STORAGE DEVELOPER CONFERENCE
SDC 22
Fremont, CA
September 12-15, 2022
BY Developers FOR Developers

A SNIA Event

Persistent Memories: Without Optane, Where Would We Be?

Thomas Coughlin, Coughlin Associates
Jim Handy, Objective Analysis

Outline

- The Rise and Fall of Intel's Optane
- Optane's Bountiful Legacy
- Future Chips and Persistence
- Meet the Memories
- How Future Processors Will Benefit
- Summary
- Q&A

The Rise and Fall of Intel's Optane

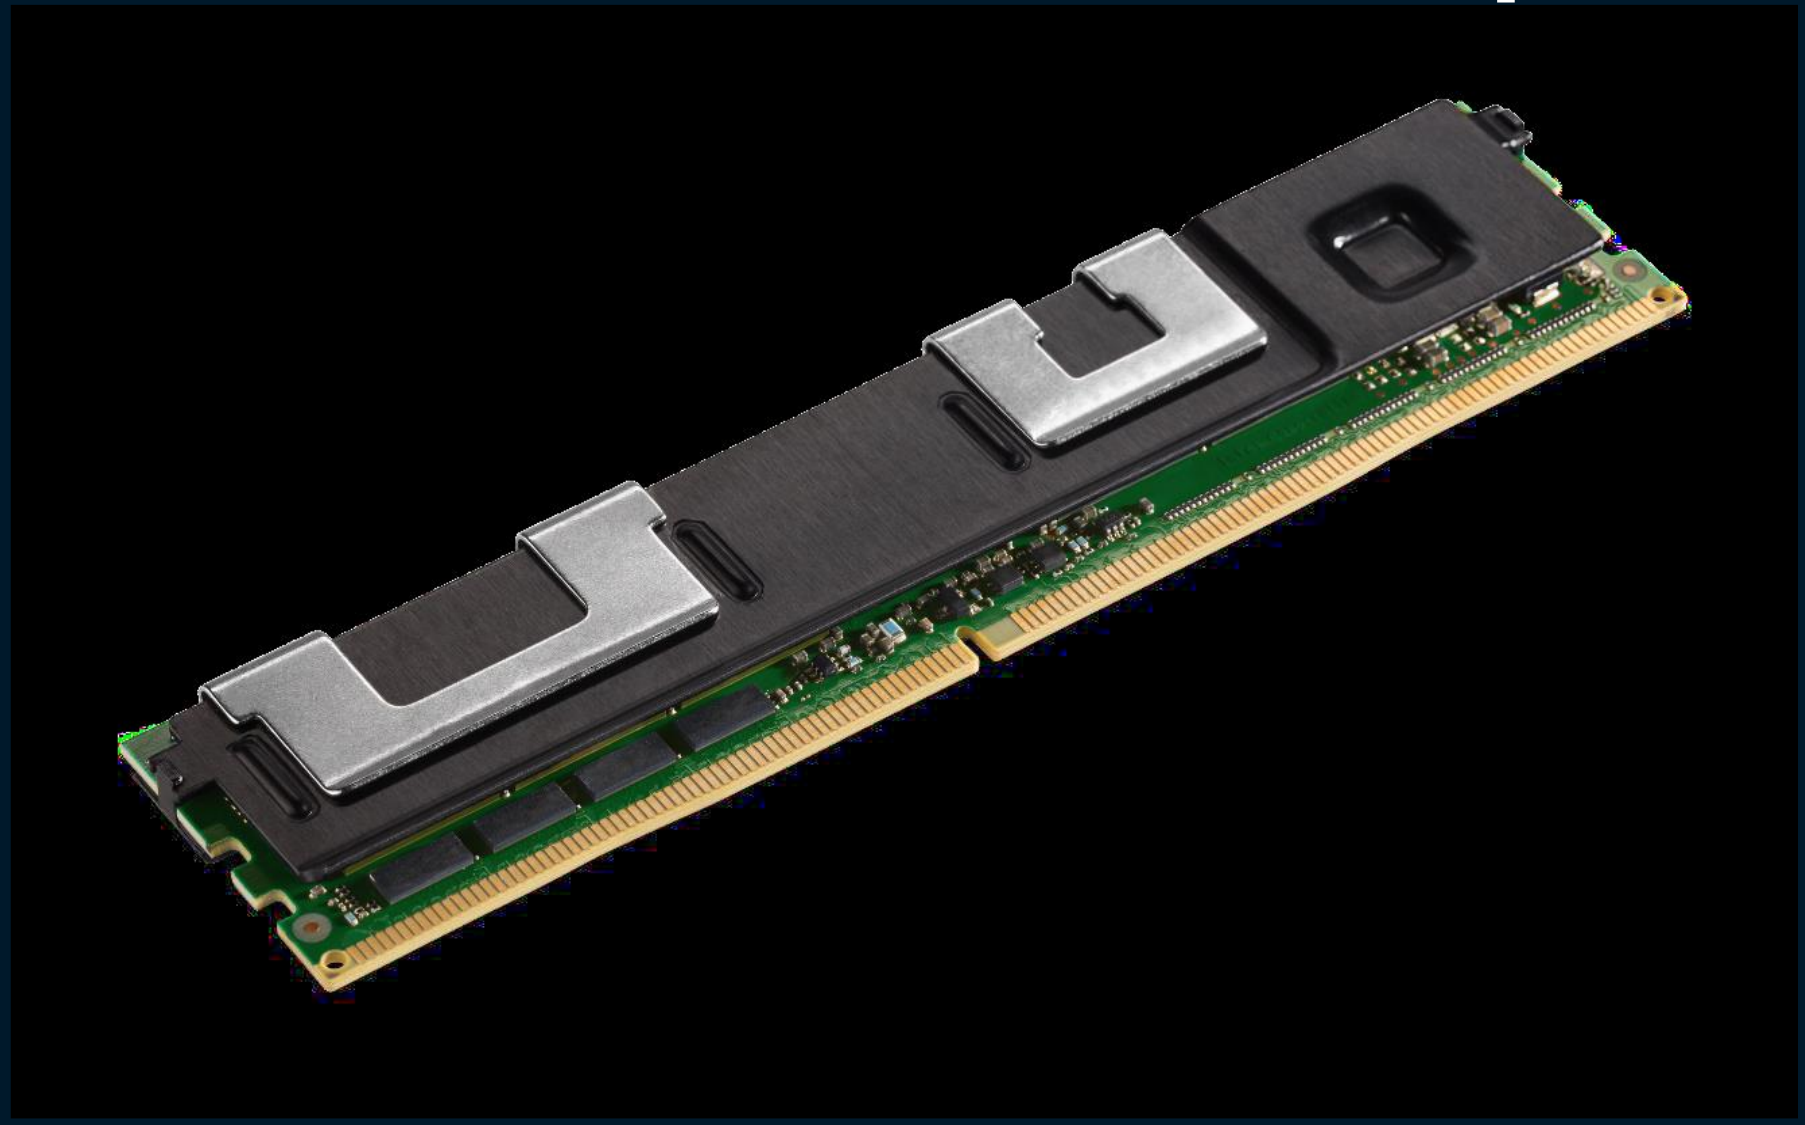

Intel Has a Long PCM History

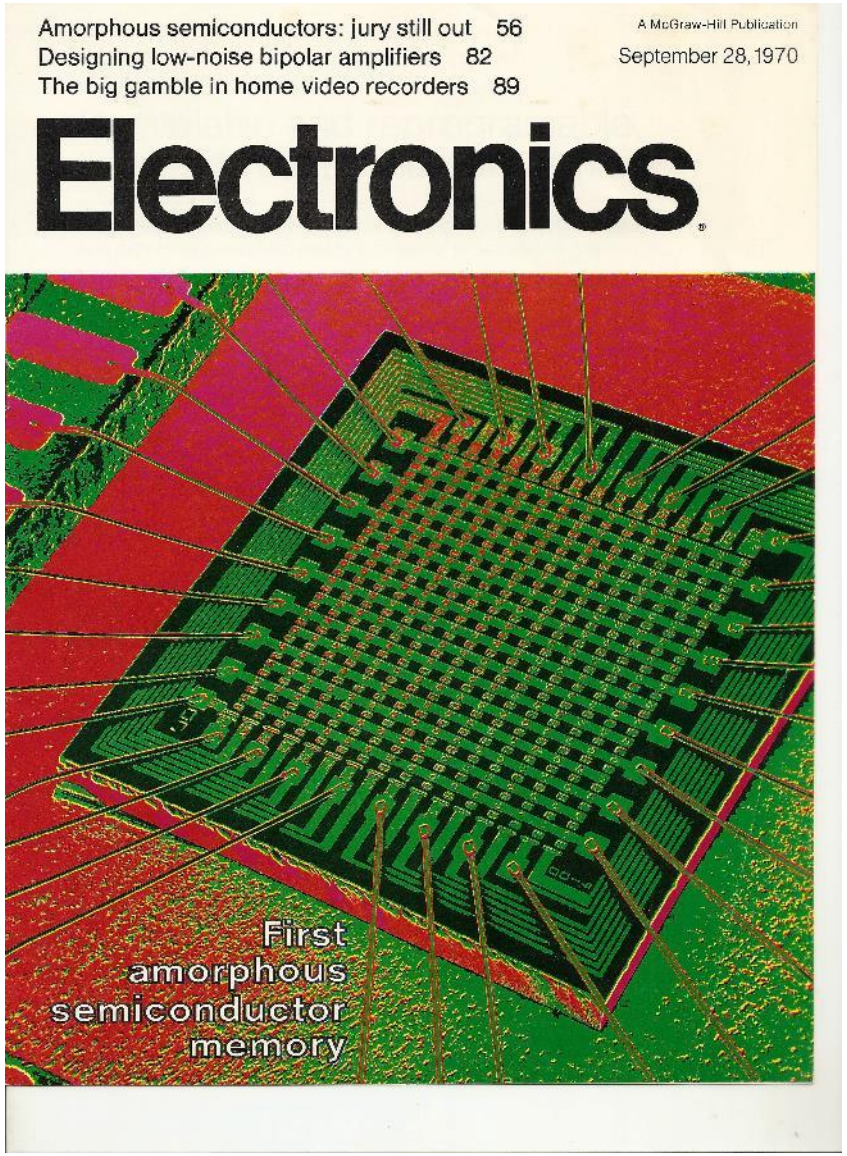

- 1970 first 256-bit chip
  - Cover story in Sept 1970 issue of Electronics
- 2007 first commercial product
  - 128Mb chip on 90nm process
  - Intended as NOR replacement
- 2015 3D XPoint
  - Based on same materials, but used different switching mechanism: “Bulk Switching”

Optane's Mission: Near-DRAM Speed at Sub-DRAM Prices

- 1,000X as fast as NAND
- 1/10th the die size of DRAM
- 1,000X NAND's endurance
- Fills a hole in the Memory/Storage Hierarchy

