Overview of Persistent Memory

FMS 2018 Pre-Conference Seminar

Mark Webb

MKW Ventures Consulting, LLC

Mark’s Presentations at FMS

- Persistent Memory Preconference Class (Monday 8:45AM)
- MRAM Developer Day: MRAM Memory Technology Roadmap and Market (Monday 11:45AM)
- Advances in Persistent Memory (PMEM 101-1, Tuesday 8:30AM)
- Annual Update on Emerging Memory Technologies (NEWM-102B-1. Tues 4:55)
- 3D Xpoint Technology Expert table (Beer/Pizza) (Tuesday 7PM-830PM)
- 3D XPoint: Current Implementations and Future Trends (NEWM-201A-1, Wed 8:30AM

Contents

- Persistent Memory Definitions
- Applications and what is shipping today
- NVM Technologies for persistent memory
- Persistent memory configurations
- Challenges and opportunities
- Revenue projections and forecasts

Mark Webb, MKW Ventures

Different Concepts of Persistent Memory

- It's a universal Non-Volatile Memory Technology (Device Geeks)
  - PCM, ReRAM, MRAM, Memristor, NVRAM
- It's a storage/memory concept (Storage Experts)
  - What if we wrote to address and didn't have to worry about data loss or storage later?
- Its BIG DATA Memory (End users)
  - I want to look at all my TBs of data like hot data

Mark Webb, MKW Ventures

A Persistent Memory Definition

- It's persistent … ie NVM (duh!). No need to worry about loss
- It's accessed like memory on memory bus
  - "Byte addressable" …. Could also be used in Block Mode
  - Anything can be virtual memory… but this is less interesting
- Speed…unclear, lets say <1us latency
  - 2018 PM Summit had some great discussions on this (WDC/Bandic)
  - Raw memory read latency on order of 100ns
- Used for data being worked on and addressed by programs. Not primarily used as Storage

How is PM Accessed

- Like DRAM: DDR4 bus. Parallel memory slots on server/PC board (Today). NVDIMM-N, NVDIMM-P or non-standard DDR4
  - Also PMoF/RDMA
- On New Bus: GenZ, OpenCAPI, Rapid-IO (coming)

- *Through NVMe/Storage bus: This is available today working with different memories but it is not my focus*

