Session R22 x103 Timeline

- High level timeline for sessionR22 x103

1. Introductions, Session overview & Enabling PM Ecosystem 10mins (10 Mins of 65mins)
2. Three Speakers for 15mins each = 45mins (55mins of 65mins)
3. About 10mins of Q&A (65mins of 65mins)

Emerging Solutions for Persistent Memory

Forum R-22 Session x-103

Speakers:	Cyril Gayot,	Western Digital
	Wendy Elsasser,	ARM
	Tom Friend,	Toshiba Memory
Chair:	Vikas Agrawal,	CachePhysics

Introductions

- Vikas Agrawal – Session Chair, Founder & CTO CachePhysics

- Vikas Agrawal is founder and CTO at CachePhysics, which is developing the intelligent self learning software for accelerating applications. Vikas is a 21-year industry veteran of the Memory and Storage industry with last 5 years driving strategy and technology leadership in office of CTO at Toshiba. Vikas has split his career in deep Engineering roles together with Product Management and strategy.

Introductions

- Wendy Elsasser – Principal Research Engineer, ARM

- Wendy Elsasser has 20 years of experience in design and architecture with a focus on memory sub-systems for the last 15 years. Wendy graduated from the Georgia Institute of Technology with a bachelors in Electrical Engineering and received a masters in Electrical Engineering from the University of Illinois. She has worked in the Research group at ARM for past 3 1/2 years, specifically working on future memory systems.

Introductions

- Tom Friend - Director, Industry Standards, Toshiba Memory America

- Responsible for a team of Technology Standards specialists with a broad range of backgrounds in Cloud Storage, SSD storage devices, Data Security, and Storage Class Memory. Standards committee engagements include T10 and T13, NVMe, PCI-SIG, SATA-IO, TCG, SFF, SNIA, OSF, IDC and JEDEC. His long background in hardware and software design has made him ideal in this challenging role.

Introductions

- Cyril Guyot - Director, Software Solutions and Algorithms, Western Digital

- Cyril's interests spans cryptography and practical security, information theory, machine learning and distributed systems. In Western Digital's Research organization, he has been leading the Software Solutions and Algorithms team in developing novel algorithms and implementations for storage systems and persistent memory architectures.

Session Overview & Agenda

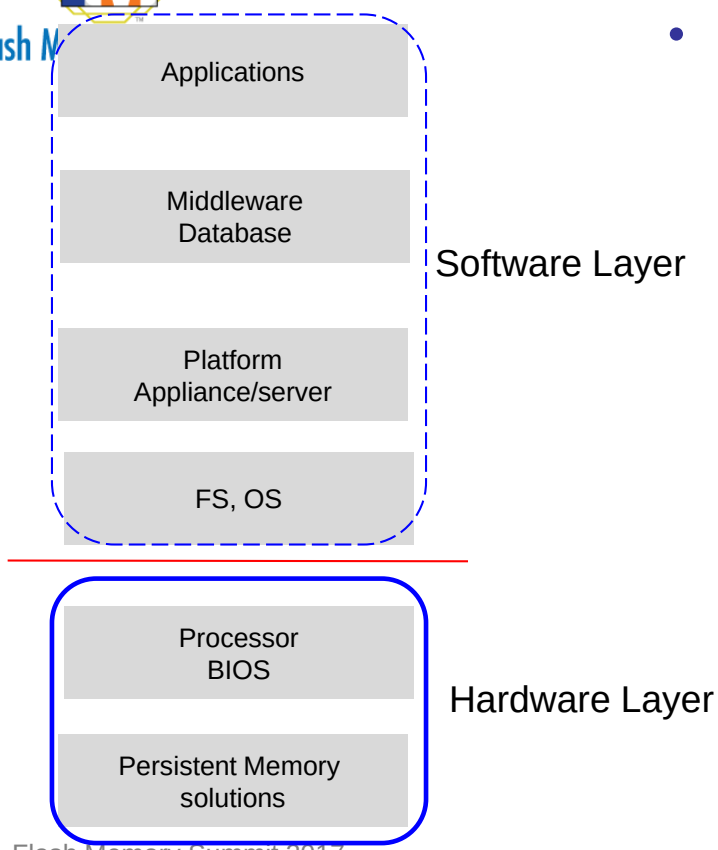

- Hardware ecosystem enabling persistent memory
  - Overview and Memory/Storage convergence – Vikas Agrawal
  - NVMe Vs NVDIMM-P implementation for persistent memory – Wendy Elsasser
  - Buses & Connectivity considerations for persistent memory – Tom Friend
  - Security considerations for persistent memory - Cyril Guyot

Flash Memory Summit

What is Persistent Memory??

