STORAGE DEVELOPER CONFERENCE
SDC 22
Fremont, CA
September 12-15, 2022
BY Developers FOR Developers

A SNIA Event

# Is Storage Orchestration a Headache? Try Infrastructure Programming

Jarosław Kogut
Tomasz Zawadzki

# Introductions

Jaroslaw Kogut
Technical Lead, Intel

Tomek Zawadzki
SPDK Core Maintainer, Intel

 

# Notices and Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing on certain dates using certain configurations and may not reflect all publicly available updates.  Reach out to Intel for configuration details.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.



 

# Infrastructure Processing Unit (IPU)

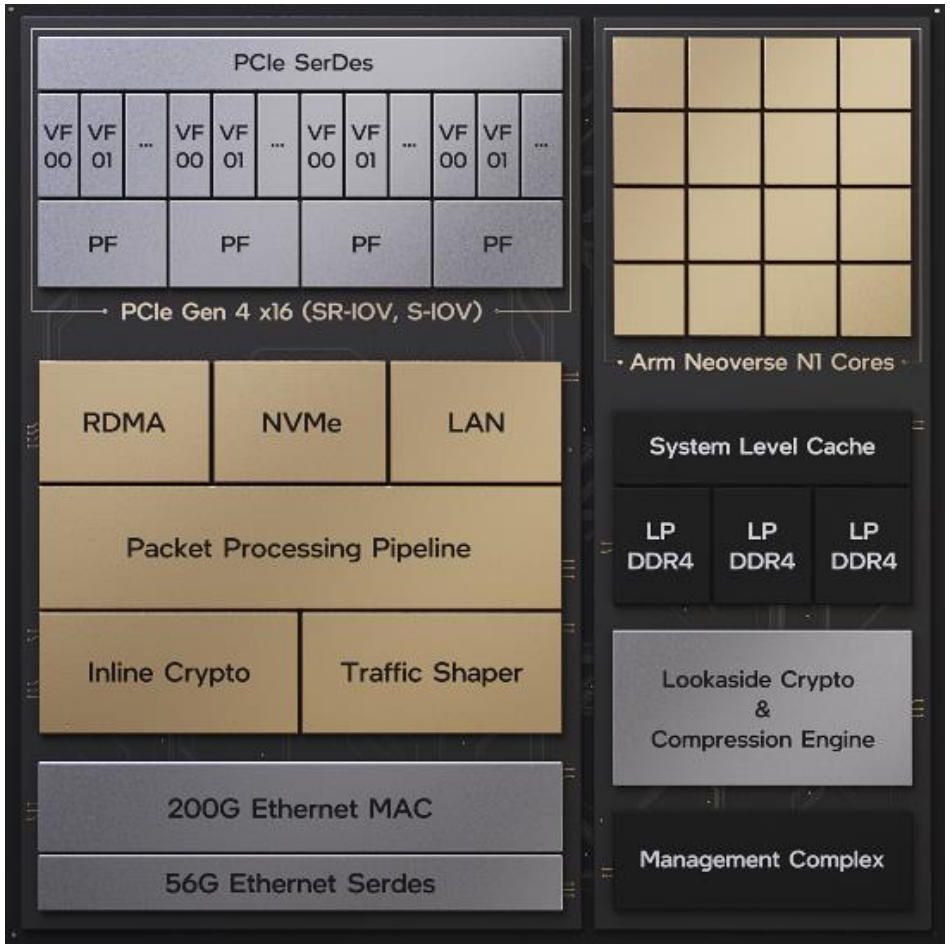


IPU ES2000

- A new class of device in data center
  - Instead of a standard NIC
  - Example: Intel® IPU ES2000
- Advantages and characteristic
  - Separation of customer workloads from infrastructure workloads
  - Programmable HW and SW
  - Support of storage use cases

 

# Storage Use Cases for Infrastructure Processing Unit (IPU)

 

# From standard NIC to IPU

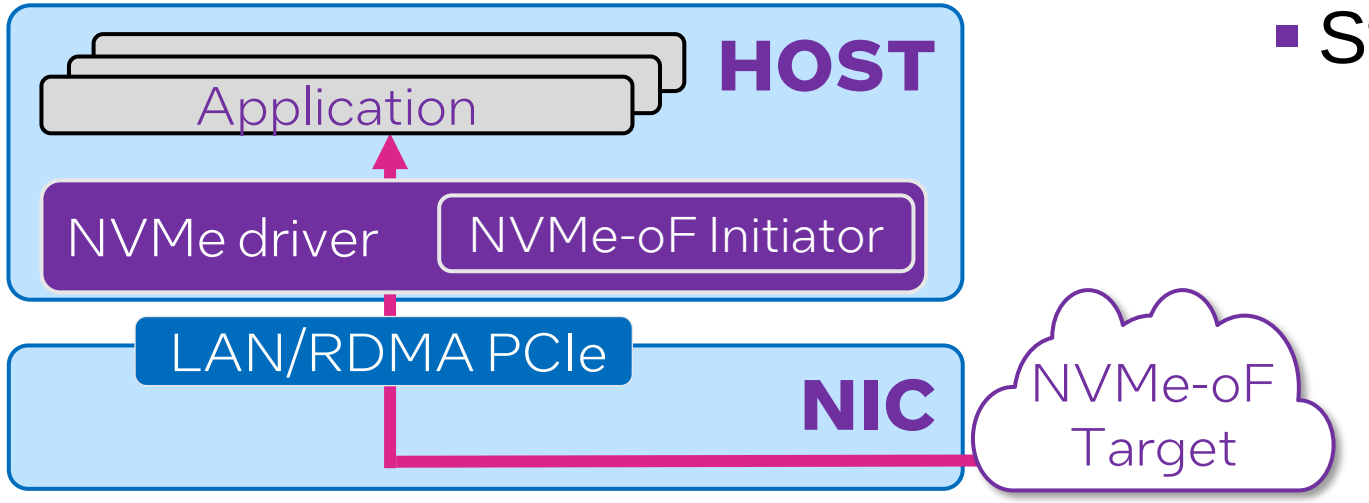


- Standard NIC
  - Host CPU cycles are consumed by NVMe-oF software stack
  - NVMe-oF software stack is in Linux kernel, Windows, user space (SPDK) on the host

