

SNIA DEVELOPER CONFERENCE



BY Developers FOR Developers

September 16-18, 2024
Santa Clara, CA

Integrating S3 Into Distributed, Multi-Protocol Hyperscale NAS

Alan Wright

Software Architect, Hammerspace

Agenda

- The problem of adapting typical enterprise data environments for HPC-class AI workloads.
- Bridging on-prem data silos with standards-based Parallel NFS.
- Extending this to cloud & edge with an S3 service
- S3 Service Details

Major Industry Trends Driving AI Storage Requirements

LLM Training
Load and Iterate in GenAI



Large, Decentralized
Data Sets



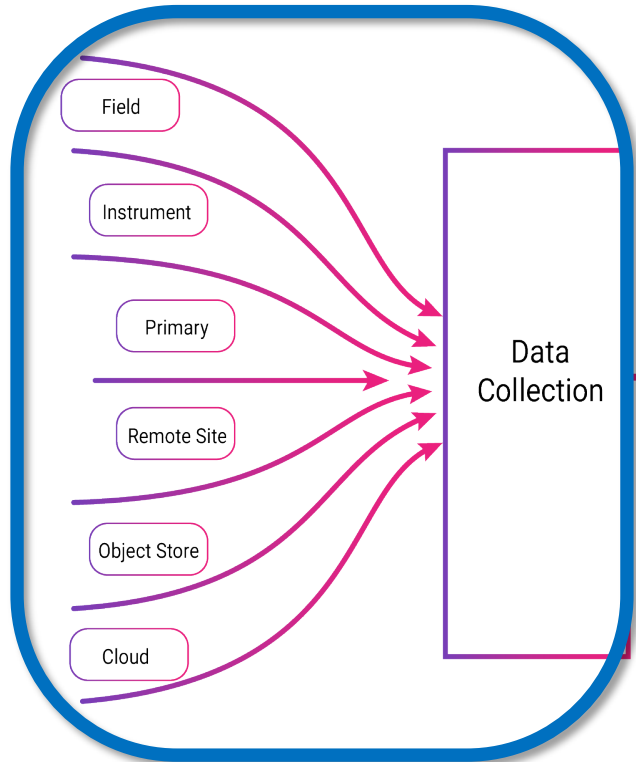
Multi-Site
Multi-Cloud
Remote AI Researchers



