



Open-E JovianDSS

KIOXIA FL6 800 GB

Enterprise SCM NVMe™

SSD



Certification report

Release date: 2022.03.10

Table of contents

Introduction	2
Tested device description	3
The following table includes the KIOXIA FL6 NVMe™ SSDs hardware specifications.	3
Test environment description	3
Raw device tests	5
4.1. Test description	5
4.2. Test results	5
4.3. Test conclusions	7
All-Flash storage test	8
5.1. Test description	8
5.2. Test results	8
5.3. Test conclusions	12
Write log device tests	13
6.1 Test description	13
6.2. Zpool with 12x 2-way mirror HDD data groups	13
6.3. Zpool with 6x RAID-Z2 data groups	15
6.4 Test conclusions	17
7. Read cache device test	17
7.1. Test description	17
7.2. Zpool with 12x 2-way mirror HDD data groups	18
7.3. Zpool with 6x Raid-Z2 HDD data groups	19
7.4. Test conclusions	20
8. HA non-shared storage cluster compatibility test	21
8.1. Functional tests	21
8.2. Test conclusions	21
9. Summary	21

1. Introduction

The following certification report aims to present the results of various compatibility and performance tests conducted on the KIOXIA FL6 NVMe™ SSDs in combination with Open-E JovianDSS, along with a description of the testing methods used. The results obtained from the testing, along with technical specifications, were used as the basis for recommendations on specific roles in which to use the tested drives in Open-E JovianDSS-based systems. The devices were tested in the following data storage scenarios:

Serving as cache devices supporting both write operations (write log) and read operations (read cache),
Data drives in the All-Flash data storage appliance.

For the above-mentioned cases, functional and performance testing was performed for single node configuration and functional testing for High Availability non-shared storage cluster configuration. Detailed descriptions for each of the tests are included in the appropriate report chapters.

2. Tested device description

The following table includes the KIOXIA FL6 NVMe™ SSDs hardware specifications.

Table 1. KIOXIA FL6 NVMe™ SSDs hardware specifications.

Product name	Kioxia FL6
Model name	KFL61HUL800G
Storage capacity	800 GB
Form factor	2.5-inch U.3 conformant
Interface	PCIe 4.0, NVMe
Flash memory type	XL-Flash
Power Loss Data Protection	Yes
Security	SIE, SED (TCG Opal/Ruby), FIPS 140-2

3. Test environment description

Hardware specifications of environments used during certification testing are included in the following table.

Table 2. Hardware specifications.

System name	Intel M50CYP2SBSTD
CPU	2x Intel(R) Xeon(R) Gold 6334 CPU @ 3.60GHz
RAM	16x Samsung M393A2K40DB3-CWE 3200 MHz 16 GB DDR4
Storage controller	Broadcom / LSI 9400 SAS HBA (for HDD)
Storage devices	<ul style="list-style-type: none">4x Kioxia FL6 KFL61HUL800G

KIOXIA FL6 800GB Enterprise SCM NVme™ SSD – Certification report

	<ul style="list-style-type: none"> 24x Toshiba AL14SEB030NY SAS-3 10k 300 GB HDD
System	Open-E JovianDSS 1.0up29r2 b46777

Described in the tables below (Table 3 and Table 4) are ZFS and FIO parameters that apply to all conducted tests.

Table 3. ZFS parameters used during tests.

Test specifications		
zvol	Size	200 GB
	Provisioning	thin
	Deduplication	disabled
	Number of data copies	1
	Compression	lz4
	Volblocksize	4 kB (random workload); 64 (sequential workload)
	Sync	always
	Logbias	latency
	ARC scope	all
	L2ARC scope	all
zpool	ARC size	~180 GB (sequential and random write); ~15 GB (other cases)
	Autotrim	enabled

Table 4. FIO parameters used for testing.

Test time	90 s
Ramp time	10 s
Test area size	50 GB
Direct IO	yes
IO engine	libaio

Note: Zvols were initialized after creation by writing random data on their entire space.

Note: All performance tests were conducted locally in a single node configuration.

4. Raw device tests

In order to be able to properly interpret the results acquired from the tests of the KIOXIA FL6 NVMe™ SSDs, it's necessary to measure their raw performance. This benchmark then allows for a comparison to be made between the actual performance obtained in the test environment with the information provided by the vendor and guarantees that there are no bottlenecks on a hardware level.

4.1. Test description

Applied test cases are described in Table 5. In all instances, every combination of thread numbers (1, 4, 8, 16) and queue depths (1, 16, 64, 128) was applied to the FIO test tool.

Table 5. Test cases description for raw disk tests

Test case	IO pattern	Read to write %	Block size
Random read	random	100/0	4 kB
Random write	random	0/100	4 kB
Sequential read	sequential	100/0	1 MB
Sequential write	sequential	0/100	1 MB
Mixed 70/30	random	70/30	4 kB

4.2. Test results

Figures 1 to 5 present the results of every test case applied to the raw device.

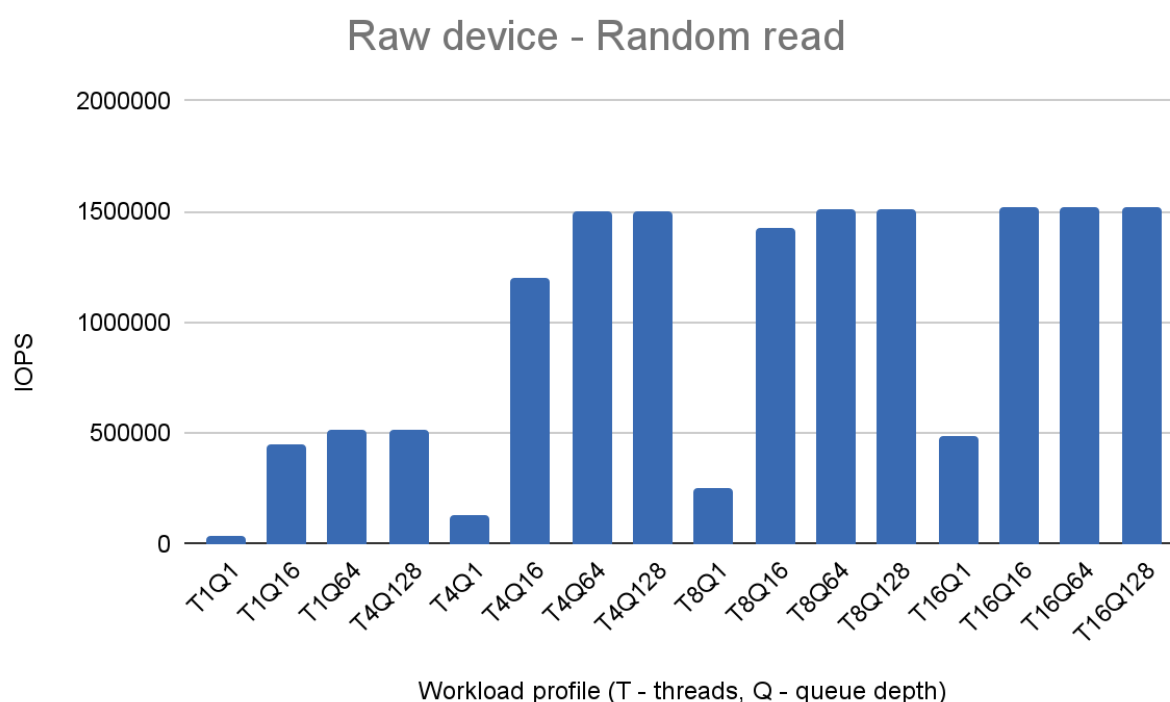


Fig. 1. Random read performance on the raw device.

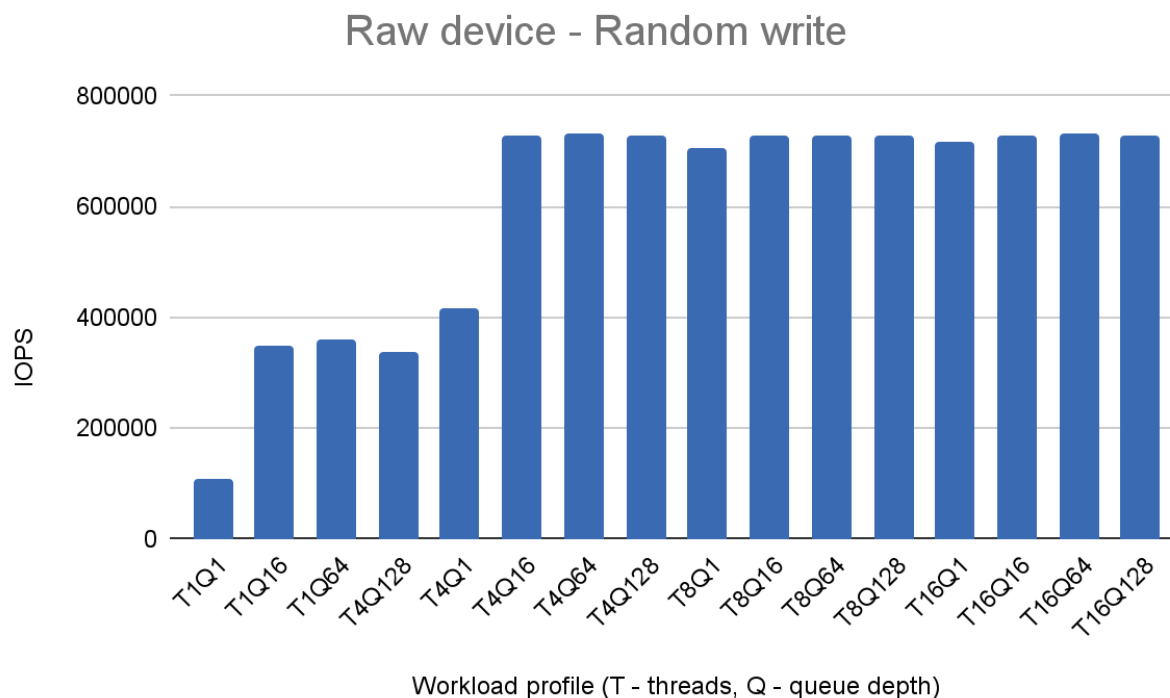


Fig. 2. Random write performance on the raw device.

