

STORAGE DEVELOPER CONFERENCE



Fremont, CA  
September 12-15, 2022

*BY Developers FOR Developers*

A **SNIA** Event

# OPI (Open Programmable Infrastructure) Overview

13-SEP-2022

Presented by

Joseph L White, OPI TSC Chair, Dell Fellow/VP

Venkat Pullela, OPI TSC, Chief of Technology, Networking, Keysight



**OPEN  
PROGRAMMABLE  
INFRASTRUCTURE  
PROJECT**



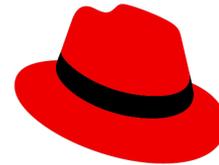
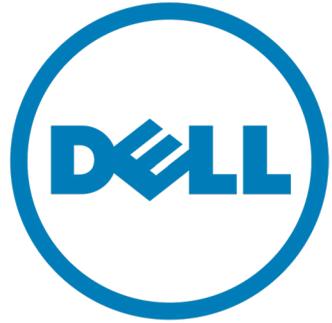
*The objective of the Open Programmable Infrastructure Project is to foster a community-driven **standards-based open ecosystem** for next generation architectures and frameworks based on **DPU and IPU technologies**.*

<https://opiproject.org>

<https://github.com/opiproject>



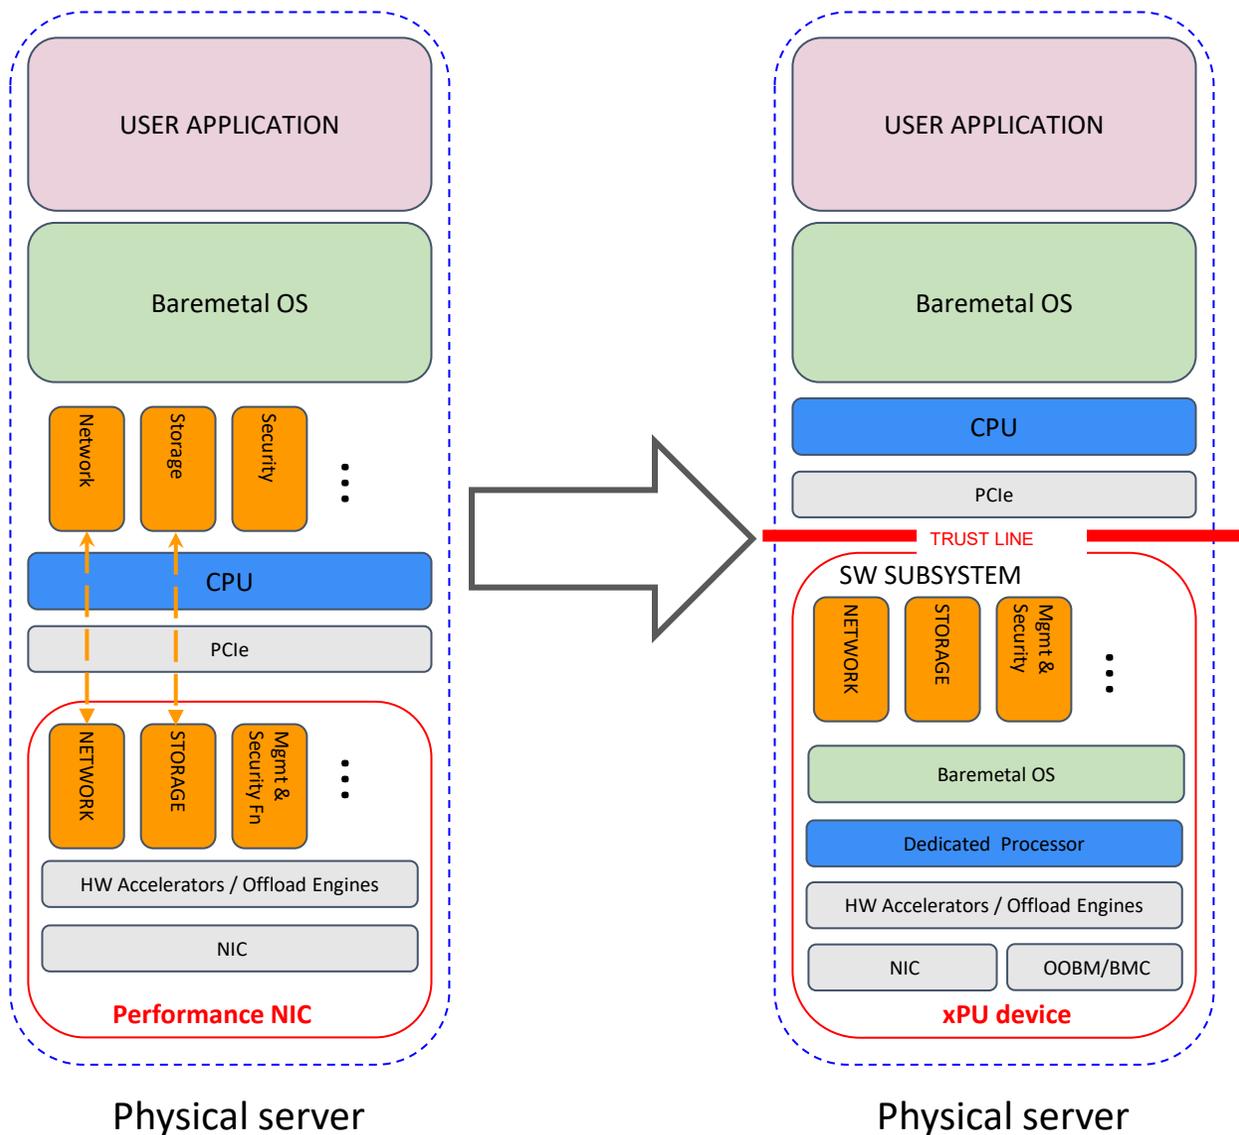
Joins



**Founding Members listed**

*Many other companies and individuals participating & contributing*

