

SPDK Vhost performance report

Release 20.01

Testing Date: March 2020

Performed by: Karol Latecki

Acknowledgments:

Seth Howell (seth.howell@intel.com)


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

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



Contents

Contents	2
Audience and Purpose.....	3
Test setup	4
Hardware configuration	4
BIOS Settings	5
Virtual Machine Settings.....	5
Kernel & BIOS Spectre-Meltdown information	6
Introduction to the SPDK Vhost target	7
SPDK Vhost target architecture	7
Test Case 1: SPDK Vhost core scaling	9
4KB Random Read Results	11
4KB Random Write Results	13
4KB Random Read-Write Results.....	15
Logical Volumes performance impact	17
LTO performance impact	18
Conclusions	19
Test Case 2: Rate Limiting IOPS per VM.....	20
Test Case 2 Results	22
Conclusions	24
Test Case 3: Performance per NVMe drive	25
Test Case 3 results.....	26
Conclusions	29
Summary	30




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

Item	Description												
Server Platform	Intel WolfPass R2224WFTZS 												
Motherboard	S2600WFT												
CPU	Intel(R) Xeon(R) Gold 6230N CPU @ 2.30GHz Number of cores 20, number of threads 40												
Memory	10 x 32GB Micron DDR4 36ASF4G72PZ-2G6H1R Total 320 GBs Memory channel population: <table border="1" data-bbox="479 1291 1429 1543"> <thead> <tr> <th>P1</th> <th>P2</th> </tr> </thead> <tbody> <tr> <td>CPU1_DIMM_A1</td> <td>CPU2_DIMM_A1</td> </tr> <tr> <td>CPU1_DIMM_B1</td> <td>CPU2_DIMM_B1</td> </tr> <tr> <td>CPU1_DIMM_C1</td> <td>CPU2_DIMM_C1</td> </tr> <tr> <td>CPU1_DIMM_D1</td> <td>CPU2_DIMM_D1</td> </tr> <tr> <td>CPU1_DIMM_E1</td> <td>CPU2_DIMM_E1</td> </tr> </tbody> </table>	P1	P2	CPU1_DIMM_A1	CPU2_DIMM_A1	CPU1_DIMM_B1	CPU2_DIMM_B1	CPU1_DIMM_C1	CPU2_DIMM_C1	CPU1_DIMM_D1	CPU2_DIMM_D1	CPU1_DIMM_E1	CPU2_DIMM_E1
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												
Operating System	Fedora 29												
BIOS	02.01.0010 (06.01.2020)												
Linux kernel version	5.4.14-100.fc30.x86_64												
SPDK version	SPDK 20.01												
Qemu version	QEMU emulator version 3.1.1 (qemu-3.1.1-2.fc30)												
Storage	OS: 1x 120GB Intel SSDSC2BB120G4 Storage: 24x Intel® P4610™ 1.6TBs (FW: VDV10152) (6 on CPU NUMA Node 0, 18 on CPU NUMA Node 1)												



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

Common settings used for all VMs used in tests.

Item	Description
CPU	2vCPU, pass through from physical host server. Explicit core usage enforced using “taskset –a –c” command. Related QEMU arguments used 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. Related 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 29
Linux kernel version	5.1.20-200.fc29.x86_64
Additional boot options in /etc/default/grub	<ul style="list-style-type: none">• Multi queue enabled: scsi_mod.use_blk_mq=1• Spectre-meltdown patches disabled: spectre_v2=off nopti



Kernel & BIOS Spectre-Meltdown information

Host server system uses 5.4.14 kernel version which is available from the DNF repository. The default Spectre-Meltdown mitigation patches for this kernel version have been left enabled.

The guest VM systems use 5.1.20 kernel version, which is available from the DNF repository. The default Spectre-Meltdown mitigation patches for this kernel version have been disabled on guest systems by adding the following in their `/etc/default/grub` file:

```
spectre_v2=off nopti
```

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 team has leveraged existing SPDK SCSI layer, DPDK Vhost library, QEMU Vhost-scsi and Vhost-user functionality in order to create the high performance SPDK userspace Vhost target.

