FMS the Future of Memory and Storage

SOLIDIGM

Cloud Storage Acceleration Layer Append Cache

Presenter: Wojciech Malikowski, Software Engineer, Solidigm

August 7, 2024

Authors: Wojciech Malikowski, Mariusz Barczak, Mateusz Kozłowski

Solidigm Disclaimers

All information provided here is subject to change without notice.

- The products described in this document may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Solidigm 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.
- Solidigm disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.
- Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase.
- Cost reduction scenarios described are intended as examples of how a given Solidigm-based product, in the specified circumstances and configurations, may affect future costs and provide cost-savings. Circumstances will vary. Solidigm does not guarantee any costs or cost reduction.
- Solidigm does not control or audit the design or implementation of third-party benchmark data or Web sites referenced in this document. Solidigm encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.
-

CSAL Evolution

2022

- CSAL for Alibaba ECS D3C Service compute-intensive big data local disk instance release
- CSAL release in SPDK 22.09
- FMS 2022 CSAL Enabling Storage Performance & Capacity Scaling for BIG DATA

2023

- FMS 2023 New Developments in Cloud Storage Acceleration Layer (CSAL), an FTL in SPDK
  Reference Storage Platform (RSP)
  CSAL based high-density design for QLCs and SLCs

2024

- EuroSys 2024 CSAL: the Next-Gen Local Disks for the Cloud
  How CSAL removes QLC's adoption challenges and provides QLC based local disk architecture in Alibaba Cloud
- FMS2024 CSAL
  1. FDP NVMe Technology Implementation and Scaling with PCIe Gen5 Cache SSD
  2. CSAL Append Cache
  3. CSAL demo of running L2P on CXL

What's CSAL?

- Open-source cloud-scale shared-nothing Flash Translation Layer (FTL bdev) in Storage Performance Development Kit (SPDK)
- CSAL provides transparent block device to the upper application
- Ultra-fast cache and write shaping tier to improve performance and endurance to scale QLC value
- Consistent performance in multi-tenant environment
- Flexible scaling of NAND performance and capacity to the user/workload needs

Cloud Storage Acceleration Layer (CSAL)

