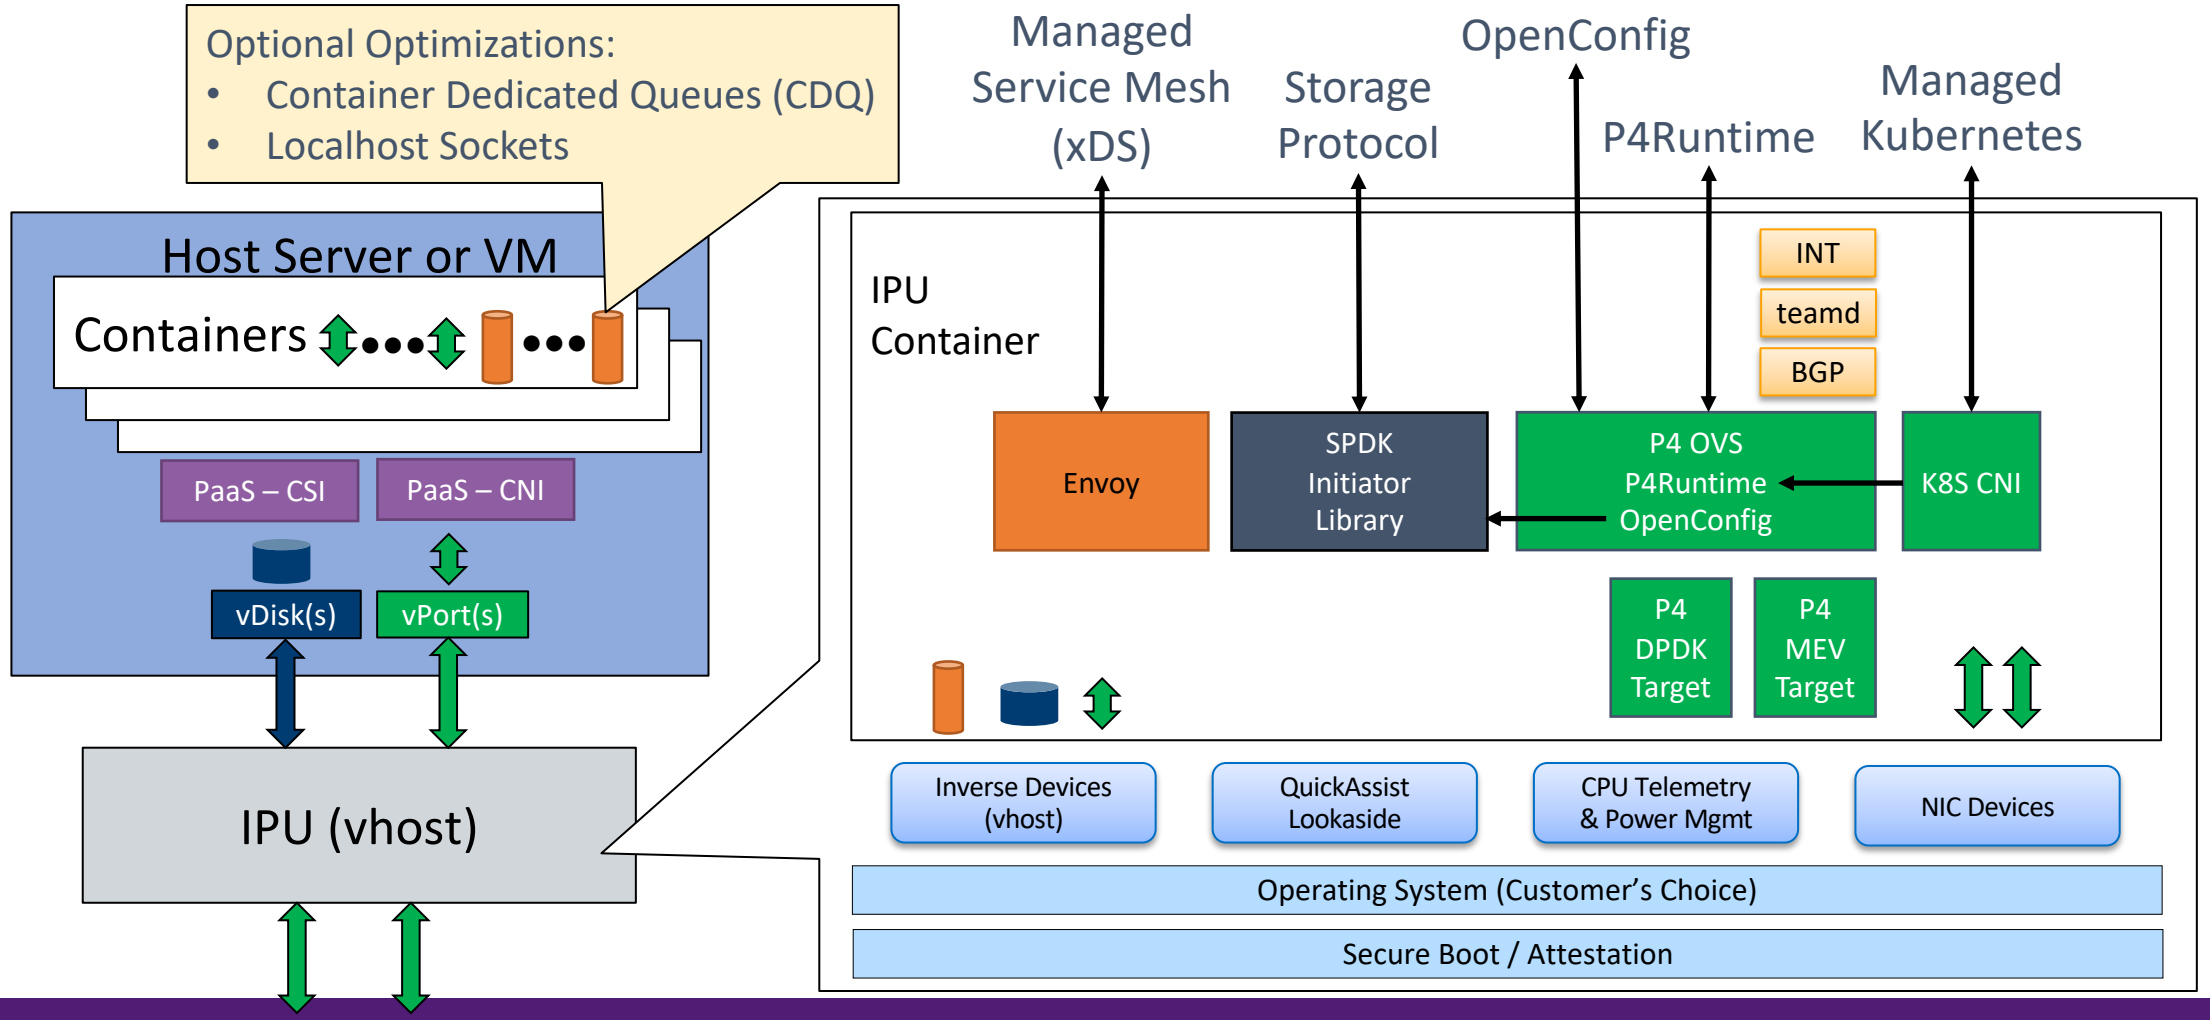

PaaS – K8s Delineations

0. Bare Metal, no IPU

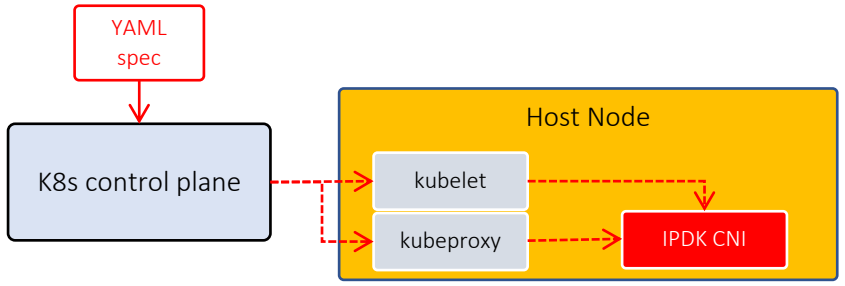

2. IPU anchored control (Managed Services)

