SNIA DEVELOPER CONFERENCE

SDC 25

By Developers FOR Developers

Hyatt Regency Santa Clara, CA
September 15-17, 2025

Open Flash Platform:

An Initiative for Open, Highly Efficient AI Storage

David Flynn
Founder & CEO

www.sniadeveloper.org

The Impending Existential Data Challenge

Current Data Storage Architectures Do Not Address the Demands of the Next Wave of AI Applications

Data Center Real Estate

Existing space constraints and new buildouts in many regions take 6+ years to execute

AI Everywhere

AI is not just a data center challenge. Edge and multi-cloud strategies create additional data chaos

Power and Cooling

The grid is already under massive strain. Power requirements expected to double in next 5 years

Compute-centrism

Previous inflection points have been driven by processor innovation; the next will require data innovation

OFP Introduces Standardized Data Storage for AI

Closed Scale-Out NAS Architecture

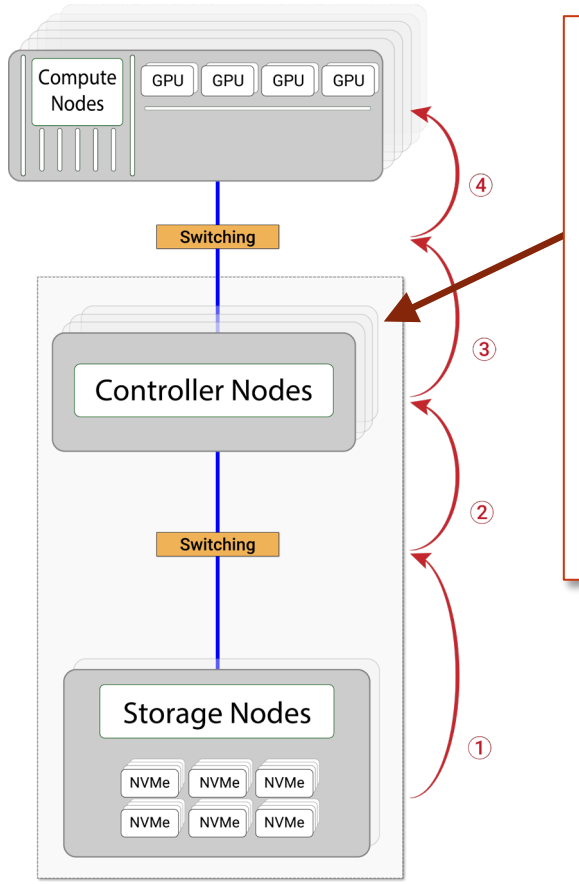

Closed Proprietary Systems

- Data funnels through Controller Nodes.
- Puts data within a proprietary silo
  - Adds latency & degrades performance
  - Limits scalability
  - Adds power, cost and complexity.

Qumulo DELL VAST NetApp HITACHI

Open Architecture

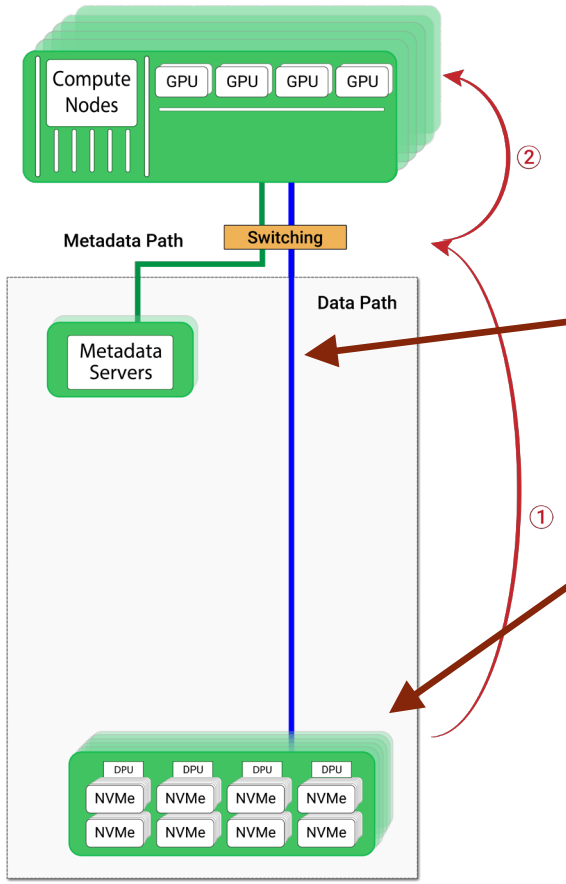

pNFS v4.2

- Separates Data Path from Control / Metadata
- Standard Linux – No proprietary client software or storage servers needed in the data path
- Ensures Extreme Linear Scalability

OFP

>10x Capacity Increase

60% Cost/Capacity Reduction

90% Power Savings

OPEN FLASH PLATFORM

SDC 25

Open vs. Closed Storage Architecture

Closed Storage Platforms	Open Storage Platform
Data services are done in-band, in the data path on purpose-built server hardware with significant computing resources. Component cost drives the need for replacement.	**Utilizes standard commodity hardware, optimized for efficiency, data services are performed by a separate, side-band control plane. Enables fail in place to reduce service overhead.**
Data is written to a storage server and then to media in a proprietary format.	**Data is written via the most direct path, from AI servers to storage media.**
Data orchestration and management are run from the storage server.	**Data orchestration and management are out of band.**
Data services are either limited by point-to-point architecture or require local client deployment and complex networking.	**Data path uses standard Linux and standard NFS, no proprietary client or proprietary storage appliance(s).**
Scale requires excessively complex backend networking or the deployment of additional power and cooling-hungry servers.	**Scaling only requires additional lightweight flash hardware units.**

Removing the Storage Server Improves Efficiency and Reliability

Storage servers lock data into silos and consume needless resources to scale with modern requirements.

OFP uses flash as a commodity, with standard, open Linux features and a streamlined data path.

1 | No Dedicated Server Layer

With OFP, intelligence (a DPU/IPU running standard Linux and NFS) is embedded directly within the flash storage sled. There's no separate, bulky, proprietary server box to buy, power, cool, or manage for the basic task of serving data.

2 | Standard, Open Software

Each OFP sled serves data using industry-standard NFS running on Linux. You're not paying a "proprietary software tax" for basic data access at the sled level.

3 | Direct Data Path

Leveraging standard pNFSv4.2 with Flex Files, the path from a client to an individual OFP sled is streamlined, bypassing the heavy processing of a traditional NAS head. This enables extreme scale-out for AI workloads.

What Is Parallel NFS with Flexible Files?

