

SPDK NVMe BDEV Performance Report Release 23.01

Testing Date: February 2023

Performed by:

Karol Latecki (karol.latecki@intel.com)

Jaroslav Chachulski (jaroslawx.chachulski@intel.com)

Acknowledgments:

James Harris (james.r.harris@intel.com)

John Kariuki (john.k.kariuki@intel.com)

Krzysztof Karas (krzysztof.karas@intel.com)

Contents

| | |
|--|----|
| Contents | 2 |
| Audience and Purpose..... | 3 |
| Test setup | 4 |
| Hardware configuration | 4 |
| BIOS Settings | 5 |
| SSD Preconditioning | 5 |
| Introduction to SPDK Block Device Layer..... | 6 |
| Test Case 1: SPDK NVMe BDEV IOPS/Core Test | 9 |
| SPDK NVMe BDEV Single Core Throughput | 10 |
| Bdevperf vs. fio IOPS/Core results..... | 12 |
| NVMe BDEV vs. Polled-Mode Driver IOPS/Core..... | 12 |
| Conclusions | 13 |
| Test Case 2: SPDK NVMe BDEV I/O Cores Scaling | 14 |
| Results | 15 |
| Conclusions | 15 |
| Test Case 3: SPDK NVMe BDEV Latency | 17 |
| Average and tail latency comparison..... | 19 |
| Linux Kernel libaio Histograms | 21 |
| Linux Kernel io_uring Histograms | 22 |
| SPDK fio Bdev Histograms | 23 |
| Performance vs. increasing Queue Depth | 24 |
| Conclusions | 26 |
| Test Case 4: IOPS vs. Latency at different queue depths..... | 27 |
| 4KiB Random Read Results..... | 29 |
| 4KiB Random Write Results | 30 |
| 4KiB Random 70%/30% Read/Write Results..... | 31 |
| Conclusions | 32 |
| Summary | 33 |
| List of tables..... | 34 |
| List of figures..... | 35 |
| References | 36 |

Audience and Purpose


This report is intended for people who are interested in comparing the performance of the SPDK block device layer vs the Linux Kernel (6.1.6-200.fc37.x86_64) block device layer. It provides performance and efficiency information between the two block layers under various test workloads.

The purpose of the report is not to imply a single “correct” approach, but rather to provide a baseline of well-tested configurations and procedures with repeatable and reproducible results. This report can be viewed as information regarding best known method/practice when performance testing the SPDK NVMe block device.

Test setup

Hardware configuration

Table 1: Hardware setup configuration

| Item | Description | | | | | | | | | | | | | | | | | | |
|----------------------|---|----|----|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| Server Platform | Ultra SuperServer SYS-220U-TNR  | | | | | | | | | | | | | | | | | | |
| Motherboard | Server board X12DPU-6 | | | | | | | | | | | | | | | | | | |
| CPU | 2 CPU sockets, Intel(R) Xeon(R) Gold 6348 CPU @ 2.60GHz Number of cores 28 per socket, number of threads 56 per socket Both sockets populated Microcode: 0xd0002e0 | | | | | | | | | | | | | | | | | | |
| Memory | 16 x 32GB SK Hynix DDR4 HMA84GR7DJR4N-XN; Total 512 GBs. Memory channel population: <table border="1"> <thead> <tr> <th>P1</th><th>P2</th></tr> </thead> <tbody> <tr> <td>CPU1_DIMM_A1</td><td>CPU2_DIMM_A1</td></tr> <tr> <td>CPU1_DIMM_B1</td><td>CPU2_DIMM_B1</td></tr> <tr> <td>CPU1_DIMM_C1</td><td>CPU2_DIMM_C1</td></tr> <tr> <td>CPU1_DIMM_D1</td><td>CPU2_DIMM_D1</td></tr> <tr> <td>CPU1_DIMM_E1</td><td>CPU2_DIMM_E1</td></tr> <tr> <td>CPU1_DIMM_F1</td><td>CPU2_DIMM_F1</td></tr> <tr> <td>CPU1_DIMM_G1</td><td>CPU2_DIMM_G1</td></tr> <tr> <td>CPU1_DIMM_H1</td><td>CPU2_DIMM_H1</td></tr> </tbody> </table> | P1 | P2 | CPU1_DIMM_A1 | CPU2_DIMM_A1 | CPU1_DIMM_B1 | CPU2_DIMM_B1 | CPU1_DIMM_C1 | CPU2_DIMM_C1 | CPU1_DIMM_D1 | CPU2_DIMM_D1 | CPU1_DIMM_E1 | CPU2_DIMM_E1 | CPU1_DIMM_F1 | CPU2_DIMM_F1 | CPU1_DIMM_G1 | CPU2_DIMM_G1 | CPU1_DIMM_H1 | CPU2_DIMM_H1 |
| P1 | P2 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_A1 | CPU2_DIMM_A1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_B1 | CPU2_DIMM_B1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_C1 | CPU2_DIMM_C1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_D1 | CPU2_DIMM_D1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_E1 | CPU2_DIMM_E1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_F1 | CPU2_DIMM_F1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_G1 | CPU2_DIMM_G1 | | | | | | | | | | | | | | | | | | |
| CPU1_DIMM_H1 | CPU2_DIMM_H1 | | | | | | | | | | | | | | | | | | |
| Operating System | Fedora 37 | | | | | | | | | | | | | | | | | | |
| BIOS | 1.4a | | | | | | | | | | | | | | | | | | |
| Linux kernel version | 6.1.6-200.fc37.x86_64 Spectre-meltdown mitigations enabled | | | | | | | | | | | | | | | | | | |
| SPDK version | SPDK 23.01 | | | | | | | | | | | | | | | | | | |
| Fio version | 3.28 | | | | | | | | | | | | | | | | | | |
| Storage | OS: 1x 250GB Crucial CT250MX500SSD1 Storage: 22x Kioxia® KCM61VUL3T20 3.2TBs (FW: 0105) (10 on CPU NUMA Node 0, 12 on CPU NUMA Node 1) | | | | | | | | | | | | | | | | | | |

BIOS Settings

Table 2: Test setup BIOS settings

| Item | Description |
|------|--|
| BIOS | VT-d = Enabled CPU Power and Performance Policy = <Performance> CPU C-state = No Limit CPU P-state = Enabled Enhanced Intel® Speedstep® Tech = Enabled Turbo Boost = Enabled Hyper Threading = Enabled |

Table 3: Test System NVMe storage setup

| Table 37: Test system NVMe Storage Setup | | |
|--|---|-------------------------------------|
| Item | Description | |
| PCIe Riser cards | “Ultra” Riser Card: AOC-2UR68G4-i2XT <ul style="list-style-type: none">• PCIe Slot 1 – x16, CPU2• PCIe Slot 2 – x8, CPU2• PCIe Slot 3 – x8, CPU2 Right-facing riser card: RSC-WR-6 <ul style="list-style-type: none">• PCIe Slot 4 – x16, CPU1 Left-facing riser card: RSC-W2-66G4 <ul style="list-style-type: none">• PCIe Slot 5 – x16, CPU2• PCIe Slot 7 – x16, CPU1 More information can be found in SYS-220U-TNR manual document . | |
| | PCIe Retimer cards | |
| NVMe Drives distribution across the system | 3 x AOC-SLG4-4E4T Installed in: <ul style="list-style-type: none">○ PCIe Retimer 1: RSC-WR-6, PCIe Slot 4 (using CPU1 PCIe Lanes)○ PCIe Retimer 2: AOC-2UR68G4-i2XT, PCIe Slot 1 (using CPU2 PCIe Lanes)○ PCIe Retimer 3: RSC-W2-66G4, PCIe Slot 5 (using CPU2 PCIe Lanes) | |
| | Nvme0 – 5 | Motherboard ports (CPU1 PCIe Lanes) |
| | Nvme6 – 9 | Motherboard ports (CPU2 PCIe Lanes) |
| | Nvme9 – 13 | PCIe Retimer 1 (CPU1 PCIe Lanes) |
| | Nvme14 - 17 | PCIe Retimer 2 (CPU2 PCIe Lanes) |
| | Nvme18 - 21 | PCIe Retimer 3 (CPU2 PCIe Lanes) |

SSD Preconditioning

An empty NAND SSD will often show read performance far beyond what the drive claims to be capable of because the NVMe controller knows that the device is empty and completes the read request successfully without performing any actual read operation on the device. Therefore, prior to running each performance test case we preconditioned the SSDs by writing 128K blocks sequentially across the namespace’s full LBA range twice to ensure the controller accesses the NAND media for each subsequent I/O. Additionally, the 4K 100% random writes performance decreases from one test to the next until the NAND management overhead reaches steady state because the wear-levelling activity increases dramatically until the SSD reaches steady state. Therefore, to obtain accurate and repeatable results for the 4K 100% random write workload, we ran the workload for 60 minutes before starting the benchmark test and collecting performance data. For a highly detailed description of exactly how to force an SSD into a known state for benchmarking see the [SNIA Solid State Storage Performance Test Specification](#).

Introduction to SPDK Block Device Layer

SPDK Polled Mode Driver

The NVMe PCIe driver is something that is usually expected to be part of the system kernel and your application would interact with the driver via the system call interface. SPDK takes a different approach. SPDK unbinds the NVMe devices from the kernel NVMe driver and binds them to a userspace NVMe driver instead. This allows a userspace application to directly access the device and its queues from userspace.

