

# Recap of 2023 Storage Technology Showcase

## About storage professionals and vertical market failure

By Philippe Nicolas | March 24, 2023 at 2:01 pm

The [Storage Technology Showcase](#) finally restarted on March 13-15, 2023 with physical presence and took place last week in Tucson, AZ. Tucson is also the Mecca for IBM tape storage.

This event was the opportunity to reconnect with vendors but also with peers from largest data sites such US research centers like Los Alamos, Lawrence Livermore, or NOAA to name a few.

Companies, exhibitors, and sponsors include Alliance IT, Brad Johns Consulting, Coldago Research, Comport, Fujifilm, Globus, Hammerspace, Horison Information Strategies, HPCwire, HPE, IBM, Insurgo, Nyriad, Spectra Logic, Twist BioScience, Vast Data and attendees were from AMD, DNAalgo, Folio Photonics, Grau Data, HPSS organization, Phison, StarFish and Versity.



The topic of the year was about “Engaging the Rising Generation of Storage Professionals” and we heard from several presenters interesting stories related to storage profiles. In fact, there is no special studies dedicated to storage, people learn storage following more classic studies around OS, architectures ... becoming system admin, architect ... and extend those positions with vendors’ trainings and certifications.

For sure, industry bodies, standards, white papers, and other independent sources feed this learning curve but clearly something must be done. Some organizations understand this erosion of skills, with teams soon retired and new technologies coming so they put in place some internships to facilitate the transition or other special projects. We invite the reader to read a few key presentations delivered on that topic during the conference, but all converge on that fact that IT professional curriculums and studies must be improved and extended with storage content.

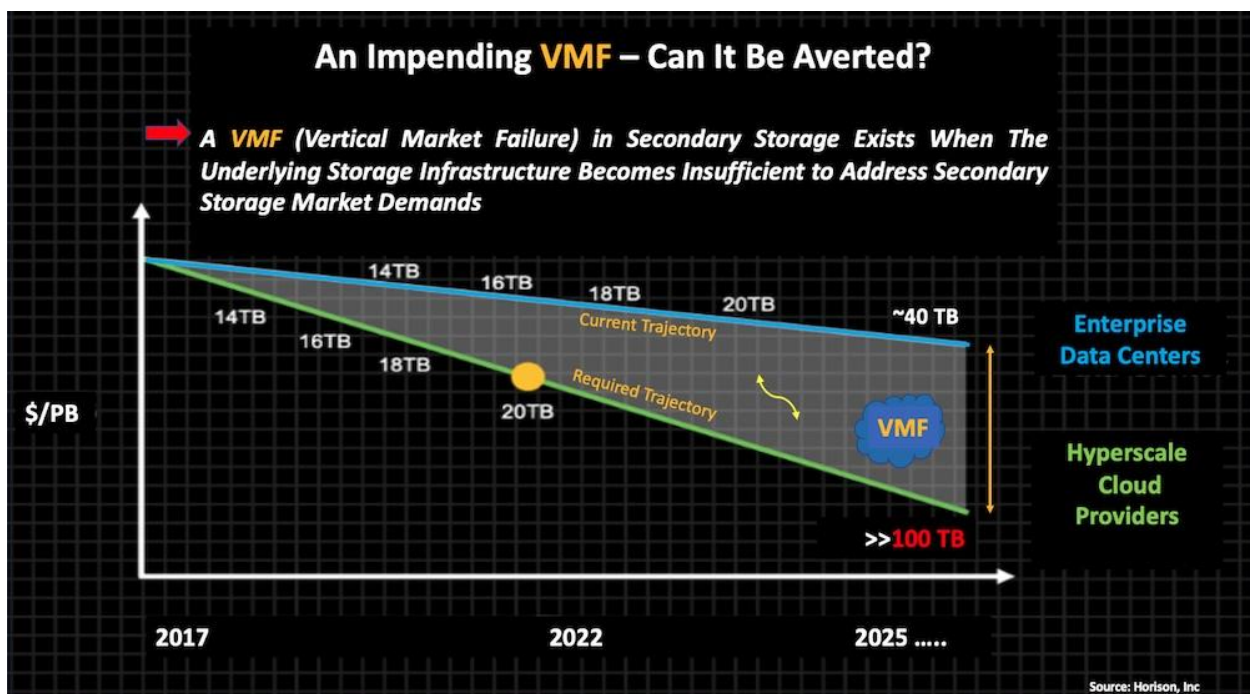
## Secondary Storage Market Scenario HSDCs and Cloud Providers are at the VMF Epicenter

- HSDC's secondary storage and archival needs are growing much faster and have become significantly different than traditional data center archive needs – *the Hyperscale Effect*.
- The HDD and tape development roadmaps are slowing.
- Relaxed LTO roadmap specifications and schedules are pushed out delaying future tape advancements.
- HDD roadmap faces challenges to increase capacity, any performance gains (IOPs) appear minimal.
- IP litigation risks, as evidenced by recent lawsuits, further impact and delay roadmap progress.
- HDD (3) and tape drive (1) supplier consolidations are concerning trends.
- Neither HDD nor tape is currently positioned by itself to effectively meet the enormous future performance and capacity demands.
- HSDCs leverage their bargaining and buying power to drive down prices impacting vendor margins, R&D.
- The secondary storage market is left with incremental investments in existing technologies and moves slowly compared to soaring secondary storage demand.

