Welcome to Flash Memory Summit 2019

Jenni Dietz
Co-Chair, SNIA SSSI

Advanced Computation and Storage, LLC

AGIGA TECH
A CYPRESS SEMICONDUCTOR COMPANY

AIC

AMD

SNIA. SSSI | SOLID STATE STORAGE

BROADCOM

CALYPSO Systems

Coughlin Associates

CRAY

Demartek

EIDETICOM

EXTEN TECHNOLOGIES

FADU

GIGABYTE

Hewlett Packard Enterprise

hyper I/O

IBM

inspur

intel

Intuitive Cognition Consulting

iol

kaminario.

Lenovo

Mellanox TECHNOLOGIES

Micron

msi

NetApp

NETINT

Objective Analysis Semiconductor Market Research

ORACLE

Quanta Computer

Rambus

SAMSUNG

SiliconMotion

SK hynix

SMART Modular Technologies

SPIN MEMORY

SunRise Memory Corporation

SUPERMICRO

TOSHIBA

TRENDFOCUS

viking TECHNOLOGY

WD Western Digital

XILINX

SNIA Solid State Reception Agenda

Monday August 5; 5:30 – 7:00 pm Room 209/210

SNIA SSSI | SOLID STATE STORAGE

- 5:40 - 5:50 pm **SNIA/SSSI Solid State and Persistent Memory Vision** **-** Jim Pappas, SNIA Vice Chairman
- 5:50 - 6:00 pm **PM&NVDIMM SIG Updates** - Jim Pappas
- 6:00 – 6:05 pm **SSD SIG updates –** Cameron Brett, SNIA SSD SIG Co-Chair
- 6:05 – 6:10 pm **Real World Workload Updates –** Eden Kim, SNIA S3 TWG Chair
- 6:10 - 6:15 pm **Alliance Activities** – JEDEC, OFA, NVM Express
- 6:15 - 6:25 pm **New SNIA Activities – Computational Storage TWG –** *Scott Shadley and Nick Adams, Co-Chairs*
- 6:25 - 6:35 pm **Emerging Memory Market Update** – *Tom Coughlin, Coughlin Associates and Jim Handy, Objective Analysis*
- 6:35 – 6:45 pm **Recap of SNIA & SSSI Activities at Flash Memory Summit and SSD Opportunity Drawing**
- 6:45 - ? **Networking**

© 2019 Storage Networking Industry Association. All Rights Reserved.

SNIA SSSI | SOLID STATE STORAGE

Solid State and Persistent Memory Vision

Persistent Memory and NVDIMM SIG Update

PM and NVDIMM SIG Highlights

- Transitioned from ecosystem development to applications
- Driving ISV adoption of Persistent Memory and NVDIMMs through targeted events and webinars
- Supporting Hackathons to develop and promote Persistent Memory and NVDIMM use cases
- Launching NVDIMM Programming Contest
- Promoting NVDIMM standardization efforts through JEDEC
- Working with the TCG and JEDEC on encryption standards

© 2019 Storage Networking Industry Association. All Rights Reserved.

NVDIMM Programming Contest

- **To drive additional development of NVDIMM-aware programs**
  - Provide systems online with SNIA branding and marketing support
  - Encourage programming and experimentation
  - Highlight success of participants
  - Launching after FMS

- **Support and focus from NVDIMM providers on software**
  - Additional content and online support for software
  - Highlight available tools
  - Evaluate scaling the program to other geographies

© 2019 Storage Networking Industry Association. All Rights Reserved.

PM and NVDIMM SIG Related Events

- ☑ Webcast: "Applications Take Advantage of PM", 1/15
- ☑ PM Summit/Hackathon, 1/24, Hyatt Santa Clara
- ☑ NVM Workshop/Hackathon, 3/10-12, UCSD
- ☑ Container World, 4/17-19, Santa Clara CC
- ☑ Open Call, Computational Storage, 5/10
- ☑ PIRL'19, 7/22-23, UCSD
- ☑ DOIT Hackathon, 7/22, Xi'an
- ☐ FMS, 8/5-8, Santa Clara CC
- ☐ SDC, 9/23-26, Hyatt Santa Clara
- ☐ JEDEC Workshop, 10/7-10, Santa Clara Marriott
- ☐ In-Memory Computing Summit, 11/13-14, Hyatt SFO
- ☐ PM Summit, 1/23/2020, Hyatt Santa Clara

© 2019 Storage Networking Industry Association. All Rights Reserved.

SSD SIG Update

Solid State Storage SIG

2019 Activities

- Revise charter and set 2019 goals
  - Become the industry authority on SSDs: technology, form factors, interfaces, nomenclature
    - With ties to NVMe and STA, we have first hand information and can influence our associations
  - Gain participation from all SIG members; don’t want observers
  - Update and change organization of form factors page
    - Breakdown tasks into small, bite sized tasks that everyone in the SIG can tackle
  - Re-use/re-purpose material from SIG member companies to become SIG content

© 2019 Storage Networking Industry Association. All Rights Reserved.

Real World Workloads Update

Join the Real World Workload Revolution!