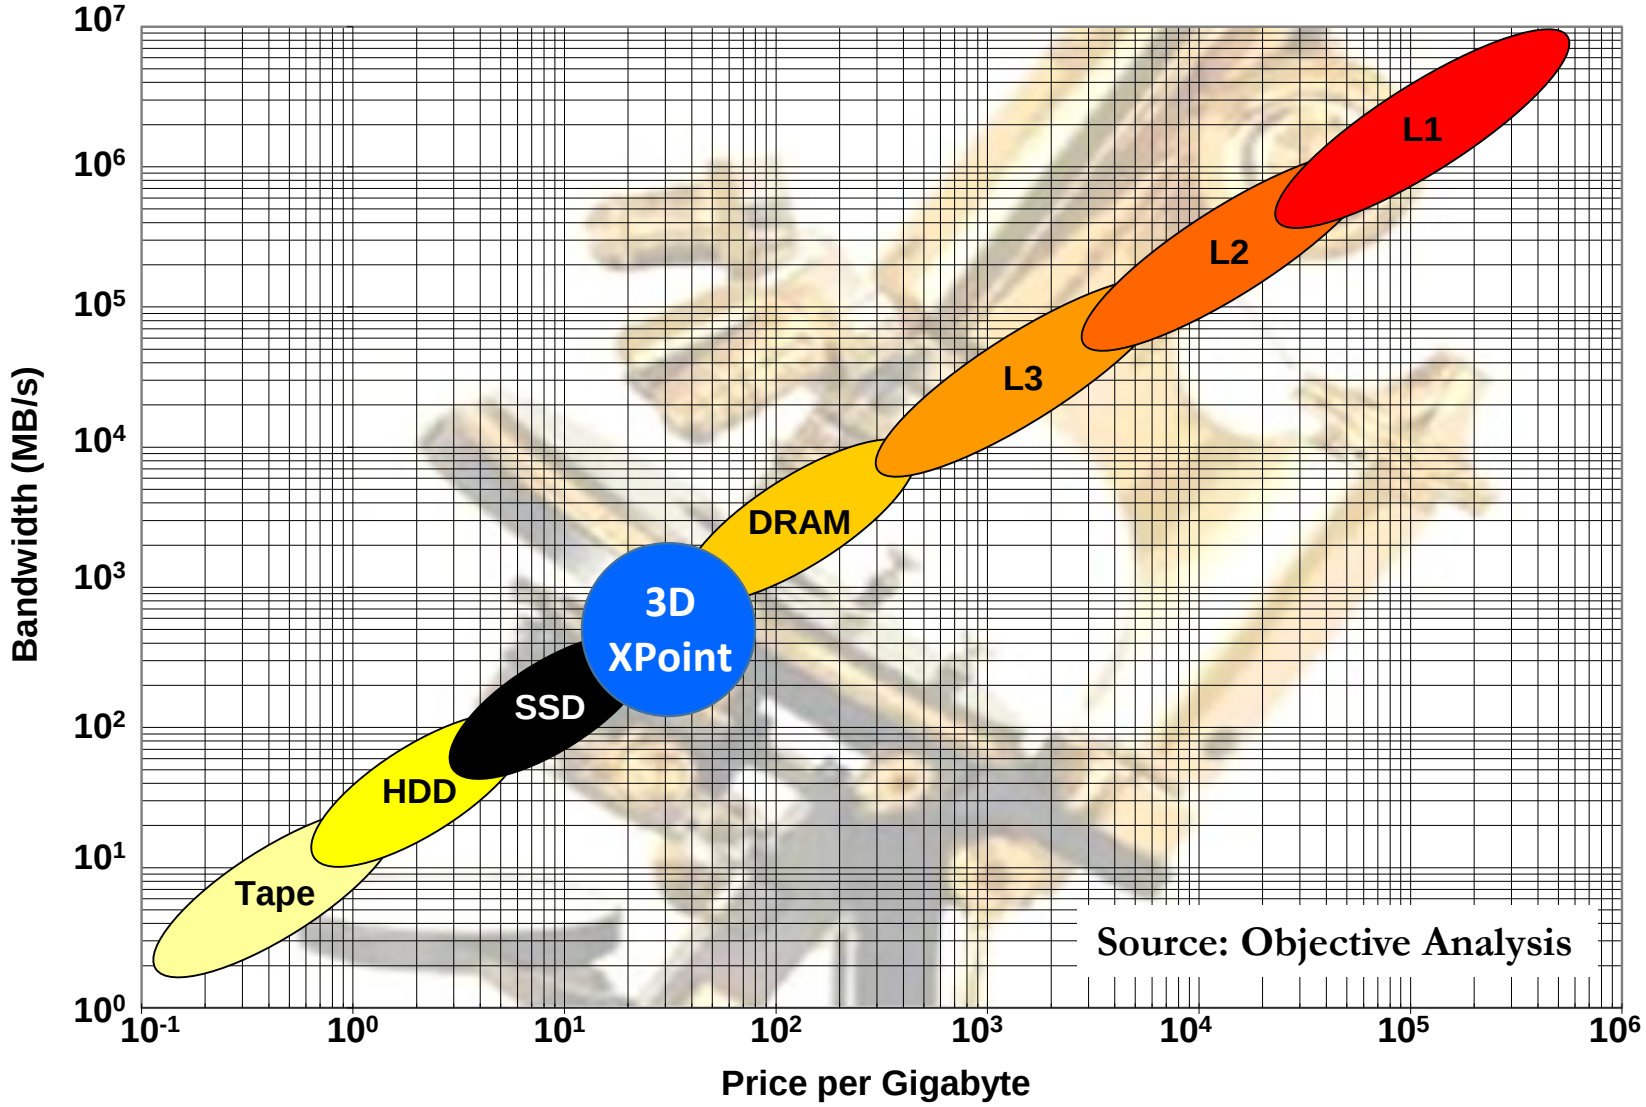

Cost Was the Hard Part

- Die size isn't everything
  - You also need the economies of scale
  - High volumes drive down production costs
- Chicken & egg problem
  - To get high volumes the price must be low
  - The cost won't be low until production volumes are high
- Intel committed to losing money to drive out the cost

NAND flash did this years ago. Let's look at that.

SLC NAND Die Sizes ~½ DRAM's

	DRAM		NAND	
	1Gb	2Gb	4Gb	8Gb
54nm		91.0mm²	98mm²	
50nm		93.1mm²		163mm²
48nm	28.6mm²			145mm²
44nm		51.6mm²		117mm²

This has bee

- At 54nm 4Gb NAND ≈ 2Gb DRAM
- At 50nm 8Gb NAND ≈ 2x2Gb DRAM
- At 48nm 8Gb NAND ≈ 4x1Gb DRAM
- At 44nm 8Gb NAND ≈ 2x2Gb DRAM