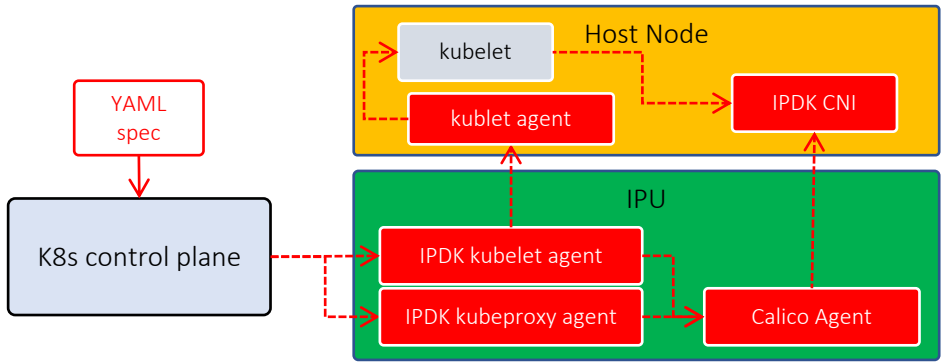

1. Trusted host anchored control (Self Managed)

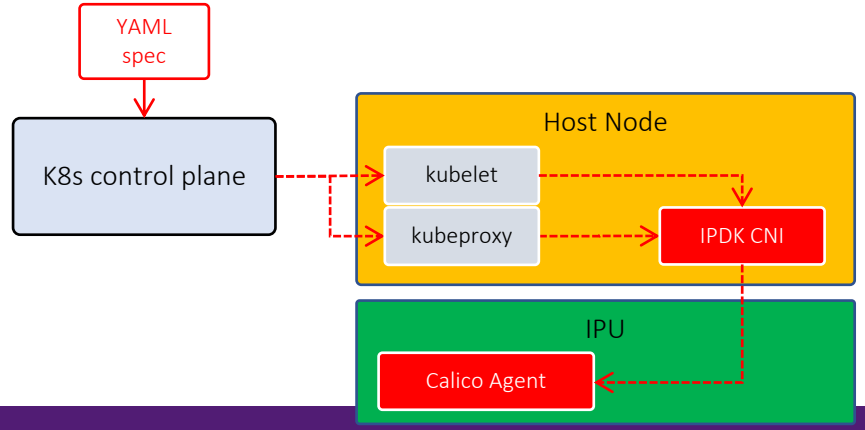

3. Split Tenant & Managed Infrastructure