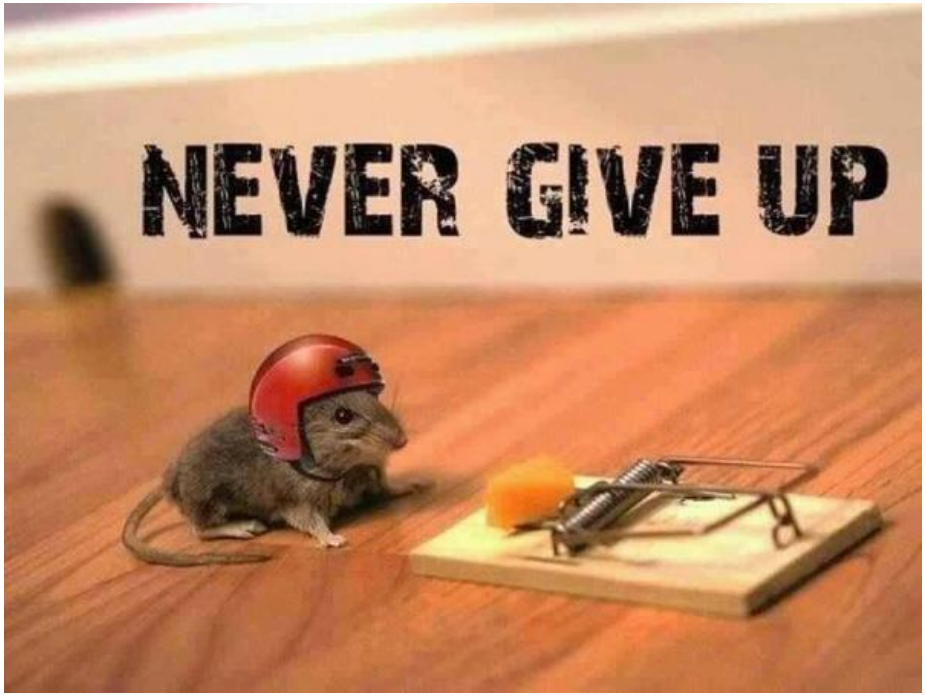

Memory and Storage Convergence

- **Volatile and non-volatile technologies are continuing to converge**

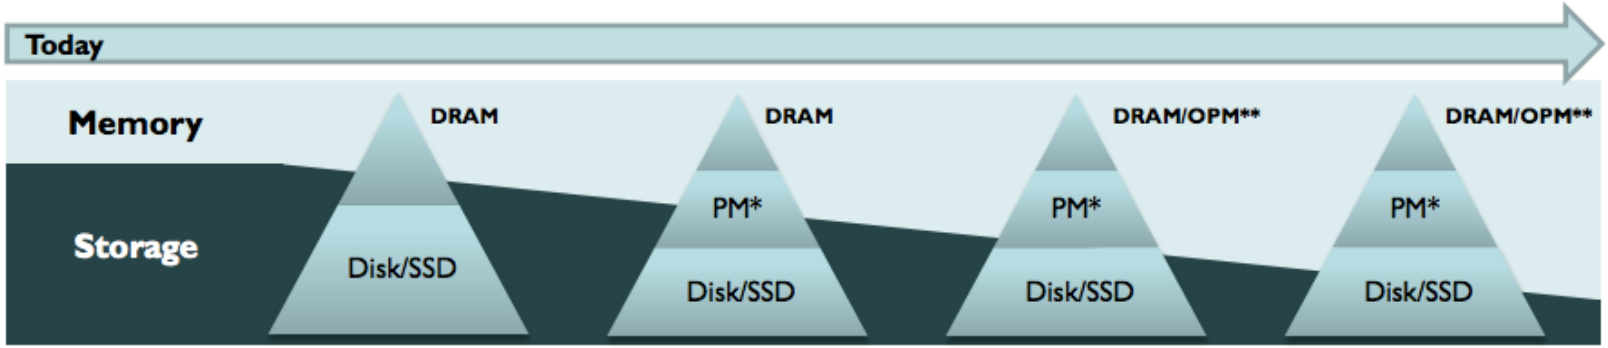

*PM = Persistent Memory

**OPM = On-Package Memory

New and Emerging Memory Technologies

HMC	3DXPoint™ Memory	Low Latency NAND
HBM	MRAM	
RRAM	PCM	Managed DRAM

Source: Gen--Z Consortium 2016

Revolutionary Hardware, Software and Algorithms arriving simultaneously

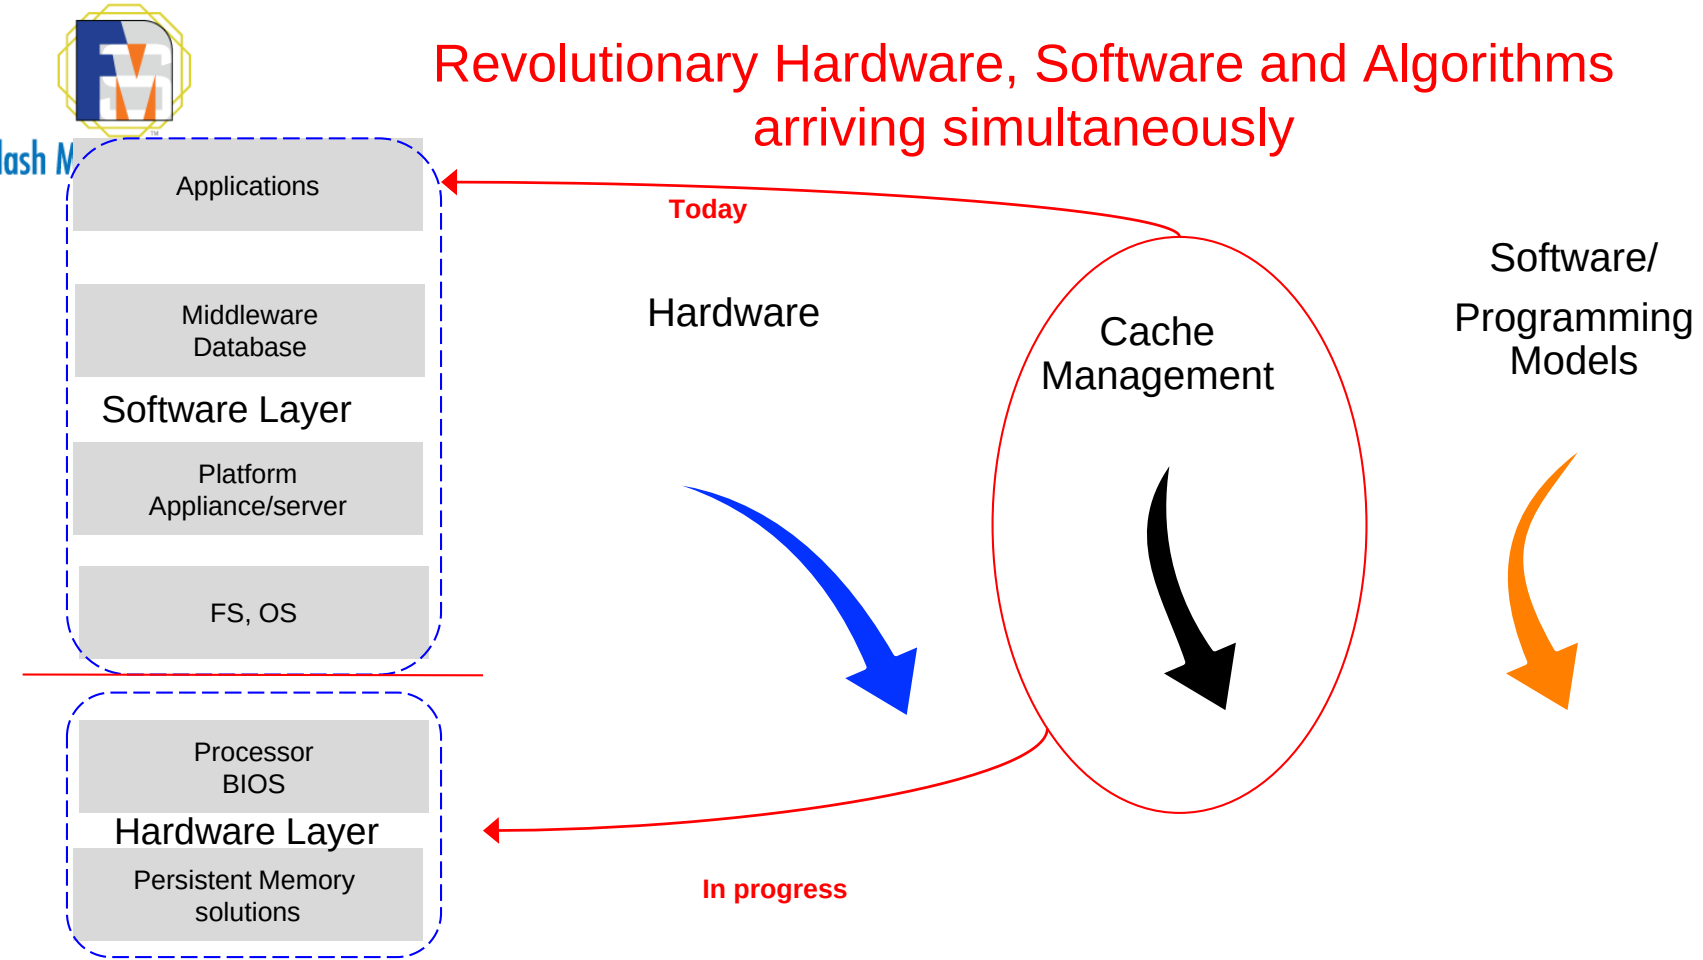

