

Nirvana ABS vs AWS EBS Benchmarking Report

Evaluation of Block Storage Performance for Cold-Read Analytics and Storage Behavior

Executive Summary

Performance Highlights

vs AWS EBS gp3

10-14x faster

on cold-read analytical queries

vs AWS EBS io2

77% faster

on IO-heavy analytical queries

Mixed-workload throughput

5× higher

20,000 sustained IOPS per VM
(80/20 read/write)

Sequential Throughput

408-430 MB

per second per VM

Assessment

Nirvana ABS delivers a clear price-performance edge, outpacing gp3 on cold reads and matching io2-class behavior at far lower cost, making it a strong default for latency-sensitive, analytics, and blockchain workloads.

Key Findings

Per Each Test



Cold-Read Analytical Performance (ClickBench)

Under forced cold-read conditions, ABS delivers ~10–14× median speedup vs gp3 on IO-heavy analytical queries while tracking closely with EBS io2, operating within 1.5× on ~77% of queries and within 2× on ~98%.



Infrastructure IO Behavior (FIO)

ABS performs strongly in real-world operating regimes, sustaining 20,000 guaranteed IOPS per VM at 4K+ IO sizes with low-millisecond latency. Sequential throughput stabilizes at ~408–430 MB/s per VM, indicating predictable performance under sustained workloads.



Price-Performance

At a normalized 1 TB volume sustaining ~10,000 IOPS, ABS costs ~\$95.7/month vs ~\$116.9 for gp3 and ~\$778 for io2, delivering ~18% lower cost than gp3 and ~8× lower cost than io2.

Benchmarking Objective

This report evaluates the performance of Nirvana Labs Accelerated Block Storage (ABS) against AWS EBS gp3 and io2, assessing their suitability for modern Web3 workloads that demand low latency, sustained throughput, predictable cold-read behavior, and strong mixed-workload IOPS.

To provide a comprehensive view, we conducted two complementary benchmark classes and compared pricing across the 3 products:

- ClickBench Hardware

Measure end-to-end analytical query performance under forced cold-read conditions.

- FIO synthetic tests

Measure raw block storage performance independent of any application layer, isolating the storage path from database or caching effects.

- Price-Performance

Compare monthly cost among Nirvana ABS, AWS EBS gp3, and AWS EBS io2.



NIRVANA LABS

ClickBench Test

ClickBench Test

Purpose

ClickBench Hardware is used to measure end-to-end analytical query performance under forced cold-read conditions, where storage latency and throughput become the primary determinants of performance.

The benchmark executes 43 standardized OLAP queries (Q0–Q42) against the hits_100m dataset and records per-query runtime. By explicitly dropping the OS page cache before each query, the test removes memory caching effects and exposes the true behavior of the underlying storage system.

In practical terms, the benchmark answers a single operational question:

How fast can each storage platform execute real analytics when data is not served from cache?

This methodology makes ClickBench particularly effective for evaluating storage platforms used in:

- ClickHouse and OLAP analytics
- Large historical scans
- Data warehouse workloads
- Blockchain indexing and archival systems

Test Configuration

Parameter	AWS EBS gp3	AWS EBS io2	Nirvana ABS
Instance Type	r7a 16xlarge	r7a 16xlarge	Identical cores and architecture, lower clock speed
vCPU	64	64	64
Memory	512 GB	512 GB	512 GB
Network Performance	25 Gbps	25 Gbps	25 Gbps
Volume Size	50 GB	50 GB	50 GB
Storage Backend	EBS gp3	EBS io2	Nirvana ABS
Filesystem	ext4	ext4	ext4
Mount Configuration	Identical across tests	Identical across tests	Identical across tests
Benchmark Tool	ClickBench Hardware	ClickBench Hardware	ClickBench Hardware
Cache Policy	Cold-run (cache dropped per query)	Cold-run (cache dropped per query)	Cold-run (cache dropped per query)

Methodology Controls

To ensure fairness and reproducibility:

- Each environment was benchmarked twice
- Only the second full run was used
- Only cold-run timings were analyzed
- CPU and memory were held constant
- Results were compared on a per-query basis

These controls isolate storage performance as the primary variable.

Results Summary

ABS vs EBS gp3

Across the ClickBench cold-run suite:

- Median speedup (all queries): ~10.5×
- Median speedup (IO-heavy queries): ~14.1×
- Worst-case speedup: ~28×

AWS gp3 exhibits repeated performance cliffs on large scan and aggregation queries (e.g., Q20–Q22, Q31–Q34), where runtimes expand into tens of seconds. ABS maintains these same queries in the sub-second to low-second range, indicating materially stronger cold-read behavior.