# Infrastructure Transformation



## xPU models

- Offload and Accelerate CPU functions
- Security Isolation
- Independent infrastructure endpoint

# Project Goals

## Open-Source and Standards for xPU/DPU/IPU Technologies

- Community-driven standards-based **open ecosystem**
- Vendor agnostic framework and architecture
- Define new APIs and Standards where needed
- Reuse already existing APIs and Standards
- Provide implementation examples + Reference Platform

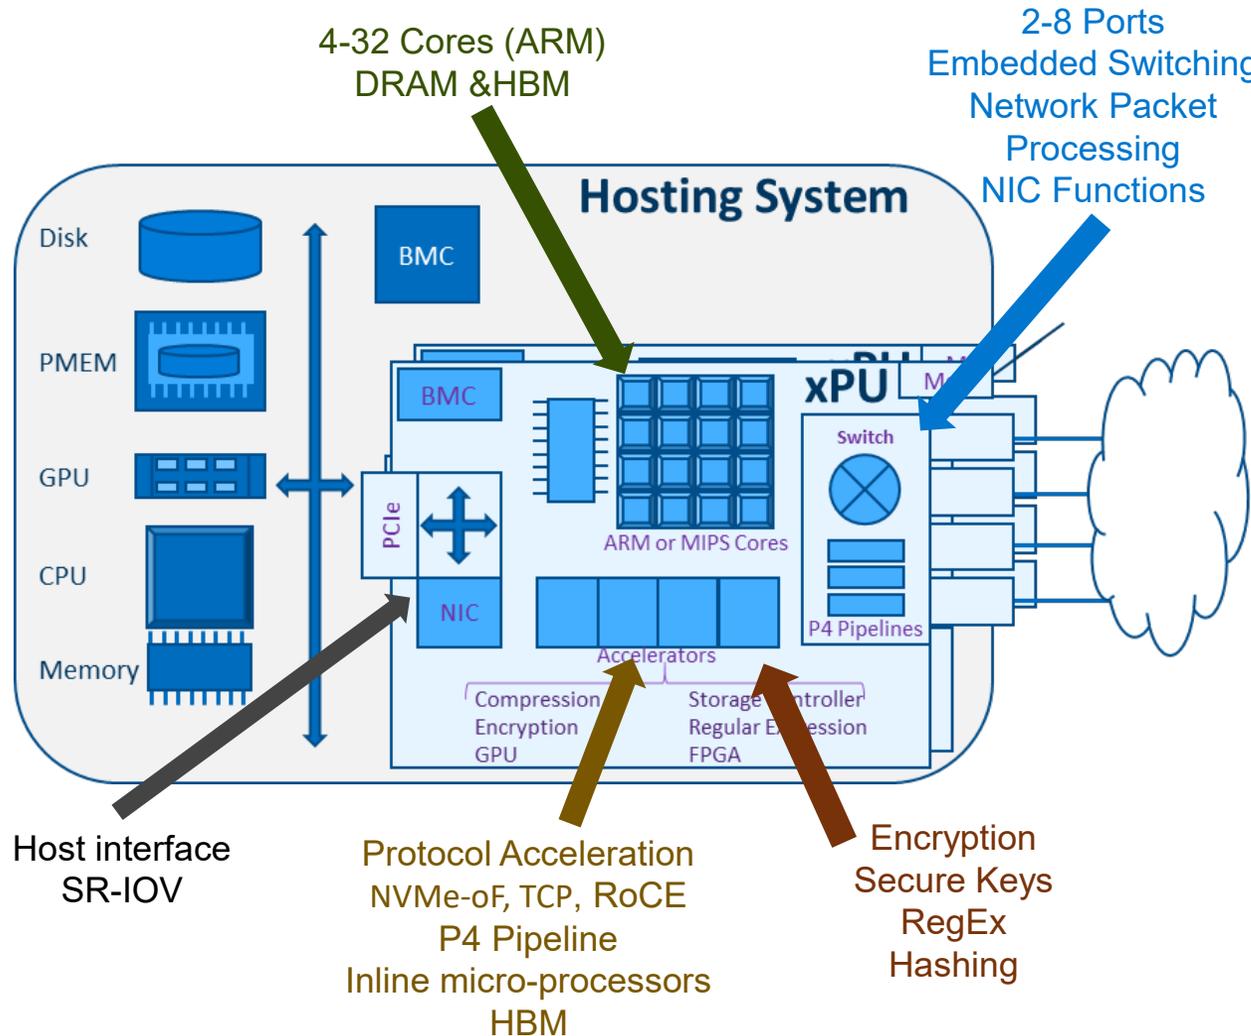


# SNIA Dictionary Definition

**DPU:** An accelerator element capable of parsing, processing, and transferring data with performance efficiency.

A Data Processing Unit (DPU) usually has a set of programmable acceleration engines that offload and improve performance for applications such as AI/ML, security, telecommunications, and storage. DPUs may also be called SmartNICs, IPU's or NAPUs.