➔ Without Changes, the Extensibility of the Current Secondary Storage Model to Effectively Meet These Demands is in Jeopardy Accelerating a VMF.

Source: Horison, Inc.

The second obvious topic is related to the VMF – Vertical Market Failure – exposed by Fred Moore in his keynote. In a nutshell, hyperscalers need goes faster than the industry pace and we realize that the HDD and tape roadmaps are slow and won't be able to address these gigantic secondary storage demands. At the same time, these players put pressure on prices and therefore impact seriously vendors' capacity to innovate. And the key question remains open, what kind of technology will be able to address these challenges in secondary storage? Some are promising but not ready yet such DNA data storage, the challenge is open ...



Among other information gathered during the conference, Fujifilm stopped Object Archive, its S3-to-Tape product. It's a surprise as that segment is pretty hot on the market with now 8 vendors Grau Data, Nodeum, PoINT Software & Systems, QStar, Quantum, Spectra Logic, StrongLink and XenData promoting a kind of modern VTL with a remote access.

We saw lots of content around tape and archiving with very large data sites with some impressive migration projects. We all know that the tape market is active with tape libraries from HPE, essentially BDT models plus the resale of Spectra Logic ones, and also strong presence from IBM, Quantum and Spectra Logic, equipped with IBM TSxx family drives, Oracle or LTO ones. As a reminder, IBM is the only manufacturer of LTO drives equipped with Western Digital head and only Fujifilm and Sony produce cartridges. The LTO roadmap has been extended up to LTO-14 still without any timing.

Hyperscalers really drive the tape library segment with the need of new design, density and modularity. Quantum has introduced the i6H already adopted by a few hyperscalers and more recently during the OCP Summit in October, IBM announced Diamondback with interesting specifications, in fact 27PB raw with 14 LTO-9 drives and 1548 tapes.



The open-source software GUFU (Grand Unified File Index) developed at Los Alamos and downloadable [here](#), continues its journey. The laboratory also started a new study around Seagate Kinetic drives for a second life, for sure you remember what Seagate developed around 2015 with these IP/Ethernet drives. The idea here is to off load CPU based workloads to a pool of drives equipped with Cortex ARM CPUs and run locally SQL queries federated at the client level. These drives exposed a block interface, NVMe in fact, and are integrated to the application

via ZFS. At the same time, they embed DuckDB, the companion of SQLite but for OLAP, and can receive SQL queries to process on the ingested data thanks to the dual host path and the special data format allowed with Parquet. This represents a new example of computational storage that goes beyond encryption or compression ...

Vast Data confirmed that the total capacity deployed passed 8EB today with one customer having 350PB+. It is good for the market to see a vendor developing a new architecture – a shared everything one – and validates NFS at scale, as many other players claimed that the issue was and is NFS. The opposite demonstration is made, of course with key elements in the equation such, persistent memory, NVMe, its network companion and fast connectivity. In other words, Vast Data, who is not a HPC storage vendor, supports and delivers great results for HPC workloads with NFS. We saw Vast Tables; it reminds me what LizardFS or RozoFS did 10 years ago but clearly with a modern approach.

We got interesting talks around RAID-6 and one of them was presented by Nyriad. The company builds storage subsystem based on HDDs coupled with RAID controllers with GPU-based erasure coding. It delivers large stripe over 100 or 200 drives with 20 parity drives for a better capacity and protection efficiency versus the multiplication of smaller RAID-6 groups.

We also discovered [Insurgo Media Services Limited](#), a Welsh tape services company, who developed dedicated tools to safely erase tapes or even kill tapes with a deep traceability approach.

And to conclude, we had a fun moment when Jakob Farmer from Starfish tried to destabilize David Flynn with uncomfortable questions at the end of his presentation related to data orchestration promoting pNFS among other things.

**Read also :**

[Special Report on Storage Technology Showcase 2020](#)

Parallel file system, tape, migration, archiving, bit of cloud and some NFS coming

[Top News](#) | [By Philippe Nicolas](#) | March 30, 2020 | News



**Storage Newsletter**

This placement is available for ad banner

Click on "Advertise with us" for more details