

Simplyblock Performance

All tests are run with 20 fio jobs (threads) and an IO-depth of 128 (except otherwise stated).

We adjust the node count per host to better reflect the actual resource usage and take advantage of NUMA cores in modern CPUs and multi-socket systems. Since each VM runs two storage nodes and each physical host has two VMs, we count only one effective node per host, one-fourth of the total node count.

According to the above rule, the x-axis of all graphs represents the number of physical hosts.

General Test Environment:

- 1 physical storage host with 2 NUMA nodes
- 2x Intel Xeon 8175 (24 cores each) at 3.9GHz
- Up to 4 storage nodes per physical host (2 per NUMA node)
- 2 NVMe devices per storage node (8 per storage host)

Test configurations:

Storage Hosts per Cluster	Storage Node per Storage Host	NVMe Devices per Cluster	CPU Cores per Cluster (Intel x64 Skylake, 2 vCPUs per Core)
1	1	2	12*
1	2	4	24*
1	4	8	48*
2	4	16	96*
3	4	24	144*
4	4	32	192*
8	4	64	384*

* This is the number of CPU cores per storage node. Per 12 cores each, 10 were assigned to the simplyblock storage services. The remaining two cores were available to the Linux kernel.

Executive Summary

Simplyblock demonstrates near-linear scalability in both IOPS and throughput across multiple cluster sizes and workload patterns. Benchmark tests using fio with 20 jobs and an I/O depth of 128 were conducted across configurations ranging from 1 to 32 nodes (1 to 8 physical hosts).

The system showed excellent scaling efficiency and consistency under random, sequential, read-heavy, and mixed workloads.

Performance Highlights

IOPS (4K Blocksize - Random Read)

- Scaling from 1 to 32 nodes (0.25 to 8 physical hosts) shows about **29x increase in IOPS**, from **~340K to ~9.9 million**.
- Simplyblock shows an almost linear scalability of **~90% per storage node**.
- Even under a suboptimal 50/50 mixed read-write workload, performance scales **from ~195K to ~5.5 million IOPS—a 28x improvement**.

Throughput (128K Blocksize - Sequential Mixed 70/30)

- Throughput scales **from ~22 Gbit/s at 1 node to almost 770 Gbit/s at 32 nodes—an almost 35x increase**.
- That's over **107 GB/s** of bandwidth, showing simplyblock's ability to fully saturate available network hardware.

Latency-Constrained Peak Performance

At higher concurrency (20 jobs, 128 IO depth), with 4 volumes:


- Achieved up to **888K read IOPS** and **29.5 Gbit/s throughput** within target latencies.
- Simplyblock maintains performance even under tighter IO depth settings (e.g., 8 or 16).

IO Depth Sensitivity


- Read IOPS remained consistent (~450K–715K) across varying IO depths from 4 to 128, indicating **stable latency and internal queuing efficiency**.
- Write IOPS increased significantly at IO depth 16 (up to **660K**), suggesting simplyblock's responsiveness to deeper pipelines.

Total IOPS Per Cluster (Blocksize: 4K)


4K - IOPS - 100% Read




4K - IOPS - 100% Write



4K - IOPS - 30% Read - 70% Write



4K - IOPS - 70% Read - 30% Write



4K - IOPS - 50% Read - 50% Write




Table 1: Random Access Pattern with 4K blocksize


Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	343,895	156,380	177,216	223,922	194,991
2	0.5	637,077	291,938	329,272	417,770	365,190
4	1	1,262,162	560,574	641,198	827,038	705,169
8	2	2,487,440	1,099,440	1,260,028	1,631,707	1,385,397
12	3	3,715,484	1,638,129	1,879,415	2,437,916	2,065,594
16	4	4,943,528	2,176,818	2,498,802	3,244,126	2,745,792
32	8	9,855,703	4,331,573	4,976,352	6,468,964	5,466,583

Table 2: Sequential Access Pattern with 4K blocksize


Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	350,901	189,253	213,315	275,190	244,174
2	0.5	666,736	354,030	399,742	503,329	447,254
4	1	1,353,539	643,371	729,121	941,129	820,323
8	2	2,691,703	1,247,903	1,415,827	1,828,608	1,587,733
12	3	4,033,805	1,849,564	2,099,428	2,714,767	2,352,779
16	4	5,375,907	2,451,224	2,783,029	3,600,925	3,117,826
32	8	10,744,314	4,857,865	5,517,432	7,145,560	6,178,012

Total IOPS Per Cluster (Blocksize 128K)


128K - IOPS - 100% Read




128K - IOPS - 100% Write



128K - IOPS - 30% Read - 70% Write



128K - IOPS - 70% Read - 30% Write



128K - IOPS - 50% Read - 50% Write




Table 3: Random Access Pattern with 128K blocksize


Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	27,927	15,202	17,207	21,573	19,151
2	0.5	47,455	29,004	39,286	44,004	40,573
4	1	86,511	56,609	83,442	88,865	83,416
8	2	164,623	111,817	171,755	178,588	169,104
12	3	242,735	167,026	260,068	268,311	254,792
16	4	320,846	222,235	348,381	358,034	340,479
32	8	633,294	443,070	701,632	716,926	683,229

