

SPDK NVMe-oF RDMA (Target & Initiator)

Performance Report

Release 21.04

Testing Date: May 2021

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)

Contents

Contents	2
Audience and Purpose.....	3
Test setup	4
Target Configuration.....	4
Initiator 1 Configuration.....	5
Initiator 2 Configuration.....	5
BIOS settings	6
Kernel & BIOS spectre-meltdown information	6
Introduction to SPDK NVMe-oF (Target & Initiator).....	7
Test Case 1: SPDK NVMe-oF RDMA Target I/O core scaling	9
4KB Random Read Results	12
4KB Random Write Results	13
4KB Random Read-Write Results.....	14
LTO Performance Impact.....	15
Large Sequential I/O Performance	18
Test Case 2: SPDK NVMe-oF RDMA Initiator I/O core scaling	20
4KB Random Read Results	23
4KB Random Write Results	24
4KB Random 70/30 Read/Write Results	25
Conclusions	26
Test Case 3: Linux Kernel vs. SPDK NVMe-oF RDMA Latency	27
SPDK vs Kernel NVMe-oF RDMA Target Results	30
SPDK vs Kernel NVMe-oF RDMA Initiator Results	31
SPDK vs Kernel NVMe-oF RDMA Kernel + Initiator Results.....	32
Conclusions	33
Test Case 4: NVMe-oF RDMA Performance with increasing # of connections	34
4KB Random Read Results	36
4KB Random Write results.....	37
4KB Random Read-Write Results.....	38
Conclusions	39
Summary	40
List of tables.....	41
List of figures.....	43
Appendix A – Test Case 1 & 2 SPDK NVMe-oF Initiator bdev configuration.....	44
Appendix B – Test Case 3 SPDK NVMe-oF Initiator bdev configuration	48
Appendix C - Kernel NVMe-oF RDMA Target configuration.....	48

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 contains performance and efficiency of the SPDK vs. Linux Kernel NVMe-oF Target and Initiator for the RDMA 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	<p>SuperMicro SYS-2029U-TN24R4T</p>   <p>Integrated Board Super X11DPU</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> <p>Microcode: 0x5003003</p>
Memory	<p>12 x 32GB Hynix HMA84GR7AFR4N-VK, DDR4, 2666MHz</p> <p>Total of 384GB</p>
Operating System	Fedora 33
BIOS	3.4
Linux kernel version	5.8.15-300.fc33.x86_64
SPDK version	SPDK 21.04
Storage	<p>OS: 1x 120GB Intel SSDSC2BB120G4</p> <p>Storage Target: 16x Intel® SSD DC P4610™ 1.6TB (FW: VDV10170) (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

Table 2: Hardware setup configuration – Initiator system 1

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) Microcode: 0x5003003
Memory	6 x 32GB Micron M393A1G40EB1-CRC, DDR4, 2933MHz Total 192GBs
Operating System	Fedora 33
BIOS	SE5C620.86B.02.01.0013.121520200651
Linux kernel version	5.8.15-300.fc33.x86_64
SPDK version	SPDK 21.04
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

Table 3: Hardware setup configuration – Initiator system 2

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) Microcode: 0x5003003
Memory	6 x 32GB Micron M393A1G40EB1-CRC, DDR4, 2933MHz Total 192GBs
Operating System	Fedora 33
BIOS	SE5C620.86B.02.01.0013.121520200651
Linux kernel version	5.8.15-300.fc33.x86_64
SPDK version	SPDK 21.04
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

Table 4: Test systems 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

Kernel & BIOS spectre-meltdown information

All three server systems use 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 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 as close to the CPU cache as possible.

The SPDK NVMe-oF target and initiator uses the Infiniband/RDMA verbs API to access an RDMA-capable NIC. These should work on all flavors of RDMA transports but are currently tested against RoCEv2. 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 individual 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 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 Ex 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).

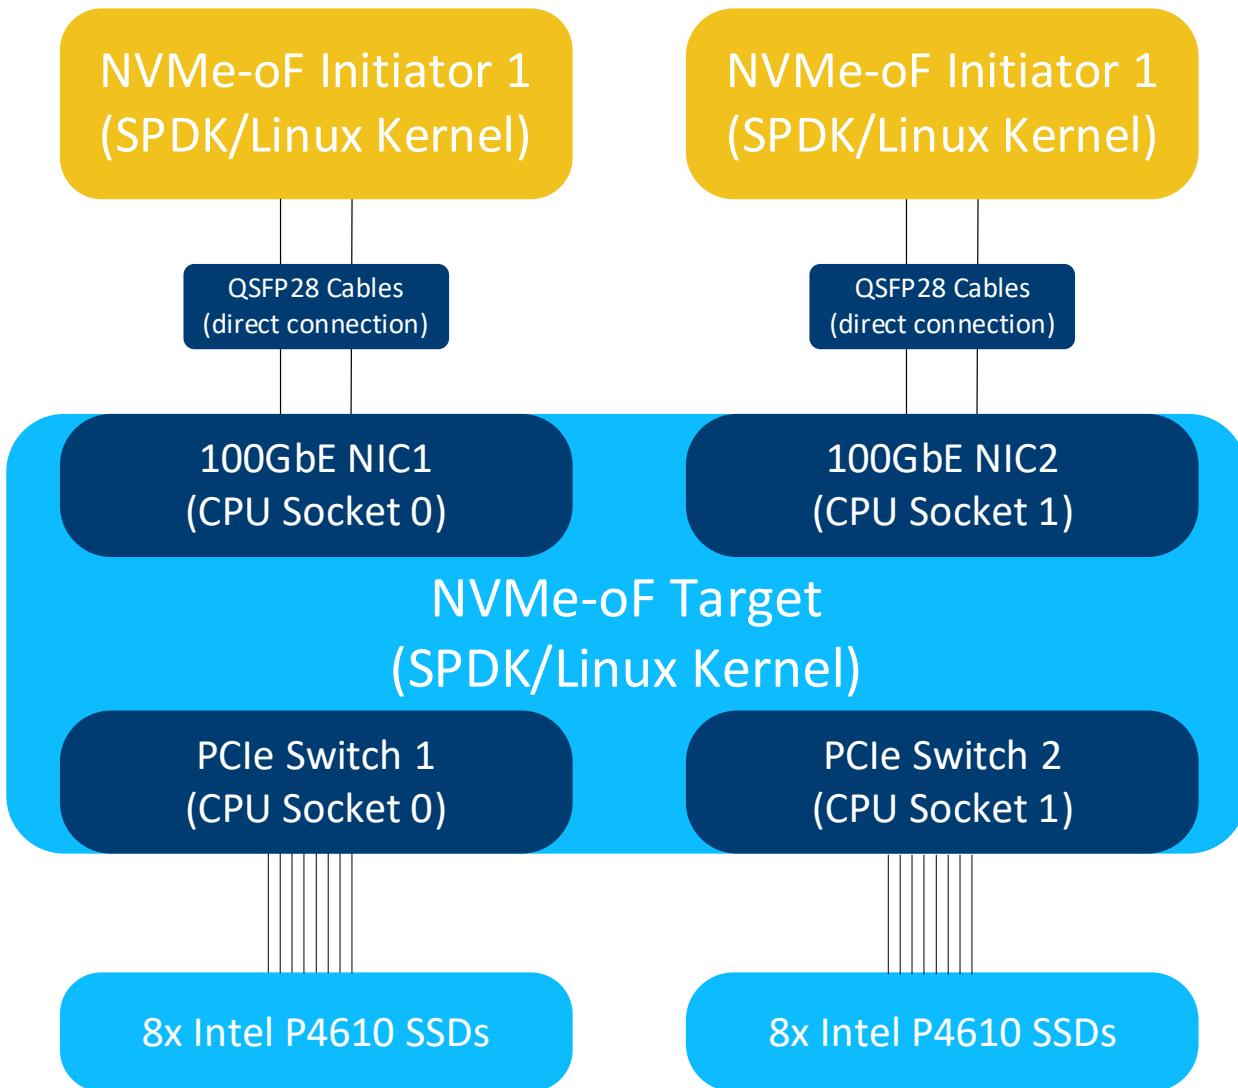


Figure 1: High-Level NVMe-oF RDMA performance testing setup

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

This test case was designed to demonstrate how the SPDK NVMe-oF target throughput in IOPS (I/O per second) scales when additional CPU cores are added to the SPDK NVMe-oF target application.

The SPDK NVMe-oF RDMA target was configured to run with 16 NVMe-oF subsystems. Each NVMe-oF subsystem ran on top of an individual bdev backed by a single Intel® SSD DC P4610 device. Each of the 2 initiators were connected to 8 individual NVMe-oF subsystems which were exposed via SPDK NVMe-oF Target over 1x 100GbE NIC. SPDK bdev FIO plugin was used to target 8 individual NVMe-oF bdevs on each of the initiators. SPDK Target Reactor Mask was configured to use up to 6 CPU cores tests while running following workloads on each initiator:

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

The table below contains more information about the test configuration. The SPDK NVMe-oF Target was configured using JSON-RPC; the table contains a 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. We preconditioned the SSDs once before running the 4KB Rand Read and 4KB Rand 70/30 Read/Write workloads to ensure that the SSDs reached their steady state where we get repeatable results. However, for the 4KB Rand Write workload we didn't precondition the NVMe devices to ensure workload saturated the network rather than being limited to the steady state performance of the SSDs which is much lower than the available network bandwidth.