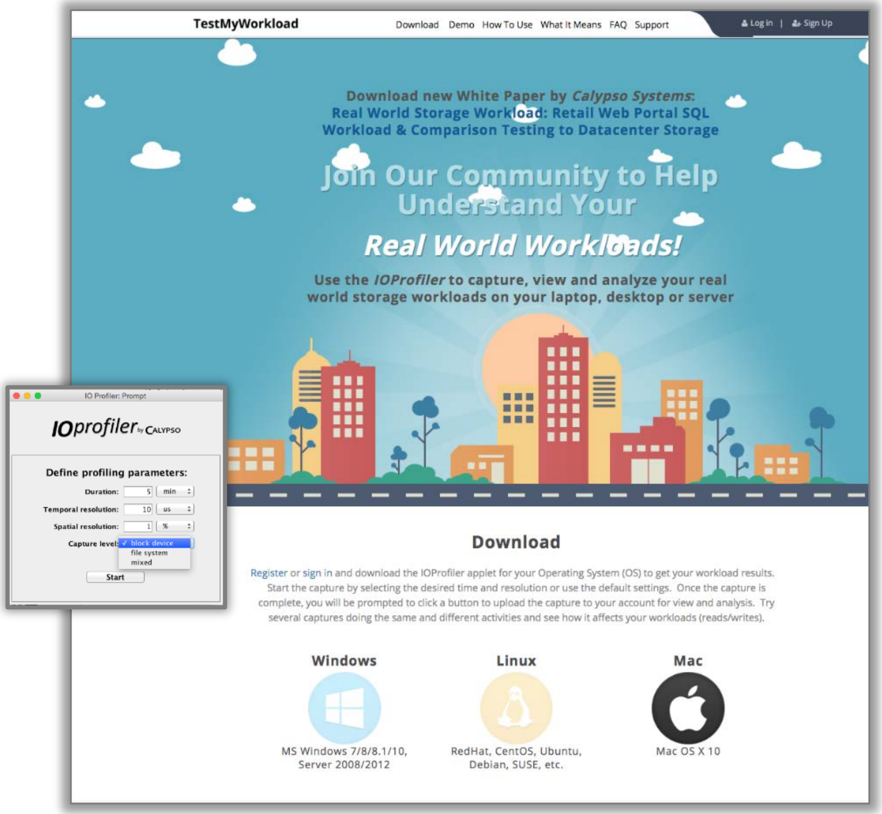

SNIA SSSI: TestMyWorkload.com

- Official SNIA SSSI IO Capture repository for IO Captures with Free IO Capture tools, reference workloads and user captures with data analytics
- Easy Capture of Real World Workloads
- IO Captures are portable, binary Tables of IO Stream statistics. No Personal or Actual User Data is captured
- IO Capture step resolution from uSec for fine grain IO analysis to long term workload capture (24 hours or longer)

© 2019 Storage Networking Industry Association. All Rights Reserved.

Datacenter & Web Portal Workloads

SNIA. SSSI | SOLID STATE STORAGE

Comparing Real World Workloads

1. Retail Web Portal: Drive0/Drive1

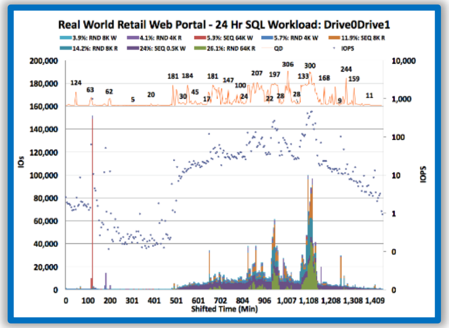

9 IO Streams = 71%
66:34 RW Mix
Retail Store Events

Ave QD = 22
Median QD = 19
Max QD = 306

2. GPS Nav Portal: Drive C

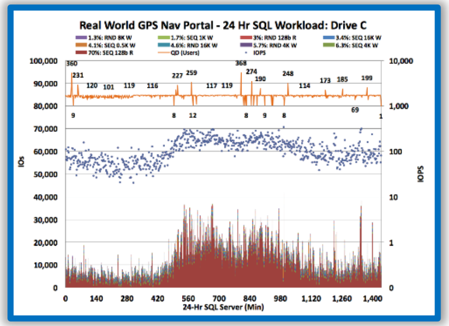

9 IO Streams = 86%
67:33 RW Mix
SEQ 128b R = 70% IOs

Ave QD = 114
Median QD = 115
Max QD = 368

3. GPS Nav Portal: Drive 0

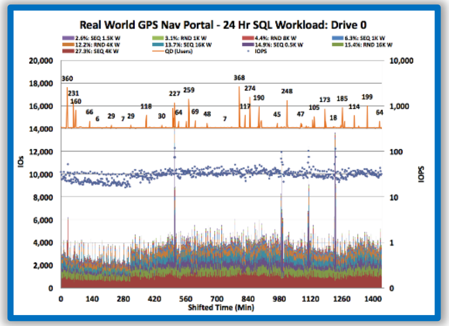

9 IO Streams = 78%
96% Write
SEQ 0.5K W Spikes

Ave QD = 15
Median QD = 8
Max QD = 368

Real World Workloads:
Constantly Changing Combinations of IO Streams & Users (QDs)

© 2019 Storage Networking Industry Association. All Rights Reserved.

How to be a Part of the Real World Workload Revolution

https://www.snia.org/sites/default/files/SSSI/SNIA_SSSI_Workload_Capture_Whitepaper_v1.0.pdf

SNIA. SSSI | SOLID STATE STORAGE

Home » Technology Communities » Solid State Storage Initiative » Knowledge Center » White Papers

White Papers

Persistent Memory White Papers

- SNIA Technical White Papers on Persistent Memory
- SNIA NVM PM Remote Access For High Availability Technical White Paper - May 2019
- SNIA Persistent Memory Atomics and Transactions
- NVDIMM Messaging and FAQ
- NVDIMM Technical Brief
- NVDIMM - Fastest Tier in Your Storage Strategy
- Non-Volatile Memory and Its Use in Enterprise Applications
- NVDIMM: Enabling Greater ROI from SSDs

Solid State Storage Performance

- How To Be a Part of the Real World Workload Revolution
- How To Be a Part of the Real World Workload Revolution (Chinese translation)
- Data Center Real World Storage Workloads
- Survey Update: Users Share Their Storage Performance Needs
- Understanding Datacentre Workload Quality of Service
- Survey: Storage Performance Needs
- The PTS User Guide
- Workload I/O Capture Program FAQ
- Workload I/O Capture Program
- SSD Performance - A Primer
- Foreword to Understanding SSD Performance Using the SNIA SSS PTS
- SSS PTS Case Study
- Solid State Storage Performance Test Specification (SSS PTS) White Paper
- Storage Performance Benchmarking Guidelines – Part 1: Workload Design
- Understanding SSD Performance Using the SSS Performance Test Specification

Solid State Storage Applications

- Flash Memory Enables Media and Entertainment Workflows
- Flash Memory in Media & Entertainment
- Two May be Better than One; Why HDDs & Flash Belong Together

Solid State Storage Essentials

- How Controllers Maximize SSD Life
- SSDs-What's Important to You?
- NAND Flash Solid State Storage–An In-depth Look at Reliability
- Solid State Storage 101
- Total Cost of Solid State Storage Ownership

SNIA. SSSI | SOLID STATE STORAGE

How To Be a Part of the Real-World Workload Revolution

Capture and Analyze Your Own Application Workload

Introduction by Tom Coughlin and Jim Handy
Authors: Eden Kim, Calypso Systems
Jim Fister, The Decision Place

© 2019 Storage Networking Industry Association. All Rights Reserved.

Alliances

JEDEC®

OPENFABRICS ALLIANCE

nvm EXPRESS™

DMTF®

© 2019 Storage Networking Industry Association. All Rights Reserved.

SNIA | COMPUTATIONAL STORAGE

What happens when Compute meets Storage?

Scott Shadley, Co-Chair, NGD Systems
Nick Adams, Co-Chair, Intel

August 2019

Compute, Meet Data

- Based on the premise that storage capacity is growing, but **storage architecture has remained mostly unchanged** dating back to pre-tape and floppy…

- How would you define changes to take advantage of Compute at Data?

