Why People Use Solid State Storage

Tom Coughlin

Coughlin Associates

www.tomcoughlin.com

Tom Coughlin

Tom Coughlin, President, Coughlin Associates is a widely respected digital storage analyst as well as business and technology consultant. He has over 37 years in the data storage industry with multiple engineering and management positions at high profile companies.

Dr. Coughlin has many publications and six patents to his credit. Tom publishes the *Digital Storage Technology Newsletter, the Media and Entertainment Storage Report*, the *Emerging Non-Volatile Memory Report* and other industry reports. Tom is also a regular contributor on digital storage for Forbes.com and other blogs.

Tom is the founder and organizer of the Annual Storage Visions Conference (www.storagevisions.com), a partner to the International Consumer Electronics Show, as well as the Creative Storage Conference (www.creativestorage.org). He has been the general chairman of the annual Flash Memory Summit, the world's largest independent storage event. He is a Senior member of the IEEE and a member of the Consultants Network of Silicon Valley (CNSV). For more information on Tom Coughlin and his publications go to www.tomcoughlin.com.

Outline

- Storage Drivers, Touch Rate and Projections
- How Many IOPS are Enough?
- Non-Volatile Memory Trends
- Solid State Drive Advances
- Storage Fabrics and Memory Centric Computing
- Why Do People Use Solid State Storage?

Storage Drivers and Touch Rate

Drivers for Storage and Memory

- Increasing storage demands—IDC 163 Zetabytes of data created by 2025 (16 ZB in 2016)
- New sources for unstructured data from media and entertainment, IoT, medicine, geo-science and big data
- Growth in local storage, storage at the edge (or the fog) and storage in large data centers (the cloud)
- There is a need for fast memory and storage to support processing and accessing this data and cheap storage to keep it for the long term

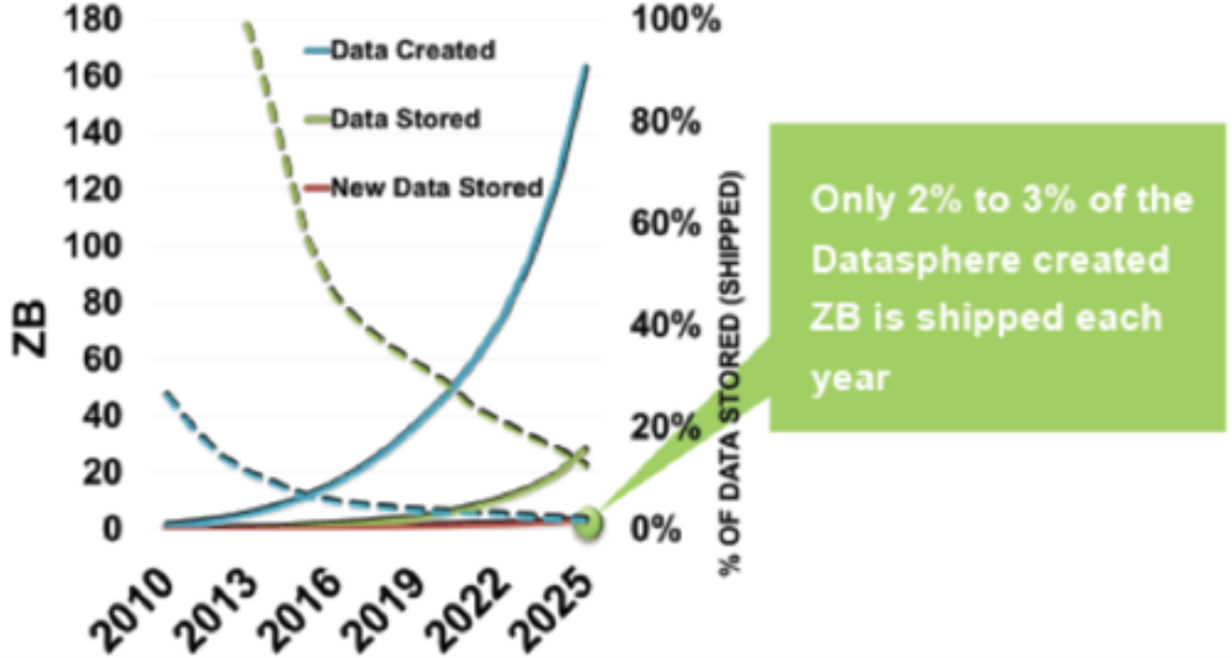

Touch rate versus response time indicating various types of uses

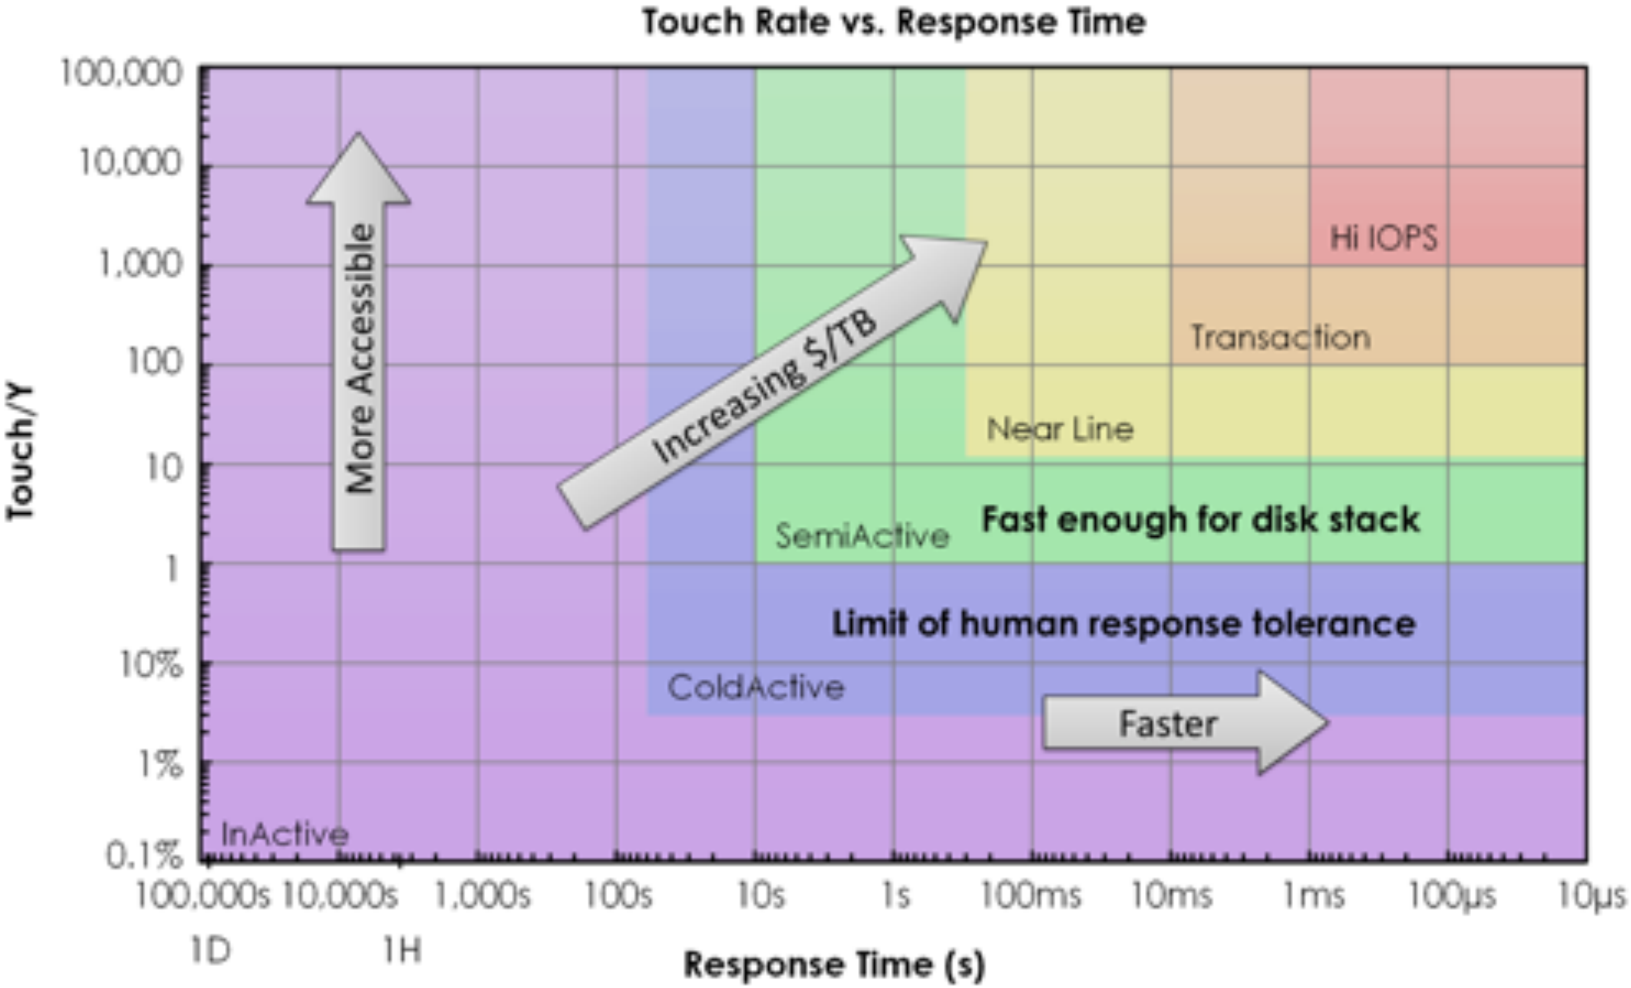

Digital storage technologies regions overlaid on the Touch Rate/Response Time chart

