HEPiX Techwatch WG : Flash Storage

First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation)

“Tape is Dead, Disk is Tape, Flash is Disk, RAM Locality is King”, 12/2006, Jim Gray, Microsoft.

Staging area

References

Following are links to articles and information that should be integrated into the document:

- Recent News (Add links to new articles for later integration in this document))
    - Micron NY Fab
    - Micron Financials
    - New Micron Data Center SSDs
    - Samsung High Capacity SSD
    - SK Hynix “AI” SSD
    - Kioxia Flash Chip
    - WD Financials“
- Full Picture” Sources
    - (Updated 9/24) Library of Congress National Digital Information Infrastructure and Preservation Program - The annual Designing Storage Architectures Meetings (DSA meeting) draws speakers from various industry and academic organizations.
    - (Updated 9/24) The presentations by Georg Lauhoff (IBM) at the DSA meetings are particularly relevant as they cover storage technology and market in detail.
    - (Updated 9/24) The Information Storage Industry Consortium (INSIC) is a consortium of academic, industry and government organizations in the field of information storage. They publish the INSIC Tape Technology Roadmap on a periodic basis.
    - (Updated 9/24) SpectraLogic Data Storage Outlook Report - Annual storage market survey from SpectraLogic, a manufacturer of tape automation

    equipment.T he most recent report is from 2023 and can be downloaded from SpectraLogic (requires registration). Older reports also appear to be available (The SpectraLogic website search tool can be used to locate reports from previous years.)

- News Sites and Market Analysts
    - (Updated 1/23) Blocks and Files is a website that provides news and analysis of the storage market
    - (Updated 1/23) TrendFocus - Market analysis company focusing on storage.
- Storage Trade Organization
    - (Updated 1/23) SNIA - The Storage Networking Industry Association is, to quote their website: “a not-for-profit global organization that leads the storage industry in developing and promoting vendor-neutral architectures, standards, and educational services that facilitate the efficient management, movement, and security of information.” SNIA sponsors multiple conferences and technical working groups and publishes numerous resources covering storage technology. Relevant SNIA sponsored conferences and meetings include the following:
        - SNIA Storage Developer Conference
        - SNIA Persistent Memory and Computational Storage Summit
        - SNIA Preview
- Standards and associated trade organizations
    - (Updated 1/23) SATA-IO.org - Custodians of the SATA standard
    - (Updated 1/23) SCSI Trade Association (SCSITA) - SAS/SATA trade organization
    - (Updated 1/23) INCITS T10 Technical Committee - Custodians of the SCSI and SAS standard
    - (Updated 1/23) Fibre Channel Industry Association - Custodians of the Fibre Channel storage area network standard
    - (Updated 1/23) Infiniband Trade Organization - Custodians of the Infiniband standard and the RoCE (RDMA over Converged Ethernet) standard
    - PCI-e standard is developed by the PCI Special Interest Group (PCI-SIG).
    - NVMexpress.org is the organization developing the NVMe including NVMeoF the standard.
        - Additional information on NVMe can be found at the NVMe Developers Day conference.
    - (Updated 1/23) Computer Express Link - Custodians of the CXL CPU to device interconnect standards.
    - IEEE 802.3 Working Group develops the Ethernet standard.
    - (Updated 1/23) Zone Storage IO - Website dedicated to the Zoned Storage (HDD) and Zoned Namespace (NVMe SSD) APIs.
- Links that need to be cleaned up
    - NVMexpress.org is the organization developing the NVMe standards
    - The Joint Electron Device Engineering Council (JEDEC) is responsible for the three NVDIMM terminology, “-F”, “-N”, and “-P”. JEDEC is also

  responsible for DDR4 and DDR5 memory standards that include support for the various NVDIMM types.
  - The ACPI specification part of the UEFI specification. UEFI support for NVDIMM support in ACPI 6.2 interface, including NVDIMM Firmware Interface Table (NFIT).
  - SNIA NVM Programming Working Group has defined a programming model for non-volatile memory, the SNIA NVM Programming Model.
  - The Persistent Memory Development Kit (PMDK) is an open source library for Linux and Windows that can be used by applications to access persistent memory in the ways outlined by the SNIA NVM Programming model. Development of the PDMK is managed by Intel. The web site for PDMK development is pmem.io
  - An overview of the state of persistent memory technology can be found a the SNIA 7th Annual Persistent Storage Summit
  - The annual Flash Memory Summit has a plethora of information on persistent memory. An archive of their proceedings since 2006 is at Flash Memory Summit Proceedings.
  - SCMandMIEC_overview_12Feb2013
  - http://smorgastor.drhetzler.com/wp-content/uploads/2014/08/SSD-Reliability.pdf
  - The Memory Guy is a web site run by an industry analyst that covers memory technology. For those with money, a quick way to get an in depth view of the memory industry can be obtained by purchasing the analyst's memory technology reports.
  - The Memory Guy's market research company is Objective Analysis. Although you need to pay for their in depth research reports, the outlines for the reports are freely available and provide an high level view of the important areas memory technology

Sources

References to sources of information that cover the technologies covered in this report.
The following are links to online resources that cover storage technology. The sites include news sites, industry/trade/user organizations, companies, and standards organizations. Much of the material in this document was drawn from these sites.

1. (Updated 1/23) DRAMeXchange - DRAM/Flash market analysis company, a part of TrendForce (not to be confused with TrendFocus). Their Weekly Research page covers news in the memory industry.
2. (Updated 1/23) SNIA - The Storage Networking Industry Association is, to quote their website: "a not-for-profit global organization that leads the storage industry in developing and promoting vendor-neutral architectures, standards, and educational services that

facilitate the efficient management, movement, and security of information." SNIA sponsors multiple conferences and technical working groups and publishes numerous resources covering storage technology. Relevant SNIA sponsored conferences and meetings include the following:
    a. SNIA Storage Developer Conference
    b. SNIA Persistent Memory and Computational Storage Summit
    c. SNIA Preview
3. (Updated 1/23) Flash Memory Summit - Exposition and conference focused on the persistent memory industry
4. (Updated 3/23) Chinese Flash Memory Summit
5.
6. (Updated 1/23) Seagate Technology - HDD manufacturer website. Occasionally posts technical information in the Investor Relations page. 2019 and 2021 Analyst Day presentations are particularly noteworthy.
7. (Updated 1/23) Western Digital - HDD manufacturer website. Their Investor Day presentations, available from their Investor Relations page, provides technical information on the state of HDD and Flash technologies (and market)
8. (Updated 1/23) Toshiba - HDD manufacturer website. Occasionally posts technical information in the Investor Relations page
9. (Updated 1/23) Micron - DRAM/NAND memory manufacturer website. Occasionally posts technical information in the Investor Relations page
10. (Updated 1/23) Kioxia - DRAM/NAND manufacturer website. Information about Kioxia memory technologies can be found on their R&D page.
11. (Updated 2/23) Samsung - DRAM/NAND manufacturer website.. Information about Samsung memory technologies can be found in their Investor Relations page.
12. (Updated 1/23) ASML - Sole manufacturer of EUV lithography equipment. Their Investor Day presentations typically provide insight into the state of DRAM and NAND technology.
13. (Updated 1/23) ISSCC - International Solid-State Circuit Conference website. ISSCC covers integrated circuits of all types including logic, communications, and memory devices. Their conference Press Kits provide a high level summary of the technical trends in each class of devices.

Suggested Outline

1. Executive Summary
2. Burning questions
    a. Will SSD match HDD in terms of $/TB
    b. What is the SSD capacity and price roadmap
    c. What is the future of SSD form factors/connectivity
        i. Are there alternatives
    d. Can you build your own low latency, high bandwidth storage system ?
    e. Is software RAID viable for NVMe storage
3. Basic Technology

    a. Cell Types
        i. Charge Trap
        ii. Floating Gate
    b. Multi-layer
    c. 3D NAND Flash
    d. Multi-level (SLC, MLC, TLC, QLC, PLC)
    e. Die stacking ?
    f. NAND Alternatives
4. Packaging and Interfaces
    a. Packaging
        i. M.2
        ii. U.2/U.3
        iii. EDSFF
    b. Interfaces
        i. NVMe/NVMeoF
        ii. SATA/SAS
        iii. CXL ?
5. Market
        i. Micron
        ii. Kioxia
        iii. Samsung
        iv. Western Digital
        v. SK Hynix/Solidigm
    b.
    c. Flash Systems
6. Software
    a. Zone Storage
    b. SPDK
    c. NVMe
    d. GPU Direct
7. Storage Systems
    a. “Exotic” low latency, high bandwidth storage systems (e.g. VAST, Weka.IO, Pure)
    b. Intel DAOS

Executive Summary

Summary of key findings, including possible impact of cost and technology evolution on HEP/NP in the future. (Single Paragraph)

1. Flash Summary

> Cost per TB continues to drop and capacities increase as manufacturers move to flash memory with more layers. Status of Penta Level Cells (PLC) is unclear. Performance continues to increase as SSD's move to PCI-e Gen 5. SSDs now available in capacities that exceed the largest HDD drives.

Burning Questions with Answers

1. Long term viability, both technical and financial, of HDD
2. Will flash subsume the HDD market
3. Long term financial viability of tape
4. Impact of cloud on the tape and HDD market

<Insert Answers Here>

Introduction

High energy (HEP) and nuclear physics (NP) experiments have entered the "exabyte" era, as they are expected to generate 100s of petabytes of data per year and run for decades. Collecting, storing and analyzing this large volume of data has historically required the use of leading edge storage technologies of all major types at relatively large scale. In recent years there has been a concern that the rate of advancement in storage technologies has not kept pace with the needs of the community, particularly in the area of cost per TB of capacity. The HEP and NP communities have been working to reduce their storage requirements, but the success of these mitigation efforts will be partially determined by the rate of storage advancement over the next decade.

This document covers the major storage technologies that are, or are expected to be in use by the HEP/NP community moving forward. Topics covered include their current state, their projected evolution, their financial viability and the relationship between the various technologies, particularly flash vs HDD and HDD vs tape.

Overview

Analysis of storage technologies starts with an analysis of the two major random access, block storage technologies in use, magnetic disk drives (HDD) and flash storage (SSD), each in isolation. This is followed by a look at the competition between HDD and SSD storage. The analysis then turns to secondary technologies that have an impact on how HDDs and SSDs are deployed in storage systems. At this point, the analysis turns to tape storage. This is then followed by a look at the competition between HDD and tape storage. The final section covers alternatives to tape for archival storage. The following summarized the major sections in this report.

1. Magnetic disk drives (HDD)
2. Flash storage (SSD)
3. Competition between HDD and flash technologies
4. Secondary technologies associated with “disk” storage
5. Tape storage
6. Competition between HDD and Tape storage
7. Alternatives to tape

Flash Storage (SSD)

Executive Summary

<Insert Content Here>

Burning Questions

<Insert Content Here>

Basic Technology

NAND Flash

NAND flash is the underlying memory technology in the majority of commercially available SSDs and Persistent Memory. This section discusses this underlying technology, leaving a discussion of SSDs, Persistent Memory, and other packaging in later sections. The basis of NAND flash is a “cell” holding storaged charge, of which there are two types: “Floating Gate” and “Charge Trap”. Flash memory chips consist of silicon chips on which cells are constructed. Early devices, 2D or Planar NAND, had cells laid out in a 2 dimensional array while current 3D NAND devices consist of multiple layers of cells. Over the years, 3D NAND devices have increased in bit densities through three mechanisms:

1. Smaller cell sizes
2. More layers
3. More bits per cell

The following graph from the 2022 ISSCC Trends Report shows the progression in bit density for NAND flash over the past two decades.

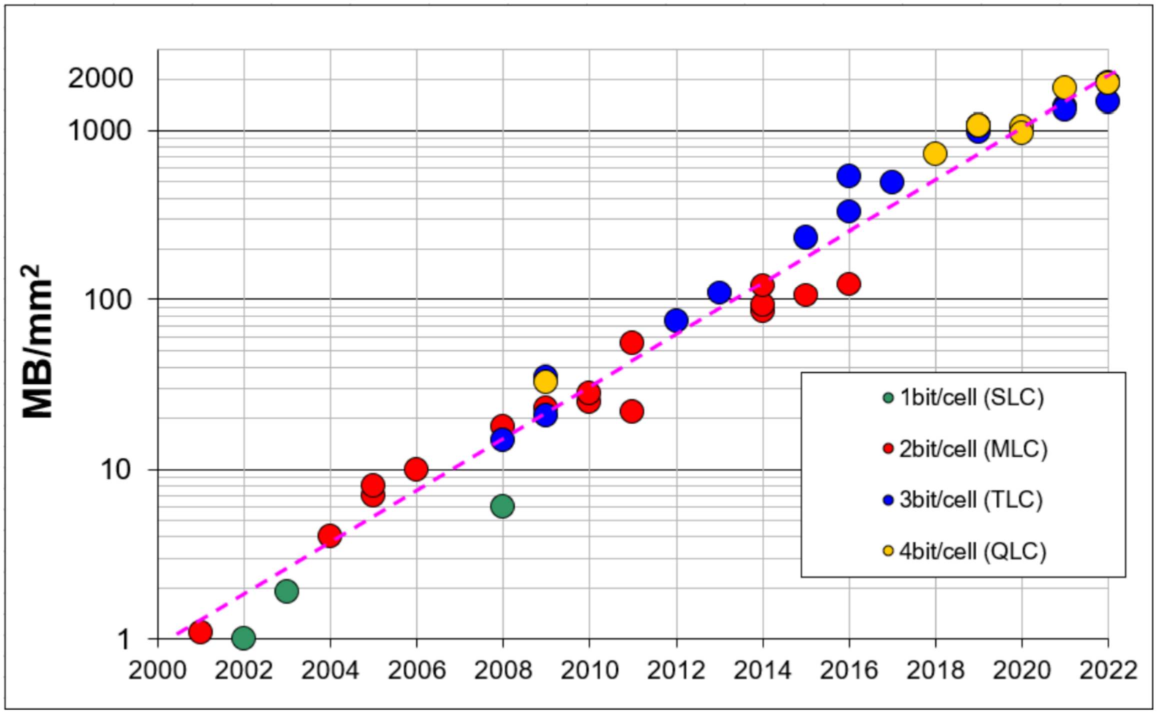

(Updated 1/23) Bit Density for NAND Flash vs Time (2022 ISSCC Trends Report)

3D NAND Flash (Layering/Stacking)

Of the three mechanisms for achieving higher capacity flash memory, the addition of layers to 3D NAND flash is the most obvious. At this point in time (Jan 2023), all major flash vendors have 3D NAND flash with more than 100 layers in volume production. Next generation flash devices with around 200 layer devices are on the horizon. In July of 2022 Micron started shipping a 1Tb 32 layer TLC 3D NAND chip. The devices is built from two 116 layer “decks” (Anandtech).