# DPU Definition Expanded



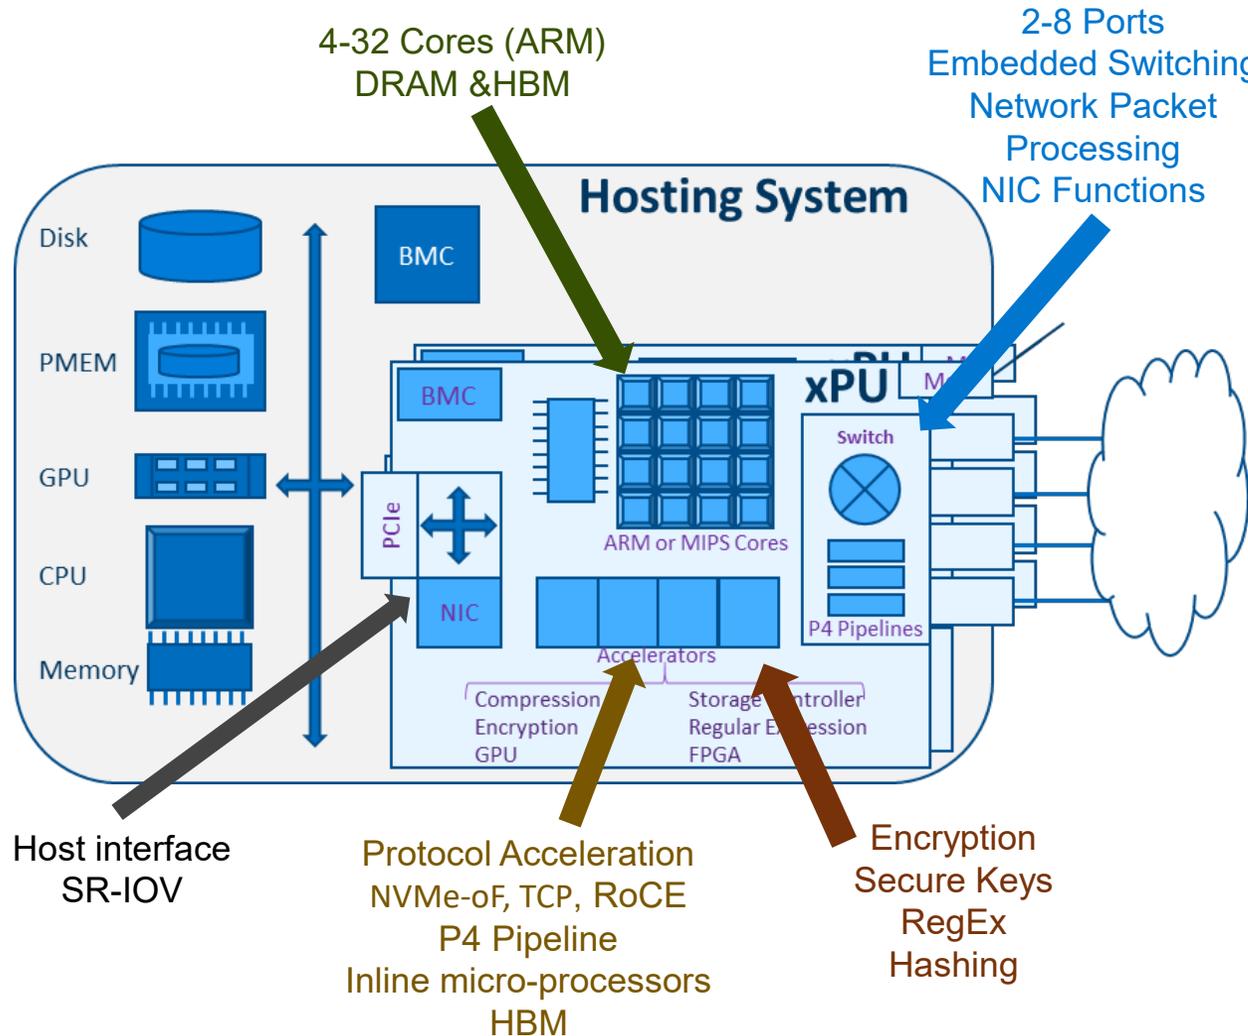
## DPU - Data Processing Unit (aka xPU)

Effectively a micro-server optimized for dataflow and packet processing providing accelerators, offload engines, & local services

Presents virtual functions to a host (looks like a NIC, GPU, etc)

