

STORAGE DEVELOPER CONFERENCE



Fremont, CA
September 12-15, 2022

BY Developers FOR Developers

A **SNIA** Event

Disaggregated NVMe/TCP Storage Using an Infrastructure Processing Unit (IPU)

Yadong Li, Principal Engineer, Intel

Lev Solomonov, Principal Architect, Lightbits Labs

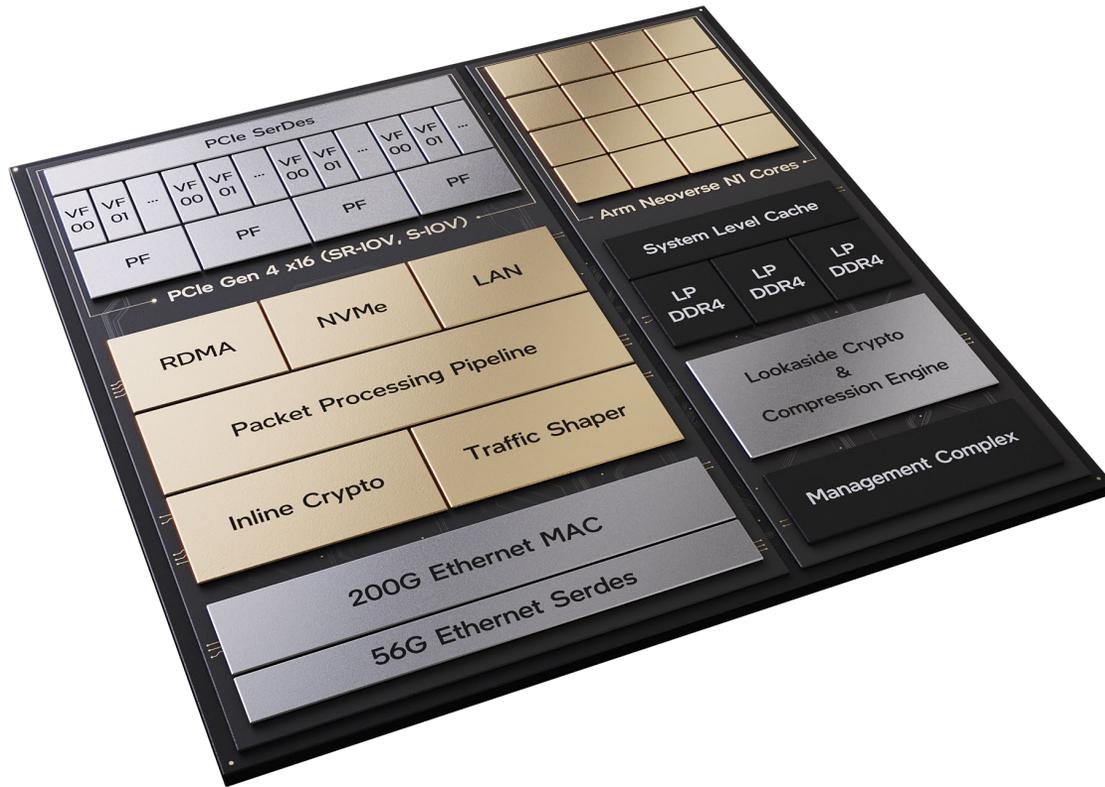
Acknowledgement: Dave Minturn (Intel), Muli Ben-Yehuda (Lightbits)

Agenda

- Overview of IPU architecture
- SPDK-based NVMe/TCP Initiator Design
- Overview of the Lightbits Cloud Data Platform
- Integration of IPU-based NVMe/TCP Initiator and Backend Storage Service
 - IPU Storage Management Agent (SMA)
 - Integration with K8s CSI node driver for orchestration
- Summary

Mount Evans

Intel's 200G IPU



Hyperscale Ready

Co-designed with **Google**

Integrated learnings from multiple gen. of FPGA sNIC/IPU

High performance under real world load

Security and isolation from the ground up

Technology Innovation

Best-in-Class Programmable Packet Processing Engine

NVMe storage interface scaled up from Intel Optane Tech

Next Generation Reliable Transport

Advanced crypto and compression accel.

