Digital Storage Technology Roadmaps & Implications for Cloud Computing *iNEMI*

2017 Roadmap Mass Data Storage TWG

Roger Hoyt, Consultant, Chair, r.hoyt@ieee.org

Tom Coughlin, Coughlin Associates, Co-Chair, tom@tomcoughlin.com

iNEMI®

Advancing manufacturing technology

Outline

- **iNEMI Overview & Roadmap Process**
- **Mass Data Storage**
  - Solid State
  - Hard Disk Drives
  - Tape
  - Optical
  - Cloud Storage
- **Summary**

iNEMI

Mission: Forecast and Accelerate improvements in the Electronics Manufacturing Industry for a Sustainable Future.

5 Key Deliverables:

- **Technology Roadmaps**
- Collaborative Deployment Projects
- Research Priorities Document
- Proactive Forums
- Position Papers

International Electronics Manufacturing Initiative (iNEMI) is an industry-led consortium of global manufacturers, suppliers, industry associations, government agencies and universities. A Non Profit Fully Funded by Member Dues; In Operation Since 1994.

Visit us at www.inemi.org

2017 Technology Working Groups

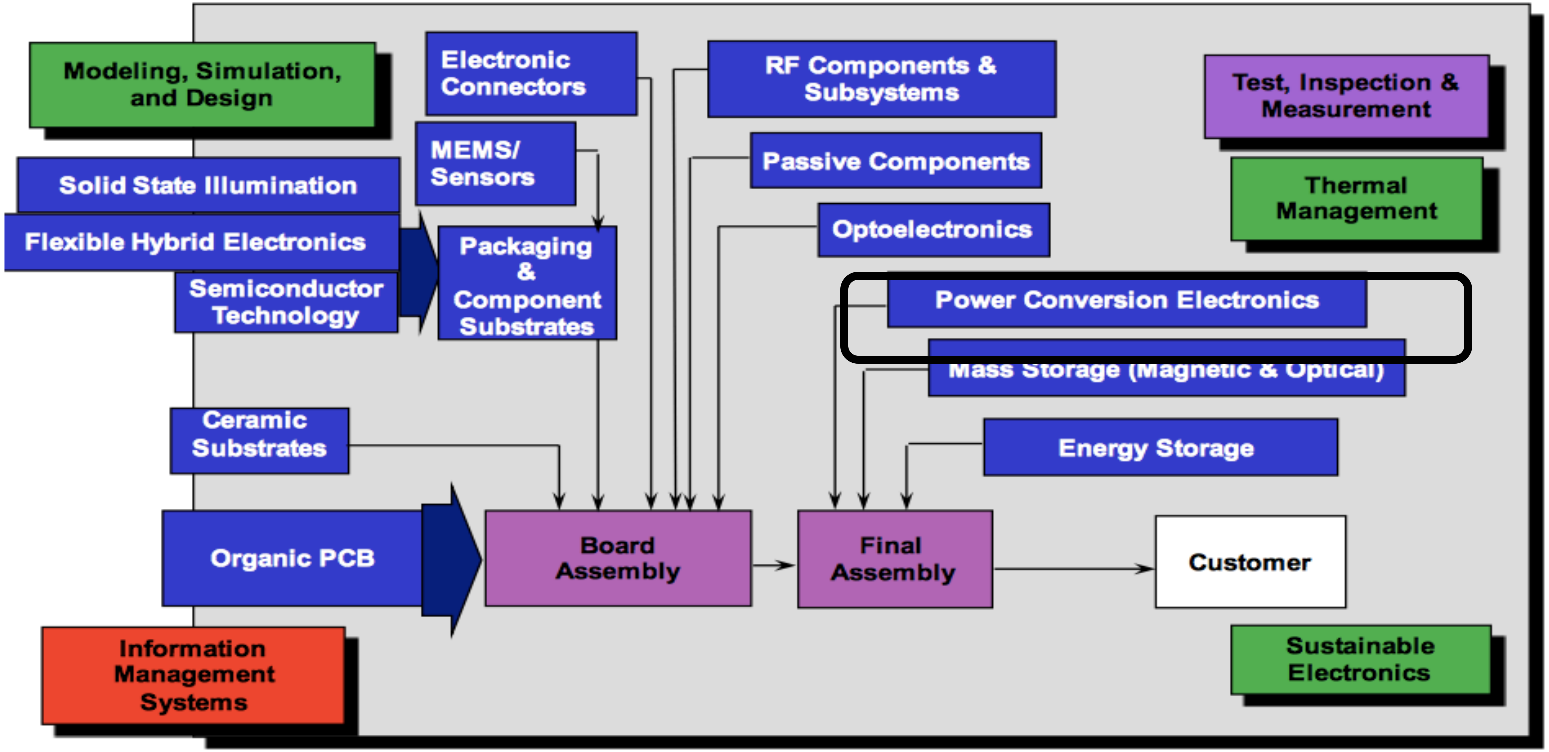

Red=Business Green=Engineering Purple=Manufacturing Blue=Component & Subsystem

Mass Data Storage Committee

Solid State

Flash/Phase Change/Ferroelectric/RRAM

Jim Handy	Objective Analysis Consultants

MRAM

Jim Deak	Multidimension Technology
Mark Johnson	Naval Research Laboratory

Hard Disk Drives

David B. Aune	Seagate/ Univ of Minnesota
Ron Dennison	Research/Development Consultants

Tape & Optical Drives and Media

Dick Zech	Advent Technologies
Tom Coughlin	Coughlin Associates
Barry Schechtman	INSIC

iNEMI

iNEMI Mass Data Storage Roadmap

A Spectrum of Technologies

Solid State/Flash

HDD

Tape

Optical

iNEMI

Touch rate versus response time indicating various types of uses

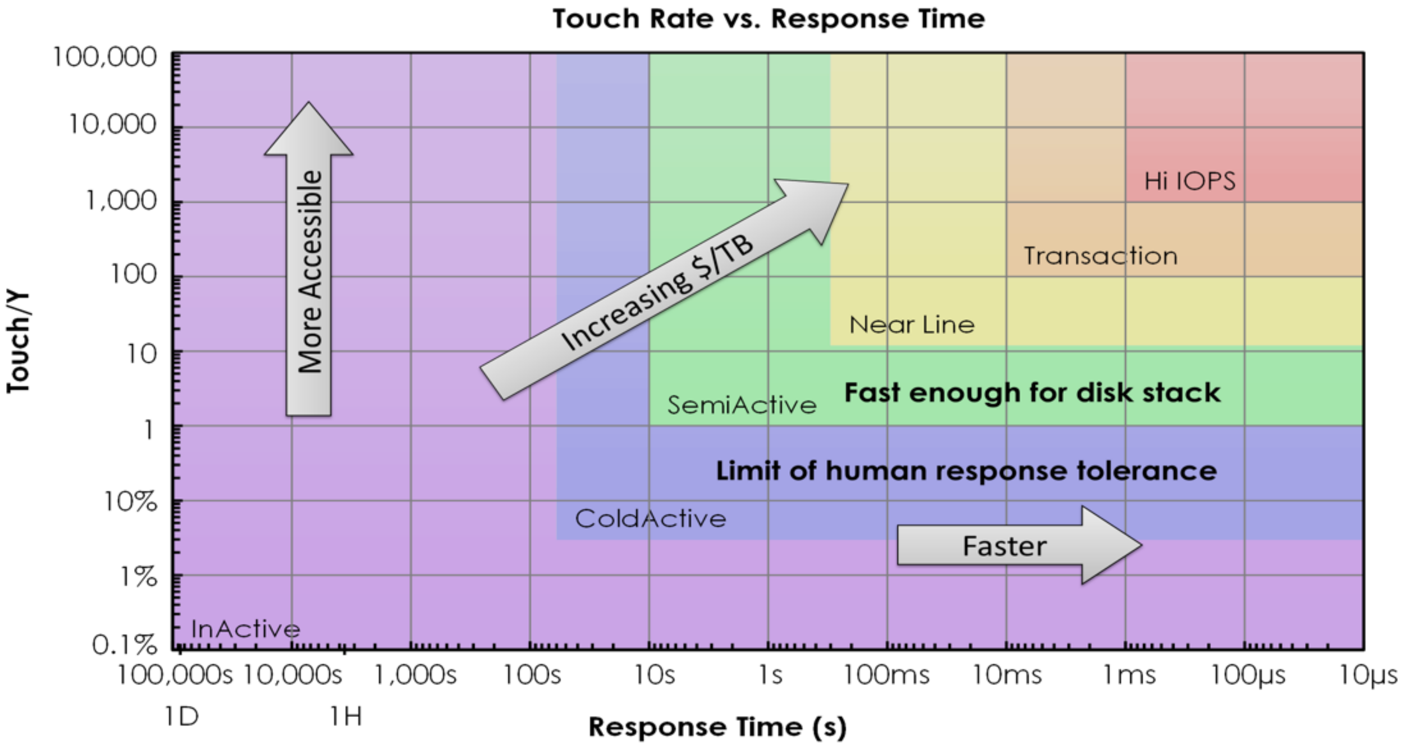

Digital storage technologies regions overlaid on the Touch Rate/Response Time chart