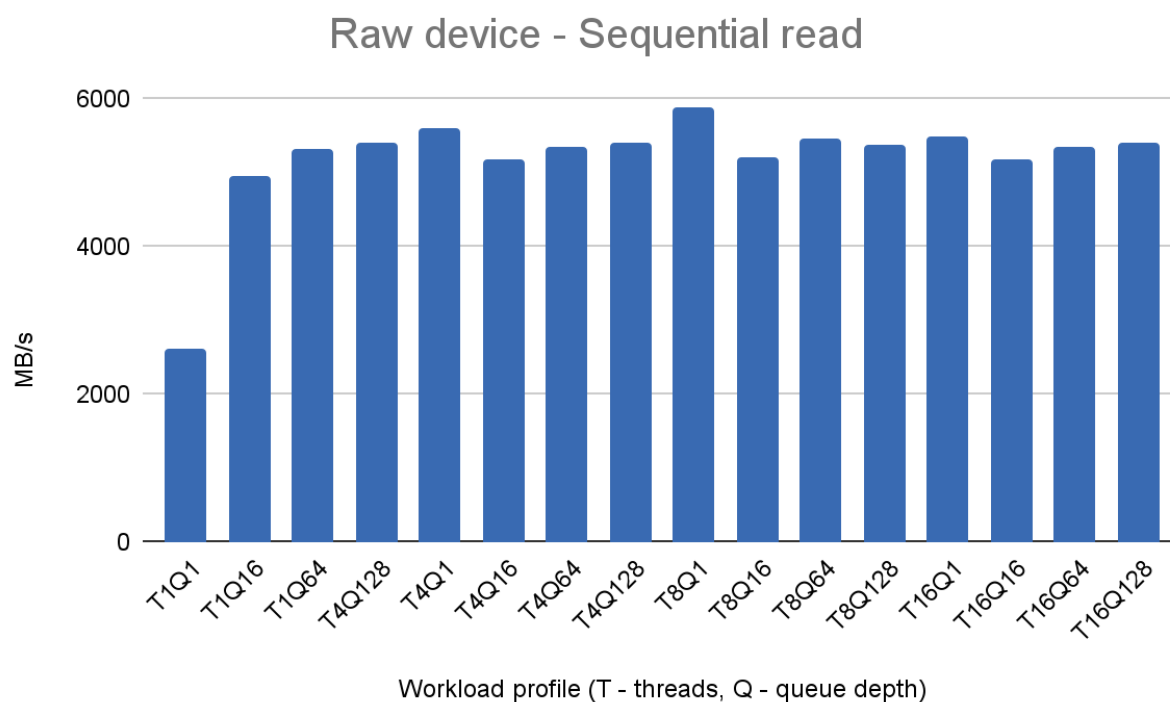


Fig. 3. Sequential read performance on the raw device.

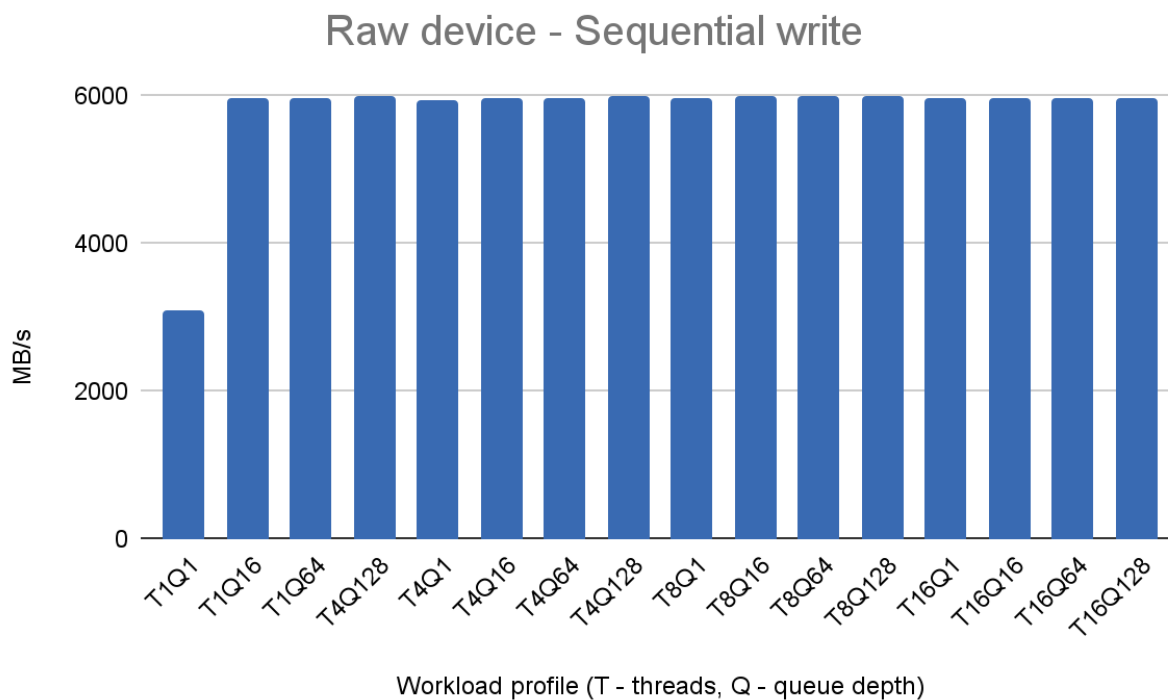


Fig. 4. Sequential write performance on the raw device.

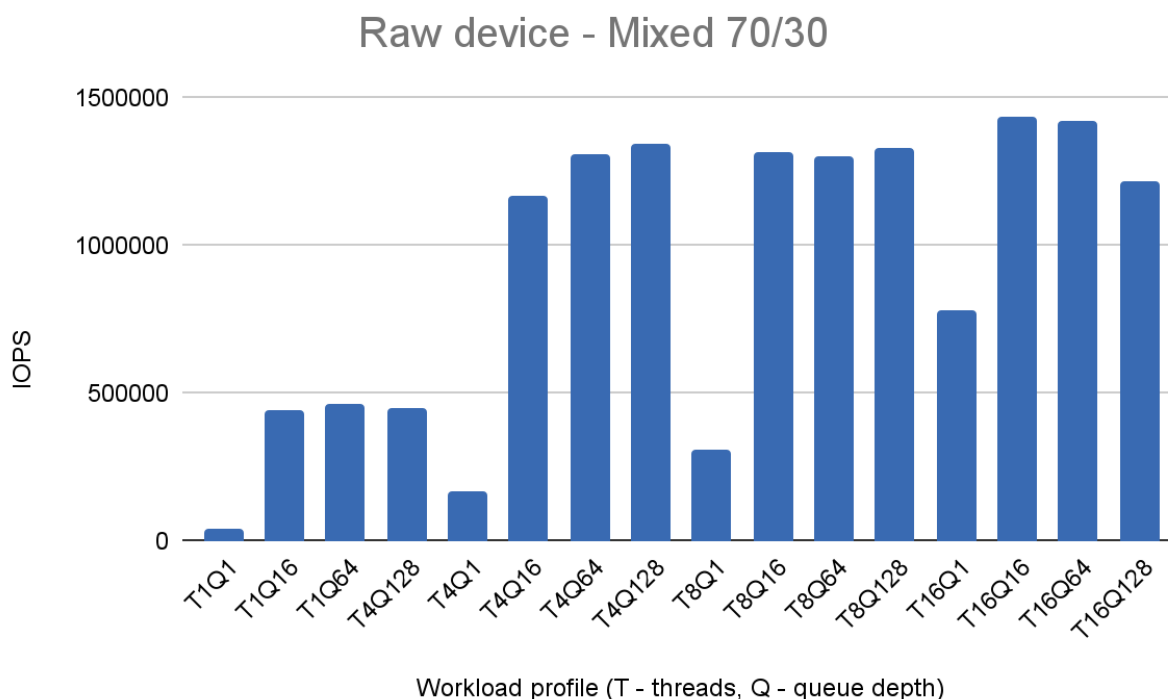


Fig. 5. Mixed 70/30 performance on the raw device.

4.3. Test conclusions

All performance tests carried out on a raw device showed proper operation at the hardware level in the test environment. The results obtained are at the same level as those declared by the manufacturer, except for the

random write test results, which turned out to be two times higher for a large number of threads and queue depths, than what was expected from the specifications.

5. All-Flash storage test

5.1. Test description

Tests of the KIOXIA FL6 NVMe™ SSDs, when used as data drives in the All-Flash storage solution powered by Open-E JovianDSS, were carried out on a zpool consisting of 4 single-drive data groups. When read tests were performed, the size of an ARC was artificially decreased (both for random and sequential cases) in order to make sure that most IOs would be served from the actual storage, not from RAM. For write workloads, the size of the ARC was left unchanged.

The measured performance was additionally compared to another NVMe SSD drive with similar specifications.

5.2. Test results

The following figures present the acquired performance results in comparison with another NVMe drive model. For random read and write workloads, latency distribution in T1Q1 and T16Q1 cases is also included.

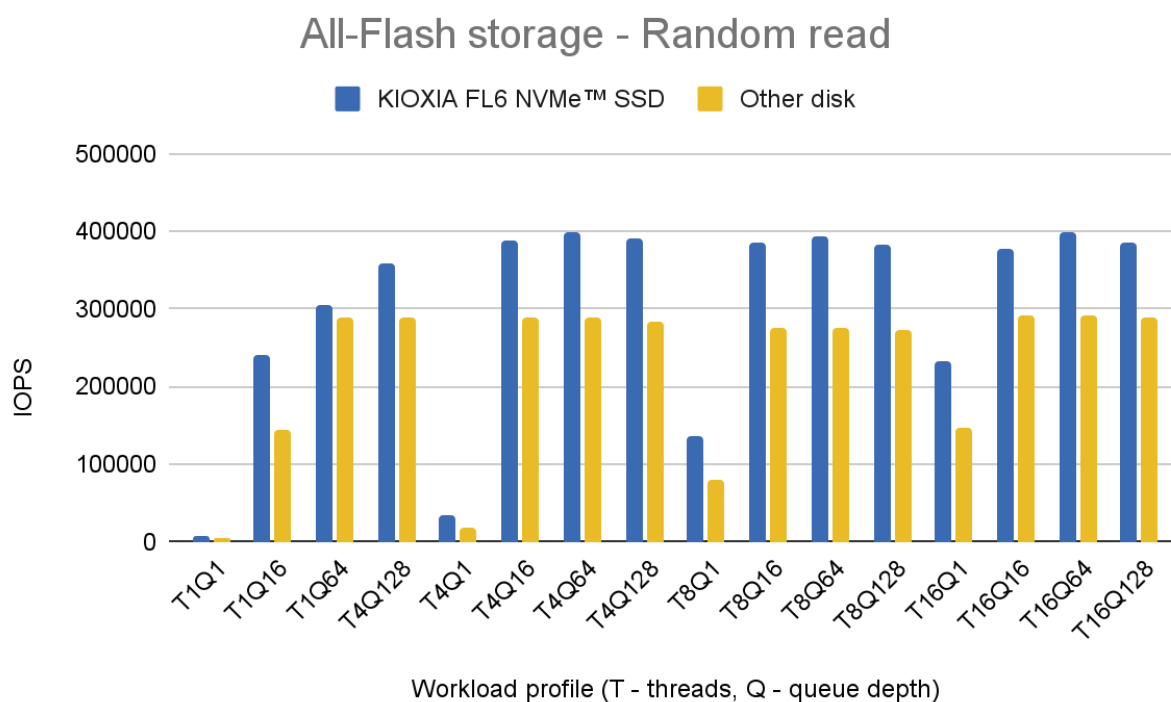


Fig. 6 Random read performance of the All-Flash storage, compared to a similar device.