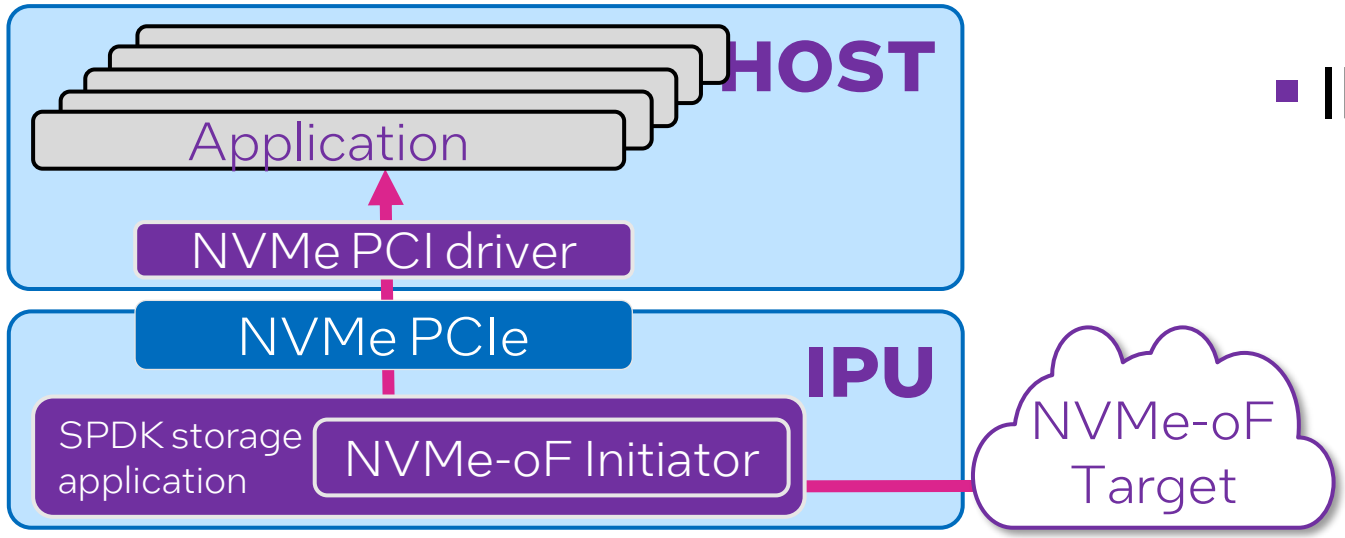


- IPU
  - Frees host CPU for customer workload
  - Host treats IPU as NVMe Controller connected over PCI

 

# Bare Metal Cloud

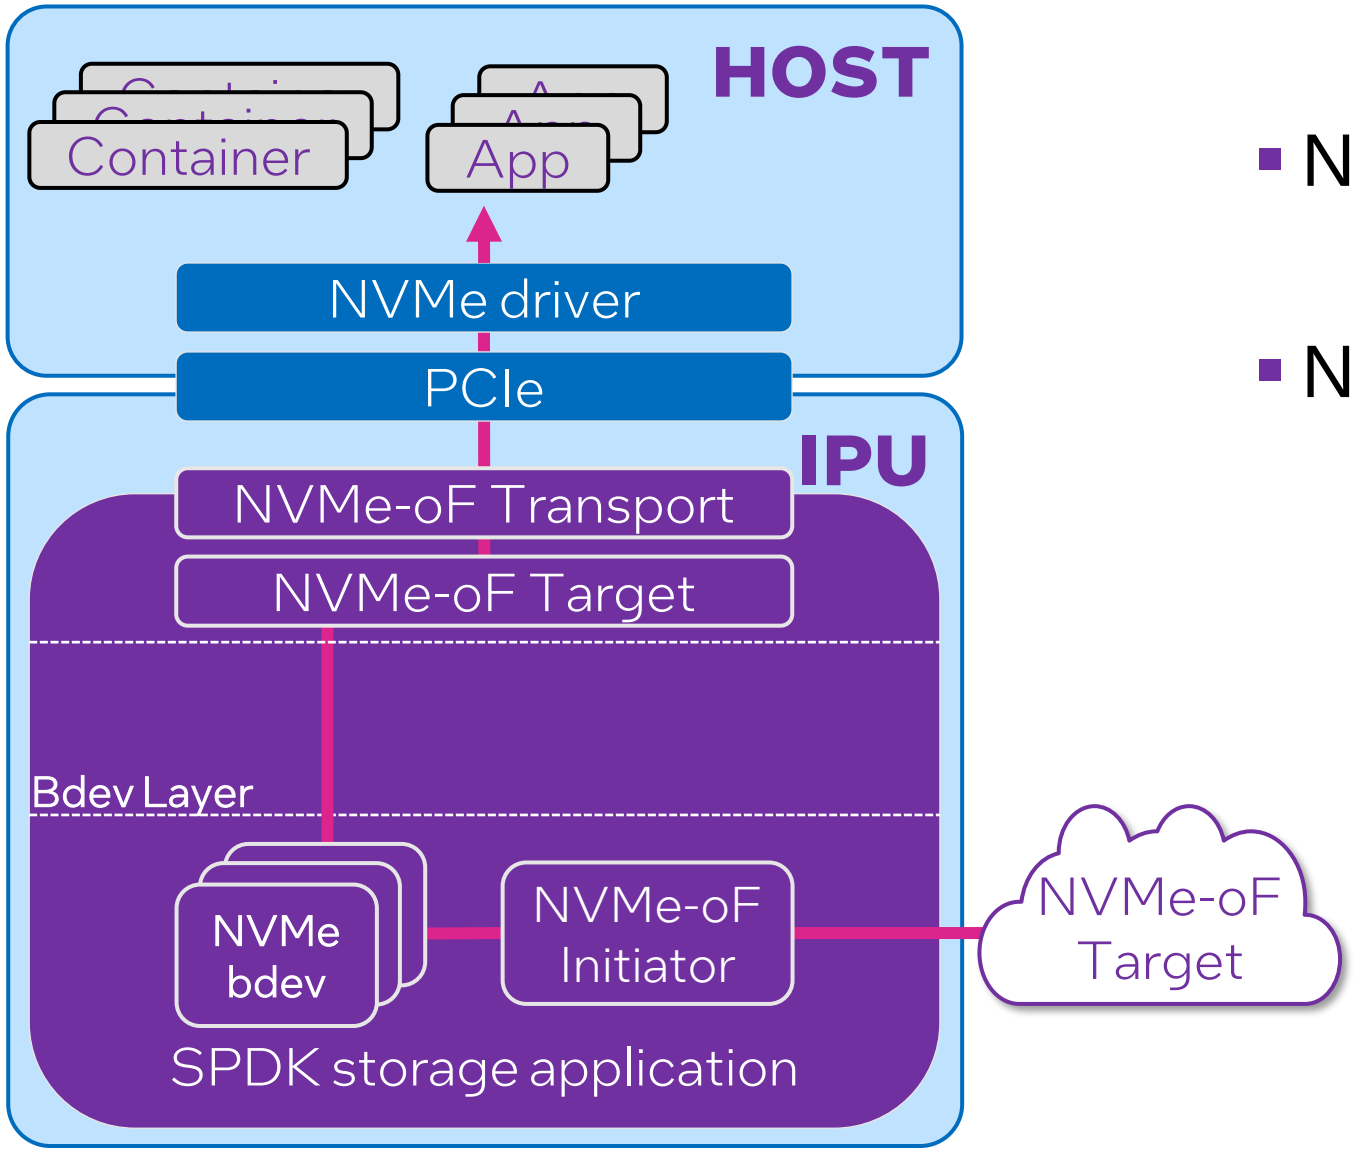


- NVMe device presented to host
- No NVMe-oF Initiator on host

 