SPDK Vhost target architecture

QEMU setups Vhost target via UNIX domain socket. The Vhost target transfers data to/from the guest VM via shared memory. 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 on example of virtio-scsi. It should be noted that there is no QEMU intervention during the I/O submission process. 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.

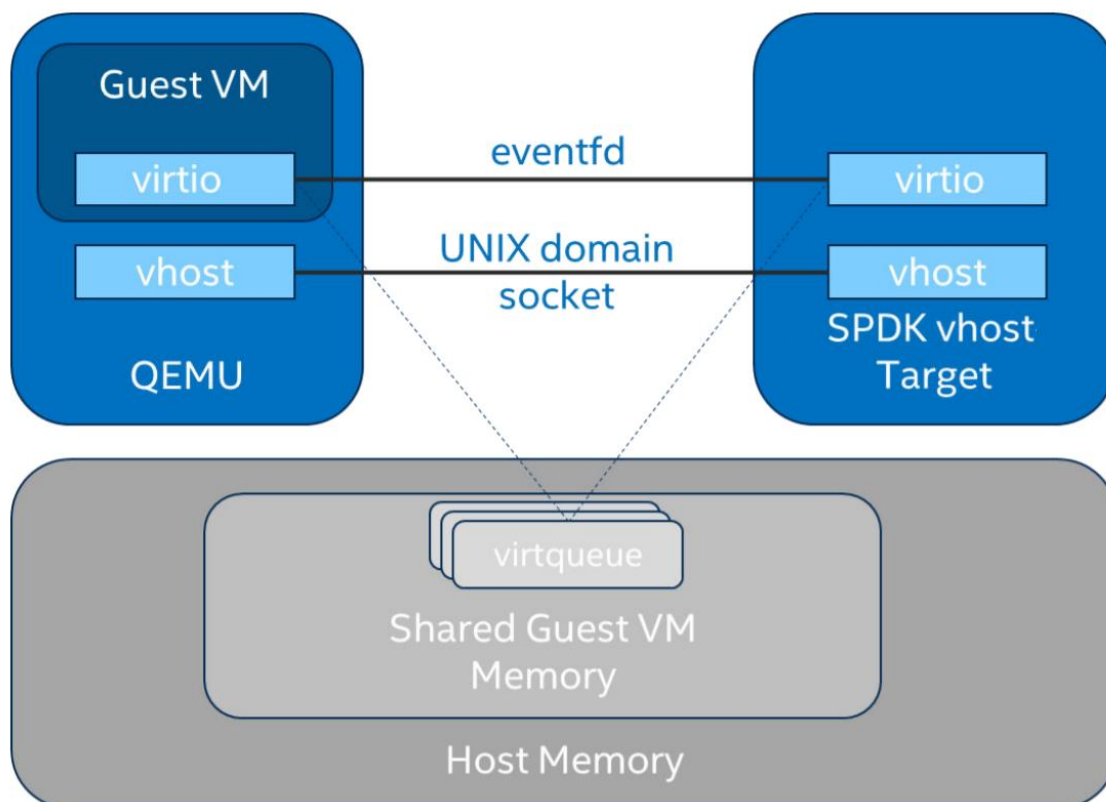



Figure 1: SPDK Vhost-scsi architecture



This report shows the performance comparisons between the traditional interrupt-driven kernel Vhost-scsi and the accelerated polled-mode driven SPDK Vhost-scsi under 4 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 in order 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. The `gtod_reduce=1` option was used to disable FIO latency measurements which allowed better IOPS and bandwidth results.

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

Item	Description
Test case	Test SPDK Vhost target I/O core scaling performance
Test configuration	<p>FIO Version: fio-3.3</p> <p>VM Configuration:</p> <ul style="list-style-type: none"> • 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. <p>SPDK Vhost target configuration:</p> <ul style="list-style-type: none"> • 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 ran with 1,2,3,4,5,6,8,10 and 12 cores for each stack-bdev combination. <p>Kernel Vhost target configuration:</p> <ul style="list-style-type: none"> - N/A



FIO configuration	<pre>[global] ioengine=libaio direct=1 thread=1 norandommap=1 time_based=1 gtod_reduce=1 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, 8, 32, 64}</pre>
--------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------



