KIOXIA

Performance Brief

Accelerating Application Workloads Using NVMe™ Namespaces with KIOXIA CM7 Series SSDs

Introduction

Applications constantly generate and write data, which creates the need for storage devices that can handle these write intensive workloads. SSDs are generally available with 1 and 3 Drive Writes Per Day[1] (DWPD) endurance that are referred to as read intensive and mixed use drives, respectively. However, some write workloads need 10 DWPD endurance or higher which are not commonly available. Higher DWPD write endurance can be achieved by utilizing NVMe namespaces to configure a lower DWPD drive to use higher amounts of overprovisioning (OP) from available capacity. This can effectively emulate the random write performance of a 3 DWPD or 10 DWPD SSD.

A namespace within the NVMe specification divides a SSD into logically separate and individually addressable storage spaces with its own input/output (I/O) queue. They are created at the controller level and appear as separate SSDs to the connected host that interacts with them as it would with local or shared NVMe targets.

Why use namespaces? When applications run simultaneously, sequential and random I/O data streams are written to and read from SSDs. A single sequential I/O stream mixes with other I/O streams (sequential or random) and becomes a random workload once it reaches the SSD. As multiple servers process these varied workloads and move data to SSDs, the activity changes from sequential or random read/write workloads into a large mix of random read/write I/Os (I/O blender effect) from the SSD perspective. This can lead to performance drops and latency delays. Using NVMe namespaces to increase the random write performance as well as the DWPD of each SSD can help systems better address these complex I/O application requirements.

This performance brief presents how a KIOXIA CM7-R Series read intensive PCIe® 5.0 enterprise NVMe SSD (1 DWPD) can perform comparable random write performance when configured to emulate a higher endurance 3 DWPD or 10 DWPD mixed use SSD by utilizing NVMe namespaces[2]. These capabilities were tested in four configurations:

1. *CM7-R SSD (1 DWPD) – 3.84 terabytes[3] (TB) with a namespace configured for 3.84 TB*
2. *CM7-V SSD (3 DWPD) – 3.2 TB with a namespace configured for 3.2 TB*
3. *CM7-R SSD (1 DWPD vs. 3 DWPD) – 3.84 TB with a namespace configured for 3.2 TB (emulating a 3 DWPD drive)*
4. *CM7-R SSD (1 DWPD vs. 10 DWPD) – 3.84 TB with a namespace configured for 1.6 TB (emulating a 10 DWPD drive)*

The tests were performed through Flexible I/O[4] (FIO) software and measured 100% sequential read/write throughput, 100% random read/write IOPS (input/output operations per second) and random IOPS using 70% read/30% write and 30% read/70% write ratios.

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 or 10 DWPD SSD using NVMe namespaces with reduced 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, which in turn enabled higher random write and mixed random 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) demonstrated comparable performance with a KIOXIA CM7-V Series SSD (3 DWPD) using NVMe namespaces (with reduced storage capacity for an individual namespace):

100% Sequential Read:
Comparable
~14.5 GB/s average

100% Sequential Write:
Comparable
~7.1 GB/s average

100% Random Read:
Comparable
~2,814K IOPS average

A KIOXIA CM7-R Series SSD also demonstrated significant write performance gains over a KIOXIA CM7-V Series SSD using NVMe namespaces to emulate a 3 DWPD or 10 DWPD drive:

100% Random Write:
From 94% to 192% Higher

Random 70% Read / 30% Write:
From 55% to 116% Higher

Random 30% Read / 70% Write:
From 76% to 200% Higher

Test Results

Workload 1: 100% Sequential Read

This test measured the sequential read throughput of one thread with a 128K block size and queue depth of 32. Measurements are in gigabytes per second (GB/s).

