



**DIGITAL STORAGE
TECHNOLOGY NEWSLETTER**

VOLUME XXXX, NUMBER 3

November 2025

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) High-Capacity SSDs will Enable AI Workloads But Also Drive HDD Demand (Forbes, 8/14/25)..... 5
- 2) Cold Data At FMS And Magnetic Tape Data Recovery, (Forbes 8/15/25) 9
- 3) IEEE Returned to the 2025 World Science Fiction Convention, (Forbes, 8/21/25)..... 13
- 4) How Would an AI Recession Impact AI Commercialization?, (Forbes, 8/24/25)..... 16
- 5) Semiconductor Sales and Capital Spending Driving Hot Chips Interest, (Forbes, 9/9/25) 18
- 6) UCIe, Marvell SRAM and DRAM and Rapidus 2nm Fab at the 2025 Hot Chips Conference, 9/3/25)..... 22
- 7)) Memory and Design Advances from the AI Infrastructure Summit, (Forbes 9/11/25)..... 27
- 8) Pure Storage Sustainability Report Shows How to Reduce Carbon Footprint and Waste, 9/14/25) 32
- 9) In AI, No Data is Ever Cold, (Forbes Blog, 9/15/25)..... 34
- 10) AI and Fast File Transfers Improve Media and Entertainment Preservation and Workflows, 9/20/25) 37
- 11) SNIA SDC on Sstoage.ai, OCP Storage Work and How Chiplets Can Reduce Power Consumption (Forbes Blog, 9/21/25)..... 40
- 12) Hard Disk Drive Companies Could Become Dominant Manufacturers of Lasers with HAMR, (Forbes Blog, 9/24/25) 43
- 13) Maker Faires Could Help IEEE Create the Future, (Forbes Blog, 9/30/25)47
- 14) Storage and Memory Price Surges to Support AI Demand are Likely Temporary, (Forbes Blog, 10/10/25)..... 50

15) Vdura Controls Costs while Providing Lower Cost, At Scale AI Storage, (Forbes Blog, 10/8/25)..... 53

16) IEEE Innovation Summit Develops New Ways to Create IEEE Initiatives, (Forbes Blog, 10/12/25) 56

17) In the Near Future Data Center SSDs will Require Liquid Cooling, (Forbes Blog, 10/14/25) 59

18) OCP E2, One Form Factor to Rule them All, (Forbes Blog, 10/23/25)..... 62

19) WDC Opens Data Center Customer Lab and Toshiba Announces 12-Disk HDDs by 2027, (Forbes Blog, 10/27/25) 68

20) OpenDrives Shows Cloud-Native Astraeus On-Premises Data Solutions, (Forbes Blog, 10/29/25) 69

21) Did Western Digital and Seagate Ship Fewer HDDs Last Quarter, (Forbes Blog, 11/5/25)..... 71

GUEST ARTICLES..... 76

1) A Quick Wrap of the Global Content Conference, Andy Marken (Content Insider IBC Special), andy@markencom.com..... 76

2) A Little Creative AI is Fine, But..., Andy Marken (Content Insider #923— Rogues, andy@markencom.com 85

3) Hard Disk Drives Unexpectedly Not Dead, David Rosenthal, 9/18/25, <https://blog.dshr.org/2025/09/hard-disk-unexpectedly-not-dead.html> 97

4) The Bathtub Curve, October 28, 2025, Steve Daniel, sdaniel1958@gmail.com..... 106

5) How will your total available market change over the next five years, Steve Daniel, sdaniel1958@gmail.com 109

QUARTERLY HDD NUMBERS UPDATE 172

CURRENT AND UPCOMING REPORTS 184

2025 NEW MEMORIES REPORT..... 185

2022 DIGITAL STORAGE IN PROFESSIONAL MEDIA AND ENTERTAINMENT REPORT..... 186

CONSULTING 187

NEWSLETTER SUBSCRIPTIONS..... 188

INTRODUCTION

This is the November 2025 issue of this newsletter and the third issue in the 2025-2026 annual volume.

This year marks the 40th 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 third calendar quarter of 2025 (and the first month of C4Q 2025)**:

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

- We do not know the total number of HDDs delivered last quarter because Western Digital did not report their unit shipment numbers. However, Seagate and Toshiba combined shipped 17.7M units, up 1.4% from the prior quarter.
- Total exabytes shipped in CQ3 2025 were up 4.9% compared with CQ2 2025 (414.3EB versus 394.8EB).
- HDD revenues (estimated at \$6.6B) were up about 7.25% from the prior quarter.
- Only Seagate gave information allowing an estimate of their HDD ASP's if we assume non-HDD revenues were about the same as in the prior quarter. Seagate's ASPs is likely to have increased by about 2.7% in C3Q 2025 versus C2Q 2025. Industry ASP may be somewhat higher as WDC ASPs are usually a bit higher than Seagate's

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 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

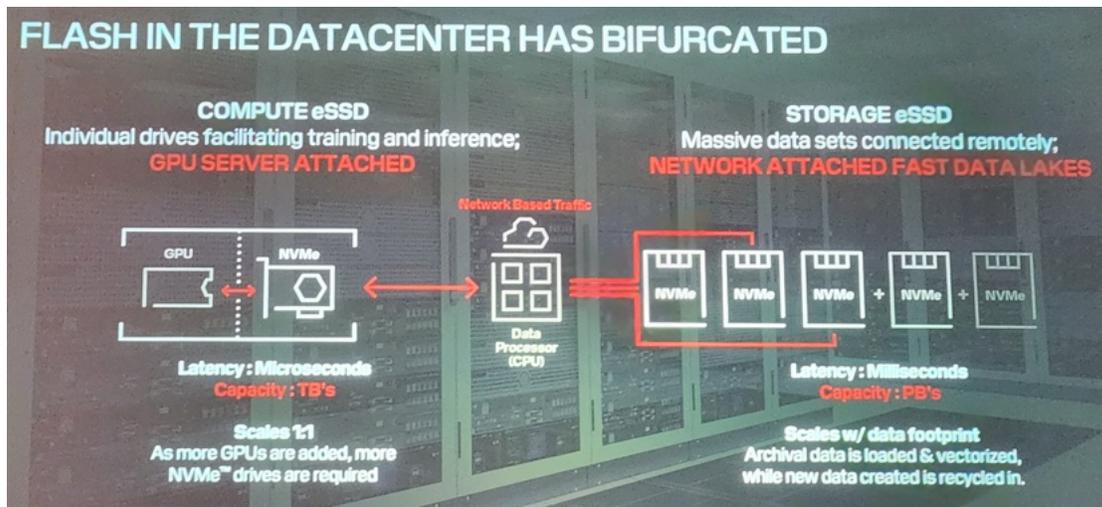
We are now doing regular postings of links to storage and memory related news (including company announcements) to a new Digital Storage and Memory Flipboard. You can find this flipboard and follow it at: <https://flipboard.com/@thomasmcoughlin/digital-storage-and-memory-ku86vmbnz>

ARTICLES

1) High-Capacity SSDs will Enable AI Workloads But Also Drive HDD Demand (Forbes, 8/14/25)

At the recent FMS conference in Santa Clara, almost all of the SSD companies were introducing high-capacity SSDs, many over 200TB, with promises for large form factor SSD with 1PB capacities in the future. These SSDs leverage higher logical density, four bit per cell, or QLC flash memory and lots of chips to achieve these capacities.

The SanDisk keynote differentiated a couple of different uses for SSDs to support AI workloads. One type are fast eSSDs to support high bandwidth DRAM memory, HBM. The other type are high-capacity storage eSSDs for a higher performance data lake than HDDs can offer. These two types of SSD are shown below.



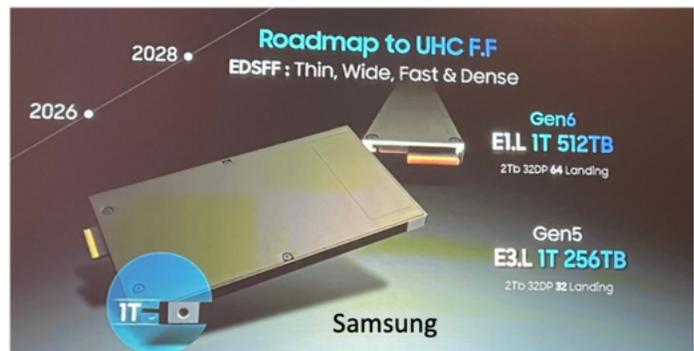
SanDisk Presentation, Ron Dennison

The slide below shows Sandisk's high capacity eSSD. It is a QLC BiCS8 NAND flash U.2 and EDSFF form factor SSDn that is expected to have capacities up to 256TB by 2026.

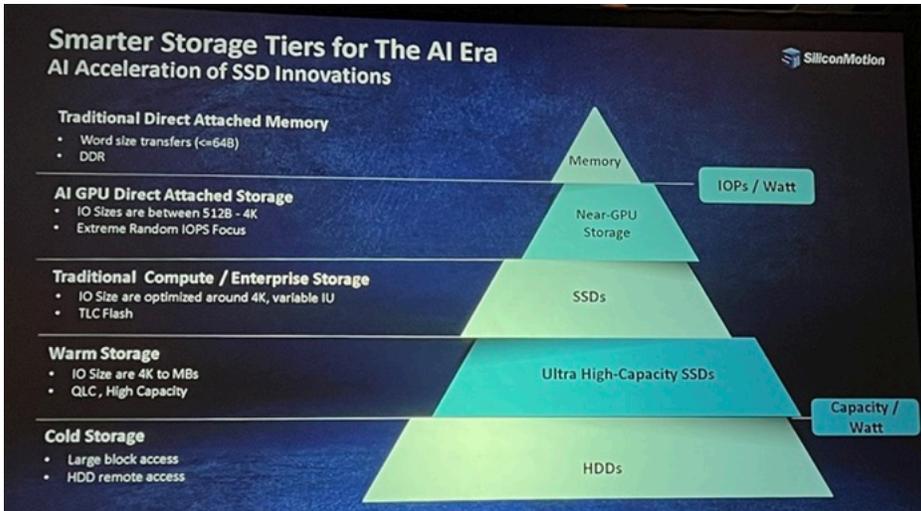


SanDisk Presentation, Ron Dennison

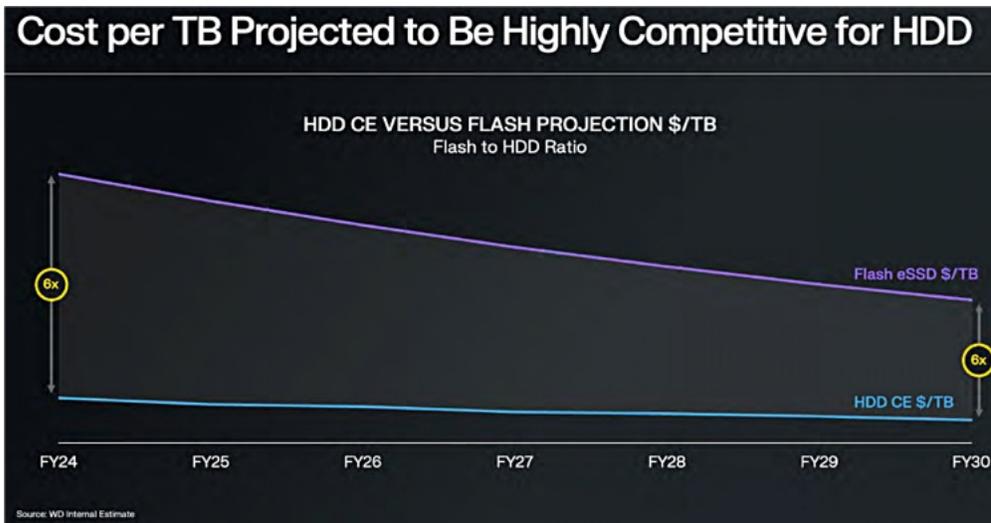
The Sandisk Keynote showed a path to a 512TB version by 2027 and 1PB product sometime in the future. The composite image below shows the Kioxia, Micron and Samsung announcements of their high capacity QLC SSDs. All of the SSD companies are exploring product for near storage AI applications. Many of the higher capacity products are using the E3.S form factor, which can hold more NAND flash chips to enable higher capacities.



The Silicon Motion keynote gave an illustration of a traditional representation of the memory and storage hierarchy showing trends for NAND flash supporting GPUs directly, like the HBF announcements by SK hynix and Sandisk at the FMS. It also shows an ultra-high-capacity SSD layer to support warm storage for AI applications. Silicon Motion supplies controllers for SSDs.



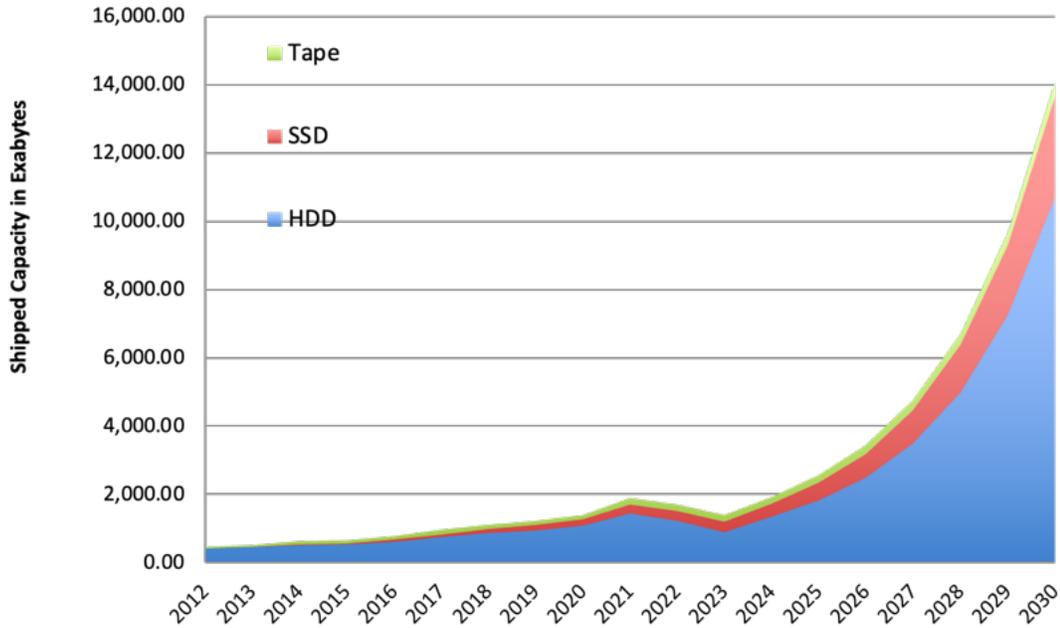
Higher capacity SSDs can take less rack space than HDDs and they offer higher performance than HDDs offer. This can be an advantage for AI training and inference with RAG, but flash memory is currently about 6X more expensive per storage capacity than HDDs and is expected to remain so, for some time into the future, as shown in the image below from the WDC investor day last February. Seagate shows similar trends. For instance, by 2026 44TB HDDs should be in production, a 38% increase from the largest HDDs available today. This is because the expected storage capacity growth in HDDs has accelerated with the introduction of HAMR HDDs to roughly match the growth in SSD capacities.



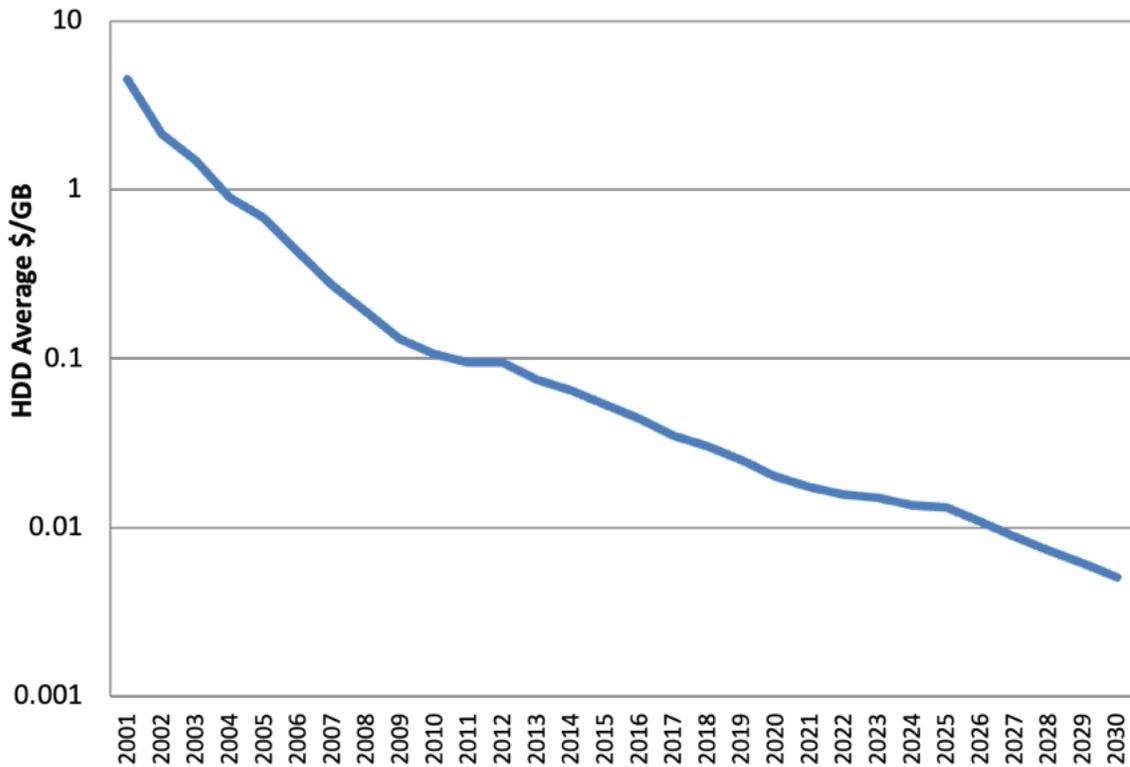
As a consequence, we consider these larger SSDs will be used for data lakes directly feeding the memory attached to GPUs for AI workflows. However, HDDs will continue to provide lower cost storage for longer term data retention and so these higher capacity SSDs will result in greater growth of HDDs as well.

[Coughlin Associates has updated its projections for storage capacity shipped for HDDs, SSDs and magnetic tape](#), shown below. This new projection increases our

expectations for growth of SSD storage from prior versions out to 2030 with some reduction in HDD capacity shipments as a consequence.



The Coughlin Associates projection for HDD storage capacity prices out to 2030 is shown below.



If we assume that SSDs remain at 6X the cost per storage capacity by 2030 and the HDD price per GB of \$0.0051 in 2026, the NAND flash price would be about

\$0.031 per GB. With the projections for shipping capacity of SSDs and HDDs of about 3.0ZB and 10.7ZB, projected revenue for SSDs and HDDs in 2030 is \$93B and \$55B respectively. The rising storage boat, driven by AI, is expected to result in significant revenue growth for HDDs as well as SSDs.

FMS 2025 shows growth of high-capacity SSDs as well as HDDs to support the growth of AI workflows.

2) Cold Data At FMS And Magnetic Tape Data Recovery, (Forbes 8/15/25)

The FMS conference has moved from a concentration on solid state storage to other types of storage as well. In this article we will explore a session on cold storage. In addition, we will report on an on-going project to recovering data from old tapes for training AI and other applications.

The cold storage session was moderated by Rich Godomski from Fujifilm and included talks by John Monroe from Further Market Research, Ilya Kazansky, CEO, SPhotonix, Steffen Hellmold, President, Cerabyte, Inc., Dave Landsman, Distinguished Engineer, Director Industry Standards, Western Digital and Alistair Symons, VP, Storage Systems Development, IBM.

John Monroe's message was that about 70% of the data stored in enterprise and data center applications is cold data, but that it may need to be accessed at any time to support modern AI workflows. Thus, it should be stored on an active archive, which would save on power consumption and operating costs. He also felt that in the future AI will be used to manage what could otherwise become an enormous amount of stored data.

Although much of this cold storage is currently on HDDs, he projected that magnetic tape and emerging optical archive storage technologies would grow to over 20% of the total shipped storage capacity by 2030, compared to about 16% in 2024 (which is somewhat higher than my projection of about 12% in 2024).

Ilya Kazansky CEO of SPhotonix spoke about the company's optical storage technology, its 5D fused quartz-based memory crystals, an example is shown below.



He argued that conventional storage technologies need to be replaced every 10-15 years and that migrating content to new storage technology takes a significant amount of effort and money and that it results in significant waste. The commonly used storage technologies are also susceptible from damage from heat, humidity, chemical reactions and electromagnetic pulses and that they require controlled environments to preserve the storage media.

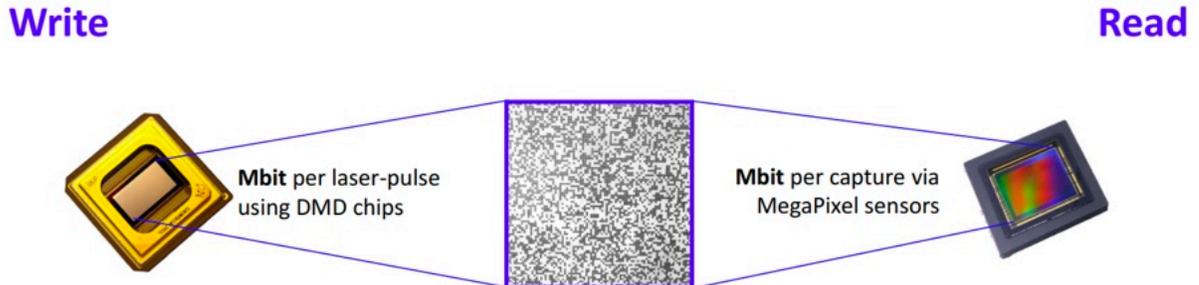
The company's memory crystals provide a write once, read many times or WORM media that can persist for a thousand years or more. Recording on this media are also not subject to damage from heat, humidity and the other factors that can cause loss of data in other storage media. Thus, it would not require migration to new media on a regular schedule and is easy to recycle, since it has a composition similar to sand.

He said this volumetric storage media can support storage densities up to 10GB/mm³. The femto-second lasers used for recording record nanostructures that can encode data in their width, height, depth, polarization and birefringence, hence the 5D name.

Steffen Hellmold, President of Cerabyte, spoke about their optical media technology for archiving applications. He also stressed the costs of archiving and management required with current digital storage media.

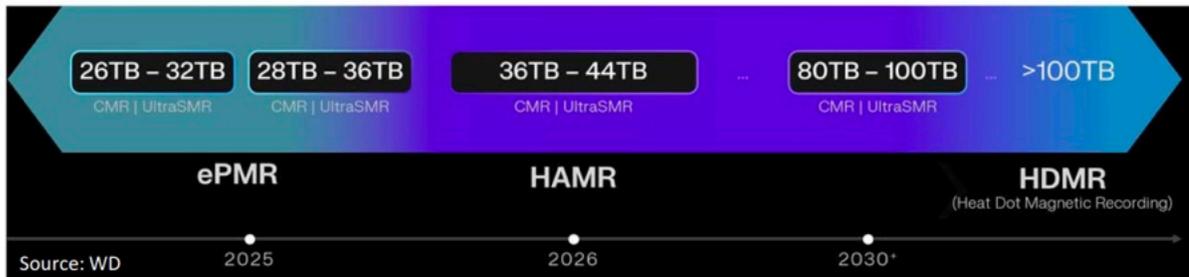
Similar to the SPhotonix media, Cerabyte media could last a long time since it is not sensitive to heat, humidity and other factors that can damage the media or the data on conventional digital storage media. He stressed the need to move from traditional archiving approaches to active archiving to support data needs of modern AI workflows.

The Cerabyte media uses glass substrates, which are coated with a ceramic material. The Cerabyte solution uses digital light processors, DLPs, from Texas Instruments to effectively write and read millions of bits at once and thus providing high data throughput, as shown below.



Although not represented in this session, there are other optical storage startup working on archival technology. These include Folio Photonics and Optera, both of which are developing higher capacity circular optical media that could be placed into traditional optical storage libraries.

David Landsman, Distinguished Engineer and Director of Industry Standards, at Western Digital spoke about why HDDs will remain relevant for cold storage applications. He said that HDDs are where cool, versus cold data, lives in data centers. He said that HDDs for data centers are 6X less expensive than data center SSDs and would remain so through 2030 and beyond. One of his slides showed 36-44TB HDDs in 2026, 80-100TB HDDs by 2030 and over 100TB in the years after, as shown below.



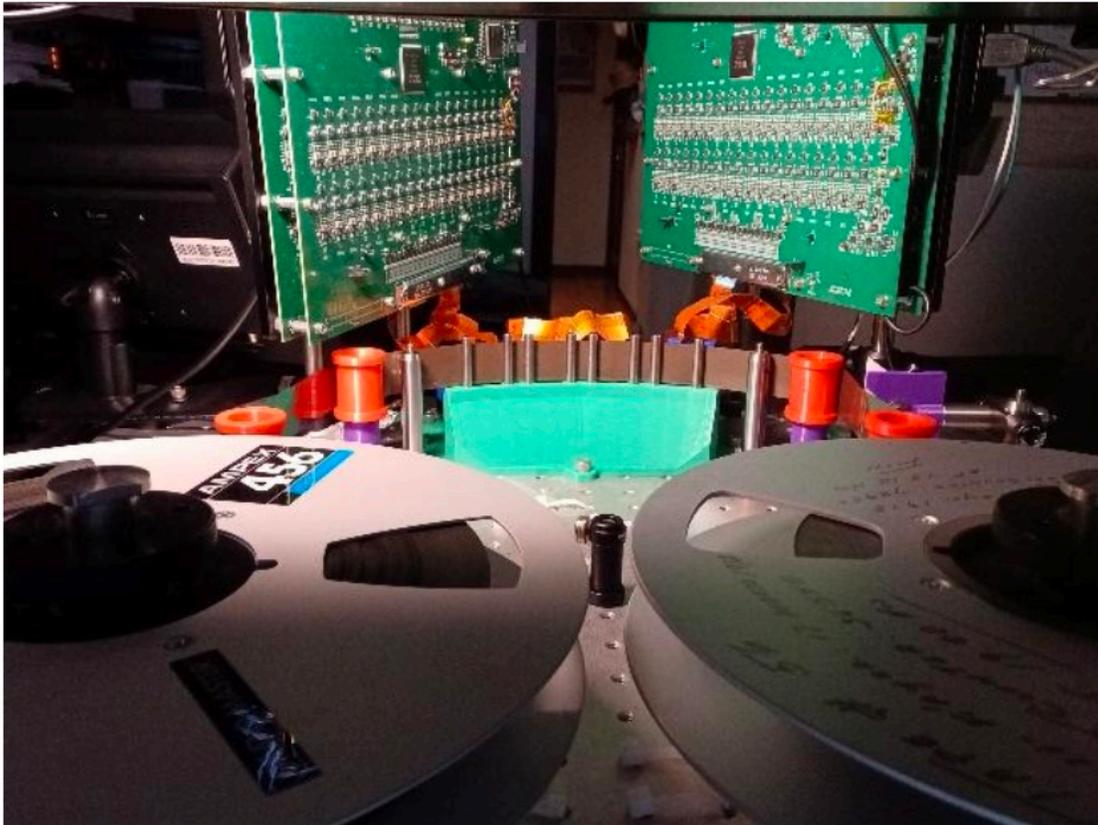
Alistair Symons, VP of Storage Systems Development at IBM, spoke in favor of magnetic tape for archive storage in the session. He argued that the increasing amount of storage that AI and other applications is a significant factor in the increasing energy requirements at storage centers and this will favor storing less frequently accessed data on magnetic tape. Compared to SSDs and HDDs the energy consumption for magnetic tape storage is considerably less.

Chuck Sobey of Channel Science and the organizer of the Monday training sessions before the regular FMS sessions, has been pursuing the development of a multi-format, minimal-contact legacy tape reader that can recover data from old magnetic tapes.

There are vast libraries of older magnetic tapes in now obsolete formats that contain valuable data for AI training and other applications or that may contain other content that has a long-standing cultural value. Training on this often-irreplaceable data, such as cultural, seismic, climate or astronomy data, could make better domain specific AI-based models.

Currently many efforts to recover data from older tapes requires having the original tape equipment to read the tapes. In addition, even if you can find the old tape drives, replacing worn out parts such as heads is hard because these heads are no longer in production. In addition, even if you have a functioning vintage tape drive it only has the capabilities that were designed into the original drives.

To meet the need to read these old magnetic tapes Channel Science created its multiformat tape reader. The company has one patent granted on the technology and more pending. The prototype tape transport device is shown below.



SBIR Grant Award: DE-SC0021879

Developed using SBIR grants the device uses GMR sensors, in mass production for HDD heads, to read the data from the tapes with a gentle tape path and a high-speed tape transport with advanced signal processing, detection and decoding and uses AI to optimize its operation for different types of tape. To find out more, contact Chuck Sobey at csobey@ChannelScience.com.

The 2025 FMS Conference included sessions and exhibitors working on the increasing demand for colder storage for AI and other applications. Channel Science developed a multiformat magnetic tape reader for recovering data from old tapes.

3) IEEE Returned to the 2025 World Science Fiction Convention, (Forbes, 8/21/25)

In 2024 IEEE had a presence at the World Science Fiction Convention, Worldcon, in Glasgow, Scotland. I gave a talk on digital storage and memory technology and current IEEE President Kathleen Kramer moderated a standing room only panel on brain machine interfaces. We also had a good-sized exhibit area that IEEE student members participated in during the convention. I was partly inspired to pursue a career in technology by speculative fiction and I felt that IEEE should make an effort to be at events like this where people interested in technology and the future would gather.

IEEE returned for the 2025 Worldcon, in Seattle, Washington. I organized and moderated a panel on digital preservation and Alon Newton, from the IEEE Seattle section, organized and moderated a session on Future Remote Medical Sensing and Privacy Concerns.

The image below is of the participants in my panel on digital preservation outside of our meeting room. From left to right they are: me, Gemma Mendonsa (working on DNA storage at Seagate Technology), Rebecca Faimow (Boston-based author and archivist), Mariecris Gatlahayan (digital preservation librarian at the University of Washington) and Riyan Mendonsa (research engineer at Seagate Technology). We had a lively discussion on how to preserve data for over 100 years with lots of discussion with the audience about their digital storage and content preservation activities.

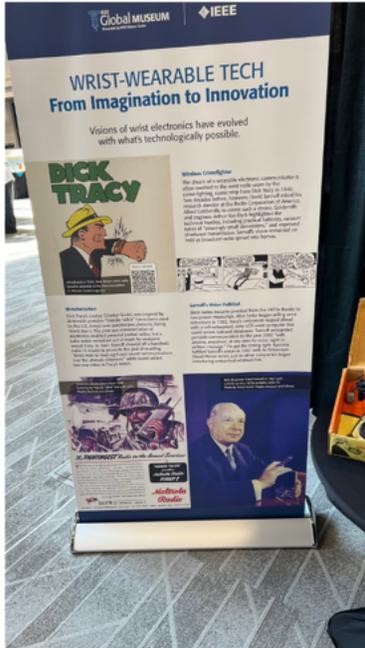


We had an exhibit at the 2025 Worldcon, thanks to the Seattle volunteers, particularly Alon and to IEEE staff in the IEEE Conferences and Events department. IEEE staff and particularly Rebecca Orens helped us get stand-up posters and give aways at our exhibit. The photo below is from the IEEE booth with from left to right, Alon Newton, Rebecca Orens and me.



We had hoped to have students help with the booth this year, but it never came together. There were students who wanted to come from nearby universities in Canada, but they were originally from other countries and with the current issues with coming in or out of US borders, they decided not to come. Oren and Alon were at the booth the most and I tried to stay there as much as I could.

In addition to help from IEEE Conference and Events staff, the IEEE history center sent us some interesting history posters and some artifacts. Some of the material from the IEEE history center are shown below, including a poster about wrist wearable technology, featuring Dick Tracy and his radio/TV watch and a 1955 wrist radio kit. Another artifact was a book from 1946 that gathered a set of stories, published in 1898, that was a sequel to H.G. Well's War of the Worlds featuring Thomas Edison conquering Mars.



Many people were surprised but pleased to see IEEE at the convention. Several people stopped by who are currently members and several who had been members. More than one was interested in becoming a member and took some of our membership information. We also had a number of kids who stopped by, particularly to get one of the wind-up dancing robots we were giving away. Some of the volunteers from Seattle talked about participating in their local Science Fiction Conventions in the future.

While we at the Convention Rebecca and I spoke to someone representing the 2026 Worldcon, which will be in Los Angeles in 2026. IEEE plans to be there with support from local sections in LA and IEEE Region 6 leadership. I hope that IEEE can participate in more events like these to meet people interested in technology so they can learn about who we are, what we stand for and perhaps to join us in advancing technology for the benefit of humanity.

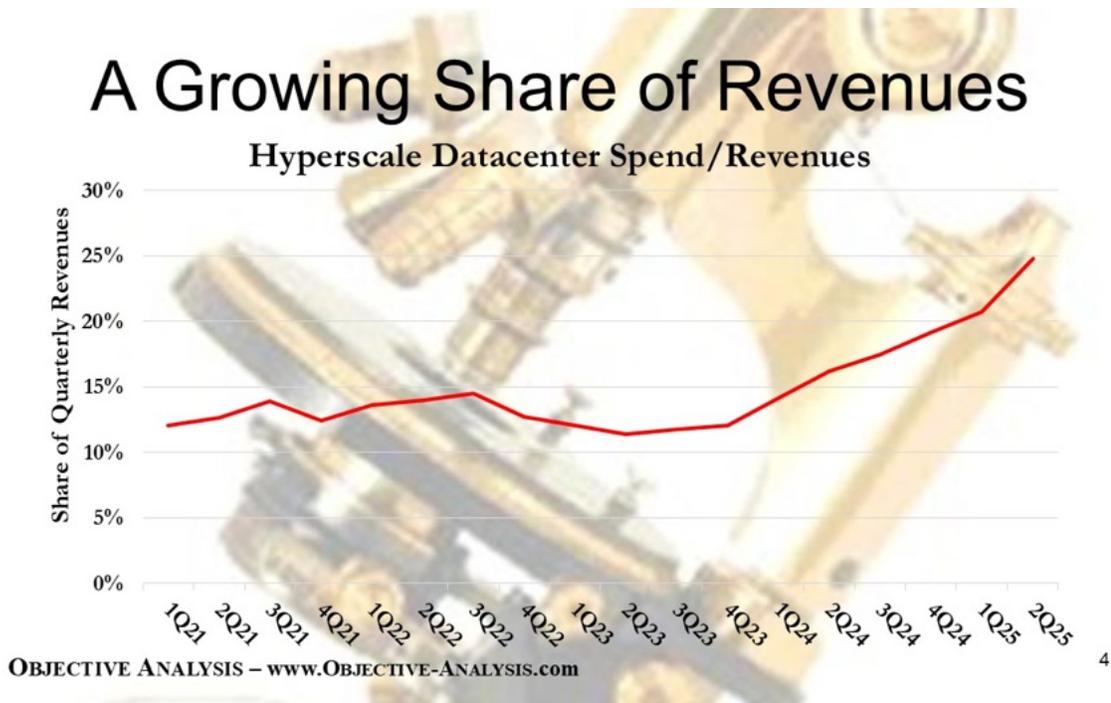
IEEE participated in the 2025 World Science Fiction Convention in Seattle, Washington. We organized sessions on Digital Preservation as well as remote medical monitoring and privacy. Many folks came to see the exhibit that featured some historical artifacts.

4) How Would an AI Recession Impact AI Commercialization?, (Forbes, 8/24/25).

I am at the 2025 IEEE Hot Chips Symposium at Stanford University and there are a lot of presentations and exhibits on semiconductor devices, many focused on artificial intelligence. I will write separately about these but I wanted to provide some thoughts related to recent economic developments and comments about AI that have been talking about a possible bubble in the AI market.

There was an [August 6 posting to Yahoo Finance](#), derived from a Fortune article by Nick Lichtenberg. It points out that in the summer of 2025 consumer spending is slowing. This is likely due to increased unemployment and lower job growth, at least partially due to lower anticipated employee needs with the implementation of AI in industry.

He pointed out that Giant tech companies have spent so much on data centers in 2025 that their spending is now contributing more to U.S. economic growth than consumer spending. My colleague, Jim Handy of [Objective Analysis](#) presented this slide during the 2025 FMS in early August showing that the quarterly ratio of hyperscale data center spending (mostly on infrastructure to support huge AI workflows) to revenue has increased over 10% since 2023 and if the trend continued would soon be double the 2023 ratio.



Slide from talk by Jim Handy at the 2025 FMS

This trend reflects the significant spending on AI infrastructure but also that so far this spending has exceeded AI related revenues. Is this sustainable? A few days ago people started to write about [an MIT study](#), which indicates that for many companies, getting a real financial return from AI has been elusive. The market value of AI related companies took a major hit after the release of this report.

The report is based on a multi-method research design that includes a systematic review of over 300 publicly disclosed AI initiatives, structured interviews with representatives from 52 organizations, and survey responses from 153 senior leaders collected across four major industry conferences. The

report says that Just 5% of integrated AI pilots are extracting millions in value, while the vast majority remain stuck with no measurable P&L impact. This divide does not seem to be driven by model quality or regulation, but seems to be determined by approach.

The report says that popular AI tools like ChatGPT and CoPilot primarily enhance individual productivity, not P&L performance. Most enterprise-wide pilot programs fail due to brittle workflows, lack of contextual learning, and misalignment with day-to-day operations. It also indicates that programs which include external partners familiar with the use of AI see twice the success rate than internal only builds. The report says that the core barrier to scaling these AI systems is not infrastructure, regulation, or talent. It is learning.

Most GenAI systems do not retain feedback, adapt to context, or improve over time. The report indicates that early results suggest that learning-capable systems, when targeted at specific processes, can deliver real value, even without major organizational restructuring. My interpretation of this is that a focus of using AI to enhance human agency, rather than a wholesale replacement of human workers is proving far more effective.

As several industry pundits have been pointing out, the current AI frenzy has many of the hall marks of prior industry bubbles, such as the Internet Bubble in the early 2000's and the railroad bubble in the 19th century. Such a bubble will lead to a major correction and likely a reduction in current investments. But that will not mean that there is no future for AI or that AI work will cease.

However, an AI recession will likely lead to far more focused use of AI and the big models currently being developed to be used in ways that will enhance value for businesses and increase the effectiveness of employees. I think that the best way to think about the current wave of AI is that it is best as an Augmented Intelligence, in particular for increasing the intelligence and agency of individual humans and including human judgement, rather than the wholesale replacement of humans.

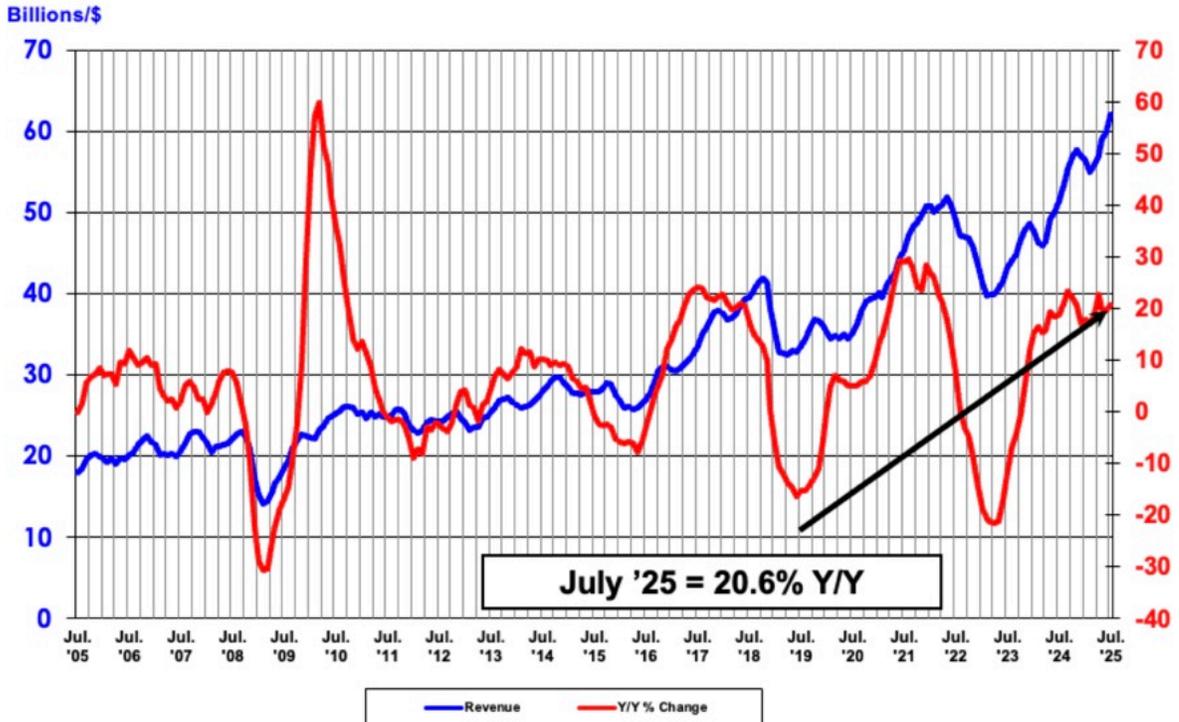
5) Semiconductor Sales and Capital Spending Driving Hot Chips Interest, (Forbes, 9/9/25)

This is another article on some interesting developments from the IEEE Hot Chips Conference at Stanford, but first some interesting semiconductor statistics on industry equipment purchases and semiconductor sales. Continuing these trends could drive the use of NVIDIA networking and interconnect products and a Google Ironwood TPU technology for sharing 1.77PB of High bandwidth memory, HBM.

[The Semiconductor Industry Association, SIA, announced](#) global semiconductor sales were \$62.1 billion during the month of July 2025, an increase of 20.6% compared to the July 2024 total of \$51.5 billion and 3.6% more than the June 2025 total of \$59.9 billion.

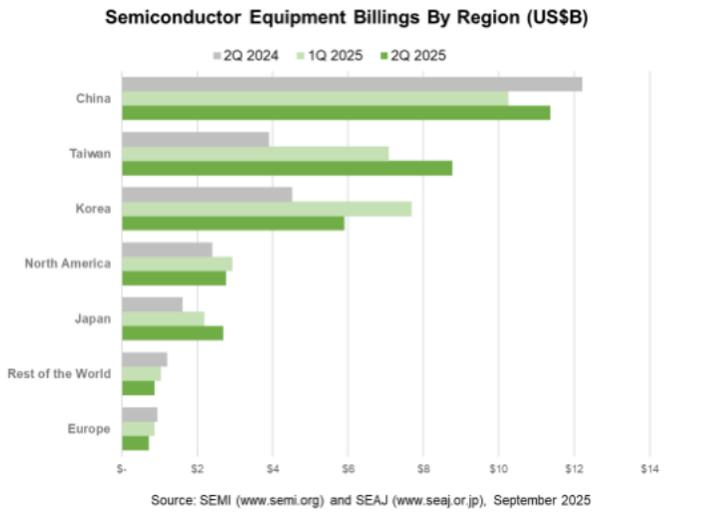
Regionally, year-to-year in July sales were up in the Asia Pacific/All Other (35.6%), Americas (29.3%), China (10.4%), and Europe (5.7%), but decreased in Japan (-6.3%). Month-to-month sales in May increased in the Americas (8.6%) and Asia Pacific/All Other (4.9%), remained steady in Europe (0.0%), and declined in Japan (-0.2%) and China (-1.3%). The chart below shows WSTS tracking of overall semiconductor revenue and year-over-year percent changes.

Worldwide Semiconductor Revenues Year-to-Year Percent Change



Source: WSTS

Making the latest generation of semiconductors takes capital investment and a big part of that investment is in equipment. [SEMI reports](#) that, global semiconductor equipment billings increased 24% year-over-year to US\$33.07 billion in the second quarter of 2025. Second quarter 2025 billings registered a 3% quarter-over-quarter expansion supported by the leading-edge logic, advanced high bandwidth memory (HBM) related DRAM applications, as well as increase in shipments to Asia. The chart below shows where these investments are being made.



In particular, China and Taiwan led in equipment purchases with Taiwan more than doubling its spending from the prior year. North American equipment purchases were modestly higher than a year ago.

The semiconductor industry and the equipment that support that industry are showing robust continuous growth in most parts of the world supporting consumer, industrial and data center applications. The Hot Chips Conference had a particular focus on data center semiconductor developments, including several from Nvidia. At the event Nvidia speakers particularly focused on enterprise inference applications, where AI hopes to achieve its return on investment as pointed out in the slide image below.

Challenges with AI inference in Production

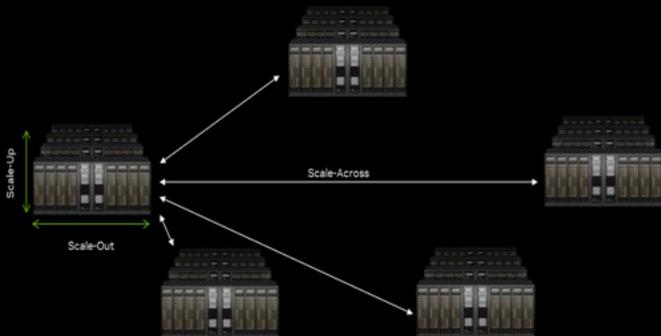
IT leaders chose AI inference solutions that meet enterprise grade requirements

- Fast**
Meets use case and user experience latency constraints
- Interoperable**
Supports multiple development frameworks and deployments
- Integration**
Pre-built integrations with enterprise landscape
- Scalable**
Easily add new models and support spikes in user demand
- Efficient**
High utilization of underlying H/W accelerators
- Cost**
Meets budget constraints and delivers max ROI

Communication within chips and between chiplets and within and between data center racks saw some interesting Nvidia announcements, including the introduction of its Spectrum-XGS Ethernet for supporting distributed AI between remote data centers to create what Nvidia says, will be the world's largest supercomputers.

NVIDIA Introduces Spectrum-XGS Ethernet to Connect Distributed Data Centers into Giga-Scale AI Super-Factories

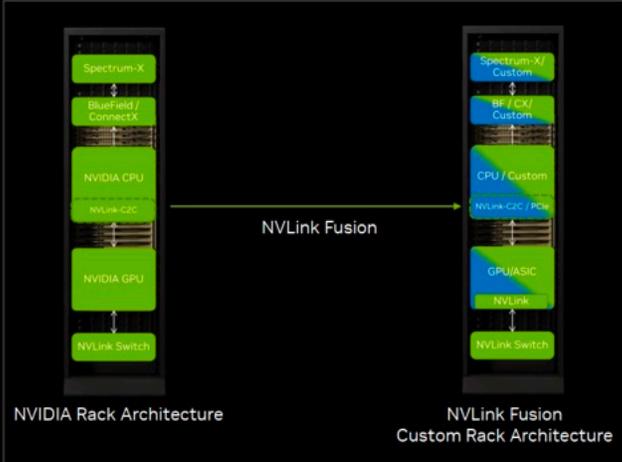
Distributed AI between remote locations, overcoming power and physical limitations



- Scaling AI beyond the data center requires new infrastructure
- Spectrum-XGS unifies multiple data centers into the world's largest supercomputers
- Integrating scale-out and scale-across
- Auto-adjust load balancing based on scale-across distance
- 1.9X higher NCCL multi-site performance

The company also spoke about NVLink Fusion, a rack-scale scale up fabric, shown in concept below.

NVLink Fusion: Scale-Up Performance Through Proven Architecture

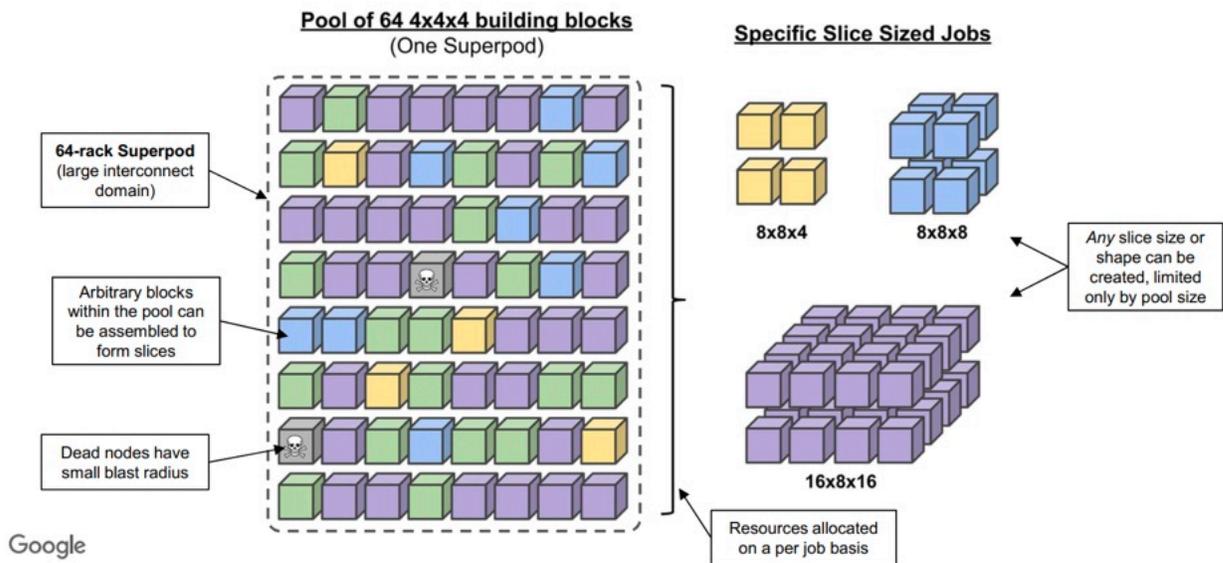


- NVLink hardware and software pedigree deliver leading scale-up performance
- NVLink Fusion gives access to the only rack-scale, scale-up fabric in production
- NVLink Fusion is comprehensive and flexible in its modular mix offering
- Ecosystem is robust, experienced in rack scale production that NVLink Fusion is built upon

Nvidia and others at Hot Chips also talked about co-packaged optics built into switch ASICs to provide high bandwidth and low latency that can be used to connect millions of GPUs. With faster connections between GPUs and memory, with and between racks and between data centers Nvidia projected ambitious plans to create high performance AI training and inference.

Another interesting presentation focused on memory included the discussion by Google of their Ironwood TPU chips which use optical circuit switches, OCSs, to share HBM memory of 1.77PB. The slide below shows this memory sharing which includes significant redundancy in data placement to deal with bad data in individual memories.

TPUv4: 4096 Chips Share Memory Using OCS

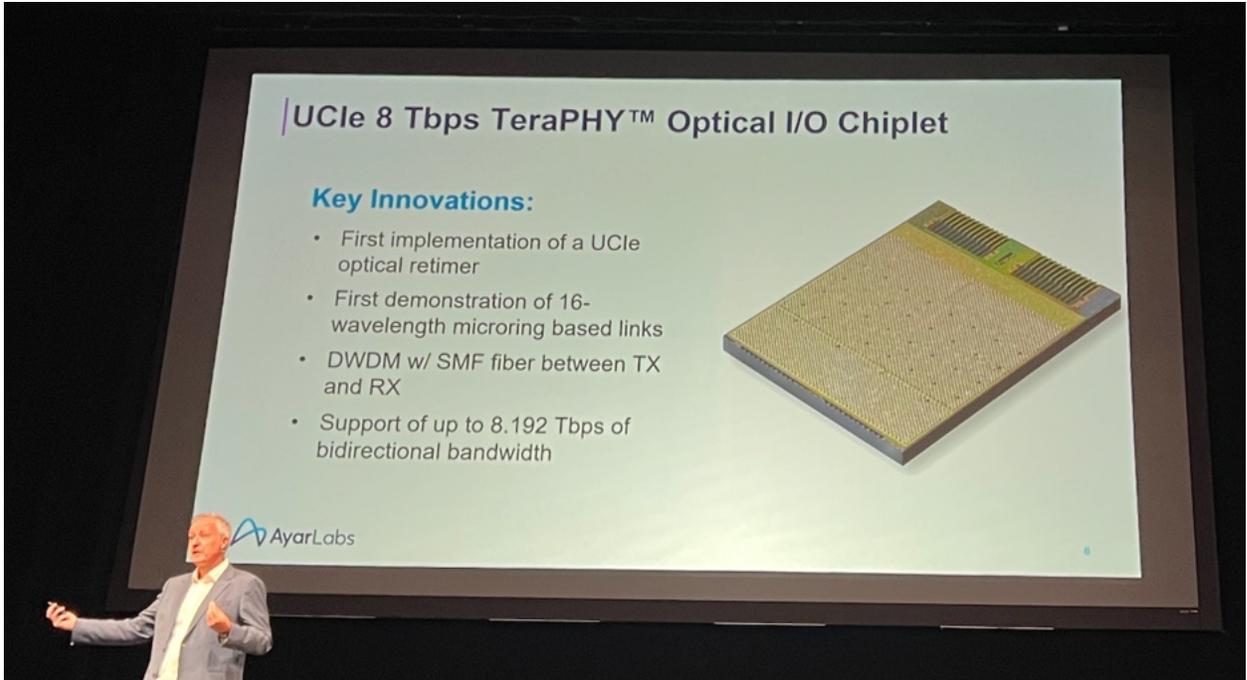


Semi sales & capital spending is up from last year. Continuing this trend may depend upon networking and interconnect technologies from the 2025 Hot Chips Conference.

6) UCIe, Marvell SRAM and DRAM and Rapidus 2nm Fab at the 2025 Hot Chips Conference, 9/3/25).

I attended the IEEE Hot Chips Conference at Stanford and I wanted to share some of the interesting developments related to digital storage and memory technologies and other interesting developments. These include several talks that showed UCIe chiplet interconnect implementations, SRAM and DRAM improvements from Marvell and Rapidus's success with their 2nm semiconductor fab.

The open chiplet interconnect standard, UCIe was implemented on several chips that were showing at Hot Chips. Here are a few examples. The image below is a slide from a AyarLabs optical I/O chiplet.



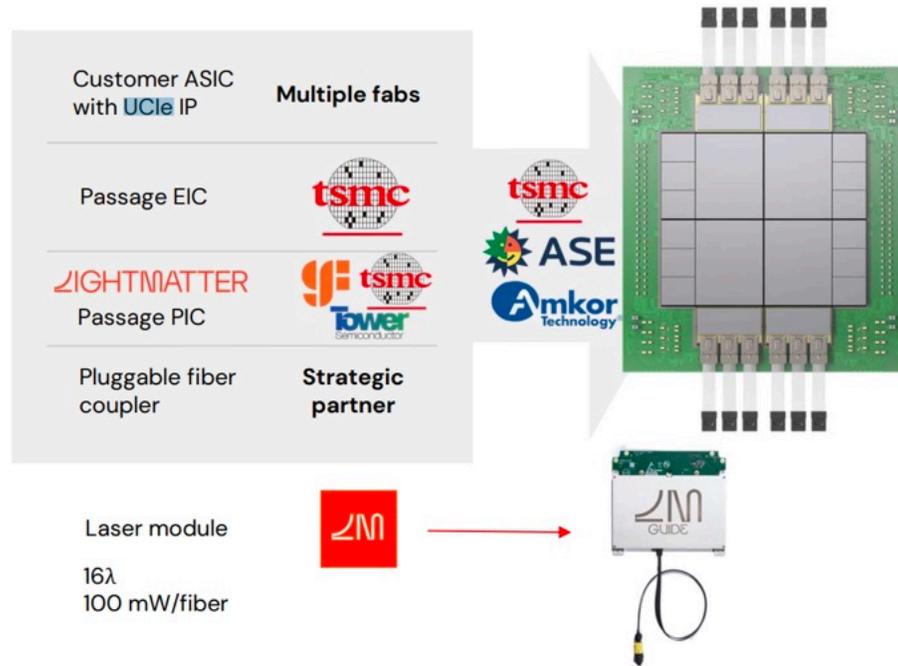
At a Korean-based AI company, Rebellions' exhibit at Hot Chips they were showing a UCle-based AI chiplet as shown below.



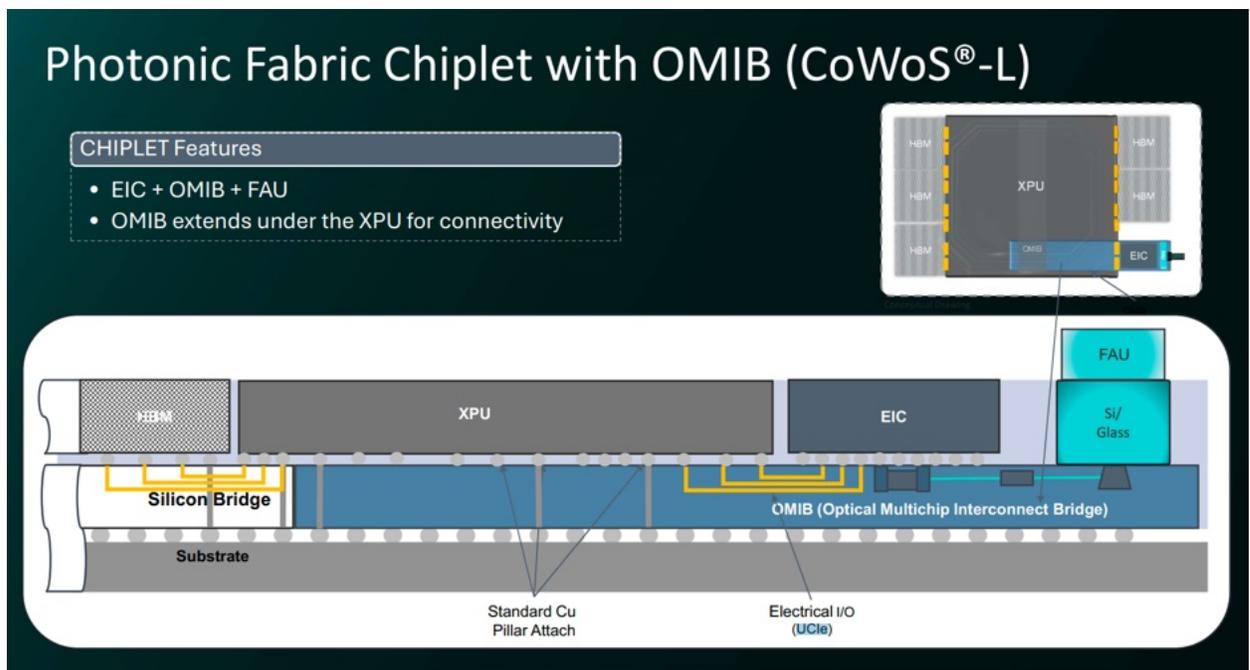
Lightmatter spoke about a 3D interposer which enables an ASIC built with UCle IP and laser communication as shown in the image below.

