

The Need for a New Storage Architecture: **Hyperscale NAS**



Awards & Recognition



- Founder and CEO, David Flynn
 - Previous Founder and CEO of Fusion-IO
- Founder and CTO, Trond Myklebust
 - Linux Kernel Maintainer for NFS Client
- 9 years of R&D investment
 - Global file system
 - Advanced metadata technologies
 - Enterprise data management services
- Funding
 - \$20 million personal investment in Hammerspace
 - \$57M Series A Institutional Fund Raise – July 2023



Hammerspace Global Data Environment

PERFORMANCE AND SCALE

HPC/AI performance for mixed IO workloads

Scale linearly from a few nodes to thousands of nodes

Hyperscale NAS

Standards-based enterprise NAS features and RAS

Software-defined and storage agnostic

AGILITY AND PROTECTION

Spans multiple locations: edge-core-cloud

Automatically move data seamlessly and transparently

Data Orchestration

Assimilate data instantly from any file or object store

Harness rich metadata to unlock business value



Benefits of Parallel NFS 4.2 with Flex Files

SIMPLE

Built into the Linux kernel.

No special software required.

~~Client~~

FAST AND EFFICIENT

Eliminates excess protocol chatter to reduce network traffic and trips (vs NFSv3).

Files can be moved while they are being accessed without interruption.


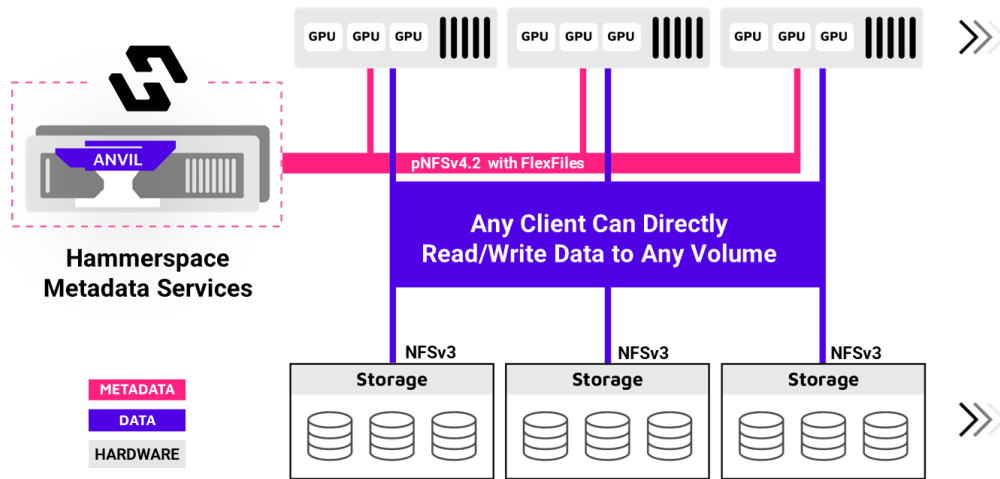
PARALLEL BY NATURE

Provides for multiple parallel network connections between client and server.

Write to multiple storage nodes synchronously.



Hyperscale NAS Advantages over Scale-Out NAS



Hyperscale NAS Architecture

- ✔ Standard Protocols
- ✔ Enterprise Features
- ✔ Single Network
- ✔ HPC Performance
- ✔ Linear Scalability


Scale-Out NAS Architecture

- ✔ Standard Protocols
- ✔ Enterprise Features
- ✘ Multiple Networks
- ✘ Lacks Performance
- ✘ Limited Scalability