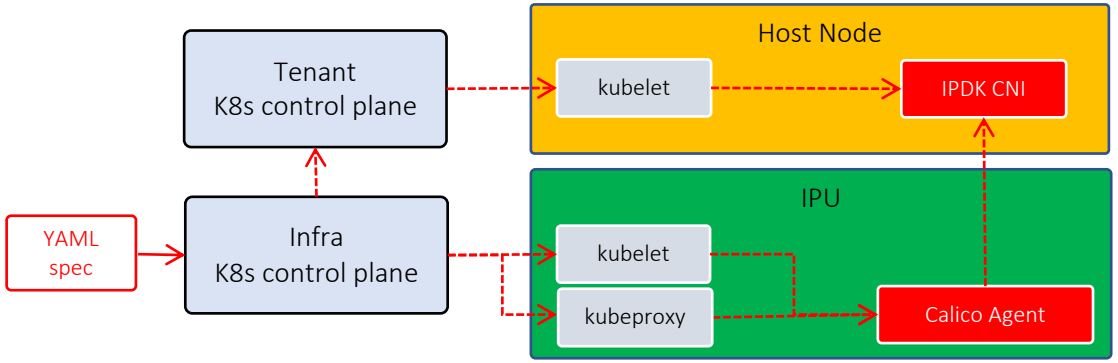

PaaS – Calico Running on IPU

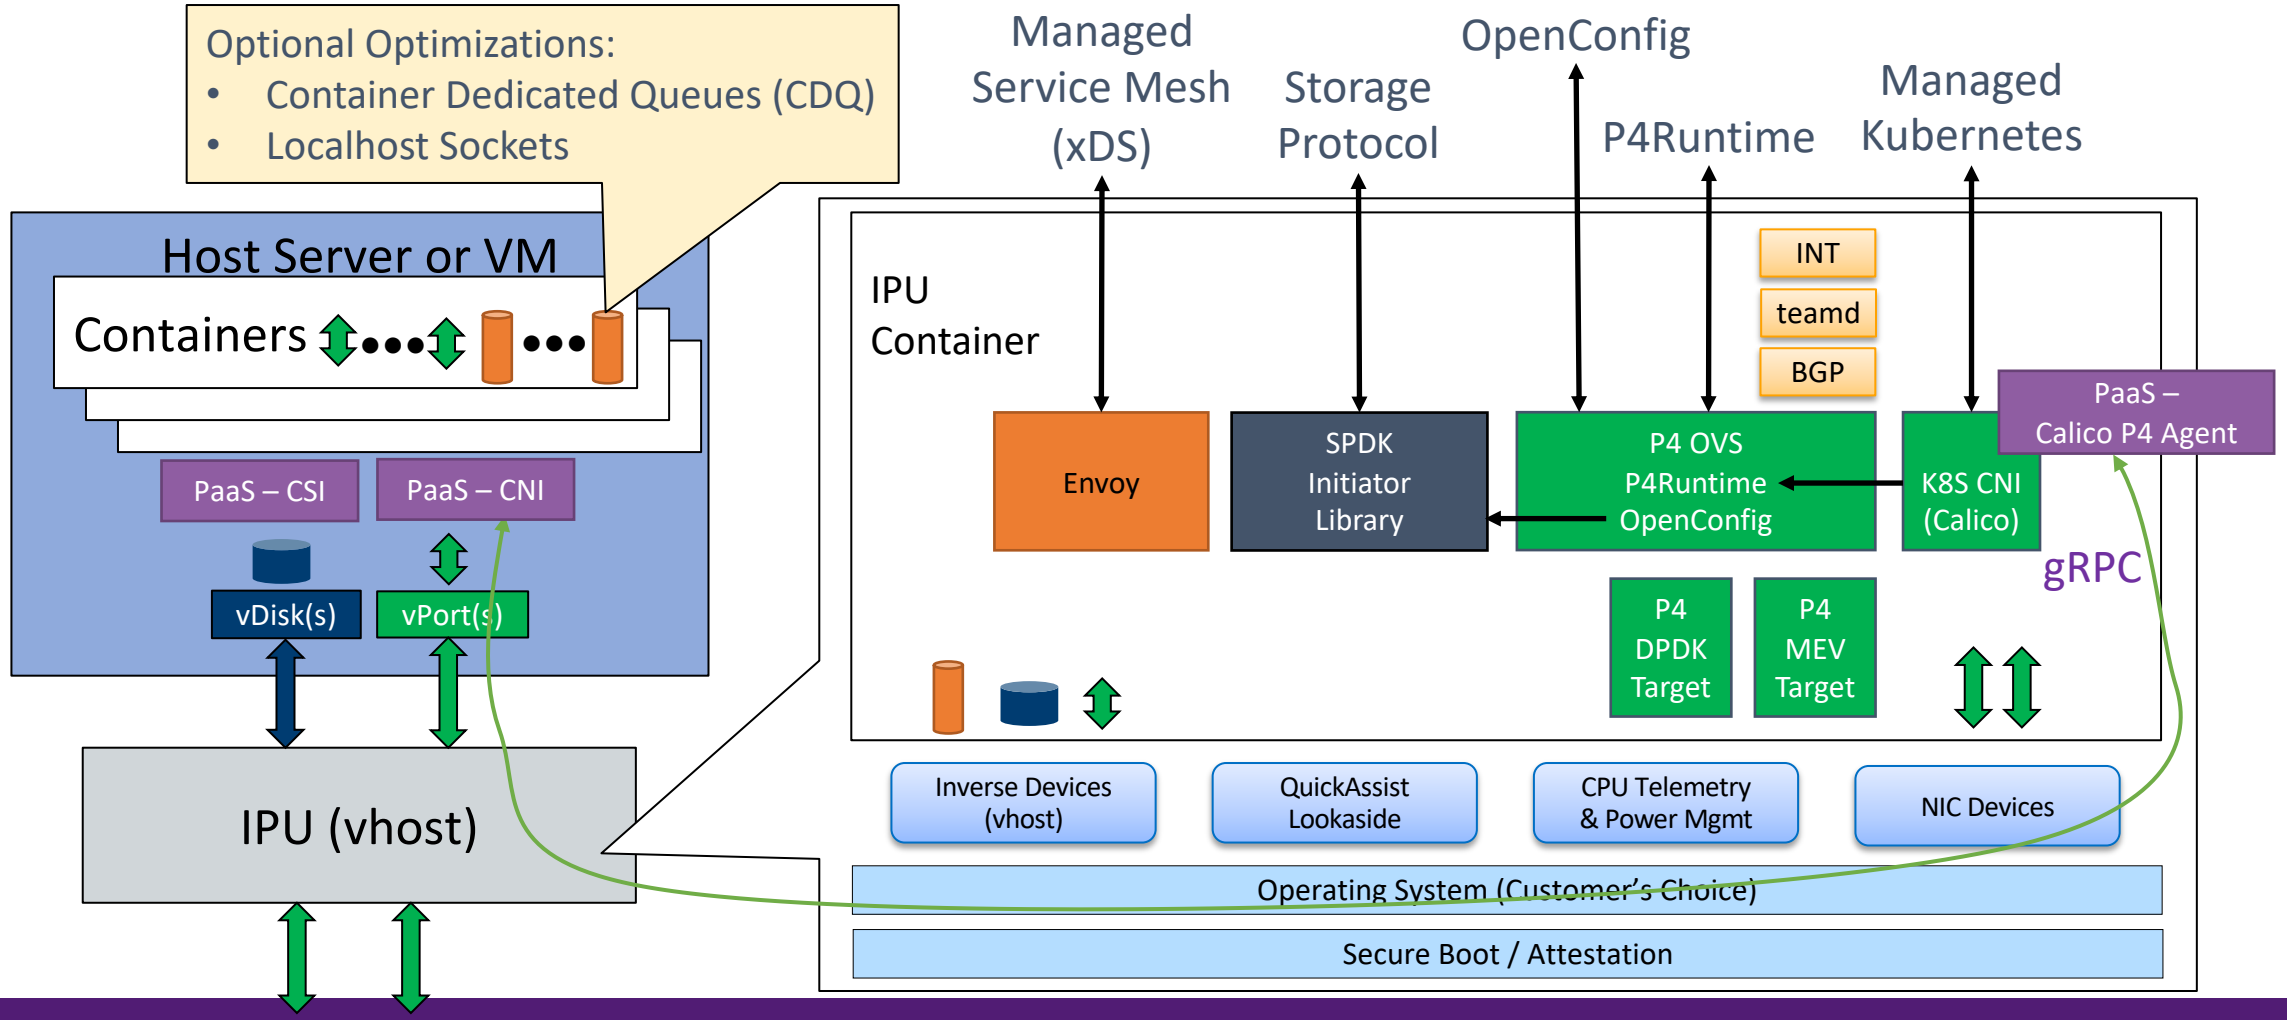

PaaS – Envoy Delineations

Current Service Mesh on Host

