



# Activating Data for AI:

Unifying Data Without Rebuilding Infrastructure

Floyd Christofferson  
VP Product Marketing -- Hammerspace

HPC User Forum  
Austin, May 5-6

# Data underpins **every AI solution...**

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

**Today's  
Enterprise  
Needs...**



**73%**  
of Enterprise data  
is siloed


**Visibility into Data**

**Unified Access**

**Automated Pipelines**

# AI Exposes Limits of Centralized Data

AI is not just compute-constrained. It is data-constrained.



## AI Changes How Data Must be Used

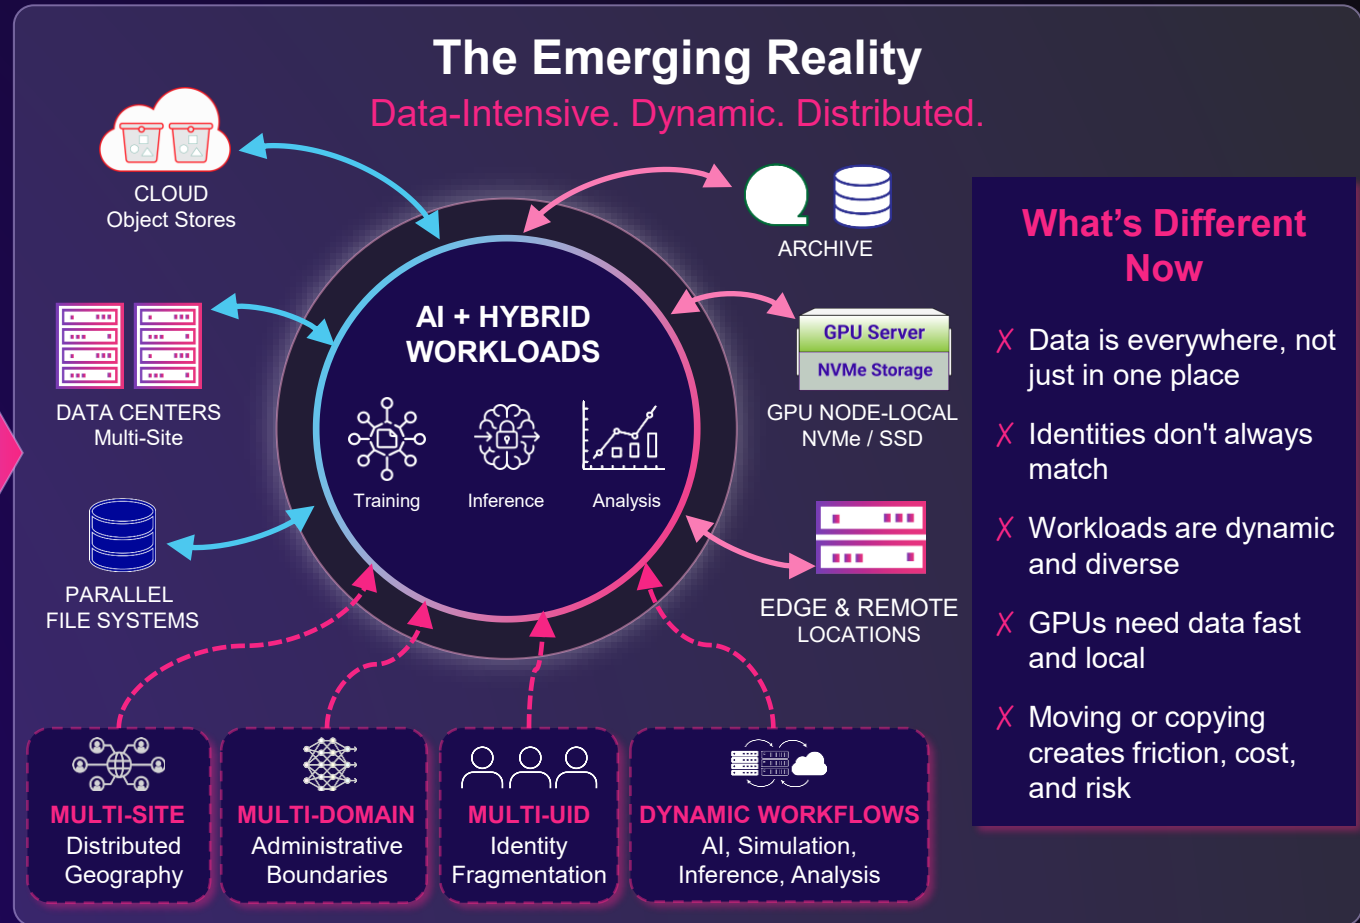
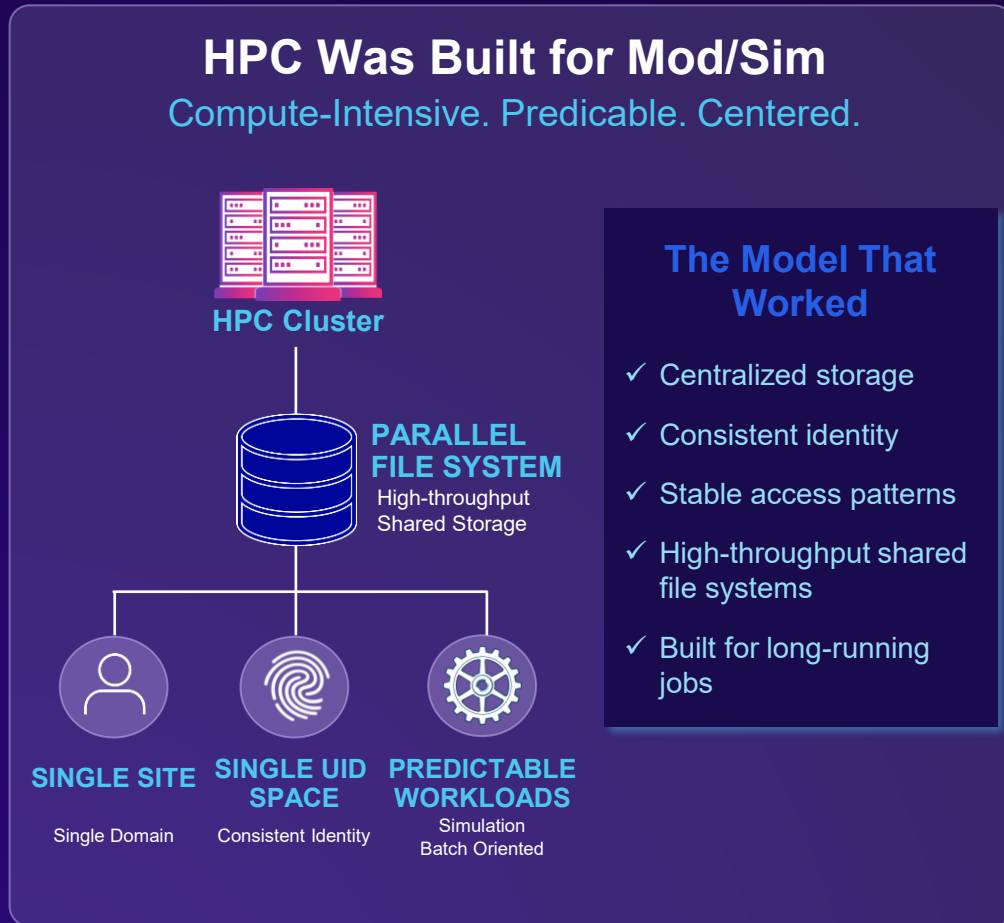
- Data stays where it is
- Compute moves to where it's available
- Workloads are no longer fixed