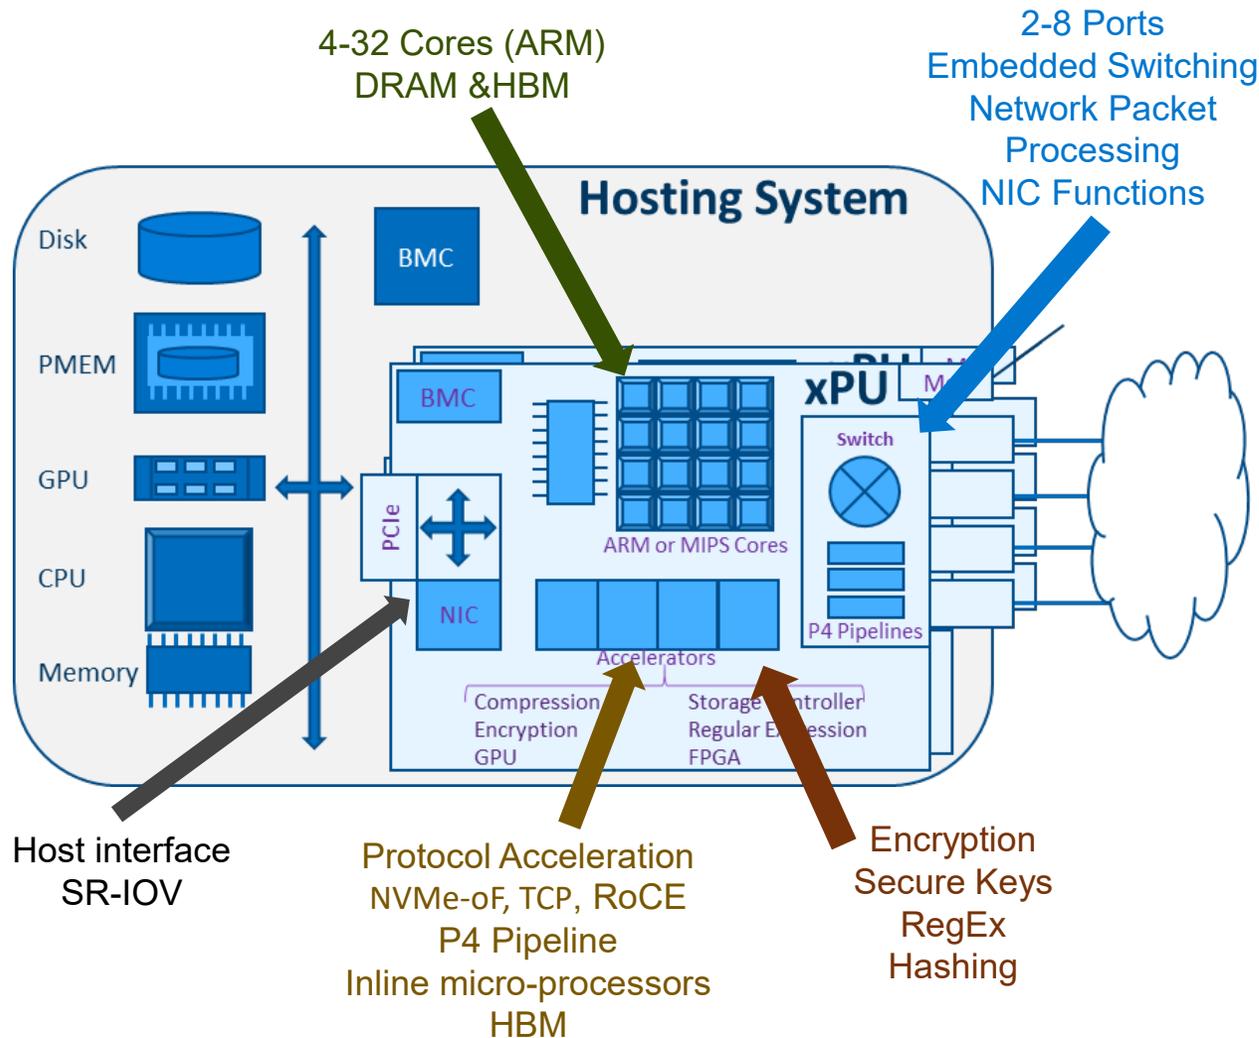
- DPU Internal Components
  - General Purpose CPU Cores with Memory
  - PCIe Interface with Local Switching
  - Network Interfaces (Data and Management) with Local Switching
  - Accelerators, Offloads Hardware, Programmable Pipelines
  - Embedded BMC
- Server Architecture
  - DPUs typically a built as a PCIe Card (>1 allowed)
  - Other instantiations like switch embedded or standalone possible
  - DPUs present conventional PCIe functions to hosting servers
  - DPUs can directly access PCIe Devices
- DPU Operating System
  - Linux (N flavors, Ubuntu/Debian is common)
  - VMware
  - proprietary
- Common Tool Chains Apply
  - System configuration and management
  - Network configuration and management
- K8s
  - container installation and management