About Parallel NFS with Flex Files

- Parallel NFS (pNFS) introduced as optional feature of NFS in 2010, enhanced in later RFCs
- Defines a standards-based parallel file system architecture using NFS
- NFSv4.2 client is part of the Linux kernel since 2019
- Already in place in every data center in the world as part of standard Linux

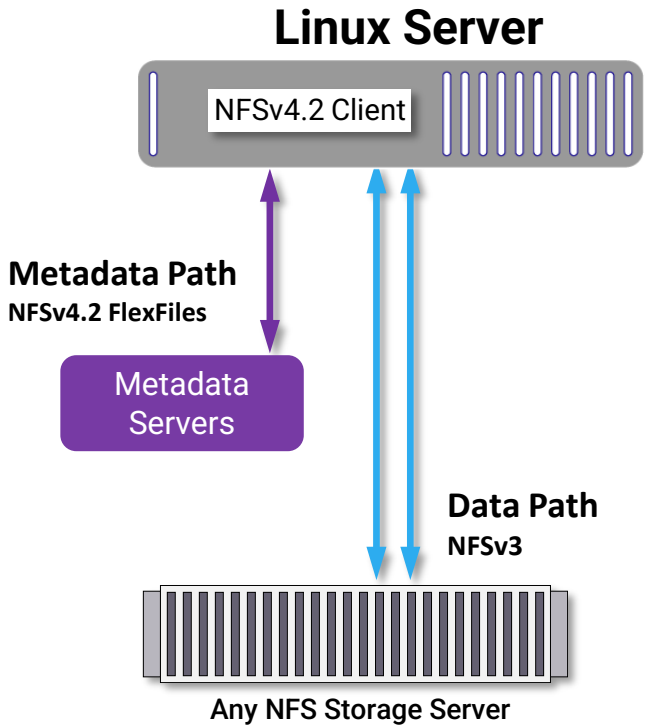

Direct data path between Linux client and storage volumes

Also provides for:

- Multiple parallel network connections between client and server
- Ability to write to multiple storage nodes synchronously (mirroring)
- Ability to move data while it is being accessed, without interruption
- Timestamp generation delegation
- Eliminates protocol chatter
- File granular access and telemetry
- TCP and RDMA Supported