The following table from blocksandfiles.com, shows the number of layers in each generation of 3D NAND flash from the major vendors.

Micron		SK Hynix		Samsung		Western Digital/Kioxia	
Generation	Layers	Generation	Layers	Generation	Layers	Generation	Layers
Gen 1	32	V3	48	V3	48	BICS 2	32

Gen 2	64	V4	72	V4	64	BICS 3	64
Gen 3	96	V5	96	V5 (2018)	9x	BICS 4	96
Gen 4 (2019)	128	V6 (2019)	128	V6 (2019)	128	BICS 5 (2020)	112
Gen 5 (2020)	176	V7 (2021)	176	V7 (2021)	176	BICS 6	162
Gen 6 (2022)	232			V8 (2021)	200	BICS 7	212
Gen 7	2xx			V9	3xx	BICS 8	
Gen 8	3xx			V10			
Gen 9	4xx			V11			

An open question with regards to additional layers, is how far this stacking can be extended. The diagram below shows the structure of a typical 3D NAND flash chip.

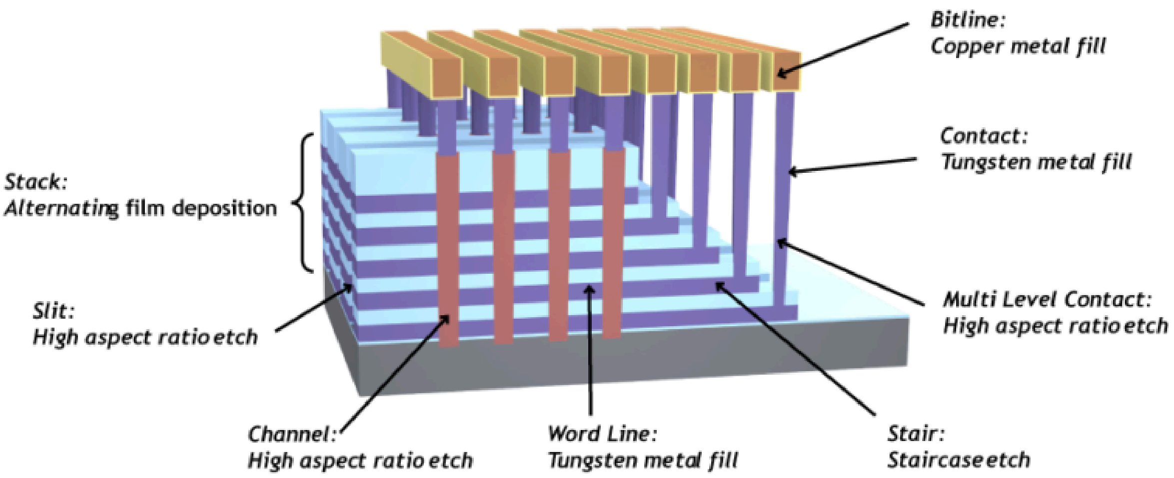

A high level overview of the steps involved in manufacturing 3D NAND can be found in the EE Journal article at the following link. Among the challenges with additional layers are the following:

1. Increased chip area relegated to the “staircase” needed to access to memory cells in each layer.
2. Manufacturing difficulties caused by increased channel aspect ratio (channel length vs. width) with more layers.
3. Increased manufacturing time, caused by the increase in the number of layers.

The following two diagrams show the challenges involved in the creation of the channels.

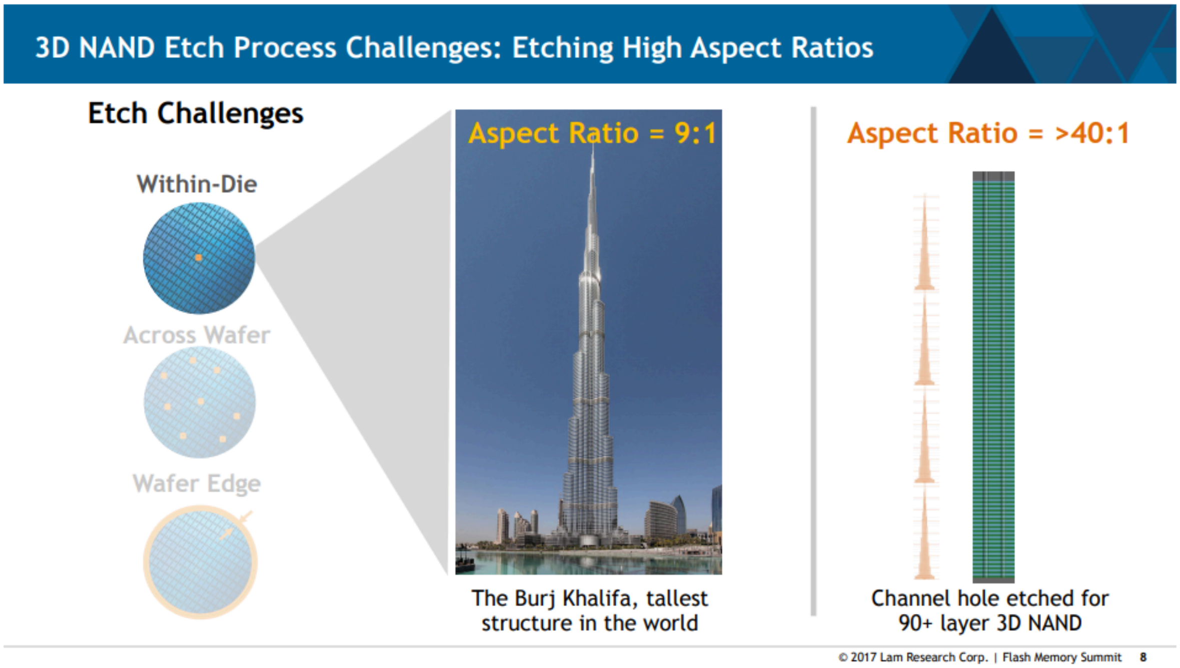

3D NAND Etch Process Challenges: Etching High Aspect Ratios

Etch Challenges

Within-Die

Across Wafer

Wafer Edge

Aspect Ratio = 9:1

The Burj Khalifa, tallest structure in the world

Aspect Ratio = >40:1

Channel hole etched for 90+ layer 3D NAND

© 2017 Lam Research Corp. | Flash Memory Summit 8

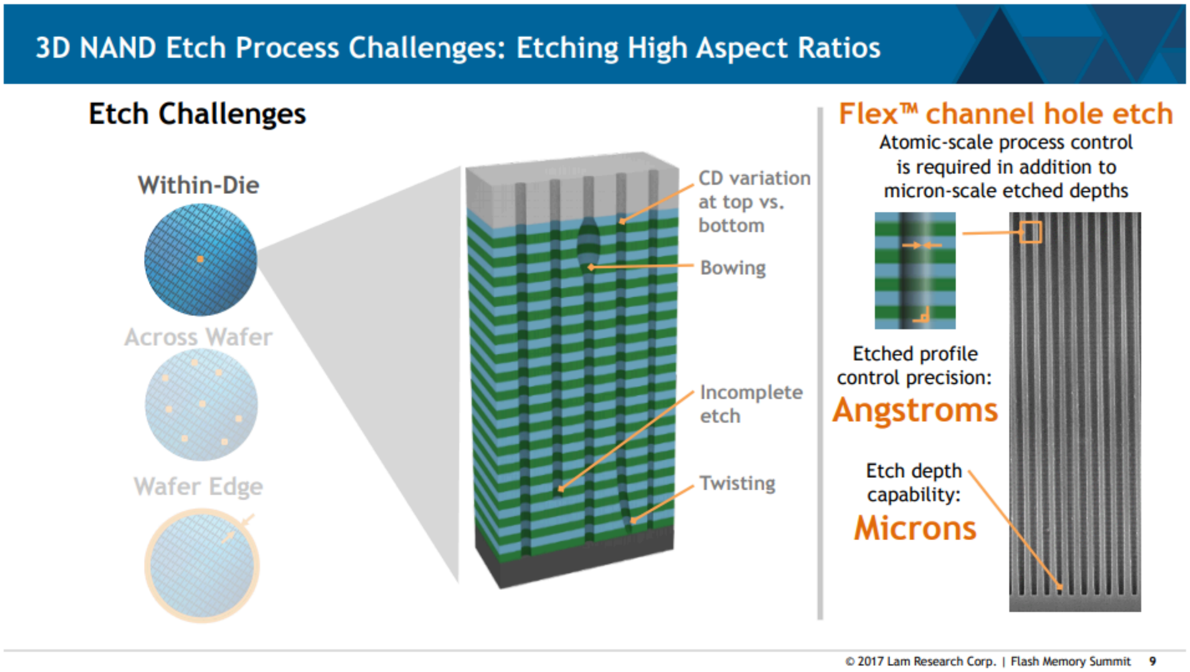

3D NAND Etch Process Challenges: Etching High Aspect Ratios

Etch Challenges

Within-Die

Across Wafer

Wafer Edge

CD variation at top vs. bottom

Bowing

Incomplete etch

Twisting

Flex™ channel hole etch

Atomic-scale process control is required in addition to micron-scale etched depths

Etched profile control precision:
Angstroms

Etch depth capability:
Microns

© 2017 Lam Research Corp. | Flash Memory Summit 9

Multi-Level NAND (TLC/QLC/PLC)

In another path to higher density flash memory is storing more bits per cell. At the present time all major manufacturers of 3D NAND have triple level cell (TLC) and quad level cell (QLC) flash chips. With the original single level cell (SLC) NAND flash chips, each cell stored on bit, corresponding to charge/no charge state. With TLC each cell stores three bits, corresponding to 8 different charge states while with QLC each cell stores four bits, corresponding to 16 different charge states.

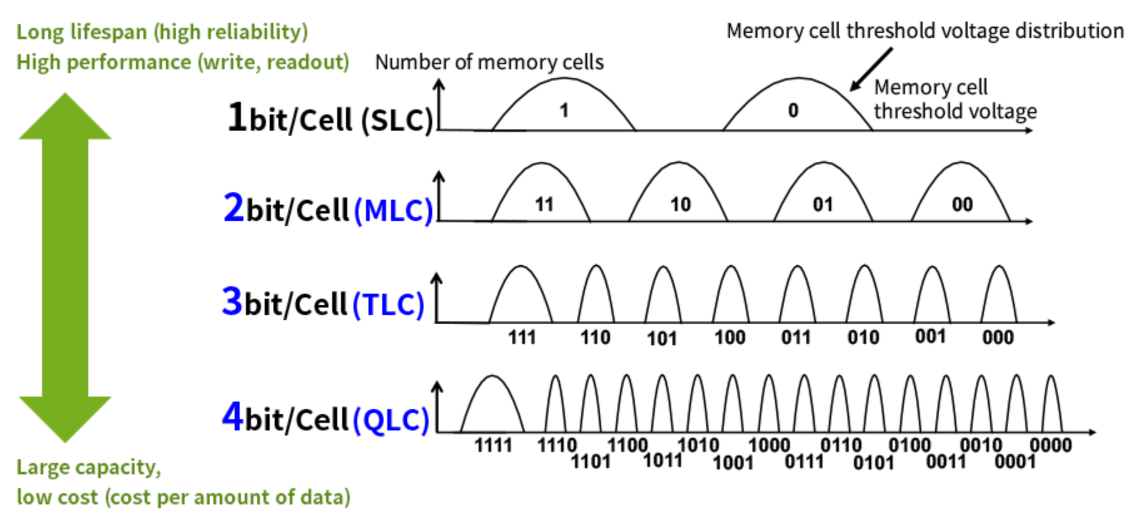

(Updated 1/23) Illustration of the differences between SLC, MLC, TLC, and QLC NAND flash cells. Kioxia

As would be expected, distinguishing more charge states is harder than fewer charge states. QLC will also have lower endurance (Program/Erase cycles) and lower performance than TLC flash. Beyond QLC, both Solidigm (formerly Intel) and Kioxia (formerly Toshiba Memory) have announced they are working on penta level cells (PLC) or 5 bits /cell. With regards to endurance, the following chart from Wells Fargo, via blocksandfiles.com, shows the reduction in endurance as a function of bits per cell and small device geometries.