technology

Yet, NAND $/GB >DRAM's Until 2004

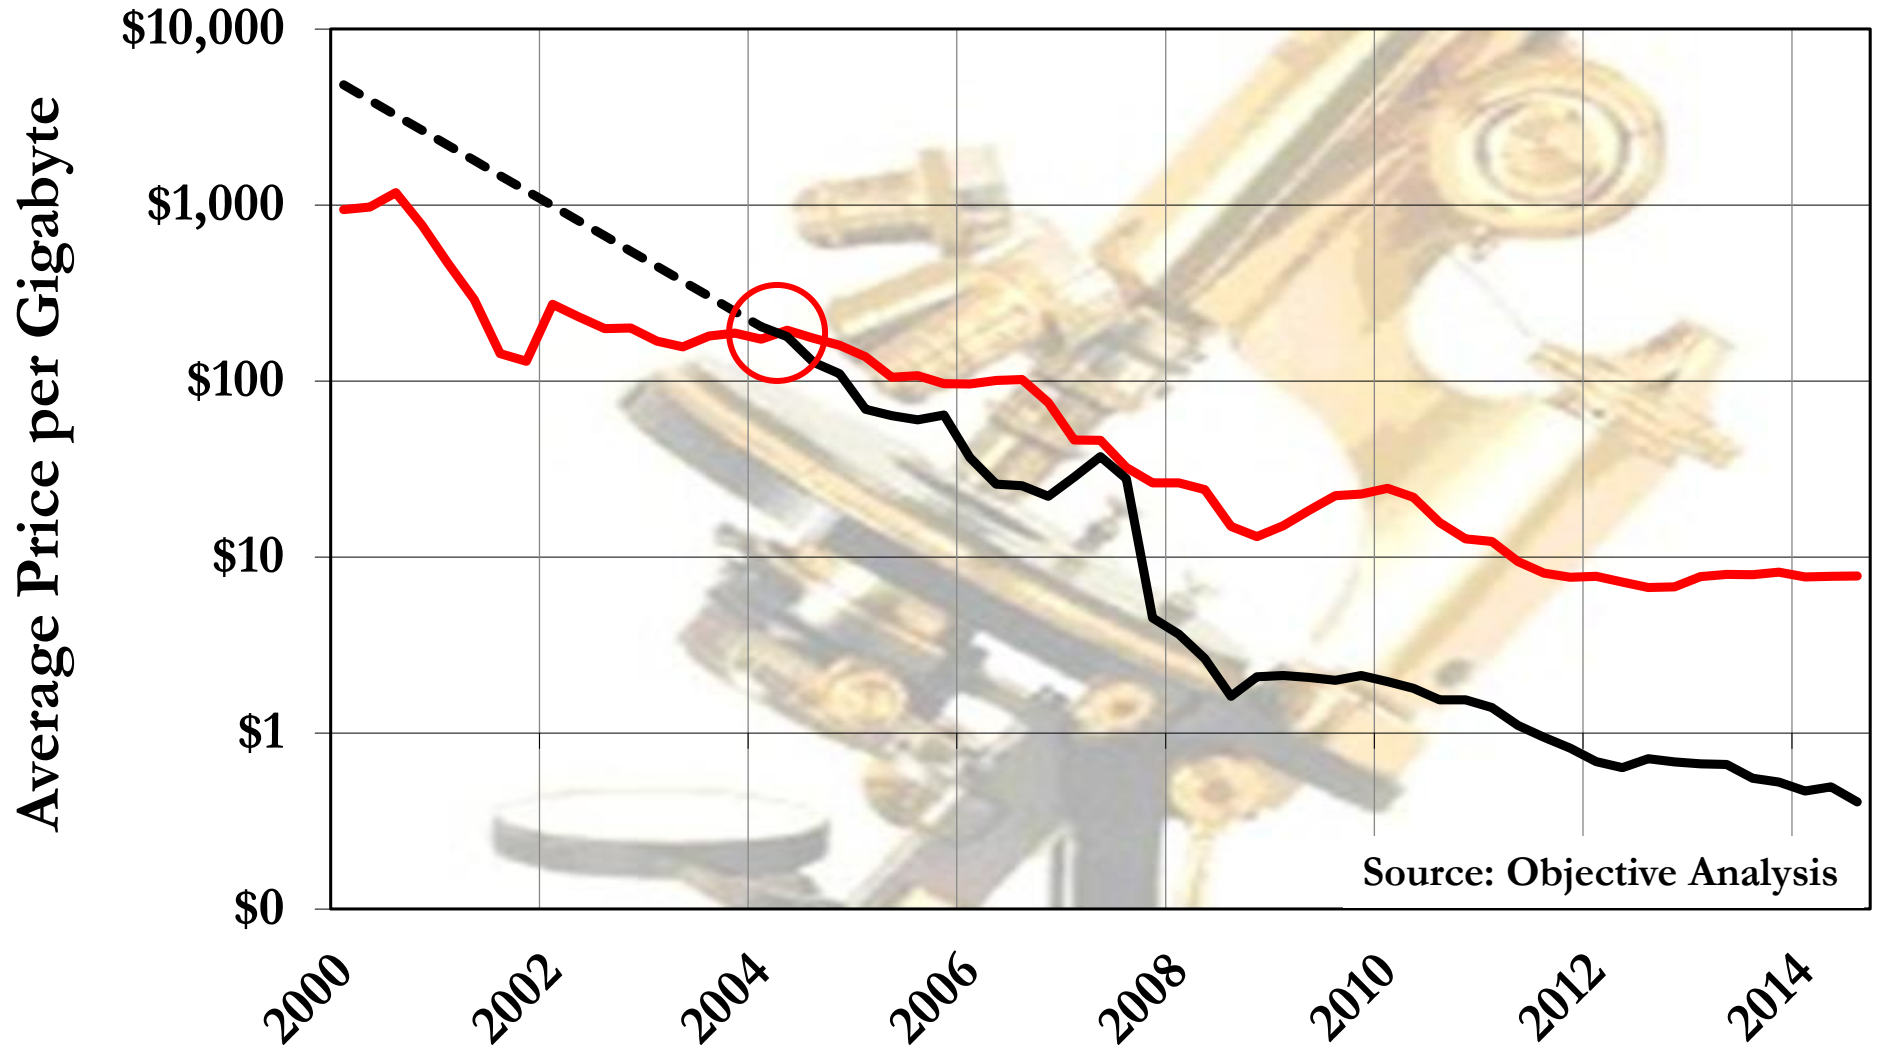

From: ***Hybrid Drives: How, Why, & When?***

In 2004 NAND TB Reached ⅓ That of DRAM

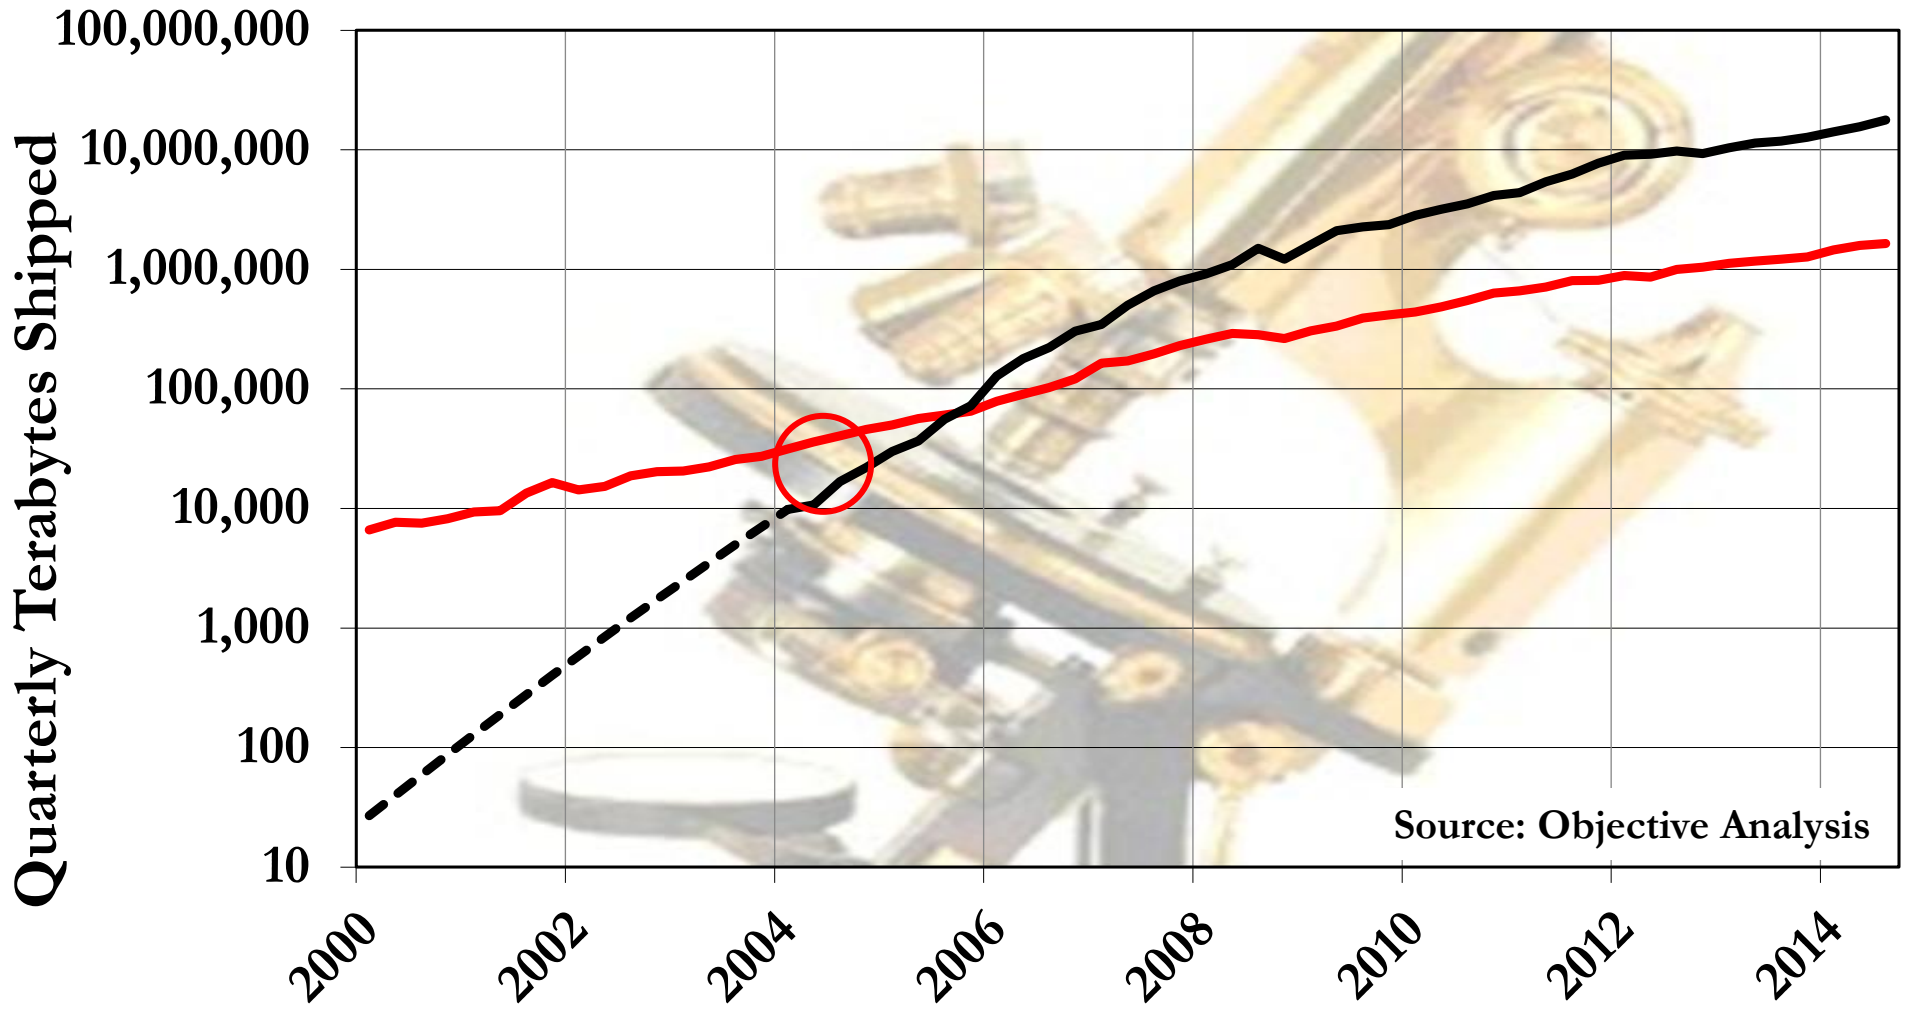

To Match DRAM Costs, Wafer Volume Must Be ~10% Of DRAM's

The Economies of Scale Cannot Be Ignored!

