VAST

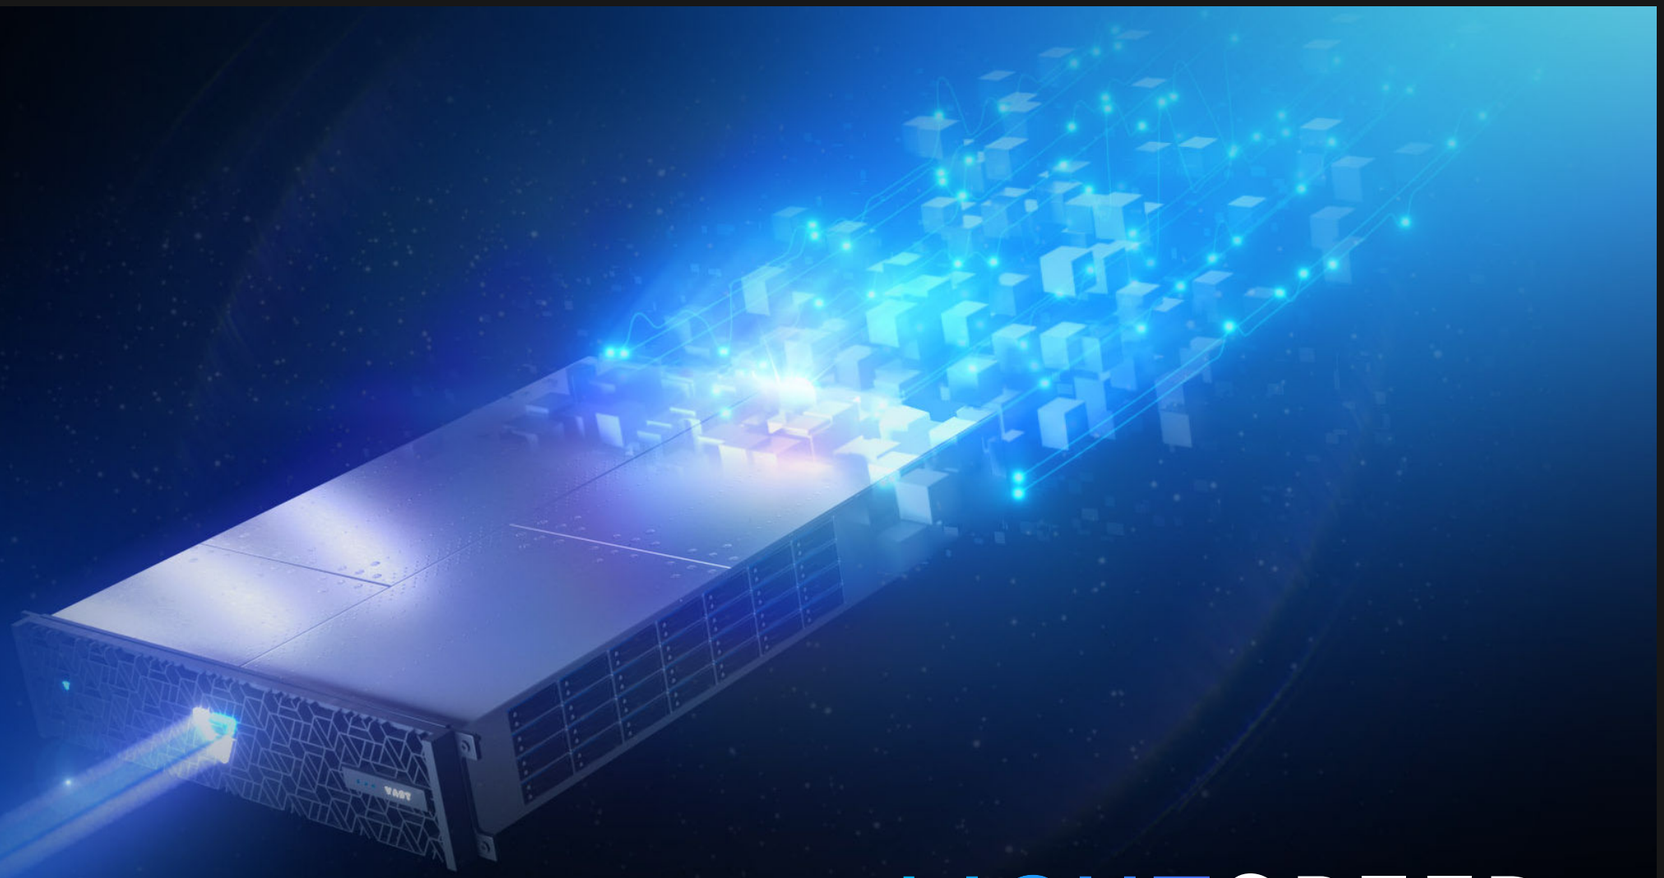

LIGHTSPEED

THE STORAGE PLATFORM FOR THE NEXT DECADE OF MACHINE INTELLIGENCE.

TABLE OF CONTENTS.

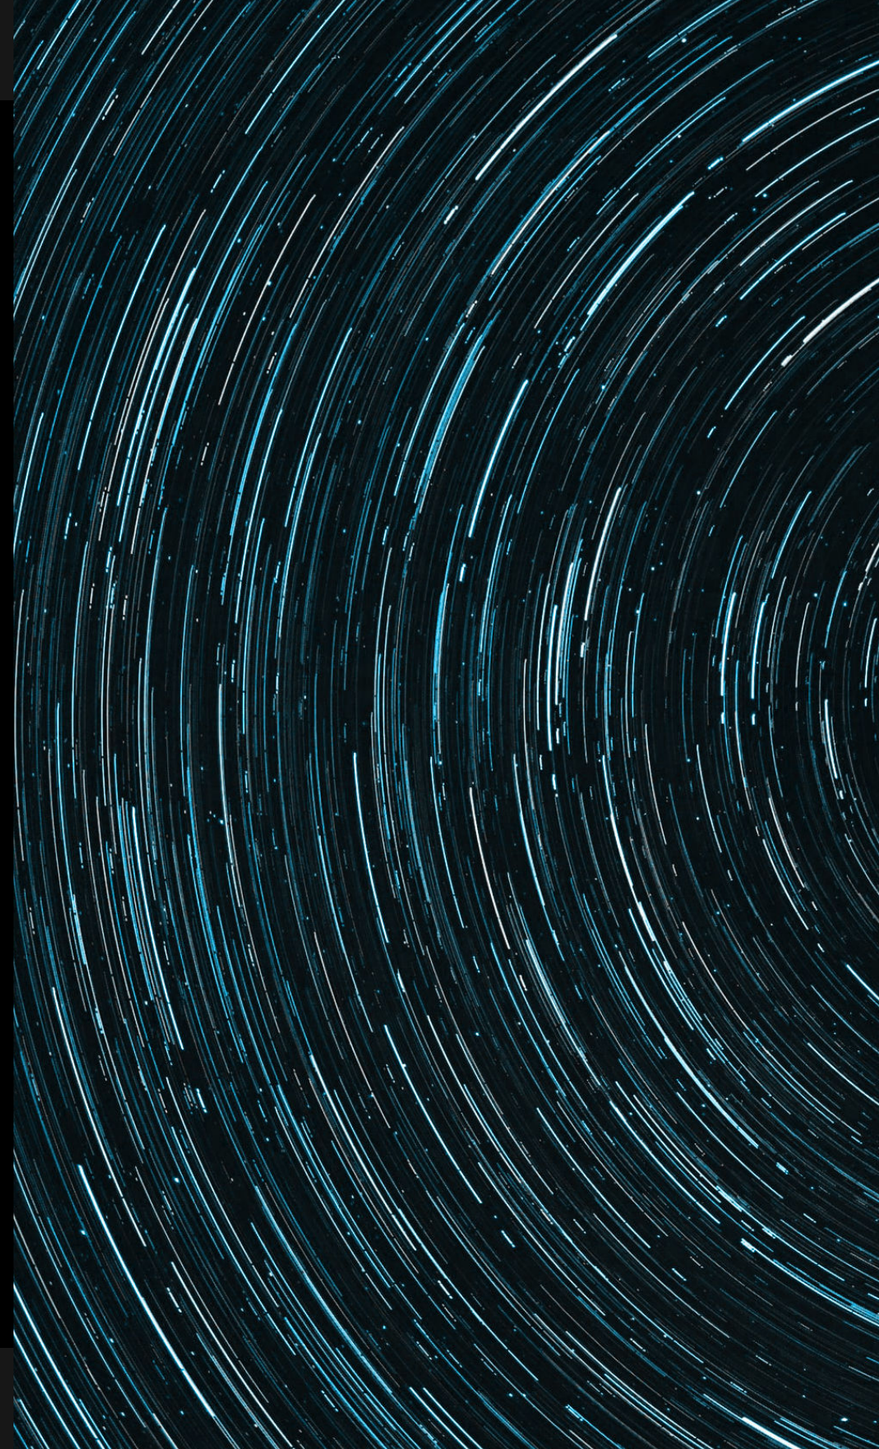

2020 AI INFRASTRUCTURE CONSIDERATIONS.

As organizations race to develop new AI-based product and service offerings, this new class of computing interacts with data in ways never-before encountered. This raises critical infrastructure planning issues that will challenge how storage has been built for decades.

