Technology, Market and Cost Trends

IC Markets

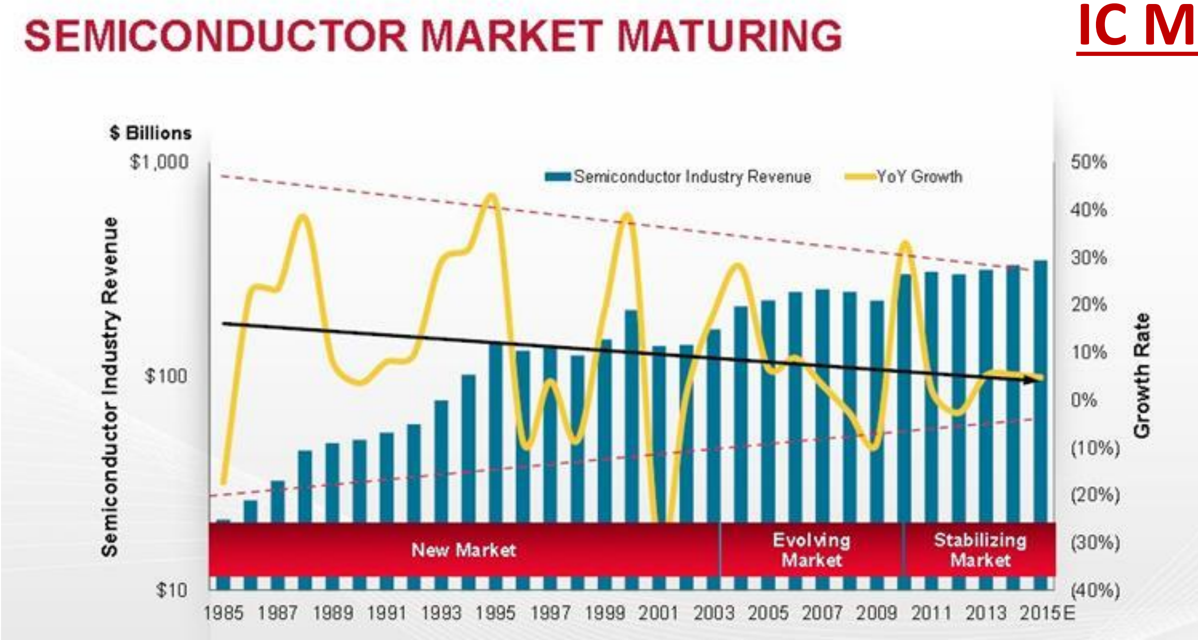

Chip market made 333 B$ revenues in 2014

Moderate growth
Stabilized market

Expect 1 Trillion ICs (integrated Circuit) to be produced per year in 2017

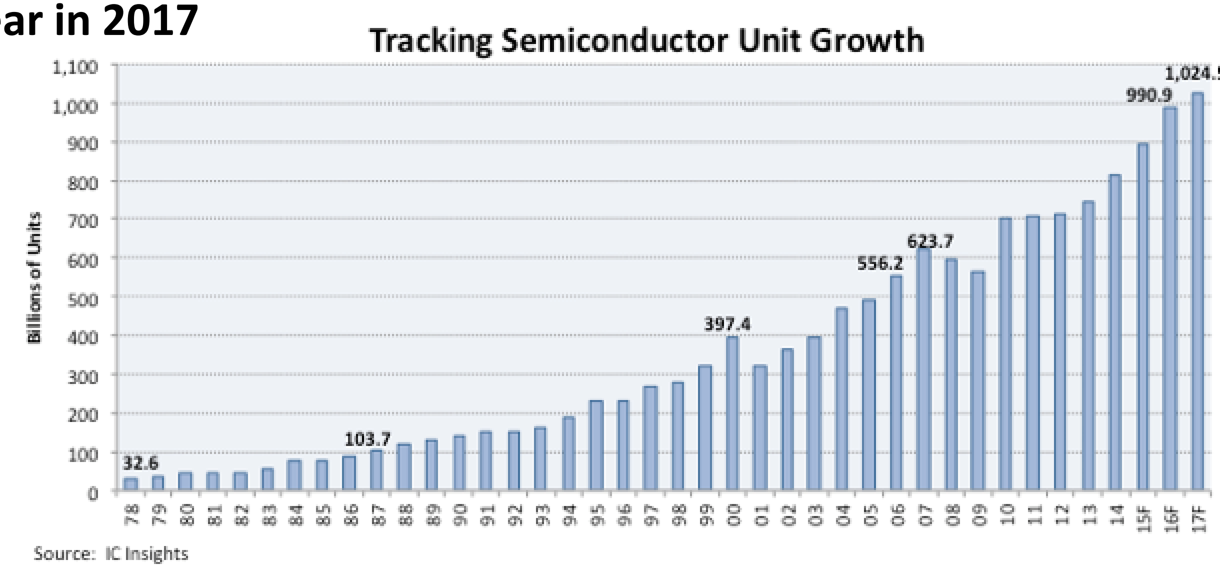

End-Use Markets

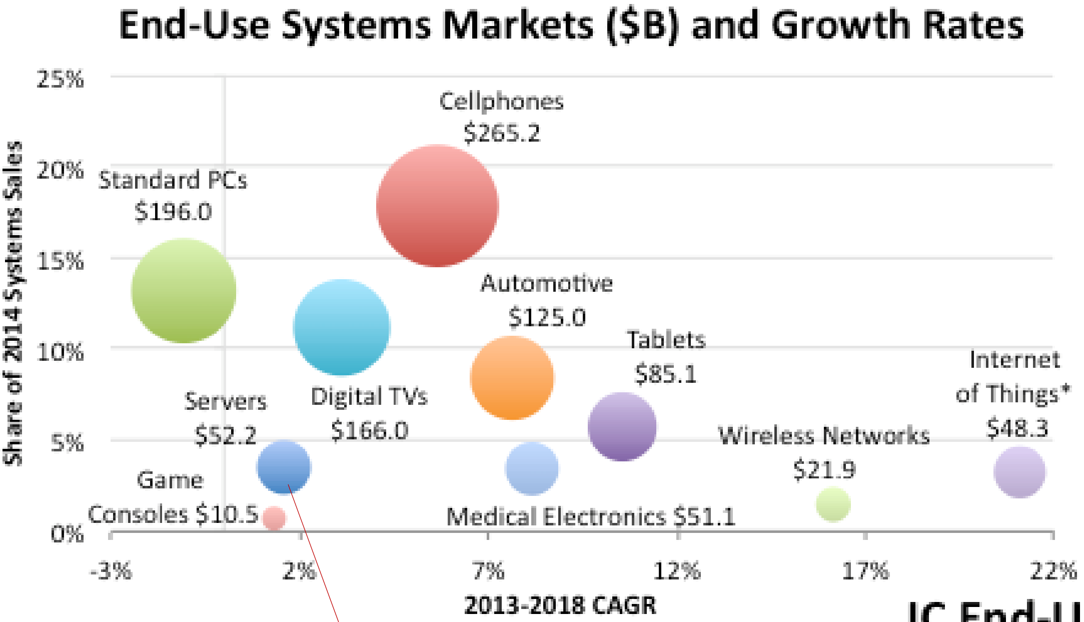

*Covers only the Internet connection portion of systems

Source: IC Insights

Electronic systems market value in 2014 was ~1.5 Trillion $

10 biggest segments
Moderate growth rates
Maturing markets

HEP is here
~15M$ out of 52B$

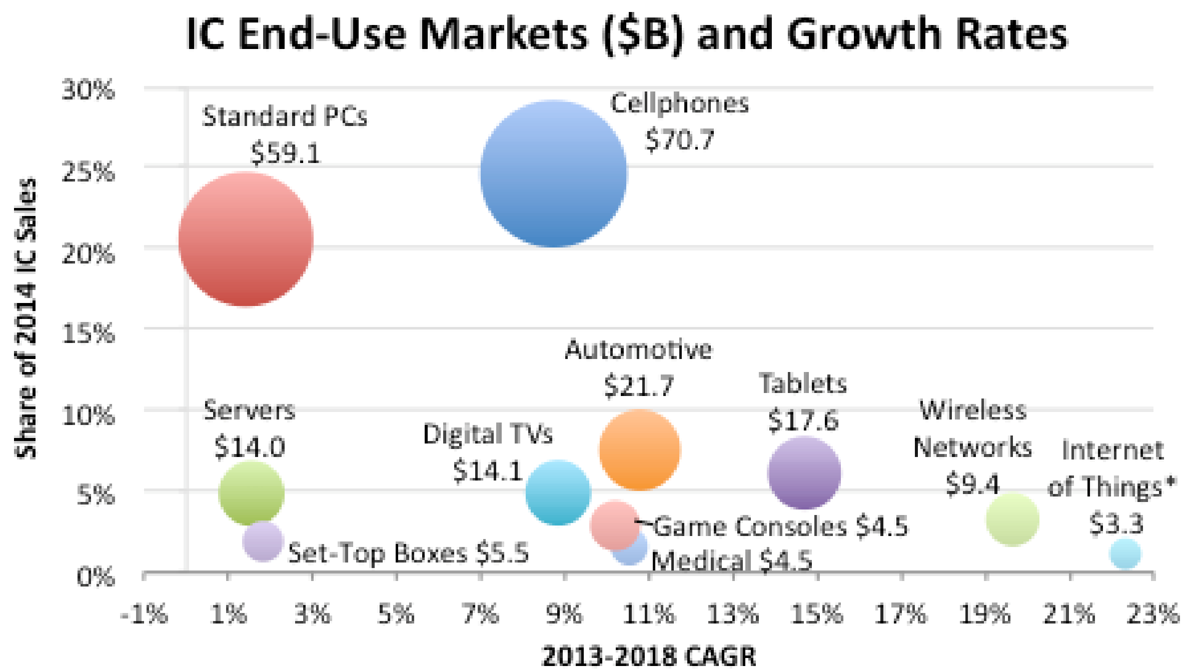

*Covers only the Internet connection portion of systems

Source: IC Insights

CAGR = Compound Annual Growth Rate

Notebook and Desktop Markets

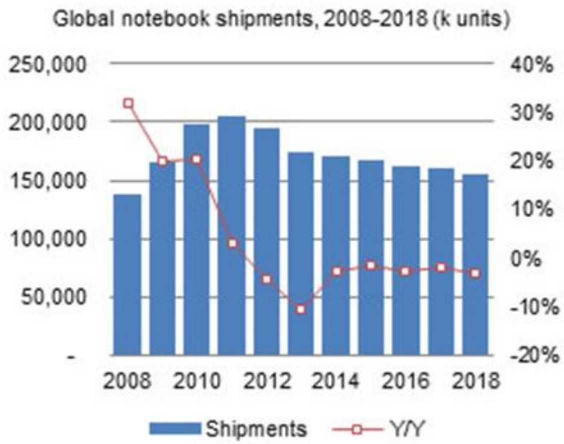

Stable markets , decreasing growth rates

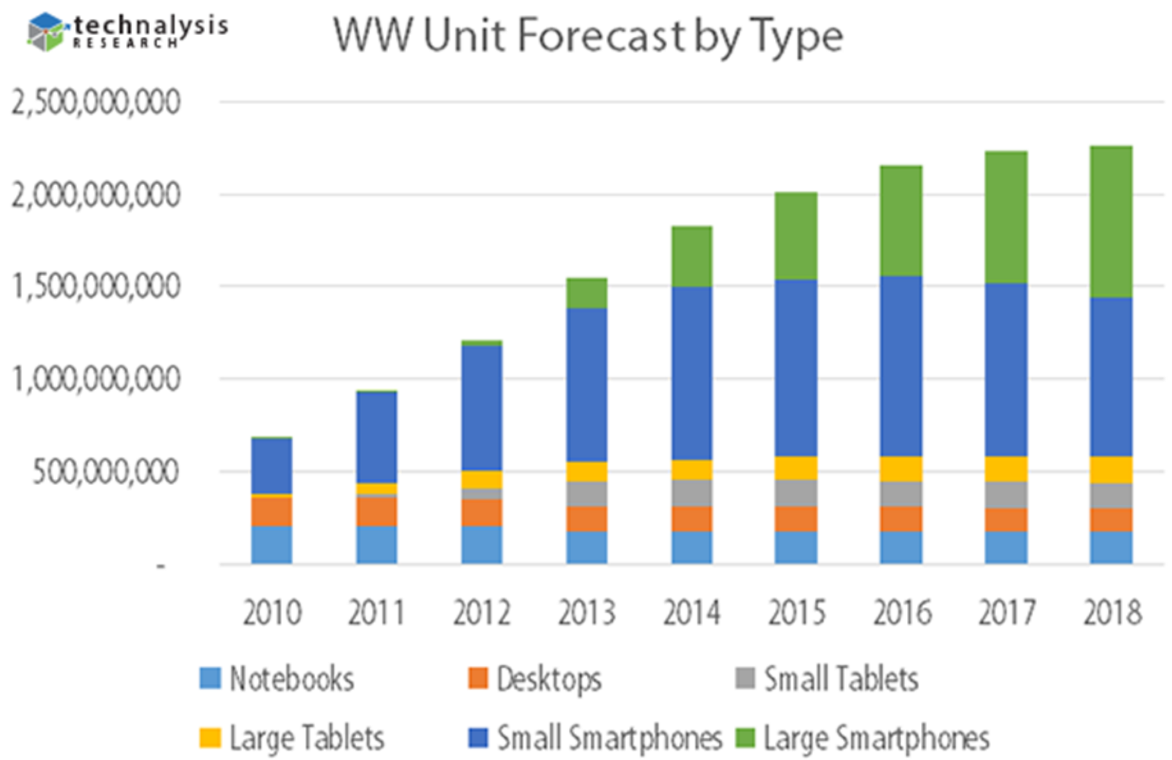

Important End-User sectors:

- Smartphones
- Tablets
- Notebooks
- Desktops
- Server
- HPC

Smartphone and Tablet Markets

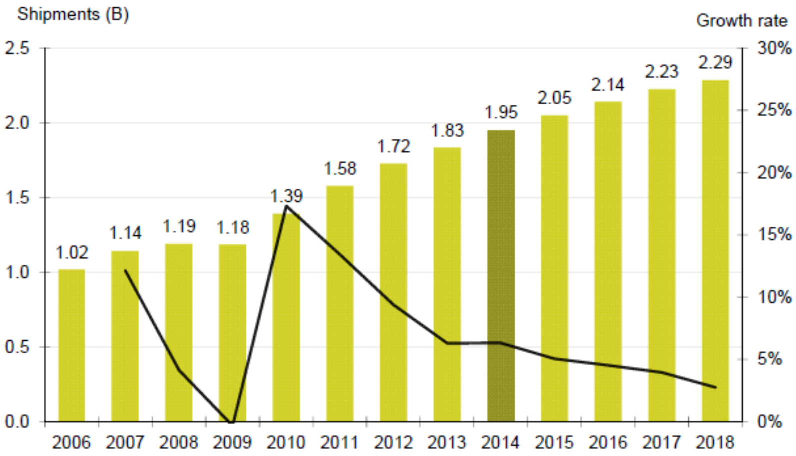

Figure 6-1. Global mobile phone shipments and growth rate, 2006-2018
Source: CCS Insight

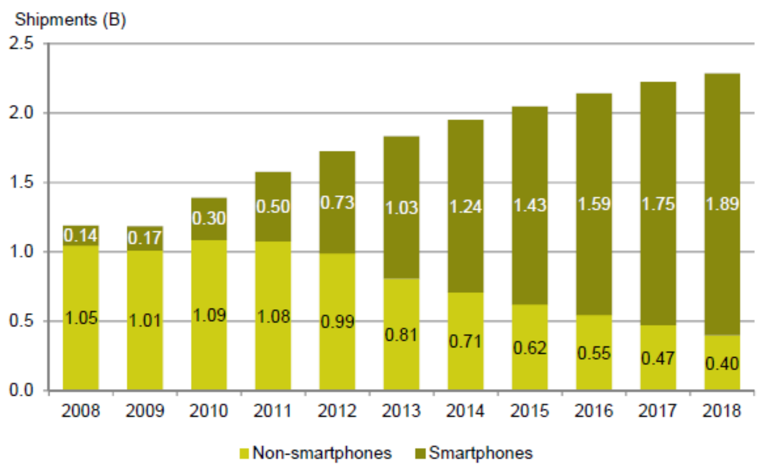

Figure 6-2. Global mobile phone shipments by type, 2008-2018
Source: CCS Insight

Smartphone install base in 2014: ~2B

Total cell phone install base 2014 : ~4.6B
Cell phone contracts 2014 : ~ 7B
PC and notebook install base 2014: ~ 3B

