Open-E JovianDSS

# Intel® Optane™ SSD P1600X 58 GB / 118 GB


**Certification report**

Release date: 2022.05.23

# Table of Contents

# 1. Introduction

The following certification report aims to present the results of various compatibility and performance tests performed on the Intel Optane P1600X NVMe Solid State Drive when used in combination with Open-E JovianDSS. A description of the testing methods used during those procedures will also be provided. The results, obtained from the aforementioned testing, along with technical specifications were used as the basis for recommendations on specific roles in which to use the device under test (DUT) in Open-E JovianDSS-based systems.
The following applications were considered during the certification process:

- boot medium device for Open-E JovianDSS,
- writelog device (ZFS SLOG device).

Not considered was usage as a read cache device or data storage medium on account of storage capacity being too small.
For all cases mentioned above, functional testing was performed for both Single node and High Availability non-shared storage cluster configurations. In addition, performance tests were also conducted for the Single node configuration variant. Detailed descriptions for each of the aforementioned are included in the appropriate report chapters.

# 2. Device Under Test description

The following table includes the Intel Optane P1600X NVMe drive hardware specifications.

Table 1. Intel Optane P1600X hardware specifications.

| | |
|---|---|
| **Product name** | Intel Optane SSD P1600X |
| **Model name** | SSDPEK1A058GA<br>SSDPEK1A118GA |
| **Storage capacity** | 58 GB<br>118 GB |
| **Form factor** | M.2 22 x 80mm |
| **Interface** | PCIe 3.0 x4, NVMe |
| **Technology** | 3D XPoint |
| **Enhanced Power Loss Data Protection** | Yes |
| **Security** | No encryption |

# 3. Test environment description

Hardware specifications for environments used during certification testing are included in the following tables. The configuration described in Table 2a was used for every Single node test. Table 2b shows the configuration used when testing High Availability non-shared storage cluster nodes. More information can be found on HA cluster compatibility testing in chapter 7.

Table 2a. Hardware specifications for single node tests.

| | |
|---|---|
| **System name** | Supermicro SYS-620U-TNR |
| **Motherboard** | Supermicro X12DPU-6 |
| **CPU** | 2x Intel Xeon Gold 6330 |
| **RAM** | 64 GB - 4x SK Hynix HMA82GR7DJR8N-XN 3200 MHz DDR 4 ECC 16 GB |
| **Storage controller** | Broadcom / LSI 9300 SAS HBA (used for HDD) |
| **Storage devices** | • 2x Intel Optane P1600X SSD 118 GB<br>• 2x Intel Optane P1600X SSD 58 GB<br>• 20x Toshiba AL13SEB900 SAS HDD 10K 900 GB<br>• 4x Intel D5-P4320 SSD 2 TB |
| **System** | Open-E JovianDSS up29r1 b44475 |

Table 2b. Hardware specifications for HA non-shared storage cluster tests.

| | |
|---|---|
| **System name** | Tarox ParX R2242i G6 Server |
| **Motherboard** | Intel S2600WFT |
| **CPU** | 2x Intel Xeon Gold 5222 |
| **RAM** | 192 GB - 16x Micron MTA18ASF2G72PDZ-3G2E1 3200 MHz DDR416GB DDR4 ECC 16 GB |
| **Storage devices** | • 2x Intel Optane P1600X SSD 118 GB<br>• Intel Optane P1600X SSD 58 GB<br>• 4x Intel D5-P4320 SSD 2 TB |
| **System** | Open-E JovianDSS up29r1 b44475 |

Tool used for performance benchmarking: Fio for Linux, v3.28.

# 4. Raw device tests

In order to be able to properly interpret the results acquired from the tests of the Intel Optane P1600X drive while it was used as a data storage device for Open-E JovianDSS, it's necessary to measure the raw performance of the device. 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 allowing for effective storage performance tests to be done at the ZFS level. The following measurements were carried out on a P1600X drive model with 118 GB storage capacity.

## 4.1. Test description

Included test cases are described in Table 3. 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. For additional information, results of the same tests performed on the Intel S4510 M.2 480 GB SATA disk are included as the disk usually recommended for the Open-E JovianDSS boot medium.

Table 3. 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 |

## 4.2. Performance results

Figures 1 to 4 present the number of IOPS acquired in every test case.

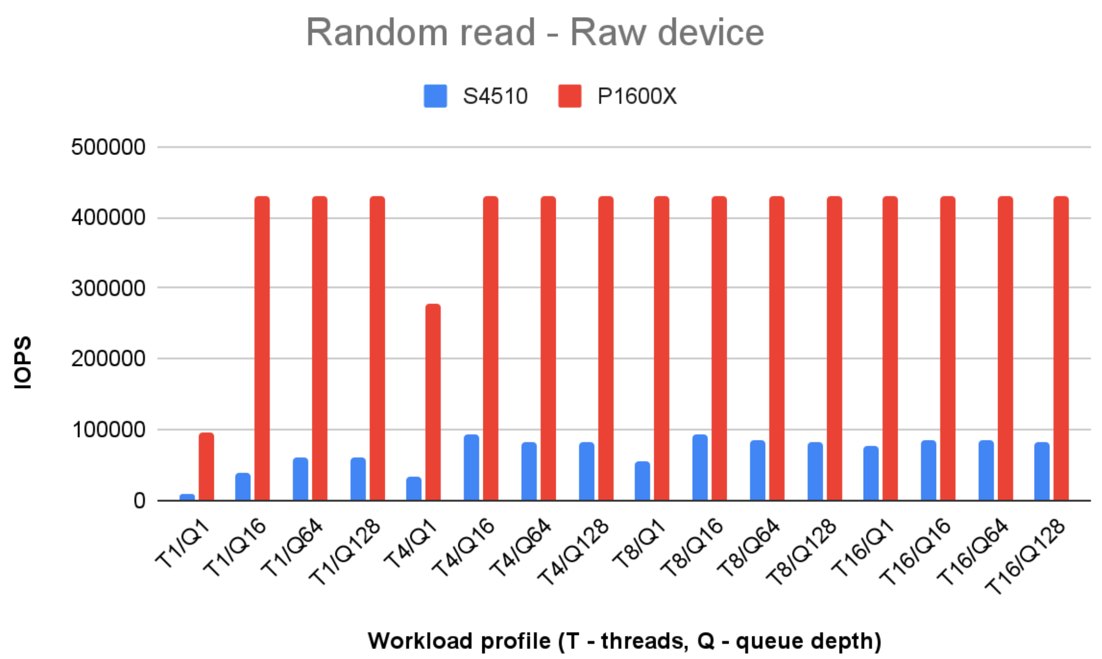


Fig. 1. Random read performance on raw device compared to S4510.

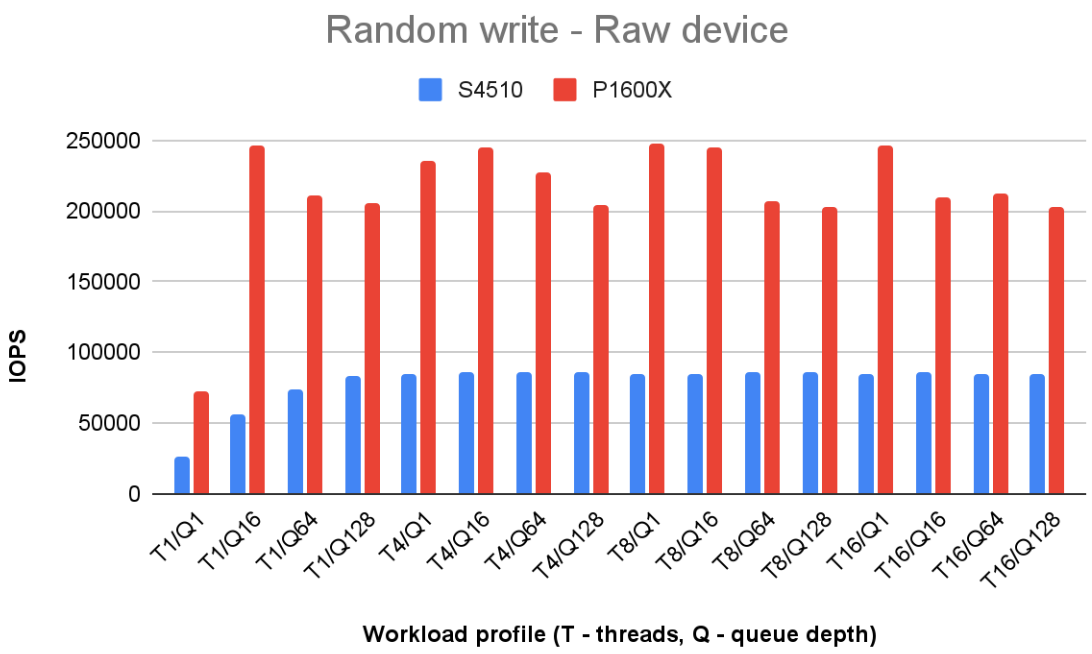


Fig. 2. Random write performance on raw devices compared to S4510.

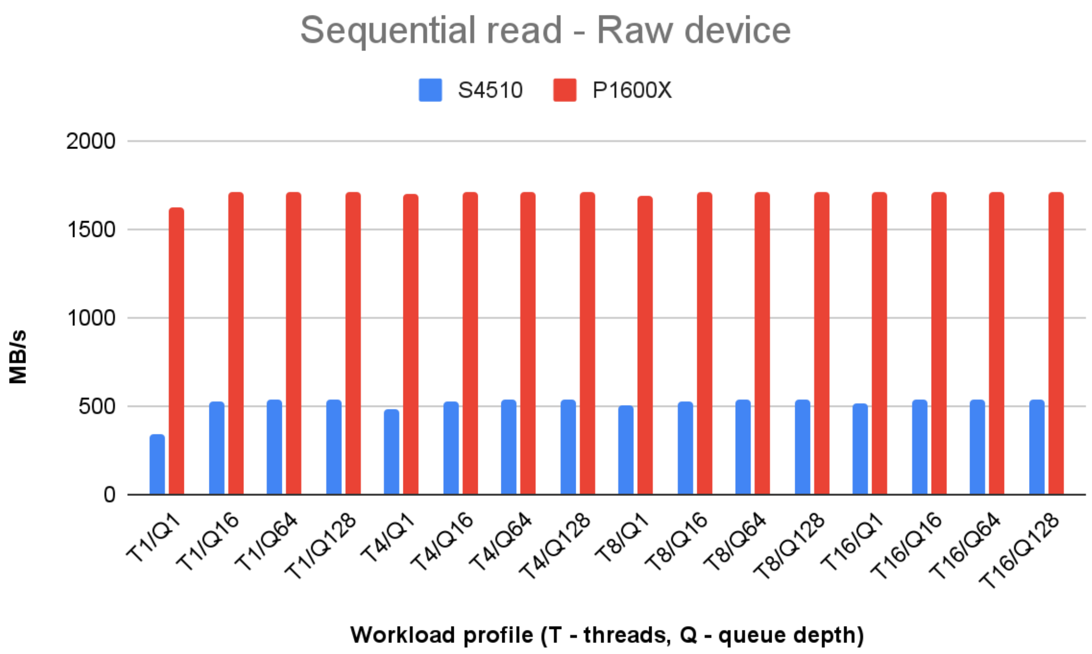


Fig. 3. Sequential read performance on raw device compared to S4510.