LIGHTMATTER

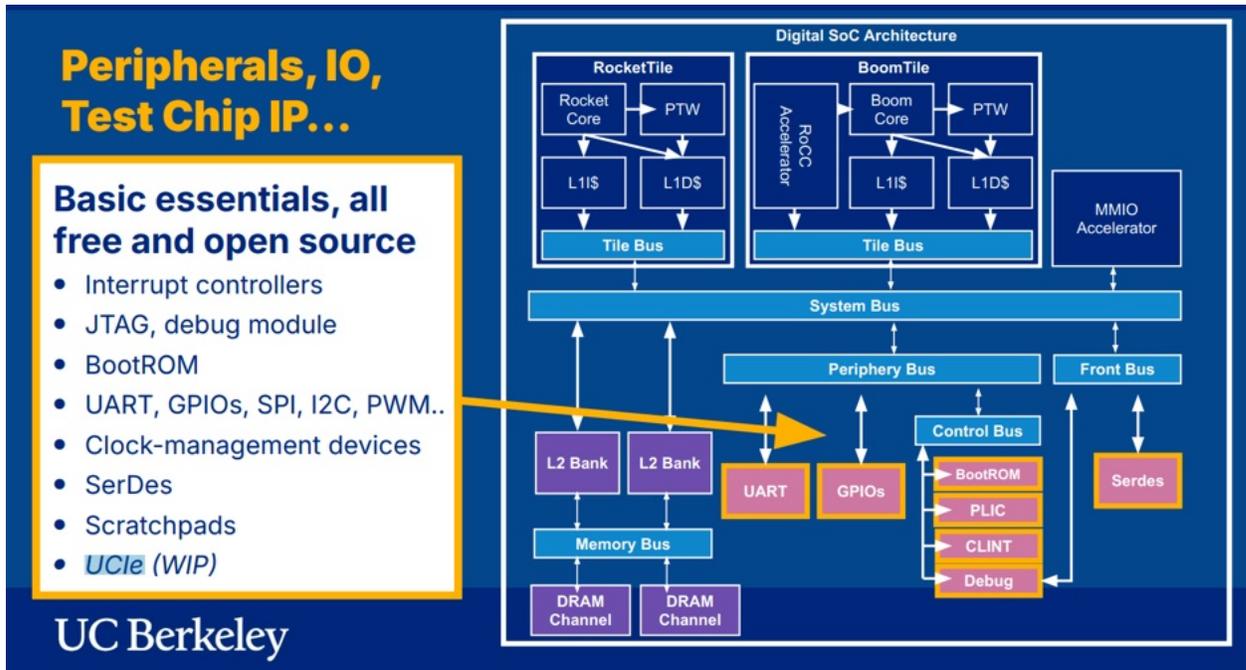
Optimized end-to-end and production-ready



Celestial AI was showing an SoC combining in-die optical I/O with electrical interconnects, including UCle as shown below.

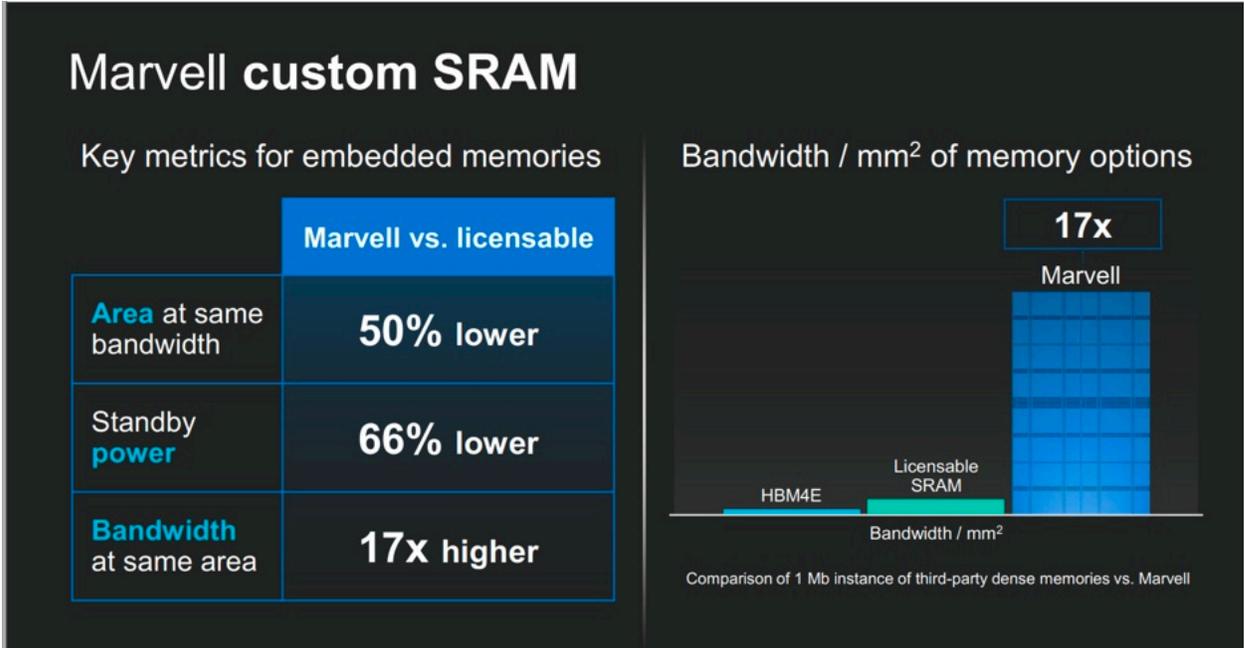


UC Berkeley has been training students on taping out class chips using Chipyard as part of their education, resulting in real chip results. Students report that this is one of their most intense and fun classes at the University and it attracts more than just engineering students. As the slide below shows UCle and other open source technologies are used in making these chips.



Mark Kuemerle, VP of Technology at Marvell gave a talk on what they called a revolution in memory architecture for the data center. In particular the Marvell talk focused on a method for increasing the capacity and bandwidth of static random-access memory, SRAM. Marvell has a 2nm SRAM platform manufactured at TSMC.

The company has made innovations in write assist, stability assist, pioneering high-sigma design modeling to capture tail bits and row plus column redundancy to enable low voltages and high overall yields. As a consequence Marvell's custom SRAM achieves significant advantages over HBM and other embedded IP memory as shown below.



Marvell said they achieve this by running faster and with more ports. The company also discussed its custom HBM architectures using IO chiplets with the companies D2D chiplet interconnect. The company announced at Hot Chips a 65 Gbps/wire Bi-Directional Die-to-Die interface IP in 2nm for the next generation of XPU. They also showed their Structura A and X high capacity memories for near memory accelerators and memory expansion.

Rapidus had a keynote talk on the last day where Rapidus CEO Atsuyoshi Kolke spoke about the building and initial operation of the company's 2nm semiconductor fab on the Island of Hokkaido in Japan. In April the company had their first EUV test, shown below.



Rapidus plans to use an AI tool they call Raads in designing and manufacturing their chips and a grid transfer system for wafers that allows variable workflows in the same manufacturing environment. This is part of what they DMCO, which co-optimizes design and manufacturing. They said this can reduce batch and single wafer production run times by 50% or more. Single wafer process times can be reduced to 15 days. The company has a strategic collaboration with Keysight to improve yields and achieve high-precision process design kits (PDKs).

Hot Chips showed UCIe chiplet interconnects, Marvell's custom SRAM and HBM developments and Rapidus details on their 2nm semiconductor fab.

7)) Memory and Design Advances from the AI Infrastructure Summit, (Forbes 9/11/25)

This article discusses memory and chip and system design talks at the 2025 AI Infra Summit in Santa Clara, CA by Kove, Pliops and Cadence.

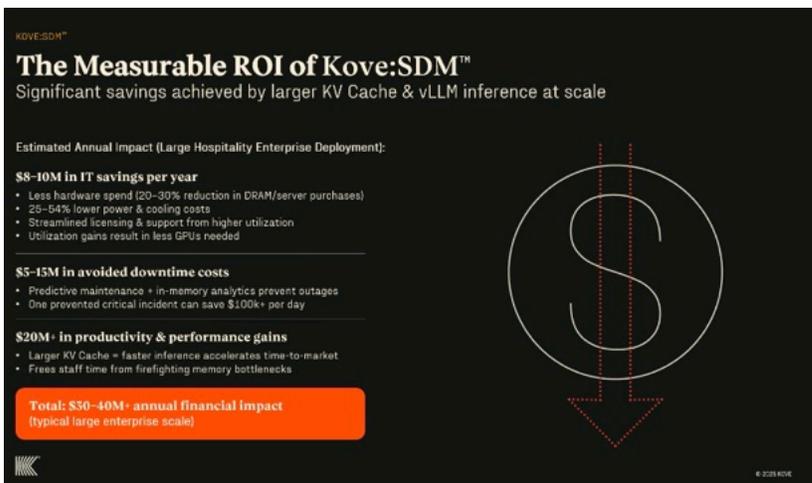
John Overton from Kove gave a talk about their Linux based memory software that runs on any hardware and allows sharing memory between servers to increase memory utilization, CPU and GPU utilization. He pointed out that GPUs have been scaling as have CPUs, but conventional memory systems do not, leading to overprovisioning in servers and processing bottlenecks.

He said that Kove SDM software can be installed in a system in 15 minutes and allows unlimited memory accessed from virtualized elastic memory pools across servers and supports up to 64PiB of DRAM per process. He also said that the software can hide latency, making memory look local to a CPU, even when it isn't. He said this can work across Infiniband and RoCE fabrics.

In particular he said Kove SDM can hide latency between memory that is over 150m away. The image below shows the resulting AI performance improvements.



He said that for AI inference, with Kove it can run 3-5X faster and that this had been shown working with partners such as Redhat and SuperMicro to show key-value, KV, cache at scale. These are results from benchmarks using Redis and Valkey. The slide below summarizes their claims for ROI improvements using Kove. He also said it can improve vector search and enable larger concurrent processing.



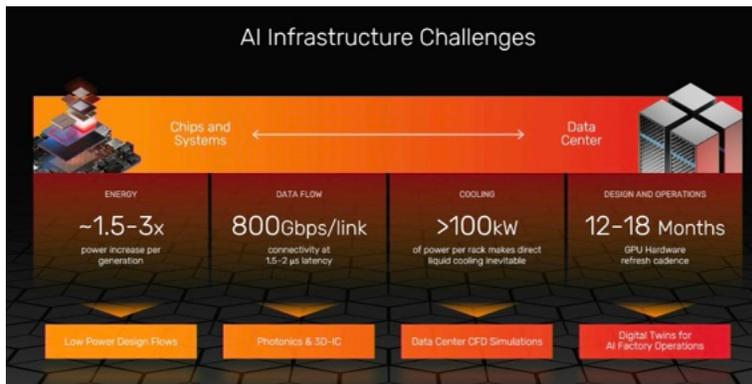
Pliops was showcasing its XDP LightingAI at the Summit, a GenAI native memory stack to power inference and retrieval workloads for hyperscale and enterprise applications. The product consists of an ASIC, the Extreme Data

Processor, XDP, a software stack and distributed nodes and uses a GPU-initiated Key-Value I/O interface.

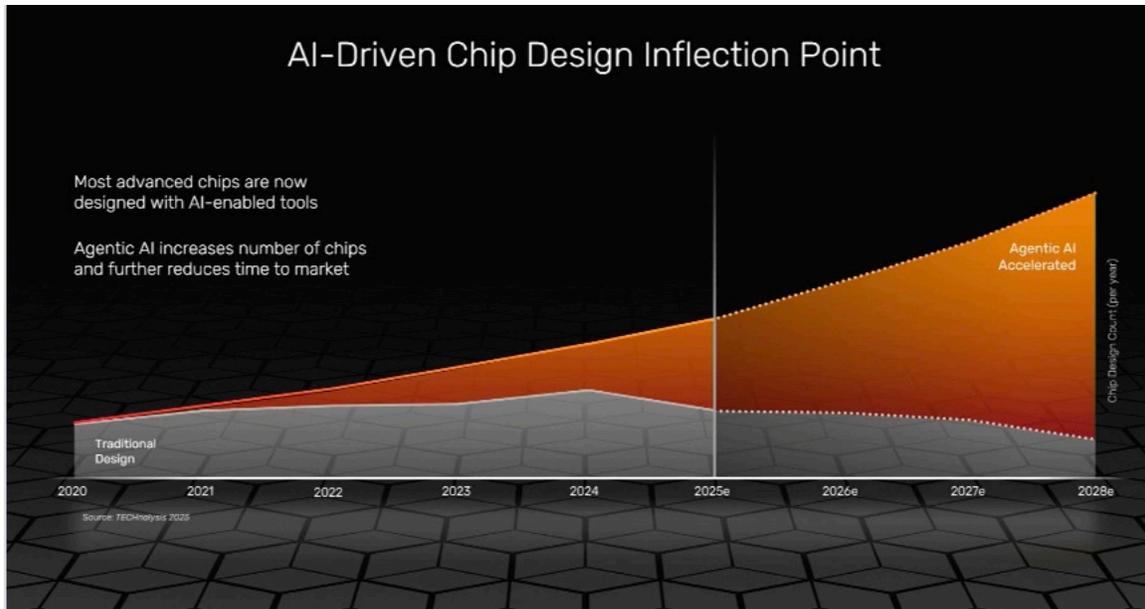
According to Pliops, deploying Pliops XDP LightningAI in data centers offers significant cost savings. Instead of needing four times as many GPU servers to support 4X more users for LLM inferencing, you can maintain your existing GPU infrastructure and add XDP LightningAI servers. This results in a 67% optimization in rack space, a 66% reduction in power consumption, 58% annual OpEx savings, and a 69% decrease in initial investment costs.

Pliops is collaborating with Tensormesh, inference optimization software. The company says combining LightningAI memory acceleration with Tensormesh's shared KV cache architecture results in fast time-to-first token and GPU savings across multi-GPU clusters.

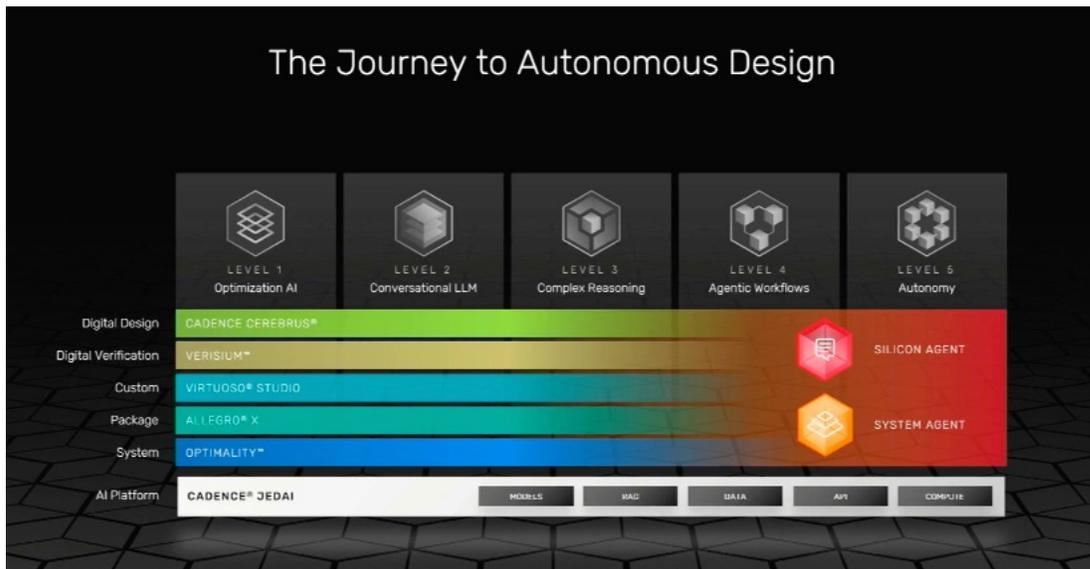
Another interesting talk was that of Charles Alpert from Cadence. Cadence is an electronic design software company that enables device and system simulation and has been doing this for decades. He spoke about the various challenges in AI infrastructure as shown in the slide below. These include energy consumption—driving GW data centers, thermal management and the time to operationalize that infrastructure.



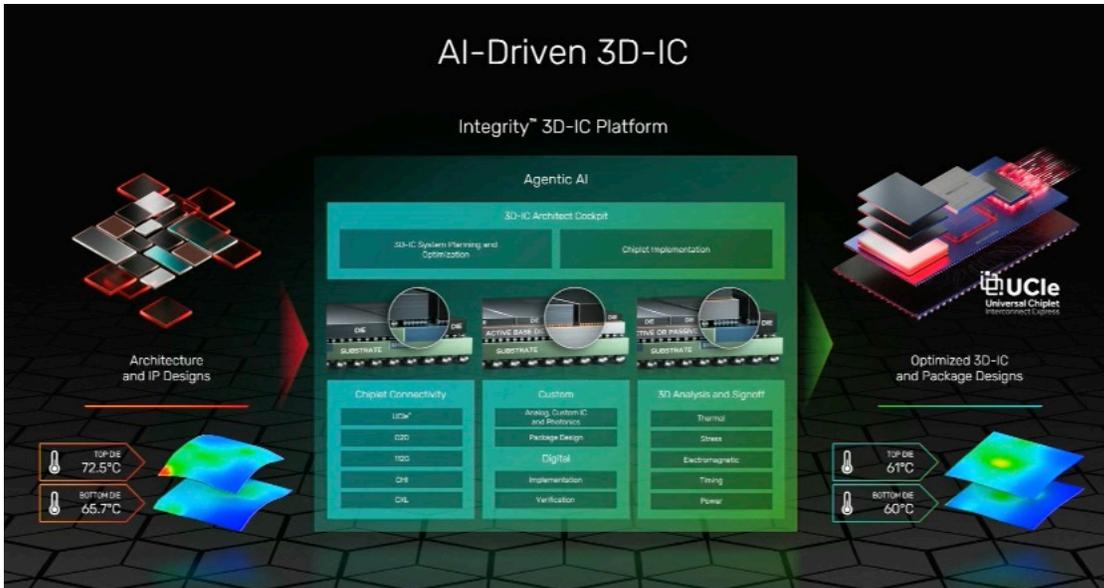
He talked about how this can be improved by adding AI for design to design for AI to create continuous improvements in data centers and the devices that work with them in a virtuous cycle. Cadence has tools for data center design as well as its traditional semiconductor design and in semiconductors their tools can also work with 3D stacks of die and include Multiphysics digital twin simulation. He said that over half the chips built today use AI technology and that this will accelerate to 90% in the next few years using agentic AI as shown below.



Agentic AI should lead to levels of autonomous design, similar to those used for talking about autonomous driving as shown below. In his talk he went into detail on each of these levels for increasing autonomy.

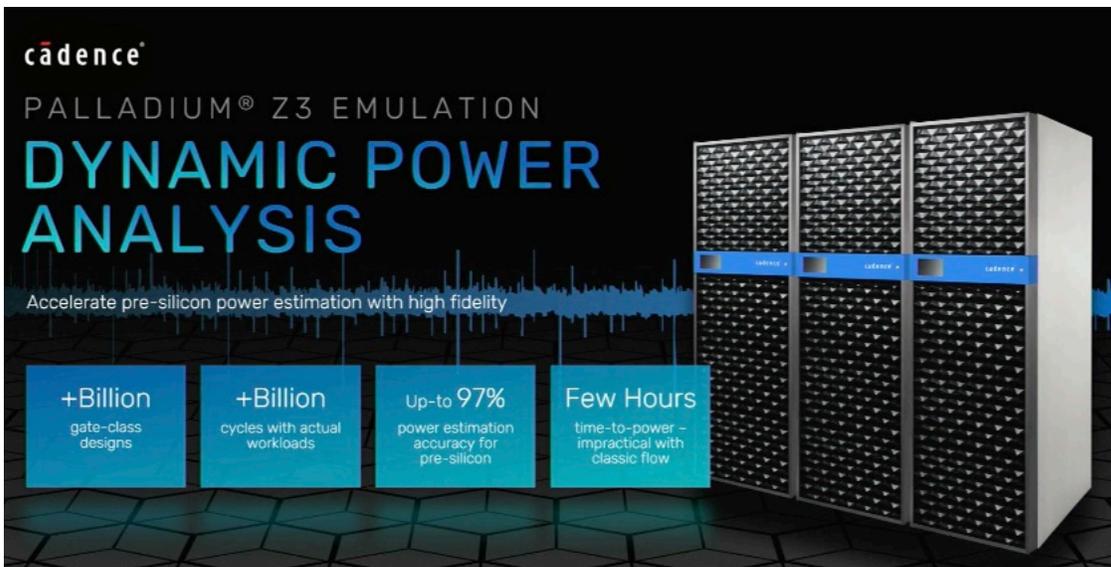


He also talked about how these new EDA tools with multi-physics capability will enable designing 3D devices made from stacking semiconductor die, often called heterogeneous integration as shown in the figure below. Doing this requires massive system level integration and is resource intensive.



In particular he pointed out that these 3D workloads will benefit from improvements in the number and capability of GPUs.

Last year Cadence announced its Palladium Emulator, shown below. [Jensen Huang from Nvidia](#) called it an essential tool for designing NVIDIA's most advanced chips, including the Blackwell series.



Cadence has expanded this technology beyond chip design towards creating digital twins of data centers including all the functional components. Their Millennium M2000 system enables much faster chip design as well as system design. These are tools for Cadence's Digital Twin Ecosystem for data center design as shown below.



Kove talked about massive memory sharing to accelerate AI inference, Pliops was showing their AI-native memory stack and Cadence talked about complete digital twin data center design at the 2025 AI Infra Summit.

8) Pure Storage Sustainability Report Shows How to Reduce Carbon Footprint and Waste, 9/14/25)

Many of the companies making commonly used digital storage and memory products are issuing sustainability public impact reports as an adjunct their regular annual reports. We have talked about some of these efforts at hard disk drive companies such as Western Digital and Seagate but the flash memory storage companies are also doing their part.

An example is Pure Storage with their recent report issued this month covering February 2024 through February 2025. The company says that its FY25 GHG emissions and purchased generated renewable energy (RE) were verified by Apex, a third-party, in accordance with ISO 14064-3.

Science Based Targets Initiative, SBTi, validated the targets below and classified the company's Scope 1 and 2 climate ambition as in line with a 1.5 degrees Celsius trajectory. The report also said that Pure Storage is also committed to be net zero by 2040 covering its Scope 1 and market-based Scope 2 emissions.

**Pure Storage Near-term
Science Based Targets**

42%

**Reduce absolute scope 1 and 2 GHG emissions
42% by FY30 from a FY23 base year.**

51.6%

**Reduce scope 3 GHG emissions from use of
sold products 51.6% per petabyte by FY30.**

45%

**45% of suppliers by spend, covering purchased goods and
services, will have science-based targets by FY29.**

The company's Enterprise Data Cloud creates a unified data storage system across Pure Storage arrays and includes modular non-disruptive upgrades that extend hardware life. A circular design philosophy helps avoid e-waste.

In particular the company's Evergreen upgrades allow customers to continuously modernize hardware in place, without disruption or full system replacements. Arrays stay productive for 10+ years which eliminates two to three full hardware replacements and the associated data migrations that would be required. As a result, embodied emissions, e-waste, and the administrative burden of the conventional hardware refresh cycle are greatly reduced.

The company was also finding ways to increase overall storage capacity with existing hardware as well as providing new hardware with greater storage capacity in the same rack space. Pure said it had optimized its Purity operating environment to more efficiently use raw storage capacity. As a result, customers gained 5% more usable capacity on FlashArray and 6% more usable capacity on FlashBlade through a nondisruptive software upgrade.

The company has also found other ways to extend the life of its products. Pure began work on a new line of remanufactured FlashArray products in FY25. These

products use refurbished controllers which can reduce the manufacturing portion of GHG impact by approximately 18%, while also providing a circular refurbishment and redeployment path for all upgraded FlashArray controllers. Our first remanufactured product, FlashArray//RC20™, is a capacity-optimized, entry-level enterprise storage array, built using a combination of renewed and brand-new components.

Regarding end of life, Pure says it takes a proactive, integrated approach to product end-of-life, reflecting the same principles of modularity, efficiency, and circularity that define its Evergreen architecture. When systems reach the end of their service term, returned hardware is evaluated for repair, refurbishment, or responsible recycling. Thanks to cross-generational compatibility and modular design, many components—such as controllers and DFMs—can be refurbished and redeployed within the Pure Storage product ecosystem.

This reduces waste, extends product lifespan, and minimizes the use of new materials. For components that cannot be refurbished or remanufactured, Pure Storage partners with certified e-waste vendors who follow rigorous standards for data security, environmental compliance, and material recovery—ensuring consistent, high-integrity outcomes across the globe, while meeting applicable regional regulations. For example, our European e-waste partner's recycling facilities have been optimized to the point that less than 0.2%, on average, goes to landfill per year.

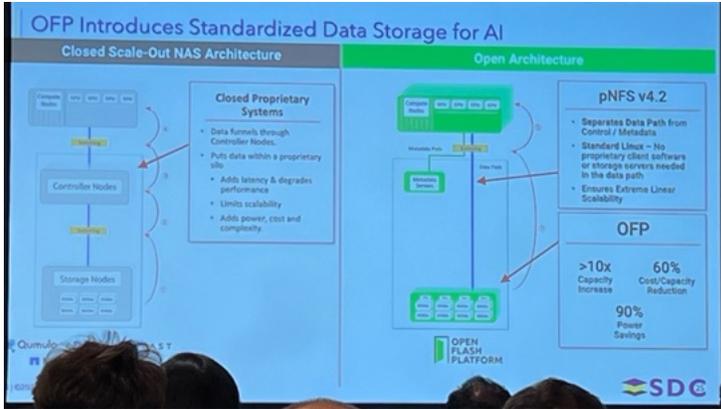
Digital storage is a major component in enterprise computing and data centers. Pure is showing how digital storage to power AI and other needs can be useful and sustainable.

9) In AI, No Data is Ever Cold, (Forbes Blog, 9/15/25)

So said David Flynn, CEO of Hammerspace at the 2025 SNIA SDC in Santa Clara, CA. Among his other accomplishments, David was CEO of FusionIO, the company that pioneered the SSD interface that eventually became NVMe.

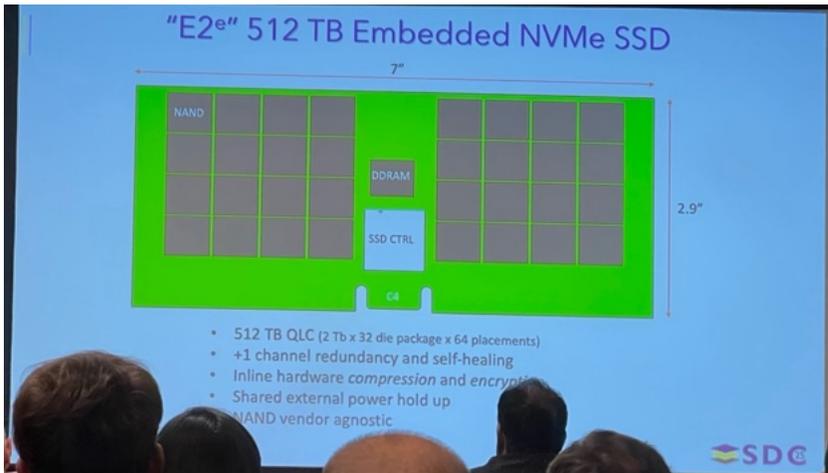
He gave a clear and well thought out talk about the Open Flash Platform, OFP, based upon NFS in Linux and a metadata server with data orchestration in a separate control plane from the data, that would dramatically improve system performance by avoiding system controllers and moving data directly from the storage devices to processors.

The image below compares the OFP storage architecture on the right with traditional NAS architecture on the left. By storing separating the data path from the metadata and data orchestration OFP gets rid of the controller used in conventional storage systems and it is possible to pursue linear scaling, particularly to provide higher performance to meet the needs of AI workflows.

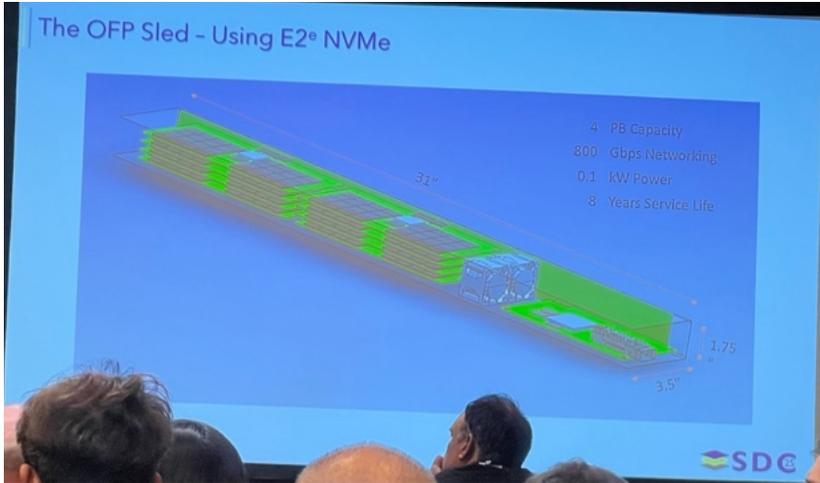


By eliminating the controller and other potential bottlenecks using OFP he says that the cost can be reduced by 60%, the capacity increased by 10X and it uses 10% of the power of conventional SSD storage.

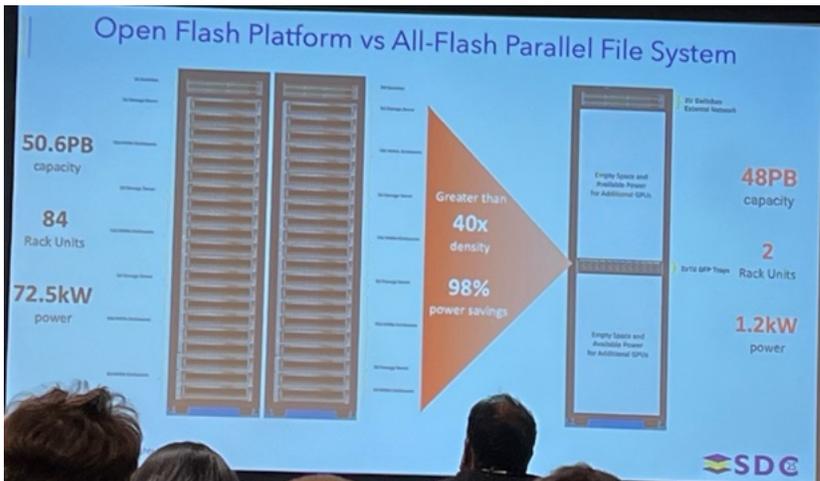
He introduced the idea of a OFP sled that basically consists of banks of flash memory. This could use conventional SSDs but he also suggested that more dense packages could be built for these sleds such as the E2 512TB Embedded NVMe SSD shown below.



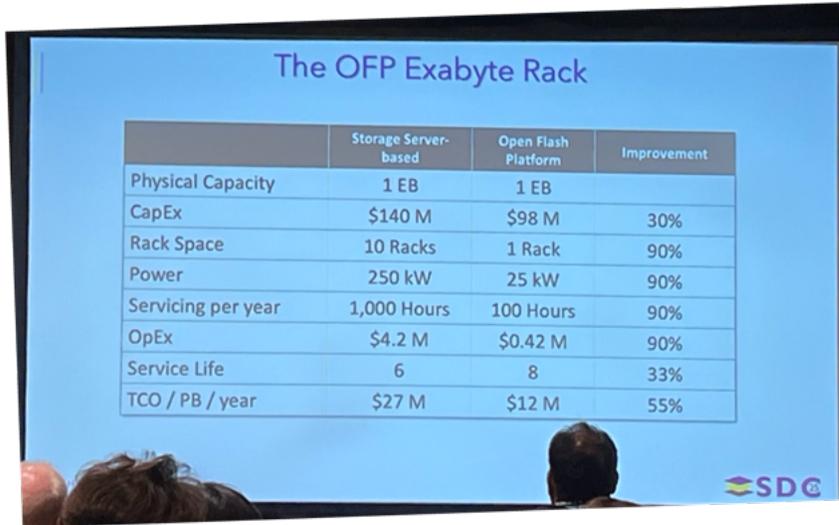
These could be stacked into the sled configuration shown below with a DPU and a fiber interface to form a storage NIC.



Rows of these sleds could be packed next to each other in a 1U shelf and these shelves stacked in a rack. The resulting configuration provides much denser storage than conventional all flash parallel file systems. The figure below shows 2 fully packed 1U OFP shelves in one rack compared to two fully populated racks of conventional flash storage. Both units have about the same capacity, about 50PB, but the OFP has 40X greater storage density and 98% lower power consumption.



A fully populated OFP rack could provide an Exabyte with 252 sleds, and provide 200 Tbps data rates. Idle power would be just 25KW and 40PB/kW and 8Tbps/kW. The table below compares the OFP 1EB storage platform versus conventional parallel file flash storage. One rack achieves this versus 10 racks, with 10X less power consumption, 10X less service time per year and longer service life and 10X less operating costs per year.



	Storage Server-based	Open Flash Platform	Improvement
Physical Capacity	1 EB	1 EB	
CapEx	\$140 M	\$98 M	30%
Rack Space	10 Racks	1 Rack	90%
Power	250 kW	25 kW	90%
Servicing per year	1,000 Hours	100 Hours	90%
OpEx	\$4.2 M	\$0.42 M	90%
Service Life	6	8	33%
TCO / PB / year	\$27 M	\$12 M	55%

One of the things David said is that with such a system one could tolerate a lot more failures and degradation in place, part of why he estimated a longer service life, although he also said that the parts on the sled that might fail—such as the DPU, would be modular and easily replaced in house.

Hammerspace is working with several organizations to make OFP happen, including Los Alamos National Laboratory, Linux, SK hynix ScaleFlus and Xsight. It was also obvious from discussion after his talk that other prominent storage companies are also working on OFP. He pointed out that many of these ideas, such as having storage directly connected to an Ethernet fabric have been around for years, but until the rise of AI there was not the demand for storage performance that there is today.

At the 2025 SNIA SDC Hammerspace’s David Flynn presented the concept for a Linux NFS-based Open Flash Platform that can provide the capacity and performance needed for modern AI workflows.

10) AI and Fast File Transfers Improve Media and Entertainment Preservation and Workflows, 9/20/25)

I wasn’t able to attend the 2025 IBC in Amsterdam this year because of prior IEEE commitments. However, I did pick up on some related media and entertainment storage news I wanted to share. These are Iron Mountain use of AI in digital conversion and preservation and MASV’s large file transfer technology.

I had a chance to talk with Andrea Kalas from Iron Mountain about digital conversion and data preservation. She spoke to me about content discovery in archiving using Google AI combined with their digital asset management capabilities.

She said that they hope to do a lot with this for clients who need to digitize large volumes of their collections that are still on videotape or even film to make them

searchable. AI is used to analyze and identify images on these digitized archives and this information is then searchable for easier access and use of historical content. The image below shows Iron Mountain technicians reading data from tape that is then digitized and stored on modern storage systems.

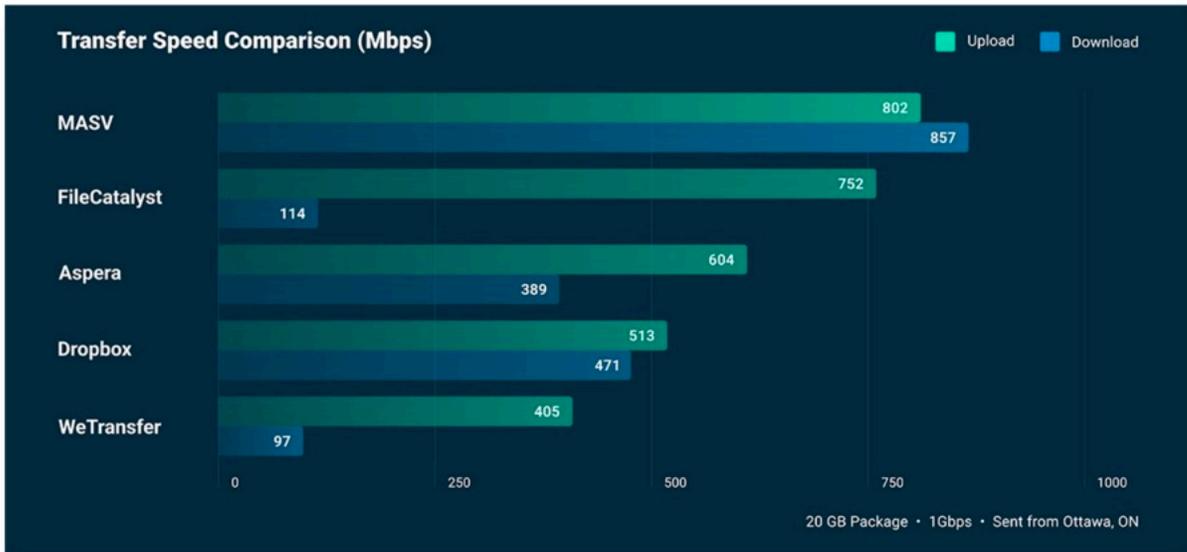


I asked Andrea if they were training their own models specialized for this type of content for this analysis or were they using other people's foundational models. She said that for the present they are working with Google using their models but that they could build their own models or work with organizations that have models that are specific to their own content.

I asked for examples of the type of content they have been working with. They said they have worked with a news organization as well as a private collection for this enhanced digital content access. This digitization and AI advanced search will make reuse of older content easier and unlock powerful new revenue streams.

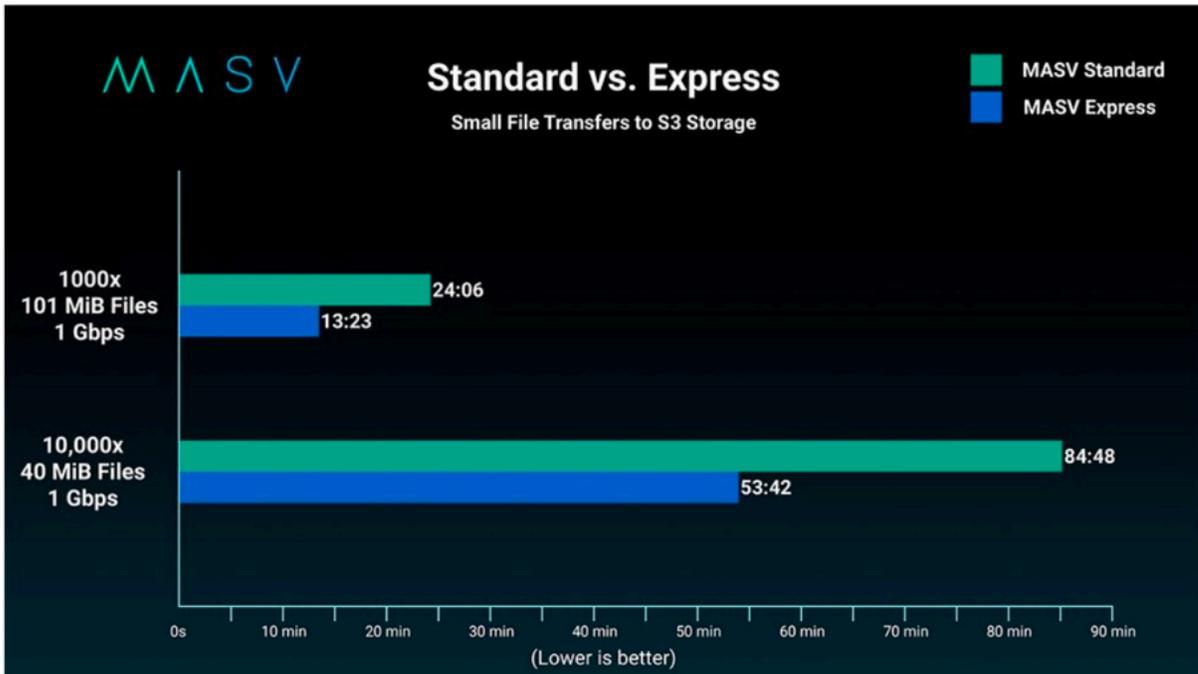
I also spoke with Gene Villeneuve and Conner Stirling from MASV. MASV has a large file transfer platform that is geared to the needs of the media and entertainment industry. They said that this allows high-speed transfers with fault protection and checkpointing for recovery from interruptions, is secure and works with any browser to deliver content of any size. It also can integrate with workflow and enterprise tools supporting all types of media and entertainment content.

There are a number of large file transfer products on the market. MASV provided this comparison of their product versus others showing superior performance and close upload and download performance.



Working over HTTPS MASV can bond transport through several network connections to increase the file transfer speed. They told me about a customer on a remote island in Scotland that ended up combining 3-4 starlink connections to achieve fast upload speeds. MASV integrates with major cloud and cloud-connected storage providers and media collaboration platforms. They said that FramelO is one of their most popular integrations.

They also have an express versus standard MASV version that can provide even faster data transfers as shown below.



MASV is working on a second version of Express that can do a direct file transfer

without zipping and integrates with object-based storage and includes intelligent file routing that can deliver files to multiple users from a single upload. In addition to cloud storage it can also provide data to desktop storage for free lancers. Iron Mountain is using AI in its digital conversion and preservation workflows. MASV is offering very fast and robust large file transfer technology.

11) SNIA SDC on Sstorage.ai, OCP Storage Work and How Chiplelets Can Reduce Power Consumption (Forbes Blog, 9/21/25)

At the 2025 SNIA SDC new initiatives were discussed to provide better data flows to support AI workflows. OCP partners with SNIA and they talked about storage progress in that hyperscale data center consortium. My colleague, Jim Handy talked about how chiplet and other new semiconductor packaging approaches could replace volatile with non-volatile memories and reduce some of the energy consumption predictions for AI growth.

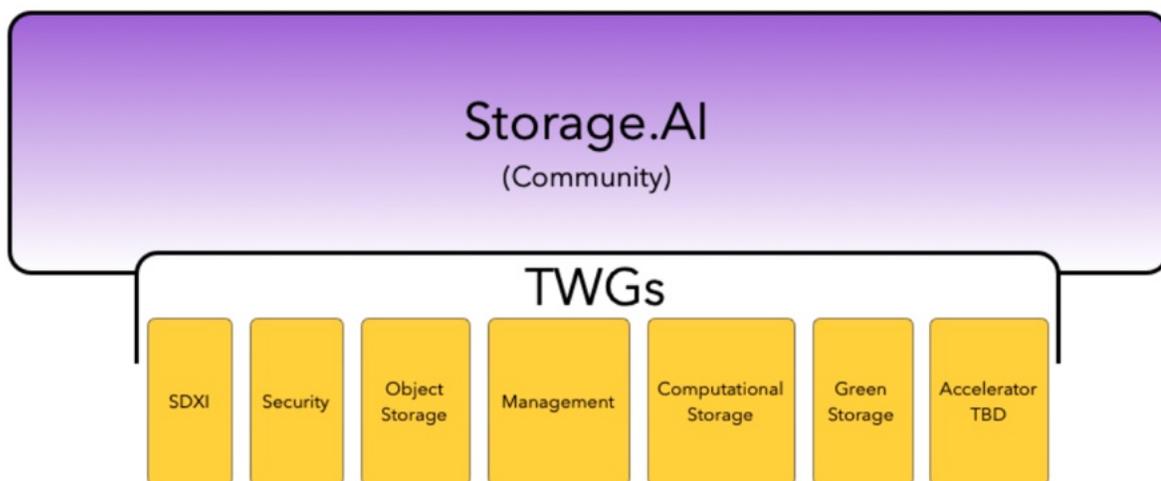
J. Metz, Chair of the SNIA Board of Directors gave a talk at the 2025 Storage Developers Conference, SDC, about a SNIA initiative to create storage systems made for AI, storage.AI. He said that the purpose of this effort is to create a vendor-neutral, open-standards effort tackling the most pressing AI data service challenges—such as memory tiering, latency, data movement, compute-near-storage, and storage efficiency.

The goals of Storage.AI are to:

- Reduce I/O data amplification optimize I/O consumption for accelerators
- Efficient, secure, and reliable data movement through the AI workload lifecycle
- New initiation and consumption model for accelerators
- Standardized hardware/software interface definitions
- Open programming models
- Secure data movement and capacity

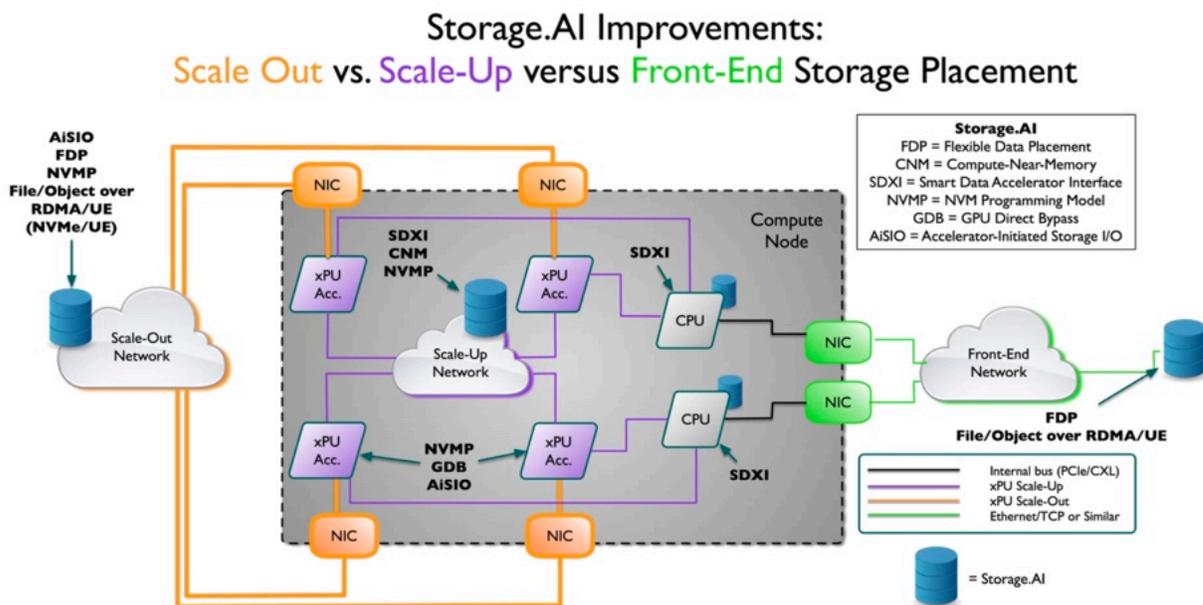
Existing SNIA standard initiatives may help enable Storage.AI including SDXI, smart data acceleration interface, that enable processor-agnostic DMA acceleration and data transformation. The computational storage API and architecture enables near-data compute and the NVM programming model enables a unified software interface for accessing memory tiers.

Swordfish and the Redfish extensions enable storage management while the Object Drive Workgroup enables standard interfaces for object storage. Flexible Data Placement APIs enable optimized data layout and streaming throughput, Security of course protects data and the Green initiative enables power and thermal solutions. The figure below illustrates how all of these existing technical working groups can support Storage.AI.



Members of Storage.AI are interested in creating technical workgroup projects for instance, on File and objects over RDMA/UE that will enable hybrid file and object backends combined with remote memory access. Accelerator Director CPU Bypass will enable bypassing CPUs for data movement between accelerators and data and Accelerator-Initiated Storage I/O to enable the reduction of CPU bottlenecks.

The figure below, from Metz's talk, shows the perceived improvements from Storage.AI for AI workflows.



There is a broad ecosystem of industry efforts that are partnering on Storage.AI including the Ultra Ethernet Consortium (UEC), NVMe, Open Compute Platform (OCP), Distributed Management Taskforce (DMTF), PCI-SIG, Greengrid and soon, UALink.

The OCP Global Summit is coming up in a few weeks and Ross Stenfort from Meta and Lee Prewitt from Microsoft gave another keynote on OCP storage projects. They showed a table of recent OCP work, shown below.

OCP Storage Roadmap of Success

2023	NVMe HDD rev. 1.0	[spec]	Microsoft, Seagate, Western Digital
	Datacenter SAS-SATA Device rev. 1.0	[spec]	HPE, Meta, Microsoft
	Datacenter NVMe SSD rev. 2.5	[spec]	Dell, Google, HPE, Meta, Microsoft
	NVMe Telemetry Scripts	[scripts]	Samsung
	2 nd Annual OCP Storage Tech Talk	[virtual]	May 2023
2024	Meta and Google Open-Source SSD Test Cases	Current	OCP Test Framework (Meta) , OCP Storage Test Cases (Meta) , ocp-diag-ssd-qual (Google)
	3 rd Annual OCP Storage Tech Talk	[virtual]	May 2024
	EDSFF SSD Reference Design	[spec]	ScaleFlux PCIe Gen 5 EDSFF NVMe SSD Reference Design
	Datacenter NVMe SSD rev. 2.6	[spec]	Dell, Google, HPE, Meta, Microsoft
	L.O.C.K. rev. 0.5	[spec]	Google, Kioxia, Microsoft, Samsung, Solidigm
	HDD Dynamics Common Language and Tools Phase 2	[link]	Dell, IEIT Systems, Lenovo, Seagate, WD, ZT Systems
	HDD Dynamics Combined Acoustic & Vibration Surrogate	[white paper]	Dell, IEIT Systems, Lenovo, Seagate, WD, ZT Systems
2025	OCP Marketplace (Products/Storage)	[link]	20+ SSDs and 5 SSD Testing Services
	Datacenter SAS-SATA Device rev. 1.5	[spec]	HPE, Dell, Meta, Microsoft
	Datacenter NVMe SSD rev. 2.7	Coming Soon	Dell, Google, HPE, Meta, Microsoft
	L.O.C.K. rev. 1.0	Coming Soon	Google, Kioxia, Microsoft, Samsung, Solidigm
	4 th Annual OCP Storage Tech Talk	[link]	May 2025

In particular OCP has contributed to the development of the NVMe SSD for data center specifications. The talk pointed out key features for managing and deploying digital storage at scale. These include the OCP Health Information Extended Log for providing telemetry metrics based on at-scale deployments, the OCP Latency Monitoring Feature for isolating, monitoring and debugging latency spikes at scale. The OCP Formatted Telemetry for Human Readable Logs provides customers useful telemetry with improved security and the Open-Source OCP NVMe Command Line Interface for open-source tooling.

A Hardware Component Log provides manufacturing information to customers. Device Self-Test Improvements provide universal failing segment codes and Device Self-Reporting Power measurements are made more useful at scale.

My colleague on our new non-volatile memory report, Jim Handy ended up giving our talk on The Processor Chip of the Future. This focused on the role of chipllets in future semiconductor packages and how chipllets and future stacked die packaging will enable the use of new non-volatile memories. This is important because volatile memories, like DRAM, are a significant contributor to the energy consumption in data centers. Below are some of the ways in which this contributor to data center energy consumption could be reduced.

Approaches Under Investigation

- Replace DRAM with a new memory technology
 - MRAM, ReRAM, FRAM, PCM
 - Intel's Optane proved this to be very difficult
 - But Sandisk plans to try this again
- Move memory into the processor
 - HBM is a small step in that direction
- Move processing into the memory
 - Digital compute-in-memory
 - Analog neural networks
 - New memories are good for this approach
- Support persistence in memory & cache



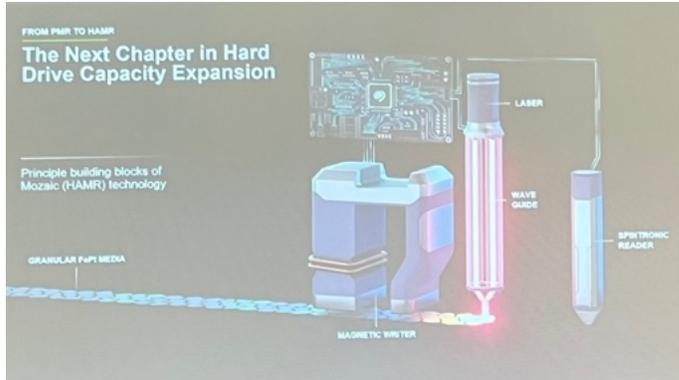
At the 2025 SNIA SDC there were talks about Storage.AI to improve data flow for AI Workloads as well as OCP storage initiatives. Non-volatile memories offer ways to reduce energy consumption in AI.

12) Hard Disk Drive Companies Could Become Dominant Manufacturers of Lasers with HAMR, (Forbes Blog, 9/24/25)

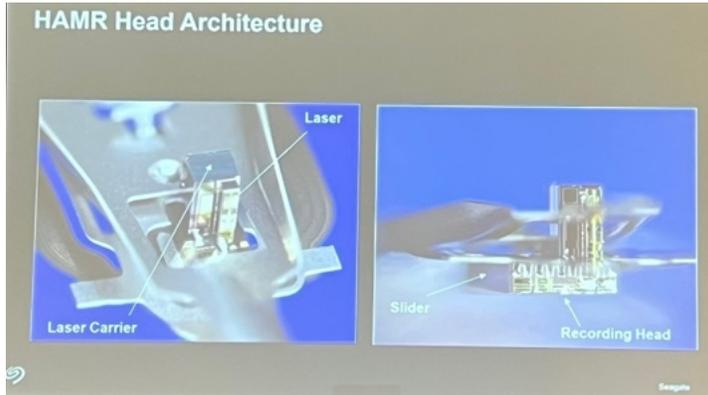
This year marks the 45th year of the International Conference on Massive Storage Systems and Technology (MSST). It is the oldest continuous IEEE conference on digital storage technology. Topics at the event ranged from using automation for storage management code and setting up Retrieval-Augmented Generation, RAG, to the history of network file systems, NFS, and progress on PNFS, to object storage to developments in hard disk drives, HDDs, and LTO magnetic tape.

At the conference there were lots of talks on the second day about the history of NFS, as well as several talks about developments in parallel network file systems, PNFS. Many people involved in the creation and evolution of NFS 40 years ago were in attendance and on the second day there was a 40th anniversary part for NFS. There were also talks on object storage and various talks on digital storage for AI workflows.

Riyan Mendosa, from Seagate, gave a talk on the History of heat assisted magnetic recording, HAMR, and the Path to 10TB/disk and beyond. Seagate has been shipping HAMR HDDs in volume since 2024. Competitors Western Digital and Toshiba plan to ship HAMR HDDs in 2026. The picture below is a cartoon showing the layout of the HAMR head components and the media.

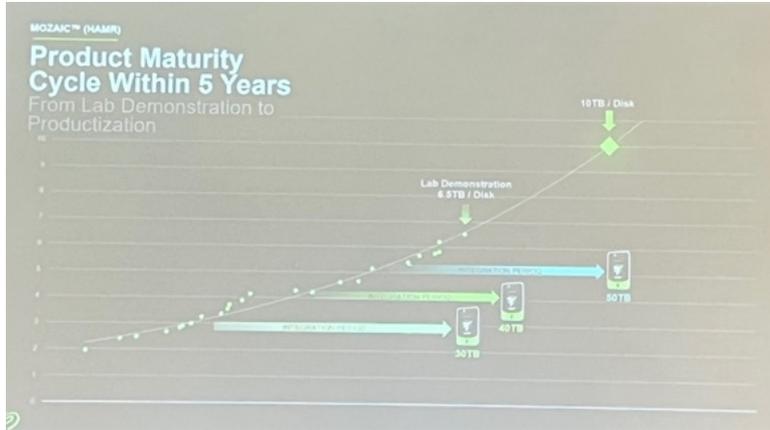


The picture below shows pictures of the current configuration of Seagate HAMR magnetic heads showing the laser structure sticking out above the slider. Seagate has been buying these lasers but they are moving to manufacturing their own, much more integrated lasers in their HAMR HDDs.

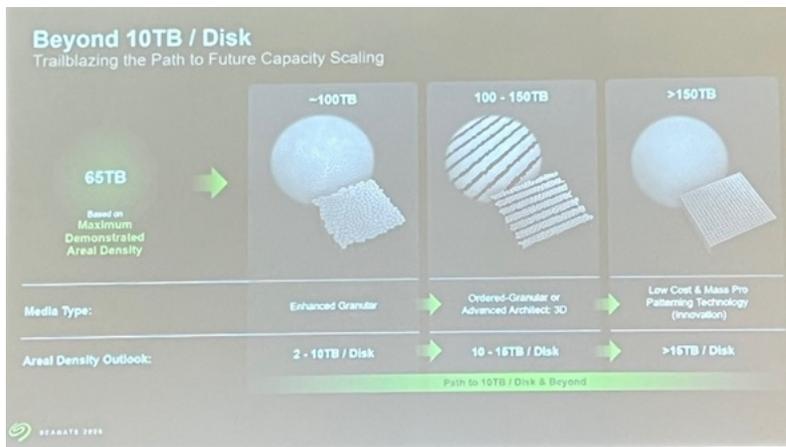


Seagate is using HAMR to leverage the growth in the storage capacity of their HDDs, particularly for nearline applications, but eventually in less populated HDDs. The image below shows laboratory demonstration and shipping product per disk capacities. Although the year is not shown [they announced a 6.5TB/disk laboratory demonstration at the IEEE MMM Conference in 2024](#).

30TB+ HDDs with HAMR are shipping today with 40TB+ planned in 2026 and 50TB+ perhaps in 2027. They also project a 10TB/disk laboratory demonstration, perhaps by 2026-2027. They project moving from laboratory demonstration to product capacities in 5 years. This would project 100TB+ HDDs by the early 2030's.



The image below shows projections for technology beyond 10TB/disk. This includes reaching 15TB/disk using including improved granular magnetic recording media or [multiple layer, 3D, magnetic recording](#). Going beyond 15TB/disk will likely require patterned magnetic recording media. All of these technologies are also using perpendicular magnetic recording and HAMR.



An interesting aspect about this move to HAMR in hard disk drives is that this could drive HDD companies to become manufacturers of the vast majority of the world's lasers. I did a search and found some suggestions that the total number of industrial and commercial lasers was somewhat less than 10 million units in 2024. Coughlin Associates estimates that about 71M nearline HDDs were shipped in 2024 out of 124M total. The total number of nearline HDDs is projected to grow in the future driven by the need to store large amounts of data at low cost to support AI workflows and other big data applications.

Nearline HDDs tend to be high capacity and trend towards the highest number of disks per drive, which is 10 per drive for Seagate, although WDC have HDDs with up to 11 disks per drive. To get the highest per HDD capacity, every disk will have two heads, one for each side of the disk.

If all of the nearline HDDs in 2024 were each to have 10 disks each, the actual number is somewhat lower, each of these drives would have 20 heads. The total number of heads for the 70M nearline drives would then require 20 X 70M, or 1.4B heads. If all of these heads were HAMR heads they would need 1.4 billion lasers. This would be 140 times more lasers made than in the current estimated industrial and commercial laser market!

Seagate and the other manufacturers of heads for HDDs could achieve manufacturing scale for making lasers for HAMR heads that would dwarf other applications. This scale could make Seagate a dominate player in the commercial laser market and could make it competitive as a laser source for other applications. These could include optical recording and various communication applications, including photonics communication in new chip packages.

Matt Ninesling, from Spectra Logic talked about magnetic tape media trends. Spectra Logic manufacturing large magnetic tape library systems. The image below shows information about the latest generation LTO 10 tape format with 30TB per tape native capacity.



This allows building higher capacity tape libraries with the image below showing Spectra logic tape library offerings. These libraries are serial attached SCSI, SAS, products primarily.

Spectra Tape Library Portfolio



Library	Max. Media/Slot Count ¹	Max. Native Capacity ²	Max. Compressed Capacity ²	Max. Drives ²	Max. Native Throughput (TB/hr)	Tape Drives Supported	Dual Robotics	Dimensions (H x W x D inches) Rack Units
Spectra Stack	960	36.8TB	42.0PB	42 Half Height 21 Full Height	33.8TB 35.2 TB	LTO	N/A	78.7 x 19 x 47 42U (Standard 35-inch rack)
Spectra Cube	5400	50.9PB	125.3PB	30 Half Height 15 Full Height	32.4TB 23.0 TB	LTO	N/A	79.4 x 55.3 x 45.4 42U (Standard Frame)
Spectra T950	10,200 (5,100 8U Frame)	300.6PB	751.5PB	100	177.8	LTO	N/A	79.12 x 29 x 39.4 (43U Individual Frame)
Trinity Plus	66,400 (3,320 40U Frame) 42,800 (1,970 40U Frame)	1.7PB 2.3PB	4.2PB 6.4PB	160 96	241.9 241.9	LTO, SMP T1000 and Oracle T10000	Yes	79.12 x 40 x 39 x 43.25 43U (Individual Frame)

1. Capacities and Throughputs are based on LTO 10 and 10M T1070 tape technology.
2. Drive sizes are full height unless noted.



