

New Age SSDs - Need for differentiations for emerging use cases

Tahmid Rahman

Director Product and Partner Marketing



the Future of Memory and Storage

SSD segment analysis



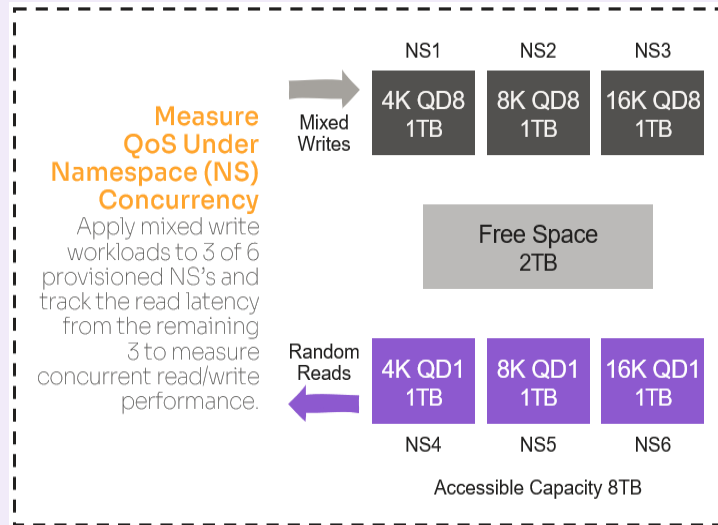
	Segments	Performance	Endurance	Workload	Best Fit	Form Factor	'26 TAM
SLC	Premium Compute	Very Write-centric	50 DWPD	HPC, Caching, Hi Freq. Trading, AI Compute	D7-P5810	U.2	10%
			3 DWPD		D7-PS1030 D7-P5620	U.2, E3.S U.2, E1.S, E3.S	
TLC	Volume Compute	Write-centric and mixed	1 DWPD	General Purpose, Emails, Decision Support	D7-PS1010 D7-P5520	U.2, E3.S U.2, E1.S, E3.S	35%
	Hot Storage	Mixed and Mainstream	1 DWPD	OLTP, Cloud Storage, Data Analytics, AI/ML	D7-PS1010 D7-P5520	U.2, E3.S U.2, E1.S, E3.S	30%
QLC	Warm Storage	Mainstream and Read-intensive	Beyond DWPD spec - TBW	CDN, Obj Store, OLAP, AI Ingest, Archival, ADAS	D5-P5336, D5-P5430	U.2 E3.X E1.L	20%
	Cold & Tepid Storage	Read Data-intensive	Very Low Write Endurance <0.1 DWPD	Archival	High Density D5-P5336	U.2 E3.X E1.L	5%

Source: Solidigm Market Research


Key Perf. Considerations (Noisy Neighbor)