The [SPDK NVMe Driver](#) is a C library that may be linked directly into an application that provides direct, zero-copy data transfer to and from NVMe SSDs. It is entirely passive, meaning that it spawns no threads and only performs actions in response to function calls from the application. The library controls NVMe devices by directly mapping the PCI BAR into the local process and performing MMIO. The SPDK NVMe driver is asynchronous, which means that the driver submits the I/O request as an NVMe submission queue entry on a queue pair and the function returns immediately, prior to the completion of the NVMe command. The application must poll for I/O completion on each queue pair with outstanding I/O to receive completion callbacks.

SPDK Block Device Layer

SPDK further provides a full block stack as a user space library that performs many of the same operations as a block stack in an operating system. The [SPDK block device layer](#) often simply called **bdev**, is a C library intended to be equivalent to the operating system block storage layer located above the device drivers in traditional kernel storage stack.

The bdev module provides an abstraction layer with common APIs for implementing block devices that interface with different types of block storage device. An application can use the APIs to enumerate and claim SPDK block devices, and then perform asynchronous I/O operations (such as read, write, unmap, etc) in a generic way without knowing if the device is an NVMe device or something else, for example Ceph RBD or malloc ramdisk block device. The SPDK NVMe bdev module can create block devices for both local PCIe-attached NVMe device and remote devices exported over NVMe-oF.

In this report, we benchmarked the performance and efficiency of the bdev for the local PCIe-attached NVMe devices use case. We also demonstrated the benefits of the SPDK approaches, like user-space polling, asynchronous I/O, no context switching etc. under different workloads.

FIO Integration

SPDK provides an [fio plugin](#) for integration with [Flexible I/O](#) benchmarking tool. The quickest way to generate a configuration file with all the bdevs for locally PCIe-attached NVMe devices is to use the `gen_nvme.sh` script with “—json-with-subsystems” option as shown in Figure 1.

```
[user@localhost spdk]$ sudo scripts/gen_nvme.sh --json-with-  
subsystems | jq  
{  
  "subsystems": [  
    {  
      "subsystem": "bdev",  
      "config": [  
        {  
          "method": "bdev_set_options",  
          "params": {  
            "bdev_io_pool_size": 65535,  
            "bdev_io_cache_size": 2048,  
            "bdev_auto_examine": true  
          }  
        },  
        {  
          "method": "bdev_nvme_attach_controller",  
          "params": {  
            "trtype": "PCIe",  
            "name": "Nvme0",  
            "traddr": "0000:1a:00.0"  
          }  
        }  
      ],  
      [...]  
    },  
    {  
      "method": "bdev_nvme_attach_controller",  
      "params": {  
        "trtype": "PCIe",  
        "name": "Nvme23",  
        "traddr": "0000:df:00.0"  
      }  
    }  
  ]  
}
```

Figure 1 : Example NVMe bdev configuration file

Add SPDK bdevs to the fio job file, by setting the *ioengine=spdk_bdev* and adding the *spdk_json_conf* parameter whose value points to the NVMe bdev configuration file.

The example fio configuration file in Figure 2, shows how to define multiple fio jobs and assign NVMe bdevs to each job. Each job is also pinned to a CPU core on the same NUMA node as the NVMe SSDs that the job will access.

Finally, to use the bdev fio plugin specify the LD_PRELOAD when running fio.

LD_PRELOAD=<path to spdk repo>/examples/bdev/fio_plugin/fio_plugin fio <fio job file>

```
[global]
direct=1
thread=1
time_based=1
norandommap=1
group_reporting=1
ioengine=spdk_bdev
spdk_json_conf=/tmp/bdev.conf

rw=randread
rwmixread=70
bs=4096
numjobs=1
runtime=300
ramp_time=60

[filename0]
iodepth=192
cpus_allowed=0
filename=Nvme0n1
filename=Nvme1n1
filename=Nvme4n1
filename=Nvme5n1
filename=Nvme6n1
filename=Nvme7n1

[filename1]
iodepth=192
cpus_allowed=21
filename=Nvme2n1
filename=Nvme3n1
filename=Nvme8n1
filename=Nvme9n1
filename=Nvme10n1
filename=Nvme11n1

[filename2]
iodepth=192
cpus_allowed=22
filename=Nvme12n1
filename=Nvme13n1
filename=Nvme14n1
filename=Nvme15n1
filename=Nvme16n1

[filename3]
iodepth=192
cpus_allowed=23
filename=Nvme17n1
filename=Nvme18n1
filename=Nvme19n1
filename=Nvme20n1
filename=Nvme21n1
```

Figure 2: Example SPDK Fio BDEV configuration file

Test Case 1: SPDK NVMe BDEV IOPS/Core Test

Purpose: The purpose of this test case was to measure the maximum performance in IOPS/Core of the NVMe block layer on a single CPU core. We used different benchmarking tools (SPDK bdevperf vs. SPDK fio bdev plugin vs SPDK NVMe perf) to understand the overhead of benchmarking tools. Measuring IOPS was key in this test case, so latency measurements were either disabled or skipped.

The following Random Read/Write workloads were used:

- 4KiB 100% Random Read
- 4KiB 100% Random Write
- 4KiB Random 70% Read 30% Write

For each workload we followed the following steps:

- 1) Precondition SSDs as described in [“Test Setup”](#) chapter.
- 2) Run each test workload: Start with a configuration that has 22 Kioxia KCM61VUL3T20 NVMe devices and decrease the number of SSDs on each subsequent run.
 - This shows us the IOPS scaling as we add SSDs till the maximum IOPS/Core is reached.
 - Starting with 22 SSDs and reducing the number of SSDs on subsequent runs eliminates having to precondition between runs because all SSDs were used in the previous run, so they should still be in a steady state.
- 3) Repeat three times. The data reported is the average of the 3 runs.

Table 4: SPDK NVMe BDEV IOPS Test configuration

| Item | Description |
|------------------------|---|
| Test case | SPDK NVMe BDEV IOPS/Core Test |
| Test configuration | fio version: fio-3.28 Number of NVMe SSDs: {1, 2, 3, 4, 5, 6, 7, 8, 10, 12, 14, 16, 18, 20, 22} SPDK_BDEV_IO_CACHE_SIZE changed from 256 to 2048 (using bdev_set_options RPC call). |
| Bdevperf configuration | spdk/test/bdev/bdevperf/bdevperf -c bdev.conf -q \${iodepth} -o \${block_size} -w \${rw} -M \${rwmixread} -t 300 -m 0 -p 0 |
| fio configuration | [global] ioengine=spdk_bdev spdk_json_conf=bdev.conf |

| | |
|-----------------------------------|---|
| | <pre> gtod_reduce=1 direct=1 thread=1 norandommap=1 time_based=1 ramp_time=60s runtime=300s bs=4k numjobs=1 </pre> |
| (Random read and mixed workloads) | <pre> rw={randread, randrw} rwmixread={100,70} iodepth={128, 192, 256} </pre> |
| (Random write workload) | <pre> rw=randwrite rwmixread=0 iodepth={32, 64, 128} </pre> |

SPDK NVMe BDEV Single Core Throughput

The first test was performed using SPDK bdevperf, which is lightweight benchmarking tool that adds minimal latency to the I/O path. The charts below show the single core IOPS results for the SPDK Block Layer with increasing number of NVMe SSDs.

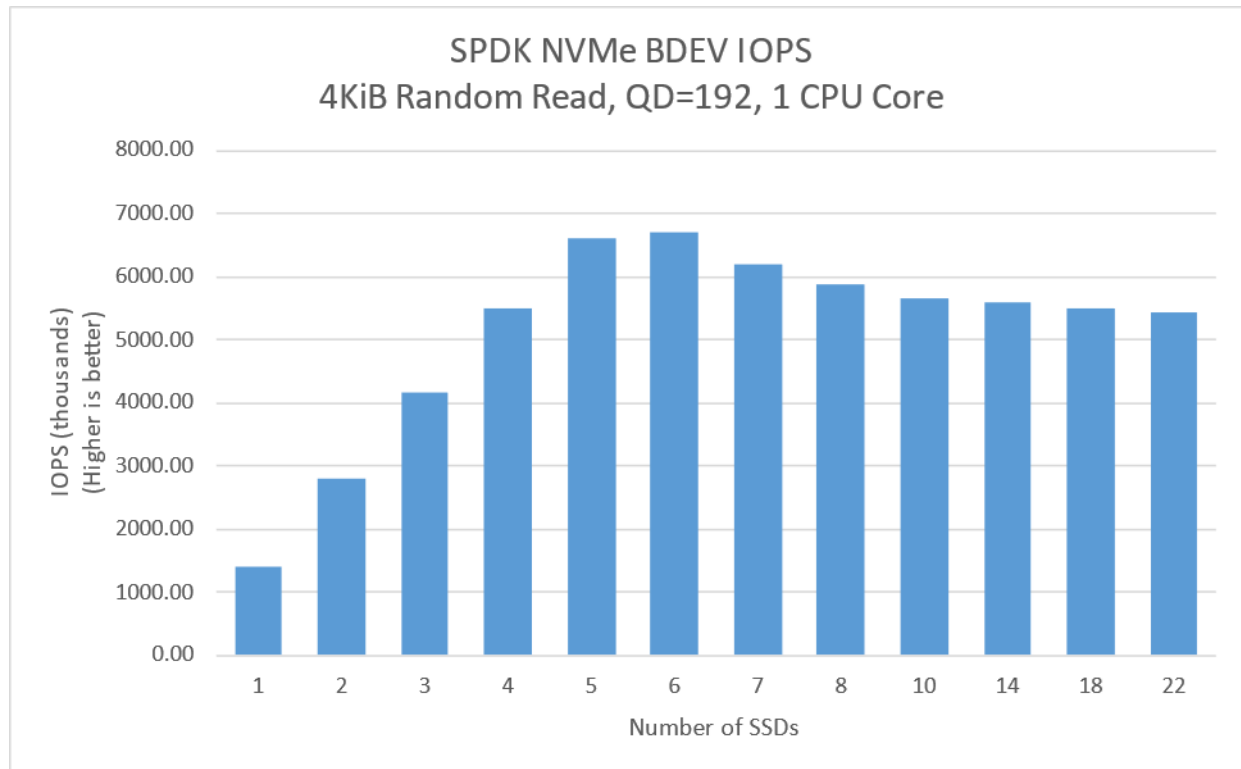


Figure 3: SPDK NVMe BDEV IOPS scalability with addition of SSDs (4KiB Random Read, 1CPU Core, QD=192, using bdevperf tool)




