Storage
IT Technology and Markets, Status and Evolution

Storage Categories

Three main areas, very few companies driving the market.

Tape:
IBM (drives) Fujitsu, Sony (tape media) ongoing patent 'war'

***NAND flash memory* (25% yearly production used for SSDs):**
Samsung, Toshiba, Western Digital, Micron, SK Hynix, Intel

Hard disks:
Western Digital, Seagate, Toshiba

→ Increasing overlap between HDD and SSD provider

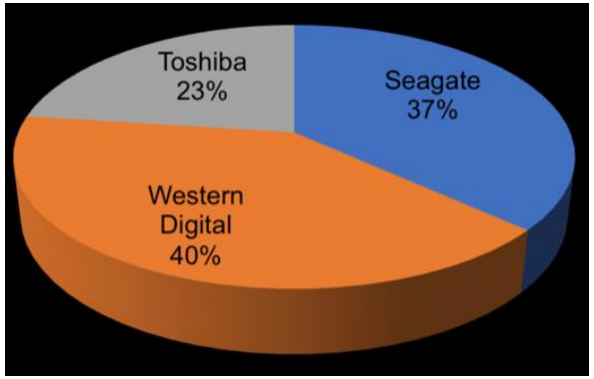

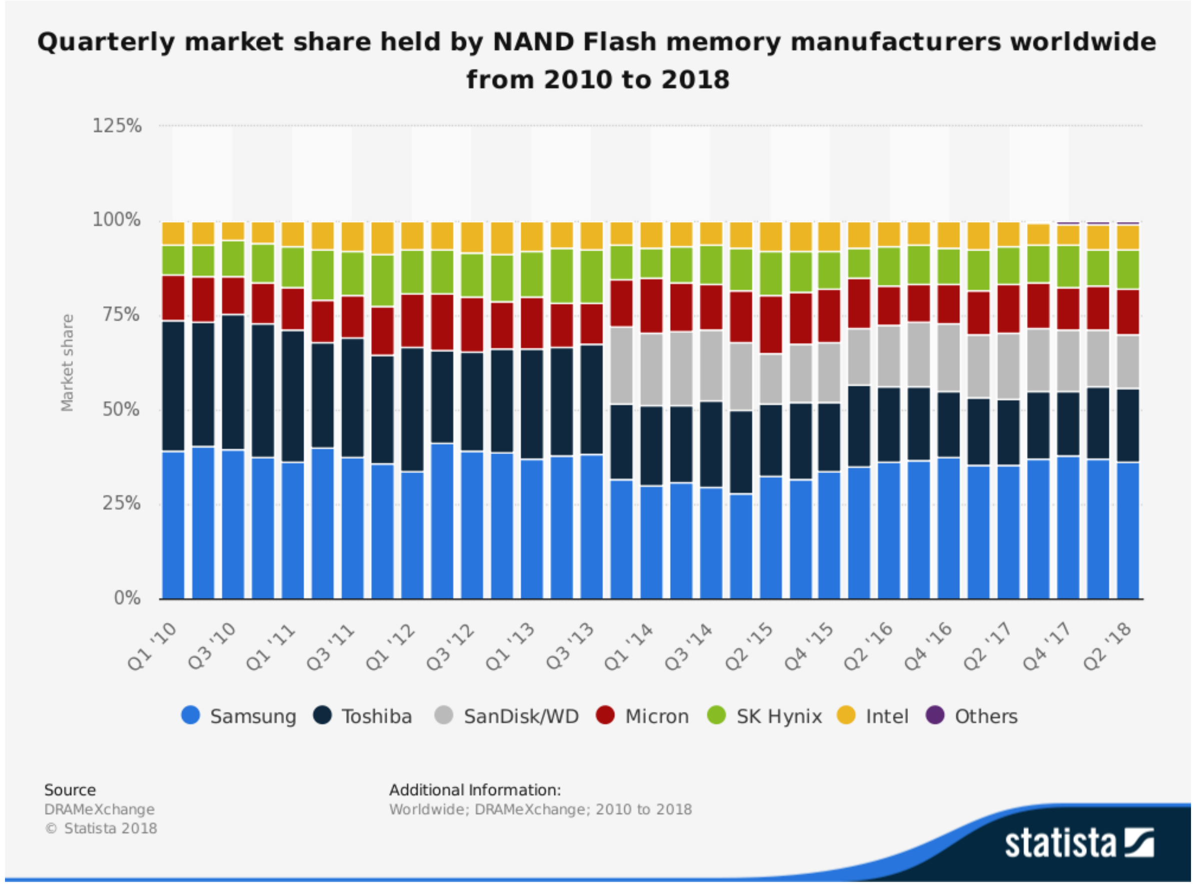

Indirect concentrations:
Limited suppliers for wafers, rare-earth magnets, fab equipment,…..
e.g. One supplier for >80% of ALL HDD spindle motors

