



SPDK NVMe-oF TCP (Target & Initiator)

Performance Report

Release 20.04

Testing Date: May 2020

Performed by: Karol Latecki (karol.latecki@intel.com)

Maciej Wawryk (maciejx.wawryk@intel.com)

Acknowledgments:

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

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

Benjamin Walker (benjamin.walker@intel.com)

Seth Howell (seth.howell@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	5
TCP configuration	6
Kernel & BIOS spectre-meltdown information	6
Introduction to SPDK NVMe-oF (Target & Initiator).....	7
Test Case 1: SPDK NVMe-oF TCP Target I/O core scaling	9
4KB Random Read Results	12
4KB Random Write Results	13
4KB 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
4KB Random Read Results	22
4KB Random Write Results	23
4KB Random Read-Write Results.....	24
Conclusions	25
Test Case 3: Linux Kernel vs. SPDK NVMe-oF TCP Latency	26
SPDK vs Kernel NVMe-oF Target Latency Results.....	29
SPDK vs Kernel NVMe-oF TCP Initiator results.....	30
SPDK vs Kernel NVMe-oF Latency Results.....	31
Conclusions	32
Test Case 4: NVMe-oF Performance with increasing # of connections	33
4KB Random Read Results	35
4KB Random Write Results	36
4KB Random Read-Write Results.....	37
Low Connections Results	38
Conclusions	39
Summary	40
Appendix A - Jumbo Frames Performance Impact.....	41
4KB Random Read with JF results	41
4KB Random Write with JF results	41
4KB Random Read-Write with JF results	42
Appendix B – Kernel NVMe-oF TCP Target configuration	43



Audience and Purpose

This report is intended for people who are interested in evaluating SPDK NVMe-oF (Target & Initiator) performance as compared to the Linux Kernel NVMe-oF (Target & Initiator). This report compares the performance and efficiency of the SPDK NVMe-oF Target and Initiator vs. the Linux Kernel NVMe-oF Target and Initiator. This report covers the TCP transport only.

The purpose of reporting these tests 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 & Initiator).

Test setup

Target Configuration

Item	Description
Server Platform	<p>SuperMicro SYS-2029U-TN24R4T</p>  
CPU	<p>Intel® Xeon® Gold 6230 Processor (27.5MB L3, 2.10 GHz)</p> <p>Number of cores 20 per socket, number of threads 40 per socket (both sockets populated)</p>
Memory	<p>12 x 32GB Hynix HMA84GR7AFR4N-VK, DDR4, 2666MHz</p> <p>Total of 384GB</p>
Operating System	Fedora 30
BIOS	3.1a
Linux kernel version	5.4.14-100.fc30
SPDK version	SPDK 20.04
Storage	<p>OS: 1x 120GB Intel SSDSC2BB120G4</p> <p>Storage Target: 16x Intel® SSD DC P4610™ 1.6TB (FW: QDV10190) (8 on each CPU socket)</p>
NIC	<p>2x 100GbE Mellanox ConnectX-5 NICs. Both ports connected.</p> <p>1 NIC per CPU socket.</p>



Initiator 1 Configuration

Item	Description
Server Platform	Intel® Server System R2208WFTZSR
CPU	Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz (35.75MB Cache) Number of cores 24 per socket, number of threads 48 per socket (Both sockets populated)
Memory	6 x 32GB Micron M393A1G40EB1-CRC, DDR4, 2933MHz Total 192GBs
Operating System	Fedora 30
BIOS	02.01.0008 03/19/2019
Linux kernel version	5.4.14-100.fc30
SPDK version	SPDK 20.04
Storage	OS: 1x 240GB INTEL SSDSC2BB240G6
NIC	1x 100GbE Mellanox ConnectX-5 Ex NIC. Both ports connected to Target server. (connected to CPU socket 0)

Initiator 2 Configuration

Item	Description
Server Platform	Intel® Server System R2208WFTZSR
CPU	Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz (35.75MB Cache) Number of cores 24 per socket, number of threads 48 per socket (Both sockets populated)
Memory	6 x 32GB Micron M393A1G40EB1-CRC, DDR4, 2933MHz Total 192GBs
Operating System	02.01.0008 03/19/2019
BIOS	3.1 06/08/2018
Linux kernel version	5.4.14-100.fc30
SPDK version	SPDK 20.04
Storage	OS: 1x 240GB INTEL SSDSC2BB240G6
NIC	1x 100GbE Mellanox ConnectX-5 Ex NIC. Both ports connected to Target server. (connected to CPU socket 0)

BIOS settings

Item	Description
BIOS <i>(Applied to all 3 systems)</i>	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:

```
# Set 256MB buffers
net.core.rmem_max = 268435456
net.core.wmem_max = 268435456
# Increase autotuning TCP buffer limits
# min, max and default settings
# auto-tuning allowed to 128MB
net.ipv4.tcp_rmem = 4096 87380 134217728
net.ipv4.tcp_wmem = 4096 65536 134217728
# MTU probing for Jumbo Frames
net.ipv4.tcp_mtu_probing = 1
```

The NIC ports were configured to use Jumbo Frames using network-scripts (/etc/sysconfig/network-scripts for RHEL-based distributions) and setting MTU=9000 for some of the test cases for comparison with standard MTU=1500.

Kernel & BIOS spectre-meltdown information

All three server systems use Fedora 5.4.14-100.fc30 kernel version available from DNF repository with default patches for spectre-meltdown issue enabled.

BIOS on all systems was updated to post spectre-meltdown versions as well.



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, TCP and Intel® Omni-Path. 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 uses the underlying transport layer API which in case of TCP are POSIX sockets. In case of RDMA-capable NICs Infiniband/RDMA verbs API is used which should work on all flavors of RDMA transports, but is currently tested against RoCEv2, iWARP, and Omni-Path NICs. 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 sixteen Intel® SSD DC P4610 SSDs which were used as block devices for NVMe-oF subsystems and two 100GbE Mellanox ConnectX®-5 NICs connected to provide up to 200GbE of network bandwidth. Each Initiator system has one Mellanox ConnectX®-5 100GbE NIC connected directly to the target without any switch.

One goal of this report was to make clear 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 libaio 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).

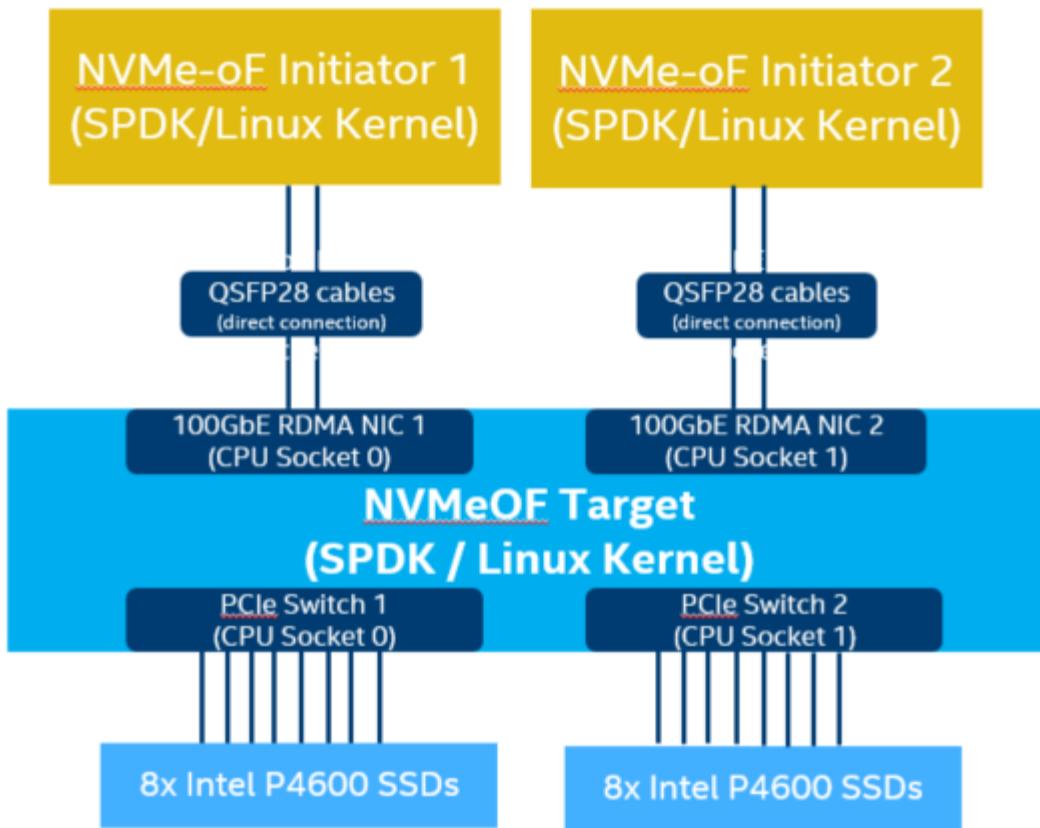


Figure 1: High-Level NVMe-oF TCP performance testing setup



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 16 NVMe-oF subsystems. Each NVMe-oF subsystem ran on top of an individual NVMe bdev backed by a single Intel P4610 device. Each of the 2 host systems was connected to 8 NVMe-oF subsystems which were exported by the SPDK NVMe-oF Target over 1x 100GbE NIC. The SPDK bdev FIO plugin was used to target 8 NVMe-oF bdevs on each of the host. The SPDK Target was configured to use 1, 4, 8, 12, 16, 24, 32 and 40 CPU cores. We ran the following workloads on each initiator:

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

We scaled the fio jobs using fio parameter numjob=3 in order to generate more I/O requests. 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 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 (4KB rand write and 4KB 70% read 30% write we ran preconditioning once before running all of the workload to ensure that NVMe devices reached higher IOPS so that we can saturate the network.

