

SPDK NVMe-oF TCP (Target & Initiator) Performance Report Release 23.01

Mellanox ConnectX-5 version

Testing Date: March 2023

Performed by:

Karol Latecki (karol.latecki@intel.com)

Jaroslav Chachulski (jaroslawx.chachulski@intel.com)

Acknowledgments:

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

Krzysztof Karas (krzysztof.karas@intel.com)

Contents

Contents	2
Audience and Purpose.....	3
Test Setup	4
Target Configuration.....	4
Initiator 1 Configuration	5
Initiator 2 Configuration	5
BIOS settings	6
TCP configuration	6
Introduction to SPDK NVMe-oF (Target & Initiator)	7
Test Case 1: SPDK NVMe-oF TCP Target I/O core scaling	9
4KiB Random Read Results.....	12
4KiB Random Write Results	13
4KiB Random Read-Write Results	14
Large Sequential I/O Performance	15
Conclusions	18
Test Case 2: SPDK NVMe-oF TCP Initiator I/O core scaling	19
4KiB Random Read Results.....	21
4KiB Random Write Results	22
4KiB Random Read-Write Results	23
Conclusions	24
Test Case 3: Linux Kernel vs. SPDK NVMe-oF TCP Latency	25
SPDK vs Kernel NVMe-oF Target Results	28
SPDK vs Kernel NVMe-oF TCP Initiator Results	29
SPDK vs Kernel NVMe-oF Kernel + Initiator Results	30
Conclusions	31
Test Case 4: NVMe-oF Performance with increasing # of connections	32
4KiB Random Read Results.....	34
4KiB Random Write Results	35
4KiB Random Read-Write Results	36
Low Connections Results	37
Conclusions	38
Summary	39
List of Figures	40
List of Tables	41
Appendix A – SPDK NVMe-oF Target Dynamic Scheduler	43
Appendix B – Test Case 1 SPDK NVMe-oF Initiator bdev configuration	46
Appendix C – Test Case 2 SPDK NVMe-oF Initiator bdev configuration	49
Appendix D – Test Case 3 SPDK NVMe-oF Initiator bdev configuration	53
Appendix E – Kernel NVMe-oF TCP Target configuration	53

Audience and Purpose

This report is intended for people who are interested in evaluating SPDK NVMe-oF (Target & Initiator) performance. This report contains SPDK NVMe-oF Target and Initiator performance characteristics and provides comparison data between SPDK and its Kernel NVMe-oF Target and Initiator counterparts. This report covers the TCP transport only.

The purpose of report is not to imply a single “correct” approach, but rather to provide a baseline of well-tested configurations and procedures that produce repeatable results. This report can also be viewed as information regarding best known method/practice when performance testing SPDK NVMe-oF Target and Initiator components.

Test Setup

Target Configuration

Table 1: Hardware setup configuration – Target system

Item	Description
Server Platform	SuperMicro® 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: 0xd000375
Memory	16 x 32GB SK Hynix HMA84GR7DJR4N-XN, DDR4, 3200MHz Total of 512GB
Operating System	Fedora 37
BIOS	1.4a
Linux kernel version	6.0.18-300.fc37.x86_64 Spectre-meltdown mitigations enabled
SPDK version	SPDK 23.01
Storage	OS: 1x 250GB Crucial CT250MX500SSD1 Storage Target: 14x Kioxia® KCM61VUL3T20 3.2TBs (FW: 0105) (6 on CPU NUMA Node 0, 8 on CPU NUMA Node 1)
NIC	4x 100GbE Mellanox ConnectX-5 MCX516A-CDAT NICs. Both ports connected. 2 NICs per CPU socket.

Initiator 1 Configuration

Table 2: Hardware setup configuration – Initiator system 1

Item	Description
Server Platform	Intel® Server System M50CYP2UR208
CPU	Intel® Xeon® Gold 6348 Processor @ 2.60GHz (42MB Cache) Number of cores 28 per socket, number of threads 56 per socket (Both sockets populated) Microcode: 0xd000375
Memory	16 x 32GB Micron 36ASF4G72PZ-3G2J3, DDR4, 3200MHz Total 512GBs
Operating System	Fedora 37
BIOS	SE5C620.86B.01.01.0007.2210270543
Linux kernel version	6.0.18-300.fc37.x86_64 Spectre-meltdown mitigations enabled
SPDK version	SPDK 23.01
Storage	OS: 1x 250GB Crucial CT250MX500SSD1
NIC	2x 100GbE Mellanox ConnectX-5 Ex MCX516A-CDAT NICs. Single port on each NIC connected to Target server. 1 NIC per CPU socket.

Initiator 2 Configuration

Table 3: Hardware setup configuration – Initiator system 2

Item	Description
Server Platform	Intel® Server System M50CYP2UR208
CPU	Intel® Xeon® Gold 6348 Processor @ 2.60GHz (42MB Cache) Number of cores 28 per socket, number of threads 56 per socket (Both sockets populated) Microcode: 0xd000375
Memory	16 x 32GB Micron 36ASF4G72PZ-3G2J3, DDR4, 3200MHz Total 512GBs
Operating System	Fedora 37
BIOS	SE5C620.86B.01.01.0007.2210270543
Linux kernel version	6.0.18-300.fc37.x86_64 Spectre-meltdown mitigations enabled
SPDK version	SPDK 23.01
Storage	OS: 1x 250GB Crucial CT250MX500SSD1
NIC	2x 100GbE Mellanox ConnectX-5 Ex MCX516A-CDAT NICs. Single port on each NIC connected to Target server. 1 NIC per CPU socket.

BIOS settings

Table 4: Test systems BIOS settings

Item	Description
BIOS (Applied to all 3 systems)	Hyper threading Enabled CPU Power and Performance Policy: <ul style="list-style-type: none"> “Extreme Performance” for Target “Performance” for Initiators CPU C-state No Limit CPU P-state Enabled Enhanced Intel® SpeedStep® Tech Enabled Turbo Boost Enabled

TCP configuration

Note that the SPDK NVMe-oF target and initiator use the Linux Kernel TCP stack. We tuned the Linux Kernel TCP stack for storage workloads over 100 Gbps NIC by settings the following parameters using `sysctl`:

```

net.core.busy_poll = 0
net.core.busy_read = 0
net.core.somaxconn = 4096
net.core.netdev_max_backlog = 8192
net.ipv4.tcp_max_syn_backlog = 16384
net.core.rmem_max = 268435456
net.core.wmem_max= 268435456
net.ipv4.tcp_mem = "268435456 268435456 268435456"
net.ipv4.tcp_rmem = "8192 1048576 33554432"
net.ipv4.tcp_wmem = "8192 1048576 33554432"
net.ipv4.route.flush = 1
vm.overcommit_memory = 1
    
```

Introduction to SPDK NVMe-oF (Target & Initiator)

The NVMe over Fabrics (NVMe-oF) protocol extends the parallelism and efficiencies of the NVMe Express* (NVMe) block protocol over network fabrics such as RDMA (iWARP, RoCE, InfiniBand™), Fibre Channel and TCP. SPDK provides both a user-space NVMe-oF target and initiator that extends the software efficiencies of the rest of the SPDK stack over the network. The SPDK NVMe-oF target uses the SPDK user-space, polled-mode NVMe driver to submit and complete I/O requests to NVMe devices which reduces the software processing overhead. Likewise, it pins connections to CPU cores to avoid synchronization and cache thrashing so that the data for those connections is kept close to the CPU.

The SPDK NVMe-oF target and initiator use the underlying transport layer API which in case of TCP are POSIX sockets. Similar to the SPDK NVMe driver, SPDK provides a user-space, lockless, polled-mode NVMe-oF initiator. The host system uses the initiator to establish a connection and submit I/O requests to an NVMe subsystem within an NVMe-oF target. NVMe subsystems contain namespaces, each of which maps to a single block device exposed via SPDK's bdev layer. SPDK's bdev layer is a block device abstraction layer and general-purpose block storage stack akin to what is found in many operating systems. Using the bdev interface completely decouples the storage media from the front-end protocol used to access storage. Users can build their own virtual bdevs that provide complex storage services and integrate them with the SPDK NVMe-oF target with no additional code changes. There can be many subsystems within an NVMe-oF target and each subsystem may hold many namespaces. Subsystems and namespaces can be configured dynamically via a JSON-RPC interface.

Figure 1 shows a high-level schematic of the systems used for testing in the rest of this report. The set up consists of three systems (two used as initiators and one used as the target). The NVMe-oF target is connected to both initiator systems point-to-point using QSFP28 cables without any switches. The target system has fourteen Kioxia® KCM61VUL3T20 SSDs which were used as block devices for NVMe-oF subsystems and four 100GbE Mellanox ConnectX®-5 NICs connected to provide up to 400GbE of network bandwidth. Each Initiator system has two Mellanox ConnectX®-5 100GbE NICs connected directly to the target without any switch.

One goal of this report was to clarify the advantages and disadvantages inherent to the design of the SPDK NVMe-oF components. These components are written using techniques such as run-to completion, polling, and asynchronous I/O. The report covers four real-world use cases.

For performance benchmarking the fio tool is used with two storage engines:

- 1) Linux Kernel io_uring engine
- 2) SPDK bdev engine

Performance numbers reported are aggregate I/O per second, average latency, and CPU utilization as a percentage for various scenarios. Aggregate I/O per second and average latency data is reported from fio and CPU utilization was collected using sar (systat).

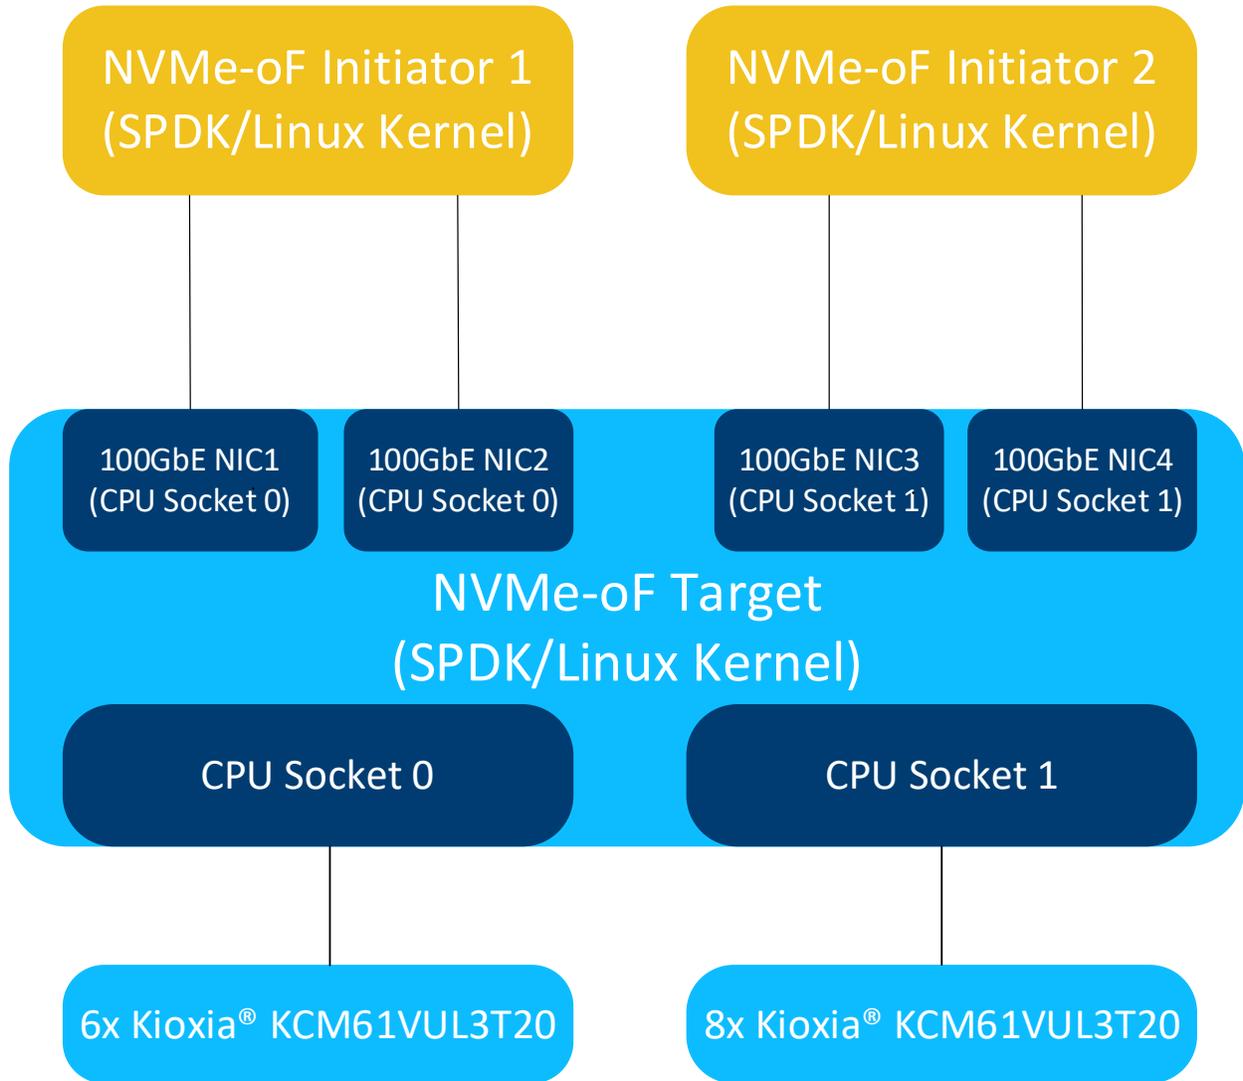


Figure 1: High-Level NVMe-oF TCP performance testing setup

Test Case 1: SPDK NVMe-oF TCP Target I/O core scaling

This test case was performed in order to understand the performance of SPDK TCP NVMe-oF target with I/O core scaling.

The SPDK NVMe-oF TCP target was configured to run with 14 NVMe-oF subsystems. Each NVMe-oF subsystem ran on top of an individual NVMe bdev backed by a single Kioxia KCM61VUL3T20 NVMe drive. Each of the 2 host systems was connected to 7 NVMe-oF subsystems which were exported by the SPDK NVMe-oF Target over 2 x 100GbE NIC. The SPDK bdev fio plugin was used to target 7 NVMe-oF bdevs on each of the host. The SPDK Target was configured to use 1, 4, 8, 12, 16, 24, 32, 40 and 48 CPU cores. We ran the following workloads on each initiator:

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

We scaled the fio jobs using fio parameter numjobs=4 in order to generate more I/O requests. When using the SPDK fio plugin it is important to note the difference between the fio I/O depth parameter and the NVMe device I/O depth because we can configure an fio job to send I/Os to more than one NVMe device and we can also scale the number of fio jobs using the numjobs parameter. The parameter values presented in the table below are actual queue depths used for each of the NVMe devices specified by the filename. These values were calculated in test based on number of fio job sections, numjobs parameter and the number of “filename” targets grouped in each of the fio job sections.

For detailed configuration please refer to the table below. The actual SPDK NVMe-oF configuration was done using JSON-RPC and the table contains the sequence of commands used by spdk/scripts/rpc.py script rather than a configuration file. The SPDK NVMe-oF Initiator (bdev fio_plugin) still uses plain configuration files.

Each workload was run three times at each CPU count and the reported results are the average of the 3 runs. For workloads which need preconditioning, 4KiB Random Read and 4KiB Random 70%/30% Read /Write we ran preconditioning once before running all of the workload to force the NVMe devices into a steady state so that we get consistent results.