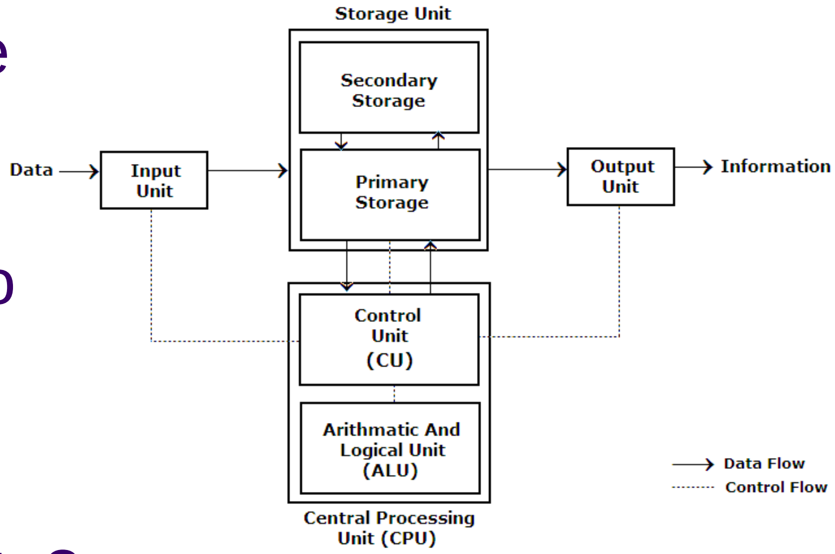

© 2019 Storage Networking Industry Association. All Rights Reserved.

The Evolution of Computational Storage

SNIA COMPUTATIONAL STORAGE

- A delicate process to build an Ecosystem
- Great ideas! Time was needed to build it
  - Many technology papers exist around:
    - “Active Disks”, “CAFS”, “Near-Data”
    - “In-Storage”, “In-Situ”, “Near-Storage”
- So did some initial products!

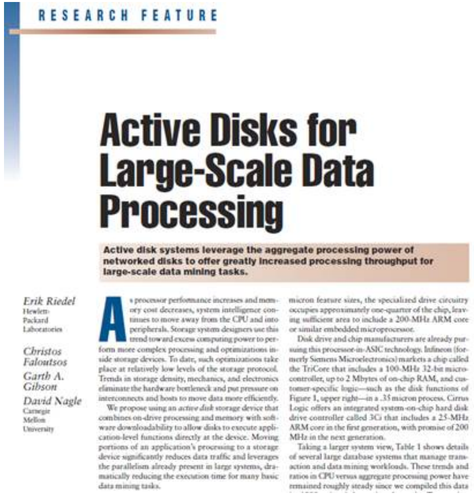

RESEARCH FEATURE

Active Disks for Large-Scale Data Processing

Active disk systems leverage the aggregate processing power of networked disks to offer greatly increased processing throughput for large-scale data mining tasks.

Erik Riedel
Hewlett-Packard Laboratories

Christos Faloutsos
Garth A. Gibson
David Nagle
Carnegie Mellon University

As processor performance increases and memory cost decreases, system intelligence continues to move away from the CPU and into peripherals. Storage system designers use this trend toward excess computing power to perform more complex processing and optimizations inside storage devices. To date, such optimizations take place at relatively low levels of the storage protocol. Trends in storage density, mechanics, and electronics eliminate the hardware bottleneck and put pressure on interconnects and hosts to move data more efficiently.

We propose using an *active disk* storage device that combines on-drive processing and memory with software downloadability to allow disks to execute application-level functions directly at the device. Moving portions of an application's processing to a storage device significantly reduces data traffic and leverages the parallelism already present in large systems, dramatically reducing the execution time for many basic data mining tasks.

micron feature sizes, the specialized drive circuitry occupies approximately one-quarter of the chip, leaving sufficient area to include a 200-MHz ARM core or similar embedded microprocessor.

Disk drive and chip manufacturers are already pursuing this processor-in-ASIC technology. Infineon (formerly Siemens Microelectronics) markets a chip called the TriCore that includes a 100-MHz 32-bit microcontroller, up to 2 Mbytes of on-chip RAM, and customer-specific logic—such as the disk functions of Figure 1, upper right—in a .35 micron process. Cirrus Logic offers an integrated system-on-chip hard disk drive controller called 3Ci that includes a 25-MHz ARM core in the first generation, with promise of 200 MHz in the next generation.

Taking a larger system view, Table 1 shows details of several large database systems that manage transaction and data mining workloads. These trends and ratios in CPU versus aggregate processing power have remained roughly steady since we compiled this data

© 2019 Storage Networking Industry Association. All Rights Reserved.

Playing Nice Together is Needed!

SNIA COMPUTATIONAL STORAGE

- Is this a solution replacing a solution?
- Complimentary work to the pyramid
- Another facet of advancement of compute
- In-Memory is needed, but some work can be offloaded all the way to storage!

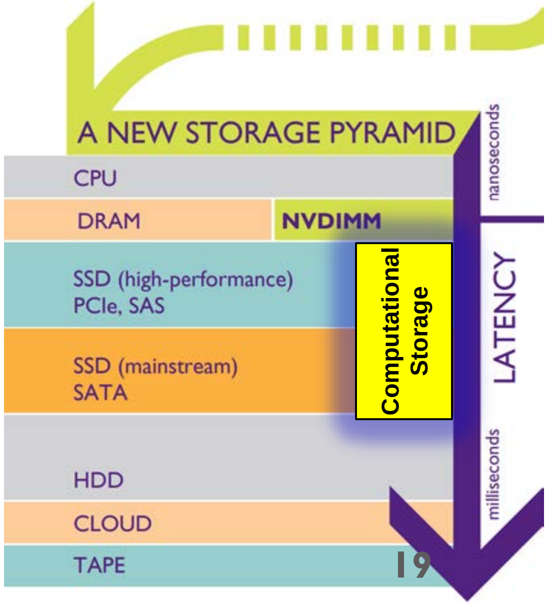

© 2019 Storage Networking Industry Association. All Rights Reserved.

So Now What?

The Progression of the TWG

40+ Participating Companies
148 Individual Members

SNIA COMPUTATIONAL STORAGE

© 2019 Storage Networking Industry Association. All Rights Reserved.

Finding a Focus and Direction

SNIA COMPUTATIONAL STORAGE

- Initial focus on a definition list to ensure we covered questions on what it is and what products can be

- Drive to a Scope and path to universal usage model
  - Today we have custom… Tomorrow Standard… Sound Familiar?

Starting the Standards Work

SNIA COMPUTATIONAL STORAGE

- Multiple F2F sessions have been focused on what we can accomplish and what we will leave for later

- Management
- Security
- Operation

Computational Storage TWG Dictionary Submissions

Computational Storage – Architectures that provide Computational Storage Services coupled to storage, offloading host processing or reducing data movement.

These architectures enable improvements in application performance and/or infrastructure efficiency through the integration of compute resources (outside of the traditional compute & memory architecture) either directly with storage or between the host and the storage. The goal of these architectures is to enable parallel computation and/or to alleviate constraints on existing compute, memory, storage, and I/O.

