



**DIGITAL STORAGE
TECHNOLOGY NEWSLETTER**

VOLUME XXXX, NUMBER 4

February 2026

SENIOR EDITOR:

Thomas M. Coughlin, Coughlin Associates, tom@tomcoughlin.com

EDITORS AND CONTRIBUTORS:

Dennis Waid, Peripheral Research Corporation, waidprc@aol.com

Stephen Daniel, Steve@DanielRG.com

Jim Handy, Objective Analysis, jim.handy@objective-analysis.com

Andy Marken, Marken Communications, andy@markencom.com

David Rosenthal, david.sh.rosenthal@gmail.com

Tom Gardner, t.gardner@computer.org

Table of Contents

INTRODUCTION	2
NEWS ARTICLES ON STORAGE AND MEMORY TOPICS	4
ARTICLES	5
1) Will AI Make Us Better or Worse, (Forbes, 11/13/25)	5
2) LTO-10 Tape Cartridge Going to 40TB and Vdura V12 Platform Uses SMR HDDs, (Forbes 11/16/25).....	7
3) IEEE Milestones at Nokia Bell Labs and for AES at Leuven, Belgium, (Forbes, 11/17/25)	9
4) CXL Supports Higher Performance and DDN Enables Sovereign AI at SC25, (Forbes, 11/21/25).	13
5) Seagate Demonstrates 7.0TB Per Disk and Discusses Ferroelectric Recording, (Forbes, 11/28/25).....	16
6) Atlas Data Storage Announces Eon 100 Synthetic DNA Storage, (Forbes, 12/6/25).....	19
7)) Micron Exits Consumer Business and Focuses on AI and Data Center Sales, (Forbes 12/8/25)	21
8) The Field Effect Transistor Centennial and Final IEEE Event, (Forbes, 12/10/25).....	23
9) MRAM Forum Features Automotive and other MRAM Applications, (Forbes Blog, 12/13/25)	25
10) Digital Storage and Memory Projections for 2026, Part 1, (Forbes, 12/19/25).....	30
11) Pixitmedia Survey Shows Media and Entertainment Content Archive Trends (Forbes Blog, 12/20/25).....	39
12) Komprise Unstructured Data Survey shows AI Driving Data Management, (Forbes Blog, 1/1/26)	44
13) Digital Storage and Memory Projections for 2026, Part 2, (Forbes Blog, 1/2/26).....	47

14) Nvidia CES AI Announcements Expand AI Capabilities and Continue Memory and Storage Shortages, (Forbes Blog, 1/7/26)	52
15) Digital Storage and Memory Feeds AI at the 2026 CES, (Forbes Blog, 1/10/26).....	57
16) Inference Context Memory Storage Platform could Drive Additional 100EB of AI Storage, (Forbes Blog, 1/16/26)	63
17) Digital Storage and Memory Projections for 2026, Part 3, (Forbes Blog, 1/21/26).....	66
18) Digital Storage and Memory Projections for 2026, Part 4, (Forbes Blog, 2/2/26).....	71
19) C4Q 2025 and 2025 HDD Industry Update, (Forbes Blog, 2/3/26)	74
20) Western Digital Innovation Event—Storage is where AI goes to Learn, 2/3/26).....	77
GUEST ARTICLES.....	83
1) 100M IOPS SSDs? You Must be Kidding, Jim Handy, Objective Analysis. Originally published in the SSD Guy blog 2/2/26	83
2) CXL is Coming—Don't Get Left Behind, Jim Handy, Objective Analysis, 2/9/26	86
3) Internet Archives Storage, David Rosenthal,.....	87
4)It's a Global Show About AI Everything, Andy Marken (Content Insider CES--Visit), andy@markencom.com	93
5) We Need a Source to Solve AI's Problems, Andy Marken (Content Insider #936—Killer AI, andy@markencom.com	104
6) How will your total available market change over the next five years, Steve Daniel, sdaniel1958@gmail.com	114
QUARTERLY HDD NUMBERS UPDATE	180
CURRENT AND UPCOMING REPORTS	191
2025 NEW MEMORIES REPORT.....	192
2022 DIGITAL STORAGE IN PROFESSIONAL MEDIA AND ENTERTAINMENT REPORT.....	194
CONSULTING	196
NEWSLETTER SUBSCRIPTIONS.....	197

INTRODUCTION

This is the February 2026 issue of this newsletter and the last issue in the 2025-2026 annual volume.

2026 marks the 41th year of this newsletter, formerly called the Rigid Disk Heads and Media Newsletter (and originally published by Peripheral Research, Inc.).

Here are some **top items impacting digital storage in the fourth calendar quarter of 2025 (and the first month of C1Q 2026)**:

- The numbers reported here for some shipment information are our best guess estimates based upon information from the HDD industry.

- We do not know for sure the total number of HDDs delivered last quarter because Western Digital and Seagate did not report their unit shipment numbers. However, based upon our analysis we estimate 33.4M HDDs shipped in the last calendar quarter of 2025, up 4.8% from the prior quarter estimate.
- Total exabytes shipped in CQ4 2025 were up 9.1% compared with CQ3 2025 (456.9EB versus 418.7EB).
- HDD revenues (estimated at \$6.6B) were up about 7.5% from the prior quarter.
- Total HDD unit shipments in 2025 are estimated at 124.3M, up 0.4% from the prior year.
- Estimated nearline HDD shipments in 2025 were 80.8M, up 13.6% from the prior year.
- Total exabyte shipments for the year were 1,620.95EB (1.61ZB), up 22% from the prior year.
- Total industry revenue in 2025 were \$23,267M (\$23.3B), up about 18% from the prior year.

You are getting this newsletter as a member of the Entertainment Storage Alliance (www.entertainmentstorage.org) As a member of the Entertainment Storage Alliance you can also purchase our published reports and other Coughlin Associates (and partner) Reports with a 10% discount. For more information on these and other reports please see the technical papers page on the Coughlin Associates web site at www.tomcoughlin.com. To order call 408-202-5098 or email tom@tomcoughlin.com.

This issue has several guest articles by various friends of the Entertainment Storage Alliance including **Steve Daniel, David Rosenthal, Jim Handy and Andy Marken**. It also includes sales flyers for the **latest Digital Storage in Media and Entertainment Report** and the recently released **Emerging Memory Report**.

Please let your friends know about this newsletter. There is a subscription form at the end of the report. Feel free to pass this newsletter to your friends in the industry.

Thanks for being a customer!

Tom Coughlin

News Articles on Storage and Memory Topics

We are now doing regular postings of links to storage and memory related news (including company announcements) to a new Digital Storage and Memory Flipboard. You can find this flipboard and follow it at: <https://flipboard.com/@thomasmcoughlin/digital-storage-and-memory-ku86vmbnz>

ARTICLES

1) **Will AI Make Us Better or Worse, (Forbes, 11/13/25)**

I was invited to participate in the 2025 Kailath Symposium and Colloquia Agenda, which celebrated the 90th birthday of Thomas Kailath, Professor Emeritus of Electrical Engineering at Stanford University. It was an amazing event and was supported by the IEEE as well as Stanford University.

The people who participated as well as those attending were a stunning array of technical luminaries and entrepreneurs. They included many IEEE Medal of Honor and ACM Turning Award winners such as Vint Cerf, Martin Hellman, John Hennessy, David Forney, Robert Kahn, Irwin Jacobs, John Cioffi, Mark Horowitz and Andy Viterbi, who is shown with me below.



Those in the audience were equally well known in engineering circles, including Donald Knuth, author of *The Art of Computer Programming*. Vint Cerf gave a great talk on the history and future of the Internet. Martin Hellman gave a fascinating talk about *The Wisdom of Follishness*. There were several video presentations honoring Professor Kailath, who also spoke. John Hennessy gave an insightful talk on *Computational Challenges on the Road to AGI*.

I participated in a panel discussion focusing on the Future of Intelligence with John Cioffi, Amin Arbabian, Vint Cerf and Andy Viterbi. We were each asked to come up with a two-minute introduction to the topic before a panel discussion. Below is what I put together as my introduction.

In my mind, intelligence is the ability to make good informed decisions, based upon available information. In that sense, at least, many creatures and things have some level of intelligence, whether it is a worm, a whale, a bird, a dog, a primate or many machines.

I also note that we humans have a long history, particularly in the last 300 years, of creating machines of various types to do repetitive and physically demanding work that is difficult for humans to do well or for long. Indeed, many of these historical developments have eliminated jobs, but also created even more new jobs.

Today's various approaches to doing AI may continue that trend, allowing people to focus on activities that play to our strengths in making judgements and decisions using factors that machines do not now have and to learn and develop solutions in entirely new ways. This could allow an increase in individual human agency as well as new ways for us to work together to achieve worthy common goals.

For instance, there are several startups that are focusing on ways that humans can leverage AI's abilities to accelerate the discovery of new physical phenomena, materials and especially scientific discovery. This is an example of how AI can help us become better at what we do.

AI could also be used to enhance human interaction and to spread democratic practices. In today's New York Times, Eric Schmidt and Andrew Sorota, in an opinion piece, discussed the dangers of governance by algorithm, or "algocracy," where public governance decisions are made by AI, e.g. to reduce corruption. In contrast they also point out that AI could be used to reinvigorate democracy, making it more responsive, more deliberative and more worthy of public trust.

We have also seen how AI can be used to enhance engagement, especially engagement that play to sensationalism and emotional reactions that can and have led to negative consequences, such as an increase in civil discord and widespread belief in false information. If used irresponsibly, AI could lead to increasing divisions between people and greater conflicts than any we have experienced, and it could threaten our own survival.

Of course, there is also the fear that as AI gets better it will replace humans in more and more fields, including white collar jobs such as engineering. However, as mentioned previously, it could create whole new jobs and new and satisfying ways of living and having a livelihood.

It is clear that AI can be an aid to advancing human potential, or for creating conflict. It can be used to enhance special interests and lead to a decline in human accountability and nobility or for helping us follow our better natures. It is up to us, as people engaged in the creation, implementation and use of AI to determine


whether AI will ultimately make us better at who we are, or further reduce human sovereignty and welfare.

IEEE supported the 2025 Kailath Symposium and Colloquia at Stanford University, attended by many technical luminaries. I participated in a panel on the Future of Intelligence.

2) LTO-10 Tape Cartridge Going to 40TB and Vdura V12 Platform Uses SMR HDDs, (Forbes 11/16/25)

This blog focuses on two recent announcements. One from the LTO Program about 40TB LTO 14 tapes being available by 2026. The other is Vdura's recent announcement of its Data Platform V12 software that improves performance and increases capacity using SMR HDDs.

Magnetic tape is an important storage technology supporting most hyperscale data centers and other enterprises. It is the lowest cost storage technology in widespread use today and represents the lowest tier of storage in most organizations. LTO tape is the dominate tape in use although IBM also makes an enterprise tape, currently with capacities up to 50TB per cartridge. The image below shows current Coughlin Associates projections for HDD, SSD and tape annual storage capacity shipments out to 2030.




The LTO Tape Program said that it was increasing the storage capacity of its LTO-10 cartridges to 40TB of native capacity by the first calendar quarter of 2026. The program often refers to the 2.5:1 compressed capacity of its LTO tapes, which would be 100TB for this updated LTO-10 tape.

The Ultrium LTO announcement quotes Stephen Bacon, VP at HPE as saying that, "AI has turned archives into strategic assets. The new 40 TB LTO-10 cartridge will

help enterprise-class organizations—across healthcare, financial services, media, research, manufacturing, the public sector and beyond—consolidate petabytes efficiently, strengthen cyber resiliency with true offline air-gapping, and keep long-term retention affordable and sustainable.”

This 33% increased capacity from the 30TB LTO-10 tape [announced earlier this year](#) is said to be due to a new Aramid base film, the substrate, that enables a thinner, smoother and longer tape in the same cartridge form factor. Longer tape means more capacity per cartridge. There were also refinements in the drive head design to allow more tape in the cartridge. Existing LTO 10 tape drives will be able to use the new higher capacity tapes.

In addition, the LTO roadmap was adjusted for LTO 11 out to LTO 14, as shown below.



[It now peaks with a 913TB compressed capacity, 365TB native capacity](#), down from earlier LTO 14 generation targets of 576TB native.

Vdura released its software Data Platform V12, featuring its elastic metadata engine, snapshot support and shingled magnetic recording, SMR, hard disk drive optimization. SMR HDDs allow increased storage capacity in hard drives by partially overwriting track on top of each other. Read performance is essentially unchanged, but writing is slower for SMR HDDs compared to conventional HDDs. Vdura says that, “The Elastic Metadata Engine enables metadata operations to scale linearly across multiple metadata nodes, supporting billions of files and objects under active use. Combined with snapshot and SMR tier support, V12 further reduces total cost of ownership while extending VDURA’s lead in AI-scale storage economics. With aggregate throughput performance increased by more

than 20%, V12 delivers unmatched efficiency for training pipelines, inference workloads, and hyperscale data movement.”

The LTO Program announced 40TB LTO-10 tapes by 2026 with an adjusted capacity roadmap out to LTO generation 14. Vdura released a new data platform allowing higher capacity per rack using SMR HDDs.

3) IEEE Milestones at Nokia Bell Labs and for AES at Leuven, Belgium, (Forbes, 11/17/25)



Technology drives the world’s economy and IEEE has an IEEE Milestone program that recognizes technical developments of at least regional importance that have helped to shape the modern world. These developments must be at least 25 years old for IEEE Milestone recognition. I was able to help represent the IEEE at two events recently. The first was the celebration of 7 new milestones at the Nokia Bell Labs in Murray Hills, New Jersey. The second was for AES, the widely used data encryption technology in Leuven, Belgium.

On October 21 there was a large IEEE milestone celebration at Nokia Bell Labs in Murray Hills, New Jersey. The event celebrated seven technologies developed in Bell Labs. Bell Labs gave birth to many technologies that have enabled our modern world. Some of these developments are explored in Jon Gertner’s book, [The Idea Lab](#).

At the Nokia Bell Labs there are 13 IEEE milestone plaques on the wall, including these new milestones plaques. Two of these were duplicates of milestones in other locations, including for Project Echo at the Robert Wilson Park in Holmdel,

New Jersey, and five were new milestones. Many of these technologies led to Nobel Prizes for the inventors. The full list of the milestones honored at this event are given below:

Super-Resolved Fluorescence Microscopy – Eric Betzig (Nobel Prize 2014)
Charge-Coupled Device – George Smith & Willard Boyle (Nobel Prize 2009)
Fractional Quantum Hall Effect – Horst Stormer, Robert Laughlin, Daniel Tsui (Nobel Prize 1998)
Molecular Beam Epitaxy – Al Cho
Convolutional Neural Networks – Yann LeCun
Development of the BELLMAC-32 Microprocessors – Duplicate Milestone (California)
Echo Project – Duplicate Milestone (Robert Wilson Park, Holmdel)

I also attended an IEEE milestone event near Leuven to celebrate AES encryption, Belgium on November 18, 2025. The image below shows me with the new IEEE Engineering Milestone plaque next to the inventors of Rijndael, which became AES, Vincent Rijmen and Joan Daemen.



Below are my remarks at the IEEE Engineering Milestone ceremony for Rijndael and the Advanced Encryption Standard (AES) from 1995 through 1998 on November 18 at the Auditorium Tweede Hoofdwet, Kasteelpark Arenberg in Heverlee, Belgium. I spoke about the IEEE and IEEE milestones in general, while Martin Bastiaans, member of the IEEE History Committee and past IEEE director, talked about some of the IEEE milestones in the Benelux section and in IEEE Region 8.

I would like to thank my colleague Luc Claesen for all his efforts in planning today's celebration, and also for his long-time dedication to IEEE and ongoing volunteer service, particularly as a champion for the IEEE Life Member community and for IEEE milestones.

I would like to share a warm welcome to all our members joining us from across the IEEE Benelux Section. I greatly appreciate the support from KU Leuven and thank vice rector Stefaan Vaes for hosting us today, as well as my colleagues – Dirk Rabaey, Wim van Etten, Patrick Dewilde, and Martin Bastiaens, who assisted in the planning of today's celebration. It was my honor to listen to the inventors, cryptographers Joan Daemen and Vincent Rijmen earlier today discussing the innovative achievements they led that resulted in this milestone.

I am so pleased to have Mohamed Ridouani, the mayor of Leuven, with us as well. IEEE milestone plaques can serve as inspiring public educational tools, helping communities understand the significance of local contributions to global progress.

IEEE is a public charity – which means that the purpose of our work is to provide education, in our case, for scientific purposes. Our activities are all tied to our mission — of advancing technology for the benefit of humanity— which help us demonstrate the great impact that IEEE, our members, and engineering and technology have on society.

These activities encompass efforts in areas such as education, ethics, humanitarian technologies, public policy, sustainability, and history. They help to increase the public's understanding and appreciation of our fields of interest and highlight not only our members' technical expertise, but also their desire to use their skills and knowledge to benefit humanity.

As the world's largest technical professional organization, it is an important part of IEEE's mission to preserve the legacy and heritage of our professions, to recognize great achievements, and to promote the importance and impact of engineering.

Our professions have a long, rich, and important history. And it is essential that our pioneering technological developments and discoveries, and the men and women behind them, are recognized and celebrated for generations to come.

The IEEE Milestone Program recognizes great moments throughout our world's long history of technical innovation. They are a symbol around the world of how electrical engineering and computing have built the modern world.

By showcasing achievements across disciplines—from computing to biomedical engineering—milestones connect past achievements to future possibilities. They remind students, researchers, and professionals that bold ideas can have lasting impact.

Milestones have been dedicated across all 10 IEEE regions, making it one of IEEE's most visible and respected programs and reinforcing IEEE's mission to advance technology for humanity by honoring innovations that truly changed the world.

The IEEE milestone program is funded by donations to the IEEE Foundation, the philanthropic partner of IEEE. The Foundation leverages the generosity of donors interested in supporting our mission activities, such as today's milestone, to support and enable IEEE efforts to enhance technology access, literacy, and education, and support the IEEE professional community.

Since our inception, IEEE has had a standing History Committee that advises the IEEE Board of Directors on matters of the legacy and heritage of IEEE, as well as administering and coordinating the IEEE milestone program. The Center maintains many useful resources for engineers, students, historians of technology, and for anyone interested in the development of electrical and computer engineering and its role in modern society.

Besides our focus today, I would encourage all of you to experience and, if appropriate, contribute to the IEEE Oral History Program, our peer-to-peer interviewing program that has collected spoken memories and personal commentaries of historical significance through recorded interviews of hundreds of leading engineers and technologists.

IEEE's commitment to historical stewardship is a strategic pillar of our public mission. As a global nonprofit, IEEE works to safeguard the stories, standards, and systems that define the enduring impact of electrical and electronic

engineering. Working to ensure that the achievements behind our greatest innovations remain visible, relevant, and instructive— and serving as a guide for ethical design and responsible leadership.

I would like to once again share my sincere gratitude to the organizers of today's event, and KU Leuven for their warm hospitality.

Visiting these historically significant locations, where local volunteers have worked with technical experts, historians, and the public to celebrate pioneering achievements and the people behind them, is truly inspiring and a testament to how technological innovation benefits humanity. They are a great way to experience the history of technology alongside the global IEEE community.


Preserving the history of our profession is vital for fostering innovation, inspiring future generations, and ensuring the safe, sustainable, and ethical development of technology that will continue to shape our world.

New IEEE milestones at Nokia Bell Labs and in Lueven, Belgium demonstrate the importance of technical innovations in enabling the modern world.

4) CXL Supports Higher Performance and DDN Enables Sovereign AI at SC25, (Forbes, 11/21/25).

The IEEE Super Computing Conference, SC25, was held in St. Louis, Missouri this last week. Digital storage and memory technologies are key elements in high performance computing, including artificial intelligence. There were some interesting digital storage and memory announcements at the show. In this piece we will explore the latest generation CXL announcement as well as DDN announcement of NVIDIA-certified solutions to support sovereign AI.


The CXL consortium maintains the specifications for the Compute Express Link interconnect for high-speed, high-capacity CPU-to-device and CPU-to-memory connections. It expands and allows pooling of server memory beyond what is possible with traditional DIMM slots. Since its initial introduction in 2019, CXL has gone through a number of updates to the specification as shown in the image below.




The consortium announced the CXL 4.0 specification at SC25. Key features of that specification are given below:

- Doubles the bandwidth to 128GTs with zero added latency
- Enables rapid data movement between CXL devices, directly improving system performance
- Maintains previously enabled CXL 3.x protocol enhancements with the 256B Flit, network flow control unit, format
- Introduces the concept of native x2 width to support increased fan-out in the platform
- Support for up to four retimers for increased channel reach
- Implements CXL bundled port capabilities
- Ability to combine device ports between Host and CXL accelerators (Type ½ devices) to increase bandwidth of the connection
- Delivers memory RAS, reliability, availability and serviceability, enhancements
- Continued full backward compatibility with CXL 3.x, 2.0, 1.1, and 1.0

The ability to logically aggregate CXL device ports, or bundled ports, between a CXL host and CXL accelerators, such as the image shown below, allows increased connection bandwidth between that host and a CXL device connection. This could be especially useful for AI applications, since GPUs are generally data bandwidth constrained and can use all the data bandwidth they can get.



