SNIA NSF | NETWORKING STORAGE

An Introduction to the OPI (The Open Programmable Infrastructure) Project

Live Webcast

January 25, 2023

10:00 am PT / 1:00 pm ET

Today's Presenters

Paul Pindell
Principal Architect, F5

Boris Glimcher
Distinguished Engineer, Dell Technologies

Mark Sanders
Distinguished Engineer,
Dell Technologies

Elad Blatt
GH Telco Networking, NVIDIA

Steven Royer
Principal Software Engineer, Red Hat

Dr. Joseph L. White
Fellow, Dell Technologies

SNIA - By the Numbers

Industry Leading Organizations

180

Active Contributing Members

2,500

IT End Users & Storage Pros Worldwide

50,000

OPEN PROGRAMMABLE INFRASTRUCTURE PROJECT

SNIA NSF | NETWORKING STORAGE

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Today's Agenda

- Introduction to OPI
- Lifecycle Provisioning
- API
- Use Cases
- Proof of Concept, Developer Platform

Introduction to OPI

Paul Pindell, F5

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/IPU-like technologies****.*

https://opiproject.org https://github.com/opiproject

https://lists.opiproject.org/g/opi

Premier

General

Project Goals

- Create community-driven standards-based open ecosystem for DPU/IPU-like technologies
- Create vendor agnostic framework and architecture for DPU/IPU-based software stacks
- Reuse existing or define a set of new common APIs for DPU/IPU-like technologies when required
- Provide implementation examples to validate the architectures/APIs

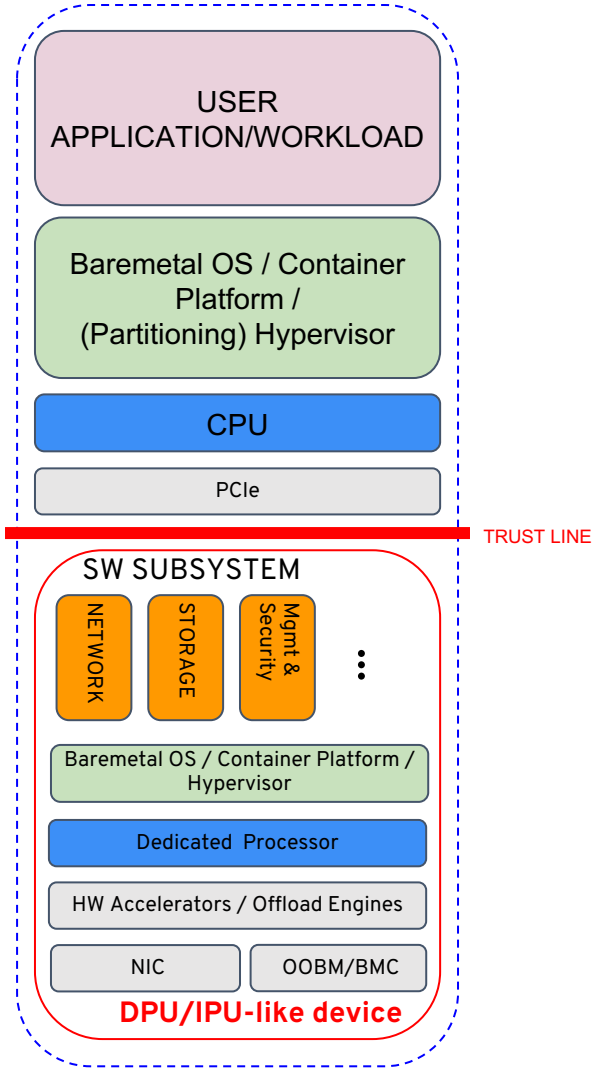

Physical server

Lifecycle Provisioning

Boris Glimcher, Dell

Lifecycle Provisioning Focus Areas

PROVISIONING INVENTORY BOOT SEQUENCING LIFECYCLE MANAGEMENT MONITORING & TELEMETRY

Main repo : https://github.com/opiproject/opi-prov-life
Slack: https://opi-project.slack.com/archives/C0342L6T7EC