Computational Storage Service (CSS) – A data service or information service that performs computation on data where the service and the data are associated with a storage device.

The Computational Storage Service may be a Fixed Computational Storage Service or a Programmable Computational Storage Service.

Fixed Computational Storage Service (FCSS) – CSS that provides a given function that may be configured and used. (Service examples: compression, RAID, erasure coding, regular expression, encryption).

Programmable Computational Storage Service (PCSS) – CSS that is able to be programmed to provide one or more CSSes. (Service examples: this service may host an operating system image, container, Berkeley packet filter, FPGA bitstream).

Computational Storage Device (CSx): A Computational Storage Drive, Computational Storage Processor, or Computational Storage Array.

Computational Storage Drive (CSD): A storage element that provides Computational Storage Services and persistent data storage.

Computational Storage Processor (CSP): A component that provides Computational Storage Services for an associated storage system without providing persistent data storage.

Computational Storage Array (CSA): A collection of Computational Storage Devices, control software, and optional storage devices.

SNIA
Advancing storage & information technology

Computational Storage Architecture and Programming Model

Version 0.1 Revision 5

Abstract: This SNIA document defines recommended behavior for software supporting Non-Volatile Memory (NVM).

This Internal Use Draft is an internal document of the Computational Storage TWG that has not been approved for release outside of the membership of the Computational Storage TWG. This draft may not represent the position of the Computational Storage Technical Working Group.

Internal Draft

April 24th 2019
For SNIA Computational Storage TWG Internal Use Only

© 2019 Storage Networking Industry Association. All Rights Reserved.

Speaking the Same Language

- Computational Storage:
  - Architectures that provide Computational Storage Services coupled to storage offloading host processing and/or reducing data movement.

- Two Foundational Constructs
  - Computational Storage Devices (CSx)
  - Computational Storage Services (CSS)

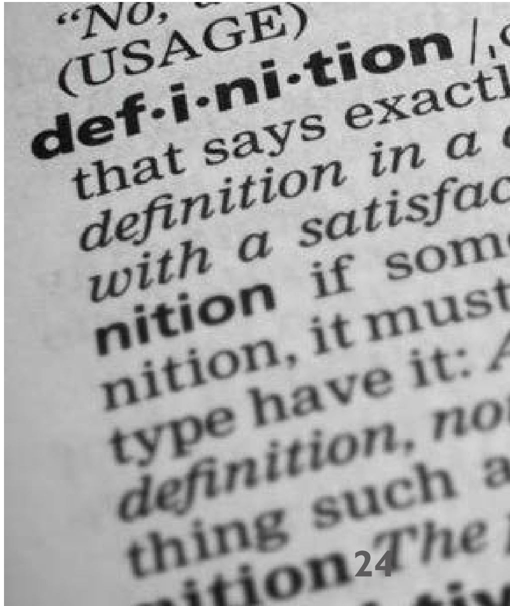

© 2019 Storage Networking Industry Association. All Rights Reserved.

