Phase Change Memory

Chief Technology Office
Carnegie Mellon, Sept 23rd, 2009

Agenda

- Background / Why PCM
- What is PCM and How it Works
- Technology Characteristics and Trends
- Application: Storage Input / Output Acceleration
- Application: Directly Addressable NVRAM

Agenda

- Background / Why PCM
- What is PCM and How it Works
- Technology Characteristics and Trends
- Application: Directly Addressable NVRAM
- Application: Storage Input / Output Acceleration

Phase Change Memory

- "New Memory" Motivation (Recognized ~1999) :
  - Current NVM (and DRAM) are becoming Electrostatics Limited
    - MOS Transistor Based Cell; Charge Storage Memory Effect
  - Starting to Encounter Physical Scaling Limitations
  - Manifesting First as Degradation in Reliability (Endurance/Retention)
- Realization:
  - Next Generation NVM will Exploit New Storage Physics
  - Significant Innovation Will take Time (History says ~10yr "Gestation")
- Response:
  - Phase Change Memory Identified as the Best Candidate
  - Research Efforts Validated Assumptions and Initiated Development
  - Phase Change Memory Moving into Production Today
- We Believe PCM Features Enable New Usage Models

NOR Scaling Challenges

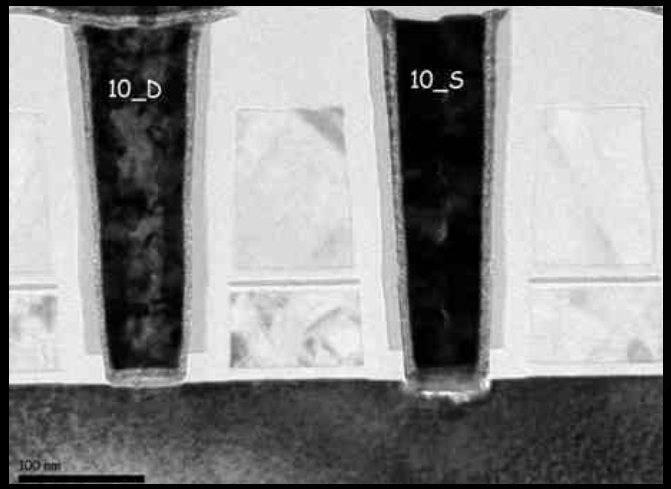

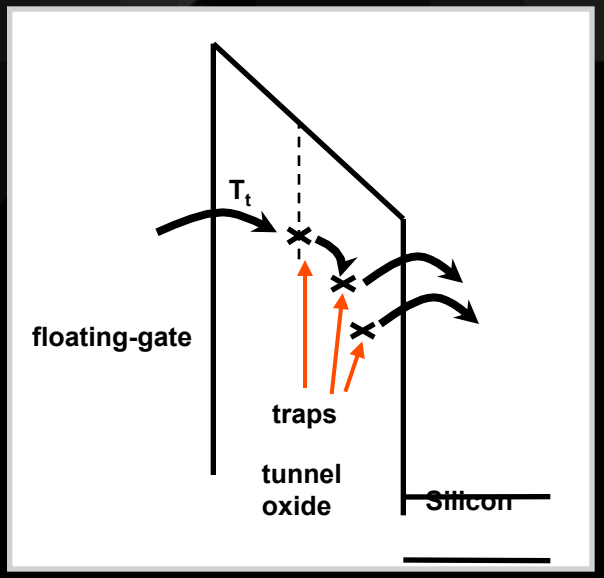

NOR (ETOX & NROM) Scaling: 3.2ev required to surmount Si-SiO2 barrier → Limits Cell Gate Length Scaling

NOR Reliability: Write / Erase Tox traps leading to TAT or de-trapping of trapped oxide electrons

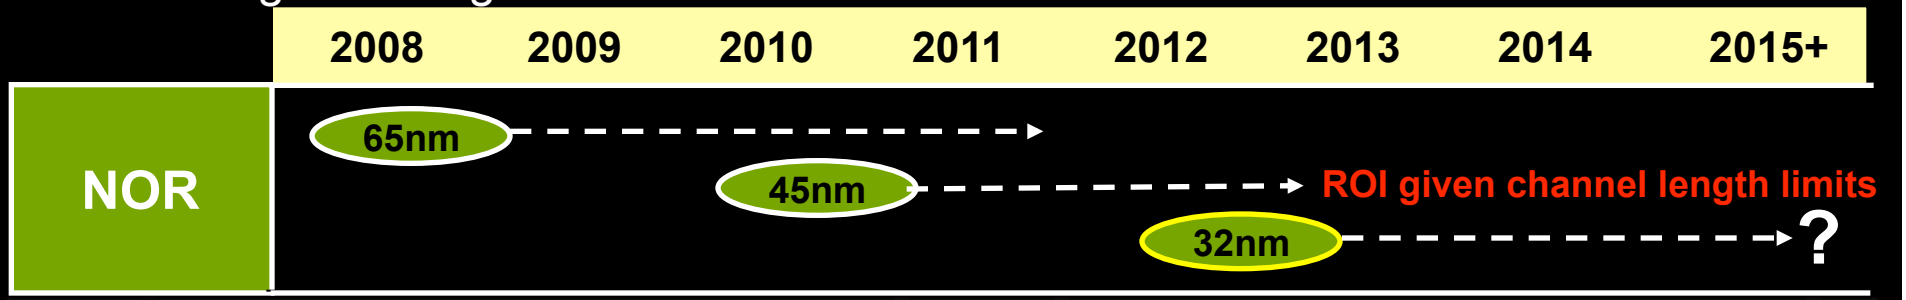

numonyx™

NAND Scaling Challenges

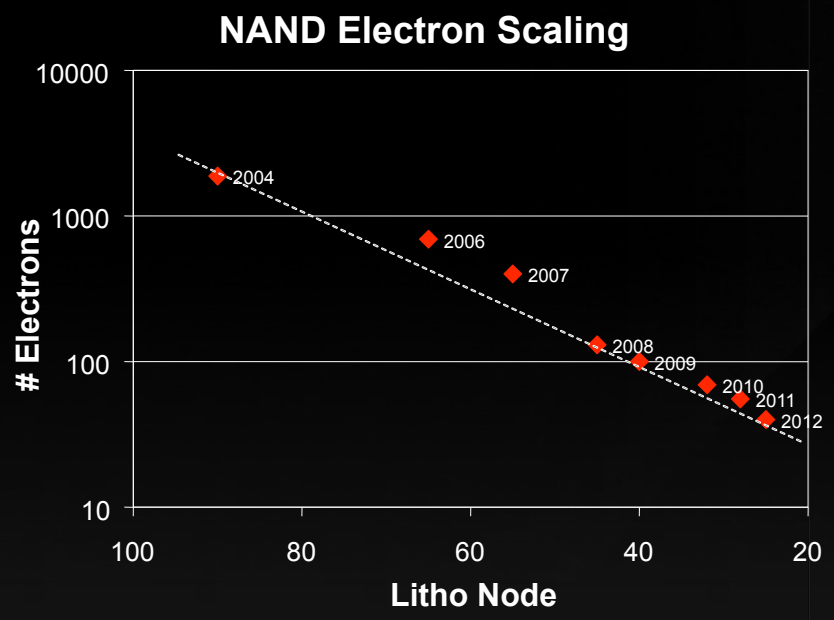

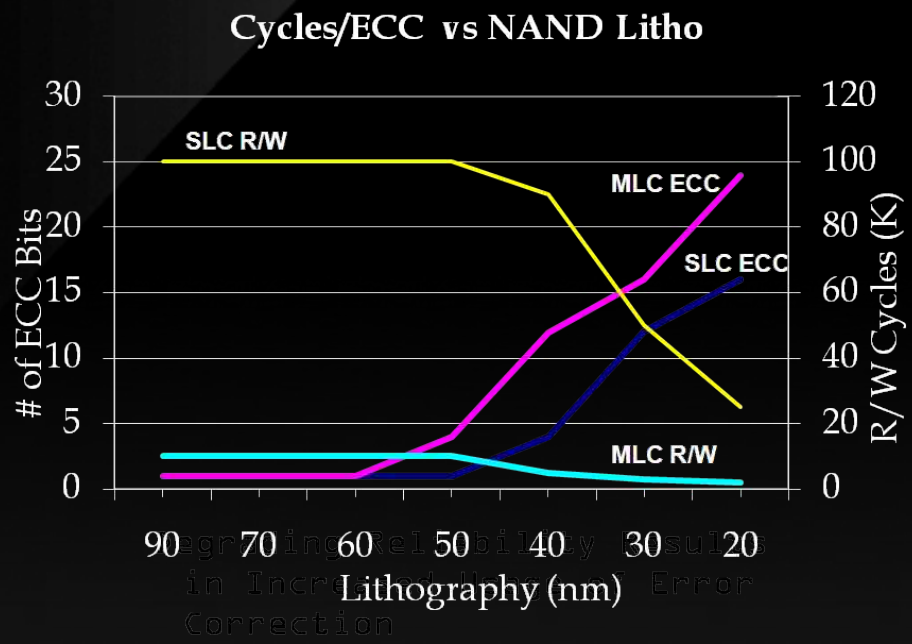

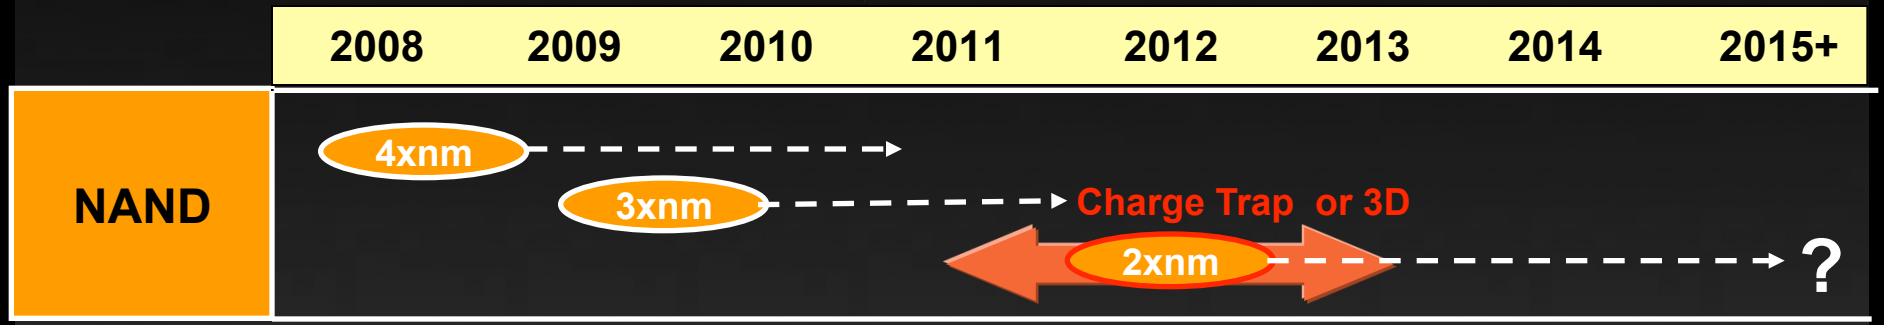

numonyx™

DRAM Scaling Challenges

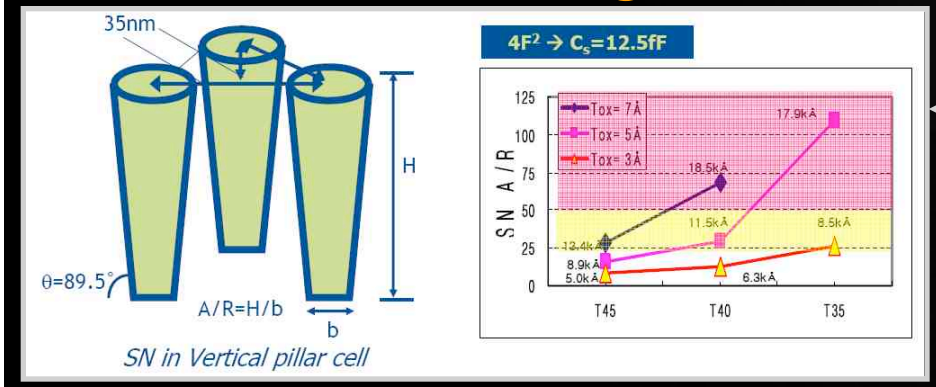

Frequent Capacitor Structure and Material Changes with Highest "known" HiK Material at ~32nm Generation

