KIOXIA

Performance Brief

Accelerating Object-Based Storage in a MinIO® Cluster using KIOXIA RM7 Series Value SAS SSDs

Introduction

Modern applications have vastly different performance requirements that are directly impacted by the storage options used to back them. For example, file-based storage organizes data within a hierarchical folder structure, while block-based storage manages data as disk sectors grouped into blocks. Though used traditionally, many of these storage options do not have the ability to handle or manage vast amounts of unstructured data.

Object-based storage (or object storage) is a data storage option where information is organized into individual units called objects, where each object contains the data itself, associated metadata and a unique identifier that allows for efficient storage and retrieval of large volumes of unstructured data. Object storage provides the ability to manage large amounts of unstructured data while providing petabyte[1] scalability by adding more storage resources as needed. As a result, object store performance is tied directly to compute, networking and storage capabilities of the hardware.

An object store is dependent on its underlying storage as slow read/write throughput and latency can be a bottleneck. In applications where many concurrent requests occur, the inability to service those requests in a timely manner can result in slow application response times and negative user experiences. This bottleneck can often occur when using legacy storage options, such as HDDs or SATA SSDs. To prevent these bottlenecks, KIOXIA RM7 Series Value SAS SSDs support twice the interface speed of SATA via a 12 gigabit[1] per second (Gb/s) SAS-3 interface that delivers higher throughput and lower latency when compared to a slower storage interface, like SATA.

MinIO® is an open-source object store designed for cloud-native and distributed applications. It is compatible with Amazon Web Services™ S3 application programming interfaces (APIs) that utilize the RESTful[2] API standard. MinIO is regarded for its scalability, high performance and ease of deployment, making it a popular choice for building private cloud storage infrastructures, data lakes and content delivery networks (CDNs). MinIO has many built-in features, such as data replication, erasure coding and encryption that make it highly suitable for a wide range of use cases.

This performance brief presents throughput and latency test results of KIOXIA RM7-R Series Value SAS SSDs using an object storage architecture. It shows the expected cluster and drive results in a SAS server configuration. The KIOXIA Innovation Lab configured an object store on four Intel® Xeon® Gold 6330 servers, creating a MinIO cluster. Each server deployed eight KIOXIA RM7-R Series Value SAS SSDs, for a total cluster of thirty-two SSDs. Once the cluster was built, the MinIO Warp™ workload generator was used to test the performance of the object store and was setup utilizing two additional servers for load generation.

Test Results Snapshot

KIOXIA RM7-R Series Value SAS SSDs delivered the following cluster and drive test results in a MinIO® cluster:

Cluster Results

Average GET Throughput

up to 23.5+ GB/s

(based on the Restore workload)

Average PUT Throughput

up to 13.2+ GB/s

(based on the Backup workload)

Average CPU Utilization

up to 86.7%

Drive Results

Average Read Latency

up to 1.07 ms

Average Write Latency

up to 5.9 ms

Average Drive Utilization

up to 89.6%

Five different workloads were tested to emulate common use cases of object stores that included AI, web server, cloud storage, restore and backup. These five workloads included up to four RESTful operation types, GET, PUT, STAT and DELETE. Each workload was run three times, and their averages were recorded. The test metrics included average GET/PUT throughput, average read/write drive latency, average drive utilization and average CPU utilization.

The test results show that the four Intel Xeon Gold 6330 servers worked in tandem with the KIOXIA RM7-R Series Value SAS SSDs to deliver high application throughput and low read/write latency. The server/SSD configuration represents an excellent MinIO cluster solution that accelerates object stores.

The test results also include a brief description of each test metric and a graphical depiction of the test results, followed by analysis. Appendix A covers the hardware and software test configuration. Appendix B covers the configuration setup and test procedures.

Workload Descriptions

Each workload utilized two hundred concurrent operations, with each load generation system pushing an object size of 10 mebibytes[3] (MiB) that included one hundred threads each. This concurrency of two hundred does not mean two hundred users, but instead, it means two hundred initiators constantly launching operations of different types, which in turn, could emulate thousands or hundreds of thousands of actual end users. The total number of objects was set to forty thousand. Each workload consisted of the following RESTful API operations.