Item	Description
Test Case	Test SPDK NVMe-oF Target I/O core scaling
SPDK NVMe-oF Target configuration	All of the commands below were executed with spdk/scripts/rpc.py script. <pre>construct_nvme_bdev -t PCIe -b Nvme0 -a 0000:60:00.0 construct_nvme_bdev -t PCIe -b Nvme1 -a 0000:61:00.0 construct_nvme_bdev -t PCIe -b Nvme2 -a 0000:62:00.0 construct_nvme_bdev -t PCIe -b Nvme3 -a 0000:63:00.0 construct_nvme_bdev -t PCIe -b Nvme4 -a 0000:64:00.0 construct_nvme_bdev -t PCIe -b Nvme5 -a 0000:65:00.0 construct_nvme_bdev -t PCIe -b Nvme6 -a 0000:66:00.0 construct_nvme_bdev -t PCIe -b Nvme7 -a 0000:67:00.0 construct_nvme_bdev -t PCIe -b Nvme8 -a 0000:b5:00.0</pre>

```

construct_nvme_bdev -t PCIe -b Nvme9 -a 0000:b6:00.0
construct_nvme_bdev -t PCIe -b Nvme10 -a 0000:b7:00.0
construct_nvme_bdev -t PCIe -b Nvme11 -a 0000:b8:00.0
construct_nvme_bdev -t PCIe -b Nvme12 -a 0000:b9:00.0
construct_nvme_bdev -t PCIe -b Nvme13 -a 0000:ba:00.0
construct_nvme_bdev -t PCIe -b Nvme14 -a 0000:bb:00.0
construct_nvme_bdev -t PCIe -b Nvme15 -a 0000:bc:00.0

nvmf_create_transport -t TCP
(creates TCP transport layer with default values:
trtype: "TCP"
max_queue_depth: 128
max_qpairs_per_ctrlr: 64
in_capsule_data_size: 4096
max_io_size: 131072
"io_unit_size": 131072,
max_aq_depth: 128
num_shared_buffers: 4096
buf_cache_size: 32,
"c2h_success": true,
"dif_insert_or_strip": false,
"sock_priority": 0
)

nvmf_subsystem_create nqn.2018-09.io.spdk:cnode1 -s SPDK001 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode2 -s SPDK002 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode3 -s SPDK003 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode4 -s SPDK004 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode5 -s SPDK005 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode6 -s SPDK006 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode7 -s SPDK007 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode8 -s SPDK008 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode9 -s SPDK009 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode10 -s SPDK0010 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode11 -s SPDK0011 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode12 -s SPDK0012 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode13 -s SPDK0013 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode14 -s SPDK0014 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode15 -s SPDK0015 -a -m 8
nvmf_subsystem_create nqn.2018-09.io.spdk:cnode16 -s SPDK0016 -a -m 8

nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode1 Nvme0n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode2 Nvme1n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode3 Nvme2n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode4 Nvme3n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode5 Nvme4n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode6 Nvme5n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode7 Nvme6n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode8 Nvme7n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode9 Nvme8n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode10 Nvme9n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode11 Nvme10n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode12 Nvme11n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode13 Nvme12n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode14 Nvme13n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode15 Nvme14n1
nvmf_subsystem_add_ns nqn.2018-09.io.spdk:cnode16 Nvme15n1

nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode1 -t tcp -f ipv4 -s 4420 -a 20.0.0.1
nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode2 -t tcp -f ipv4 -s 4420 -a 20.0.0.1
nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode3 -t tcp -f ipv4 -s 4420 -a 20.0.0.1

```



	<pre> nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode4 -t tcp -f ipv4 -s 4420 -a 20.0.0.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode5 -t tcp -f ipv4 -s 4420 -a 20.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode6 -t tcp -f ipv4 -s 4420 -a 20.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode7 -t tcp -f ipv4 -s 4420 -a 20.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode8 -t tcp -f ipv4 -s 4420 -a 20.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode9 -t tcp -f ipv4 -s 4420 -a 10.0.0.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode10 -t tcp -f ipv4 -s 4420 -a 10.0.0.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode11 -t tcp -f ipv4 -s 4420 -a 10.0.0.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode12 -t tcp -f ipv4 -s 4420 -a 10.0.0.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode13 -t tcp -f ipv4 -s 4420 -a 10.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode14 -t tcp -f ipv4 -s 4420 -a 10.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode15 -t tcp -f ipv4 -s 4420 -a 10.0.1.1 nvmf_subsystem_add_listener nqn.2018-09.io.spdk:cnode16 -t tcp -f ipv4 -s 4420 -a 10.0.1.1 </pre>
SPDK NVMe-oF Initiator - FIO plugin configuration	<p>BDEV.conf</p> <p>[Nvme]</p> <pre> TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode1" Nvme0 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode2" Nvme1 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode3" Nvme2 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode4" Nvme3 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode5" Nvme4 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode6" Nvme5 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode7" Nvme6 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode8" Nvme7 </pre> <p>FIO.conf</p> <p>[global]</p> <pre> 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, 64, 128, 192, 256} time_based=1 numjobs=3 ramp_time=60 runtime=300 [filename0] filename=NvmeOn1 [filename1] filename=Nvme1n1 [filename2] filename=Nvme2n1 [filename3] filename=Nvme3n1 [filename4] filename=Nvme4n1 [filename5] filename=Nvme5n1 [filename6] filename=Nvme6n1 [filename7] filename=Nvme7n1 </pre>

4KB Random Read Results

Test Result: 4KB 100% Random Read IOPS, QD=128

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	1722.00	440.8	4681.6
4 cores	7229.05	1850.6	1114.3
8 cores	13540.35	3466.3	594.6
12 cores	17537.77	4489.7	458.1
16 cores	19215.89	4919.3	417.8
24 cores	19980.47	5115.0	401.9
32 cores	20263.73	5187.5	396.7
40 cores	21009.05	5378.3	382.1

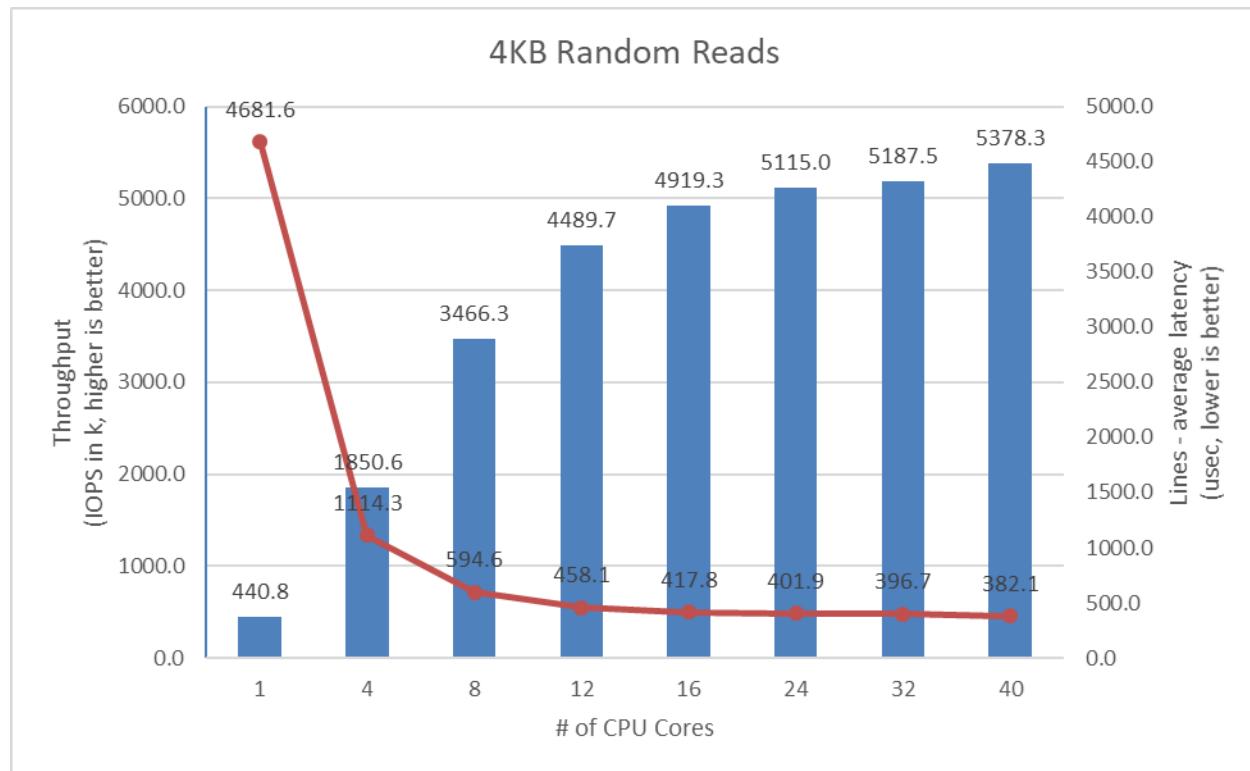


Figure 2: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Read workload at QD = 128

4KB Random Write Results

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

Test Result: 4KB 100% Random Writes IOPS, QD=128

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	1077.8	275.9	11182.0
4 cores	4411.8	1129.4	2718.0
8 cores	8374.6	2143.9	1403.9
12 cores	11160.5	2857.1	1053.8
16 cores	13097.6	3353.0	888.6
24 cores	14936.3	3823.7	780.3
32 cores	15727.2	4026.1	746.4
40 cores	16347.9	4185.0	719.1