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

Item	Description
Test Case	SPDK NVMe-oF Target I/O core scaling
SPDK NVMe-oF Target configuration	All the commands below were executed with <code>spdk/scripts/rpc.py</code> script. Construct NVMe bdevs: <code>construct_nvme_bdev -t PCIe -b Nvme0 -a 0000:60:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme1 -a 0000:61:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme2 -a 0000:62:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme3 -a 0000:63:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme4 -a 0000:64:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme5 -a 0000:65:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme6 -a 0000:66:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme7 -a 0000:67:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme8 -a 0000:b5:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme9 -a 0000:b6:00.0</code> <code>construct_nvme_bdev -t PCIe -b Nvme10 -a 0000:b7:00.0</code>

	<pre> construct_nvme_bdev -t PCle -b Nvme11 -a 0000:b8:00.0 construct_nvme_bdev -t PCle -b Nvme12 -a 0000:b9:00.0 construct_nvme_bdev -t PCle -b Nvme13 -a 0000:ba:00.0 construct_nvme_bdev -t PCle -b Nvme14 -a 0000:bb:00.0 construct_nvme_bdev -t PCle -b Nvme15 -a 0000:bc:00.0 Create a RDMA transport layer: nvmf_create_transport -t RDMA (creates RDMA transport layer with default values: trtype: "RDMA" 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) Create NVMe-oF subsystems and add NVMe bdevs as namespaces: for i in \$(seq 1 16); 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 rdma \ -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</p> <pre> [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_json_conf=/tmp/spdk/bdev.conf thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth={1, 8, 16, 32} time_based=1 ramp_time=60 runtime=300 [filename0] filename=Nvme0n1 </pre>

	[filename1] filename=Nvme1n1 [filename2] filename=Nvme2n1 [filename3] filename=Nvme3n1 [filename4] filename=Nvme4n1 [filename5] filename=Nvme5n1 [filename6] filename=Nvme6n1 [filename7] filename=Nvme7n1
--	---

4KB Random Read Results

Table 6: SPDK NVMe-oF RDMA Target Core Scaling results, Random Read IOPS, QD=64

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	4554.58	1166.0	877.5
2 cores	10503.74	2689.0	380.3
3 cores	16222.23	4152.9	245.4
4 cores	20806.02	5326.3	191.5
5 cores	21483.92	5499.9	185.8
6 cores	22935.18	5871.4	174.8

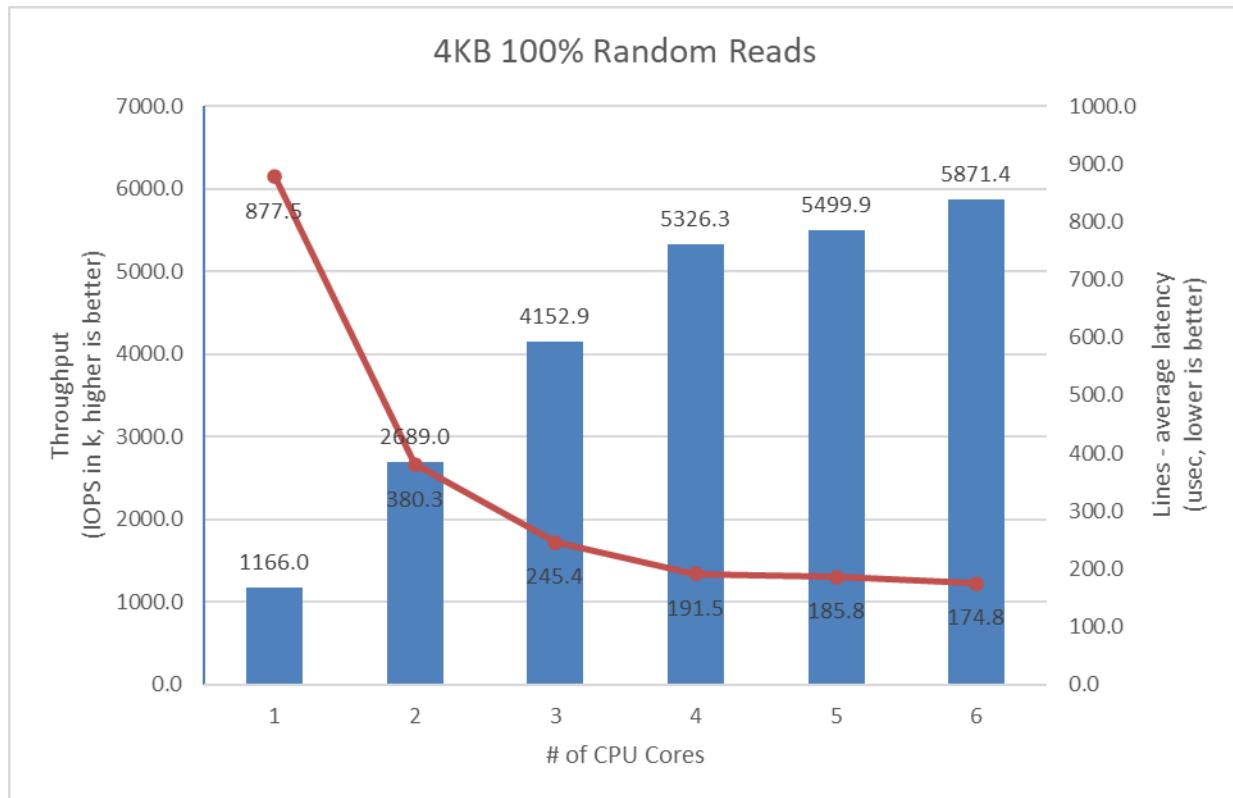


Figure 2: SPDK NVMe-oF RDMA Target I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Read workload at QD = 64

4KB Random Write Results

Table 7: SPDK NVMe-oF RDMA Target Core Scaling results, Random Write IOPS, QD=32

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	6201.09	1587.5	319.9
2 cores	13179.06	3373.8	150.3
3 cores	18799.37	4812.6	104.0
4 cores	22925.61	5869.0	85.2
5 cores	22357.23	5723.4	88.6
6 cores	22965.21	5879.1	86.3

Note that the SSDs were not preconditioned for the 4K random write workload because that would limit the workload performance to the SSDs steady state performance.

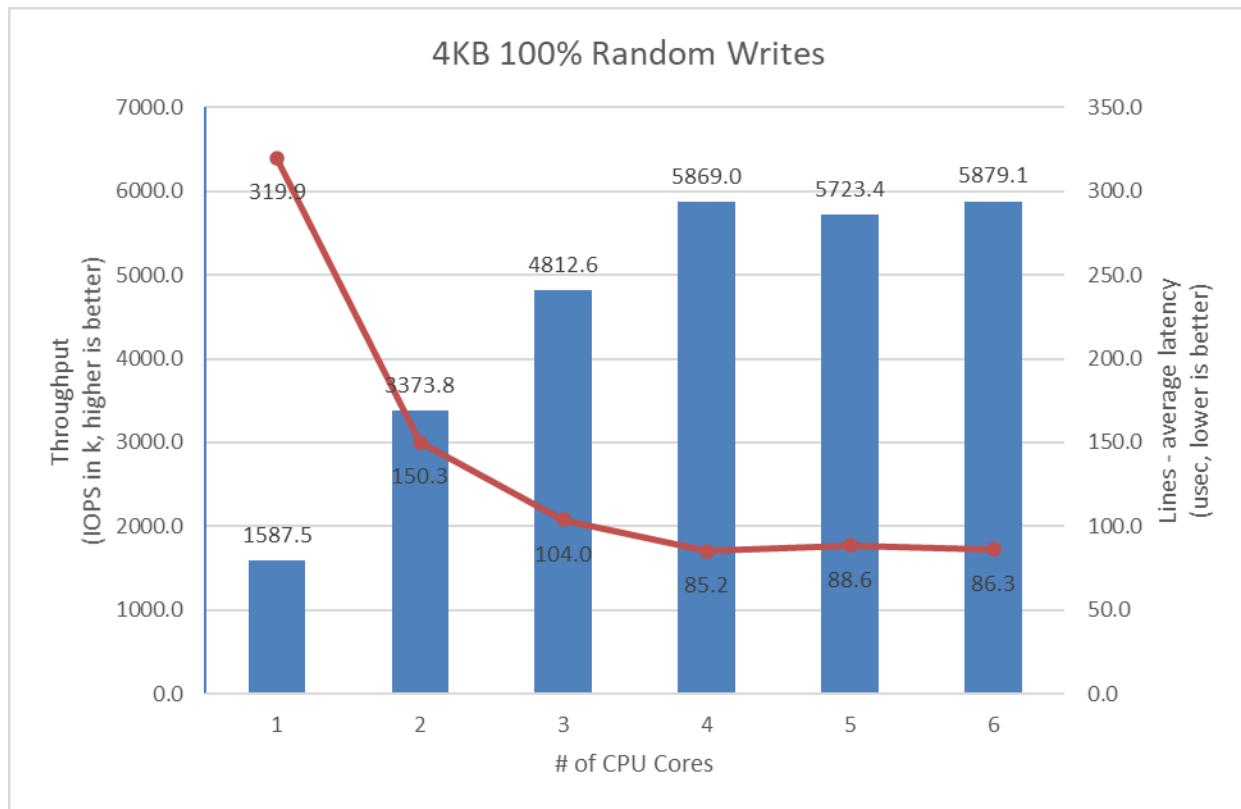


Figure 3: SPDK NVMe-oF RDMA Target I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Write Workload at QD=32

4KB Random Read-Write Results

Table 8: SPDK NVMe-oF RDMA Target Core Scaling results, Random Read/Write 70%/30% IOPS, QD=128

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	4544.51	1163.4	1759.4
2 cores	10117.43	2590.1	796.4
3 cores	15755.24	4033.3	506.7
4 cores	21035.88	5385.2	380.1
5 cores	22440.20	5744.7	356.3
6 cores	23840.05	6103.0	335.7

