STORAGE PERFORMANCE DEVELOPMENT KIT STATE OF THE PROJECT

Jim Harris

Principal Software Engineer

Data Center Group

NOTICES AND DISCLAIMERS

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

Some results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance..

Intel processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Optimization Notice: 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. Notice Revision #20110804.

The benchmark results may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

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 visit www.intel.com/benchmarks.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com.

The cost reduction scenarios described are intended to enable you to get a better understanding of how the purchase of a given Intel based product, combined with a number of situation-specific variables, might affect future costs and savings. Circumstances will vary and there may be unaccounted-for costs related to the use and deployment of a given product. Nothing in this document should be interpreted as either a promise of or contract for a given level of costs or cost reduction.

No computer system can be absolutely secure.

*Other names and brands may be claimed as the property of others.

AGENDA

- Current State of the Project
- Looking Forward
- Call to Action

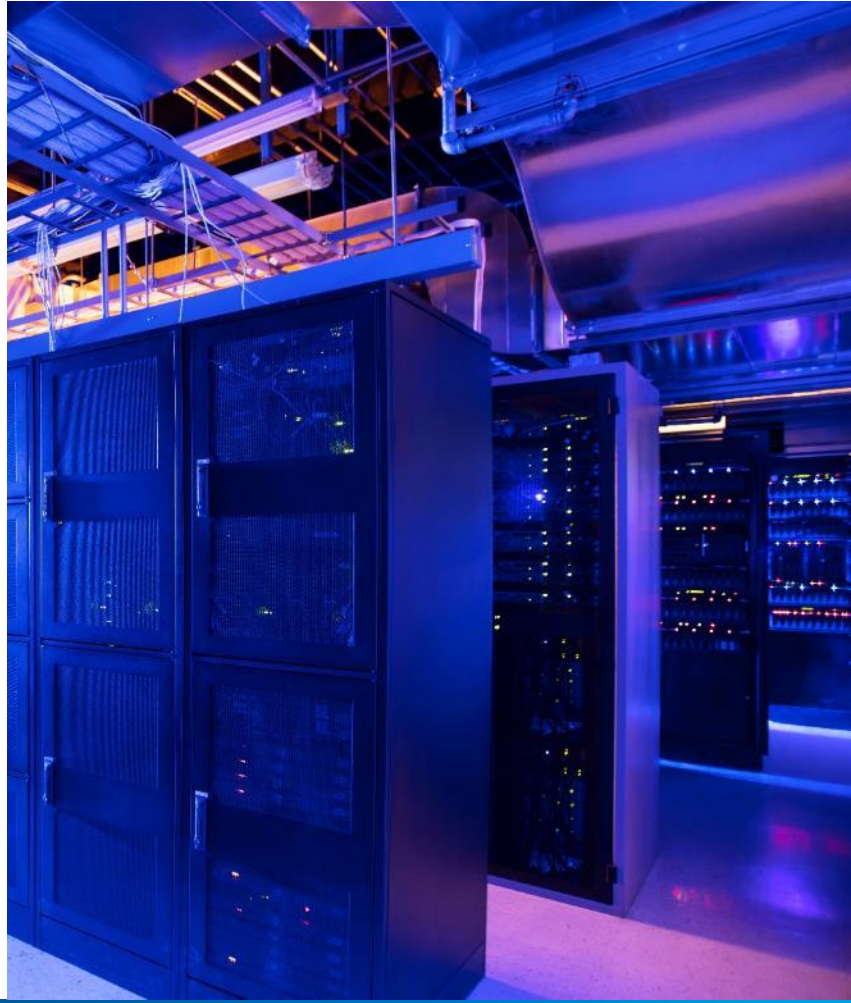

- Current State of the Project
- Looking Forward
- Call to Action

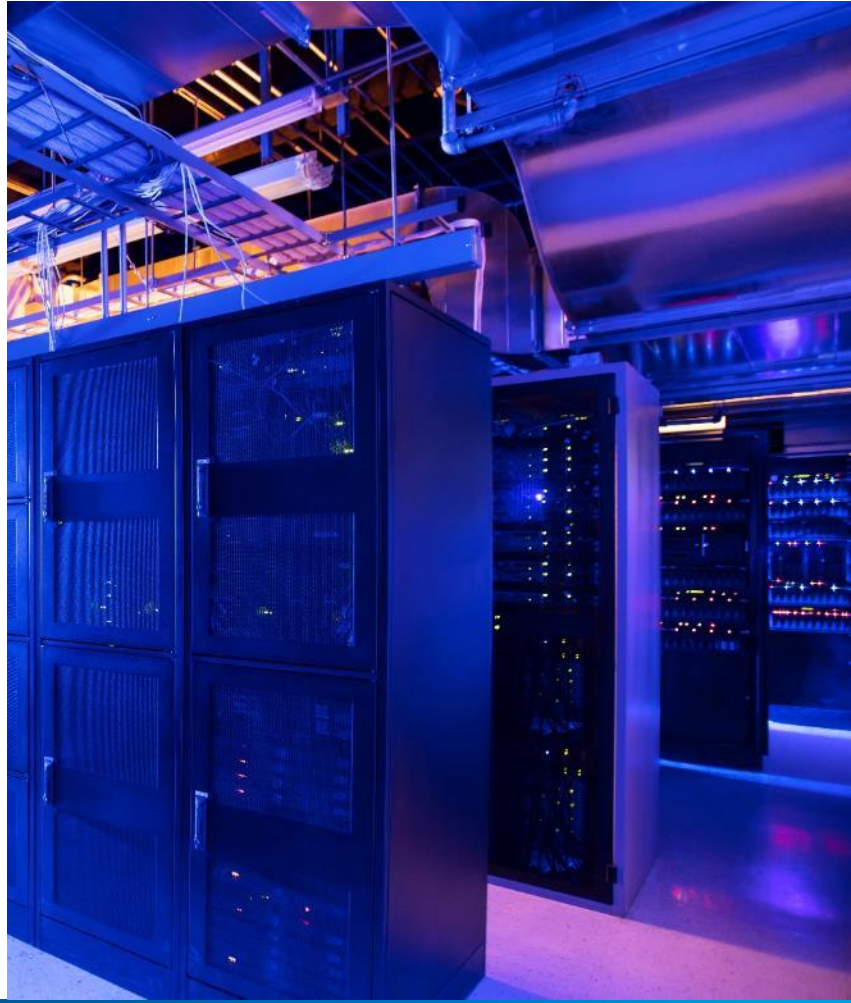

ONE YEAR AGO…

SPDK delivering exciting features and performance, but…

Was difficult for people outside Intel to contribute

No transparency on work going on inside Intel

Internal automated test framework

SPDK needed to truly become an open source community – not just an Intel open source code repository!

SINCE THEN...