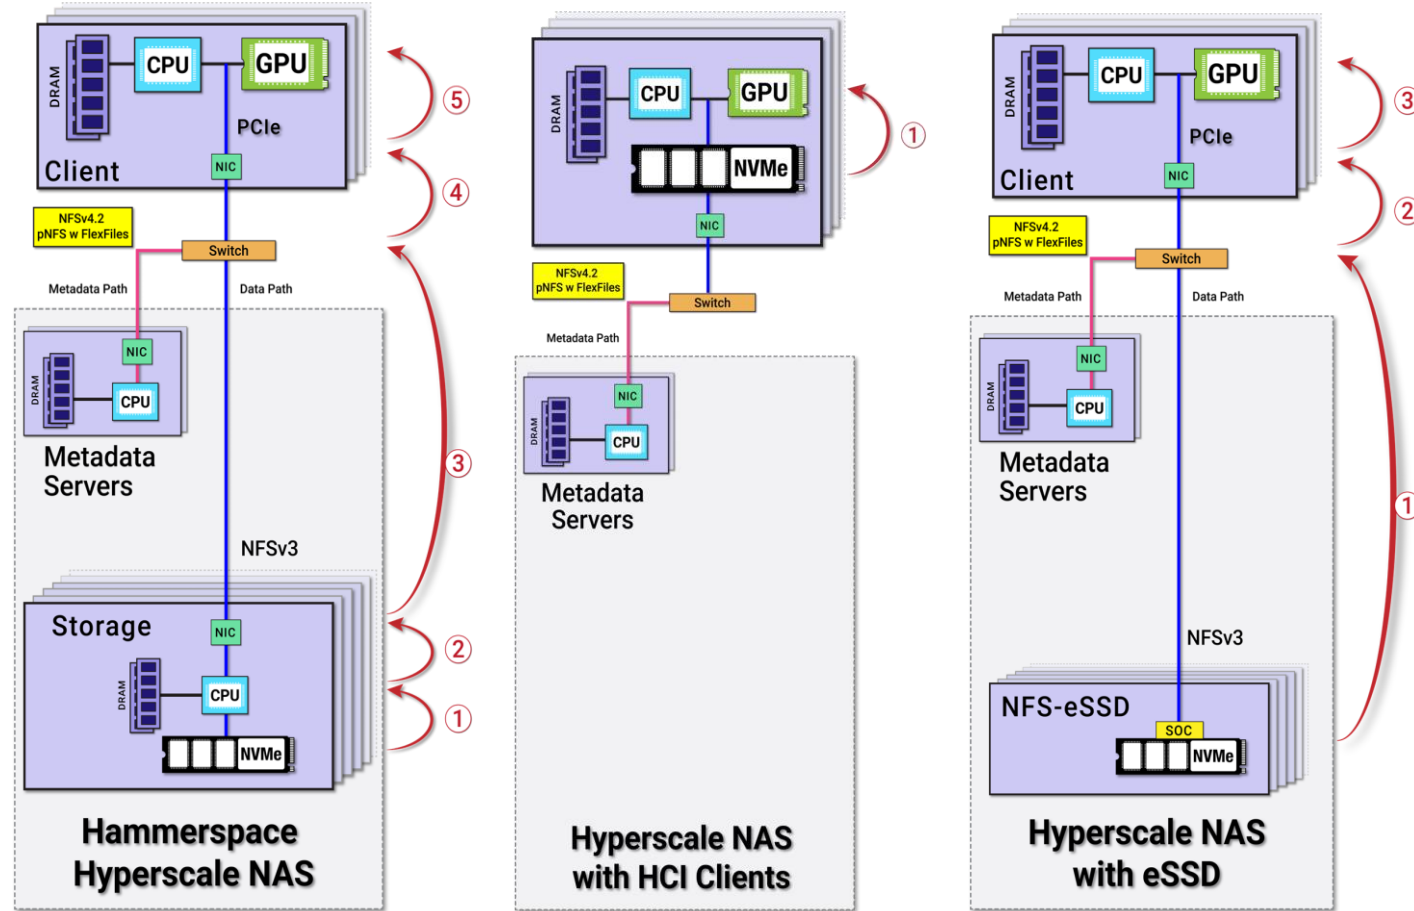


File Storage Architecture Comparison

Scale-Out NAS Architectures



Hyperscale NAS Architectures



2x
Reduction in
Servers

2x
Reduction in
Networking

2x
Reduction in
Latency