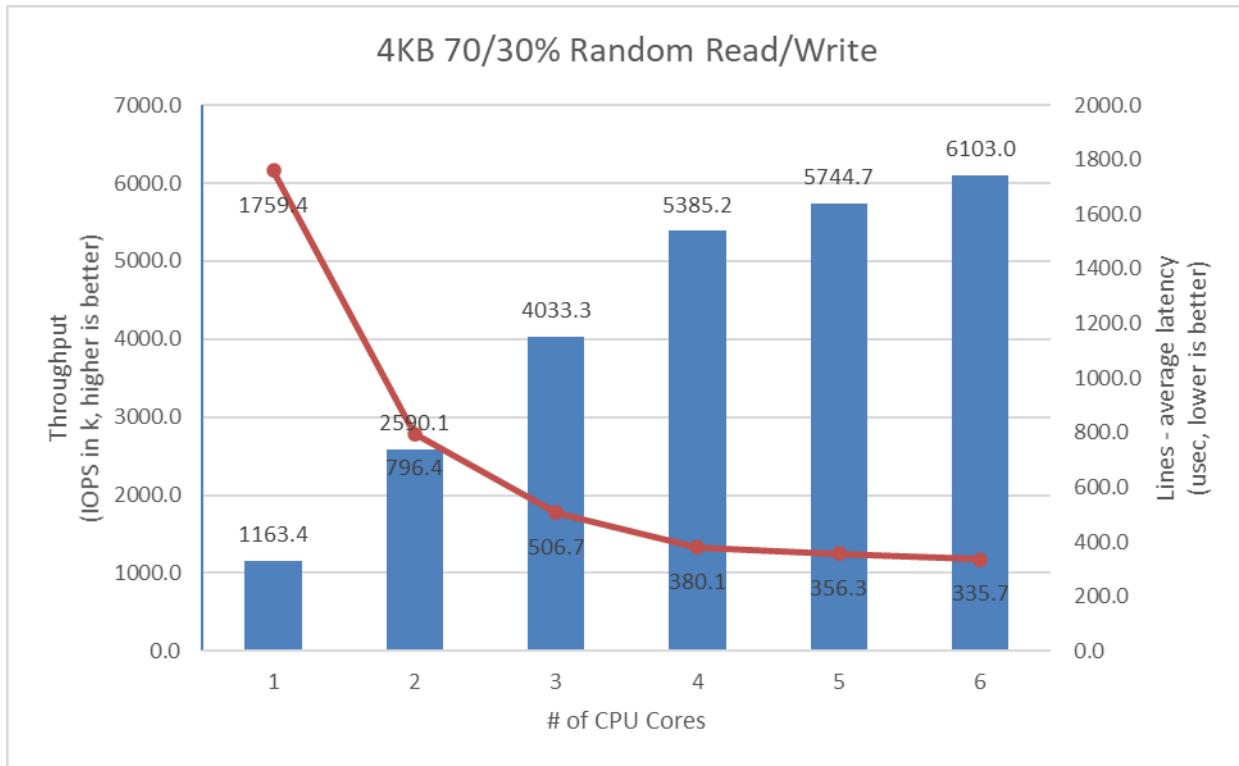


Figure 4: SPDK NVMe-oF RDMA Target I/O core scaling: IOPS vs. Latency while running 4KB Random 70/30 Read/Write workload at QD=128

LTO Performance Impact

Below results present performance impact of using SPDK “—with-lto” build option, which enables link-time optimization.

Table 9: LTO impact for SPDK NVMe-oF RDMA Target Core Scaling performance – Bandwidth

Workload	# of Cores	Bandwidth (MBps)		
		LTO Disabled	LTO Enabled	LTO Impact
4KB Randread QD=64	1 core	4554.58	4593.31	0.85%
	2 cores	10503.74	10794.04	2.76%
	3 cores	16222.23	16361.94	0.86%
	4 cores	20806.02	21293.40	2.34%
	5 cores	21483.92	22282.97	3.72%
	6 cores	22935.18	23562.31	2.73%
4KB Randwrite QD=32	1 core	6201.09	6379.18	2.87%
	2 cores	13179.06	13291.73	0.85%
	3 cores	18799.37	19532.05	3.90%
	4 cores	22925.61	23100.20	0.76%
	5 cores	22357.23	22586.04	1.02%
	6 cores	22965.21	22718.48	-1.07%
4KB RandRW QD=128	1 core	4544.51	4679.57	2.97%
	2 cores	10117.43	10504.65	3.83%
	3 cores	15755.24	16058.06	1.92%
	4 cores	21035.88	21754.64	3.42%
	5 cores	22440.20	22991.60	2.46%
	6 cores	23840.05	24508.93	2.81%

Table 10: LTO impact for SPDK NVMe-oF RDMA Target Core Scaling performance - Throughput

		Throughput (IOPS k) – Higher is better		
Workload	# of Cores	LTO Disabled	LTO Enabled	LTO Impact
4KB Randread QD=64	1 core	1165.97	1175.88	0.85%
	2 cores	2688.95	2763.27	2.76%
	3 cores	4152.89	4188.65	0.86%
	4 cores	5326.34	5451.11	2.34%
	5 cores	5499.88	5704.44	3.72%
	6 cores	5871.40	6031.95	2.73%
4KB Randwrite QD=32	1 core	1587.48	1633.07	2.87%
	2 cores	3373.84	3402.68	0.85%
	3 cores	4812.64	5000.20	3.90%
	4 cores	5868.95	5913.65	0.76%
	5 cores	5723.45	5782.03	1.02%
	6 cores	5879.09	5815.93	-1.07%
4KB RandRW QD=128	1 core	1163.39	1197.96	2.97%
	2 cores	2590.06	2689.18	3.83%
	3 cores	4033.33	4110.86	1.92%
	4 cores	5385.18	5569.18	3.42%
	5 cores	5744.68	5885.84	2.46%
	6 cores	6103.05	6274.28	2.81%

Table 11: LTO impact for SPDK NVMe-oF RDMA Target Core Scaling performance – Average Latency

Workload	# of Cores	Avg. Latency (usec) – Lower is better		
		LTO Disabled	LTO Enabled	LTO Impact
4KB Randread QD=64	1 core	877.47	870.05	-0.84%
	2 cores	380.33	369.79	-2.77%
	3 cores	245.45	244.90	-0.22%
	4 cores	191.51	187.24	-2.23%
	5 cores	185.78	179.47	-3.40%
	6 cores	174.82	170.00	-2.76%
4KB Randwrite QD=32	1 core	319.92	311.09	-2.76%
	2 cores	150.25	154.39	2.75%
	3 cores	103.97	99.94	-3.88%
	4 cores	85.17	84.88	-0.34%
	5 cores	88.62	87.64	-1.10%
	6 cores	86.33	87.39	1.23%
4KB RandRW QD=128	1 core	1759.41	1708.41	-2.90%
	2 cores	796.35	766.83	-3.71%
	3 cores	506.66	503.51	-0.62%
	4 cores	380.09	365.27	-3.90%
	5 cores	356.32	348.57	-2.18%
	6 cores	335.66	327.17	-2.53%

Large Sequential I/O Performance

128KB block size I/O tests were performed with sequential I/O workloads at queue depth 8. The rest of the FIO configuration is similar to the 4KB test case in the previous part of this document. We used iodepth=8 (iodepth=1 for write workload) because higher queue depth resulted in negligible bandwidth gain and a significant increase in the latency.

Table 12: SPDK NVMe-oF RDMA Target Core Scaling results, 128KB Sequential Read IOPS, QD=8

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	24971.77	199.8	640.5
2 cores	24971.77	199.8	640.5
3 cores	24971.77	199.8	640.5
4 cores	24971.77	199.8	640.5

Table 13: SPDK NVMe-oF RDMA Target Core Scaling results, 128KB Sequential Write IOPS, QD=1

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	22978.94	183.8	86.7
2 cores	23202.38	185.6	85.9
3 cores	23322.82	186.6	85.4
4 cores	23241.49	185.9	85.7

Table 14: SPDK NVMe-oF RDMA Target Core Scaling results, 128KB Sequential 70% Read 30% Write IOPS, QD=4

# of Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	24350.03	194.8	328.2
2 cores	24530.34	196.2	325.8
3 cores	24664.32	197.3	324.0
4 cores	24558.58	196.5	325.4

Conclusions

1. For all tested workloads throughput scales up and latency decreases almost linearly with the addition of I/O cores up to 4 cores. Adding more CPU cores results in small performance gain as network is almost saturated.
2. For large sequential I/Os, a single CPU core Target saturated the network bandwidth. The SPDK NVMe-oF Target running on 1 core does close to 22-25 GBps of bandwidth in all tested workloads. Therefore, adding more CPU cores did not result in increased performance for these workloads because the network was saturated.
3. Enabling Link-Time Optimization build option for SPDK slightly increased performance by up to 3.9%

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

This test case was designed to demonstrate how the SPDK NVMe-oF initiator throughput in IOPS (I/O per second) scales when additional CPU cores are added to the SPDK NVMe-oF initiator.

The SPDK NVMe-oF RDMA Target was configured using 6 cores; all the other configurations are similar to test case 1. The SPDK bdev FIO plugin was used to target 8 NVMe-oF bdevs on each of the 2 initiators. The following workloads were executed on both of the initiators using fio in client-server mode.

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

Depending on the number of initiators and initiator cores, we varied the number of NVMe-oF subsystems exported by the target as shown below, instead of exposing all 16 NVMe-oF subsystems.

1 core: 1 initiator running on a single core connecting to 4 subsystems.

2 cores: 2 initiators, each running on a single core. Each initiator connected to 4 subsystems.

3 cores: 2 initiators, the first one running on 1 core and other one running on 2 cores. Initiator 1 connected to 4 subsystems and initiator 2 connected to 8 subsystems.

4 cores: 2 initiators, both running on 2 cores each. Both initiators connected to 8 subsystems.

This was done to avoid a situation where single initiator would connect to all 16 target subsystems, which resulted in unnatural latency spike in the 1 CPU core test case.

