


**Architectures for  
Explosive Data Growth**  
Efficiently Handling AI, Analytics  
and Distributed Data



**DDN is the Most Efficient Data Solution for AI and HPC, driving 10x\* more data in a fraction of the energy of other systems.**



\* Not a STAC benchmark



## DDN's Modern Data Storage is Optimized for the AI Era

DDN Virtualization  
Technology **Eliminates**  
cables, switches & servers

Large Capacity QLC  
Drives **Doubles**  
Capacity per Watt

End-to-End Parallel  
Datapath **Doubles**  
Power/Performance

Workload Focused  
Performance Optimization  
**Speeds up Applications**


100% linear scale out  
Software **Removes**  
costly silos



## DDN Reduces Complexity, Cost and Risk

Fully-integrated and optimized data platforms are easy to deploy and manage at-scale!

- IB or Ethernet Connected
- Everything you need in 2RU
- Easy to Scale Linearly

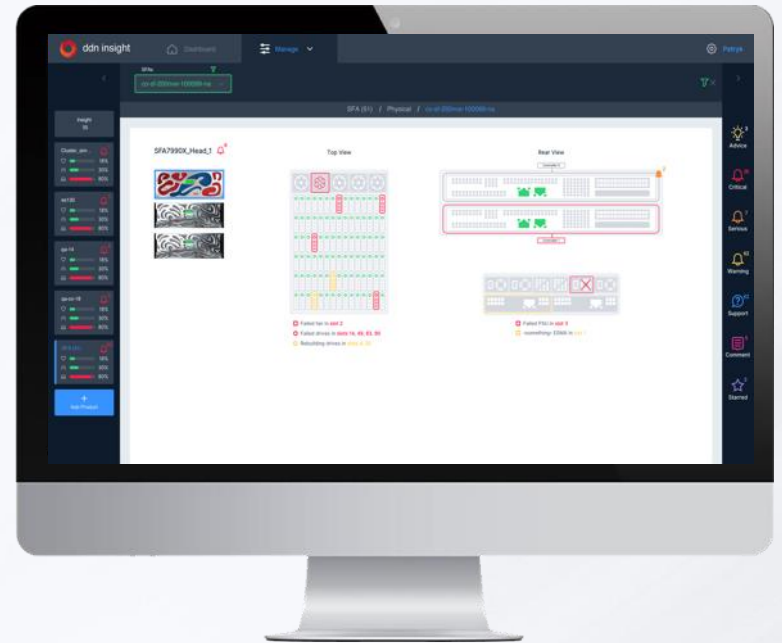


**SIMPLIFIED  
STACK**



## DDN Makes Massive Scale Easy to Manage

- Fast resolution of issues with DDN AutoSupport and advanced analytics
- Pinpoint HW and SW problems fast with DDN Scalable Monitoring and Health Management
- Workflow visibility allows Administrators to identify rogue workloads easily
- Deploy, Upgrade, Expand and Configure 100s PB with DDN SW Orchestration
- Accelerated Storage means much fewer servers and switches and cables to manage