4KB Random Read Results

Table 1: 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	0.97
		BLK / Lvol Bdev	0.76
2	12	SCSI / Split NVMe Bdev	2.47
		SCSI / Lvol Bdev	1.62
		BLK / Lvol Bdev	1.58
3	18	SCSI / Split NVMe Bdev	3.08
		SCSI / Lvol Bdev	2.27
		BLK / Lvol Bdev	2.13
4	24	SCSI / Split NVMe Bdev	4.07
		SCSI / Lvol Bdev	2.88
		BLK / Lvol Bdev	3.27
5	30	SCSI / Split NVMe Bdev	4.53
		SCSI / Lvol Bdev	3.37
		BLK / Lvol Bdev	3.85
6	36	SCSI / Split NVMe Bdev	5.16
		SCSI / Lvol Bdev	3.75
		BLK / Lvol Bdev	4.34
8	36	SCSI / Split NVMe Bdev	5.23
		SCSI / Lvol Bdev	4.48
		BLK / Lvol Bdev	4.96
10	36	SCSI / Split NVMe Bdev	5.30
		SCSI / Lvol Bdev	5.14
		BLK / Lvol Bdev	6.05
12	36	SCSI / Split NVMe Bdev	4.90
		SCSI / Lvol Bdev	4.64
		BLK / Lvol Bdev	5.49

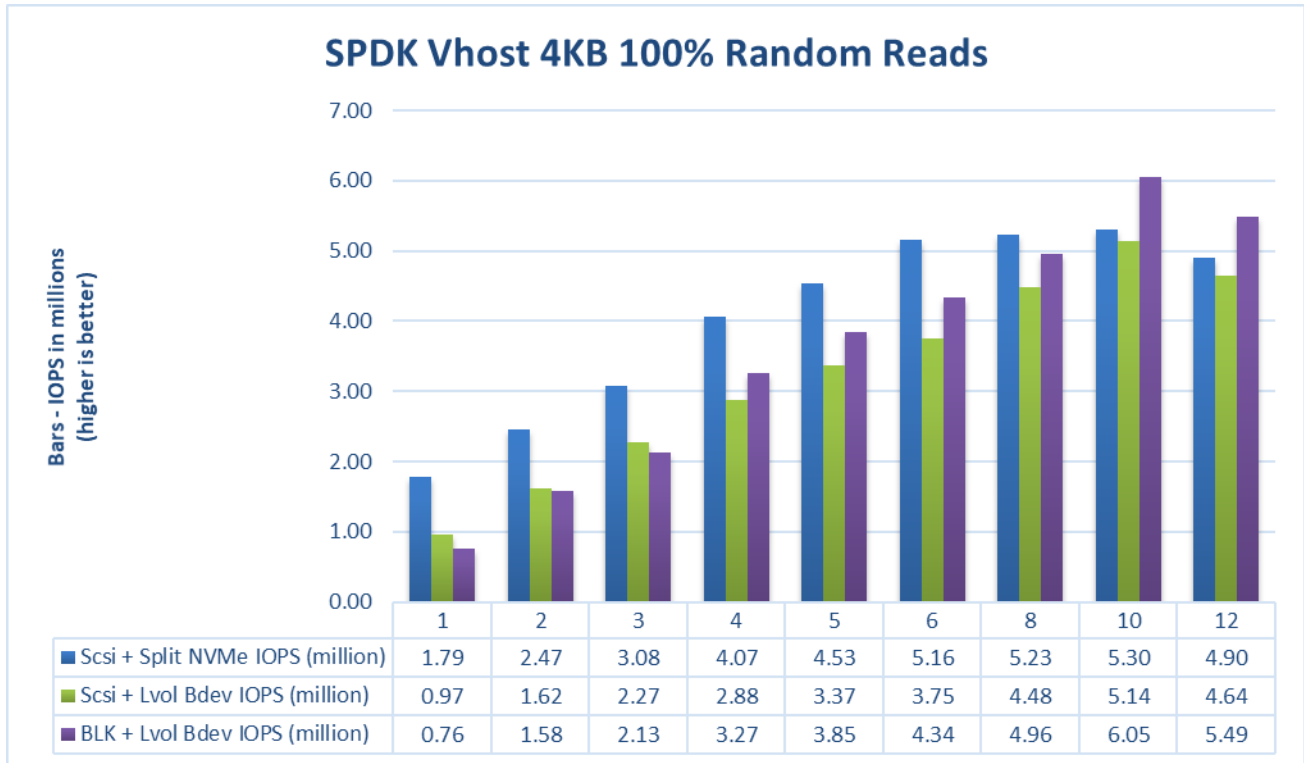


Figure 2: Comparison of performance between various SPDK Vhost stack-bdev combinations for 4KB Random Read QD=64 workload