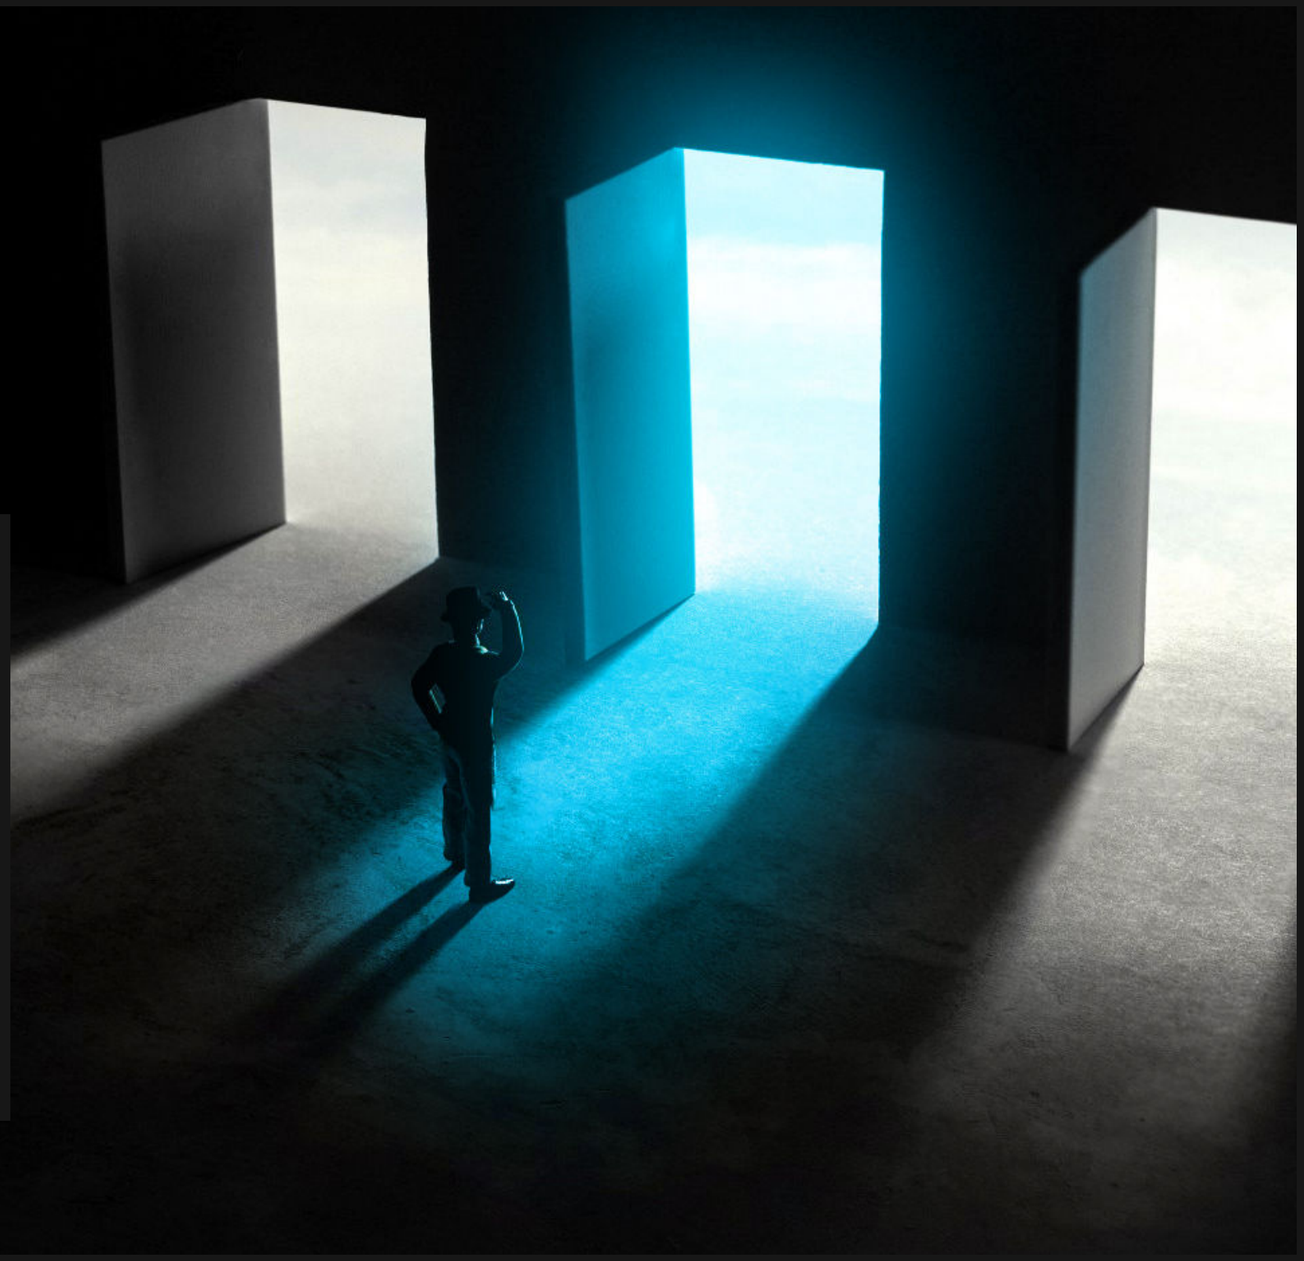

STORAGE TIERING DOESN'T WORK FOR AI.

AI is all about random I/O... and you can't prefetch a random read.

Hard drive accesses can result in 50x slower I/O for AI applications.