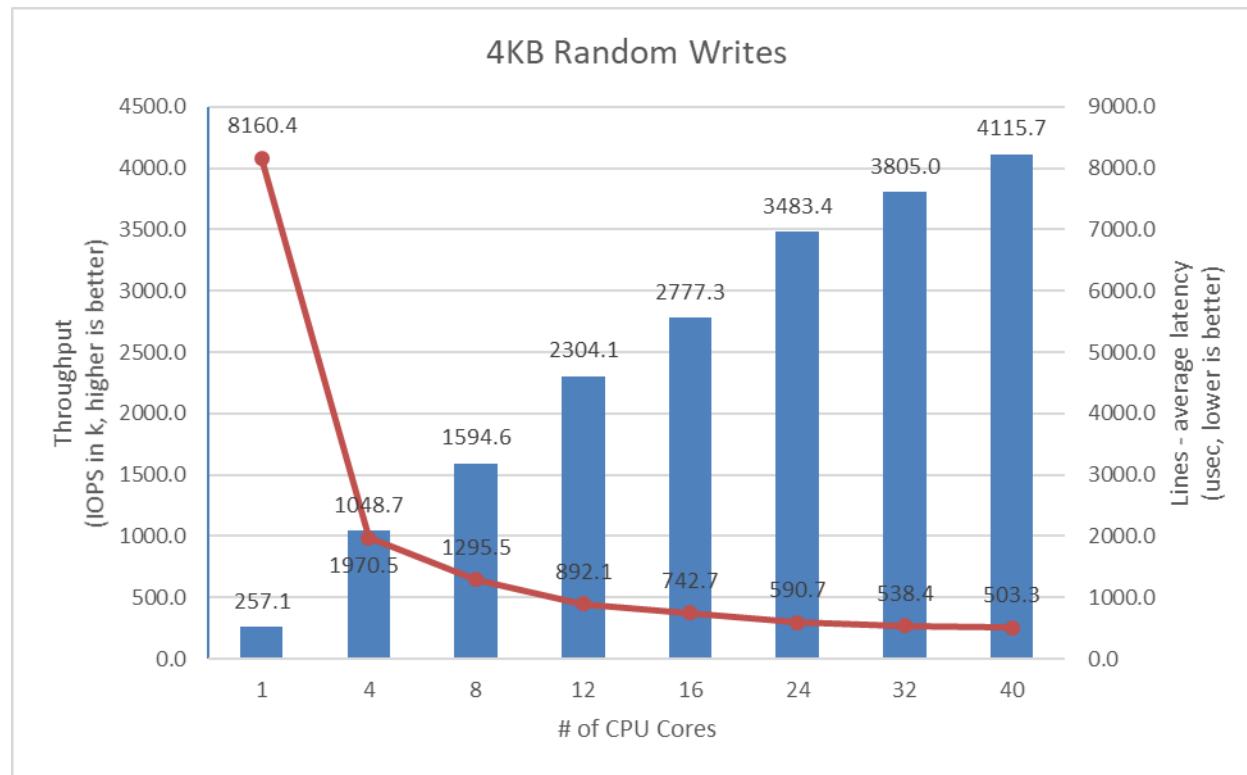


Figure 3: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Write Workload at QD=128

4KB Random Read-Write Results

Test Result: 4KB Random Read/Write 70%/30% IOPS, QD=192

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	1392.7	356.5	8612.4
4 cores	5091.1	1303.3	2353.3
8 cores	10836.6	2774.2	1104.3
12 cores	14411.7	3689.4	827.3
16 cores	16953.5	4340.1	703.2
24 cores	19527.6	4999.1	609.5
32 cores	19310.5	4943.5	619.9
40 cores	20285.5	5193.1	586.9

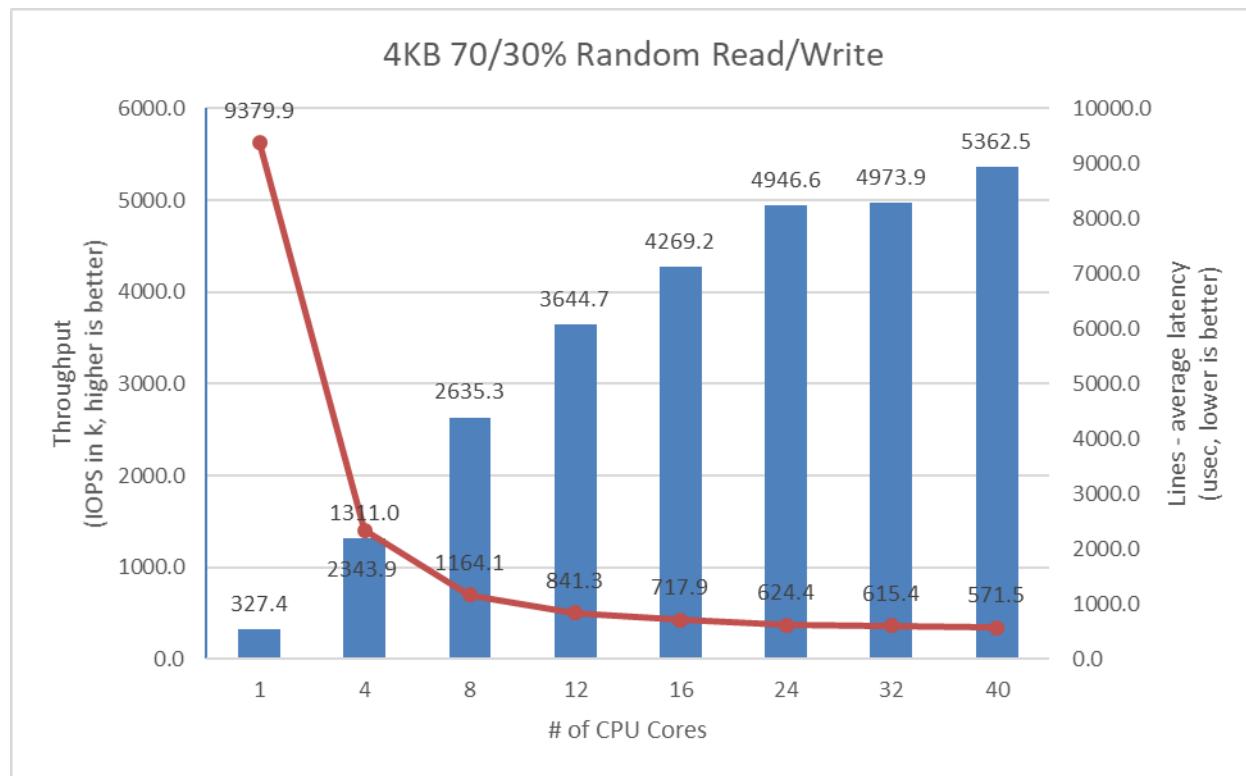


Figure 4: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 4KB Random 70/30 Read/Write workload at QD=192

Large Sequential I/O Performance

We measured the performance of large block I/O workloads by performing sequential I/Os of size 128KBs at queue depth 8. We used iodepth=8 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 4KB test case in the previous part of this document.

Test Result: 128KB 100% Sequential Reads, QD=8

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	14196.15	113.6	2280.6
4 cores	21375.31	171.0	1508.6
8 cores	21215.90	169.7	1520.2
12 cores	21208.85	169.7	1521.1
16 cores	20992.81	167.9	1537.5
24 cores	21246.98	170.0	1518.0
32 cores	20929.98	167.4	1543.4
40 cores	21041.93	168.3	1533.4



Figure 5: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KB 100% Sequential Read Workload at QD=8 and initiator FIO numjobs=2

Test Result: 128KB 100% Sequential Writes, QD=8

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	1901.39	15.2	16860.5
4 cores	7539.40	60.3	4250.2
8 cores	13388.85	107.1	2414.8
12 cores	17192.48	137.5	1867.9
16 cores	17102.31	136.8	1891.1
24 cores	18852.72	150.8	1702.0
32 cores	18129.80	145.0	1782.7
40 cores	17156.05	137.2	1882.7

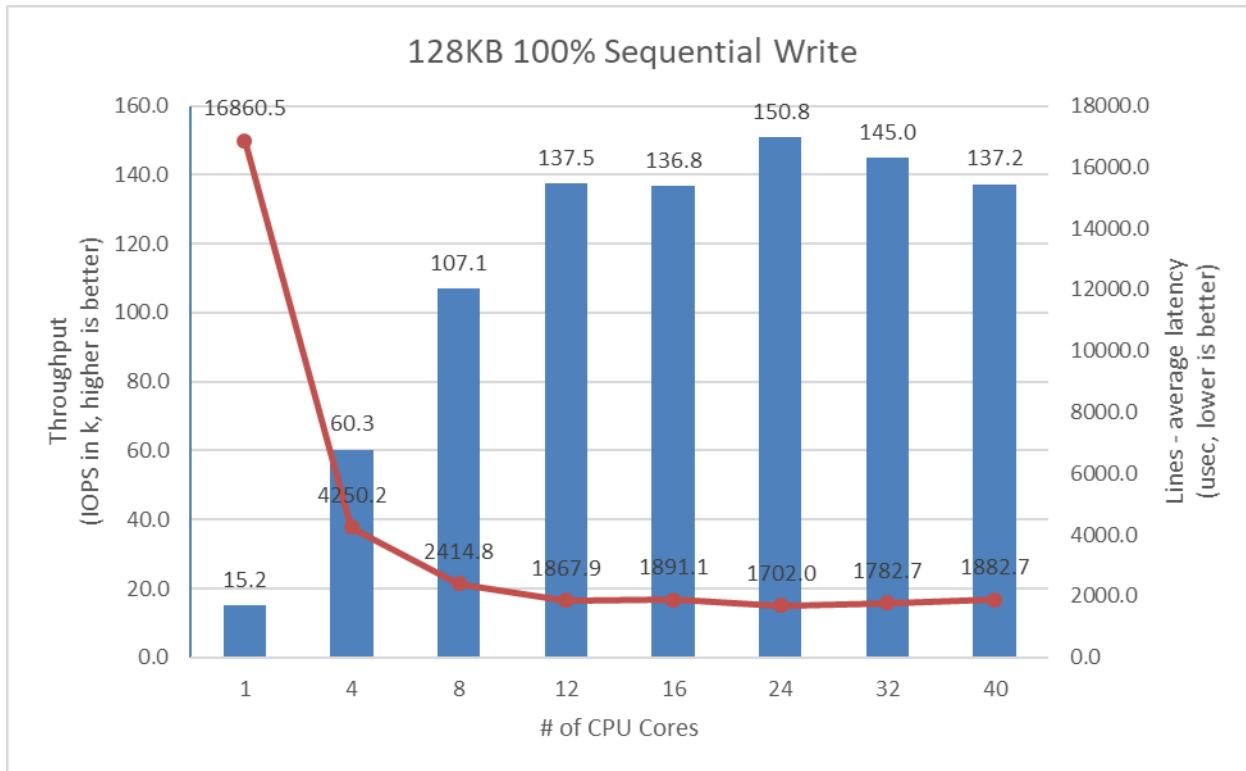


Figure 6: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KB 100% Sequential Write Workload at QD=8 and Initiator FIO numjobs=2

Test Result: 128KB Sequential 70% Reads 30% Writes, QD=8

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	4734.76	37.9	6805.6
4 cores	15219.84	121.8	2106.5
8 cores	21005.72	168.0	1527.4
12 cores	21763.08	174.1	1498.6
16 cores	23489.04	187.9	1379.2
24 cores	22821.30	182.6	1429.7
32 cores	28165.54	225.3	1135.9
40 cores	26099.50	208.8	1234.5