4KB Random Write Results

Table 2: 4KB 100% Random Write IOPS, QD=32

# of CPU cores	# of VMs	Stack / Backend	IOPS (millions)
1	6	SCSI / Split NVMe Bdev	1.73
		SCSI / Lvol Bdev	1.15
		BLK / Lvol Bdev	1.25
2	12	SCSI / Split NVMe Bdev	2.48
		SCSI / Lvol Bdev	1.43
		BLK / Lvol Bdev	2.22
3	18	SCSI / Split NVMe Bdev	3.42
		SCSI / Lvol Bdev	2.17
		BLK / Lvol Bdev	2.42
4	24	SCSI / Split NVMe Bdev	3.42
		SCSI / Lvol Bdev	2.95
		BLK / Lvol Bdev	3.52
5	30	SCSI / Split NVMe Bdev	4.14
		SCSI / Lvol Bdev	3.77
		BLK / Lvol Bdev	4.11
6	36	SCSI / Split NVMe Bdev	5.52
		SCSI / Lvol Bdev	3.92
		BLK / Lvol Bdev	4.79
8	36	SCSI / Split NVMe Bdev	4.72
		SCSI / Lvol Bdev	4.68
		BLK / Lvol Bdev	5.09
10	36	SCSI / Split NVMe Bdev	4.33
		SCSI / Lvol Bdev	4.45
		BLK / Lvol Bdev	4.51
12	36	SCSI / Split NVMe Bdev	4.33
		SCSI / Lvol Bdev	4.34
		BLK / Lvol Bdev	5.39




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 3: 4KB Random 70% Read 30% Write IOPS, QD=64

# of CPU cores	# of VMs	Stack / Backend	IOPS (millions)
1	6	SCSI / Split NVMe Bdev	1.36
		SCSI / Lvol Bdev	1.25
		BLK / Lvol Bdev	1.47
2	12	SCSI / Split NVMe Bdev	2.36
		SCSI / Lvol Bdev	1.65
		BLK / Lvol Bdev	2.15
3	18	SCSI / Split NVMe Bdev	2.76
		SCSI / Lvol Bdev	2.01
		BLK / Lvol Bdev	2.42
4	24	SCSI / Split NVMe Bdev	3.73
		SCSI / Lvol Bdev	3.01
		BLK / Lvol Bdev	3.37
5	30	SCSI / Split NVMe Bdev	4.14
		SCSI / Lvol Bdev	3.31
		BLK / Lvol Bdev	3.74
6	36	SCSI / Split NVMe Bdev	5.13
		SCSI / Lvol Bdev	3.69
		BLK / Lvol Bdev	4.19
8	36	SCSI / Split NVMe Bdev	4.48
		SCSI / Lvol Bdev	4.46
		BLK / Lvol Bdev	5.24
10	36	SCSI / Split NVMe Bdev	4.69
		SCSI / Lvol Bdev	4.60
		BLK / Lvol Bdev	5.25
12	36	SCSI / Split NVMe Bdev	4.47
		SCSI / Lvol Bdev	4.51
		BLK / Lvol Bdev	5.25




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 NVMe and Logical Volumes which allows to compare their impact on the performance.

Table 4: Logical Volumes performance impact for 4KB Random 70% Read 30% Write IOPS, QD=64 workload


