

The Industry is Building AI Factories



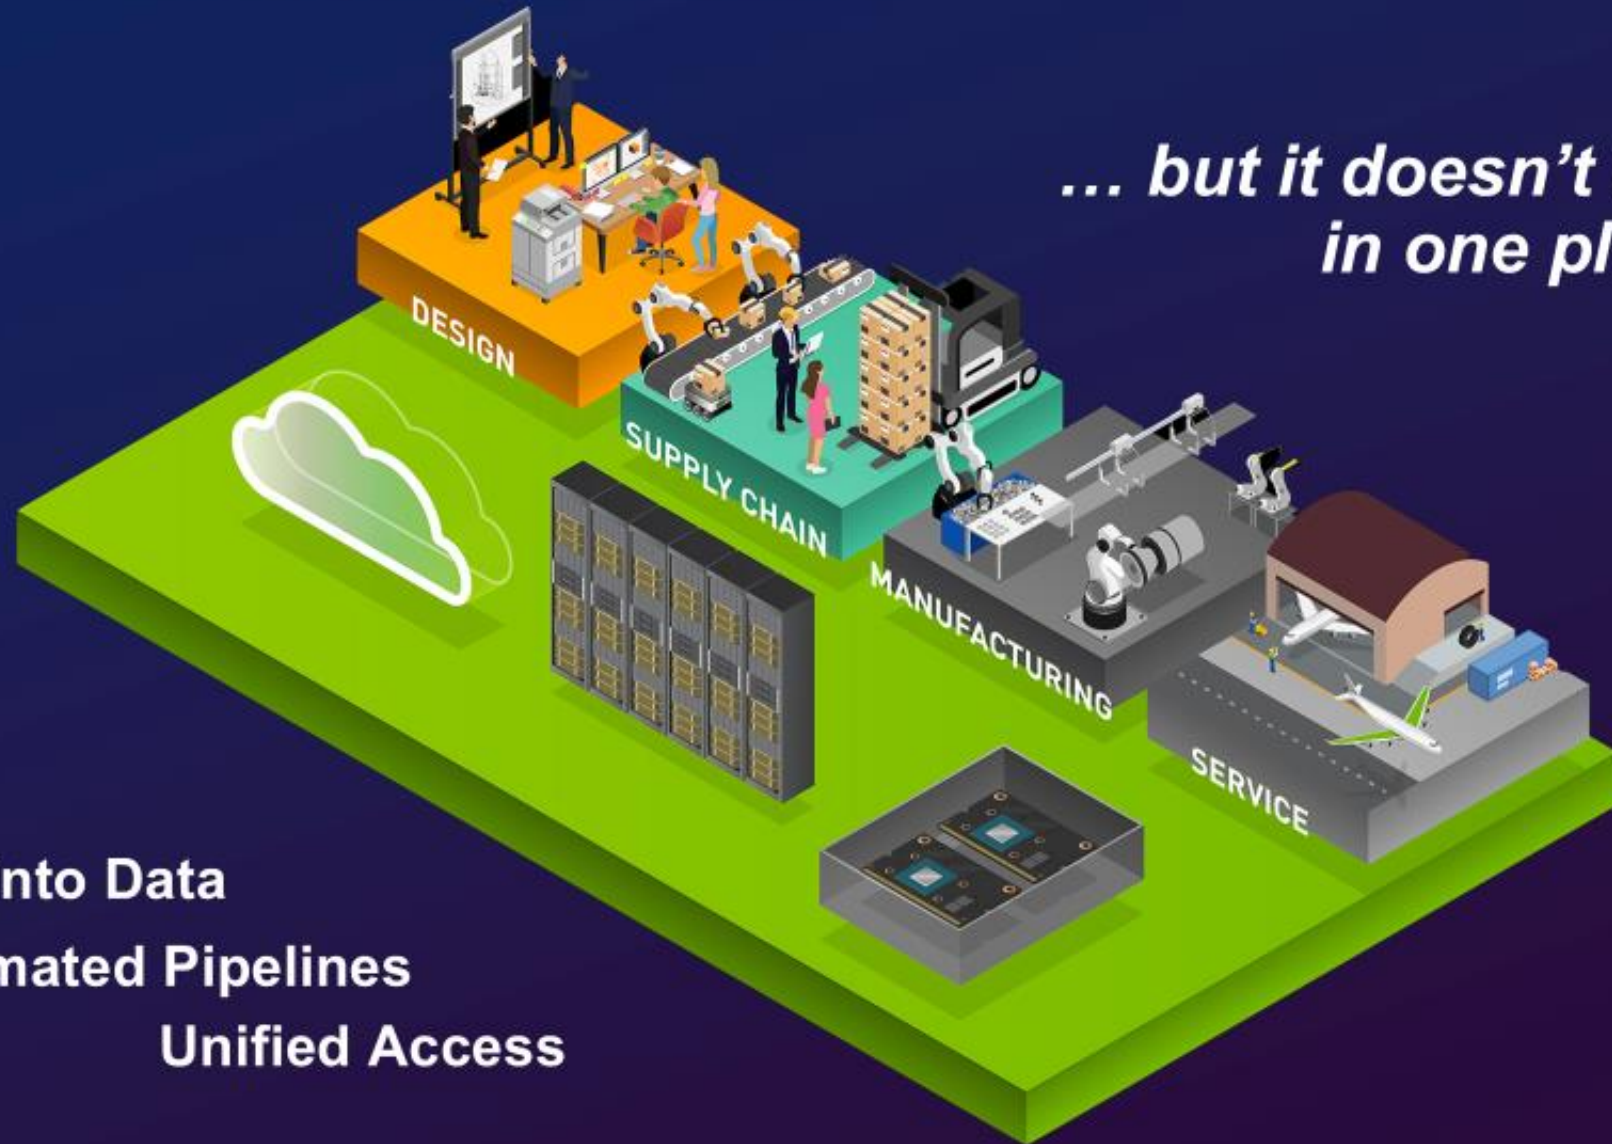
But Something is Missing

Data underpins every AI solution...