Provisioning

- Secure Zero-touch provisioning (sztp RFC8572 standard)
  - Adding new "**Bootstrap**" server
- Security considerations
  - DPU Validation by the network
  - Network Validation by the DPU
  - Artifacts Validation (signed)
- Auxiliary RFCs
  - IEEE 802.1AR - Secure Device Identity
  - RFC8366 - A Voucher Artifact for Bootstrapping Protocols
- Implementing SZTP client/agent:
  - https://github.com/opiproject/sztp

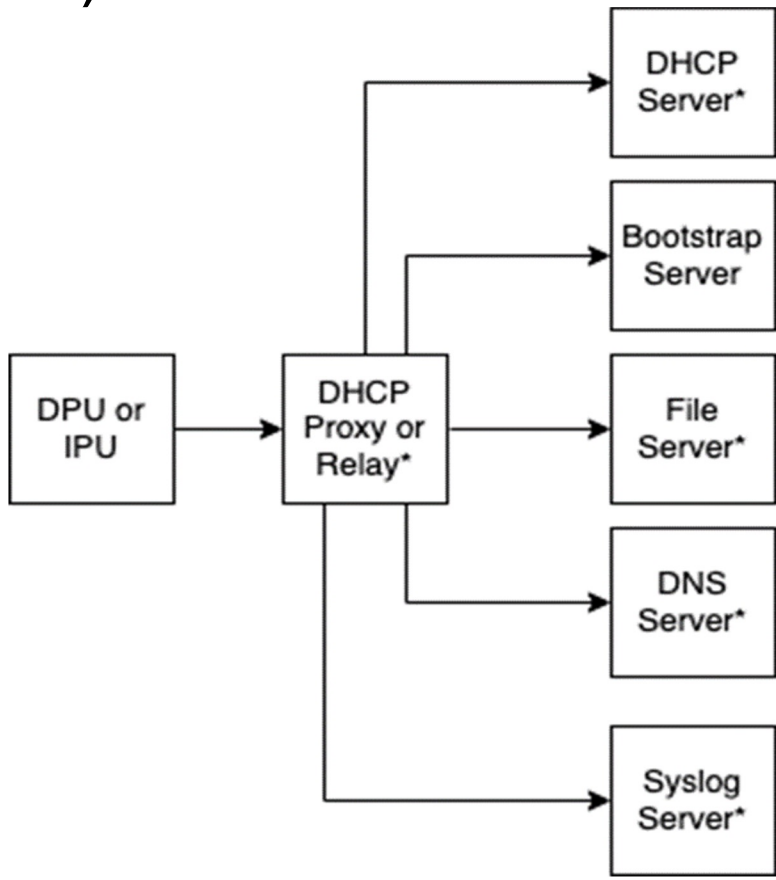

Inventory

- Local inventory
  - DMI/SMBIOS standard
    - https://www.dmtf.org/standards/smbios
  - Compliance
    - https://github.com/opiproject/smbios-validation-tool
- Remote network inventory
  - Redfish standard
    - https://www.dmtf.org/standards/redfish
  - gRPC protocol buffers
    - https://github.com/opiproject/opi-smbios-bridge
- Remote host (OS and BMC)
  - NCSI, PLDM, VPD, …
  - WIP


```
$ dmidecode -t system

Handle 0x0001, DMI type 1, 27 bytes
System Information
        Manufacturer: https://www.mellanox.com
        Product Name: BlueField SoC
        Version: 1.0.0
        Serial Number: Unspecified System Serial Number
        UUID: 2e3bc1d1-e205-4830-a817-968ed1978bac
        Wake-up Type: Power Switch
        SKU Number: Unspecified System SKU
        Family: BlueField
```


