KIOXIA

Performance Brief

Optimizing NVMe® Namespace Sizes to Accelerate Application Workloads with KIOXIA CM7 Series SSDs

Introduction

Write intensive applications require high random write performance from storage media such as SSDs. The amount of overprovisioning (OP) in SSDs has a direct impact on random write performance. Excessive write workloads will also wear out the flash memory on SSDs much faster relative to other workloads, especially on SSDs with lower Drive Writes Per Day[1] (DWPD) endurance. For these workloads, SSD endurance that is higher than 3 DWPD is required but no longer commonly available. Lastly, lower DWPD drives have less OP, which can negatively affect SSD performance.

Since most workloads have some element of randomness, random write operations are prevalent in today's data centers. Increasing random write performance on SSDs while providing a higher level of DWPD endurance is a real issue that can be addressed through NVMe namespaces.

A namespace within the NVMe specification divides an SSD into logically separate and individually addressable storage spaces with their own input/output (I/O) queue. Each namespace appears to the connected host as a distinct SSD. The host interacts with the individual namespaces in the same way as it would with local or shared NVMe storage targets. Adjusting NVMe namespace sizes can increase the performance of locally attached SSDs while emulating higher DWPD endurance. Any additional capacity that is not allocated to a namespace goes into the SSD OP pool, which enables increases in random write performance. KIOXIA CM7 Series SSDs support these capabilities enabling them to use higher amounts of OP from available capacity.

This performance brief presents how a KIOXIA CM7-R Series read intensive enterprise NVMe SSD (1 DWPD) can provide comparable performance with a higher endurance mixed use KIOXIA CM7-V Series SSD (3 DWPD) using NVMe namespaces. It also presents how a KIOXIA CM7-R Series SSD configured with a 6.4 terabyte[2] (TB) NVMe namespace can deliver significant random write performance gains over a KIOXIA CM7-R Series SSD configured with a 7.68 TB NVMe namespace. The brief also presents varying NVMe namespace sizes and their effect on random write performance. The testing included three configurations:

1. *CM7-R SSD (1 DWPD) – 7.68 TB with a namespace configured for 7.68 TB*
2. *CM7-R SSD (1 DWPD vs. 3 DWPD) – 7.68 TB with a namespace configured for 6.4 TB (to emulate a 3 DWPD drive)*
3. *CM7-V SSD (3 DWPD) – 6.4 TB with a namespace configured for 6.4 TB*

The tests were performed through Flexible I/O[3] (FIO) software and measured 100% sequential read/write throughput, 100% random read/write IOPS (input/output operations per second), random 70% read/30% write IOPS, and 100% random write IOPS of varying namespace sizes.

The test results show that a KIOXIA CM7-R Series SSD (1 DWPD) demonstrated comparable sequential read/write and random read performance when configured to emulate a 3 DWPD SSD using NVMe namespaces (reducing storage capacity for an individual namespace). With KIOXIA CM7 Series SSDs, flash memory not provisioned for a namespace was utilized by the SSD OP pool, enabling higher random write performance.

The results presented include a brief description of each workload test, a graphical depiction of the test results and an analysis. Appendix A covers the hardware and software configuration – Appendix B covers the configuration set-up and test procedures.

Test Results Snapshot

A KIOXIA CM7-R Series SSD (1 DWPD) using NVMe namespaces to emulate a 3 DWPD drive demonstrated comparable performance with a KIOXIA CM7-V Series SSD (3 DWPD):

100% Sequential Read Throughput:

Comparable ~14.5 GB/s average

100% Sequential Write Throughput:

Comparable ~7.1 GB/s average

100% Random Read IOPS:

Comparable ~2,536K IOPS average

A KIOXIA CM7-R Series SSD configured with a 6.4 TB NVMe namespace demonstrated significant write performance gains over the same KIOXIA CM7-R Series SSD configured with a 7.68 TB NVMe namespace:

100% Random Write IOPS:

88% Higher

Random 70% Read/ 30% Write IOPS:

59% Higher

Test Results

Workload 1:
100% Sequential Read Throughput

This test measured the sequential read throughput where read data of a specific size was ordered one after the other from a logical block address (LBA). This workload included 1 CPU thread with a 128 kibibyte[4] (KiB) block size and queue depth of 32. Measurements are in gigabytes per second (GB/s).