- *GET - a request to download an object from MinIO®*
- *PUT - a request to upload an object to MinIO*
- *STAT – a request to retrieve object information from MinIO*
- *DELETE – a request to remove an object from MinIO*

Workload 1: Artificial Intelligence Use Case
AI utilizes object storage to store and access massive amounts of data needed for training machine learning models. Object storage provides scalability, flexible metadata capabilities and the ability to handle large unstructured datasets efficiently, which in turn, enables fast retrieval of specific data points during the training and inference processes. In essence, object storage essentially acts as a central data lake for AI workloads. This workload consists of 90% GET operations for accessing the data, and 10% PUT operations related to metadata updates and checkpoints.

Workload 2: Web Server Use Case
Web servers utilize object storage to store and retrieve website content. This data can include images, cascading style sheet files and JavaScript® files. This workload consists of 70% GET operations for accessing the data and 30% PUT operations related to metadata updates and log files.

Workload 3: Cloud Storage Use Case
The cloud storage workload is intended to represent multiple users connecting to a cloud storage service, and are sharing and syncing files across multiple devices. These users are uploading, retrieving and deleting files, while the metadata for the files are being utilized for data analytics services. This workload consists of 40% GET operations, 10% STAT operations, 40% PUT operations, and 10% DELETE operations.

Workload 4: Restore Use Case
The restore workload simulates the downloading of data that has been saved in an object store to restore data that has previously been backed up. For this workload, the object store is storing the data from a backup system, which in turn, accesses the object store and downloads the data that was requested. This workload consists of 100% GET operations.

Workload 5: Backup Use Case
The backup workload simulates backing up data from a local system to an external location, which provides an outside source of the data in case of data loss or corruption. The workload consists of 100% PUT operations.

Test Results[4]

<u>Cluster Level Results:</u>

Cluster Level Test Metric 1: Average GET/PUT Throughput
This test measures the average total throughout during testing for each of the five workloads using KIOXIA RM7-R Series Value SAS SSDs in the MinIO cluster. The results are an average from three test runs to demonstrate the GET and PUT throughput performance that can be expected. The results are recorded in megabytes[1] per second (MB/s).

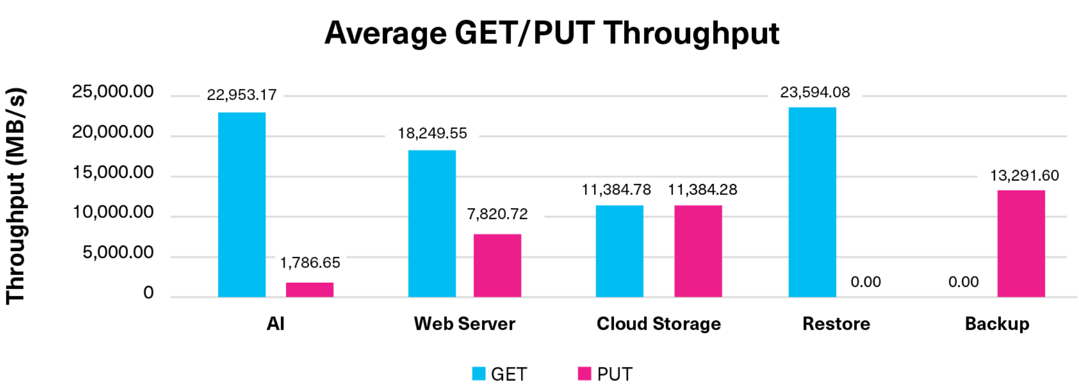

Cluster Level Test Metric 2: Average CPU Utilization

CPU utilization determines the amount of computing tasks that are being performed by the server CPU. Low CPU utilization means that the processor is not being used efficiently. The results are an average from three test runs to demonstrate the CPU utilization that can be expected. The results are recorded in percentage of usage.