Historical Memory/Storage vs PM

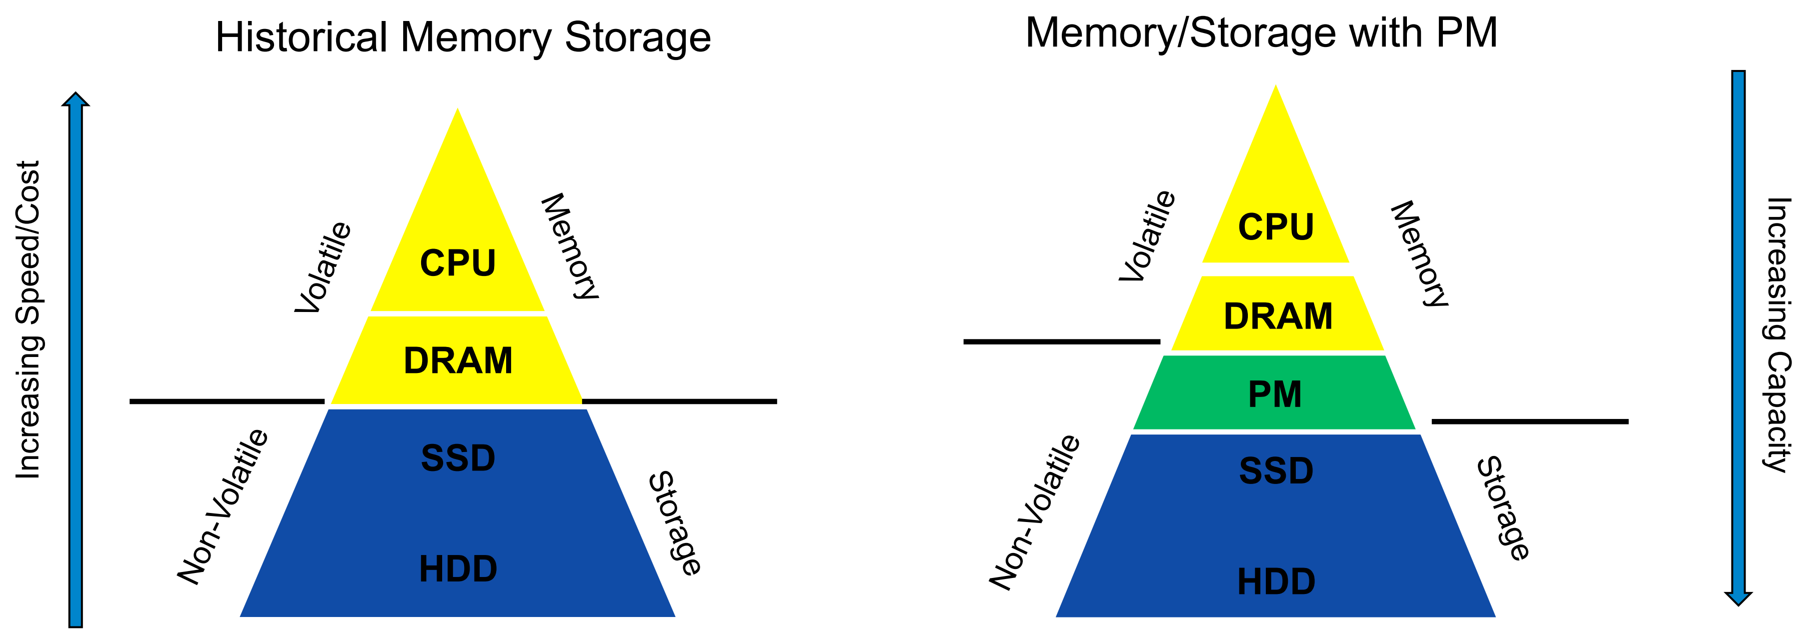

Mark Webb, MKW Ventures

How to work with 1TB of Data

OVERSIMPLIFIED!

196GB DRAM+8TB HDD	196G DRAM+1TB NVMe SSD+ 8TB HDD	1.5TB Persistent Memory+ 8TB HDD

196GB DRAM+8TB HDD

Memory

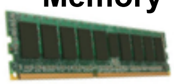

Storage

196G DRAM+1TB NVMe SSD+ 8TB HDD

Memory

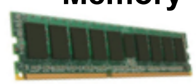

Storage

1.5TB Persistent Memory+ 8TB HDD

Memory

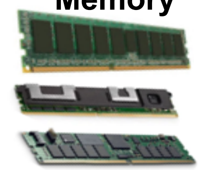

Storage

How to work with 1TB of Data

OVERSIMPLIFIED!

196GB DRAM+8TB HDD

- Data is on HDD
- Load part of it in to Memory
- Swap out blocks of data as needed until done
- Memory access times 30ns
- HDD access time mS
- Hope no power lost during work

196G DRAM+1TB NVMe SSD+ 8TB HDD

- Data is on HDD
- Load it all to SSD
- Load part of it in to Memory
- Swap out blocks of data with SSD until done
- Perhaps treat SSD as memory
- Memory access times 30ns
- SSD access time 10uS
- Hope no power lost during work

1.5TB Persistent Memory+ whatever

- Data is on HDD/SSD
- Load it all to persistent memory
- Complete work on data
- Leave it there or store it to SSD/HDD
- Memory access time is 30-150ns
- No SSD/HDD access needed
- If power lost, you are good

Mark Webb, MKW Ventures

Persistent Memory Applications … It's Here Today

- Server DIMMS/Main Memory for Compute and Servers
- RAM requirements where max speed is needed and memory cannot be lost due to outage.
- Tremendous ecosystem and standards work supporting this
- Log file, networks, quick start ups and quick restarts.
- 16GB applications growing to 64GB
- Multiple Suppliers here at FMS
- Still relatively low volume and penetration (<5% of servers)

Mark Webb, MKW Ventures

What's Shipping Today

- NVDIMM-N is classic version of persistent memory DIMM
  - Addressed just like DRAM in a DIMM
  - Backed to NAND periodically or when power lost
  - Typical NVDIMM is 16G DRAM plus 32G of SLC NAND with control and capacitor/battery
  - Appears as 16GB of DRAM at DRAM speed
  - Costs more than DRAM and does not provide increased capacity.
- NVDIMM-F is version with only Flash. "SSD on Memory bus"
  - NAND on DRAM Bus has always shipped in limited volume
  - Much slower than DRAM at higher capacity

