intel.

SPDK Vhost Performance Report Release 22.09

Testing Date: October 2022

Performed by:

Karol Latecki (karol.latecki@intel.com)

Jaroslaw Chachulski (jaroslawx.chachulski@intel.com)

Acknowledgments:

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

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

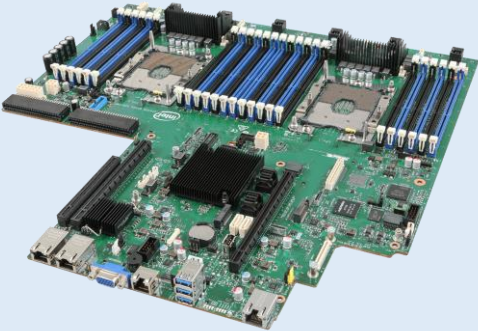# *Contents*

Audience and Purpose

This report is intended for people who are interested in looking at SPDK Vhost scsi and blk stack performance and comparison to its Linux kernel equivalents. It provides performance and efficiency comparisons between SPDK Vhost-scsi and Linux Kernel Vhost-scsi software stacks under various test cases.

The purpose of this report is not to imply a single correct approach, but rather to provide a baseline of well-tested configurations and procedures that produce repeatable and reproducible results. This report can also be viewed as information regarding best known method when performance testing SPDK Vhost-scsi and Vhost-blk stacks.

Test setup

Hardware configuration

Table 1: Hardware setup configuration

Item	Description
Server Platform	Intel WolfPass **R2224WFTZS**
	Server board **S2600WFT**
Motherboard	S2600WFT
CPU	2 CPU sockets, Intel(R) Xeon(R) Gold 6230N CPU @ 2.30GHz
	Number of cores 20 per socket, number of threads 40 per socket
	Both sockets populated
	Microcode: 0x4003003
Memory	12 x 32GB Micron DDR4 36ASF4G72PZ-2G9E2
	Total 384 GBs
	Memory channel population:

P1	P2
CPU1_DIMM_A1	CPU2_DIMM_A1
CPU1_DIMM_B1	CPU2_DIMM_B1
CPU1_DIMM_C1	CPU2_DIMM_C1
CPU1_DIMM_D1	CPU2_DIMM_D1
CPU1_DIMM_E1	CPU2_DIMM_E1
CPU1_DIMM_F1	CPU2_DIMM_F1

Item	Description
Operating System	Fedora 33
BIOS	SE5C620.86B.02.01.0015.032120220358

Linux kernel version	5.11.20-200.fc33.x86_64
SPDK version	SPDK 22.09
Qemu version	QEMU emulator version 5.1.0 (qemu-5.1.0-9.fc33)
Storage	**OS:** 1x 120GB Intel SSDSC2BB120G4
	Storage: 24x Intel® P4610™ 1.6TBs (FW: VDV10170) (6 on CPU NUMA Node 0, 18 on CPU NUMA Node 1)

BIOS Settings

Table 2: Test platform BIOS settings

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

Virtual Machine Settings

Table 3: Guest VM configuration

Item	Description
CPU	2vCPU, pass through from physical host server. Explicit core usage enforced using "taskset –a –c" command. QEMU arguments for starting the VM: -cpu host -smp 1
Memory	4 GB RAM. Memory is pre-allocated for each VM using Hugepages on host system and used from appropriate NUMA node, to match the CPU which was passed to the VM. QEMU arguments: -m 4096 -object memory-backend-file,id=mem,size=4096M,mem-path=/dev/hugepages,share=on,prealloc=yes,host-nodes=0,policy=bind
Operating System	Fedora 33
Linux kernel version	5.9.16-200.fc33.x86_64
Additional boot options in /etc/default/grub	Multi queue enabled: scsi_mod.use_blk_mq=1

Introduction to the SPDK Vhost target

SPDK Vhost is a userspace target designed to extend the performance efficiencies of SPDK into QEMU/KVM virtualization environments. The SPDK Vhost-scsi target presents a broad range of SPDK-managed block devices into virtual machines. SPDK community has leveraged existing SPDK SCSI layer, DPDK Vhost library, QEMU Vhost-scsi and Vhost-blk functionality to create the high performance SPDK userspace Vhost target.

SPDK Vhost target architecture

QEMU sets up the Vhost target via UNIX domain socket. QEMU pre-allocates huge pages for the guest VM to enable DMA by the Vhost target. The guest VM submits I/O directly to the Vhost target via virtqueues in shared memory as shown in Figure 1. The Vhost target transfers data to/from the guest VM via shared memory. The Vhost target then completes I/O to the guest VM via virtqueues in shared memory. There is a completion interrupt sent using eventfd which requires a system call and a guest VM exit. It should be noted that there is no QEMU intervention during the I/O submission process.
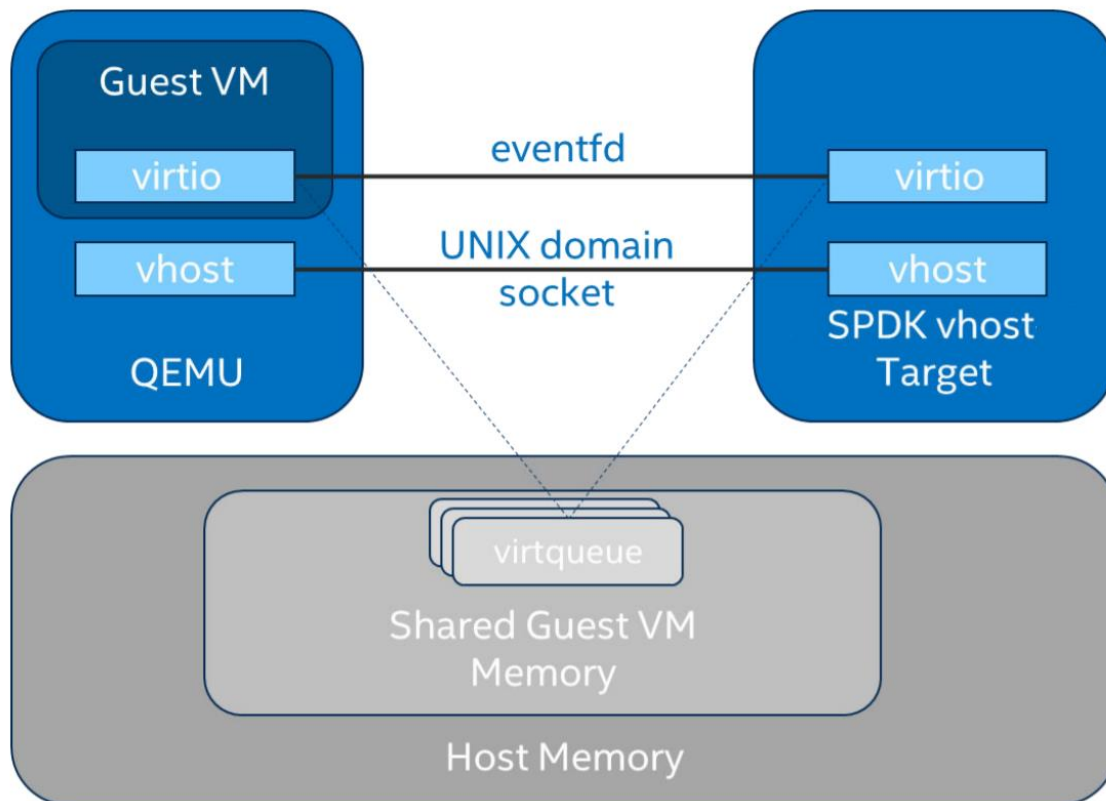

Figure 1: SPDK Vhost-scsi architecture

This report shows the performance comparisons between the traditional interrupt-driven Linux Kernel Vhost-scsi and the accelerated polled-mode SPDK Vhost-scsi under 3 different test cases using local

NVMe storage. Additionally, the SPDK Vhost-blk stack is included in the report for further comparison with the SCSI stack.

Test Case 1: SPDK Vhost Core Scaling

This test case was performed to understand aggregate VM performance with SPDK Vhost I/O core scaling. We ran up to 36 virtual machines, each running following FIO workloads:

- 4KB 100% Random Read

- 4KB 100% Random Write

- 4KB Random 70% Read / 30 % Write

We increased the number of CPU cores used by SPDK Vhost target to process I/O from 1 up to 12 and measured the throughput (in IOPS) and latency. The number of VMs between test runs was not constant and was increased by 6 for each Vhost CPU added, up to a maximum of 36 VMs. VM number was not increased beyond 36 because of the platform capabilities in terms of available CPU cores.