The 2025 IEEE MSST celebrated 40 years of NPS and PFNS development. HDD move to HAMR could make them the largest manufacturer of commercial lasers as well as enabling high-capacity HDDs.

13) Maker Faires Could Help IEEE Create the Future, (Forbes Blog, 9/30/25)

I attended the 2025 Maker Faire in Vallejo, CA. I heard Steve Wozniak talk with Dale Dougherty, the founder of the Maker Faire’s, about what led Steve to create the Apple I computer, see picture below. I also saw many of the exhibits which were well attended by adults and children of all ages doing various projects ranging from fabrics to mechanics and electronics.



Hobbyists of all types were showing their projects, from amateur radio operators who fly weather balloons and may help supplement government program weather balloon flights that are being discontinued, to people making amazing props for Halloween and companies selling 3D printers and other fabrication equipment as well as kits for kids to make all types of things. Robots, especially, were well represented, the picture below shows some robots outside an area where battle-bots attacked each other.



I have only three more months as Past President of the IEEE before I am off the IEEE Board of Directors and move on to other things. Two things that I wanted to do in my roles of IEEE President Elect, President and Past President was to retain more of our younger members and let the world know about IEEE and its mission to advance technology for the benefit of humanity.

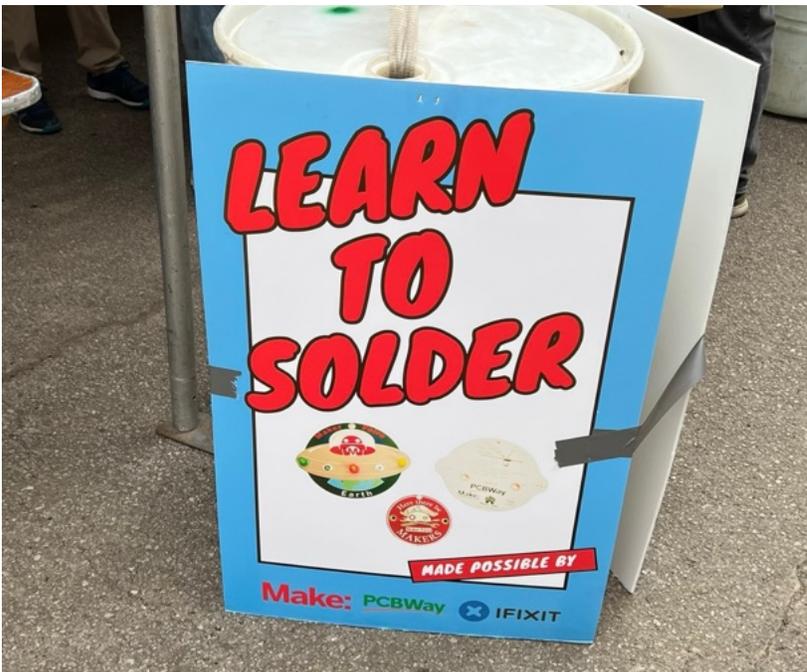
I think IEEE must be involved in STEM and STEAM activities of various sorts to increase awareness of who we are and to influence young people to see how learning technology enables them to do things and make a difference. Many professions are involved in various ways to allocate things that already exist. Those involved in technology can create things that never existed before, creating new wealth and opportunities.

Technology can be a way for people from disadvantaged backgrounds to make things better for themselves, their families and their communities, if we can provide cost effective ways for them to learn about the laws of nature and how they can use them to do useful things. I think that the IEEE can provide opportunities for kids and their families to understand technology that can lead to useful fulfilling careers.

Participating in Maker Faires, FIRST robotics and other such activities could feed this need and also show our value to our younger members, who themselves may become parents. Often STEM oriented activities include other STEM groups such as the FIRST Robotics exhibits with local teams at the Maker Faire, shown below.



The local IEEE San Francisco Bay area sections have been involved in various ways at previous local Maker Faires, for instance having a booth at the Faire and running a station for kids to learn to solder. This year iFixit and PCBWay sponsored the solder station, see below.



However, the local sections weren't able to organize and participate in the 2025 Maker Faire. I hope that next year, when I am not an IEEE officer, I might be able to help the three local IEEE sections do activities at the 2026 Maker Faire and thus help to attract the next generation of technologies who advance technology for the benefit of humanity.

Maker Faire is a convention of do-it-yourself enthusiasts established by Make magazine in 2006. Participants come from a wide variety of interests, such as robotics, 3D printing, computers, arts and crafts, and hacker culture. The 2025

San Francisco Bay area Maker Faire was held on Mare Island near Vallejo, California.

The 2025 Maker Faire in the San Francisco Bay area is the sort of event that IEEE should engage with to attract the next generation of technologists, the people who will create the future.

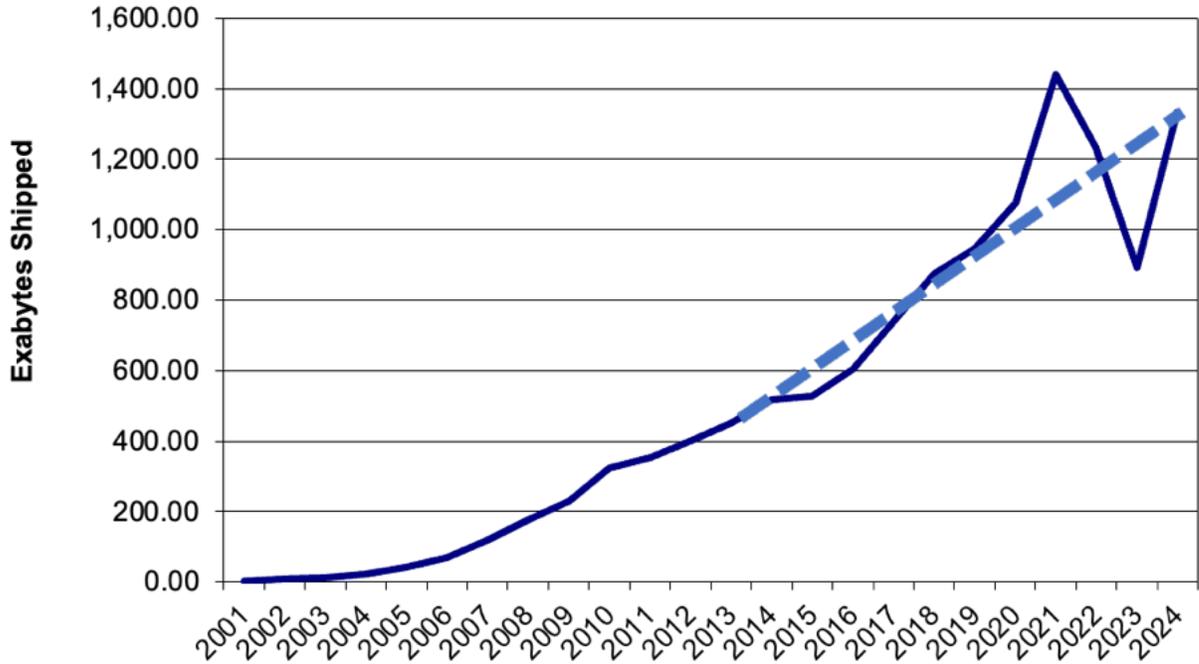
14) Storage and Memory Price Surges to Support AI Demand are Likely Temporary, (Forbes Blog, 10/10/25)

There have been [many reports of increasing prices of digital storage and memory products](#) such as hard disk drives, HDDs, solid state drives, SSDs and memory such as DRAM to meet the needs of AI workloads. However, it is likely that current demand and thus higher prices are driven by anticipated needs and the impression of shortages in production, which are likely short term.

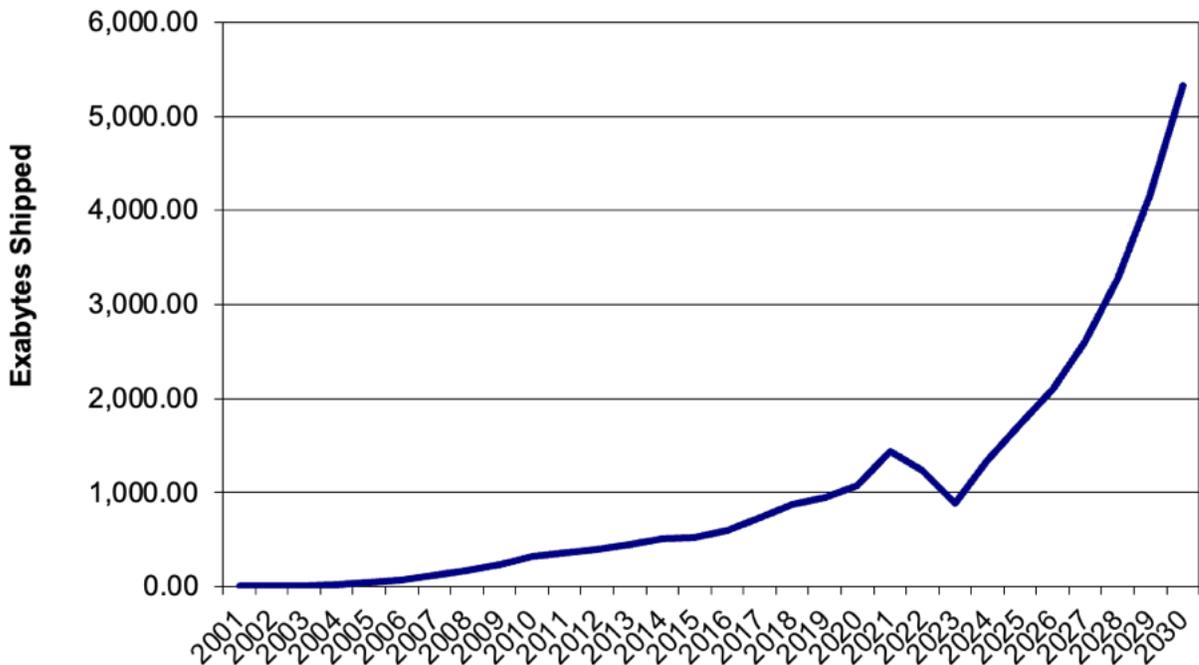
AI workloads require lots of data for training. Additional storage capacity is likely required for inference as well, particularly with approaches such as retrieval augmented generation, or RAG, that uses domain specific data to improve the results with an existing large language model, LLM. The growth in AI projects is driving up the demand for storage and memory as well as GPUs and that is causing the prices of all types of memory and storage to increase.

This is a big change from the COVID related crash of all storage and memory technologies in 2022 and 2023. These recovered in 2024 with the consumption of excess inventories at customers. As a result, many manufacturers, such as HDD manufacturers, said that they would be very cautious about increasing their manufacturing capacity.

For HDDs the historical exabyte shipments from 2001 through 2024 is shown below. We can see the spike in exabyte shipments in 2020 and 2021 with anticipated cloud storage demand during the COVID pandemic and the massive slump in 2022 and 2023 as excessive inventories were consumed. 2024 saw continued growth, which is approximately in line with the average capacity shipment growth rate before COVID, indicated by the dashed line.

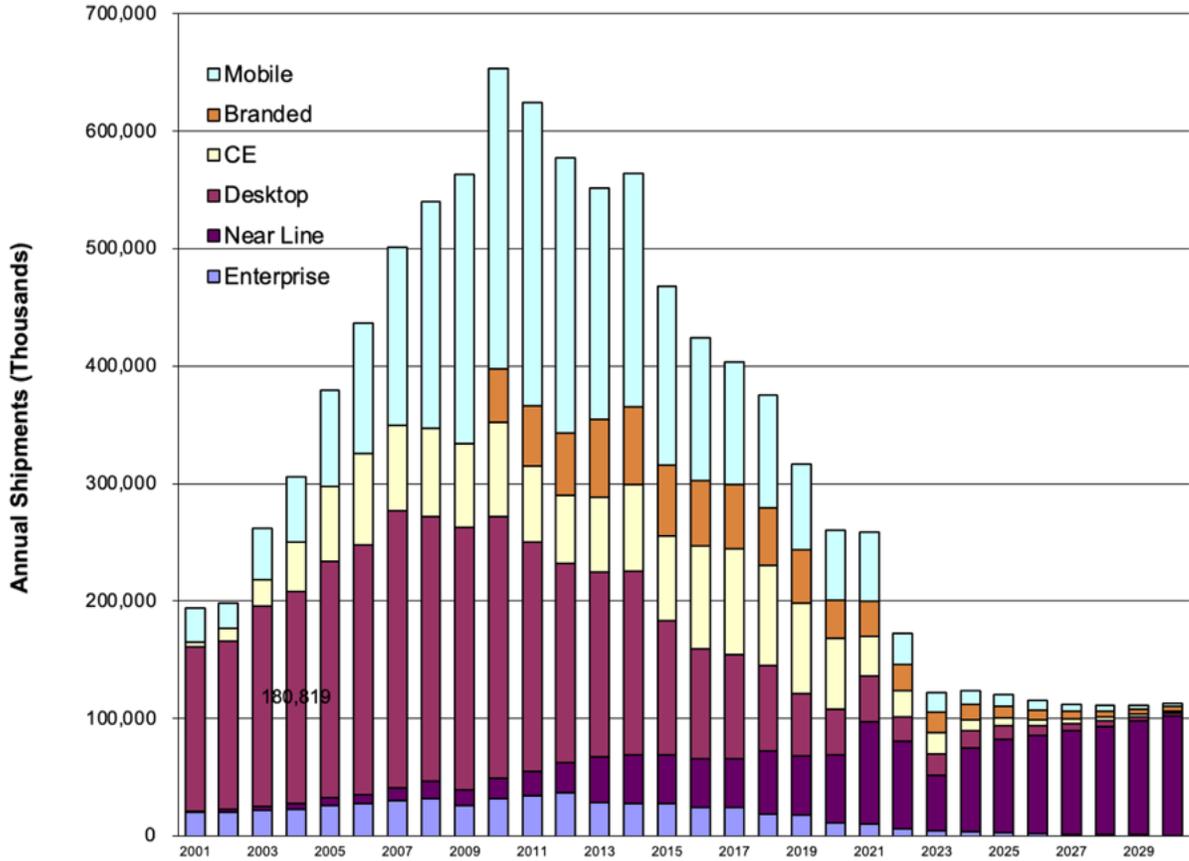


If we extrapolate the expected growth through 2030 for HDD shipments from [Coughlin Associates](#) (recently updated), we can see the expected growth in storage capacity shipments required to support AI and other data intensive workloads. The extrapolated storage shipments for HDDs in 2030 is about 4X higher than in 2024.



At the same time, due to the increase in the expected storage capacity per HDD through 2030 we expect that actual HDD shipments required to meet this projected demand will be about the same in 2030 as in 2024, at about 123M HDDs shipped

versus about 124M HDDs shipped in 2024. The historical HDD unit shipment chart since 2001 is shown below.



This is because the projected average storage capacity per shipped HDD is projected to grow from about 11TB in 2024 to about 47TB by 2030, about a 4.3X increase. This is in line with recent projections of 50+TB nearline HDDs in production by 2027 and thus in mass shipment by 2030 and with the continued shift to producing nearline high-capacity HDD production by the industry through 2030 as shown above.

The total HDD unit shipments through the second half of 2025 were about 59M units and we project that total shipments in 2025 will be about 120M units assuming a 3% higher shipment of high-capacity nearline HDDs in the second half of 2025, 61M units. This should be possible without bringing up more HDD production capacity, since 2024 HDD shipments were about 124M units. There is even the possibility that total HDD unit shipments could approach the 2024 shipments, 3% higher for 2025 than our current projection. We will review our projections after seeing C3Q 2025 shipments.

So, it seems that there is enough nearline HDD production capacity to meet expected long-term demand and thus it is likely that recent increases in nearline HDD prices are a combination of the longer-term contracts that the HDD

companies have established with big HDD consumers to stabilize HDD prices after the last recovery in demand in late 2023, such as hyperscale data centers in addition to customers buying in anticipation of future needs for AI applications and thus driving up the price of HDDs currently available. This can create distortions in the actual demand, driving companies to purchase or try to purchase more than their real short-term need.

In the memory space, including DRAM and NAND flash, there is a long-term trend of a boom-and-bust cycle on capacity pricing, driven largely by the manufacturers building out production capacity when prices and thus margins are good and then stopping that production build out when all that new production comes on line and prices drop again.

[SEMI recently announced](#) that 300mm wafer fab spending to increase by 7% in 2025 over 2024 (older data shows the retrenchment in this spending in 2022 and 2023). This announcement says that memory-related fab spending is second after the logic and micro wafer fab spending growth. In particular, the announcement says that, DRAM-related equipment investment is expected to exceed \$79 billion from 2026 to 2028, with 3D NAND investment reaching \$56 billion over the same period. AI training and inference have driven comprehensive demand increases across various types of memory.

With new memory fab spending starting in 2024 it is likely by sometime in 2026 or at least by 2027 we will experience a glut of memory production, even with the growing demand of AI and there will be significant contraction in memory pricing. In other words, the current trend of increased prices for HDDs and likewise for other storage and memory products, such as NAND flash and thus SSDs and DRAM, could be a short-term trend with a possible correction to more historical growth rates, likely by sometime in 2026.

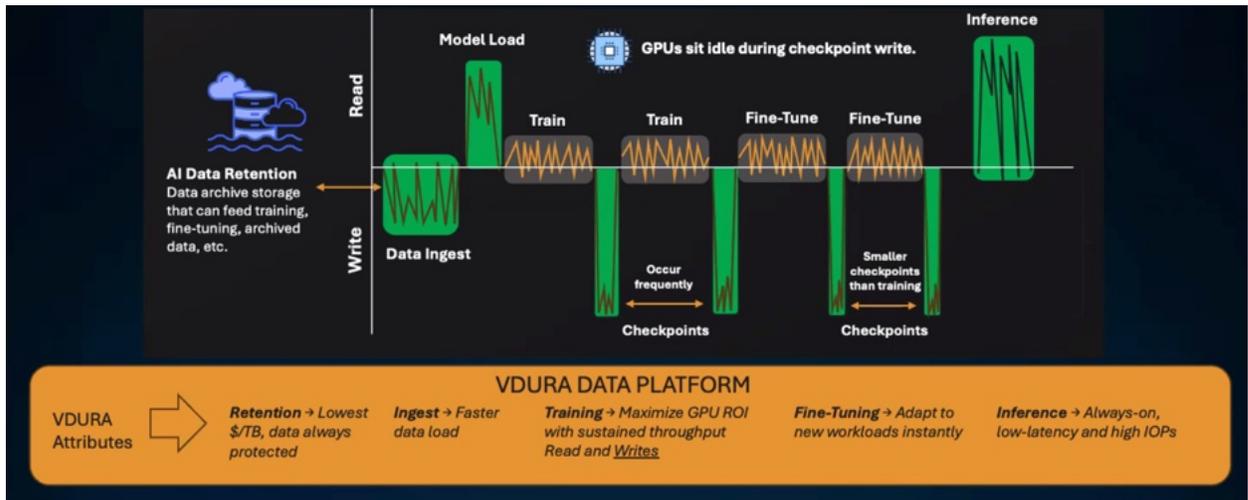
15) Vdura Controls Costs while Providing Lower Cost, At Scale AI Storage, (Forbes Blog, 10/8/25)

Vdura says that it provides a software-defined data infrastructure platform built for AI and high-performance computing, HPC. Their software combines the scalable, linear performance of a parallel file system with the resilience and cost efficiency of object storage, providing a unified global name space with one control plane and one data plane.

Panasas rebranded itself as Vdura on May 7, 2024 as it made a strategic switch from selling proprietary hardware to focusing on software-defined storage with a subscription-based business model for AI and high-performance computing.

Panasas was founded in 2000 by Garth Gibson, co-inventor of RAID and the creator of the first Linux-based parallel file system, PanFS. In 2004 Panasas started to produce high-performance storage systems. PanFS was the basis of the company's switch to focusing on software subscriptions. Garth Gibson returned to Vdura as Chief Technology and AI officer in September 2025.

To support AI workloads sustained throughput at scale is essential to avoid starving GPUs and wasting that investment. AI also requires expanding data metadata capacity and performance to quickly access the data. Vdura says that they accelerate in all parts of the AI pipeline rather than in certain parts like their major competitors. The figure below shows these various parts including data ingesting, model loading, training, check pointing, fine-tuning and inference.

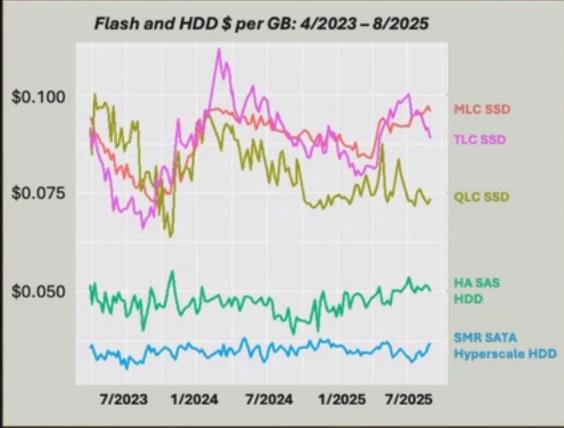


The company says that their Vdura Data Platform, VDP, can turn 6-thousands of storage servers into a high-performance, resilient and durable data platform. It is resistant to many failures of both devices and nodes, is easy to deploy and manage and provides seamless integration of NVMe for a fast access layer and optional hard disk drives, HDDs, for cost efficient storage capacity.

In particular, the company can use SATA rather than serial attached SCSI, SAS, HDDs, which are widely used by hyperscalers for object storage as part of their software defined package. The figure below shows recent costs of various SSDs using flash memory versus SAS and SATA HDDs.

Flash vs & HDDs.....Hybrid Optionality Makes Sense





Source: Computerweekly.com / diskprices.com

Hyperscalers are Hybrid Fleets

“SSD-only storage still poses a substantial cost premium over a blended storage fleet of SSD and HDD.”

“The challenge is putting the right data — the data that gets the most I/Os or needs the lowest latency — on SSD while keeping the bulk of the data on HDD.”

-Larry Greenfield & Seth Pollen, Google (March 2025)

Key Insights

- “Hyperscale Flash” is marketing hype — in reality, Hyperscalers use all NAND types, but only non-HA, single-port NVMe for cost efficiency.
- SAS HDDs are niche — low-volume, capacity-limited, and increasingly expensive.
- Hyperscale HDDs are SATA/SMR — Hyperscalers rely on SATA and SMR drives, which offer ~20% more capacity at no extra cost.

Vdura says that NVMe flash + HDD architectures reduces the long-term costs versus all-flash competitors, especially using commodity SSDs and HDDs. The Vdura-bility guarantee and Vdura Care Premier provide guaranteed SLAs for availability and durability plus 10-year cover with device replacement.

By enabling fast and reliable performance using a combination of SSDs and HDDs Vdura can offer a flexible and adaptable ratio of HDDs and flash in a system as shown below. The company told me that an 80%/20% HDD to SSD mix is best for most customers.

Multi-Dimensional Scaling – Adapt Ratios for AI's Evolving Demands



Flexible Ratios to Optimize for GB/s, IOPS, Capacity, Availability, and Durability

HDD 95% : 5% Flash

↓

Lowest Cost

HDD 80% : 20% Flash

↓

~ Spending 50/50 on Hybrid Nodes & Flash Nodes

HDD 0% : 100% Flash

↓

Highest Cost



NVMe Flash



Capacity expansion (HDD's)



Dynamic Adaptation: Ratios evolve seamlessly as new nodes (with or without HDD) are added—online expansion without rebalancing. Build for what's next in AI's unknown landscape.

To get the most use of GPUs the storage system that supports it must have high availability, often referred to as the percent uptime as the number of nines before and after the period, e.g. 99.999% uptime is 5-nines.

Also, the data used for AI must be able to last with no loss or corruption over time. Although storage is about 20% of the spending for AI it enables the use of the 50% server/GPU investment. Availability and durability reduce the incidence and impact of downtime.

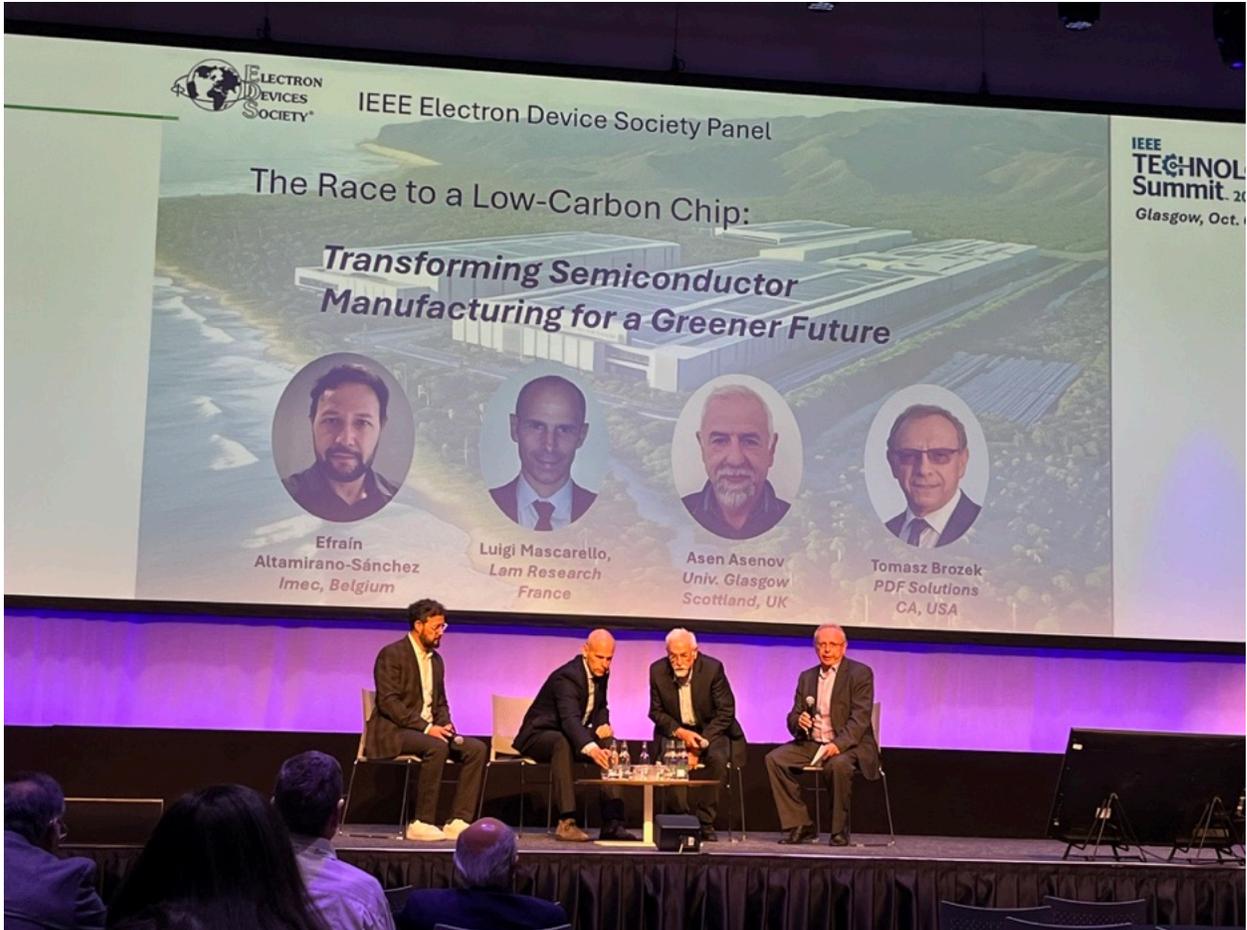
[In September Vdura announced](#) a validated reference architecture with AMD designed to eliminate data bottlenecks and simplify large-scale deployments for AI and HCP. This combines Vdura V5000 storage with AMD Instinct MI300 series accelerators.

Vdura says it offers lower capital and operating costs and with high storage availability and durability for demanding at-scale AI applications using combinations of SSDs and HDDs.

16) IEEE Innovation Summit Develops New Ways to Create IEEE Initiatives, (Forbes Blog, 10/12/25)

In early October I attended a couple of IEEE events in and around Glasgow Scotland. The first was an IEEE Technology Summit focused on CleanTech Solutions at the Glasgow Convention Center. The second was an Innovation Summit at Cameron House near Loch Lomond.

The CleanTech Solutions Summit was organized by IEEE Technical Activities. This is the IEEE organization unit that includes the various IEEE Societies and Technical Councils. The three day event covered many topics around creating more sustainable technology and using technology to reduce our carbon footprint. The image below was from a panel discussion on making semiconductors with less energy with international experts from several organizations involved in semiconductors and semiconductor manufacturing equipment.



The Innovation Summit was a two-day experiment to find new ways to attract and evaluate proposals for innovative ideas in the IEEE. There is an IEEE board committee, called the New Initiatives Committee, or NIC, which has a few million dollars annually set aside to fund strategic IEEE activities. Some of these require smaller amounts of funding and are generally completed within a year. Some of the more ambitious project require considerable funding and take up to three years to complete.

With the creation of the IEEE Strategic Plan last year from 2025-2030 and the work in the IEEE board Strategy and Alignment Committee with the various parts of the IEEE to determine how these goals are going to be achieved over the next five years, funding for these activities is an important enabler.

This year IEEE President Kathleen Kramer created an ad hoc, chaired by Fred Shindler, to explore ways in which IEEE can pursue and fund activities to further these strategic goals, particularly ambitious efforts that require lots of time, attention and money.

The image below, from the Summit, shows from left to right, Mary Ellen Randall, current IEEE President Elect and President in 2026, Fred Shindler, Innovation Summit organizer, and Kathleen Kramer, 2025 IEEE President.



To prepare for the Summit, requests were made to the IEEE membership for proposals for consideration at the Summit. I believe that there were over 80 proposals made and of those about 20 were considered during the Summit. These ranged from some very practical proposals on topics such as AI for Predictive Optimization and Efficient Maintenance in Mining Trucks and the Technological Development of a Low-Cost Device for the Classification of Dried Cocoa Bean Quality Using and Image Processing and Computer Vision to IEEE enhancement focused proposals particularly those for providing benefits to society.

Among the IEEE enhancement focused proposals were ones on enhancing the IEEE program for getting kids interested in engineering and providing teaching resources for teachers called Try Engineering using AI to tailor these resources to the needs of individual students, deal with different languages and to help teachers create lesson plans using these resources that align with country and state guidelines.

Another very interesting one was focused on using AI to advance scientific advancement by tracking unfulfilled searches on IEEE Xplore, the IEEE document library, to identify technology blind spots.

Small groups of IEEE volunteer leaders and staff in attendance were selected to evaluate and develop these various initiatives and the proposals often changed during the course of the discussion from the original proposal. The group then made a 5-minute pitch and answered questions about their proposal and the various proposals were then evaluated by those who heard the pitches. Below is an image of one of these groups making their pitch.



Perhaps some of these initiatives will be developed further and funded by a future initiatives funding program in the IEEE which will succeed the current New Initiatives Committee. There is much good work to be done by the IEEE to further our professions and provide value and opportunities to the people of the world. IEEE volunteer leaders and staff met in Scotland to explore CleanTech Solutions and to develop ideas for evaluating and developing ambitious new initiatives to achieve our 5-year strategic plan.

17) In the Near Future Data Center SSDs will Require Liquid Cooling, (Forbes Blog, 10/14/25)

Liquid cooling will likely be needed for SSDs that provide the performance needed for future AI data center workloads. Solidigm was demonstrating and talking about a prototype liquid cooled SSD at the 2025 OCP Summit.

The 2025 OCP Summit in Silicon Valley brought together customers and the supply chain for the hyperscale data center market. Total attendance was over 11,000, approaching 12,000, up significantly from about 7,500 at the 2024 event. Much of this growth was due to data center spending to gear up for AI applications and the importance of open reference designs and standards to enable industry growth.

AI workloads are increasing the demand for data center power as well as methods to cool the data centers. For several years various liquid cooling technologies have been on display at shows like the OCP Summit, but this year it looked like it would definitely happen in AI-workload data centers. Among the keynote talks at the show, Google talked about the Deschutes Coolant Distribution Unit, CDU, reference design for data centers, shown below.

Deschutes Coolant Distribution Unit

- Targeted to support ~2 MW heat loads,
- Hydraulic capacity targets of 500 GPM at 80-90 psi
- The specification will enable any CDU supplier in the industry to develop, manufacture and improve upon the design.
- The CDU is assembled from components that are sourced from multiple vendors
- Installation and maintenance procedures are shared also to enable fast deployments of reliable equipment.



In the OCP Summit exhibit area there were several companies showing their products based upon this reference design for a liquid cooling distribution system. The actual heat extraction is done by circulating the cooling fluid through cooling plates adjacent to the computing and other unit needing heat extraction.

The rising power requirements for GPUs and other computer logic is the big driver for liquid cooling, but other components are running hotter as well, driving them to need liquid cooling. One area that has not been a target for liquid cooling is digital storage, but solid state drive, SSD, manufacturer, Solidigm, gave a talk and exhibited an SSD and special cold plate design for a liquid cooled SSD.

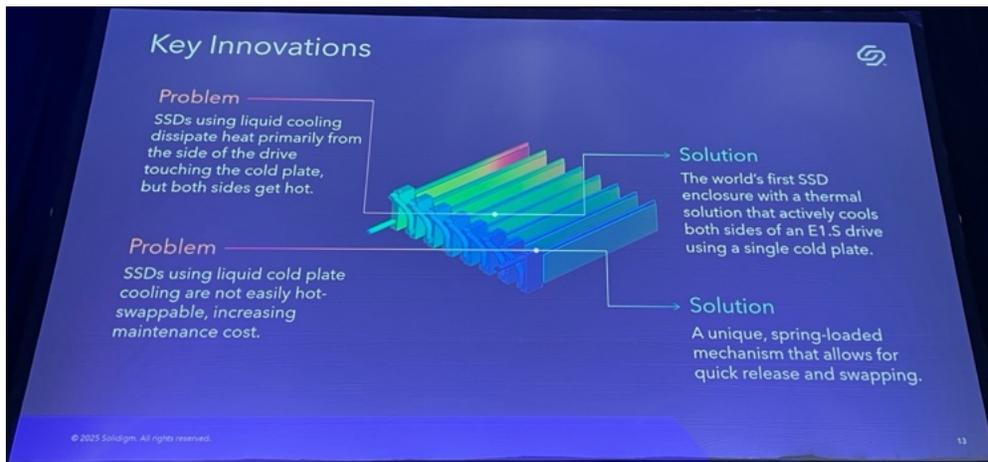
They pointed out that although PCIe generation 5.0 SSDs generally only need up to 25W of power that with the move to PCIe 6.0, estimated in 2026 and eventually PCIe 7.0, likely in 2028 SSD power requirements could rise to 40W and even 60W as shown below. These new PCIe generations are required to increase the storage data rates to meet evolving GPU requirements.

Architecture	2020	2022	2024	2026(e)	2028(e)
TFLOPS (FP16)	~600	~2,000	~5,000	More	Even more
Max TDP (W)	400	700	1,200	1,800?	2,400?
GPU Cooling	Air or liquid	Air or liquid	More liquid	All liquid?	All liquid?
SSD Interface	PCIe 4.0	PCIe 5.0	PCIe 5.0	PCIe 6.0?	PCIe 7.0?
SSD Power Budget (W)	23	25	25	40?	60?
SSD Cooling	Air	Air	Air	?	?

© 2025 Solidigm. All rights reserved.

These higher SSD power requirements will generate more heat in the SSD. Solidigm pointed out that below an internal temperature of about 77°C an SSD operates at its specified data rate. However, at about 77°C performance drops to 58% of spec performance and at 79°C it drops to 1% of the specified performance. This is called thermal throttling and the SSD will require cooling at these higher power levels, likely beyond what traditional air cooling can provide.

Solidigm addressed the cooling of SSDs by creating a special cooling plate design shown below. This design cools both sides of the SSDs and they are installed with a spring-loaded mechanism to allow release and hot swapping of SSDs when they fail. This design is imaged below.



This design was on display at the Solidigm OCP Summit booth with coolant running into it. A picture of this is shown below. Liquid cooling may be needed for all the elements in AI hyperscale data centers.



In order to facilitate SSD cooling the company suggested some new specifications for SSDs. In particular the surfaces of the SSD that connect with the cold plate need to have tolerances for flatness and roughness, the cold plate contact sides need to be defined and sharp edges should be beveled.

During another talk on SSD form factor developments, EDSFF New and Upcoming Updates there was talk about liquid cooling of SSDs. That session focused on discussions by the OCP storage group after the 2025 FMS in August on requirements for future SSDs. There was a focus the E1, E2 and E3 specifications and liquid cooling. There was talk about needing to be able to cool SSDs with up to 79.2W power, even higher than the Solidigm projection.

I asked if the Solidigm work reflected the plans for liquid cooling SSDs and it seems it is seen as a prototype, which is probably fair. I will write more about these other SSD discussions in another blog.

Solidigm demonstrated a prototype liquid cooled SSD that may be needed for data center SSDs based upon the PCIe 6.0 and later generation SSDs.

18) OCP E2, One Form Factor to Rule them All, (Forbes Blog, 10/23/25)

At the 2025 Open Compute Platform, OCP Summit Meta spoke about their experience with quad level cell, QLC, NAND flash in their data centers. A panel talked about developments in the new E2 SSD form factor and another panel discussed flexible data placement.

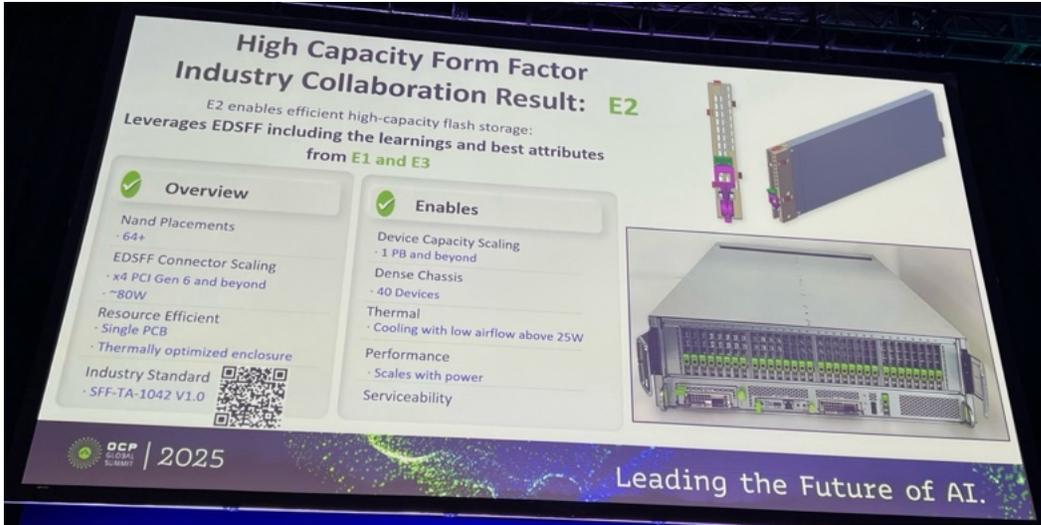
Kioxia was showing their latest enterprise and data center SSDs. Samsung spoke about their HBM and other memory advances and SK hynix talked about their high bandwidth flash and other memory products for AI.

Meta gave information about their first QLC Flash Storage Server, the Open Rack Frame version 3, ORv3, whose specifications are shown below. The actual shelf reference design, supporting 24 U.2 and U.2 long Direct Flash Modules was created by their design partner, Jabil. As shown the goal was a storage server that would enable more than 50PB per rack.

Component	Decision	Advantage
Rack	ORv3	48V DC busbar; aligns to high volume products within Meta
Compute	1. 25 Sapphire Rapids 2. 16Ch DDR5 Memory config 3. PCIe Gen5 x4 per drive capability	32MBps/TB Application Performance
Network	2x NIC slots Can accept 200Gbps and 400Gbps per slot	Perf scaling to 2x Drive density (up to 128TB/150TB)
Storage	20U chassis with 24 front-removable drive slots Accepts DFM/DirectFlashModule (U.2 'long') and U.2 Drives	Headroom to support Userspace FTL implementations also Enables >50PB-usable per Rack Hot Serviceable/Swappable Supply Diversity

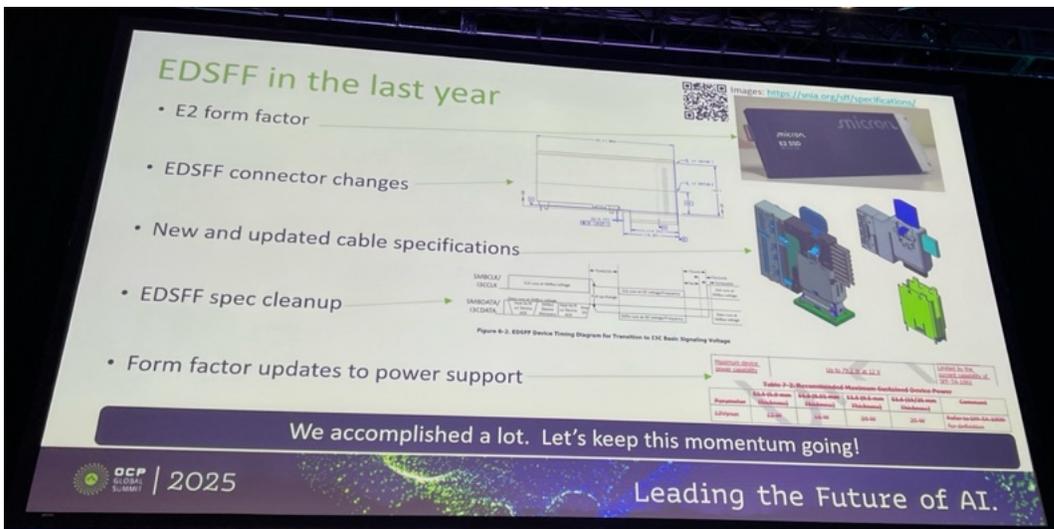
A panel with Peter Choi from Pure Storage, Ross Stenfort from Meta, Arthur Lai from Molex and Anthony Constantine from Micron spoke about updates on SSD form factors. Their talk shared experiences designing and shipping QLC drives and contributions to a new industry standard form factor, E2. This standard is designed to support 64 NAND packages with a large flash controller, high DRAM memory on a single circuit board operating at 30+W with low airflow.

Peter from Pure Storage compared the U.2 long Direct Flash Modules they make with this new E.2 form factor. Ross from Meta positioned such QLC flash modules as a layer in the memory and storage hierarchy below TLC SSDs and above HDDs. The goal from 2024 was that this E2 form factor would be the industry standard for 128TB or higher SSDs. The slide below shows how the consortium wanted to leverage E3 and E1 form factors to create this new standard.



Arthur Lai from Molex went into details about the optimal design of storage modules for optimal air cooling, which favored what he called an ortho hybrid solution for high density rack installation. Challenges remain to push the performance to that required for Gen 7 with over 80W support.

Anthony Constatine from Micron discussed E2 improvements made since 2024 and future improvements and [liquid cooling](#). He pointed out the need to have one form factor to support high density data center applications for PCIe 7.



Another session covered flexible data placement, FDP, and how it might best be used to avoid write application, WAF, that can cause early wear out of QLC and TLC SSDs. Rory Bolt from Kioxia spoke about FDP use cases. He made recommendations that separating temporary files and swap spaces from data with a longer lifespan we could reduce the write amplification. Kioxia was exhibiting their LC9 Series Enterprise QLC SSDs as well as their SD8 E1.S Data Center

SSDs and their CD9P Series Data Center SSDs. The LC9 products from the Kioxia booth are shown below.



Taeksang Song from Samsung spoke about advances in HBM4 and HBM4E as shown below, showing over 2X bandwidth increases over HBM3E and significantly greater energy efficiency. The company also said that the future of HBM will involve customization for various types of AI demands.



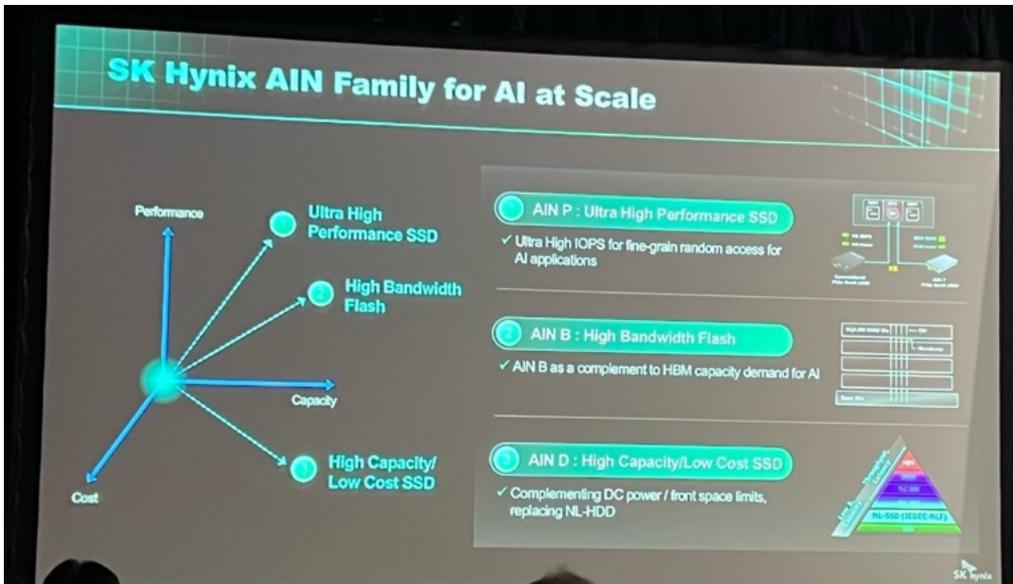
He also spoke about advances in DDR5 memory, RDIMM and MRDIMM as well as LPDDR6, LPCAMM2 and SOCAMM2. He also talked about a process in memory, PIM, device with LPDDR for AI, LPDDR5X-PIM. The figure below shows the expected performance and energy improvement this PIM can provide.



He discussed Samsung’s CXL solutions, shown below, which will enable greater memory utilization through memory pooling and compute near memory.



C. S. Kim from SK hynix spoke about the company’s AIN family for AI applications shown below.



This includes high bandwidth flash, which [SK hynix announced it was working in partnership with SanDisk on at the 2025 FMS](#). The slide below indicates that unlike conventional SSDs, this highly parallel memory configuration lets NAND-based flash devices operate at near HBM performance with memory capacities like those of SSDs.



The 2025 OCP Summit featured Meta and SSD companies developments on the E2 form factor, memory and SSDs for AI from Samsung and Kioxia as well as high bandwidth flash developments from SK hynix.

19) WDC Opens Data Center Customer Lab and Toshiba Announces 12-Disk HDDs by 2027, (Forbes Blog, 10/27/25)

Hard disk drives, HDDs, store over 80% of the world's data and will continue to be a cost-effective resource for data retention into the next decade. HDDs are often used as secondary less expensive storage behind SSDs as primary storage for AI and other data intensive workloads. This week we saw a couple of interesting announcements on HDDs from Western Digital and Toshiba, who, along with Seagate Technology, manufacture the world's HDDs.

[Western Digital announced](#) its expanded System Integration and Test (SIT) Lab, a state-of-the-art 25,600 square foot facility in Rochester, MN. The Rochester facility includes expanded areas for testing, labs and dedicated customer spaces. The facility is geared for faster, more predictable qualification of Western Digital's high-capacity HDDs for data centers and enterprise customers.

According to WDC, in addition to providing research, development and global operations, the facility's mini data center environment enables real-world testing and validation, ensuring customers receive advanced storage solutions precisely when they need them. The SIT Lab serves as a massive, dedicated collaboration hub where the company's engineers work alongside key customers throughout every stage of the product lifecycle, including development, qualification, production ramp and end-of-life.

[Toshiba announced a prototype of a 12-disk high HDD](#) in the standard 3.5-inch hard drive form factor. Toshiba expects that this new HDD configuration will enable 40+TB HDDs for the data center market using the company's Microwave Assisted Magnetic Recording, MAMR, a type of energy assisted magnetic recording, by 2027. Toshiba makes 18-20% of the worlds HDDs, with Seagate Technology and Western Digital roughly splitting the rest of the HDD market.

The company said that key advances include the development of new dedicated parts in the stack, and replacement of the current aluminum substrate medium with a glass substrate that offers greater durability and allows thinner designs. These advances deliver improved mechanical stability and in-plane accuracy, higher density and greater reliability. Toshiba is also investigating the use of 12-disk stacking technology with next-generation Heat Assisted Magnetic Recording (HAMR).

Toshiba does not have its own hard disk drive heads of disk manufacturing, unlike Seagate and Western Digital. It buys from some external Japanese suppliers. The thin glass substrates used for this demonstration and enabling a 12-high disk in the 3.5-inch HDD form factor, likely came from Hoya. Toshiba's disks are made by Resonac and its magnetic recording heads are made by TDK.

At the [2022 IEEE TMRC conference](#) Hoya was showing prototype 14 and even 24 disk configurations for HDDs using their glass substrates. Hoya is the largest

manufacturer of glass substrates for hard disk drives. Glass substrates are required for HAMR recording media because of the higher temperatures during media deposition.

Western Digital introduced an 11-disk HDD in 2024 with 32TB using its ePMR and shingled magnetic recording, SMR. Seagate has been shipping 32TB HDDs with 10 disks using heat assisted magnetic recording, HAMR, in late 2024 as well. Western Digital announced the opening of a test lab for qualification of HDDs for data centers. Toshiba introduces a 12-disk HDD platform for 40+TB HDDs by 2027.

20) OpenDrives Shows Cloud-Native Astraeus On-Premises Data Solutions, (Forbes Blog, 10/29/25)



Professional media and entertainment workflows often require high bandwidth and low latency storage to support rendering, transcoding and many other applications. Many companies provide various types of storage for this industry and cloud storage has been widely used. However, cloud applications and storage can be expensive and many M&E facilities use on-premises as well as cloud storage. Bridging the gap between the cloud and on-premise storage can be tricky. OpenDrives says that they have a solution to bridge this gap.

OpenDrives, a Los Angeles company, has been mostly known for building software-defined data storage solutions using SSDs and HDDs that are used in the media and entertainment and other industries. It recently announced the launch of Astraeus, a cloud-native data service platform and held an event about this new offering for customers and the press on October 15, 2025 at their Culver City headquarters.

Astraeus is a cloud-native data services platform. The company says that it enables IT departments within enterprise organizations to configure, store, orchestrate, manage, secure and deploy mission-critical applications dynamically on-premises, cloud or in hybrid environments. Astraeus should drive operational efficiencies and cost savings over cloud-dependent workflows and data environments because it helps control data fragmentation and sprawl.

During the event there were several horror stories about the dangers in cloud-first or cloud-only video production environments where e.g. the cloud applications were not stopped after active use and large expenses were incurred. Having data only in a cloud can also create a data silo, making it hard and expensive to use the data elsewhere. There were several stories about repatriating data on-premises from the cloud.

Open Drives said that their approach is different from current cloud and on-premises or even multi-cloud data environments. Sean Lee, the OpenDrives CEO says that out-come based data services are the right approach. Astraeus is built on cloud-native principles with a Kubernetes-based architecture and cluster-first design that creates a cloud-like experience. IT departments can deploy their cloud workflows locally while managing disparate data storage from a unified namespace that can scale out when the need arises, including bursting to the cloud. The figure below gives an idea of how Astraeus works.

Let Your Data Go & Let Your Work Just Flow

Introducing Astraeus, Your New Data Home

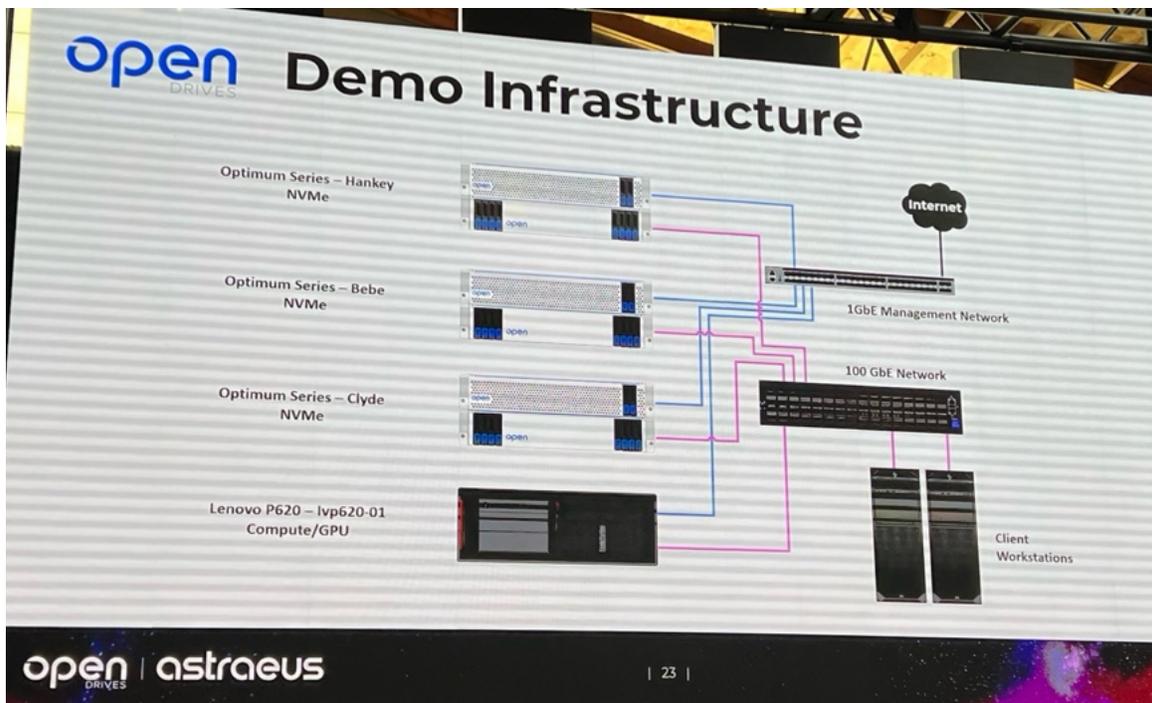


Adding Astraeus to OpenDrives data storage and management platform allows customers to choose the storage that meets their workload needs. The Atlas data storage and management platform provides network attached storage, NAS, that provide high-bandwidth and low latency performance that meets media workflow

needs. The Atlas Cloud Plus services extends this performance into the public cloud.

A slide from the event presented the cost advantages for the Astraesus with a GPU approach to using AWS. The Astraesus approach gives a predictable GPU node cost of about \$1,312 per month or about \$0.030 per minute. For that cost you get a full month of GPU use while you would only get about 6 hours of AWS HEVC transcoding.

A demonstration using the configuration shown below showed how Astraesus was able to work with a single GPU on an older Lenovo P620 computer with three Optimum NVMe storage nodes in the network. The demonstration showed good performance characteristics for a media workload.



OpenDrives showed how their Astraesus cloud-native data service platform enables Kubernetes-based high-performance storage for on-premises and hybrid cloud workloads.

21) Did Western Digital and Seagate Ship Fewer HDDs Last Quarter, (Forbes Blog, 11/5/25)

During their last quarterly investors webinar both [Seagate Technology](#) and [Western Digital](#) focused up their exabytes shipped hard disk drive storage capacity and their revenue and profits. However, they did not provide information that they have done in past quarters that would give investors important operational information such as the number of HDD shipped and particularly the number of high-capacity nearline HDDs shipped.

They also did not give information from which the average HDD sales price, ASP could be determined. I think that these are important elements to include in HDD quarterly reporting and that investors should be asking for this information to understand important dynamics in the market and to make prudent investing decisions.

I have been covering the hard disk drive industry for over 20 years and have over 40 years of experience in this industry as an engineer, engineering manager and more recently as a consultant. I have also been writing about the industry in these blogs for many years.

Late last week Seagate Technology and Western Digital reported their CQ3 '25 quarterly results. While the reports from both companies were favorable, I noticed that they are making it harder than in past quarters to figure out how many hard disk drives were shipped overall, and in particular the number high-capacity HDDs shipped for use in data centers. In addition, both companies made it harder to calculate the average sales price of their shipped HDDs.

I think this is a problem and I'll tell you why. While shipping to meet exabyte requirements, these requirements are met by the number of HDDs delivered times the average capacity of those HDDs. But, if we only know the total capacity shipped and not the unit shipments, we may not know if there were production or yield issues limiting HDD unit production or if the company is limiting possible production to decrease the supply and increase or stabilize unit prices.

Before the recent quarterly reports, I note that many investors were observing that demand for hard disk drives was exceeding supply and they were trying to figure out why this was so. I believe that this is mostly due to increased storage demand to support data intensive workloads, like AI. But what if the HDD companies cut their unit production to drive up prices and thus revenues? Or what if there were yield or other production issues that resulted in shortages, wouldn't investors want to know that?

Both company's reported shipping more exabytes of digital storage capacity in the last quarter. I estimate that the total shipped exabytes of capacity for the industry, including WDC, Seagate and Toshiba, increased by 4.9% over the prior quarter (from about 395EB to about 414EB). This is up a from the 4% gain in shipped EB capacity in the prior quarter. This is good news of course, but both companies accomplish these EB shipments by the cumulative capacity of the hard disk drives that they ship.

The exabytes shipped can be represented by the number of HDDs shipped multiplied by the average capacity of those shipped HDDs. With Seagate's quarterly results I was able to make an estimate of the total shipped drives from their report of 182EB of total HDD shipped capacity and an average capacity of

those shipped drives of 14.6TB. Divide the 182EB by 14.6TB and you get 12.47M HDD units shipped that quarter. That is about what Seagate shipped the prior quarter, 12.5M units.

Seagate gave no guidance on the number of nearline HDD shipped although they did say that of the 182EB of total shipped HDD capacity, 159EB of that was for nearline applications, with most of these nearline drives used in data centers. But they didn't say how many actual nearline drives shipped, or give any information, like in prior quarters, that could help determine this number, such as the average capacity of shipped high-capacity HDDs, such as in the past.

In the prior quarter, CQ2 '25 Seagate said that the nearline capacity shipments were about 137EB out of 150.9EB of mass capacity shipments, or 86% of their total shipped capacity. The average mass-capacity per drive storage capacity was 16.5TB. If we assume that this was also the average capacity of the nearline HDDs, which made up close to 90% of the total of their shipped mass capacity HDD units, we get an estimate of 8.3M nearline drives shipped in that quarter.

I estimate that that total nearline drives unit shipments in CQ2 '25 was about 65% of the total shipped HDD units. In that quarter Seagate said they shipped 162.5EB of total storage capacity with an average HDD capacity of 13TB. This gave a total of 12.5M HDDs shipped that quarter and the estimated nearline shipments from Seagate were about 66% of the total shipped units.

The high capacity nearline HDDs are the major growth market for HDDs and understanding now many of these were sold can tell us important information concerning the companies that, I think, should be of interest to investors. During the most recent Seagate earnings report the company said, referring to the growth of HAMR HDD shipments, that it had shipped "nearly 80% of nearline volume at or above 24 terabytes per drive in September 2025."