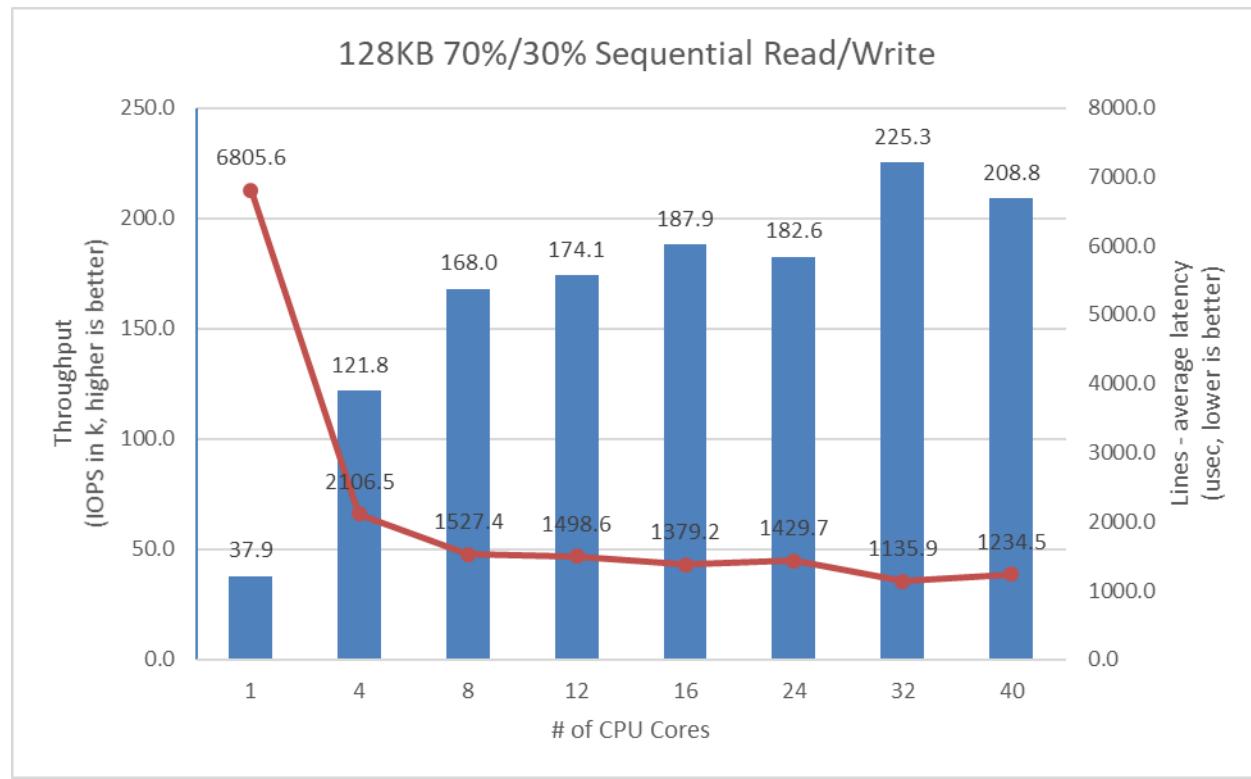


Figure 7: SPDK NVMe-oF TCP Target I/O core scaling: IOPS vs. Latency while running 128KB Sequential 70% Read 30% Write Workload at QD=8 and Initiator FIO numjobs=2

Conclusions

1. The SPDK NVMe-oF TCP Target IOPS throughput scales up linearly with addition of CPU cores for 4KB Random workload up to 100 GbE network speeds. The target uses just 8 CPU cores to serve at 100Gbps for the 4KB Rand Reads and 4KB 70/30 Rand Read/Write workloads. Beyond 100 Gbps the IOPS gains that are non-linear. The peak performance of 5.1 million IOPS is reached at 24 CPU cores, which is close to saturating 200Gb network link.
2. The target uses up to 16 CPU cores to serve 4KB Random Writes at 100 Gbps and the maximum performance observed was 4.1 million IOPS at 40 CPUs. There was no steady state or saturation reached during this test because we did not test with more than 40 CPU cores.
3. The best trade-off between CPU efficiency and network saturation was observed when the Target was configured with 16 CPU cores. The performance we achieved fully saturated a 100Gbps NIC connection between Target and Initiator for all tested workloads.
4. For the 4KB Random Write workload, we saturated the NVMe drives, which if preconditioned, would max out at about 3.2 million IOPS. Not preconditioning the drives allowed us to artificially increase their throughput and serve more IO requests than usual.
5. 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 171 Gbps was observed at 4 CPU cores. For the 128K Sequential Writes, the peak throughput of 150 Gbps was observed at 24 CPU cores. For the 128K Sequential 70/30 Read/Write workload the peak throughput of 225 Gbps was observed at 32 CPU cores.
6. Benefits of recently introduced MSG_ZEROCOPY are visible in case of 128K Sequential Read workload where we were able to reach 65% of line rate with a single CPU core.



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 100 Gbps network link.

The SPDK NVMe-oF TCP Target was configured similarly as in test case 1, using 20 cores. We used 20 CPU cores based on results of the previous test case which show that the target can easily serve over 3 million IOPS, that is enough IOPS to saturate 100 Gbps network connection

The SPDK bdev FIO plugin was used to target 16 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 40 CPUs. For detailed FIO job configuration see table below. FIO was run with following workloads:

- 4KB 100% Random Read
- 4KB 100% Random Write
- 4KB 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 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.

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 20 CPU cores.
SPDK NVMe-oF Initiator 1 - FIO plugin configuration	BDEV.conf [Nvme] TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode1" Nvme0 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode2" Nvme1 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode3" Nvme2 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode4" Nvme3 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode5" Nvme4 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode6" Nvme5 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode7" Nvme6 TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode8" Nvme7 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode1" Nvme8 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode2" Nvme9 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode3" Nvme10 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode4" Nvme11 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode5" Nvme12

	<pre> TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode6" Nvme13 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode7" Nvme14 TransportId "trtype:TCP adrfam:IPv4 traddr:10.0.1.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode8" Nvme15 </pre> <p>FIO.conf</p> <p>For 1 CPU initiator configuration:</p> <pre> [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={1,32, 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 filename=Nvme14n1 filename=Nvme15n1 </pre>
	<p>FIO.conf</p> <p>For X*4 CPU (up to 40) initiator configuration:</p> <pre> [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, 32, 64, 128, 192} time_based=1 ramp_time=60 runtime=300 numjobs=X </pre>



```
[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
filename=Nvme11n1
[filename3]
filename=Nvme12n1
filename=Nvme13n1
filename=Nvme14n1
filename=Nvme15n1
```

4KB Random Read Results

Test Result: 4KB 100% Random Read, QD=128

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	969.40	248.2	4115.4
4 cores	3918.60	1003.2	1017.5
8 cores	6911.35	1769.3	574.8
12 cores	10029.83	2567.6	395.4
16 cores	10566.34	2705.0	375.1
24 cores	9089.45	2326.9	441.6
32 cores	8643.55	2212.7	461.3
40 cores	7802.50	1997.4	519.2

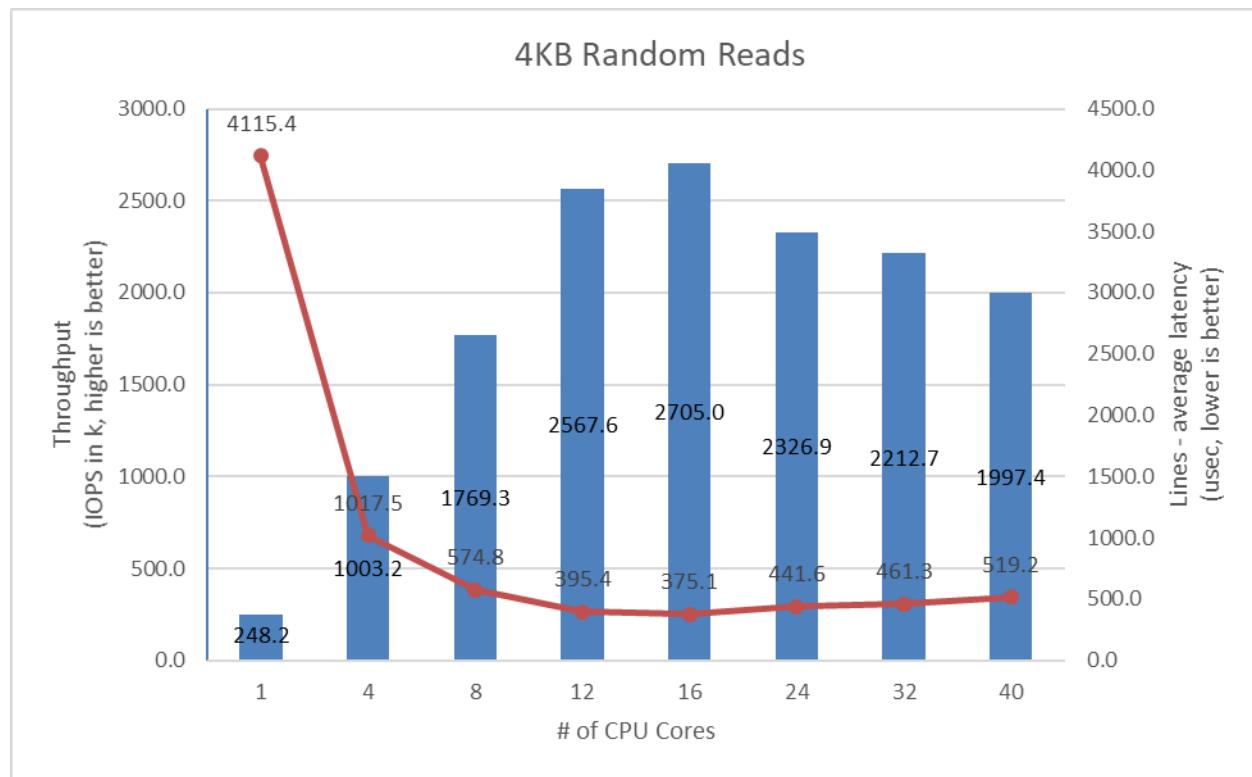


Figure 8: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Read QD=128 workload