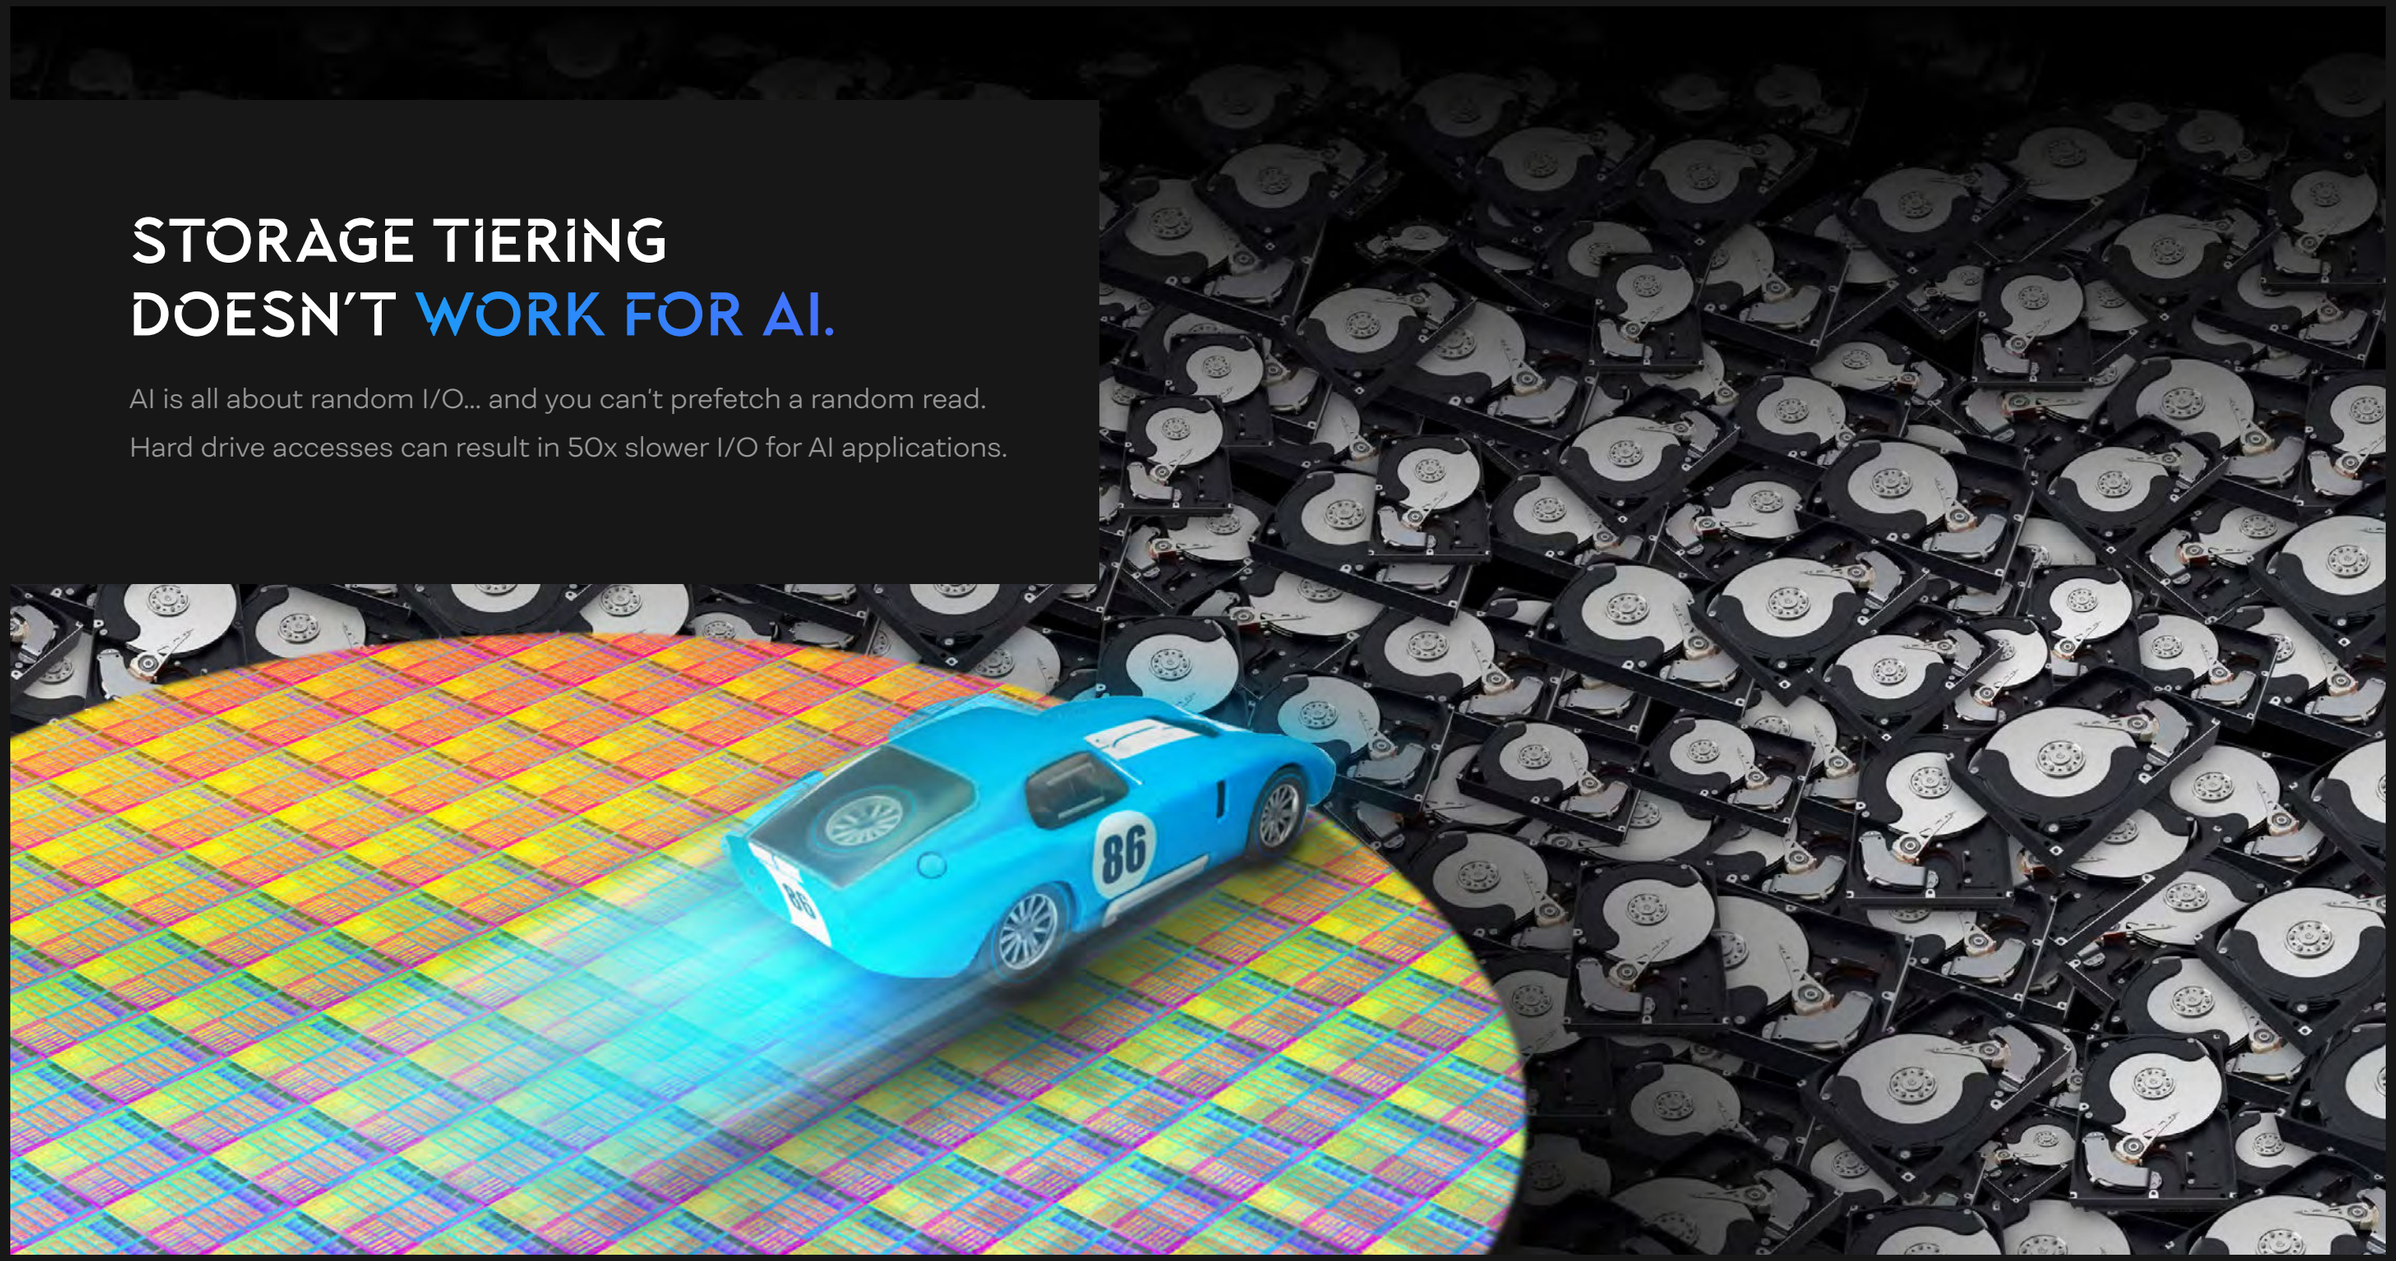

AI WORKLOADS DEMAND FLASH PERFORMANCE.

Industry leaders are now establishing flash as a critical requirement to successful AI workload deployments.