2x
Reduction in
Power

2x
Reduction in
Rack Space

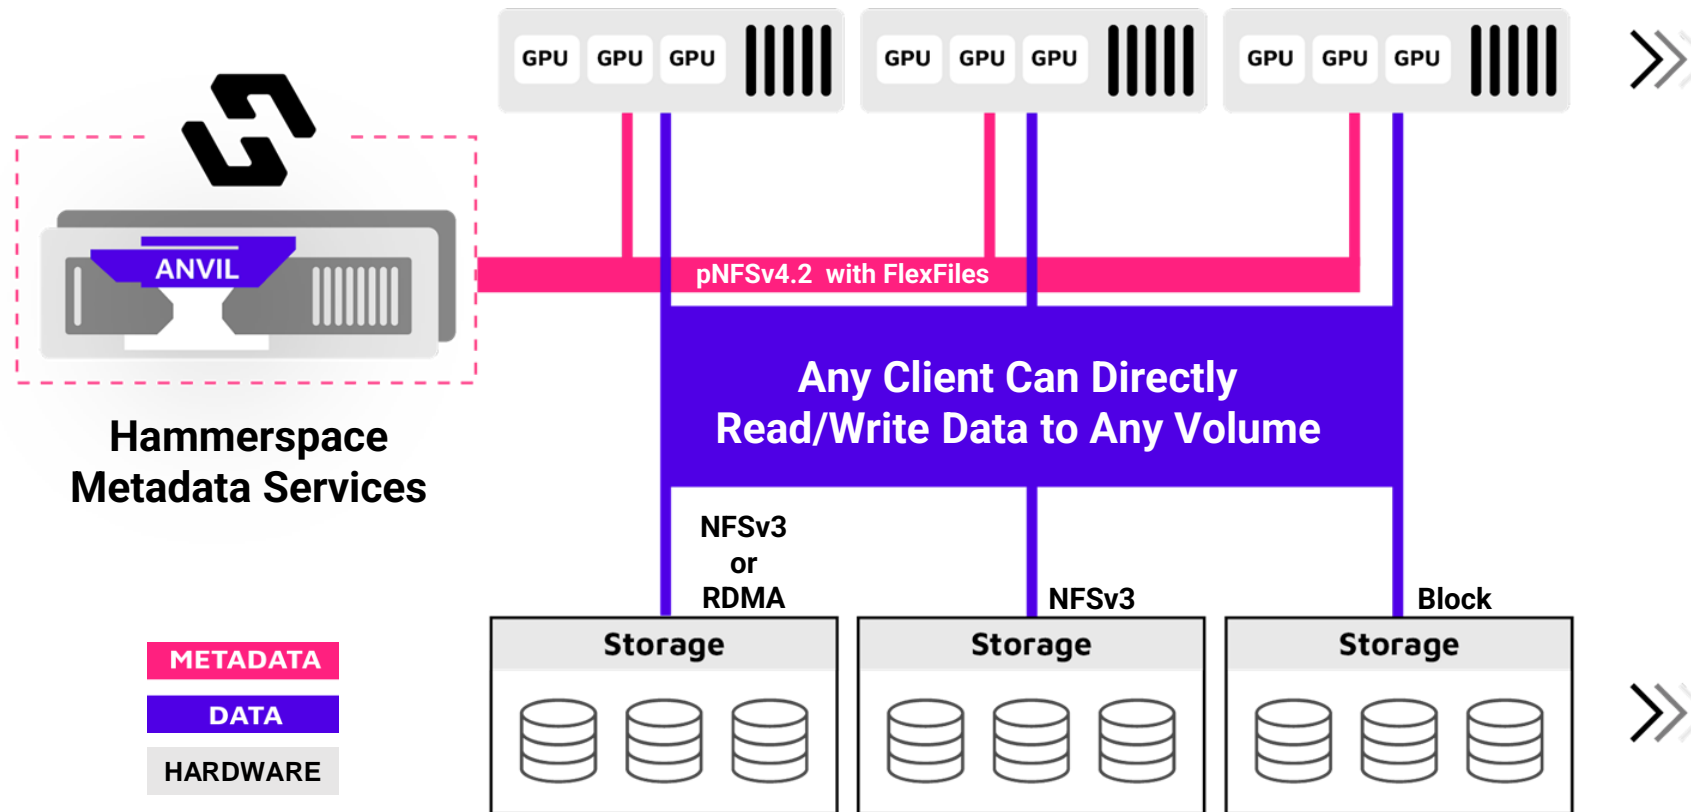


Hyperscale NAS Architecture

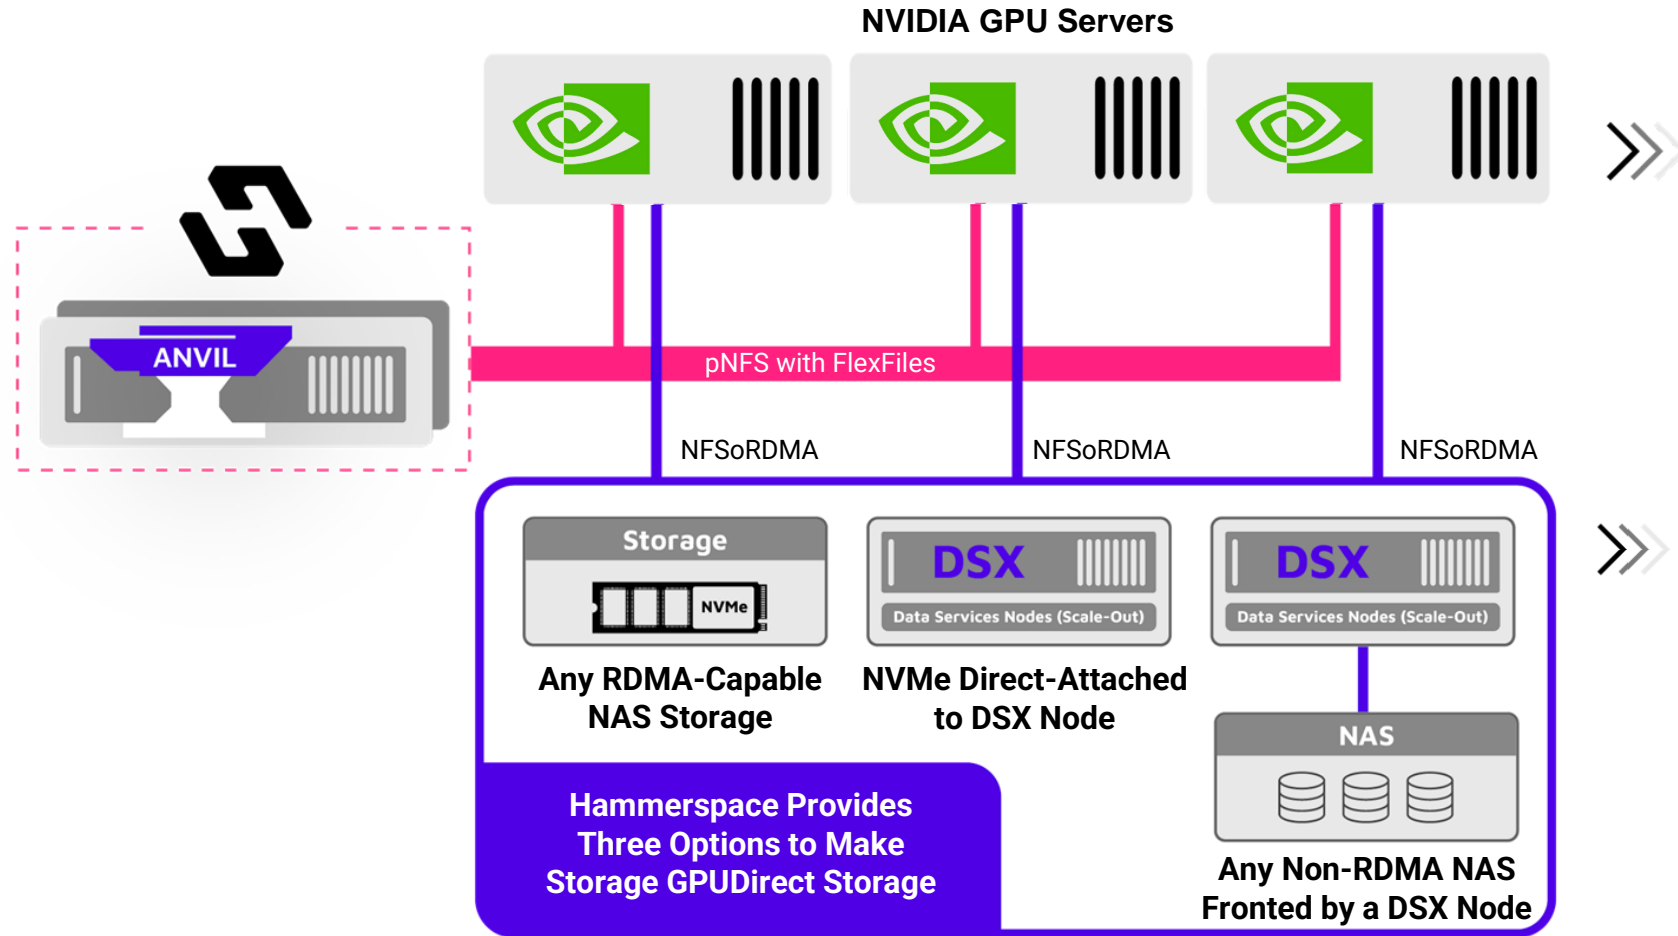
- HPC Parallel Performance
- Linear scale to 1000s of nodes