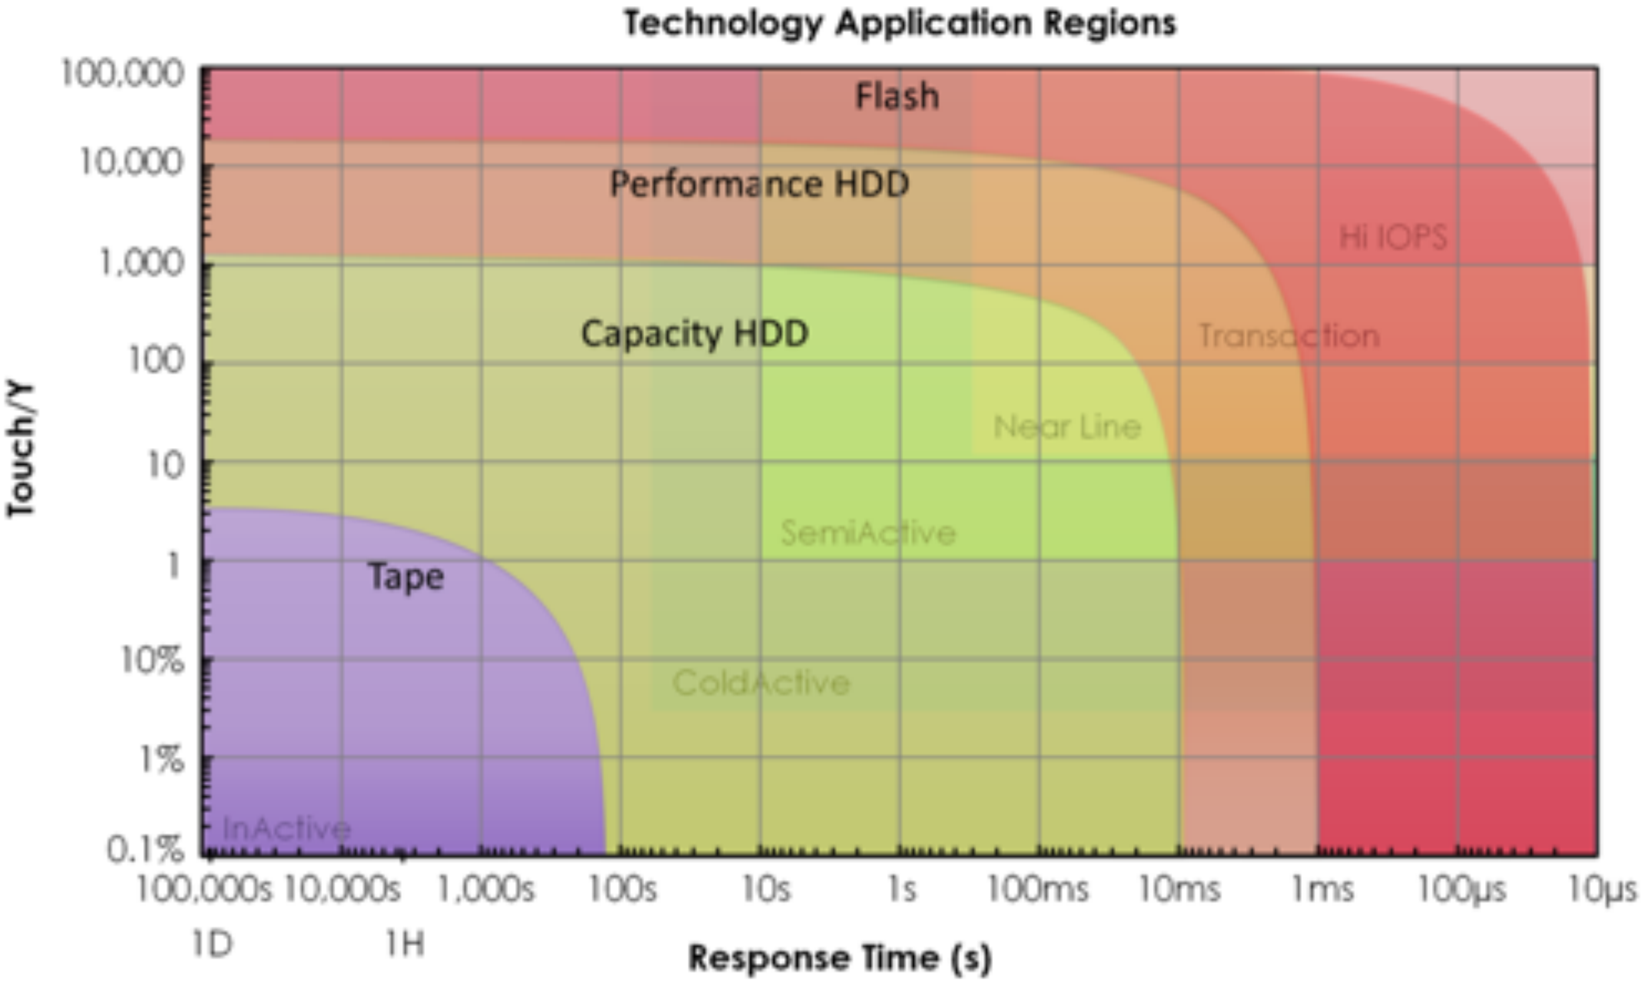

HDD-Flash tiering/caching touch rate chart

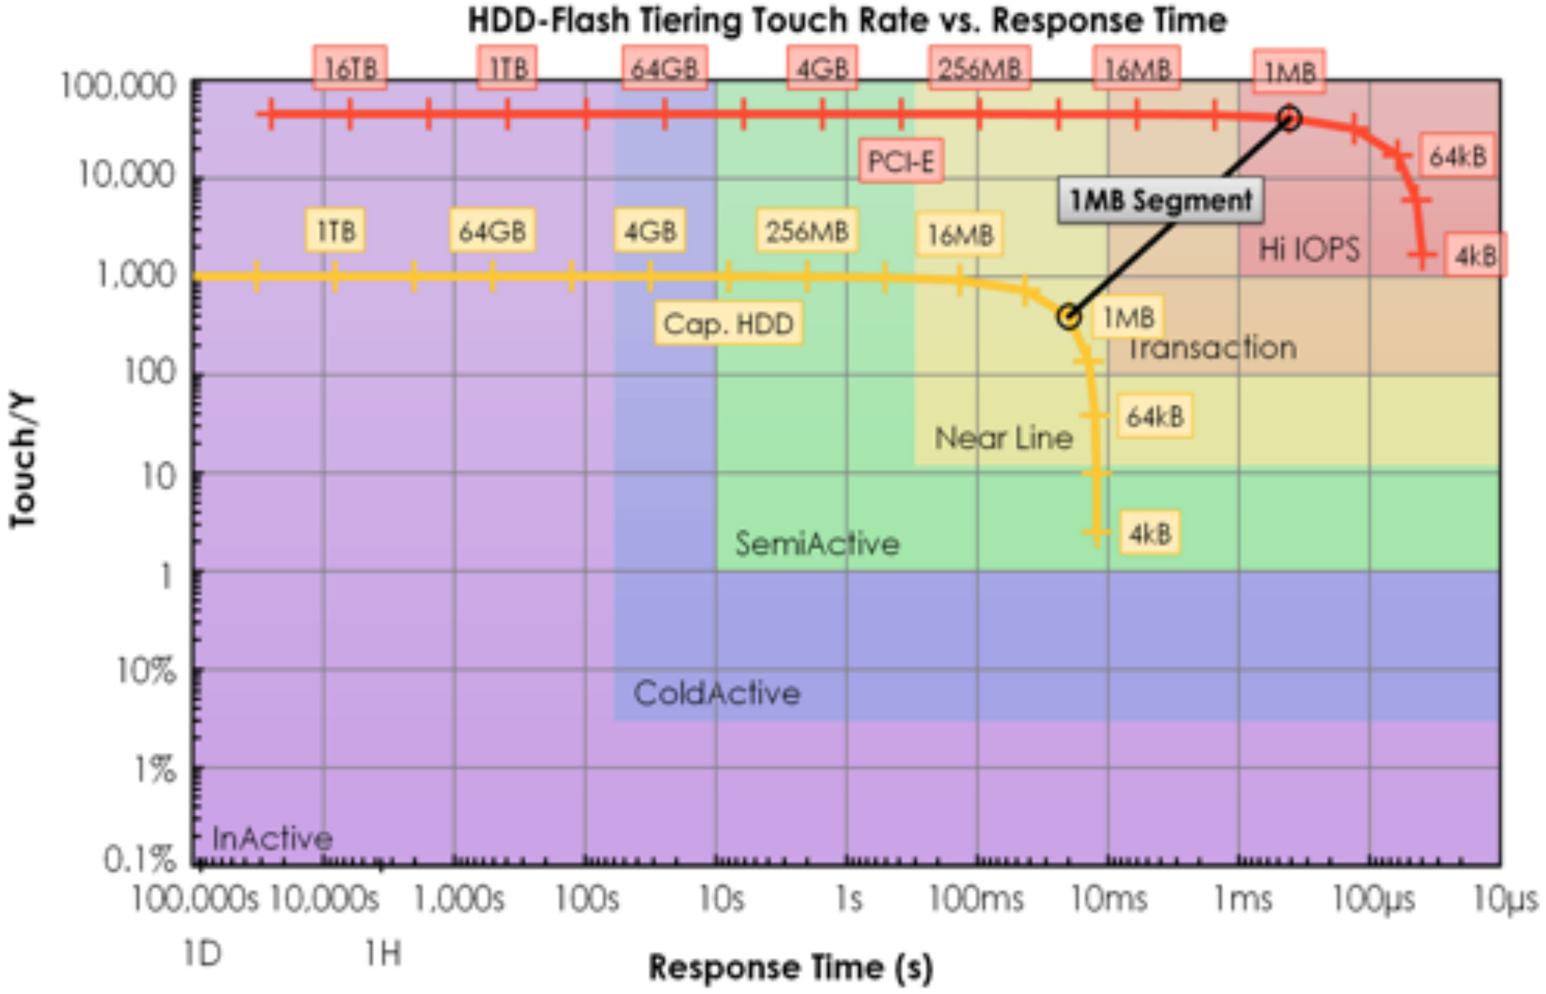

© 2017 Coughlin Associates

Digital Storage Capacity Projections

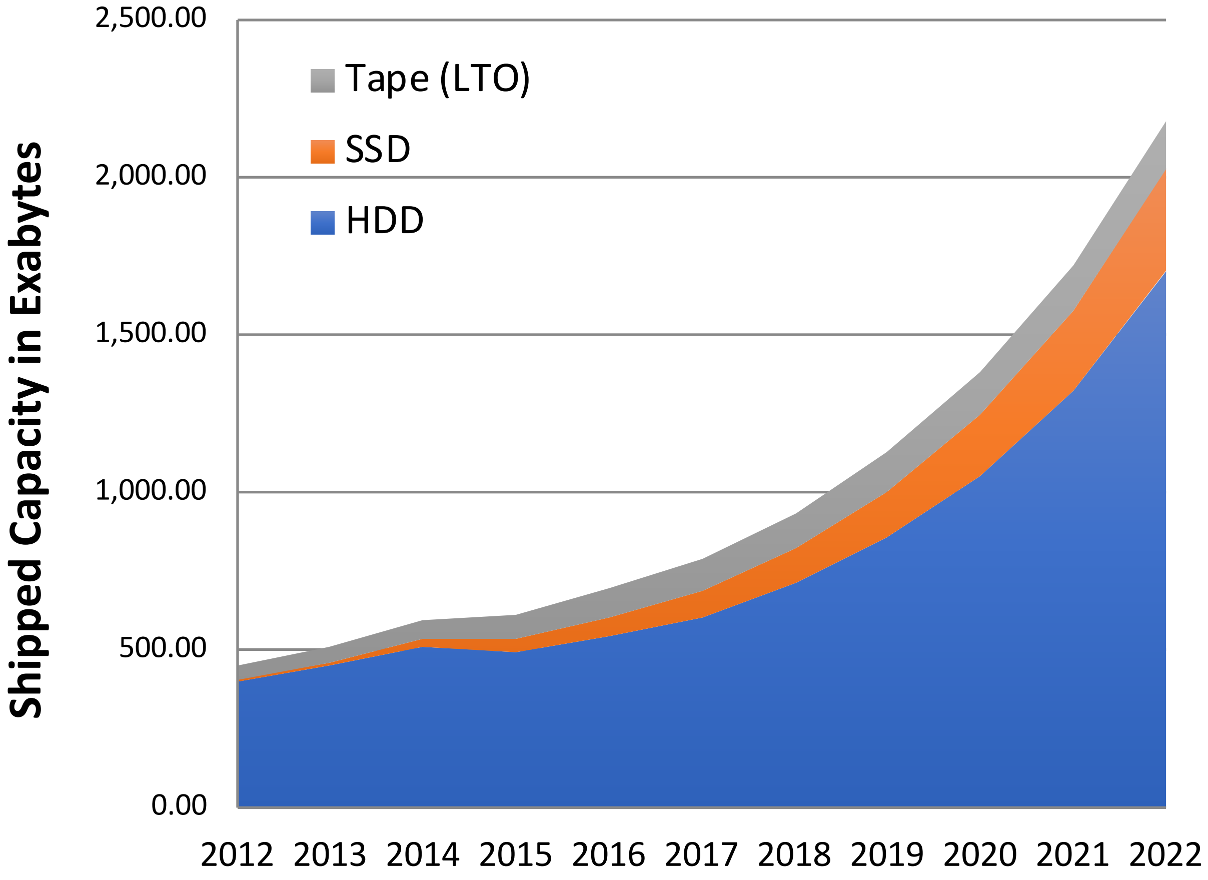

- The growth and processing of data will lead to the use of many types of digital storage
- SSDs will dominate for high performance storage and higher total revenue
- HDDs will be high capacity and used for colder storage
- Magnetic tape will be used by some organizations for the lowest cost (currently <1 cent/GB)

How Many IOPS are Enough?

Our Survey
(Objective Analysis and Coughlin Associates)

- Ongoing. Take our survey at: http://TinyURL.com/IOPSsurvey
- IT participants participating
- Asks for IOPS, capacity and latency needs
  - Also their primary applications
- Some results are in a SNIA SSSI white paper
- We compared results from 2012 to those in the last few years (in 2016)