Boot Sequencing

- Boot Sequencing
  - Host waits for DPU to complete booting
    - Host halted
    - DPU sends “boot complete” notification
    - Timeout in case no notification
  - DPU reboots
    - PCIe DPC hostplug and surprise removal issues
    - Host needs notification from DPU
  - Host reboots
    - DPU needs notification from Host (i.e., free resources)
  - Crashes
    - both DPU and Host can panic

Lifecycle Management

- Required actions:
  - Reboot a DPU
  - Update:
    - DPU FW
    - DPU OS and Bootloaders
    - DPU SW/Apps
  - Recover DPU to a known state
    - i.e. factory reset
    - Via network & via Host
  - Debugging of the DPU

Monitoring & Telemetry

- Open Telemetry (OTEL) standard
  - https://opentelemetry.io/
- Supported data sources
  - Traces
  - Metrics
  - Logs
- Example OTEL collector with SPDK and System monitoring
  - https://github.com/opiproject/opi-prov-life/tree/main/examples
- Recommended OTEL collector deploy option:
  - On every DPU close to data generation
  - Increased redundancy, Enrichment, Filtering, Batching, Sanitization, …

API

Mark Sanders, Dell

API Focus

- Create a Multi-Vendor Open API definition and adoption for:
  - Storage Services
  - Network Services
  - Security Services
  - AI/ML
  - Telemetry
  - System and Lifecycle Management
- Define control/management facing APIs
  - Inventory Discovery Service and Capabilities Advertisement API
  - Configuration API
  - Consumption API

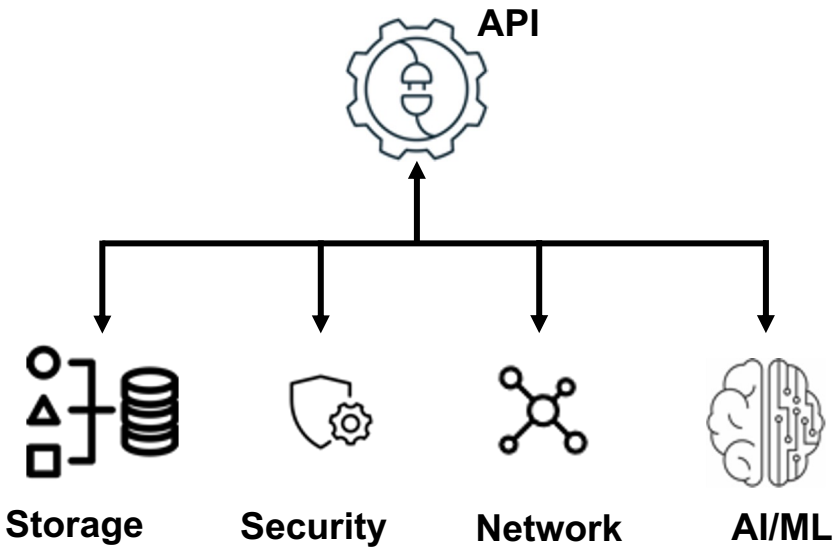

Main repo : https://github.com/opiproject/opi-api
Slack: https://opi-project.slack.com/archives/C0344KMEAKB

API Mechanism

- Utilize gRPC for the config/control interface
- Use API Gateway
  - Direct delivery of gRPC messages to appropriate shim layer
  - gRPC to REST translation
  - Support gNMI and gNOI
- Expose VF/PF for the data consumption interface