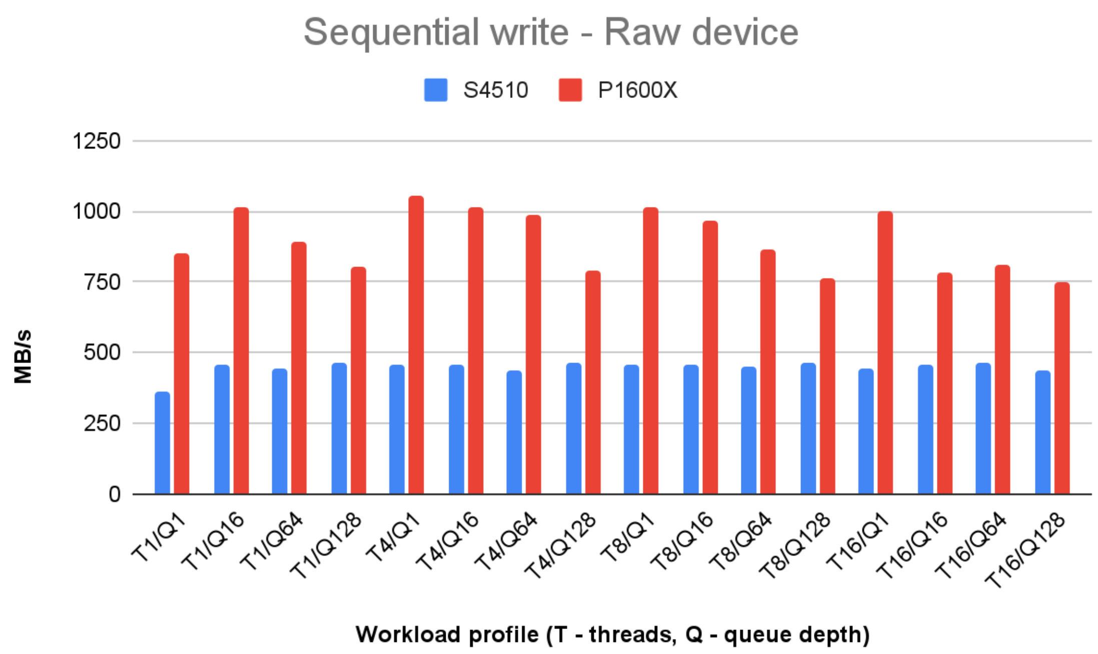


Fig. 4. Sequential write performance on raw device compared to S4510.

### 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 a level similar to those declared by the manufacturer. Compared to the S4510 SSD drive, the DUT outperformed it in both random and sequential IO pattern tests, which is directly caused by it having a different storage architecture and data transmission interface.

## 5. Boot medium test

Tests of the Intel Optane P1600X drive, when used as a boot medium for Open-E JovianDSS, included both functional and performance aspects. A device with 58 GB capacity was chosen for further evaluation. It was considered for this role because of the following characteristics:

- form factor: M.2 NVMe interface popular for boot drive applications,
- relatively small storage space - adequate for system partitions but not as a storage drive,
- suitable write endurance,
- exceptional declared performance results.

### 5.1. Functional tests

Both the single as well as mirrored configurations of the boot medium were tested for full compatibility. Installation was carried out using a regular, built-in installer of Open-E JovianDSS. Table 4 presents all of the functional aspects that were taken into consideration for the tested device. The outcome for each functional aspect is also listed.

Table 4. Functional test results of the Intel P1600X 58 GB when used as a boot medium in single and mirrored disk configurations.

| Functional aspect | Result |
|---|---|
| Boot medium storage space | passed |
| Open-E JovianDSS installation process | passed |
| System stability | passed |
| Updating the system | passed |
| Single boot medium failure simulation in mirrored drive configuration | passed |

## 5.2. Performance tests

In addition to functional tests, performance measurements for the single drive boot medium were also carried out on Open-E JovianDSS' local partition using fio. All test configurations are presented in Table 5. Every combination of thread (1, 4, 8, 16) and queue depth number (1, 16, 64, 128) was used for the generated load. Results of those tests are presented in figures 5-9. The performance of the P1600X was also compared to the previously mentioned S4510 drive.

Table 5. Test cases description for boot medium performance 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 | random | 70/30 | 4 kB |

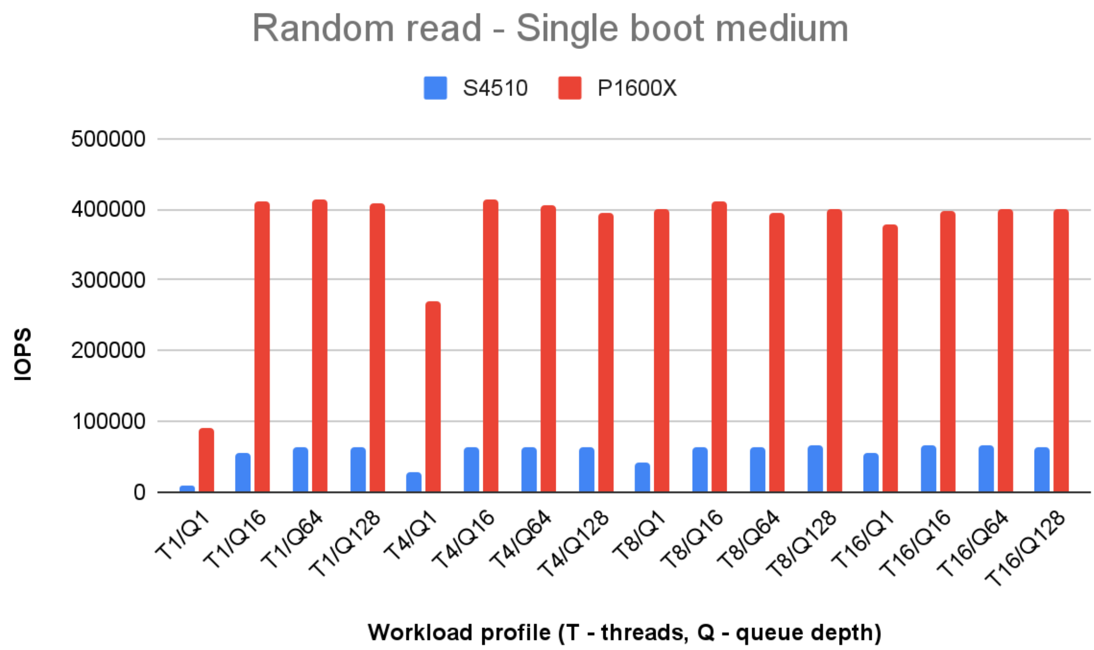


Fig. 5. Random read performance on a single boot medium, compared to the S4510.

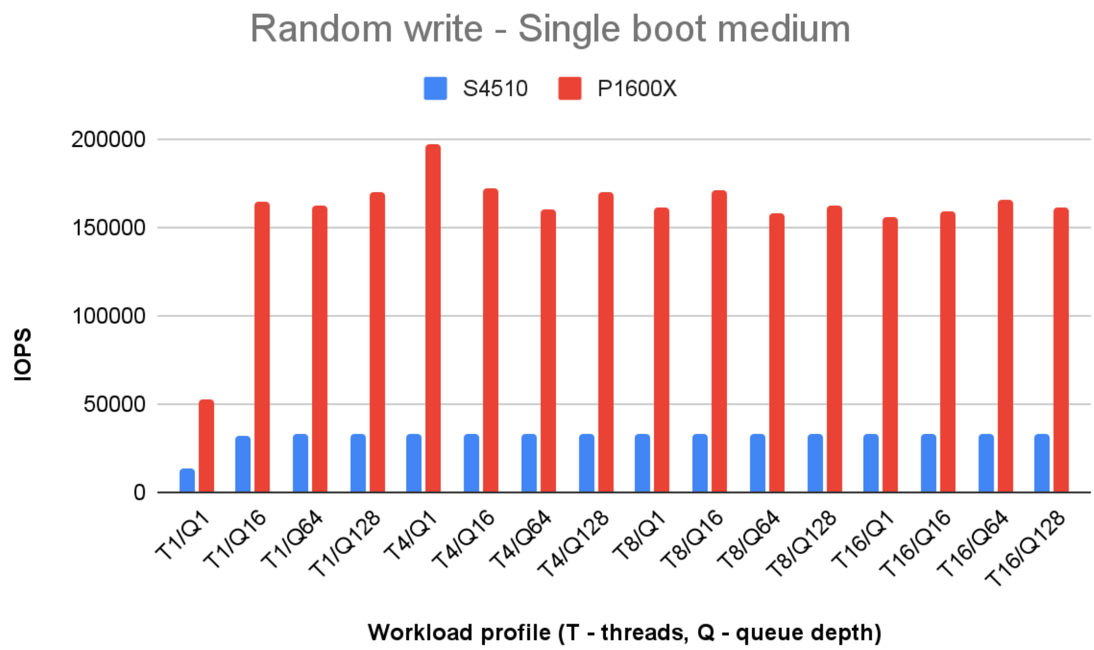


Fig. 6. Random write performance on a single boot medium, compared to the S4510.

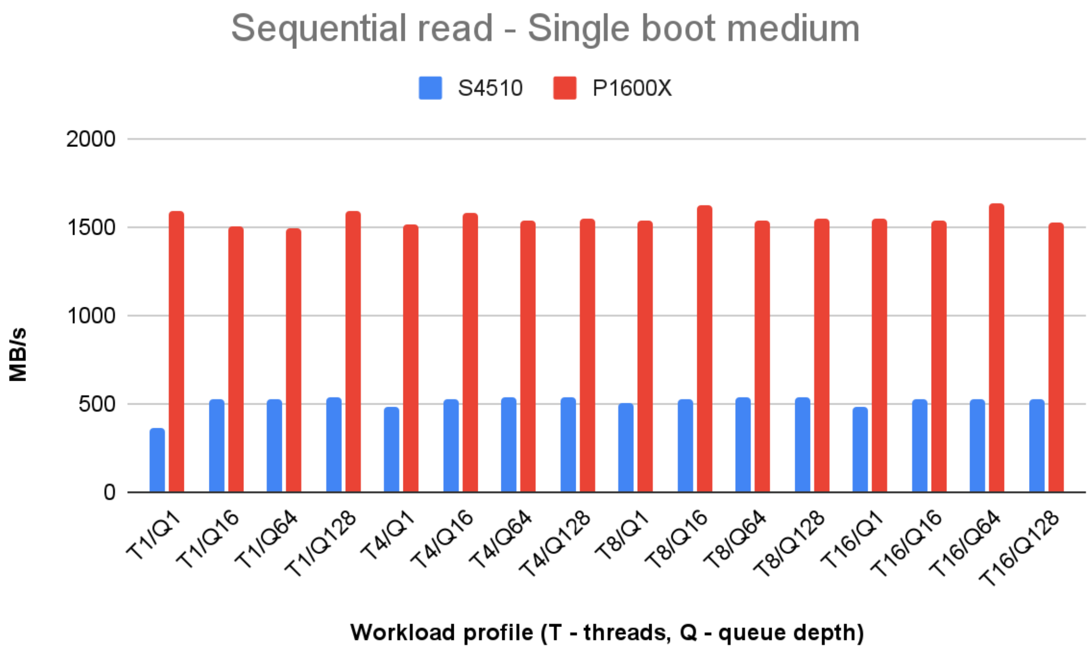


Fig. 7. Sequential read performance on a single boot medium, compared to the S4510.

