THE MICRON 7000 SERIES SSDs BRING STRONG PERFORMANCE TO MINIO OBJECT STORAGE

Data center IT and cloud managers want the fast, low-latency and consistent performance of NVMe™ storage for their active object stores — and they don't want to break their budget. The Micron 7000 series SSD,[1] combined with MinIO®, meets these needs head on.

Micron lab testing shows that the 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

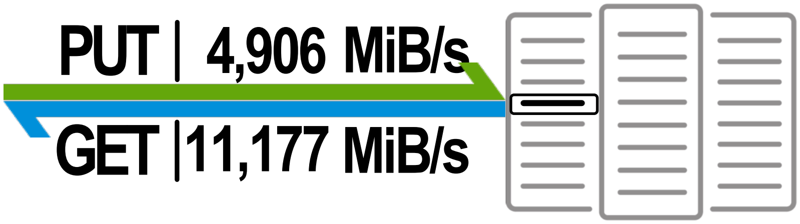

Fast Facts

Micron 7000 Series SSDs Deliver Exceptional Storage Performance and Extreme Value

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

The 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.

MinIO Is a Flexible Solution for the Hybrid Cloud

Micron chose MinIO for these tests because MinIO 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 ten Micron 7000 series SSDs delivered 4,906 MiB/s peak PUT performance and 11,177 MiB/s of peak GET performance.

Imagine what you can do with an entire cluster.

1. Micron 7300 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/ on June 21, 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 use puts object stores in a new spotlight with file system-like performance.

Measurement of 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 standard hardware. In this test, we set up a single node using standard AMD EPYC servers equipped with Micron 7000 series SSDs. This single-node performance may enable MinIO to operate as the primary storage for a diverse set of workloads ranging from Apache Spark™ to Presto®, TensorFlow and H2O.ai. It can also be used as an alternative to Apache Hadoop HDFS.[5]

We used a single MinIO node with ten 7.68TB Micron 7000 series SSDs and connected them to two Warp[6] client nodes using NVIDIA® Mellanox® ConnectX®-6 Dx (running at 100 Gb/s)[7] (Figure 1).

We tested object PUTs (write throughput) and object GETs (read throughput) using thread counts ranging from 40 to 120 and object sizes ranging from 4 MiB to 64 MiB. Performance results are measured in MiB/s (higher is better).

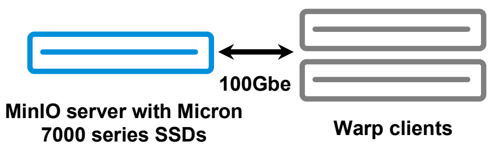

Figure 1: Test Configuration

Peak PUT Performance: 4,906 MiB/s

The increased adoption of SSDs, especially those with NVMe, fostered a change in perspective on object storage performance.

Table 1 shows that this single instance of MinIO reaches a maximum average throughput of 4,906 MiB/s when performing PUTs using 32 MiB object size and 120 threads. Table 1 also shows that the choice of object size can significantly affect peak PUT performance: The lowest PUT performance with this configuration is just over half the maximum. The lowest performance is observed using 4 MiB object size and 120 threads.

	Object Size						
#Threads	**4 MiB**	**8 MiB**	**16 MiB**	**24 MiB**	**32 MiB**	**48 MiB**	**64 MiB**
40	2,719	3,141	3,639	4,728	4,646	4,132	4,124
60	2,673	3,076	3,882	4,305	4,441	4,542	4,358
80	2,725	3,045	3,870	4,101	4,397	4,329	4,340
100	2,702	3,016	3,906	4,509	4,552	4,358	4,458
120	2,618	2,959	3,936	4,784	**4,906**	4,812	4,473

Lower **Higher**

Table 1: S3 PUT Performance (MiB/s)

5. Suitability for your needs may differ.
6. https://github.com/minio/warp
7. https://www.nvidia.com/en-us/networking/ethernet/connectx-6-dx/

Table 2 shows that peak PUT performance (4,906 MiB/s using 32MiB object size and 120 threads) corresponds to a maximum CPU utilization of 80.65% and that CPU utilization shows similarly high values across most tested object sizes and thread counts (the minimum CPU utilization is approximately 72% of maximum).