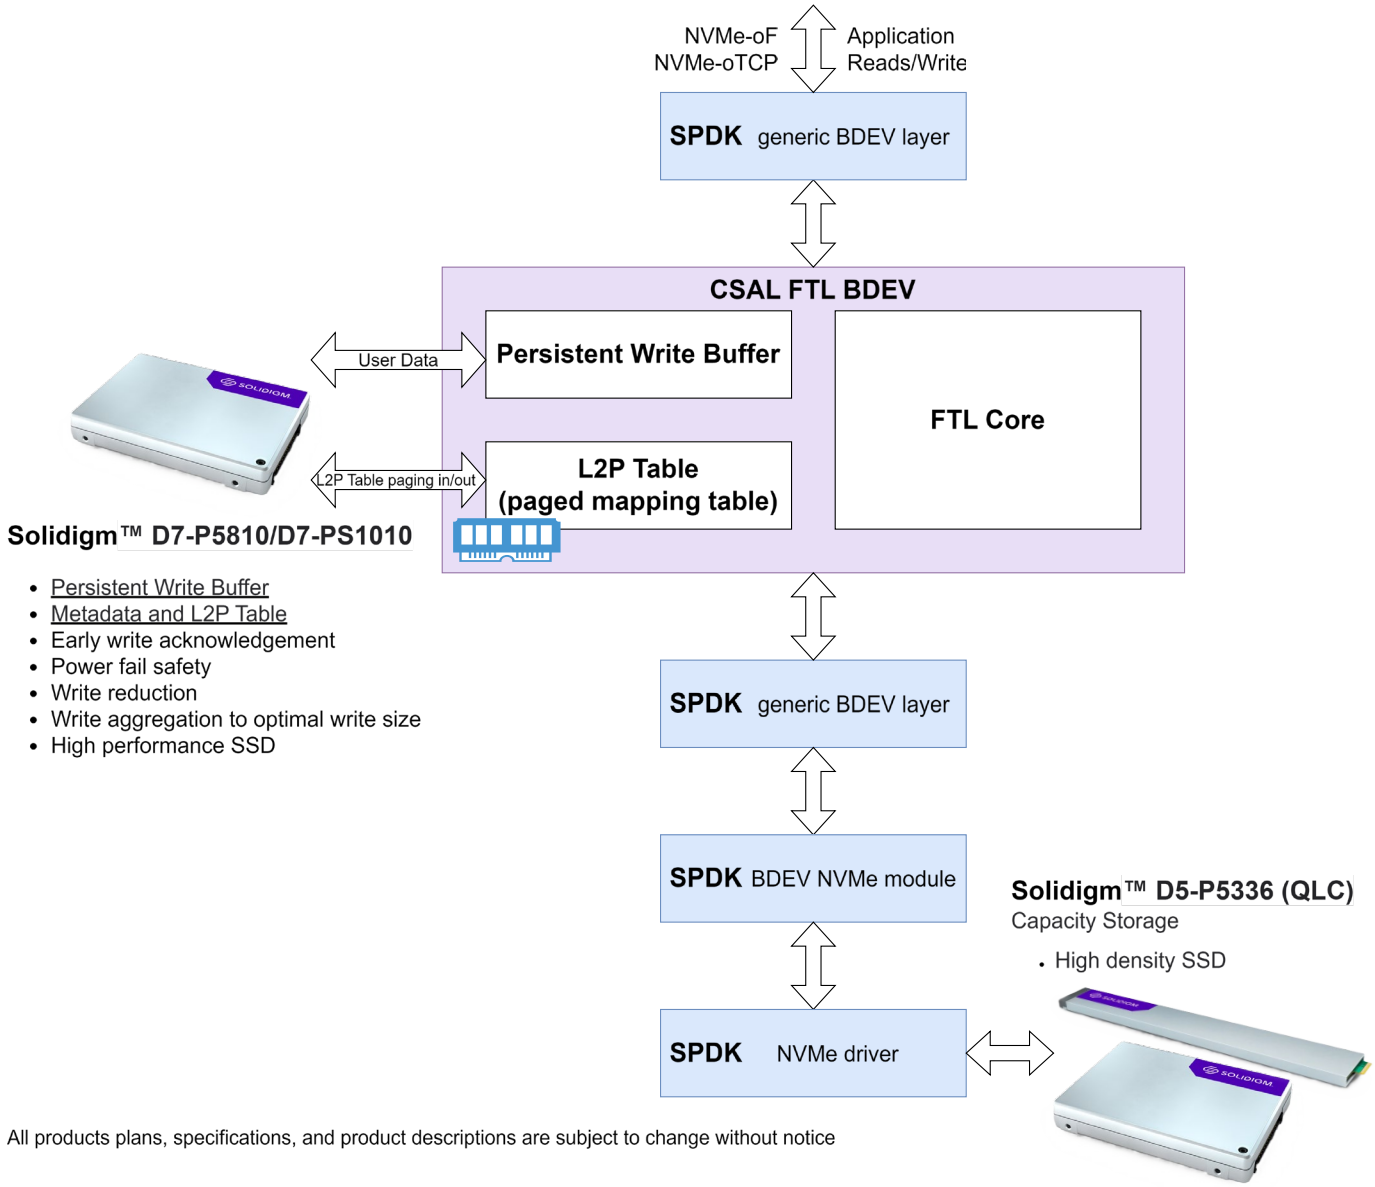

All products plans, specifications, and product descriptions are subject to change without notice

Introducing CSAL append cache

- Optimized use of NAND media as a cache device
- Reduction in write amplification factor (WAF)
- Faster cache lookup times