... but it doesn't live in one place

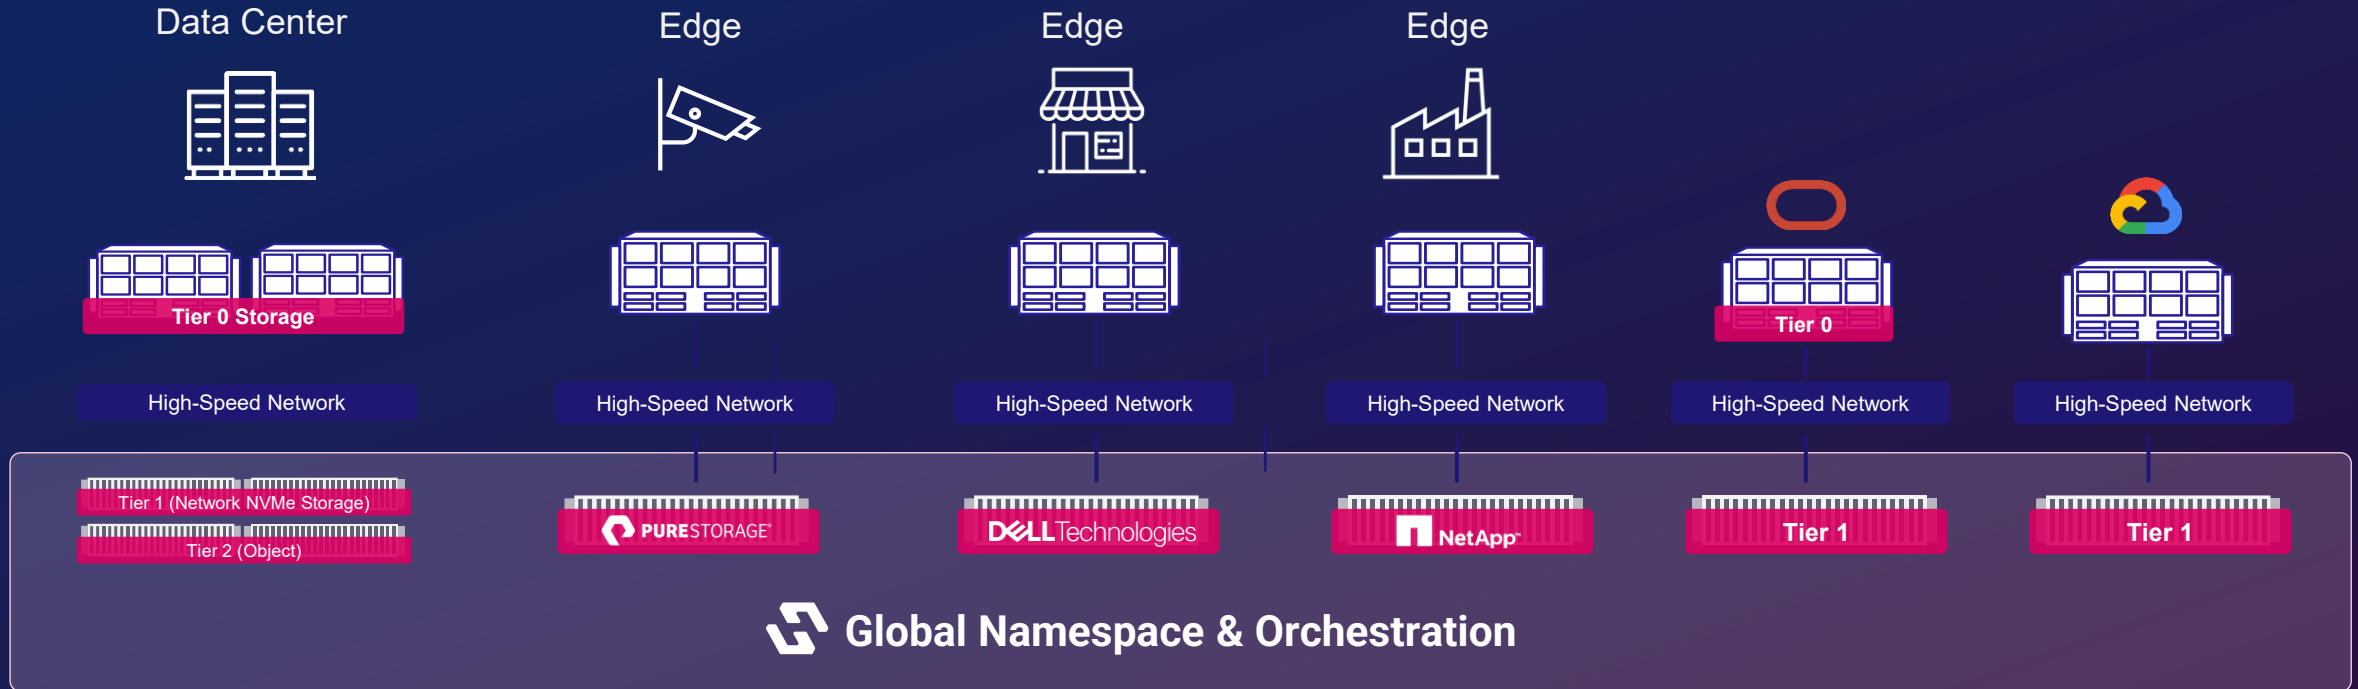
Tomorrow's
Enterprise
Needs...