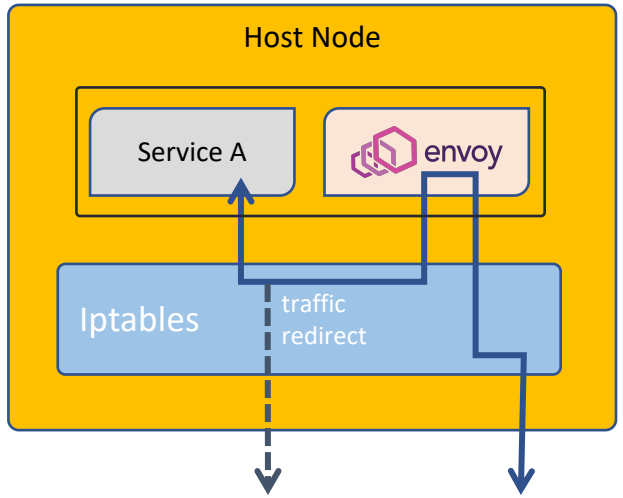

Service Mesh on IPU worker node

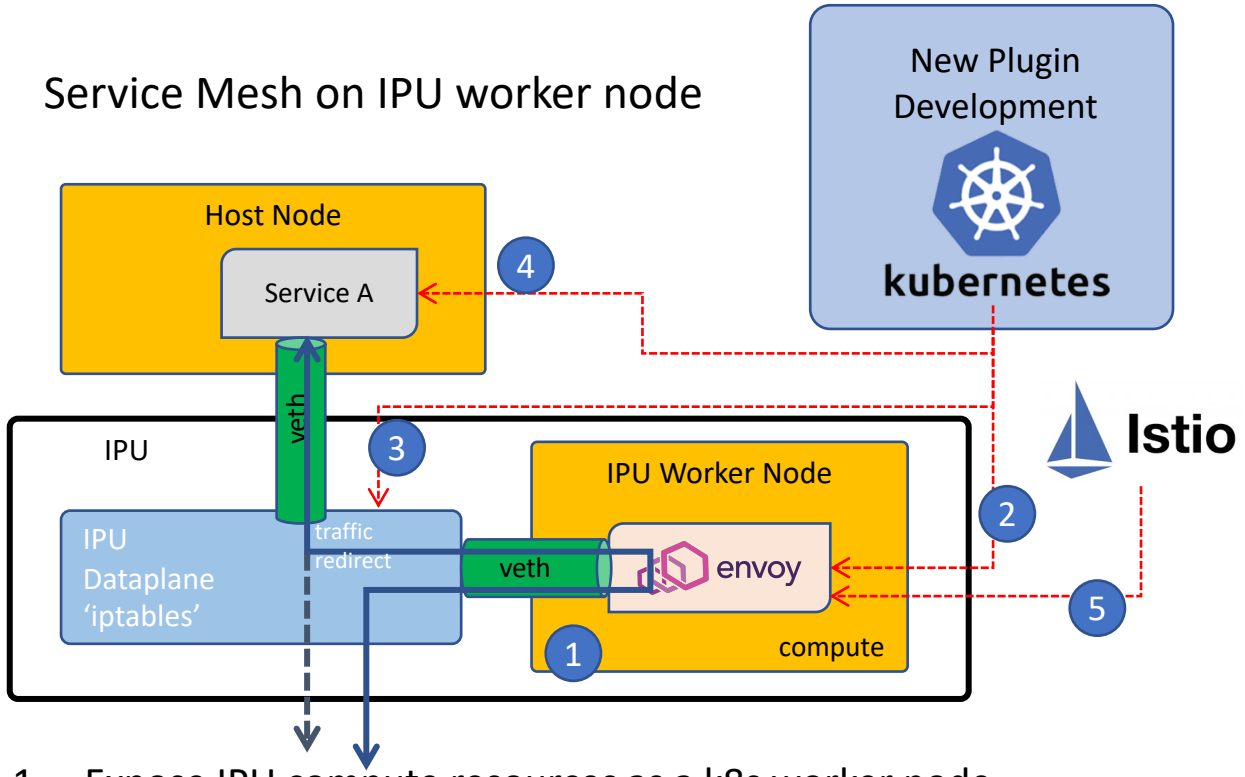

1. Expose IPU compute resources as a k8s worker node
2. Modify k8s orchestration to inject Envoy to IPU Worker Node
3. Configure IPU dataplane offload for 'iptables' traffic redirect
4. Orchestrate tenant service
5. Istio configuration of Envoy

