One Micron 6600 ION SSD or 16 HDDs in the spin cycle? A throughput, latency, and power efficiency wakeup call.

Rethinking storage foundations

Modern data center architectures are under pressure as capacity scales faster than performance. As HDDs grow denser, mechanical limits can increasingly constrain throughput, latency, and power efficiency, forcing architects to rethink how storage tiers support AI, analytics, and large-scale data platforms.

This paper presents a detailed comparison between a single 245TB Micron 6600 ION NVMe™ SSD and a similar-capacity 256TB HDD array (16x 16TB Seagate® EXOS® HDDs).[1] Using controlled Flexible I/O Tester (fio)[2] benchmarks, throughput, latency, power efficiency, per-TB performance, and rack-scale density are evaluated across multiple block sizes and read/write mixes. Across all tested workloads herein, block sizes, and concurrency levels, the 1x SSD leads in throughput, latency, and power/energy efficiency.

The sections that follow analyze these behaviors across the representative workloads in Table 1 and highlight how storage characteristics shape performance, efficiency, and scalability under real-world access patterns.

Workload mapping

We represent different application profiles using four read/write ratio mixes and two request sizes. All tests use fio **randrw** (random read/write) access, with 256KB representing moderate-size application requests and 4MB representing large-block transfers commonly seen in streaming, backup/restore, and large-file workflows. Note: Block sizes are shown as KB/MB for readability; fio was configured with 256KB and 4MB (Table 2). Terminology: In this brief, read-heavy refers to 70/30 R/W, while read-dominant refers to 90/10 R/W. Table 1 summarizes how these read/write mixes and request sizes map to common use cases.[3]

Common use case	Workload type	Block size[4]	Workload	Read %	Write %
Logging, moderate-sized ingest	Write-heavy	256KB	30/70 R/W	30%	70%
Large-sized file ingestion		4MB	30/70 R/W	30%	70%
General-purpose server	Balanced	256KB	50/50 R/W	50%	50%
Media processing		4MB	50/50 R/W	50%	50%
Analytics, reporting	Read-heavy	256KB	70/30 R/W	70%	30%
Streaming, restoring backups		4MB	70/30 R/W	70%	30%
Content delivery, caching	Read-dominant	256KB	90/10 R/W	90%	10%
		4MB	90/10 R/W	90%	10%

Table 1: Workload mapping

1. All capacity statements refer to rated capacity; formatted capacity will be less; 1GB = 1 billion bytes. The single 245TB Micron 6600 ION SSD configuration is abbreviated “1x SSD,” and the 256TB HDD array (16x 16TB HDDs) configuration is abbreviated “16x HDD” for brevity. We chose 16x 16TB HDDs because they yield an array size that closely matches the 1x SSD size and provide an architectural building-block analysis, ensuring a reasonable number of HDD spindles are used. Different HDDs may yield different results.
2. See “fio – Flexible I/O tester rev. 3.4.1 Overview and history” (fio.readthedocs.io). Test configuration summarized in Table 2.
3. Block sizes and I/O patterns are based on Micron Data Center Workload Engineering research.
4. Block sizes are in MB and throughput in MB/s per the unit convention in the Appendix (fio reports in base 2; all data was converted to base 10 (SI) before calculations).

Write-heavy: Logging and ingest (30/70 R/W)

I/O profile: 256KB (logging, ingest); 4MB (large file ingest). This profile represents logging, ingestion, and backup creation. It is akin to a continuous, append-oriented stream from applications, infrastructure, security systems, or sensors.

Performance (Figure 1)

256KB: The 1x SSD scaled well: at a 256KB block size, it climbed (1,687→4,168) MB/s as numjobs rose (1→8, horizontal axis),[5] while the 16x HDD barely scaled (235→735) MB/s.

Details: Low concurrency (1,687 / 235) ≈ 7.2x. High concurrency (4,168 / 735) ≈ 5.7x.