Query	Nirvana ABS	AWS gp3
Q20 large scan	0.79 s	20.2 s
Q21 aggregation scan	0.84 s	22.8 s
Q22 heavy analytical query	1.48 s	41.9 s
Q23 large group aggregation	1.05 s	19.7 s

Interpretation: gp3 is optimized for general workloads but becomes a bottleneck under cold analytical access patterns.

ABS eliminates the major cold-read performance cliffs observed on gp3 and delivers materially more predictable analytical performance.




Chart notes

- Each bar represents the execution time of a ClickBench cold-read query
- Lower values indicate better performance
- Blue: AWS EBS gp3
- Red: Nirvana Labs ABS
- CPU, memory, and dataset were identical across tests
- OS page cache was dropped before each query

Results Summary

ABS vs EBS io2

iAWS io2 performs strongly as the premium storage tier.
ABS remains highly competitive:

- ~77% of queries run within 1.5× of io2
- ~98% run within 2×
- ABS is faster than io2 on a subset of workloads

Query	Nirvana ABS	AWS gp3
Q20	0.79 s	0.59 s
Q21	0.84 s	0.79 s
Q22	1.48 s	1.78 s
Q36	0.046 s	0.091 s

These results show that ABS frequently performs very close to io2 and occasionally faster depending on the workload.

Interpretation: ABS delivers io2-class analytical performance while targeting improved price-performance efficiency.

ABS delivers io2-class analytical performance with consistent cold-read behavior, positioning it as a strong price-performance alternative.




Chart notes

- Each bar represents the execution time of a ClickBench cold-read query
- Lower values indicate better performance
- Blue: AWS EBS io2
- Red: Nirvana Labs ABS
- CPU, memory, and dataset were identical across tests
- OS page cache was dropped before each query

Overall Query Performance

The heatmap below shows per-query execution time (seconds) across the ClickBench cold-run suite (Q0–Q42). Lower values indicate better performance and appear in green on the chart.

Across the full suite, Nirvana ABS significantly outperforms AWS gp3 and closely tracks AWS io2 performance.

- Columns compare storage backends across identical compute environments
- Color scale:
 - Green = faster execution
 - Yellow = moderate
 - Red = slower execution
- OS page cache was dropped before each query to expose true storage behavior
- Disk size was normalized to 50 GB across all tests

QUERY	AWS r7a. 16xlarge	AWS r7a. 16xlarge	Nirvana Labs	Nirvana Labs
	64 vCPU / 512GB RAM / 50GB EBS GP3	64 vCPU / 512GB RAM / 50GB EBS IO2	24 vCPU / 64GB RAM / 50GB ABS	64 vCPU / 512GB RAM / 50GB ABS
Q0	0.001	0.001	0.001	0.001
Q1	0.053	0.023	0.011	0.025
Q2	0.367	0.023	0.022	0.034
Q3	1.48	0.039	0.04	0.093
Q4	1.5	0.077	0.063	0.104
Q5	2.605	0.124	0.127	0.19
Q6	0.003	0.003	0.002	0.003
Q7	0.013	0.012	0.01	0.015
Q8	1.623	0.129	0.185	0.177
Q9	2.465	0.192	0.217	0.241
Q10	1.296	0.09	0.083	0.136
Q11	1.243	0.071	0.079	0.101
Q12	2.655	0.165	0.225	0.202
Q13	4.195	0.191	0.281	0.264
Q14	2.678	0.15	0.205	0.194
Q15	1.255	0.091	0.194	0.131
Q16	4.357	0.318	0.638	0.42
Q17	3.818	0.214	0.446	0.292
Q18	7.164	0.688	1.261	0.88
Q19	0.705	0.039	0.027	0.067
Q20	20.238	0.59	0.351	0.794
Q21	22.799	0.746	0.305	0.84
Q22	41.893	1.779	0.927	1.478
Q23	0.328	0.467	0.2	0.366
Q24	5.273	0.123	0.116	0.209
Q25	2.561	0.08	0.085	0.117
Q26	5.488	0.119	0.111	0.21
Q27	1.57	0.057	0.067	0.087
Q28	17.096	0.466	0.598	0.729
Q29	0.059	0.033	0.031	0.051
Q30	1.585	0.102	0.139	0.15
Q31	7.812	0.209	0.266	0.353
Q32	7.565	0.731	1.513	1.062
Q33	19.709	0.825	1.162	1.052
Q34	19.702	0.813	1.113	1.132
Q35	0.131	0.072	0.133	0.094
Q36	0.071	0.091	0.058	0.046
Q37	0.035	0.036	0.029	0.032
Q38	0.052	0.054	0.028	0.033
Q39	0.113	0.148	0.105	0.078
Q40	0.023	0.02	0.01	0.018
Q41	0.021	0.019	0.01	0.017
Q42	0.014	0.013	0.006	0.009