Figure 4: SPDK NVMe BDEV IOPS scalability with addition of SSDs (4KiB Random Write, 1CPU Core, QD=32, using bdevperf tool)

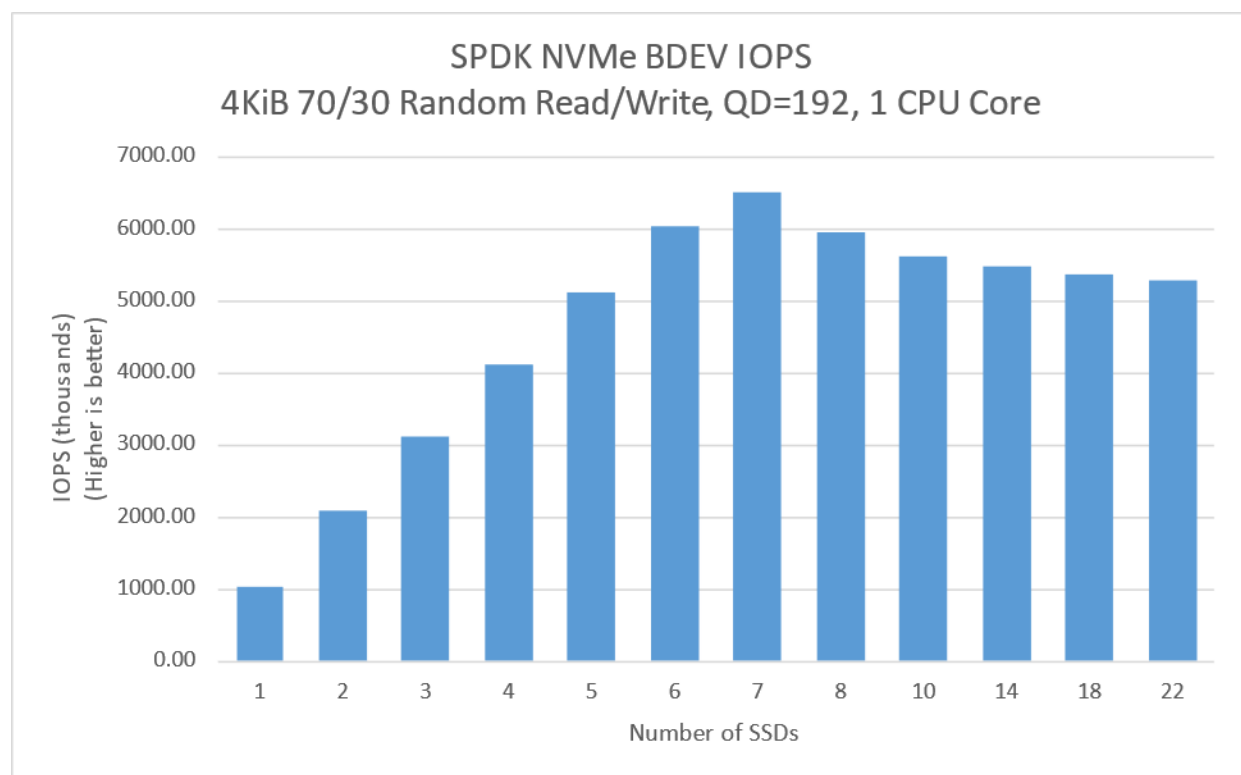


Figure 5: SPDK NVMe BDEV IOPS scalability with addition of SSDs (4KiB 70/30 Random Read/Write, 1CPU Core, QD=192, using bdevperf tool)

Bdevperf vs. fio IOPS/Core results

SPDK provides the bdevperf benchmarking tool that provides minimal capabilities needed to define basic workloads and collects a limited amount of data. The fio benchmarking tool provides a lot of great features to enable users to quickly define workloads, scale the workloads and collect many data points for detailed performance analysis, however, at cost of higher overhead. This test compares the performance in IOPS/core of bdevperf vs. the fio benchmarking tool with the SPDK bdev plugin.

Table 5: IOPS/Core performance; SPDK fio bdev plugin vs SPDK bdevperf (Blocksize=4KiB, 1 CPU Core)

| Workload | SPDK fio bdev Plugin (IOPS, thousands) | SPDK bdevperf (IOPS, thousands) | Performance gain |
|---|---|------------------------------------|---------------------|
| 4KiB Random Read, QD=192, 5 SSDs | 3417.84 | 6624.73 | 93.8% |
| 4KiB Random Write, QD=32, 10 SSDs | 2638.50 | 7842.75 | 197.2% |
| 4KiB 70/30 Random Read/Write, QD=192, 7 SSDs | 2927.39 | 6498.77 | 122.0% |

The overhead of the benchmarking tools is important when you are testing a system that is capable of millions of IOPS/Core. Using a benchmarking tool that has minimal overhead like the SPDK bdevperf yields up to 197% more IOPS/Core than fio.

NVMe BDEV vs. Polled-Mode Driver IOPS/Core

In this test case, we compared the throughput of the NVMe BDEV with that of the polled-mode driver. How to read this data? The SPDK block layer provides several key features at a cost of approximately 5.1% and 7.2% more CPU utilization for Random Read and Random Write workloads. If you are building a system with many SSDs that is capable of millions of IOPS, you can take advantage of the block layer features at the cost of approximately 1 additional CPU core for every 20 I/O cores for Random Read workload and 1 additional CPU core for every 14 I/O cores for Random Write workload. Comparison was done using SPDK bdevperf and nvmeperf test tools.

Table 6: SPDK NVMe Bdev vs SPDK NVMe PMD IOPS/Core (Blocksize=4KiB, 1 CPU Core)

| Workload | SPDK Bdevperf (IOPS, thousands) | SPDK Nvmeperf (IOPS, thousands) | Performance gain |
|--------------------------------------|------------------------------------|------------------------------------|---------------------|
| 4KiB Random Read, QD=192, 5 SSDs | 6624.73 | 6964.06 | 5.1% |
| 4KiB Random Write, QD=32, 10 SSDs | 7842.75 | 8414.56 | 7.2% |

Conclusions

1. The SPDK NVMe block device module adds approximately 5.1% and 7.2% overhead compared to using only the SPDK NVMe polled-mode driver without the block device module for Random Read and Random Write workloads respectively.
2. Performance scales linearly with addition of NVMe SSDs up to 5 NVMe SSDs for Random Read workload, reaching around 6.6 million IOPS.
3. Performance scaling is linear for Random Write workload up to 10 NVMe SSDs, reaching around 7.8 million IOPS.
4. Performance scales linearly with addition of NVMe SSDs up 7 SSDs for Random Read/Write workload, reaching around 6.5 million IOPS.
5. For all workloads there is a noticeable performance degradation with addition of more NVMe SSDs after peak performance point has been reached.
6. The IOPS for the 4 KiB Random Write workload exceeded the expected NVMe SSDs maximum throughput. We suspect this is due to imperfect preconditioning process, which wears off over time. The results, however, were repeatable for several test runs.

Test Case 2: SPDK NVMe BDEV I/O Cores Scaling

Purpose: The purpose of this test case is to demonstrate the I/O throughput scalability of the NVMe BDEV module with the addition of more CPU cores to perform I/O. The number of CPU cores used was scaled as 1, 2, 3, 4, 5 and 6.