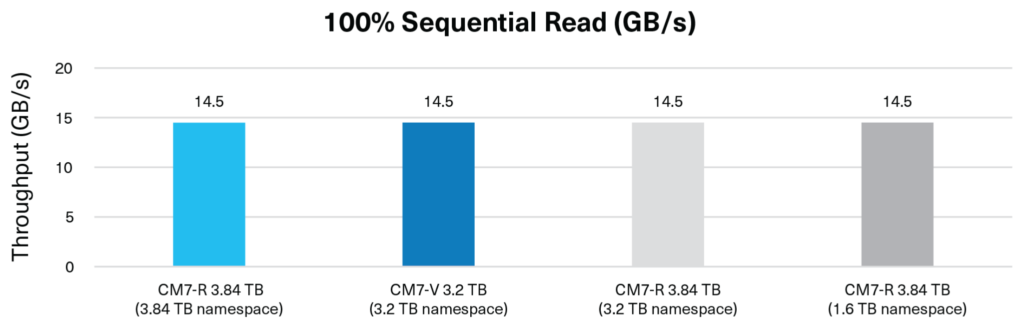

Workload 2: 100% Sequential Write

This test measured the sequential write throughput of one thread with a 128K block size and queue depth of 32. Measurements are in GB/s.

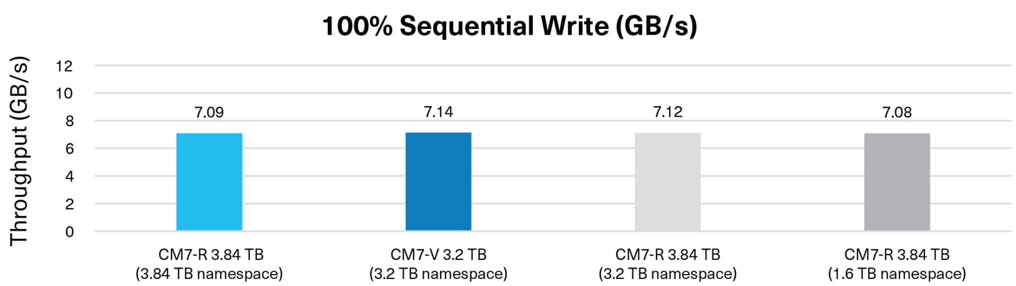

Workload 3: 100%: Random Read

This test measured the random read performance of 16 threads with a 4K block size and queue depth of 32 per thread for an effective I/O queue depth of 512. Measurements are in IOPS.

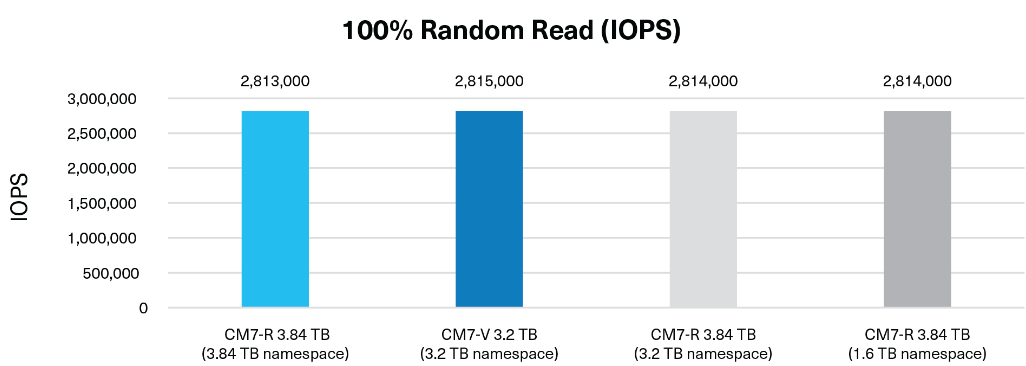

Workload 4: 100% Random Write

This test measured the random write performance of 8 threads with a 4K block size and queue depth of 4 per thread for an effective I/O queue depth of 32. Measurements are in IOPS. The percentage gains for each namespace were compared against the KIOXIA CM7-R 3.84 TB namespace and not against the prior measurement.