FIO was run in client-server mode. FIO client was run on the host machine and distributed jobs to FIO servers run on each VM. This allowed us to start the FIO jobs across all VMs at the same time.

Results in the table and charts represent aggregate performance (IOPS and average latency) seen across all the VMs. The results are average of 3 runs.

Table 4: SPDK Vhost Core Scaling test configuration

Item	Description
Test case	Test SPDK Vhost target I/O core scaling performance
Test configuration	**FIO Version**: fio-3.19 **VM Configuration**: • Common settings are described in the **Virtual Machine Settings** chapter. • Number of VMs: variable (6 VMs per 1 Vhost CPU core, up to 36 VMs max). • Each VM has a single Vhost device as a target for the FIO workload. This is achieved by sharing SPDK NVMe bdevs by using either a Split NVMe vbdev or Logical Volume bdev configuration. **SPDK Vhost target configuration:** • Test were run with both the Vhost-scsi and Vhost-blk stacks. • The Vhost-scsi stack was run with Split NVMe bdevs and Logical Volume bdevs. • Vhost-blk stack was run with Logical Volume bdevs. • Tests were performed with 1, 2, 4, 6, 8, 10 and 12 Vhost cores for each stack-bdev combination. **Kernel Vhost target configuration:** - N/A
FIO configuration	[global] ioengine=libaio

direct=1
thread=1
norandommap=1
time_based=1
gtod_reduce=0
ramp_time=60s
runtime=240s
numjobs=1
bs=4k
rw=randrw
rwmixread=100 (100% reads), 70 (70% reads, 30% writes), 0 (100% writes)
iodepth={1, 32, 64}

4KB Random Read Results

Table 5: SPDK Vhost core scaling results, 4KB 100% Random Reads IOPS, QD=64

# of CPU cores	# of VMs	Stack / Backend	IOPS (millions)
1	6	SCSI / Split NVMe Bdev	1.79
		SCSI / Lvol Bdev	1.58
		BLK / Lvol Bdev	1.69
2	12	SCSI / Split NVMe Bdev	2.95
		SCSI / Lvol Bdev	2.63
		BLK / Lvol Bdev	2.90
4	24	SCSI / Split NVMe Bdev	4.92
		SCSI / Lvol Bdev	4.25
		BLK / Lvol Bdev	4.58
6	36	SCSI / Split NVMe Bdev	6.45
		SCSI / Lvol Bdev	5.27
		BLK / Lvol Bdev	5.79
8	36	SCSI / Split NVMe Bdev	7.05
		SCSI / Lvol Bdev	6.30
		BLK / Lvol Bdev	7.06
10	36	SCSI / Split NVMe Bdev	6.93
		SCSI / Lvol Bdev	6.43
		BLK / Lvol Bdev	7.22
12	36	SCSI / Split NVMe Bdev	6.99
		SCSI / Lvol Bdev	7.19
		BLK / Lvol Bdev	7.84

SPDK Vhost 4KB 100% Random Reads

Bars - IOPS in millions (higher is better)

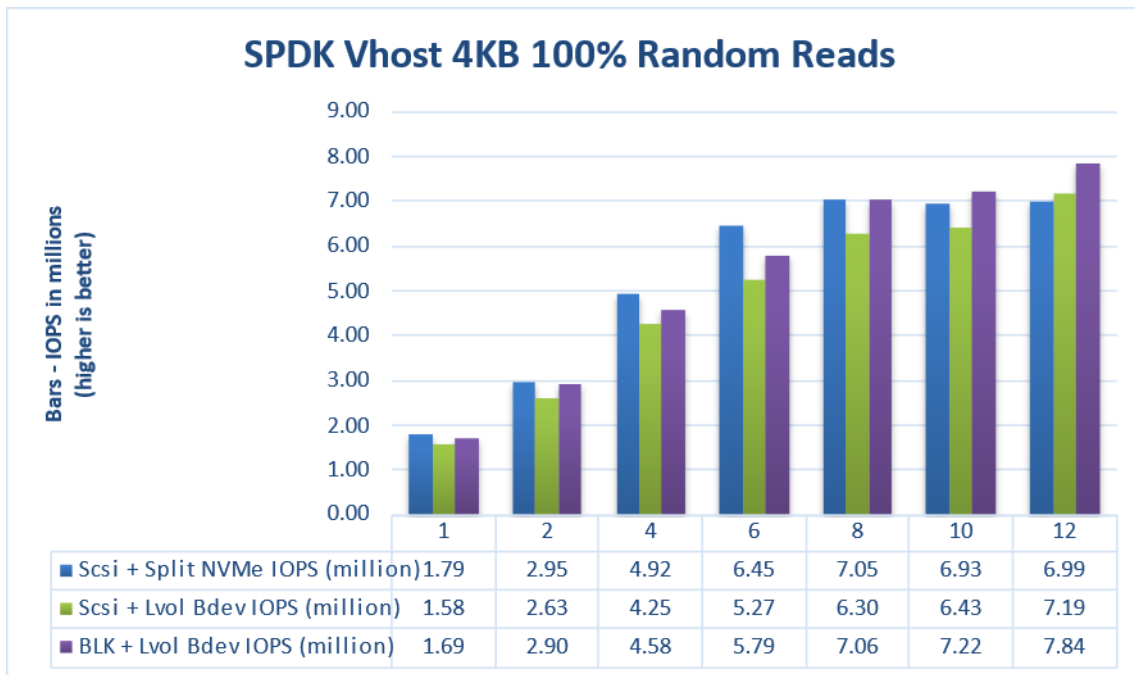|	1	2	4	6	8	10	12
Scsi + Split NVMe IOPS (million)	1.79	2.95	4.92	6.45	7.05	6.93	6.99
Scsi + Lvol Bdev IOPS (million)	1.58	2.63	4.25	5.27	6.30	6.43	7.19
BLK + Lvol Bdev IOPS (million)	1.69	2.90	4.58	5.79	7.06	7.22	7.84

Figure 2: Comparison of performance between various SPDK Vhost stack-bdev combinations for 4KB Random Read QD=64 workload

4KB Random Write Results

Table 6: SPDK Vhost core scaling results, 4KB 100% Random Write IOPS, QD=32

# of CPU cores	# of VMs	Stack / Backend	IOPS (millions)
1	6	SCSI / Split NVMe Bdev	1.68
		SCSI / Lvol Bdev	1.59
		BLK / Lvol Bdev	1.65
2	12	SCSI / Split NVMe Bdev	2.97
		SCSI / Lvol Bdev	2.73
		BLK / Lvol Bdev	2.97
4	24	SCSI / Split NVMe Bdev	4.89
		SCSI / Lvol Bdev	4.54
		BLK / Lvol Bdev	4.64
6	36	SCSI / Split NVMe Bdev	6.02
		SCSI / Lvol Bdev	5.59
		BLK / Lvol Bdev	6.05
8	36	SCSI / Split NVMe Bdev	6.63
		SCSI / Lvol Bdev	6.36
		BLK / Lvol Bdev	6.96
10	36	SCSI / Split NVMe Bdev	6.42
		SCSI / Lvol Bdev	6.28
		BLK / Lvol Bdev	6.86
12	36	SCSI / Split NVMe Bdev	6.57
		SCSI / Lvol Bdev	6.71
		BLK / Lvol Bdev	7.36