Test Workloads: We use the following Random Read/Write mixes

- 4KiB 100% Random Read
- 4KiB 100% Random Write
- 4KiB Random 70% Read 30% Write

Table 7: SPDK NVMe BDEV I/O Cores Scalability Test

| Item | Description |
|-------------------------|---|
| Test case | Test SPDK NVMe BDEV I/O Cores Scalability Test |
| Test configuration | Number of CPU Cores: 1, 2, 3, 4, 5, 6 Number of NVMe SSDs: 5 per each CPU Core used in test, up to maximum of 22 NVMe SSDs NUMA optimization: CPUs for test were selected in a way to match NVMe drives distribution across platform NUMA nodes. |
| Bdev perf configuration | <code>spdk/test/bdev/bdevperf/bdevperf --json bdev.conf \ -q 128 -o 4096 -w randrw -M \${MIXREAD} \ -t 300 -m \${CORE_MASK} -p \${PRIMARY_CORE}</code> |

Results

Table 8: SPDK NVMe BDEV I/O Cores Scalability Test (4KiB 100% Random Read IOPS at QD=192; 4KiB 100% Random Write IOPS at QD=32; 4KiB 70/30 Random Read/Write IOPS at QD=192)

| CPU Cores | NVMe SSDs | IOPS (thousands) | | |
|-----------|-----------|--------------------|--------------------|--------------------------------|
| | | Random Read QD=192 | Random Write QD=32 | 70/30 Random Read/Write QD=192 |
| 1 | 5 | 6648.75 | 2478.65 | 4741.60 |
| 2 | 10 | 13316.76 | 6428.50 | 9827.52 |
| 3 | 15 | 19669.48 | 10695.46 | 14906.14 |
| 4 | 20 | 26104.93 | 14909.26 | 20011.20 |
| 5 | 22 | 29116.93 | 17761.53 | 22329.52 |
| 6 | 22 | 29912.52 | 18445.44 | 22680.27 |

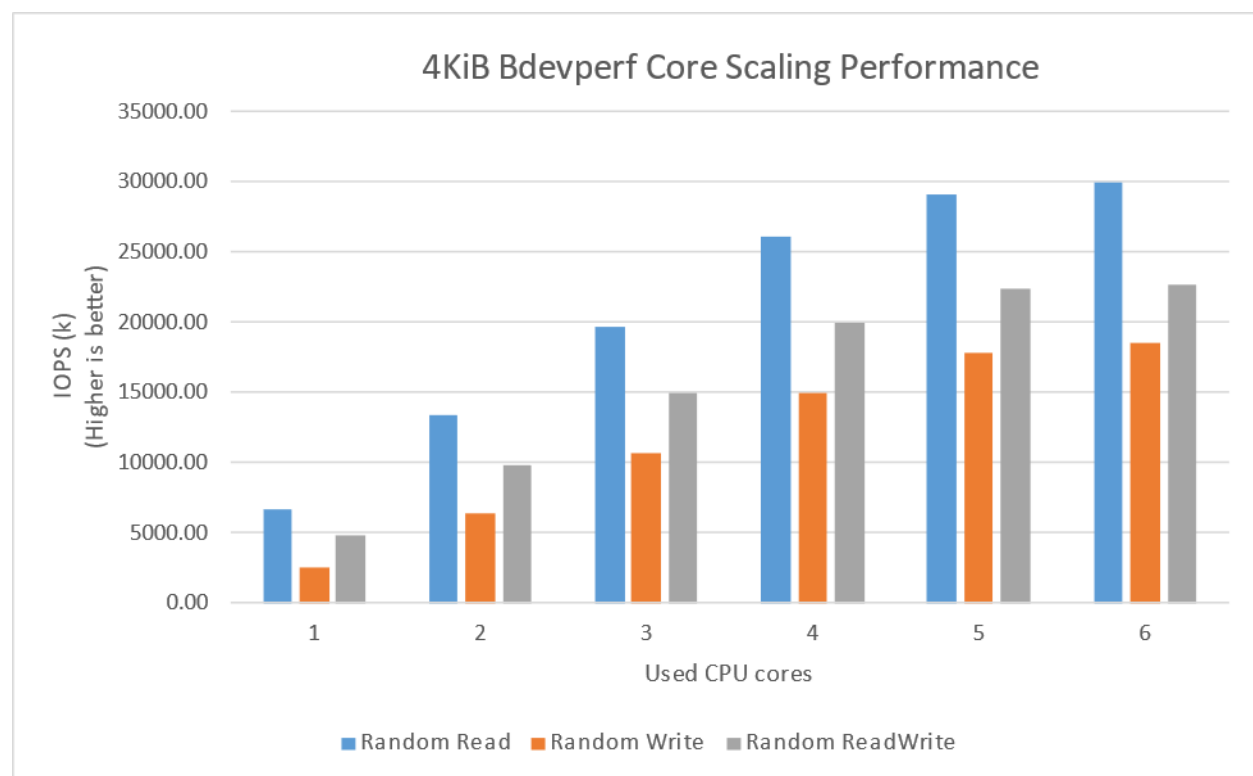


Figure 6: SPDK NVMe BDEV I/O Cores Scalability (4KiB 100% Random Read IOPS at QD=192; 4KiB 100% Random Write IOPS at QD=32; 4KiB 70/30 Random Read/Write IOPS at QD=192)

Conclusions

1. The IOPS for 4 KiB Random Read workload scales up linearly with addition of I/O cores until NVMe drives used for test are saturated.
2. The IOPS for 4 KiB Random Read/Write workloads scale up linearly with the addition of I/O cores up to 5 I/O cores. Increasing the number of cores to 6 does not result in performance improvement.

3. The IOPS for the 4 KiB Random Write workload scale up linearly. The IOPS exceeded the expected NVMe SSDs throughput for this workload which is about 7.7M IOPS. We suspect this is due to imperfect preconditioning process, which wears off over time. However, the results were repeatable and showed SPDK's high scalability with addition of I/O cores.

Test Case 3: SPDK NVMe BDEV Latency

This test case was carried out to understand latency characteristics while running SPDK NVMe bdev and its comparison to Linux Kernel NVMe block device layer. We used SPDK fio bdev plugin instead of the SPDK Bdevperf tool, as it allowed us to gather detailed latency metrics. fio was ran for 15 minutes targeting a single block device over a single NVMe drive. This test compares consistency between latency of the SPDK and Linux Kernel block layers over time in a histogram. The Linux Kernel block layer provides I/O polling capabilities to eliminate overhead such as context switch, IRQ (Interrupt Request) delivery delay and IRQ handler scheduling. This test case includes a comparison of the I/O latency for the Kernel vs. SPDK.

Test Workloads: We use the following workloads:

- 4KiB 100% Random Read
- 4KiB 100% Random Write

Table 9: SPDK NVMe BDEV Latency Test

| Item | Description |
|--|--|
| Test case | Test SPDK NVMe BDEV Latency Test |
| Test configuration | Fio version: fio-3.28 Number of CPU cores: 1 Number of NVMe SSDs: 1 |
| SPDK NVMe Driver Configuration | ioengine=spdk_bdev |
| Linux Kernel Default (libaio) Configuration | ioengine=libaio |
| Linux Kernel io_uring | ioengine=io_uring System NVMe block device configuration: echo 0 > /sys/block/nvme0n1/queue echo 0 > /sys/block/nvme0n1/rq_affinity echo 2 > /sys/block/nvme0n1/nomerges echo -1 > /sys/block/nvme0n1/io_poll_delay |
| fio configuration (common part) | [global] direct=1 thread=1 time_based=1 norandommap=1 group_reporting=1 rw={randread randwrite} bs=4096 |

| | |
|---|--|
| | runtime=900 ramp_time=120 numjobs=1 log_avg_msec=15 write_lat_log=/tmp/tc3_lat.log |
| fio configuration (SPDK specific) | [global] ioengine=spdk_bdev spdk_conf=/tmp/bdev.conf [filename0] iodepth=1 cpus_allowed=0 filename=Nvme0n1 |
| fio configuration (Linux Kernel common) | [global] ioengine={libaio io_uring} [filename0] iodepth=1 cpus_allowed=0 filename=/dev/nvme0n1 |
| fio configuration (Linux Kernel io_uring specific) | [global] fixedbufs=1 hipri=1 registerfiles=1 sqthread_poll=1 |

The Linux block layer implements I/O polling on the completion queue. Polling can remove context switch overhead, IRQ delivery and IRQ handler scheduling overhead[1].

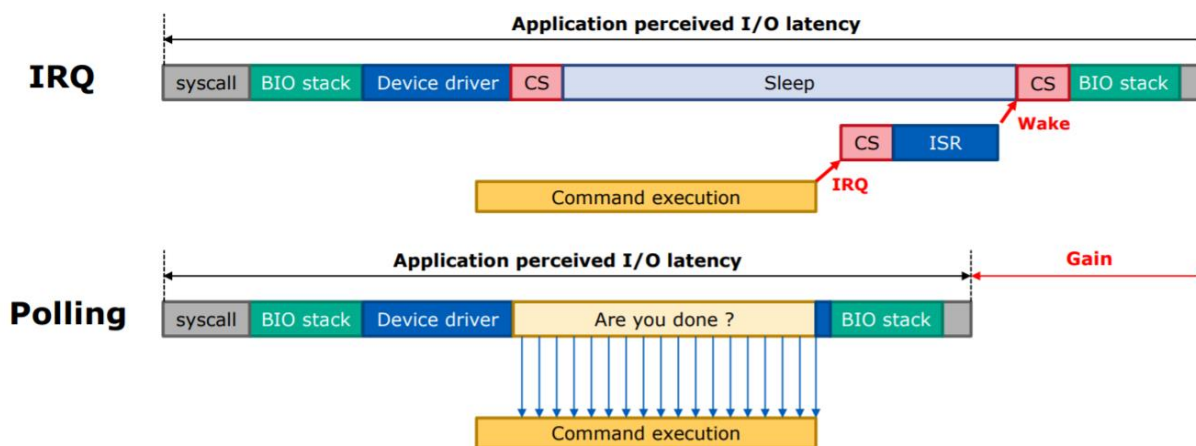


Figure 7: Linux Block Layer I/O Optimization with Polling. Source [1]

Furthermore, the Linux block I/O polling provides a mechanism to reduce the CPU load. In the *Classic Polling* model, the CPU spin-waits for the command completion and utilizes 100% of a CPU core [1]. There's also an adaptive hybrid polling which reduces the CPU load by putting the I/O polling thread to sleep for about half of the command execution time, but the polling thread must be woken up before the I/O completes with enough heads-up time for a context switch[1]. Hybrid polling mode was not used for testing in this document.