New Cell Transistor Every Generation

New Periphery Transistor Every Gen

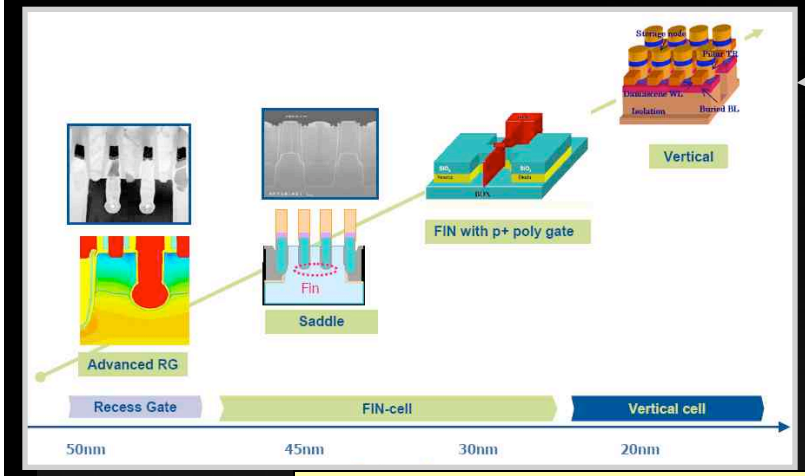

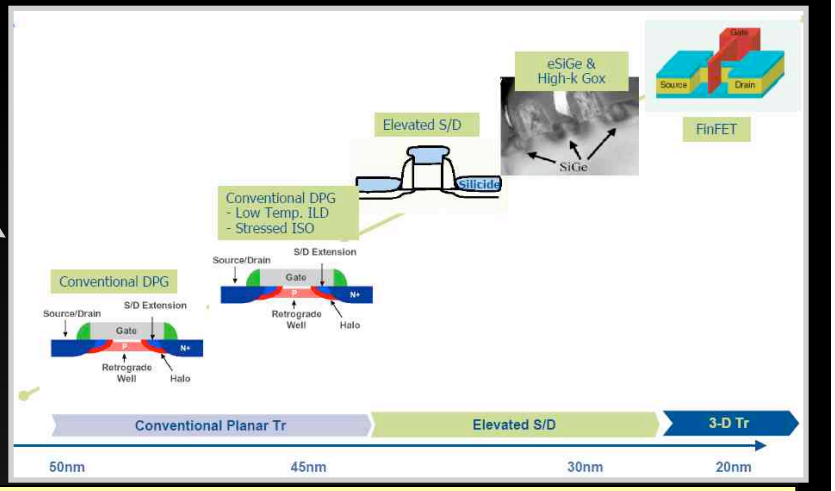

	2008	2009	2010	2011	2012	2013	2014	2015+
RAM	57nm →							
		45nm →			High-k / FinFET			
					36nm →			No known solution
							2Xnm →	

numonyx™

Theoretical Chip Cost Factors

Silicon Cost Component		SLC PCM	DRAM	SLC NAND
Die Size	Cell Size (F^2)	5.5	6.0	5.0
	4G Prod Example	1.0x	1.2x	1.0x
Wafer Complexity	Total Process Mask Count	~35	~34	30
	300mm cost structure	1.2x	1.2x	1.0x
Theoretical Die Cost Summary		1.2x	1.4x	1.0x

- PCM will be cheaper than DRAM at lithography parity
- PCM scales to lower densities better than NAND
- 300mm infrastructure required to get lowest cost
- PCM attributes can also save cost at system level

High Density PCM vs RAM Comparisons

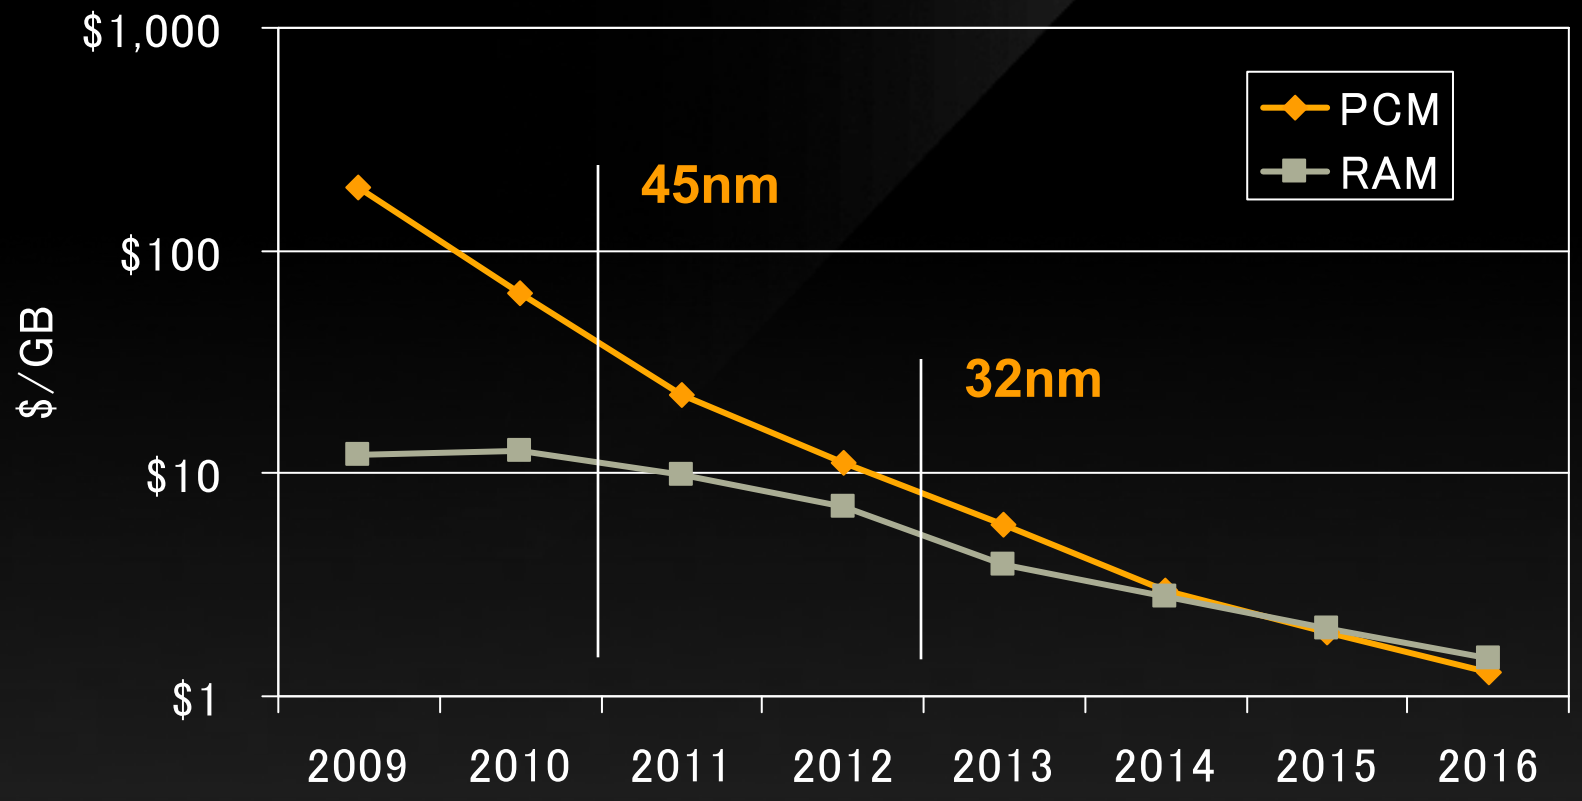

Source: Gartner Q209 (RAM), Numonyx (PCM)*

**1-8Gbit Density Products*

Landscape: Where do we go first?

Trying to figure out who will pay what $/GB

Storage

(No OS change,
No chipset changes)

Modules

...

Chips

Main Memory

(OS changes)

2010

2015

Agenda

- Background / Why PCM
- What is PCM and How it Works
- Technology Characteristics and Trends
- Application: Directly Addressable NVRAM
- Application: Storage Input / Output Acceleration

Phase Change Memory Mechanisms

- Storage
  - GST: Germanium-Antimony-Tellurium Chalcogenide glass
  - Cell states varying from amorphous (high resistance) to crystalline (low resistance) states
- Read Operation
  - Measure resistance of the GST
- Write Operation
  - Heat GST via current flow (Joule effect)
  - Time at critical temperature determines cell state