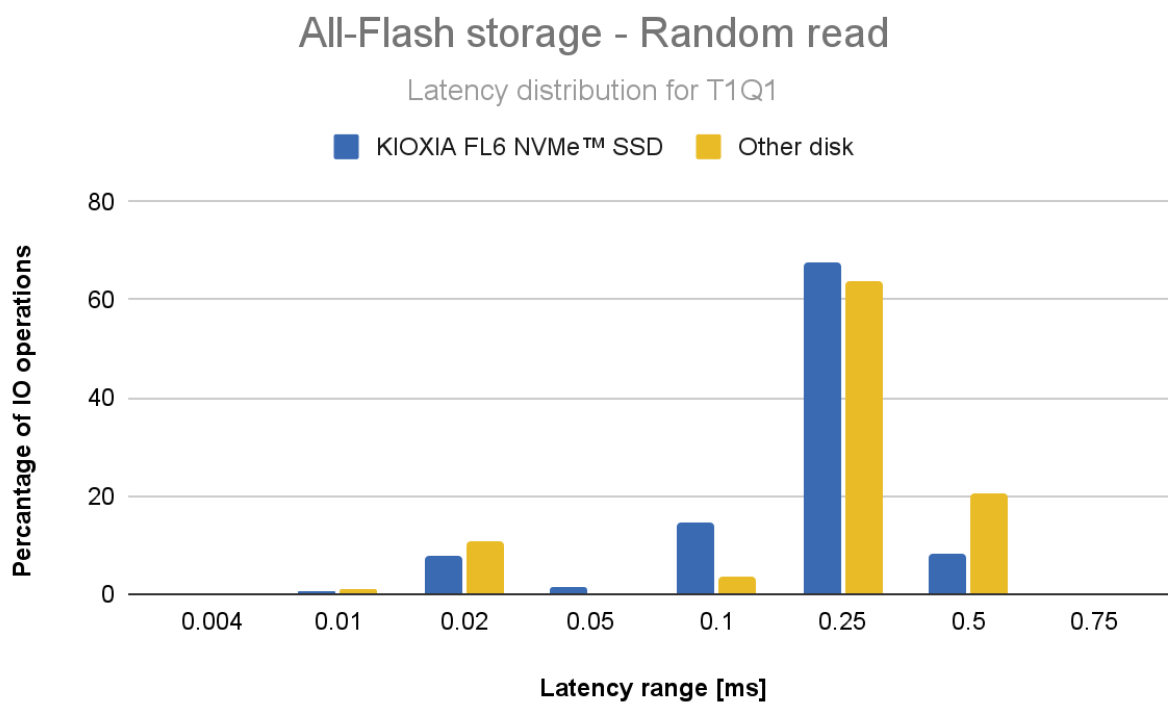


Fig. 7 Random read latency distribution for T1Q1 on the All-Flash storage, compared to a similar device.

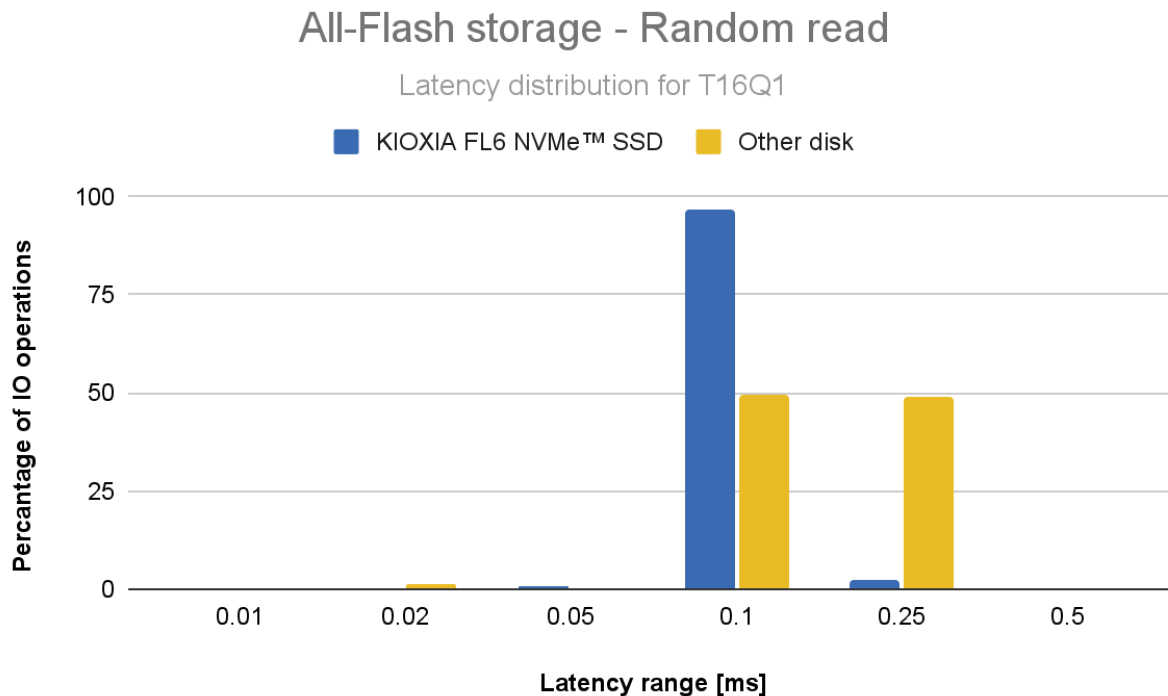


Fig. 8 Random read latency distribution for T16Q1 on the All-Flash storage, compared to a similar device.

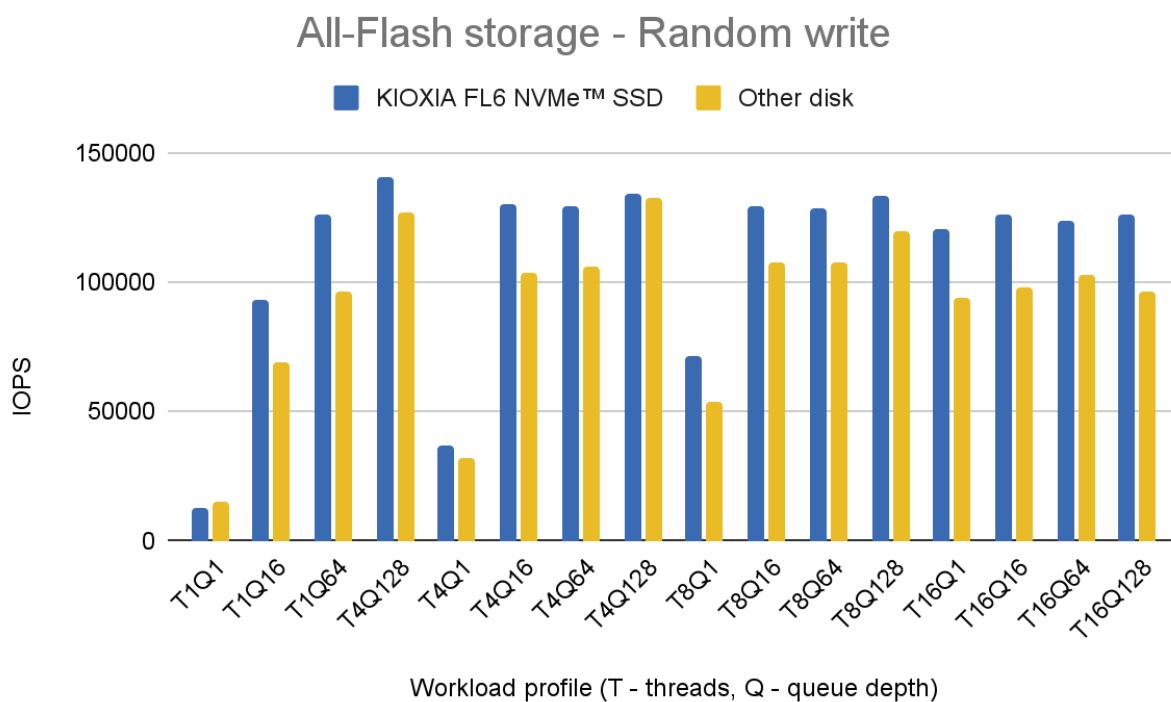


Fig. 9 Random write performance of the All-Flash storage, compared to a similar device.

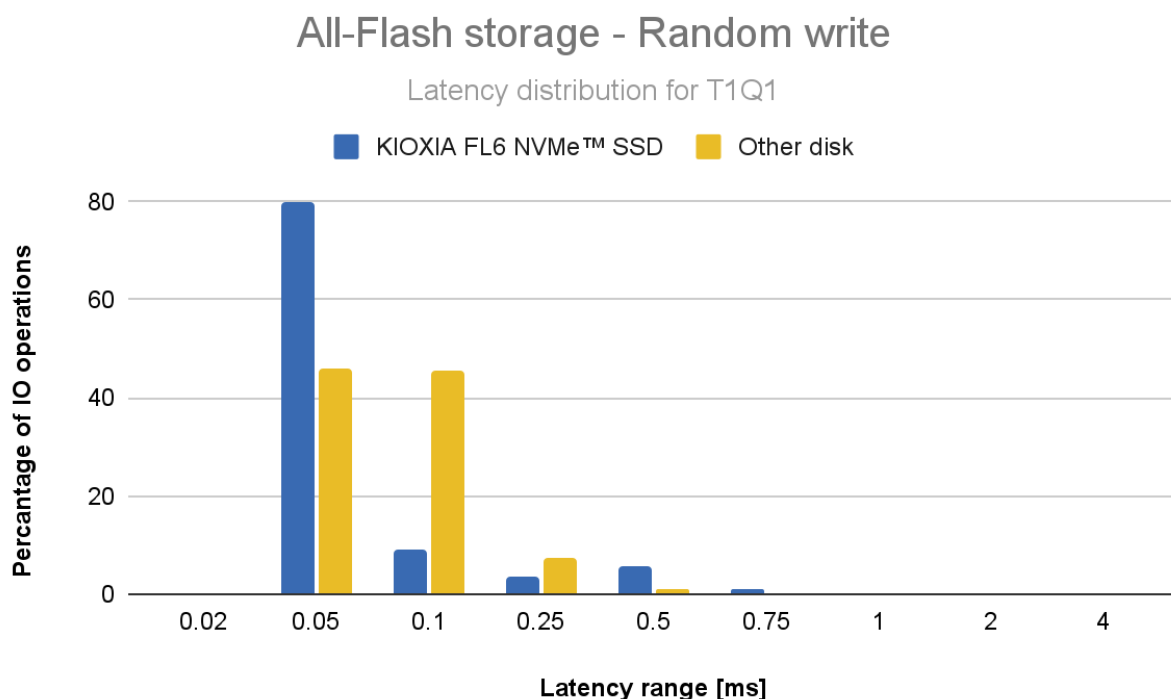


Fig. 10 Random write latency distribution for T1Q1 on the All-Flash storage, compared to a similar device.

