OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

May 22, 2025

Evaluating DAOS as AI-Optimized Storage

Michael J. Brim

Oak Ridge National Laboratory

U.S. DEPARTMENT of ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

Leadership-class computing is increasingly used for Scientific Machine Learning and Artificial Intelligence

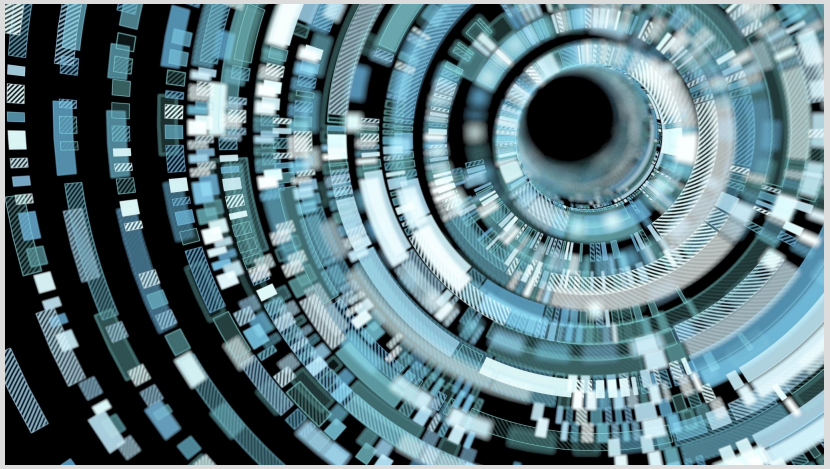

Scientific ML/AI workloads require either ***high metadata operation rates*** for datasets containing lots of small files, or ***high random read IOPS*** from large files

- Rates only grow as workloads are scaled and application demands increase

Lustre, GPFS, and other Parallel File Systems were designed for ***large sequential reads and writes***, not high-rate metadata or random reads

- Random read patterns lack temporal and spatial locality, and are unlikely to benefit from existing PFS caching strategies in clients or servers

OLCF has used node-local storage to enable scientific AI/ML

Summit and Frontier both utilized node-local NVMe storage

- NVMe provides high metadata and random read performance
- Small capacity and independent namespace per node requires additional data management work
    - Data sharding can negatively impact convergence and model quality
- Node-local performance scales with node count, limiting storage performance for smaller AI/ML jobs

Hypothesis: Isolating the AI/ML storage from the Modeling and Simulation storage may provide better performance to both workloads

- A shared AI/ML storage system eliminates usability and capacity limitations of node-local storage while still improving performance versus the PFS.
- As MTBF decreases with increased leadership-class HPC system size, prioritizing PFS I/O for mod/sim workloads is important.

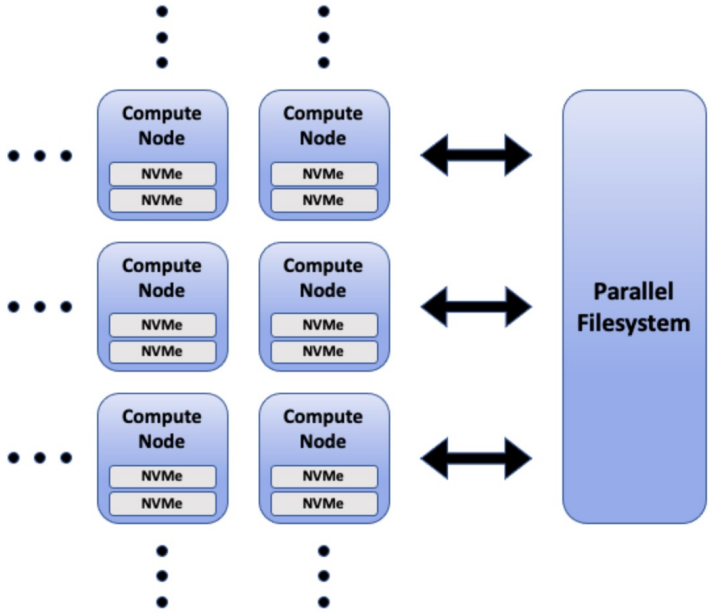

The NVMes on Frontier are local to each node.

Our AI-Optimized Storage (AOS) evaluation seeks to understand storage system performance for Scientific ML/AI

Testbed Storage Systems and HPC

400 Gbps NDR Infiniband Network

VAST Storage (Release 5.2.0-sp10-1631657)

- 8 CNodes (2 CBox), each with 100Gbps NIC
- 2 DBox, each with 8 SCM and 22 15.3 TB SSD

DAOS Storage (v2.6.2)

- 8 servers, each with dual Intel Xeon Gold 6338 (32-core), 512 GB RAM, two 200Gbps Mellanox ConnectX-7, and ten 3.8 TB NVMe
- DAOS Config: 16 engines, 320 targets, MD-on-SSD

HPC - Quokka cluster

- 16 nodes, each with dual Intel Xeon Gold 5418Y (24-core), 256 GB RAM, 200Gpbs Mellanox ConnectX-7
- Slurm 22.05, OpenMPI 4.17, DAOS CLI version 2.6.2, libdaos v2.7.0

Benchmark Software

elbencho v3.0-5 (https://github.com/breuner/elbencho)

- distributed (non-MPI) and threaded
- configured for random reads of entire dataset per epoch/iteration (x4 iterations)
  - small files dataset: whole-file reads of 32KiB and 1MiB files
  - large file dataset: block-aligned reads using 64MiB and 512MiB blocks

Deep Learning I/O Benchmark (https://github.com/argonne-lcf/dlio_benchmark)

- base software used by MLPerf Storage Benchmark
- models: resnet50, cosmoflow

DISCLAIMER: All performance results are the result of initial single run experiments. Draw conclusions at your own risk.