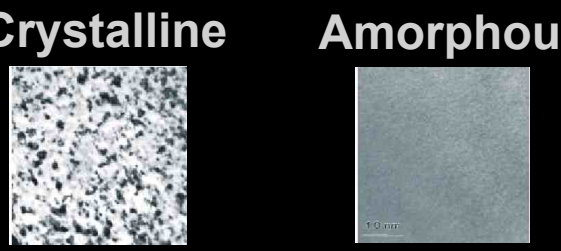

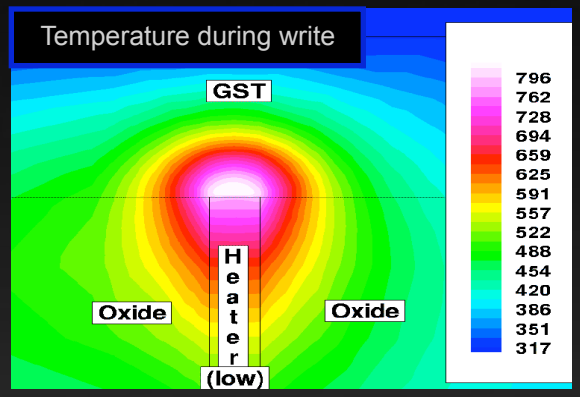

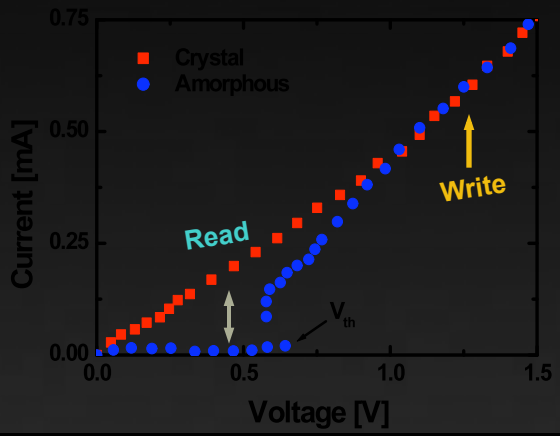

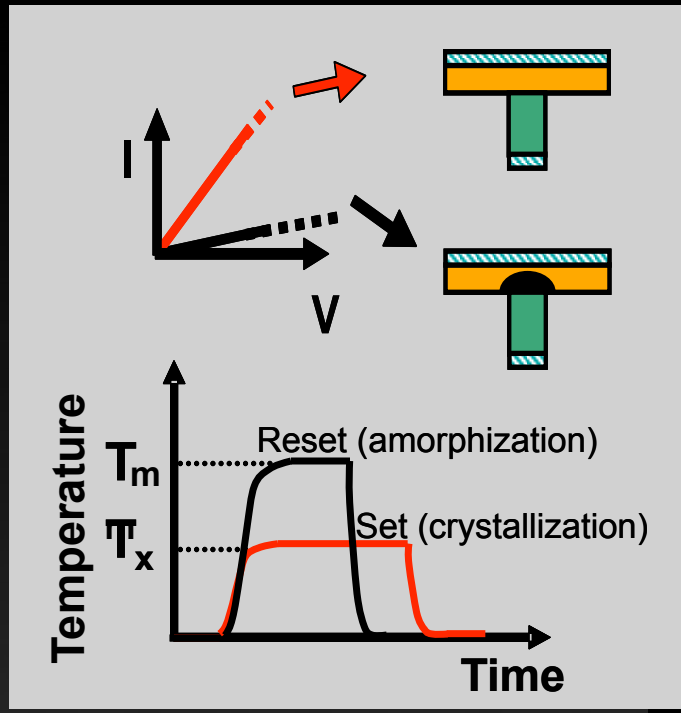

Summary of Key PCM Reliability Metrics

Retention decoupled from endurance

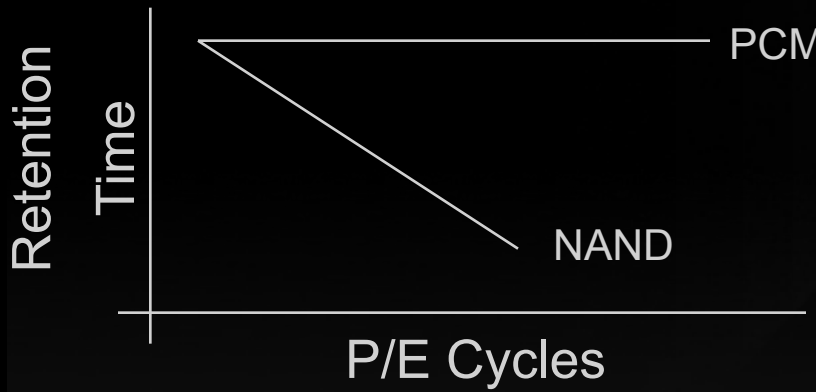

No issue with Read disturb

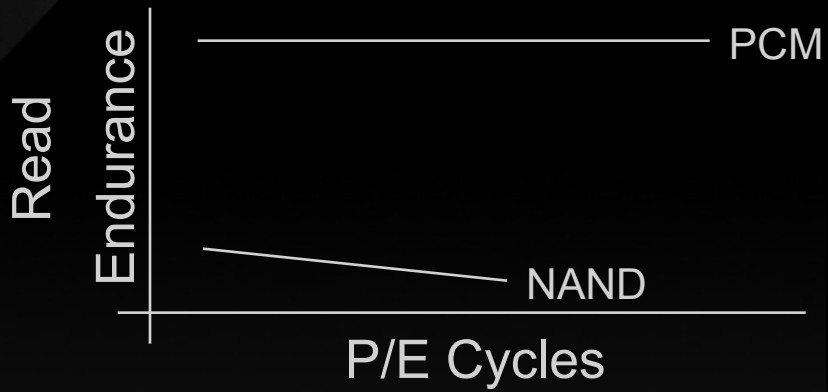

Result:

- Greatly simplified management
  - *No software overhead for some apps*
- Post-stress shelf-life problem solved
- Endurance can be increased significantly through management of bit errors
  - *Data is stable once successfully written*

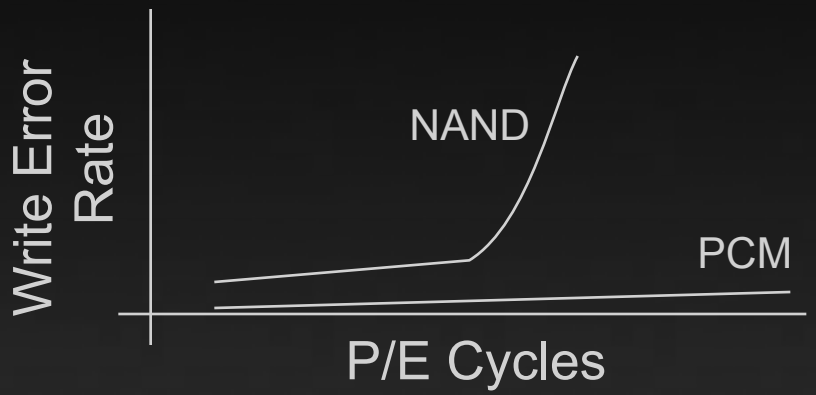

128Mb Reliability Data

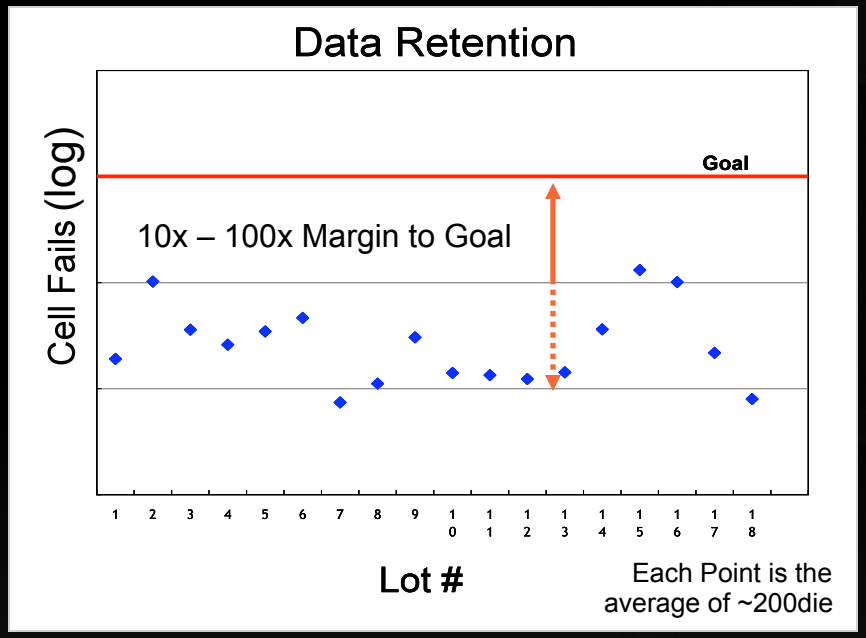

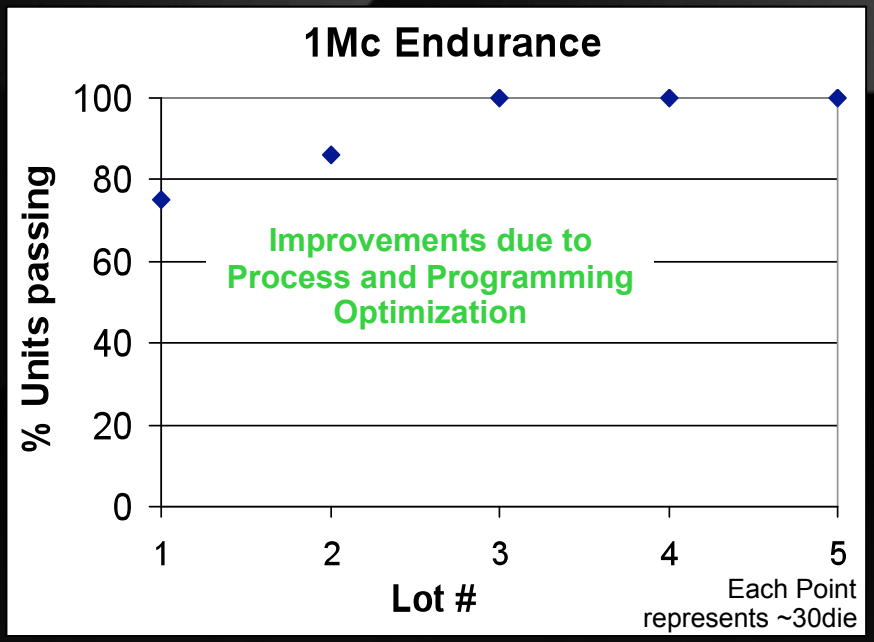

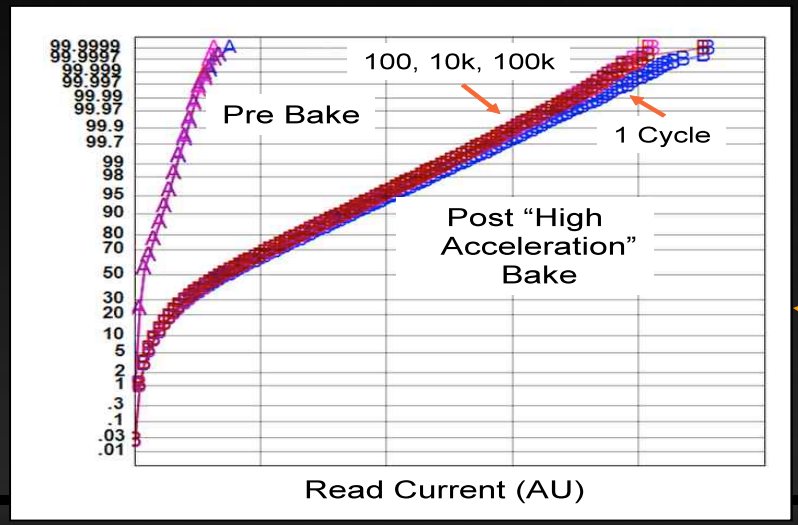

Data Retention and Endurance Look Good

Retention is Independent of Endurance

numonyx™

Use Condition Sensitivities by Technology

	NAND	NOR/NROM	PCM
Program / Erase Cycles	SBC: 1e5 MLC: 1e4 Trending Down	1e5 Long term downward trend	1e6 Trending up
Write Amplification	Very Important for data applications	Important	N/A
Cycling Temperature	Important	Important	Unimportant < 85C
Retention Temperature	Important	Important	Unimportant < 85C
Reads between cycles	SBC: 1e6 MLC: 1e5 Trending down	>1e13 / page Trending down slightly	Infinite
Partial Page Program Cycles	1-4, trending to 1	N/A	N/A
Tolerable Bad Block Rate	2%	N/A	N/A
System ECC Capability	Important	N/A	N/A
Solder Temperature	Not Important	Not Important	Important