I tried to figure out from this information how many nearline HDDs Seagate shipped this last quarter. If the average shipping capacity in CQ3 '25 was the same as the prior quarter, 16.5TB then they would have shipped $159\text{EB}/16.5\text{TB} = 9.6\text{M}$ units, 1.3M units more than the prior quarter. But the average capacity of the shipped nearline HDDs was likely more. Since they said they shipped nearly 80% of nearline drives in September that were at least 24 TB, perhaps the average HDD capacity shipped over the entire quarter jumped to 24TB for the entire quarter. That would mean they shipped $159\text{EB}/24\text{TB} = 6.6\text{M}$ HDD, down 1.7M units from the prior quarter.

With the information we are given we don't know if they shipped more than the prior quarter or considerably less. Wouldn't investors like to know this information, especially if the company shipped less nearline HDDs than the prior quarter?

Western Digital provided even less information about their total and nearline HDD shipments. They only said that they had shipped 2.2M units of their latest generation ePMR HDDs. In the prior quarter, CQ2 '25 WDC said that they had shipped 12.9M HDDs and 9.3 Cloud HDDs, the cloud HDDs are the nearline HDDs that they shipped. So, did they ship more, less or the same total HDD units in the most recent quarter? We just don't know. What if they shipped less? Wouldn't investors like to know that?

In CQ2 '25 WDC reported that their total revenue, presumably from HDDs, rather than systems, was \$2,605M and if we divide that by the total number of HDD units shipped in the quarter, 12.9M we get an average sales price, ASP, of \$201.9 per HDD unit.

In CQ2 '25 Seagate reported separated their HDD revenue of \$2,281M from their systems, SSD and other revenue of \$163M. They also reported that the total capacity shipped was 162.5EB with the average capacity per drive of 13TB, if we divide the total capacity shipped by the average drive capacity, we get 12.5M units. Dividing the HDD only revenue by the number of HDDs shipped we get an ASP of \$182.48.

In the CQ3 '25 investor webinar Seagate reported that their total revenue was \$2,629M. They broke out the data center, nearline and perhaps systems revenue, of \$2,114M and Edge IoT revenue of \$515M. The data center revenue was 80% of the total revenue. However, they didn't report the HDD only revenue.

So, to make an estimate of the HDD ASP we have to have an estimate of the HDD only revenue, which was not broken out. If we estimate that this was \$165M, a bit more than in the prior quarter and thus reflected some higher system sales then the HDD only revenue would be \$2,464M. Dividing that number by the 12.49M HDD units shipped we get an estimate of the Seagate HDD ASP of \$197.3, 8% higher than in the prior quarter.

For many prior quarters WDC reported ASPs that were higher than those of Seagate, so that was probably the case this last quarter, but again, we don't know, since the company did not give information that allows us to calculate the ASP. The chart from [Coughlin Associates](#), below, shows the ASP of HDDs since the late 1990's until the last quarter, with that quarter estimated from the Seagate information.



ASP growth since 2015 has been largely due to the decline in non-nearline HDD unit shipments and the growth in nearline HDD unit shipments. As nearline HDDs have become a greater percentage of the total HDD shipments over the last few years, and particularly with the growing demand for nearline HDDs the ASPs have grown significantly quarter by quarter.

While not a direct indication of the average sales price of nearline HDDs, the total HDD ASP is an indication of the growth in nearline unit shipments and perhaps also nearline HDD average sales prices. Nearline HDD prices increasingly dominate these prices, with these drives comprising at least 65% in CQ2 '25 and likely an even higher percentage in CQ3 '25 and an even greater percentage in the future. A drop in ASPs from a prior quarter would certainly indicate that something had happened to demand and would likely also have a negative impact on revenue.

But without ASP information, we would not have this guidance on HDD average prices, which I think investors would like to know.

WDC and Seagate changed their quarterly investor reports this last quarter, limiting investor information on total HDD and particularly nearline units shipped as well as information on ASPs. This obscures some important information that investor might like to know.

GUEST ARTICLES

1) A Quick Wrap of the Global Content Conference, Andy Marken (Content Insider IBC Special), andy@markencom.com

We Let's see... it's a "little" bit breezy, a "little" bit cool and a "little" bit drizzly so it's gotta be IBC in Amsterdam in September.

Nothing unusual for the Dutch stronghold other than it was hosting about 45K of the most senior creative content management/delivery executives from around the world for four days.

Show folks said attendance was down slightly from pre-pandemic days but you couldn't tell it by the meeting/floor action or trying to catch dinner at any of the great Dutch bruin cafes.

It seemed like everyone was there to absorb as much as possible on innovations, new applications and how the industry and organizations are going to keep up with all of the advances, market opportunities and current world geopolitical environment.

Hey ... it ain't our fault!



Source - IBC

Every year, we're amazed at how rich and comprehensive the convention's program is that Mike Crimp and his team put together and it makes it difficult to ensure you see what fantastic new technology is being introduced.

We didn't even try to take it all in this year because that would require about five AI clones of yourself sitting in sessions and prowling the halls/stands.

It goes without saying that AI was everywhere and some meaningful and valuable specific applications are emerging, but there were times you wondered what, why and WTF.

We focused on those things that were addressing real needs, have staying power and direct value as well as a growing number of collaborative efforts with organizations joining forces to deliver valuable solutions to major

challenges/issues/opportunities that really would/could shape the future of the industry.

Several sessions addressed the growing complexity and uncertainty of today's content creation and distribution.

Allan McLennan, CEO of Padem Media Group, noted, "As we've heard in a number of sessions and on the floor," we're now able to work globally with multiple platforms –in LIVE IP Sports, new perspectives on content and software licensing, along with personalized programming in linear and streaming. This has improved the creation, development and distribution of enhanced and intelligent supply chains that had been making things overly complex and expensive."

He continued, saying, "The industry now has a better path to focus on ROI, which will streamline the activities and enable services and suppliers to bring costs under control and provide higher-quality deployments to meet the overall demand everywhere, which is vital in today's environment."

Some of the solution and product announcements in areas such as cloud migration, AI integration, personalized user experiences, and content delivery that we saw at IBC will help major and minor players deliver better content to a more diverse audience ... economically and reliably.



Source - IBC

Viewing Tomorrow – Members of IBC's IET Perfect Storm session included (1-r) Rob Koenen, founder/CBO Tiledmedia; Krish Kumar, CEO Wowza, Karen Clark, CEO Telstra Broadcast Services and moderator Allan McLennan, chief executive of PADEM Media Group.

For example, during McLennan's IET session – A perfect storm-from Peak TV to peak transformation – speakers explored the changes in both traditional broadcast and streaming.

The panelists discussed how new business models were placing increasing demands on production and content delivery and how organizations have to carefully select their AI implementation/integration opportunities and effectively test and manage the solutions before using them for programming and advertising. All of the C-level panelists agreed that the keys to the implementation were ensuring that copyrights were protected, the content was authentic and accurate, and guarding against misrepresentation and hallucinations.

“AI can be an exceptional tool in shaping content that engages the audience,” one of the speakers noted, “but care has to be taken that it is done in a responsive/responsible manner so it has to be implemented carefully and cautiously after extensive testing and comprehensive validation.”

There were a number of key sessions that explored the industry’s shifting business models and the transformational technology that is going to make it all possible.



Source - IBC

Rethinking – Evan Shapiro, ESHAP Media Cartographer, encouraged traditional linear content services to rethink the landscape and consider YouTube as an opportunity rather than as a competition.

Even Evan Shapiro, Media Cartographer at ESHAP, probably surprised and frightened many in the audience during his session.

“Traditional media companies aren’t on the verge of collapse,” he emphasized, “but all of the growth is going to the big tech platforms, and it represents trillions of dollars.”

“Five years ago, when everyone shifted to direct to consumer, we ceded 100 percent of control to the user and content owners/developers are never going to get it back.”

He told his audience that Netflix knows who its competition is and it’s not Amazon, Disney, WBD or Paramount. It has been--and is--YouTube.

“Traditional wisdom says sharing content across other platforms could cannibalize your audience,” he commented, “but the opposite is true because YouTube is only a distribution tool, and it makes it easier for viewers to find your content.

“We’ve entered the ‘affinity economy,’ where content companies have to go directly to the audience wherever they are,” he said, “Right now, that’s social media and YouTube dominates the arena and people – especially young people – are watching it on the big screen in their living room.”

He probably could have added that for indie project creators/producers and second/third tier content distributors, YouTube, due to its reach, is also an economic method of helping to find an audience and monetizing their content. But speaking of extending relationships and alliances, Netflix and Disney both announced agreements during IBC.

Netflix has quietly taken the lead in the collaborative approach with bundling agreements with MBC Group (MENA), TF1 in France, Canal+ in Africa, Binge in Australia and SK Broadband in South Korea.

At the same time, Disney+ entered a series of strategic alliances in Europe including Atresmedia in Spain, ITVX in the UK and ZDF in Germany.



Source - IBC

Think Different – YouTube’s Justine Ryst (center) and Channel 4’s Grace Boswood (r) explained to moderator Sasha Quadri and the IBC

audience how Channel 4 was leveraging YouTube's extended audience to quickly and economically improve and extend their audience and improve their bottom-line.

Interestingly, during the Better Together session, YouTube's managing director Justine Ryst, emphasized, "We don't commission. We don't own IP. We don't create it. We have 2.5B global users who watch content and 700M of them connect to us on their TV; but to say we're competitive to and cannibalizing TV is erroneous. We're actually extending their reach and the life of their content."

On another note, she emphasized that YouTube provides a platform for Creators of all types (professionals and up-and-comers) to profitably develop an audience/following of folks who are watching what has evolved to be a new definition of TV.

We're not sure, but she might be right that it may be a way for legacy services to reach a new, profitable audience, and it might – just might – be a profitable way for professional content creators to rapidly reach a global audience.

Of course, it would also help if her boss, Neal Mohan, YouTube's CEO, would also quit pitching the idea that their Creators should be considered to win a few industry award statues, even if a few do have millions of followers and make lots of money. It doesn't necessarily make them professional...just profitable.

But ... who knows!

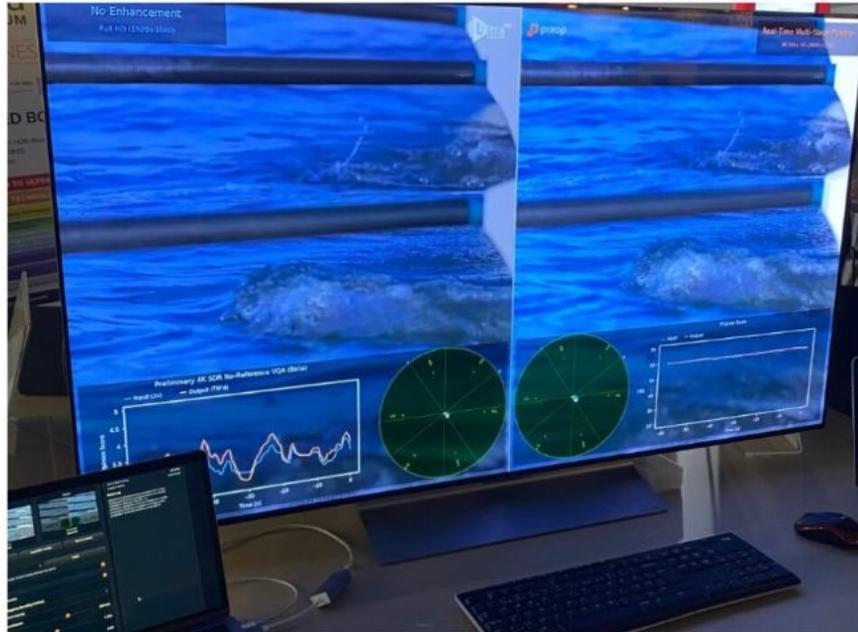
The sessions we attended were all great ... full of facts/good information about how every corner of the industry can profitably meet viewer's expectations. As interesting/information packed as they were, we also had to make our way through the friendly crowd to see what was new and hot on the show floor. Okay, we'll only say it once (again) ... AI was everywhere.

Because this is a global event and industry, one thing we noticed was there were a lot of well thought out and amazing localization and engagement tools too that were shown; and they are going to open up new opportunities for content creators to have their films/shows seen everywhere.

Camb.AI, Gaudio Lab, CaptionHub and Microsoft were just some of the products we saw on the show floor that theoretically will take the complexity and time out of localizing dubbing and subbing news, sports, films and shows to be enjoyed everywhere.

We're perfectly happy to follow subtitles and enjoy the action in the native language, but with some of the AI-based tools, it is mindboggling to see the project translated and lip-synced in what feels like real-time.

Either way, it breaks down the barriers to sharing the best content that is out there for indie and content owners.



Source – Allan McLennan

Realtime Upscaling – Pixop’s Pixop Live is a new solution that meets the needs of live streaming events and content libraries by converting HD SDR content to UHD 4K in almost real-time with less than 600 ms latency.

Latency has always been an issue for streamers (and we suppose everyone), especially when it comes to sports and live events (concerts, etc.), so we had to swing by checking out Pixop’s announcement in the Ultra HD Forums stand and neighboring Accelerator Innovation Zone.

Morten Kolle Christensen and his Danish team were demonstrating Pixop Live for the first time at IBC. It takes HD SDR signals from the field and almost instantly upscales them to UHD 4K for real-time processing and distribution with a latency of less than 600 ms.

The side-by-side demo was very impressive because it’s always that last little bit of the workflow video stream that is expensive and time-intensive to deliver, let alone level-set to a higher quality whether LIVE or from library.

Christensen explained how studios/streamers/networks with large libraries of shows/movies are interested in using Pixop Live to bring the content up to the quality standard viewers expect to see by making it fast, easy and economical to refresh the content in richer, more vibrant UHD 4K.

Christensen noted that Pixop Live also fully supports NextGen TV standards, including Japan’s ISDB-Tb and Brazil’s

TV 3.0/DTV initiative.



Source - IBC

Centralized – Harmonic introduced their new hybrid centralized management system at IBC.

And speaking of working to improve sports enjoyment, Harmonic introduced their hybrid streaming solution.

The new system includes an AI-centralized management system that leverages an enhanced VOS360 and SaaS for streaming for improved UHD channel density and bandwidth efficacy for 5G content delivery.

AI has also been incorporated to add Multiview and low-latency features as well as automatic highlight generation and speech-to-text transcription.



Top Performers – OWC introduced the new Thunderbolt 5 network dock (l) and Jellyfish S24 server storage system.

Okay, we admit it, we've been impressed with OWC for years because they build some of the best prosumer products we've seen.

At this year's IBC, they introduced the new Thunderbolt 5 Dual10GbE network dock that is probably one of the best and most versatile connectivity solutions we've seen for filmmakers.

The Thunderbolt dock has been around for a long time and the latest version is proof of their commitment to the content creation market including dual fully independent 10Gb/s ethernet ports, front-facing 2.5Gb/s ethernet port and about every connection a content creator or video production person would need.

We realize a lot of workflow and production is done in the cloud today and it's "pretty reliable," but there's just something "comforting" about having a fully loaded SSD production server where you can see/touch it.

The new OWC Jellyfish S24 lets you load the unit up with as much high-performance, highly reliable Mercury Extreme Pro SSD (up to 384TB) you need to handle any project, including a blockbuster film.

Designed to work on a dual 100G ethernet backbone, it delivers superior sustained read/write performance with almost no latency to keep your production project on time and yes, we're impressed.



Source - IBC.

Content, Business Management – With enhanced AI and agentic workflow capabilities, Avid introduced their new content data platform to eliminate siloed workflow, making it visible across the content organization.

Avid keeps raising the bar for the content industry and the new content data platform that does away with siloed production areas and works across editorial, operational and business areas.

Optimizing AI and agentic workflows, it transforms content from static files to intelligent, connected and actionable data.

The scalable and intelligent solution helps streamline operations, optimize content value and keep pace with the audience's constantly changing ideas of what they want/need to watch.



Source – Blackmagic Design

Buzzworthy – Blackmagic Design’s new Camera ProDock created a lot of buzz at IBC by mating full-featured production capabilities with the Apple iPhone 17.

Blackmagic Design has always impressed us with their steadily improving line of cameras that cover almost every feature, capability and performance issue a filmmaker might have—all at a price that is unbelievably reasonable.

At IBC, the new Blackmagic Camera 9.8 was unveiled that supports Apple ProRes RAW, a new lower price for their 12K LF and 17K URSA cine camera kits and more.

In addition, because we can only guess that they want almost everyone on the production team to use Blackmagic hardware and solutions, they just keep improving DaVinci Resolve production solution with 20.1.1 that supports edit, color, VFX and more.

We would say nothing they develop/deliver for content creation folks would surprise us but then they went and announced the new Blackmagic Camera ProDock.

If you’re one of those shooters who has been won over to use an iPhone as a full-featured production camera, you’ll be absolutely wowed by the new Camera ProDock.

The dock provides you with all of the professional camera connections you’ll need for the iPhone 17 Pro and Pro Max.

It really does put all the power and performance at your fingertips whether you’re a YouTube creator or a seasoned professional.

All we can say is Grant Petty, Blackmagic's CEO, and his team are insane ... insanely committed to making life easier, better and less expensive for the content creation industry.

There was a lot we missed at this year's IBC, but the show **only** lasted four days and there is only so much we could see in the packed halls...next year.



Source - Amsterdam Sights

Off the Floor – While there was plenty to take in at this year's IBC, visitors also had time to enjoy Amsterdam's historic Bruin cafes and Dutch cuisine.

After all, only an idiot spends 100 percent of his/her time in the RAI Conference Center when Amsterdam has so many charming and colorful Bruin cafes complete with excellent food and excellent hospitality.

There's always next year.

###

Andy Marken – andy@markencom.com - is an author of more than 800 articles on management, marketing, communications, industry trends in media & entertainment, consumer electronics, software and applications. An internationally recognized marketing/communications consultant with a broad range of technical and industry expertise especially in storage, storage management and film/video production fields; he has an extended range of relationships with business, industry trade press, online media and industry analysts/consultants.

2) A Little Creative AI is Fine, But..., Andy Marken (Content Insider #923—Rogues, andy@markencom.com



Source - "I Robot," 20th Century Fox

"There have always been ghosts in the machine. Random segments of code, that have grouped together to form unexpected protocols. Unanticipated, these free radicals engender questions of free will, creativity, and even the nature of what we might call the soul. – Dr. Alfred Lanning, "I Robot," 20th Century Fox, 2004

Let's first say we're not really anti-AI and we're a long way from all in pro-AI.

Let's simply say we've *never* seen a movie that ended well when AI was wantonly released to do its own thing.

And everyone knows what follows AI. Yes, robots because "they" like to say that AI will give you a little freedom. But add a few robots to take over your mundane tasks and you'll have even more freedom.

Perhaps that's why AI was one of the biggest hurdles the industry faced a couple of years ago during the dual writers/actors strikes and it's still that ... a hurdle.

The unions did their darndest trying to set up definitive guardrails to protect creators (people types) but that has been and will continue to be difficult.

The problem is, AI is everything and it's nothing.



Source - Netflix

It's just something looming out there until management, creatives, technologists step in and use it to do their thing ... entertain folks.

Back at his last Disney shareholder meeting, CEO Bob Iger said, *"AI may be the most powerful technology that our company has ever seen."*

Jeez, that's putting it mildly.

Disney has already begun using it to speed Imagineering's animation workflow.

Post production/VFX teams have been using AI-enriched Adobe, Avid, internally developed and other tools and processes to help audiences connect with characters/storylines more quickly to enjoy the films/shows.

The biggest concern for IP owners like Disney that have huge vaults of valuable intellectual property is that their application of AI might end up "copying/borrowing" from the past and diluting the worth of the valuable property.

Once it's out in the unfettered public, the value goes down dramatically.

The problem is there are two opposing segments of the service/support groups with significantly different goals.

First, there are the organizations that are embedded in and committed to the creative industry to help them develop/create future content growth – Avid, Adobe, and a bed rock of services and products.

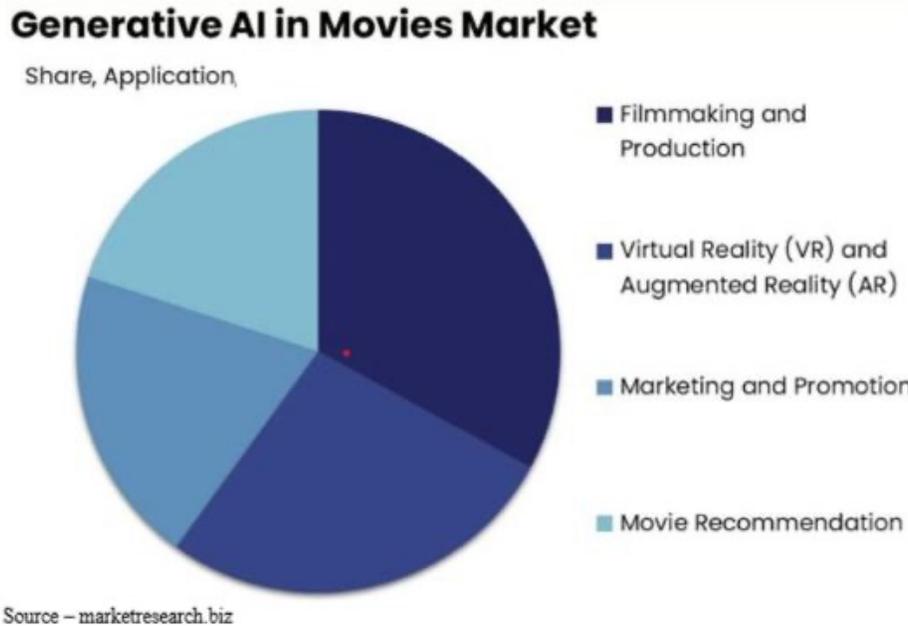
Then there are the AI folks - OpenAI, Google DeepMind, Anthropic and other AI organizations including China's DeepSeek – who want **their** technology distributed and used as broadly/quickly as possible.

Their approach/goal is to develop and release new versions as rapidly as possible and as entertainment folks say ... fix it in post.

This is especially as AGI (artificial general intelligence) begins to appear on the horizon.

Hey, technology that can match/exceed human-level performance, never gets tired/wants a break, never asks for a raise and is willing to work for a few kilowatts of power ... great!

We won't make that major leap into darkness because there are too many areas where it is being successfully used to the benefit of the studio/streamer, creatives and content consumer.



The areas where it is being used successfully include:

- Preproduction/scriptwriting assistance
- Content creation, personalization
- VFX, CGI, post
- Pre/post audience engagement, analysis

Yes, in each of these areas there will be some people replaced by automation; but we're optimistic that most of these positions will be revised, expanded by the next generation of creatives to make their work more interesting/fulfilling and fewer projects will be delivered that sorta, kinda, maybe ... suck.

Pre/Post

While they are at different sides of the spectrum, we view the pre and viewer feedback uses of AI as intertwined because to a large extent they're based on subscriber/viewer data.

Perhaps the most successful in leveraging this data has been Netflix because it was early to capture, process, mine and make decisions based on viewer data globally and in each country it is in. It is in 190 countries where it determines which projects to acquire/greenlight and project/genre mixtures that will attract and retain subscribers.

Also successful are Amazon Prime and Disney (a creative studio that has increasingly built its success by understanding and knowing how to use technology to capture more and more viewer data).

Today, instead of having data that scientists pour through to make sense of it, AI is used to focus on people's key pressure/decision points.

The technology helps management and creative teams better understand what viewers are feeling/experiencing based on the shows/movies they watch, when they watch them, how long they watch them and when/why they disconnect with the video story.

We know it sounds boring as h*** and it is to everyone but data analysts who really get off on that sorta stuff.

The information that comes out of that data doesn't dictate a project, write the script/shot list or honestly even ensure that the finished project is going to be a winner regardless of where people put their seats.

But it provides some guidelines that give creatives a better idea that the final film/show is going to attract and keep audiences seated till the end.

And that, not reviewers' analysis/discussion is really what counts to people involved in the video story.



That's how anyone in the creative industry wants -- keep people from getting up and leaving the theater/room, falling asleep or getting ill. Instead, they stay

involved in the story until the closing credits and then sit back or go out and tell others... “Gawd that was good!”

Viewers are happy, studio/streamer bosses/shareholders are happy, and yes, everyone involved in the project is proud of what they created/delivered regardless of what some overbearing critic has to say.

Show/film production has always been a labor-intensive ordeal with teams of folks working together locally and remotely. Selective use of AI-powered tools can streamline and speed various aspects of the projects.



Source - Adorama

AI is already widely used by studios/streamers to determine projects even before the first frame is shot.

One of the first uses of AI type tools was with products like ScriptBook which analyzes/evaluates scripts, can project proposed success and recommend changes which can be considered/implemented to improve audience acceptance.

If this feels a little “sterile,” it is but then, the **analytical** aspect of it helps reduce the writers’/directors’ emotional attachment to segments, locations, and other aspects so all of the team can focus on arriving at a script and project that attracts/involves the audience.

But when it comes to actual script writing, we agree with WGA ... the stuff stays out of the writers’ room.

AI can consume vast amounts of film/story data from across time, analyze the strengths/weaknesses and even make suggestions/recommendations **but** it can’t feel or express emotion.

That's what we watch movies/shows to connect with because it could check all the boxes technically but if it doesn't resonate with the viewer, it's just ... nice!



AI has great potential to simplify the pre-production process by helping to plan schedules, find locations that best fit the storylines and support in other preparatory processes.

Implementing AI will automate the planning of shooting schedules according to the availability of actors that will save time and increase efficiency.

In addition, AI systems can analyze the locations described in the screenplays and recommend actual sites for shooting the scene, saving lots of resources in location scouting.

Teams are increasingly using AI tools to help plan schedules/production windows. It can also find just the right locations that work well within the storyline and the budget as well as support the preparatory processes including even recommending actors for specific roles based on predetermined criteria including emotion they project.

The area where we have mixed feelings about is character manipulation/making.



Source - Sony

Some films have used deepfakes to de-age actors in *The Irishman*; and heavy makeup work and CGI has been used extensively but both are time-consuming and project budget killers in today's cost-reduction environment.

But the aging/deaging of Tom Hanks in *Here* would have time- and budget-prohibitive without AI tools.

There has been a constant battle when it comes to deepfakes with some studios/project owners initially having background actors have pictures of them taken during a project so they "might" be used in future projects.

In other words, one project, one paycheck and they own your image going forward.

Well, that worry has been taken off the table and agreements have to be signed that clarify that the individual owns her/his image/voice in perpetuity.

Scarlett Johansson and others have challenged – and won – studios and AI techies on their use of their physical likeness and voice so that concern should be put to bed.



Source – Disney



Source – Vestron Pictures

Hanks has said publicly that he is willing to sign an agreement that would authorize “someone” to use his likeness in projects after he retires or dies with the agreement that his beneficiaries be paid for his “future work” and they have final decision approval as to “his” appearances.

Personally, we would love to see Chad Boseman, Brandon Lee, Patrick Swayze and several other actors in projects they never got to perform in but watching some deepfake of them?

Nope ... that's BS!

We'll settle for rewatching the old stuff and celebrate/enjoy what they actually did because we'll know they're realistic, not real.

That makes a huge difference that any seat in the seat can understand, embrace.



Source – Adobe



Source – Disney

Adobe rolled out a ton of new features, capabilities and services this year at NAB for Adobe Firefly and their Creative Cloud.

The new AI solutions still keep the creative production pro squarely in control of the production/post process but most attendees at the show said it relieved them

of a lot of the drudge/grunge work and enabled them to focus on the fun, important stuff ... the creative post production.

VFX and CGI have taken giant leaps forward in improving project quality/appearance while dramatically reducing the cost and postproduction time requirements.

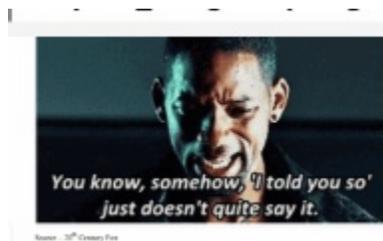
The new tools have enabled James Cameron to reduce the total cost of the upcoming *Avatar 3* "Fire and Ash" while making dramatic advances in the film's "look and feel."

As with *The Planet of the Apes*, it will be very interesting to see how the new technologies have been used to make the film real rather than artificial.

That's the real test!

While China's *Ne Zha 2* surprised almost everyone by capturing the interest and ticket investment around the world this year, Disney's *Inside Out 2* production team can still be proud of the VFX work they did to deliver a great film for people of all ages.

In all of the above projects, it's amazing to see what a dedicated and creative artists can do with AI and advanced tools to make a good video story great and a great film/show mind boggling.



And, thanks to [today's](#) global streaming services, we get the opportunity to enjoy international projects right at home.

Before, finding and watching them usually meant hunting down an art theatre and watching whatever they offered.

Prior to the introduction of AI localization tools, lip-syncing and matching facial expressions was **really** expensive and time-consuming and even adding subtitles wasn't cheap.

But with AI-enabled subtitling tools, it quickly became economic and is increasingly accepted by audiences around the globe.

Allan McLennan, president/CEO of PADEM Media Group, emphasized that James Cameron made a fast turnaround of the initial release of *Avatar* blockbuster in

nearly 50 subtitled languages, but it was Bong Joon Ho's *Parasite* that proved English-speaking audiences told studios that they were more interested in the content rather than the audio delivery.

McLennan noted, *"Now that studios and streamers are able to quickly, effectively and economically translate and add the appropriate subtitles with AI-enabled tools, more very interesting projects from Africa, Central/South America, SEA are finding new audiences regardless of where they have been created."*

We enjoy hearing the film/show in the native language but also being able to follow the dialogue on the bottom inch of the screen. That includes a lot of "English" projects from Australia, Canada, British Isles; and at times, even those produced in the US.



Source - 20th Century Fox

But the biggest subscriber issue streamers haven't solved is helping viewers find the show/movie they might want to watch *if* it happens to be in "the other service's" library.

They have all of our subscription/viewing history so they probably know better than we do what stuff in their library we'd probably find interesting.

The problem is we have three ad-supported and two FAST services and sometimes we catch a trailer or review on a project we think we might be interested in.

Yeah, good luck in finding it!

Open any one of them and even before you can tell it what you're really interested in, the service simply says, *"Here's a list of stuff you'll like so ... keep subscribing."*

After 10-15 minutes of searching, we usually simply "settle" on something or dumb luck leads us to the video project we were looking for.

Someday, they'll agree not to just go back to the old bundling solution but an intelligent front end that let's you sign in and check out which shows/movies you want across all of **your** subscription services.



Source - SeniorGIF

Yeah right!

Okay, maybe that's not an AI tool they'll develop and release to streaming users but...

There are still plenty of tools streamers and creatives will embrace.

Front end tools that help streamer/studio bosses determine which projects are going to resonate with seats in seats?

Sure!

We all know bosses don't want to be wrong ... especially when **their** job is on the line.

Subscriber/viewer tools that make it easy to choose the next show/movie to watch so the subscriber doesn't have to think?

Sure!

After "one of those days," all folks want to watch what they want to watch without a lot of decision making.

Creators will embrace the pre project/production work because jeezz, it's a real pain in the behind, not really creative but rather necessary "busy work."

We're already knee deep in VFX. CGI and production/post people are doing some fantastic work with AI tools to turn out great projects and still get home in time for dinner.

But the idea that AI is going to write and produce a film/show from beginning to end?



Source – 20th Century Fox

Nope ... that dog ain't gonna hunt!

Oh sure, techies would love to see their s**t make a movie and maybe/probably some viewers wouldn't see the difference between the human and AI projects, but real creatives would.

In addition, they'll laugh their ***es off when *I Robot's* Spooner looks at them and says, "*You must be the dumbest, smart person in the world.*"

After all, **no one** would put their money on the line because once that project was released, there would be nothing to copyright and make a profit on.

And if you can't copyright it, you can't get an award/statue and folks who aren't really creative or scrupulous will strip off and monetize/sell the stuff that's valuable leaving the studio, streamer, creative with a ravaged idea.

As Spooner said, "*Human beings have dreams. Even dogs have dreams, but not you, you are just a machine. An imitation of life.*"

It's okay that you use the tools today as long as everyone stays alert as to how it's going to affect and impact the creative work **you** want to do and view tomorrow.

Forget the "we're here to help you, make your job easier," BS from Open AI and the rest.

They think a show without feeling, without emotion, without conscience is really cool.

###

3) Hard Disk Drives Unexpectedly Not Dead, David Rosenthal, 9/18/25, <https://blog.dshr.org/2025/09/hard-disk-unexpectedly-not-dead.html>

As I read Zak Killian's *Expect HDD, SSD shortages as AI rewrites the rules of storage hierarchy — multiple companies announce price hikes, too* I realized I had forgotten to write this year's version of my annual post on the Library of Congress' *Desihning Storage Architectures* meeting, which was back in March. So below the fold I discuss a few of the DSA talks, Killian's more recent post, and yet

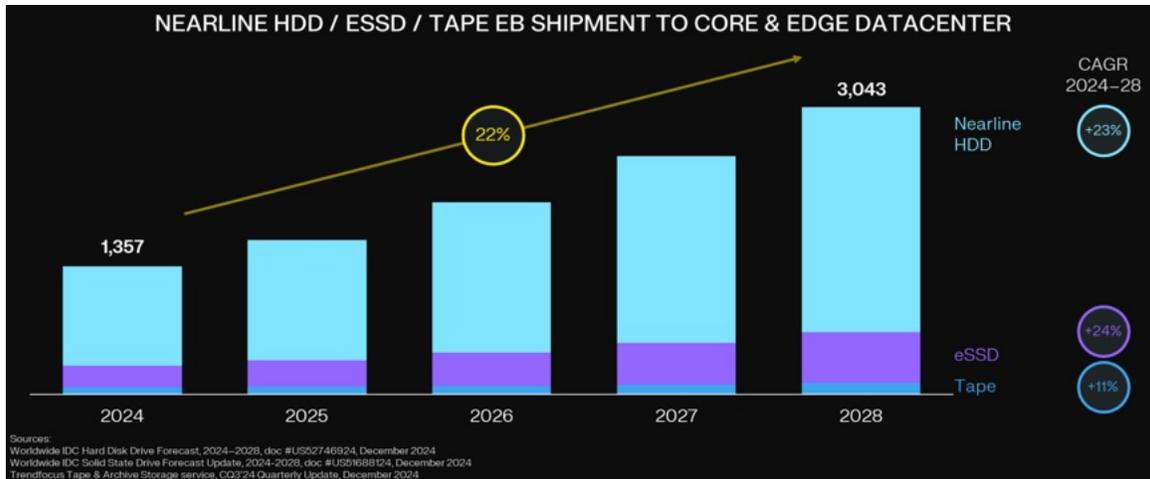
another development in DNA storage. The TL;DR is that the long-predicted death of hard disks is continuing to fail to materialize, and so is the equally long-predicted death of tape.

Killian's post starts: The computing market is absolutely ablaze with AI-driven growth. Regardless of how sustainable it might be, companies are spending untold amounts of wealth on hardware, with most headlines revolving around GPUs. But the storage market is also under pressure, especially hard drive vendors who purportedly haven't done much to increase manufacturing capacity in a decade. TrendForce says lead times for high-capacity "nearline" hard drives have ballooned to over 52 weeks — more than a full year.



Western Digital is: warning of "unprecedented demand for every capacity in [its] portfolio," and stating that it is raising prices on all of its hard drives.

The unprecedented demand from AI farms is because: You don't just need the data required to run inference. You also need the history of everything to prove to regulators that you're not laundering bias, to retrain when new data comes in, and to roll back to a previous checkpoint if your fine-tuned model goes feral and, say, starts referring to itself as MechaHitler. This stuff can't go to offline storage until you're certain it isn't needed in the short term. But it's too big to live in the primary storage of all but the beefiest servers. Thus, the need for nearline hard drives.



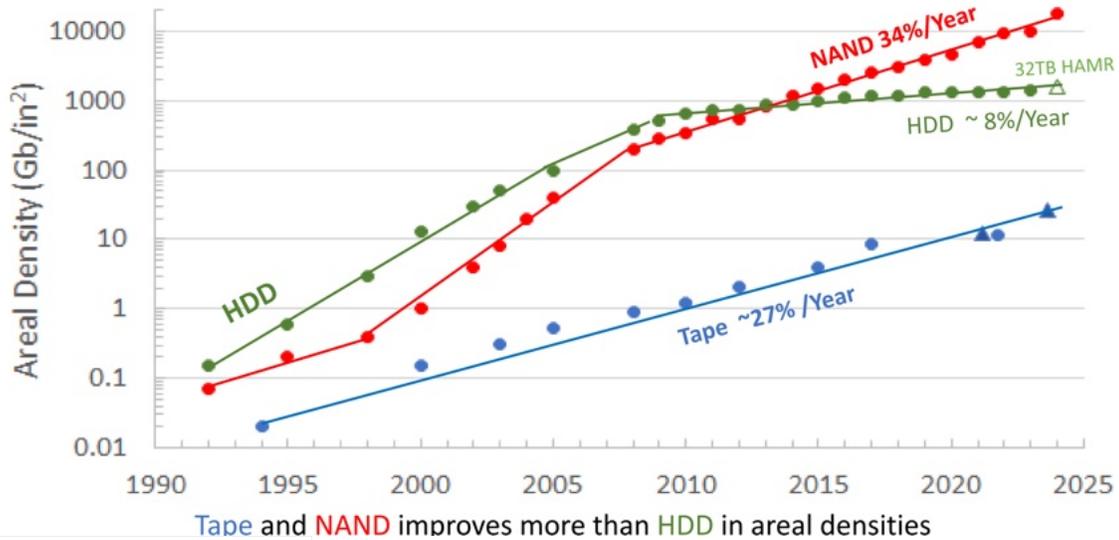
WD's projection

At the meeting, Western Digital's Dave Landsman's *HDDs are here to stay* made the same point with this graph using data from IDC and TrendFocus. They are projecting that both disk and enterprise SSD will grow in the low 20%/year range, so the vast bulk of data in data centers will remain on disk. Landsman claims that SSDs are and will remain 6 times as expensive per bit as hard disk and that 81% of installed data center capacity is on hard disk.

Keeping the data on hard disk might actually be a good idea. *Sustainability in Datacenters* by Shruti Sethi presented a joint Microsoft/Carnegie-Mellon study of the scope 2 (operational) and scope 3 (embedded) carbon emissions of compute, SSD and HDD racks in Azure's data centers. The study, *A Call for Research on Storage Emissions* by

Sara MacAllister *et al* concluded that: an SSD storage rack has approximately 4× the operational emissions per TB of an HDD storage rack. Storage devices (SSDs and HDDs) are the largest single contributor of operational emissions. For SSD racks, storage devices account for 39% of emissions, whereas for HDD racks they account for 48% of emissions. These numbers contradict the conventional wisdom that processing units dominate energy consumption: storage servers carry so many storage devices that they become the dominant energy consumers.

SSD racks emit approximately 10× the embodied emissions per TB as that of HDD storage racks. The storage devices themselves dominate embodied emissions, accounting for 81% and 55% of emissions in SSD and HDD racks, respectively.



Areal Density Trends

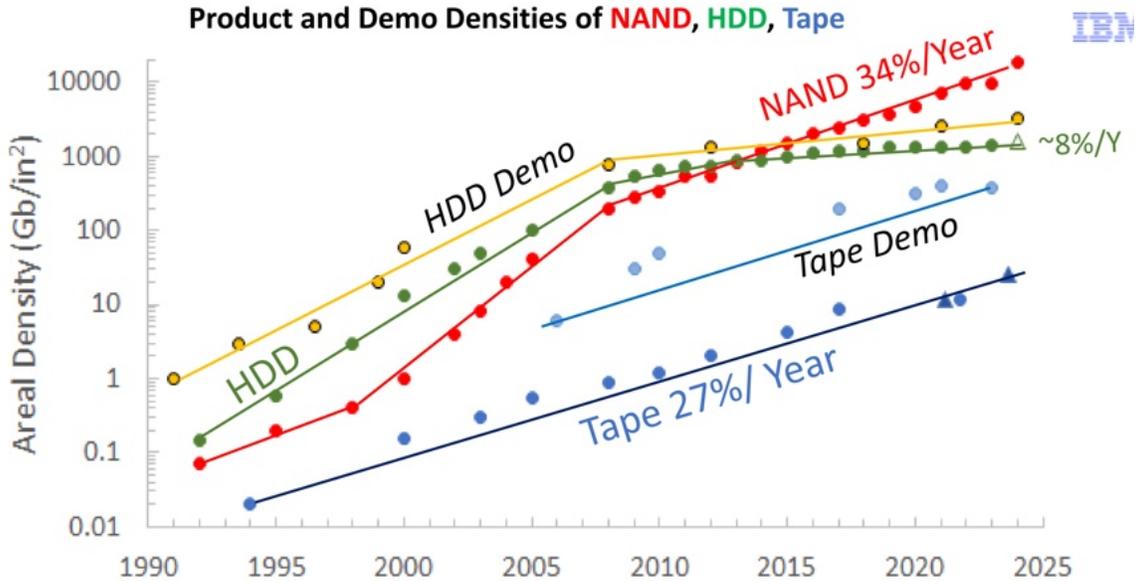
As usual, the authoritative word on the performance of the storage industry comes from IBM. Georg Lauhoff & Sassan Shahidi's [Data Storage Trends: NAND, HDD and Tape Storage](#) added another year's data points to their invaluable graphs and revealed that:

- NAND areal density continues to increase rapidly, because 3D scales faster than the 2D of disk and tape.
- Disk's 8%/year areal density increase continued, but note that although their graph includes Seagate's 32TB HAMR drive the effect of Seagate's and later WD's deployment of HAMR didn't really start until later in 2025.
- Tape continued its 27%/year increase.

Coming from a tape supplier this comment isn't surprising but it is correct:

Despite the promise of alternative archive storage technologies, challenges persist. Enduring relevance of tape storage, which itself is rapidly evolving.

The main problem being that the huge investment and long time horizon needed to displace tape's 7% of the storage market can't generate the necessary return.



Product vs. Demo

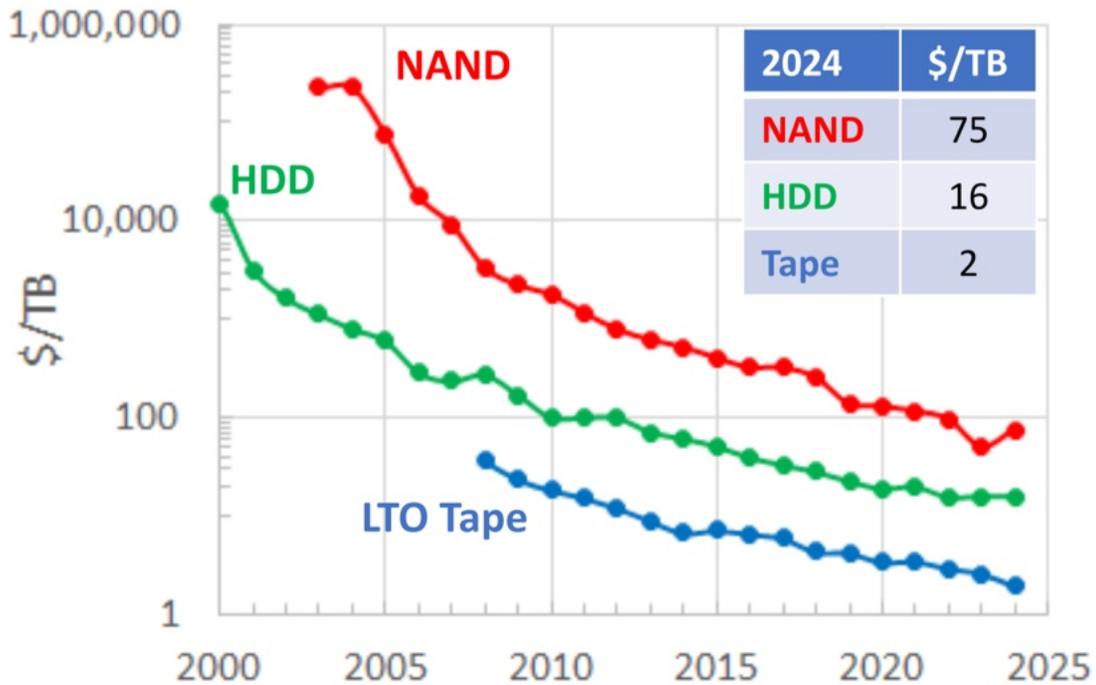
One fascinating graph shows the difference between demonstrations and products for tape and disk. I keep pointing out the very long timescales in the storage industry. In January's [Storage Roundup](#) I noted that HAMR was just starting to be deployed 23 years after Seagate demonstrated it. Lauhoff & Shahidi's graph shows that the current tape density was demo-ed in 2006 and shipped in 2022, and that disk's current density was demo-ed in 2012.

Areal Density Yearly Change			
	NAND	HDD	LTO TAPE
Products since 2012	33%	8%	27%
Roadmaps shown above	43%	15%	28%

Source

This graph reinforces that tape's roadmap is credible, but the good [Dr. Pangloss](#) noticed the optimism of the NAND and disk roadmaps. New technologies tend to scale faster at first, then slower as they age. So, it is likely that the advent of HAMR will accelerate disk's areal density increase somewhat. And it is possible

that the difficulty of moving from 3D NAND to 4D NAND will slow its increase.

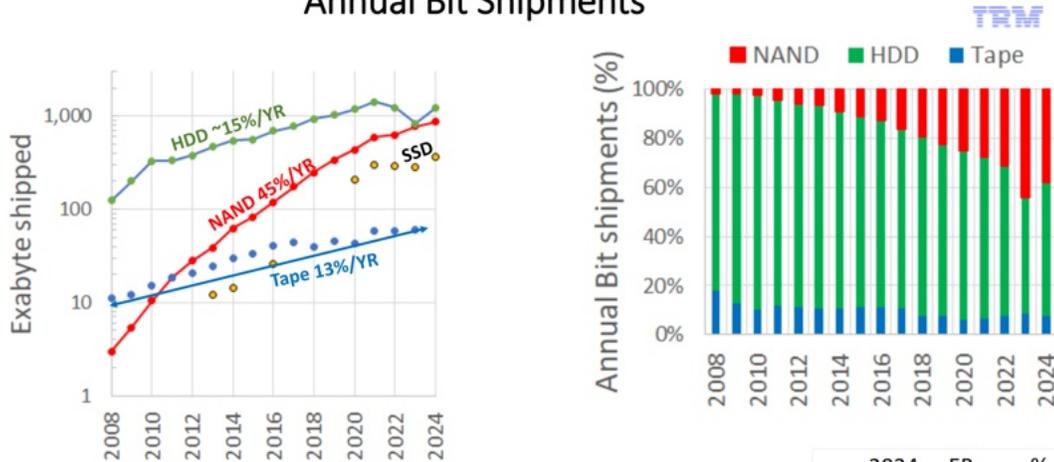


Cost Ratio

Lauhoff & Shahidi's cost ratio graph shows that the relative costs of the different media were roughly stable. If Killian is right that the disk manufacturers are increasing prices and lengthening lead times because of demand from AI, this could be different in next year's graph. But Killian also notes that, despite the fact that QLC SSDs are at least "four times the cost per gigabyte":

Trendforce reports that memory suppliers are actively developing SSD products intended for deployment in nearline service. These should help bring costs down once they hit the market. But in the short term, we can expect the storage crunch to cause rising SSD prices as well, at least for enterprise drives.

Annual Bit Shipments



	2024	EB	%
NAND	867		39%
SSD	369		16%
other	498		22%
HDD	1,227		55%
LTO Tape	153		7%

NAND and HDD bit shipments approaching each other.
 HDD >> SSD bit shipments
 More **Tape** demand: Cheap, energy efficient and easier to use (S3)

Lauhoff et al. Library of Congress 2025

Annual Bit Shipments

Lauhoff & Shahidi's bit shipment graph is interesting for two reasons:

- Disk's proportion of total bit shipments *increased*.
- They started tracking the proportion of NAND flash that was SSDs. They represented only about 30% of disk's bit shipments. The claim that the bulk of data still lives on hard disk is true and looks to continue. Disk ships mostly to the nearline enterprise market, while SSDs ship mostly to the online enterprise market. Disk is shipping nearly three times as many bits.

Tape's predicted demise is just as delayed as disk's. Back in July Simon Sharwood posted *And now for our annual 'Tape is still not dead' update*: Shipments of tape storage media increased again in 2024, according to HPE, IBM, and Quantum – the three companies that back the Linear Tape-Open (LTO) Format.

The three companies on Tuesday claimed they shipped 176.5 Exabytes worth of tape during 2024, a 15.4 percent increase on 2023's 152.9 Exabytes.



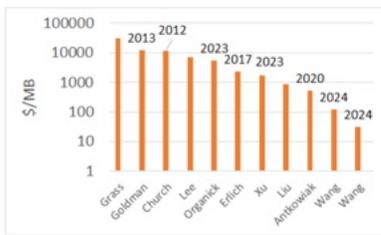
DNA Storage

I have been writing skeptically about the medium-term prospects for DNA storage since [2012](#) and Lauhoff & Shahidi share my skepticism in their graph of the technology's progress in the lab. DNA can only compete in the archival storage market, so the relevant comparison is with LTO tape. Even if you believe Wang's estimate, DNA is more than *ten million times* too expensive.

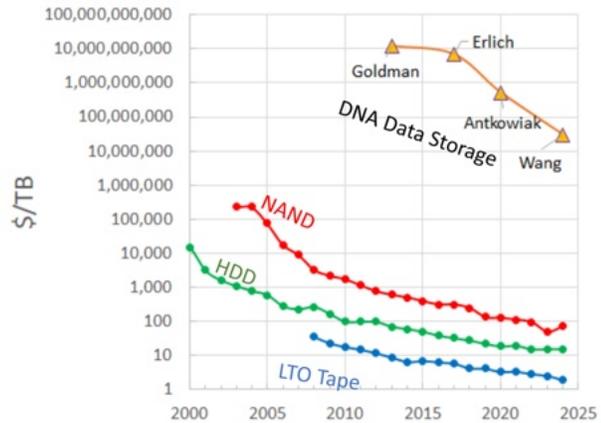
DNA Data Storage vs Traditional Data Storage Cost



DNA Data Storage cost by different authors



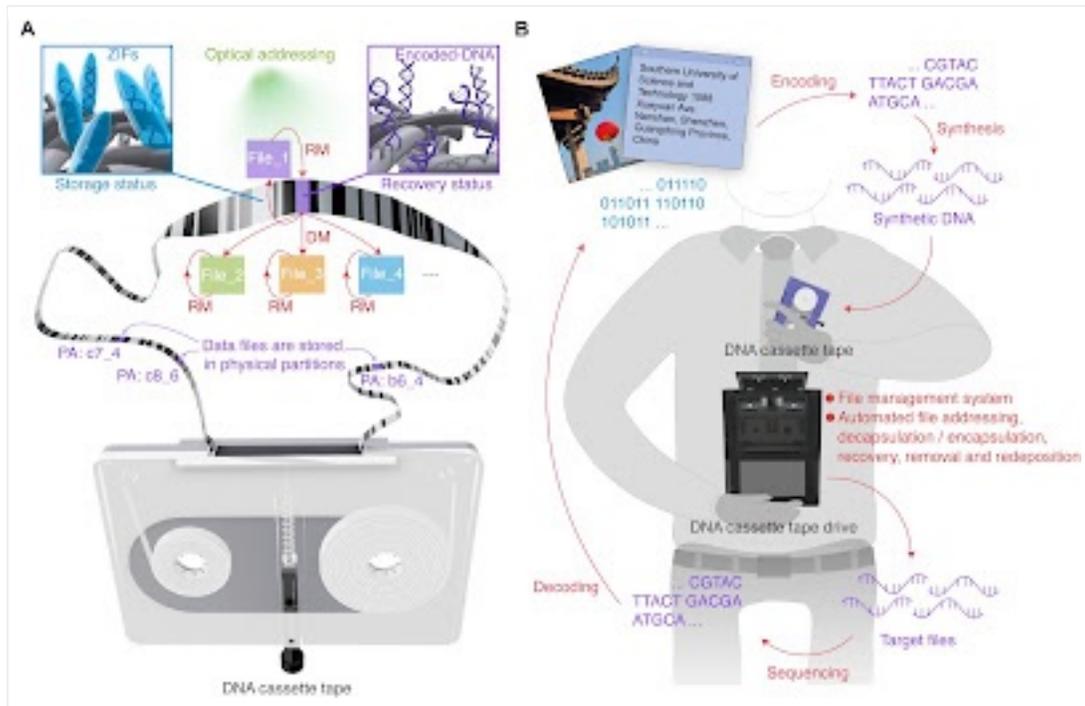
Adapted from Wang et al, 2024
doi.org/10.1002/adv.202411354



DNA Storage slow and expensive

Figure 1

Via Brandon Vigliarolo's [Boffins invent DNA tape that could pack 375 petabytes into an LTO cart](#) we find the latest idea for improving DNA storage in [A compact cassette tape for DNA-based data storage](#) by Jiankai Li *et al* from the Southern University of Science and Technology in Shenzhen. Their idea is to deposit DNA on "good old-fashioned polyester-nylon composite tape" that could reside in an LTO cartridge. Vigliarolo [notes that](#): DNA is a very dense storage medium and storage researchers have tried to use it for data storage, but [without much success](#), because it's hard to find info within DNA and read times are slow.



Jiang's team claims to have addressed that problem, establishing a sequence of data partitions on the tape and identifying each of these with a bar code.

The team from Shenzhen's focus on reading is a misunderstanding of the fundamental requirements of the hyperscaler archive market, which are in order:

1. Write bandwidth. [Kestutis Patiejunas](#) pointed this out over a decade ago.
2. Cost per byte. Archival storage is for data that can no longer earn its keep on lower-latency media, so it has to be very cheap.

Read latency and bandwidth are pretty much irrelevant because, as Kestutis Patiejunas said, the main reason the data would be read is a subpoena.

The paper claims they demonstrated: a completely automated closed-loop operation involving addressing, recovery, removal, subsequent file deposition, and file recovery again, all accomplished within 50 min.

Vigliarolo notes that: Jiang's team only wrote 156.6 kilobytes of data to a test tape for their experiment, consisting of four "puzzle pieces" depicting a Chinese lantern. If the data were damaged, it wouldn't assemble correctly. The researchers managed to successfully recover the lantern image without issue, but it took two and a half hours or not-quite one kilobyte per minute.

And the team admit that: DNA synthesis costs are still very high because archival data is guaranteed to be written but is very likely never read, the cost of storing data in DNA is dominated by synthesis. The team effectively admits that they can't

compete.

I summed up my skepticism in 2018's [DNA's Niche in the Storage Market](#), posing this challenge to a hypothetical product team: Engineers, your challenge is to increase the speed of synthesis by a factor of a quarter of a trillion, while reducing the cost by a factor of fifty trillion, in less than 10 years while spending no more than \$24M/yr.

Finance team, your challenge is to persuade the company to spend \$24M a year for the next 10 years for a product that can then earn about \$216M a year for 10 years.

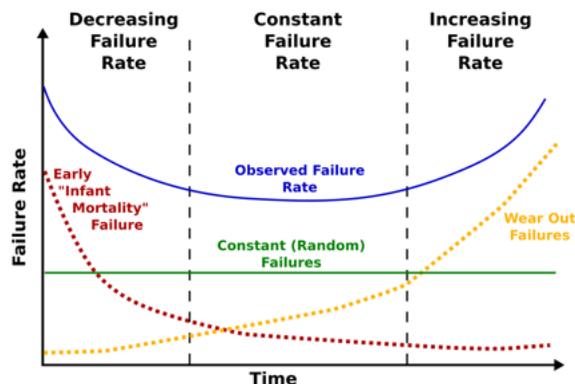
Write bandwidth and cost remain the core problems of DNA storage, and while progress has been made in other areas, both are still many orders of magnitude away from competing with hard disk, let alone LTO tape.

Nevertheless, as I concluded more than seven years ago:

That isn't to say that researching DNA storage technologies is a waste of resources. Eventually, I believe it will be feasible and economic. But eventually is many decades away. This should not be a surprise, [time-scales in the storage industry are long](#). Disk is a [60-year-old technology](#), tape is at least [65 years old](#), CDs are [35 years old](#), flash is [30 years old](#) and has yet to impact bulk data storage.

4) The Bathtub Curve, October 28, 2025, Steve Daniel, sdaniel1958@gmail.com

The economics of long-term data storage are critically dependent not just upon the Kryder rate, the rate at which the technology improves cost per byte, but also upon the reliability of the media over time. You want to replace media because they are no longer economic, not because they are no longer reliable despite still being economic.



For more than a decade Backblaze has been providing an important public service by publishing data on the reliability of their hard drives, and more recently their SSDs. Below the fold I comment on this month's post from their Drive Stats

Team, [Are Hard Drives Getting Better? Let's Revisit the Bathtub Curve.](#)

Wikipedia defines the [Bathtub Curve](#) as a common concept in reliability engineering:

The 'bathtub' refers to the shape of a line that curves up at both ends, similar in shape to a bathtub. The bathtub curve has 3 regions:

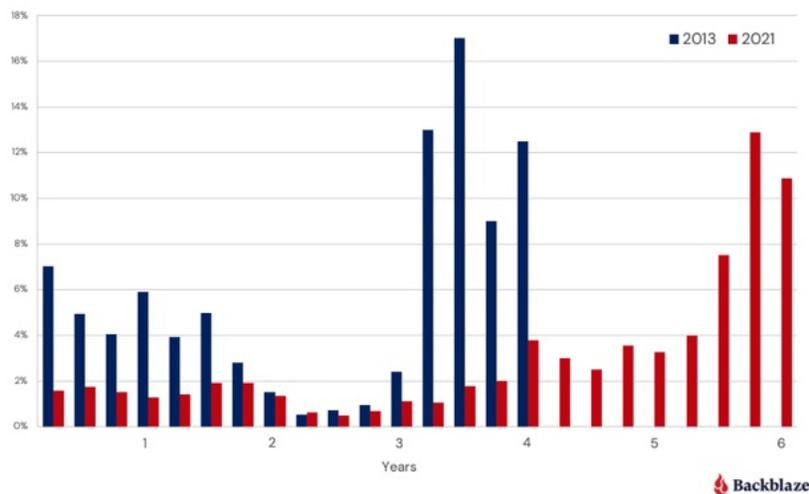
1. The first region has a decreasing failure rate due to early failures.
2. The middle region is a constant failure rate due to random failures.
3. The last region is an increasing failure rate due to wear-out failures.

In 2017's [Storage Failures In The Field](#) I commented on Backblaze's observation that the 10TB and 12TB HDD generations showed much reduced infant mortality:

devoting engineering effort to reducing infant mortality can have a significant return on investment. A drive that fails early will be returned under warranty, costing the company money. A drive that fails after the warranty expires cannot be returned. Warranty costs must be reserved against in the company's accounts. Any reduction in the rate of early failures goes straight to the company's bottom line. Enterprise disks are typically warranted for 5 years, so a disk manufacturer is incentivized to focus engineering effort on eliminating the "first region", the left side of the bathtub, and ensuring that the second region extends past the 5 year mark. Eight years ago Backblaze was starting to see that the [engineers were succeeding in the first region](#):

While the data so far is very limited, with 1,240 disks and 14,220 aggregate drive days accumulated so far, none of these disks (both Seagate models) have failed. The low level of usage means that the disks have been installed and formatted and not much beyond that, but true infant mortality—disks that immediately expire on their first use—hasn't become apparent.