Replacement market
Stabilized market

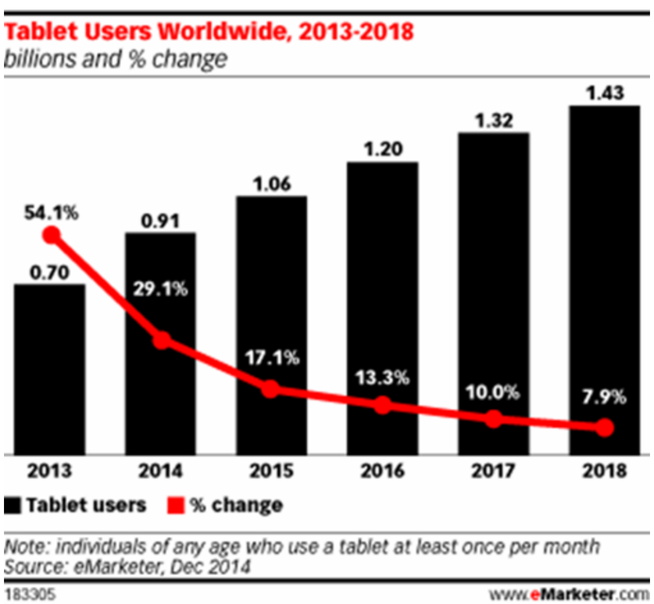

Compute Server Market Evolution

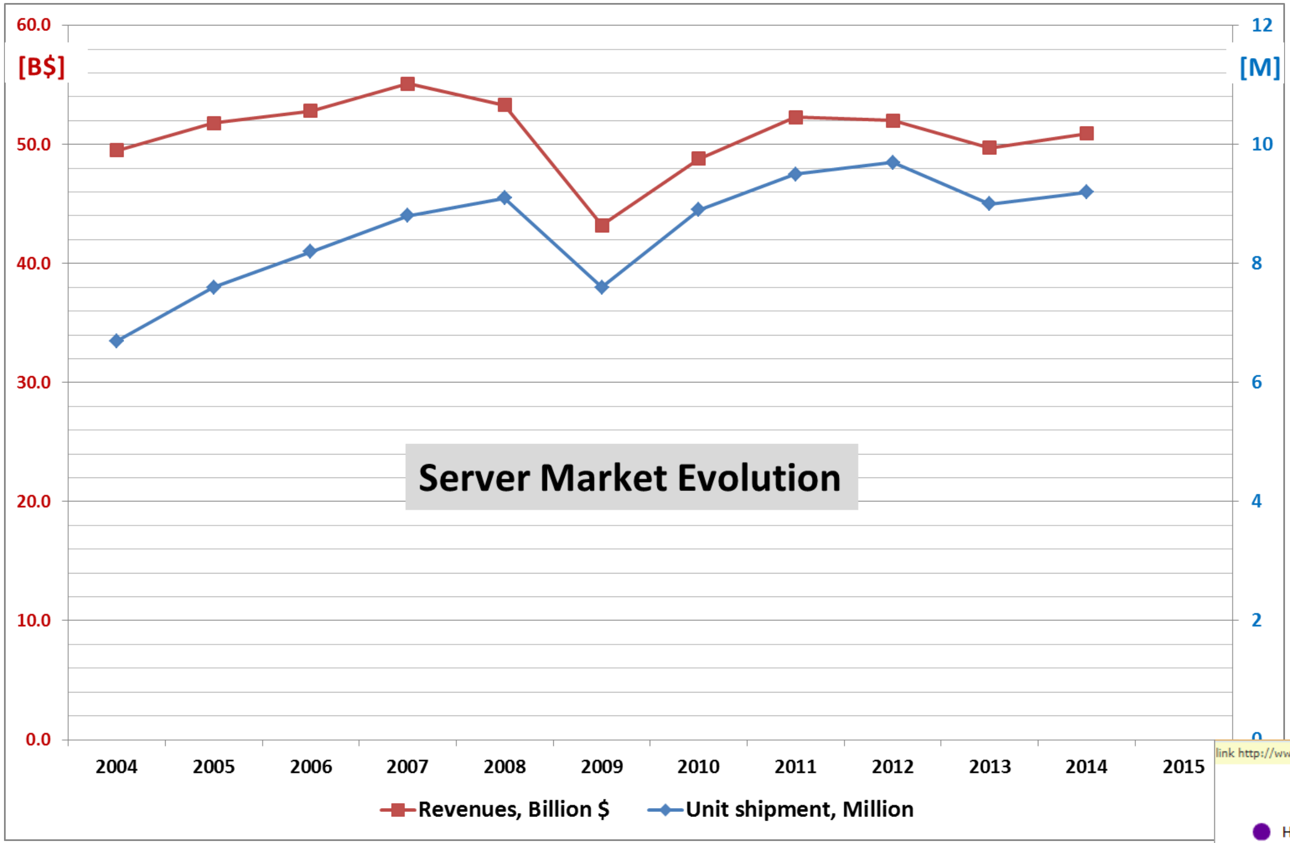

The HPC market is
Much smaller:
~11B$ yearly revenues
~140000 units sale

Very profitable market and stable, INTEL >98% share
(small share of IBM, ORACLE, AMD)

Mature replacement market

ODM original design manufactures with increasing market share
Special for hyperscale centers (Google, Facebook, etc.)

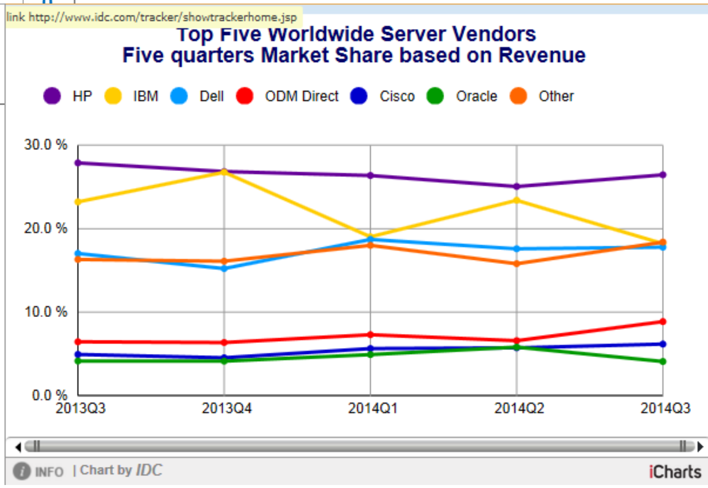

Leading Players

2014F Top 20 Semiconductor Sales Leaders ($M)

2014F Rank	2013 Rank	Company	Headquarters	2013 Total	2014 Total	2014/2013 % Change
1	1	Intel	U.S.	48,321	51,368	6%
2	2	Samsung	South Korea	34,378	37,259	8%
3	3	TSMC*	Taiwan	19,935	25,088	26%
4	4	Qualcomm**	U.S.	17,211	19,100	11%
5	5	Micron + Elpida	U.S.	14,294	16,614	16%
6	6	SK Hynix	South Korea	12,970	15,838	22%
7	8	TI	U.S.	11,474	12,179	6%
8	7	Toshiba	Japan	11,958	11,216	-6%
9	9	Broadcom**	U.S.	8,219	8,360	2%
10	10	ST	Europe	8,014	7,374	-8%
11	11	Renesas	Japan	7,975	7,372	-8%
12	12	MediaTek + MStar**	Taiwan	5,723	7,142	25%
13	14	Infineon	Europe	5,260	6,151	17%
14	16	NXP	Europe	4,815	5,625	17%
15	13	AMD**	U.S.	5,299	5,512	4%
16	17	Sony	Japan	4,739	5,192	10%
17	15	Avago + LSI**	Singapore	4,979	5,087	2%
18	19	Freescale	U.S.	3,977	4,548	14%
19	20	UMC*	Taiwan	3,940	4,300	9%
20	21	Nvidia**	U.S.	3,898	4,237	9%
		Top 20 Suppliers		237,379	259,562	9%
		Top 20 Suppliers Excluding Foundries		213,504	230,174	8%

*Foundry **Fabless

Source: IC Insights' Strategic Reviews Database

Very few companies can effort large R&D spending and the investments for IC fabrication units

TSMC and Samsung have started to build new fabs at a cost of ~16 B$ per unit
Takes 2 years to build

2014 Top Semiconductor R&D Spenders

2014 Rank	2013 Rank	Company	Region	IDM	FABLESS	FOUNDRY	2013 Semi Sales ($M)	2013 R&D Exp ($M)	2013 R&D/Sales (%)	2014 Semi Sales ($M)	2014 R&D Exp ($M)	2014 R&D/Sales (%)	2014/2013 % Change in R&D
1	1	Intel	Americas	•			48,321	10,611	22.0%	51,400	11,537	22.4%	9%
2	2	Qualcomm	Americas		•		17,211	3,395	19.7%	19,291	5,501	28.5%	62%
3	3	Samsung	Asia-Pac	•			34,378	2,820	8.2%	37,810	2,965	7.8%	5%
4	4	Broadcom	Americas		•		8,219	2,486	30.2%	8,428	2,373	28.2%	-5%
5	7	TSMC	Asia-Pac			•	19,935	1,623	8.1%	24,976	1,874	7.5%	15%
6	5	Toshiba	Japan	•			11,958	2,040	17.1%	11,040	1,820	16.5%	-11%
7	6	ST	Europe	•			8,014	1,816	22.7%	7,384	1,520	20.6%	-16%
8	9	Micron	Americas	•			14,294	1,487	10.4%	16,814	1,430	8.5%	-4%
9	14	MediaTek + MStar	Asia-Pac		•		5,723	1,110	19.4%	7,032	1,430	20.3%	29%
10	10	Nvidia	Americas		•		3,898	1,323	33.9%	4,348	1,362	31.3%	3%
		Top 10 Total					171,951	28,711	16.7%	188,523	31,812	16.9%	11%

Source: Company reports, IC Insights' *Strategic Reviews* database

Market Dominance

Only a few large companies are dominating the various components markets

Processors	INTEL, Qualcomm, Samsung, AMD
Graphics	INTEL, Nvidia, AMD
Hard Disk Drives	Western Digital, Seagate, Toshiba
DRAM memory	Samsung, SK Hynix, Micron
NAND Flash memory	Samsung, Toshiba, SanDisk, Micron, Hynix, INTEL
Solid State Disks	Samsung, INTEL, SanDisk, Toshiba, Micron
FPGA	Xilinx, Altera (currently being bought by INTEL)
Tape Storage	HP, Fuji, IBM, SpectraLogic ORACLE, IBM

RoI Return-on-Investment is the keyword
Few companies capable of large scale investments, majority fabless companies
Favour evolutionary (adiabatic) changes of technology
Clear bias against 'disruptive' new technologies
(memristor, holographic storage, DNA storage, quantum computing, non-volatile memory, etc.)

e.g. Yearly revenues: Samsung 209 B$ INTEL 56 B$

Processor Technology I

Shrinking by a factor 2 every 2 years. 65nm node in 2006 --> 14nm node in 2014

The '14nm node' is a process name, not a description of the real feature sizes.
On a 14nm chip there are NO 14nm structures
There is no standard or a detailed definition
Still very, very small feature sizes

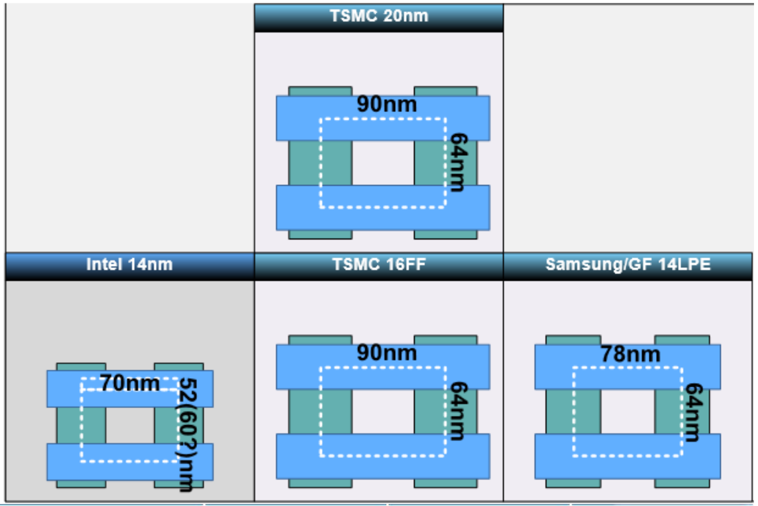

Minimum Feature Size

	22 nm Node	14 nm Node	Scale
Transistor Fin Pitch	60	42	.70x
Transistor Gate Pitch	90	70	.78x
Interconnect Pitch	80 nm	52 nm	.65x

Intel Has Developed a True 14 nm Technology with Good Dimensional Scaling

Foundry Technology Node:	14 nm	10 nm	7 nm
Gate length, L_G	25 nm	20 nm	15 nm
Fin width, W_{fin}	~10 nm	~8 nm	~6 nm
Equivalent oxide thickness	0.9 nm	0.85 nm	0.8 nm

Transistor Fin Improvement

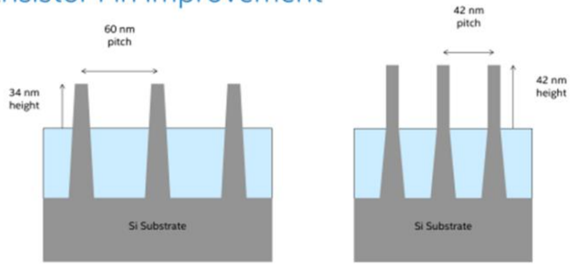

22 nm Process

14 nm Process

Taller and Thinner Fins for Increased Drive Current and Performance

X-SEM Images

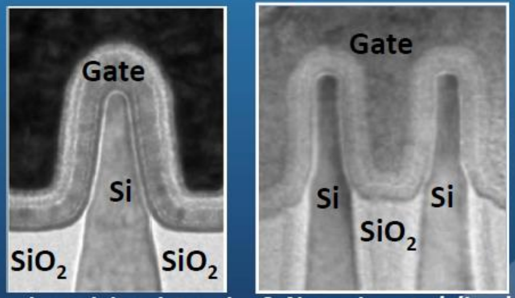

C. Auth *et al.* (Intel Corp.) *VLSI Symp.* 2012

S. Natarajan *et al.* (Intel Corp.) *IEDM* 2014

Processor Technology II

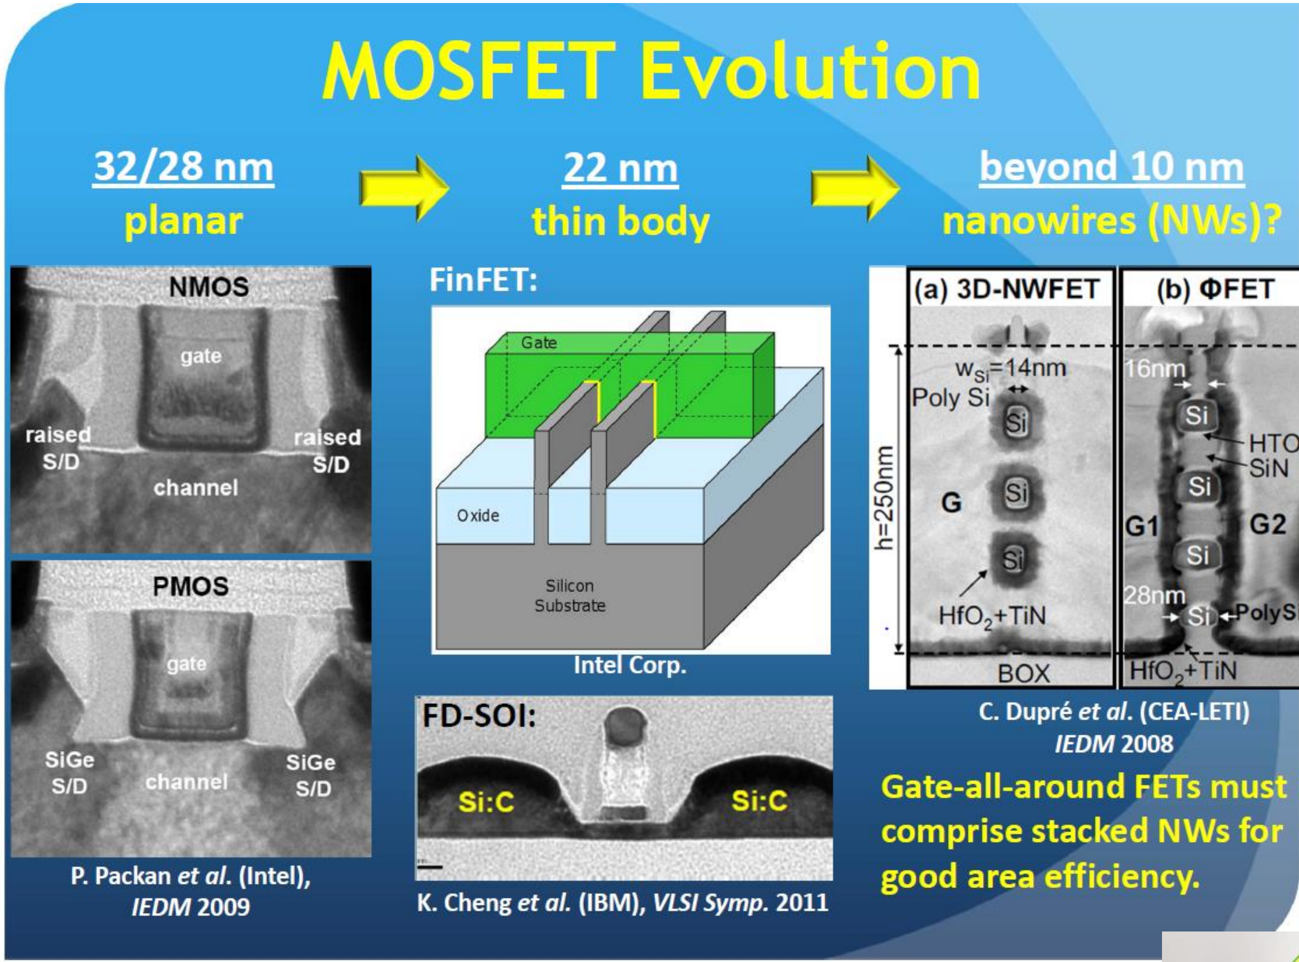

Very sophisticated lithography techniques, double patterning

Still using 193 nm light source
EUV Extreme Ultraviolet not yet in production

3D-FinFET transistor

Leakage current reduction

2014 - 2015

INTEL (x86) → 14 nm process node
Samsung (ARM) → 16 nm process node

2016 → 10nm process node

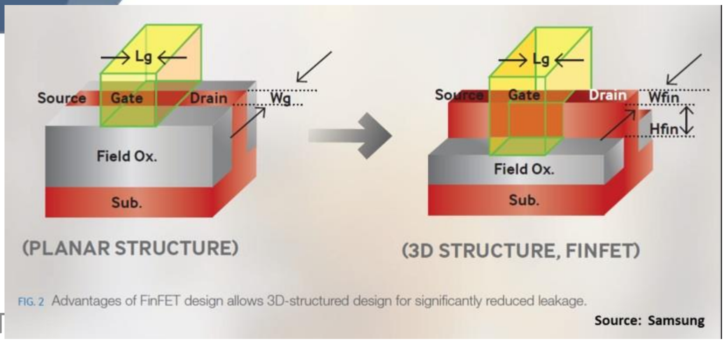

Processor Technology, Moore's Law

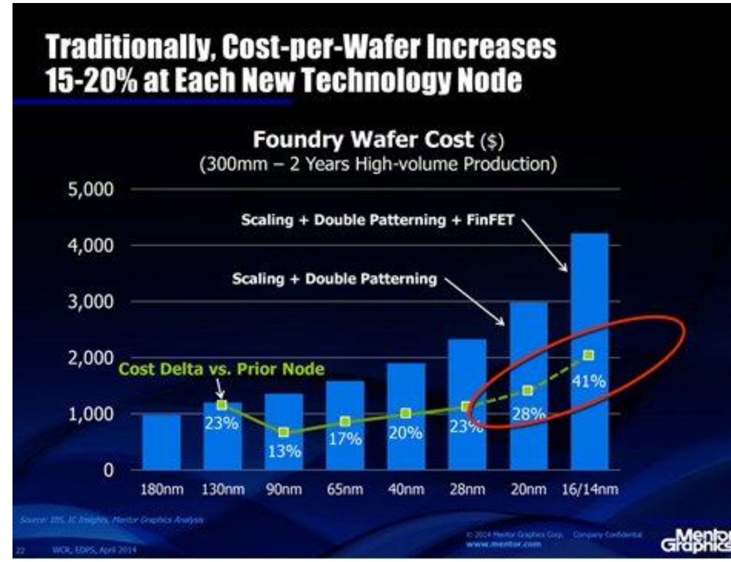

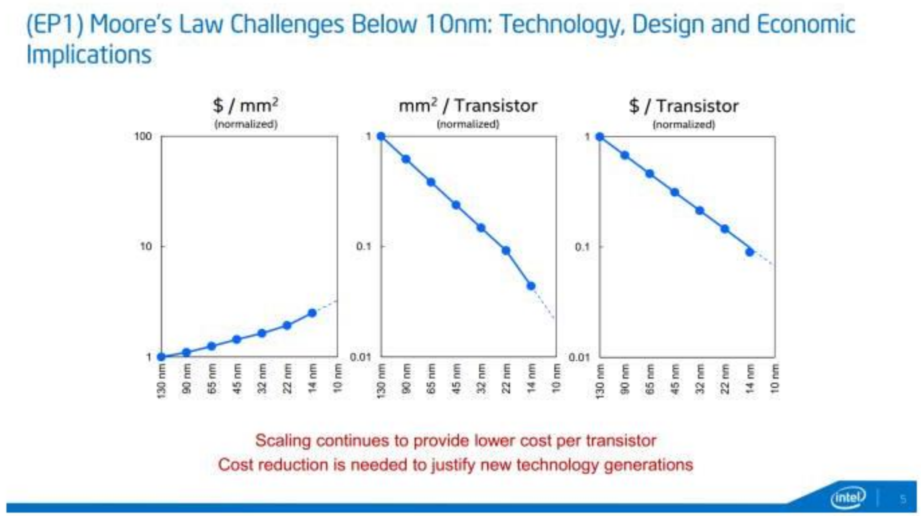

Quite some discussion in 2014 about the end of Moore's Law

Moore's Law is about the production cost of transistors not about the sales cost of processors

INTEL claims to overcome this up to the 10nm node scale

Processor Technology, architecture

- **Kept the pipeline stages at 14 for the last few generations**
- **Stable frequencies around 3+- 0.5 GHz**
- **Number of cores per processor is increasing in a linear fashion, 1-2 per year**
  market volumes, best price/performance → 2/4-cores in smartphones, 4-cores in notebook+desktops, 8-cores in servers
  high end, smaller volumes → octo-core in smartphones (actually this is 2 x 4, big-little concept), 6-cores in desktops, 18-cores in Xeon servers, 32-cores Oracle SPARC M7
- **Increase vector length and sophistication of SIMD operations, steady IPC increase**
- **Haswell running with up to 32 Instructions per Cycle (IPC)**

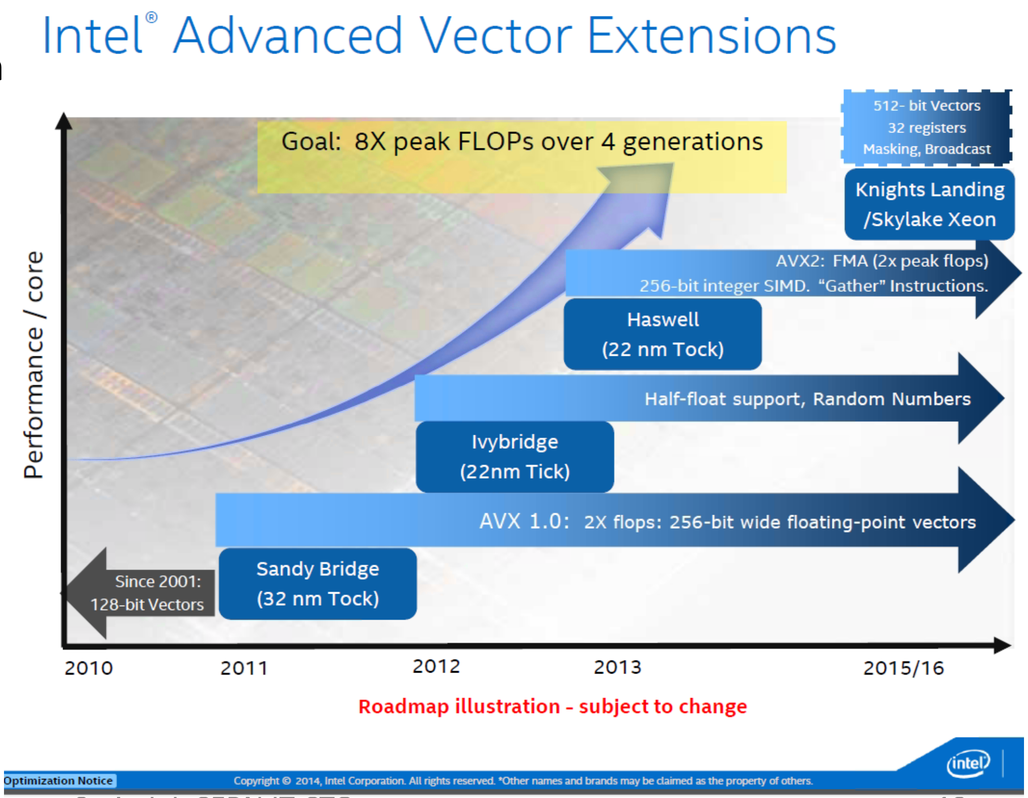

Processor Technology, prices

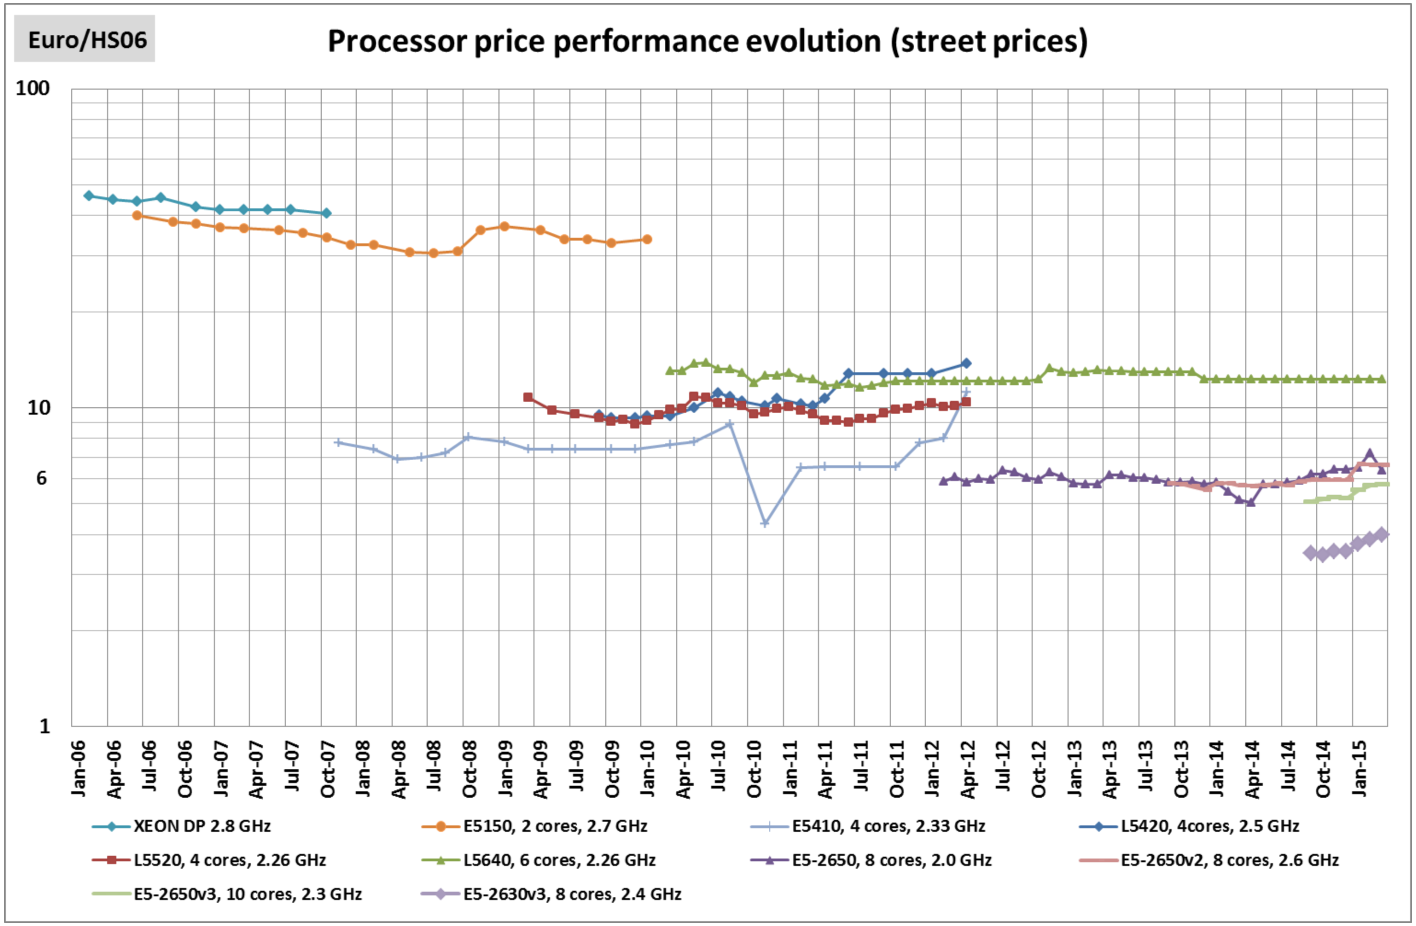

Processors from CERN purchases
Flat prices per processor generation
Server processor prices are more defined by the market then the technology

INTEL data centre group results for Q4 2014 : Revenue = 4.1 B$ Profit= 2.2B$
(~5 M server processors) highly profitable market

CPU Server Cost Evolution

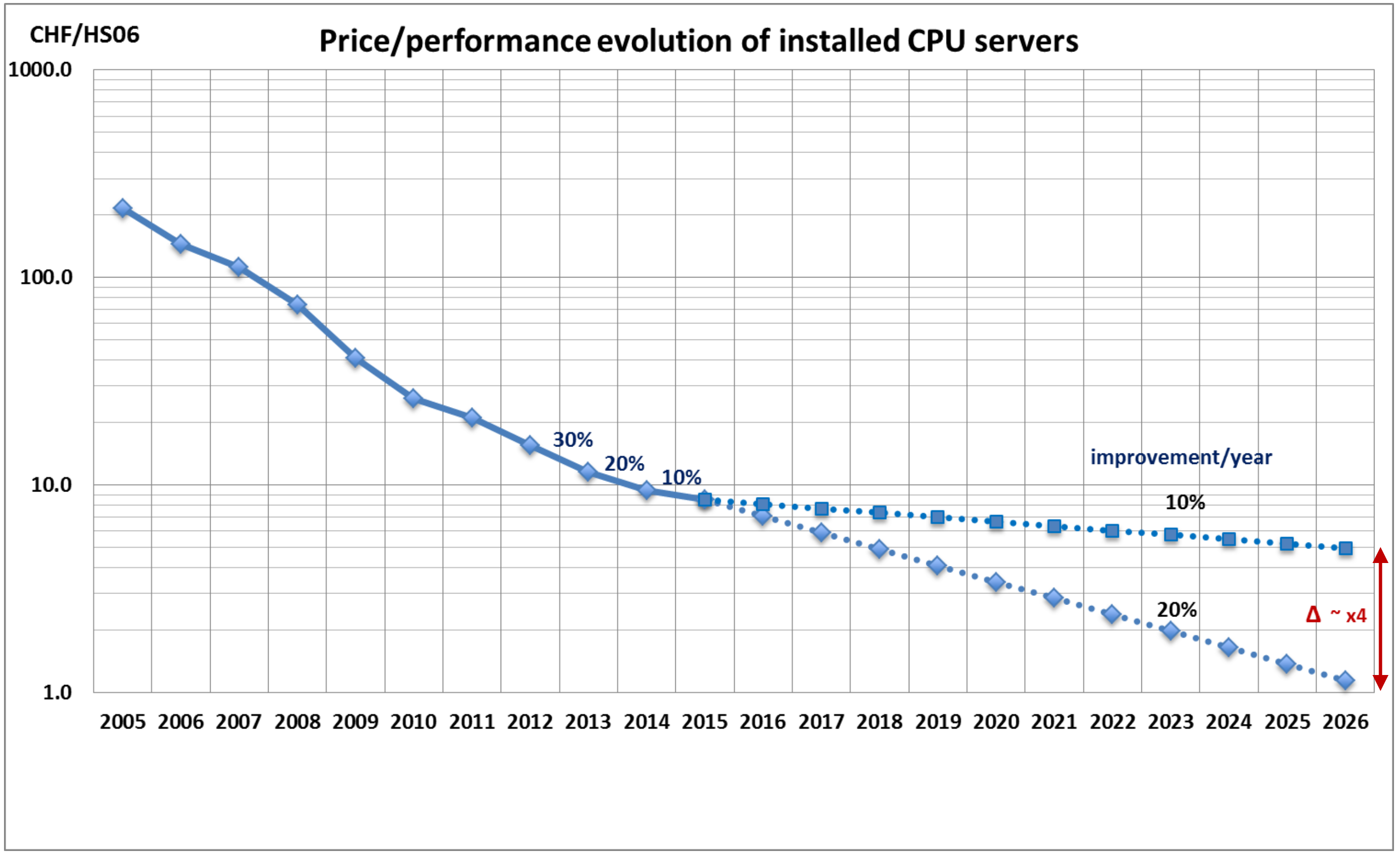

2015 to 2026
Improvement = factor 7.5
At 20% increase/year

CERN purchases, server nodes, latest version e.g. dual Haswell E5-2630v3, 64 GB memory, 1 Gbit NIC , 2 x 2TB disks
Network costs are not included, 10% effect

Purchase cycles are not directly overlapping with technology cycles

Possible Architecture changes: move to 10 Gbit, SSD disks, SMT on or off

Micro Server Developments

- Cavium,   48-core server chips based on ARM (ThunderX SoCs)
- Gigabyte server motherboard released using X-Gene 1 (AppliedMicro), 8-core ARMv8 45 W 2.4 Ghz
- HP Moonshot,   AppliedMicro X-Gene ARM processors
- Calxeda went bust in early 2014
- AMD is very late with their ARM product
- Many INTEL product releases

Facebook just dropped ARM plans in favour the new INTEL XEON D server chips
(ARM power advantage diminishing, software porting is the issue)

New generation of Windows Surface Tablet has dropped ARM

INTEL 'supported' 40 million tables with x86 processors in 2014 (4.2 B$ contra-revenue !)
(comparison: AMD stock market value is about 4 B$)

Game changer most likely only if and when Samsung buys AMD
→ R&D investments

New Processing Architectures I

- Micron's Automata Processor reconfigurable, massive parallelism; for bioinformatics, pattern recognition, data analytics and image processing

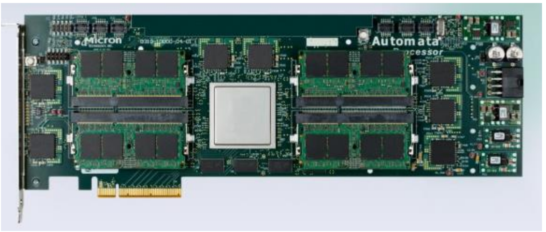

- Optalysys, Laser plus liquid crystal spatial light modulators
  UK technology company

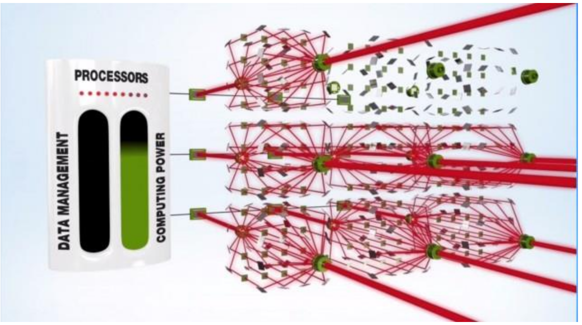

- IBM research, neuromorphic chips
  4096 cores, 1 million neuron, 5.4 B transistors, 72 mW

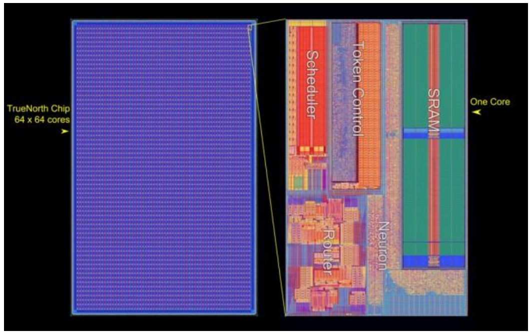

- Qualcomm cognitive compute Platform (Zeroth),
  along the Snapdragon 820 ARM architecture
  deep learning for smartphones

- D-Wave Quantum Computing (Maybe !, still controversial)

New Processing Architectures II

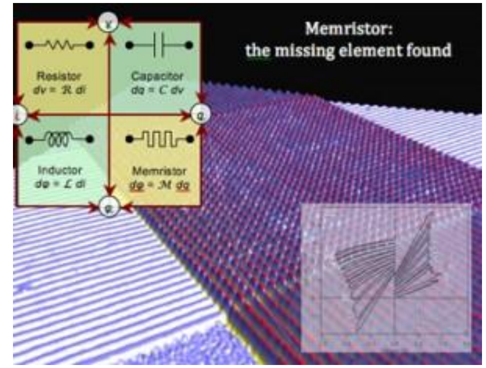

The Machine
based on silicon photonics interconnects and memristors
as active components (HP)
Completely different programming model: Linux++
Started in 2012, prototype in 2016
Memristor concept from 1971, implemented in HP Labs (2008)

DARPA initiative
Petaflops On Desktops: Ideas Wanted For Processing
Paradigms That Accelerate Computer Simulations
Includes the use of analogue circuits

DIGITS DevBox from NVIDIA, GPU based, special libraries → deep learning applications

Soft Machines , Variable Instruction Set Computing (VISC) virtual cores implemented in hardware

GPU processing and Markets

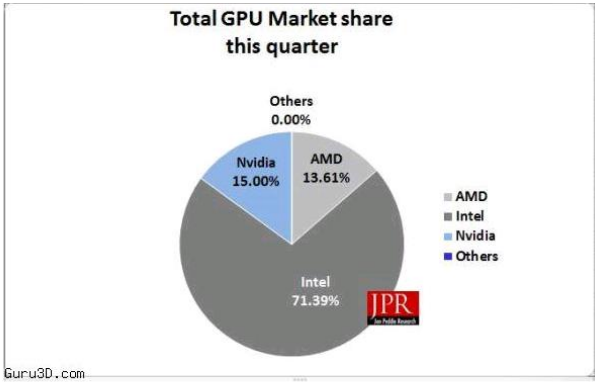

Q4 2014

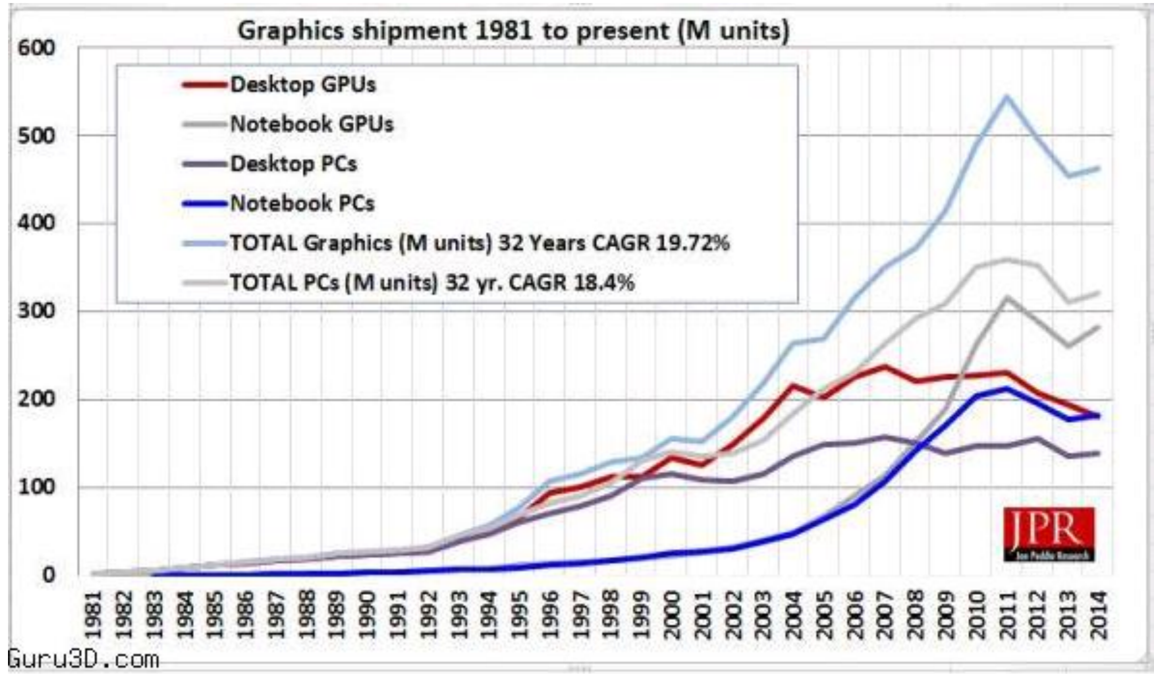

450 M GPUs sold per year, compared to ~10000 very high end GPUs (HPC)

GPU technology still at the 28nm level

AIB = Add-in-boards
Discrete graphics cards

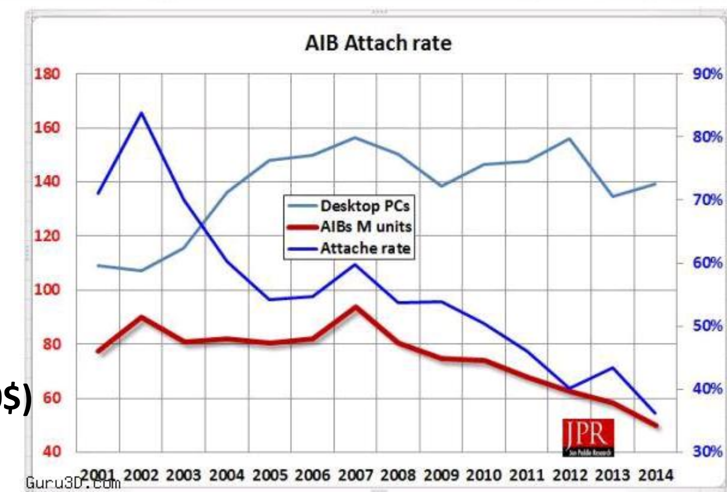

Most likely skip the 20nm step and move directly into 16nm

16 B$ fab investment from TSMC

Latest 28nm cards from Nvidia:
Titan X (8B transistors, 3000 cuda cores, 8 TF SP, 0.2 TF DP, 1000$)
K80 (14B transistors, 5000 cuda cores, 8.7 TFlops SP, 2.9 TFlops DP, 7000$)

Constant decrease of discrete graphic card sales
CPU+GPU integrated from INTEL increasing

Split between gaming and HPC market

Tape Storage I

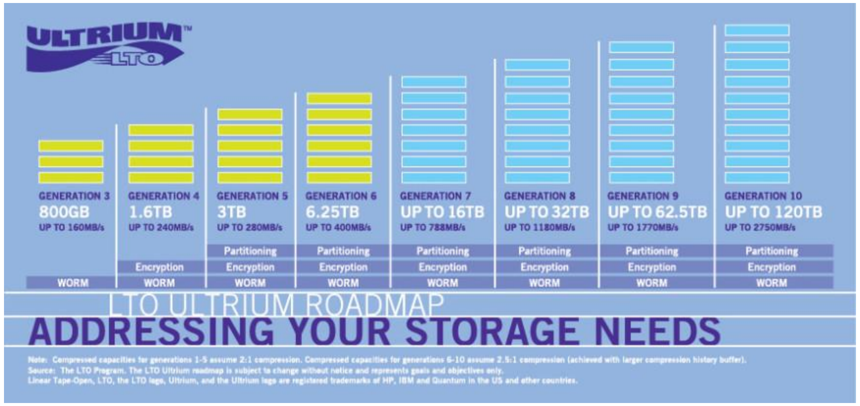

- ❑ LTO has > 96% of the market , (LTO-6, 2.5TB Cartridges)
- ❑ Enterprise tapes (ORACLE- 8.5TB, IBM – 10TB) niche products
- ❑ TDK&Maxwell stopped producing tapes
- ❑ R&D looks okay, 220 TB (IBM/Fuji) and 185 TB (Sony) tape in the labs
- ❑ LTO roadmap lately extended to 10 generations, but steady decrease of revenues
- ❑ LTO 6 capacity was reduced (3.2 → 2.5 TB)

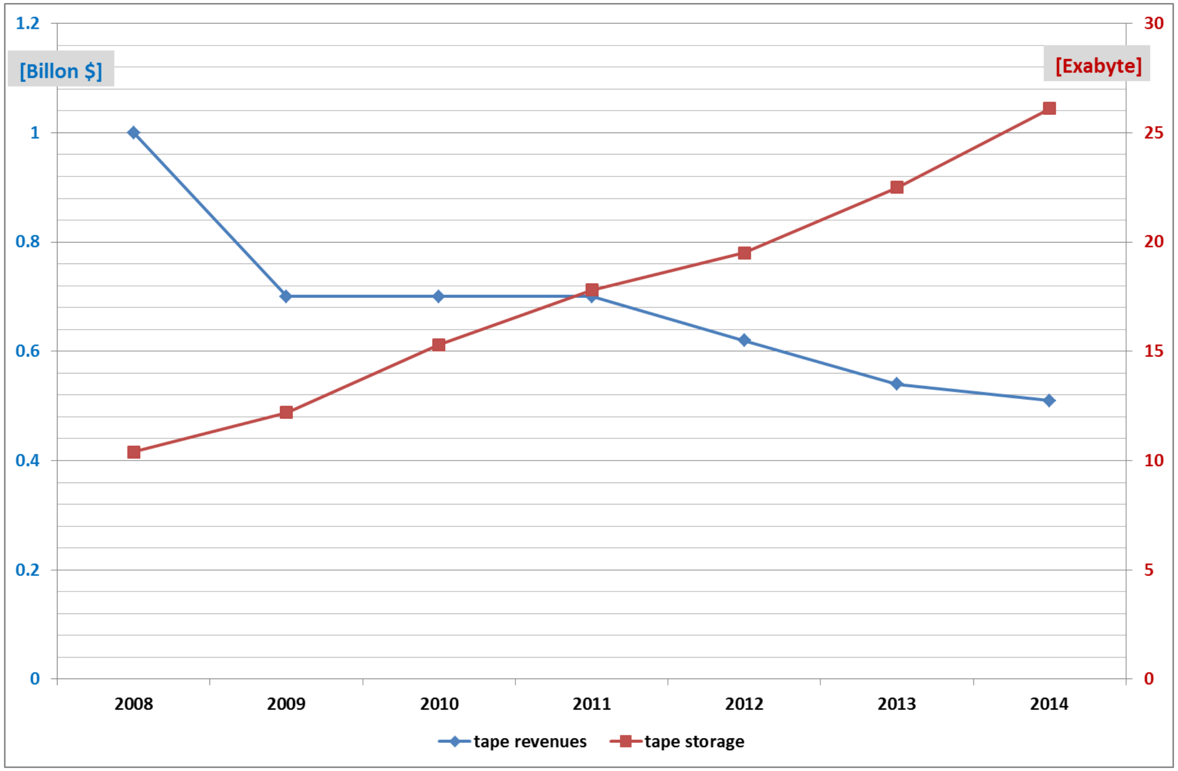

Source: Santa Clara Group

Tape Storage II

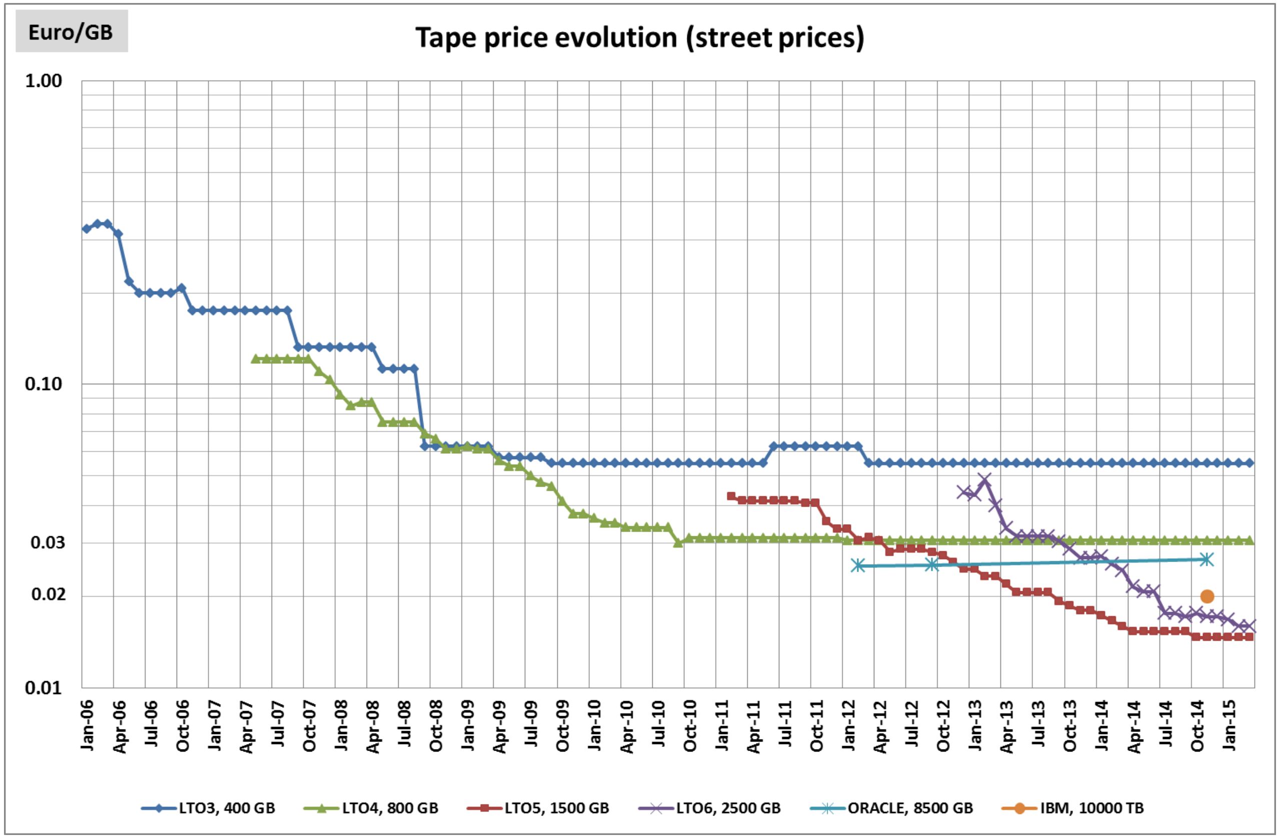

Assuming a constant evolution of the LT0 technology, with a new Generation every two years
→ 2025
192 TB tape x32 cost improvement

3 years 50 TB tape x8

LTO approaching 1 cent/GB, steady cost decrease
Enterprise more expensive, but can be re-used with next generation

Size difference (LTO6 2.5 TB, IBM/Oracle 8.5-10 TB) == infrastructure cost difference (silos, drives, maintenance)

Storage Components: DRAM Memory I

Memory production has moved from 25/28nm to 20nm in 2014

The same companies produce NAND and DRAM
Shifting capacities
Weak PC market, stable server market
Reduced capacity
→ Volatile DRAM prices

Focus on speed improvement especially in the low-power memory formobile devices

DRAM market size ~42 B$ in 2014

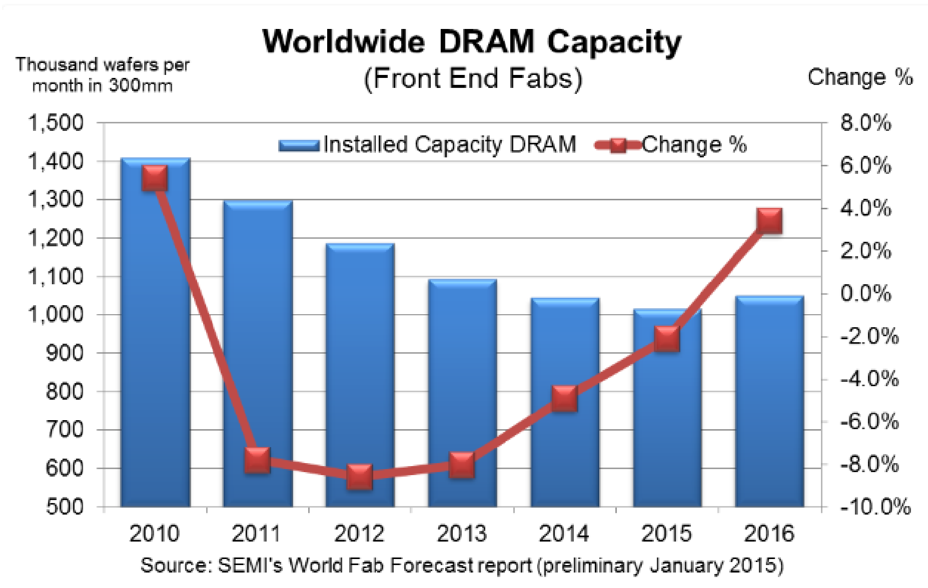

Figure 3: Worldwide DRAM capacity for Front End facilities in 300mm equivalent wafers per month and change rate in percent (Source: SEMI, 2015)

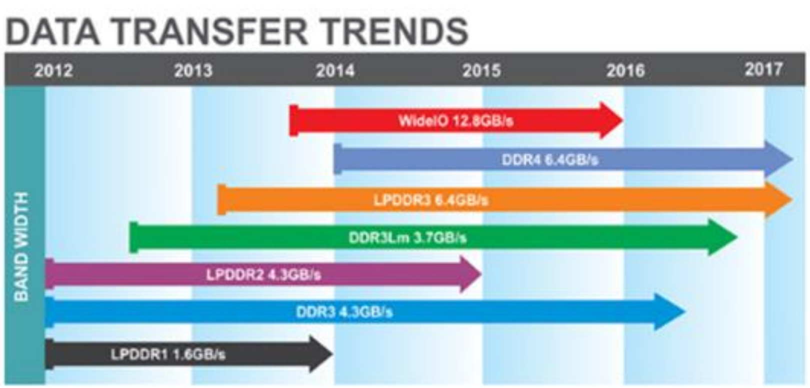

Source: Techinsights

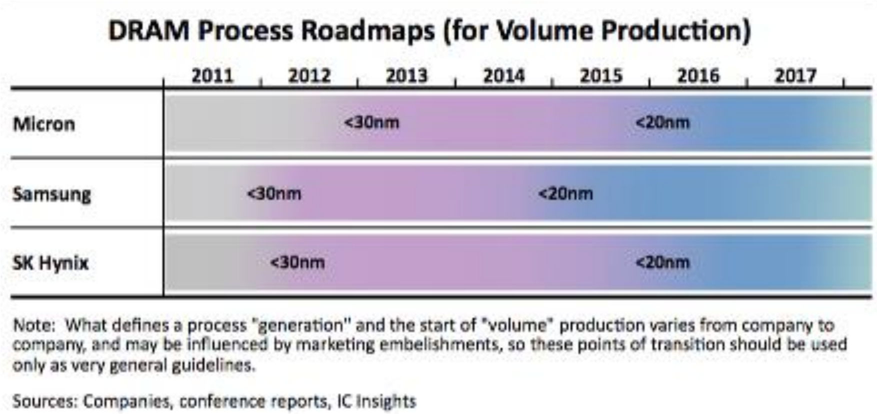

Storage Components: DRAM Memory II

3D memory delayed, coming this year,
solves data transfer issues, density

Microns Hybrid Memory Cube concept
factor 15 memory speed improvements

Focused on the server and HPC
area. Memory wall problem

Nvidia new Pascal GPU technology
in 2016 will use memory stacks

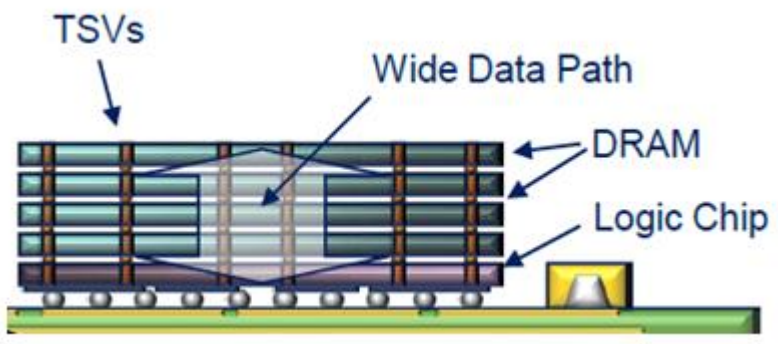

Memory stack
TSV Through Silicon Via

Storage Components: DRAM Memory III

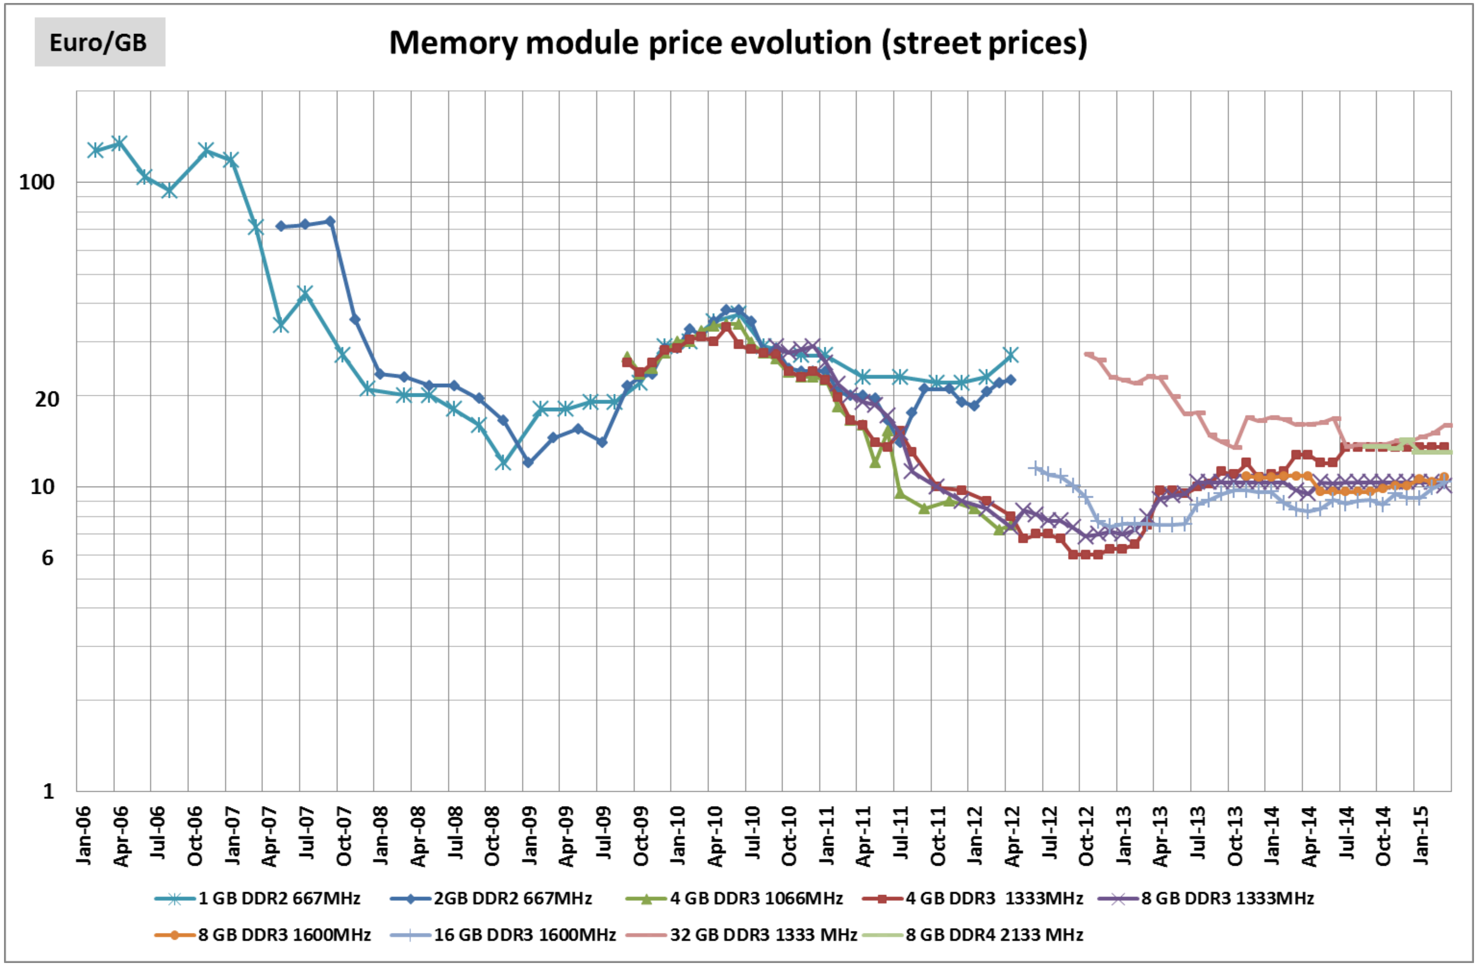

Volatile memory DRAM market

Side effects: Apple will consume 25% of the worldwide DRAM production in 2015
→ Shift to mobile DRAM, some shortage in PC RAM and server RAM expected

Storage Components: NAND Flash Memory I

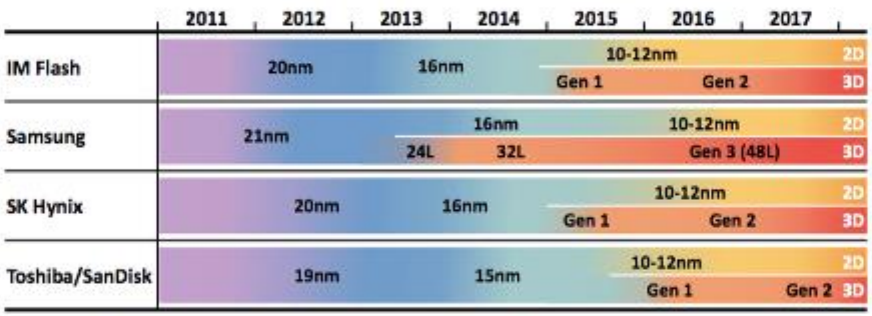

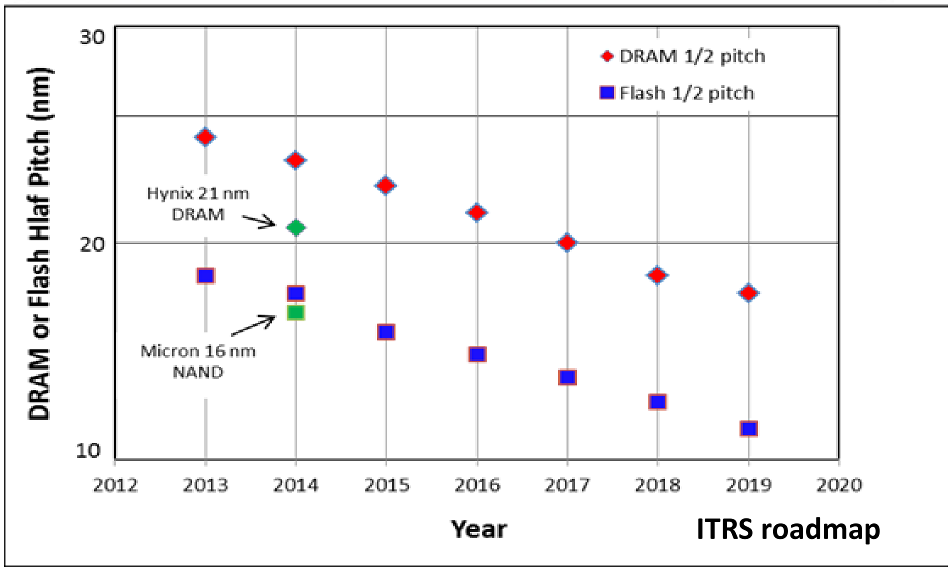

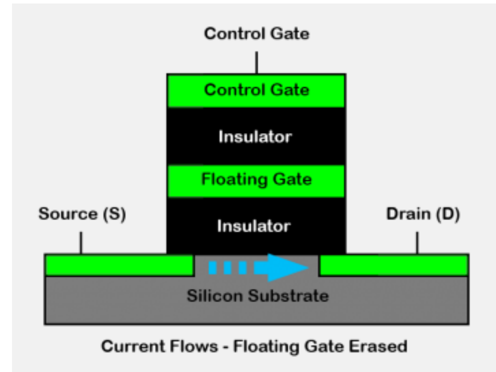

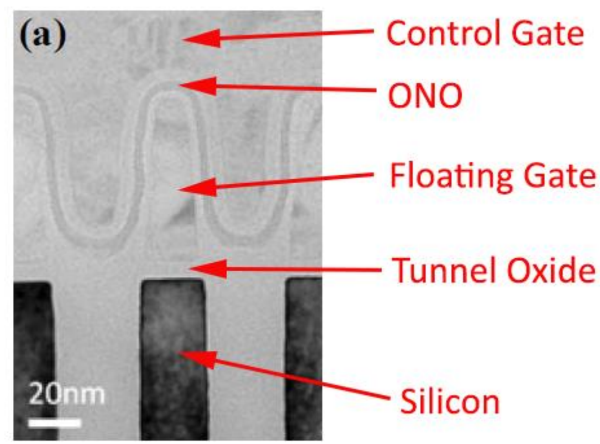

Micron has moved to 15nm technology
3D-NAND flash 128 Gbit chips

Commercially the limit for 2D flash is 15nm

Storage Components: NAND Flash Memory II

SLC 1bit/cell 100000 cycles
MLC 2 bit/cell 5000 cycles
TCL 3 bit/cell 1000 cycles

Sandisk BiCS 3D-NAND

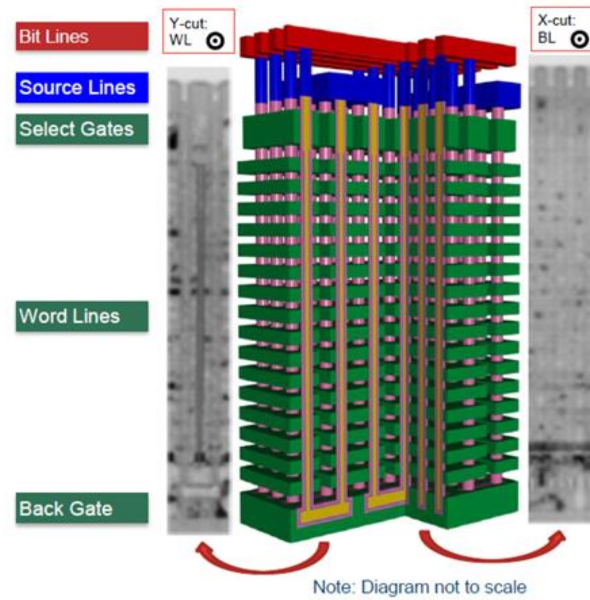

FG Limitation : Number of Electrons

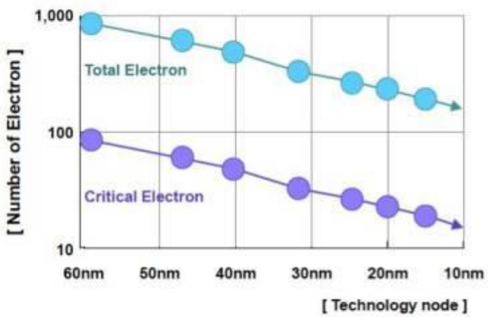

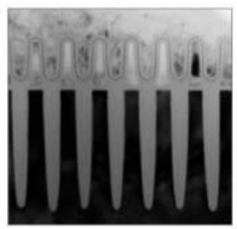

How to Manage 10 electrons in sub-1xnm design rule?

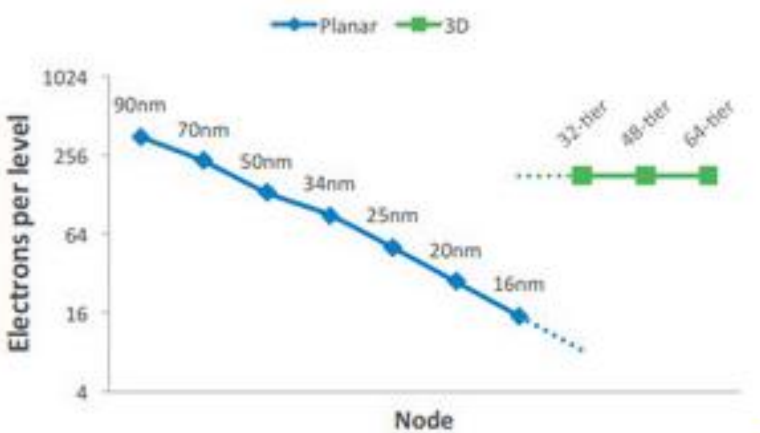

Move to 3D and increase 2D structures

INTEL/Micron have produced 32 layer 3D-NAND

Samsung already shipping products
V-NAND 32 levels 32nm production node

Toshiba is moving to 48 layers

NAND Flash Market

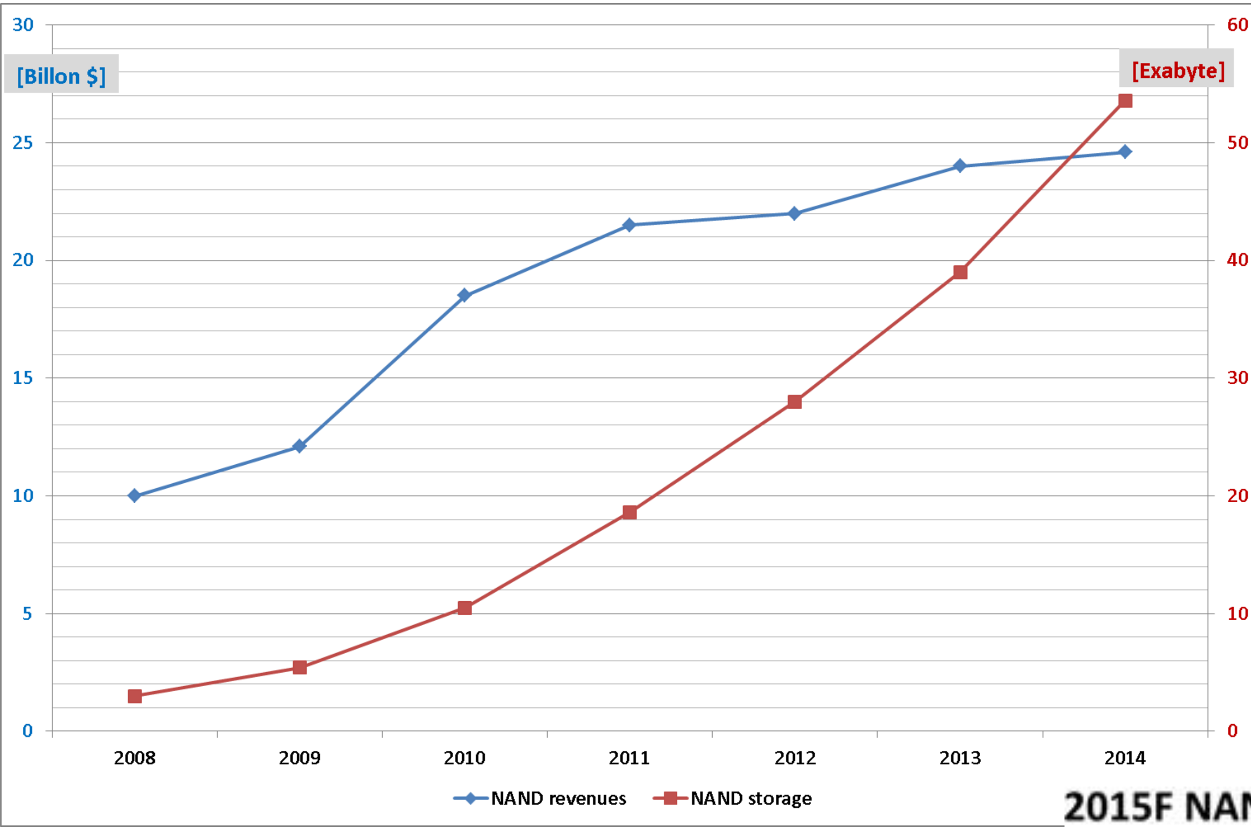

Revenues are becoming flat

2015F NAND Market by Application ($27.2B)

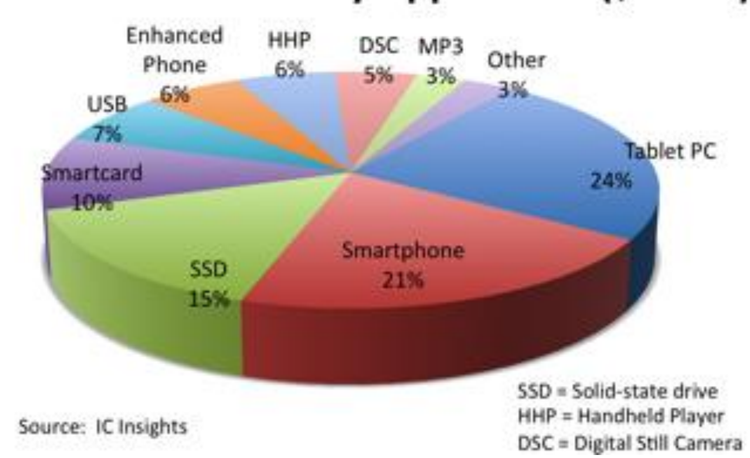

Only 15% of the yearly NAND capacity is for SSDs

Storage Components: Non-Volatile Memory I

Contenders :

3 types of MRAM (Magnetoresistive RAM)
Spin-Transfer-Torque, field driven, magneto thermal

PCRAM (Phase-Change RAM)

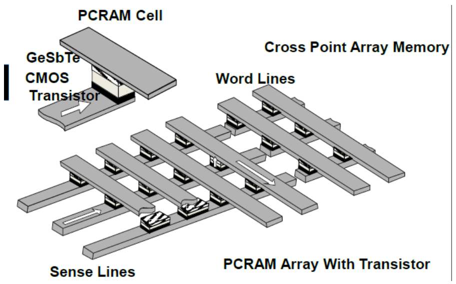

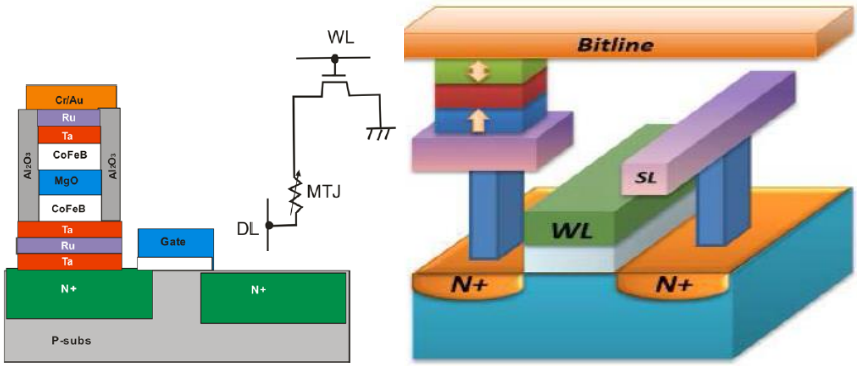

ReRAM/RRAM   (Resistive RAM)
CBRAM (Conductive Bridge RAM)

Memristor

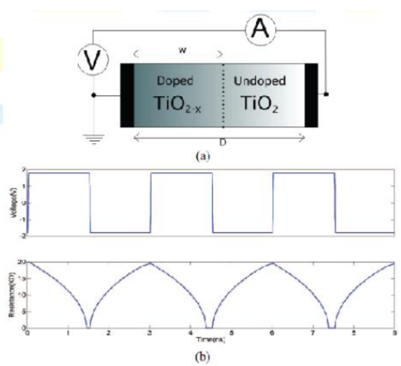

FIGURE 1. (a) Characterizing the memristor and (b) change of resistance when a 3.6 V p–p square wave is applied.

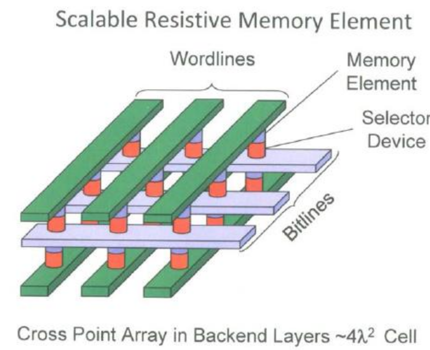

Source: Flash Memory Summit 2013

Storage Components: Non-Volatile Memory II

NVM market in 2014 is 65M$
Comparison: DRAM 42 B$, NAND 25B$
Expected to rise to 7 B$ in 2020

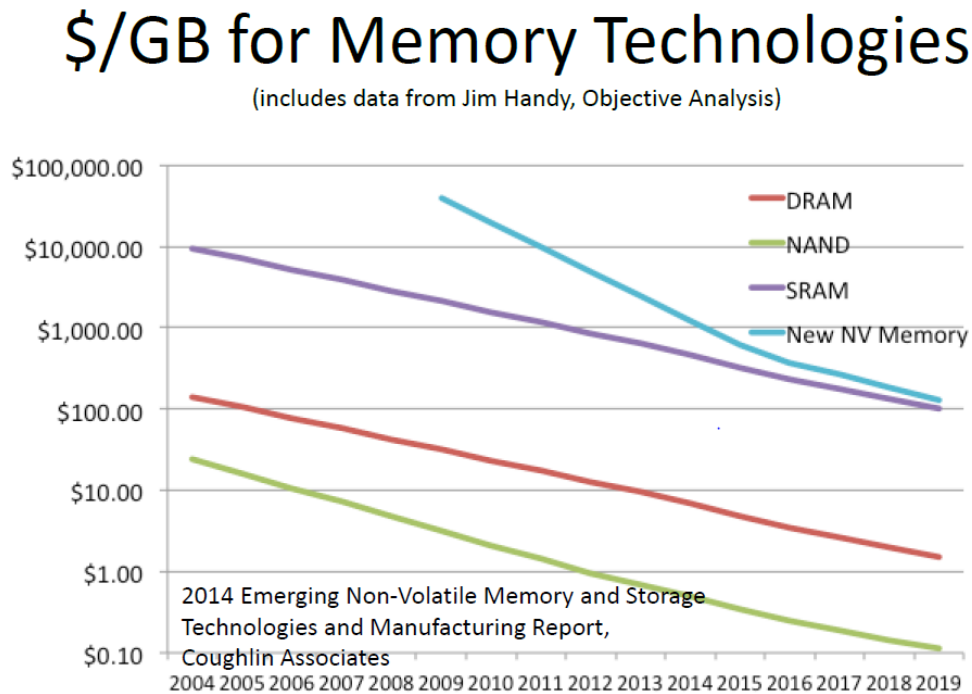

Everspin is producing MRAM since 2008
64 Mb chips in 90nm technology

Micron/Sony have just shown 27nm 16 Gbit CBRAM

Micron, the main PCM memory promoter dropped this activity in 2014 focused on 3D-NAND

Complicated and ‘disruptive’ fabrication process

Storage Components: Hard-Disk-Drives I

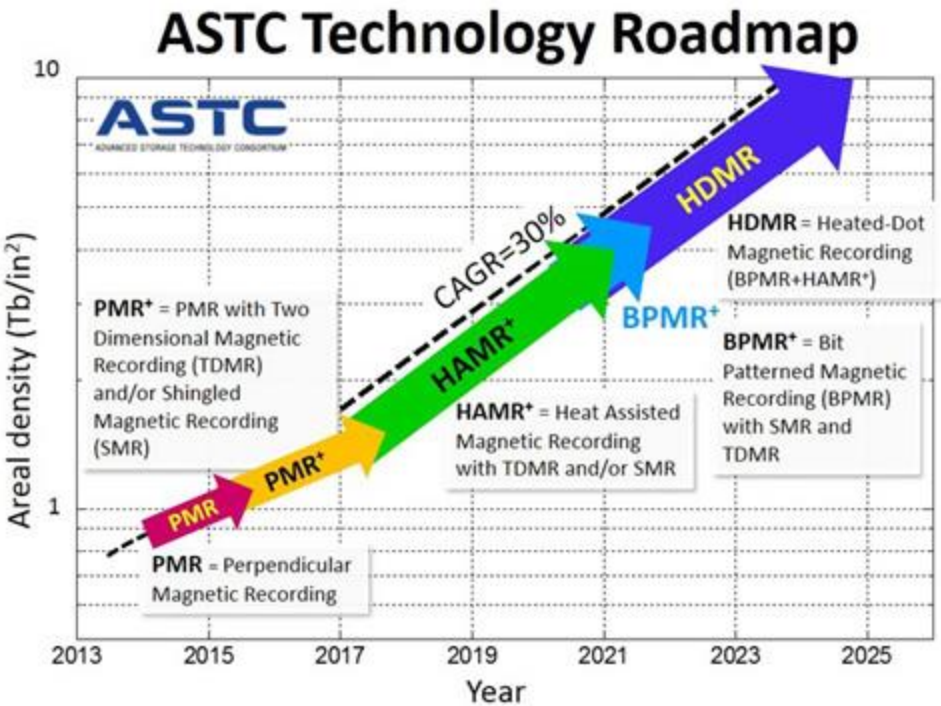

100 TByte drives in 2025 (possible)

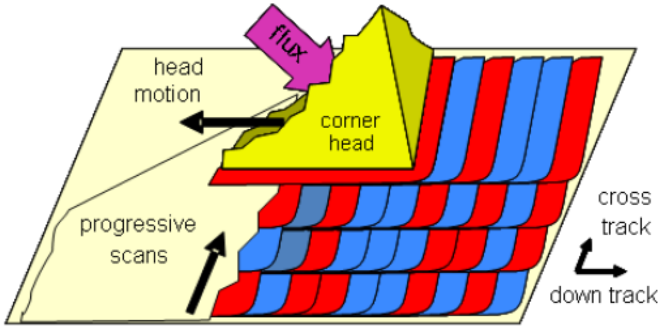

SMR

- PMR at it's limit , current drives at 0.75 Tbit/in2, max is about 1 Tbit/in2
- The density increase rate has slowed down considerably over the last years
- Shingled Magnetic Recording (1D, 2D) now in the market (e.g. 8 TB Seagate drives) extends the limit to 1.5 – 2 Tbit/in2 → increased surface density Good read, but restricted write performance. Sophisticated controller
- More platters per disk, Helium filled (e.g. 6 TB HGST) drives) → increased volume density
- HAMR prototypes already shown 3 years ago (Seagate 1 Tbit/in2), but very sparse information about the current roadmaps. Introduction in 2017 !?

no principle technology problems, HAMR and BPMR are sophisticated and very expensive

Storage Components: Hard-Disk-Drives II

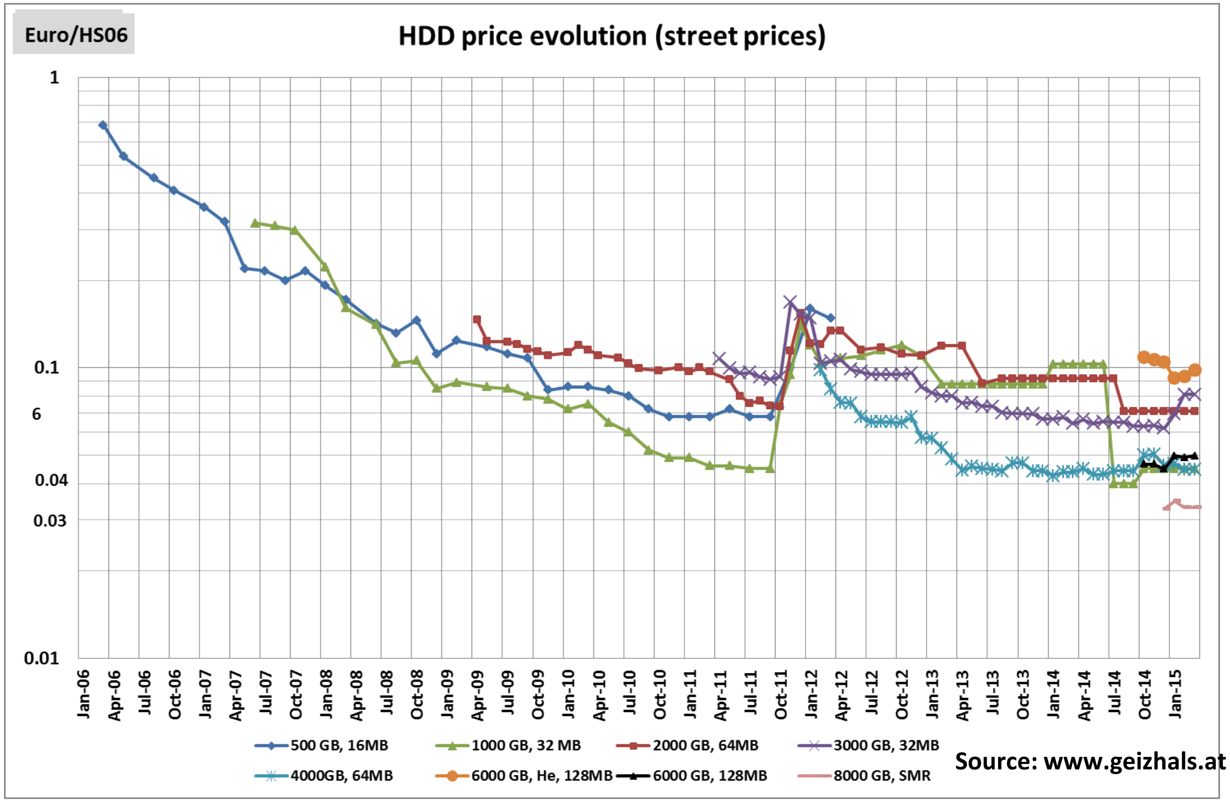

Source: www.geizhals.at

‘Thailand’ crisis end of 2011
Price recovery period
was very long (artificial !?)

Raw disk price evolution of server disks
(CERN purchase)

Consumer disk price evolution

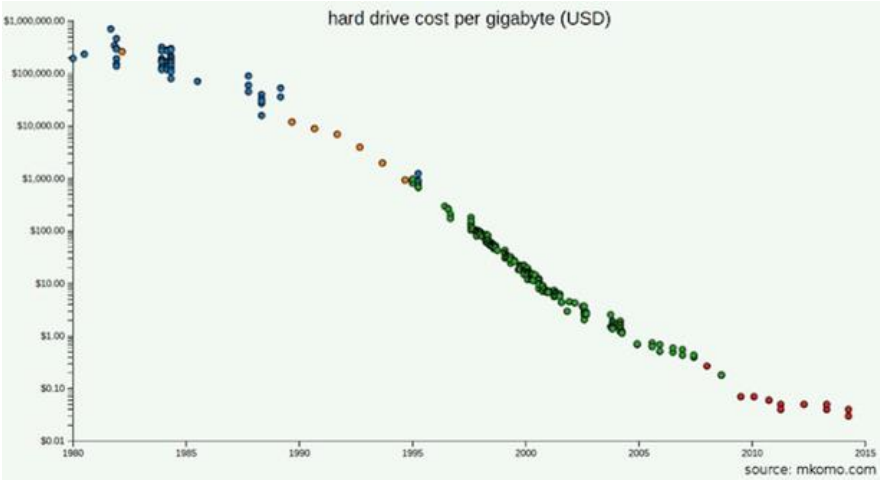

Decreasing price/space
improvement rate

Storage Components: Hard-Disk-Drives III

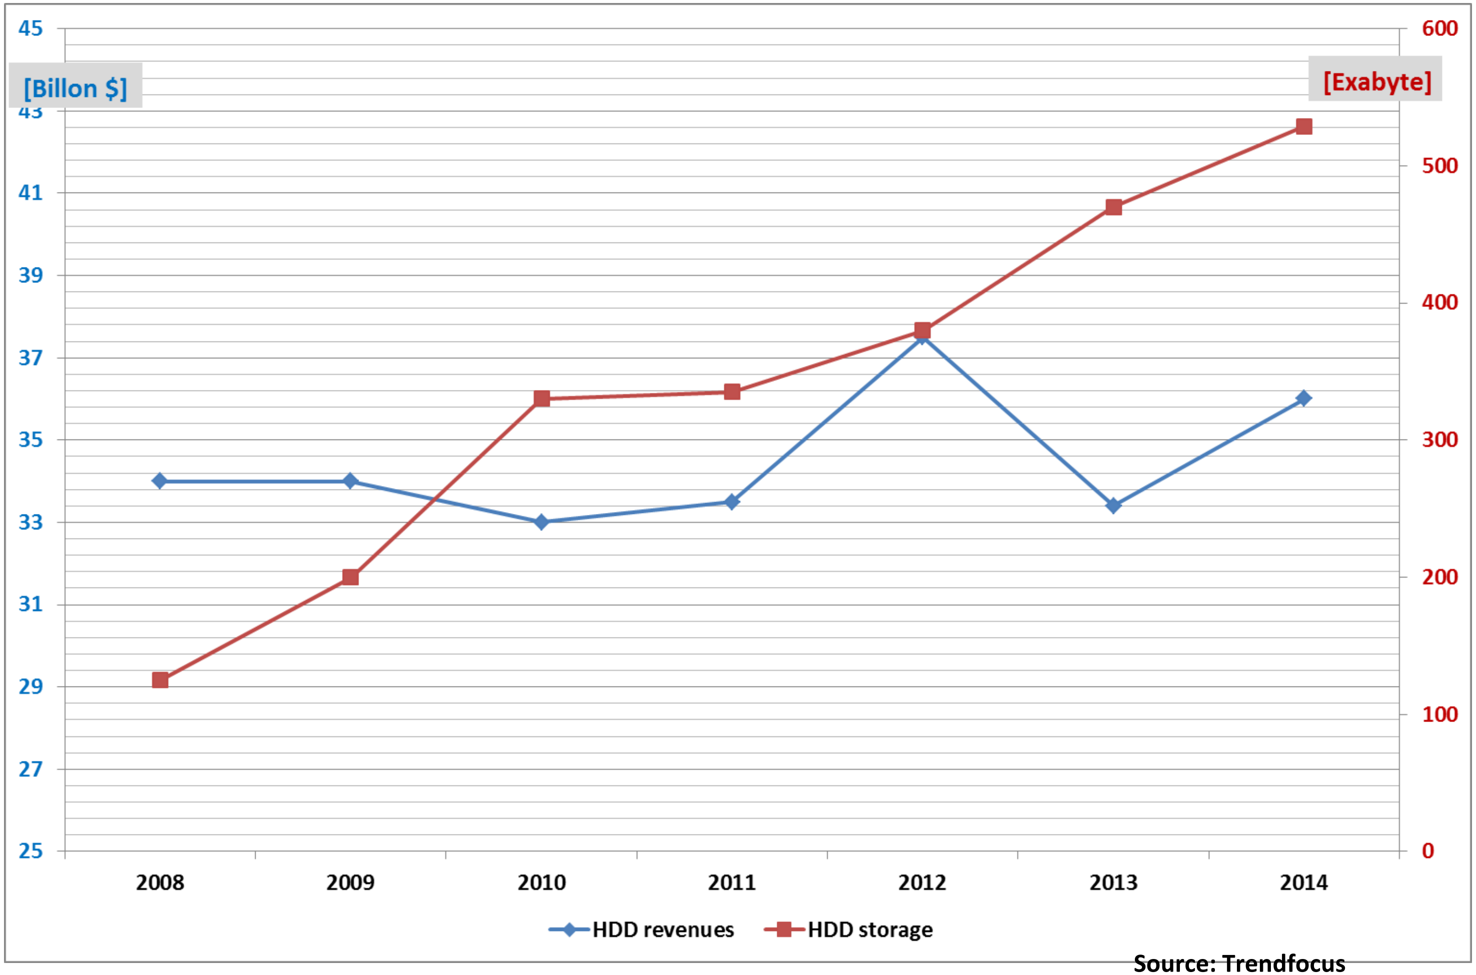

Source: Trendfocus

564 million HDDs sold in 2014
The market for server level disks is only 13% of the total

Revenue increase in 2012 due to the 'Thailand' crisis in 2011
Steady, but slower yearly increase in total space shipped

Storage Components: Hard-Disk-Drives IV

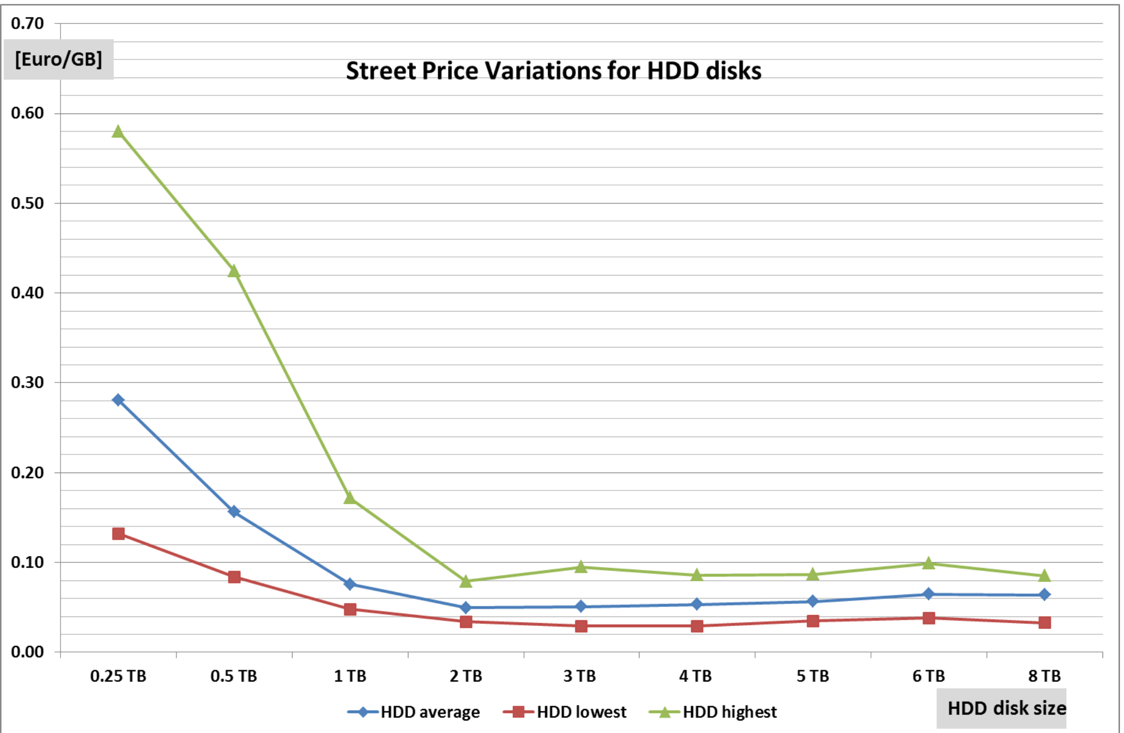

HDD cost variation of a factor 3 for the same disk size (performance, reliability)

Cost/GB difference between HDD and SSD = factor 3 to 25
Disk size dependent

Source: www.geizhals.at

Storage Server Cost Evolution

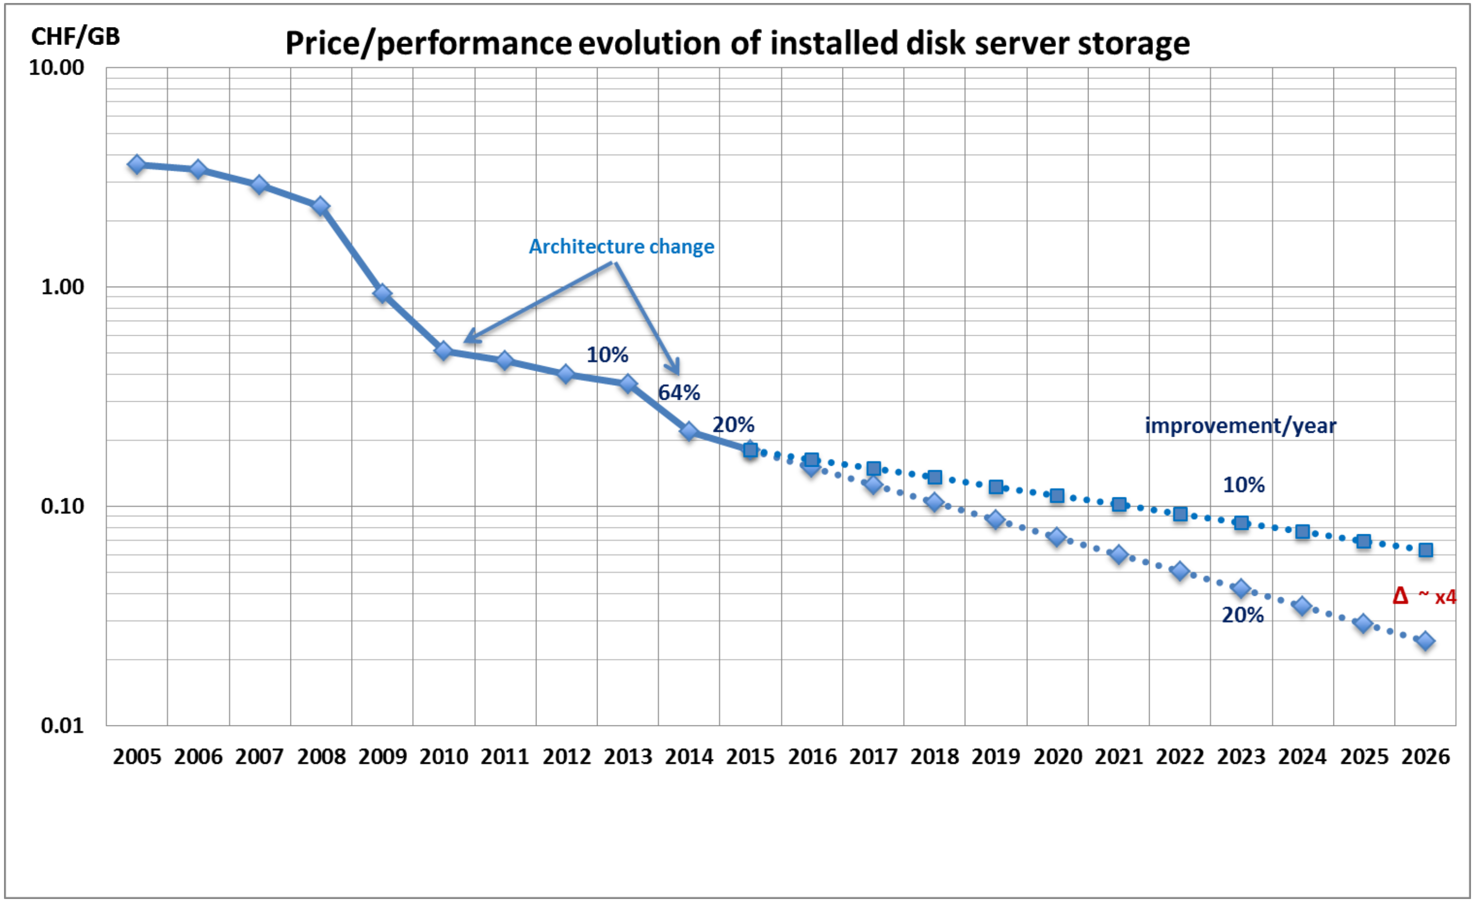

2015 to 2026
Improvement = factor 9
At 20% growth rates

CERN purchases of disk servers: costs defined by component costs, economy of scale (homogeneity !) and the Architecture (also software dependent)

Architecture changes during the last years:

- **RAID5 → RAID1**
- **Integrated disk server → CPU frontend with SAS attached JBOD array**
- **RAID1 → software data replication**
- **One array per server → two arrays per server**

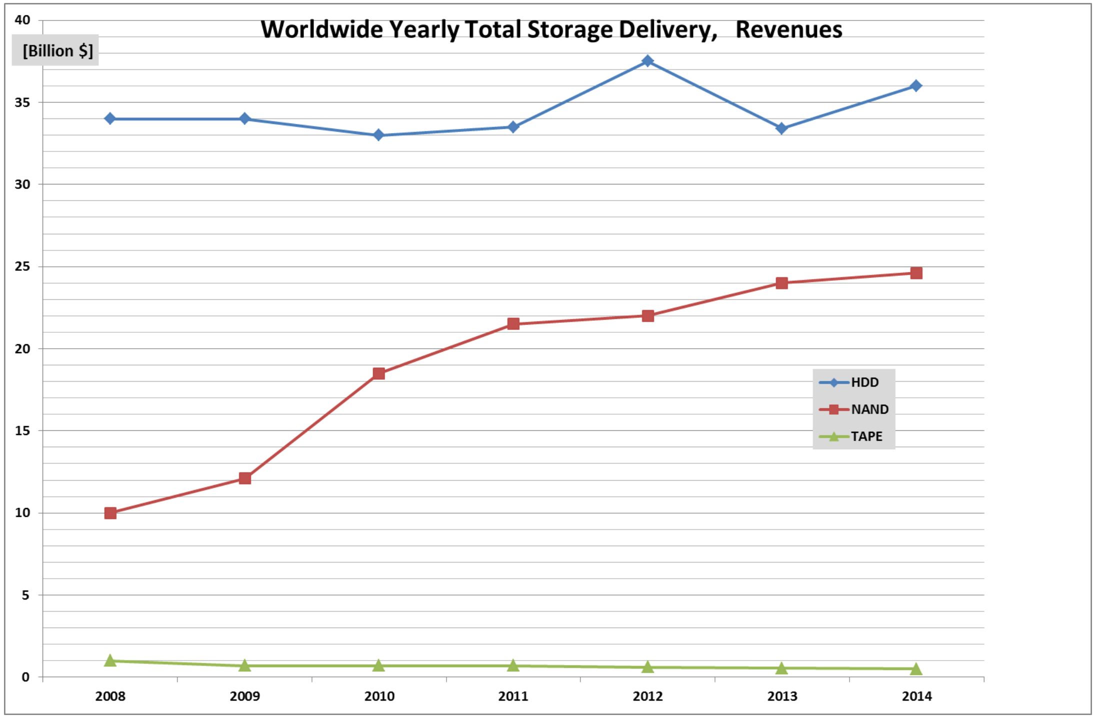

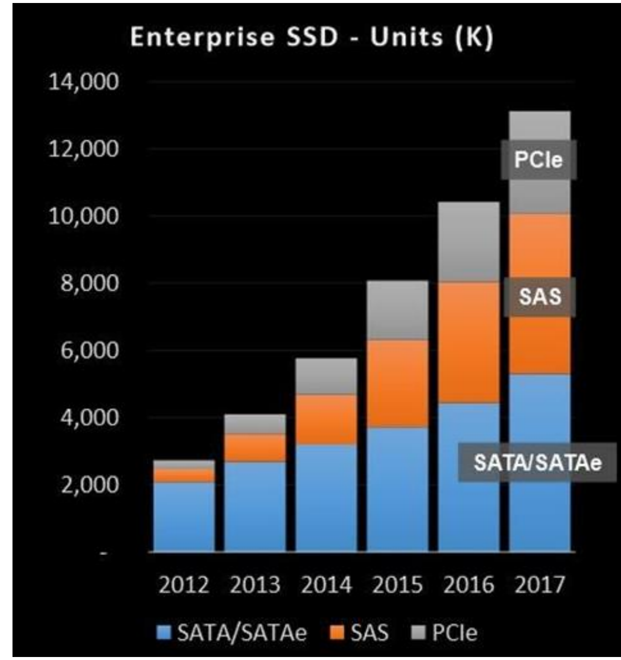

75 million enterprise HDDs in 2014

15% of the NAND storage is used for SSDs

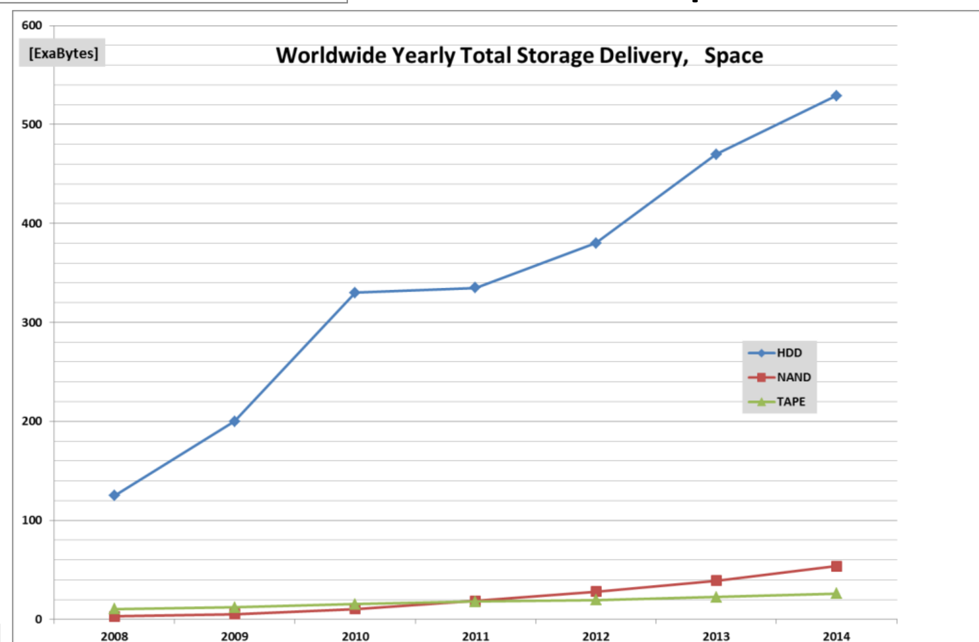

To yearly deliver the 530 Exabytes of HDD storage with SSDs would require an investment of ~0.5 T$ in NAND fabrication

The replacement of HDDs by SSDs will take quite some time

Component Capacity vs Time (best case)

IBM

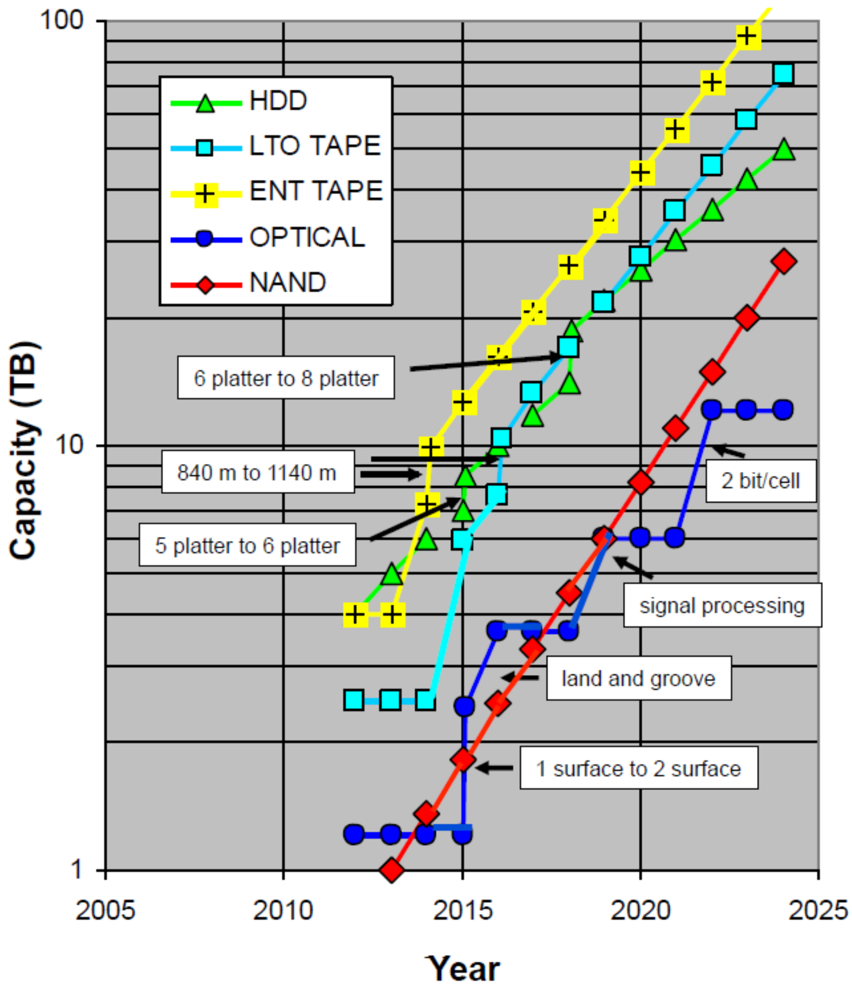

- Component capacity scaled from 2013 data using best case areal density growths from 6 year history
- 5X spread in capacity
- HDD – convergence with TAPE
- OPTICAL – significant capacity lag relative to TAPE and HDD

2025

200 TB enterprise tape
100 TB LTO tape
60 TB HDD
25 TB SSD

Not a direct relationship to costs

14 R. Fontana, G. Decad -- LOC -- September 22, 2014

Back-of-an-Envelope Calculations, component savings

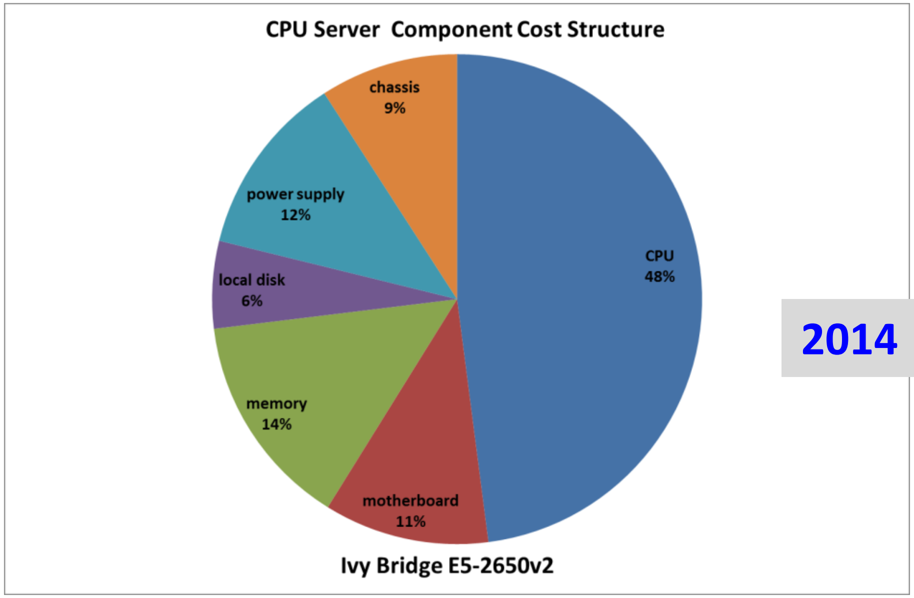

2014 → 2015

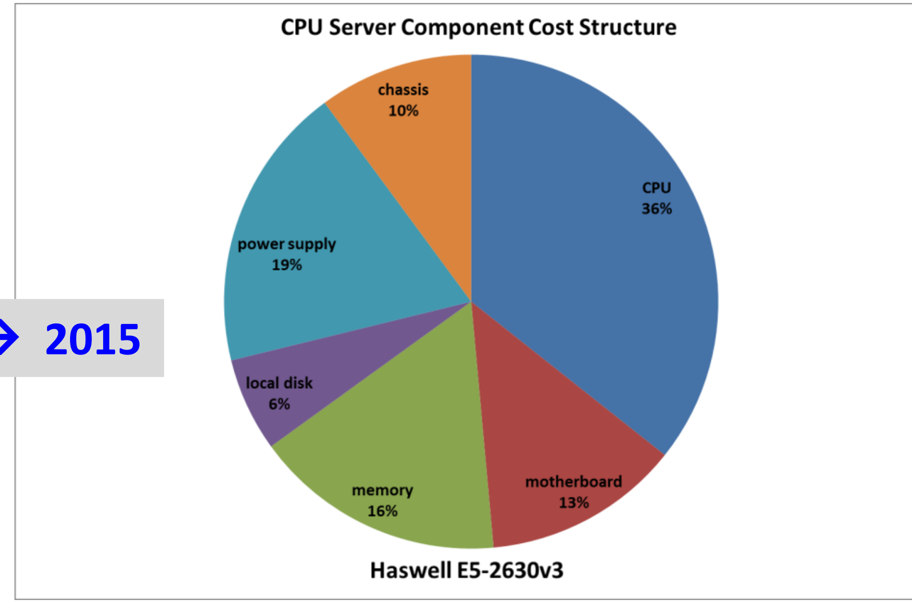

- Dominant part is the CPU, still getting best price/performance processors including infrastructure costs
- Sweet spot is still dual processors with medium frequencies ~(~2.5 GHz)
- The usual question about the relation of HepSpec and real HEP code.....
- Reducing memory by a factor 2 could create costs savings of 7-8%
- SMT increases performance by 20-25% while increasing memory costs by 7-8%, still a gain

  → local disk performance issues cost increase with SSDs
- Lower 'quality' of memory, ECC?, MHz ? → HepSpec is sensitive to memory features at the 10% level , HEP code ?
- Quad server packaging better than Blade server (also operational issues)
- Open Compute Project architecture (racks, power, server); pilot on the way; savings seem to be small
- Desktop, processor+GPU, lower price/performance but single proc, no ECC, operational aspects --> gain 30% ?
- Maybe new microservers later --> gain 30%?

Not much to gain here, 10% level

Back-of-an-Envelope Calculations, power savings

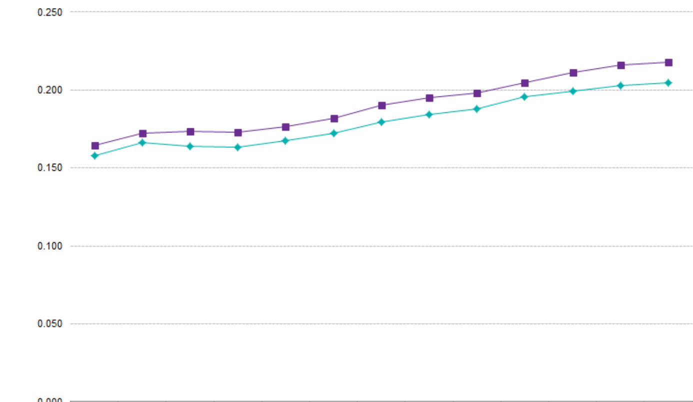

	2008s1	2008s2	2009s1	2009s2	2010s1	2010s2	2011s1	2011s2	2012s1	2012s2	2013s1	2013s2	2014s1
EU-28*	0.158	0.166	0.164	0.164	0.167	0.173	0.180	0.184	0.188	0.195	0.199	0.203	0.205
EA*	0.165	0.172	0.174	0.173	0.177	0.182	0.190	0.195	0.198	0.205	0.211	0.216	0.218

Average electricity price development in Europe, 2008-2014, Euro/kWh
Increase is ~4.5% per year

Relative energy costs of a CPU server:
Dual processor, 64 GB memory, 2 local disks → 3500,- Euro
4 years lifetime
300 W under full load, 80% efficiency, PUE of 1.7,

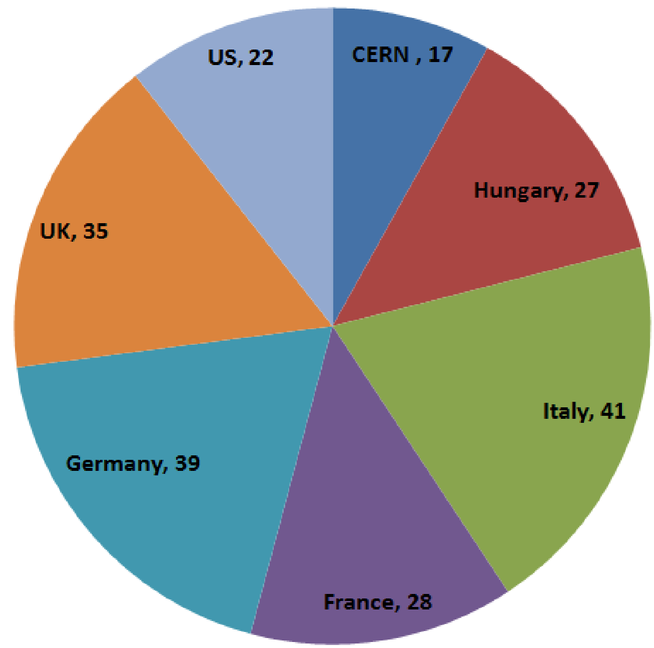

$$\frac{\text{Energy costs}}{\text{(Purchase costs + Energy costs)}}$$

Electricity cost varies by more than a factor 2 within Europe.
US costs are up to a factor 3 cheaper

→ Cutting the energy consumption by a factor 2 saves between 10 and 20% of the total cost

e.g. the cost for energy of a CPU server is 39% of the total costs in Germany

Back-of-an-Envelope Calculations, processor architecture savings

Cost and performance of various processor and accelerators

	Gflops SP	Gflops DP	cost [Euro]	power [W]		Gflops DP/ Euro	Gflops DP/ Watt	
Intel E5-2630v3 8x2.4 GHz	600	300	720	85		0.42	3.53	← Reference
Intel E5-2650v3 10x2.3 GHz	740	370	1250	105		0.30	3.52	
Intel E5-2690v3 12x2.6 GHz	1000	500	2150	135		0.23	3.70	
Xeon Phi, knights corner, 16GB	2416	1208	3500	270		0.35	4.47	
Xeon Phi, knights landing, 16GB	7000	3000	3500	300		0.86	10.00	← Price unknown, assumption
Nvidia GeForce Titan X	7000	200	1000	250		0.20	0.80	
Nvidia Tesla K40	4290	1430	5500	235		0.26	6.09	
Nvidia Tesla K80	8740	2910	7000	300		0.42	9.70	
Radeon firepro S9150	5070	2530	3500	235		0.72	10.77	
Altera Arria® 10 FPGAs 16 GB		1500	3000	50		0.50	30.00	

Assuming the code can use 100% of the Instructions per Cycle (IPC)

- **Price/performance gain of maybe a factor 2 for the new Xeon Phi**
- **Power/performance gain of a factor 9 for the Altera FPGA == costs saving of up to 35% (see previous slide)**
- **Savings are reduced due to fact that the processors/accelerators are only 30-40% of the total system (cost and power)**

Microsoft and Baidu bought Altera FPGA PCIe boards for their search servers, Microsoft also uses Xeon Phi.
HPC GPUs, Xeon Phi, HPC FPGAs are niche products with sales of ~10000 units per year.

Detailed investigations of the new ARM (HP Moonshot) and power8 servers have shown that they are not yet a real competition http://lvalsan.web.cern.ch/lvalsan/processor_benchmarking/presentation/#/future_work
→ At least a factor 5 worse in terms of price/performance and a factor 2 worse in power/performance

A Haswell processor can do up to 32 instruction per cycle, HEP code uses about 1

Back-of-an-Envelope Calculations, storage component savings

CERN disk server: CPU server with SAS attached JBOD array

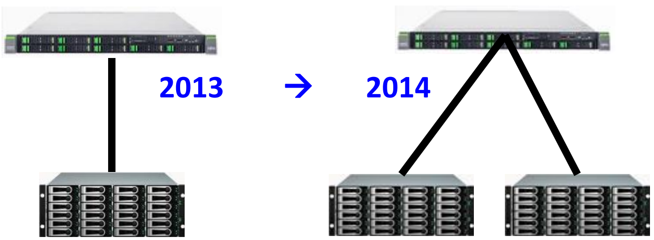

200 TB RAW capacity
100 TB usable → 0.2 Euro/GB

Infrastructure and architecture 'overhead' =~ factor 7

Cheapest server disk today is the 8 TB Seagate SMR (0.03 Euro/GB)

1440 TB RAW capacity
1152 TB usable → 0.06 Euro/GB

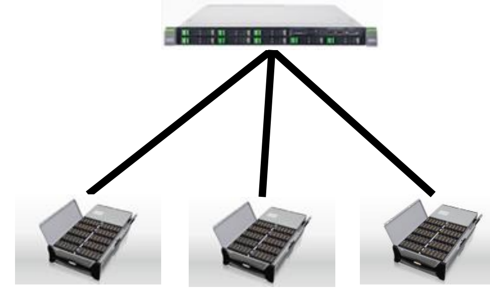

Example: 'improve' the storage costs by a factor 3:

4 TB server disk ~0.05 Euro/GB → 8 TB SMR ~0.03 Euro/GB (low-end desktop 6 TB)
Dual 24-bay disk tray → three 60-bay disk trays per frontend
RAID0 / data replica → Erasure code, data increase by 1.25 instead of 2

This improves the space costs but reduces considerable the IO capabilities. But how much IO do we actually need ?
(Application, data management, data distribution dependent)
Much more tuning between application and hardware needed…..

Redefine our notion of storage space
→ Storage space plus performance

Split
MC+processing facilities -- analysis facilities

different IO architecture based on Seagate Kinetic object drive model or the HGST Open Ethernet drive

FLAPE
Flash+Tape

Summary

Semiconductor Component and end-user markets are stabilizing.
Saturation effects seen nearly everywhere, moving to 'replacement' markets

Very few companies dominating the market: technology evolution , not revolution

Moore's Law validity being debated. 3D technology helps.
Expect still continuous price/performance improvements, but lower levels

Server market is small compared to the consumer market, stable and highly profitable Market --> high prices. Microservers show in principle potential, but currently overrated

Way to improve price/performance beyond the technology --> architecture

Should not talk about disk, SSD or tape but rather storage units (space+performance)

There will be processing and storage technologies in 2025 and most likely not too different from today, but estimating the cost is pretty difficult.
So.. You will get what you get (equal or rather lower budget than today)……