Barriers to Implementing AI Projects

Data Aggregation Issues

- Siloed in disparate locations
- Difficulties assembling large data sets

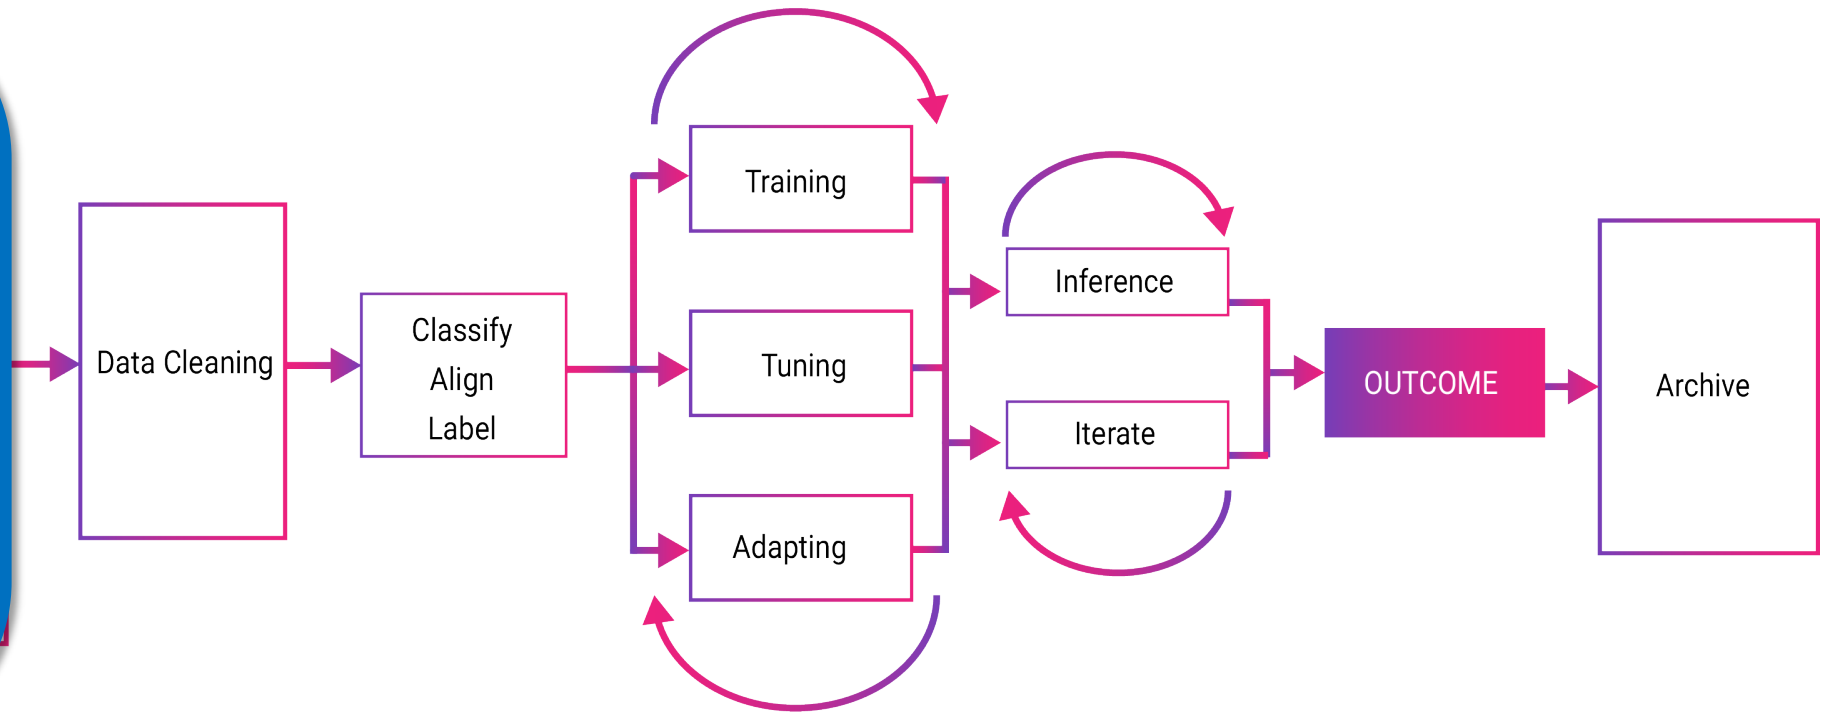


Poor Tech Infrastructure

- Existing storage systems lack performance to feed GPUs

GPUs Not Close to Data

49% of companies expect to run AI projects both in-cloud and on-prem by 2025¹




AI Use Cases Need Global High-Performance Data Access

- Prior to AI use cases, unstructured data was managed hierarchically.
- This resulted in data spread across multiple storage silos, data centers, and/or clouds