Table 5: SPDK NVMe-oF TCP Target Core Scaling test configuration

Item	Description
Test Case	Test SPDK NVMe-oF Target I/O core scaling
SPDK NVMe-oF Target configuration	All the commands below were executed with spdk/scripts/rpc.py script. Enable zero-copy send on Target side before initializing all other subsystems. sock_impl_set_options --impl=posix --enable-zero-copy-send-server (note: zero-copy for Client side is disabled by default) Construct NVMe bdevs: bdev_nvme_attach_controller -t PCIe -b Nvme0 -a 0000:17:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme1 -a 0000:18:00.0

	<pre> bdev_nvme_attach_controller -t PCIe -b Nvme2 -a 0000:65:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme3 -a 0000:66:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme4 -a 0000:67:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme5 -a 0000:68:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme6 -a 0000:98:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme7 -a 0000:99:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme8 -a 0000:9a:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme9 -a 0000:9b:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme10 -a 0000:e3:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme11 -a 0000:e4:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme12 -a 0000:e5:00.0 bdev_nvme_attach_controller -t PCIe -b Nvme13 -a 0000:e6:00.0 Create a TCP transport: nvmf_create_transport -t TCP { "trtype": "TCP", "max_queue_depth": 128, "max_io_qpairs_per_ctrlr": 127, "in_capsule_data_size": 4096, "max_io_size": 131072, "io_unit_size": 131072, "max_aq_depth": 128, "num_shared_buffers": 8192, "buf_cache_size": 32, "dif_insert_or_strip": false, "c2h_success": true, "sock_priority": 0, "abort_timeout_sec": 1 } Create NVMe-oF subsystems and add NVMe bdevs as namespaces: for i in \$(seq 1 14); do nvmf_subsystem_create nqn.2018-09.io.spdk:cnode\${i} -s SPDK00\${i} -a -m 8 nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode\${i} Nvme\${((i-1))}n1 done Add listeners to NVMe-oF Subsystems: i=1 ips=(20.0.0.1 20.0.1.1 10.0.0.1 10.0.1.1) for ip in \${ips[@]}; do for j in \$(seq 1 4); do nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode\${i} -t tcp \ -f ipv4 -s 4420 -a \${ip} ((i++)) done done </pre>
SPDK NVMe-oF Initiator - fio plugin configuration	<p>BDEV.conf: See appendix A.</p> <p>fio.conf [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_conf=/tmp/spdk/bdev.conf thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k</p>

<pre>iodepth={192, 256, 384} time_based=1 numjobs=4 ramp_time=60 runtime=300 [filename0] filename=Nvme0n1 [filename1] filename=Nvme1n1 [filename2] filename=Nvme2n1 [filename3] filename=Nvme3n1 [filename4] filename=Nvme4n1 [filename5] filename=Nvme5n1 [filename6] filename=Nvme6n1</pre>

4KiB Random Read Results

Table 6: SPDK NVMe-oF TCP Target Core Scaling results, Random Read IOPS, QD=384

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	639.9	20.97	8389.7
4 cores	2420.7	79.32	2216.7
8 cores	5835.4	191.21	913.2
12 cores	8326.8	272.85	636.0
16 cores	9958.9	326.33	529.8
24 cores	10663.8	349.43	492.7
32 cores	10958.9	359.10	480.0
40 cores	11127.3	364.62	472.1
48 cores	11163.2	365.80	470.4

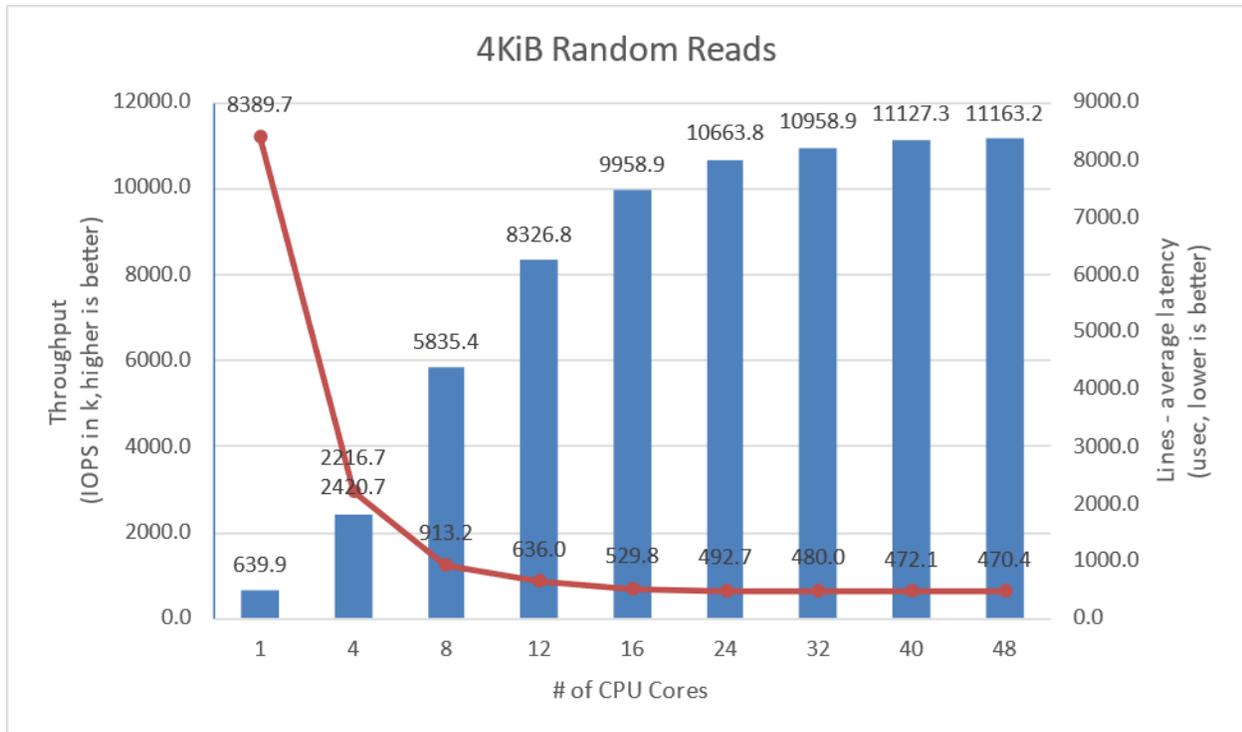


Figure 2: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Read workload at QD = 384

4KiB Random Write Results

Disks were not preconditioned for this test case, which allows for higher IOPS numbers.

Table 7: SPDK NVMe-oF TCP Target Core Scaling results, Random Write IOPS, QD=128

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	387.9	12.71	4632.4
4 cores	1599.2	52.40	1115.0
8 cores	3004.4	98.45	590.7
12 cores	4240.8	138.96	417.8
16 cores	4970.8	162.88	355.1
24 cores	5692.9	186.55	309.7
32 cores	5919.9	193.98	298.0
40 cores	5879.0	192.64	300.3
48 cores	5920.4	194.00	298.0

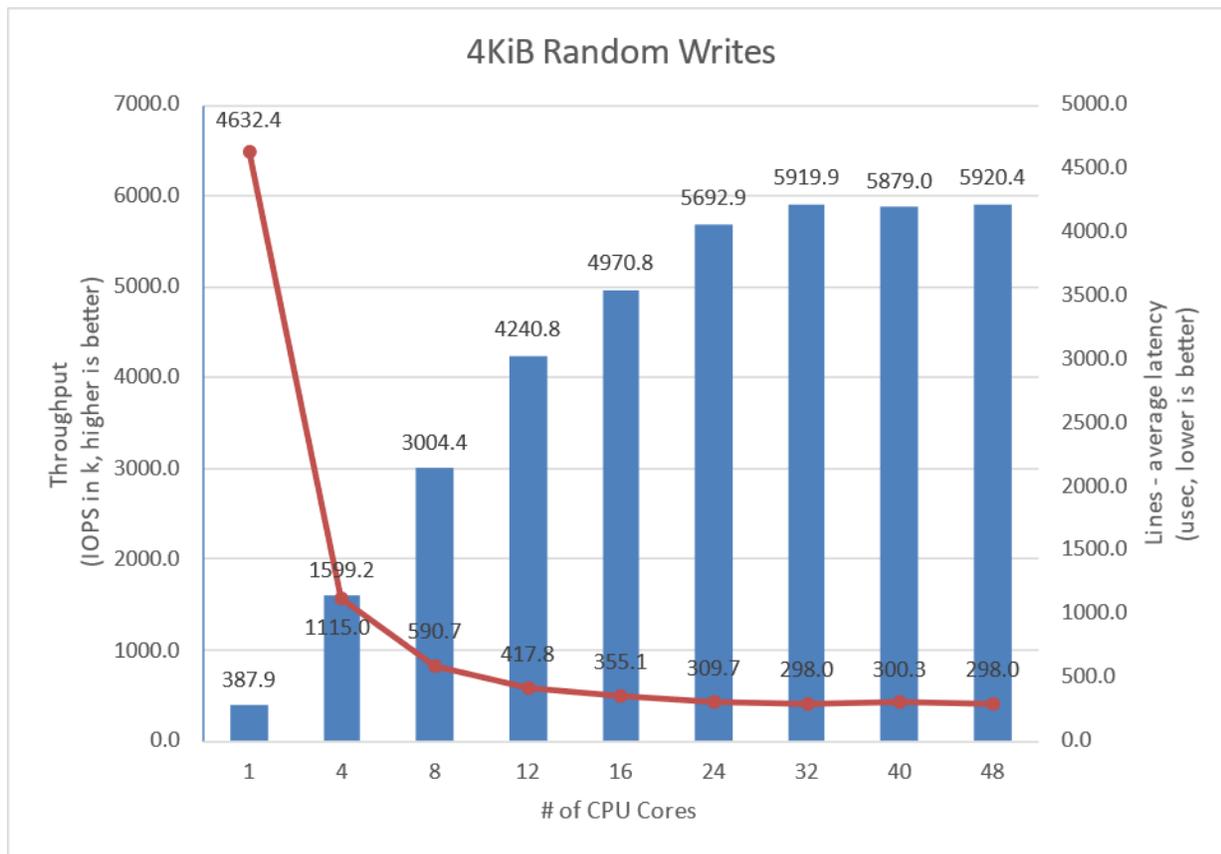


Figure 3: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Write Workload at QD=384

4KiB Random Read-Write Results

Table 8: SPDK NVMe-oF TCP Target Core Scaling results, Random Read/Write 70%/30% IOPS, QD=384

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	451.5	14.79	11903.6
4 cores	1879.9	61.60	2855.6
8 cores	3974.2	130.23	1348.7
12 cores	5823.1	190.81	919.9
16 cores	7544.0	247.20	709.3
24 cores	9145.7	299.69	584.1
32 cores	9955.0	326.21	535.6
40 cores	9975.1	326.86	534.7
48 cores	9895.0	324.24	538.8

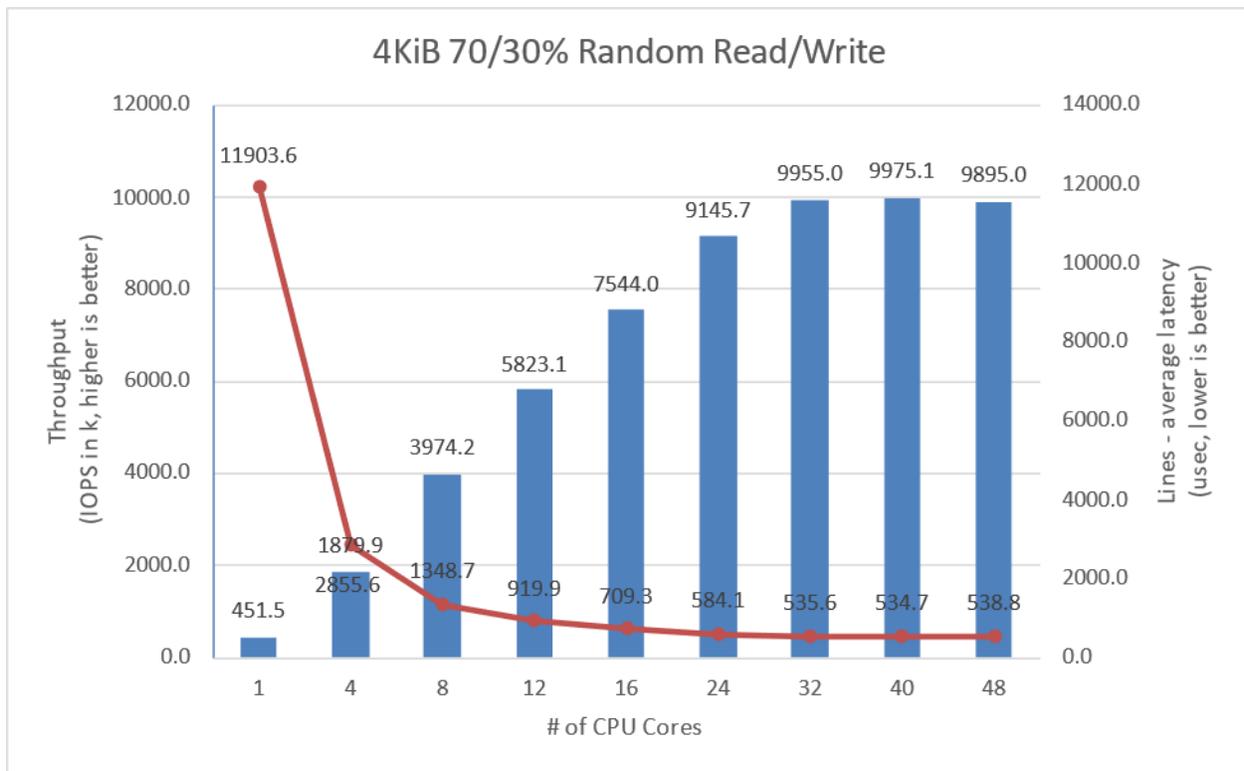


Figure 4: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KiB Random 70/30 Read/Write workload at QD=384

Large Sequential I/O Performance

We measured the performance of large block I/O workloads by performing sequential I/Os of size 128KiBs at queue depth 32. We used iodepth=32 because higher queue depth resulted in negligible bandwidth gain and a significant increase in the latency. The rest of the fio configuration is similar to the 4KiB test case in the previous part of this document.