PaaS – Envoy on IPU

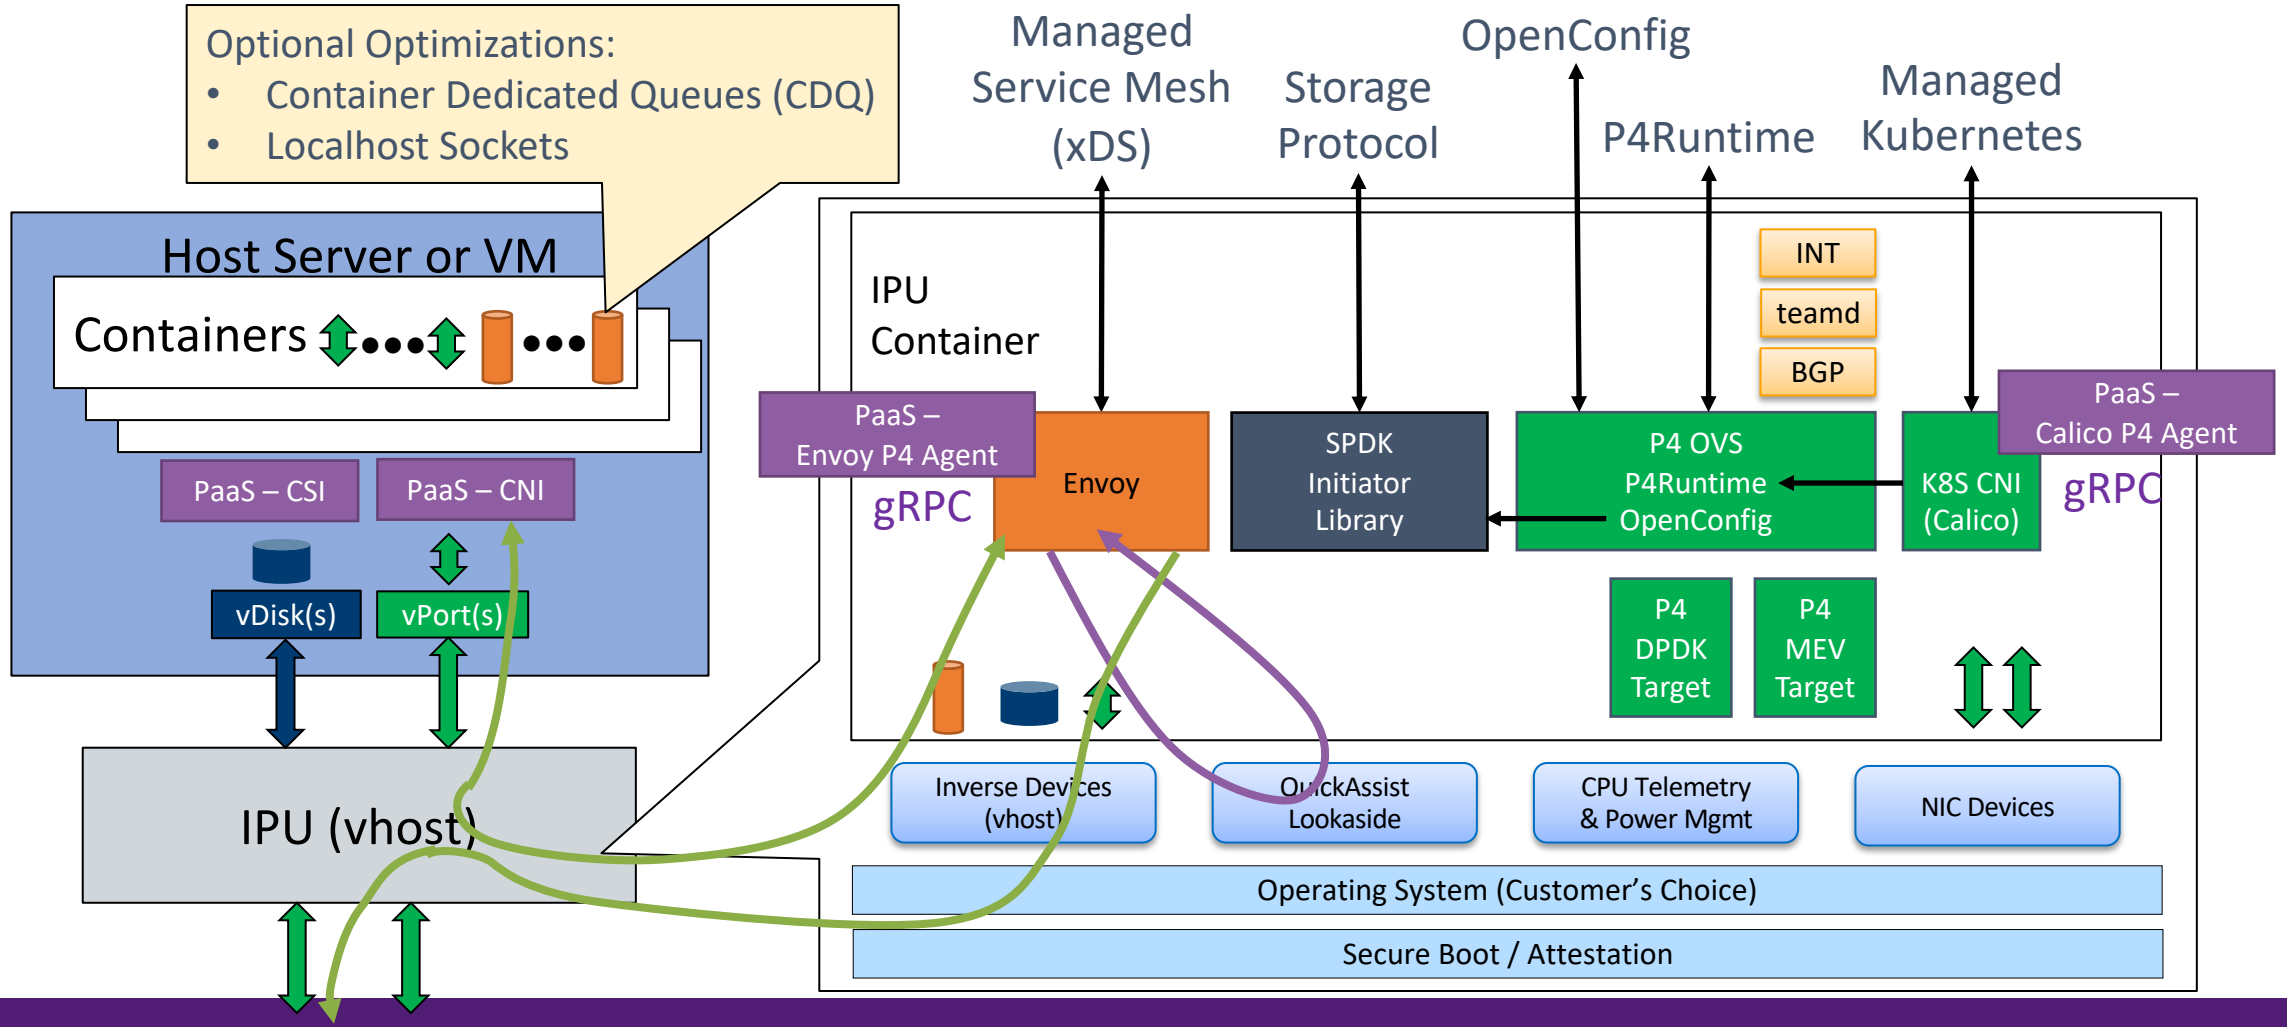

PaaS – Multi-Tenancy (Service:Sidecar)