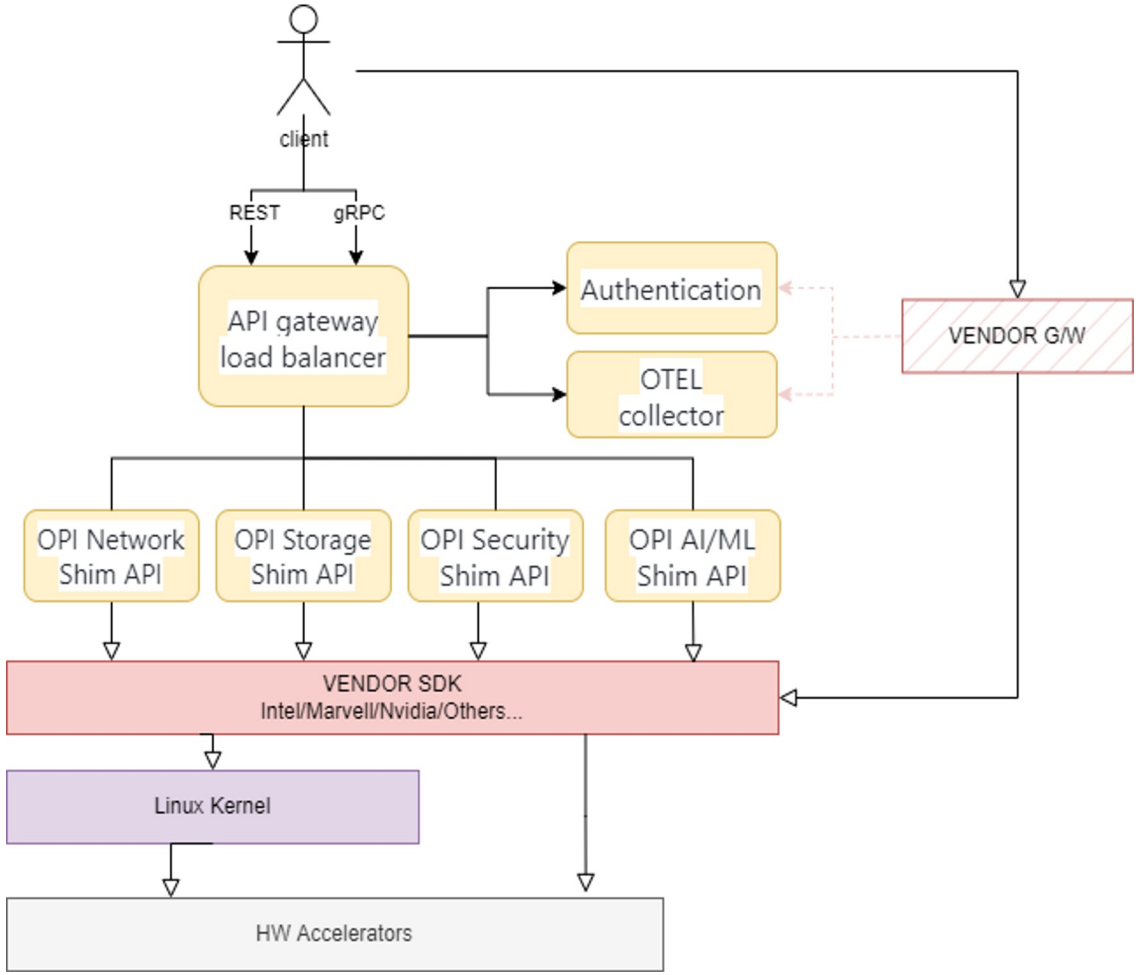

Storage

- Identified 3 sets of APIs for Storage
  - Front End
    - Host Facing APIs (emulated NVMe devices, Virtio-blk, Virtio-scsi, …)
  - Middle End
    - Storage Services APIs (RAID protection, Compression Encryption, …)
  - Back End
    - Network Facing APIs
    - Connects to devices locally (pcie) or remotely (iSCSI, RDMA, TCP, IPv4/v6, …)
- Reference Implementation: OPI SPDK Bridge
  - https://github.com/opiproject/opi-spdk-bridge
- Vendor Reference Implementations
  - opi-marvell-bridge
  - opi-nvidia-bridge