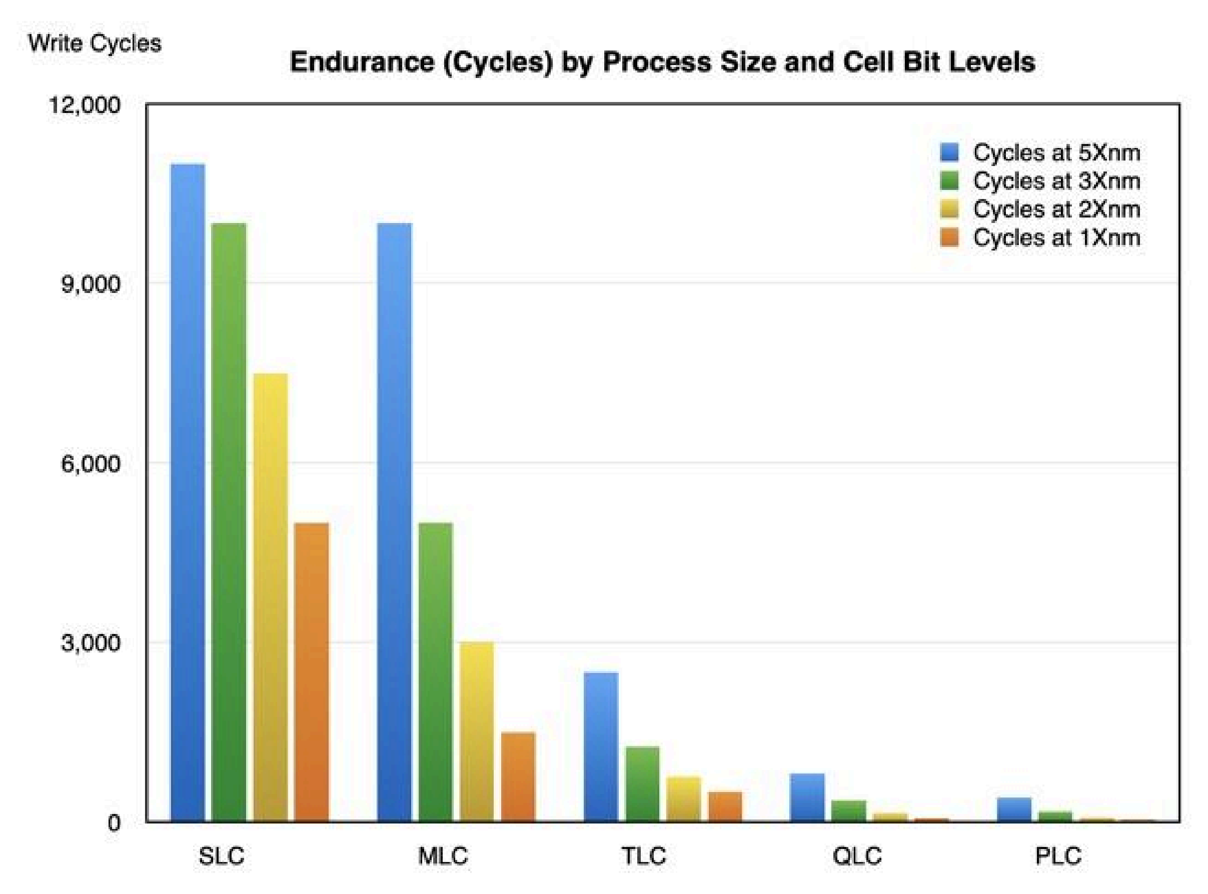

Wells Fargo, via blocksandfiles.com. Reduction in endurance as a function of bits per cell and small device geometries.

Taken in isolation, NAND flash technology is of little use to the HEP community, rather, it is the packaging and presentation of this technology that is of use. After obtaining an initial toe hold in the storage market through Solid State Disks (SSDs), new packaging and presentations of NAND flash are appearing that are poised to revolutionize storage.

Challenges for NAND FLASH

The direction of most new flash products is largely towards higher capacity. The reduction in the endurance and retention of higher capacity has been mitigated through a massive increase in ECC overhead in terms of bits and processing power. In addition, a reduction of IOPS (at the storage interface, i.e. SAS) could be observed (compared to previous generation MLC or lower devices). Asymmetric performance (write slower than reads) continues or even slightly grows.

(Updated 1/23) A niche market that is being pursued by Samsung is low latency flash, branded by Samsung as Z-NAND. This product is meant to compete with Intel and Micron Technology’s 3D XPoint non volatile memory. Kioxia recently entered this market with their XL-Flash memory.

3D XPoint (Optane) (DEAD)

(Updated 1/23) 3D XPoint (Optane) was a non volatile memory technology developed by Intel and Micron Technology that is claimed to have lower latency and higher endurance than NAND. Micron positions 3D XPoint between NAND flash and DRAM, being faster but more expensive than NAND, but cheaper and higher capacity than DRAM. 3D Xpoint memory was available in two forms, SSDs and modified DDR4 DIMMs. Utilization of the later required explicit support by the CPU. Optane DIMM support was provided in specific models of processors from Intel. For various reasons, Optane failed to gain market acceptance, resulting in Micron exiting the business in 2021. Intel followed suit in late 2022.

Emerging Memory Technologies

NAND flash is the non-volatile solid state memory of choice at this point in time. There are other non-volatile solid state memory technologies that are in development or in early stages of commercial product deployment that are in some ways similar to 3D XPoint. Except for some types of Magnetoresistive (MRAM) and Carbon Nanotube (NRAM) memory, these emerging memory technologies fit between DRAM and flash in the memory hierarchy as show below in the diagram from a presentation by James Singer of Everspin Technologies at the Flash Memory Summit.

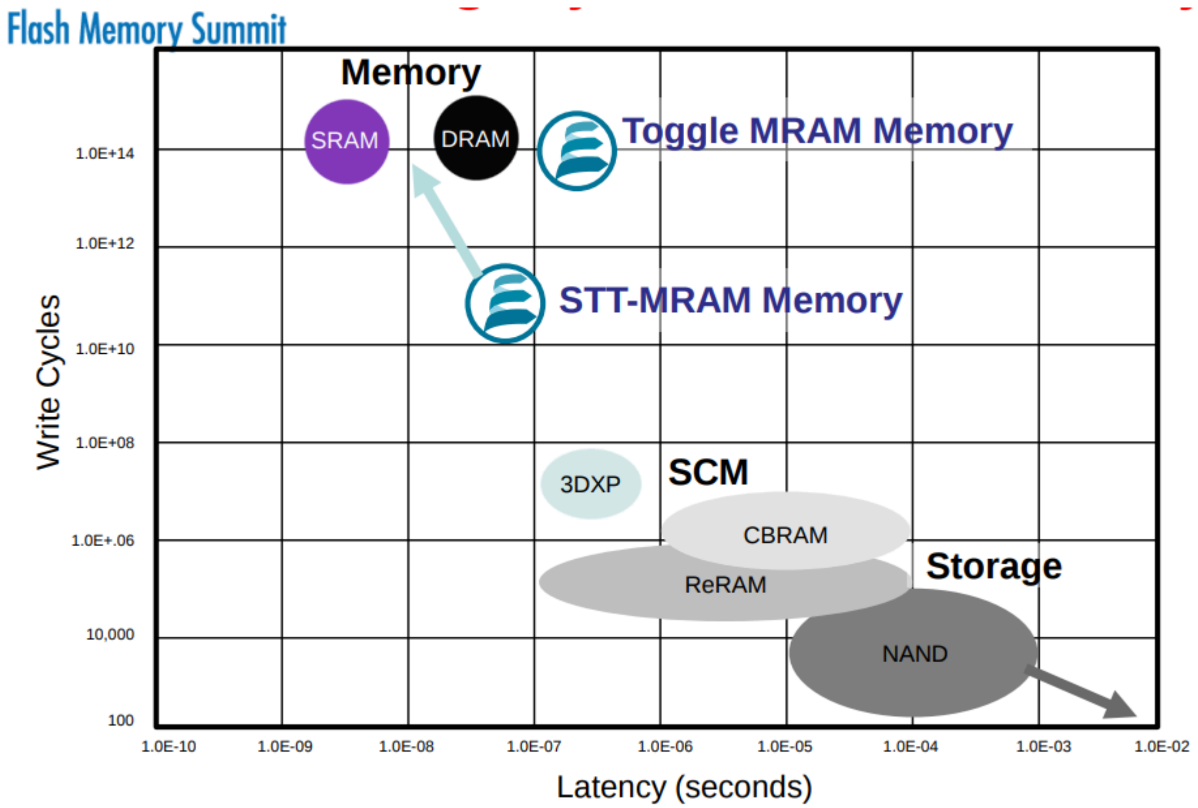

An incomplete list of the classes of emerging memory in development, in no particular order, is as follows:

1. PCM (Phase Change Memory)
2. STT-MRAM (Spin-Transfer Torque Magneto resistive RAM)
3. MRAM (Magneto resistive RAM)
4. FeRAM (Ferroelectric RAM)
5. NRAM (Nanotube RAM)
6. ReRAM (Resistive RAM)

A table of the salient characteristics of different emerging memory technologies from a presentation by Jim Handly and Tom Coughlin of Objective Analysis at the 2018 SNIA Persistent Memory Summit is shown below:

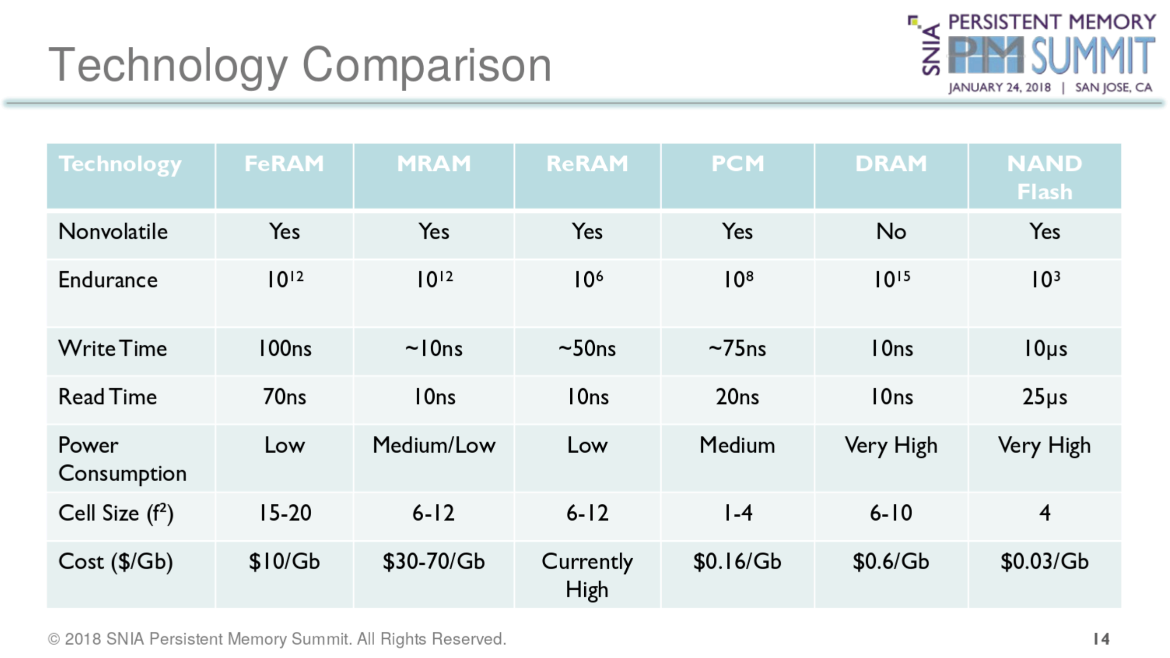

Technology Comparison

SNIA PERSISTENT MEMORY PM SUMMIT
JANUARY 24, 2018 | SAN JOSE, CA

Technology	FeRAM	MRAM	ReRAM	PCM	DRAM	NAND Flash
Nonvolatile	Yes	Yes	Yes	Yes	No	Yes
Endurance	10^{12}	10^{12}	10^{6}	10^{8}	10^{15}	10^{3}
Write Time	100ns	~10ns	~50ns	~75ns	10ns	10µs
Read Time	70ns	10ns	10ns	20ns	10ns	25µs
Power Consumption	Low	Medium/Low	Low	Medium	Very High	Very High
Cell Size (f^2)	15-20	6-12	6-12	1-4	6-10	4
Cost ($/Gb)	$10/Gb	$30-70/Gb	Currently High	$0.16/Gb	$0.6/Gb	$0.03/Gb

© 2018 SNIA Persistent Memory Summit. All Rights Reserved.

14

Capacity trends for the various emerging memory technologies, from the 2022 ISSCC Trends report which summarizes key trends from the presentations at the IEEE International Solid-State Circuits Conference (ISSC), is shown in the following graph.