Software

SW/HW/Accel co-design

P4 Studio based on Barefoot technology

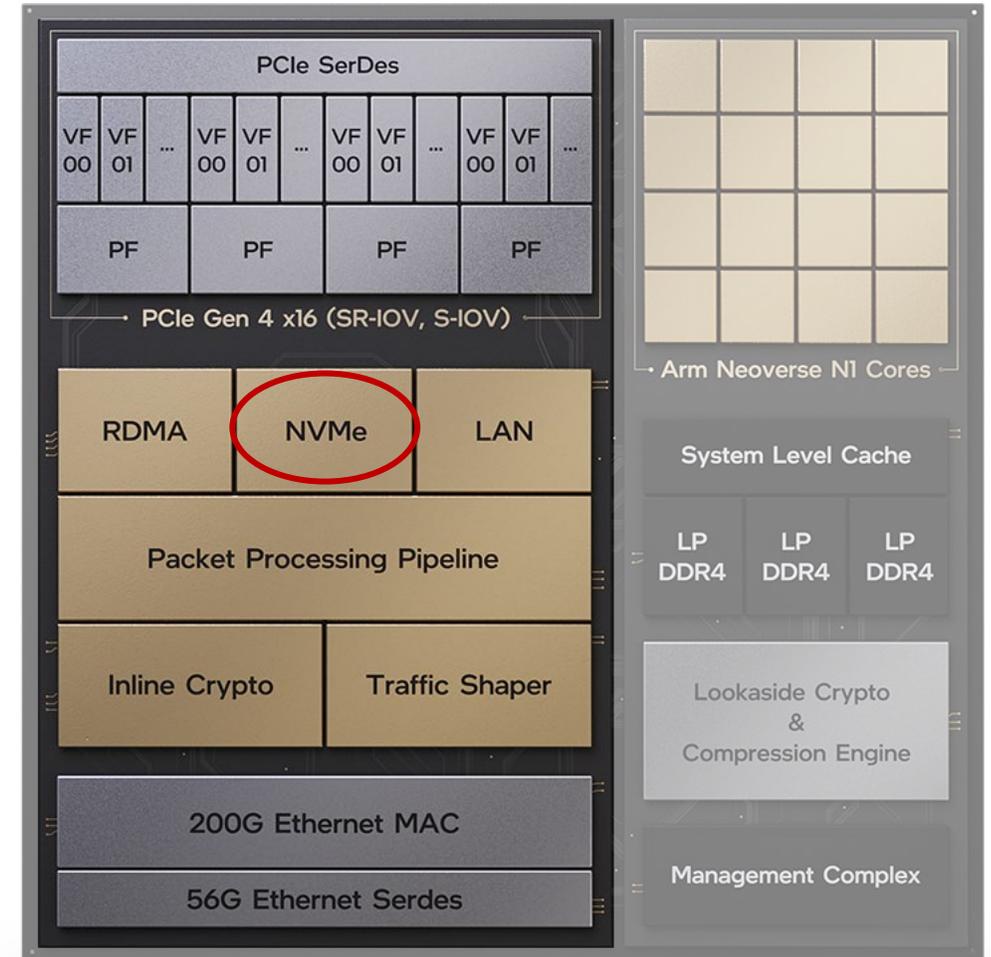
Linux OS leveraging DPDK, SPDK & IPDK eco-systems

VMWare's Project Monterey for telco & enterprise

Mount Evans - NVMe Overview

NVMe Initiator HW

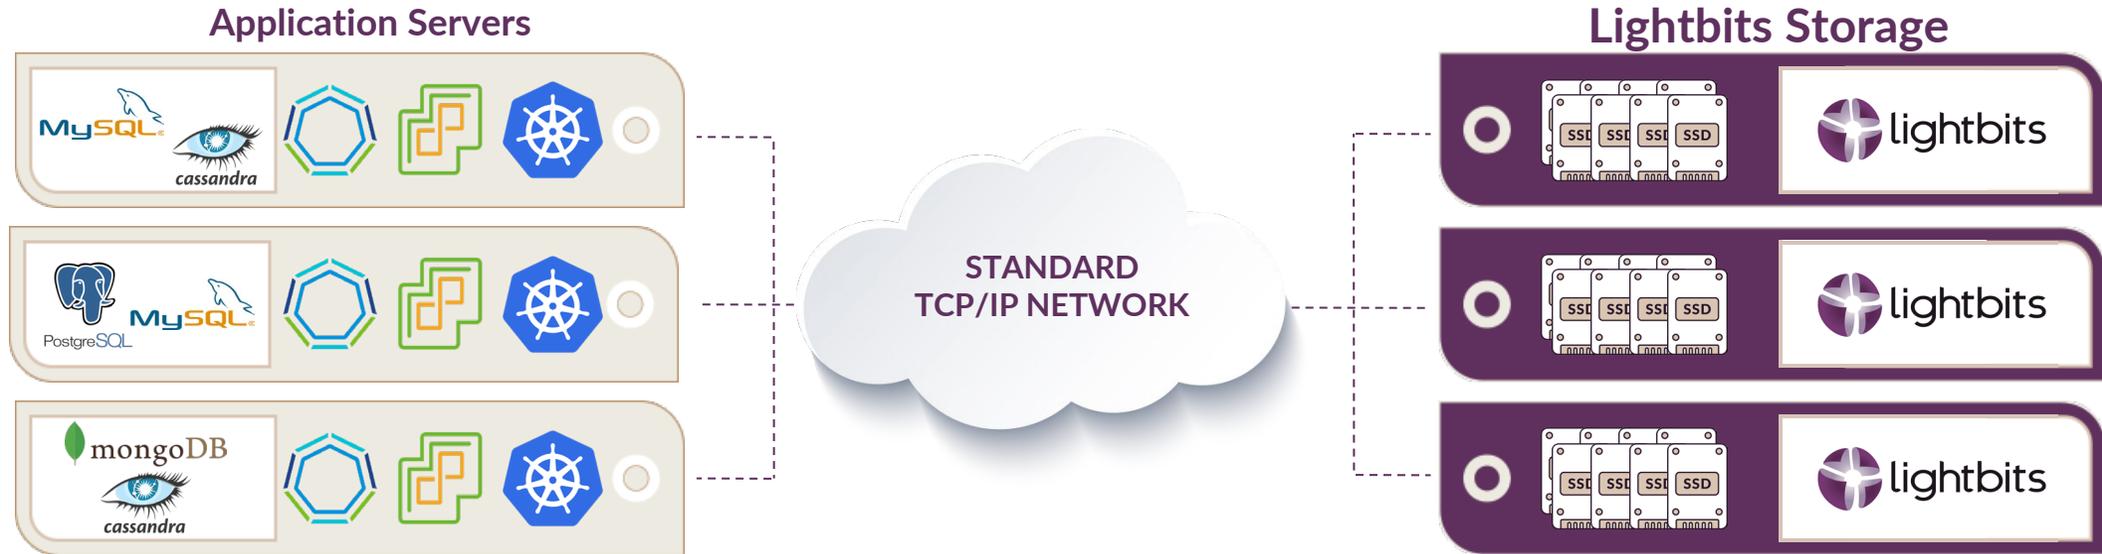
- Multi-Host support (4 PFs for Hosts)
- SR-IOV support
- Inline data-at-rest crypto (AES-XTS 256b)
- E2E data integrity protection
- IOPS and BW limiting per VF (or Queue Group)
- Live migration support
- LCE for lookaside crypto, compression and CRC offloads, support chained ops