Groundbreaking Intelligent Caching Technology Now Ready to Deploy

- Just in last 5 years, many firsts demonstrated
- Examples
  - Online efficiency prediction (CachePhysics, Coho, Google)
  - Online non-disruptive cache tuning (MIT, CachePhysics)
  - Continuously self-optimizing cache algorithms (MIT)
  - Cache-Enabled QoS (latency guarantees) (FIU, CachePhysics)
  - Multi-tier QoS (CachePhysics)

Intelligent self-learning Caching Technology

- Opportunity
  - Hardware software partnership for unprecedented self-tuning, cost-performance-QoS aware systems

Emerging NVM Interfaces

Wendy Elsasser

arm

Incorporating NVM into your system

Numerous attachment points; varied media characteristics

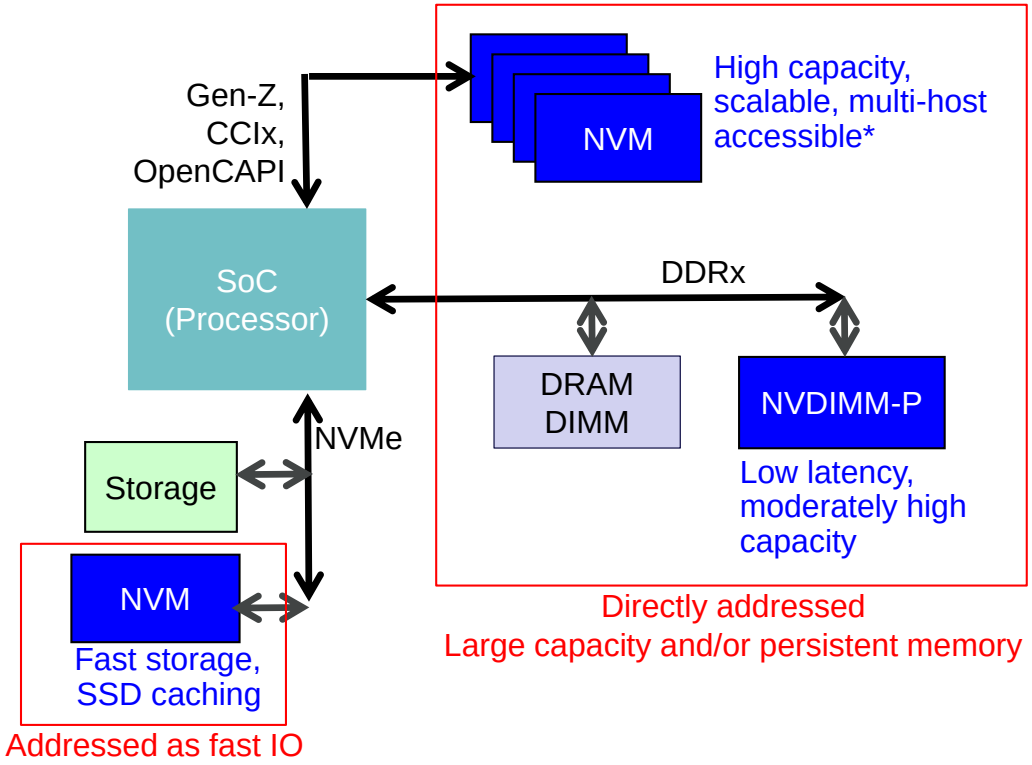

Managing NVM as main memory

Higher latency, lower bandwidth

- Additional buffering required

Lower Endurance

- Requires wear-leveling, ECC
- Processor architecture techniques could also reduce number of writes