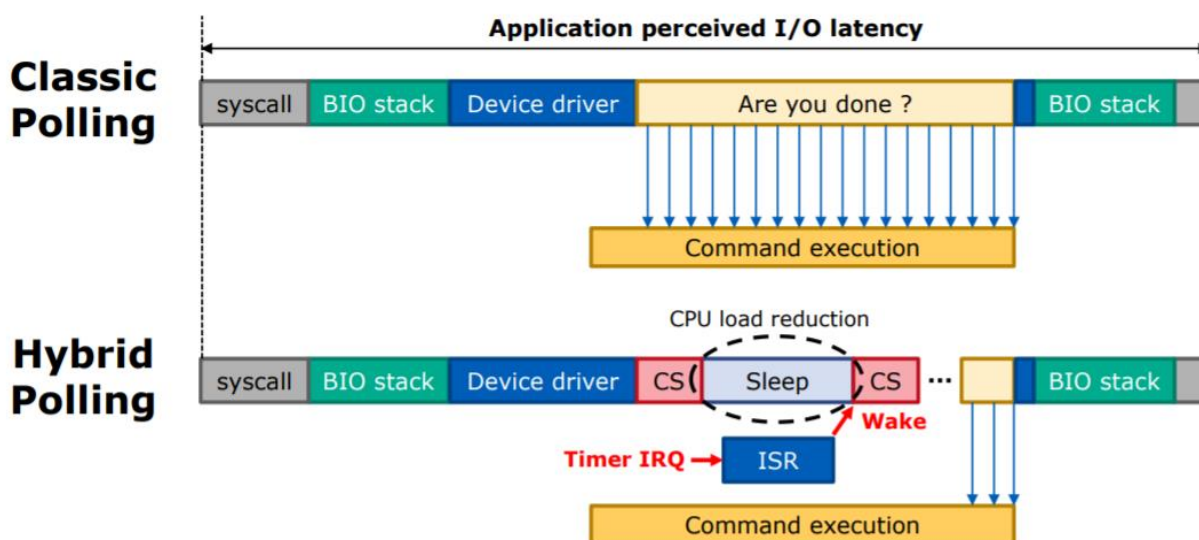


Figure 8: Linux Block I/O Classic and Hybrid Polling latency breakdown. Source [1]

The data in tables and charts compares the I/O latency for a various 4KiB workloads performed using the SPDK bdev vs. Linux block layer I/O model libaio and io_uring with polling mode enabled.

Average and tail latency comparison

Table 10: SPDK bdev vs. Linux Kernel latency comparison (4KiB Random Read, QD=1, runtime=900s)

| Latency metrics (usec) | SPDK Fio BDEV Plugin | Linux Kernel (libaio) | Linux Kernel (io_uring) |
|----------------------------|----------------------|-----------------------|-------------------------|
| Average | 73.38 | 76.99 | 74.34 |
| P90 | 82.43 | 85.50 | 83.46 |
| P99 | 83.46 | 86.53 | 84.48 |
| P99.99 | 148.48 | 154.62 | 164.86 |
| Stdev | 7.03 | 7.04 | 9.02 |
| Average submission latency | 0.15 | 0.98 | 0.00 |
| Average completion latency | 73.23 | 75.93 | 74.30 |

Table 11: SPDK bdev vs. Linux Kernel latency comparison (4KiB Random Write, QD=1, runtime=900s)

| Latency metrics (usec) | SPDK Fio BDEV Plugin | Linux Kernel (Default libaio) | Linux Kernel (io_uring) |
|----------------------------|----------------------|-------------------------------|-------------------------|
| Average | 5.47 | 8.99 | 6.68 |
| P90 | 5.34 | 7.97 | 6.69 |
| P99 | 6.11 | 8.03 | 7.01 |
| P99.99 | 8.77 | 11.46 | 9.02 |
| Stdev | 2.15 | 0.75 | 2.63 |
| Average submission latency | 0.18 | 1.00 | 0.00 |
| Average completion latency | 5.29 | 7.91 | 6.61 |




Figure 9: SPDK bdev vs Linux Kernel Latency comparison (4KiB Random Read)




Figure 10: SPDK bdev vs Linux Kernel Latency comparison (4KiB Random Write)

Linux Kernel libaio Histograms




Figure 11: Linux Kernel (Default libaio) 4KiB Random Read Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)




Figure 12: Linux Kernel (Default libaio) 4KiB Random Write Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)

Linux Kernel io_uring Histograms

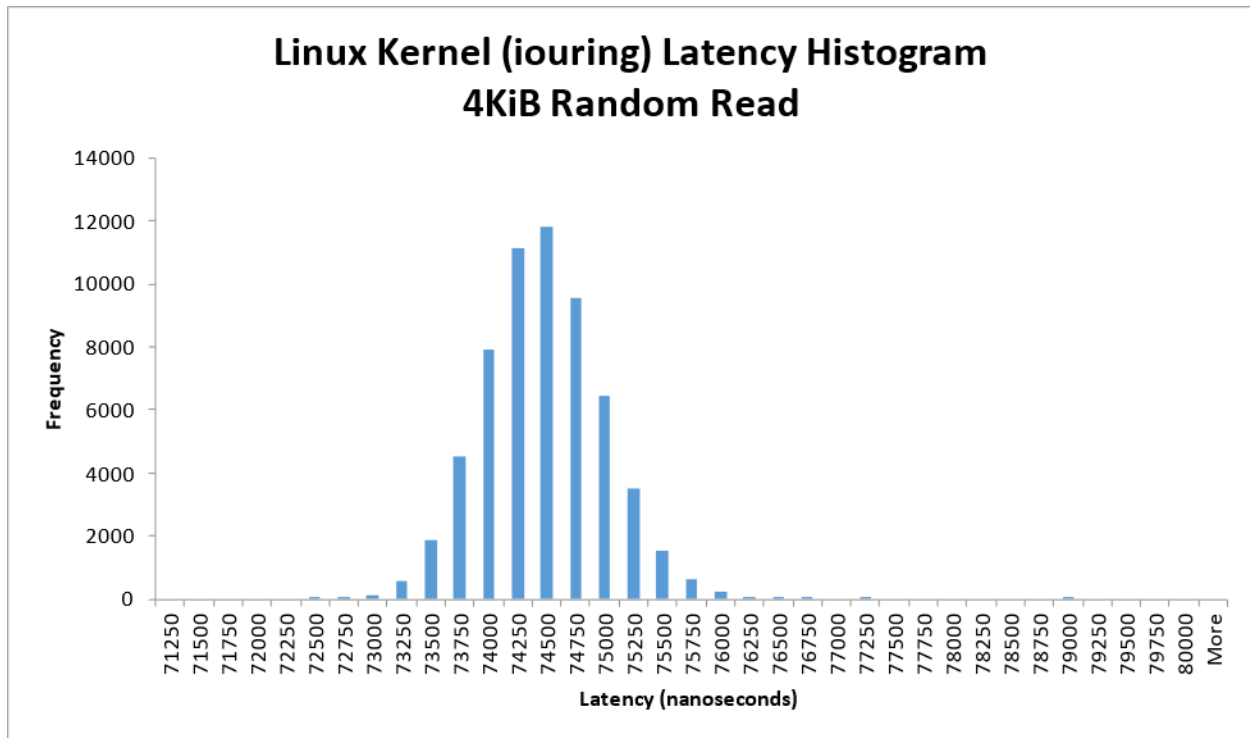


Figure 13: Linux Kernel (io_uring polling) 4KiB Random Read Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)