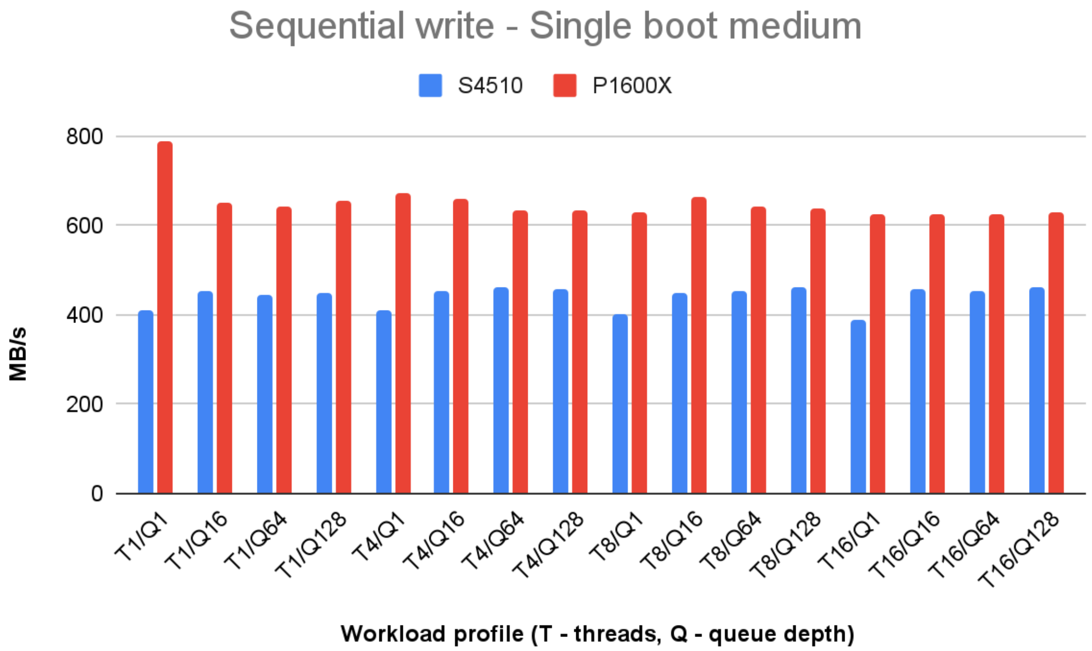


Fig. 8. Sequential write performance on a single boot medium, compared to the S4510.

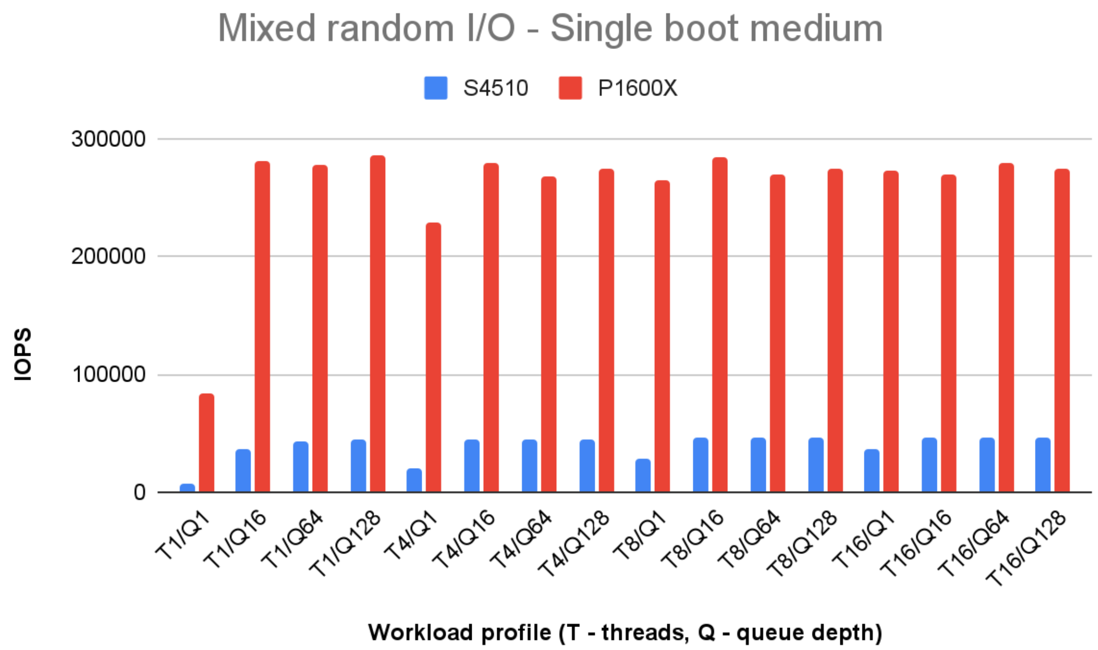


Fig. 9. Mixed random read and write performance on a single boot medium, compared to the S4510.

### 5.3. Test conclusions

In order to check boot medium functional compatibility with Open-E JovianDSS, several tests were conducted. The results are summarized in Table 4. No faulty behaviour was found during any test. Following the functional tests, a performance analysis was executed directly on Open-E JovianDSS' partition. The tested device showed much higher performance compared to the S4510 drive in all test cases, with no issues found when using the device as an Open-E JovianDSS boot medium in either single or mirrored drive configurations.

## 6. Writelog device tests

In order to evaluate how the Intel Optane P1600X drive operated as a writelog for Open-E JovianDSS (ZFS SLOG device), several zpool configurations were assembled using HDD or SSD drives. Measurements included the number of IOPS as well as the latency distribution under a fio generated load for zpools with and without a writelog device, which was then added for each zpool in a single drive configuration. A model with a capacity of 118 GB was chosen for this.

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

- relatively high declared write performance,
- decent write endurance (suitable for workloads with data writes smaller than 600 GB/day over warranty period),
- low latency.

### 6.1. Zpool with 10x 2-way mirror HDD data groups

A zpool consisting of 20 Toshiba AL13SEB900 SAS HDD 10K disks, with a capacity of 900 GB, was assembled in a 2-way mirror configuration for each data group created. The test parameters are shown in Table 6. Figure 10 presents the IOPS results of a random write test case, with the latency distribution for T1/Q1 and T16/Q1 combinations being presented in figures 11 and 12 respectively. Figure 13 shows results for sequential write tests with the latency presented in Fig. 14 and 15.

Table 6. 10x 2-way mirror HDD zpool test parameters.

| Test case | IO pattern | Read to write % | Zvol block size | Data block size |
|---|---|---|---|---|
| Random write | random | 0/100 | 4 kB | 4 kB |
| Sequential write | sequential | 0/100 | 64 kB | 1 MB |

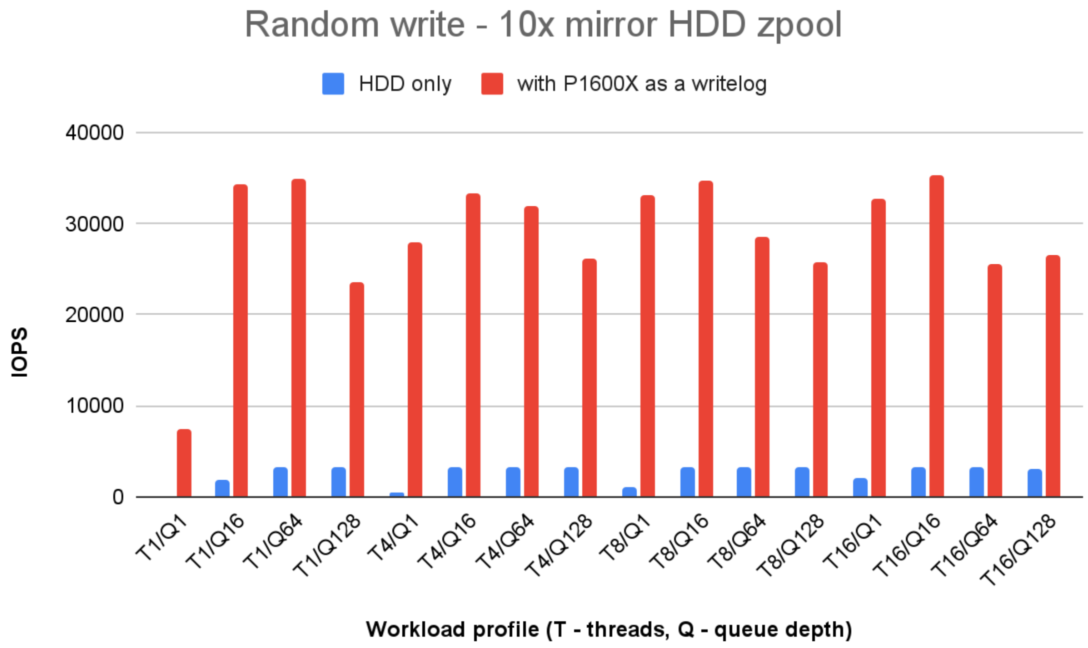


Fig. 10. Random write performance of a 10x mirror HDD zpool, with and w/o writelog.