# Key characteristics of DPU based architectures



- Capable of booting a general-purpose OS
- Domain-specific HW acceleration capabilities
- Software-defined device functions
  - allow the software components to be flexibly deployed
  - define the device's functions that are presented to the host
  - Offloading complete software subsystems, (eg Networking or Storage stack)
  - Control planes
- Security isolation from the host at the hardware-level
- Unique network identity
- Management
  - Capable of being managed as part of the hosting server (through BMC or hosting OS)
  - Capable of being directly managed (out-of-band) separately from the hosting server
  - Capable of managing the hosting server

# DPU Use Cases



## Common Industry Acceleration & Offload Use Cases

- Network Switching
- Network Connectivity
- Gateway
- Storage Connectivity including NVMe/TCP, NVMe/RoCE
- Storage Services
- Expose Hosting System Resources
- Security (Firewall, DPI, Key Management, Intrusion Detection/Protection, Host Isolation)
- Telemetry Collection and Processing
- Hypervisor
- CNF/NFV Hosting
- Provide Accelerators/Co-processing to Hosting system
- Boot and provisioning

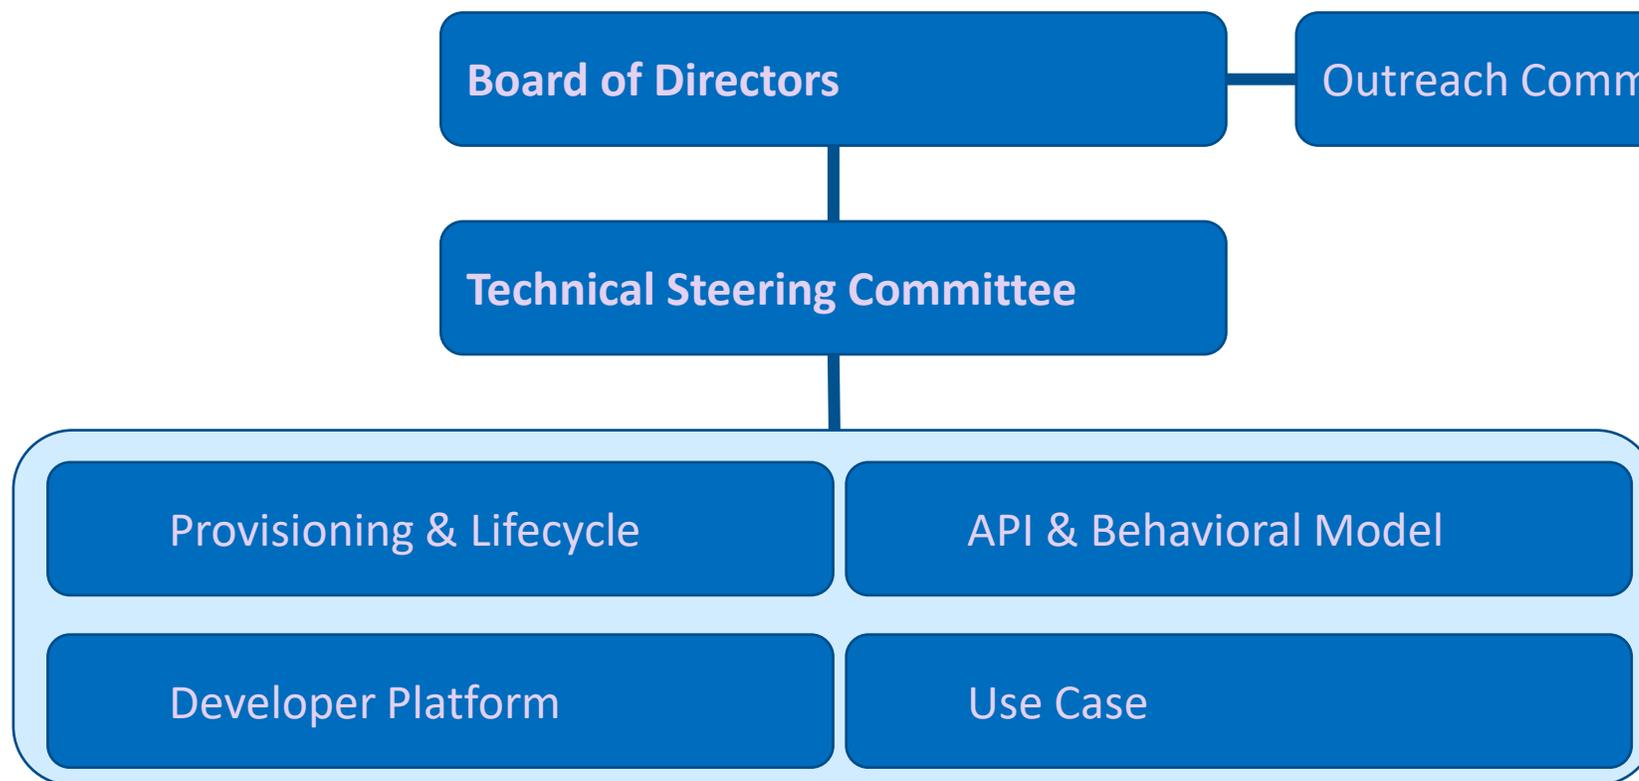
## OPI Compliant Devices Minimum Expectations

