

SPDK: STATE OF THE PROJECT



Jim Harris

*Principal Engineer
Intel*

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

AGENDA

- 01** **State of the Community**
Community Participation, Major Milestones

- 02** **State of the Technology**
Key Features added over the Last Year

- 03** **State of the Future**
New Focus Areas and Use Cases for SPDK

AGENDA

01

State of the Community

Community Participation, Major Milestones

02

State of the Technology

Key Features added over the Last Year

03

State of the Future

New Focus Areas and Use Cases for SPDK

COMMUNITY TIMELINE

2013: SPDK starts as Intel internal project

Sep 2015: NVMe driver open sourced on GitHub

Jun 2016: NVMe-oF (RDMA) target released

Sep 2016: First SPDK summit in PRC

Dec 2017: First core maintainer outside of Intel

Nov 2018: NVMe-oF (TCP) target released along with specification

RECENT COMMUNITY MILESTONES

Two new core maintainers

- Alexey Marchuk (Mellanox/NVIDIA)
- Tomek Zawadzki (Intel)

Developer Meetup hosted by Nutanix (November 2019)

Virtual US Forum (June 2020)

- 630+ attendees over three days

Mellanox and Broadcom CI brought online

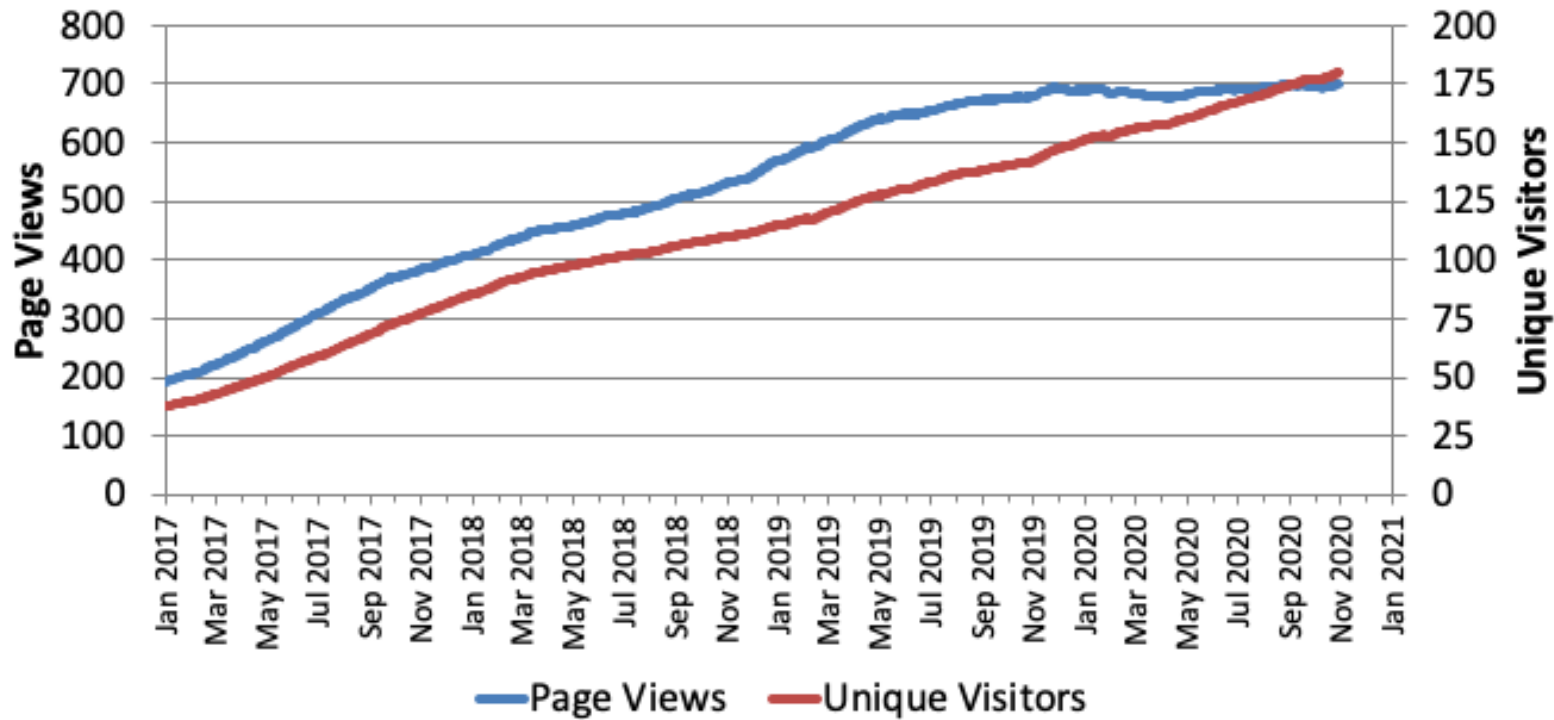
First LTS release (v20.01)