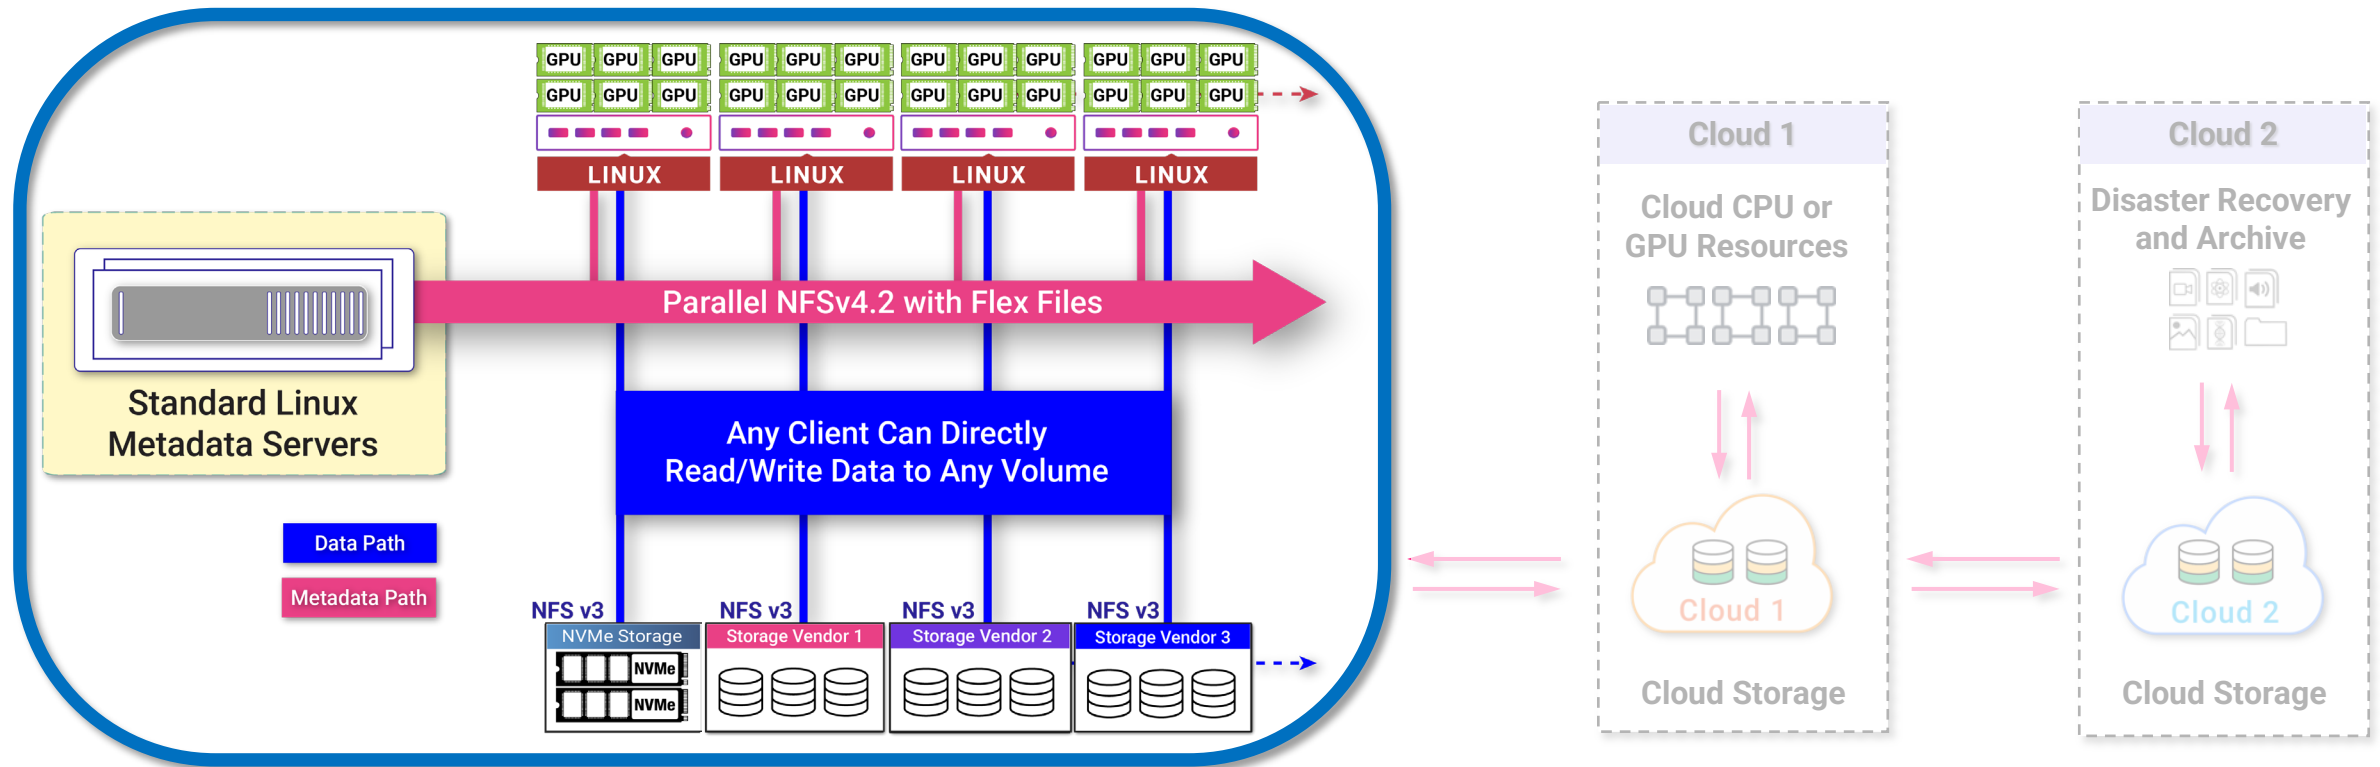
Data Aggregation Issues

- Siloed in disparate locations
- Difficulties assembling large data sets



Parallel NFSv4.2 Bridges On-Prem Silos

- pNFS v4.2 with Flex Files provides HPC-class performance compatible with existing NFSv3 storage.
- This bridges on-premises silos, and enables high-performance to accelerate even existing storage.