# Virtualized Cloud - NVMe

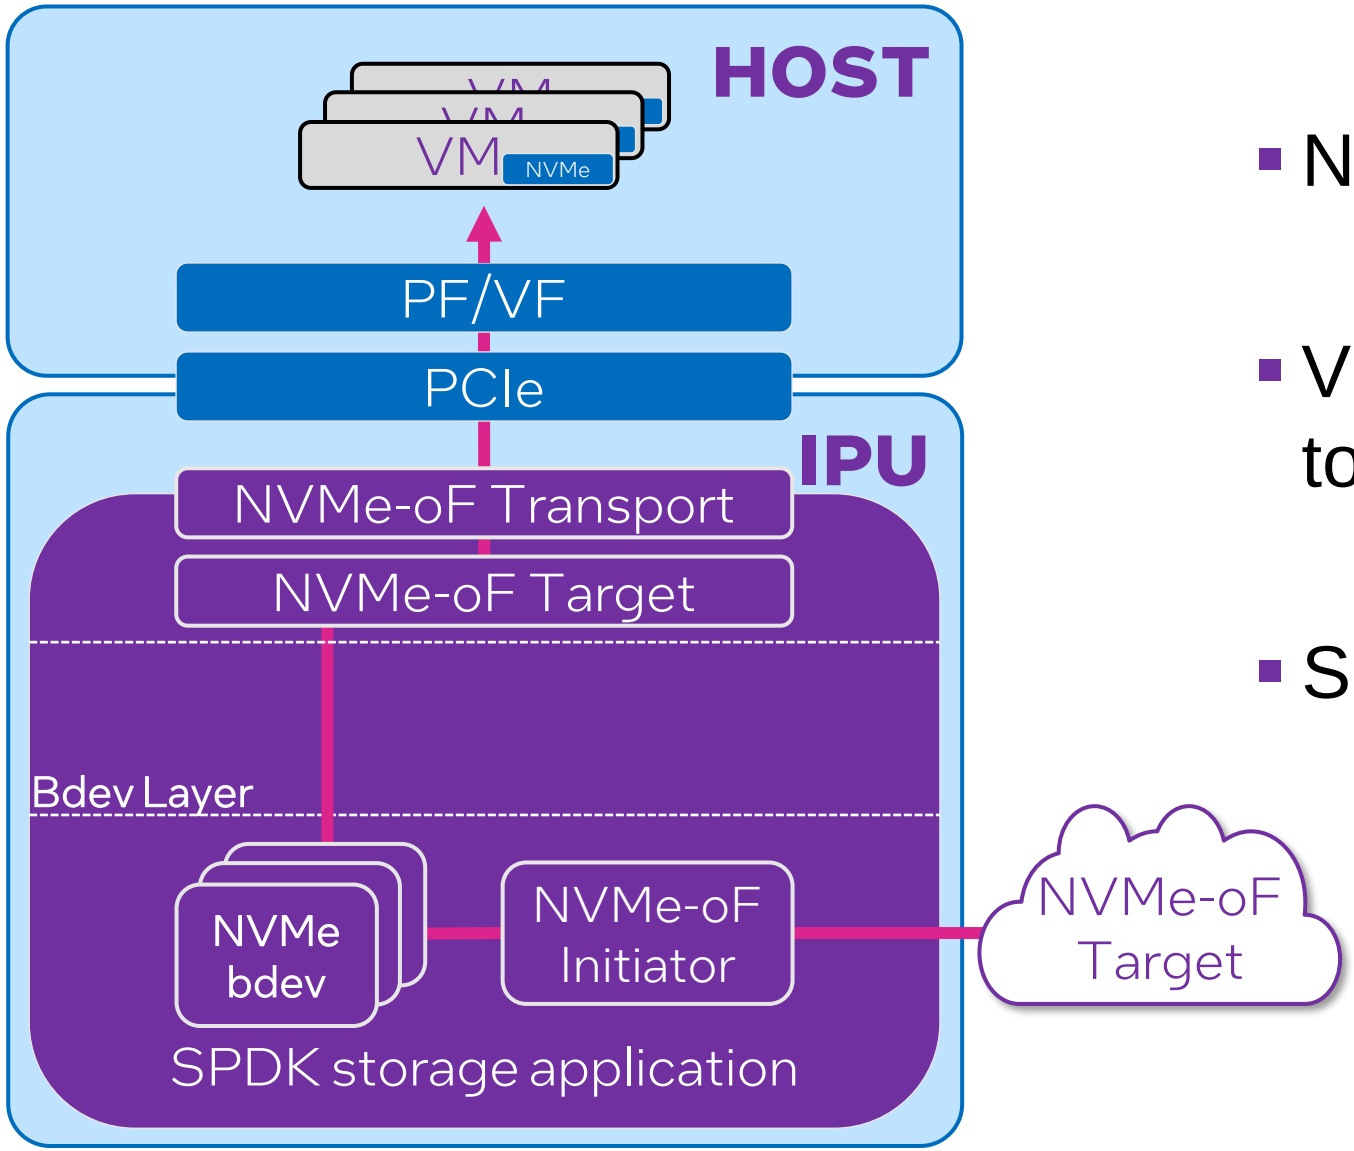


- NVMe device presented to VMs
- Virtual Functions attached directly to VMs
- Support for Live Migration!

 

