



# Techwatch WG Restart

WLCG Grid Deployment Board Monthly Meeting  
February 14, 2024

Presenter: Shigeki Misawa  
Scientific Data and Computing Center (SDCC)  
Brookhaven National Laboratory

# Introduction

- HEP/NP experiments are heavily dependent on timely advancements in compute, network, and storage technologies to analyze their data.
- Changes in the trajectory of these advances will have an impact on current and future research
- The HEPiX Techwatch WG was created in 2018 to monitor trends in technology and provide periodic reports to the HEPiX community on the state of technology
  - Outgrowth of previous work by Bernd Panzer and Helge Meinhard
- The group has been dormant since the start of the Covid 19 epidemic with the last meeting in Feb 2020 and the last major presentation at CHEP 2019

# Restarting Techwatch

- New co-conveners: Andrea Chierici, Shigeki Misawa and Andrea Sciabà
- First meeting of a reconstituted Techwatch group in 17/01/2024
  - Introduce the purpose of the group and deliverables as originally envisioned by Bernd and Helge as well as the technologies of interest.
- Significant discussions on the nature of the working group
  - Mid term or long term view of technology ?
    - Long term is considerably more speculative. E.g., what's beyond CMOS
  - Component vs system level technology focus or both ?
  - Role of commercially available, for fee, research reports, if any ?
    - Previous group utilized publicly available information, e.g. news sites, press releases, freely available conference presentation
  - Extent of focus on market analysis in addition to technical overviews
  - Mechanics of how the group should be organized and operate.

# Areas of Interest of Previous Techwatch 2018 Group

- General market trends
- Server market
- CPUs and accelerators
  - Semiconductor process technology
  - Chiplets and die interconnect
  - x86, ARM and Power CPUs
  - GPU
  - AI/ML processors
- Memory
  - State of static and dynamic RAM
- Storage
  - Disk
  - Flash (solid state)
  - Tape
- Network
  - WAN
  - LAN

**Partitioning of subject areas for the new Techwatch group are being worked out**

# Proposed Subgroup/Subject Area Mapping

- Market Trends Subgroup
- Server and data center infrastructure
- CPUs, accelerators, and APUs
- Memories, busses and interconnects
- Online Storage (HDD/SSD)
- Offline Storage (Archive disk, tape)
- Network


# Deliverables and Audience

- Presentations at HEPiX meetings
  - Current expectation is one or more presentation from the working group at each meeting
  - Presentation topic - TBD
- Presentations at other venues
- Reports
  - “Executive” summaries
    - Short documents for consumption by project planners
    - Identify technologies worthy of more in depth investigation to the HEPiX board
  - In-depth reports
    - A longer document providing details on current state of technologies, major milestones and decision points in the evolution of technology
    - Targeted at the entire HEP/NP community
  - Live documents
    - Expectation is that all deliverables will be updated on a periodic basis

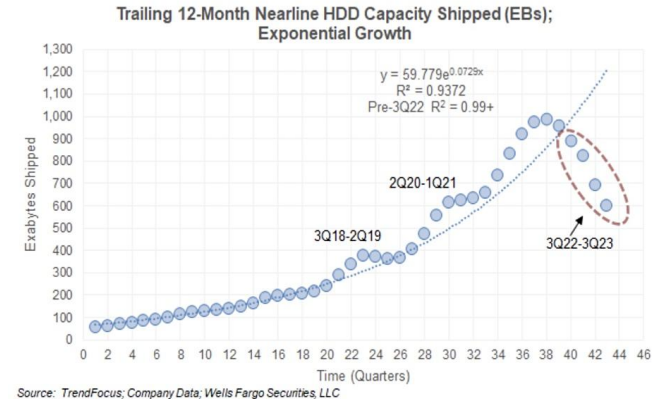


# Preview of Updates in HDD Storage

- HAMR drives have finally arrived
  - Seagate Mosaic 3+ 30TB HAMR drives shipping in Q1 2024
  - No word on HAMR products from Western Digital or Toshiba
    - Latest WD CMR drive is 22TB
- ~50% of exabytes shipped by Western Digital are SMR drives
- Majority of exabytes shipped and revenue are nearline HDD
  - But market was down for 2023