Table 9: SPDK NVMe-oF TCP Target Core Scaling results, 128KiB Sequential Read IOPS, QD=16

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	149.2	156.41	1513.6
4 cores	358.1	375.52	625.0
8 cores	358.8	376.25	623.8
12 cores	358.8	376.28	623.8

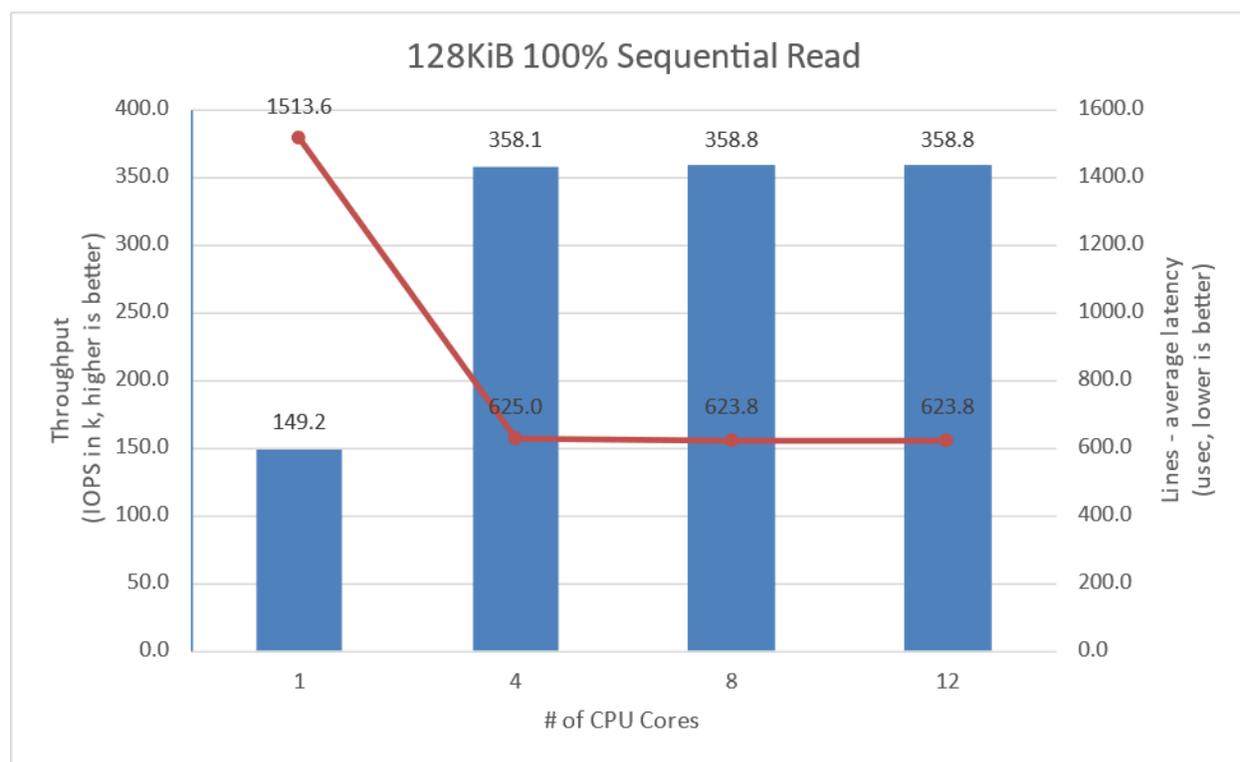


Figure 5: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KiB 100% Sequential Read Workload at QD=16 and initiator fio numjobs=4

Table 10: SPDK NVMe-oF TCP Target Core Scaling results, 128KiB Sequential Write IOPS, QD=32

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	23.4	24.53	19158.4
4 cores	81.5	85.50	5500.7
8 cores	140.4	147.24	3192.5
12 cores	196.8	206.38	2278.5
16 cores	234.3	245.73	1911.2
24 cores	287.5	301.48	1568.8
32 cores	290.1	304.23	1544.7
40 cores	294.2	308.45	1547.0
48 cores	322.4	338.07	1392.2

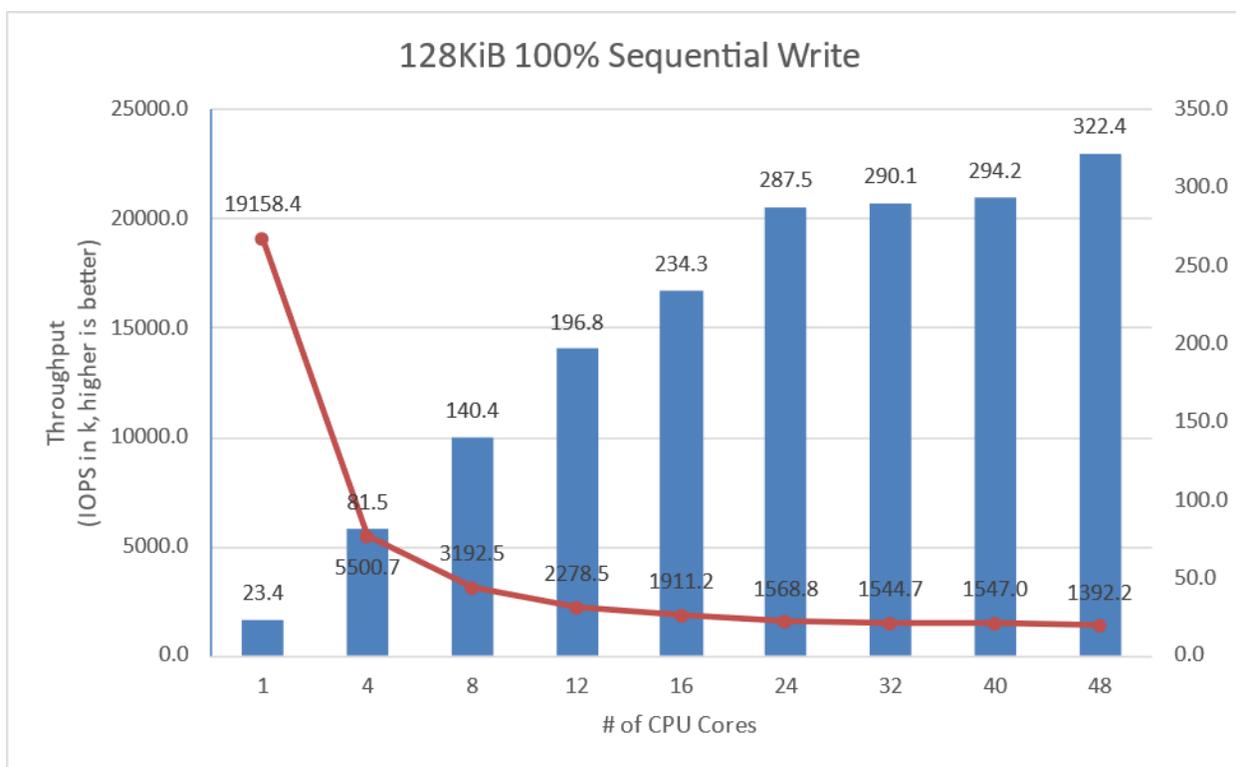


Figure 6: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KiB 100% Sequential Write Workload at QD=32 and Initiator fio numjobs=4

Table 11: SPDK NVMe-oF TCP Target Core Scaling results, 128KiB Sequential 70% Read 30% Write IOPS, QD=32

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	57.2	60.02	7842.1
4 cores	186.2	195.26	2412.6
8 cores	293.0	307.25	1529.1
12 cores	386.0	404.76	1160.4
16 cores	417.2	437.46	1073.3
24 cores	448.3	470.12	998.5
32 cores	470.2	493.04	951.9

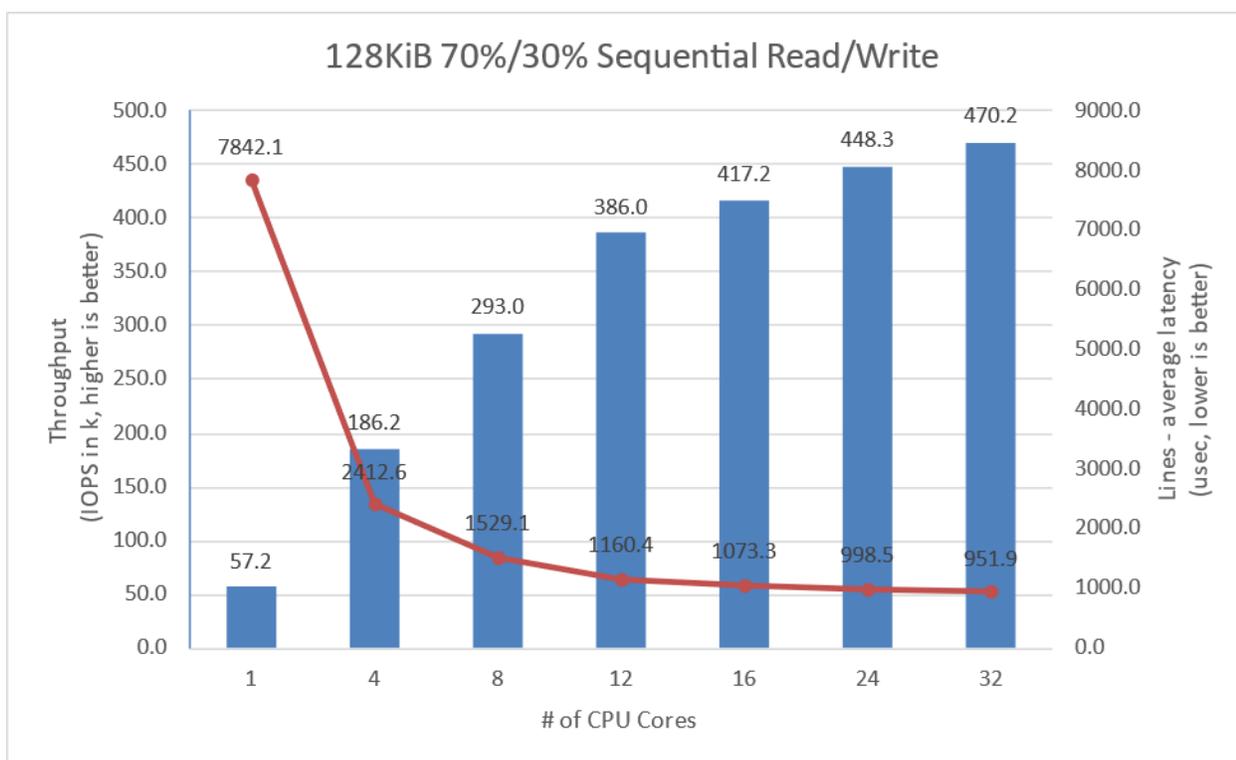


Figure 7: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KiB Sequential 70% Read 30% Write Workload at QD=32 and Initiator fio numjobs=4

Conclusions

1. The SPDK NVMe-oF TCP Target IOPS throughput scales up linearly with addition of CPU cores for 4KiB Random Read workload up to 16 CPU cores, reaching 326 Gbps bandwidth and 9.9 million IOPS. Adding more CPUs to Target configuration results in non-linear performance gains peaking at about 11 million IOPS at 32 CPU cores, reaching 400GbE network link saturation.
2. For the 4KiB Random Write workload 100GbE link is saturated at 8 CPU cores. Performance scales up almost linearly to 16 CPU cores reaching 4.9 million IOPS. Adding more IO cores results in non-linear performance improvement which peaks at 32 CPU cores reaching 200 GbE network link saturation.
3. 4KiB Random Read-Write workload throughput scales linearly up to 24 CPU cores reaching 9.15 million IOPS. Increasing the number of CPU cores beyond 24 results in non-linear performance improvement that peaks at about 10 million IOPS at 32 CPU cores.
4. The best trade-off between CPU efficiency and network saturation is when the Target is configured with 24CPU cores. The performance we achieved with these configurations allows for full (or nearly) saturation of a 400Gbps link between Target and Initiators for all Random Read and Random Read/Write workloads, and 200Gbps link saturation for Random Write workload.
5. For the 4KiB Random Write workload, we did not precondition the drives. If preconditioned, the NVMe drives would max out at about 4.9 million IOPS. Not preconditioning the drives allowed us to artificially increase their throughput and serve more IO requests than usual.
6. The throughput of large block workloads scaled up with addition of CPU cores reaching peak performance at different CPU core counts. For the 128K Sequential Reads workload, the peak throughput of about 376 Gbps was observed at 4 CPU cores. For the 128K Sequential Writes, the throughput scaled linearly to 301 Gbps at 24 cores and reached peak performance of 338 Gbps at 48 CPU cores. For the 128K Sequential 70/30 Read/Write workload, the scaling was linear up to 12 CPU cores, reaching 404 Gbps. Beyond 12 cores the scaling was not linear, reaching beyond 400 Gbps due to bi-directional characteristic of the workload.

Test Case 2: SPDK NVMe-oF TCP Initiator I/O core scaling

This test case was performed in order to understand the performance of SPDK NVMe-oF TCP Initiator as the number of CPU cores is scaled up.

The test setup for this test case is slightly different than the set up described in [introduction chapter](#), as we used just a single SPDK NVMe-oF TCP Initiator. The Initiator was connected to Target server with two 100 Gbps network links.

The SPDK NVMe-oF TCP Target was configured similarly as in test case 1, using 24 cores. We used 24 CPU cores based on results of the previous test case which show that the target can easily serve about 6 million IOPS for all workloads, which is enough IOPS to saturate 200 Gbps network connection.

The SPDK bdev fio plugin was used to target 14 individual NVMe-oF subsystems exported by the Target. The number of CPU threads used by the fio process was managed by setting the fio job sections and numjobs parameter and ranged from 1 to 48 CPUs. For detailed fio job configuration see table below. fio was run with following workloads:

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

It is important to note that fio IO depth parameter values presented in the table below are actual queue depths used for each of the connected subsystem. These values were calculated in test based on number of fio job sections, numjobs parameter and the number of “filename” targets grouped in each of the fio job sections.

Table 12: SPDK NVMe-oF TCP Initiator Core Scaling test configuration

Item	Description
Test Case	Test SPDK NVMe-oF TCP Initiator I/O core scaling
SPDK NVMe-oF Target configuration	Same as in Test Case #1, using 24 CPU cores.
SPDK NVMe-oF Initiator 1 - fio plugin configuration	<p>BDEV.conf See appendix B.</p> <p>fio.conf For 1 CPU initiator configuration: [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_conf=/tmp/spdk/bdev.conf thread=1 group_reporting=1 direct=1</p> <p>norandommap=1 rw=randrw</p>

	<pre> rwmixread={100, 70, 0} bs=4k iodepth={1,32, 64, 128, 192} time_based=1 ramp_time=60 runtime=300 numjobs=1 [filename0] filename=Nvme0n1 filename=Nvme1n1 filename=Nvme2n1 filename=Nvme3n1 filename=Nvme4n1 filename=Nvme5n1 filename=Nvme6n1 filename=Nvme7n1 filename=Nvme8n1 filename=Nvme9n1 filename=Nvme10n1 filename=Nvme11n1 filename=Nvme12n1 filename=Nvme13n1 </pre>
	<pre> fiio.conf For X*4 CPU (up to 48) initiator configuration: [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_conf=/tmp/spdk/bdev.conf thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth={128, 192, 256, 384} time_based=1 ramp_time=60 runtime=300 numjobs=X [filename0] filename=Nvme0n1 filename=Nvme1n1 filename=Nvme2n1 filename=Nvme3n1 [filename1] filename=Nvme4n1 filename=Nvme5n1 filename=Nvme6n1 filename=Nvme7n1 [filename2] filename=Nvme8n1 filename=Nvme9n1 filename=Nvme10n1 [filename3] filename=Nvme11n1 filename=Nvme12n1 filename=Nvme13n1 </pre>

4KiB Random Read Results

Table 13: SPDK NVMe-oF TCP Initiator Core Scaling results, 4KiB Random Read IOPS, QD=256

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	442.9	14.51	8059.5
4 cores	2058.9	67.47	1707.3
8 cores	3204.0	104.99	1101.3
12 cores	4442.8	145.58	792.4
16 cores	5216.5	170.93	669.3
24 cores	5695.9	186.64	615.2
32 cores	5684.3	186.26	620.4
40 cores	5687.5	186.37	621.3
48 cores	5684.5	186.27	622.3