Overview of the Lightbits Cloud Data Platform

Lightbits Software Defined Storage

Enterprise Features, at NVMe Speeds, on Any Cloud

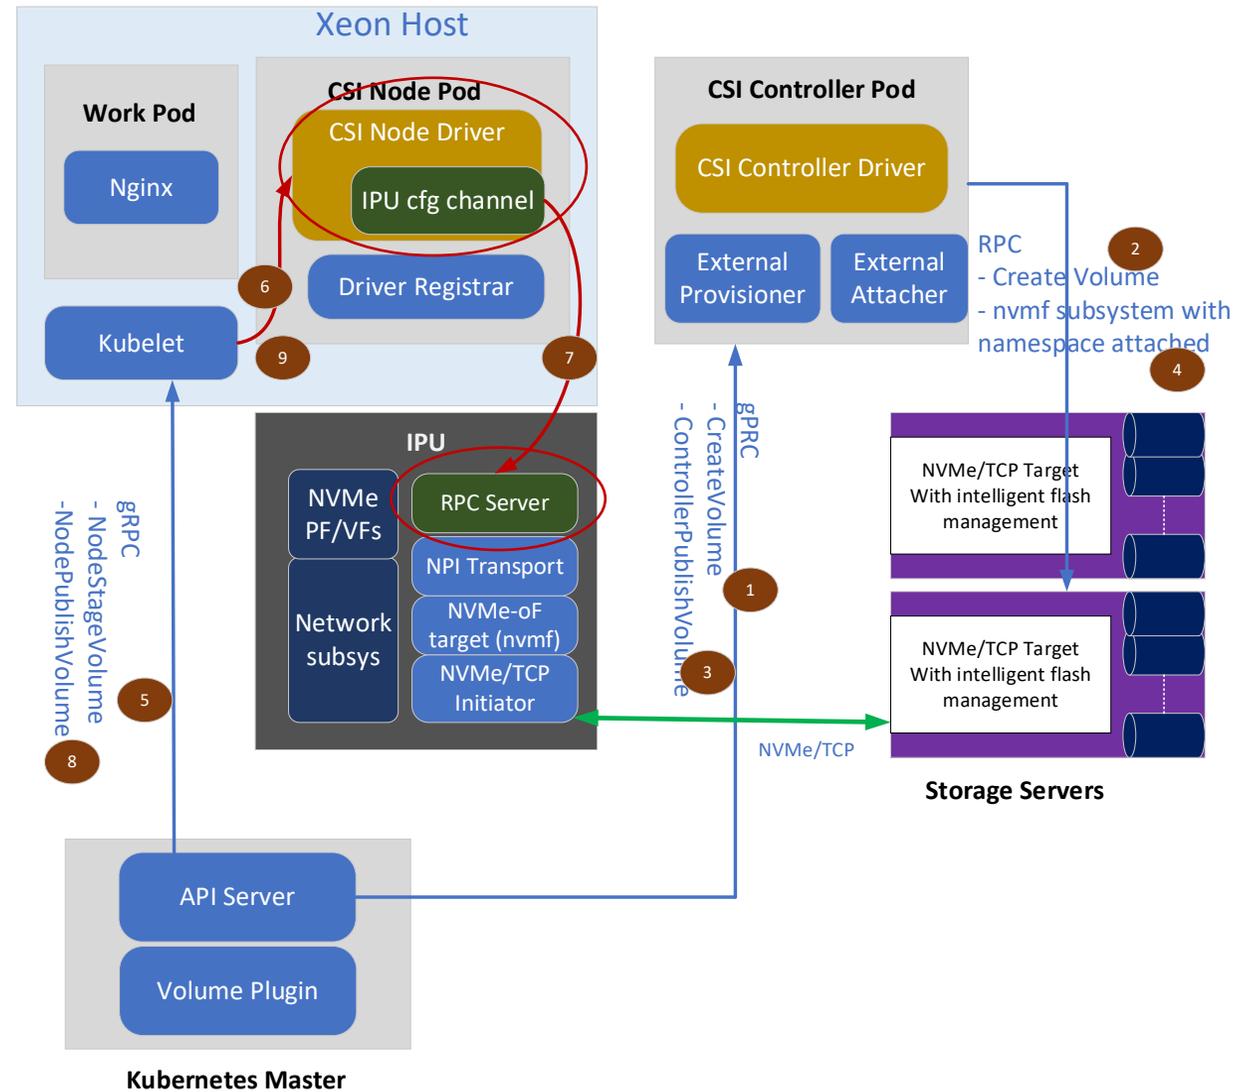


- **Flexible:** Scale storage and compute independently
- **Lower TCO:** Increased flash endurance, data reduction
- **High performance:** High IOPS, consistent low latency
- **Multi-tenant:** Single storage cluster can service multiple heterogeneous app environments
- **Easy:** Existing TCP/IP network, run on Intel/AMD/ARM servers
- **High availability:** SSD-level eRAID, volume replication
- **IPU-ready:** Applications servers with Intel IPUs to free up host CPUs

