Persistent Memory is Not What You Think!

Jim Handy

OBJECTIVE ANALYSIS – Semiconductor Market Research

OBJECTIVE ANALYSIS

Profound Analysts

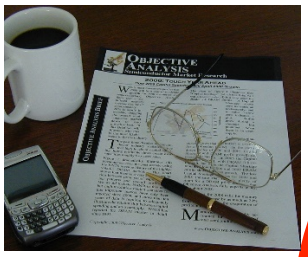

Reports & Services

Custom Consulting

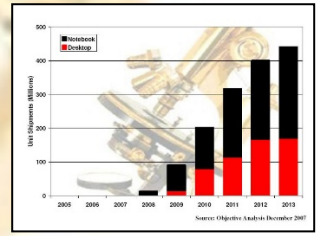

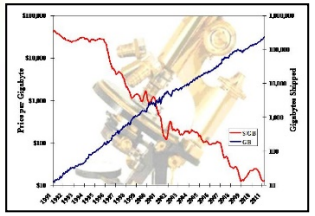

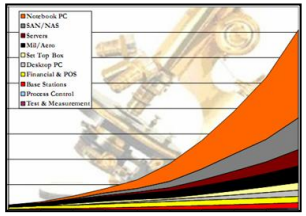

Objective Analysis

Semiconductor Forecast Accuracy

Year	Forecast	Actual
2008	Zero growth at best.	-3%
2009	Growth in the mid teens	-9%
2010	Should approach 30%	32%
2011	Muted revenue growth: 5%	0%
2012	Revenues drop as much as -5%	-2.7%
2013	Revenues increase nearly 10%	4.9%
2014	Revenues up 20%+	9.9%
2015	Revenues up ~10%	-0.2%
2016	Revenues up ~10%	1.1%
2017	Revenues up ~20%	TBD

Fallacies

- "Instant-On"
- Needs No Software Support
- Needs No Firmware Support
- CPUs Don't Need to Change
- PM Will Replace HDDs
- PM Will Replace Flash
- PM Will Replace DRAM
- All PM Is The Same

Making Sense of Persistent Memory

- Why is it needed?
  - Where does it fit?
- How will it impact other technologies?
  - Hardware
    - HDD, SSD, & DRAM
  - Software/Firmware
  - Processors
- What new chip technologies are coming?

Making Sense of Persistent Memory

- Why is it needed?
  - Where does it fit?
- How will it impact other technologies?
  - Hardware
    - HDD, SSD, & DRAM
  - Software/Firmware
  - Processors
- What new chip technologies are coming?