Applications: 2012

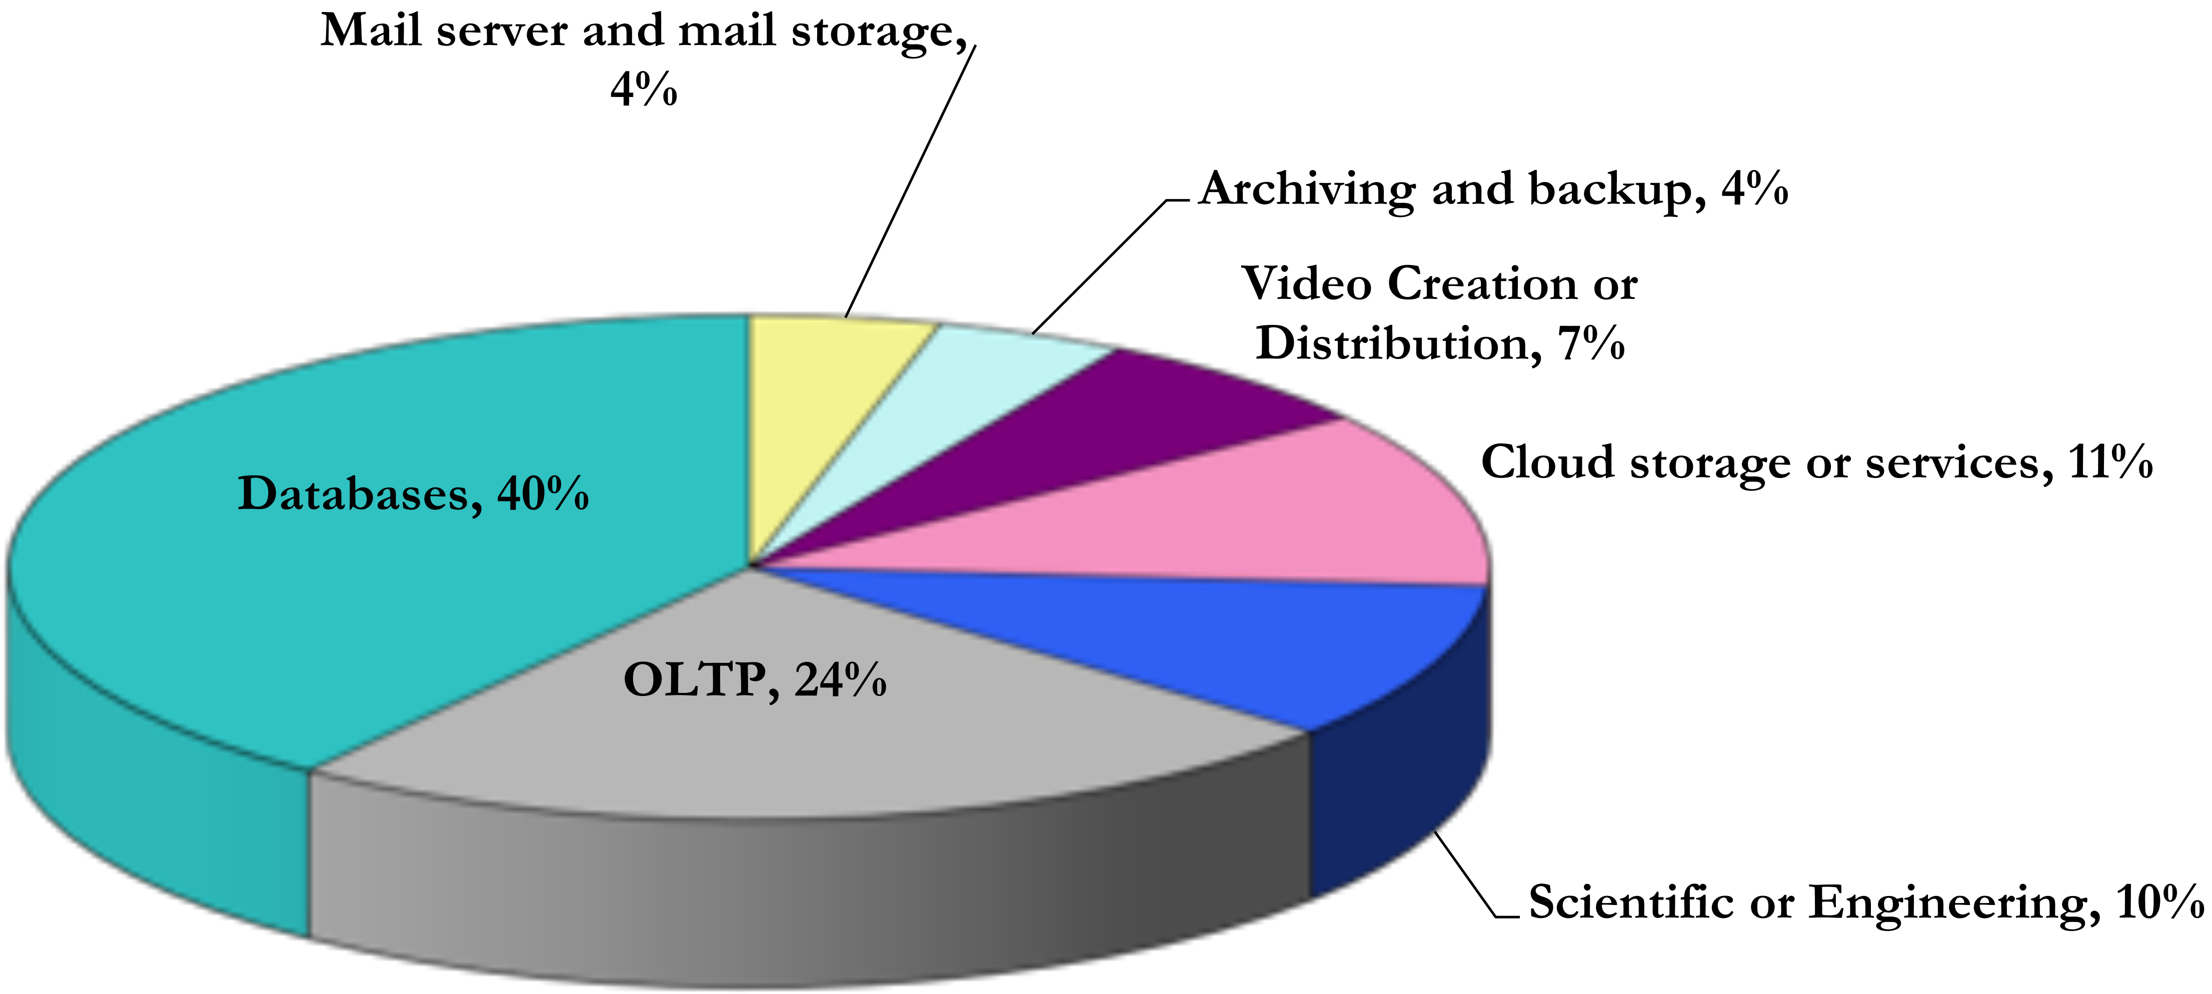

Applications: 2016

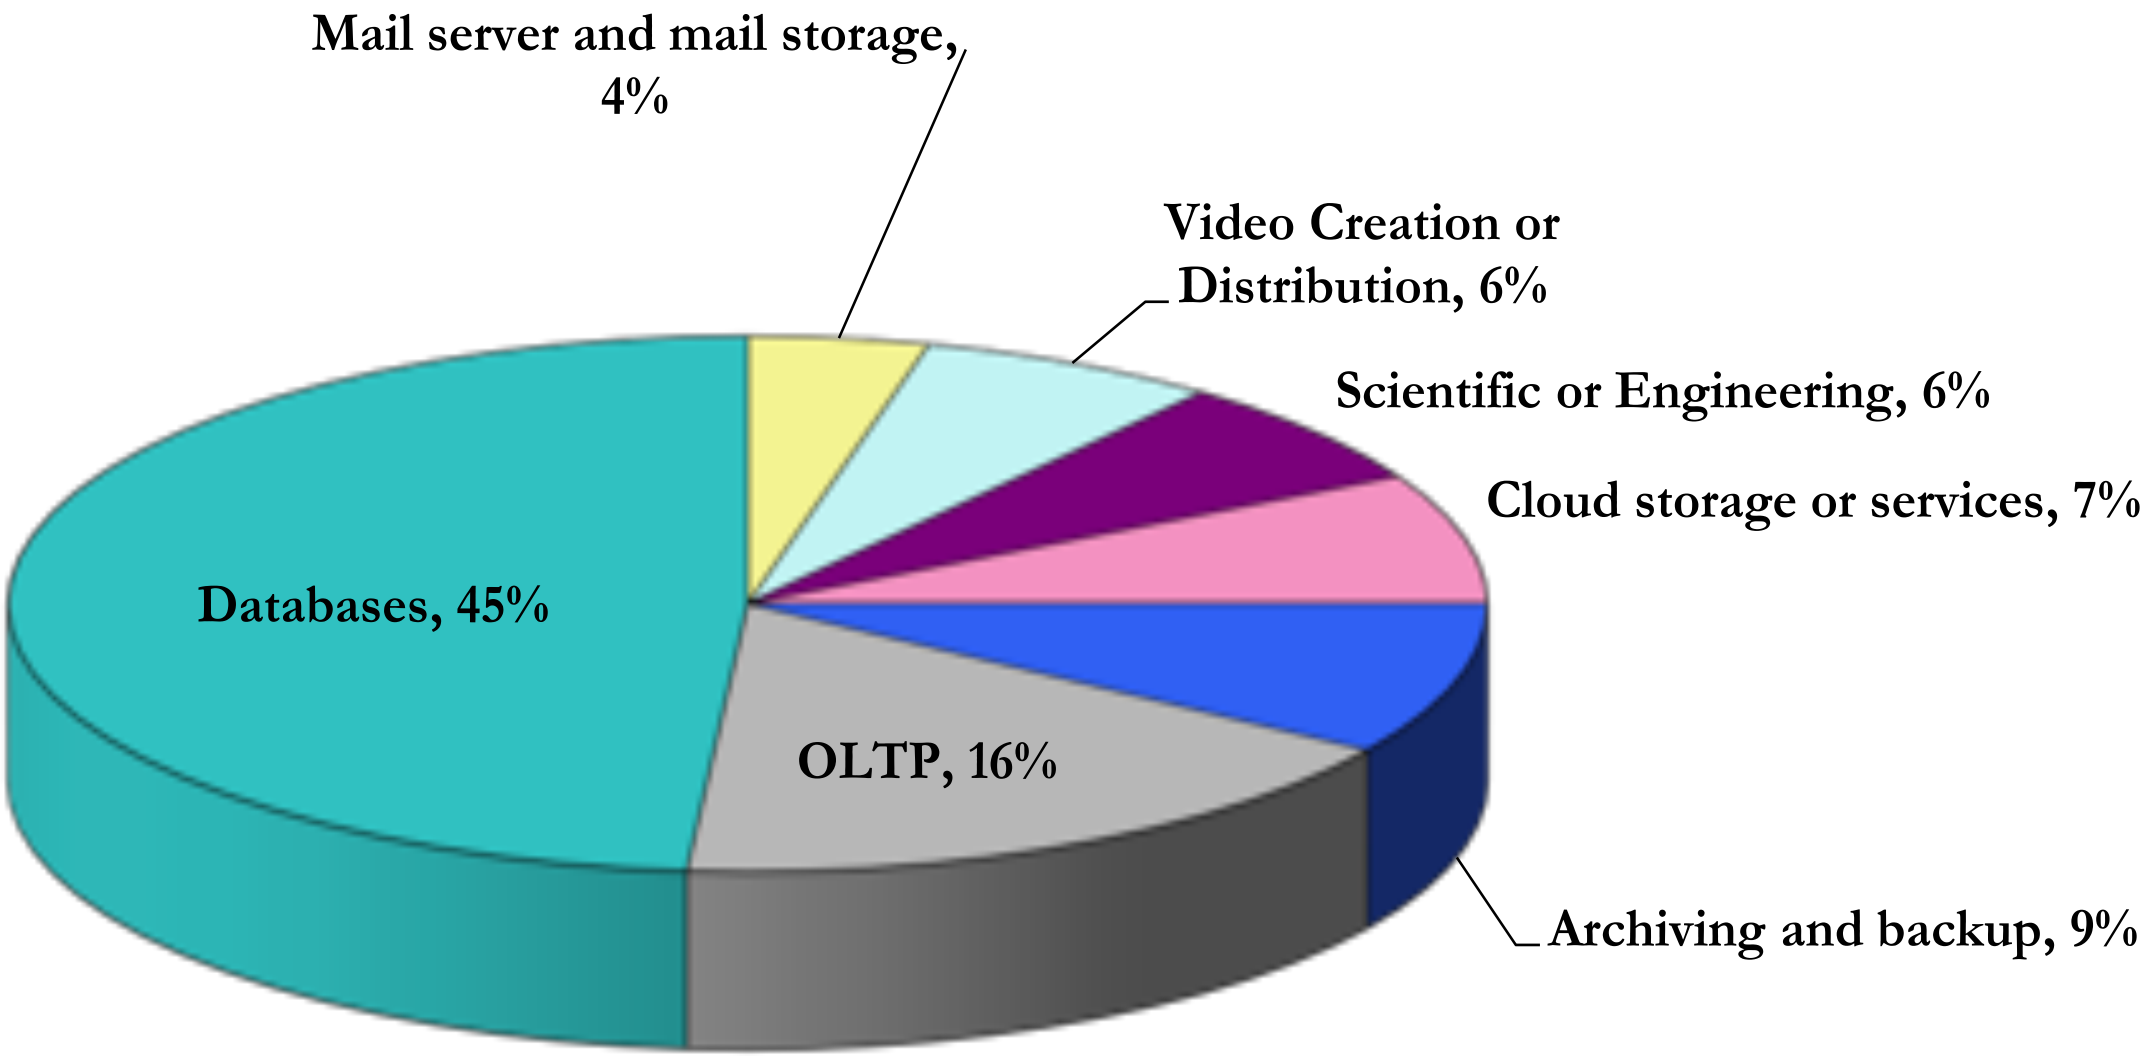

IOPS Required for Dominant Application

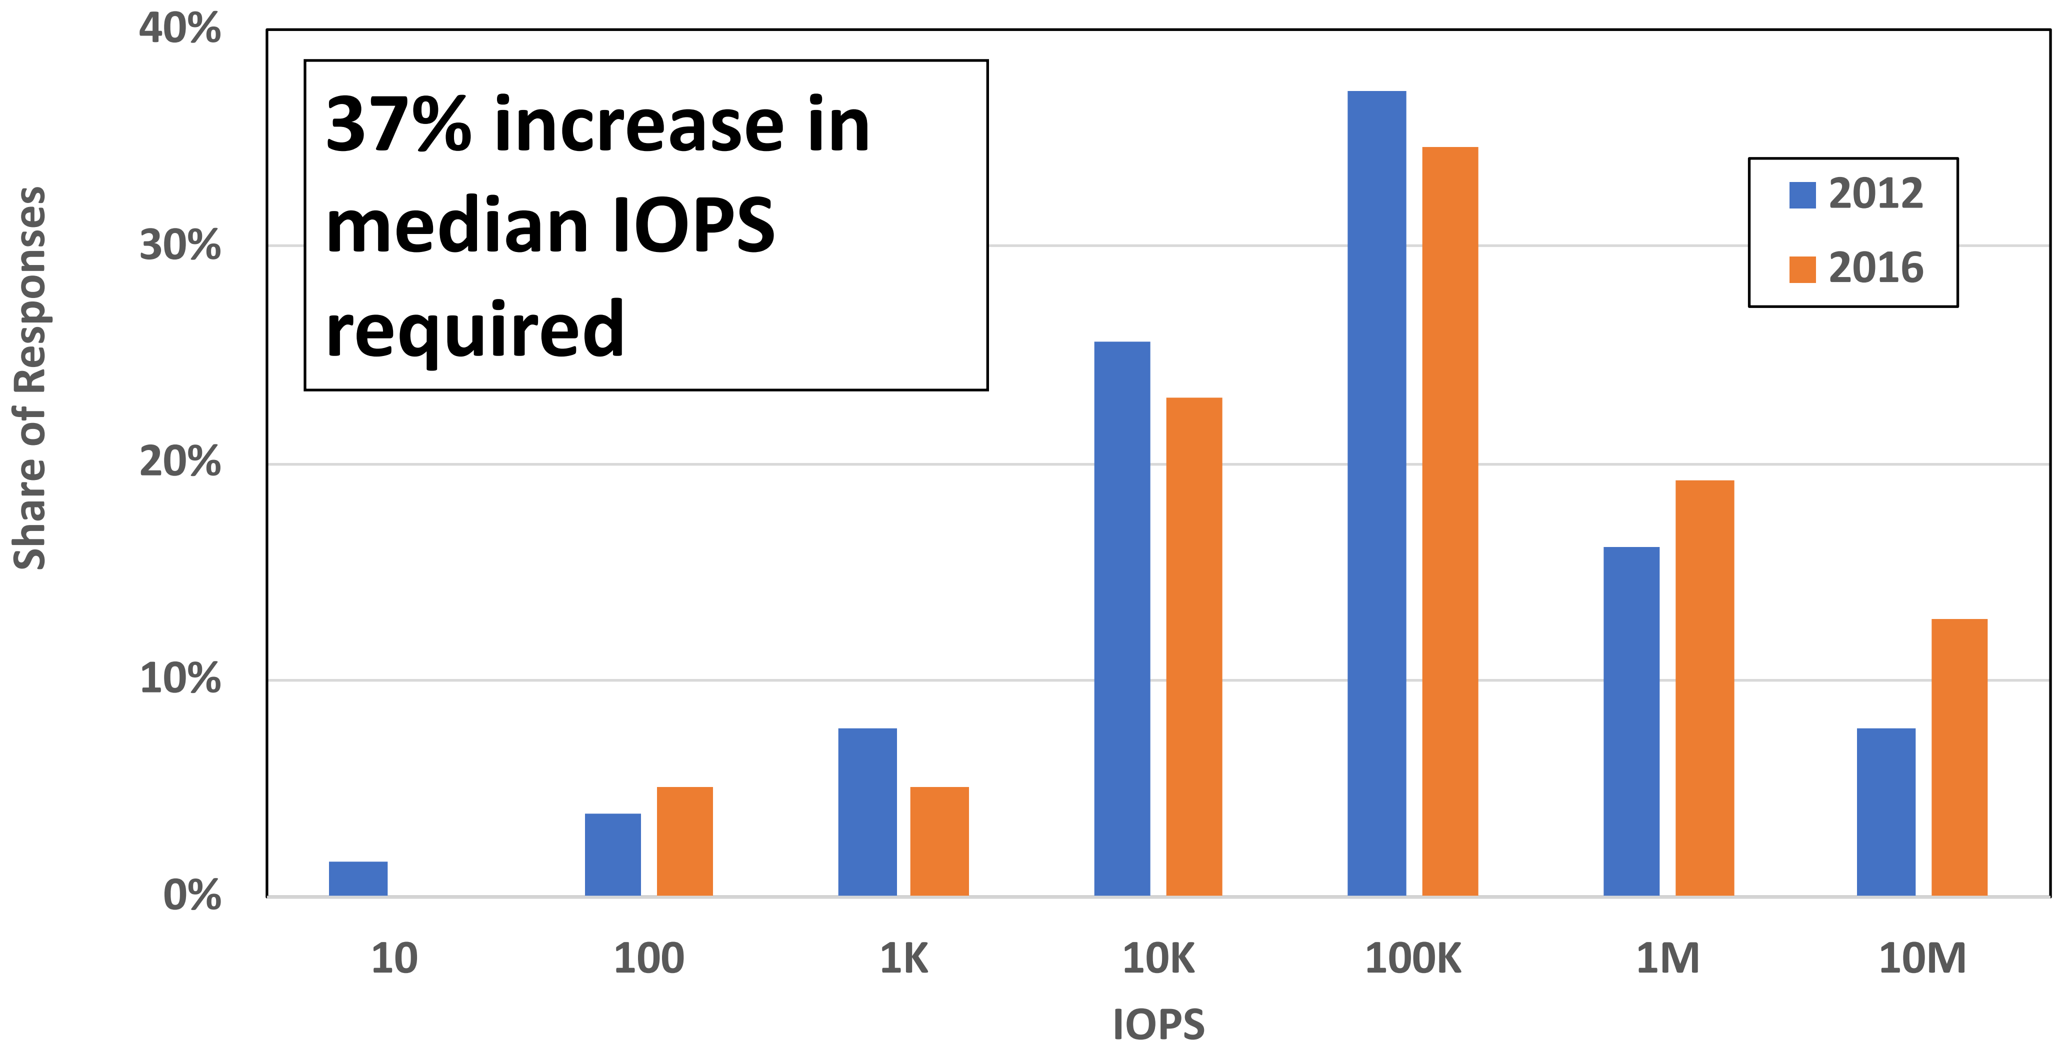

Capacity Required

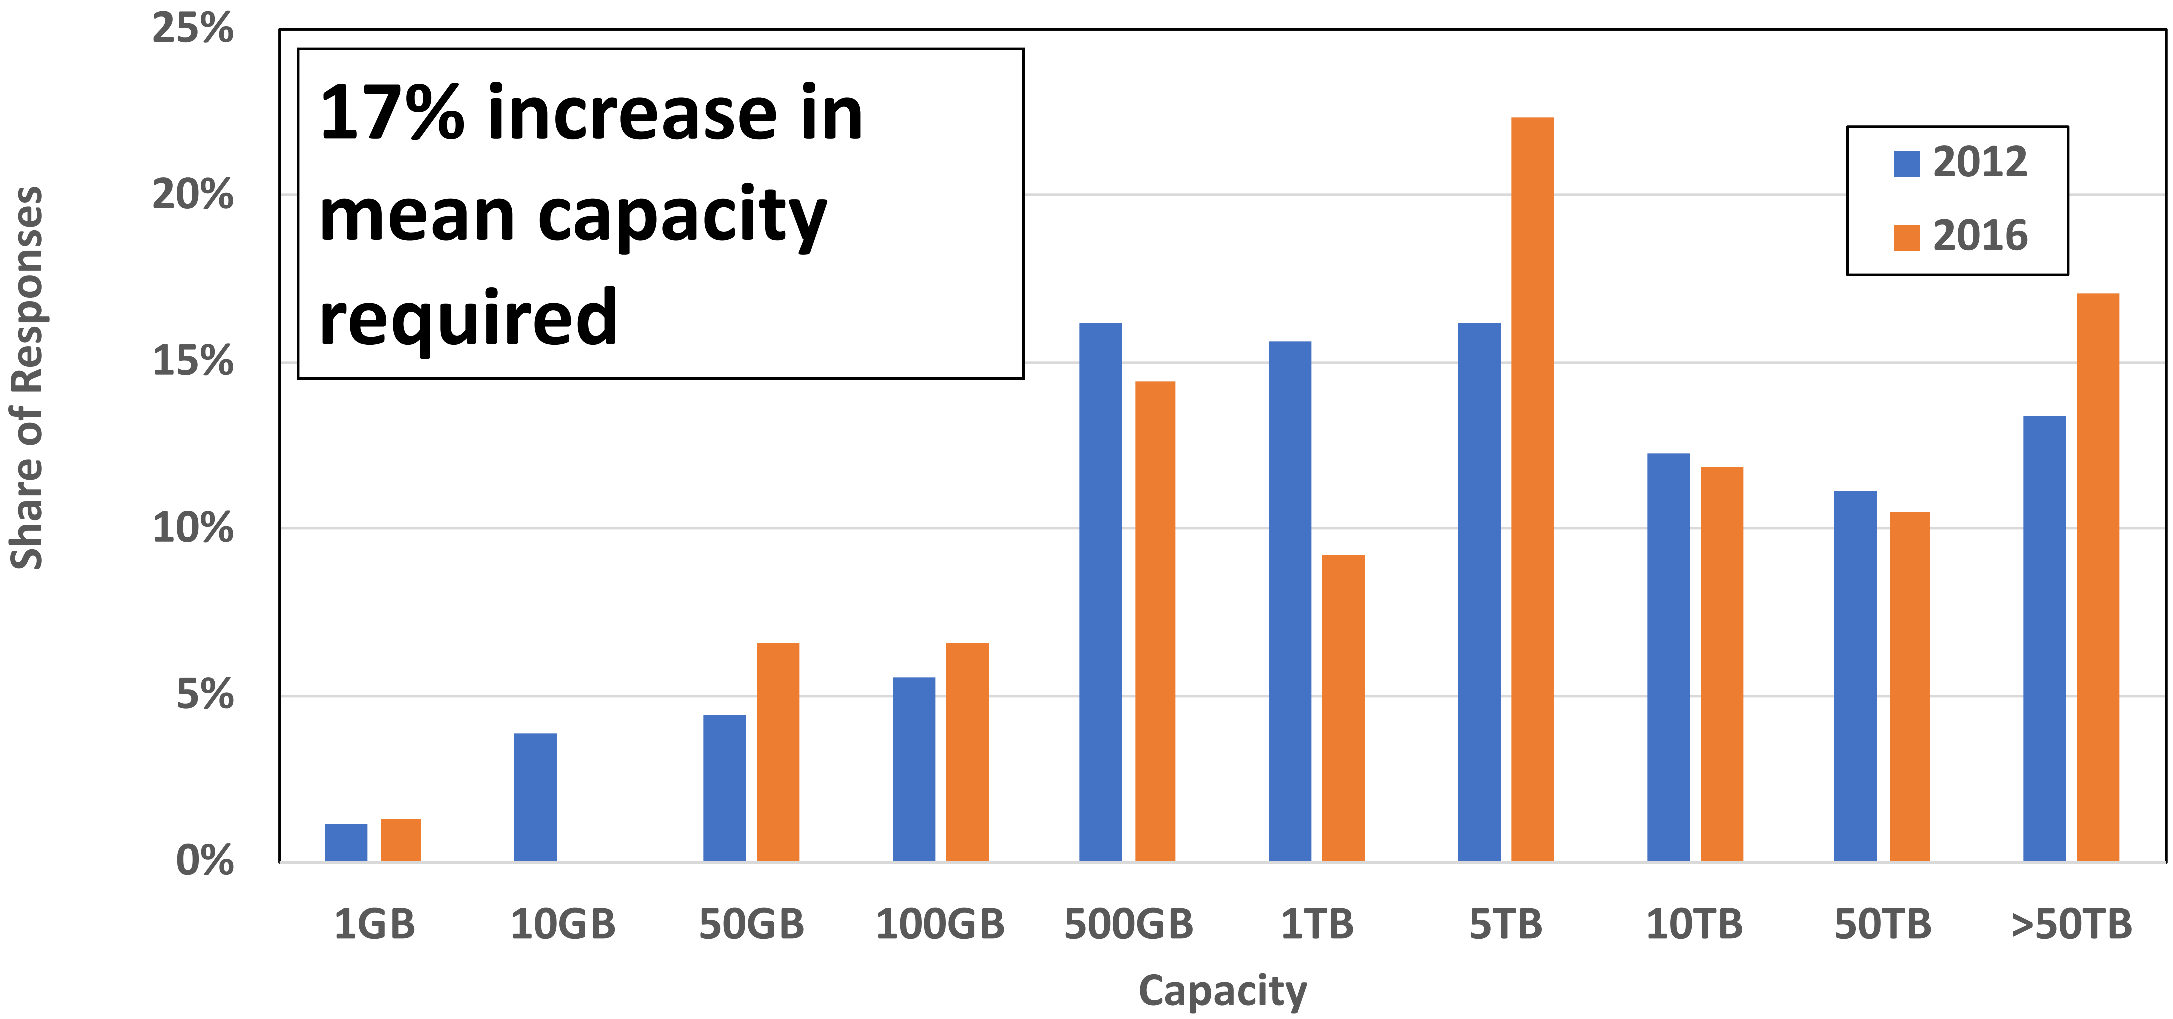

Other Hardware IOPS Bottleneck

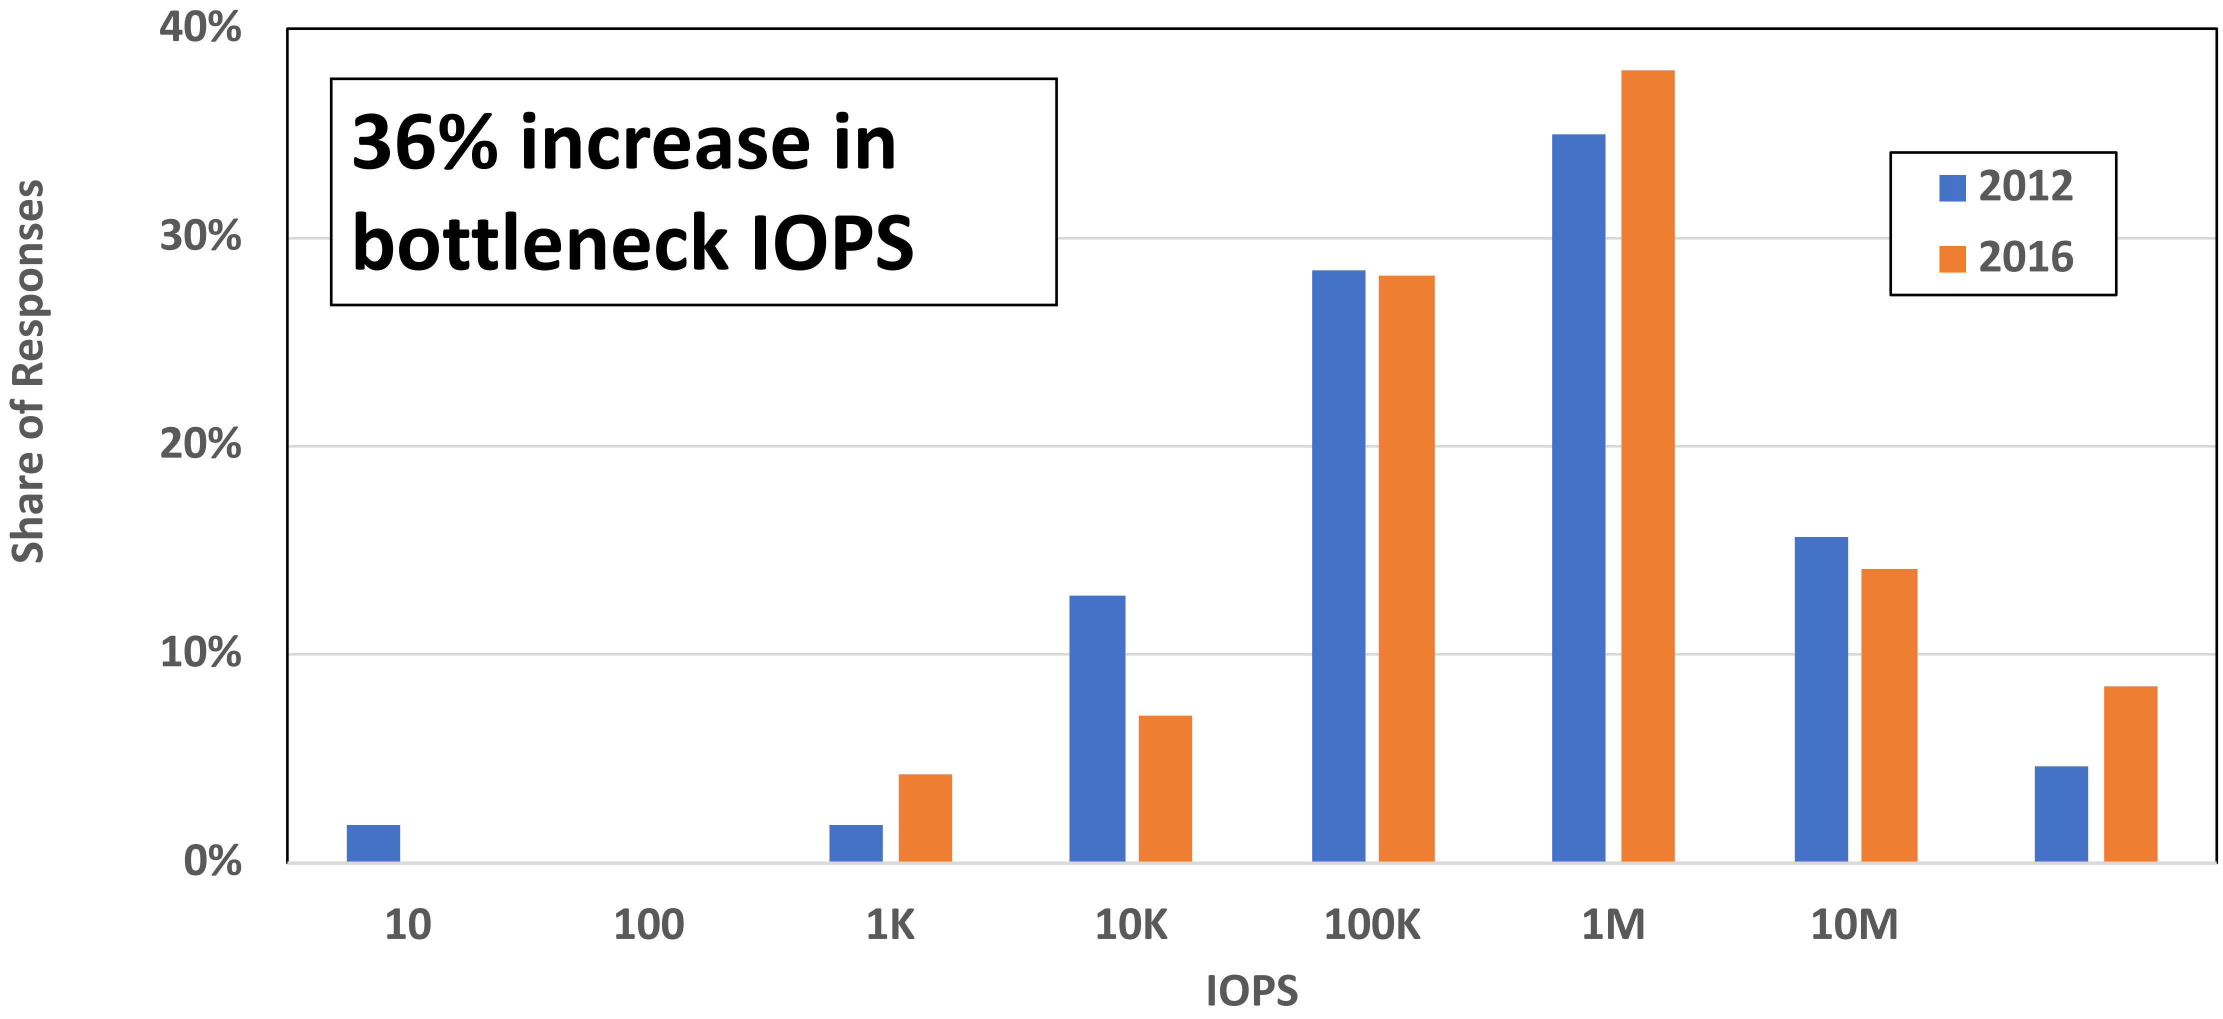

Fastest Latency the System Can Use

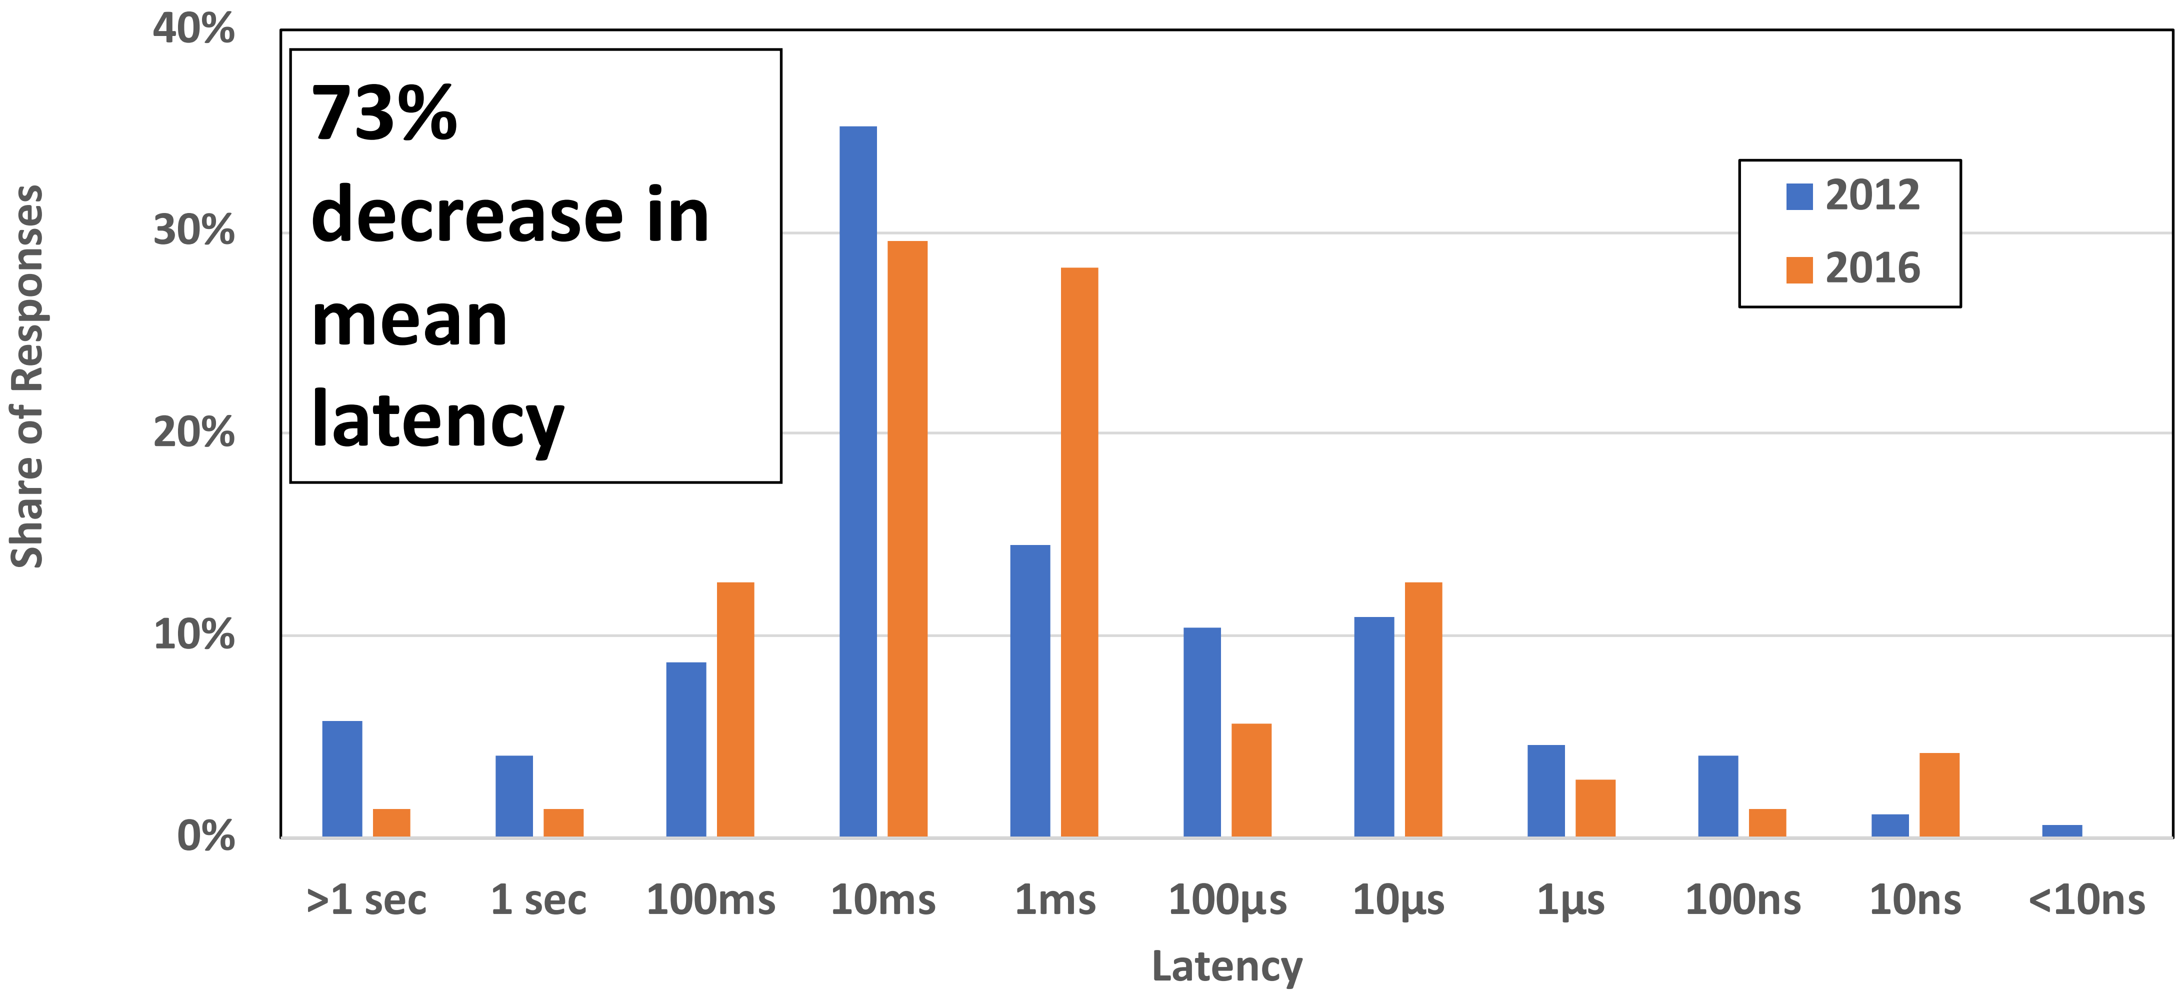