## What Breaks At Scale

- Copy sprawl
- Idle GPU's
- Power & cost waste
- Deployment delays

# The Reality Has Changed

From HPC Built for Mod/Sim to AI + Hybrid Workflows



**⚠ This is not just a Performance Problem. It Is a Data Orchestration Problem**

# The Problem of Distributed Data with Storage-Centric Infrastructure

## Data Sprawl

Unstructured data is siloed and distributed

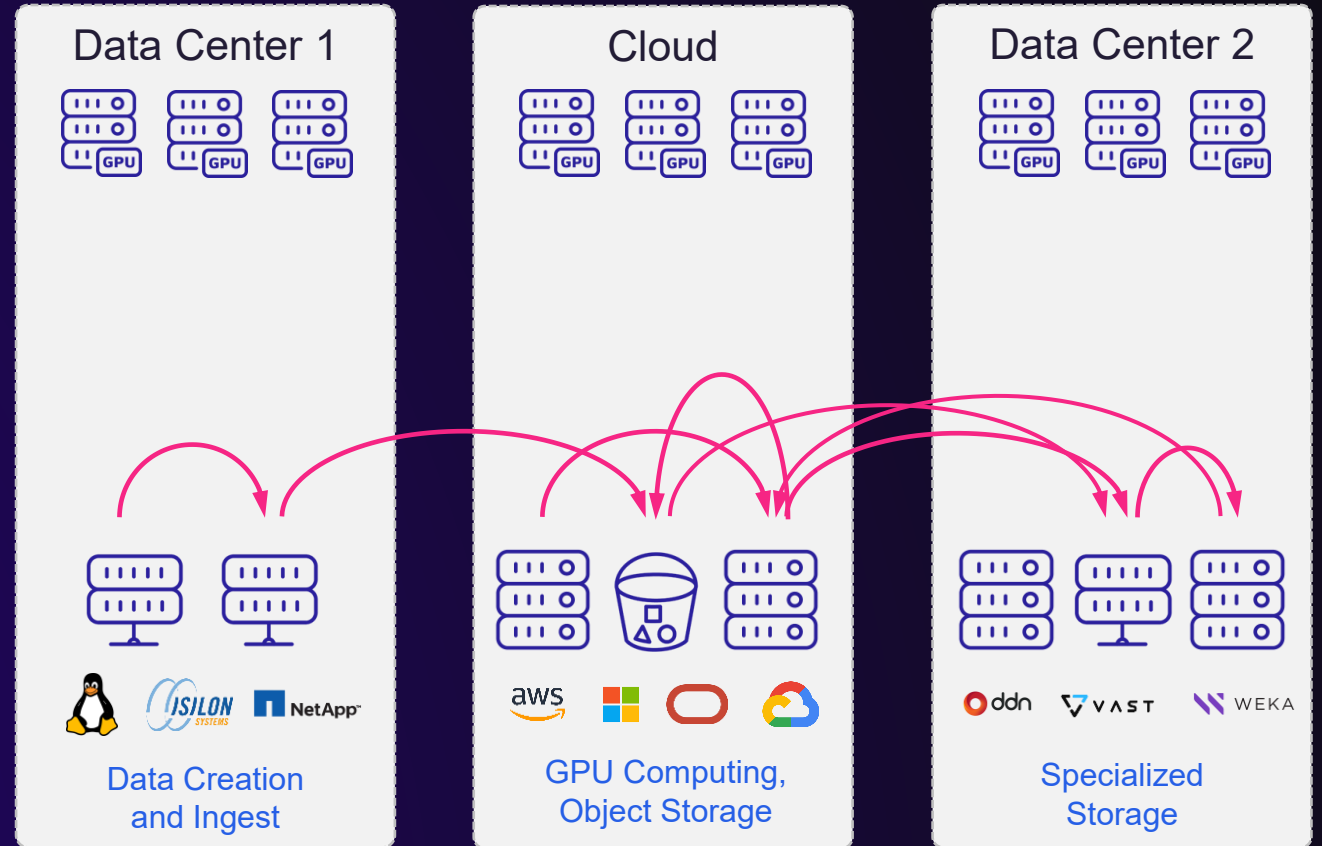
## Data Gravity

Need to bring data to compute

## Performance

Legacy storage not built for efficient speed

- ❌ Data copy sprawl
- ❌ Manual data management
- ❌ Too many storage systems
- ❌ Cumbersome migrations
- ❌ Performance constraints
- ❌ Power constraints
- ❌ Lack cloud agility
- ❌ High costs
- ❌ Idle GPU resources
- ❌ AI project delays
- ❌ Productivity impacted
- ❌ Not AI-Ready