Integration of IPU-based NVMe/TCP Initiator and Backend Storage Service

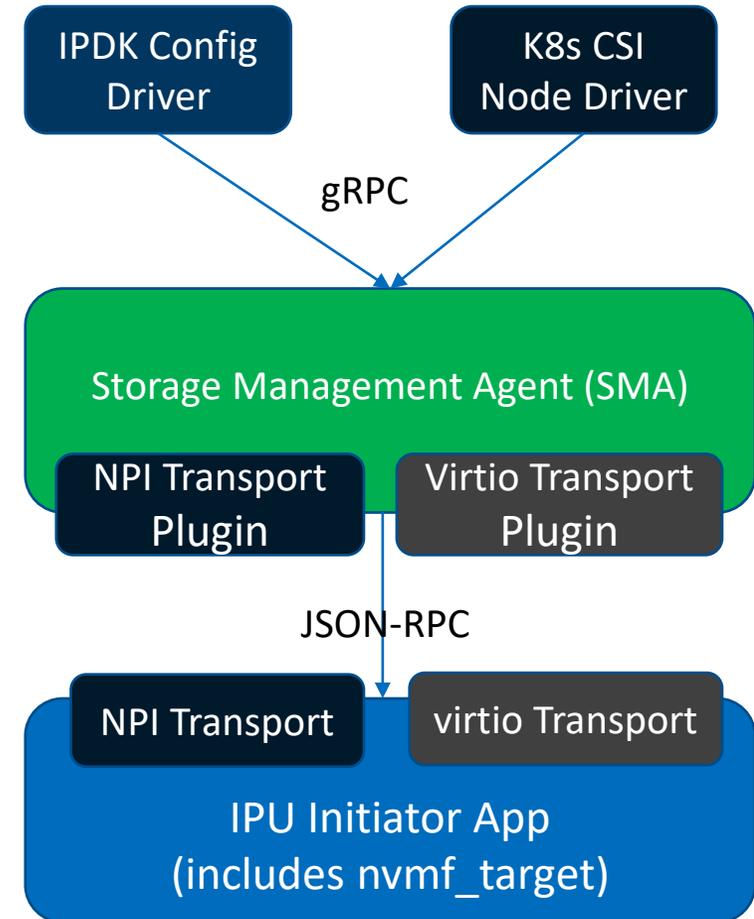
K8s CSI Components

- Integration of IPU NVMe/TCP Initiator management with K8s CSI Node Driver enables a full SDS solution