The figure below details the increased capabilities of CXL with the release of generation 4. As shown, generation 3 and 4 provide capabilities for improved data center and wide area network, WAN, computing. This includes memory pooling where various types of memory, including DRAM but could include non-volatile memories as well as various local computing accelerators can be shared between processors and servers.



DDN was also at SC25 and announced Sovereign AI Blueprints based upon NVIDIA-aligned reference designs. The sovereign-by-design architectures are aimed at helping governments and other organizations that need to restrict where their data resides deploy storage for artificial intelligence applications.

These solutions are based upon DDN's unified data-intelligence software stack for NVIDIA certified systems. The company says that they are a part of sovereign AI programs in Europe, India and other countries in Asia. They keep data within defined border, governed under local policies and protected with verifiable encryption and access control. The company says that they deliver more than 99% sustained GPU utilization and up to 40% lower energy consumption in a data center.

The CXL consortium announced its 4th generation interconnect with up to 128GT/s performance. DDN announced support for sovereign AI.

5) Seagate Demonstrates 7.0TB Per Disk and Discusses Ferroelectric Recording, (Forbes, 11/28/25)

There has been a fair amount of coverage in the last couple of days related to a presentation Seagate Technology gave in [October about a 6.9TB per HDD disk laboratory demonstration](#). This appears to have been first reported by guru3D. The [seminar abstract below](#) was used as the reference to this advancement.



The 290th Special CMSM seminar



The Future of Hard Disk Drive Storage Technology

Dr. Stephanie Hernandez

*Senior Engineering Director,
Advanced Data Storage Device Group,
Seagate Technology*

Date & Time: 15:00 - 16:00, October 21st (Tue), 2025.

Place: 2nd Conference Room, Main Bldg. 1F, Sengen.


Abstract:

Heat assisted magnetic recording (HAMR) technology has been officially released at capacities of over 3 TB/disk. Recent laboratory experiments of HAMR demonstrate 6.9 TB/disk. This presentation will explore ideas to maximize the ultimate areal density capability of HAMR to 10 TB/disk and beyond, including novel system architectures and reader designs. It will also introduce ferroelectric recording as a potential successor to magnetic recording. And finally, AI techniques for design and materials discovery for future HDD technology will be presented.


In March of 2024 [Research groups from NIMS, Seagate Technology, and Tohoku University made a demonstration of multi-level recording using a three-dimensional magnetic recording medium](#). This demonstration showed a path beyond the roughly 10Tbpsi areal density that could be achieved with heat assisted magnetic recording, HAMR. It was believed that 3-4 multiple recording levels could be possible using this technology.

The article also referenced increasing the magnetic recording areal density from a then shipping 1.5Tbpsi to 4.0Tbpsi and beyond. At the time 32TB 10-disk HDDs were shipping, implying 32TB capacities require about 1.5Tbpsi. So, a 4.0Tbpsi HDD could provide a 10-disk 85TB HDD. The image below shows a comparison of current HAMR recording to a two-layer recording.

a Conventional (2-level) recording



b Next-generation (multi-level) recording




Seagate Technology referred to this demonstration during a May 2025 analyst event as showing [that they had achieved a laboratory magnetic recording demonstration that would allow 6.5TB per disk](#), or 65TB for a 10-disk HDD. Based upon the calculation above the areal density for this demonstration is then about 3Tbpsi, or $1.5 \times 6.5 / 3.2$. The recent announcement of 6.9TB per disk probably implies an areal density of about 3.2Tbpsi, and may also utilize multi-layer recording, although that is not clear from the October announcement.


At that May analyst event Seagate showed the following chart giving projections for laboratory demonstrations of laboratory demonstrations and a rough idea of when these demonstration storage capacities would be achieved in shipping products. This chart implies that 65TB 10-disk drives could be likely available in the early 2030's. Of course with 11 and even 12 disk products announced by WDC and Toshiba respectively, such HDD capacities could be available before 2030.

Product Maturity Cycle Within 5 Years

From Lab Demonstration to Productization



The chart below is an updated one showing current shipping product maximum areal densities compared to the maximum laboratory demonstration areal density with this adjustment for laboratory magnetic recording areal densities. We estimate that this latest laboratory demonstration at 6.9TB per disk is about 1.9X higher than for the current shipping product.



Another interesting thing mentioned in the October paper abstract was the use of ferroelectric recording as a potential successor to magnetic recording. Ferroelectric solid state non-volatile memory devices have been around for

decades and arguably these have been produced in the highest volume, particular in RFID readable fare cards, particularly in Asia.

However, the materials used in the older ferroelectric memory devices contained elements like lead that are not conducive to semiconductor manufacturing processes. More recently hafnium oxide has been found to have a ferroelectric state that could be used for a solid-state memory. Hafnium oxide is commonly used as a higher dielectric constant insulator in modern solid-state electronics and so this material could be more acceptable in a modern semiconductor manufacturing facility.

Presumably, a ferroelectric recording technology from Seagate would extend the conventional HDD technology with disks, heads and actuators, but somehow use ferroelectric recording approaches, perhaps like those used in ferroelectric solid-state memories, rather than conventional ferromagnetic recording as in conventional magnetic recording.

A search found [information on such a recording approach](#) for a ferroelectric recording drive from 2016. The abstract for that paper says, "Ferroelectric probe data storage has been proposed as a novel data storage method in which bits are recorded based on the polarization directions of individual domains. These bits are subsequently read by scanning nonlinear dielectric microscopy. The domain walls of typical ferroelectric materials are quite thin: often only several times the lattice constant, which is advantageous for high-density data storage. In this work, high-density read/write (R/W) demonstrations were conducted using a hard disk drive-type test system, and the writing of bit arrays with a recording density of 3.4 Tbit/in.² was achieved. Additionally, a series of writing and reading operations was successfully demonstrated at a density of 1 Tbit/in.²"

I am making inquiries to see if I can learn more about how this ferroelectric recording drive would work and will be glad to share it, when I have that information.

Seagate demonstrated magnetic recording technology enabling 6.9TB per disk, likely with multi-level recording. The demonstration also mentioned a potential move from magnetic to ferroelectric recording in the future.

6) Atlas Data Storage Announces Eon 100 Synthetic DNA Storage, (Forbes, 12/6/25).

The value of digital data of all types has always been important, but with the growth of today's AI models and the data centers that support them, the need for data has exploded. This is an important driver for the current shortages of digital storage and memory and the significant price increases in these technologies.

Today this data is stored on solid state drives, hard disk drives, magnetic tape and optical storage. Magnetic tape and optical storage are the primary ways that data

still having value but not being actively accessed is stored. Magnetic tape continues to evolve and there are at least six optical storage startups active in developing technology for long term data storage. In addition, there are a few startups working on long term digital storage on synthetic DNA, including Atlas Data Storage.

Atlas Data Storage was spun out from Twist Bioscience in May 2025 with \$155M in seed funding. The company recently announced Atlas Eon 100, which it said is the first scalable DNA storage service. Below is an image supplied by Atlas of its Eon 100.



Atlas Eon 100 uses synthetic DNA, that is, artificially created DNA sequences that do not derive from a living creature and presumably would have no biological function. The company said that this technology is:

- Compact - 1,000 times denser than magnetic media
- Encoded in a universal, time-tested format
- Efficient to copy
- Easy to transport
- Unparalleled durability

Twist had been working to create a viable DNA storage system for over 10 years. Bill Banyai, Atlas founder said that, “...*We intend to offer new solutions for long-term archiving, data preservation for AI models, and the safeguarding of heritage and high-value content*”

Atlas officially showcased the Atlas Eon 100 at the Association of Moving Image Archivists, AMIA, Conference on December 3 in Baltimore, Maryland.

There are a few startups, including Atlas, working on archival storage using synthetic DNA. DNA can last for a long time with little degradation if kept in a

cooler environment without humidity. The main issues with DNA storage relate to the speed with which DNA base pairs can be written and read.

Data rates for writing and reading DNA are nowhere near other current storage technologies, but the companies working on this technology point out that DNA writing and reading speed have been increasing since the original genomic work early in the 2000's, driven by the developments of DNA-based medical developments. If these trends continue, DNA data storage could reach speeds that start to approach other technologies in another ten years.

DNA writing has another feature that could be of value for archivists, a single write can produce millions of copies of the same synthetic DNA strands. This creates a natural way to create data redundancy, which can help in recovering data during later reads, but it also allows storing copies of the data in many places, improving the odds that it will survive over the long term.

Atlas Data Storage, spun out from Twist Bioscience in May 2025 announced its Atlas Eon 100 scalable synthetic DNA storage service. Synthetic DNA is one of the technologies vying to support long term data preservation.

7)) Micron Exits Consumer Business and Focuses on AI and Data Center Sales, (Forbes 12/8/25)

This last week [Micron announced that it would discontinue its Crucial consumer products brand](#) and focus its efforts on memory for the AI and data center markets by February 2026. The Crucial brand was started in 1996. This closes almost three decades of Micron selling into consumer markets [note that Micron later clarified that it would continue to sell memory to PC manufacturers although not to individuals].

Micron's decision comes at a time when data center expansions to support AI workloads has caused shortages of all digital storage and memory products. Micron has been making headwinds in selling its high bandwidth memory, HBM, with stacks of DRAM in a package in the GPU market. It also comes after Micron's memory fabrication expansion plans ran into some headwinds.

The shortages of storage and memory to support AI data center build-outs has caused price increases for every type of memory and storage technology. Many of the hyperscale data center companies are making contracts with suppliers, like Micron, to purchase these technologies well into the future. This is causing increased prices for consumer products which use these memory and storage technologies and it is also making it harder to supply the consumer market.

Sumit Sadana, EVP and Chief Business Officer at Micron said that the "surging demand for memory and storage solutions in the AI-driven data center market" is behind this decision. Micron and other solid state memory companies have generally supported both consumer and enterprise customers even though

consumer products generally have lower profit margins. It will be interesting to see if other memory companies follow suit.

Another factor in Micron's decision may be that its plans for factory expansion to provide more memory production at its Clay, NY megafab have run into [2-3 year construction delays](#) due to a combination of labor shortages and supply chain complications. The first of these facilities is now pushed out to open in late 2030. Micron is also building a new semiconductor fab in Boise, Idaho, shown below.



It may be that the slower expansion of its memory production capacity in the near future was another reason for Micron to exit the consumer market in favor of the more lucrative enterprise and data center market.

Micron's exit will leave a considerable hole in the consumer DRAM and NAND flash memory business. Overall [Micron is the third largest supplier of DRAM](#) behind Samsung and SK hynix. On the other hand, Micron's exit from the consumer market should strengthen the market share of Samsung, Sandisk and other companies.


Some milestone's in Crucial's consumer product history include:

- **1996:** Crucial Technology was founded, allowing everyday consumers to buy memory upgrades direct from a major manufacturer.
- **2008:** The first Crucial-branded SSD was debuted with capacities up to 64GB.

- **2013:** Crucial was a pioneer in introducing DDR4 memory technology to the market.
- **2018:** The P1 series SSD was launched, setting new performance benchmarks for consumer products.

Micron announced that it will exit the consumer memory business by February 2026 and focus on the more lucrative enterprise and data center markets.

8) The Field Effect Transistor Centennial and Final IEEE Event, (Forbes, 12/10/25)



This last week I had my last two public appearances as the IEEE Past President. My term ends at the end of December, although I still have a few on-line meetings before December 31. This last weekend I participated in the IEEE New Era AI World Leaders Summit in Seattle and starting on Monday I was at the 2025 International Electron Devices Meeting, IEDM, in San Francisco.


On Monday morning, December 8 I presented two IEEE technical field awards at a morning event with probably close to 2,000 people attending. The first of these awards was for IEEE Fellow Sayeef Salahuddin from UC Berkeley for pioneering contributions to the physics of ferroelectrics and integrated ferroelectric devices, who received the Andrew S. Grove Award sponsored by the Electron Devices Society, EDS, which put on the 2025 IEDM.

The Fredrick Phillips Award, sponsored by the IEEE Frederik Phillips Award Fund, recognizes outstanding accomplishments in the management of research and development resulting in effective innovation in the electrical and electronics industry. Deirdre Hanford received this award for visionary leadership in electronic design automation for secure and energy-efficient microelectronics.

The 2025 IEDM celebrated the 100th anniversary of the invention of the field effect transistor, or FET. In 1925, Polish-American physicist Julius Lilienfeld, who was

an experimental pioneer in field electron emission, filed a patent in Canada for what we would now call a field effect transistor device and later filed a US patent on that invention. This was a three terminal device to control current using a gate electrode.

The US patent image of the device is shown below.



The term, transistor, was later coined by technologists at Bell Labs in the late 1940's to describe the devices developed there after the invention of the point contact transistor by John Bardeen, Walter Brattain, and William Shockley in 1947.

Lilienfeld later invented a metal oxide semiconductor transistor, what would now be called a MOSFET. Unfortunately, the materials Lilienfeld worked with, copper sulfide semiconductors were not the best materials for this application and the high purity semiconductor germanium and silicon was not available until the 1940's. In addition, it was not until the 1940's and 1950's that the role of surface states was well enough understood to make such FET's commercially viable.

John Atalla and Dawon Kahng at Bell Labs developed the first working MOSFET in 1959-1960 using thermally grown silicon dioxide as the insulator on a silicon semiconductor.

At the 2025 IEDM the Electron Devices Society, EDS, had a set of posters near their exhibit that presented the history of semiconductor technology, shown below.



IEDM organizers also put on an invited lunch that focused on the history of transistor technology and future prospects for this technology featuring talks by Hiroshi Iwai, NYCU/Inst. Sci. Tokyo and Bob Johnson from Gartner. An evening session titled, FETs at 100: Yesterday's Sarks, Tomorrow's Arc—Human Wisdom Meeting AI Insights that featured a presentation generated by ChatGPT as well as human speakers, Serge Beisemans from imec, Chenming Hu from UC Berkeley, John Wu from AMD, Daewon Ha from Samsung Electronics and H. S. Philip Wong from TSMC and Stanford.

In my last official physical meetings as IEEE Past President, I attended an AI conference in Seattle and participated in the 2025 IEDM and its celebration of the 100th anniversary of the Field Effect Transistor.

9) MRAM Forum Features Automotive and other MRAM Applications, (Forbes Blog, 12/13/25)

At the end of the 2025 IEEE IEDM, Samsung and other companies sponsored the 17th annual MRAM Forum. The Forum featured speakers on MRAM commercial technology and related developments from TSMC, GlobalFoundries, Samsung, Everspin, NXP, Renesas, Sony, Numen, LAM Research, HProbe, Raytheon, IBM and Tohoku University. There was also a panel discussion at the end that included two companies with chips using MRAM, Applied Brain Research and Alif Semiconductor. Let's look at a few tidbits from the Forum. There was also a meeting of a nascent MRAM consortium the night before the MRAM Forum.

Allen Wang from TSMC spoke about high-performance MRAM for Advancing Automotive Innovation. He showed an advanced logic roadmap for automotive

applications that indicated that 3nm products were on track to be automotive qualified by the end of 2025. He showed 16nm MRAM technology passing qualification that allow 4 cycles of solder reflow as shown below.


Robust N16 MRAM Technology	
Node	16nm
Application	Flash-like MRAM
Cell Size(um ²)	0.025
Memory Density	168Mb
Reflow Capability (260C/30s)	4 Cycles
Endurance (cycles)	>1E6
Retention	>150C_20yr
Chip Failure Rate	<1ppm

Achievement:


- > Excellent Short BER/Reflow BER after MTJ Film/Etch/Process co-optimization
- > Meet Auto-G1 spec w/ CFR <1ppm
- > Achieve high yield (>90%) for high memory density MRAM
- > Successfully pass qualification!!

He also spoke about RAM-like and Flash-like MRAM at 16nm (RAM-like has 20ns write pulses and 10^{15} write cycles of endurance with retention times of more than one day, while Flash-like has less than 100ns write pulses and 10^6 write endurance with retention times of more than 20 years. These can be scaled down to 5nm.


Siddharth Gupta from GlobalFoundries also spoke about MRAM in automotive applications with 22nm-FDSOI technology. He addressed the various needs and challenges for automotive MRAM, shown below. He spoke about various architectural factors and also designing for test time reduction, bit protection and magnetic immunity.



Jeong-Heon Park from Samsung addressed automotive MRAM as well as, and in particular 8nm eMRAM and a scheme for wafer level magnetic shielding (that was also shown in an MRAM session at the IEDM). The figure below shows various applications for eMRAM as a function of latency and endurance.



He and the following Everspin speaker addressed magnetic field vulnerability of MRAM in various applications. Samsung showed the following slide regarding the various magnetic fields in IoT and edge devices while Everspin gave a demonstration of the magnetic field immunity. The basic message is that with good designs with spacing between MRAM and strong magnetic field sources or with magnetic shielding, magnetic field sensitivity can be addressed.




Thomas Jew from NXP showed work that NXP had discussed at the 2023 MRAM Forum and showing their more recent developments with the company's S32K5 MCU family using 16nm MRAM from TSMC for automotive applications. The S32KS with embedded MRAM and low-leakage SRAM supports write speeds up to 15MB/s and higher durability up to 1M write cycles and 20-year data retention and also supports post-quantum cryptography for long term security. The MRAM memory can flexibly support code and data storage.

Tomoya Saito from Renesas spoke on embedded MRAM MCUs for edge AI applications. In particular he spoke about inference processing with a neural

processing unit, NPU, that uses on-chip MRAM and SRAM. The device supports weight compression to improve speed and reduce memory demands.

Masanori Hosami from Sony spoke on recent developments in Sony's use of eMRAM for image sensor applications. He pointed out that Sony spoke about STT-MRAM at the 2005 IEDM. He showed a slide, below, with examples of various embedded devices and MCUs that include MRAM memory.



He also spoke about ReRAM, particularly Weebit memories ReRam product and talked about what MRAM needed to do in order to survive, comparing MRAM with Weebit's ReRAM product. Sony's primary interest is the use of non-volatile memories in imaging applications.

MRAM could help realize a simple camera system that allows consolidating prior external flash and SRAM into a single MRAM memory and using the tunable features in MRAM to provide fast and lower retention memory with slower and higher retention memory for various steps in image processing, especially for in-device motion detection, dewarping and region of interest detection.

Max Simmons from Numem spoke about its AI memory engine. He seems an opportunity to replace and supplement some SRAM and external cache DRAM. The image below shows Numem's view of MRAM advantages for various IoT and edge applications.

MRAM & ITS APPLICATIONS					
MRAM Key Benefits for Edge AI and Physical AI	WEARABLES	AUTOMOTIVE	Laptop	ROBOTICS	Smart Camera
Longer Battery Life No cell refresh	☒	☒	☒		☒
Instant-On Non-Volatile	☒	☒	☒	☒	☒
Longer Product Life High Endurance	☒	☒		☒	
Smooth AI Workload Very low latency	☒	☒	☒	☒	☒
Better TCO Unified Memory	☒	☒	☒	☒	☒

Xiaoyu Kang from LAM Solutions spoke about the etching challenges with the magnetic tunnel junctions, MTJs, used in MRAM. He described how the LAM Kyber Platform, shown below can provide MRAM etching and encapsulation in a single cluster system.

Lam Solutions: Integrated Kyber® Platform

Non-reactive, multi-angle Ion Beam Etch

IBE

- Noble gases (Ar, Ne, Xe, Kr), O₂
- Full beam control
- High (2000V) and Low (10V) energy
- Full range of tilt angle: 0-90°

PECVD Encapsulation (Versa G)

PECVD Encapsulation

- 65-90% step coverage
- In situ pre-treatment (Ox/Red)
- Low particles: <10ea @>40nm

Integrated System


Integrated System

- IBE & PECVD system
- No air exposure of MTJ
- No H₂O permeation
- 1 PECVD module, supports >2 IBE
- Option to add Kryo® RIE for mask open

Kang said that the Kyber ion beam etching tool has been used in making MRAM devices for over 5 years and that it had an extensive installed base in nine sites for use in MRAM and other types of memory including RRAM, DRAM and NAND flash. For MTJ this tool provides improved MTJ etching strategies at tight pitch and he showed up to 97% yield on 50nm pitch MRAM from the Lam-imec JDP and said that further improvements are possible.