PNFS v4.2 at Scale: Meta's AI Research SuperCluster

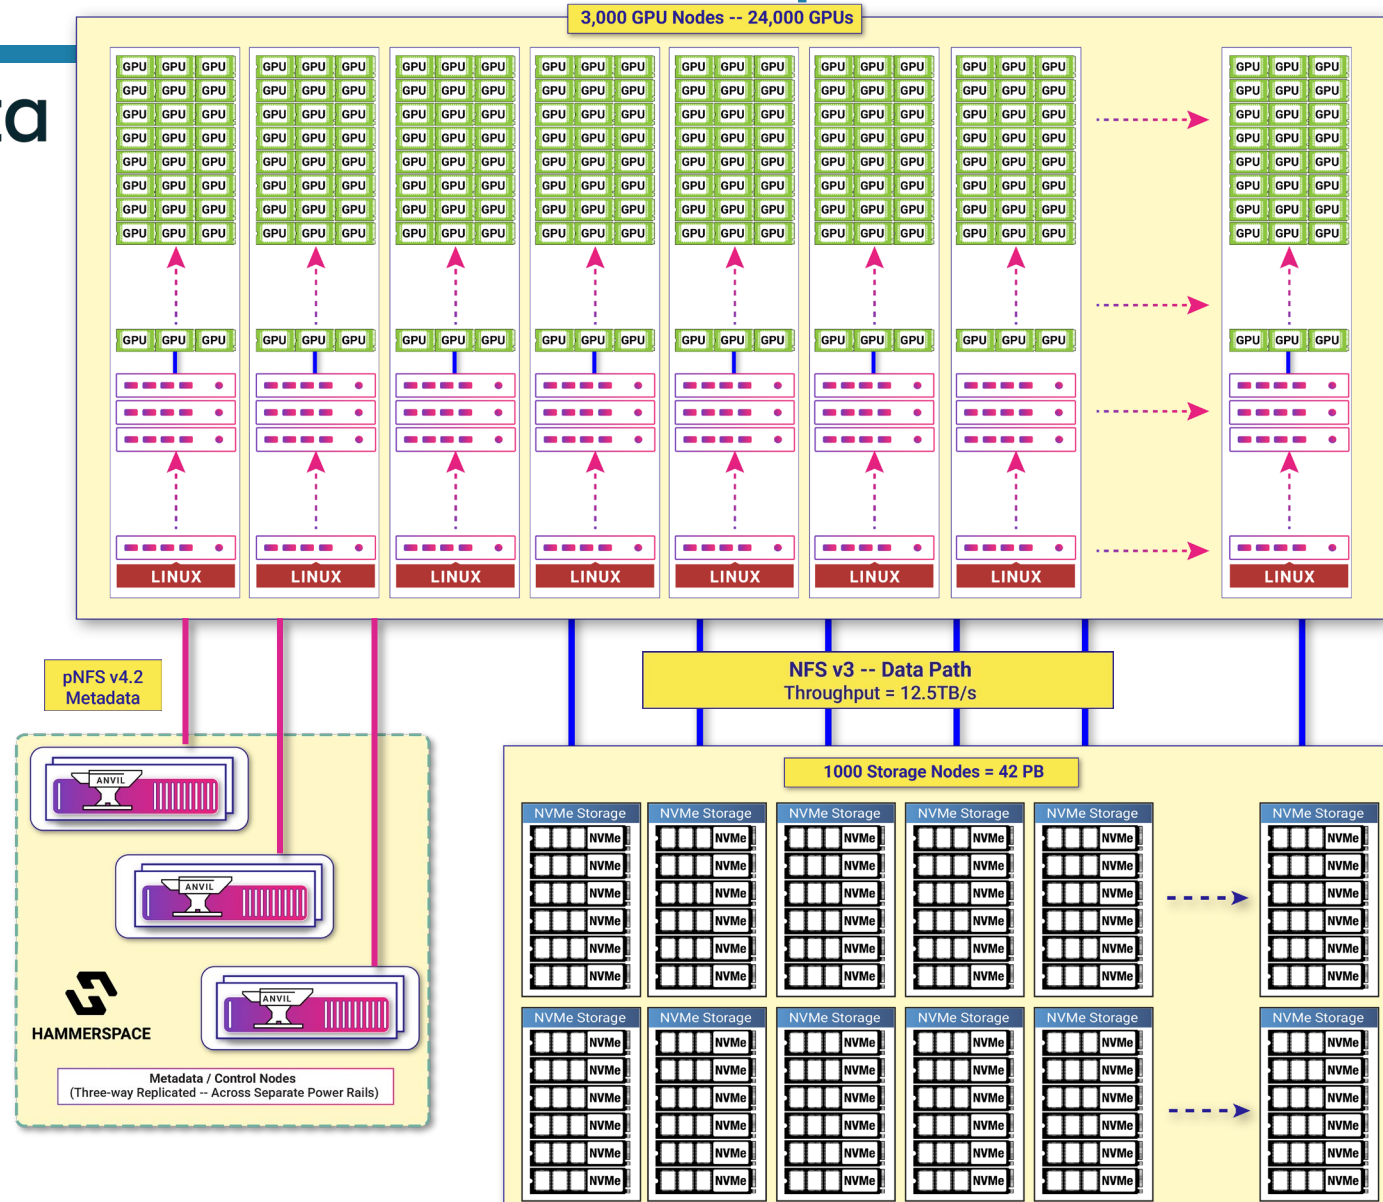


About the Customer

- Meta's AI Research Super Cluster
- Powering Llama 2 & 3 LLMs
- Massive performance and scale demands
- Evaluated leading storage vendors


Hammerspace Solution based on pNFS v4.2

- Triple redundancy on metadata nodes
- 42PB across existing 1,000+ node 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**