SPDK PATCH COMMIT STATISTICS

Version Range	Total Patches	Total Committers	Patches from Top 10 Committers
v16.08 .. v17.10	2592	59	2246 (87%)
v17.10 .. v18.10	3074	89	2227 (72%)
v18.10 .. v19.10	3709	97	2516 (68%)
v19.10 .. v20.10	3582	76	2350 (66%)

GITHUB VISITORS

SPDK Github Traffic 52 Week Running Daily Averages



Unique Visitors +26% YoY
Page Views +3% YoY

AGENDA

01 State of the Community
Community Participation, Major Milestones

02 State of the Technology
Key Features added over the Last Year

03 State of the Future
New Focus Areas and Use Cases for SPDK

EVEN BETTER PERFORMANCE!

NVMe Driver (PCIe)

- >14M IO/s (1 Intel Xeon CPU core)

NVMe Fabrics Target (TCP)

- 4KiB Random Read: 3.5M IO/s (8 Intel Xeon CPU cores)
 - 87% improvement over last 9 months
- 128KiB Sequential Read: >14GiB/s (1 Intel Xeon CPU core)

Logical Volume performance

- New metadata layout for thin provisioning
- I/O path optimizations (up to 10% better perf/core)

NVMe driver configuration: Intel Server S2600WFT, Intel(R) Xeon(R) Platinum 8280L CPU @ 2.70GHz, 192GB DDR4, 6x Memory Channels per socket, 1 16GB 2667 DIMM per channel, Fedora 27, Linux Kernel 5.0.0-rc6. BIOS: Intel® Speed Step enabled, Intel® Turbo Boost Technology enabled, Hyper Threading enabled, p-states enabled. 4x Intel® DC P5800X 503GB SSD. SPDK commit 41b7f1ca2189. I/O workload: 512B random read, QD=128 per SSD. Data collected at Intel Storage Lab 09/12/2019 using SPDK perf tool

NVMe/TCP Target Configuration: SuperMicro SYS-2029U-TN24R4T, 2S Intel Xeon Gold 6230 Processor, 12x32GB DDR4 2666MHz DRAM, Fedora 30, Linux kernel 5.4.14-100.fc30, SPDK 20.04, 16x Intel SSD DC P4610 SSD, 2x100GbE Mellanox ConnectX-5 NIC

INTEL® DSA

idxd driver for Intel® Data Streaming Accelerator (DSA)

- Available in future Intel processors
- Replaces Intel® QuickData Technology

New “accel” framework

- APIs for copy, fill, dualcast, crc32c operations
- Supports operation batching
- Offload to Intel® DSA if available, otherwise use CPU