Validation Signal: CPU-Bound Query Convergence

Some ClickBench queries are CPU-bound, meaning their runtime is determined by computation rather than storage performance.

When this happens, all platforms should produce nearly identical runtimes if the compute environments are properly normalized.

The results below show that these queries execute in the same millisecond range across gp3, io2, and ABS, confirming that compute resources were comparable. Millisecond-scale queries converge across all platforms, confirming that compute environments were comparable and that the observed performance differences are driven by storage behavior.

Query	Nirvana ABS	AWS io2
Q20	0.79 s	0.59 s
Q21	0.84 s	0.75 s
Q22	1.48 s	1.78 s
Q36	0.046 s	0.091 s

Cold-Read Analytical Performance

Conclusion



01

ABS delivers ~10–14× median speedup vs AWS EBS gp3 on IO-heavy analytical queries under forced cold-read conditions.



02

ABS closely tracks AWS EBS io2 performance, operating within 1.5× on ~77% of queries and within 2× on ~98% of the ClickBench suite.



03

AWS gp3 exhibits severe cold-read performance cliffs on large scan and aggregation queries (e.g., Q20–Q22, Q31–Q34), where runtimes expand into tens of seconds.



04

ABS avoids these latency spikes, maintaining stable low-latency behavior even on the largest analytical scans.



05

CPU-bound queries converge across all platforms, confirming that compute environments were properly normalized and that storage performance drives the observed differences.

Overall, ABS materially reduces the cold-read bottlenecks seen on gp3 while delivering io2-class analytical performance across the majority of real workloads.



NIRVANA LABS

FIO Synthetic Storage Test

FIO Synthetic Storage Test

Purpose

FIO (Flexible I/O Tester) was used to measure the raw block storage behavior of ABS independent of any application layer. While ClickBench evaluates end-to-end database performance, FIO isolates the storage path itself and exposes the fundamental limits of the system.

The objective of this test suite was to understand how ABS behaves under controlled I/O patterns that simulate common infrastructure workloads. Specifically, the testing evaluated:


- Achievable IOPS ceilings
- Latency characteristics across block sizes
- Bandwidth limits per VM
- Sensitivity to very small I/O
- Scaling under mixed read/write pressure

This allows us to separate true storage behavior from database or caching effects.

FIO SYNTHETIC

Test Environment and Methodology

The primary test was executed across a 16-VM cluster using LUNs mounted via iSCSI through the Portworx CSI driver. This intentionally exercises the full production data path:



All tests used direct I/O to bypass OS page cache and expose true storage performance.

Key configuration

- **ioengine:** libaio
- **Direct I/O:** enabled
- **Numjobs:** 4 per VM
- **iodepth:** 32
- **Runtime:** 15 seconds (baseline) with extended steady-state validation runs
- **Block sizes:** 512B through 1M
- **Workload mix:** read, write, and 80/20 mixed

In addition to the distributed cluster runs, two validation passes were performed:

- **Post-tuning reruns to verify stability**
- **Single-VM isolation tests to measure raw per-node capability**

Results Overview

Small Random Writes ($\leq 2K$)

Overall

- Small-write sensitivity exists
- It is workload-specific
- Unlikely to impact most Web3 or database deployments

Initial Small-Write Results

- Very small random writes produced $\sim 2.9K$ – $3.0K$ IOPS per VM
- Latency observed: ~ 43 – 45 ms

Rerun Results

- After longer steady-state runs, performance improved up to $\sim 85\%$ (512B case)
- Indicates no hard architectural ceiling
- Early results were partially influenced by warm-up effects

Why This Is Not Representative of Production

- Most databases batch writes
- RocksDB flushes at $\sim 4K$ or larger
- Blockchain clients rarely sustain 512B write patterns

Why This Happens

- Extremely small network writes become packet-rate bound rather than media-bound
- The system spends more time processing interrupts and packets than writing to flash
- This is a known characteristic of network-attached block storage, not a limitation of the flash media

Results Overview

4K Random Writes

Performance improved dramatically once block size reached 4K

System delivered >20K IOPS per VM

Observed latency: ~3.5 ms

Implication for Real Workloads

- ABS performs strongly in the most relevant operating range
- Well suited for:
 - Databases
 - Blockchain validators
 - State engines

Why 4K Matters?

