



Open. Together.




OCP
SUMMIT

Storage

Toward the large deployment of open channel SSD

Feng Zhu, Yu Du, Wentao Zhang, Sheng Qiu, Yan Liu, Shu Li

Alibaba Group



Open. Together.

Agenda

Introduction to AliFlash

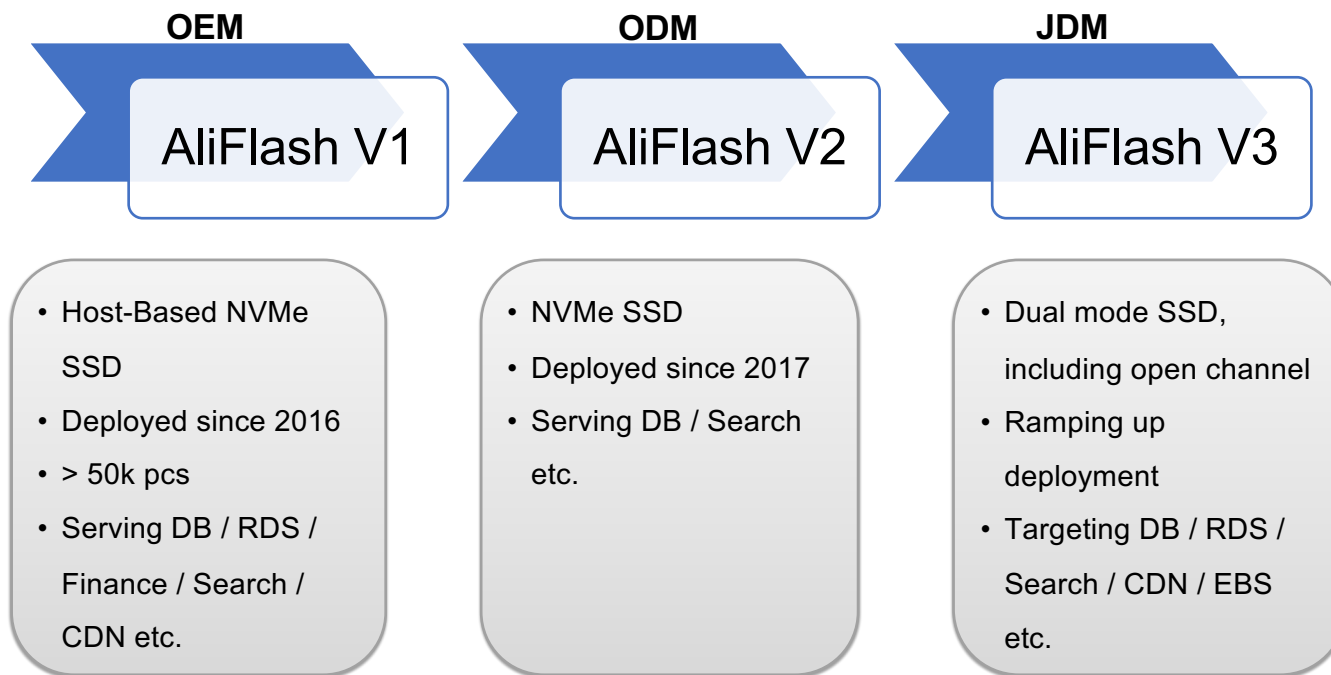
Replica SSD based on open channel

Multi-stream SSD based on open channel



Open. Together.