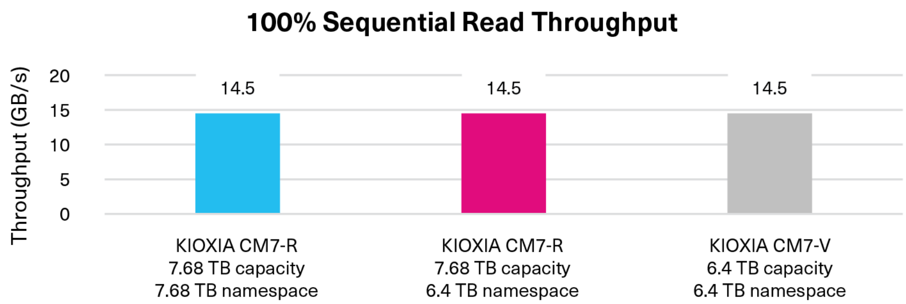

Workload 2:
100% Sequential Write Throughput

This test measured the sequential write throughput where write data of a specific size was ordered one after the other from an LBA. This workload included 1 CPU thread with a 128 KiB block size and queue depth of 32. Measurements are in GB/s.

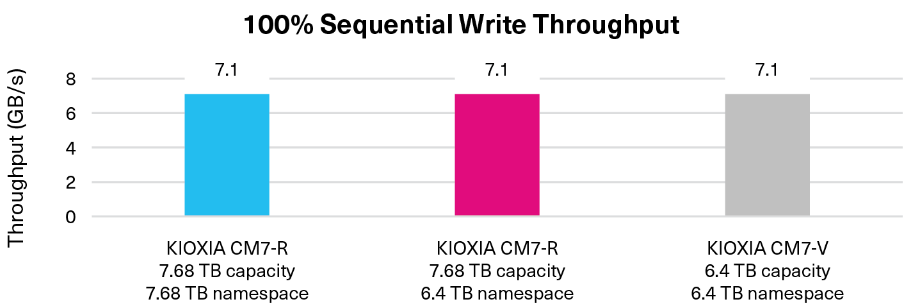

Workload 3:
100%: Random Read IOPS

This test measured the random read performance where read data of a specific size was ordered randomly from an LBA. This workload included 16 CPU threads with a 4 KiB block size and queue depth of 32 per CPU thread for an effective I/O queue depth of 512. Measurements are in IOPS.

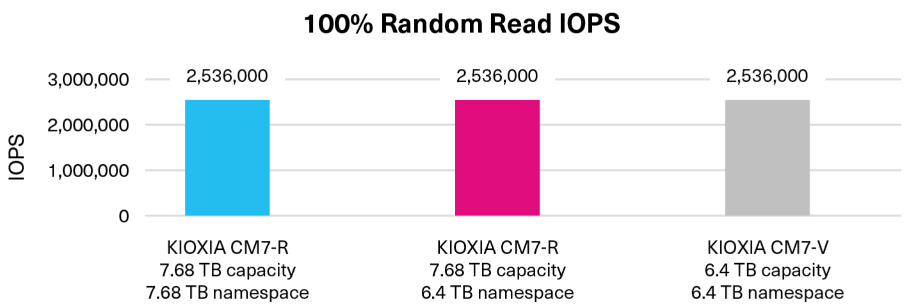

Workload 4:
100%: Random Write IOPS

This test measured the random write performance where write data of a specific size was ordered randomly from an LBA. This workload included 8 CPU threads with a 4 KiB block size and queue depth of 4 per CPU thread for an effective I/O queue depth of 32. Measurements are in IOPS.