Email Discussions

http://SPDK.IO

Main Web Presence

IRC

Real Time Chat w/
Development Community

Weekly Calls
Multiple Annual Meetups

Trello

Backlog and
Ideas for Things to Do

GerritHub™

GitHub

Code Reviews & Repo

Jenkins

Continuous
Integration

COMMUNITY GROWTH

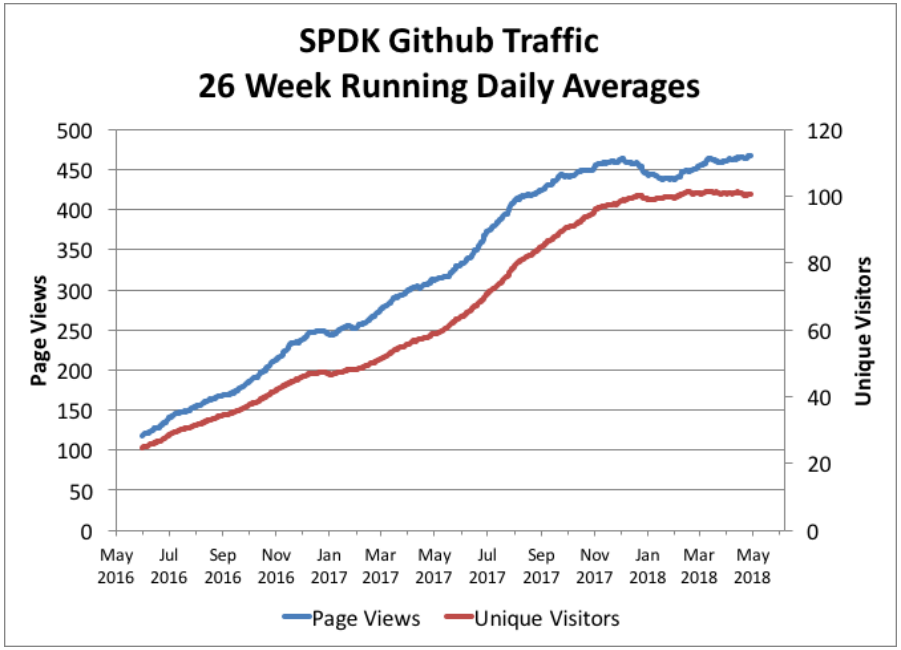

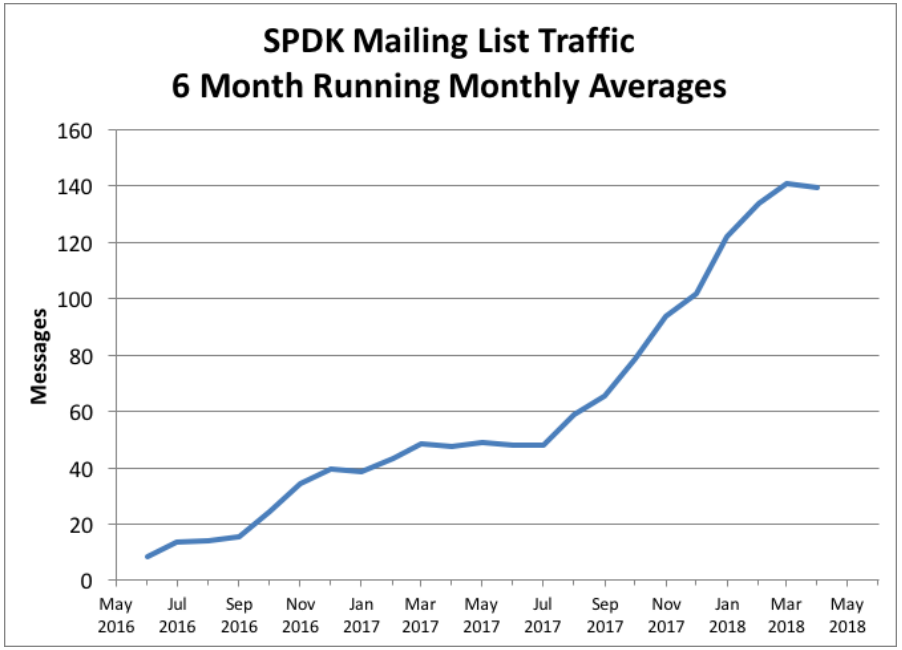

GitHub Traffic +60% Mailing List Traffic +200%

Intel internal data based on public available GitHub and mailing list data.

SPDK PATCH COMMIT STATISTICS

SPDK Version	Total Commits	Non-Intel Committers	Non-Intel Commits	Non-Intel Commit %
18.04	867	11	79	9.1%
18.01	775	14	79	10.2%
17.10	753	8	17	2.2%
17.07	685	7	11	1.6%
17.03	483	8	18	3.7%

SPDK PATCH COMMIT STATISTICS

SPDK Version	Total Commits	Non-Intel Committers	Non-Intel Commits	Non-Intel Commit %
18.04	**867**	**11**	**79**	**9.1%**
18.01	775	14	79	10.2%
17.10	753	8	17	2.2%
17.07	685	7	11	1.6%
17.03	**483**	**8**	**18**	**3.7%**