[Western Digital roller coaster continues as Seagate brings down the HAMR – Blocks and Files](#)



Source: TrendFocus; Company Data; Wells Fargo Securities, LLC

[Nearline drives will be last HDD holdout by 2028 – Blocks and Files](#)

# Preview of Updates in Flash Storage

- ~ 200+ Layer 3D NAND flash chips from all five major vendors
  - Roadmaps for more layer beginning to appear
- Viability of penta level cells unclear
- PCI-e Gen 5 SSDs now available
- Samsung and SK Hynix dominate
- Western Digital spinning off flash business
- Total revenues recovering from dip in late 2022/early 2023

3D Layer Cake

Micron		Samsung		SK hynix		SK hynix Solidigm		Western Digital/Kioxia		YMTC	
Generation	Layers	Generation	Layers	Generation	Layers	Generation	Layers	Generation	Layers	Generation	Layers
Gen 1	32	V3	48	V3	48	Gen 1	32	BICS 2	48	Gen 1	32
Gen 2	64	V4	64	V4	72	Gen 2	64	BICS 3	64	Xtacking 1 Gen 2	64
Gen 3	96	V5	96	V5	96	Gen 3	96	BICS 4	96		
Gen 4	128	V6	128	V6	128	Gen 4	144	BICS 5	112	Gen 3 Xtacking 2	128
Gen 5	176	V7	176	V7	176	Gen 5 (QLC 1H '23 & PLC?)	192	BICS 6 (Q1 2023)	162	Gen 2 2022 2H	196
Gen 6 (End 2022)	232	V8 (2022)	236	V7(2022 Q3)	238					Gen 6?	238
Gen 7	3xx	V9 (2024)	3xx	V8 (2023/4)	300			BICS 8 (2023)	>212	Gen 5	5xx?
Gen 8	4xx	V10	4xx	V9 (2025/5)	500+			BICS 9	300+	Gen 6	1,000?
Gen 9	5xx	V11	5xx	V10 (2030)	800+			BICS 10	400+		
		V7 (2030)	1,000								

[SK hynix breezes past 300-layer 3D NAND mark – Blocks and Files](#)




# Preview of Updates in Archive Storage

- Magnetic Tape
  - Strategy change at IBM for enterprise drives
    - TS1170 - 50TB / cartridge. No backward compatibility
  - IBM Diamondback “library in a rack” targets cloud and **Redundant Array of Independent Libraries (RAIL)**
    - Divergence of cloud and enterprise markets ?
  - Total LTO cartridges shipped has been declining, total exabytes shipped is flat
- Optical disk dead
  - Panasonic and Sony discontinued Archival Disc drives and libraries.
- On the horizon
  - Cerabyte “ceramic nano-memory” - Data etched in material via laser or particle beam
  - Folio Photonics - No news since 2022




LTO MEDIA UNIT SHIPMENTS\*



[https://www.lto.org/wp-content/uploads/2022/04/LTO-Ultirum-2021-Media-Shipment-Report-Slides\\_FINAL-1.pdf](https://www.lto.org/wp-content/uploads/2022/04/LTO-Ultirum-2021-Media-Shipment-Report-Slides_FINAL-1.pdf)

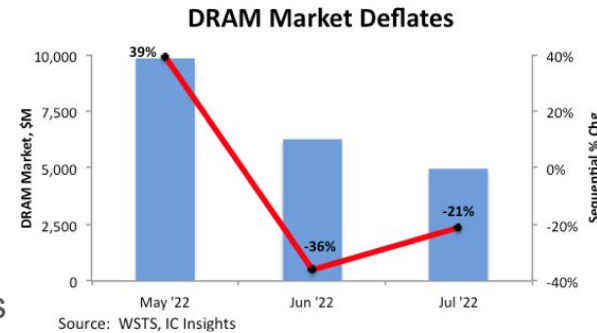
TOTAL CAPACITY BY CY\*\* (EB COMPRESSED)



[https://www.lto.org/wp-content/uploads/2022/04/LTO-Ultirum-2021-Media-Shipment-Report-Slides\\_FINAL-1.pdf](https://www.lto.org/wp-content/uploads/2022/04/LTO-Ultirum-2021-Media-Shipment-Report-Slides_FINAL-1.pdf)