Storage Management Agent

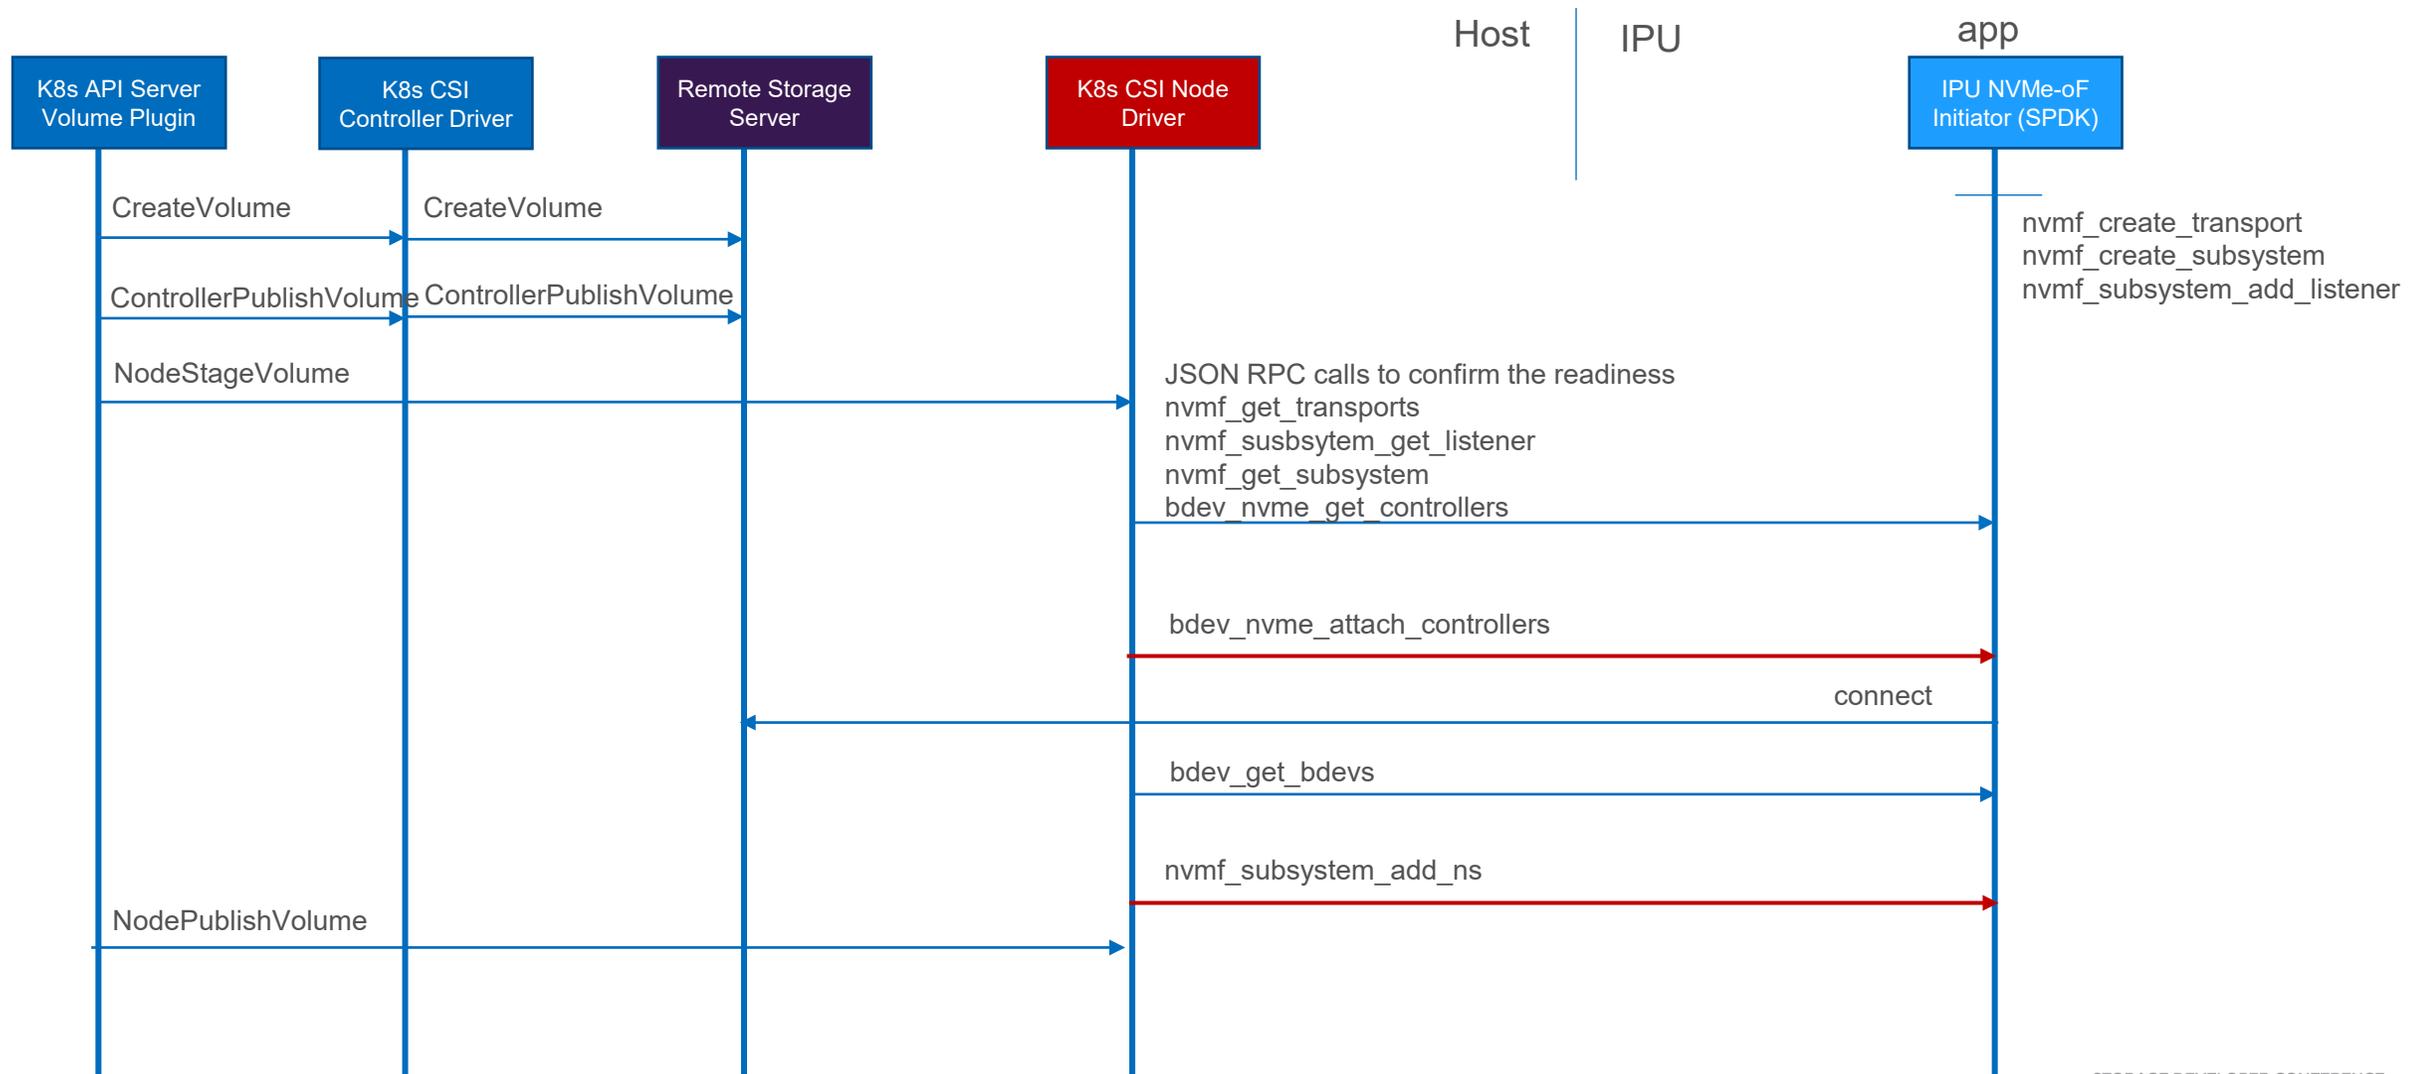
- Storage Management Agent (SMA) is a service providing a gRPC interface for orchestrating SPDK applications.
 - Provides a secure and authenticated configuration channel for IPU storage applications.
 - Abstracts out low-level details of SPDK JSON RPCs.
 - Simplifies integrations with K8s and OpenStack.
- For details, please refer to <https://spdk.io/doc/sma.html>