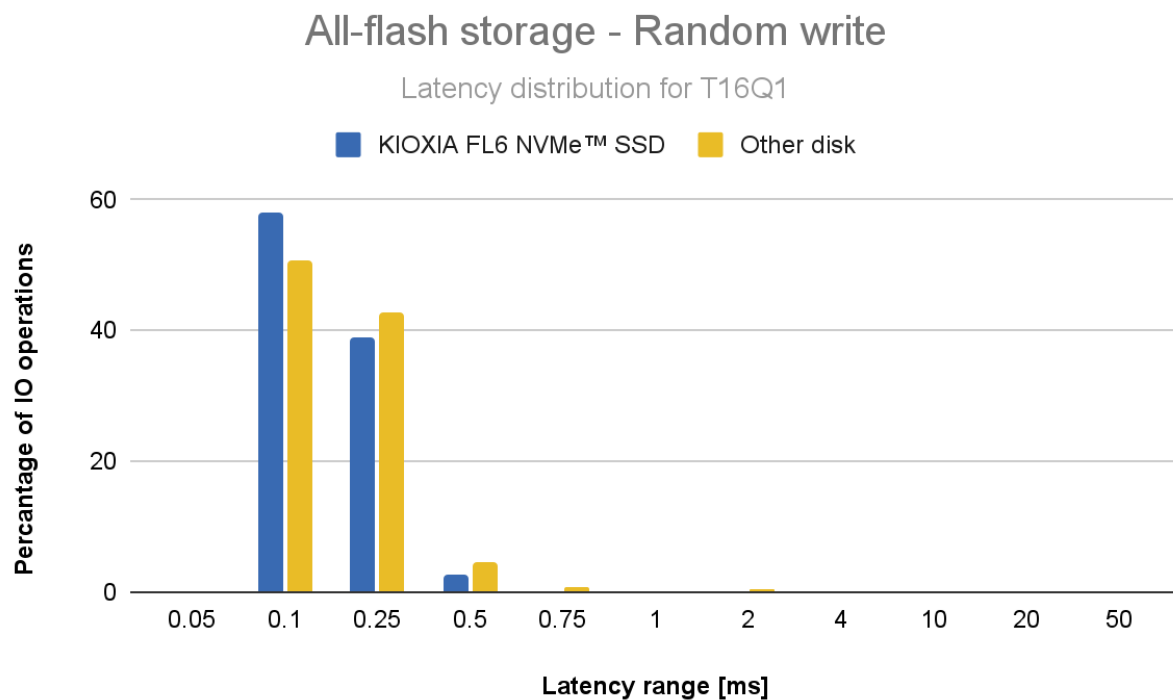


Fig. 11 Random write latency distribution for T16Q1 on the All-Flash storage, compared to a similar device.

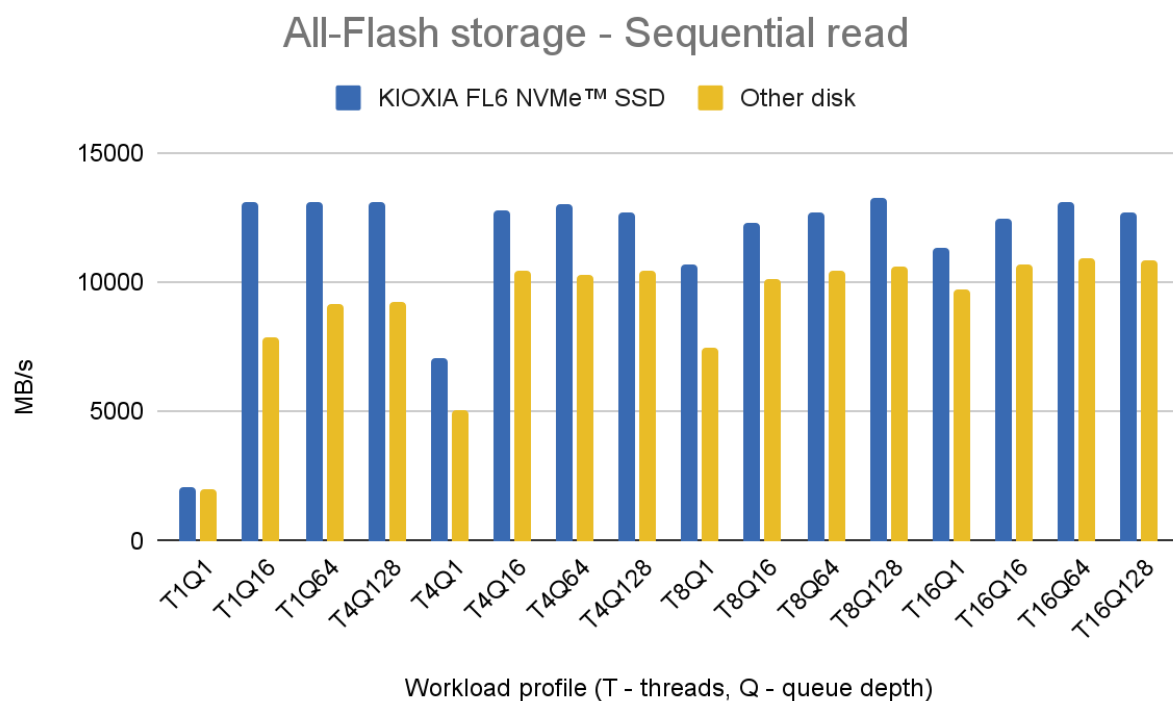


Fig. 12 Sequential read performance on the All-Flash storage, compared to a similar device.

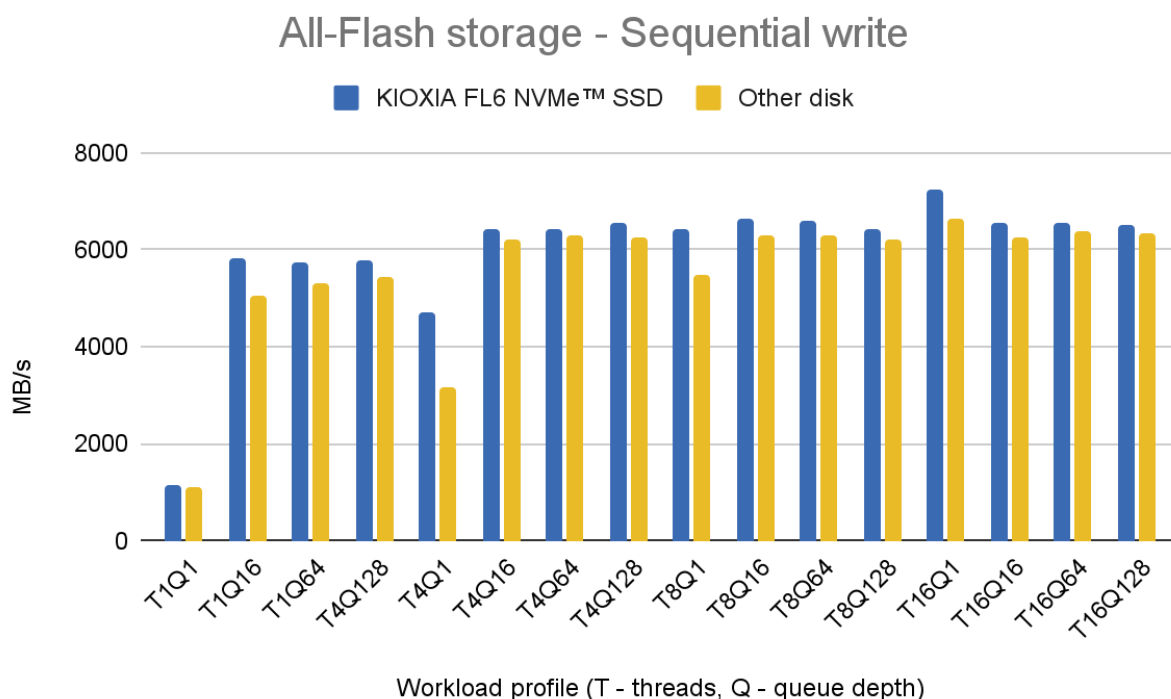


Fig. 13 Sequential write performance on the All-Flash storage, compared to a similar device.

