Micron® 7000 Series SSDs Bring Strong Performance to **MinIO®** Object Storage

As data center scale out their infrastructure utilizing multi-node clusters for data sensitive applications, performance often suffers from file system overhead and data movement between nodes. IT and cloud managers require the fast, low-latency and predictable performance of NVMe™ storage for their active object stores without breaking their budget. Micron 7000 series SSDs,[1] combined with MinIO, meet these needs head on.

Lab testing shows that Micron 7000 series SSDs provide excellent storage for MinIO object storage cluster nodes that use AMD EPYC™ CPUs by demonstrating exceptional peak PUT and GET performance[2] in a single MinIO node.[3]

Single-Node MinIO Peak Performance

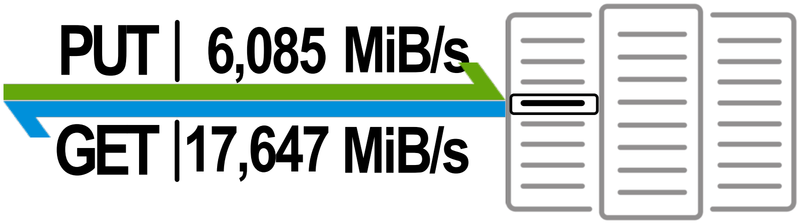

Fast Facts

Micron 7000 Series SSDs Deliver Exceptional Storage Performance

Micron optimized the 7000 series SSDs to bring NVMe performance to an approachable price point, helping customers meet their performance and budget demands.

Micron 7000 series leverages the low power consumption and price-performance efficiencies of 3D NAND technology, delivering fast NVMe throughput for performance-focused object stores like MinIO.

AMD™ CPUs: The Right Fit for Performance

Micron selected AMD EPYC™ CPUs because it enables the I/O, memory capacity, and memory bandwidth to accomplish what is needed with MinIO and the Micron 7000 series.

MinIO Is a Flexible Solution for the Hybrid Cloud

Micron chose MinIO for these tests because it supports a broad range of use cases across an immense number of environments, including in the public cloud, in the private cloud and at the edge.

MinIO has been cloud-native since its inception. With a focus on performance and scalability, MinIO can deliver on a range of cloud-based use cases — from artificial intelligence, machine learning, analytics and backups/restores to modern web and mobile apps.[4]

MinIO and Micron 7000 Series SSDs Offer an Amazing Combination of Speed and Scale

In our tests, a single MinIO node with six Micron® 7000 series SSDs delivered 6,085 MiB/s peak PUT performance and 17,647 MiB/s of peak GET performance.

Imagine what is possible with an entire cluster.

1. Micron 7400 Pro SSD with NVMe used in this document
2. Performance is understood to mean average throughput, typically measured in mebibytes per second (MiB/s). PUT represents write throughput, while GET represents read throughput.
3. Single-node results may not reflect complete cluster results. Actual results may vary.
4. Information retrieved from https://min.io/ June, 2021.

Performance Needed by Object Stores

Object stores have become broadly used, and their role is changing. Object stores enable simple, rapid expansion to manage larger data environments, and they can integrate with a wide range of applications (via Microsoft® Azure® and Amazon S3 APIs). The increased adoption of SSDs, especially those with NVMe, has yielded a significant increase in object store performance.

The increased performance puts object stores in a new spotlight by enabling file system-like performance.

Measuring MinIO Performance

Customers choose the Micron 7000 series SSDs to expand the benefits of NVMe across their data center; they choose MinIO due to its extreme read and write speeds with readily available, standards-based servers. For this test, a single node using a standard AMD EPYC™ server equipped with Micron® 7000 series SSDs was set up. This single-node performance shows how the Micron 7400 SSD can enable MinIO to operate as the primary storage for a diverse set of workloads ranging from Apache Spark™ to Presto®, TensorFlow to H2O.ai platforms, an alternative to Apache™ Hadoop® HDFS and other potential use cases.[5]

The test configuration used a single MinIO node with six 7.68TB Micron® 7000 series SSDs and connected them to two client nodes executing the Warp[6] benchmark using NVIDIA® Mellanox® ConnectX®-6 Dx running at 200 Gb/s[7] (Figure 1).

Testing procedures consisted of object PUTs (write throughput) and object GETs (read throughput) using a broad range of thread counts (#Threads) and object sizes. Performance results are measured in MiB/s (higher is better).