Reliability Concerns for Large Arrays

- Persistent Errors
  - DRAM and SRAM susceptible to silent data corruption of memory cells caused by radiation effects
  - PCM storage element immune to these radiation effects
  - With a thermally activated storage mechanism, PCM cells will tend to change state at high temperatures. Techniques similar to “Scrubbing” used in high-end DRAM applications can serve to extend the PCM operating range.
- Non-Persistent Errors
  - Logic and interface errors expected to be common mode between DRAM and PCM
  - Use of ECC to fix these errors, however, is straightforward
- PCM does not require refresh of idle banks
  - DRAM refresh rates expected to increase with scaling

PCM may prove to be much more manageable at very large densities

Agenda

- Background / Why PCM
- What is PCM and How it Works
- Technology Characteristics and Trends
- Application: Directly Addressable NVRAM
- Application: Storage Input / Output Acceleration

Performance & Density Comparisons

Circa 2011, 45nm Silicon

Attributes	DRAM	PCM	NAND	HDD
Non-Volatile	No	Yes	Yes	Yes
Idle Power	~100mW/GB	~1 mW/GB	~10 mW/GB	~10W/TB
Erase / Page Size	No / 64Byte	No / 64Byte	Yes / 256KB	No / 512Byte
Write Bandwidth*	~GB/s per die	50-100 MB/s per die	5-40 MB/s per die	~200MB/s per drive
Page Write Latency	20-50 ns	~1 us	~500 us	~5 ms
Page Read Latency	20-50 ns	~50 ns	~25 µs	~5 ms
Endurance	∞	$10^6 \rightarrow 10^7$	$\rightarrow 10^4$	∞
Maximum Density	4Gbit	4Gbit	64Gbit	2TByte

500x

500x

1000x

16x

*Assumes multiple banks

Numonyx PCM in Compute

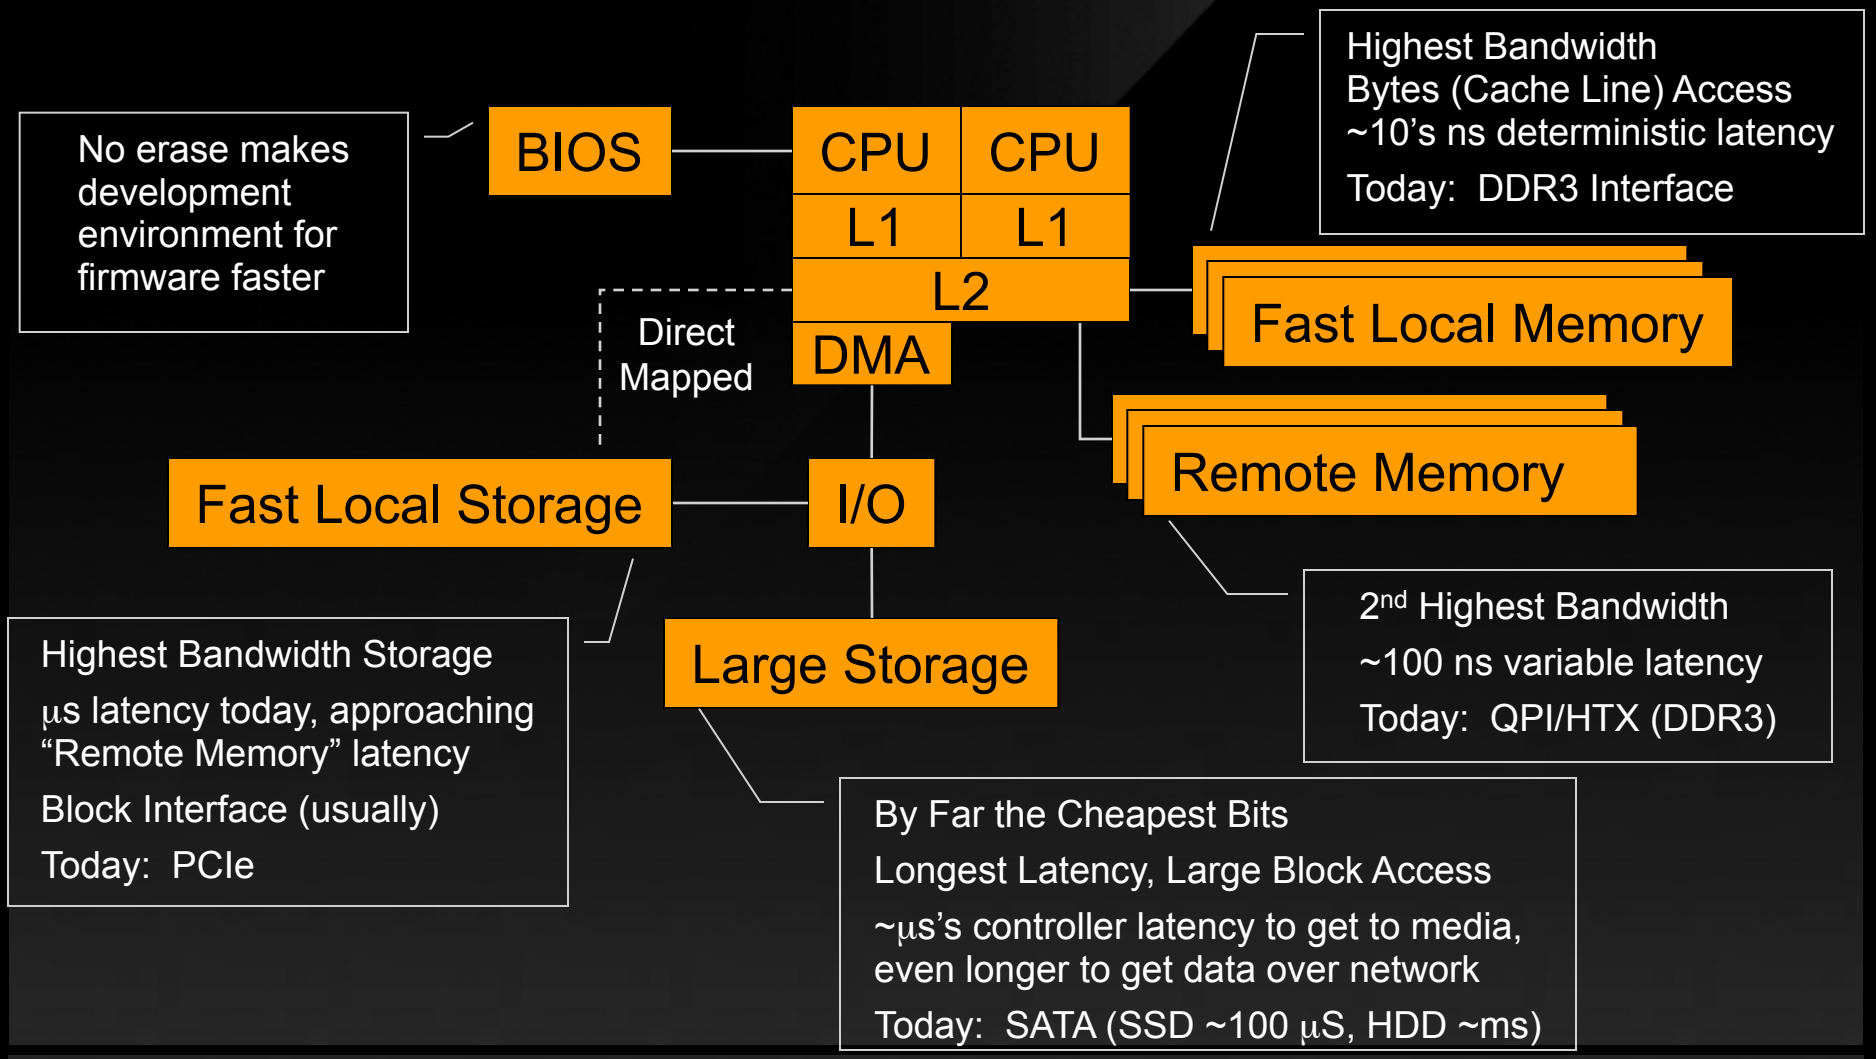

Numonyx PCM in Compute

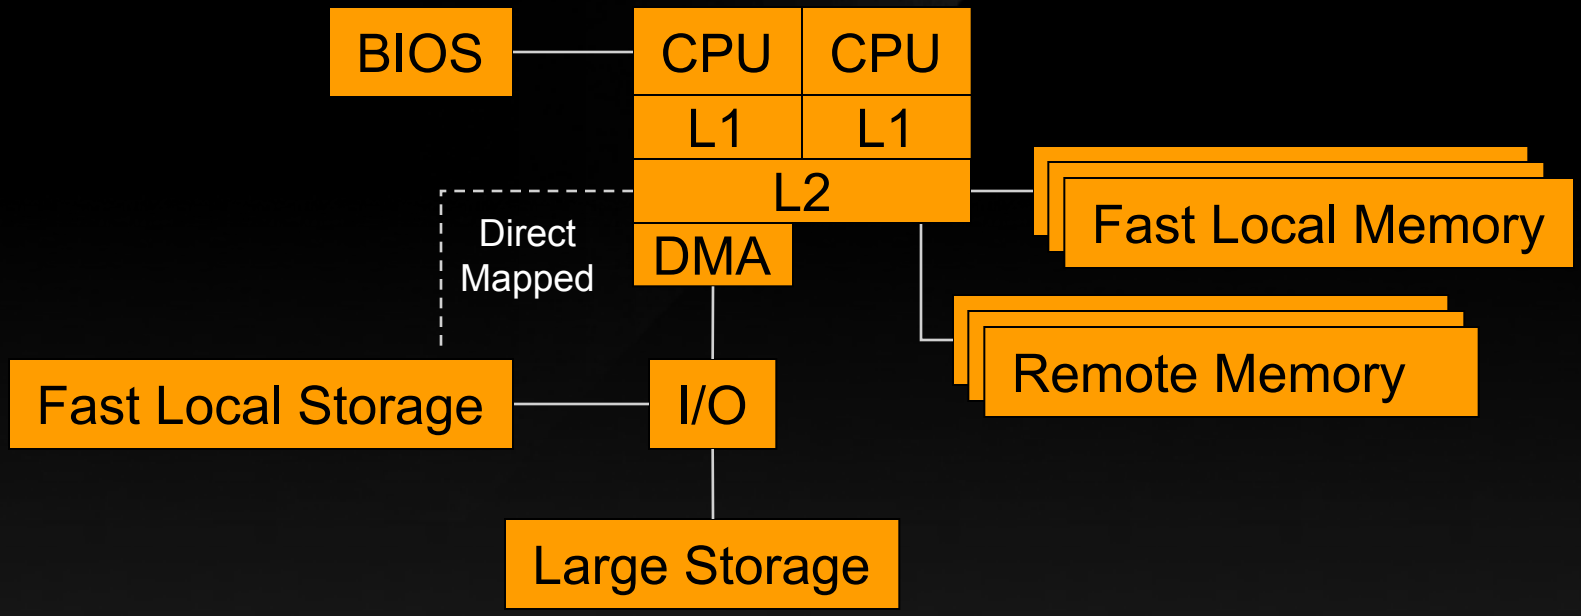

Numonyx PCM in Memory

PCM Fit

System Latency	Local is Lowest
System Access	Bytes (Cache Line)
Hardware Changes	Yes, DDRx-NVM
Software Changes	Yes, for write, NV

