Petabyte-Capacity SSDs by 2030: Is It Possible, and How?

FMS the Future of Memory and Storage

DapuStor

The Growing Demand for Massive Storage

Annual Data Growth [1]

40%

Driven by AI, big data, and HPC workloads

Enterprise Storage Growth [2]

150%

By 2026, from ~150PB to over **300PB** is projected.

Hot Data Percentage [1]

30%

Hot data is actively used and requires immediate, fast access.

1. Huawei, *Data Storage 2030* 2. Penta, *Data s torage requirements to increase due to AI*

Why Traditional HDDs Can't Keep Up

Limited I/O Performance

HDDs offer <200 IOPS and ~250MB/s throughput due to mechanical limitations, while SSDs deliver orders of **magnitude better** performance.

High Latency

Millisecond-scale seek times make HDDs unsuitable for random access patterns common in AI training and analytics workloads [3].

Scaling Limitations

Achieving 1PB requires **~50 HDDs**, consuming significant rack space, power, and introducing many failure points.

Power Inefficiency

Modern SSDs achieve **3.4×** more terabytes per watt and can reduce power consumption by up to 84% compared to HDD solutions [4].

3. DDN, *NLP and AI: The Role of Data Storage in Powering Generative AI Tools*

4. Businesswire, *Solidigm D5-P5336 SSD improves power and space efficiency for critical IT infrastructure*

The Promise of Petabyte-Scale SSDs

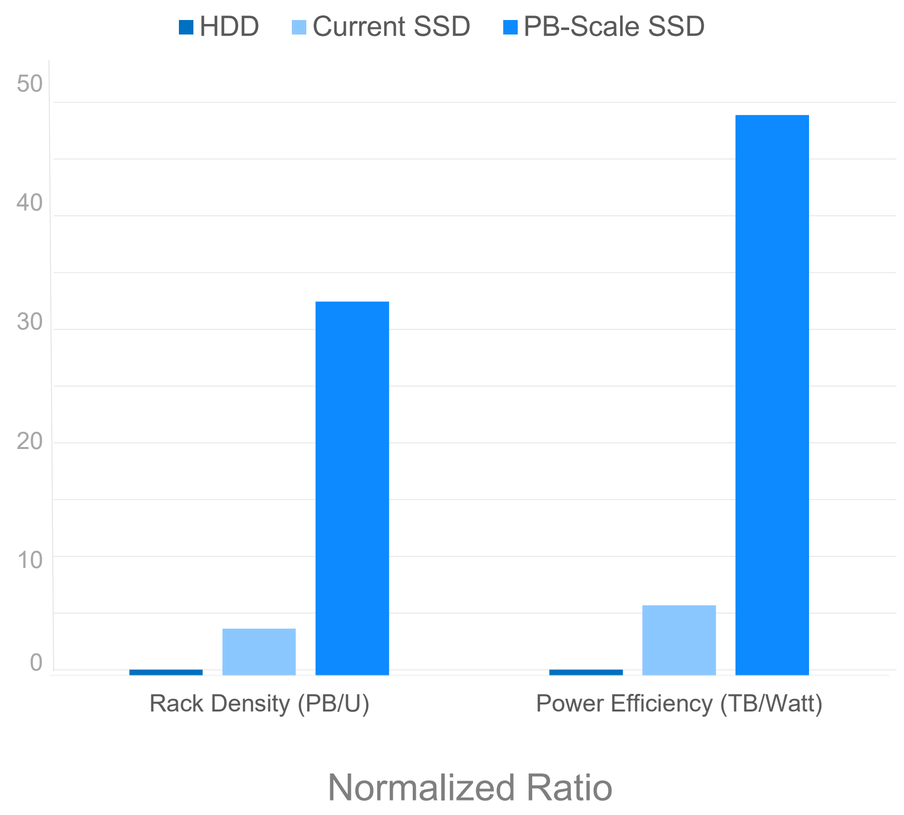

AI Model Training Revolution [5]

Entire training datasets can reside on a single device, eliminating network bottlenecks and accelerating model development.

Cloud Provider Efficiency

4× density improvement means data centers can handle growth without proportional physical expansion.

Scientific Computing Advances

Fast checkpointing and access to massive datasets enables breakthroughs in climate modeling, genomics, and astronomy.