- IPU dataplane used for traffic redirection to Envoy proxy and offload of CNI
- 1:1 Service Mesh (Envoy) containers
  - Resource allocation & life cycle
  - CDQ enable scheduling & BW allocations
  - Separation of mTLS Certificates
- No restriction on Istio deployment
  - Support various deployments such as multiple Envoy instances, soft multi-tenancy, & RBAC namespaces.
- Additional services (OpenTelemetry) can be orchestrated
- Optional to use Istio soft multi-tenancy

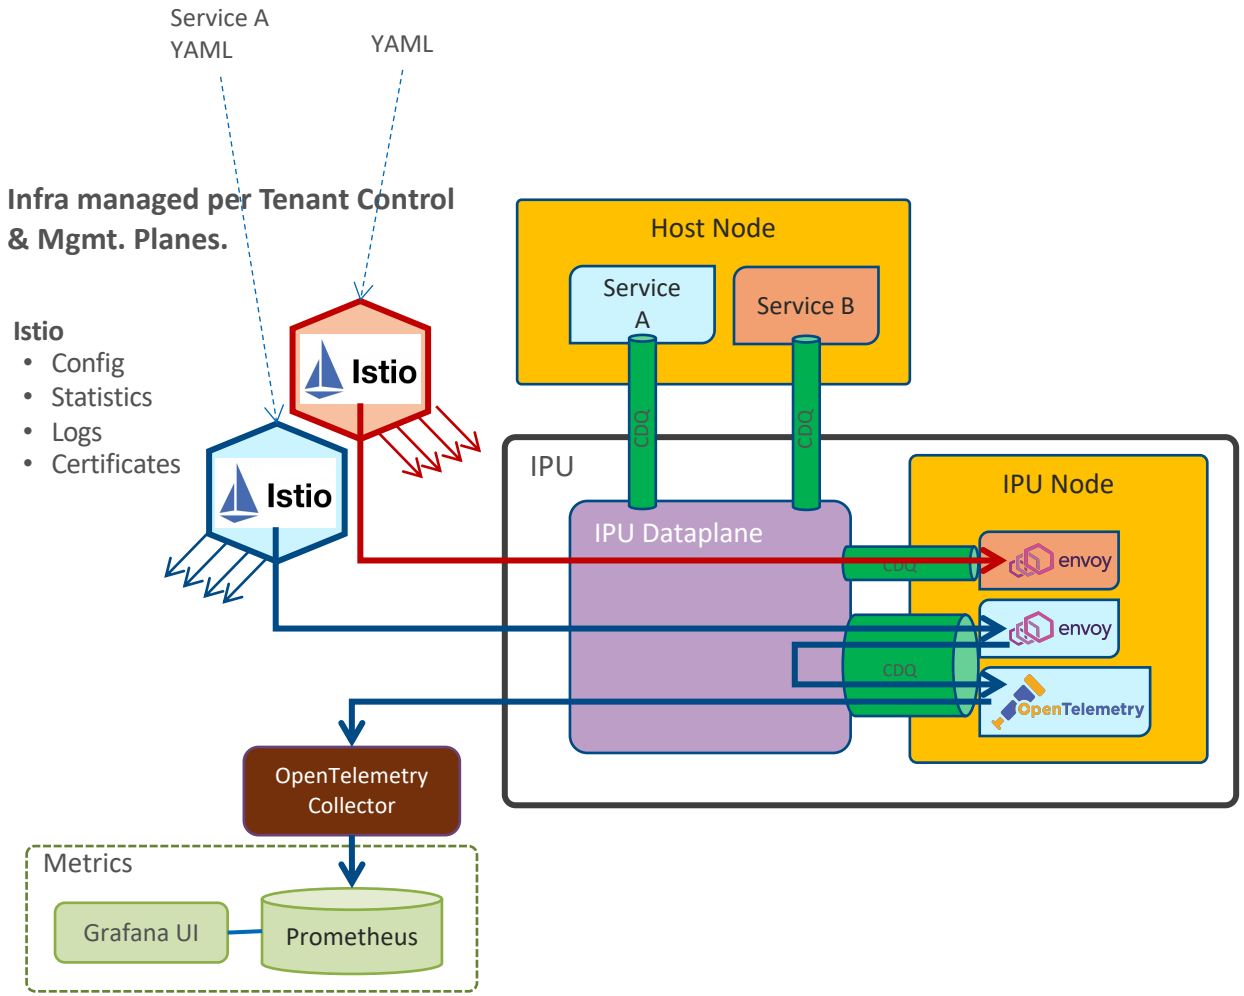

Thank you.

IPDK.io: Infrastructure Programmer Development Kit
Collaborate with the community on Github & Slack

ipdk