Table 4: Sequential Access Pattern with 128K blocksize


Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	27,757	15,128	17,191	21,266	19,100
2	0.5	47,441	28,794	38,264	43,766	40,538
4	1	86,808	56,124	80,410	88,765	83,414
8	2	165,543	110,784	164,702	178,764	169,166
12	3	244,277	165,445	248,994	268,764	254,918
16	4	323,011	220,105	333,286	358,763	340,670
32	8	637,949	438,748	670,454	718,759	683,679

Total Throughput (Mbit/s) Per Cluster (Blocksize 4K)


4K - Throughput (Mbit/s) - 100% Read




4K - Throughput (Mbit/s) - 100% Write



4K - Throughput (Mbit/s) - 30% Read - 70% Write



4K - Throughput (Mbit/s) - 70% Read - 30% Write



4K - Throughput (Mbit/s) - 50% Read - 50% Write




Table 5: Random Access Pattern with 4K blocksize


Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	11,264	5,122	5,804	7,334	6,386
2	0.5	20,860	9,560	10,782	13,680	11,958
4	1	41,310	18,353	20,990	27,073	23,084
8	2	81,401	35,992	41,244	53,408	45,348
12	3	121,582	53,626	61,516	79,794	67,610
16	4	161,764	71,259	81,788	106,179	89,872
32	8	322,870	141,768	162,953	211,934	178,916

Table 6: Sequential Access Pattern with 4K blocksize


Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	11,493	6,199	6,986	9,013	7,998
2	0.5	21,833	11,594	13,088	16,483	14,647
4	1	44,302	21,064	23,868	30,812	26,852
8	2	88,090	40,853	46,343	59,863	51,965
12	3	132,006	60,547	68,717	88,870	76,999
16	4	175,921	80,241	91,091	117,877	102,034
32	8	352,128	158,617	180,153	233,720	201,840

Total Throughput (Mbit/s) Per Cluster (Blocksize 128K)


128K - Throughput (Mbit/s) - 100% Read




128K - Throughput (Mbit/s) - 100% Write



128K - Throughput (Mbit/s) - 30% Read - 70% Write



128K - Throughput (Mbit/s) - 70% Read - 30% Write



128K - Throughput (Mbit/s) - 50% Read - 50% Write




Table 7: Random Access Pattern with 128K blocksize

Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	29,274	15,939	18,041	22,615	20,074
2	0.5	49,734	29,675	41,083	45,940	42,491
4	1	90,654	57,147	87,167	92,590	87,325
8	2	172,494	112,091	179,335	185,890	176,993
12	3	254,334	167,035	271,503	279,190	266,661
16	4	336,174	221,979	363,671	372,490	356,329
32	8	663,534	441,755	732,343	745,690	715,001

Table 8: Sequential Access Pattern with 128K blocksize

Node Count	Host Count	RW (100/0)	RW (0/100)	RW(30/70)	RW (70/30)	RW (50/50)
1	0.25	29,099	15,861	18,022	22,296	20,027
2	0.5	49,732	29,091	39,592	46,402	42,626
4	1	90,998	55,551	82,732	94,614	87,824
8	2	173,530	108,471	169,012	191,038	178,220
12	3	256,062	161,391	255,292	287,462	268,616
16	4	338,594	214,311	341,572	383,886	359,012
32	8	668,722	425,991	686,692	769,582	720,596

Max IOPS by Latency Target (jobs=20, I/O-Depth=128)


IOPS with Latency Target (bs=4K, Read=300μs, Write=500μs)



Number of Volumes	Read IOPS	Write IOPS
1	205948	159506
2	459417	283540
4	888768	527846


Max Throughput by Latency Target (jobs=20, I/O-Depth=128)

Throughput with Latency Target (bs=4K, Read=300μs, Write=500μs)



Number of Volumes	Read Throughput (Mbit/s)	Write Throughput (Mbit/s)
1	6839	5256
2	15096	9370
4	29483	17399

Max IOPS by I/O Depth + Latency Target (jobs=10, volumes=4)



Glossary: Storage Benchmarking Terms

IOPS (Input/Output Operations Per Second)

The number of discrete read/write operations processed per second is key for workloads with many small, random access patterns.

Throughput

Measures the total data volume moved over time (e.g., MB/s or Gbit/s). Important for bandwidth-heavy applications like media or analytics.

Latency

Time it takes to complete a single I/O operation. Low latency is crucial for responsive applications.

IO Depth

The number of simultaneous I/O operations the system handles. Higher depth = more parallelism.

Blocksize

The unit size of each I/O request. Small blocks (e.g., 4K) simulate database loads; large blocks (e.g., 128K) are typical of media streaming or backup.

Read-Write Mix

Defines the workload profile (e.g., 100% reads, 70/30 reads/writes). System performance can vary widely depending on this ratio.