4MB: 1x SSD rose (2,854→4,068) MB/s while the 16x HDD varied little (2,652→2,671) MB/s.

Details: Low concurrency (2,854 / 2,652) ≈ 1.1x. High concurrency (4,068 / 2,671) ≈ 1.5x.

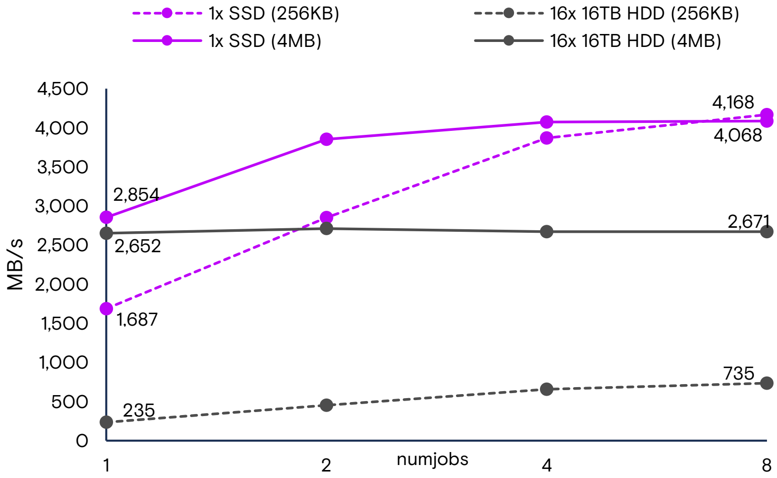

Figure 1: Write-heavy workload performance

Power efficiency (Figure 2)

MB/s per watt: How efficiently storage converts power into usable data movement in real time. At 256KB and 4MB, the 1x SSD stayed essentially constant (150.5→145.6), while the 16x HDD varied (3.1→11.7). That yields (150.5 / 3.1) ≈ 48x at 256KB and (145.6 / 11.7) ≈ 12x at 4MB.

TB/kWh: The volume of data processed per unit of electrical energy. At 256KB, the 1x SSD delivered 541.7 TB/kWh versus 11.2 for the 16x HDD (541.7 / 11.2 ≈ 48x difference); at 4MB, it again led 521.9 versus 42.2 (521.9 / 42.2 ≈ 12x difference).

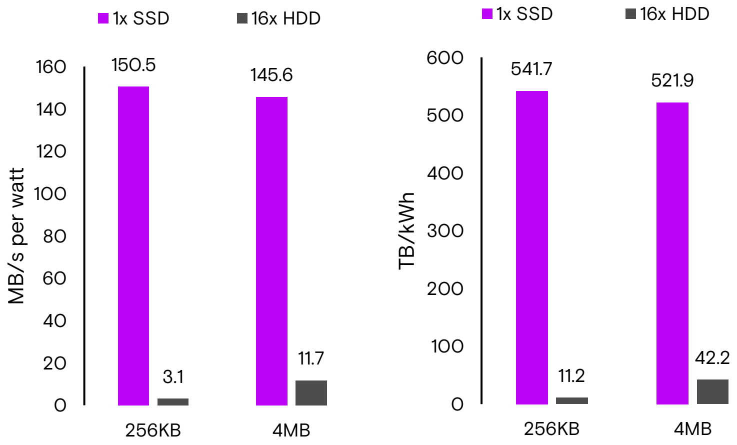

Figure 2: Write-heavy workload power and energy efficiency

Read latency (Figure 3)[6]

256KB: At 256KB, 1x SSD showed 0.89ms while 16x HDD reached 1,607ms, a difference of (1,606.80 / 0.89) ≈ 1,808x.

4MB: At 4MB, SSD latency stayed low at 3.18ms, whereas HDD latency was quite high at 1,850.90ms, a difference of (1,850.90 / 3.18) ≈ 582x.