AliFlash @Alibaba




AliFSC V1



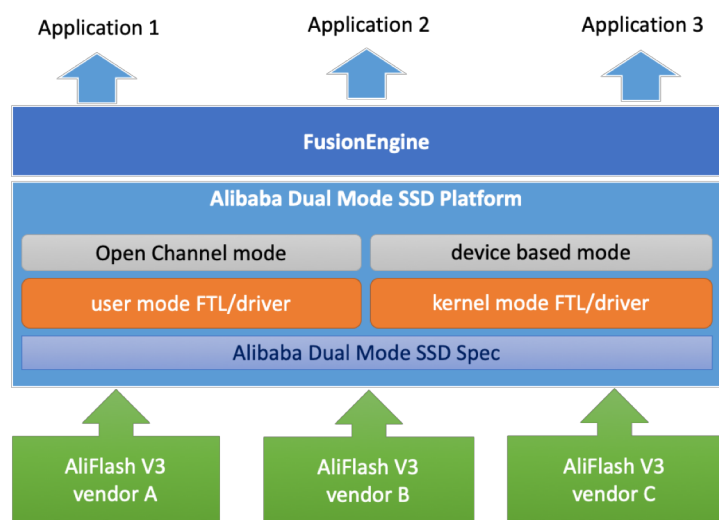
Alibaba's first self-developed SSD controller AliFSC V1.

- Support Open Channel mode
- PCIe Gen 3 x 4
- NVMe 1.3
- 16 NAND channels , drive capacity up to 16TB
- 3 dual core Arm Cortex R5
- 4KB random read IOPS up to 800k
- Support all major vendors 3D TLC/QLC NAND

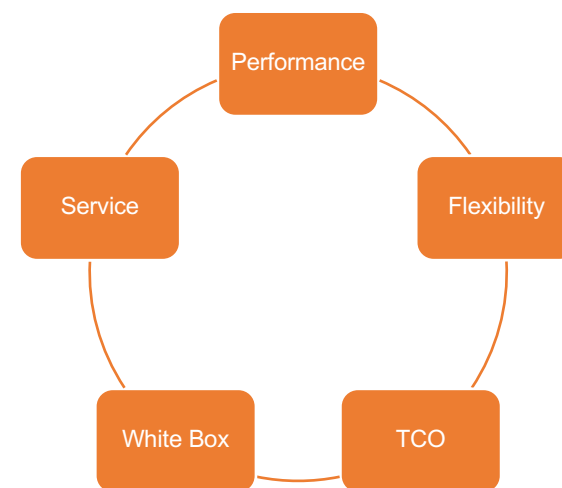


Open. Together.

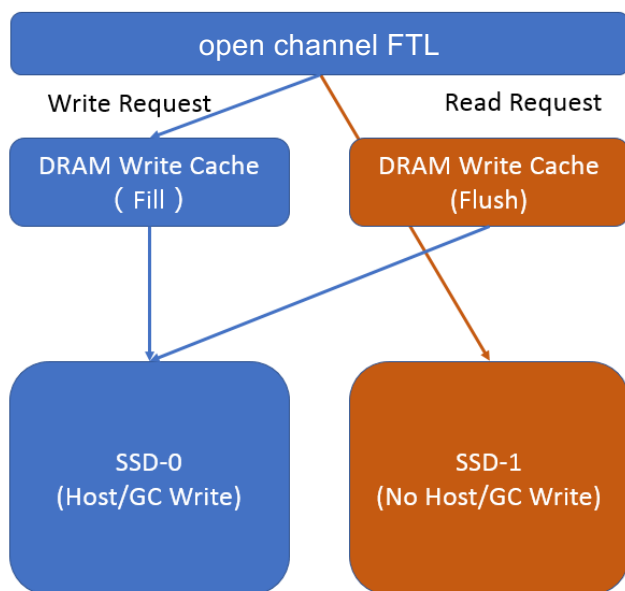
Dual Mode SSD



- Demand for HW/SW co-optimization
- Diversified, fast-changing workloads
- Continuous pressure for TCO and supply improvement
- “white box” of I/O path - more determinism
- Quick response to issues and requests