- Standards-Based
- Choice of networking

- Software-defined
- Storage-agnostic



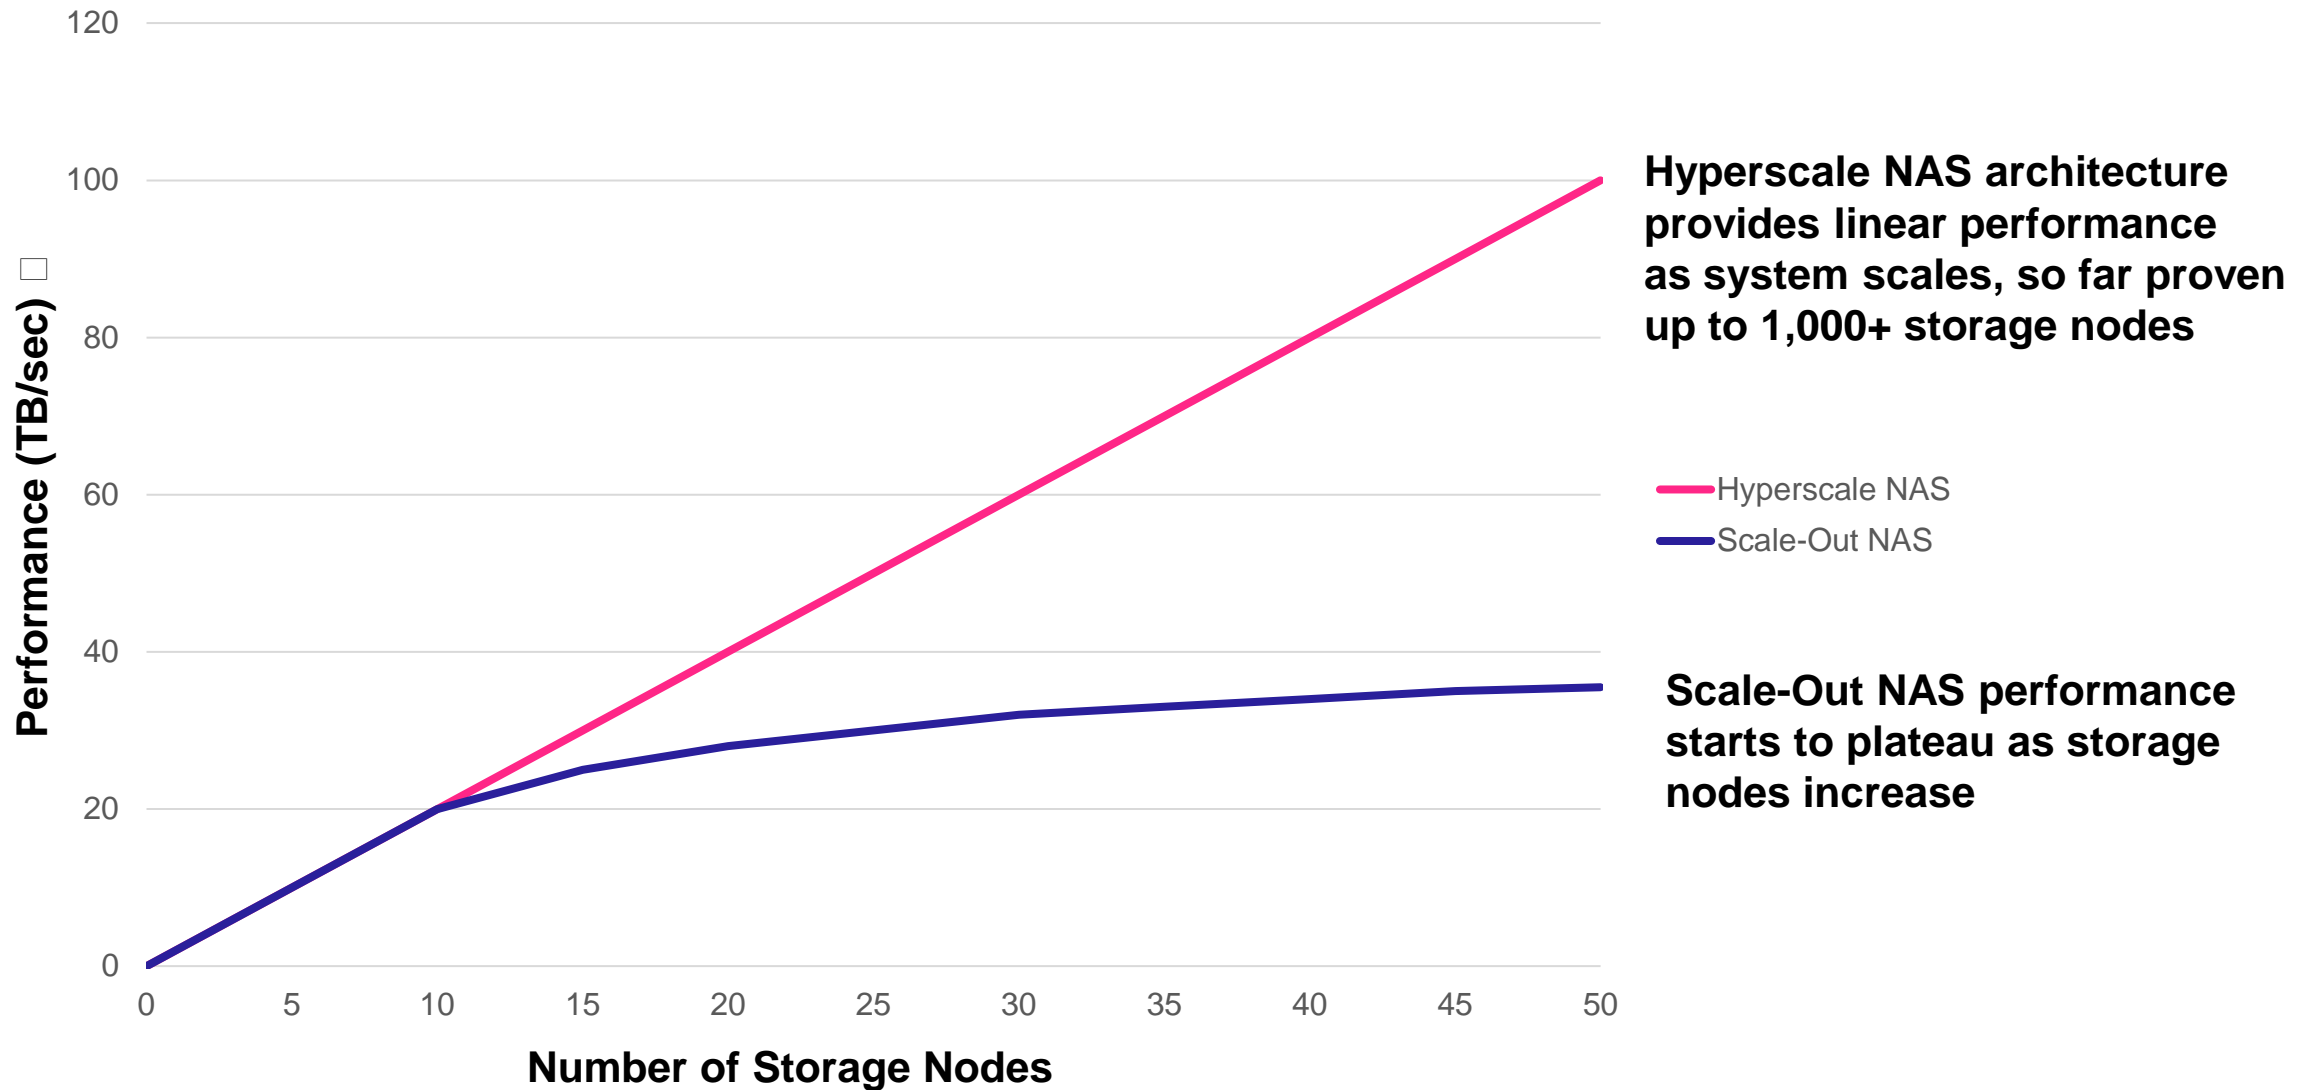
Hammerspace Can Make Any Storage GPUDirect



NVIDIA GPUDirect Storage (GDS) uses RDMA to streamline the path between GPU and storage, to improve throughput and reduce latencies