Workload	# of CPU cores	# of VMs	Vhost SCSI + Split NVMe IOPS (millions)	Vhost SCSI + Lvol IOPS (millions)	Lvol to NVMe performance (%)
4KB 100% Random Read	1	6	1.79	0.97	54.17%
	2	12	2.47	1.62	65.67%
	3	18	3.08	2.27	73.52%
	4	24	4.07	2.88	70.71%
	5	30	4.53	3.37	74.37%
	6	36	5.16	3.75	72.67%
	8	36	5.23	4.48	85.68%
	10	36	5.30	5.14	97.01%
4KB 100% Random Write	1	6	1.73	1.15	66.40%
	2	12	2.48	1.43	57.76%
	3	18	3.42	2.17	63.49%
	4	24	3.42	2.95	86.13%
	5	30	4.14	3.77	91.00%
	6	36	5.52	3.92	70.90%
	8	36	4.72	4.68	99.19%
	10	36	4.33	4.45	102.86%
4KB 70% Random Read 30% Random Write	1	6	1.36	1.25	92.07%
	2	12	2.36	1.65	69.65%
	3	18	2.76	2.01	72.85%
	4	24	3.73	3.01	80.66%
	5	30	4.14	3.31	79.87%
	6	36	5.13	3.69	71.84%
	8	36	4.88	4.46	91.42%
	10	36	4.69	4.60	98.10%
12	36	4.47	4.51	101.01%	

LTO performance impact

Selected test cases were re-run with LTO (Link Time Optimization) enabled for SPDK compilation. This should positively impact overall SPDK performance. The following comparison was done using SPDK Vhost SCSI with Logical Volume bdevs.

Table 5: LTO performance impact for 4KB Random 70% Read 30% Write IOPS, QD=64 workload

Workload	# of CPU cores	# of VMs	IOPS (millions) LTO Disabled	IOPS (millions) LTO Enabled	Lvol to NVMe performance (%)
4KB 100% Random Read	1	6	0.97	1.20	23.50%
	2	12	1.62	1.54	-4.61%
	3	18	2.27	2.22	-2.07%
	4	24	2.88	3.23	12.35%
	5	30	3.37	3.29	-2.29%
	6	36	3.75	3.83	2.06%
	8	36	4.48	4.77	6.55%
	10	36	5.14	5.44	5.97%
4KB 100% Random Write	12	36	4.64	4.95	6.67%
	1	6	1.15	1.22	6.18%
	2	12	1.43	1.88	31.08%
	3	18	2.17	2.71	24.70%
	4	24	2.95	3.64	23.47%
	5	30	3.77	3.98	5.67%
	6	36	3.92	4.33	10.61%
	8	36	4.68	4.34	-7.38%
4KB 70% Random Read 30% Random Write	10	36	4.45	4.26	-4.35%
	12	36	4.34	4.35	0.31%
	1	6	1.25	1.16	-7.28%
	2	12	1.65	1.89	14.96%
	3	18	2.01	2.22	10.70%
	4	24	3.01	3.08	2.14%
	5	30	3.31	3.43	3.50%
	6	36	3.69	4.15	12.39%
4KB 70% Random Read 30% Random Write	8	36	4.46	4.40	-1.23%
	10	36	4.60	4.88	6.09%
	12	36	4.51	4.39	-2.72%



Conclusions

1. SPDK Vhost SCSI performance when using Split NVMe bdevs for backend is noticeably better than the same setup with Logical Volume bdevs. It scales near linearly up to 6 CPU cores and achieves peak performance at this point for all workloads. For 4KB Random Read workload increasing the number of cores does not result in performance improvement, whereas for Random Write and Random Read/Write workloads minor performance degradations are observed.
2. SPDK Vhost SCSI using Logical Volume backend devices performance scales near linearly up to 8 CPU cores. Increasing the number of cores does not result in performance improvement.
3. SPDK Vhost BLK using Logical Volume backend devices performance scales near linearly up to 8 CPU cores. Increasing the number of cores does not result in performance improvement.
4. Using Logical Volumes as part of testing setup has a large impact on the overall performance. Performance drop varies from 15-25% when Vhost process uses multiple CPU up to 50% (!) for a single CPU Vhost process. Interestingly – Logical Volume impact is negligible beyond the point where Vhost process has hit the bottleneck (beyond 8 CPUs).
5. Using LTO increases SPDK Vhost performance by about 5-10% percent and up to 30% in some of the best cases.

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 were chosen as the rate limiter using linux cgroups.

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.

