Micron® 6600 ION NVMe™ SSDs: AI data lake storage built to keep GPUs fed

Rethinking AI data lakes: One NVMe SSD or 16 hard drives?

AI data lakes are often treated as “cheap and deep” storage, huge-capacity systems that have often been built from spinning disk, after disk, after disk. When we needed more capacity or performance, we added more disks.[1]

Was that approach ever optimal? No, not really.

There is a better way. Today, we can increase storage capacity while simplifying the architecture, improving performance, and boosting power efficiency. This change is fundamental and architectural. Modern lake workloads are no longer dominated by simple, sequential data I/O. They are more complex, more concurrent, and more demanding.

This technical brief analyzes storage performance across AI data lake pipelines using measured benchmark data from a single Micron® 6600 ION 245TB NVMe™ SSD and a legacy building block of scale-out storage, a 256TB HDD array (16x 16TB Seagate® EXOS® HDDs).[2]

Five workloads representing five AI data lake pipeline stages were tested: data ingest, data preparation, object storage, training data loading, and bulk data transfer.

The goal was not to declare a universal winner, but to provide a practical framework for choosing the right storage architectural building blocks from which to create AI data lakes (scaling each building block to meet data lake capacity requirements in production). A data lake is often considered the most favorable use case for HDDs in AI.

Micron 6600 ION NVMe SSD (245TB)

Key findings

SSDs are the AI data lake winners

SSD performance and power efficiency help ensure that storage isn’t the limiting factor in AI data lake results. In testing, 1x 245TB Micron 6600 SSD compared to 16x 16TB HDDs delivered up to:

4.4x
Data loading

In MLPerf® Storage v2.0 (U-Net3D workload) data loading, a single SSD delivered 4.4x the throughput of 16x HDDs (12,738 vs. 2,906 MB/s).

45.6x
Object storage

MinIO S3 data-access operations showed a single SSD delivering up to 45.6x the throughput of 16x HDDs (8,473 vs. 186 MB/s)

8.6x
Data preparation

ETL preprocessing tests showed an 8.6x throughput increase with 1x SSD (5,415 vs. 632 MB/s), accelerating a time-consuming stage.

282x
Power efficiency

1x SSD delivered up to 282x the power efficiency (MB/s per watt) of the 16x HDD, helping reduce data center operating costs.

Across the AI data pipeline, the Micron 6600 ION SSD delivered the bandwidth and power efficiency AI data lakes demand.

micron.com/6600-ION

1. For additional performance from HDDs, we might short-stroke them. See “Flex Dynamic Recording” (usenix.org).
2. Unformatted. 1GB = 1 billion bytes. Formatted capacity is less. In this document, the 1x 245TB Micron 6600 ION SSD configuration is referred to as “1x SSD,” and the 16x 16TB configuration is referred to as “16x HDD” for brevity.

What is an AI data lake?

An AI data lake is a centralized, high-capacity repository designed to store raw and processed data for machine learning pipelines at petabyte scale. Unlike traditional data warehouses that require data to be structured before storage, data lakes accept data in its native format: structured tables, semi-structured JSON/Parquet files, unstructured images, video, and text—enabling flexible access patterns for diverse AI workloads.[3]

Industry analyst research treats the data lakehouse as a mature platform category, evaluating vendors and enterprise capabilities across a defined set of criteria.[4] The shift from traditional storage to data lakes is driven by three factors: the need to store diverse data types, the requirement for high-throughput parallel access during model training, and the economics of storing petabytes of training data cost-effectively.

Data lake storage architecture components

While we may think of a car as a single machine, we know it's more complex than that. It is built from multiple integrated systems that work together. Data lakes are similar.

Storage layer: The data lake foundation. S3-compatible object storage (AWS S3, MinIO, Ceph) that holds the data lake's datasets, checkpoints, and intermediate outputs. It needs high throughput for sequential training reads and for random/mixed ETL I/O. We measured this layer using the MinIO S3 benchmark.

Compute layer: Processing engines (like Apache Spark, Dask, and Ray) read from and write to the storage layer during data preparation and feature engineering. These engines can generate mixed read/write I/O patterns with varying block sizes. These are the patterns we measured with ETL benchmarks.

Metadata and catalog layer: Delta Lake, Iceberg, and Hudi add transactional metadata to table-formatted data stored in object storage, enabling ACID transactions, schema evolution, and time travel.