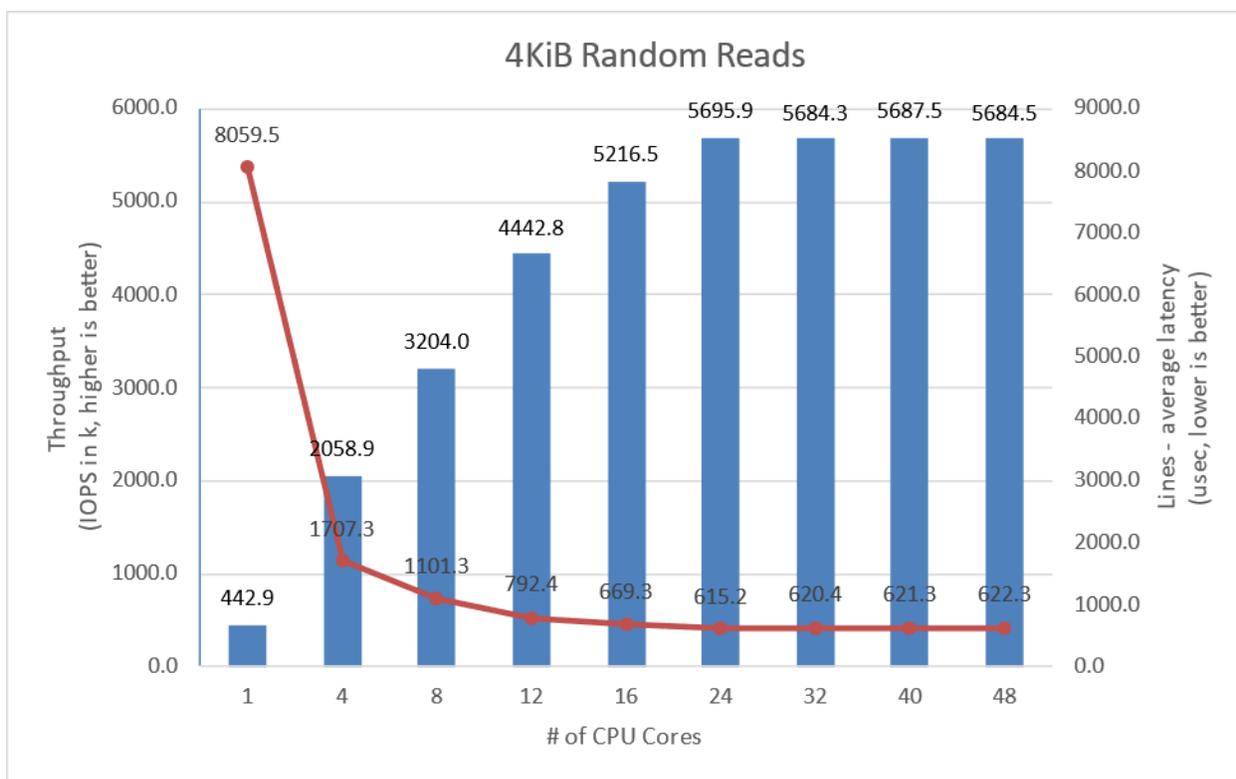


Figure 8: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Read QD=256 workload

4KiB Random Write Results

Table 14: SPDK NVMe-oF TCP Initiator Core Scaling results, 4KiB Random Write IOPS, QD=128

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	623.4	20.43	1986.0
4 cores	2778.3	91.04	566.2
8 cores	3802.1	124.59	453.8
12 cores	3676.3	120.47	476.3
16 cores	3476.1	113.90	507.2
24 cores	3429.5	107.38	516.5
32 cores	3279.6	107.47	542.9
40 cores	3292.2	107.88	538.1
48 cores	3251.3	106.54	551.5

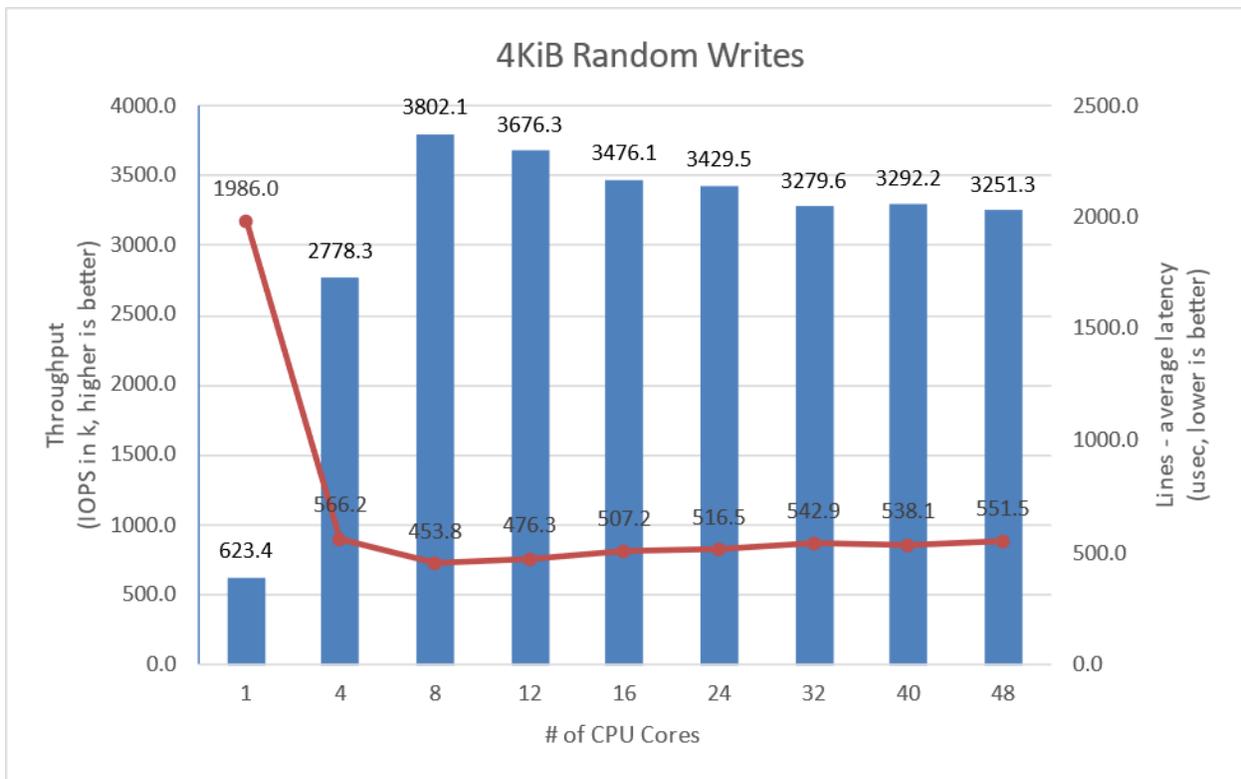


Figure 9: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Write Workload at QD=128

4KiB Random Read-Write Results

Table 15: SPDK NVMe-oF TCP Initiator Core Scaling results, 4KiB Random 70%/30% Read/Write IOPS, QD=256

# of Cores	Throughput (IOPS k)	Bandwidth (Gbps)	Avg. Latency (usec)
1 core	452.6	14.83	7877.2
4 cores	2199.4	72.07	1620.4
8 cores	3724.5	122.04	949.6
12 cores	4736.4	155.20	741.5
16 cores	5767.2	188.98	607.2
24 cores	6414.0	210.17	548.2
32 cores	6339.2	207.72	557.1
40 cores	6422.3	210.45	550.4
48 cores	6138.5	201.15	577.4

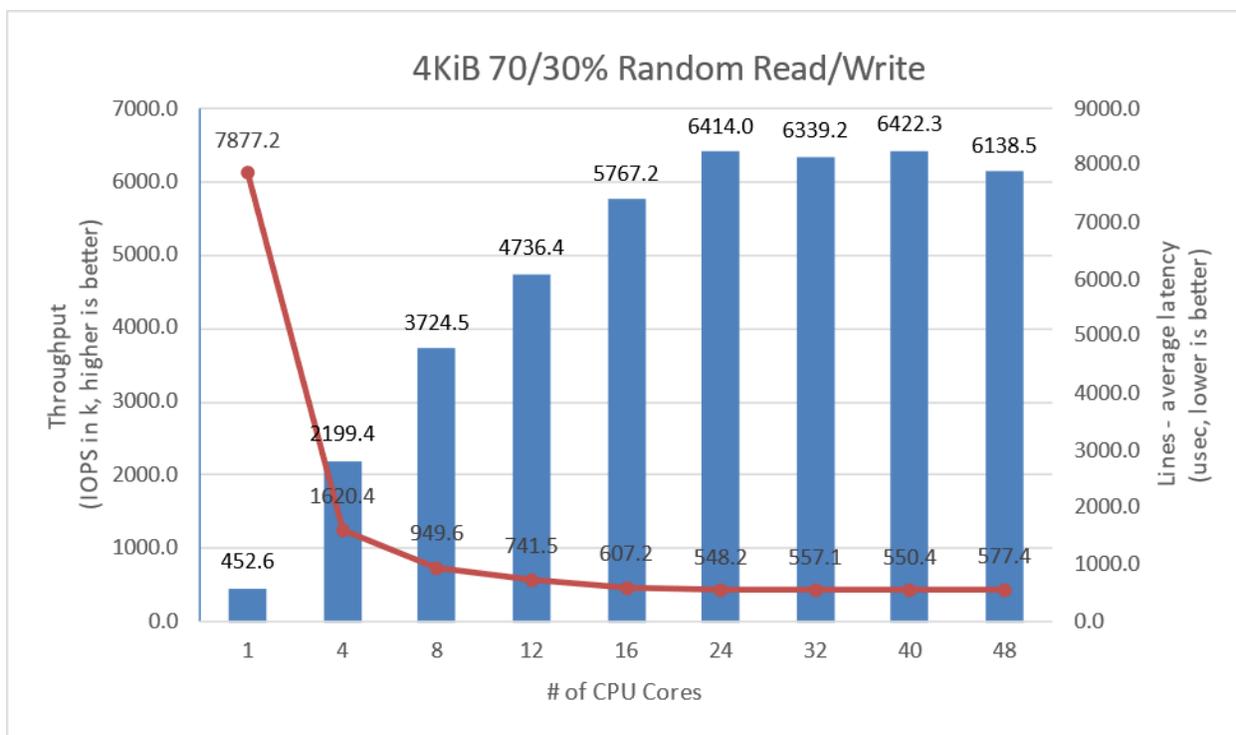


Figure 10: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KiB Random 70% Read 30% Write Workload at QD=256

Conclusions

1. For the 4KiB Random Read workload, the SPDK NVMe-oF TCP Initiator performance scales linearly up to 12 CPU cores. Increasing the number of CPU cores beyond 12 CPU results in non-linear performance improvement, which peaks at 24 CPU cores with 5.7 million IOPS, reaching 200 Gbps network link saturation.
2. In case of 4KiB Random Write workload, performance scaling is non-linear. It rapidly scales to 2.7 million IOPS when transitioning to 4 CPU cores and reaches 100 Gbps network link saturation. Further CPU scaling results in non-linear performance change, and peaks at 8 CPU cores at 3.8 million IOPS. 200 Gbps network link is not saturated.
3. Mixed Random Read-Write workload performance scales linearly up to 16 CPU cores, reaching 5.77 million IOPS. Increasing the number of cores further to 24 allows to reach peak performance of 6.41 million IOPS. Beyond this point increasing number of cores results in noticeable performance degradation.

Test Case 3: Linux Kernel vs. SPDK NVMe-oF TCP Latency

This test case was designed to understand latency characteristics of SPDK NVMe-oF TCP Target and Initiator vs. the Linux Kernel NVMe-oF TCP Target and Initiator implementations on a single NVMe-oF subsystem. The average I/O latency and p99 latency was compared between SPDK NVMe-oF (Target/Initiator) vs. Linux Kernel (Target/Initiator). Both SPDK and Kernel NVMe-oF Targets were configured to run on a single core, with a single NVMe-oF subsystem on top of a *Null Block Device*. The null block device (bdev) was chosen as the backend block device to eliminate the media latency during these tests.

For this test case a Linux Kernel feature called aRFS (Accelerated Receive Flow Steering) was used. For just RFS network packets are forwarded depending on the location of the application receiving the packets. For accelerated version, RFS effect is amplified by enabling hardware support for the packet forwarding process. In aRFS packets are directed to a CPU core which is local to the thread of application receiving the packet. Both RFS and aRFS are available in most Linux distributions but need to be configured before using. Steps for enabling aRFS in this test are described in table below.

Table 16: Linux Kernel vs. SPDK NVMe-oF TCP Latency test configuration

Item	Description
Test Case	Linux Kernel vs. SPDK NVMe-oF Latency
Test configuration	
SPDK NVMe-oF Target configuration	<p>All below commands are executed with <code>spdk/scripts/rpc.py</code> script.</p> <pre> nvmf_create_transport -t TCP (creates TCP transport layer with default values: { "trtype": "TCP", "max_queue_depth": 128, "max_io_qpairs_per_ctrlr": 127, "in_capsule_data_size": 4096, "max_io_size": 131072, "io_unit_size": 131072, "max_aq_depth": 128, "num_shared_buffers": 8192, "buf_cache_size": 32, "dif_insert_or_strip": false, "c2h_success": true, "sock_priority": 0, "abort_timeout_sec": 1 } bdev_null_create Nvme0n1 10240 4096 nvmf_subsystem_create nqn.2018-09.io.spdk:cnode1 -s SPDK001 -a -m 8 nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode1 Nvme0n1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode1 -t tcp -f ipv4 -s 4420 -a 20.0.0.1 </pre>
Kernel NVMe-oF Target configuration	<p>Target configuration file loaded using <code>nvmf-cli</code> tool.</p> <pre> { "ports": [{ "addr": { "adrfam": "ipv4", "traddr": "20.0.0.1", </pre>

	<pre> "trsvcid": "4420", "trtype": "tcp" }, "portid": 1, "referrals": [], "subsystems": ["nqn.2018-09.io.spdk:cnode1"] }], "hosts": [], "subsystems": [{ "allowed_hosts": [], "attr": { "allow_any_host": "1", "version": "1.3" }, "namespaces": [{ "device": { "path": "/dev/nullb0", "uuid": "621e25d2-8334-4c1a-8532-b6454390b8f9" }, "enable": 1, "nsid": 1 }], "nqn": "nqn.2018-09.io.spdk:cnode1" }] } </pre>
fiio configuration	
SPDK NVMe-oF Initiator fio plugin configuration	<p>BDEV.conf See appendix C.</p> <p>fio.conf [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_conf=/tmp/spdk/bdev.conf thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth=1 time_based=1 ramp_time=60 runtime=300</p> <p>[filename0] filename=Nvme0n1</p>
Kernel initiator configuration	<p>Device config Done using nvme-cli tool. modprobe nvme-fabrics nvme connect -n nqn.2018-09.io.spdk:cnode1 -t tcp -a 20.0.0.1 -s 4420</p> <p>fio.conf [global]</p>

	<pre>ioengine=io_uring thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth=1 time_based=1 numjobs=1 ramp_time=60 runtime=300 [filename0] filename=/dev/nvme0n1</pre>
aRFS Configuration	
aRFS Configuration	<p>Enable ntuple feature in the NIC driver and check it's status:</p> <pre>\$ ethtool -K eth3 ntuple on</pre> <p>Disable Linux Kernel IRQ balancer</p> <pre>\$ service irqbalance stop</pre> <p>Ensure that NIC's IRQ affinity is spread across all cores:</p> <pre>\$ set_irq_affinity.sh eth3 \$ show_irq_affinity.sh eth3</pre> <p>(Mellanox utility scripts available on Github.com)</p> <p>Configure the RFS global and per-queue flow table entries. This needs to be done for every NIC interface taking part in the test.</p> <pre>echo 32768 > /proc/sys/net/core/rps_sock_flow_entries for r in /sys/class/net/eth3/queues/rx-*/rps_flow_cnt; do echo 512 > \$r done</pre>

SPDK vs Kernel NVMe-oF Target Results

This following data was collected using the Linux Kernel initiator against both SPDK & Linux Kernel NVMe-oF TCP target.