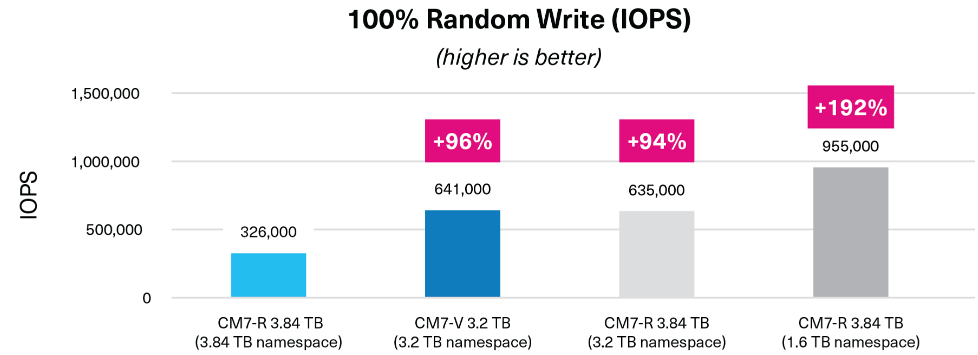

Workload 5: Random 70% Read / 30% Write

This test measured mixed 70% read / 30% write random performance of 8 threads with a 4K block size and queue depth of 4 per thread for an effective I/O queue depth of 32. Measurements are in IOPS. The read/write workloads for each namespace were separated to show individual IOPS performance. The percentage gains for each namespace were compared against the KIOXIA CM7-R 3.84 TB namespace and not against the prior measurement, and included the aggregate read/write measurement.

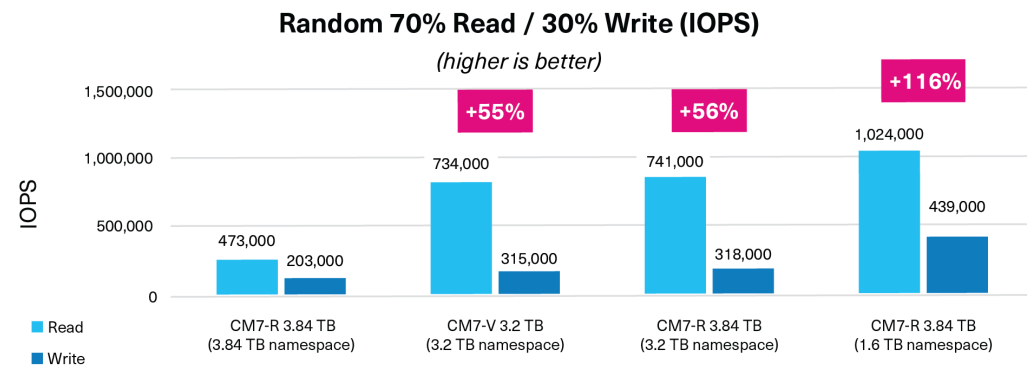

Workload 6: Random 30% Read / 70% Write

This test measured mixed 30% read / 70% write random performance of 8 threads with a 4K block size and queue depth of 4 per thread for an effective I/O queue depth of 32. Measurements are in IOPS. The read/write workloads for each namespace were separated to show individual IOPS performance. The percentage gains for each namespace were compared against the KIOXIA CM7-R 3.84 TB namespace and not against the prior measurement, and included the aggregate read/write measurement.