Tape Storage I

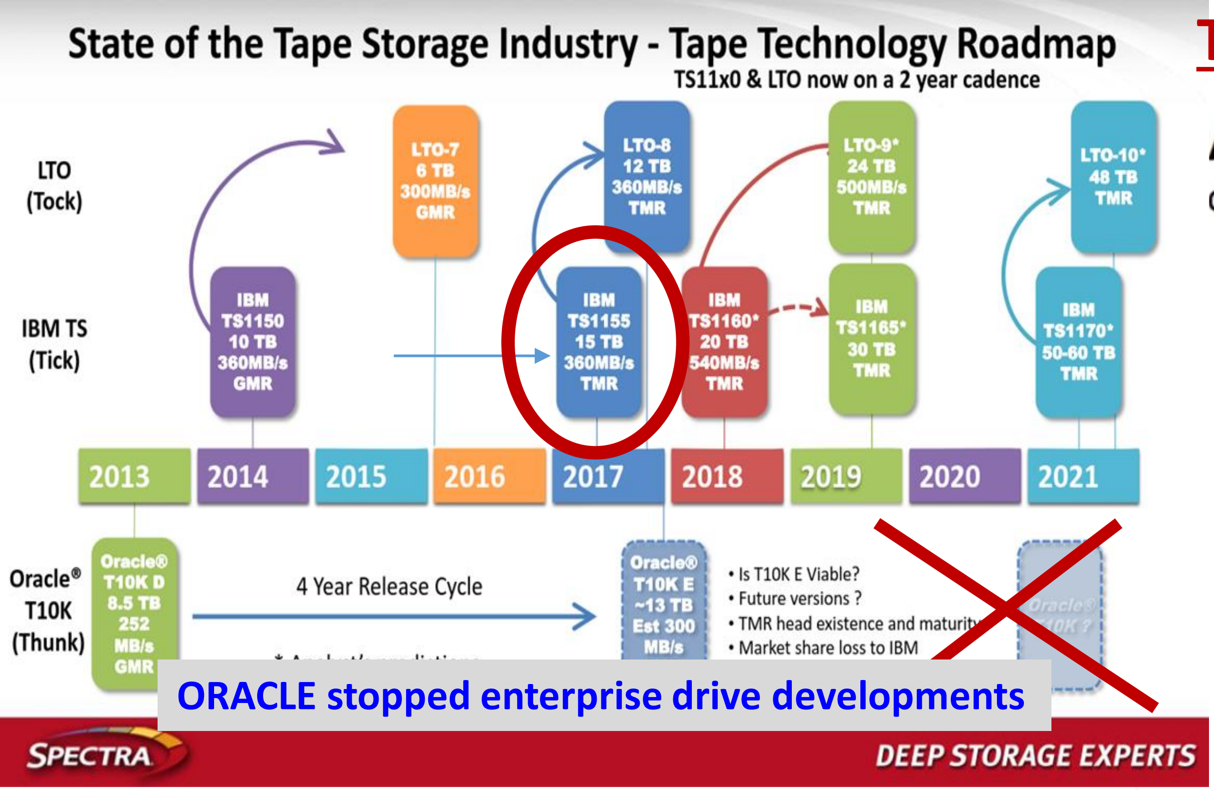

Quite some headroom for density improvements, x10 compared to HDD

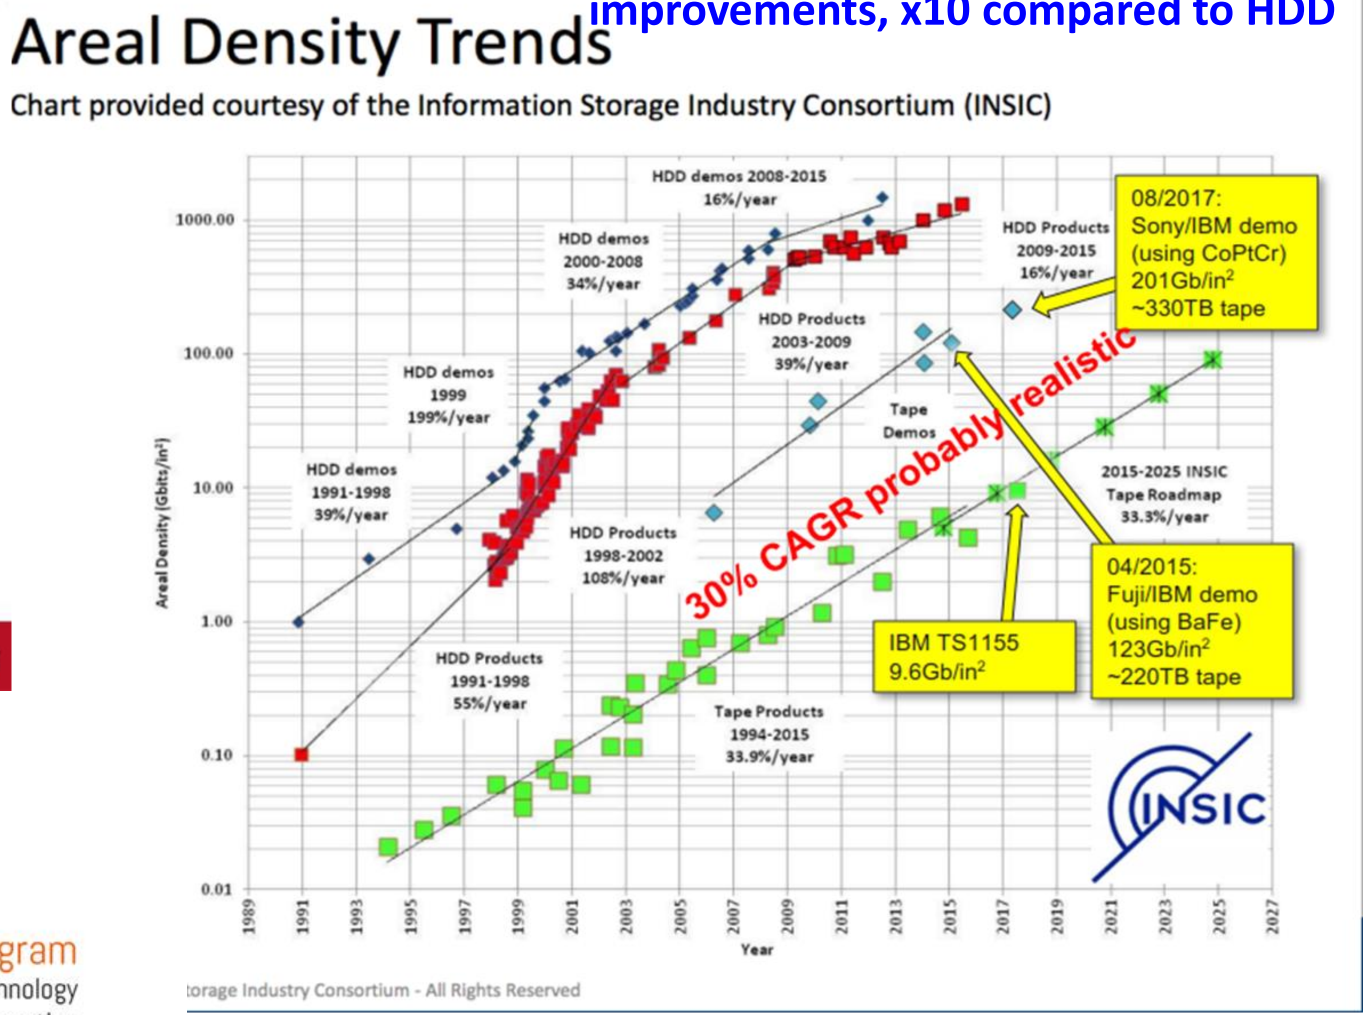

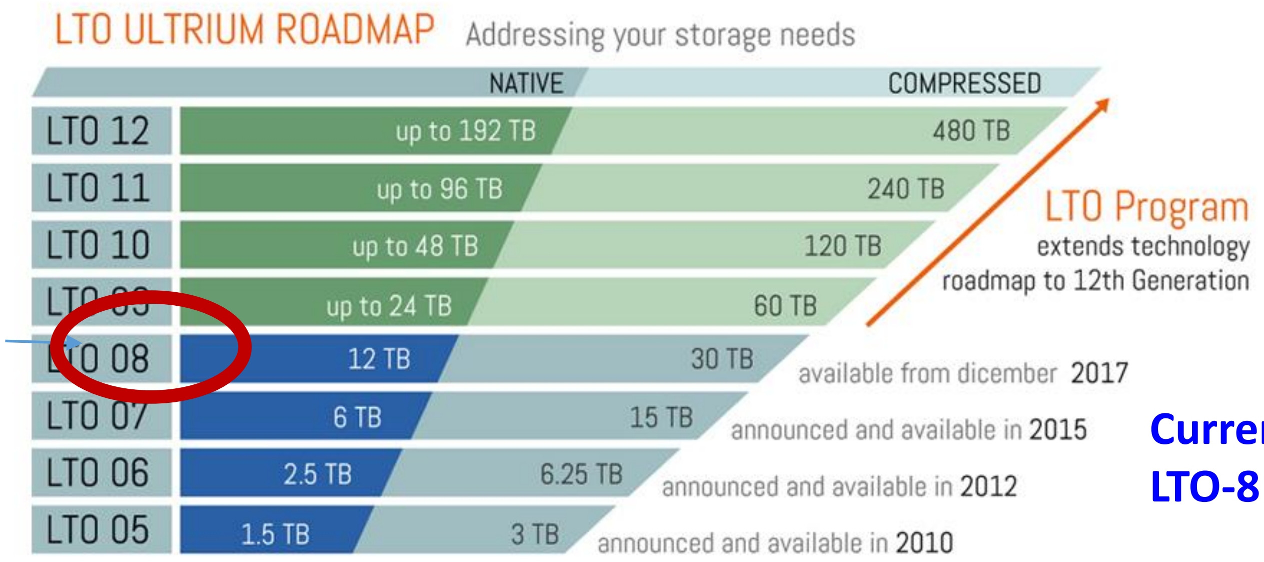

Technology change to Tunnel Magnetoresistive heads (used already in HDDs) for IBM TS1155 and LTO-8

Current generation
LTO-8 (12 TB) , TS1155 (15 TB)

Tape Storage II

Unit Shipments: Calendar Year

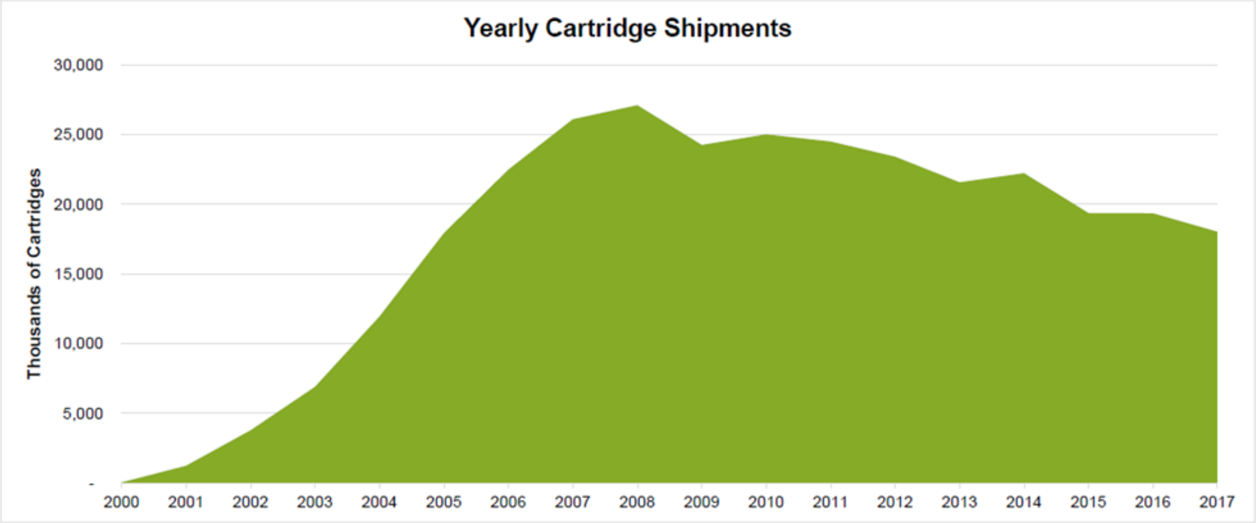

LTO tape market domination >95%
Enterprise tapes ~4%

44 EB of tape media in 2017 compared to 750 EB HDD
Linear increase in EB sold per year

Declining media shipment since 10 years

Factor 2 decrease in #drives sold over the last 4 years
Only IBM left for LTO and Enterprise drive heads

Only two suppliers of media: Fujifilm and Sony
Fujifilm only supplier in the US (patent 'war')
(currently heavy shortage of LTO-8 tapes in the US)

