



SPDK: STATE OF THE PROJECT

Jim Harris

Principal Software Engineer

Data Center Group



By Edmund Garman from Salem, Oregon, USA - Waikiki, Honolulu Panorama
CC BY 2.0, <https://commons.wikimedia.org/w/index.php?curid=39574339>

AGENDA

- Community Growth
- Feature Growth
- Looking Forward



SPDK PATCH COMMIT STATISTICS

Version Range	Total Commits	Non-Intel Committers	Non-Intel Commit %
v1.0 .. v16.12	1360	13	2.6%
v16.12 .. v18.01	2696	31	4.6%
v18.01 .. v19.01	3083	51	14.7%

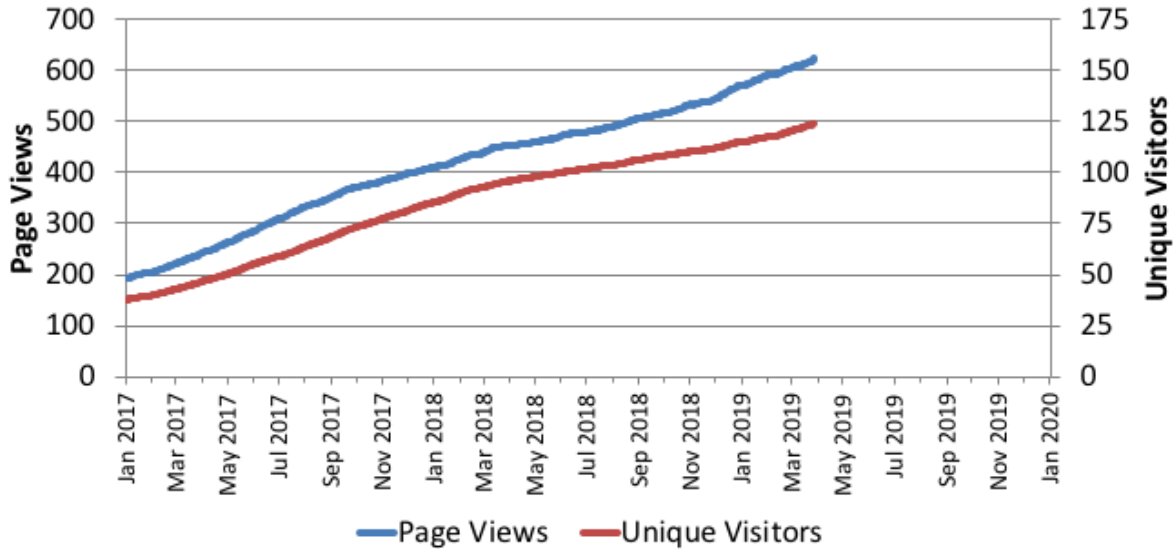
SPDK PATCH COMMIT STATISTICS

Version Range	Total Commits	Non-Intel Committers	Non-Intel Commit %
v1.0 .. v16.12	1360	13	2.6%
v16.12 .. v18.01	2696	31	4.6%
v18.01 .. v19.01	3083	51	14.7%