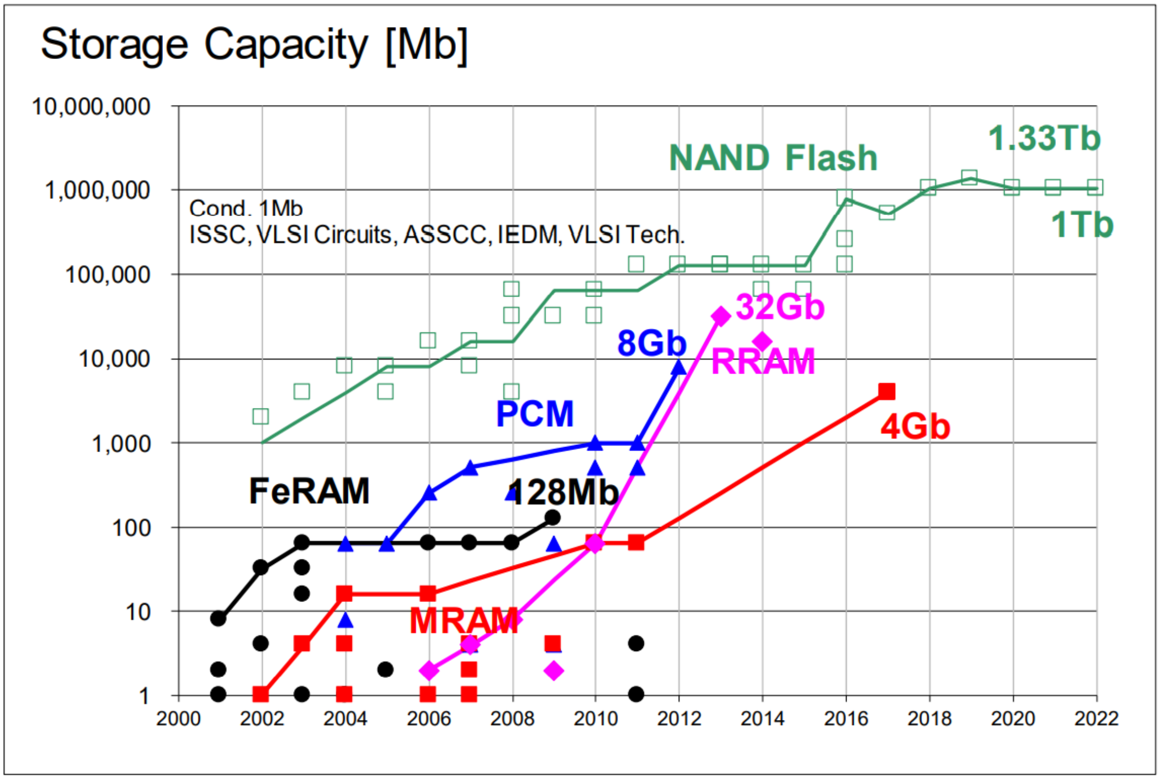

(Updated 1/23) Storage Capacity Trends in Non Volatile Memory Technology (2022 ISSCC Trends Report)

The following graph from the 2021 ISSCC Trends report shows read/write bandwidth trends for the various emerging memory technologies.

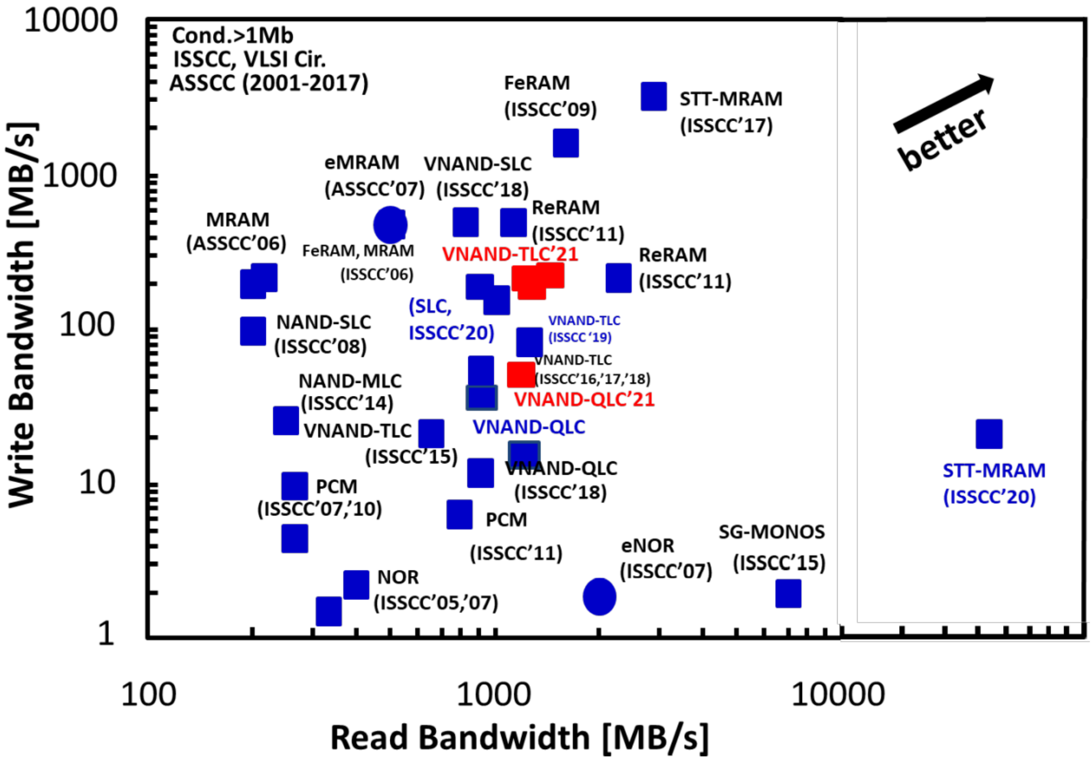

(Updated 1/23) Write Bandwidth Trends in Non Volatile Memory Technology (2021 ISSCC Trends Report)

The following chart from Tech Insights Emerging Memory Technology/Product Roadmap shows the various players in the different emerging memory fields.

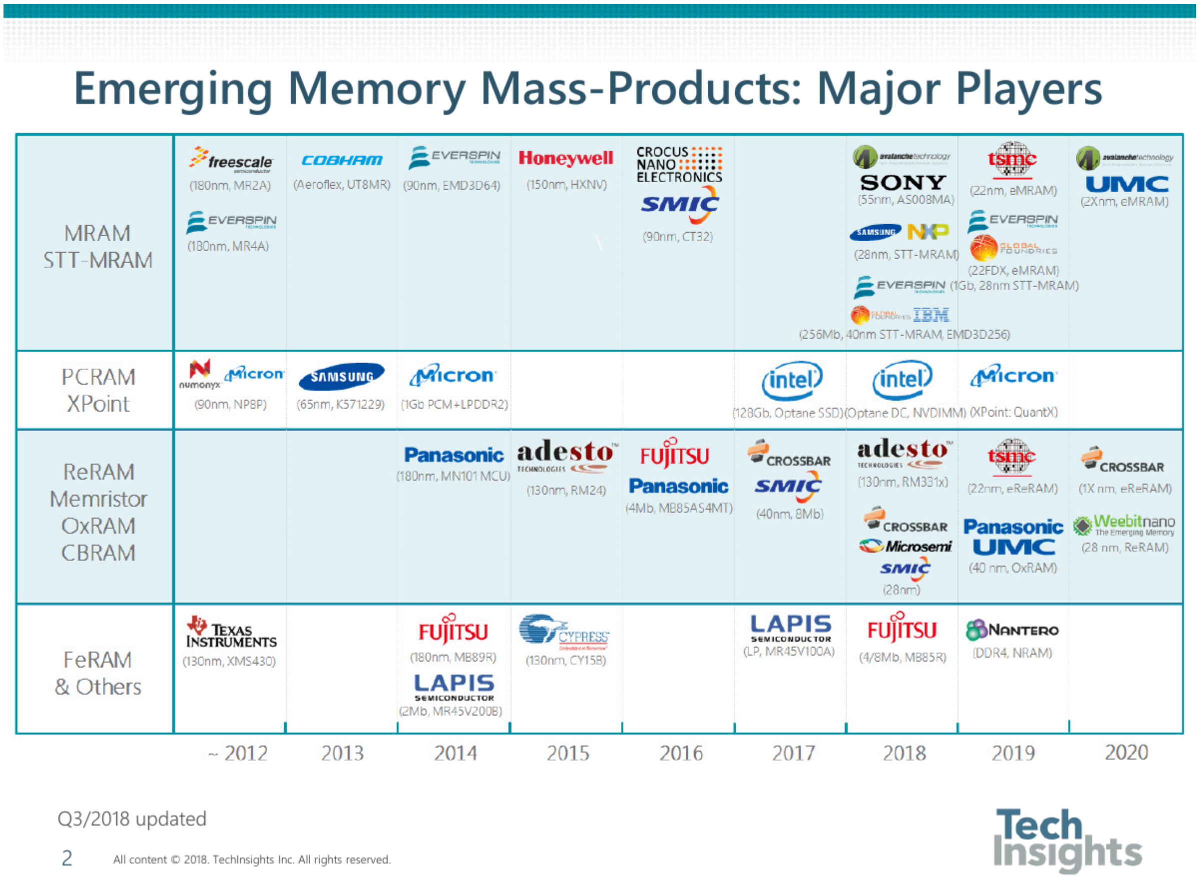

Emerging Memory Mass-Products: Major Players

	~ 2012	2013	2014	2015	2016	2017	2018	2019	2020
MRAM STT-MRAM	freescale semiconductor (180nm, MR2A) EVERSPIN TECHNOLOGIES (180nm, MR4A)	COBHAM (Aeroflex, UT8MR)	EVERSPIN TECHNOLOGIES (90nm, EMD3D64)	Honeywell (150nm, HXNV)	CROCUS NANO ELECTRONICS SMIC (90nm, CT32)		avalanche technology SONY (55nm, AS008MA) SAMSUNG NXP (28nm, STT-MRAM) EVERSPIN TECHNOLOGIES (1Gb, 28nm STT-MRAM) GLOBALFOUNDRIES IBM (256Mb, 40nm STT-MRAM, EMD3D256)	tsmc (22nm, eMRAM) EVERSPIN TECHNOLOGIES GLOBALFOUNDRIES (22FDX, eMRAM)	avalanche technology UMC (2Xnm, eMRAM)
PCRAM XPoint	numonyx Micron (90nm, NP8P)	SAMSUNG (65nm, K571229)	Micron (1Gb PCM+LPDDR2)			intel (128Gb, Optane SSD)	intel (Optane DC, NVDIMM)	Micron (XPoint: QuantX)	
ReRAM Memristor OxRAM CBRAM			Panasonic (180nm, MN101 MCU)	adesto TECHNOLOGIES (130nm, RM24)	FUJITSU Panasonic (4Mb, MB85AS4MT)	CROSSBAR SMIC (40nm, 8Mb)	adesto TECHNOLOGIES (130nm, RM331x) CROSSBAR Microsemi SMIC (28nm)	tsmc (22nm, eReRAM) Panasonic UMC (40 nm, OxRAM)	CROSSBAR (1X nm, eReRAM) Weebitnano The Emerging Memory (28 nm, ReRAM)
FeRAM & Others	TEXAS INSTRUMENTS (130nm, XMS430)		FUJITSU (180nm, MB89R) LAPIS SEMICONDUCTOR (2Mb, MR45V200B)	CYPRESS (130nm, CY15B)		LAPIS SEMICONDUCTOR (LP, MR45V100A)	FUJITSU (4/8Mb, MB85R)	NANTERO (DDR4, NRAM)	

Q3/2018 updated

2 All content © 2018. TechInsights Inc. All rights reserved.

TechInsights

Market Overview

For enterprise SSDs, prices have been dropping as a result of slowing demand. According to TrendForce combined revenue from the major enterprise SSD vendors dropped by 28.7% in the 3rd quarter of 2022 and are expected to continue falling in Q4. The market is also transitioning from PCI-e Gen 3 to Gen 4 based SSD, which may also be causing a drop in demand. Enterprise SSD market share by supplier, according to TrendForce via DRAMeXchange weekly research from December 5, 2022 is shown below:

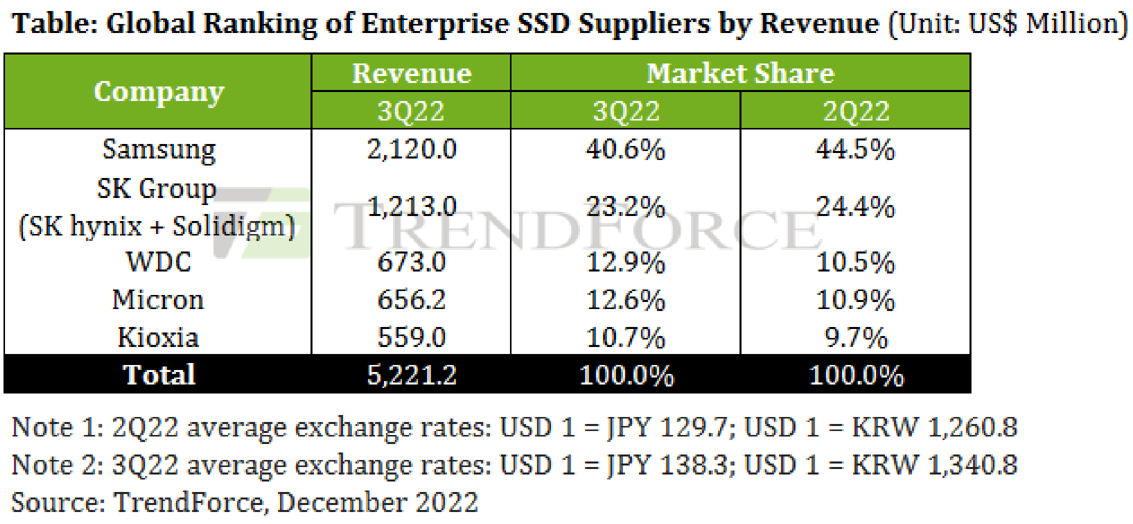

Table: Global Ranking of Enterprise SSD Suppliers by Revenue (Unit: US$ Million)

Company	Revenue	Market Share	
	3Q22	3Q22	2Q22
Samsung	2,120.0	40.6%	44.5%
SK Group (SK hynix + Solidigm)	1,213.0	23.2%	24.4%
WDC	673.0	12.9%	10.5%
Micron	656.2	12.6%	10.9%
Kioxia	559.0	10.7%	9.7%
Total	5,221.2	100.0%	100.0%

Note 1: 2Q22 average exchange rates: USD 1 = JPY 129.7; USD 1 = KRW 1,260.8
Note 2: 3Q22 average exchange rates: USD 1 = JPY 138.3; USD 1 = KRW 1,340.8
Source: TrendForce, December 2022

(Updated 1/23) Enterprise SSD market share by supplier (TendForce via DRAMeXchange)

Like the enterprise SSD market, the broader flash market also dropped in 2022Q3 with quarter over quarter revenues falling 24.3% according to TrendForce,. The drop in global flash revenue, like the SSD market, have been attributed to falling demand. This table from TrendForce (via blocksandfiles.com) shows the breakdown in revenues by vendor.