IOPS by Form Factor

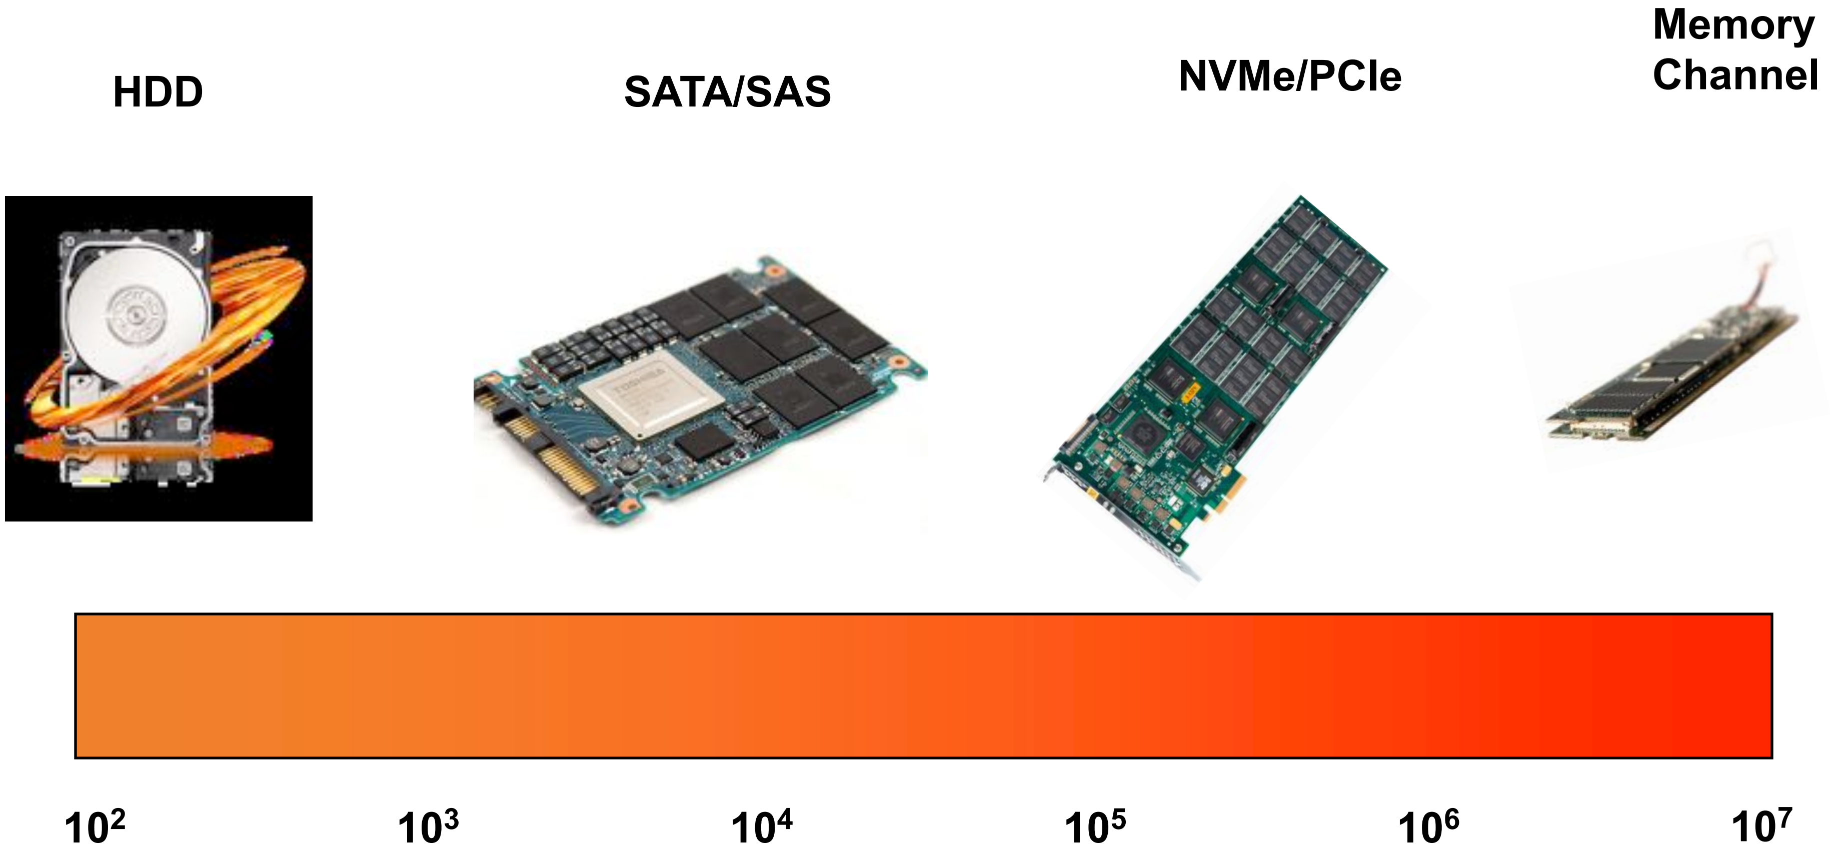

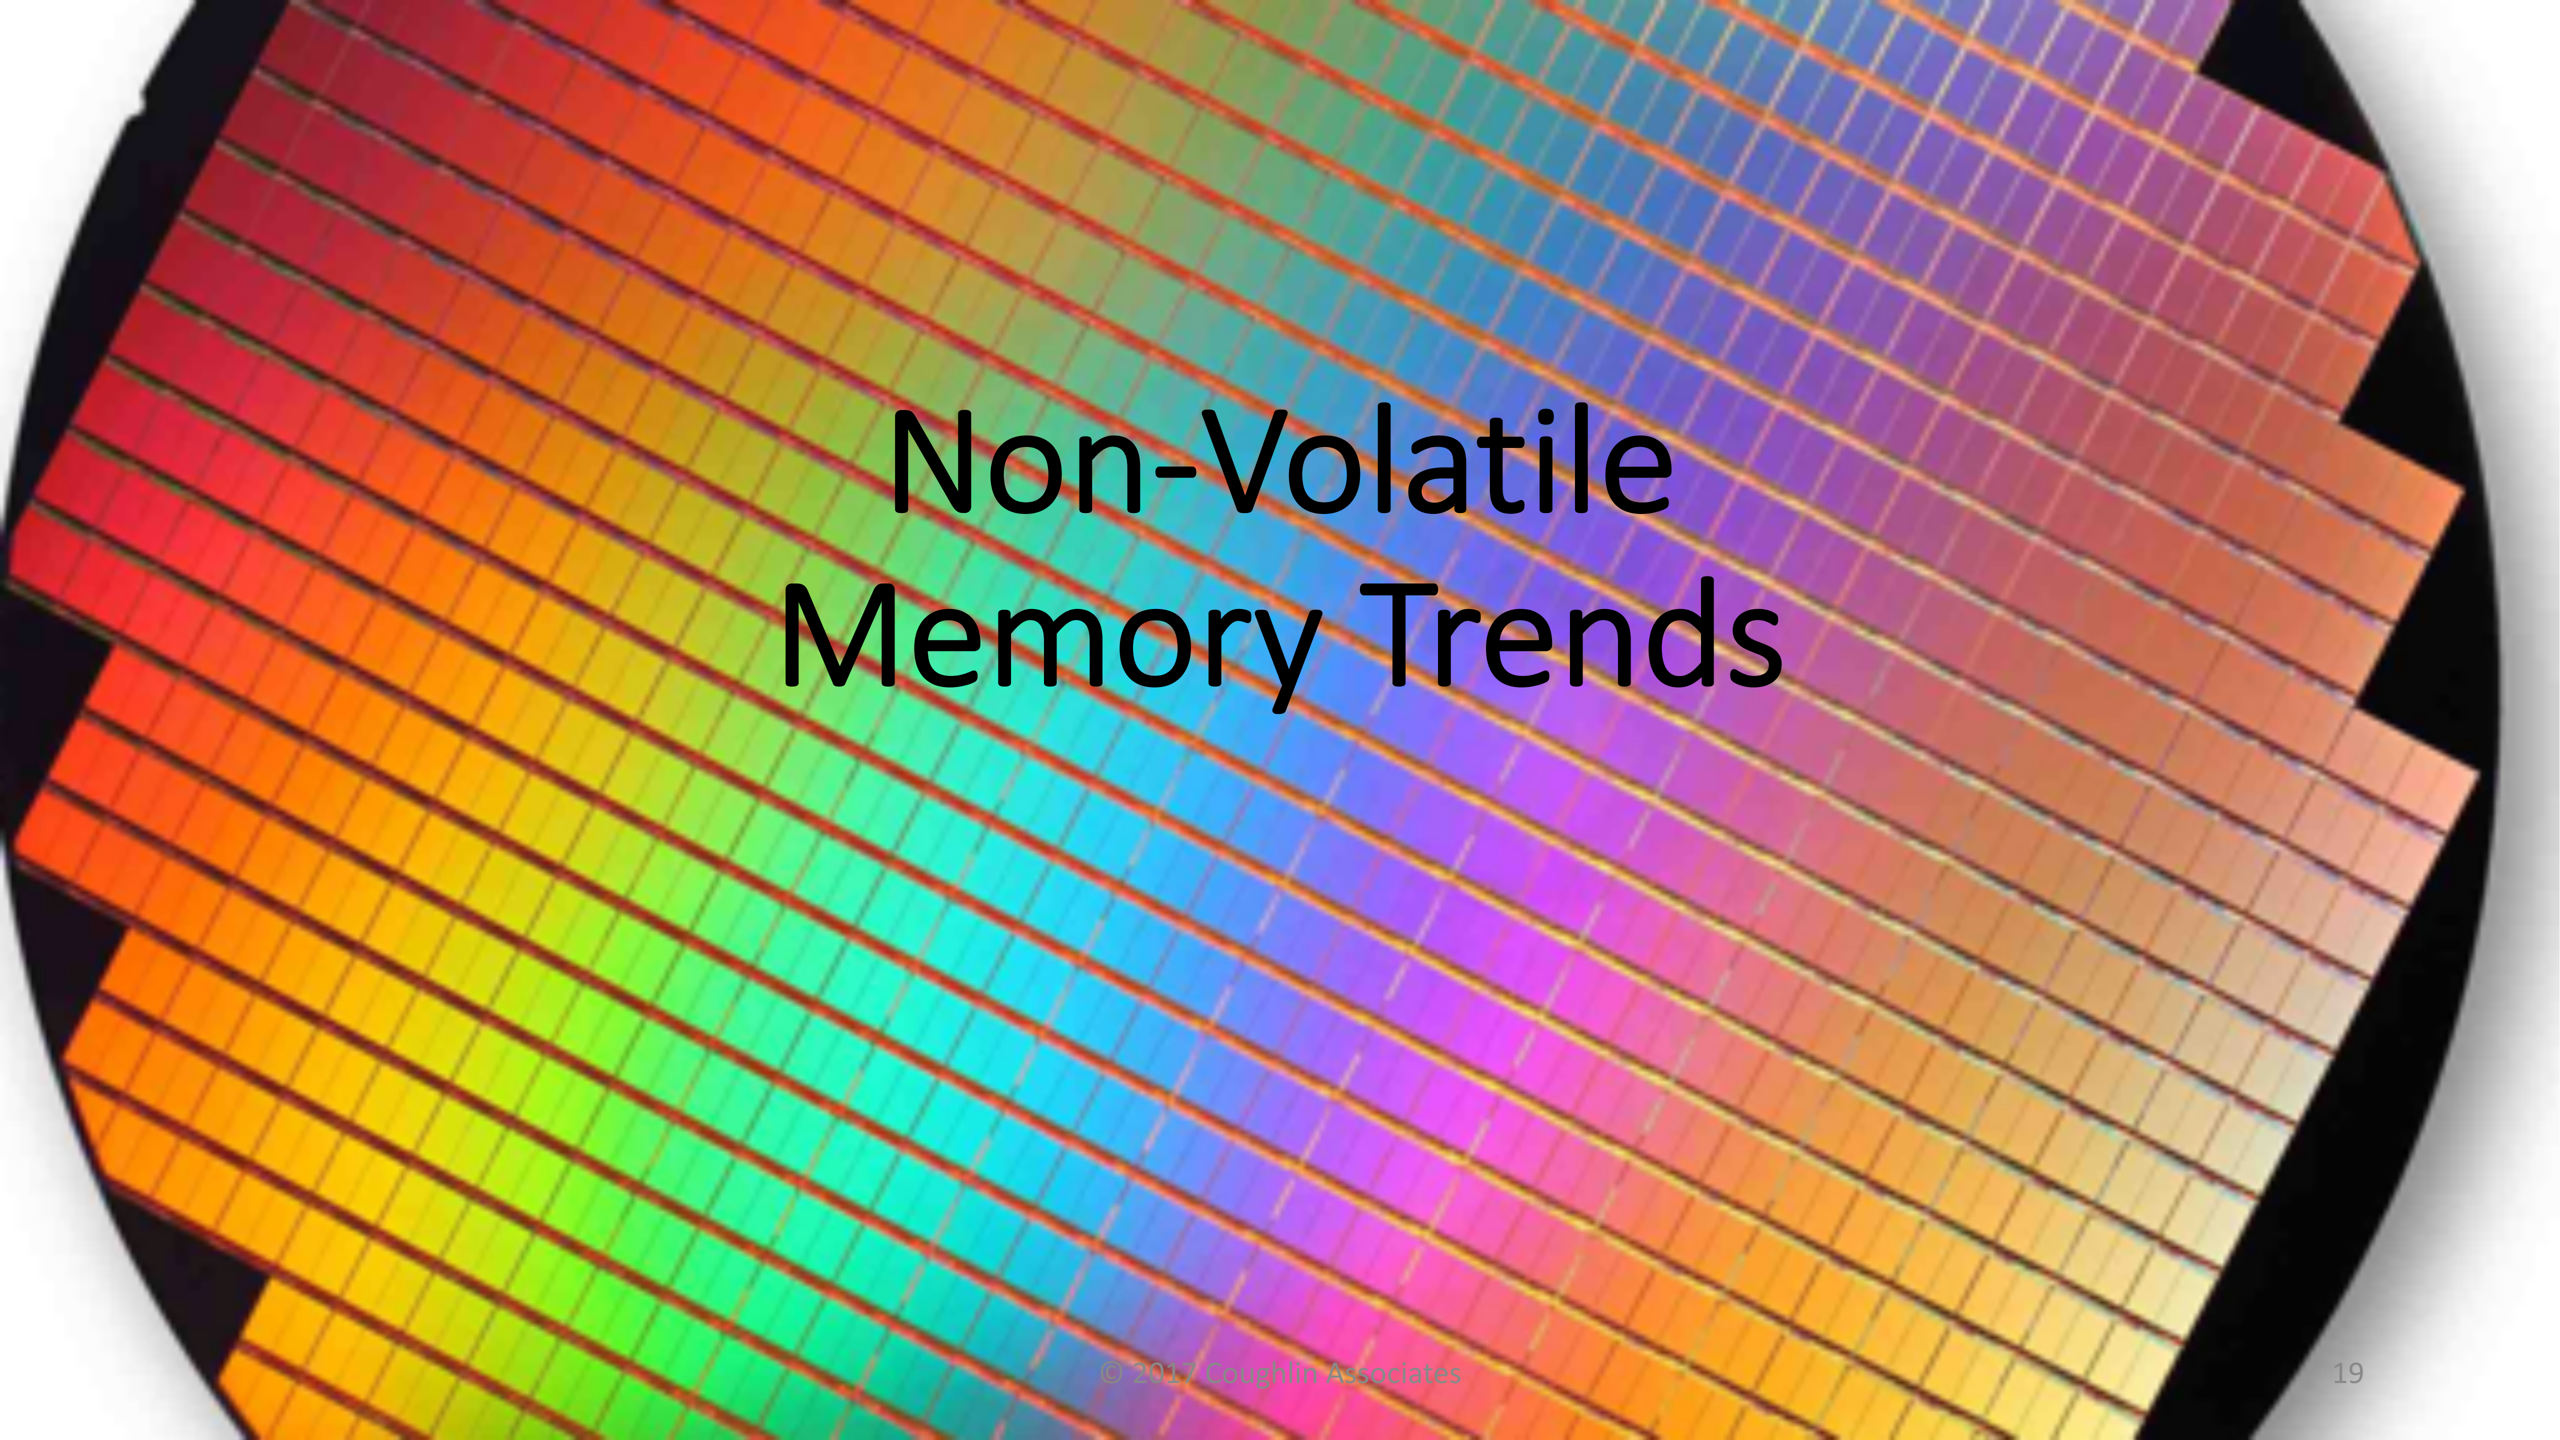

Non-Volatile Memory Trends

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

iNEMI Mass Data Storage Roadmap

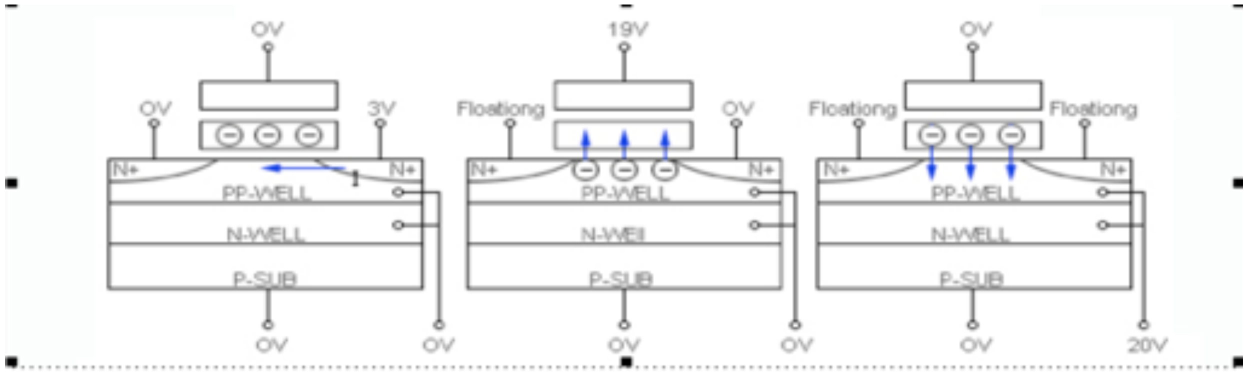

Figure 4. Cross-Section of planar floating-gate flash memory cell

Left: Reading, Middle: Programming, Right: Erasing

(Source: Samsung Semiconductor Company)