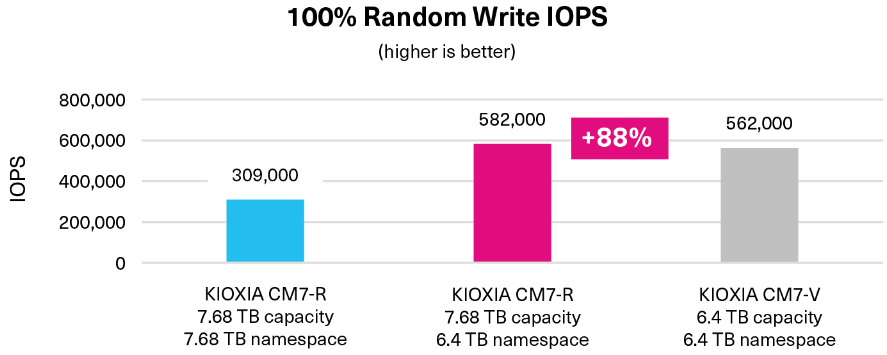

When the 7.68 TB KIOXIA CM7-R Series SSD (1 DWPD) namespace size was reduced from 7.68 TB to 6.4 TB, random write IOPS performance increased 88%. When compared with the 6.4 TB KIOXIA CM7-V Series SSD (3 DWPD), the deviation in performance was 3.5%.

Workload 5:
Random 70% Read/30% Write IOPS

This test measured random 70% read/30% write IOPS performance where mixed data of a specific size was ordered randomly from an LBA. This workload included 8 CPU threads with a 4 KiB block size and queue depth of 4 per CPU thread for an effective I/O queue depth of 32. Measurements are in IOPS. The read/write workloads for each namespace are separated to show individual random read and write IOPS performance.

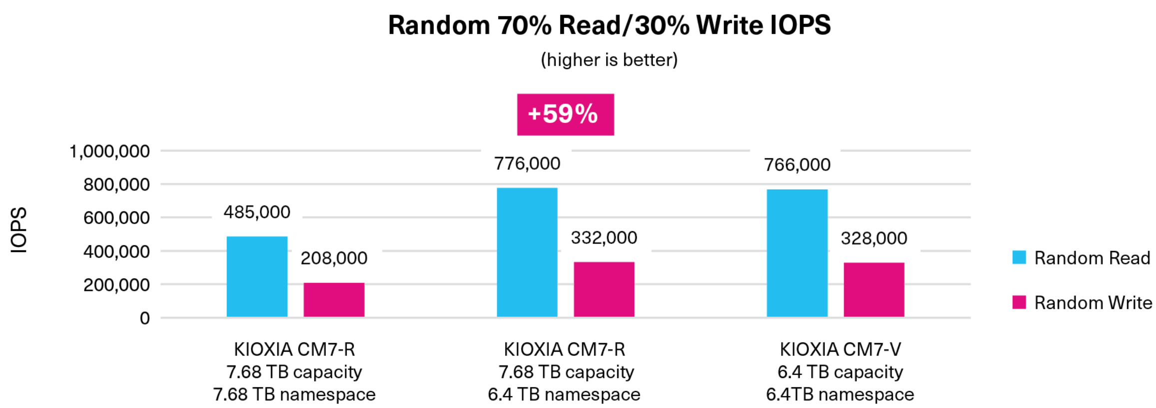

When the 7.68 TB KIOXIA CM7-R Series SSD (1 DWPD) namespace size was reduced from 7.68 TB to 6.4 TB, random 70% read/30% write IOPS performance increased 59%. When compared with the 6.4 TB KIOXIA CM7-V Series SSD (3 DWPD), the deviation in performance was slightly above 1%.

Workload 6:
100% Random Write IOPS of Varying Namespace Sizes

This test measured a 100% random write workload against varying namespace sizes to show the effect on random write performance. This workload included 8 CPU threads with a 4 KiB block size and queue depth of 4 per CPU thread for an effective I/O queue depth of 32. Measurements are in IOPS. As namespace sizes increase, random write performance decreases.

