TechArena®

2025 Storage Demands Call for a Change

TechArena
May 2025

EXECUTIVE SUMMARY

Data is having its time in the spotlight. What was once seen as an organizational burden has shifted to a strategic priority with storage requirements shifting from a place to hold data that might be used at some point to one that demands immediate access to fuel competitive advantage. This “need it now” requirement has become table stakes for advanced analytics and all flavors of AI, becoming critically important to application viability. With this change, the industry has accelerated data platform innovation starting with software architectures for distributed data access and extending to accelerated control of data arrays, all based on a foundation of high capacity, high bandwidth, low latency storage media technology to fuel distributed data control and analysis.

This paper dives into the foundational role storage media is taking in the modern data center including a look at a major opportunity available to savvy IT operators to utilize innovative media to deliver increased storage density and improved energy efficiency to free up power for accelerated compute. Those planning on next generation storage infrastructure deployments will find useful data for comparison of traditional hard drives (HDD) and solid-state drives (SSD) as well as a look at how to approach the various types of SSD drives on the market.

THE AI ERA DATA CENTER

AI data centers are placing incredible pressure on innovation, challenging every element of the compute platform to advance at a pace faster than Moore's Law. As large tech players deploy AI clusters to deliver to every element of the AI pipeline,

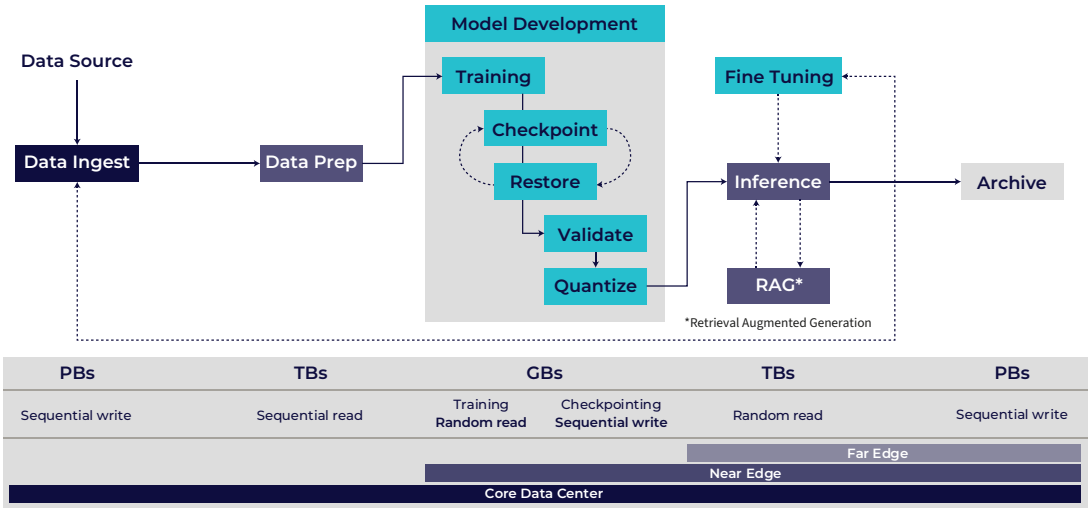

Fig. 1: AI Data Pipeline[1]

the compute density of infrastructure is increasing as well as underlying power demands. In fact, data center power demand is expected to grow more than 3X by 2030.[1] With utility grids already tapped, many operators are seeking as many power saving solutions as possible to squeeze more accelerated compute into existing data centers as well as make the most out of every square foot of data center buildout. When we consider that NVIDIA is forecasting its Vera Rubin platforms, targeted for delivery in 2027, to hit 600 kw per rack vs 120 kw per rack draw of today's Grace Blackwell platforms, we see that the power challenge Is escalating at an exponential rate.

A BRIEF HISTORY OF STORAGE MEDIA

Traditional tiered storage models have utilized a mixture of high-performance SSDs, spinning HDDs, and tape to house organizational data within the cloud, across on-prem data centers, and at the edge. Companies have weighed investment across this continuum between cost, performance, and long-term resiliency. Storage media has not been immune to the accelerated pace of innovation, with SSD technology gaining new capability from scaling NAND layer counts, a market transition to quad-level cell (QLC) technology, and advancements in PCIe Gen5 NVMe controller technology [graph on side-by-side compare of technology in 2020 and 2025] including AI assisted storage control. These advancements have tipped the scales to an advantage for the technology and more tiers of a distributed storage environment. Now even more tiers have been introduced, including new use cases

utilizing fast NAND technology as a distant memory solution for large scale distributed storage arrays. While the performance advantages are clear with scaling capacity and bandwidth, what may be the most important capability of these new storage beasts is an ability to sip limited power per PB through increased storage density.

HOW MUCH DATA ARE WE TALKING ABOUT?

Just how much data is being consumed by advanced AI clusters? With large language model (LLM) training growing to millions of data points on average, data platform leaders are specifying increased capacity networked storage for GPU platforms, with platform configurations scaling to 32 PB and beyond for an NVIDIA DGX rack. [Insert table with network storage provisioning across vendors].

