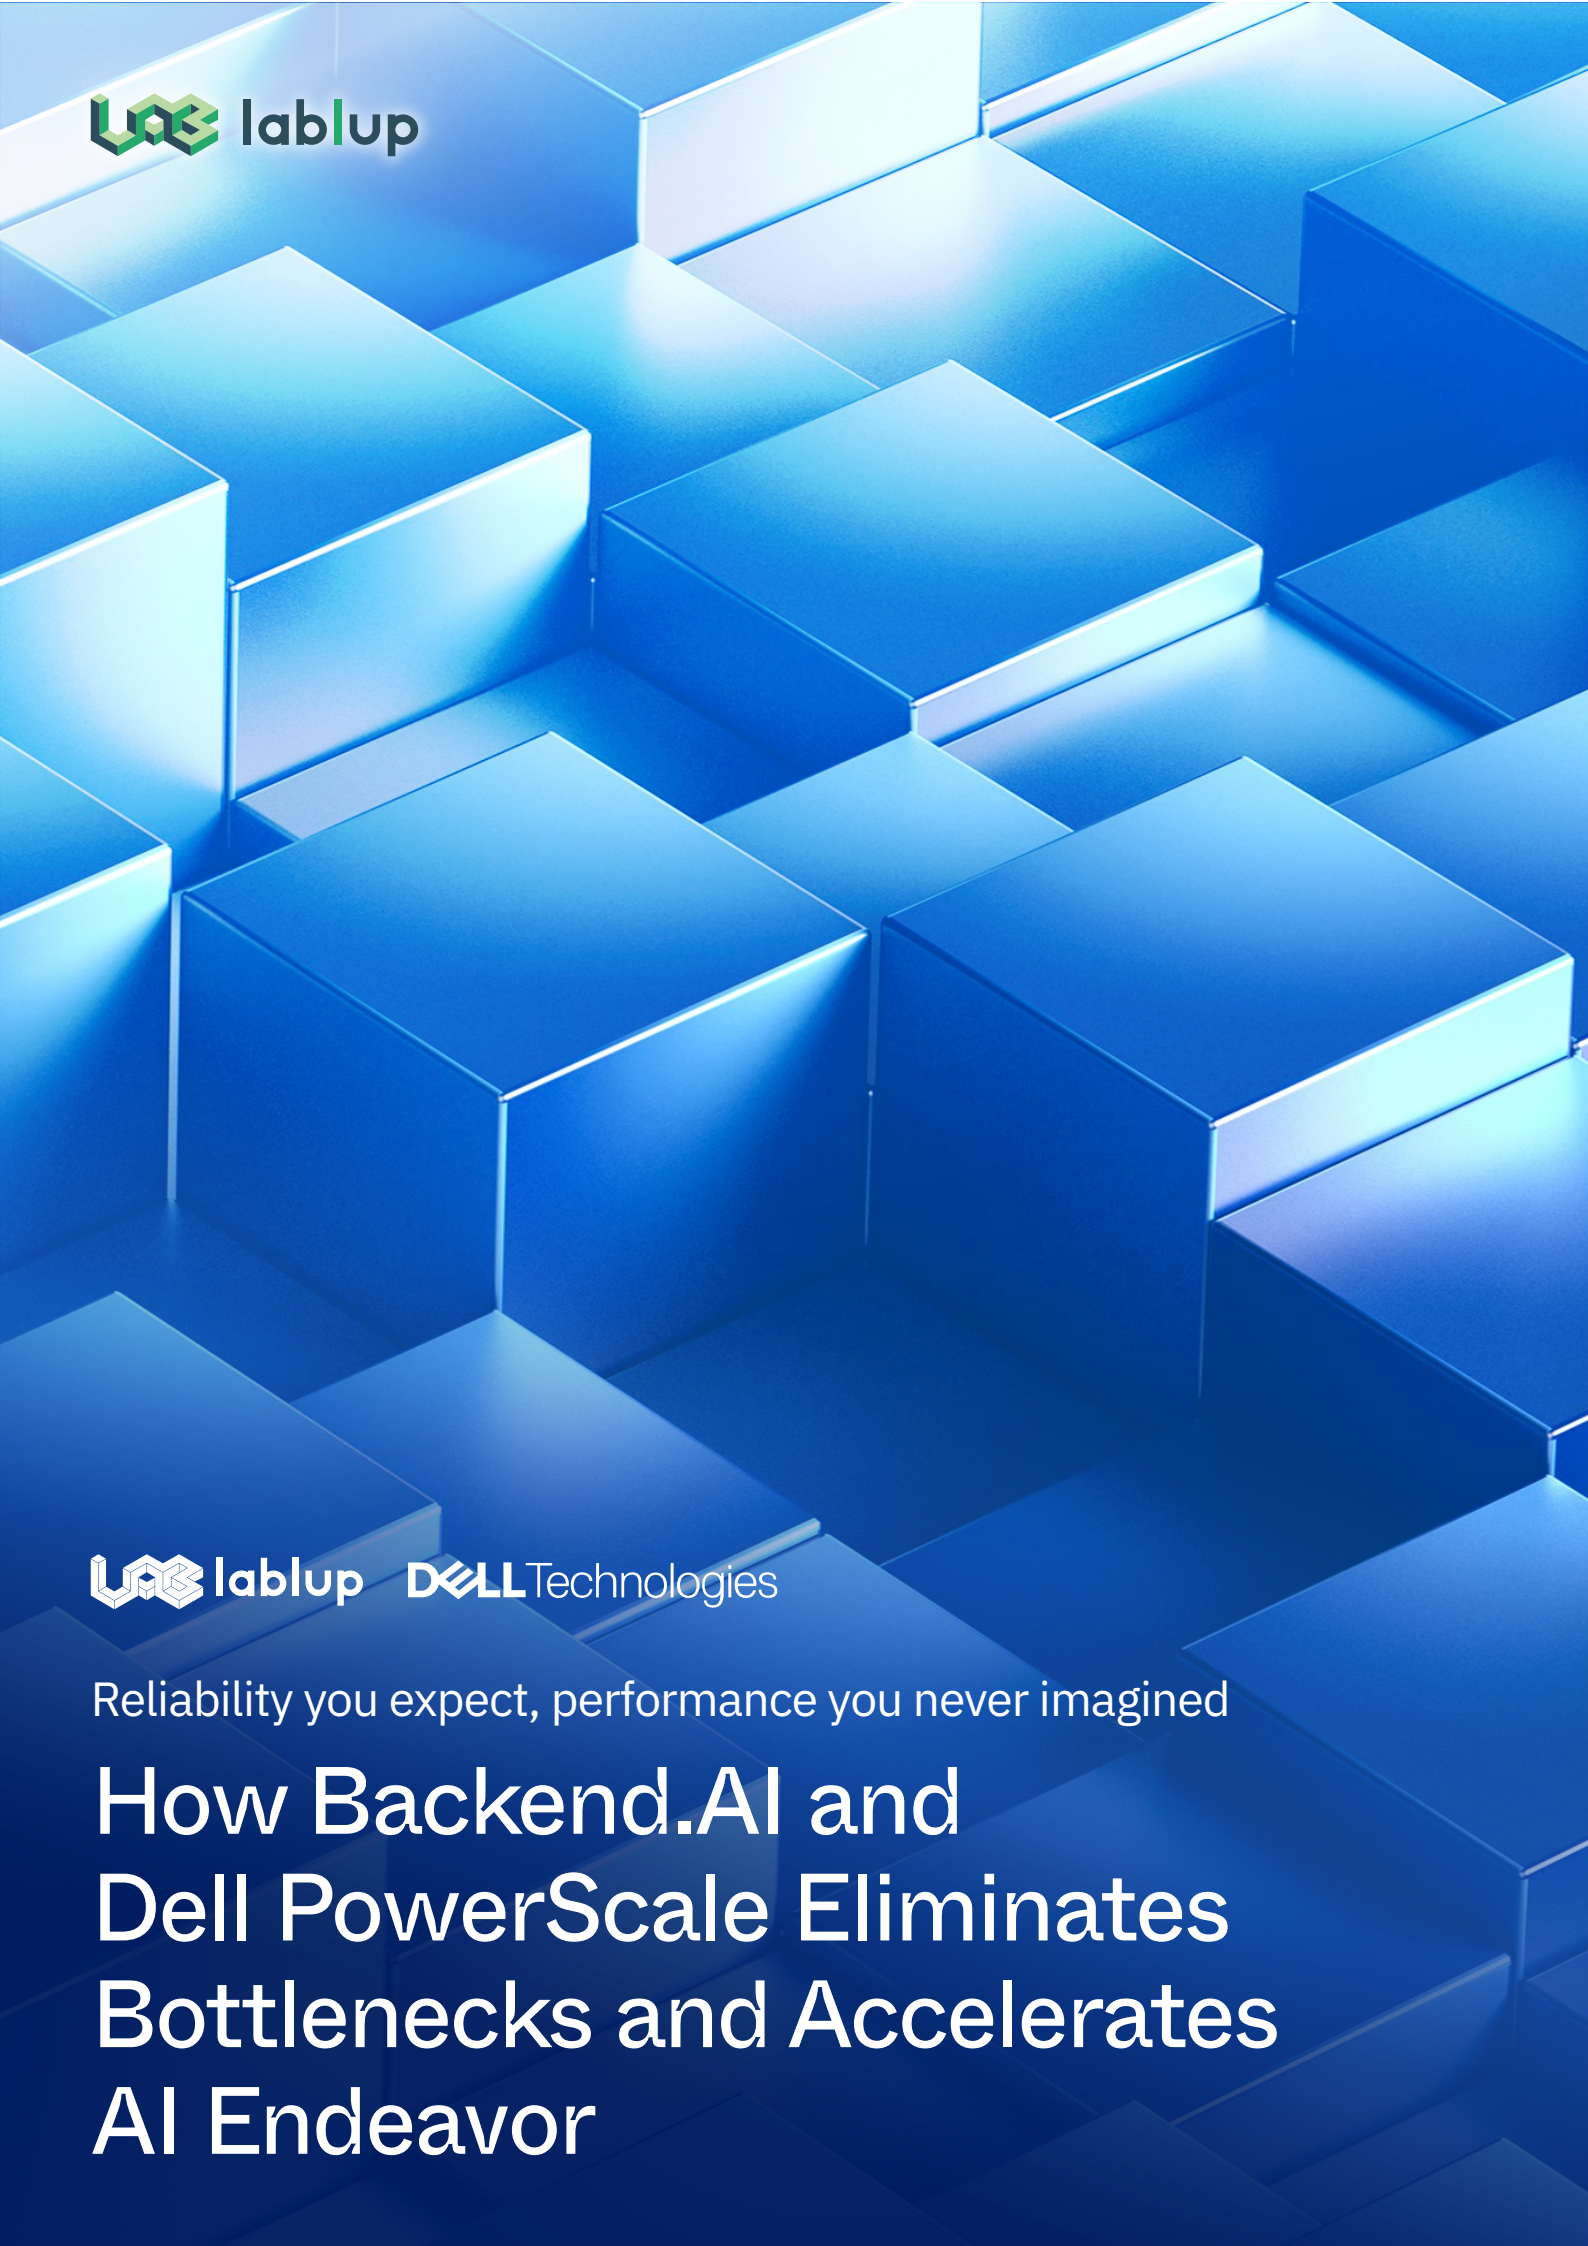
lablup
lablup DELL Technologies
Reliability you expect, performance you never imagined
How Backend.AI and Dell PowerScale Eliminates Bottlenecks and Accelerates AI Endeavor