Solidigm™ D7-PS1010 7.68TB vs in-Market PCIe 5.0 7.68TB SSD

FIO Noisy Neighbor Test Overview



Random Read Avg. Latency (µs)
Lower is Better



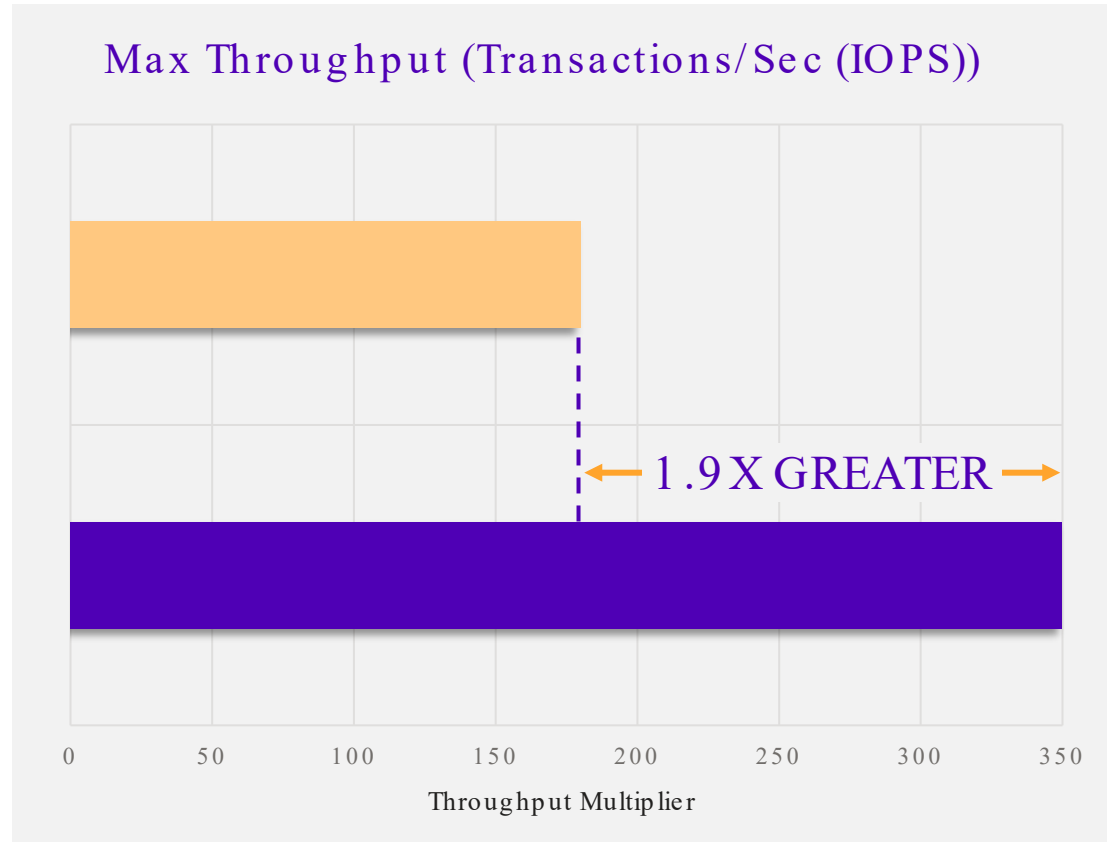

■ Solidigm D7-PS1010 ■ Solidigm D7-P5520 (Gen 4) ■ Competitor 2, Prod #2

Latency response under mixed use case is key



Key Perf. Considerations (Scalability w/ SLA)

Solidigm™ D7-PS1010 7.68TB vs in-Market PCIe 5.0 7.68TB SSD running Aerospike Certification Tool (ACT)



■ Solidigm D7-PS1010

■ Samsung PM1743

Need to meet latency SLA...