Non-volatility

- Security concerns (even if persistency is not required)

Larger capacity

- Stresses translation

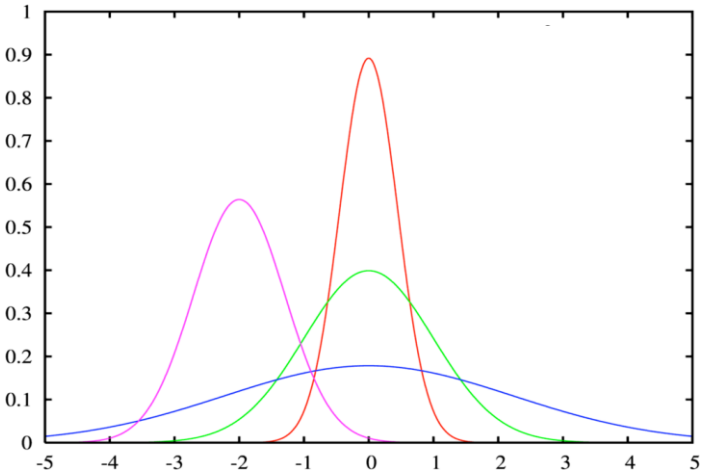

Variable latency and tail distribution

- Requires non-deterministic, agnostic protocol

Protocol requirements (read cmds)

Managing NVM Characteristics with NVDIMM-P

- Agnostic, non-deterministic protocol
  - Support various high capacity technologies
  - Optimal media management on the DIMM including light FTL like services
- Longer read latency
  - Buffer read commands
  - Support out of order completion
  - Unique read ID per command, returned with data response
  - Caching could be incorporated to manage longer latency media

NVDIMM-P resides and shares same channel as DDR

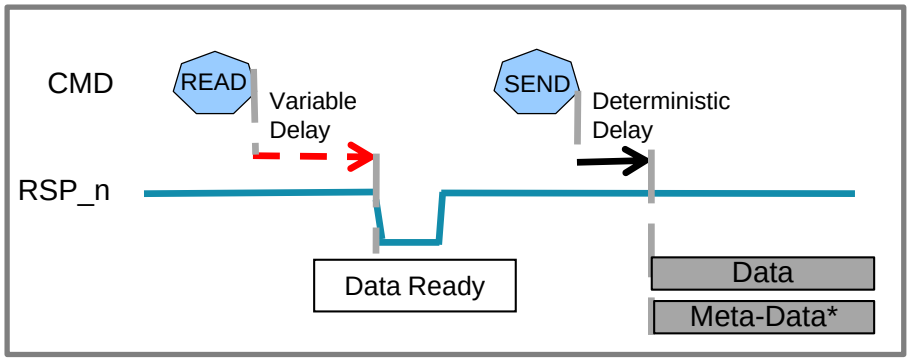

*Meta-Data consists of
{ID, Write Credits, Channel ECC, Poison, …}

Protocol requirements (write cmds)

Managing NVM Characteristics with NVDIMM-P

- Longer write latency
  - Write data buffered
  - Flow control ensures buffers don't overflow
  - Write credits returned in meta-data of read and status read commands
- Explicit write response is only required for persistent operations
  - Not required for normal writes
  - No need for write IDs

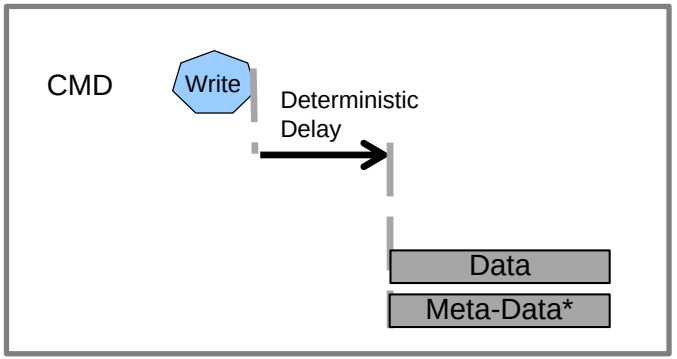

*Meta-Data consists of
{Channel ECC, Poison, …}

What about persistence?