The OFP Sled – Using U.2 NVMe

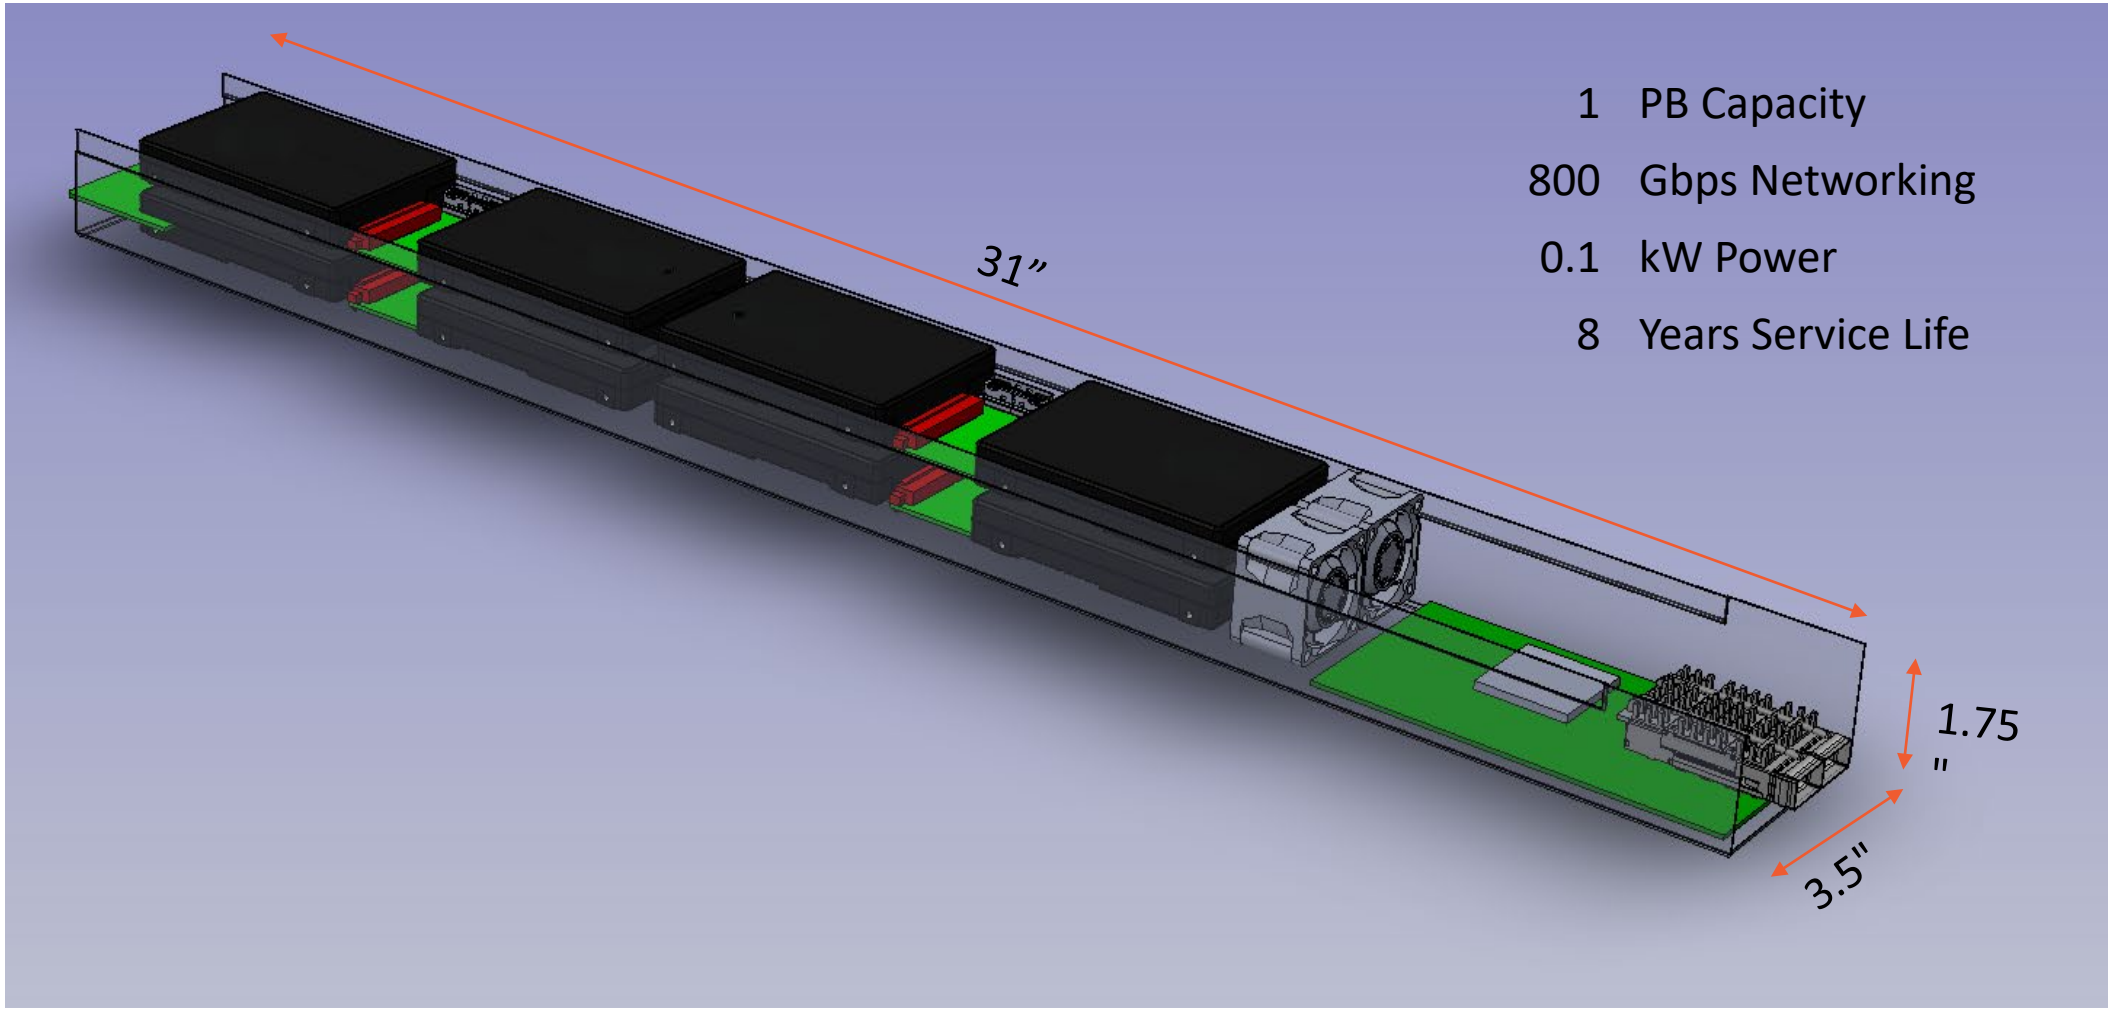

The OFP Sled – Using E2e NVMe

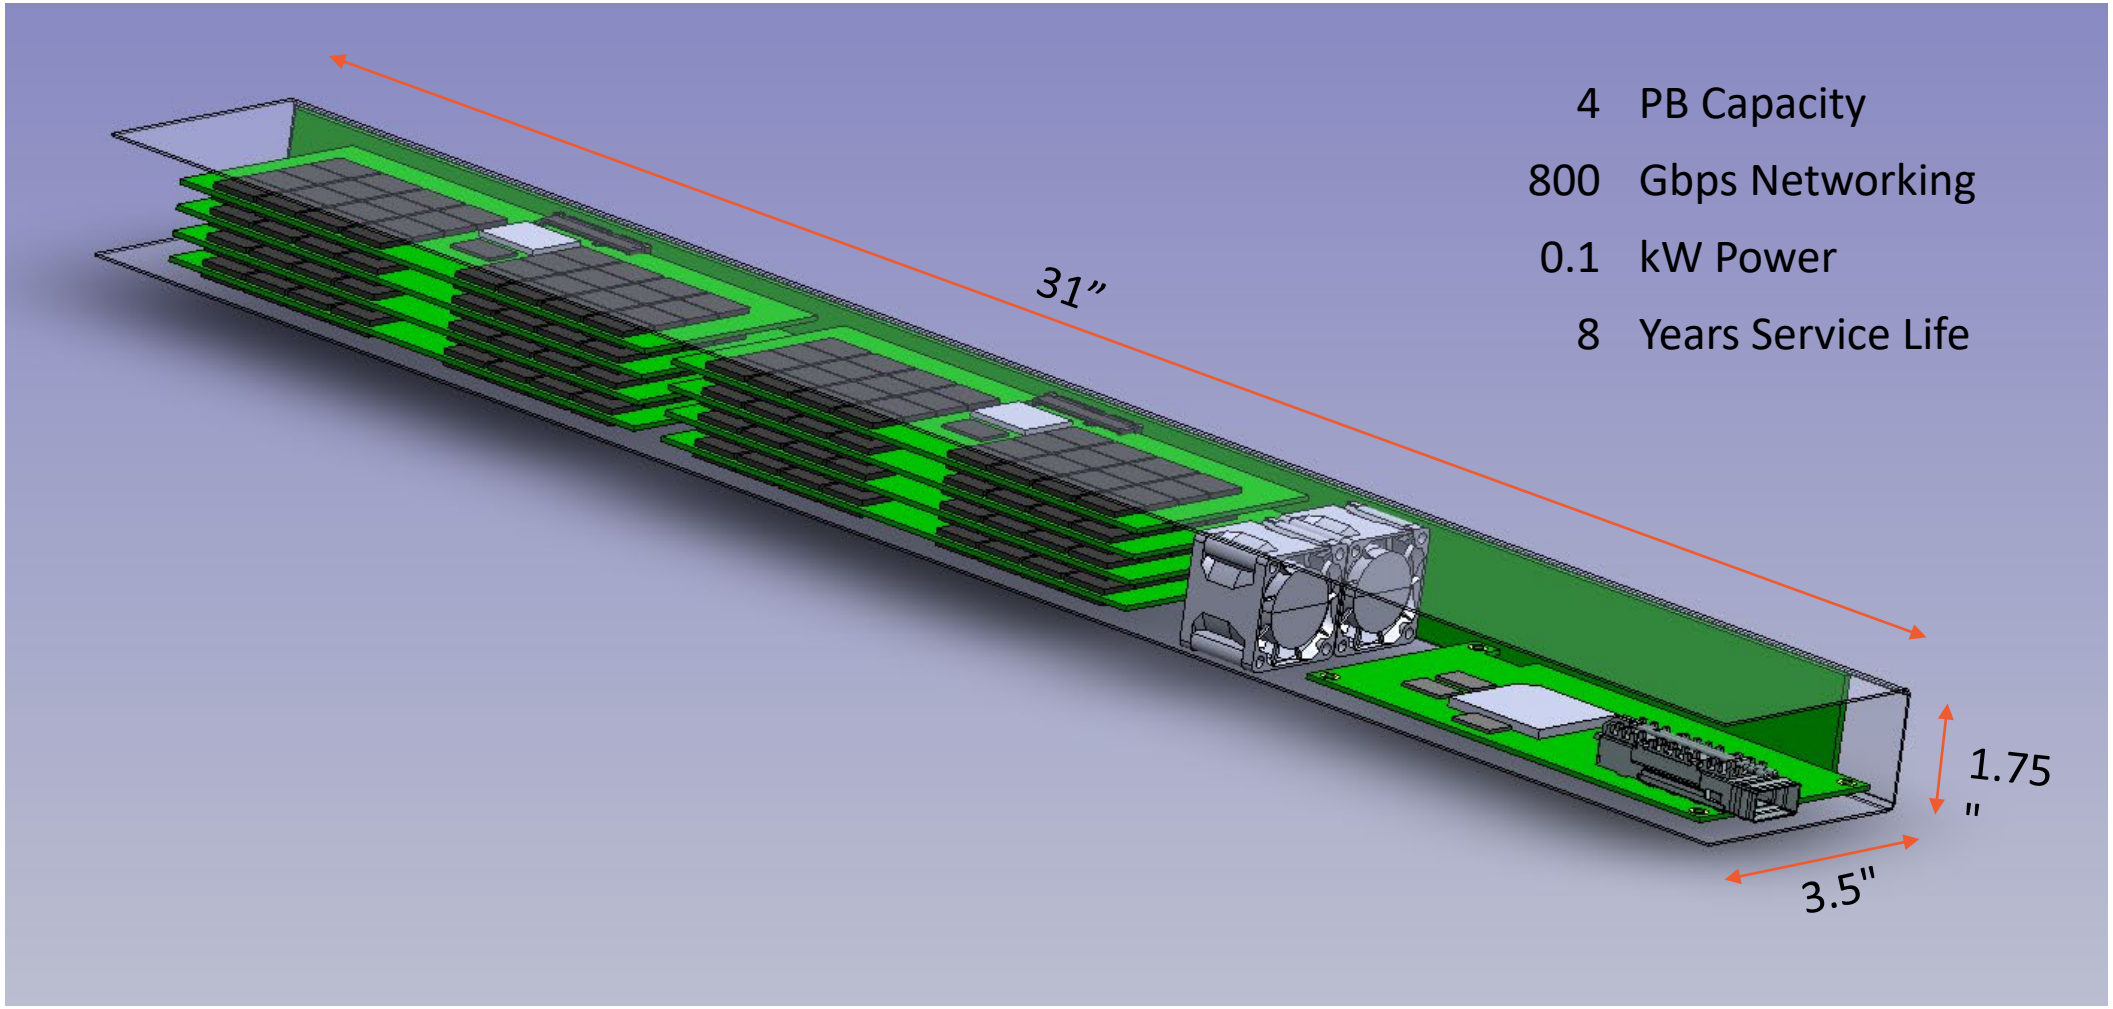

The 1U 21" OFP Tray