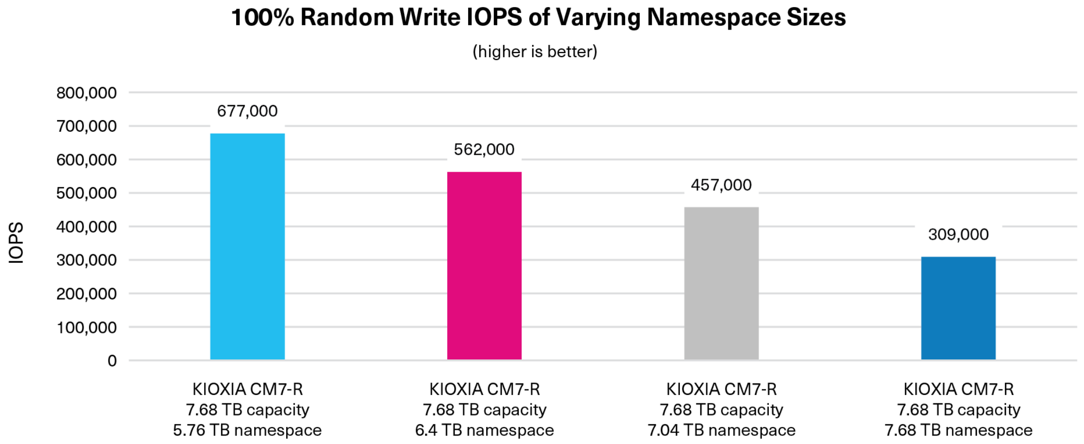

Analysis

The 7.68 TB KIOXIA CM7-R Series SSD (1 DWPD) performed comparably with the higher endurance 6.4 TB KIOXIA CM7-V Series SSD (3 DWPD) when provisioned with smaller namespaces for 100% sequential read/write and 100% random read workloads.

When the 100% random write and 70% read/30% write workloads were tested, the 7.68 TB KIOXIA CM7-R Series SSD configured with a 6.4 TB namespace delivered higher random write performance when compared with the 7.68 TB KIOXIA CM7-R Series SSD configured with a 7.68 TB namespace. Flash memory not provisioned for a namespace was utilized by the drive's OP pool.

When the 100% random write IOPS test was performed on the 7.68 TB KIOXIA CM7-R Series SSD, it included varying namespace sizes ranging from 5.76 TB to 7.68 TB. As the NVMe namespace sizes got smaller, random write IOPS performance increased, validating that the amount of OP in a drive has a direct impact on random write performance and DWPD endurance.

Summary

By adjusting the namespace size on SSDs, users are able to achieve write performance that reflects various DWPD endurance ratings. Since almost all modern workloads are inherently random, many applications can also benefit from the performance boost that can be achieved by adjusting namespace sizes, particularly as it relates to random write performance.

A KIOXIA CM7-R Series SSD with 1 DWPD endurance demonstrated higher random write performance when the NVMe namespace was reduced from 7.68 TB to 6.4 TB. Flash memory not provisioned for a namespace was utilized by the SSD OP pool and delivered higher 100% random write and random 70% read/30% write workload test results. Unused extra capacity was not stranded, but instead, was utilized to gain extra performance. The KIOXIA CM7-R Series SSD with 1 DWPD endurance also demonstrated comparable 100% sequential read/write and 100% random read performance when configured to emulate a 3 DWPD KIOXIA CM7-V Series SSD using NVMe namespaces (with reduced storage capacity for an individual namespace).

FIO Test	7.68 TB Capacity 7.68 TB Namespace	7.68 TB Capacity 6.4 TB Namespace	6.4 TB Capacity 6.4 TB Namespace	Namespace Adjustment Result
100% Random Write *(in IOPS)*	309K	582K	562K	Increased performance
Random 70% Read/30% Write *(in IOPS)*	485K *(read)* 208K *(write)* 693K *(total)*	776K *(read)* 332K *(write)* 1,108K *(total)*	766K *(read)* 328K *(write)* 1,094K *(total)*	Increased performance
100% Sequential Read *(in GB/s)*	14.5	14.5	14.5	Comparable performance
100% Sequential Write *(in GB/s)*	7.1	7.1	7.1	Comparable performance
100% Random Read *(in IOPS)*	2,536K	2,536K	2,536K	Comparable performance

This capability of adjusting NVMe namespace sizes where smaller namespaces can lead to better random write performance also applies to KIOXIA CD8/CD8P Series and KIOXIA XD7P Series data center NVMe SSDs.

