

# Hammerspace

## *Global Data Platform*

David Flynn – Chief Executive Officer

April 3, 2024



**HAMMERSPACE**

*the Future of Memory and Storage*

# Hammerspace Is Like Magic




“Hammerspace” is an extra-dimensional space that is instantly accessible and infinite in size

# Goal to Radically Improve How Data Is Used

AI is forcing a long-overdue industry reckoning to implement radical changes in how data is used and preserved

FROM: Data Placed to Rest in Storage

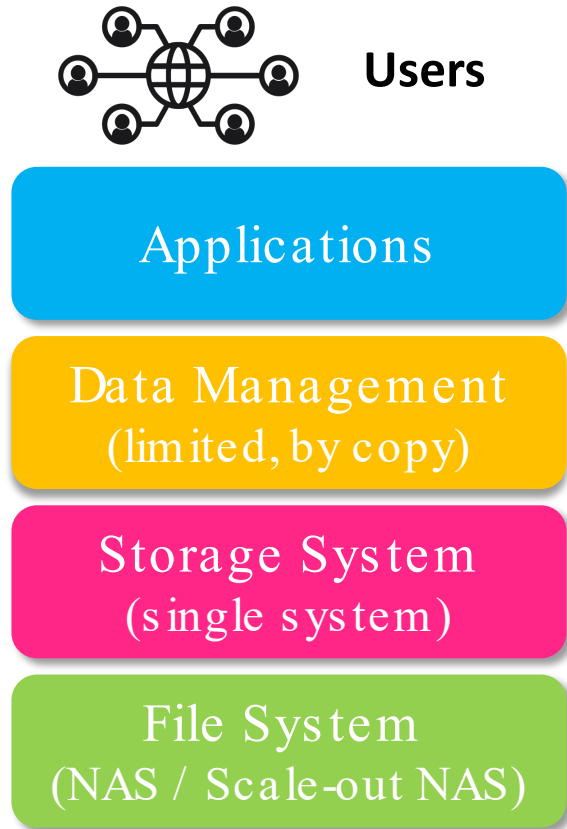


TO: Data in Motion Across a Global Data Environment



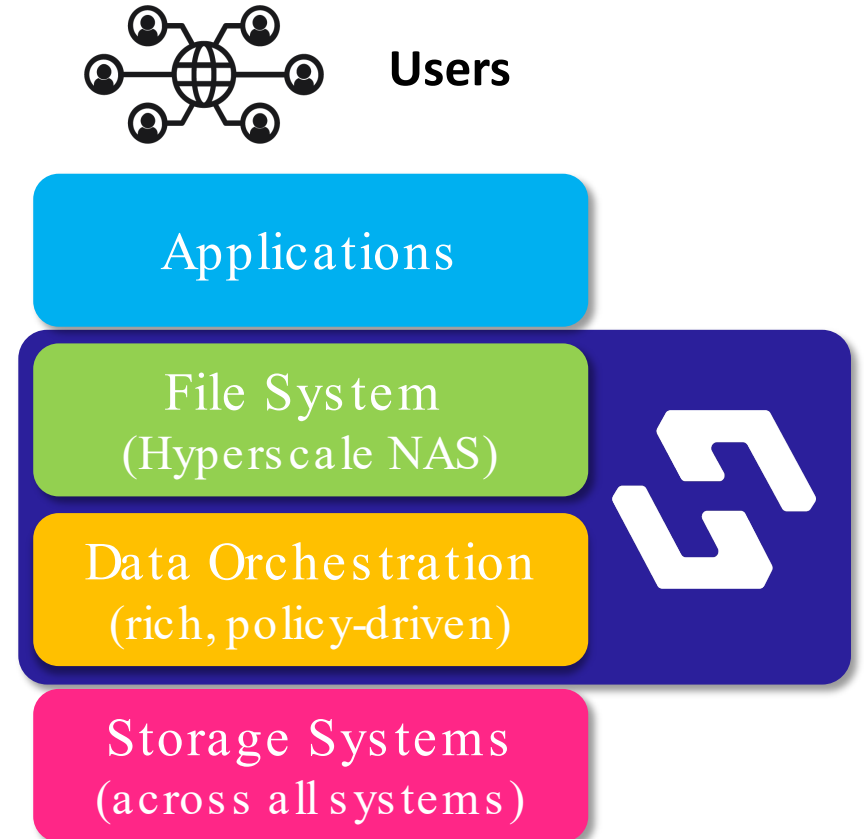
# A Fundamentally Different, Data-Centric Architecture