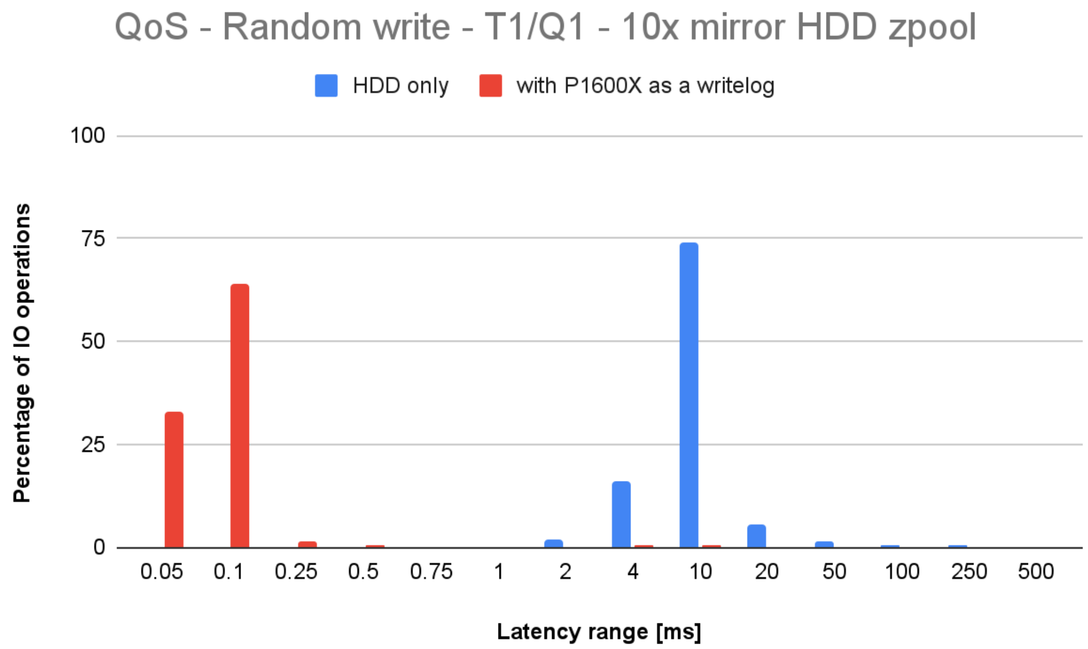


Fig. 11. Latency distribution for a random write test, done on a 10x mirror HDD zpool (T1/Q1), with and without writelog.

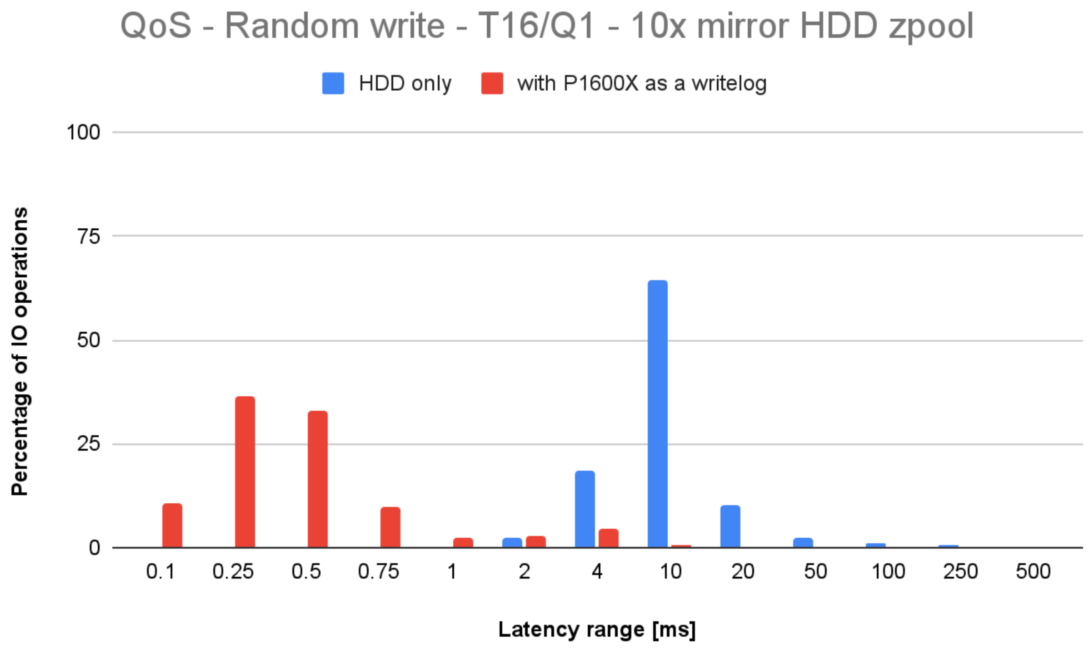


Fig. 12. Latency distribution for a random write test, done on a 10x mirror HDD zpool (T16/Q1), with and without writelog.

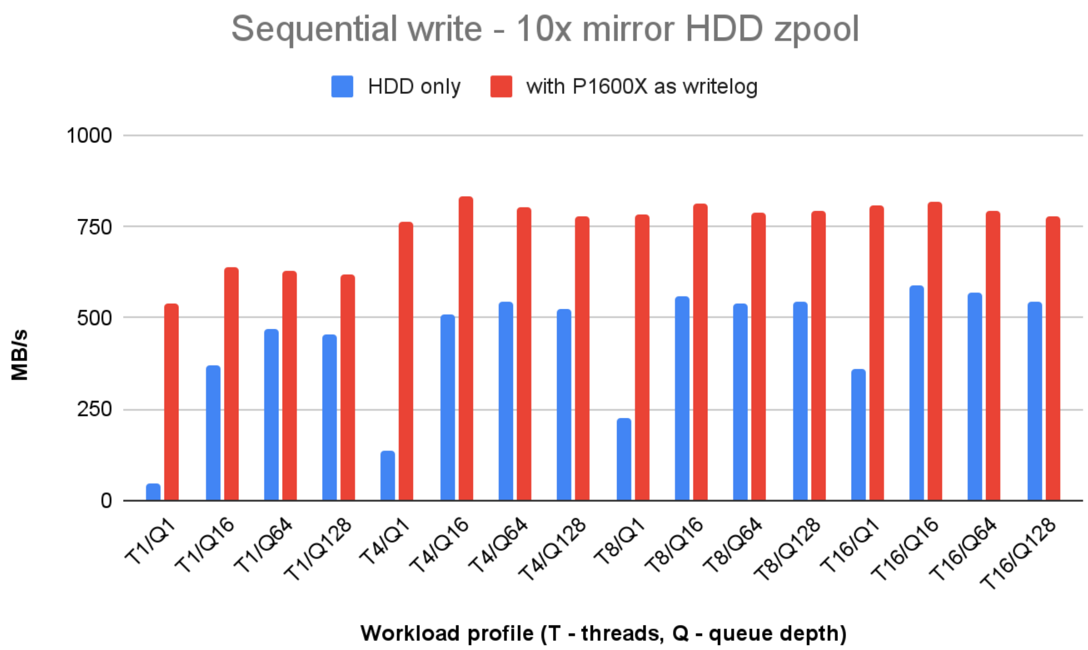


Fig. 13. Sequential write performance of a 10x mirror HDD zpool, with and w/o writelog.


Fig. 14. Latency distribution for a sequential write test, done on a 10x mirror HDD zpool (T1/Q1), with and without writelog.

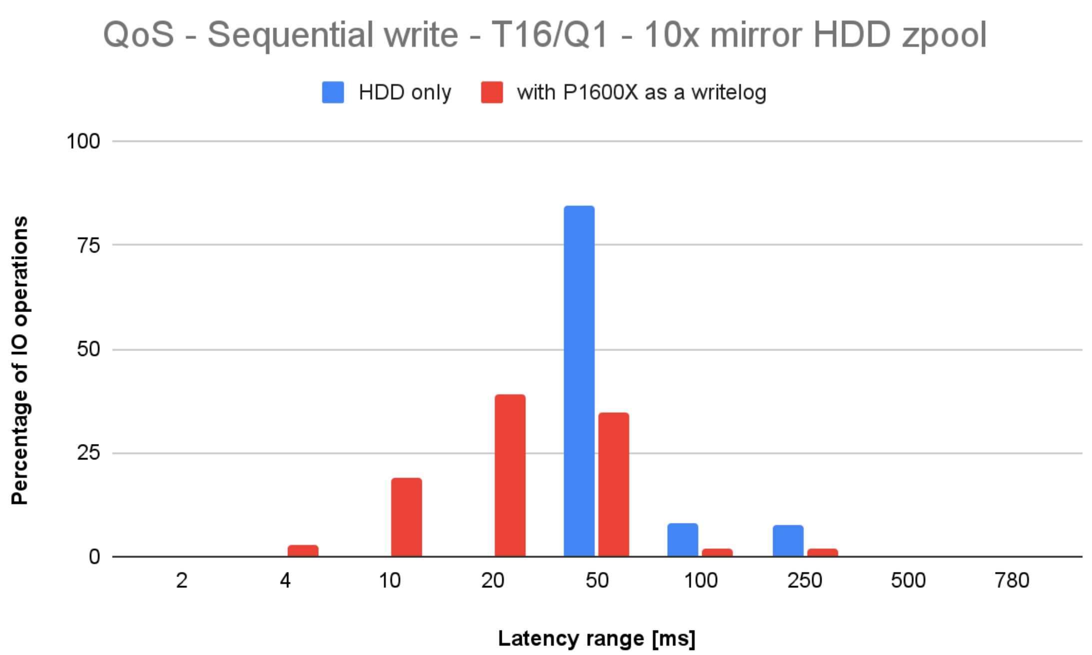


Fig. 15. Latency distribution for a sequential write test, done on a 10x mirror HDD zpool (T16/Q1), with and without writelog.

## 6.2. Zpool with 20x single HDD data groups

In this test, the zpool consisted of 20 single-disk data groups assembled from Toshiba AL13SEB900 SAS HDD 10K drives. Test parameters are shown in Table 6. Figure 16 presents the results of the random write test case with latency distributions for the T1/Q1 and T16/Q1 combinations presented in figures 17 and 18 respectively. Figure 19 shows the results for the sequential write test with the latency results being presented in Fig. 20 and 21.

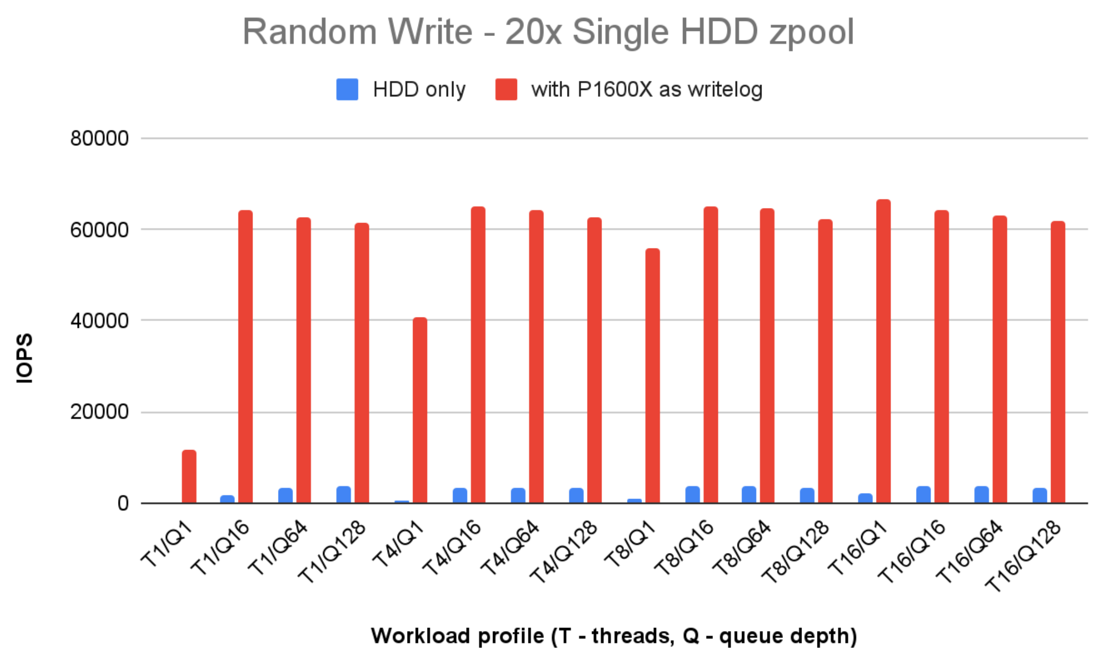


Fig. 16. Random write performance of a 20x single HDD zpool, with and w/o writelog.