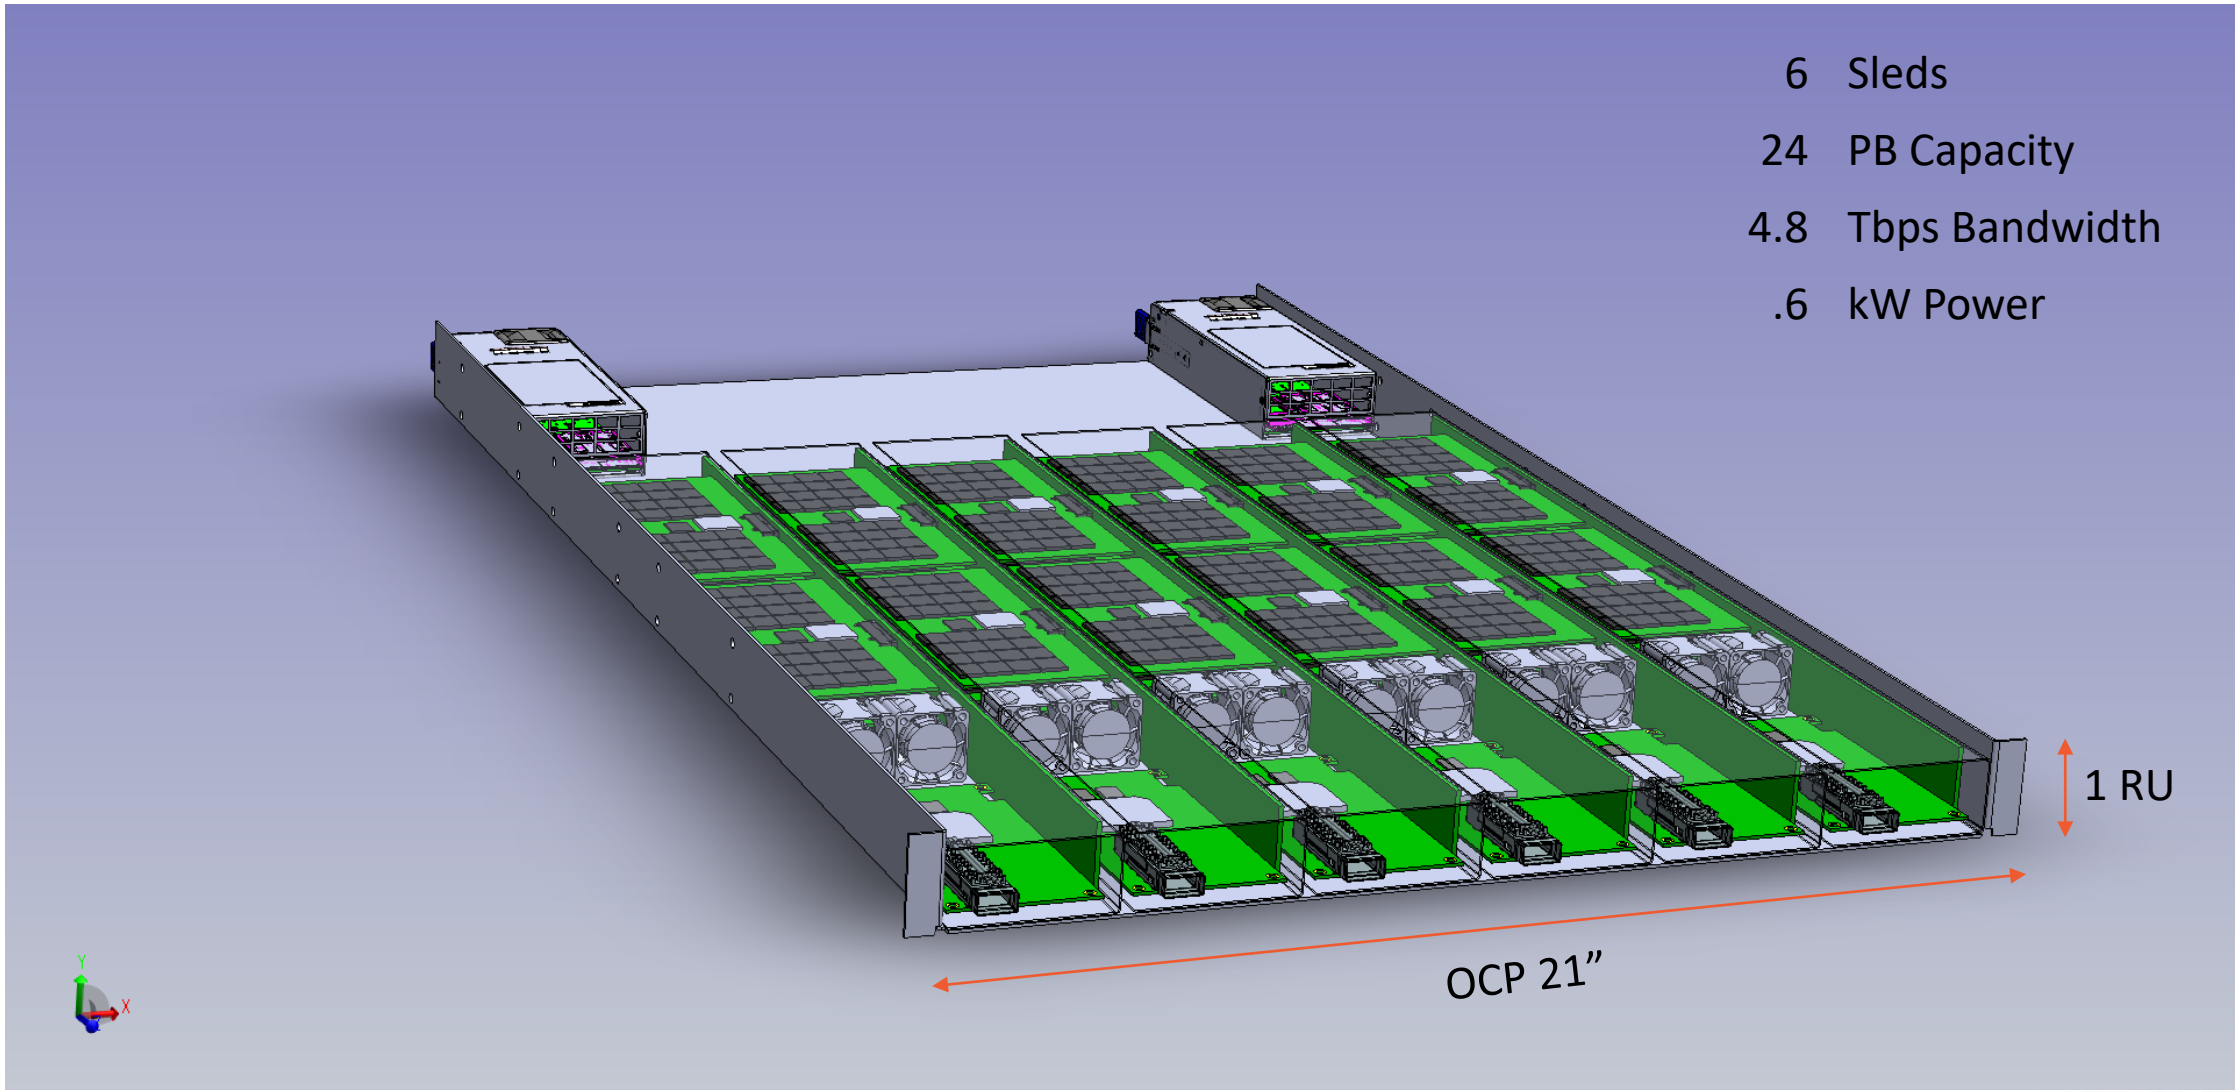

Open Flash Sled – Granular Expansion

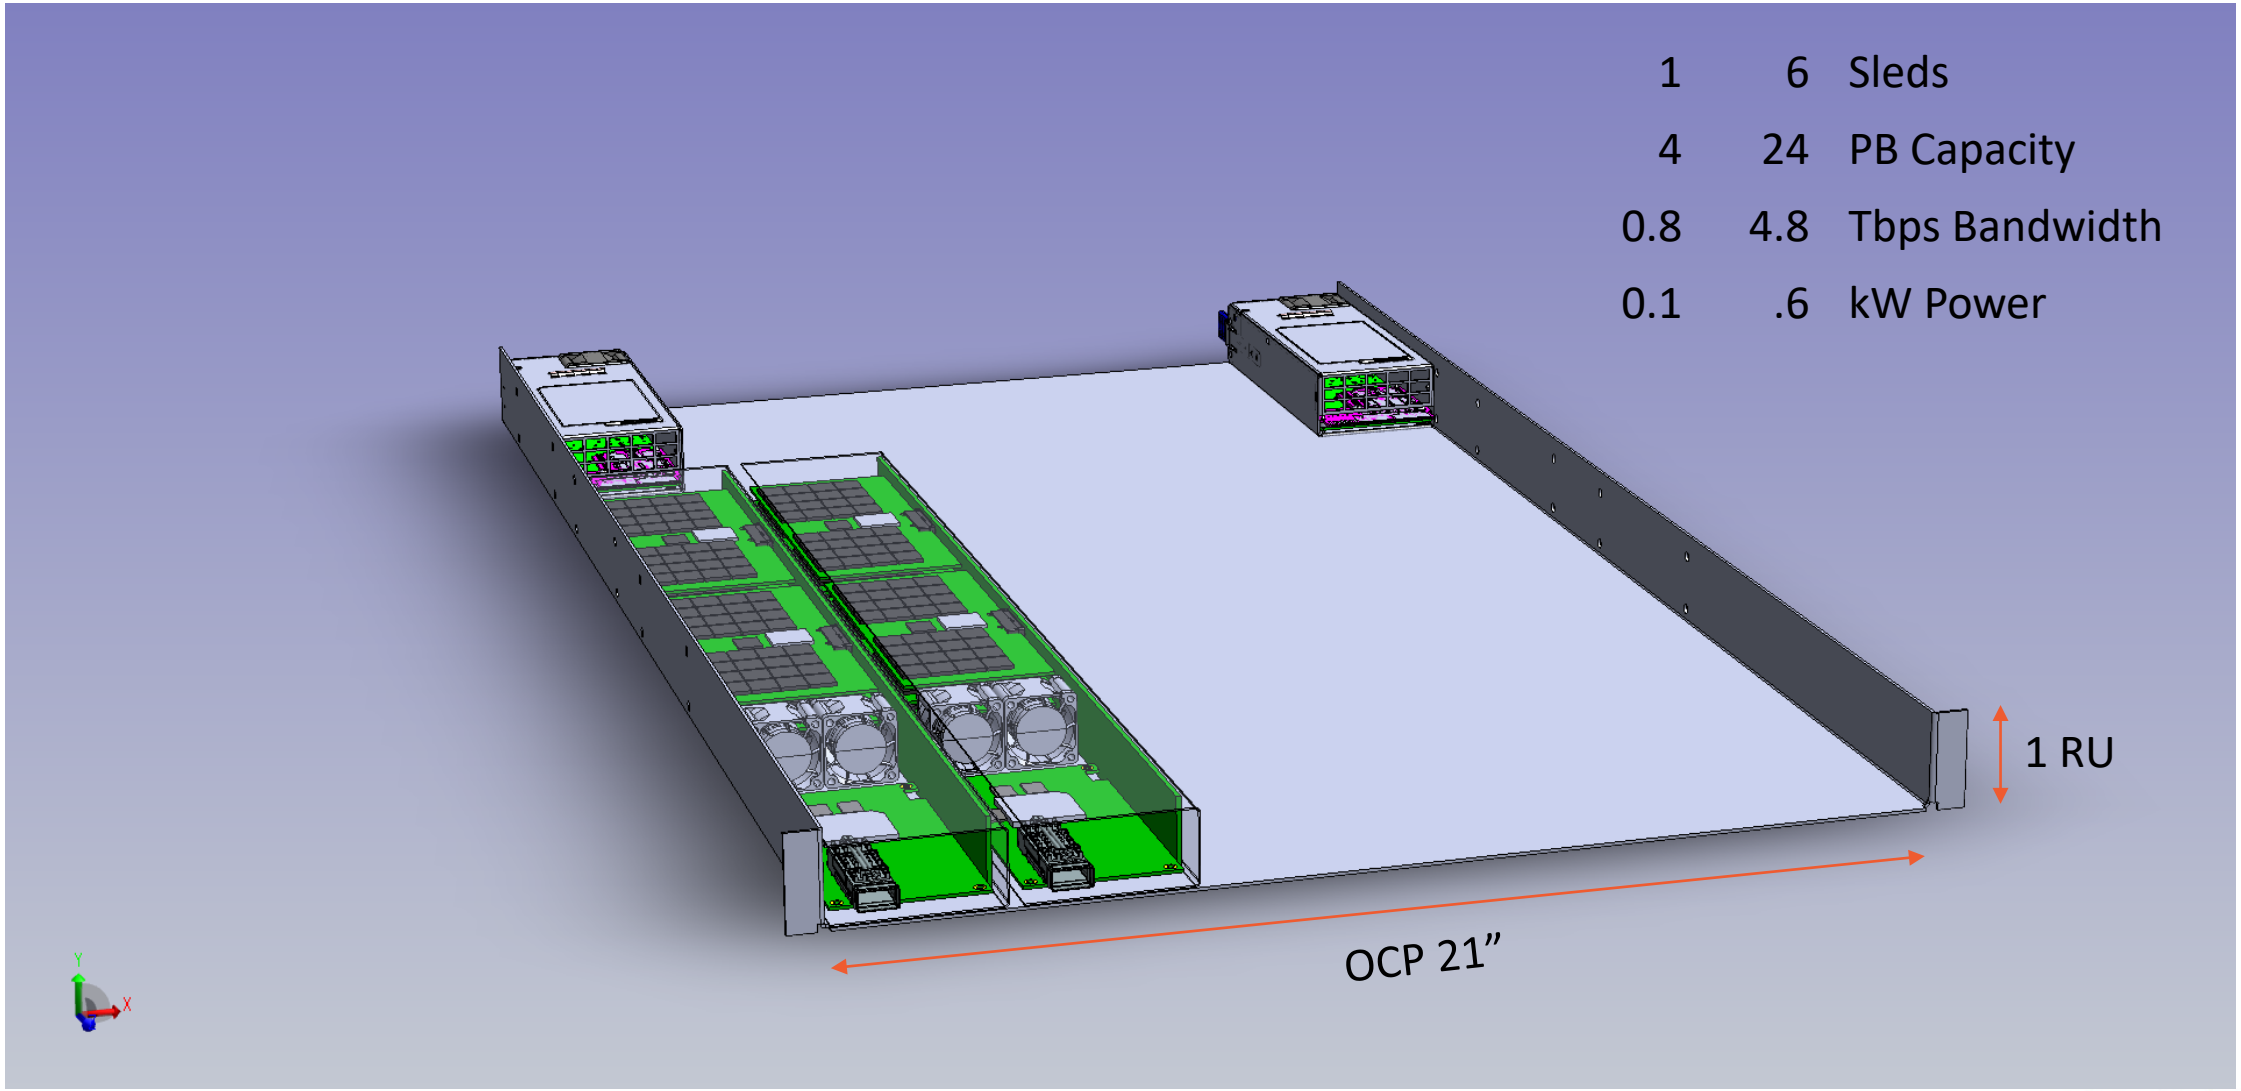

OFP Delivers Incredible Simplicity and Efficiency