AI/ML Reads of Small File Datasets - 32KiB Files - First Epoch

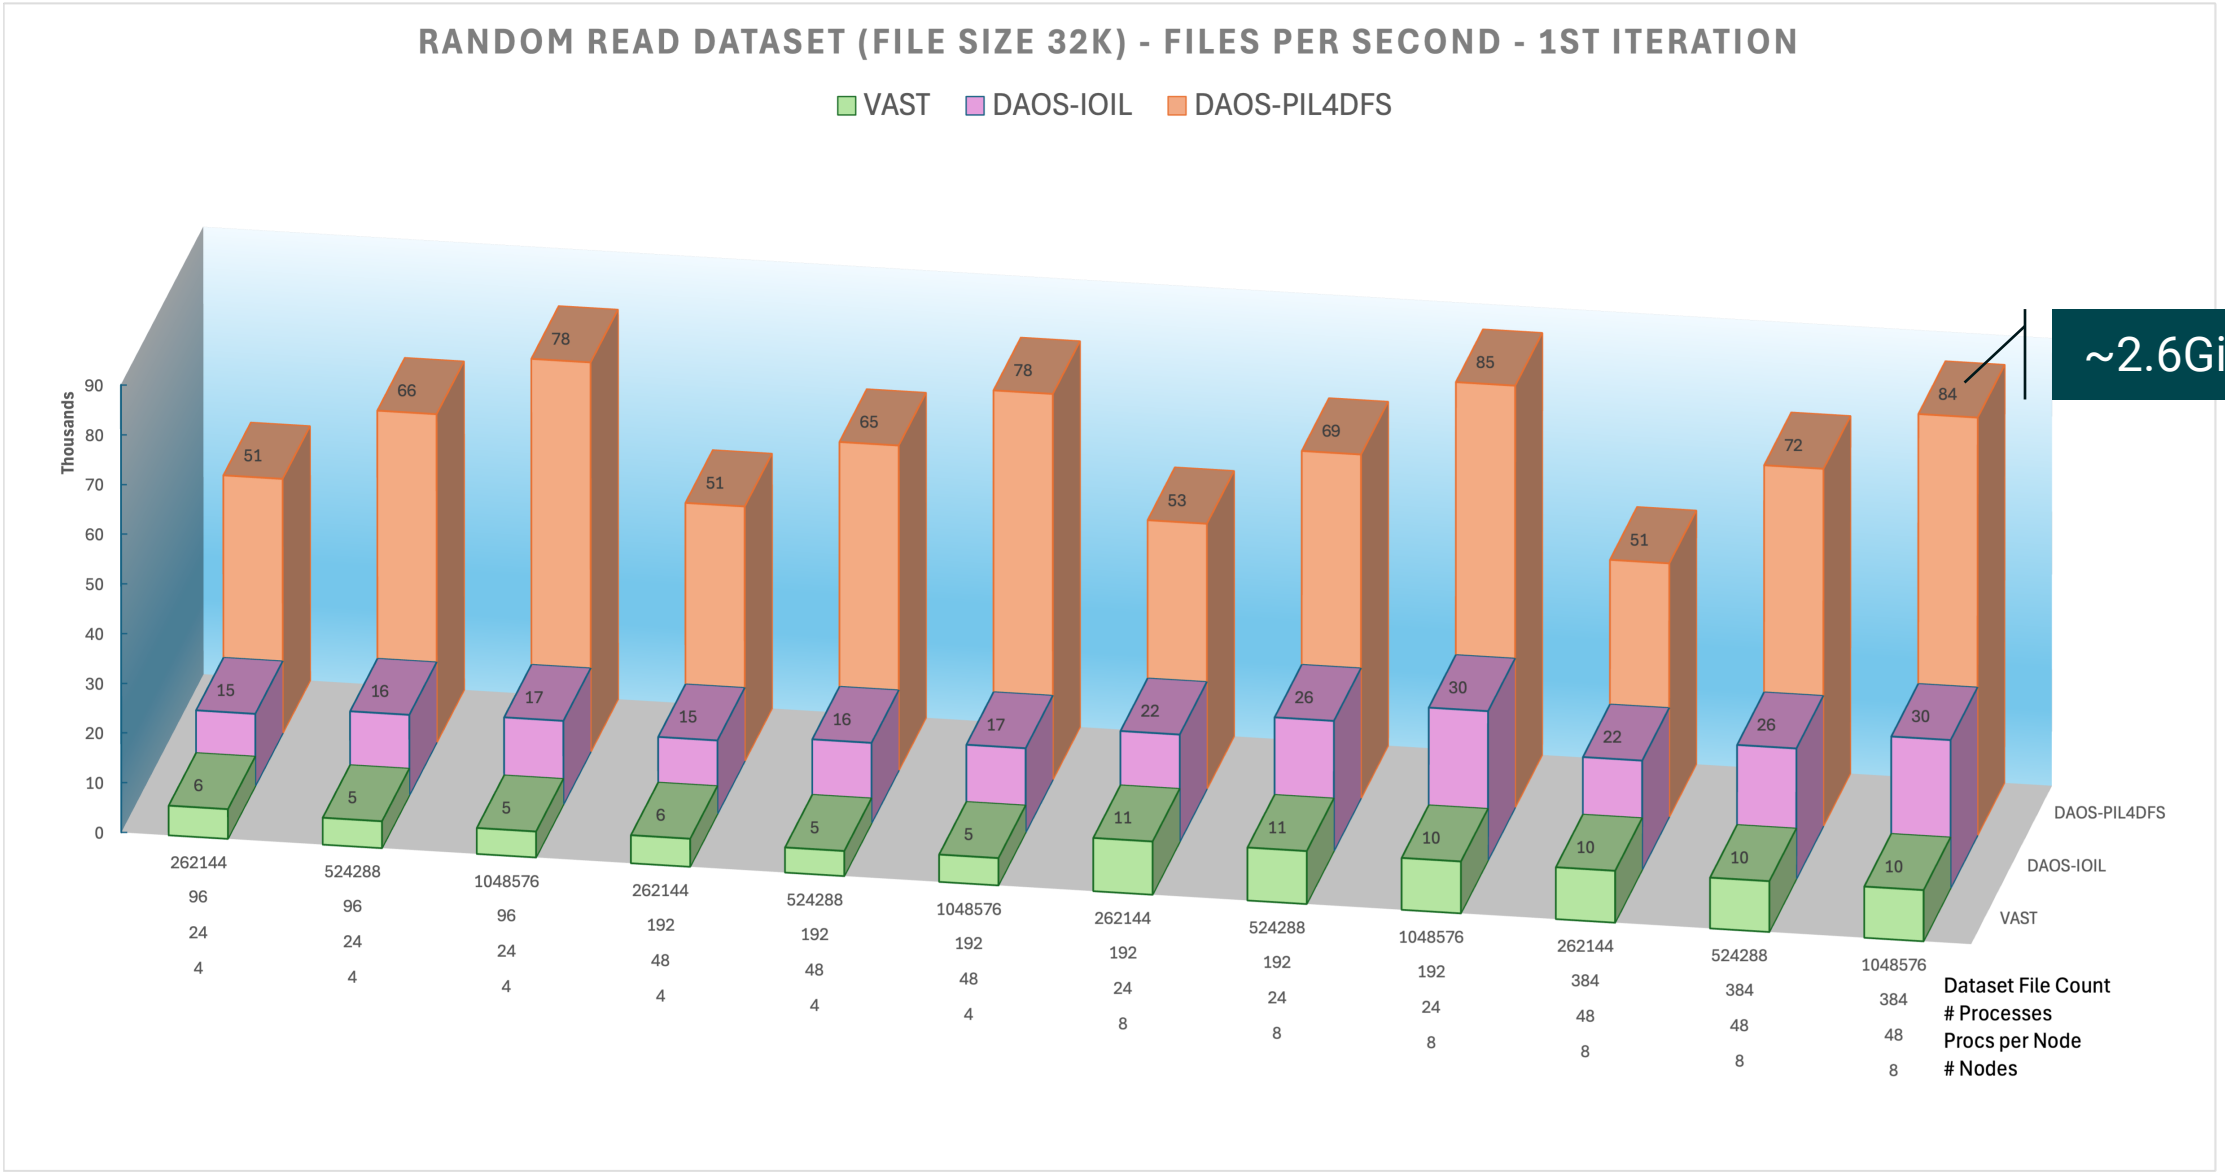

AI/ML Reads of Small File Datasets - 32KiB Files - Later Epochs