## Storage-Centric Approach



**✗** Data is Trapped in Silos

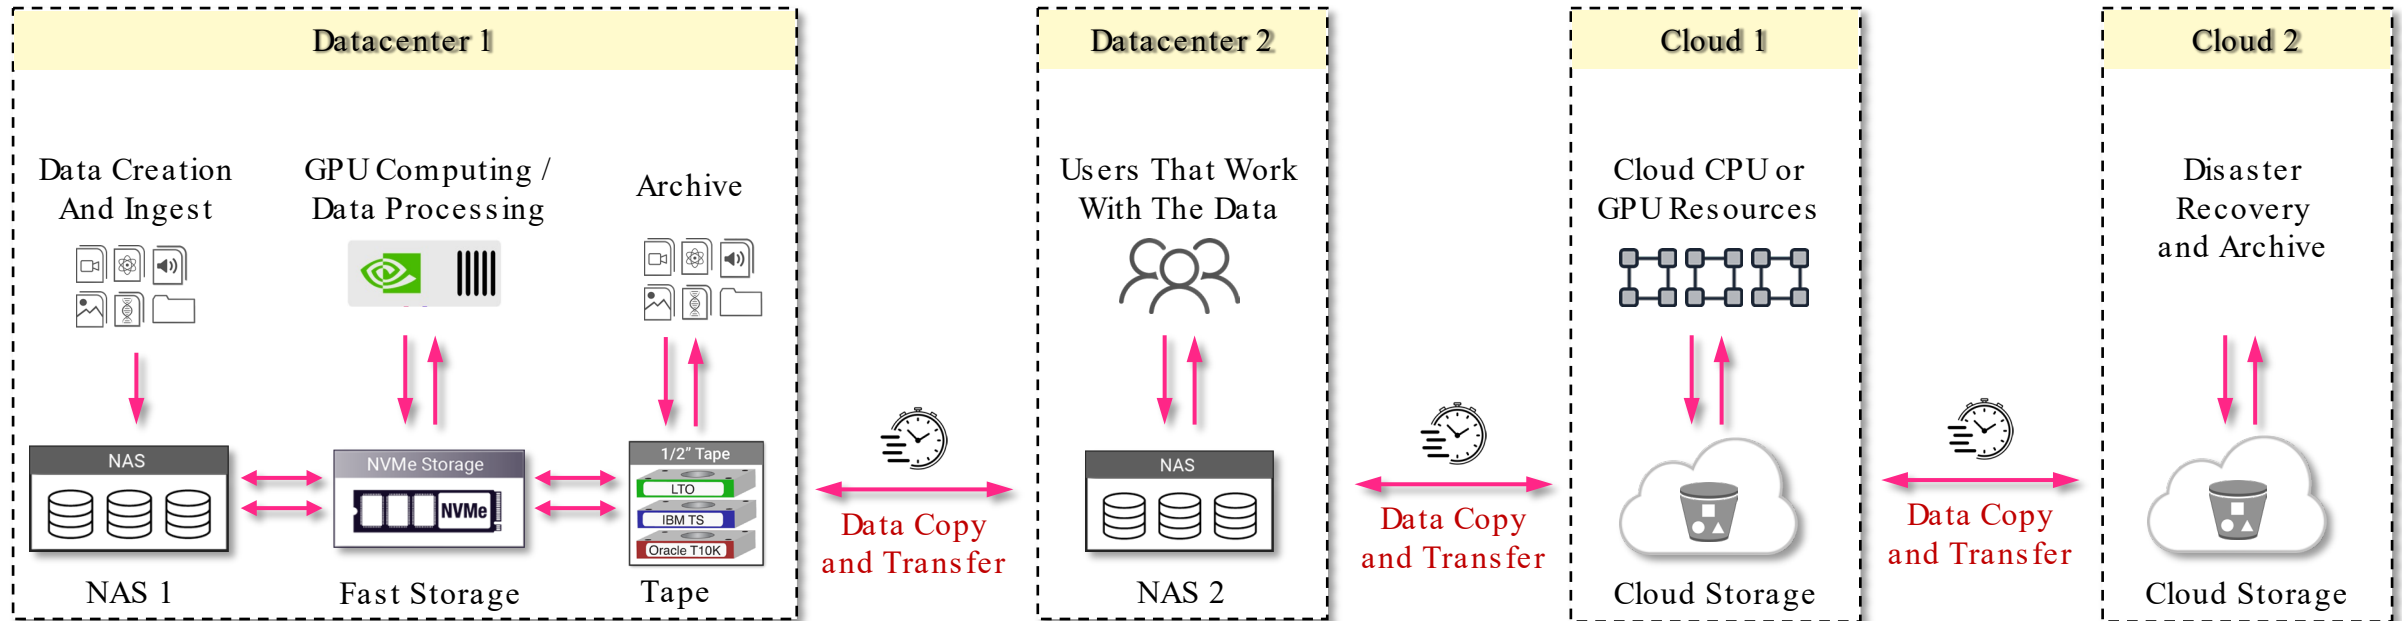
## Hammerspace Data-Centric Approach



**✓** Data Becomes a Global Resource

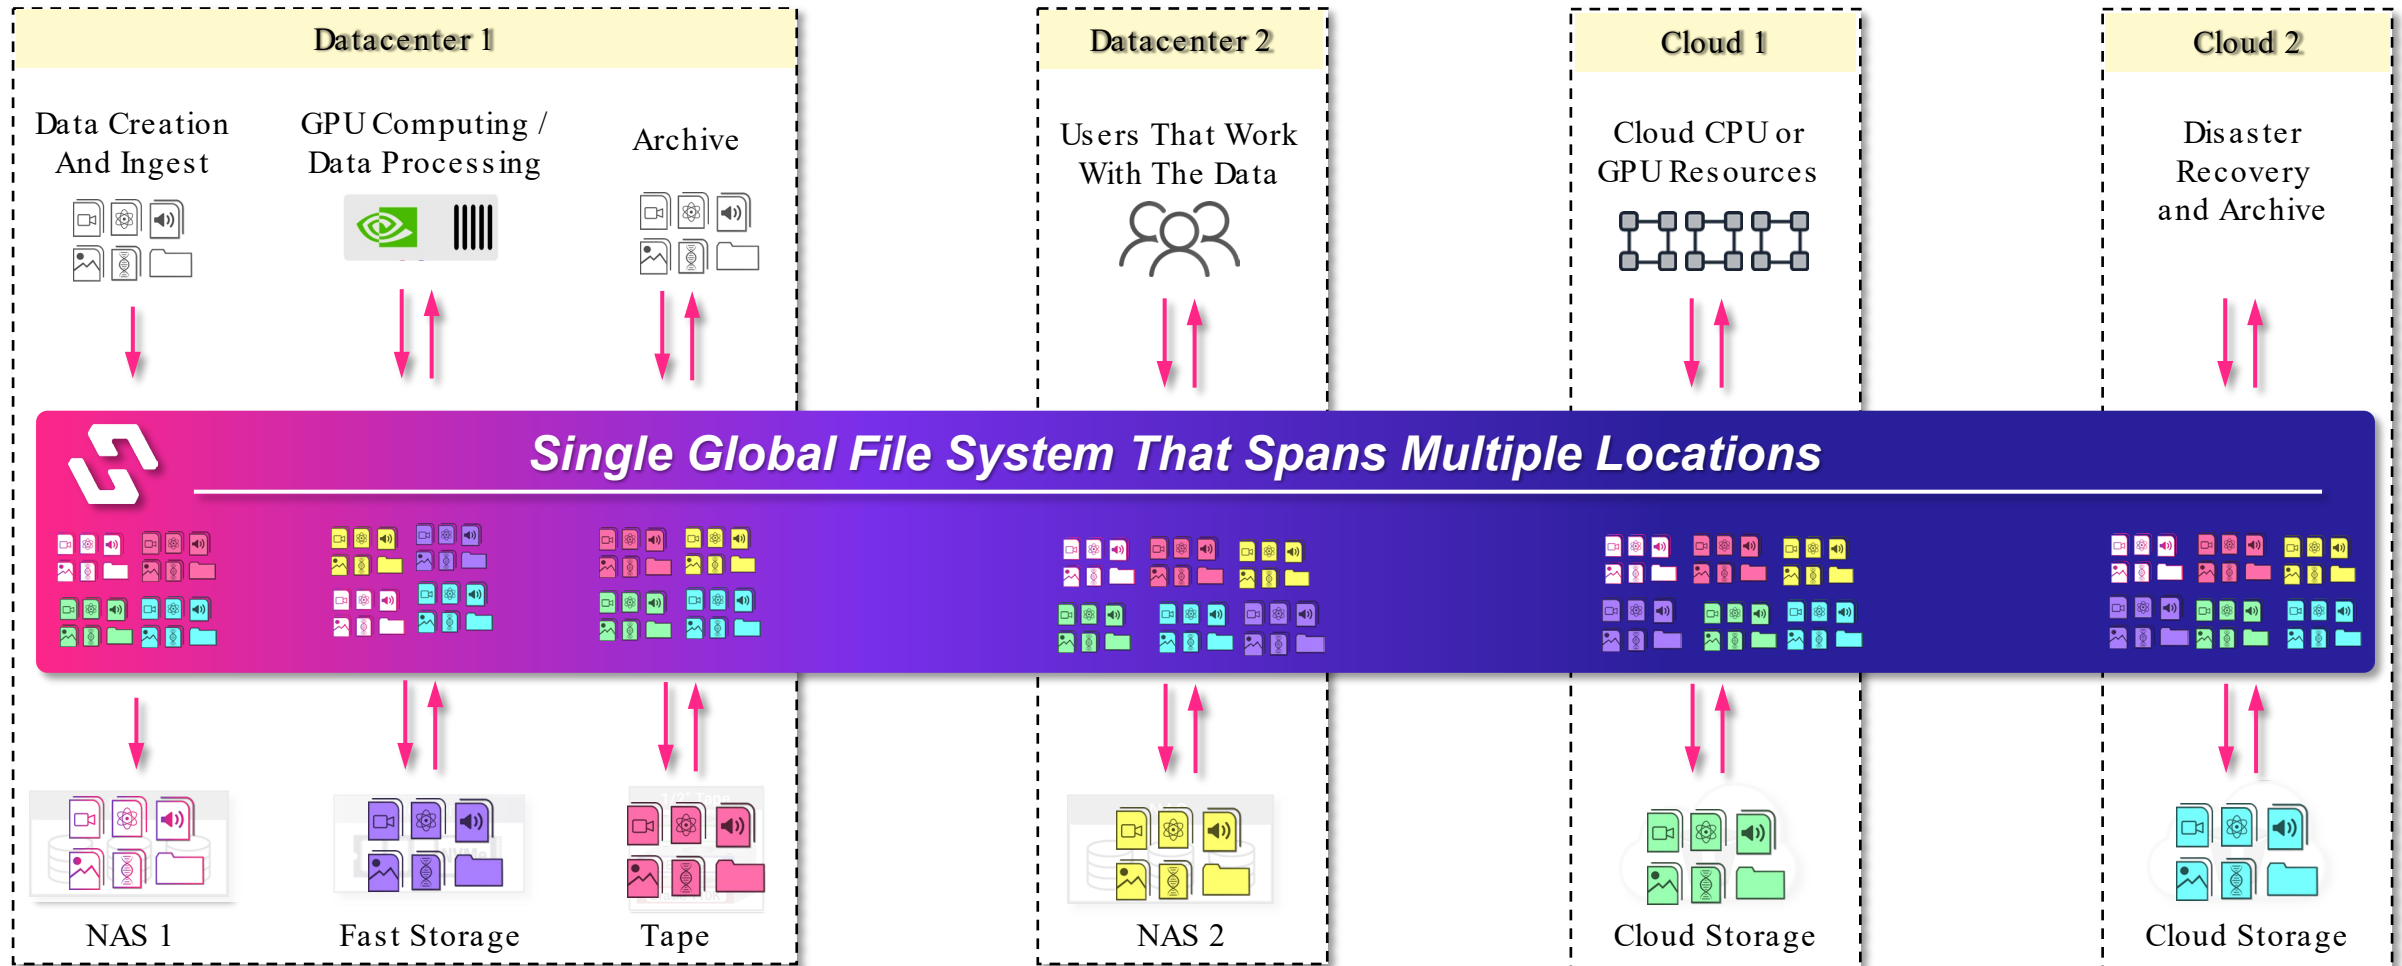
# A New Solution Was Inevitable

- ❌ Valuable data trapped in silos
- ❌ Data copy sprawl impacts governance and security
- ❌ Getting data to global users takes too long
- ❌ Lack performance to keep GPUs utilized
- ❌ Infrastructure is not ready for AI
- ❌ Lack agility to use elastic cloud resources