SPDK Vhost 4KB 100% Random Writes

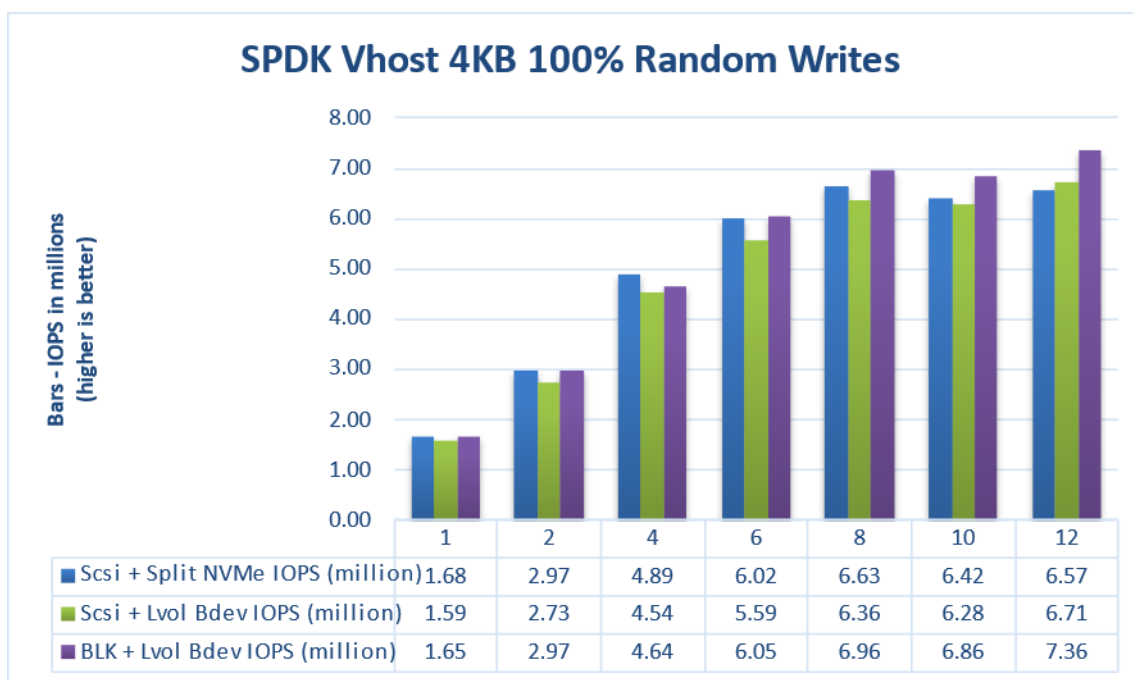|	1	2	4	6	8	10	12
Scsi + Split NVMe IOPS (million)	1.68	2.97	4.89	6.02	6.63	6.42	6.57
Scsi + Lvol Bdev IOPS (million)	1.59	2.73	4.54	5.59	6.36	6.28	6.71
BLK + Lvol Bdev IOPS (million)	1.65	2.97	4.64	6.05	6.96	6.86	7.36

Figure 3: Comparison of performance between various SPDK Vhost stack-bdev combinations for 4KB Random Write QD=32 workload

4KB Random Read-Write Results

Table 7: SPDK Vhost core scaling results, 4KB Random 70% Read 30% Write IOPS, QD=64

| # of CPU cores	# of VMs	Stack / Backend	IOPS (millions)
1	6	SCSI / Split NVMe Bdev	1.65
		SCSI / Lvol Bdev	1.52
		BLK / Lvol Bdev	1.57
2	12	SCSI / Split NVMe Bdev	2.81
		SCSI / Lvol Bdev	2.51
		BLK / Lvol Bdev	2.78
4	24	SCSI / Split NVMe Bdev	4.64
		SCSI / Lvol Bdev	4.15
		BLK / Lvol Bdev	4.51
6	36	SCSI / Split NVMe Bdev	6.13
		SCSI / Lvol Bdev	5.34
		BLK / Lvol Bdev	5.81
8	36	SCSI / Split NVMe Bdev	6.30
		SCSI / Lvol Bdev	5.85
		BLK / Lvol Bdev	6.39
10	36	SCSI / Split NVMe Bdev	6.05
		SCSI / Lvol Bdev	5.95
		BLK / Lvol Bdev	6.25
12	36	SCSI / Split NVMe Bdev	6.25
		SCSI / Lvol Bdev	6.27
		BLK / Lvol Bdev	6.72
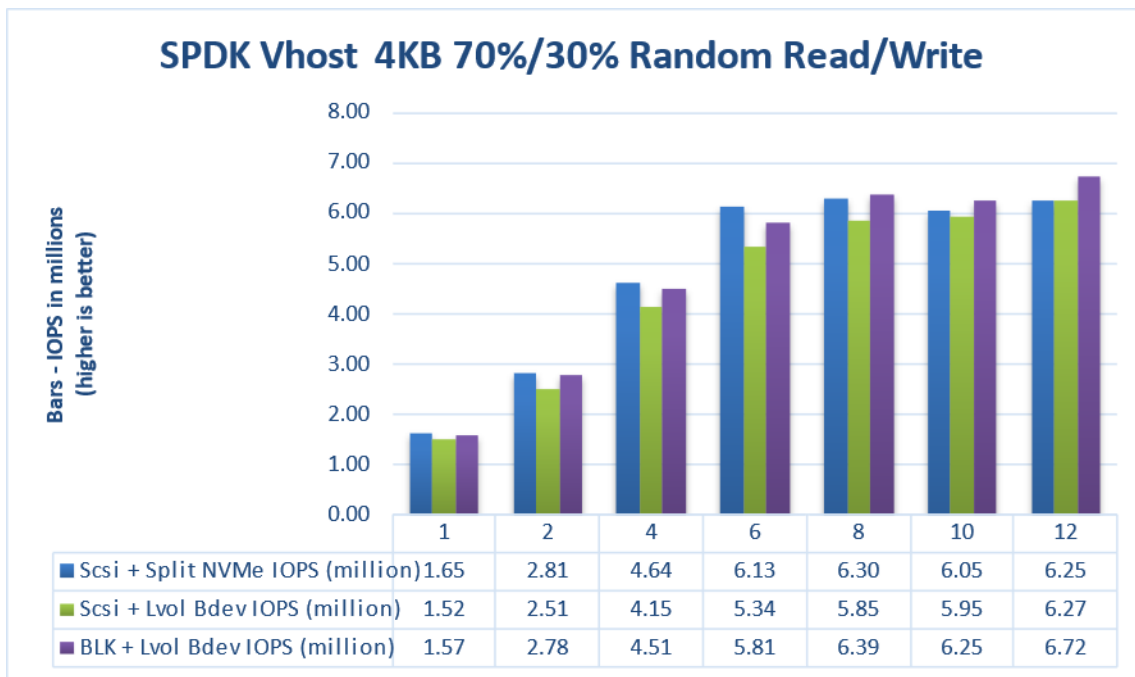
Figure 4: Comparison of performance between various SPDK Vhost stack-bdev combinations for 4KB Random 70% Read 30% Write QD=64 workload

Logical Volumes performance impact

The SPDK Vhost SCSI tests were run using two bdev backends – Split NVMes and Logical Volumes. Both "Split NVMe Bdevs" and "Logical Volume Bdevs" allow to logically partition NVMe SSDs, the latter being more flexible in configuration. Here we measure the overhead of extra flexibility afforded by Logical Volumes.

Table 8: Logical Volumes performance impact for SPDK Vhost SCSI

Workload	# of CPU cores	# of VMs	Vhost SCSI + Split NVMe IOPS (millions)	Vhost SCSI + Lvol IOPS (millions)	Lvol Impact (%)
4KB 100% Random Read	1	6	1.79	2.63	-11.96%
	2	12	2.95	4.25	-10.97%
	4	24	4.92	5.27	-13.58%
	6	36	6.45	6.30	-18.35%
	8	36	7.05	6.43	-10.69%
	10	36	6.93	7.19	-7.30%
	12	36	6.99	2.63	2.88%
4KB 100% Random Write	1	6	**1.68**	**1.59**	-5.24%
	2	12	**2.97**	**2.73**	-8.11%
	4	24	**4.89**	**4.54**	-7.16%
	6	36	**6.02**	**5.59**	-7.22%
	8	36	**6.63**	**6.36**	-4.10%
	10	36	**6.42**	**6.28**	-2.18%
	12	36	**6.57**	**6.71**	2.11%
4KB 70% Random Read 30% Random Write	1	6	**1.65**	**1.52**	-7.58%
	2	12	**2.81**	**2.51**	-10.86%
	4	24	**4.64**	**4.15**	-10.48%
	6	36	**6.13**	**5.34**	-13.03%
	8	36	**6.30**	**5.85**	-7.23%
	10	36	**6.05**	**5.95**	-1.56%
	12	36	**6.25**	**6.27**	0.30%

Packed Ring performance impact

Selected test cases were re-run to show benefits of using Packed Rings as an option when configuring SPDK Vhost BLK controllers. For this, an optional parameter "—packed_ring" must be used when creating a SPDK Vhost BLK controller. Packed Ring feature requires QEMU 4.2.0 or later.