Visibility into Data
Automated Pipelines
Unified Access



Making AI Anywhere a Reality

The Data Foundation for Generative, Agentic and Physical AI



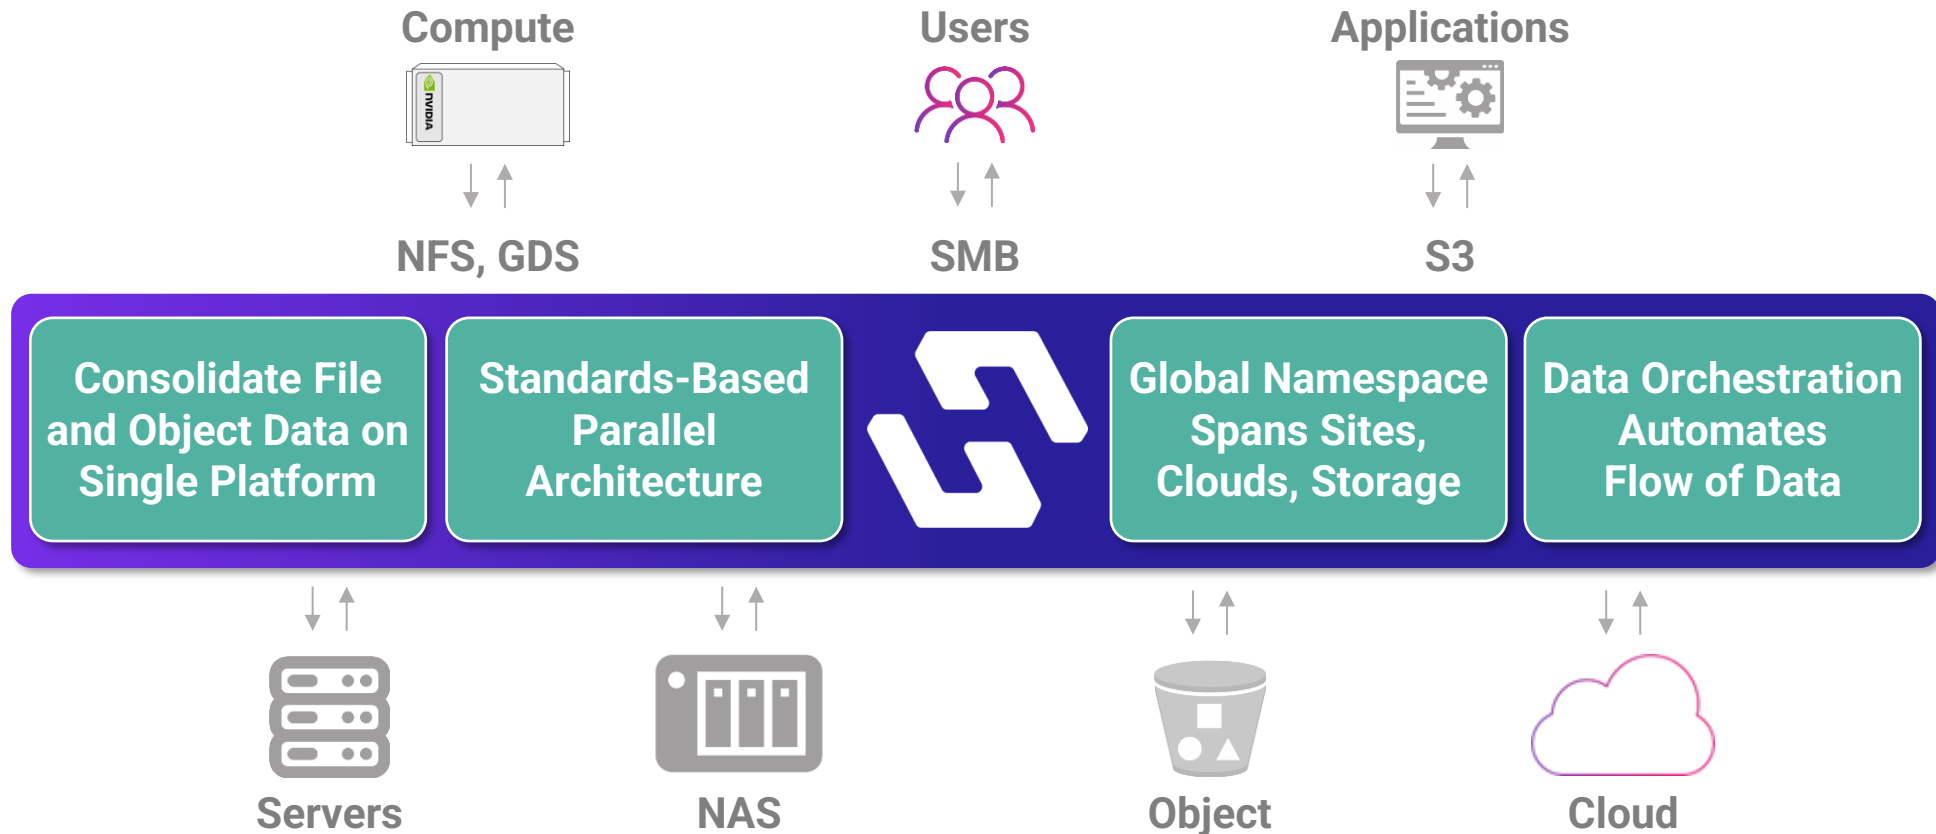
**Data-centric
Not storage-centric**

**Open, automated,
intelligent**

**Transcends location,
provider, & protocol**

The Hammerspace Global Data Platform

High Performance Data Access Using Industry Standard Protocols



Software-Defined and Storage-Agnostic



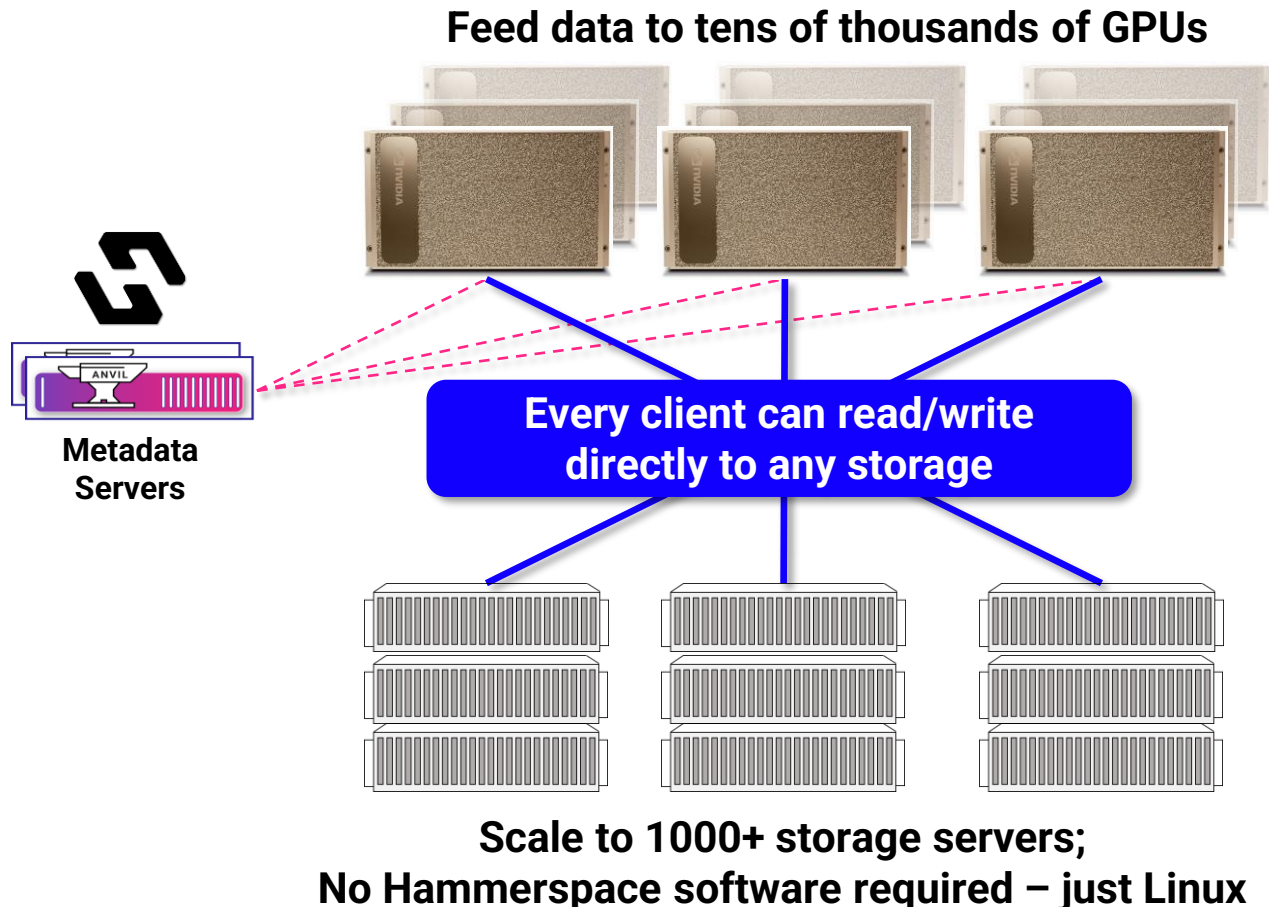
Hammerspace Tier 0



**GPU server local storage typically goes unused.
Today 100s of TBs. Tomorrow 2 PB+.**