5. Businesswire, *Solidigm Extends AI Portfolio Leadership with the Introduction of 122TB Drive, the World's Highest Capacity PCIe SSD*

Key Technological Barriers

01. NAND Density Scaling Limits

Current cutting-edge NAND is ~232 layers; reaching 500-1000 layers presents challenges in wafer stress, etching uniformity and yield.

02. DRAM Metadata Requirements

Traditional SSDs require ~1GB of DRAM per 1TB of storage for 4KB mapping. A petabyte drive would need ~1TB of DRAM just for mapping tables.

03. Controller Architecture

Managing the parallelism of thousands of flash dies requires rethinking controller design to handle heat dissipation and throughput.

04. Economic Feasibility

Flash remains ~6× more expensive per TB than HDD through mid-2020s, challenging cost justification for petabyte-scale devices.

Enabling Innovation 1: Hardware-Accelerated Compression

Compression Method	Throughput	Latency (4KB)
CPU (Xeon 8458p， 44 cores)	4.9GB/s	~70μs
In-storage ASIC Engine	12.1GB/s	~5μs
Performance Gain	2.5×	14×

Modern hardware compression adds negligible latency compared to typical SSD access times (sub100 μs) while providing substantial capacity gains.

How It Works

Hardware compression engines in SSDs compress data on the fly, effectively multiplying storage capacity while reducing I/O traffic.

Key Benefits [6]

- Effectively doubles usable capacity
- Reduces write amplification
- Increases effective write throughput
- Transparent to host systems

6. Wang et.al., *Reviving In-Storage Hardware Compression on ZNS SSDs through Host-SSD Collaboration*, HPCA'25

Enabling Innovation 2: CXL Memory Expansion

CXL-Based Memory Expansion [7]

The Compute Express Link (CXL) standard allows SSDs to use external (host-side) memory for metadata storage, eliminating the need for massive onboard DRAM.

Minimal Performance Impact

CXL memory adds only 150-250ns of additional latency - negligible for storage operations that typically take microseconds.

Shared Resource Efficiency

Multiple PB-scale SSDs can share a common memory pool, improving resource utilization and simplifying system design.

7. Peng et.al., *XHarvest: Rethinking High-Performance and Cost-Efficient SSD Architecture with CXL-Driven Harvesting*, ISCA'25

Timeline to Petabyte SSDs

By 2030, we expect to see initial PB-scale SSDs in production, with further capacity growth beyond 1PB in the early 2030s.

2023-2024

122TB

122TB SSDs enter production using 200+ layer QLC NAND

2025-2026

250TB

Introduction of PLC (5 bits per cell) NAND; drives reach ~250TB capacity

2027-2028

500TB

500+ layer NAND technology; SSDs approach 500TB raw capacity

2029-2030

1 Petabyte

First petabyte SSDs debut using combined technologies: 1000-layer NAND, PLC storage, hardware compression, and CXL metadata management

SSD Manufacturer Perspectives

Advanced Flash R&D

HDDs offer <200 IOPS and ~250MB/s throughput due to mechanical limitations, while SSDs deliver orders of magnitude better performance.

Computational Features

Integrate hardware compression, deduplication, and CXL interface capabilities into controller designs.

Standards Collaboration

Work with industry groups to refine the NVMe protocol for large namespace management and CXL memory standardization.

Reliability at Scale

Develop intra-drive redundancy techniques and AI-driven predictive maintenance to ensure data reliability at petabyte scale.

Recommendations for Data Center Architects

" Petabyte-scale SSDs are not a question of " if " but " when". Organizations that prepare for this shift by modernizing storage strategies and investing in the right technologies will gain a competitive edge in harnessing their data. "

Plan for High-Density Storage

Leverage compression, deduplication, and CXL interfaces to maximize density.

Modernize Data Management

Adopt software for efficient drive-level compression and object storage.

Embrace CXL Memory Techniques

Use CXL memory modules to bridge CPU and storage needs in next-gen systems.

📌 *Focus on* ***TCO*** *and* ***Sustainability*** *to justify next-gen storage investments.*

DapuStor

Advanced Storage for AI

See You at **Booth #805**

LinkedIn

Website

FMS
the **Future** *of* **Memory** *and* **Storage**