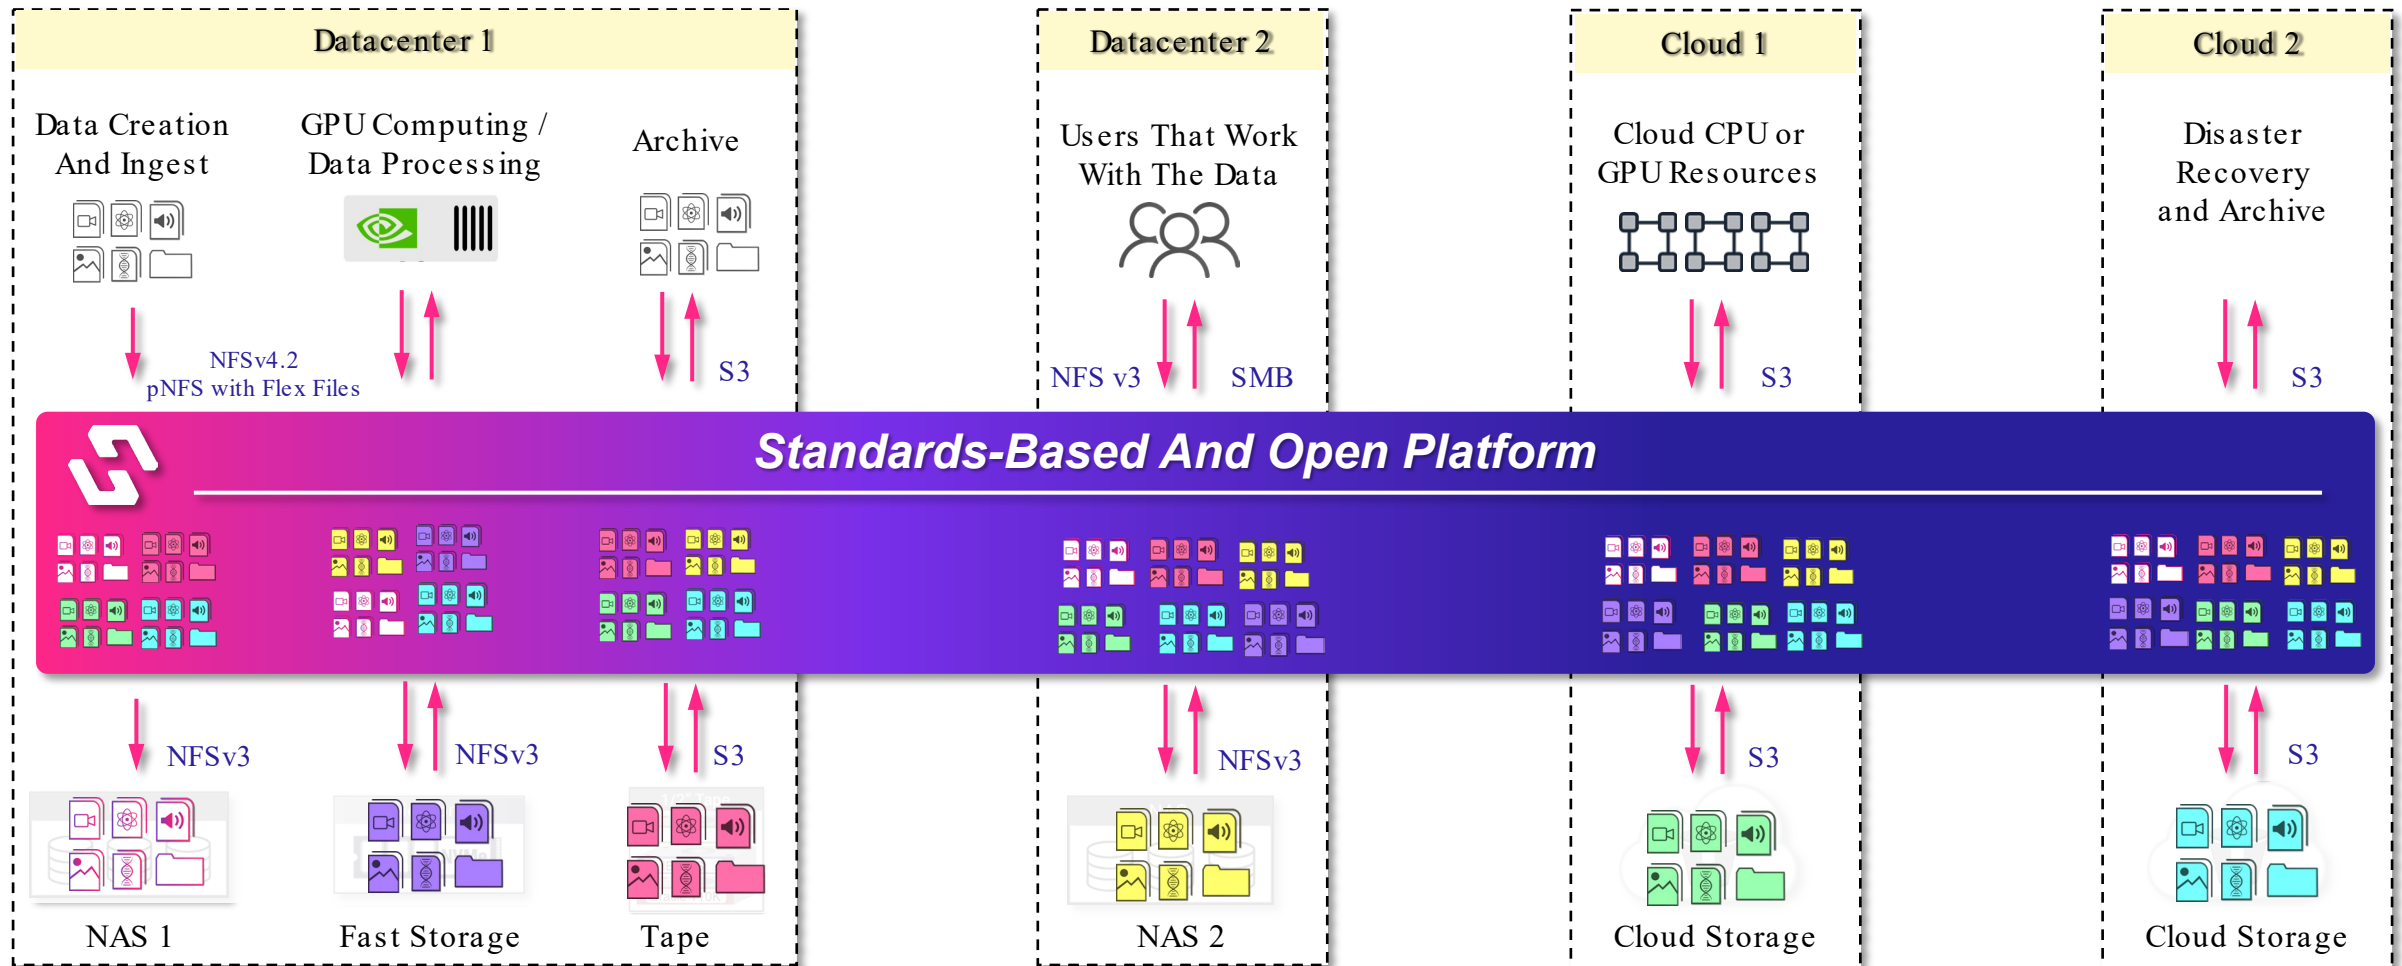
# Hammerspace Changes the Paradigm

*Hammerspace virtualizes the underlying storage infrastructure, so all authorized users and applications can access the same data simultaneously.*



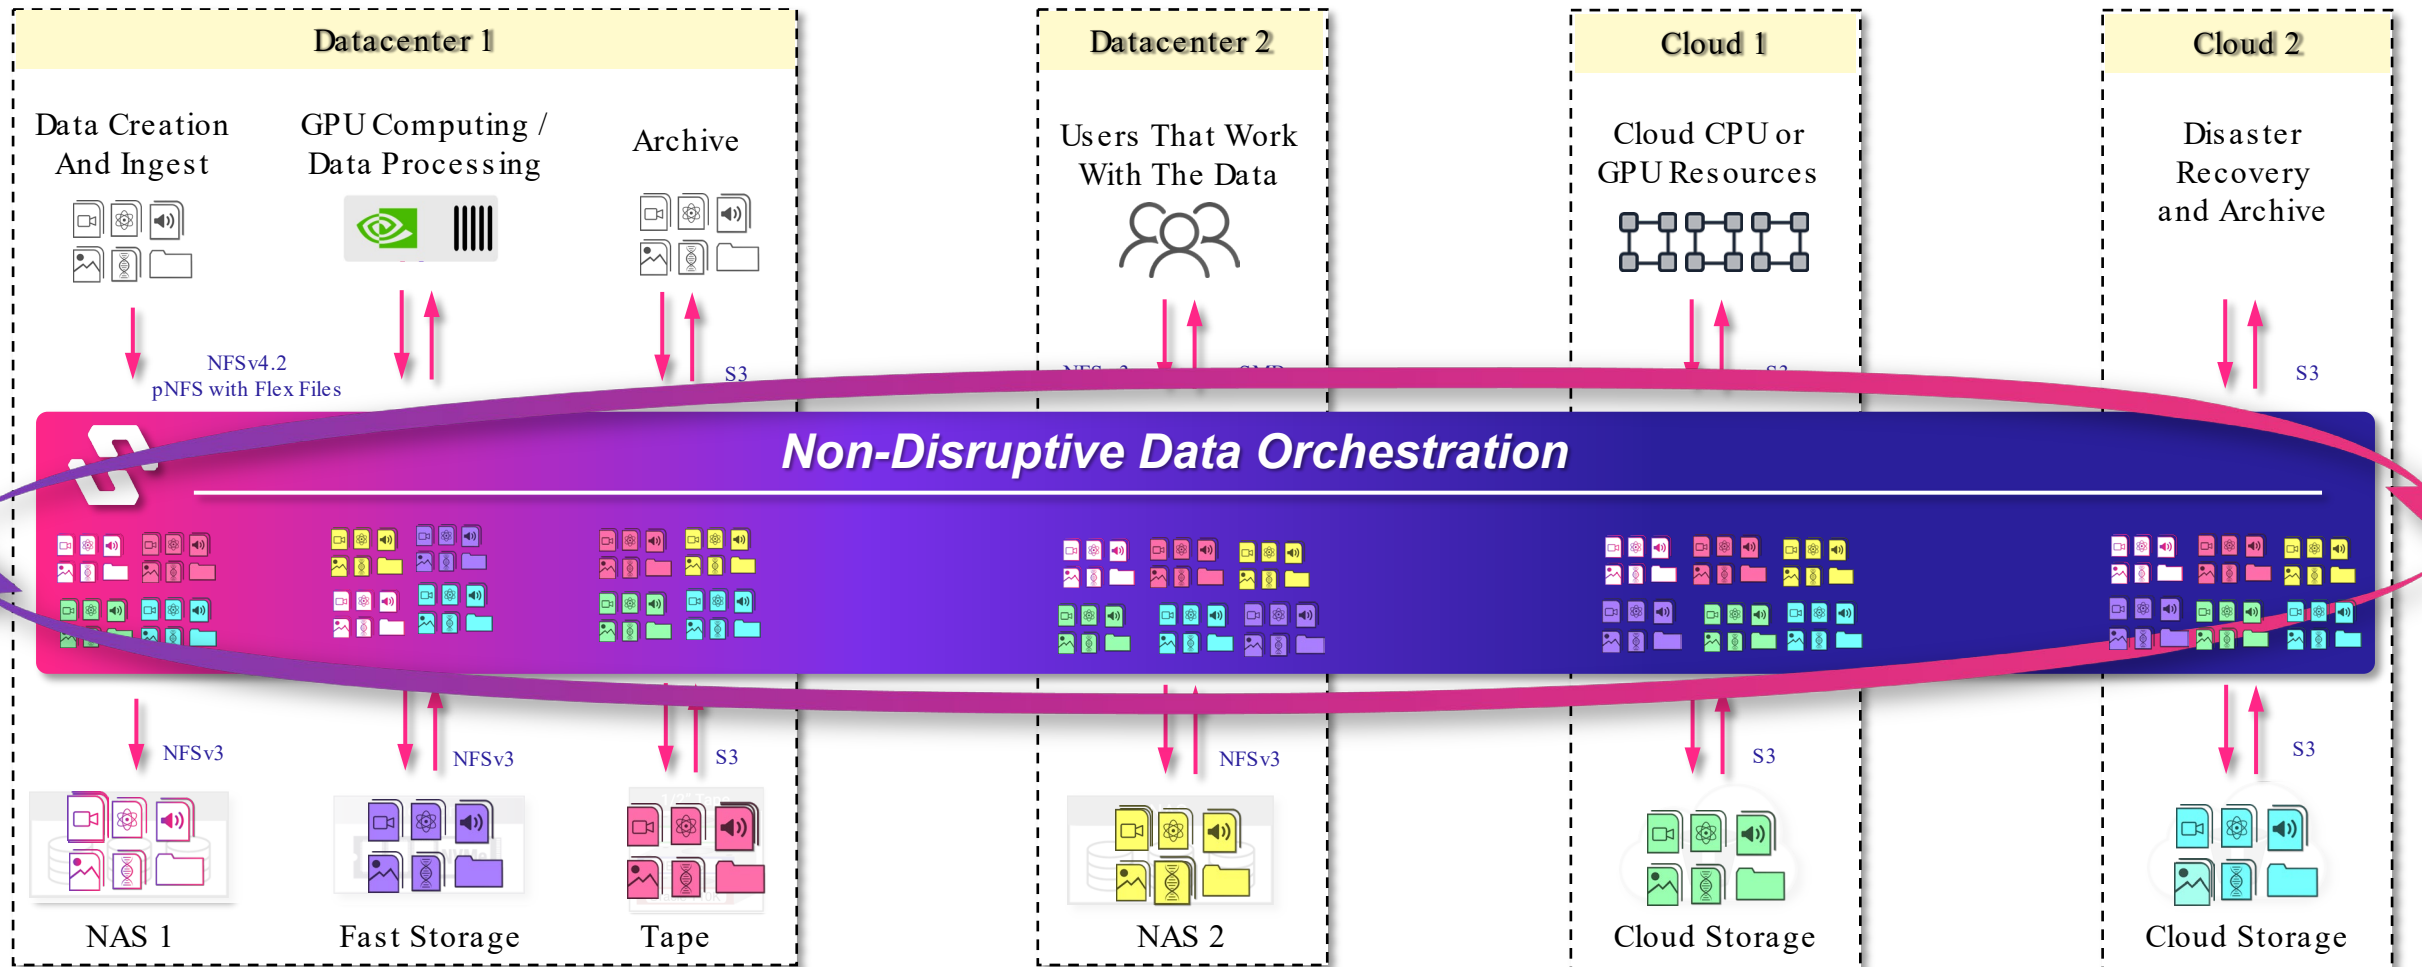
# Hammerspace Changes the Paradigm

*Hammerspace provides fast data access using industry-standard protocols, including a direct data path between Linux clients and storage for high-speed data processing.*