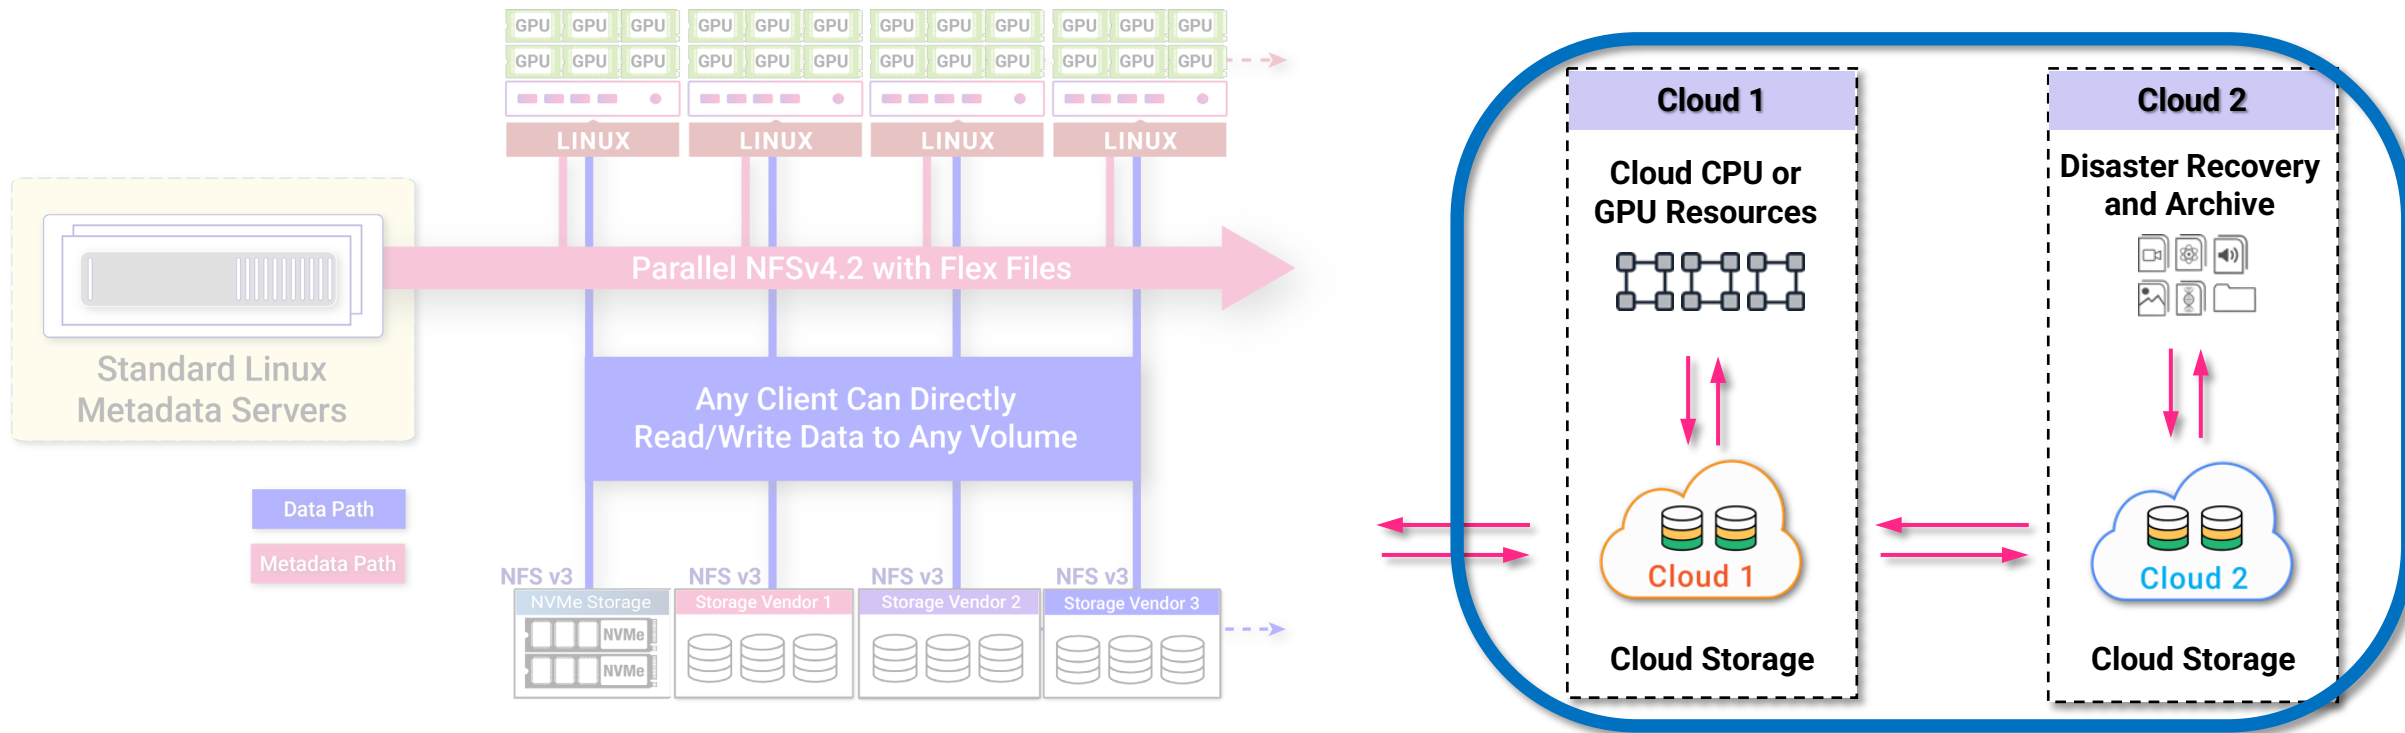
Solving the Problem of Remote Data

- Data needed for AI pipelines is often in remote sites, in archives, or the cloud.
- Such datasets may be in HDFS, or other storage that can be accessible via S3.
- This required a way to seamlessly bridge S3 storage into high-performance pNFS environments.



Solving the Problem of Remote Data

- Data needed for AI pipelines is often in remote sites, in archives, or the cloud.
- Such datasets may be in HDFS, or other storage that can be accessible via S3.
- This required a way to seamlessly bridge S3 storage into high-performance pNFS environments.