Item	Description
Test case	Test rate limiting IOPS/VM to 10000 IOPS
Test configuration	<p>FIO Version: fio-3.3</p> <p>VM Configuration:</p> <ul style="list-style-type: none"> • Common settings are described in the Virtual Machine Settings chapter. • Total of 24 / 48 / 72 VMs • 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. <ul style="list-style-type: none"> ○ 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 and write) <p>SPDK Vhost target configuration:</p> <ul style="list-style-type: none"> • Test were run with both Vhost-scsi and Vhost-blk stacks. • The Vhost-scsi stack was run with Split fvNVMe bdevs and Logical Volume bdevs. • The Vhost-blk stack was run with Logical Volume bdevs. • Test were run with 4 CPU cores (NUMA optimized). <p>Kernel Vhost-scsi configuration:</p> <ul style="list-style-type: none"> • Cgroups were used to limit the Vhost process to 4 cores. • NUMA optimization were not explored.
FIO configuration run on each VM	<pre>[global] ioengine=libaio direct=1 rw=randrw</pre>



	<pre>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</pre>
--	-------------------------------------------------------------------------------------------------------------------------------------------------

Test Case 2 Results

Table 6: 4KB 100% Random Reads QD=1

# of VMs	Stack	Backend bdev	IOPS (k)	Avg Lat. (usec)
24 VMs	SPDK-SCSI	Split NVMe	239.41	98.51
	SPDK-SCSI	Logical Volume	239.04	98.65
	SPDK-BLK	Logical Volume	239.16	98.64
	Kernel-SCSI	Partitioned NVMe	150.36	161.67
48 VMs	SPDK-SCSI	Split NVMe	457.00	101.33
	SPDK-SCSI	Logical Volume	456.24	101.51
	SPDK-BLK	Logical Volume	467.55	100.44
	Kernel-SCSI	Partitioned NVMe	142.30	344.77
72 VMs	SPDK-SCSI	Split NVMe	663.70	106.08
	SPDK-SCSI	Logical Volume	652.68	107.41
	SPDK-BLK	Logical Volume	662.29	104.79
	Kernel-SCSI	Partitioned NVMe	221.09	349.33




Figure 5: 4KB 100% Random Reads IOPS and latency, QD=1, throttling = 10k IOPS



Table 7: 4KB 100% Random Writes QD=1

# of VMs	Stack	Backend bdev	IOPS (k)	Avg Lat. (usec)
24 VMs	SPDK-SCSI	Split NVMe	240.00	97.90
	SPDK-SCSI	Logical Volume	240.00	97.88
	SPDK-BLK	Logical Volume	240.00	97.86
	Kernel-SCSI	Partitioned NVMe	134.75	186.18
48 VMs	SPDK-SCSI	Split NVMe	473.33	97.62
	SPDK-SCSI	Logical Volume	479.99	97.67
	SPDK-BLK	Logical Volume	480.00	97.65
	Kernel-SCSI	Partitioned NVMe	147.00	349.65
72 VMs	SPDK-SCSI	Split NVMe	713.27	97.48
	SPDK-SCSI	Logical Volume	713.22	97.46
	SPDK-BLK	Logical Volume	709.91	97.50
	Kernel-SCSI	Partitioned NVMe	256.54	283.03





Figure 6: 4KB 100% Random Writes IOPS and latency, QD=1, throttling = 10k IOPS



Conclusions

1. VMs using SPDK Vhost exposed devices were able to achieve the expected IOPS result.
2. SPDK Vhost was able to serve IO at the desired level for an increasing number of VMs.
3. Average latencies were up to 3.7x times better for Random Read and Random Write workloads in SPDK Vhost when compared to Kernel Vhost.

Note: The Kernel-Vhost process was not NUMA-optimized for this scenario.



Test Case 3: Performance per NVMe drive

This test case was performed in order 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.

Item	Description
Test case	Test SPDK Vhost target I/O core scaling performance
Test configuration	<p>FIO Version: fio-3.3</p> <p>VM Configuration:</p> <ul style="list-style-type: none"> • 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. <p>SPDK Vhost target configuration:</p> <ul style="list-style-type: none"> • 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. <p>Kernel Vhost target configuration:</p> <ul style="list-style-type: none"> • The Vhost process was run on a single, physical CPU core using cgroups.
FIO configuration	<pre>[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</pre>