Siamak Salimy from Hprobe announced a couple of new test systems. The H3DM-XLL magnetic instrument is shown as part of an Hprobe test system in the image

below. This tester can measure the influence of uniform magnetic fields of over 500mT from any angle and from -40C to 200C on magnetic devices, such as MRAM.



Hprobe also introduced its MC16X multiprobe pulse generator. This system allows testing of magnetic devices with more than two terminals such as various spin orbit torque (SOT) MRAM and cross-point MRAM structures. The system supports 16-256 pulsing channels with pulse widths of 200ps-100ns and supporting fall and rise times from 80ps to 300ps.


At the 2025 MRAM Forum major foundries and chip designers talked about automotive and other applications and ways to deal with magnetic field sensitivity of MRAM devices. LAM and HProbe spoke about MRAM manufacturing equipment and new test equipment.

10) Digital Storage and Memory Projections for 2026, Part 1, (Forbes, 12/19/25)

For several years I have written a series of articles outlining my projections for various digital storage and non-volatile memory technologies for the coming year and beyond. This will be the first of my projection articles for 2026, focusing on magnetic recording technology, particularly hard disk drives (HDDs) and magnetic tape.

2025 turned out to be a boom year for all types of digital storage and memory technologies, driven mostly to satisfy storage and memory demand for data centers in order to support various AI workloads. HDDs are no exception as cost effective secondary storage is needed to store AI model logs, training data and other information for longer term storage. The three hard disk drive companies, Seagate, Western Digital and Toshiba have either reported on or are assumed to have long term contracts with major data centers going out as far as 2027. There are reports that [HDD contract prices jumped roughly 4% quarter over quarter in Q4 2025.](#)


Our current projections are for 2025 revenues to be at least 23% higher than in 2024. However, the HDD companies are being cautious in their expansion plans, likely guided by their long-term contracts with major data centers. Below we show our optimistic projections for total HDD revenue going out to 2030 as well as a more cautious model with possible corrections if current AI data center spending undergoes a correction, starting in 2026. Note that this kind of correction would be more in line with the longer industry revenue trends.



Note that given the current cost difference between HDD and SSDs (likely remaining greater than 4X out to 2030, at least), even with the higher capacity QLC NAND flash eSSDs, we don't believe that this will have a significant impact on the revenue projections. Rather we think the higher capacity eSSDs will support AI workflows and their data generation that will eventually end up being stored for the long term on less expensive storage, such as HDDs. However, we do expect continuation of the long-term decline in demand for legacy smaller capacity HDD applications continuing into 2026 and later years.


We estimate nearline high-capacity HDDs unit shipments in 2025 will be about 67% of the total HDDs, up about 10% of the total HDD shipments in 2024. Most of the shipped storage capacity is in nearline HDDs. For instance, Seagate reported that over 87% of their shipped capacity was for nearline drives in C3Q 2025. By 2030 we estimate that over 90% of all shipping HDDs will be for nearline applications. The nearline drives have higher capacity and sell for higher prices,

as a consequence, as the percentage of nearline HDD grew the average sales price of HDDs has been generally increasing since 2015 as shown below.




Also, these nearline drives are also in short supply in 2025 and this has also contributed to higher ASPs this year. Note that reporting changes by Seagate and Western Digital will make it harder to track ASPs for HDDs in the future.

Results for the final quarter of 2024 have not been announced, but my estimates project that overall HDD unit shipments in 2025 will be about the same as in 2024, at about 124M units, perhaps ending the long term reduced HDD shipment trends since the peak shipments in 2010. The chart below shows the history of total HDD unit shipments through 2024 with high, medium and low shipments projected for 2025 through 2030.



[Note that with the changes in Seagate and Western Digital quarterly reporting it may be difficult to get unit shipment information in the future.](#) Not having this information will leave investors with additional risk since they won't know actual product production information.

HDD capacity shipments are estimated to increase about 34% from 2024 to 2025, to about 1.79 ZB. We project total HDD capacity shipments will be about 6.5ZB by 2030. HDD industry capacity shipment history and projections out to 2030 is shown in the chart below.



Market share for HDD capacity shipments from the three HDD manufacturers in 2025 is estimated at about 47% for Western Digital, 42% for Seagate and 11% for Toshiba. Through C2Q 2025 unit shipment market share in 2025 was 43% for Western Digital, 40% for Seagate and 17% for Toshiba. Because of unit shipment report changes by Seagate and Western Digital we don't have good unit shipment data for C3Q 2025.

The HDD companies indicate that they are entering into long term agreements with major customers to ensure meeting the customer's anticipated demand and to prevent manufacturing excess products. Some of these contracts are out to 2027.

Increasing storage capacity per drive results in lower prices per bit of data (\$/TB). For high-capacity enterprise and data center HDDs, helium sealed drives with more disks allows higher capacities. Currently these drives are available with up to 12 disks (with WDC's 32TB 11-disk SMR drive introduction in 2024 and with Toshiba's announcement of upcoming 12-disk drives in 2026). The advantage of more disks in a drive is higher capacity, but it also comes with additional costs for the added disks and heads.


HDD capacities can also be increased by shingling the tracks, partially writing new tracks over existing tracks, although as the drive gets full, writing new data requires moving good data to another location and then writing the new and good data again (a process somewhat like the erase and write process for SSDs). This type

of shingled magnetic recording or SMR drive is most effective for write-seldom applications and WDC continues to say that they have a robust market demand for SMR HDDs.

HDD capacities can also be increased by increasing the track density or the linear density of the information recorded on the concentric tracks. Track density increases had slowed down up until recently, however, energy assisted magnetic recording where energy is applied to very stable magnetic recording media (having high magnetic coercivity) to allow it to be written on, is boosting storage capacity areal density growth.

WDC is currently using what they call ePMR, which stabilizes the written bits and which they call a form of energy assisted recording. This is used in their latest nearline HDDs, including the 11-disk products, which combined with their OptiNAND technology and SMR, achieved 32TB capacities. Toshiba says that they are using a form of micro-wave assisted magnetic recording, MAMR, in their high-capacity HDDs but that future 30+TBHDDs will include heat assisted magnetic recording, HAMR, as well. Toshiba also announced that they will develop 12-disk HDD products.


Seagate has been shipping HAMR HDDs for a few years now and in their latest quarter they said that they had shipped over 1M units (out of an estimated 12.5M units), with 32TB and 36TB with or without SMR. By 2026 we expect that Seagate's percentage of HDDs shipping with HAMR will be much higher, although it will be difficult to check unless the company provides that information. In May 2025 [I published the chart, shown below](#) on Seagate's laboratory demonstrations and product maturity, showing about 5 years for product maturity from laboratory demonstrations.



The chart shows 40+TB HDDs by 2026 and 50+TB drives by 2028 with 100TB HDDs possible by the early 2030's.

Both WDC and Toshiba are planning to introduce HAMR HDDs for the nearline industry by next year, although volume production will likely remain behind Seagate for at least a year beyond that. In order to match these HDD capacities WDC will need to either ramp up its own HAMR HDD program or it, along with Toshiba, will need to buy HAMR heads from TDK, and perhaps HAMR media from Resonac (formerly known as Showa Denko). Both companies are believed to be working together to develop HAMR products, although there has been little publicity recently on HAMR technology development.

Seagate recently demonstrated 6.9TB per disk technology, boosting the industry's laboratory demonstration. The chart below is our latest tracking developments for the highest capacity (in areal density) laboratory (technology) demonstrations and shipping HDD products.




Seagate and WDC have also shipped some dual actuator HDDs (where the rotary actuator is split into two effective actuators on the same rotating point) in the past few years, but this technology hasn't taken off. It is expected that these dual actuator HDDs (which Seagate called its Mach-2 HDDs) will become standard fare on higher capacity HDDs sometime in the future, although the HDD companies have backed of saying when. I expect that the higher price of dual actuator HDDs is a factor in the slower introduction of these drives. However, the major customers will likely eventually make the switch to multi-actuator HDDs, since they can


increase performance and, more importantly, deduce the rebuild time for failed drives.

The figure below shows the most recent HDD technology roadmap from the Advanced Storage Research Consortium (ASRC), from the 2022 IEEE TMRC conference. ASRC is HDD industry consortium that provides projections on storage density growth (areal density is the average amount of information stored per area on the HDD disks). This image shows projections for over 10 Tbpsi HDDs by the mid-2030's (which should provide over 200TB HDDs) using a combination of technologies that includes HAMR. Note that e-PMR and MAMR don't show the capacity growth potential that HAMR does.

ASRC HDD Technology Roadmap - 2022



The figure below shows the history and our projections for declining price of HDD storage in \$/GB, with projections out to 2030.




Magnetic tapes also use magnetic recording, but the areal density of magnetic tape is many generations behind HDDs. This gives magnetic tape significant future technology for increasing storage capacity in the future. However magnetic tape is different than HDDs since it uses flexible media and multiple tracks are written or read at a time. This means that magnetic tape, once mounted in a tape drive, has data rates much greater than HDDs. HDDs are increasingly used for nearline storage while magnetic tape is used to store colder data and particularly for archiving data. The primary magnetic tape technology is LTO, with an estimated 80% of the market with IBM enterprise tape providing the balance.

In August of 2023 IBM announced its TS1170 magnetic tape drive, supporting storage capacities up to 50TB native and higher capacities with compression. The product has a native data rate of 400 MB/s and has 12Gb SAS and 16Gb Fibre Channel interface options. These drives support new tape cartridge media using Strontium Ferrite magnetic particles. The JF media is produced by Fuji Film, who also makes LTO media. This 50TB tape cartridge technology was a 2.5X increase from prior generation 20TB products.

LTO magnetic tape is the most used with LTO 10 tape the highest capacity LTO tape shipping today. LTO 10 was introduced in May 2025 with 30TB native storage capacity. [In November the LTO consortium announced that LTO-10 was available with 40TB native storage capacity](#) by moving to a thinner tape and thus more tape area on the cartridge. Both products support 400MB/s data transfer speeds.

The LTO tape consortium also adjusted its LTO roadmap as shown below. The roadmap goes to LTO generation 14 with up to 365TB of native capacity, down from the earlier target of 576TB of native storage capacity. This roadmap is shown below.



If new LTO generations are introduced about every three years, which has been typical, and if each generation continues to follow this roadmap, then LTO 14 could be available in about 2036, with 365TB native tape cartridge capacity at about the same time as possible 200TB HDDs.

Tape manufacturers have been promoting magnetic tape as a more sustainable storage media, since little energy is used when magnetic tape cartridges are sitting on shelves. Fujifilm estimates that about 43% less carbon emissions result from the use of magnetic tape compared to HDDs.

The total market for magnetic tape media, drives and libraries is probably over \$2B. IBM is the sole manufacturer of magnetic tape drives and Fujifilm and Sony are the only manufacturers of magnetic tape media. We estimate that total tape capacity shipments in 2025 will be 220EB and that this will climb to over 400EB by 2030.

HDD shipping capacity and revenue are projected to be over 6.5EB and between \$25B and \$35B by 2030. We project magnetic tape will ship over 400EB by 2030.

11) Pixitmedia Survey Shows Media and Entertainment Content Archive Trends (Forbes Blog, 12/20/25)

One of the markets for digital storage and memory that I have covered for several years is the media and entertainment industry. I have done a report on this for

several years and plan to release a new one in 2026. Archiving of media content is an important element of this because historical media assets can be remonetized and thus has ongoing value.

Pixitmedia recently concluded a short, 10 question, survey of 330 digital media professionals about their perceptions on digital media archiving, in aggregate and then broken down by market segments for the participants. You can access the report [here](#).

Pixitmedia is now part of Datacore. Pixitmedia provides digital storage and software. The company is focused on using AI tools to help users understand, organize and manage digital media assets.

The survey participants were from various professional media markets including TV and video broadcasting, TV and film studios, sports teams and franchises, OTT TV and video services, enterprise video, publishing, house of worship and digital service providers.

We will compare and supplement some of their findings with [results from the Coughlin Associates Digital Storage in Media and Entertainment Survey](#).


The Pixitmedia survey shows the need for active archives that accelerate time to market through unified data management, enhance productivity via efficient media search and discovery, optimized resources with scalable remote collaboration and sustaining profitability by orchestrating the entire media value chain.

Here are some of the general findings from the Pixitmedia survey, and Coughlin Associates survey results where appropriate.

Almost one-third (32%) of large M&E companies expect their content libraries to exceed 5PB within the next 12 months, up from 17% today. In the 2025 Coughlin Associates survey, almost 36% had more than 2,000 hours of media content in their archives. In addition, in the Coughlin Associates survey, 57% of the respondents said that their annual archive growth rate was >6%. 23% of the Pixitmedia survey participants anticipate a 5-9% increase in content library value over the next 12 months.

Three quarters (76%) expect to allocate at least moderate higher technology budgets in the next 12 months, broken down by business size in the figure below.

Media Archiving Workflow budgets in the next 12 months by M&E business size




Source: Pixitmedia's State of Media Archiving Survey, n=330 (September 2025)

Preferences for on-premises workflows will increase from 25% to 31% over the next year in the Pixitmedia survey. In the Coughlin Associates survey archiving in a private or public cloud was down from 19% in 2023 to 11% in 2025. It appears that there is a larger trend to use more on-premise versus cloud storage for many elements in media workflows.

Many M&E companies are gravitating towards on-premises media archiving in the short to medium term, in order to regain control over sensitive data. 31% of the Pixitmedia survey respondents said that on-prem was their preferred deployment, up from 25% in the past year.

At the same time, hybrid deployments continue to dominate, preferred by 40% of respondents for 2026, compared to use by 42% in 2025. Hybrid architectures combine local performance with remote accessibility, improving editorial and post-production collaboration while supporting agile, distributed creative teams. The hybrid model also provides strong data resilience through automated, offsite backups and disaster recovery. The chart below shows these trends.




Source: Pixitmedia's State of Media Archiving Survey, n=330 (September 2025)

A unified media archiving platform replaces proprietary closed systems with integrated interoperable platforms to reduce complexity and the total cost of ownership. Eighty-five percent plan to migrate to a unified media archiving workflow. Almost half (45%) of Pixitmedia respondents cite automation, remote collaboration, and metadata enrichment as top workflow challenge

However, as of 2025, only 6% of M&E companies have fully migrated to a single platform. But the following chart shows strong support for more migration to a unified media archiving system in the future.


Global: Migration road-map to a single unified media archiving platform, 2025–2027



Source: Pixitmedia's State of Media Archiving Survey, n=330 (September 2025)

AI is widely used in professional media for enhancing metadata, automation and content discovery. Almost half (45%) of respondents cited operational productivity as the primary benefit of AI/ML integration, followed by automation of manual processes (22%) and cost savings (29%). Among those emphasizing cost benefits, more than two-fifths (41%) pointed to deduplication and anomaly reduction as key outcomes as shown in the chart below.

Global: Benefits achieved via AI/ML pre-configured with media archiving workflow, 2025




Source: Pixitmedia's State of Media Archiving Survey, n=330 (September 2025)

Pixitmedia media archiving survey shows increased demand for active unified media archives that use AI for management and content enhancement, balance on-premises versus cloud storage, and provide for a return on investment for growing storage spending.

12) Komprise Unstructured Data Survey shows AI Driving Data Management, (Forbes Blog, 1/1/26)


Komprise, an unstructured data management software company, recently released their [fifth annual survey on unstructured data management](#) with over 300 IT and storage leaders participating in the 2025 survey. The survey shows some significant growth trends in unstructured data capacity and spending growth in 2026. Much of this growth is driven by unstructured data to support AI workloads. As the graphic below shows how much unstructured data the survey participants were managing with 74% having more than 5PB. Komprise says that this is a 57% increase over the results from the 2024 survey.

How much unstructured data are you currently managing?




Furthermore, as shown in the pie chart below, 85% of the survey participants anticipate spending more on storage in 2026 than in 2025.

Will you spend more on data storage in 2026 v. 2025?



The next graphic shows topic storage priorities for the next twelve months that will be driving this storage capacity and spending growth. Optimizing storage costs is the top priority at 64%, but it is closely followed by 61% for data preparation and classification for AI. Cloud migration was the third highest priority at 54%. Supporting the use of storage for AI, 39% also reported supporting AI use cases in production as a storage priority.

What are your top data storage priorities in the next 12 months?



Some other results from the survey show the role of AI and other applications as well as general data management. The top technical challenges for unstructured data management included classifying data for AI at 58% and moving data without disruption at 53%. The top business challenge for unstructured data management is to reduce the data risk from AI for 62% of the participants. The greatest data concern for generative AI is security and in particular concern about corporate data leakage for 46% of the participants.

The top challenge in preparing data for AI workloads is classifying and tagging that data according to 56% of the survey participants. This is up from 41% in 2024. As a result the leading future requirement for unstructured data management is data classification and tagging for 61% of the participants. The second leading challenge is data governance and security concerns at 46%. 47% of the survey

participants worry about their departments lacking visibility into their storage spend and data use.

Some additional leading survey future requirements for unstructured data management were analytics and reporting at 60% and sensitive data detection at 57%. The top skills gaps from the survey participants were AI data management at 62%, up from 43% in 2024 followed by cloud storage strategies at 60% and data security and compliance at 49%.

Planning for 2026, 40% of survey participants plan to increase their IT budgets in 2026 to pay for AI. This was on 30% of the participants in 2024 planning on increasing their IT budgets for 2025. Also, in order to meet security and all these AI requirements, 64% of the IT leaders will make investments in upgrading their data storage and management platforms in 2026. There were 53% planning to do this in 2025 in the 2024 survey.

58% of the survey participants are creating an internal task force of IT, security, legal and other departments in order to develop a corporate AI strategy. Nearly half of the survey participants said that they will be adding staff, particularly IT infrastructure leaders to focus on developing their AI foundation at 53% and 49% said that they were hiring engineers and developers with AI expertise.

The Komprise unstructured data management survey shows that 74% have 5PB or more of stored data and 85% of participants expect to spend more on data storage in 2026 than in 2025.

13) Digital Storage and Memory Projections for 2026, Part 2, (Forbes Blog, 1/2/26)

This is the second in a set of four blogs about projections for digital storage and memory for the following year that we have been doing for several years. Our first blog focused on the latest developments and projections for magnetic recording (HDDs and magnetic tape). This blog focuses on status and projections for various types of solid-state non-volatile memory and storage. We will talk about the latest developments in NAND flash memory, DRAM and non-volatile memories like MRAM, RRAM, FRAM and PCM.

Memory and storage technologies recovered from a severe recession in 2022 and 2023 and a significant recovery in 2024. 2025 has seen demand grow even more with shortages reported for all the major storage and memory technologies. Much of this growth is to support data for AI workflows and has particularly hit consumer and other solid-state memory and storage users with higher prices and low availability.

For instance, manufacturers have prioritized DRAM use in high-profit high bandwidth memory, HBM. This led to soaring prices for both consumer and server DRAM, despite growth in overall capacity shipments. As of [Q3 2025 DRAM](#)

[contract prices increased by about 172% year over year.](#) Trendforce also estimated a 35% increase in average selling prices for 2025.

The same is true for NAND flash. [Kingston said that the price of NAND flash for them had increased by 246% in 2025.](#) Overall, some reports indicate that some NAND prices more than doubled in under six months in late 2025.

Some solid-state memory manufacturers are planning on investments to increase production in 2026. For instance, [CXMT in China is planning on an IPO on the Shanghai market to increase their DRAM production capacity](#) to support Chinese DRAM demand. Micron has also reported on expanding its production capacities for DRAM and NAND flash in Boise, Japan and with partners, in Taiwan.

Kioxia and its partner Sandisk as well as YMTC plan to increase their production capacity of NAND flash. But there are also reports that [NAND demand in 2026 will grow by 20-22% year over year while supply is projected to rise by 15-17%, leading to shortages in 2026.](#)


Solid state memory technologies such as DRAM and NAND flash are made in expensive fabrication facilities and expansion of production capacity takes a significant time to come online. As a consequence, with continuing demand growth and slower increases in production we expect increasing prices in 2026 for these memory and storage products.

TrendForce [projects 12.8% server growth from 2025 to 2026.](#) This will continue to drive DRAM and NAND flash demand and likely higher prices in 2026. However, if there is any reduction in AI spending in 2026 these trends could quickly reverse.

SSDs are primary storage for frequently used data in most data centers and in particular for AI workflows, and this is driving demand for these products. Active data lives on SSDs with colder data kept on hard disk drives, HDDs, and magnetic tape. The SSDs companies are also trying to introduce high-capacity QLC NAND storage products, with four bits per NAND memory cell, with capacities of over 200TB in 2026 to displace HDDs as secondary storage for less active data. [The images below show large QLC NAND eSSD announcements from the 2025 FMS last August.](#)




Higher capacity SSDs can take less rack space than HDDs and they offer higher performance than HDDs offer. This can be an advantage for AI training and inference with RAG, but flash memory is currently about 6X more expensive per storage capacity than HDDs and is expected to remain so, for some time into the future, as shown in the image below from the WDC investor day in February 2025.