KIOXIA CM7 Series SSD Product Info

The latest generation KIOXIA CM7 Series enterprise NVMe SSDs support Enterprise and Datacenter Standard Form Factor (EDSFF) E3.S and 2.5-inch[5] form factors and are compliant with the NVMe 2.0 and PCIe® 5.0 specifications. These SSDs are available in two configurations: CM7-R Series for read intensive applications (1 DWPD, up to 30.72 TB capacities) and CM7-V Series for higher endurance mixed use applications (3 DWPD, up to 12.80 TB capacities). Additional features include a dual-port design for high availability applications, Flash Die Failure Protection and security options[6].

KIOXIA CM7 Series SSD[7]

Additional KIOXIA CM7 Series SSD information is available here.

Appendix A

Hardware/Software Test Configuration

Server Information	
Server	Supermicro® AS-1125HS-TNR
Number of Servers	1
Number of CPU Sockets	2
CPU	AMD EPYC™ 9534
Number of CPU Cores	64
CPU Frequency	2.45 gigahertz (GHz)
Total Memory	768 GB DDR5 DRAM
Memory Frequency	DDR5-4800

Operating System Information	
Operating System	Ubuntu®
Version	22.04.3
Kernel	5.15.0-78-generic

Test Software Information	
Test Software	FIO
Version	3.28

SSD Information		
Model	KIOXIA CM7-R Series	KIOXIA CM7-V Series
Interface	PCIe 5.0 x4	PCIe 5.0 x4
Number of Devices	1	1
Form Factor	2.5-inch (U.2)	2.5-inch (U.2)
Capacity	7.68 TB	6.4 TB
DWPD	1 (5 years)	3 (5 years)
Active Power	up to 25 watts	up to 25 watts

KIOXIA

Appendix B

Configuration Set-up/Test Procedures

Configuration Set-up

The Ubuntu® 22.04.3 Linux® operating system was installed on the Supermicro® AS-1125HS-TNR server.

FIO version 3.28 test software was installed on the OS.

One 7.68 TB KIOXIA CM7-R Series SSD (1 DWPD) was installed into the Supermicro AS-1125HS-TNR server.

One 6.4 TB KIOXIA CM7-V Series SSD (3 DWPD) was installed into the Supermicro AS-1125HS-TNR server.

Test Procedures

Three test configurations included:

- One 7.68 TB KIOXIA CM7-R Series SSD with a namespace configured for 7.68 TB
- The same 7.68 TB KIOXIA CM7-R Series SSD with a namespace configured for 6.4 TB
- One 6.4 TB KIOXIA CM7-V Series SSD with a namespace configured for 6.4 TB

The following tests were run on the three test configurations and results were recorded for each:

- 100% Sequential Read Throughput
- 100% Sequential Write Throughput
- 100% Random Read IOPS
- 100% Random Write IOPS
- Random 70% Read/30% Write IOPS

Four namespace sizes (5.76 TB, 6.4 TB, 7.04 TB and 7.68 TB) were configured on a 7.68 TB KIOXIA CM7-R Series SSD and the 100% random write workload test was run on each configuration.

NOTES:

[1] Drive Write(s) Per Day (DWPD): 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.

[2] 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.

[3] Flexible I/O (FIO) is a free and open source disk I/O tool used both for benchmark and stress/hardware verification. The software displays a variety of I/O performance results, including complete I/O latencies and percentiles.

[4] One kibibyte (KiB) is equal to 1,024 bytes (or 2^10 bytes).

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

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

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