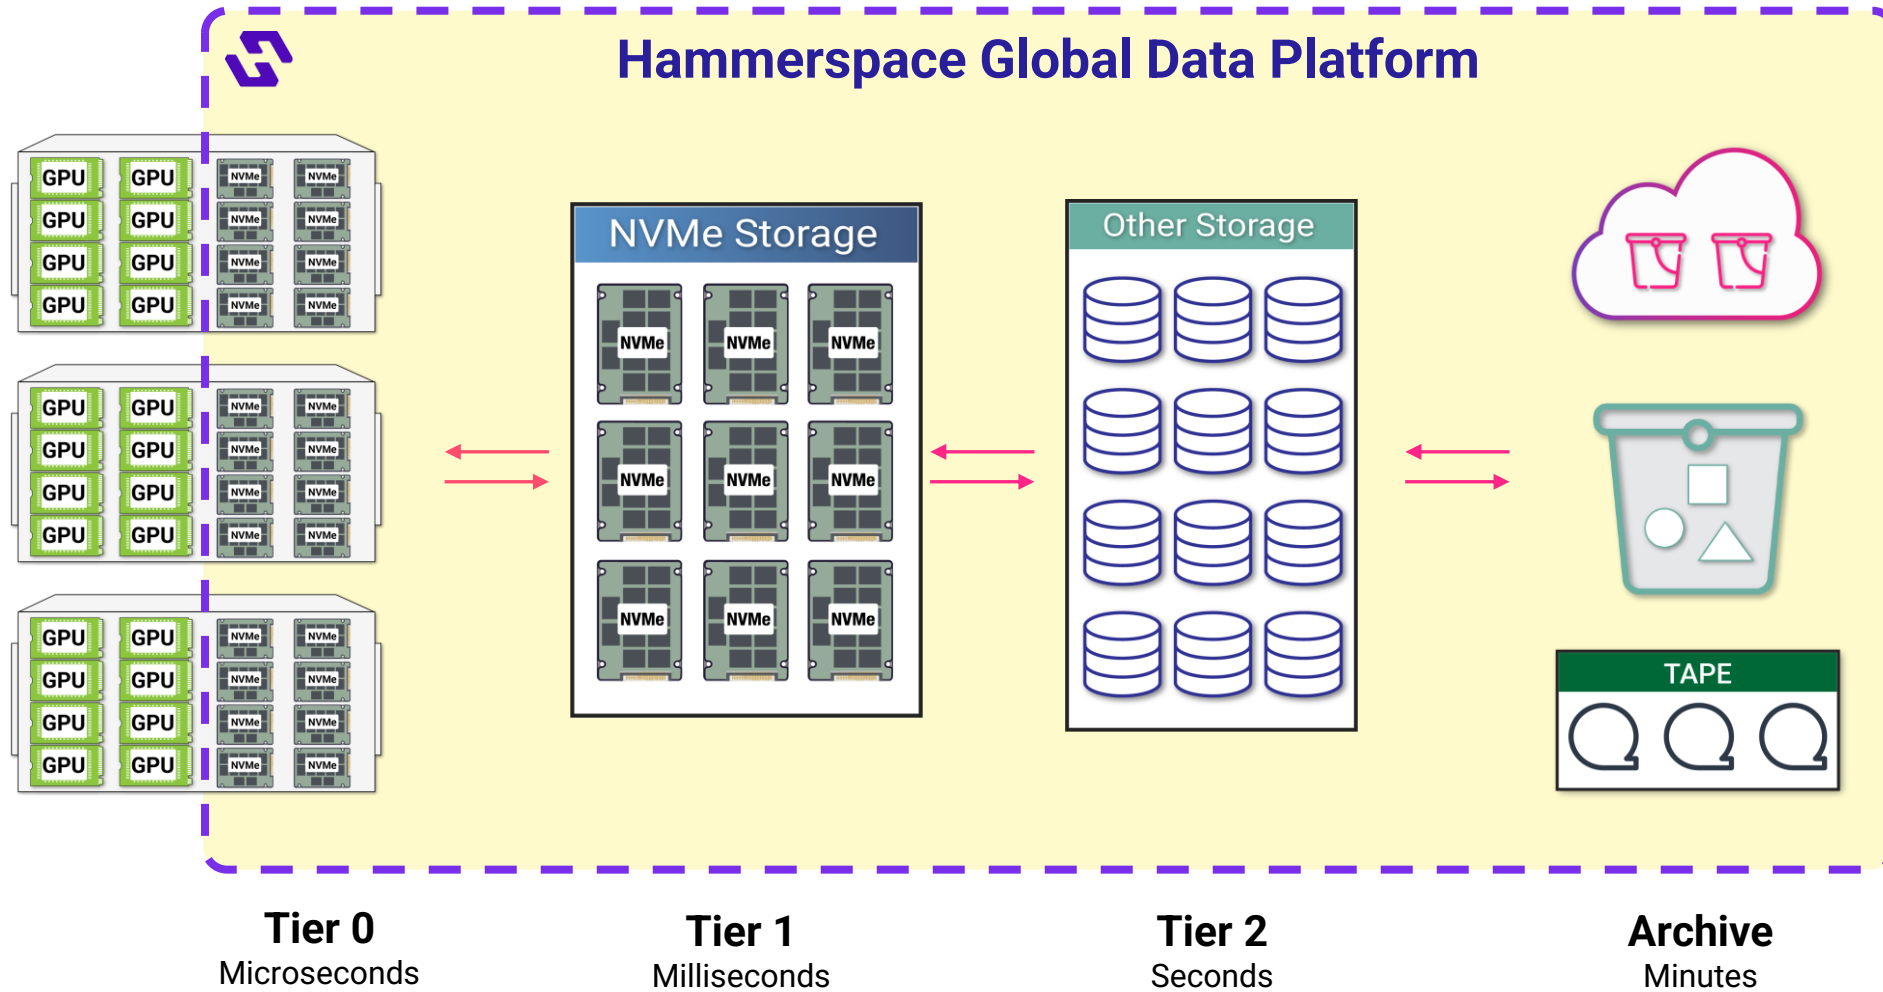


Hammerspace and NVIDIA: Tier 0



- ✓ **GPUDirect Storage Certified**
- ✓ **Direct Data Path:** Separate metadata layer (control plane) from data layer
- ✓ **Performance and Scale for BasePOD and SuperPOD Deployments**
- ✓ **Standards-based:** no proprietary client software
- ✓ **Lowest Latency Tier 0:** Use the local NVMe storage in the GPU servers for lowest-latency “tier 0” storage
- ✓ **Increase GPU utilization** by increasing checkpointing performance by 3.7x vs external storage