4KB Random Write Results

Test Result: 4KB 100% Random Write, QD=128

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	2132.44	545.9	1162.1
4 cores	6415.37	1642.3	596.9
8 cores	8142.53	2084.5	492.5
12 cores	8252.47	2112.6	490.5
16 cores	7743.54	1982.3	514.5
24 cores	7792.53	1994.9	515.7
32 cores	7214.93	1847.0	553.9
40 cores	7180.27	1838.1	564.7



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

4KB Random Read-Write Results

Test Result: 4KB 70% Random Read 30% Random Write, QD=128

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	1112.96	284.9	3574.3
4 cores	4792.10	1226.8	830.1
8 cores	8457.04	2165.0	466.6
12 cores	11378.04	2912.8	347.7
16 cores	11212.10	2870.3	353.7
24 cores	10726.29	2745.9	373.5
32 cores	10465.16	2679.1	382.4
40 cores	9994.99	2558.7	405.1

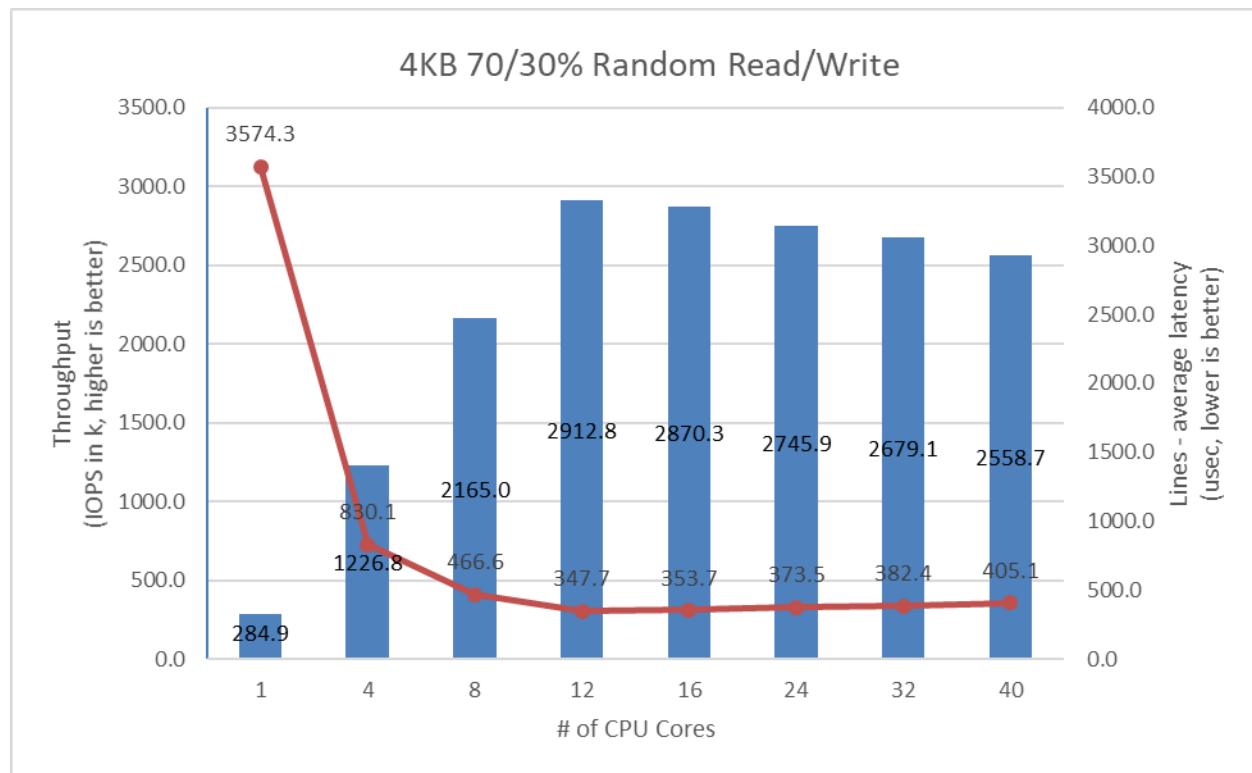
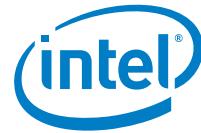


Figure 10: SPDK NVMe-oF TCP Initiator I/O core scaling: IOPS vs. Latency while running 4KB Random 70% Read 30% Write Workload at QD=128



Conclusions

1. For the 4KB Random Read workload, the SPDK NVMe-oF TCP Initiator performance scales linearly up to 12 CPU cores. The peak performance of 2.7 million IOPS was reached at 16 CPU cores, which saturates 100Gb link. Increasing the number of CPU cores beyond 16 CPU cores results in performance degradation.
2. In case of 4KB Random Write workload, performance scales linearly up to 8 CPU cores, reaching peak performance of 2.1M IOPS. Increasing the number of Initiator cores beyond 8 cores does not improve results.
3. Mixed Random Read-Write workload performance scales linearly up to 12 CPU cores, reaching 2.9M IOPS and saturating the network link.

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.

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 spdk/scripts/rpc.py script.</p> <pre>nvmf_create_transport -t TCP (creates TCP transport layer with default values: trtype: "TCP" max_queue_depth: 128 max_qpairs_per_ctrlr: 64 in_capsule_data_size: 4096 max_io_size: 131072 io_unit_size: 8192 max_aq_depth: 128 num_shared_buffers: 4096 buf_cache_size: 32) construct_null_bdev 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 nvmet-cl tool.</p> <pre>{ "ports": [{ "addr": { "adrfam": "ipv4", "traddr": "20.0.0.1", "trsvid": "4420", "trtype": "tcp" }, "portid": 1, "referrals": [], "subsystems": ["nqn.2018-09.io.spdk:cnode1"] }], "hosts": [], "subsystems": [{ "allowed_hosts": [], "attr": { "allow_anonymous": 1 } }] }</pre>



	<pre> "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>
FIO configuration	
SPDK NVMe-oF Initiator FIO plugin configuration	<p>BDEV.conf</p> <p>[Nvme]</p> <pre>TransportId "trtype:TCP adrfam:IPv4 traddr:20.0.0.1 trsvcid:4420 subnqn:nqn.2018-09.io.spdk:cnode1" Nvme0</pre> <p>FIO.conf</p> <p>[global]</p> <pre>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</pre> <p>[filename0]</p> <pre>filename=Nvme0n1</pre>
Kernel initiator configuration	<p>Device config</p> <p>Done 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</pre> <p>FIO.conf</p> <p>[global]</p> <pre>ioengine=libaio 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</pre>



	[filename0] filename=/dev/nvme0n1
--	--------------------------------------

SPDK vs Kernel NVMe-oF Target Latency Results

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

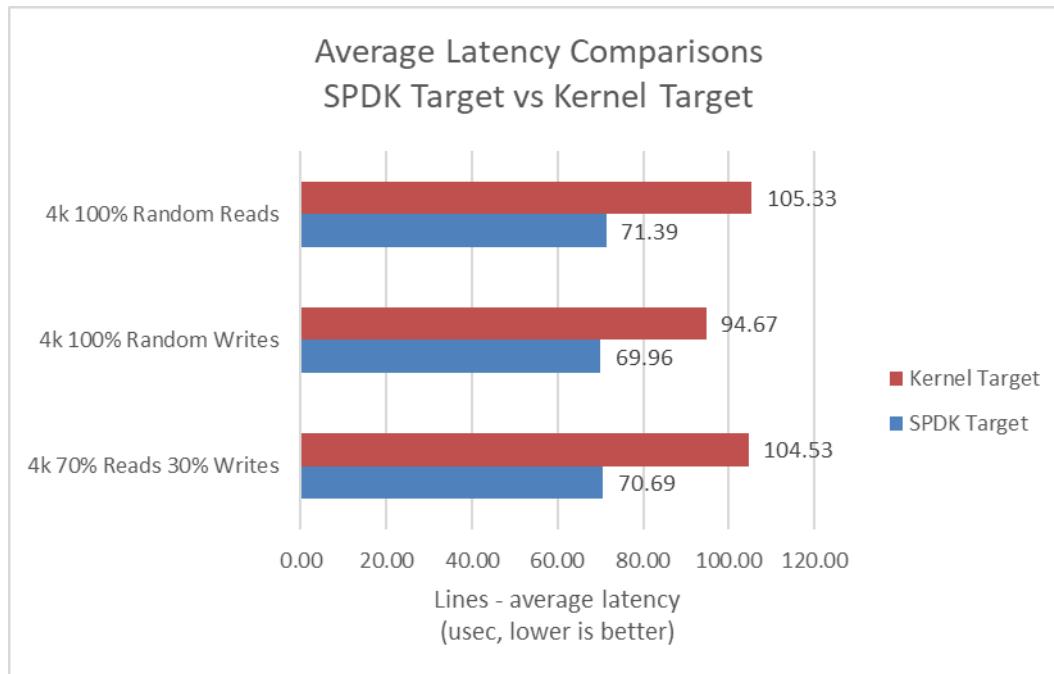


Figure 11: SPDK vs. Kernel NVMe-oF TCP Average I/O Latency for various workloads run using the Kernel Initiator

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)
4KB 100% Random Reads IOPS	71.39	13631	61.0	74.2	163.5	210.6
4KB 100% Random Writes IOPS	69.96	13918	58.6	111.4	170.3	224.9
4KB 100% Random 70% Reads 30% Writes IOPS	70.69	13748	62.0	149.9	193.1	249.9

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)
4KB 100% Random Reads IOPS	105.33	9328	162.1	176.5	241.3	1034.9
4KB 100% Random Writes IOPS	94.67	10366	88.2	113.2	198.3	770.7
4KB 100% Random 70% Reads 30% Writes IOPS	104.53	9393	144.0	194.5	232.8	329.3