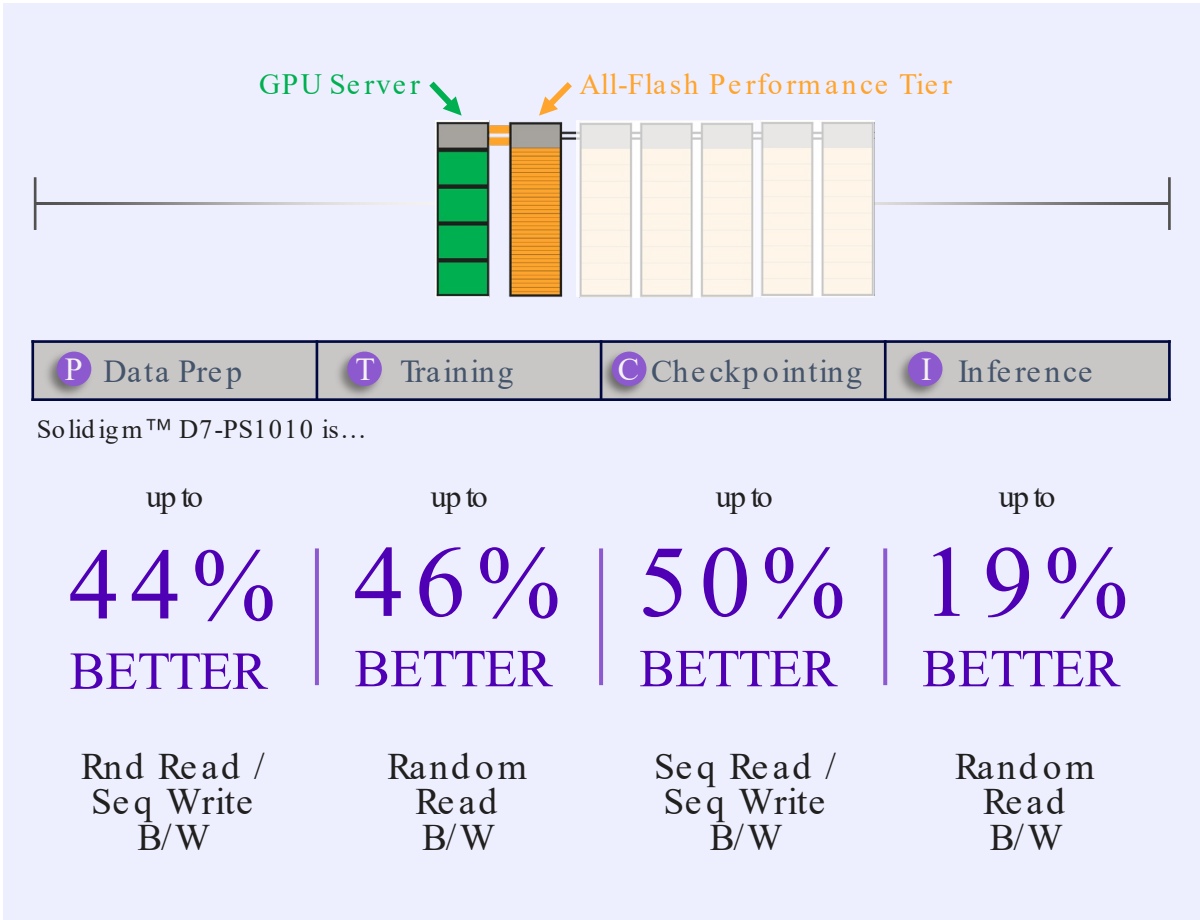
...while delivering higher throughput

¹ Source – Solidigm. See appendix 'Real Time Big Data' for details

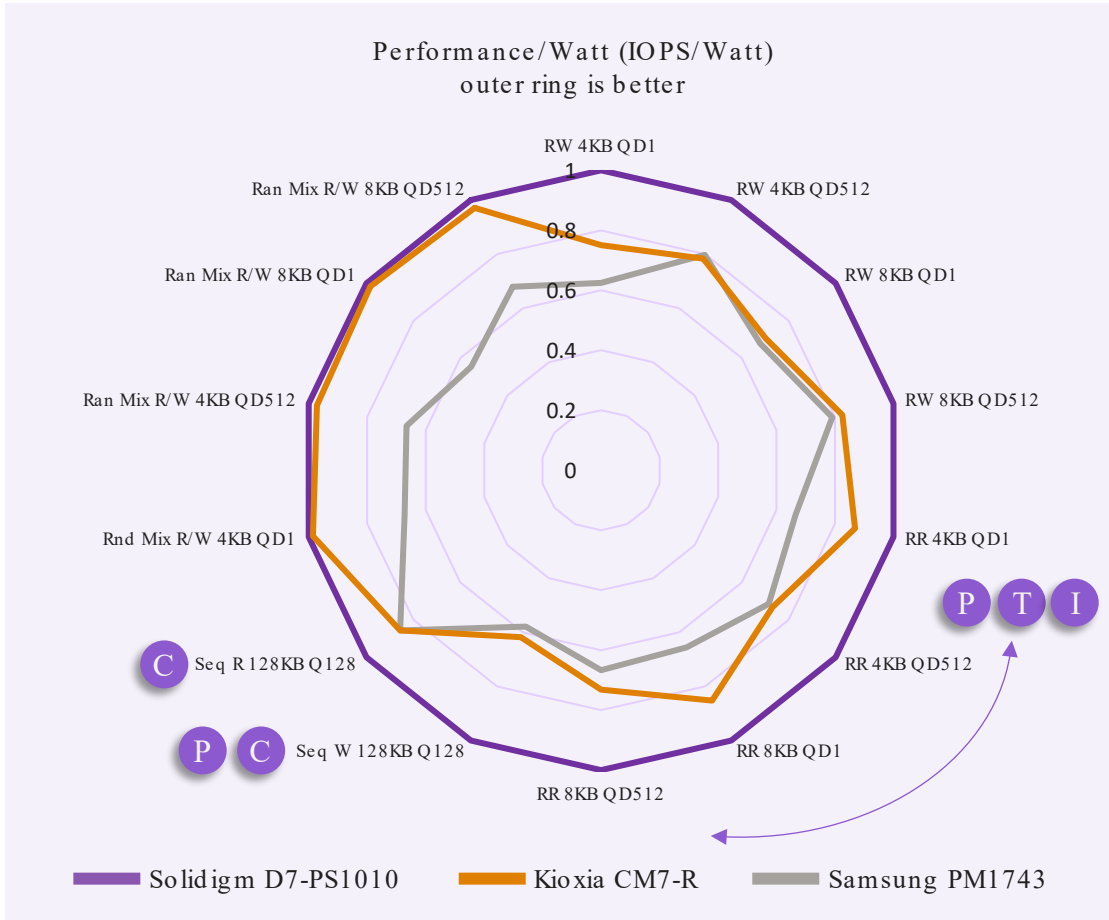


Performance + power efficiency –key for AI

Class-leading AI Performance




Class-leading Power Efficiency for AI



¹ Source – Solidigm. See appendix 'Efficient performance for energy-sensitive workloads' for details



What EDSFF best suited for high density? –beyond U.2




Form Factor	Density/unit	Cost Overhead	Drives/2U	TB/2U	Thermal power	Adoption trend
E3.S 1T	1.0x	Baseline	44-48	1.0x	25W	++++
E3.L 1T	1.5x	++	44-48	1.5x	40W	++
E3.S 2T	2.0x	+++	22-24	1.0x	40W	+
E3.L 2T	3.0x	++++	22-24	1.5x	70W	+
E1.L	1.5x	++	64	2.0x	25W	+++

- ❑ E3.S
 - 76mm x 112.75mm x 7.5mm
 - Target to support powers: 20W – 25W
 - Optimized for primary NAND storage in servers
- ❑ E3.S, 2T
 - 76mm x 112.75mm x 16.8mm
 - Target to support powers: 35W – 40W
 - Support for higher power devices like CXL based SCM
- ❑ E3.L
 - 76mm x 142.2mm x 7.5mm
 - Target to support powers: up to 40W
 - Support for higher capacity NAND storage
- ❑ E3.L, 2T
 - 76mm x 142.2mm x 16.8mm
 - Target to support powers: up to 70W
 - Support for higher power devices such as FPGAs/accelerators



High Density QLC - Endurance Limiter?

QLC SSD	Estimated life (100% Rand, W, 100% duty cycle)	Estimated life (100% Seq. W, 100% duty cycle)
61.44 TB	3.07 Yrs 62 PBW	2.25 Yrs 203 PBW
122 TB	~ 6 Yrs 124 PBW	> 4 Yrs



High Density QLC, still takes YEARS to wear out under synthetic workload



* Preliminary estimation for a future 122.88 TB SSD calculations done using published endurance and performance specs for released products with IU aligned writes.