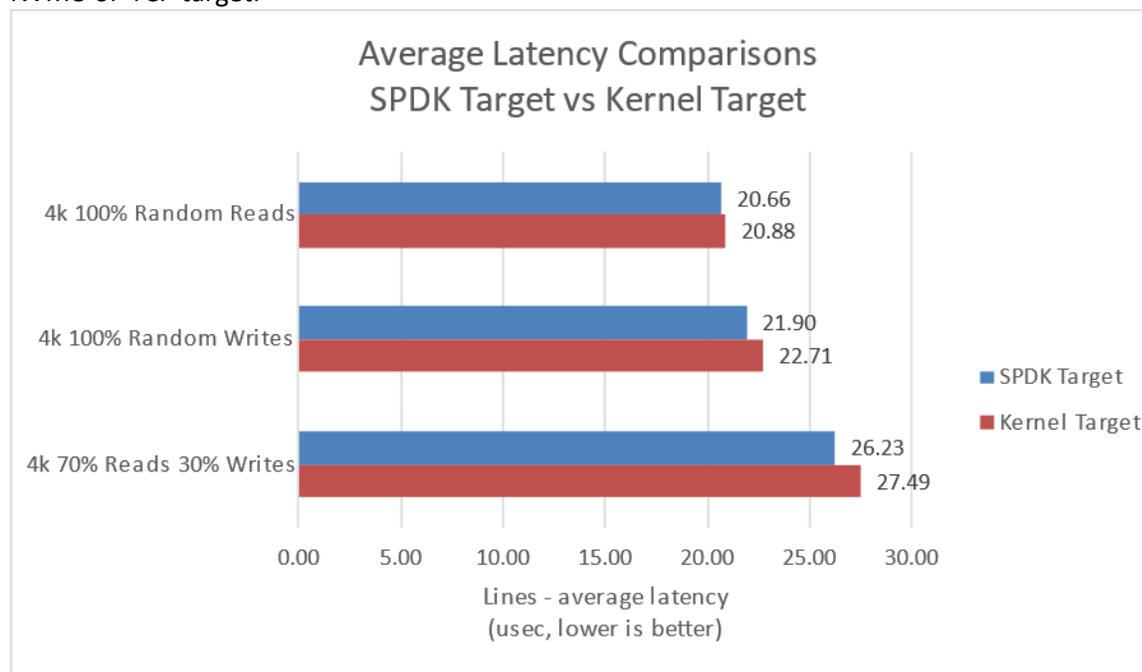


Figure 11: SPDK vs. Kernel NVMe-oF TCP Target Average I/O Latency for various workloads run using the Kernel Initiator

Table 17: SPDK NVMe-oF Target Latency and IOPS at QD=1, Null Block Device

Access Pattern	Avg. Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KiB 100% Random Reads IOPS	20.66	47499	32.8	36.6	88.7	184.7
4KiB 100% Random Writes IOPS	21.90	44964	35.6	49.6	83.5	169.6
4KiB 100% Random 70% Reads 30% Writes IOPS	26.23	37587	60.4	107.5	190.5	222.1

Table 18: Linux Kernel NVMe-oF Target Latency and IOPS at QD=1, Null Block Device

Access Pattern	Avg. Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KiB 100% Random Reads IOPS	20.88	47097	22.1	28.0	58.5	162.8
4KiB 100% Random Writes IOPS	22.71	43351	27.9	31.0	60.2	159.4
4KiB 100% Random 70% Reads 30% Writes IOPS	27.49	35882	58.8	97.0	179.2	215.7

SPDK vs Kernel NVMe-oF TCP Initiator Results

This following data was collected using Kernel & SPDK initiator against an SPDK target.

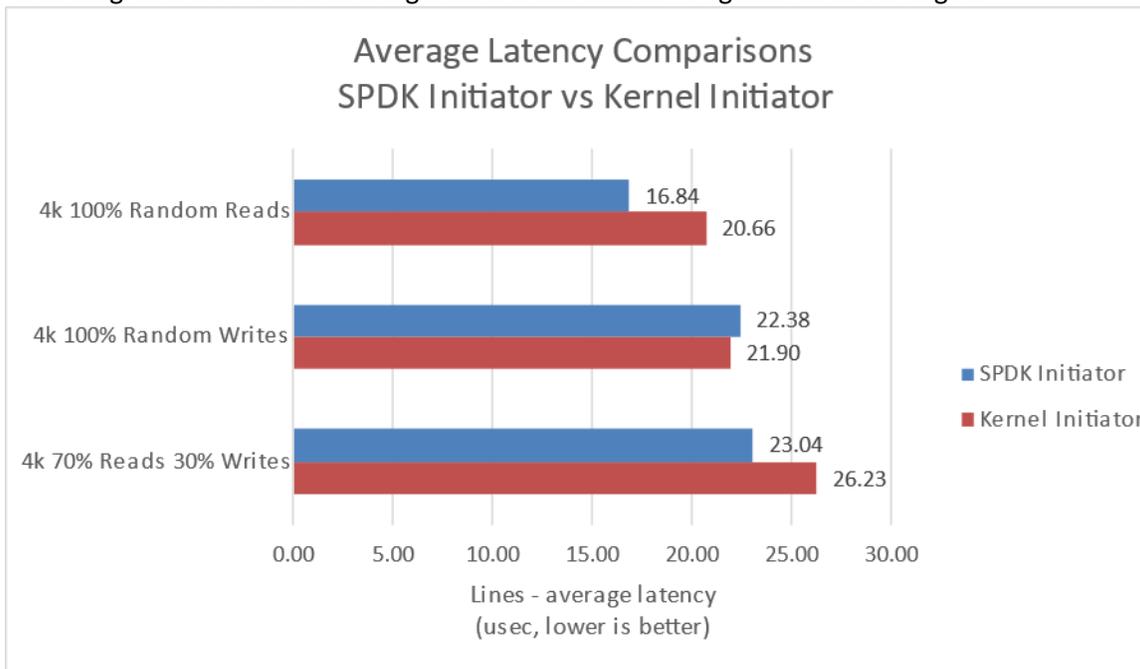


Figure 12: SPDK vs. Kernel NVMe-oF TCP Initiator Average I/O Latency for various workloads against SPDK Target

Table 19: SPDK NVMe-oF Initiator Latency and IOPS at QD=1, Null Block Device

Access Pattern	Avg. Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KiB 100% Random Reads IOPS	16.84	58639	25.0	33.7	51.3	166.2
4KiB 100% Random Writes IOPS	22.38	44256	36.4	56.9	85.8	170.3
4KiB 100% Random 70% Reads 30% Writes IOPS	23.04	42933	54.9	92.0	121.9	174.8

Table 20: Linux Kernel NVMe-oF Initiator Latency and IOPS at QD=1, Null Block Device

Access Pattern	Avg. Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KiB 100% Random Reads IOPS	20.66	47499	32.8	36.6	88.7	184.7
4KiB 100% Random Writes IOPS	21.90	44964	35.6	49.6	83.5	169.6
4KiB 100% Random 70% Reads 30% Writes IOPS	26.23	37587	60.4	107.5	190.5	222.1

SPDK vs Kernel NVMe-oF Kernel + Initiator Results

Following data was collected using SPDK Target with SPDK Initiator and Linux Target with Linux Initiator.

Table 21: SPDK NVMe-oF Latency and IOPS at QD=1, Null Block Device

Access Pattern	Avg. Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KiB 100% Random Reads IOPS	16.84	58639	25.0	33.7	51.3	166.2
4KiB 100% Random Writes IOPS	22.38	44256	36.4	56.9	85.8	170.3
4KiB 100% Random 70% Reads 30% Writes IOPS	23.04	42933	54.9	92.0	121.9	174.8

Table 22: Linux Kernel NVMe-oF Latency and IOPS at QD=1, Null Block Device

Access Pattern	Avg. Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KiB 100% Random Reads IOPS	20.88	47097	22.06	28.03	58.45	162.82
4KiB 100% Random Writes IOPS	22.71	43351	27.95	31.02	60.16	159.40
4KiB 100% Random 70% Reads 30% Writes IOPS	27.49	35882	58.78	96.97	179.18	215.72

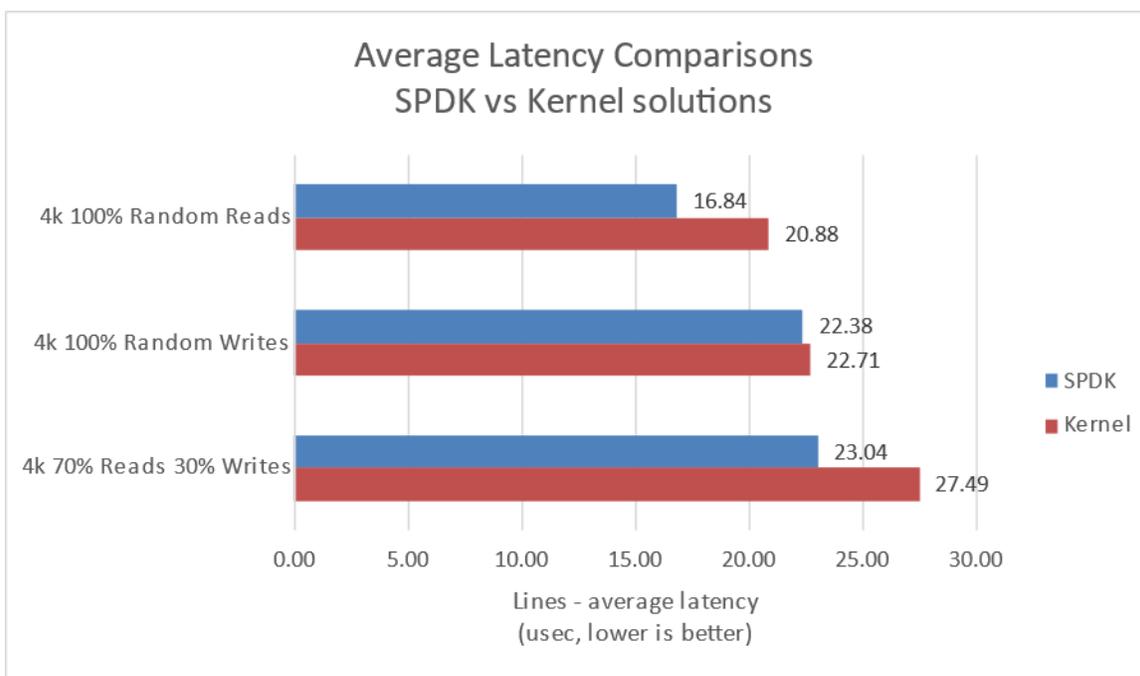


Figure 13: SPDK vs. Kernel NVMe-oF TCP solutions Average I/O Latency for various workloads

Conclusions

1. SPDK NVMe-oF TCP Target reduces the average latency by up to 1.62 usec. vs. Linux Kernel NVMe-oF TCP Target, which eliminates up to 4.59% of software overhead.
2. SPDK NVMe-oF TCP Initiator reduces the average latency by up to 3.82 usec. vs. Linux Kernel NVMe-oF TCP Initiator, which eliminates up to 18.5% of software overhead.
3. SPDK NVMe-oF TCP Target and Initiator reduce the average latency by up to 4.45 usec. vs. Linux Kernel NVMe-oF TCP Target and Initiator, which eliminates up to 16.20% of software overhead.

Test Case 4: NVMe-oF Performance with increasing # of connections

This test case was performed in order to understand throughput and latency capabilities of SPDK NVMe-oF Target vs. Linux Kernel NVMe-oF Target under increasing number of connections per subsystem. The number of connections (or I/O queue pairs) per NVMe-oF subsystem were varied and corresponding aggregated IOPS and number of CPU cores metrics were reported. The number of CPU cores metric was calculated from % CPU utilization measured using sar (systat package in Linux). The SPDK NVMe-oF Target was configured to run on 24 cores, 14 NVMe-oF subsystems (1 per Kioxia NVMe SSD) and 2 initiators were used both running I/Os to 7 separate subsystems using Kernel NVMe-oF initiator. We ran the following workloads on the host systems:

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

Table 23: NVMe-oF Performance with increasing number of connections test configuration

Item	Description
Test Case	NVMe-oF Target performance under varying # of connections
SPDK NVMe-oF Target configuration	Same as in Test Case #1, using 24 CPU cores.
Kernel NVMe-oF Target configuration	Target configuration file loaded using nvmet-cli tool. For detail configuration file contents please see Appendix D .
Kernel NVMe-oF Initiator #1	<p>Device config Performed using nvme-cli tool.</p> <pre>modprobe nvme-fabrics nvme connect -n nqn.2018-09.io.spdk:cnode1 -t tcp -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode2 -t tcp -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode3 -t tcp -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode4 -t tcp -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode5 -t tcp -a 20.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode6 -t tcp -a 20.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode7 -t tcp -a 20.0.1.1 -s 4420</pre>
Kernel NVMe-oF Initiator #2	<p>Device config Performed using nvme-cli tool.</p> <pre>modprobe nvme-fabrics nvme connect -n nqn.2018-09.io.spdk:cnode8 -t tcp -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode9 -t tcp -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode10 -t tcp -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode11 -t tcp -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode12 -t tcp -a 10.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode13 -t tcp -a 10.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode14 -t tcp -a 10.0.1.1 -s 4420</pre>
fio configuration (used on both initiators)	<p>fio.conf [global] ioengine=io_uring thread=1 group_reporting=1</p>

<pre>direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth={128, 192, 384} time_based=1 ramp_time=60 runtime=300 numjobs={1, 4, 8, 12, 16} [filename1] filename=/dev/nvme0n1 [filename2] filename=/dev/nvme1n1 [filename3] filename=/dev/nvme2n1 [filename4] filename=/dev/nvme3n1 [filename5] filename=/dev/nvme4n1 [filename6] filename=/dev/nvme5n1 [filename7] filename=/dev/nvme6n1</pre>
--

The number of CPU cores used while running the SPDK NVMe-oF target was 24, whereas for the case of Linux Kernel NVMe-oF target there was no CPU core limitation applied.

The metrics in the graph represent relative efficiency in IOPS/core which was calculated based on total aggregate IOPS divided by total CPU cores used while running that specific workload. For the case of Kernel NVMe-oF target, total CPU cores was calculated from % CPU utilization which was measured using sar utility in Linux.

4KiB Random Read Results



Figure 14: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KiB 100% Random Reads QD=384 using the Kernel Initiator

Table 24: Linux Kernel NVMe-oF TCP Target: 4KiB 100% Random Reads, QD=384

Connections per subsystem	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	131.93	4026.3	1559.4	22.3
4	268.34	8189.1	767.4	37.0
8	322.39	9838.5	646.6	58.0
12	295.32	9012.3	596.0	63.2
16	302.81	9241.1	705.8	67.6

Table 25: SPDK NVMe-oF TCP Target: 4KiB 100% Random Reads, QD=384

Connections per subsystem	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	248.41	7580.9	850.3	30.0
4	258.21	7880.1	847.2	30.7
8	235.07	7173.7	749.1	32.7
12	239.40	7306.0	939.1	32.5
16	248.41	7580.9	850.3	30.0

4KiB Random Write Results

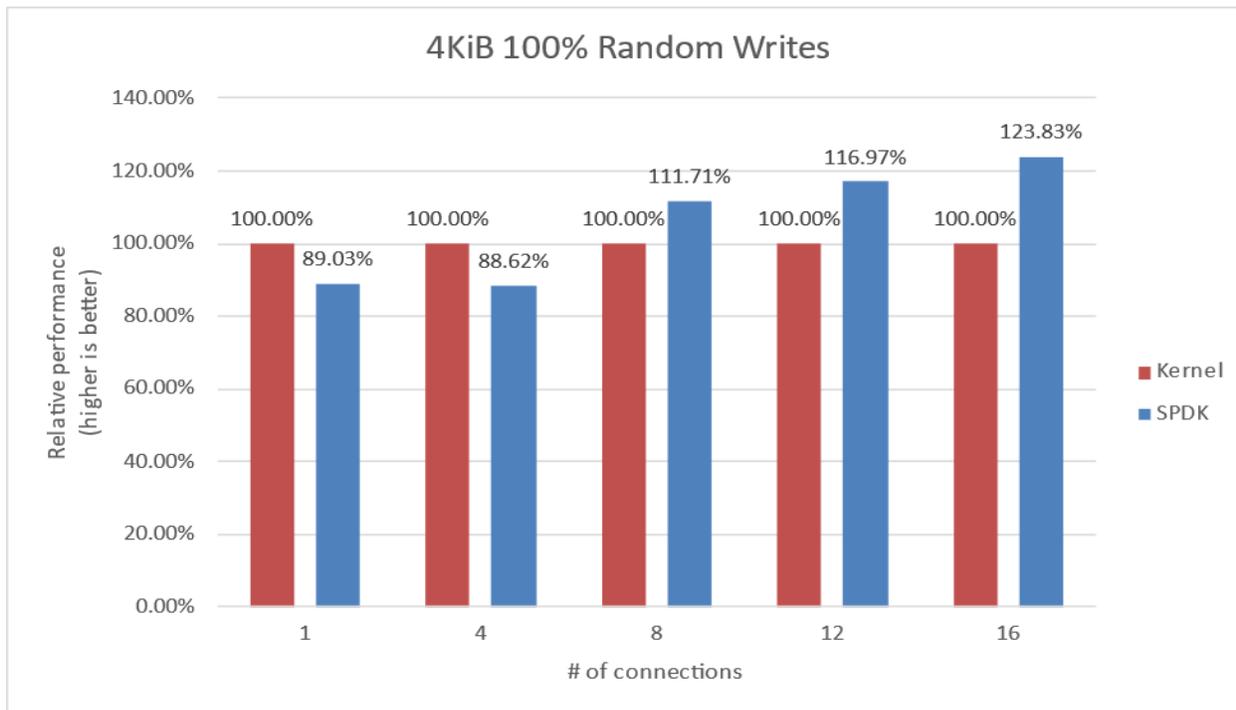


Figure 15: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KiB 100% Random Writes QD=192 using the Kernel Initiator