Write latency (Figure 3)

256KB: The 1x SSD maintained millisecond-scale latency under load, while the HDD remained roughly two orders of magnitude higher. At 256KB, 1x SSD showed a latency of 0.29ms, while 16x HDD showed 298.60ms, a difference of (298.6 / 0.29) ≈ 1,030x.

4MB: At 4MB, 1x SSD was 10.19ms while 16x HDD reached 3,770.40ms, a difference of (3,770.40 / 10.19) ≈ 370x.

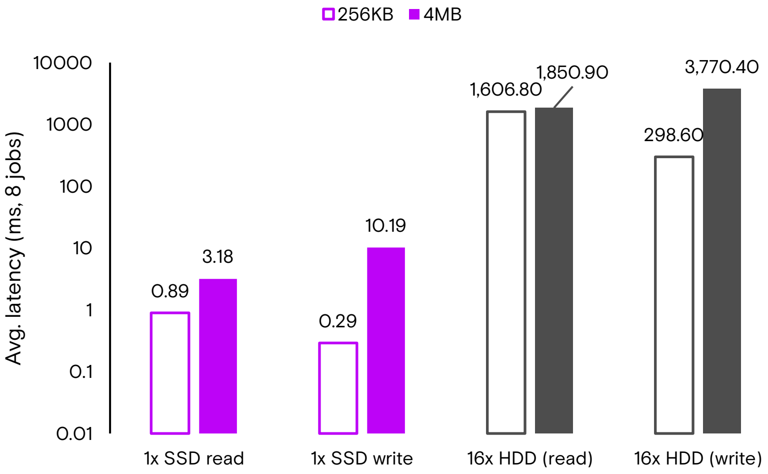

Figure 3: Write-heavy workload latency

---

[5] The fio parameter "numjobs" controls the number of parallel job instances fio launches to generate load. Increasing numjobs scales the workload by running multiple independent threads or processes with identical parameters. See "fio - Flexible I/O tester rev. 3.41" (fio.readthedocs.io) for more information.

[6] Latency is defined as fio-reported average I/O completion latency (submission-to-completion), in ms, tied to the same workload definitions and numjobs. Benchmark methodology and fio parameters are listed in Table 2. Reported latencies are averages.

Balanced: General purpose servers and media processing (50/50 R/W)

I/O profile: 256KB general purpose; 4MB media processing. This profile represents a common operating mode for shared repositories and multi-tenant services.

Performance (Figure 4)

256KB: The 1x SSD scales from 1,418→4,636 MB/s, while the 16x HDD rose from 170→654 MB/s.
Details: Low concurrency (1,418 / 170) ≈ 8.3x. High concurrency (4,636 / 654) ≈ 7.1x.

4MB: Throughput again diverged with parallelism for 1x SSD: it scaled from 3,236 to 5,254, while 16x HDD was nearly flat (2,519→2,516).
Details: Low concurrency (3,236 / 2,519) ≈ 1.3x. High concurrency (5,254 / 2,516) ≈ 2.1x.

Power efficiency (Figure 5)

MB/s per watt: Block size helped improve HDD performance by 16x, but it did not improve power efficiency. The 1x SSD stayed high (175.8→189.6 MB/s per watt for 256KB and 4MB, respectively). 16x HDD rose slightly (2.7→11.01 MB/s per watt) yet remained lower by (175.8 / 2.7) ≈ 65.1x at 256KB and (189.6 / 11.01) ≈ 17.2x at 4MB.

TB/kWh: Block size helped 16x HDD, but it didn't close the gap: 1x SSD delivered 632.2→682.8 TB/kWh (for 256KB and 4MB, respectively), while 16x HDD showed 9.8→34.1 TB/kWh. That corresponds to (632.2 / 9.8) ≈ 64.5x at 256KB and (682.8 / 34.1) ≈ 20.0x at 4MB.