# From Storage-Centric to **Data-Centric**

## Unified Data Plane

Make all file and object data immediately accessible everywhere

## Automated Data Services

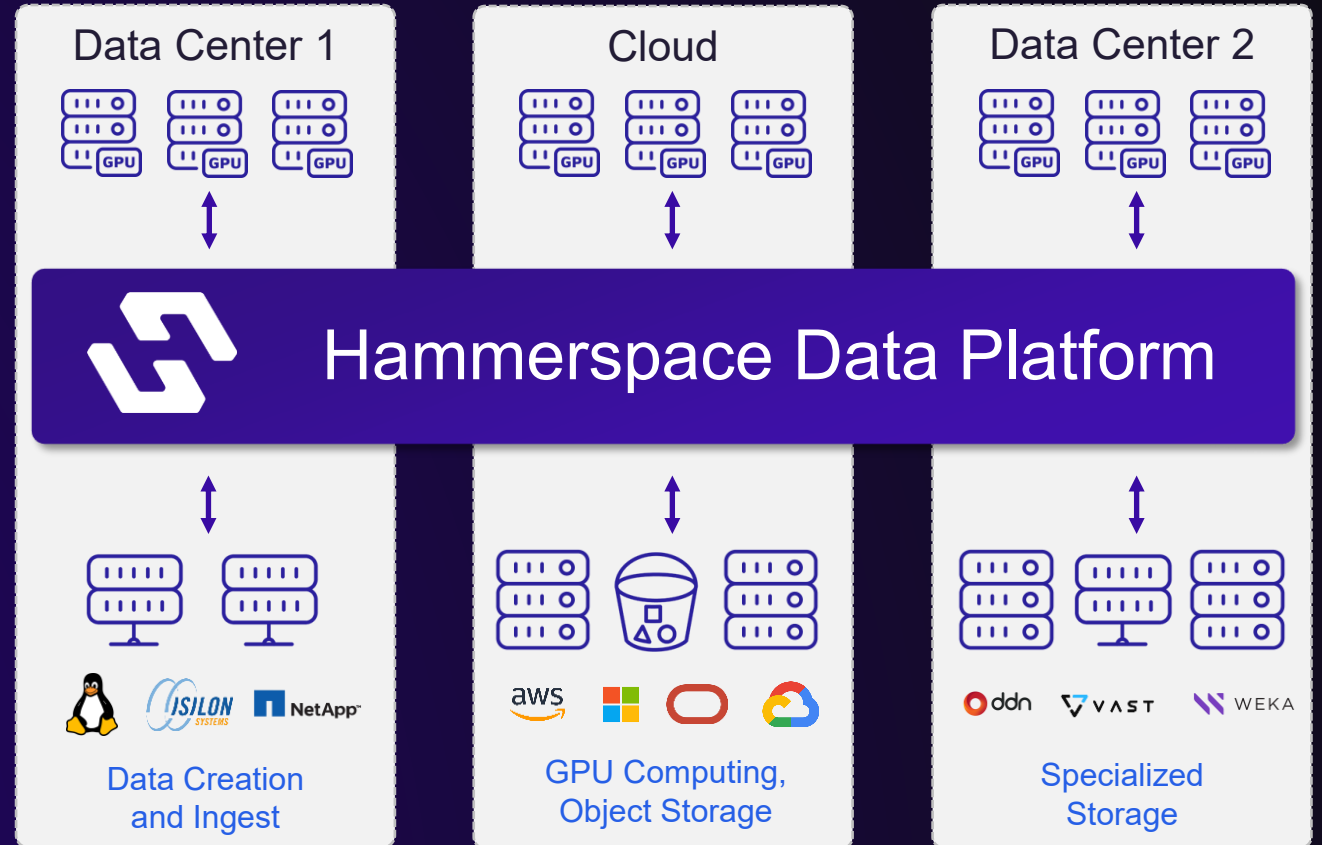
Automate data mobility, data protection, data migrations