Note: Drives were not pre-conditioned while running 100% Random write I/O Test

Table 26: Linux Kernel NVMe-oF TCP Target: 4KiB 100% Random Writes, QD=192

Connections per subsystem	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	133.20	4064.9	767.4	31.1
4	183.20	5590.9	580.4	33.7
8	227.63	6946.6	474.7	51.2
12	228.30	6967.2	386.7	58.3
16	216.60	6610.3	522.5	56.1

Table 27: SPDK NVMe-oF TCP Target: 4KiB 100% Random Writes, QD=128

Connections per subsystem	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	118.94	3629.7	853.9	31.2
4	151.24	4615.6	739.3	31.4
8	159.79	4876.5	716.2	32.2
12	154.33	4709.7	570.5	33.7
16	159.22	4858.9	728.7	33.3

4KiB Random Read-Write Results

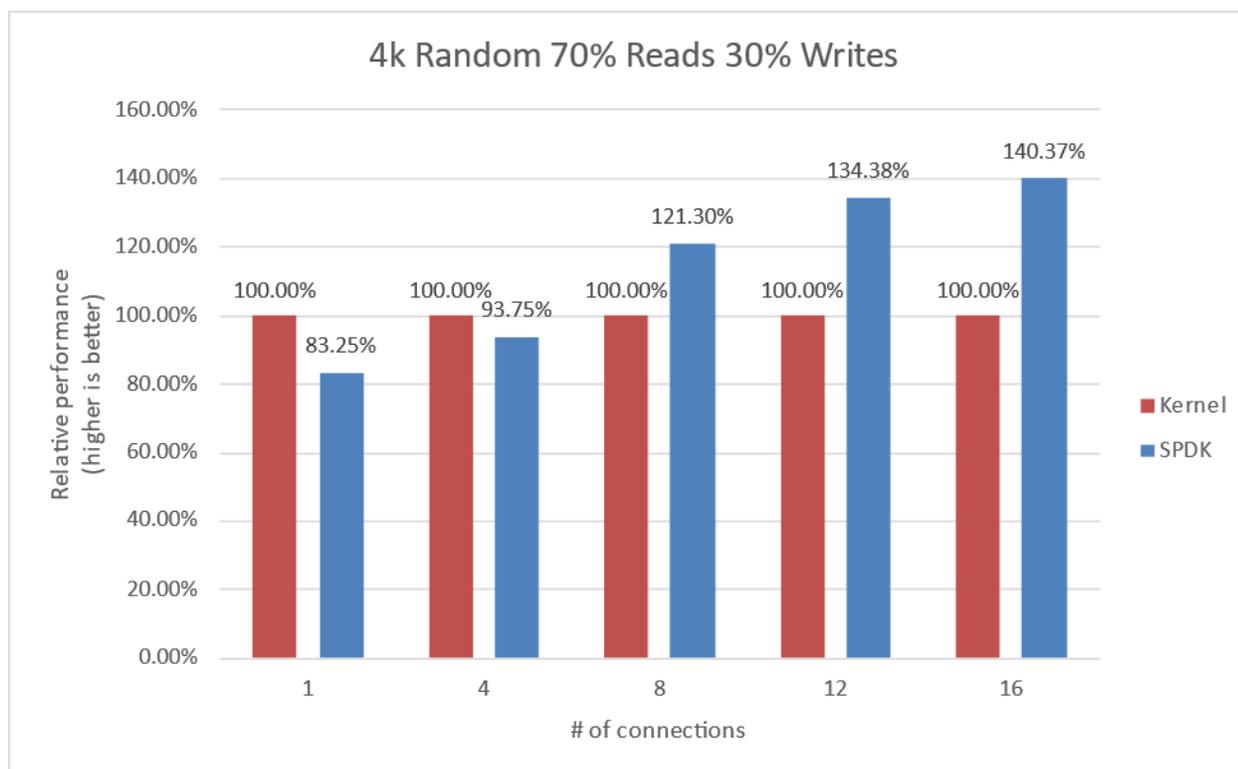


Figure 16: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KiB Random 70% Reads 30% Writes QD=192 using Kernel Initiator

Table 28: Linux Kernel NVMe-oF TCP Target: 4KiB 70% Random Read 30% Random Write, QD=192

Connections per subsystem	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	129.81	3961.4	790.9	24.0
4	192.00	5859.5	544.2	32.0
8	237.25	7240.4	440.8	53.2
12	222.21	6781.4	395.7	61.7
16	218.76	6676.0	489.1	62.0

Table 29: SPDK NVMe-oF TCP Target: 4KiB 70% Random Read 30% Random Write, QD=128

Connections per subsystem	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	133.49	4073.7	768.4	29.6
4	173.03	5280.3	613.3	30.8
8	175.93	5369.1	620.0	32.5
12	159.37	4863.5	552.3	32.9
16	167.45	5110.3	673.4	33.8

Low Connections Results

During testing it was observed that relative efficiency of SPDK Target is about 50-60% of Kernel Target, because SPDK uses a fixed number of CPU cores that was configured at start up and does not have a mechanism to decrease the number of I/O cores on the fly if the SPDK target does not need all of the CPU resources.

The test cases with 1 connection per subsystems were re-run with SPDK using only 4 CPU cores.

Table 30: SPDK & Kernel NVMe-oF TCP Target relative efficiency comparison for various workloads, QD=128, 1 connection per subsystem

Workload	Target	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
Random Read	Linux	93.42	2850.9	628.3	10.9
	SPDK	69.80	2130.2	841.2	6.4
Random Write	Linux	91.14	2781.4	644.0	13.0
	SPDK	46.56	1420.8	1261.1	7.1
Random Read/Write	Linux	98.27	2998.9	597.3	12.6
	SPDK	58.88	1797.0	997.0	6.9

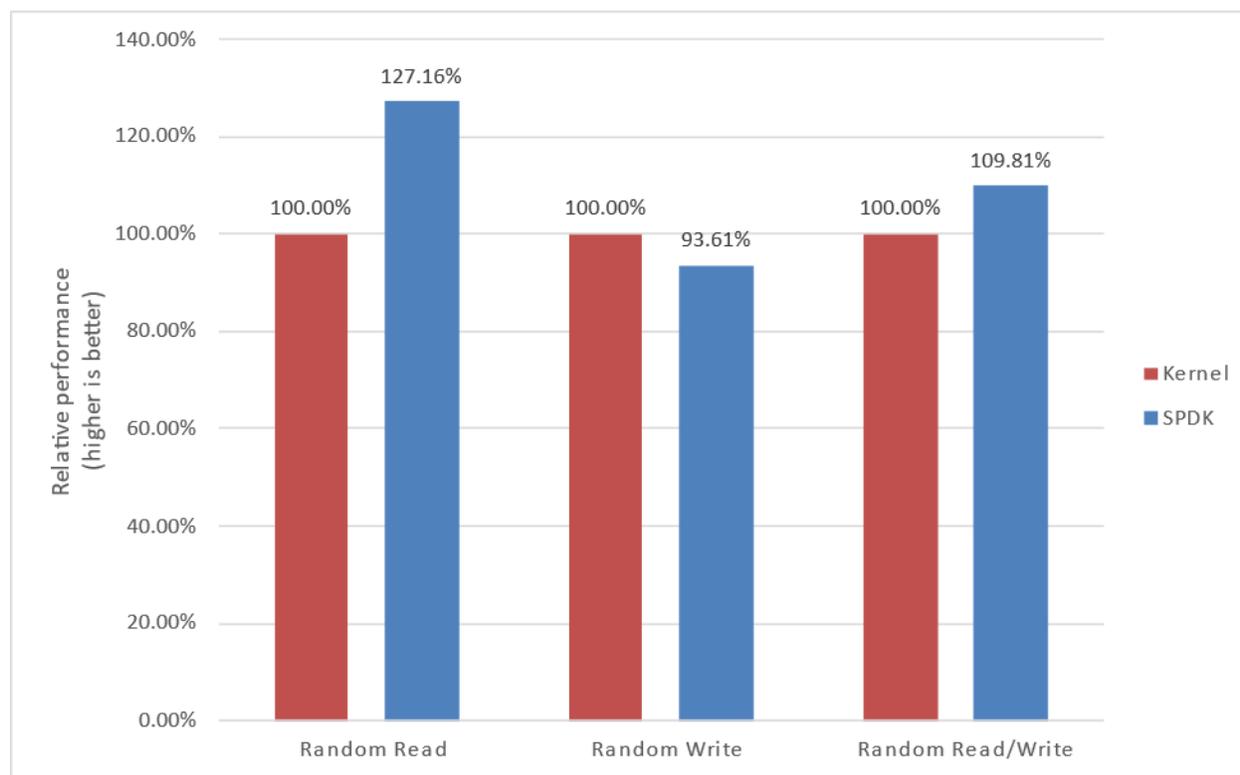


Figure 17: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for various workloads, 1 connection per subsystem and reduced number of SPDK Target CPU Cores (4)

Conclusions

1. The Linux Kernel NVMe-oF TCP target relative efficiency in IOPS/Core was better than SPDK when there was low number of connections per subsystem because the SPDK NVMe-oF target uses a fixed number of CPU cores when configured with the static scheduler. Therefore, we re-run the test cases with 1 connection per subsystem but lowered the number of I/O cores used by the SPDK Target to 4 and added the results to the tables which show a relative performance better than Linux Kernel NVMe-oF TCP target up to 1.23x for Random Read and 1.11x times for Random Read/Write workloads. Random Write workload relative performance was on similar level.
2. The performance peaked for all workloads at 8 connections per subsystem for both SPDK and Kernel NVMe-oF TCP Target.
3. The SPDK NVMe-oF TCP target relative efficiency in IOPS/Core was up to 1.27x, 1.23x and 1.4x times better than the Linux Kernel NVMe-oF target for Random Read, Random Write and Random Read/Write workloads respectively.

Summary

This report showcased performance results with SPDK NVMe-oF TCP target and initiator under various test cases, including:

- I/O core scaling
- Average I/O latency
- Performance with increasing number of connections per subsystems

It compared performance results while running Linux Kernel NVMe-oF (Target/Initiator) against the accelerated polled-mode driven SPDK NVMe-oF (Target/Initiator) implementation.

Throughput scales up and latency decreases almost linearly with the scaling of SPDK NVMe-oF target I/O cores when serving 4KiB random workloads. The SPDK NVMe-oF target saturates a 400 Gbps network link using 32 CPU cores for the 4KiB Random Read and 200 Gbps for Random Write workload at 24 CPU cores. The IOPS scalability remains close to linear for all workloads until the results are close to saturating network link (or NVMe drives throughput in case of Random Write workload).

For the SPDK NVMe-oF TCP Initiator running Random Read and Random Read/Write workloads the IOPS throughput scales almost linearly with addition of CPU cores until the network was saturated. Further increasing the number of CPU cores results in performance degradation. A single initiator was able to saturate a 200Gb link for these workloads.

For the NVMe-oF TCP latency comparison, the SPDK NVMe-oF Target and Initiator average latency is up to about 4.59% and 18.5% lower than their Linux Kernel counterparts respectively when testing against null block device-based backend

The SPDK NVMe-oF TCP Target performed up to 1.64 times better w.r.t IOPS/core than Linux Kernel NVMe-oF target while running 4KiB 100% Random Read workload with increasing number of connections per NVMe-oF subsystem.

This report provides information regarding methodologies and practices while benchmarking NVMe-oF using SPDK, as well as the Linux Kernel. It should be noted that the performance data showcased in this report is based on specific hardware and software configurations and that performance results may vary depending on the hardware and software configurations.

List of Figures

Figure 1: High-Level NVMe-oF TCP performance testing setup	8
Figure 2: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Read workload at QD = 384.....	12
Figure 3: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Write Workload at QD=384.....	13
Figure 4: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KiB Random 70/30 Read/Write workload at QD=384	14
Figure 5: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KiB 100% Sequential Read Workload at QD=16 and initiator fio numjobs=4	15
Figure 6: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KiB 100% Sequential Write Workload at QD=32 and Initiator fio numjobs=4	16
Figure 7: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KiB Sequential 70% Read 30% Write Workload at QD=32 and Initiator fio numjobs=4.....	17
Figure 8: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Read QD=256 workload	21
Figure 9: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KiB 100% Random Write Workload at QD=128.....	22
Figure 10: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KiB Random 70% Read 30% Write Workload at QD=256	23
Figure 11: SPDK vs. Kernel NVMe-oF TCP Target Average I/O Latency for various workloads run using the Kernel Initiator.....	28
Figure 12: SPDK vs. Kernel NVMe-oF TCP Initiator Average I/O Latency for various workloads against SPDK Target.....	29
Figure 13: SPDK vs. Kernel NVMe-oF TCP solutions Average I/O Latency for various workloads	30
Figure 14: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KiB 100% Random Reads QD=384 using the Kernel Initiator	34
Figure 15: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KiB 100% Random Writes QD=192 using the Kernel Initiator	35
Figure 16: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KiB Random 70% Reads 30% Writes QD=192 using Kernel Initiator	36
Figure 17: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for various workloads, 1 connection per subsystem and reduced number of SPDK Target CPU Cores (4).37	
Figure 18: IOPS/Core Performance comparison of SPDK NVMe-oF Target static and dynamic scheduler for 4 KiB Random Read workload, 1 connection per subsystem, 24 CPU cores SPDK Target.....	44
Figure 19: IOPS/Core Performance comparison of SPDK NVMe-oF Target static and dynamic scheduler for 4 KiB Random Write workload, 1 connection per subsystem, 24 CPU cores SPDK Target	45
Figure 20: IOPS/Core Performance comparison of SPDK NVMe-oF Target static and dynamic scheduler for 4 KiB Random 70/30 Read/Write workload, 1 connection per subsystem, 24 CPU cores SPDK Target	45