SPDK ARCHITECTURE

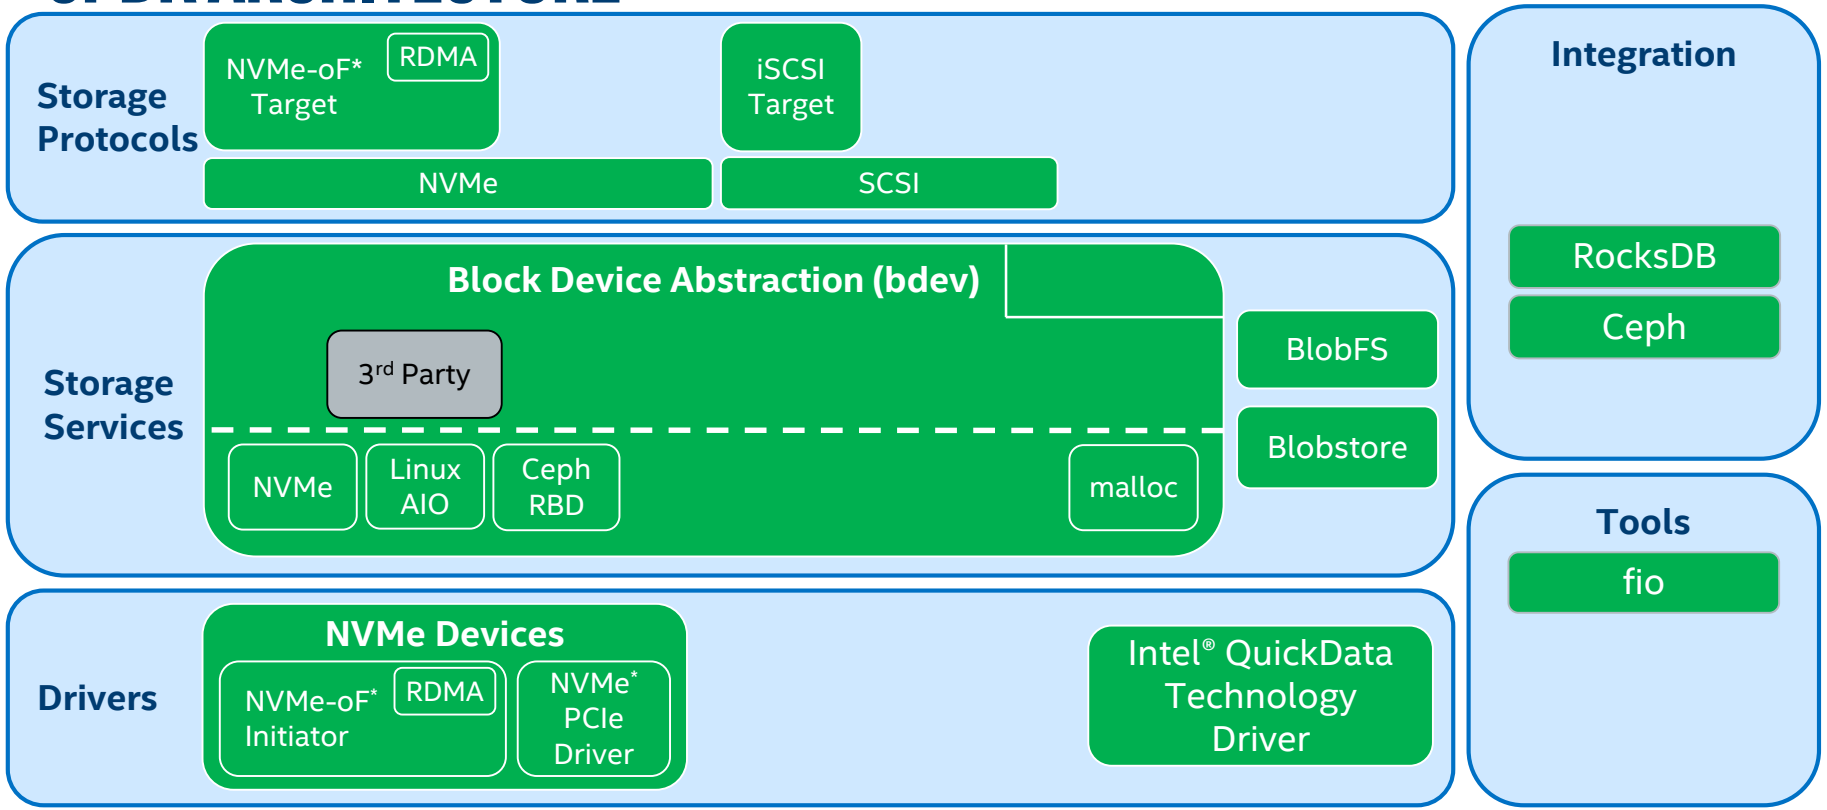

SPDK 17.03

Added since 17.03

SPDK ARCHITECTURE

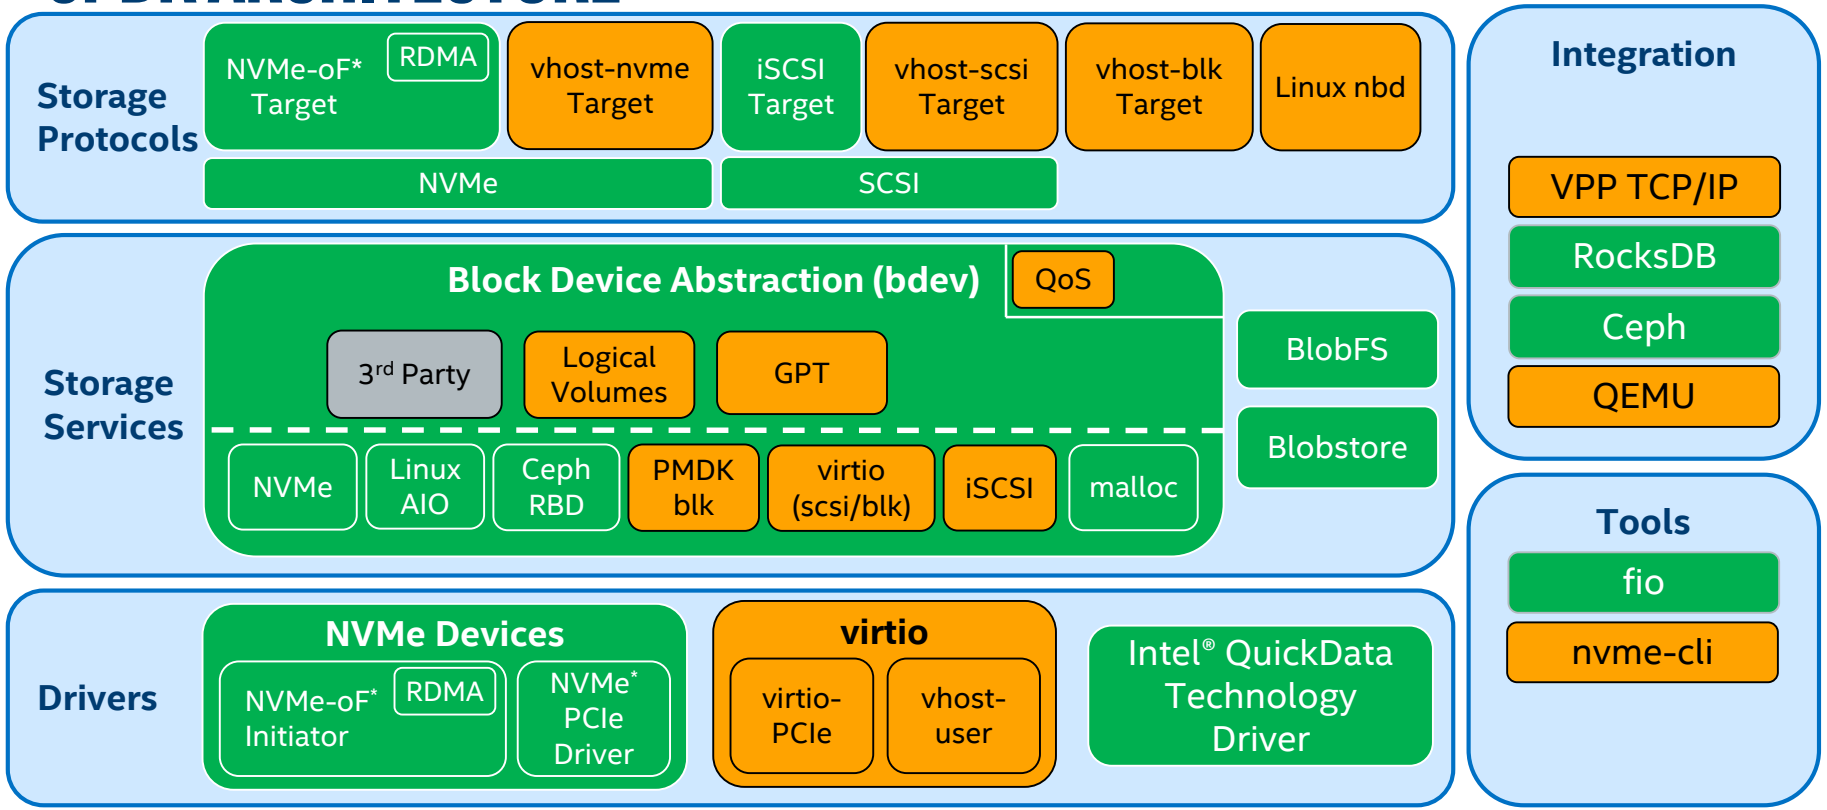

VHOST

Provide SPDK-based storage to QEMU-based VMs

- virtio-scsi
- virtio-blk
- nvme (experimental)