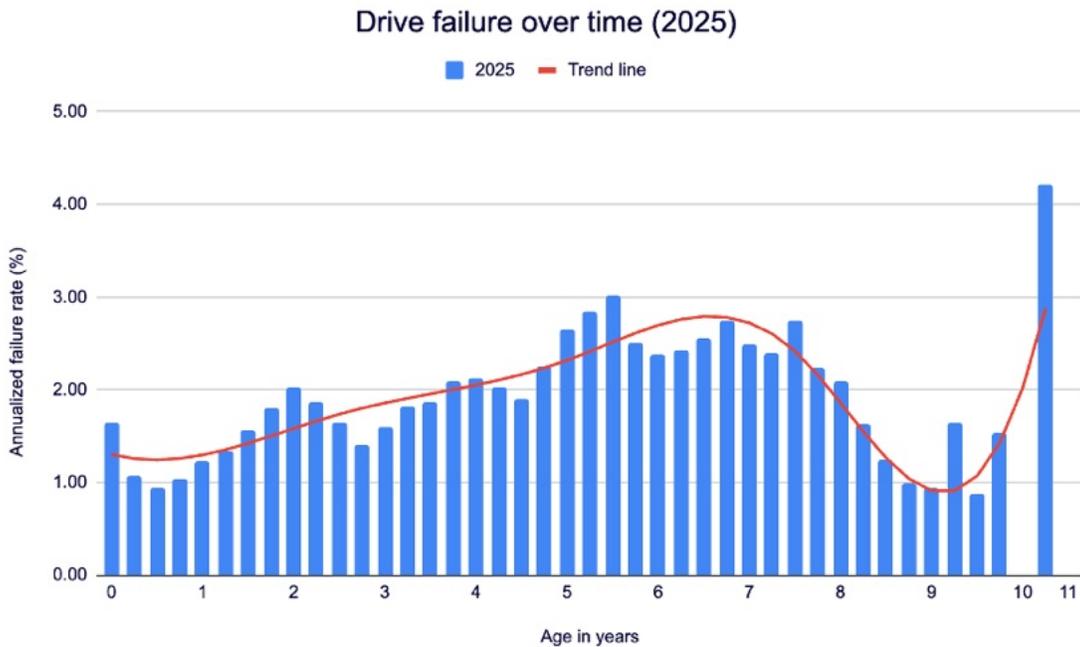
A Comparison of Backblaze Drive Failure Rates Over Time



Four years later in [Drive Failure Over Time: The Bathtub Curve Is Leaking](#) Klein had some much stronger evidence:

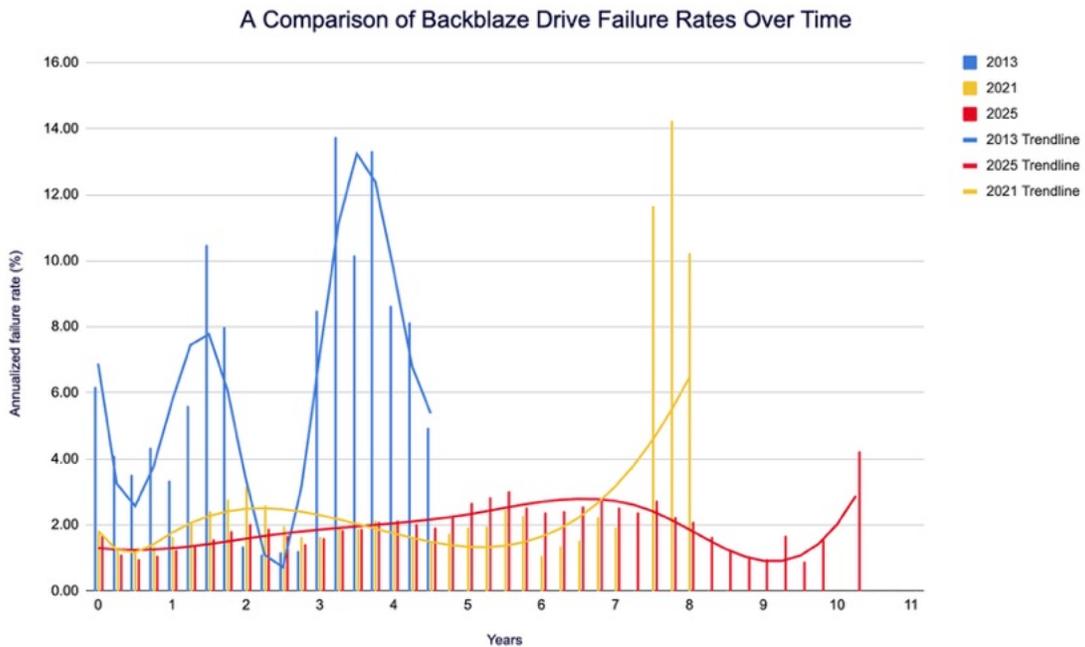
The left side of the bathtub, the area of “decreasing failure rate,” is dramatically lower in 2021 than in 2013. In fact, for our 2021 curve, there is almost no left side of the bathtub, making it hard to take a bath, to say the least. We have reported how [Seagate breaks in and tests their newly manufactured hard drives](#) before shipping in an effort to lower the failure rates of their drives. Assuming all manufacturers do the same, that may explain some or all of this observation.

Note that the engineers hadn't quite succeeded in the second region, as the 2021 failure rate for years 4 and 5 was noticeably higher than for younger drives. But they had succeeded in pushing the major increase in failures out beyond 5.5 years. Everything before that was less than 4% Annual Failure Rate (AFR).



Another four years later in the current graph things have changed quite a bit:

- For the first 5 years the AFR is under 2%.
- For the first 5 years the AFR gradually increases, the first region has been completely eliminated.
- For the three years after a 5-year warranty expires, the AFR is under 3%.
- For the next two years the AFR drops, bottoming out around 1%.
- Everything out to 10 years has an AFR under 3%.
- The third region starts with a spike at 10.5 years.
- Even at 10.5 years the AFR is only just over 4%



Backblaze's final graph puts all three sets of data in one graph and shows a dramatic improvement in drive longevity:

- Drives in 2013 were dramatically less reliable than in later years, both because their AFRs were consistently higher and because their AFR hit 13% after 3 years.
- By 2021 the engineers had kept the AFR around 2% through the warranty, but the drives wore out rapidly in the 7th year.
- Now, the drives are only showing signs of beginning to wear out at 10.5 years.

That is some serious engineering at the long end! And at the short end things are great:

we see that the drive failure rates on the front end of the curve are also incredibly low—when a drive is between zero and one years old, we barely crack 1.30% AFR. The left side of the bathtub is really gone, improving the manufacturers' margins. But if I eyeball the 2025 graph's first 20 quarters, I estimate the AFR averages 1.6%, which implies that over the 5-year warranty 8% of the drives failed. Clearly, the engineers still have work to do.

Remember to check back in 2029 when Backblaze plans to return to this issue.

5) How will your total available market change over the next five years, Steve Daniel, sdaniel1958@gmail.com



A Forecast of the United States Business Economic Demographic Structure November 2025 Update

- [Executive Summary: Challenges](#)
- [BEDD: A Total Available Market History and Forecast 1998-2026](#)
- [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](#)
- [\$\infty/\Delta\$ About Daniel Research Group](#)

Executive Summary: Challenges

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

1. Implemented and announced future economic disruptive US government policies regarding Tariffs, Deportations, and US Government spending and support services.
2. When will these policies be fully enacted and for how long.

Essentially, we are asking four questions

1. What has the administration announce or implied it will or may do?
2. What can they legally do?
3. What has been implemented?
4. What will be permanent and what may be reversed in the future?

Furthermore, the recent Federal Government shut down severely limits 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:
 - GPD 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

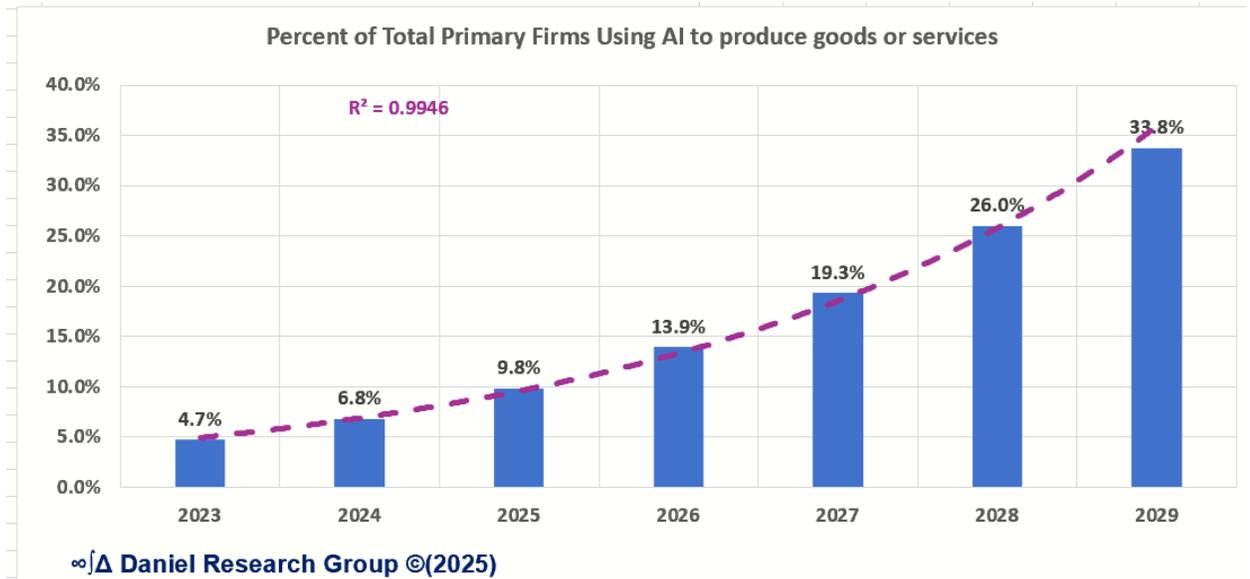
Structural demographic changes will continue to exert downward pressure on labor force participation:

- An aging population and accelerated retirement rates

- Slower growth in new immigration
These factors are expected to constrain workforce expansion and productivity.

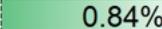
Artificial Intelligence Adoption

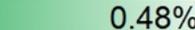
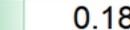
While AI adoption in business operations is progressing more slowly than initially projected, the technology is still reshaping the labor market. The following forecast table and charts is bases on date for the last published Census Bureau Business Trends and Outlook Survey.



Percent of Primary Firm Planning on Using AI to Produce Goods or Services							
Industry	2023	2024	2025	2026	2027	2028	2029
Forestry, Fishing, Hunting, Agriculture	3.3%	3.5%	3.7%	3.9%	4.1%	4.3%	4.6%
Mining	1.1%	3.0%	7.8%	18.8%	38.8%	63.3%	82.5%
Utilities	2.8%	5.4%	10.1%	18.1%	30.3%	46.2%	62.8%
Construction	1.5%	2.5%	4.3%	7.2%	11.8%	18.9%	28.8%
Manufacturing	2.9%	4.1%	6.0%	8.5%	12.0%	16.7%	22.7%
Wholesale	2.7%	4.3%	6.8%	10.6%	16.0%	23.5%	33.2%
Retail	3.3%	4.6%	6.4%	8.8%	11.9%	16.0%	21.1%
Transportation, Warehousing	1.8%	2.5%	3.5%	5.0%	7.0%	9.7%	13.3%
Information	17.7%	22.6%	28.5%	35.1%	42.5%	50.1%	57.8%
Finance, Insurance	7.0%	10.9%	16.5%	24.3%	34.2%	45.7%	57.7%
Real Estate, Rental, Leasing	7.8%	10.4%	13.8%	18.1%	23.3%	29.5%	36.6%
Professional, Scientific, Technical Services	11.7%	16.8%	23.4%	31.7%	41.3%	51.7%	61.9%
Management of Companies, Enterprises	10.5%	14.3%	19.1%	25.1%	32.2%	40.3%	48.9%
Administrative Support, Waste Management, Remediation Services	4.6%	6.6%	9.3%	13.1%	18.0%	24.3%	32.0%
Education	8.4%	12.4%	18.0%	25.4%	34.6%	45.1%	56.0%
Health Care, Social Assistance	4.8%	7.4%	11.4%	17.1%	24.8%	34.5%	45.7%
Arts, Entertainment, Recreation	4.3%	6.6%	9.9%	14.6%	21.0%	29.2%	39.1%
Accommodation, Food Services	1.8%	2.5%	3.3%	4.4%	5.9%	7.8%	10.3%
Other Services	2.3%	3.4%	5.0%	7.2%	10.4%	14.8%	20.6%
Total	4.7%	6.8%	9.8%	13.9%	19.3%	26.0%	33.8%

Employment Current Trends

ADP Employment Monthly Percent Change			
Industry	Sep		Oct
Trade, transportation, and utilities	 0.35%		 0.84%
Education and health services	 0.28%		 0.49%
Professional and business services	 -0.05%		 0.39%
Grand Total	 -0.09%		 0.22%
Natural resources and mining	 -0.22%		 0.11%
Financial activities	 -0.32%		 0.07%
Manufacturing	 -0.38%		 -0.05%
Construction	 -0.26%		 -0.06%
Information	 -0.57%		 -0.07%
Other services	 -0.67%		 -0.45%
Leisure and hospitality	 -0.70%		 -0.72%

Revelio Lab Employment Monthly Percent Change			
Industry	Sep		Oct
Agriculture	 0.48%		 0.18%
Financial activities	 0.08%		 0.10%
Education and health services	 -0.11%		 0.08%
Transportation and warehousing	 0.14%		 0.03%
Construction	 0.14%		 0.02%
Professional and business services	 -0.08%		 0.00%
Leisure and hospitality	 0.39%		 0.00%
Total nonfarm	 0.02%		 -0.01%
Other services	 -0.06%		 -0.04%
Manufacturing	 -0.04%		 -0.04%
Wholesale trade	 0.05%		 -0.05%
Retail trade	 0.09%		 -0.06%
Information	 -0.17%		 -0.07%
Government	 -0.09%		 -0.10%
Utilities	 0.03%		 -0.11%
Mining and logging	 -0.19%		 -0.24%

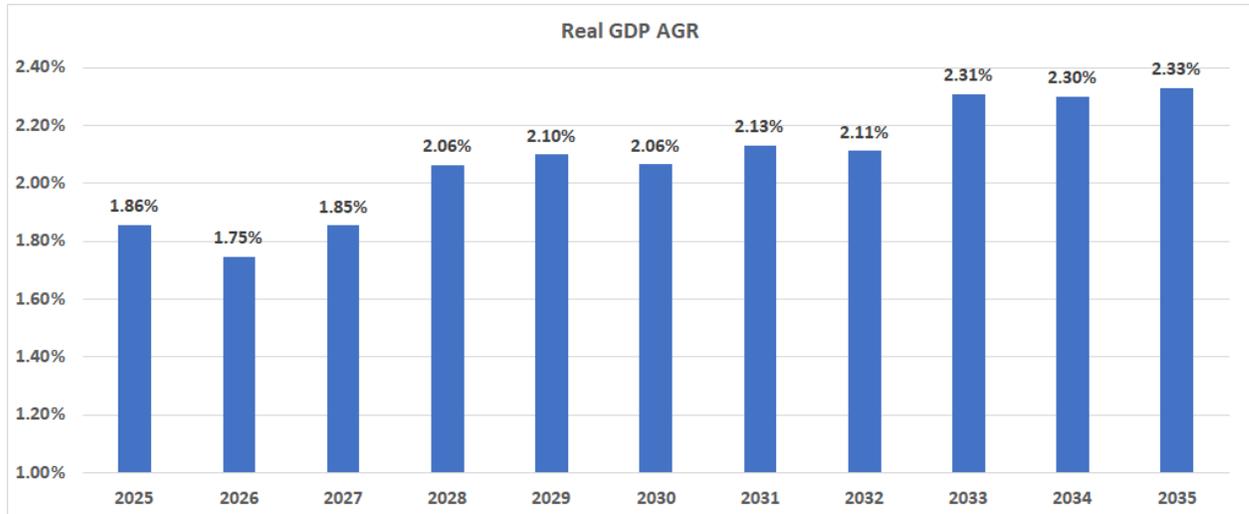
Current Forecasts Five-Year Economic Outlook

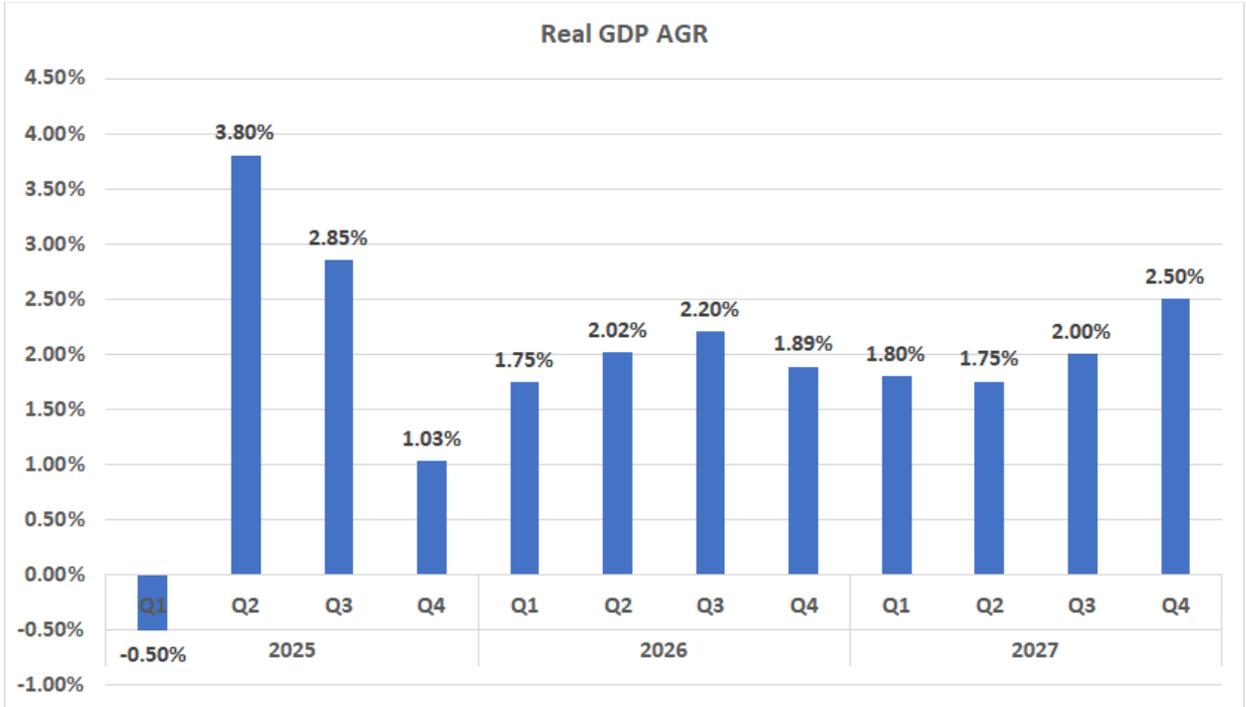
The **DRG BEDD** model forecasts an essentially stagnant economy over the next five years characterized by very low employment growth, and a reduction in the number of businesses.

Sector Sub-Sector Industry Size Class	Metric	Total Total (All) Total							Change	Trends
		2023	2024	2025	2026	2027	2028	2029	24 - '29	
Firms (K)	AGR	6,985.675	7,073.506	7,130.136	7,172.302	7,232.551	7,315.324	7,437.660	364.154	
	Annual Change	2.3%	1.3%	0.8%	0.6%	0.8%	1.1%	1.7%	1.1%	
Primary Firms (K)	AGR	6,344.934	6,336.551	6,306.196	6,263.220	6,235.890	6,223.432	6,238.963	-97.588	
	Annual Change	0.9%	-0.1%	-0.5%	-0.7%	-0.4%	-0.2%	0.2%	-0.3%	
Employees (K)	AGR	159,483	160,219	161,313	161,856	162,384	163,109	164,413	4,194	
	Annual Change	2.6%	0.5%	0.7%	0.3%	0.3%	0.4%	0.8%	0.5%	
Establishments (K)	AGR	8,939.865	9,056.397	9,064.919	9,061.000	9,069.668	9,107.957	9,195.677	139.280	
	Annual Change	2.2%	1.3%	0.1%	0.0%	0.1%	0.4%	1.0%	0.4%	
Annual Payroll (M)	AGR	11,153	11,815	12,434	13,109	13,846	14,649	15,560	3,745	
	Annual Change	6.5%	5.9%	5.2%	5.4%	5.6%	5.8%	6.2%	5.8%	
Employees per Primary Firm	AGR	25.1	25.3	25.6	25.8	26.0	26.2	26.4	1.07	
	Annual Change	1.7%	0.6%	1.2%	1.0%	0.8%	0.6%	0.5%	0.7%	
Establishments per Primary Firm	AGR	1.41	1.43	1.44	1.45	1.45	1.46	1.47	0.045	
	Annual Change	1.3%	1.4%	0.6%	0.6%	0.5%	0.6%	0.7%	0.6%	
Annual Payroll (1,000) per Employee	AGR	69.9	73.7	77.1	81.0	85.3	89.8	94.6	20.90	
	Annual Change	3.8%	5.5%	4.5%	5.1%	5.3%	5.3%	5.4%	5.3%	

© Daniel Research Group ©(2025)

Consensus GDP and UER





Structural Changes US Economy Structural Changes, 2025 Employees

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Size Code	Total		
Size Class	(All)		
Metric	Employees		
Industry	2024 to 2025	CAGR	
Forestry, Fishing, Hunting, Agriculture	-31,347	-1.4%	
Mining	-3,849	-0.6%	
Utilities	14,753	1.7%	
Construction	77,009	0.9%	
Manufacturing	-105,635	-0.8%	
Wholesale	407	0.0%	
Retail	52,701	0.3%	
Transportation, Warehousing	116,454	1.5%	
Information	-9,213	-0.3%	
Finance, Insurance	6,873	0.1%	
Real Estate, Rental, Leasing	59,271	2.4%	
Professional, Scientific, Technical Services	10,945	0.1%	
Management of Companies, Enterprises	36,718	1.4%	
Administrative Support, Waste Management, Remediation Services	-64,118	-0.7%	
Education	210,421	1.4%	
Health Care, Social Assistance	383,505	1.5%	
Arts, Entertainment, Recreation	160,082	4.8%	
Accommodation, Food Services	238,247	1.7%	
Other Services	-15,017	-0.2%	
Public Administration	-14,181	-1.7%	
National Defense	-29,304	-1.9%	
Total	1,094,722	0.7%	

∞|Δ Daniel Research Group © (2025)

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Industry	(All)		
NAICS	Total		
Metric	Employees		
Size Class	2024 to 2025	CAGR	
Less than 5	53,994	0.7%	
5 to 9	24,641	0.3%	
10 to 19	107,210	1.0%	
20 to 99	230,248	0.9%	
100 to 499	68,981	0.3%	
500 to 999	62,439	0.7%	
1000 to 2499	47,570	0.4%	
2500 to 4999	12,718	0.2%	
5000 to 9999	20,363	0.2%	
10000 or more	466,558	1.0%	
Total	1,094,722	0.7%	

∞|Δ Daniel Research Group ©(2025)

Less than 100	485,074	0.8%	
100 or More	609,648	0.6%	

∞|Δ Daniel Research Group © (2025)

44.3% in business with less than 100 employees

US Economy Structural Changes, 2025 Businesses

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Size Code	Total		
Size Class	(All)		
Metric	Primary Firms		
Industry	2024 to 2025		CAGR
Forestry, Fishing, Hunting, Agriculture	-2,717		-0.9%
Mining	345		1.6%
Utilities	-11		-0.2%
Construction	-9,441		-1.2%
Manufacturing	-4,270		-1.8%
Wholesale	-3,439		-1.3%
Retail	-396		-0.1%
Transportation, Warehousing	-6,618		-3.6%
Information	-537		-0.9%
Finance, Insurance	-1,033		-0.5%
Real Estate, Rental, Leasing	3,066		1.0%
Professional, Scientific, Technical Services	-9,625		-1.2%
Management of Companies, Enterprises	-402		-7.1%
Administrative Support, Waste Management, Remediation Services	-8,146		-3.8%
Education	-3,047		-2.3%
Health Care, Social Assistance	5,779		0.8%
Arts, Entertainment, Recreation	8,980		5.8%
Accommodation, Food Services	5,650		1.0%
Other Services	-4,496		-0.6%
Public Administration	5		0.2%
National Defense	-2		-33.3%
Total	-30,355		-0.5%
∞Δ Daniel Research Group © (2025)			

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Industry	(All)		
NAICS	Total		
Metric	Primary Firms		
Size Class	2024 to 2025		CAGR
Less than 5	7,758		0.2%
5 to 9	-1,028		-0.1%
10 to 19	-4,900		-0.7%
20 to 99	-27,162		-5.1%
100 to 499	-4,418		-4.9%
500 to 999	-266		-2.8%
1000 to 2499	-194		-3.3%
2500 to 4999	-74		-3.9%
5000 to 9999	-23		-2.2%
10000 or more	-48		-4.5%
Total	-30,355		-0.5%
∞Δ Daniel Research Group ©(2025)			
Less than 100	-29,750		-0.4%
100 or More	-605		-4.6%
∞Δ Daniel Research Group © (2025)			

99.0% of the loss in Business with less than 100 Employees

US Economy Structural Changes, 2024-2029 Employees

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Size Code	Total		
Size Class	(All)		
Metric	Employees		
Industry	2024 to 2029	CAGR	
Forestry, Fishing, Hunting, Agriculture	31,637	0.3%	
Mining	20,367	0.7%	
Utilities	2,005	0.0%	
Construction	127,896	0.3%	
Manufacturing	243,250	0.4%	
Wholesale	170,393	0.6%	
Retail	174,889	0.2%	
Transportation, Warehousing	252,809	0.7%	
Information	241,700	1.5%	
Finance, Insurance	213,818	0.6%	
Real Estate, Rental, Leasing	102,302	0.8%	
Professional, Scientific, Technical Services	705,768	1.0%	
Management of Companies, Enterprises	127,445	1.0%	
Administrative Support, Waste Management, Remediation Services	518,210	1.1%	
Education	221,367	0.3%	
Health Care, Social Assistance	265,389	0.2%	
Arts, Entertainment, Recreation	10,415	0.1%	
Accommodation, Food Services	411,402	0.6%	
Other Services	-22,898	-0.1%	
Public Administration	-38,911	-0.9%	
National Defense	414,984	4.8%	
Total	4,194,237	0.5%	
∞Δ Daniel Research Group © (2025)			

US Economy Structural Changes			
Sector	(All)		
Sub-Sector	(All)		
Industry	(All)		
NAICS	Total		
Metric	Employees		
Size Class	2024 to 2029	CAGR	
Less than 5	326,554	0.9%	
5 to 9	186,179	0.5%	
10 to 19	478,555	0.9%	
20 to 99	-66,953	-0.1%	
100 to 499	-249,691	-0.2%	
500 to 999	127,219	0.3%	
1000 to 2499	119,062	0.2%	
2500 to 4999	-63,032	-0.2%	
5000 to 9999	393,858	0.8%	
10000 or more	2,942,486	1.2%	
Total	4,194,237	0.5%	
∞Δ Daniel Research Group ©(2025)			
Less than 100	674,644	0.3%	
100 or More	3,519,593	0.6%	
∞Δ Daniel Research Group © (2025)			

16.1% in business with less than 100 employees.

US Economy Structural Changes, 2024-2029 Businesses

US Economy Structural Changes				
Sector	(All)			
Sub-Sector	(All)			
Size Code	Total			
Size Class	(All)			
Metric	Primary Firms			
Industry	2024 to 2029	CAGR		
Forestry, Fishing, Hunting, Agriculture	14,864	1.0%		
Mining	4,284	3.7%		
Utilities	30	0.1%		
Construction	-59,212	-1.6%		
Manufacturing	-4,172	-0.4%		
Wholesale	-6,490	-0.5%		
Retail	-2,390	-0.1%		
Transportation, Warehousing	-39,221	-4.7%		
Information	4,900	1.5%		
Finance, Insurance	-4,479	-0.4%		
Real Estate, Rental, Leasing	-1,560	-0.1%		
Professional, Scientific, Technical Services	-26,210	-0.6%		
Management of Companies, Enterprises	-1,547	-6.1%		
Administrative Support, Waste Management, Remediation Services	-16,996	-1.6%		
Education	3,229	0.5%		
Health Care, Social Assistance	-34,903	-1.0%		
Arts, Entertainment, Recreation	14,519	1.8%		
Accommodation, Food Services	47,391	1.7%		
Other Services	10,225	0.3%		
Public Administration	151	1.3%		
National Defense	-1	-3.6%		
Total	-97,588	-0.3%		

© Daniel Research Group © (2025)

US Economy Structural Changes				
Sector	(All)			
Sub-Sector	(All)			
Industry	(All)			
NAICS	Total			
Metric	Primary Firms			
Size Class	2024 to 2029	CAGR		
Less than 5	72,296	0.4%		
5 to 9	3,776	0.1%		
10 to 19	-20,821	-0.6%		
20 to 99	-130,230	-5.5%		
100 to 499	-20,417	-5.0%		
500 to 999	-1,030	-2.2%		
1000 to 2499	-701	-2.5%		
2500 to 4999	-247	-2.7%		
5000 to 9999	-35	-0.7%		
10000 or more	-179	-3.6%		
Total	-97,588	-0.3%		

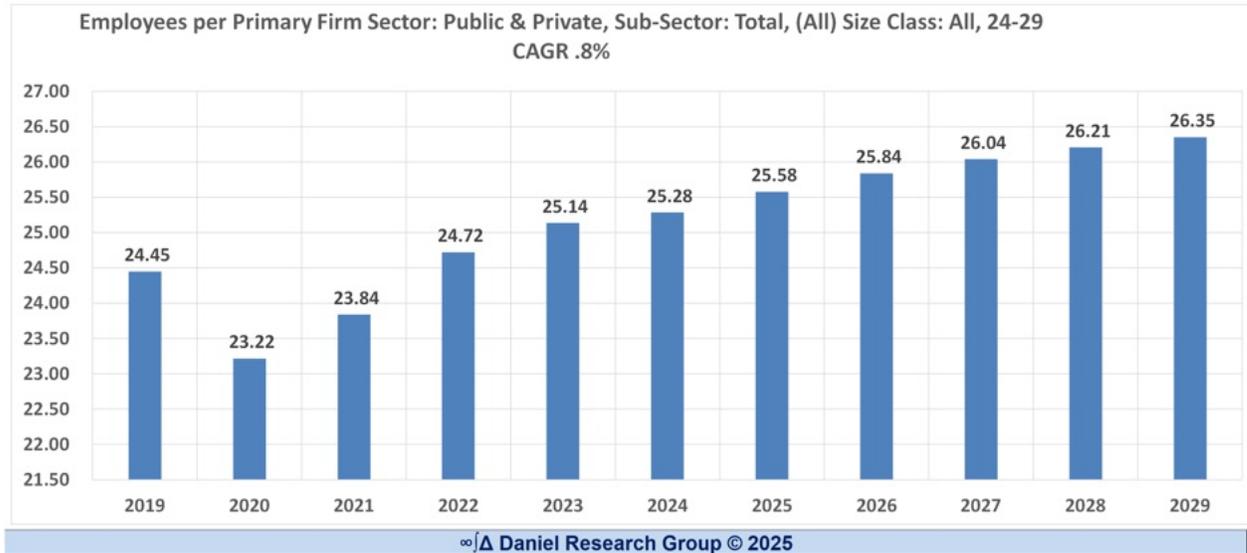
© Daniel Research Group © (2025)

Less than 100	-95,396	-0.2%		
100 or More	-2,192	-4.5%		

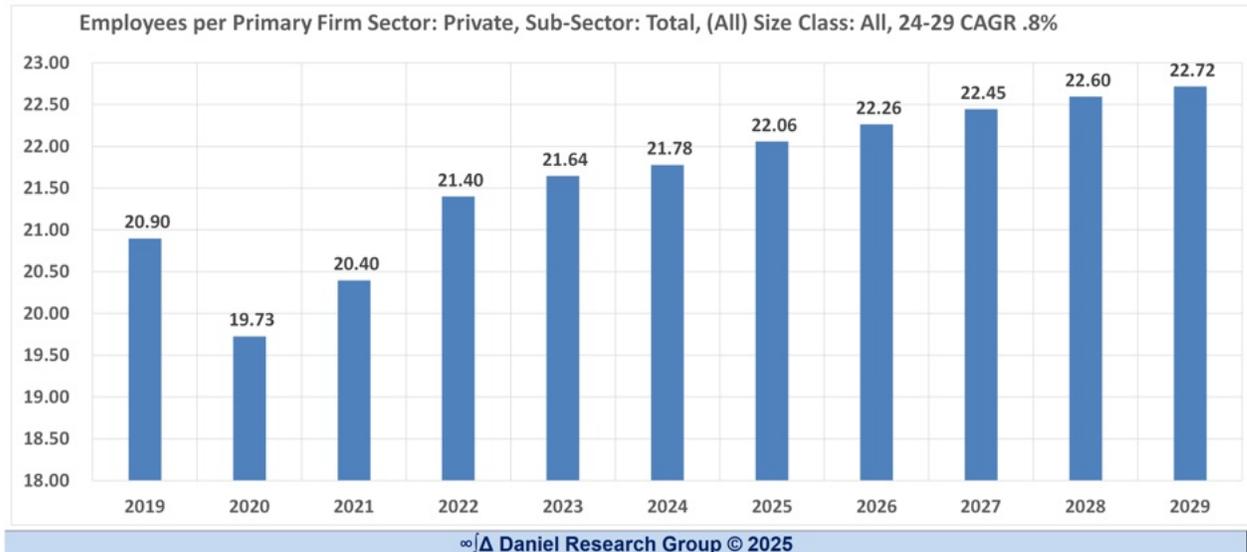
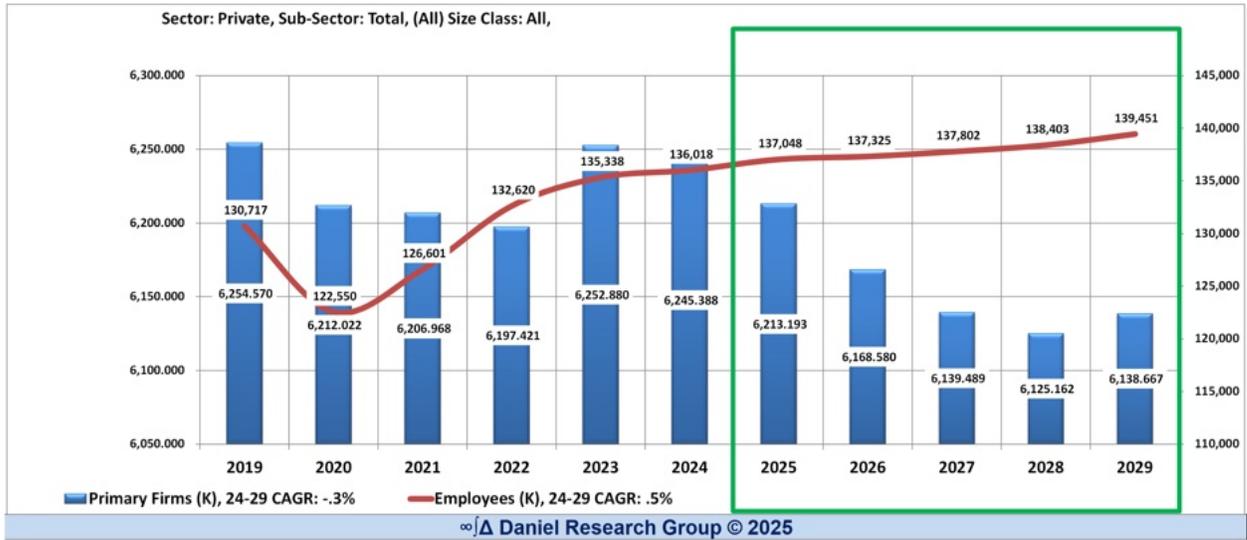
© Daniel Research Group © (2025)

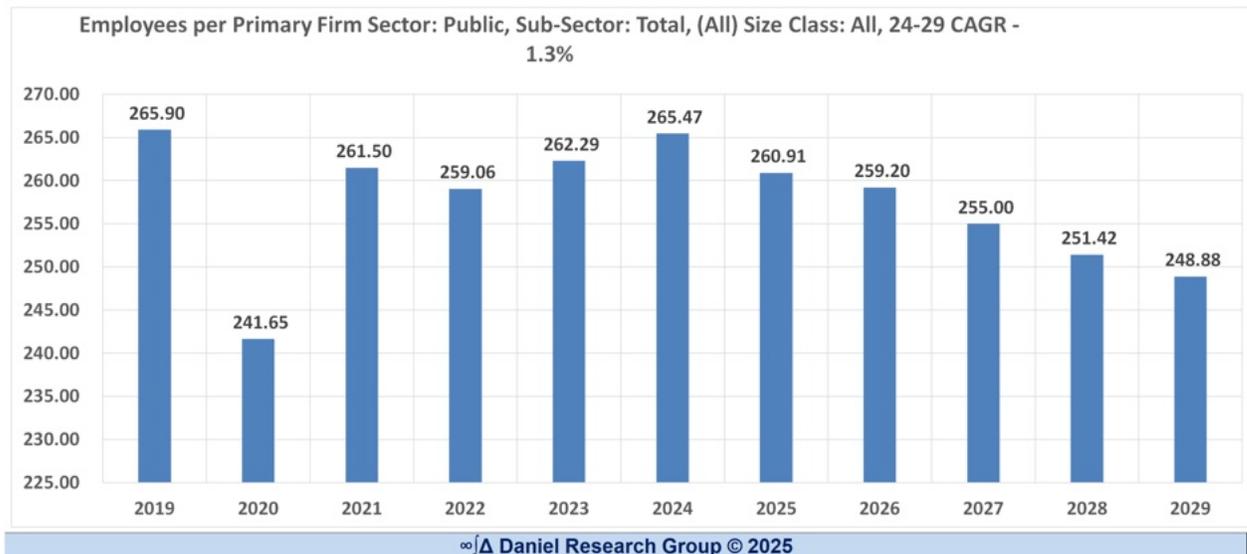
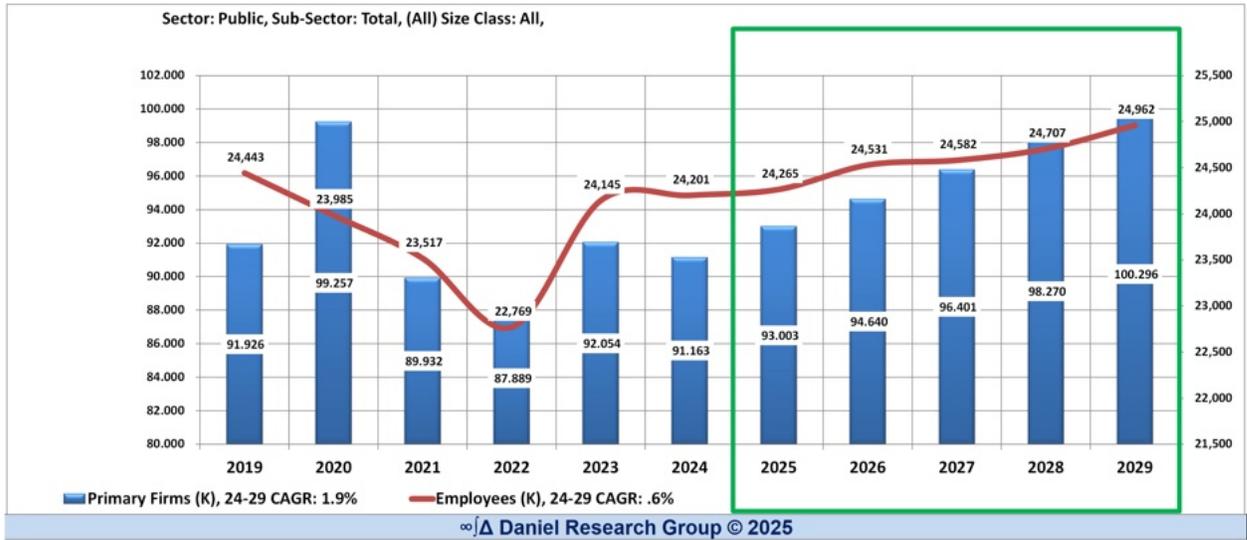
98.1% of the loss in Business with less than 100 Employees

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



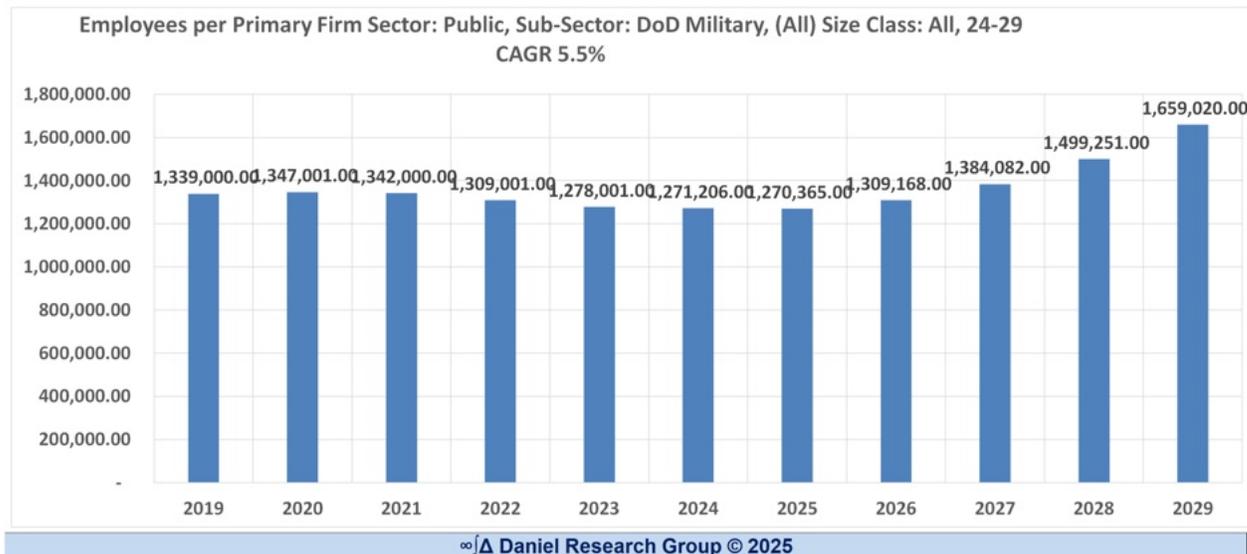
By Sector



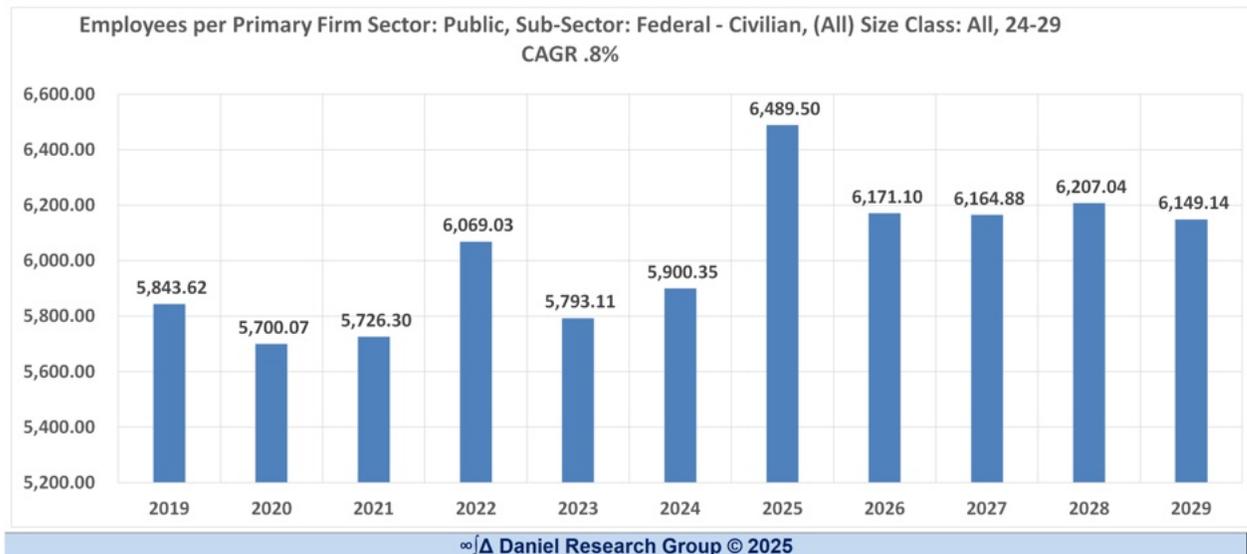
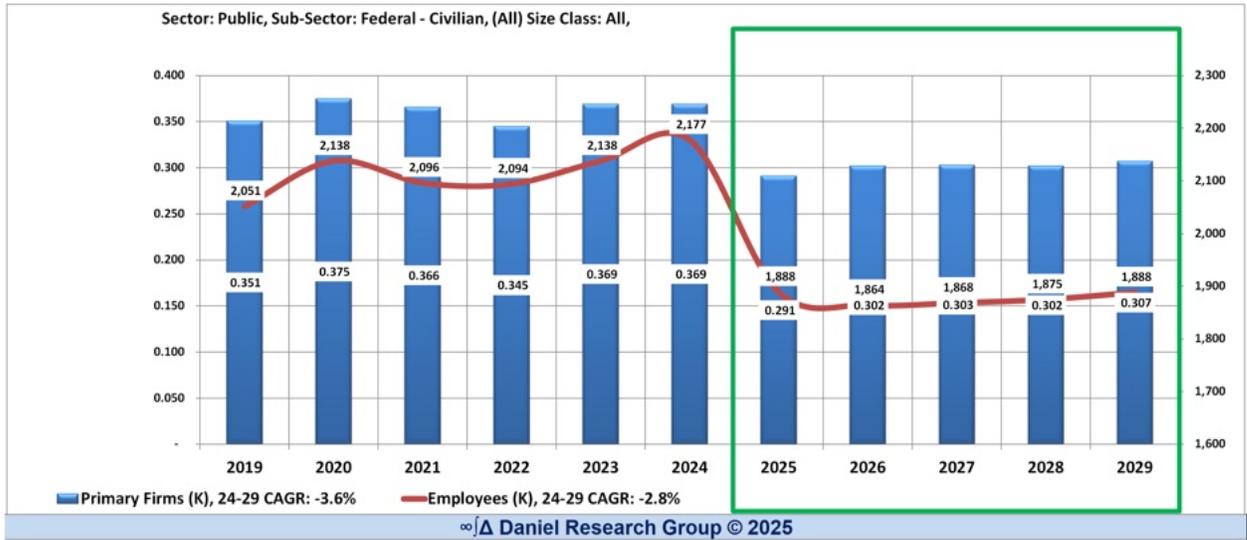


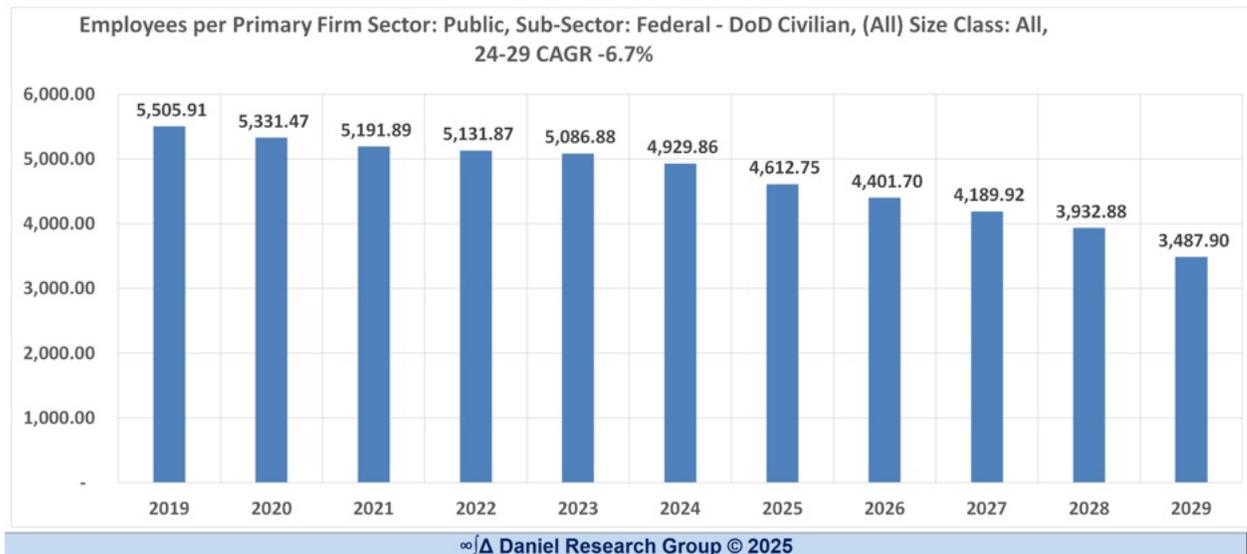
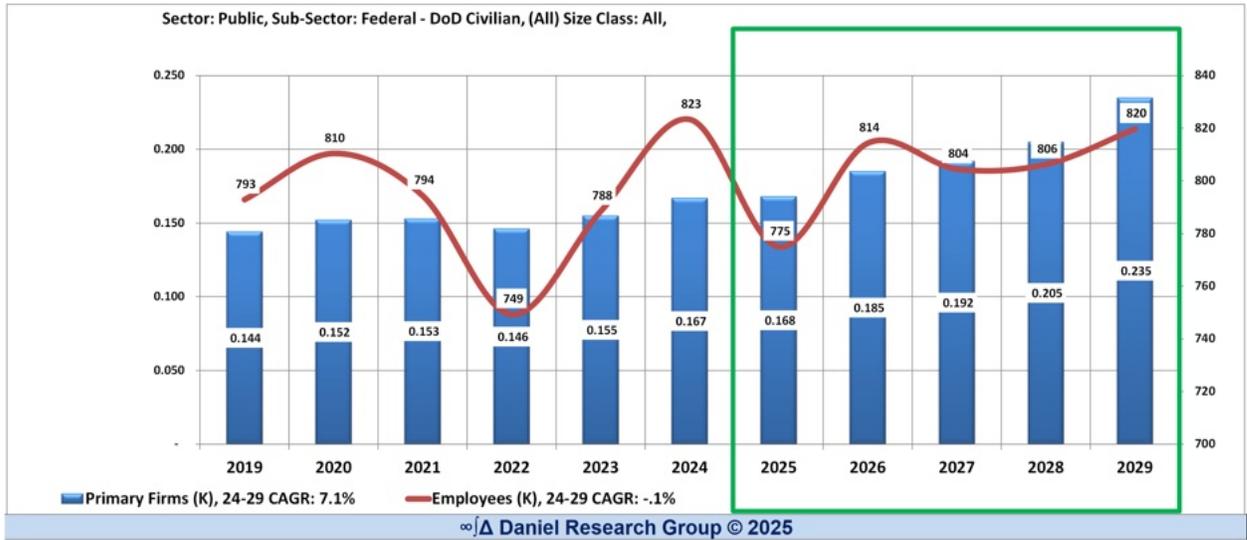
Primary Firms for Governments are buying entities such as Agencies or Departments.

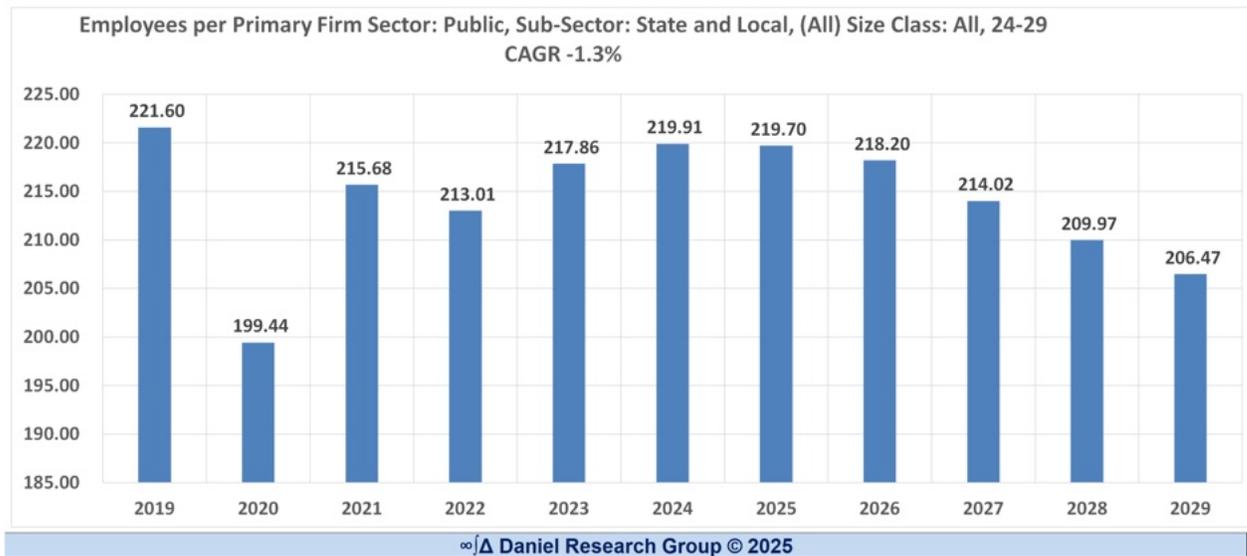
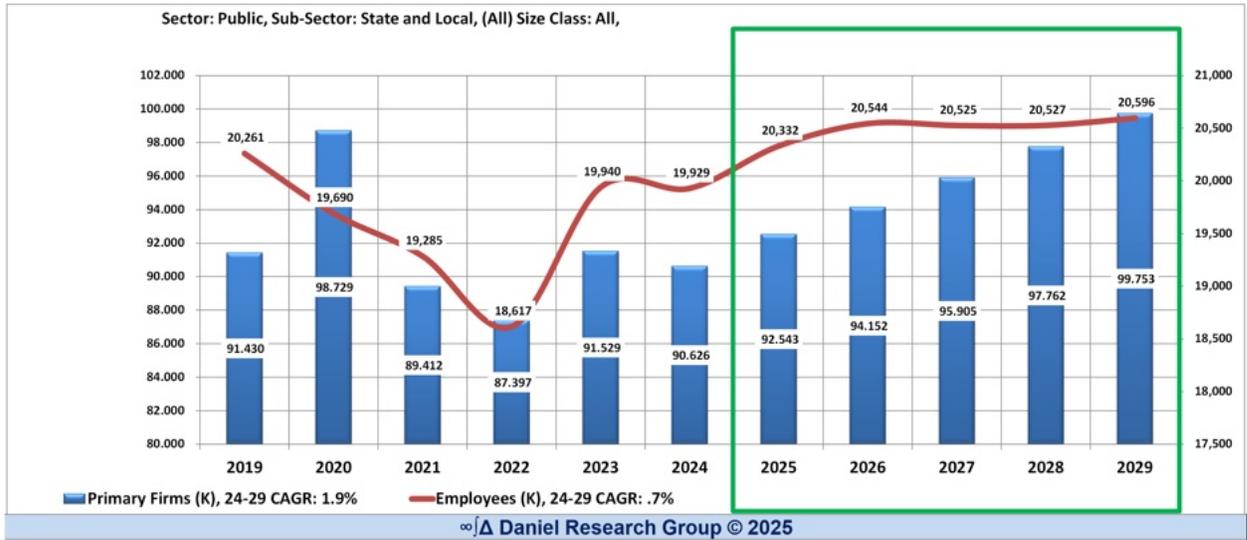
By Sub-Sector – Public Sector



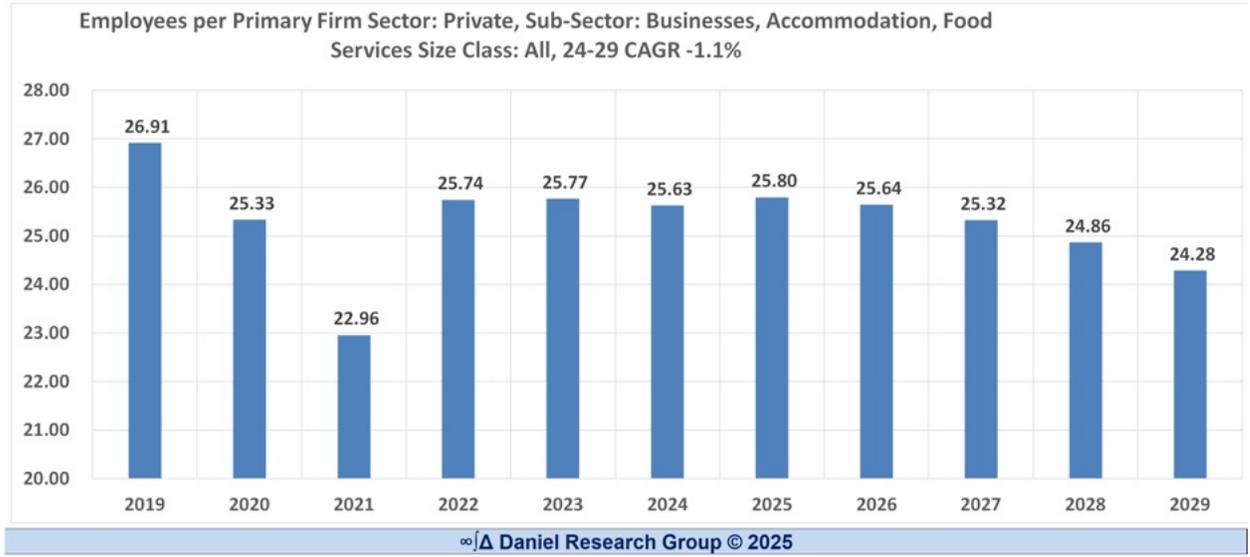
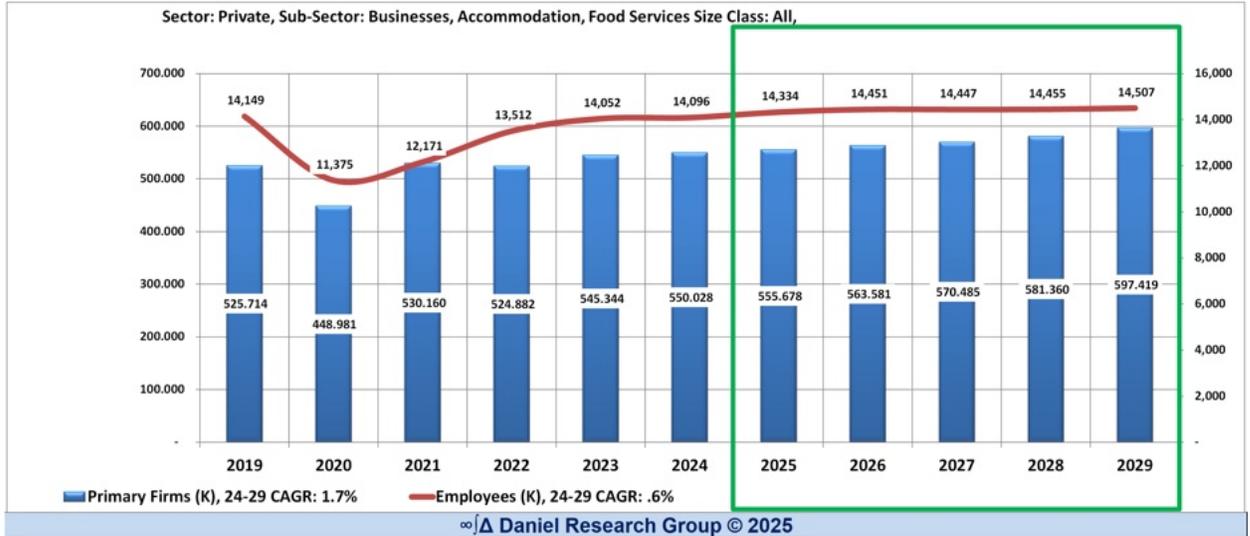
There is only one (1) primary firm for the DoD Sub-Sector, the DoD.