Table 9: Packed Ring performance impact on SPDK Vhost BLK controllers.

Workload	# of CPU cores	# of VMs	IOPS (millions) Split Ring	IOPS (millions) Packed Ring	Avg. Latency (usec) Split Ring	Avg. Latency (usec) Packed Ring	Packed Ring IOPS impact (%)	Packed Ring Avg. Latency impact (%)
4KB 100% Random Read QD=64	1	6	1.69	1.68	227.47	228.13	-0.21%	0.29%
	2	12	2.90	2.96	264.22	259.11	2.18%	-1.94%
	4	24	4.58	4.64	331.38	334.38	1.44%	0.91%
	6	36	5.79	5.82	396.92	395.53	0.51%	-0.35%
	8	36	7.06	6.96	326.31	331.56	-1.31%	1.61%
	10	36	7.22	7.54	316.70	308.39	4.42%	-2.62%
	12	36	7.84	7.77	292.77	296.35	-0.86%	1.22%
4KB 100% Random Write QD=32	1	6	1.65	1.54	117.29	127.67	-6.97%	8.85%
	2	12	2.97	2.99	129.01	128.24	0.83%	-0.59%
	4	24	4.64	4.90	167.47	156.90	5.60%	-6.31%
	6	36	6.05	5.99	190.92	193.37	-1.06%	1.28%
	8	36	6.96	7.05	165.07	163.29	1.27%	-1.08%
	10	36	6.86	7.05	167.62	163.79	2.68%	-2.28%
	12	36	7.36	7.57	155.69	152.18	2.91%	-2.25%
4KB 70% Random Read 30% Random Write QD=64	1	6	1.57	1.56	244.23	246.09	-0.87%	0.76%
	2	12	2.78	2.84	276.19	269.48	2.12%	-2.43%
	4	24	4.51	4.66	341.96	328.42	3.28%	-3.96%
	6	36	5.81	5.80	395.44	400.33	-0.21%	1.24%
	8	36	6.39	6.55	361.50	352.64	2.64%	-2.45%
	10	36	6.25	6.40	367.28	359.71	2.50%	-2.06%
	12	36	6.72	6.86	343.02	337.82	2.05%	-1.52%

Conclusions

1. For SPDK Vhost SCSI performance with split NVMe bdevs, we measured 1.68 million IOPS on one Vhost core for the 4KB 100% Random Write workload. The single Vhost core IOPS for 4 KB Random Read and 4KB Random 70/30 Read/Write were 1.79 million and 1.65 million IOPS respectively. For all workloads, the IOPS scaled near linearly with addition of I/O processing cores up to 6 CPU cores. Peak performance was achieved at 8 CPU cores for all workloads. Further increasing the number of cores does not result in performance improvement or it is not significant.

2. For SPDK Vhost SCSI with Logical Volume backend devices, we measured about 1.5 million IOPS on one Vhost core for all 3 workloads. Performance scaled near linearly with addition of I/O processing cores up to 8 CPU cores for Random Read and up to 6 CPU cores for other workloads. Increasing the number of I/O processing cores further results in non-linear IOPS gains. Peak performance is reached at 12 CPU cores for all workloads.

3. For SPDK Vhost BLK with Logical Volume backend devices, we measured about 1.6 million IOPS on one Vhost core for all workloads. Performance scaled near linearly up to 8 CPU cores for Random Read workload, and up to 6 CPU cores for Random Read/Write and Random Write workloads. Increasing the number of cores improves performance further, but the gains are not linear.

4. Using Logical Volumes has an impact of 5-18% lower IOPS when Vhost saturates the I/O cores at 6 or less CPUs. Further increasing SPDK Vhost CPU cores allow Logical Volumes to perform better, and their performance is on par with Split NVMe Bdevs (less than 10% difference).

5. For some workloads there is a slight performance drop when Vhost is run with 10 or 12 CPU cores. The platform has 80 CPU threads available, and 36 VMs use 72 CPU threads. Therefore, when 10 or 12 are used for the Vhost process there is not enough left to accommodate all the VMs. Some of the VMs share CPU threads, thus becoming less efficient.

6. Using Packed Ring option instead of default Split Ring mode for SPDK Vhost BLK controllers results in minor performance improvement.

Test Case 2: Rate Limiting IOPS per VM

This test case was geared towards understanding how many VMs can be supported at a pre-defined Quality of Service of IOPS per Vhost device. Both read and write IOPS were rate limited for each Vhost device on each of the VMs and then VM density was compared between SPDK & the Linux Kernel. 10K IOPS per VM were chosen as the rate limiter using Linux cgroups2.

Each individual VM was running FIO with the following workloads:

- 4KB 100% Random Read

- 4KB 100% Random Write

The results in tables are average of 3 runs.

Table 10: Rate Limiting IOPS per VM test case configuration

Item	Description
Test case	Test rate limiting IOPS/VM to 10000 IOPS

Test configuration	**FIO Version:** fio-3.19
	
	VM Configuration:
	
	• Common settings are described in the **Virtual Machine Settings** chapter.
	• Number of VMs: 24, 48, and 72
	• Each VM has a single Vhost device which is one of equal partitions of NVMe drive. Total number of partitions depends on run test case.
	    ○ For 24 VMs: 24xNVMe * 1 partition per NVMe = 24 partitions
	    ○ For 48 VMs: 24xNVMe * 2 partitions per NVMe = 48 partitions
	    ○ For 72 VMs: 24xNVMe * 3 partitions per NVMe = 72 partitions
	• Devices on VMs were throttled to run at a maximum of 10k IOPS (read or write)
	
	SPDK Vhost target configuration:
	• Test were run with both Vhost-scsi and Vhost-blk stacks.
	• The Vhost-scsi stack was run with Split NVMe bdevs and Logical Volume bdevs.
	• The Vhost-blk stack was run with Logical Volume bdevs.
	• Test were run with the Vhost target using 4 CPU cores (NUMA optimized).
	
	Kernel Vhost-scsi configuration:
	• Cgroups were used to limit the Vhost process to 4 cores.
	• NUMA optimizations were not explored.

FIO configuration run on each VM	[global] ioengine=libaio direct=1 rw=randrw

	rwmixread=100 (100% reads), 0 (100% writes) thread=1 norandommap=1 time_based=1 runtime=300s ramp_time=10s bs=4k iodepth=1 numjobs=1	17

Test Case 2 Results

Table 11: 4KB 100% Random Reads QD=1 rate limiting test results, 4 Vhost CPU cores