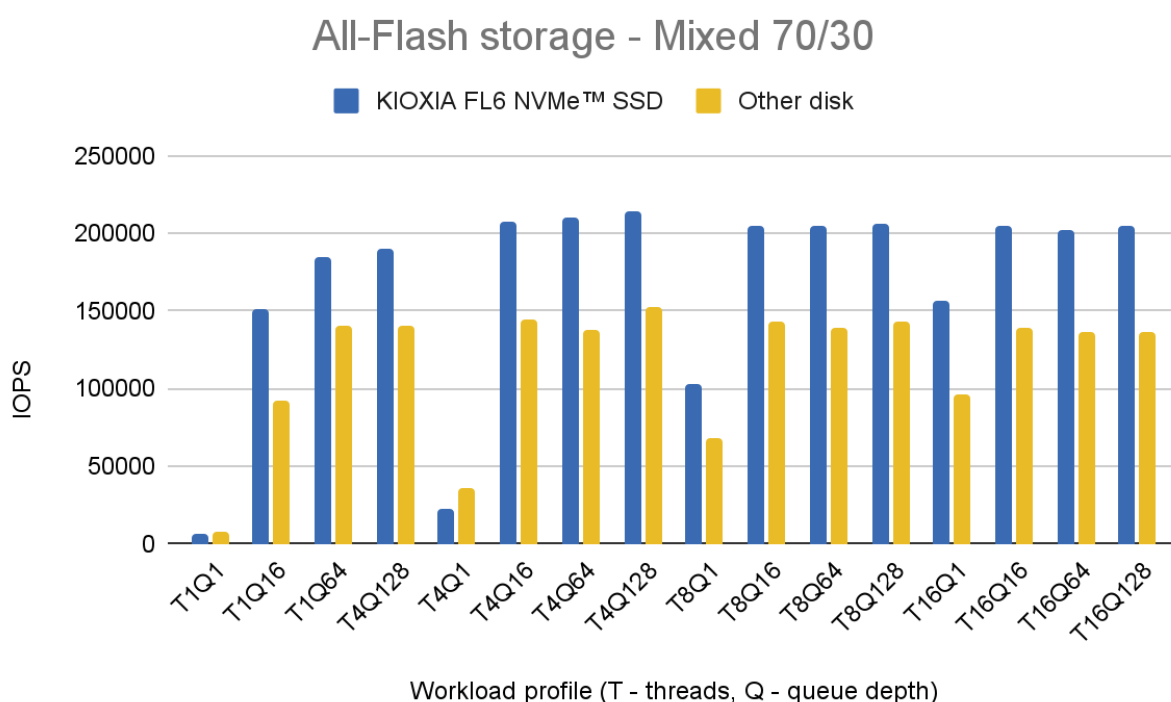


Fig. 14 Mixed 70/30 workload performance on the All-Flash storage, compared to a similar device.

5.3. Test conclusions

The KIOXIA FL6 NVMe SSDs were tested as data drives in the All-Flash storage appliance on a single node setup. The acquired performance results indicate that all the requirements for high-performing All-Flash storage have been met by providing high throughput for sequential workload and large IOPS number accompanied with low latency for

random access patterns. Additional comparison with similar NVMe drives shows the advantage of KIOXIA drives in all mentioned cases.

6. Write log device tests

6.1 Test description

In order to evaluate how the KIOXIA FL6 NVMe™ SSDs operate as a write log (ZFS SLOG) device in Open-E JovianDSS, two zpool configurations were assembled using HDD drives in a 2-way mirror or RAID-Z2 data groups. Measurements included the number of IOPS as well as the latency distribution under FIO-generated load for zpools with and without a write log in a single drive configuration.

The following characteristics qualify the tested device to be considered as a write log:

- relatively high declared write performance,
- sufficient write endurance,
- low latency.

Tests were performed on the following pool architectures: 12x 2-way mirror HDD and 6x RAID-Z2 HDD (4 disks per data group).

6.2. Zpool with 12x 2-way mirror HDD data groups

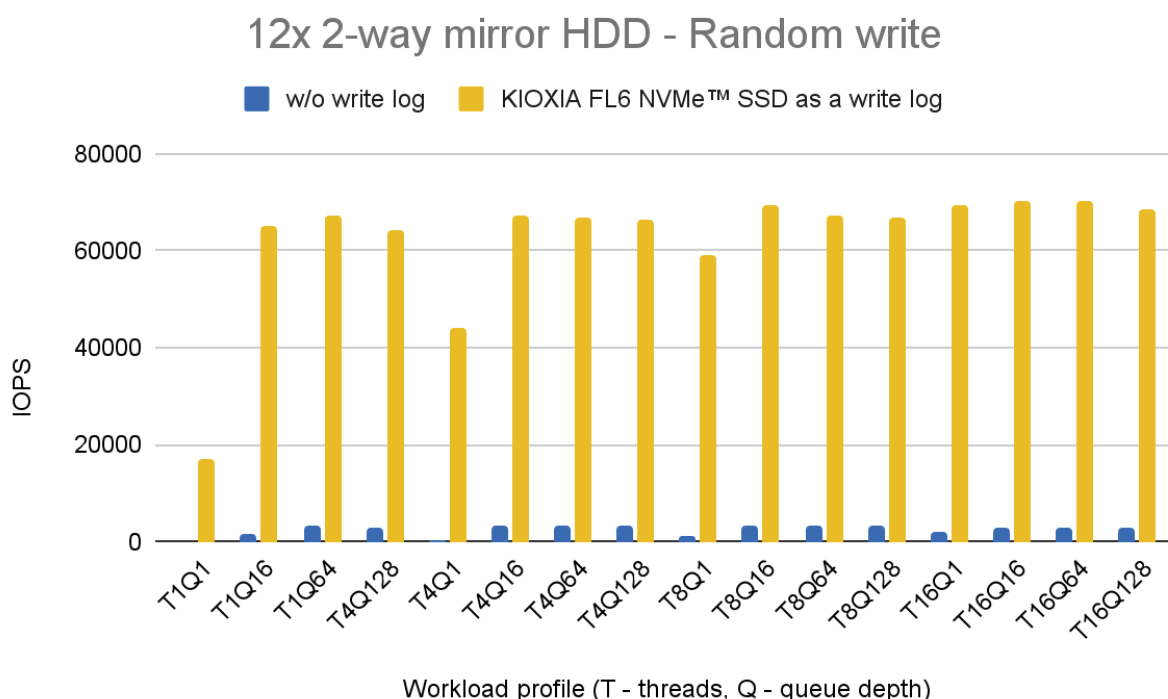


Fig. 15 Random write performance with 12x 2-way mirror HDD.

12x 2-way mirror HDD - Random write

Latency distribution for T1Q1

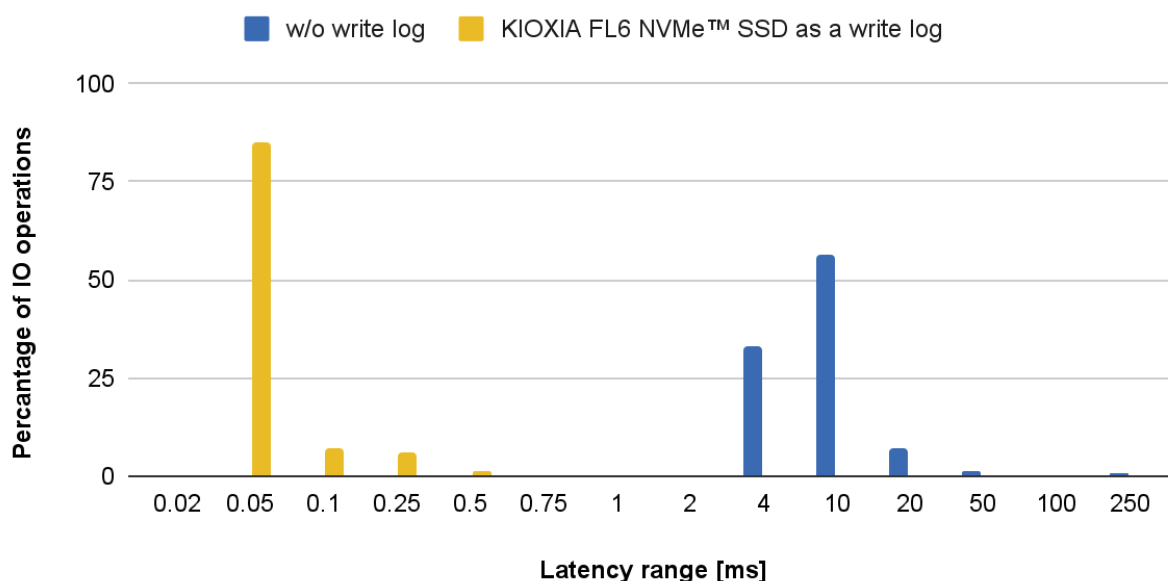


Fig. 16 Random write latency with 12x 2-way mirror HDD.

12x 2-way mirror HDD - Random write

Latency distribution for T16Q1

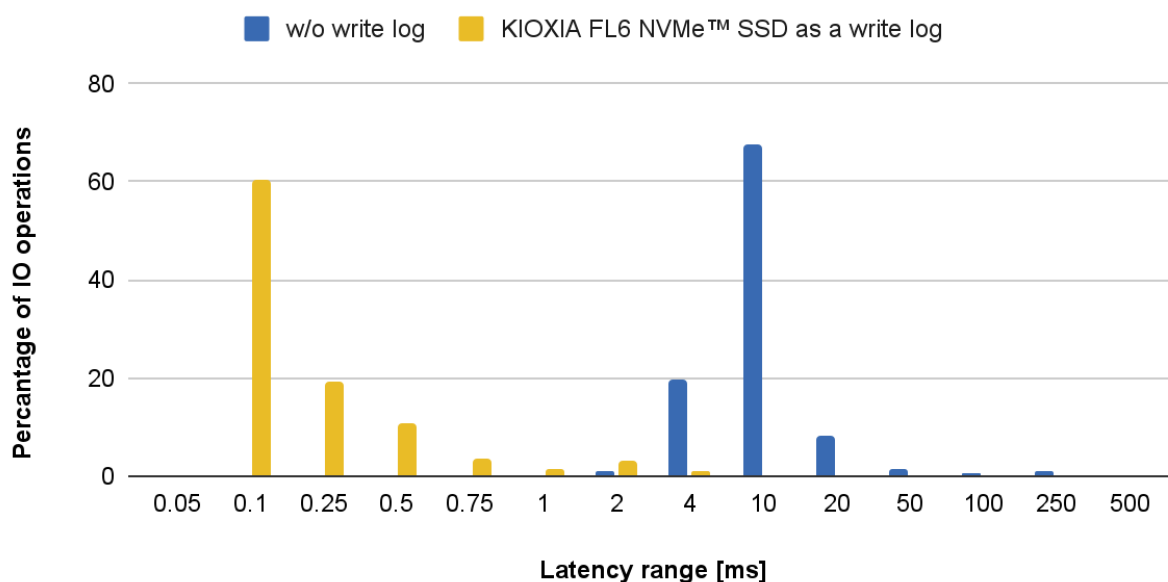


Fig. 17 Random write latency with 12x 2-way mirror HDD.