# Hammerspace Changes the Paradigm

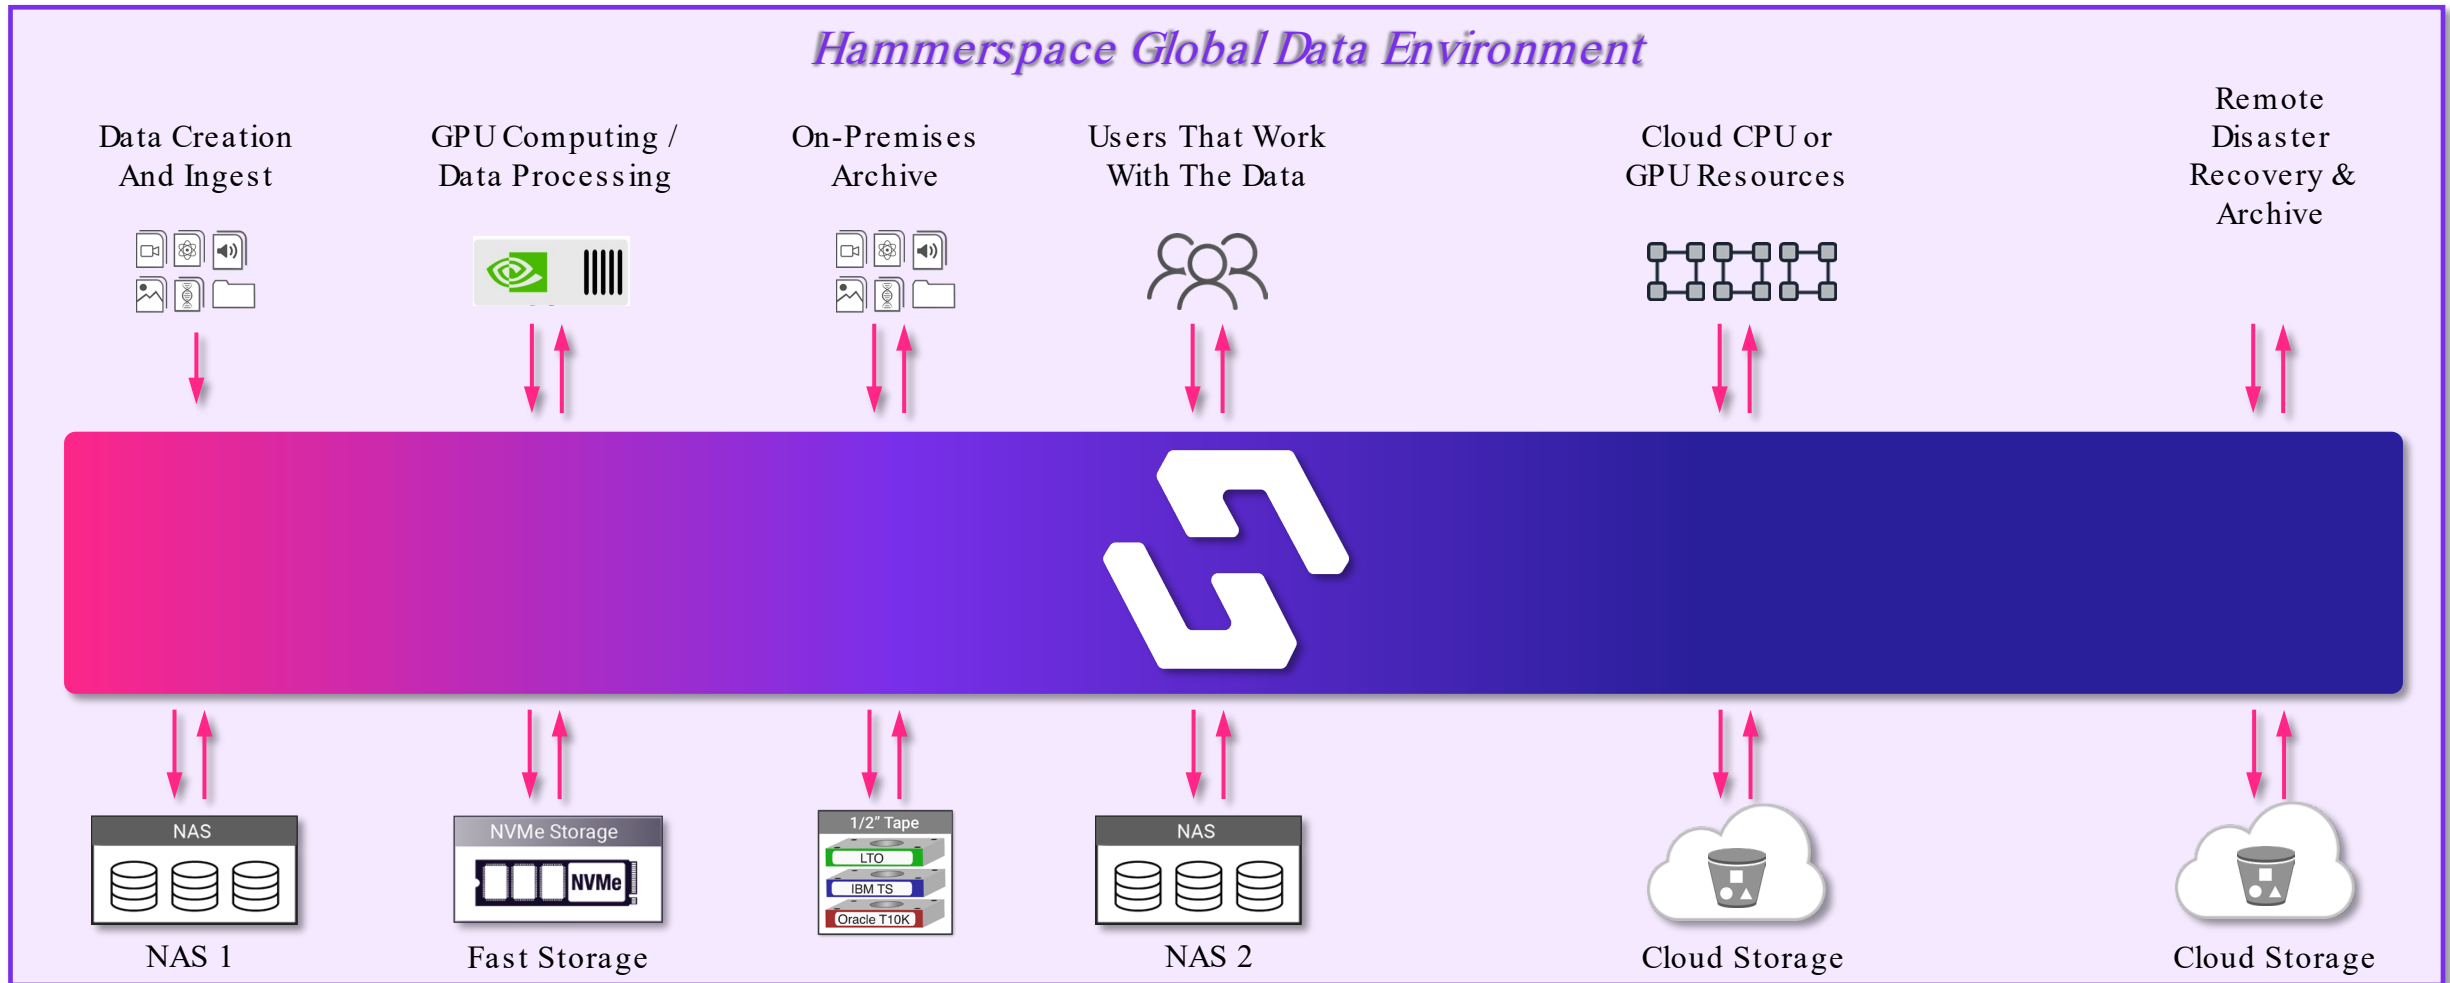
Hammerspace automates data orchestration behind the scenes, transparent to users even on live data.





# Hammerspace Global Data Environment


*Unify and automate unstructured data across any data center, any cloud, anywhere*



# Standards-Based, Efficient Architectures


# Hyperscale NAS for Data Intensive Workloads

**Hyperscale NAS** is a fundamentally different NAS architecture urgently needed to power AI initiatives and the GPU computing boom.



### Scale-Up NAS

- Introduced 1990s
- Single storage controller
- For Corporate data



### Scale-Out NAS

- Introduced early 2000s
- Multiple storage controllers
- For Unstructured Data
- Good for scaling capacity



### HYPERSCALE NAS

- Parallel file system + Enterprise NAS
- Standards-based
  - pNFS client in Linux
  - Networking
  - Storage Interfaces
- Linear scale to thousands of nodes
- Software-defined
- Storage-agnostic



