SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers

Hyatt Regency Santa Clara, CA
September 15-17, 2025

CSAL with Core Scaling for RAID5F: Revolutionizing Cloud Storage Performance and Reliability

Mariusz Barczak, Sarika Mehta
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.
- © Solidigm. Solidigm and the Solidigm logo are trademarks of Solidigm in the United States and other countries. Other names and brands may be claimed as the property of others.

Agenda

- **CSAL Feature Development**
- **CSAL Architecture Overview**
- **CSAL with Core Scaling for RAID5F**
- **CSAL Cache**

CSAL Feature Development

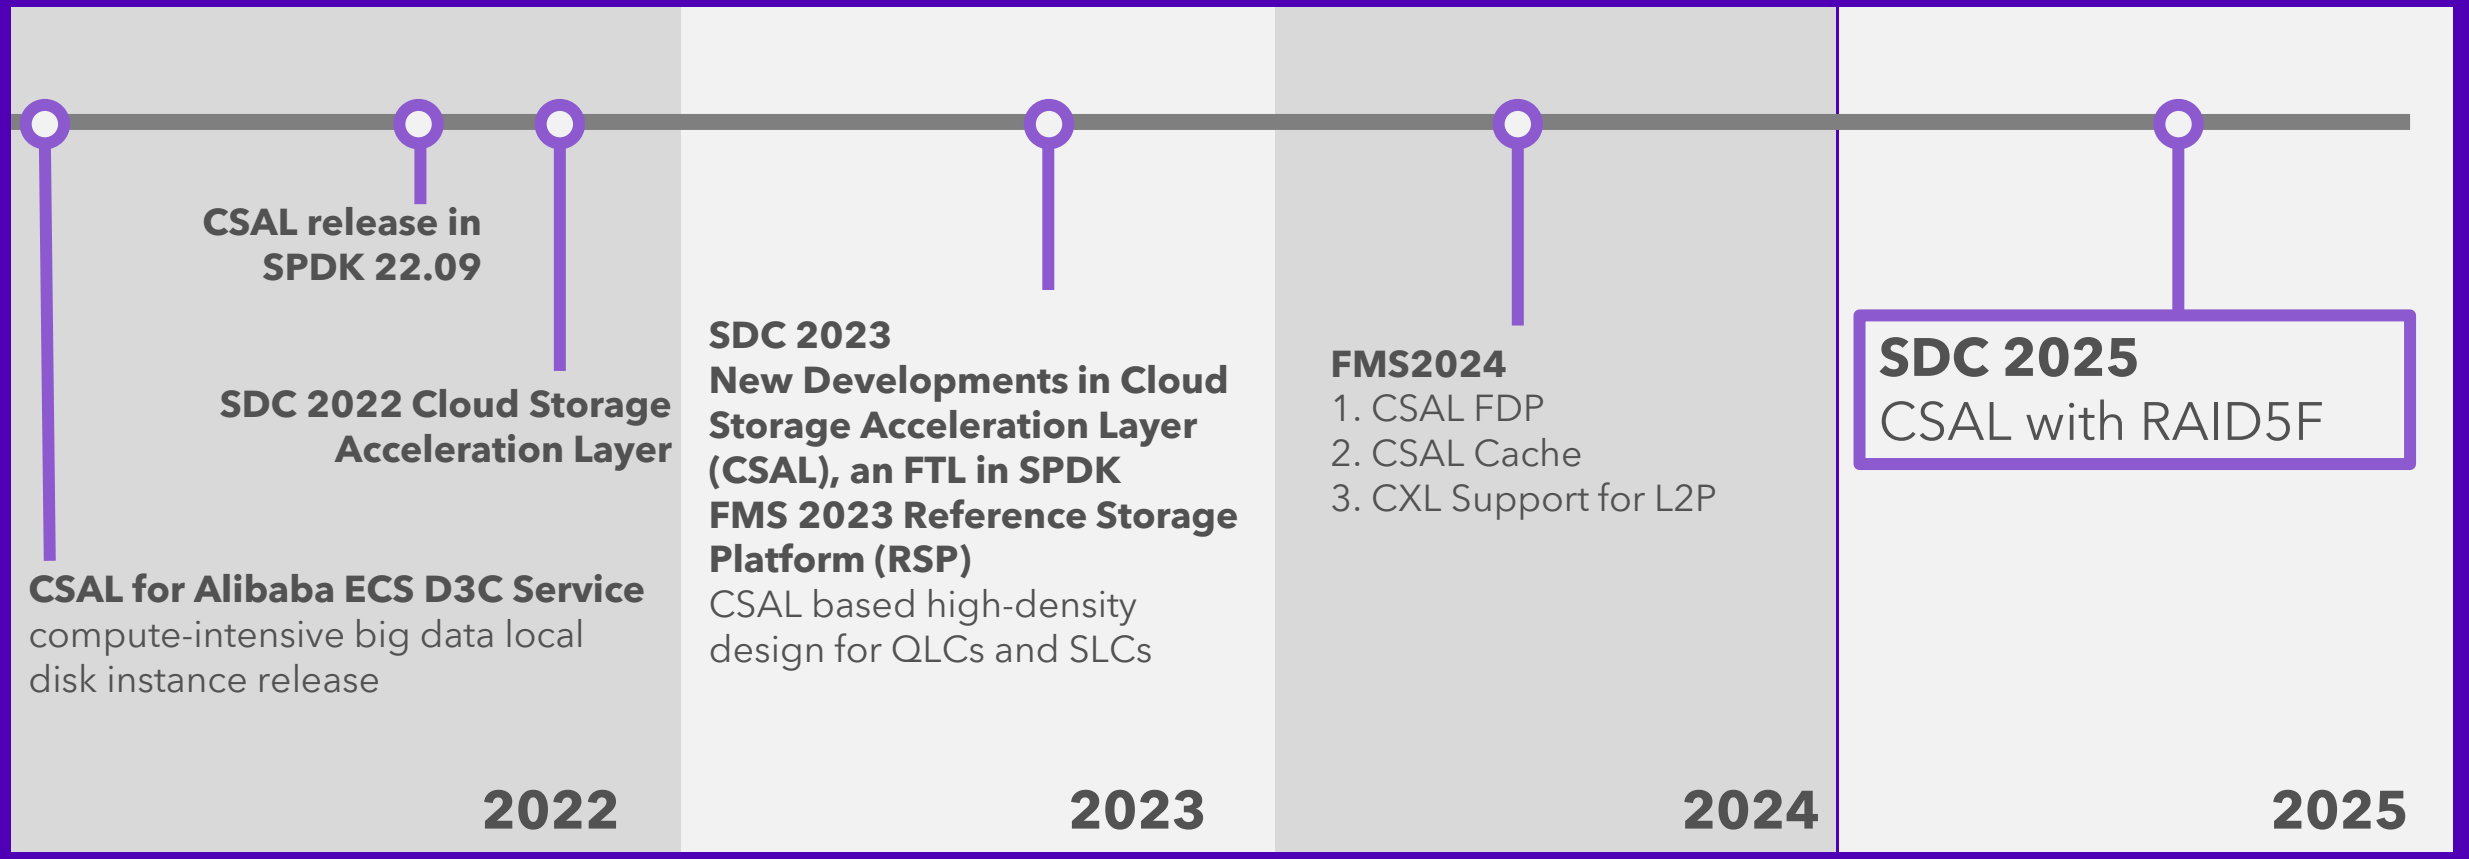

Cloud Storage Acceleration Layer (CSAL)

- **Built on SPDK**
  - Modular extensions enhancing SPDK's core
- **Write shaping**
- **CSAL Cache**
- **RAID0/1/10/5F**
- **Features in planning**
  - data placement, compression, deduplication, management, analytics

