ACADEMIC RESEARCH BRIEFING

Modern Data Storage for AI & HPC Research

Flash volatility, GPU pipeline economics, and the hyperscaler playbook for university HPC and AI clusters.

PREPARED FOR

Academic Researchers, HPC Directors, and Research Computing Leads

RMACC HPC Symposium · May 2026

AGENDA

Today's discussion in four acts.

What university HPC and AI teams need to know about flash, fleets, and pipelines.

01 **The Flash Crisis**

Q1-Q2 2026 was the worst stretch for SSD pricing on record. What changed and why it matters for research budgets.

02 **GPU Pipeline Economics**

Why one storage tier cannot serve ingest, training, checkpointing, fine-tuning, and inference.

03 **Mixed Fleet, the Hyperscaler Answer**

How Google, Meta, and Microsoft already run storage — and how VDURA HYDRA brings it to academic clusters.

04 **Proof Points from Research Institutions**

Goethe University. ZUSE Institute Berlin. Santa Clara University.

VDURA AT A GLANCE

Modern data storage for AI & HPC.

Built on 25 years of parallel file system DNA — born in academic research.

1,000+	2.7 TB/s	12 9s	6+ 9s	60%+
Production deployments across 50+ countries	All-flash throughput per rack	Durability with multi-level erasure coding	Availability in production	Lower TCO vs. all-flash architectures

HERITAGE

Founded 1999 as Panasas. Co-invented the parallel file system and pNFS. Deployed at NASA, RTX, Boeing, Airbus, and dozens of universities. Rebranded VDURA in 2024 with software-defined pivot. Exabytes managed.

LEADERSHIP

Ken Claffey (CEO, ex-Seagate, scaled ClusterStor to 40% of top supercomputers). Garth Gibson (CTAIO, co-invented RAID at Berkeley, former CEO Vector Institute). Eddie White (Board, Google Cloud). Craig Bernero (Board, ex-EMC/Dell).

RESEARCH COMMUNITY

Trusted across the research enterprise.

Universities, national labs, and research institutes managing exabytes of irreplaceable data.

Universities & Higher Ed

Goethe Frankfurt, Santa Clara, and more across the U.S., EMEA, and APAC.

National & Federal Research Labs

NASA, NOAA, U.S. Department of Defense, and tier-1 federal program partners.

Independent Research Institutes

ZUSE Institute Berlin, Mass General Brigham, GSI, and global biomedical centers.

Industrial R&D

Boeing, Airbus, RTX, and energy-sector exploration teams.

Exabytes upon exabytes managed. Tens of millions of cumulative runtime hours. Decades of research-grade trust.

ACT ONE

The Flash Crisis

Everyone is buying GPUs. Almost no one priced in what feeds them — and the cost just snapped.

01 / 04

THE MACRO SETUP

Storage was a 10% line item. Not anymore.

AI spending is on an order-of-magnitude trajectory — storage is taking a much bigger slice.

$35B → **$236B**

Neocloud / GPU cloud market, 2024

Projected Neocloud size by 2031

Storage is no longer peripheral to AI infrastructure. It is foundational — and academic budgets feel it first.

Source: Industry research; AI Journal, Ken Claffey, March 2026.

Power eclipsed silicon

AI racks pull 40–250 kW vs. 10–15 kW for conventional research compute.

Storage share is climbing

Trending to 20–30%+ of build cost in all-flash designs — procurement cycles can't keep pace.

GPU economics are storage-gated

5,000 GPUs at 98% availability = 876,000 GPU-hours of lost research compute per year.

NVMe flash prices broke loose in Q4 2025. They have not come back.

Quarterly price for a 30 TB enterprise drive — SSD vs HDD.