2022 Optane PB Optimistically 1/30th of DRAM

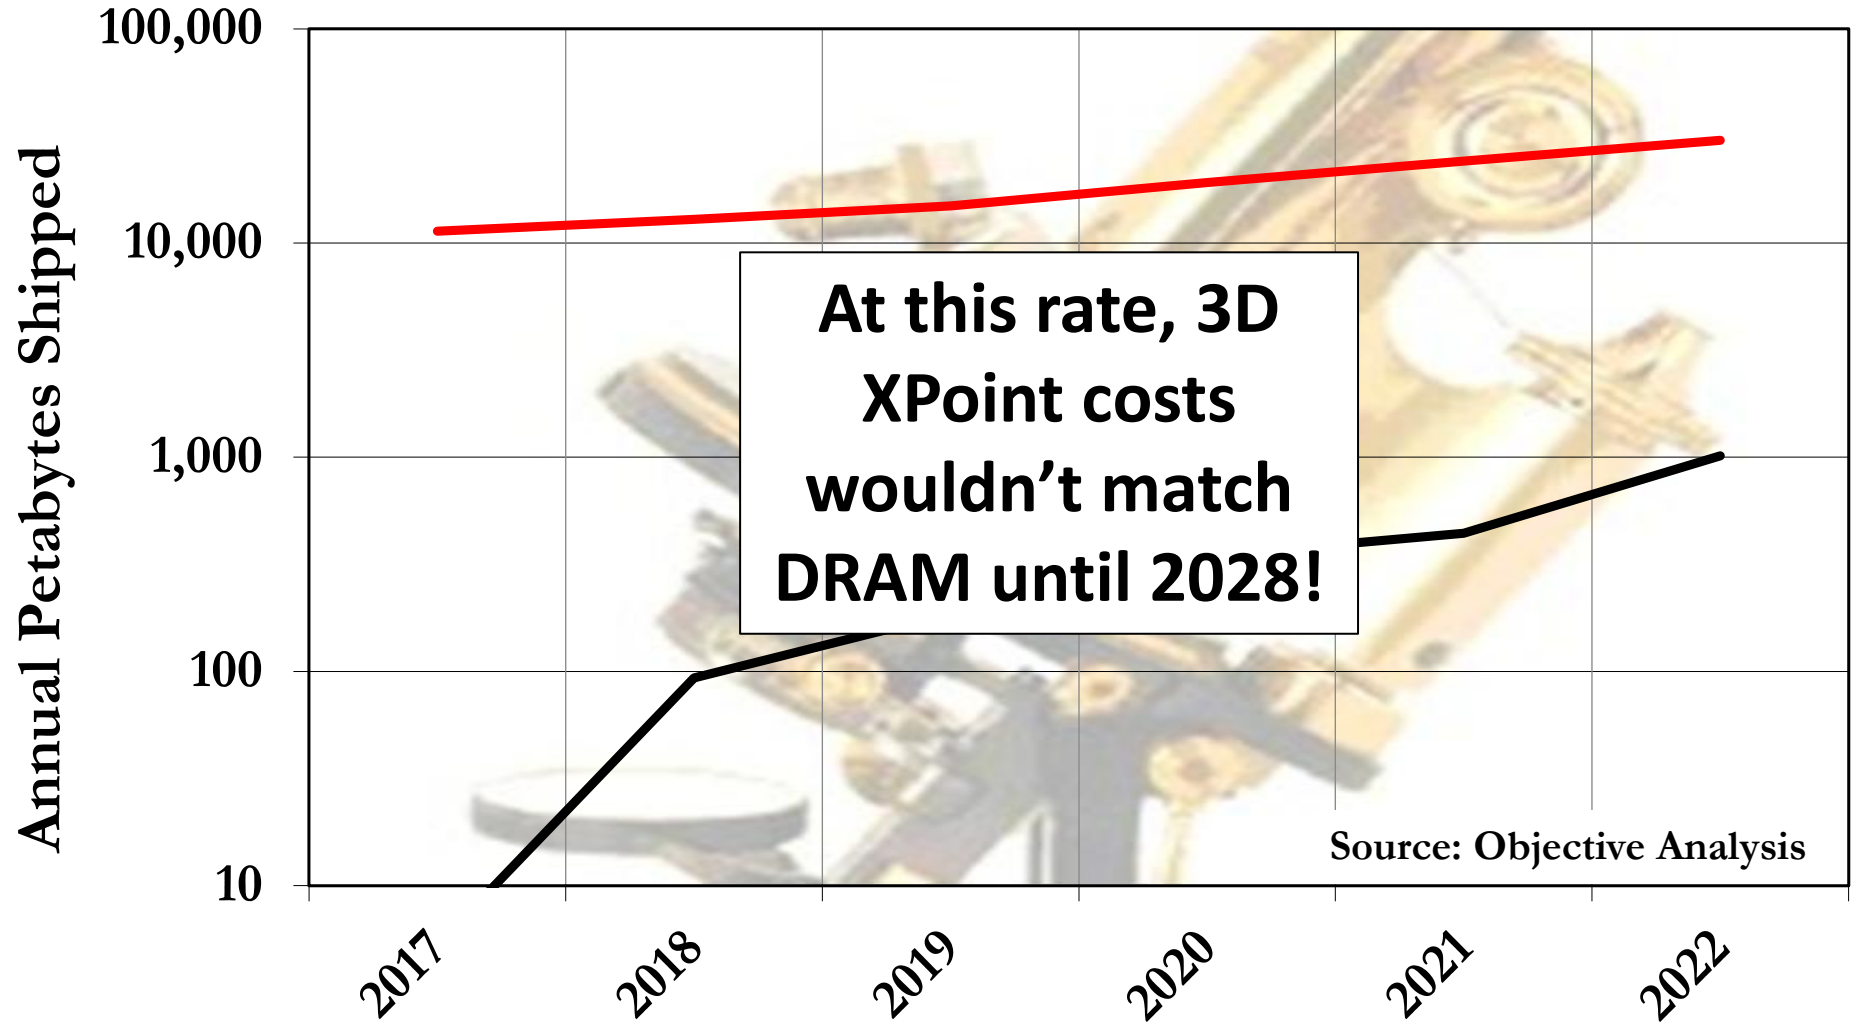

An Optane Timeline

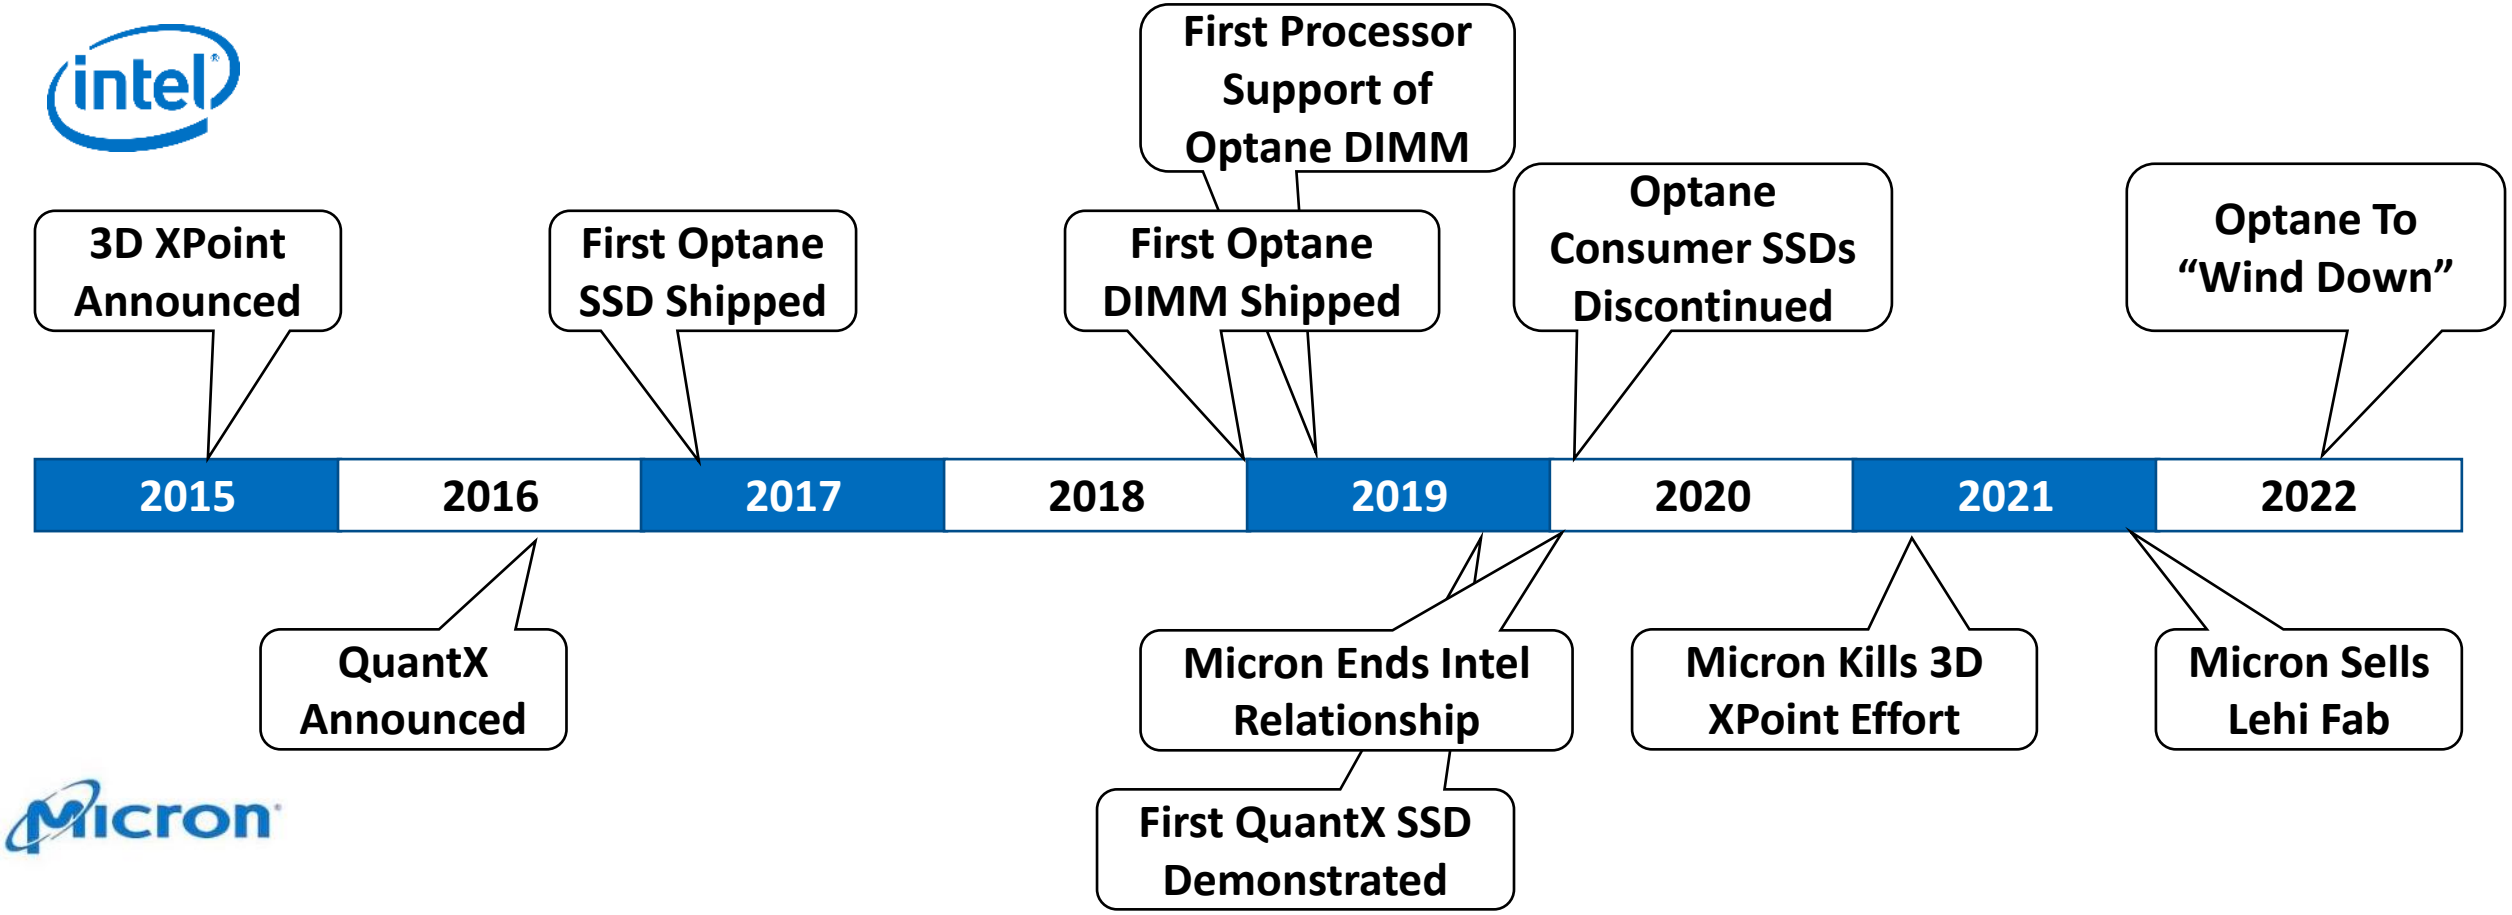

Intel's Optane Losses

- >$7B in Intel losses
  - Micron losses ~$400M/quarter
- Optane SSDs intended to drive early volume
  - Users disliked the price
  - Volume never really ramped
- Optane DIMM was delayed

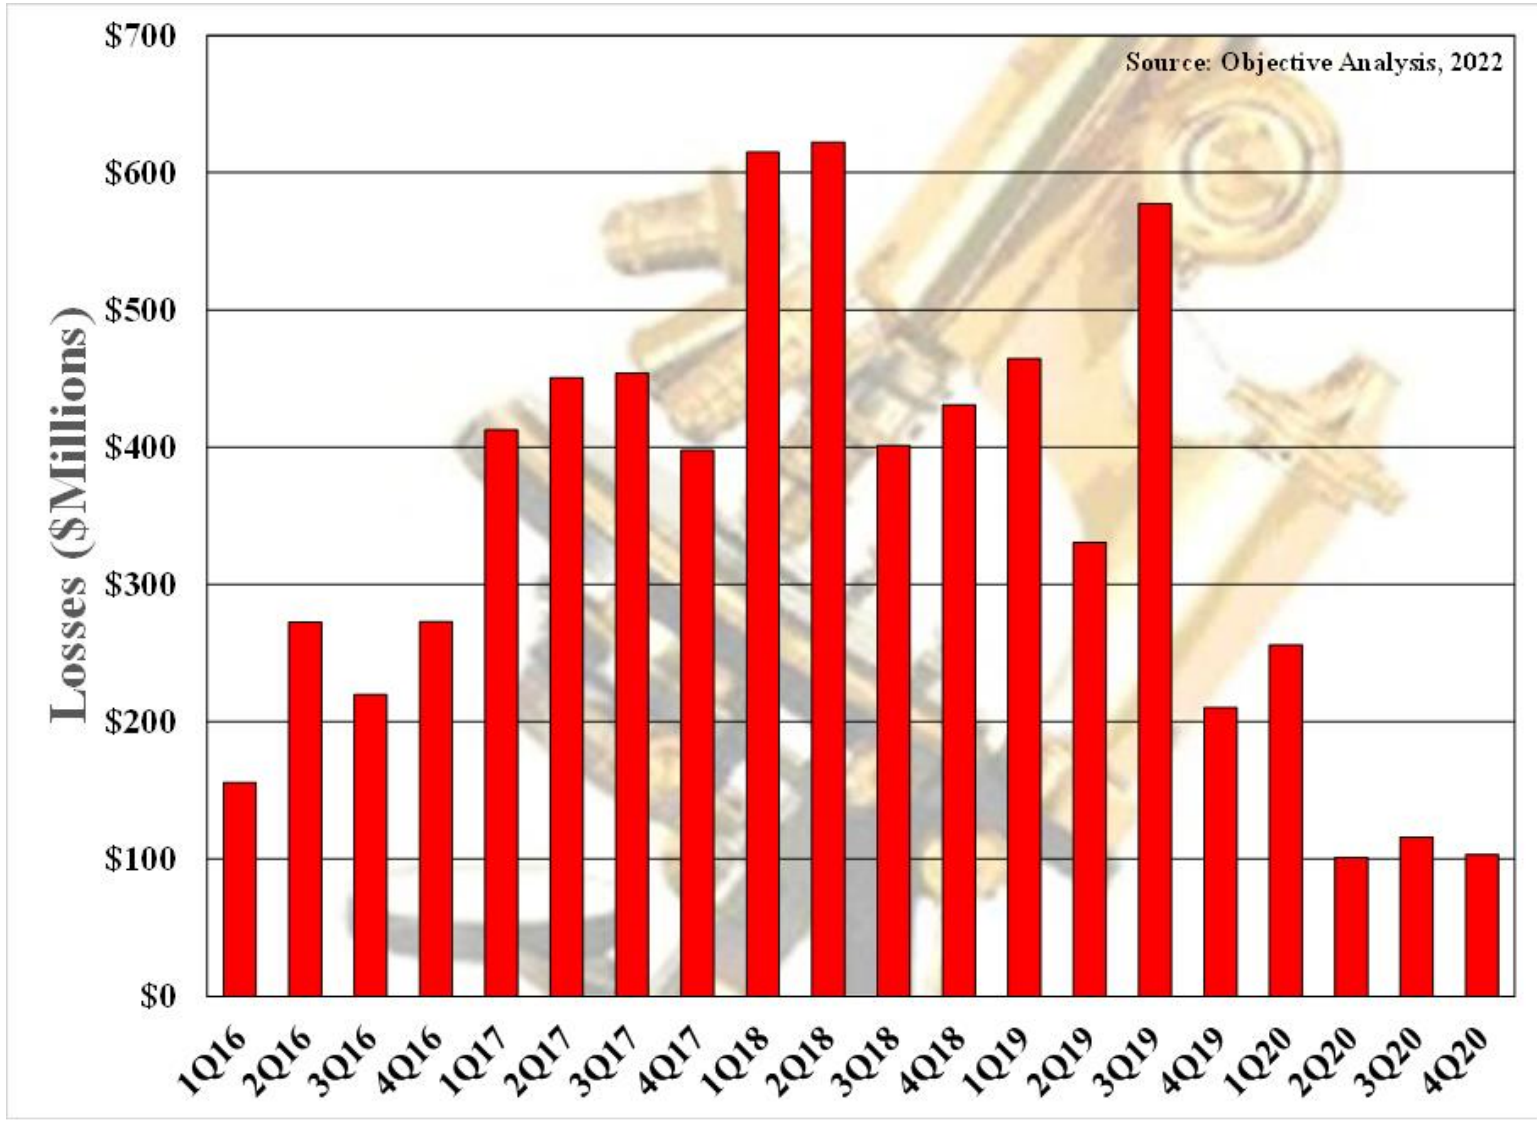

Lessons to Be Learned

- Expect losses until volume ramps
  - A small die size doesn't matter if volume is too low
- Losses might be larger than expected to reach a low enough price point
- Supporting elements may delay adoption
  - Optane's processor support was delayed
  - Application software wasn't widely available

Optane’s Bountiful Legacy

Intel Gave the Industry a Great Gift

What Did We Get from Optane?