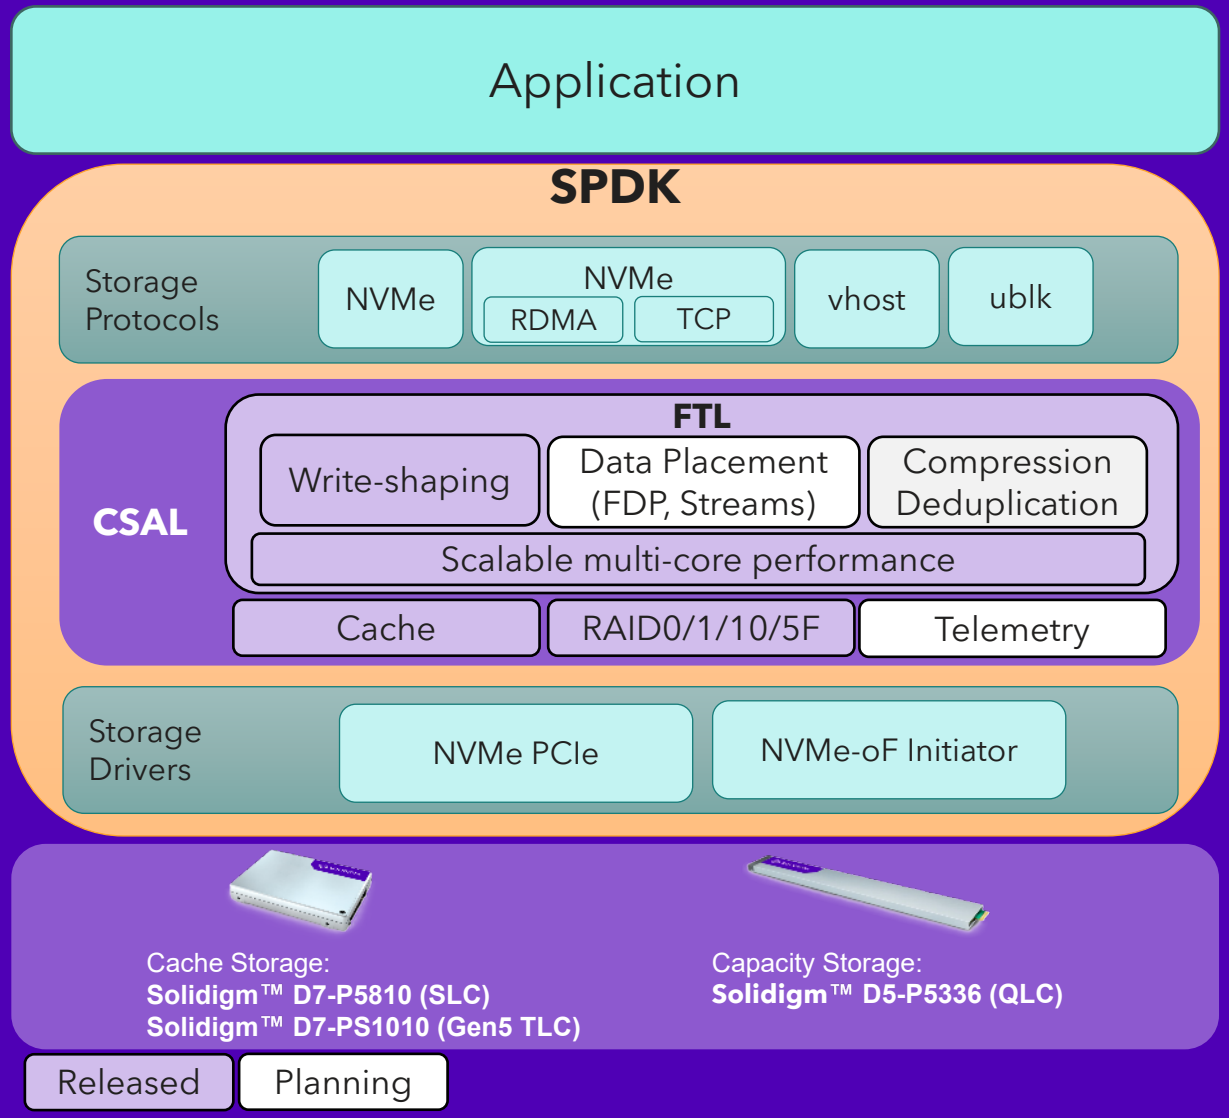

CSAL with Core Scaling for RAID5F with write-hole closure

Write Shaping

- 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)