List of Tables

Table 1: Hardware setup configuration – Target system	4
Table 2: Hardware setup configuration – Initiator system 1	5
Table 3: Hardware setup configuration – Initiator system 2	5
Table 4: Test systems BIOS settings	6
Table 5: SPDK NVMe-oF TCP Target Core Scaling test configuration	9
Table 6: SPDK NVMe-oF TCP Target Core Scaling results, Random Read IOPS, QD=384	12
Table 7: SPDK NVMe-oF TCP Target Core Scaling results, Random Write IOPS, QD=128	13
Table 8: SPDK NVMe-oF TCP Target Core Scaling results, Random Read/Write 70%/30% IOPS, QD=384	14
Table 9: SPDK NVMe-oF TCP Target Core Scaling results, 128KiB Sequential Read IOPS, QD=16	15
Table 10: SPDK NVMe-oF TCP Target Core Scaling results, 128KiB Sequential Write IOPS, QD=32	16
Table 11: SPDK NVMe-oF TCP Target Core Scaling results, 128KiB Sequential 70% Read 30% Write IOPS, QD=32	17
Table 12: SPDK NVMe-oF TCP Initiator Core Scaling test configuration	19
Table 13: SPDK NVMe-oF TCP Initiator Core Scaling results, 4KiB Random Read IOPS, QD=256	21
Table 14: SPDK NVMe-oF TCP Initiator Core Scaling results, 4KiB Random Write IOPS, QD=128	22
Table 15: SPDK NVMe-oF TCP Initiator Core Scaling results, 4KiB Random 70%/30% Read/Write IOPS, QD=256	23
Table 16: Linux Kernel vs. SPDK NVMe-oF TCP Latency test configuration	25
Table 17: SPDK NVMe-oF Target Latency and IOPS at QD=1, Null Block Device	28
Table 18: Linux Kernel NVMe-oF Target Latency and IOPS at QD=1, Null Block Device	28
Table 19: SPDK NVMe-oF Initiator Latency and IOPS at QD=1, Null Block Device	29
Table 20: Linux Kernel NVMe-oF Initiator Latency and IOPS at QD=1, Null Block Device	29
Table 21: SPDK NVMe-oF Latency and IOPS at QD=1, Null Block Device	30
Table 22: Linux Kernel NVMe-oF Latency and IOPS at QD=1, Null Block Device	30
Table 23: NVMe-oF Performance with increasing number of connections test configuration	32
Table 24: Linux Kernel NVMe-oF TCP Target: 4KiB 100% Random Reads, QD=384	34
Table 25: SPDK NVMe-oF TCP Target: 4KiB 100% Random Reads, QD=384	34
Table 26: Linux Kernel NVMe-oF TCP Target: 4KiB 100% Random Writes, QD=192	35
Table 27: SPDK NVMe-oF TCP Target: 4KiB 100% Random Writes, QD=128	35
Table 28: Linux Kernel NVMe-oF TCP Target: 4KiB 70% Random Read 30% Random Write, QD=192	36
Table 29: SPDK NVMe-oF TCP Target: 4KiB 70% Random Read 30% Random Write, QD=128	36

Table 30: SPDK & Kernel NVMe-oF TCP Target relative efficiency comparison for various workloads, QD=192, 1 connection per subsystem37

Table 31: SPDK NVMe-oF TCP Target static and dynamic scheduler performance comparison for various workloads, QD=1-128, 1 connection per subsystem43

Appendix A – SPDK NVMe-oF Target Dynamic Scheduler

SPDK supports scheduling of lightweight threads by use of schedulers which are provided as plugins. By default, SPDK uses “static” scheduler where threads are distributed round-robin among application reactors and no thread is ever moved. The available “dynamic” scheduler allows moving threads between application reactors – active threads are distributed equally between reactors taking desired `cpu_mask` into consideration, while idle threads are moved to the main core. In case a thread becomes active again it is once again moved to appropriate reactor. Thank to this approach “dynamic” scheduler can be used for power saving and reduction of CPU utilization.

The number of CPU cores metric was calculated from % CPU utilization measured using `sar` utility (`systat` package in Linux). The “average power consumption” metric refers to the overall, system-wide power draw of the Target system during the test. It was calculated by using `ipmitool` to extract power statistics from sensors exposed by the system’s BMC.

The test cases from Test Case 4 (“NVMe-oF Performance with increasing number of connections”) with 1 connection per subsystem were re-run with SPDK with 24 CPU cores and scheduler set to both “static” and “dynamic” for comparison. We used a target configuration that scales to saturate 100 Gbps and observed that dynamic scheduling reduced the number of CPU cores used by the target at a slight cost of IOPS performance. For high queue workloads CPU utilization and IOPS are similar for both schedulers, except for Random Read workload where Dynamic Scheduler still shows lower CPU consumption.

Table 31: SPDK NVMe-oF TCP Target static and dynamic scheduler performance comparison for various workloads, QD=1-128, 1 connection per subsystem

Workload	Scheduler	Queue Depth	Bandwidth (Gbps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores	Average Power Consumption [W]
Random Read	Static	1	1.60	48.9	284.8	25.3	679.41
		8	12.75	389.2	287.5	25.8	688.22
		32	40.26	1228.8	364.3	26.2	700.09
		64	68.14	2079.5	430.6	26.5	719.79
		128	97.14	2964.5	604.3	25.9	744.54
	Dynamic	1	0.74	22.7	614.1	2.4	570.01
		8	6.82	208.2	537.6	2.5	585.02
		32	28.78	878.3	509.9	5.1	601.39
		64	51.32	1566.3	571.9	6.5	619.45
		128	76.88	2346.3	763.7	8.0	637.09
Random Write	Static	1	7.21	220.1	63.2	26.4	689.11
		8	14.44	440.7	253.8	26.0	693.89
		32	47.10	1437.4	311.4	27.1	723.53
		64	71.86	2193.1	408.3	26.9	747.02
		128	91.60	2795.3	640.9	27.4	775.66

Random Read/Write	Dynamic	1	0.92	28.0	498.1	2.4	578.07
		8	7.37	224.8	498.0	3.8	592.68
		32	31.16	951.0	471.2	6.7	630.61
		64	49.77	1518.8	590.6	9.6	653.97
		128	67.66	2064.9	869.4	13.2	685.65
	Static	1	1.72	52.6	264.7	25.3	680.82
		8	12.94	394.9	283.3	25.9	690.55
		32	42.28	1290.2	347.0	26.4	707.62
		64	70.38	2147.9	416.9	26.9	732.08
		128	101.66	3102.3	577.4	27.4	765.6
Dynamic	1	0.81	24.6	565.9	2.5	577.55	
	8	6.98	213.0	536.1	2.6	584.6	
	32	32.62	995.5	449.8	6.4	615.07	
	64	52.67	1607.4	557.4	8.5	634.15	
	128	79.32	2420.7	740.3	11.7	667.14	

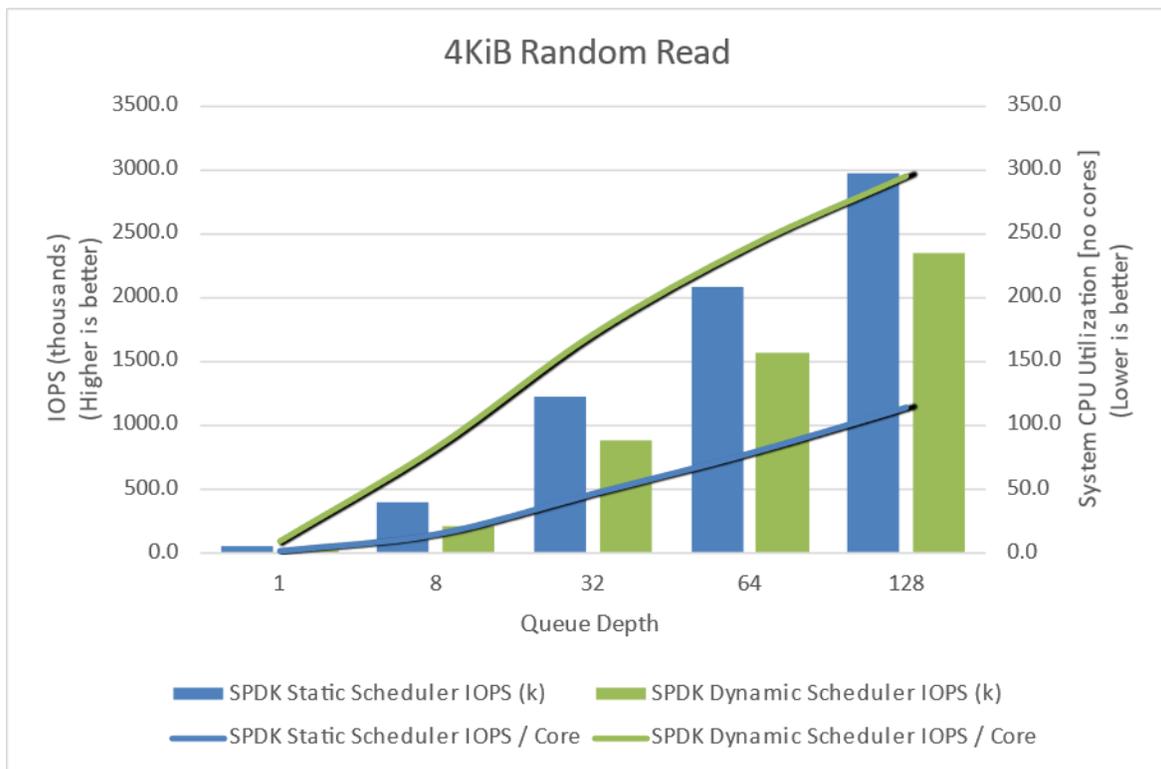


Figure 18: IOPS/Core Performance comparison of SPDK NVMe-oF Target static and dynamic scheduler for 4 KiB Random Read workload, 1 connection per subsystem, 24 CPU cores SPDK Target

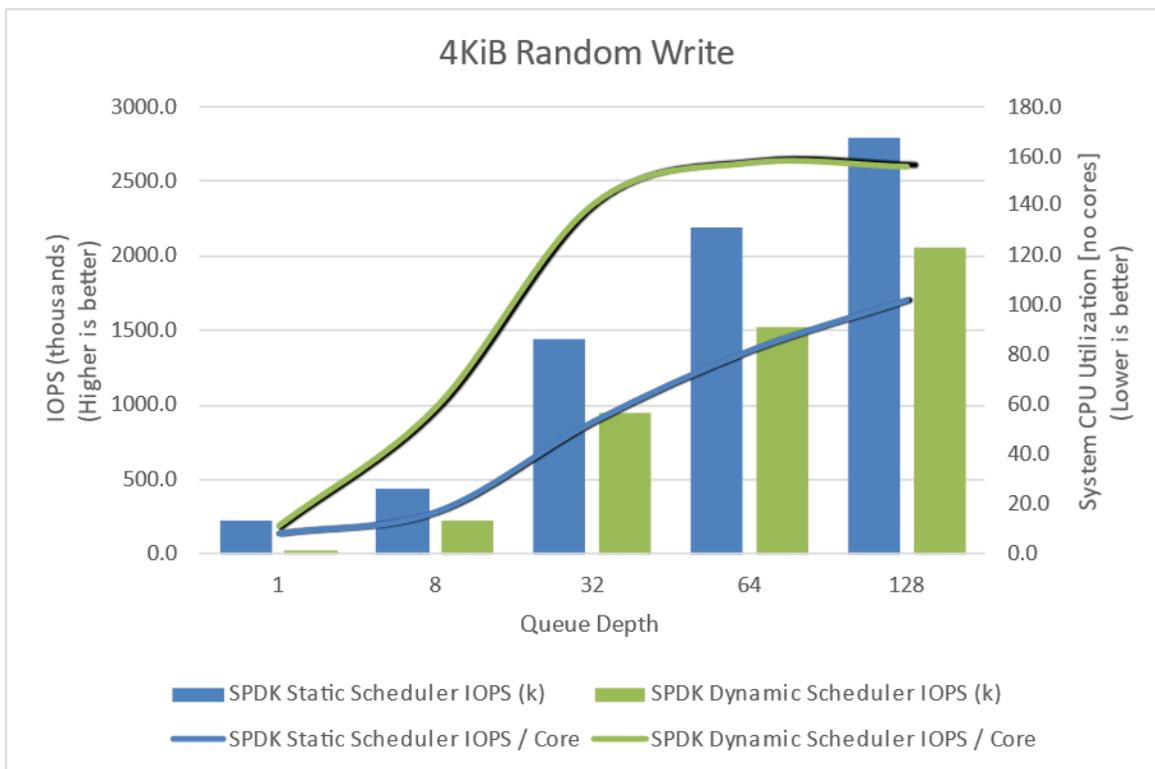


Figure 19: IOPS/Core Performance comparison of SPDK NVMe-oF Target static and dynamic scheduler for 4 KiB Random Write workload, 1 connection per subsystem, 24 CPU cores SPDK Target

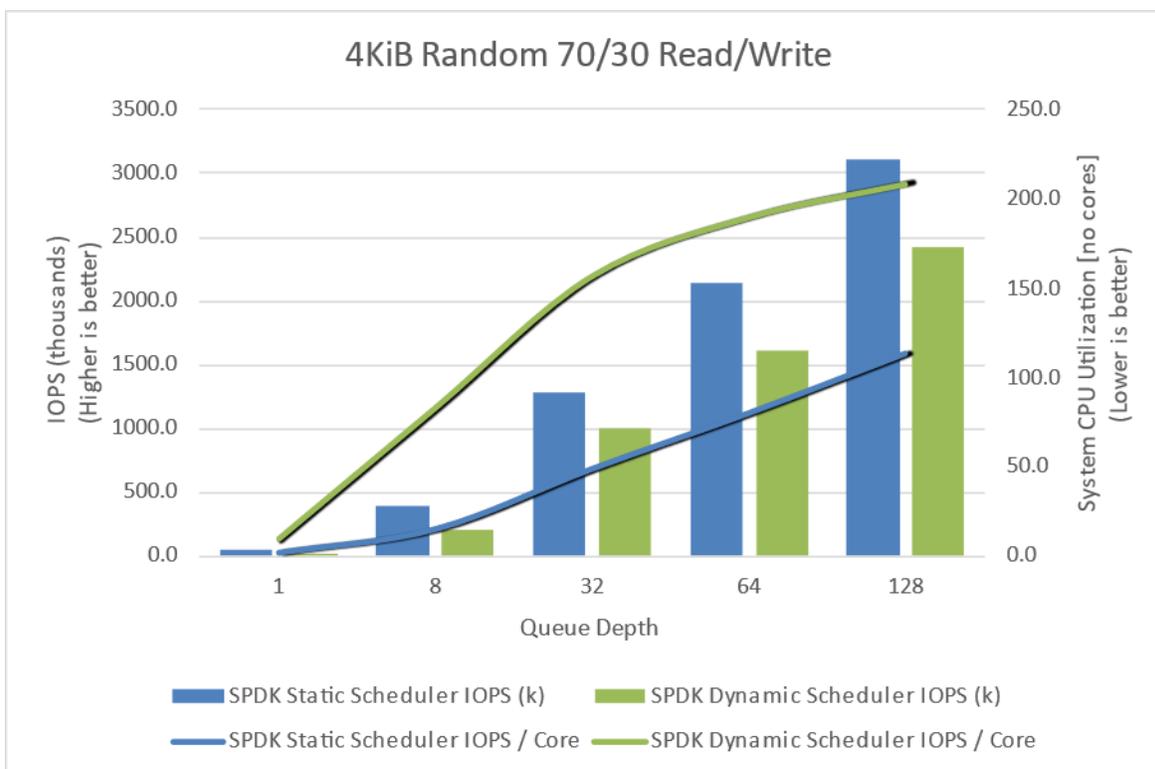


Figure 20: IOPS/Core Performance comparison of SPDK NVMe-oF Target static and dynamic scheduler for 4 KiB Random 70/30 Read/Write workload, 1 connection per subsystem, 24 CPU cores SPDK Target

