

OPI Event – Open D/IPU API

Need for Common Interface Framework

DELLTechnologies

Multiple Vendors

Solutions and Ecosystems

- Vendors

- NVIDIA
- Pensando
- Intel
- Marvell
- Xsight
- Fungible
- AMD/Xilinx
- ...

- Cards and Silicon

- Bluefield, Bluefield2
- DSC-25/100, DSC-200
- Mt Evans
- CN9xxx
- Octeon 10
- FC50, FC100, FC200
- ...

- Ecosystems

- DOCA
- IPDK
- Various Marketplaces

DPU/IPU Roulette

Multiple Behaviors, Multiple Interfaces, Multiple Frameworks

- High Level Behavioral Models
 - SDXI (SNIA)
 - DASH
 - Redfish (DMTS)
 - OpenBMC
 - OpenConfig
 - Internal Vendor System Level
 - DOCA
 - IPDK
 - Low Level
 - Vendor SDKs
 - Pipeline interfaces
 - DPDK
 - SPDK
 - RegEx
 - Operating Environment (Linux)
 - Container
 - OS Acceleration (eBPF)
 - Service Specific
 - Networking
 - Storage
 - Security
 - Gateway
 - Accelerator (AI/ML)
 - Telemetry
 - Lifecycle Management
- Need a community agnostic API set to provide a common behavioral interface set for the DPU/IPU multi-vendor ecosystem

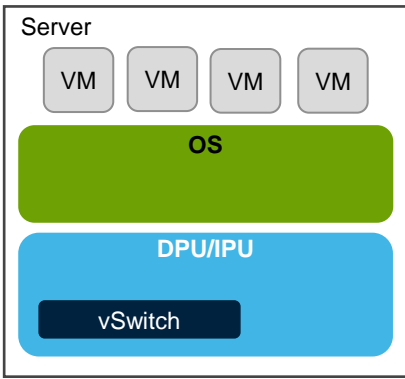
Key Use Cases for D/IPU

0 Better / Distributed Security (cross-cutting use-case)

1

Network Offload and Disaggregation

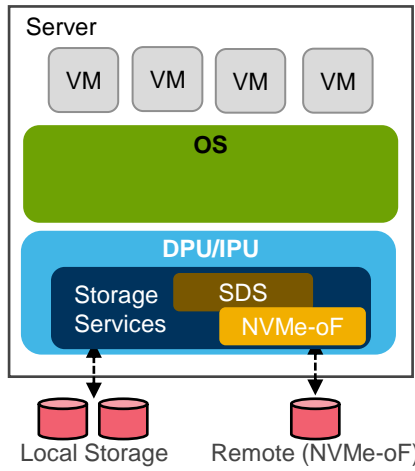
Virtualize the ToR and offload network functions