NVDIMM-N, NVDIMM-F

NVDIMM-N

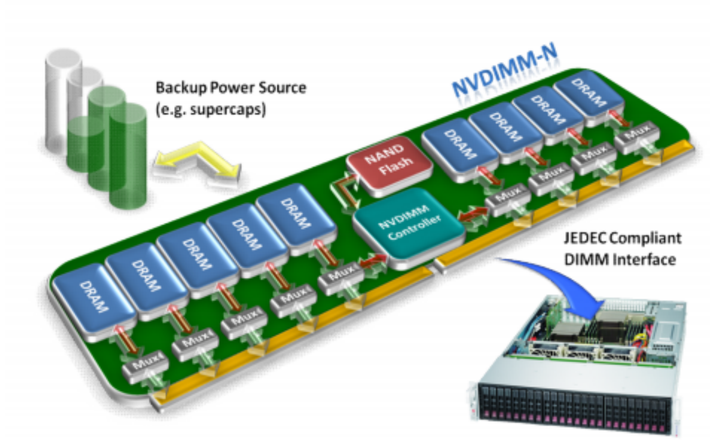

NVDIMM-F

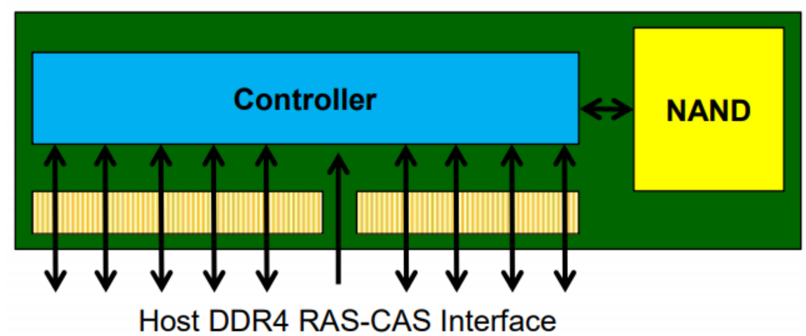

Mark Webb, MKW Ventures

High Density Server DIMMs

- Future Apps: Large databases where loading and swapping portions is not efficient.
  - Size of SSD (Terabyte) with memory bus speed (Mark's definition)
  - This is a major revenue Focus
- Anything where faster loading, faster analysis provides monetary return to pay for it
- Examples:
  - Financial database/transaction processing ($/mS metrics available)
  - VMs that are currently memory limited (10x more VMs/Server)
  - Video/entertainment/Animation (Huge databases, PM Summit)
  - Similar to applications currently using high performance NVMe SSDs

Persistent Memory Applications (MORE)

CE/Mobile Devices (Potential Revenue)

  - Smaller density replacing Capacitor/battery backed DRAM, replacing SRAM/DRAM/Flash. CE device optimization
- For cost-speed reasons, these applications often optimize NAND and DRAM and HDD in gaming/CE systems
- Potential to create a memory system that is fast enough and allows less chips, faster overall speed, better reliability.
- For Many apps, lower density is OK enabling more media (memory types) options
  - 16M SRAM+1G DRAM+8G NAND could use MRAM for aspects.
  - 2G DRAM+16G NAND could go to ReRAM/PCM-3D Xpoint

Mark Webb, MKW Ventures

Memory Types/Media

	Latency	Density	Cost	HVM ready
DRAM	*****	***	***	*****
NAND	*	*****	*****	*****
MRAM	*****	*	*	***
3DXP	***	****	****	****
ReRAM	***	****	****	**
NRAM	***	**	**	*
Other	***	**	**	*

Combined Today (DRAM, NAND)

Alone or Combined In future (MRAM, 3DXP, ReRAM)

Notes: NOR/SRAM and low density Not in Included (Small), Low density FeRAM not included