Memory/Storage Speed Gap

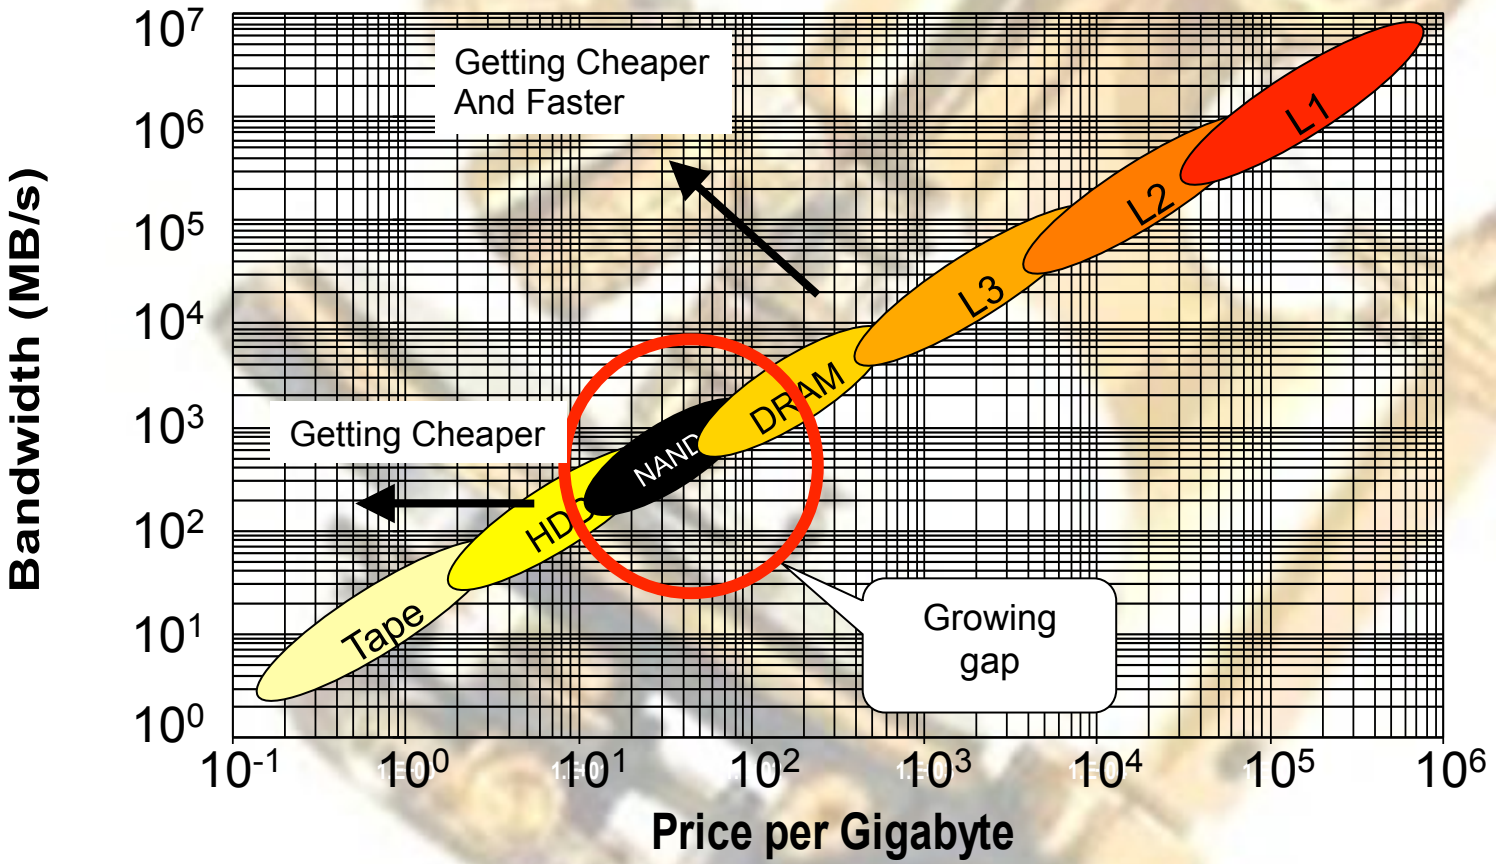

Cost/Performance MUST Fit The Hierarchy!

