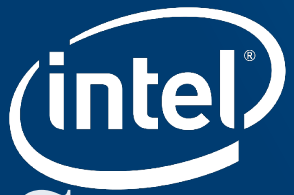

Storage performance development kit: USING DPDK TO Accelerate Storage services

Jim Harris

Principal Software Engineer

Intel Data Center Group

Legal Disclaimer

The Challenge: Media Latency

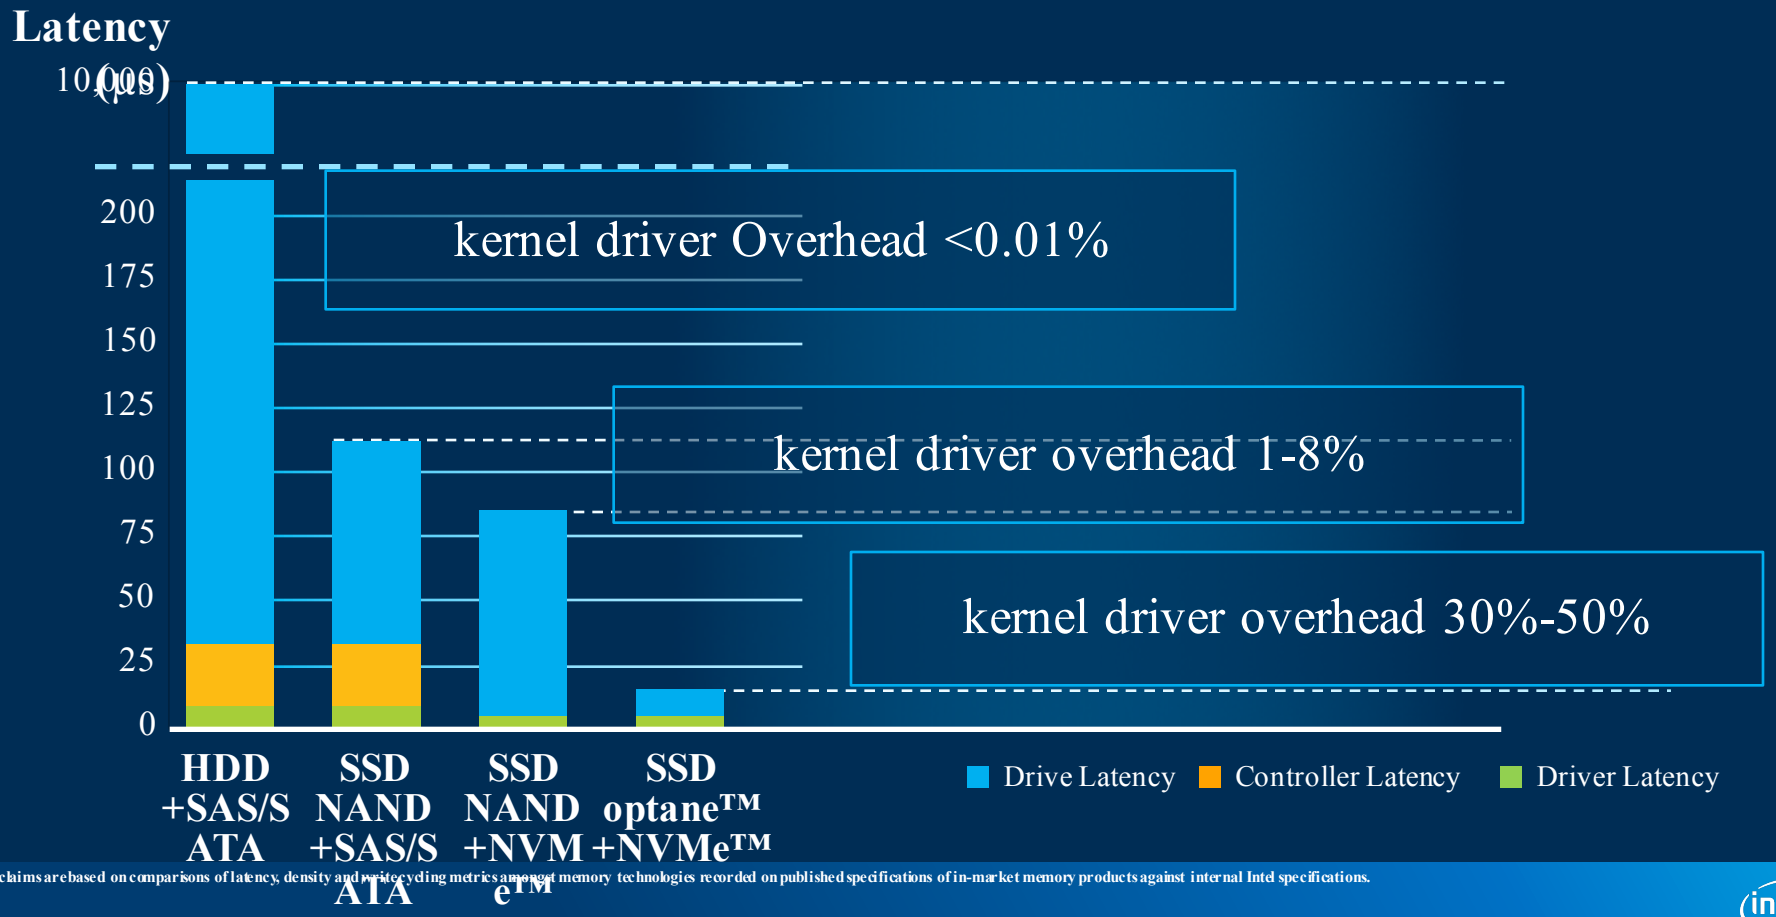

Technology claims are based on comparisons of latency, density and write cycling metrics amongst memory technologies recorded on published specifications of in-market memory products against internal Intel specifications.

Storage Performance Development Kit (SPDK)

What is SPDK?

- Userspace polled-mode drivers, libraries and applications for storage, storage networking and storage virtualization
- Leverages DPDK
- Started in 2013, open sourced in 2015
- BSD licensed
- http://SPDK.io

NVM Express* Driver Throughput Scalability