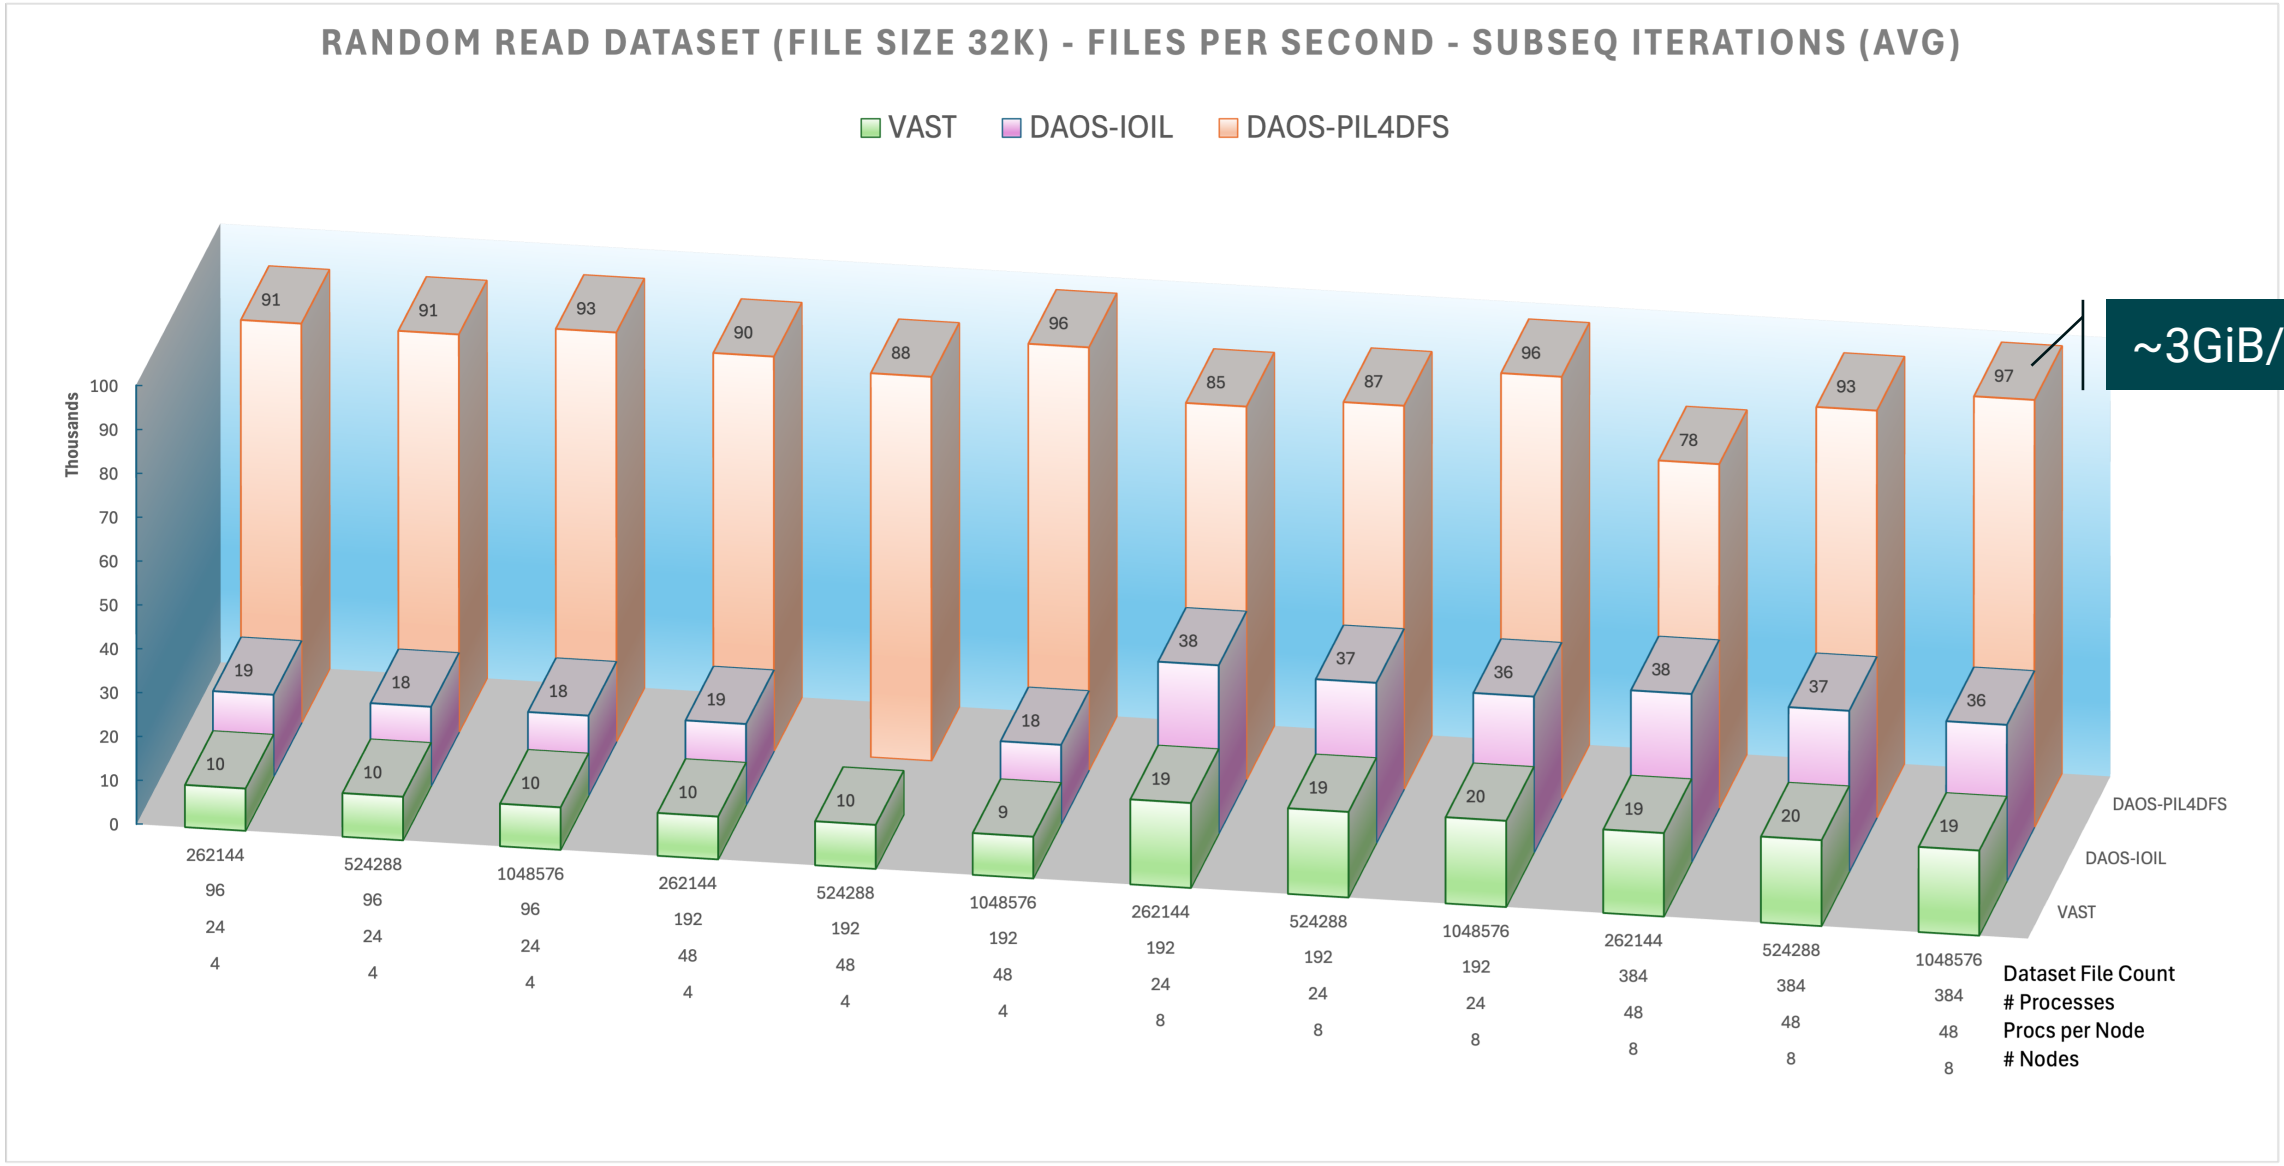

AI/ML Reads of Small File Datasets - 1MiB Files - First Epoch

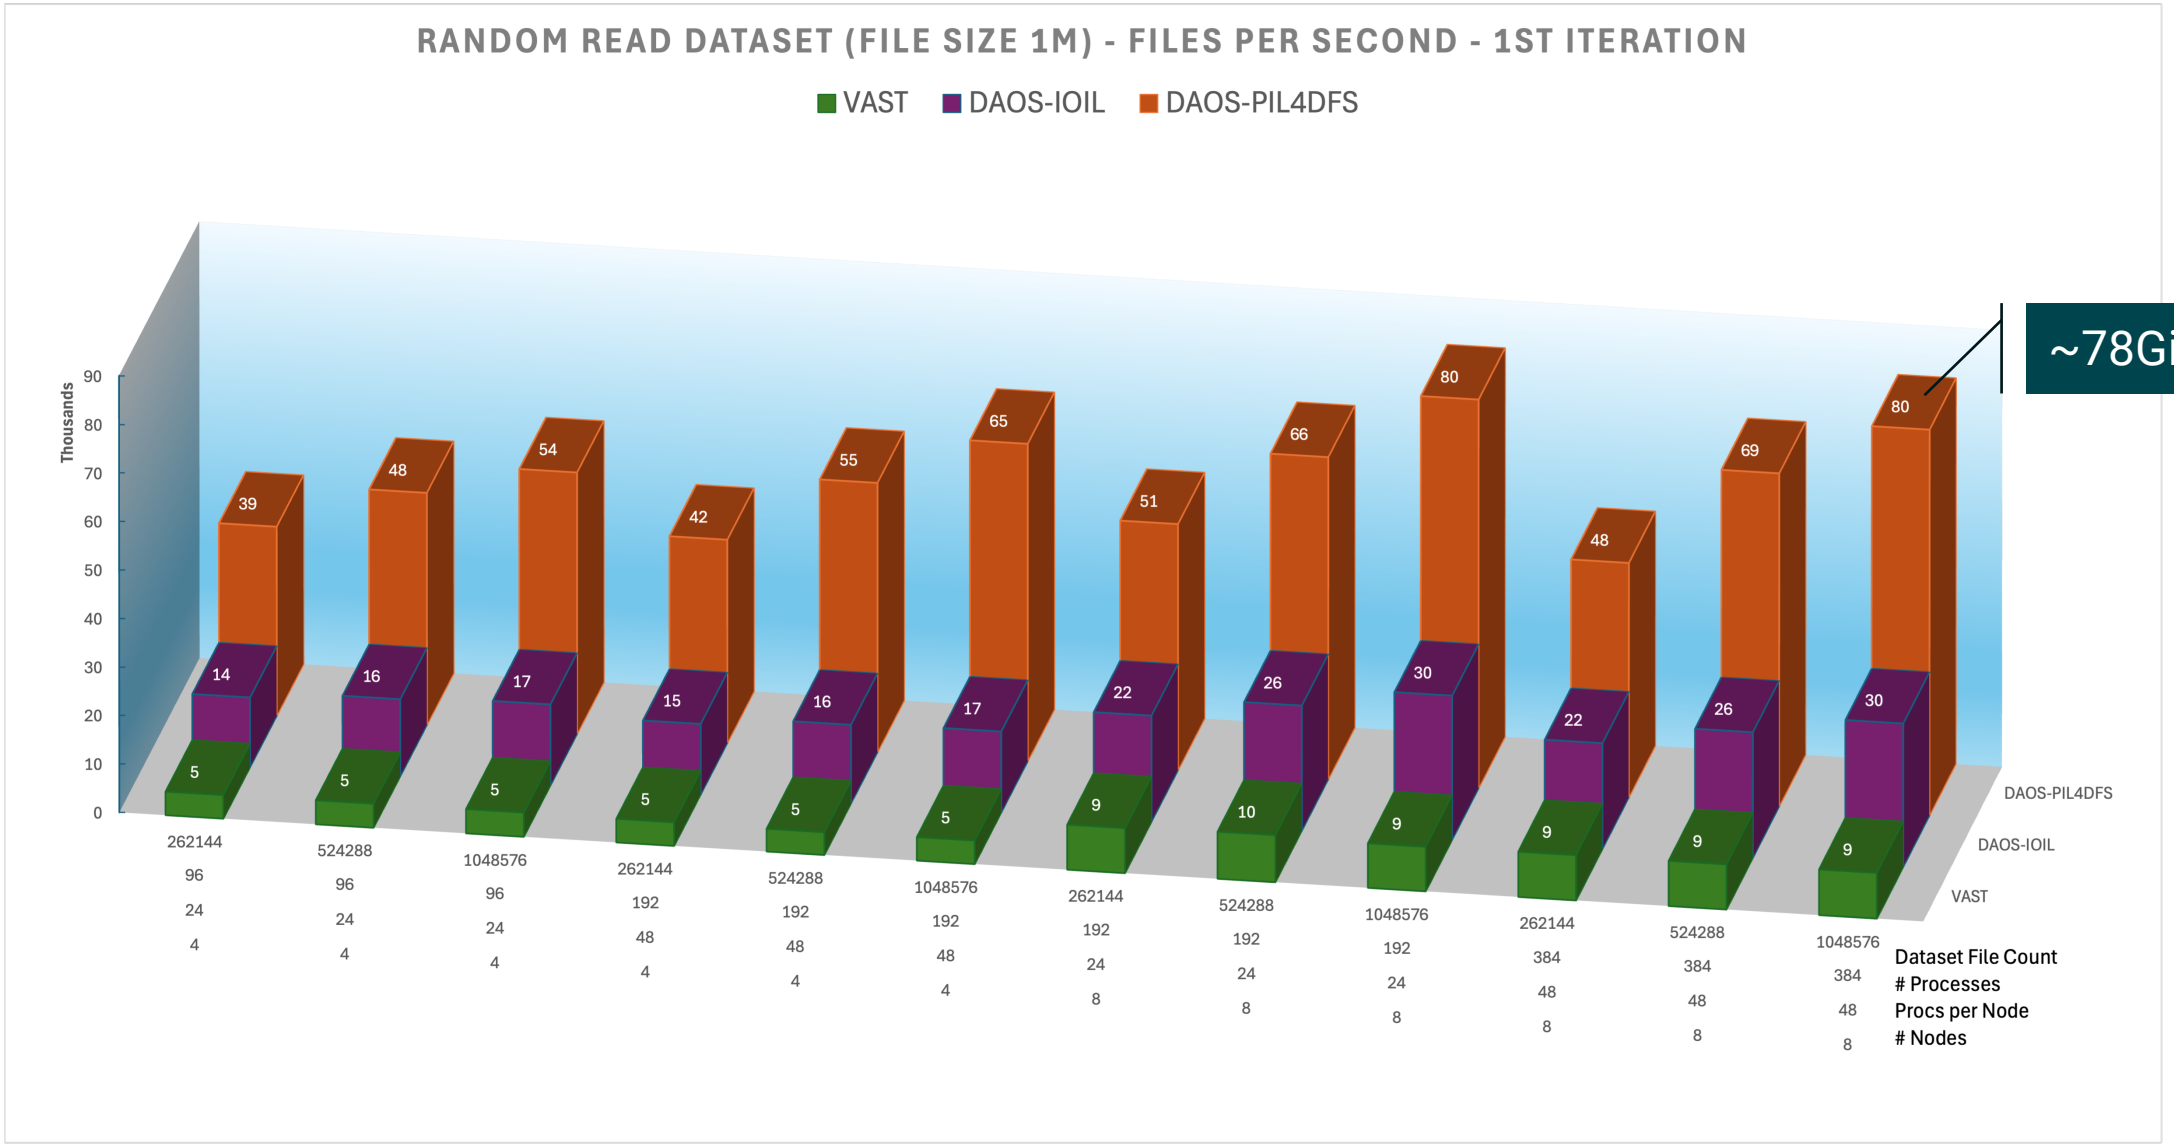