# of VMs	Stack	Backend bdev	IOPS (k)	Avg Lat. (usec)
24 VMs	SPDK-SCSI	Split NVMe	239.96	99.45
	SPDK-SCSI	Logical Volume	239.97	99.45
	SPDK-BLK	Logical Volume	239.96	99.47
	Kernel-SCSI	Partitioned NVMe	177.95	133.84
48 VMs	SPDK-SCSI	Split NVMe	479.82	99.11
	SPDK-SCSI	Logical Volume	479.71	99.13
	SPDK-BLK	Logical Volume	479.82	99.14
	Kernel-SCSI	Partitioned NVMe	208.67	232.93
72 VMs	SPDK-SCSI	Split NVMe	682.52	104.45
	SPDK-SCSI	Logical Volume	665.70	107.12
	SPDK-BLK	Logical Volume	684.15	104.16
	Kernel-SCSI	Partitioned NVMe	292.63	249.17
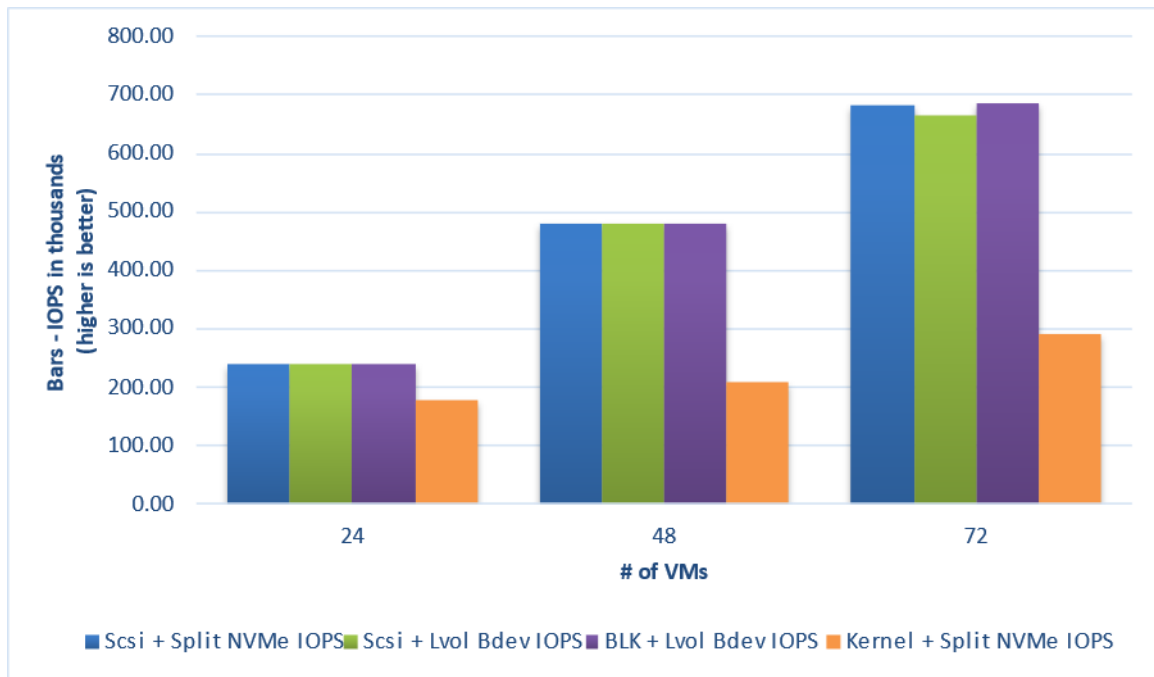
Figure 5: 4KB 100% Random Reads IOPS, QD=1, throttling = 10k IOPS, 4 Vhost CPU cores

Table 12: 4KB 100% Random Writes QD=1 rate limiting test results

# of VMs	Stack	Backend bdev	IOPS (k)	Avg Lat. (usec)
24 VMs	SPDK-SCSI	Split NVMe	239.99	99.46
	SPDK-SCSI	Logical Volume	239.99	99.45
	SPDK-BLK	Logical Volume	240.00	99.46
	Kernel-SCSI	Partitioned NVMe	235.48	101.31
48 VMs	SPDK-SCSI	Split NVMe	479.99	99.22
	SPDK-SCSI	Logical Volume	479.99	99.24
	SPDK-BLK	Logical Volume	479.98	99.23
	Kernel-SCSI	Partitioned NVMe	261.34	182.52
72 VMs	SPDK-SCSI	Split NVMe	719.93	99.11
	SPDK-SCSI	Logical Volume	719.96	99.14
	SPDK-BLK	Logical Volume	719.95	99.12
	Kernel-SCSI	Partitioned NVMe	322.70	221.85
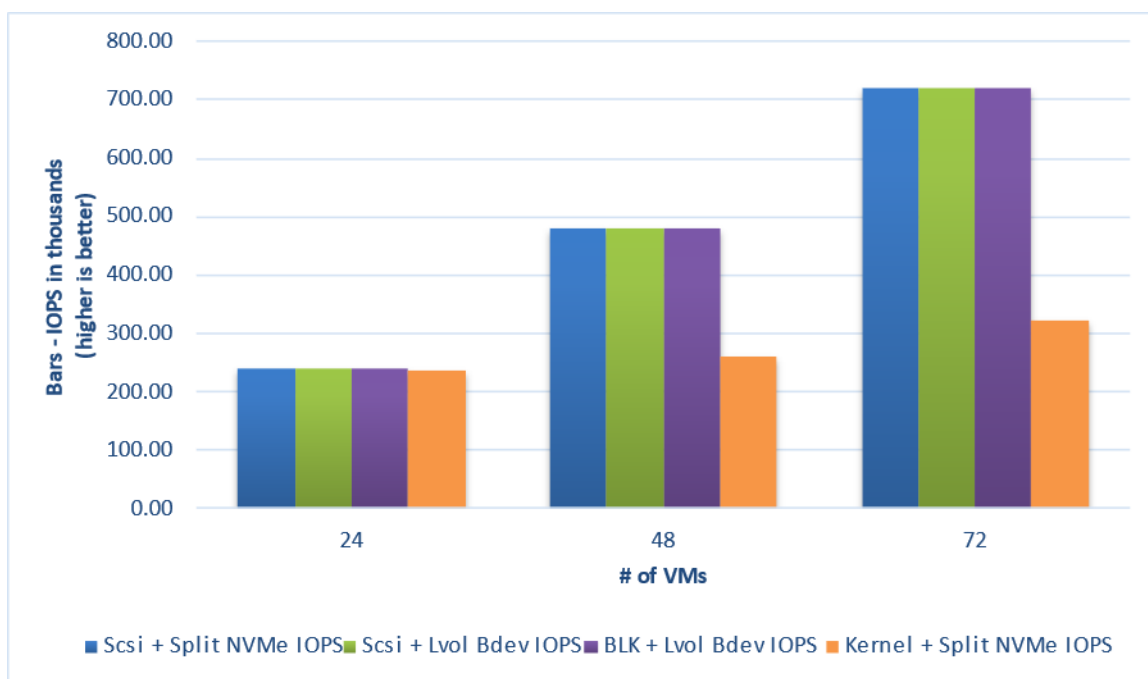
Figure 6: 4KB 100% Random Writes IOPS, QD=1, throttling = 10k IOPS, 4 Vhost CPU cores

Conclusions

1. Using just 4 I/O processing cores, the SPDK vhost served 10,000 IOPS/VM to up to 48 VMs for 4 KB Random Read workload and 72 VMs for the 4 KB Random Write workload.

2. The Kernel Vhost was not able to serve IO at 10K IOPS/VM with just 4 I/O processing cores to 24 VMs for the 4KB Random Read workload. For the 4K Random Write the Kernel Vhost target served 10K IOPS/VM to up to 24 VMs

3. Average latencies were up to 2.4x times better for Random Read and up to 2.24x times better for Random Write workloads with the SPDK Vhost when compared to Kernel Vhost.

Test Case 3: Performance per NVMe drive

This test case was performed to understand performance and efficiency of the Vhost scsi and blk process using SPDK vs. Linux Kernel with a single NVMe drive on 2 VMs. Each VM had a single Vhost device which is one of two equal partitions of an NVMe drive. Results in the table represent performance (IOPS, avg. latency & CPU %) seen from the VM. The VM was running FIO with the following workloads:

- 4KB 100% Random Read

- 4KB 100% Random Write

- 4KB Random 70% Read 30% Write

The results in tables are average of 3 runs.

Table 13: Performance per NVMe drive test case configuration

Item	Description
Test case	Test SPDK Vhost target I/O core scaling performance
Test configuration	**FIO Version:** fio-3.19 **VM Configuration**: • Common settings are described in the **Virtual Machine Settings** chapter. • 2 VMs were tested • Each VM had a single Vhost device which was one of two equal partitions of a single NVMe drive. **SPDK Vhost target configuration:** • The SPDK Vhost process was run on a single, physical CPU core. • The Vhost-scsi stack was run with Split NVMe bdevs and Logical Volume bdevs. • The Vhost-blk stack was run with Logical Volume bdevs. **Kernel Vhost target configuration**: • The Vhost process was run on a single, physical CPU core using cgroups.
FIO configuration	[global] ioengine=libaio direct=1 rw=randrw rwmixread=100 (100% reads), 70 (70% reads, 30% writes), 0 (100% writes) thread=1 norandommap=1 time_based=1 runtime=240s ramp_time=60s bs=4k iodepth=1 / 8 / 32 / 64 numjobs=1