Recent NAND flash capacity increases have focused on smaller diameter memory cells, closer memory cells and particularly on using more bits per cell, rather than more NAND flash layers. 2025 saw the introduction of NAND flash products with more than 300 layers. We don't expect big increases in layer numbers in 2026 as the cost reductions with more layers are much less than in the past.

The higher capacity SSDs being offered do have somewhat lower operating power and allow considerably more storage in a given rack space, compared to HDDs, while providing higher performance, which can be an advantage in some applications. However, the QLC flash has lower endurance and thus is best used in write seldom applications. As a consequence, I don't believe they will displace too much of the demand for nearline storage using HDDs. I expect instead that they will create a new layer in the SSD storage hierarchy.

Another interesting development at the 2025 FMS in August 2025 was the announcement by Sandisk and SK hynix of an Memorandum of Understanding to create standards for high bandwidth flash, HBF, with NAND flash in similar packaging as DRAM in HBM to supplement HBM for AI applications. An image of a flash based, HBF module architecture is shown below.



In addition to providing more memory using NAND flash in HBM-like packages, NAND-flash is also a non-volatile memory while DRAM is a volatile memory. That means that DRAM requires regular refreshes of the data it contains and this consumes energy. By substituting some non-volatile memory for what otherwise would be volatile memory in GPUs and AI applications it may be possible to reduce the energy requirements for AI applications in data centers.


This could open up more AI opportunities for AI applications in energy constrained data centers, thus democratizing the development of AI applications. It could also help to reduce the projected power requirements for hyperscale and other data centers to support AI development.

The market for non-volatile memories such as magnetic random access memory, MRAM, resistive RAM, ReRAM, ferroelectric RAM, FRAM, and phase change


memory, PCM, continues to develop, particularly for embedded memories. In our 2024 report on non-volatile memories, [A Deep Look at New Memories](#), we looked into developments and projections. The latest report will be available in January 2026.

The first widespread use of emerging memories is for embedded memory on a CMOS logic chip, replacing NOR flash, which has reached its scaling limit at 28nm, and is today frequently being replaced with MRAM and ReRAM. Single-transistor MRAM cells are now competing with multi-transistor static random access memory, SRAM, to dramatically reduce the number of memory transistors in order to provide a lower-cost, higher-density solution. Several enterprise, industrial and consumer devices currently use MRAM as an embedded memory, and this trend is growing.

At the 2025 IEEE MRAM Forum, at the end of the 2025 IEDM conference in San Francisco, sponsored by Samsung and other companies and organized by the IEEE Magnetics Society and the IEEE Electron Devices Society, Masanori Hosami from Sony spoke on recent developments in Sony's use of eMRAM for image sensor applications. He pointed out that Sony spoke about STT-MRAM at the 2005 IEDM. He showed a slide, below, with examples of various embedded devices, MCUs and standalone devices that include MRAM memory.



TSMC spoke about their 16nm MRAM technologies for automotive applications. NXP gave an update on their automotive products using MRAM technology. Jeong-Heon Park from Samsung addressed automotive MRAM as well as, and in particular 8nm eMRAM and a scheme for wafer level magnetic shielding (that was also shown in an MRAM session at the IEDM). The figure below, from the Samsung talk, shows various applications for eMRAM as a function of latency and endurance.



Everspin had a demo on magnetic field immunity for MRAM and there were presentations on when this could be an issue and how to mitigate it so customers can use MRAM in their products.


Based upon considerations of economy of scale, that is, the ability to drive up production volume over time, the upcoming New Memory report projects that total baseline new memory annual shipping capacity will rise from an estimated 1.4PTB in 2024 to 13.5EB in 2035. Total new memory baseline revenues are expected to increase from \$565M in 2024 to about \$171B by 2035. The bulk of this rapid revenue growth will be supported by new memory technologies' displacement of SRAM, NOR flash and some DRAM.

AI is driving demand and higher prices for DRAM and NAND into 2026. Products using non-volatile memories to replace NOR and SRAM in embedded applications are growing.

14) Nvidia CES AI Announcements Expand AI Capabilities and Continue Memory and Storage Shortages, (Forbes Blog, 1/7/26)

At the 2026 CES Nvidia CEO and IEEE Medal of Honor Winner, Jensen Huang announced availability of Vera Rubin AI computing gear and new context aware memory. The capabilities of Vera Rubin will enable faster, less expensive AI training but the increased demand for data will drive additional demand for digital memory and storage, continuing shortages into 2026.

During his talk Huang spoke about Rubin, now in production at Nvidia that the company described as an extreme-co-designed six-chip AI platform. He also introduced Alpamayo, an open reasoning model for autonomous vehicles. He spoke about Alpamayo and the various other open AI models that the company had trained and made available to the industry for a broad range of domains. The image below shows these open models.




The Rubin computing platform is named after American astronomer Vera Rubin, who did early work leading to the theory of dark matter. Rubin succeeds the prior Blackwell computer architecture. It is a six-chip platform with the following specifications:

- Rubin GPUs with 50 petaflops of NVFP4 inference and with 22TB/s HBM4 bandwidth and 3.6GB/s NVLink bandwidth per GPU
- Vera CPUs engineered for data movement and agentic processing with 1.tTB of system memory and 1.8TB/s NVLink bandwidth
- NVLink 6 scale-up networking with 3.6TB/s per GPU all-to-all bandwidth
- Spectrum-X Ethernet Photonics scale-out networking with 102.4Tb/s scalee out switching
- ConnectX-9 SuperNICs with 800 Gb/s Ethernet and a programmable RDMA and data path accelerator
- BlueField-4 DPUs, an 800 Gb/s DPU for SmartNIC and storage processing

Another innovation is that the Vera Rubin compute tray has no cables, hoses or fans and can be assembled in 5 minutes.


He also introduced an AI-native storage architecture, called Nvidia Inference Context Memory Storage enabled on the BlueField 4. This is an AI-native KV-

cache tier to boost long-context inference and enabling five-times higher tokens per second, five-times better performance per dollar and five-times better power efficiency. The images below show the resulting Rubin compared to the prior generation Blackwell performance.



Rubin takes $\frac{1}{4}$ the time to train, resulting in ten times as many tokens generated per second at a token cost that is $\frac{1}{10}$ as Blackwell.

Huang also said that the resulting Nvidia computing system, an example of which is shown below, is liquid cooled and that the water used for cooling runs is at 45°C or 113°F. Cooling with warm water will have a big impact on AI data center water and power requirements.



The new context aware memory is built for long-context-processing agentic AI systems and enables very fast long-term and short-term memory. By extending an AI agents long-term memory this enables high-bandwidth sharing of data context across clusters of rack-scale AI systems and enables the five times higher power efficiency and token generation.


However, running Rubin AI data centers will create even greater demand for digital storage and memory. Nvidia wants to use 16-high DRAM stack HBM, with DRAM already in short supply and the data rates and resulting additional data consumption for generating tokens and achieving these cost savings will require both greater storage performance and costs. This could drive additional demand for NAND flash, certainly for SSDs, including eSSDs, but could also drive demand for high bandwidth flash, NBF to help provide data to the HBMs connected to the GPUs.

In addition to solid-state storage, the need for more data for training and for storing various types of data from training will also drive demand for hard disk drives and perhaps also magnetic tape for longer term storage. Hard disk drives were also in short supply at the end of 2025.

So, I believe, if Nvidia's Rubin goes into production and demand for AI continues, digital storage and memory will remain a gating item, at least through 2026. The NAND and DRAM manufacturers are starting to invest in new fabs, but these will take many months to come on line. HDD manufacturers can increase their production capacity with less investment, but particularly if new head manufacturing is required, this would also take several months to come online.

Thus storage and memory shortages and higher prices will likely continue through this year.

Also, during the 2026 CES, Jensen Huang was announced as the winner of the 2026 IEEE Medal of Honor, IEEE's highest award. During a short ceremony on January 6 he gave a short speech that thanked the IEEE, spoke about how as an IEEE student member he met his wife, also an IEEE student member, in a lab, talked about his experience at Nvidia and finishing by saying that the award was really for all the people who worked at Nvidia and helped to make their products a reality. The image below shows Jensen during his talk.



The IEEE Medal of Honor and many other IEEE medals will be presented during the IEEE Honors Event April 24, 2026 in New York City.

Nvidia's Rubin AI architecture will drive even higher demand for storage and memory, bringing continued shortages and higher prices into 2026. Jensen Huang is the 2026 IEEE Medal of Honor winner.


15) Digital Storage and Memory Feeds AI at the 2026 CES, (Forbes Blog, 1/10/26)



The 2026 CES show is a great place to find out about the latest in technologies whether applied to consumer products, data centers or other applications. There were interesting introductions, such as Nvidia's Inference Context Memory Storage platform announcement using their BlueField 4, which we will write about separately. Here are some of the other digital storage and memory related things I uncovered at the CES show.


Samsung and LG spoke about how they are integrating AI into all of their consumer product lines. AI needs data for training and inference and digital storage and memory technologies are where that data lives. Let's look at storage and memory products from the CES.

I visited Kioxia in the Venetian Conference Center meeting room. They were showing NAND flash memory and SSD products for automotive, consumer and industrial applications. This included their high-speed automotive UFS 4.1 products. They also showed their line of DRAM-less BG7 client generation 8 3D-flash products, type 2242, 2230 and 2280 GB, PCIe 4.0 and NVMe 2.0d compliant, with storage capacities up to 2,048 GB, shown below.



The BG7 SSDs provide random read/write up to 1M IOPS and sequential read performance up to 7GB/s, a 10% and 16% improvement over the BG6 products. BG7 SSDs have a 67% improvement in sequential write performance compared to BG6. These SSDs use Kioxia's CMOS directly bonded to array, CBA, technology.

Kioxia also had on display a generation 10 BiCS NAND wafer, shown below with 332 layers with even higher read/write performance and power efficiency than the prior generation 8 products.



SanDisk was also at the show. I saw them at the Pepcom event for CES press and industry analyst attendees. Sandisk has rebranded the former WD_BLACK and WD Blue NVMe SSDs for gamers, creators and professionals as Sandisk Optimus. The internal storage solutions include Optimus for content creators, Optimus GX for gamers and Optimus GX Pro for developers, professionals and AI PC, workstations and high-end systems. These products can be seen in the photo from the Pepcom Digital Experience below.



Phison, a manufacturer of SSD controllers and its own line of SSDs was promoting its E37T PCIe Gen5 controller and client products using this controller. This controller provides up to 4.8GT/s performance. It is a DRAM-less design targeted at notebooks, handheld devices and other compact systems that use the M.2 2280, 2242 or 2230 form factor. The Phison E27T and E28 PCIe Gen5 controller are shown in the image below.



I visited Miin Chyou Wu, CEO of flash memory manufacturer Macronix at the show. He told me that the company is still working on its 3D NOR flash and there were

posters in their suite about their eMMC memory and serial NOR flash shown below.



Other World Computing, OWC, made two announcements at the 2026 CES. These included their 8TB Envoy Ultra Thunderbolt 5 SSD and their ThunderBlade X12 Expansion providing up to 192TB of Thunderbolt 5 RAID storage. The later product is shown in the image below.



The ThunderBlade X12 provides up to 6.6GB/s peak and 5.99GB/s sustained write speeds. It has twelve 16TB TLC NVME M.2 SSDs and uses OWC's SoftRAID. These products are geared towards media and entertainment professionals.

In addition to its own products, OWC hosted Strada's peer-to-peer collaboration platform for video content professionals, allowing collaboration with local drives anywhere in the world. OWC has invested in the Strada technology.

Some other storage products at the 2026 CDES included UGREEN NAS equipment at Showstoppers and an interesting reintroduction of Conner as a storage brand. The original Conner Peripherals was a Colorado-based company started by Finis Conner, which developed the first 3.5-inch form factor HDDs. The Conner logo was on a number of portable NAND flash products from a new Conner, like the image shown below.




Conner also was showing small external SSDs, card readers for MicroSD cards and hybrid devices combining storage with battery charging.

Although I didn't visit them at the show, SK hynix was showing 16 high DRAM die HVM4 48GB products at the CES. There were other companies, including VAST and Micron, who I wasn't able to meet at the show but will talk with afterwards to get the latest news.

Besides NAND and NOR flash products, some of the new non-volatile memories were also on display at the 2026 CES, including products from NXP and Ambiq. At the NXP pavilion outside the Central Hall of the Las Vegas Convention Center I was able to see how their S32KS 16nm automotive 3-MRAM MCU is being used

to enable various levels of vehicle control integration. [Thomas Jew from NXP spoke about this at the 2025 MRAM Forum last January.](#)



I also visited Ambiq and talked with one of its founders, Scott Hansen, at the 2026 CES. Ambiq has been using new non-volatile memory technologies in their consumer-oriented and industrial AI trained SoC products. At the show they were showcasing the introduction of a Neural Processing Unit, NPU, for real-time AI at the edge. This SoC uses their subthreshold power optimized technology, SPOT. The NPU is an Arm Ethos-U85 and its enabled computer vision, multilingual speech recognition and sensory models on the device.

The company had demonstrations in their suite of many applications using their products including smart cameras, AR glasses and other wearables as well as industrial edge and robotics. They also discussed their partnerships with Bravechip for smart rings and with Ronds for intelligent industrial monitoring solutions. I was told that their 22nm products use up to 4MB of embedded MRAM from TSMC but that their smaller lithographic products will be using TSMC's RRAM technology for on-board non-volatile memory.


I also met with resistive RAM developer, Weebit, at the CES. They spoke to me about their recent licensing agreement with Texas Instruments, TI, for the use of their ReRAM technology in TI's advanced process nodes for embedded processing semiconductors. The agreement includes IP licensing, technology transfer, design and qualification of Weebit ReRAM in TI's process technology. We also discussed their steady progress in their technology development and

demonstrations over the last few years. They are certainly a new non-volatile memory company to keep an eye on.

The 2026 CES featured lots of digital storage and memory technologies including SSDs, NAS products, DRAM and HBM as well as new non-volatile memory embedded products.

16) Inference Context Memory Storage Platform could Drive Additional 100EB of AI Storage, (Forbes Blog, 1/16/26)

At the 2026 CES Jensen Huang made many announcements associated with the introduction of Nvidia's Rubin Vera AI architecture and the various open-source AI foundational models that the company created for various applications, including Alpamayo for automotive. One interesting announcement from NVIDIA was about an Inference Context Memory Storage Platform or ICMSP, using the Nvidia BlueField-4 Network intelligent Interface Card or NIC. This is described as [a new kind of AI-native storage infrastructure designed for gigascale inference in order to accelerate and scale agentic Ai](#). An image of the BlueField-4 NIC is shown below.



So, what is context memory? During an interaction with an AI system data is generated about the interaction with the AI. This is the context of that interaction. If that context information is saved it can make future interactions with the AI more consistent, coherent and personalized. It allows AI to remember details across conversations, learn unique patterns, and understand complex, multi-turn interactions by storing relevant data beyond the immediate prompt. By storing this context information in long term storage this context can be retained for future interactions.

Besides the additional usefulness of context memory storage, it can also reduce the new calculation requirements by an AI system for individual queries, since data is recovered from storage rather than regenerated or kept in expensive and limited HBM, thus saving energy and allowing more efficient use of GPU processing and memory by using the context from prior interactions. This data is in the form of key-value, or KV, cache.

I spoke with Kevin Deierling, an executive from Nvidia, about the BlueField smart ethernet NIC, or digital processing unit, DPU. He told me that the ICMSp is a network storage system that can consist of SSDs and/or HDDs for storing and accessing the context memory data. It thus comprises a new tier of storage between traditional enterprise storage and the HBM DRAM, which holds the data being processed by GPUs.


The Nvidia ICMSp will store 16TB of storage per GPU and this can enable petabytes of shared context across a GPU cluster to support very large workloads. Throughput is targeted at 800Gb/s through the BlueField-4 board. The ICMSp retains an interesting form of data to be stored in that, if needed, it can be regenerated, unlike the data typically stored in enterprise storage systems. This means that traditional data retention requirements, such as redundancy, can be relaxed and still meet the needs for context memory storage. Thus a 4-9's reliability might be acceptable versus 9-9's required in conventional enterprise storage.

NVIDIA said that ICMSp products from AIC, Cloudian, DDN, Dell Technologies, HPE, Hitachi Vantara, IBM, Nutanix, Pure Storage, Supermicro, VAST Data and WEKA will be available in the second half of 2026.

In addition to Kevin, I also spoke with Phil Manetz, VP of GTM Execution at VAST and Jeremy Werner SVP & GM Core Data Center Business Unit at Micron about their plans and observations about memory, storage and the ICMSp. Phil spoke about the shortages in all types of memory and storage this year.

He also pointed out that 16TB per GPU in an ICMSp could easily result in an additional 100 Exabytes of context memory data being stored, putting additional requirements for storage, particularly solid-state NAND storage. He said that the

ICMSP could provide a premium inference experience for customers. The image below shows a conceptual drawing of VAST's implementation of an ICMSP.



VAST allows adding policies in their system, such as providing premium user experiences. Phil said that VAST has an advantage for this types of storage through their use of erasure coding with an overhead ratio of only about 3% with n+4 redundancy. They also have an extensive capability for data reduction where only the differences in otherwise very similar files are stored.

In addition to their ICMSP VAST also has what they call a flash reclamation program that can reuse a company's existing SSD storage using VAST and they have a soft launch currently with this offering.

Jeremy Werner from Micron said that this is a very good time to be in the memory and storage business and spoke about the company's investments in additional production capacity in Boise and New York state. This should result in 3.6M square feet of DRAM fabrication in the US. He said that the storage and memory hierarchy is getting more specialized layers and spoke also about trends for memory disaggregation.

He said that the company's Gen 6 NVMe SSD is in qualification and that at the 2025 Super Compute Conference they demonstrated 230M IOPS in a single storage server and that the company is working on additional innovations like [Storage Next](#).


Micron is also looking at the large amount of storage required for context KV cache storage. He mentioned the company's 245TB E3L form factor SSDs that it is introducing this year. He foresaw overall DRAM supply growth in the high teens to 20% this year but that demand will be much higher in 2026, and this will lead to higher prices and some constraints on AI data center buildouts.

Nvidia's inference context memory storage initiative based upon the BlueField-4 DPU will drive even greater demand for storage to support higher quality and more efficient AI inference experience.

17) Digital Storage and Memory Projections for 2026, Part 3, (Forbes Blog, 1/21/26)

In late 2024 we did [a piece on developments in optical digital storage and expectations for the following year](#). This piece continues this, looking at developments with the many optical storage startups as well as two DNA storage startup focusing on providing new options for digital archiving and data preservation.


As stated in the 2024 article, data centers often use a combination of different digital storage technologies to optimize performance requirements and the costs of storing information. This leads to the use of various types of digital storage technology trading off cost versus performance. The chart below, from my colleague, Jim Handy, of Objective Analysis shows a log-log plot of storage performance versus capacity price for various memory and digital storage technologies.



In archival storage, data is kept that still has potential value but which is very seldom accessed, for instance for various types of legal compliance or for historical value. This layer generally requires lower costs for storage than warmer storage layers and this is the realm where magnetic tape and optical disc storage have their greatest use and where new archival storage startups could find a home,

especially with the expected growth in archived data, including raw data, to support AI training and inference.

In August 2025 John Monroe and Brad Jones, long time storage industry market researchers wrote a [white paper projecting enormous active archive digital storage](#). In their projections out to 2050 they project significant growth in digital storage shipments with some emerging technology dominating shipping storage capacity by that time, as shown in the figure below. For comparison I note that Coughlin Associates projections are for about 8.3ZB (8,300 EB) of shipping storage capacity in 2030, up from about 2.5ZB of shipped digital storage in 2025.



Source: Furthur Market Research and Brad Johns Consulting (August 2025)

As this white paper points out, the vast majority of stored data is archived data. This then will drive demand for storage technologies to support this archival data growth. Magnetic tape and hard disk drives continue to grow in storage capacity and lower storage costs but as secondary and archive storage demand grows, there may be opportunities for new storage media to play a role, particularly if they can provide higher performance, lower cost, including ongoing energy costs, and greater longevity than current archive storage media. Depending upon developments, extensions of current or new storage technologies could dominate in the future.


Optical storage libraries are being sold as an alternative to magnetic tape storage and with the technologies being introduced by numerous startup optical storage startups, this use could increase. [Coughlin Associates wrote a white paper on current archive storage technologies, including these newer optical storage technologies](#). Let's look at the latest developments in optical storage from PiqIFilm, Folio Photonix, SPhotonix, Cerabyte and Optera Data.