DELIVERING A SPACE AND COST-EFFECTIVE ARCHITECTURE

PRODUCTION PROPOSAL | \$1.6M ENERGY COST SAVINGS OVER 3-YEARS

So lid ig m™
D7-P5520 SSD
So lid ig m™
D5-P5430 SSD



= 10 racks



= 3 racks

4.41 PB workload

Average Power: 145 kW

Max Power: 202 kW

Current Annual Energy Spend: \$727,200*

4.41 PB workload

Average Power: 22 kW

Max Power: 46 kW

Estimated Annual Energy Spend: \$164,529*

*assuming \$300 per kw inclusive of power, cooling, space, security, facility etc.

Wrap up

- Storage needs are differentiated and requires deep understanding of requirement while designing recommending SSDs – QLC and TLC-based SSDs will have unique fit
- Density scaling with lower TCO will dictate the future adoption of form factor beyond U.2
- It is the “fundamental” mixed usage SLA that still dominates the performance discussion
- AI pipeline workload are defining new boundary conditions on how to think of performance



Appendix: Efficient performance for energy-sensitive workloads



Source - Solidigm

Performance Measurements:

System Config for Solidigm D7-PS1010 and Samsung PM1743: Server: Dell Power Edge, Mainboard:095DFK, BIOS: A03, CPU: Intel®Xeon®Gold 6426Y, CPU Sockets 2, Cores:32, DDR5 64GB, OS: Kernal version:3.10.0-862.el7.x86 64 Used System Config 3: Server: SuperMicro ASG-2115S-NE332R, Mainboard: Super H13SSF, BIOS: 5.27, CPU: AMD EPYC 9374F, CPU Sockets 1, Cores:32, DDR5 128GB, OS: Kernal version:3.10.0-862.el7.x86 64, Storage interface: E3.S

System Config for Kioxia CM7: Server: SuperMicro AS -2015CS-TNR, Mainboard: Super H13SSW, BIOS: 5.27, CPU: AMD EPYC 9374F, CPU Sockets 1, Cores:32, DDR5 128GB, OS: Kernal version:3.10.0-862.el7.x86 64, Storage interface: U.2


System Config for Samsung PM9D3a: Server: Inventec-G225-B6, Mainboard: Xiangyang MLB, BIOS: 5.27, CPU: AMD EPYC 9Y24, CPU Sockets 2, Cores:96, DDR5 128GB, OS: Kernal version:3.10.0-862.el7.x86 64, Storage interface: U.2

Power Measurements:

System Config for Solidigm™ D7-PS1010, Samsung PM1743 and Kioxia CM73: Server: SuperMicro ASG-2115S-NE332R, Mainboard: Super H13SSF, BIOS: 5.27, CPU: AMD EPYC 9374F, CPU Sockets 1, Cores:32, DDR5 128GB, OS: Kernal version:3.10.0-862.el7.x86 64, Storage interface:E3.S

System Config for Samsung PM9D3a: Server:Inventec-G225-B6, Mainboard: Xiangyang MLB, BIOS: 5.27, CPU: AMD EPYC 9Y24, CPU Sockets 2, Cores:96, DDR5 128GB, OS: Kernal version:3.10.0-862.el7.x86 64, Storage interface:U.2







Appendix: Performance optimized for real-world workloads

Source – Solidigm and SSD Vendors as Noted. Comparing random mixed performance measurements for Solidigm™ D7-PS1010 versus in-market PCIe 5.0-based SSDs shown below as of July 15, 2024:

Samsung PM1743 - <https://semiconductor.samsung.com/us/ssd/enterprise-ssd/pm1743/>

Kioxia CM7-R - <https://americas.kioxia.com/en-us/business/ssd/enterprise-ssd/cm7-r.html>

Test Configurations:

System Config for Solidigm D7-PS1010 and Samsung PM1743: Server: Dell Power Edge, Mainboard:095DFK, BIOS: A03, CPU: Intel®Xeon®Gold 6426Y, CPU Sockets 2, Cores:32, DDR5 64GB, OS: Kernal version:3.10.0-862.el7.x86 64 Used