Test Case 3 results

SPDK Vhost-Scsi

Table 14:Performance per NVMe drive IOPS and latency results, SPDK SCSI stack

Access pattern	Backend	QD	Throughput (IOPS)	Avg. latency (usec)
4k 100% Random Reads	Split NVMe	1	24829.41	79.97
4k 100% Random Reads	Split NVMe	8	170781.80	92.86
4k 100% Random Reads	Split NVMe	32	440347.86	145.11
4k 100% Random Reads	Split NVMe	64	561230.44	227.71
4k 100% Random Reads	Lvol	1	24831.94	79.82
4k 100% Random Reads	Lvol	8	170250.60	93.51
4k 100% Random Reads	Lvol	32	440847.41	144.95
4k 100% Random Reads	Lvol	64	564306.29	226.40
4k 100% Random Writes	Split NVMe	1	115538.24	16.91
4k 100% Random Writes	Split NVMe	8	344421.30	47.09
4k 100% Random Writes	Split NVMe	32	435466.32	146.68
4k 100% Random Writes	Split NVMe	64	468421.22	274.35
4k 100% Random Writes	Lvol	1	112581.02	17.02
4k 100% Random Writes	Lvol	8	347167.88	46.63
4k 100% Random Writes	Lvol	32	446125.64	144.13
4k 100% Random Writes	Lvol	64	462097.90	277.10
4k 70%/30% Random Read Writes	Split NVMe	1	31499.25	62.97
4k 70%/30% Random Read Writes	Split NVMe	8	164240.33	96.35
4k 70%/30% Random Read Writes	Split NVMe	32	320792.40	198.96
4k 70%/30% Random Read Writes	Split NVMe	64	372622.50	344.45
4k 70%/30% Random Read Writes	Lvol	1	31391.62	62.96
4k 70%/30% Random Read Writes	Lvol	8	166538.08	95.65
4k 70%/30% Random Read Writes	Lvol	32	323623.26	194.87
4k 70%/30% Random Read Writes	Lvol	64	381515.71	335.39

SPDK Vhost Performance Report
Release 22.09

SPDK Vhost-Blk

Table 15: Performance per NVMe drive IOPS and latency results, SPDK BLK stack

Access pattern	Backend	QD	Throughput (IOPS)	Avg. latency (usec)
4k 100% Random Reads	Lvol	1	24952.43	79.48
4k 100% Random Reads	Lvol	8	172920.18	92.37
4k 100% Random Reads	Lvol	32	444078.21	143.97
4k 100% Random Reads	Lvol	64	570932.78	223.79
4k 100% Random Writes	Lvol	1	115465.22	17.30
4k 100% Random Writes	Lvol	8	347589.02	46.55
4k 100% Random Writes	Lvol	32	440554.61	145.68
4k 100% Random Writes	Lvol	64	470142.54	272.96
4k 70%/30% Random Read Writes	Lvol	1	31578.01	63.02
4k 70%/30% Random Read Writes	Lvol	8	167397.88	93.23
4k 70%/30% Random Read Writes	Lvol	32	324513.62	202.36
4k 70%/30% Random Read Writes	Lvol	64	392410.28	325.35

Kernel Vhost-Scsi

Table 16: Performance per NVMe drive IOPS and latency results, Kernel Vhost-Scsi

Access pattern	Backend	QD	Throughput (IOPS)	Avg. latency (usec)
4k 100% Random Reads	NVMe	1	20752.16	95.96
4k 100% Random Reads	NVMe	8	149623.76	109.66
4k 100% Random Reads	NVMe	32	363324.63	175.27
4k 100% Random Reads	NVMe	64	451597.65	283.19
4k 100% Random Writes	NVMe	1	70425.18	28.07
4k 100% Random Writes	NVMe	8	238329.93	68.23
4k 100% Random Writes	NVMe	32	387295.41	162.69
4k 100% Random Writes	NVMe	64	441972.62	289.09
4k 70%/30% Random Read Writes	NVMe	1	25944.69	76.55
4k 70%/30% Random Read Writes	NVMe	8	130327.59	121.22
4k 70%/30% Random Read Writes	NVMe	32	281273.03	226.20
4k 70%/30% Random Read Writes	NVMe	64	354091.53	360.86
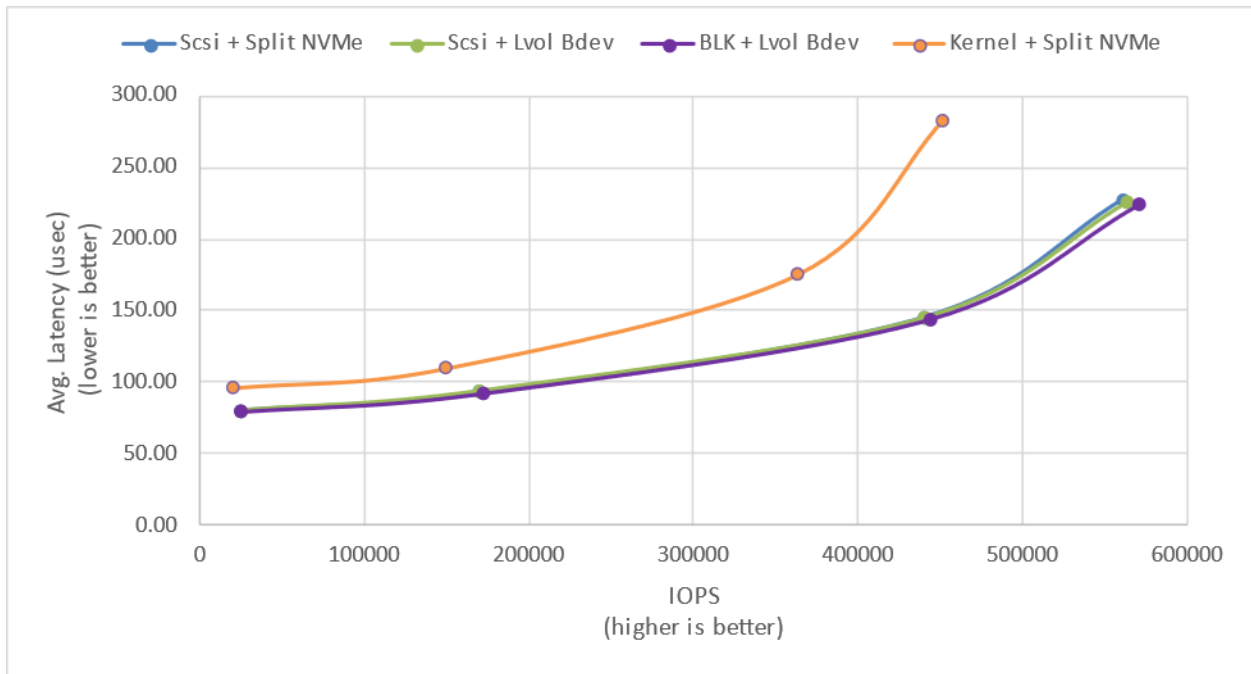
Figure 7: 4KB 100% Random Reads IOPS and latency
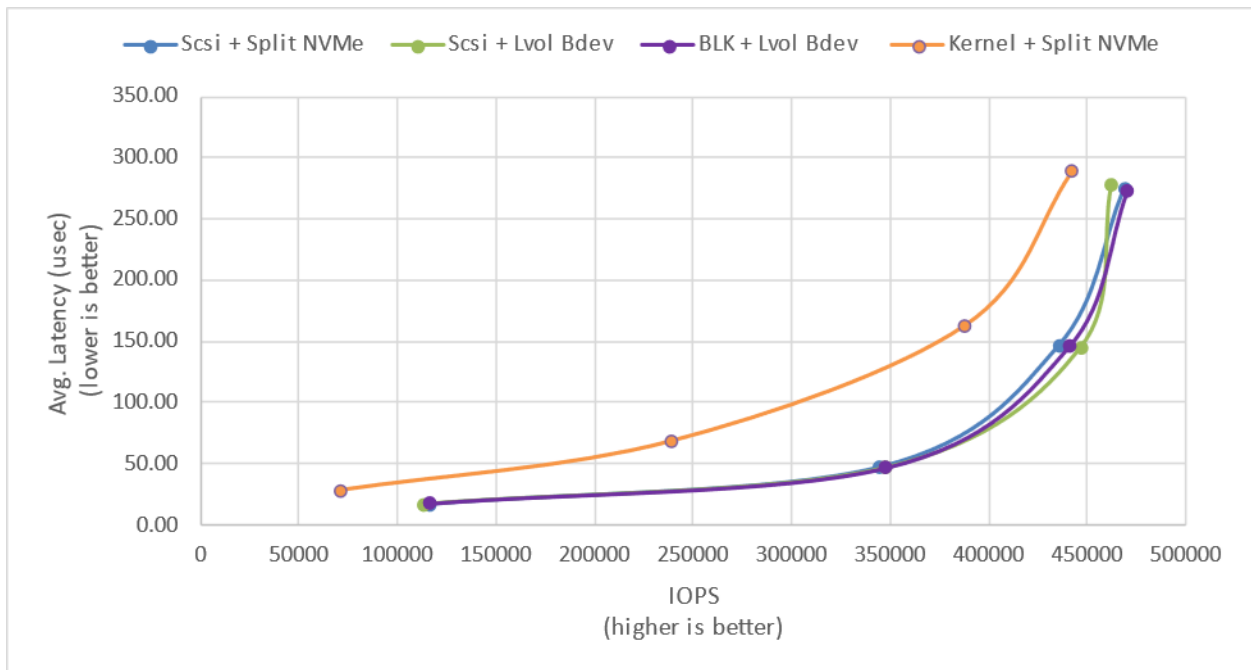
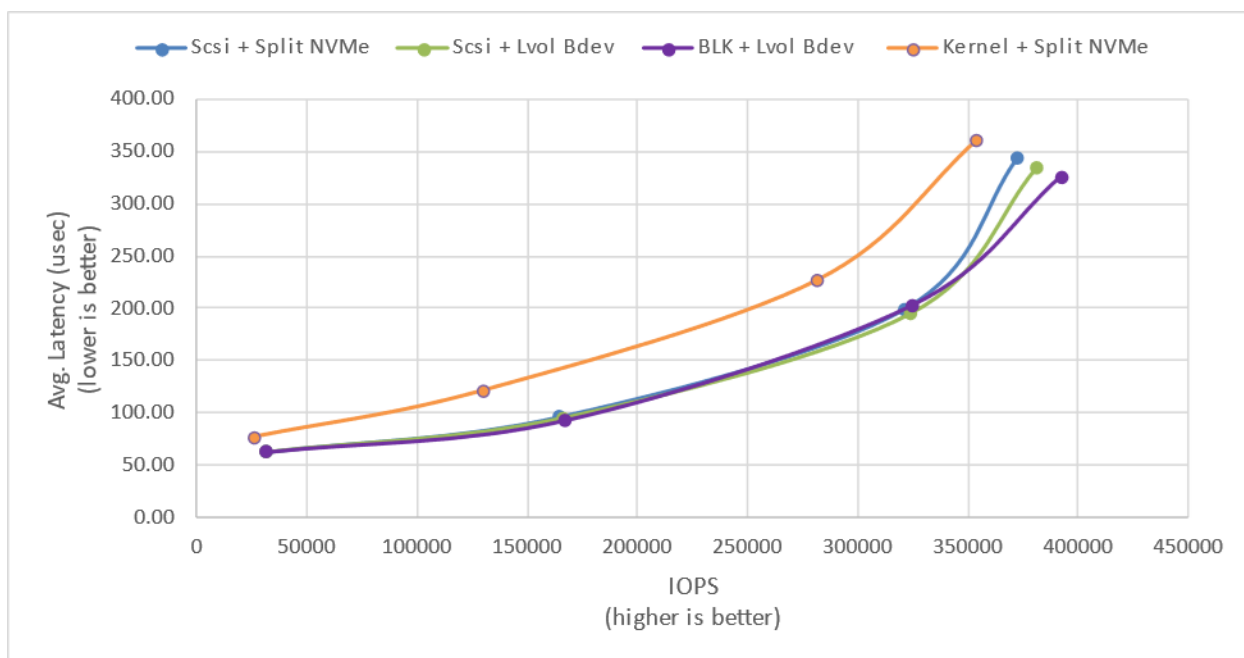*Figure 8: 4KB 100% Random Writes IOPS and latency*