TLC SSD vs HDD	QLC SSD vs HDD	SSD price Y/Y	HDD price Y/Y
26.6x	**21.3x**	**+568%**	**+55%**
Q2 '26 — 30 TB drives	Q2 '26 — 30 TB drives	Q2 '25 → Q2 '26	Same window — stable curve

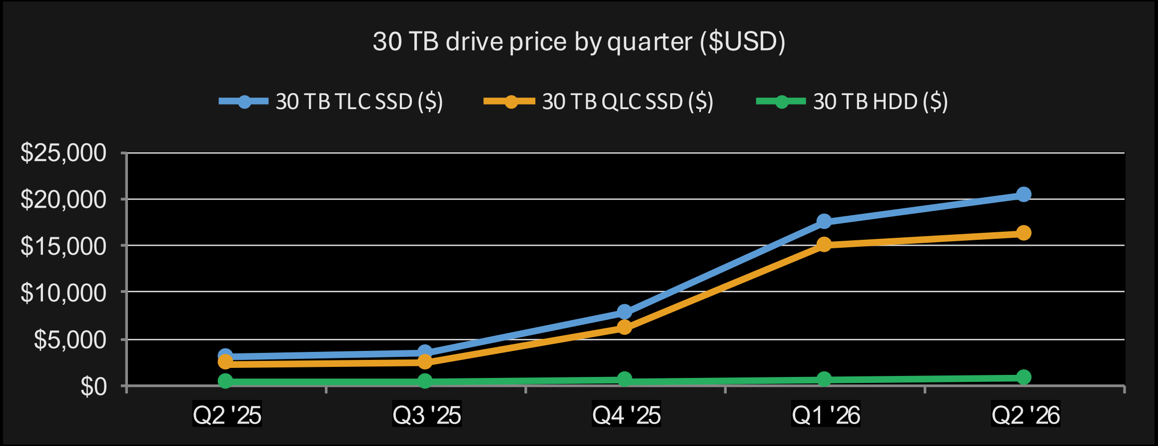

RATIO BY QUARTER

Quarter	TLC vs HDD	QLC vs HDD
Q2 '25	6.2x	4.9x
Q3 '25	7.0x	5.2x
Q4 '25	13.4x	10.7x
Q1 '26	26.2x	22.6x
Q2 '26	26.6x	21.3x

SO WHAT: Flash and HDD diverged in Q4 2025 — and the gap keeps widening. **Mixed-fleet architectures get better every quarter, all-flash bets get more painful.**

Source: VDURA Storage Economics Optimizer (vdura.com), live data through April 2026.

INDEPENDENT VALIDATION

Forbes is telling the story.

An independent industry analyst named the gap — in print, in April 2026.

INNOVATION > ENTERPRISE TECH

SSD Storage Capacity Prices Are Over 20 Times HDD Storage Capacity Prices

By Thomas Coughlin, Contributor. ⓘ Covering Digital Storage Technology & ...

Follow Author

Published Apr 16, 2026 at 04:01pm EDT, Updated Apr 19, 2026 at 10:26am EDT

FORBES · INNOVATION · ENTERPRISE TECH
APRIL 16, 2026

Industry press has been talking about this for months. Mainstream press has now picked up the story.

WHY IT MATTERS

>20x

SSD vs HDD per-TB prices, per Forbes (Apr 16, 2026).

+25%

Additional SSD price increase expected in Q2 2026.

+40%

Everpure (Pure Storage) list-price hike announced April 2026.

BUDGETS UNDER PRESSURE

What flash volatility means for a research budget.

All-flash purchasing assumes a stable cost curve. The market just removed that assumption.

ALL-FLASH BUYER

A 25% NAND swing breaks the budget.

— ~90% of an all-flash bill of materials is the SSD itself

— The vendor doesn't make the NAND — they pass the swing through

— Stalled procurements, deferred research, missed grant windows

VDURA MIXED-FLEET BUYER

A 25% NAND swing shifts the ratio.

— Performance on flash, capacity on HDD economics

— No rebuy of the namespace when the NAND market moves

— Predictable $/GB across multi-year research procurements

"Hyperscalers — Google, Meta, Microsoft, AWS — never bought all-flash. The mixed-fleet playbook is the answer." — Ken Claffey, VDURA blog, April 24, 2026

ACT TWO

GPU Pipeline Economics

Sizing flash for the workload, not the brochure. How research clusters keep GPUs fed without overpaying.