and non-QEMU host processes

- containers

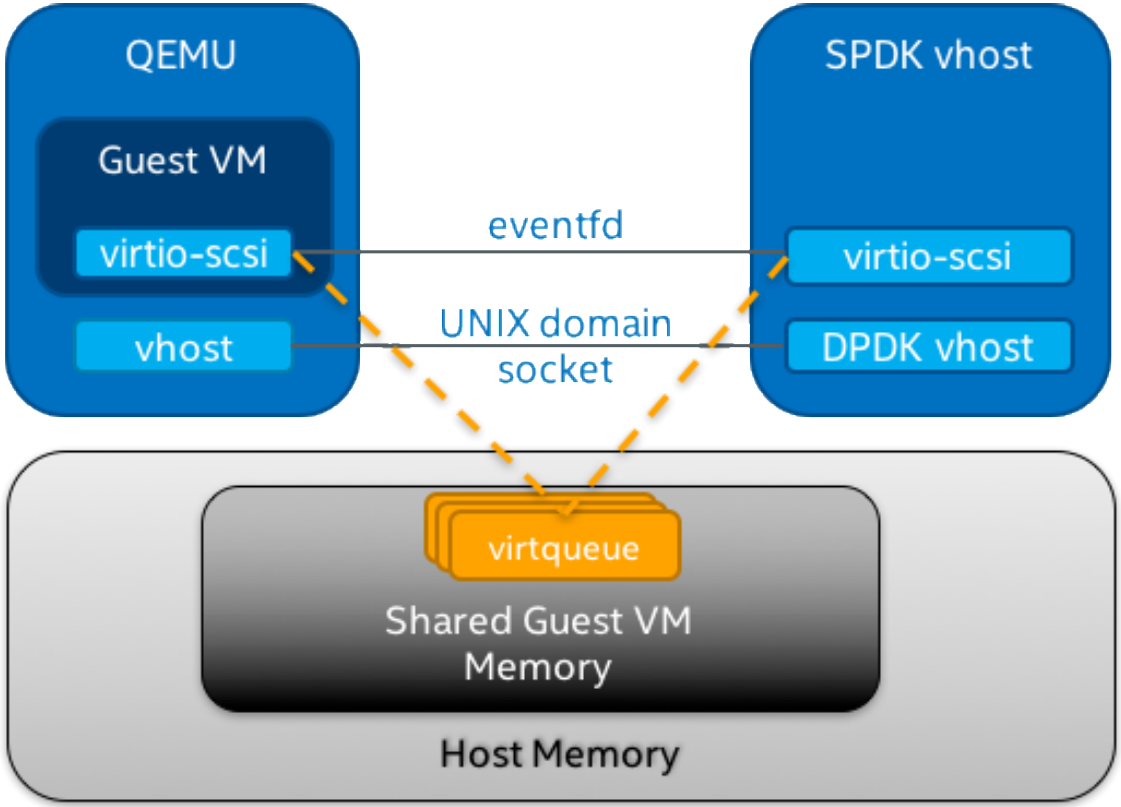

LOGICAL VOLUMES

Dynamic partitioning

Thin provisioning

Clones and snapshots

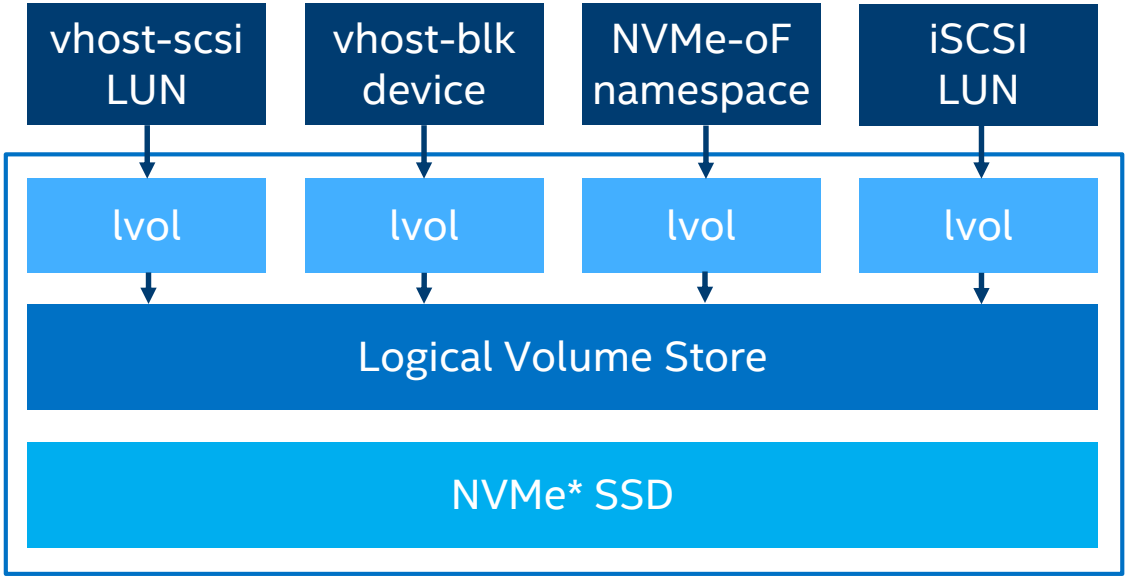

USERSPACE TCP/IP

FD.io – The Fast Data Project

- VPP – Vector Packet Processing