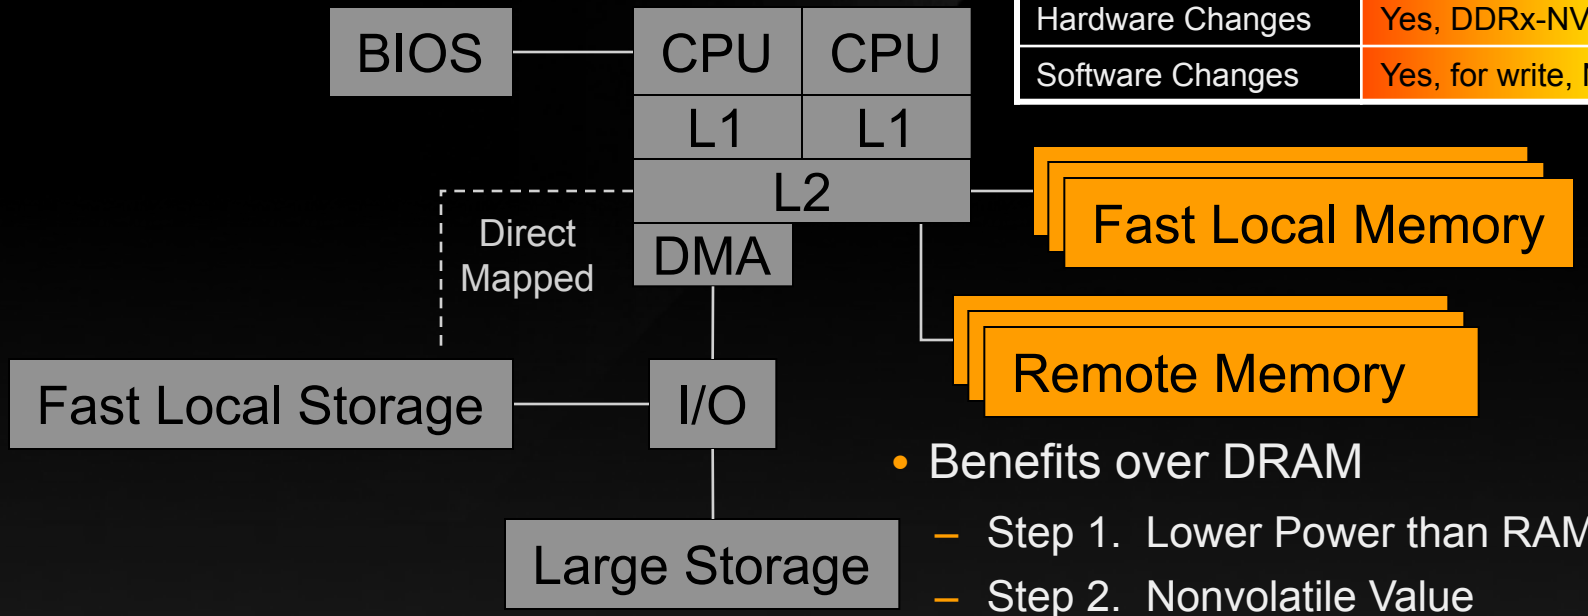

- Benefits over DRAM
  - Step 1. Lower Power than RAM
  - Step 2. Nonvolatile Value
    - Reduced Software Cache density
    - Simplified Software PLR
    - Fast, efficient, access of only necessary bytes
  - Tradeoffs
    - Write Bandwidth, Write Latency vs. DRAM
    - Time/TTM (OS and controller changes required)

Power and Energy

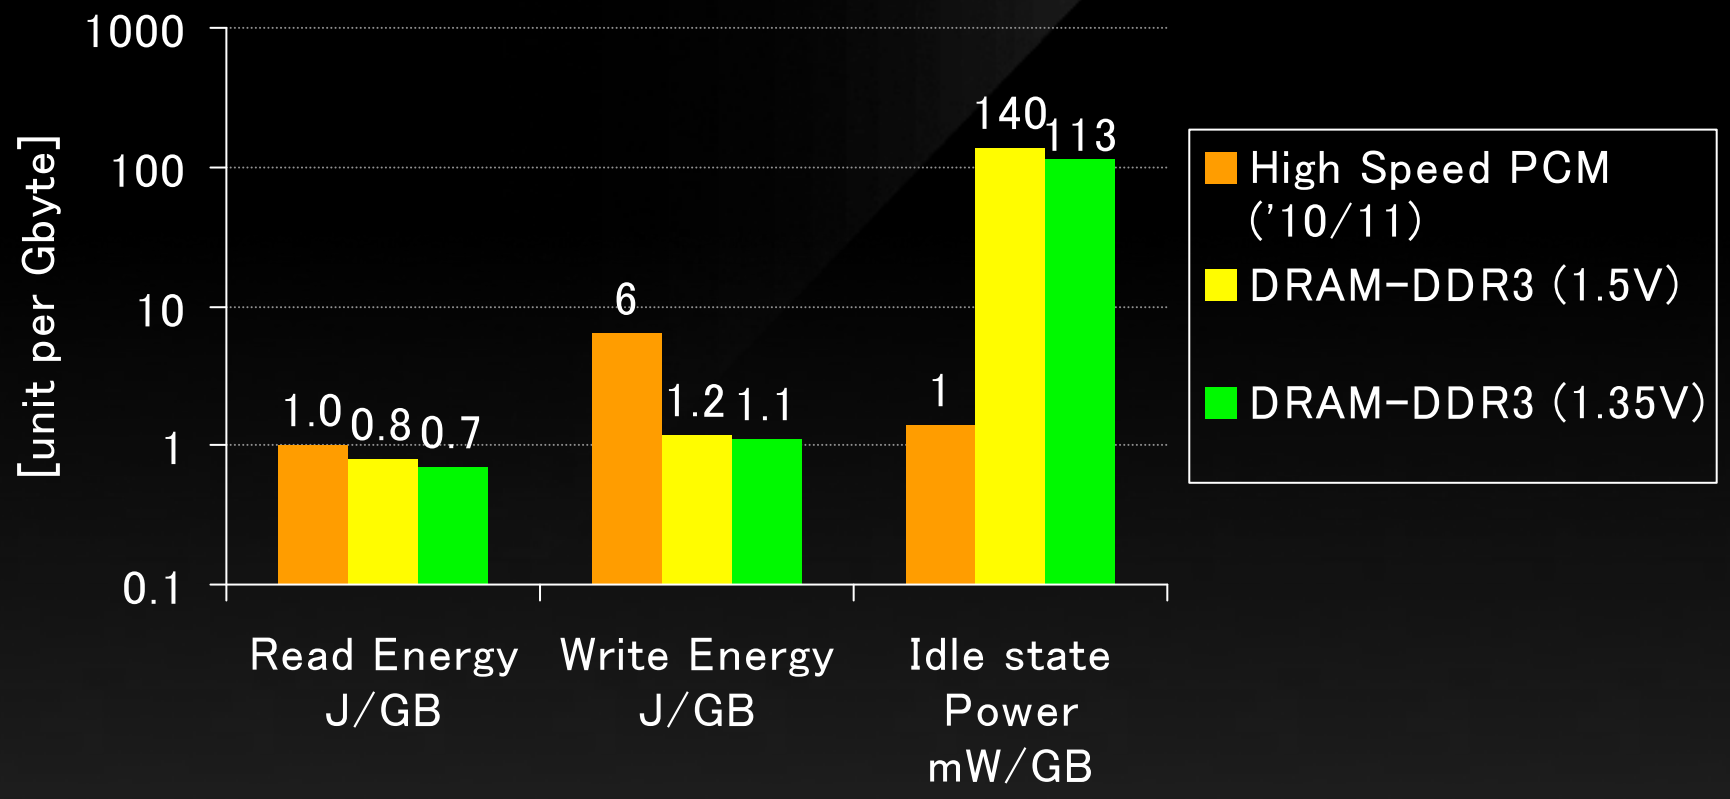

Energy Crossover Example

Assumptions: 25.6GB/s Bandwidth, 50% bus utilization; 70% read, 30% write)