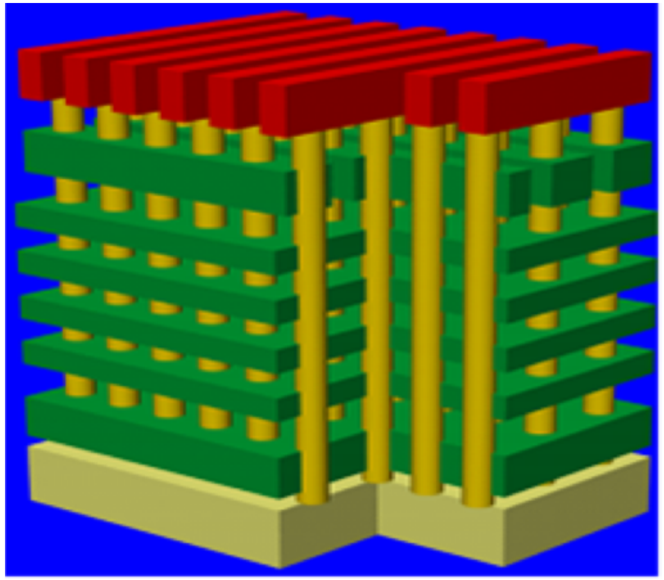

Table 1. ITRS NAND Flash Chip Roadmap

	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 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)

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

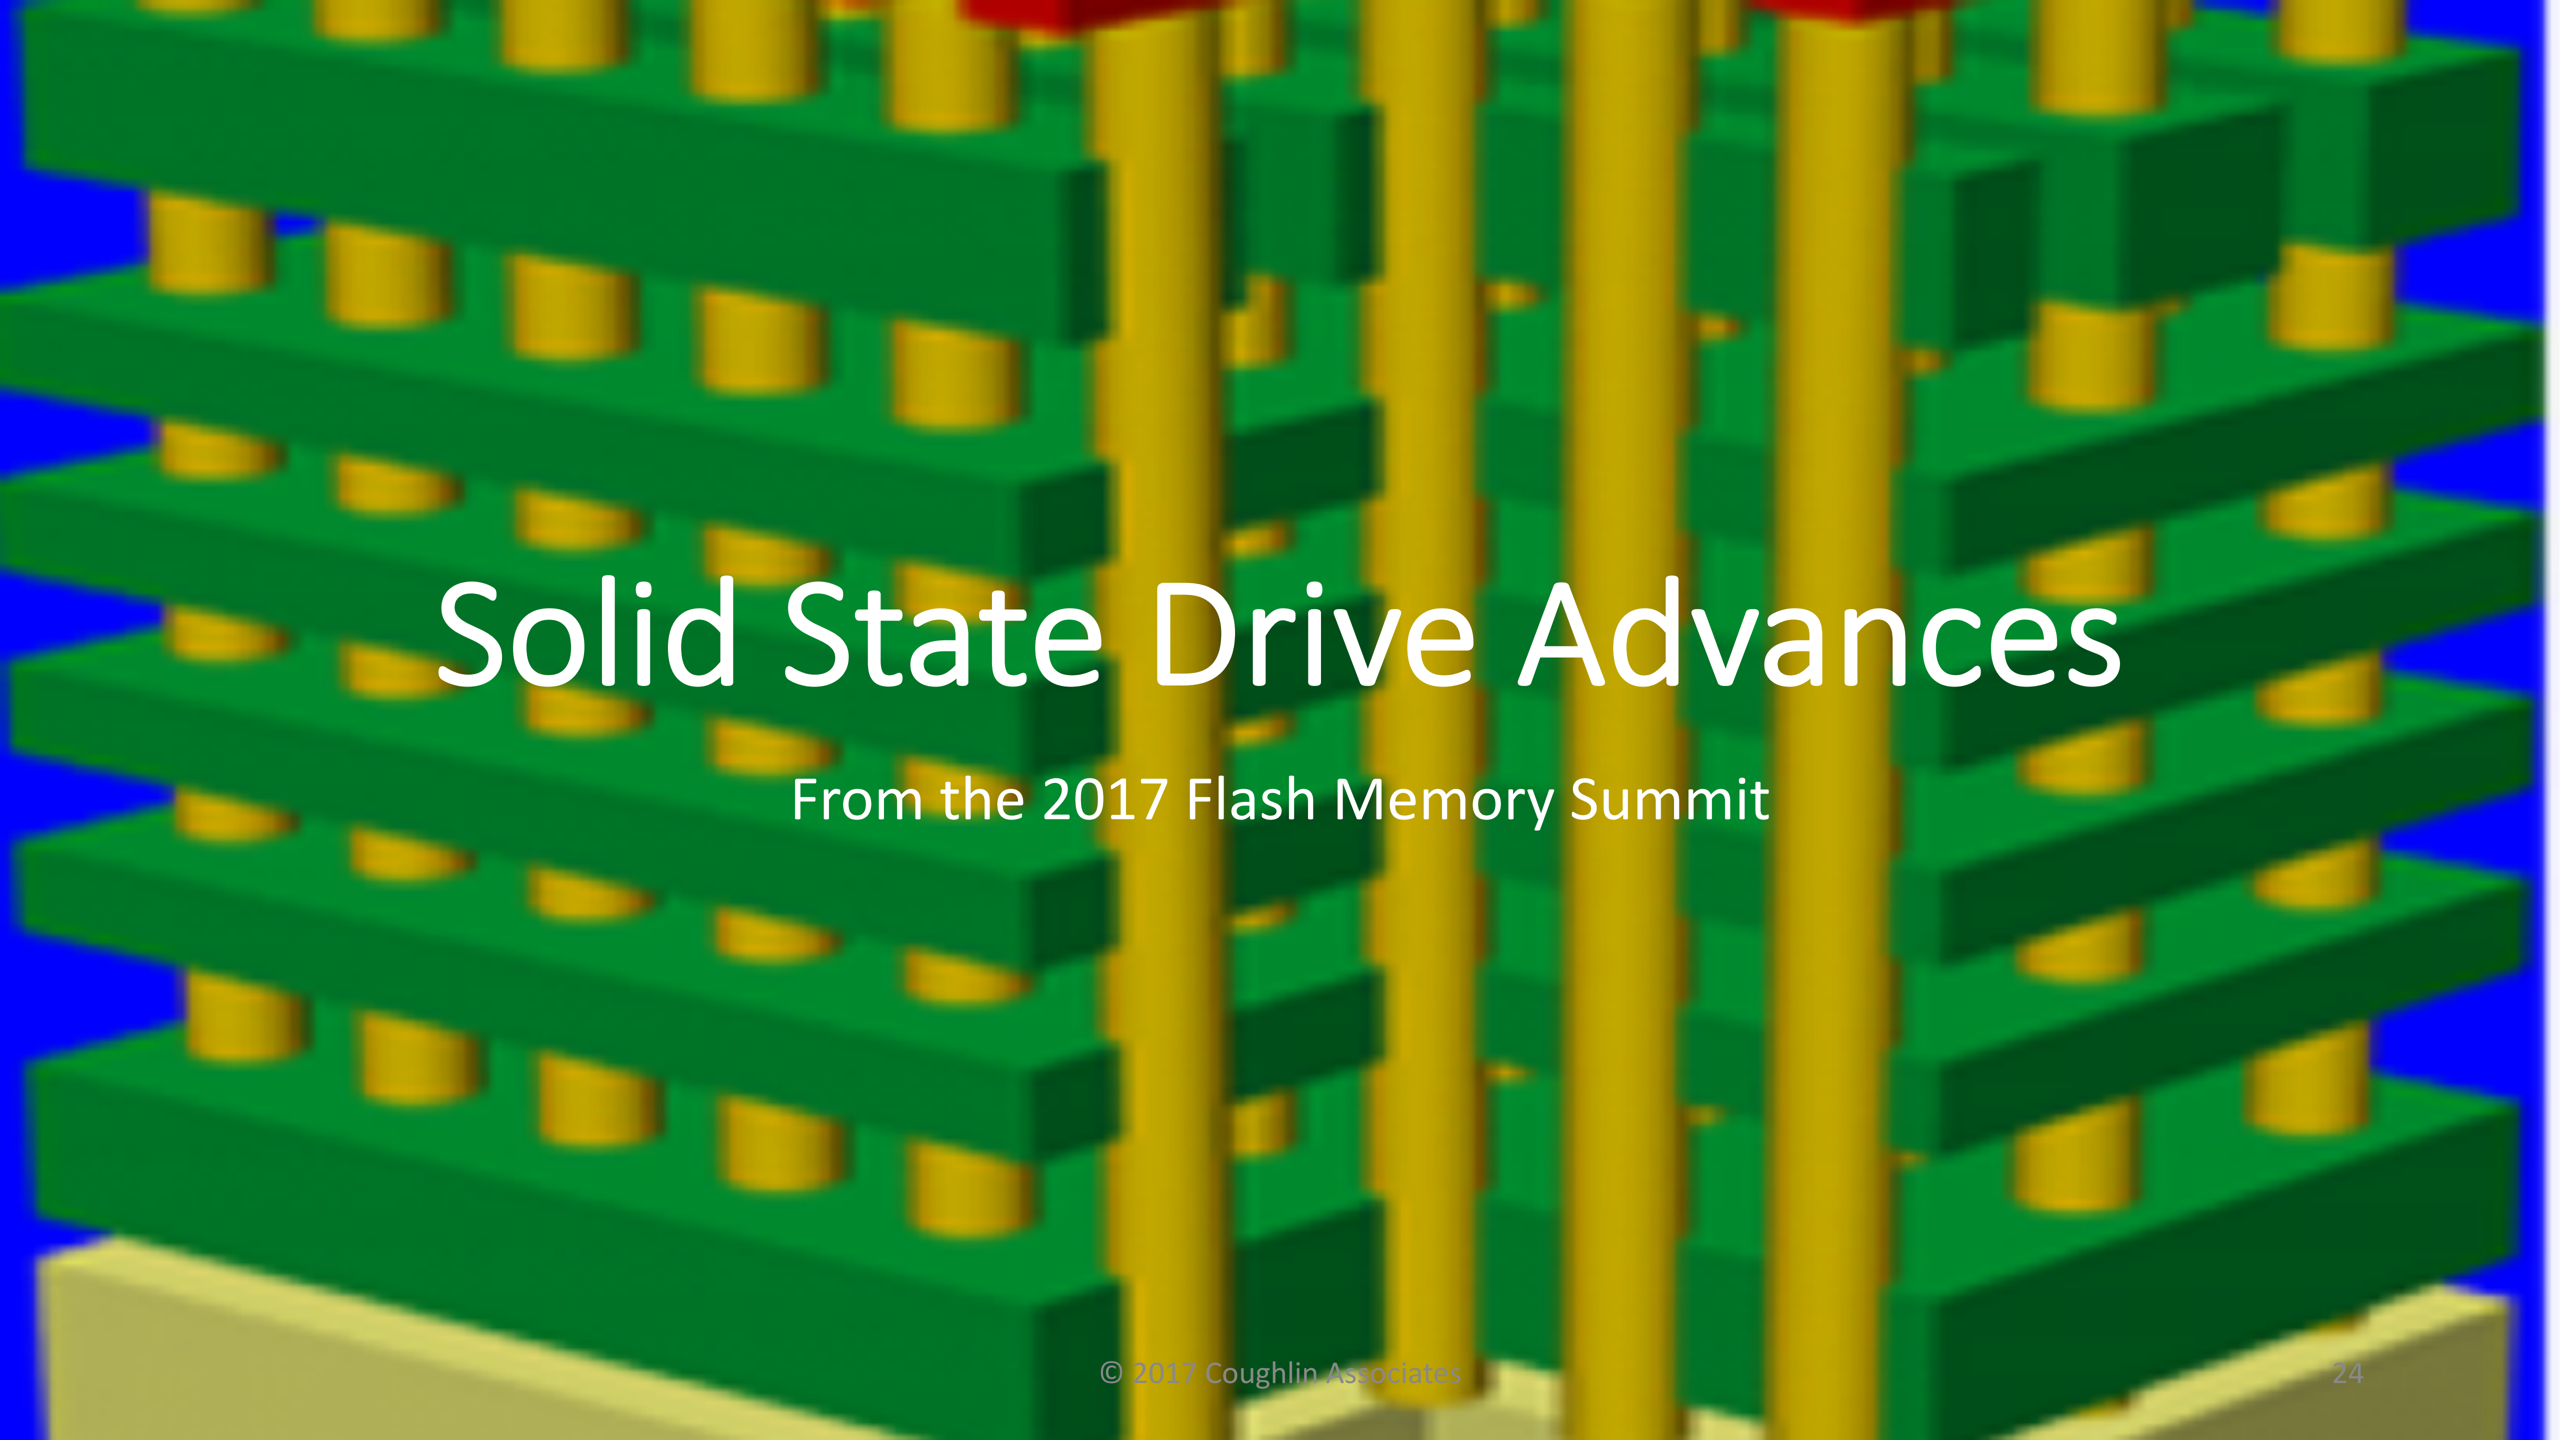

Solid State Drive Advances

From the 2017 Flash Memory Summit

The move to Non-Volatile Storage

- NVMe PCIe-based storage interfaces avoid much of the overhead of hard disk drives—designed for fast solid state storage
- 2D to 3D manufacturing transition is underway with 3D yields improving and achieving cost parity with 2D by 2018
- 3D flash fab investments are over $10 B US per plant-many being build—now at 96 layers
- Shortage of flash memory throughout 2017 and into 2018 is due to yield issues with the 3D flash transition

3D Flash

- Shipments up to 64 layer—common by end of year
- Announcements up to 96 layers—1 Tb per die
- Technology projections of hundreds of layers
- Announced quad-level cells for higher density

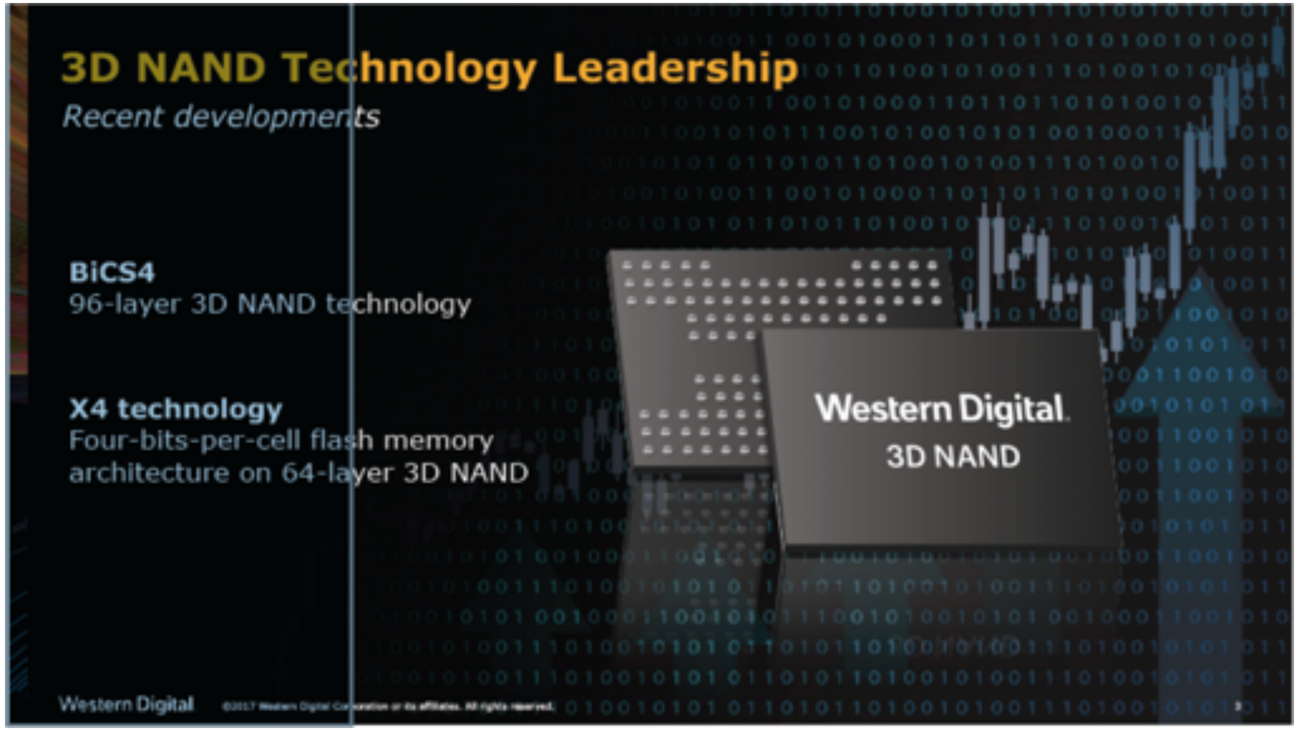

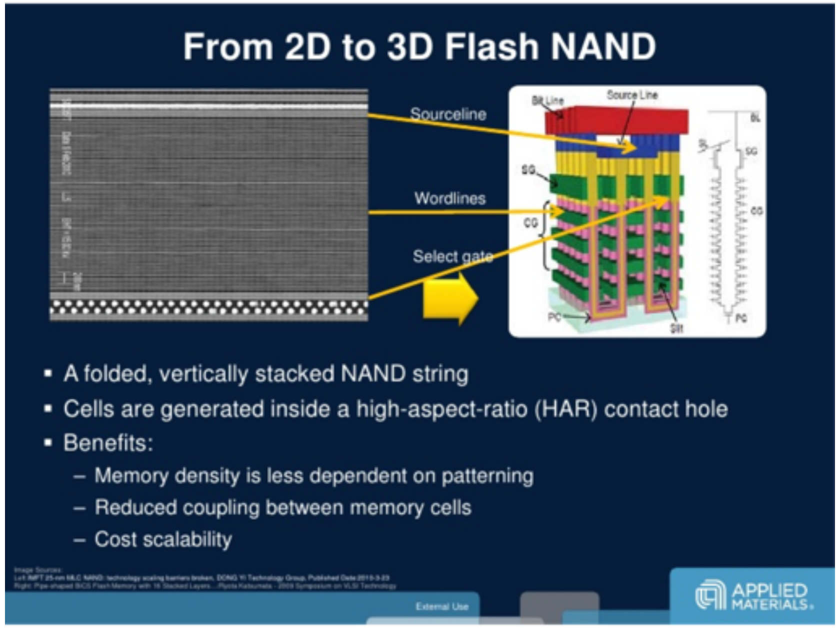

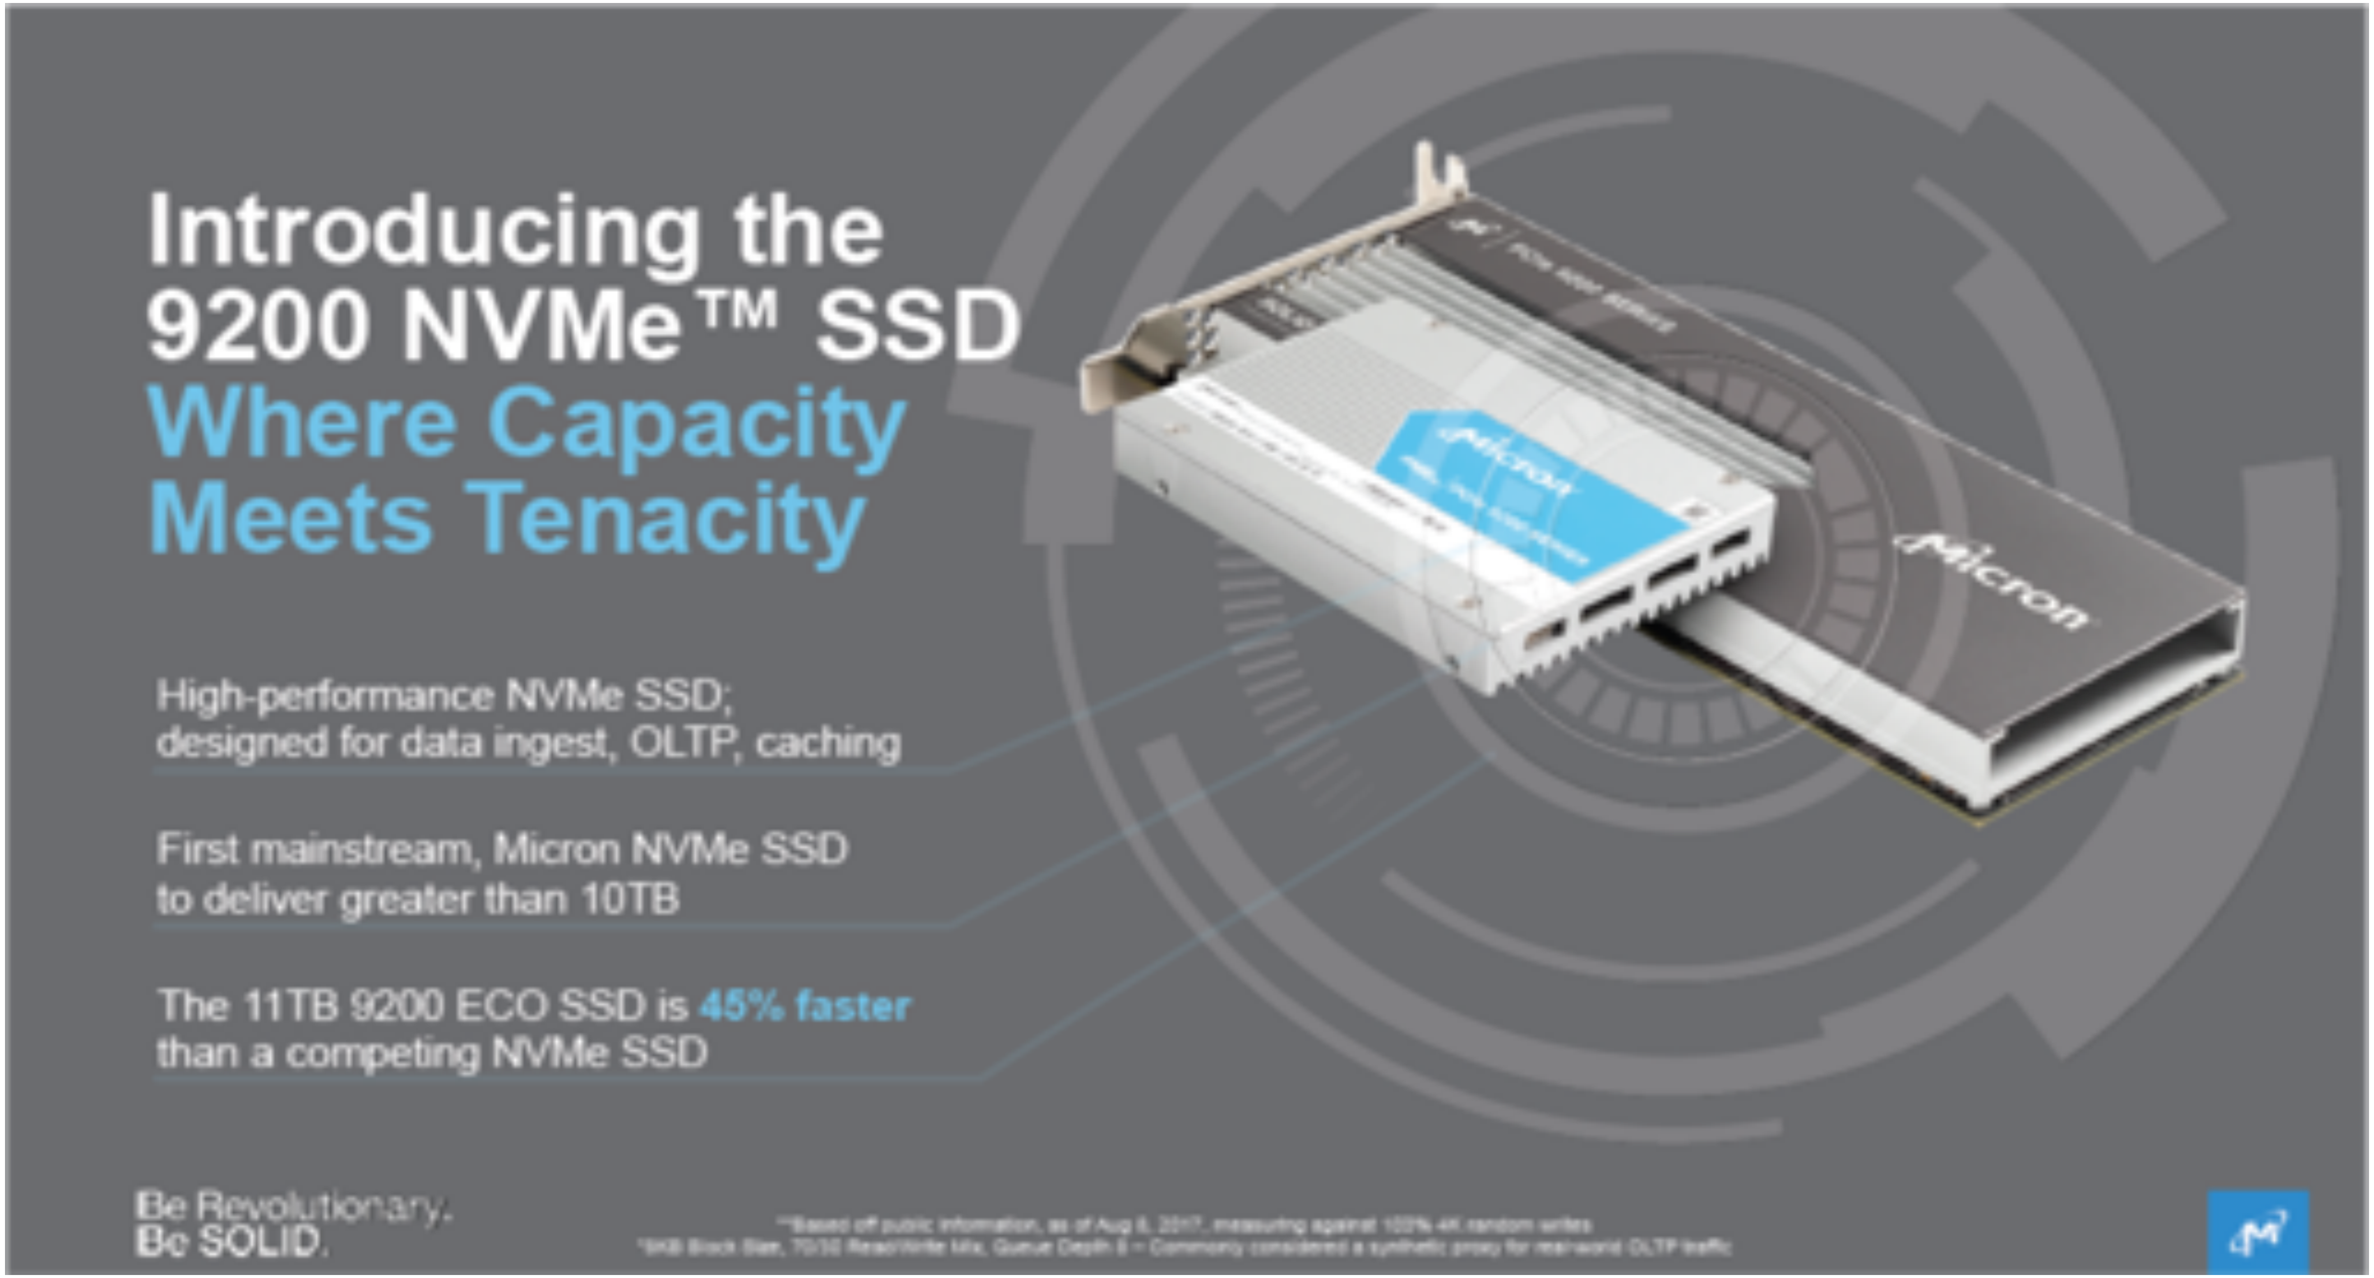
Introducing the
9200 NVMe™ SSD
Where Capacity
Meets Tenacity
High-performance NVMe SSD;
designed for data ingest, OLTP, caching
First mainstream, Micron NVMe SSD
to deliver greater than 10TB
The 11TB 9200 ECO SSD is 45% faster
than a competing NVMe SSD
Be Revolutionary.
Be SOLID.

Intel Dual Port NVMe SSDs

1U Form Factor NVMe SSDs

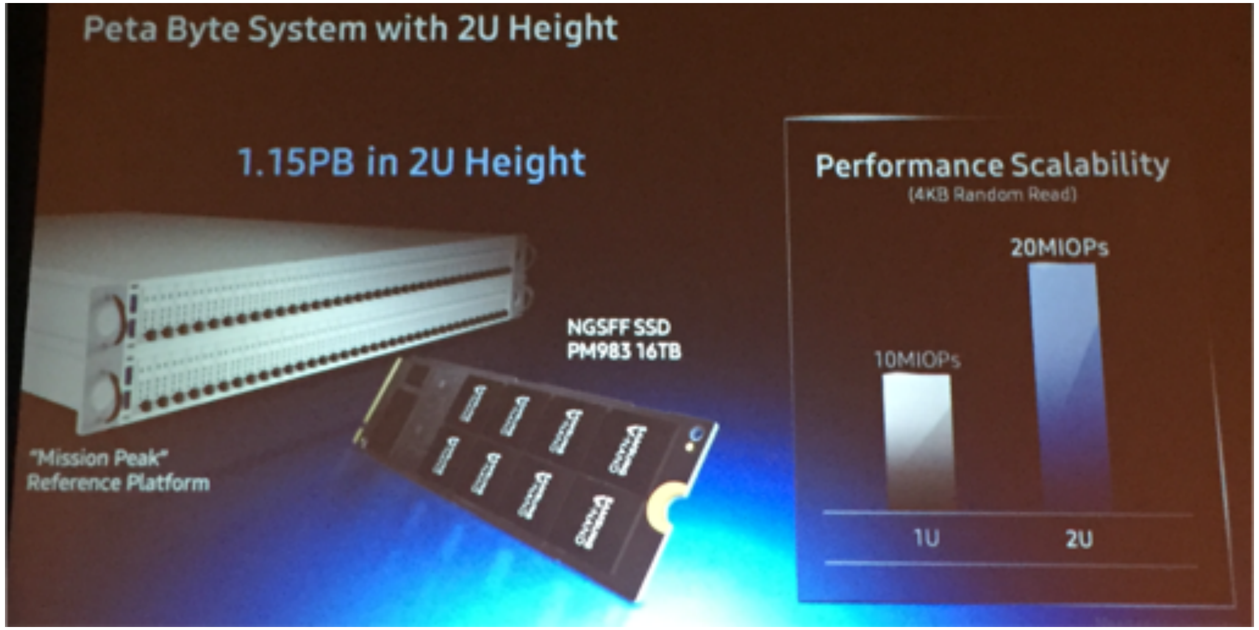

New technologies

Fastest & highest capacity SSD

64TB NVMe SSD
Over 13,000 MB/s of throughput

Seagate Nytro Q-Boost™ features

IO Determinism — Nytro Q-Boost Demo →

7x reduced average read latency

16x reduced max read latency

Multi-Stream — Nytro Q-Boost Demo →

4x Improved WAF & NAND endurance

3x Faster performance

longsys
The world's smallest SSD
60GB~480GB
FORESEE
PCIe BGA SSD
11.5mm
13mm

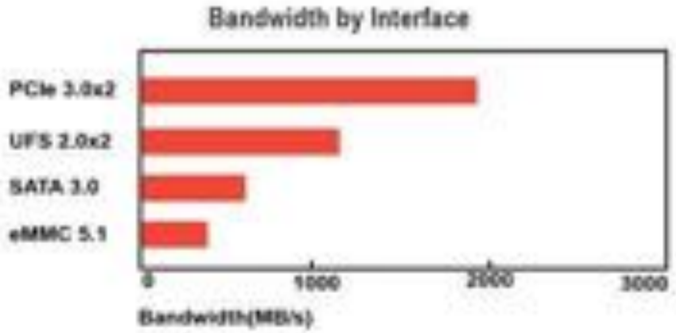

Bandwidth by Interface
PCIe 3.0x2
UFS 2.0x2
SATA 3.0
eMMC 5.1
0
1000
2000
3000
Bandwidth(MB/s)

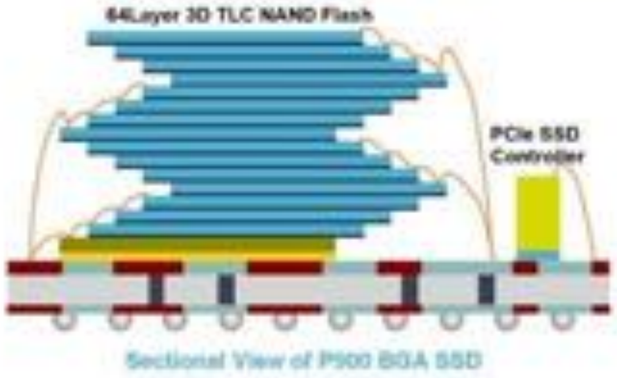

64Layer 3D TLC NAND Flash
PCIe SSD Controller
Sectional View of P900 BGA SSD

Samsung's 128 TB SSD

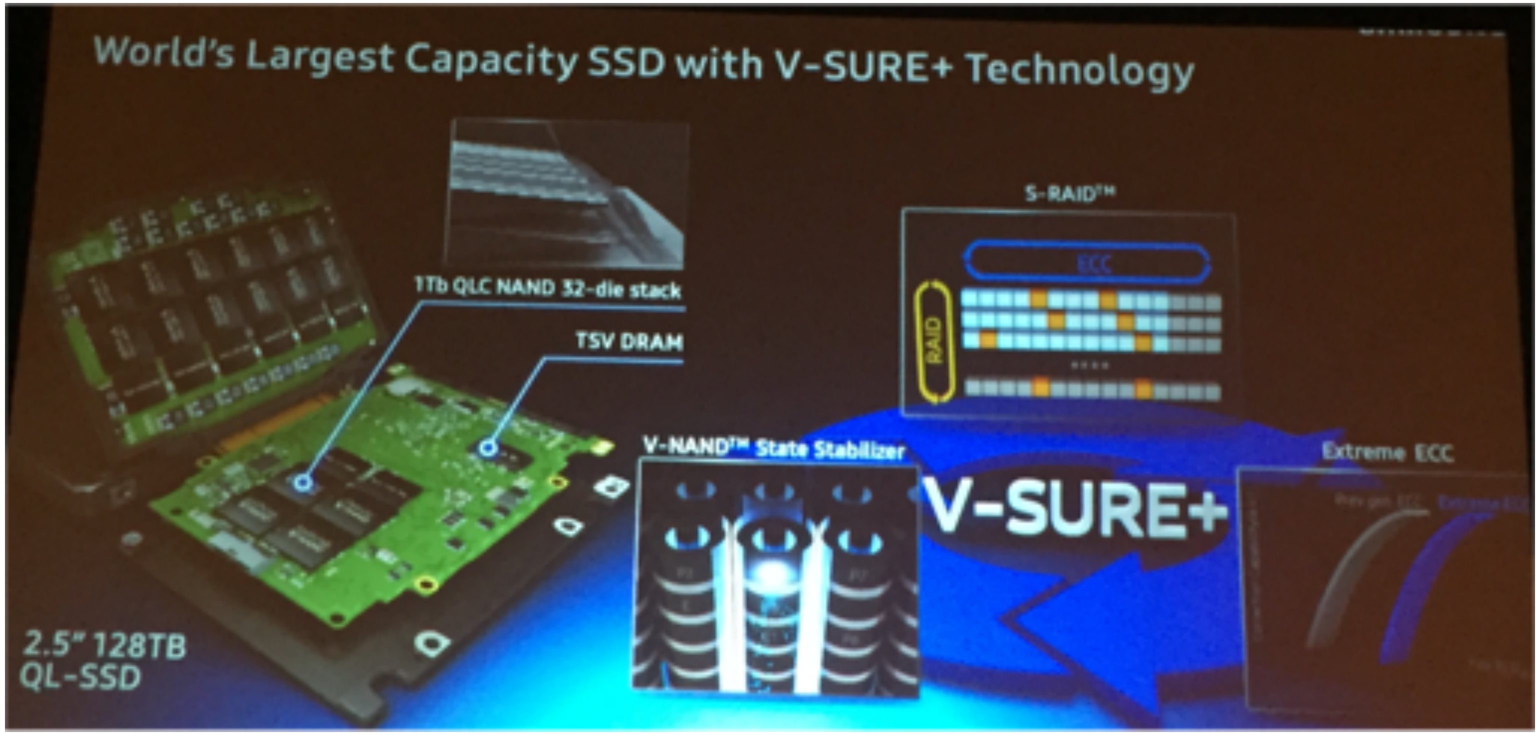

- Samsung's latest V-NAND chip is a a 1Tb V-NAND chip, available next year.
- This will enable 2TB of memory in a single V-NAND package by stacking 16 1Tb dies
- Using 32-die stack of 1 Tb QLC NAND the company was showing a 2.5" form factor 128 TB SSD.