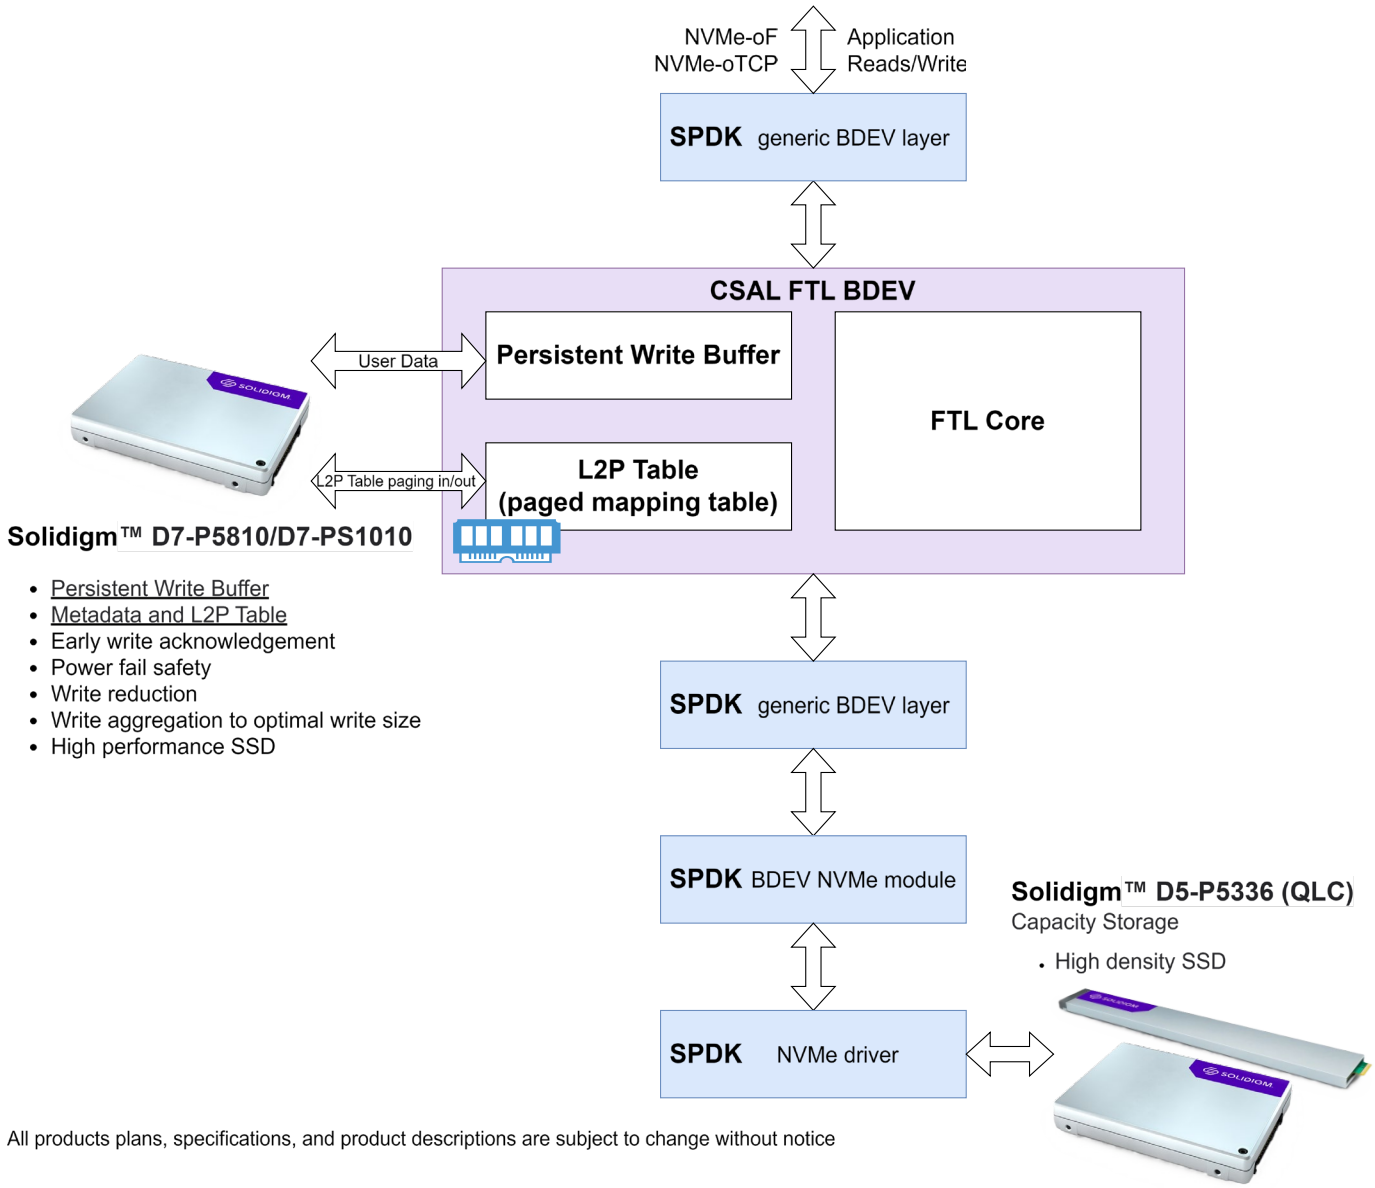

All products plans, specifications, and product descriptions are subject to change without notice

The Problems with Traditional RAID5

RAID 5 writes are often not aligned to a full stripe.

- Involves updating:
  - One or more data blocks
  - Parity block
- Requires read-modify-write cycle:
  - Read old data + parity
  - Compute new parity
  - Write new data + new parity
- Write-hole - Inconsistency risk between data and parity during update if power loss or crash occurs
  - Old parity with new data
  - New parity with old data

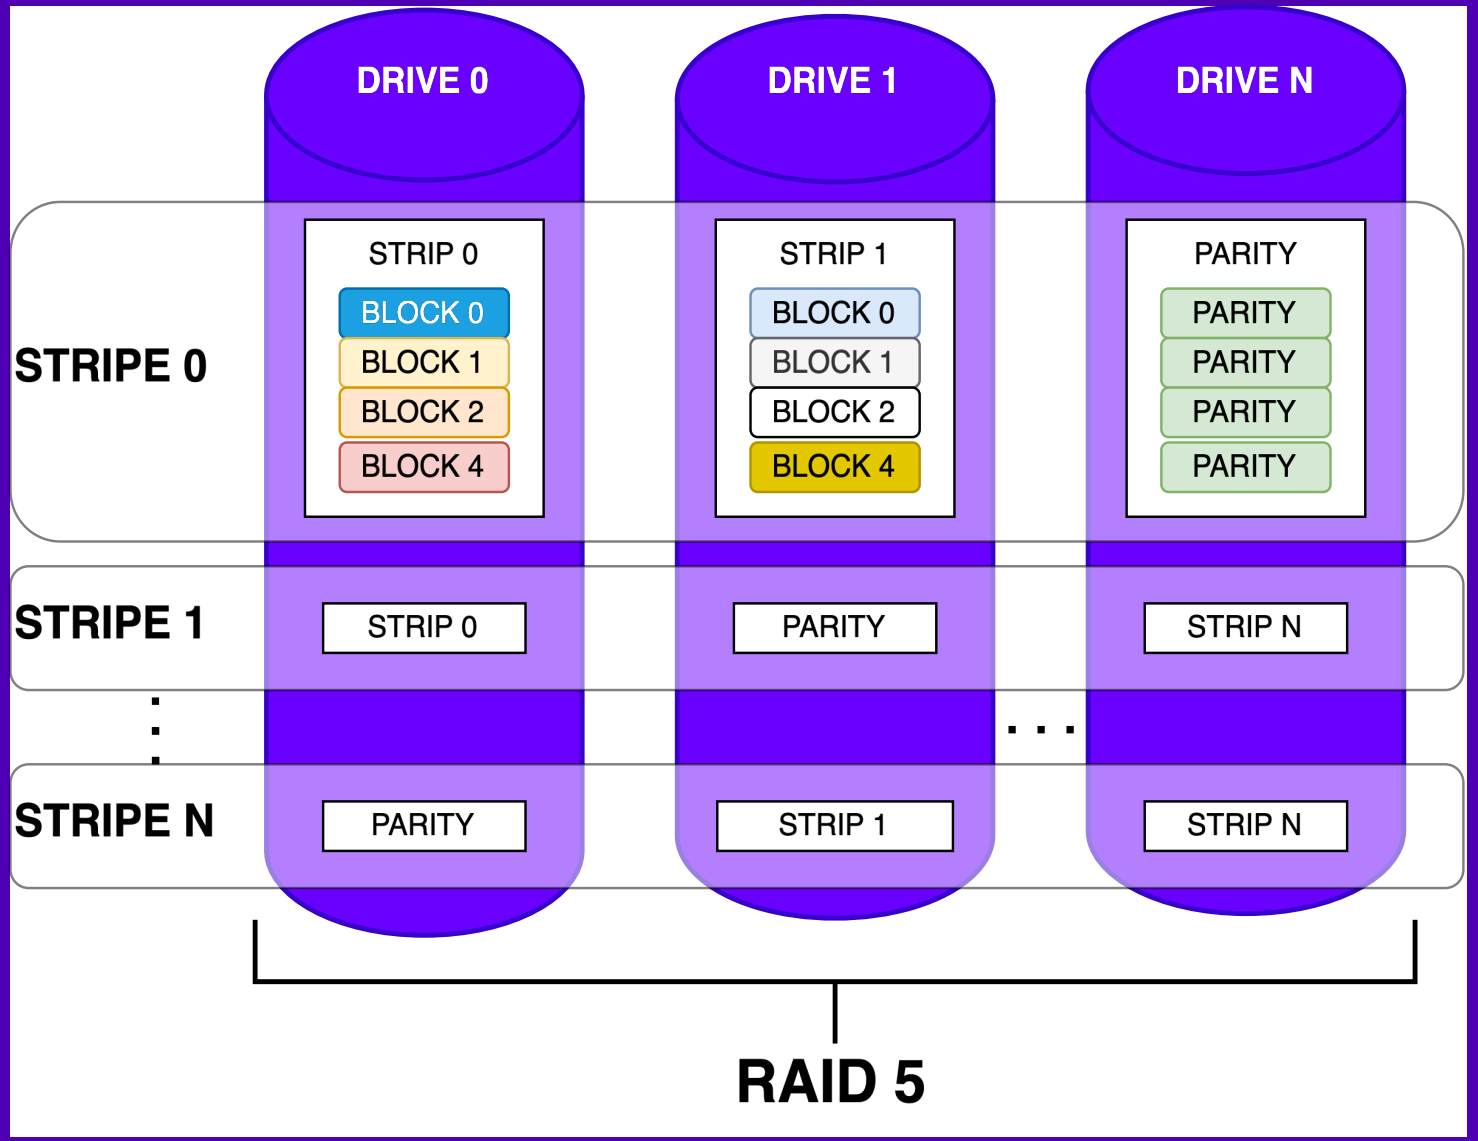

Introducing CSAL RAID5F

- Combining CSAL FTL write shaping future with new RAID5F BDEV
- Full-stripe atomic write
  - lack of costly read-modify-write
  - better performance
- No write hole
  - lack of silent data corruption
- Lower WAF
  - increased NAND endurance
- No RAID journaling required