Span Tier 0 to Archive on a Single Data Platform



Customer Story: George Washington University

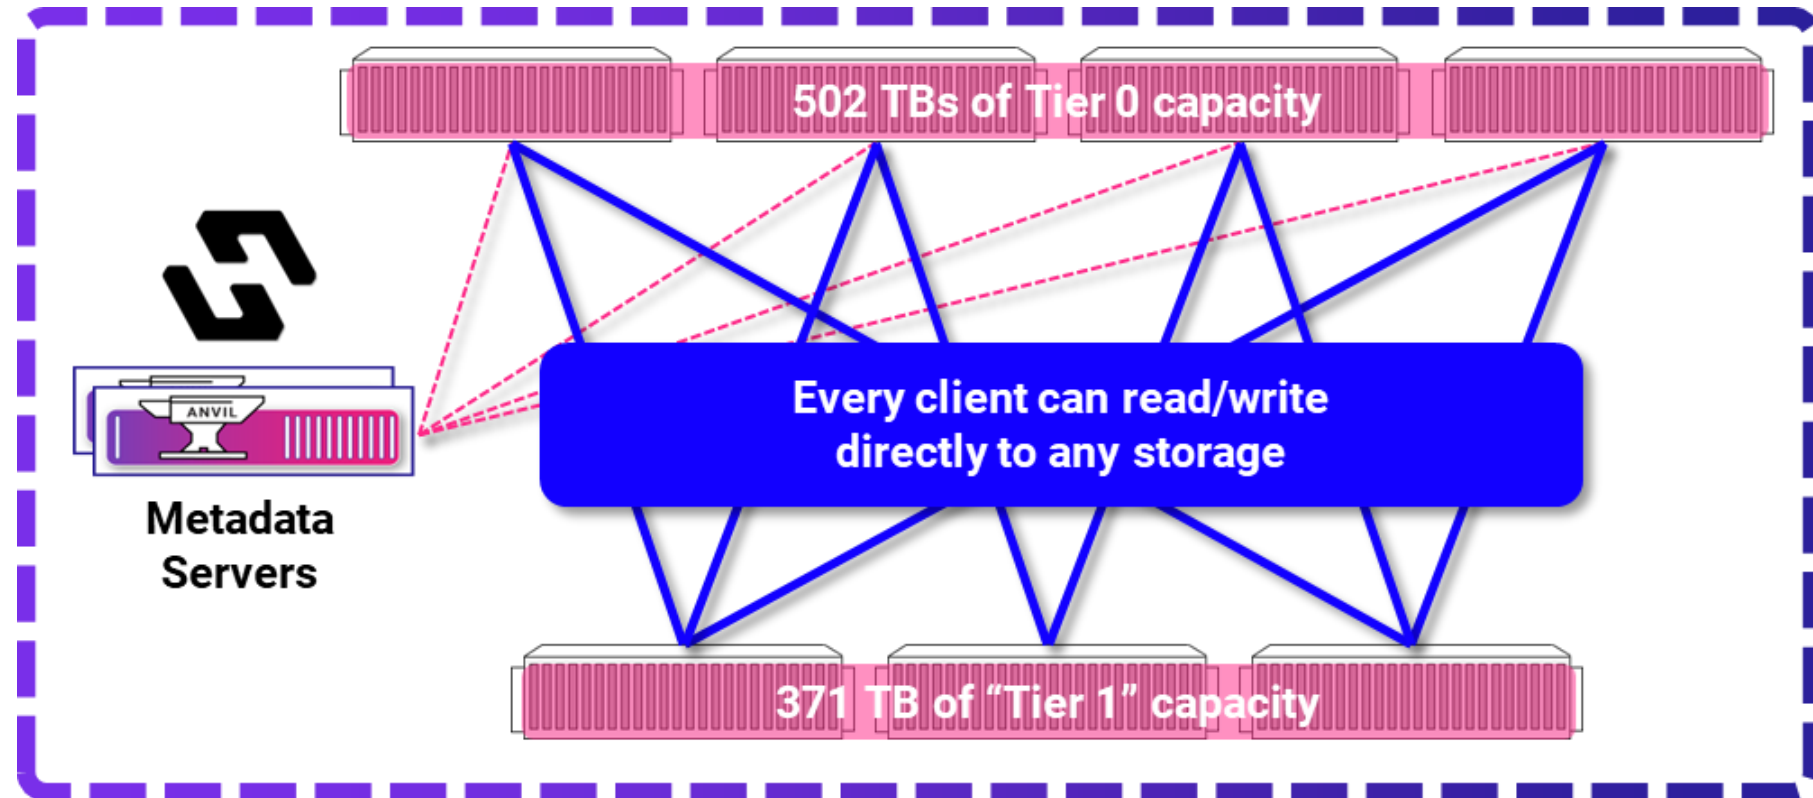


Research university

New AI initiative

500 TBs "Tier 0"

400 TBs "Tier 1"



Cost and Power Savings of Tier 0 (10 PB example)