Total Capacity Shipped: Calendar Year

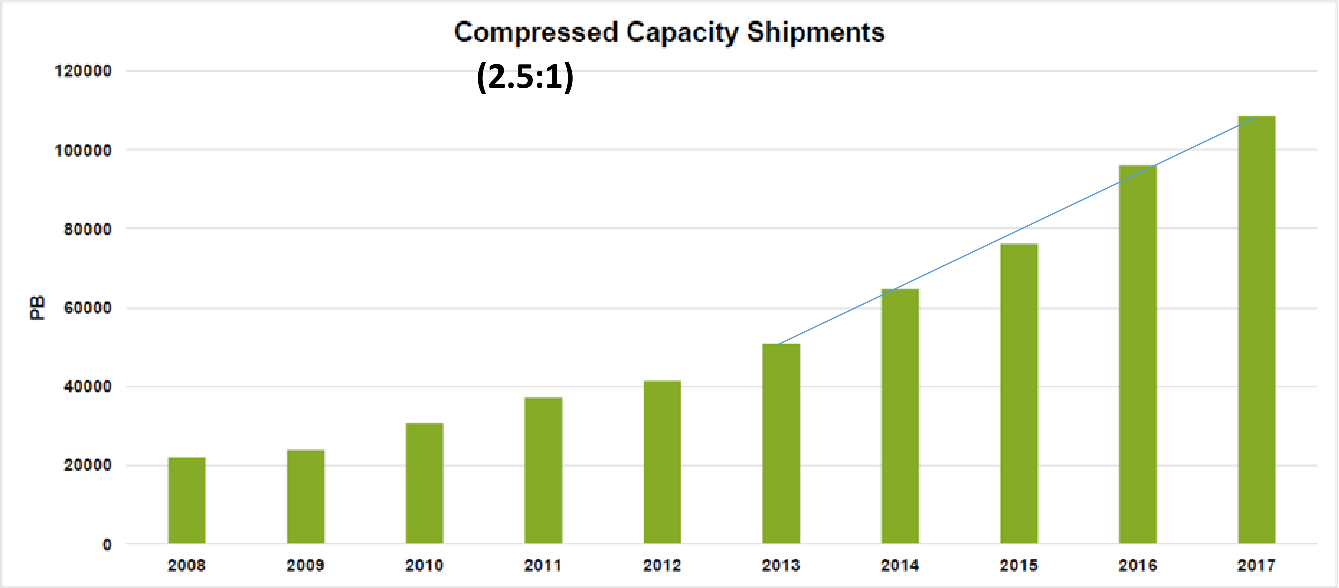

NAND Storage I

NAND storage is based on a charged trap flash architecture
== storing electrons in a SiN layer

Small structure sizes leads to higher error rates, less erase cycles
→ Move from 2D to 3D with an increase in structure sizes (20nm → back to 40nm

3D prototypes in 2007 -- 2017 80% of all NAND production uses 3D

64 layers in the market, 96-layers production started

Multi-level cells provide cheaper storage, but less erase cycles

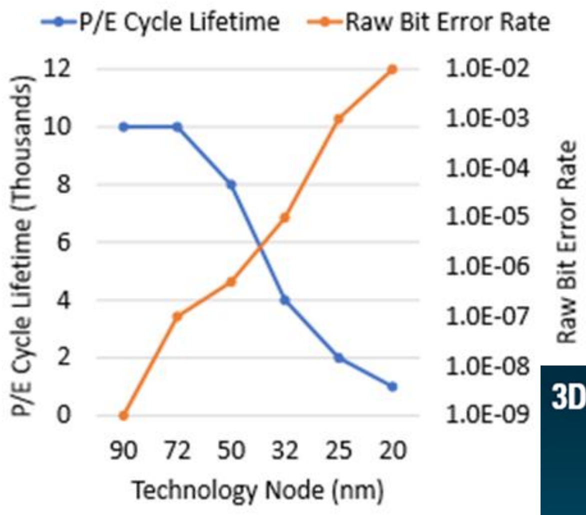

(a) Decreasing flash reliability

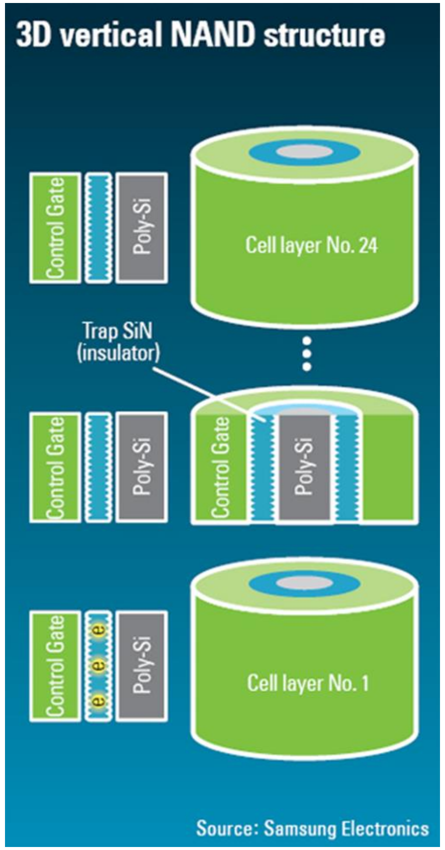

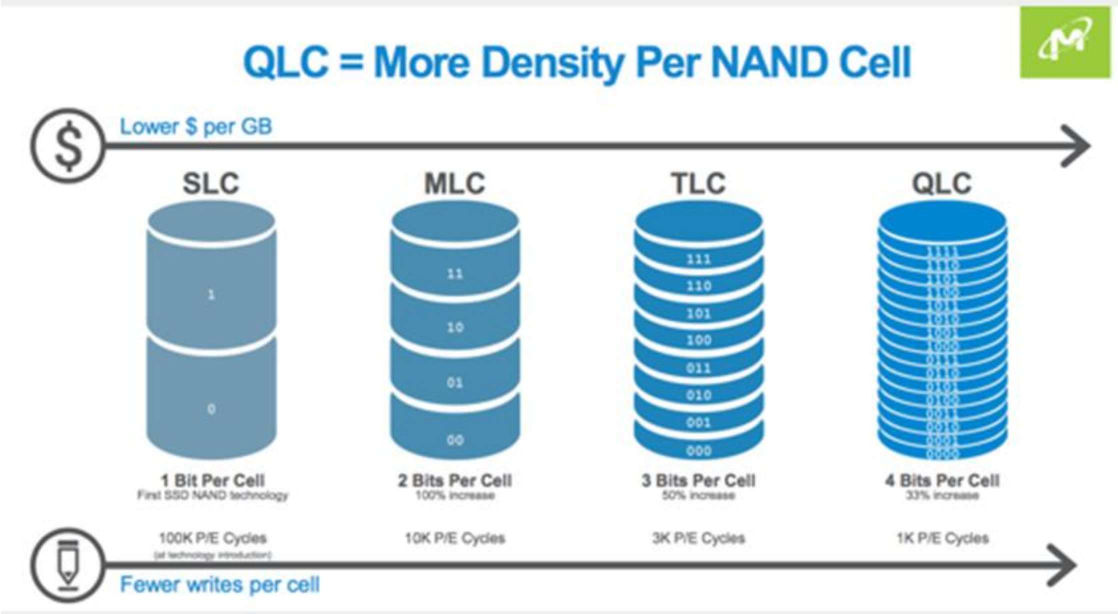

NAND Endurance – Figures courtesy of Micron.

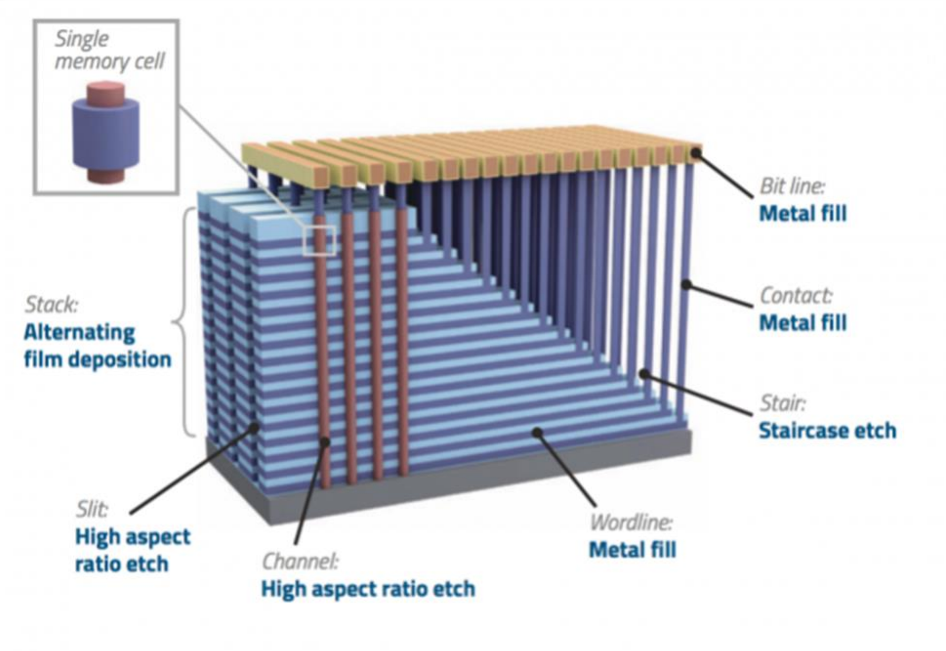

FIGURE 1. 3D NAND architecture showing some of the most challenging and critical deposition and etch processes.