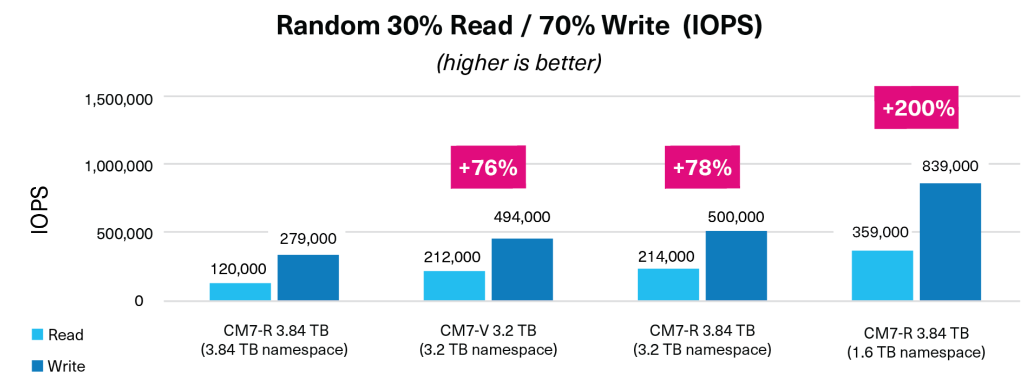

Analysis

The 3.84 TB KIOXIA CM7-R Series SSD (1 DWPD) performed comparably with the higher endurance 3.2 TB KIOXIA CM7-V Series SSD (3 DWPD) when provisioned with smaller namespaces for sequential read/write and random read workloads. These results are validated from the sequential read/write and random read FIO tests performed.

When the 100% random write and mixed workloads were tested, the 3.84 TB KIOXIA CM7-R Series SSD configured with a 3.2 TB namespace and a 1.6 TB namespace delivered higher random write performance when compared with the 3.2 TB KIOXIA CM7-V Series SSD. Flash memory not provisioned for a namespace was utilized by the drive's OP pool.

When the 3.84 TB KIOXIA CM7-R Series SSD (1 DWPD) was configured for a 1.6 TB namespace (emulating 10 DWPD effective endurance), the random 70% read / 30% write workload delivered a whopping 1,463K IOPS or three times the aggregate performance of the 3.84 TB KIOXIA CM7-R Series SSD. Proper configuration of NVMe namespaces can increase random write performance, better utilize unused SSD capacity and enable a 1 DWPD drive to perform comparably to a higher endurance drive.

Summary

Today's workloads are typically mixed. NVMe namespaces provide the ability to adjust the random write performance of SSDs to better address mixed and I/O blender effected workloads, providing maximum performance simply by giving up unused capacity. This also enables drives to be available for other applications and workloads, maximizing their use as well as their lifecycles. Additionally, namespaces can provide higher effective DWPD and help offset the limited availability of 10 DWPD SSD endurance options that served high write workloads in the past.

The test results demonstrate that the 3.84 TB KIOXIA CM7-R Series SSD with a 3.84 TB namespace performed comparably in random write performance to the 3.2TB KIOXIA CM7-V Series SSD and when the KIOXIA CM7-R Series SSD namespace was reduced to 3.2 TB capacity (to emulate a 3 DWPD drive). Random write and mixed workload performance were improved even more when the 3.84 TB KIOXIA CM7-R Series SSD used a namespace of 1.6 TB (to emulate a 10 DWPD drive). It also showed that the flash memory not provisioned for a namespace was utilized by the drive's OP pool resulting in higher write performance to the 100% random write workload and the mixed workloads tested.

FIO Test	3.84 TB Capacity 3.84 TB Namespace	3.2 TB Capacity 3.2 TB Namespace	3.84 TB Capacity 3.2 TB Namespace	3.84 TB Capacity 1.6 TB Namespace
100% Sequential Read (in GB/s)	14.5	14.5	14.5	14.5
100% Sequential Write (in GB/s)	7.09	7.14	7.12	7.08
100% Random Read (in IOPS)	2,813K	2,815K	2,814K	2,814K
100% Random Write (in IOPS)	326K	641K	635K	955K
Random 70% Read / 30% Write (in IOPS)	473K (read) 203K (write) 676K (total)	734K (read) 315K (write) 1,049K (total)	741K (read) 318K (write) 1,059K (total)	1,024K (read) 439K (write) 1,463K (total)
Random 30% Read / 70% Write (in IOPS)	120K (read) 279K (write) 399K (total)	212K (read) 494K (write) 706K (total)	214K (read) 500K (write) 714K (total)	359K (read) 839K (write) 1,198K (total)

The testing principals where smaller namespaces can equal better random write performance would also apply to other KIOXIA NVMe SSDs such as the latest KIOXIA CD8/CD8P Series and KIOXIA XD7P Series data center NVMe SSDs.