~\$3-4M

Cost savings of external flash storage hardware

~\$500K

Energy cost savings over three years

~3M kWh

Reduced energy consumption over three years

~3-5 Racks

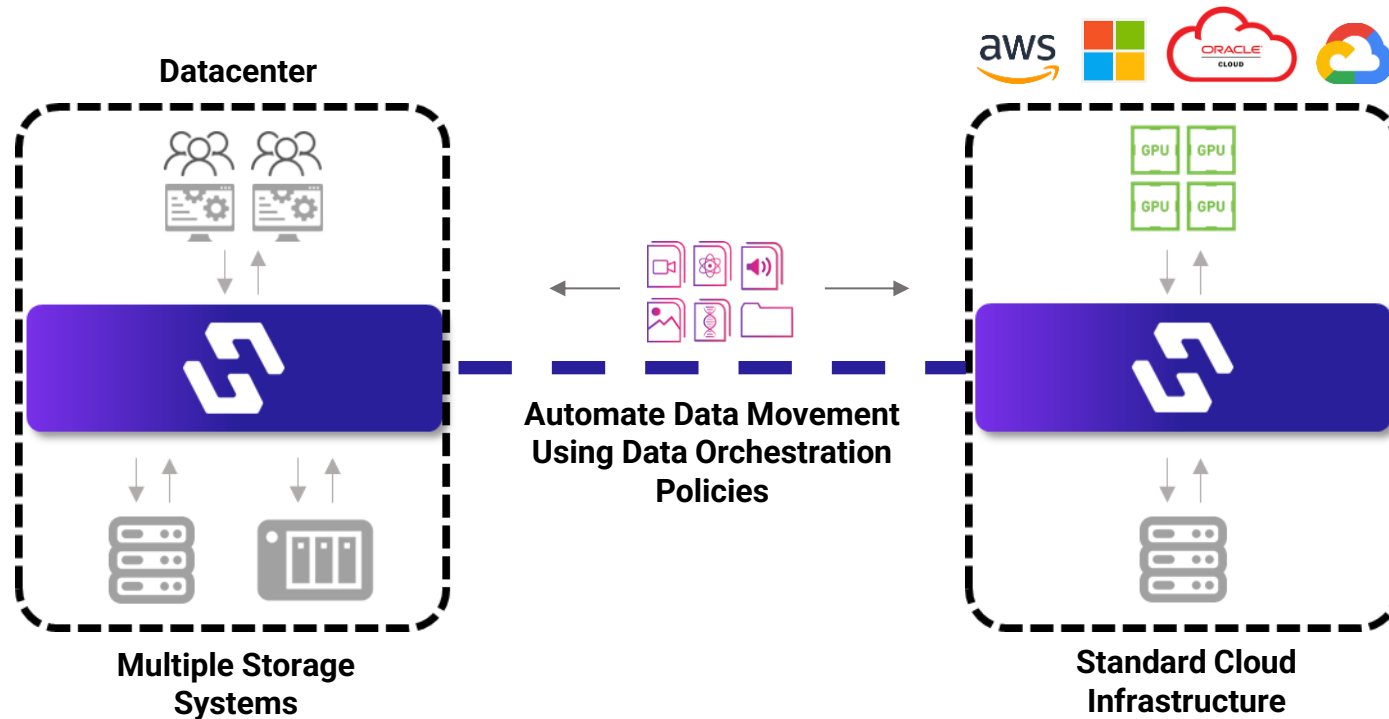
Data center space gained

Source: Tier 0 Savings Estimator at <https://hammerspace.com/tier-0/>



Orchestrate Between On-Prem and GPU Clouds

Single Global Namespace that Spans Sites, Clouds, Storage

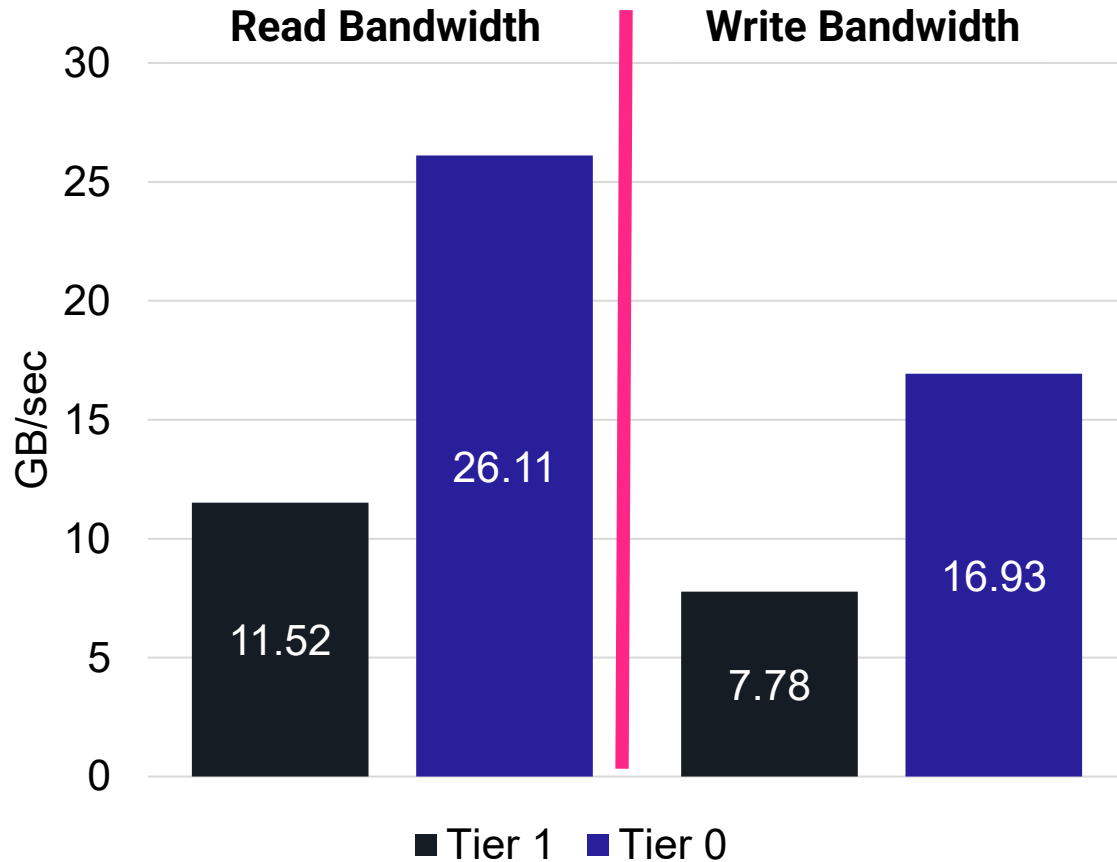


- ✓ **Bring the data to the GPUs** wherever they are located
- ✓ Take advantage of the **same parallel architecture** in the cloud, including Tier 0
- ✓ **Gain cloud agility** to use NVIDIA GPUs in the cloud(s)
- ✓ **Automate data movement** between on-prem, cloud, and multi-cloud environments
- ✓ Hammerspace is in **all four major US clouds**: AWS, Azure, GCP, OCI*