SPDK vs Kernel NVMe-oF TCP Initiator results

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

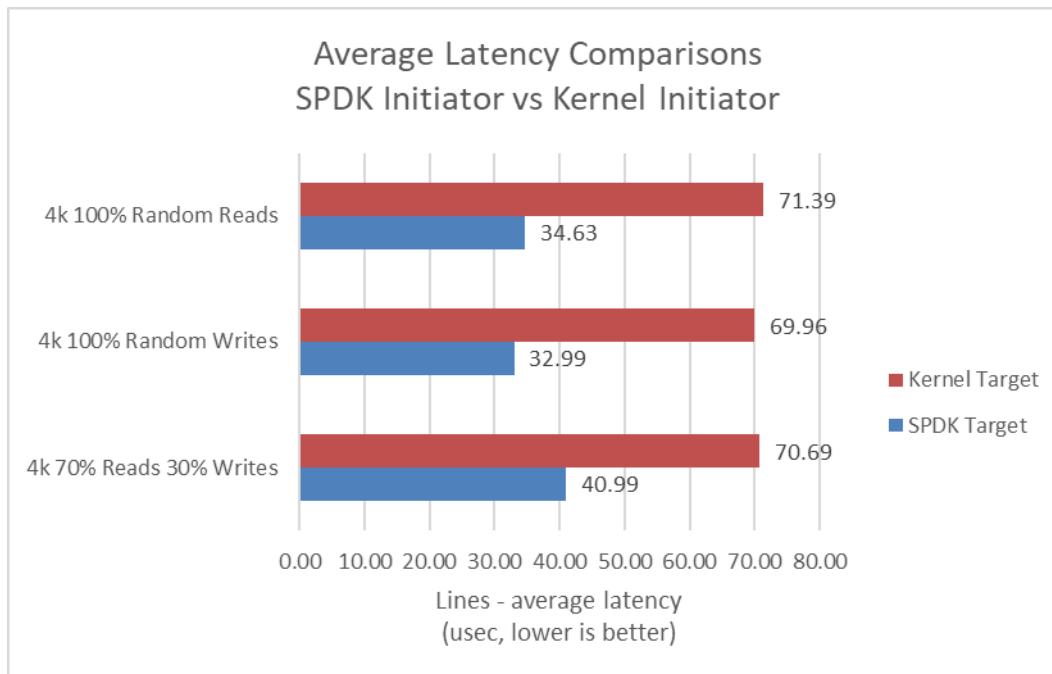


Figure 12: SPDK vs. Kernel NVMe-oF TCP Average I/O Latency for various workloads against SPDK Target

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)
4KB 100% Random Reads IOPS	34.63	28741	50.8	59.0	72.9	169.0
4KB 100% Random Writes IOPS	32.99	30162	44.6	57.8	71.9	171.0
4KB 100% Random 70% Reads 30% Writes IOPS	40.99	24308	84.7	117.5	132.4	205.7

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)
4KB 100% Random Reads IOPS	71.39	13631	61.0	74.2	163.5	210.6
4KB 100% Random Writes IOPS	69.96	13918	58.6	111.4	170.3	224.9
4KB 100% Random 70% Reads 30% Writes IOPS	70.69	13748	62.0	149.9	193.1	249.9



SPDK vs Kernel NVMe-oF Latency Results

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

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)
4KB 100% Random Reads IOPS	34.63	28741	50.8	59.0	72.9	169.0
4KB 100% Random Writes IOPS	32.99	30162	44.6	57.8	71.9	171.0
4KB 100% Random 70% Reads 30% Writes IOPS	40.99	24308	84.7	117.5	132.4	205.7

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)
4KB 100% Random Reads IOPS	105.33	9328	162.13	176.47	241.32	1034.92
4KB 100% Random Writes IOPS	94.67	10366	88.24	113.15	198.32	770.73
4KB 100% Random 70% Reads 30% Writes IOPS	104.53	9393	143.99	194.49	232.82	329.25

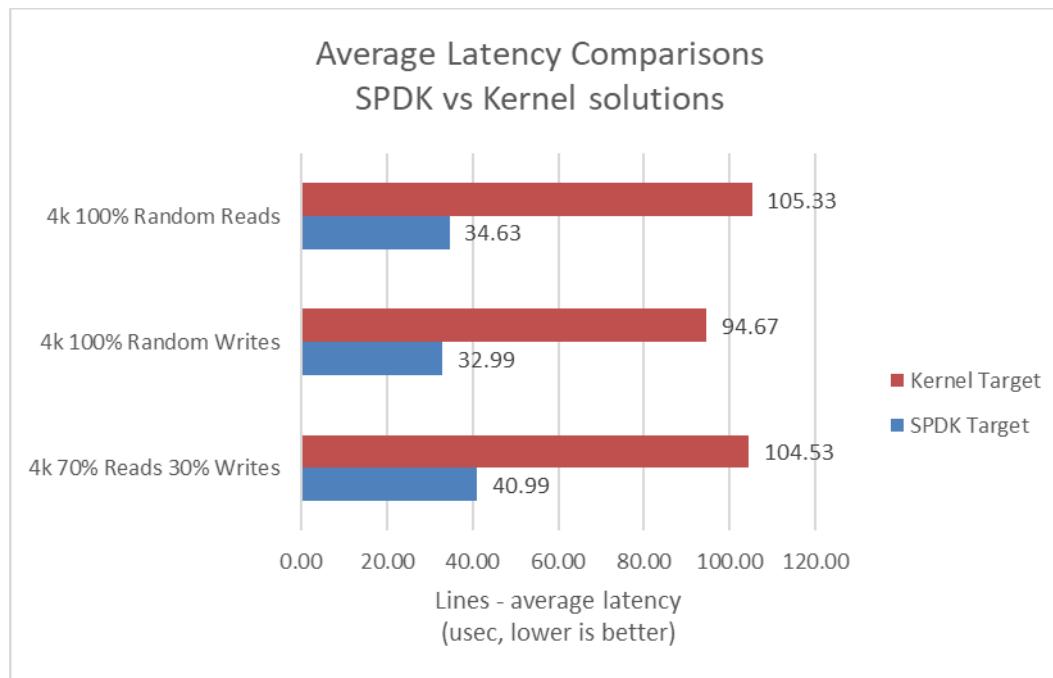


Figure 13: SPDK vs. Kernel NVMe-oF TCP solutions Average I/O Latency for various workloads

Conclusions

1. SPDK NVMe-oF Initiator reduces the average latency by up to 37 usec vs. the Linux Kernel NVMe-oF Initiator, which eliminates up to 50% NVMe-oF software overhead.
2. The SPDK NVMe-oF Target reduces the NVMe-oF average round trip I/O latency (reads/writes) by up to 34 usec vs. the Linux Kernel NVMe-oF target, which eliminates up to 33% NVMe-oF software overhead.
3. The SPDK NVMe-oF TCP target and initiator reduced the average latency by up to 67% vs. the Linux Kernel NVMe-oF target and initiator.
4. The SPDK NVMe-oF TCP target reduces tail latencies (99th percentile) by over 60% for the 4KB Random Read and 4KB Random Read/Write 70/30 workloads.
5. The SPDK NVMe-oF Initiator reduces the p99 latency by 16% and 24% for the 4KB random reads and write workloads respectively.



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 30 cores, 16 NVMe-oF subsystems (1 per Intel P4600) and 2 initiators were used both running I/Os to 8 separate subsystems using Kernel NVMe-oF initiator. We ran the following workloads on the host systems:

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

Item	Description
Test Case	NVMe-oF Target performance under varying # of connections
SPDK NVMe-oF Target configuration	Same as in Test Case #1, using 30 CPU cores.
Kernel NVMe-oF Target configuration	Target configuration file loaded using nvmet-clt tool. For detail configuration file contents please see Appendix A.
Kernel NVMe-oF Initiator #1	Device config Performed using nvme-clt tool. 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 nvme connect -n nqn.2018-09.io.spdk:cnode8 -t tcp -a 20.0.1.1 -s 4420
Kernel NVMe-oF Initiator #2	Device config Performed using nvme-clt tool. modprobe nvme-fabrics 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 10.0.0.1 -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 nvme connect -n nqn.2018-09.io.spdk:cnode15 -t tcp -a 10.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode16 -t tcp -a 10.0.1.1 -s 4420
FIO configuration (used on both initiators)	FIO.conf [global]

	ioengine=libaio thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth={8, 32, 64, 128} time_based=1 ramp_time=60 runtime=300 numjobs={1, 4, 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 [filename8] filename=/dev/nvme7n1
--	--

The number of CPU cores used while running the SPDK NVMe-oF target was 30, whereas for the case of Linux Kernel NVMe-oF target there was no CPU core limitation applied.

The metrics in the graph represent relative performance 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.

4KB Random Read Results

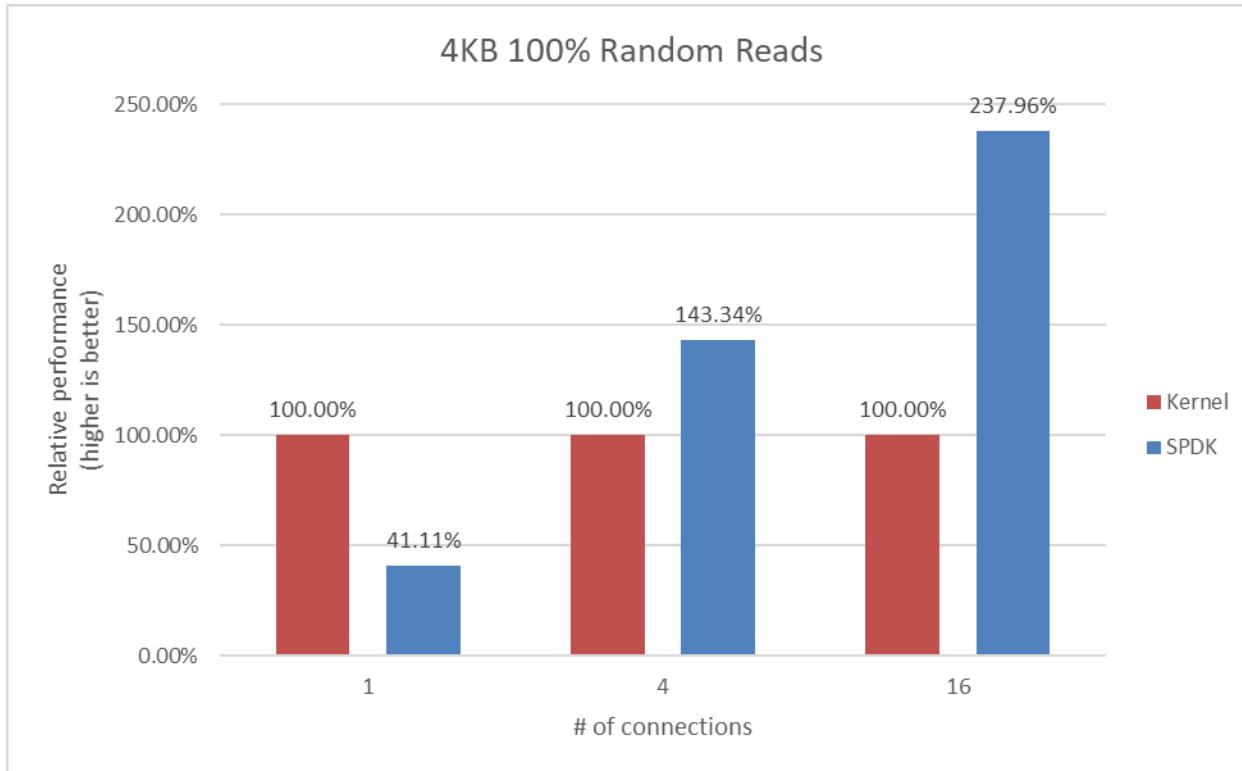


Figure 14: Relative Performance Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KB 100% Random Reads using the Kernel Initiator