Fig. 17. Latency distribution for a random write test, done on a 20x single HDD zpool (T1/Q1), with and without writelog.

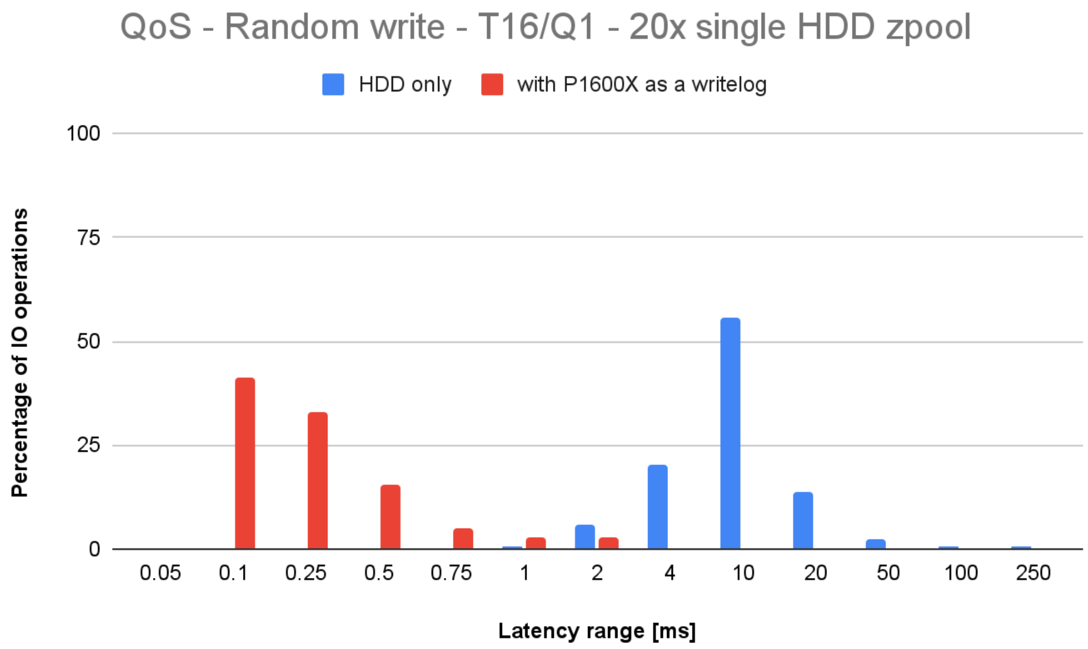


Fig. 18. Latency distribution for a random write test, done on a 20x single HDD zpool (T16/Q1), with and without writelog.

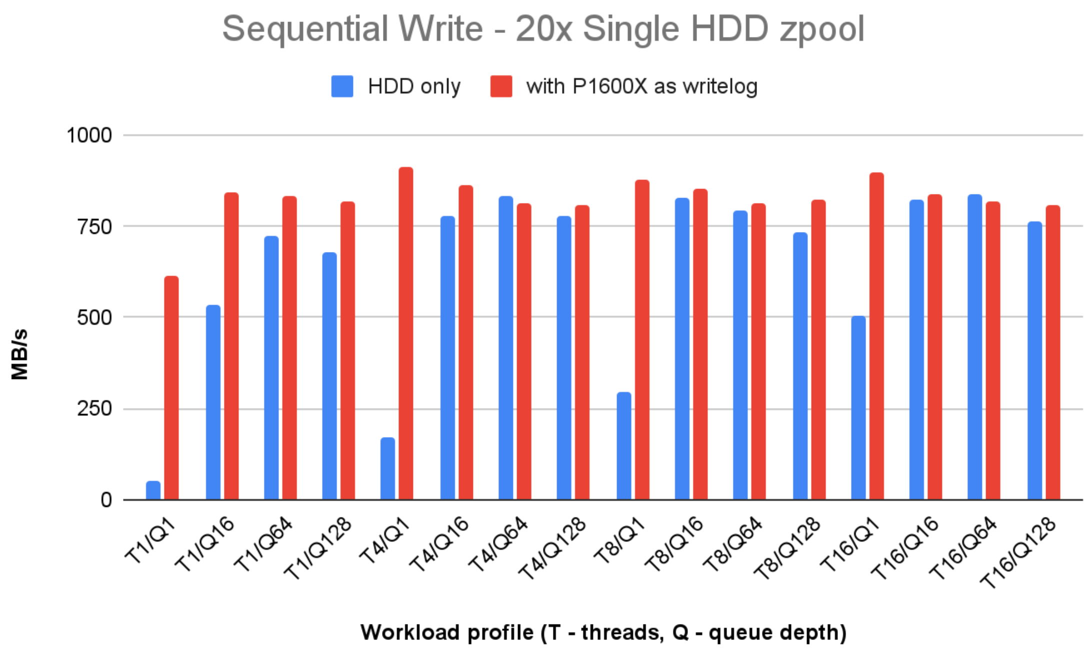


Fig. 19. Sequential write performance of a 20x single HDD zpool, with and w/o writelog.

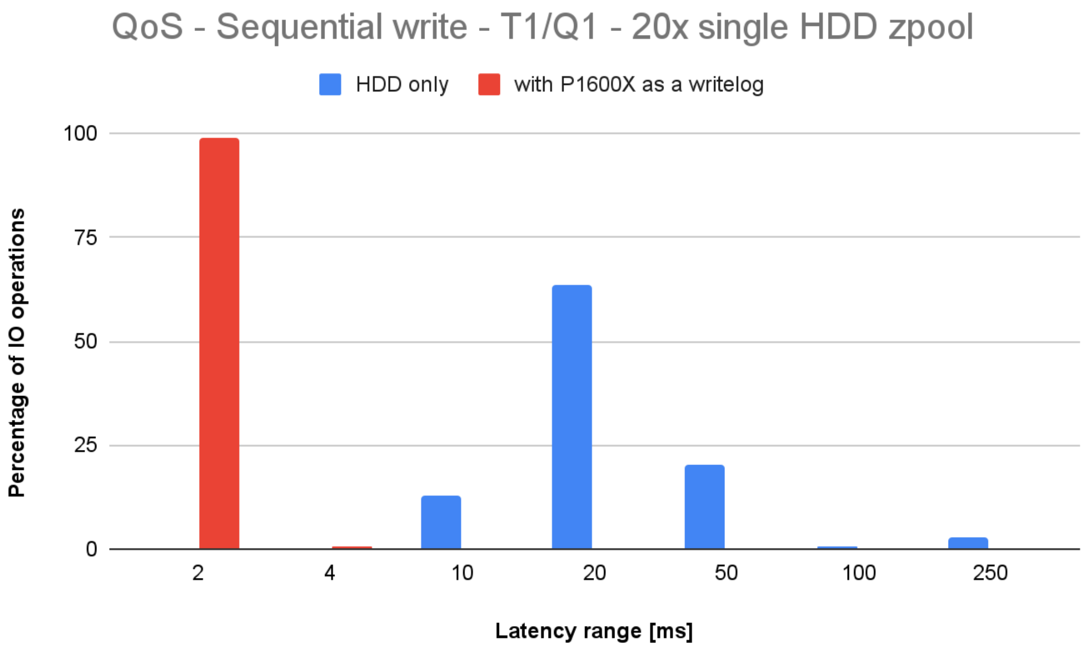


Fig. 20. Latency distribution for a sequential write test, done on a 20x single HDD zpool (T1/Q1), with and without writelog.

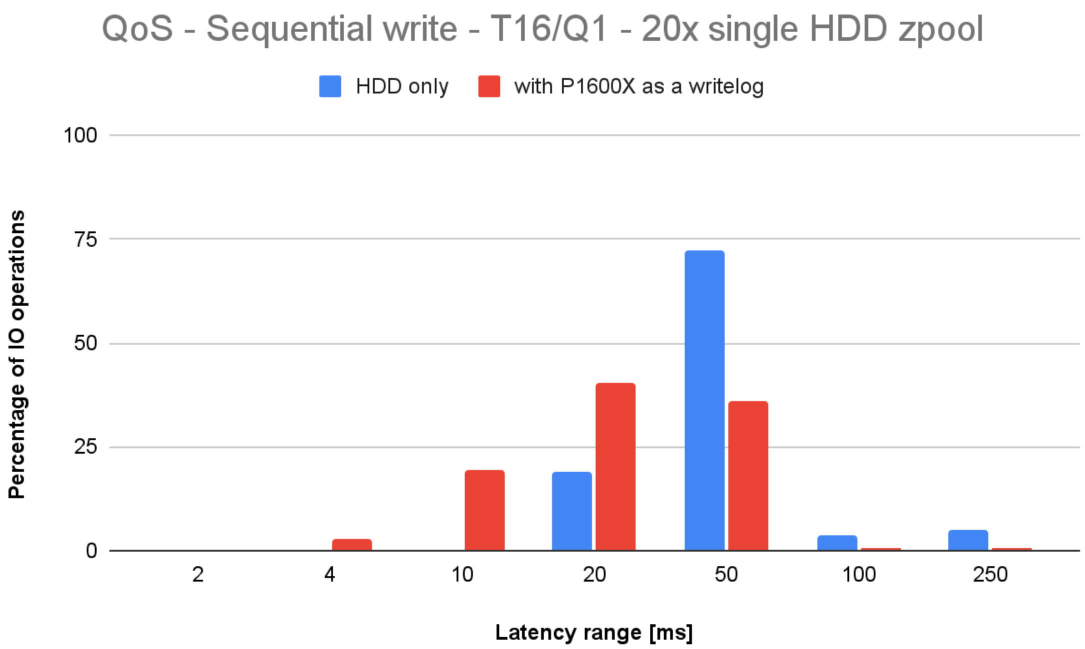


Fig. 21. Latency distribution for a sequential write test, done on a 20x single HDD zpool (T16/Q1), with and without writelog.