lablup DELL Technologies

Solution Brief

Reliability you expect, performance you never imagined

How Backend.AI and Dell PowerScale Eliminates Bottlenecks and Accelerates AI Endeavor

The combination of Dell PowerScale and Lablup Backend.AI delivers a powerful and intuitive AI platform for data-intensive next-generation AI workloads. Dell PowerScale's high-performance scale-out storage infrastructure, integrated with Backend.AI's intelligent orchestration, automates and optimizes large-scale data processing and GPU resource utilization.
This enables enterprises to easily deploy, manage, and scale generative AI solutions with greater efficiency and simplicity

Why enterprise must address bottlenecks

Enterprise customers running data-intensive workloads such as AI and HPC are increasingly required to process massive datasets with speed and agility. In environments involving large-scale AI model training and inference, high-volume analytics, and real-time data processing, it is essential to minimize bottlenecks between storage and compute infrastructure to fully leverage GPU performance.

Traditional network file system (NFS) environments restrict throughput because a single NFS mount can only direct traffic to one storage node. As a result, when multiple GPU nodes simultaneously read and write large volumes of data, performance degradation occurs and individual client throughput is limited. This leads to GPU idle time as they wait for data, causing significant resource waste and reducing the overall efficiency of the AI pipeline.

When large volumes of data are transferred to GPUs, repeatedly passing through the CPU and system memory causes unnecessary data copying, increased latency, and CPU overhead, which degrade overall system performance. NVIDIA GPUDirect Storage (GDS) addresses this by providing a direct data path between storage and GPU memory, allowing data to move without CPU intervention. This reduces latency and maximizes bandwidth.

To effectively manage these direct data paths across diverse storage nodes and GPU environments, respond swiftly to failures, and ensure optimal data flow, a systematic infrastructure management system that fully supports GPUDirect Storage is essential.

Lablup Backend.AI:
AI infrastructure operating platform for relentless scale

Backend.AI is an AI infrastructure operations platform purpose-built for data-intensive workloads. By supporting multi-node and multi-tenant environments through the Sokovan™ orchestrator, Backend.AI enables simultaneous large-scale distributed training and inference services. Its container-level GPU fractional virtualization (Fractional GPU™) technology empowers customers to maximize GPU utilization and operate their infrastructure with greater efficiency and flexibility.

Backend.AI provides broad compatibility with a wide range of AI accelerators and high-speed storage solutions such as Dell PowerScale, enabling seamless scalability across both cloud and on-premises environments. With robust MLOps integration, high availability (HA), fine-grained resource control, automated experimentation and deployment, as well as advanced security and fault tolerance features, Backend.AI delivers a reliable solution that maximizes efficiency and productivity for enterprise AI infrastructure operations.

Dell PowerScale:
Trusted high-performance storage for data-intensive workloads

Dell PowerScale is a storage solution that delivers exceptional scalability and efficiency for large-scale data environments. Powered by the OneFS operating system, PowerScale enables seamless node additions and non-disruptive expansion to multiple petabytes. PowerScale maximizes storage resource utilization, eliminates data silos, and provides optimal storage efficiency and industry-leading reliability.

PowerScale supports a wide range of protocols—including NFS, SMB, HDFS, and S3—and features an architecture optimized for modern workloads such as AI, big data, and cloud environments. With automated data tiering, inline data reduction, and robust data protection, enterprises can flexibly respond to data growth while achieving both operational efficiency and cost savings.

Lablup and Dell: Paving the way for a future without data bottlenecks

Benchmark: Multipath + GPUDirect Storage on Backend.AI and Dell PowerScale

This benchmark was conducted to evaluate the performance improvement achieved when both Dell PowerScale Multipath and GPUDirect Storage are enabled simultaneously. Using cuCIM, NVIDIA's computer vision acceleration library for GPUs, the test compared TIFF image read and write performance with and without Multipath Driver and GPUDirect Storage enabled.