2018
+65% Non-Intel Committers
3x Non-Intel Commits

COMMUNITY GROWTH

SPDK Github Traffic 52 Week Running Daily Averages



+30% YoY
+218% since Jan'17

RECENTLY ADDED FEATURES

SPDK 18.04

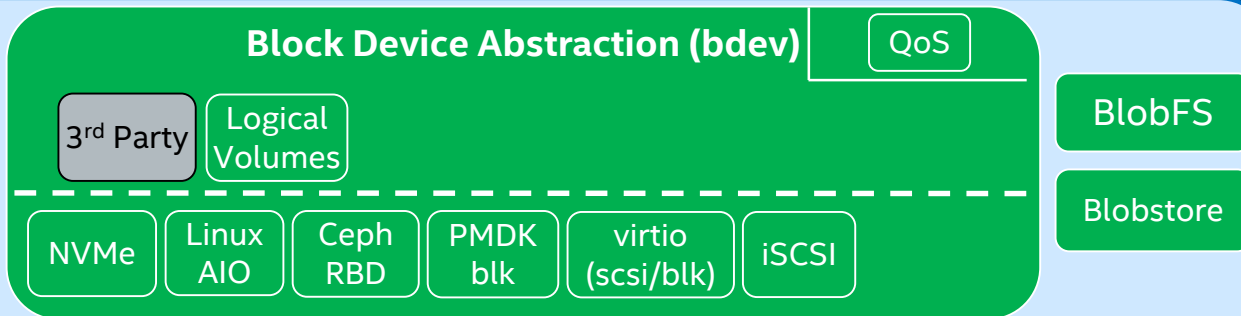
Added since 18.04

In Progress


Storage Protocols



Storage Services



Drivers



Integration

VPP TCP/IP

RocksDB

Ceph

QEMU

Tools

fio

nvme-cli

RECENTLY ADDED FEATURES

SPDK 18.04

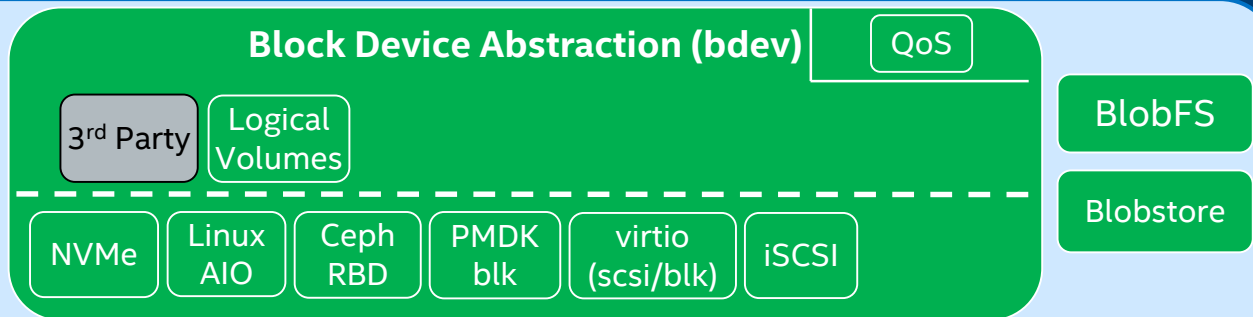
Added since 18.04

In Progress


Storage Protocols



Storage Services



Drivers



Integration

VPP TCP/IP

RocksDB

Ceph

QEMU

Tools

fio

nvme-cli

NVME/TCP



NVMe TP ratified November 2018

SPDK added TCP transport for

- NVMe driver
- NVMe-oF target

Supports alternative TCP stack implementations

RECENTLY ADDED FEATURES