intel.

Figure 9: 4KB 70%/30% Random Read/Write IOPS and latency

Conclusions

1. SPDK Vhost-scsi with NVMe Split bdevs has lower latency and higher throughput than Kernel Vhost-scsi in majority of workload / queue depth combinations.

Summary

This report compared performance results while running Vhost-scsi using traditional interrupt-driven kernel Vhost-scsi against the accelerated polled-mode driven SPDK implementation. Various local ephemeral configurations were demonstrated, including rate limiting IOPS, performance per VM, and maximum performance from an underlying system when comparing kernel vs. SPDK Vhost-scsi target implementations.

In addition, performance impacts of using SPDK Logical Volume Bdevs and the SPDK Vhost-blk stack were presented.

This report provided information regarding methodologies and practices while benchmarking Vhost-scsi and Vhost-blk using both SPDK and 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.

Appendix A - Vfio-User Performance

Vfio-user is a framework that allows implementing PCI devices in userspace. Clients, such as Qemu, talk to the server over a UNIX socket using vfio-user protocol. For more information see: libvfio-user.

SPDK has added a vfio-user transport to the nvmf target. This allows presenting a virtual NVMe PCIe device to a client, such as a Qemu VM. This is analogous to how the SPDK vhost target presents virtual virtio PCIe devices to a Qemu VM.

The performance comparison between SPDK Vhost stacks versus vfio-user transport layer was done using oracle/qemu_fork (vfio-user-irqmask2 branch), which supports running Qemu guest VMs with "vfio-user-pci" devices.

In SPDK v22.09, the vfio-user transport performance is roughly comparable with SPDK Vhost-SCSI results with increasing number of SPDK target application CPU cores assigned. There is one exception for data point where 4 CPU cores are used to run SPDK target application, where using Vfio-User transport layer yields noticeably worse performance results than using SPDK Vhost-SCSI. We were not able explain this discrepancy and investigations are still a work in progress.

SPDK Vhost 4KB 100% Random Reads

	1	2	4	6	8	10	12
■ Scsi + Split NVMe IOPS (million)	1.83	3.04	4.91	6.42	6.90	6.58	6.76
■ Vfio + Split NVMe IOPS (million)	2.00	3.36	3.36	5.58	6.20	6.82	7.17
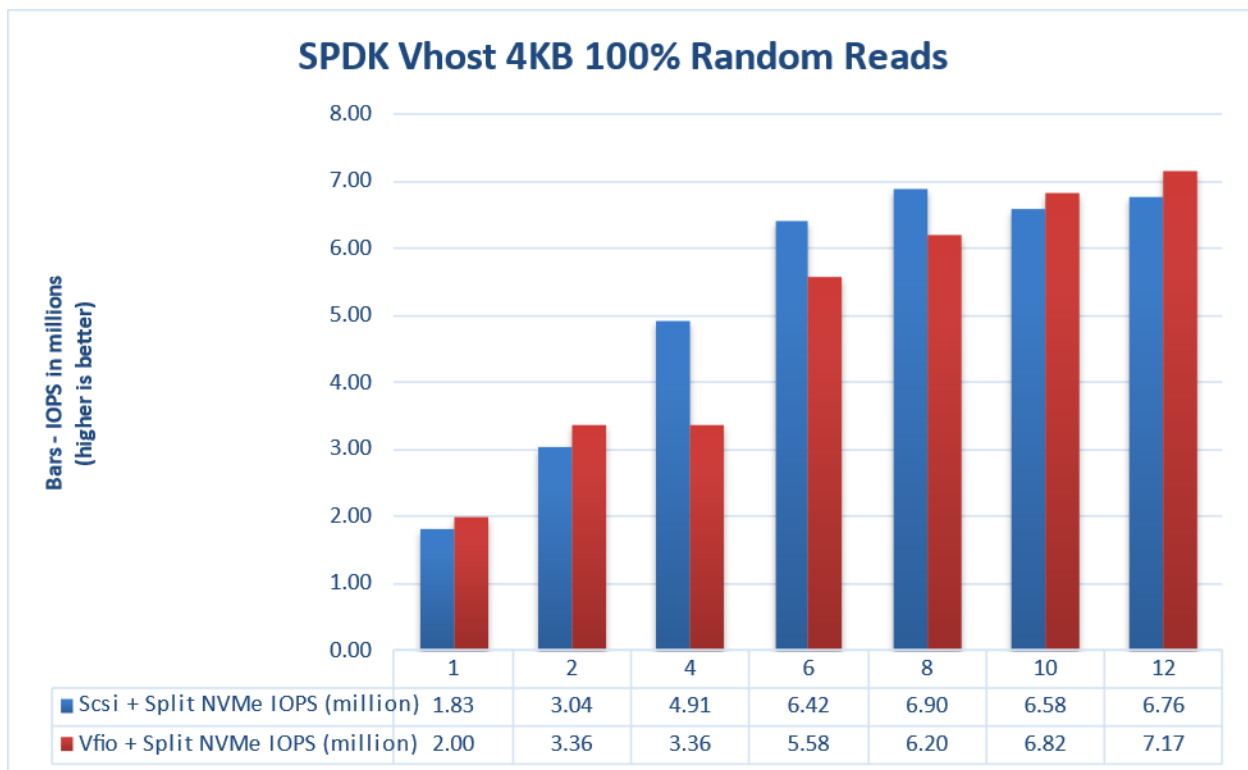
Figure 10: Comparison of performance between SPDK Vhost-SCSI and Vfio-User transport for 4KB Random Read QD=64 workload