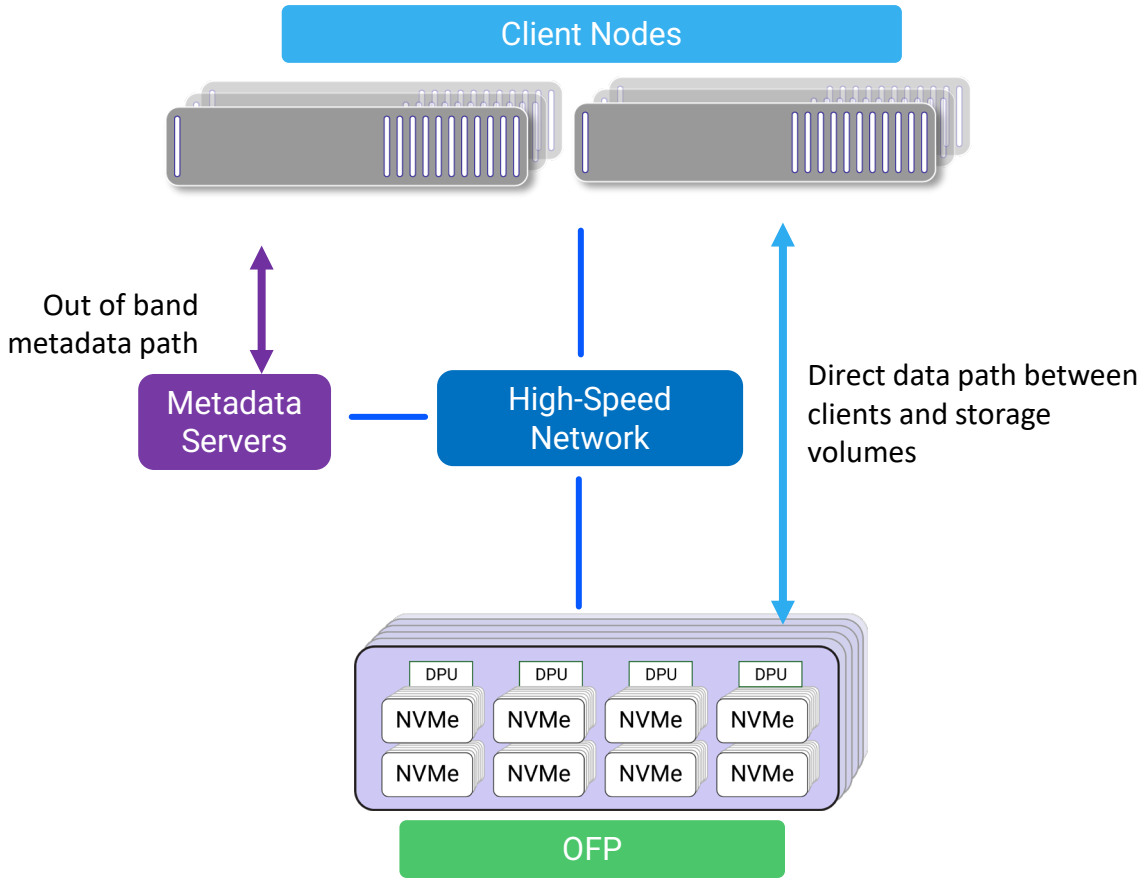

- Multiple parallel network connections between client and server
- Ability to write to multiple storage nodes synchronously (mirroring)
- Ability to move data while it is being accessed, without interruption
- Timestamp generation delegation
- Eliminates protocol chatter
- File granular access and telemetry
- TCP and RDMA Supported