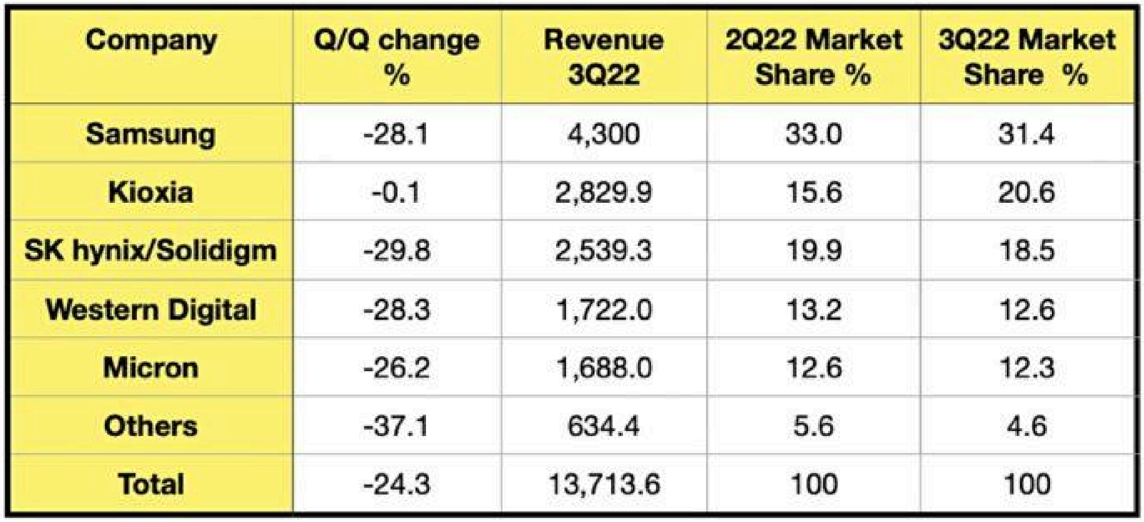

Company	Q/Q change %	Revenue 3Q22	2Q22 Market Share %	3Q22 Market Share %
Samsung	-28.1	4,300	33.0	31.4
Kioxia	-0.1	2,829.9	15.6	20.6
SK hynix/Solidigm	-29.8	2,539.3	19.9	18.5
Western Digital	-28.3	1,722.0	13.2	12.6
Micron	-26.2	1,688.0	12.6	12.3
Others	-37.1	634.4	5.6	4.6
Total	-24.3	13,713.6	100	100

(Updated 1/23) Q over Q change in flash revenue by vendor (TrendForce via blocksandfiles.com)

Market share according to TrendForce via DRAMeXchange

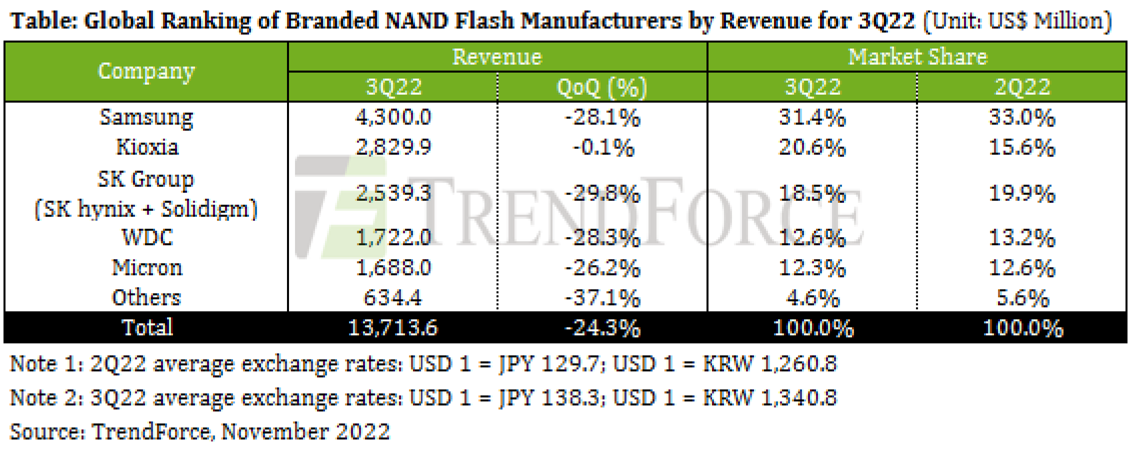

Table: Global Ranking of Branded NAND Flash Manufacturers by Revenue for 3Q22 (Unit: US$ Million)

Company	Revenue		Market Share	
	3Q22	QoQ (%)	3Q22	2Q22
Samsung	4,300.0	-28.1%	31.4%	33.0%
Kioxia	2,829.9	-0.1%	20.6%	15.6%
SK Group (SK hynix + Solidigm)	2,539.3	-29.8%	18.5%	19.9%
WDC	1,722.0	-28.3%	12.6%	13.2%
Micron	1,688.0	-26.2%	12.3%	12.6%
Others	634.4	-37.1%	4.6%	5.6%
Total	13,713.6	-24.3%	100.0%	100.0%

Note 1: 2Q22 average exchange rates: USD 1 = JPY 129.7; USD 1 = KRW 1,260.8

Note 2: 3Q22 average exchange rates: USD 1 = JPY 138.3; USD 1 = KRW 1,340.8

Source: TrendForce, November 2022

(Updated 1/23) Global flash market share by supplier (TendForce via DRAMeXchange)

With pandemic induced demand and supply chain issue waning, the external factors that may be influencing the market include efforts of government agencies to control inflation, the war in Ukraine, ongoing COVID issues in China, diversification/”reshoring” of supply chains and US-China relations. The 2022 CHIPS and Science Act in 2022 by the US Congress, a bill to encourage preserving/growing semiconductor production in the US, is a prime example of actions that will have an effect on the memory market.

An unknown moving forward is the effects of market consolidation. In January of 2022 SK Hynix purchased Intel’s NAND and SSD business. A potential change with even bigger impact is the possible merger of Koxia (formerly Toshiba Memory) and Western Digital. Their combined revenue would rival the market leader Samsung. US trade restriction on NAND technology has the potential to severely affect YMTC.

Utilization of NAND memory by market sector is shown in the following graph by Micron.

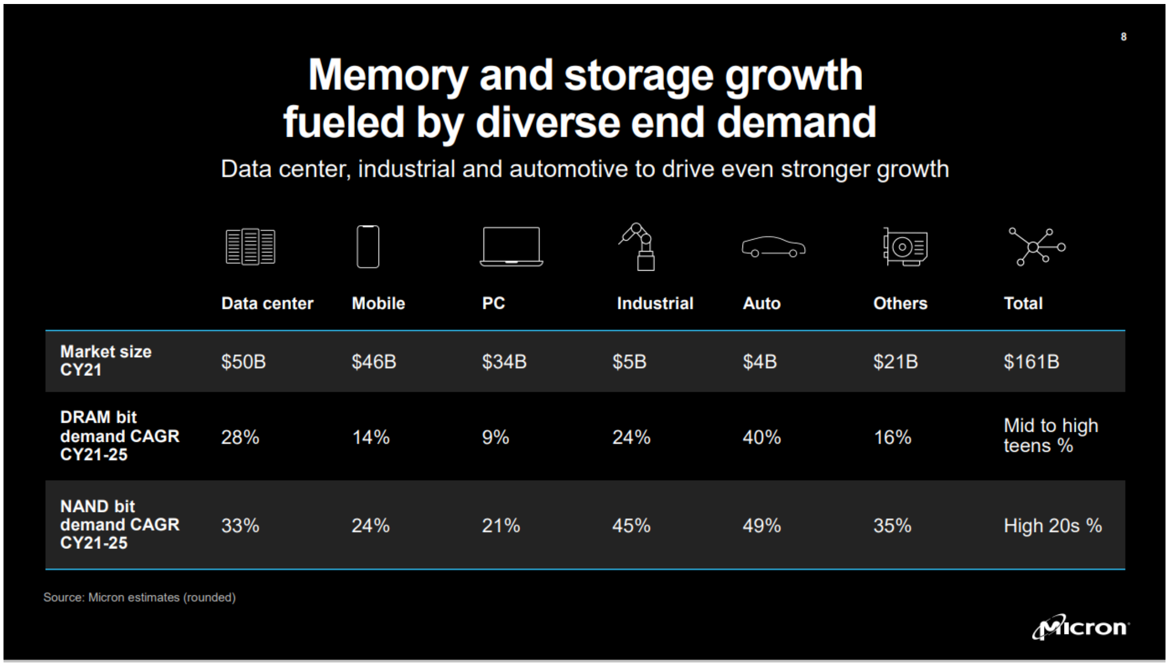

8

Memory and storage growth fueled by diverse end demand

Data center, industrial and automotive to drive even stronger growth

	Data center	Mobile	PC	Industrial	Auto	Others	Total
Market size CY21	$50B	$46B	$34B	$5B	$4B	$21B	$161B
DRAM bit demand CAGR CY21-25	28%	14%	9%	24%	40%	16%	Mid to high teens %
NAND bit demand CAGR CY21-25	33%	24%	21%	45%	49%	35%	High 20s %

Source: Micron estimates (rounded)

2022 Micron Analyst Day Presentation on NAND (and DRAM) Market Size and Growth Rate, by sector.

On the investment front, according to a 2018 presentation by Micron, successive generations of NAND technology is requiring more capital investment and is providing lower percentage growth in capacity.

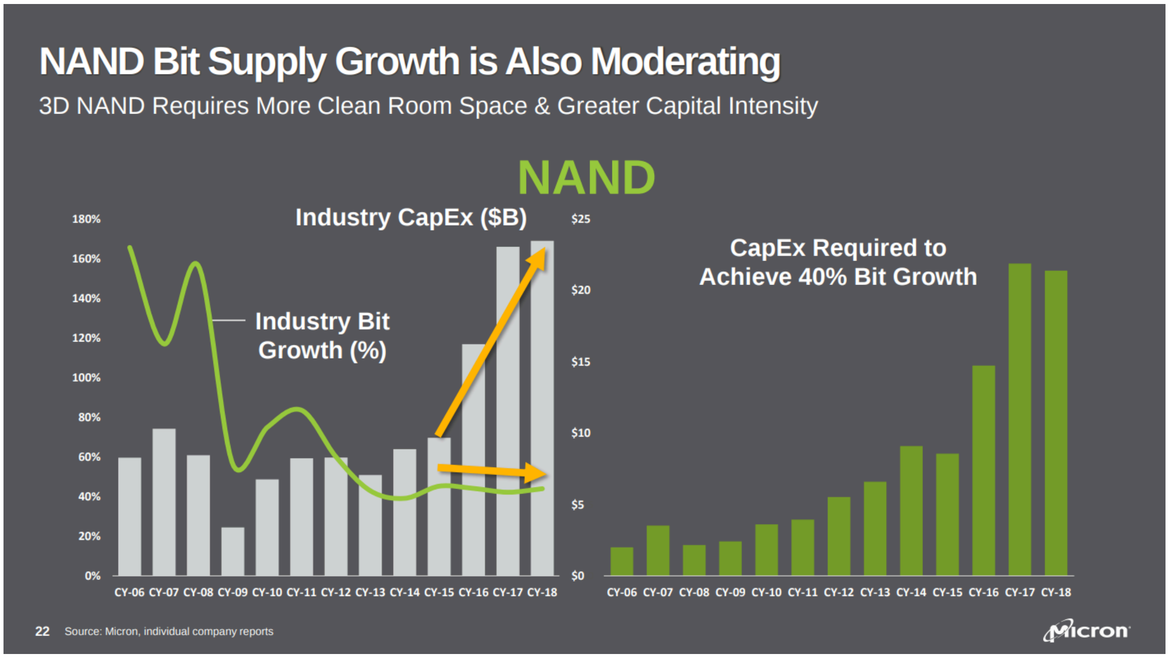

NAND Bit Supply Growth is Also Moderating

3D NAND Requires More Clean Room Space & Greater Capital Intensity

22 Source: Micron, individual company reports

[2018 Micron Investor/Analyst Day Presentation](#) on Capital Expenses for NAND Manufacturing

All Flash Array Market Status

All flash array 1Q2022 revenue by vendor from [blocksandfiles.com](#)

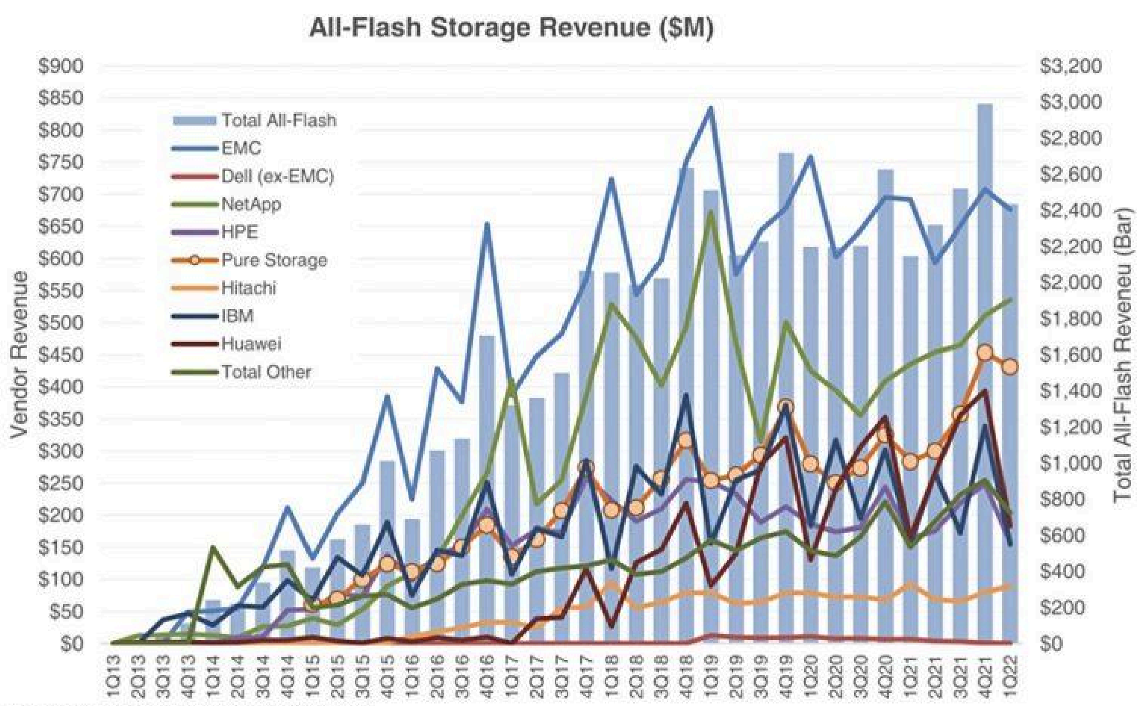

Source: Gartner; Wells Fargo Securities, LLC

(Updated 1/23) All Flash Array Storage Revenue through 1Q 2022 (Gartner via blocksandfiles.com)

All flash array 2Q2020 market share by vendor from blocksandfiles.com

- % market, units/PB shipped, annual growth rate
- pricing evolution, (discount rates)
- vendors and their status and market share, sales, net revenues
- impact on HEP

Form Factors

Solid State Disk (SSDs)

Solid State disks (SSDs) are non-volatile memory technologies (e.g., 3D-XPoint, NAND flash) that have been packaged in standard 3.5” and 2.5” magnetic hard disk form factors, as well as the more recent M.2 physical form factor. For the purposes of this document, the SSD designation also implies communication interfaces that are used for normal HDDs, namely:

1. Serial ATA (SATA)
2. Serial Attached SCSI (SAS)
3. Fibre Channel (?)
4. InfiniBand (running SCSI protocol on top)
5. Ethernet (running SCSI or FC on top, i.e. iSCSI, FCoE)

SSD’s benefit from all of the engineering work that has gone into SATA, SAS, and Fibre Channel, such as hot swap, boot support, and software stacks. With the advent of NVMe, the future of these “conventionally attached” SSD’s appears to be cloudy. A market brief by Horizon Technology states that sales of NVMe based SSD’s have exceeded sales of SATA/SAS connected SSD’s. This is reinforced in market data from IDC via the talk NVMe 2.0 Specifications: The Next Generation of NVMe Technology” by Peter Onufryk at the 2021 Storage Developer Conference.

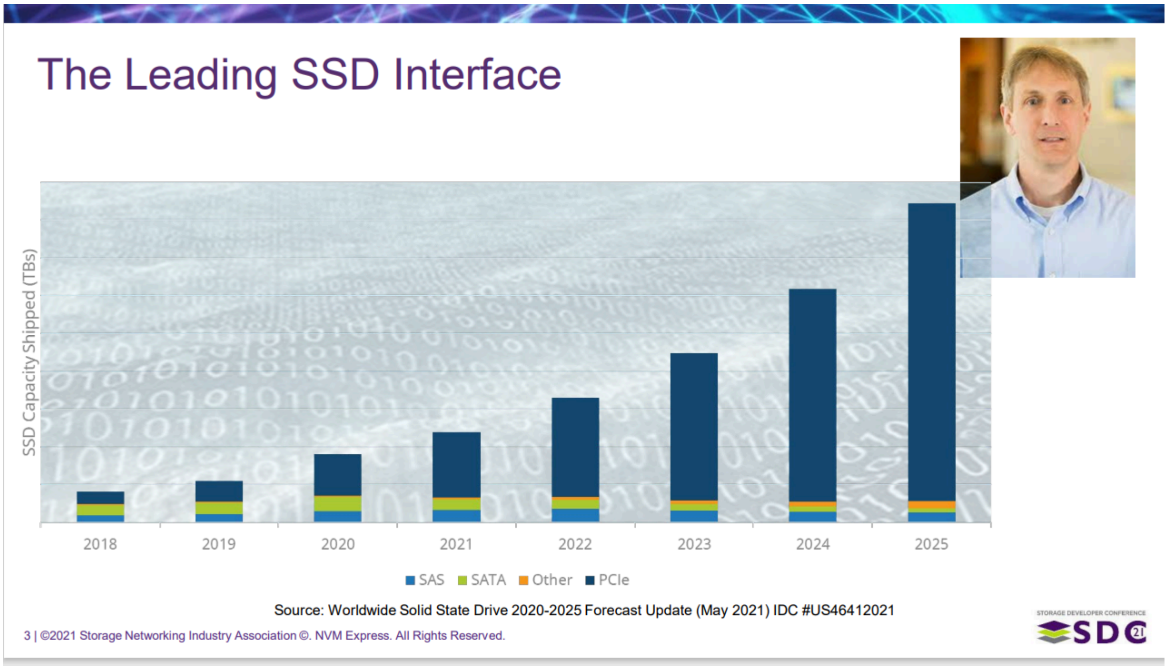

The Leading SSD Interface

Source: Worldwide Solid State Drive 2020-2025 Forecast Update (May 2021) IDC #US46412021

3 | ©2021 Storage Networking Industry Association ©. NVM Express. All Rights Reserved.

(Updated 1/23) Enterprise SSD Capacity shipped broken down by interface (IDC via presentation by Onufryk at the 2021 Storage Developer Con

This suggests that conventionally attached SSD's were a convenient transition technology that enabled the infiltration of NAND flash into the storage market, without new infrastructure and thus avoiding the time and costs associated with rolling out said infrastructure. Although not strictly related to SAS/SATA, today's systems deploying a larger number of SSDs connected through SAS/SATA encounter Linux driver limitations in terms of IOPS channeled to the physical SSDs. These limits might continue because of no/low interests in re-design the SCSI stack and a clear focus on the NVMe protocol to connect to high IOPS devices.

NVMe SSD

Although a variation of the above mentioned SSD's, Non Volatile Memory Express (NVMe) SSDs are sufficiently different that they warrant a separate discussion. The fundamental difference between NVMe "disks" and "regular" SSDs is that NVMe disks directly connect to the PCI-e bus (as transport) and NVMe is used as the access protocol (Wikipedia) substantially increasing bandwidth and reducing latency compared to SSDs connected through the traditional transports and protocols. These changes allow NVMe devices to more fully utilize the capabilities of their underlying non-volatile memory technology. However, by abandoning traditional storage interfaces, a substantial amount of development is needed to reclaim some of the capabilities provided by these interfaces, e.g., hot swap, enclosure services, etc. (addressed by U.2 standard and SES over NVMe).

The following graph from the Stanford EE380 EECS Colloquium presentation by Rick Coulson of Intel shows the benefits of NVMe attached SSDs relative to legacy attached SSDs.

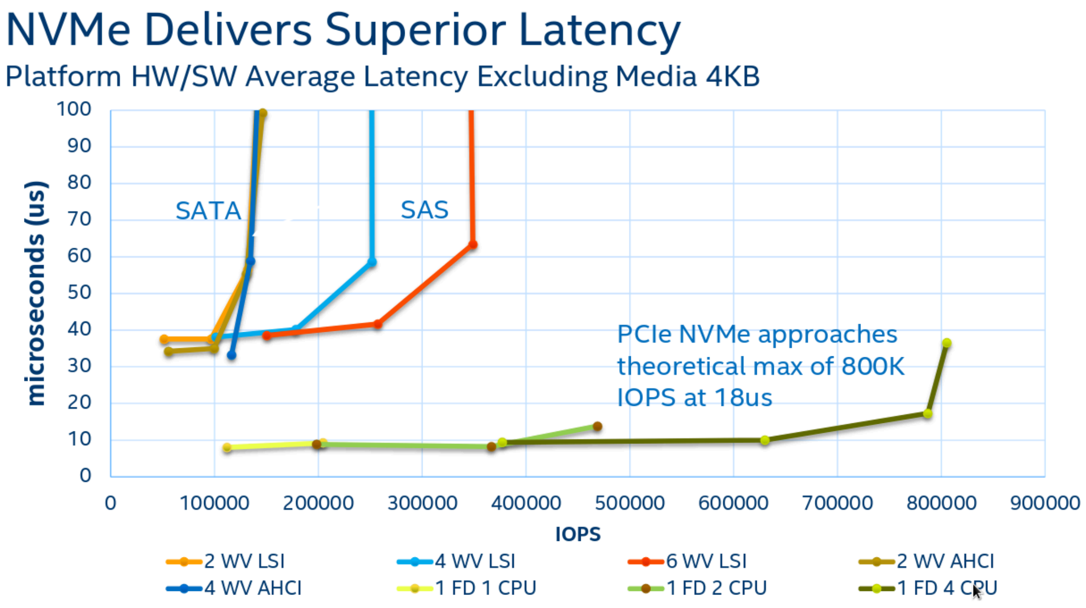

Persistent Memory Technology (NVDIMM)

Persistent memory is non-volatile memory that is directly connected to the memory bus on the CPU. This memory falls somewhere in between DRAM/

Interconnect

1. Interconnect Technology
    a. SATA/SATA Express
    b. SAS
    c. Fibre Channel
    d. PCI-e/NVMe/NVMeoF
    e. Infiniband
    f. Ethernet
    g. Physical Form Factors

SATA

Serial ATA (SATA) has been the interconnect of choice for consumer HDDs and looks to keep that position in the market. In the past it has been the interconnect of choice for SSDs, but consumer SSD are rapidly transitioning to NVMe in conjunction with the adoption of the M.2 form factor. PCI-e looks to be the interconnect of choice for SSD, excluding extremely cost sensitive and low performance applications. For these latter cases, SATA is likely to remain the interconnect of choice. The SATA standard continues to be updated, with the most recent version being SATA 3.5a; however, maximum link speed has remained at 6Gbps. The bulk of the updates to the SATA 3.x standard have been mostly feature additions, not performance enhancements.

At this point in time, SATA seems to be relegated to low cost/low performance HDD device connectivity, but even this position may be tenuous given the addition of support for HDDs in the latest NVMe revision, NVMe base specification 2.0c.

SAS

Serial Attached SCSI (SAS), has been the interconnect of choice for “enterprise” HDD’s. In recent years, SAS has branched out to HDD expansion chassis connectivity and direct attached storage connectivity (DAS), particularly for JBOD deployments. For enterprise HDD connectivity and HDD DAS connectivity SAS continues to be the interconnect of choice. For DAS, this is the case as SAS supports SAS, SATA, and PCI-e (through SAS Express). Market penetration for SAS connected SSD is unclear. The slide shown below, from 2019 Storage Developer Conference presentation by Cameron Brett of the SCSI Trade Association shows the historical and projected future market share for connectivity types to “enterprise” storage drives.\

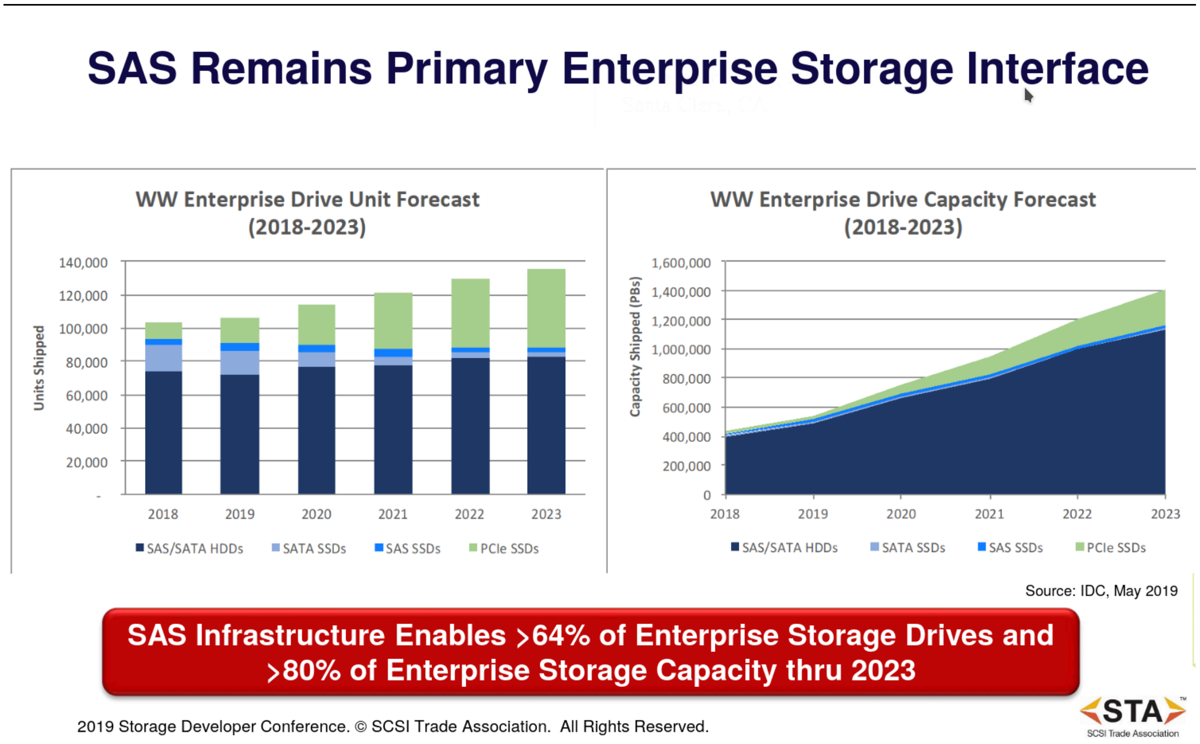

SAS Remains Primary Enterprise Storage Interface

Source: IDC, May 2019

SAS Infrastructure Enables >64% of Enterprise Storage Drives and >80% of Enterprise Storage Capacity thru 2023

2019 Storage Developer Conference. © SCSI Trade Association. All Rights Reserved.

SAS Express, SAS with PCI-e support appears to be a “transition” technology (with potentially limited life) for SSDs. The official SAS/SCSI roadmap was more recently updated in 2021. Underlying 24 Gbps SAS chip sets started appearing in 2019 and end user products started appearing in 2020. Moving forward, with the advent of NVMe, the open question is : What is the future of SAS ? The best guess at this point is that SAS will be relegated to “legacy” enterprise HDD connectivity and DAS connectivity, but like SATA, NVMe is looking to replace SAS as the storage interconnect standard of choice.

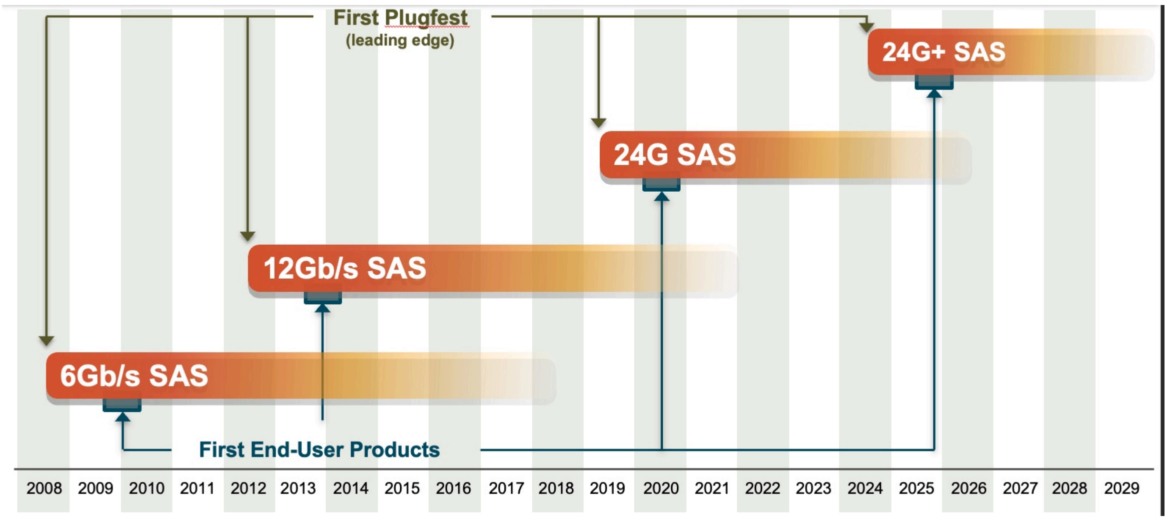

(Updated 1/23) SAS Roadmap v4 as of 9/2021.

Fibre Channel

The official Fibre Channel roadmap, now at version 23, was updated in 2020. Development and market availability of higher performance Fibre Channel versions (generations) have been pushed out by one to two years. The bandwidth chart is shown below

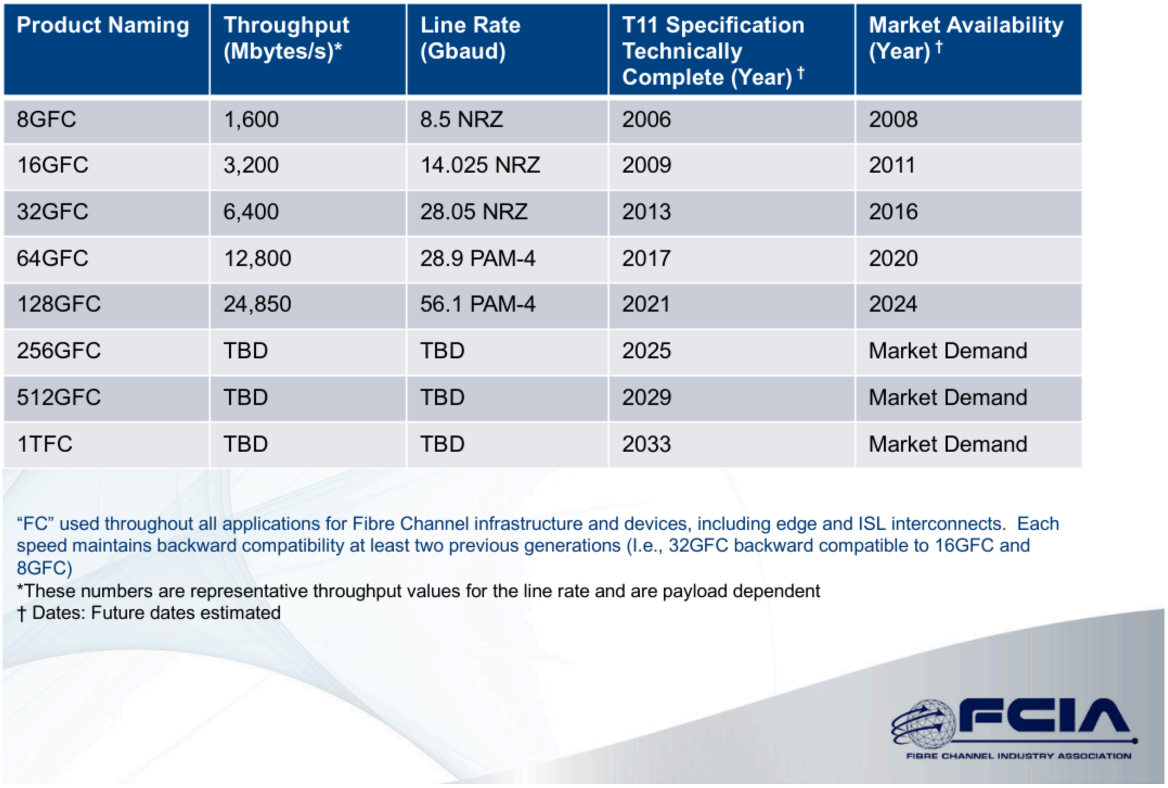

Product Naming	Throughput (Mbytes/s)*	Line Rate (Gbaud)	T11 Specification Technically Complete (Year) †	Market Availability (Year) †
8GFC	1,600	8.5 NRZ	2006	2008
16GFC	3,200	14.025 NRZ	2009	2011
32GFC	6,400	28.05 NRZ	2013	2016
64GFC	12,800	28.9 PAM-4	2017	2020
128GFC	24,850	56.1 PAM-4	2021	2024
256GFC	TBD	TBD	2025	Market Demand
512GFC	TBD	TBD	2029	Market Demand
1TFC	TBD	TBD	2033	Market Demand

“FC” used throughout all applications for Fibre Channel infrastructure and devices, including edge and ISL interconnects. Each speed maintains backward compatibility at least two previous generations (I.e., 32GFC backward compatible to 16GFC and 8GFC)

*These numbers are representative throughput values for the line rate and are payload dependent

† Dates: Future dates estimated

Fibre Channel Industry Association Fibre Channel Product Roadmap

The end of 2018 saw the introduction of 64G Fibre Channel (“Gen 7”) 64Gb chip and HBAs and 2019 saw the introduction of “Gen 7” switches. Bandwidth and connectivity options (single mode, multi mode, NRZ/PAM4) mirror those for Ethernet.

(Updated 1/23) [Fibre Channel Roadmap]() as of 2020

The predicted adoption cycles for the various generations of Fibre Channel from the Fibre Channel Industry Association are shown below:

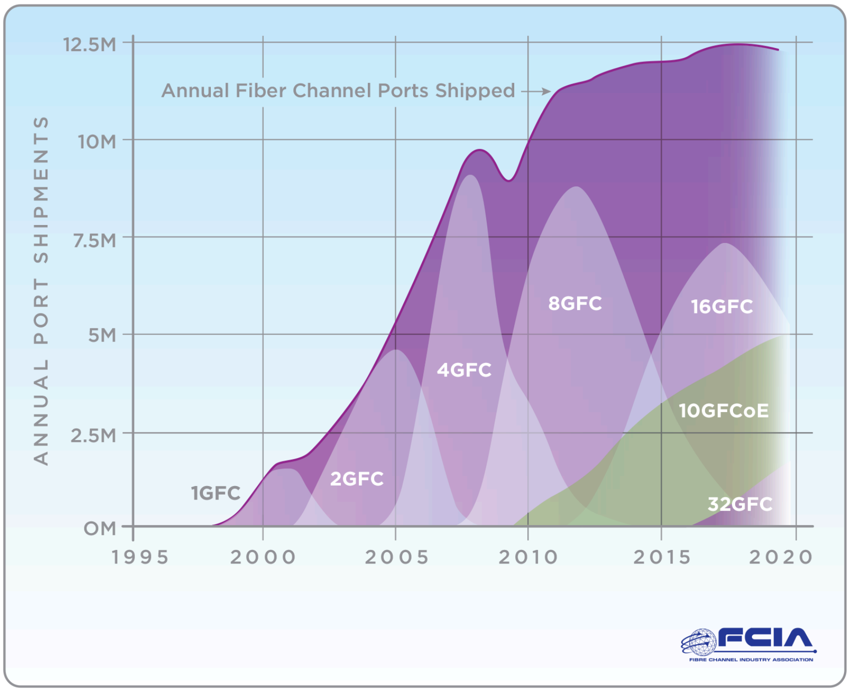

(Updated 1/23) Predicted Adoption Cycle for Fibre Channel Generations

With only three vendors, Broadcom (switches/chips/HBA), Cisco (switches), and Marvel (HBAs), the Fibre Channel business shows all the signs of a mature market and technology. Fibre Channel continues to be the interconnect of choice for Storage Area Networks and may yet have a role with NVMe-oF.

For HEP, Fibre Channel continues to be a critical interconnect technology, with FC adaptors being necessary within tape environments and represent a significant fraction of the tape mover cost.

NVMe/PCI-e/NVMe-oF

NVMe-oF storage connectivity - Future connectivity to “exotic”, high performance storage ?

The feature development roadmap for NVMe from the NVMExpress website is shown in the image below. More recent updates may be available from the upcoming 2020 Storage Developers Conference.

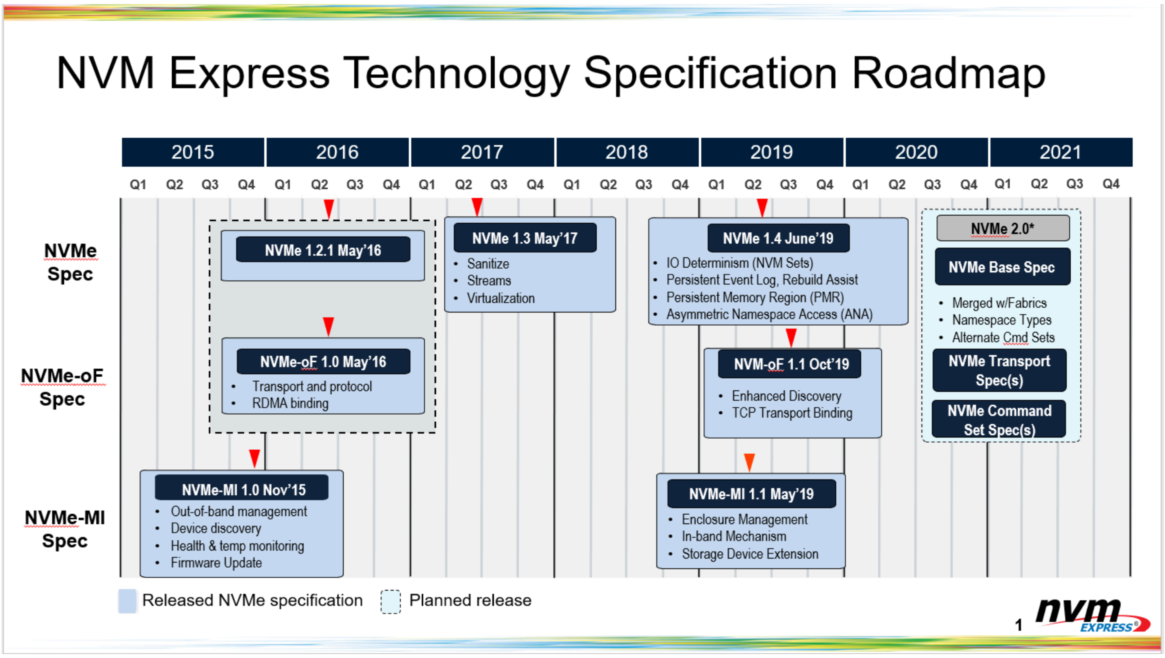

NVM Express Technology Specification Roadmap

NVMe 1.2.1 May'16

NVMe 1.3 May'17
- Sanitize
- Streams
- Virtualization

NVMe 1.4 June'19
- IO Determinism (NVM Sets)
- Persistent Event Log, Rebuild Assist
- Persistent Memory Region (PMR)
- Asymmetric Namespace Access (ANA)

NVMe 2.0*

NVMe Base Spec
- Merged w/Fabrics
- Namespace Types
- Alternate Cmd Sets

NVMe Transport Spec(s)

NVMe Command Set Spec(s)

NVMe-oF 1.0 May'16
- Transport and protocol
- RDMA binding

NVM-oF 1.1 Oct'19
- Enhanced Discovery
- TCP Transport Binding

NVMe-MI 1.0 Nov'15
- Out-of-band management
- Device discovery
- Health & temp monitoring
- Firmware Update

NVMe-MI 1.1 May'19
- Enclosure Management
- In-band Mechanism
- Storage Device Extension

▭ Released NVMe specification ⬚ Planned release

(Updated 1/23) NVM Express Technology Specification Roadmap

One of the goals of the NVMe developers is to have the NVMe software interface become the “lingua franca” of storage connectivity, regardless of the underlying physical layer, as shown in the diagram below from the SDC NVMe presentation.