- 4K aligns with common database page sizes
- Matches typical blockchain state update sizes
- At this size:
 - Packet overhead is amortized
 - Queue efficiency improves
 - The storage pipeline becomes fully utilized

Results Overview

Mixed Workload (80/20 Read/Write)

The mixed 4K workload is the most representative of real production environments. Under this pattern, ABS achieved approximately:


- Guaranteed 20K IOPS (+bursting enabled for short periods) per VM
- ~1.3 ms average latency
- ~6.6 GB/s aggregate throughput

This is a strong result and reflects the behavior of real systems such as:

- Blockchain state databases
- Indexing pipelines
- Analytical ingestion workloads
- RocksDB compaction

Because reads are highly efficient and writes are partially amortized, the system sustains high throughput under realistic pressure.

This is the most important positive signal in the FIO suite.



Sequential Throughput Behavior

408-430 MB/s per VM

Consistent throughput plateau in large-blocks sequential tests



Observation

The curve flatlines cleanly rather than showing erratic throttling.



Interpretation

- Stable per-VM data-patch ceiling
- Not flash saturation or burst-credit behavior



Operational Implication

- Predictable and repeatable bandwidth delivery
- More valuable in production than short-lived peak bursts.

Post-Tuning Validation

Follow-up testing was performed to verify whether earlier small-write behaviour represented a structural limitation or simply a matter of tuning and runtime conditions. The results were conclusive.

IOPS Improvement

Significant gains observed at sub-4K block sizes once the pipeline was properly saturated during longer runtime tests.

Reduced Latency

Steady-state latency decreased measurably, confirming that runtime duration and queue depth materially influence observed behaviour.

No Architectural Caps

No evidence of hard architectural limits was found. ABS benefits from steady-state load and behaves consistently under sustained pressure.



Small sequential writes improved substantially once the pipeline was fully saturated confirming that earlier limitations were operational, not structural.

Single-VM Isolation Results

To separate cluster effects from raw storage capability, additional tests were executed on a single VM. This isolated environment removed the influence of networked paths and virtualisation layers, exposing the true performance ceiling of the underlying ABS media and backend.

20K

Guaranteed IOPS

At 4K block size, with burst capability enabled for short periods above baseline.

2.2ms

Average Latency

Across typical IO patterns, with sub-millisecond latency achieved at smaller IO sizes.

GB/s

Multi-GB/s Bandwidth

Achieved at larger block sizes, confirming strong sequential throughput under sustained load.

These results confirm that the underlying ABS media and backend are not the primary constraint. Distributed cluster results primarily reflect the realities of networked storage paths and virtualisation layers, rather than any fundamental storage weakness.

Workload Implications

Mapping FIO benchmark patterns to real-world operations reveals a strong alignment between ABS performance characteristics and the demands of modern infrastructure.



Block Proposal & WAL Writes

4K write workloads operate safely within the efficient performance regime, making ABS well-suited for write-ahead logging and block proposal tasks.



Database Compaction & Mixed Workloads

Compaction and mixed I/O patterns perform strongly. Initial sync throughput is bounded by the per-VM bandwidth ceiling but remains within expected ranges for network block storage.



Tiny Packet Workloads

Sub-4K packet workloads are less efficient, as expected for network-attached storage. In practice, this is mitigated by batching strategies implemented in real blockchain clients.

Overall, the storage profile aligns well with the I/O patterns that dominate modern Web3 infrastructure where $\geq 4K$ block sizes and sustained mixed I/O are the norm.

Conclusion

Across all test conditions, ABS demonstrates strong, stable, and predictable performance that aligns with the real-world I/O demands of database, AI, blockchain, analytical workloads.

- **Strong 4K and Larger Block Performance**

Delivers stable performance aligned with database and blockchain I/O patterns, sustaining ~20,000 IOPS per VM under 80/20 read/write conditions roughly 5× typical per-volume limits.

- **Predictable Sequential Throughput**

Sequential bandwidth reaches ~408–430 MB/s per VM, demonstrating well-bounded and consistent behaviour under sustained load.