Linux Kernel NVMe-oF TCP Target: 4KB 100% Random Reads, QD=64

Connections per subsystem	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	2798.71	716.5	2859.4	11.5
4	7632.99	1954.0	1050.1	39.9
16	8164.88	2090.2	976.6	59.0

SPDK NVMe-oF TCP Target: 4KB 100% Random Reads, QD=64

Connections per subsystem	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	3008.21	770.1	2658.6	30.0
4	8232.06	2107.4	972.2	30.0
16	9877.27	2528.6	807.3	30.0

4KB Random Write Results

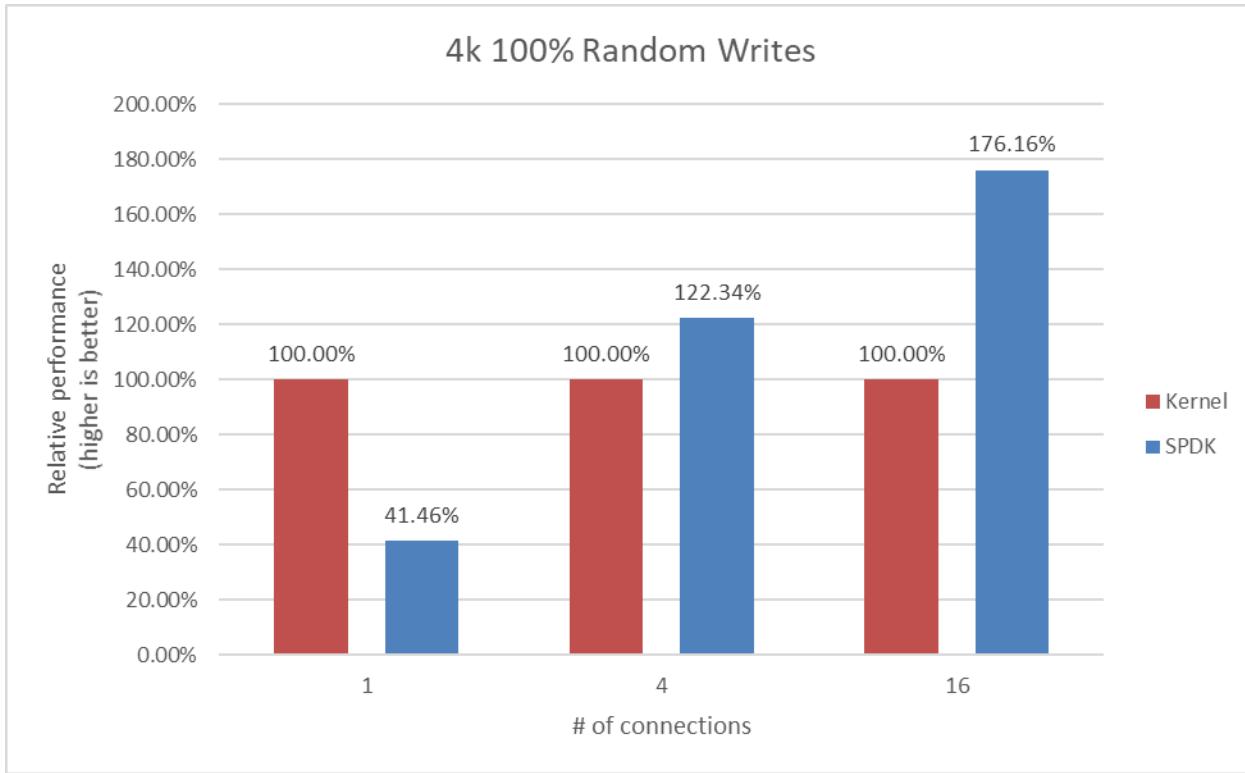


Figure 15: Relative Performance Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KB 100% Random Writes Workload

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

Linux Kernel NVMe-oF TCP Target: 4KB 100% Random Writes, QD=64

Connections per subsystem	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	3441.44	881.0	2327.8	12.4
4	7904.98	2023.7	1009.2	40.0
16	9693.40	2481.5	822.2	62.8

SPDK NVMe-oF TCP Target: 4KB 100% Random Writes, QD=64

Connections per subsystem	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	3448.95	882.9	2321.1	30.0
4	7255.16	1857.3	1099.5	30.0
16	8162.30	2089.5	976.2	30.0

4KB Random Read-Write Results

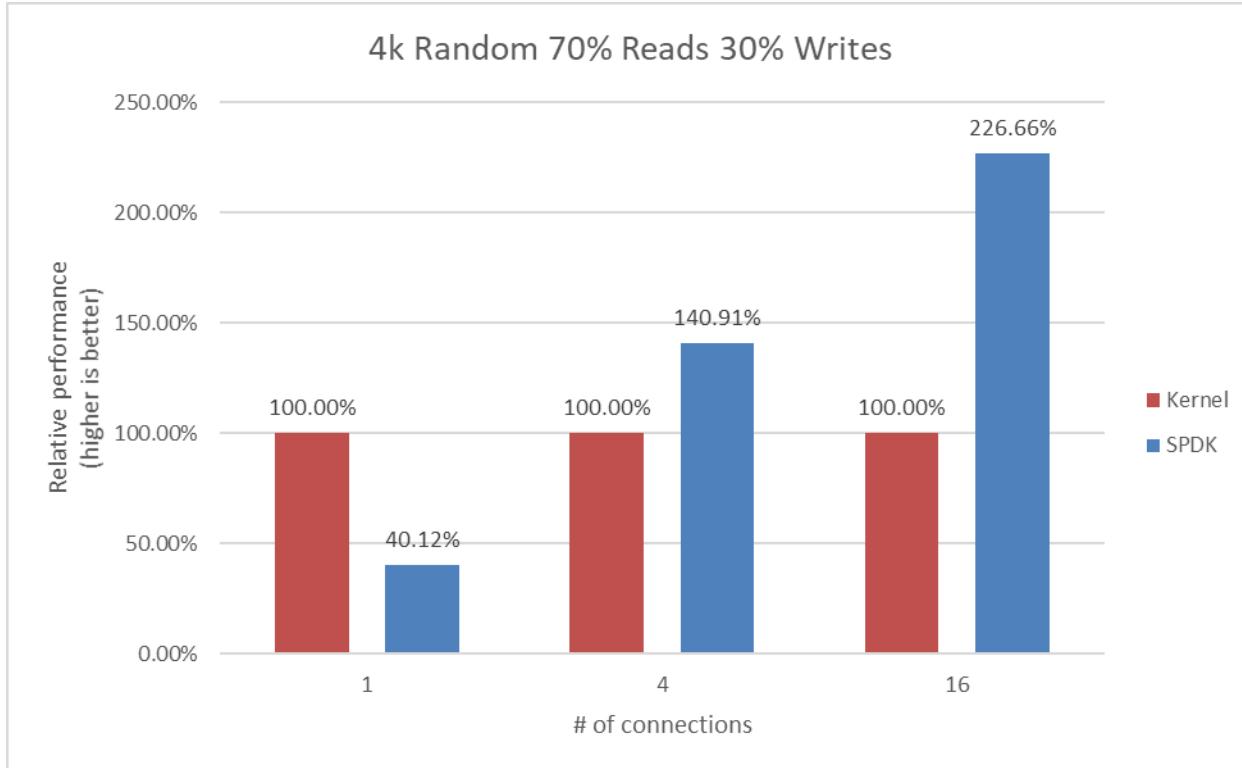


Figure 16: Relative Performance Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for 4KB Random 70% Reads 30% Writes Workload

Linux Kernel NVMe-oF TCP Target: 4KB 70% Random Read 30% Random Write, QD=64

Connections per subsystem	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	2824.84	723.2	2766.3	11.6
4	7144.76	1829.1	1073.9	39.7
16	8425.32	2156.9	886.0	61.2

SPDK NVMe-oF TCP Target: 4KB 70% Random Read 30% Random Write, QD=64

Connections per subsystem	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
1	2931.50	750.5	2638.5	30.0
4	7609.36	1948.0	1009.4	30.0
16	9363.87	2397.1	829.2	30.0

Low Connections Results

During testing it was observed that relative performance of SPDK Target is about 60-70% of Kernel Target performance 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.

SPDK & Kernel NVMe-oF TCP Target relative performance comparison for various workloads, QD=64, 1 connection per subsystem.