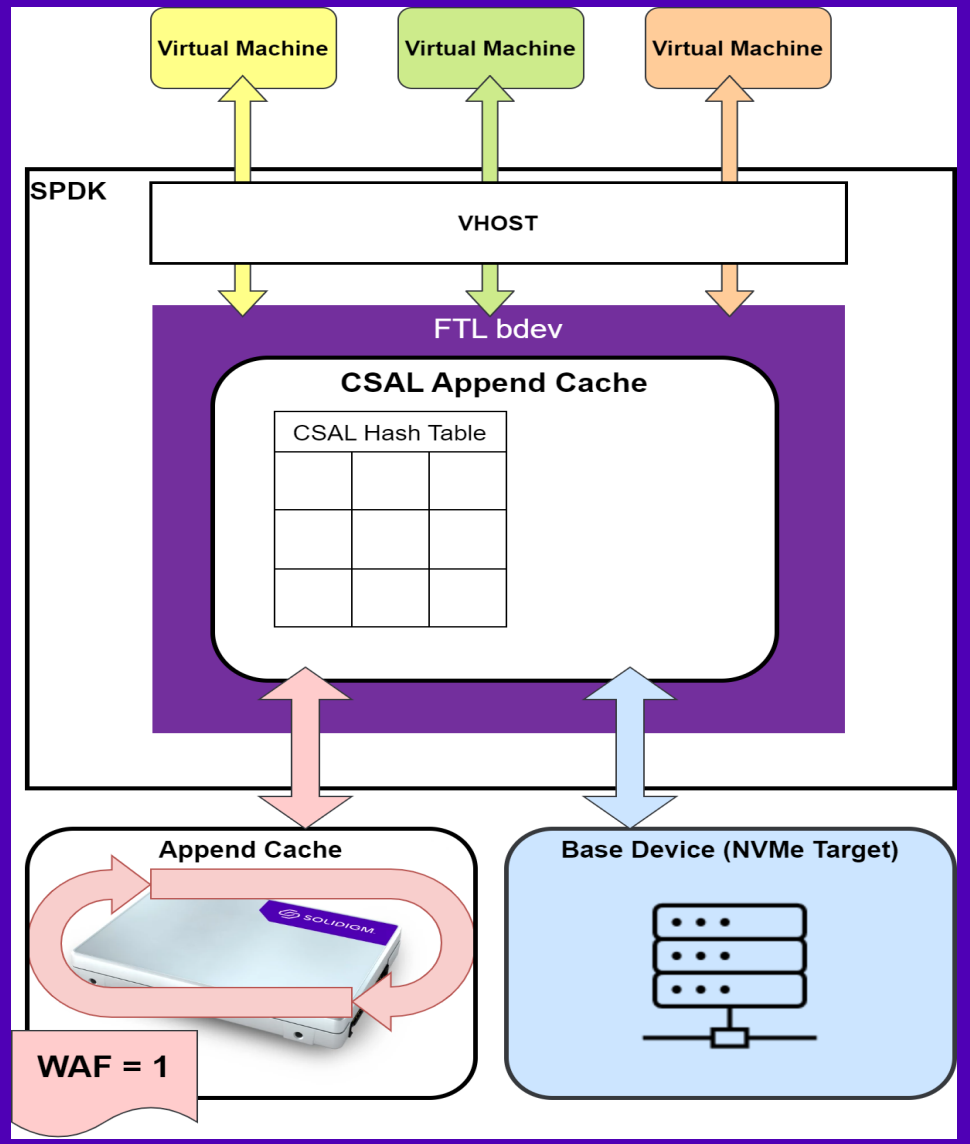

Addressing Key Technical Challenges

Challenges:

- Write Amplification Factor (WAF)
- Cache Fragmentation
- Cache Lookup Latency

Our solution:

- No Random Workload to Cache Device: Ensures consistent performance
- Append-Only Access Pattern: Minimizes complexity and wear
- Efficient Mapping and Lookup: One hash item for larger I/Os
- Shared nothing architecture

A Legacy Cache Hash Table

CPU Cache Miss Latency	~100 ns
IO Size	64 kiB
Cache Line Size	4 kiB
Cache Lookups	16
CPU Cache Misses	80
Cache Look time	~8000 ns
Number of cache IOs	16