# Virtualized Cloud – Virtio

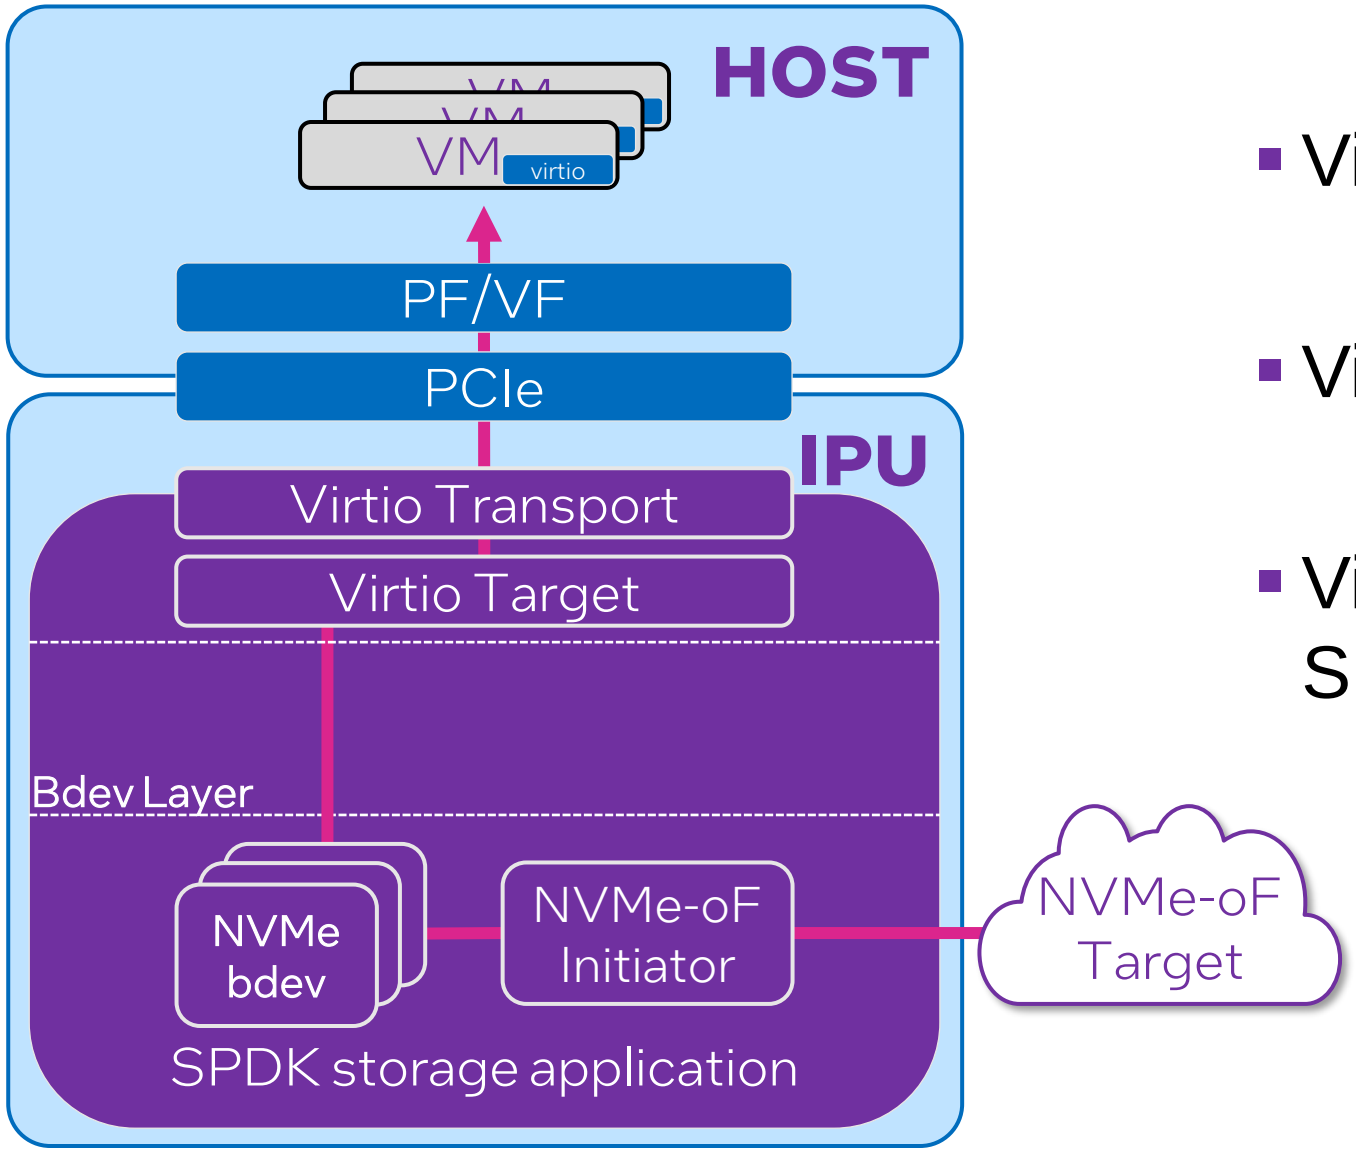


- Virtio device presented to VMs
- Virtio target used on IPU
- Virtio-blk abstraction added in SPDK 22.05

 

# Advantages of Managing Storage and Features on IPU

 