Table 15: SPDK NVMe-oF RDMA Initiator Core Scaling test configuration

Item	Description
Test Case	SPDK NVMe-oF RDMA Initiator I/O core scaling
SPDK NVMe-oF Target configuration	Same as in Test Case #1, using 6 CPU cores.
SPDK NVMe-oF Initiator 1 - FIO plugin configuration	BDEV.conf See Appendix A FIO.conf [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_json_conf=/tmp/spdk/bdev.conf thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw

	<pre>rwmixread={100, 70, 0} bs=4k iodepth={32, 64, 128, 256} time_based=1 ramp_time=60 runtime=300 {filename_section} FIO.conf filename section for 1 & 2 CPU test run [filename0] filename=Nvme0n1 filename=Nvme1n1 filename=Nvme2n1 filename=Nvme3n1 FIO.conf filename section for 3 & 4 CPU test run [filename0] filename=Nvme0n1 filename=Nvme1n1 filename=Nvme2n1 filename=Nvme3n1 [filename1] filename=Nvme4n1 filename=Nvme5n1 filename=Nvme6n1 filename=Nvme7n1</pre>
SPDK NVMe-oF Initiator 2 - FIO plugin configuration	<p>BDEV.conf See Appendix A</p> <p>FIO.conf Similar as Initiator 1. The only differences are were in the filename section which are shown below.</p> <p>FIO.conf filename section for 1 CPU test run N/A (only Initiator 1 is used)</p> <p>FIO.conf filename section for 2 CPU test run [filename0] filename=Nvme0n1 filename=Nvme1n1 filename=Nvme2n1 filename=Nvme3n1</p> <p>FIO.conf filename section for 3 CPU test run [filename0] filename=Nvme0n1 filename=Nvme1n1 filename=Nvme2n1 filename=Nvme3n1</p> <p>FIO.conf filename section for 4 CPU test run [filename0] filename=Nvme0n1 filename=Nvme1n1</p>

	<p>filename=Nvme2n1 filename=Nvme3n1 [filename1] filename=Nvme4n1 filename=Nvme5n1 filename=Nvme6n1 filename=Nvme7n1</p>
--	--

4KB Random Read Results

Table 16: SPDK NVMe-oF RDMA Initiator Core Scaling results, 4KB Random Read IOPS, QD=64, SPDK Target 6 CPU Cores

# of Initiator CPU Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	5929.86	1518.04	546.42
2 cores	16628.18	4256.81	213.03
3 cores	18381.87	4705.76	206.34
4 cores	22334.59	5717.65	177.96
5 cores	22998.59	5887.64	172.71
6 cores	21287.81	5449.67	186.60
7 cores	22552.75	5773.50	176.87
8 cores	20975.51	5369.73	189.90

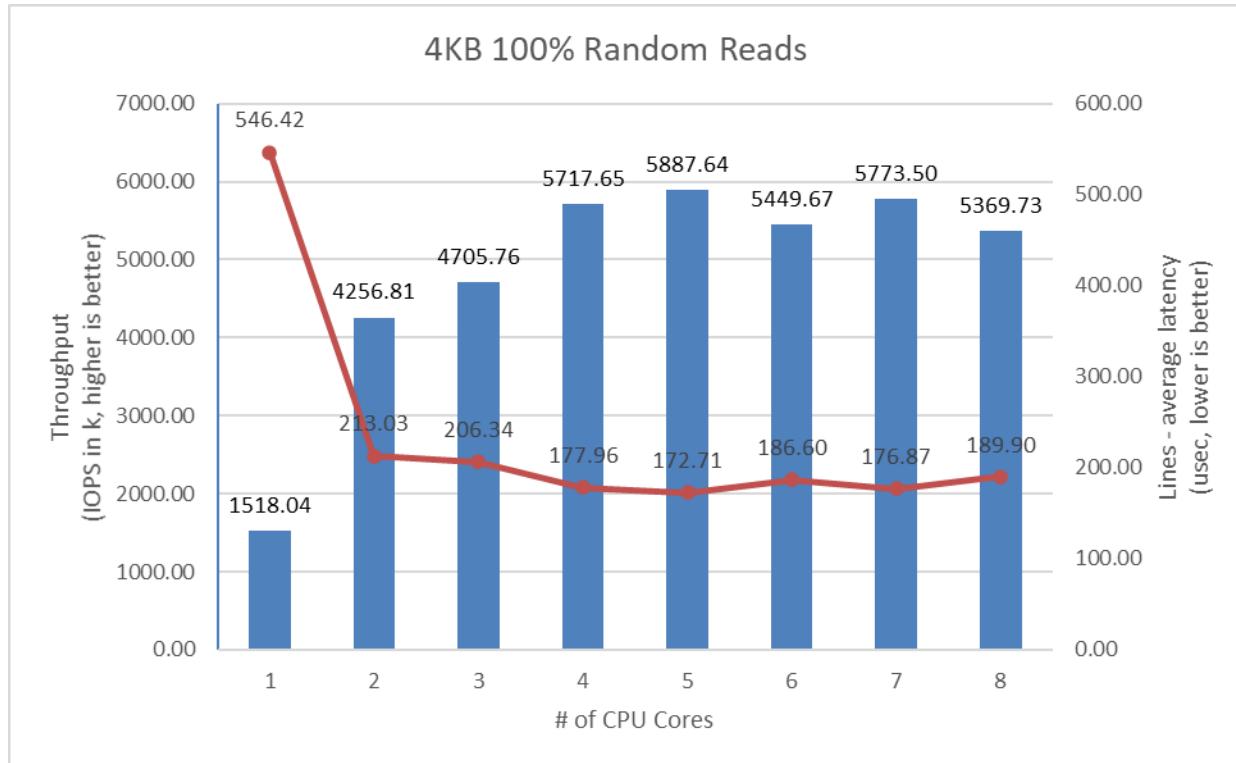


Figure 5: SPDK NVMe-oF RDMA Initiator I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Read QD=64 workload

4KB Random Write Results

Note: The SSDs were not pre-conditioned before running the 100% Random Write test cases. This allowed the throughput to scale to the 2x 100GbE network bandwidth when testing with 3 and 4 CPU cores rather than limiting the workload performance to the storage bottleneck (which is approx. 3.2M IOPS).

Table 17: SPDK NVMe-oF RDMA Initiator Core Scaling results, 4KB Random Write IOPS, QD=32, SPDK Target 6 CPU Cores

# of Initiator CPU Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	6358.65	1627.81	251.11
2 cores	14082.50	3605.12	114.30
3 cores	18461.85	4726.23	99.02
4 cores	22833.44	5845.36	83.75
5 cores	22943.31	5873.49	83.55
6 cores	22951.68	5875.63	84.67
7 cores	22667.59	5802.90	86.77
8 cores	22609.28	5787.97	87.41

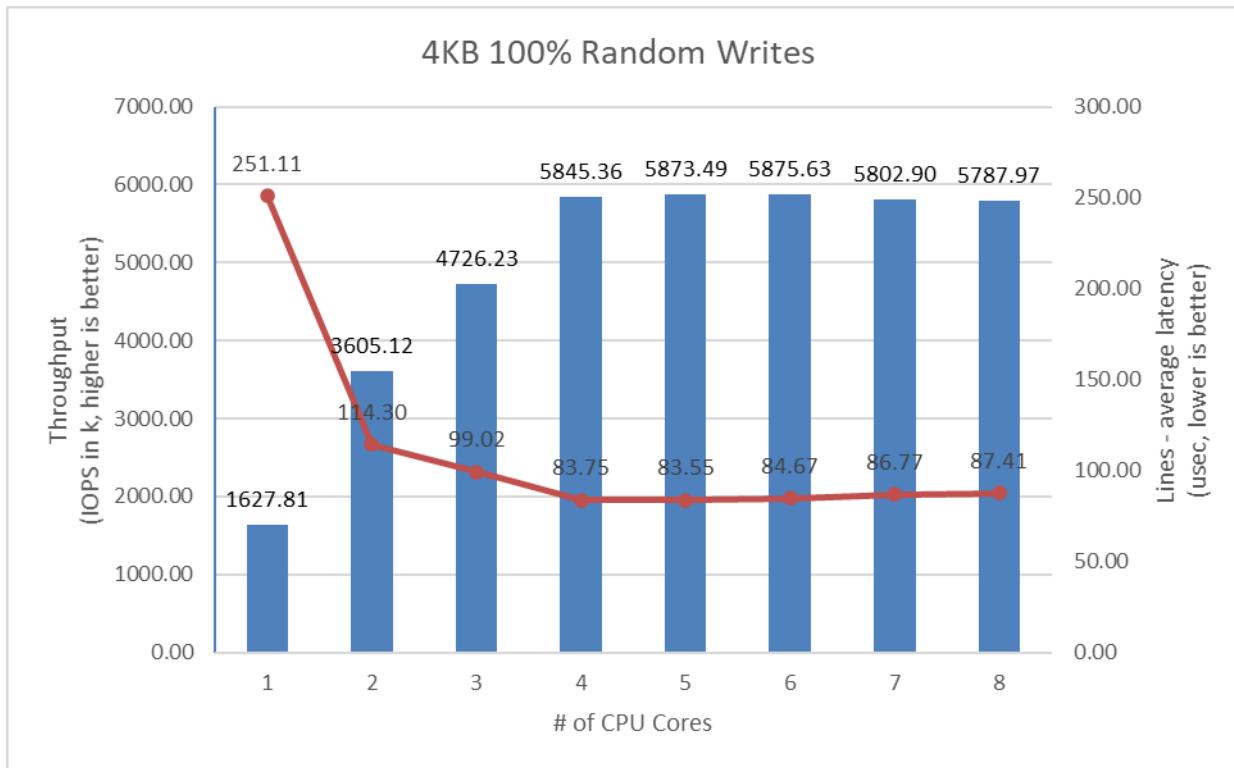


Figure 6: SPDK NVMe-oF RDMA Initiator I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Write Workload at QD=32

4KB Random 70/30 Read/Write Results

Table 18: SPDK NVMe-oF RDMA Initiator Core Scaling results, 4KB Random 70%/30% Read/Write IOPS, QD=128, SPDK Target 6 CPU Cores

# of Initiator CPU Cores	Bandwidth (MBps)	Throughput (IOPS k)	Avg. Latency (usec)
1 core	4824.61	1235.09	1343.11
2 cores	10864.25	2781.24	617.51
3 cores	16686.21	4271.66	465.72
4 cores	23002.91	5888.74	338.30
5 cores	22068.68	5649.57	358.30
6 cores	22450.40	5747.29	352.90
7 cores	21017.54	5380.48	381.61
8 cores	21171.40	5419.87	377.84