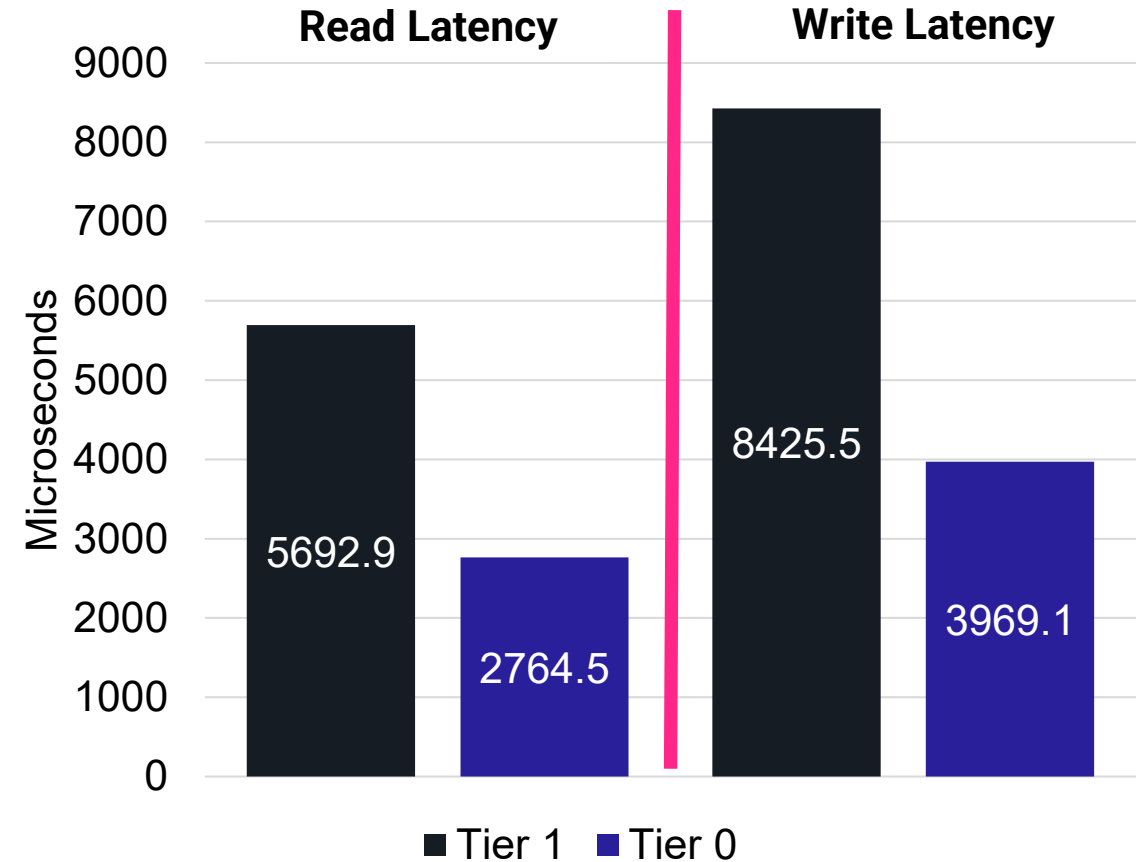


Performance Results of Tier 0 in Oracle Cloud (OCI)

Tier 0 Provides ~2.5x More Read Bandwidth, ~2x More Write Bandwidth




Tier 0 Reduces Read Latencies by ~51%, Write Latencies by 51%



Analysis based on benchmark testing using Flexible I/O (fio) tester. Clients: 2. Files: 16. Filesize: 50GB. Direct: True. Blocksize: 1MB. IO Depth: 2. IO Engine: libaio. Number of Jobs: 1 (Per File). Run Time=300. Workloads: 100% Sequential Read, 100% Sequential Write, 100% Sequential 50/50 Read/Write Mix, Iterations: 3. Results were averaged.