2

Storage Offload and Disaggregation

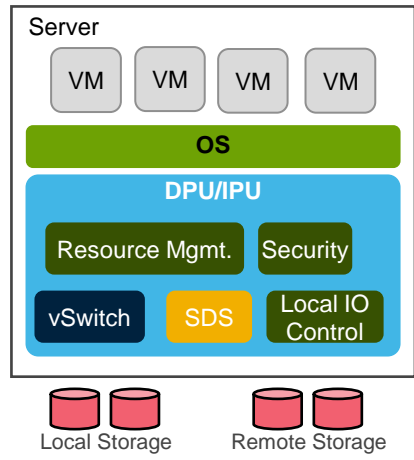
Offload storage functions



3

Bare Metal

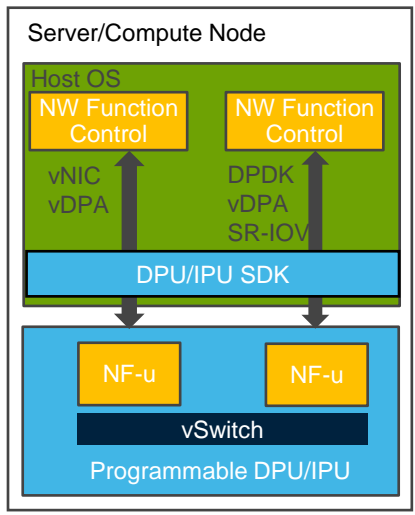
Utilize as a Bare Metal Controller



4

Edge/5G

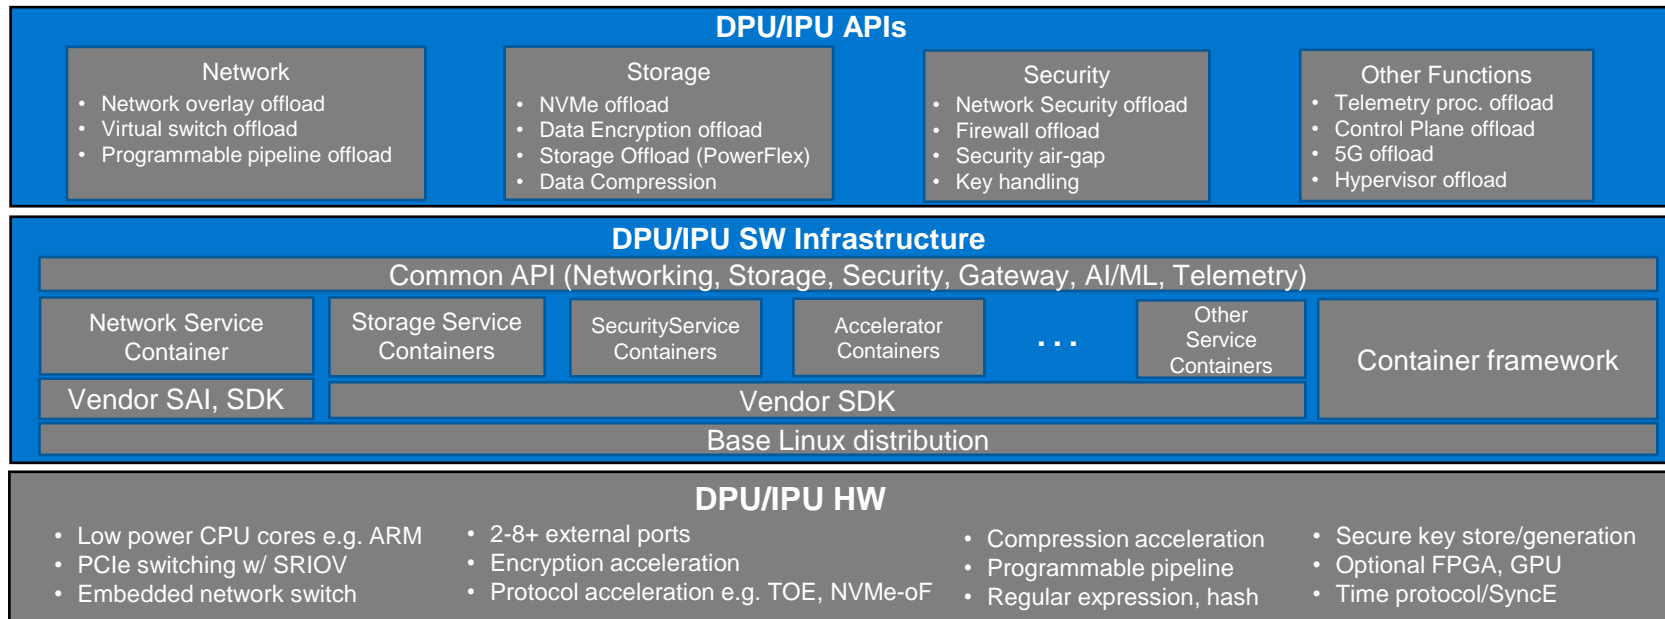
Offload I/O intensive functions



Need for an Open API for D/IPU

- Define standard mechanisms for Service Deployment
- Support of a Multi-Vendor Open D/IPU API definition and adoption for
 - Storage Services
 - Network Services
 - Security Services
 - AI/ML
 - Telemetry
 - System and Lifecycle Management
- Reuse Existing or define new common APIs for Configuration, Management and Consumption