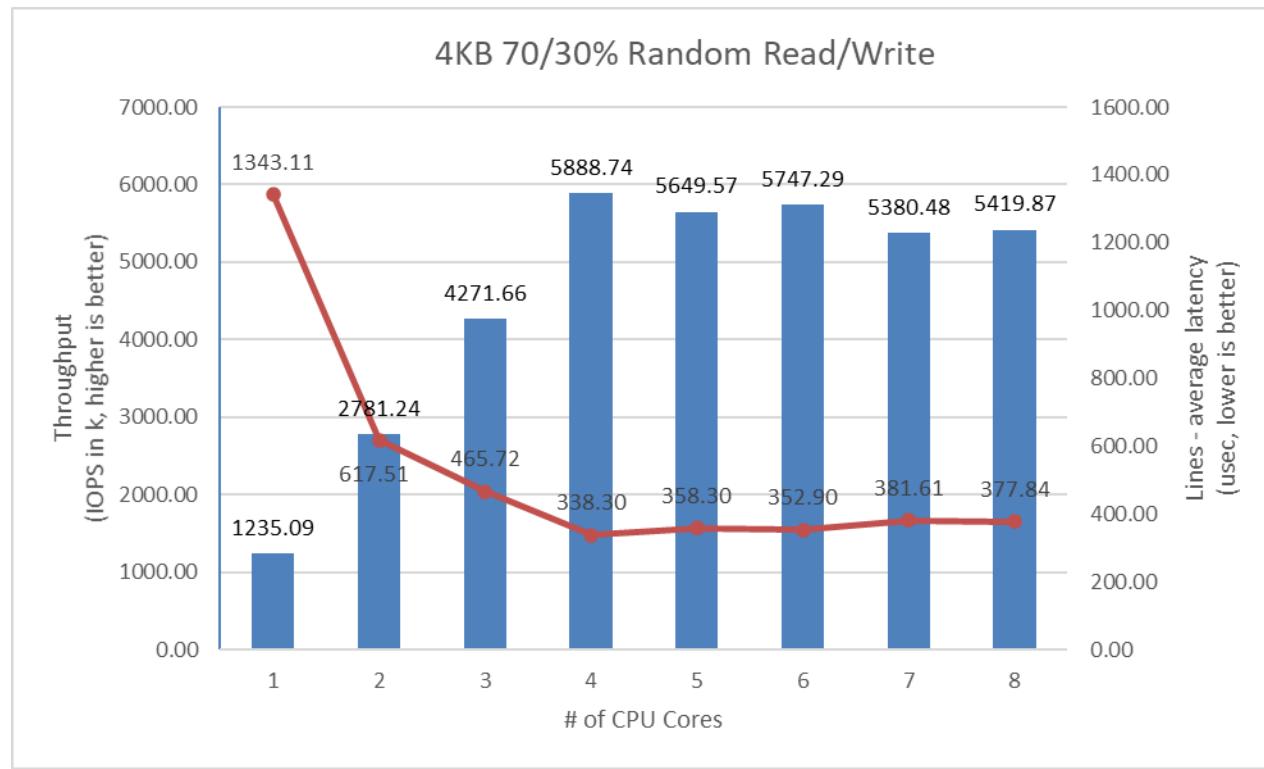


Figure 7: SPDK NVMe-oF RDMA Initiator I/O core scaling: IOPS vs. Latency while running 4KB Random 70% Read 30% Write Workload at QD=128

Conclusions

1. Random Read workload scaling was not linear when increasing the number of initiator CPU cores.
2. Random Write and Random Read-Write IOPS scaling was linear or close to linear as we increased the number of initiator cores up to 4 CPU cores.
3. For all workloads beyond 4 CPU cores there was no significant change in IOPS because the 200 Gbps network was saturated.
4. Peak performance of 5.8 million IOPS was reached at 5 Initiator CPU cores for Random Read workload, 5.8 million at 4 CPU cores for Random Write workload and for mixed Random Read/Write workload.

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

This test case was designed to understand latency characteristics of SPDK NVMe-oF RDMA Target and Initiator vs. the Linux Kernel NVMe-oF RDMA 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 containing a *Null Block Device*. The null block device (bdev) was chosen as the backend block device to eliminate the media latency during these tests.

Table 19: Linux Kernel vs. SPDK NVMe-oF RDMA Latency test configuration

Item	Description
Test Case	Linux Kernel vs. SPDK NVMe-oF RDMA Latency
Test configuration	
SPDK NVMe-oF Target configuration	<p>The following commands are executed with spdk/scripts/rpc.py script to configure the SPDK NVMe-oF target.</p> <pre>nvmf_create_transport -t RDMA (creates RDMA transport layer with default values: trtype: "RDMA" 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 rdma -f ipv4 -s 4420 -a 20.0.0.1</pre>
Kernel NVMe-oF Target configuration	<p>The following target configuration file loaded using nvmet-cl tool.</p> <pre>{ "ports": [{ "addr": { "adrfam": "ipv4", "traddr": "20.0.0.1", "trsvcid": "4420", "trtype": "rdma" }, "portid": 1, "referrals": [], "subsystems": [...] }] }</pre>