AI/ML Reads of Small File Datasets - 1MiB Files - Later Epochs

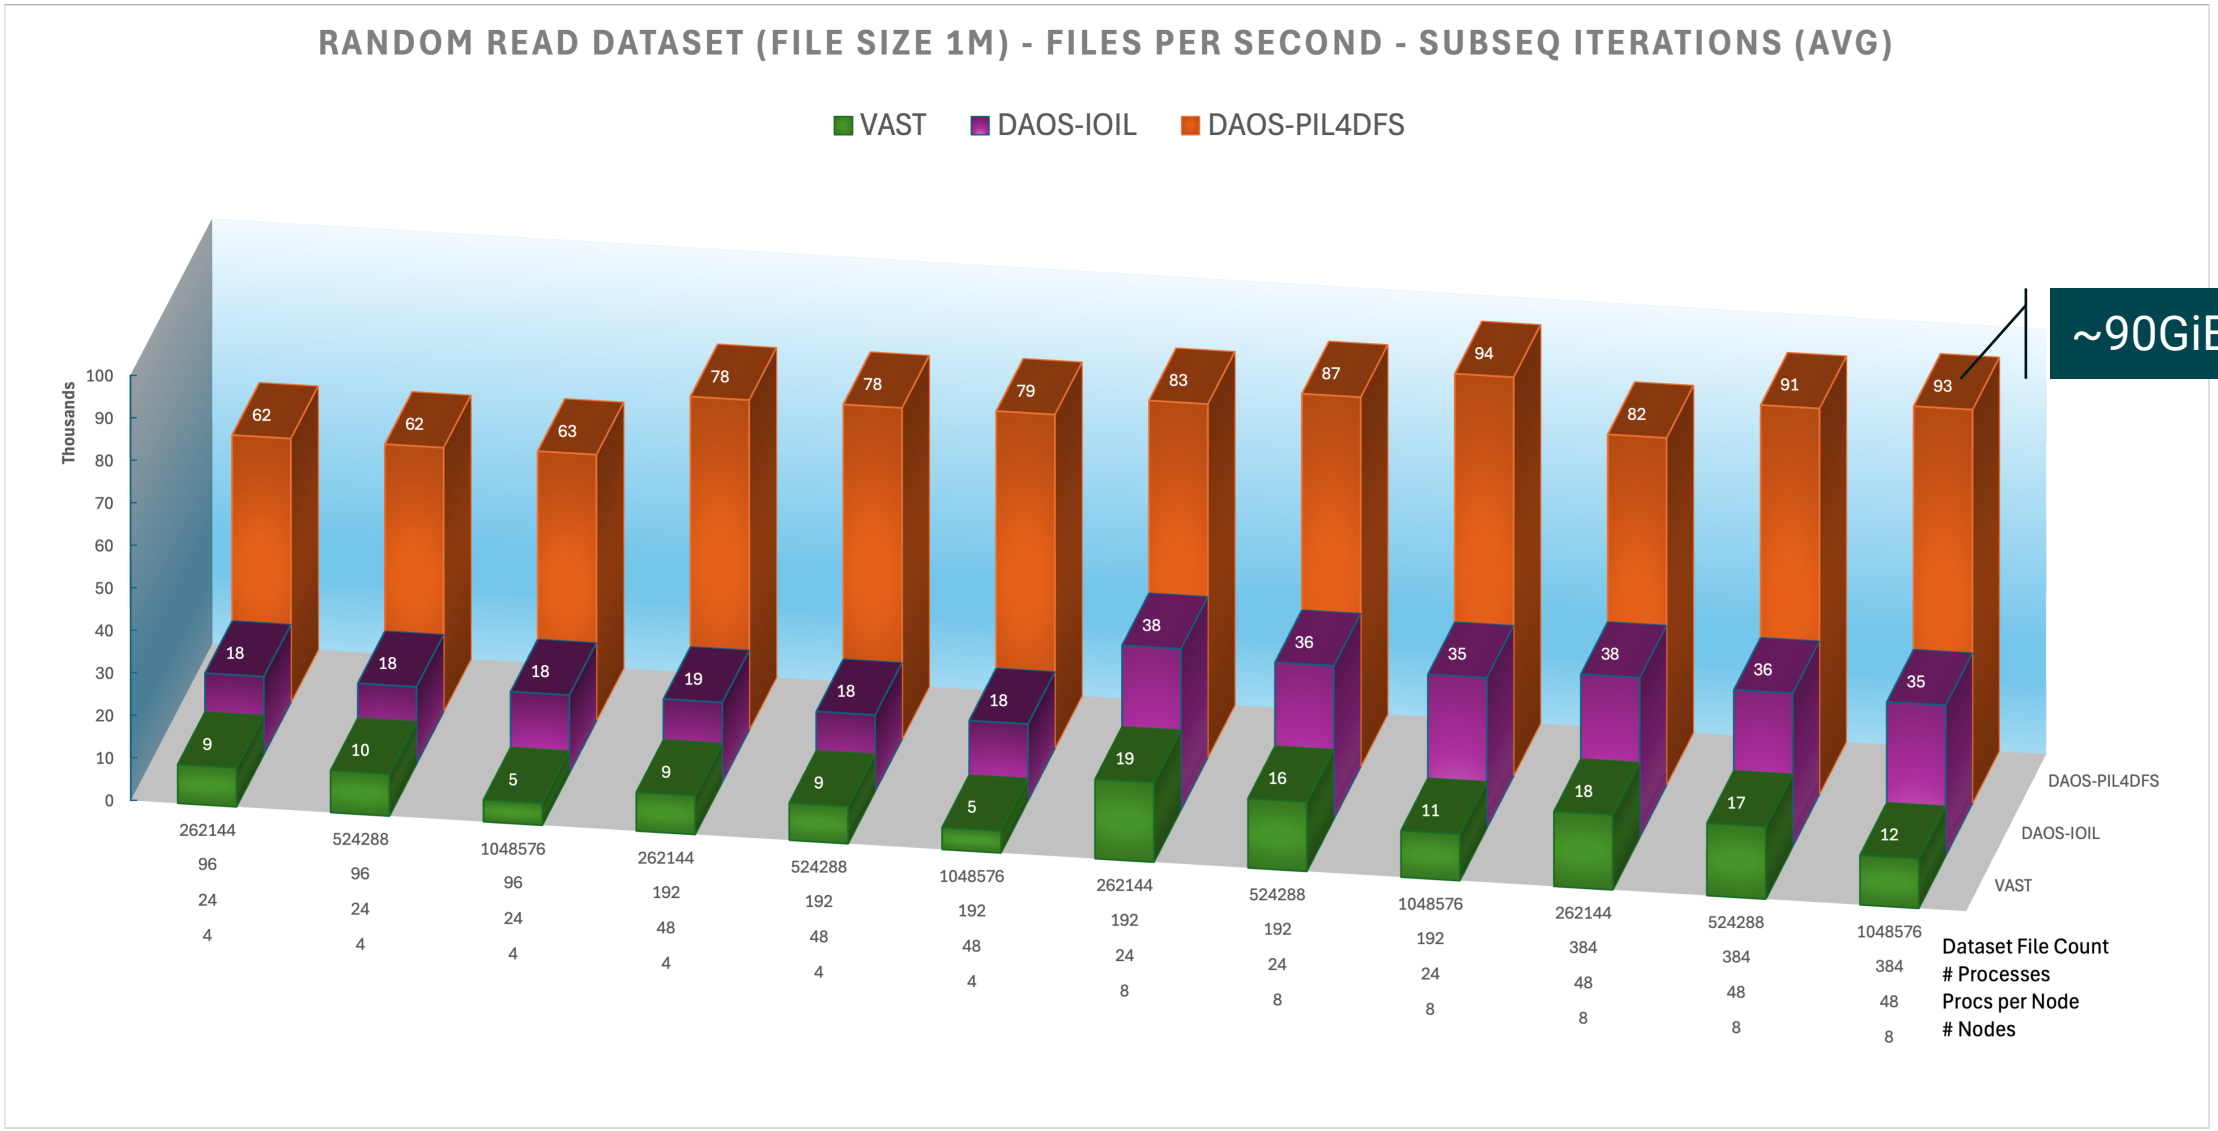

AI/ML Reads of Large File Datasets - First Epoch