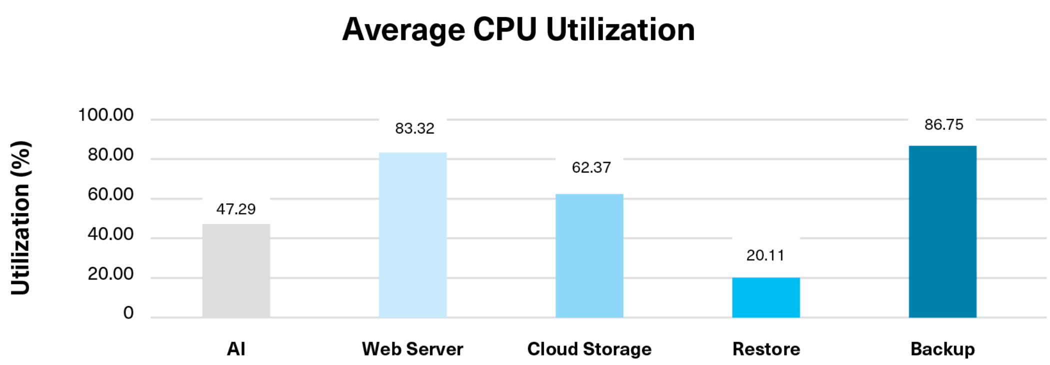

The workloads were distributed evenly across all four Intel® Xeon® Gold 6330 servers in the cluster. For CPU utilization, there was no evidence of complete saturation of the compute resources, and the average CPU utilization was 86.75% for backup workloads.

AI workloads, such as training models in machine learning, need to have compute resources, such as accelerators and GPUs, to be constantly fed with data. The cluster level test results demonstrate high GET throughput from the MinIO® cluster, which allows for large amounts of data to be populated into accelerator/GPU memory quickly so that these compute resources constantly have data to train on.

Web servers and cloud storage services need to be able to handle many user requests simultaneously. The cluster level test results demonstrate high throughput for both GET and PUT operations. These results highlight the MinIO clusters' ability to perform different operations.

For backups and restores, the MinIO cluster needs to be able to handle large data throughput to minimize the amount of time it takes to perform these operations. In instances when applications crash, it is essential to backup that data periodically to minimize data loss. Meeting business recovery time objectives (RTOs) and recovery point objectives (RPOs) are critical to maximize application uptime and prevent revenue loss. Higher PUT throughput could be used to run more frequent backups to meet RPOs. Higher GET throughput allows for faster restores, which are critical to get the data back quickly to meet RTOs. The MinIO cluster was able to support high GET and PUT throughput to support these restore and backup processes.

Drive Level Results:

Drive Level Test Metric 1: Average Read/Write Drive Latency

Latency is the measure of the total time required for a subsystem, or one of its components, to process a single storage operation or data request. The time it takes for the data to begin moving from one system to another may greatly affect application performance and the overall user experience. Read latency is the total time for data to be returned from a storage device following a read request from the host. Write latency is like read latency, but for writing data, and it is the total time for a storage device to complete writing the data following instruction from the host. The results are an average from three test runs to demonstrate the read/write latency that can be expected. The results are recorded in milliseconds (ms).