Read latency (Figure 6)

256KB: At 256KB, 1x SSD read latency was 0.7ms, while 16x HDD read latency reached 1,373.00ms, a difference of (1,373.00 / 0.7) ≈ 1,961x.

4MB: At 4MB, 1x SSD read latency increased to 4.07ms, whereas 16x HDD read latency rose to 2,129.30ms. Even with larger transfers, SSD maintained a latency advantage of (2,129.30 / 4.07) ≈ 523x, indicating HDD read latency still escalated sharply at scale.

Write latency (Figure 6)

256KB: For 256KB write, 1x SSD write latency was 0.14ms, compared to 198.4ms for the 16x HDD, a difference of (198.4 / 0.14) ≈ 1,420x.

4MB: At 4MB, 1x SSD write latency rose to 8.48ms, while 16x HDD write latency surged to 4,659.50ms. Although block size increased latency for both, 16x HDD remained (4,659.50 / 8.48) ≈ 550x the 1x SSD value. This showed that larger I/O amplification increased HDD write latency far more than it did for SSDs.

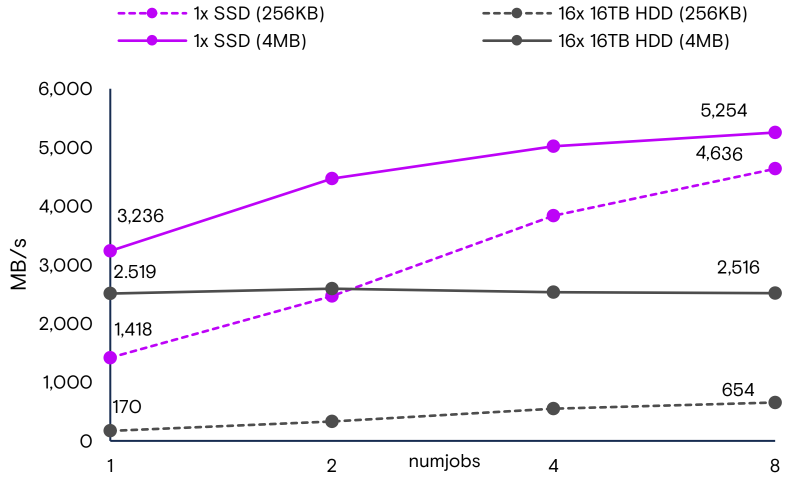

Figure 4: Balanced workload performance

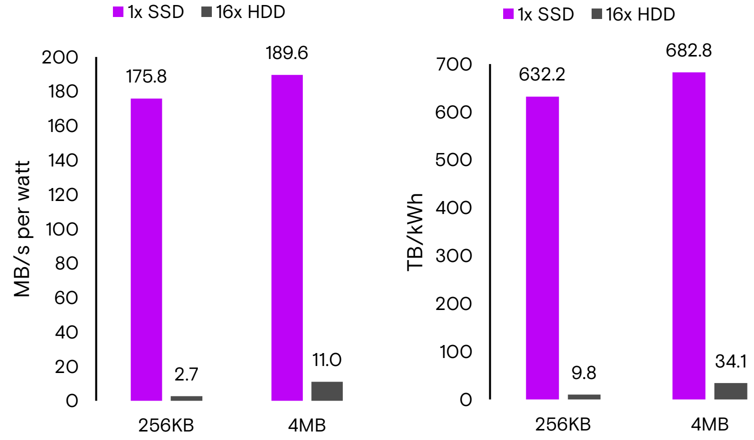

Figure 5: Balanced workload power and energy efficiency

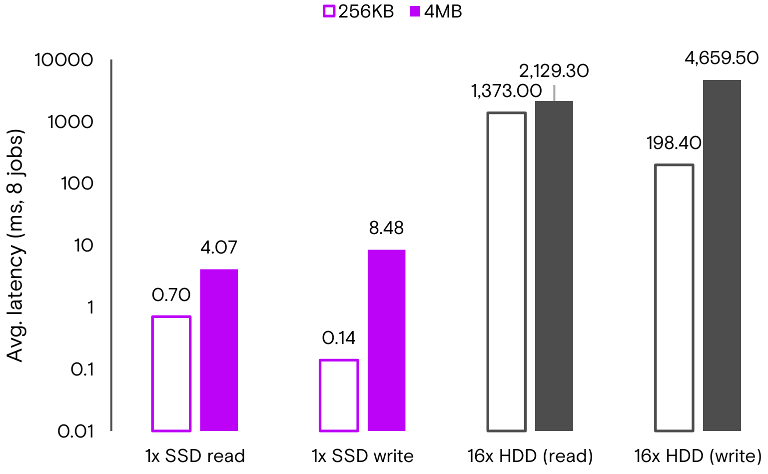

Figure 6: Balanced workload latency

Read-heavy: Analytics, reporting, streaming, and backup (70/30 R/W)

I/O profile: 256KB (analytics, reporting), 4MB (streaming, reading backups with updates). It can be thought of as an analytics read stream with a small percentage of simultaneous updates or inserts.