- Presence of their own general purpose processor
- The ability to boot a general purpose OS
- Domain-specific HW acceleration capabilities
- Software-defined device functions that allow the software components deployed to them to define the device's functions that are presented to the host
- Offloading of whole software subsystems, such as the Networking or Storage stack, including their control planes
- Strict security isolation from the host on the hardware-level
- Unique network identity
- Management
  - Capable of being managed as part of the hosting server (through BMC or hosting OS)
  - Capable of being directly managed (out-of-band) separately from the hosting server
  - Capable of managing the hosting server

# OPI Scope

Platform	API	Device Monitoring
<ul style="list-style-type: none"><li>○ Device Discovery</li><li>○ Zero Touch</li><li>○ Zero Trust</li><li>○ Inventory</li><li>○ Lifecycle &amp; Updates</li></ul>	<ul style="list-style-type: none"><li>○ Storage</li><li>○ Network</li><li>○ Security</li><li>○ AI/ML Interface</li></ul>	<ul style="list-style-type: none"><li>○ Open Telemetry (OTEL)</li><li>○ Metrics</li><li>○ Logs</li><li>○ Tracing</li></ul>
<h3>Developer Platform</h3> <ul style="list-style-type: none"><li>○ Real devices &amp; emulation</li><li>○ CI/CD pipeline</li></ul>		

# OPI Overall Structure

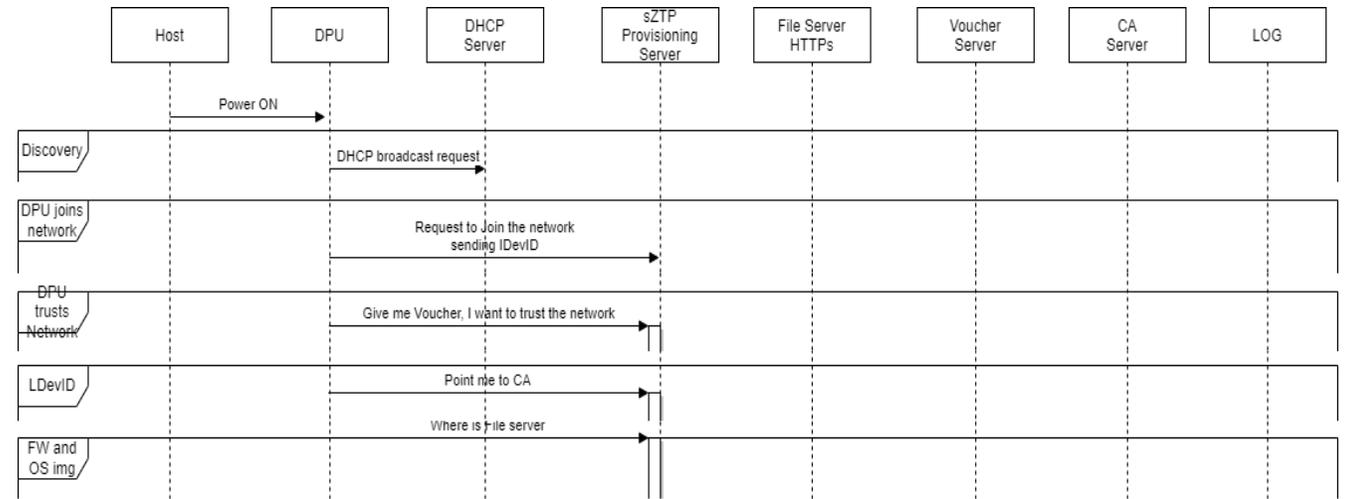


## OPI Technical Deliverables

- Open-Source Projects
- Specifications/Standards
- Reference Platforms
- Test Suites & Cases
- POC/Prototypes

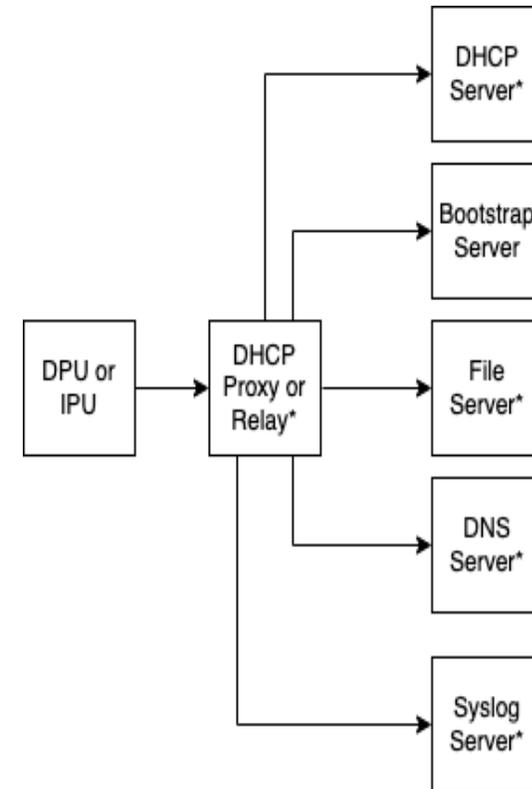
# Provisioning and Lifecycle Working Group