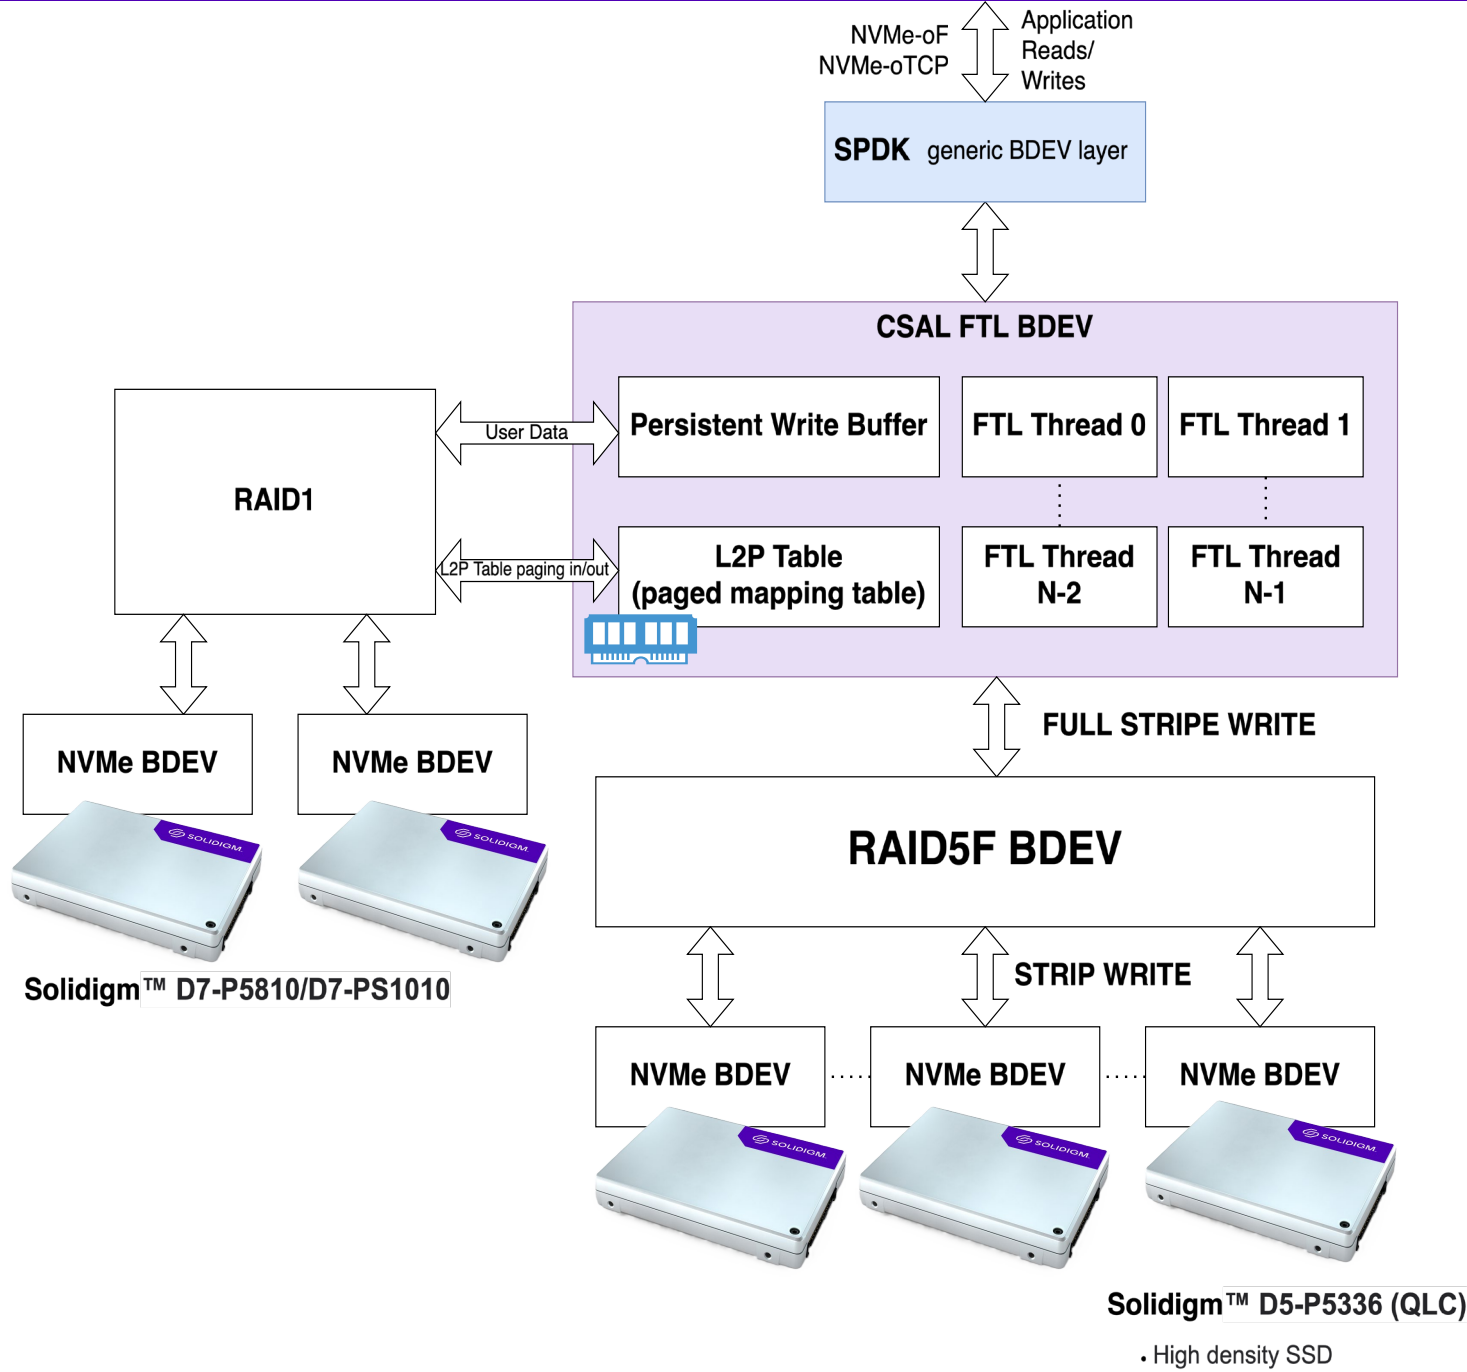

Evaluation - CSAL RAID5F vs Linux Software RAID (MDRAID)

- Workloads & Parameters:
  - Sequential Write:
    - Block size: 128K
    - I/O depth: 128
    - Jobs: 8
  - Random Write:
    - Block size: 4 KiB, 64 KiB
    - I/O depth: 128
    - Jobs: 8
- Preconditioning:
  - Full-drive sequential write pass across all partitions before test (to eliminate fresh-block bias)

- Test Environment:
  - Storage Backend (RAID5 with strip size 128K):
    - 9 × Solidigm D5-P5336 (QLC, 15.36 TB each)
    - Total usable capacity: ~123 TB
  - Write Cache (RAID10):
    - 8 × Solidigm D7-P5810 (SLC, 800GB each)
    - 5% of total usable capacity
  - Server: AMD EPYC 9534 (64 cores)
    - 8 cores dedicated for RAID logic
  - Operating System: Fedora 41
  - Kernel: 6.13.4-200.fc41.x86_64