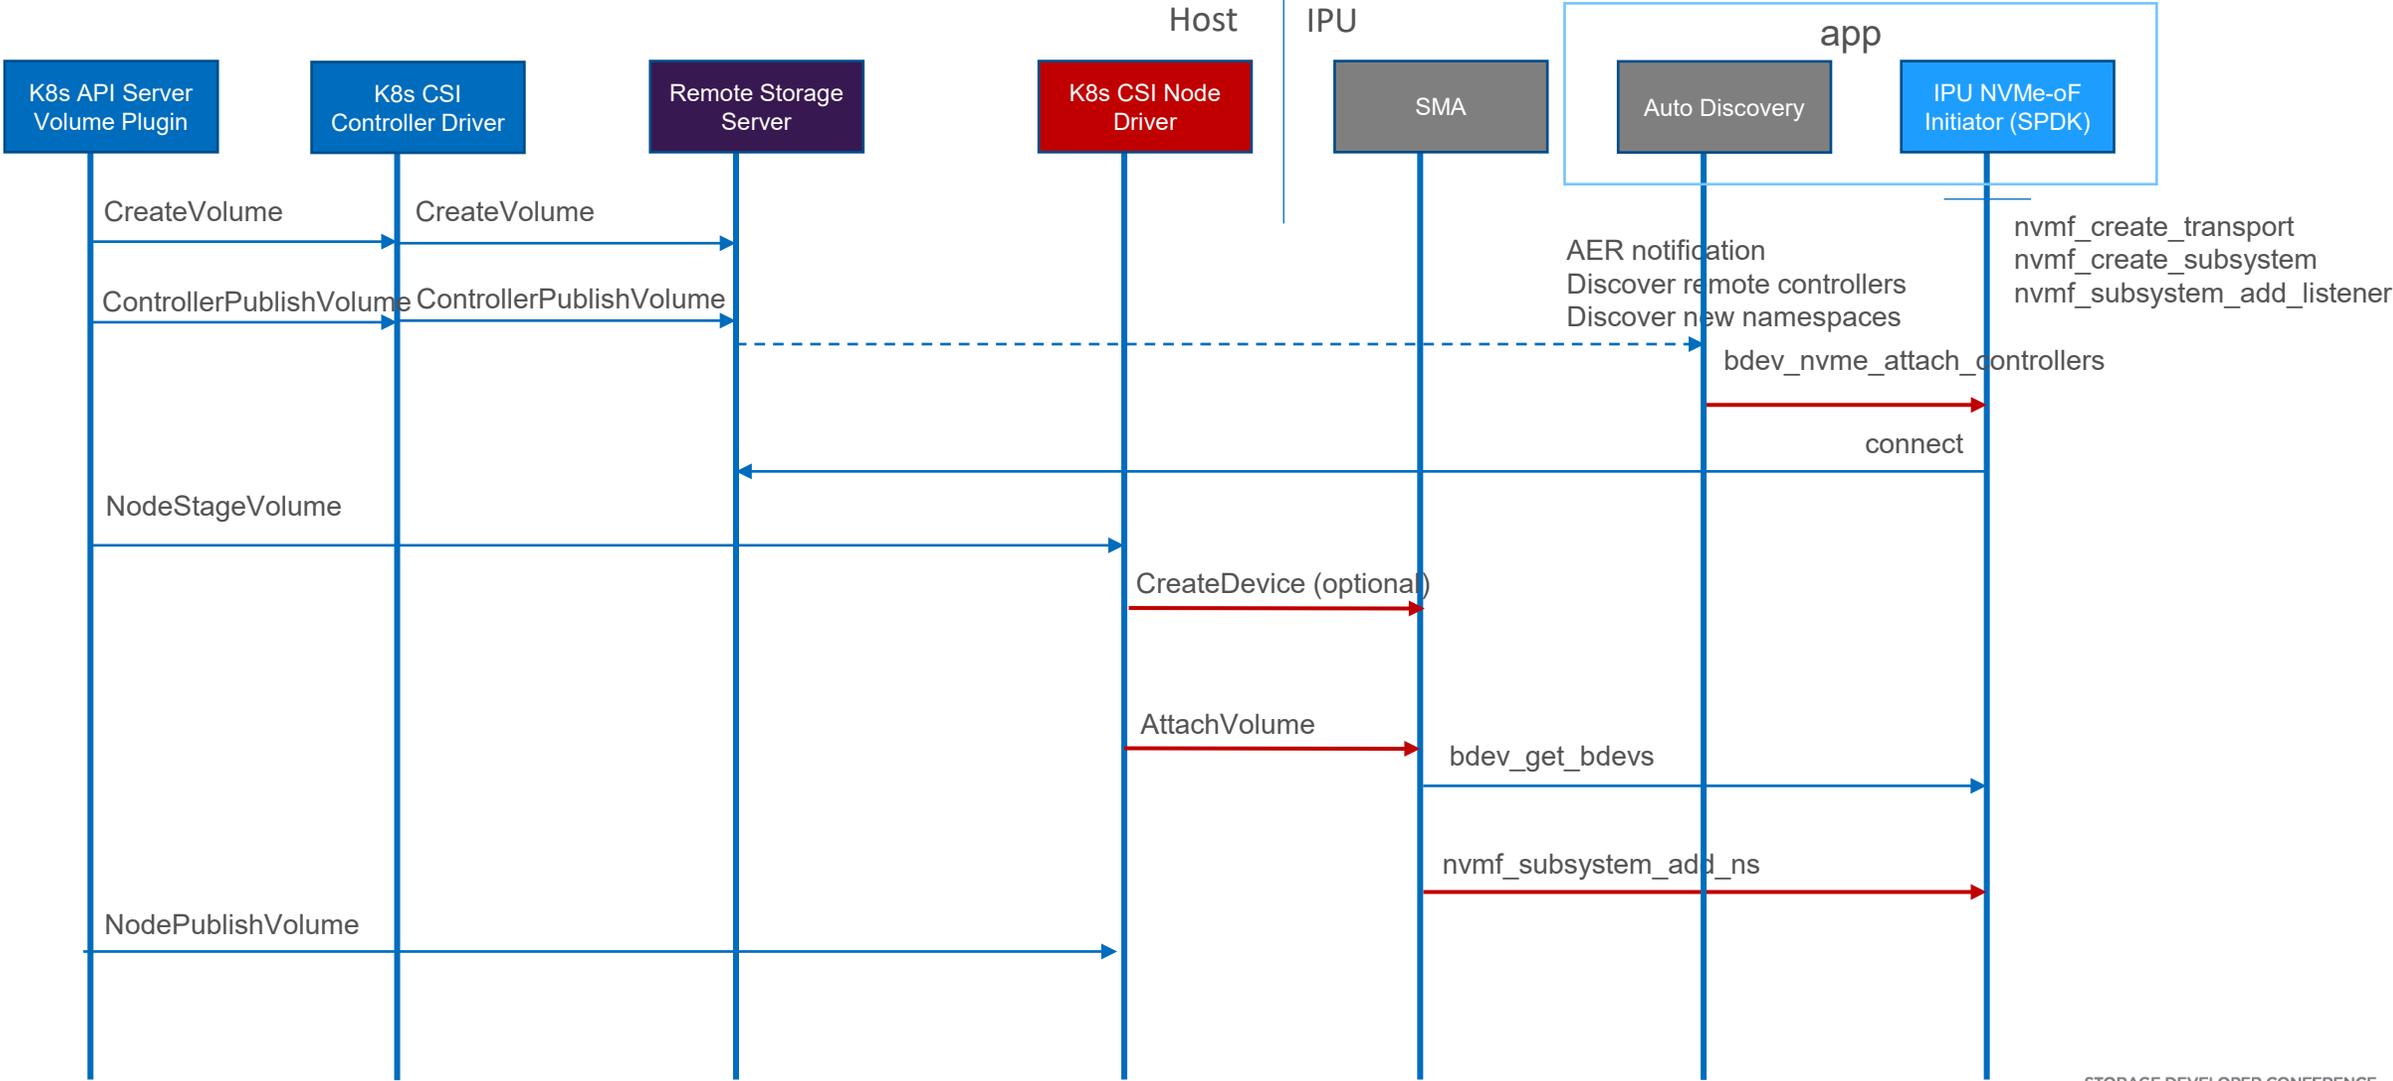
SMA API

- Basic configurations (available in SPDK 22.05 release)
 - CreateDevice
 - DeleteDevice
 - AttachVolume
 - DetachVolume
- Volume based configurations
 - SetVolumeQos, SetDeviceQos
 - SetVolumeE2EProtection
 - SetVolumeEncryptionKeys
 - GetVolumeStats
- Live Migration API
 - SaveDeviceState
 - RestoreDeviceState