# Previews in Logic and Memory Since 2018

- All major logic foundries (TSMC/Samsung/Intel) now have EUV in production
  - Intel last with “Intel 4”
  - Next transition is high NA EUV or multi pattern EUV
- Advanced packaging increasingly important
  - Chiplets/die connectivity
  - Power delivery and cooling
- DRAM memory
  - CPU’s transitioning from DDR4 to DDR5 memory
  - HBM3 introduced in 2022. Higher bandwidth “HBM3E” products introduced in 2023
  - All major manufacturers except Micron have transitioned to EUV
  - DRAM market recovering from collapse in late 2022/early 2023



[DRAM Market Deflates Cyclical Downturn Looms \(icinsights.com\)](https://www.icinsights.com)

# Preview of Advances in CPUs Since 2018

- Segmentation in CPU product lines
  - HPC (higher frequency)
    - Intel “Granite Rapids”
    - AMD “Genoa” (Zen4)
  - Cloud (more cores)
    - Intel “Sierra Forest”
    - AMD “Bergamo” (Zen4c)
  - Edge (lower power)
- ARM resurgence
  - Ampere Ultra/Ampere One
  - NVidia Grace
  - Amazon Graviton 4
  - Microsoft Cobalt 100
- Single package GPU/CPU systems (“APUs”)
  - AMD Mi300A, Nvidia Grace Hopper
- In package memory
  - AMD 3D V-Cache (SRAM on die)
  - Intel Xeon Max (HBM2E DRAM)
  - Nvidia Grace (LPDDR5X DRAM)




# Previews in Accelerators Since 2018

- GPU accelerators and related technology
  - Nvidia H200 released in 2023
  - AMD Mi300X released in 2023
  - Intel Data Center Max in 2023
  - Broadcom announces support for AMD Infinity Fabric in its next generation PCI-e switch chips. This allows interconnection of more AMD GPUs, similar to NVSwitch for Nvidia GPUs.
- AI Accelerators
  - Captive processors from Amazon (Trainium2), Google (TPU v5) and Microsoft (Maia 100)
  - Intel Gaudi2, SambaNova SN40L, and others.


# Preview of Network Technology

- Transition to 400GbE (4x100Gbs) in progress
- 800GbE (8x100Gbs) specification released in 2020
- Cloud adoption of higher bandwidth Ethernet outpaces rate in the enterprise
- Ultra Ethernet Consortium formed to make Ethernet more competitive with Infiniband for AI workloads
- Co-packaged optics in the works
  - Reduce power consumption

## PATH TO SINGLE LANE




[2023 Ethernet Roadmap - Ethernet Alliance](#)



Cloud Link-Speed Forecast

**Cloud:** All about 100G+



Enterprise Link-Speed Forecast

**Enterprise:** Mix of 10G, 25G, 100G



# Some Technology Open Questions

- What is the future \$/TB for HDD with the arrival of HAMR and will multi-actuator drives change the trajectory ?
- How healthy is the tape market and will the public clouds use of tape make things worse or better ?
- Are APUs, single package GPU + CPU, a better fit for HEP/NP compared to discrete GPU + CPU.
- How viable are both fully custom and Neoverse derived ARM CPUs in the open market given the clouds use of internally sourced CPUs ?
- Is there a place for AI/ML processors in the open market?



Goal of the TechWatch WG is to answer these and other questions about technology that are of interest to the community

# Next Steps

- Develop schedule for future presentations
  - HEPiX meetings (next HEPiX is April 2024/Paris)
  - Identify other venues
- Determine collaborative tools needed to create and maintain deliverables
- Organize subgroups and begin work on gathering information on subject areas
- Work on updating the TechWatch WG website
  - [Technology Watch Working Group \(hepix.org\)](https://w3.hepix.org/techwatch.html) (<https://w3.hepix.org/techwatch.html>)

# Call for Participants

- We need people to volunteer to lead subgroups or contribute information.
- Meeting times and agenda are posted in CERN indico.
  - <https://indico.cern.ch/category/10621/>
- Techwatch mailing is available for communication.
  - [hepix-techwatch-wg@hepix.org](mailto:hepix-techwatch-wg@hepix.org)

**NEXT Meeting is Wednesday February 21 @ 16:00 CET/10:00 AM EST**

**See you there !**