From: Solid State Drives in the Enterprise

Memory/Storage Speed Gap

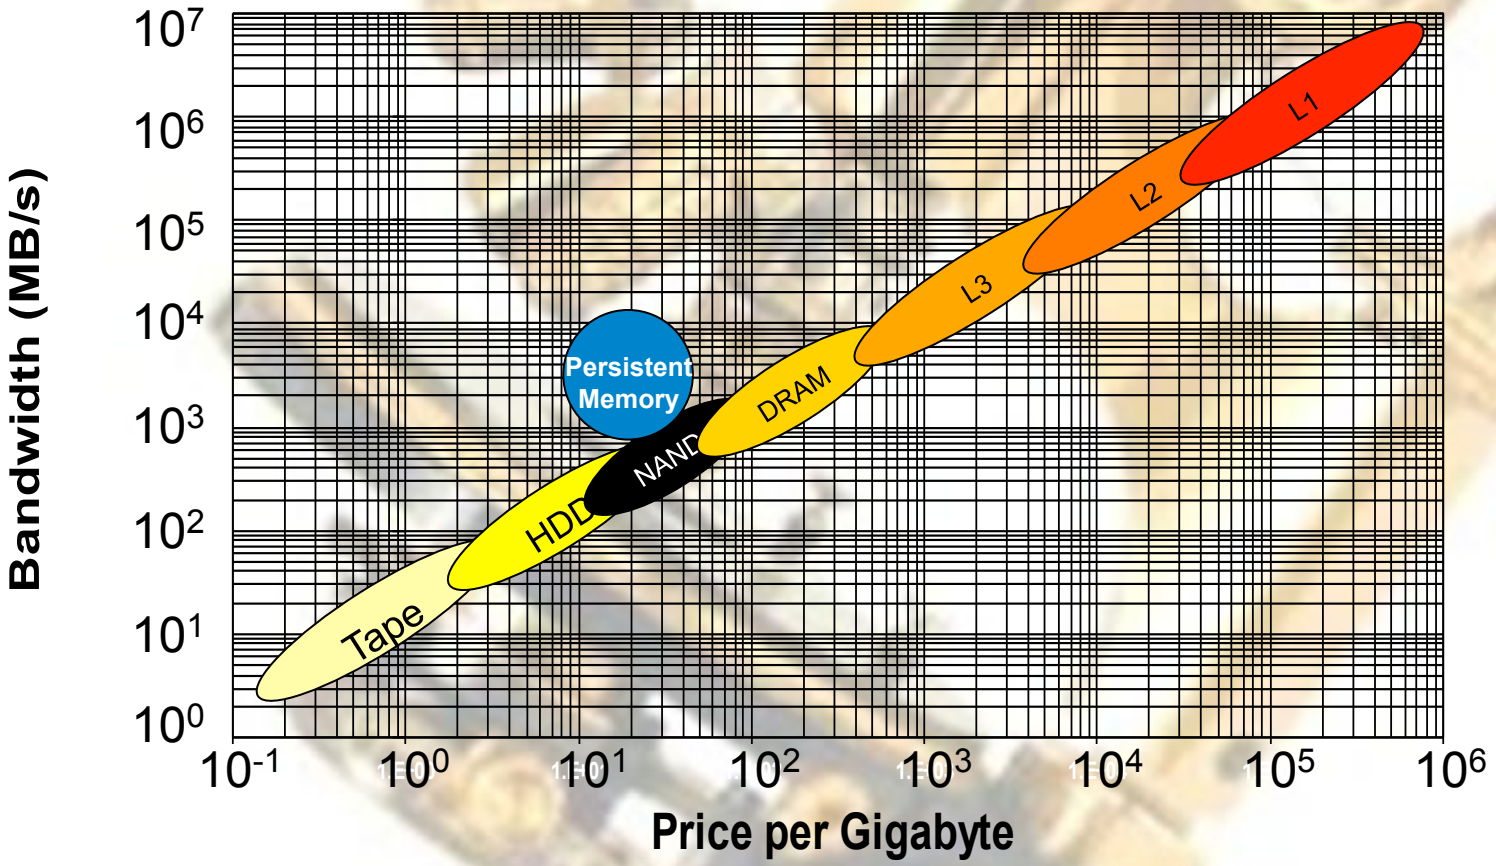

From: Solid State Drives in the Enterprise

Persistent Memory Implications

- Retains data during a power loss
  - Faster recovery of state before power down
  - Can be used to securely store data
    - Removes need for slow storage from memory

Persistent Memory Implications

- Retains data during a power loss
  - Faster recovery of state before power down
  - Can be used to securely store data
    - Removes need for slow storage from memory
- Lower latencies than SSDs
  - Accelerates system performance
- Lower power & cost than DRAM
  - Improves cost/performance ratio

Is PM Memory or Storage?

- Can function as either or both!
  - It's a huge memory but is not persistent
  - It's storage but is expensive and really fast!
- Both improve system performance
  - "Large Memory" model works today
  - Software support required if its storage

Making Sense of Persistent Memory

- Why is it needed?
  - Where does it fit?
- How will it impact other technologies?
  - Hardware
    - HDD, SSD, & DRAM
  - Software/Firmware
  - Processors
- What new chip technologies are coming?

Will PM Kill Off HDD/SSD?