Top Level Framework View



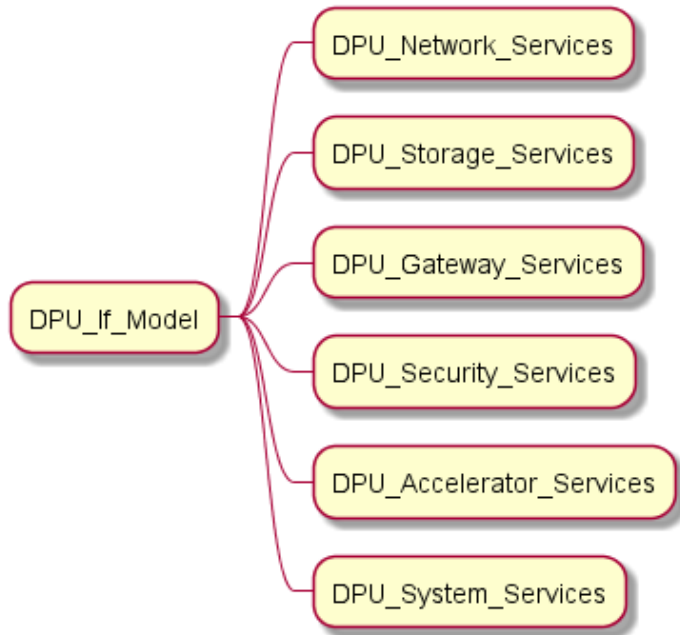
API Scope

<ul style="list-style-type: none">• System<ul style="list-style-type: none">• Systems Management & Lifecycle<ul style="list-style-type: none">• (Redfish)(OpenBMC?)(etc)• Monitoring, Metering, & Telemetry
<ul style="list-style-type: none">• Operating System (Linux)<ul style="list-style-type: none">• Standard Linux Libraries and packages• Container and Application Hosting
<ul style="list-style-type: none">• Hardware (PCIe...)<ul style="list-style-type: none">• Virtual Function Mapping• Offload Configuration
<ul style="list-style-type: none">• Low Level APIs<ul style="list-style-type: none">• Micro-Code in Data Flow Processing Cores• P4 Packet Processing Pipelines• Leverage commonly used APIs<ul style="list-style-type: none">• DPDK, SPDK, EBPF
<ul style="list-style-type: none">• Vendor Unique API & SDK<ul style="list-style-type: none">• <i>These are NOT common/Open APIs</i>• DOCA, ASAP2, SNAP

<ul style="list-style-type: none">• Storage<ul style="list-style-type: none">• Networked Storage<ul style="list-style-type: none">• NVMe/TCP• NVMe/RoCE(RDMA)• Storage Services<ul style="list-style-type: none">• RAID/Erasure Coding/etc• Compression• SDXI Offload	<ul style="list-style-type: none">• Networking<ul style="list-style-type: none">• SONiC<ul style="list-style-type: none">• OpenConfig (includes BGP, etc)• SAI implementation by the DPU• Policing and QoS and SLA• Multi-tenant Overlay• Host facing NIC Configurations• OVS
<ul style="list-style-type: none">• Gateway<ul style="list-style-type: none">• Connection Tracking• Load Balancing• NAT• Tunnels	<ul style="list-style-type: none">• Security<ul style="list-style-type: none">• Policy & Filters• Crypto Offloads• Secure Storage<ul style="list-style-type: none">• keys, secrets, attestation, ...• Key Management• Network security offload<ul style="list-style-type: none">• (TLS, IPSec)• RegEx matching

- Other Groupings such as Accelerators, AI/ML, etc.