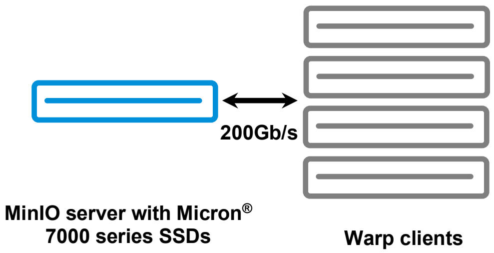

Figure 1: Test Configuration

Peak PUT Performance: 6,085 MiB/s

Table 1 provides PUT throughput performance numbers for various object size and number of threads. A single instance of MinIO reaches a maximum average throughput of 6,085 MiB/s when performing PUTs using 128 MiB object size and 140 threads. Object size and #Threads can significantly affect peak PUT performance. With this configuration, peak PUT performance is more than double the minimum PUT performance with the lowest performance observed using 4 MiB object size and 140 threads. Higher PUT throughput is achieved using larger object sizes indicated by the darker green shading..

For each #Threads, PUT performance generally increases with larger object size.

	Object Size									
#Threads	**4MiB**	**8MiB**	**16MiB**	**24MiB**	**32MiB**	**48MiB**	**64MiB**	**96MiB**	**128MiB**	**160MiB**
40	2,941	4,156	4,623	5,177	5,215	5,397	5,359	5,520	5,705	5,684
60	2,854	4,127	5,172	5,172	5,195	5,428	5,373	5,517	5,844	5,821
80	2,828	4,150	4,558	5,157	5,183	5,393	5,389	5,532	5,948	5,927
100	2,817	4,118	4,528	5,159	5,179	5,410	5,381	5,531	6,000	5,999
120	2,826	4,105	4,519	5,152	5,207	5,406	5,381	5,524	6,042	6,035
140	2,811	4,099	4,529	5,182	5,186	5,416	5,353	5,540	**6,085**	6,055

Lower — Higher

Table 1: S3 PUT Performance (MiB/s, higher is better)

Optimize Your Configuration

If your workload is predominantly large object sizes (≥ 64MiB), using a high #Threads may provide greater PUT performance. If, however, your workload is predominantly small object sizes (4MiB, 8MiB), using a small #Threads can provide greater PUT performance.

Table 2 shows that peak PUT CPU use (79.70%) was observed with 96MIB objects at #Threads = 40. When the object size ≥ 64MiB, CPU use varies little across all #Threads.

	Object Size									
#Threads	4MiB	8MiB	16MiB	24MiB	32MiB	48MiB	64MiB	96MiB	128MiB	160MiB
40	54.27%	56.60%	65.21%	71.63%	73.07%	76.49%	78.62%	**79.70%**	74.55%	74.16%
60	51.95%	54.61%	64.17%	69.41%	71.91%	75.46%	77.12%	78.20%	75.67%	75.56%
80	48.46%	52.96%	62.00%	67.75%	70.74%	73.58%	75.66%	78.62%	74.91%	76.35%
100	47.79%	51.42%	59.53%	66.30%	69.08%	72.40%	74.64%	76.20%	75.17%	75.24%
120	46.25%	49.43%	59.48%	64.52%	67.60%	71.50%	74.70%	77.17%	74.09%	74.27%
140	45.11%	48.28%	57.66%	63.94%	66.66%	71.33%	73.74%	75.79%	73.10%	74.27%

Lower — Higher

Table 2: S3 PUT CPU Use

Peak GET Performance: 17,647 MiB/s

GET performance was also measured peak using the same system configuration. For this testing, the maximum #Threads was reduced to 120 along with the range of object sizes reduced to ≤ 64MiB for this test. Testing observations indicated that increasing either beyond these values yielded no additional performance increase. Table 3 summarizes peak S3 GET performance results.

#Threads	4MiB	8MiB	16MiB	24MiB	32MiB	48MiB	64MiB
40	**17,647**	17,174	16,860	16,716	16,699	16,652	16,637
60	17,465	16,987	16,678	16,547	16,537	16,508	16,472
80	17,416	16,886	16,642	16,459	16,440	16,375	16,396
100	17,365	16,811	16,515	16,427	16,380	16,299	16,303
120	17,365	16,770	16,449	16,361	16,321	16,236	16,226

Lower — Higher

Table 3: S3 GET Performance (MiB/s)

Peak GET performance was consistently high; maximum and minimum GET performance differed by approximately 9%. Peak GET performance reached 17,647 MiB/s with 4 MiB objects.