CSAL Hash Table

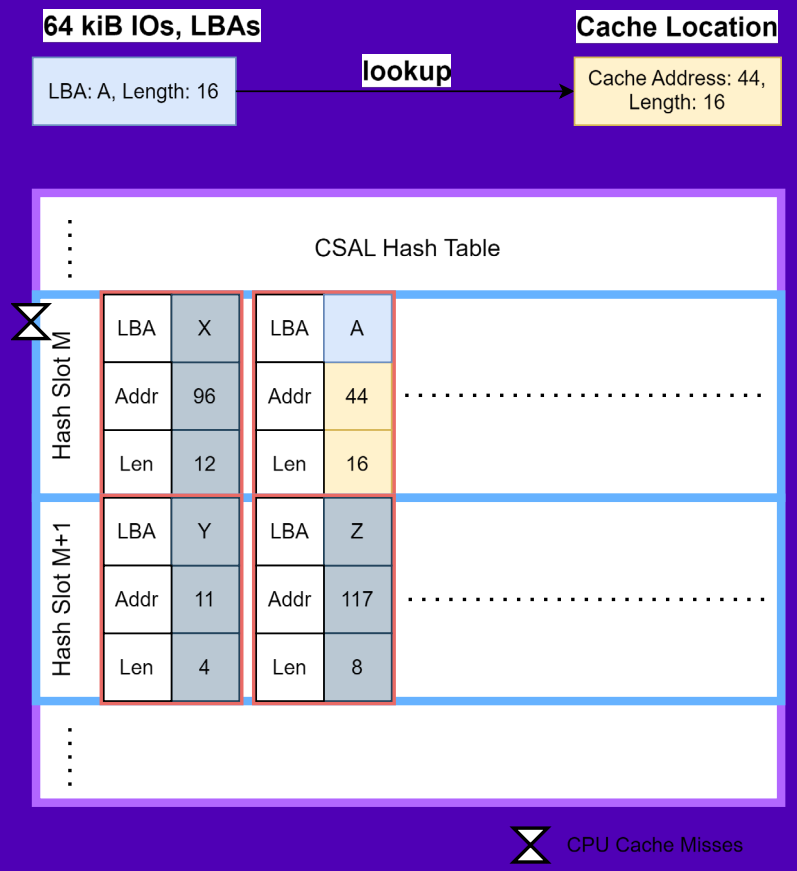

	Legacy Cache	CSAL Append Cache
CPU Cache Miss Latency	~100 ns	
IO Size	64 kiB	
Cache Line Size	4 kiB	N/A
Cache Lookups	16	1
CPU Cache Misses	80	1
Cache Look time	8000 ns	100 ns
Number of cache IOs	16	1

Evaluation

Comparing CSAL with Open CAS (Open Cache Acceleration Software)

Workload and setup:

- 8 jobs
- Zipfian 1.2 distribution – strong cache locality
- Read write ratio 70:30
- 128 IO depth
- 64K block size

Precondition: Fill all partitions executing sequential write

- Test setup and config:
  - Software: CSAL SPDK, Open-CAS kernel
  - Solidigm Drives: D5-P5316 (QLC, 15 TB), D7-P5520 (TLC, 800 GB)
  - Server: Intel® Xeon® Platinum 8380
  - Operating system: Fedora 39, kernel version: 6.8.7-200.fc39.x86_64