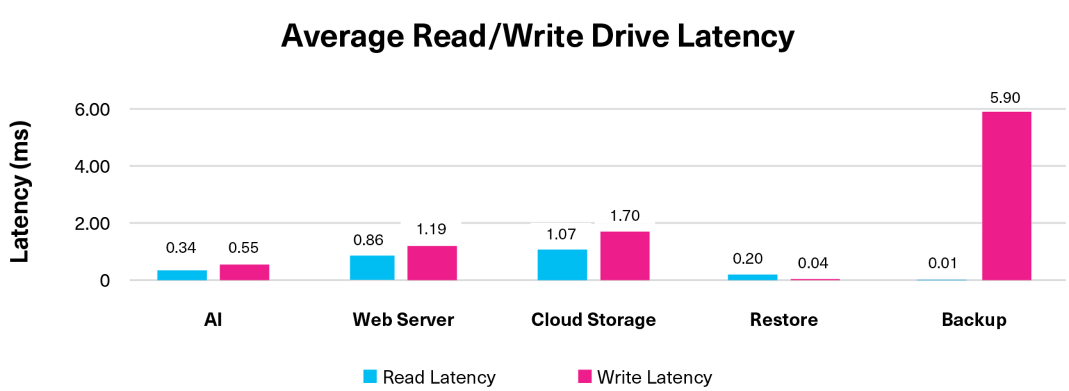

Drive Level Test Metric 2: Average Drive Utilization

Drive utilization is defined as the percent of time that the drive is working on any active read or write I/O. High drive utilization indicates that the drives in the MinIO® cluster are always processing read and write requests, while low drive utilization indicates that the drives are in an idle state. The results are an average from three test runs to demonstrate the drive utilization that can be expected. The results are recorded in percentage of usage.

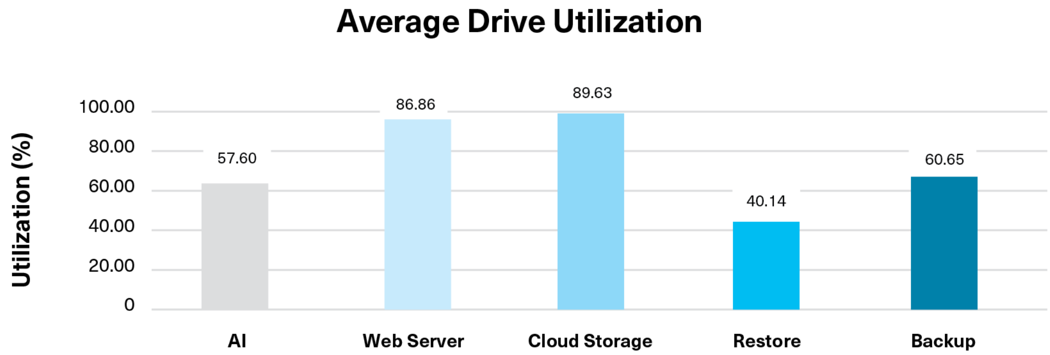

The workloads were distributed evenly across all thirty-two KIOXIA RM7-R Series Value SAS SSDs in the cluster. The ability for each drive to focus on smaller subsets of read and write I/O ensured that the storage subsystem was not the cause of bottlenecks, as seen from the drive utilization not being saturated at 100% across any workload. The drives were able to service requests very quickly due to the low average read/write latencies.

The read/write latency for the AI workload had sub 1 millisecond results showing that underlying KIOXIA RM7-R Series Value SAS SSDs were able to respond quickly to the requests for these objects.

Web servers and cloud storage services both had low read/write latencies which support the workload's need to handle multiple requests simultaneously.

For backups and restores, quick response times of the underlying drives were needed to maintain RTO and RPO goals. The backup and restore workloads completed quickly due to quick response times delivered by the underlying KIOXIA RM7-R Series Value SAS SSDs.

Test Analysis

An encapsulation of the six test metrics for the five use case workloads yielded the following results:

	Cluster Results			Drive Results		
Workload	**GET Throughput** *(in MB/s)*	**PUT Throughput** *(in MB/s)*	**CPU Utilization** *(in %)*	**Read Latency** *(in ms)*	**Write Latency** *(in ms)*	**Drive Utilization** *(in %)*
AI	22,953.17	1,786.65	47.29	0.34	0.55	57.60
Web Server	18,249.55	7,820.72	83.32	0.86	1.19	86.86
Cloud Storage	11,384.78	11,384.28	62.37	1.07	1.70	89.63
Restore	23,594.08	0.00	20.11	0.20	0.04	40.14
Backup	0.00	13,291.60	86.75	0.01	5.90	60.65

The MinIO cluster was configured with a 100 Gb/s network, which was equivalent to a network throughput maximum of 12.5 gigabytes[1] per second (GB/s). As a result, the maximum network bandwidth that was achievable by the two load generation systems was equal to 25 GB/s. As evident from the GET/PUT throughput test results, network saturation became an issue for four of the five tests when the aggregate throughput was about 23.5 GB/s. The remaining difference between the theoretical maximum network bandwidth and the throughput was due to the overhead that occurred at the network protocol level.