CSAL RAID5F vs MDRAID Write Performance Comparison

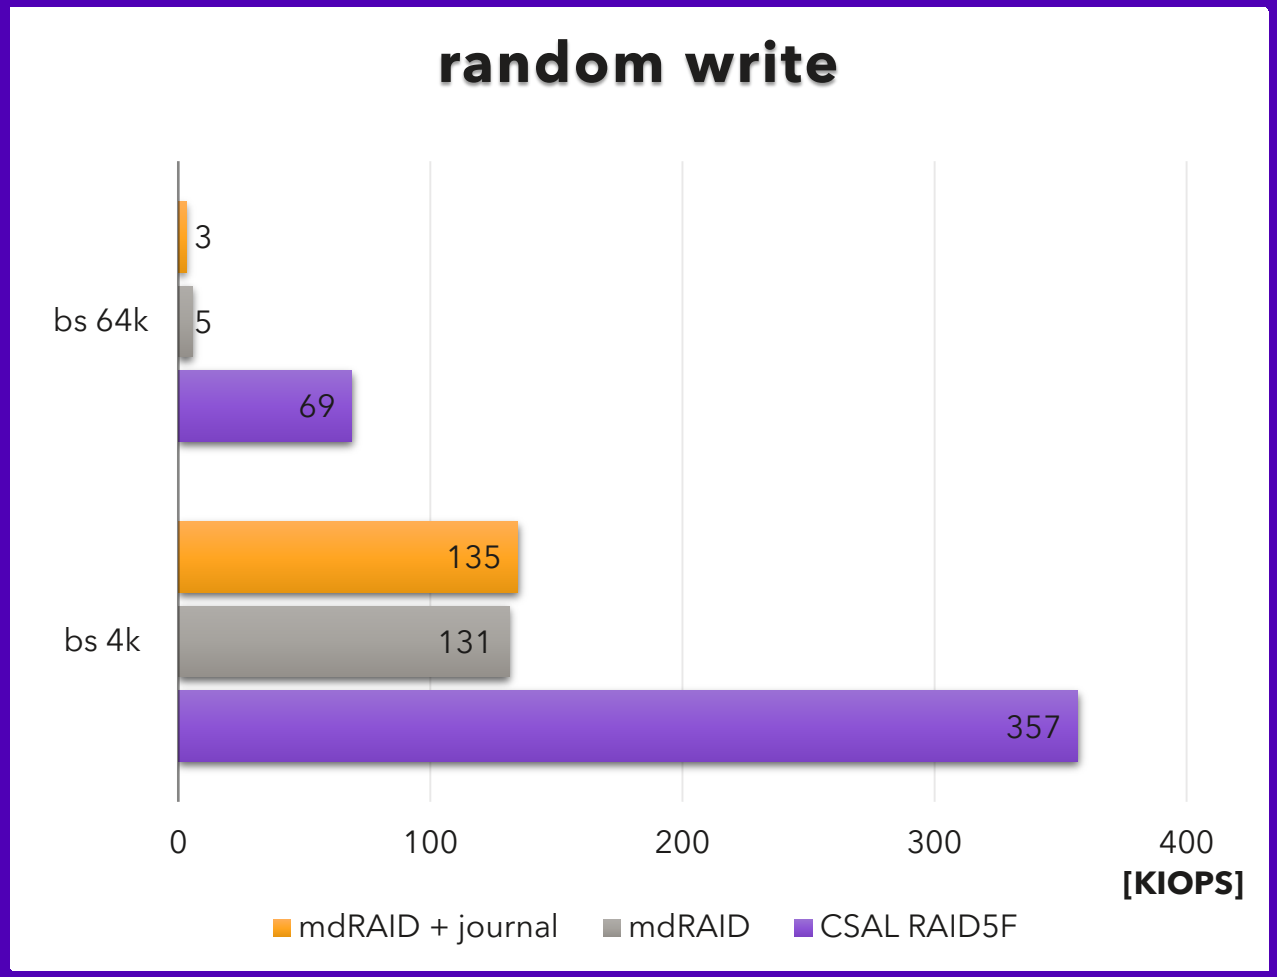

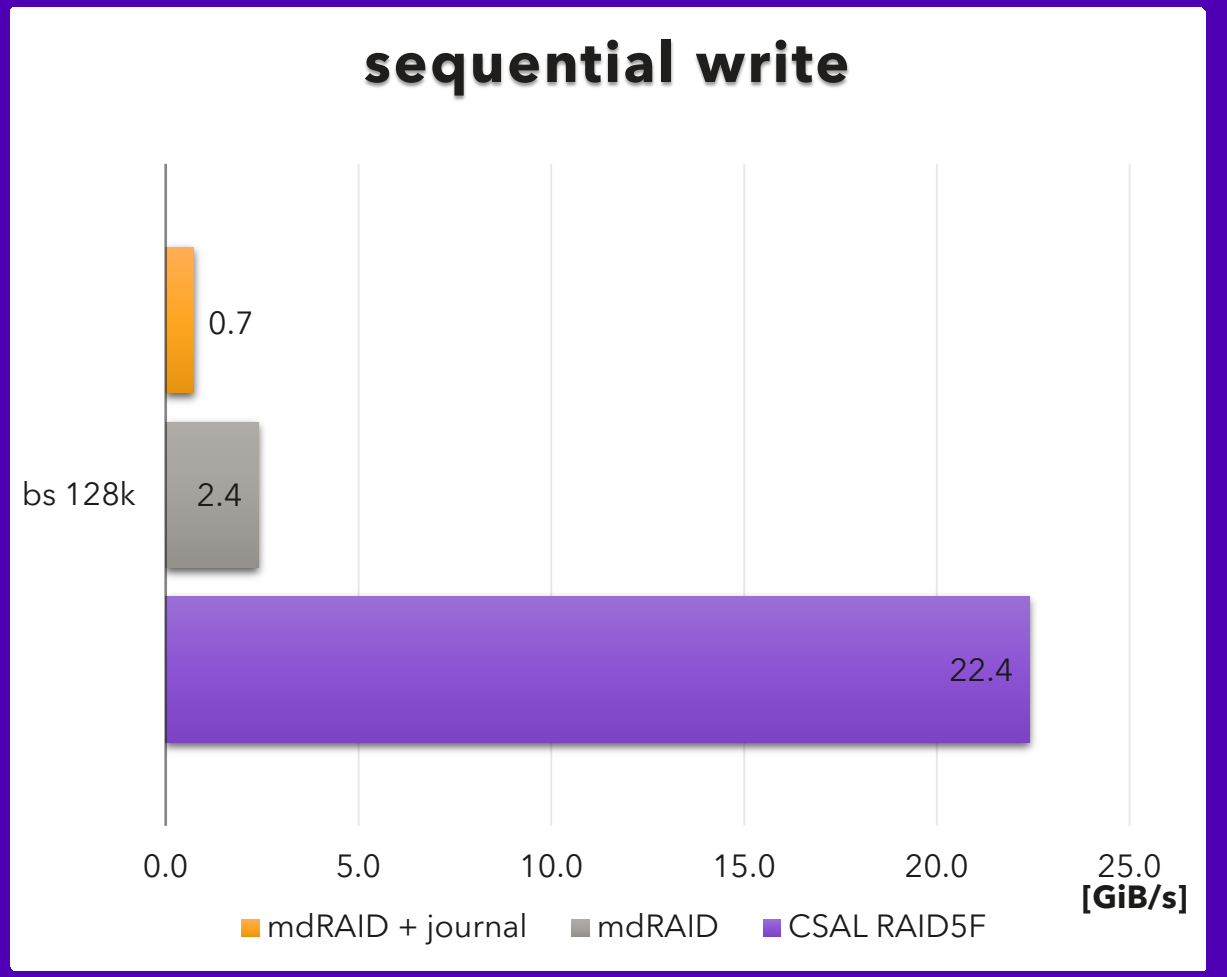

Unlock 4×–20× better performance with CSAL RAID5F, compared to traditional Linux kernel software RAID.

CSAL RAID5F Summary

CSAL offers a feature-rich, robust, QLC-friendly RAID solution – unlocking high-density, high-performance deployments.