Hyperscale NAS Delivers Linear Performance



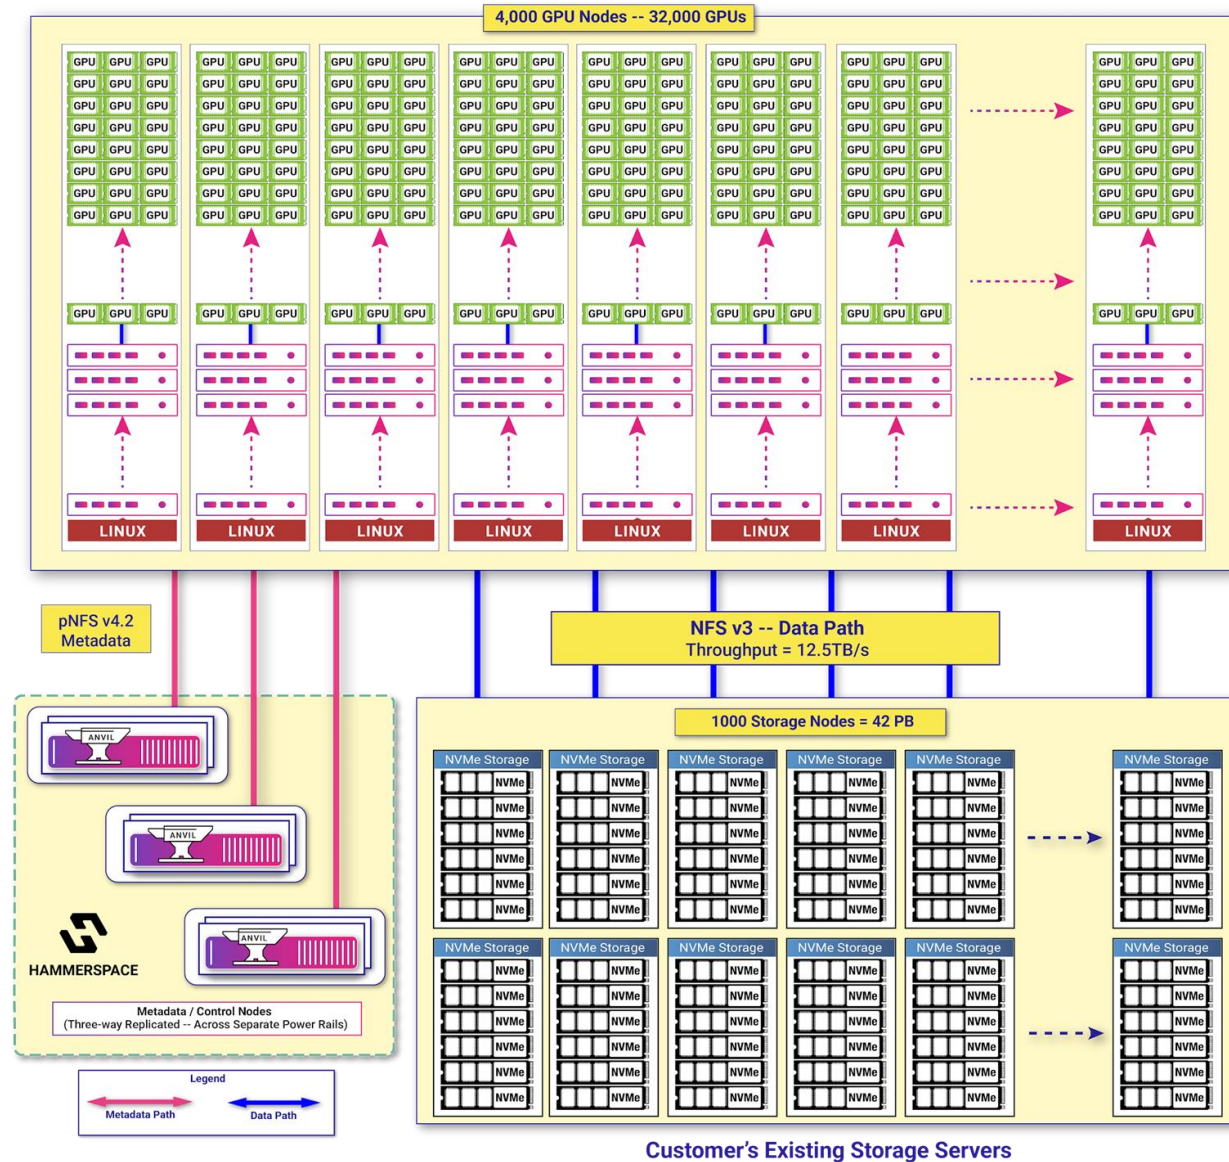
Hyperscale NAS at META for AI Training

About the Customer

- Largest web property in the world
- Training LLMs and other Gen AI models
- Massive performance and scale demands
- Evaluated every storage vendor