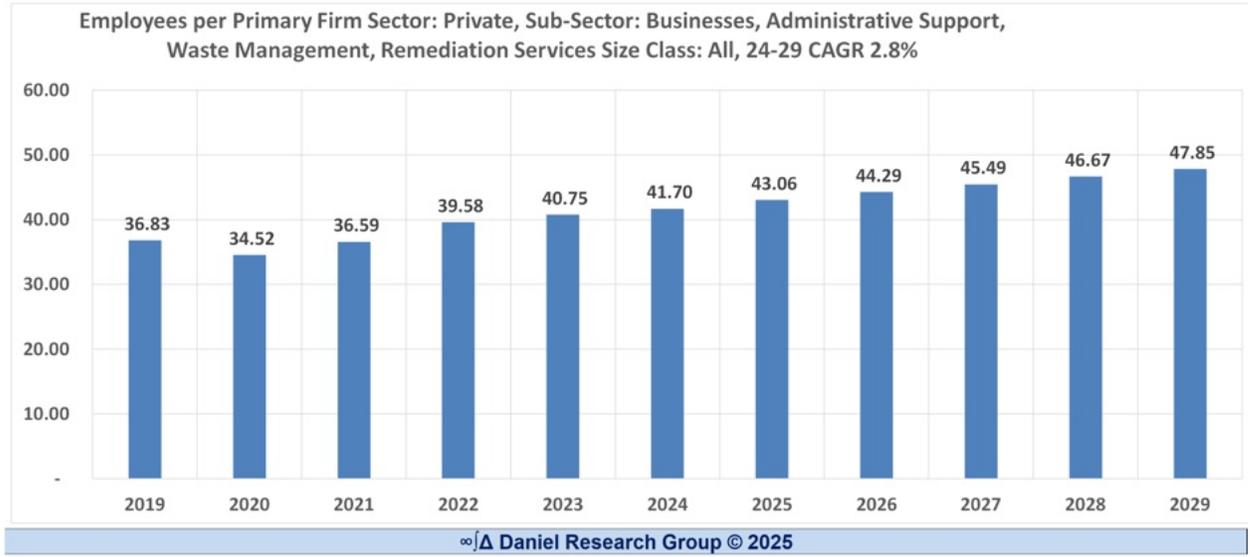


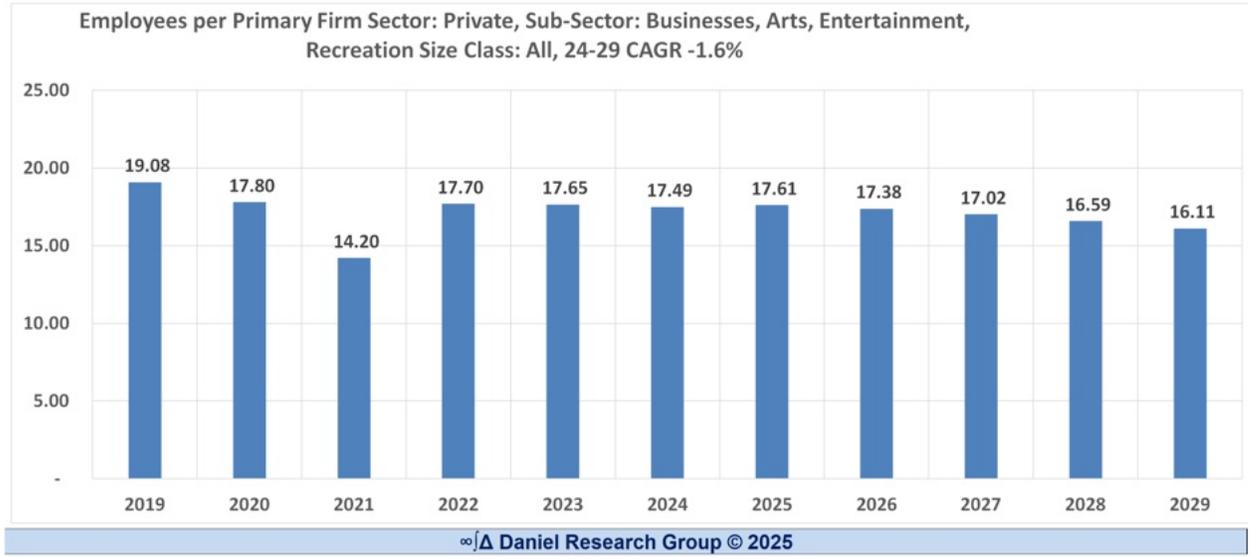
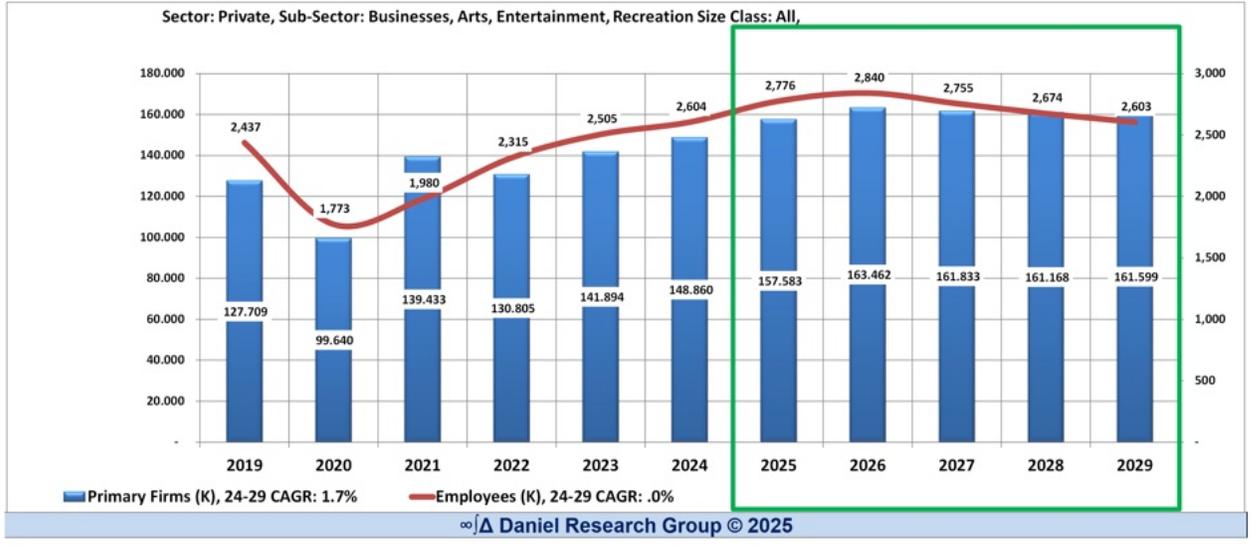


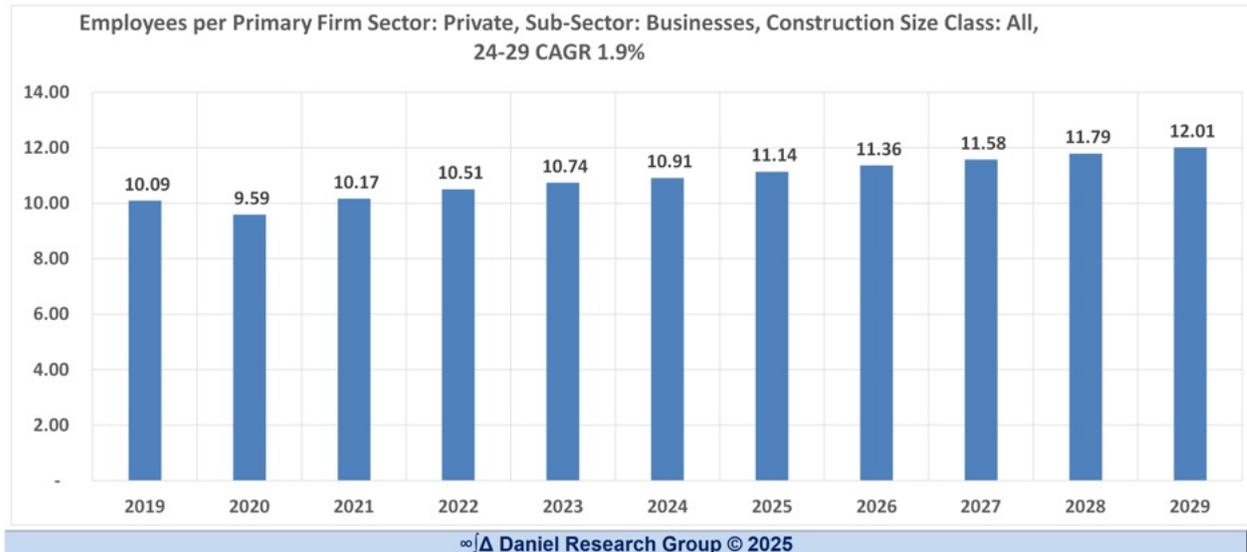


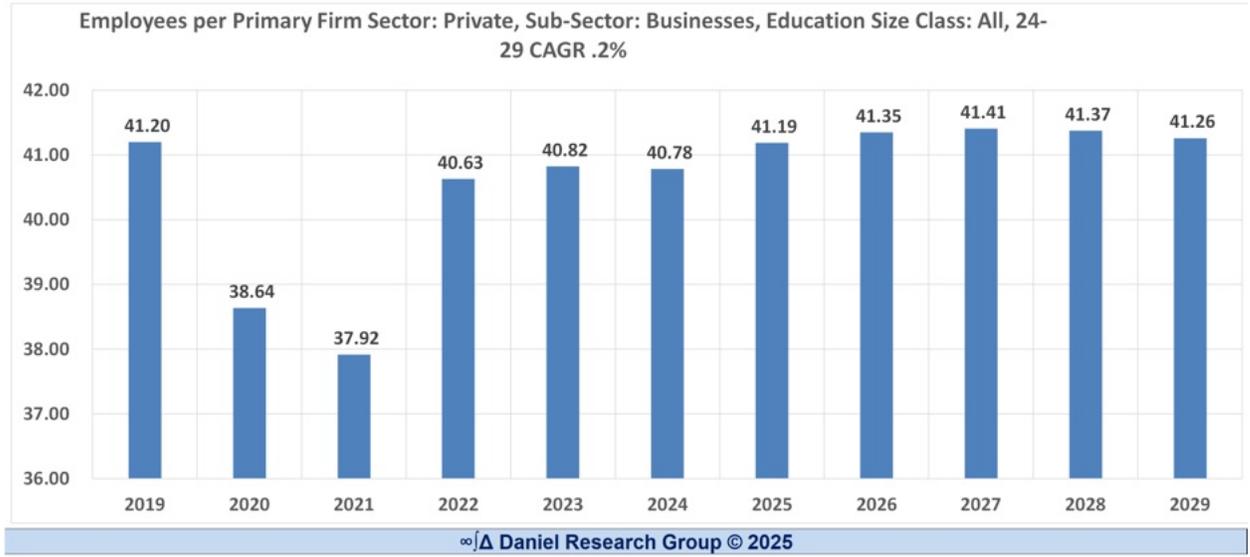
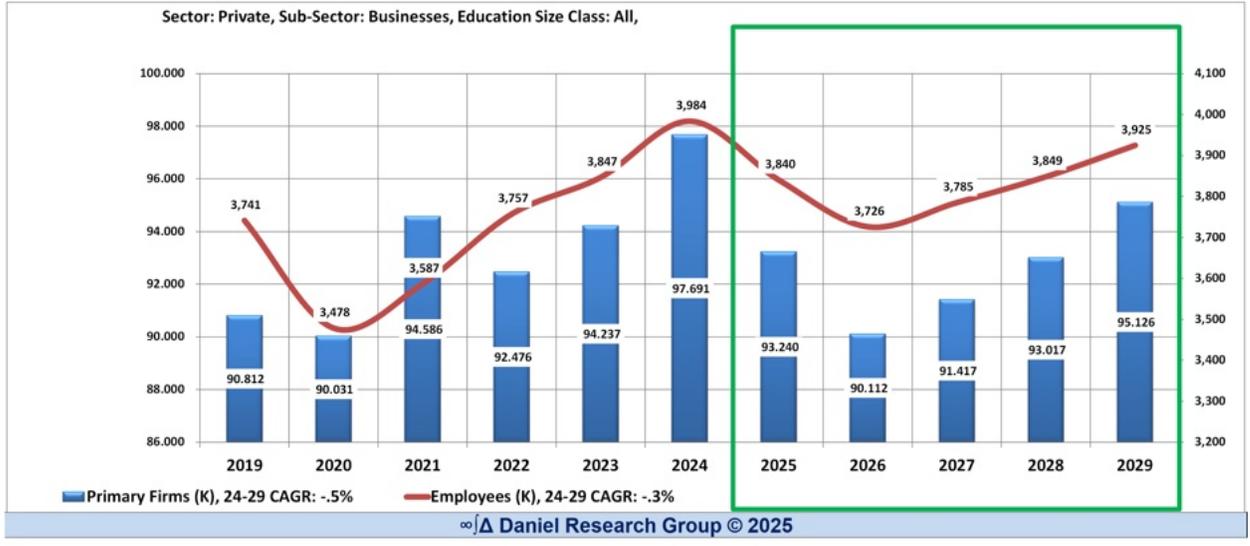
By Industry – Private Sector

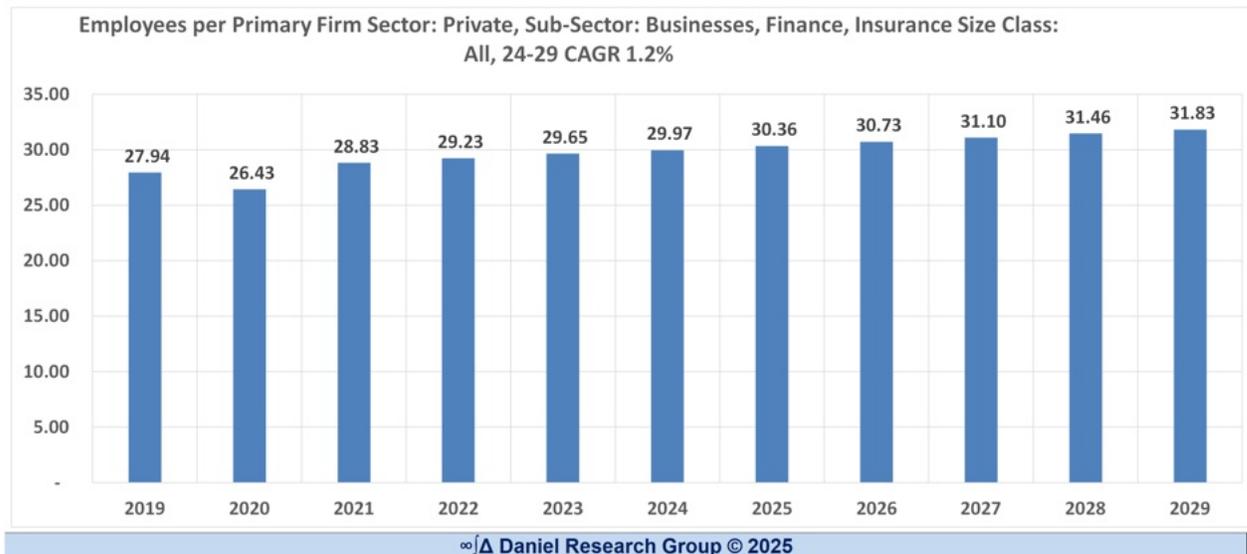
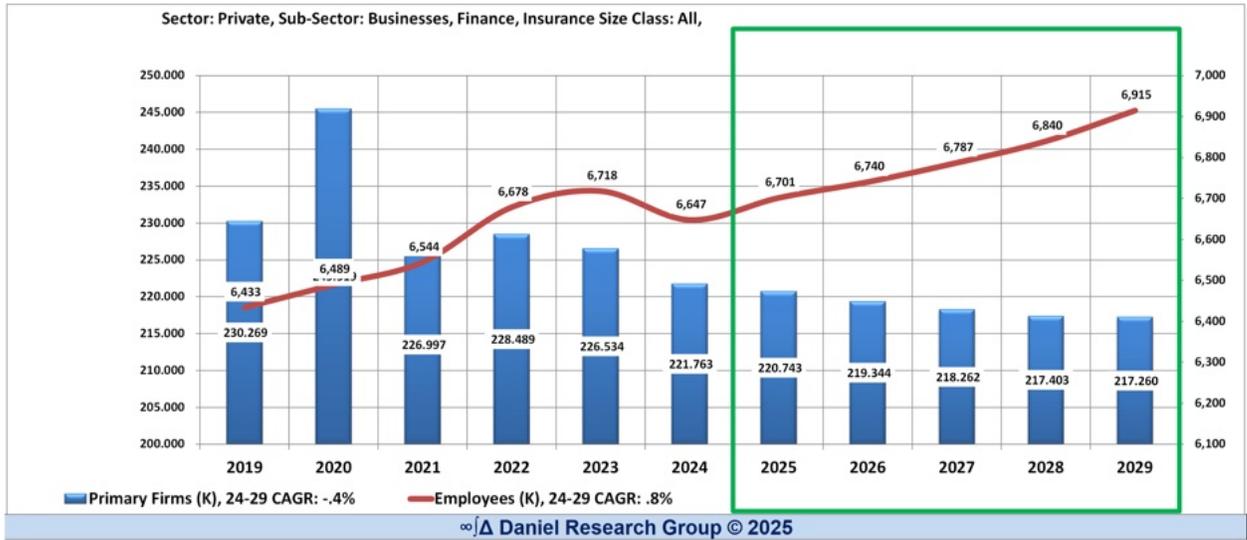


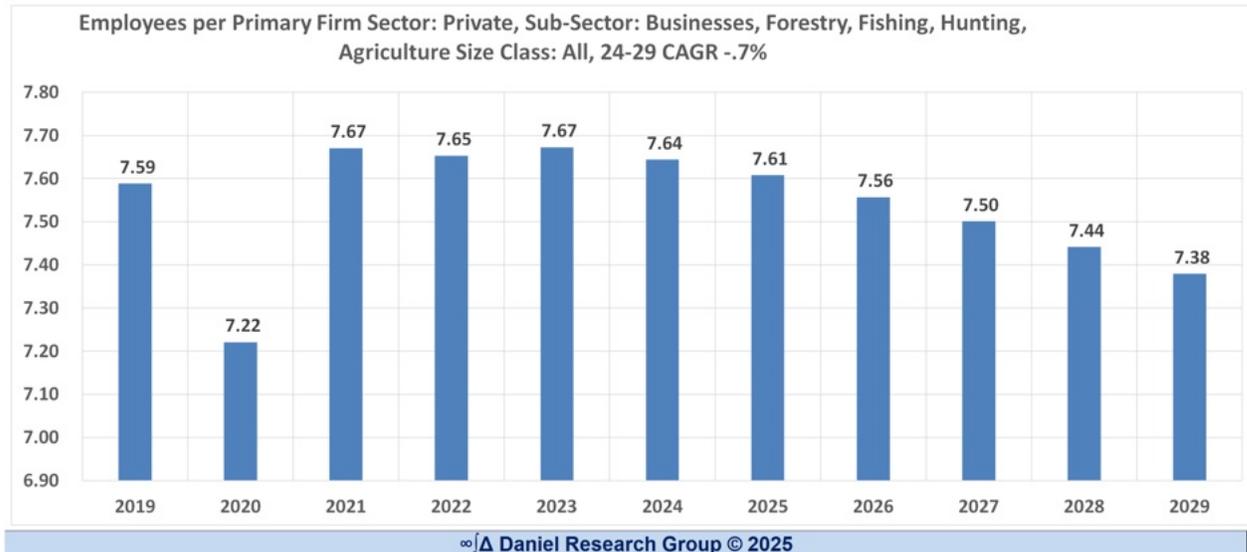
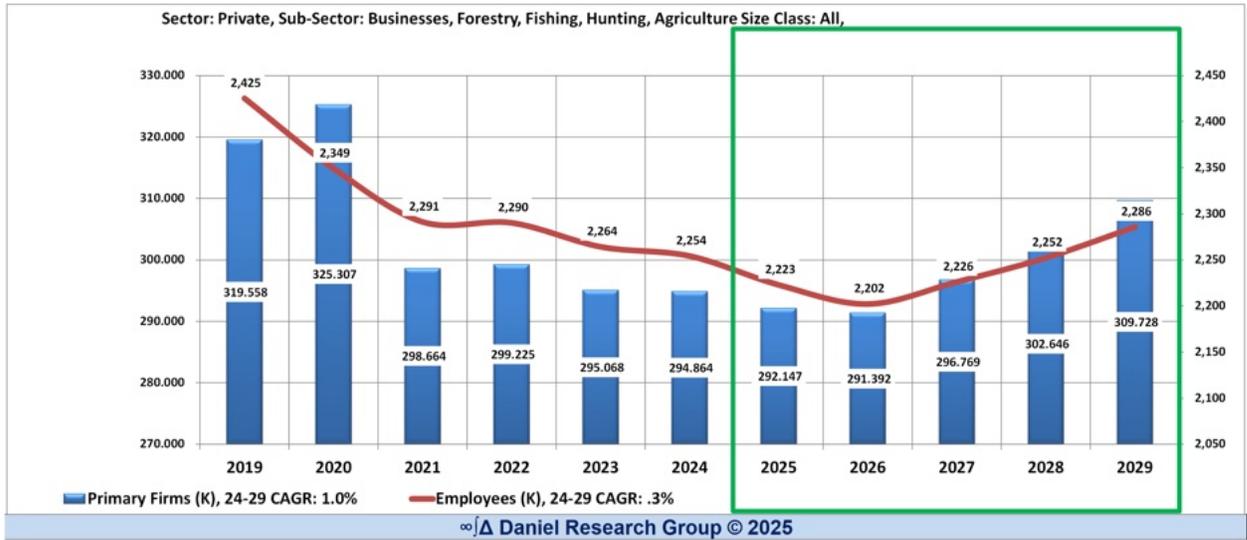


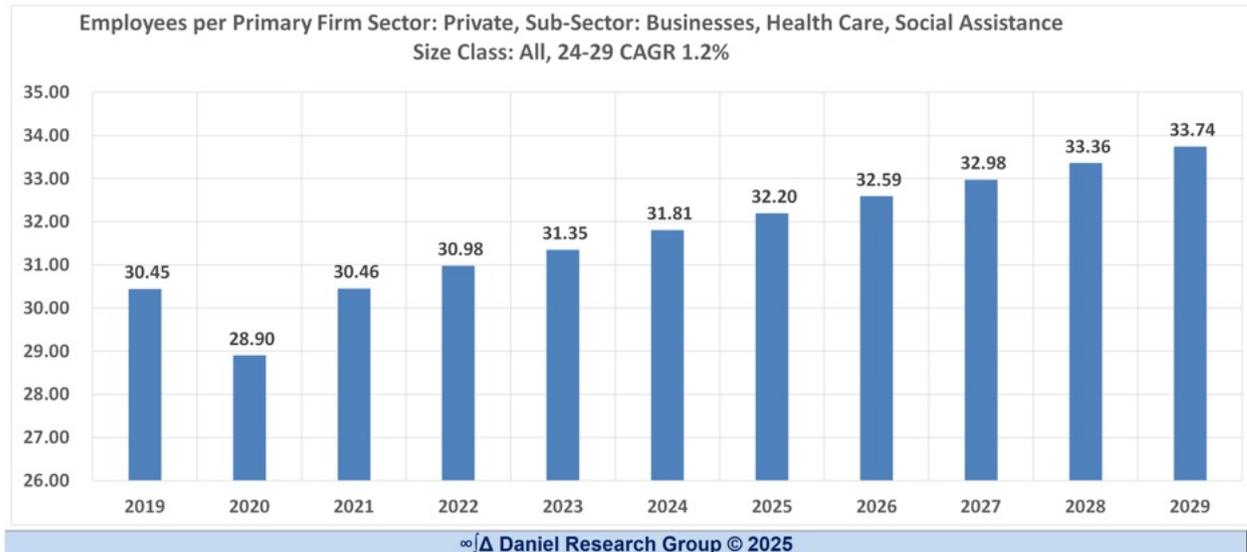
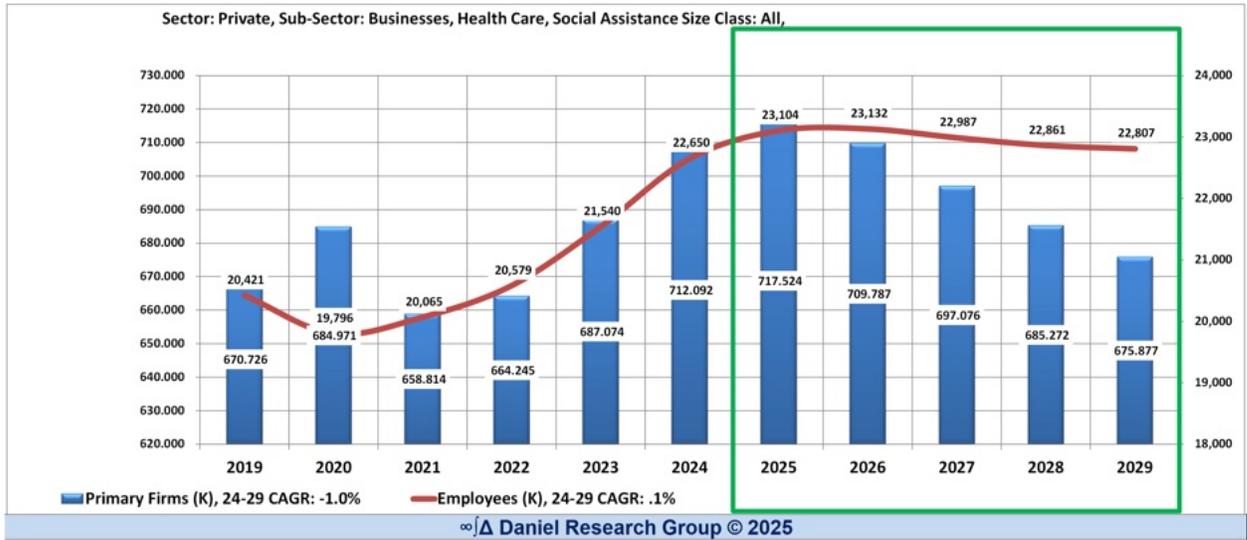


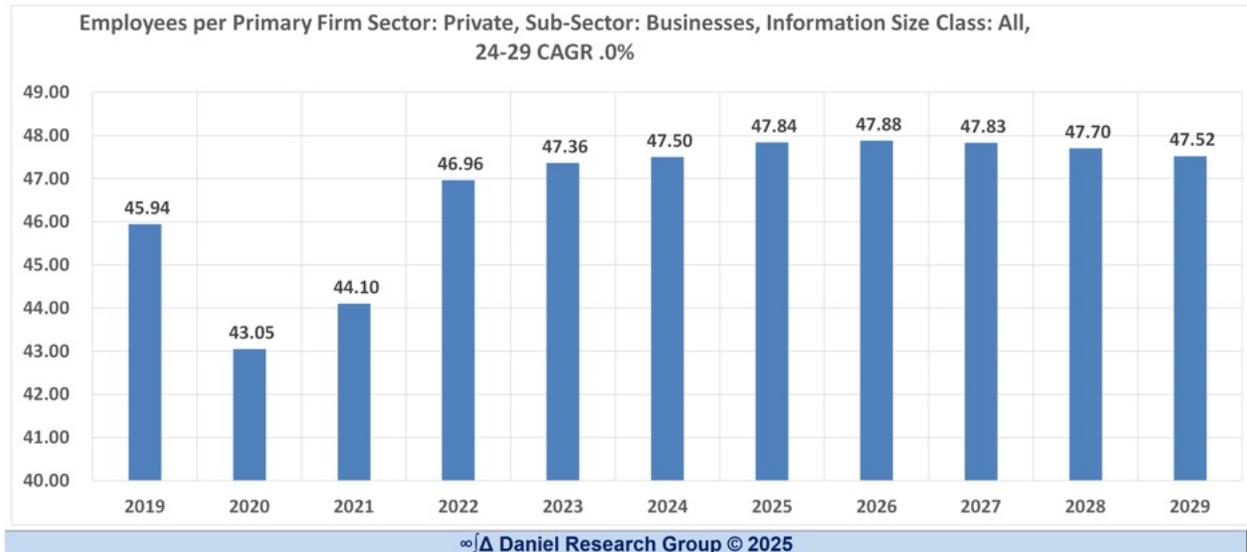
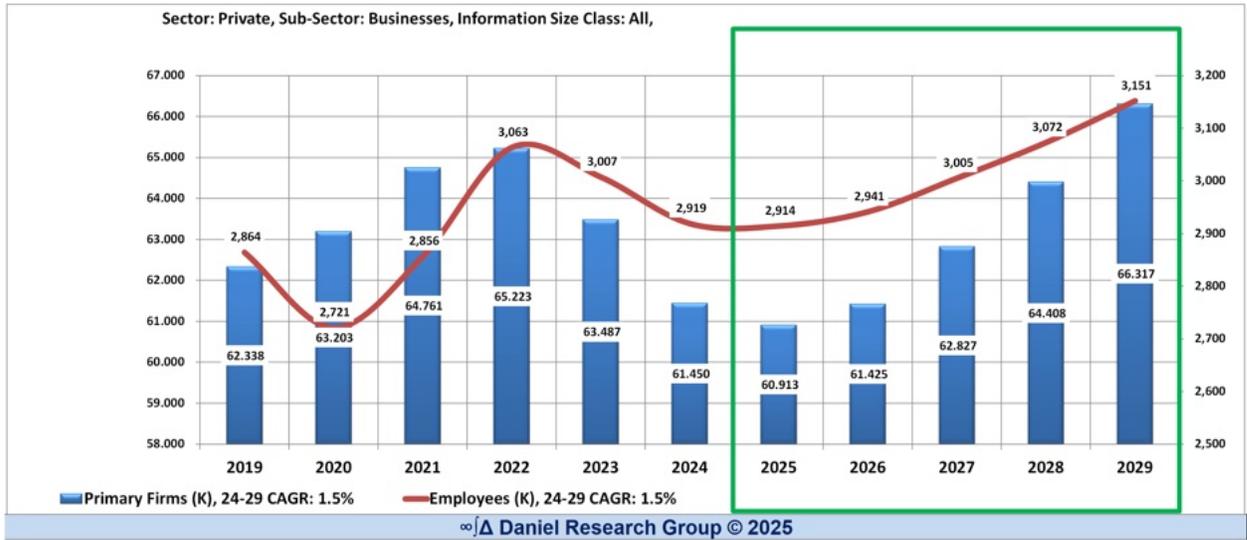


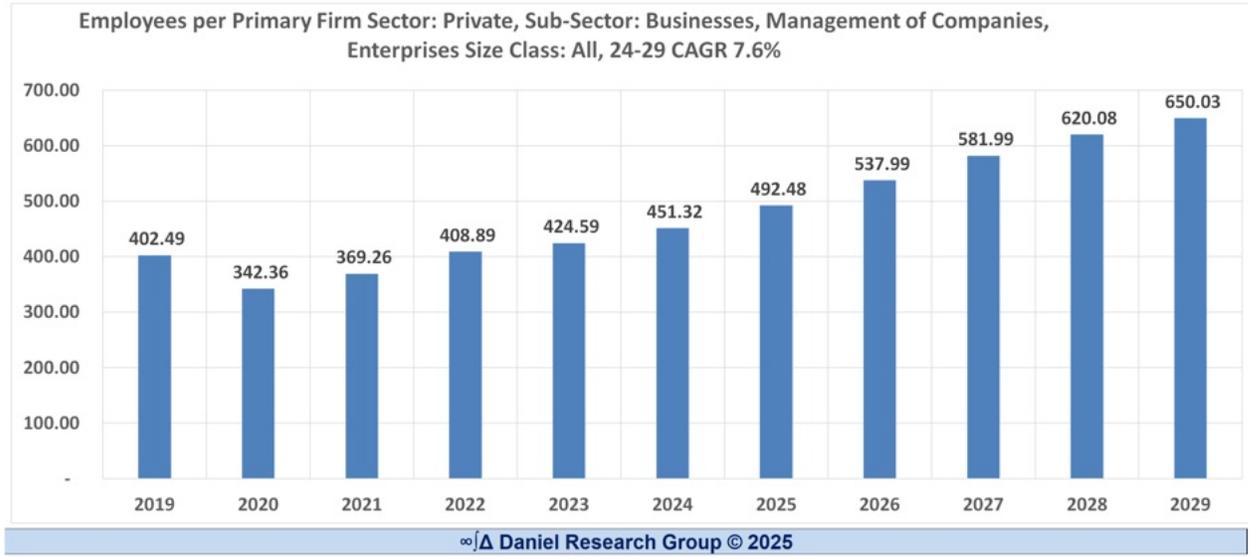
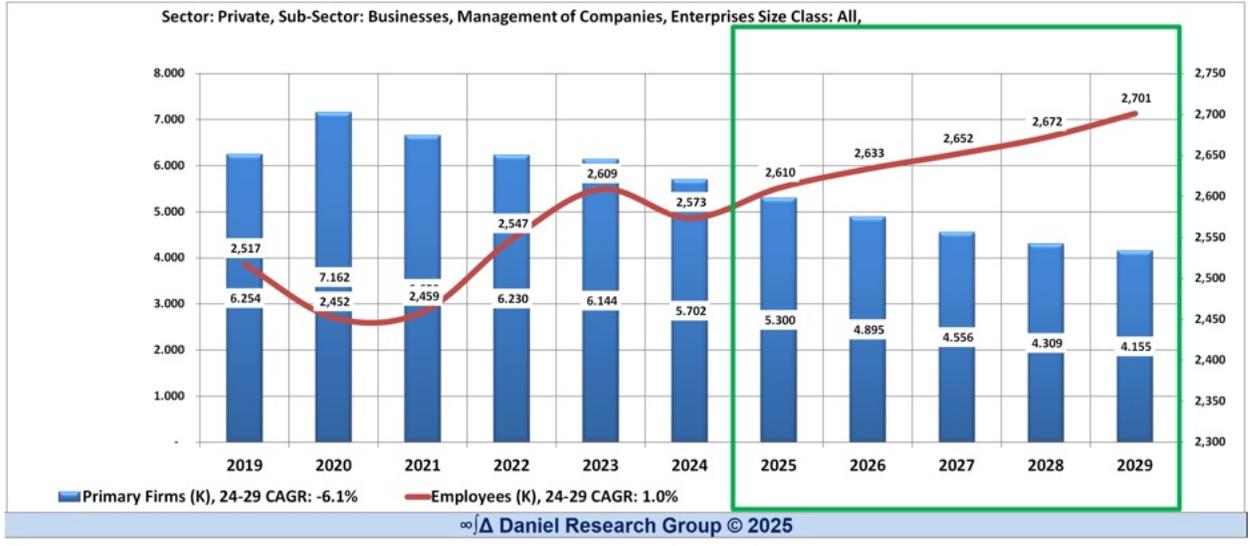


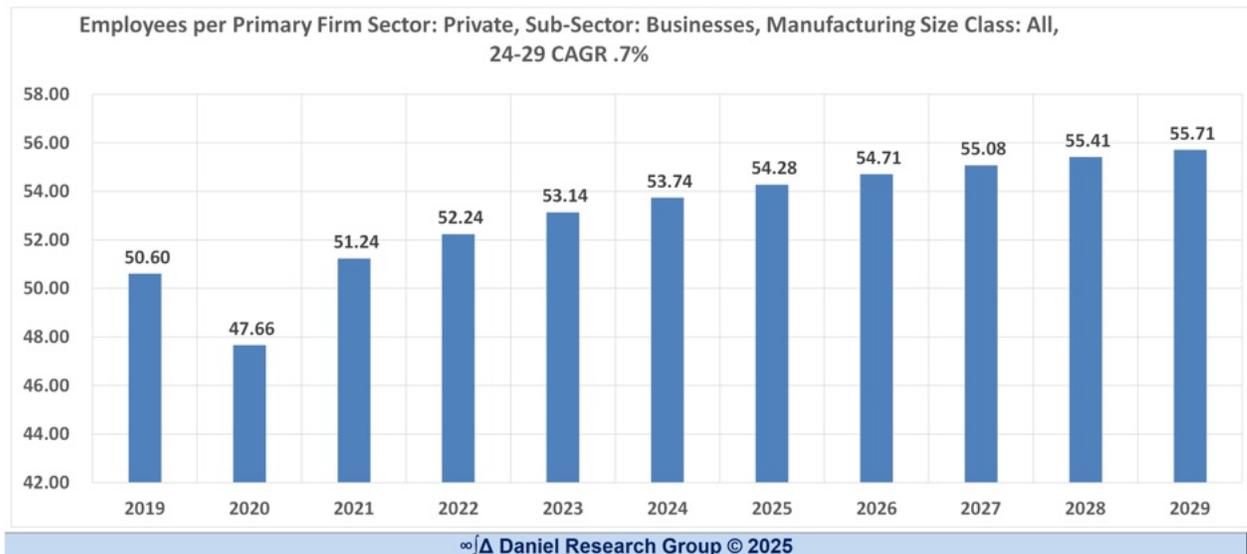
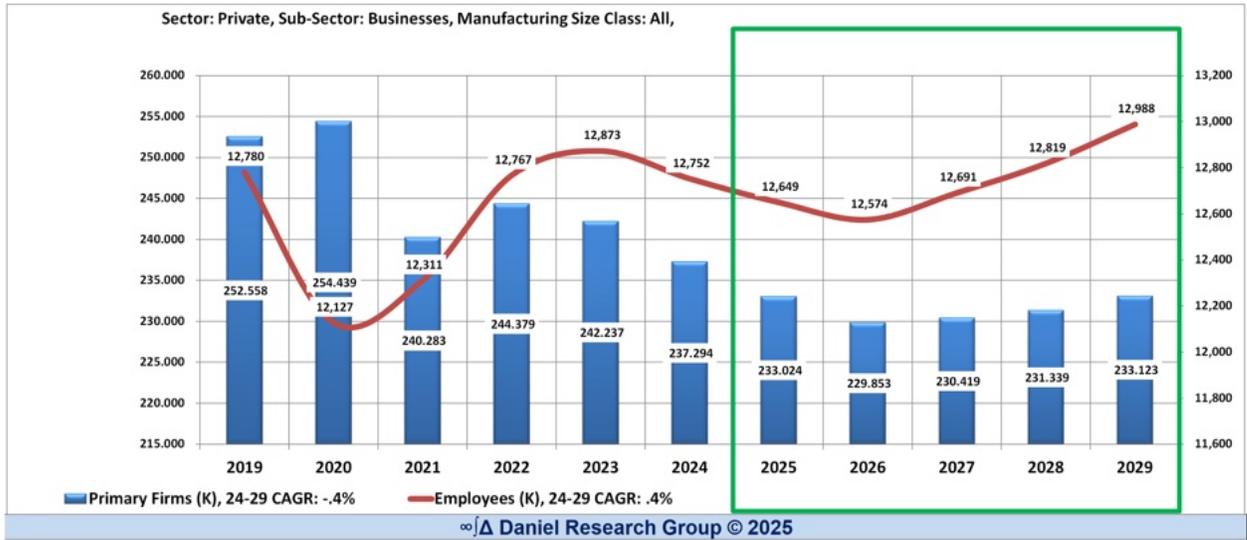


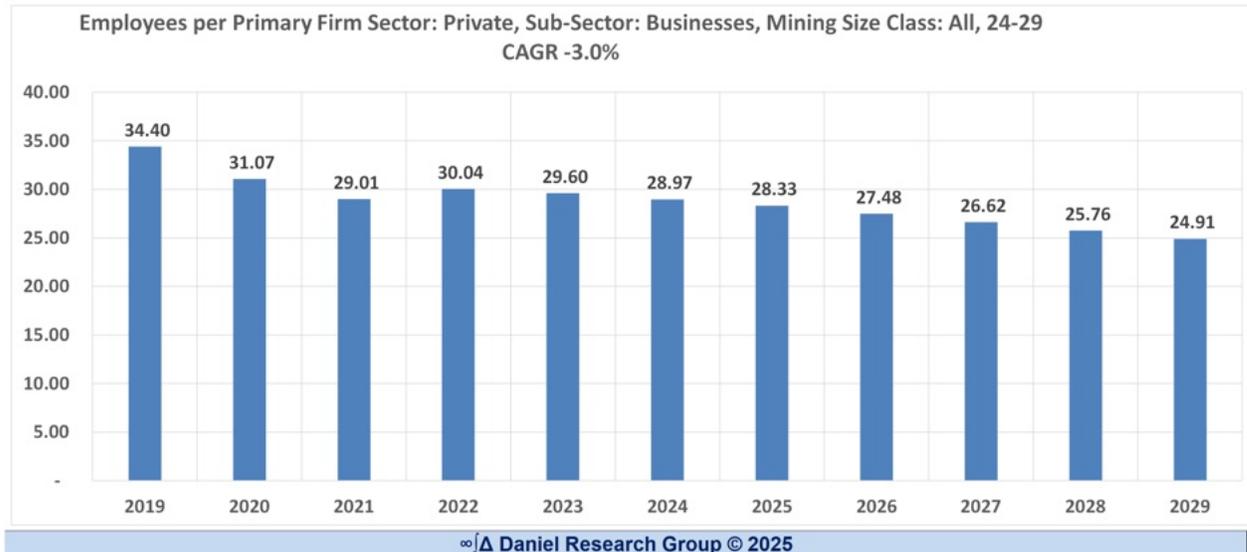
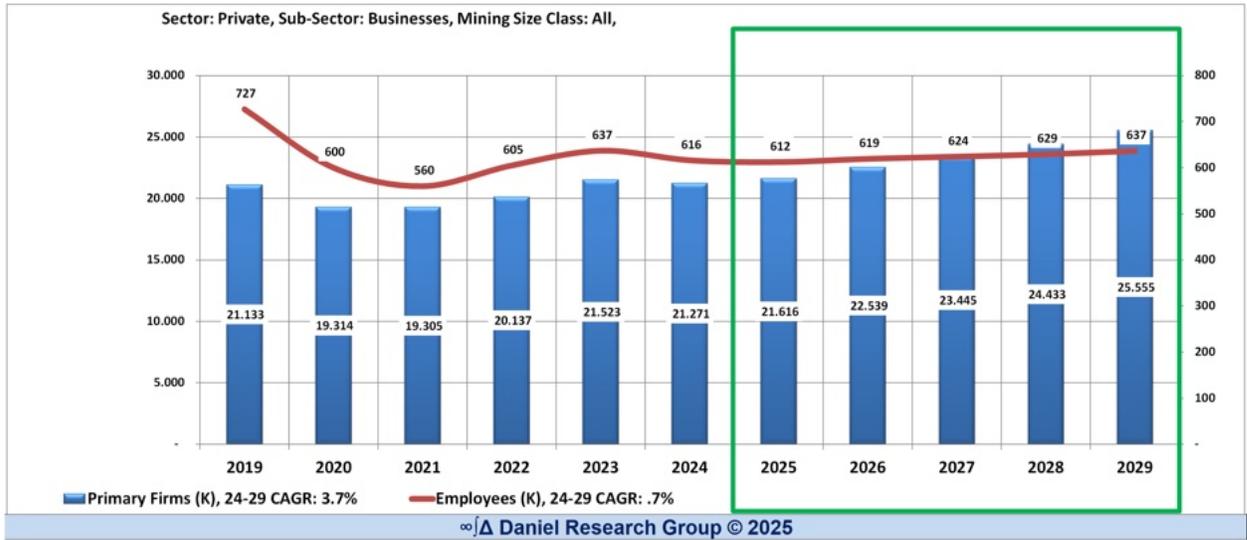


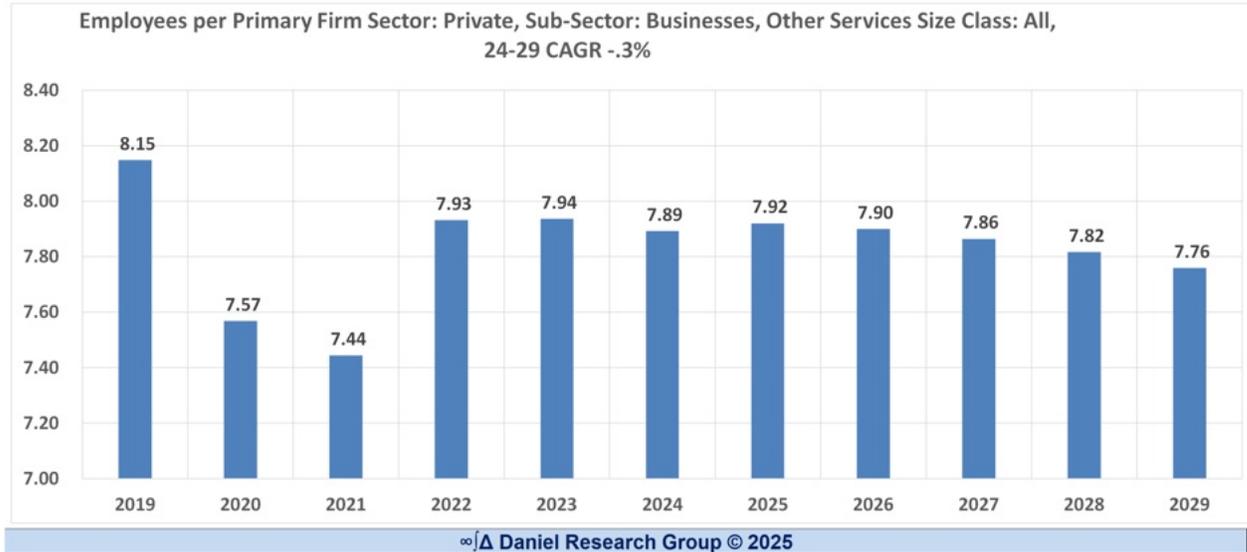
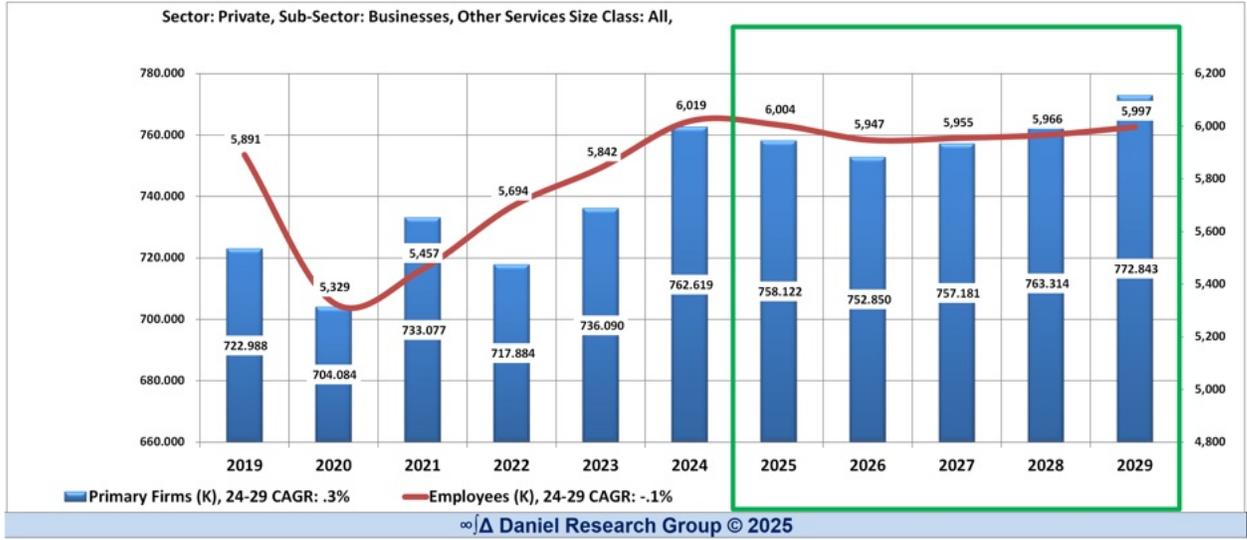


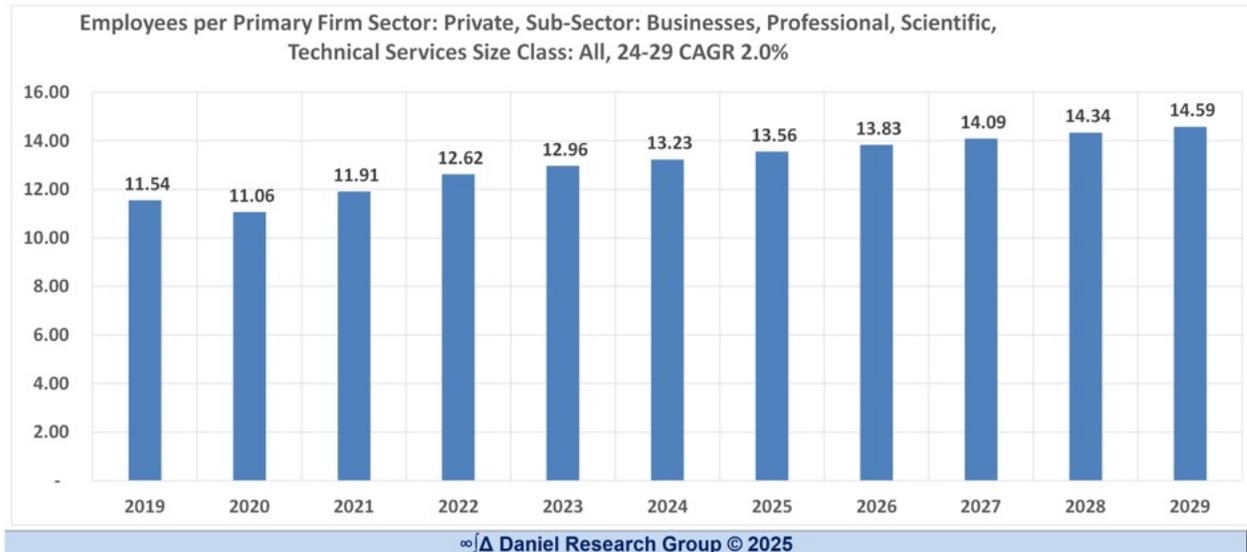
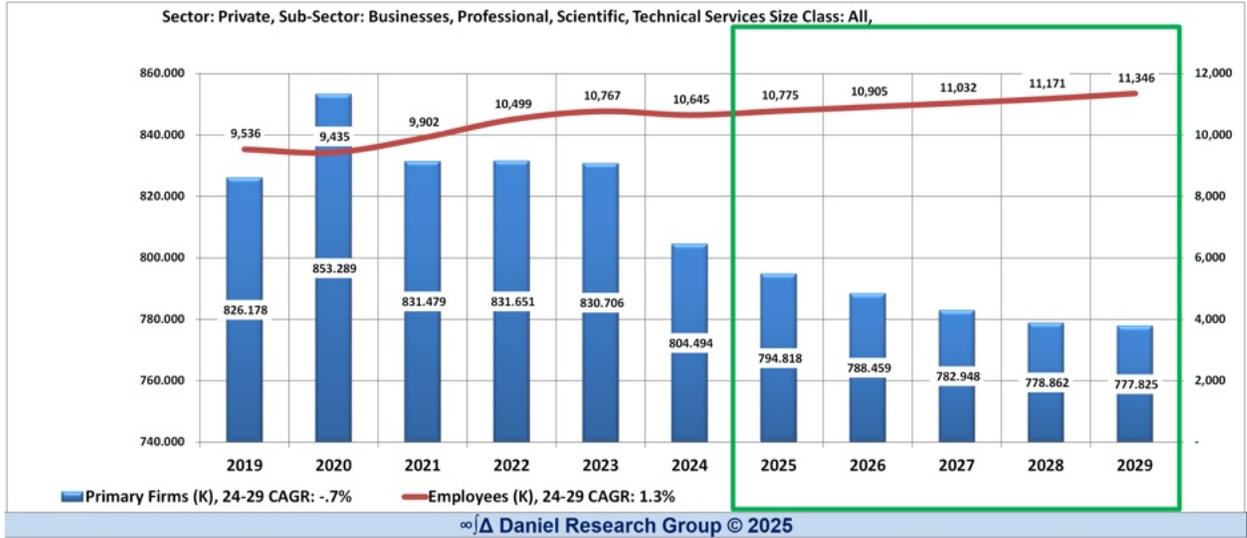


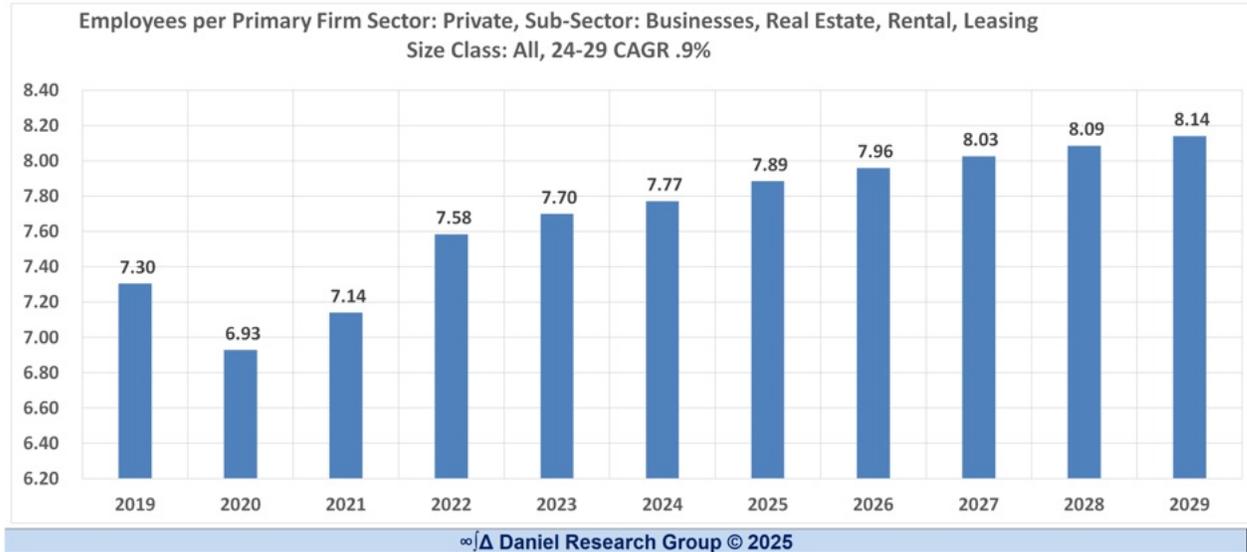


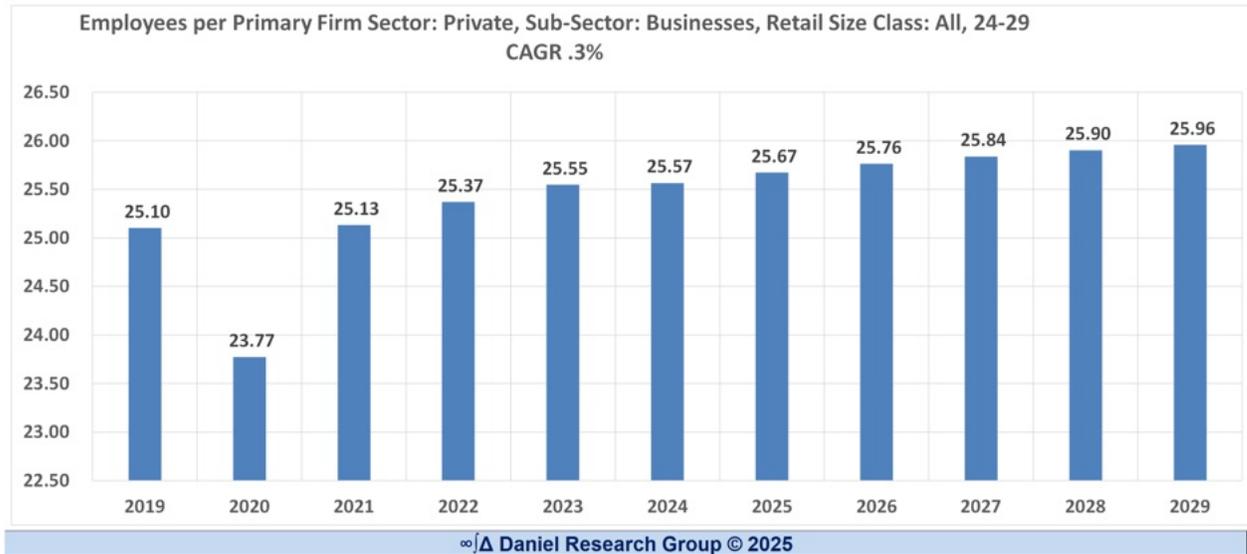
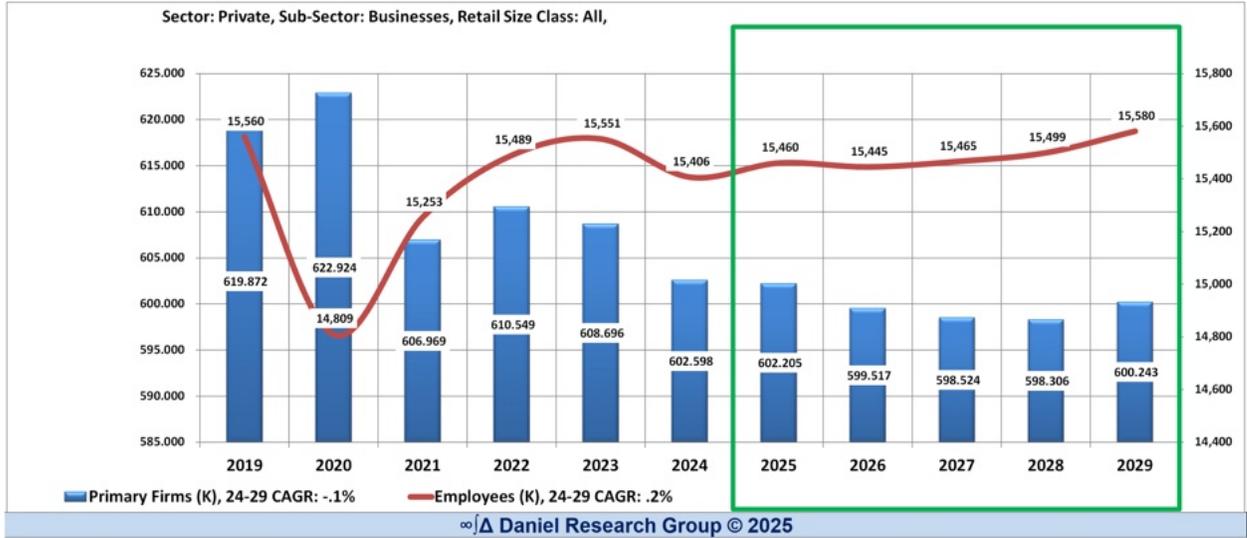