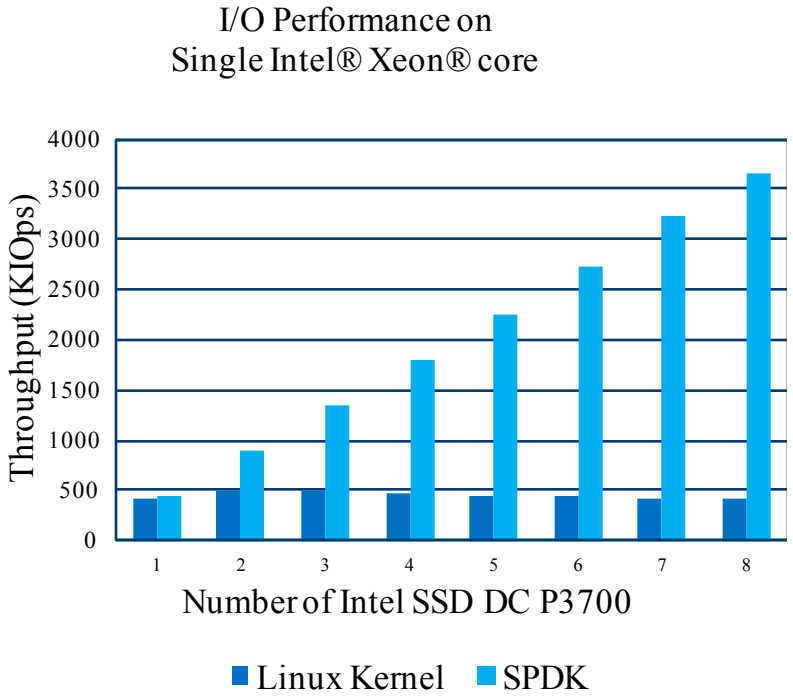

- Systems with multiple NVM Express* (NVMe) SSDs capable of millions of I/O per second
- Results in many cores of software overhead with kernel-based interrupt-driven driver model
- SPDK enables:
  - more CPU cycles for storage services
  - lower I/O latency

SPDK saturates 8 NVMe SSDs with a single CPU core!

System Configuration: 2x Intel® Xeon® E5-2695v4 (HT off), Intel® Speed Step enabled, Intel® Turbo Boost Technology disabled, 8x 8GB DDR4 2133 MT/s, 1 DIMM per channel, CentOS* Linux* 7.2, Linux kernel 4.10.0, 8x Intel® P3700 NVMe SSD (800GB), 4x per CPU socket, FW 8DV101H0, I/O workload 4KB random read, Queue Depth: 128 per SSD, Performance measured by Intel using SPDK perf tool, Linux kernel data using Linux AIO