- New programming paradigm
- New processor instructions
- New approach to 2-speed memory
- New Near-Memory bus concept
- New approach to memory expansion
- New thinking about security concerns with persistent memory

New Programming Paradigm

The SNIA NVM Programming Model, SNIA.org/PM

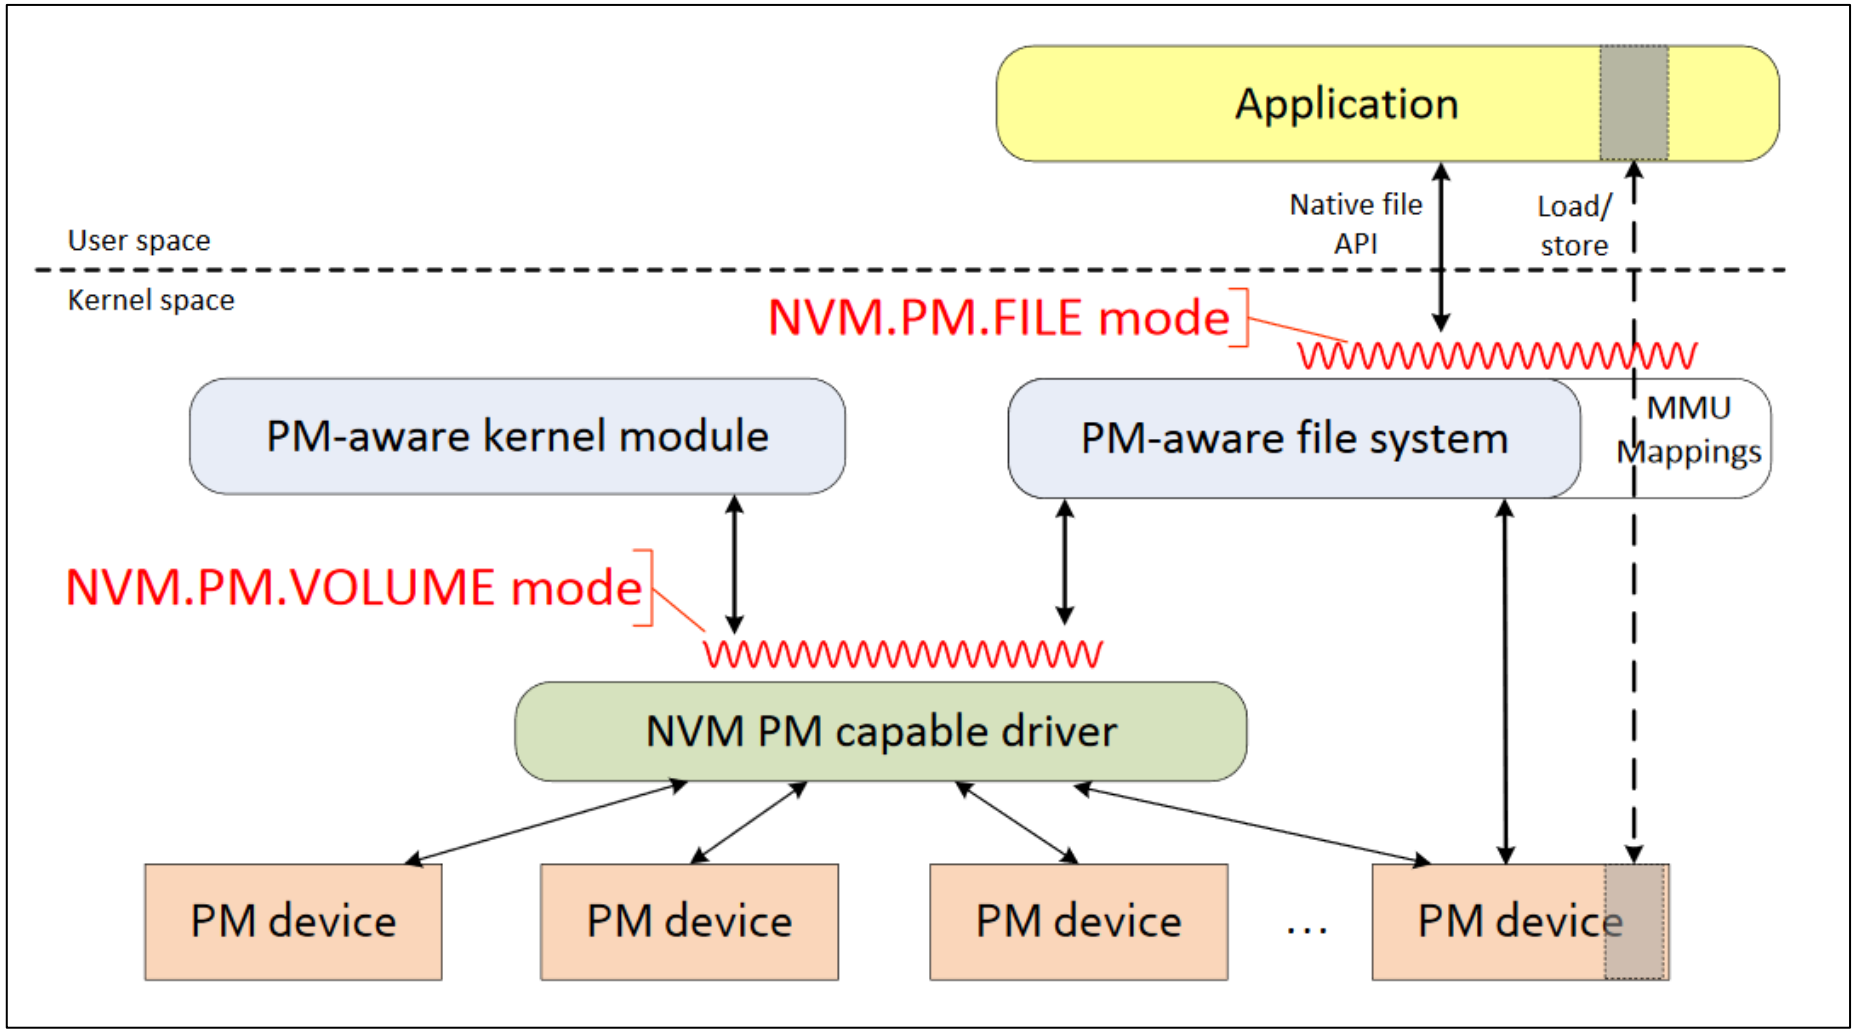

New Processor Instructions

- Brings processor cache into the power fail domain (i.e. persistence)
  - Everything from top to bottom is persisted (except for the registers)
  - Both move data into the power-fail protected ADR domain
- CLFlushOpt
  - Flush Cache Line Optimized – Invalidate a specific cache line, and if that line is dirty flush it back to memory
- CLWB
  - Cache Line Write Back – If a specific cache line is dirty write it back to memory

New Approach to 2-Speed Memory

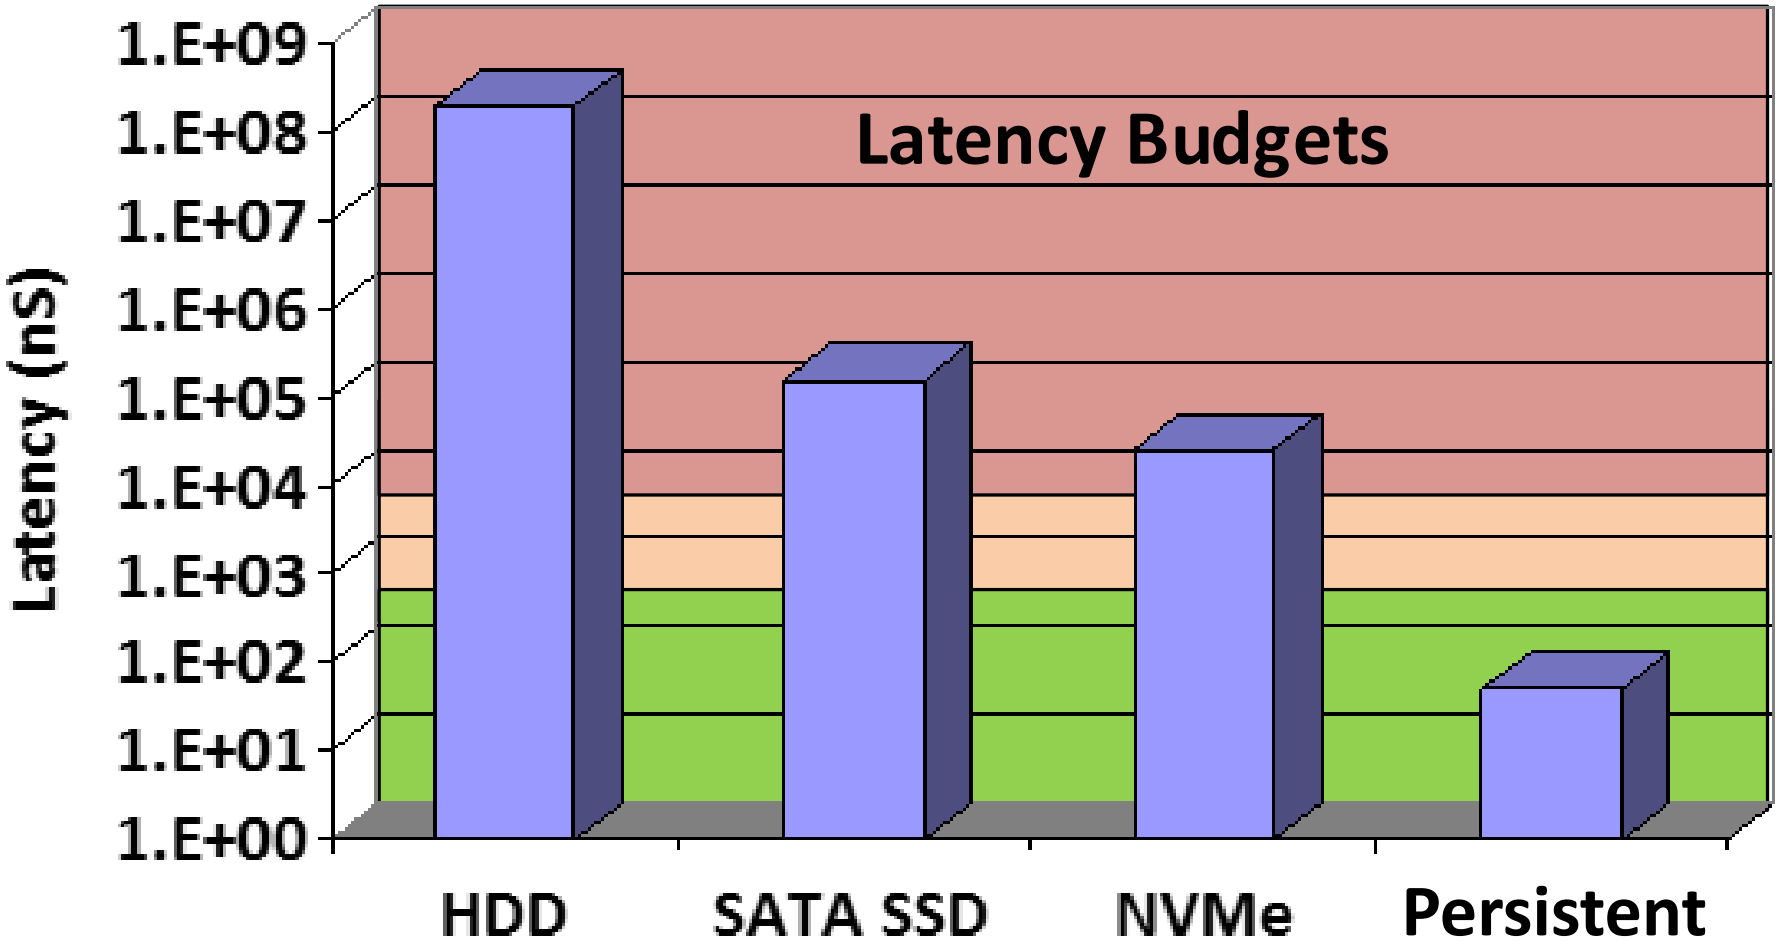

Context Switch

Polling used when a context switch would be too slow

New Near-Memory Bus Concept

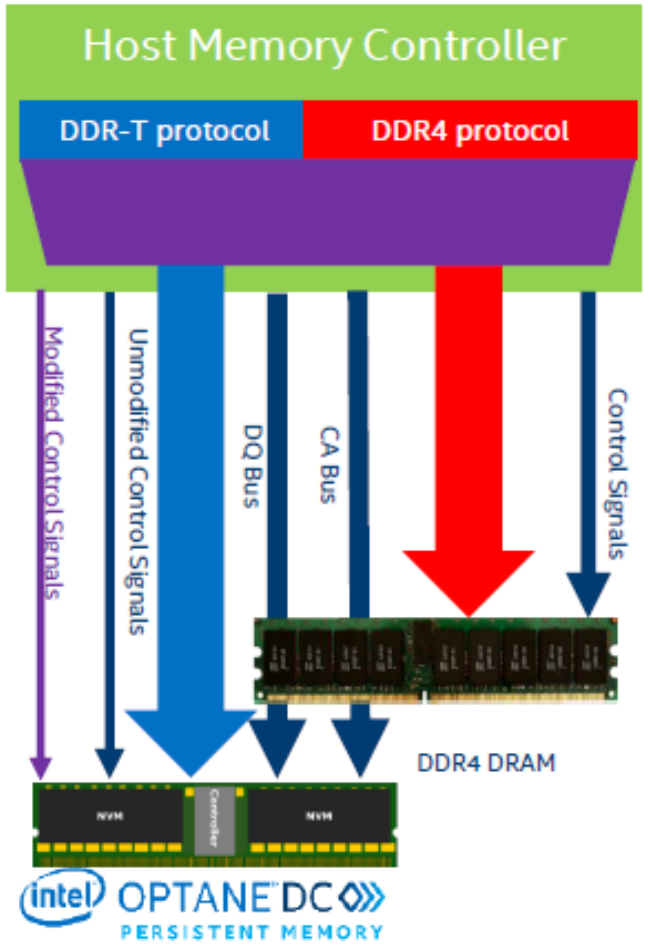

- Handles both fast & slow memory
  - Transactional protocol for Optane
- Operates on standard DDR4 interface
  - "Modified Control Signals" added to unassigned pins
  - All timing, signaling, protocol otherwise unmodified
- DRAM and Optane share the same sockets