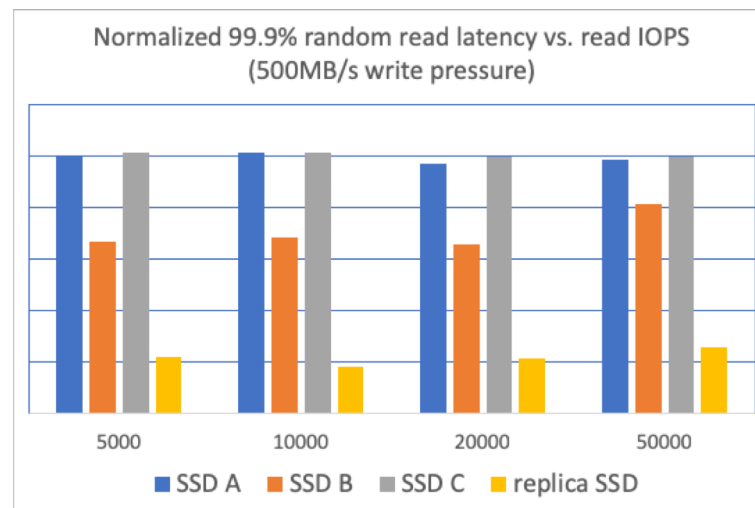



Replica SSD through open channel



- Build upon open channel SSD. Transparent to host software.
- Physical partitioning the NVM resources to two sets SSD-0 and SSD-1
- Each LBA has dual copies on the same drive
- Read and write are dispatched to different set

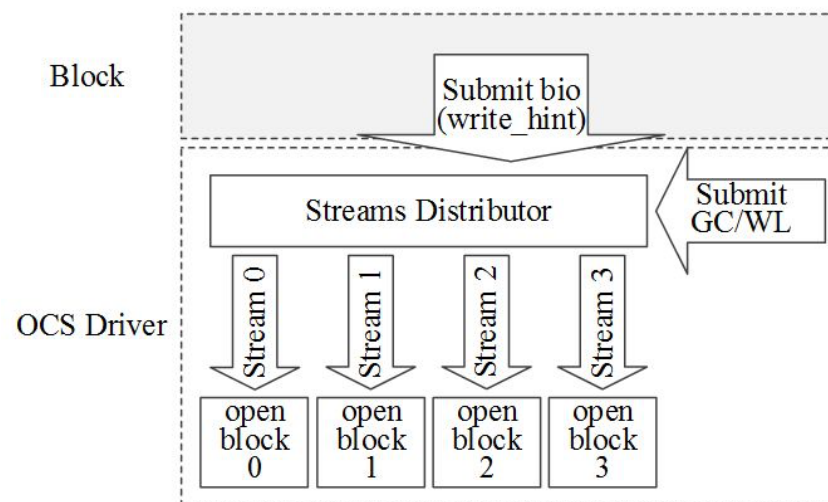
Replica SSD results



- Read-only like latency achieved

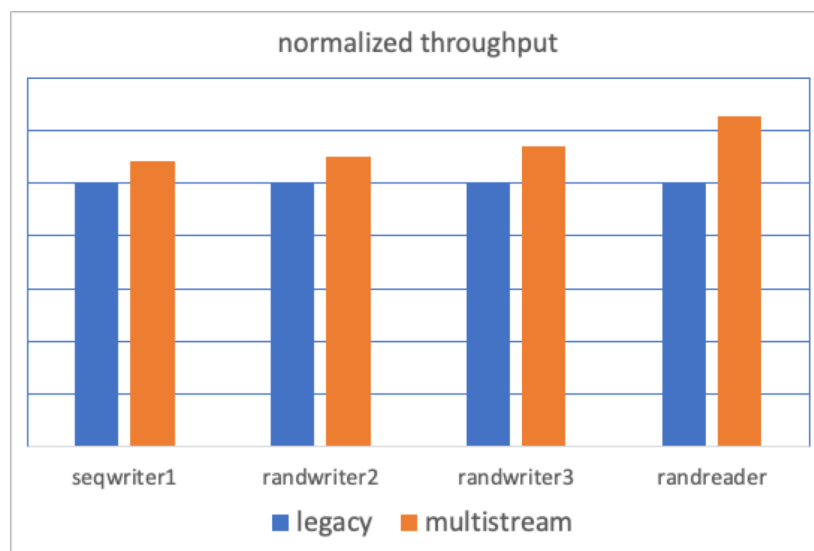
Multi-stream through open channel

- Multi-stream to improve WA and QoS:
Separate the data with different life cycles into different streams
- Simplified FTL to handle the address collision