<Prognostications about future of NVMe and Flash

https://www.forbes.com/sites/tomcoughlin/2020/05/15/nvme-for-all-data-center-storage/#56ed87077e05
https://blocksandfiles.com/2020/05/27/nvme-universal-block-storage-access-protocol/

Summary necessary>

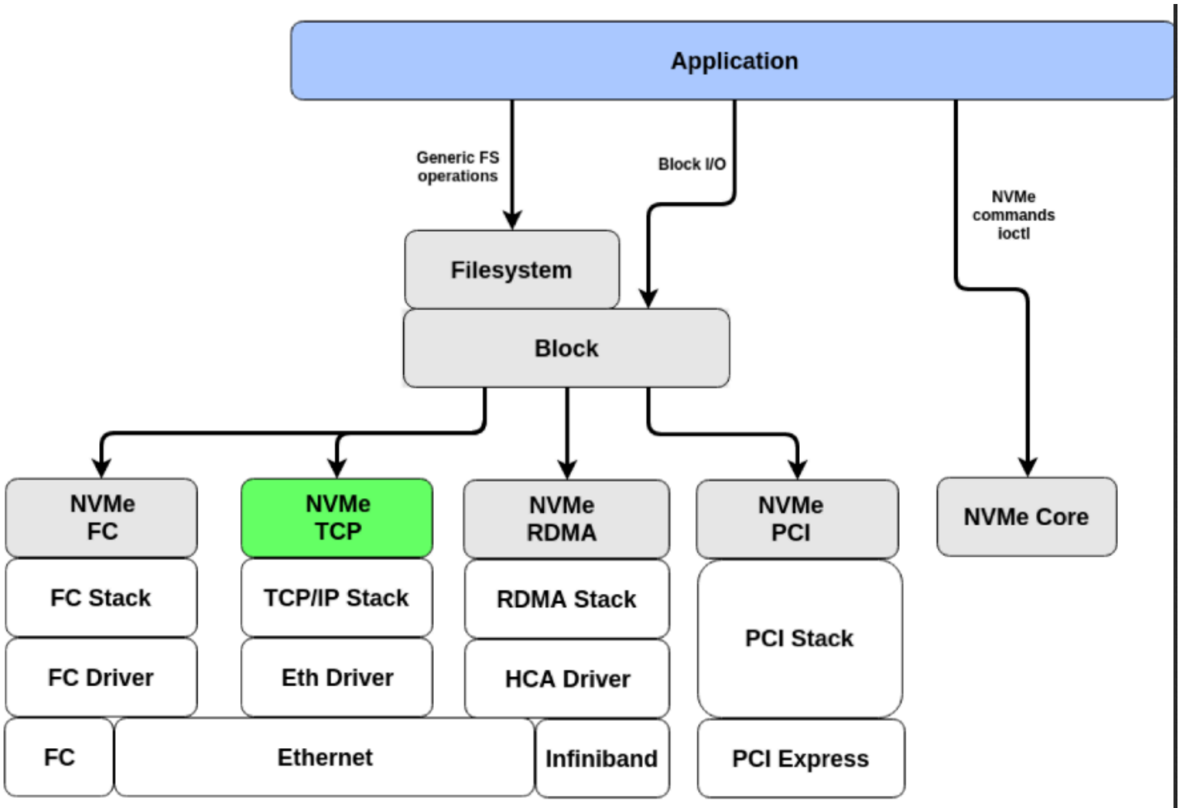

Source: Kam Eshghi (Lightbits Labs)

The following graph from a presentation made by G2M Research at the 2018 NVMe Developer Days conference shows the expected growth and segmentation of the NVMe storage market.

Infiniband

The major development in Infiniband with regards to storage is the development and adoption of GPU Direct.

Ethernet

As an alternative to FC, the newest IBM enterprise drives (TS1155, TS1160) support RoCE (RDMA over Converged Ethernet, Wikipedia) on Windows-based movers. LTO drives do not support RoCE. RoCE-based tape driver support for RedHat Linux and derivatives is not yet available (expected for RHES8 around Q3 2019).

Physical Form Factors

The rapid adoption of SSD in the marketplace has resulted in the development of new physical form factors beyond the HDD standard 2.5” form factor. NVMe SSD can now be found in the following form factors.

1. M.2
2. U.2
3. PCI-e cards (aka “AIC” or Add In Card)
4. EDSFF (Intel “ruler” physical form factor and interface)

The previously announced NGSFF/NF1 form factor from Samsung, a competitor to EDSFF, appears to have lost traction in the marketplace.

M.2

The M.2 form factor is found mostly in laptop computers and more recently on the motherboard of server and desktop computers. This is mostly due to its small physical size.16mm x 20mm, and 22mm x (30mm, 80mm, 110mm). The primary issue with M.2 is that it does not support hot swap and is designed for low power environments

U.2

The U.2 is defined by the PCI-e SFF-8639 module specification, but colloquially references an SSD with a standard 2.5” HDD form factor but with the U.2 connector.

EDSFF

EDSFF is a family of form factors targeted at two applications, E1 for 1U servers and E3 for 2U servers. These form factors break free from the standard HDD form factors, allowing for better thermal management. Both E1 and E2 provide for both short and long form factors, much in the way that PCIe cards come in full length and half length variants. The following two slides from D. Bae and J. Kim (Samsung) presentation “Challenges and Effects of EDSFF-based NVMe-oF Storage Solution” at the SNIA 201 Storage Developer Conference

EDSFF E1.X

- 1RU optimized, Offers various thickness

System Density ← → Power Capacity

Capacity ↓

E1.S up to 8 NAND Landing					
Size	31.5 x 111.49 x **5.9mm**	31.5 x 111.49 x **8.01mm**	33.75 x 118.75 x **9.5mm**	33.75 x 118.75 x **15mm**	33.75 x 118.75 x **25mm**
Recommended Power(W)	12W	16W	20W	20W	25W

E1.L up to 16 NAND Landing		
Size	38.4 x 318.75 **x 9.5mm**	38.4 x 318.75 x **18mm**
Recommended Power(W)	25W	40W

(Updated 1/23) EDSFF E1 Form Factors from Bae and Kim (Samsung)

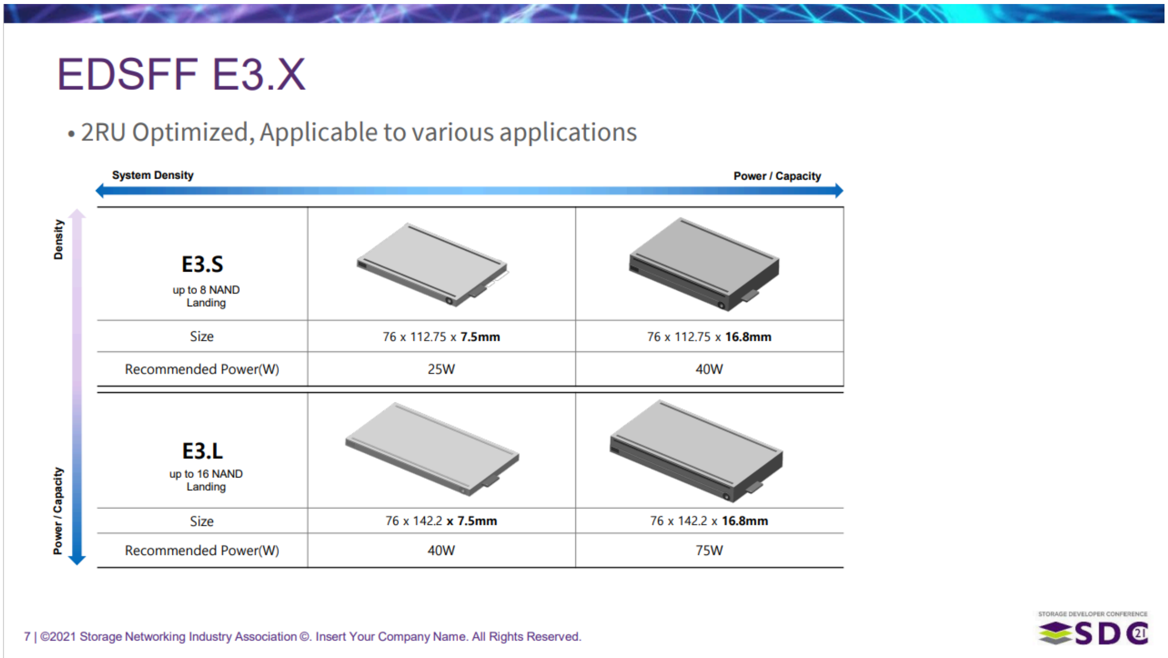

EDSFF E3.X

- 2RU Optimized, Applicable to various applications

System Density ← → Power / Capacity

Density ↕ Power / Capacity

E3.S up to 8 NAND Landing		
Size	76 x 112.75 x **7.5mm**	76 x 112.75 x **16.8mm**
Recommended Power(W)	25W	40W

E3.L up to 16 NAND Landing		
Size	76 x 142.2 **x 7.5mm**	76 x 142.2 x **16.8mm**
Recommended Power(W)	40W	75W

(Updated 1/23) EDSFF E1 Form Factors from Bae and Kim (Samsung)

Specifications for the M.2, U.2, and EDSFF form factor can be found at the SNIA Solid State Drive Form Factors web page.

Market Overview

Organization and Access

Advances in storage hardware is just one driver of improvements in capacity, cost, and performance of storage services. Two other drivers are advances in system architecture and improvements in software that take advantage of the new hardware.

Tape/HDD Competition

Overall, there is little competition in the tape market (media and drives) and concerns remain about the sustainability of a contracting market dominated by a single technology provider. Despite these issues, cost per TB for tape continues to drop and it still remains the most cost effective storage media for large data volumes.

The following graph from FujiFilms presentation at the 2022 Flash Memory Summit, using data derived from “The Escalating Challenge of Preserving Enterprise Data” report from Further Market Research..

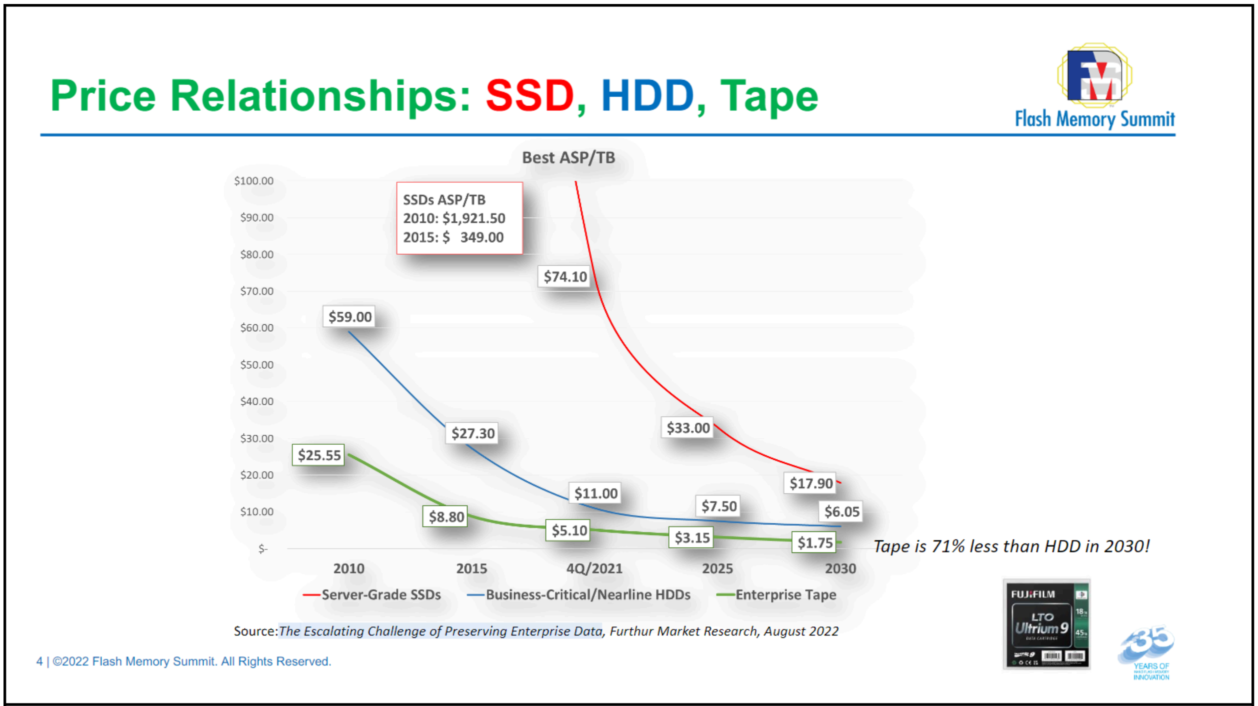

Price Relationships: SSD, HDD, Tape

Flash Memory Summit

Best ASP/TB

Tape is 71% less than HDD in 2030!

Source: *The Escalating Challenge of Preserving Enterprise Data, Furthur Market Research, August 2022*

(Updated 1/23) Acquisition cost per TB for SSD, HDD, and Tape media. (Further Market Research via FujiFilms presentation at the 2022 Flash Memory Summit)