Managing ordering requirements

- Persistent memory requires strict ordering
  - Data persisted without going through the SW stack
  - Known state is required to recover from power failures without corrupting data
- The interface must manage a point-of-persistency (PoP)
  - Host processor must know when PoP is achieved
  - Could be per cache line eviction, software barrier, epoch

Persistent support on the interface

Multiple ways to handle persistence

1. Battery, super-cap backed DIMM
2. Explicit write persistent command – Persist per write group ID
3. FLUSH command – FLUSH all or per write group ID

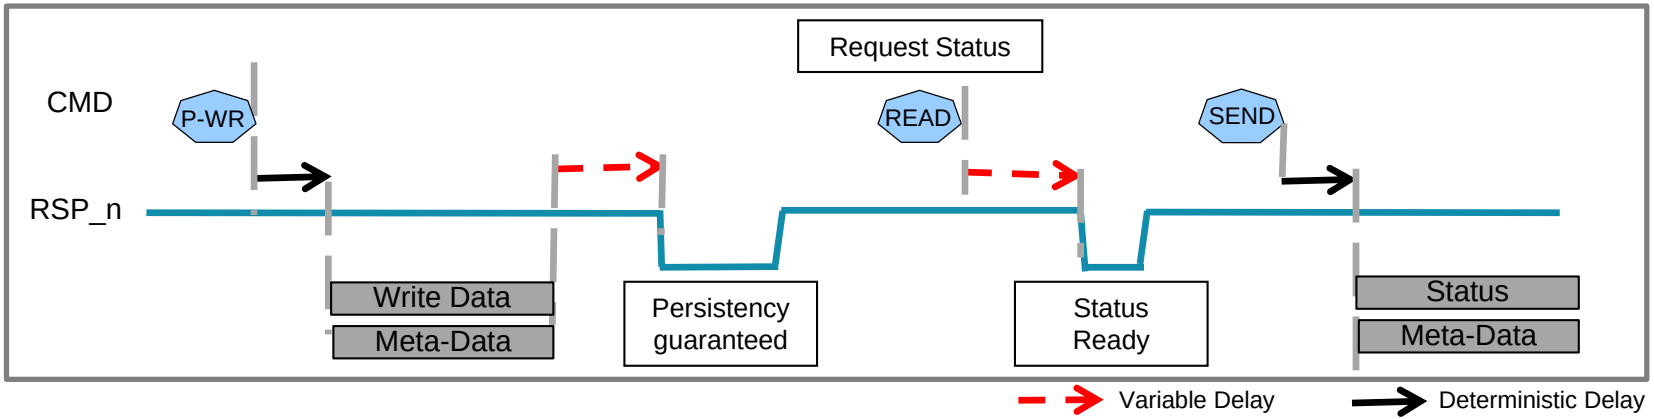

Where do we go from here?

- Emerging NVM is creating opportunities to redefine the memory sub-system
  - Transformative capacity
  - Directly addressable persistent memory

- JEDEC is actively defining NVDIMM-P protocols as one key attachment point
  - Key items have been balloted to enable IP development
  - Will continue to flush out details to finalize v1.0 specification (early 2018 projected)
  - **Become more active in JEDEC for more details and to help steer the future!**

Thank You!

Persistent Memory

Connectivity & Performance

Tom Friend, Toshiba Memory America

Its all about connections

Wanted: high speed bus for long term relationship with persistent memory. Must be able to respond quickly, carry large loads and work in a flexible manner. You will be judged by how many different vendors you can talk to. Extra points given for great feedback to those around you and for keeping secrets.