Samsung Z-NAND

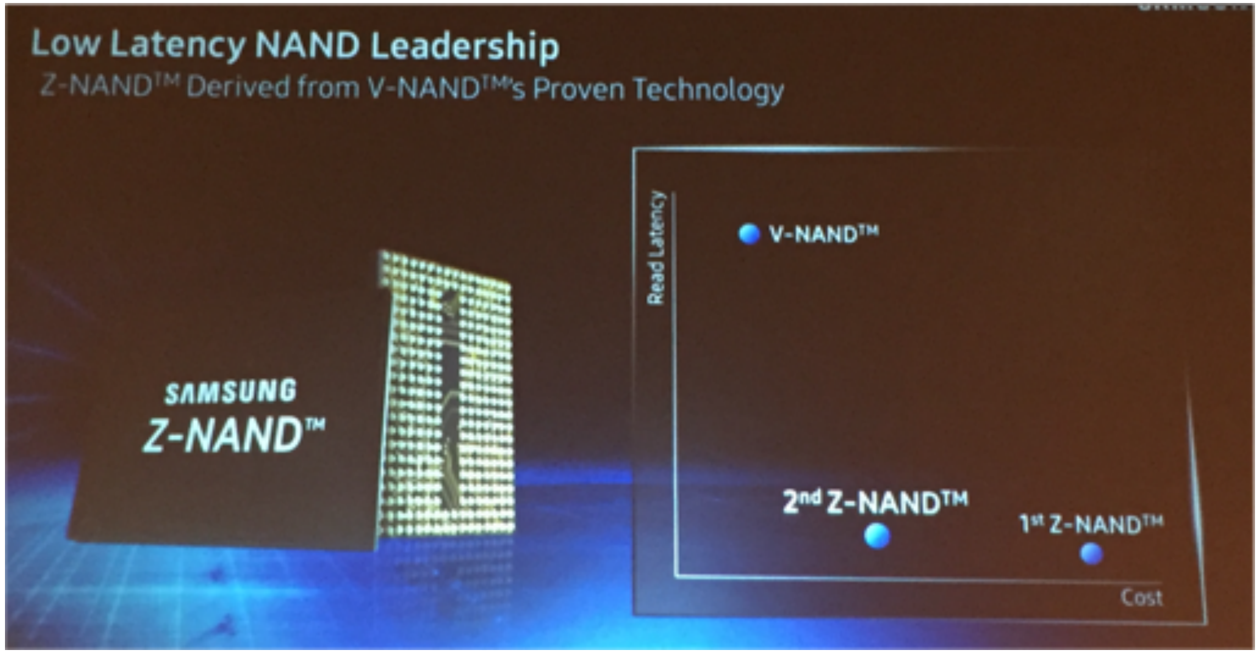

- Samsung introduced its Z-SSD technology at FMS 2016.
- In 2017, its first product, the SZ98 Z-SSD (15 microseconds of read latency)
- The Z-SSD is intended for data centers and enterprise systems dealing with extremely large, data-intensive tasks
- Samsung says that at the application level, the use of Samsung's Z-SSDs can reduce system response time by up to 12 times, compared to using NVMe SSDs.

Key Value SSD

- Samsung introduced a Key Value SSD. KVS is a way to organize data in storage for more rapid access that is commonly used in object storage systems.
- Samsung’s Key Value assigns a ‘key’ or specific location to each ”value,” or piece of object data – regardless of its size, enabling direct addressing of a data location
- Samsung’s Key Value technology enables SSDs to scale-up (vertically) and scale-out (horizontally) in performance and capacity.

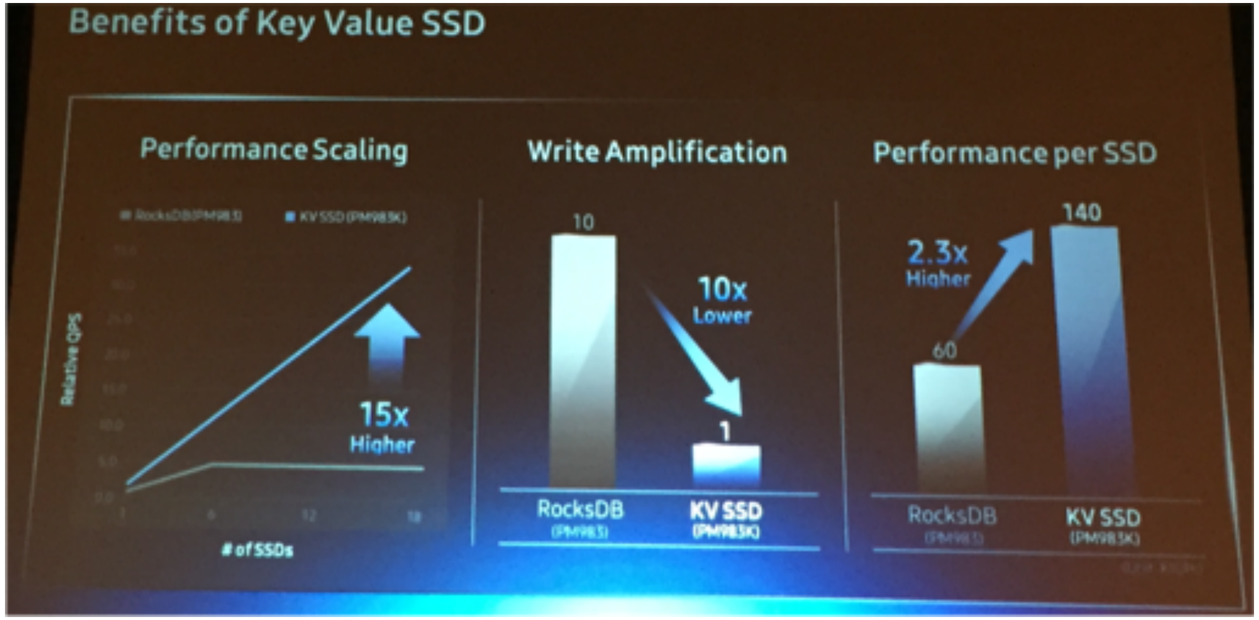

When data is read or written, a Key Value SSD can reduce redundant steps, which leads to faster data inputs and outputs, as well as increasing TCO and significantly extending the life of an SSD

More Generations of V-NAND

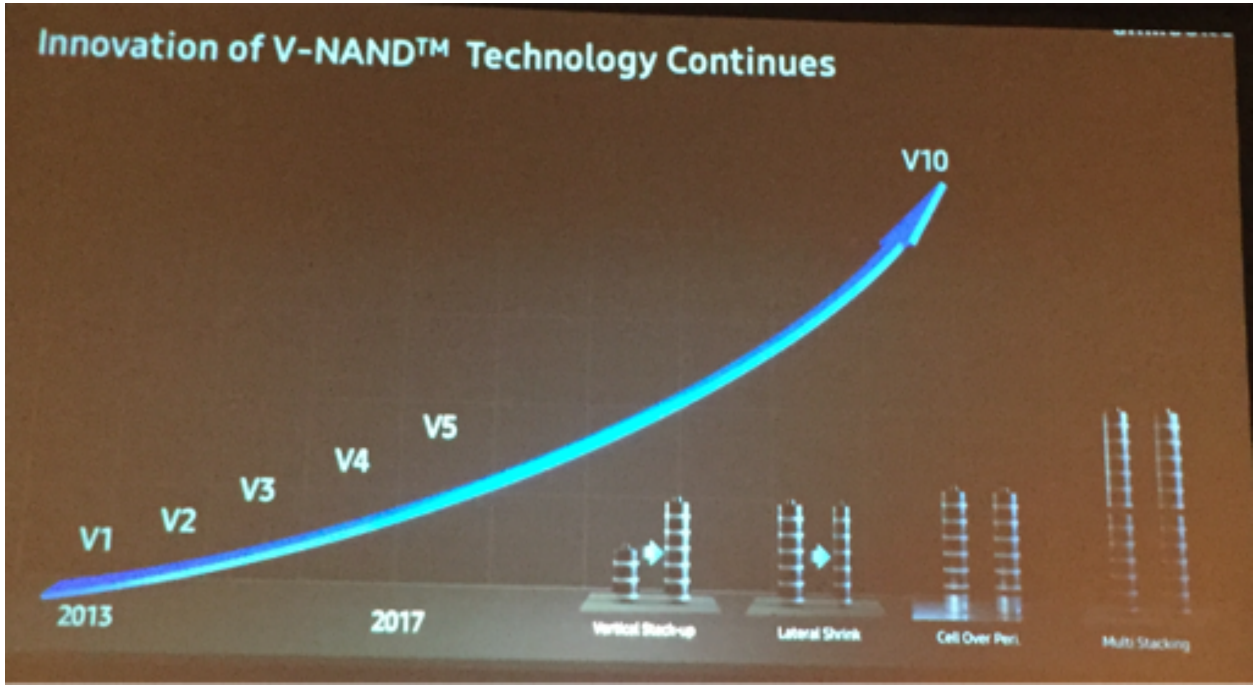

- Samsung was presented their V-NAND roadmap showing--at least 5 more generations of V-NAND beyond their announced 5th generation.
- These higher capacity products will leverage three technological developments.
- These include continued vertical stacks (more 3D layers), lateral shrink (smaller features) and cell over peripheral (building the cell layers over the supporting CMOS logic.

Storage Fabrics and Memory Centric Computing

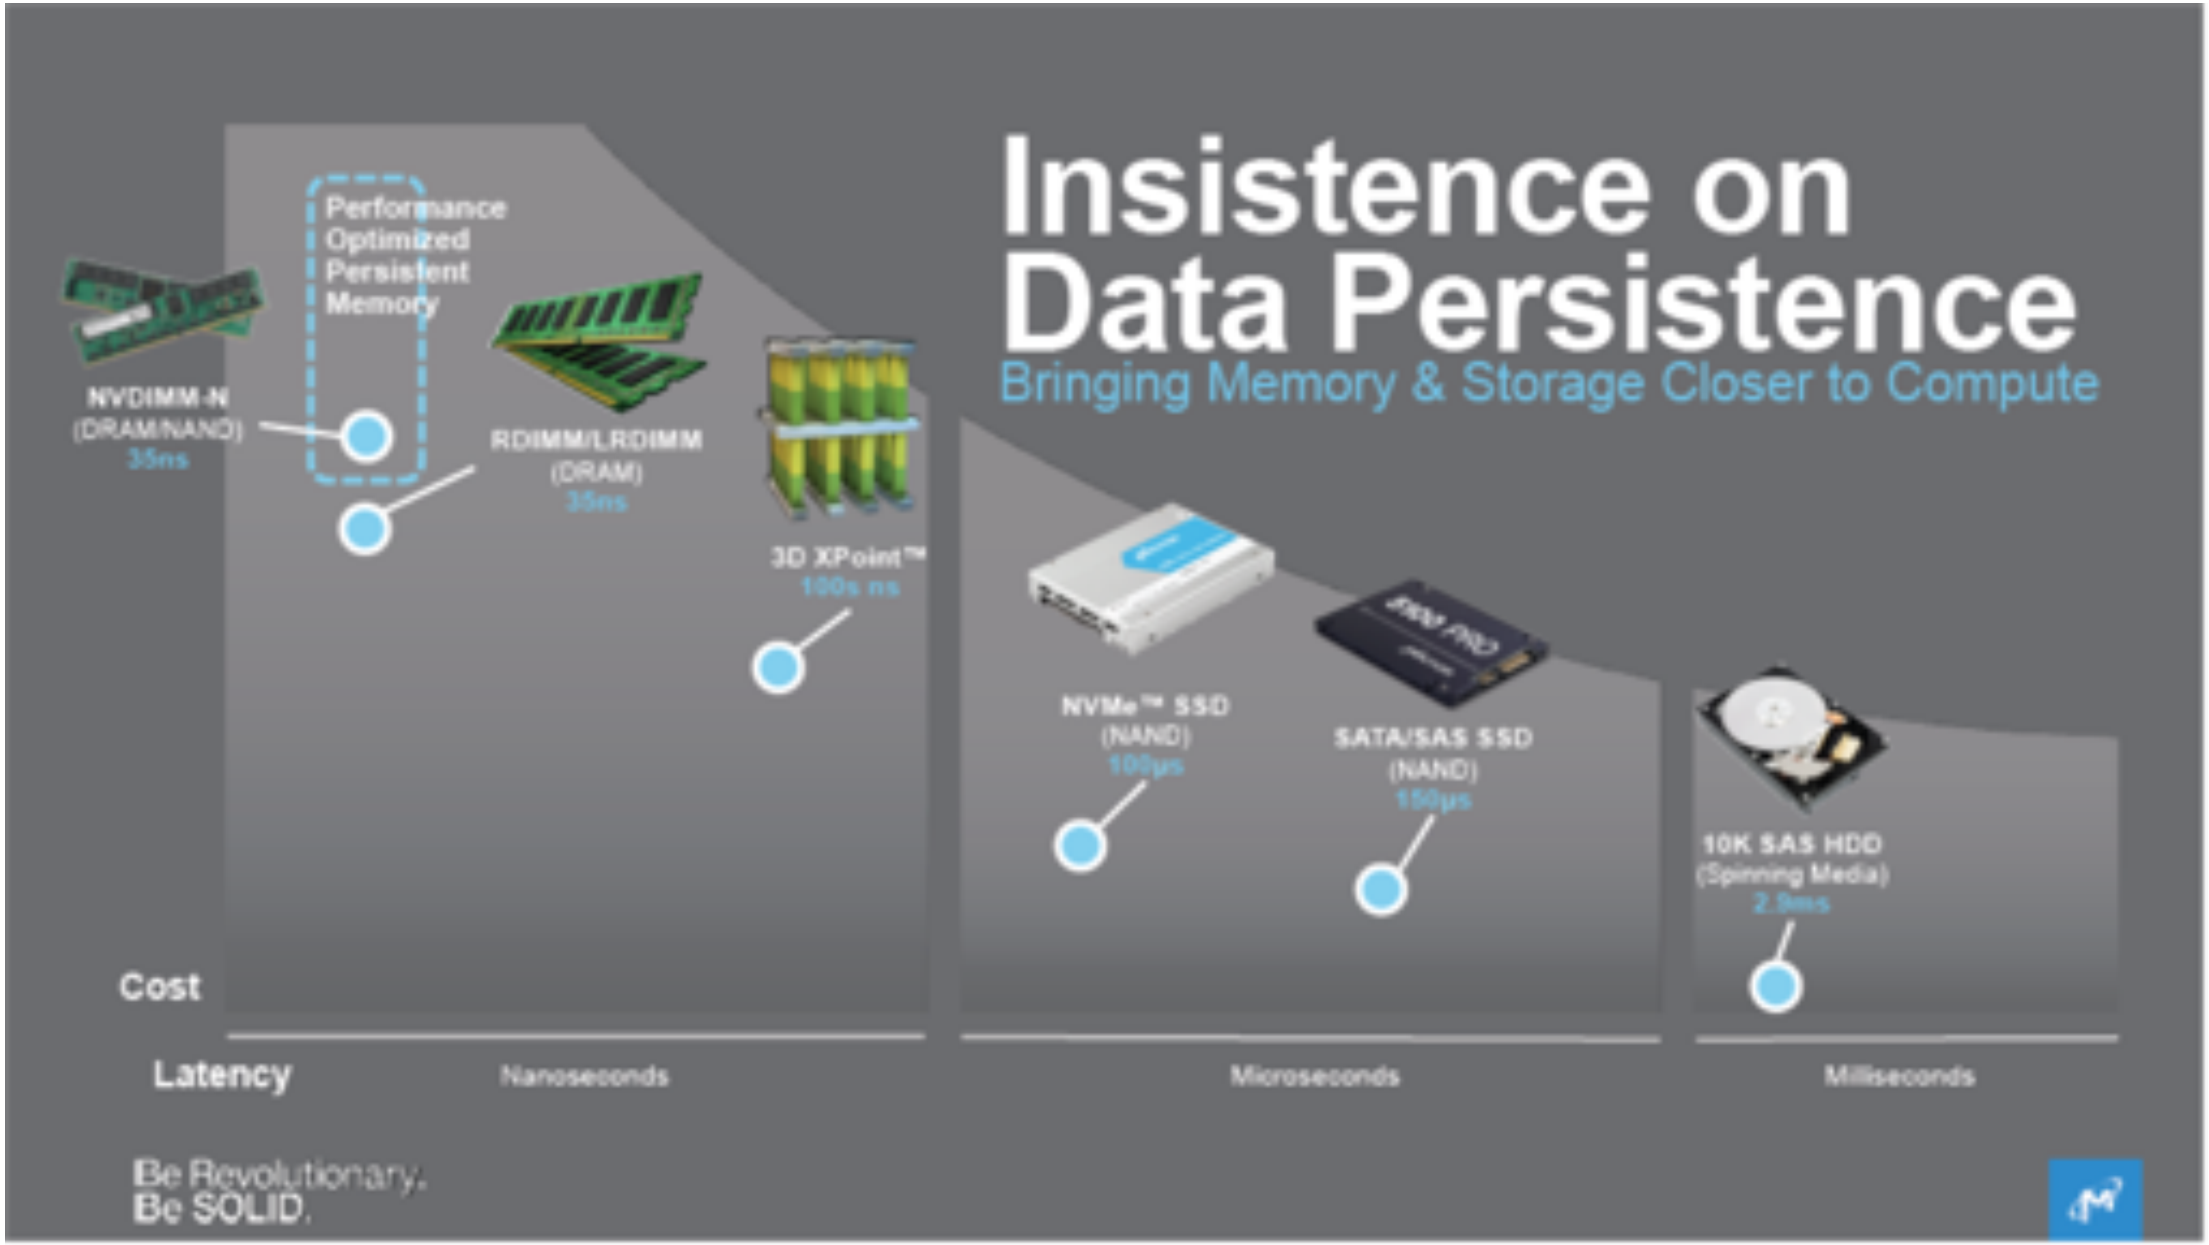

Insistence on
Data Persistence
Bringing Memory & Storage Closer to Compute
Performance
Optimized
Persistent
Memory
NVDIMM-N
(DRAM/NAND)
35ns
RDIMM/LRDIMM
(DRAM)
35ns
3D XPoint™
100s ns
NVMe™ SSD
(NAND)
100μs
SATA/SAS SSD
(NAND)
150μs
10K SAS HDD
(Spinning Media)
2.9ms
Cost
Latency
Nanoseconds
Microseconds
Milliseconds
Be Revolutionary.
Be SOLID.

NVMe over Fabric

- Mellanox announced their BlueField SoC for accelerating NVMe over Fabrics (200 Gb/s of throughput and more than 10 million IOPS in a single SoC device)
- A French company, Kalray said that it has released a high-performance NVMe-oF target controller for enabling NVMe-based (JBOF) array boxes.