Benchmark #1

Read & Write Benchmark

gausian filter + global kernel scope

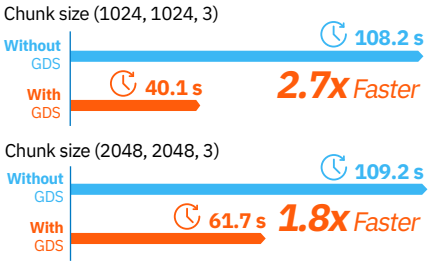

This benchmark involved reading TIFF images, decompressing them, applying various filters by tile size, and writing the uncompressed parameters to disk in Zarr format.

For a chunk size of (1024, 1024, 3):
Without GDS: 108.2 seconds
GDS enabled: 40.1 seconds
Approximately 2.7x faster

For a chunk size (2048, 2048, 3):
Without GDS: 109.2 seconds
GDS enabled: 61.7 seconds
Approximately 1.8x faster

Benchmark #2

Write Benchmark

kvikio-pwrite + float32 data format

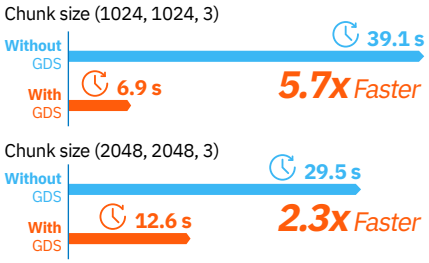

This benchmark was conducted by writing TIFF images to disk as uncompressed parameters in Zarr format, excluding the reading and preprocessing steps included in benchmark_roundtrip.py.

For a chunk size of (1024, 1024, 3):
Without GDS: 39.1 seconds
GDS enabled: 6.9 seconds
Approximately 5.7x faster

For a chunk size of (2048, 2048, 3):
Without GDS: 29.5 seconds total
GDS enabled: 12.6 seconds total
Approximately 2.3x faster

Benchmark #3

Compress + Write Benchmark

uint16 / LZ4 Compressed with nvCOMP

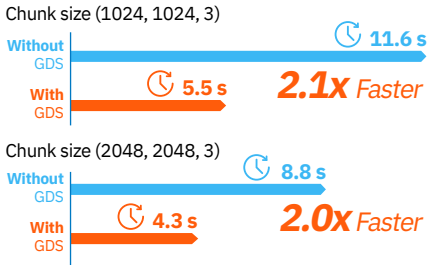

This benchmark compares the performance of writing TIFF images to disk as parameters in Zarr format, with and without LZ4 compression, using NVIDIA's compression acceleration API, nVCOMP.

For a chunk size of (1024, 1024, 3):
Without GDS: 11.6 seconds
GDS enabled: 5.5 seconds
Approximately 2.1x faster

For a chunk size of (2048, 2048, 3):
Without GDS: 8.8 seconds
GDS enabled: 4.3 seconds
Approximately 2.0x faster

Result and Outcomes

When Backend.AI is combined with GPUDirect Storage and Dell PowerScale Multipath, tasks are completed at least 1.8x and up to 5.7x faster. Even when using LZ4 compression and nvCOMP acceleration, a notable performance improvement is maintained, demonstrating that optimization benefits persist even with added compression steps. In environments with frequent data movement or significant I/O bottlenecks, implementing GDS maximizes GPU resource utilization and effectively increases overall pipeline throughput. Enhanced data processing speed unlocks new possibilities for leveraging complex and large-scale datasets—use these advancements as a springboard for new technologies, services, and research initiatives.

Benchmark Conditions: PowerScale F600

Storage: 6 nodes with 6 100GbE client connections
Client: 2 H100 fGPU on 1 HGX H100 node with single 100GbE connection
without GDS: NFS mount without multipath and RDMA option enabled
with GDS: NFS mount with Dell Multipath driver version 4.0.30 (nconnect=60, 6 remoteports)

Explore the full potential of Backend.AI and Dell PowerScale
Contact Lablup or visit backend.ai for more information.