Appendix B – Test Case 1 SPDK NVMe-oF Initiator bdev configuration

Initiator system 1

```
{
  "subsystems": [
    {
      "subsystem": "bdev",
      "config": [
        {
          "method": "bdev_nvme_attach_controller",
          "params": {
            "name": "Nvme0",
            "trtype": "tcp",
            "traddr": "20.0.0.1",
            "trsvcid": "4420",
            "subnqn": "nqn.2018-09.io.spdk:cnode0",
            "adrfam": "IPv4"
          }
        },
        {
          "method": "bdev_nvme_attach_controller",
          "params": {
            "name": "Nvme1",
            "trtype": "tcp",
            "traddr": "20.0.0.1",
            "trsvcid": "4420",
            "subnqn": "nqn.2018-09.io.spdk:cnode1",
            "adrfam": "IPv4"
          }
        },
        {
          "method": "bdev_nvme_attach_controller",
          "params": {
            "name": "Nvme2",
            "trtype": "tcp",
            "traddr": "20.0.0.1",
            "trsvcid": "4420",
            "subnqn": "nqn.2018-09.io.spdk:cnode2",
            "adrfam": "IPv4"
          }
        },
        {
          "method": "bdev_nvme_attach_controller",
          "params": {
            "name": "Nvme3",
            "trtype": "tcp",
            "traddr": "20.0.0.1",
            "trsvcid": "4420",
            "subnqn": "nqn.2018-09.io.spdk:cnode3",
          }
        }
      ]
    }
  ]
}
```

```
    "adrfam": "IPv4"
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme4",
    "trtype": "tcp",
    "traddr": "20.0.1.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode4",
    "adrfam": "IPv4"
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme5",
    "trtype": "tcp",
    "traddr": "20.0.1.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode5",
    "adrfam": "IPv4"
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme6",
    "trtype": "tcp",
    "traddr": "20.0.1.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode6",
    "adrfam": "IPv4"
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme7",
    "trtype": "tcp",
    "traddr": "20.0.1.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode7",
    "adrfam": "IPv4"
  }
}
]
]
```

Initiator system 2

```
{
  "subsystems": [
    {
      "subsystem": "bdev",
```

```
"config": [  
  {  
    "method": "bdev_nvme_attach_controller",  
    "params": {  
      "name": "Nvme0",  
      "trtype": "tcp",  
      "traddr": "10.0.0.1",  
      "trsvcid": "4420",  
      "subnqn": "nqn.2018-09.io.spdk:cnode0",  
      "adrfam": "IPv4"  
    }  
  },  
  {  
    "method": "bdev_nvme_attach_controller",  
    "params": {  
      "name": "Nvme1",  
      "trtype": "tcp",  
      "traddr": "10.0.0.1",  
      "trsvcid": "4420",  
      "subnqn": "nqn.2018-09.io.spdk:cnode1",  
      "adrfam": "IPv4"  
    }  
  },  
  {  
    "method": "bdev_nvme_attach_controller",  
    "params": {  
      "name": "Nvme2",  
      "trtype": "tcp",  
      "traddr": "10.0.0.1",  
      "trsvcid": "4420",  
      "subnqn": "nqn.2018-09.io.spdk:cnode2",  
      "adrfam": "IPv4"  
    }  
  },  
  {  
    "method": "bdev_nvme_attach_controller",  
    "params": {  
      "name": "Nvme3",  
      "trtype": "tcp",  
      "traddr": "10.0.0.1",  
      "trsvcid": "4420",  
      "subnqn": "nqn.2018-09.io.spdk:cnode3",  
      "adrfam": "IPv4"  
    }  
  },  
  {  
    "method": "bdev_nvme_attach_controller",  
    "params": {  
      "name": "Nvme4",  
      "trtype": "tcp",  
      "traddr": "10.0.1.1",  
      "trsvcid": "4420",  
      "subnqn": "nqn.2018-09.io.spdk:cnode4",  
      "adrfam": "IPv4"  
    }  
  }  
]
```

```
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme5",
      "trtype": "tcp",
      "traddr": "10.0.1.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode5",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme6",
      "trtype": "tcp",
      "traddr": "10.0.1.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode6",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme7",
      "trtype": "tcp",
      "traddr": "10.0.1.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode7",
      "adrfam": "IPv4"
    }
  }
]
}
```

Appendix C – Test Case 2 SPDK NVMe-oF Initiator bdev configuration

```
{
  "subsystems": [
    {
      "subsystem": "bdev",
      "config": [
        {
          "method": "bdev_nvme_attach_controller",
          "params": {
            "name": "Nvme0",
            "trtype": "tcp",
            "traddr": "20.0.0.1",
            "trsvcid": "4420",
            "subnqn": "nqn.2018-09.io.spdk:cnode0",
            "adrfam": "IPv4"
          }
        }
      ]
    }
  ]
}
```

```
    },
    {
      "method": "bdev_nvme_attach_controller",
      "params": {
        "name": "Nvme1",
        "trtype": "tcp",
        "traddr": "20.0.0.1",
        "trsvcid": "4420",
        "subnqn": "nqn.2018-09.io.spdk:cnode1",
        "adrfam": "IPv4"
      }
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme2",
      "trtype": "tcp",
      "traddr": "20.0.0.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode2",
      "adrfam": "IPv4"
    }
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme3",
    "trtype": "tcp",
    "traddr": "20.0.0.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode3",
    "adrfam": "IPv4"
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme4",
    "trtype": "tcp",
    "traddr": "20.0.0.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode4",
    "adrfam": "IPv4"
  }
},
{
  "method": "bdev_nvme_attach_controller",
  "params": {
    "name": "Nvme5",
    "trtype": "tcp",
    "traddr": "20.0.0.1",
    "trsvcid": "4420",
    "subnqn": "nqn.2018-09.io.spdk:cnode5",
    "adrfam": "IPv4"
  }
},
{
```

```
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme6",
      "trtype": "tcp",
      "traddr": "20.0.0.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode6",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme7",
      "trtype": "tcp",
      "traddr": "20.0.0.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode7",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme8",
      "trtype": "tcp",
      "traddr": "20.0.1.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode8",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme9",
      "trtype": "tcp",
      "traddr": "20.0.1.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode9",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
      "name": "Nvme10",
      "trtype": "tcp",
      "traddr": "20.0.1.1",
      "trsvcid": "4420",
      "subnqn": "nqn.2018-09.io.spdk:cnode10",
      "adrfam": "IPv4"
    }
  },
  {
    "method": "bdev_nvme_attach_controller",
    "params": {
```

```
        "name": "Nvme11",
        "trtype": "tcp",
        "traddr": "20.0.1.1",
        "trsvcid": "4420",
        "subnqn": "nqn.2018-09.io.spdk:cnode11",
        "adrfam": "IPv4"
    }
},
{
    "method": "bdev_nvme_attach_controller",
    "params": {
        "name": "Nvme12",
        "trtype": "tcp",
        "traddr": "20.0.1.1",
        "trsvcid": "4420",
        "subnqn": "nqn.2018-09.io.spdk:cnode12",
        "adrfam": "IPv4"
    }
},
{
    "method": "bdev_nvme_attach_controller",
    "params": {
        "name": "Nvme13",
        "trtype": "tcp",
        "traddr": "20.0.1.1",
        "trsvcid": "4420",
        "subnqn": "nqn.2018-09.io.spdk:cnode13",
        "adrfam": "IPv4"
    }
},
{
    "method": "bdev_nvme_attach_controller",
    "params": {
        "name": "Nvme14",
        "trtype": "tcp",
        "traddr": "20.0.1.1",
        "trsvcid": "4420",
        "subnqn": "nqn.2018-09.io.spdk:cnode14",
        "adrfam": "IPv4"
    }
},
{
    "method": "bdev_nvme_attach_controller",
    "params": {
        "name": "Nvme15",
        "trtype": "tcp",
        "traddr": "20.0.1.1",
        "trsvcid": "4420",
        "subnqn": "nqn.2018-09.io.spdk:cnode15",
        "adrfam": "IPv4"
    }
}
]
}
```

Appendix D – Test Case 3 SPDK NVMe-oF Initiator bdev configuration

```
{
  "subsystems": [
    {
      "subsystem": "bdev",
      "config": [
        {
          "method": "bdev_nvme_attach_controller",
          "params": {
            "name": "Nvme0",
            "trtype": "tcp",
            "traddr": "20.0.0.1",
            "trsvcid": "4420",
            "subnqn": "nqn.2018-09.io.spdk:cnode0",
            "adrfam": "IPv4"
          }
        }
      ]
    }
  ]
}
```

Appendix E – Kernel NVMe-oF TCP Target configuration

Example Kernel NVMe-oF TCP Target configuration for Test Case 4.

```
{
  "ports": [
    {
      "addr": {
        "adrfam": "ipv4",
        "traddr": "20.0.0.1",
        "trsvcid": "4420",
        "trtype": "tcp"
      },
      "portid": 1,
      "referrals": [],
      "subsystems": [
        "nqn.2018-09.io.spdk:cnode1"
      ]
    },
    {
      "addr": {
        "adrfam": "ipv4",
        "traddr": "20.0.0.1",
        "trsvcid": "4421",
        "trtype": "tcp"
      },
      "portid": 2,
      "referrals": [],
      "subsystems": [
        "nqn.2018-09.io.spdk:cnode2"
      ]
    }
  ]
}
```

```
]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "20.0.0.1",
    "trsvcid": "4422",
    "trtype": "tcp"
  },
  "portid": 3,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode3"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "20.0.0.1",
    "trsvcid": "4423",
    "trtype": "tcp"
  },
  "portid": 4,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode4"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "20.0.1.1",
    "trsvcid": "4424",
    "trtype": "tcp"
  },
  "portid": 5,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode5"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "20.0.1.1",
    "trsvcid": "4425",
    "trtype": "tcp"
  },
  "portid": 6,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode6"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
```

```
    "traddr": "20.0.1.1",
    "trsvcid": "4426",
    "trtype": "tcp"
  },
  "portid": 7,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode7"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "20.0.1.1",
    "trsvcid": "4427",
    "trtype": "tcp"
  },
  "portid": 8,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode8"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.0.1",
    "trsvcid": "4428",
    "trtype": "tcp"
  },
  "portid": 9,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode9"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.0.1",
    "trsvcid": "4429",
    "trtype": "tcp"
  },
  "portid": 10,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode10"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.0.1",
    "trsvcid": "4430",
    "trtype": "tcp"
  },
  "portid": 11,
```

```
"referrals": [],
"subsystems": [
  "nqn.2018-09.io.spdk:cnode11"
]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.0.1",
    "trsvcid": "4431",
    "trtype": "tcp"
  },
  "portid": 12,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode12"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.1.1",
    "trsvcid": "4432",
    "trtype": "tcp"
  },
  "portid": 13,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode13"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.1.1",
    "trsvcid": "4433",
    "trtype": "tcp"
  },
  "portid": 14,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode14"
  ]
},
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.1.1",
    "trsvcid": "4434",
    "trtype": "tcp"
  },
  "portid": 15,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode15"
  ]
},
},
```

```
{
  "addr": {
    "adrfam": "ipv4",
    "traddr": "10.0.1.1",
    "trsvcid": "4435",
    "trtype": "tcp"
  },
  "portid": 16,
  "referrals": [],
  "subsystems": [
    "nqn.2018-09.io.spdk:cnode16"
  ]
},
],
"hosts": [],
"subsystems": [
  {
    "allowed_hosts": [],
    "attr": {
      "allow_any_host": "1",
      "version": "1.3"
    },
    "namespaces": [
      {
        "device": {
          "path": "/dev/nvme0n1",
          "uuid": "b53be81d-6f5c-4768-b3bd-203614d8cf20"
        },
        "enable": 1,
        "nsid": 1
      }
    ],
    "nqn": "nqn.2018-09.io.spdk:cnode1"
  },
  {
    "allowed_hosts": [],
    "attr": {
      "allow_any_host": "1",
      "version": "1.3"
    },
    "namespaces": [
      {
        "device": {
          "path": "/dev/nvme1n1",
          "uuid": "12fcf584-9c45-4b6b-abc9-63a763455cf7"
        },
        "enable": 1,
        "nsid": 2
      }
    ],
    "nqn": "nqn.2018-09.io.spdk:cnode2"
  },
  {
    "allowed_hosts": [],
    "attr": {
      "allow_any_host": "1",
      "version": "1.3"
    }
  }
]
```

```
    },
    "namespaces": [
      {
        "device": {
          "path": "/dev/nvme2n1",
          "uuid": "ceae8569-69e9-4831-8661-90725bdf768d"
        },
        "enable": 1,
        "nsid": 3
      }
    ],
    "nqn": "nqn.2018-09.io.spdk:cnode3"
  },
  {
    "allowed_hosts": [],
    "attr": {
      "allow_any_host": "1",
      "version": "1.3"
    },
    "namespaces": [
      {
        "device": {
          "path": "/dev/nvme3n1",
          "uuid": "39f36db4-2cd5-4f69-b37d-1192111d52a6"
        },
        "enable": 1,
        "nsid": 4
      }
    ],
    "nqn": "nqn.2018-09.io.spdk:cnode4"
  },
  {
    "allowed_hosts": [],
    "attr": {
      "allow_any_host": "1",
      "version": "1.3"
    },
    "namespaces": [
      {
        "device": {
          "path": "/dev/nvme4n1",
          "uuid": "984aed55-90ed-4517-ae36-d3afb92dd41f"
        },
        "enable": 1,
        "nsid": 5
      }
    ],
    "nqn": "nqn.2018-09.io.spdk:cnode5"
  },
  {
    "allowed_hosts": [],
    "attr": {
      "allow_any_host": "1",
      "version": "1.3"
    },
    "namespaces": [
      {
```

```
    "device": {
      "path": "/dev/nvme5n1",
      "uuid": "d6d16e74-378d-40ad-83e7-b8d8af3d06a6"
    },
    "enable": 1,
    "nsid": 6
  }
],
"nqn": "nqn.2018-09.io.spdk:cnode6"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme6n1",
        "uuid": "a65dc00e-d35c-4647-9db6-c2a8d90db5e8"
      },
      "enable": 1,
      "nsid": 7
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode7"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme7n1",
        "uuid": "1b242cb7-8e47-4079-a233-83e2cd47c86c"
      },
      "enable": 1,
      "nsid": 8
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode8"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme8n1",
        "uuid": "f12bb0c9-a2c6-4eef-a94f-5c4887bbf77f"
      }
    }
  ]
}
```

```
    },
    "enable": 1,
    "nsid": 9
  }
],
"nqn": "nqn.2018-09.io.spdk:cnode9"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme9n1",
        "uuid": "40fae536-227b-47d2-bd74-8ab76ec7603b"
      },
      "enable": 1,
      "nsid": 10
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode10"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme10n1",
        "uuid": "b9756b86-263a-41cf-a68c-5c7b23c7a6eb"
      },
      "enable": 1,
      "nsid": 11
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode11"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme11n1",
        "uuid": "9d7e74cc-97f3-40fb-8e90-f2d02b5fff4c"
      },
      "enable": 1,
      "nsid": 12
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode12"
}
```

```
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode12"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme12n1",
        "uuid": "d3f4017b-4f7d-454d-94a9-ea75ffc7436d"
      },
      "enable": 1,
      "nsid": 13
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode13"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme13n1",
        "uuid": "6b9a65a3-6557-4713-8bad-57d9c5cb17a9"
      },
      "enable": 1,
      "nsid": 14
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode14"
},
{
  "allowed_hosts": [],
  "attr": {
    "allow_any_host": "1",
    "version": "1.3"
  },
  "namespaces": [
    {
      "device": {
        "path": "/dev/nvme14n1",
        "uuid": "ed69ba4d-8727-43bd-894a-7b08ade4f1b1"
      },
      "enable": 1,
      "nsid": 15
    }
  ],
  "nqn": "nqn.2018-09.io.spdk:cnode15"
```

```
    },  
    {  
      "allowed_hosts": [],  
      "attr": {  
        "allow_any_host": "1",  
        "version": "1.3"  
      },  
      "namespaces": [  
        {  
          "device": {  
            "path": "/dev/nvme15n1",  
            "uuid": "5b8e9af4-0ab4-47fb-968f-b13e4b607f4e"  
          },  
          "enable": 1,  
          "nsid": 16  
        }  
      ],  
      "nqn": "nqn.2018-09.io.spdk:cnode16"  
    }  
  ]  
}
```

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.