# The Problem of Scale

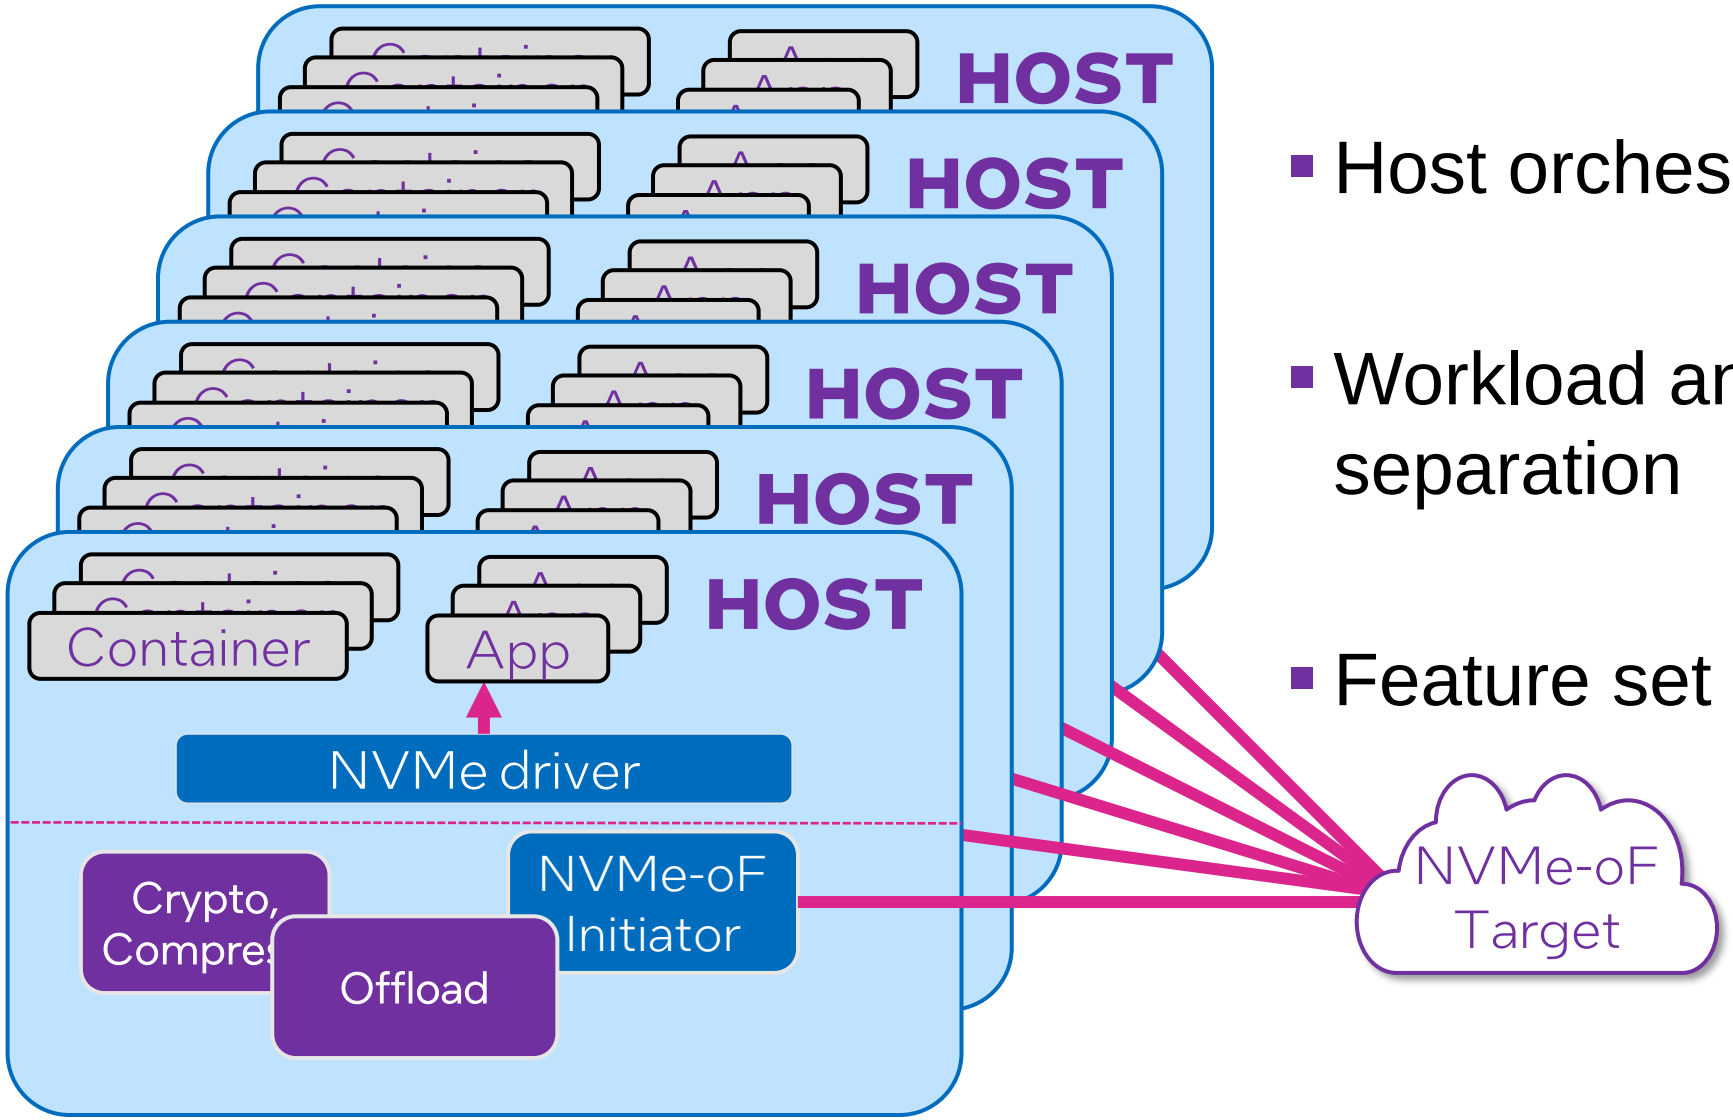


- Host orchestration
- Workload and infrastructure separation
- Feature set evolution on IPU

 

# Features Moved from Host to IPU

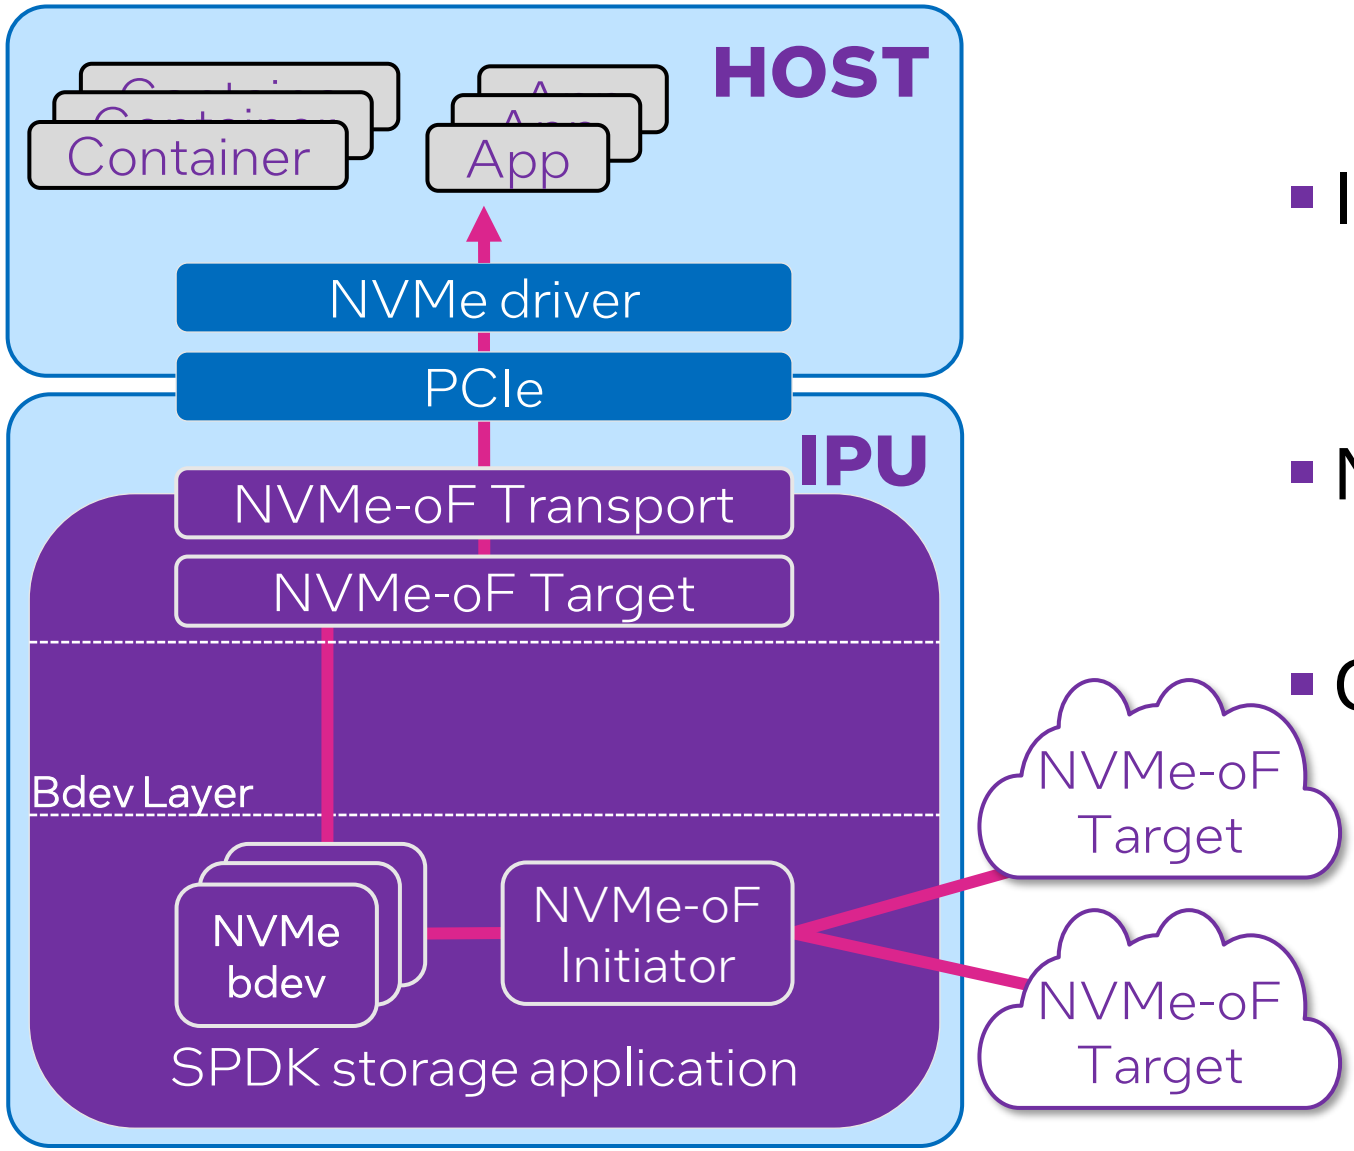


- Interchangable backends
  - NVMe-oF, iSCSI, more…
- NVMe-oF multipath and discovery
- Quality of service

 