### 6.3. Zpool with 2x 2-way mirror QLC NAND SSD data groups

In addition to HDD disks, tests with QLC NAND SSD drives were also carried out to check the influence of the P1600X when used as a writelog device for an all-flash storage solution. In this case 4 Intel P4320 2 TB drives were used to create a zpool in a 2x 2-way mirror configuration. Figure 22 presents the results of the performance for random writes on the zpool both with and without a writelog device. Latency distribution for both the T1/Q1 and T16/Q1 combinations are included in Figures 23 and 24. Sequential write results are shown in Figure 25 along with the latencies, in Fig. 26-27 respectively. Test parameters are described in Table 6.

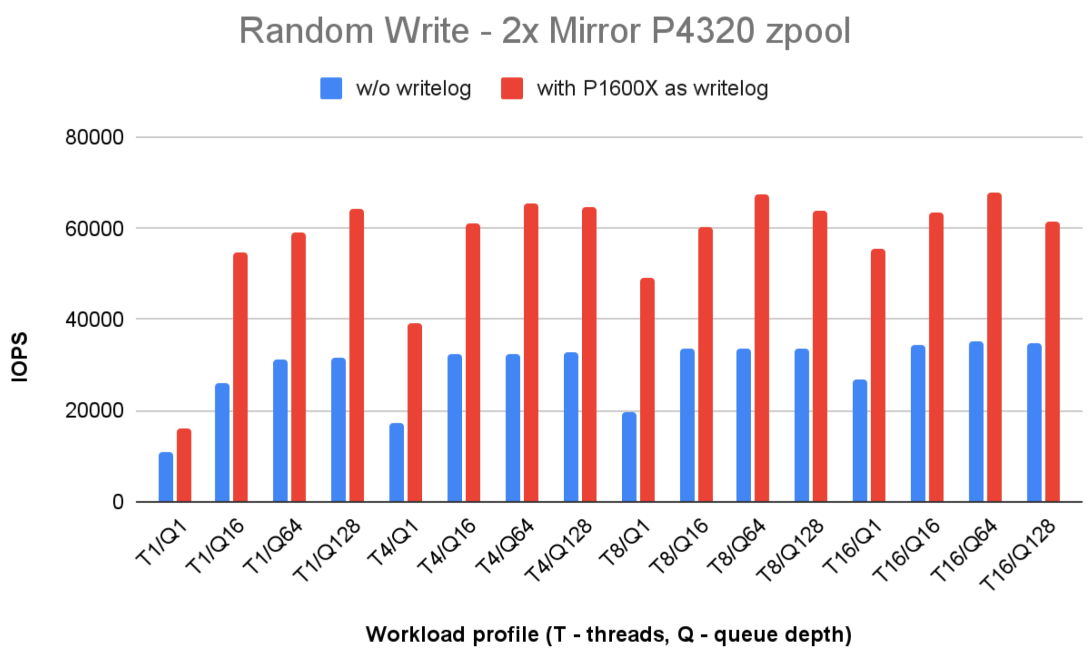


Fig. 22. Random write performance of a 2x 2-way mirror SSD zpool, with and w/o writelog.

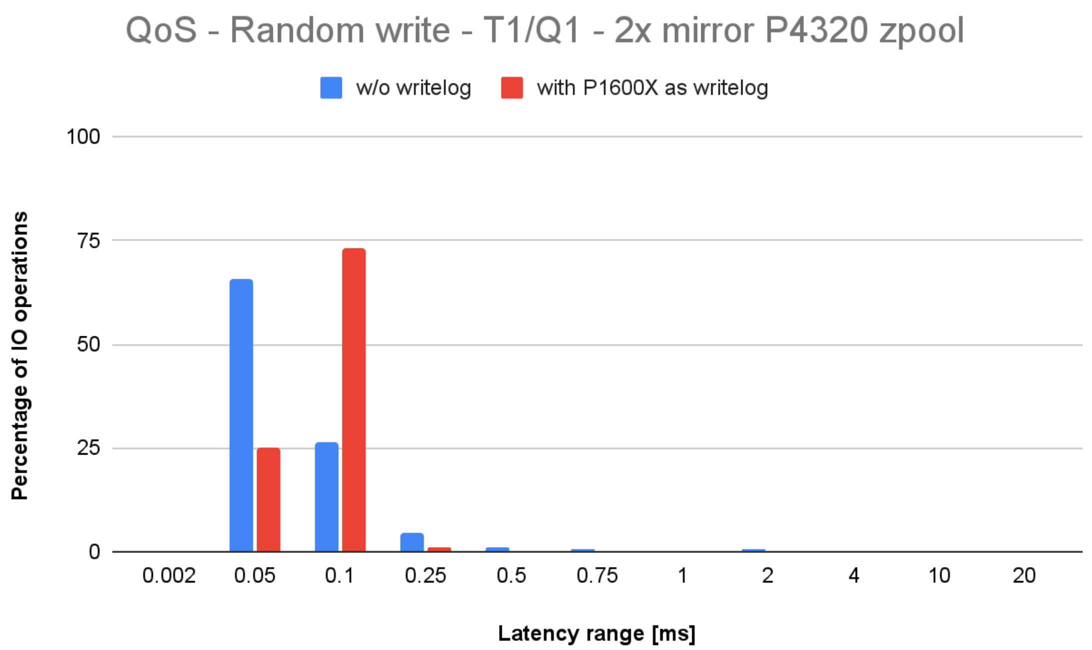


Fig. 23. Latency distribution for a random write test, done on a 2x 2-way mirror SSD zpool (T1/Q1), with and without writelog.

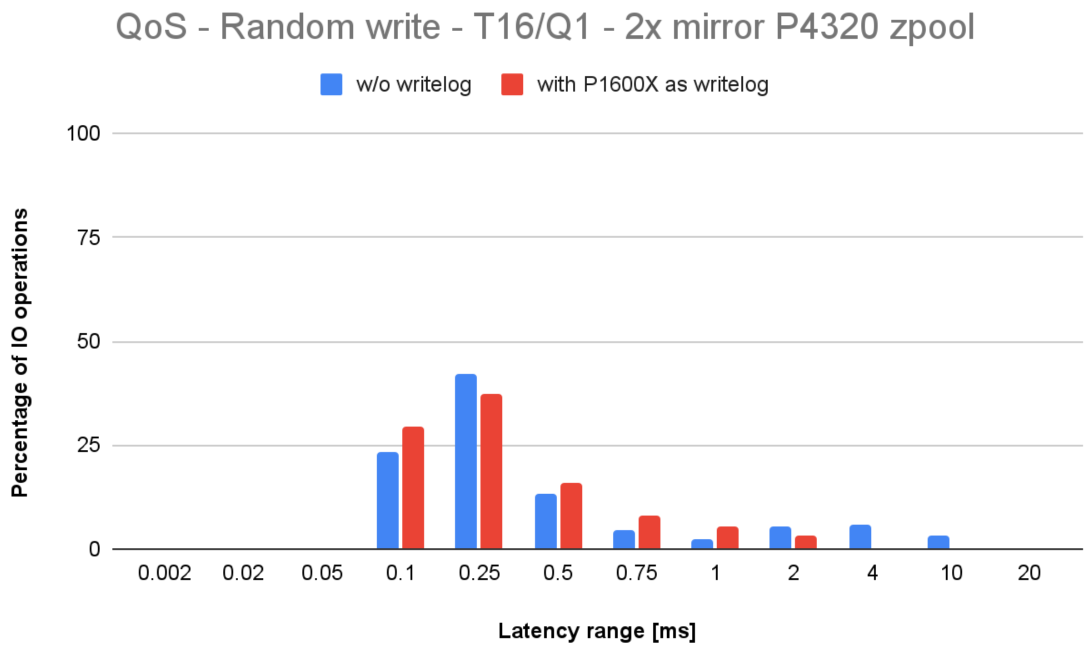


Fig. 24. Latency distribution for a random write test, done on a 2x 2-way mirror SSD zpool (T16/Q1), with and without writelog.

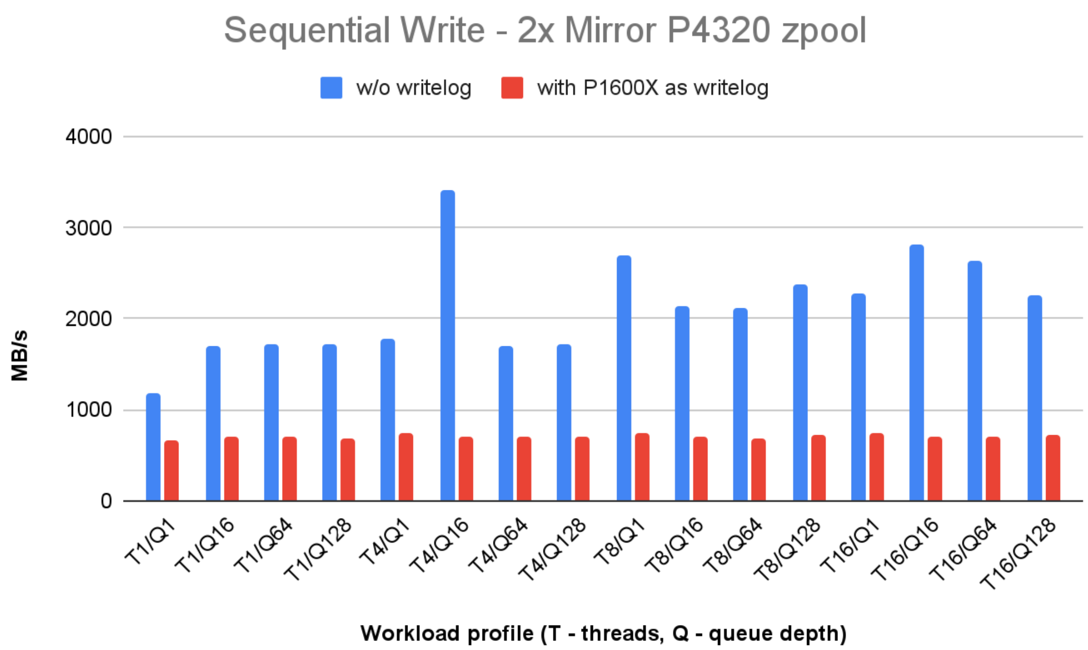


Fig. 25. Sequential write performance of a 2x 2-way mirror SSD zpool, with and w/o writelog.