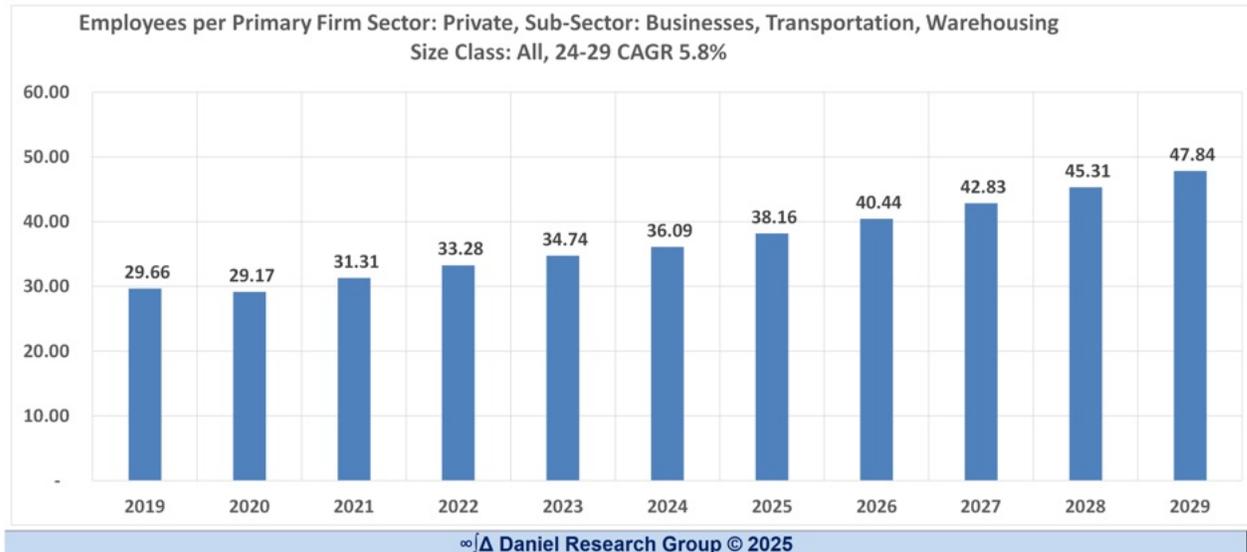
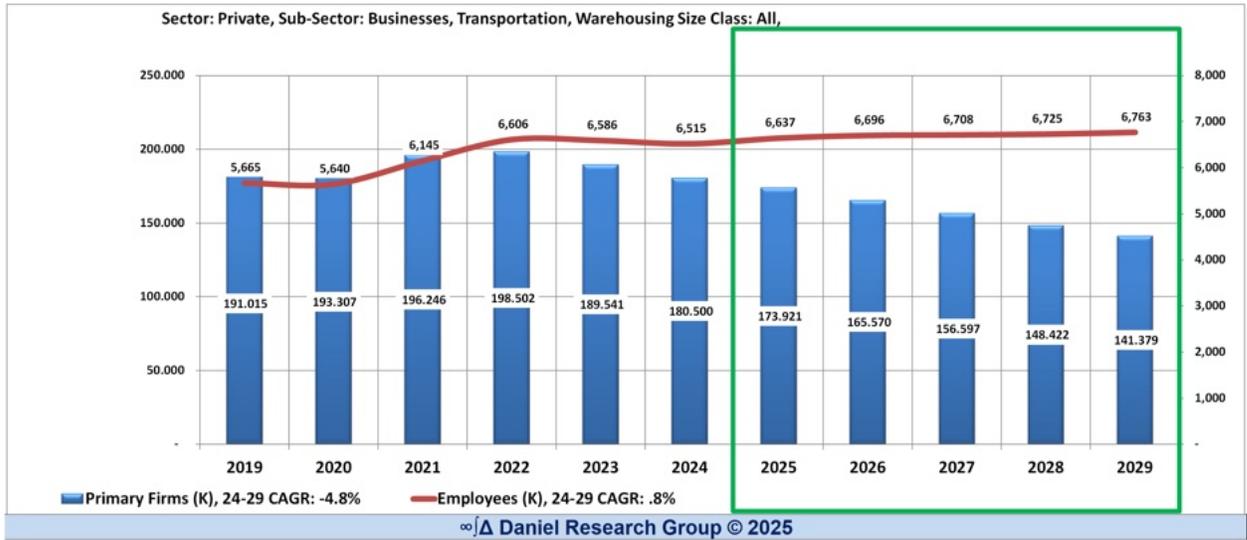


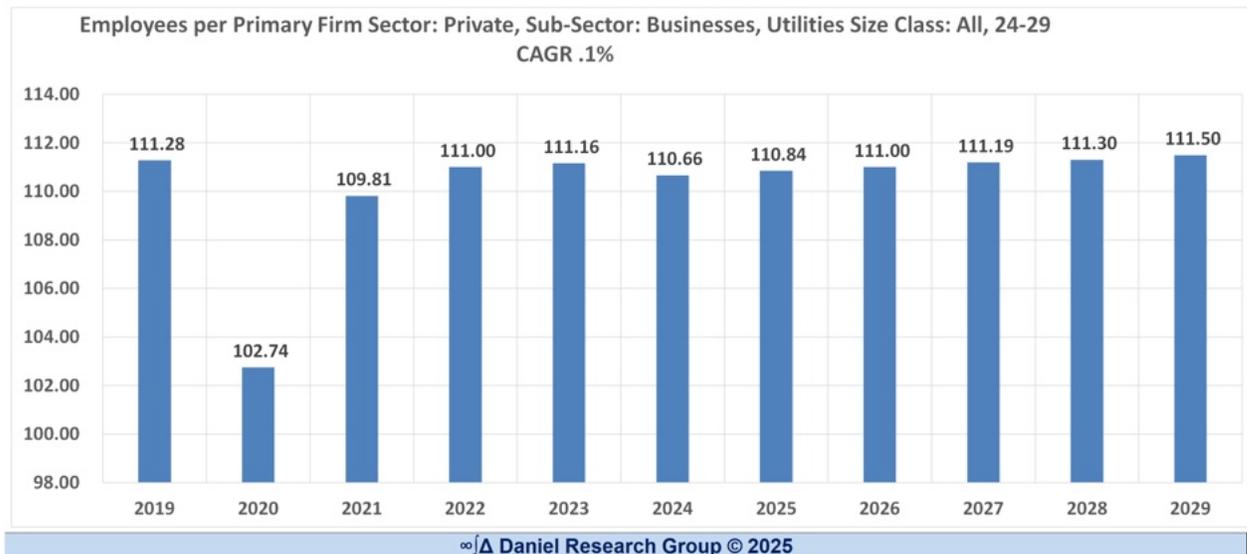
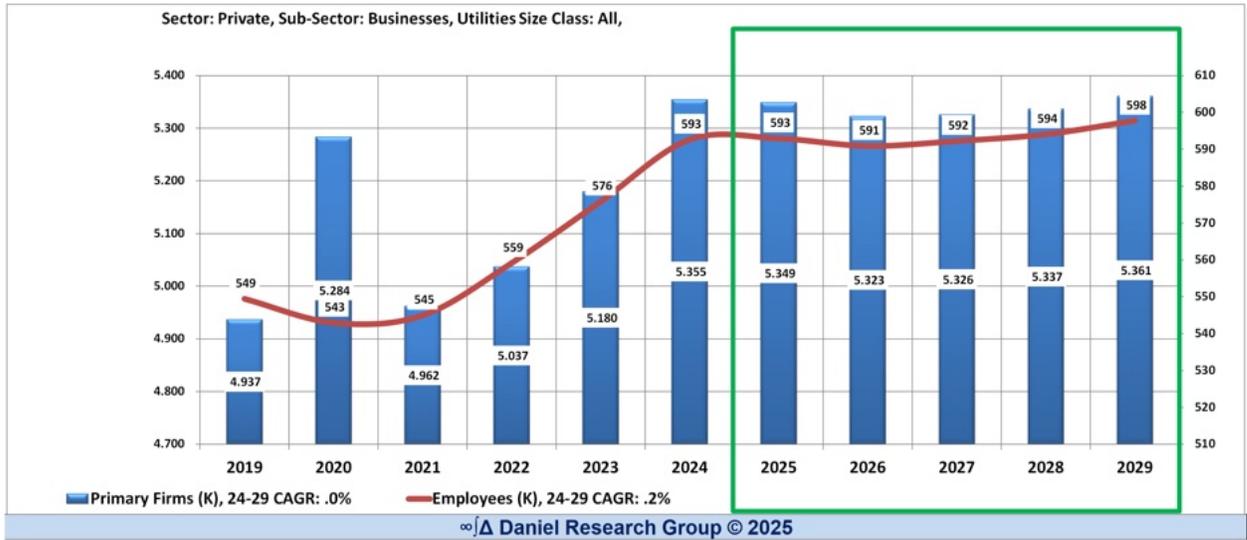


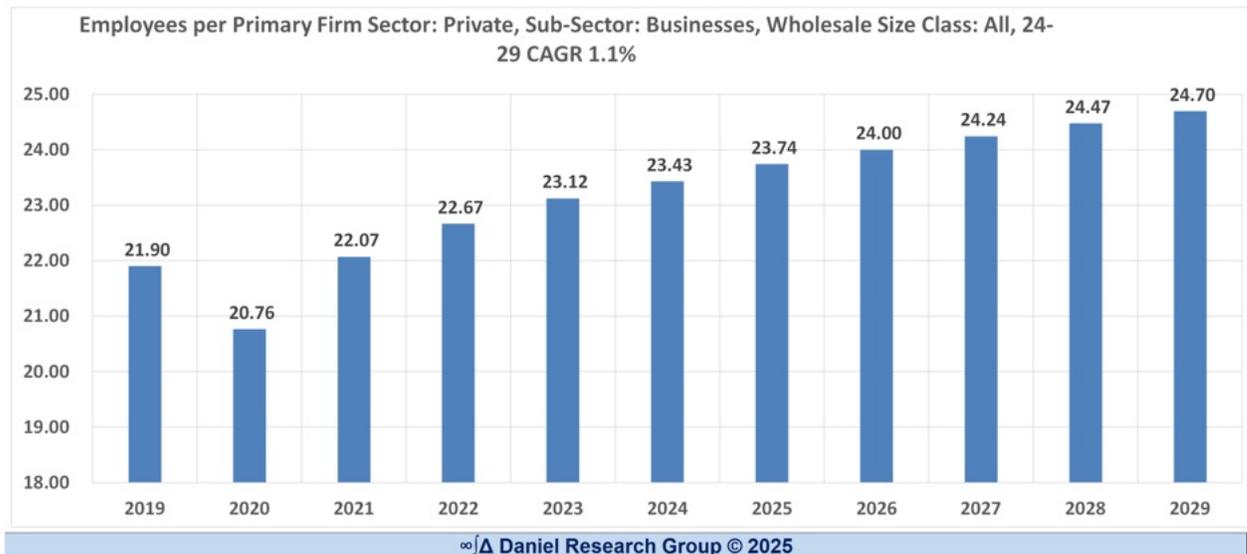
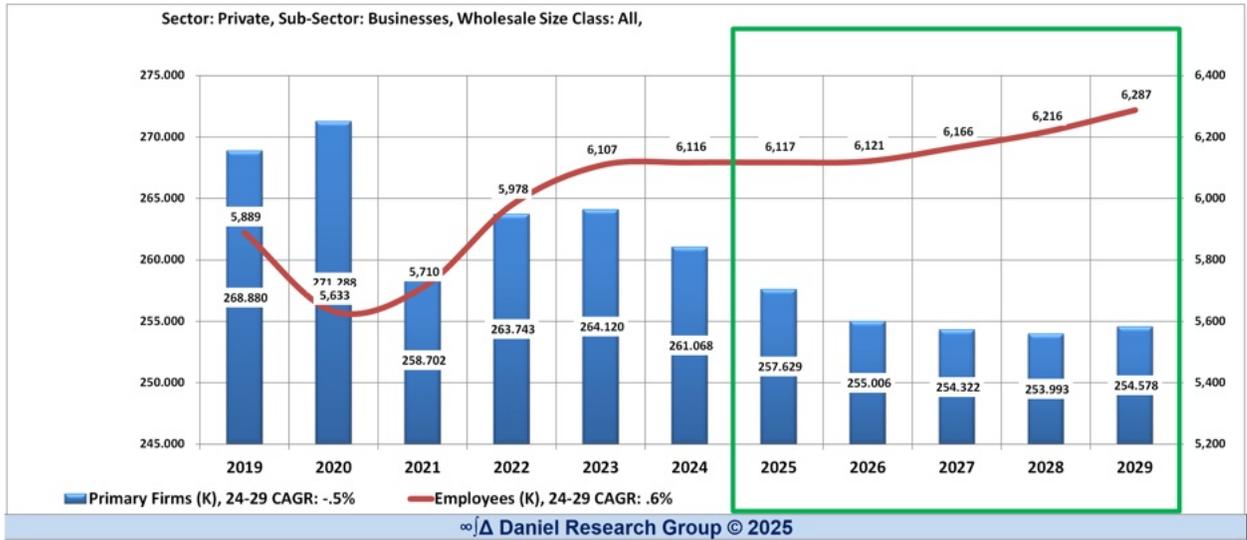




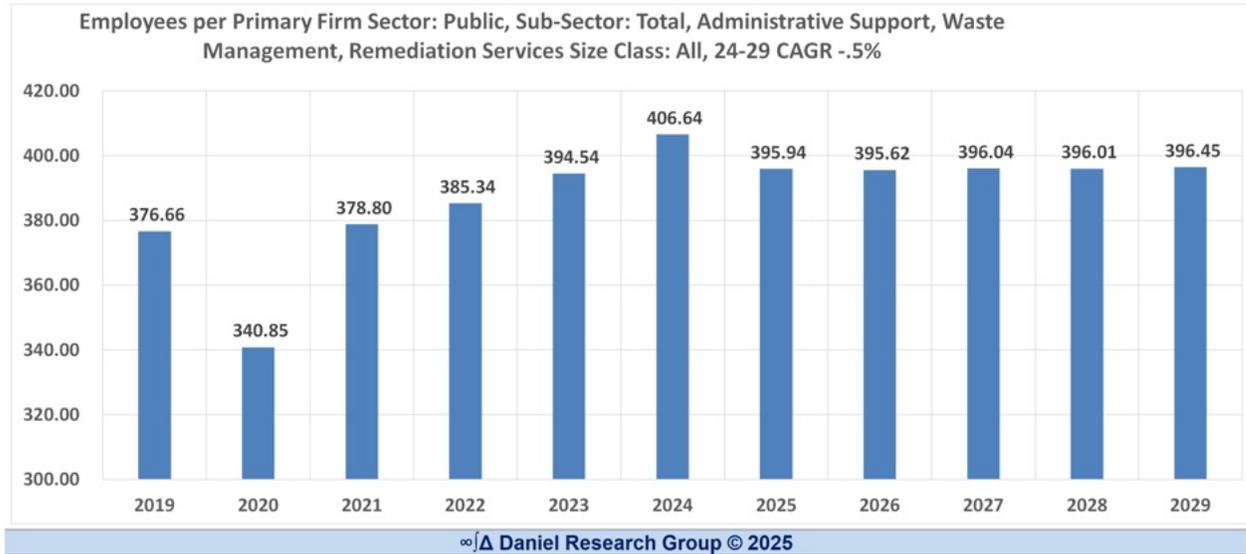
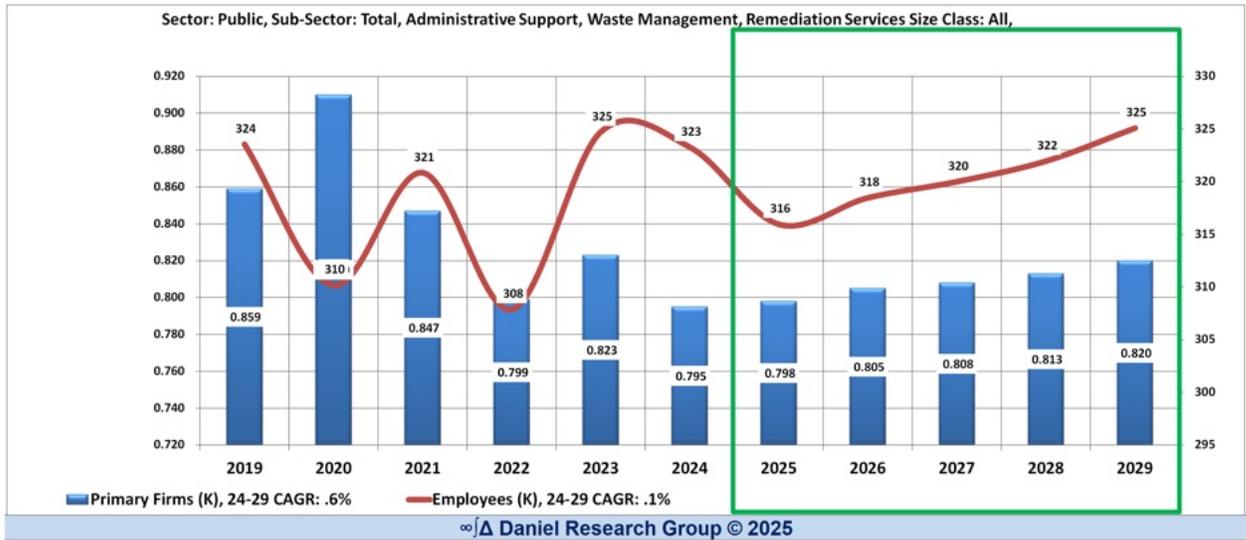


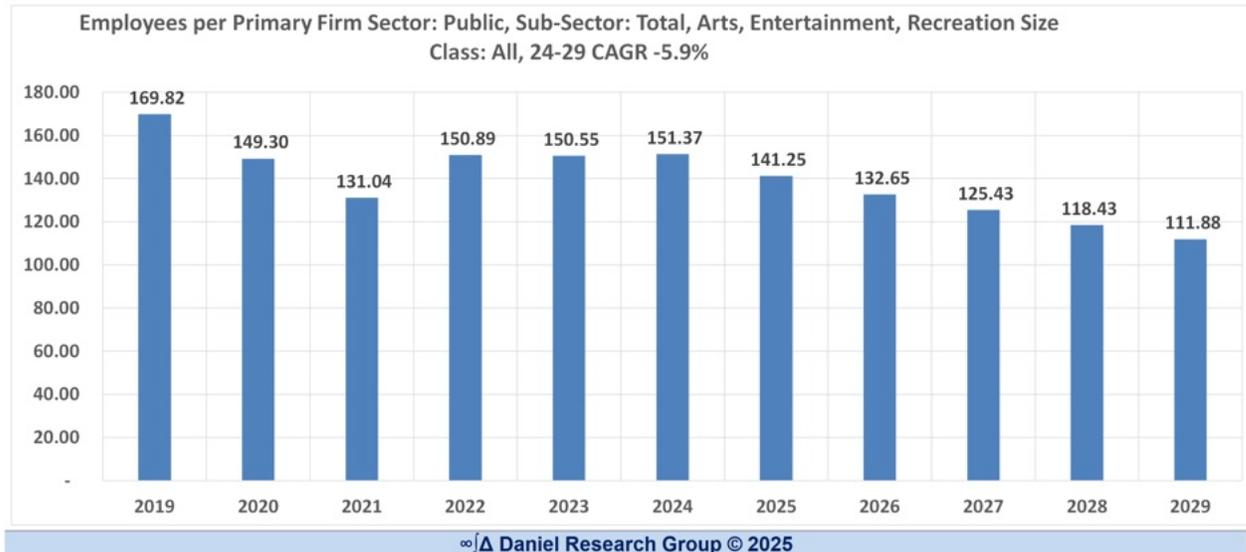


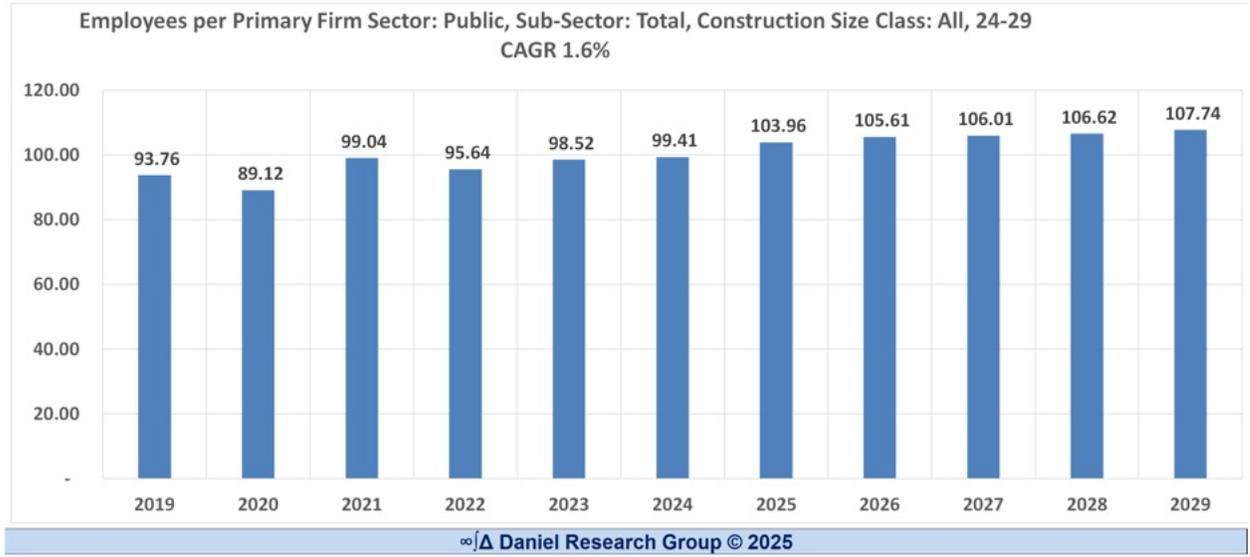


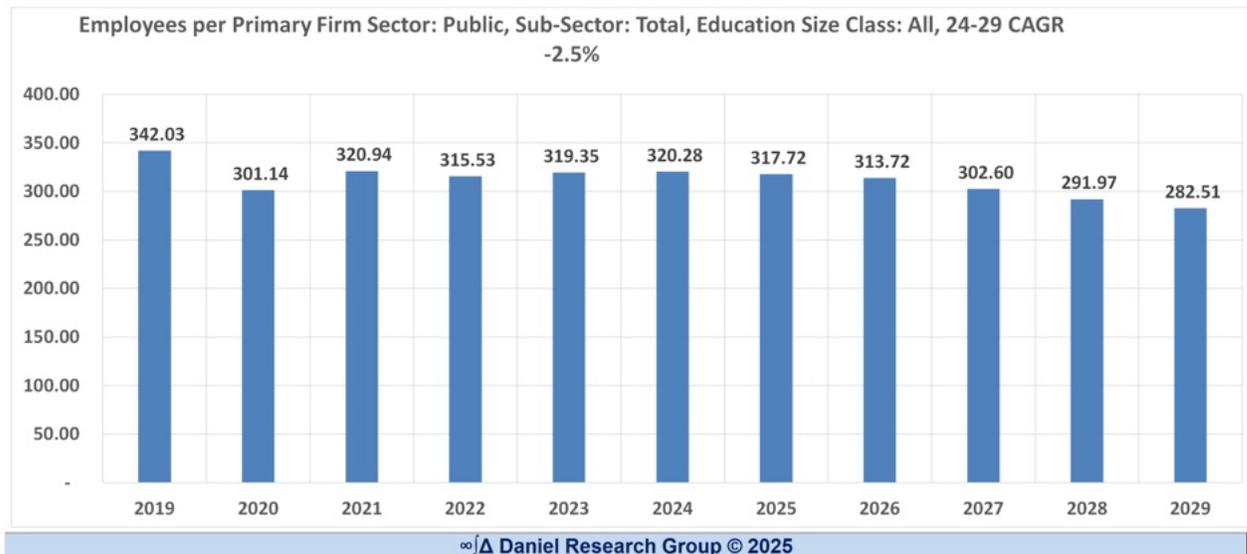
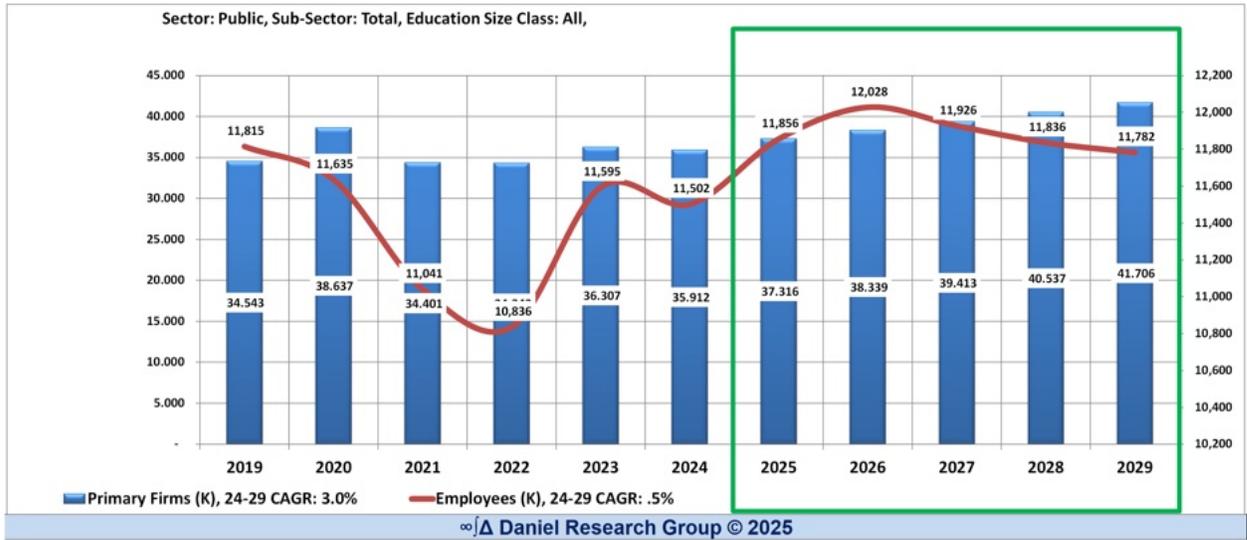


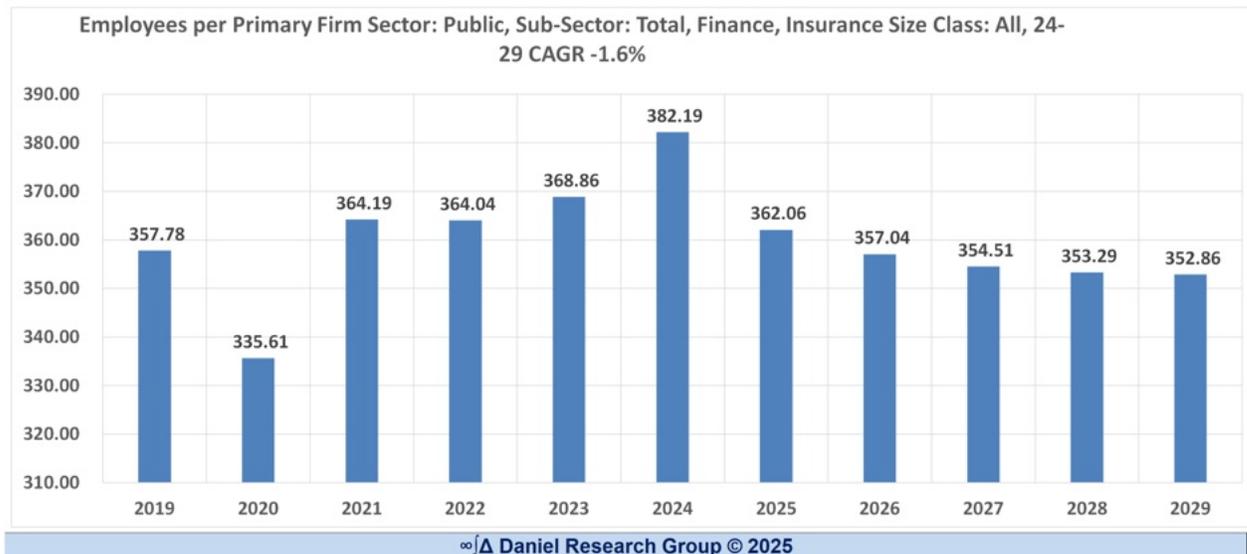
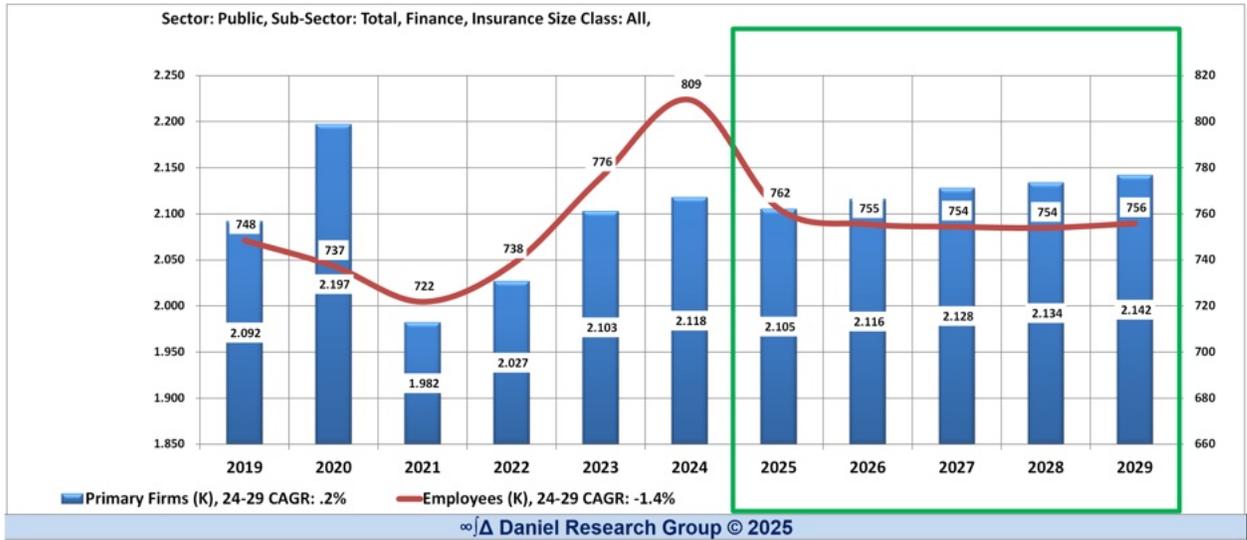
By Significant Industry – Public Sector

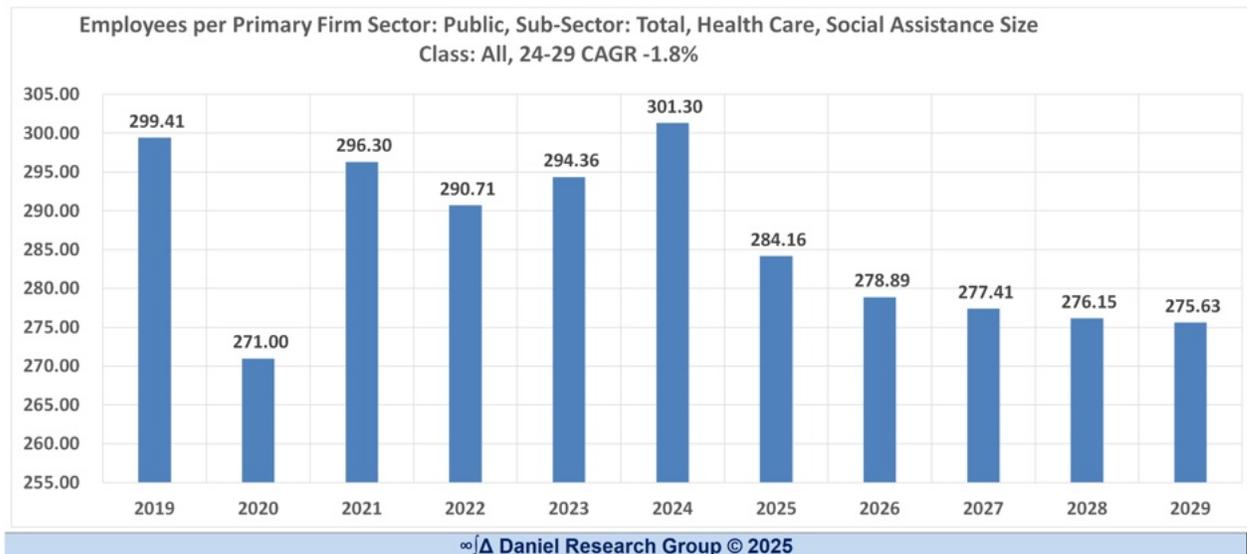
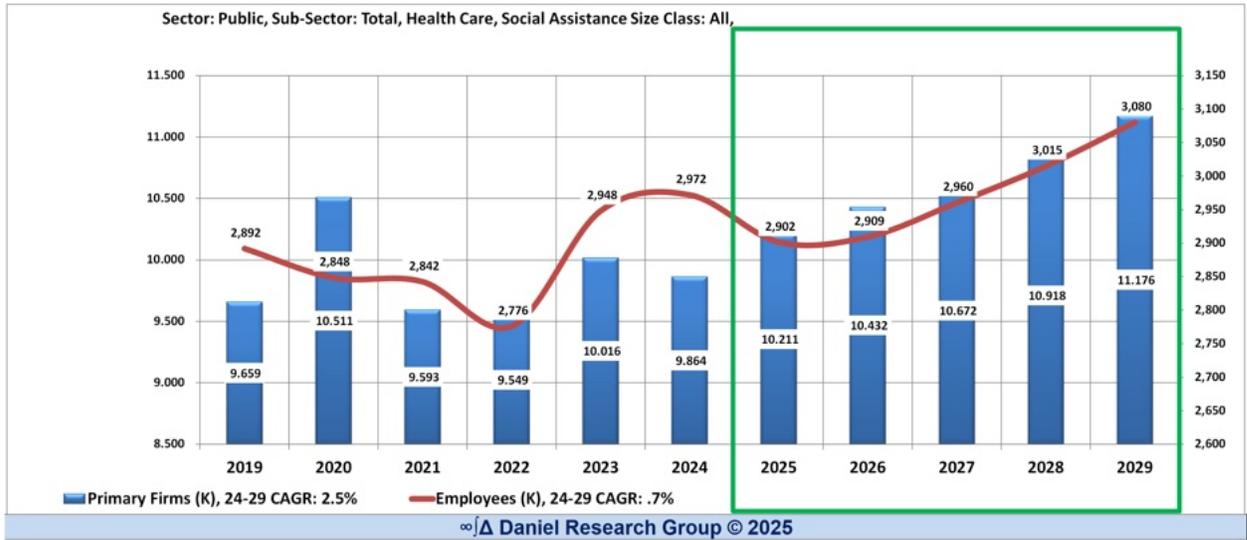


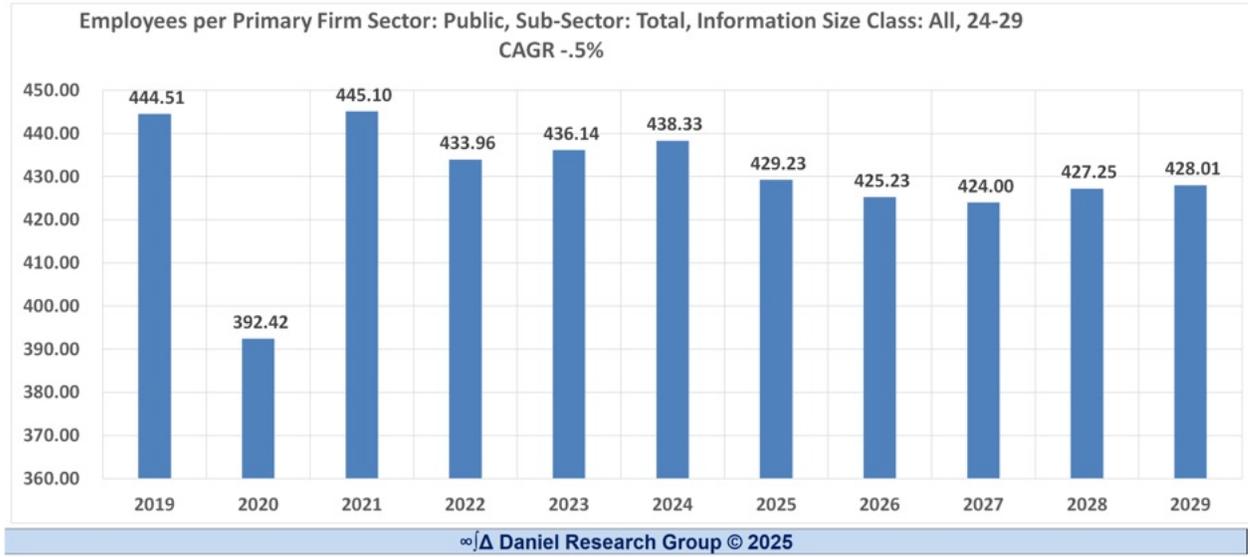


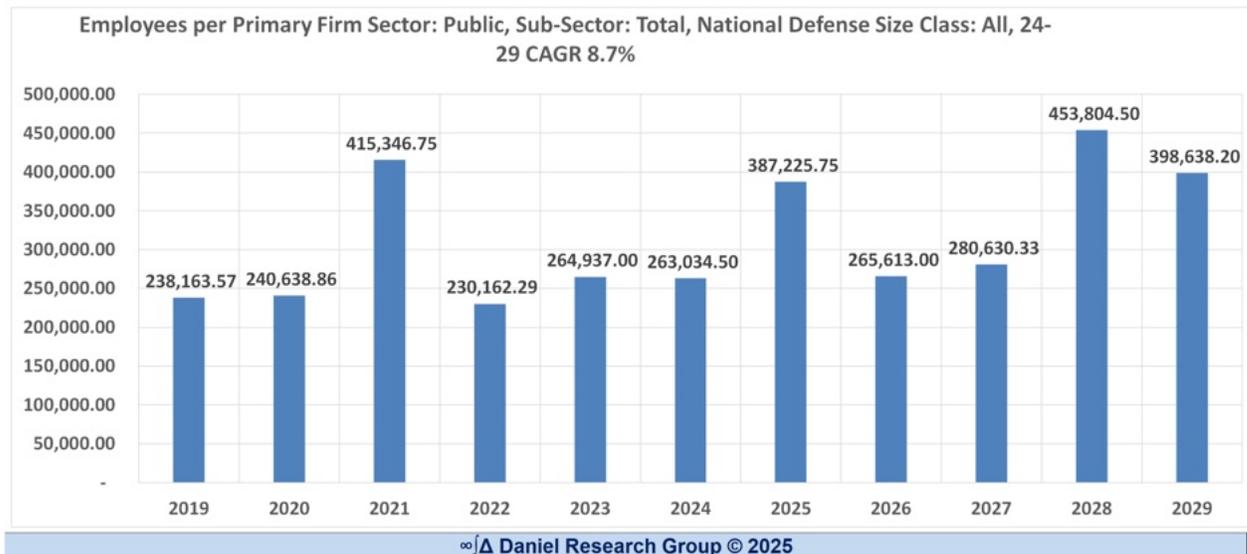


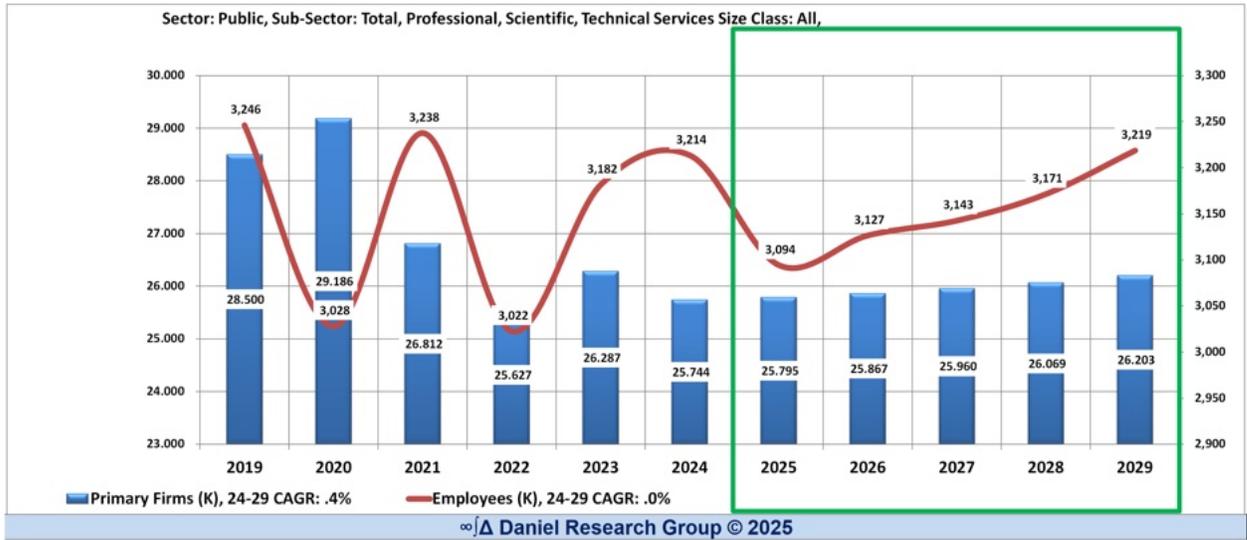


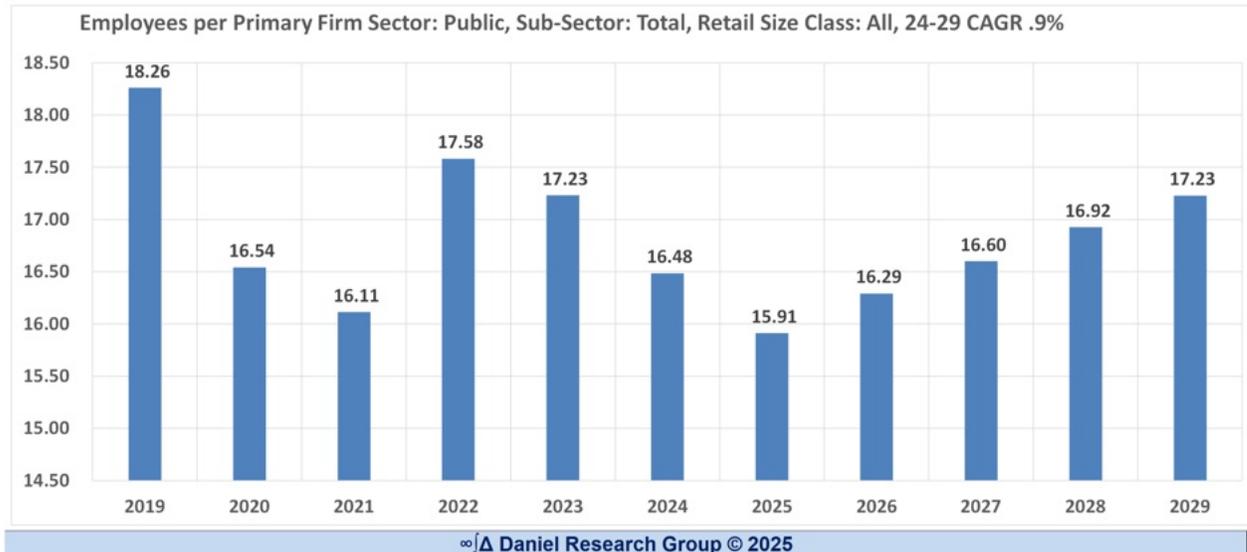
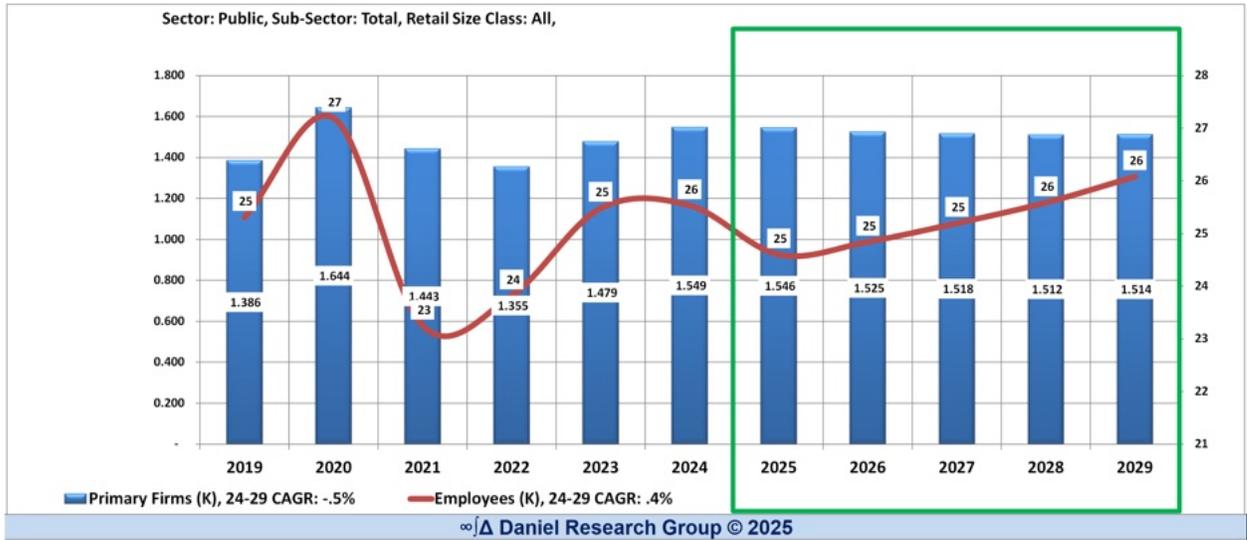


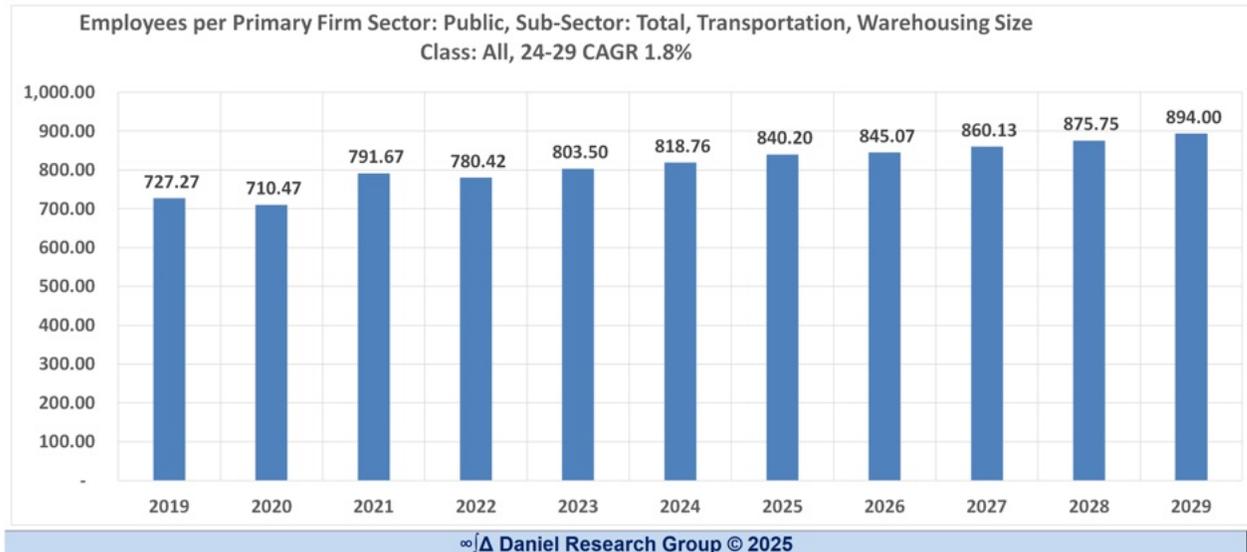




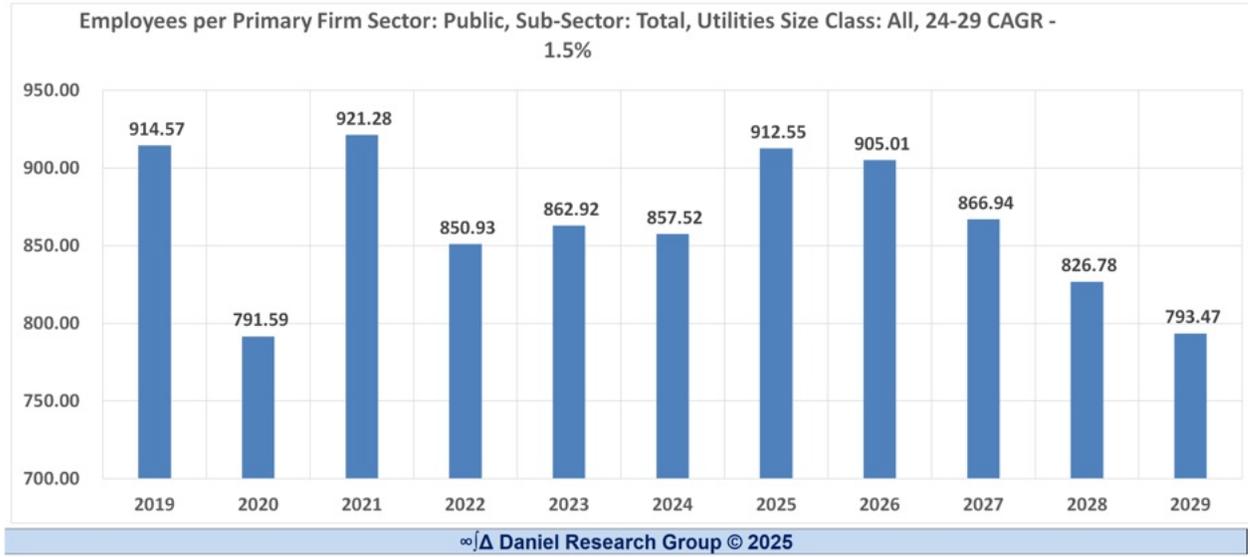


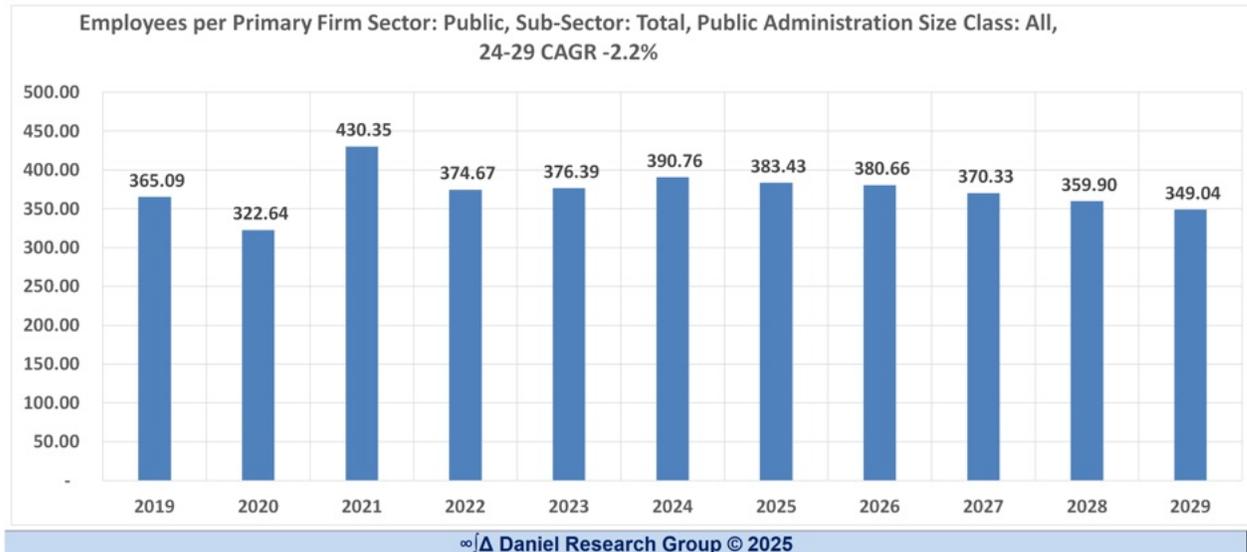
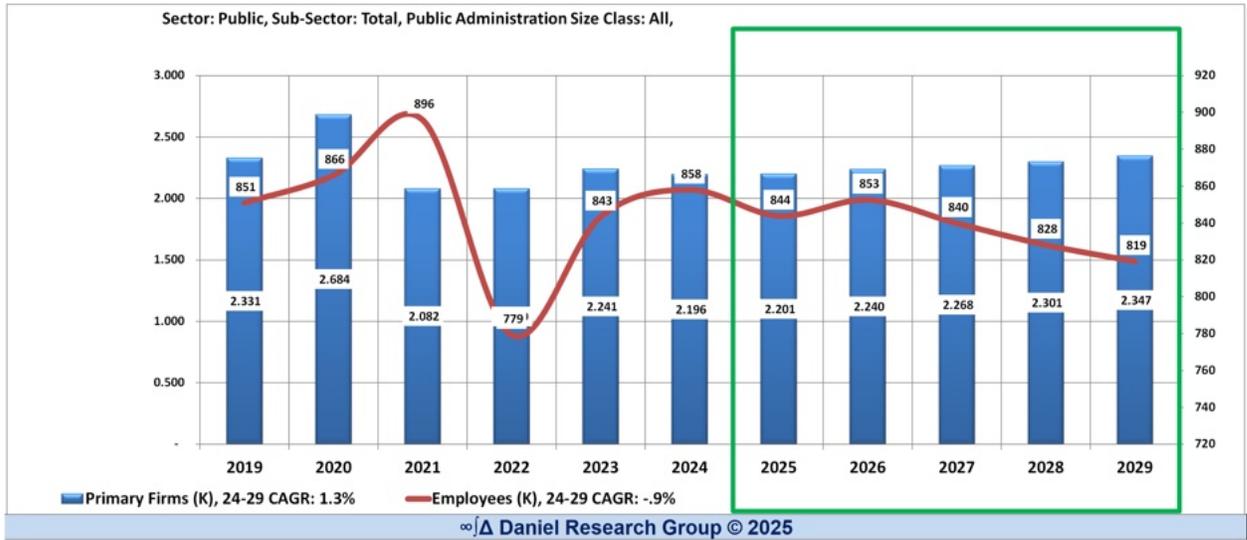




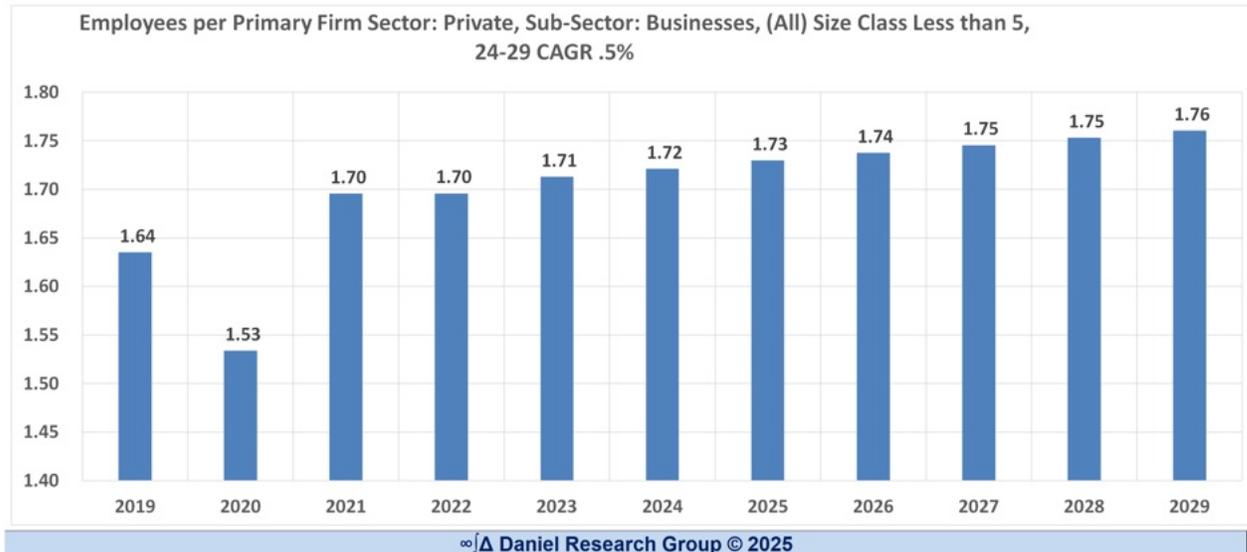
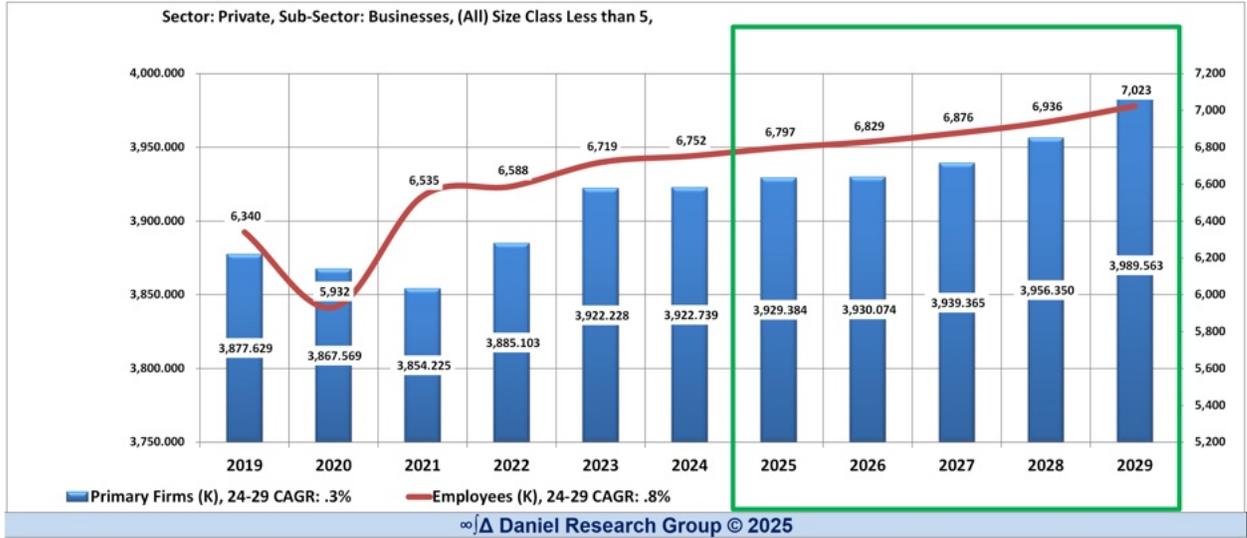