## Accelerated Data Orchestration

Deliver unmatched performance across every aspect of the data pipeline

## Silo-Less Data Storage

Leverage data where it is, plus add extreme performance for AI workloads



# For Enterprise AI Adoption - Silos is Only the First Problem

## Equally Important is Reducing Complexity

### Speed-Bumps to Launching AI Projects

- Multiple Complex Steps
- Added Risk → Governance
- Huge upfront investment before ROI is known
- Requires Different Skillset than traditional IT



# A Turn-key, Real-Time AI Data Pipeline for Enterprise

## *Taming Data Chaos*

- Builds on Hammerspace core capabilities.
- Helps unlock their unstructured data for AI
- Removes steps for data collection, cataloging, and processing from data scattered across the enterprise
- Helps everyday enterprises start their AI journey by solving the biggest challenge to every project – AI Data readiness
- It's an agentic platform for AI Adoption



## HAMMERSPACE AI Data Platform

Use Data In Place

Orchestrated Data Movement

New Data Detection

Automated Transformation

Fast Data Access

Governance, Sovereignty, and Security

# AI-Ready Data in Days, Not Months

# The AI Data Platform for Accelerating Outcomes

*Fast times to token are good, but time to VERY First token is essential*

## Traditional AI Solutions:




## Hammerspace AIDP Solution:



**Over 10x Reduction** in time spent compared to traditional AI Storage Solutions

# Simplifies Transforms Existing Siloed Data into Unified AI-Ready Intelligence



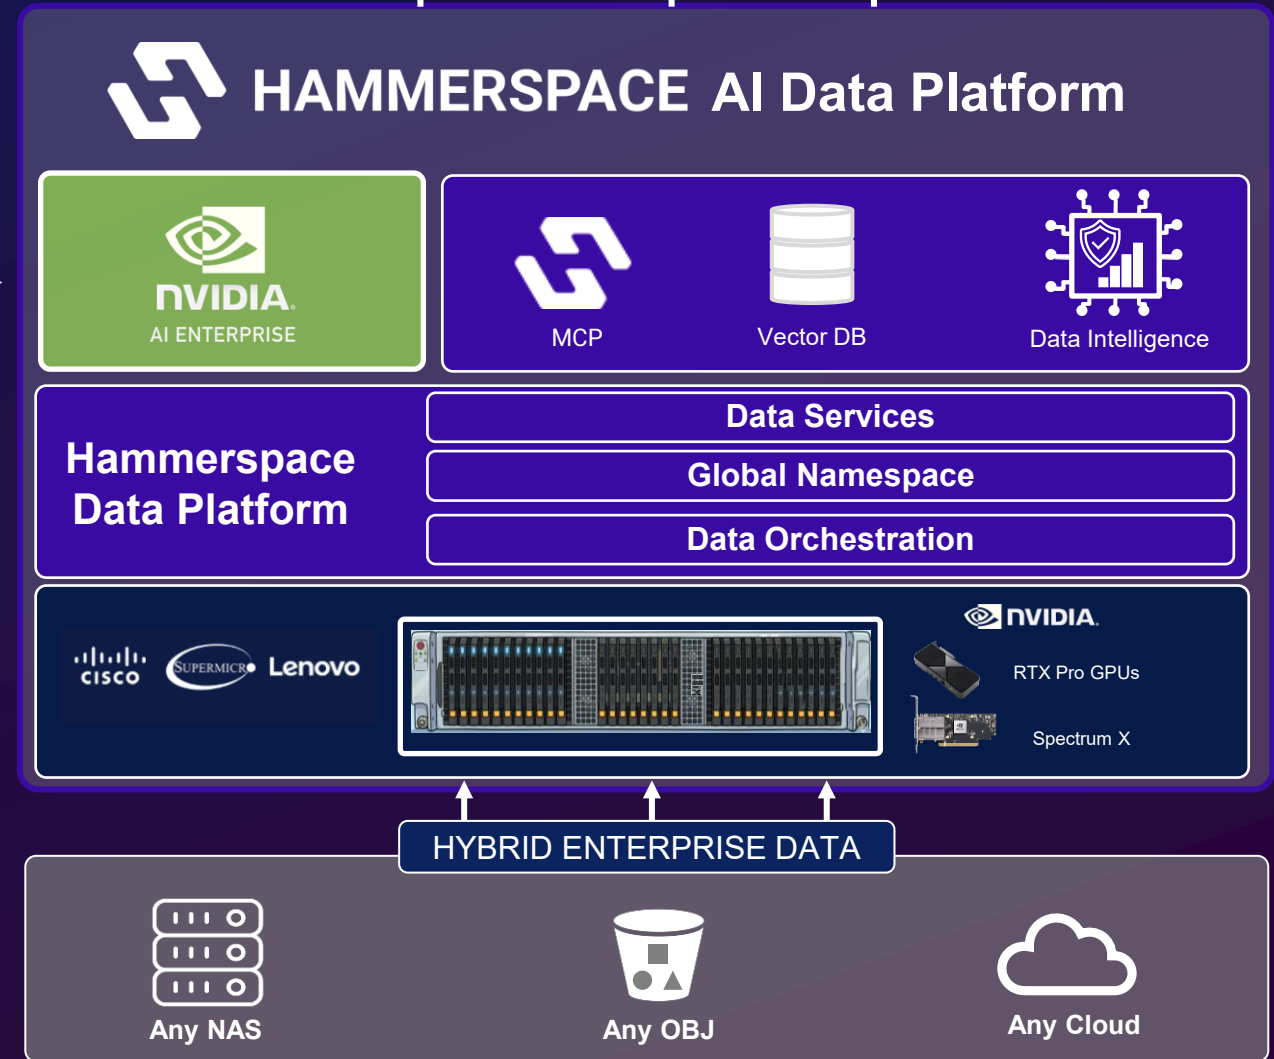
## Adds Key AI-Specific Features:

- Hammerspace MCP Server
- Vector DB & Data Intelligence Software
- Includes NVIDIA AI Enterprise Software


Built Upon Hammerspace Core Solution

## Fully Integrated With Hardware:

- Cisco, Supermicro, Lenovo
- Supports Any Third-Party Hardware



# AIDP Plain Language Query Interface



**System Status**  
● System health  
System health

**AI Readiness**  
● NIM, Milvus, NV-Ingest OK  
NIM Nemo Embedding, Milvus, NV-Ingest

**MCP Status**  
● Pod running, ● UI up  
Pod running, UI

**Vector DB**  
● 10 collections, 11.4k embeddings  
Collections and embeddings

**Hammerspace Status**  
● Anvil reachable  
Anvil reachable, files in transit

**AIDP Console**  
Intelligent testing powered by NVIDIA NIM

[Diagrams & Metrics](#) [File Discovery](#) [Debug](#) [Events](#) [Logs](#)

Describe your planning step:

e.g. 'Discover files in /mnt/hub, place on DSX, check alignment' or 'Secuvy BOM workflow: Product Launch Update'

[Run Agent](#) [Clear](#) [Engage Debug](#)

**EXAMPLE PLANNING STEPS**

- Discovery flow: dropped files → DSX
- Discovery flow: dropped files → DSX → alignment
- Secuvy BOM workflow: Product Launch Update (Secuvy-0001) → AIDP Cluster 1 → internal\_prod vector DB
- Dataset test-0021 → AIDP Cluster 1 → internal\_prod (tag collectionid=test-0021)

# Hammerspace is the Data Fuel of AI



## US GOVERNMENT



Department of  
**DEFENSE**



Department of  
**ENERGY**



## AT-SCALE AI OPERATORS

 **Meta**

**zoom**



## NEOCLOUDS



US Government  
**NEOCLOUD**

 **Cirrascale®**

# Thank You!

[www.HAMMERSPACE.com](http://www.HAMMERSPACE.com)