Mark Webb, MKW Ventures

The Latency Spectrum and Gaps ~2015

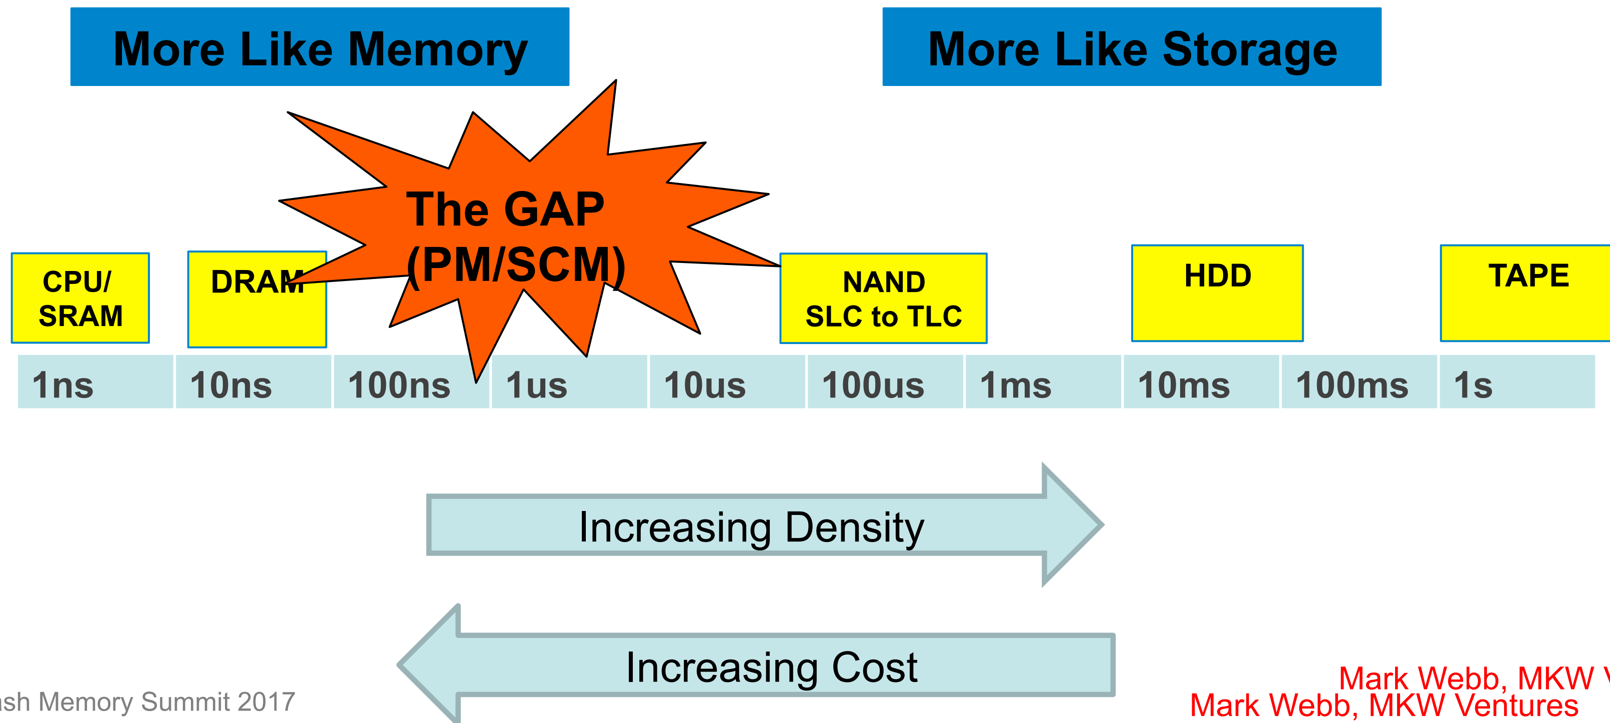

Mark Webb, MKW Ventures
Mark Webb, MKW Ventures

The Latency Spectrum and Gaps
Future

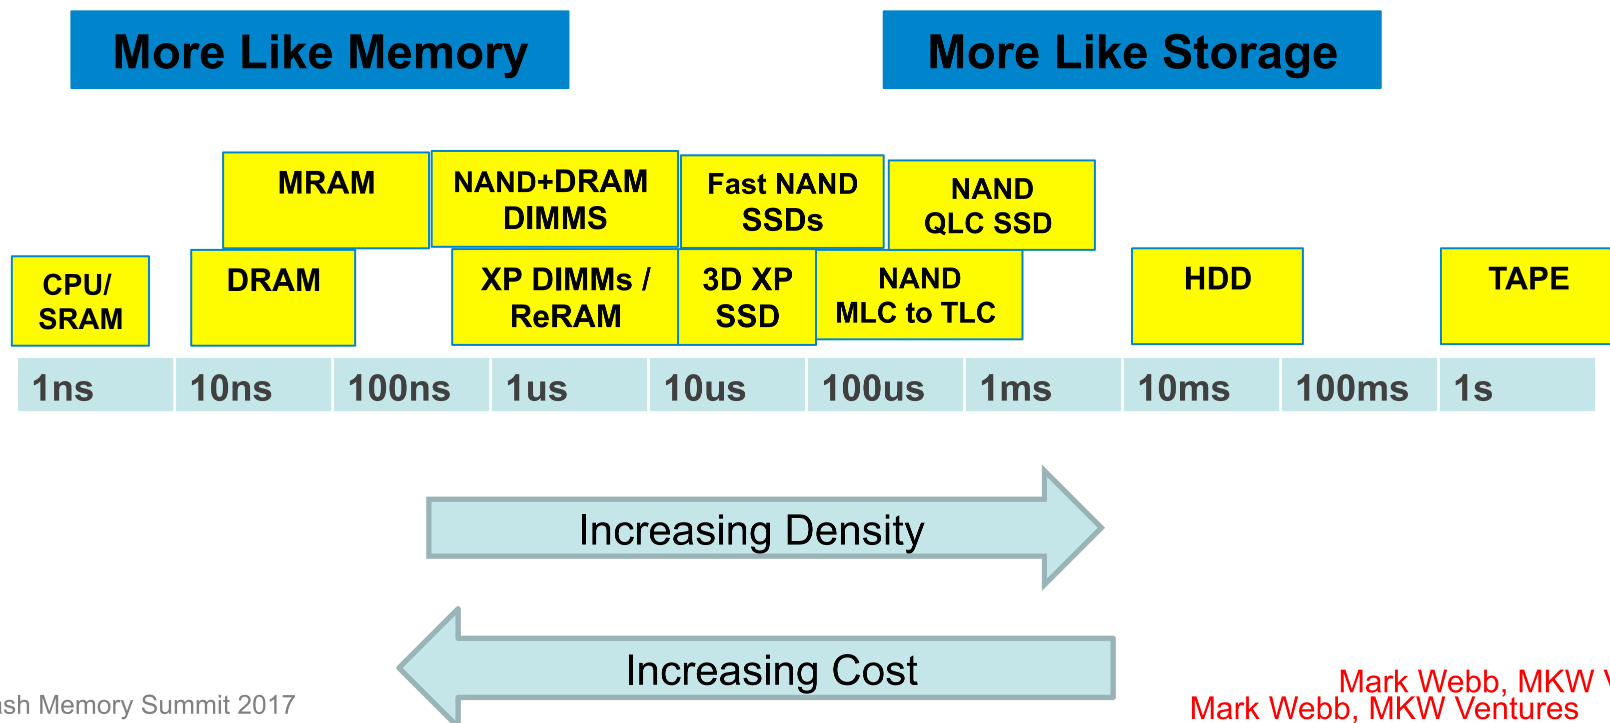

Mark Webb, MKW Ventures
Mark Webb, MKW Ventures

Coming Persistent Memory/ SCM Technologies

- NVDIMM-N meets the specs but is very expensive and density <=DRAM
- Optane Persistent Memory (DIMM) will be dominant PM very quickly
  - Better density than DRAM, lower cost
  - But slower speed, cycling limitations mean tradeoffs.
- ZNAND/Fast NAND: slower than DRAM, cycling limitations (good for SSDs)
- MRAM: Much more expensive than DRAM (But close on speed)
- ReRam: Slower than DRAM, Cycling limitations (Like Optane)
- Other Memories have potential but are much less Mature/low density
- **DRAM "replacement" isn't the way to persistent memory market growth !**