Open Flash Platform vs All-Flash Scale-Out NAS

50.16PB
capacity

40
Rack Units

29.1kW
power

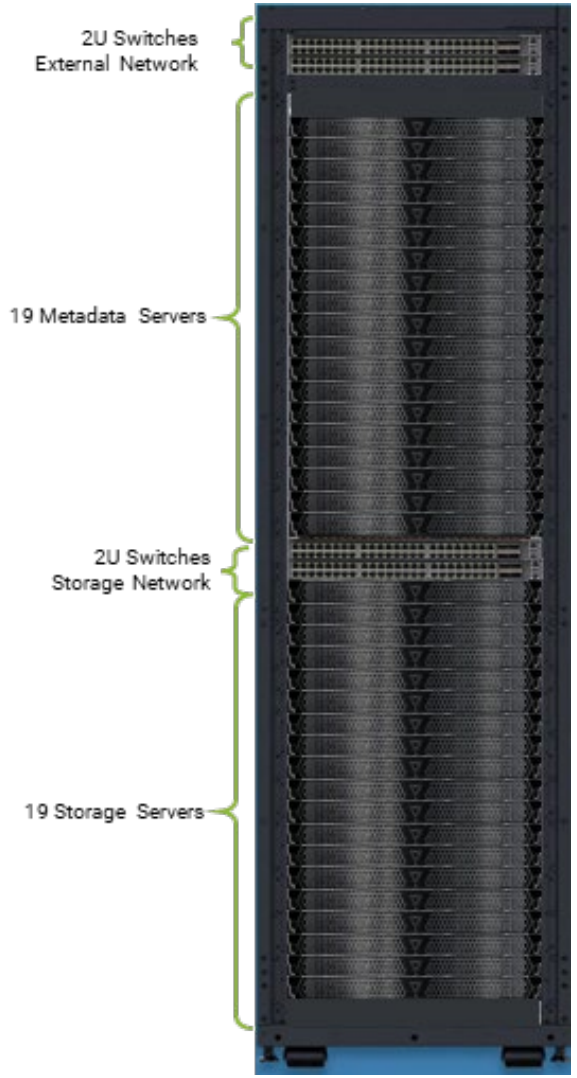

Greater than
10x
density

96%
power savings

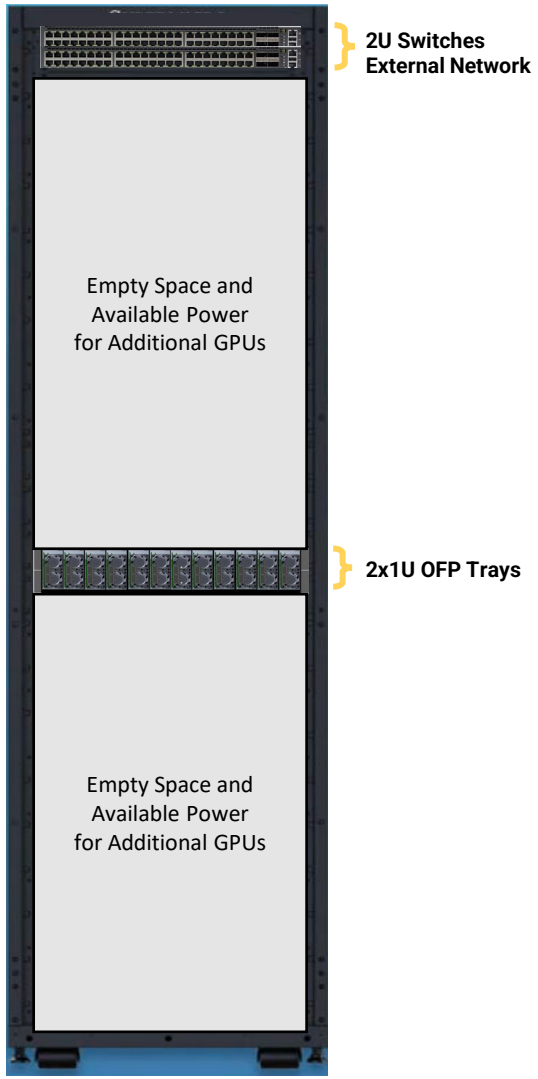

48PB
capacity

2
Rack Units

1.2kW
power

Open Flash Platform vs All-Flash Parallel File System

50.6PB
capacity

84
Rack Units

72.5kW
power

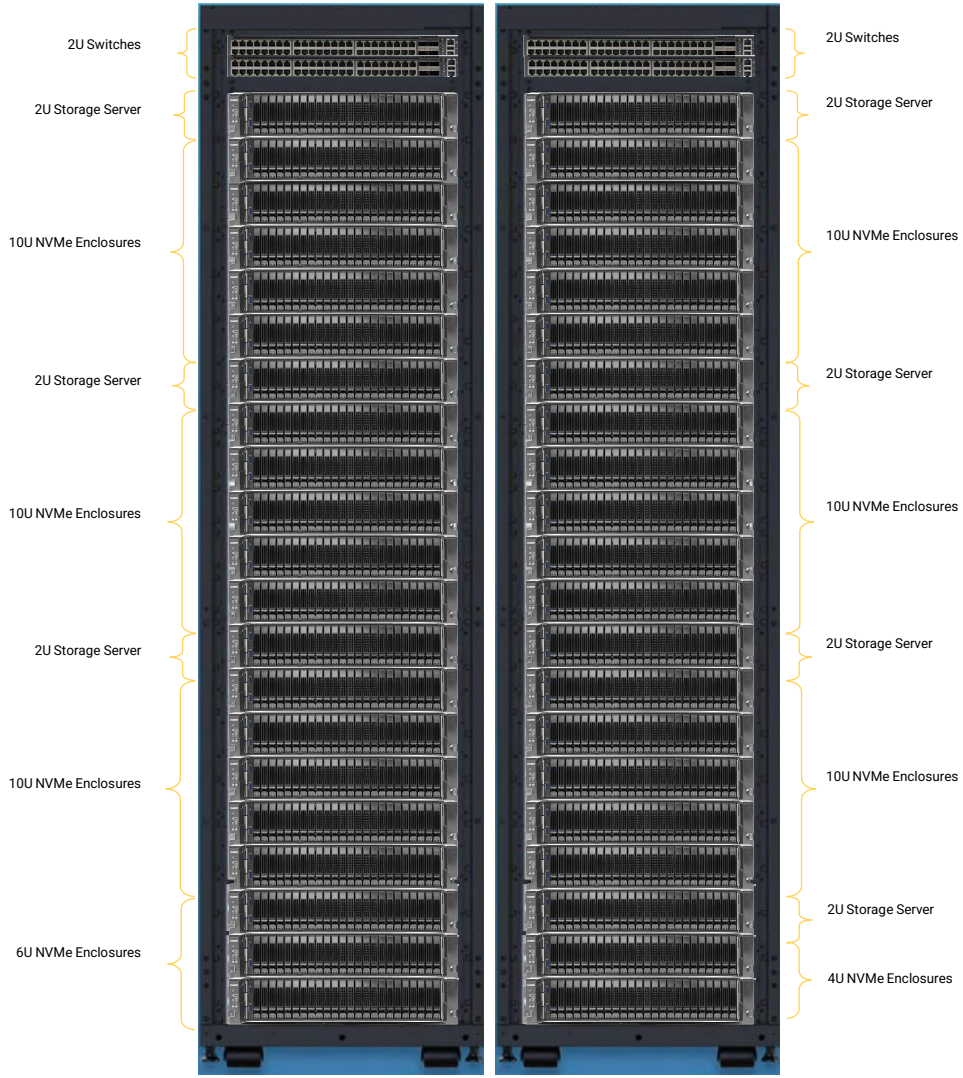

Greater than
40x
density

98%
power savings

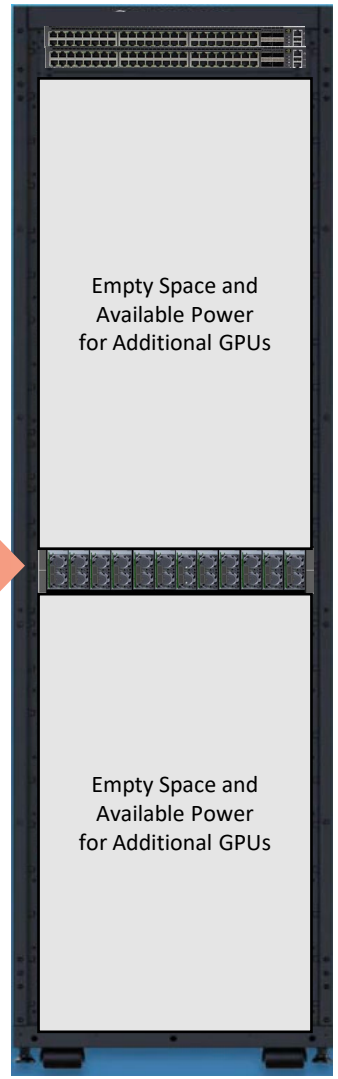

48PB
capacity

2
Rack Units

1.2kW
power

The OFP Exabyte Rack

OCP V3 Rack

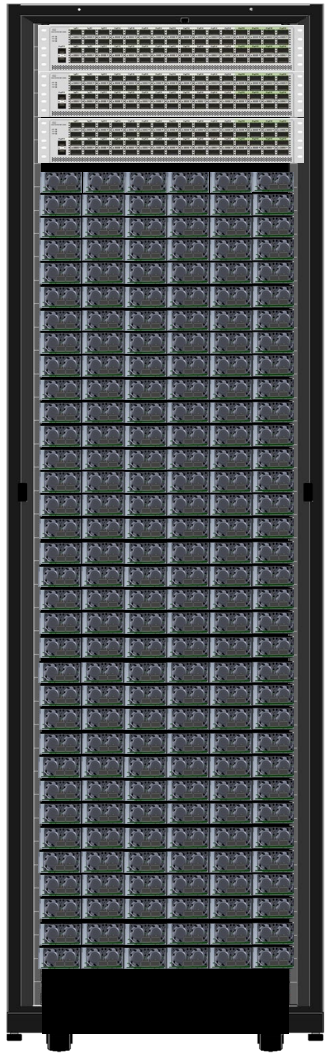

252 OFP Sleds

1 EB Capacity

200 Tbps Bandwidth

25 kW Idle Power

40 PB per kW

8 Tbps per kW

The OFP Exabyte Rack

	Storage Server-based	Open Flash Platform	Improvement
Physical Capacity	1 EB	1 EB	
CapEx	$140 M	$98 M	30%
Rack Space	10 Racks	1 Rack	90%
Power	250 kW	25 kW	90%
Servicing per year	1,000 Hours	100 Hours	90%
OpEx	$4.2 M	$0.42 M	90%
Service Life	6	8	33%
TCO / PB / year	$27 M	$12 M	55%

OFP Initiative Participants

"Power efficiency isn't optional; it's the only way to scale AI. Period. The Open Flash Platform removes the shackles of legacy storage, making it possible to store exabytes using less than 50 kilowatts, vs yesterday's megawatts. That's not incremental, it's radical."