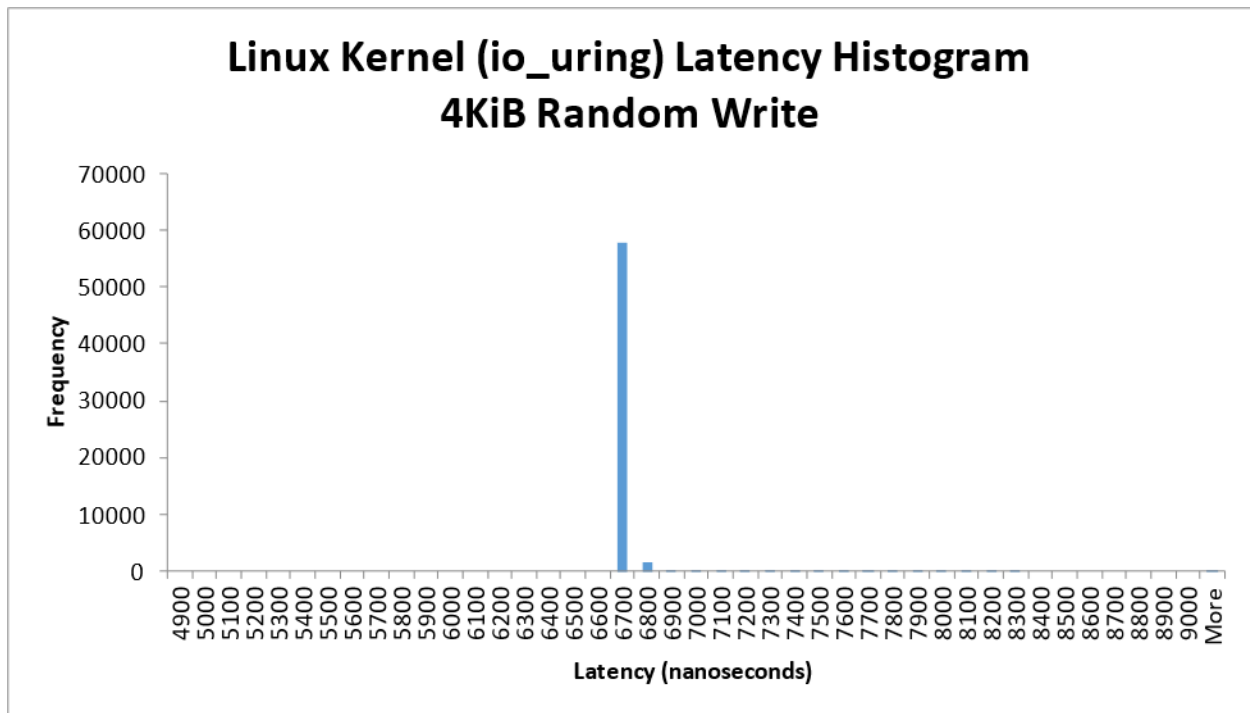


Figure 14: Linux Kernel (io_uring polling) 4KiB Random Write Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)

SPDK fio Bdev Histograms

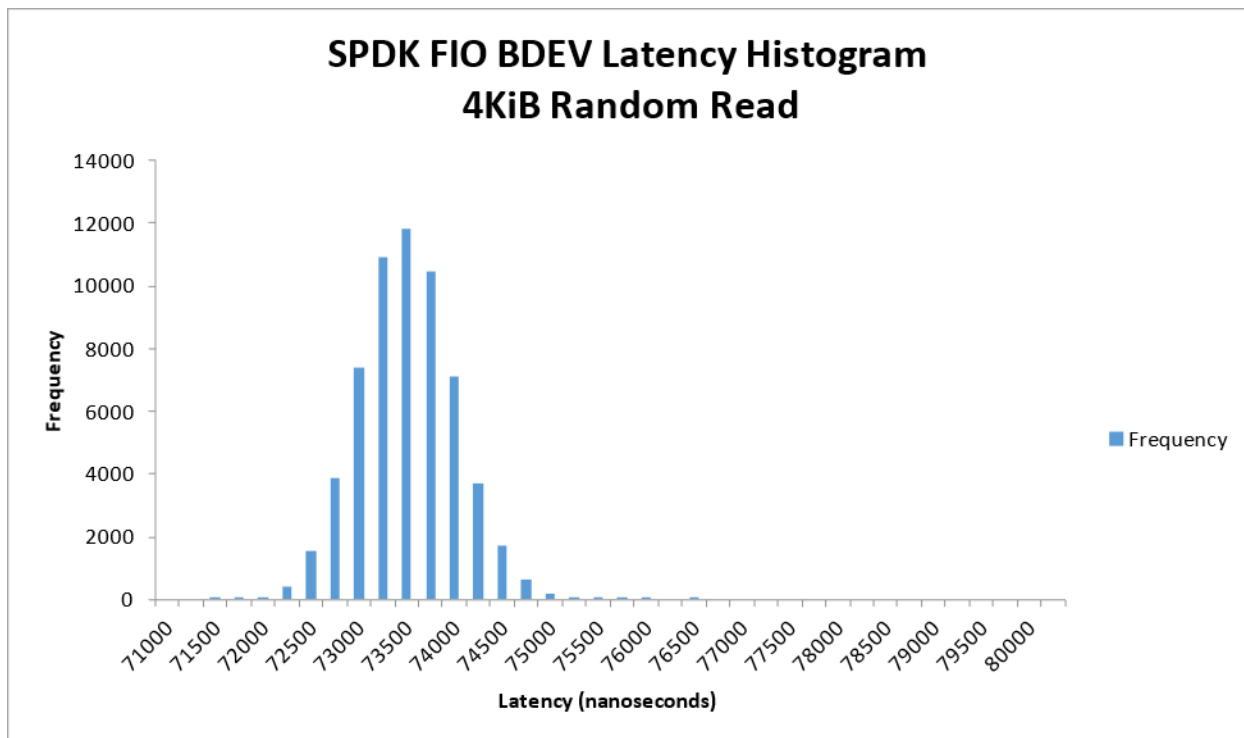


Figure 15: SPDK BDEV NVMe 4KiB Random Read Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)




Figure 16: SPDK BDEV NVMe 4KiB Random Write Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)

Performance vs. increasing Queue Depth

Purpose: Understand the performance in IOPS and average latency of SPDK vs. the Linux io_uring polling and libaio block layer as the queue depth increases by powers of 2 from 1 to 512 for single NVMe SSD and single CPU Core.

Table 12: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 1 NVMe SSD, 1 CPU Core, Numjobs=1)

| | SPDK | | Linux Kernel (Default libaio) | | Linux Kernel (io_uring polling) | |
|-----|---------|------------------|-------------------------------|------------------|---------------------------------|------------------|
| QD | IOPS | Avg. Lat. (usec) | IOPS | Avg. Lat. (usec) | IOPS | Avg. Lat. (usec) |
| 1 | 13587 | 73 | 12946 | 77 | 13409 | 74 |
| 2 | 27109 | 74 | 25807 | 77 | 26756 | 75 |
| 4 | 53966 | 74 | 51435 | 78 | 53186 | 75 |
| 8 | 106805 | 75 | 101685 | 78 | 105059 | 76 |
| 16 | 209022 | 76 | 198421 | 80 | 204769 | 78 |
| 32 | 399457 | 80 | 374053 | 85 | 388814 | 82 |
| 64 | 725692 | 88 | 493829 | 129 | 699578 | 91 |
| 128 | 1172164 | 109 | 492440 | 260 | 1056231 | 121 |
| 256 | 1400104 | 183 | 493029 | 519 | 1319359 | 194 |
| 512 | 1399856 | 366 | 493311 | 1038 | 1486734 | 344 |

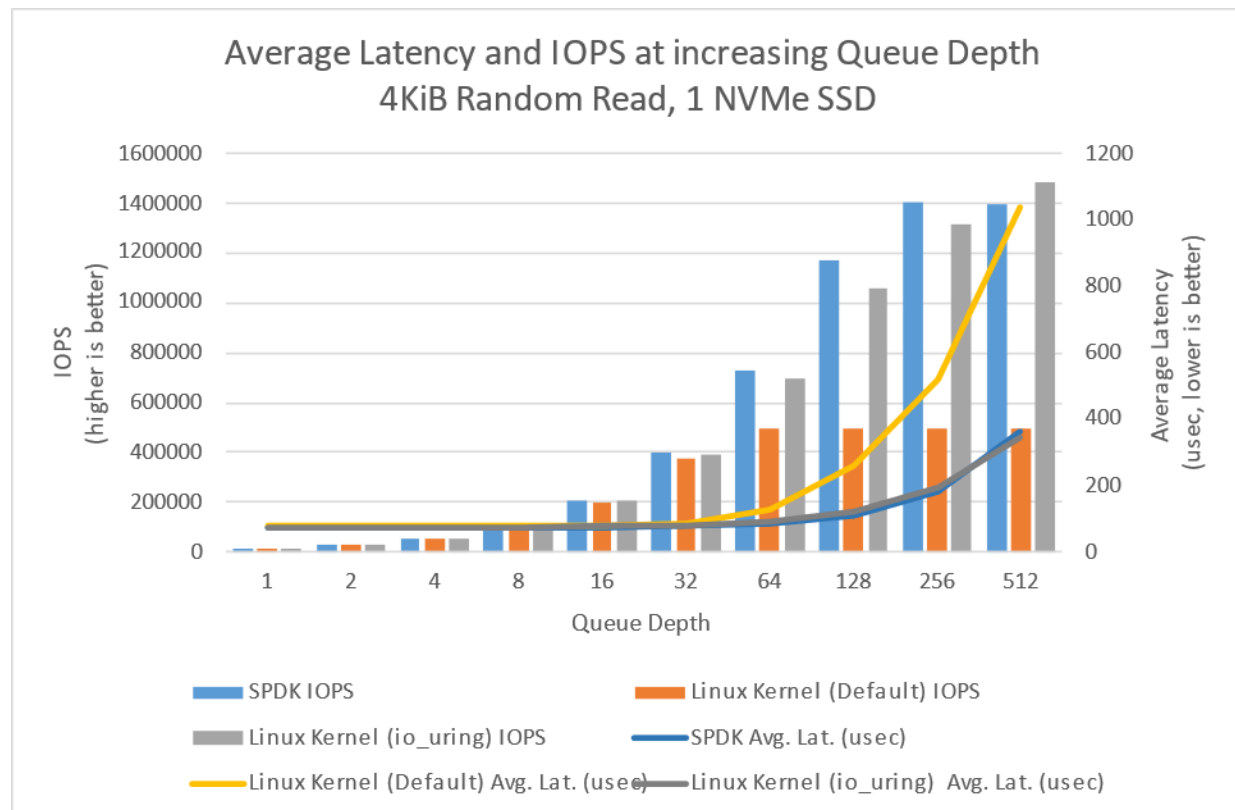


Figure 17: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 1 NVMe SSD, 1 CPU Core, numjobs=1)