VM Storage Acceleration

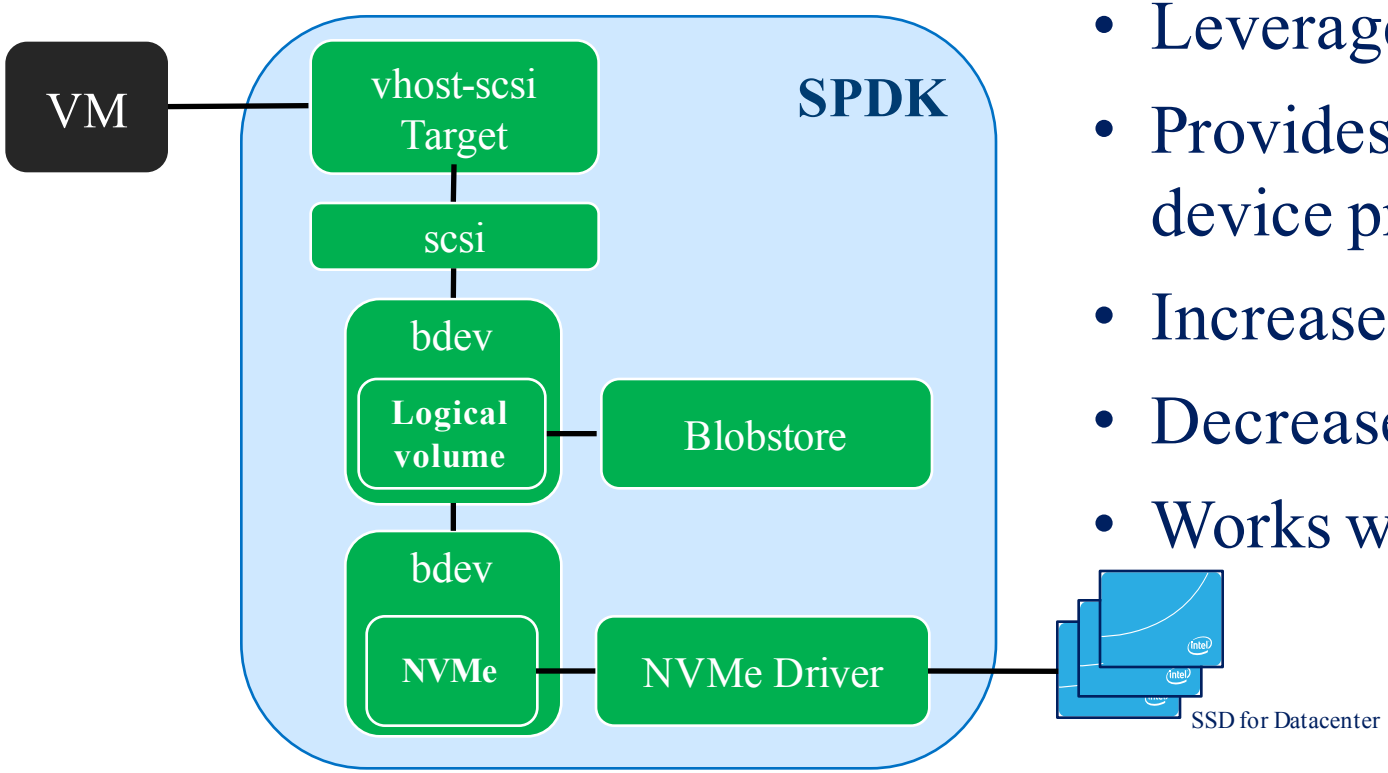

- Leverages DPDK vhost
- Provides dynamic block device provisioning
- Increase VM Density
- Decrease Guest Latency
- Works with KVM/QEMU