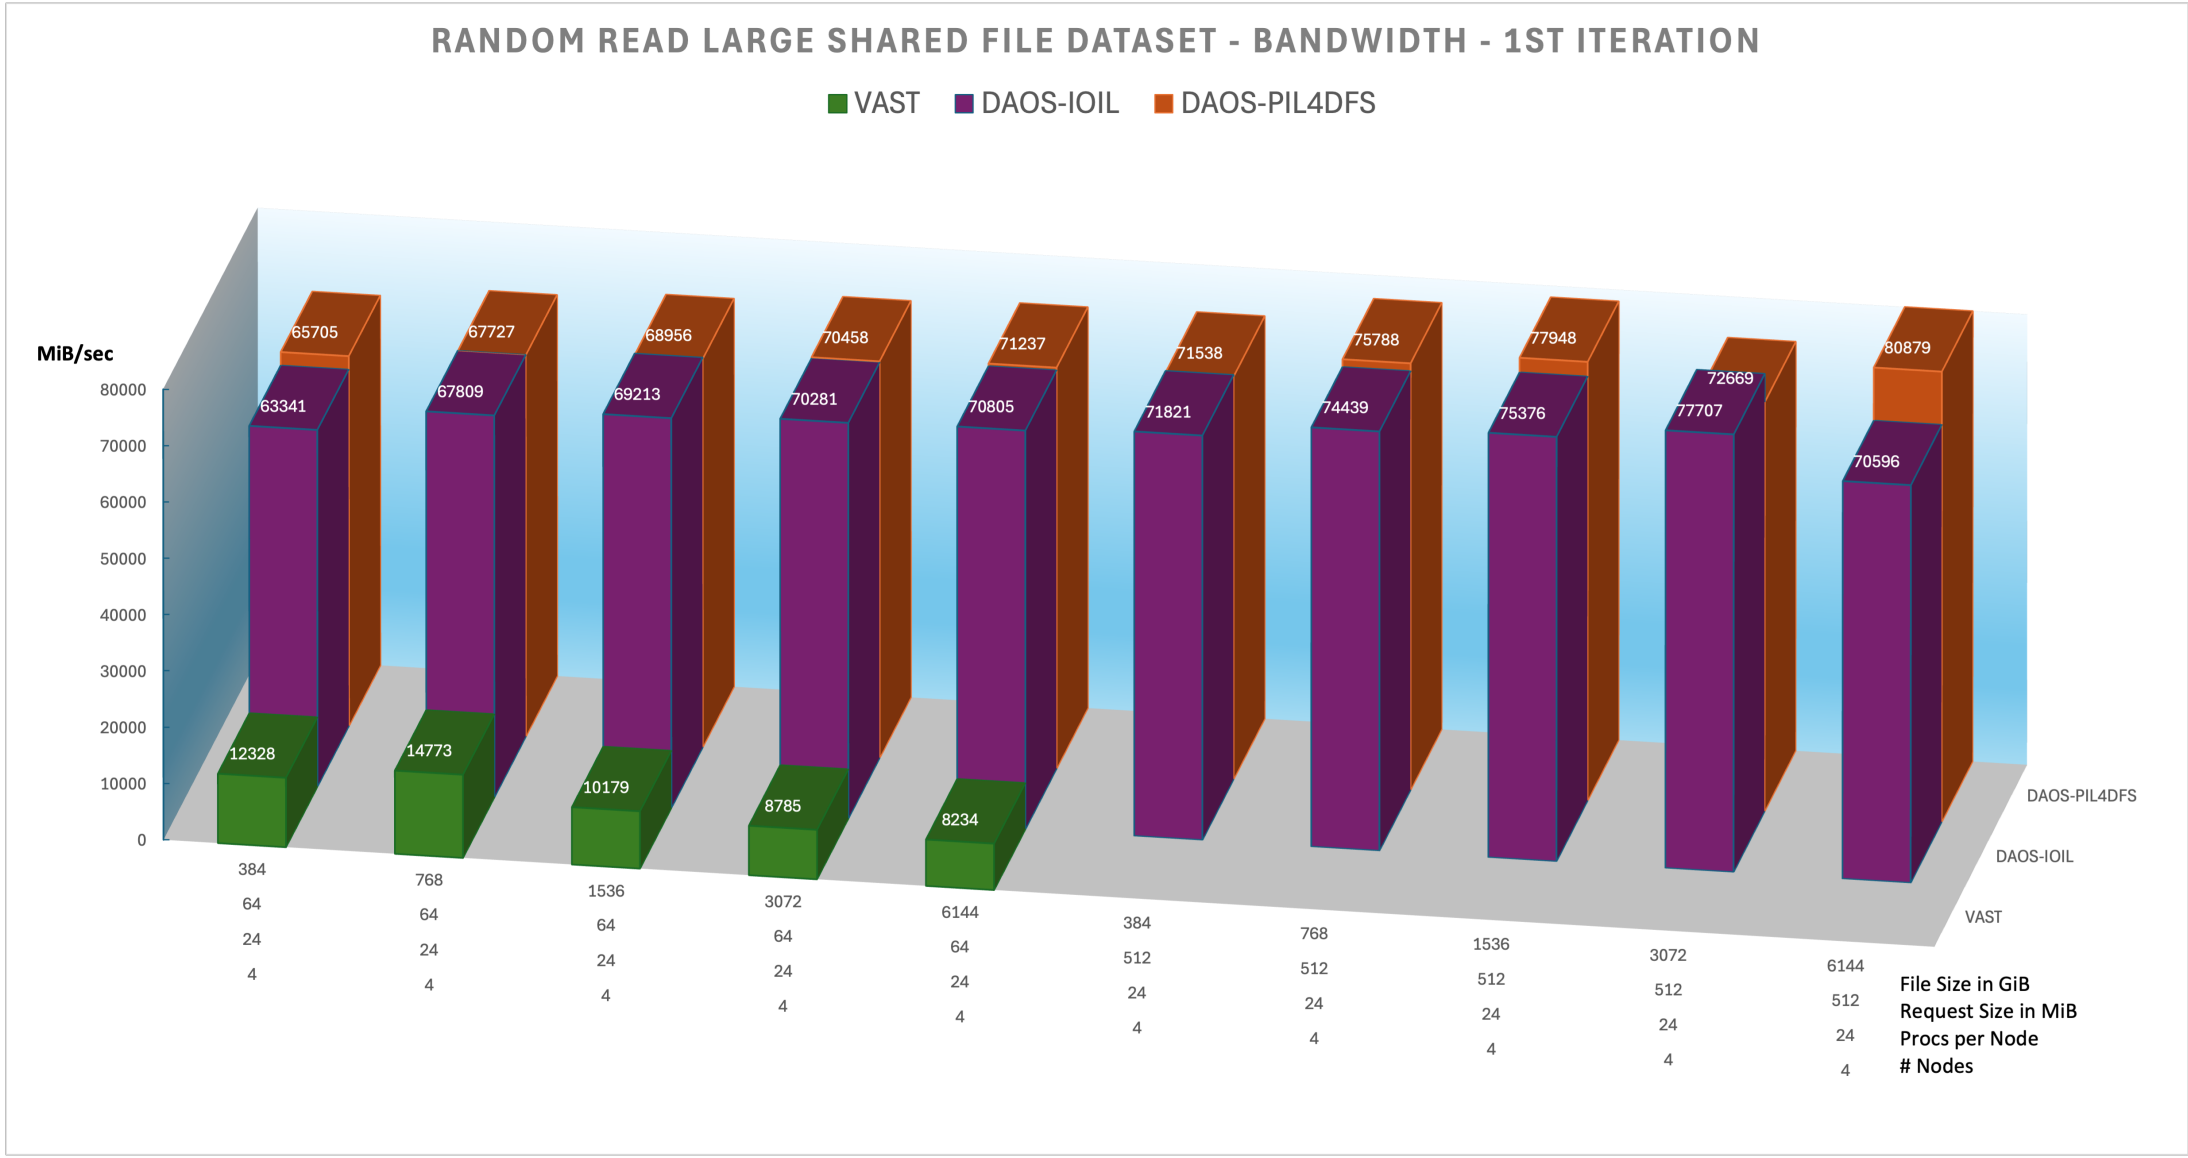

AI/ML Reads of Large File Datasets - Later Epochs

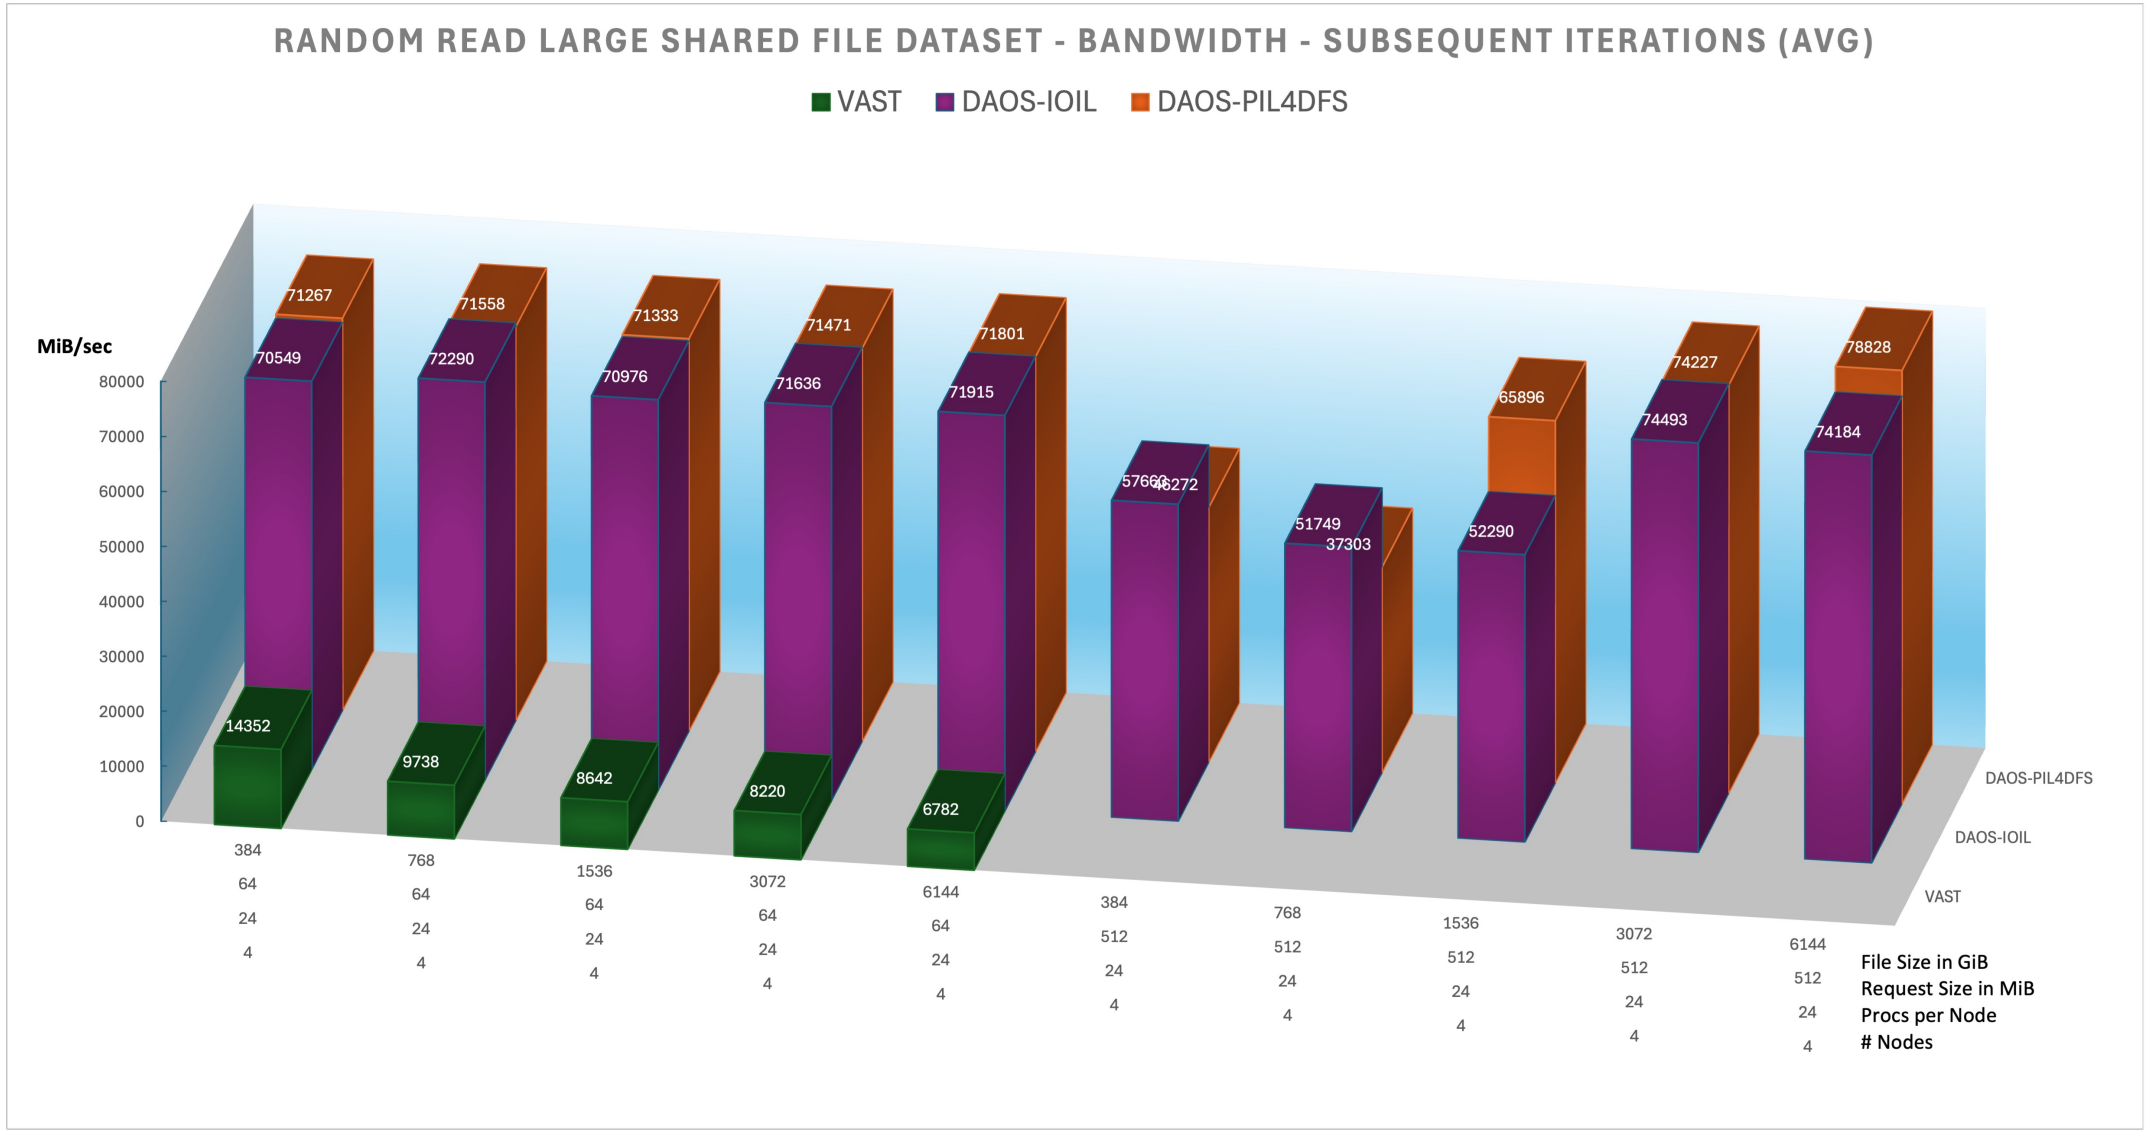