Test Case 3 results

SPDK Vhost-Scsi

Table 8: IOPS and latency results, SCSI stack

Access pattern	Backend	QD	Throughput (IOPS)	Avg. latency (usec)
4k 100% Random Reads	Split NVMe	1	23985.91	82.67
4k 100% Random Reads	Split NVMe	8	166351.27	95.12
4k 100% Random Reads	Split NVMe	32	391621.33	161.44
4k 100% Random Reads	Split NVMe	64	395548.91	320.17
4k 100% Random Reads	Lvol	1	24248.69	81.32
4k 100% Random Reads	Lvol	8	164909.25	95.61
4k 100% Random Reads	Lvol	32	390708.91	162.91
4k 100% Random Reads	Lvol	64	390217.1	325.45
4k 100% Random Writes	Split NVMe	1	101605.51	18.56
4k 100% Random Writes	Split NVMe	8	361553.21	41.89
4k 100% Random Writes	Split NVMe	32	371895.55	175.39
4k 100% Random Writes	Split NVMe	64	383264.53	515.07
4k 100% Random Writes	Lvol	1	101823.96	18.81
4k 100% Random Writes	Lvol	8	385071.51	39.48
4k 100% Random Writes	Lvol	32	383425.89	165.53
4k 100% Random Writes	Lvol	64	388458.71	327.56
4k 70%/30% Random Read Writes	Split NVMe	1	30930.88	63.531
4k 70%/30% Random Read Writes	Split NVMe	8	153491.57	102.11
4k 70%/30% Random Read Writes	Split NVMe	32	354861.51	178.878
4k 70%/30% Random Read Writes	Split NVMe	64	350505.13	362.404
4k 70%/30% Random Read Writes	Lvol	1	29826.77	65.902
4k 70%/30% Random Read Writes	Lvol	8	175921.71	89.538
4k 70%/30% Random Read Writes	Lvol	32	291590.97	218.792
4k 70%/30% Random Read Writes	Lvol	64	355379.54	359.351



SPDK Vhost-Blk

Table 9: IOPS and latency results, BLK stack

Access pattern	Backend	QD	Throughput (IOPS)	Avg. latency (usec)
4k 100% Random Reads	Lvol	1	24252.64	80.87
4k 100% Random Reads	Lvol	8	168075.39	94.28
4k 100% Random Reads	Lvol	32	417881.84	151.94
4k 100% Random Reads	Lvol	64	436141.57	292.16
4k 100% Random Writes	Lvol	1	99252.55	18.42
4k 100% Random Writes	Lvol	8	409656.43	37.49
4k 100% Random Writes	Lvol	32	410330.56	189.05
4k 100% Random Writes	Lvol	64	417830.56	301.66
4k 70%/30% Random Read Writes	Lvol	1	30779.79	63.702
4k 70%/30% Random Read Writes	Lvol	8	171064.08	91.875
4k 70%/30% Random Read Writes	Lvol	32	299984.77	213.531
4k 70%/30% Random Read Writes	Lvol	64	362972.27	353.793

Kernel Vhost-Scsi

Table 10: IOPS and latency results, Kernel Vhost-Scsi

Access pattern	Backend	QD	Throughput (IOPS)	Avg. latency (usec)
4k 100% Random Reads	NVMe	1	12917.92	164.86
4k 100% Random Reads	NVMe	8	70868.25	301.6
4k 100% Random Reads	NVMe	32	219263.46	291.18
4k 100% Random Reads	NVMe	64	214580.54	594.4
4k 100% Random Writes	NVMe	1	58274.64	34.86
4k 100% Random Writes	NVMe	8	110475.42	142.75
4k 100% Random Writes	NVMe	32	127096.91	392.31
4k 100% Random Writes	NVMe	64	163445.64	708.5
4k 70%/30% Random Read Writes	NVMe	1	19711.14	99.857
4k 70%/30% Random Read Writes	NVMe	8	70368.22	315.175
4k 70%/30% Random Read Writes	NVMe	32	154684.98	580.636
4k 70%/30% Random Read Writes	NVMe	64	198266.89	578.135




Figure 7: 4KB 100% Random Reads IOPS and latency




Figure 8: 4KB 100% Random Writes IOPS and latency





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 all 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.



DISCLAIMERS

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

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

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information go to <http://www.intel.com/performance>

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

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

Your costs and results may vary.

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

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