Current Instances of Computational Storage

SNIA COMPUTATIONAL STORAGE

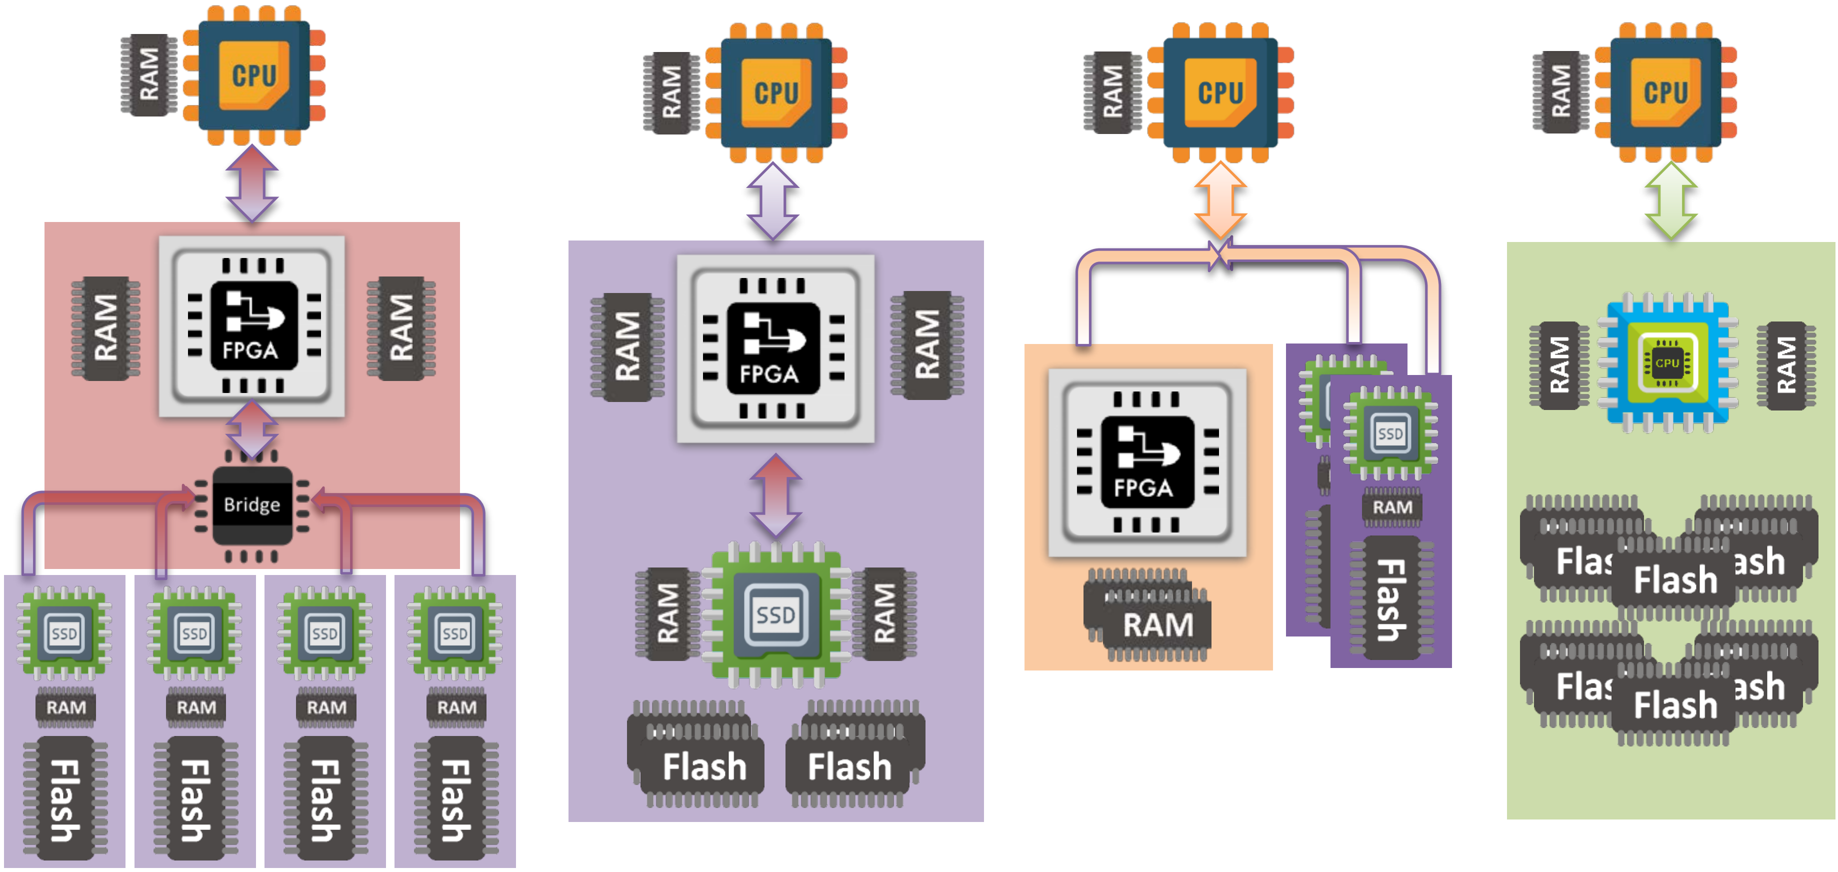

© 2019 Storage Networking Industry Association. All Rights Reserved.

Computational Storage Devices (CSx)

SNIA COMPUTATIONAL STORAGE

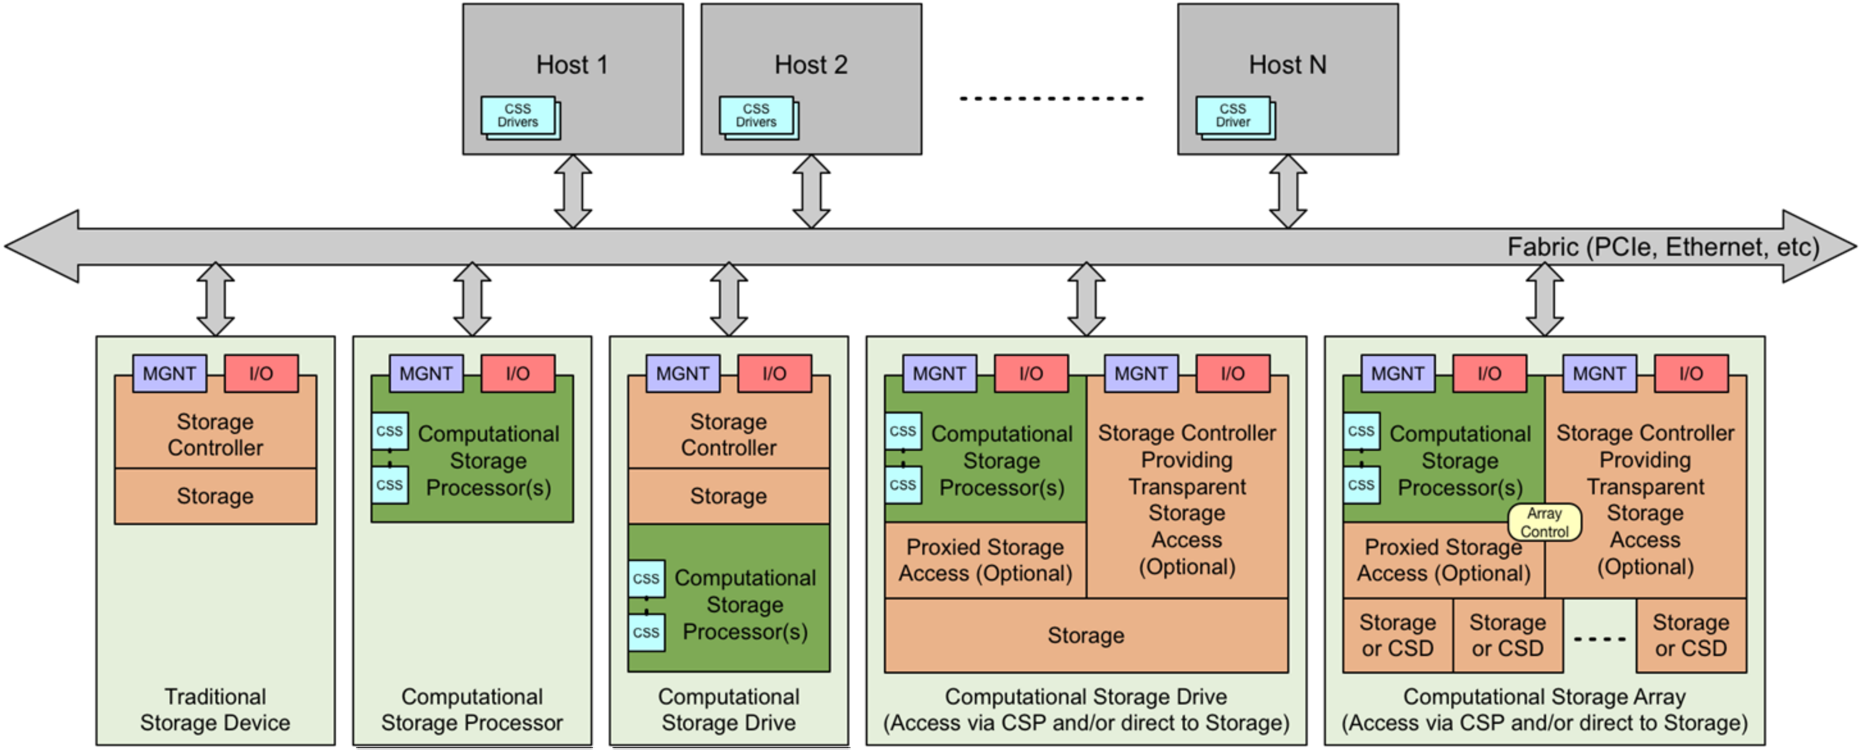

© 2019 Storage Networking Industry Association. All Rights Reserved.

Computational Storage Services

- Fixed Computational Storage Service (FCSS)
  - CSS that is well-defined
  - Consumable by the Host Agent for a well-defined purpose
  - Examples: Compression, RAID, Erasure Coding, or Encryption

- Programmable Computational Storage Service (PCSS)
  - Configured by the Host Agent to provide one or more CSSes
  - Examples: May host an Operating System image, Container, Berkeley Packet Filter, or FPGA Bitstream

© 2019 Storage Networking Industry Association. All Rights Reserved.