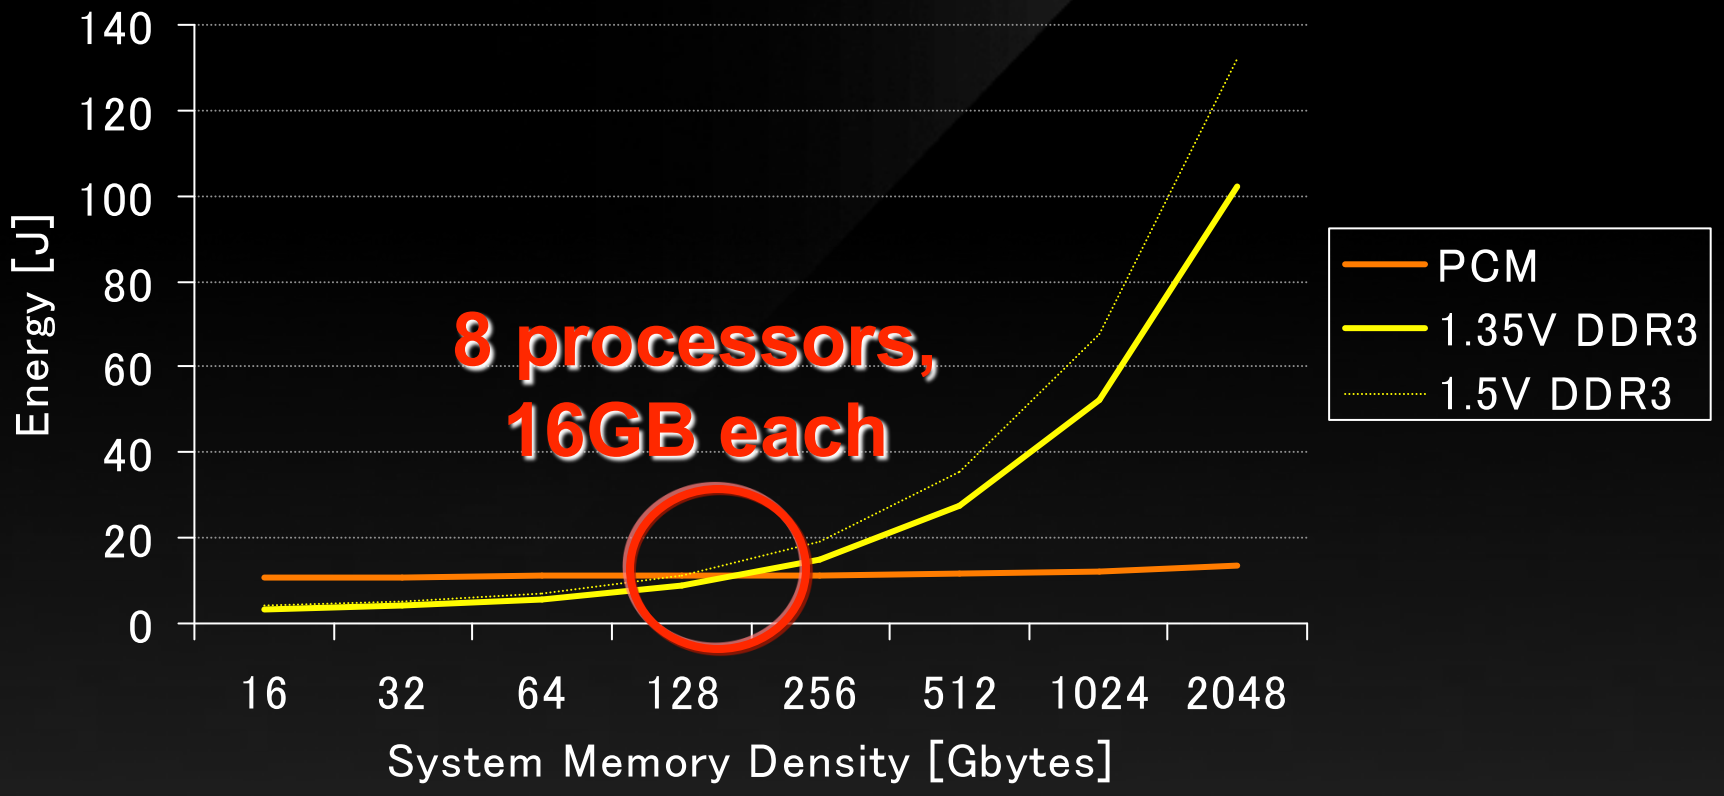

numonyx™

Latency: Random Read

Numonyx PCM in Compute

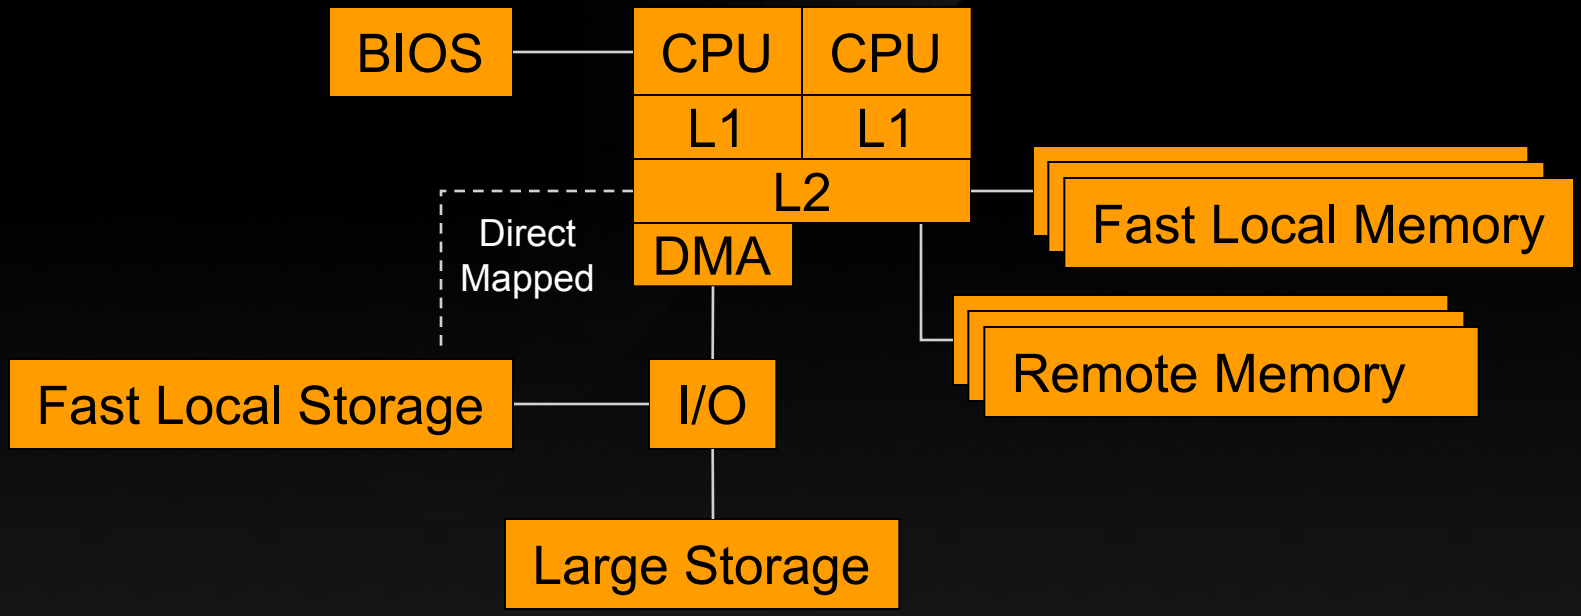

Numonyx PCM in Storage

PCM Fit

System Latency	µs >> 100ns
System Access	≥4kB (typically)
Hardware Changes	No
Software Changes	No* (not optimized)

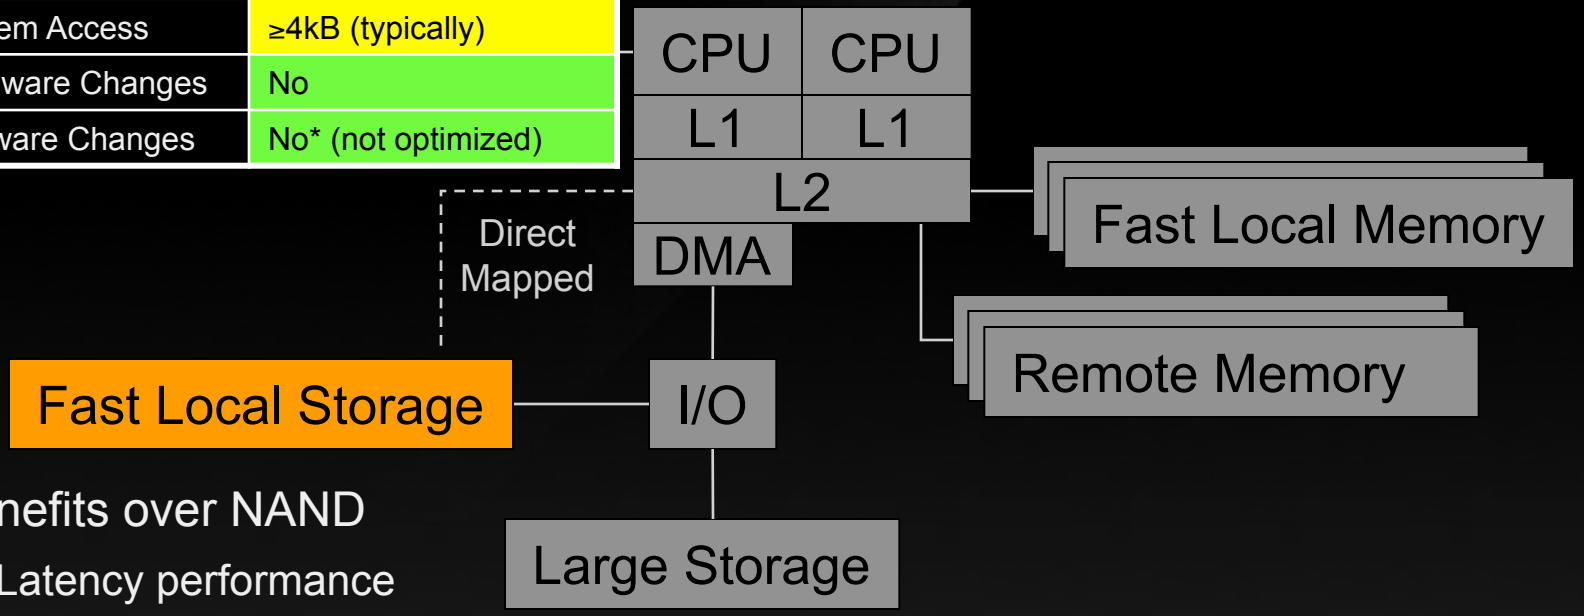

- Benefits over NAND
  - Latency performance
  - Physical (Bandwidth per Chip)
  - Reliability and Scaling
  - Capable of Smaller IO for meta data
  - Tradeoffs
    - Density and $/GB vs NAND