- **Sub-4K Sensitivity is Expected**

Tiny write sensitivity is a common characteristic of network-attached block storage driven by packet-rate and interrupt overhead not a storage deficiency.

- **Backend Capability Confirmed**

Single-VM validation proves the ABS backend is capable and consistent. Observed cluster limits reflect networked storage path realities, not fundamental weakness.




NIRVANA LABS

Price Comparison

Price Comparison

We compared Nirvana ABS against AWS gp3 and AWS io2 using a normalized workload: a 1 TB volume sustaining 10,000 IOPS — representative of common blockchain and database workloads.

The comparison highlights effective monthly cost, including baseline performance, and metering model behavior. As shown in the table, ABS bundles a much higher baseline IOPS and avoids step-function pricing, resulting in more predictable and often lower total cost at production scale.



Cost Breakdown

Feature	ABS	AWS gp3	AWS io2
Storage price	\$0.0935/GB-month	\$0.08/GB-month	\$0.125/GB-month
Baseline IOPS	20,000 (included)	3,000	N/A
Metering model	Unlimited	Capped	Provisioned
Cost above baseline	TBC	\$0.005 / IOPS-month	Tiered: 0-32K: \$0,065 32k-64K: \$0,046 >64K: \$0,032
Max IOPS	600,000	80,000	256,000
Throughput	300 MB/s included	125 MB/s free \$0.040 / MB/s-month over	Tied to IOPS
Performance model	Sustained	Burst-based	Provisioned
1TB volume - 10,000 sustained IOPS (monthly cost)	~\$95.7	~\$116.9	~\$778.0
Breakdown	Storage 1,024 x \$0.00935 = \$95.744 IOPS Included (20,000 baseline) Total: \$95.7 / mo	Storage 1,024 x \$0.08 = \$81.92 IOPS Included: 3,000 Extra IOPS needed: 10,000 - 3,000 = 7,000 Extra IOPS cost: 7,000 x \$0.005 = \$35.00 Total: \$81,92 + \$35,00 = \$116,92 ≈ 116,9 / mo	Storage 1,024 x \$0.125 = \$128.00 IOPS 10,000x\$0.065 (0-32K tier) = \$650.00 Total: \$128.00 + \$650.00 = \$778.00 / mo Note: excludes required Nitro-based compute costs

Conclusion

ABS delivers the strongest price-performance

~18%

cheaper than gp3

~8x

cheaper than io2

At a normalized workload of:

1TB storage + ~10,000 sustained IOPS

Performance Parity Without Extra Compute

- ABS maintains comparable analytical performance to io2 without requiring higher-tier Nitro instances.
- Avoids \$250-\$500+ /month compute overhead needed to fully realized provisioned IOPS performance.

Final Assessment


Nirvana ABS

Price-Performance Leadership for Modern Infrastructure

Based on combined ClickBench and FIO results, Nirvana ABS demonstrates a strong price-performance advantage for latency-sensitive and analytics-heavy workloads.

It materially outperforms gp3 under cold-read conditions while delivering io2-class behavior at a significantly lower monthly cost.

For teams running blockchain infrastructure, indexing pipelines, or high-throughput data systems, ABS represents a compelling default storage choice where predictable performance and cost efficiency are both critical.




Workload Fit


ABS is especially strong for:



OLAP Databases




High IOPS Database




ETL pipelines



AI Pipelines



Vector Databases




Gaming



Blockchain
Infrastructure

When Nirvana Makes Sense



Built for Real-Time Workloads

Infrastructure purpose-engineered for latency-sensitive and always-on systems, powering blockchain, AI, and real-time applications at production scale.

Cloud without oversharing

Dedicated hardware means no noisy neighbours. Your workload gets full, uncontested access to the resources it's allocated — every time.

Built-In Cost Efficiency

Over 50% cost savings compared to hyperscaler equivalents. Performance is included by default — not locked behind premium add-ons or tier upgrades.

Alternative to Hyperscalers

Enterprise-grade performance without hyperscaler complexity, opaque pricing, or vendor lock-in. The power of a major cloud — with none of the bloat.



Get in Touch

The Modern Cloud for Real-time Workloads

High-IOPS Block Storage.

Learn more at [Nirvana Labs](#)

[Nirvana Cloud](#) | [Pricing](#) | [Blog](#) | [Docs](#) | [Twitter](#) | [Telegram](#) | [LinkedIn](#)