Training and serving layer: ML frameworks (PyTorch, TensorFlow, JAX) that read training data from the lake and write model artifacts back. We used the MLCommons® MLPerf® Storage v2.0 3D U-Net (U-Net3D) workload to measure the I/O demands of this layer.[5]

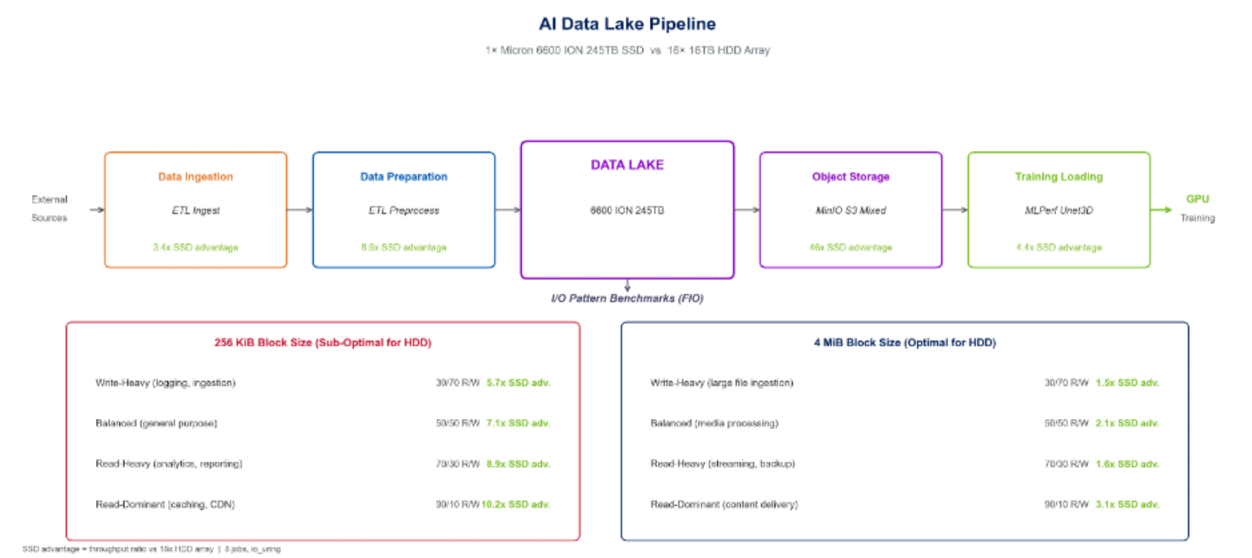

Figure 1: AI data architectural overview (courtesy Micron Data Center Workload Engineering

3. In this paper, 'data lake' refers to lake-based architectures broadly, including data lakehouses where warehouse-grade governance and performance are layered on top of lake storage." also see "What is a data lake?" (sap.com) for additional details on AI data lakes.
4. "Key Takeaways from the 2025 State of the Data Lakehouse Report: Navigating the AI Landscape" (forrester.com).
5. The results shown were run in Micron's data center workload engineering lab and are not official MLPerf Storage results. See the MLCommons home page to learn more about MLCommons, its charter, and the benchmarks it develops. This benchmark is hereafter abbreviated "MLPerf UNet3D" for brevity.

Data movement is critical

Optimizing data lake storage is a balancing problem. AI accelerator performance (one side of the balance) has been scaling faster than the data lake bandwidth needed to feed it (the other side of the balance). And HDDs are on the wrong side of the balance. Data movement, not simple math, can govern pipeline speed.

To start to rebalance, storage must sustain high throughput across sequential, random, and mixed I/O to keep accelerators utilized. This report quantifies the storage I/O gap between 1x NVMe SSDs (1x SSD) and 16x 16TB HDDs (16x HDD) across multiple AI data lake stages using measured, production-representative workloads.

Classic HDD-based storage rarely fails an AI data lake on capacity; but can fail on behavior. HDDs can appear acceptable in narrow, sequential lanes, but exhibit large swings as access shifts to mixed I/O and metadata-heavy requests—potentially leading to deeper queues, retries, buffering, and missed service-level agreements (SLAs).

Measuring SSD success

We measured five AI data lake data stages using common benchmarks for illustration: loading from the AI data lake into the training pipeline (using MLPerf UNet3D), S3-compatible object storage accessed via cloud-native APIs, data preparation, data ingest, and bulk sequential transfer.

We then compared performance, latency, power efficiency, and response time for each, closely matching the capacity of a single 245TB Micron 6600 ION SSD to an array of 16x 16TB HDDs (256TB).