Write perfomance comparison

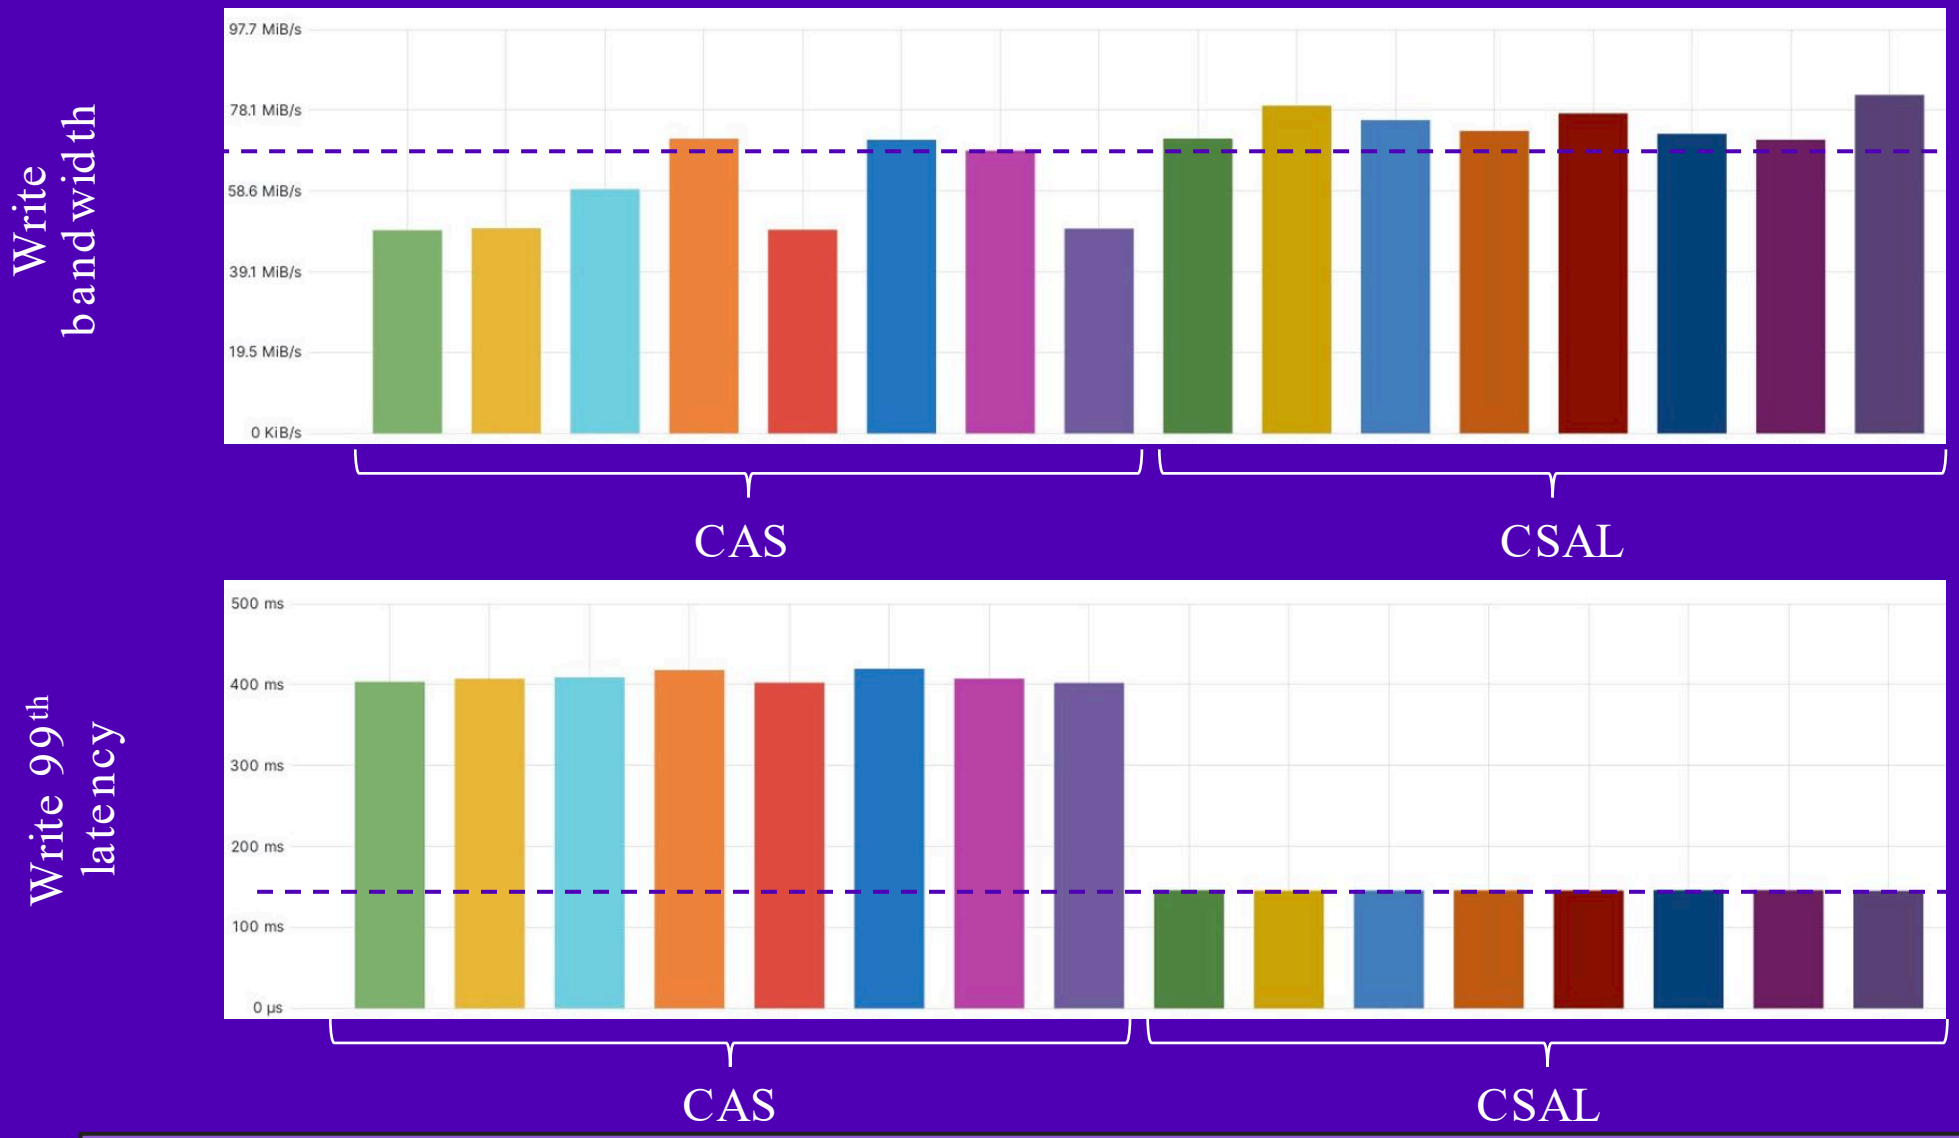

Delivering consistently higher bandwidth and significantly improving the write latency QoS