Latency: Random Read

Latency: Random Read

Latency: Random Write

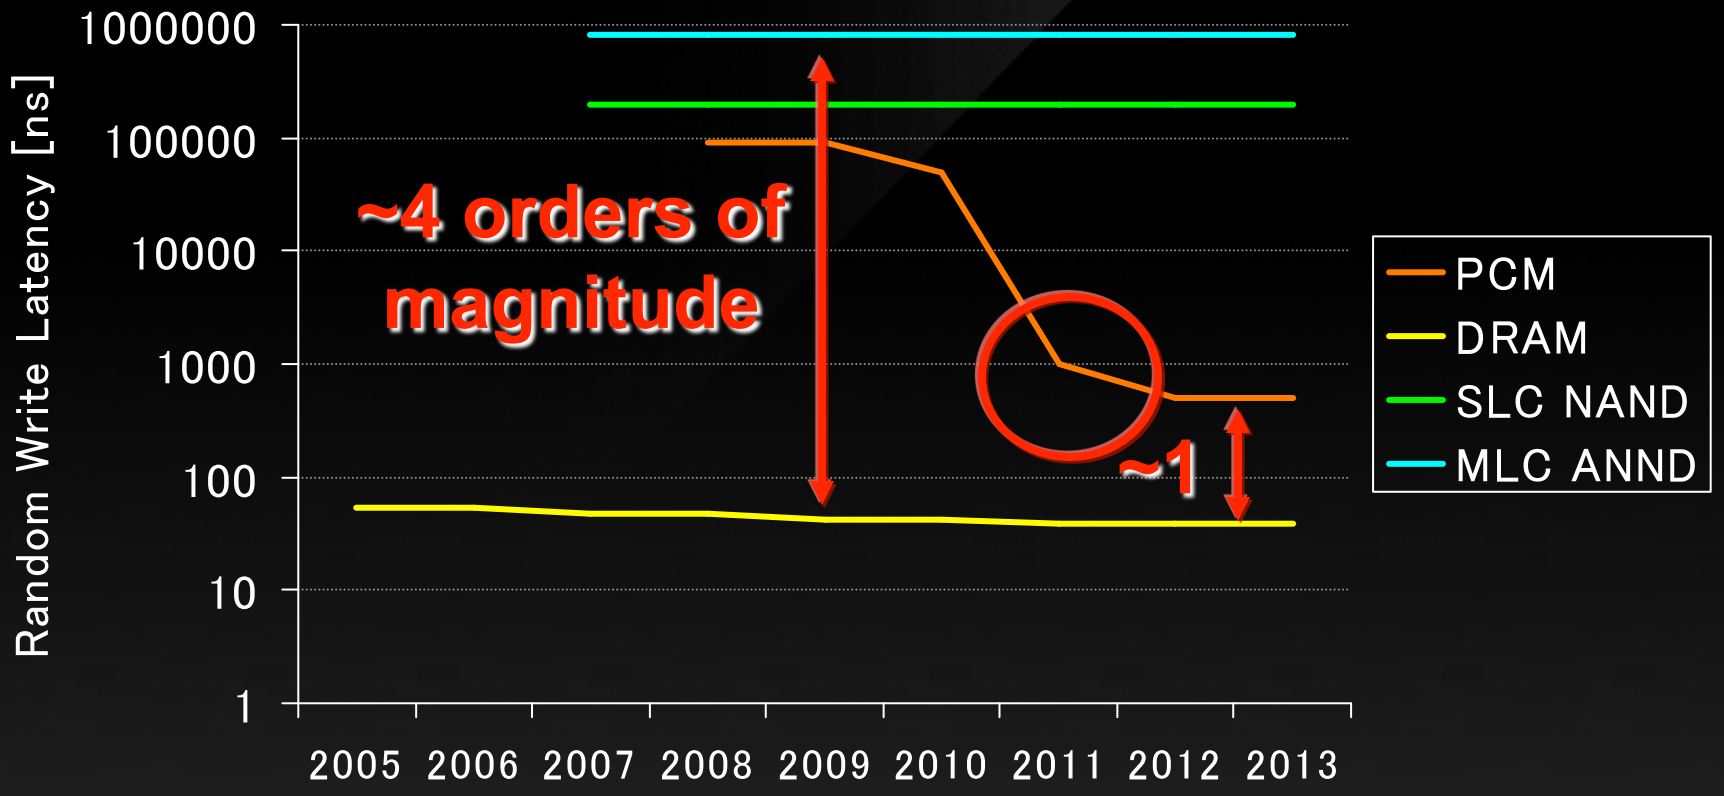

PCIe Form Factor – Pre-Production Vehicle

*mechanical concept

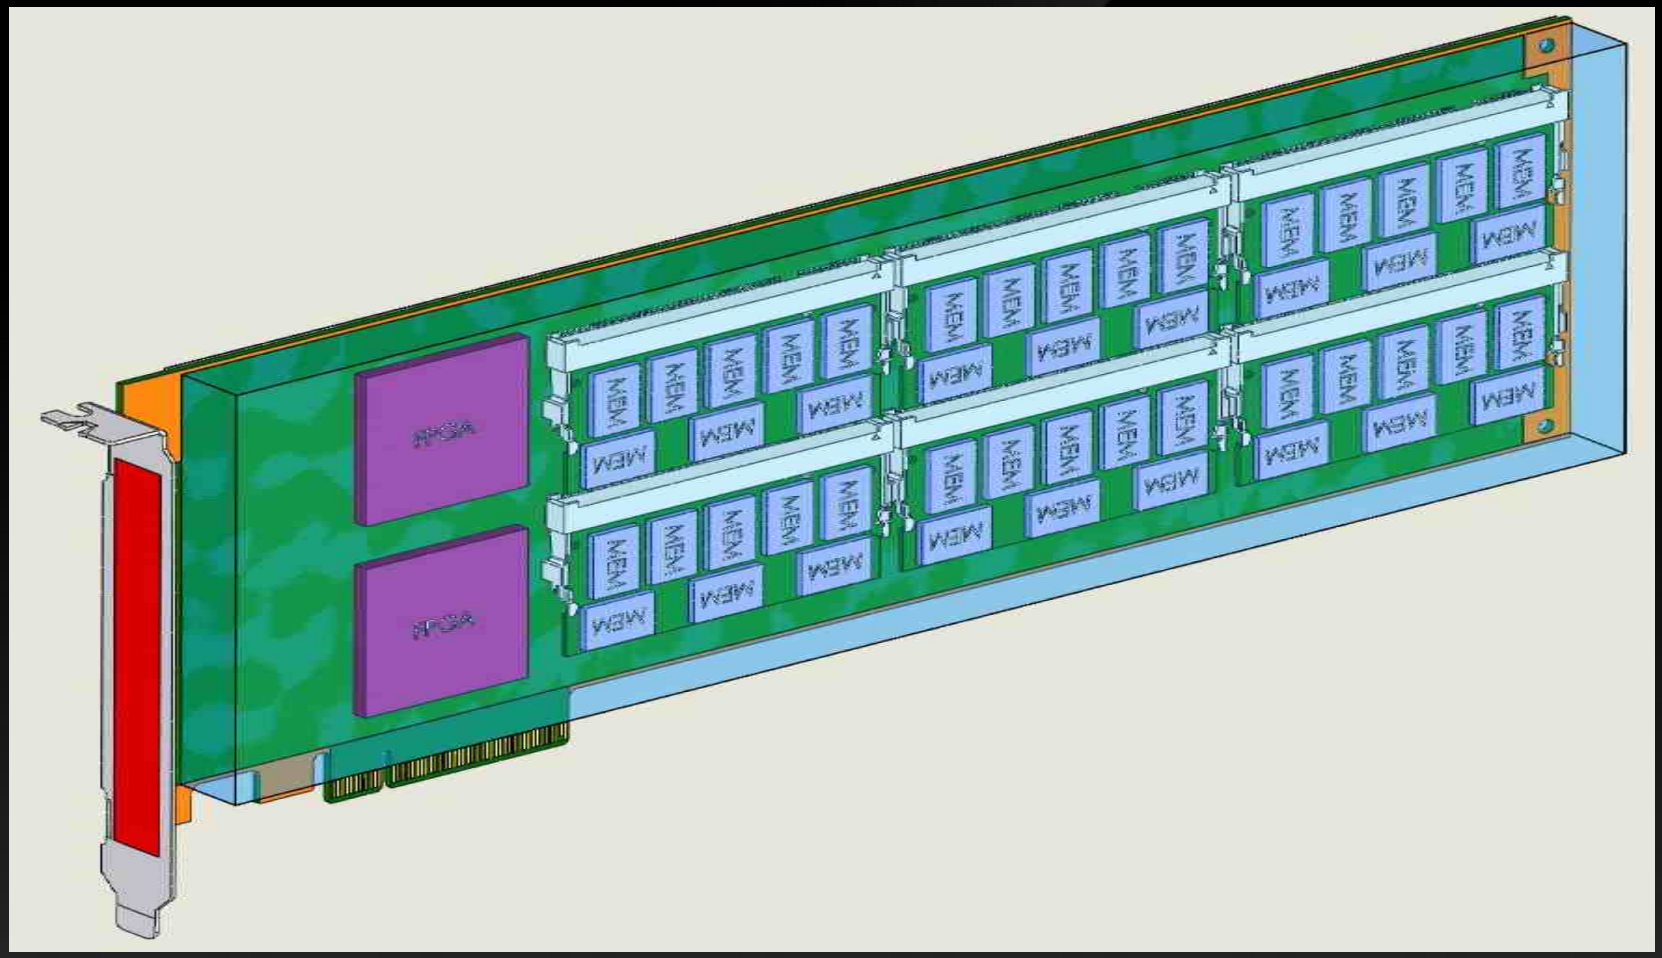

Bandwidth Per Chip: Read

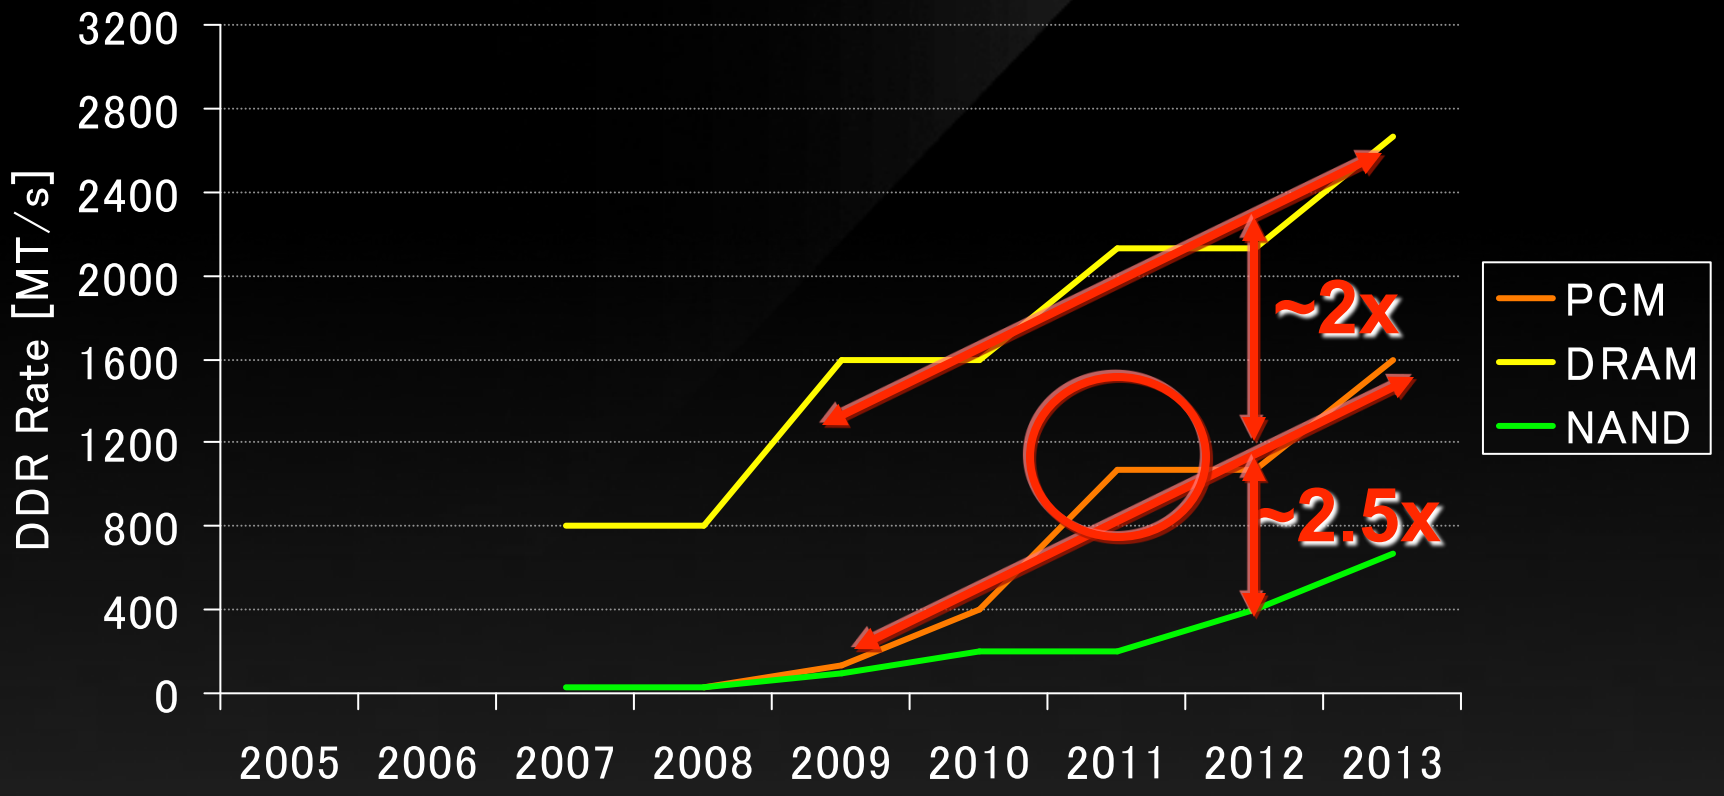

Bandwidth per Chip: Write

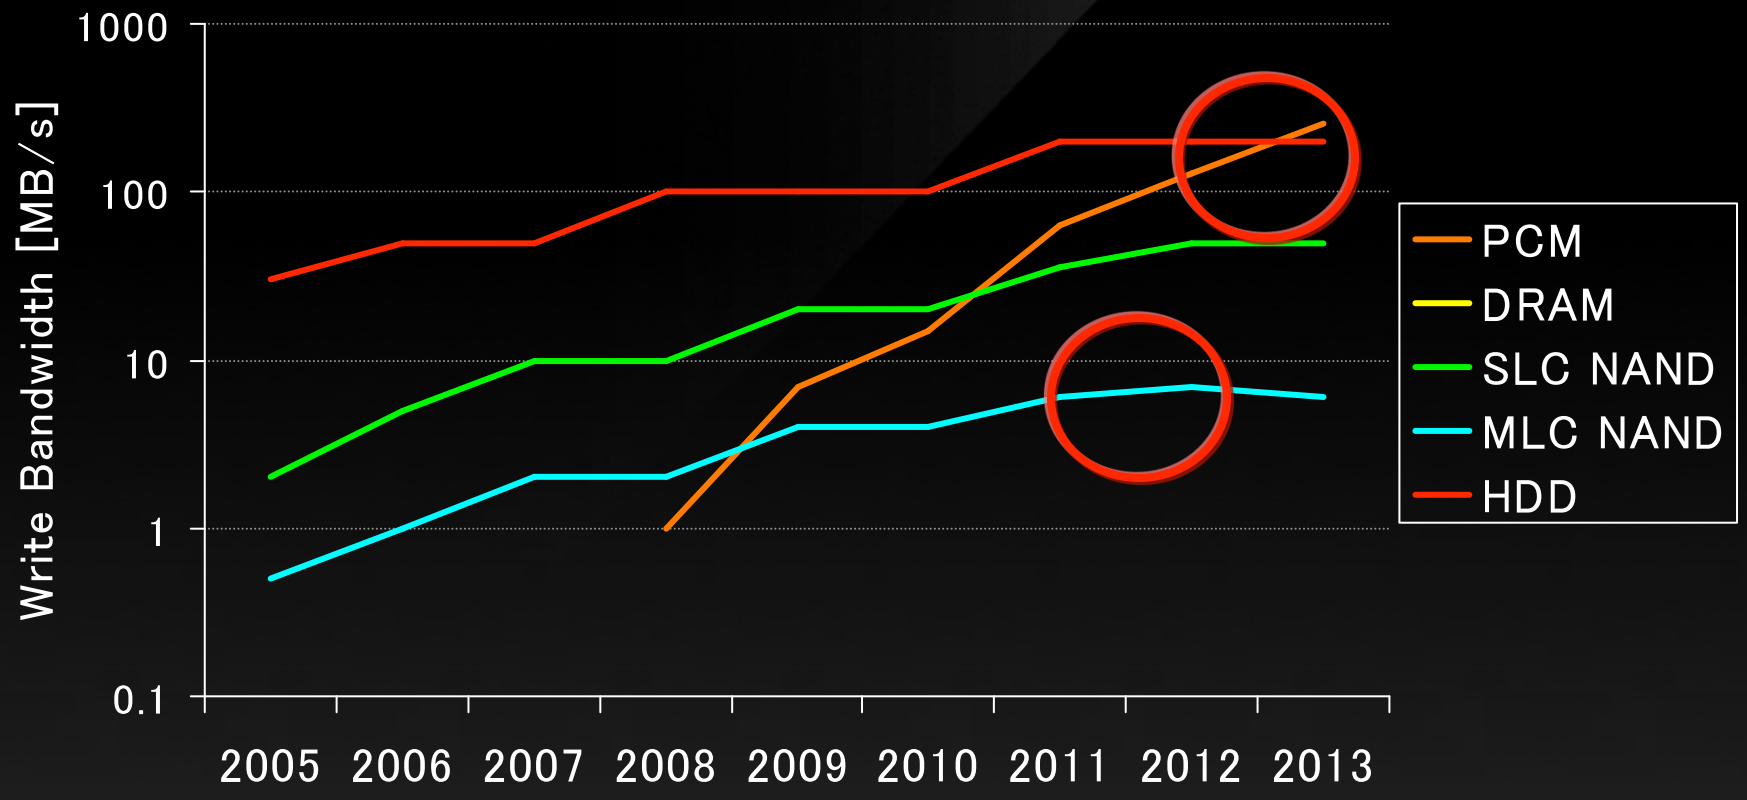

First (Easiest) Use Model: IO cache Cards

	2010
Min Density	64GB
Max Density	512GB
Interface	PCIe 2.0x8
Read Bandwidth	4.0 GB/s
Write Bandwidth	400 MB/s
Input Voltage	Up to 12V
Power	20W
Read Latency	5μS (hw)
Write Latency	150μS
Physical Dimensions	Full Size PCIe Card
Temperature	0°to +55°C

Numonyx PCM in Compute

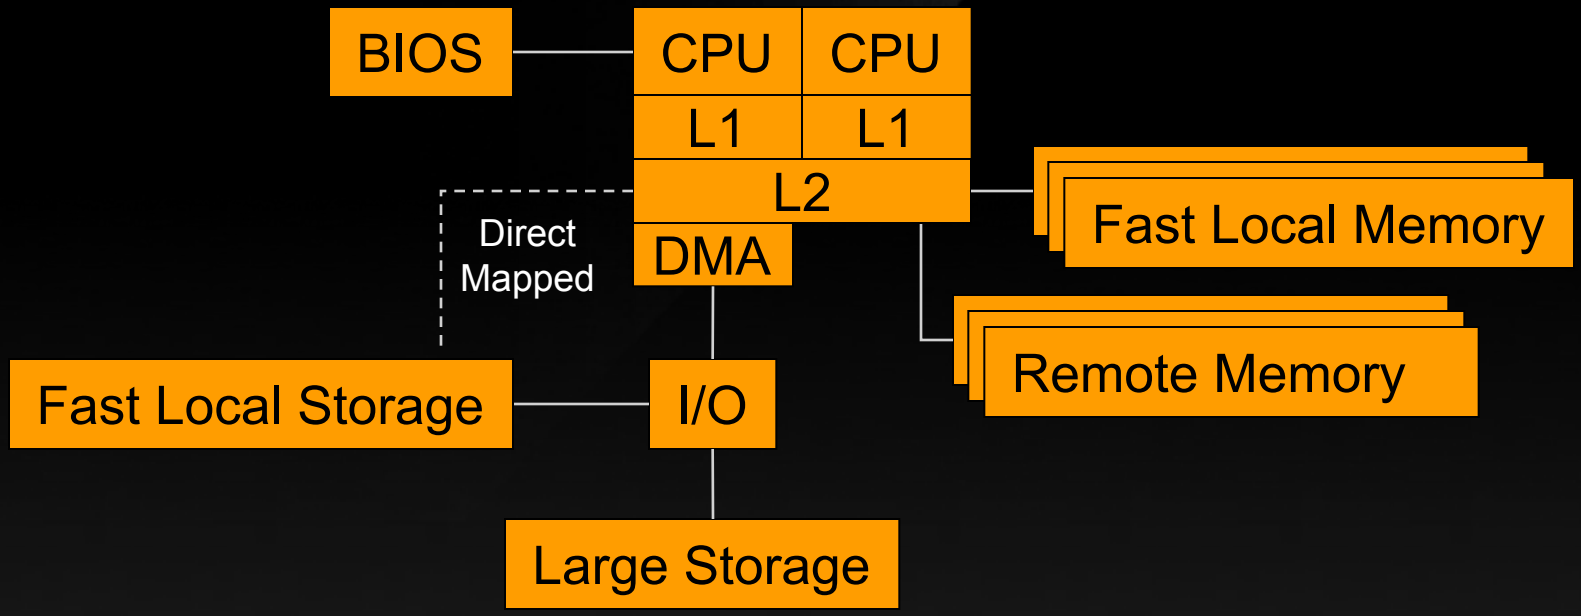

Numonyx PCM in Storage

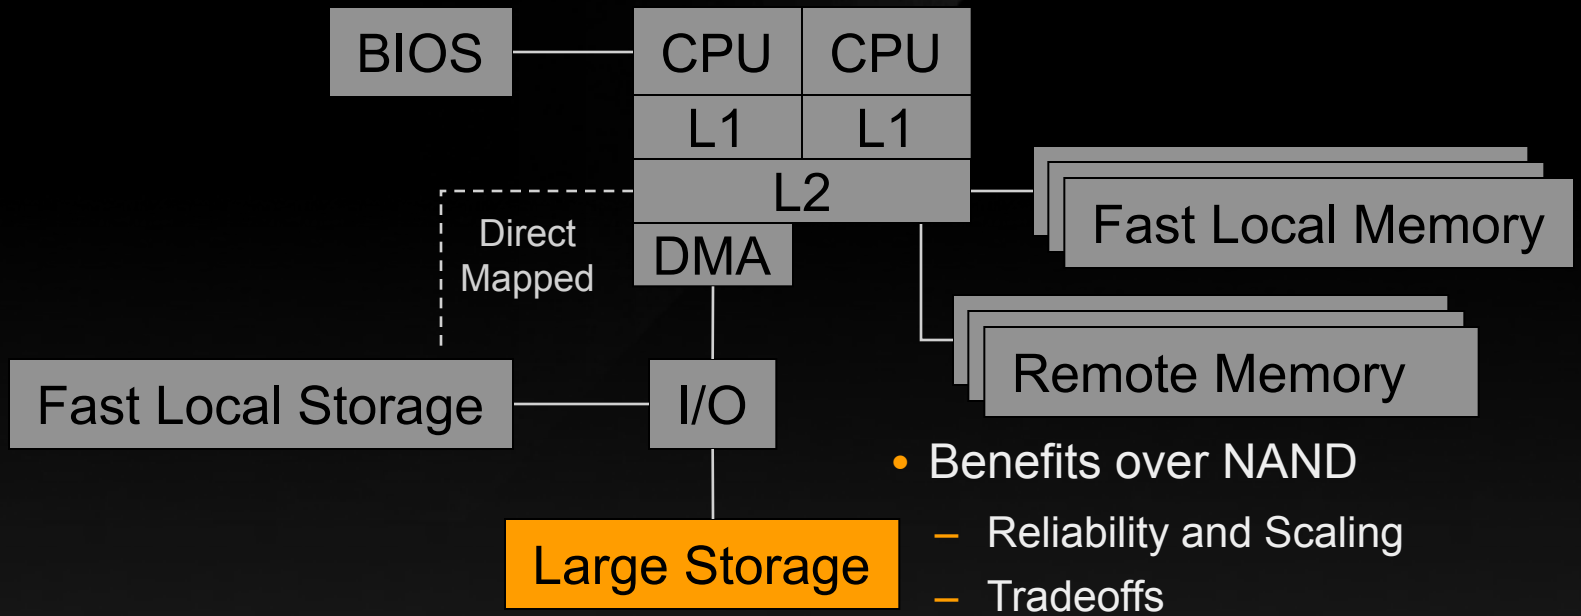

- Benefits over NAND
  - Reliability and Scaling
  - Tradeoffs
    - Density, $/GB vs NAND

System Latency	Highest
System Access	Really Large Blocks
Hardware Changes	No
Software Changes	No* (not optimized)

Random Performance also seems valuable

- “Even in HPC, high IOPS is becoming a necessity. Multicore computing means structured data access is becoming randomized, since I/O now tends to be performed across multiple threads. In particular, metadata is becoming the choke point for large clustered file systems since it tends to be accessed randomly.”

Source: July 01, 2009, “A Trio of HPC Offerings Unveiled at ISC” by Michael Feldman, HPCwire Editor

In Summary

- We think PCM Power has value in Memory
  - But it will take a while for software to catch up
- We think PCM Latency has value in a subset of storage
  - Up to ~50-100x faster than the SSD's today
  - Small IO access performance (for databases) gets more interesting as Fast Storage and Remote Memory Latencies converge
- We think PCM will out scale (outlast) existing technologies
  - More opportunities will open up if technologies start to have difficulty

Questions?

numonyx™

innovative. memory. solutions.