The backup workload achieved a maximum throughput over 13.2 GB/s, but showed low drive utilization and high CPU utilization, demonstrating that the bottleneck was not occurring at the drive level. Upgrading the network to higher bandwidth capabilities, as well as adding additional nodes for compute resources, can enable the drives to perform at higher drive utilization while still maintaining low latency.

Summary

This performance brief presented the test results of an object store in a MinIO® cluster comprised of four Intel® Xeon® Gold 6330 servers, each with eight KIOXIA RM7-R Series Value SAS SSDs, and showcased cluster and drive results that can be expected in this configuration. Testing showed that the four servers worked in tandem with the KIOXIA RM7-R Series Value SAS SSDs to deliver high application throughput and low read/write latencies. The combination of Intel Xeon Gold 6330 servers with KIOXIA RM7-R Series Value SAS SSDs represents an excellent MinIO cluster solution that accelerates object stores.

KIOXIA RM7 Series Value SAS SSD Product Info

The latest generation KIOXIA RM7 Series Value SAS SSDs support the 12 Gb/s SAS-3 interface and are available in two 2.5-inch[5] configurations: RM7-R Series for read-intensive applications (1 DWPD[6], up to 7.68 terabyte[1] (TB) capacities) and RM7-V Series for higher endurance mixed-use applications (3 DWPD, up to 3.84 TB capacities). Security options[7] are available for both configurations.

Additional KIOXIA RM7 Series Value SAS SSD information is available at the KIOXIA 'Life After SATA' site.

KIOXIA RM7 Series SSD[8]

Appendix A

Hardware/Software Test Configuration

Server Information	
No. of Servers	4
CPU Information	
CPU Model	Intel® Xeon® Gold 6330
No. of Sockets	2
No. of Cores	28
Frequency	2.0 GHz
Memory Information	
Memory Type	DDR4
Memory Speed	DDR4-2000
Total Memory Size	128 GB

Load Generation Server Information	
No. of Servers	2
CPU Information	
CPU Model	Intel Xeon Gold 6354
No. of Sockets	2
No. of Cores	18
Frequency	3.0 GHz
Memory Information	
Memory Type	DDR4
Memory Speed	DDR4-3200
Total Memory Size	128 GB

Operating System Information	
Operating System (OS)	Linux® Ubuntu®
OS Version	24.04.1

Software Information	
Database Software Model	MinIO®
Version	minio version RELEASE.2025-01-20T14:49:07Z
Test Software Model	MinIO Warp™
Version	1.0.6

SSD Information	
SSD Model	KIOXIA RM7-R Series
Form Factor	2.5-inch
Interface	SAS-3
Interface Speed	12 Gb/s SAS
No. of SSDs	8 per server / 32 total
SSD Capacity	3.84 TB
Drive Writes per Day	1
Active Power	up to 9 watts

Appendix B

Configuration Setup/Test Procedures

Configuration Setup

Four Intel® Xeon® Gold 6330 servers were initialized.

Linux® Ubuntu® v24.04.1 was installed on all the four servers.

MinIO® was installed on all of the four servers.

- Reference MinIO official documentation for the commands: https://min.io/docs/minio/linux/index.html.
- Reference official MinIO Warp™ GitHub™ for installation commands.
- Reference MinIO kernel level tuning utilizing the tuned package: https://github.com/minio/minio/tree/master/docs/tuning

An XFS® file system was placed on each drive in the cluster and mounted.

Eight KIOXIA RM7-R Series Value SAS SSDs were mounted in each server for a total of thirty-two using this scheme /mnt/disk1/minio ... /mnt/disk8/minio for the four servers.

The /etc/default/minio file that was generated with the MinIO install was modified as follows:
Edit the MINIO_VOLUMES line to point to the servers in the cluster and their respective drives, i.e. https://minio{1...4}.example.net:9000/mnt/disk{1...4}/minio and change the MINIO_SERVER_URL to point to one of the servers in the cluster so it can act as the host for the MinIO cluster.