Optane and DRAM modules look nearly identical to the end user

New Approach to Memory Expansion

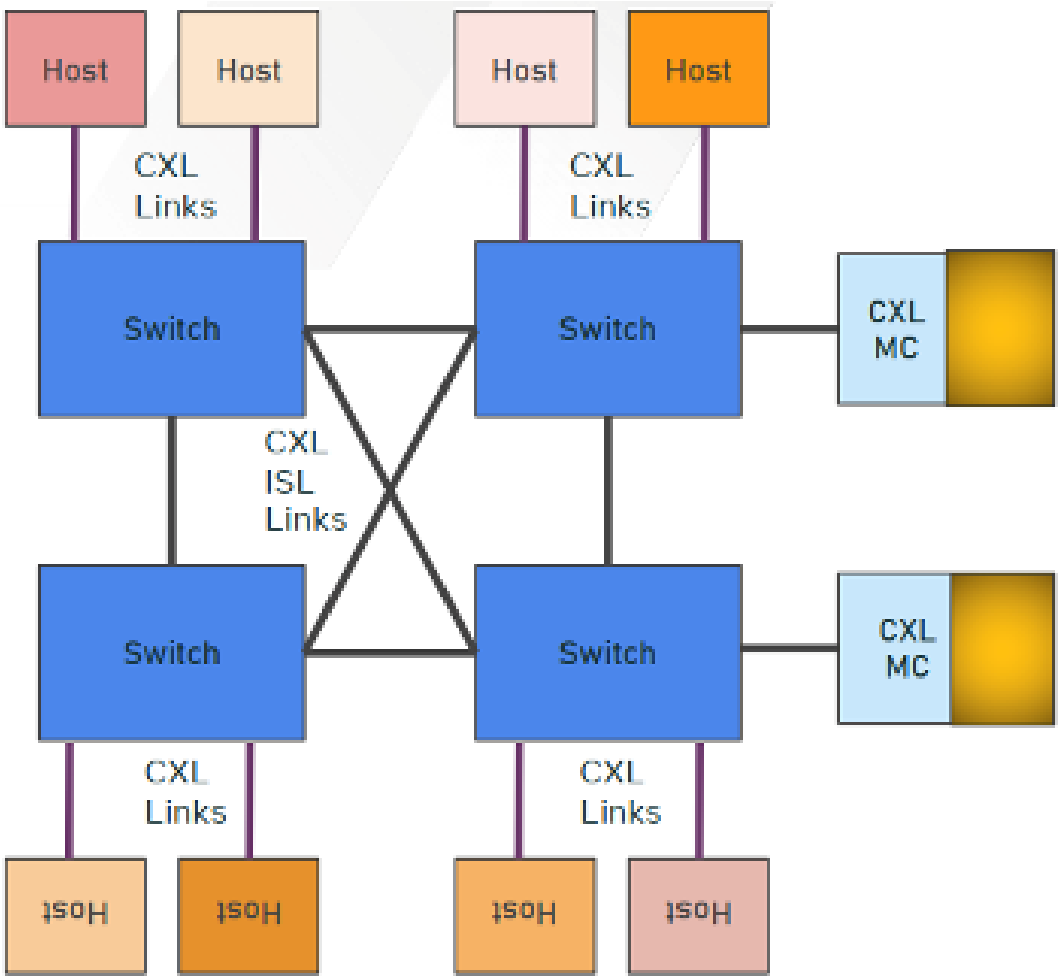

- Near Memory at CPU
- Far Memory on CXL
- CXL to support multiple Far Memory configurations
  - Large Memory
  - Memory Pools
  - Memory Sharing
    - Used for trading messages
  - Memory Fabrics

New Thinking about Security Concerns with Persistent Memory

- Security for persistence has been addressed in SSDs & HDDs
  - Physical destruction
  - Secure erase software
  - AES encryption
- How to bring this to persistent memory?
  - AES encryption
    - Memory Mode: Key is lost on power fail
    - App Direct Mode: Key stored on module. Module locked on power fail. Passcode re-enables
  - Secure cryptographic erase and over-write when decommissioned