Early and commonly referenced data lake architectures assume large, throughput-oriented access patterns, which align with HDDs' best-case performance characteristics.[6]

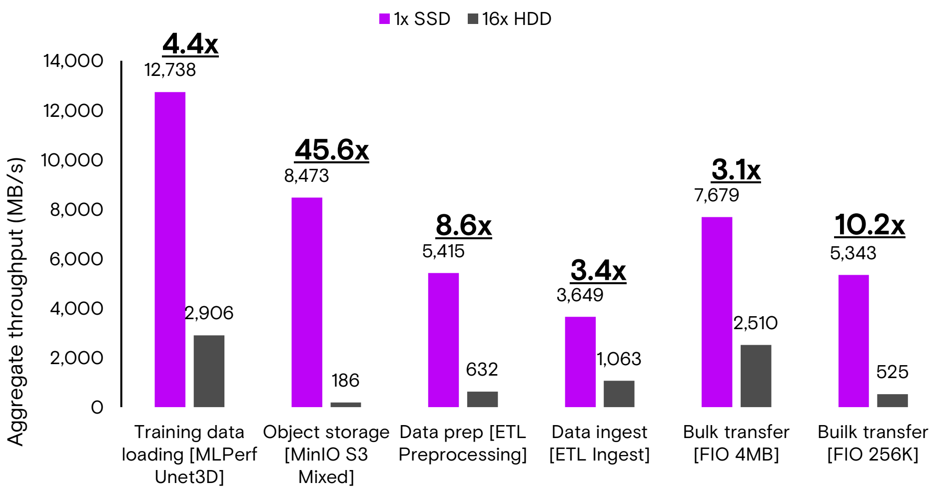

Figure 2: AI data lake stages, measured results, and benchmarks used.

Loading training data: MLPerf Unet3D
Storage throughput often affects whether GPUs remain busy or starve for data.[7] The MLPerf Storage benchmark simulates realistic patterns of loading training data. At 12,738 MB/s measured throughput, a single SSD can supply multi-GB/s training data streams to multiple host-attached accelerators, whereas a 16-HDD array delivered only 2,906 MB/s. This deficit can create a persistent bottleneck that wastes expensive GPU compute time.

SSD advantage: 4.4x the 16x HDD throughput (12,738 / 2,906 ≈ 4.4).

Object storage: MinIO S3 mixed
S3-compatible object storage with mixed GET/PUT operations via MinIO Warp is a model for cloud-native AI data lake access patterns.[8]

The MinIO benchmark results reveal that the 16x HDD delivered only 186 MB/s under realistic S3 workloads, as object storage operations often involve small metadata lookups,[9] small random reads,[10] and mixed read/write patterns that HDDs typically do not handle well.
SSD advantage: 45.6x the 16x HDD throughput (8,473 / 186 ≈ 45.6).

6. "Understanding the evolution of data lake" (capitalone.com) and "How to Model Hard Drive Performance in the Cloud" (seagate.com).
7. "GPUs Are Fast, I/O is Your Bottleneck" (alluxio.io).
8. "Architecture of an object storage system" (min.io) and the warp CLI references (min.io).
9. "Request rate and access distribution guidelines" (docs.cloud.google.com).
10. "Storage Benchmarking with Deep Learning Workloads" (newtraell.cs.uchicago.edu).

Data preparation: ETL preprocessing
The ETL data preparation stage operates within the data lake, cleaning, normalizing, and transforming raw data into model-ready inputs while storing intermediate and final results back in the lake.[11] In practice, this stage may represent the most time-intensive portion of the AI pipeline, accounting for approximately 60% to 80% of total project time.[12] Here, 1x SSD reached 5,415 MB/s, while 16x HDDs reached just 632 MB/s.
SSD advantage: 8.6x the 16x HDD throughput (5,415 / 632 ≈ 8.6).

Data ingest: ETL ingest
This is write-heavy data loading into the data lake, like ingesting raw datasets from external sources.
Data ingest is a write-heavy (10/90 read/write) workload that loads raw datasets from external sources into the data lake. 1x SSD performance reached 3,649 MB/s while 16x HDD performance was 1,063 MB/s during testing. For organizations ingesting terabytes of new data daily, this bottleneck may directly limit how quickly new data becomes available for training.
SSD advantage: 3.4x the 16x HDD throughput (3,649 / 1,063 ≈ 3.4).

Bulk data transfer: FIO 4MB and 256KB
Large-block I/O models dataset uploads, checkpoint saves, and data migration between storage tiers. This type of I/O is often considered an optimal case for HDDs, in which they may approach peak throughput. For bulk operations like dataset migration between storage tiers, model checkpoint saves, and large-file transfers, 1x SSD demonstrated significantly higher throughput.
SSD advantage: 4MB: 1x SSD: 3.1x the 16x HDD throughput (7,679 / 2,510 ≈ 3.1); 256KB: 1x SSD offered about 10.2x the 16x HDD throughput (5,343 / 525 ≈ 10.2).

The performance results detailed above show that storage media characteristics may be a concern if they propagate through the pipeline. When the lake can sustain higher throughput and lower tail latency, upstream compute spends less time stalled, and downstream stages can start sooner. In practice, this changes how the platform scales, enabling smaller "drive count to hit bandwidth," with fewer parallel spindles to manage and fewer moving parts to keep steady under mixed I/O.