SCALING NAS CAPACITY

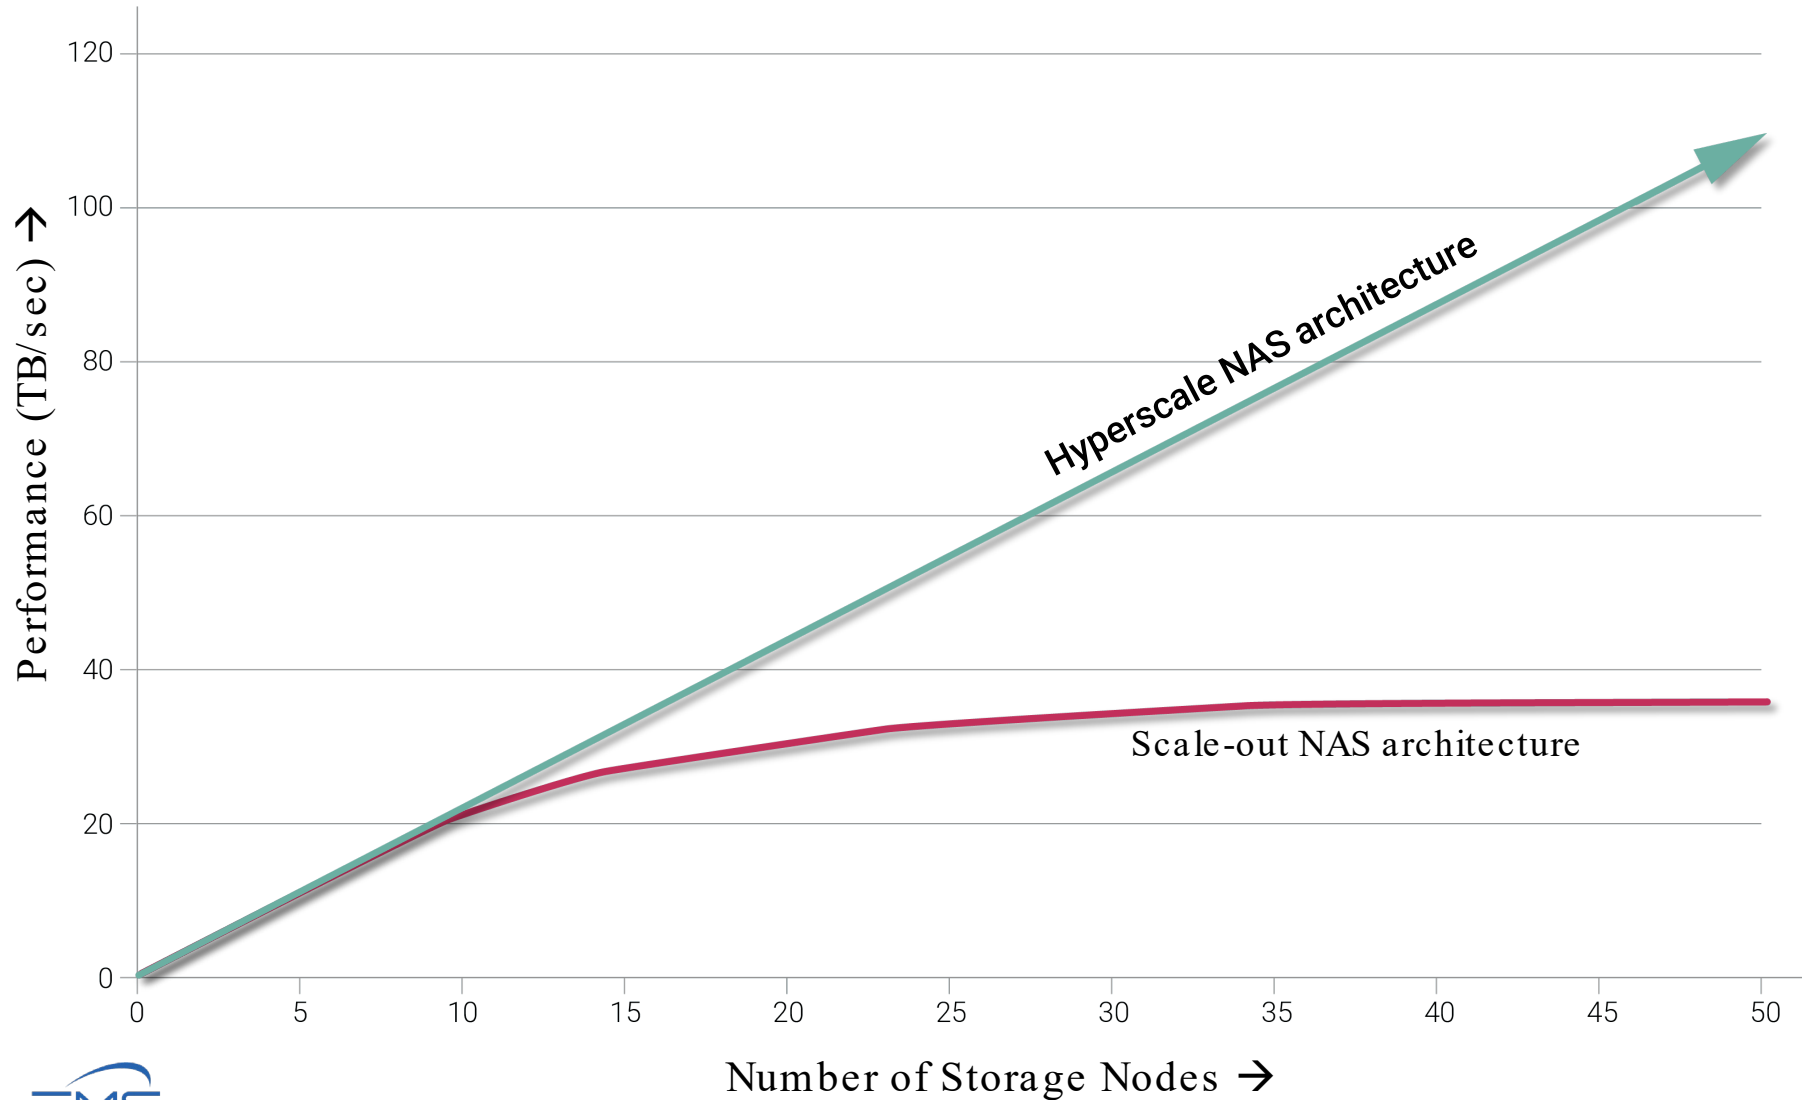
SCALING NAS PERFORMANCE

MOSTLY HUMAN-GENERATED DATA

MOSTLY MACHINE-GENERATED DATA



# Hyperscale NAS Delivers Linear Performance Scalability



Hyperscale NAS increases performance linearly as node count increases. So far proven at over 1,000+ storage nodes.

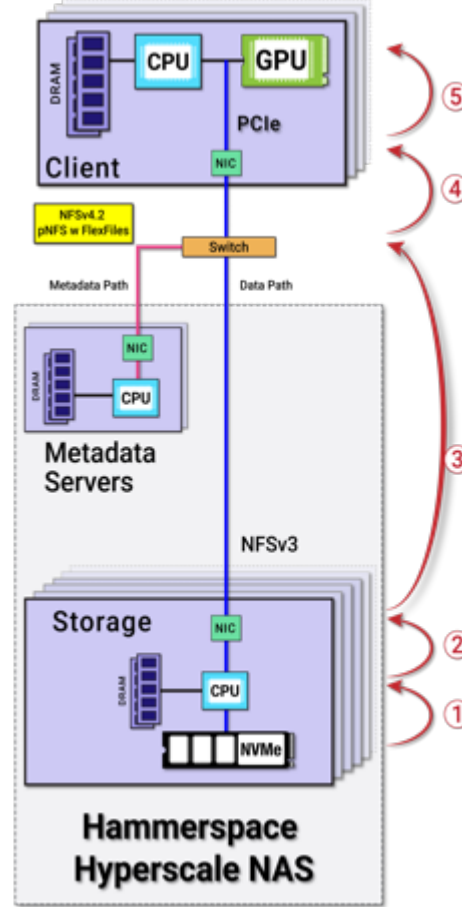

Scale-Out NAS performance plateaus as storage nodes increase



# 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

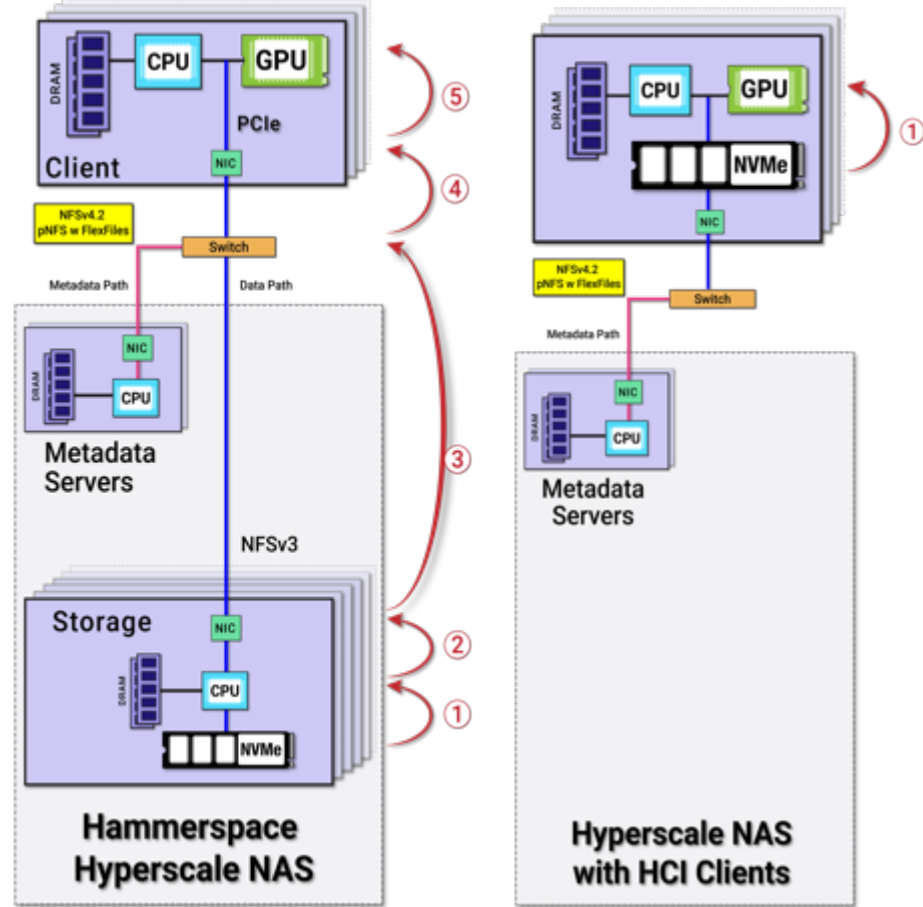
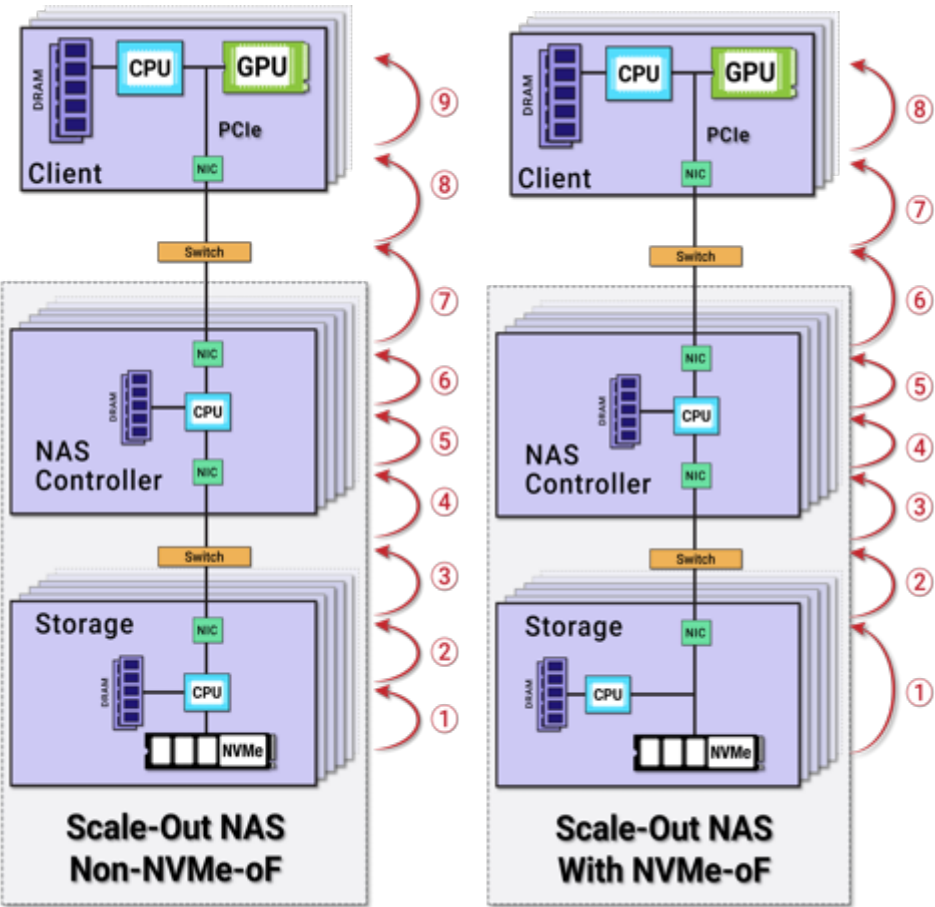
**NVIDIA GPUDirect**