# Storage Acceleration

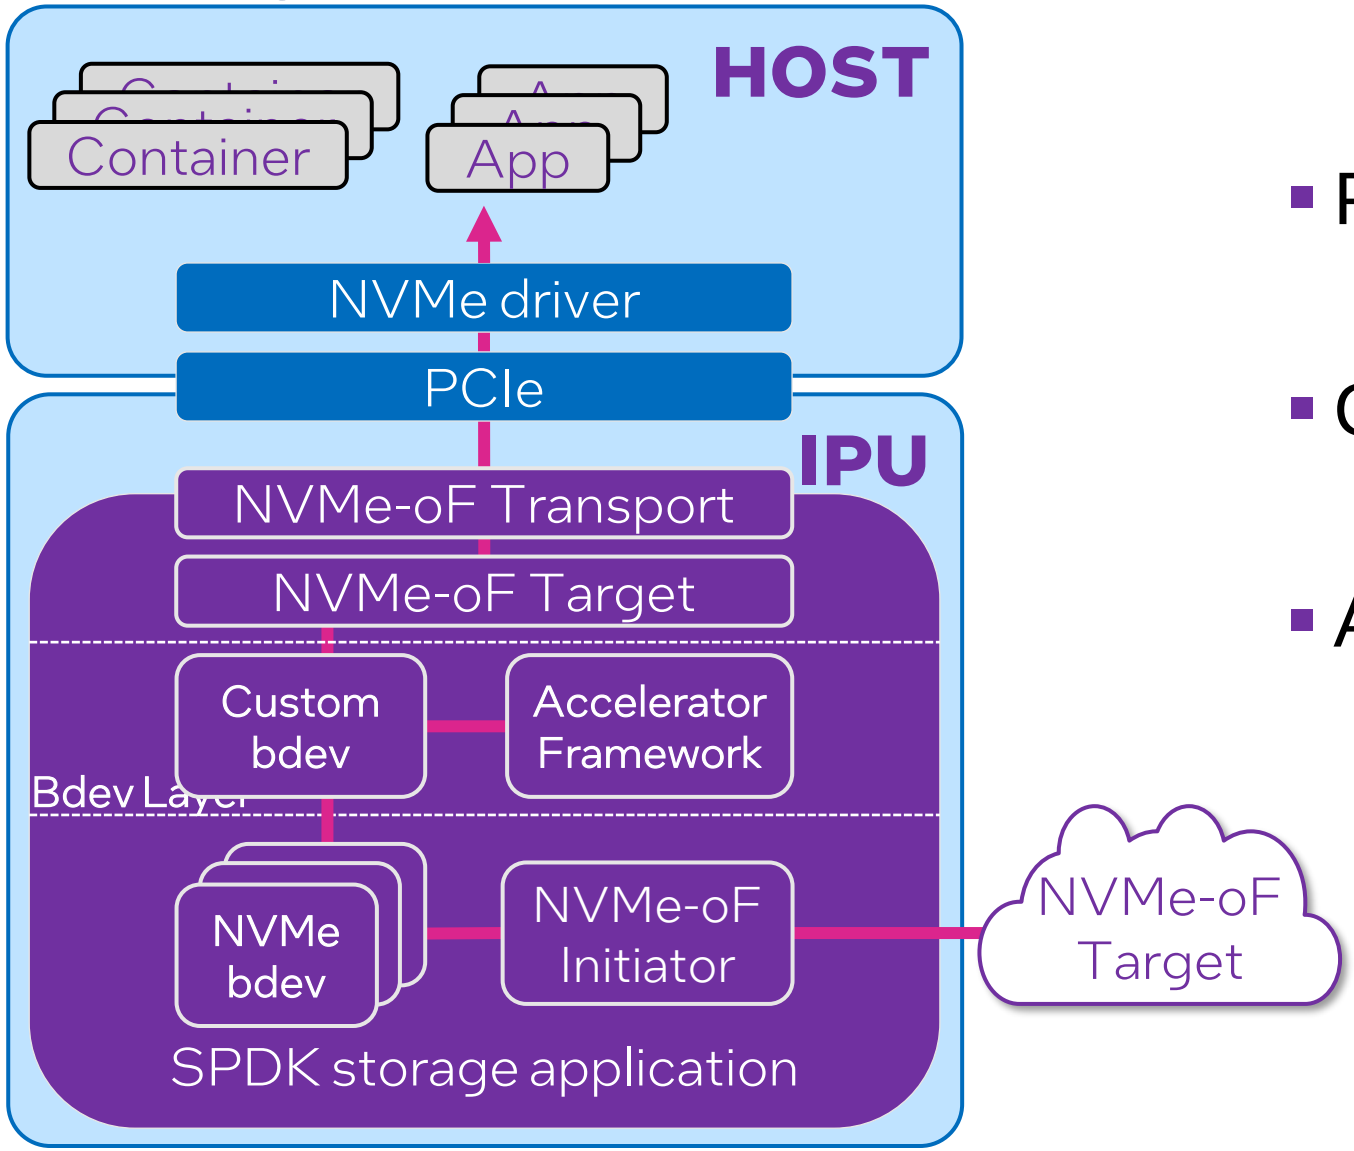


- Programable HW via SW
- Custom bdevs
- Accelerator Framework
  - Encryption
  - CRC
  - Compression

 

# Target Agnostic Frameworks for Storage Infrastructure Programing

Storage Performance Development Kit (SPDK)
Infrastructure Programmer Development Kit (IPDK)
Open Programmable Infrastructure (OPI)

 

# Storage Performance Development Kit Architecture

- User-space Tools, Libraries, Drivers, & Applications
- Open Source & BSD Licensed
- Optimized for bleeding edge storage solutions