NAND Storage II

- ~60 B$ market
- NAND prices increased over the last 18 month, high request for smartphones and SSDs (Apple buys 20% of the world-wide NANDs),
- 2018 price trend seems to change now, -10% for Q3/Q4 expected → 3 new Chinese fabs will start production this year
- 4-bit cells are now feasible with 3D : ECC code easier with 2D cell size increased ; first products by Intel+Micron
- investment 3D fabrication process is up to 5x higher than 2D,
  ~10B$ investments needed for new fabrication facility
- Technical challenges: > 64 layers show exponential scaling problems (current density, cell uniformity)
  a wafer stays up to 3 month in the fab before the >100 defect-free layers are done
- Density improvements are now linear, adding 8/16/32 layers

Figure: NAND Flash Factories Map in 2020

Source: DRAMeXchange, Jan., 2018

Hard Disk Storage I

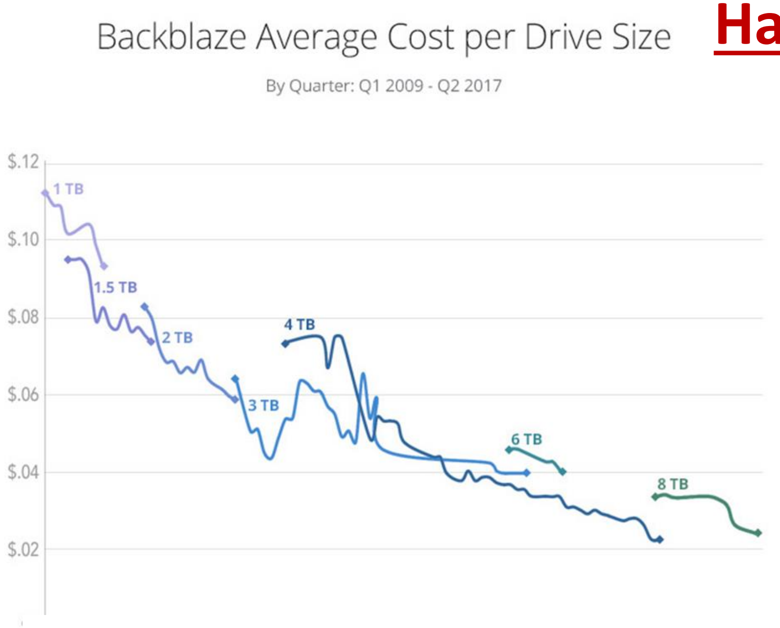

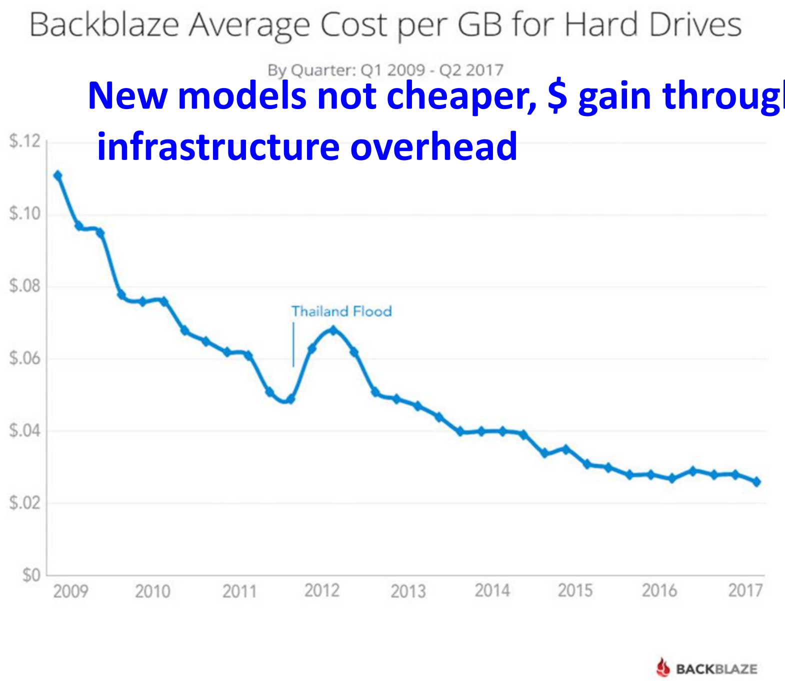

New models not cheaper, $ gain through less infrastructure overhead

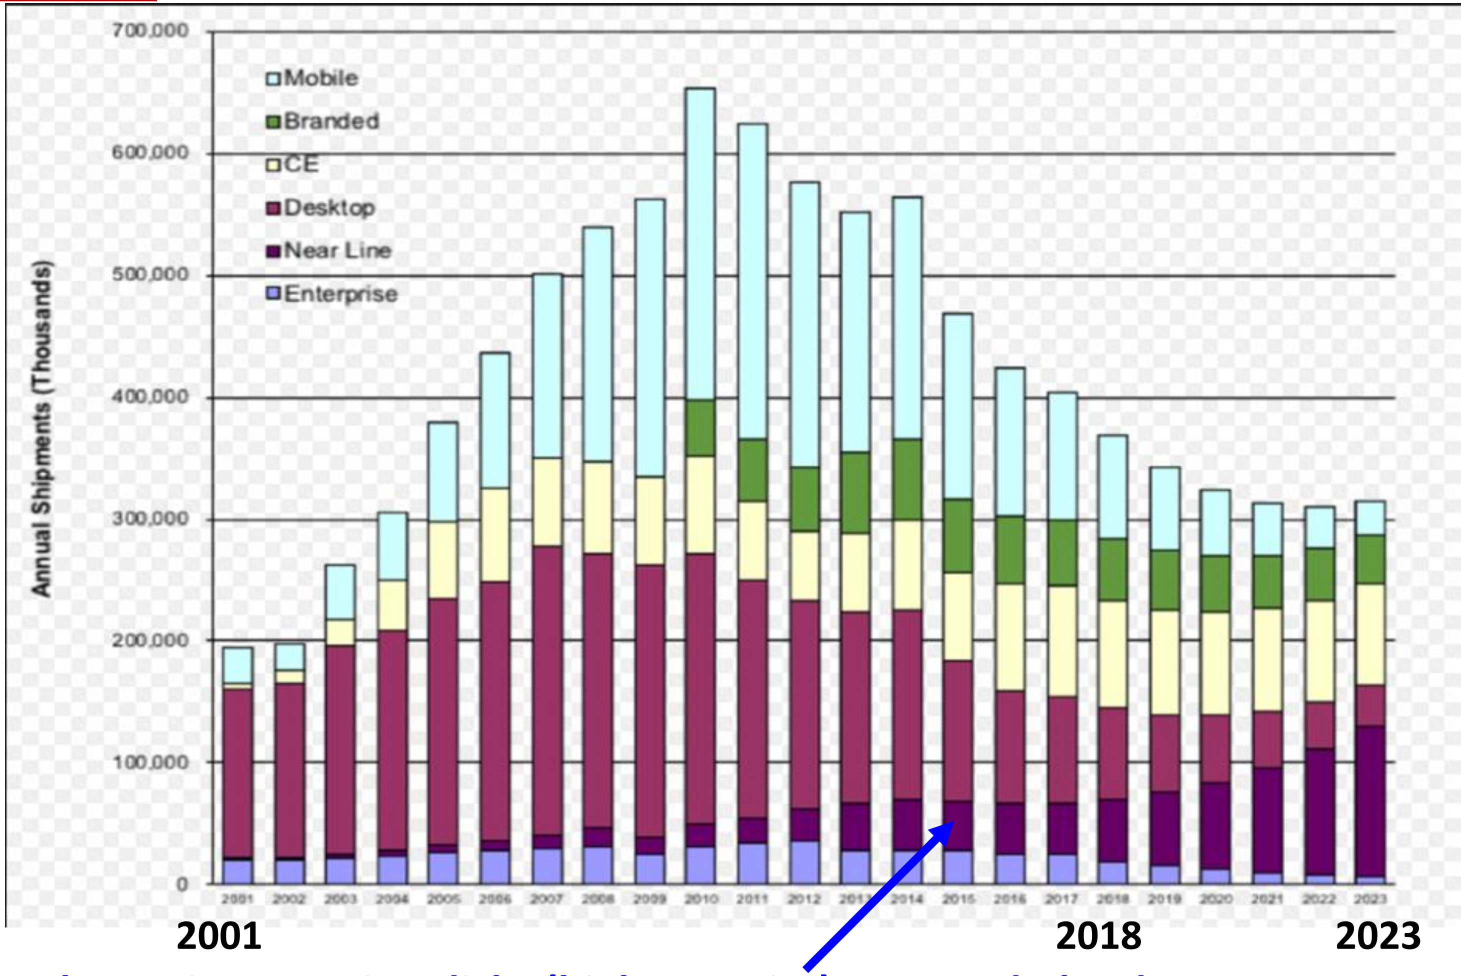

Only growth rate in Near Line disks (high capacity), HEP and Cloud Storage area (50% of 800 EB delivered, 15% of total units shipped, >40% of revenues)