> "Storage systems, such as those that support Lustre and burst buffers, are not designed for and often perform poorly for these read-heavy, random access workloads"

AI for Science Report
US Department of Energy

Go to the report

> "The newer NVMe SSDs and storage-class memory can ameliorate the low-bandwidth and high-latency challenges associated with neural networks, while the quad-level cell (QLC) drives offer flexibility to balance costs relative to sheer volume of data to be stored"

Prepare Your Storage and Data Management Strategy for the Impact of Artificial Intelligence Workloads
Gartner

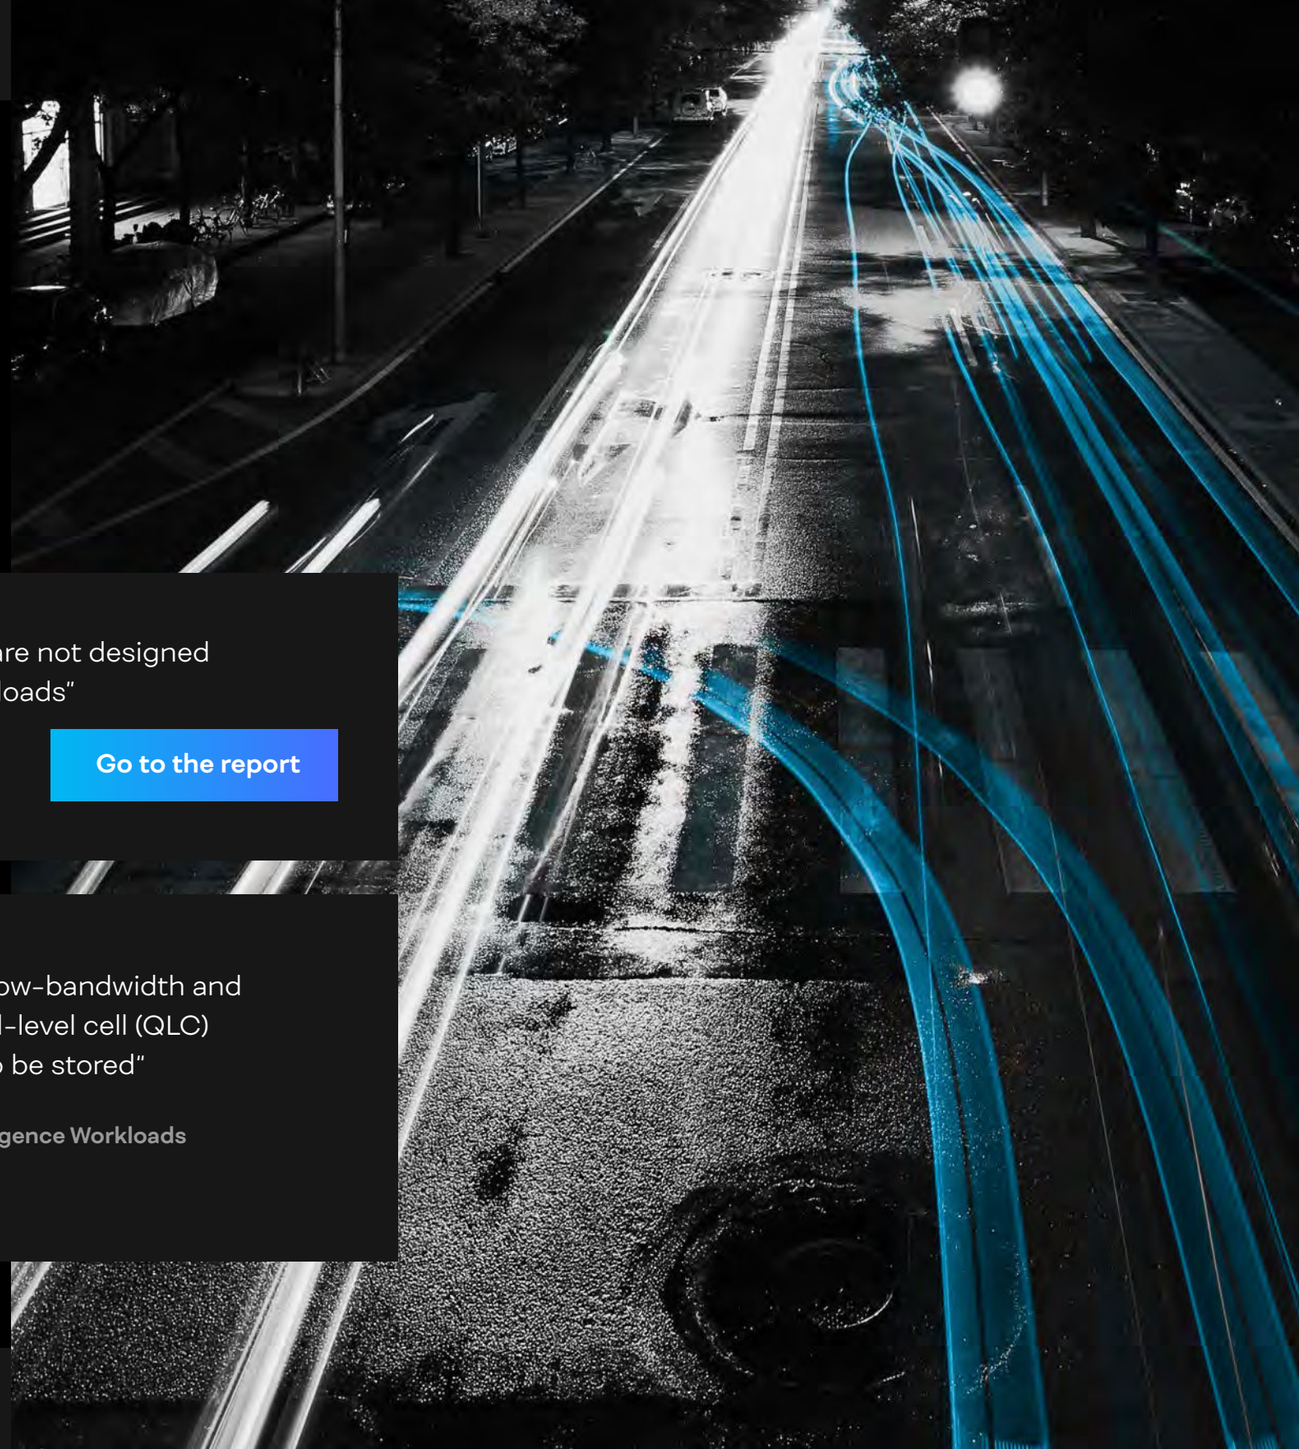

DECADES OF STORAGE COMPROMISE IS AT ODDS WITH THE AI IMPERATIVE.

Storage teams have wrestled with the tradeoff between cost, performance and simplicity for decades. What's come from this series of compromises is a complex, tiered storage hierarchy that is not suitable for the dynamic and evolving needs of data science organizations.

UNLEASH THE DATA.

Machine learning and deep learning applications are built upon the principle that we now have accelerated computers and vast reserves of data to actually train large neural networks. The classic tiered storage model of the last 30 years cannot keep up with data scientists' need to train and retrain on any piece of data at a moment's notice. AI data accesses happen randomly and at very small increments, making the concept of an 'archive' obsolete.

UNLEASH THE GPUS.

GPUs pack more than 150x more performance into an AI computer as compared to a high-end CPU-based server. This compute concentration creates an equal concentration of I/O, where AI systems can now ingest as much as 200GB/s of training and inference data per system. Classic NAS systems have historically peaked performance at 2GB/s, the maximum possible from a single TCP stream. This misalignment of capabilities has forced IT organizations into more exotic solutions.

UNLEASH THE ORGANIZATION.

As AI moves out of R&D to be operationalized by the enterprise, IT organizations have been challenged to take ownership of AI infrastructure because the file and object storage protocols that have been use, until now, have been exotic and foreign to general IT practitioners. Simple tools are required if AI is going to be moved out of the shadows and deployed under the oversight of global IT organizations.

HPC FILESYSTEMS ARE COMPLICATED.

For years, HPC centers have avoided NAS and deployed proprietary file system drivers to provide their applications the streaming bandwidth and scalability they need. These systems specialize in performance and scale, but can require PhDs to deploy and maintain. While they are fast for classic HPC simulation workloads, HPC file systems have been difficult for AI organizations to adopt for many reasons:

- These products were designed for streaming and not the IOPS and metadata load that AI applications exhibit
- These systems deploy a custom kernel driver on a file system client, which creates software dependencies and can make kernel upgrades very problematic
- HPC storage systems have never received the mass market adoption that NAS systems have enjoyed, and they see less polish and product hardening than NAS appliances, making them fragile and difficult to maintain
- These systems have limited suitability outside HPC, which makes it difficult to put an end to shadow AI initiatives

> "Upgrading the OS kernel is very challenging, because it implies upgrading our parallel file system... which is not unlike a suicide pact."
>
> **- Former Lustre User / Now VAST Customer**


“THE STORAGE ARCHITECTURE
OF THE FUTURE.
IDC, 2020
INTRODUCING
UNIVERSAL STORAGE.
Designed to democratize the use of flash for all data, VAST's Universal Storage was born in the age of machine intelligence to deliver a level of ease of use, scalability and performance to power the intelligent applications of the next decade.

PIONEERING THE NEW STORAGE PARADIGM.

VAST Data's Universal Storage defies conventional storage categorization to break long-standing tradeoffs and combine capabilities never-before found in a single storage platform.