Deep Learning I/O (DLIO) Benchmark Results

- Resnet50 Training Dataset
  - 137 GiB
  - 1,024 137MiB files
  - 1,251 samples per file
  - 5 epochs
- Cosmoflow Training Dataset
  - 1.4 TiB
  - 524,288 2.7MiB files
  - 1 sample per file
  - 5 epochs

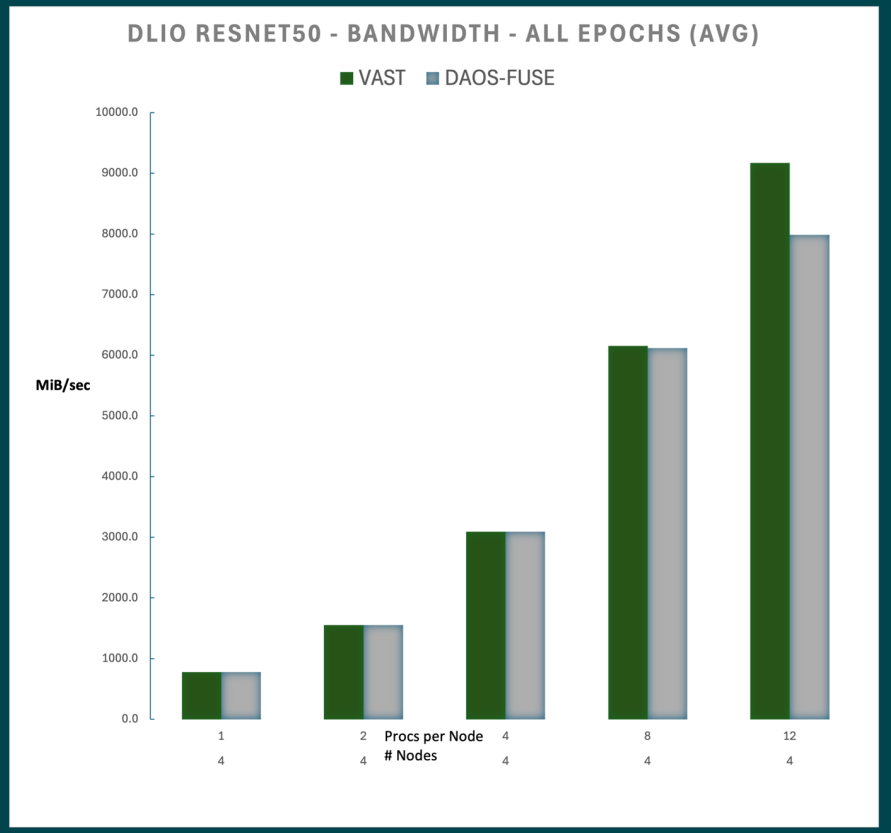

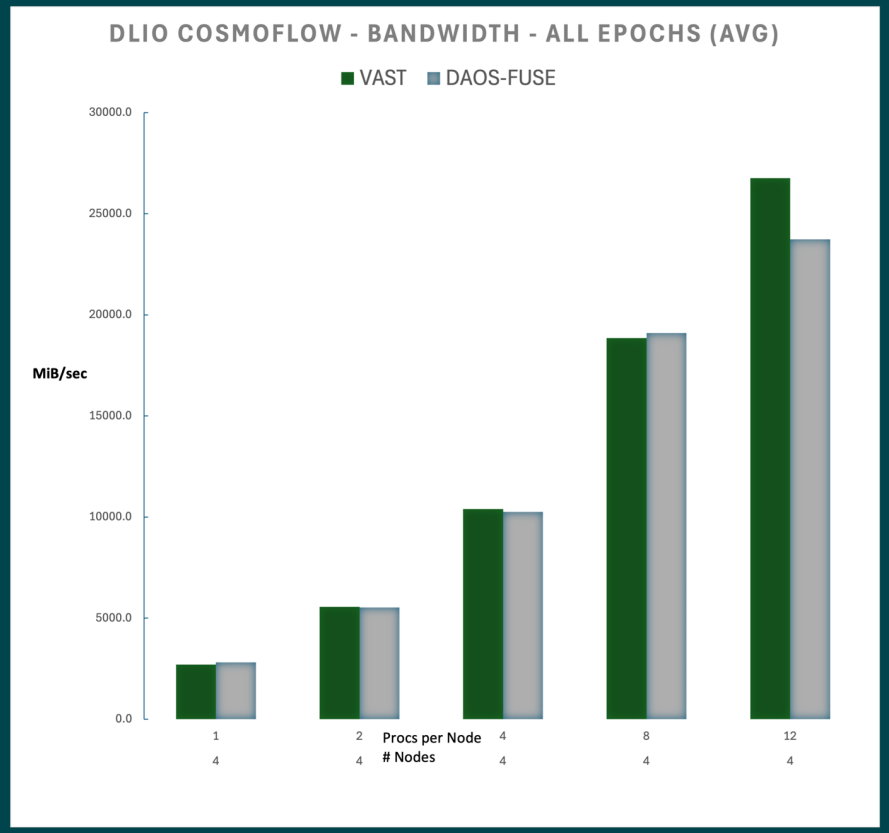

Early results support DAOS suitability for Scientific ML/AI

More to Come! Paper, AOS Evaluation Software

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.