SPDK 18.04

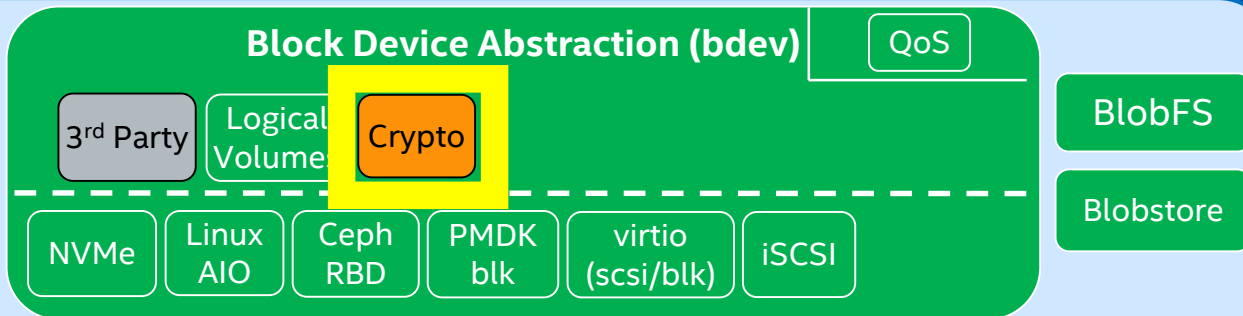
Added since 18.04

In Progress

Storage Protocols



Storage Services



Drivers



Integration

VPP TCP/IP

RocksDB

Ceph

QEMU

Tools

fio

nvme-cli

CRYPTO

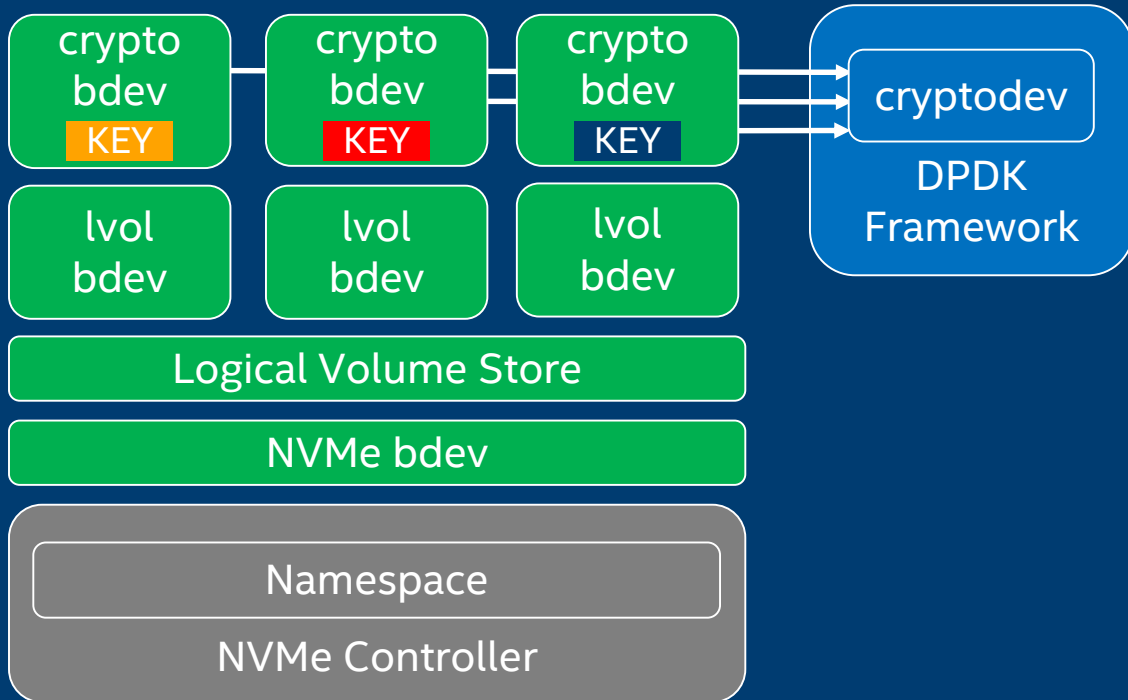


Hardware Accelerators

- Intel Quickassist

Optimized Software

- ISA-L




RECENTLY ADDED FEATURES

Storage Protocols



Storage Services



Drivers



SPDK 18.04

Added since 18.04

In Progress

Cinder

RocksDB

Ceph

QEMU

Tools

fio

nvme-cli

CINDER

Provision ephemeral storage using SPDK

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

Kubernetes integration in progress



RECENTLY ADDED FEATURES

SPDK 18.04

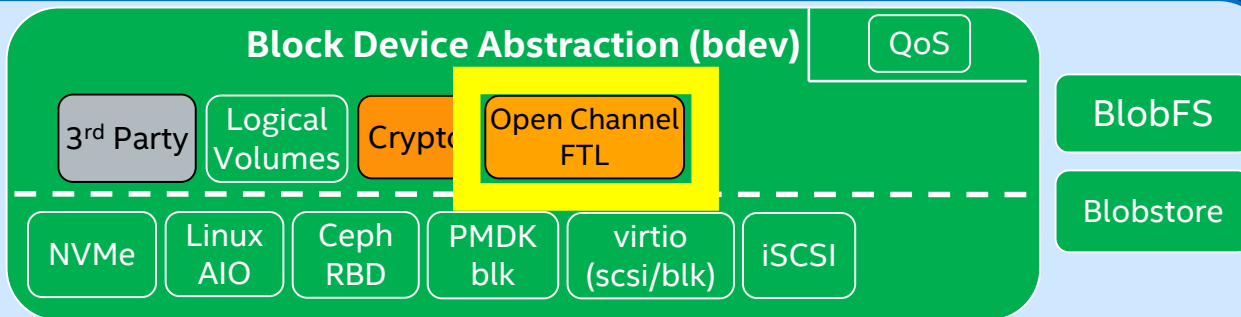
Added since 18.04

In Progress

Storage Protocols



Storage Services



