

Persistent Memory in 2020: Introduction



Flash Memory Summit

Dave Eggleston
Intuitive Cognition Consulting




Compute accelerates, memory...not so much

Adapting to the Changing Data Landscape


COMPUTE DEMAND ACCELERATING

CPU Core Growth Projection Over Time

DRAM IS NOT SCALING








LARGE MEMORY SYSTEMS GROWING



Source for "Compute Demand Accelerating and Large Memory Systems Growing": Intel Market Intelligence Group - DRAM Market Monitor Q1-20" by Yole Development
Source for DRAM is Not Scaling: "3D NAND Technology - Implications for Enterprise Storage Applications" by J.Yoon (IBM), 2015 Flash Memory Summit



DRAM! (But...)




1. More cores, more channels, DDR bandwidth/core has... slowed down?!



DRAM! (But...)



HPC System 1	
Cores	56
Memory	112GB
Memory Power	50W




HPC System 2	
Cores	112
Memory	1.8TB
Memory Power	~700W!


2. More cores, more memory = Order of magnitude more Power!



DRAM! (But...)



Possible DDR5 Server Subsystem



3. More cores, more channels = Substantially more DDR board complexity!

Can PM deliver where DRAM is falling short?

1. Bandwidth
2. Power
3. Complexity


And...

4. Cost/GB

Plus what application(s) does
PM accelerate?




PM Media Options



- PM media choices have been broad, but seeing consolidation



PM Media: Who is Doing What



- Only PCRAM has shown the characteristics necessary to deliver the PM capacity, cost, and performance.
- Memory giants are now consolidating around PCRAM.
- MRAM is a good embedded NVM for foundries – just not suitable for high capacity PM.

Source: Industry announcements & scuttlebutt



PM Media: Optane Series 100 to 200


	100 Series	200 Series	% Improvement
DIMM Capacities	128/256/512 GB	←	
DDR Frequencies	2666, 2400, 2133, 1866 MT/sec	2666 MT/sec	
Performance shown for 256GB DIMM @ 15W (Best performing DIMM)			
Endurance 100% Writes 256B	363 PBW	497 PBW	37%
Endurance 100% Writes 64B	91 PBW	125 PBW	37%
BW 100% Read 256B	6.8 GB/sec	8.1 GB/sec	19%
BW 100% Read 64B	1.75 GB/sec	2.03 GB/sec	16%
BW 100% Write 256B	2.3 GB/sec	3.15 GB/sec	37%
BW 100% Write 64B	0.58 GB/sec	0.79 GB/sec	36%

- 200 series Optane focused on endurance and BW improvements



PM Value Propositions: TCO, Throughput, Speed

Intel Optane PMem Delivering Real World Benefits




Performance results are based on testing as of dates in configuration and may not reflect all publicly available security updates. See backup for configuration details.
For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.



DB App: PM + DRAM Outperforms DRAM alone

DRAM-like Performance

Sysbench QPS on MySQL



- Plus a 50%+ cost savings for PM + DRAM system!



AI/ML App: PM acceleration

Case Study: AI / Machine Learning – Facebook's DLRM

Background

Customer Type: AL / ML Customer

Business Challenge:

- Dynamic and Scalable Production Inferencing

Platform:

- Innovative Big Memory Computing platform for leveraging persistent memory for real-time, AI/ML and Advanced Analytics and extensible to all memory – Centric workloads.


Software:

- Software Defined Architecture extracting performance benefits of cutting edge hardware supporting workload portability to truly compute anywhere with the memory speeds.



Penguin, MemVerge and Intel Solution

~400GB swap on NVMe is on, 20480 samples evaluated



Result:


- Customer has state of the art AI/ML Big Memory Platform that is can scale and deliver performance when Data is Greater than Memory
- Achieved flexible software defined platform Big Memory Computing capabilities and poised for future dynamic model and data growth

- 10x inferencing acceleration, NO App rewrite!



PM Attachment: DDR, OMI, GenZ, CCIX, CXL...

01100101 00100000 01100011 01101111 01101101 01100101 01110011 00100000 01101000 01101111 00100000 01110001 01110101 01100001 01101110 01110100 01110101 01101101 01100111 01100100 01101000 01101001 01101110 01101100 01101100 01101001 01101110 01100011



CXL Compute Express Link™


OpenCAPI™

GEN Z

- DDR (now) and CXL (future) for PM attachment




PM Form Factor: DDR DIMM to E1.S



E1.S



DDR4 DIMM



- E1.S form factor as the PM successor to the DDR DIMM



PM: \$/GB versus DRAM

	PMEM	DRAM
1 x 512GB	 \$13.86/GB	
1 x 256GB	 \$7.02/GB	 \$18.94/GB
1 x 128GB	 \$4.00/GB	 \$13.67/GB
1 x 64GB		 \$7.65/GB
1 x 32GB		 \$8.43/GB
1 x 16GB		 \$9.37/GB

August 2020 prices from online resellers. Prices vary widely.

- PM priced ~30% to 50% of DRAM (\$/GB basis)



PM Introduction Summary



- DRAM challenges ahead
- PM can address bandwidth, power, complexity, cost
- PM media consolidation around PCRAM
- 2nd generation PM is here now with improvements
- DB and AI/ML applications accelerated using PM + DRAM
- PM priced attractively versus DRAM
- PM form factor from DDR DIMM to E1.S



Stay tuned for details at
www.snia.org/pm-summit

- 9th annual Summit – April 21-22, 2021
- Deep dive into the latest memory and storage developments
- New for 2021
 - Expanded to two days – 30+ sessions
 - Virtual - stream live or watch on demand

Thank You!



Flash Memory Summit

Everything You Need To Know
For Success