Define the Scope & Prioritize

SNIA COMPUTATIONAL STORAGE

- **Management**
  - **Discovery.** Identify and determine the capabilities and functions.
  - **Configuration**. Parameters for initialization, operation, and/or resource allocation
  - **Monitoring**. Reporting mechanisms for events and status
- **Security**
  - **Authentication**. Host Agent to CSx and CSx to Host Agent.
  - **Authorization**. Mechanism for secure data access and permissions control.
  - **Encryption**. Mechanisms to perform computation on encrypted data.
  - **Auditing**. Mechanisms to generate and retrieve a secure log.
- **Operation**
  - Mechanisms for the CSx to store and retrieve data.
  - Host Agent interaction may be explicit or transparent.

© 2019 Storage Networking Industry Association. All Rights Reserved.

In Summary – Call to Action

- Computational Storage is a Real Market
  - Customers are deploying today
- Solutions exist and will continue to grow
  - Making the interface ‘uniform’ helps adoption
- Standardizing the host interaction is vital
  - We NEED more Support from Users/SW Solutions
- Working across the industry will be crucial

© 2019 Storage Networking Industry Association. All Rights Reserved.

SNIA COMPUTATIONAL STORAGE

Thank You!!

www.SNIA.org/Computational

Emerging Memory Market Update

Tom Coughlin, Coughlin Associates
Jim Handy, Objective Analysis

5 August 2019

Who Wants Persistent Memory?

1) If it Costs MORE than DRAM

- That's NVDIMM-N!
  - High-availability systems
  - Financial databases
  - Some hyperscale applications
- MRAM DIMMs work, too

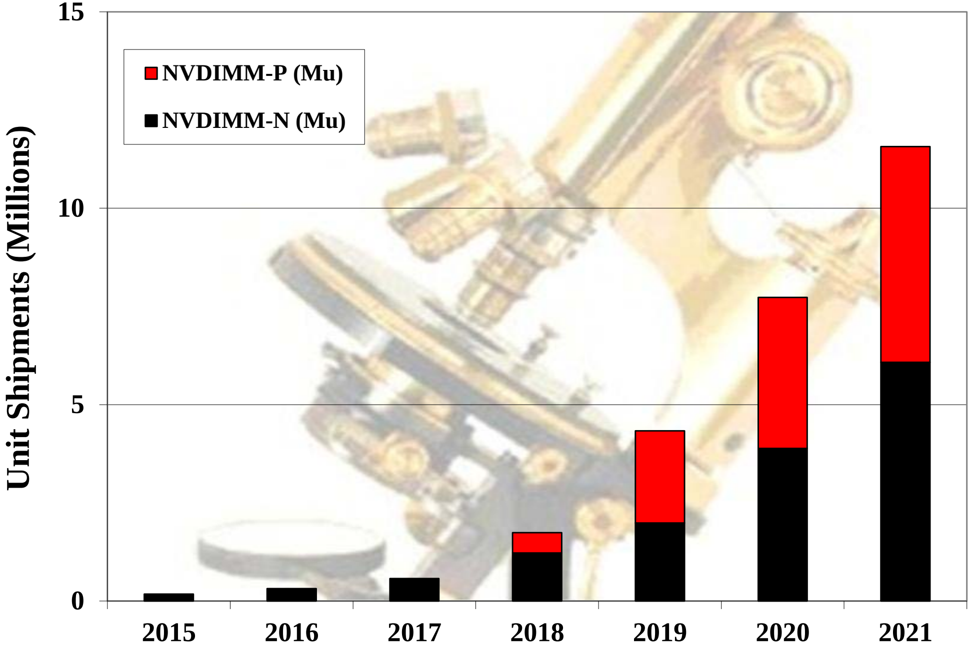

© 2019 Storage Networking Industry Association. All Rights Reserved.

Who Wants Persistent Memory?
2) If it Costs LESS than DRAM

- That's 3D XPoint!
  - Everybody will want it!
  - It's improves cost/performance
    - Persistence is of secondary importance
  - This will drive its success
    - Persistent apps come later

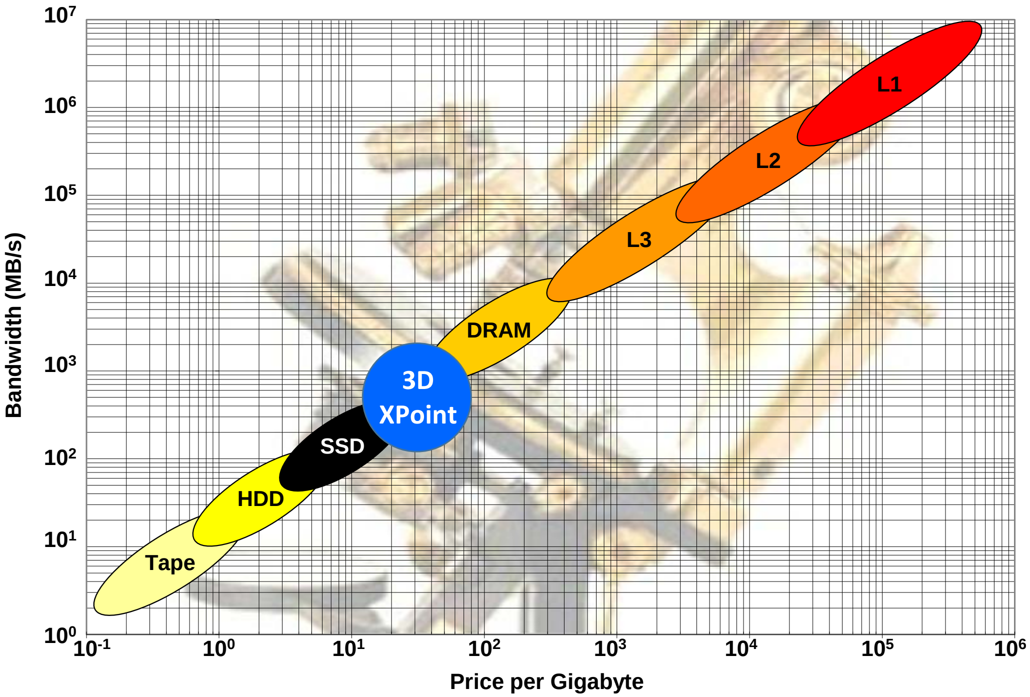

© 2019 Storage Networking Industry Association. All Rights Reserved.

Who Wants Persistent Memory?
3) If NOR becomes Unavailable in Foundries

- Below 28nm NOR stops scaling
  - Need an alternative NVM
    - MRAM & ReRAM viable candidates
- SRAM doesn't scale well
  - Emerging memories could displace SRAM caches
  - Caches could be the early adopter
    - What if caches become persistent???

© 2019 Storage Networking Industry Association. All Rights Reserved.