Table 13: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 1 NVMe SSD, 1 CPU Core, numjobs=1)

| | SPDK | | Linux Kernel (Default libaio) | | Linux Kernel (io_uring polling) | |
|-----|--------|------------------|-------------------------------|------------------|---------------------------------|------------------|
| QD | IOPS | Avg. Lat. (usec) | IOPS | Avg. Lat. (usec) | IOPS | Avg. Lat. (usec) |
| 1 | 176572 | 5 | 107893 | 9 | 145753 | 7 |
| 2 | 348908 | 6 | 200523 | 10 | 287227 | 7 |
| 4 | 669244 | 6 | 457902 | 9 | 544934 | 7 |
| 8 | 866243 | 9 | 465985 | 17 | 849417 | 9 |
| 16 | 827450 | 19 | 466421 | 34 | 836855 | 19 |
| 32 | 794421 | 41 | 464923 | 69 | 846508 | 38 |
| 64 | 778083 | 82 | 468686 | 136 | 832796 | 77 |
| 128 | 805460 | 159 | 465182 | 275 | 812585 | 158 |
| 256 | 762116 | 336 | 465748 | 549 | 781118 | 328 |
| 512 | 733573 | 698 | 466029 | 1098 | 759843 | 675 |




Figure 18: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 1 NVMe SSD, 1 CPU Core, numjobs=1)

Conclusions

1. Polling hardware for completion instead of relying on interrupts lowers both total latency and its variance.
2. SPDK NVMe Bdev average latency was up to 4.7% and 39.2% lower than Linux Kernel libaio, for Random Read and Random Write workloads respectively.
3. SPDK NVMe Bdev average latency was up to approximately 1.30% and 18.10% lower than Linux Kernel io_uring for Random Read and Random Write workloads respectively.
4. Frequency buckets for 4KiB Random Write at QD=1 workload were so narrow that it was decided to present the results using 100ns as an interval unit for x-axis.
5. For 4KiB Random Read workload all test engines scaled linearly up to QD=32 queue depth. Beyond this value:
 - a. SPDK NVMe Bdev scaling became non-linear and peaked at QD=256, reaching 1.4 million IOPS and saturating NVMe drive.
 - b. Kernel io_uring scaling became non-linear and peaked at QD=512, reaching 1.48 million IOPS and saturating NVMe drive.
 - c. Kernel libaio peaked at QD=64 reaching approximately 493k IOPS. Increasing queue depth did not improve throughput.
6. For 4KiB Random Write workload, the IOPS for the libaio scaled linearly up to QD=4 and the IOPS for the io_uring and SPDK engines scaled linearly up to QD=8. Further increasing the queue depth resulted in minor performance degradation.

Test Case 4: IOPS vs. Latency at different queue depths

Purpose: This test case was performed in order to understand throughput & latency trade-offs with varying queue depth while running SPDK vs. Kernel NVMe block layers.

Results in the table represent performance in IOPS and average latency for the SPDK and Linux Kernel NVMe block layers. We limited both the SPDK and Linux NVMe block layers to use the same number of CPU Cores.

Test Workloads: We use the following Random Read/Write mixes

- 4KiB 100% Random Read
- 4KiB 100% Random Write
- 4KiB Random 70% Read 30% Write

Table 14: SPDK NVMe BDEV Latency Test at different Queue Depths configuration

| Item | Description |
|--|--|
| Test case | Test SPDK NVMe BDEV Latency Test at different Queue Depths |
| Test configuration | fio version: fio-3.28 Number of CPU cores: 12 Number of NVMe SSDs: 22 |
| Linux Kernel io_uring NVMe block device configuration | <pre>echo 0 > /sys/block/nvme0n1/queue echo 0 > /sys/block/nvme0n1/rq_affinity echo 2 > /sys/block/nvme0n1/nomerges echo -1 > /sys/block/nvme0n1/io_poll_delay</pre> |
| fio configuration (common part) | <pre>[global] direct=1 thread=1 time_based=1 norandommap=1 group_reporting=1 rw={randread randwrite randrw} rwmixread={100 0 70} bs=4096 runtime=240 ramp_time=60 numjobs=1</pre> |
| fio configuration (SPDK specific) | <pre>[global] ioengine=spdk_bdev spdk_conf=/tmp/bdev.conf</pre> |

| | |
|---|---|
| | <pre> [filename0] iodepth={2, 4, 8, 16, 32, 64, 128, 256, 512, 1024}* cpus_allowed=0 filename=Nvme0n1 filename=Nvme1n1 [filename1] iodepth={2, 4, 8, 16, 32, 64, 128, 256, 512, 1024}* cpus_allowed=1 filename=Nvme2n1 filename=Nvme3n1 [...] [filename11] iodepth={2, 4, 8, 16, 32, 64, 128, 256, 512, 1024}* cpus_allowed=11 filename=Nvme20n1 filename=Nvme21n1 * - - actual iodepth parameter value used in test; this was multiplied by the number of "filename" objects in job section to achieve desired queue depth value per NVMe SSD (e.g. QD=256 in this case is QD=128 per SSD) </pre> |
| fio configuration (Linux Kernel common) | <pre> [global] ioengine={libaio io_uring} [filename0] iodepth={2, 4, 8, 16, 32, 64, 128, 256, 512, 1024}* cpus_allowed=0 filename=/dev/nvme0n1 filename=/dev/nvme1n1 [filename1] iodepth={2, 4, 8, 16, 32, 64, 128, 256, 512, 1024}* cpus_allowed=1 filename=/dev/nvme2n1 filename=/dev/nvme3n1 [...] [filename11] iodepth={2, 4, 8, 16, 32, 64, 128, 256, 512, 1024}* cpus_allowed=11 filename=/dev/nvme20n1 filename=/dev/nvme21n1 * - - actual iodepth parameter value used in test; this was multiplied by the number of "filename" objects in job section to achieve desired queue depth value per NVMe SSD (e.g. QD=256 in this case is QD=128 per SSD) </pre> |
| fio configuration (Linux Kernel io_uring specific) | <pre> [global] fixedbufs=1 hipri=1 registerfiles=1 sqthread_poll=1 </pre> |

4KiB Random Read Results

Table 15: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 22 NVMe SSDs, 12 CPU Cores)

| | SPDK | | Linux Kernel (Default libaio) | | Linux Kernel (io_uring polling) | |
|-----|--------------------|---------------------|----------------------------------|---------------------|------------------------------------|---------------------|
| QD | IOPS (millions) | Avg. Lat. (usec) | IOPS (millions) | Avg. Lat. (usec) | IOPS (millions) | Avg. Lat. (usec) |
| 1 | 0.30 | 74 | 0.28 | 77 | 0.30 | 74 |
| 2 | 0.60 | 74 | 0.57 | 77 | 0.59 | 74 |
| 4 | 1.19 | 74 | 1.12 | 78 | 1.17 | 75 |
| 8 | 2.35 | 75 | 2.21 | 79 | 2.31 | 76 |
| 16 | 4.60 | 76 | 4.20 | 84 | 4.50 | 78 |
| 32 | 8.78 | 80 | 5.40 | 130 | 8.45 | 83 |
| 64 | 15.85 | 88 | 5.65 | 249 | 13.91 | 101 |
| 128 | 25.31 | 110 | 5.67 | 497 | 18.62 | 151 |
| 256 | 29.75 | 186 | 5.64 | 998 | 22.98 | 245 |
| 512 | 28.91 | 384 | 5.65 | 1994 | 24.18 | 466 |




Figure 19: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 22 NVMe SSDs, 12 CPU Cores)

4KiB Random Write Results

Table 16: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 22 NVMe SSDs, 12 CPU Cores)

| QD | SPDK | | Linux Kernel (Default libaio) | | Linux Kernel (io_uring polling) | |
|-----|--------------------|---------------------|----------------------------------|---------------------|------------------------------------|---------------------|
| | IOPS (millions) | Avg. Lat. (usec) | IOPS (millions) | Avg. Lat. (usec) | IOPS (millions) | Avg. Lat. (usec) |
| 1 | 3.64 | 6 | 2.11 | 10 | 3.21 | 7 |
| 2 | 7.30 | 6 | 4.38 | 10 | 6.50 | 7 |
| 4 | 13.70 | 6 | 5.45 | 16 | 11.12 | 8 |
| 8 | 17.58 | 9 | 5.46 | 32 | 14.74 | 12 |
| 16 | 17.15 | 20 | 5.44 | 64 | 16.45 | 21 |
| 32 | 16.98 | 40 | 5.45 | 129 | 16.92 | 41 |
| 64 | 16.85 | 83 | 5.45 | 258 | 16.74 | 84 |
| 128 | 16.40 | 171 | 5.46 | 516 | 16.62 | 169 |
| 256 | 16.10 | 349 | 5.43 | 1036 | 16.28 | 346 |
| 512 | 15.87 | 708 | 5.36 | 2103 | 15.91 | 708 |




Figure 20: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 22 NVMe SSDs, 12 CPU Cores)