Current memory I/O- DDR

- DDR is perfect for DRAM
- Synchronous design, works for predictable latency devices
- CPU memory controllers designed for DDR

What PM needs

- Variable timing
- Sideband management channel
- Low latency, high bandwidth protocol

JEDEC is working on it

- DDR-5 might be “it”
- Let’s make it work for DRAM first

NVM Architecture

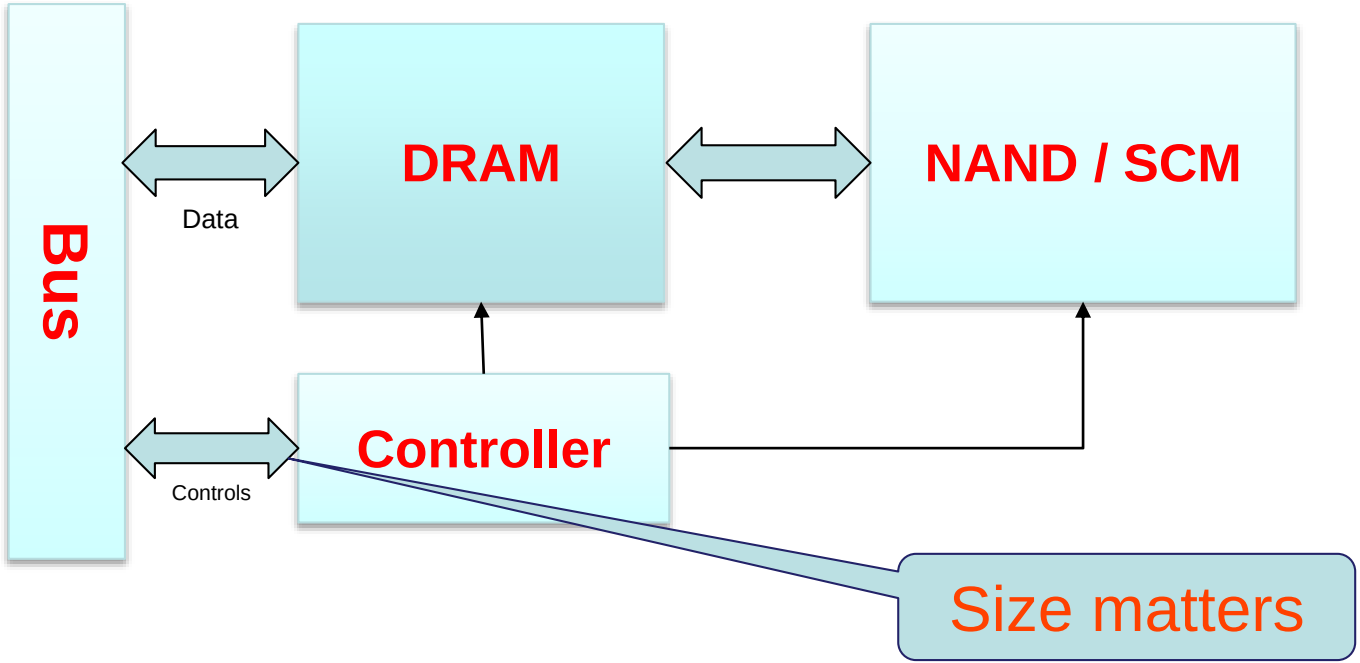

Many interfaces for attaching PM

Interface	Notes	Speed	When	Suitable
DDR4	Available today	25.6GBs	Now	
HBM2	Faster DDR	256GBs	2016	☹
HBM3	Really fast DDR	512GBs	2020?	☹
PCIe 3.0	Readily available	16GBs	2010	☺
PCIe 4.0	Soon	32GBs	2017?	☺
PCIe 5.0	Next generation	64GBs	2020?	☺
OpenCAPI	POWER9 CPUs	50GBs	2018?	☺
Gen-Z	Storage Class Memory	112GBs	2019?	☺
CCIX	Cache coherent	50GBs	2018?	☺
Fabrics	Disaggregated access via network	10GB	2016	☺

Persistent Memory

Security considerations

Cyril Guyot, Western Digital

Security Threats

- Unauthenticated access to data at rest

Theft

Disposal

- Application isolation

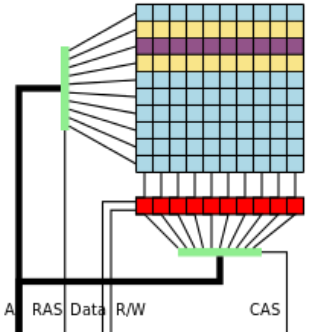

Rowhammer

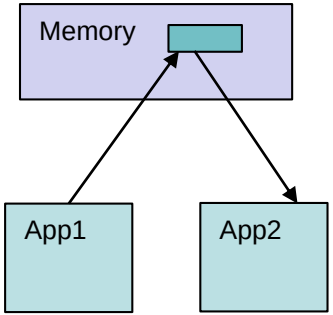

Multi-tenancy

Common threats: Storage/Persistent Memory

- Security of data at rest (block ciphers for confidentiality and pseudo-integrity: AES-XTS, others?…)
- Authentication (configuration/locking/unlocking)
- Storage security standardization well covered
- TCG Storage Workgroup (Opal SSC etc…)