	<pre> "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>
FIO configuration	
SPDK NVMe-oF Initiator FIO plugin configuration	<p>BDEV.conf See Appendix B.</p> <p>FIO.conf</p> <pre> [global] ioengine=/tmp/spdk/examples/bdev/fio_plugin/fio_plugin spdk_json_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 [filename0] filename=NvmeOn1 </pre>
Kernel initiator configuration	<p>Device config The following configuration was performed using nvme-cli tool. modprobe nvme-fabrics nvme connect -n nqn.2018-09.io.spdk:cnode1 -t rdma -a 20.0.0.1 -s 4420</p> <p>FIO.conf</p>

```
[global]
ioengine=libaio
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

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

SPDK vs Kernel NVMe-oF RDMA Target Results

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

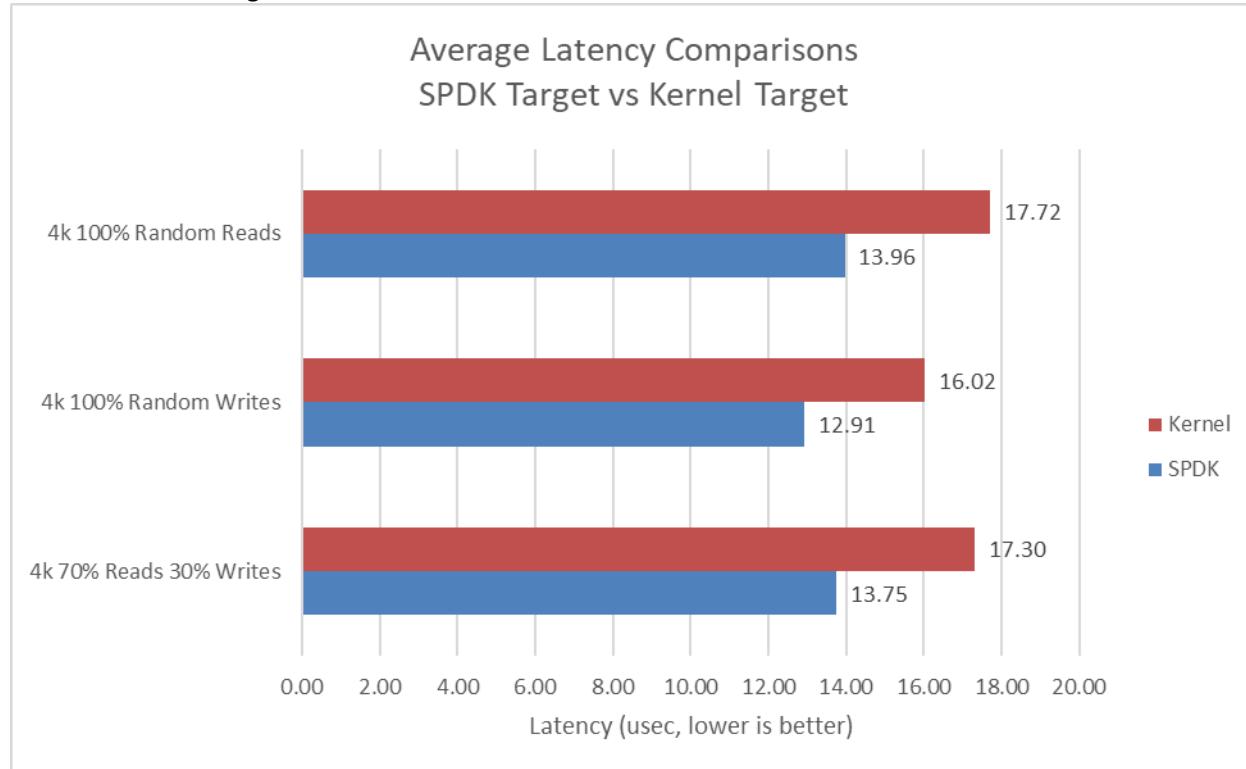


Figure 8: SPDK vs. Kernel NVMe-oF RDMA Target average I/O latency for various workloads run using the Kernel Initiator

Table 20: SPDK NVMe-oF RDMA Target Latency and IOPS at QD=1, Null Block Device

Access Pattern	Average Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KB 100% Random Read	13.96	70398	21.6	27.7	36.8	136.2
4KB 100% Random Write	12.91	75897	20.3	26.8	34.2	134.8
4KB 70/30% Random Read/Write	13.75	71302	21.4	27.4	37.3	136.7

Table 21: Linux Kernel NVMe-oF RDMA Target Latency and IOPS at QD=1, Null Block Device

Access Pattern	Average Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KB 100% Random Read	17.72	55691	18.5	31.4	32.4	34.6
4KB 100% Random Write	16.02	61323	17.1	29.4	30.7	34.1
4KB 70/30% Random Read/Write	17.30	56922	17.8	29.4	31.8	34.3

SPDK vs Kernel NVMe-oF RDMA Initiator Results

The following data was collected using the Linux Kernel and SPDK NVMe-oF RDMA initiator against an SPDK NVMe-oF RDMA target.

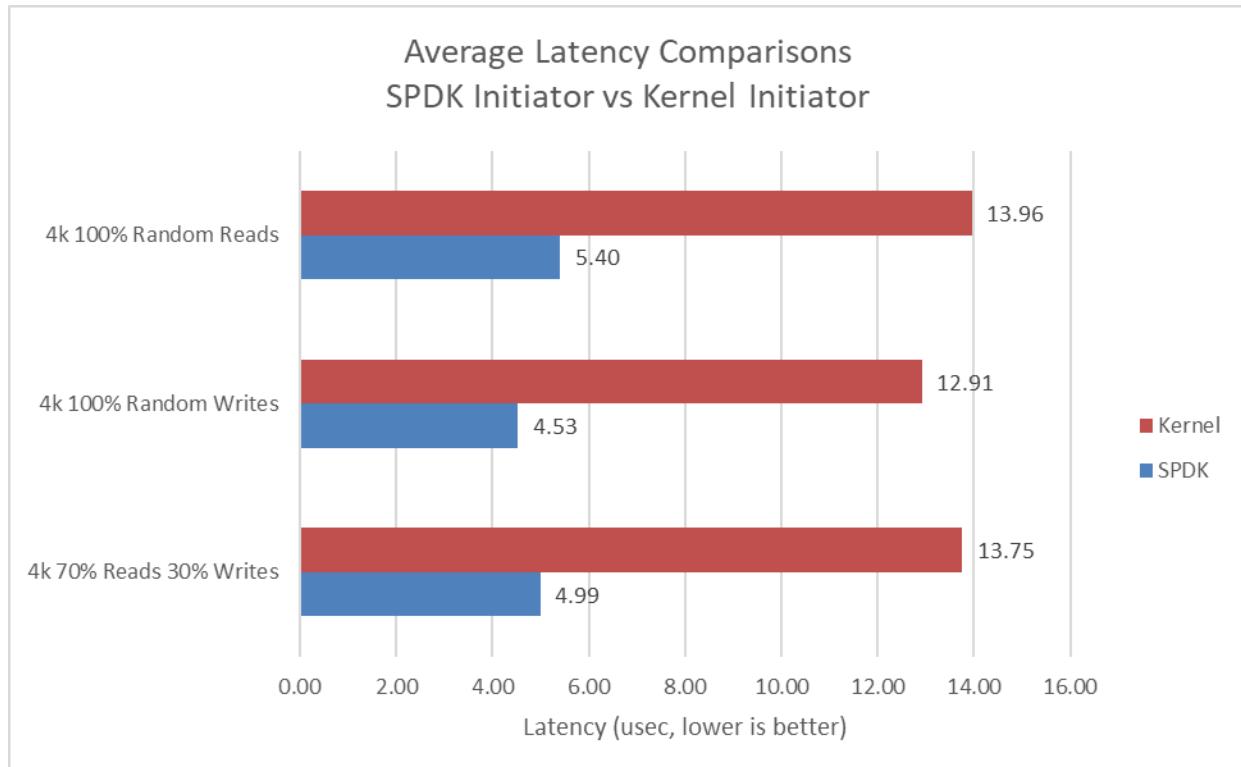


Figure 9: SPDK vs. Kernel NVMe-oF RDMA Initiator average I/O latency for various workloads against SPDK Target

Table 22: SPDK NVMe-oF RDMA Initiator Latency and IOPS at QD=1, Null Block Device

Access Pattern	Average Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KB 100% Random Read	5.40	179705	5.7	13.7	19.1	49.4
4KB 100% Random Write	4.53	212339	4.8	17.8	20.6	41.6
4KB 70/30% Random Read/Write	4.99	192774	5.3	18.8	21.0	48.6

Table 23: Linux Kernel NVMe-oF RDMA Initiator Latency and IOPS at QD=1, Null Block Device

Access Pattern	Average Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KB 100% Random Read	13.96	70398	21.6	27.7	36.8	136.2
4KB 100% Random Write	12.91	75897	20.3	26.8	34.2	134.8
4KB 70/30% Random Read/Write	13.75	71302	21.4	27.4	37.3	136.7

SPDK vs Kernel NVMe-oF RDMA Kernel + Initiator Results

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

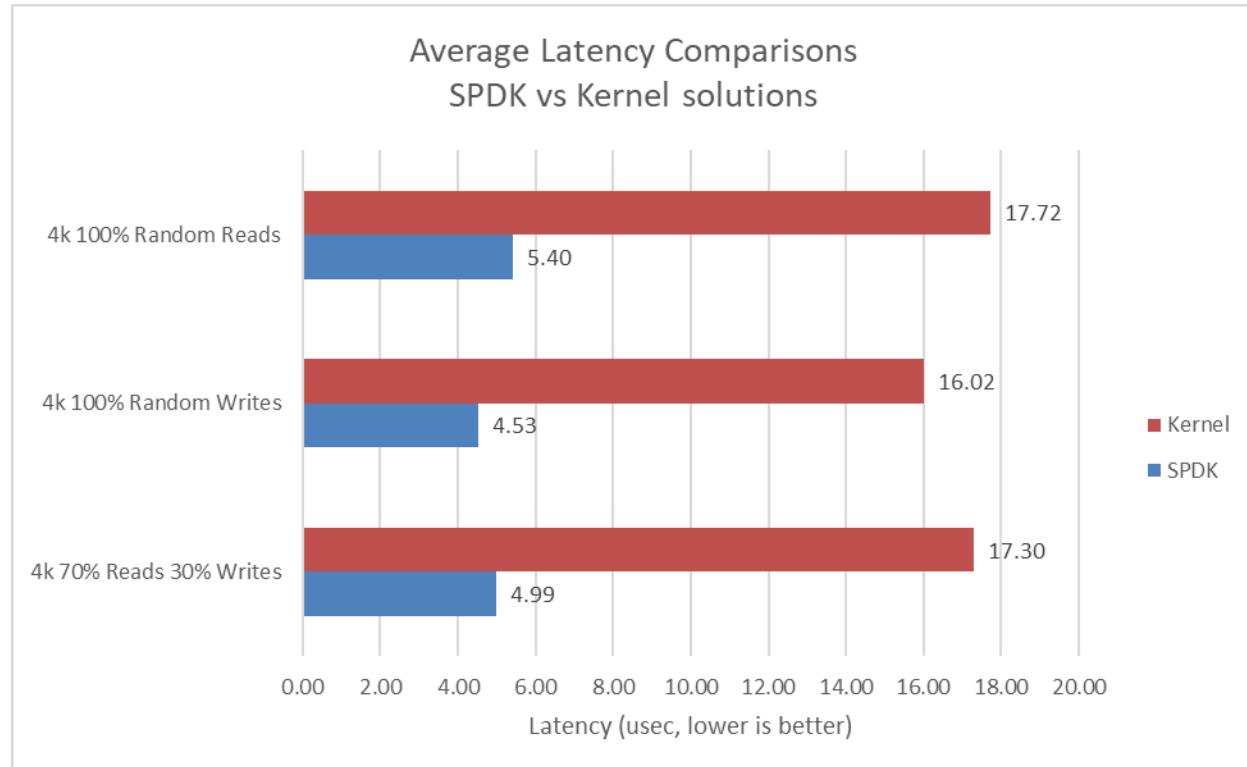


Figure 10: SPDK vs. Kernel NVMe-oF RDMA solutions average I/O Latency for various workloads against SPDK Target

Table 24: SPDK NVMe-oF RDMA Latency and IOPS at QD=1, Null Block Device

Access Pattern	Average Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KB 100% Random Read	5.40	179705	5.7	13.7	19.1	49.4
4KB 100% Random Write	4.53	212339	4.8	17.8	20.6	41.6
4KB 70/30% Random Read/Write	4.99	192774	5.3	18.8	21.0	48.6

Table 25: Linux Kernel NVMe-oF RDMA Latency and IOPS at QD=1, Null Block Device

Access Pattern	Average Latency (usec)	IOPS	p99 (usec)	p99.9 (usec)	p99.99 (usec)	p99.999 (usec)
4KB 100% Random Read	17.72	55691	18.5	31.4	32.4	34.6
4KB 100% Random Write	16.02	61323	17.1	29.4	30.7	34.1
4KB 70/30% Random Read/Write	17.30	56922	17.8	29.4	31.8	34.3

Conclusions

1. For the RDMA transport, the SPDK NVMe-oF Target reduces the NVMe-oF average round trip I/O latency (reads/writes) by up to 3.76 usec vs. the Linux Kernel NVMe-oF target used in Fedora 33 5.8.15 setup. This is entirely software overhead, therefore, using the SPDK NVMe-oF target reduces the NVMe-oF software overhead by approximately 21% vs. the Linux Kernel NVMe-oF target.
2. The SPDK NVMe-oF initiator reduces the NVMe-oF software overhead by up to 2.8x times vs. the Linux Kernel NVMe-oF Initiator for the RDMA transport.
3. The SPDK NVMe-oF RDMA target and initiator reduced the average latency by up to 71% vs. the Linux Kernel NVMe-oF RDMA target and initiator.

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

This test case was designed to demonstrate the throughput and latency of the SPDK NVMe-oF RDMA Target vs. Linux Kernel NVMe-oF RDMA Target under increasing number of connections per subsystem. The number of connections (or I/O queue pairs) per NVMe-oF subsystem were varied, we measured the aggregated IOPS and number of CPU cores used by each target. The number of CPU cores metric was calculated from %CPU utilization measured using sar (systat package in Linux). The SPDK NVMe-oF RDMA Target was configured to run on 8 CPU cores, export 16 NVMe-oF subsystems (1 per Intel P4610) and 2 initiators were used both running the I/O workloads below to 8 separate subsystems using Kernel NVMe-oF RDMA initiator.

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

Table 26: NVMe-oF RDMA Performance with increasing number of connections test configuration

Item	Description
Test Case	NVMe-oF RDMA Target performance with increasing # of connections
SPDK NVMe-oF Target configuration	Same as in Test Case #1, using 8 CPU cores.
Kernel NVMe-oF Target configuration	Target configuration file loaded using nvmet-clt tool. For detailed configuration file contents please see Appendix C .
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 rdma -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode2 -t rdma -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode3 -t rdma -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode4 -t rdma -a 20.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode5 -t rdma -a 20.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode6 -t rdma -a 20.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode7 -t rdma -a 20.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode8 -t rdma -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 rdma -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode10 -t rdma -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode11 -t rdma -a 10.0.0.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode12 -t rdma -a 10.0.0.1 -s 4420

	<pre>nvme connect -n nqn.2018-09.io.spdk:cnode13 -t rdma -a 10.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode14 -t rdma -a 10.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode15 -t rdma -a 10.0.1.1 -s 4420 nvme connect -n nqn.2018-09.io.spdk:cnode16 -t rdma -a 10.0.1.1 -s 4420</pre>
FIO configuration (used on both initiators)	<pre>FIO.conf [global] ioengine=libaio thread=1 group_reporting=1 direct=1 norandommap=1 rw=randrw rwmixread={100, 70, 0} bs=4k iodepth={8, 16, 32, 64, 128} time_based=1 ramp_time=60 runtime=300 numjobs={1, 4, 16, 32} [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</pre>

The SPDK NVMe-oF Target was configured to use 8 CPU cores, whereas we did not limit the number of CPU cores for the Linux Kernel NVMe-oF target. The graph below shows the relative efficiency in terms of IOPS/core which was calculated by dividing the total aggregate IOPS by the 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 11: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF RDMA Target IOPS/Core for 4KB 100% Random Reads QD=64, using the Kernel Initiator

Table 27: Linux Kernel NVMe-oF RDMA Target: 4KB 100% Random Reads, QD=64

Connections per subsystem	Bandwidth (MBps)	IOPS (k)	Avg. Latency (usec)	# CPU Cores
1	15317.95	3921.4	261.1	9.8
4	14319.60	3665.8	278.9	22.1
16	13747.36	3519.3	290.3	31.6

Table 28: SPDK NVMe-oF RDMA Target: 4KB 100% Random Reads, QD=64

Connections per subsystem	Bandwidth (MBps)	IOPS (k)	Avg. Latency (usec)	# CPU Cores