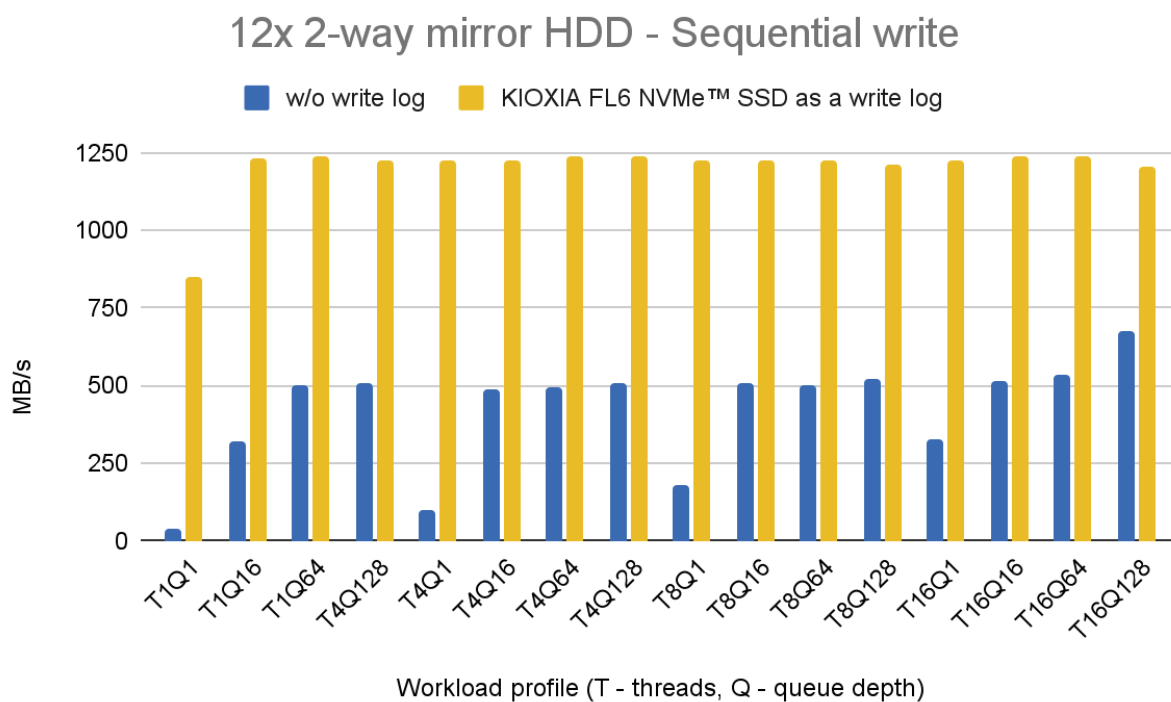


Fig. 18 Sequential write performance with 12x 2-way mirror HDD

6.3. Zpool with 6x RAID-Z2 data groups

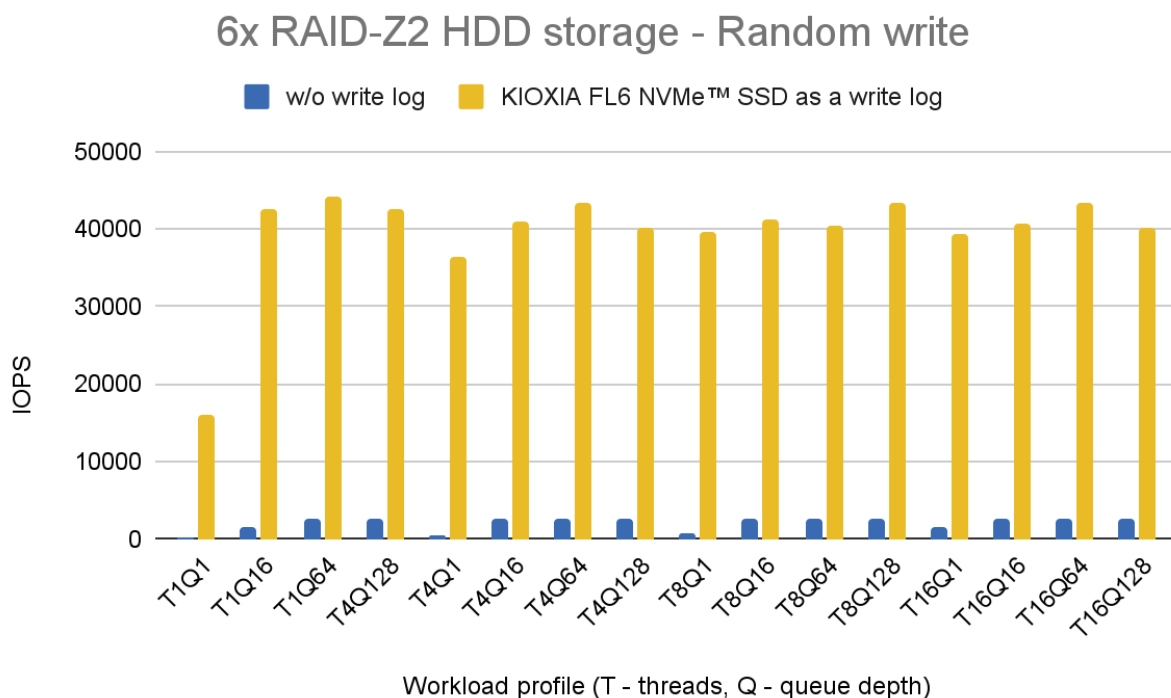


Fig. 19 Random write performance with 6x RAID-Z2 HDD storage.

6x RAID-Z2 HDD storage - Random write

Latency distribution for T1Q1

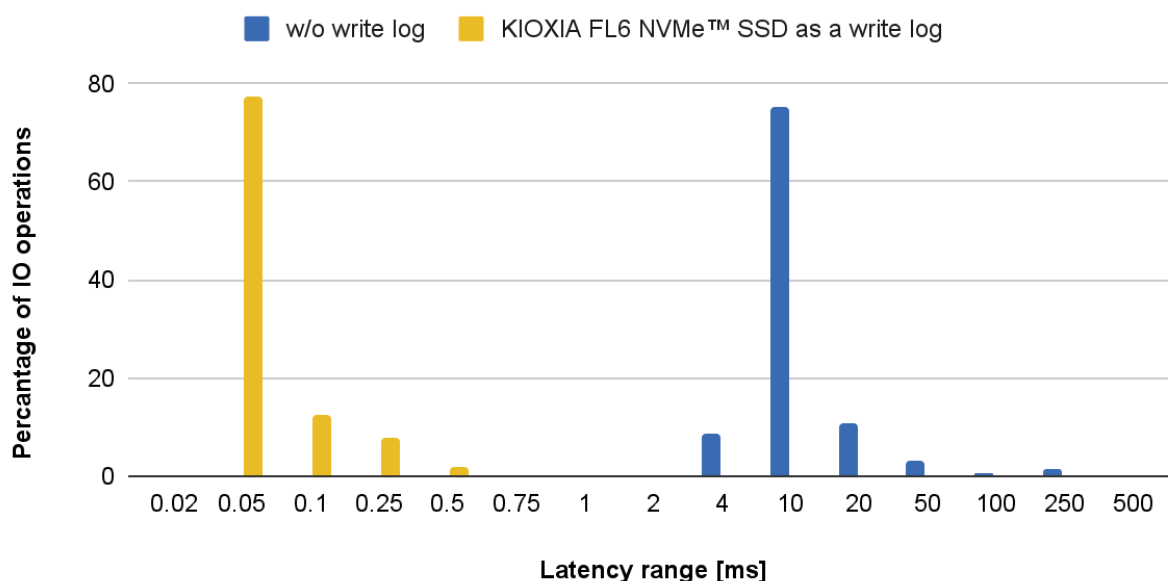


Fig. 20 Random write latency distribution for T1Q1 with 6x RAID-Z2 HDD storage.

6x RAID-Z2 HDD storage - Random write

Latency distribution for T16Q1

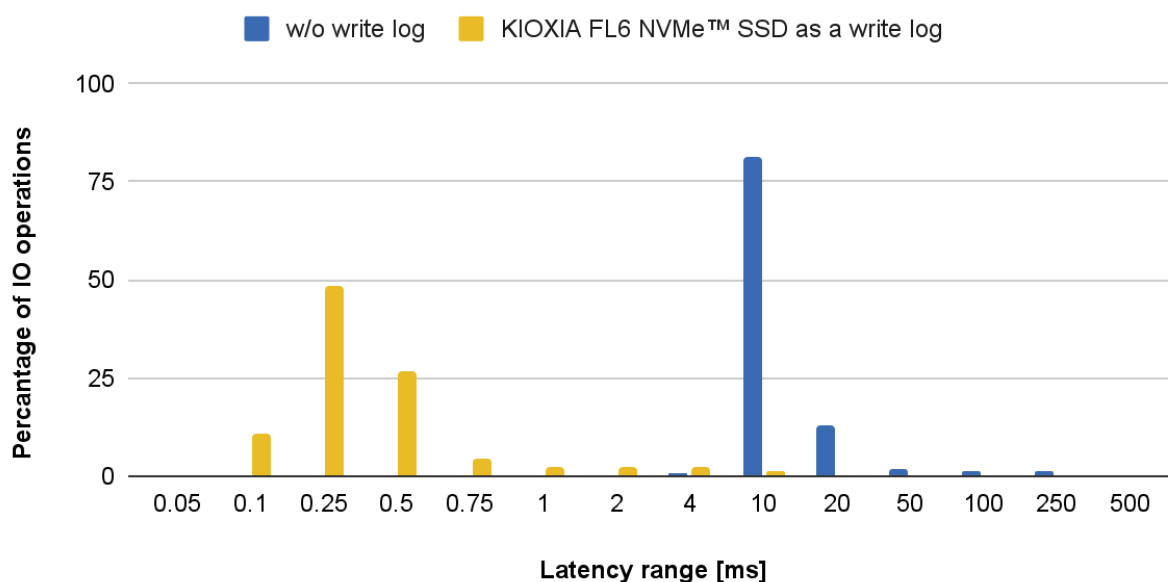


Fig. 21 Random write latency distribution for T16Q1 with 6x RAID-Z2 HDD storage.