## In 2RU and 2KW **DDN** Delivers:

**90GB/s**

*READS*

**3M**

*IOPS*

**65GB/s**

*WRITES*

Up to **720TB**


*NVMe TLC CAPACITY*



# Optimizations for AI

# GPT3 training using 128 DGX A100s and DDN storage

At-Scale NLP with DDN: Megatron-LM on NVIDIA SELENE

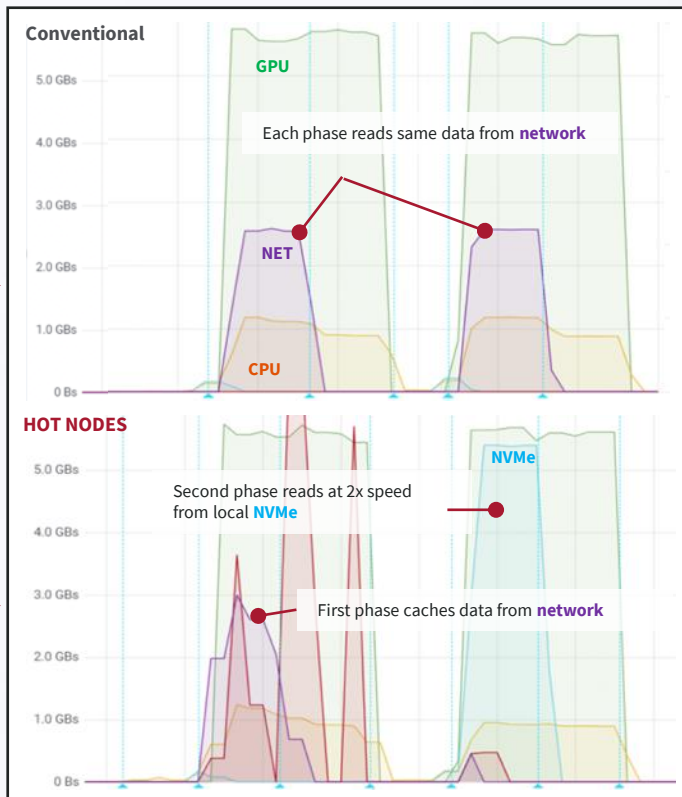


- 13B parameters in model (2020). **Today, models are 40-50X larger.**
- **Read data set at beginning of every training job:**
  - Up to **1 TB/s read** from shared DDN storage during first iteration
- **DDN Hot Nodes makes distributed training process more efficient:**
  - Cache reads in local storage, transparent to application
  - Shared storage available for checkpoints, ingest and other cluster activities



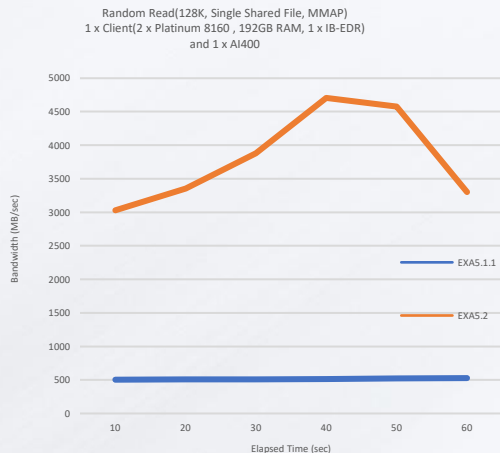
# Compute Caching for Multi-Epoch Training

- ResNet50 benchmark on DGX-A100 + AI400 without cacheing
  - Each phase reads same data from network (purple)
  - Compute runs in parallel with IO (CPU orange, GPU green)
- 
- ResNet50 with chacheing on internal NVMe devices
    - First phase also reads from network (purple)
      - Total data read volume is similar, second read from RAM
    - Computation also reads from network/RAM while files copied
    - Write to cache storage on NVMe (red)
    - Second phase reads from NVMe at double bandwidth (cyan)
    - GPU usage (green) the same, CPU usage (orange) lower
      - No network/server load on second and later runs




# How to Make Shared File Operations FAST

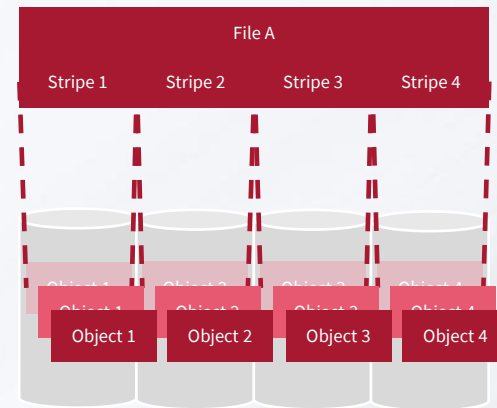
1 Let's ensure MMAP is Fast




2 Let's ensure 1,2,4 threads is fast



3 Let's scale



# DDN Delivers the highest Performing Storage Systems on the Planet




\* Not a STAC benchmark





ddn

# Parallel Filesystems can max-out the hardware even with just one file



## vs All-Flash NAS solutions was faster in 13 of 17 mean response time Antuco benchmarks



- was faster in 13 of 17 mean response time Antuco benchmarks, including:
- 6x speedup in 50-**multi-user intervalized stats** (STAC-M3.β1.50T.STATS-UI.TIME)
- 5x faster in **10-user aggregate stats** (STAC-M3.β1.10T.STATS-AGG.TIME)
- 4.9x in **single-user intervalized stats** (STAC-M3.β1.1T.STATS-UI.TIME)