Mark Webb, MKW Ventures

Example of Cost Challenges

- 2018 estimated Cost (not price) per Bit (DRAM RDIMM=1x)
  - MRAM: 5x
  - NVDIMM-N: 1.6x
  - ReRam (today): 0.75x
  - 3D Xpoint (today): 0.55x
  - Fast SLC NAND (today): 0.15x
- DRAM+ReRAM/Optane/NAND is lower cost/bit, more capacity at “similar” performance

Mark Webb, MKW Ventures

DRAM/NVM Combinations

- NVDIMM-P Supports multiple memories and hybrid systems
- Coming solutions are some DRAM merged with lots of NVM.
  - Lower cost, near DRAM performance, managed endurance
- 3D-Xpoint persistent memory combines DRAM DIMMs and 3D Xpoint DIMMs with processor/memory controller managing data
  - ~5:1 Xpoint:DRAM ratio, manage data for performance/endurance
- Netlist HybriDIMM/Xitore: DRAM and Fast NAND on DIMM
- Z-NAND and solutions from All NAND and NVDIMM vendors will use similar architecture
  - Cheaper than DRAM, Lots of memory, Managed endurance

Mark Webb, MKW Ventures

Big Data vs Fast Data vs Persistent Data (Cost question to ponder)

- Big Data: Data analyzed is getting bigger
- Fast Data: Data Analyzed is needed faster.
- Do I want my memory faster, cheaper or persistent??
  - 200GB of DRAM at $2000
  - 150GB of Fast NVDIMM-N at $2000 (hypothetical)
  - 1000GB of slower PM at $3000 (hypothetical)
  - 1000GB of slower DRAM (volatile) at $2500 (hypothetical)
- The answer will determine which products take off

Mark Webb, MKW Ventures

Predictions for Market

- All of these options will be provided to end users
  - NVDIMM-N/P, Optane DIMM, Hybrid DIMM, Z-NAND/Fast NAND combined with DRAM, etc
  - Some Proprietary, Some open, with the usual arguments why
- If Persistent memory is important, Certain architectures will become standard and grow faster leading toward "High Revenue"
- If we are having "what's possible" discussions at end of 2019, Market will be much, much lower than middle revenue☹....

Mark Webb, MKW Ventures

Persistent Memory

Revenue Growth "Guess-timate"

Year	Revenue Middle	Revenue High	Requirements to meet <u>Middle</u>
2020	$2.0B	$3.0B	Optane, NVDIMM must takeoff ASAP
2022	$3.9B	$7.0B	Persistent memory is in all compute areas. Multiple bus options evaluated
2025	$7.0B	$10B	Multiple new memories allow utilization in mobile, server, PCs

NOTES:
Revenue "low" is too depressing to show. I'm an optimistic guy
NVDIMM+SCM/NVRAM standalone memory only. Virtual memory on storage bus not included
NVDIMM could be DRAM+NAND, Fast NAND, SCM
Embedded PM is difficult to measure revenue

Mark Webb, MKW Ventures

Mark's Summary

- Persistent memory is here today, but it is just a start
- To grow, we need to be cost effective.
  - DRAM replacement by expensive tech won't work broadly
  - Memory that is too slow won't work broadly
  - Neither DRAM nor NAND are getting replaced.
- DRAM + NAND/SCM will be the PM future
  - Includes Optane Persistent Memory which requires DRAM
- Revenue could grow 30% CAGR if technologies deliver to commitment dates

Flash Memory Summit

Where should we attach persistent memory?

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

	CPU BUS: PARALLEL	CPU BUS: SERIAL	SERIAL PERIPHERAL BUS
Physical interface	DIMM	DIMM/other	PCIe
Logical interface	Non-standard DDR4, NVDIMM-P	DMI for Power 8, CCIX, OpenCAPI 3.1, Rapid-IO, gen-Z	NVMe, DC express*
Pros	• Low latency • High bandwidth • Power proportional • Coherent through memory controller	• High bandwidth • Significant pin reduction • Higher memory bandwidth to CPU • Coherent through memory controller, or in some cases can even present lowest point of coherence	• Standardized • CPU/platform independent • Latency low enough for storage • Easy RDMA integration • Hot pluggable
Cons	• CPU memory controller has to implement specific logical interface • Not suited for stochastic latency behavior • Not hot pluggable • BIOS needs change	• CPU memory controller has to support • May have higher power consumption	Higher latency (~1us)

Western Digital.

Mark Webb, MKW Ventures