Utilize VCL (VPP Communications Library) for accelerated TCP processing

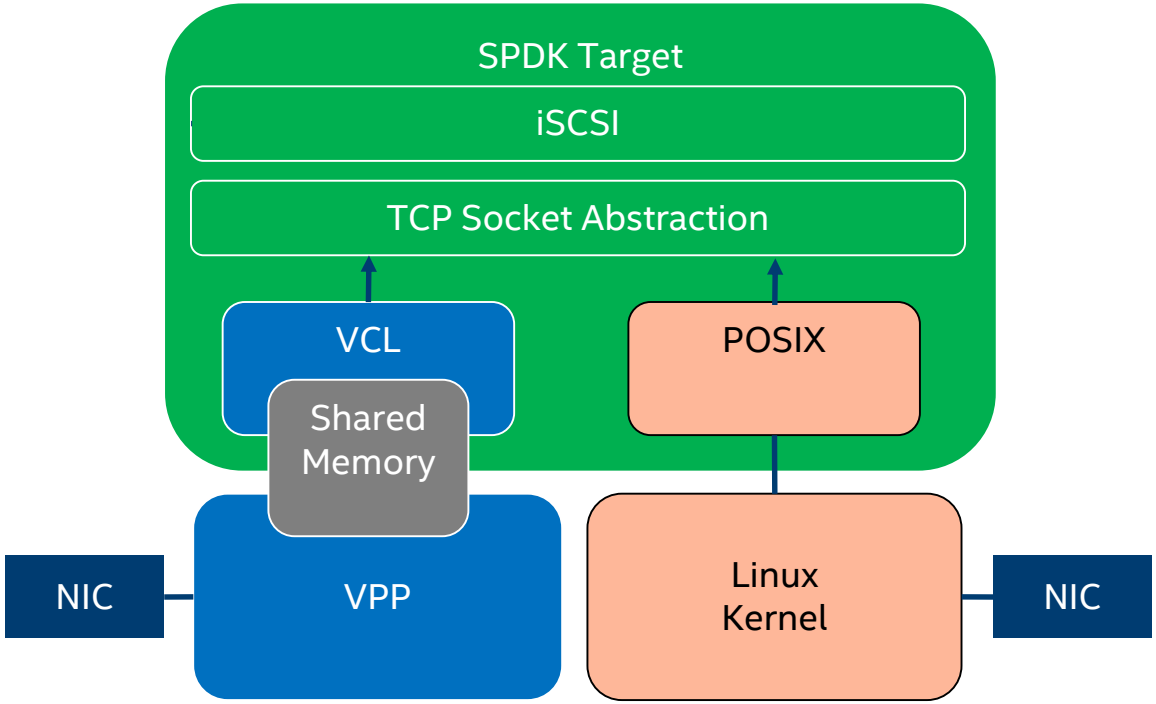

- Current State of the Project
- Looking Forward
- Call to Action

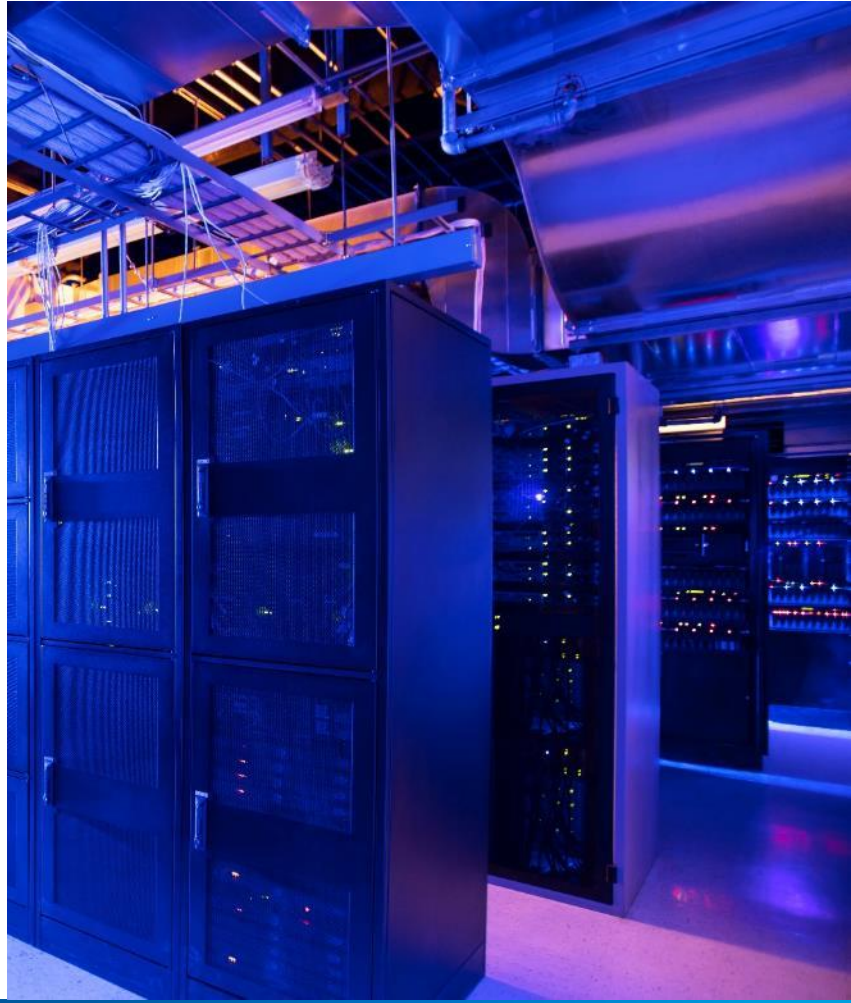

SPDK ARCHITECTURE

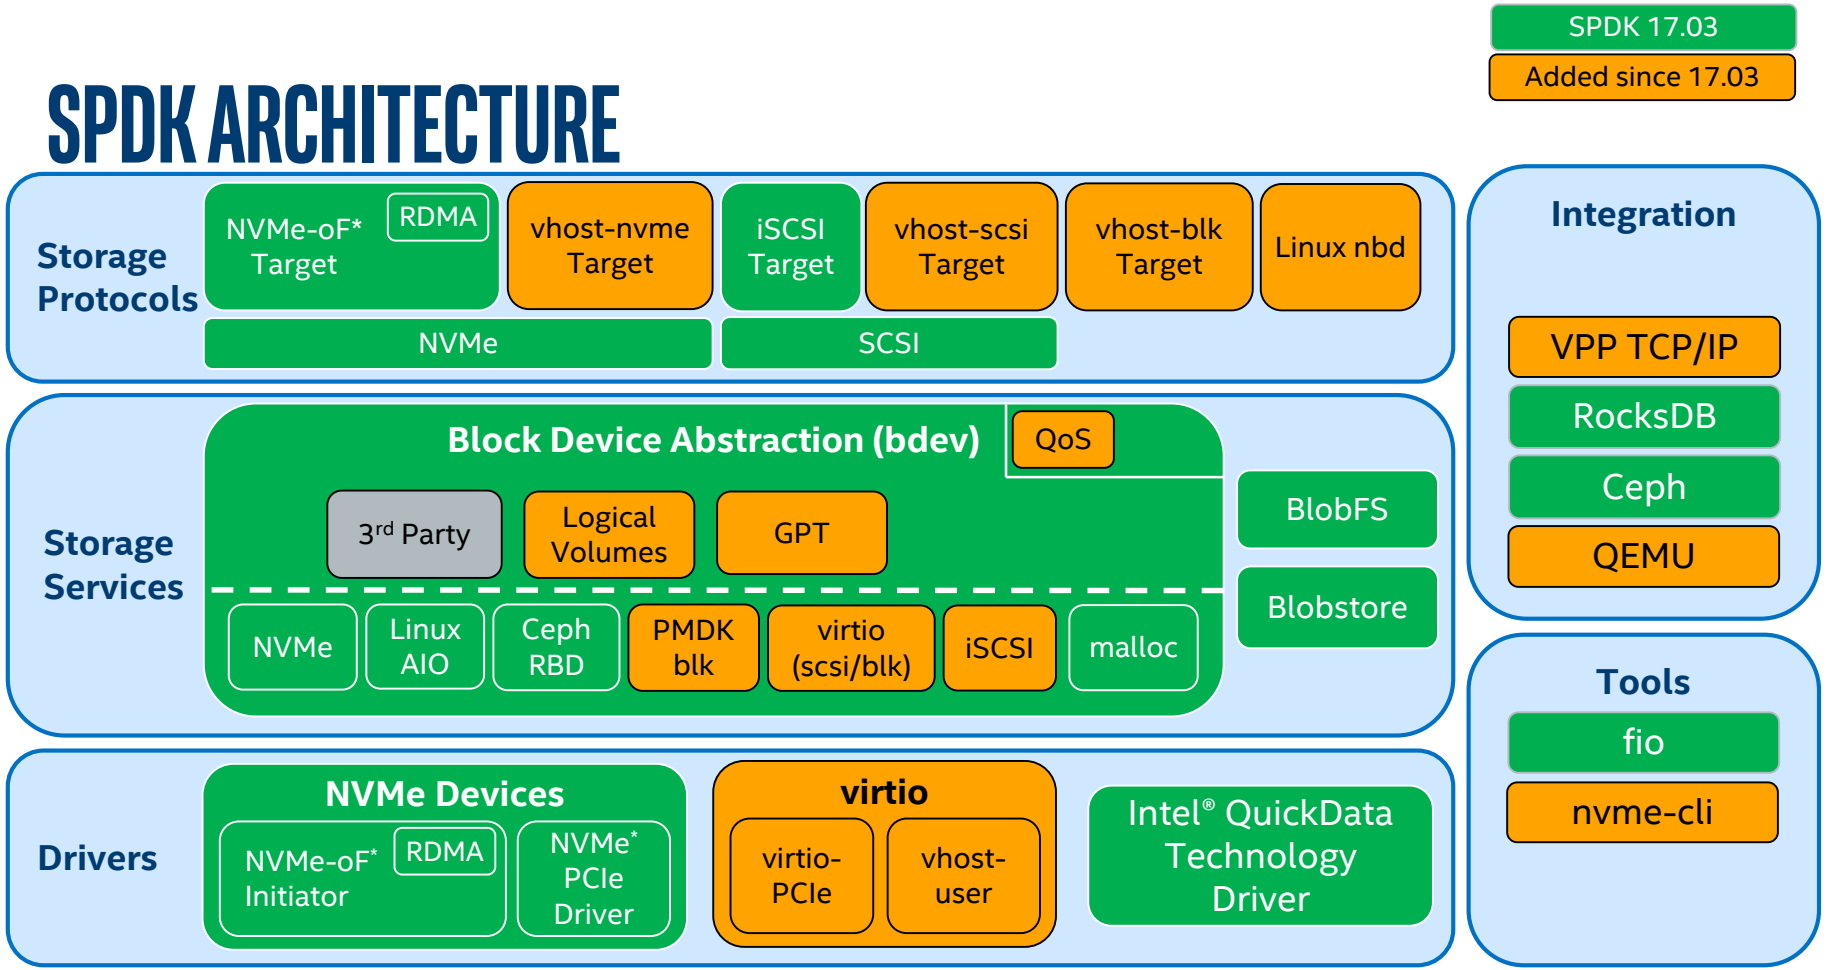