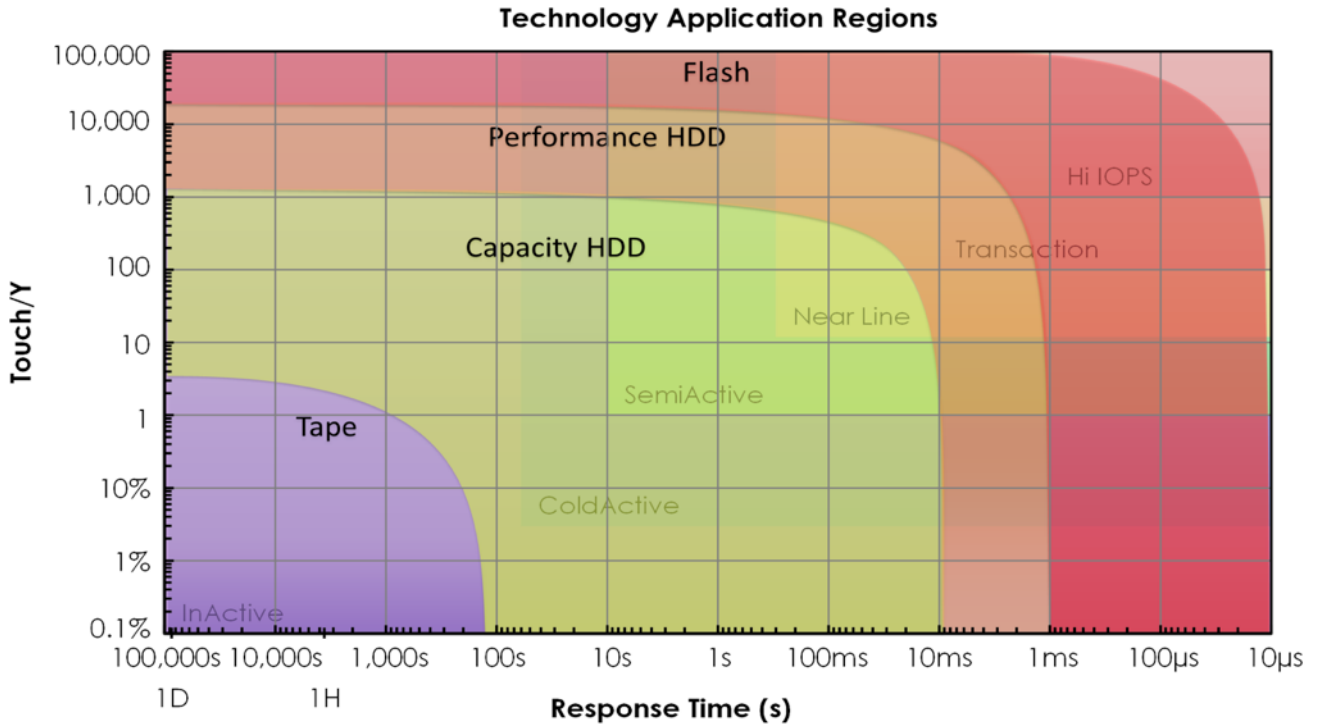

Mass Data Storage Chapter Status

Solid State Storage/Memory

iNEMI Solid State Memory Trends

- **Flash**
  - Scaling Limits lead to conversion from planar to 3D.
  - Market moving from displacement (i.e. photographic film) to new applications (SSDs in PC and servers)
- **MRAM**
  - Evolution of Next Generation to spin torque switching
  - Growth of Applications
  - New players, partnerships and Everspin IPO
- **Phase Change**
  - Newly-defined application creates 3D Xpoint
- **RRAM**
  - Some positioning as competing against 3D XPoint
  - Otherwise viewed as an eventual NAND replacement

Source: Electricity Storage Association

iNEMI Mass Data Storage Roadmap

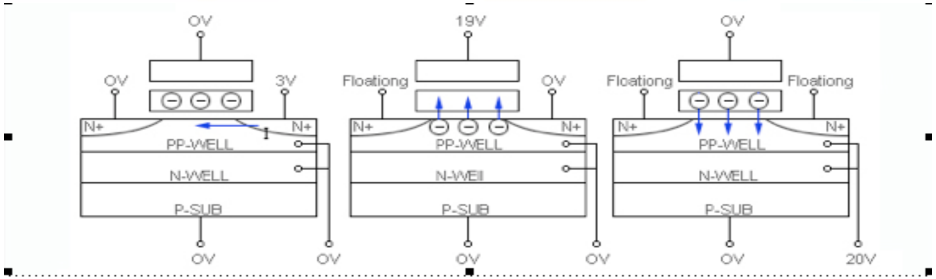

Figure 4. Cross-Section of planar floating-gate flash memory cell

Left: Reading, Middle: Programming, Right: Erasing

(Source: Samsung Semiconductor Company)

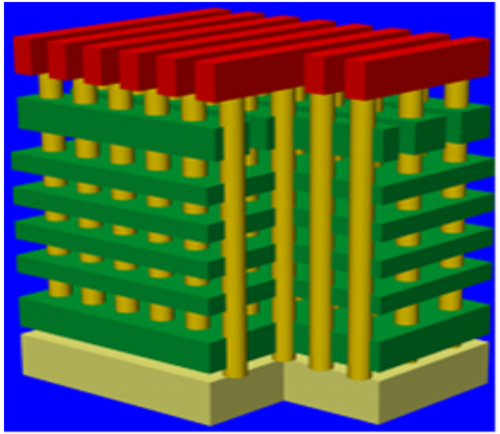

Table 1. ITRS NAND Flash Chip Roadmap

	2013	2015	2017	2019	2021	2023	2025	2028
Density (SLC/MLC)	64G/128G	128G/256G	256G/512G	512G/1T	512G/1T	1T/2T	2T/4T	4T/8T
Planar Process	18nm	15nm	13nm	11nm	9nm	8nm	8nm	8nm
3D Layers	16-32	16-32	16-32	32-64	48-96	64-128	96-192	192-384
3D Process	64nm	54nm	45nm	30nm	28nm	27nm	25nm	22nm

Source: ITRS, 2013

iNEMI

Flash memory roadmap

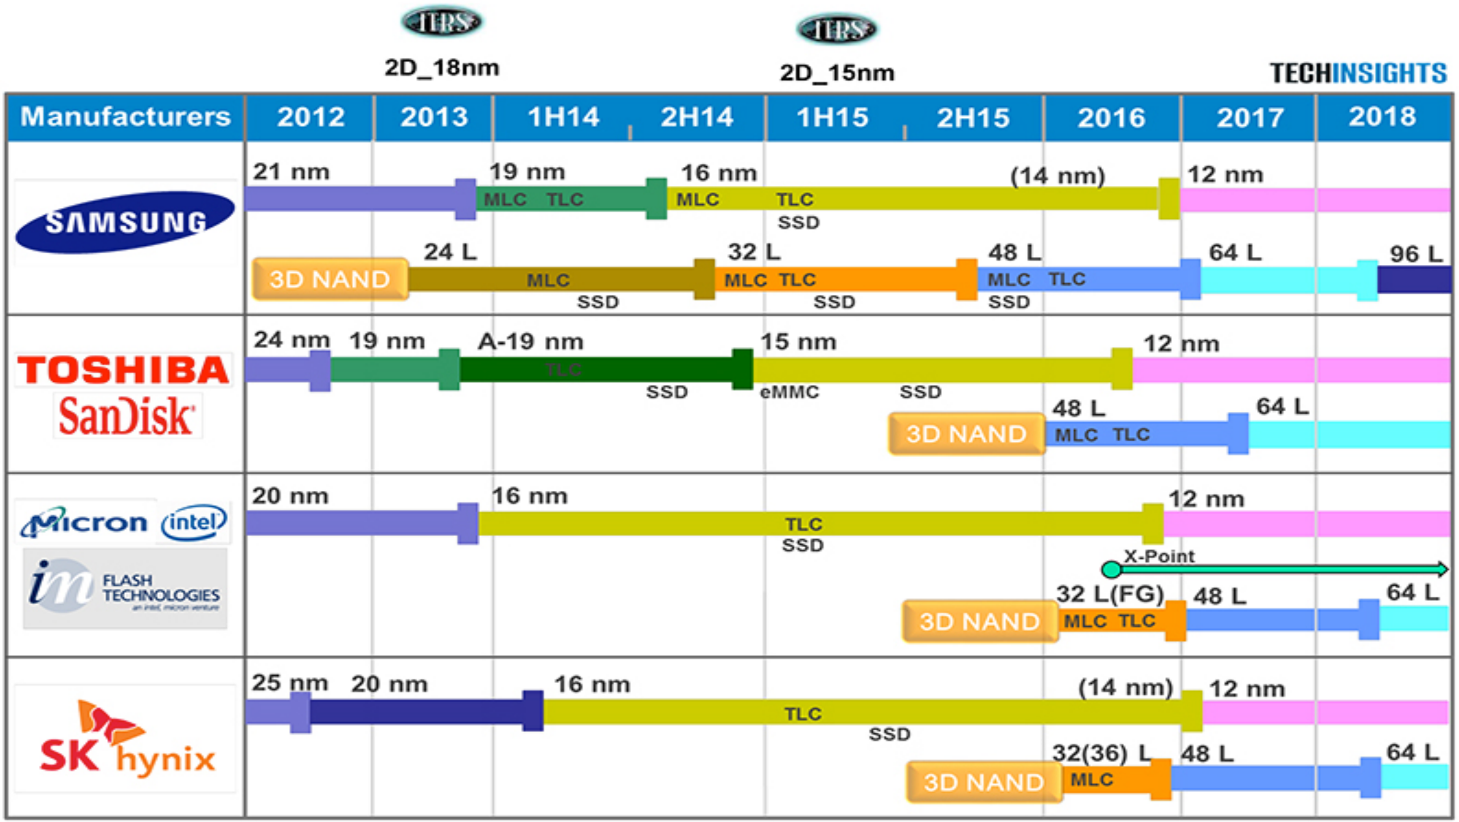

iNEMI

iNEMI Mass Data Storage Roadmap

Table 6. Attributes of Different Memory Technologies

	SRAM	DRAM	Flash	FRAM	MRAM	ReRAM
Read Speed	Fast	Medium	Medium	Fast	Fast	Medium
Write Speed	Fast	Medium	Slow	Fast	Medium	Medium
Array Efficiency	High	High	Medium	Medium	High	High
Scalability	Good	Limited	Limited	Limited	Medium	Good
Cell Density	Low	High	High	Medium	Medium	High
Volatile?	Yes	Yes	No	No	No	No
Endurance	Infinite	Infinite	Limited	Limited	Infinite	Limited
Current Consumption	Low/High	High	Low	Low	Low	Low
Low-Voltage	Yes	Limited	Limited	Limited	Yes	Yes
Process Complexity	Low	Medium	Medium	Medium	Complex	Medium

(Source: Objective Analysis)

3D Xpoint Memory

- This is a non-volatile memory that is faster and has higher endurance than flash memory, while slower than DRAM/SRAM
- 3D XPoint is a Phase Change Memory
- Intel and Micron are introducing products with this technology in Q2 2017