SPDK vhost Performance

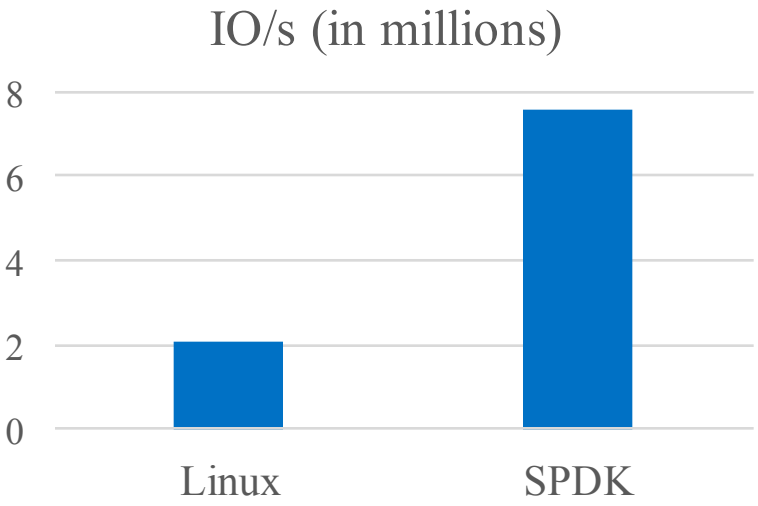

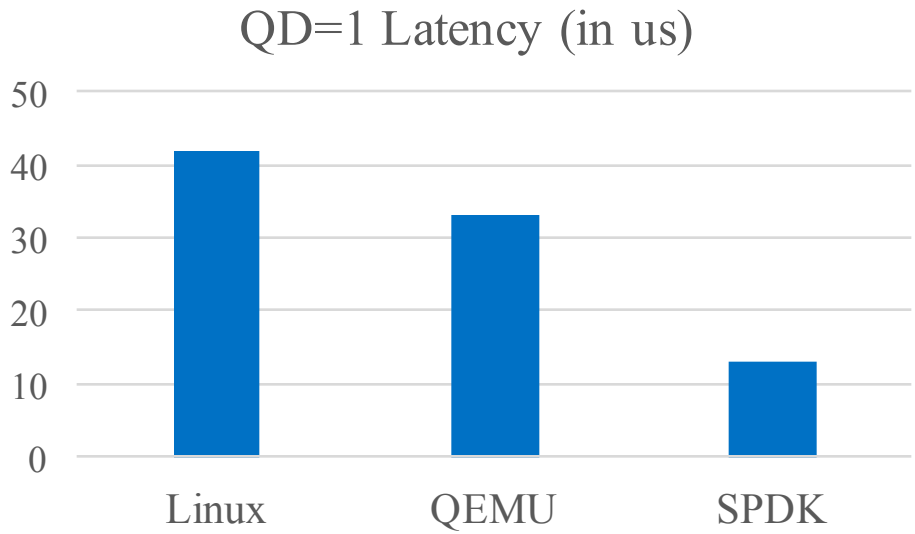

SPDK up to 3x better efficiency and latency

System Configuration: 2S Intel® Xeon® Platinum 8180: 28C, E5-2699v3: 18C, 2.5GHz (HT off), Intel® Turbo Boost Technology enabled, 12x16GB DDR4 2133 MT/s, 1 DIMM per channel, Ubuntu® Server 16.04.2 LTS, 4.11 kernel, 23x Intel® P4800x Optane SSD – 375GB, 1 SPDK lvolstore or LVM lvgroup per SSD, SPDK commit ID c5d8b108f22ab, 46 VMs (CentOS 3.10, 1vCPU, 2GB DRAM, 100GB logical volume), vhost dedicated to 10 cores

As measured by: fio 2.10.1 – Direct=Yes, 4KB random read I/O, Ramp Time=30s, Run Time=180s, Norandommap=1, I/O Engine = libaio, Numjobs=1

Legend: Linux: Kernel vhost-scsi QEMU: virtio-blk dataplane SPDK: Userspace vhost-scsi

Storage Performance Development Kit (SPDK)