- ❑ Discovery & Provisioning
- ❑ Inventory
- ❑ Boot sequencing
- ❑ Lifecycle & Updates
- ❑ Monitoring & Telemetry



# Device Discovery and Provisioning

- ❑ Security first (mutual trust)
  - ❑ sZTP & FIDO
- ❑ Zero-Touch
  - ❑ Plug & Play
- ❑ Monitoring all the way
  - ❑ OTEL
- ❑ Multiple use cases
  - ❑ Challenges

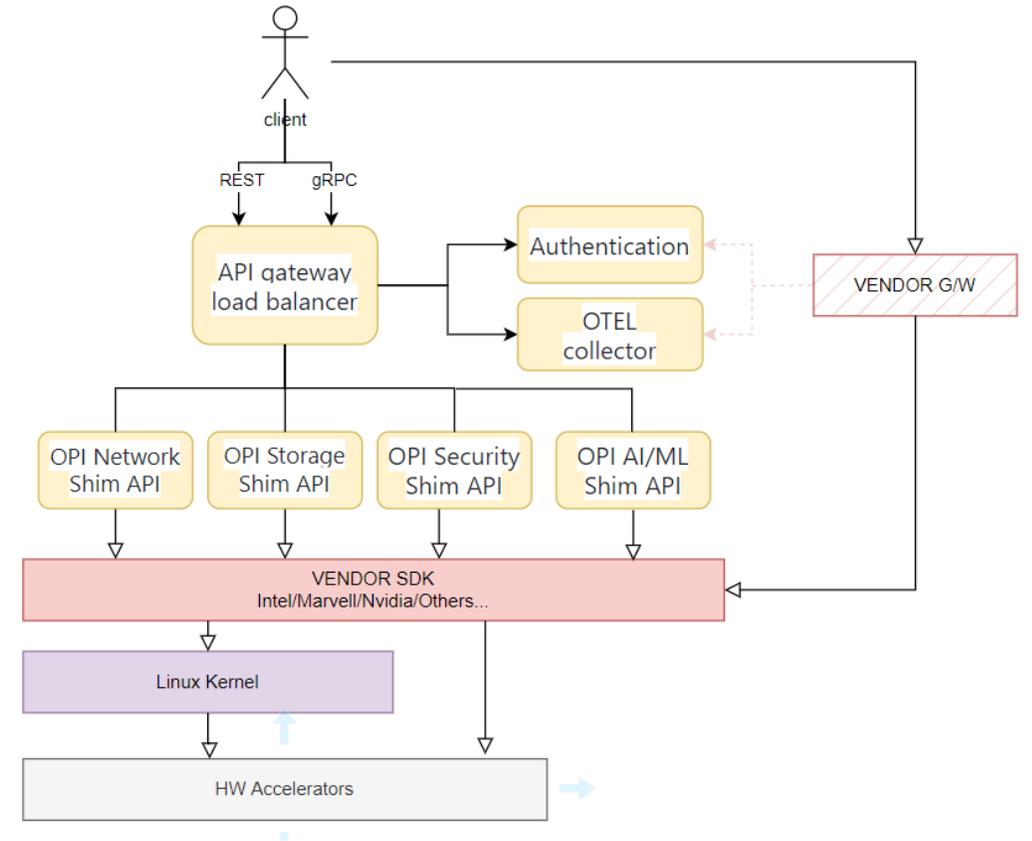


# Monitoring & Telemetry via OTEL

- ❑ OPI adopted [OTEL](#) for xPUs
- ❑ Single integration with OTEL instead of with multiple systems
- ❑ Supports Traces, Metrics, Logs
- ❑ OPI mandates only OTEL [Specification](#)
  - ❖ Not OTEL SDK, OTEL Collector
- ❑ Micro-Aggregator in xPUs, Marco-Aggregator across xPUs
- ❑ Common Metrics across xPU vendors

# API & Behavioral Model Working Group

- ❑ Object models
- ❑ Host & Management facing APIs
- ❑ Taxonomy for Services
- ❑ Re-use industry standard APIs
- ❑ Reference Orchestration Client



# DPU Open APIs

## • System

- Systems Management & Lifecycle
  - (Redfish, BMC, etc.)
- Monitoring, Metering, & Telemetry

## • Operating System (Linux)

- Standard Linux Libraries and packages
- Container and Application Hosting
- Leverage commonly used APIs
  - DPDK, SPDK, EBPF

## • Hardware (PCIe...)

- Virtual Function Mapping
- Offload Configuration

## • Low Level (likely Vendor specific APIs)

- Micro-Code in Data Flow Processing Cores
- P4 Packet Processing Pipelines

## • Vendor Unique API & SDK

- *These are NOT common/Open APIs*
- ASAP2, SNAP

## • Storage

- Networked Storage
  - NVMe/TCP
  - NVMe/RoCE(RDMA)
- Storage Services
  - RAID/Erasure Coding/etc
- Compression
- SDXI Offload