Key Advantages:

- No RMW overhead
- Built-in write hole protection
- Scales across multiple CPU cores
- 4×–20× better performance vs MDRAID with journal
- Improved SSD endurance

CSAL Cache

Introducing CSAL Cache

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

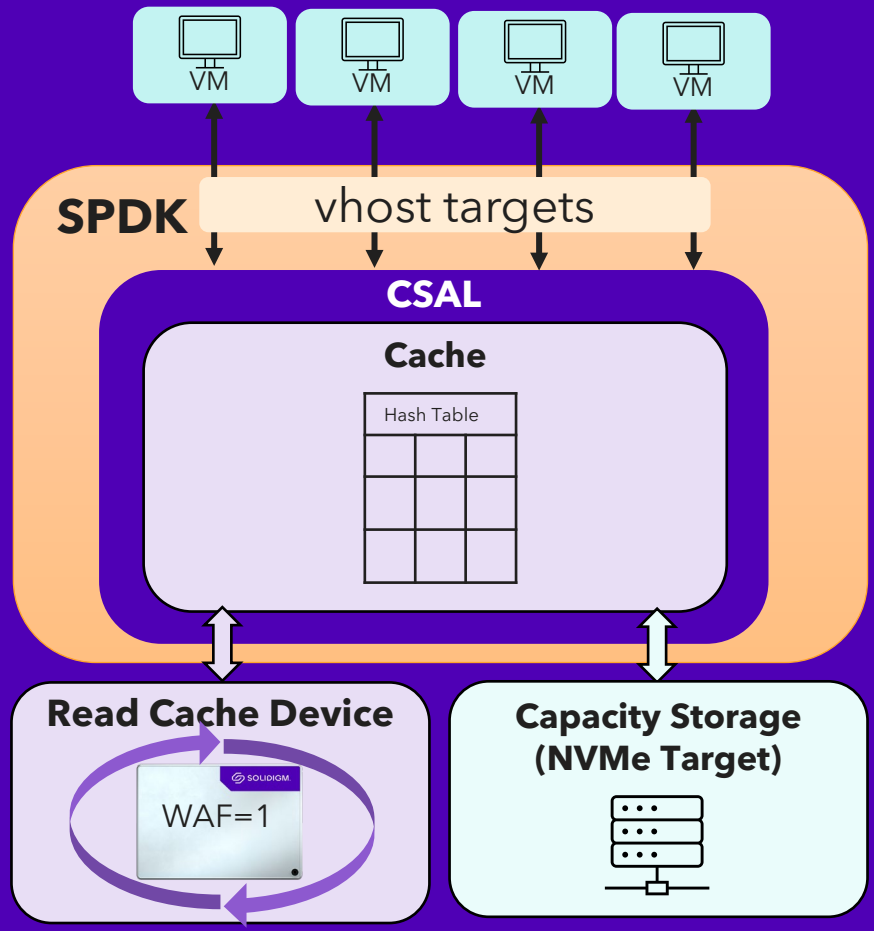

Addressing Technical Challenges of Traditional Cache

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 cache SSD wear-out
- **Efficient Mapping and Lookup:** One hash item for larger I/Os
- **Shared nothing architecture:**

  Faster cache lookups

A Legacy Cache Hash Table Lookup Simulation

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

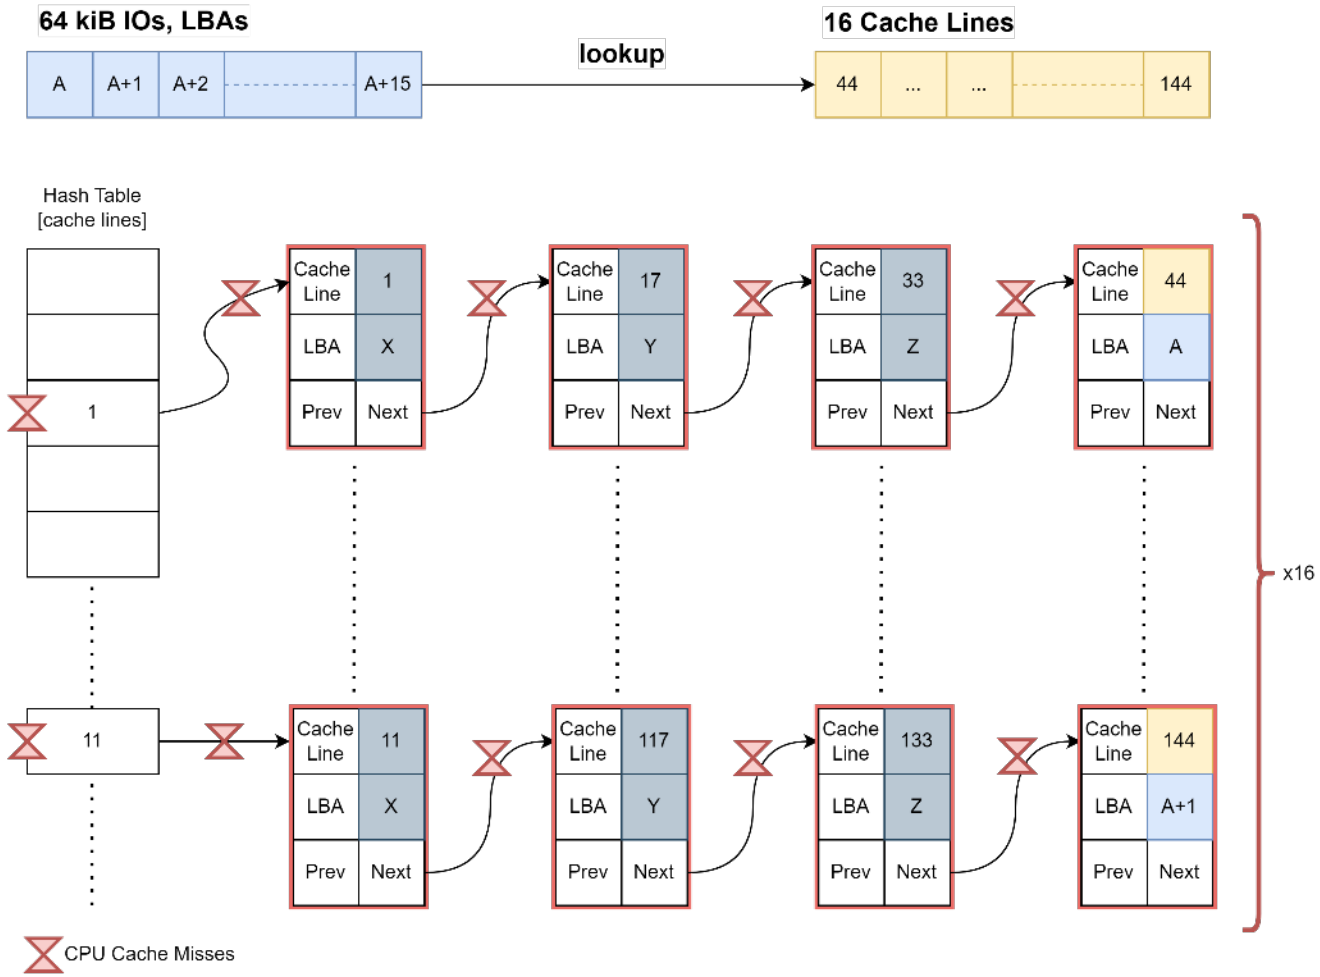

CSAL Hash Table Lookup

	Legacy Cache	**CSAL 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**
Metadata DRAM footprint for 4TiB cache	~84GB	**~2.5GB**