02 / 04

HOW MUCH STORAGE DO YOUR GPUs ACTUALLY NEED?

NVIDIA specifies storage throughput per GPU. Workload determines the tier.

OPTION 1 — NVIDIA ENHANCED

Text-Based AI / LLM Training

0.5 GB/s read / 0.25 GB/s write

per GPU — lower throughput requirement, less flash needed to saturate compute.

OPTION 2 — NVIDIA VISUALIZATION

Physical AI / Omniverse / Video

4 GB/s read / 2 GB/s write

per GPU — 8x the read throughput. Flash sizing and architecture efficiency become critical.

"How do I keep storage at ~10% of my budget while still saturating every GPU and have enough capacity?"

~10% of AI infrastructure budget on all-flash a year ago → **20–30%** in 2026 — unsustainable for most budgets.

Source: SemiAnalysis; VDURA workload analysis.

How much flash? Same 25 PB system, four configurations.

Dial performance and $/TB without changing software, namespace, or operations overhead.

All Flash 100% NVMe	Performance 50% SSD / 50% HDD	Cost-Optimized 20% SSD / 80% HDD	Deep Capacity 2% SSD / 98% HDD
6.8x relative price	**4.0x** relative price	**1.9x** relative price	**1.0x** relative price
Throughput **2,405 GB/s**	Throughput **1,147 GB/s**	Throughput **670 GB/s**	Throughput **210 GB/s**
IOPS **78 M**	IOPS **42 M**	IOPS **21.6 M**	IOPS **7.2 M**
Power **54.7 kW**	Power **35.3 kW**	Power **26.0 kW**	Power **14.0 kW**

Each system: 25 PB capacity. Throughput, IOPS, and power scale with the flash ratio. Three-year subscription pricing reference.

Buy SSDs for performance. Use HDDs for capacity.

Flash capacity = cluster checkpoint size × number of checkpoints retained. Nothing more.

RECOMMENDED

3.84 TB SSD (TLC)

$2,500

Same GB/s

Buy for performance

Same GB/s as larger drives at <25% of the cost.

MODERATE

30 TB SSD

$12,000

Same GB/s

Diminishing returns

Same GB/s, 4.8x the cost. Pay only if rack space drives you here.

AVOID

120 TB SSD (QLC)

$47,000

Same GB/s

Capacity trap

Same GB/s, almost 19x the cost. Pure premium media for cold data.

Key insight · A 3.84 TB TLC SSD delivers ~the same GB/s as a 30 TB SSD for under 25% of the price.

ACT THREE

Mixed Fleet,
the Hyperscaler Answer.

What Google, Meta, and Microsoft already do — finally available to academic AI clusters.

03 / 04

VDURA HYDRA Architecture

UNIFIED DATA PLATFORM FOR AI & HPC WORKLOADS

CLIENT WORKLOADS

HPC	**Instruments**	**GPU**	**Container**	**TPU / VM**
x86	& Sensors	ARM	Orchestration	Cloud

DIRECTFLOW CLIENT LIBRARY — POSIX-COMPLIANT PARALLEL ACCESS

Data & Metadata Separated
Independent scaling of control and data planes

File-Level Network Erasure Coding
Multi-level N+2 · up to 12 nines durability

PARALLEL ACCESS · LINEAR SCALABILITY

CONTROL PLANE

Metadata Services

VeLO	**VeLO**	**VeLO**
Director 1	Director 2	Director 3

DATA PLANE

Storage Services

VPOD	**VPOD**	**VPOD**	**VPOD**	**VPOD**	**VPOD**
Node 1	Node 2	Node 3	Node 4	Node 5	Node 6

COMMODITY HARDWARE · SHARED-NOTHING ARCHITECTURE

Single-port NVMe SSDs (TLC/QLC) · Single-port SATA HDDs (CMR/SMR/HAMR) · Standard servers

AI PIPELINE COVERAGE

Every workload asks for something different.

From ingest to inference, one namespace and the right tier for each stage.