There are several optical recording media, which include PiqIFilm that uses black and white, negative silver-halide film on a polyester base. This method is unique

in that it can include recorded human readable instructions and file format and source code for reading data, in addition to digital data. This archive method is being used in some scientific, engineering and historical archives. PiqlFilm is also the initiator of the Arctic World Archive (AWA), a repository for world memory located in Northern Norway on the Arctic Ocean.

Folio Photonics is a company which has developed a multilayer optical recording system for archiving applications which can result in photothermal recording similar to Blu-ray optical discs with either reflective or fluorescent recording technologies. Their technology uses photosensitive dyes dispersed in a co-extruded polymer matrix to create a reflective or fluorescent optical media. This photosensitive material has a strong optical absorption at 405nm, the wavelength used for conventional Blu-ray discs.


Microsoft and the University of Southampton explored volumetric archive recording in fused silica using a fast laser, starting several years ago. This approach is a write once method with a media that should be stable for 100's of years and was shown to store up to 360TB of data on a 5-inch glass plate. The Southampton researchers started a company, SPhotonix, to commercialize this technology. The image below shows the SPhotonix optical storage concept.



The SPhotonix 5D Memory Crystal was an [important element in the plot of the last Mission Impossible movie](#). The company received additional funding in November 2025 and I am acting as an advisor to SPhotonix. A 5D Memory Crystal containing a copy of Isaac Asimov's Foundation Trilogy was launched into orbit in 2018 during the SpaceX Falcon Heavy launch, in the glove compartment of the red Tesla Roadster. In June 2025 a 5D Memory Crystal containing images of the Chauvet Cave painting from 32,000 years ago were launched into earth's part of The Exploration Company's first Nyx space capsule. SPhotonix offers custom storage on its 5D Memory Crystals.

Another interesting optical recording contender is Cerabyte, founded in 2020 and with headquarters in Silicon Valley. Cerabyte uses sputtered 10nm ceramic layers on a glass substrate. Data is written encoded in an array of data matrices using a 2-D digital micro mirror with up to 2 million elements simultaneously written by femtosecond laser pulses in the UV spectrum with a write speed of 1GB/s and with less than 1W average power. Reading is done at GB/s data rates using high-

speed image sensors and parallel high speed image processing for decoding. The image below shows Cerabyte's digital recording concept.




Both reading and writing are done across the square substrate by scanning the microscope optics using high-speed XY stages kept in focus using a piezo driven auto focus system enabling random data access. The 9X9 cm media sheets can be recorded on both sides and stacked in cartridges for robotic access similar to that used for conventional optical and magnetic tape libraries. Cerabyte projects media costs below \$1/TB by 2030. Cerabyte has significant investments from several players in the digital storage industry, including Pure Storage and Western Digital.

Optera Data's storage technology takes advantage of changes in the optical absorption/emission characteristics of its recording media at several adjacent laser frequencies during writing and then by reading these changes. These changes are referred to as spectral holes. I have acted as an advisor for Optera Data.

They do this using an optical media which consists of a mixture of nano-particles with many particles lying within the write/read laser spot and the nano-particles have different but adjacent optical emission/absorption frequencies.

The combination of these nano-particles, in which spectral sensitivities are close to and partly overlapping each other, results in a combined spectral emission profile that is called a "top-hat" fluorescence emission profile, that is, the light emitted by the different nano-particles, combine together to make a pattern like that shown in (a) in the figure below.



With this media, tuning the write laser frequency to match that of one of the nanoparticle frequencies creates a spectral hole where the nanoparticle emission can be diminished as shown in (b) in the figure below. These spectral holes can have a depth that depends upon the level of laser energy during the spectral hole writing as shown in the figure. Data can be encoded in both the frequency and depth of these spectral holes.

Optera believes that, short term, 1TB discs are feasible with particulate media and medium term, a thin film single layer write once archival disc with high volume manufacturing costs of \$1/10TB (\$0.10/TB) is possible before the end of the decade. Longer term (say within a decade or so), these costs could be reduced even further, if this technology were implemented as a volumetric recording technology (perhaps even having 10X lower cost, \$0.01/TB).

Let's briefly look at some DNA storage developments in 2025. In May of 2025 [Atlas Data Storage was spun out from Twist Bioscience with \\$155M in seed funding](#). In November the company announced its Eon 100 Synthetic DNA storage. However, this approach appears far from a commercial product and DNA storage still suffers from slow write and read speeds. It is projected that because of human genomic medicine development the speeds for reading and writing DNA could increase dramatically over the next decade, making archival data storage using synthetic DNA commercially viable.

In January 2025 another synthetic DNA storage startup, Iridia (I am advisor to this company) contributed to [a lunar human monument with historical data on the company's synthetic DNA on board the Dragonfly Aerospace Blue Ghost Lander](#). Projections of new optical storage archiving systems promise 1PB raw capacity optical cartridges by the 2030's, compared to the possible 365TB raw capacity cartridge projections for LTO Gen 14. Optical storage and even DNA storage could be a significant contender for the digital archive market.

18) Digital Storage and Memory Projections for 2026, Part 4, (Forbes Blog, 2/2/26)


This is my fourth and last blog on digital storage and memory projections for 2026. The first two articles focused on digital storage and memory devices including magnetic tape, HDDs, SSDs as well as NAND, DRAM and emerging non-volatile memories. The third was an update on optical storage, since several optical storage startups targeting the archiving and digital preservation market have come to the forefront in recent years.

This article will focus on developments in storage systems and software and their use in workflows as well as additional insights on the future of storage devices and systems.

Demand for all forms of digital storage and memory exploded last year, mostly associated with the expansion of data centers to support AI deployments. 2025 saw shortages in all types of digital storage and memory that impacted all markets, data centers, enterprise, client and consumer markets. Many storage and memory companies have been leery to expand their production, concerned with the inventory build-up during Covid that resulted in the glut and severe price drops in 2022-2023.

Other companies, such as Micron and China's YMTC are expanding their production. At the same time Micron's Crucial Product family will no longer support BYI memory purposes as the company tried to maximize its production of higher margin memory such as HBM's to support AI workloads.

[IDC projected a 6.3% growth in external OEM storage systems in 2026 versus 2025](#) in October 2025, with spending at about \$37.7B, to build out smaller AI and other infrastructure. The figure below was from the October announcement.



In a [report in October 2028 IDC](#) said that organizations increased spending on compute and storage hardware infrastructure for AI by 166% year-over-year in the second quarter of 2025, reaching \$82B.

[In December 2025 IDC said that worldwide external enterprise storage systems market revenue increased 2.1% during the third quarter of 2025](#), reaching \$8B in vendor revenue. Revenue associated with All Flash Arrays showed more dynamic by growing 17.6% YoY while Hybrid Flash Arrays and Hard Disk Drives Arrays declined 9.8% and 6.3% respectively. AI deployments are the biggest driver for storage and memory growth.

But AI isn't just for the cloud, for many reasons, companies will want to keep data and some of its processing in house. Larry O'Conner, Founder and CEO of Other World Computing said that, "In 2026, more teams are going to rediscover the joy of having their data and workflows close to where the work actually happens. Not because the cloud is bad. The cloud is a great tool. It is just not the right answer for everything, especially when you are talking about performance, predictable costs, and keeping control of your own data."

With such a growing market for AI, it is natural to expect companies to develop particular approaches to address different market needs. Paul Aubrey, Director of Product Management at NetApp Instacluster said that, "The next frontier in AI/ML isn't about building bigger models, it's about making smaller ones work together. The rise of Model Context Protocol (MCP) and agentic frameworks will turn AI into a composable ecosystem of reusable, discoverable micro-agents. Organizations will deploy fleets of ML models, each powering specialized classification, prediction, and recommendation tasks, each behind MCP endpoints that plug directly into the agent mesh."

The growth of AI has led to growing needs for compute and networking but also storage and memory. Mark Seamans, SVP of Global Marketing at Penguin Solutions said, “The demand for GPUs and compute has been covered extensively, but 2026 will be the year that the industry reckons with the fact that compute horsepower is not the critical gating factor in AI infrastructure success. Rather, it will be memory and specialty components that determine timelines for AI deployments. With numerous components in high demand and facing significant backlog, there will again be heightened emphasis on supply chain.”

Furthermore, more efficient use of memory could grow, stimulated by the shortages due to AI workloads. Andy Mills, VP of Advanced Product Development at Penguin said that, “Emerging technologies that enable disaggregation of memory from compute, such as CXL or advances in OMA-based optical networking, will see increased interest and investment.”

Growth in AI workloads has led to both shortages in digital storage and memory, and other technologies to support that growth, but also increasing projections for longer term storage of raw data used for training as well as logs and other data that can be used to retrain and understand AI models, for instance in liability and compliance cases. But does that mean we need to keep all the data?

Eric Polet, Director of Product Marketing at Arcitecta said that there are costs that need to be considered in deciding what to keep and what to delete, “The next decade of data management will be defined by data sustainability. Organizations will increasingly treat storage growth as an environmental liability, not just a cost center, and will demand tools that assess the energy footprint embedded in every byte. New storage technologies, such as glass, DNA, ceramic, etc., will help, but the real gains will come from smarter curation, not ever-expanding capacity. Institutions that succeed will strike a balance between retention and restraint, ensuring that their data footprint is not only useful but also sustainable. Sustainable storage doesn’t mean deleting everything. It means knowing what can be deleted and building policies and lifecycles around that knowledge.”

The increasing use of inference will increase the need for new types of storage. [At the 2026 CES Jensen Huang from Nvidia announced the Inference Context Memory Storage Platform](#), a Key-Value store supported by a number of enterprise storage hardware manufacturers and based upon Nvidia’s BlueField-4 NIC. Retaining coherence between AI inference sessions allows more accurate results as well as faster and lower cost inference. This creates a new storage tier as well as demand for higher performance network bandwidth.

While AI workloads have boosted demand for all-flash storage, this demand has resulted in higher SSD storage prices, changing the financial dynamics versus HDDs. Gal Naor, CEO at StorOne said, “Flash pricing has surged dramatically since 2024, and this structural shift is expected to become the primary influence on storage planning in 2026. The cost gap between flash and HDD continues to

widen, now reaching ratios of 1:10 or higher forcing organizations to reconsider 'All-Flash' strategies that no longer align with financial reality... This new economic environment is driving a transition toward intelligently tiered hybrid architectures. Flash must be reserved for genuine Tier-1 workloads delivering high performance where it is actually needed, while cold and infrequently accessed data moves automatically to lower-cost media."

David Mosely, CEO of Seagate Technology said that, "AI is reshaping hard drive demand by elevating the economic value of data and data storage. AI inferencing, for example, consumes and generates large volumes of data. AI-generated video content (from social media platforms to content delivery networks and online marketing) is a major source of unstructured data and is driving considerable demand for hard drives."

It's not only memory, SSDs and HDDs that are benefiting from AI storage and memory demand, so is magnetic tape as a cost effective, lower energy storage alternative. Skip Levens, Product Leader and AI Strategist at Quantum said that, "In 2026 AI and HPC data center buildouts will hit a non-negotiable limit: they cannot get more power into their data centers. Build-outs and expansions are on hold and power-hungry GPU-dense servers are forcing organizations to make hard choices. Yet modern tape libraries can be the smart pivot to solve two pressing problems - first by returning as much as 75% of power to the power budget to 'spend' on GPUs and servers, while also keeping massive data sets nearby on highly efficient and reliable tape technology. So, the next wave of AI scale will get lean to get stronger - and tape will make it possible."

2026 promises increased demand for memory, storage devices, systems and software to support the growth of AI data processing. This will drive the growth of all types of memory and storage technology.

19) C4Q 2025 and 2025 HDD Industry Update, (Forbes Blog, 2/3/26)

Data is in for the fourth calendar quarter 2025 HDD industry from the February 2026 [Digital Storage Technology Newsletter](#).

Seagate and Western Digital are no longer providing guidance on their total HDD unit shipments or data to provide average sales prices directly, but based upon our estimates total HDD unit shipments in the CQ4 2025 are 33.4M, up 4.8% from our prior quarter estimate.

Total exabytes shipped in CQ4 2025 were up 9.1% compared with CQ3 2025 (456.9EB versus 418.7EB). HDD revenues (estimated at about \$6.6B) were up about 7.5% from the prior quarter. We estimate about 22.5M nearline HDDs were shipped in the last quarter, up 5.9% from the prior quarter.

Total HDD unit shipments in 2025 are estimated at 124.3M, up 0.4% from the prior year. Estimated nearline HDD shipments were 80.8M, up 13.6% from the prior

year. Total exabyte shipments for the year were 1,620.95EB (1.61ZB), up 22% from the prior year. Total industry revenue in 2025 were \$23,267M (\$23.3B), up about 18% from the prior year.

Total HDD shipments in C2Q 2025 were up about 6.0% compared with C1Q 2025 (30.4M versus 28.7M). Total exabytes shipped in C1Q 2025 were up about 12.6% from C1Q 2025. HDD revenues (estimated at \$6.2B) were up about 17.4% from the prior quarter. Nearline HDD shipments were up 13% from the prior quarter.

Seagate said their overall revenue was \$2,825M for CQ4 2025. Of that, we estimate that non-HDD revenue was about \$169M, slightly more than in CQ3 2025 (\$166M). We thus estimate Seagate's total HDD revenue at \$2,656 with data center revenues representing 79% of overall Seagate revenue for the quarter (probably combined system and HDD data center revenue). Total storage capacity shipped was 190EB with 165EB for Nearline HDDs (86.8%), the balance of 25EB was shipped for non-nearline applications. The average HDD capacity is estimated at 14.0TB. The company shipped about 14M units in the quarter with an ASP I estimate at \$195.28.

For the year Seagate total revenue was \$9.4B. Total storage capacity shipped was 678.1EB with an estimated 50.6M HDDs shipped.

Western Digital said that their overall revenue was \$3,017M in CQ4 2025 and that cloud revenue was \$2,673M. They didn't break out HDD only revenue, but we assume about \$101.6M for non-HDD revenue. They also said that they shipped 215B of total HDD storage in CQ3 2025 with 192EB (89.3% of the total) shipped for nearline applications. 23EB were shipped for non-nearline applications. WDC gave little guidance on the total number of HDDs shipped, or of nearline HDD unit shipments, only saying that they had shipped 3.5M units of their latest generation ePMR HDDs. We estimate that the company shipped about 14.1M units in the CQ4 2025 quarter. Likewise we estimate that their HDD ASP was about \$206.8.

For the year Western Digital total revenue was \$10.3B. Total storage capacity shipped was 761EB with and estimated 52.6M HDDs shipped.


For 2025 Toshiba total HDD revenue is estimated at \$3.5B. Total storage capacity shipped was 51.86EB with 21.1M HDD shipped.

Although, neither Seagate or Western Digital provided average sales price data, ASP, we estimate a 3.3% average sales price increase or HDDs from C3Q 2025 to C4Q 2025.

We saw a 8.4% average sales price (ASP) increase for HDDs from C3Q 2025 to C4Q 2025. HDD ASPs are now at levels about 25% higher than seen at the end of 1998 due to the increased demand for high capacity, high ASP, nearline HDDs and the increase in the percentage of these drives shipped as shown below.




The pie chart below shows 2025 total exabyte shipment market share by company to the left and total revenue shipment estimates by company to the right.



Unit shipment market share estimates are 42.3% for WDC, 40.8% for Seagate and 17% for Toshiba.


We note that the growth of capacity-oriented nearline drives for enterprise and hyperscale applications with more heads and media per drive and the move to 11-disk and 12-disk HDDs will result in more components per drive and thus higher head and media demand than in the past. We estimate that this could lead to new investments in head manufacturing capacity by 2027, probably including integrated laser fabrication for HAMR production expansion.

AI and other data intensive workloads are expected to increase shipping storage capacity demand in 2026 and into the future. Our latest projections for HDD capacity shipments out to 2030 are shown below.



HDD unit shipments, exabyte shipments and revenue were up significantly in C4Q 2025 and in 2025 as a whole. Continued growth is expected in 2026.


20) Western Digital Innovation Event—Storage is where AI goes to Learn, 2/3/26)




Iving Tan, WD CEO, with new WD Branding in background

So said Ahmed Shihab, WD Chief Product Officer during the 2026 Western Digital Innovation Day event on February 3, 2026. WD gave guidance for considerable growth through 2030 driven by a strong digital storage market and the company's technology advances to make higher capacity, higher performance and lower energy consumption HDDs to fill multiple data center storage tiers. This article will explore the technical developments that WD presented at their innovation event.

In Ahmed's presentation he referred to digital storage tiers as an analogy to different types of printed content storage in a physical library as shown below. Data centers currently store more than 80% of their content on HDDs and WD believes that HDDs will continue to provide the majority of data center data.




WD also spoke about the different types of digital storage required to support AI workloads with the visual shown below. This image shows HDDs filling most of these AI storage requirements and this tied into the capabilities of the company's technology roadmap.



Below is Western Digital's technology roadmap out to 2032. The company has aggressive plans for increasing HDD storage capacity with its current ePMR as


well as HAMR technologies and with higher disk counts, a couple of ways they plan to introduce higher performance HDDs, introduction of a lower power HDD and also storage platform developments. We will explore all of these developments below.



WD has two hyperscale customers qualifying their next generation HAMR drives and as shown in the roadmap plans to offer both ePMR and HAMR drives at similar capacity points. The image below is of a 40TB ePMR, Ultrastar HDD as well as a 40TB HAMR HDD that they were showing.




By some areal density increases as well as higher disk counts the company plans to continue its ePMR HDDs out to 60TB before moving to only HAMR HDDs for higher capacities. WD said that this provided their customers with a way to gracefully transition from the current ePMR products to HAMR products. The roadmap shows 100+TB HDDs by 2029. The slide below shows how the company plans to do this.



WD's recent acquisition of IP and personal to create a smaller footprint vertical emitting integrated laser technology that are more easily tested to improve drive yields will assist with this roadmap. The company is also working to increase areal density, likely mostly with the introduction of HAMR. With a move to glass substrate disks for HAMR, disk thicknesses can decrease and with smaller heads, WD plans to increase the number of disks per HDD to 14 from their current 11 disks. 14 disks each with 10TB per disk gives a 140TB HDD.

I note that Seagate is currently shipping 10-disk HAMR HDDs with 40+TB capacity. Higher disk count means more disks and heads and could require higher bill of materials cost, but with the continuation of ePMR with less expensive aluminum substrates, higher disk count BOM costs could be moderated compared to a glass substrate HAMR drive. This could be one reason to continue to ship ePMR drives for a couple more years.

In addition to higher areal densities and higher platter counts WD plans to take two approaches to increase HDD performance. The first approach is to have more than one HDD head and channel active at a time. The images WD showed indicate that this is being done with pairs of heads on the opposite sides of disks. As shown in the slide below, using two heads at once for writing and reading will give a 2X bandwidth increase. By going to 4, 6 and 8 active heads/channels the HDD bandwidth can be increased by 4, 6 and 8 X higher than with the single head/channel on conventional HDDs.




This approach can saturate interconnects and match the capabilities of SSDs for some applications. However, data rates writing and reading on HDD tracks is one factor in HDD performance. Another factor is the time it takes for the heads to move from one track to another. This is the seek time. WD plans to reduce the seek time by late 2027 into 2028, from the roadmap, by moving to dual actuator HDDs.

The image below from the Innovation Day shows two separate actuator arms for addressing different sets of disks. This is different from prior split actuator HDD designs that were previously shown by WD and Seagate. This general approach was used in the early 1990's in Conner Peripherals Chinook HDD. WD says that this will be a cost-effective approach to reduce HDD seek time.



WD also plans to introduce higher capacity and lower HDDs for colder storage tiers with lower rotation rate products, their power optimized products, by 2028. By slowing the rotation rate data access performance slows but it will use less energy and it allows about a 10% increase in the storage capacity.


In addition to providing higher capacity, higher performance HDDs, WD also presented its plan for a more intelligent storage platform incorporating flash memory as well as HDD with a single open API, as shown below.



Western Digital's Innovation Day showed their plans to introduce 100+TB HAMR HDDs with 2X or higher performance by 2029 as well as ways to fill multiple AI data center storage tiers.

GUEST ARTICLES

1) 100M IOPS SSDs? You Must be Kidding, Jim Handy, Objective Analysis. Originally published in the SSD Guy blog 2/2/26




Back in August, at the [FMS show](#), keynoters surprised their audiences with promises of upcoming 100 million IOPS SSDs. Given that today's highest-performing SSDs provide well under ten million IOPS, how will SSD makers reach this very challenging goal?

We can gain some insight into this from a little information in a paper published this past November for OCP 2025 by researchers at [ScaleFlux](#), [Nvidia](#), and

Stanford University: [From Minutes to Seconds: Redefining the Five-Minute Rule for AI-Era Memory Hierarchies](#).