5. Suitability for your needs may differ.
6. More details about the Warp benchmark is available from this Github: site https://github.com/minio/warp
7. More details about the NVIDIA® Mellanox® ConnectX®-6 Dx network adapters used is available from this NVIDIA® site: https://www.nvidia.com/en-us/networking/ethernet/connectx-6-dx/

Optimize Your Configuration

With peak GET performance of 17,647 MiB/s, Client-to-MinIO-server network bandwidth is critical. As noted, we used 200 Gb/s network connection. Ensure your network offers sufficient bandwidth for your configuration.

In our tests, a single MinIO node showed excellent PUT and GET performance using just six Micron® 7000 series SSDs for primary data storage. This object store building block delivered 6,085 MiB/s peak PUT performance and 17,647 MiB/s of peak GET performance: file system-like performance with the versatility of object storage.

Summary

Object stores have transformed from hard drive-based data dumps to modern architectures, like SSD-based data lakes and SSD-powered business-critical asset repositories. This change is fueled by a new generation of NVMe SSDs like the Micron® 7000 series, performance-focused storage software like MinIO and revolutionary CPUs like the AMD EPYC™.

The results of our testing demonstrates how far that transformation has come.

These tests used an example MinIO cluster node to illustrate the power of modern SSDs with NVMe, CPUs and cluster software. The results show how controlling the investment needed for all-flash infrastructure by using mainstream NVMe storage devices like Micron® 7000 series SSDs can be done without compromising performance.

Hardware Configuration

Hardware	Details
Server	Supermicro AS -1114S-WN10RT
Processor	Single-Socket AMD EPYC™ 7F72 (24-Core)
Memory	256GB Micron DDR4-3200
Server Storage	6x Micron® 7400 PRO 7.68TB SSDs with NVMe
Boot Drive	Micron 7000-series 960GB M.2 SSD with NVMe
MinIO Version	RELEASE.2021-08-05T22-01-19Z
OS	CentOS 8.4.2105
Kernel	4.18.0-305.3.1.el8.x86_64
Client Network Interface	NVIDIA® Mellanox® ConnectX®-6 Dx (running at 200 Gb/s)
Client Switch	NVIDIA® Mellanox® SN4700 running Onyx 3.9.2110

Table 4: MinIO Server Configuration

Hardware	Details
Server	Supermicro AS -1114S-WN10RT
Processor	Single-Socket AMD EPYC™ 7F72 (24-Core)
Memory	256GB Micron DDR4-3200
Boot Drive	Micron 7000-series 960GB M.2 SSD with NVMe
Warp Version	0.4.5
OS	CentOS 8.3.2011
Kernel	4.18.0-240.15.1.el8_3.x86_64
Client Network Interface	NVIDIA® Mellanox® ConnectX®-6 Dx (running at 200 Gb/s)

Table 5: Warp Client Configuration

Software Configuration

Setting	Value
fs.file-max	4194303
vm.swappiness	1
vm.vfs_cache_pressure	10
vm.min_free_kbytes	1000000
net.core.rmem_max	268435456
net.core.wmem_max	268435456
net.core.rmem_default	67108864
net.core.wmem_default	67108864
net.core.netdev_budget	1200
net.core.optmem_max	134217728
net.core.somaxconn	65535
net.core.netdev_max_backlog	250000

Setting	Value
net.ipv4.tcp_rmem	67108864 134217728 268435456
net.ipv4.tcp_wmem	67108864 134217728 268435456
net.ipv4.tcp_low_latency	1
net.ipv4.tcp_adv_win_scale	1
net.ipv4.tcp_max_syn_backlog	30000
net.ipv4.tcp_max_tw_buckets	2000000
net.ipv4.tcp_tw_reuse	1
net.ipv4.tcp_fin_timeout	5
net.ipv4.conf.all.send_redirects	0
net.ipv4.conf.all.accept_redirects	0
net.ipv4.conf.all.accept_source_route	0
net.ipv4.tcp_mtu_probing	1

Table 6: Software Configuration

About the Warp Benchmark

Micron used the Warp benchmark (https://github.com/minio/warp) to test the PUT and GET performance of this S3-compatible object store (a single MinIO node). Warp is a distributed benchmark, allowing multiple client systems to test performance against an S3-compatible cluster. Note that in these tests, we used a single cluster node to evaluate its configuration's suitability as a MinIO cluster node building block.

The benchmark returns average throughput, fastest object throughput, slowest object throughput and 50% median throughput; it does not return average or percentile latency. Objects in the object store are refreshed for each test (performing cleanup and then reloading).

Rev. A 02/2022 CCM004-676576390-11589