SPDK ARCHITECTURE

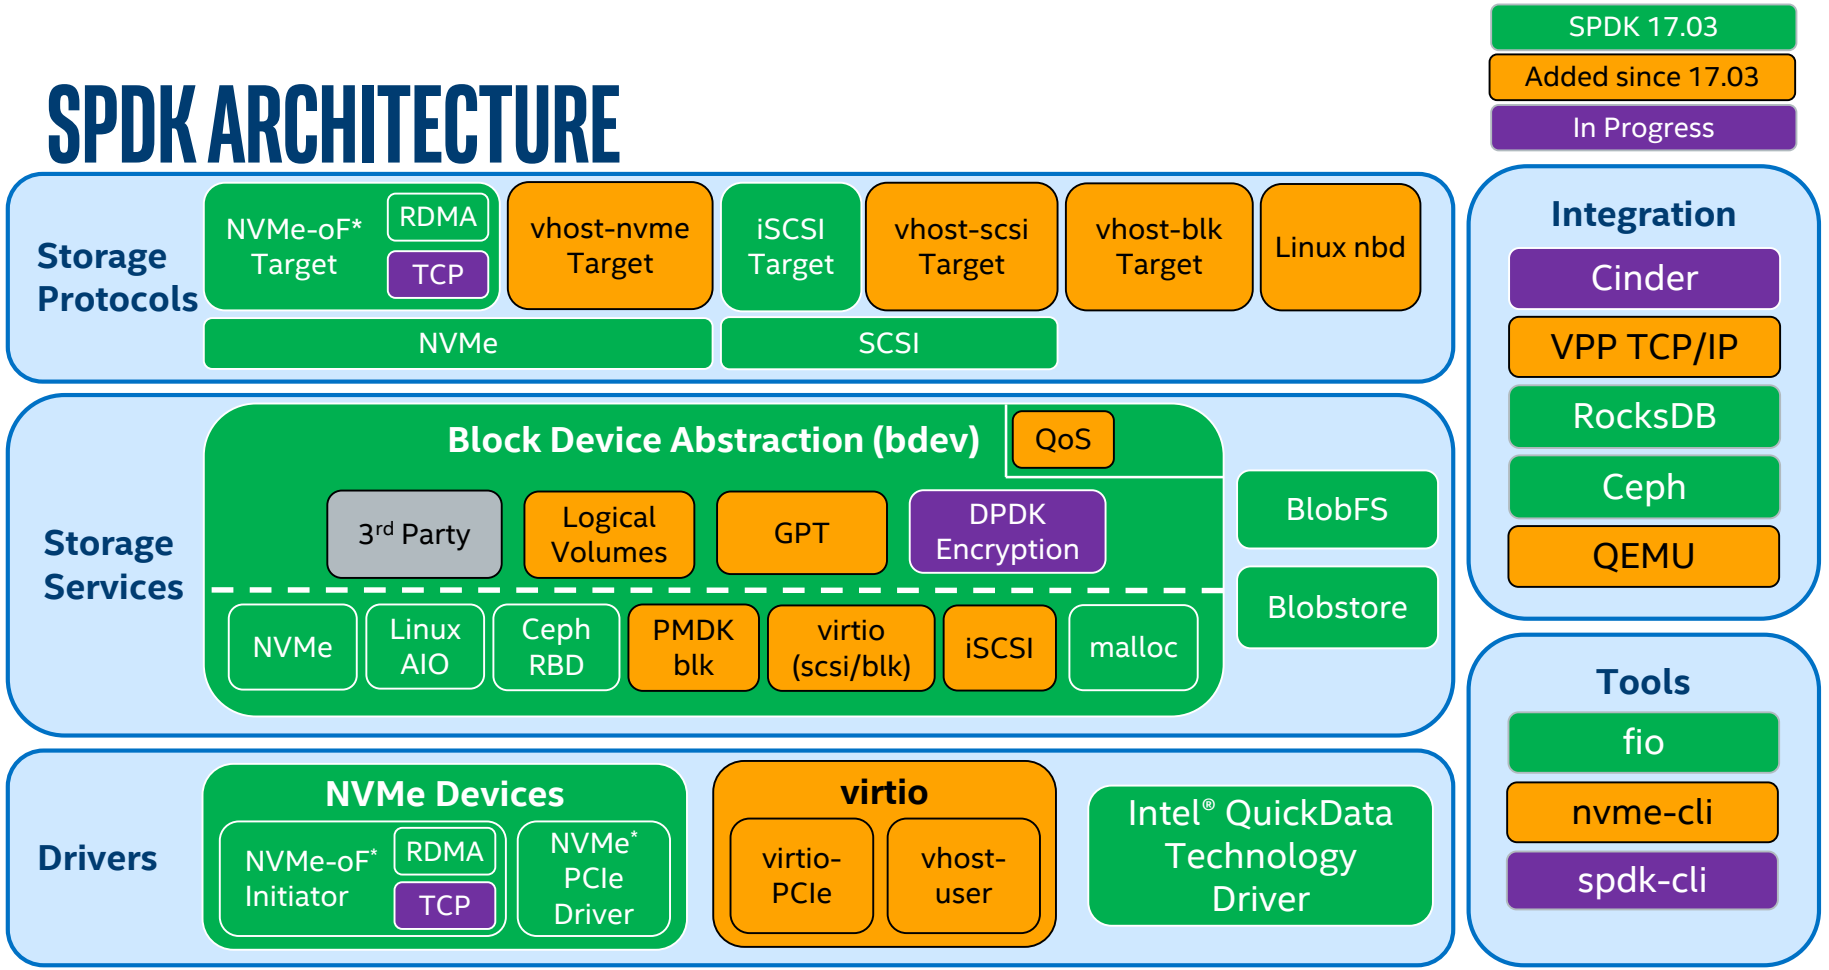

ORCHESTRATION

Cinder plugin for SPDK

- Work in progress

Provision ephemeral storage using SPDK:

- NVMe-oF target
- Logical volumes
- Dynamic configuration via JSON-RPC

Enables SPDK ephemeral storage provisioning for OpenStack!