Read perfomance comparison

Read bandwidth

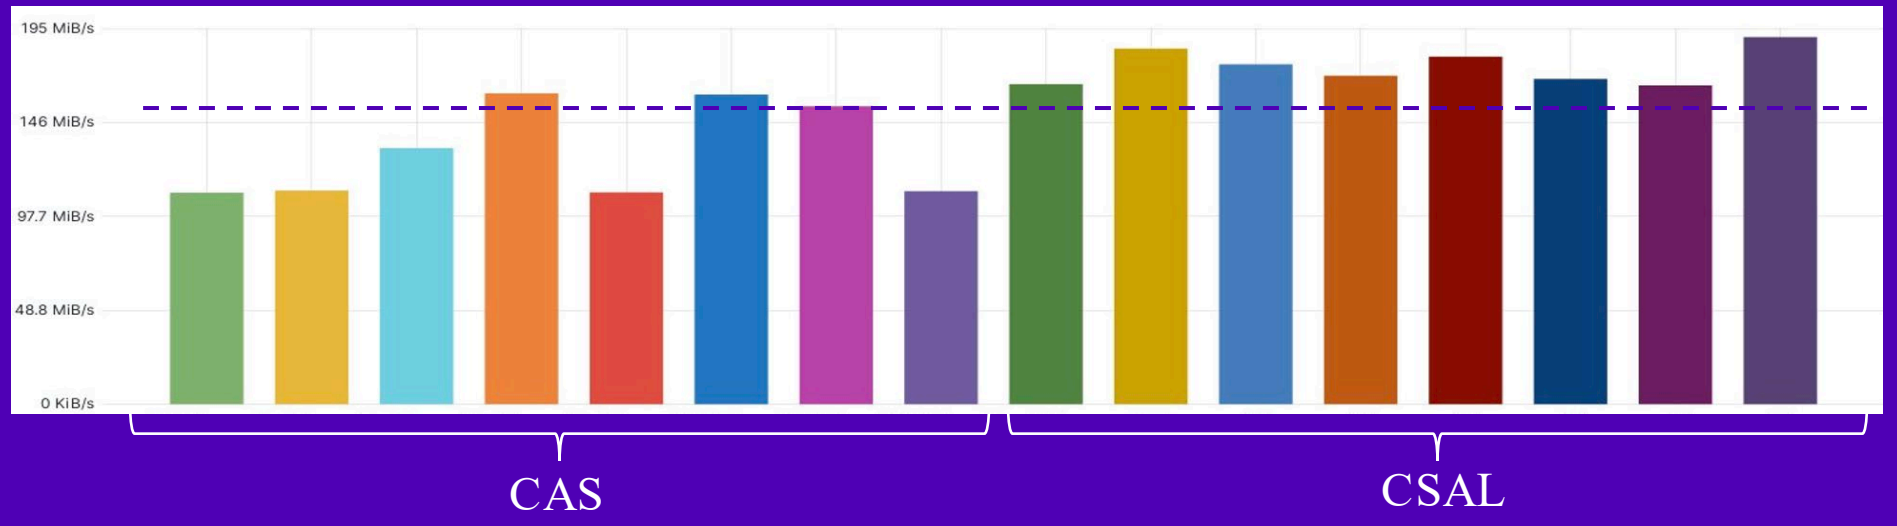

Read 99th latency

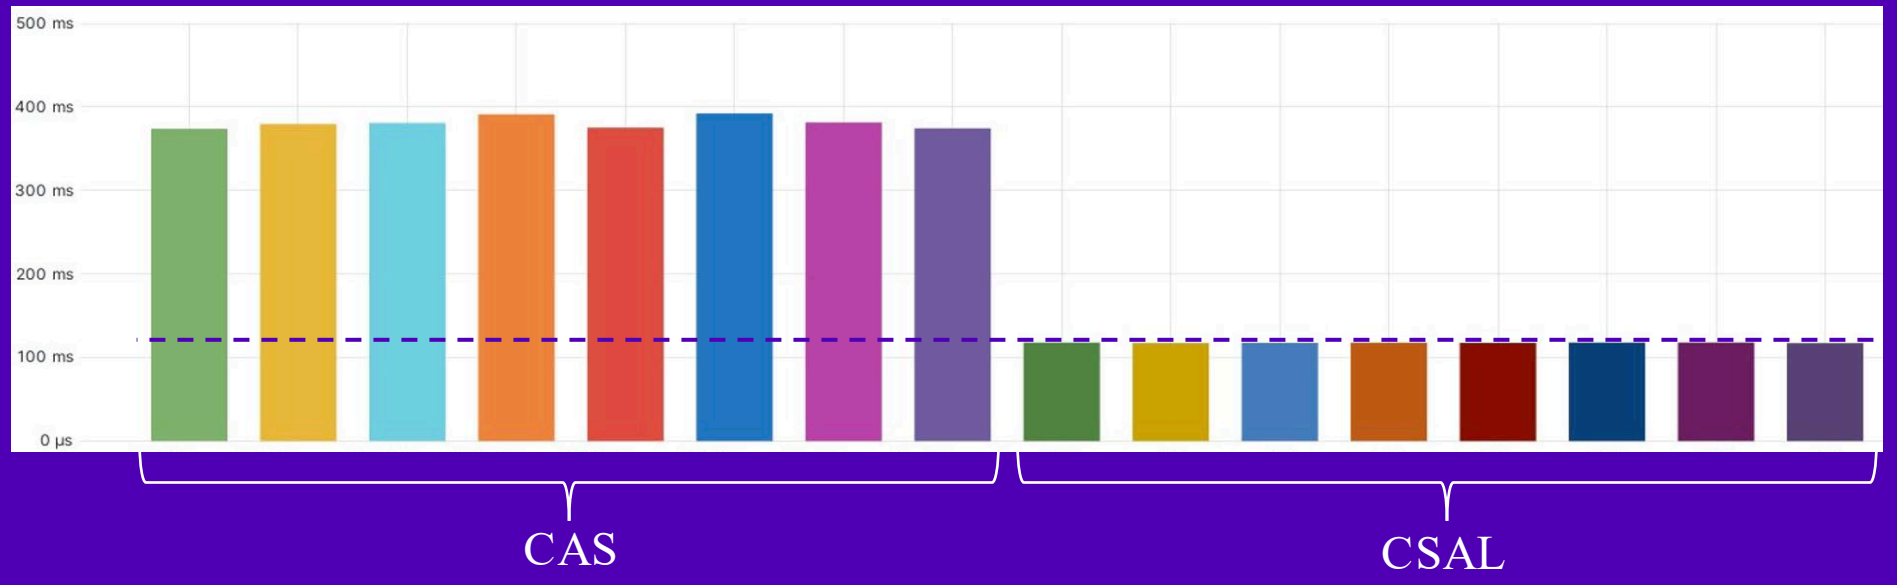

Delivering consistently higher bandwidth and significantly improving the read latency QoS

Cache device write amplification factor comparison between Open-CAS and CSAL append cache