The paper illustrates that NAND can come closer to being used as a memory layer than as a storage layer in [the memory-storage hierarchy](#) when the SSD's NAND is managed differently than is done in modern SSDs. Part of this is due to the fact that certain widely-accepted beliefs might need reconsidering, and the rest comes from fundamental architectural changes to the SSDs.



Reconsidering Widely-Accepted Beliefs

The researchers challenge of one widely-accepted belief is reflected in the title of their paper. For those who haven't run across it, [Jim Gray's 5-minute rule](#) was first devised in the 1980s and has seemed to stand the test of time ever since. People run updated numbers through Gray's formula, using HDDs, SSDs, and various

sorts of memory, and it always seems to produce a result similar to that found in the first paper.

Gray's formula is a very simple calculation, based on hard drive (HDD) and DRAM prices and datasheet maximum performance, and it finds that you should keep data in memory (DRAM) if it's going to be needed within 5 minutes, but data that's unused for 5 minutes or longer should be moved out of memory and into storage. The authors of the new paper noticed that Gray's equation may have been just a little bit too simple, and they added in several other important factors, including wasted server resources and the typical behavior of SSDs and DRAM, and these shed a completely new light on the matter. It's an interesting read.

Two other widely-accepted beliefs are that memory is memory and storage is storage, and that hardware and software should follow established delineations and cannot be tuned to work together as well as possible.

The new paper starts with the assumption that the researchers should consider redesigning anything that can be redesigned. This includes hardware as well as software, and the delineation between memory and storage, and it doesn't simply assume that today's SSD designs are the optimum for AI applications. Instead, they decided to use hardware/software co-design to optimize both sides to the benefit of the overall system, while employing a dual-resolution data structure to provide approximate results initially, on a data subset, before refining the results by evaluating the full-resolution values of only the most likely candidates. Finally, they chose to not really think of SSDs as storage, but instead consider them a slower alternative to DRAM, using flash to provide cheaper fast data access than DRAM could provide.

A New SSD Architecture


This approach made the researchers to explore ways to make SSDs with faster access, so that its performance was more closely matched to the DRAM it replaced.

They decided not to limit their exploration to today's SSDs, but chose to look at another architecture, the Storage-Next SSD that is currently being co-developed by Nvidia, SSD suppliers, and other companies [including Marvell](#). This SSD uses a design that is optimized for smaller transfers than the 4k-byte blocks standard to SSDs, since a new algorithm shrinks the size of the data that the GPU manipulates to reduce the computational workload.

Marvell explains that today's SSDs don't scale IOPS for data transfers smaller than 4KB, and this causes the PCIe bus to be underutilized, starving the GPU and wasting cycles. The current approach to this problem is to deploy a number of SSDs in parallel, which increases both the system's cost and power, but even this approach has difficulty meeting performance targets.

Marvell explains that the limited performance of 4kB-optimized SSDs is mainly the result of their error correction algorithm. One of the Storage-Next SSD's architectural changes is to optimize error correction for smaller data transfers.

In a recent presentation Nvidia showed a conceptual chart to illustrate how a Storage-Next SSD would perform in comparison to the industry's current highest-IOPS SSD:



Algorithm and data structure co-design for IOPS efficiency

Note the caption: “Algorithm and data structure co-design for IOPS efficiency.” this co-design involves using small 512B representations of larger pieces of data, and performing initial identification of those, choosing which are closest to the target value. Once this step is complete then the processor performs a match on much wider data, a 6KB original vector, but only for those records that passed the first screening. Today's approach compares every 6KB record using their full data width. While the new approach requires two steps, one for the 512B mini-representations of the data and one for a subset of 6KB original vectors, the net result is significantly faster processing using only the items that are the most likely fit, and ignoring all the rest.

In short: All the software must be reconfigured, and the matching process must be performed twice, once at low precision and once at full precision. By doing this, and by using an SSD optimized for 512B transfers, the cost/performance of the system can be dramatically improved.

I should note that other factors also help. Part of the Storage-Next SSD's speed advantage comes from the fact that it will use the PCIe 7.0 interface. The rest is provided by a re-architecture of the SSD, including the error correction for smaller transfers.

There is also talk of using speed-optimized NAND flash chips. This is quite a departure from the industry's current direction, since NAND chips are nearly always cost-optimized at the expense of speed. Some of the leading NAND flash makers have demonstrated speed-optimized flash chips for the past decade or more, particularly Samsung with its Z-NAND and SK hynix's 4D NAND. Two tricks these chips use is that they are based on SLC flash, which is significantly faster than MLC, TLC, or QLC flash (see this post's first chart), and these special SLC chips have a larger number of internal planes. Planes are like miniature flash chips within a single flash chip, and they improve performance, but increase the die size

and therefore cost. The trick is to find the right balance of cost/performance to optimize the total system's cost/performance. I'll leave that to the data center designers.


One of [Objective Analysis'](#) strengths is to understand and identify potential winning new technologies before they become big. The Storage-Next SSD is an example of a technology that has the potential to significantly change the SSD market. If you'd like to benefit from our understanding of new technologies please contact us to explore ways that we can work together to strengthen your corporate strategy.

2) CXL is Coming—Don't Get Left Behind, Jim Handy, Objective Analysis, 2/9/26

Although it has been overshadowed by recent explosive HBM sales, Compute Express Link, or CXL, remains significantly important for both the DRAM market and for hyperscale computing systems. CXL brings computer systems a very broad range of features aimed at enabling computing architectures that were previously hard to implement. Objective Analysis explores this technology and its potential in a new report titled: "CXL Looks for the Perfect Home."

CXL is here now, and is ready to help system architects go beyond DDR's speed issues to achieve significantly better memory bandwidth while decreasing power consumption. Those who understand CXL's technical and business outlook will be able to benefit greatly from its adoption. The report provides this understanding with its thorough, yet clear, technology and market analysis.


Figure 4. Sharing Multiple Memory Types via a CXL 2.0 Switch



Source: Objective Analysis

This report explains how this emerging technology will enjoy rapid growth to nearly \$3.4B by 2028 as computing systems are configured to take advantage of its many features. It explores important questions about this budding new technology, while it exhaustively examines CXL technology and its market to give an inside view of how the technology will grow, and what must be done to allow this growth to take place.

Figure 23. CXL Memory Module Revenue and Exabyte Shipment Forecast



The 60-page report provides the basis for a very solid grasp of CXL and the way that it is likely to be used, as well as its impact on the markets that surround it. This strong background leads to a very reasonable forecast based on field interviews of CXL memory module suppliers and users, and provides clear strategies for successful participation in this business. With its 27 figures and eight tables this report is essential reading to anyone whose business may be impacted by CXL.

To create this study Objective Analysis performed thorough and exhaustive research into the technology and the events leading up to CXL's introduction, including vendor and user interviews, and briefings from standards bodies. Forecasts are based on the company's rigorous and unique forecasting methodology that has produced the most consistently-accurate forecasts in the semiconductor industry.

Objective Analysis' CXL report will be your foundation for a successful CXL strategy. The report can be purchased for immediate download at <http://Objective-Analysis.com/Reports.html#CXL>.

3) Internet Archives Storage, David Rosenthal,

Bruce Li introduces his *The Long Now of the Web: Inside the Internet Archive's Fight Against Forgetting*:

This report delves into the mechanics of the Internet Archive with the precision of a teardown. We will strip back the chassis to examine the custom-built PetaBox servers that heat the building without air conditioning. We will trace the evolution of the web crawlers—from the early tape-based dumps of Alexa Internet to the sophisticated browser-based bots of 2025. We will analyze the financial ledger of this non-profit giant, exploring how it survives on a budget that is a rounding error


for its Silicon Valley neighbors. And finally, we will look to the future, where the "Decentralized Web" (DWeb) promises to fragment the Archive into a million pieces to ensure it can never be destroyed.

It is long, detailed, comprehensive and well worth reading in full. Below the fold I comment on the part about storage.



Internet Archive 10/29/96

This picture shows Brewster Kahle with the Internet Archive's very first storage, tape drives holding crawls from [Alexa Internet](#), the Web traffic analysis company founded by Brewster and Bruce Gilliat in 1996 and, three years later, acquired by Amazon for \$250M.



Source

The first custom-designed Internet Archive storage system was version 1 of the PetaBox, a vividly red-painted 1U system introduced in 2004. This picture shows a rack of them. I'm familiar with them because the LOCKSS Program used them for some years.

They were great! Li explains:

The first PetaBox rack, operational in June 2004, was a revelation in storage density. It held 100 terabytes (TB) of data—a massive sum at the time—while consuming only about 6 kilowatts of power. To put that in perspective, in 2003, the entire Wayback Machine was growing at a rate of just 12 terabytes per month. The LOCKSS Program eventually transitioned to 4U systems similar to those then in use at Backblaze. Around the same time the Archive also moved to 4U systems for their PetaBox.

Specification	Generation 1 (2004)	Generation 4 (2010)	Current Generation (2024-2025)
Capacity per Rack	100 TB	400 TB	~1.4 PB (1,400 TB)
Drive Count	~40-80 drives	240 drives (2TB each)	~360+ drives (8TB+ each)
Power per Rack	6 kW	~6-8 kW	~6-8 kW
Heat Dissipation	Utilized for building heat	Utilized for building heat	Utilized for building heat
Processor Arch	Low-voltage VIA C3	Intel Xeon E7-8870 (10-core)	Modern High-Efficiency x86
Cooling	Passive / Fan-assisted	Passive / Fan-assisted	Passive / Fan-assisted

PetaBox specifications

As disk technology progressed the Archive continued evolving the system:

The fourth-generation PetaBox, introduced around 2010, exemplified this density. Each rack contained 240 disks of 2 terabytes each, organized into 4U high rack mounts. These units were powered by Intel Xeon processors (specifically the E7-8870 series in later upgrades) with 12 gigabytes of RAM. The architecture relied on bonding pair of 1-gigabit interfaces to create a 2-gigabit pipe, feeding into a rack switch with a 10-gigabit uplink.



Brewster Kahle & PetaBoxes

Over time older, smaller drives were phased out and replaced by newer, bigger drives:

By 2025, the storage landscape had shifted again. The current PetaBox racks provide 1.4 petabytes of storage per rack. This leap is achieved not by adding more slots, but by utilizing significantly larger drives—8TB, 16TB, and even 22TB drives are now standard. In 2016, the Archive managed around 20,000 individual disk drives. Remarkably, even as storage capacity tripled between 2012 and 2016, the total count of drives remained relatively constant due to these density improvements.

This is indeed an impressive story and Li writes:

The trajectory of the PetaBox is a case study in Moore's Law applied to magnetic storage.

Here I have to correct Li. Moore's Law applies to silicon such as solid-state storage; it is *Kryder's Law* that applies to magnetic storage such as the hard disks in the PetaBox.

Li notes the Archive's innovative approach to cooling the racks:

The Archive's primary data center is located in the Richmond District of San Francisco, a neighborhood known for its perpetual fog and cool maritime climate. The building utilizes this ambient air for cooling. There is no traditional air conditioning in the PetaBox machine rooms. Instead, the servers are designed to run at slightly higher operational temperatures, and the excess heat generated by the spinning disks is captured and recirculated to heat the building during the damp San Francisco winters.

This "waste heat" system is a closed loop of efficiency. The 60+ kilowatts of heat energy produced by a storage cluster is not a byproduct to be eliminated but a resource to be harvested. This design choice dramatically lowers the Power Usage Effectiveness (PUE) ratio of the facility, allowing the Archive to spend its limited funds on hard drives rather than electricity bills.

In the unlikely, for San Francisco, event that the day is too hot, less-urgent tasks can be delayed, or some of the racks can have their clock rate reduced, disks put into sleep mode, or even be powered down. Redundancy means that the data will be available elsewhere.

Brian Wilson, CTO of BackBlaze pointed out back in 2014 ago that in their long-term storage environment "Double the reliability is only worth 1/10th of 1 percent cost increase". And thus that the moral of the story was "design for failure and buy the cheapest components you can".

The Archive has followed this advice for a long time, as Li explains:

With over 28,000 spinning disks in operation, drive failure is a statistical certainty...

The PetaBox software is designed to be fault-tolerant. Data is mirrored across multiple machines, often in different physical locations (including data centers in Redwood City and Richmond, California, and copies in Europe and Canada).¹² Because the data is not "mission-critical" in the sense of a live banking transaction, the Archive can tolerate a certain number of dead drives in a node before physical maintenance is required.

Here is Brian Wilson's analysis:


Replacing one drive takes about 15 minutes of work. If we have 30,000 drives and 2 percent fail, it takes 150 hours to replace those. In other words, one employee for one month of 8 hour days. Getting the failure rate down to 1 percent means you save 2 weeks of employee salary - maybe \$5,000 total? The 30,000 drives costs you \$4m.

Back in 2006 I started asking *How Hard Is "A Petabyte for a Century"?* and concluding that it was *really* hard. By analogy between "bit rot" and radioactivity, I estimated that keeping a petabyte for a century with a 50% chance of no bit flips required a bit half-life a hundred million times the age of the universe. But that wasn't the fundamental problem:

The basic point I was making was that even if we ignore all the evidence that we can't, and assume that we could actually build a system reliable enough to preserve a petabyte for a century, we could not prove that we had done so. No matter how easy or hard you think a problem is, if it is impossible to prove that you have solved it, scepticism about proposed solutions is inevitable.

We have to assume that, despite their best efforts, the Archive will over time lose some data. This isn't a problem for two reasons:

- The Archive's collection is, fundamentally, a *sample* of cyberspace. Stuff missing is vastly more likely to be because it wasn't collected than because it was collected but subsequently lost.
- Driving the already low probability of loss down further gets exponentially more expensive. The more of the Archive's limited resources is devoted to doing so, the less can be devoted to collecting and storing more stuff. Thus devoting more to eliminating data loss results in *less* data surviving.



Since at least 2011 I have written repeatedly about the fact that preserving data for the long term is and economic, not a technical problem. With a lavish budget it is simple; the smaller the budget the harder it gets. Li correctly points out that the Archive's budget, in the range of \$25-30M/year, is vastly lower than any comparable website:

By owning its hardware, using the PetaBox high-density architecture, avoiding air conditioning costs, and using open-source software, the Archive achieves a storage cost efficiency that is orders of magnitude better than commercial cloud rates.

But in making this point he uses a *flawed comparison*:

Consider the cost of storing 100 petabytes on Amazon S3. At standard rates (~\$0.021 per GB per month), the storage alone would cost over \$2.1 million per month. The Internet Archive's entire annual operating budget—for staff, buildings, legal defense, and hardware—is less than what it would cost to store their data on AWS for a year.

First, Li ignores bandwidth charges, which as one of the most-visited sites on the Web would dwarf the storage charges. Second, Li compares S3, intended for and priced as instantly accessible, replicated data for the back-end of mission-critical online services, with the Archive's storage, intended for the long-term storage of data, most of which is accessed very rarely. Amazon has products aimed at this market. As an example, writing a Petabyte to Amazon Glacier, storing it for a decade, and then reading it out at last year's prices would cost just under \$160K. This too is a misleading comparison, Glacier is intended for *archival* data that has a much lower probability of access than the Archive's.

The key to making storage affordable is *tiering*, moving data that hasn't been accessed recently down the storage hierarchy to cheaper media. Kestutis Patiejunas gave an excellent talk about this back in 2014. I discussed it in *More on Facebook's "Cold Storage"*. Facebook was in the happy position of having an extremely accurate model of the access patterns to each of the small number of



How Will Your Total Available Market Change Over the Next Five Years?

A Forecast of the United States Business Economic Demographic Structure February 2026 Update

- [Executive Summary: Challenges](#)
- [BEDD: A Total Available Market History and Forecast 1998-2030](#)
- [DRG Forecasting Approach Amid Policy Uncertainty](#)
- [Key Forecast Assumptions](#)
- [Employment Trend 2025 Forecast](#)
- [Current Forecast](#)
- [Structural Changes](#)
- [Business Economic Demographic Forecast](#)
- [Informing Source](#)
- [Data Access and Licensing](#)
- [∞|Δ About Daniel Research Group](#)

Executive Summary: Challenges

These are the times that try men's souls." Today we face several related forecasting challenges.

Implemented and announced future economic disruptive US government policies regarding Tariffs, Deportations, and US Government spending and support services.

When will these policies be fully enacted and for how long.

Essentially, we are asking four questions

What has the administration announce or implied it will or may do?

What can they legally do?

What has been implemented?

What will be permanent and what may be reversed in the future?

Furthermore, Federal Government shutdowns severely limit the availability of real-time economic data.

BEDD

Business leaders require forecasts of their total available market, even in heightened uncertainty times to support strategic decisions concerning asset allocation and investments, produce and service offering, and market strategy. **BEDD (Business Economic Demographic Database)** is a Total Available Market (TAM) database containing the history from 1998 and forecast to 2029 of Employees, Businesses, Establishments and Payroll in the United States by Sector, Sub-Sector, Industry and Employee Company Size.

Census Bureau State and Local employment and payroll data by function (products or services provided) are mapped to industries that provide the same products or services. Federal Government employment and payroll data by agency or department is mapped to industries according to the products or services they provide. Public Administration and National Defense are "industries" unique to the public sector.

BEDD answer three essential questions:

1. **Who are the potential buyers and users?**
2. **How large is the market opportunity?**
3. **How will the market evolve over time?**

DRG Forecasting Approach Amid Policy Uncertainty

A broad consensus among leading economists suggests that the current administration's stated economic and domestic policy agenda could result in elevated inflation, rising unemployment, and a potential recession. However, the extent to which these outcomes materialize depends on two key variables: the administration's commitment to implementing its proposals, and the ability of legislative bodies, the judiciary, and market participants to moderate or counteract those initiatives. The forecast reflects current economic conditions and trends as reported from several authoritative sources, including:

- The Congressional Budget Office's (CBO) March 10-Year Employment Forecast
- The Bureau of Labor Statistics' (BLS) September 2024 10-Year Projections
- Historical data from the Census Bureau's SUSB database (1998–2022)
-

Note: in the absence of the latest JOB Situation Report from the Bureau of Labor Statistics due to the recent US Government shut down, DRG has derived industry level growth estimates from the number of recognized credible sources.

- Current employment trend data and statistics from the following sources
 - ADP
 - Indeed
 - Revelio Labs
 - Challenger, Gray & Christmas
- Department of Defense Fiscal Year 2026 Green Book
- The Budget Lab at Yale
- Current forecast and analysis from additional financial, labor, and academic sources focusing on:
 - GDP Growth
 - Employment Trends
 - Tariff Rates and Effects
 - Deportations
 - Reduction of Federal Government Spending, Services, and Employment
 - Artificial Intelligence

Key Forecast Assumptions

Tariffs

All tariffs in effect as of November 10, 2025, are assumed to remain unchanged throughout the forecast period.

Deportation Trends

An estimated 11 million undocumented immigrants currently reside in the United States, with approximately 7 million participating in the labor force across key industries.

Industry	Percent Undocumented Immigrant Employees
Forestry, Fishing, Hunting, Agriculture	12.7%
Mining	3.6%
Utilities	1.0%
Construction	13.7%
Manufacturing	5.4%
Wholesale	5.5%
Retail	3.9%
Transportation, Warehousing	5.5%
Information	2.5%
Finance, Insurance	2.7%
Real Estate, Rental, Leasing	2.5%
Professional, Scientific, Technical Services	4.7%
Management of Companies, Enterprises	2.5%
Administrative Support, Waste Management, Remediation Services	6.0%
Education	2.0%
Health Care, Social Assistance	2.0%
Arts, Entertainment, Recreation	4.0%
Accommodation, Food Services	7.1%
Other Services	6.5%

- In 2025, deportations are projected to reach 400,000 individuals.
- This figure is expected to rise steadily, reaching approximately 1 million annually by 2029.
- This trend may significantly impact labor availability in sectors reliant on undocumented workers.

Federal Employment, Spending, and Services

The ongoing decline in federal employment, budgetary outlays, and public services is expected to persist through 2026.

- This contraction is likely to produce negative multiplier effects across the broader economy, particularly in employment and small business formation.


Demographic Shifts

Structural demographic changes will continue to exert downward pressure on labor force participation:

- An aging population and accelerated retirement rates
 - Slower growth in new immigration
- These factors are expected to constrain workforce expansion and productivity.

Artificial Intelligence Adoption

While AI adoption in business operations is progressing more slowly than initially projected, the technology is still reshaping the labor market. The following forecast table and charts is based on data for the last published Census Bureau Business Trends and Outlook Survey.