Who Wants Persistent Memory?
4) If the World Continues to Embrace AI

- Neural Nets Like PM
  - Inference Engines
  - Edge Applications
  - Vision Applications
- PM Provides a path to training
  - Requires linear storage
  - A big research focus

© 2019 Storage Networking Industry Association. All Rights Reserved.

Persistent Memory Types

MRAM

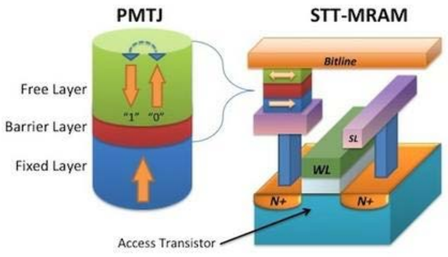

PCM

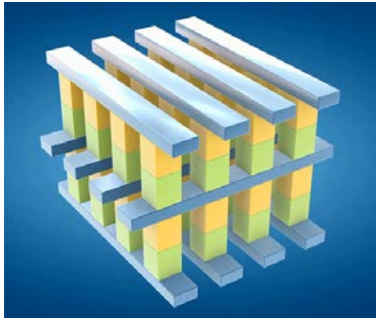

ReRAM

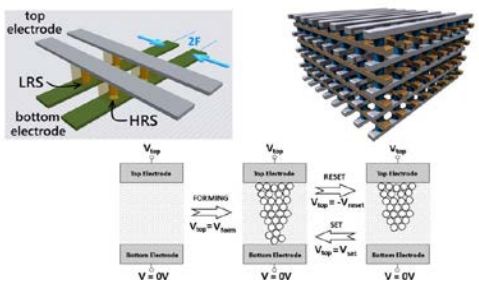

FRAM

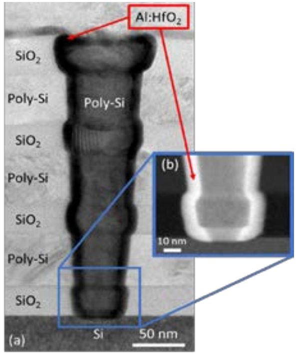

© 2019 Storage Networking Industry Association. All Rights Reserved.

Comparing the Technologies

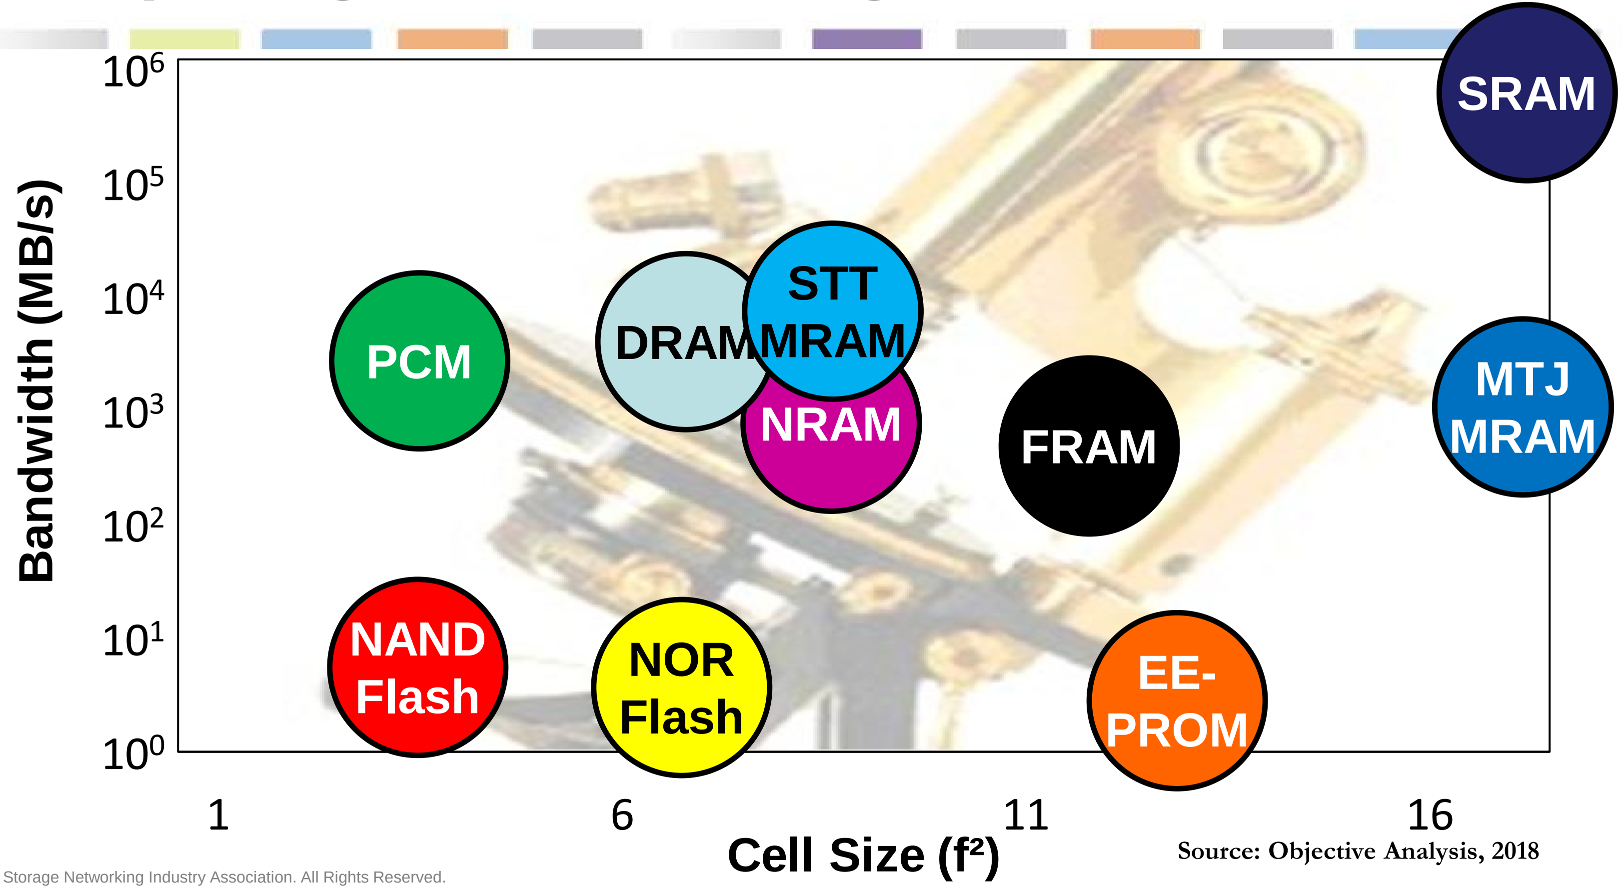

Source: Objective Analysis, 2018

© 2019 Storage Networking Industry Association. All Rights Reserved.