Create a Hybrid / Multi-Cloud Global Data Environment



Customer Story: Jellyfish Pictures

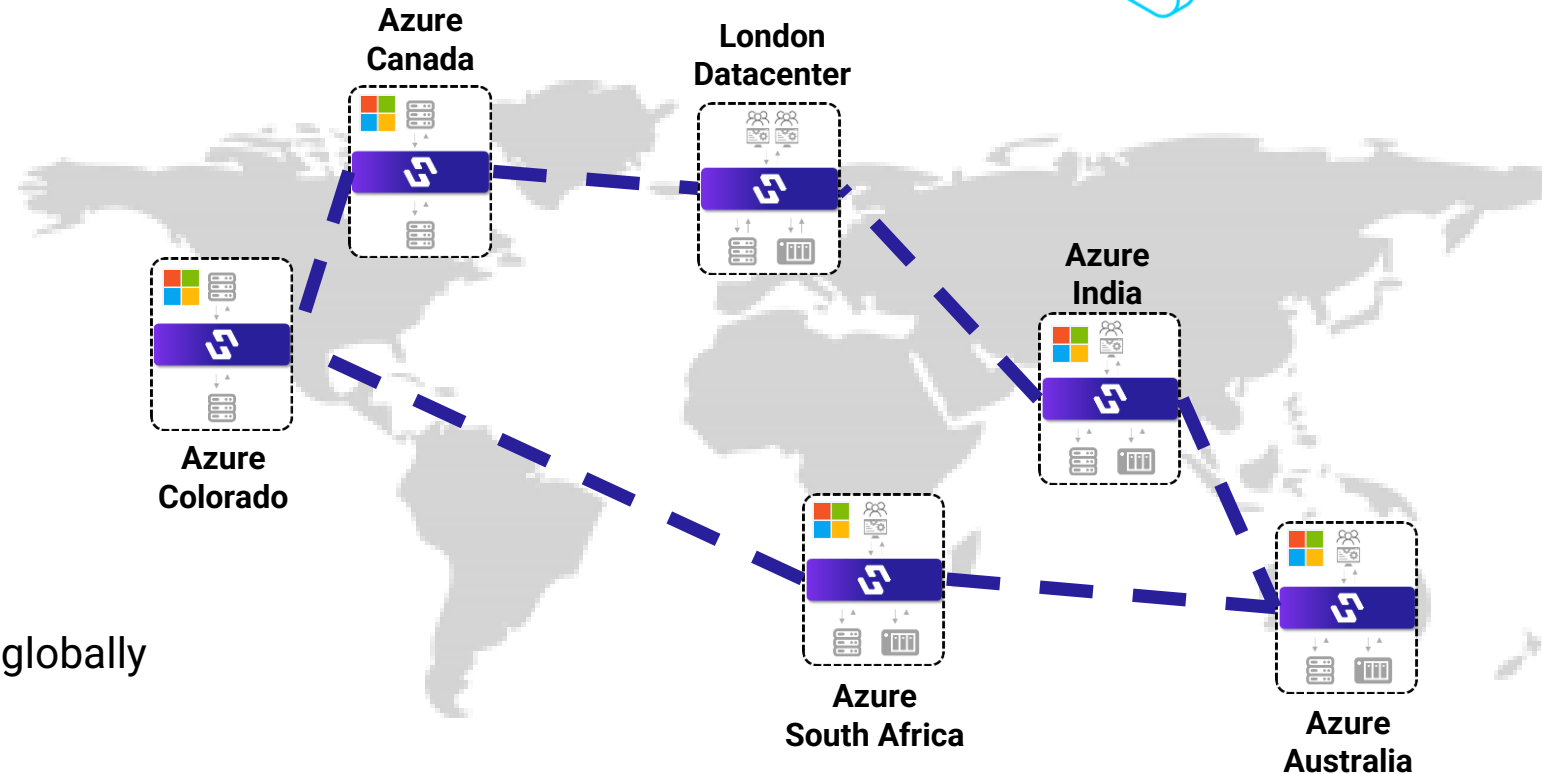


Customer and Requirements

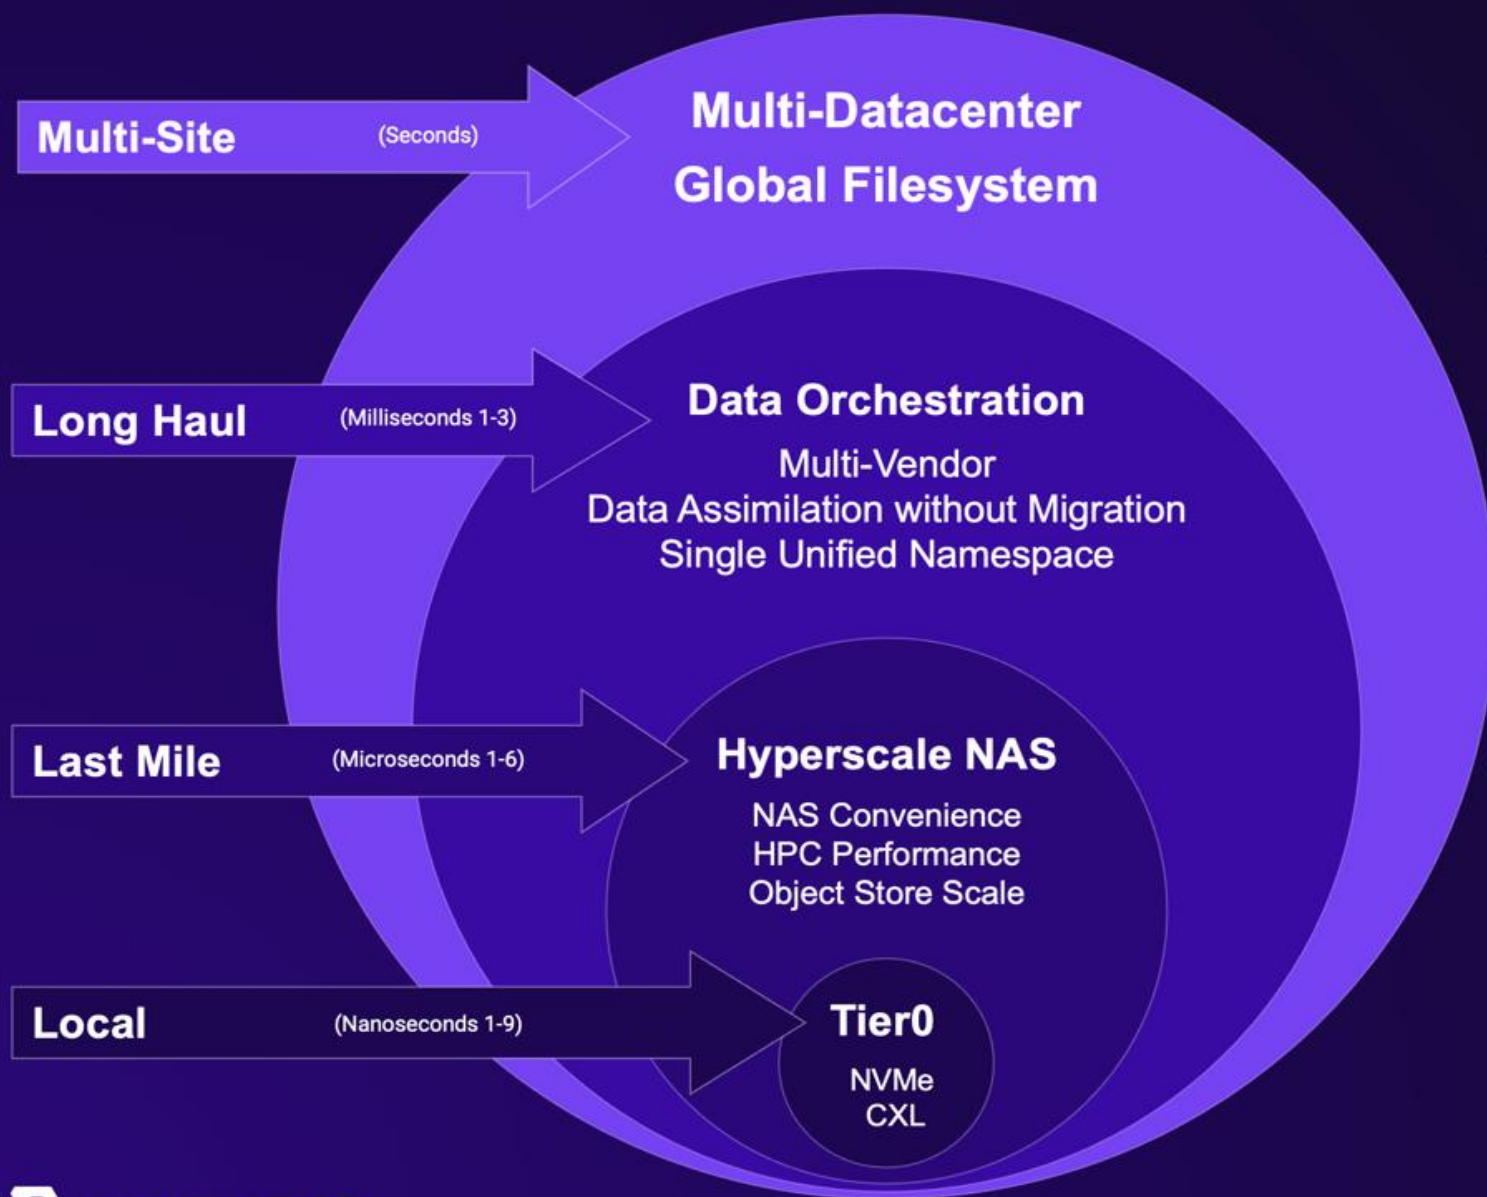
- Leading global VFX and Animation Studio
- 300+ artists around the world
- Virtual since 2019
- Storage in multiple locations and clouds
- Sharing files globally was a “real issue”

Hammerspace Solution

- Global file system spans 5 continents
- 100M+ files globally accessible
- Hammerspace orchestrates data movement globally
- On-demand burst render
- Saved \$3M+ in render costs
- Enabled access to global talent pool



Hammerspace Superpowers Solve Multiple Problems



Enterprise Data Is Scattered Across Silos

-- Hammerspace Bridges Them All --

- **Global Access and Control:**
 - Unified across multiple sites & clouds
- **Standards-based:**
 - No proprietary clients
- **No Expensive Migrations:**
 - Use data in place on any storage
- **Integrated Performance Where Needed:**
 - HPC-class performance with Enterprise NAS features
- **Ultra-fast Performance for AI:**
 - Activate underutilized NVMe within on-prem & cloud-based GPU clusters

Hammerspace Supporting AI Everywhere