Desktop, Mobile, Enterprise HDDs replaced by SSDs

Price/space evolution flattening, Seagate and WD are closing fabs

Hard Disk Storage II

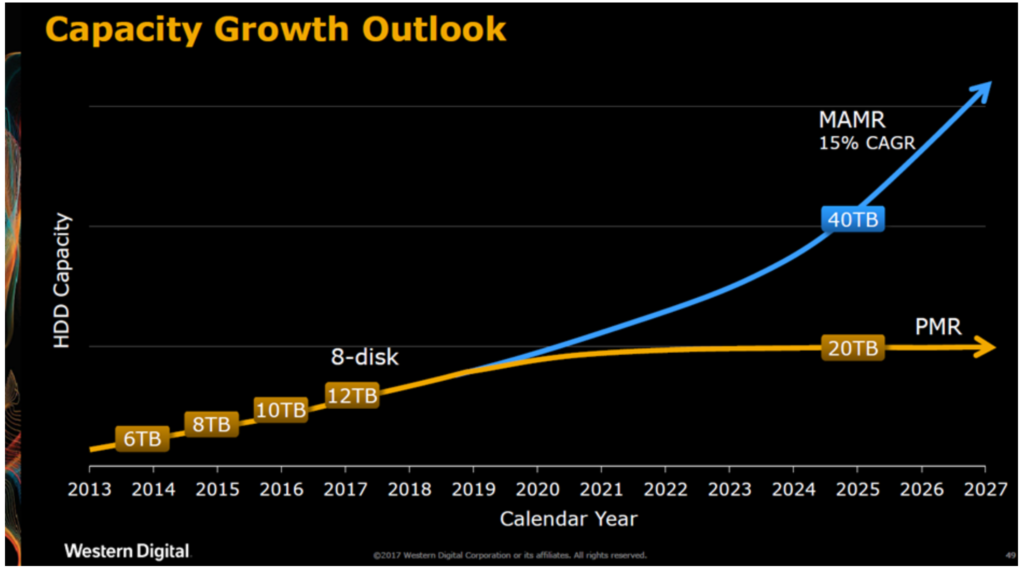

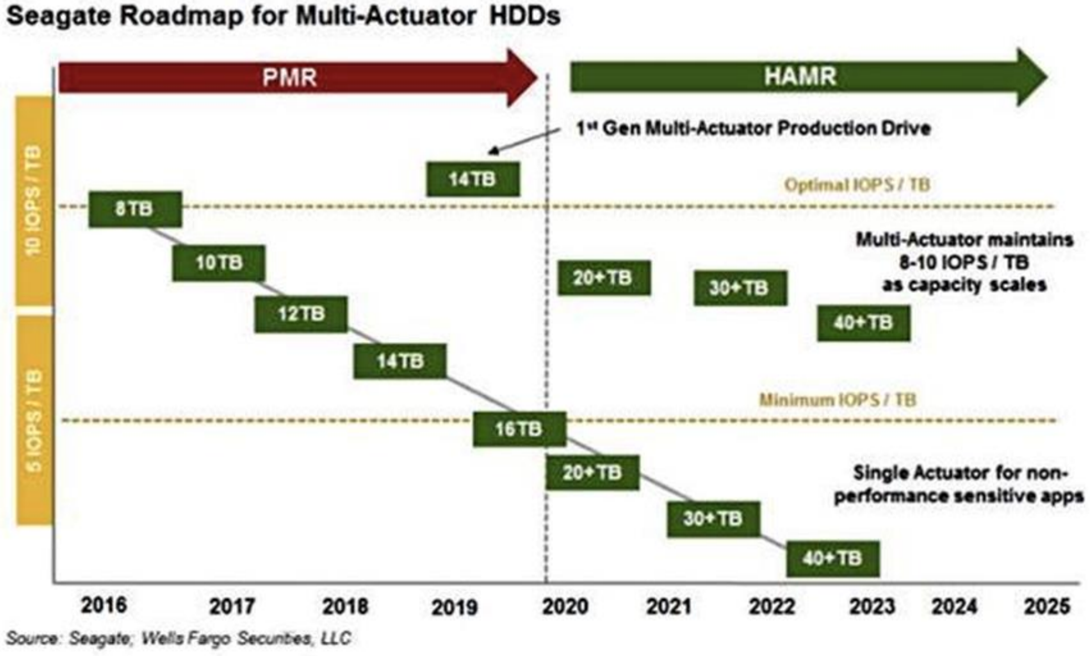

9 platter in one disk 14 TB capacity today He filled
Max with SMR is probably around 20 TB per HDD

The market introduction of these new technologies has already 'slipped' by several years (complicated, expensive)