Performance (Figure 7)

256KB: 1x SSD scaled from 1,255→5,216 MB/s, while the 16x HDD grew 126→386 MB/s.

Details: Low concurrency (1,255 / 126) ≈ 10x. High concurrency (5,216 / 386) ≈ 13.5x.

4MB: 1x SSD delivered 2,449→3,971 MB/s versus 2,379→2,448 MB/s for 16x HDD.

Details: Low concurrency (2,449 / 2,379) ≈ 1.0x (essentially the same). High concurrency (3,971 / 2,448) ≈ 1.6x.

Power efficiency (Figure 8)

MB/s per watt: Across 256KB→4MB, the 1x SSD holds a clear orders-of-magnitude advantage (221.2→137.4 MB/s per watt) over the 16x HDD values (2.4→10.7). That is (221.2 / 2.4) ≈ 92.2x at 256KB and (137.4 / 10.7) ≈ 12.8x at 4MB.

TB/kWh: At 256KB, the 1x SSD showed 795.7 vs 8.6 (795.7 / 8.6 ≈ 92.5x), and at 4MB, that comparison was 494.7 vs 38.6 (494.7 / 38.6 ≈ 12.8x).

Read latency (Figure 9)

256KB: The 1x SSD read latency was very low at 0.50ms, while the 16x HDD array was 1,222.20ms, a difference of (1,222.20 / 0.50) ≈ 2,444x.
4MB: 1x SSD read latency remained single-digit, at just 4.32ms, while the 16x HDD reached 3,310.60ms, a difference of (3,310.60 / 4.32) ≈ 766x.

Write latency (Figure 9)

256KB: 1x SSD maintained sub-millisecond latency of 0.06ms, while the 16x HDD reached 123.6ms, a difference of (123.6 / 0.06) ≈ 2,060x.
4MB: The gap was again clear: 1x SSD 4MB write latency was 17.78ms, while 16x HDD reached 3,907.50ms, a difference of (3,907.50 / 17.78) ≈ 220x.

