
CERN

Computing Evolution: Technology and Markets

‘Beyond the LHCb Phase-1 Upgrade’ Workshop, Elba (IT), 28 – 31 May 2017

Helge Meinhard / CERN
Presenting material prepared by Bernd Panzer-Steindel / CERN

Outline

- Semiconductor market
- Device market
- Processors
- Hard Disk
- Solid-State Disks
- Memory
- Tapes
- Server

- Summary
- References

General Market (1)

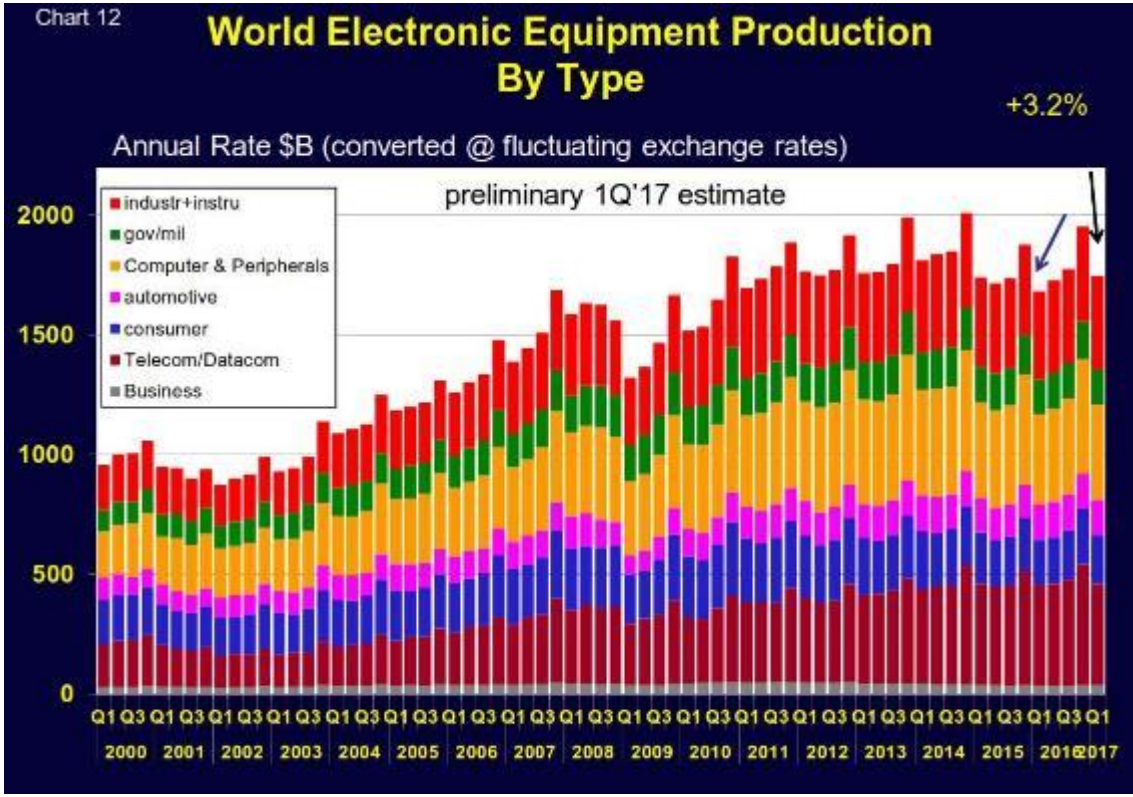

Few companies dominating the markets

Server CPUs	Intel (99%)
FPGA	Xilinx (49%), Intel (38%)
GPU	Intel (68%), Nvidia (18%), AMD (14%)
Hard disks	Western Digital (41%), Seagate (37%), Toshiba (22%)
Tape drives	IBM
Tape media	Fujifilm, Sony
NAND	Samsung (35%), Toshiba (20%), Western Digital, Micron
DRAM	Samsung (50%), Hynix (25%), Micron/Intel (19%)

Electronic equipment production is essentially flat (market saturation)

General Market (2)

Revenues have increased due to high prices for NAND and DRAM

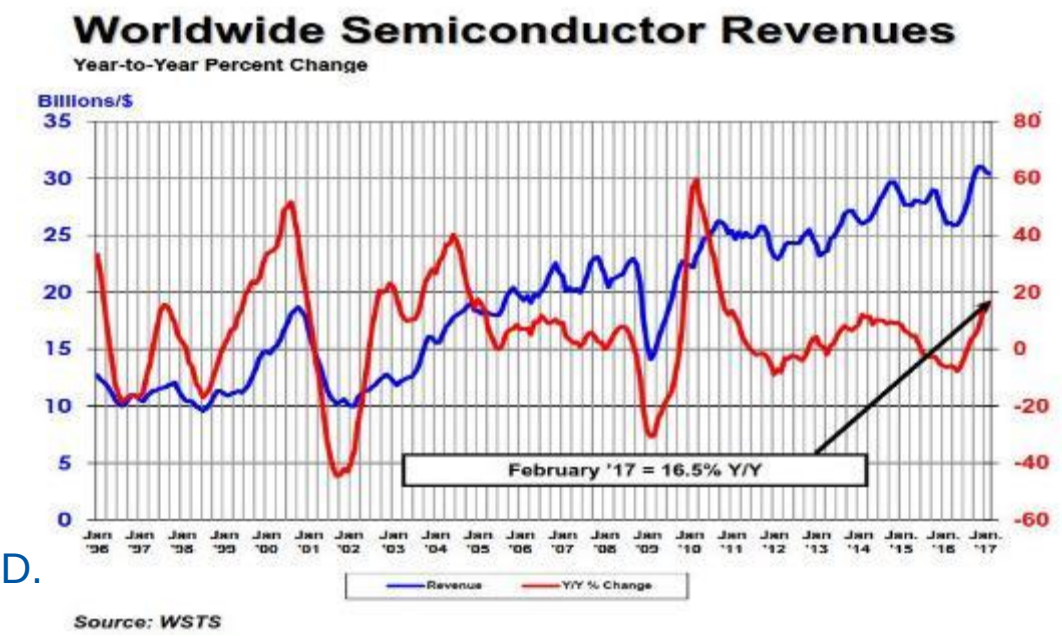

Computing market only small part
~18 B ARM processors,~ 0.3 B Intel/AMD.

Expect 1 trillion semiconductor units shipped in 2018.

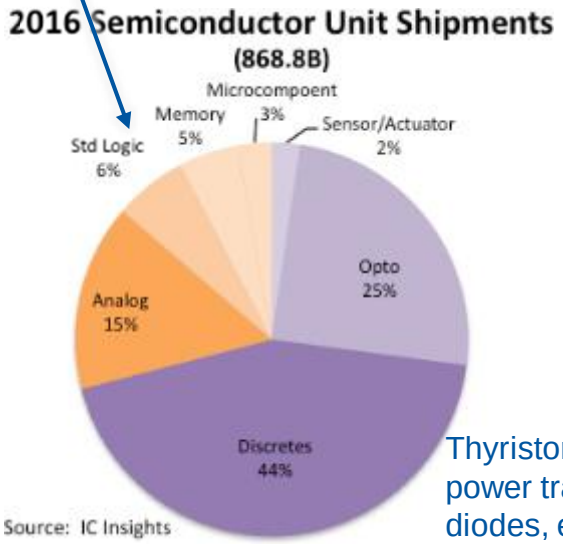

Thyristors, power transistors, diodes, etc.

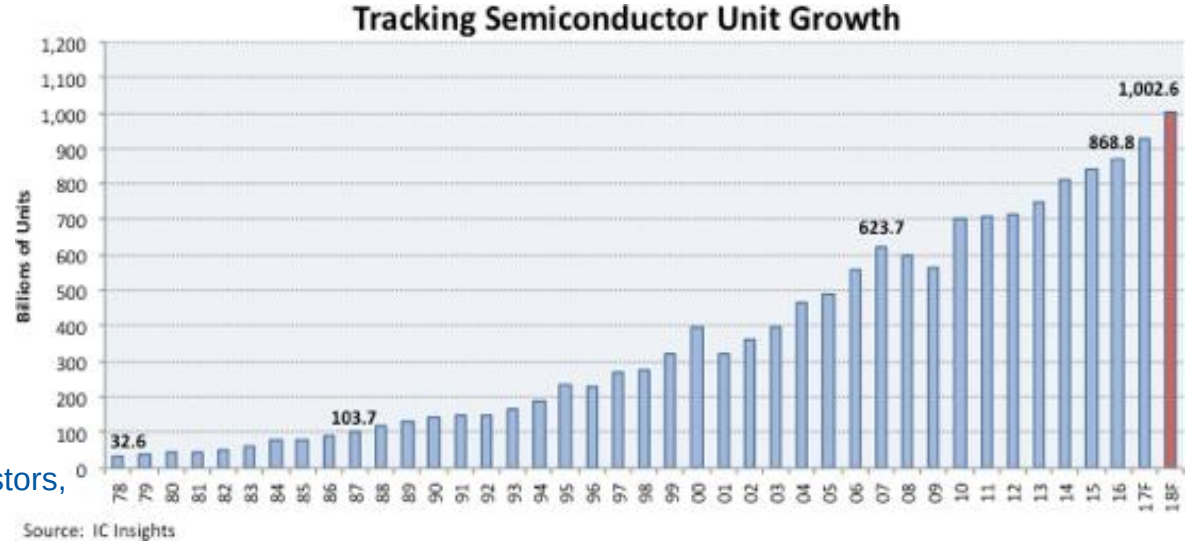

Device Markets (1)

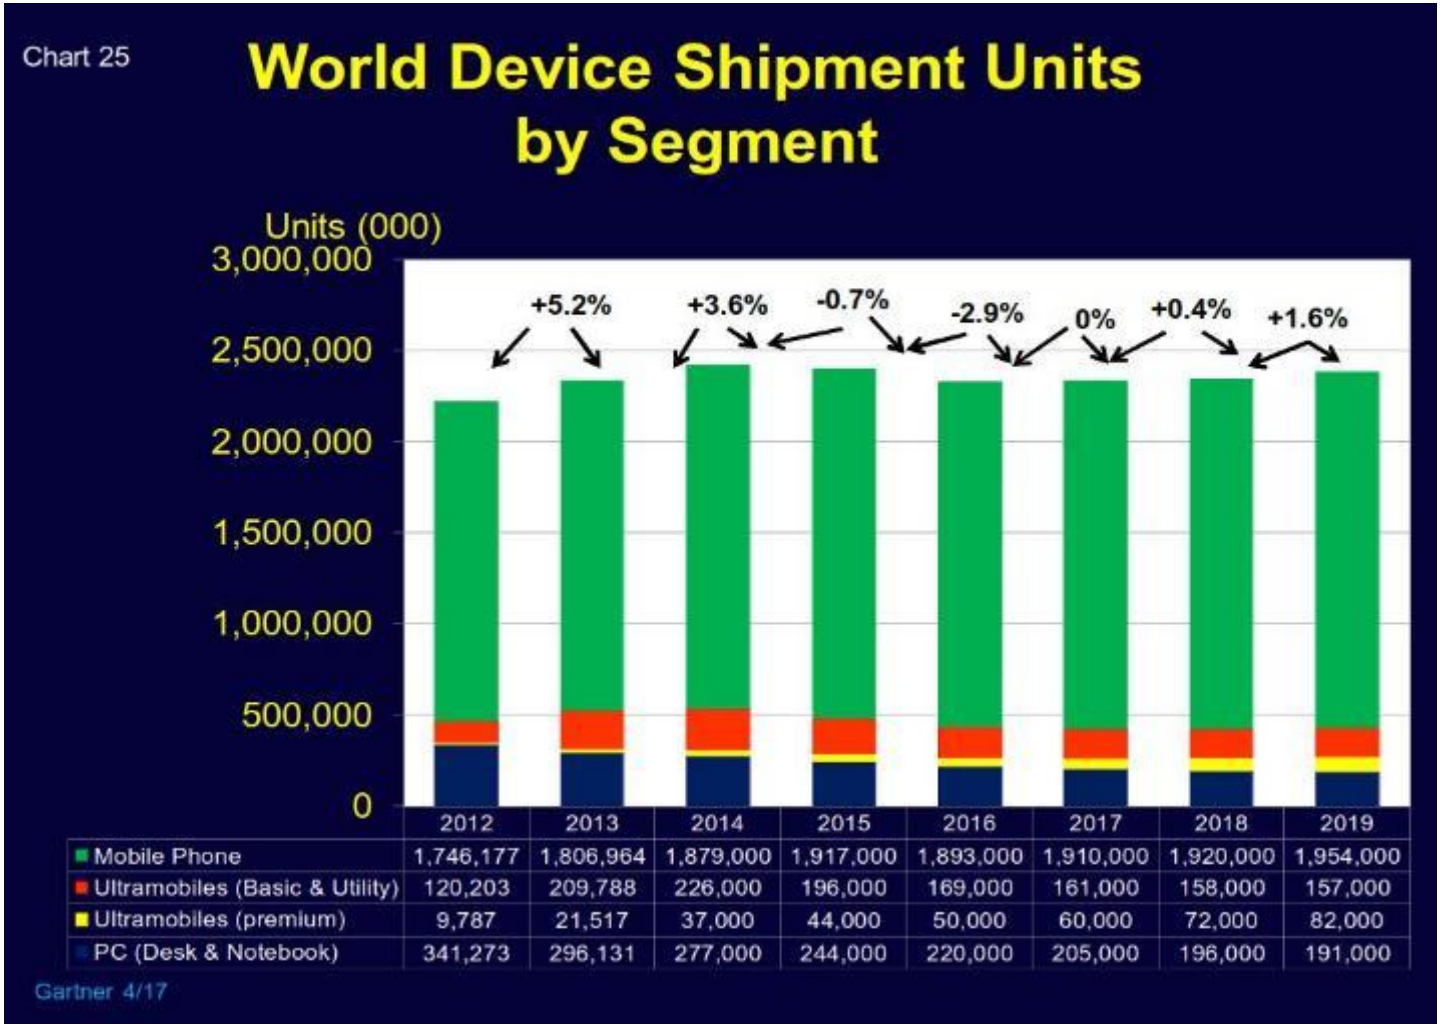

Market saturation: minimal or negative growth rates
Longer product lifetimes

Device Markets (2)

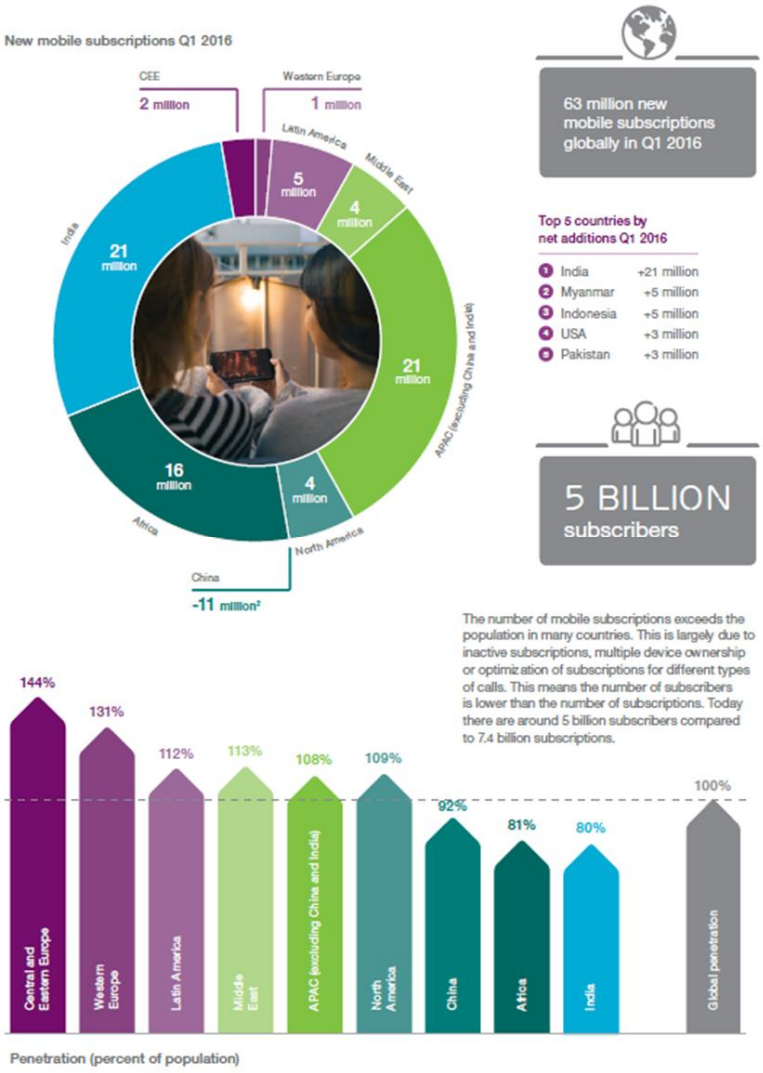

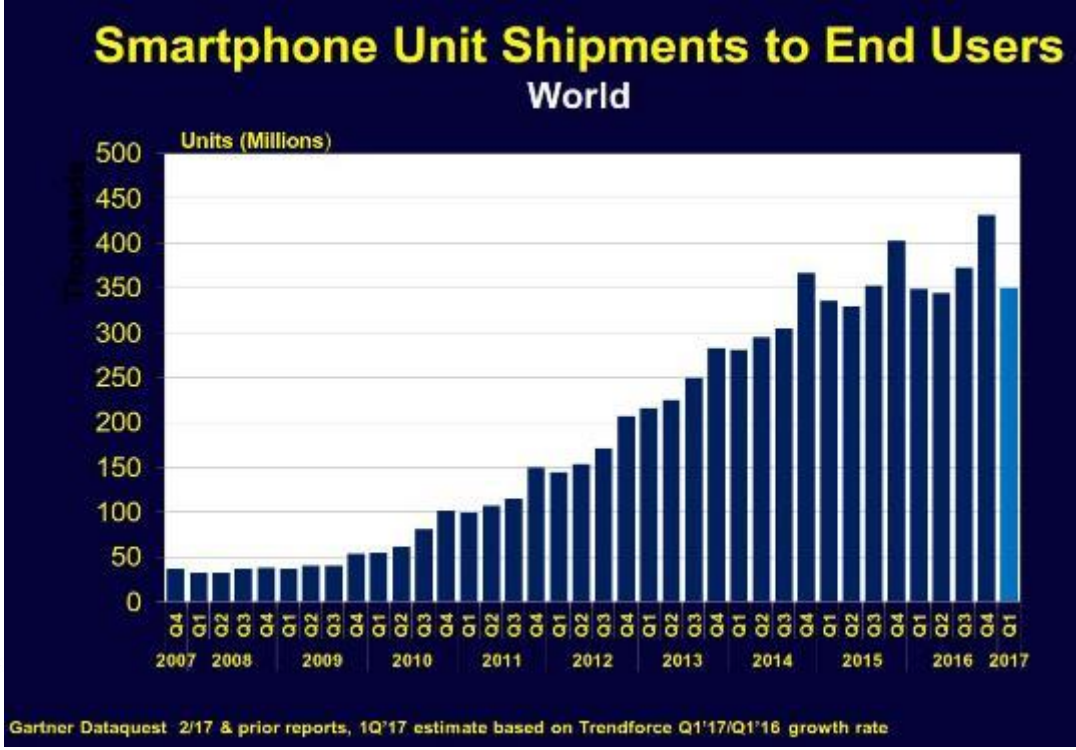

Saturation:
7.3 B phone subscriptions world-wide – more than the population

Replacement bump expected in 2018

Processors (1)

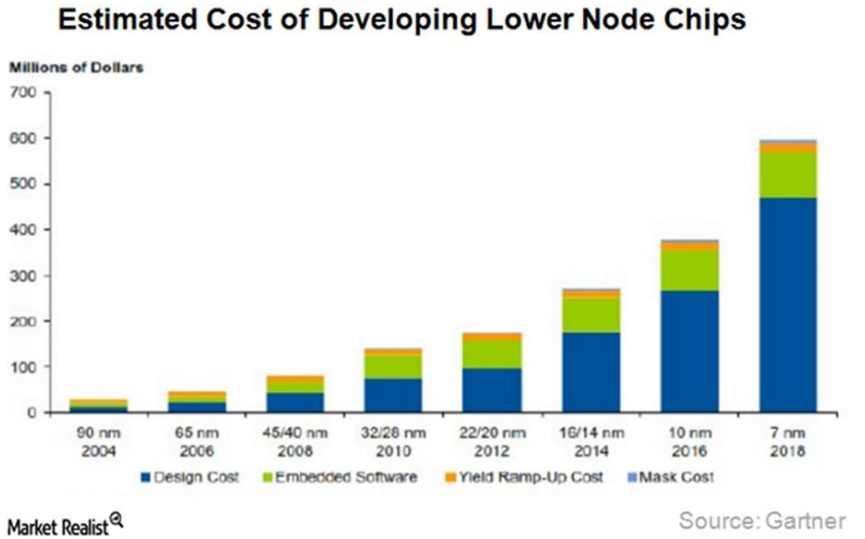

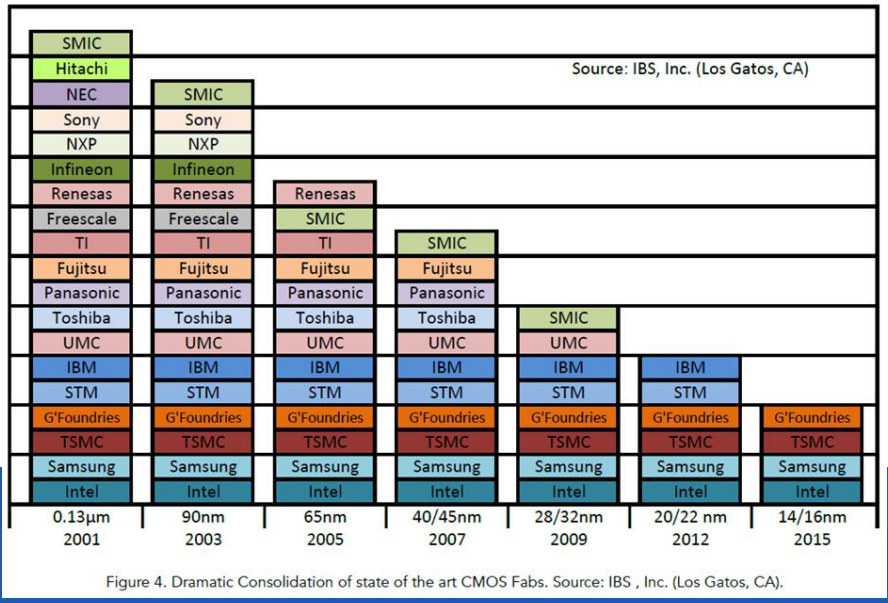

Figure 4. Dramatic Consolidation of state of the art CMOS Fabs. Source: IBS , Inc. (Los Gatos, CA).

Non-linear costs for development

- Only four companies able to fabricate 14 nm chips
- 10 nm Samsung fab costs $14 B

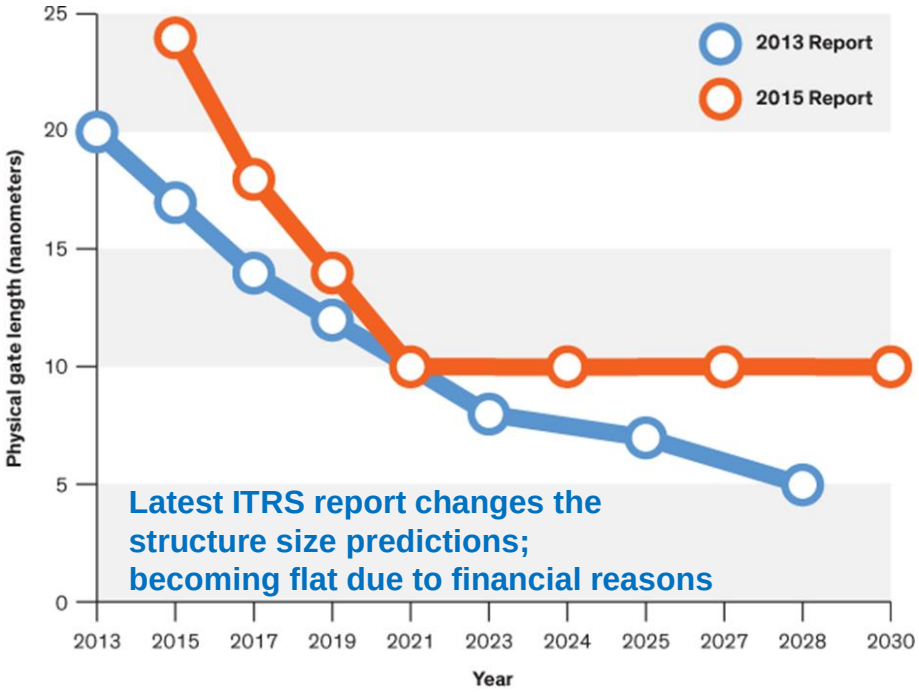

Processors (2)

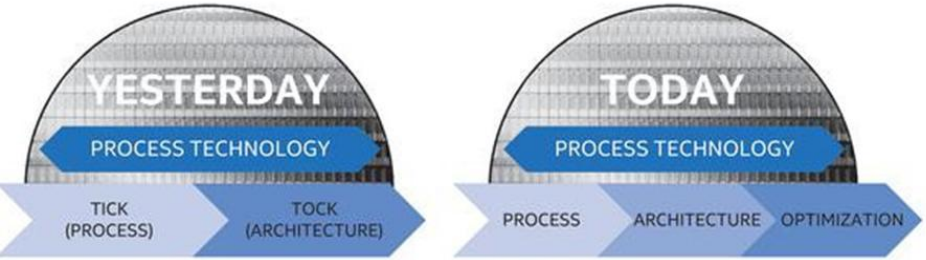

Intel moved from 2-year cycle to 3 years or more

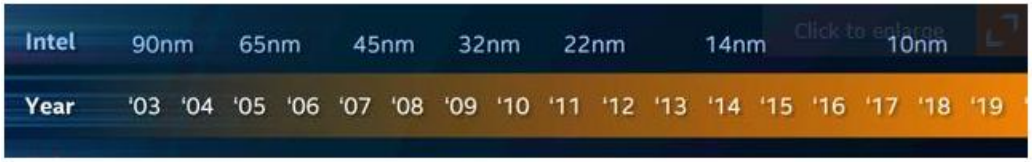

16/14 nm finFET Comparison — TECHINSIGHTS

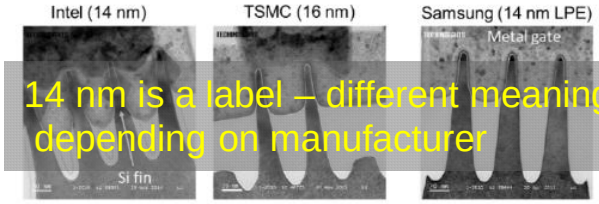

14 nm is a label – different meaning depending on manufacturer

Feature	Intel	TSMC	Samsung
Gate length (nm)	24	33	30
Min contacted gate pitch (nm)	70	90	78
Fin height under gate (nm)	42	37	37
Fin pitch (nm)	43	45	49
Min metal pitch (nm)	52	70	67

- Intel transistors are smaller than TSMC or Samsung

The ConFab

#TheConFab2016

Incubation Time

- Strained Silicon
  - 1992->**2003**
- HKMG
  - 1996->2007
- Raised S/D
  - 1993->2009
- MultiGates
  - 1997->2011

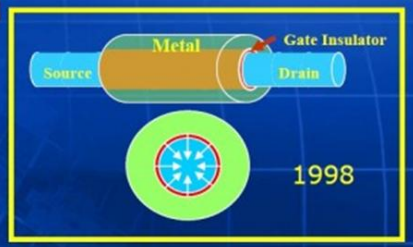

~ 12-15 years

Decrease of feature size goes along with new material technologies

R&D à production needs 12-15 years

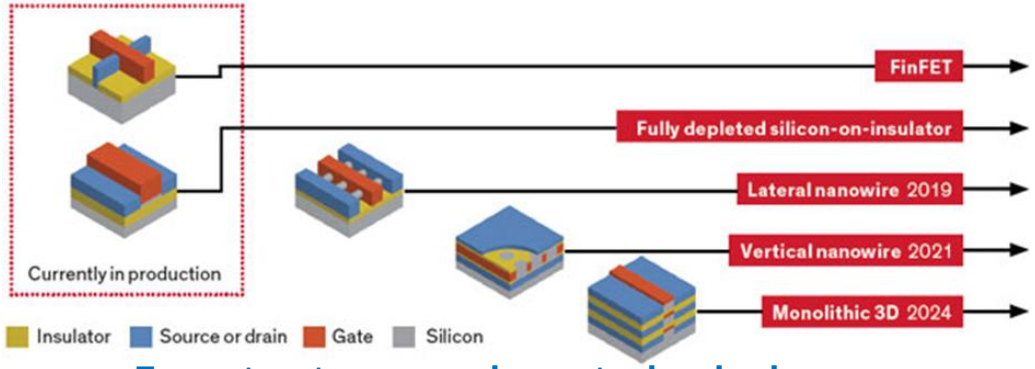

7nm structures need new technologies: nanowires and non-silicon material

Accelerators: GPU (1)

Embedded market shares (CPU+GPU):
Intel 68%, Nvidia 18%, AMD 14%
Discrete GPU cards: Nvidia 77%, AMD 23%

New products announced:

- Nvidia Volta: 12 nm, 21 B transistors, 15 TFlops SP (Q1 2018)
- AMD Vega: 14 nm, 12 B transistors, 12.5 Tflops SP (Q3 2017)

Focus: high-end Gamer (DP and FP16 artificially reduced)

Professional workstation cards and HPC: small niche, ~2 million cards per year (compared to 350 million total GPUs)

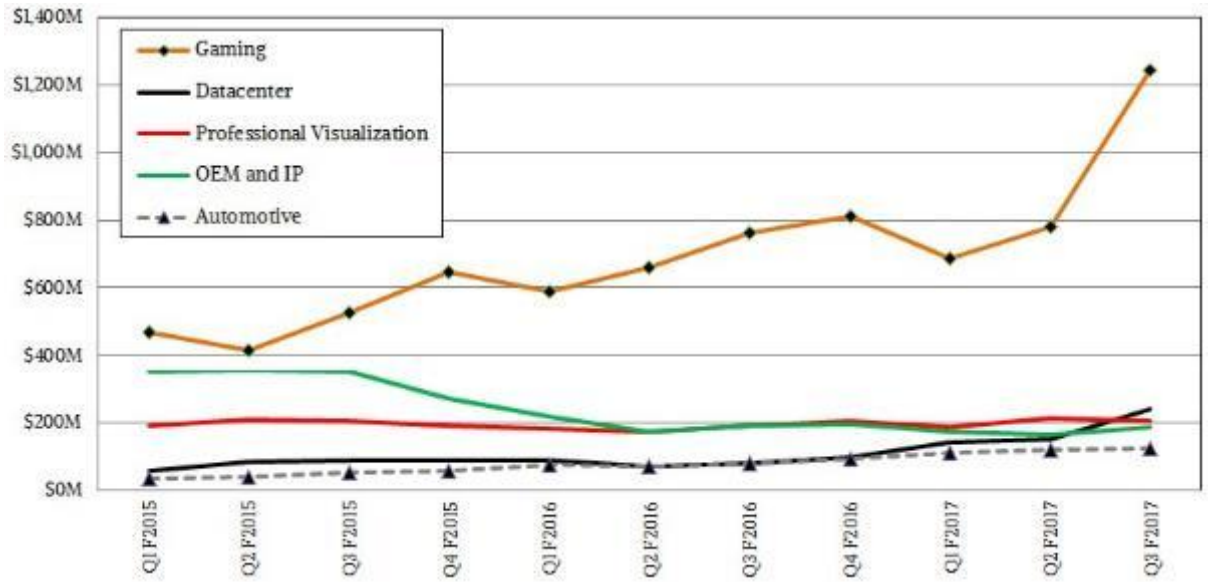

Nvidia Revenues

Accelerators: GPU (2)

- New focus for graphic cards: machine learning
- Move to FP16 and even INT8 architectures, less precision à 8 bit processing !
- Google TPU Tensor Processing Unit
- New start-ups with special processor designs:
  e.g. KnuEdge, Nervana (just bought by Intel), SpiNNaker, Eyeriss, P-Neuro, NeuRAM3
  - Essentially not usable as general purpose processors
  - Online?!
- Intel changing strategies also for their KnightsXX processors, 'forking' models (increase FP16 and decrease DP)

  ~100k units per year, very small market
- Qualcomm plans to add neuromorphic chips into the smartphone

Hard Disks (1)

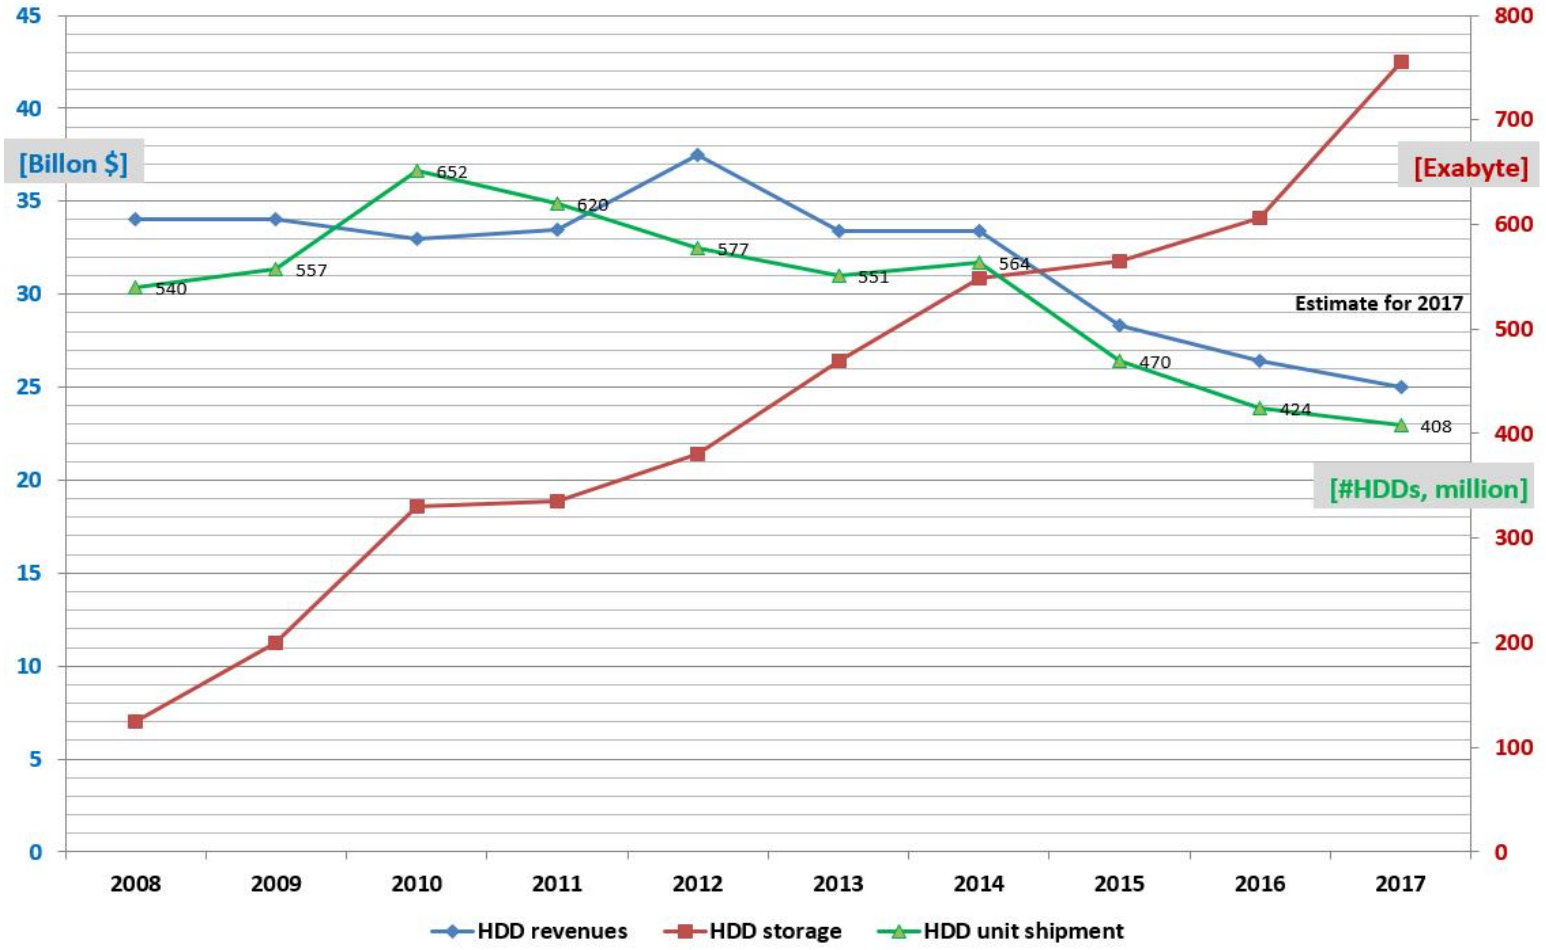

Continuous decrease in revenues

Shipments decreasing, Seagate just closed one of their major production fabs

Pressure from SSDs in the high end enterprise drive market

Hard Disks (2)

Combining bit density (30% annual growth rate) and volume density (number of platters, helium) → 100 TB in 2025 conceivable

Areal density improvement dropped from ~40% to 16% per year

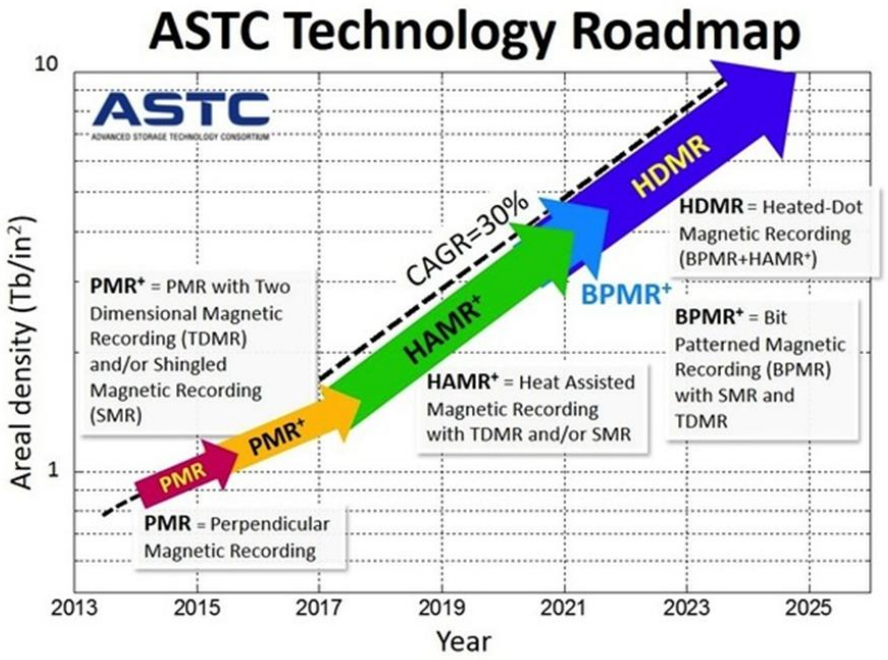

Areal Density Trends

Chart provided courtesy of the Information Storage Industry Consortium (INSIC)

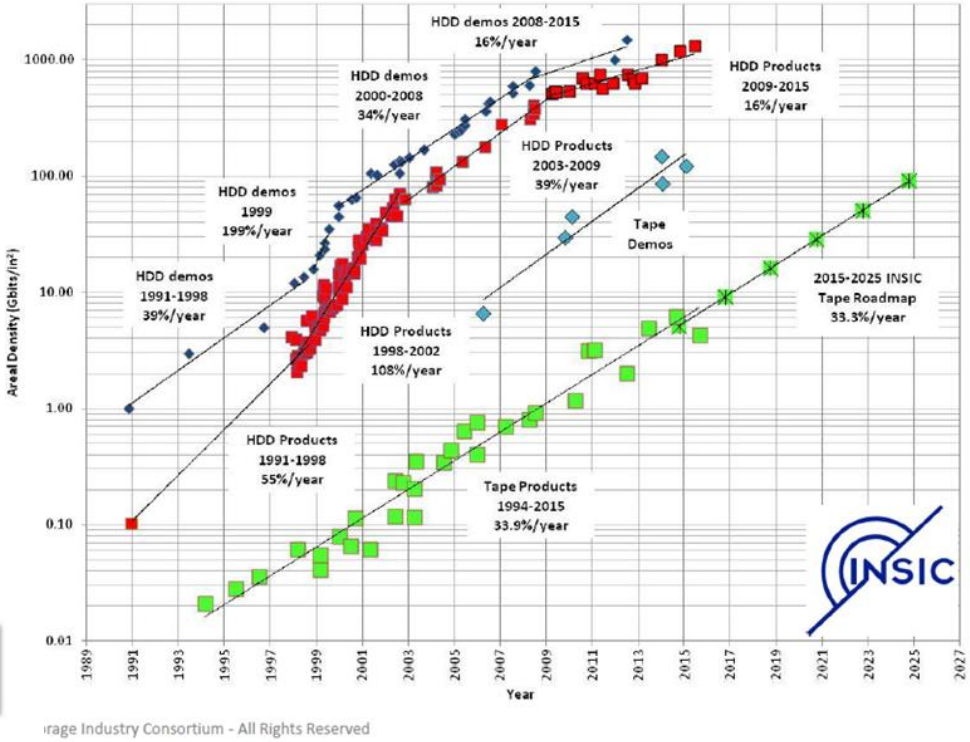

PMR limit at 1 TbPSI
SMR adds ~25%, market small
HAMR should provide 5 TbPSI

HAMR delayed, production in 2018

Hard Disks (3)

Focus on infrastructure cost reduction
i.e. Helium drives, more platters per drive

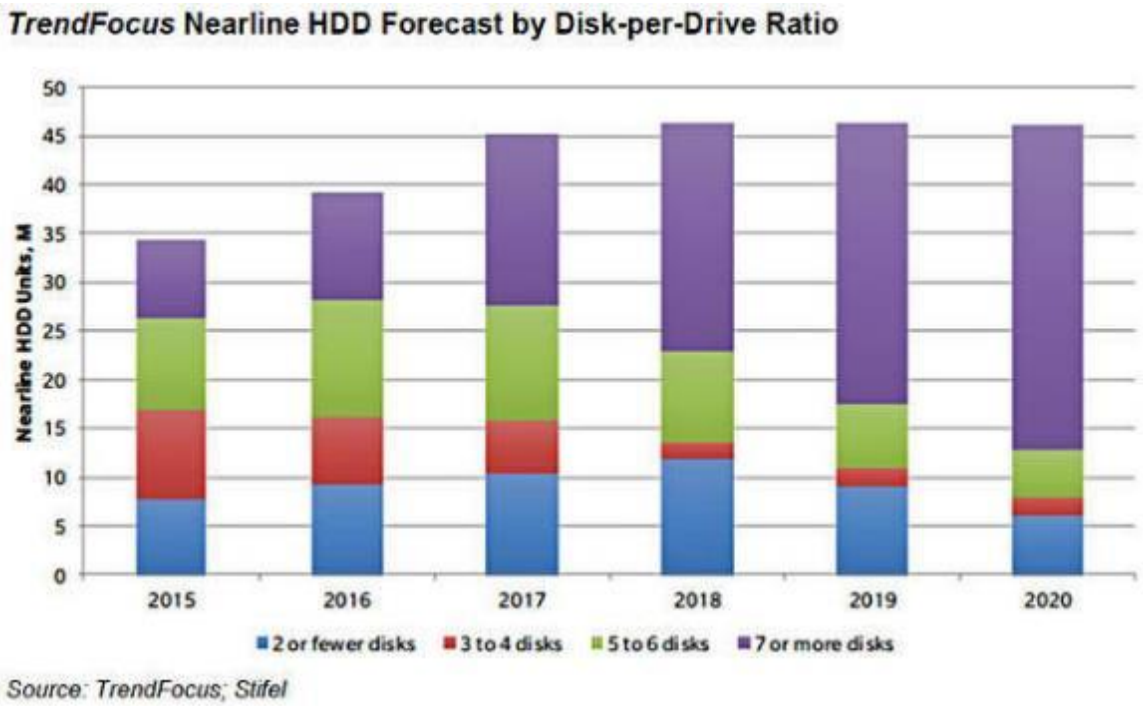

To reduce costs, at CERN, we are looking into

- Desktop drives
- Multi PB disk server
- Erasure encoding

Google paper: interest in much bigger drives (> 3.5”)
Amazon patent: separate mechanics from electronics (steering boards outside of the drive)

Solid-State Disks (1)

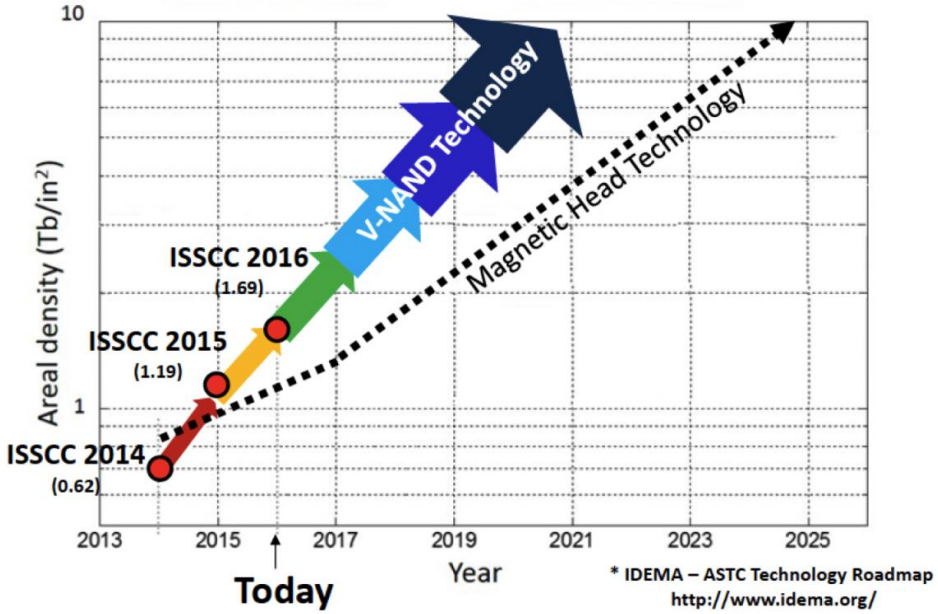

3D NAND – scaling in the third dimension

- 2D NAND scaling beyond 16nm/15nm is uneconomical.
- 3D NAND adds additional layers for scaling in place of 2D lithographic scaling.
- Bit density is continuing to scale with the potential for terabit NAND die.

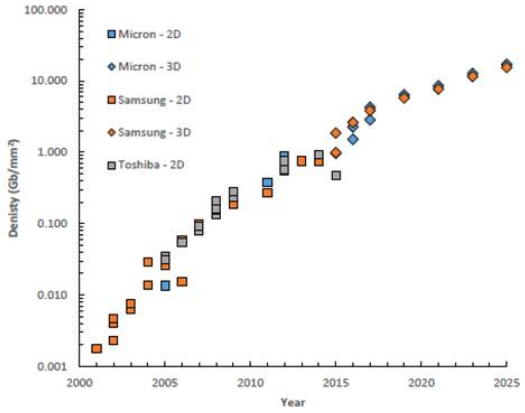

NAND density has surpassed HDD density

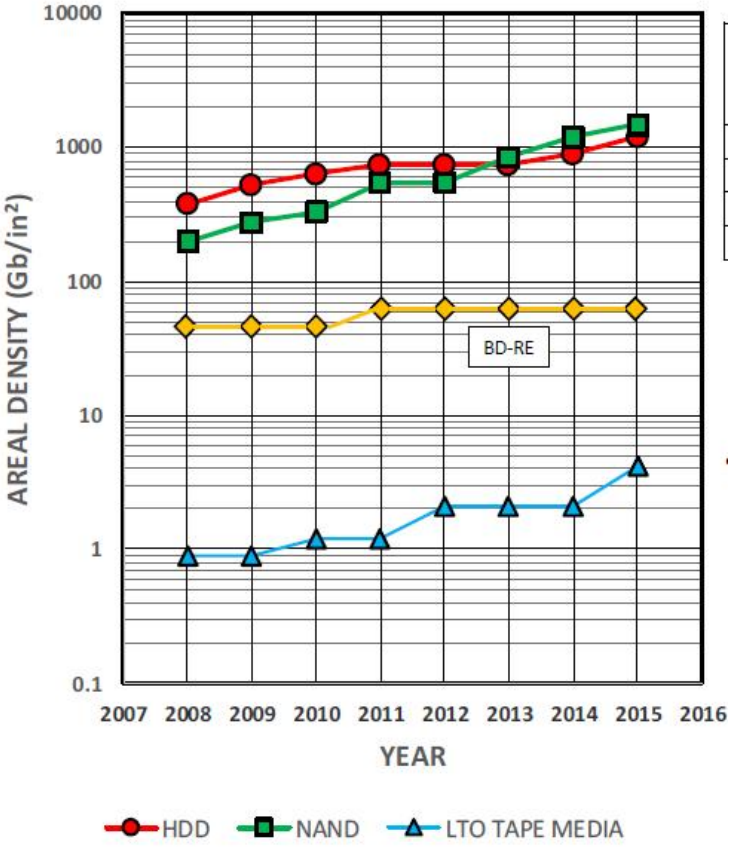

IC KNOWLEDGE LLC

Solid-State Disks (2)

Source: Yangtze Memory, Samsung Securities

NAND:

- 2D scaling came to an end 2 years ago
- Feature size manufacturing lags slightly behind processor structure sizes, <20nm today
- 3D: 64-layer 3D NAND in production; 72-layer expected end of 2017
- >30% price increases since last year, expect to last until end 2017
- Yield of 3D NAND improving, expect >50% of all shipments in Q4 2017 to be 3D instead of 2D

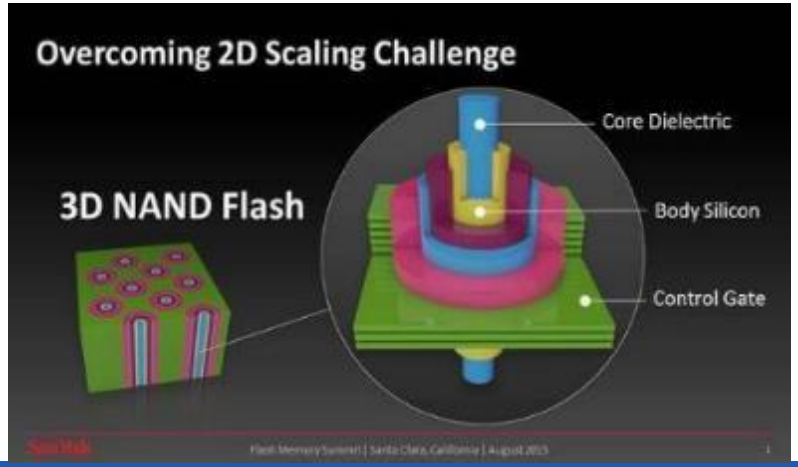

Solid-State Disks vs. Hard Disks (1)

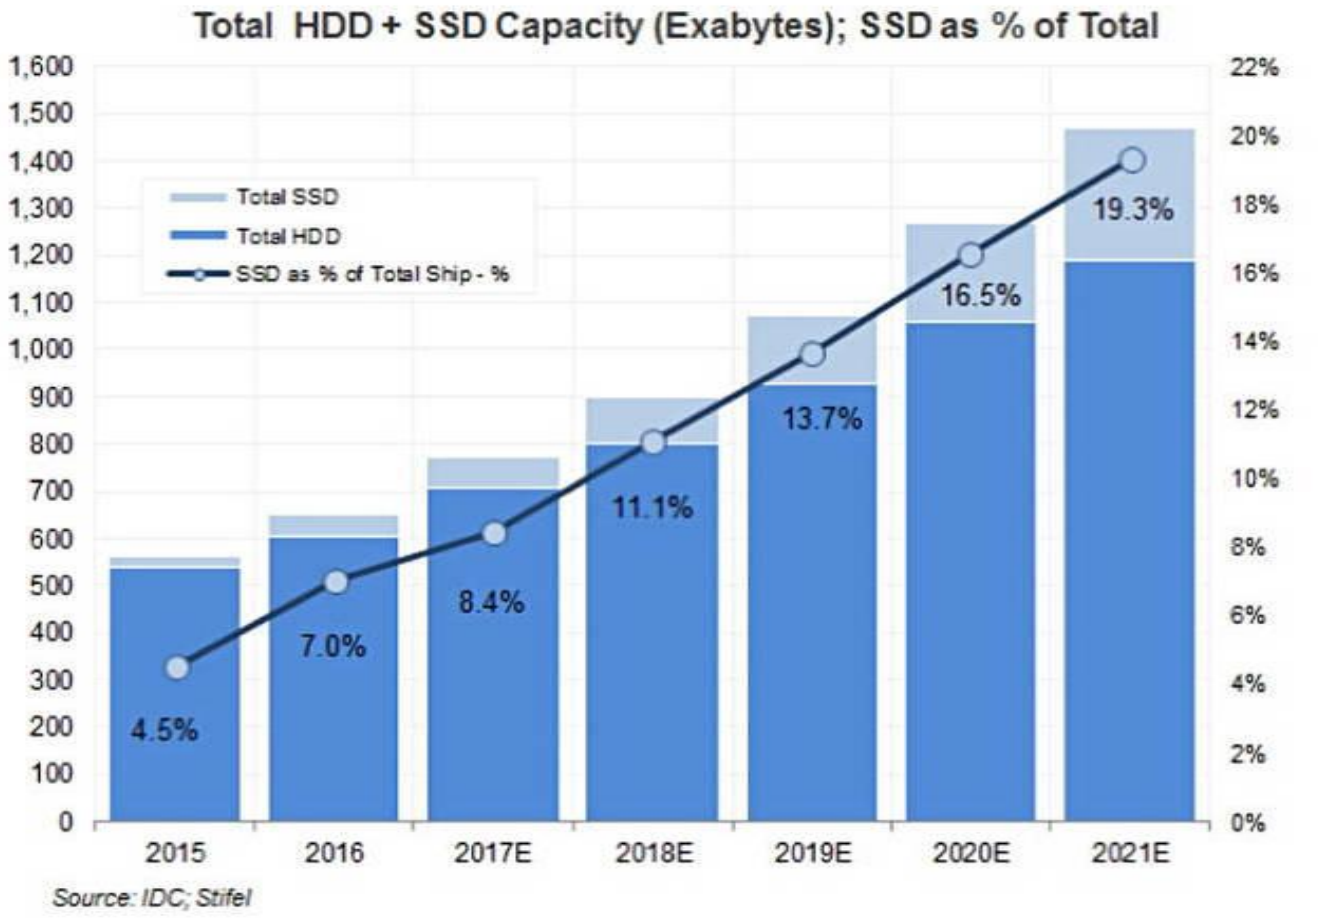

16 times more HDD capacity than SSD was shipped in 2016

NAND Fab investment of 100-200 B$ necessary to achieve HDD ExaByte deliveries

Solid-State Disks vs. Hard Disks (2)

SSDs not cost effective for HDD capacity replacement

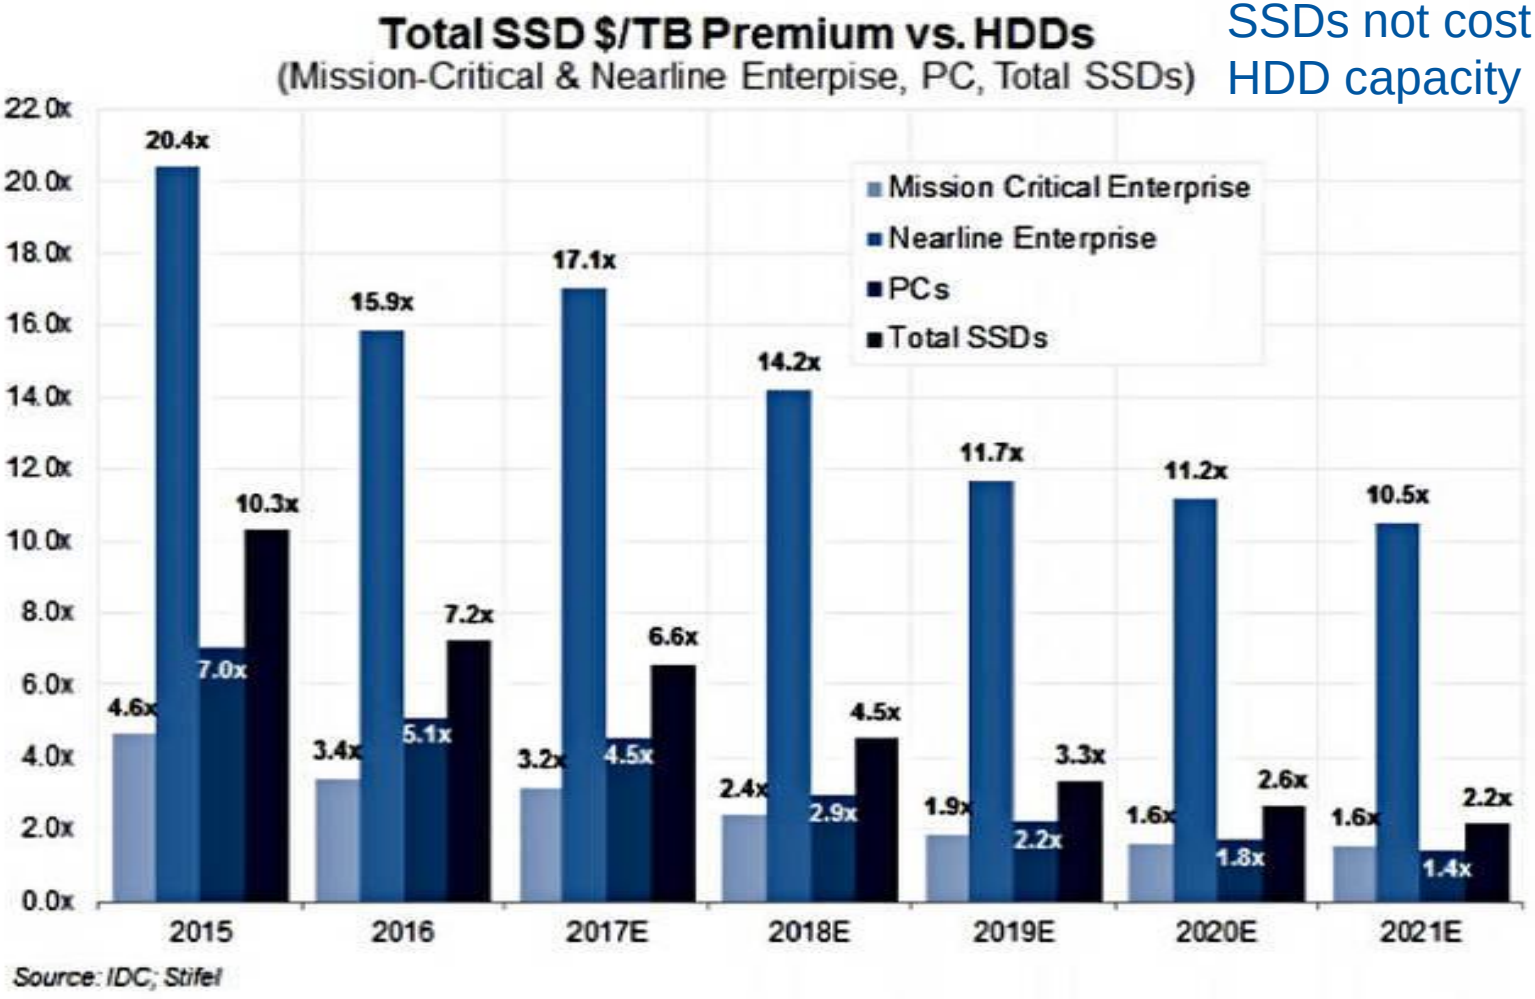

Currently supply shortage of NAND
à SSD price increases

Nearline = capacity drives
à SSDs not foreseen for large scale storage

Memory: DRAM

Limited future improvements on performance and energy efficiency

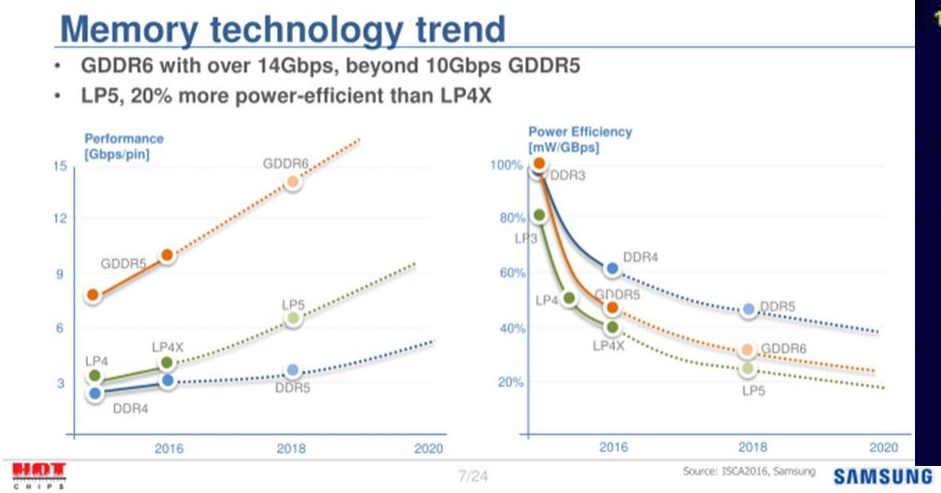

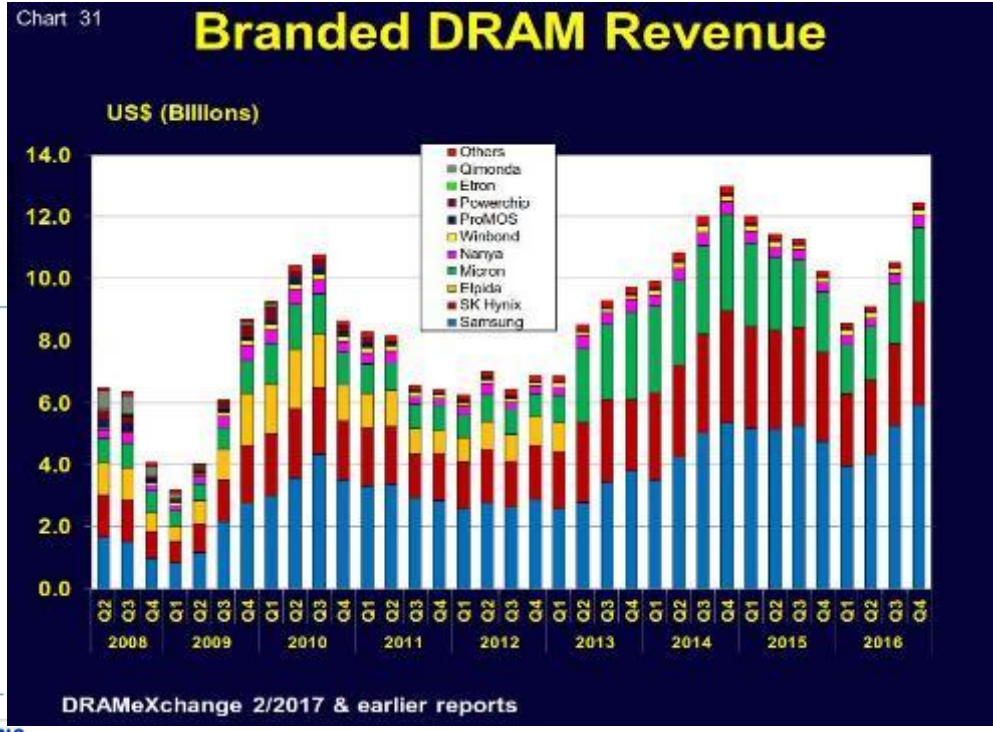

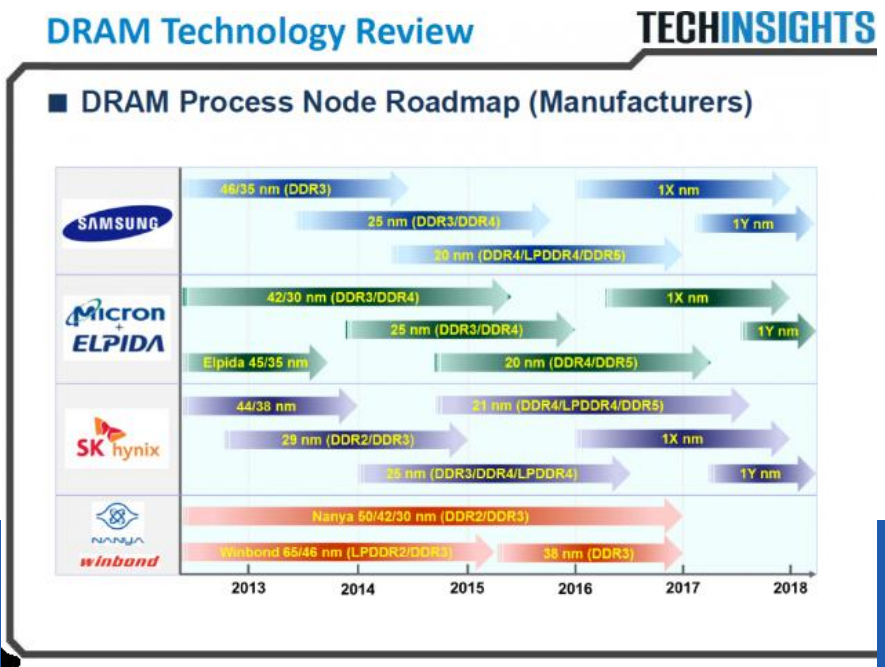

Volatile market, supply shortage started second half 2016, > 50% price increase until now, will be ongoing until end 2017

2 Chinese companies will enter the DRAM market in 2017

New Memory Technologies

- 3D Xpoint: new technology from Intel and Micron, presumably a variant of Phase Change Memory

  Specs are changing:

  Announcement 2015: 1000x faster, 1000x endurance, 10x denser than NAND

  IDF 2016: 10x faster, 3x endurance, 4x denser than NAND

  first products (Optane) announced, usage currently limited to high end HDD caching

- Memristors: developed since 2008; HPE now collaborating with SanDisk (ReRAM)

- Spin torque MRAM in larger production units available (Everquest + Globalfoundries)

  Low density and high price

- Nantero just received extra funding for their Carbon Nanotube NVRAM, exists since 17 years, no product yet

- RRAM or ReRAM , various new categories being developed: Oxide RAM (OxRAM), Conductive-Bridge RAM (CBRAM) or Self-Rectifying Cells (SRC)

à But… NAND fab investments are high, extended technology lifetime with 3D, hard to replace in the short term

Magnetic Tapes (1)

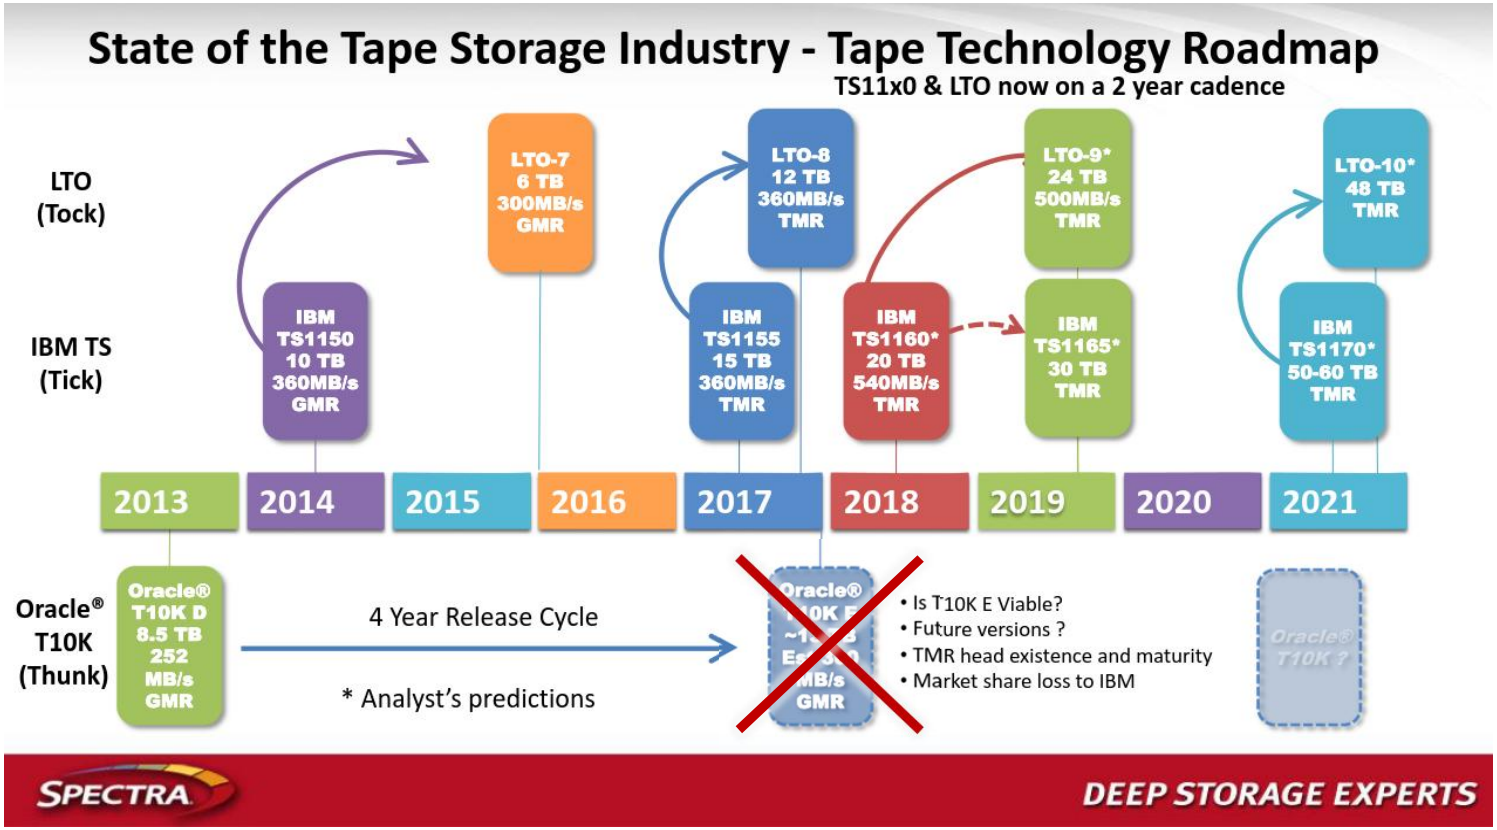

- Enterprise drives:
  Oracle 2017: stopped Enterprise production !!
  IBM end 2017: 10 TB à 15 TB (just announced, TS1150)
  LTO-8 end 2017: 12 TB
- Technology in the lab:
  Fujifilm 154 TB, Sony 185 TB, IBM 220 TB

Technology change to
Tunnel Magnetoresistive heads
(used already in HDDs)

Magnetic Tapes (2)

~10 EB LTO capacity shipped per quarter
à Comparted to 12 EB SSDs, 36 EB NAND and 160 EB HDDs

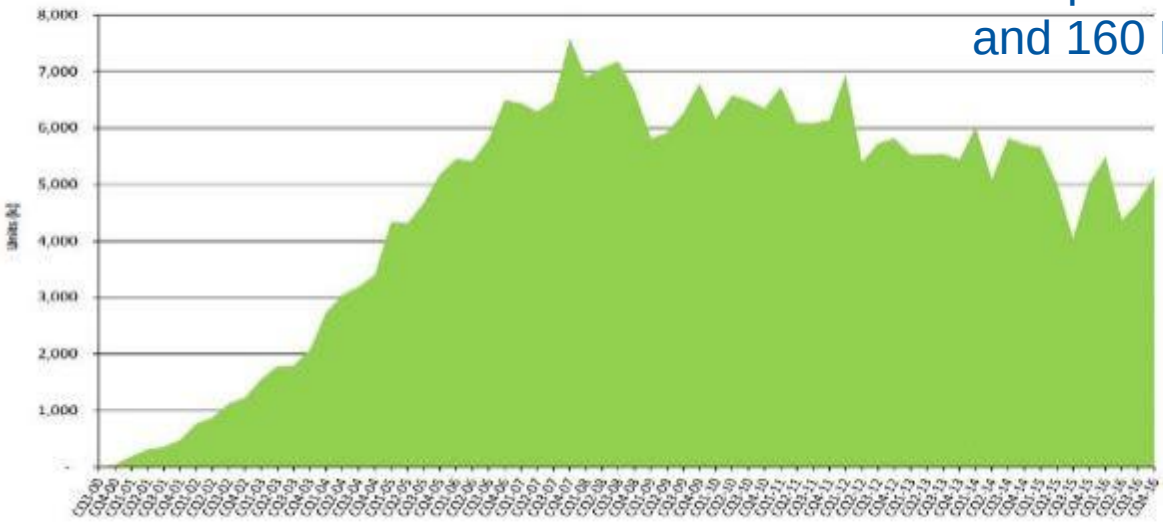

Steady decrease of LTO tapes shipped and revenues

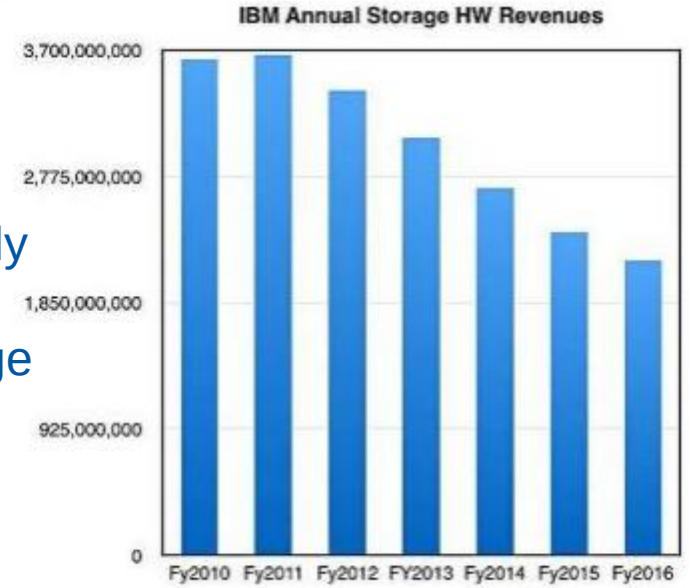

Future of tapes ?

All tape drive heads are now produced by IBM only

à Steady decrease of revenues in the IBM storage group

Servers (1)

- Server market is saturated: flat revenues and unit shipments during the last 2 years
- High profit market
- Single vendor: Intel, 99% market share
- Several initiatives to change that:
- OpenPower (IBM): consortium with many members
  - But revenues still going down, little impact so far
  - Announcement of POWER9 might help
- ARM server:
  - AppliedMicro , Qualcomm, Cavium: new high end products Announcements for 2H2017 (third ARMv8 Wave 2017-2018), First two waves had little impact
  - Phytium (China), “Mars” processor
- AMD with new processor design (Zen), Naples for servers ~Q3 2017
- Fujitsu ARM-powered supercomputer
  - Add large vector instructions to the ARM design
  - Aimed for 2020, now ~2022

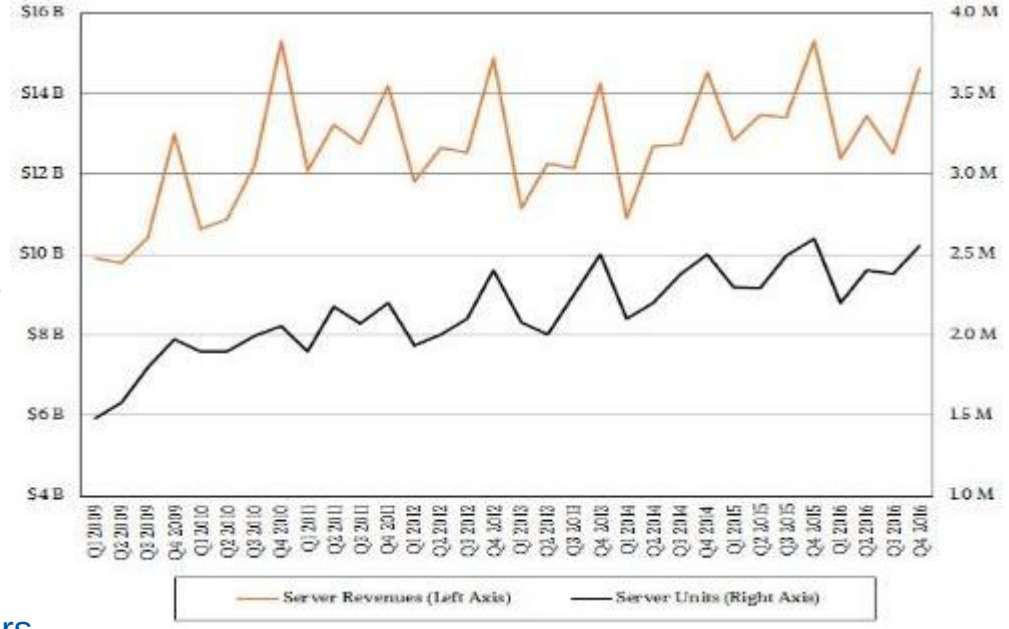

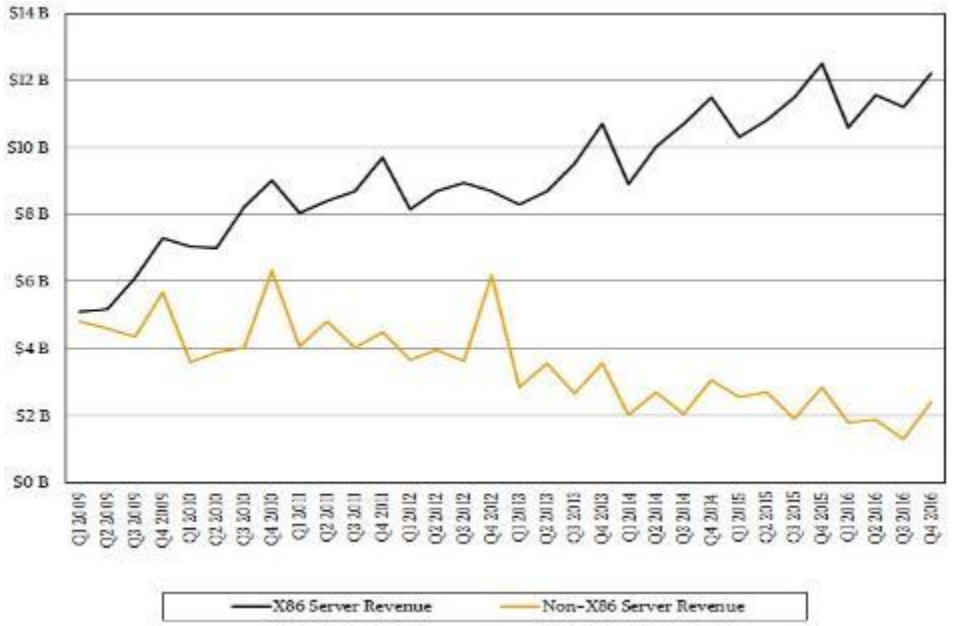

Servers (2)

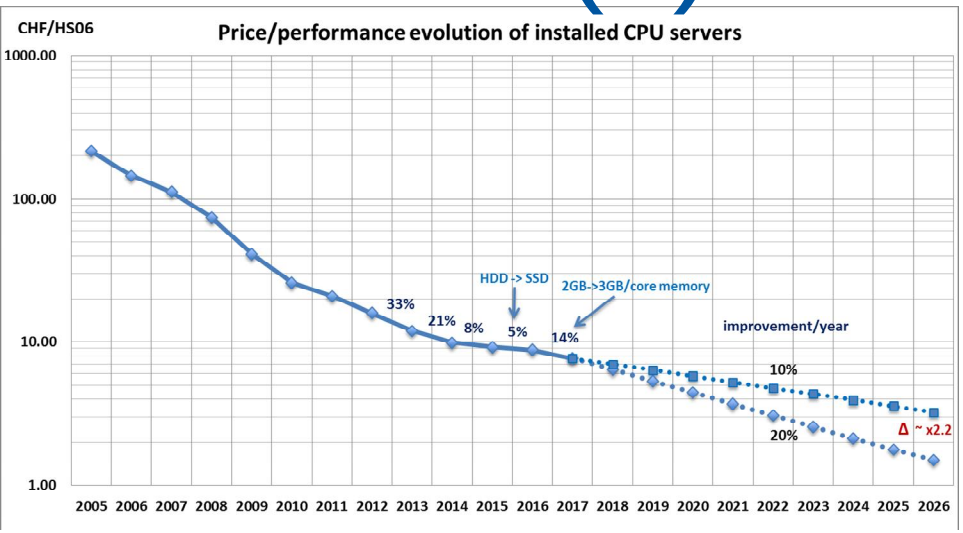

Preliminary extrapolation of CPU and disk server costs (based on CERN procurements)

Pessimistic and reasonable improvement extrapolations

Influence of changing software and hardware architecture requirements to be taken into account (programs, data model, data centre, …)

e.g. CERN moves from 2 to 3 GB/core (+8% cost), driven by experiment usage AND technology boundary conditions

- Moore's Law and Kryder's Law are slowing down
  - 18 months à >= 3 years
- Real cost/performance evolution driven by financial and market aspects rather than technology

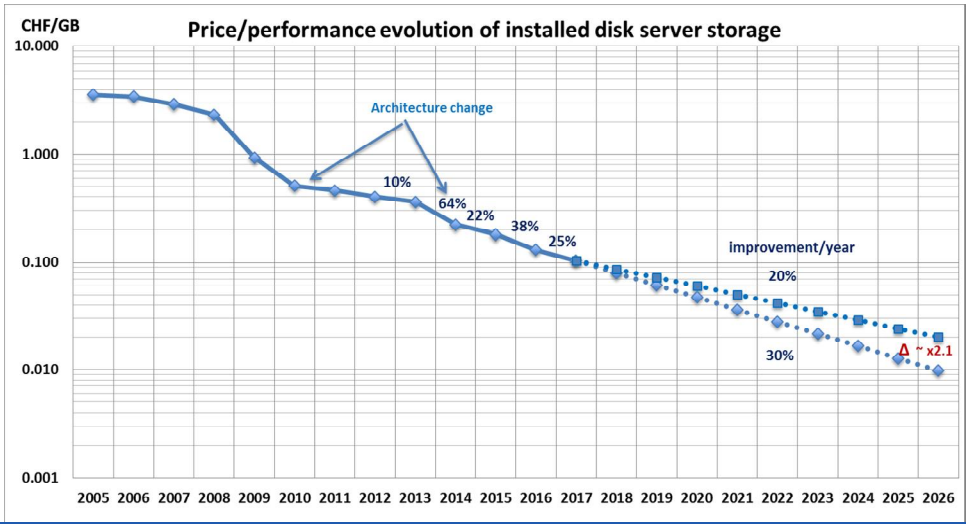

Summary (1)

- Device markets (smartphones, tablets, PCs, notebooks, servers, HPC) saturated or even negative growth
  - Replacement market
- Moore's Law in trouble, financial issues
  - Not clear how this effects price/performance evolution
  - So far okay for CPU and disk servers
- Technology improvements still continuing, but requires high CAPEX

  End-product price tag evolution more complicated
- Market dominance of few companies increases, competition diminishing

Summary (2)

- Technology alone unlikely to solve the computing problem at HL-LHC and beyond
  - Not much more to be expected than minor contributions

References

More details and a (long) list references:

https://twiki.cern.ch/twiki/bin/view/Main/TechMarketPerf


CERN