Seagate HAMR first products now in 2020

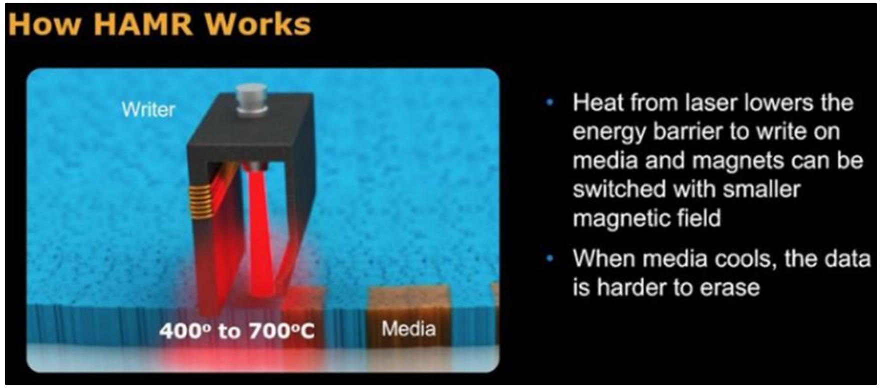

Western Digital new density approach: MAMR production in 2019

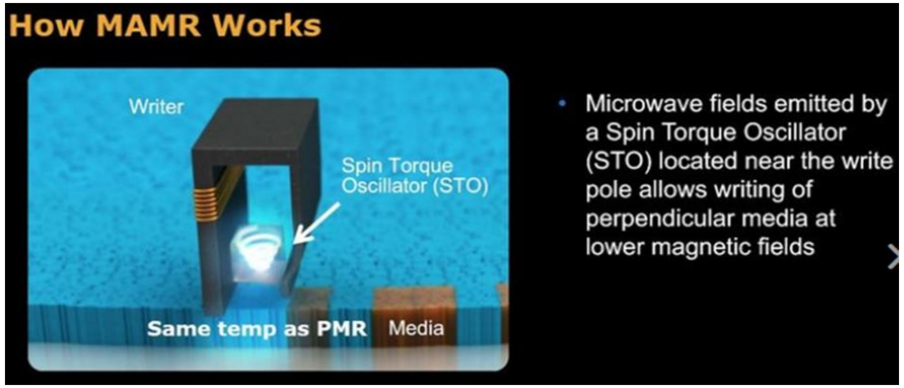

Seagate: multiple actuators per HDD to keep IOPS/TB constant

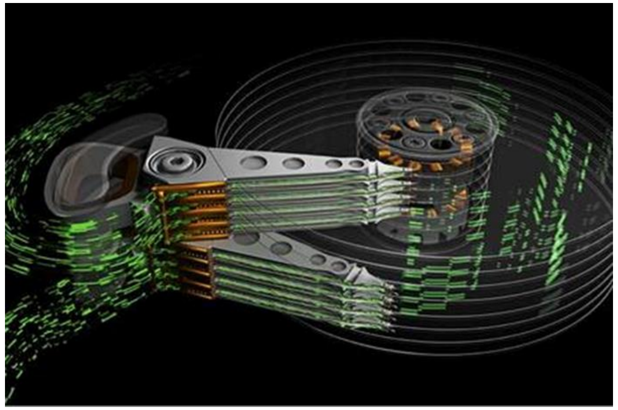

Solid State Disk Storage

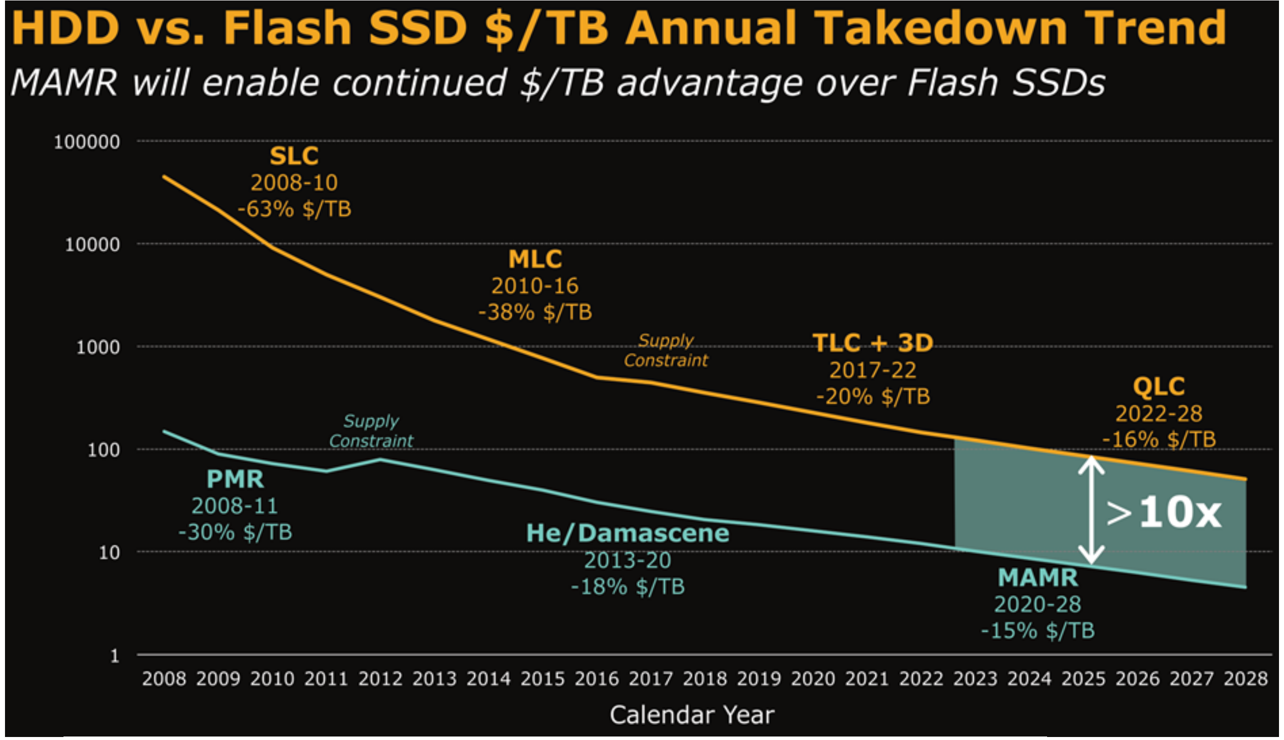

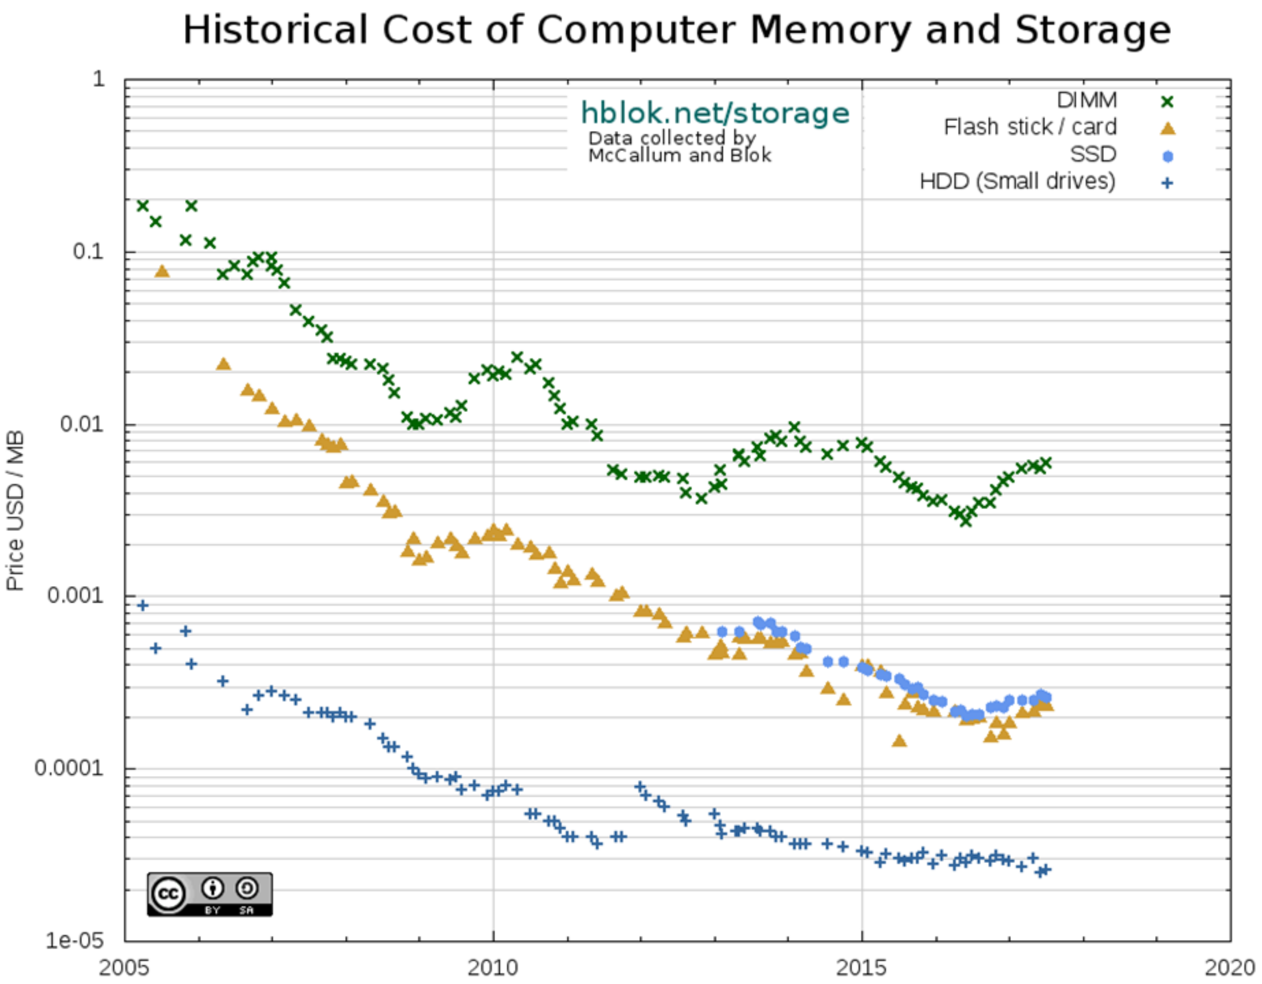

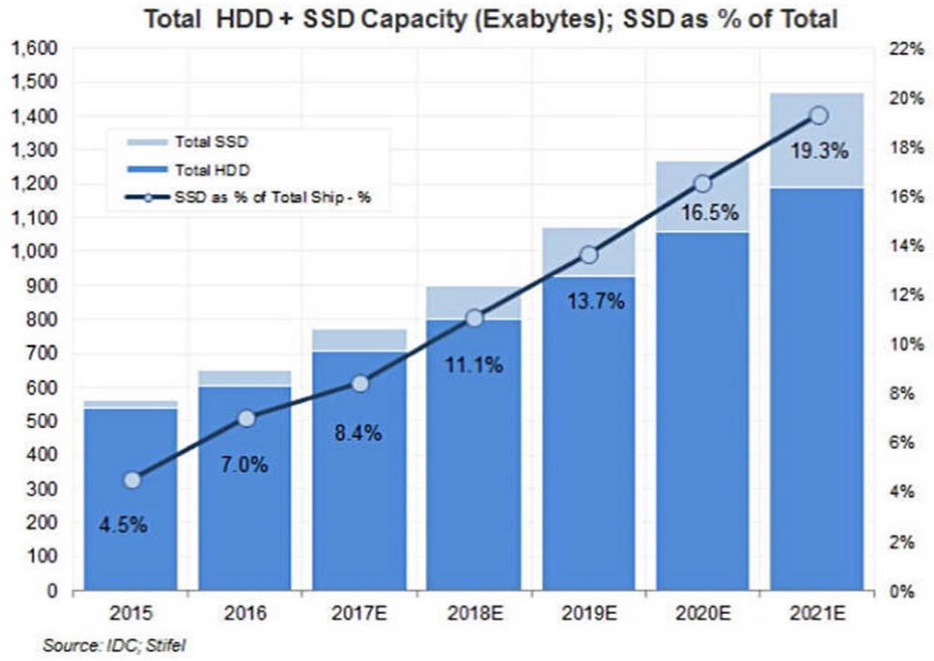

SSD versus HDD, Price difference in capacity drives will stay high for the foreseeable future

(depends on the NAND fab and market evolution in China….)