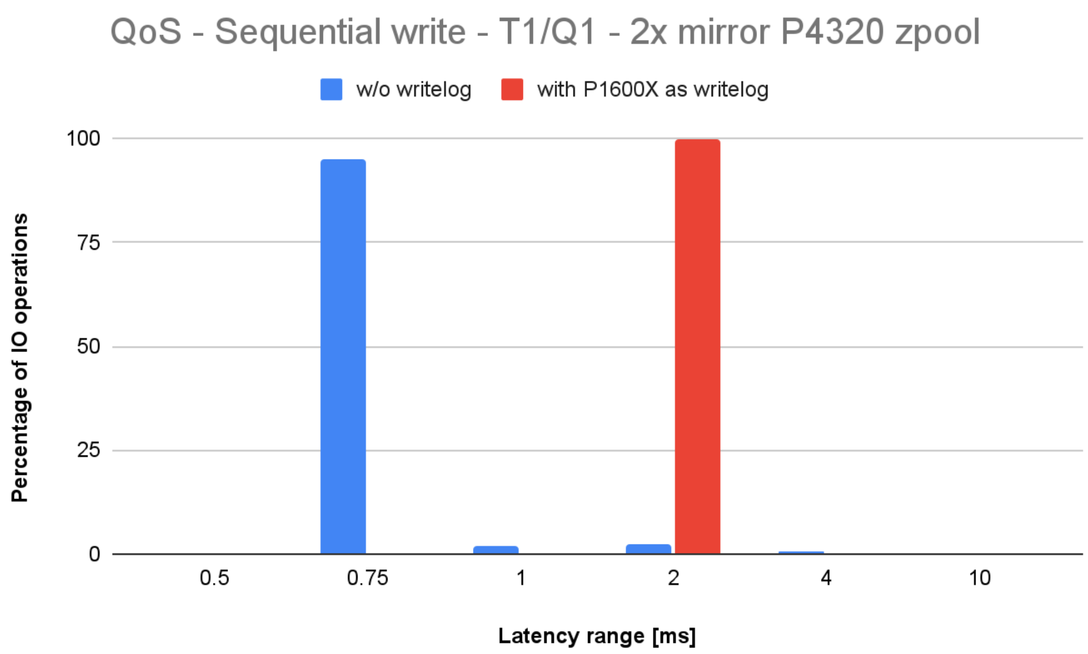


Fig. 26. Latency distribution for a sequential write test, done on a 2x 2-way mirror SSD zpool (T1/Q1), with and without writelog.

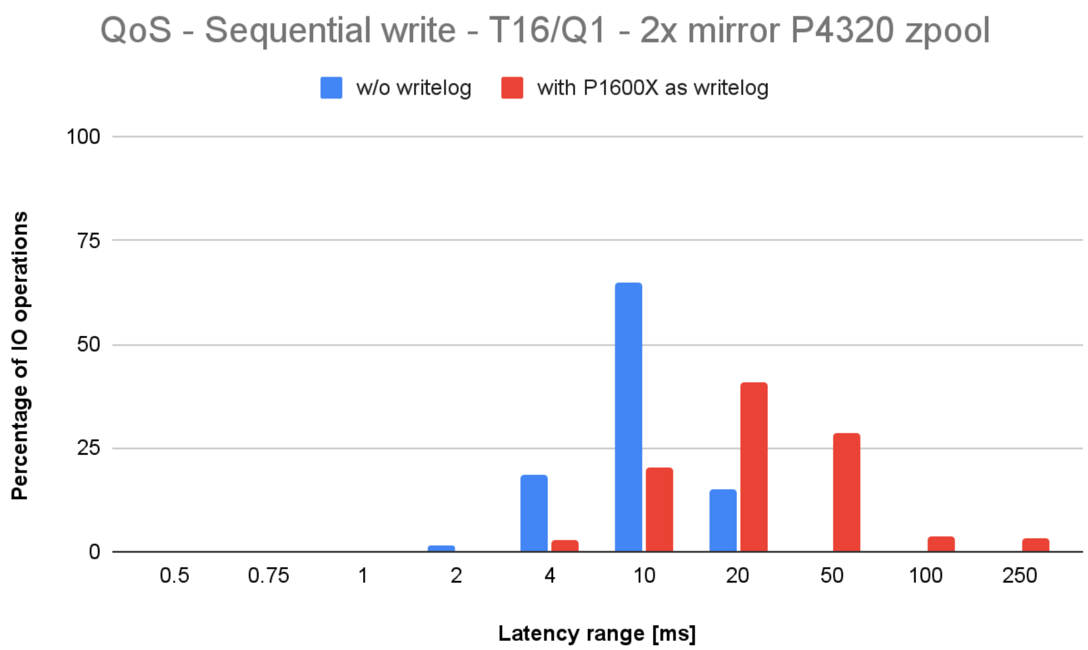


Fig. 27. Latency distribution for a sequential write test, done on a 2x 2-way mirror SSD zpool (T16/Q1), with and without writelog.

## 6.4. Zpool with 4x single QLC NAND SSD data groups

Another test including 4 Intel P4320 NAND SSD drives was carried out on a zpool created with 4 single-disk data groups. Similarly to the previous test, the IOPS and latency were measured under a load generated by fio on a zpool with and without Intel P1600X as a writelog device. Test parameters are shown in Table 6. Figures 28 and 31 present the IOPS results for random and sequential write both with and without the writelog respectively.

Latency distributions for T1/Q1 and T16/Q1 thread/queue depth combinations are presented in Figures 29-30 and 32-33.

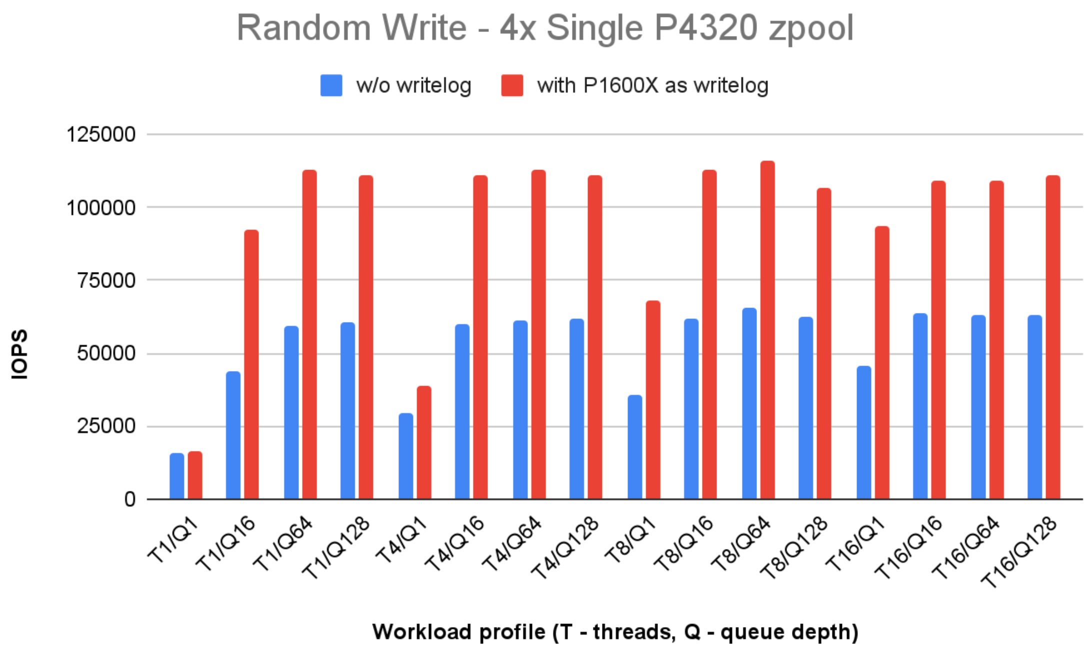


Fig. 28. Random write performance of a 4x single SSD zpool, with and w/o writelog.

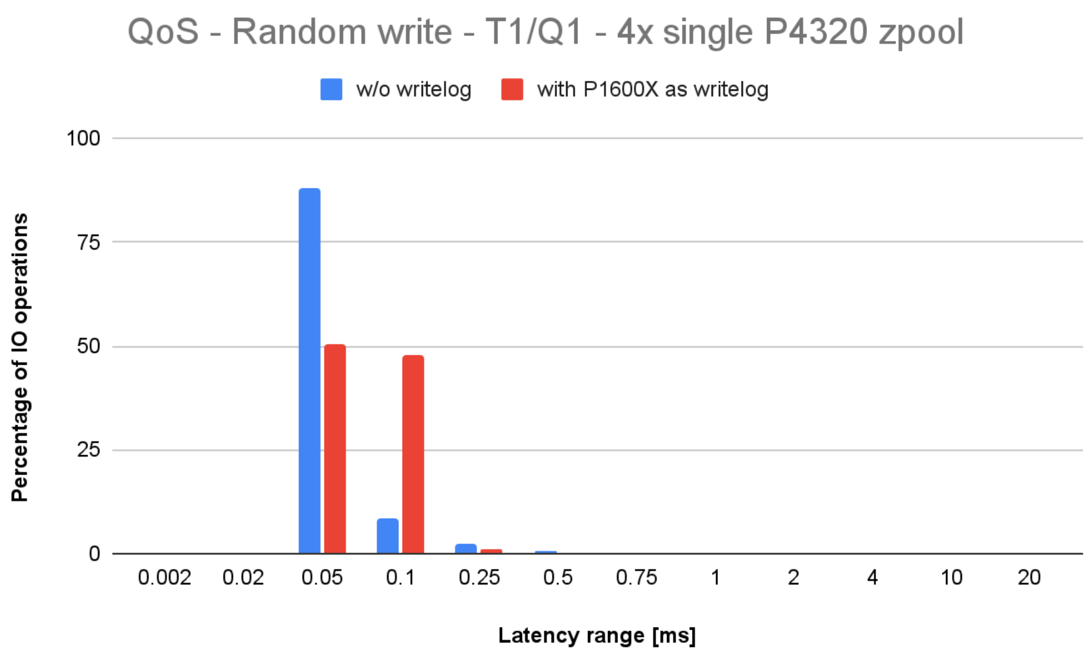


Fig. 29. Latency distribution for a random write test, done on a 4x single SSD zpool (T1/Q1), with and without writelog.

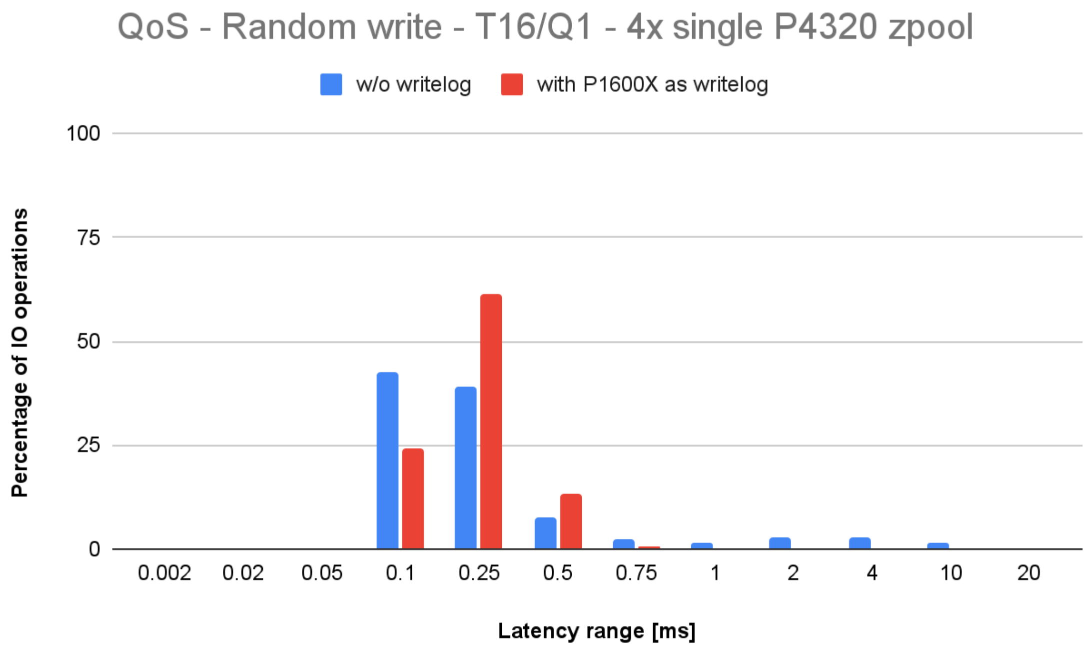


Fig. 30. Latency distribution for a random write test, done on a 4x single SSD zpool (T16/Q1), with and without writelog.