Hammerspace Solution

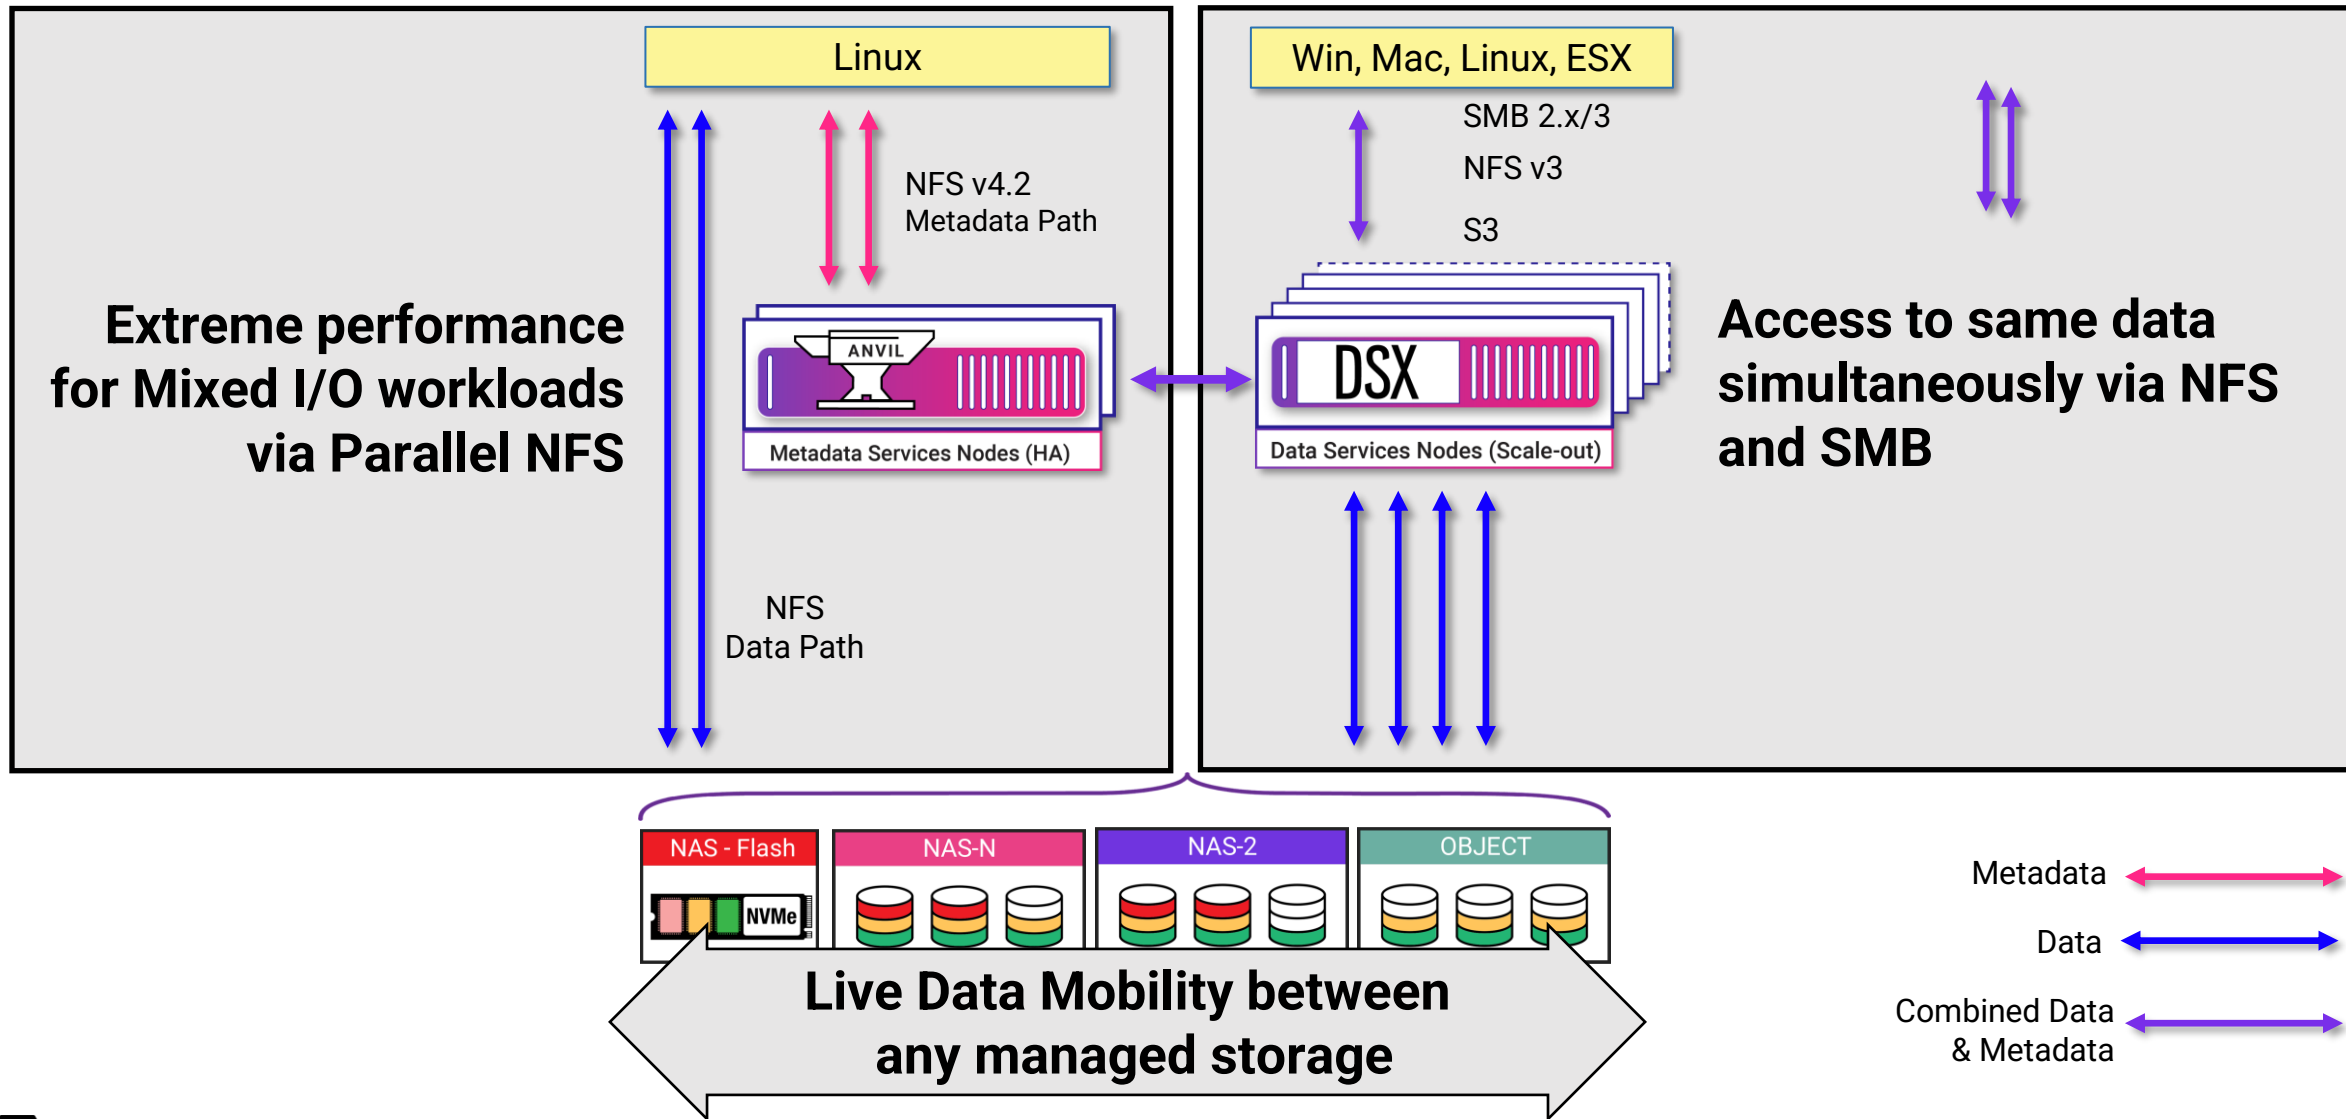
- No vendor even came close to Hammerspace's capabilities
- 1,000 node Hammerspace storage cluster
- Feeding 32,000 GPUs, soon to be 350,000, then 1M
- Aggregate performance of 12.5TB/sec (100Tb/sec)
- Everything is standards-based and plug-n-play
- Customer was able to use existing OCP storage servers