**Participate/Learn More** [https://SPDK.io](https://SPDK.io)

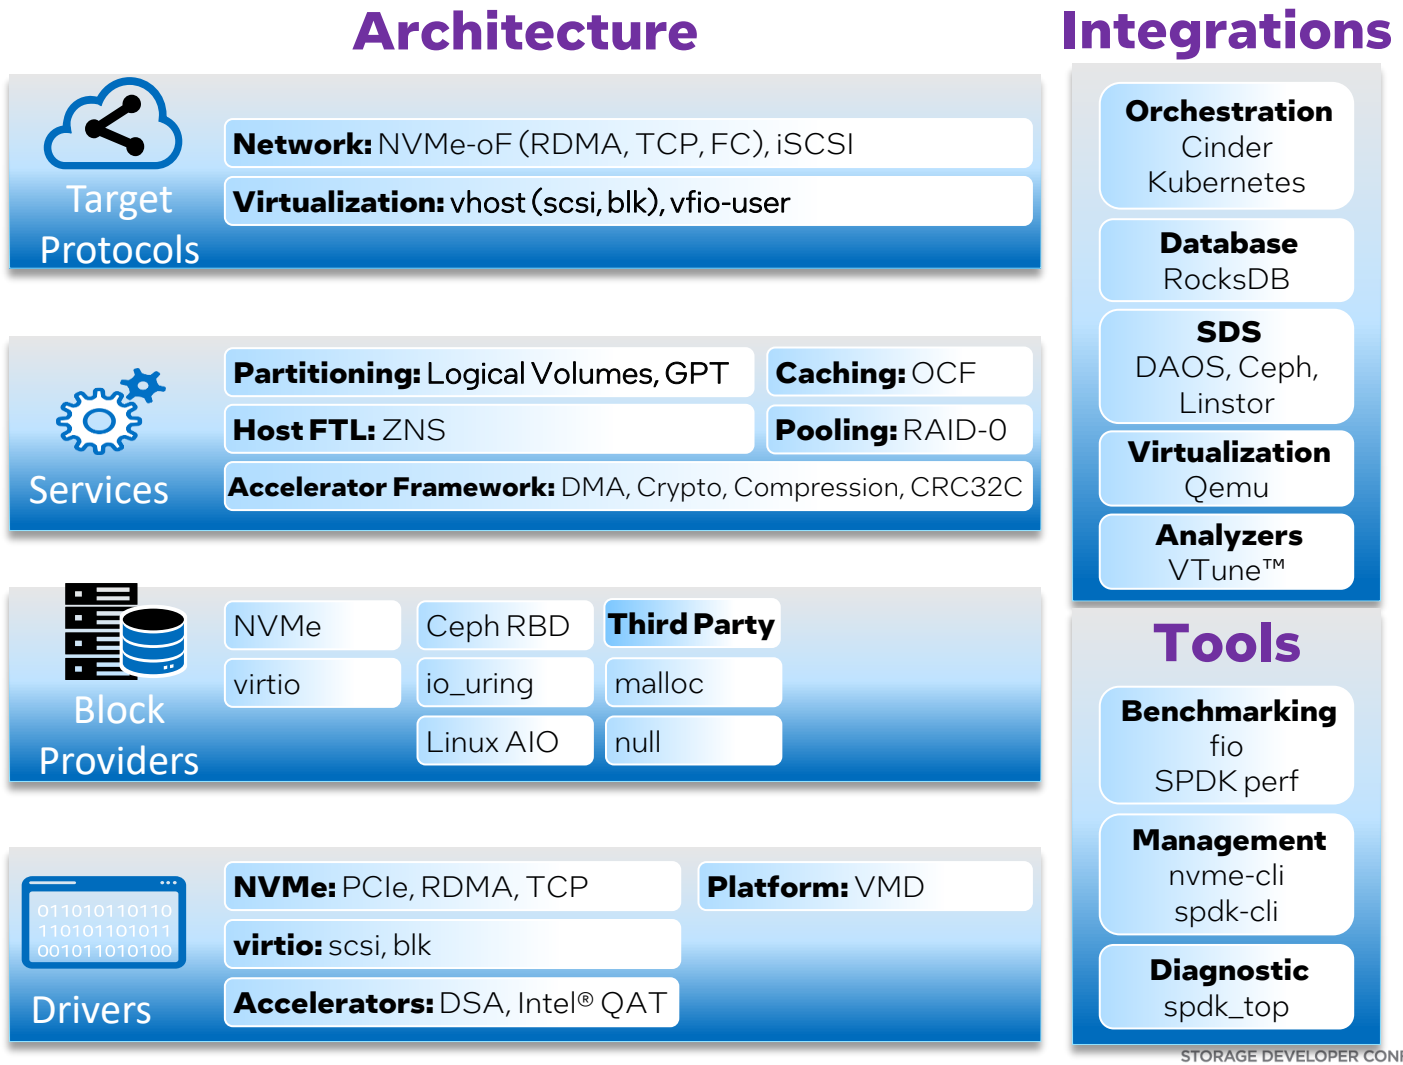


 