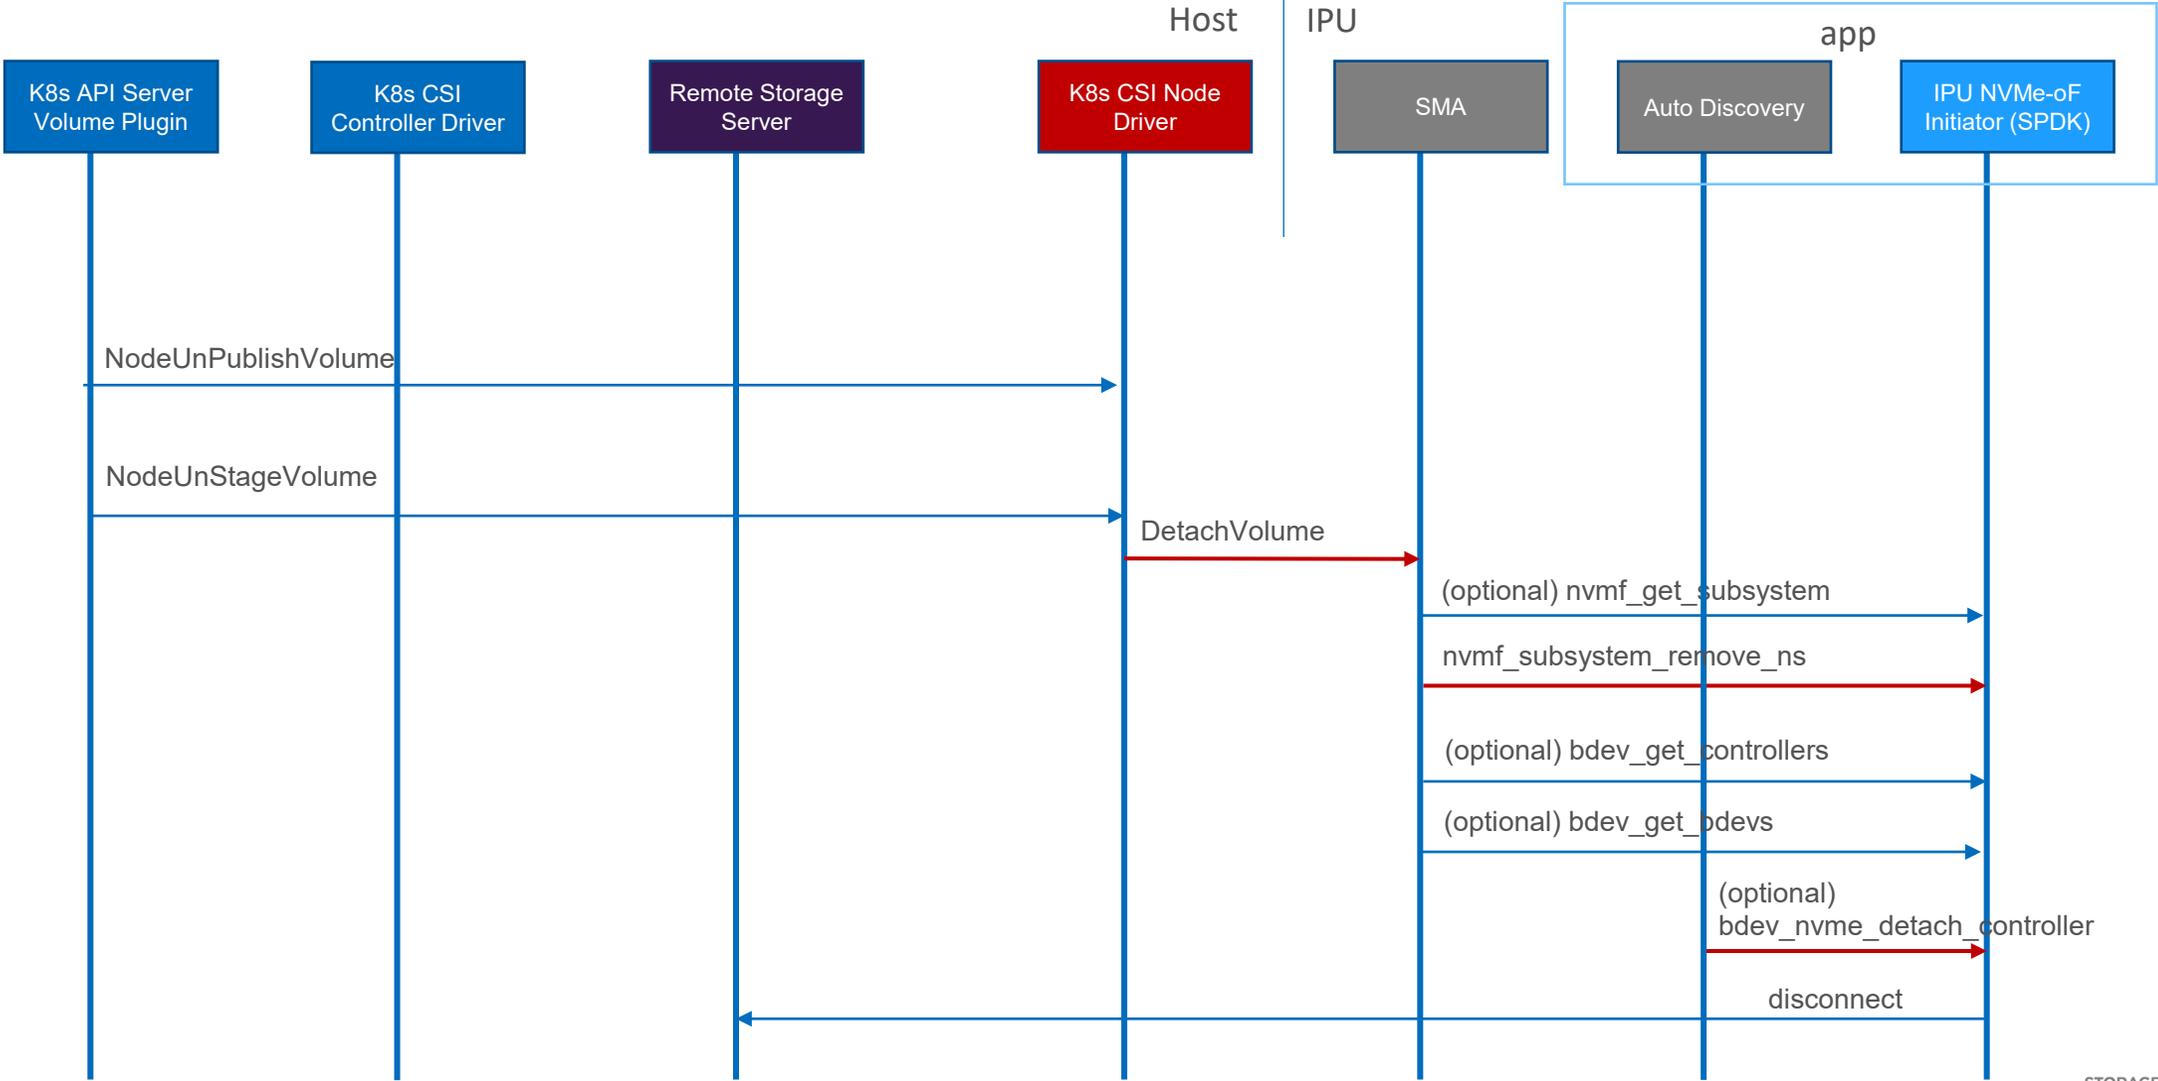
Previous Approach – Integration with SPDK RPCs



SMA based Approach - Attach Volume



SMA Based Approach - Detach Volume



Work in Progress

- SMA API for Data-at-Rest Encryption and QoS configuration
- SMA API for live migration support
- SR-IOV VF orchestration and management
- SMA API usage in OPI (Open Programmable Infrastructure) project

Summary

- IPU based NVMe/TCP initiator enables bare-metal hosting and storage disaggregation solution.
- Lightbits Cloud Data Platform provides a high performance SDS solution based on NVMe/TCP.
- IPU + Lightbits storage service provides a true end2end SDS solution.
- SMA simplifies integration with IPU, we welcome any enhancement and contribution to the SMA API implementation.



Please take a moment to rate this session.

Your feedback is important to us.