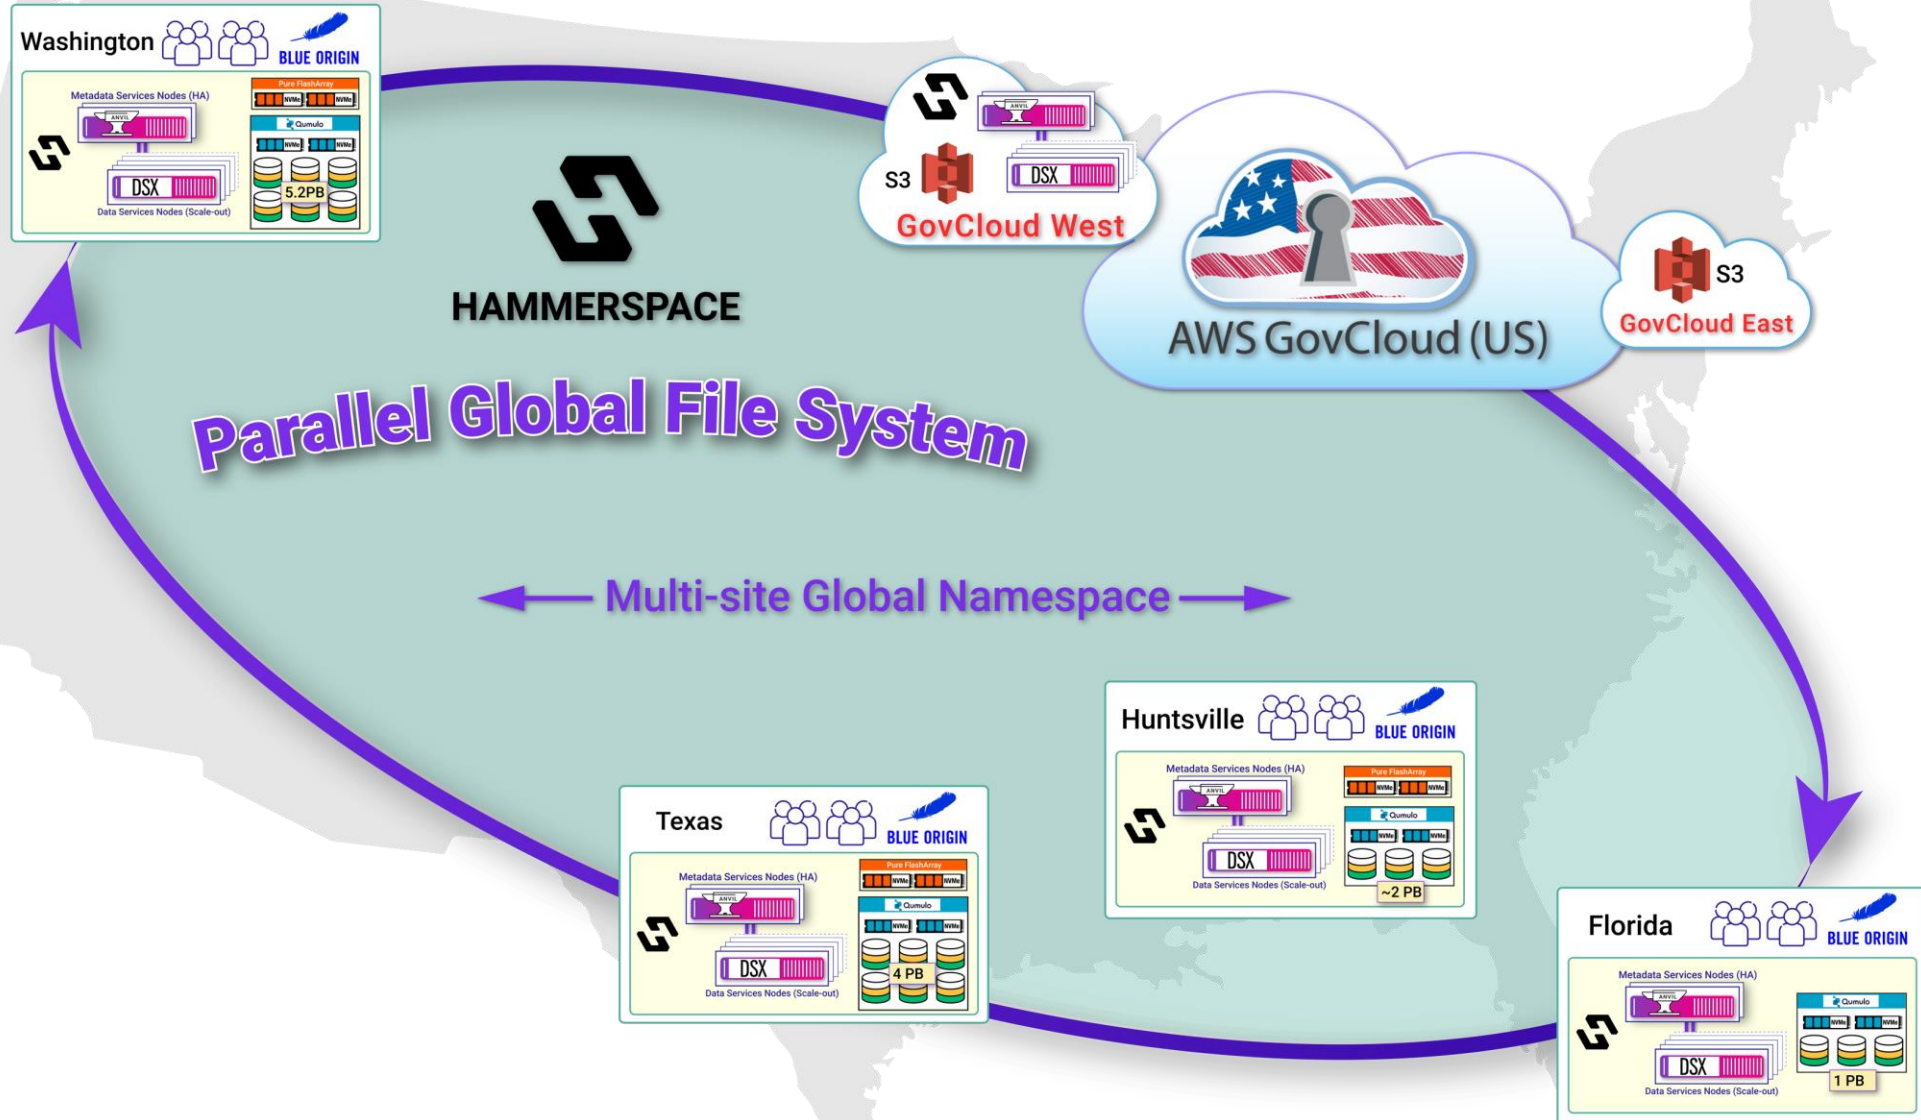
Customer's Existing Storage Servers



Hammerspace Flexible Multi-Protocol Access

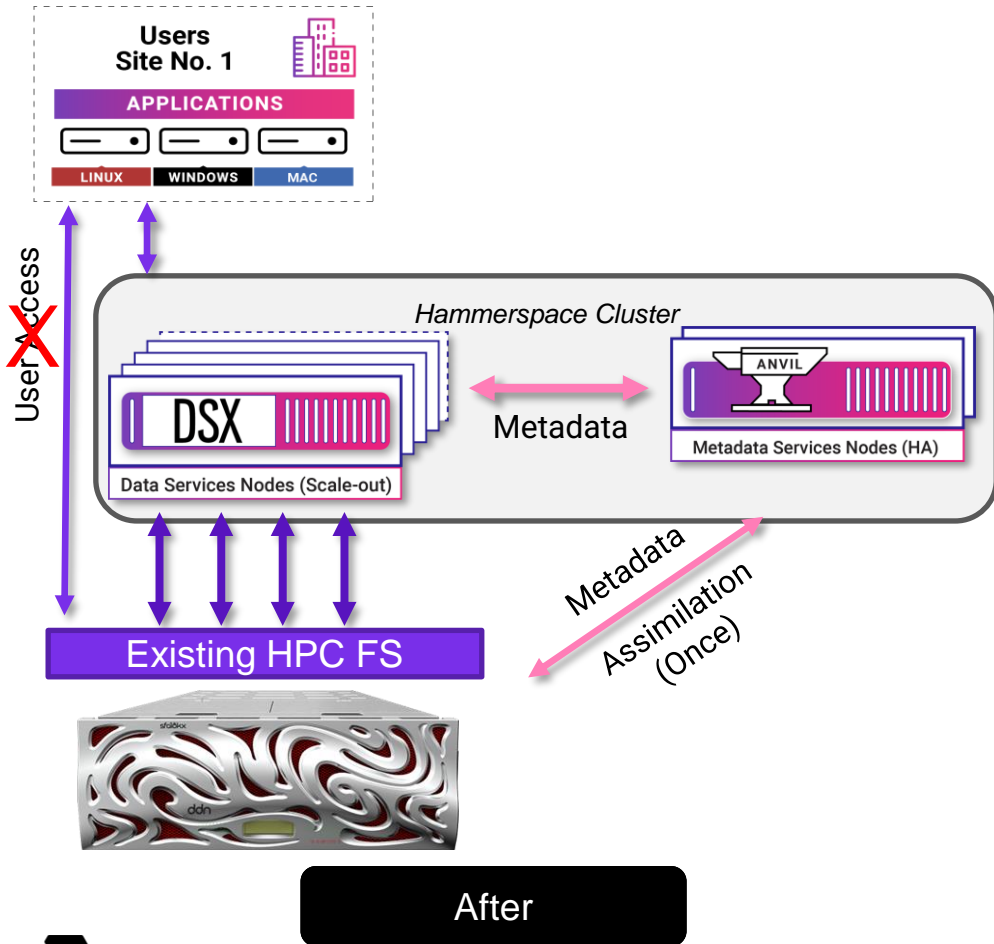


Data Powering Space Flight

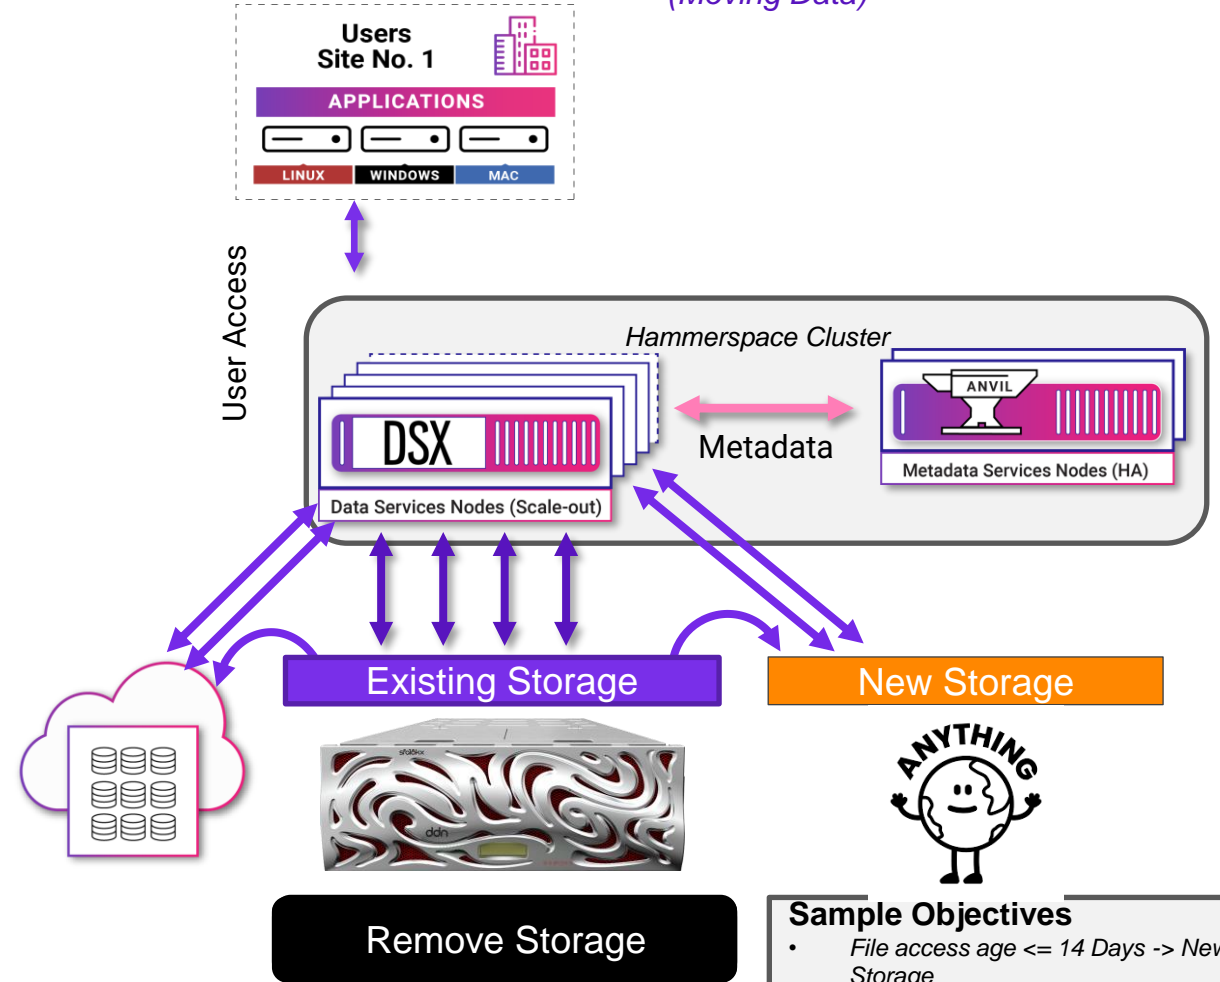


Assimilating (Metadata) / Migrating (Data) Process

Assimilating Existing Storage *(Moving Metadata)*



Adding/Migrating To New Storage *(Moving Data)*



Sample Objectives

- File access age ≤ 14 Days \rightarrow New Storage
- File access age ≥ 90 Days \rightarrow Cloud

Note: File-level granularity



Hyperscale NAS Architecture Benefits

Meet Extreme Performance and Scale Requirements

Serve thousands of users and tens of thousands of GPUs concurrently.

Scale to hundreds of storage nodes, Exabytes of capacity, and billions of files.

Reduce Required Hardware Footprint and Costs

Deploy on any off-the-shelf Linux Server. Efficient software architecture means fewer storage nodes and fewer network ports required.

Simple Standards-Based Connectivity

NFSv4.2 Client provides parallel performance built into the Linux OS. Access same data simultaneously via NFSv3 and SMB.

Simplify Data Governance with Enterprise Data Services

Easily protect data with snapshots and clones, recover data with versioning and undelete, manage permissions, and more.



Q&A

Primary Contacts:

Robert Renzoni

Dir Technical Sales Federal

Robert.Renzoni@Hammerspace.com

(571)213-5699

Additional Resources:

<https://hammerspace.com/>



Thank you



HAMMERSPACE