MRAM Developments

SNIA COMPUTATIONAL STORAGE

- All major foundries on-board:
  - Samsung, TSMC, GLOBALFOUNDRIES, UMC…
- Cost reduction efforts underway
  - Current: Back-end, between metal layers
  - Plan: Earlier in the process, deeply integrated
- New tools drive increased capital spending

© 2019 Storage Networking Industry Association. All Rights Reserved.

Growth in MRAM

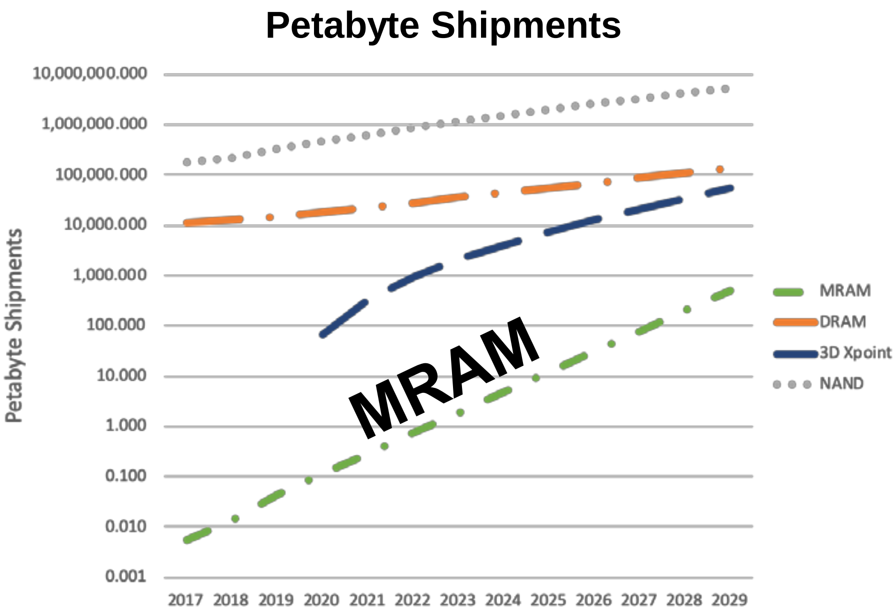

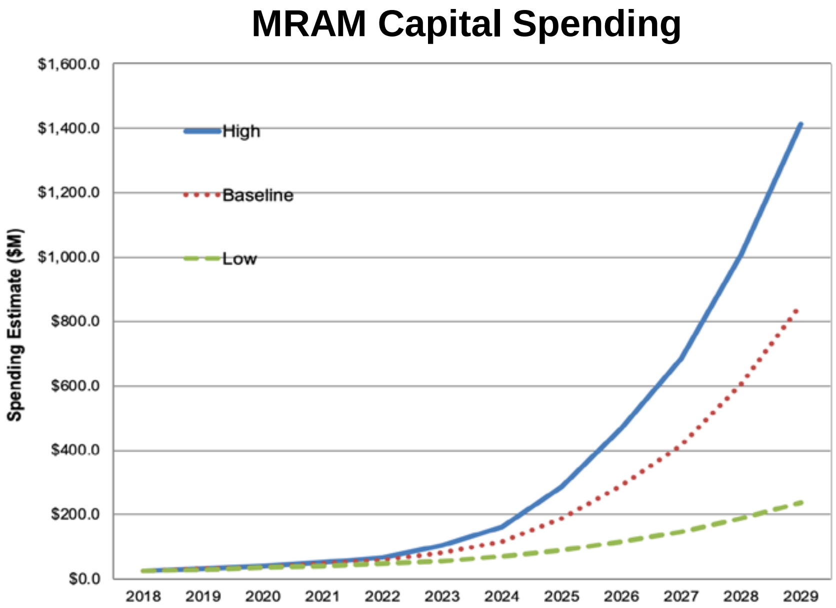

Source: ***Emerging Memories Ramp Up***, Coughlin Associates & Objective Analysis, 2019

© 2019 Storage Networking Industry Association. All Rights Reserved.

Emerging Memory Report

SNIA COMPUTATIONAL STORAGE

- Covers all major emerging memory technologies and companies
- Describes major driving applications
- Persistent memory forecasts (both embedded and stand-alone)
- Projections for capital investments
- **Now Available!**

https://tomcoughlin.com/tech-papers/

https://Objective-Analysis/reports/#Emerging

© 2019 Storage Networking Industry Association. All Rights Reserved.

Emerging Memory/AI Workshop

Stanford University, 29 August 2019

SNIA COMPUTATIONAL STORAGE

- One-day workshop
- All emerging memories:
  - MRAM, ReRAM, PCM, etc.
- Various AI applications
- Register at
  - https://EMAI19.sites.Stanford.edu

© 2019 Storage Networking Industry Association. All Rights Reserved.

Thank You!

Tom Coughlin
(408) 202-5098
Tom(at)tomcoughlin.com

Jim Handy
(408) 356-2549
Jim.Handy(at)Objective-Analysis.com

Coughlin Associates
Data Storage Consulting

OBJECTIVE ANALYSIS
Semiconductor Market Research

Upcoming Persistent Memory Events

Pick up your conference registration discount card at **SNIA booth 820**

Register at www.jedec.org

Complimentary registration now open at **snia.org/pm-summit**

Join SNIA at FMS 2019

Network with SNIA JEDEC OPENFABRICS ALLIANCE

Exhibit Hall Booth 803 (JEDEC) & 820 (SNIA & OFA)
Tuesday 4:00 pm – 7:00 pm
Wednesday 12:00 pm – 7:00 pm
Thursday 10:00 am – 2:30 pm

Beer & Pizza/Meet With the Experts
Tuesday 7:30 pm – 9:00 pm
Ballrooms A-C

Persistent Memory Track SNIA
Sponsored by SNIA, JEDEC, & OFA
Tuesday & Wednesday
8:30 am - 10:50 am & 3:40 pm – 6:00 pm

Computational Storage Track
Sponsored by SNIA
Thursday
8:30 am – 10:50 am & 2:10 pm – 5:00 pm

SNIA 2019 PERSISTENT MEMORY
HACKATHON
Tuesday and Wednesday
8:30 am to 7:00 pm
Great America Ballroom Lobby

SNIA | PERSISTENT MEMORY
Persistent Memory Meetup
Hosted by SNIA
Wednesday, 7:00 pm
Great America Ballroom K

SNIA Standards

Key Value Storage
NVME-201A-1 Wednesday 8:30 am GAMR 2

SNIA Swordfish™
SOFT-201A-1 Wednesday 8:30 am Ballroom F

Form Factors
SSDS-201B-1 Wednesday 9:45 am GAB J

Thank You!