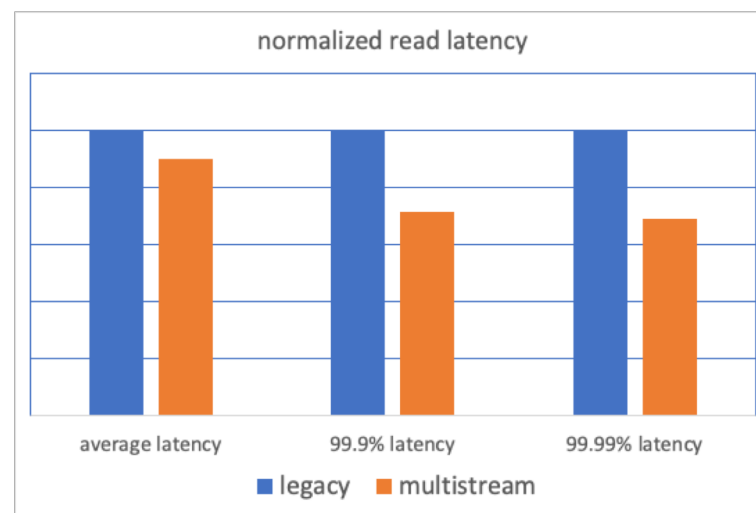
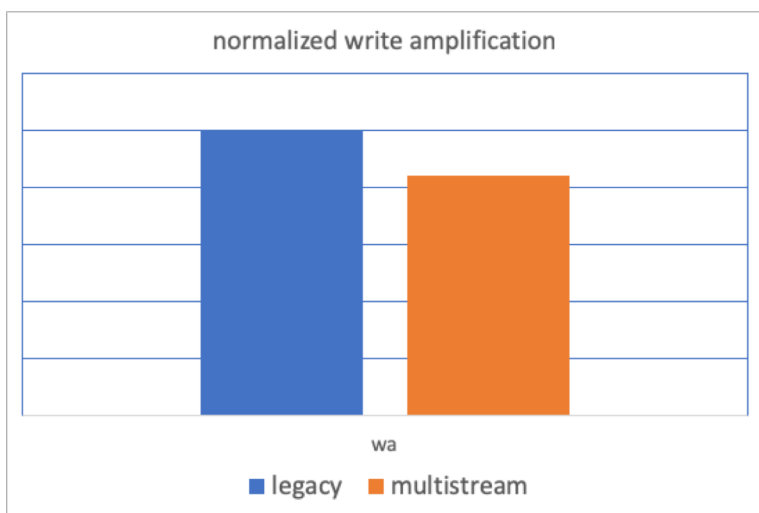


Multi-stream results

- Seqwriter1: sequential write (0, 32GB) --> log write
- Randwriter2: random write (64GB, 96GB) --> meta data update
- Randwriter3: random write (128GB, 640GB) --> data update
- Reader: random read (128GB, 640GB) --> data read



Multi-stream results



Alibaba Open Channel Ecosystem



Alibaba is collaborating with major vendors in industry to build an ecosystem for Open Channel SSD



Share development & debug resources



Reduce time & complexity for SSD qualification in Alibaba



Massive deployment in 2019



Open. Together.

What's Next

Further I/O scheduling and QoS optimization

QLC deployment in open channel

On-line disk defect monitoring and prediction

...



Open. Together.



Open. Together.

OCP Global Summit | March 14–15, 2019