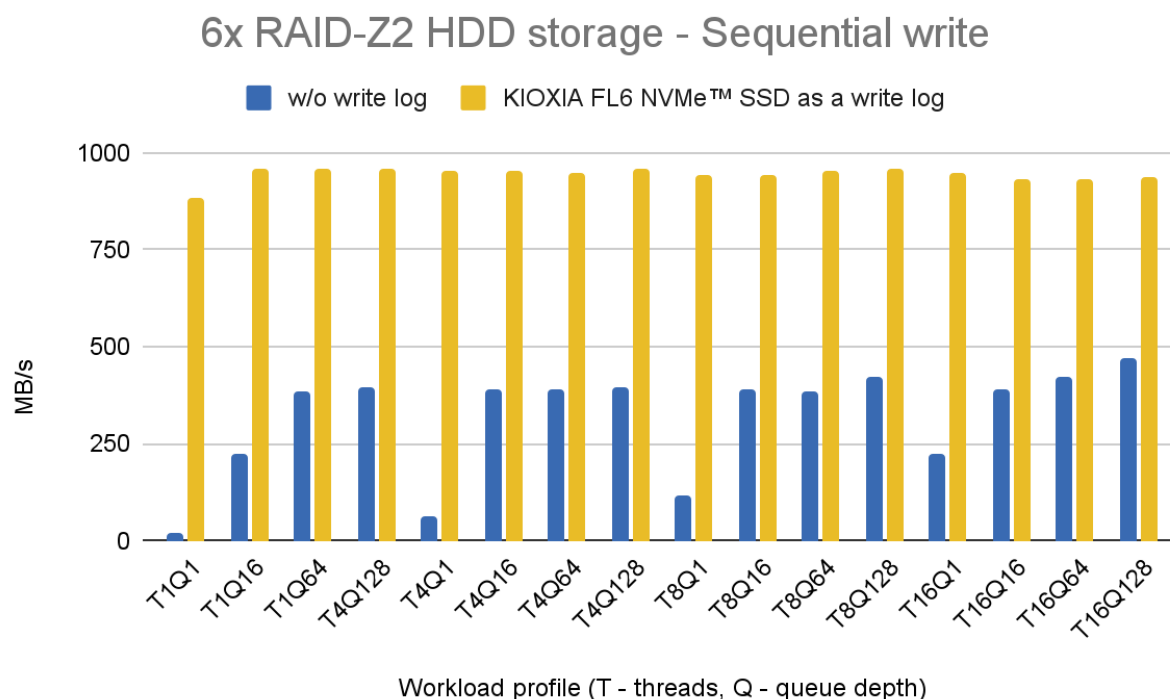


Fig. 22 Sequential write performance with 6x RAID-Z2 HDD storage.

6.4 Test conclusions

Random write IO results for both performance and latency improved significantly for zpools with an added write log over those without one. As for the sequential write test, results for the 12x 2-way mirror HDD zpool indicate a notable improvement both in throughput and latency. In the case involving throughput with a 6x RAID-Z2 HDD zpool a significant increase in performance was observed as well.

The tests have indicated that KIOXIA FL6 NVMe™ SSDs powered by Open-E JovianDSS data storage software outperform a similar NVMe SSD counterpart in terms of random reads, writes, and in sequential reads (while sequential writes achieved comparable results).

7. Read cache device test

7.1. Test description

To evaluate the KIOXIA FL6 NVMe™ SSDs serving as a read cache device for Open-E JovianDSS, a zpool consisting of 12 mirrored HDD drives was assembled, as well as a zpool with 6x RAID-Z2 HDD data groups. In order to lessen the influence of ZFS' ARC cache on the results, its size was artificially decreased to around 15 GB. This ensured that most of the cache hits would originate from the read cache device (ZFS L2ARC). Before performing actual measurements, fio was run until more than 99% of data was read from the read cache.

For better comparison, tests with standard ARC and no read cache were also carried out.

7.2. Zpool with 12x 2-way mirror HDD data groups

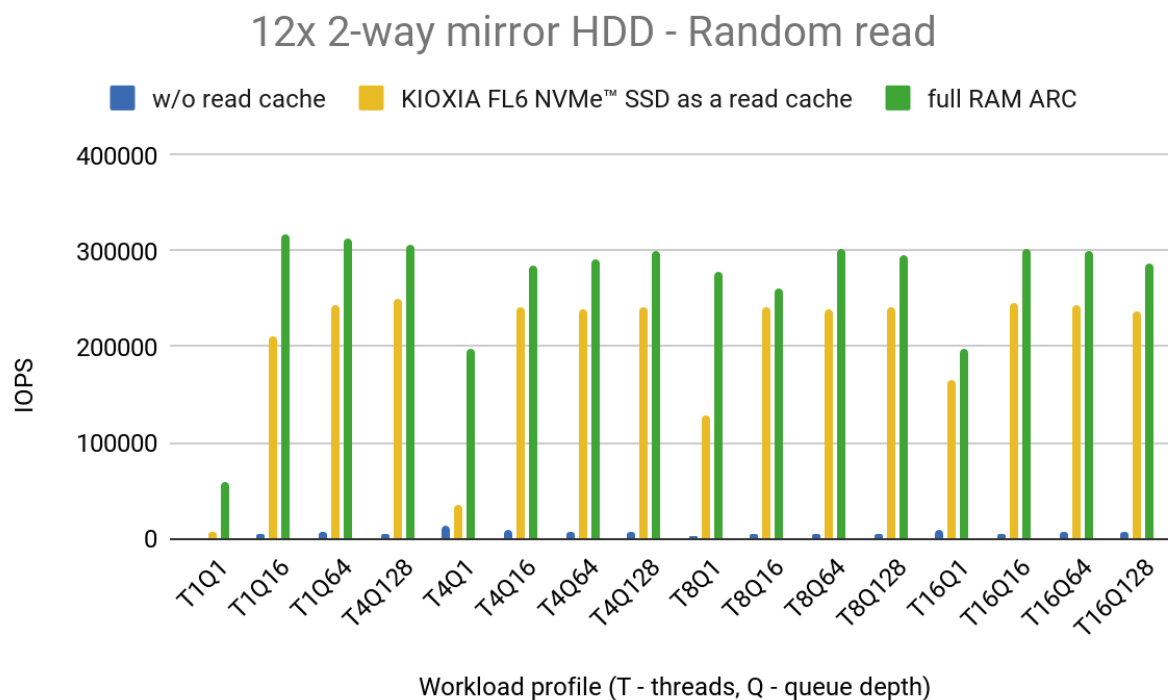


Fig. 23 Random read performance with 12x 2-way mirror HDD.

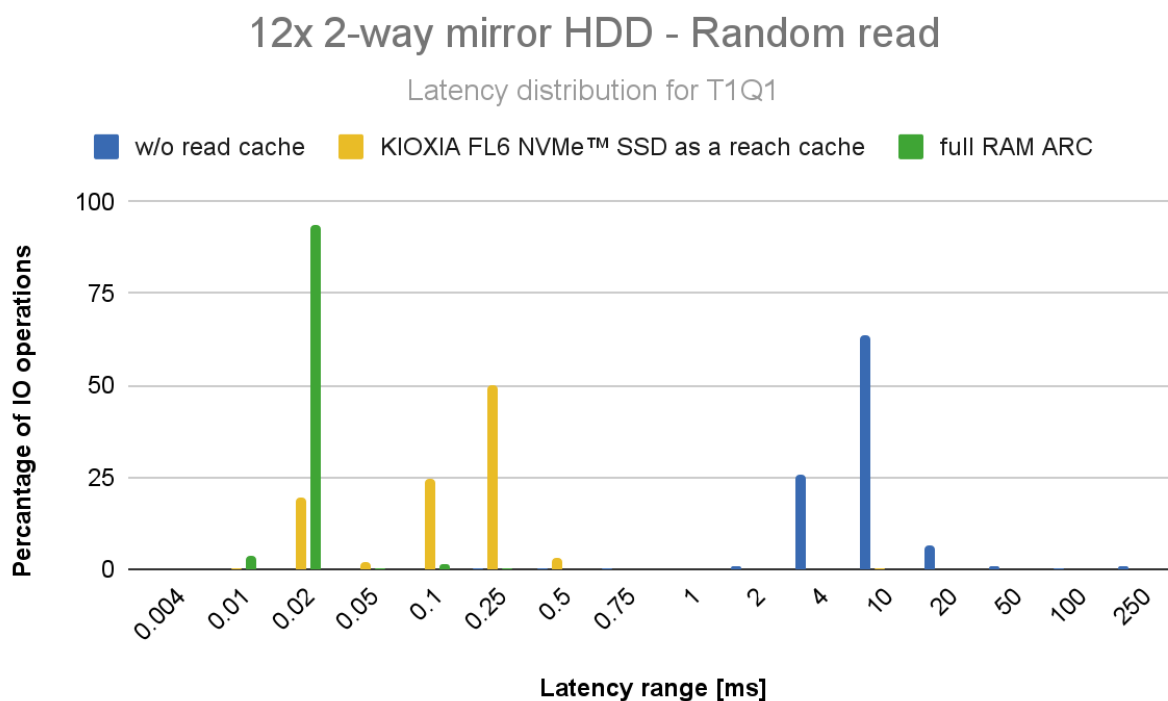


Fig. 24 Random read latency with 12x 2-way mirror HDD.

12x 2-way mirror HDD - Random read

Latency distribution for T16Q1

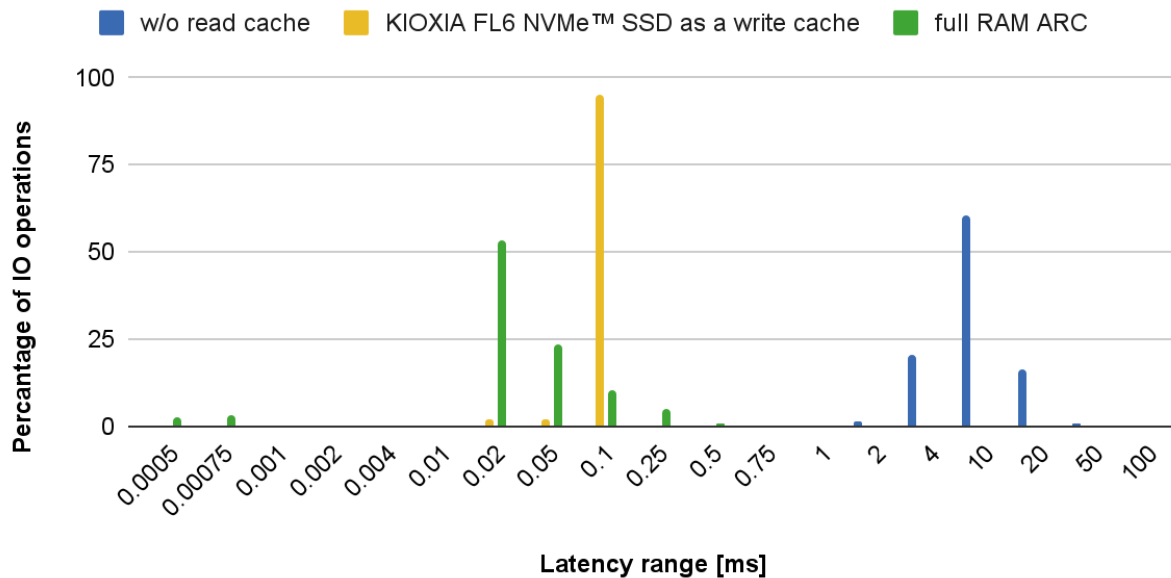


Fig. 25 Random read latency distribution for T1Q1 with 12x 2-way mirror HDD.