# Storage Management Agent (SMA)

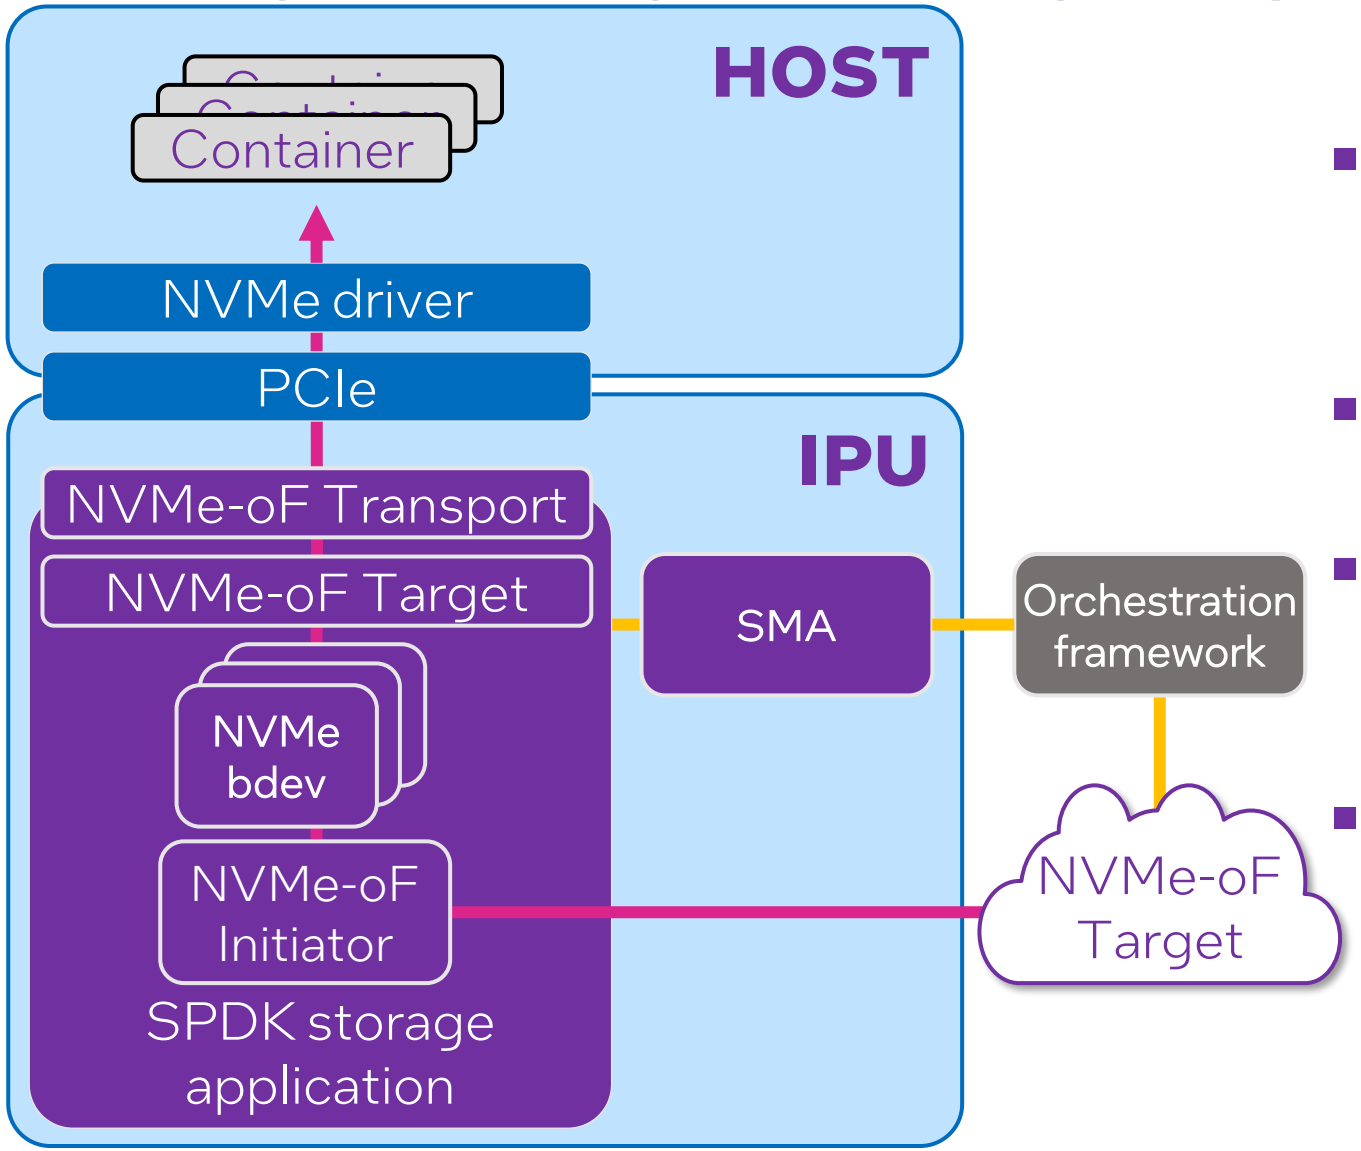


- Service running on an IPU managing SPDK application
- To hide details of any given IPU
- Can be adopted to work with any orchestration framework
- Added in SPDK 22.05, supports:
  - NVMe host interface
  - Virtio-blk host interface
  - NVMe-oF volumes

 

# IPDK Overview

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

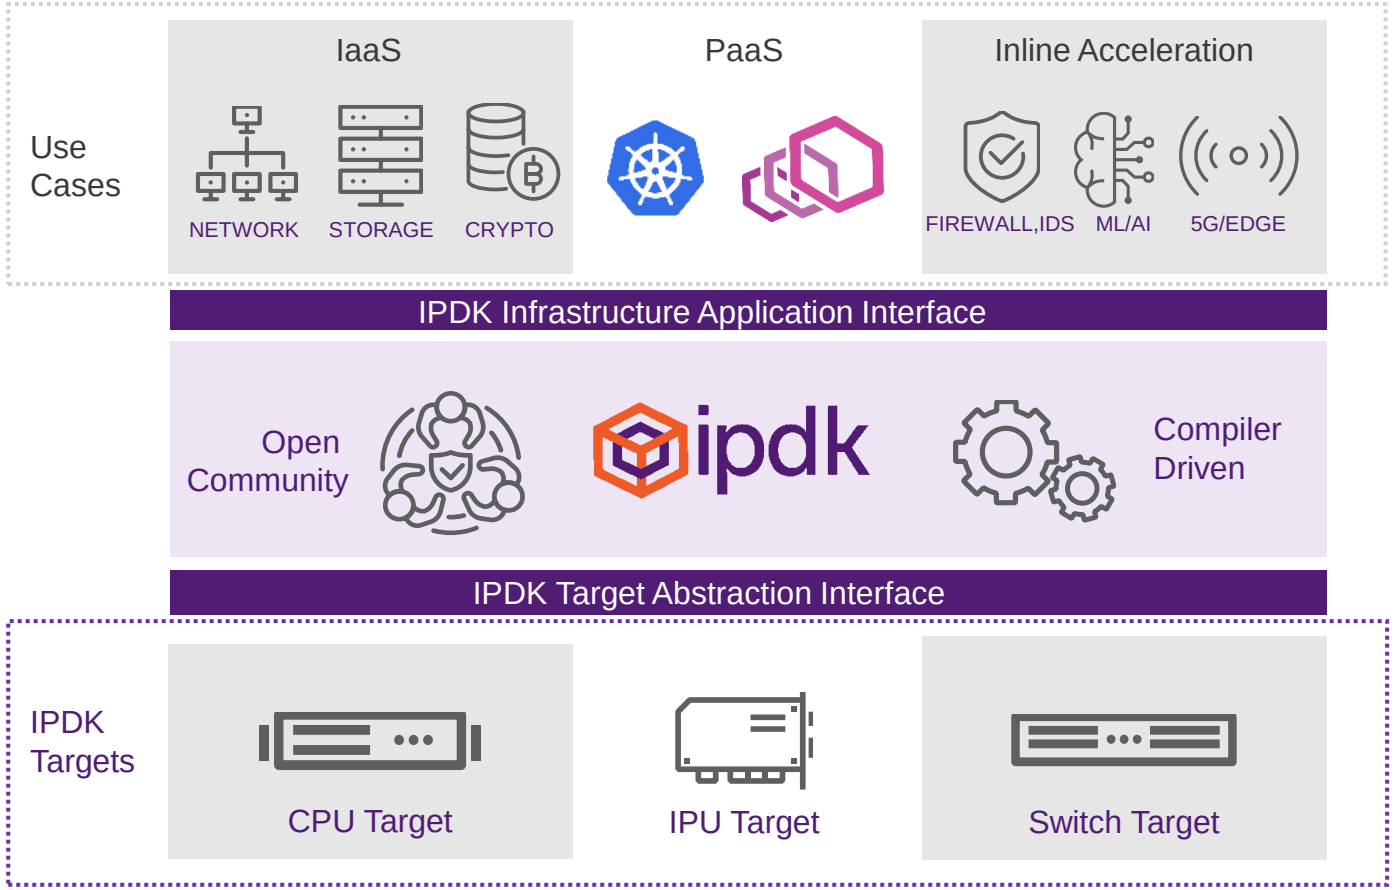


 

# OPI overview


- Linux Foundation project
- Motivation
  - Approach for all devices: DPU, IPU
  - Drive open standards for new class of cloud native infrastructure
- What else
  - Common programing model for infrastructure devices
  - OPI sub-groups e.g., OPI API, Use Cases

 

# OPI, IPDK and storage topics

- IPDK 22.07
  - Recipes how to manage IPU like devices
    - Refence environment running target, IPU and host in containers
    - Host: virtio-blk device
    - IPU software: SPDK application with SMA API
    - Target: NVMe-oF TCP target


Website : https://ipdk.io
GitHub : https://github.com/ipdk-io
Slack established

- OPI (announced June 21, 2022)
  - OPI API for storage discussion (Slack/GitHub)


Website : https://opiproject.org
GitHub : https://github.com/opiproject
Slack established

 

# OPI, IPDK next steps

## IPDK and OPI coexistence

- IPDK is official sub-project of OPI
  - Develop everything under OPI umbrella
- Releases
  - Drive to establish OPI releases
  - IPDK & OPI releases shall be aligned
- API
  - OPI shall drive the definition of APIs

**Opportunity:**
Just established initiatives to drive open and creative software ecosystem for the new class of devices like DPU/IPU.

## Storage topics

- More use cases
  - Bare metal, virtualization, containers
  - NVMe, virtio-blk, …
- Orchestration examples
  - CSI plugins
- Recipes
  - more examples how to manage IPU/DPU like devices
- API
  - OPI API, SPDK SMA API

 

# Key Takeaways

 

# What About That Headache

## Hardware

- IPU/DPU like devices allow to move infrastructure software/solutions from host cores to IPU cores/hardware

- Infrastructure software may be easy extendable without impacting consumers' software

## Software

- Can we manage IPU/DPU like device in common way?
- Can we deliver software IPU/DPU that will be managed by common software?

- Leverage benefits of SPDK, IPDK/OPI
- Participate/influence these communities!

 

# Please take a moment to rate this session.

Your feedback is important to us.

 