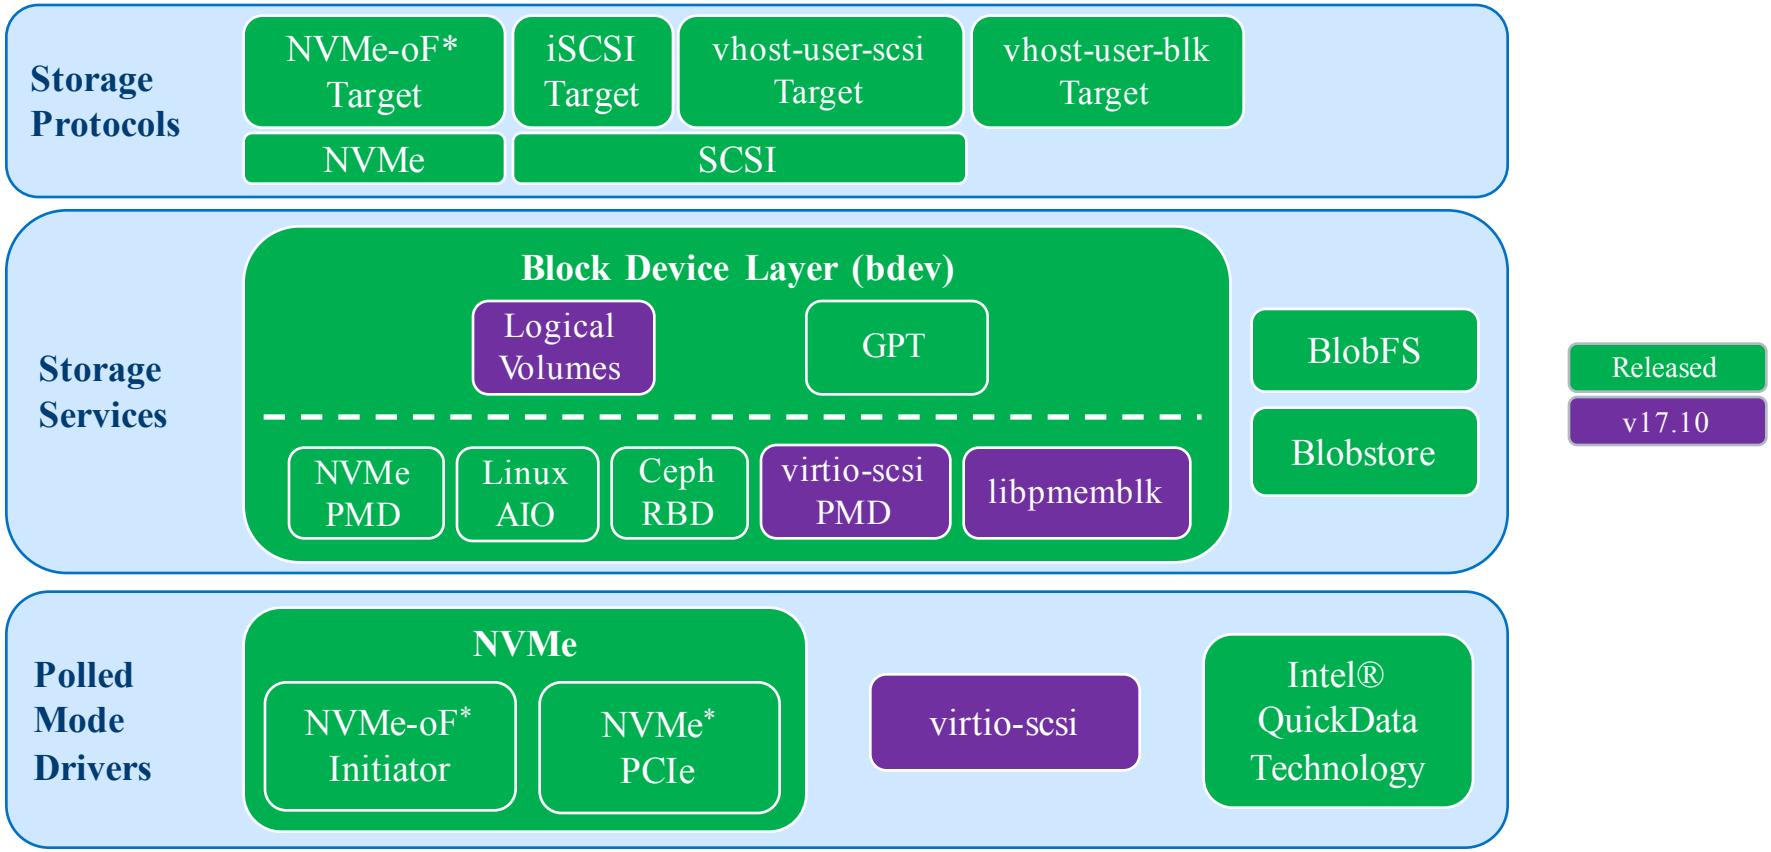

DPDK Key Features for SPDK

- Threads
- PCIe Device Management
- Memory Management
- Rings, Mempools
- Multi-Process
- vhost

Storage v. Packet Processing

- PCIe Device Hotplug
- Runtime v2phys Translation
- vhost VM Boot
- Storage is Endpoint Focused

SPDK Community

Home Page : http://www.SPDK.io/

Github : https://github.com/spdk/spdk

Trello : https://trello.com/spdk

GerritHub : https://review.gerrithub.io/#/q/project:spdk/spdk+status:open

IRC : https://freenode.net/ we're on #spdk

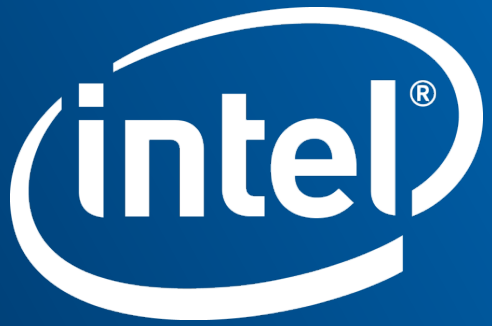
intel®