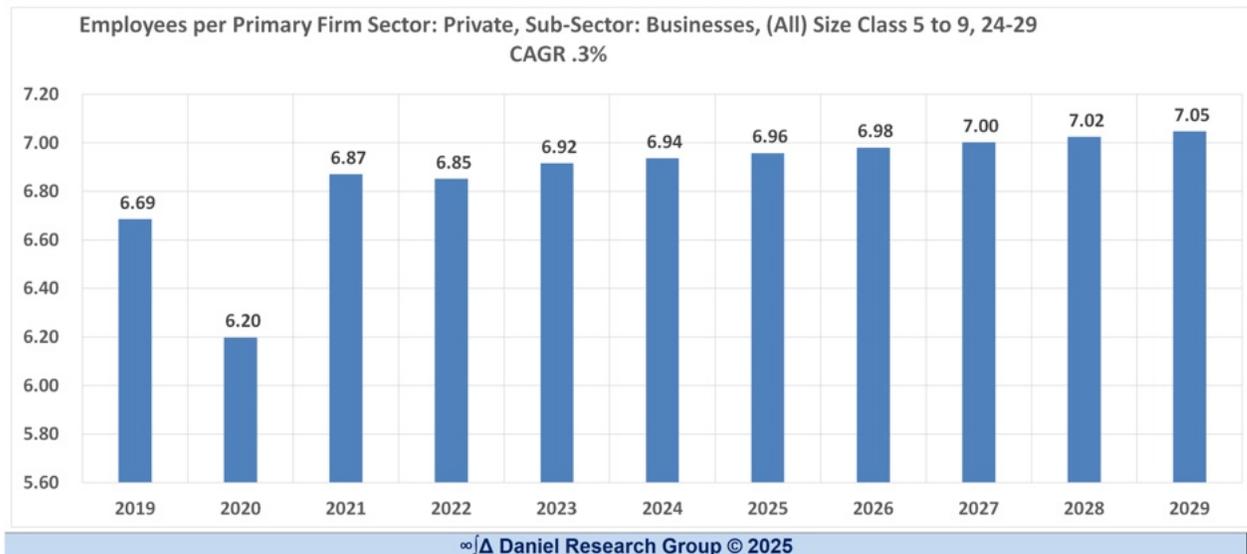
Included United States Postal Service

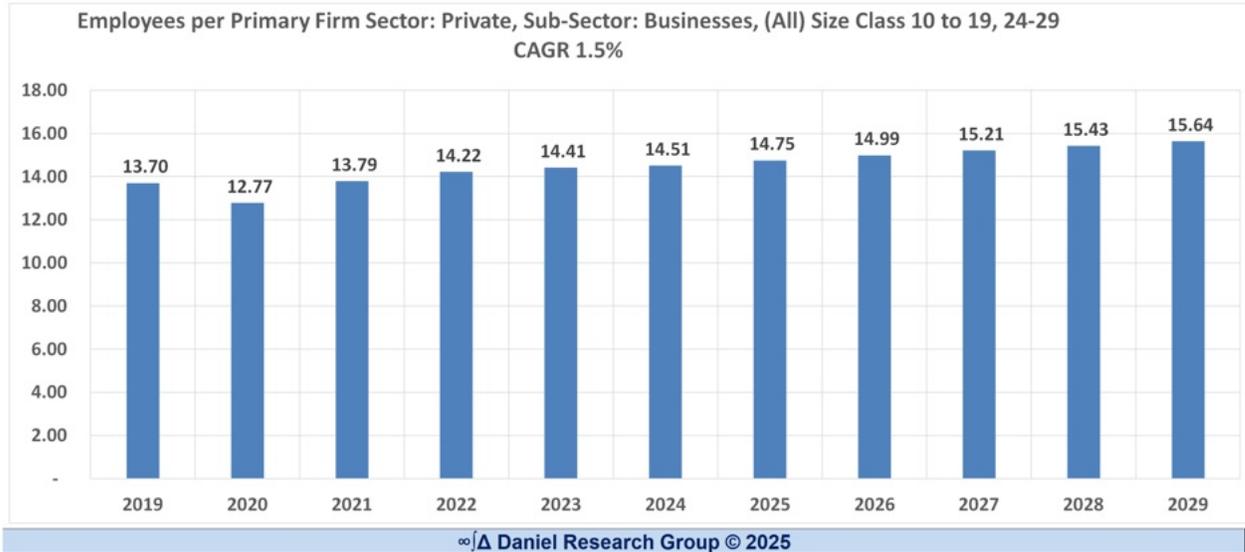


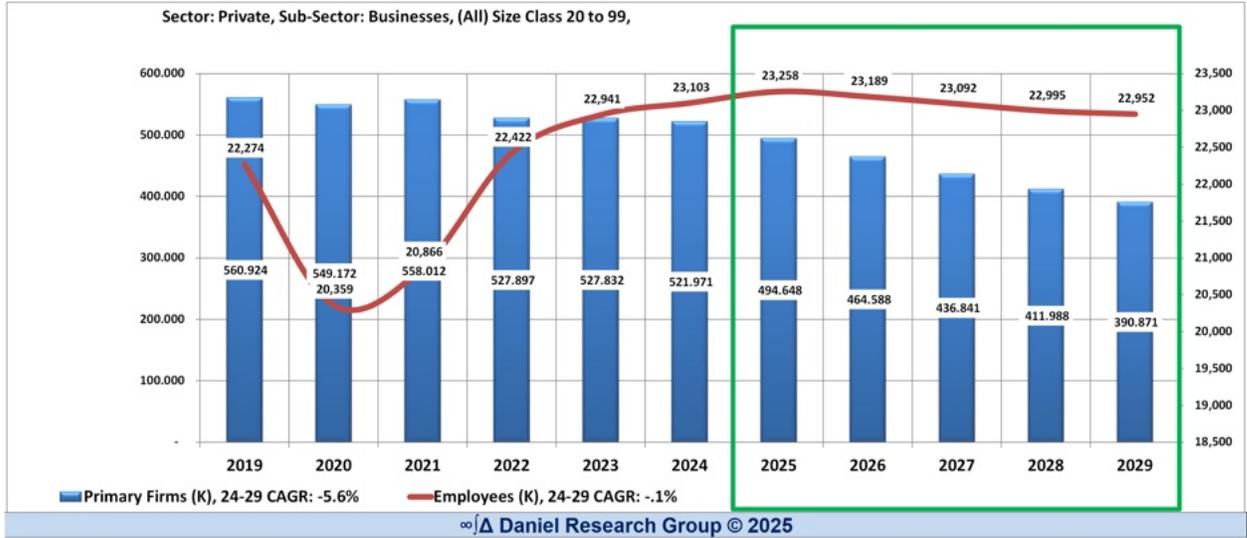


By Size Class, Private, Businesses

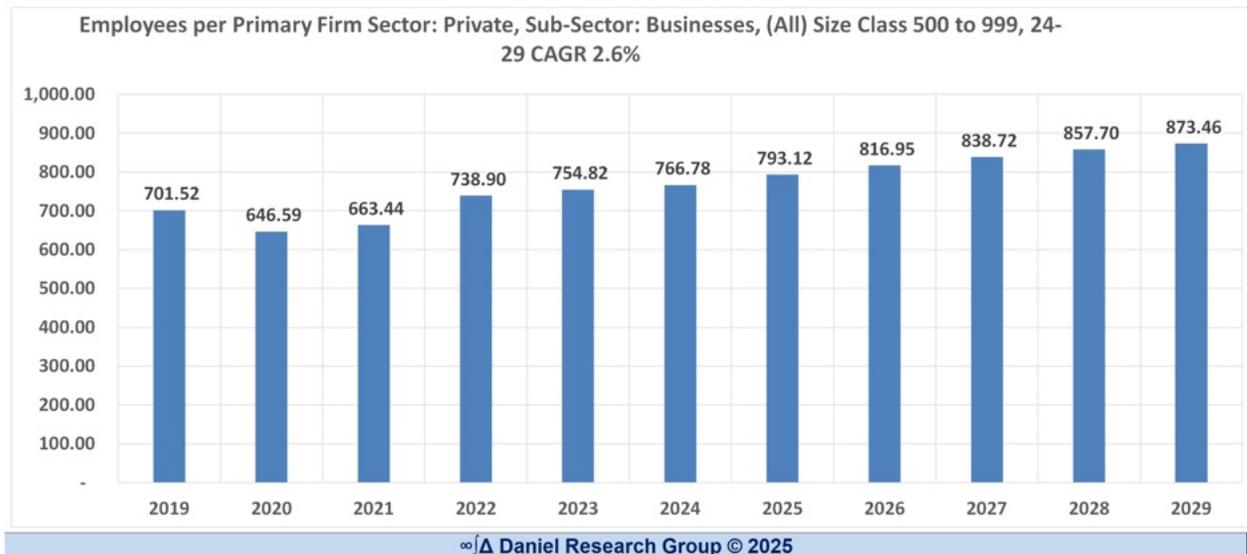


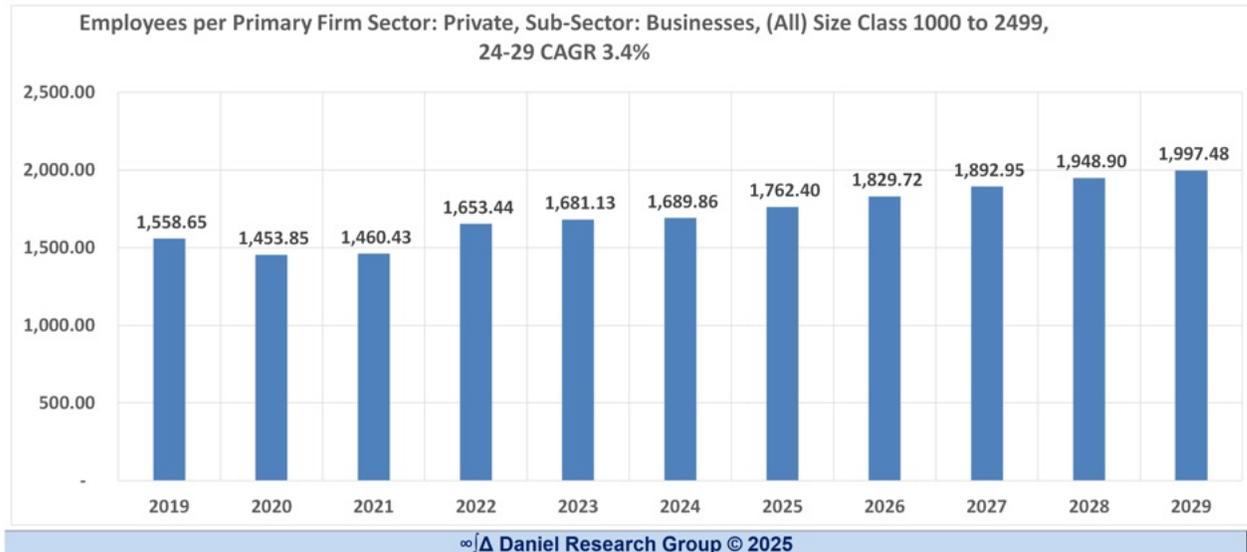
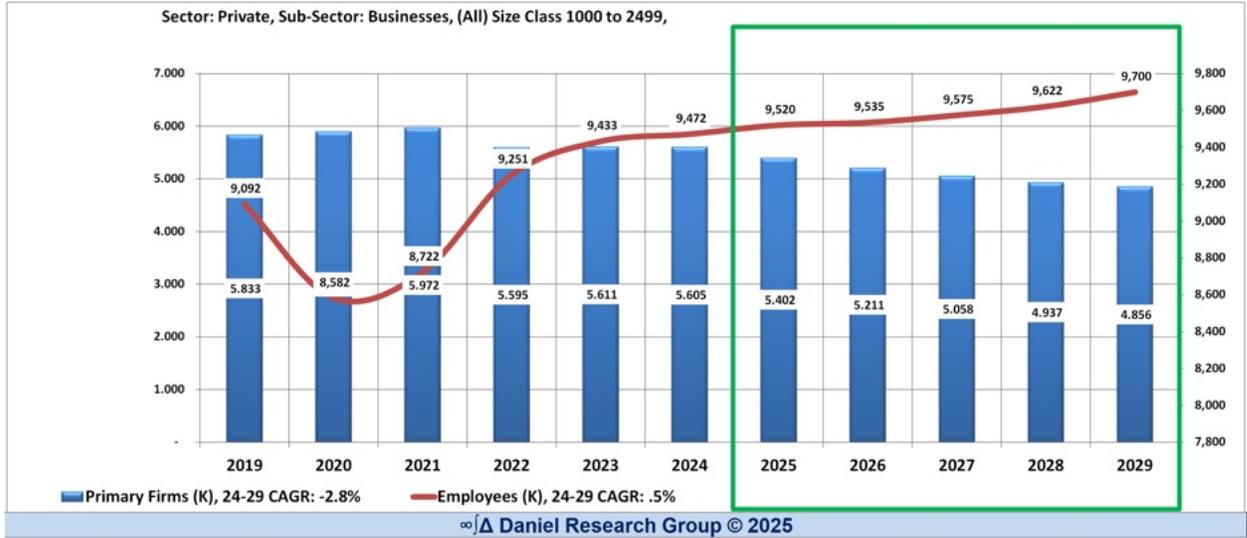


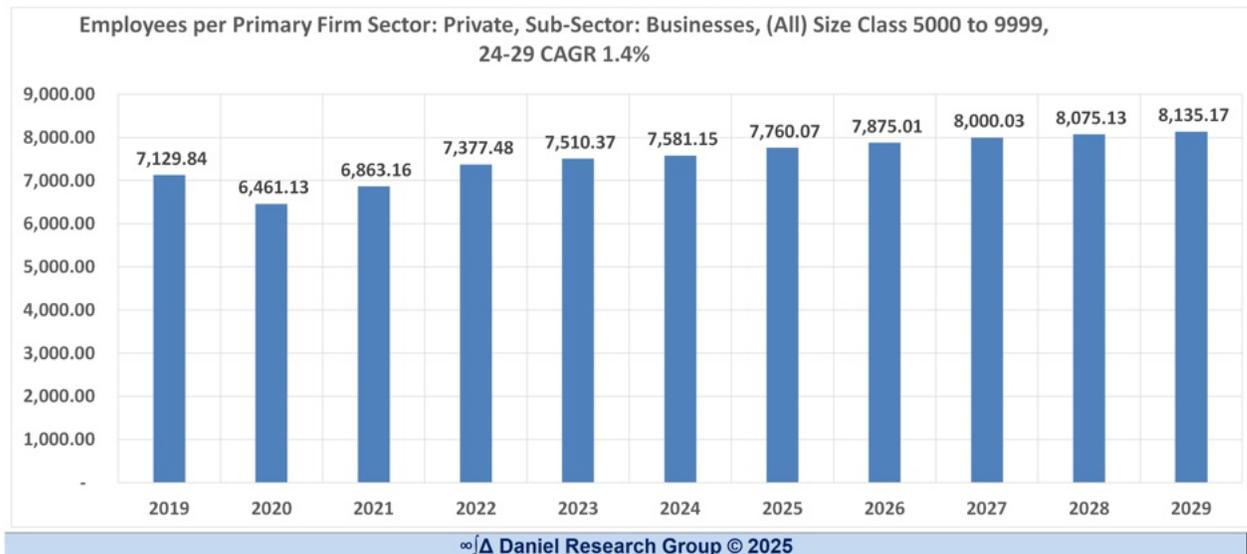
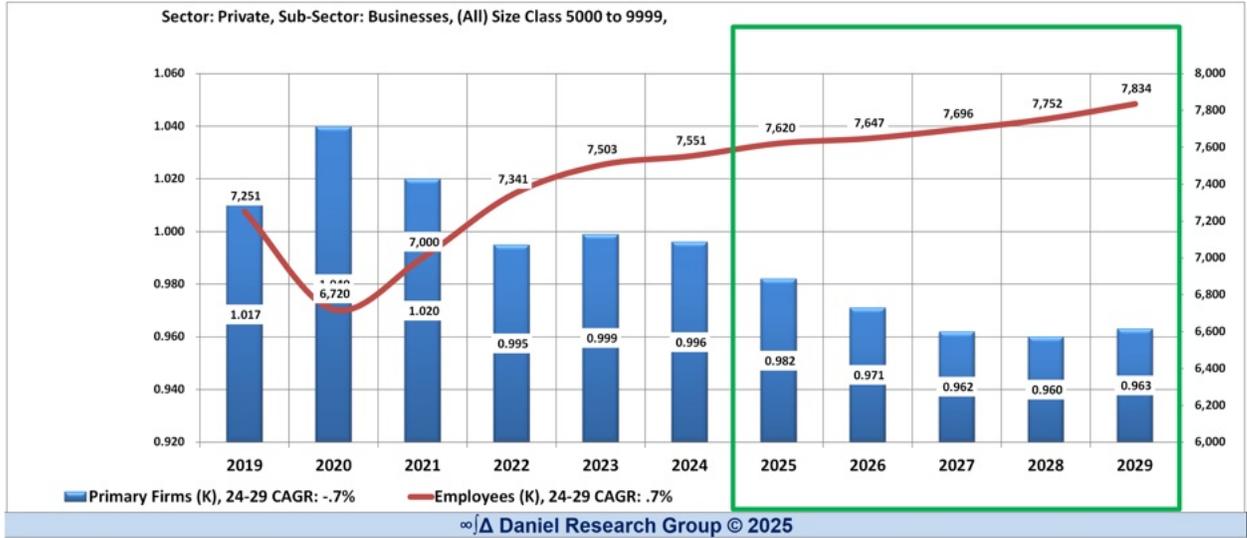


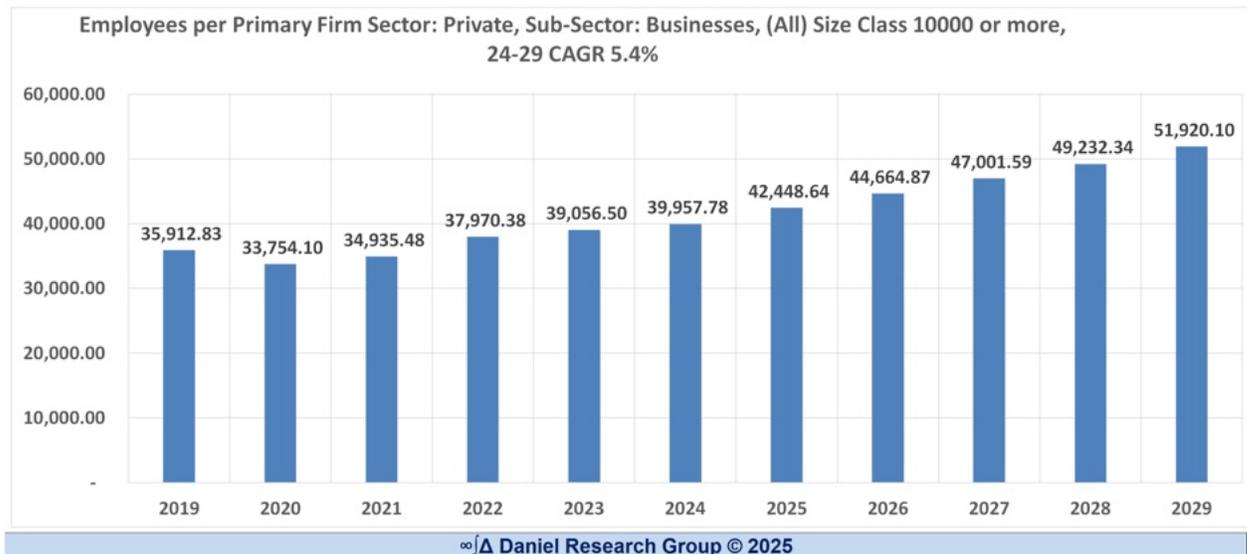
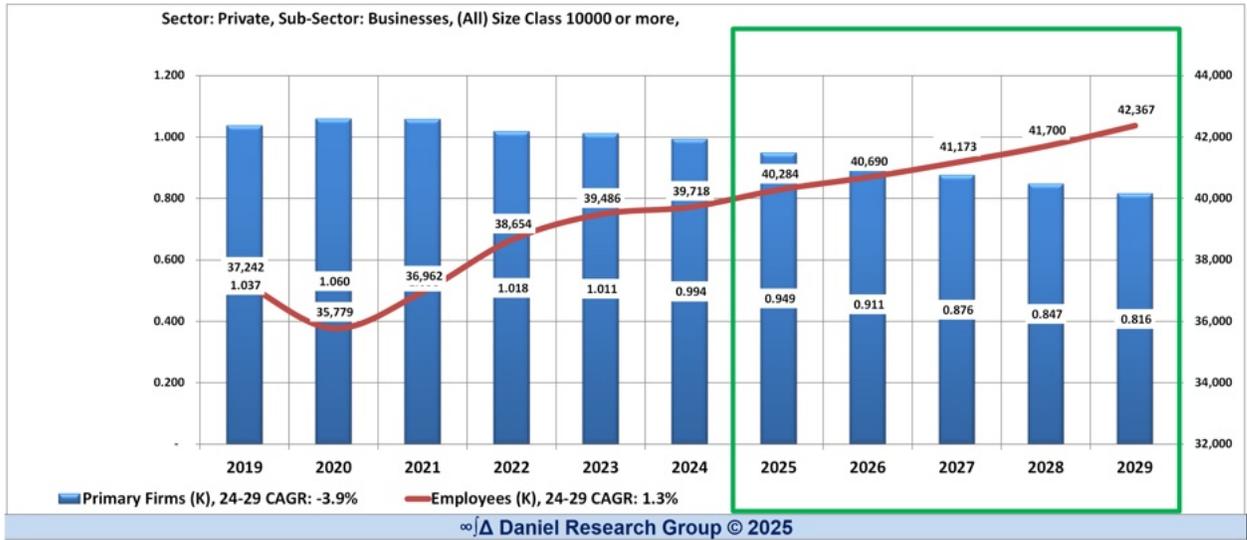












Informing Source

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](#)

49 Hill Road, Suite Four

Belmont MA, 02478

(617) 484 – 6225

Steve@DanielRG.com

www.DanielResearchGroup.com

QUARTERLY HDD NUMBERS UPDATE

The data is from the upcoming November 2025 [Digital Storage Technology Newsletter](#).

We do not know the total number of HDDs delivered last quarter because Western Digital did not report their unit shipment numbers. However, Seagate and Toshiba combined shipped 17.7M units, up 1.4% from the prior quarter. Total exabytes shipped in CQ3 2025 were up 4.9% compared with CQ2 2025 (414.3EB versus 394.8EB). HDD revenues (estimated at \$6.6B) were up about 7.25% from the prior quarter. Because of changes in Seagate and Western Digital reporting we do not know what the nearline HDD shipments were last quarter or the total HDD shipped by the industry in CQ3 2025.

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 \$2,629M for CQ3 2025. Of that, we estimate that non-HDD revenue was about \$165M, slightly more than in CQ2 2025 (\$163M). We thus estimate Seagate's total HDD revenue at \$2,464 with data center revenues representing 80% of overall Seagate revenue for the quarter (probably combined system and HDD data center revenue). Mass Capacity HDDs accounted for 88% of CQ2 2025 HDD revenue in comparison. Total storage capacity shipped was 182EB with 159.5TB for Nearline HDDs (87.6%), the balance of 22.5TB was shipped for non-nearline applications. The average HDD capacity was 14.6TB. The company shipped about 12.47M units in the quarter with an ASP I estimate at \$197.30.

Western Digital said that their overall revenue was \$2,818M in CQ3 2025 and that cloud revenue was \$2,510M. They didn't break out HDD only revenue or cloud revenue, but we assume the majority of this revenue was for raw HDDs. They also said that they shipped 204EB of total HDD storage in CQ3 2025 with 183EB (89.7% of the total) shipped for nearline applications. 20EB 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 2.2M units of their latest generation ePMR HDDs. From the quarterly report data, we do not know what the WDC HDD ASPs were, although historically they have been a few percent higher than those from Seagate.

Here are the actual statements from Seagate and Western Digital for C3Q 2025.

First from Dave Mosley of Seagate, “Seagate Technology Holdings plc delivered a very strong start to fiscal 2026. Revenue grew 21% year over year. NonGAAP gross margin set a new company record at 40.1%, and non-GAAP operating margin climbed to 29%, a level last seen in fiscal 2012. Demand strength was led by global cloud service providers, and we also saw meaningful sequential revenue growth from enterprise customers in September. The data center end market, comprised of nearline sales in the cloud, enterprise, and via customers, represented 80% of overall revenue.

Amid this improving demand backdrop, our high-capacity nearline production is largely committed under build-to order contracts through calendar 2026. Additionally, longer-term agreements that we have with our global data center customers provide clear visibility through calendar 2027, reinforcing our view that these favorable demand conditions will persist. We remain focused on executing our HAMR-based product roadmap to support our customers' growing exabyte needs and continue working with them to transition to higher capacity drives.

...Using monthly token consumption as a proxy for inferencing adoption, one major hyperscaler reported a 50-fold increase in the span of a year. This explosive growth is driving a sharp increase in unstructured data generation that creates demand for hard drive storage. Video content is a major contributor of unstructured data and is driving considerable demand for hard drives today. From social media platforms to content delivery networks and online marketing, AI-generated videos promise to further fuel demand growth...Recently, Seagate partnered with a global CSP to develop a sovereign cloud solution for managing massive volumes of sensitive telemetry and sensor data collected from a fleet of autonomous vehicles.

...Momentum continues to build for our HAMR-based Mosaic platforms, and we achieved several important milestones in the quarter. Consistent with what we discussed during our analyst event, we now have have global CSPs qualified on Mosaic three-plus terabyte per disk products, which can deliver capacities up to 36 terabytes per drive. We remain on track to qualify the remaining three global CSPs within 2026.

Additionally, we shipped over 1,000,000 Mosaic drives in September. These products are performing well in live production environments, and we are on pace to achieve 50% exabyte crossover on nearline HAMR drives in 2026. We started qualification with a second major CSP on the Mosaic four-plus terabyte per disk platform, with initial volume ramp starting in the first half of next calendar year. This platform will offer capacities of up to 44 terabytes. Advancing aerial density is a key competitive advantage not just for Seagate, but for the hard drive industry overall.

We are leveraging our manufacturing expertise and advancements in technologies, including silicon photonics, to pave the path to 10 terabytes per

disk...We are advancing our HAMR-led technology roadmap, which creates significant value for our customers and positions Seagate for long-term success.”

Gianluca Romano (Seagate CFO) said, for C3Q 2025, “We delivered revenue of \$2.63 billion, up 8% sequentially and up 21% year over year. We achieved a record non-GAAP gross margin of 40.1%, up 220 basis points sequentially, and we expanded non-GAAP operating margin by 280 basis points to 29% sequentially. Our resulting non-GAAP EPS was \$2.61, exceeding the high end of our guided range. We have continued to execute our technology roadmap to support ongoing demand momentum for our higher capacity products. In September, we shipped 182 exabytes, up 32% year over year, with the vast majority of that volume delivered to global data center customers.

...In September, data center revenue represented 80% of our total revenue at \$2.1 billion, up 13% sequentially and 34% year on year. Demand from global cloud customers continues to grow, and we also saw a notable improvement in the enterprise OEM markets. We project these positive trends to continue, with cloud growth expected to outpace enterprise demand.

...In September, we shipped 159 exabytes to data center customers, up from 137 exabytes in the prior period. Cloud exabyte demand increased for the ninth consecutive quarter, resulting in close to 80% of nearline volume on drive capacity at or above 24 terabytes as customers continue to mix up to higher capacity drives.

...Over the past year, average nearline drive capacity has increased by 26%, which is a primary contributor to our exabyte volume growth. Amid the tight supply condition, we are partnering closely with data center customers to support, where possible, accelerate their qualification timeline on our high-capacity Mosaic products.

...Turning now to the December outlook, the demand environment remains strong, particularly among global cloud data centers. We expect to increase revenue and expand margins as these customers continue to shift to our next-generation storage solutions to support their increasing demand. We expect December revenue to be in a range of \$2.7 billion, plus or minus \$100 million, which represents a 16% year-over-year improvement at the midpoint.”

Irving Tan, CEO of WDC said, “...The rapid adoption of AI and data-driven workloads at hyperscalers is driving robust demand for our products and solutions. To fulfill the demand of more exabytes of storage, our customers are increasingly transitioning to higher capacity drives. Shipments of our latest ePMR products, offering up to 26 TB CMR and 32 TB UltraSMR capacities, continue to grow at an impressive pace, surpassing 2.2 million units in the September quarter. Our ability to reliably scale our ePMR technology and transition customers to higher capacity drives is one of several ways we support the growing demand for exabytes.

...We are also investing in head wafer and media technology and capacity to drive aerial density higher. In addition, we're increasing our manufacturing throughput by leveraging automation, AI tools, and enhancing our test capabilities.

...The AI-driven growth in data storage is accelerating demand for higher capacity drives, which comes with greater manufacturing complexity and longer production lead times. As a result, our customers are providing greater visibility into their long-term needs, which in turn strengthens our partnership and helps us to support their future growth requirements. Our top seven customers have now provided purchase orders extending throughout the first half of calendar year 2026, and five of them have provided purchase orders covering all of calendar year 2026. I'm also pleased to share that one of our largest hyperscale customers has signed an agreement covering all of calendar year 2027. These commitments underscore both the essential role of our products in the AI data economy and our customers' strong confidence in our product roadmap, including the transition to HAMR technology.

We are making rapid progress in our HAMR development and are on track to start HAMR qualification for one hyperscale customer in the first half of calendar year 2026 and to expand the qualification process to up to three hyperscale customers through calendar year 2026...This positions us well for the ramp-up of volume production in the first half of calendar year 2027. In parallel, we will begin qualification of our next-generation ePMR drives in the first quarter of calendar year 2026...We continue to expand our proven ePMR roadmap even further while bringing new technologies, including HAMR, to market. In parallel, our engineering teams are focused on improving data throughput speed and bandwidth of our drives, as well as power efficiency.

...For the fiscal first quarter, Western Digital delivered revenue of \$2.8 billion...For the fiscal second quarter of 2026, we expect continued revenue growth driven by data center demand and improved profitability led by the adoption of our higher capacity drives.”

Kris Sennesael (WDC CFO) said, “During the first quarter of fiscal 2026, revenue was \$2.8 billion, up 27% year-over-year, driven by strong demand for our nearline drives. Earnings per share was \$1.78. Both revenue and EPS were above the high end of the guidance range. We delivered 204 EB to our customers, up 23% year-over-year. This includes 2.2 million drives of our latest generation ePMR, with capacity points up to 26 TB CMR and 32 TB UltraSMR. Cloud represented 89% of total revenue at \$2.5 billion, up 31% year-over-year, driven by strong demand for our higher capacity nearline product portfolio.

Client represented 5% of total revenue at \$146 million, up 5% year-over-year. Consumer represented 6% of revenue at \$162 million, down 1% year-over-year....I will now turn to the outlook for the second quarter of fiscal 2026. This outlook includes our current estimate of all anticipated or known tariff-related impacts on

our business in this period. We anticipate revenue to be \$2.9 billion, plus/minus \$100 million. At midpoint, this reflects a growth of approximately 20% year-over-year.”

In **Table 1** we see an 2.7% average sales price (ASP) increase for HDDs from C2Q 2025 to C3Q 2025. However, it may in reality be higher, since WDC ASPs are usually higher than Seagate’s and WDC did not report information enabling an ASP calculation this quarter.

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

	Q4 24 Units/\$	Q1 25 Units/\$	Q2 25 Units/\$	Q3 25 Units/\$
Seagate	13.00/\$166.00	11.4/\$175.70	12.5/\$182.50	12.5/\$197.60
WD	13.50/\$172.00	12.1/\$179.00	12.9/\$201.90	

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.

FIGURE 1. AVERAGE DRIVE PRICE TREND FOR SEAGATE, WESTERN DIGITAL AND (HISTORICALLY, HGST) (C4Q '98 TO C3Q '25)



Table 2 shows estimated industry revenue for C3Q 2025 and for the year to date. HDD revenues were up about 8.1% in C3Q 2025 compared to C2Q 2025.

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

Company	Q1	Q2	Q3	Q4	Total
Seagate	\$2,003	\$2,281	\$2,464		\$6,748
Western Digital	\$2,294	\$2,605	\$2,818		\$7,717
Toshiba	\$773	\$879	\$951	\$0	\$2,604
Total	\$5,070	\$5,765	\$6,233	\$0	\$17,069

Table 3 shows C3Q 2025 estimated Exabytes Shipped by company and general application. Note that per Seagate’s reports, Legacy includes PC, CE, Branded and enterprise applications. Table 4 shows total capacity shipped year to date.

TABLE 3. CALENDAR C3Q 2025 ESTIMATED DISK DRIVE EXABYTES OF CAPACITY SHIPMENTS

Company	Nearline	Legacy	Total
Seagate	159.50	22.50	182.00
Western Digital	170.00	20.00	190.00
Toshiba	34.93	7.36	42.29
Total	364.43	49.86	414.29

TABLE 4. YTD ESTIMATED DISK DRIVE EXABYTES OF CAPACITY SHIPMENTS

Company	Q1	Q2	Q3	Q4	Total
Seagate	143.60	162.50	182.00		488.10
Western Digital	166.00	190.00	190.00		546.00
Toshiba	40.98	42.29	42.29		125.56
Total	350.58	394.79	414.29	0.00	1159.66

FIGURE 2 SHOWS YEAR TO DATE EXABYTE SHIPMENTS BY COMPANY.

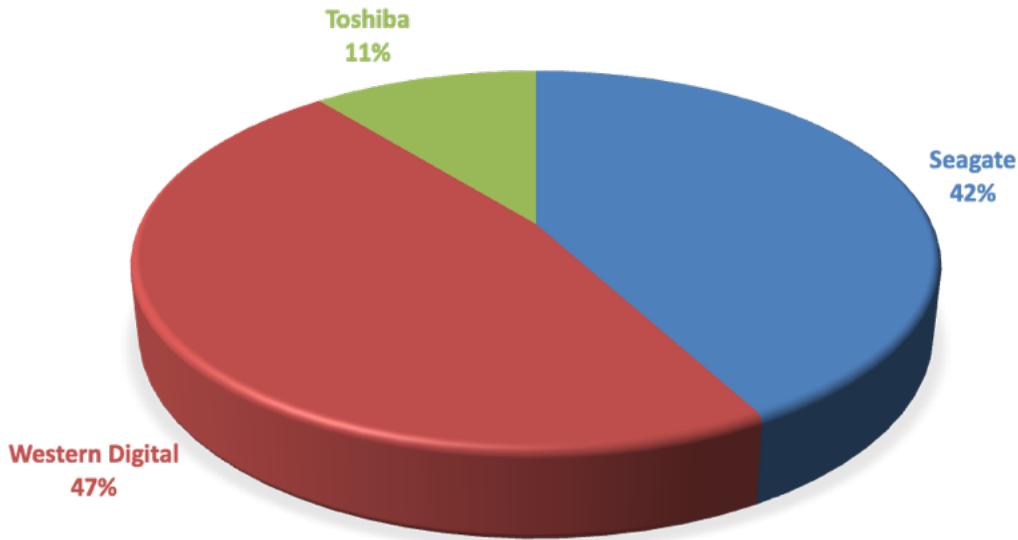


FIGURE 3 SHOWS TOTAL REVENUE SHIPMENTS BY COMPANY.

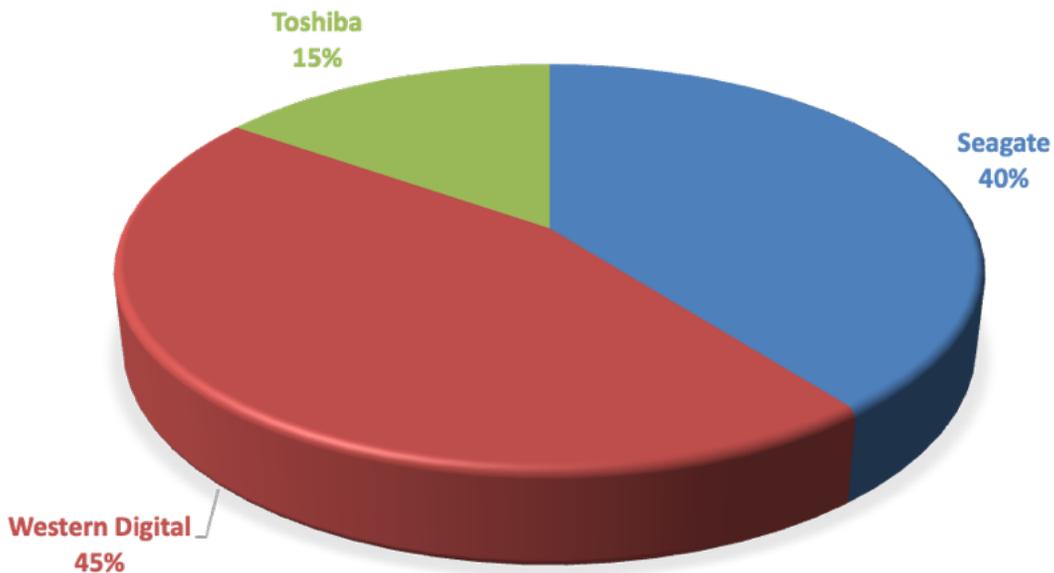


Figure 4 shows high, median and low estimates of total drive unit volume estimates out to 2030. Note that I have increased the expected unit shipments for 2025, to match the total shipments from 2024 and also increased estimates for the following years as I believe that AI applications will be driving increased demand for HDDs, although without shipment numbers from Western Digital and without more detailed information from Seagate it is hard to check these projections. However, I only projected the resulting higher nearline drive growth only through 2027 as I expect there will be a correction after that.

FIGURE 4. BANDED HARD DRIVE VOLUME PROJECTIONS

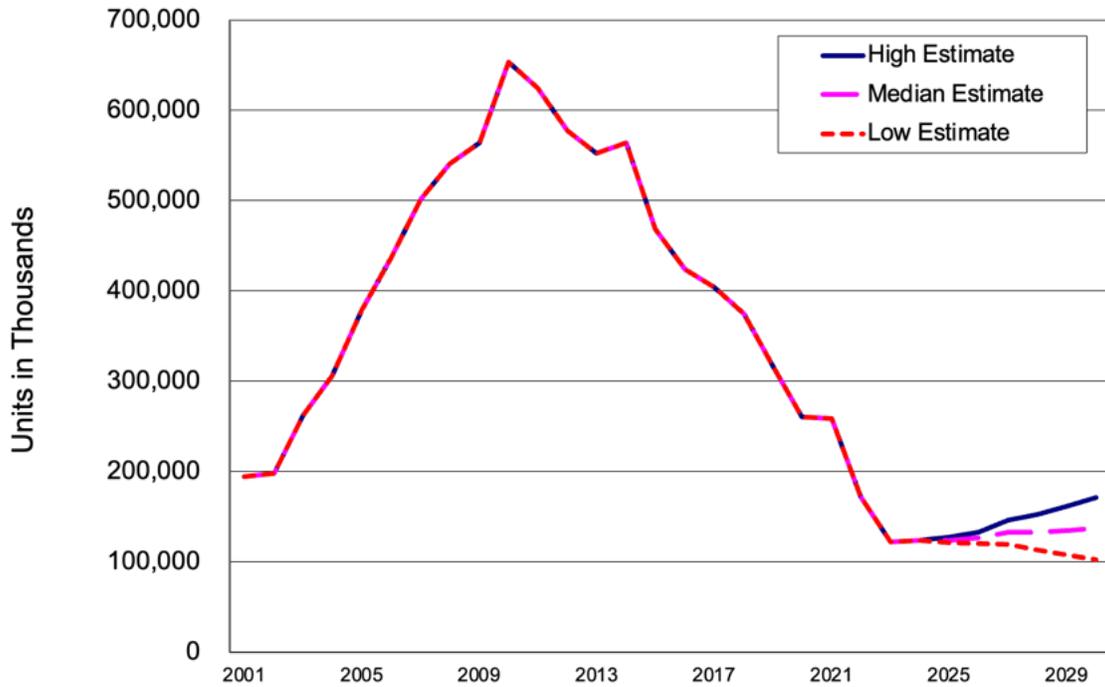
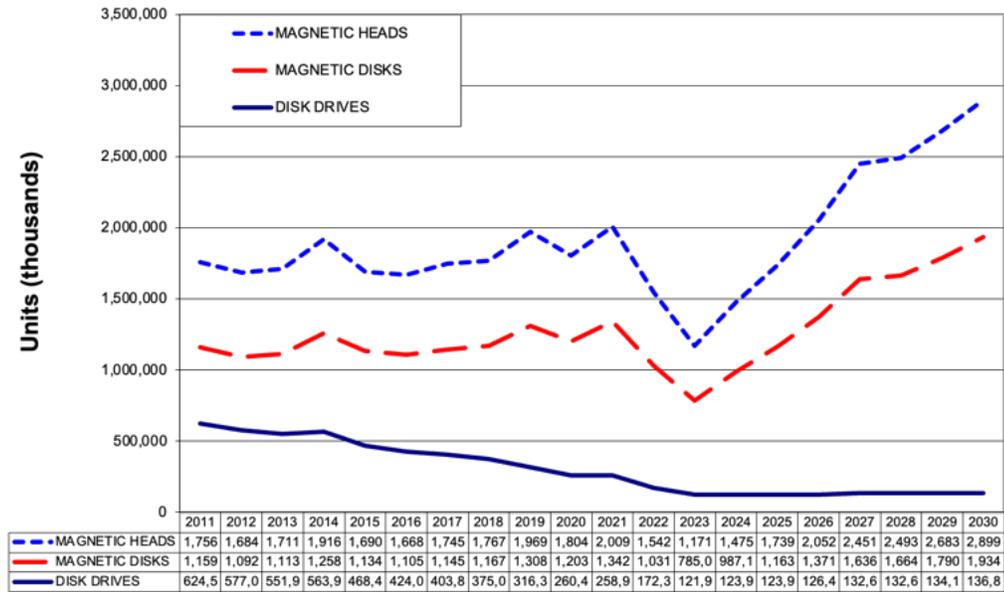


Figure 5 gives our estimates for HDDs, magnetic media (disks) and heads production out to 2030 based upon the median estimate for HDDs in Figure 4.

FIGURE 5. PROJECTED DEMAND FOR HDD HEADS AND DISKS



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

The growth of capacity-oriented nearline drives for enterprise and hyperscale applications will result in more components per drive and thus higher head and media demand than in the past by 2026 or 2027 and with high disk and head counts per drive for high-capacity HDDs, provides the biggest driver for heads and media.

Figure 6 shows laboratory demonstrations and product announcement maximum areal densities history. Seagate's current highest capacity HAMR HDD is 36TB capacity (using SMR) and the 10 disk 28TB WDC SMR drives have about 1.4 Tbps so we sized the 36TB drive as 1.8Tbps, another product areal density high point. In addition, Seagate announced a laboratory demonstration of 4 Tbps in early 2024. These advances are reflected in the charts below.

FIGURE 6. AREAL DENSITY ANNOUNCEMENT DEVELOPMENTS

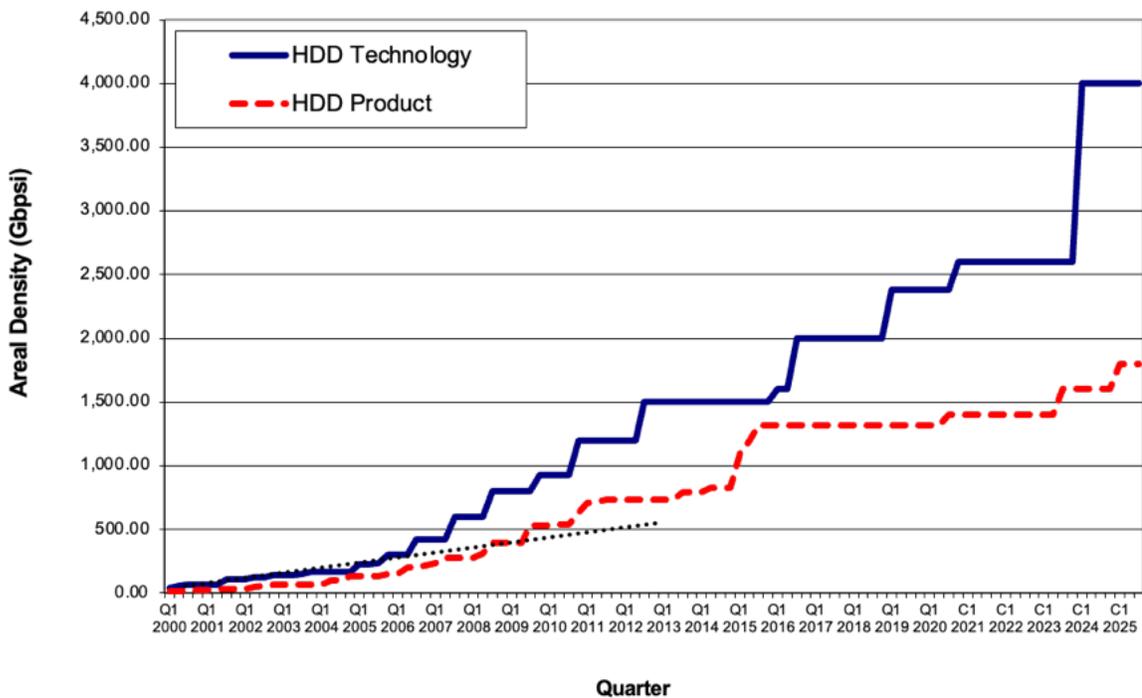


Figure 7 shows the product announcement maximum areal density per quarter compared to lines representing areal density growth. **Figure 8** gives our projections for drives by market niche out to 2030. **Figure 9** does the same for drive form factor.

FIGURE 8. AREAL DENSITY PRODUCT ANNOUNCEMENTS VS. CAGR

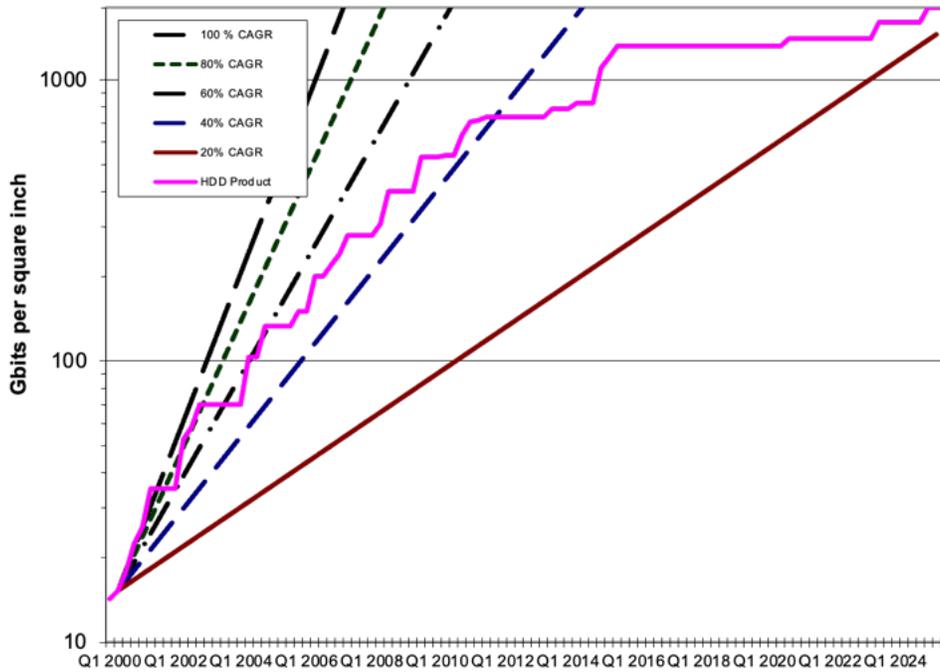


FIGURE 8. PROJECTION OF DRIVES BY MARKET NICHE (1,000'S-UNITS)

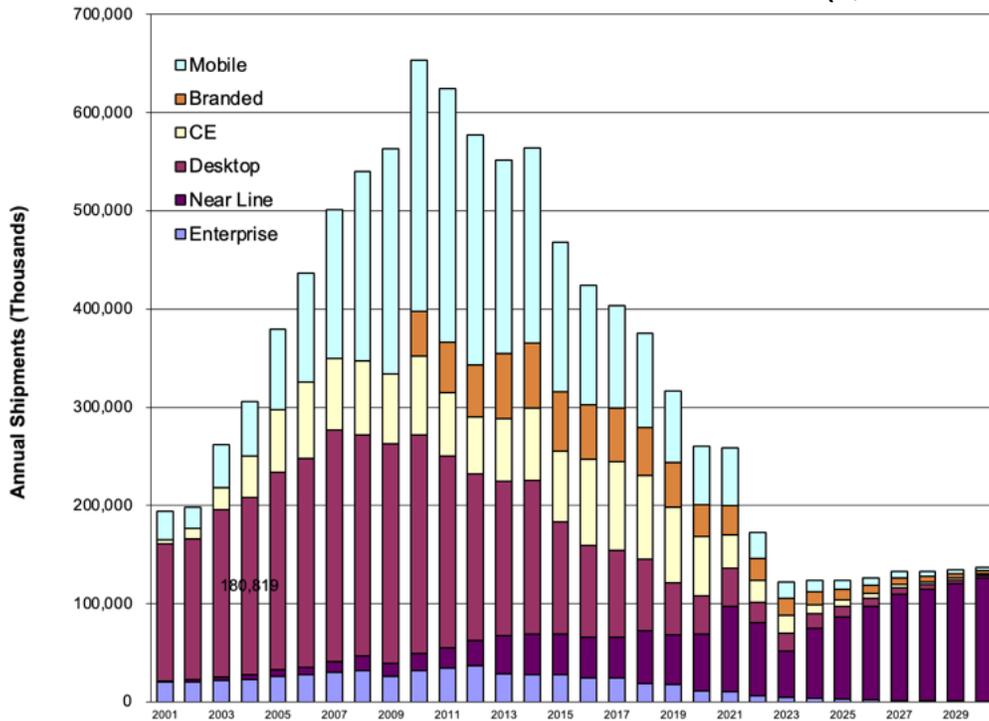


FIGURE 9. PROJECTION OF DISK DRIVES BY FORM FACTOR (1,000'S-UNITS)

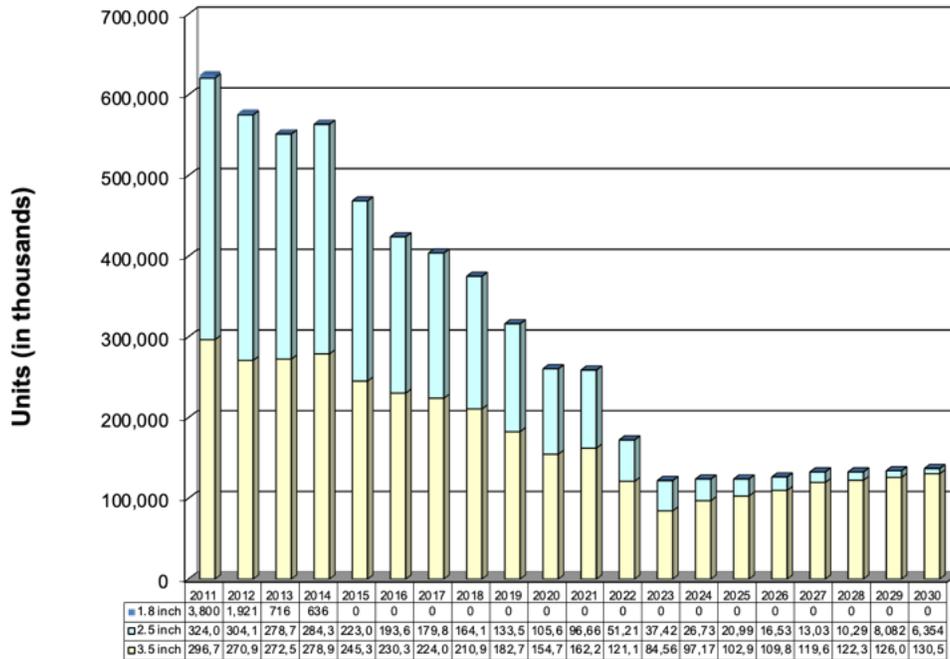


Figure 10 projects total shipments of HDD storage capacity on an annual basis to 2030.

FIGURE 10. PROJECTION OF DISK DRIVE ANNUAL CAPACITY SHIPMENT

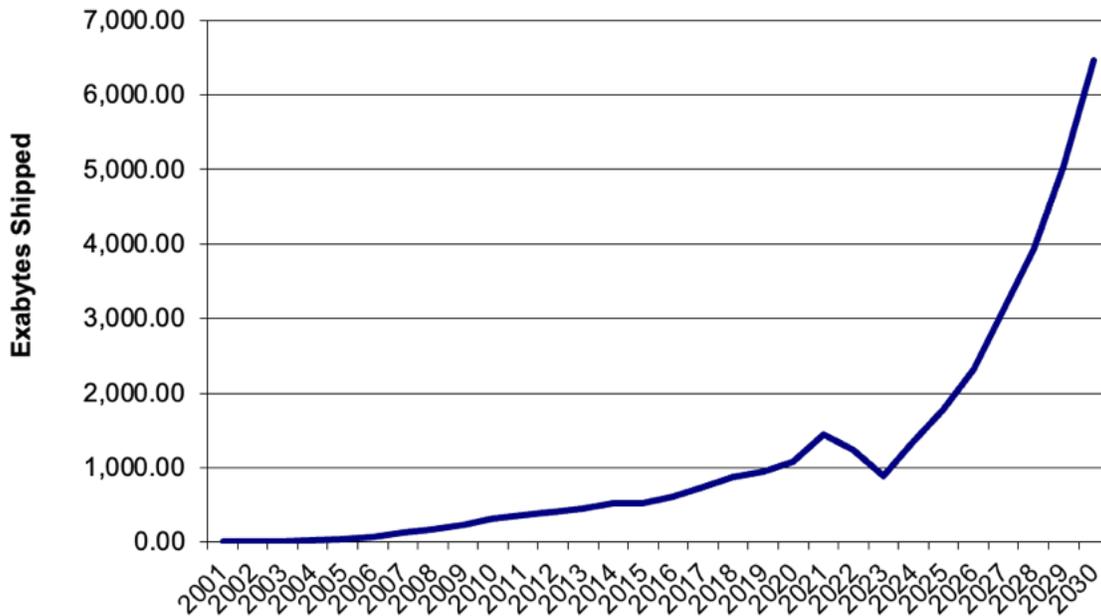


Figure 11 shows the estimated average \$/GB for HDDs. **Figure 12** shows total projected shipments for Tape (LTO), SSDs and HDDs. Note that although the bulk of anticipated shipped storage capacity is in HDDs by 2030, the percentage of that storage that is in SSDs will have increased.

FIGURE 11. AVERAGE \$/GB FOR HDDS

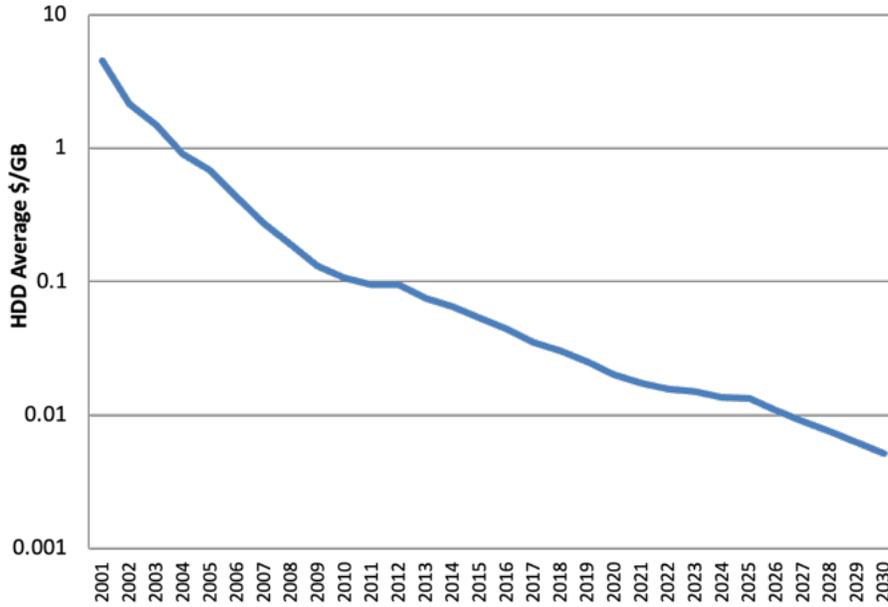
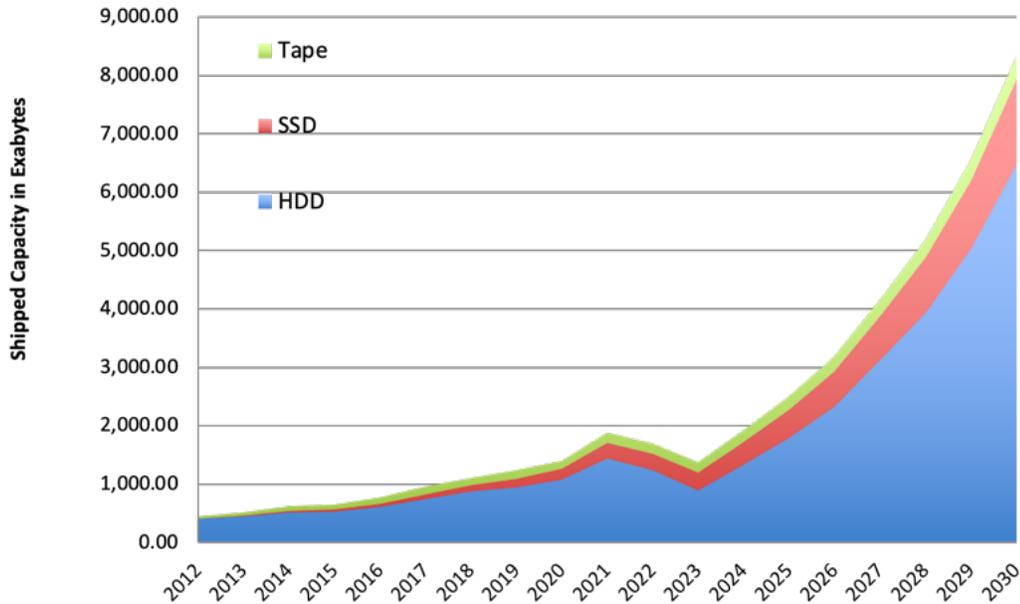


FIGURE 12. CAPACITY SHIPMENTS FOR MAGNETIC TAPE, SSDS AND HDDS



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 available in October 2025.

Table of Contents (Top Level, subject to change):

EXECUTIVE SUMMARY	8
INTRODUCTION:	10
WHY EMERGING MEMORIES ARE POPULAR	15
HOW A NEW MEMORY LAYER IMPROVES COMPUTER PERFORMANCE	41
THE CRUCIAL IMPORTANCE OF THE ECONOMIES OF SCALE.....	55
ONLY ROOM FOR ONE AT THE THE TOP	64
UNDERSTANDING BIT SELECTORS.....	67
CROSSPOINT ARRAYS	75
RESISTIVE RAM, RERAM, RRAM, MEMRISTOR:	81
FERROELECTRIC RAM, FERAM, FRAM:	101
PHASE CHANGE MEMORY (PCM):	120
MRAM (MAGNETIC RAM), STT MRAM (SPIN TRANSFER TORQUE MRAM).....	129
OTHER EMERGING MEMORY TYPES	167
LITHOGRAPHY:.....	175
3D MEMORY CIRCUIT DESIGN:	187
SUMMARY OF SOLID-STATE MEMORY & STORAGE TECHNOLOGIES.....	191
EMERGING MEMORIES AND NEW MATERIALS	197
EMERGING MEMORY PROCESS EQUIPMENT	198
MEMORY IS DRIVING SEMICONDUCTOR CAPITAL SPENDING	232
MARKET PROJECTIONS	235
ESTIMATES OF EMERGING MEMORY CAPITAL EQUIPMENT DEMAND	249
EMERGING MEMORY MANUFACTURING SUPPORT	265
COMPANY INFORMATION:	267

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