PAI	Percent of Primary Firm Planning on Using AI to Produce Goods or Services									
NAICS	Industry	2023	2024	2025	2026	2027	2028	2029	2030	
11	Forestry, Fishing, Hunting, Agriculture	2.5%	3.5%	4.9%	6.8%	9.3%	12.7%	17.0%	22.5%	
21	Mining	1.7%	3.2%	6.0%	11.1%	19.5%	32.1%	47.9%	64.2%	
22	Utilities	3.1%	5.7%	10.0%	17.2%	27.8%	41.8%	57.1%	71.3%	
23	Construction	1.3%	2.5%	4.8%	9.2%	16.9%	28.9%	44.8%	61.9%	
31	Manufacturing	2.4%	4.1%	6.8%	11.1%	17.5%	26.7%	38.4%	51.6%	
42	Wholesale	2.3%	4.2%	7.7%	13.6%	22.8%	35.7%	51.0%	66.2%	
44	Retail	2.8%	4.5%	7.2%	11.2%	17.1%	25.2%	35.5%	47.3%	
48	Transportation, Warehousing	1.5%	2.4%	4.1%	6.7%	10.9%	17.1%	25.9%	37.2%	
51	Information	16.4%	22.5%	30.0%	38.8%	48.4%	58.1%	67.2%	75.2%	
52	Finance, Insurance	6.2%	10.8%	18.1%	28.7%	42.4%	57.3%	71.0%	81.7%	
53	Real Estate, Rental, Leasing	6.8%	10.3%	15.1%	21.8%	30.3%	40.4%	51.4%	62.2%	
54	Professional, Scientific, Technical Services	10.9%	16.7%	24.7%	34.9%	46.7%	58.9%	70.1%	79.3%	
55	Management of Companies, Enterprises	8.8%	14.4%	22.7%	33.8%	47.1%	60.8%	72.9%	82.4%	
56	Administrative Support, Waste Management, Remediation Services	4.1%	6.5%	10.2%	15.7%	23.4%	33.3%	45.0%	57.2%	
61	Education	7.5%	12.3%	19.5%	29.5%	42.1%	55.7%	68.6%	79.1%	
62	Health Care, Social Assistance	4.2%	7.4%	12.4%	20.3%	31.3%	44.9%	59.3%	72.3%	
71	Arts, Entertainment, Recreation	3.9%	6.5%	10.7%	17.2%	26.4%	38.2%	51.6%	64.8%	
72	Accommodation, Food Services	1.5%	2.4%	3.9%	6.3%	9.9%	15.2%	22.8%	32.6%	
81	Other Services	2.0%	3.3%	5.5%	8.8%	14.0%	21.4%	31.3%	43.3%	
99	Total	4.2%	6.8%	10.8%	16.7%	24.9%	35.3%	47.1%	59.0%	


NAICS	Industry	% Using AI for Production in 2026	NAICS	Industry	% Using AI for Production in 2030
11	Information	38.8%	55	Management of Companies, Enterprises	82.4%
21	Professional, Scientific, Technical Services	34.9%	52	Finance, Insurance	81.7%
22	Management of Companies, Enterprises	33.8%	54	Professional, Scientific, Technical Services	79.3%
23	Education	29.5%	61	Education	79.1%
31	Finance, Insurance	28.7%	51	Information	75.2%
42	Real Estate, Rental, Leasing	21.8%	62	Health Care, Social Assistance	72.3%
44	Health Care, Social Assistance	20.3%	22	Utilities	71.3%
48	Arts, Entertainment, Recreation	17.2%	42	Wholesale	66.2%
51	Utilities	17.2%	71	Arts, Entertainment, Recreation	64.8%
52	Administrative Support, Waste Management, Remediation Services	15.7%	21	Mining	64.2%
53	Wholesale	13.6%	53	Real Estate, Rental, Leasing	62.2%
54	Retail	11.2%	23	Construction	61.9%
55	Mining	11.1%	56	Administrative Support, Waste Management, Remediation Services	57.2%
56	Manufacturing	11.1%	31	Manufacturing	51.6%
61	Construction	9.2%	44	Retail	47.3%
62	Other Services	8.8%	81	Other Services	43.3%
71	Forestry, Fishing, Hunting, Agriculture	6.8%	48	Transportation, Warehousing	37.2%
72	Transportation, Warehousing	6.7%	72	Accommodation, Food Services	32.6%
81	Accommodation, Food Services	6.3%	11	Forestry, Fishing, Hunting, Agriculture	22.5%
99	Total	16.7%	99	Total	59.0%

Employment Current Trends

ADP Employment Change, NSA January 2028	
Industry	Change
Education and health services	32,000
Information	21,000
Financial activities	-11,000
Manufacturing	-16,000
Natural resources and mining	-18,000
Other services	-27,000
Construction	-38,000
Professional and business services	-290,000
Leisure and hospitality	-297,000
Trade, transportation, and utilities	-1,528,000
Total	-2,172,000

Revelio Labs Janaury Employment Change SA	
Industry	Change
Education and health services	41,500
Financial activities	8,800
Information	700
Utilities	400
Construction	0
Mining and logging	-100
Unclassified	-100
Other services	-900
Wholesale trade	-1,100
Transportation and warehousing	-3,200
Manufacturing	-5,100
Professional and business services	-10,600
Retail trade	-12,800
Leisure and hospitality	-14,400
Government	-16,400
Agriculture	-209,800
Total	-223,100
Non-Farm SA	-13,300

JOLTS 2026 Forecast





Current Forecasts Five-Year Economic Outlook

The **DRG BEDD** model forecasts an essentially stagnant economy over the next five years characterized by very low employment growth, and a reduction in the number of businesses.

Sector	Sub-Sector	Industry	Size Class	Metric	Total	2024	2025	2026	2027	2028	2029	2030	Change	Trends
					(All)								24 - '29	
					Total									
Firms (K)	AGR					7,073.759	7,143.098	7,290.589	7,237.320	7,294.965	7,472.814	7,684.184	541.086	
	Annual Change					1.3%	1.0%	2.1%	-0.7%	0.8%	2.4%	2.8%	1.3%	
Primary Firms (K)	AGR					6,336.815	6,311.739	6,434.594	6,237.657	6,204.905	6,279.958	6,365.578	53.839	
	Annual Change					-0.1%	-0.4%	1.9%	-3.1%	-0.5%	1.2%	1.4%	-0.3%	
Employees (K)	AGR					160,219	162,407	163,495	165,152	166,776	168,390	169,997	7,590	
	Annual Change					0.6%	1.4%	0.7%	1.0%	1.0%	1.0%	1.0%	1.0%	
Establishments (K)	AGR					9,041.600	9,143.402	9,370.883	9,253.399	9,321.560	9,511.254	9,758.279	614.877	
	Annual Change					1.3%	1.1%	2.5%	-1.3%	0.7%	2.0%	2.6%	1.0%	
Annual Payroll (M)	AGR					11,814	12,550	13,178	13,920	14,729	15,588	16,501	3,951	
	Annual Change					5.9%	6.2%	5.0%	5.6%	5.8%	5.8%	5.9%	5.8%	
Employees per Primary Firm	AGR					25.3	25.7	25.4	26.5	26.9	26.8	26.7	0.97	
	Annual Change					0.8%	1.8%	-1.3%	4.2%	1.5%	-0.2%	-0.4%	1.3%	
Establishments per Primary Firm	AGR					1.43	1.45	1.46	1.48	1.50	1.51	1.53	0.084	
	Annual Change					1.4%	1.5%	0.5%	1.9%	1.3%	0.8%	1.2%	1.3%	
Annual Payroll (1,000) per Employee	AGR					73.7	77.3	80.6	84.3	88.3	92.6	97.1	19.79	
	Annual Change					5.3%	4.8%	4.3%	4.6%	4.8%	4.8%	4.9%	4.8%	
© Daniel Research Group ©(2026)														

Consensus GDP and UER



Structural Changes

US Economy Structural Changes, 2026

Employees

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Size Code	Total		
Size Class	(All)		
Metric	Employees		
Industry	2026 to 2027	CAGR	
Forestry, Fishing, Hunting, Agriculture	18,311	0.8%	
Mining	1,322	0.2%	
Utilities	14,525	1.6%	
Construction	186,979	2.1%	
Manufacturing	31,393	0.2%	
Wholesale	35,220	0.6%	
Retail	-48,665	-0.3%	
Transportation, Warehousing	67,143	0.9%	
Information	17,196	0.5%	
Finance, Insurance	11,672	0.2%	
Real Estate, Rental, Leasing	14,913	0.6%	
Professional, Scientific, Technical Services	193,557	1.4%	
Management of Companies, Enterprises	12,729	0.5%	
Administrative Support, Waste Management, Remediation Services	41,119	0.4%	
Education	259,970	1.6%	
Health Care, Social Assistance	393,981	1.5%	
Arts, Entertainment, Recreation	79,203	2.3%	
Accommodation, Food Services	252,423	1.7%	
Other Services	66,327	1.1%	
Public Administration	-15,283	-1.7%	
National Defense	22,945	1.4%	
Total	1,656,980	1.0%	
∞/Δ Daniel Research Group © (2026)			

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Industry	(All)		
NAICS	Total		
Metric	Employees		
Size Class	2026 to 2027	CAGR	
Less than 5	116,906	1.5%	
5 to 9	102,314	1.3%	
10 to 19	189,723	1.8%	
20 to 99	178,187	0.7%	
100 to 499	88,878	0.4%	
500 to 999	67,872	0.8%	
1000 to 2499	87,176	0.7%	
2500 to 4999	10,860	0.1%	
5000 to 9999	90,799	1.0%	
10000 or more	724,265	1.5%	
Total	1,656,980	1.0%	
∞/Δ Daniel Research Group ©(2025)			
Less than 100	676,008	1.1%	
100 or More	980,972	1.0%	
∞/Δ Daniel Research Group © (2026)			

40.8% in business with less than 100 employees

US Economy Structural Changes, 2026 Businesses

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Size Code	Total		
Size Class	(All)		
Metric	Primary Firms		
Industry	2026 to 2027		CAGR
Forestry, Fishing, Hunting, Agriculture	-2,815		-1.0%
Mining	980		4.6%
Utilities	-187		-3.1%
Construction	-48,979		-6.3%
Manufacturing	-4,441		-1.9%
Wholesale	-4,761		-1.9%
Retail	-10,083		-1.7%
Transportation, Warehousing	-23,423		-12.8%
Information	-2,246		-3.6%
Finance, Insurance	-4,837		-2.2%
Real Estate, Rental, Leasing	-18,641		-5.7%
Professional, Scientific, Technical Services	-29,983		-3.6%
Management of Companies, Enterprises	-508		-8.3%
Administrative Support, Waste Management, Remediation Services	-12,288		-5.8%
Education	-3,814		-2.8%
Health Care, Social Assistance	-19,463		-2.6%
Arts, Entertainment, Recreation	-1,373		-0.8%
Accommodation, Food Services	-5,816		-1.0%
Other Services	-4,266		-0.5%
Public Administration	7		0.3%
National Defense			
Total	-196,937		-3.1%
∞Δ Daniel Research Group © (2026)			

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Industry	(All)		
NAICS	Total		
Metric	Primary Firms		
Size Class	2026 to 2027		CAGR
Less than 5	-105,295		-2.6%
5 to 9	-24,399		-2.3%
10 to 19	-19,929		-3.0%
20 to 99	-42,104		-7.9%
100 to 499	-4,985		-5.6%
500 to 999	-605		-6.1%
1000 to 2499	532		8.9%
2500 to 4999	-46		-2.4%
5000 to 9999	-25		-2.3%
10000 or more	-81		-7.5%
Total	-196,937		-3.1%
∞Δ Daniel Research Group ©(2025)			
Less than 100	-196,712		-3.0%
100 or More	-225		-4.8%
∞Δ Daniel Research Group © (2026)			

99.9% of the loss in Business with less than 100 Employees

US Economy Structural Changes, 2025-2030 Employees

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Size Code	Total		
Size Class	(All)		
Metric	Employees		
Industry	2026 to 2030		CAGR
Forestry, Fishing, Hunting, Agriculture	143,529		1.6%
Mining	-10,845		-0.4%
Utilities	53,041		1.4%
Construction	407,950		1.1%
Manufacturing	590,303		1.1%
Wholesale	284,854		1.1%
Retail	295,651		0.5%
Transportation, Warehousing	139,079		0.4%
Information	222,562		1.7%
Finance, Insurance	294,288		1.0%
Real Estate, Rental, Leasing	45,959		0.5%
Professional, Scientific, Technical Services	738,500		1.3%
Management of Companies, Enterprises	92,614		0.9%
Administrative Support, Waste Management, Remediation Services	394,443		1.0%
Education	462,548		0.7%
Health Care, Social Assistance	1,168,723		1.1%
Arts, Entertainment, Recreation	82,524		0.6%
Accommodation, Food Services	558,916		1.0%
Other Services	330,018		1.3%
Public Administration	-67,118		-1.9%
National Defense	274,433		4.1%
Total	6,501,972		1.0%
∞Δ Daniel Research Group © (2026)			
US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Industry	(All)		
NAICS	Total		
Metric	Employees		
Size Class	2026 to 2030		CAGR
Less than 5	504,770		1.6%
5 to 9	438,448		1.3%
10 to 19	738,458		1.7%
20 to 99	329,325		0.3%
100 to 499	193,191		0.2%
500 to 999	247,804		0.7%
1000 to 2499	317,920		0.7%
2500 to 4999	66,335		0.2%
5000 to 9999	465,416		1.2%
10000 or more	3,200,305		1.6%
Total	6,501,972		1.0%
∞Δ Daniel Research Group ©(2025)			
Less then 100	2,204,192		0.9%
100 or More	4,297,780		1.0%
∞Δ Daniel Research Group © (2026)			

33.9% in business with less than 100 employees.



US Economy Structural Changes, 2025-2030 Businesses

US Economy Structural Changes				
Sector	(All)			
Sub-Sector	(All)			
Size Code	Total			
Size Class	(All)			
Metric	Primary Firms			
Industry	2026 to 2030		CAGR	
Forestry, Fishing, Hunting, Agriculture		14,206	1.2%	
Mining		2,231	2.5%	
Utilities		339	1.4%	
Construction		-56,186	-1.9%	
Manufacturing		-3,912	-0.4%	
Wholesale		-6,290	-0.6%	
Retail		-5,996	-0.2%	
Transportation, Warehousing		-37,820	-5.6%	
Information		3,286	1.3%	
Finance, Insurance		-6,068	-0.7%	
Real Estate, Rental, Leasing		-15,290	-1.2%	
Professional, Scientific, Technical Services		-39,128	-1.2%	
Management of Companies, Enterprises		-569	-2.4%	
Administrative Support, Waste Management, Remediation Services		-15,153	-1.8%	
Education		4,168	0.7%	
Health Care, Social Assistance		-29,448	-1.0%	
Arts, Entertainment, Recreation		19,195	2.7%	
Accommodation, Food Services		47,449	2.0%	
Other Services		55,950	1.7%	
Public Administration		22	0.2%	
National Defense		-2	-9.6%	
Total		-69,016	-0.3%	
∞Δ Daniel Research Group © (2026)				



US Economy Structural Changes				
Sector	(All)			
Sub-Sector	(All)			
Industry	(All)			
NAICS	Total			
Metric	Primary Firms			
Size Class	2026 to 2030		CAGR	
Less than 5		3,877	0.0%	
5 to 9		-1,057	0.0%	
10 to 19		-2,964	-0.1%	
20 to 99		-61,214	-3.0%	
100 to 499		-7,569	-2.2%	
500 to 999		-619	-1.6%	
1000 to 2499		648	2.6%	
2500 to 4999		-37	-0.5%	
5000 to 9999		-6	-0.1%	
10000 or more		-75	-1.8%	
Total		-69,016	-0.3%	
∞Δ Daniel Research Group ©(2025)				
Less than 100		-68,927	-0.2%	
100 or More		-89	-1.8%	
∞Δ Daniel Research Group © (2026)				

99.1% of the loss in Business with less than 100 Employees


Business Economic Demographic Forecast Top-Line Tables and Charts Total US Economy





By Sector





Primary Firms for Governments are buying entities such as Agencies or Departments.







By Sub-Sector – Public Sector





There is only one (1) primary firm for the DoD Sub-Sector, the DoD.















By Industry – Private Sector











































































By Significant Industry – Public Sector



















































Included United States Postal Service








By Size Class, Private, Businesses































Informing Source

The tables and charts presented in this report are derived from **Daniel Research Group's (DRG)** proprietary **Business Economic Demographic Database (BEDD)**, developed through our U.S. Business Economic Demographic Models. These models are informed by a comprehensive array of data inputs from authoritative U.S. government agencies, including:

- Congressional Budget Office
- U.S. Census Bureau
- Bureau of Economic Analysis
- Bureau of Labor Statistics
- Department of Labor
- Department of Defense
- Office of Personnel Management
- Federal Reserve
- Yale Budget Lab
- ADP
- Indeed
- Revelio Labs
- Challenger, Gray & Christmas

In addition to these public sources, **DRG** incorporates selected data from industry, economic, financial, and academic institutions to ensure a robust and multidimensional foundation for analysis.

DRG synthesizes these diverse inputs into a unified historical database and applies advanced trend analysis and forecasting techniques. These include both quantitative methodologies and **DRG's** proprietary qualitative assessments of macroeconomic dynamics and structural shifts. The BEDD is updated monthly and supplemented as new data becomes available.

Data Access and Licensing

DRG's U.S. Business Demographic Baseline and Forecast Database, along with our Consumer Electronics Market History and Forecast Database, are available to clients through a variety of licensing and delivery options:

- Full database licensing
- Custom data subsets tailored to specific client needs
- Integration into bespoke market sizing and forecasting models
- Granular data access by industry taxonomy or other metrics upon request

Updates to the Baseline and Forecast are issued periodically in alignment with new releases from U.S. government data sources.

For details regarding availability, pricing, or methodology of the **DRG** U.S. Business Demographic Baseline and Forecast or the **DRG** U.S. Personal Device Forecast (PCs, Tablets, Mobile Phones), please contact Daniel Research Group directly.

To receive email notifications about new report releases or DRG analytical updates, please email: Steve@DanielRG.Com

∞Δ About Daniel Research Group

Daniel Research Group (**DRG**) provides consulting and market research services to organizations whose offerings are technology-based or technology-enabled. Our core expertise lies in delivering strategic insights and actionable solutions for clients requiring:

- Market Forecasting
- Segmentation Analysis
- Market Share Modeling
- Custom Market Modeling Tools

DRG supports these engagements with a full suite of traditional market research capabilities, including:

- Quantitative and qualitative surveys
- Focus groups
- Consumer and enterprise demographic analysis
- Expert input from technology and industry specialists

In addition to delivering data-driven recommendations, **DRG** frequently develops custom models and provides training to empower clients with the tools and knowledge needed for ongoing strategic planning.

Steve Daniel, President



[Daniel Research Group](http://www.DanielResearchGroup.com)

49 Hill Road, Suite Four

Belmont MA, 02478

(617) 484 – 6225

Steve@DanielRG.com

www.DanielResearchGroup.com

QUARTERLY HDD NUMBERS UPDATE

We do not know for sure the total number of HDDs delivered last quarter because neither Seagate or Western Digital did not report their unit shipment numbers. However, we have made estimates for shipments from these companies based upon the data they provided and past history.

Estimated HDD unit shipments in the CQ4 2025 are 33.3M, up 4.8% from our prior quarter estimate. Total exabytes shipped in CQ4 2025 were up 9.1% compared with CQ3 2025 (456.9EB versus 418.7EB). HDD revenues (estimated at about \$6.6B) were up about 7.5% from the prior quarter. Because of changes in Seagate and Western Digital reporting we do not know for sure what the nearline HDD shipments were last quarter, but we estimate about 22.5M were shipped, up 5.9% from the prior quarter.

Total HDD unit shipments in 2025 are estimated at 124.3M, up 0.4% from the prior year. Estimated nearline HDD shipments were 80.8M, up 13.6% from the prior year. Total exabyte shipments for the year were 1,620.95EB (1.61ZB), up 22% from the prior year. Total industry revenue in 2025 were \$23,267M (\$23.3B), up about 18% from the prior year.

I thought that focusing only on net revenue and exabytes shipped obscured important operational data that investors would like to know and wrote a forbes.com blog about this, [Did Western Digital and Seagate Ship Fewer HDDs Last Quarter](#).

Because of changes in reporting by Western Digital and Seagate we will not be reporting estimates for HDD shipments for various applications and form factors.

Seagate said their overall revenue was \$2,825M for CQ4 2025. Of that, we estimate that non-HDD revenue was about \$169M, slightly more than in CQ3 2025 (\$166M). We thus estimate Seagate's total HDD revenue at \$2,656 with data center revenues representing 79% of overall Seagate revenue for the quarter (probably combined system and HDD data center revenue). Total storage capacity shipped was 190EB with 165EB for Nearline HDDs (86.8%), the balance of 25EB was shipped for non-nearline applications. The average HDD capacity is estimated at 14.0TB. The company shipped about 14M units in the quarter with an ASP I estimate at \$195.28.