Workload	Target	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)	# CPU Cores
Random Read	Linux	2798.71	716.5	2859.4	11.5
	SPDK	3006.69	769.7	2660.6	4.0
Random Write	Linux	3441.44	881.0	2327.8	12.4
	SPDK	3395.05	869.1	2355.5	4.0
Random Read/Write	Linux	2824.84	723.2	2766.3	11.6
	SPDK	3034.55	776.8	2550.7	4.0

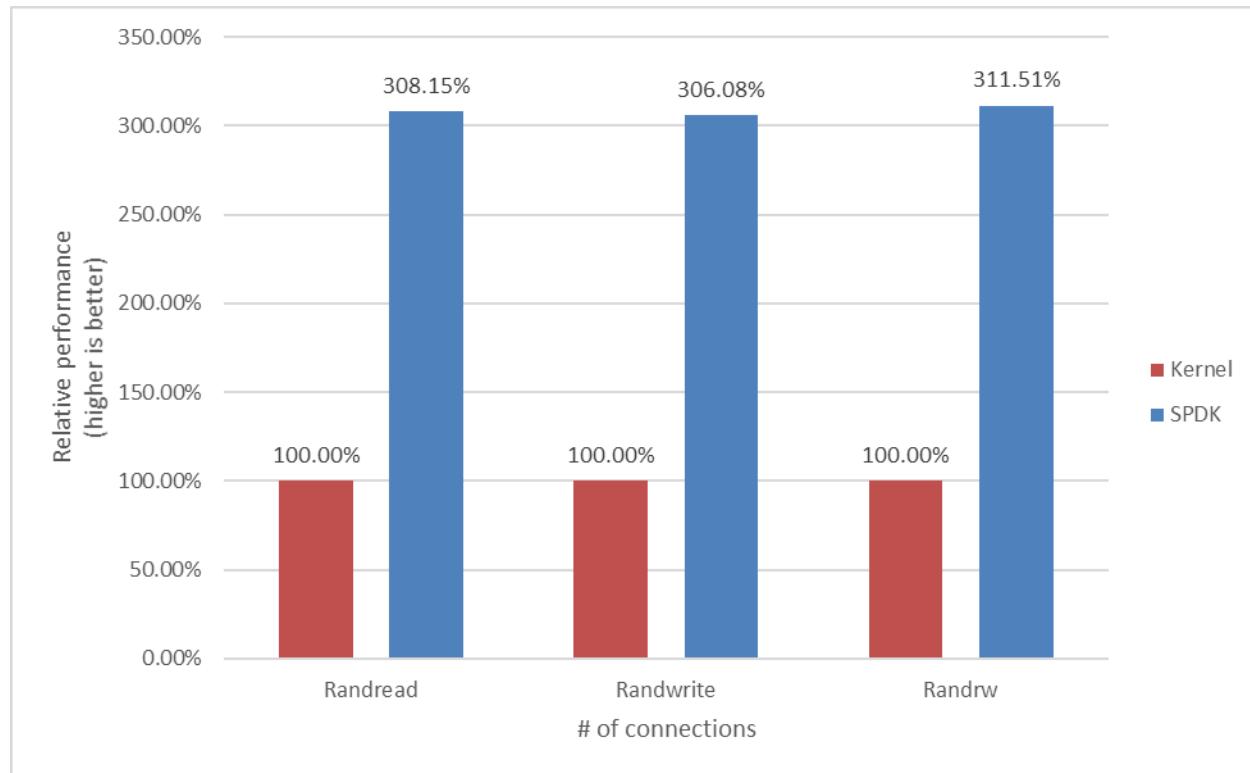


Figure 17: Relative Performance Comparison of Linux Kernel vs. SPDK NVMe-oF Target IOPS/Core for various workloads, 1 connection per subsystem and reduced number of SPDK Target CPU Cores (4)



Conclusions

1. The Linux Kernel NVMe-oF TCP target relative performance in IOPS/Core was better than SPDK when there was just one connection per subsystem because the SPDK NVMe-oF target uses a fixed number of CPU cores and there is no mechanism to dynamically decrease the number of CPU cores if the target does not need all the cores. Therefore,, we re-run the test cases with 1 connections 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 of about 3 times better than Linux Kernel NVMe-oF TCP target.
2. The SPDK NVMe-oF TCP target relative performance in IOPS/Core improved as we increased the number of connections per subsystem. It was between 1.2 - 1.4 times better than the Linux Kernel NVMe-oF target at 4 connections per NVMe-oF subsystem and up to 2.4 time better at 16 connections per NVMe-oF subsystem

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 4KB random workloads. The SPDK NVMe-oF target saturates a 100 Gbps network link using just 8 CPU cores for the 4KB Random Read and 4KB Random Read/Write 70/30 workloads, and 12 - 16 CPU cores for the 4KB Random Write workload. Beyond 100 Gbps, the IOPS scalability is non-linear for 4KB Random Read and 4KB Random Read/Write 70/30 workloads. It took 24 CPU cores for the target to almost saturate 200Gbps network link.

For the SDPK NVMe-oF TCP Initiator, the IOPS throughput scales almost linearly with addition of CPU cores until the network was almost saturated, however, as we got closer to network saturation it was observed that the throughput scaling becomes non-linear. A single initiator was able to saturate a 100Gb link.

For the NVMe-oF TCP latency comparison, the SPDK NVMe-oF Target and Initiator average latency is up to 33% and 50% lower than their Linux Kernel counterparts respectively when testing against null block device-based backend.

The SPDK NVMe-oF TCP Target performed up to 2.4 times better w.r.t IOPS/core than Linux Kernel NVMe-oF target while running 4KB 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.



Appendix A - Jumbo Frames Performance Impact

Test case 1 (SPDK NVMe-OF TCP Target Core Scaling test) was also re-run using 9000 Bytes MTU configured for RNIC interfaces to measure performance impact of using Jumbo Frames in the test setup.

Use of Jumbo Frames (MTU=9000B) has a negative impact on SPDK NVMe-OF performance, which is best visible on example of Random Write workload. Performance degradation introduced by Jumbo Frames was between 10-20% in that case.

4KB Random Read with JF results

# of Cores	Bandwidth (MBps)			Throughput (IOPS k)			Avg. Latency (usec)		
	1500MTU	9000MTU	Change [%]	1500MTU	9000MTU	Change [%]	1500MTU	9000MTU	Change [%]
1 core	1722.00	1662.86	-3.43%	440.82	425.68	-3.43%	4681.56	4847.09	3.54%
4 cores	7229.05	6732.66	-6.87%	1850.63	1723.55	-6.87%	1114.27	1197.86	7.50%
8 cores	13540.35	13137.49	-2.98%	3466.32	3363.19	-2.98%	594.64	612.21	2.95%
12 cores	17537.77	17120.39	-2.38%	4489.66	4382.81	-2.38%	458.08	469.78	2.55%
16 cores	19215.89	18636.64	-3.01%	4919.26	4770.97	-3.01%	417.80	431.19	3.21%
24 cores	19980.47	19937.97	-0.21%	5114.99	5104.11	-0.21%	401.89	402.93	0.26%
32 cores	20263.73	19931.43	-1.64%	5187.51	5102.44	-1.64%	396.72	403.41	1.69%
40 cores	21009.05	20458.83	-2.62%	5378.31	5237.45	-2.62%	382.07	393.39	2.96%

4KB Random Write with JF results

# of Cores	Bandwidth (MBps)			Throughput (IOPS k)			Avg. Latency (usec)		
	1500MTU	9000MTU	Change [%]	1500MTU	9000MTU	Change [%]	1500MTU	9000MTU	Change [%]
1 core	1004.40	885.18	-11.87%	257.12	226.60	-11.87%	8160.38	9240.99	13.24%
4 cores	4096.54	3536.73	-13.67%	1048.71	905.40	-13.67%	1970.51	2271.25	15.26%
8 cores	6229.13	6479.76	4.02%	1594.65	1658.81	4.02%	1295.48	1230.41	-5.02%
12 cores	9000.32	8307.36	-7.70%	2304.07	2126.68	-7.70%	892.13	966.93	8.38%
16 cores	10848.89	9497.90	-12.45%	2777.31	2431.45	-12.45%	742.65	843.98	13.64%
24 cores	13607.00	11329.58	-16.74%	3483.38	2900.36	-16.74%	590.66	707.74	19.82%
32 cores	14863.44	12437.63	-16.32%	3805.03	3184.03	-16.32%	538.37	647.46	20.26%
40 cores	16077.03	13657.01	-15.05%	4115.71	3496.19	-15.05%	503.30	588.33	16.89%

4KB Random Read-Write with JF results

# of Cores	Bandwidth (MBps)			Throughput (IOPS k)			Avg. Latency (usec)		
	1500MTU	9000MTU	Change [%]	1500MTU	9000MTU	Change [%]	1500MTU	9000MTU	Change [%]
1 core	1278.98	1180.41	-7.71%	327.41	302.17	-7.71%	9379.85	10163.71	8.36%
4 cores	5121.32	4855.00	-5.20%	1311.05	1242.87	-5.20%	2343.89	2470.77	5.41%
8 cores	10294.05	9153.94	-11.08%	2635.27	2343.40	-11.08%	1164.14	1309.28	12.47%
12 cores	14237.13	13082.18	-8.11%	3644.69	3349.03	-8.11%	841.27	916.45	8.94%
16 cores	16676.66	16021.43	-3.93%	4269.21	4101.47	-3.93%	717.86	747.33	4.11%
24 cores	19322.65	18605.51	-3.71%	4946.59	4763.00	-3.71%	624.36	644.44	3.22%
32 cores	19429.20	18204.58	-6.30%	4973.86	4660.36	-6.30%	615.41	660.05	7.25%
40 cores	20947.21	18460.95	-11.87%	5362.47	4725.99	-11.87%	571.53	650.50	13.82%



Appendix B – 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": []  
        }  
    ]  
}
```



```
"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"  
    }  
}
```



```
        "trtype": "tcp"
    },
    "portid": 13,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode13"
    ]
},
{
    "addr": {
        "adrfrm": "ipv4",
        "traddr": "10.0.1.1",
        "trsvcid": "4433",
        "trtype": "tcp"
    },
    "portid": 14,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode14"
    ]
},
{
    "addr": {
        "adrfrm": "ipv4",
        "traddr": "10.0.1.1",
        "trsvcid": "4434",
        "trtype": "tcp"
    },
    "portid": 15,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode15"
    ]
},
{
    "addr": {
        "adrfrm": "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": "33333333-3333-3333-3333-333333333333"  
            },  
            "enable": 1,  
            "nsid": 4  
        }  
    ],  
    "nqn": "nqn.2018-09.io.spdk:cnode4"  
},  
]
```

```
        "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-5cfb23c7a6eb"
      },
      "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"
},
{
  "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"
}
]
```



DISCLAIMERS

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

For more information go to <http://www.intel.com/performance>

Intel® AES-NI requires a computer system with an AES-NI enabled processor, as well as non-Intel software to execute the instructions in the correct sequence. AES-NI is available on select Intel® processors. For availability, consult your reseller or system manufacturer. **For more information, see <http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni/>**

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

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.

§