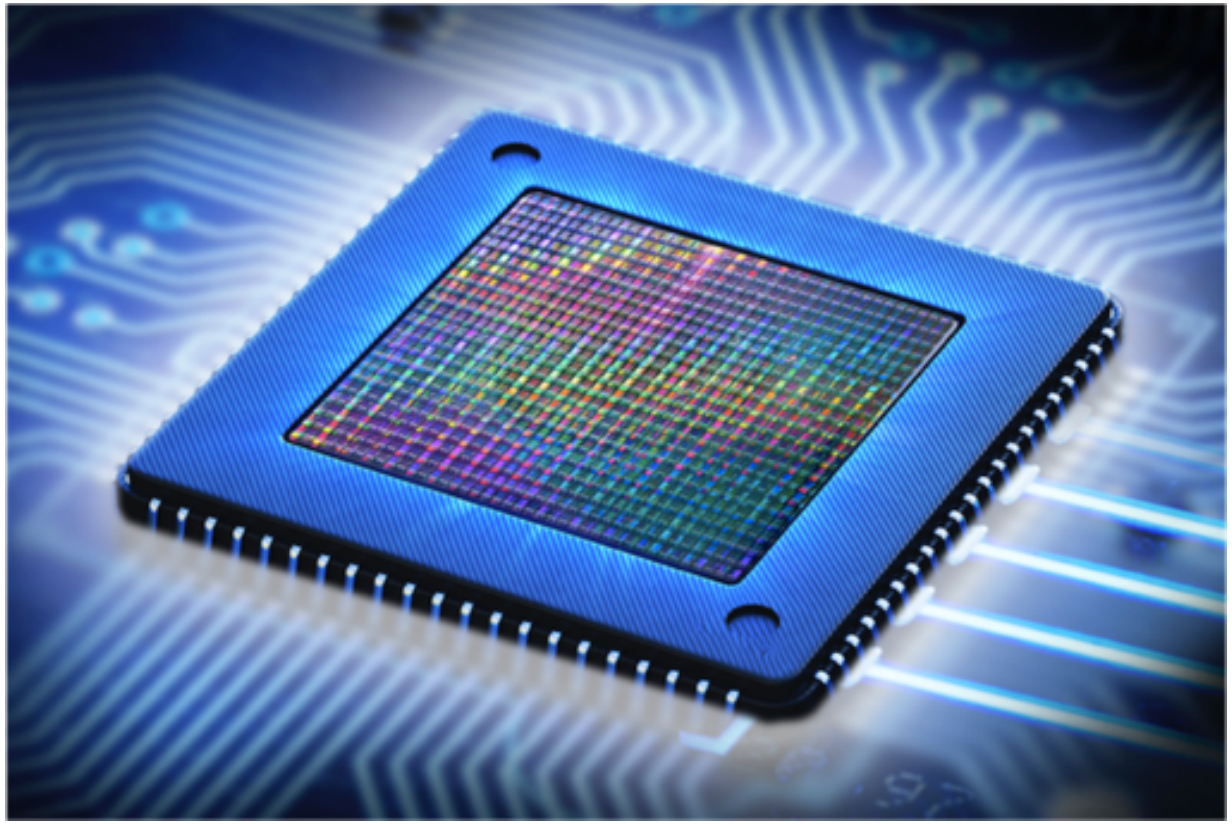

MRAM and PRAM

- MRAM
  - Everspin shipped over 70 M MRAM Chips. Company has partnership with Global Foundries, who is building 300 mm wafers and targeting embedded memory applications
  - Samsung--plans to ship STT MRAM product samples by 2018.
  - Seagate was showing an Everspin MRAM boot SSD at the 2017 FMS
- PRAM
  - Intel says their Optane NVMe products will ship this year.
  - Micron planning to introduce DIMM-based 3D XPoint product

PM Needs New Software

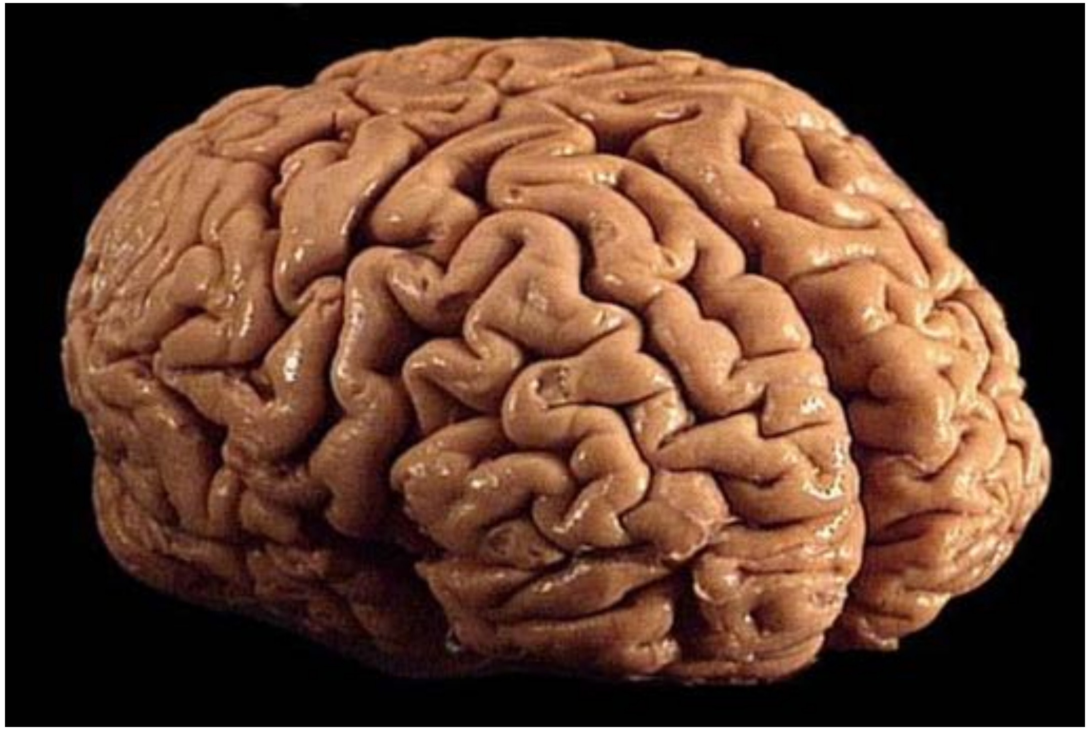

- Move away from “Storage vs. Memory” approach
  - Store at the byte level, not blocks
  - Avoid the storage stack
  - Avoid things like flash translation

The SNIA Persistent Memory Programming Model

- https://www.snia.org/PM

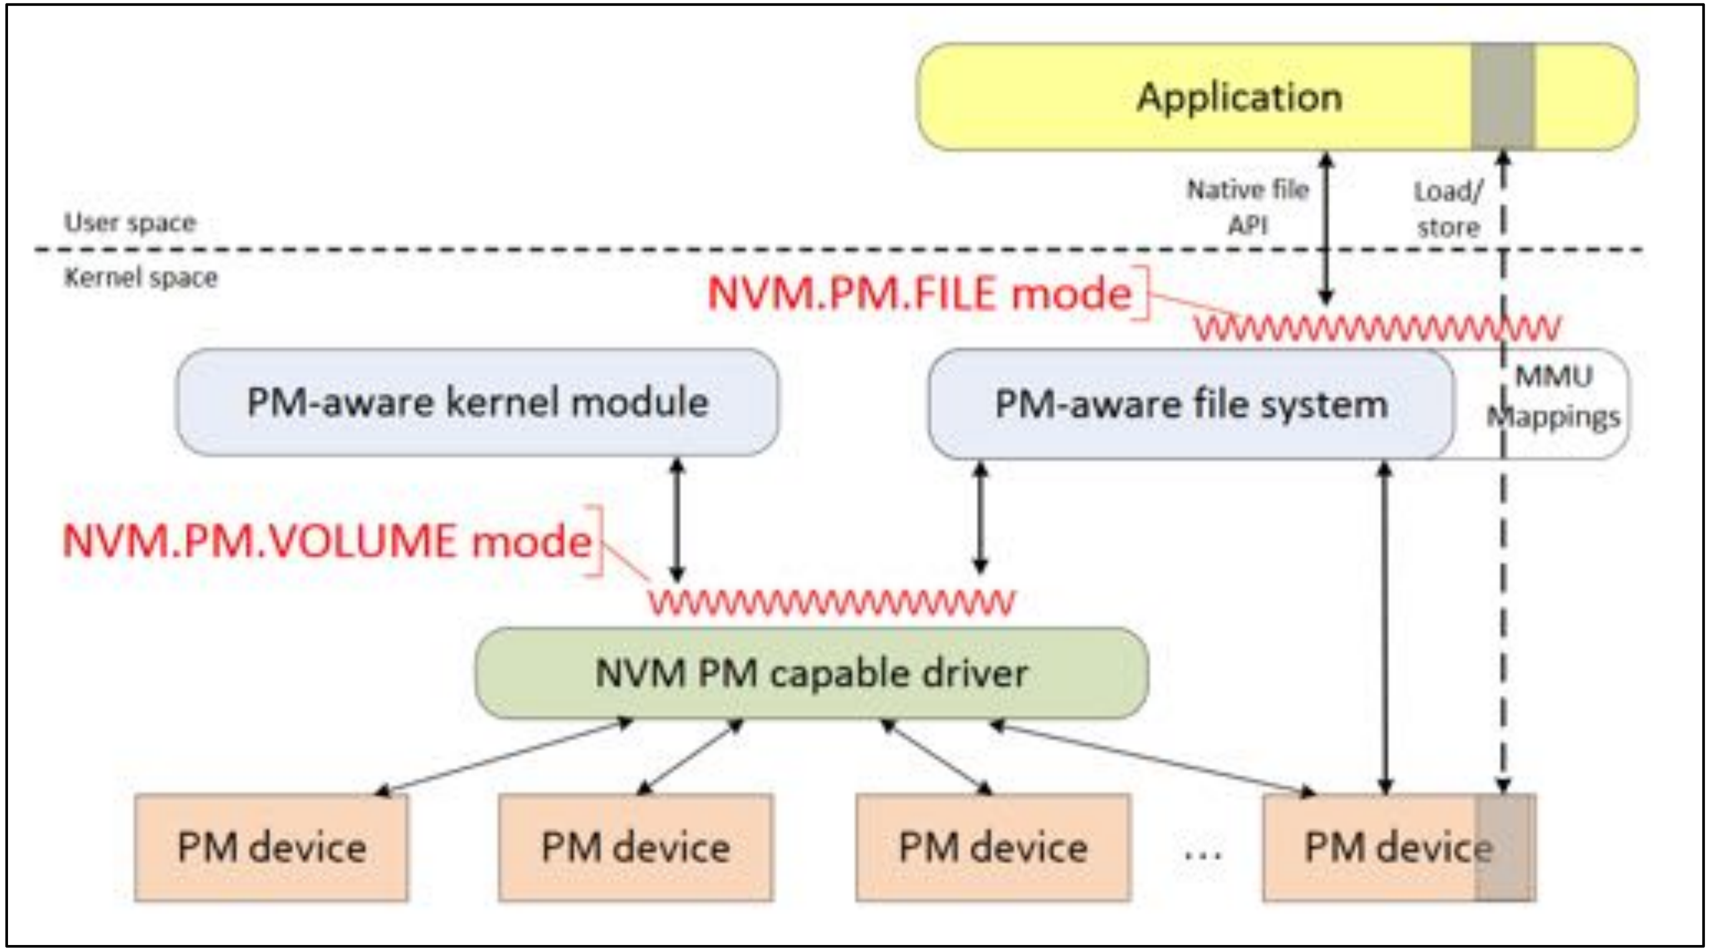

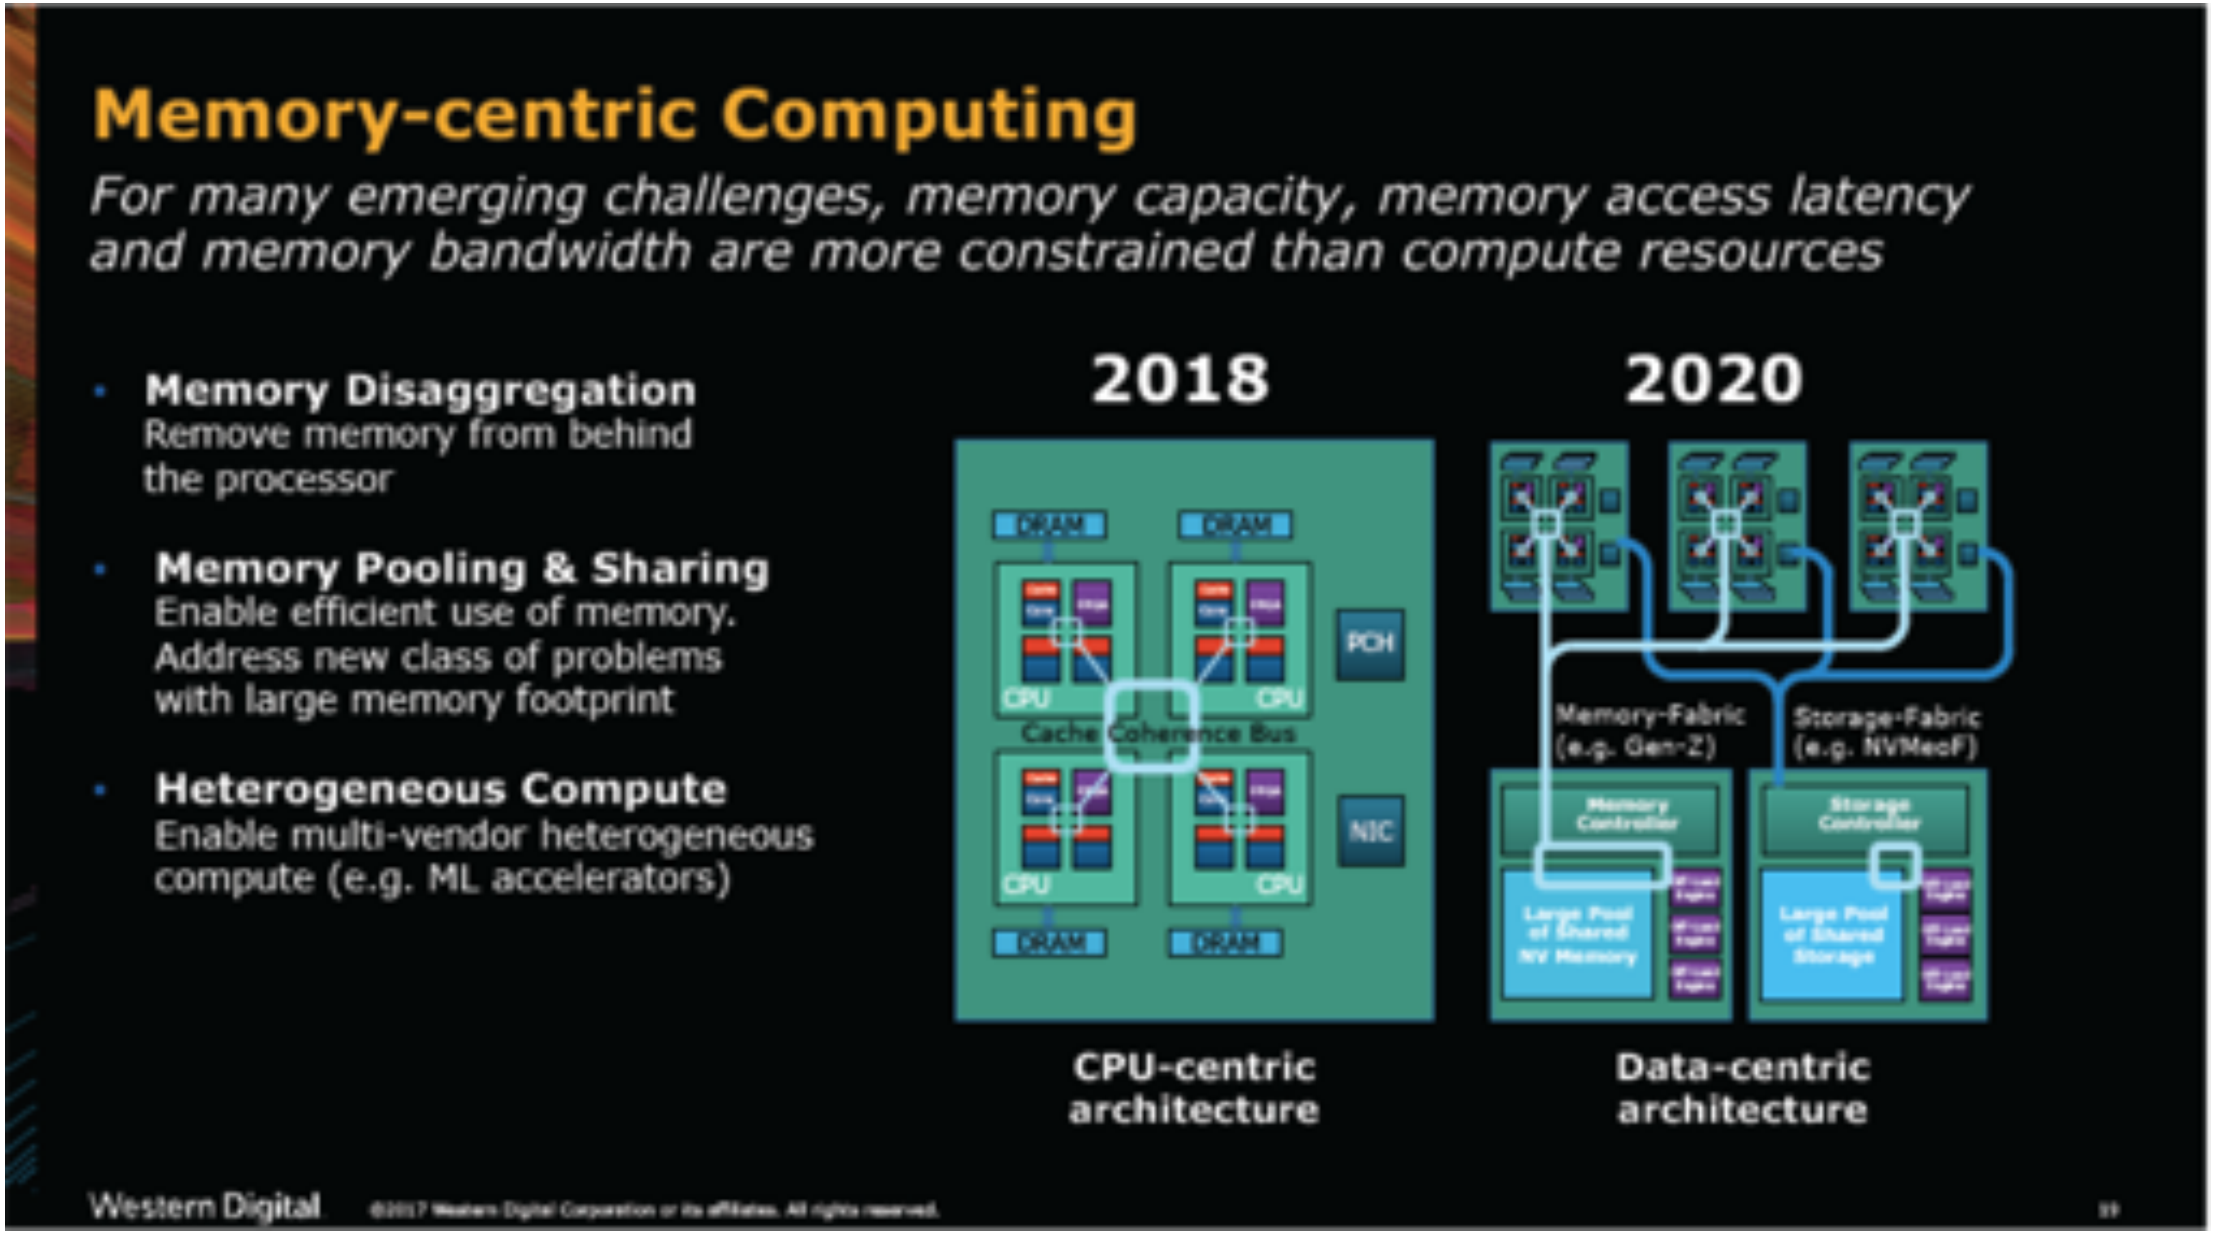

Memory-centric Computing
For many emerging challenges, memory capacity, memory access latency and memory bandwidth are more constrained than compute resources
Memory Disaggregation
Remove memory from behind the processor
Memory Pooling & Sharing
Enable efficient use of memory. Address new class of problems with large memory footprint
Heterogeneous Compute
Enable multi-vendor heterogeneous compute (e.g. ML accelerators)
2018
2020
DRAM
DRAM
CPU
CPU
PCH
Cache Coherence Bus
CPU
CPU
NIC
DRAM
DRAM
Memory-Fabric
(e.g. Gen-Z)
Storage-Fabric
(e.g. NVMeoF)
CPU-centric architecture
Data-centric architecture
Western Digital

Why People use Solid State Storage?

- It is robust—so good in mobile consumer devices
- It is fast—so people like it in some enterprise applications and client computers
- NVMe and NVMe-oF allow building storage devices and network storage without restrictions from HDD interfaces
- The groundwork is being laid for memory centric computing
- There will be many types of storage technology used for several more years


STORAGE VISIONS 2017
AN ENTERTAINMENT STORAGE ALLIANCE™ EVENT
ENTERTAINMENT STORAGE ALLIANCE™
October 16, 2017 at The Embassy Suites, Milpitas, CA
Dr. Tom Coughlin is a highly respected technology and market analyst and author and the Founder and Chairman of the Storage Visions Conference.
Storage Growth!
2016
16ZB
2025
163ZB
Register Now
to get conference and room discounts
The Storage Visions® 2017 Conference Theme:
New Visions for Digital Storage !
Join Us at our NEW VENUE; The Embassy Suites, Milpitas, CA and Help Us Celebrate Our 16th Anniversary
JOIN CEO's, Industry Leaders, manufacturers and end users at Storage Visions 2017! FOCUS ON:
· Big Data Analytics
· Video (M&E and Surveillance)
· Medical and Genomic Data
· Cloud and Fog Computing Solns
· In Memory Processing
· Exabyte Video Projects
· Artificial Intelligence
· Internet of Things
· Big Science
· New Memory/Storage
· Hyperscale
· Gb/s Cars
Register at www.storagevisions.com
WWW.STORAGEVISIONS.COM

References

- 2017 FMS Keynote Presentations by Intel, Micron, Samsung, Seagate, Western Digital
- 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