4KiB Random 70%/30% Read/Write Results

Table 17: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB 70/30 Random Read/Write, 22 NVMe SSDs, 12 CPU Cores)

| QD | SPDK | | Linux Kernel (Default libaio) | | Linux Kernel (io_uring polling) | |
|-----|-----------------|------------------|-------------------------------|------------------|---------------------------------|------------------|
| | IOPS (millions) | Avg. Lat. (usec) | IOPS (millions) | Avg. Lat. (usec) | IOPS (millions) | Avg. Lat. (usec) |
| 1 | 0.46 | 16 | 0.44 | 17 | 0.46 | 16 |
| 2 | 0.91 | 16 | 0.87 | 17 | 0.93 | 16 |
| 4 | 1.78 | 16 | 1.69 | 17 | 1.78 | 16 |
| 8 | 3.40 | 17 | 3.20 | 18 | 3.35 | 17 |
| 16 | 6.17 | 19 | 4.90 | 24 | 5.97 | 20 |
| 32 | 10.14 | 23 | 5.39 | 43 | 8.58 | 27 |
| 64 | 14.89 | 31 | 5.53 | 85 | 11.16 | 42 |
| 128 | 18.27 | 51 | 5.57 | 168 | 13.47 | 70 |
| 256 | 19.40 | 288 | 5.55 | 1014 | 14.94 | 377 |
| 512 | 19.37 | 579 | 5.55 | 2028 | 17.27 | 652 |




Figure 21: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB 70/30 Random Read/Write, 22 NVMe SSDs, 12 CPU Cores)

Conclusions

1. SPDK NVMe BDEV fio plugin reached up to around 29.75 million IOPS for Random Read workload at Queue Depth = 256. This is close to 30 million IOPS result measured in Test Case 2 - I/O Cores Scaling using Bdevperf.
2. For the 4KiB Random Write workload SPDK NVMe BDEV Fio plugin and Kernel io_uring had similar performance. We observed noticeable performance drop for both I/O engines when increasing queue depth, after reaching peak performance.
3. SPDK NVMe BDEV fio plugin reached up to 19.4 million IOPS for Random Read/Write workload at Queue Depth = 256, which is lower than the 22.6 million IOPS we measured in Test Case 2 - I/O Cores Scaling using Bdevperf.
4. The results for the Random Write workload exceeded what the 22NVMe SSDs are capable of (around 7.7M IOPS). This is probably due to imperfect preconditioning process, which wears off over time. However, these results were repeatable and still show SPDK's high scalability with increase in the I/O requests.
5. The Kernel libaio IO engine achieved maximum performance of up to 5.67M IOPS with 12 CPU cores and was unable to saturate platforms NVMe disks throughput. Peak performance was reached at QD=128 for Random Read and Random Read/Write workloads and at QD=8 for Random Write workload. Beyond these queue depth values there was no IOPS improvement, but the latency increased.
6. The Kernel io_uring engine reached a peak performance of 24.18 million IOPS at Queue Depth = 512 for Random Read workload, 16.92 million at QD = 32 for Random Write and 17.72 million at QD = 512 for Random Read/Write workload. However, when we looked at htop we noticed that io_uring was using 24 CPU cores; When we configured the sqthread_poll parameter to eliminate system calls, io_uring starts a special kernel thread that polls the shared submission queue for I/O submitted by the fio thread. Therefore, in terms of CPU efficiency we measured up to 1.01M IOPS/Core for io_uring vs up to about 2.48M IOPS/Core for the SPDK NVMe bdev fio plugin. The Submission Queue Polling blog provides more information about how to eliminate system calls with io_uring.

Summary

1. SPDK NVMe BDEV Block Layer using SPDK Bdevperf benchmarking tool can deliver up to 7.8 million IOPS on a single Intel® Xeon® Gold 6348 CPU Core with Turbo Boost enabled.
2. The SPDK NVMe BDEV IOPS scale linearly with addition of CPU cores. We demonstrated up to 29.9 million IOPS on just 5 CPU cores (Intel® Xeon® Gold 6348 with Turbo Boost enabled).
3. The SPDK NVMe BDEV has lower QD=1 latency than the Linux Kernel NVMe block driver for small (4KiB) blocks.
 - a. SPDK BDEV latency was 4.69% and 39.17% lower than Linux Kernel libaio latency for Random Read and Random Write workloads.
 - b. SPDK BDEV latency was about 1.30% lower than Linux Kernel io_uring latency for Random Read workload and 18.10% lower for Random Write workload.
4. SPDK NVMe Bdev Fio reaches up to 29.75 million IOPS with an average latency of around 186 usec while using 12 CPU cores at queue depth of 256. With the same fio workloads Kernel io_uring and Kernel libaio reach up to 22.98 million (using 24 cores: 12 for fio and 12 for submission queue polling) and 5.64 million IOPS respectively.

List of tables

| | |
|--|----|
| Table 1: Hardware setup configuration | 4 |
| Table 2: Test setup BIOS settings | 5 |
| Table 3: Test System NVMe storage setup | 5 |
| Table 4: SPDK NVMe BDEV IOPS Test configuration..... | 9 |
| Table 5: IOPS/Core performance; SPDK fio bdev plugin vs SPDK bdevperf (Blocksize=4KiB, 1 CPU Core)..... | 12 |
| Table 6: SPDK NVMe Bdev vs SPDK NVMe PMD IOPS/Core (Blocksize=4KiB, 1 CPU Core) | 12 |
| Table 7: SPDK NVMe BDEV I/O Cores Scalability Test | 14 |
| Table 8: SPDK NVMe BDEV I/O Cores Scalability Test (4KiB 100% Random Read IOPS at QD=192; 4KiB 100% Random Write IOPS at QD=32; 4KiB 70/30 Random Read/Write IOPS at QD=192) | 15 |
| Table 9: SPDK NVMe BDEV Latency Test | 17 |
| Table 10: SPDK bdev vs. Linux Kernel latency comparison (4KiB Random Read, QD=1, runtime=900s) | 19 |
| Table 11: SPDK bdev vs. Linux Kernel latency comparison (4KiB Random Write, QD=1, runtime=900s) | 19 |
| Table 12: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 1 NVMe SSD, 1 CPU Core, Numjobs=1)..... | 24 |
| Table 13: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 1 NVMe SSD, 1 CPU Core, numjobs=1) | 25 |
| Table 14: SPDK NVMe BDEV Latency Test at different Queue Depths configuration..... | 27 |
| Table 15: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 22 NVMe SSDs, 12 CPU Cores)..... | 29 |
| Table 16: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 22 NVMe SSDs, 12 CPU Cores) | 30 |
| Table 17: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB 70/30 Random Read/Write, 22 NVMe SSDs, 12 CPU Cores)..... | 31 |

List of figures

| | |
|--|----|
| Figure 1 : Example NVMe bdev configuration file..... | 7 |
| Figure 2: Example SPDK Fio BDEV configuration file | 8 |
| Figure 3: SPDK NVMe BDEV IOPS scalability with addition of SSDs (4KiB Random Read, 1CPU Core, QD=192, using bdevperf tool) | 10 |
| Figure 4: SPDK NVMe BDEV IOPS scalability with addition of SSDs (4KiB Random Write, 1CPU Core, QD=32, using bdevperf tool) | 11 |
| Figure 5: SPDK NVMe BDEV IOPS scalability with addition of SSDs (4KiB 70/30 Random Read/Write, 1CPU Core, QD=192, using bdevperf tool) | 11 |
| Figure 6: SPDK NVMe BDEV I/O Cores Scalability (4KiB 100% Random Read IOPS at QD=192; 4KiB 100% Random Write IOPS at QD=32; 4KiB 70/30 Random Read/Write IOPS at QD=192)..... | 15 |
| Figure 7: Linux Block Layer I/O Optimization with Polling. Source [1] | 18 |
| Figure 8: Linux Block I/O Classic and Hybrid Polling latency breakdown. Source [1] | 19 |
| Figure 9: SPDK bdev vs Linux Kernel Latency comparison (4KiB Random Read)..... | 20 |
| Figure 10: SPDK bdev vs Linux Kernel Latency comparison (4KiB Random Write) | 20 |
| Figure 11: Linux Kernel (Default libaio) 4KiB Random Read Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec) | 21 |
| Figure 12: Linux Kernel (Default libaio) 4KiB Random Write Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec) | 21 |
| Figure 13: Linux Kernel (io_uring polling) 4KiB Random Read Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec) | 22 |
| Figure 14: Linux Kernel (io_uring polling) 4KiB Random Write Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec) | 22 |
| Figure 15: SPDK BDEV NVMe 4KiB Random Read Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)..... | 23 |
| Figure 16: SPDK BDEV NVMe 4KiB Random Write Average Latency Histogram (QD=1, Runtime=900s, fio, sampling interval = 15msec)..... | 23 |
| Figure 17: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 1 NVMe SSD, 1 CPU Core, numjobs=1) | 24 |
| Figure 18: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 1 NVMe SSD, 1 CPU Core, numjobs=1) | 25 |
| Figure 19: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Read, 22 NVMe SSDs, 12 CPU Cores)..... | 29 |
| Figure 20: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB Random Write, 22 NVMe SSDs, 12 CPU Cores) | 30 |
| Figure 21: Performance at increasing Queue Depth; SPDK NVMe BDEV vs Linux Default libaio vs Linux io_uring polling (4KiB 70/30 Random Read/Write, 22 NVMe SSDs, 12 CPU Cores)..... | 31 |

References

[1] Damien Le Moal, "I/O Latency Optimization with Polling", Vault – Linux Storage and Filesystem Conference – 2017, May 22nd, 2017.

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at [www.Intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates.

Your costs and results may vary.

No product or component can be absolutely secure.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.