Storage power efficiency analysis

Performance is only half of the problem. Fixed facility power envelopes and limited grid expansion increasingly constrain AI infrastructure. The question becomes how much data lake work you can deliver per provisioned watt.

Industry and policy guidance increasingly treat energy efficiency as a gating factor for AI deployment and scaling, not a "nice to have," which makes power efficiency the next logical lens for interpreting the 1x SSD versus 16x HDD results.

Figure 3 shows a consistent, order-of-magnitude efficiency advantage for 1x SSD. Note that the x-axis in this figure is on a log scale.

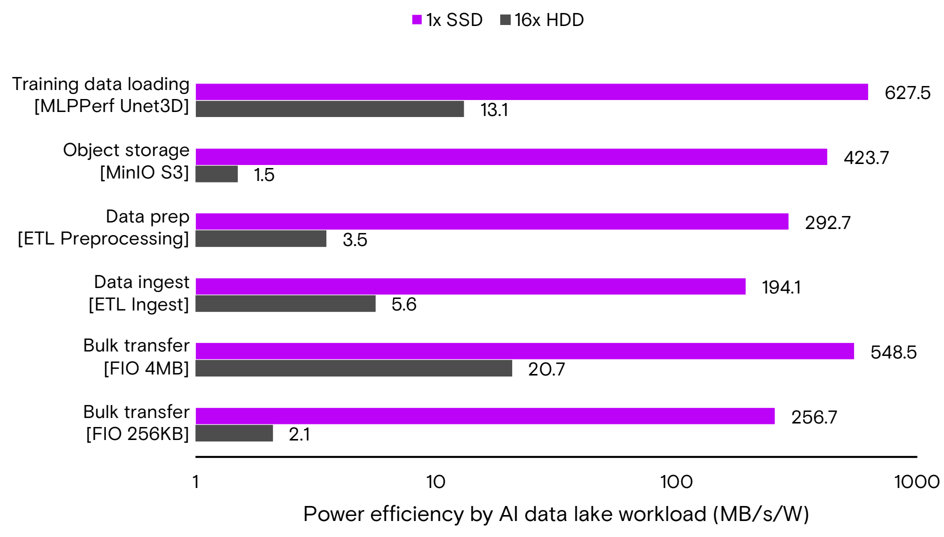

Figure 3: AI data lake workload power efficiency (log scale)

Across all measured workloads, 1x SSD delivered 26x to 282x the power efficiency of the 16x HDD, across the access patterns and application stages.[13]

11. "Data lake zones and containers" (learn.microsoft.com).
12. "Introducing: Power BI data prep with dataflows" (powerbi.microsoft.com).
13. Bulk transfer [FIO 4MB]: 548.5 / 20.7 ≈ 26; Object storage [MinIO S3]: 423.7 / 1.5 ≈ 282.

This consistent leadership indicates a fundamental architectural advantage, not a workload-specific effect.

There is also a clear separation and no efficiency overlap. The 1x SSD lowest-observed power efficiency (194.1 MB/s/W for data ingest [ETL ingest]) exceeded the 16x HDD highest (20.7 MB/s/W see on bulk transfer [FIO 4MB]) by more than 9x, further distinguishing the two storage architectures.[14]

Conclusion

Across every tested stage of the AI data lake pipeline, a single 245TB Micron 6600 ION SSD delivered dramatically superior performance compared to 16x 16TB HDDs. The advantages range from 3.1x (bulk transfer, FIO 4MB) to 45.6x (object storage, MinIO S3 mixed) in throughput, with 26x to 282x improvement in power efficiency.

Three critical findings emerged:

HDDs did not sustain peak throughput under real AI workloads: Only two of six workloads (MLPerf Unet3D (2,906 MB/s for the array ≈ 181 MB/s for each HDD) and FIO 4MB (2,510 MB/s for the array ≈ 157MB/s for each HDD) allow HDDs to operate anywhere near their measured peaks (still well below their rated individual HDD throughput of 261+ MB/s).

SSDs help eliminate storage as a pipeline bottleneck: With up to 12,738 MB/s per drive, SSDs help ensure that storage isn't the limiting factor in AI training, data preparation, or data access.

The SSD advantage accelerates with real-world complexity: As workloads involve smaller blocks and mixed I/O, the SSD performance advantage reaches up to 45.6x.

For AI data lake deployments, the choice is clear: Micron 6600 ION NVMe SSDs deliver the throughput and power efficiency required to keep modern AI pipelines running at full speed.

14. Calculated as (191.4 / 20.7) = 9.4 ≈ 9.

micron.com/6600-ION