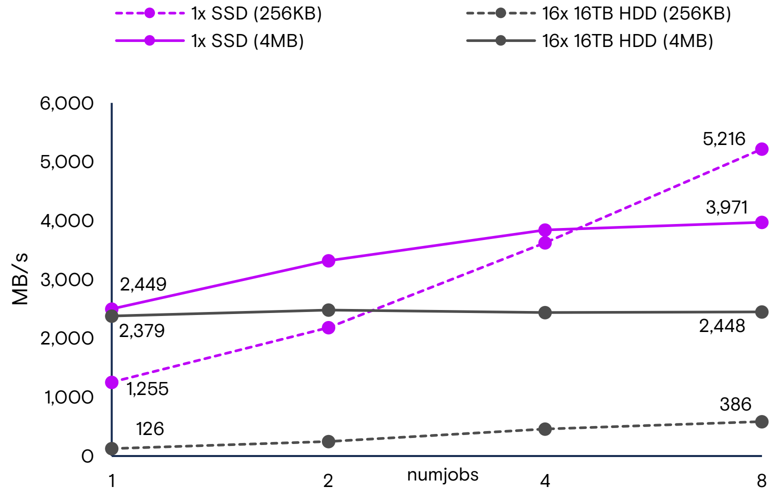

Figure 7: Read-heavy workload performance

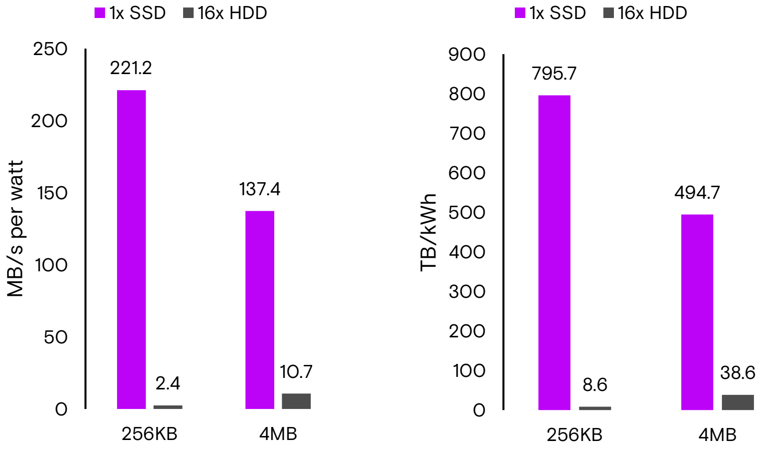

Figure 8: Read-heavy workload power and energy efficiency

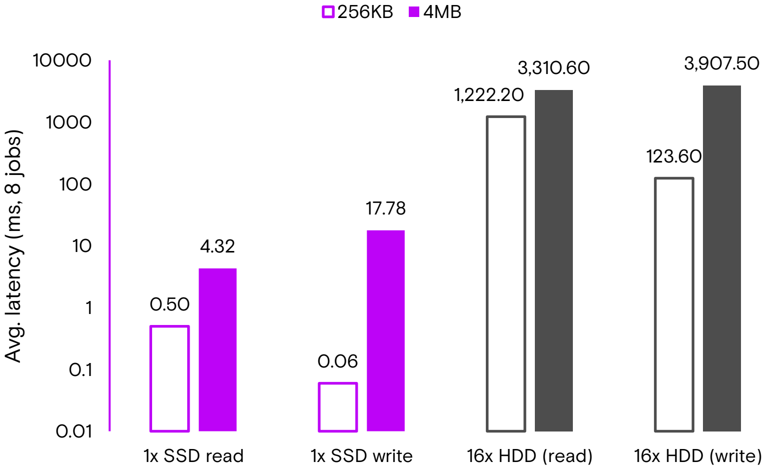

Figure 9: Read-heavy workload latency

Read-dominant: Content delivery and caching (90/10 R/W)

I/O profile: 256KB (content delivery, cache read), 4MB (content delivery). This is mostly read and can be thought of as a streaming service that shows popular broadcasts, with only occasional updates when new content is added.

Performance (Figure 10)

256KB: The 1x SSD performance climbed steadily (1,106→5,343) MB/s, while the 16x HDD also climbed, but not as steeply (98→525 MB/s).

Details: Low concurrency: (1,106 / 98) ≈ 11.3x. High concurrency (5,343 / 525) ≈ 10.2x.

4MB: 1x SSD grew from 3,645→7,679 MB/s, while the 16x HDD grew little (2,276→2,510) MB/s.

Details: Low concurrency (3,645 / 2,276) ≈ 1.6x. High concurrency (7,697 / 2,510) ≈ 3.1x.