This is where the vast difference of storage media density comes into play. Utilizing standard HDD technology, storage racks can support approximately 360 TB per system or 6.5 PB per storage rack meaning that it would take approximately 5 racks of spinning disk to deliver storage for a single rack of compute.

When we move to the latest 122 TB QLC SSDs, the numbers change dramatically, with delivery of up to 3PB per system and an eye opening 54 PB per rack, reducing the entire rack count. This is why we've previously written on TechArena how this 122 TB introduction may have been the death of HDDs for the vast majority of data center applications. The gap in capacity per drive is just becoming enormous with no real opportunity for legacy technology to advance.

REAL WORLD DEPLOYMENT IMPACT

How does this translate to the IT operator desire to pack more compute? Let's break it down. When comparing a realistic mixed platform configuration of SSD performance tier and HDD object tier vs. an all-SSD solution utilizing 122 TB QLC drives and aim at a mid-point of 16 PB per DGX rack, we see that the "all QLC" solution mops the floor with the traditional tiered approach. The results deliver a nine-to-1 storage rack reduction. Storage power drops by 90% vs the traditional configuration, and that power savings opens the door to deploying 50% more compute capacity for the complete configuration. What's the impact? Faster time to train, faster deployment of inferencing solutions, and the opportunity for a competitive edge in the market.

WHAT ABOUT COST?

Many vendors publish TCO models in this space, but truth be told, all calculators do not integrate the same vectors of cost providing confusing outcomes. TechArena recommends utilizing a TCO model that includes energy costs within equations. When taking this approach, as well as considering HDD short stroking and inferior error reduction techniques, today's performance tier sways to benefit an SSD implementation. A comparison of 122 TB vs. an all-HDD array provides a clear winner with 36% better TCO.

TCO Value When Solving for 100PB Object Storage Solution

	all-HDD Array 3.5" 24TB HDD @ 70% Util.	**Solidigm D5-P5336** U.2 122.88TB @ 95% Util.	
# of capacity drives	17,858	**1,714**	**10x fewer drives**
Total # of servers	1,489	**72**	**20x fewer servers**
Total racks	83	**4**	**20x smaller rack footprint**
TBe/W	5.5	**20.5**	**3.7x better power density**
5-year energy cost	$11.9M	**$780K**	**15x lower energy cost**
5-year total cost	$35.4M	**$22.5M**	**36% BETTER TCO**

Fig. 1: TCO Value[2]

FAST MOVERS WITH 122 TB TECHNOLOGY

Of course, it's important to know that technology is deployable when planning for an upcoming storage acquisition. While 122 TB SSDs were introduced to the market by Solidigm in the second half of 2024, mainstream platforms featuring the technology are now ramping in the market.

Among the first system level storage companies to ship this new 122TB SSD is Dell and their PowerScale platform, a new entrant for the product family that already utilizes Solidigm's high-capacity family of 30TB and 60TB class drives. This next level of innovation from Dell continues to showcase the value of the object storage market need for high-capacity drives. Dell is not alone. We've seen many storage vendors, server OEMs, and end customers across the data center arena voice their strong desire for access to this new capacity, and TechArena expects to see a flood of solutions in the market in 2025.

THE TECHARENA TAKE

So, what do we make of this major advancement in storage technology? The demand for better data management has never been more acute, and the case for SSD adoption has never been clearer. Savvy IT operators should be demanding high-capacity drives from their storage vendor of choice and moving them swiftly through qualification processes to take advantage of the efficiency gains delivered. Data center capacity is too precious, and IT demand for AI compute too strong to not act on this innovation opportunity.

REFERENCES

1. Source: Solidigm, 2025

2. Source: Solidigm, 2025

   TCO calculations based on internal SolidigmTCO estimator tool. Public tool @ https://estimator.solidigm.com/ssdtco/index.htm will support Solidigm D5-P5336 TCO calculations after product launch.

   SOLIDIGM™ D5-P5336 VALUE VS. ALL-HDD ARRAY
   All-QLC configuration:
   Capacity - Solidigm™ D5 P5336, 122.88TB, U.2, 7000 MB/s throughput, 25W average active write power, 5W idle power, 95% capacity utilization, RAID 1 mirroring, calculated duty cycle 8.9%. Solidigm D5-P5336 122.88TB Standard Distribution Cost used for calculations. Consult your Solidigm salesperson for latest pricing.

   All-HDD configuration:
   Capacity – Seagate EXOS X20 20TB SAS HDD ST18000NM007D (datasheet), 9.8W average active power, 5.8W idle power, 70% short-stroked throughput calculated to 500 MB/s; Hadoop Triplication, 20% duty cycle. Price based on ServerSupply 20TB pricing as of July 10, 2023, $0.0179/GB

   Key Common Cost Component 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

TechArena 2025