ACCELERATORS

DPDK Framework

- Hardware Accelerators
- Optimized Software

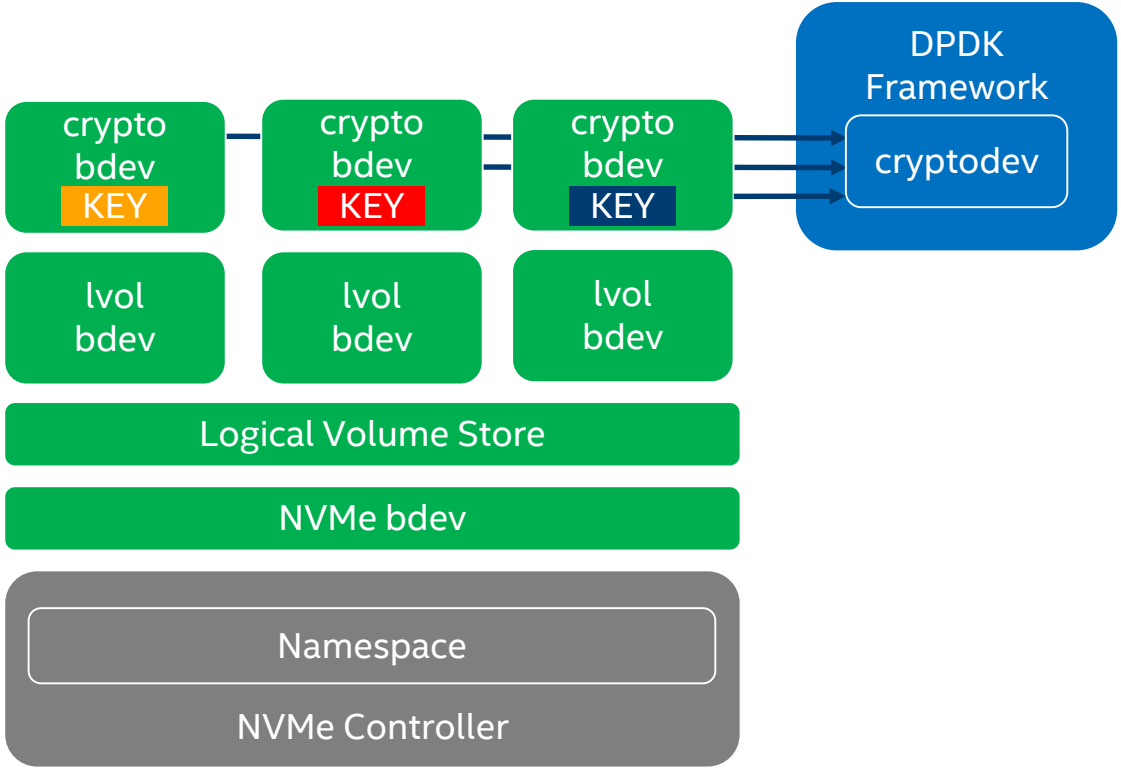

OPEN CHANNEL

Extend nvme driver API

Ideal environment for Open Channel based applications

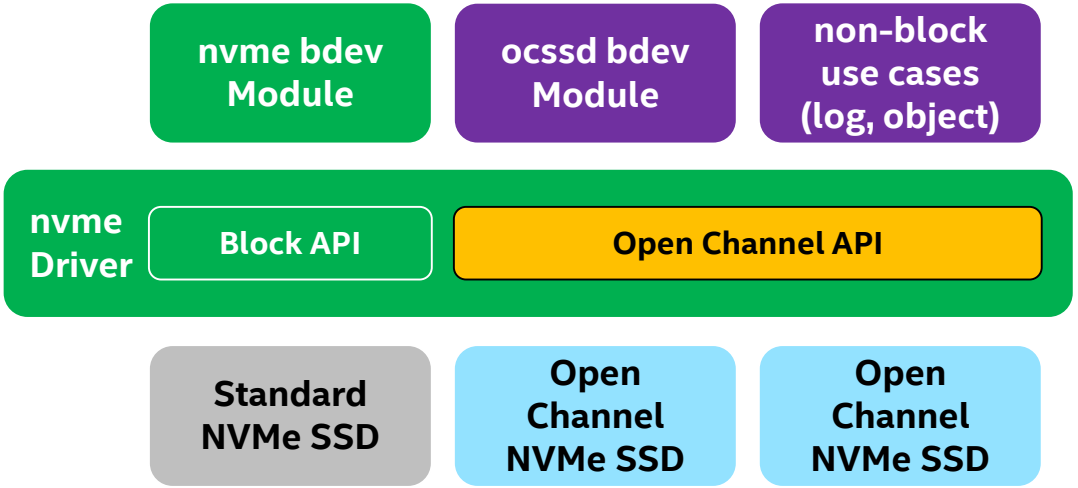

CONFIGURATION

Full JSON-RPC Configuration

spdkcli - Terminal-based user interface

```
% sudo scripts/spdkcli.py                                                    (git)-[master]-
SPDK CLI v0.1

/> ls
o- / ..................................................................................... [...]
  o- bdevs ..................................................................................... [...]
  | o- AIO ..................................................................................... [Bdevs: 0]
  | o- Logical_Volume .......................................................................... [Bdevs: 4]
  | | o- 2f0293da-2106-40c2-8837-57a86a8be9b4 ..................... [lvs0/lvol0, Size=1.0G, Not claimed]
  | | o- 47c8d645-3f10-4f08-acfb-0731b7c21aa4 ..................... [lvs0/lvol3, Size=40.0G, Not claimed]
  | | o- 89859d4a-d6cc-4bcc-821d-29a52eb1fc75 ..................... [lvs0/lvol2, Size=20.0G, Not claimed]