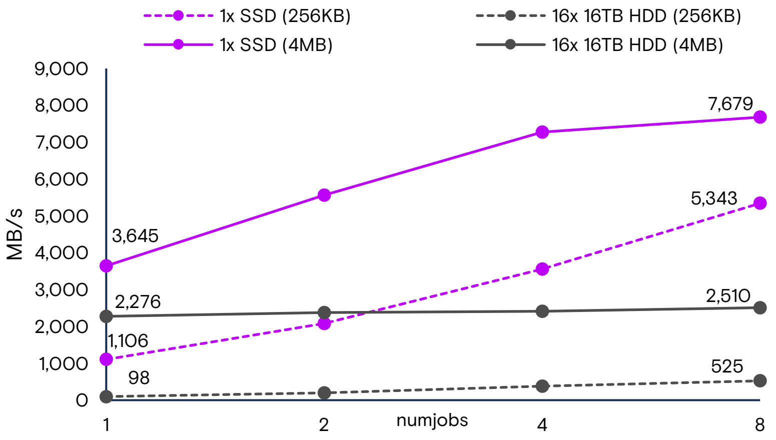

Figure 10: Read-dominant workload performance

Power efficiency (Figure 11)

MB/s per watt: At 256KB, the 1x SSD showed over a two-order-of-magnitude difference compared to the 16x HDD MB/s per watt values (256.7 / 2.1 ≈ 122x). At 4MB, that difference was (286.3 / 11.1) ≈ 25.8x.

TB/kWh: At 256KB, the 1x SSD showed 933.7 vs 7.7 (933.7 / 7.7 ≈ 121.3x), and at 4MB, that comparison was 1,058.4 vs 40.0 (1,058.4 / 40.0 ≈ 26.5x).

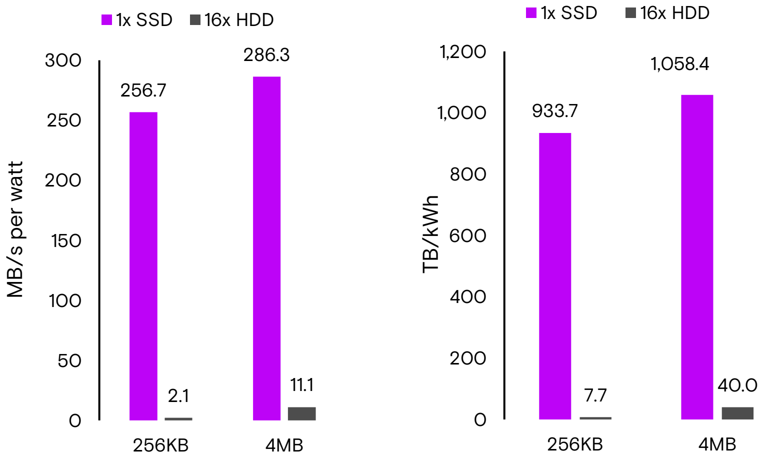

Figure 11: Read-dominant workload power and energy efficiency

Read latency (Figure 12)

256KB: At 256KB, 1x SSD read latency remained extremely low at 0.4ms, while the 16x HDD showed 1,114.50ms. This represented a latency advantage of (1,114.50 / 0.4) ≈ 2,786x for SSDs and illustrated how small-block reads dramatically penalize HDD arrays under parallel load in this testing.

4MB: At 4MB, 1x SSD read latency was 4.29ms, while 16x HDD read latency reached 3,544.30ms. Even with larger transfers, SSD maintained a decisive advantage, delivering (3,544.30 / 4.29) ≈ 826x lower latency than HDD at the same concurrency.

Write latency (Figure 12)

256KB: For 256KB writes, 1x SSD achieved 0.05ms latency, whereas the 16x HDD recorded 74.7ms, a difference of (74.7 / 0.05) ≈ 1,494x.