# 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

**NVIDIA GPUDirect**

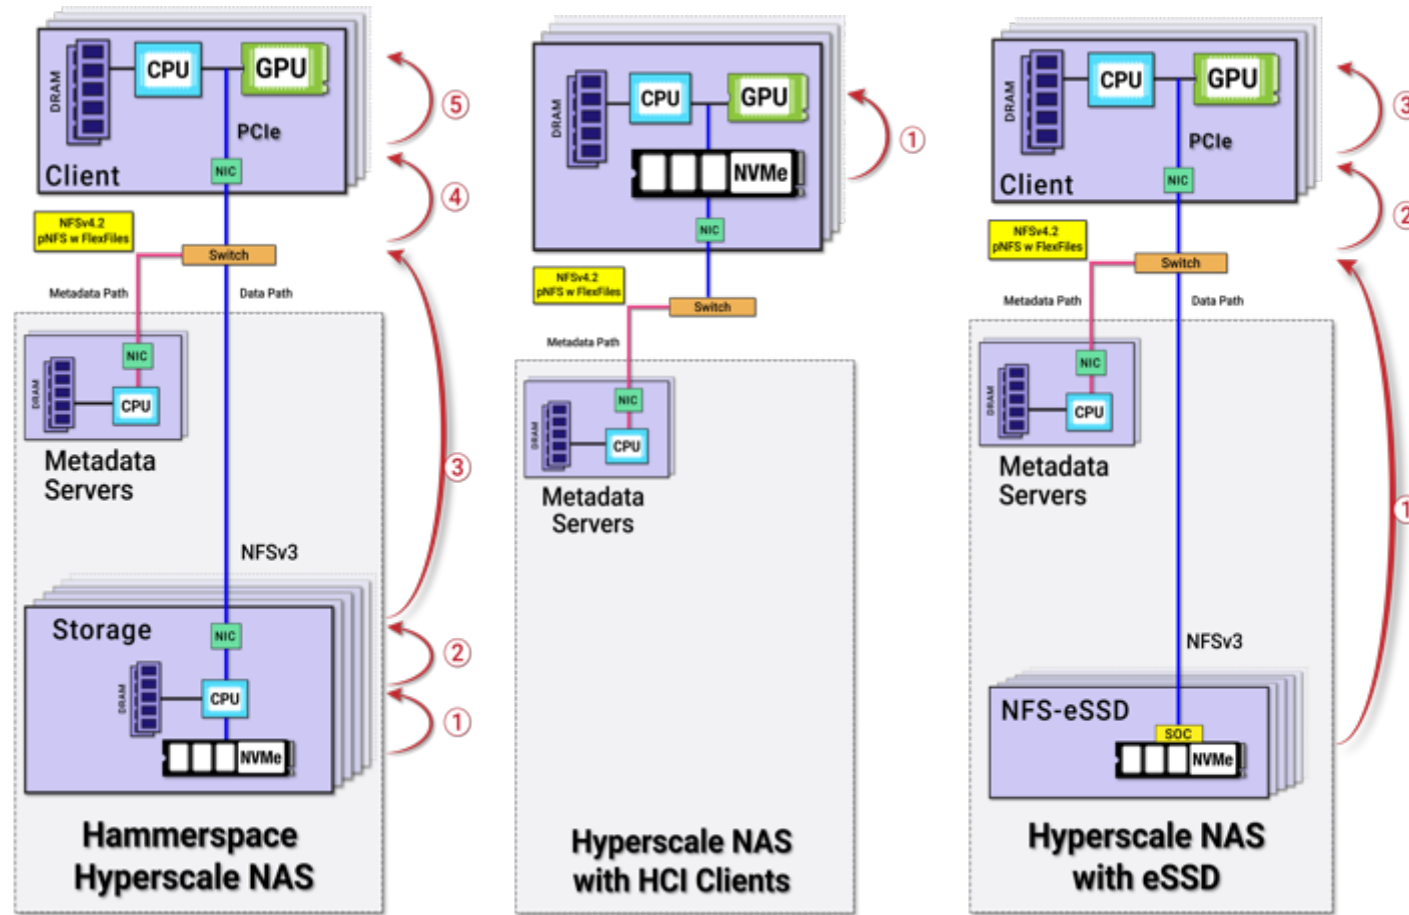


# File Storage Architecture Comparison

## 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

**NVIDIA GPUDirect**


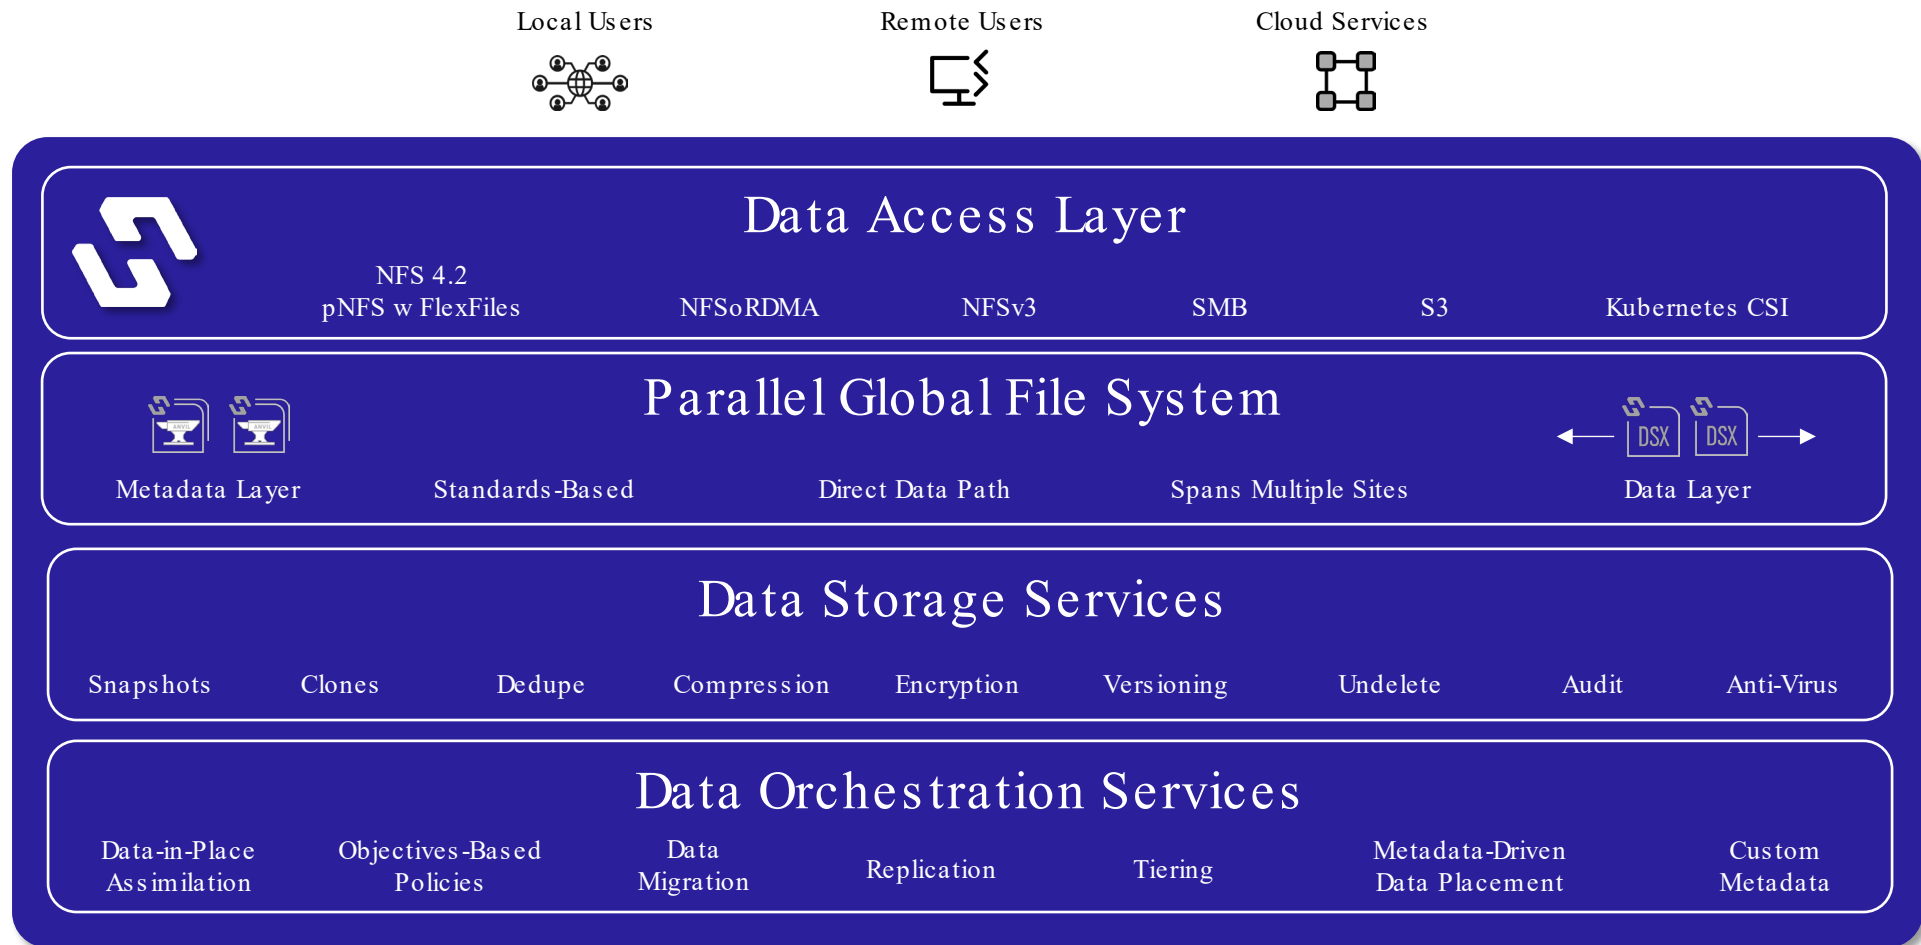




# Architecture



# Software Architecture

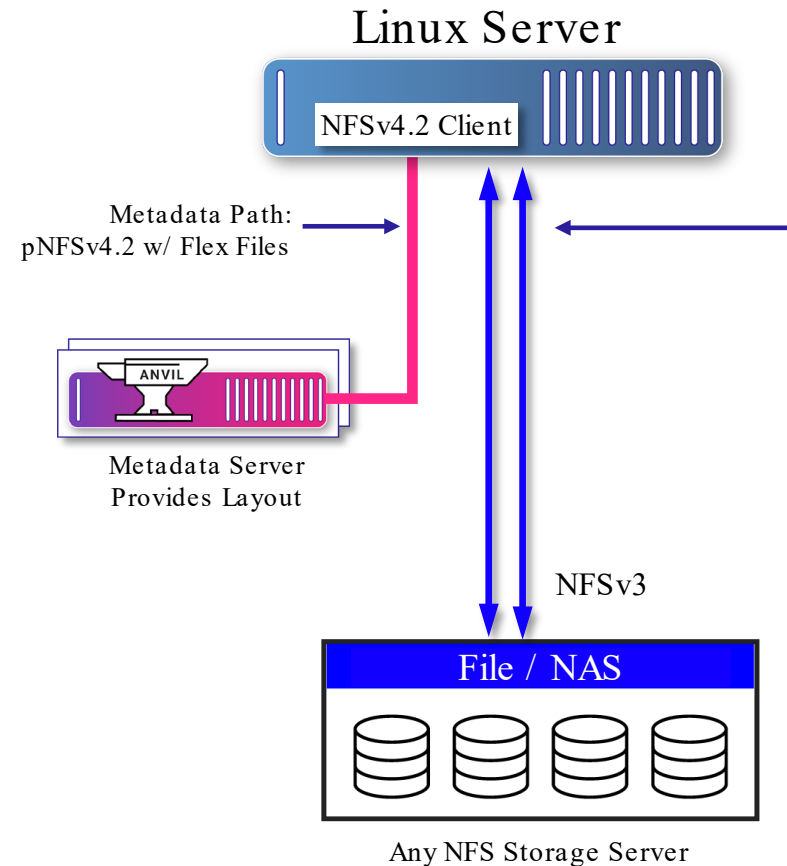


Runs on New or Existing Storage from All Leading Vendors

# First Standards-Based Parallel File System Using NFS

## About Parallel NFS with Flex Files

- Parallel NFS (pNFS) introduced as optional feature of NFS in 2010, enhanced in later RFCs
- Flex Files layout type engineered by Hammerspace and added to NFSv4.2 spec in 2018
- Defines a standards-based parallel file system architecture using NFS
- Architecture requires NFSv4.2 client which is part of the Linux kernel
- Hammerspace Global Parallel File System is the first to leverage this architecture