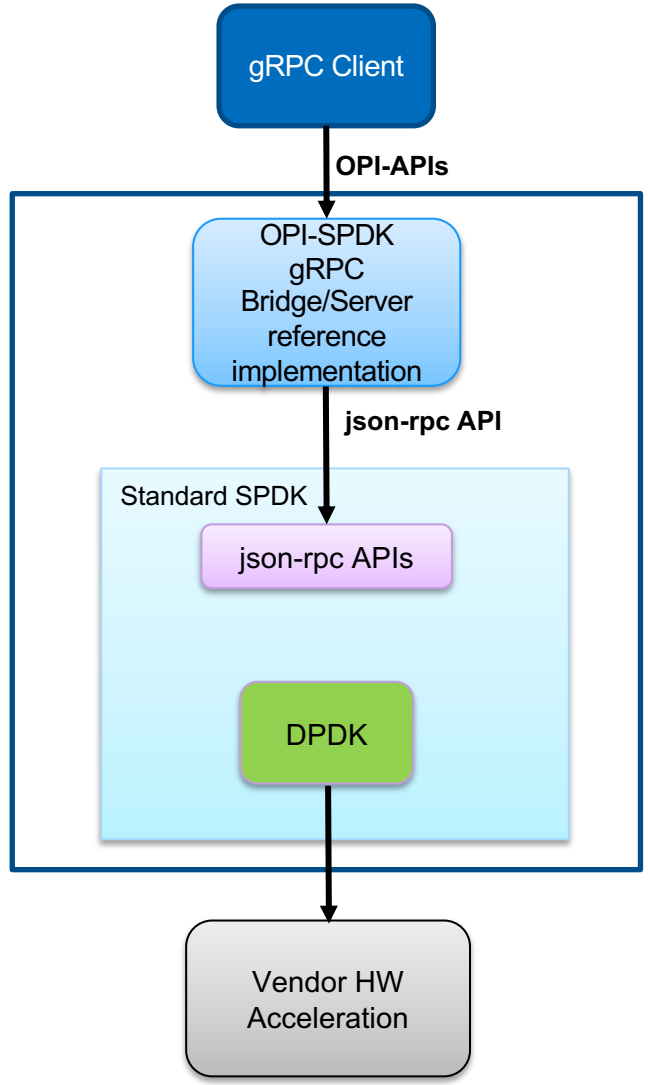

Networking

- Create a common API framework and extensions for specific areas
  - Cloud
  - Telco
  - K8s
- Support network service capabilities
  - OVS, SONiC, VPP, …
- Leverage existing API models
  - OpenConfig, OVS DB, CNI, …
  - https://github.com/opiproject/opi-api/tree/main/network