Drivers



Integration

Cinder

VPP TCP/IP

RocksDB

Ceph

QEMU

Tools

fio

nvme-cli

OPEN CHANNEL BLOCK FTL

Host FTL enabling smart data placement


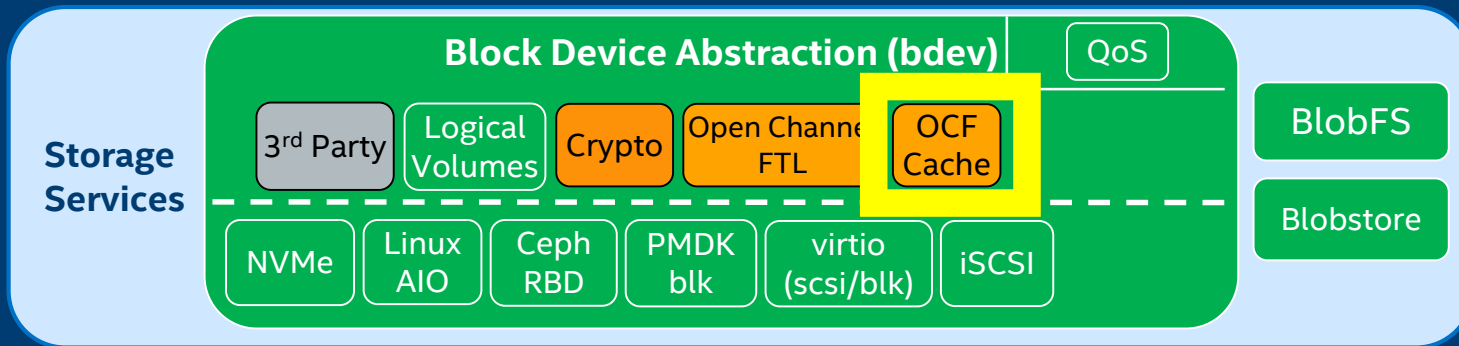
- Based on OC2.0 specification

Block FTL support added to bdev nvme module

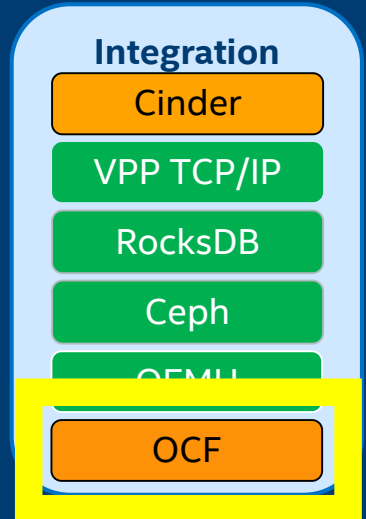
Long term goal: Zoned Namespace API

- With ZNS/OC adapters

RECENTLY ADDED FEATURES



- SPDK 18.04
- Added since 18.04
- In Progress



OPEN CAS FRAMEWORK

Block Caching Library

Open Sourced in Q4'18

SPDK integration added v19.01

RECENTLY ADDED FEATURES

SPDK 18.04

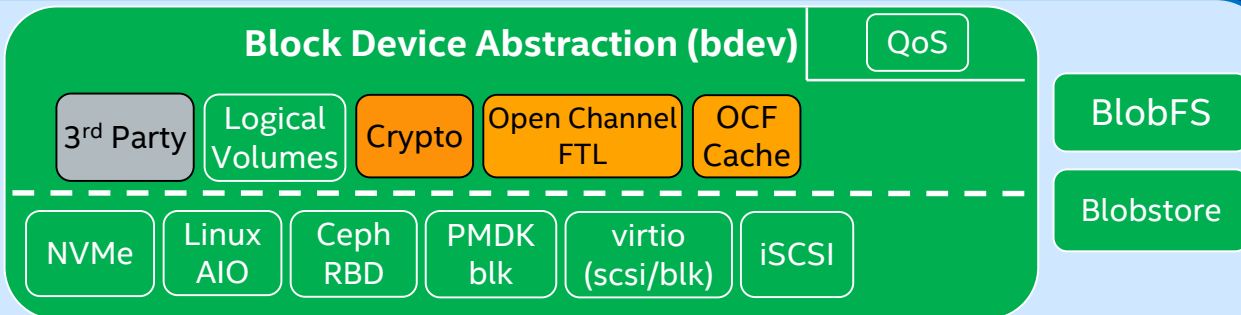
Added since 18.04

In Progress

Storage Protocols



Storage Services



Drivers



Integration

Cinder

VPP TCP/IP

RocksDB

Ceph

QEMU

OCF

Tools

fio

spdk-cli

spdk-cli

SPDK-CLI

```
0- / ..... [..]
  0- bdevs ..... [..]
    | 0- aio ..... [Bdevs: 0]
    | 0- error ..... [Bdevs: 0]
    | 0- iscsi ..... [Bdevs: 0]
    | 0- logical_volume ..... [Bdevs: 0]
    | 0- malloc ..... [Bdevs: 0]
    | 0- null ..... [Bdevs: 0]
    | 0- nvme ..... [Bdevs: 0]
    | 0- pmemblk ..... [Bdevs: 0]
    | 0- rbd ..... [Bdevs: 0]
    | 0- split_disk ..... [Bdevs: 0]
    | 0- virtioblk_disk ..... [Bdevs: 0]
    | 0- virtioscsi_disk ..... [Bdevs: 0]
  0- lvol_stores ..... [Lvol stores: 0]
  0- nvme ..... [..]
    0- subsystem ..... [Subsystems: 1]
      | 0- nqn.2014-08.org.nvmeexpress.discovery .... [st=Discovery, Allow any host]
        | 0- hosts ..... [Hosts: 0]
        | 0- listen_addresses ..... [Addresses: 0]
    0- transport ..... [Transports: 0]
/>
```