The MinIO service was started through systemd on all nodes in the cluster.

The MinIO Client was installed to check the cluster setup and to assure that it was working properly.

- Reference the MinIO installation guide: https://min.io/docs/minio/linux/reference/minio-mc.html.

MinIO Warp was installed on the two load generation servers.

- Reference the MinIO installation guide: https://github.com/minio/warp.

A virtual machine (VM) was created to launch MinIO Warp.

- MinIO Warp was installed on the VM.

MinIO Warp was run on both the load generation servers using the command 'run./warp client.'

Each load generation server used two hundred and eighty threads with an object size of four million to move data within the MinIO cluster.

Test Procedures

The cluster level results were recorded for the following test metrics using the five workload use cases:

- *Average GET Throughput (in MB/s)*
- *Average PUT Throughput (in MB/s)*
- *Average CPU Utilization (in percentage of use)*

The drive level results were recorded for the following test metrics using the five workload use cases:

- *Average Read Latency (in ms)*
- *Average Write Latency (in ms)*
- *Average Drive Utilization (in percentage of use)*

For each of these six tests, three total runs were performed and the average of the three runs was calculated.

NOTES:

[1] Definition of capacity: Kioxia Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes, a terabyte (TB) as 1,000,000,000,000 bytes and a petabyte (PB) as 1,000,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 2^{30} bits = 1,073,741,824 bits, 1GB = 2^{30} bytes = 1,073,741,824 bytes, 1TB = 2^{40} bytes = 1,099,511,627,776 bytes and 1PB = 2^{40} bytes = 1,125,899,906,842,624 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

[2] RESTful systems are stateless and separate the concerns of client and server. Those based on HTTP rely on predefined methods. Common predefined methods include GET, PUT, DELETE and POST.

[3] A mebibyte (MiB) is equal to 1,048,576 bytes.

[4] Read and write speed may vary depending on the host device, read and write conditions and file size.

[5] 2.5-inch indicates the form factor of the SSD and not the drive's physical size.

[6] DWPD: Drive Write(s) Per Day. One full drive write per day means the drive can be written and re-written to full capacity once a day, every day, for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.

[7] Optional security feature compliant drives are not available in all countries due to export and local regulations.

[8] The product image shown is a representation of the design model and not an accurate product depiction.

TRADEMARKS:

Amazon Web Services is a trademark of Amazon Technologies, Inc. GitHub is an exclusive trademark registered in the United States by GitHub, Inc. Intel and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. JavaScript is a registered trademark of Oracle and/or its affiliates. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries. MinIO is a registered trademark of MinIO Corporation. MinIO Warp is a trademark of MinIO Corporation. Ubuntu is a registered trademark of Canonical Ltd. XFS is a registered trademark of Silicon Graphics International Corporation or its subsidiaries in the United States and/or other countries. All other company names, product names and service names may be trademarks or registered trademarks of their respective companies.

DISCLAIMERS:

KIOXIA America, Inc. may make changes to specifications and product descriptions at any time. The information presented in this performance brief is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Any performance tests and ratings are measured using systems that reflect the approximate performance of KIOXIA America, Inc. products as measured by those tests. Any differences in software or hardware configuration may affect actual performance, and KIOXIA America, Inc. does not control the design or implementation of third party benchmarks or websites referenced in this document. The information contained herein is subject to change and may render inaccuracies for many reasons, including but not limited to any changes in product and/or roadmap, component and hardware revision changes, new model and/or product releases, software changes, firmware changes, or the like. KIOXIA America, Inc. assumes no obligation to update or otherwise correct or revise this information.

KIOXIA America, Inc. makes no representations or warranties with respect to the contents herein and assumes no responsibility for any inaccuracies, errors or omissions that may appear in this information.

KIOXIA America, Inc. specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. In no event will KIOXIA America, Inc. be liable to any person for any direct, indirect, special or other consequential damages arising from the use of any information contained herein, even if KIOXIA America, Inc. are advised of the possibility of such damages.

KIOXIA