  | | o- a9fe328b-a9f5-42c8-a888-f290a1f3a5e7 ..................... [lvs0/lvol1, Size=10.0G, Not claimed]
  | o- Malloc .................................................................................. [Bdevs: 2]
  | | o- Malloc0 ..................................................... [Size=512.0M, Not claimed]
  | | o- Malloc1 ..................................................... [Size=512.0M, Not claimed]
  | o- NVMe .................................................................................... [Bdevs: 1]
  |   o- Nvme0n1 ................................................................ [Size=3.6T, Claimed]
  o- lvol_stores ................................................................... [Lvol stores: 1]
    o- lvs0 ..................................................................... [Size=3.6T, Free=3.6T]
/>
```

FUTURE WORK

- Compression
- Database Integration
- Block Device Aggregation
- Mirroring/Replication
- Fibre Channel
- Deduplication

- Current State of the Project
- Looking Forward
- Call to Action

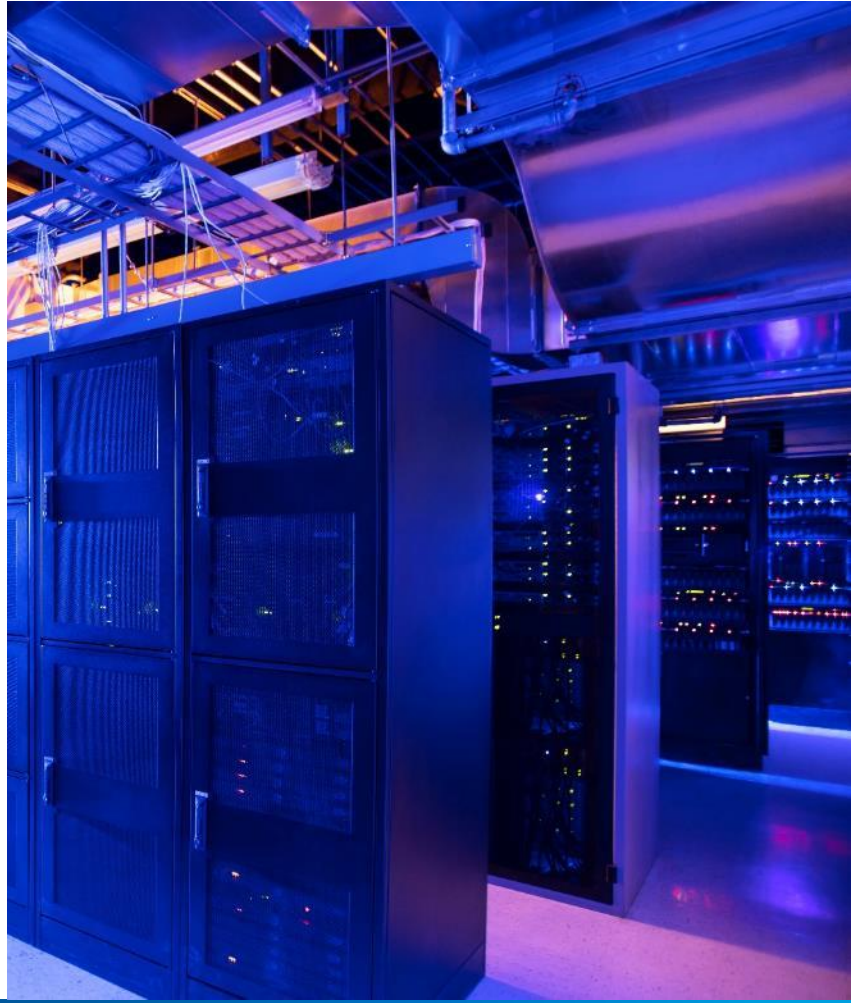

CALL TO ACTION

Learn today (and tomorrow) from SPDK experts

Join the community!

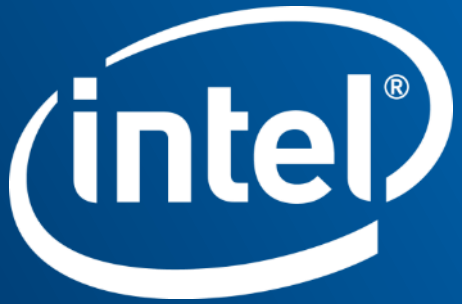
intel®