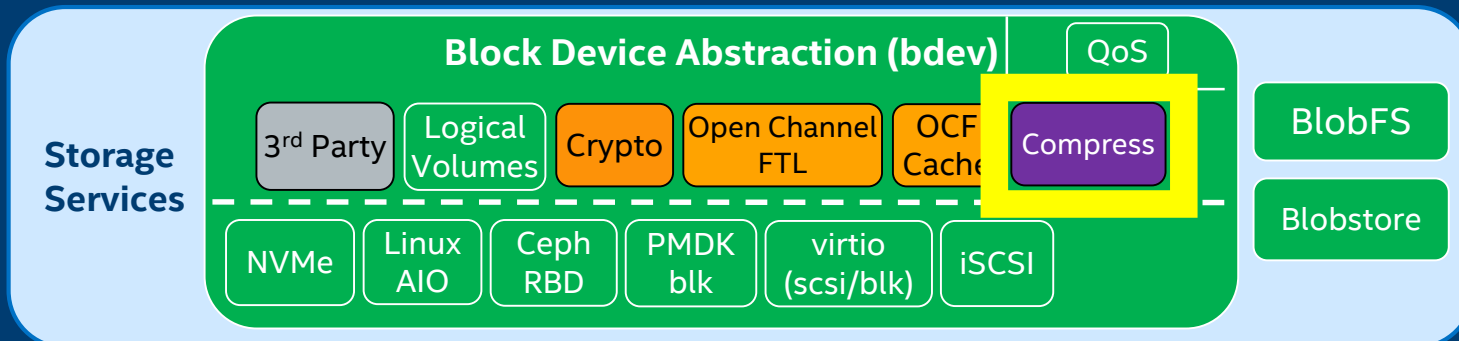
LOOKING FORWARD

Usage Models

- Rack Scale Architecture
- Smart NICs

New SPDK Features

SPDK ARCHITECTURE



- SPDK 18.04
- Added since 18.04
- In Progress

- Integration**
- Cinder
 - VPP TCP/IP
 - RocksDB
 - Ceph
 - QEMU
 - OCF

- Tools**
- fio
 - nvme-cli
 - spdk-cli

BLOCK COMPRESSION



DPDK Framework for compression algorithms

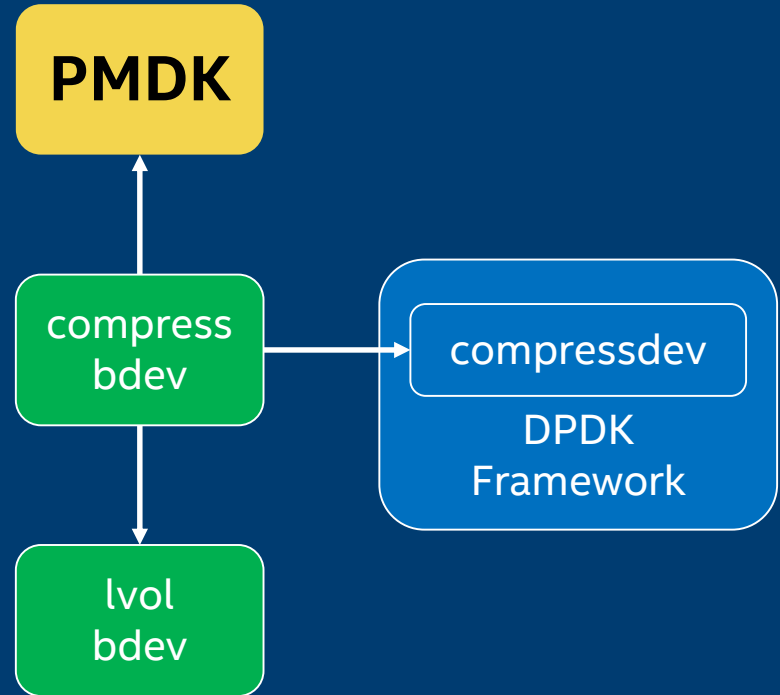
- Accelerators or Optimized SW

PMDK for persistent memory access


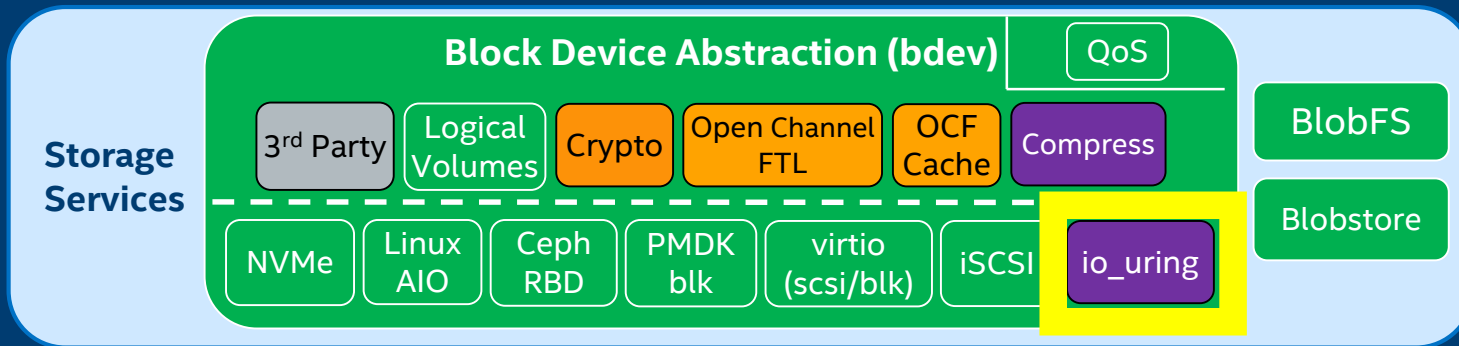
- For compression metadata

SPDK logical volumes for storage

- For realizing capacity savings



SPDK ARCHITECTURE



- SPDK 18.04
- Added since 18.04
- In Progress

Integration

- Cinder
- VPP TCP/IP
- RocksDB
- Ceph
- QEMU
- OCF

Tools

- fio
- nvme-cli
- spdk-cli

IO_URING

New Linux kernel interface for asynchronous I/O

- Targeted for Linux 5.1
- Significant efficiency improvements compared to Linux AIO

SPDK io_uring bdev module



FUTURE WORK

Smart Scheduling

Block Indirection


Zoned Namespaces

Replication and Parity

Even better performance!

CALL TO ACTION

Take advantage of talks and demos to learn more about SPDK
and PMDK and VTune™ Amplifier
Join the SPDK Community



BACKUP