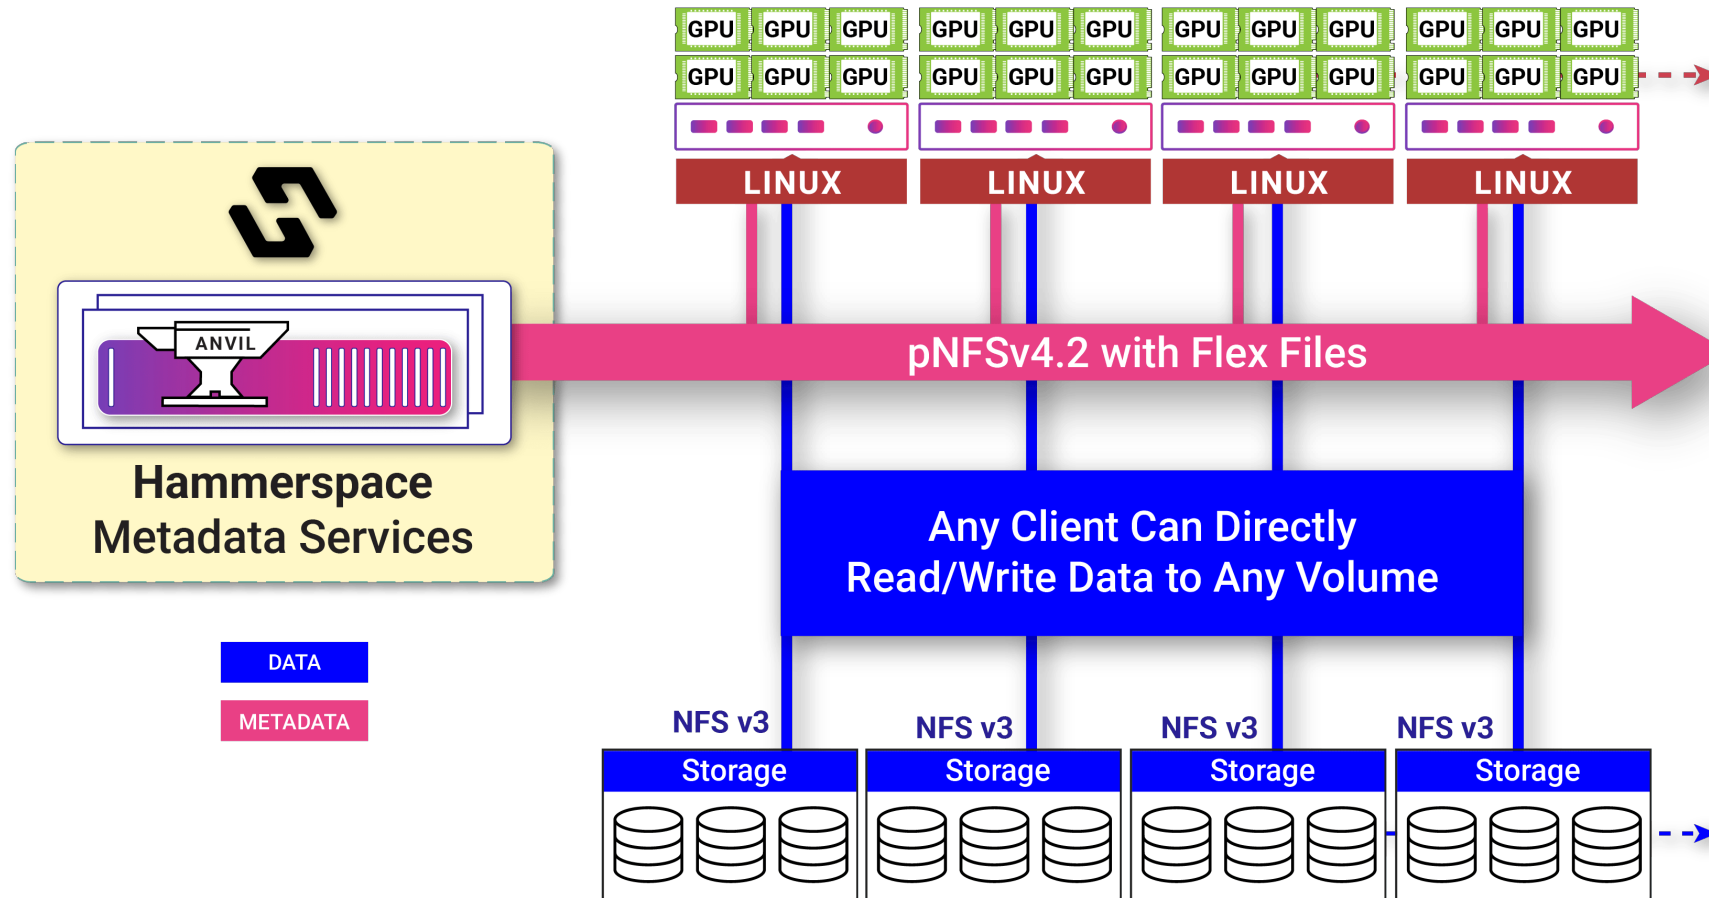


Direct data path between Linux client and storage volumes

Also provides for:

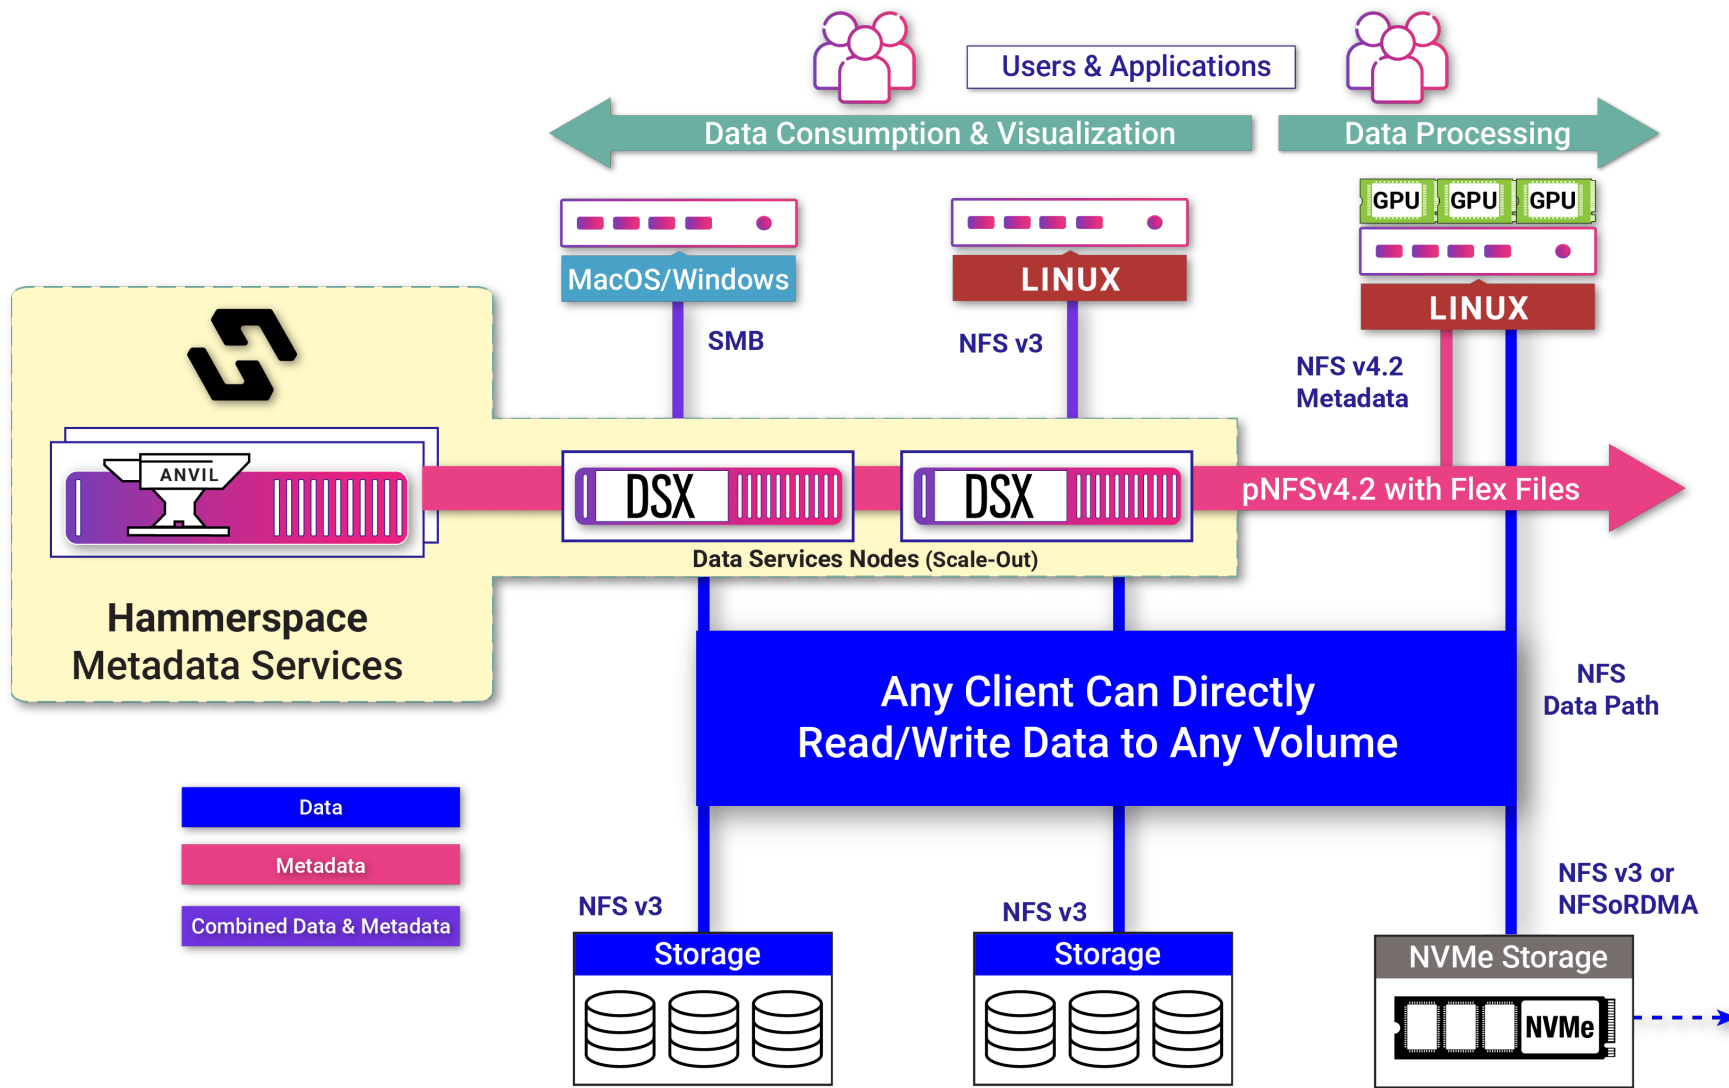
- Multiple parallel network connections between client and server
- Ability to write to multiple storage nodes synchronously
- Ability to move data while it is being accessed without interruption
- Eliminates protocol chatter
- File granular access and telemetry
- TCP or RDMA Supported