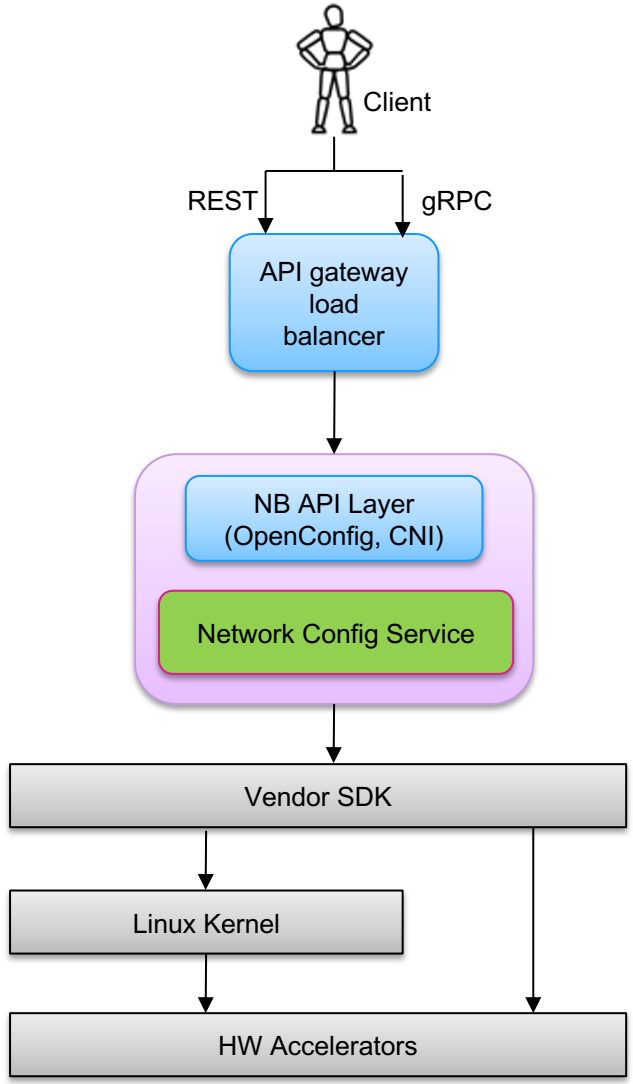

Security

- Address the security capabilities
  - Security Offloads (TLS, IPSec)
  - Crypto Offloads
  - Secure Storage and Key Functions
  - Policy and Filters

- Current focus is on IPSec API on top of strongSwan
  - https://github.com/opiproject/opi-api/tree/main/security
  - https://github.com/opiproject/opi-strongswan-bridge

Use Cases

Elad Blatt, NVIDIA

Use Case Focus:

- Create an open channel with the end users and potential deployment partners.
- Share the work being done by the subgroups and get feedback from the end users and deployment partners.
- Encourage more end users and deployment partners to contribute, take part, and join.

- Alignment to project goal(s)

- By engaging with and seeking input from end users and deployment partners, the Use Case subgroup will help by validating the implementation examples, architectures, and APIs being worked on in the other subgroups.
- This will also help build the open community to include additional contributors and members.

Use Case discussions:

We have identified 3 main groups of features that have generated the most interest :

1) Routing/Switching – OVS , BGPs , SR , Encap/decap etc
2) Security – WAF, NGFW, IPSEC , DDOS etc
3) Storage – NVMe oF , oTCP , RDMA , SPDK support etc

Initial Use cases defined :

Security – Basic FW , rule-based filtering .

Storage - Nvme/Pcie to Nvme/Tcp bridge.

Routing – OVS and looking into K8 networking .

Proof of Concept, Developer Platform

Steven Royer, Red Hat

PoC Focus

- Simulation environment and CI
  - https://github.com/opiproject/opi-poc/blob/main/integration/README.md
- Lab definition
  - Building a lab to be used for real hardware environments
- https://github.com/opiproject/opi-poc