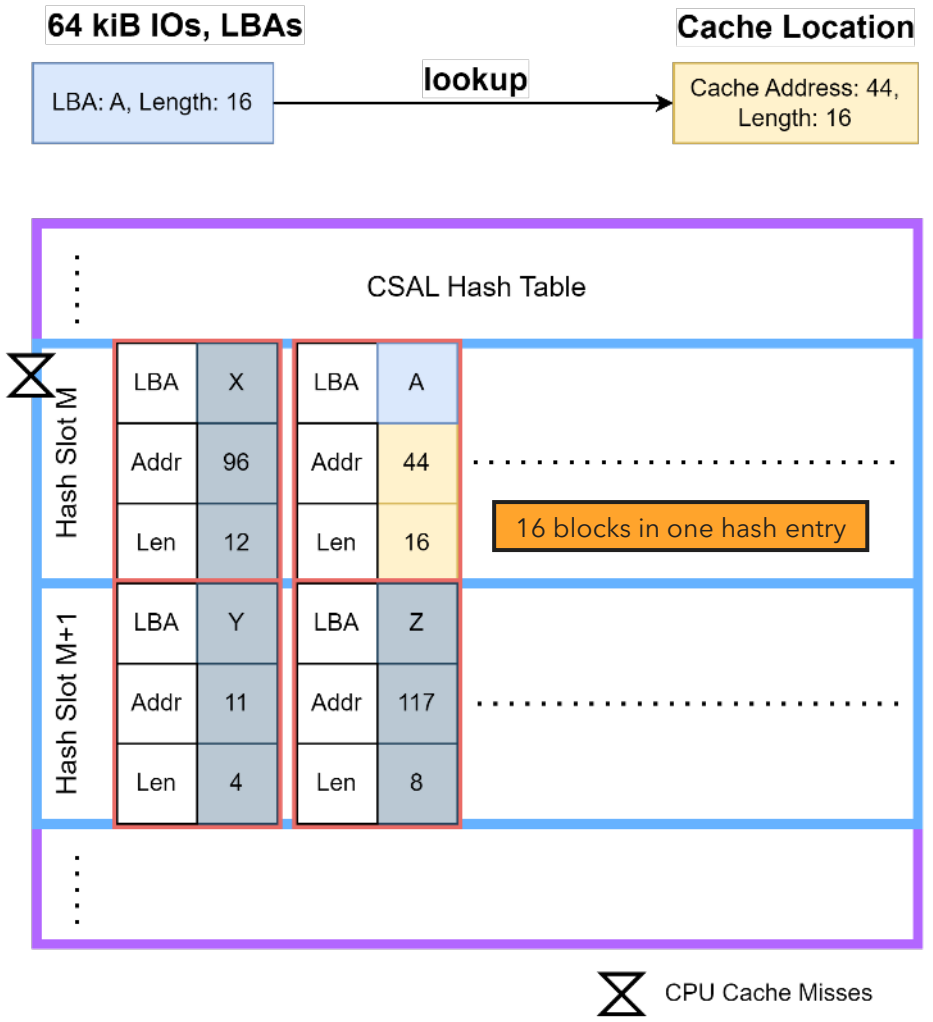

Comparing CSAL Cache with Legacy Cache

- Workload and setup:
  - Random Read
    - 8 jobs
    - zipfian 0.8 distribution – good cache locality
    - 128 IO depth
    - 64K block size
- Precondition:
  - Fill all partitions executing sequential write

- Test setup and config:
  - Software:
    - CSAL SPDK
    - Open-CAS Linux
  - Storage backend device:
    - Solidigm™ D5-P5336 61TB QLC
  - Cache Device:
    - Solidigm™ D7-PS1010 3.8T TLC
  - Server:
    - AMD EPYC™ 9354 32-Core
    - 1 core for CSAL
  - Operating system: Fedora 41
  - kernel version: 6.16.3-100.fc41.x86_64