- Tier-1 NVMe Performance
- Tier-5 Archive Economics
- Low-Latency, High-Bandwidth NAS & Object Storage
- Exabyte-Scale Storage Cluster

For AI environments, Universal Storage eliminates the barriers to flash adoption and ensures that training data is available in real time, all with a simple NAS experience that delivers best-in-class GPU performance.

INTRODUCING THE VAST DATA DASE ARCHITECTURE.

DISAGGREGATED, SHARED EVERYTHING

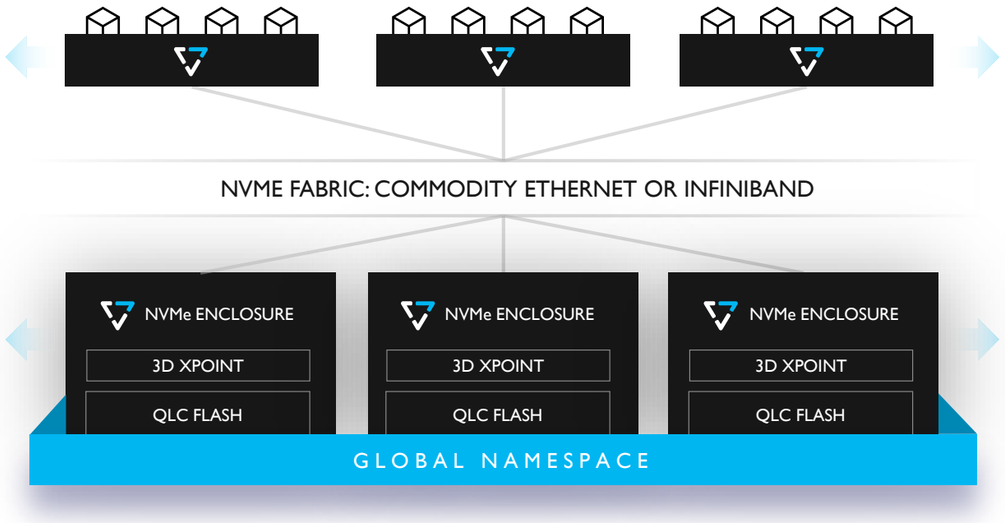

Stateless containers power VAST's system, these high-availability processes serve file and object protocols from the multi-protocol namespace. The global, shared-everything, view of the NVMe media allows VAST containers to scale massively and execute new global algorithms that enhance the efficiency of the cluster and make flash affordable for all data.

New low-latency storage networking helps build scalable clusters over commodity Ethernet or InfiniBand fabrics.

The state of the VAST cluster is stored in highly-available, high-density storage enclosures. These NVMe 'JBOFs' pack 675TB of raw flash capacity in just 2U of space. Cluster data is stored in a pool of low-cost QLC flash and metadata is stored in new 3D XPoint persistent memory. Scale to up to 1,000 enclosures in a single cluster for exabytes of capacity when factoring in the data reduction benefit of VAST's global compression.

VAST DELIVERS NEW STORAGE INNOVATION TO DEMOCRATIZE FLASH FOR ALL DATA.

NEW QLC FLASH SAVES 85% VS. LEGACY FLASH NAND

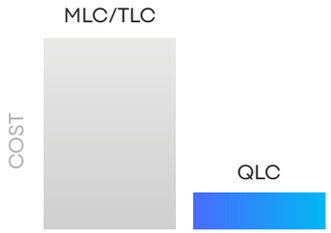

10-YEAR QLC WARRANTY

NEW ERASURE CODES SAVE 30% WHILE INCREASING RESILIENCE

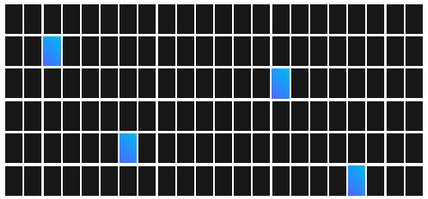

ONLY 2.7% OVERHEAD

NEW GLOBAL COMPRESSION TOOLS TO REDUCE FILE/OBJECT DATA

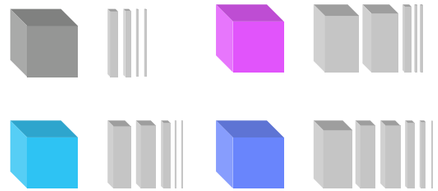

2:1 DATA REDUCTION (AVERAGE)

VAST Data has pioneered three innovations to deliver solutions as low as 1/10th the cost of legacy all-flash storage systems, making it possible to deploy your entire training dataset on flash.

Learn More About VAST Economics

ALL NEW FOR 2020:
LIGHTSPEED

VAST combines a new hardware offering, blueprints for scalability and forthcoming support for accelerated GPU computing to establish Universal Storage as the platform of choice for intelligent applications.