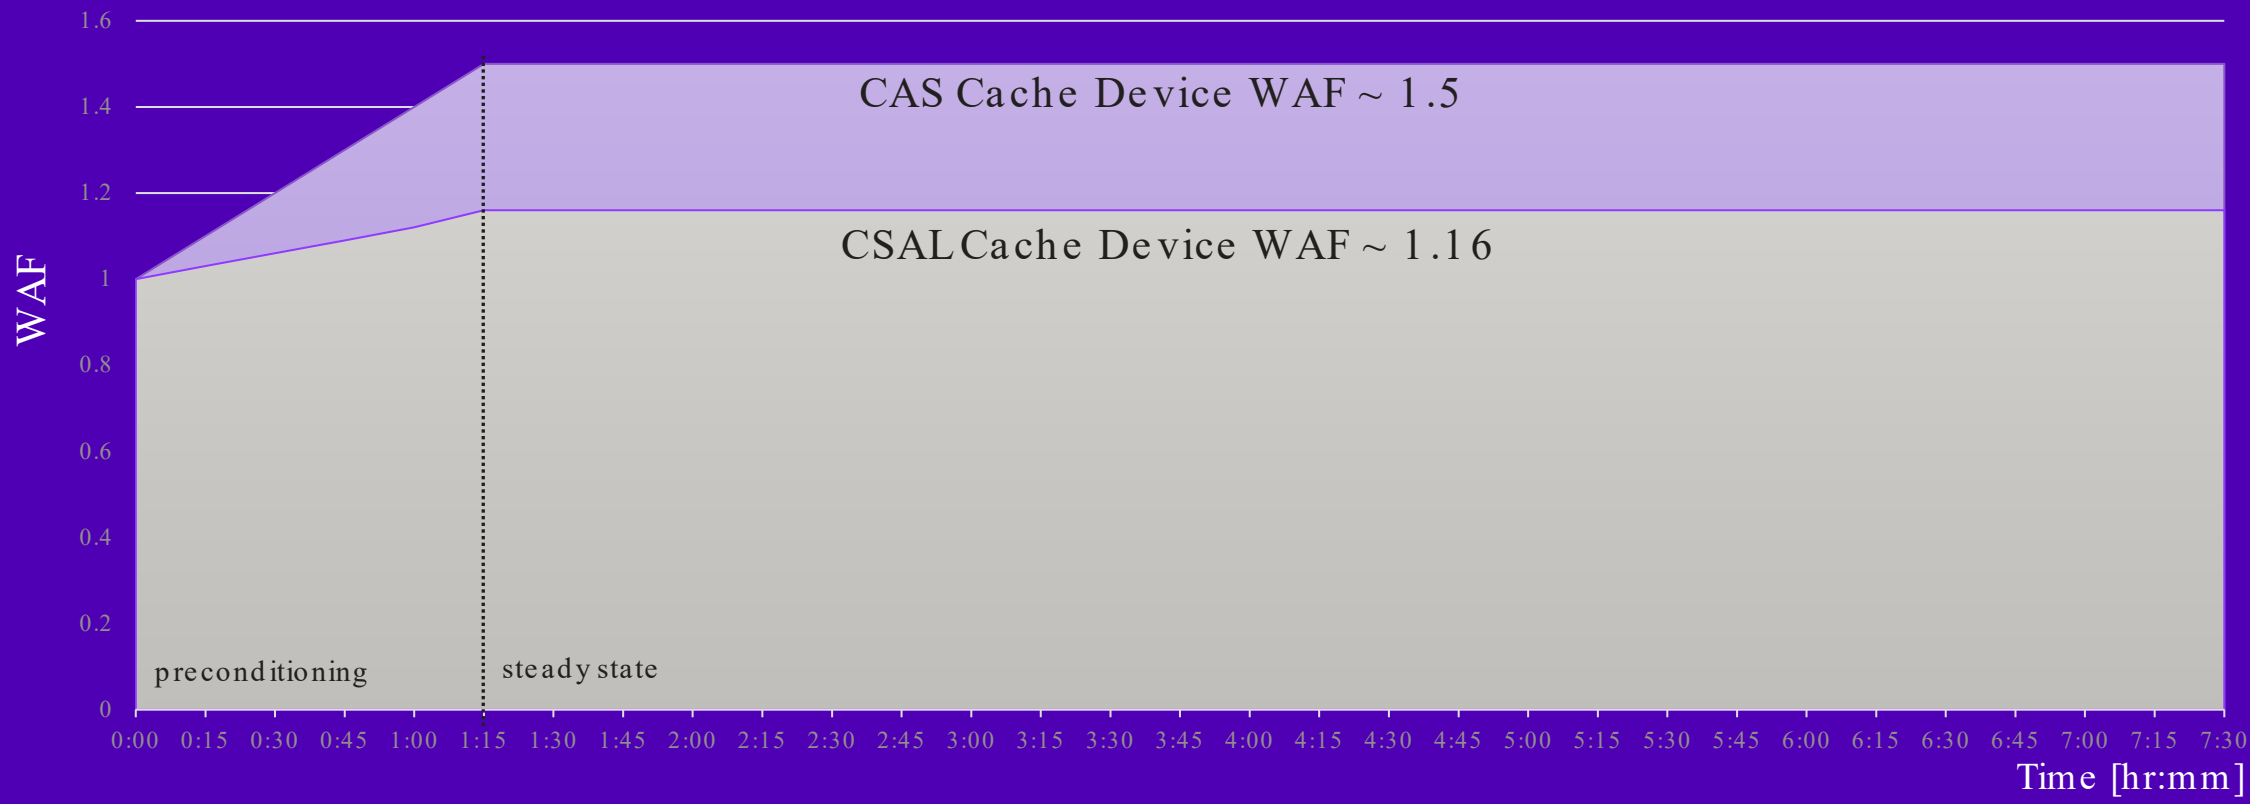

CSAL append cache manages SSD WAF level ~ 1.0

Conclusion

Benefits of CSAL when using cache append mode

- Delivers consistently higher write and read bandwidth
- Improving cache tail latency (better QoS)
- CSAL append cache manages SSD WAF level of ~1.0 thus increases the cache endurance

Q&A
Cloud Storage Acceleration Layer Append Cache

Contact us:
dl_csal@solidigm.com