## • Gateway

- Connection Tracking
- Load Balancing
- NAT
- Tunnels

## • Networking

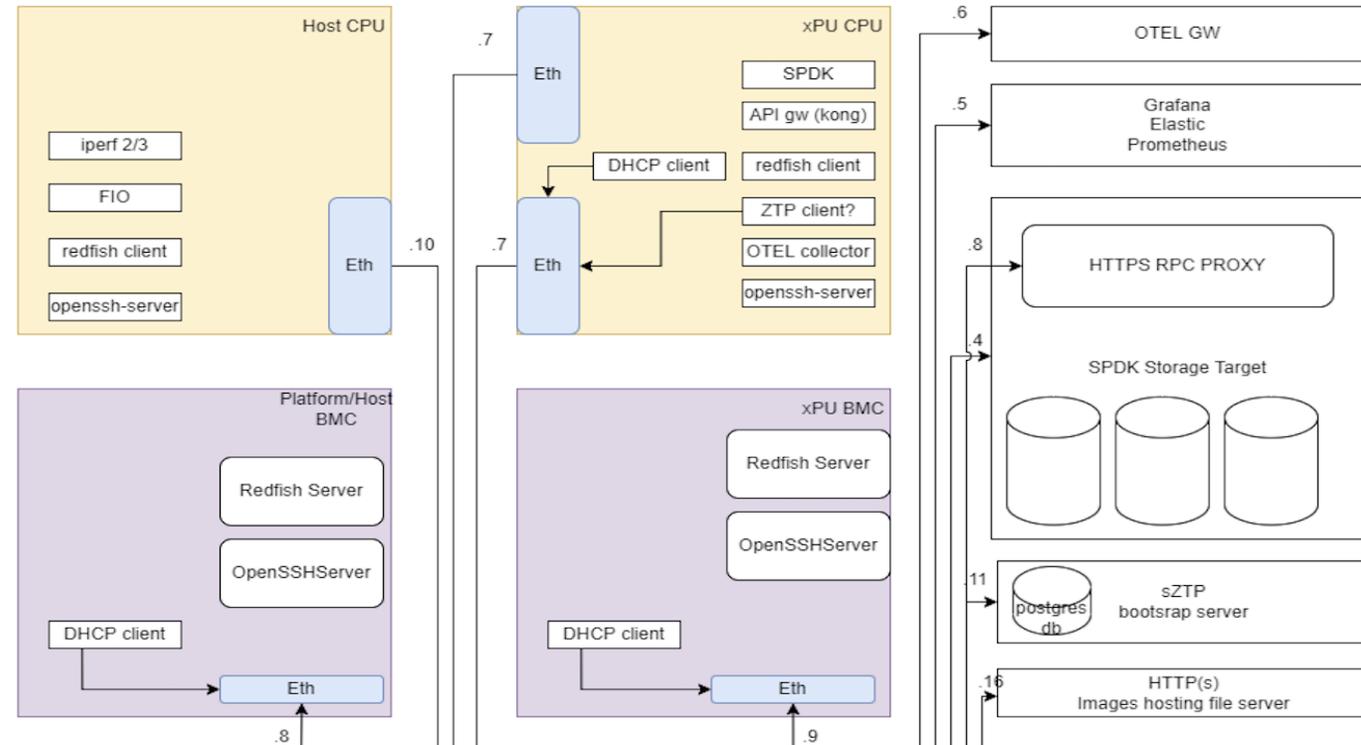
- SONiC
  - OpenConfig (includes BGP, etc)
  - SAI implementation by the DPU
- Policing and QoS and SLA
- Multi-tenant Overlay
- Host facing NIC Configurations
- OVS

## • Security

- Policy & Filters
- Crypto Offloads
- Secure Storage
  - keys, secrets, attestation, ...
- Key Management
- Network security offload
  - (TLS, IPSec)
- RegEx matching

# Developer Platform Working Group

- ❑ Multi-Vendor Lab
  - ❑ Considering UNH
- ❑ Virtual & Hardware POCs
- ❑ Simulation Environment
- ❑ CI/CD



# Use Cases Working Group

- **Initial Use Cases**
  - NVMe/PCIe to NVMe/TCP bridge
  - Basic Firewall with rule-based filtering
- **General High Interest Areas**
  - Storage
  - Security
  - Networking
  - AI/ML

# Key Takeaways and Call to Action

- ❖ Industry interest for developing common xPU APIs is strong
  - ❖ Customers, xPU Vendors, Software Vendors, Solution Providers
- ❖ Immediate Relevance to Storage and Storage Networking
- ❖ **Brand new effort so Join Now!**
  - ❖ We need input and contributions across the working groups



# Join the Project

Anyone can participate and contribute to the OPI Project

1. **To Participate**, check out the [OPI Mailing List](#), and the [OPI Slack channels](#).
  - a. Join the subgroup lists and channels in which you would like to participate.
  - b. Join the subgroup meetings via the invites found [here](#).
2. **Contribute** by following the steps [here](#) on GitHub.
3. **Become a Member** and support the OPI Project at the Linux Foundation [link](#).
  - a. Open Programmable Infrastructure would not exist without the support of the member organizations.



# Please take a moment to rate this session.

Your feedback is important to us.



**OPEN  
PROGRAMMABLE  
INFRASTRUCTURE  
PROJECT**