Slowdown of yearly price improvements in all areas

Storage Comparisons I

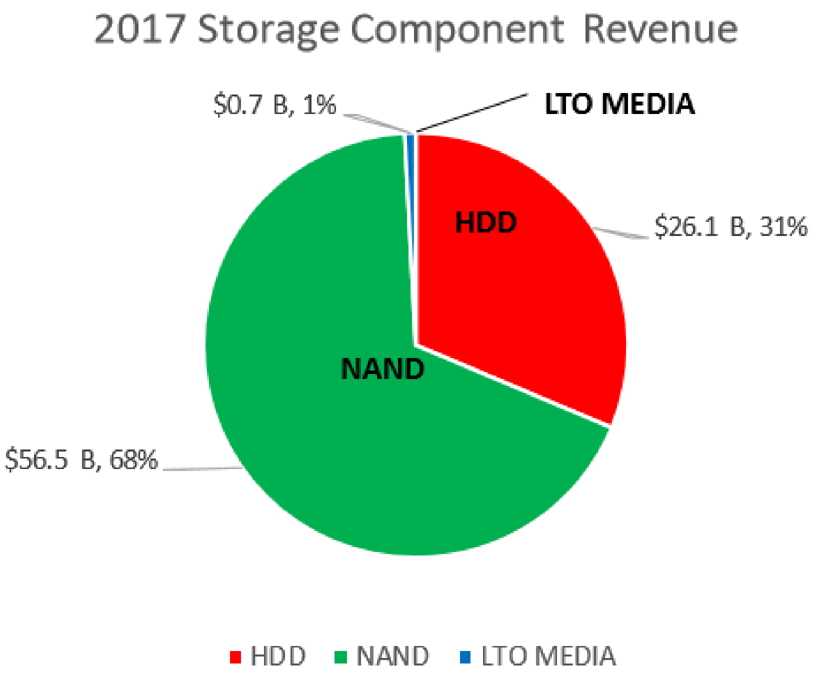

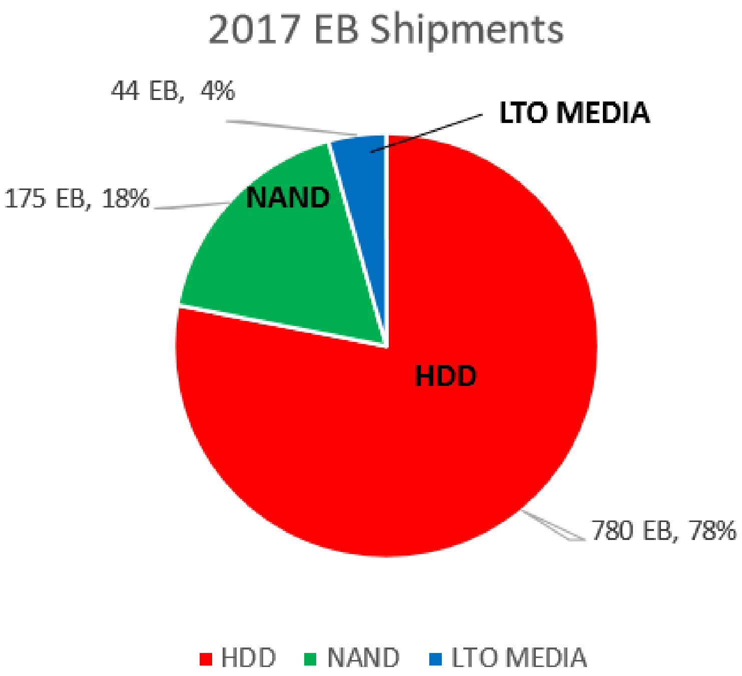

R. Fontana, G. Decad IBM Systems
5/15/2018

10 Year Storage Landscape

25% of the NAND capacity is for SSDs

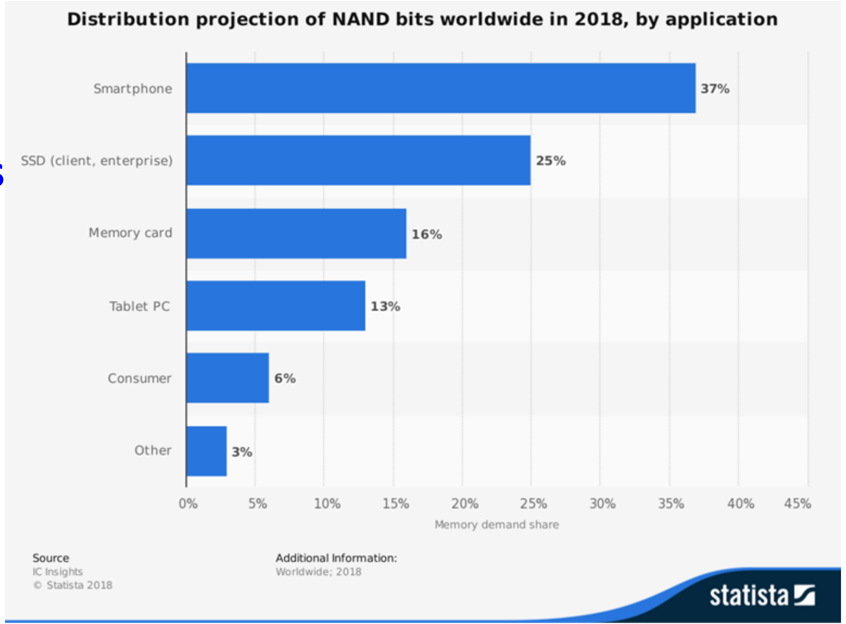

Revenues for HDD and tape are steadily decreasing,
while NAND revenues are increasing over the last years.

Storage Comparisons II

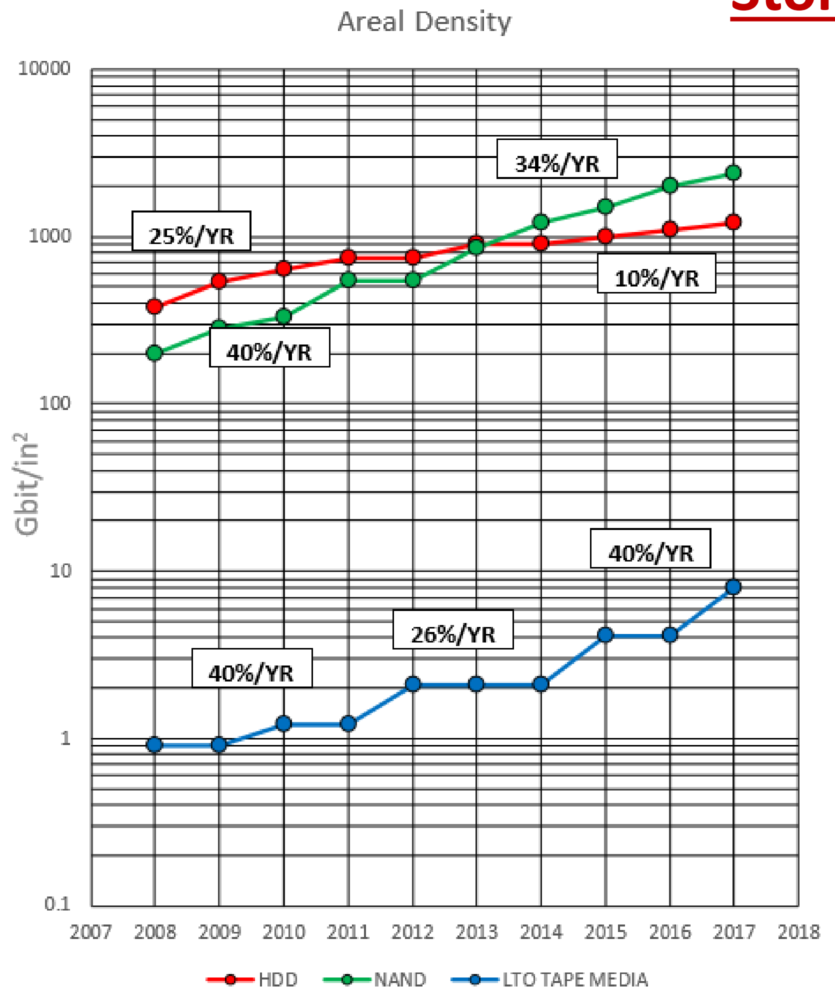

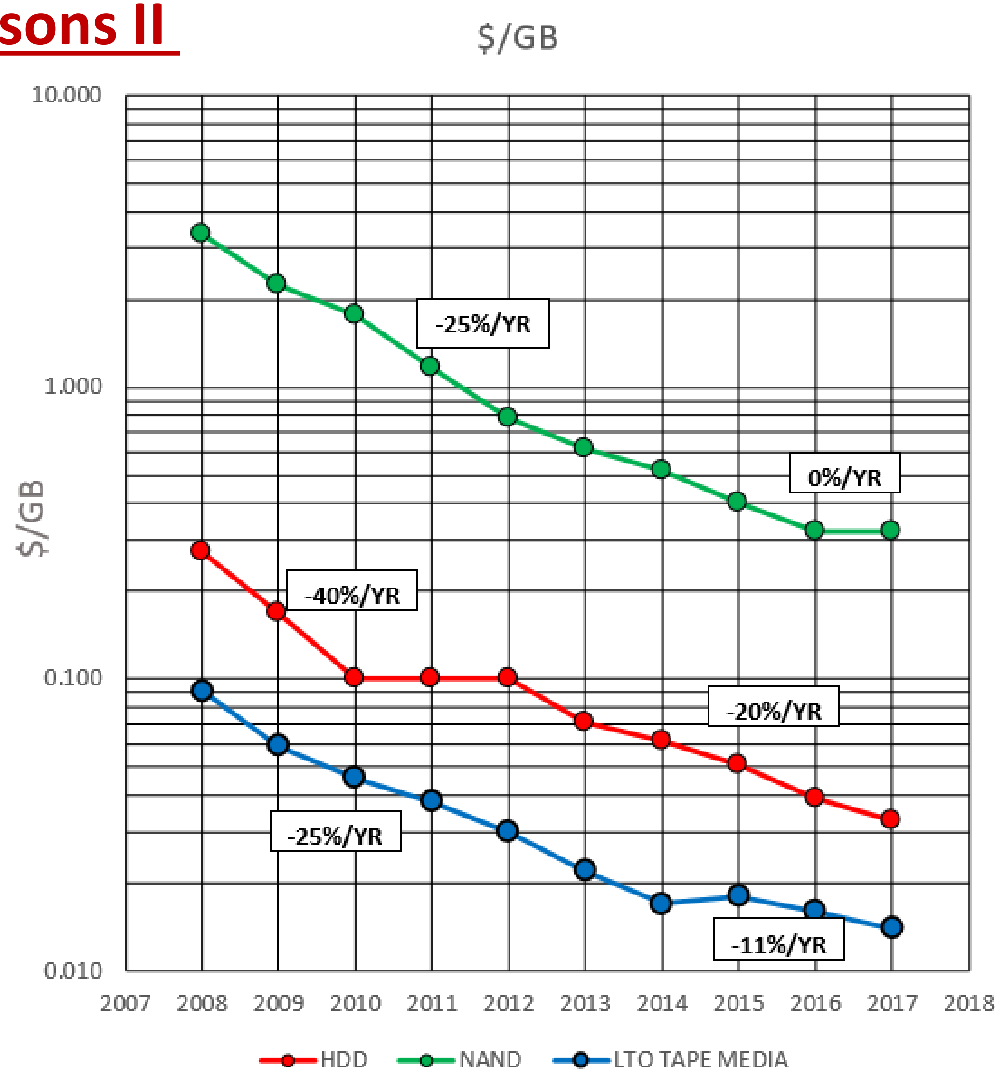

General slowdown, technological developments in the lab are still progressing fast,
Market issues are driving the slowdown
High cost investments are needed