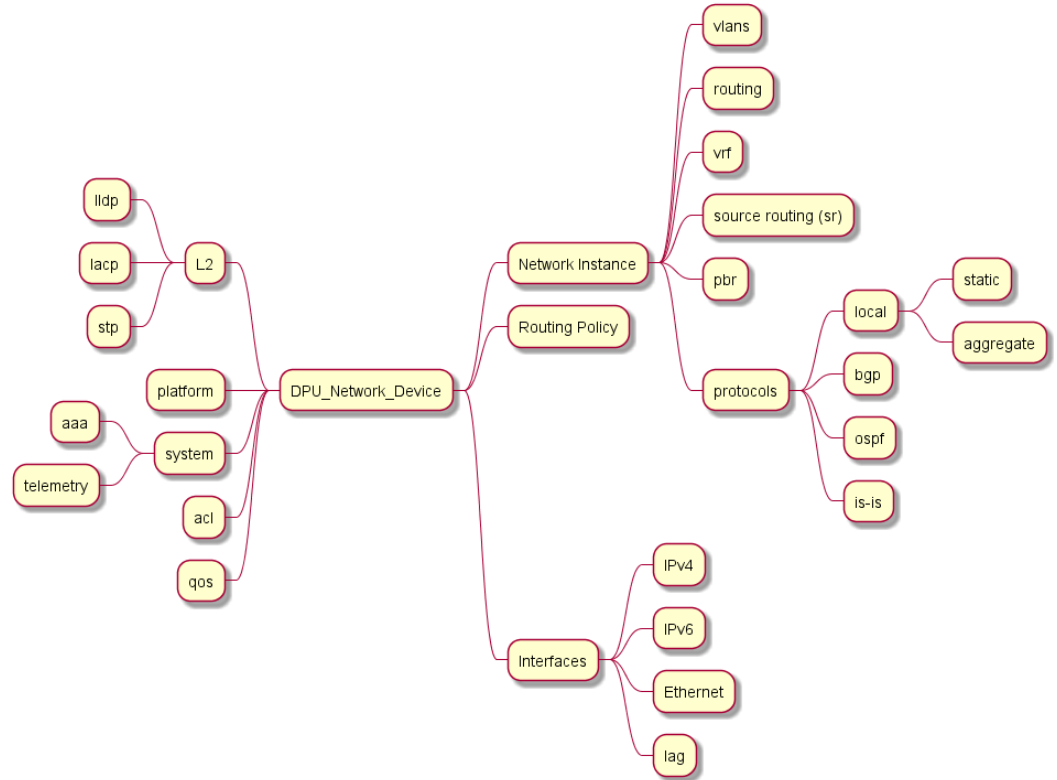
Open Programmable Infrastructure API Model



- Conceptual initial view of the API breakdown of services available in the D/IPU platforms
- Ability to support service chaining to provide the desired operation

Network Services API

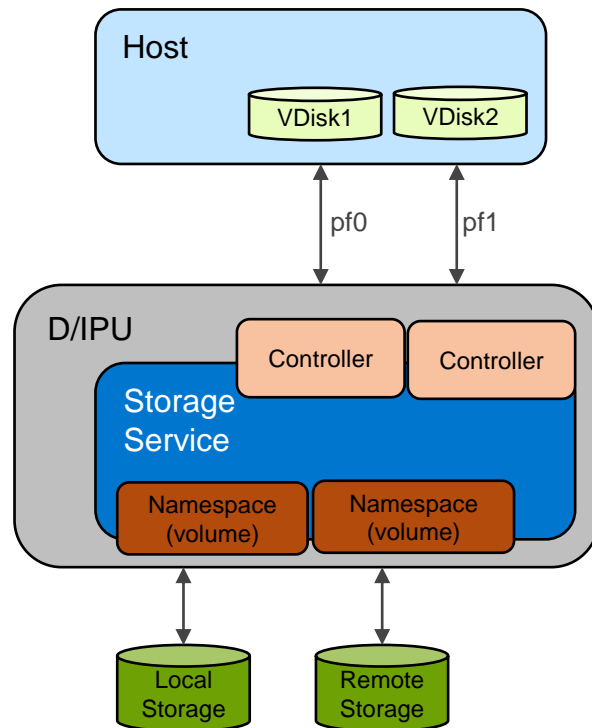
- Utilize the OpenConfig model for network service configuration
- Aligns with many Networking OS environments for configuration and management



Storage Service Example

High Level View (One of Many)

- Considerations for Storage API
 1. Setup Network Interface
 - IP Address (IPv4/IPv6), QoS, VxLAN, etc
 2. Create NVMe Subsystem (optional)
 3. Create the Controller
 - For each PF/VF
 4. Create the Namespace for the local or remote storage
 - PCIe, RDMA, TCP
 5. Attach the Namespace to the Controller



DELLTechnologies