INGEST	TRAIN	CHECKPOINT	FINE-TUNE	INFERENCE	ARCHIVE
Sequential writes Massive volume	High-BW reads + writes	Burst writes Low latency	Mixed I/O Small files	Random reads Ultra-low latency	Cold sequential Density-driven
HDD-tiered	**NVMe flash**	**NVMe flash**	**VeLO + flash**	**VeLO + flash**	**HDD**

THE INFERENCE INVERSION

Inference is no longer the small cousin of training. Read-heavy concurrency, multi-model serving, and KV-cache traffic punish slow metadata. If storage cannot serve billions of inode operations per second, GPUs sit idle — and so does the science.

ACT FOUR

Proof from

Research Institutions

Three deployments, three architectures, one platform. Goethe University. ZUSE Institute Berlin. Santa Clara University.

04 / 04

Goethe University Frankfurt.

20 PB V5000 deployment for Europe's largest AMD GPU cluster.

Prof. Volker Lindenstruth and the VDURA team during system acceptance, March 2025.

ENVIRONMENT

- AI physics applications
- Largest AMD GPU cluster in Europe
- NVIDIA NDR200 networking

DEPLOYMENT

- Phase 1: 20 PB, >100 GB/s (delivered)
- ~2 PB flash + 18 PB HDD
- Phase 2: >100 PB (2026)

WHY VDURA

"After a thorough evaluation, we selected VDURA's Data Platform and V5000 for its superior price/performance and durability — which all-flash solutions couldn't meet."

— Prof. Volker Lindenstruth

Goethe University Frankfurt

ZUSE Institute Berlin.

Tier-1 German research institute selected V5000 for an AI supercluster.

ENVIRONMENT

- AI physics applications
- Largest NVIDIA GPU cluster in Berlin
- NVIDIA NDR200 networking

PHASED ROLLOUT

- Phase 1: 12 PB, >100 GB/s (delivered)
- ~1 PB flash + 11 PB HDD
- Phase 2: >30 PB (2026)
- Phase 3: 100 PB by end of 2026

KEY SUCCESS FACTORS

- AI engine complementing IBM Storage Scale
- True parallel FS uniting flash speed with HDD economics
- EB scalability — perf, capacity, metadata scale linearly
- End-to-end encryption + multi-tenancy across Berlin institutes

"We deployed the new system within 3 hours — one of the most incredible milestones we gathered."

— Carsten Schäuble, Head of IT & Data Services, Zuse Institute

Santa Clara University.

Accelerating AI-powered research — traffic-incident captioning, stream-flow prediction.

ARCHITECTURE

DIRECTOR NODES

ALL-NVMe FLASH NODES

Sub-millisecond latency · rapid model iteration

HDD CAPACITY EXPANSION

One namespace · start all-flash, scale flexibly

RESEARCH WORKLOADS

- Traffic incident captioning (Woven by Toyota & NVIDIA)
- AI-driven stream-flow prediction for Santa Clara Valley Water District
- Supermicro GPU servers · NVIDIA HGX A100 8-way · 400 Gb NDR InfiniBand

KEY SUCCESS FACTORS

- Sub-millisecond latency for rapid model iteration
- Only 0.5 FTE to manage — stays constant as the system scales
- Built-in encryption aligns with campus security standards

THE GROWTH NARRATIVE

Start all-flash. Scale flexibly.

Add more flash for peak performance, or expand bulk capacity seamlessly — in one namespace, with no architectural rework.

Linear flash expansion enables performance and capacity to grow on demand as research data scales from hundreds of TB to multi-PB.

Walkthrough by Prof. David C. Anastasiu

DISCUSSION

What we'd love your perspective on.

Your experience as researchers and HPC leads will sharpen ours.

Q1 **How is flash volatility affecting your multi-year research procurements?**

Most 5-year research-cluster plans were built on a smooth NAND cost curve. Has that broken in your environment too?

Q2 **Where in your AI/HPC pipeline is storage the bottleneck right now?**

Ingest at the instrument? Checkpointing during long-running training? Inference for shared lab tooling?

Q3 **What would mixed-fleet storage need to do to earn its place next to (or in front of) your existing parallel file system?**

We're curious where and why other storage solutions win in academic environments — and where they don't.