KIOXIA CM7 Series SSD Product Info

KIOXIA recently launched its latest generation CM7 Series SSDs that support Enterprise and Datacenter Standard Form Factor (EDSFF) E3.S and 2.5-inch[5] form factors and compliant with the NVMe 2.0 and PCIe 5.0 specifications. These SSDs are available in two configurations: Read intensive 1 DWPD with capacities up to 30.72 TB and higher endurance mixed use 3 DWPD with capacities up to 12.80 TB. Additional features include a dual-port design for high availability applications, flash die failure protection to maintain reliability in case of a die failure and security options[6]. Additional KIOXIA CM7 Series SSD information is available here.

KIOXIA CM7 Series SSD[7]

Appendix A

Hardware/Software Test Configuration

Server Information	
Server	HPE® ProLiant® DL360 Gen 11
Number of Servers	1
Number of CPU Sockets	2
CPU	Intel® Xeon® Gold 6426Y
Number of CPU Cores	16
CPU Frequency	2.5 GHz
Total Memory	128 GB DDR5 DRAM
Memory Frequency	4.8 GHz
Operating System Information	
Operating System	Ubuntu®
Version	22.04.2
Kernel	5.15.0-76-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	EDSFF E3.S[8]	EDSFF E3.S[8]
Capacity	3.84 TB	3.2 TB
DWPD	1 (5 years)	3 (5 years)
Active Power	up to 24 watts	up to 24 watts

Appendix B

Configuration Set-up/Test Procedures

Configuration Set-up

- The Ubuntu 22.04.2 Linux® operating system was installed on the HPE DL360 Gen 11 server.
- FIO benchmark software version 3.28 was installed on the OS.
- One 3.84 TB KIOXIA CM7-R Series SSD at 1 DWPD was installed into the HPE DL360 Gen 11 server.
- One 3.20 TB KIOXIA CM7-V Series SSD at 3 DWPD was installed into the HPE DL360 Gen 11 server.

Test Procedures

- Each of the six FIO performance tests was run on the KIOXIA CM7-R Series SSD using 3.84 TB capacity, 1 DWPD and a 3.84 TB namespace size. The test results was recorded.
- Each of the six FIO performance tests were run on the KIOXIA CM7-V Series SSD using 3.2 TB capacity, 3 DWPD and a 3.2 TB namespace size. The test results was recorded.
- Each of the six FIO performance tests were run on the KIOXIA CM7-R Series SSD using 3.84 TB capacity, 1 DWPD and a 3.2 TB namespace size to emulate a 3 DWPD drive. The test results was recorded.
- Each of the six FIO performance tests were run on the KIOXIA CM7-R Series SSD using 3.84 TB capacity, 1 DWPD and a 1.6 TB namespace size to emulate a 10 DWPD drive. The test results was recorded.

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] The test comparison emulates a 3 DWPD SSD and a 10 DWPD SSD and has no warranty impact on the KIOXIA CM7-R Series SSD or the KIOXIA CM7-V Series SSD tested. In addition, KIOXIA America, Inc. does not offer any warranty or extended warranty for these products to achieve 10 DWPD endurance.

[3] 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.

[4] 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.

[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.

[8] For this comparison, the CM7-R Series and CM7-V Series EDSFF E3.S form factor had no performance impact to the tests.

TRADEMARKS:

HPE and ProLiant are registered trademarks of Hewlett-Packard Enterprise Company and/or its affiliates. Intel and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries. NVMe is a registered or unregistered trademark of NVM Express, Inc. in the United States and other countries. PCIe is a registered trademark of PCI-SIG. Ubuntu is a registered trademark of Canonical Ltd. 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.