4MB: At 4MB, 1x SSD write latency increased to 4.24ms, while 16x HDD write latency escalated to 2,157.30ms. Although both devices exhibited higher latency at larger block sizes, SSD remained (2,157.30 / 4.24) ≈ 509x lower, indicating that block size increased the 16x HDD latency far more than it did the 1x SSD latency.

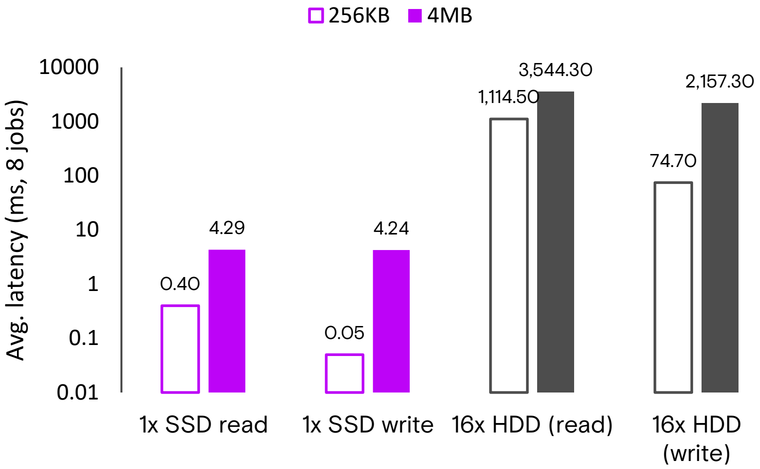

Figure 12: Read-dominant workload latency

Conclusion

Across representative workload mixes, the practical differentiator is not a single metric. It is both simpler and more complex than that.

It’s whether a storage architecture exhibits bounded behavior by design, remaining predictable under contention and mixed I/O. Real-world demand is getting more complex, messier, and more power-constrained. That means more threads, more jobs, more competing demands, and more simultaneous reads and writes. Architectures that assume best-case access patterns may pay for it later in scheduling complexity, inflated safety margins, and expensive “just in case” overbuild.

Long-term design implications must anchor architectural decisions in workload envelopes, not device datasheets. Design around the worst-case credible concurrency and mixed I/O conditions, then choose the architecture and its essential elements to keep latency predictable, throughput quick, and power efficiency within business bounds across that envelope.

Storage isn’t a passive component—it governs platform behavior in many ways. Media characteristics propagate upward as pacing, queues, and variability, which the software stack treats as instability and compensates for with retries, buffering, and overprovisioning.

Consistent storage performance enables predictable system-level behavior, allowing compute resources to sustain high utilization and scaling decisions to be capacity-driven rather than compensatory. From an architectural perspective, storage should contribute to platform efficiency rather than define its operational limits.

At rack scale, that SSD performance, power and energy efficiency, and consistency compound into higher utilization and better performance-per-watt.

micron.com/6600-ION

Appendix: fio parameters

All benchmarks were performed using fio (Flexible I/O Tester) with identical parameters across both storage configurations. Each workload was run five times, and the results presented are the average of all five runs.

Parameter	Value
fio version	3.41
I/O engine	io_uring
Block sizes (base 2, fio reported)	256 KiB, 4 MiB
Block sizes (base 10, converted)	256KB, 4MB
I/O pattern	Random Read/Write (randrw)
Direct I/O	Yes (O_DIRECT)
I/O depth	64
Ramp time	300 seconds (5 minutes)
Runtime	600 seconds (10 minutes)
Number of jobs	1, 2, 4, 8
Runs per configuration	5
Reporting method	Average of 5 runs

Table 2: fio parameters used

Unit convention: fio reports throughput in MiB/s (mebibytes per second, base-2). All throughput values in this report have been converted to SI MB and KB (megabytes per second, base-10) for readability by multiplying by 1.048576 (1 MiB = 1,048,576 bytes = 1.048576 MB). Similar conversions were performed for KiB. GB/s values use base-10 gigabytes.