For the year Seagate total revenue was \$9.4B. Total storage capacity shipped was 678.1EB with an estimated 50.6M HDDs shipped.

Western Digital said that their overall revenue was \$3,017M in CQ4 2025 and that cloud revenue was \$2,673M. They didn't break out HDD only revenue, but we assume about \$101.6M for non-HDD revenue. They also said that they shipped 215B of total HDD storage in CQ3 2025 with 192EB (89.3% of the total) shipped

for nearline applications. 23EB were shipped for non-nearline applications. WDC gave little guidance on the total number of HDDs shipped, or of nearline HDD unit shipments, only saying that they had shipped 3.5M units of their latest generation ePMR HDDs. We estimate that the company shipped about 14.1M units in the CQ4 2025 quarter. Likewise we estimate that their HDD ASP was about \$206.8.

For the year Western Digital total revenue was \$10.3B. Total storage capacity shipped was 761EB with and estimated 52.6M HDDs shipped.

For 2025 Toshiba total HDD revenue is estimated at \$3.5B. Total storage capacity shipped was 182EB with 21.1M HDD shipped.

Here are the actual statements from Seagate and Western Digital for C4Q 2025. First from Dave Mosley of Seagate,

“Seagate closed out calendar 2025, with a record-breaking quarter, driven by sequential revenue growth across nearly all end markets....We expanded non-GAAP gross margin above 42%, supported by the execution of our pricing strategy, along with an improving mix of our high-capacity drives as HAMR shipments ramp... 2025 also solidified HAMR technology as a long-term enabler of mass capacity storage. We ended the year shipping 3 terabyte per disk Mozaic-based HAMR products to our first CSP customer. And by year's end, quarterly HAMR shipments exceeded 1.5 million units and have continued to ramp. Mozaic 3 HAMR drives are now qualified with all of the major U.S. CSP customers and qualifications for our second-generation Mozaic 4 terabyte perisk products are tracking well to plan.

These developments align with our long-term areal density road map that extends to 10 terabytes per disk, which we expect to deliver early in the next decade. We continue to operate in an exceptionally strong demand environment, particularly within the data center end markets. In the December quarter, we saw sustained demand growth for our high capacity nearline drives across global cloud data centers as well as continued improvement from the enterprise edge. Based on our build-to-order pipeline, we anticipate these positive demand trends will continue for some time.

Our nearline capacity is fully allocated through calendar year 2026, and we expect to begin accepting orders for the first half of calendar year 2027 in the coming months. Further out, demand visibility is strengthening based on the long-term agreements in place with major cloud customers through calendar '27. Additionally, multiple cloud customers are discussing their demand growth projections for calendar '28, underscoring that supply assurance remains their highest priority. We will continue to meet strengthening demand through our strategy to maintain supply discipline and satisfy exabyte growth through areal density advancements and without increasing unit production volume.

In the December quarter, our average nearline drive capacities rose by 22% year-over-year, approaching 23 terabytes per drive, with those sold to cloud customers averaging

significantly higher. At the same time, revenue per terabyte sold has remained relatively stable, reflecting the effectiveness of our pricing strategy... Video applications continue to drive significant demand for hard drives with platforms like YouTube witnessing 20 million video uploads daily, up from just 2 million 3 years ago.

With the deployment of AI agents at the edge, where untapped data often resides, we believe the stage is set for a sustained and meaningful increase in data generated and stored that will support inferencing, continuous training and also maintain model integrity...

Hard drives are essential to these architectures, anchoring the mass capacity data tier that stores the vast majority of exabytes from storing the checkpoint data sets used to train and maintain model integrity to supporting vector databases that provide the context necessary for accurate inference results and agentic AI performance... We expect to begin the ramp of Mozaic 4 later this quarter and have multiple CSPs qualified in the coming months in line with our plans. We continue to set the pace for the industry, recently demonstrating 7 terabytes per disk capability in our labs."

Gianluca Romano (Seagate CFO) said, for C4Q 2025, "Seagate delivered another quarter of strong year-over-year revenue growth and set new record profitability metrics in the December quarter, underscoring the durability of data center demand trends...

To that end, we shipped 190 exabytes in the December quarter, up 26% year-over-year, while keeping overall unit capacity relatively flat. The data center market accounted for 87% of our shipment volume, supported by ongoing demand momentum from global cloud customers and sequential growth across the enterprise OEM markets. We shipped 165 exabytes in the data center market, up 4% sequentially and 31% year-over-year. Data center revenue grew at roughly the same pace totaling \$2.2 billion for the quarter, up 5% sequentially and 28% year-on-year...

Average cloud nearline capacity increased to nearly 26 terabytes in the December quarter and will continue to grow with the ramp of HAMR-based Mozaic products... The edge IoT market made up the remaining 21% of revenue at \$601 million, supported by anticipated seasonal improvement for consumer products in the VIA client market. We project the broader VIA market to grow over time with the largest growth contribution coming from VIA nearline products that are captured as part of our data center end market...

We currently project other income and expenses to remain relatively flat in the March quarter... We expect March quarter revenue to be in the range of \$2.9 billion, plus or minus \$100 million, which represent a 34% year-over-year improvement as a midpoint.”

Irving Tan, CEO of WDC said, “...We are doing this by continuing to focus on increasing our drives areal density and accelerating our HAMR and ePMR road maps as well as upshifting our customers to accelerate adoption of higher capacity drives and UltraSMR technology. This last quarter, we shipped over 3.5 million units of our latest generation ePMR products, offering up to 26 terabyte CMR and 32 terabyte by UltraSMR capacities, representing strong confidence and adoption by our customers. We have also started qualification of our HAMR and next-generation ePMR products, each with a different hyperscale customer....In addition, we continue to accelerate our HAMR innovation...

we recently acquired intellectual property assets and talent that will help us in the development of our internal laser capabilities. Also, this past quarter, in partnership with software ecosystem partners, we announced our UltraSMR-enabled JBOD platforms, expanding UltraSMR adoption to a broader customer set... We have firm purchase orders with our top seven customers through calendar year 2026. We also have in place robust commercial agreements with three of our top five customers, two through calendar year 2027 and one through calendar year 2028...

In keeping with our strategy to incubate new growth vectors based on our intellectual property and core capabilities, last month, we announced a strategic investment in Qolab, which combines our expertise in material science and precision manufacturing with Qolab's breakthrough approach to quantum hardware design. Working with Qolab, we aim to advance next-generation nanofabrication processes that improve qubit performance, reliability and scalability. Looking ahead, we see our positive momentum continuing and we will remain focused on supporting our customers' exabyte storage requirements while completing qualifications and launching our next-generation HAMR and ePMR drives.”

Kris Sennesael (WDC CFO) said, “During the second quarter of fiscal 2026, revenue was \$3 billion, up 25% year-over-year, driven by strong demand for our nearline drives. Earnings per share was \$2.13. Both revenue and EPS were above the high end of the guidance range. We delivered 215 exabytes to our customers, up 22% year-over-year. This includes over 3.5 million drives or 103 exabytes of our latest generation ePMR with capacity points up to 32 terabytes.

Cloud represented 89% of total revenue at \$2.7 billion, up 28% year-over-year, driven by strong demand for our higher capacity nearline product portfolio. Client represented 6% of total revenue at \$176 million, up 26% year-over-year.

Consumer represented 5% of revenue at \$168 million, down 3% year-over-year...I will now turn to the outlook for the third quarter of fiscal 2026.

We anticipate revenue to be \$3.2 billion, plus/minus \$100 million. At midpoint, this resects a growth of approximately 40% year-over-year. Gross margin is expected to be between 47% and 48%. We expect operating expenses in the range of \$380 million to \$390 million. Interest and other expenses are anticipated to be approximately \$50 million. The tax rate is expected to be approximately 16%.”

In **Table 1** we estimate an 3.3 % average sales price (ASP) increase for HDDs from C3Q 2025 to C4Q 2025. But this ASP is not certain because of the changes in Seagate and WDC reporting don’t provide this information explicitly.

TABLE 1. AVERAGE DISK DRIVE SALES PRICES (ASP) AND UNIT VOLUMES (MILLIONS/\$ PER UNIT)

	Q1 25 Units/\$	Q2 25 Units/\$	Q3 25 Units/\$	Q3425 Units/\$
Seagate	11.4/\$175.70	12.5/\$182.50	13.1/\$188.04	13.6/\$195.28
WD	12.1/\$179.00	12.9/\$194.98	13.5/201.30	14.1/206.80

The multiyear trends in ASPs are shown in **Figure 1**. HDD ASPs are now at levels much higher than seen at the end of 1998 due to the increased demand for high capacity, high ASP, HDDs and the increase in the percentage of these drives shipped.

Table 2 shows estimated industry revenue for 2025. HDD revenues were up about 7.5% in C4Q 2025 compared to C3Q 2025.

TABLE 2. YEAR TO DATE ESTIMATED HDD INDUSTRY REVENUE (\$MILLIONS)

Company	Q1	Q2	Q3	Q4	Total
Seagate	\$2,003	\$2,281	\$2,464	\$2,656	\$9,404
Western Digital	\$2,166	\$2,515	\$2,717	\$2,915	\$10,314
Toshiba	\$750	\$863	\$933	\$1,003	\$3,549
Total	\$4,919	\$5,660	\$6,114	\$6,574	\$23,267

Table 3 shows C4Q 2025 estimated Exabytes Shipped by company and general application. **Table 4** shows total capacity shipped year to date. **Figure 5** shows a pie chart of exabyte shipments by company and **Figure 6** shows a pie chart of estimated HDD revenue by company.

TABLE 3. CALENDAR C4Q 2025 ESTIMATED DISK DRIVE EXABYTES OF CAPACITY SHIPMENTS

Company	Nearline	Legacy	Total
Seagate	165.00	25.00	190.00
Western Digital	192.00	23.00	215.00
Toshiba	44.05	7.81	51.86
Total	401.05	55.81	456.86

TABLE 4. YTD ESTIMATED DISK DRIVE EXABYTES OF CAPACITY SHIPMENTS

Company	Q1	Q2	Q3	Q4	Total
Seagate	143.60	162.50	182.00	190.00	678.10
Western Digital	166.00	190.00	190.00	215.00	761.00
Toshiba	40.98	42.29	46.72	51.86	181.85
Total	350.58	394.79	418.72	456.86	1620.95

FIGURE 1. AVERAGE DRIVE PRICE TREND FOR SEAGATE, WESTERN DIGITAL AND (HISTORICALLY, HGST) (C4Q '98 TO C4Q '25)



FIGURE 5. YEAR TO DATE EXABYTE SHIPMENTS BY COMPANY.




FIGURE 6. TOTAL REVENUE SHIPMENTS BY COMPANY.




Figure 7 shows high, median and low estimates of total drive unit volume estimates out to 2030.

FIGURE 7. BANDED HARD DRIVE VOLUME PROJECTIONS





Figure 8 gives our estimates for HDDs, magnetic media (disks) and heads production out to 2030 based upon the median estimate for HDDs in Figure 7.

FIGURE 8. PROJECTED DEMAND FOR HDD HEADS AND DISKS



This shows the significant decline in head and media demand in 2023 with increases in 2024 and beyond as nearline demand increased requiring more heads and media.

The growth of capacity-oriented nearline drives for enterprise and hyperscale applications, particularly AI, will result in more components per drive and thus higher head and media demand than in the past by 2027 and with high disk and head counts per drive, provides the biggest driver for heads and media.

Figure 9 shows laboratory demonstrations and product announcement maximum areal densities history. Seagate's current highest capacity HAMR HDD is 36TB capacity (using SMR) and the 11 disk 32TB WDC SMR drives have about 1.4 Tbps/psi so we sized the 36TB drive as 1.8Tbps/psi, another product areal density high point. In addition, Seagate announced a laboratory demonstration of 7TB per disk in 2025. These advances are reflected in the charts below.

FIGURE 9. AREAL DENSITY ANNOUNCEMENT DEVELOPMENTS




Figure 10 shows the product announcement maximum areal density per quarter compared to lines representing areal density growth.

FIGURE 10. AREAL DENSITY PRODUCT ANNOUNCEMENTS VS. CAGR




Figure 11 projects total shipments of HDD storage capacity on an annual basis to 2030.

FIGURE 11. PROJECTION OF DISK DRIVE ANNUAL CAPACITY SHIPMENT




Figure 12 shows the estimated average \$/GB for HDDs. **Figure 13** shows total projected shipments for Tape (LTO), SSDs and HDDs. Note that although the bulk

of anticipated shipped storage capacity is in HDDs by 2030, the percentage of that storage that is in SSDs will have increased.

FIGURE 12. AVERAGE \$/GB FOR HDDS





FIGURE 13. CAPACITY SHIPMENTS FOR MAGNETIC TAPE, SSDS AND HDDS



Coughlin Associates

Current and Upcoming Reports

**Prices are for PDF copy (each report)
Contact us for multiple report discounts**

Digital Entertainment Series

- **2022 Digital Storage in Professional Media and Entertainment Report, \$7,500 (individual copy) Update coming in 2026**

Storage Infrastructure Series

- **2025 Emerging Memory Report, \$7,500 (individual copy)**

To order these (or any other of our reports) or to order Entertainment Storage Alliance membership, including the **Digital Storage Technology Newsletter** (with a 10% discount on our reports) call 408-202-5098 or visit the Coughlin Associations web site at <http://www.tomcoughlin.com/techpapers.htm>.

2025 New Memories Report

This report, jointly produced by Objective Analysis and Coughlin Associates, provides an exhaustive look at emerging memory technologies and their interaction with standard memories, both as discrete devices and in embedded applications (the memories within logic chips like ASICs and MCUs). The report provides a well of technical information, market dynamics, forecasts, and competitive analyses of the leading companies. Forecasts show how the markets will grow not only for the technologies themselves, but also for the capital equipment used to produce them. Read this to understand the competitive landscape and market drivers for these new memories, and to learn how to profit from tomorrow's market. **Report now available.**

Table of Contents:

EXECUTIVE SUMMARY	ERROR! BOOKMARK NOT DEFINED.
INTRODUCTION	ERROR! BOOKMARK NOT DEFINED.
WHY NEW MEMORY TECHNOLOGIES ARE POPULAR	ERROR! BOOKMARK NOT DEFINED.
HOW A NEW MEMORY LAYER CAN IMPROVE COMP. PERFORMANCE	ERROR! BOOKMARK NOT DEFINED.
THE CRUCIAL IMPORTANCE OF THE ECONOMIES OF SCALE	ERROR! BOOKMARK NOT DEFINED.
ONLY ROOM FOR ONE AT THE TOP	ERROR! BOOKMARK NOT DEFINED.
UNDERSTANDING BIT SELECTORS	ERROR! BOOKMARK NOT DEFINED.
CROSSPOINT ARRAYS	ERROR! BOOKMARK NOT DEFINED.
ALL YOU HAVE TO DO IS	ERROR! BOOKMARK NOT DEFINED.
RESISTIVE RAM, RERAM, RRAM, MEMRISTOR	ERROR! BOOKMARK NOT DEFINED.
FERROELECTRIC RAM, FERAM, FRAM	ERROR! BOOKMARK NOT DEFINED.
PHASE CHANGE MEMORY (PCM, PRAM, PCRAM)	ERROR! BOOKMARK NOT DEFINED.
MRAM(MAGNETIC RAM), STT MRAM(SPIN TRANSFER TORQUE MRAM)	ERROR! BOOKMARK NOT DEFINED.
OTHER NEW MEMORY TYPES	ERROR! BOOKMARK NOT DEFINED.
LITHOGRAPHY	ERROR! BOOKMARK NOT DEFINED.
3D MEMORY CIRCUIT DESIGN	ERROR! BOOKMARK NOT DEFINED.
SUMMARY OF SOLID-STATE MEMORY & STORAGE TECHNOLOGIES	ERROR! BOOKMARK NOT DEFINED.
NEW MEMORY TECHNOLOGIES AND NEW MATERIALS	ERROR! BOOKMARK NOT DEFINED.
NEW MEMORY PROCESS EQUIPMENT	ERROR! BOOKMARK NOT DEFINED.
MEMORY IS DRIVING SEMICONDUCTOR CAPITAL SPENDING	ERROR! BOOKMARK NOT DEFINED.
MARKET PROJECTIONS	ERROR! BOOKMARK NOT DEFINED.
ESTIMATES OF NEW MEMORY CAPITAL EQUIPMENT DEMAND	ERROR! BOOKMARK NOT DEFINED.

NEW MEMORY MANUFACTURING SUPPORT	ERROR! BOOKMARK NOT DEFINED.
COMPANY INFORMATION	ERROR! BOOKMARK NOT DEFINED.
METHODOLOGY	ERROR! BOOKMARK NOT DEFINED.
FURTHER READING	ERROR! BOOKMARK NOT DEFINED.
FIGURES.....	ERROR! BOOKMARK NOT DEFINED.
TABLES.....	ERROR! BOOKMARK NOT DEFINED.
AUTHORS.....	ERROR! BOOKMARK NOT DEFINED.

Order at: <https://tomcoughlin.com/tech-papers/>
or <https://Objective-Analysis.com/reports/#Emerging>



2022 Digital Storage in Professional Media and Entertainment Report

This updated and expanded report is the fifteenth annual comprehensive reference document on this topic. The report analyzes requirements and trends in worldwide data storage for entertainment content acquisition; editing; archiving and digital preservation; as well as digital cinema; broadcast; satellite; cable; network; internet and OTT and VOD distribution. Capacity and performance trends as well as media projections are made for each of the various market segments. Industry storage capacity and revenue projections include direct attached storage, cloud (including object storage), real-time as well as near-line network storage.

ORDER FORM FOR THE 2022 DIGITAL STORAGE FOR MEDIA AND ENTERTAINMENT REPORT (PDF)

NAME: _____
TITLE: _____
COMPANY: _____
ADDRESS: _____
CITY: _____ STATE: _____
ZIP: _____
TELEPHONE: _____
FAX: _____
E-MAIL: _____

Individual License \$7,500

☐ Visa ☐ Mastercard ☐ American Express

Credit Card Number: _____
Expiration Date: _____
Signature: _____

Make checks payable to: Coughlin Associates
Mail to 1665 Willowmont Ave., San Jose, CA 95124
Telephone: 408-202-5098 Fax: 866-374-6345
Order On-Line at: <https://tomcoughlin.com/tech-papers/>
Email: info@tomcoughlin.com

Consulting

Tom Coughlin is available for technical and market consulting on digital storage devices, systems and applications. He is also available for research and writing on digital storage and applications. Clients have included well know as well as start-up storage companies, storage component suppliers and storage users.

About Tom Coughlin



Tom Coughlin, President, Coughlin Associates is a widely respected storage analyst and consultant. He has over 40 years in the data storage industry with multiple engineering and management positions at high profile companies.

Tom has many publications and six patents to his credit. Tom is also the author of Digital Storage in Consumer Electronics: The Essential Guide, now in its second edition with Springer.

Tom has also written Blogs on digital storage topics for GLG, POST Magazine and Forbes.com. Coughlin Associates provides market and technology analysis (including reports on several digital storage technologies and applications and a newsletter) as well as Data Storage Technical Consulting services.

Tom is active with SMPTE, SNIA, the IEEE and other professional organizations. He is an IEEE Fellow and was President of the IEEE in 2024. Tom is the founder and organizer of the Storage Visions Conference as well as the Creative Storage Conference. Tom was chairman of the annual Flash Memory Summit for 11 years. He is currently the FMS General Chair. He is a Leader in the Gerson Lehrman Group Councils of Advisors and a member (and board member) of the Consultants Network of Silicon Valley (CNSV). For more information go to www.tomcoughlin.com.

NEWSLETTER SUBSCRIPTIONS

The Digital Storage Technology Newsletter is a Quarterly Publication of Coughlin Associates. It is published quarterly. Subscriptions for 2026-2027 are \$250 and can be made using the form below.

You can also subscribe here: <https://tomcoughlin.com//product/digital-storage-technology-newsletter/>

✂ Cut or tear here ✂
here ✂

✂ Cut or tear

ORDER FORM

NAME: _____

TITLE: _____

COMPANY: _____

ADDRESS: _____

CITY: _____

STATE: _____ ZIP: _____

TELEPHONE: _____ FAX: _____

E-MAIL: _____

☐ Quarterly PDF Newsletter Fee: \$250 (no shipping or tax)

☐ Check Enclosed

☐ Bill Me - P.O. Number _____

☐ Visa ☐ Mastercard

☐ American Express

Credit Card Number: _____

Expiration Date: _____

Signature: _____

Make checks payable to: **Coughlin Associates.**

Mail to: 1665 Willowmont Ave., San Jose, California 95124

Tel: 408-202-5098 Fax: 866-374-6345