R. Fontana, G. Decad IBM Systems
5/15/2018

100 PB storage example I

Processing Cluster
5000 nodes

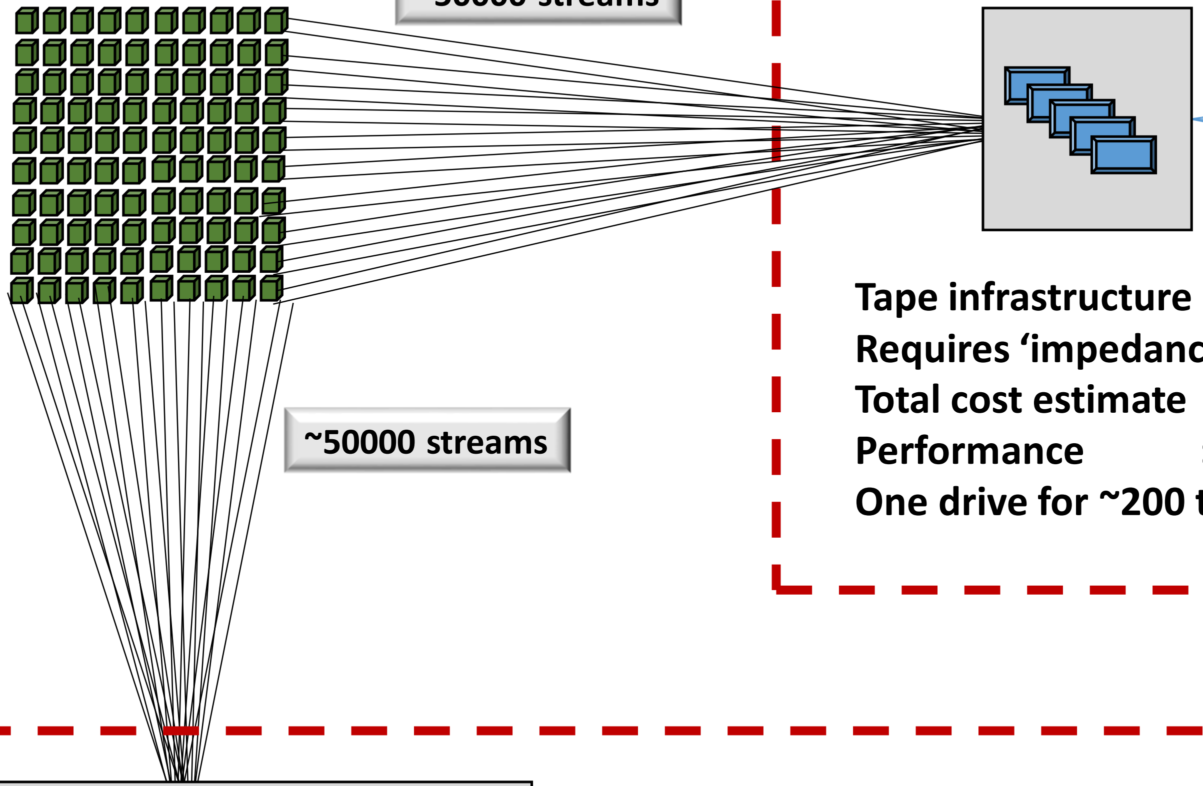

1 PB Storage Cache (SSD)

~40 streams

IBM tape library: 40 LTO tape drives, 100 PB tape media

Tape infrastructure requires disk front-end storage
Requires 'impedance' matching between clients and tape drives → SSDs needed
Total cost estimate : ~1.7 MCHF
Performance : ~ 10 GB/s 116 days to read 100 PB
One drive for ~200 tapes

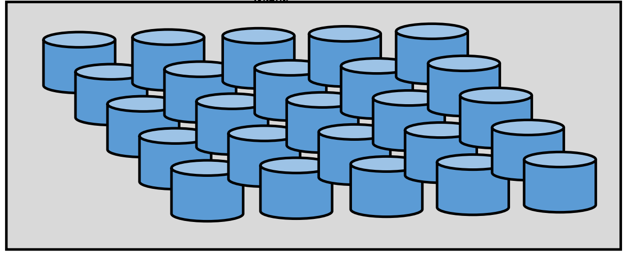

100 PB Storage Cluster (HDD)

100 PB disk storage distributed across 44 disk server
Total cost estimate : ~3.5 MCHF
Performance : ~440 GB/s
3 days to read 100 PB
One server for ~200 disks

100 PB storage example II

Quite some issues to be considered to optimize the storage costs:

- Experiments requests storage space – performance is only indirectly included
  → planning requires optimization of both together: Storage capacity AND storage performance (sequential and random I/O),
- Bare media costs for HDD, SSD and tape have limited relevance → infrastructure multiplication factors, performance differences
- Redundancy level for the various areas (mirrored, server mirrored, replication, 'erasure-code' level, etc.); failure rates
  → cost effects and performance implications
- Storage cost optimization side effects → processing job efficiencies, single point of failures
- Careful consideration of the 'impedance' between the storage areas and the client processing clusters
- Different lifecycles of the storage parts
- Complexity of job+data management infrastructure, sites and experiments
- Site specific boundary conditions
- Specific I/O profiles for different applications: T0 CDR - processing - re-processing – analysis
- Headroom levels
- Operational costs in general
- TCO of small scale specific I/O facilities versus large scale general purpose storage facilities
- Taking into account technology/market developments (e.g. evolution of HDD and tape media sizes)
- ………………

→

Requires holistic view of the full storage architecture (tape+SSD+HDD)
Flexible center storage mixture (tape, HDD, SSD)

Tape is much cheaper than disk is a bit too simplistic statement

Summary

- Technology progress per se is good, but obstacles ahead (NAND, HDD)
- Key computing markets in the hand of very few companies
- Price/performance advances are slowing down
- HDD still key storage for the foreseeable future, SSDs not cost effective for capacity
- Have to closely watch the tape development
- There will be NO relevant new storage technologies in the market in the coming few years (e.g. DNA storage)
- Holistic view needed for the storage architecture, careful combination of SDD, HDD and Tape to optimize pure storage needs and high throughput I/O

References

http://www.digitaltvnews.net/?p=30009
https://www.statista.com/chart/12798/global-smartphone-shipments/
http://www.transformingnetworkinfrastructure.com/topics/virtualization/articles/437082-almost-13-billion-cloud-waste-predicted-2018.htm
https://www.gartner.com/newsroom/id/3845563
https://www.semiconductors.org/news/2018/01/02/global_sales_report_2017/global_semiconductor_sales_increase_21.5_percent_year_to_year_in_november/
https://www.qstar.com/index.php/ltfs-linear-tape-file-system/
http://techinsights.com/about-techinsights/overview/blog/samsung-18-nm-dram-cell-integration-qpt-and-higher-uniformed-capacitor-high-k-dielectrics/
https://www.extremetech.com/computing/249075-foundry-futures-tsmc-samsung-globalfoundries-intel-gear-7nm-beyond
https://www.quora.com/When-is-Samsung-going-to-produce-7nm-and-is-it-comparable-to-Intel-10nm-in-terms-of-performance
https://www.neogaf.com/threads/glofo-7nm-details-revealed-the-most-likely-process-node-for-the-next-gen-consoles.1402083/
http://www.icinsights.com/news/bulletins/Semiconductor-Shipments-Forecast-To-Exceed-1-Trillion-Devices-In-2018/
https://www.forbes.com/sites/tomcoughlin/2018/02/05/hdd-growth-in-nearline-markets/2/#66cf02aa3e39
http://www.theregister.co.uk/2018/03/21/seagate_to_drop_multiactuator_hamr_in_2020/
https://www.theregister.co.uk/2017/12/19/seagate_disk_drive_multi_actuator/
http://www.tomshardware.com/news/seagate-wd-hamr-mamr-20tb,35821.html
https://www.extremetech.com/computing/266031-ibms-power9-dent-x86-server-market-oems-prep-new-systems-emphasize-gpu-compute
https://www.nextplatform.com/2018/03/01/server-market-booms-last/
https://www.backblaze.com/blog/hard-drive-cost-per-gigabyte
https://www.theregister.co.uk/2017/12/22/bit_price_decline_and_computestorage_closeness_to_send_enterprise_flash_use_sky_high/
https://hblok.net/blog/posts/2017/12/
https://www.businesswire.com/news/home/20180314005070/en/Record-Breaking-Amount-Total-Tape-Capacity-Shipments
https://www.dramexchange.com/WeeklyResearch/Post/5/4871.html
https://www.snia.org/sites/default/files/PM-Summit/2018/presentations/14_PM_Summit_18_Analysts_Session_Oros_Final_Post_UPDATED_R2.pdf
https://marketrealist.com/2018/01/micron-optimistic-falling-nand-prices
https://www.nextplatform.com/2018/02/09/just-large-can-nvidias-datacenter-business-grow/
https://www.eetimes.com/author.asp?section_id=36&doc_id=1331415