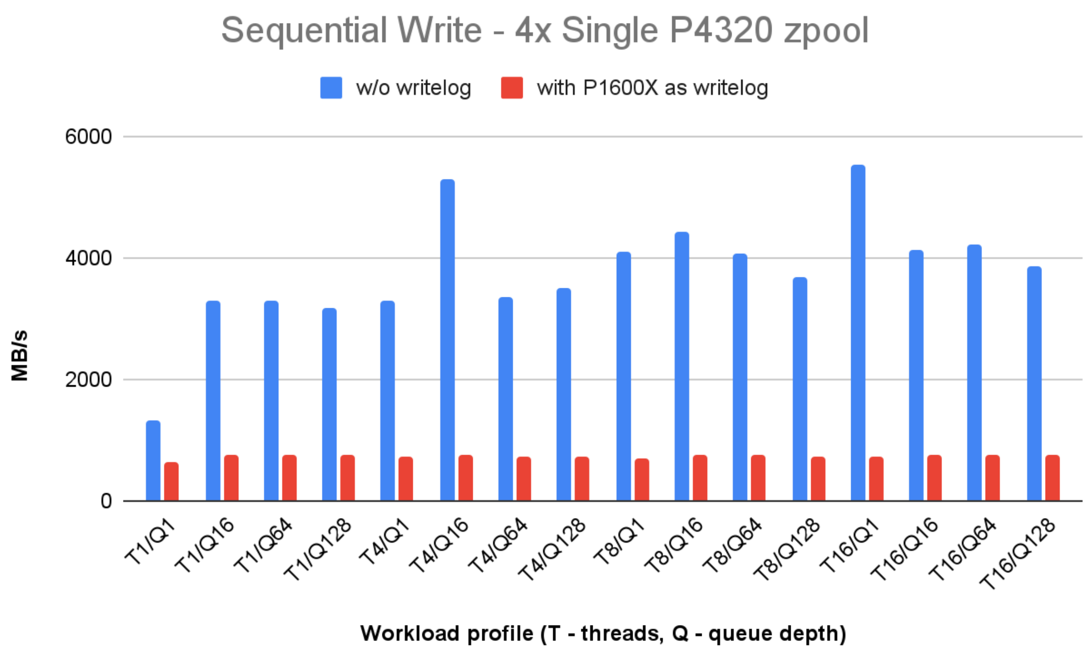


Fig. 31. Sequential write performance of a 4x single SSD zpool, with and w/o writelog.

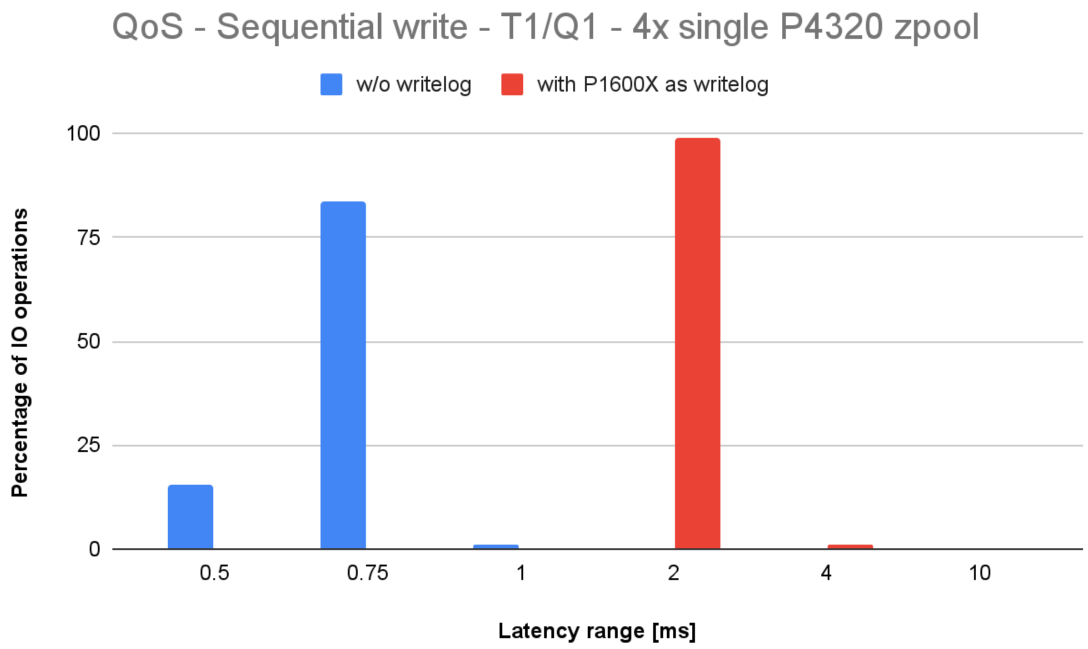


Fig. 32. Latency distribution for a sequential write test, done on a 4x single SSD zpool (T1/Q1), with and without writelog.

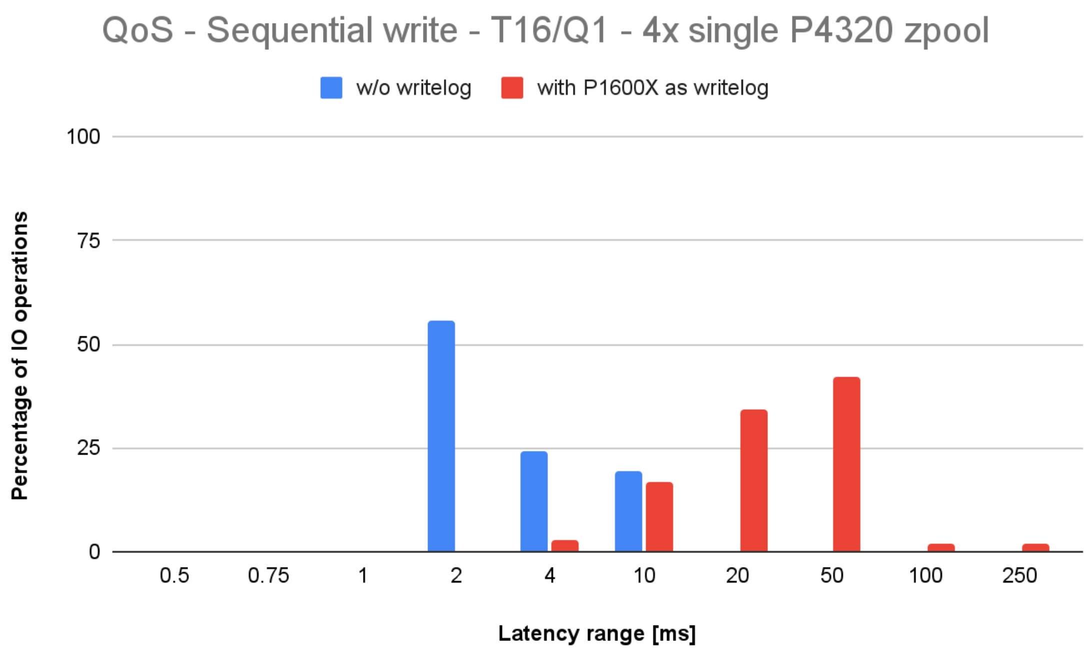


Fig. 33. Latency distribution for a sequential write test, done on a 4x single SSD zpool (T16/Q1), with and without writelog.

## 6.5. Test conclusions

Random write IO results for both performance and latency improved significantly for zpools with an added writelog over those without one, which can be seen in Figures 10-12 and 16-18. As for the sequential write test, results for the 10x 2-way mirror HDD zpool indicate a notable improvement both in throughput (Figure 13) and latency (Figures 14-15). In the case involving throughput with a 20x single HDD zpool (Fig. 19), a significant

performance improvement can be observed for queue depth equal to 1. In all other cases, this performance gain was marginal or non-existent.
Similarly to the above described tests, the 4x Intel P4320 drives were used as storage devices for a zpool consisting of SSD drives. In both the 2x 2-way mirror and the 4x single drive configurations, nearly doubled IOPS results were observed for all random writes (except in the 1T/1Q instance) in test cases with an added writelog as opposed to those without it (Figures 22 and 28). The latency distribution presented in Figures 23-24 and 29-30 may at first glance appear to suggest that the non-writelog configuration had a better result, but upon closer inspection, the average latency is actually higher in that case. A small percentage of write operations can take a particularly long time to complete, increasing the overall latency and decreasing total throughput. When using a sequential write IO pattern however, both zpool configurations seem to suffer from adding a writelog in terms of throughput (Figures 25, 31) and latency (Figures 26-27 and 32-33). Depending on the type of workload, it may be beneficial to use an Intel P1600X as a writelog device and configure each individual zvols' logbias option with a "Write log device (Latency)" setting for random heavy IO loads and "In Pool (Throughput)" for use cases depending largely on sequential IOs. However, this is a workaround which does not provide simultaneously improved sequential and random write performance. This setup will allow random write requests to utilize the benefits of a writelog device or let the sequential write requests bypass the writelog device for greater throughput but not both at the same time.

# 7. HA Storage Cluster compatibility test

In order to ensure proper operation of the Intel D7-P5510 drive in Open-E JovianDSS High Availability Storage cluster environments, various compatibility tests were performed.

## 7.1. Functional tests

All essential and critical cluster mechanisms were examined for correct operation with the tested drives used as boot medium (58 GB drive model) and writelog device (118 GB). They are summarized in Table 7.

Table 7. Results for the HA 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 |

## 7.2. Test conclusions

Compatibility of the tested device with essential HA cluster operations was extensively checked. None of the test cases described in Table 7 showed any undesirable behaviour, indicating full compatibility with Open-E JovianDSS in cluster configurations.

# 8. Summary

The Intel Optane P1600X drive was comprehensively tested for full functional compatibility with Open-E JovianDSS. Performance characteristics were also tested in several use cases. Both the Single node and HA cluster operations were taken into consideration. The tests were designed to find any abnormalities in the device used, regardless of whether it was used as a boot medium or as a writelog for Open-E JovianDSS. Given the results achieved in testing, the examined device can now safely be added to the Hardware Certification List and granted "Certified by Open-E" status.