1	17866.43	4573.8	223.7	8.0
4	14545.90	3723.7	274.6	8.0
16	13564.50	3472.5	294.2	8.0

4KB Random Write results

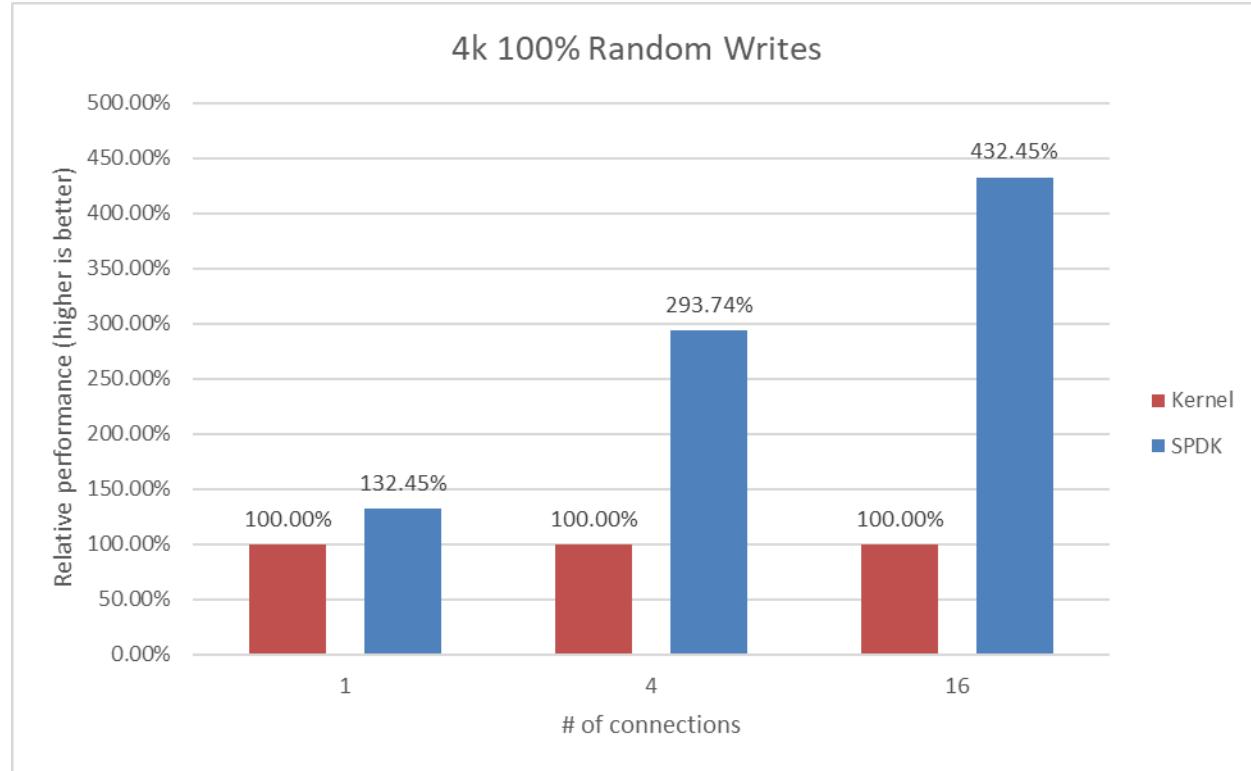


Figure 12: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF RDMA Target IOPS/Core for 4KB 100% Random Writes QD=64 workload, using Kernel Initiators

Note: The SSDs were not pre-conditioned before running 100% Random Write I/O test.

Table 29: Linux Kernel NVMe-oF RDMA Target: 4KB 100% Random Writes, QD=64

Connections per subsystem	Bandwidth (MBps)	IOPS (k)	Avg. Latency (usec)	# CPU Cores
1	20643.36	5284.7	193.6	10.1
4	12435.72	3183.5	338.6	10.4
16	12878.14	3296.8	310.1	24.6

Table 30: SPDK NVMe-oF RDMA Target: 4KB 100% Random Writes, QD=64

Connections per subsystem	Bandwidth (MBps)	IOPS (k)	Avg. Latency (usec)	# CPU Cores
1	21218.52	5431.9	188.3	8.0
4	22779.43	5831.5	175.3	8.0
16	17940.28	4592.7	222.2	8.0

4KB Random Read-Write Results

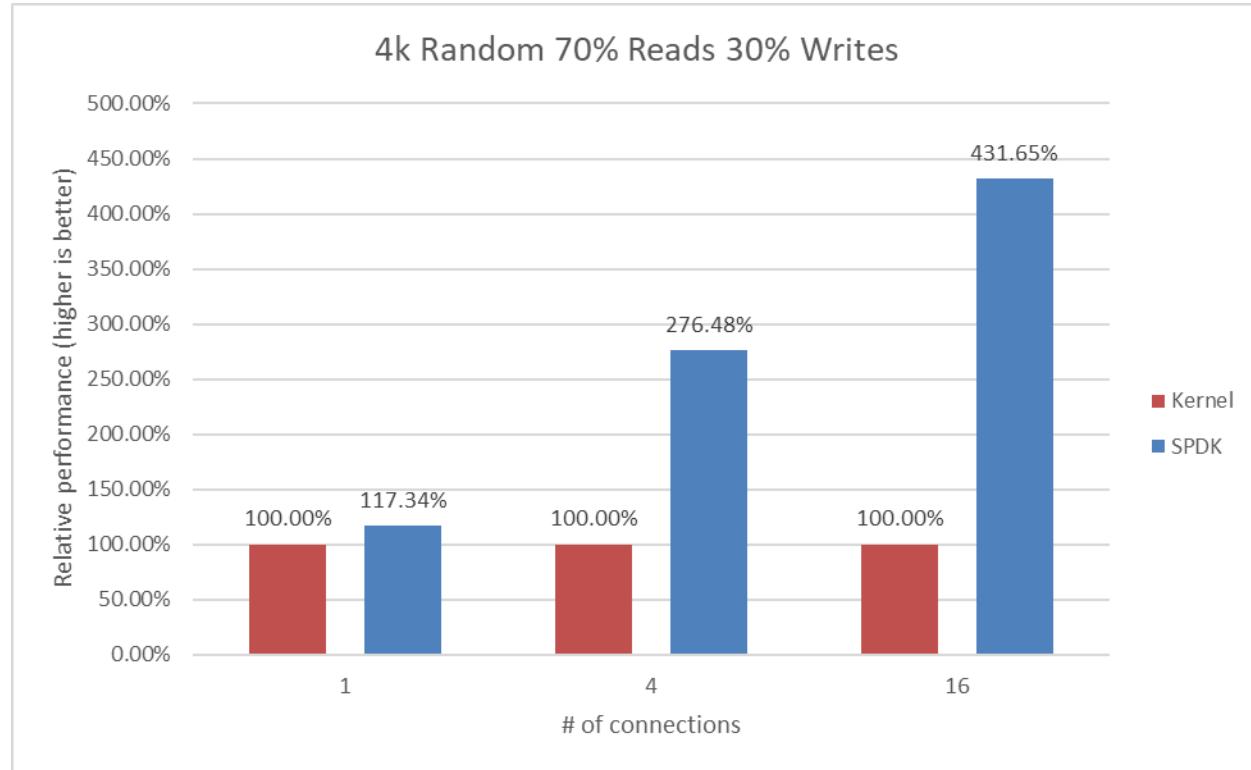


Figure 13: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF RDMA Target IOPS/Core for 4KB Random 70% Reads 30% Writes QD=64 Workload, using Kernel Initiators

Table 31: Linux Kernel NVMe-oF RDMA Target: 4KB 70% Random Read 30% Random Write, QD=64

Connections per subsystem	Bandwidth (MBps)	IOPS (k)	Avg. Latency (usec)	# CPU Cores
1	16157.37	4136.3	247.6	9.9
4	18682.35	4782.7	213.7	20.9
16	16698.33	4274.8	238.7	35.9

Table 32: SPDK NVMe-oF RDMA Target: 4KB 70% Random Read 30% Random Write, QD=64

Connections per subsystem	Bandwidth (MBps)	IOPS (k)	Avg. Latency (usec)	# CPU Cores
1	15383.96	3938.3	261.7	8.0
4	18353.27	4698.4	217.6	8.0
16	15967.51	4087.7	249.7	8.0

Conclusions

1. When the SPDK NVMe-oF Target was configured with 8 CPU cores the performance peaked at 1 connection for 4KB Random Read workload and 4 connections for Random Write and Random Read/Write workloads. Increasing the number of connections per subsystem beyond these values resulted in higher average latency and lower IOPS.
2. The performance for the Linux Kernel NVMe-oF Target peaked at 1 connection per subsystem for Random Read and Random Write workloads. Increasing the number of connections only increases the latency and CPU utilization.
3. The SPDK NVMe-oF target shows up to 4.3x more IOPS/Core relative to the Linux Kernel NVMe-oF target as the number of connections per subsystem increased.

Summary

This report showcased performance results with SPDK NVMe-oF RDMA 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 the Linux Kernel NVMe-oF RDMA (Target/Initiator) against the accelerated polled-mode driven SPDK NVMe-oF RDMA (Target/Initiator) implementation. Like in the last report, throughput scales up and latency decreases almost linearly with the scaling of SPDK NVMe-oF target and initiator cores.

It was also observed that the SPDK NVMe-oF Target average latency is up to 3.76 usec lower than Kernel when testing against null bdev based backend. The advantage of SPDK is even greater when comparing NVMe-oF Initiators: the SPDK NVMe-oF RDMA average latency is by up to 2.8x times lower than Kernel initiator.

The SPDK NVMe-oF Target performed up to 4.3 times better w.r.t IOPS/core than Linux Kernel NVMe-oF target while running 4KB 100% random read/write 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 different hardware and software configurations.

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 RDMA Target Core Scaling test configuration	9
Table 6: SPDK NVMe-oF RDMA Target Core Scaling results, Random Read IOPS, QD=64	12
Table 7: SPDK NVMe-oF RDMA Target Core Scaling results, Random Write IOPS, QD=32	13
Table 8: SPDK NVMe-oF RDMA Target Core Scaling results, Random Read/Write 70%/30% IOPS, QD=128	14
Table 9: LTO impact for SPDK NVMe-oF RDMA Target Core Scaling performance – Bandwidth	15
Table 10: LTO impact for SPDK NVMe-oF RDMA Target Core Scaling performance - Throughput	16
Table 11: LTO impact for SPDK NVMe-oF RDMA Target Core Scaling performance – Average Latency	17
Table 12: SPDK NVMe-oF RDMA Target Core Scaling results, 128KB Sequential Read IOPS, QD=4	18
Table 13: SPDK NVMe-oF RDMA Target Core Scaling results, 128KB Sequential Write IOPS, QD=1	18
Table 14: SPDK NVMe-oF RDMA Target Core Scaling results, 128KB Sequential 70% Read 30% Write IOPS, QD=4	18
Table 15: SPDK NVMe-oF RDMA Initiator Core Scaling test configuration	20
Table 16: SPDK NVMe-oF RDMA Initiator Core Scaling results, 4KB Random Read IOPS, QD=64, SPDK Target 6 CPU Cores	23
Table 17: SPDK NVMe-oF RDMA Initiator Core Scaling results, 4KB Random Write IOPS, QD=32, SPDK Target 6 CPU Cores	24
Table 18: SPDK NVMe-oF RDMA Initiator Core Scaling results, 4KB Random 70%/30% Read/Write IOPS, QD=128, SPDK Target 6 CPU Cores	25
Table 19: Linux Kernel vs. SPDK NVMe-oF RDMA Latency test configuration	27
Table 20: SPDK NVMe-oF RDMA Target Latency and IOPS at QD=1, Null Block Device	30
Table 21: Linux Kernel NVMe-oF RDMA Target Latency and IOPS at QD=1, Null Block Device	30
Table 22: SPDK NVMe-oF RDMA Initiator Latency and IOPS at QD=1, Null Block Device	31
Table 23: Linux Kernel NVMe-oF RDMA Initiator Latency and IOPS at QD=1, Null Block Device	31
Table 24: SPDK NVMe-oF RDMA Latency and IOPS at QD=1, Null Block Device	32
Table 25: Linux Kernel NVMe-oF RDMA Latency and IOPS at QD=1, Null Block Device	32
Table 26: NVMe-oF RDMA Performance with increasing number of connections test configuration	34
Table 27: Linux Kernel NVMe-oF RDMA Target: 4KB 100% Random Reads, QD=64	36

Table 28: SPDK NVMe-oF RDMA Target: 4KB 100% Random Reads, QD=64	36
Table 29: Linux Kernel NVMe-oF RDMA Target: 4KB 100% Random Writes, QD=64	37
Table 30: SPDK NVMe-oF RDMA Target: 4KB 100% Random Writes, QD=64.....	37
Table 31: Linux Kernel NVMe-oF RDMA Target: 4KB 70% Random Read 30% Random Write, QD=64	38
Table 32: SPDK NVMe-oF RDMA Target: 4KB 70% Random Read 30% Random Write, QD=64.....	38

List of figures

Figure 1: High-Level NVMe-oF RDMA performance testing setup	8
Figure 2: SPDK NVMe-oF RDMA Target I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Read workload at QD = 64	12
Figure 3: SPDK NVMe-oF RDMA Target I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Write Workload at QD=32	13
Figure 4: SPDK NVMe-oF RDMA Target I/O core scaling: IOPS vs. Latency while running 4KB Random 70/30 Read/Write workload at QD=128.....	14
Figure 5: SPDK NVMe-oF RDMA Initiator I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Read QD=64 workload.....	23
Figure 6: SPDK NVMe-oF RDMA Initiator I/O core scaling: IOPS vs. Latency while running 4KB 100% Random Write Workload at QD=32	24
Figure 7: SPDK NVMe-oF RDMA Initiator I/O core scaling: IOPS vs. Latency while running 4KB Random 70% Read 30% Write Workload at QD=128	25
Figure 8: SPDK vs. Kernel NVMe-oF RDMA Target average I/O latency for various workloads run using the Kernel Initiator	30
Figure 9: SPDK vs. Kernel NVMe-oF RDMA Initiator average I/O latency for various workloads against SPDK Target	31
Figure 10: SPDK vs. Kernel NVMe-oF RDMA solutions average I/O Latency for various workloads against SPDK Target	32
Figure 11: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF RDMA Target IOPS/Core for 4KB 100% Random Reads QD=64, using the Kernel Initiator.....	36
Figure 12: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF RDMA Target IOPS/Core for 4KB 100% Random Writes QD=64 workload, using Kernel Initiators	37
Figure 13: Relative Efficiency Comparison of Linux Kernel vs. SPDK NVMe-oF RDMA Target IOPS/Core for 4KB Random 70% Reads 30% Writes QD=64 Workload, using Kernel Initiators	38

Appendix A – Test Case 1 & 2 SPDK NVMe-oF Initiator bdev configuration

Initiator system 1

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

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

Initiator system 2

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