S3 Service

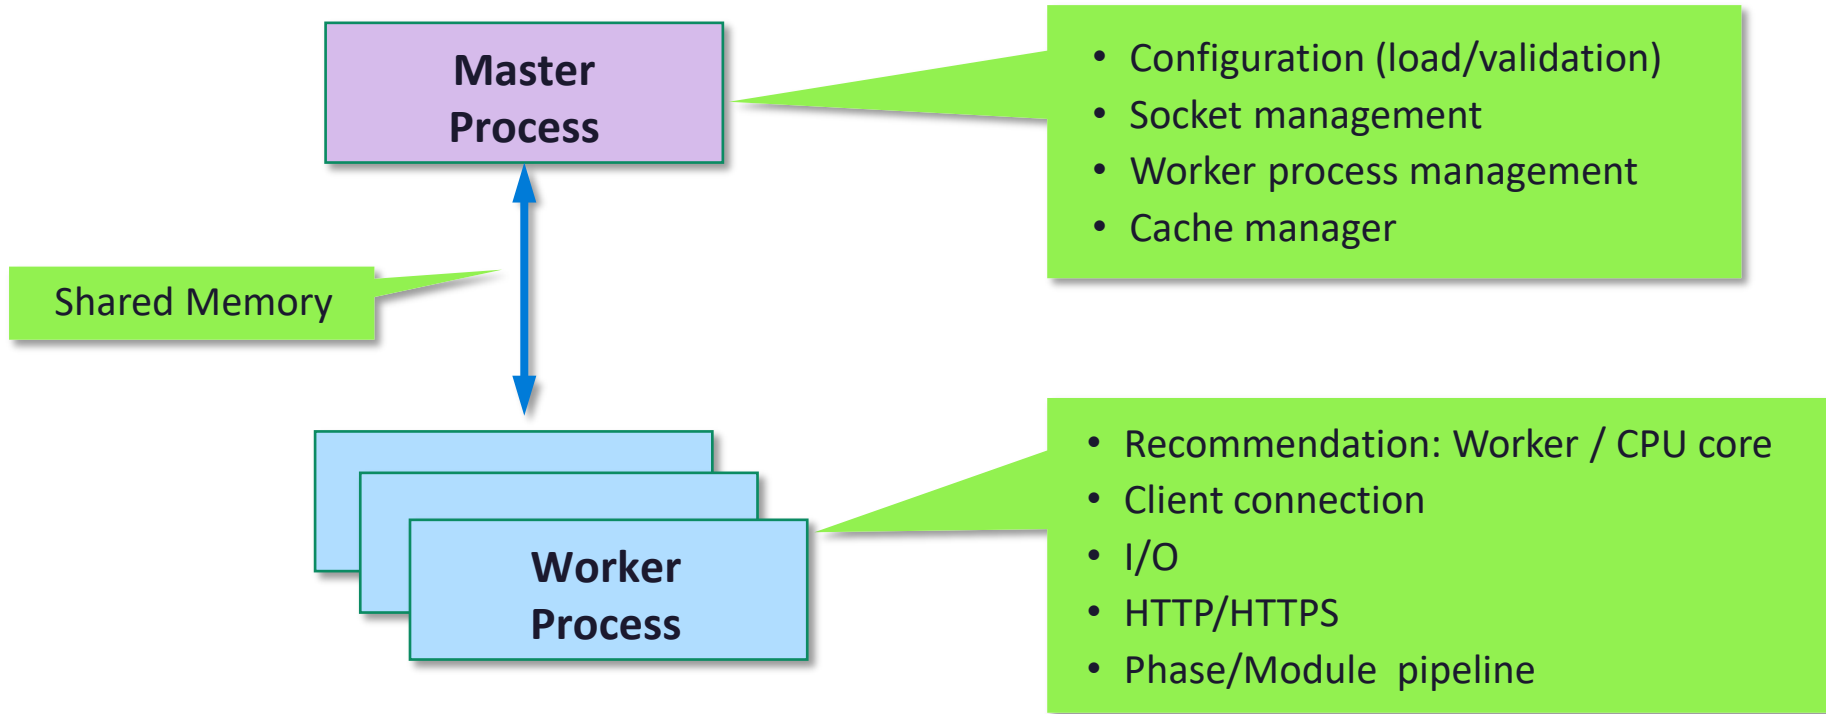
nginx web server

- <https://nginx.org/en/>
- Pronounced “Engine X”
- Open Source web server
- HTTP/HTTPS
- Multiple web sites
- High concurrency
- Multi-process
- Single threaded
- Asynchronous connections