SPDK Vhost 4k 100% Random Writes

	1	2	4	6	8	10	12
■ Scsi + Split NVMe IOPS (million)	1.76	3.08	5.07	6.14	6.79	6.36	6.52
■ Vfio + Split NVMe IOPS (million)	1.95	3.31	3.81	6.16	6.10	6.66	7.42
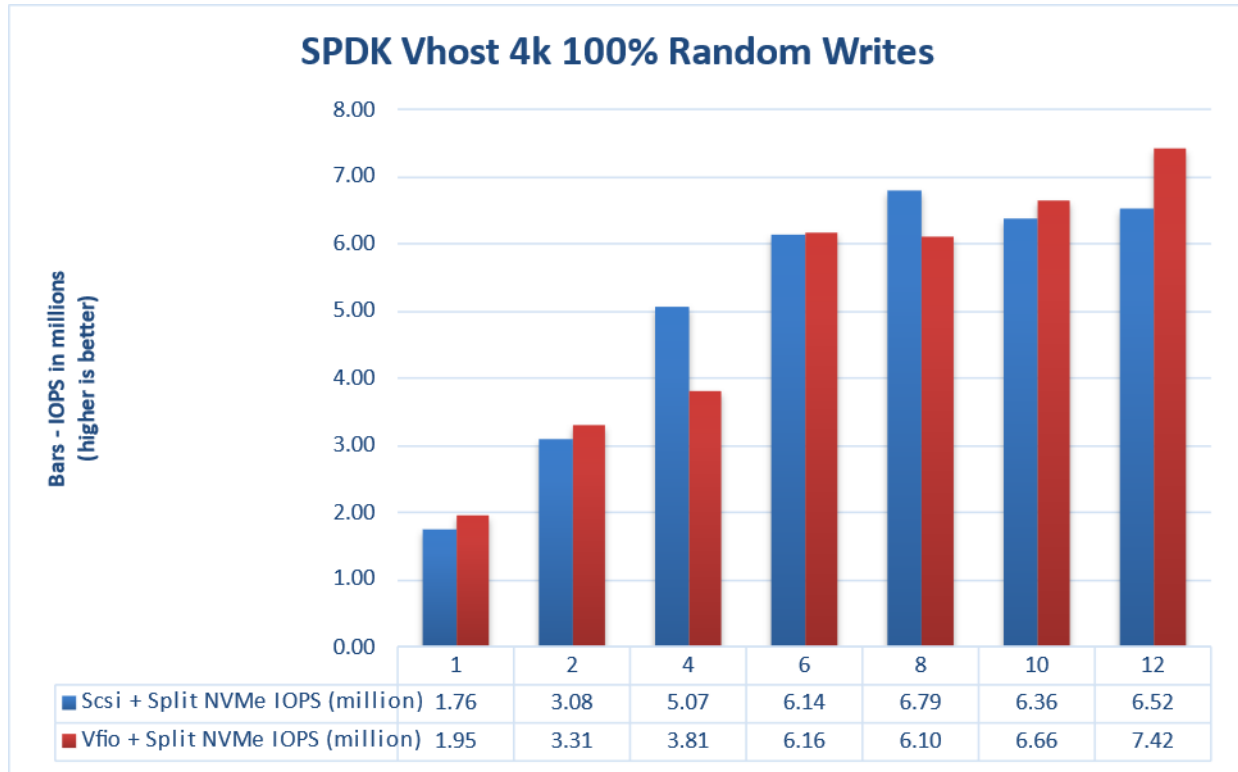
Figure 11: Comparison of performance between SPDK Vhost-SCSI and Vfio-User transport for 4KB Random Write QD=32 workload

SPDK Vhost 4KB 70%/30% Random Read/Write

	1	2	4	6	8	10	12
■ Scsi + Split NVMe IOPS (million)	1.66	2.85	4.71	6.17	6.25	5.79	6.10
■ Vfio + Split NVMe IOPS (million)	1.77	3.23	3.60	6.28	5.87	6.50	6.71
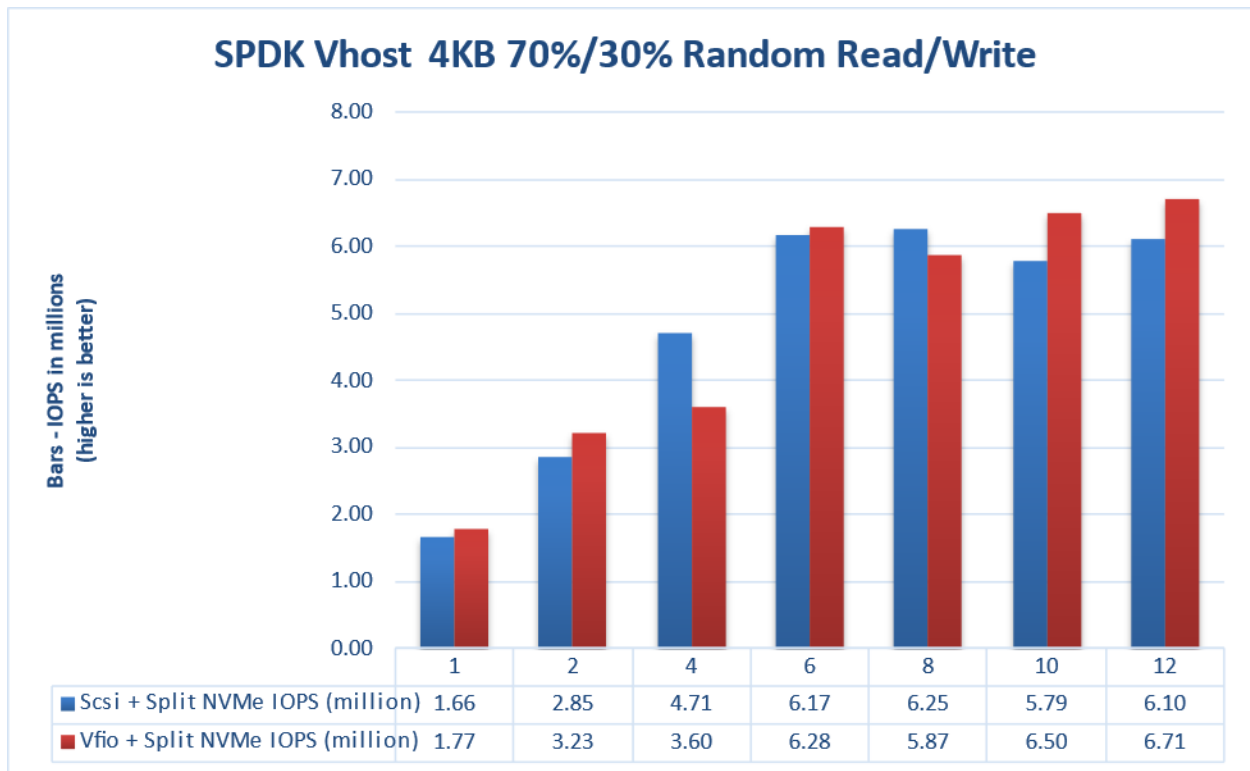
Figure 12: Comparison of performance between SPDK Vhost-SCSI and Vfio-User transport for 4KB Random Read/Write QD=64 workload

List of Tables

List of Figures

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. See backup for configuration details. No product or component can be absolutely secure.

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.

§