Future SoCs Will Include Persistence

Industry Dynamics are Forcing this Change

Embedded NOR & SRAM Scaling Has Stopped

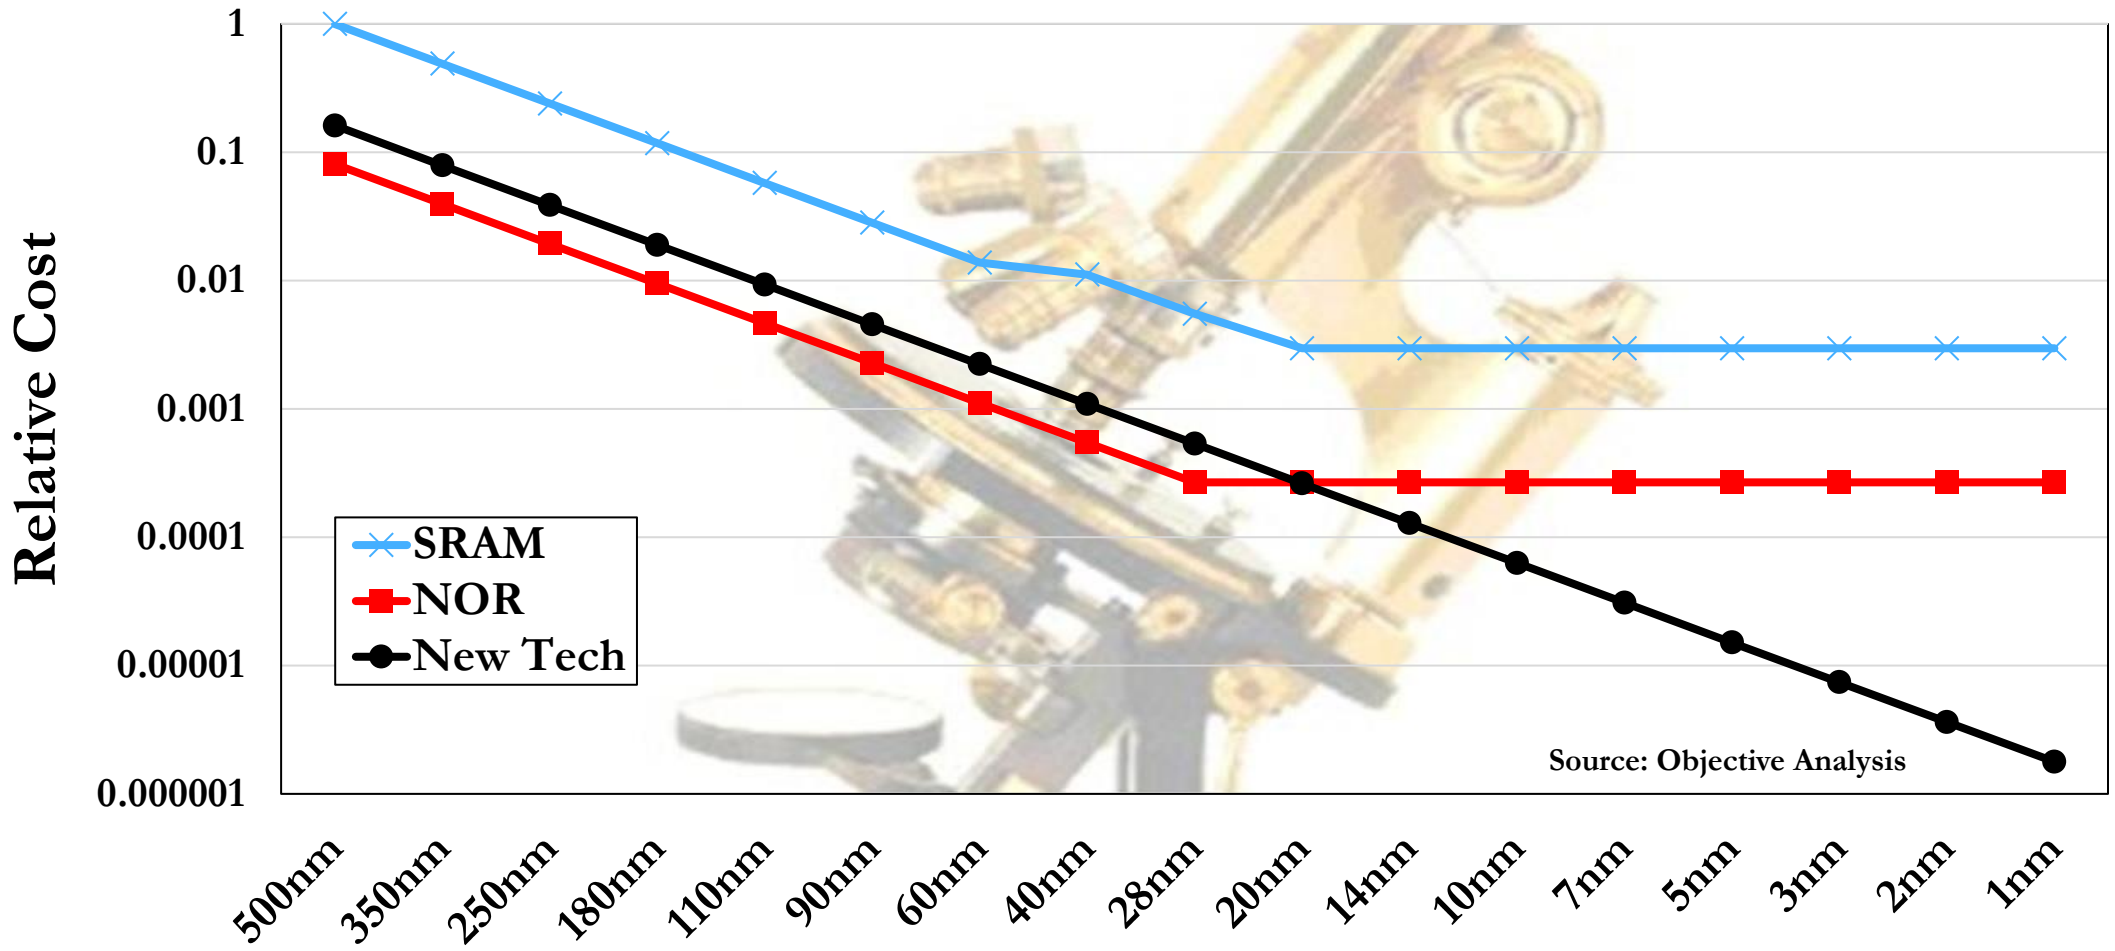

Impact on Processors

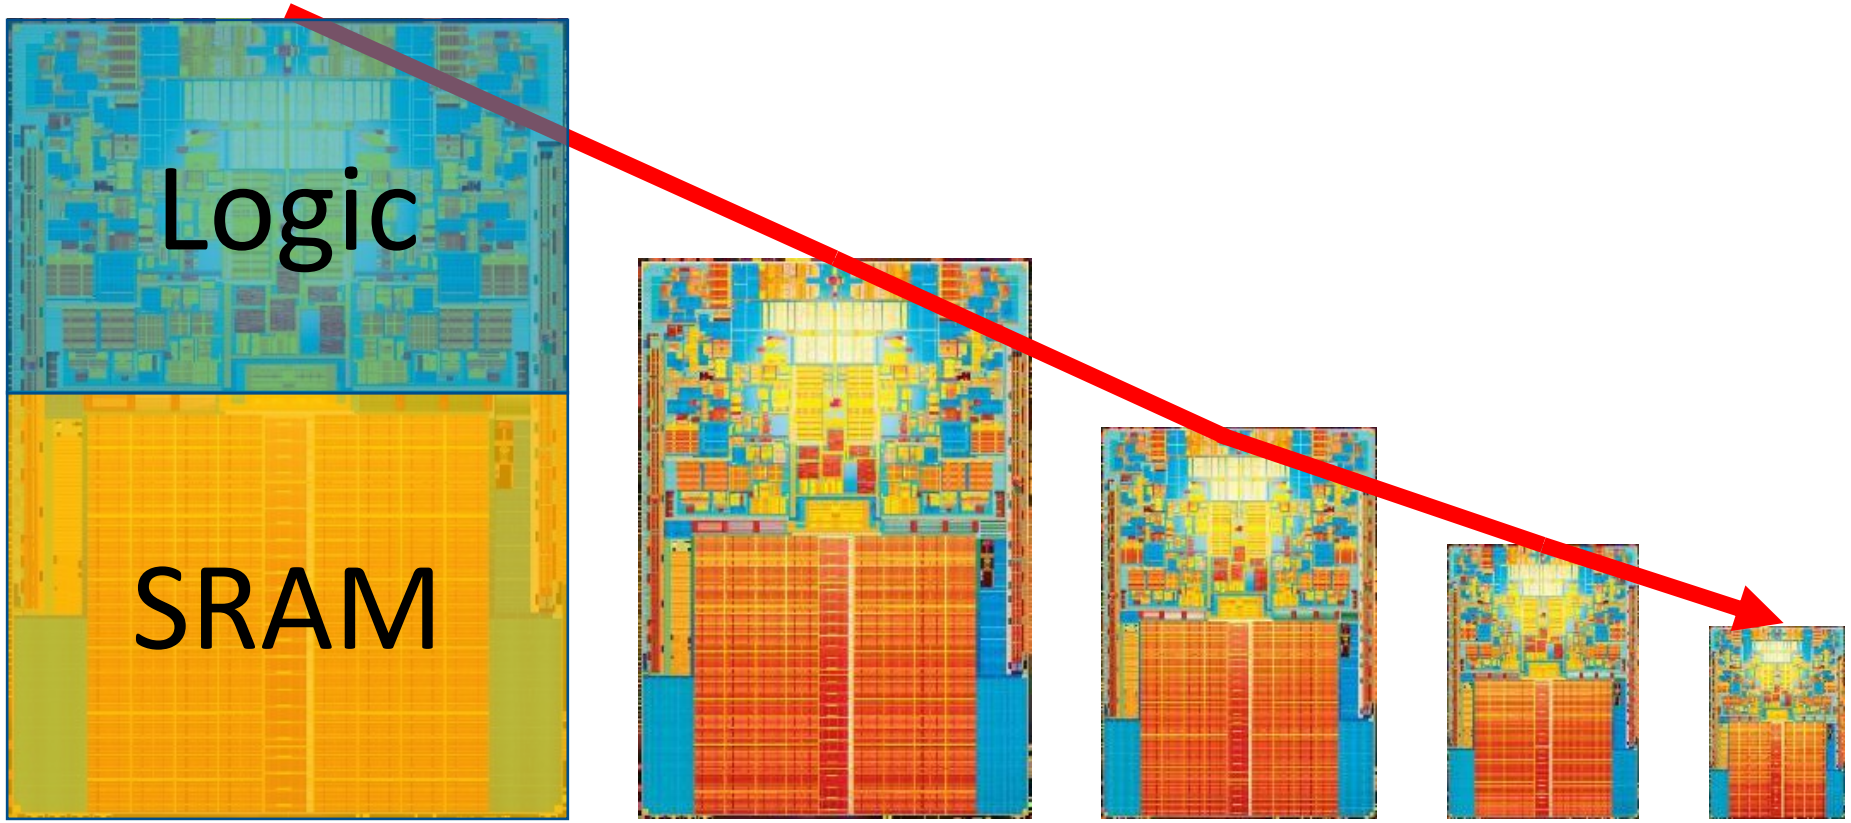

Impact on Processors

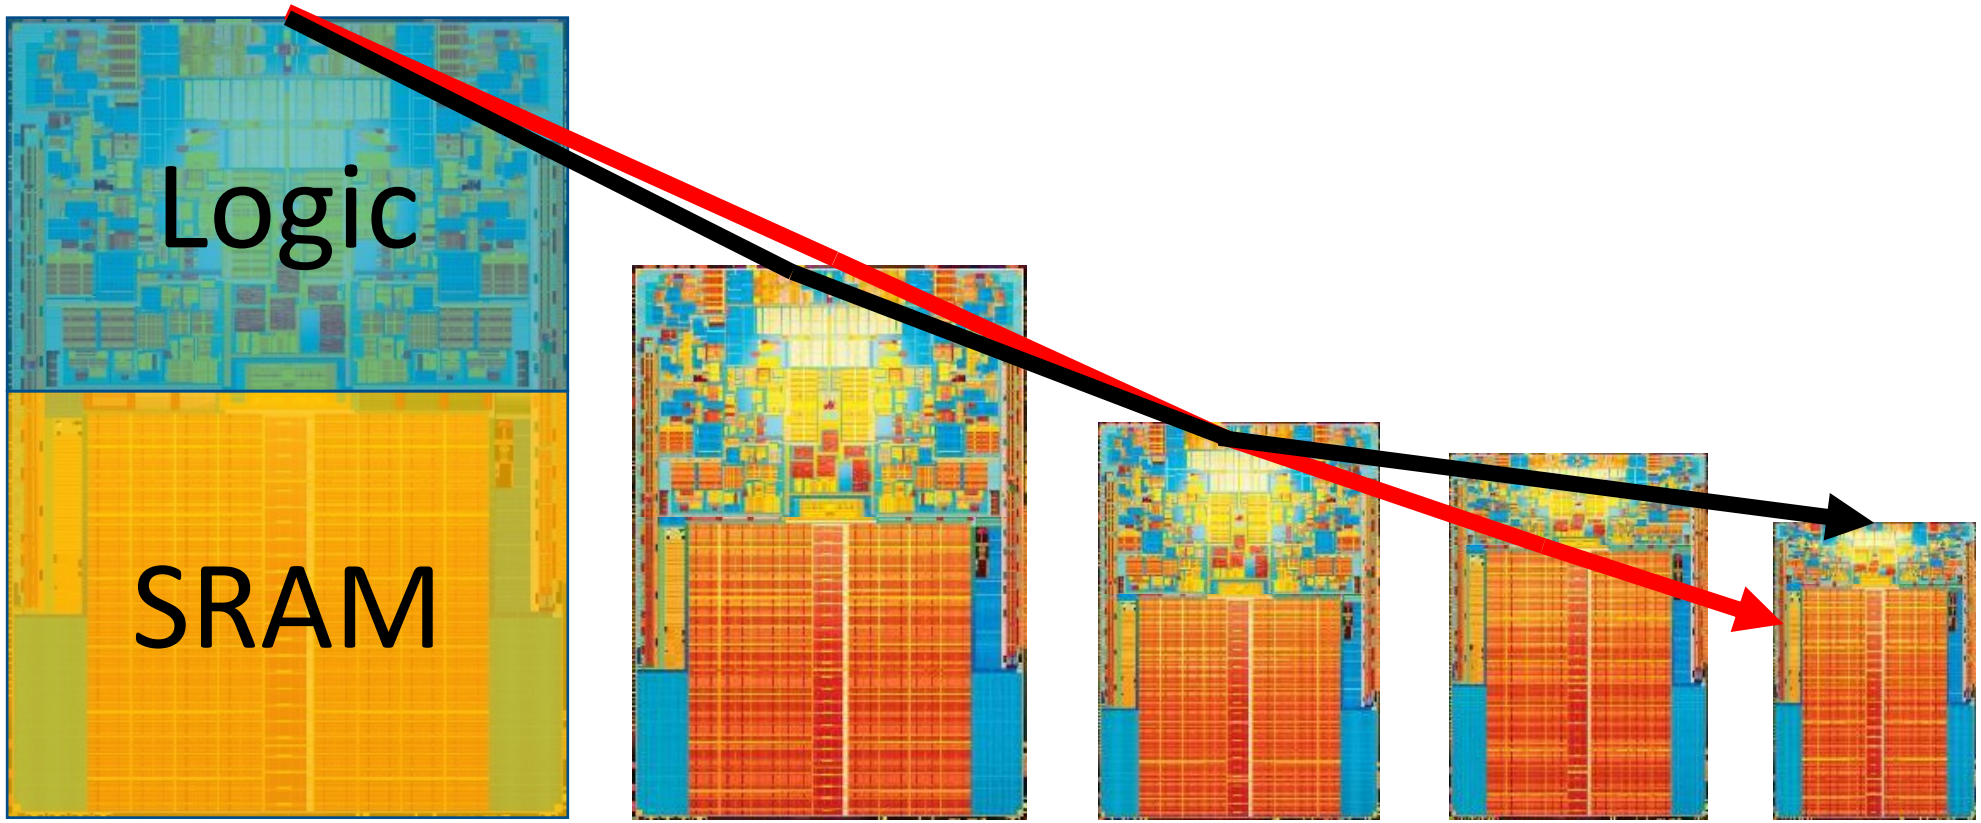

NOR & SRAM Scaling will Stop

- An emerging memory will take over as the embedded memory on SoCs
- That emerging memory will be persistent:
  - MRAM
  - ReRAM
  - FRAM
  - PCM

Optane's legacy will make this persistence useful

Meet the Memories

Candidates for Persistent Memory

MRAM

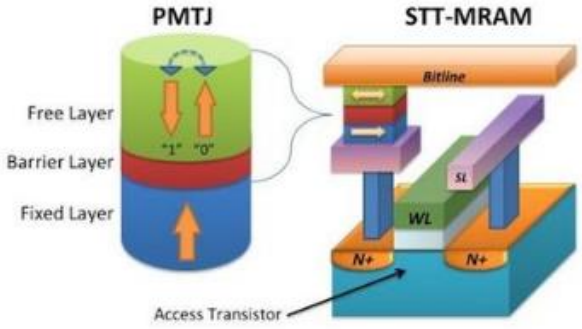

PCM

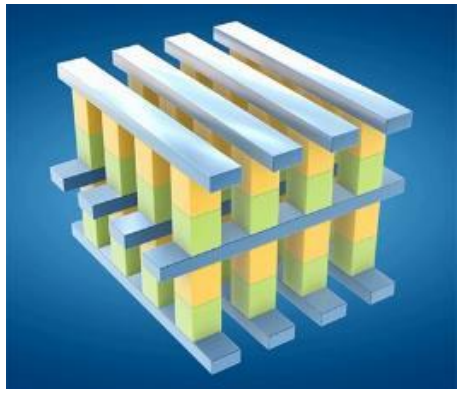

ReRAM

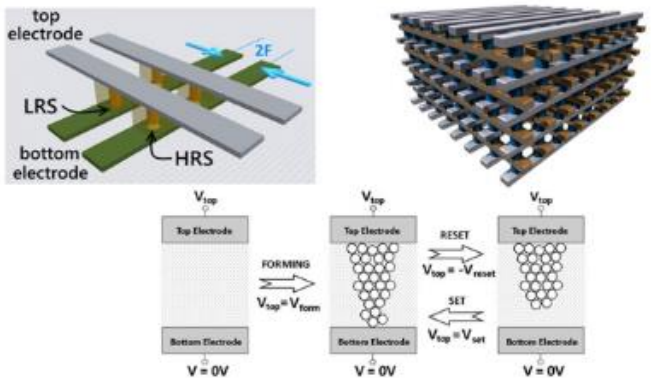

FRAM

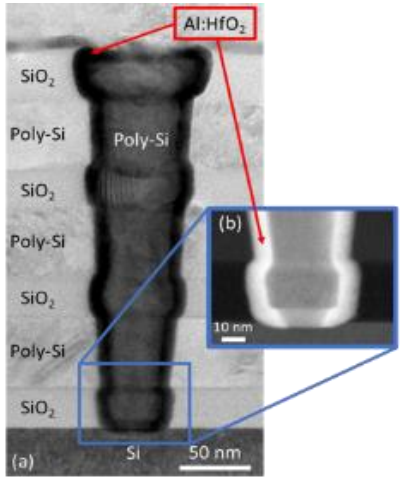

Comparing Memory TEchnologies

	Established Memory Types				Emerging Memory Types				
	SRAM	DRAM	NOR Flash	NAND Flash	FRAM	ReRAM	Toggle MRAM	STT	PCM
Nonvolatile?	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cell Size	40-500f^2	6-10f^2	10f^2	0.03-5f^2	15-32f^2	4-8f^2	16-32f^2	40-160f^2	2-8f^2
Read Time	1-100ns	30ns	10ns	50ns	20-50ns	10-20ns	3-20ns	3-15ns	5-20ns
Write Time	1-100ns	50ns	10^{5-7}ns	10^{6-8}ns	50ns	10^{2-5}ns	10-20ns	3-15ns	>30ns
Endurance	∞	∞	10^5	10^3	10^{12}	10^{5-6}	10^{15}	10^{15}	10^{12}
Write Energy	Low	Low	High	Med	Low	Low	Somewhat High	Low	Low
Write Voltage	None	2	6-8	12	2-3	1.2	3	1.5	1.5-3