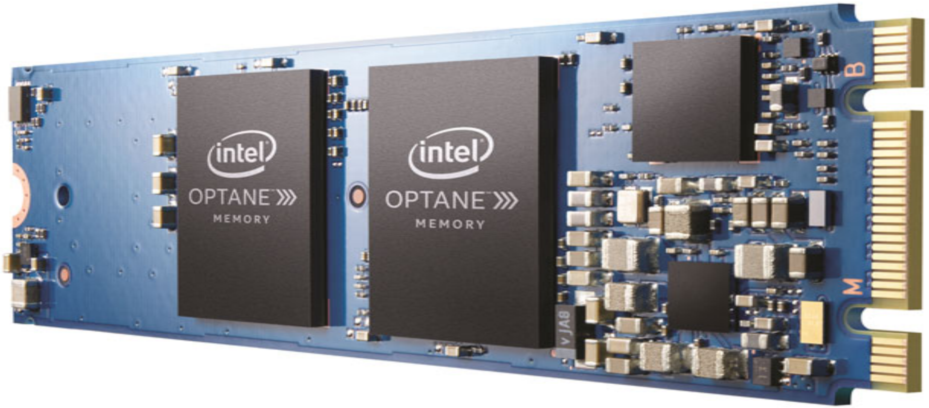

MRAM Uses

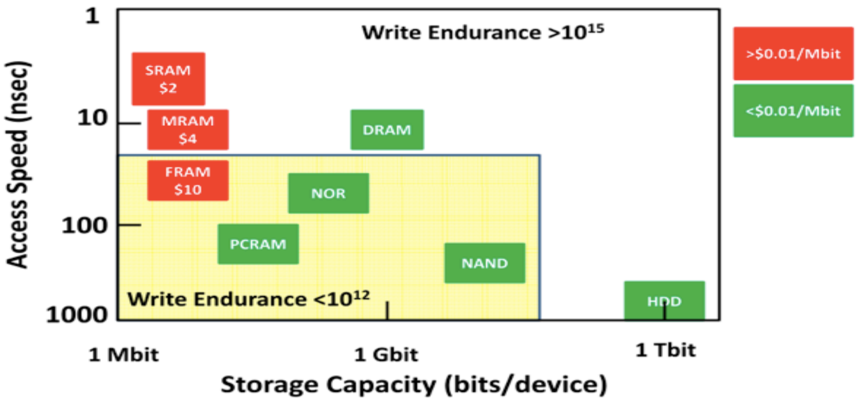

Figure 25– Cost breakdown of mass storage technologies.

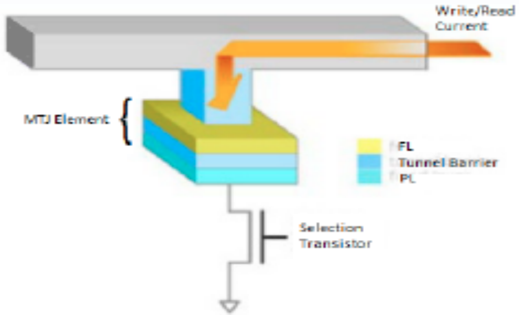

Figure 13 – General design and operation of a STT MRAM cell.

- MRAM --high write endurance and sub 10 ns access speed.
- Cost is not high compared to SRAM, and the access time is close enough to SRAM that it is already an interesting replacement for some SRAM.
- Given lower prices with higher volume, it could possibly replace DRAM in many applications.
- Embedded as well as stand alone applications

iNEMI

MRAM Roadmap

Metric	2013	2015	2017	2019	2023
density (Mbit)	16	64	256	1024	4064
technology	toggle	STT, in-plane	STT, in-plane	STT, perpendicular	STT, perpendicular
die size (cm^2)	0.58	0.58	0.58	0.9	1
density (Mbit/cm^2)	28	110	441	1138	4064
array efficiency	0.75	0.75	0.75	0.75	0.75
array element size (um^2)	1.577	0.394	0.099	0.059	0.018
cell efficiency (f^2)	38	22	22	18	16
required litho resolution (nm)	102	67	33	29	17
wafer size (mm)	300	300	300	300	450
dice/wafer	2316	2316	2316	1492	3022
wafer cost ($)	3000	3000	3000	3000	4000
est. production cost/die ($)	1.30	1.30	1.30	2.01	1.32
est. production cost/Gbit ($)	46.96	11.74	2.94	1.77	0.33
Performance	**2013**	**2015**	**2017**	**2019**	**2023**
Write / read time (ns)	25	20	10	10	5
Data rate (write or read limit) (MHz)	150	400	500	500	500
Energy to write 1 bit (picojoule)	200	2	2	2	1
Energy to read 1 bit (picojoule)	100	2	1	1	1

Mass Data Storage Chapter Status

Hard Disk Drives

iNEMI Hard Disk Drive Trends

- Industry consolidation leaves 3 players in industry
- Areal density growth has slowed from early in decade—target is about 15% CAGR going forward
- Shingled Magnetic Recording (SMR) provides growth path but not for frequently overwritten data—need energy assisted recording
- Heat Assisted Magnetic Recording (HAMR) keeps getting pushed out—now 2018 or 2019 introduction into products
- New native Ethernet interface drives from Seagate (Kinetic) and WD.
- Interesting applications driven HDDs with additional intelligence from HDD companies (e.g. WD Labs)

HDD Companies (a $30B annual industry)

- **HDD Companies**
  - **A $30B annual industry**

Seagate We turn on ideas	TOSHIBA Leading Innovation >>>
WD Western Digital®	HGST a Western Digital company

HDD Shipment Projections by Application
(Coughlin Associates, 2016)

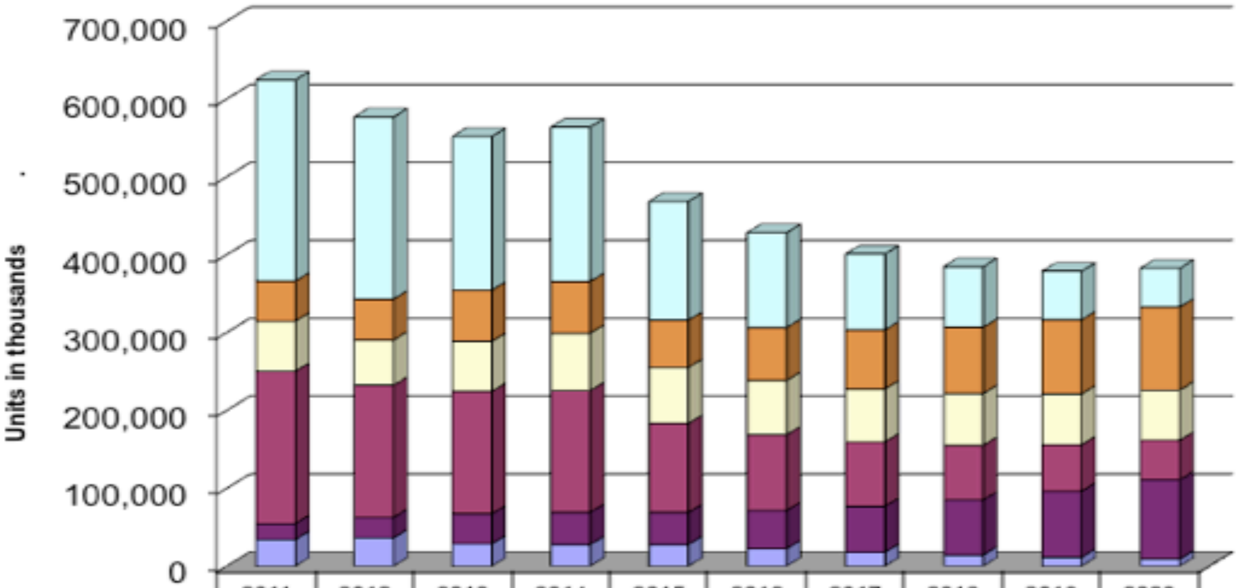

	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
Mobile	258,300	233,900	197,300	198,400	152,400	121,920	97,536	78,029	62,423	49,938
Branded	51,600	52,600	65,900	66,400	60,600	67,872	76,017	85,139	95,355	106,798
CE	64,212	58,143	64,282	73,777	72,041	69,812	67,942	66,403	65,174	64,239
Desktop	195,700	169,800	157,200	156,300	114,500	97,325	82,726	70,317	59,770	50,804
Enterprise ATA	20,400	26,100	38,500	41,300	40,900	49,080	58,896	70,675	84,810	101,772
Enterprise	34,300	36,500	28,800	27,800	28,000	22,400	17,920	14,336	11,469	9,175

iNEMI

ASTC HDD ROADMAP

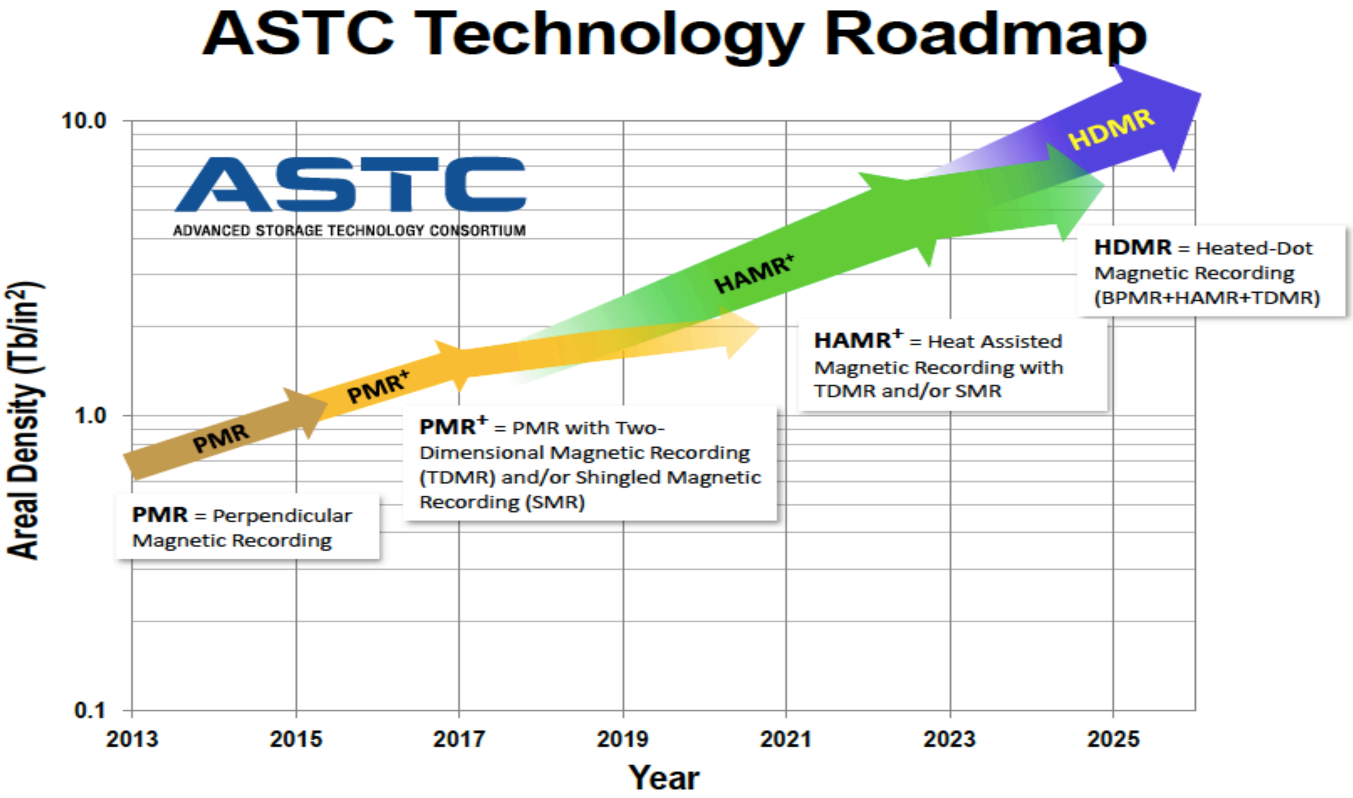

Cost Comparison HDD and SSD

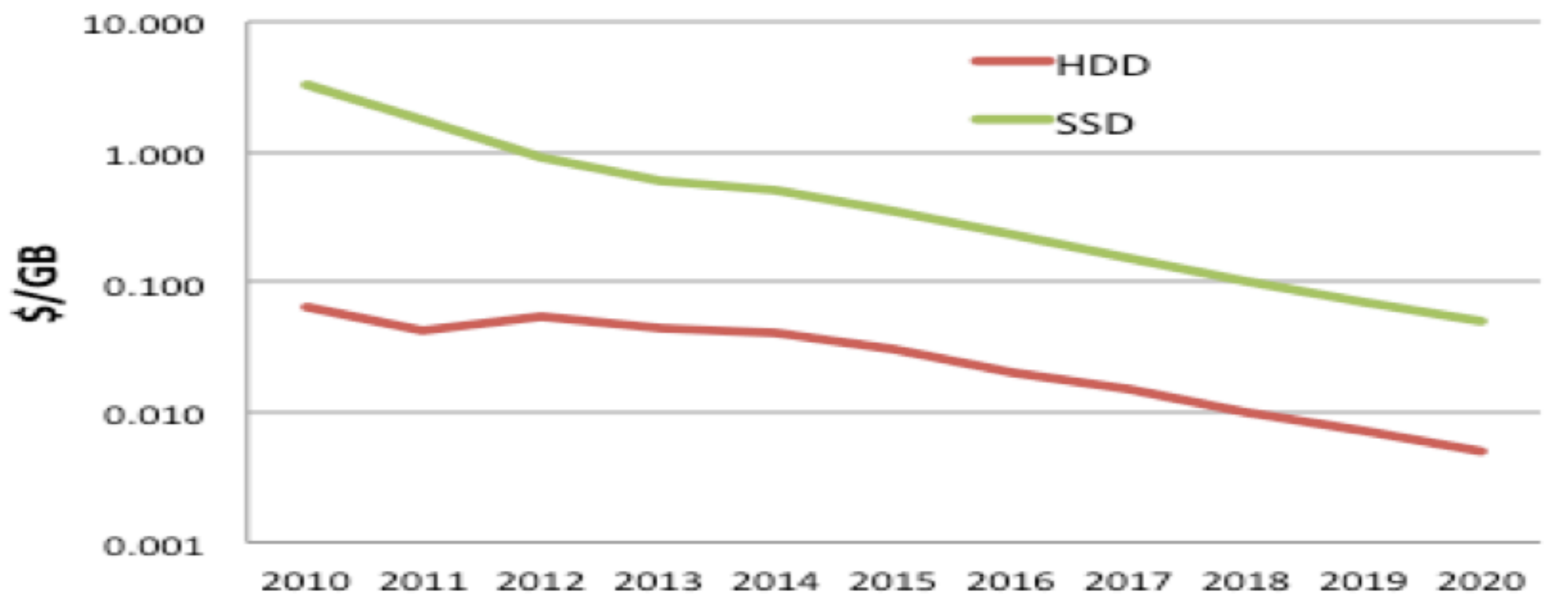

Figure 31 Raw Storage Average Retail Price vs. Time

(Source: Coughlin Associates, 2016)

iNEMI HDD Roadmap (1)

	Unit	2013	2015	2017	2019	2023	2027
Industry Metrics							
Form Factor (dominant form factor is bold)	inches	3.5, **2.5**, 1.8	3.5, **2.5**, 1.8	3.5, **2.5**,	3.5, **2.5**,	3.5, **2.5**	**2.5**, 1.8
Capacity	GB	120-6,000	250-10,000	300-14,000	500-20,000	700-34,000	1,500-60,000
Market Size	units (M)	552	468	401	379	420	500
Cost/TB (avg.)	$/TB	<50	<30	≤10	<5	<1	<0.5
Design/Performance							
Areal Density	Gb/in^2	>700	>900	>1,000	>1,600	>4,800	>10,000
Rotational Latency	ms.	2-12	2-12	2-12	2-12	3-12	3-12
Seek Time*	ms.	3-5	3-5	3-5	2-5	1.5-5	1-4
RPM		4.2-15K	4.2-15K	4.2-15K	4.2-10K	4.2-10K	4.2-10K
Data rate	Mb/sec	10-2,500	10-2,700	12-2,800	14-3,200	20-6,400	40-10,000
Power *Key Component Requirements*	watts	1-10	1-10	1-10	0.7-9	0.5-8	0.3-6
Read Head	type	TMR	TMR	TMR	TMR/CPP	CPP-CCP	CPP-CCP
Slider	type & size (% of micro, 3.86 mm^3)	5%	5%	5%	5%	<5%	<5%
Fly ht [remove?]	nm	<4	<4	<4	<4	<4	<4
Disk	type	AlMg, Glass	AlMg, Glass,	AlMg, Glass, High Temp Glass	AlMg, Glass, New Substrate, High Temp Glass	AlMg, Glass, New Substrate, High Temp Glass	AlMg, Glass, New Substrate, High Temp Glass
Disk Static Coercivity	Oe	4,500-5500	5,000-6,000	5,000-6,500	5,000-20,000	6,000-40,000	20,000-50,000

iNEMI HDD Roadmap (2)

	Unit	2013	2015	2017	2019	2023	2027
Magnetic Recording Technology		Perpendicular, SMR	Perpendicular, SMR	Perpendicular, SMR, TDMR	Perpendicular, HAMR, SMR, TDMR	Perpendicular, HAMR, Patterned Media, SMR, TDMR	Perpendicular, HAMR, Patterned Media, SMR, TDMR
Electronics/Channel	type	LDPC Iterative GPR (Turbo), Pattern Dependent Noise Predictive GPR	LDPC Iterative GPR (Turbo)	LDPC Iterative GPR (Turbo), TDMR	LDPC Iterative GPR (Turbo), TDMR	Soft ECC, TDMR	Soft ECC, TDMR
Channel Bandwidth	MHz	80-2,000	80-2,000	80-2,000	80-2,000	>2,000	>4,000
SNR	dB	<20	<20	<20	<20	<17	<15
Actuator	type	Conventional /Micro, DSA	Conventional /Micro, DSA	Conventional /Micro, + DSA	Conventional/ Micro, + DSA	Conventional/ Micro, + DSA	Conventional/ Micro, + DSA
Spindle	type	Fluid	Fluid	Fluid	Fluid	Fluid	Fluid

*Seek time is one third full stroke seek time and does not include microactuator local track

Sealed Helium Drives

Figure 43 HGST a Western Digital Company and Seagate - 10TB Sealed Helium Drive

- Key benefits: 35%- 50% higher capacity, 20% lower idle power, 45% better watts/TB, 30% quieter operation, 4° C cooler operation, 50g lighter weight
- Shipment of production models of these drives commenced in 2013. Seagate and WD are now shipping 12 TB models with 14-20 TB announced.
- After about 2 years field use He-drives may have improved reliability (2M hours)

iNEMI

Mass Data Storage Chapter Status

Tape

Magnetic Tape Trends

- LTO is dominant
- Oracle and IBM make enterprise tape
- With change to Barium Ferrite tape in LTO reliability increased, perhaps minimizing the advantage for enterprise tape
- Current high is 10 TB half-inch tape cartridges but LTO 8 with 12 TB native should be announced by end of 2017, start of 2018

Tapes and Libraries

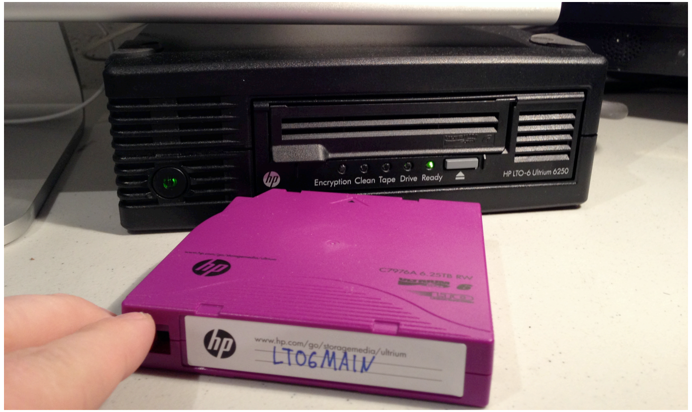

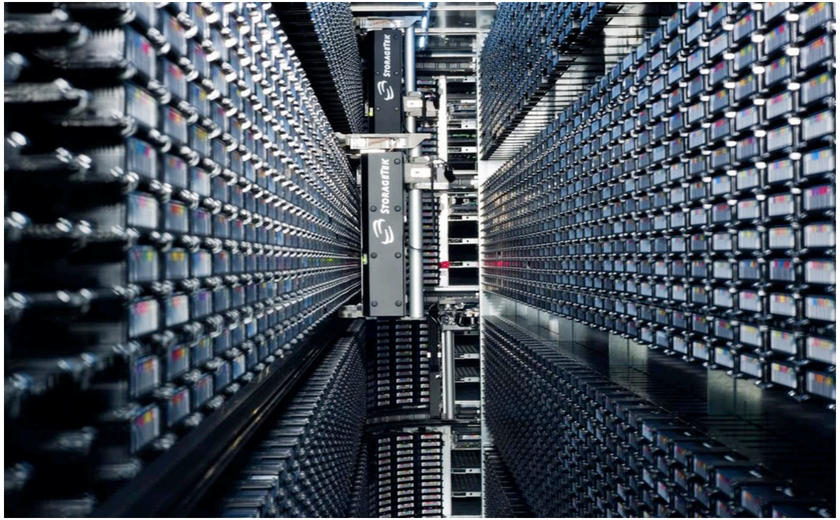

Tapes sometimes used for local copies but more often as part of a robotic library system

Tape Storage Technology Roadmap

	Unit	2015	2017	2019	2025	2029
Form Factor F/HH=Full/Half Height	inch	5.25 FH, 5.25 HH, 3.5	5.25 FH, 5.25 HH, 3.5	5.25 HH,,3.5	5.25 HH,,3.5	5.25 HH,,3.5
Volumetric Density	GB/in^3	400	700	1,300	7,200	21,000
Cartridge capacity (native)	TB	6-10 TB	10-24 TB	24-48 TB	100-300 TB	300-700 TB
Areal Density	Gb/in^2	3.1-8.0	6.5-16.0	16.0-32.0	50.0-100.0	140.0-350.0
Data Rate	MB/s/drive	250-360	400-600	600-800	1,600-3,000	4,000-6,000
Tape Speed (for data)	meters/sec	3-6	3-6	4-7	6-8	8-10
Head tracking precision required	+/- μm	0.2	<0.1	<0.1	<0.1	<0.1
Key Requirements						
Heads	type	MR/GMR	MR/GMR	GMR/TMR	TMR	TMR
Number of data channels	Number	16-32	16-36	32-64	32-64	64-128
Detection channel	type	E^2 PRML, LDPC	E^2 PRML, LDPC TURBO-CODE	E^2 PRML, LDPC TURBO-CODE	E^2 PRML, LDPC TURBO-CODE	E^2 PRML, LDPC TURBO-CODE
Magnetic film	type	multi-layer metal particle metal film, barium ferrite film	multi-layer metal particle metal film, barium ferrite film	multi-layer metal particle metal film, barium ferrite film	multi-layer metal particle metal film, barium ferrite film	multi-layer metal particle metal film, barium ferrite film
Tape/media thickness	μm (micron)	5.2	<5	<5	<4	<4
Media substrate material	type	PEN Aramid*/ adv. polymer	PEN Aramid*/ adv. polymer	PEN Aramid*/ adv. polymer	Aramid*/ adv. polymer	Aramid*/ adv. polymer

50 TB Native Capacity in Next Decade

LTO Roadmap

	LTO-3	LTO-4	LTO-5	LTO-6	LTO-7	LTO-8	LTO-9	LTO-10
Shipment Year	2005	2007	2010	2013	2015	TBD	TBD	TBD
Native Capacity	400GB	800GB	1.5TB	2.5TB	6.0 TB	Up to 12.8TB	Up to 25TB	Up to 50TB
Compressed Capacity	800GB	1.6TB	3.0TB	6.25TB	15 TB	Up to 32TB	Up to 62.5TB	Up to 125TB
Native Transfer Rate	80 MB/s	120 MB/s	140 MB/s	160 MB/s	300 MB/s	Up to 472 MB/s	Up to 708 MB/s	Up to 1100 MB/s
Compressed Transfer Rate	160 MB/s	240 MB/s	280 MB/s	400 MB/s	750 MB/s	Up to 1180 MB/s	Up to 1770 MB/s	Up to 2750 MB/s

2015 INSIC Tape Roadmap

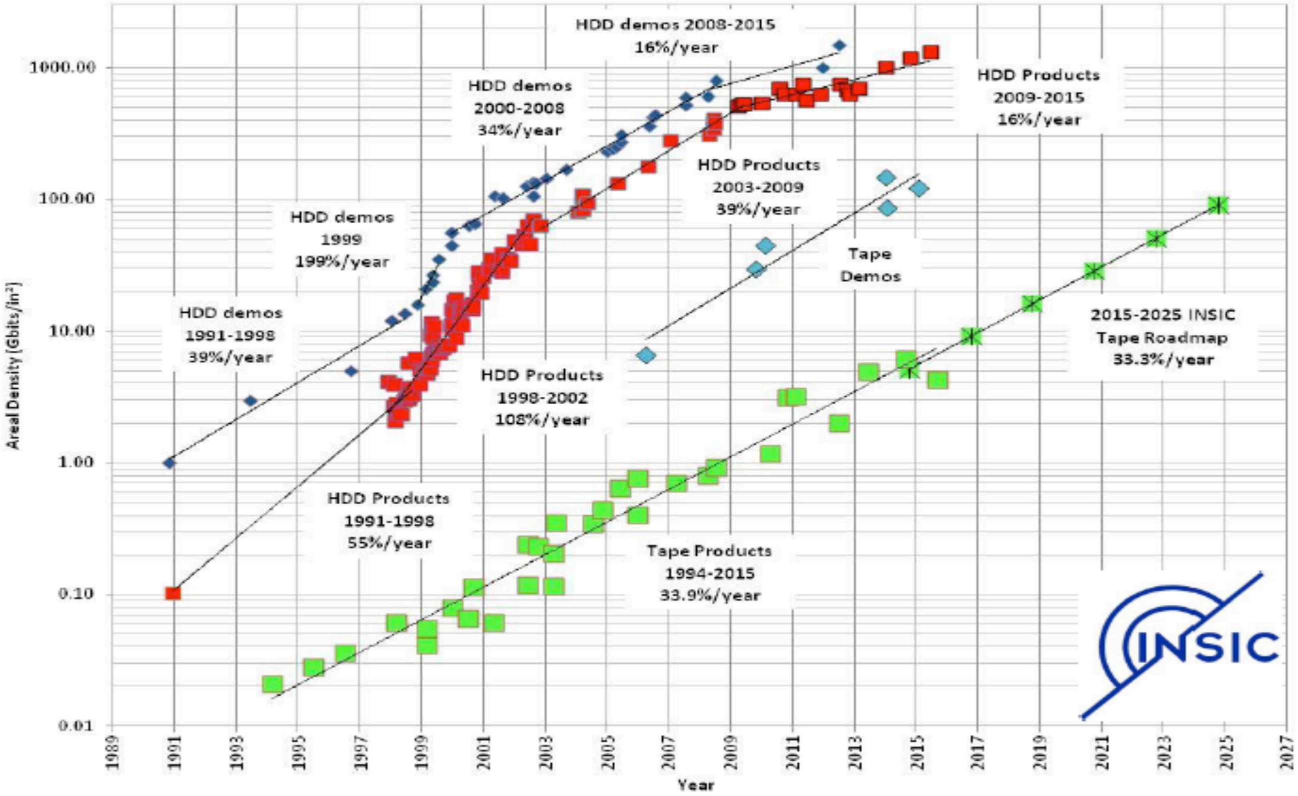

iNEMI

Tape versus HDDs

Exhibit 1 — Comparing the Average TCO for Disk and Tape for Archiving

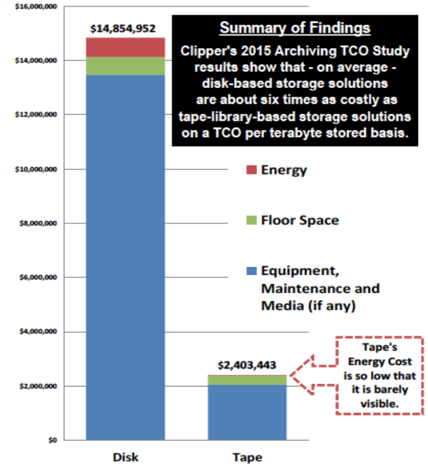

Source: *The Clipper Group*

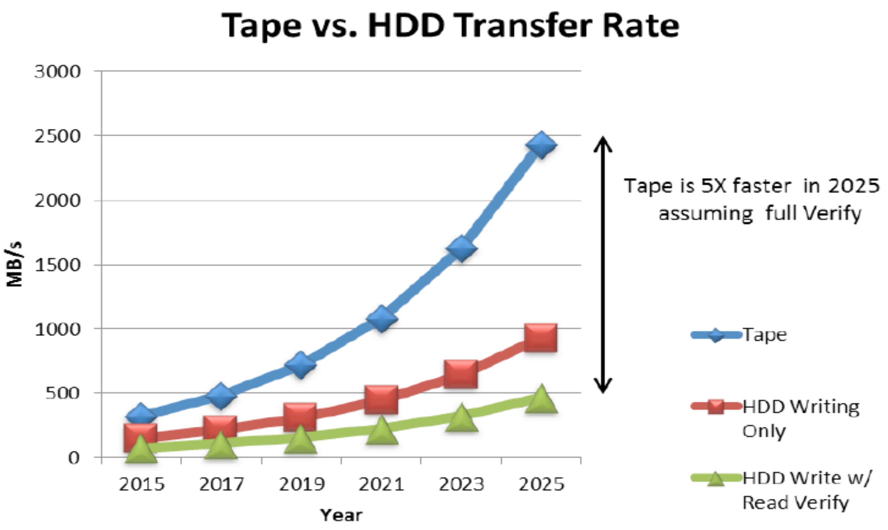

Ref: INSIC Tape Systems & Applications 2015 Roadmap

iNEMI

Mass Data Storage Chapter Status

Optical Discs

Optical Storage Trends

- Optical disc volume is in consumer products but unit shipments are declining with increase of electronic content
- Write-once Blu-ray disc library technologies from Panasonic and Sony
- Technology roadmap to 1 TB/5.25” disc, but this may require holographic recording or many recording levels

Optical Storage Technologies

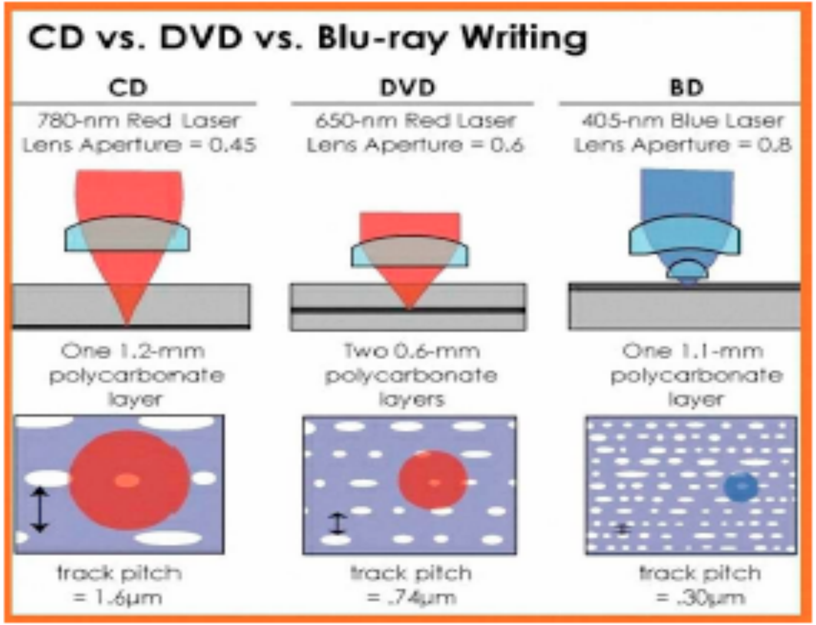

Figure 66. Decrease in laser wavelength and spot size

Blu ray Disc Roadmap

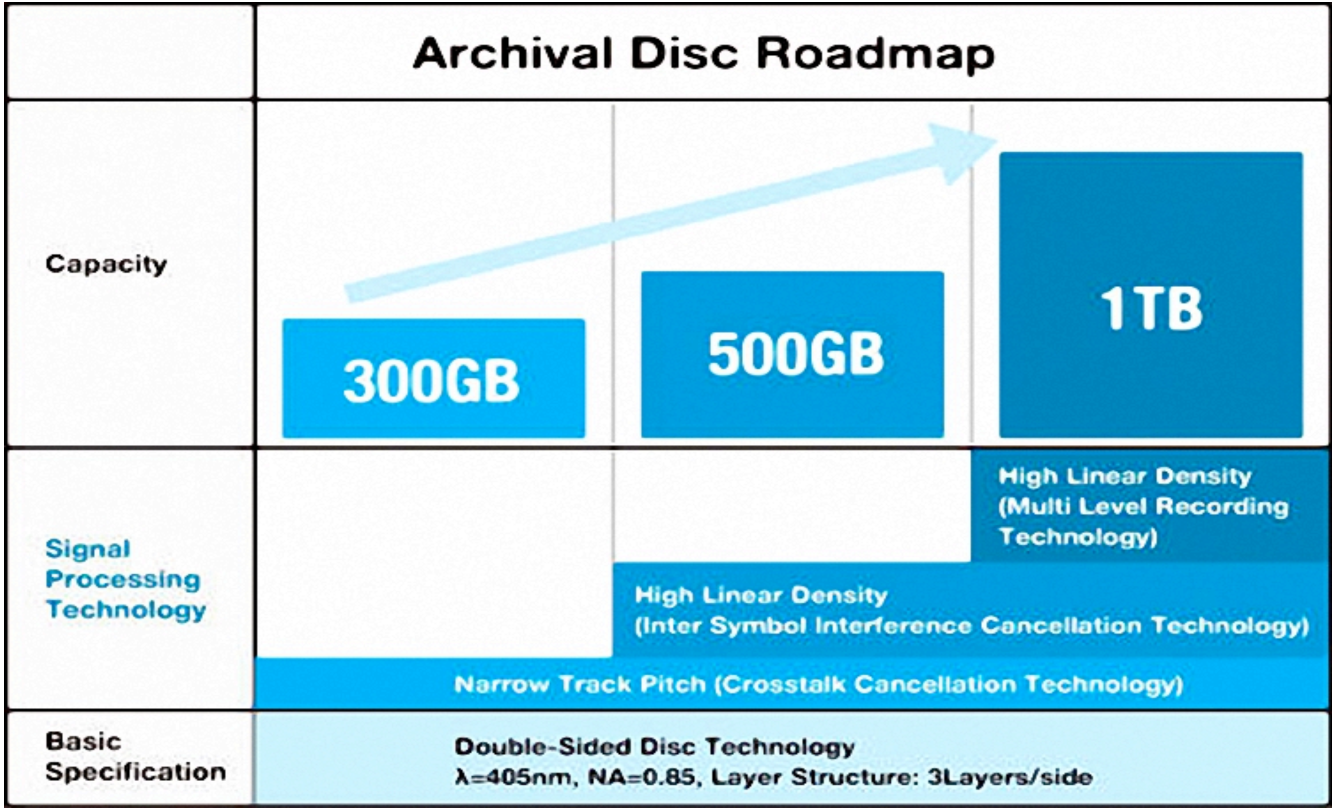

Optical Roadmap Attributes

Table 1. Optical Storage Component/Subsystem Attributes

Component	2015	2016	2018	2020	2024	Comments
Laser Wavelength (nm)	375-650	256-405	237-352	219-306	TBD	Through 2012, convergence on 405nm is expected. Beyond about 2015, UV lasers and media must be planned, if not implemented.
Laser Power (mW)	3-60	3-120	3-150	3-180	TBD	Recording speed and recording layer sensitivity are the pacing factors. Historically, this has been the range in laser powers for each generation.
Objective Lens NA	0.60-1.5	0.65-2.5	0.85-3.0	0.85-3.2	TBD	Assumes the introduction of NFR to obtain NA >1.
Disc Types	replicated, WO, RW	replicated, WO, RW	replicated, WO, RW,	replicated, WO, RW	TBD	Media types will probably stay the same. WO will survive the roadmap period (CD-R is currently the biggest selling type of optical medium).
Recording Layers	1-4	1-8	1-12	1-16	TBD	Each side of the disc. By 2015, areal densities will be so high that cartridge media will probably be required.
Data Encoding + Read Channel	RLL/PRML	TBD	TBD	TBD	TBD	Multi-level, multi-layer, NFR, and combinations will require significant coding and signal processing as 100 Gb/in^2 areal densities are approached and exceeded.

iNEMI

Sony Everspan Robotic Disc Library

Mass Data Storage Chapter Status

The Cloud

iNEMI

The cloud and the fog

- The Cloud refers to compute resources, including storage, located in large data centers
- The Fog refers to local networks that connect thing (e.g. IoT) together
- Local fog networks may connect to the Internet

iNEMI

The Role of the Fog

Source: IOT at the Network Edge, http://www.nojitter.com/post/240172079/iot-at-the-network-edge

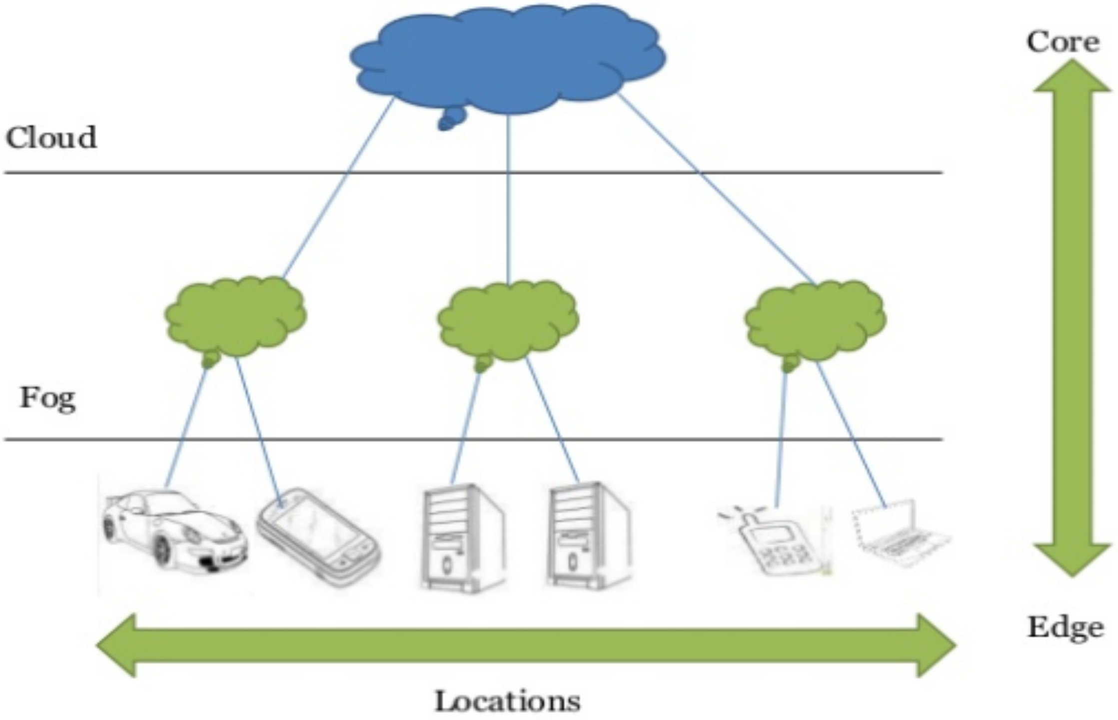

Connected Cars and Smart Cities

From: Tao Zhang, Cisco Distinguished Engineer, Co-Founder and Board Director of OpenFog Consortium

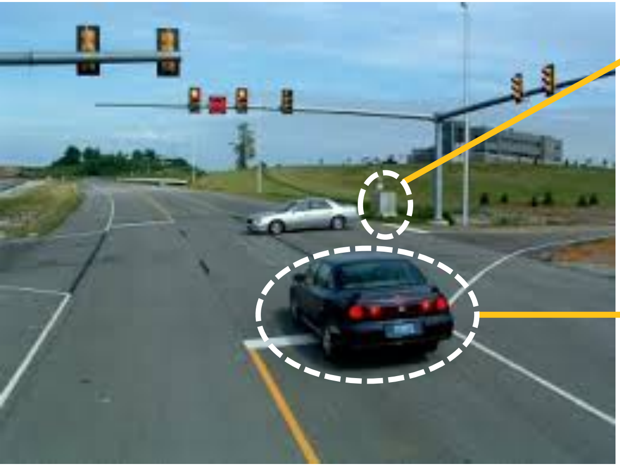

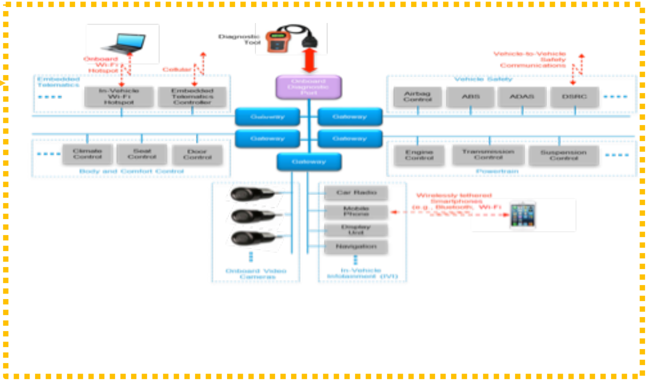

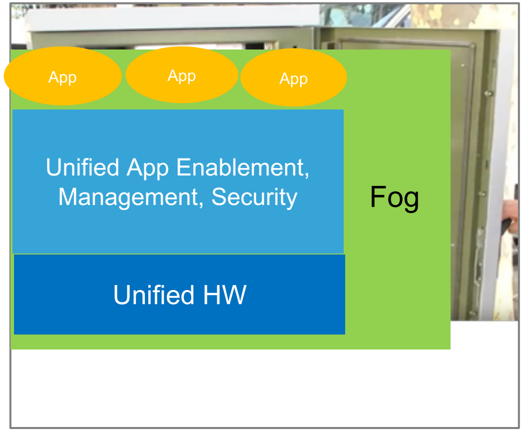

Driving Factors in Cloud Storage

- The scale of the data centers for large cloud facilities drives hyperscale computing (and storage) architectures
- This includes SDX, including storage and virtualization to get the greatest equipment utilization
- Control of energy use (especially dealing with heat) are a big factor in cost—green matters
- Because of the range of services expected from a full-service cloud provider they need a wide range of equipment—including storage tiers

Databases

- Large data sets
- Random traffic
- High I/O load
- Early SSD adopter
  - Previously used DRAM SSDs
- Some load the entire DB on flash memory

Archiving & Backup

- Snapshots and replication gaining momentum
  - Both require high-speed storage
- Business continuity places high demands on storage
- Active archives growing faster than passive archives (favoring HDDs rather than tape)

Cloud Storage/Services--Virtualization

- The "IO Blender"
  - Many streams
  - Scrambled I/O
  - Highly random
- Suits SSDs better than HDDs for rapid access
- Many VM and VDI systems using flash cache to meet demand speed needs

Image courtesy of Waring Corp.

IOPS Required for Dominant Application

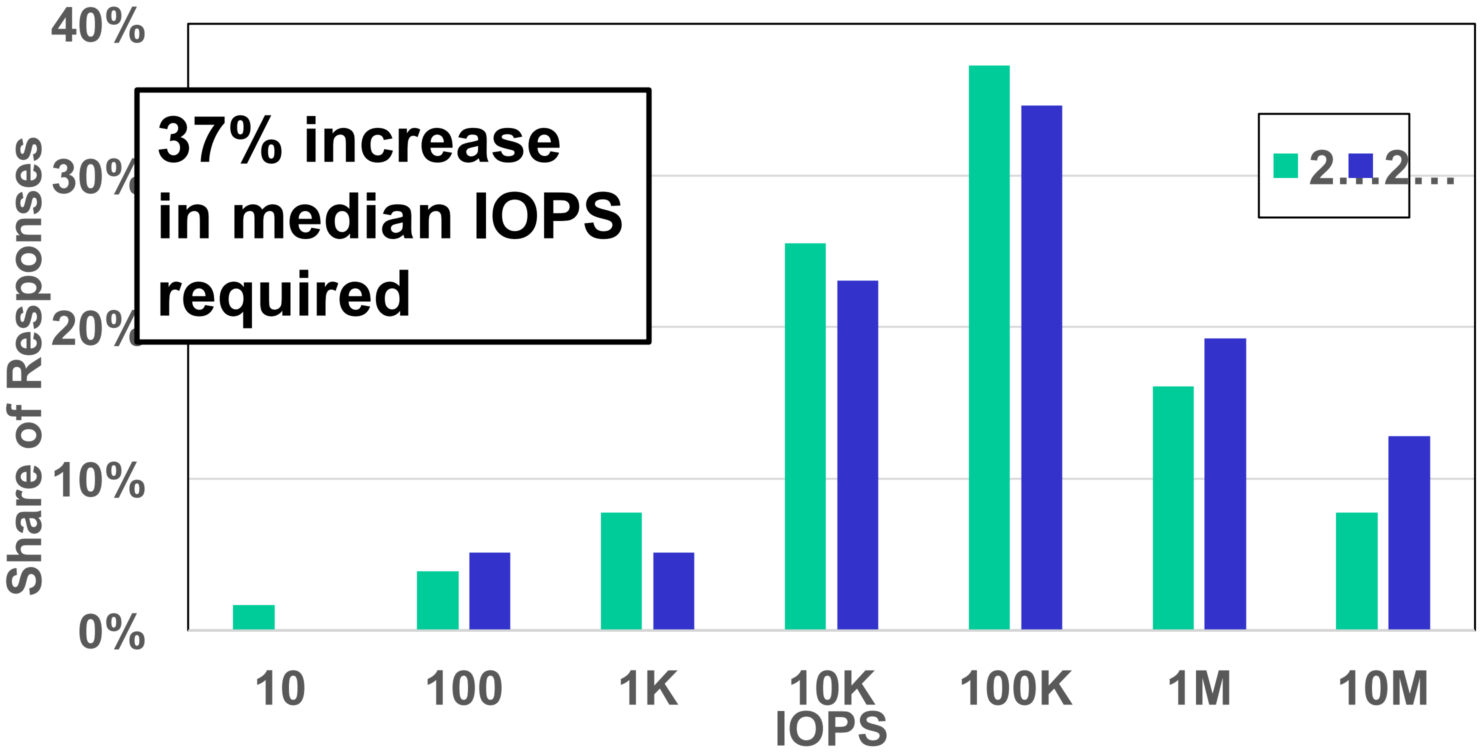

Capacity Required

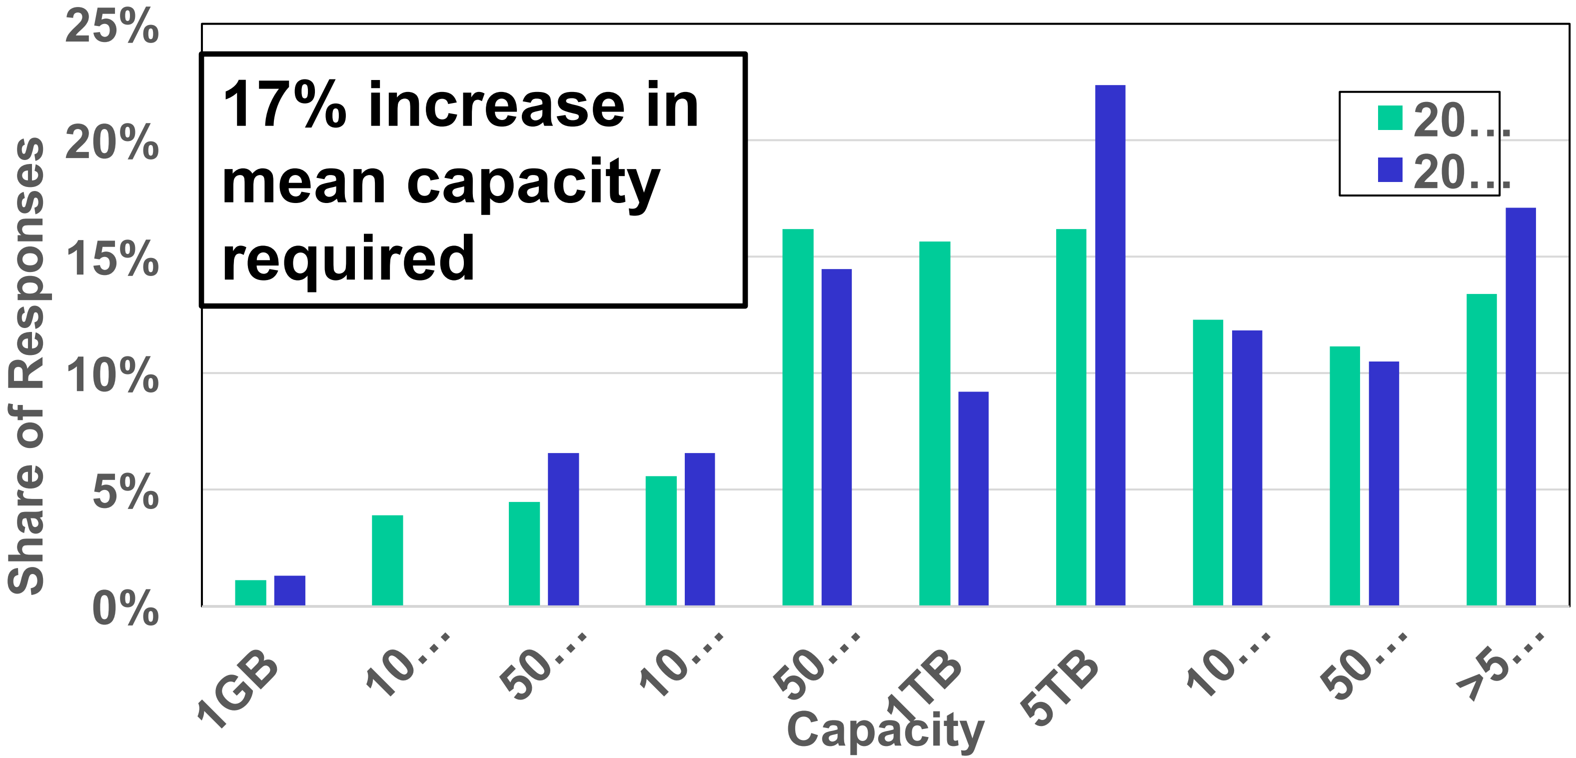

Other Hardware IOPS Bottleneck

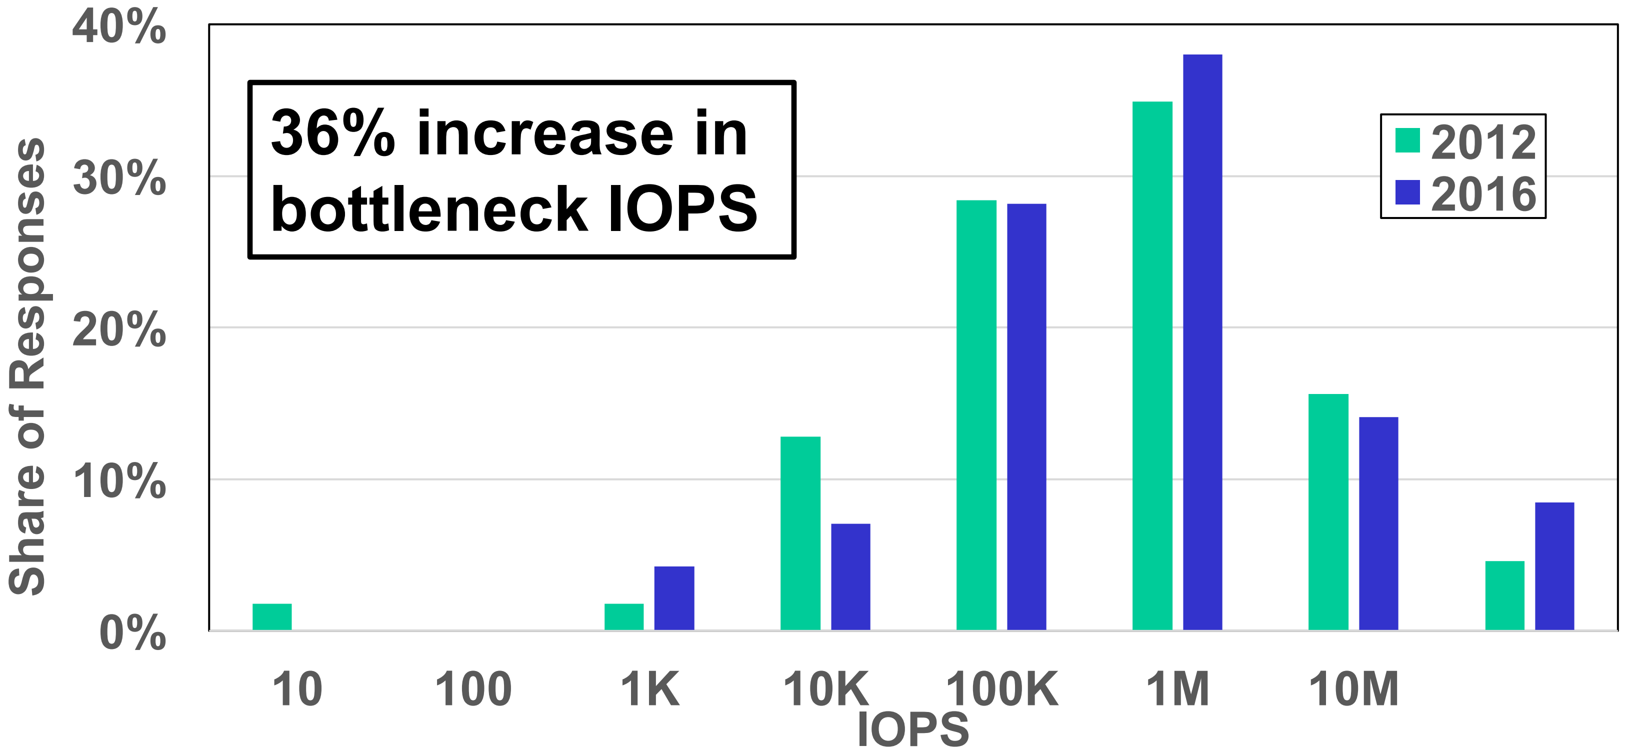

Fastest Latency the System Can Use

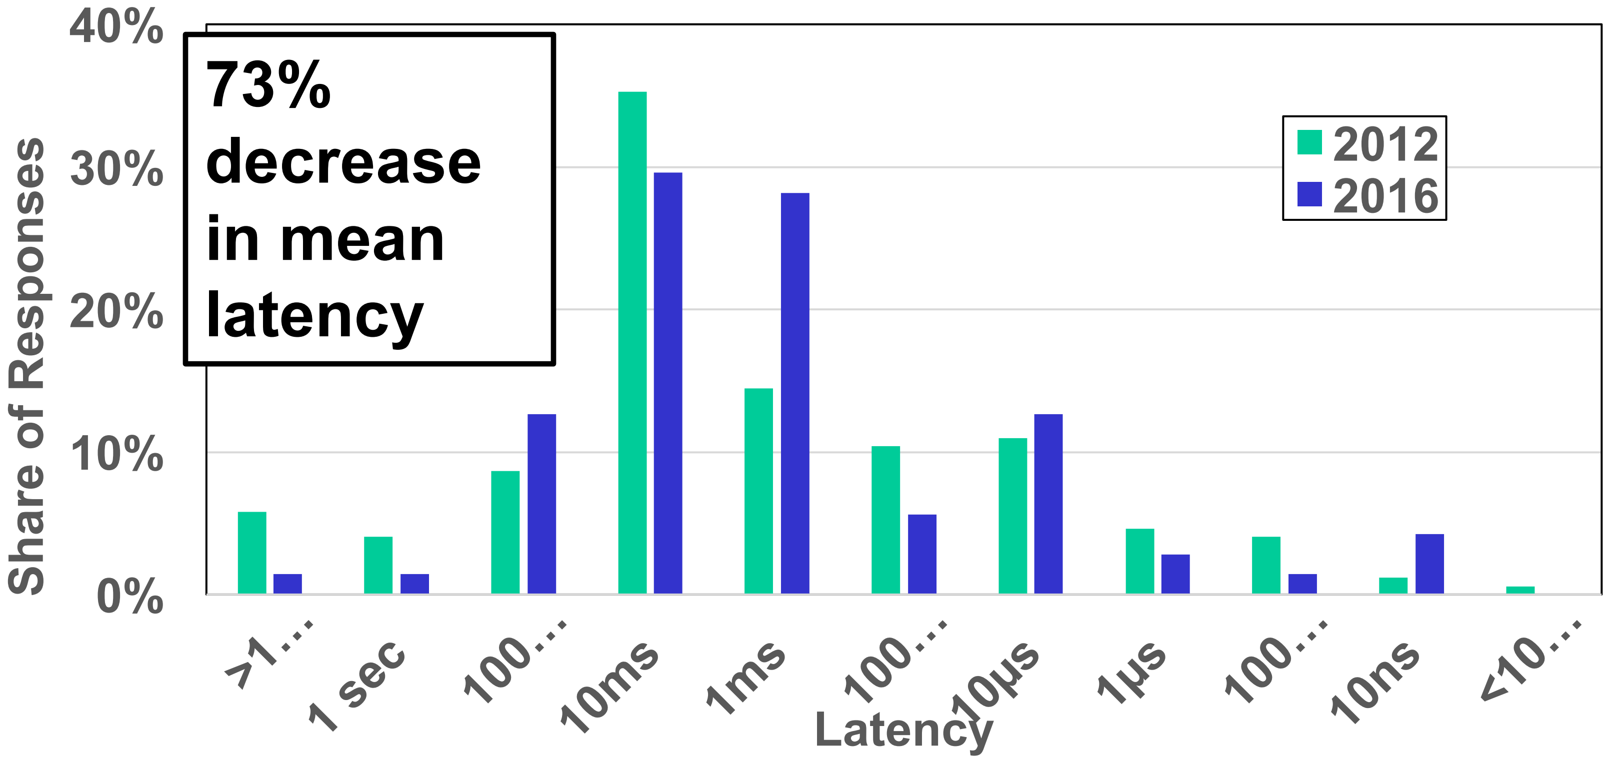

Cloud Storage Tiers

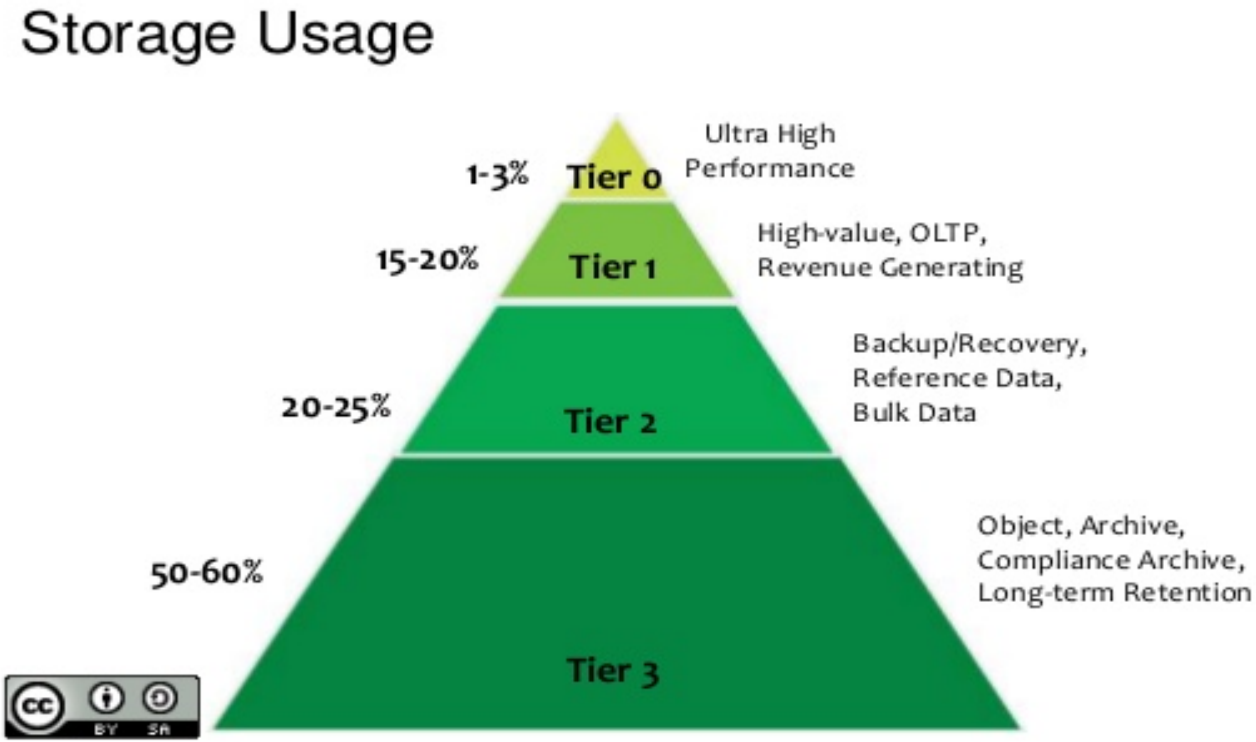

iNEMI

Storage Devices Used in the Cloud and Fog

- **Cloud**
  - DRAM and perhaps emerging memories (e.g. 3D Xpoint and MRAM) (high-performance tier)
  - SSDs and all-flash arrays (performance tier)
  - Capacity HDDs (capacity tier)
  - Tape or Optical Discs (archive tier)

- **Fog**
  - SSDs or flash memory are favored because of their reliability under more harsh conditions—such as street corners.

Conclusions

- Solid State/Flash storage migration to 3D architecture
- Introduction of 3D X-point / phase change technology
- MRAM continuing to find applications – suitable for IoT
- HDD near term volumes down due to
  - Migration from Client/Server to Mobile/Cloud systems
  - Higher utilization of available HDD capacity
  - Areal density will continue to grow at ~15% CAGR
- Tape will continue to be cost effective and show capacity progress
- Optical will continue to find niche applications
- Cloud and fog storage driven by enterprise and IoT trends and will utilize many different types of storage

References

- 2017 iNEMI Mass Storage Roadmap, iNEMI
- 2016 How Many IOPS Do You Really Need Report, Coughlin and Handy, http://www.tomcoughlin.com/techpapers.htm
- Touch Rate: A metric for analyzing storage system performance, Steven Heltzer and Tom Coughlin, 2015, http://www.tomcoughlin.com/techpapers.htm