7.3. Zpool with 6x RAID-Z2 HDD data groups

6x RAID-Z2 HDD storage - Random read

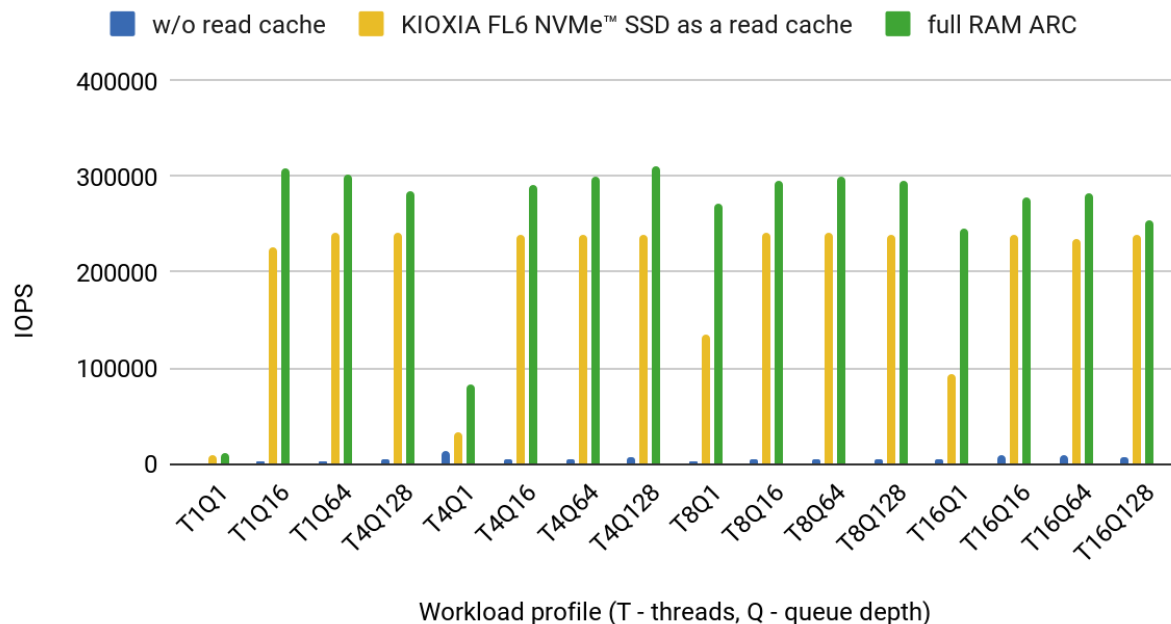


Fig. 26 Random read performance with 6x RAID-Z2 HDD storage.

6x RAID-Z2 HDD storage - Random read

Latency distribution for T1Q1

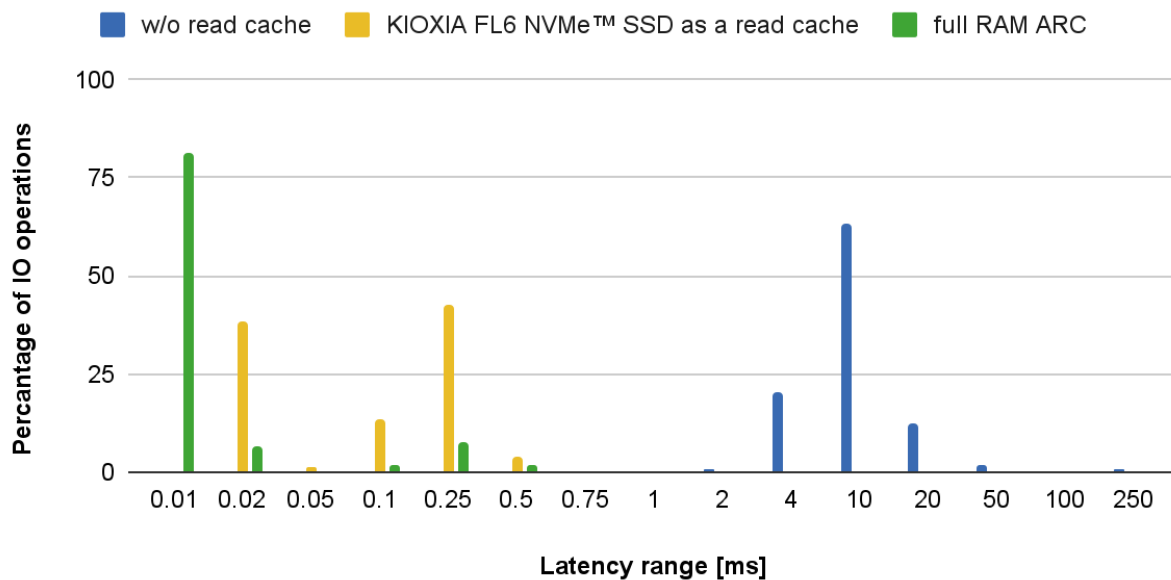


Fig. 27 Random read latency distribution for T1Q1 with 6x RAID-Z2 HDD storage.

6x RAID-Z2 HDD storage - Random read

Latency distribution for T16Q1

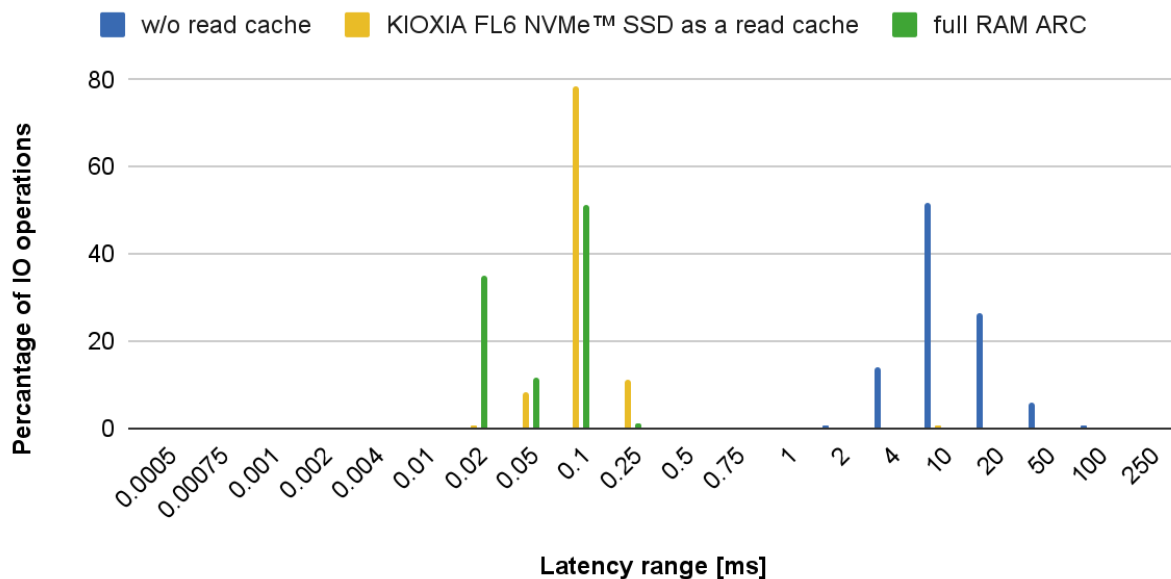


Fig. 28 Random read latency distribution for T16Q1 with 6x RAID-Z2 HDD storage.

7.4. Test conclusions

Cases with reduced ARC size (HDD zpools with and without read cache) and standard ARC size (for HDD-only zpools) were compared. Performance and latency test results turned out to be the best for the standard ARC size case (the majority of reads originating from RAM). Performance for read cache was worse than for RAM but on the other hand notably better than for regular HDDs.

Despite the fact that the performance with read cache is worse than that with RAM-based cache, it is not a significant difference. At lower costs, more cache space can be applied due to the fact that the cost of RAM is much higher than the cost of the KIOXIA FL6 NVMe™ SSDs per gigabyte.

8. HA non-shared storage cluster compatibility test

In order to ensure proper operation of KIOXIA FL6 NVMe™ SSDs in Open-E JovianDSS High Availability non-shared storage cluster environments, various compatibility tests were performed.

8.1. Functional tests

All essential and critical cluster mechanisms were examined for correct operation with the tested drives used as data drives, read cache, and write log devices. They are summarized in the table below.

Tab. 5 Results for the HA non-shared storage cluster compatibility test.

Tested functionality	Result
Manual Failover	passed
Automatic Failover triggered after network failure	passed
Automatic Failover triggered after system shutdown	passed
Automatic Failover triggered after system reboot	passed
Automatic Failover triggered after system power-off	passed
Automatic Failover triggered after I/O failure	passed

8.2. Test conclusions

Compatibility of the tested device with essential HA non-shared cluster operations was extensively checked. None of the test cases described in Table 5 showed any undesirable behavior, indicating full compatibility with Open-E JovianDSS cluster configurations.

9. Summary

The KIOXIA FL6 800GB NVMe™ SSDs were comprehensively tested for full functional compatibility with Open-E JovianDSS for both single node and HA non-shared storage cluster architectures. Performance measurements were also carried out in a single node configuration with various storage arrangements. The tests were designed to find any abnormalities in the devices, regardless of whether they were used as cache or as data drives for Open-E JovianDSS.

The tests have indicated that KIOXIA FL6 NVMe™ SSDs powered by Open-E JovianDSS data storage software outperform a similar NVMe SSD counterpart in terms of both random reads and writes, as well as in sequential reads (while sequential writes achieved comparable results). Also, the KIOXIA FL6 NVMe™ SSD-based read cache is

especially recommended for setups with expected massive read loads due to high performance, high endurance, low latency, and effective support for heavy workloads. As a second-level read cache, the KIOXIA drives delivered a surge in performance (250K IOPS), in comparison to the setup with the read operations executed directly from HDD data disks (the average result for this scenario is about 1.3K IOPS). While using KIOXIA FL6 NVMe™ SSD as a write log, the storage pool delivered a significantly improved performance with higher, consistent IOPS (about 18 times higher for random write operations, and 2,5 times higher for sequential writes operations), than without write log together with lower latency

All test results proved that KIOXIA FL6 NVMe™ SSDs can be considered a reasonable and competitive choice. It is recommended to use Open-E JovianDSS with the KIOXIA FL6 NVMe™ SSDs as data drives for All-Flash storage solutions. Given the high performance and low latencies achieved in testing, the examined device can safely be added to the Open-E Hardware Certification List and granted “Certified by Open-E” status.