System Config for Kioxia CM7: Server: SuperMicro AS -2015CS-TNR, Mainboard: Super H13SSW, BIOS: 5.27, CPU: AMD EPYC 9374F, CPU Sockets 1, Cores:32, DDR5 128GB, OS: Kernal version:3.10.0-862.el7.x86 64, Storage interface: U.2

Workload Descriptions:

General Purpose Server	namespace1= 80% device size with 32K Sequential Read QD32, namespace2= 20% device size with 32K Random Read QD32
Email and UCC	4K Random 70R/30W QD512
OLTP	8K Random 70R/30W QD512
Server based Storage	4K Random Read QD512
Aerospike Database Load	"Nx" load - 1x load (2000 reads/sec and 1000 writes/sec) times N; 128K read/write and 1.5K read; https://github.com/aerospike/act
Multi-tenant Latency-sensitive Workload	4K Random Read QD1 (Avg. Read Latency us), 4K Random Write QD64 (Rate limited to 800MB/s)
Decision Support	16K Random Write QD6





Appendix: Noisy neighbor/multitenancy

Source – Solidigm and SSD Vendors as Noted. Comparing random mixed performance measurements for Solidigm™ D7-PS1010 versus an in-market PCIe 5.0-based SSD, as-tested July, 2024:


Samsung PM1743 - <https://semiconductor.samsung.com/us/ssd/enterprise-ssd/pm1743/>

Performance Measurements:

System Config for Solidigm D7-PS1010 and Samsung PM1743: Server: Dell Power Edge, Mainboard:095DFK, BIOS: A03, CPU: Intel® Xeon®Gold 6426Y, CPU Sockets 2, Cores:32, DDR5 64GB, OS: Kernal version:3.10.0-862.el7.x86_64 Used



EDSFF E3 variants



E3 support for
x4, x8 or x16
PCIe
connection

An E3 double-thickness fits into 2 thin slots and a short E3 fits into a long slot

- E3.S
 - 76mm x 112.75mm x 7.5mm
 - Target to support powers: 20W – 25W
 - Optimized for primary NAND storage in servers
- E3.S, 2T
 - 76mm x 112.75mm x 16.8mm
 - Target to support powers: 35W – 40W
 - Support for higher power devices like CXL based SCM
- E3.L
 - 76mm x 142.2mm x 7.5mm
 - Target to support powers: up to 40W
 - Support for higher capacity NAND storage
- E3.L, 2T
 - 76mm x 142.2mm x 16.8mm
 - Target to support powers: up to 70W
 - Support for higher power devices such as FPGAs and accelerators

Key Feature considerations

	Features	Solidigm™ D7-PS1010 / PS1030	Customer Experience
Standards	NVMe Base Specification NVMe-MI Specification	NVMe v2.0 NVMe-MI v1.2	Compliant and/or certifiable with today's predominant industry standard design and security, manageability protocols
	TCG-OPAL	Version 2.02	
	OCP Compliance	OCP v2.0r21	
	FIPS 140-3 Level 2	Certifiable	
Security	Secure Boot & Firmware Sign	Yes	Modern features to secure hardware and data at rest and in motion
	Format NVM & Sanitize Erase (User/Block and Crypto)	Yes	
	Device Attestation	Yes	
Management	Out-of-Band NVMe-MI (Basic + MCTP)	SMBus + PCIe-VDM	In-band and side channel interface access to deliver maximum capability and flexibility for monitoring and control
	In-Band NVMe-MI	Yes (All Mandatory)	
	Power Management	5 States	
	PCIe Rx Measurement (TP4119)	Yes	
	Weighted Round Robin	Yes	
Perf. + Recovery	Latency Monitoring	Yes (OCP Standard)	Capabilities for enhancing performance, logging and data recovery
	Scatter Gather List	Yes	
	Persistent Event Log	Yes	
	Data Recovery (Device, System, Host)	Yes (OCP Standard)	