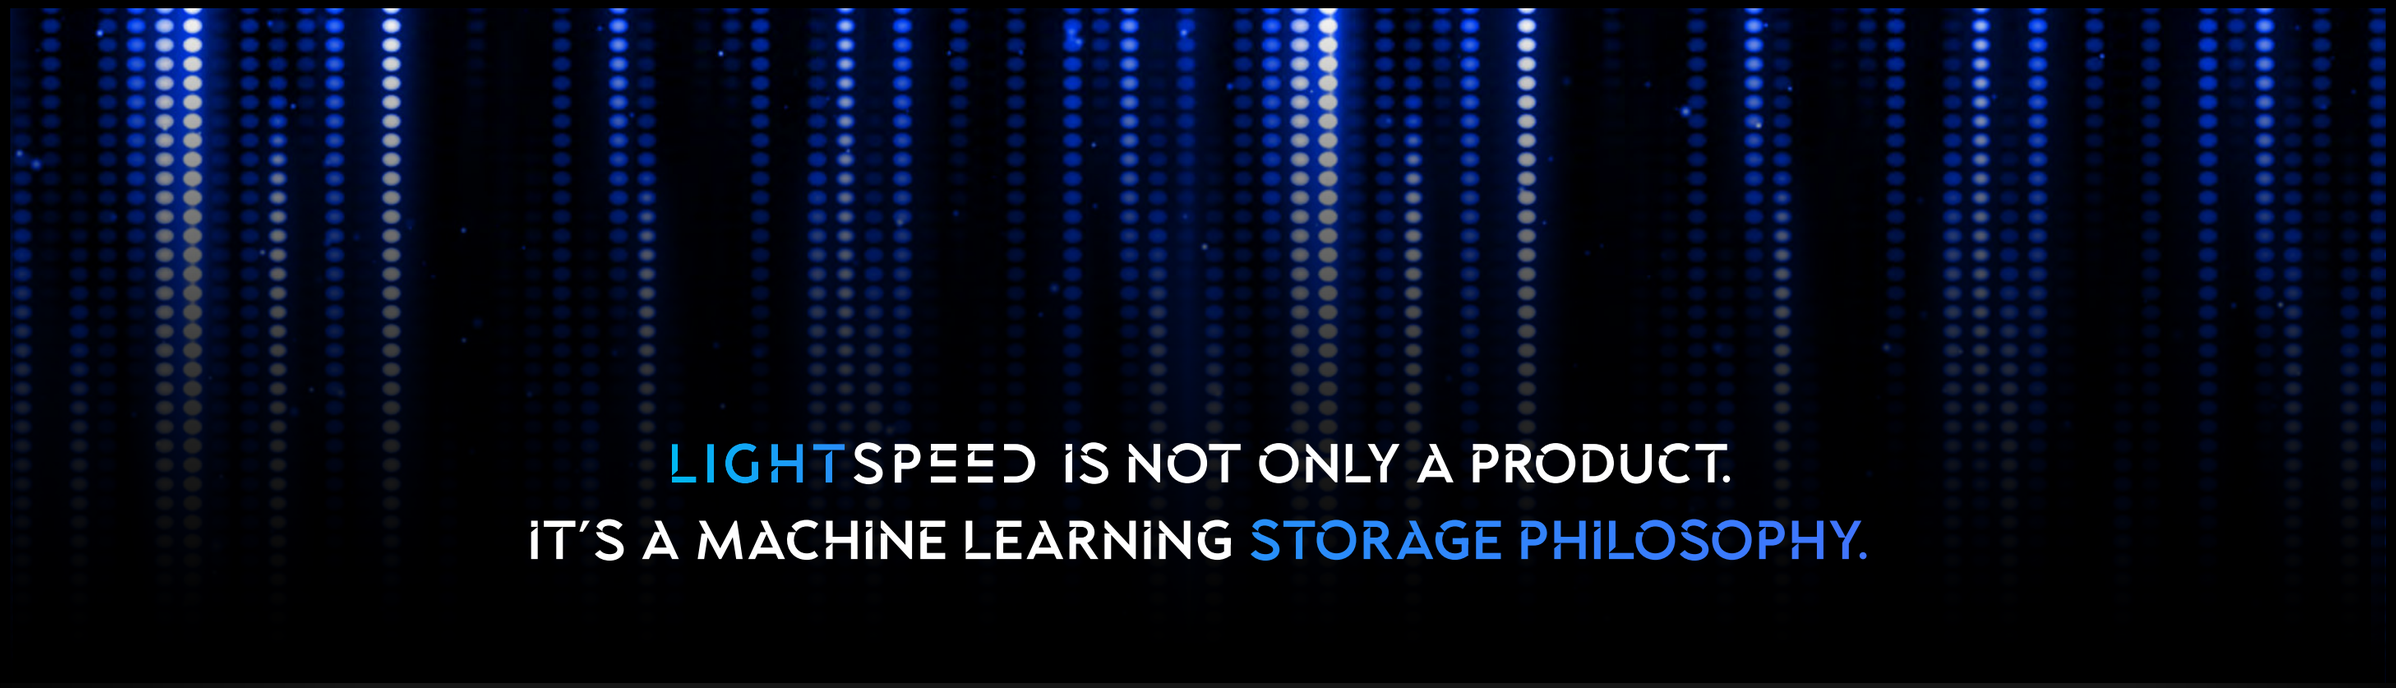

THE LIGHT-TOUCH SIMPLICITY OF NAS + THE SPEED TO SUPPORT ANY AI AMBITION

Coupled with the scale & economics of VAST's award-winning DASE™ architecture

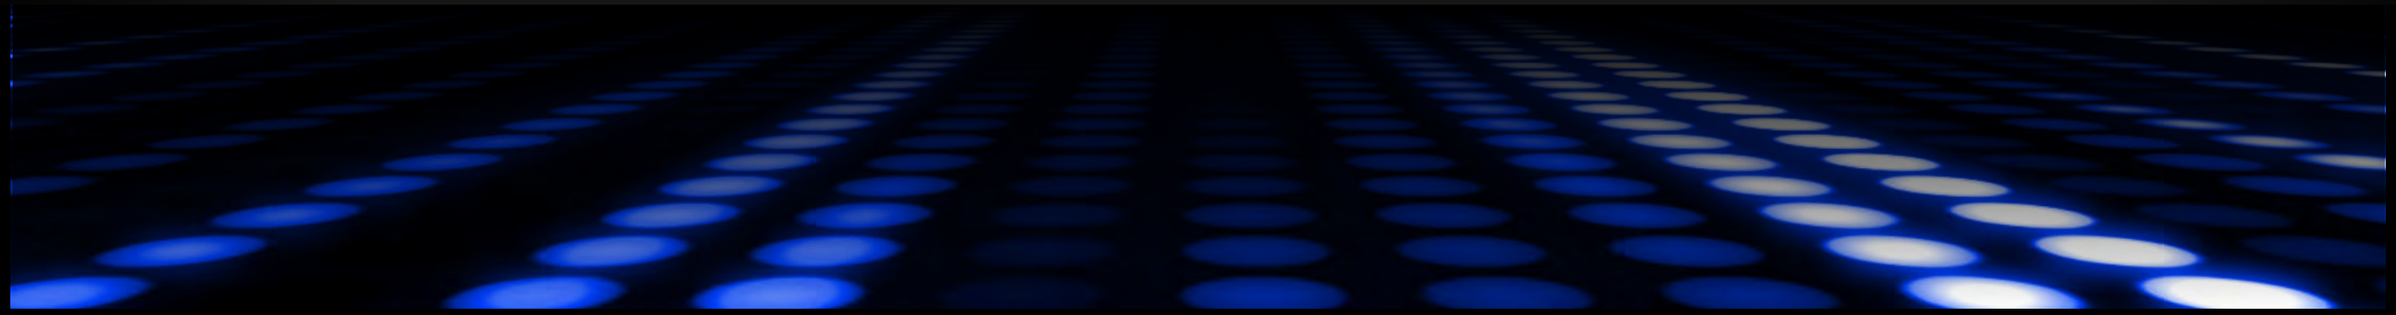

NEW LIGHTSPEED NVME ENCLOSURES

New LightSpeed enclosures deliver 2x the bandwidth of previous-generation systems. With this added performance, Universal Storage clusters can deliver 2x the AI workload I/O throughput.

LightSpeed enclosures work in Ethernet or InfiniBand fabrics to tightly integrate into any high-bandwidth AI cluster environment.

These 2U machines feature leading storage capacity (over 1PB effective at 2:1 data reduction) and the IOPS of over 3,000 HDDs.

2X
BANDWIDTH

Up to 40GB/s
Per Enclosure

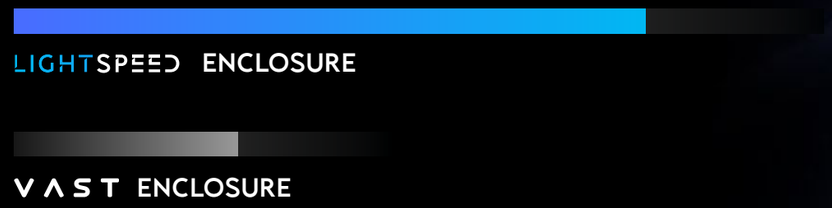

DGX-A100 BENCHMARKING.

TAKING NFS TO THE NEXT LEVEL

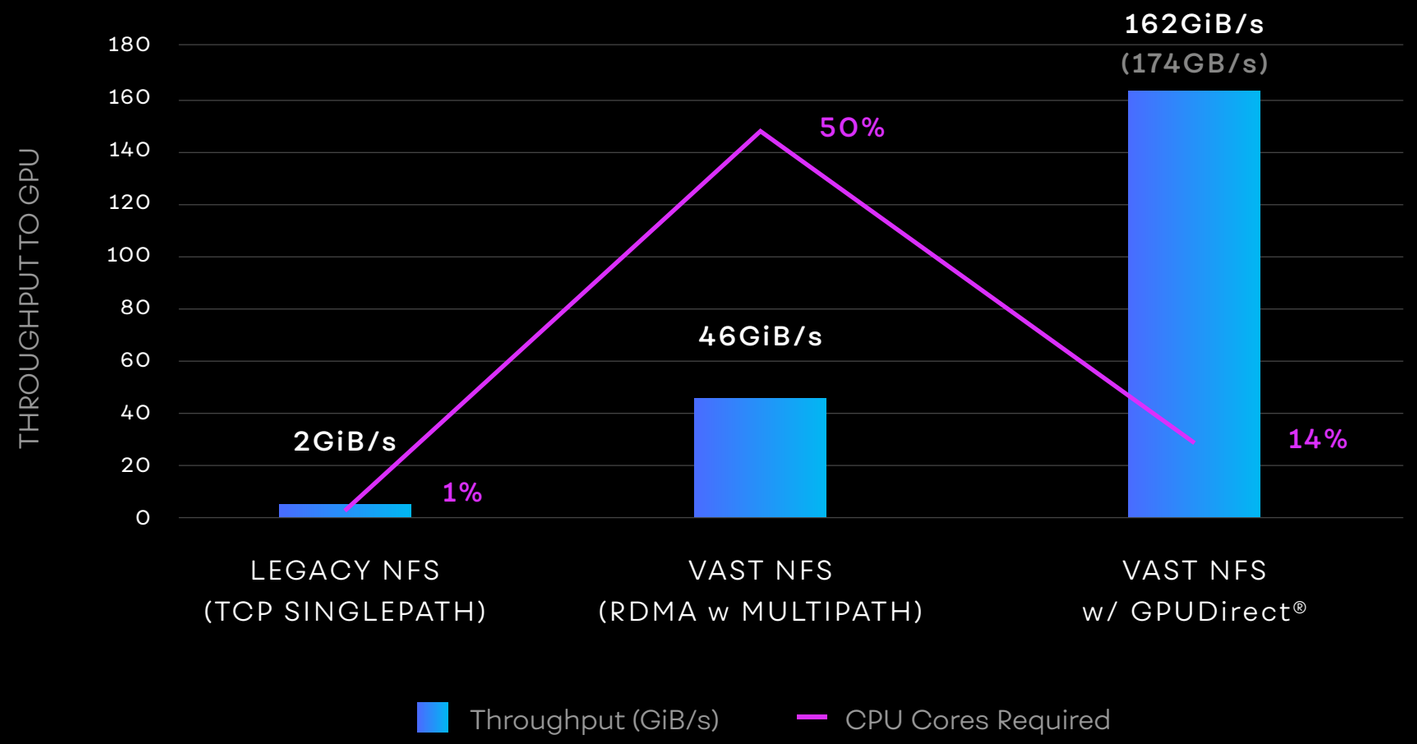

SATURATE YOUR GPUS

FREE YOUR X86 CORES

NFS = SIMPLE AI

TESTED WITH GDSIO • 1 x DGX-A100 • 5 x VAST CBOX SERVER CHASSIS & 5 x VAST LIGHTSPEED ENCLOSURES • 4MB I/O SIZE • 4GB FILE SIZE • 96 THREADS x 8GPUs

ACCELERATED SUPPORT FOR GPUS.