Simulation

Real services run in containers via docker-compose

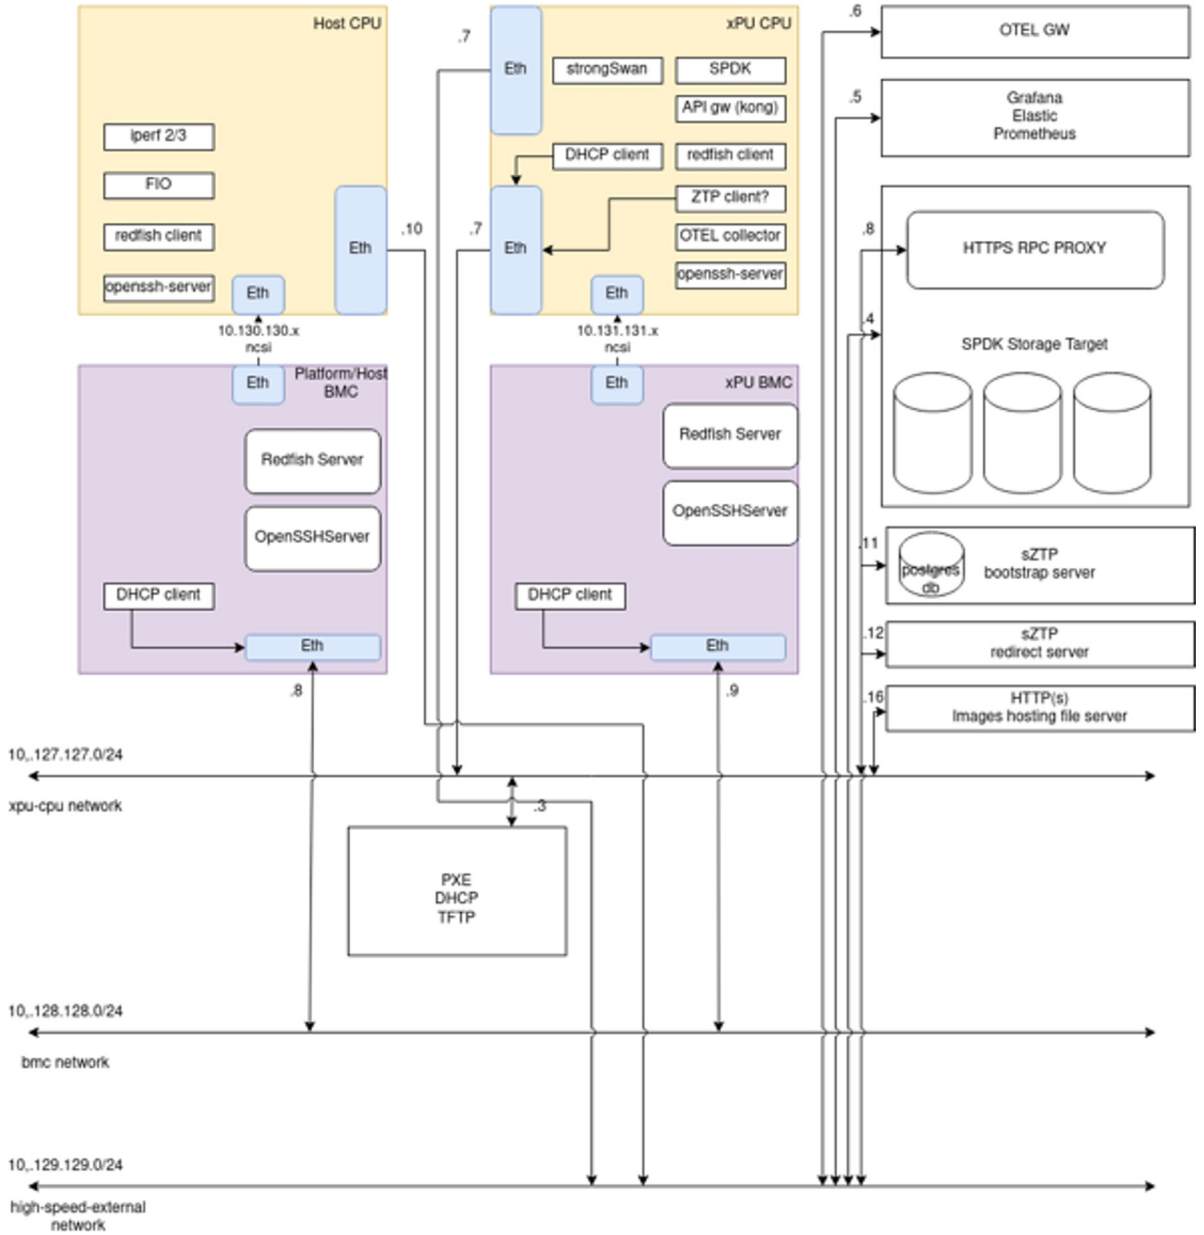

Join the Project

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 project 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.

Q&A

After this Webcast

- Please rate this webcast and provide us with your feedback
- This webcast and a copy of the slides are available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webcast, including answers to questions we couldn't get to today, will be posted on our blog at https://sniansfblog.org/
- Follow us on Twitter @SNIANSF

Thank You