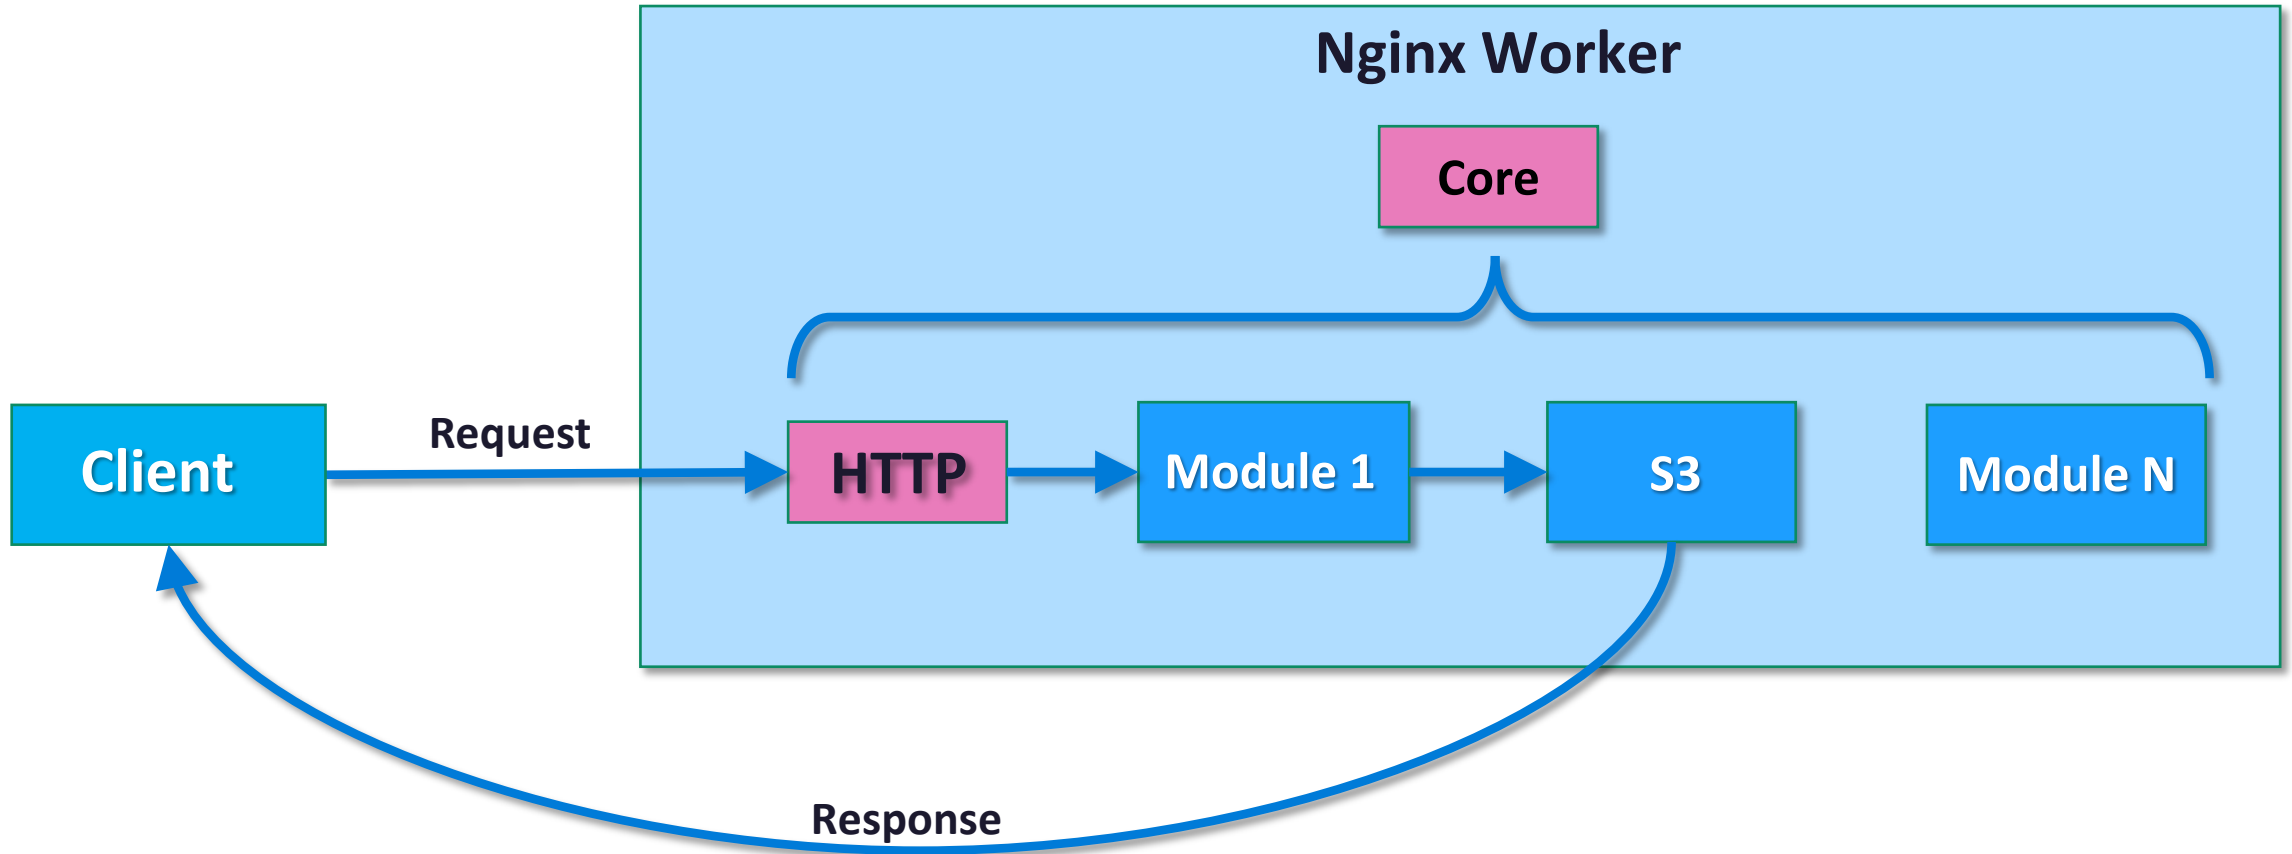
S3 Service

- nginx module
- S3 API (<https://aws.amazon.com/s3/>)
 - Buckets, Objects, Versioning, ACLs, Tagging etc.
- File System Integration
 - shared multiprotocol namespace
 - snapshots, versions,
 - ACLs, tags, objectives etc.
- Active Directory Integration
 - users/groups