- Hao Zhong, CEO & Co-Founder of ScaleFlux

"Flash will be the next driving force for the AI era. To unleash its full potential, storage systems must evolve. We believe that open and standards-based architectures like OFP can maximize the full potential of flash-based storage systems by significantly improving power efficiency and removing barriers to scale."

- Hoshik Kim, SVP, Head of Memory Systems Research, SK Hynix

"In AI, moving data fast is just as critical as storing it. Open architectures powered by advanced DPUs turn the network into an engine for agility and efficiency."

- Eric Vallone, VP Business Development, XSIGHT Labs

"Agility is everything for AI — and only open, standards-based storage keeps you free to adapt fast, control costs, and lower power use."

- Gary Grider, Director of HPC, Los Alamos National Lab (LANL)

"Open, standards-based solutions inevitably prevail. By delivering 10x greater capacity density and a 50% lower TCO, OFP accelerates that inevitability."

- David Flynn, Founder and CEO, Hammerspace

"pNFS inside a DPU makes efficient AI storage a reality: fully Open Flash Platform, Linux-native file storage that scales without adding power-hungry servers or complex network fabric."

- Trond Myklebust, NFS Client Kernel Maintainer

Key Takeaways

The long-term solution for AI data must be designed for efficiency, reliability, locality, with sufficient performance to maximize application resources.

Data must be centrally managed, but widely accessible and potentially highly distributed.

Storage resources must be significantly more basic, efficient, lower cost, and open.

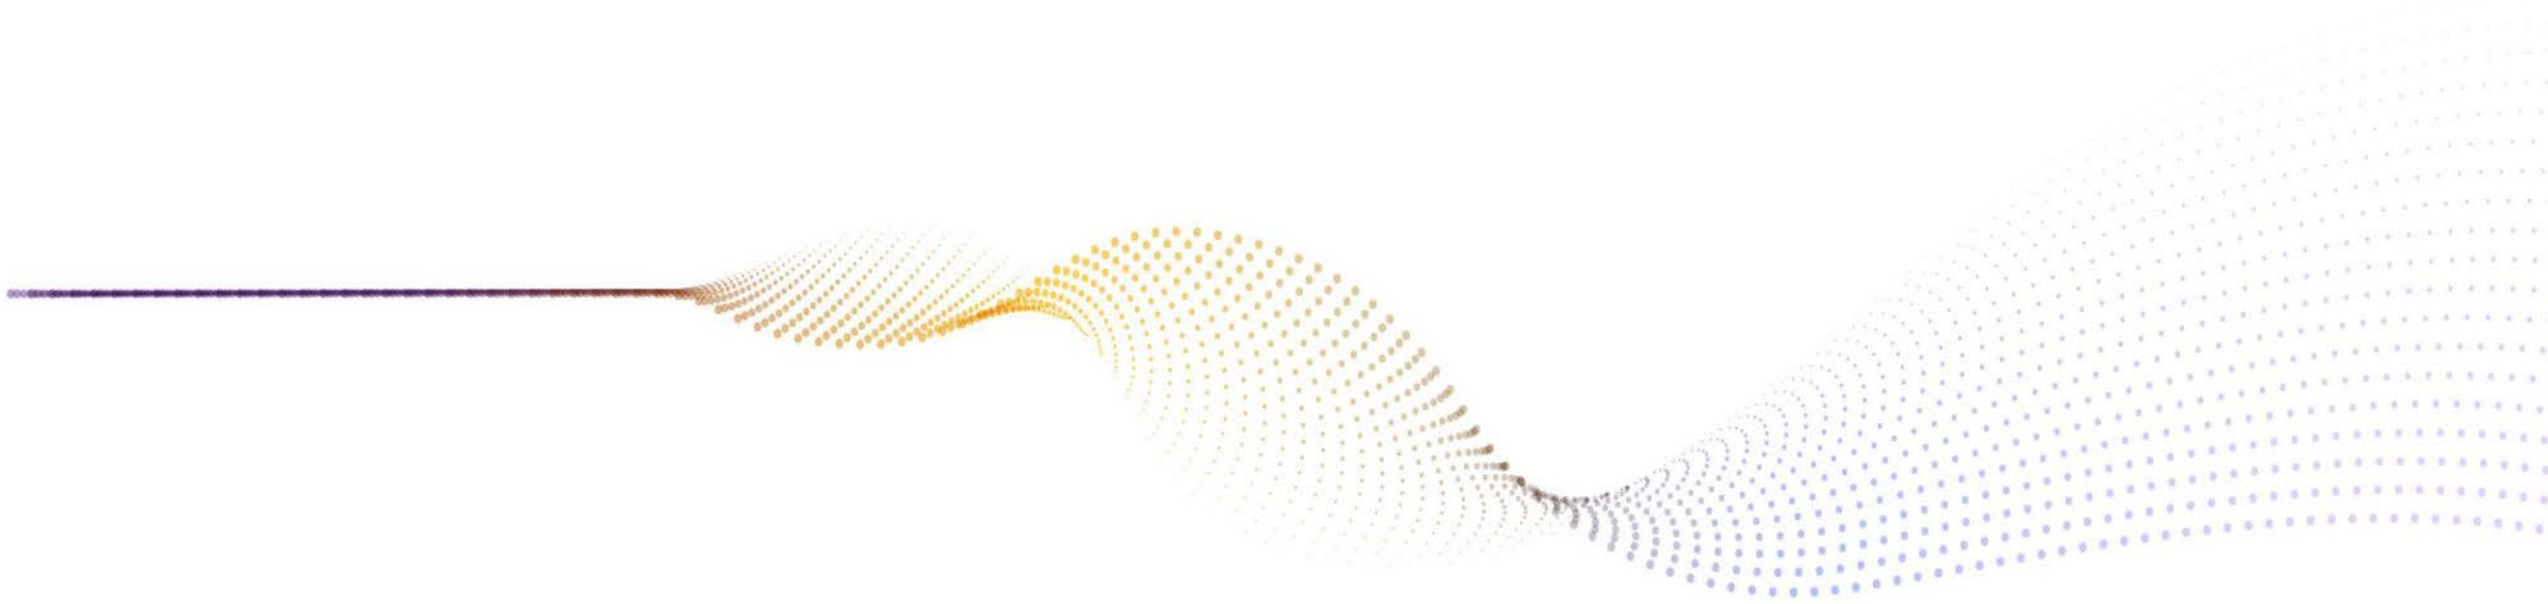

Thank you for attending!

Please remember to rate this session. You get access the presentations at
http://sniadeveloper.org/conference