	Object Size						
#Threads	4 MiB	8 MiB	16 MiB	24 MiB	32 MiB	48 MiB	64 MiB
40	60.27%	64.00%	59.83%	72.07%	74.92%	61.79%	63.79%
60	58.87%	64.62%	63.63%	67.50%	72.35%	69.68%	67.14%
80	58.72%	66.14%	66.00%	65.14%	74.11%	66.69%	66.99%
100	58.90%	67.89%	63.91%	69.24%	76.91%	67.65%	68.99%
120	58.06%	67.19%	64.88%	73.77%	**80.65%**	73.73%	69.59%

Lower Higher

Table 2: S3 PUT CPU Use

Optimize Your Configuration

If your workload is predominantly fixed object sizes, tuning the thread count for that object size may provide greater PUT performance.

Peak GET Performance: 11,177 MiB/s

We also measured peak GET performance using the same system configuration, range of thread counts and object sizes (Table 3).

	Object Size						
#Threads	4 MiB	8 MiB	16 MiB	24 MiB	32 MiB	48 MiB	64 MiB
40	11,136	**11,179**	11,170	11,138	11,128	11,113	11,105
60	11,162	11,179	11,167	11,131	11,122	11,101	11,094
80	11,161	11,177	11,165	11,127	11,115	11,094	11,090
100	11,166	11,173	11,163	11,125	11,110	11,094	11,086
120	11,161	11,171	11,161	11,124	11,108	11,090	11,084

Lower Higher

Table 3: S3 GET Performance (MiB/s)

Note that peak GET performance was consistently high. The difference between maximum and minimum GET performance was less than 1%. Peak GET performance reached 11,179 MiB/s with 8 MiB objects and 40 threads. Minimum GET performance reached 11,084 MiB/s using 64 MiB objects and 120 threads, the equivalent of 93.76 Gb/s and 92.97 Gb/s, respectively. These results approach the 100 Gb/s theoretical maximum of the network adapters.

Optimize Your Configuration

Client-to-MinIO-server network bandwidth may affect GET performance. If your workload is predominantly read transactions, lower-power server platforms may be an option. Client CPU resources are typically more important than the server CPU resources for read performance.

Summary

Object stores have transformed — from hard drive-based data dumps to SSD-based data lakes and to SSD-powered business-critical asset repositories. This change has been 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 this series of tests demonstrate 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. They show controlling the investment needed for all-flash infrastructure by using SSDs like the Micron 7000 series SSDs can be done without compromising on performance.

Hardware Configuration

The Micron 7000 series SSD offers a robust complement of proven security features built over generations of Micron data center SSDs.

Hardware	Details
Server	Supermicro AS -1114S-WN10RT
Processor	Single-Socket AMD EPYC 7F72 (24-Core)
Memory	256GB Micron DDR4-3200
Server Storage	10x Micron 7300 PRO 7.68TB NVMe SSDs
Boot Drive	Micron 7300 PRO 960GB M.2 NVMe SSD
MinIO Version	RELEASE.2021-04-06T23-11-00Z
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 100 Gb/s)
Client Switch	NVIDIA Networking SN2700 running Cumulus Linux 4.1.33

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 7300 PRO 960GB M.2 NVMe SSD
Warp Version	0.3.43
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 100 Gb/s)

Table 5: Warp Client Configuration

System 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).

© 2021 Micron Technology, Inc. All rights reserved. All information herein is provided on an "AS IS" basis without warranties of any kind. Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Micron Technology, Inc., is not responsible for omissions or errors in typography or photography. Micron, the Micron logo and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.
Rev. A 08/2021 CCM004-676576390-11555