NFS HAS NOW EVOLVED TO TAKE AI TO LIGHTSPEED

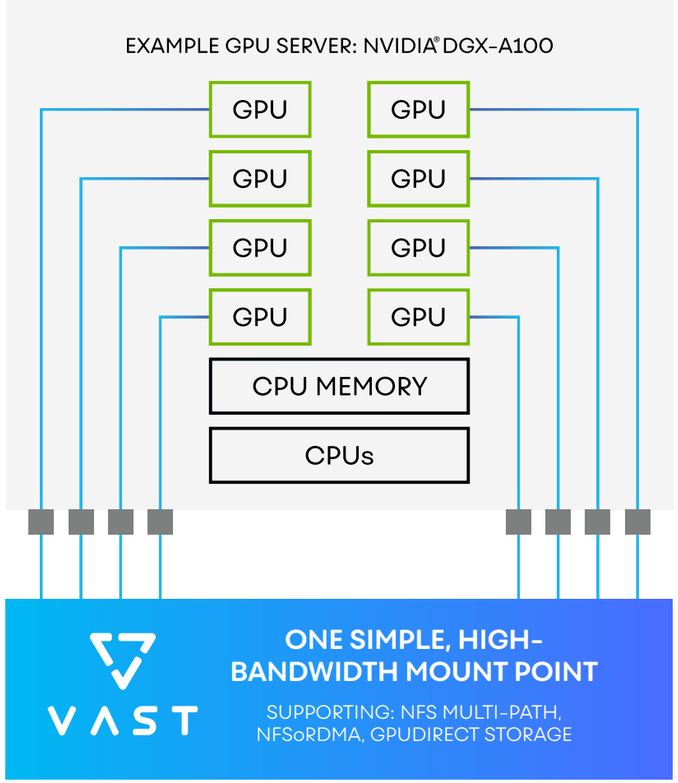

Most enterprise storage systems adhere to the standard NFS/TCP model that has been the core offering of scale-out NAS providers for decades. New AI machines need more than the 2GB/s of read throughput that a standard TCP NFS client can provide. This single-stream, non-RDMA limitation is why organizations abandon NFS in favor of parallel file systems and chose performance over simplicity.

In 2020, VAST has combined support for a number of NFS accelerations that are in the Linux kernel to make it possible to get best-in-class throughput for AI applications and GPU computers without requiring customers to choose complex parallel file systems.

- NFS over RDMA provides line-speed access over Ethernet or Infiniband
- NFS NConnect federates I/O over many streams and ports
- NVIDIA GPUDirect® enables VAST servers to directly place data in GPU memory, bypassing CPU memory bottlenecks to increase throughput, eliminate CPU overhead and lower latency for GPU I/O

VAST and NVIDIA have already demonstrated more than 88GB/s of performance using GPUDirect Storage on NVIDIA DGX-2 machines, delivering leadership-class performance with just NFS.

LIGHTSPEED CLUSTER ARCHITECTURES.

Universal Storage clusters can scale to meet the need of any AI ambition. These sample cluster architectures can help to build balanced architectures that combine the speed needed to keep AI processors working with all-flash capacity needed for deep learning.

EXAMPLE GPU CLUSTER ARCHITECTURES

LIGHTSPEED

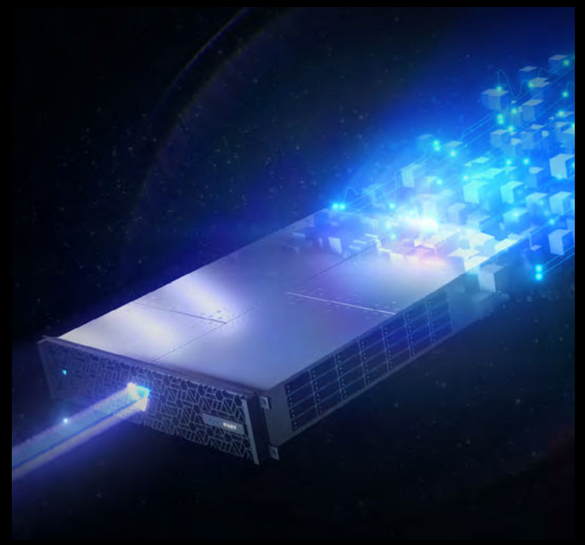

40 GB/S | 1 PB* | 400K IOPS

1x LightSpeed Enclosure
1x VAST Server Chassis
Right-sized to provide balanced I/O
for up to 16 GPU clients

PENTAGON

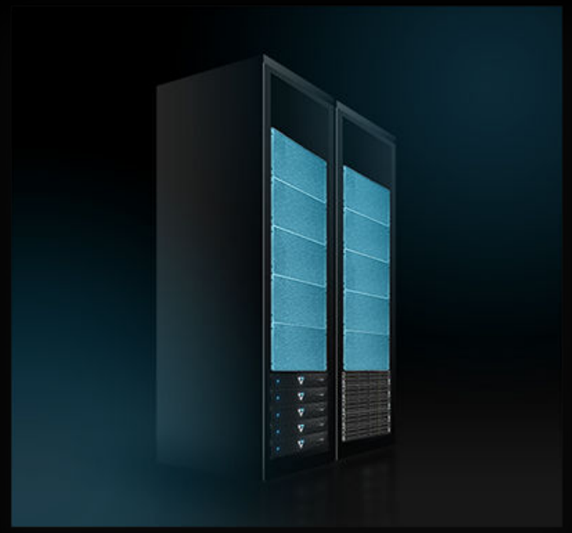

200 GB/S | 5 PB* | 2M IOPS

5x LightSpeed Enclosures
5x VAST Server Chassis
Right-sized to provide balanced I/O
for up to 80 GPU clients

DECAGON

400 GB/S | 10 PB* | 4M IOPS

10x LightSpeed Enclosure
10x VAST Server Chassis
Right-sized to provide balanced I/O
for up to 160 GPU clients

*Assumes 2:1 Data Reduction.

LINE

TWO ENCLOSURE CONFIGURATION

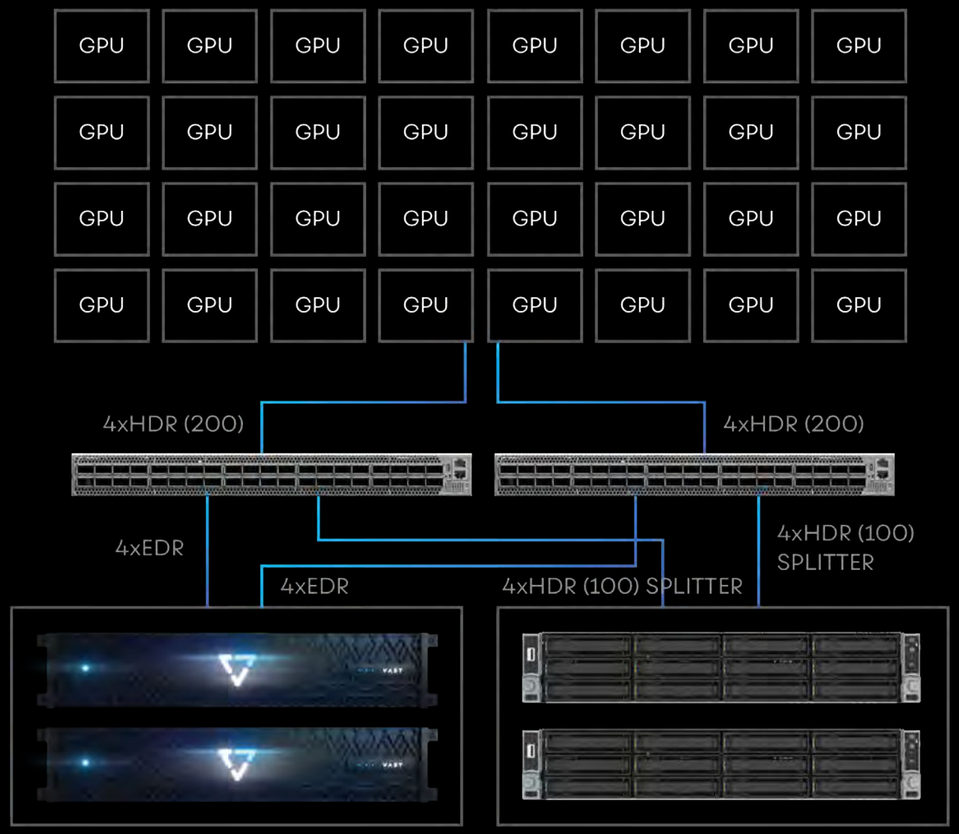

Configuration:

- 32 GPU Clients
- VAST 2x2 (2 LightSpeed Enclosures)
  - 80GB/s Read
  - 10GB/s Write
- 26RU (~31RU with gaps)

Networking Requirements:

- GPUs: 8xHDR (200)
- ISL: 4xHDR (200)
- DBOX: 8xEDR
- CBOX: 8xHDR (100) splitters

PENTAGON

FIVE ENCLOSURE CONFIGURATION

COMPUTE RACKS

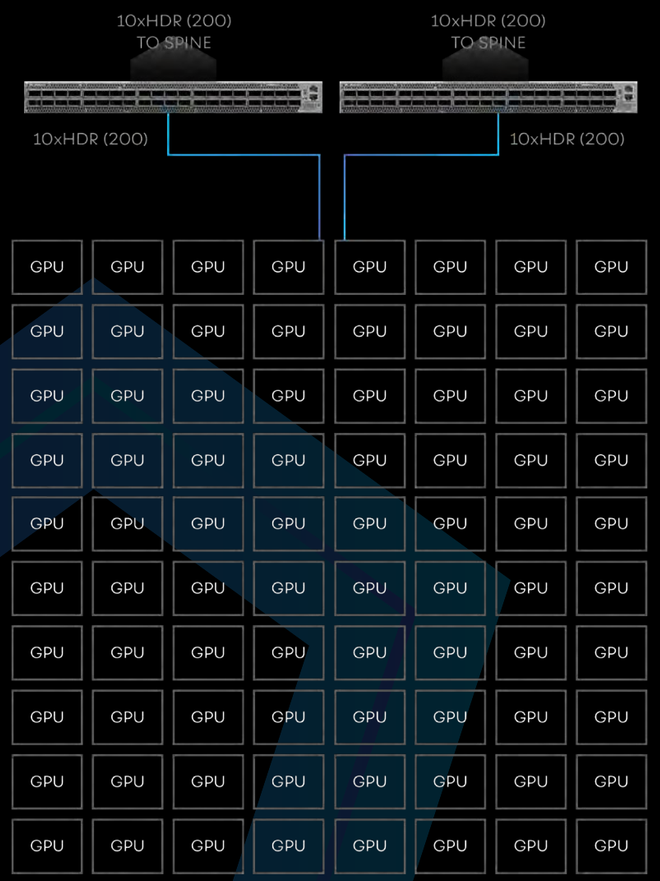

STORAGE RACKS

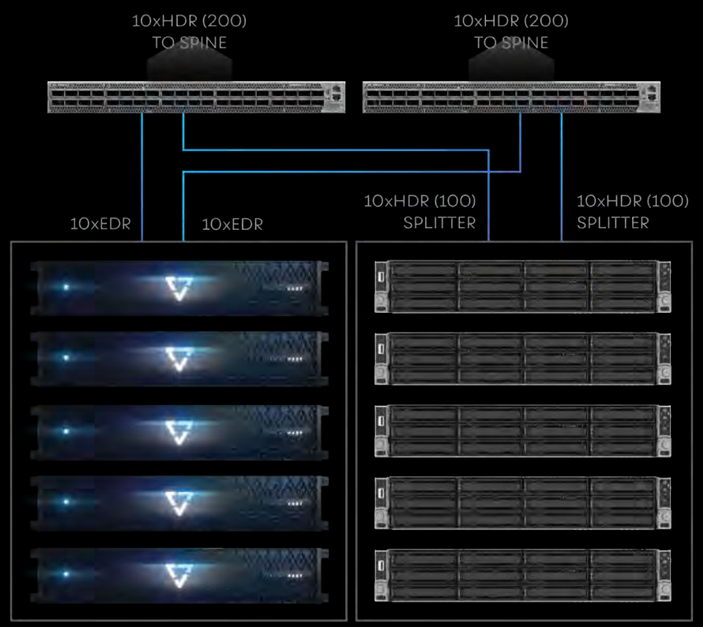

Configuration:

- 80 GPU Clients
- VAST 5x5 (5 LightSpeed Enclosures)
  - 200GB/s Read
  - 25GB/s Write
- • 60RU (~31RU with gaps)

Networking Requirements:

- GPUs: 20xHDR (200)
- UPLINKS: 40xHDR (200)
- DBOX: 20xEDR
- CBOX: 20xHDR (100) splitters

DECAGON TEN ENCLOSURE CONFIGURATION

COMPUTE RACKS

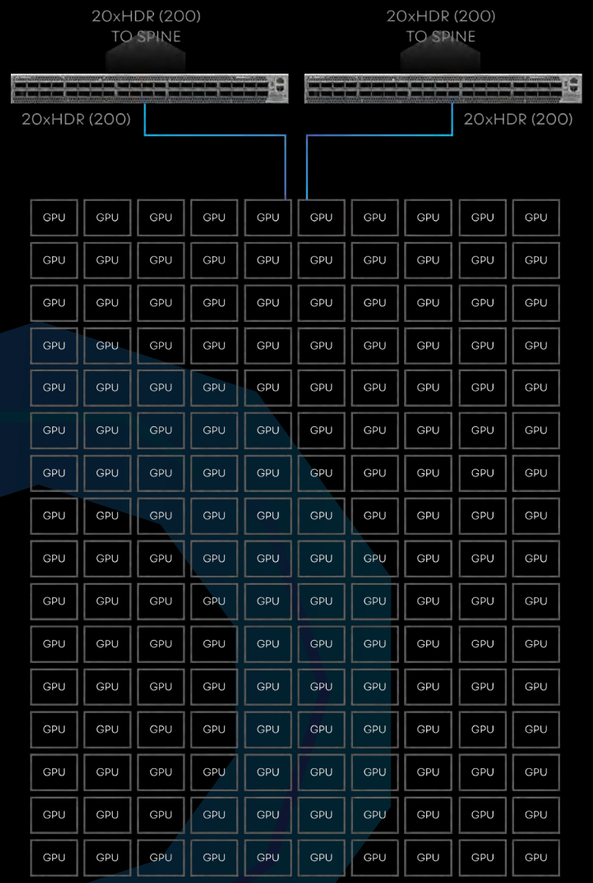

STORAGE RACKS

Configuration:

- 160 GPU Clients
- VAST 10x10 (10 LightSpeed Enclosures)
  - 400GB/s Read
  - 50GB/s Write

Networking Requirements:

- GPUs: 40xHDR (200)
- UPLINKS: 60xHDR (200)
- DBOX: 40xEDR
- CBOX: 40xHDR (100) splitters

Athinoula A. Martinos Center For Biomedical Imaging

MASSACHUSETTS GENERAL HOSPITAL

HARVARD MEDICAL SCHOOL

LIGHTSPEED SUCCESS STORY

The Athinoula A. Martinos Center for Biomedical Imaging at Massachusetts General Hospital is one of the world's premier research centers devoted to development and application of advanced biomedical imaging technologies. The Center is affiliated with both Harvard Medical School and MIT. Together, faculty and researchers are developing first-of-a-kind tools and applying them to solve challenges in neuroscience, oncology, cardiology and other clinical domains.

In 2020, the Martinos Center chose VAST Data to help accelerate and streamline their medical image research. With the rapidly growing need to capture and record every bit of MRI and PET scan data, the Martinos Center selected VAST Data's Universal Storage power four new NVIDIA DGX A100 systems and to help transition them to an all-flash platform for image-based AI.

LIGHTSPEED BENEFITS

- All-Flash For Less Than The Cost of Tiered Flash & HDD
- Storage savings that provided budget to buy more GPUs
- A simple all-flash NAS platform to eliminate AI application complexity

LIGHTSPEED

SIMPLICITY, PERFORMANCE AND SCALE TO POWER
THE NEXT DECADE OF MACHINE INTELLIGENCE

Learn more at vastdata.com

VAST