Nginx Process Model

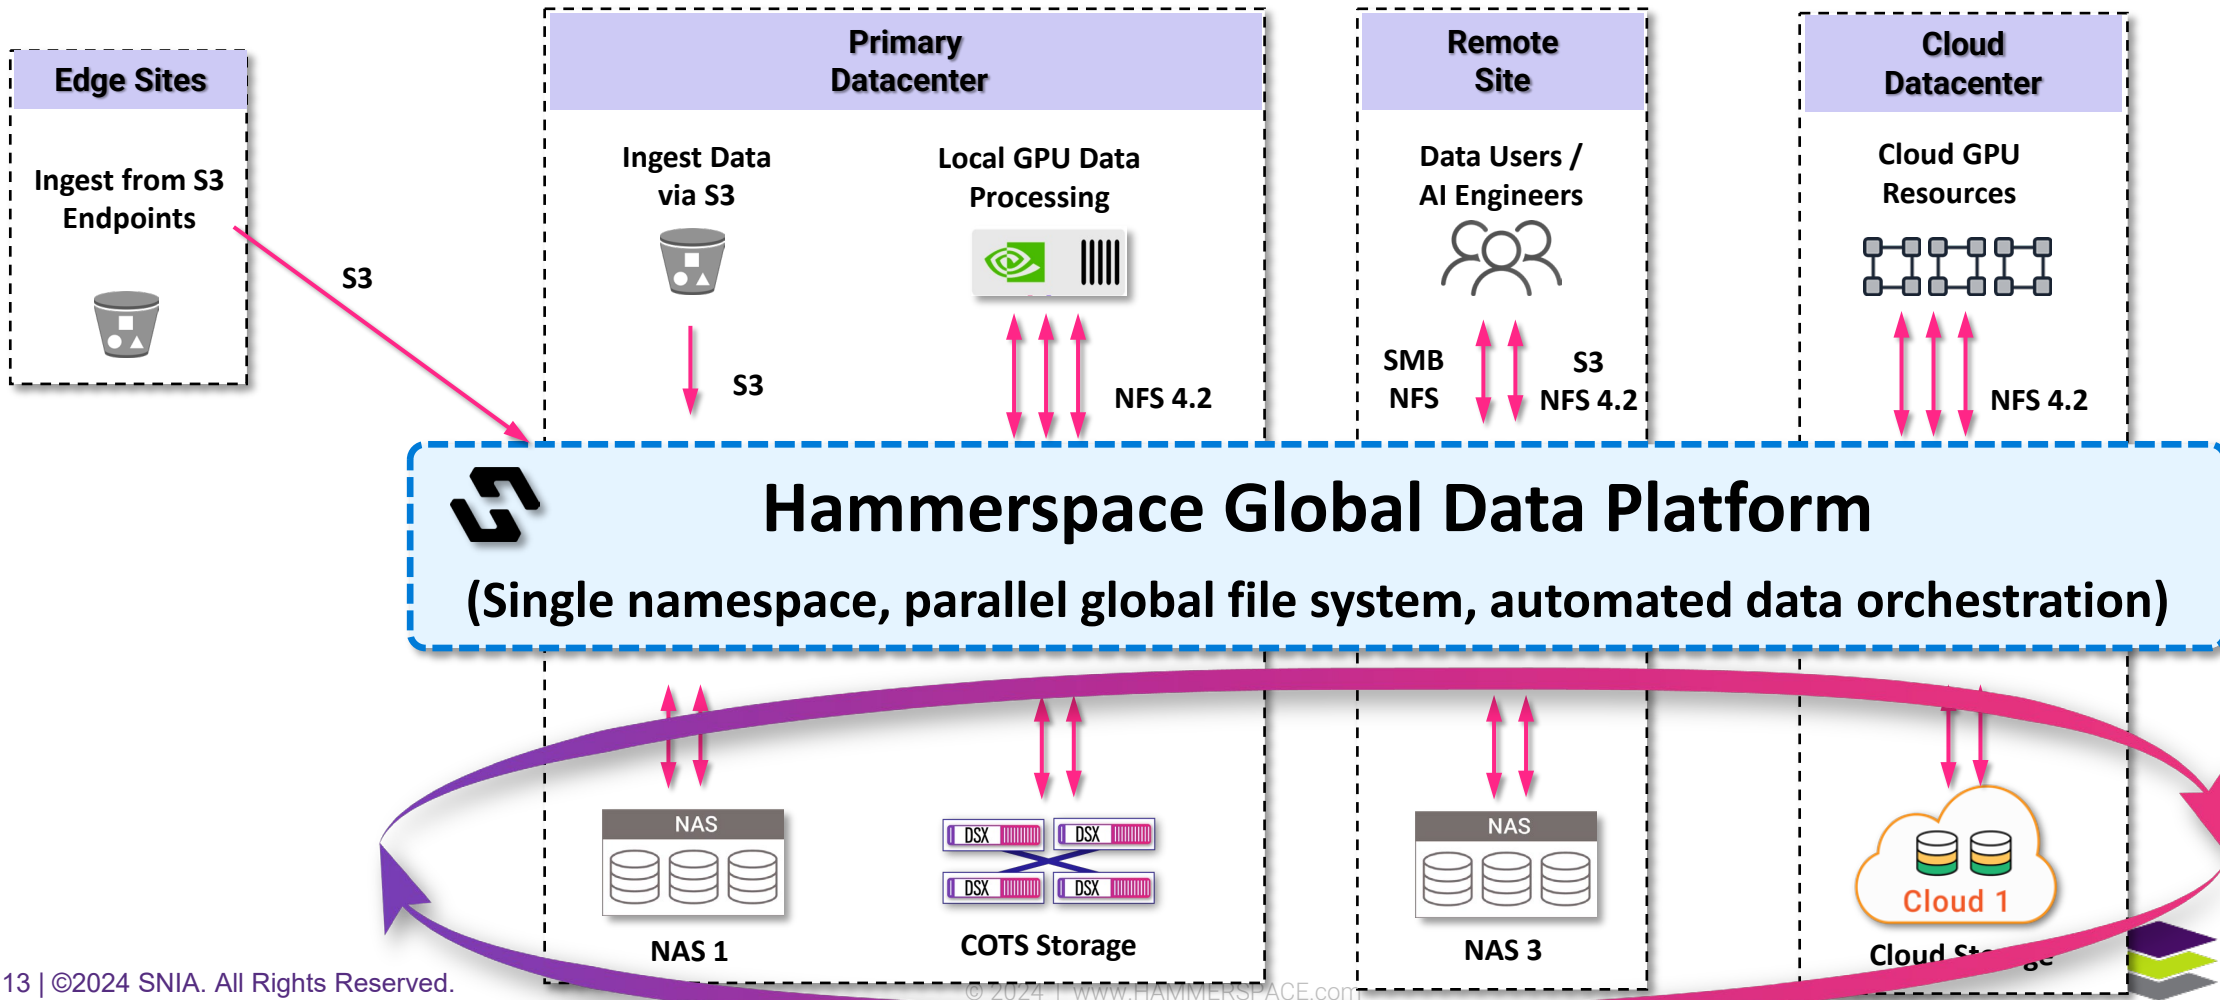


Nginx Module Chain



Global Multi-protocol Data Access Across All Storage

- The S3 service ties remote storage into a unified global data environment.
- Automated data orchestration enables high-performance use cases across incompatible silos.



Thank You

Please take a moment to rate this session.

Your feedback is important to us.