Integration into NVMe-oF and iSCSI libraries later in future releases

NVMe-oF PLUGGABLE TRANSPORTS

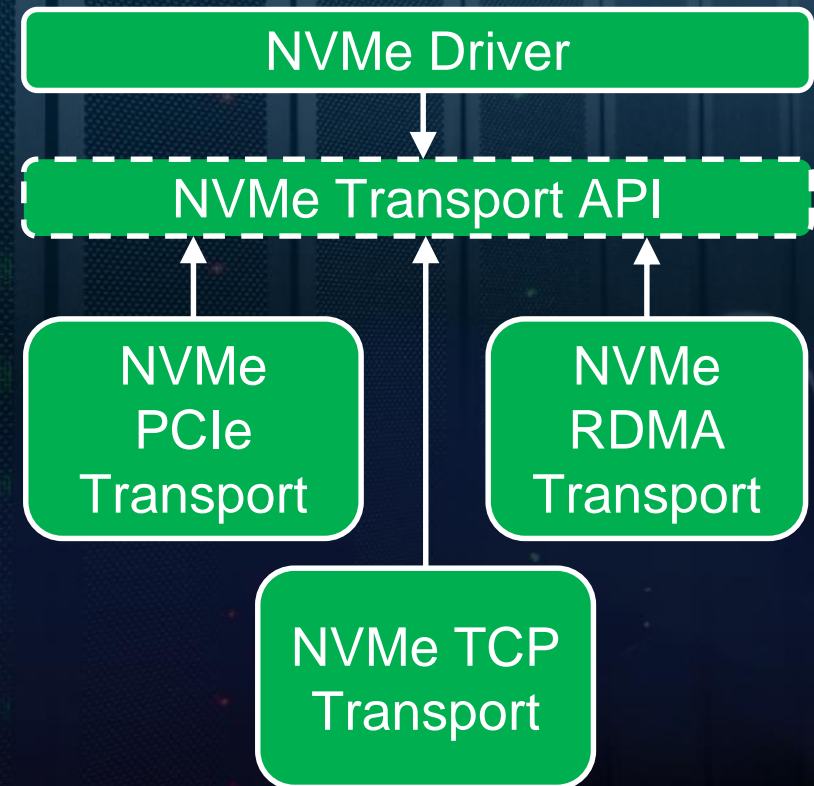
NVMe in SPDK supports multiple transports

- Host driver and target libraries

FC transport added for NVMe target in 19.07 release

Out-of-tree custom transports also supported

More “fabrics” transports on the way!



NVMe CHARACTER DEVICES

SPDK provides a framework and utilities for SSD management

- nvme/identify, JSON-RPC interfaces

SPDK v19.10 added NVMe character device support

- Uses Linux CUSE to create kernel char devices for SPDK-owned SSDs!

Enables use of standard Linux utilities with SPDK

- nvme-cli, smartctl

INTEL VMD

Available on root ports of Intel Xeon Scalable processors

Enables better SSD hotswap reliability and LED management

Native userspace SPDK VMD driver added in v19.07 release

POWER EFFICIENCY

spdk_thread scheduler framework

- Use busy/idle heuristics to move spdk_threads between CPU cores
- Enable putting cores to sleep when not needed

Intel® Speed Select Technology

- Enable CPU frequency scaling based on current workload

Interrupt/blocking mode support

- Lower performance requirements
- When polled mode cores not feasible

AGENDA

01 State of the Community
Community Participation, Major Milestones

02 State of the Technology
Key Features added over the Last Year

03 **State of the Future**
New Focus Areas and Use Cases for SPDK

NVME-OF – WITHOUT THE FABRICS!

SPDK NVMe-oF target has gradually evolved over time

- Common code for general NVMe target functionality
- Pluggable transports for PCIe, RDMA, TCP, etc.

Now well suited for additional use cases!

- Virtualization (vfio-user)
- SmartNIC

COMPUTATIONAL STORAGE

Classic use case of moving compute closer to the storage

- Fixed function (compression, hashing, erasure coding)
- Programmable functions

Expect SPDK to be at the center of standardization efforts

MORE PERFORMANCE, EFFICIENCY AND SCALABILITY

NVMe-oF

- Scaling to 1000s of connections and namespaces

Dynamic thread scheduler

- Based on heuristics of current workload

DIF/DIX

Header and Data Digests (NVMe, iSCSI)

CALL TO ACTION

Attend the SPDK (and PMDK and Intel VTune) talks today!

Keep up to date on what's happening with SPDK

- Blog
- Performance Reports

Jump in!

- Slack
- Mailing List
- Community Meetings

The Intel logo is centered in the upper half of the image. It features the word "intel" in a lowercase, white, sans-serif font. A small blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®). The background is a dark blue, slightly blurred image of server racks with glowing lights.

intel®

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum