FMS
the Future of Memory and Storage

SOLIDIGM

QLC in the Datacenter: Making the Optimal Choices

FMS

Aug. 08, 2024

Agenda

- Data Growth Trend
- Challenges facing Data Infrastructure
- General Workload Types within the Cloud and Enterprise space
- Product vs. Targeted Applications
- General Cost Analysis
- Solidigm QLC Portfolio

Storage Needs Are Not Abating

147ZB created in '23, growing to 181ZB in '25[1]

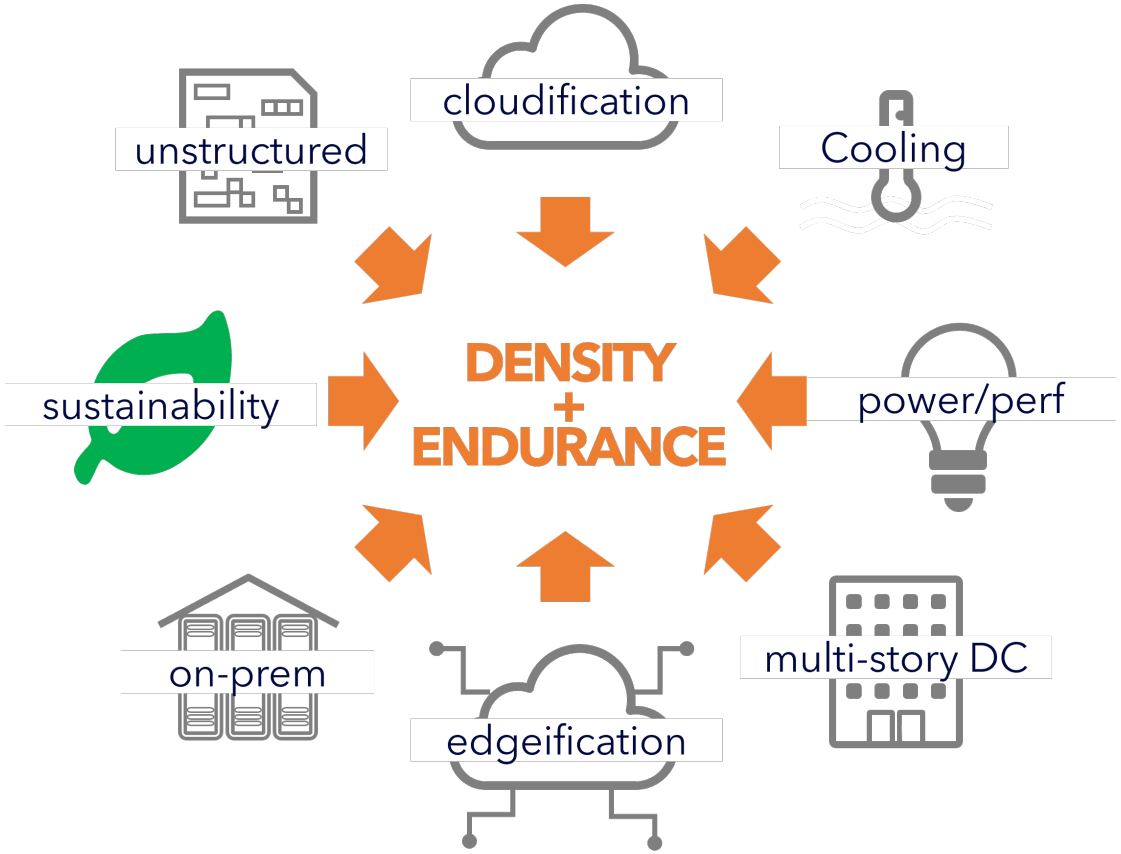

1. Source: https://explodingtopics.com/blog/data-generated-per-day

Source: https://www.domo.com/learn/infographic/data-never-sleeps-11

Storage Challenges are Intensifying

the Future of Memory and Storage

	Data Center	Edge Infrastructure	Edge
	Public Cloud Enterprise Data Center Regional Data Center	Regional data centers Micro data centers Access edge	On-prem servers Devices / IoT Laptops and PCs
Size and Weight			
Power	12000	400	
Serviceability			
Operational Efficiency			
Uptime and Reliability			

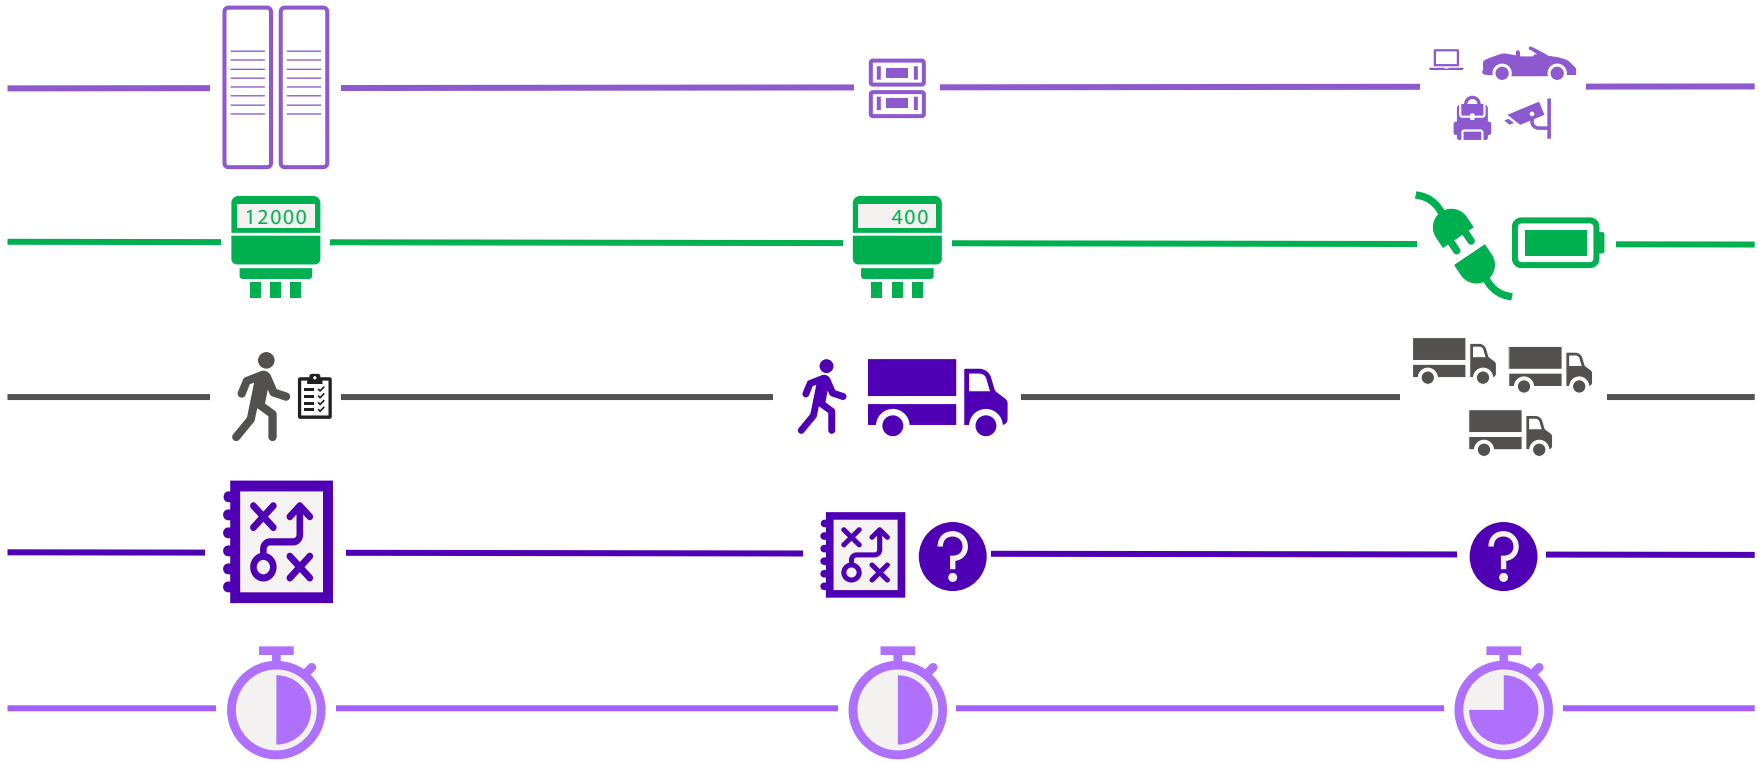

Source: https://www.domo.com/learn/infographic/data-never-sleeps-11

QLC Adoption Growth and Acceptance

Dell Technologies Bolsters Dell PowerStore with Storage Performance, Resiliency and Efficiency Advancements[1]

- "….QLC-based storage: Delivers enterprise-class performance at a lower cost per terabyte when compared to triple-level cell (TLC) models. Customers can start with as few as 11 QLC drives and scale up to 5.9 petabytes of effective capacity per appliance…"

VAST Data Leverages the Value of Solidigm SSDs to Redefine Storage[2]

QLC today leveraged on AFF C-SERIES

"Today's AI workloads are no longer latency sensitive they are bandwidth sensitive. Solidigm QLC SSDs deliver the performance, capacity and endurance needed to scale beyond for today's most data intensive AI usages."

~Ed Fiore, Vice President & Chief Architect

Data Direct Networks (DDN) has about 48 A1400x2 arrays supporting Nvidia's largest SuperPODs. Adding Solidigm QLC delivers both the best performance for AI usages, as well as extreme capacity. It uses 61.44 TB QLC SSDs enabling 1.45 TB capacity in a 2RU x 24-slot chassis, doubling capacity per watt compared to other 30 TB SSDs from other suppliers."

~Senior Vice President of Product, James Coomer

1. https://www.dell.com/en-us/dt/corporate/newsroom/announcements/detailpage.press-releases~usa~2024~05~20240521-dell-technologies-bolsters-dell-powerstore-with-storage-performance-resiliency-and-efficiency-advancements.htm#/filter-on/Country:en-us
2. https://www.solidigm.com/content/dam/solidigm/en/site/products/technology/vast-data-customer-story/documents/vast-data-solidigm-case-study.pdf
https://www.netapp.com/blog/qlc-all-flash-arrays-data-center/
3. https://www.netapp.com/blog/qlc-all-flash-arrays-data-center/

Characterizing QLC against data patterns and block sizing

QLC selection must include a review of the combination of capacity, workload pattern, duty cycle and performance needs

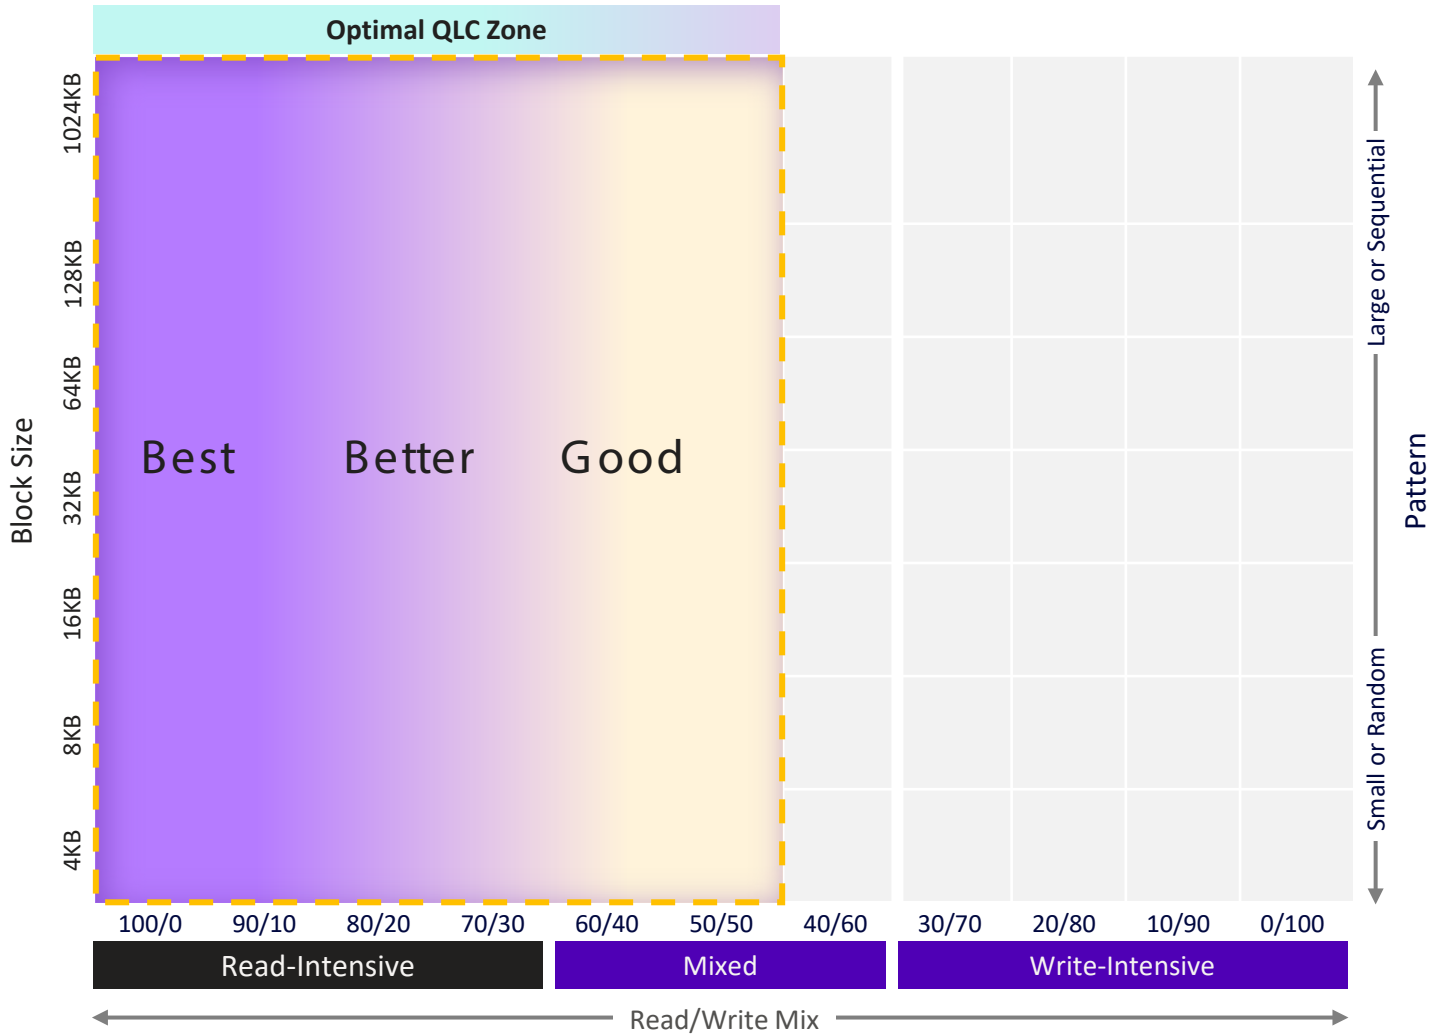

Elements of QLC: Indirection Unit (IU) options for broader deployment

Solidigm D5-P5430 — 4KB IU

All writes go directly to drive

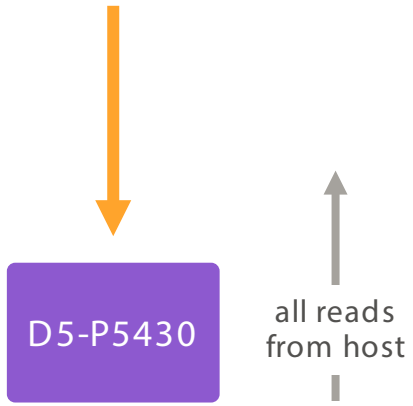

Drop-in value TLC replacement

Solidigm D5-P5336 — 16KB IU

≥16KB writes go directly to drive

OR

non 16KB-aligned write options

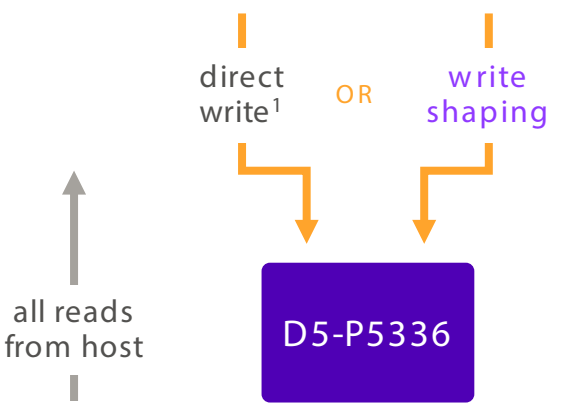

Improved ease-of-adoption[2]

1. Direct write is optional for non-16KB-aligned writes. For further guidance please evaluate possible performance and endurance impacts using the endurance estimator tool: https://estimator.solidigm.com/ssdendurance/
2. Comparing Indirection Unit (IU) size: D5-P5336 IU=16KB.

Analysis of Workloads- Areas for QLC

Type	Workload/App[1]	Average R/W and Pattern Mix	Average Block Size	Avg. Queue Depth	QLC Leverage?
Enterprise Based	HPC	70/30 to 90/10; Seq Read (data load) ;Seq Write (de-stage)	16KB and higher	0-32	D5-P5430 or D5-P5336
	General Purpose	80/20; 80% sequential; 20% random	8-32KB	0-32	D5-P5430
	Database	70/30; primarily sequential	64KB-1MB	Varies	D5-P5430
	Decision Support System (DSS)	90/10 to 100/0; 80% sequential, 20% random	Large Block (up to 1MB)	0-4	D5-P5336
	Cloud Compute	65/35; ~80% Random	8KB	0-32	With Write Shaping
Cloud Based	Content Delivery Network (CDN)	Up to 95% Read; Random Reads; Seq Writes when present (writes limited to low use periods)	Very Large Block (majority ≥128KB)	0-64	D5-P5336
	Social	68.7/31.3; Random	87.6KB - object 11.3KB - metadata	0-32	D5-P5430 / D5-P5336
	eCommerce	67/33; Primarily Random (driven by high concurrency)	Small Block (2-8KB or 16KB)	0-32	With Write Shaping
	Server-based Storage	70/30; Random	4KB	0-16	With Write Shaping
AI Specific (Core or Near Edge)	Ingest	100% writes; Sequential, though multi-streams present as random	Varies; Block sizes aligned to framework	32-128	Suboptimal
	Training	95/05; Random	Block sizes aligned to frameworks	0-64	Suboptimal
	Checkpointing	~100% Writes; Mostly sequential	32-128KB	0-64	D5-P5430 or D5-P5336
	Inference	Read intensive, writes during model feedback; Mostly Sequential	4-128KB	0-64	D5-P5430 or D5-P5336
	Data Lakes	Mixed from 100% Read to 100% write Variable random to sequential	Variable	0-32	Hybrid: SLC cache+ D5-P5336

1 Source: Solidigm, based on a combination of Intel analysis and publicly available storage workload research material.

Dedicated Swim Lanes Target Applications

Endurance Swimlanes[1]		Target Applications and Usages Examples			
>140 PBW (Ran. or Seq.) 50 DWPD	D7-P5810 Very Write-centric	❑Caching ❑AI Training		❑HPC	
65/134 PBW (Ran./Seq.) 3 DWPD	D7-P5620/PS1030 Write-centric and mixed	❑Caching ❑High Freq. Trading	❑Cloud Compute	❑OLTP (small block, high duty cycle)	
28/134 PBW (Ran./Seq.) 1 DWPD	D7-P5520/PE1010 Mixed and mainstream	❑AI ingest ❑AI preparation ❑Database	❑Decision Support Systems	❑eCommerce ❑OLTP (variable block, low duty cycle)	❑Cloud Storage ❑Data Analytics ❑Email and UCC
32/105 PBW (Ran./Seq.) 0.5+ DWPD	D5-P5430 Mainstream and read-intensive		❑OLAP		❑General Purpose Server ❑Server-based Store ❑VDI
65/213 PBW (Ran./Seq.) 0.2+ DWPD	D5-P5336 Capacity optimized	❑Content Delivery Network	❑Adv. Driver Asst. Sys. ❑AI/ML data pipelines ❑Ceph ❑Object Storage		

1. All DWPD values assumes a 100% random write workload

QLC and TCO Considerations

- Operational Analysis of Data Center system costs must drive spending decisions

Inputs[1]:

Price/GB
Duty cycle
Capacity utilization %
Error Encoding
Data Reduction
B/W
AFR
Power

Outputs:

TB/W
Number of drives
Number of servers
Number of racks
Power Cost
Total Cost/Yr

1. Key inputs listed, but not exhaustive.

Higher Density Delivers an Improved Sustainability Footprint

Solidigm D5-P5336 vs all-HDD in a 10 Petabyte Object Storage Solution

Solution-level Physical Rack Footprint

all-HDDs
24x 24TB OP'd to
90% usable capacity

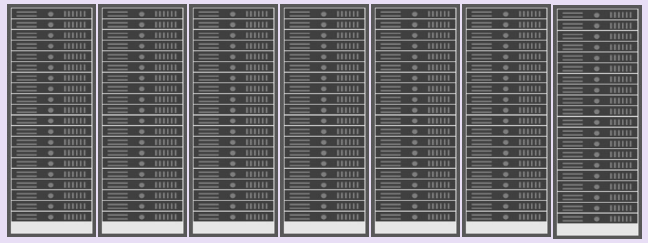

7 Racks

Solidigm
D5-P5336
24x U.2 61.44TB

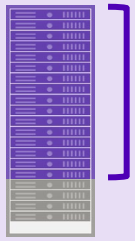

16x 2U in
1 Rack

Solution-level Energy Usage Footprint

all-HDDs
24x 24TB OP'd to
90% usable capacity

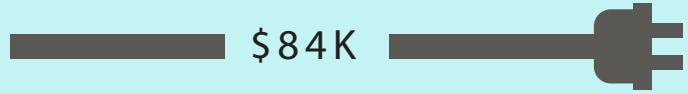

$84K

Solidigm
D5-P5336
24x U.2 61.44TB

$22K

74% TCO improvement

Massive density delivers meaningful sustainability benefits

Expanded Higher Density QLC Portfolio

Summary

- Storage challenges are driving a need for high throughput, high density, low power solutions (i.e., TB/w vs. BW/w)
- Storage matters – more data, higher XPU utilization, server TCO optimization is critical
- Not all QLC storage is equally up to the task depending on application
- Solidigm has the portfolio and partnerships to help!

Appendix 2: Solidigm D5-P5336 TCO Calculations

SOLIDIGM D5-P5336 VALUE vs. 24TB HDD

Storage Solution deployment parameters assume 42U rack capacity, 33U available for storage, 2U servers @ 12x 3.5" HDDs, and 24x P5336 SSDs per server, respectively. 90% usable capacity assumed for both drives. HDD refresh cycle = 4 years. Calculated duty cycles to deliver equivalent throughput per TB: 20% for HDD array, 4.4% for all-QLC solution. RAID 1 mirroring used for QLC; all-HDD uses Hadoop triplication.
All-QLC configuration:
Solidigm D5-P5336, 61.44TB, 90% capacity utilization, 7000 MB/s throughput, 24W average active power, 5W idle power; See details at Soldigm D5-P5336 QLC SSD (solidigm.com)
HDD configuration:
Seagate EXOS X24 24TB 3.5" SAS HDD ST120000NM002D (datasheet), 90% capacity utilization, 'short-stroked' throughput estimated at 400 MB/s; 9.8W average active power, 6.5W idle power
Key common assumptions : Power Cost = $0.15/KWHr; PUE factor = 1.60; Empty Rack Purchase Cost = $1,200; System Cost = $10,000; Rack Cost for Deployment Term = $171,200. Calculations based on Solidigm TCO estimations as of July 2024. TCO calculations based on internal Solidigm TCO tool.