```

```

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

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

Appendix B – 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",  
                        "trsvid": "4420",  
                        "subnqn": "nqn.2018-09.io.spdk:cnode0",  
                        "adrfam": "IPv4"  
                    }  
                }  
            ]  
        }  
    ]  
}
```

Appendix C - Kernel NVMe-oF RDMA Target configuration

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

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

```
        "nqn.2018-09.io.spdk:cnode2"
    ],
},
{
    "addr": {
        "adrifam": "ipv4",
        "traddr": "20.0.0.1",
        "trsvcid": "4422",
        "trtype": "rdma"
    },
    "portid": 3,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode3"
    ]
},
{
    "addr": {
        "adrifam": "ipv4",
        "traddr": "20.0.0.1",
        "trsvcid": "4423",
        "trtype": "rdma"
    },
    "portid": 4,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode4"
    ]
},
{
    "addr": {
        "adrifam": "ipv4",
        "traddr": "20.0.1.1",
        "trsvcid": "4424",
        "trtype": "rdma"
    },
    "portid": 5,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode5"
    ]
},
{
    "addr": {
        "adrifam": "ipv4",
        "traddr": "20.0.1.1",
        "trsvcid": "4425",
        "trtype": "rdma"
    },
    "portid": 6,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode6"
    ]
},
{
    "addr": {
```

```

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

```

```
"portid": 11,
"referrals": [],
"subsystems": [
    "nqn.2018-09.io.spdk:cnode11"
]
},
{
    "addr": {
        "adrfam": "ipv4",
        "traddr": "10.0.0.1",
        "trsvid": "4431",
        "trtype": "rdma"
    },
    "portid": 12,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode12"
    ]
},
{
    "addr": {
        "adrfam": "ipv4",
        "traddr": "10.0.1.1",
        "trsvid": "4432",
        "trtype": "rdma"
    },
    "portid": 13,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode13"
    ]
},
{
    "addr": {
        "adrfam": "ipv4",
        "traddr": "10.0.1.1",
        "trsvid": "4433",
        "trtype": "rdma"
    },
    "portid": 14,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode14"
    ]
},
{
    "addr": {
        "adrfam": "ipv4",
        "traddr": "10.0.1.1",
        "trsvid": "4434",
        "trtype": "rdma"
    },
    "portid": 15,
    "referrals": [],
    "subsystems": [
        "nqn.2018-09.io.spdk:cnode15"
    ]
}
```

```

        },
        {
            "addr": {
                "adrfam": "ipv4",
                "traddr": "10.0.1.1",
                "trsvcid": "4435",
                "trtype": "rdma"
            },
            "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-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"
}
]
```

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at 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.

§