# About the Hyperscale NAS Architecture



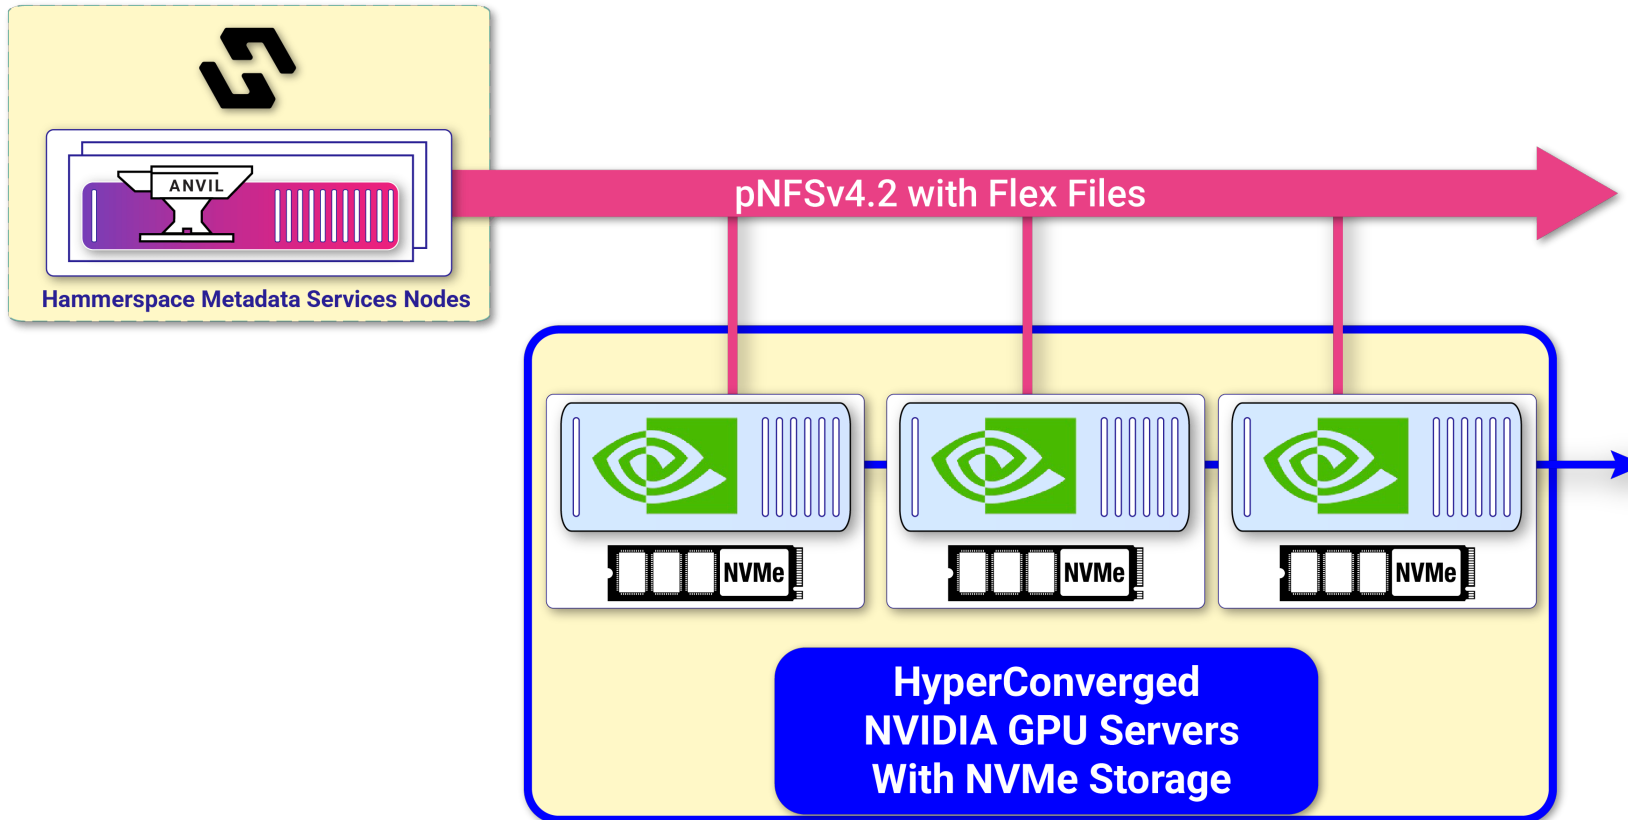
- Extreme Parallel Performance
- Linear scale to >1000s of nodes
- Standards-Based
- Choice of networking
- Software-defined
- Storage-agnostic

# Architecture Extends to Address Multiple Use Cases



- Read/write data using standard protocols
- Scale access and performance by adding more DSX Nodes.
- Use any vendor storage including NAS, block, object, cloud, and even tape.

# Tightly Coupled in a Hyperconverged Architecture



- Supports Hyperconverged GPU servers with NVMe storage for greater performance & network efficiency



# Customer Stories


# Customer Story: Aerospace Company

## About the Customer, Business Goals, Requirements

- Aerospace company building and testing reusable rockets and engines
- Remote sites producing large amounts of data
- NAS Storage at each site filling up
- Engineers in various offices and work from home
- 11,000 users across 9 locations
- Moving data to be “close to the engineer” is a challenge

## Solution and Why Hammerspace

- Hammerspace Global Data Environment spans locations
- Engine test data now proactively made available at multiple sites
- Local data stays on fast storage, then moved to AWS S3
- Data always available
- Transition from legacy storage to Hammerspace was transparent to users



# Customer Story: VFX Studio

## About the Customer, Business Goals, Requirements

- Leading global Visual FX and Animation Studio
- 300 artists in two UK locations, remote work around the world
- Virtual since 2019, storage in multiple locations and clouds
- Sharing large files globally to keep artists productive was a “real issue”

## Solution and Why Hammerspace

- Hammerspace orchestrates data movement globally
- Both on-premises and Azure cloud environments
- Integrated with Jellyfish’s production scheduler: Autodesk Shotgrid
- 100M+ files globally accessible for read/write access in minutes
- Saved 30% in initial render costs by moving files to lower cost clouds
- Enabled access to global talent pool of artists



JELLYFISH PICTURES

“Hammerspace is a core part of the Jellyfish strategic vision, helping us expand our global workforce in a highly competitive industry, increasing our productivity to meet and exceed our clients’ expectations while greatly reduce costs on multiple levels.”

Jeremy Smith, CTO, Jellyfish Pictures



# Customer Story: LLM Training at Hyperscale

“What Hammerspace does is pure magic.”

-Paul Saab, Principal Engineer Meta

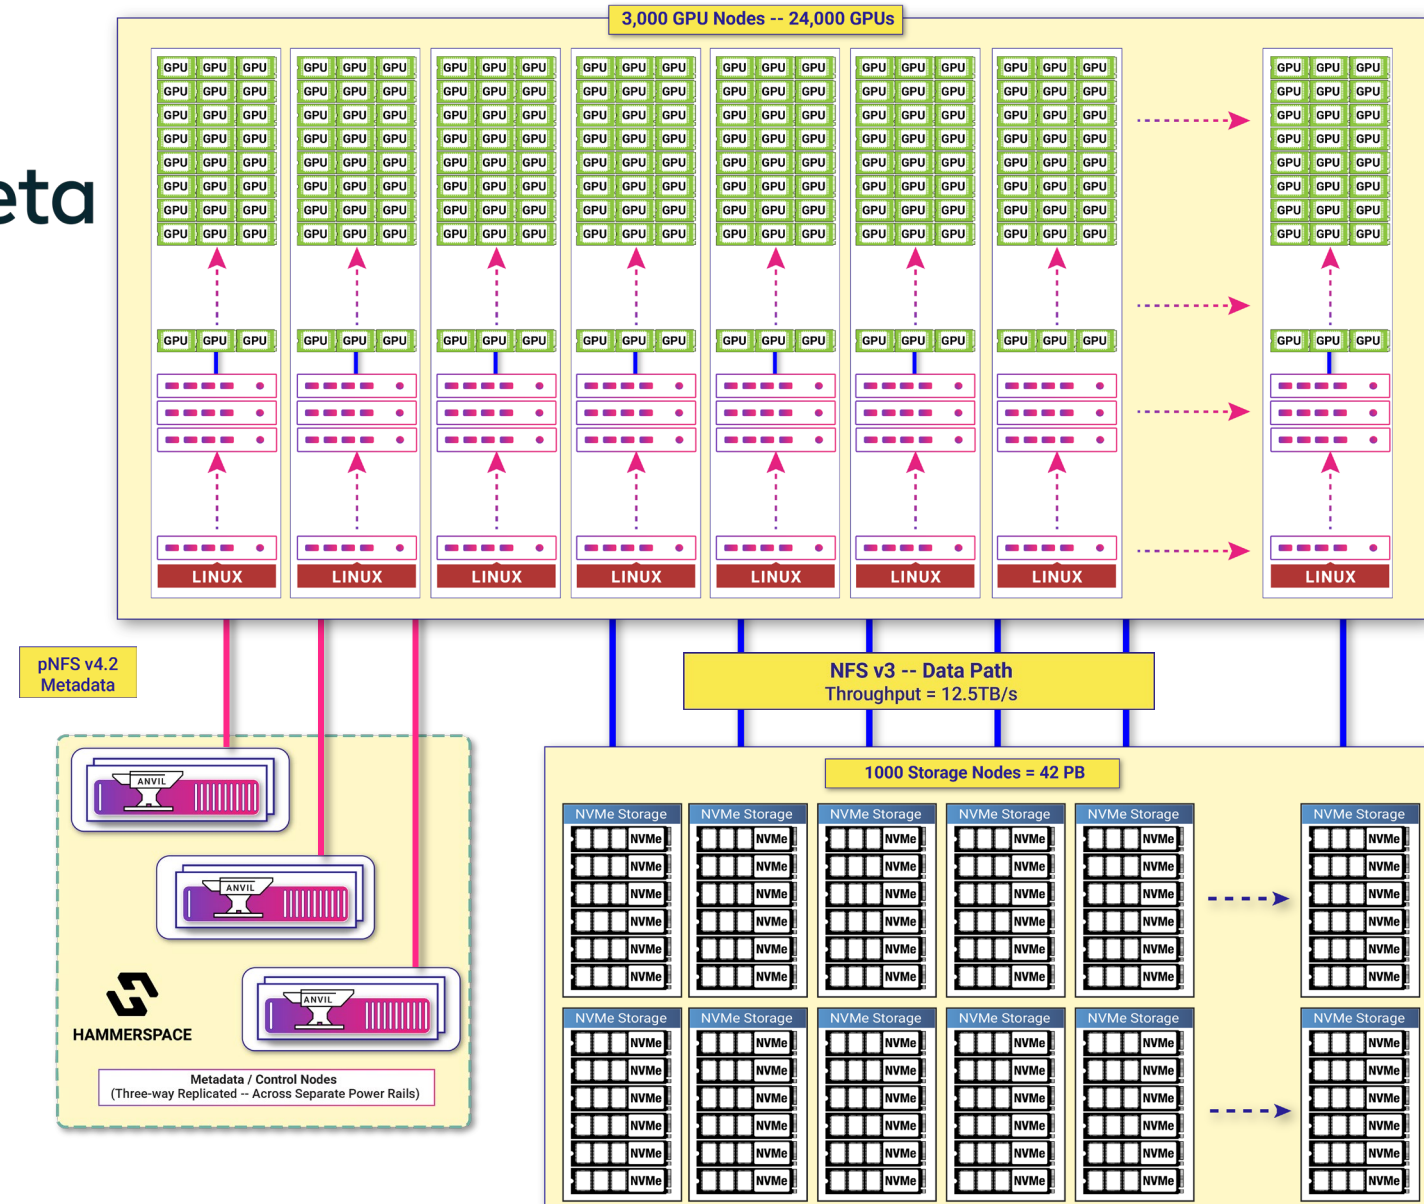
## About the Customer

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



## Hammerspace Solution

- No vendor even came close to Hammerspace’s capabilities
- 1,000+ node Hammerspace storage cluster
- Feeding 24,000 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



# Addressing Requirements of AI and GPU Computing

Deep Learning

LLM Training

Generative AI

Scientific Computing

Video and Image Rendering

Data Analytics

Extreme Parallel Performance to Feed GPUs



Hyperscale NAS Architecture

Gather and Assemble Existing Data



Data-in-Place Assimilation

Efficiently Move Data to Compute Resources



Data Orchestration

Multi-Site, Hybrid-Cloud, Multi-Cloud



Global File System

Metadata-Driven Workloads and Pipelines



Programmable Metadata

Data Security and Governance



Advanced Data Services



Thank you



HAMMERSPACE