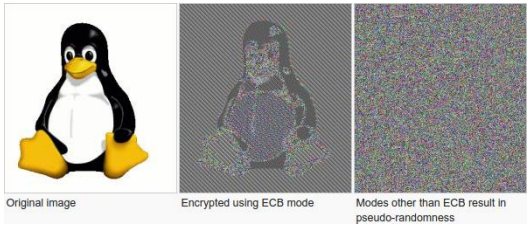

Storage vs Persistent Memory

- Security: How does persistent memory differ from storage? What are the drivers?

Characteristics	Storage	Persistent memory
Latency	ms / us	us / ns
Bandwidth	MB/s / GB/s	GB/s / xxGB/s / xxxGB/s
Access granularity	Sectors: 512B / 4KiB	Cache lines (64B), atomic instructions (down to 1B)
Interface semantics	More rich and more flexible	Less rich and less flexible

What PM needs for security

- Low-latency/high-bandwidth confidentiality and integrity ← Latency/Bandwidth/ Interface semantics
- Fast power-cycle data erasure (DRAM emulation) ← Latency/Bandwidth/ Interface semantics
- Side-channel resilience (Rowhammer) ← Fine-granularity
- Denial of service resilience (Endurance exhaustion) ← Fine-granularity
- Fine-grained access control ← Fine-granularity

TCG/JEDEC/SNIA working on it

- TCG Persistent memory workgroup: defining interfaces and properties
- Please contribute!

JEDEC Membership

- Become a JEDEC member & join with other industry leaders to set the direction for **NVDIMM** and **persistent memory** standards

- Special new member company discount: **50% off annual dues** for the first year of membership*

- **Stop by the JEDEC booth to learn more! #803**

**Discount available only to companies who have never been a JEDEC member.*

Flash Memory Summit 2017 – Exhibit Floor

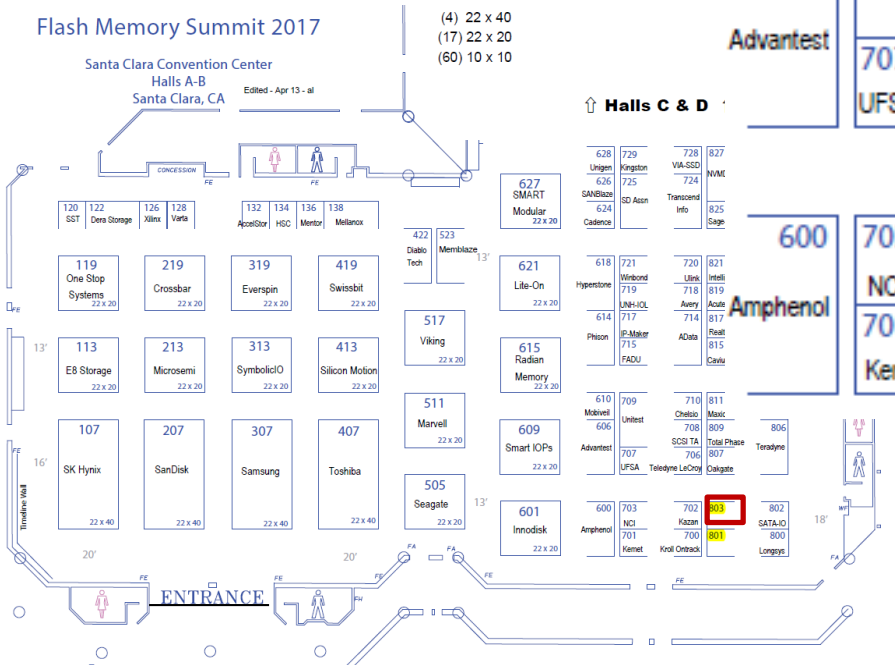

Mobiveil 606	Unitest	Chelsio 708	Maxiotek 809	Viavi Sys 806
Advantest	707 UFSA	SCSI TA 706 Teledyne LeCroy	Total Phase 807 Oakgate	Teradyne
600 Amphenol	703 NCI	702 Kazan	803 JEDEC	802 SATA-IO
	701 Kemet	700 Kroll Ontrack	801	800 Longsys

Flash Memory Summit

QUESTIONS