Performance Results Comparison

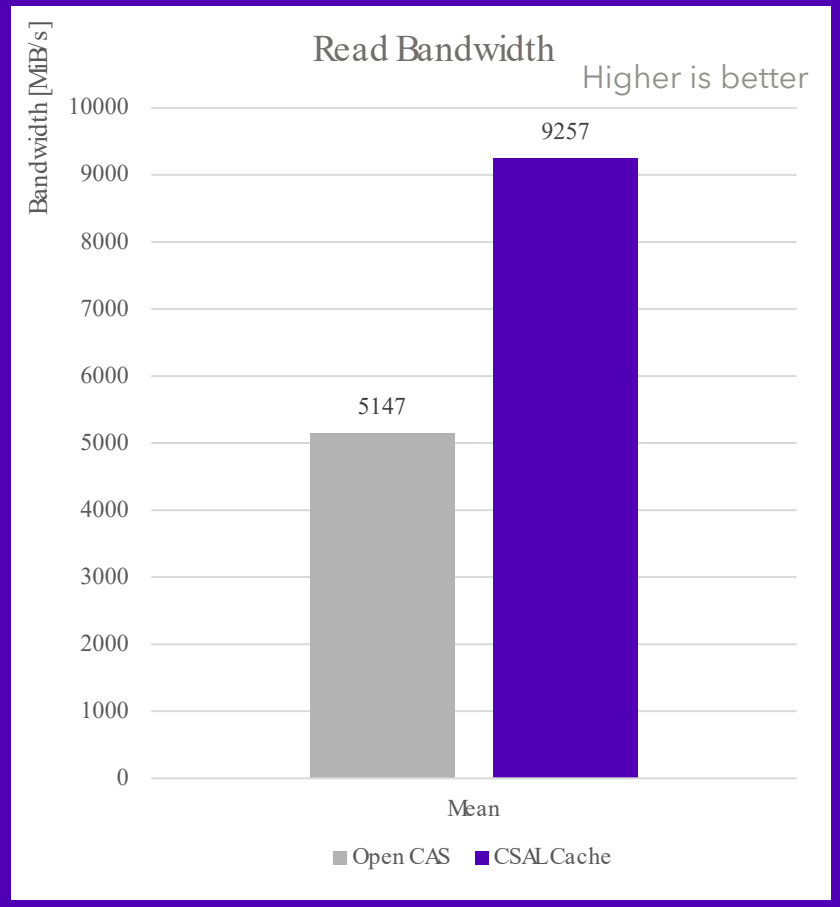

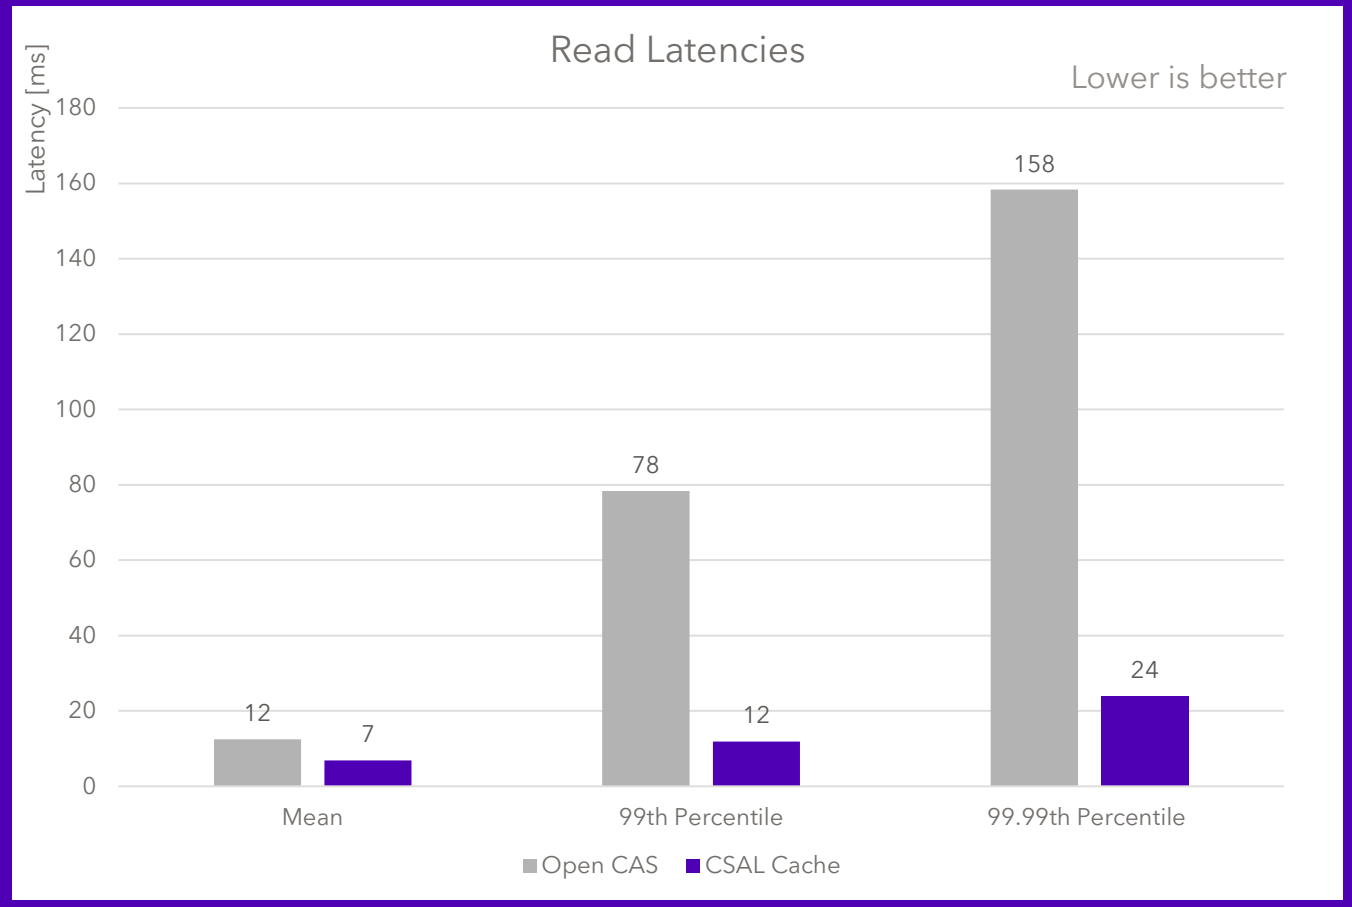

- ~80% Higher Read Bandwidth

- ~71% improved average latency
- ~5.5x better Two 9s and Four 9s latency

CSAL Cache Summary

CSAL Cache delivers a high-performance cache solution with enhanced cache SSD endurance.

Key Advantages

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

Conclusion

Flexible, Scalable and Reliable High-Capacity Storage Solution

- Solves Key Storage Bottlenecks
  - Endurance Management
  - Persistent Write Cache
  - High-performance Read Cache
  - RAID5 with Write Hole Closure
- Optimized for High-Density QLC Storage
- Delivers Consistent Performance
- Manages Implementation Complexity
- Tuneable Design to Meet Customer Needs

SOLIDIGM

Q&A
Cloud Storage Acceleration Layer

Contact us:
dl_csal@solidigm.com

SDC 25

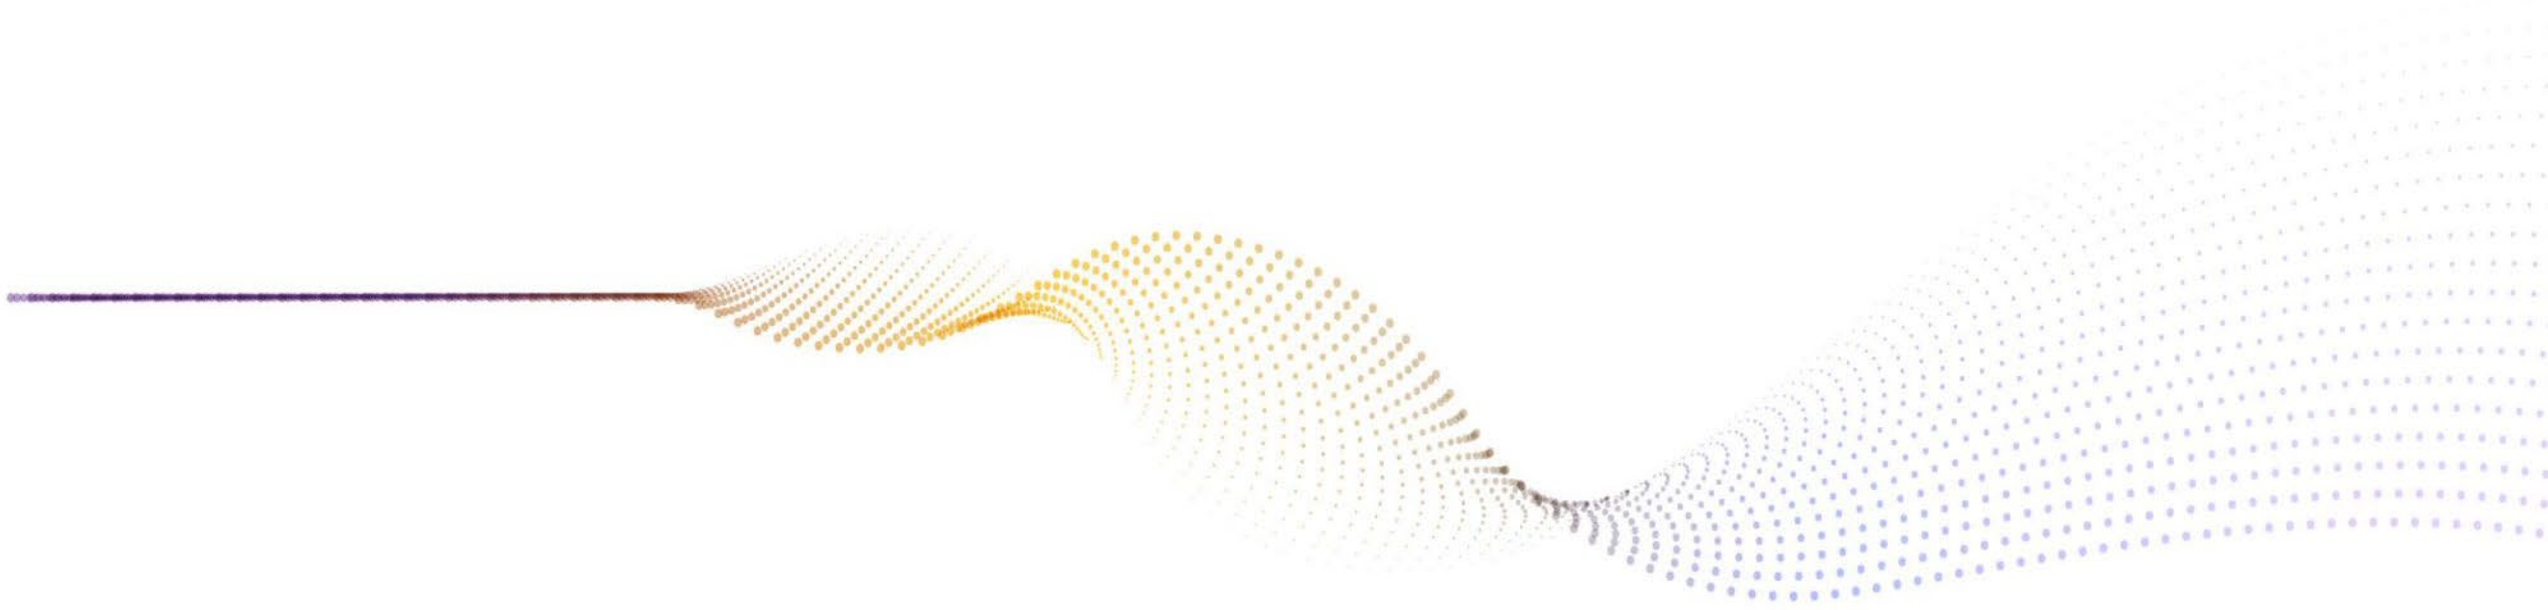

Thank you for attending!

Please remember to rate this session. You get access the presentations at
http://sniadeveloper.org/conference