MRAM

- Everspin
  - Partnership with Global Foundries
  - Used in IBM's FlashCore modules in FlashSystem arrays
- Renesas (Formerly IDT)
  - 8Mb, SPI
- Avalanche and Honeywell are shipping some MRAM for mil/aero applications
- Leading foundries starting to ship embedded MRAM
  - GlobalFoundries
  - TSMC
  - Samsung
  - Others

Source: Everspin

Source IBM

Source: TechInsights

PCM (3D XPoint)

- The second-oldest emerging memory (1970)
- Intel Optane products
  - NVMe shipped in 2017
  - DIMMs in 2018
  - In 2021 Intel stopped consumer Optane SSDs
  - In late July 2022 Intel announced “winding down” Optane
- Micron abandoning 3D XPoint
  - Sold Lehi fab

ReRAM

- Adesto has shipped CBRAM chips for several years
  - Dialog Semiconductor acquired Adesto June 2020
  - Will license CBRAM technology to GLOBALFOUNDRIES
    - GF will first offer as an embedded, option on its 22FDX platform
    - GF Plans to extend to other platforms.
- Cerfe Labs (Arm spin-out)
  - Correlated electron materials (CeRAM)
  - Licensed from Symetrix.
- Others (Weebit, Mitsubishi, Fujitsu, Panasonic, Winbond, Honeywell,…)
- Foundry support (GLOBALFOUNDRIES, TSMC, others)

FRAM

- The oldest emerging memory (1955)
- The highest-shipped emerging memory
- Finding new life with new materials

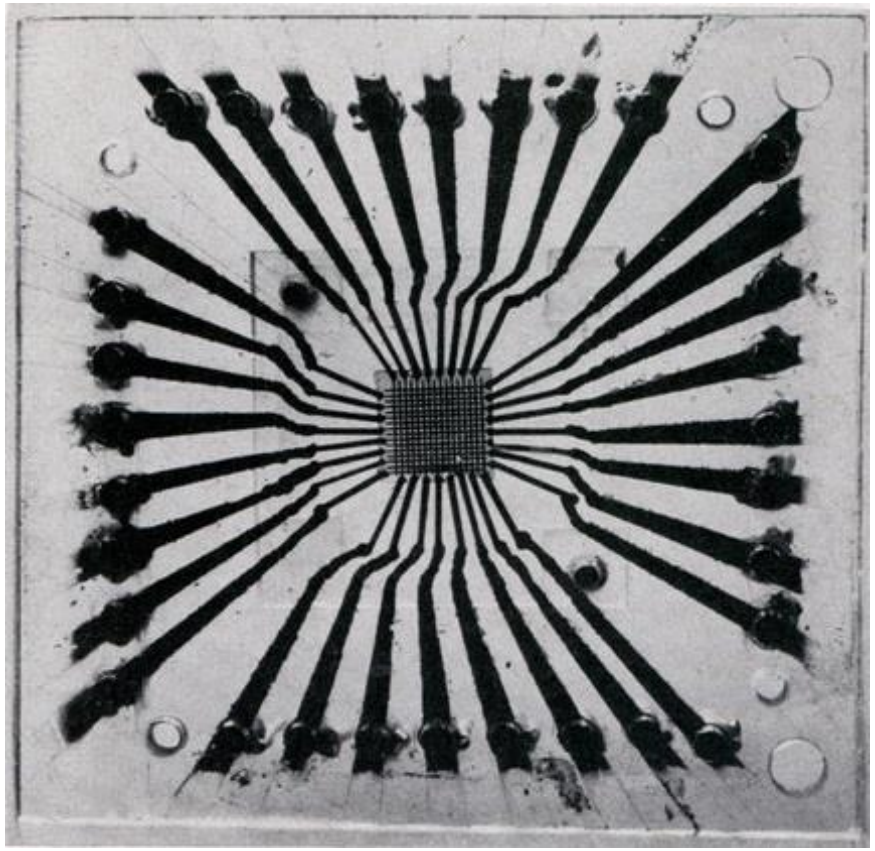

How Optane's Legacy Helps

- Caches & even registers will migrate to an emerging memory technology
  - Persistence is moving closer to the CPU
  - Eventually persistence will be within the CPU
  - SNIA"s NVM Programming Model will be sidely used at this point
- Mixed memory speeds to support both SRAM & emerging memory

Cache	Today	5 Years	10 Years
Level 1	Fast SRAM	Fast SRAM	Fast MRAM
Level 2	Slower SRAM	MRAM	Slow MRAM
Level 3	SRAM or DRAM	ReRAM	ReRAM

New Report: Emerging Memories Enter the Next Phase

OBJECTIVE ANALYSIS
Semiconductor Market Research

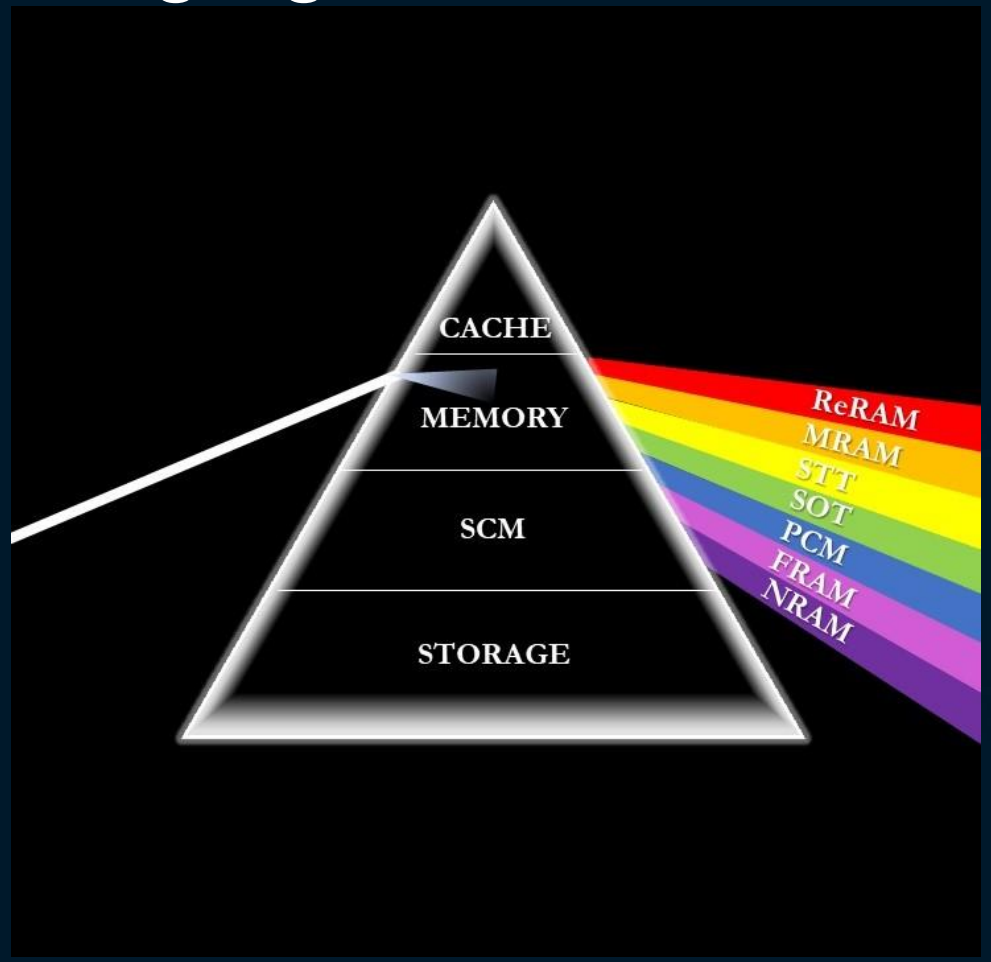

Coughlin Associates

Now Available!

http://www.tomcoughlin.com/techpapers.htm
https://Objective-Analysis.com/reports/#Emerging

How Future Processors Will Benefit

Using the Chiplet Approach

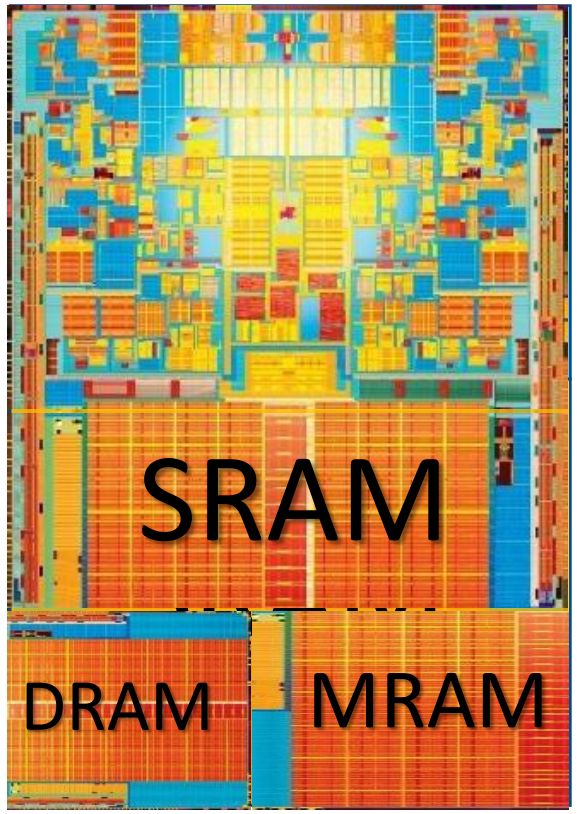

- Logic process for CPU & L1 cache
- SRAM process for L2 cache
- DRAM process for L3 cache
- MRAM process for persistent cache

Optimized Processes

Larger "Effective" Die Sizes

Persistence on the CPU

What About the Economies of Scale?

- Chiplets are not commodity memories
  - Embedded SRAM cost is much greater than cost of discrete SRAM
  - Example: Current cost of SRAM portion of CPU is ½ the cost of the CPU chip
    - If the chip costs $200 to produce, and ½ the chip is SRAM, then the SRAM costs $100
    - Server cache is 64KB L1, 1MB L2, and 1.5MB L3 = 2.6MB
      - SRAM cost would be $100/2.6MB = $38/MB
      - A 4 megabit SRAM chip (½MB) retails for $6 = $12/MB
        - (DRAM is currently $3/GB, or $0.003/MB – Three hundredths of a cent!)
  - A memory chiplet will reduce the cache's cost
    - But it can still be orders of magnitude more costly than DRAM and NAND and eventually even more expensive than a non-volatile memory, like MRAM

The Point? Persistent Caches are Coming

- Persistence will require new software support
  - That support is already in the SNIA NVM Programming Model
- Many other factors will be impacted
  - Instruction set, 2-Speed Memory, Security
  - These have also already been resolved
- The ecosystem is ready for persistent caches when they arrive
- Chiplets will make that sooner, rather than later

Summary

Summary

- Optane couldn’t harness the economies of scale
- Optane effort generated a great legacy
- CXL opens new vistas in data center architecture
- Emerging memories are here, and they’re persistent
- Future processors will have persistent cache, and later, registers
- Optane’s legacy will benefit tomorrow’s processors
- Chiplets will accelerate this transition

Questions?