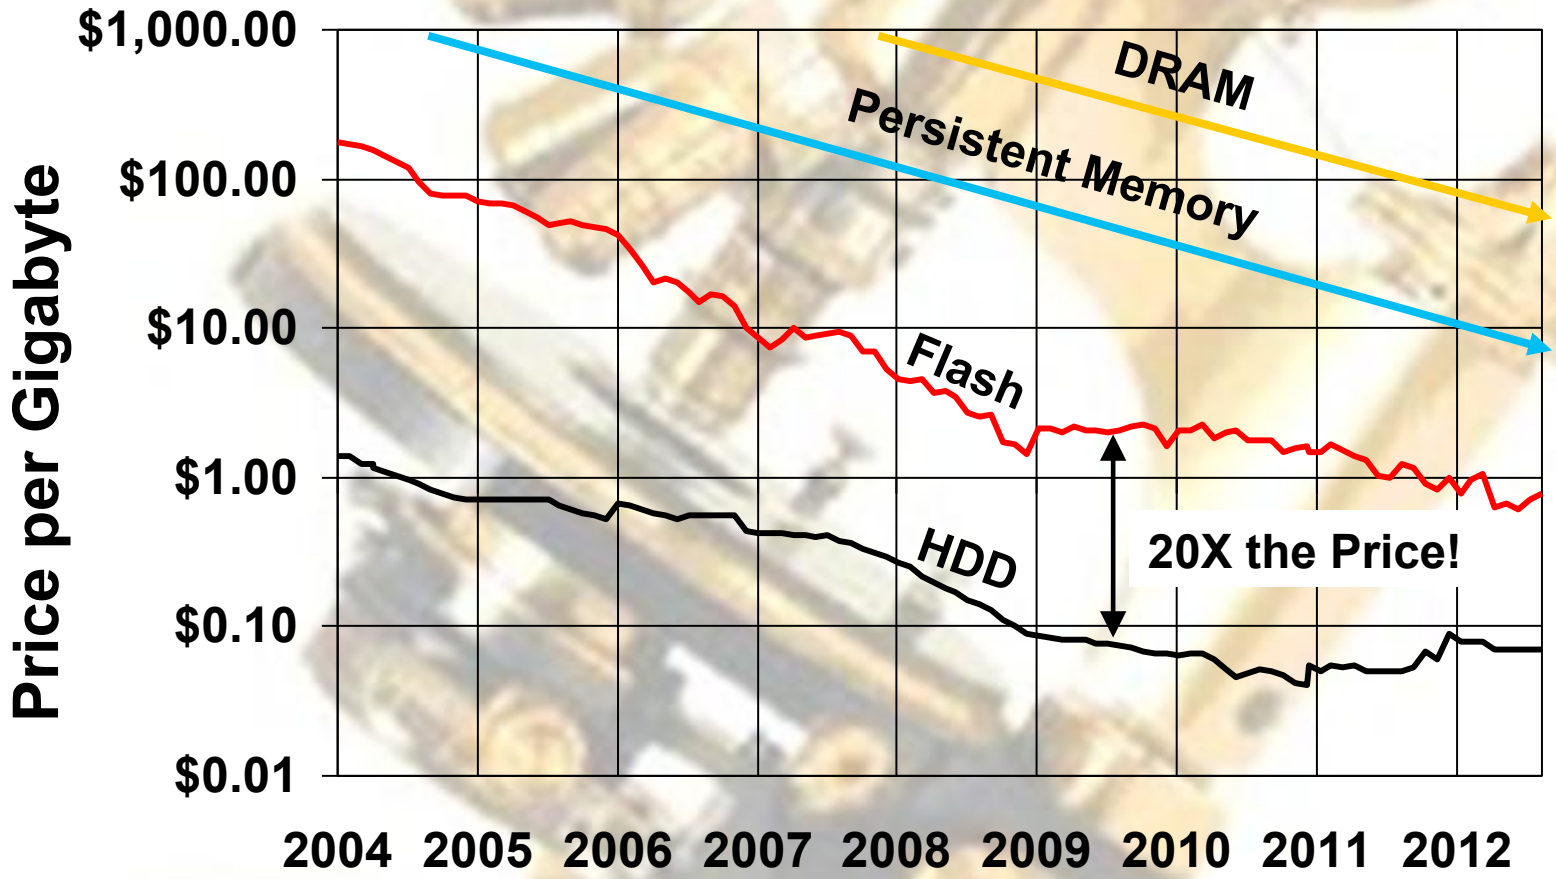

Memory/Storage Hierarchy

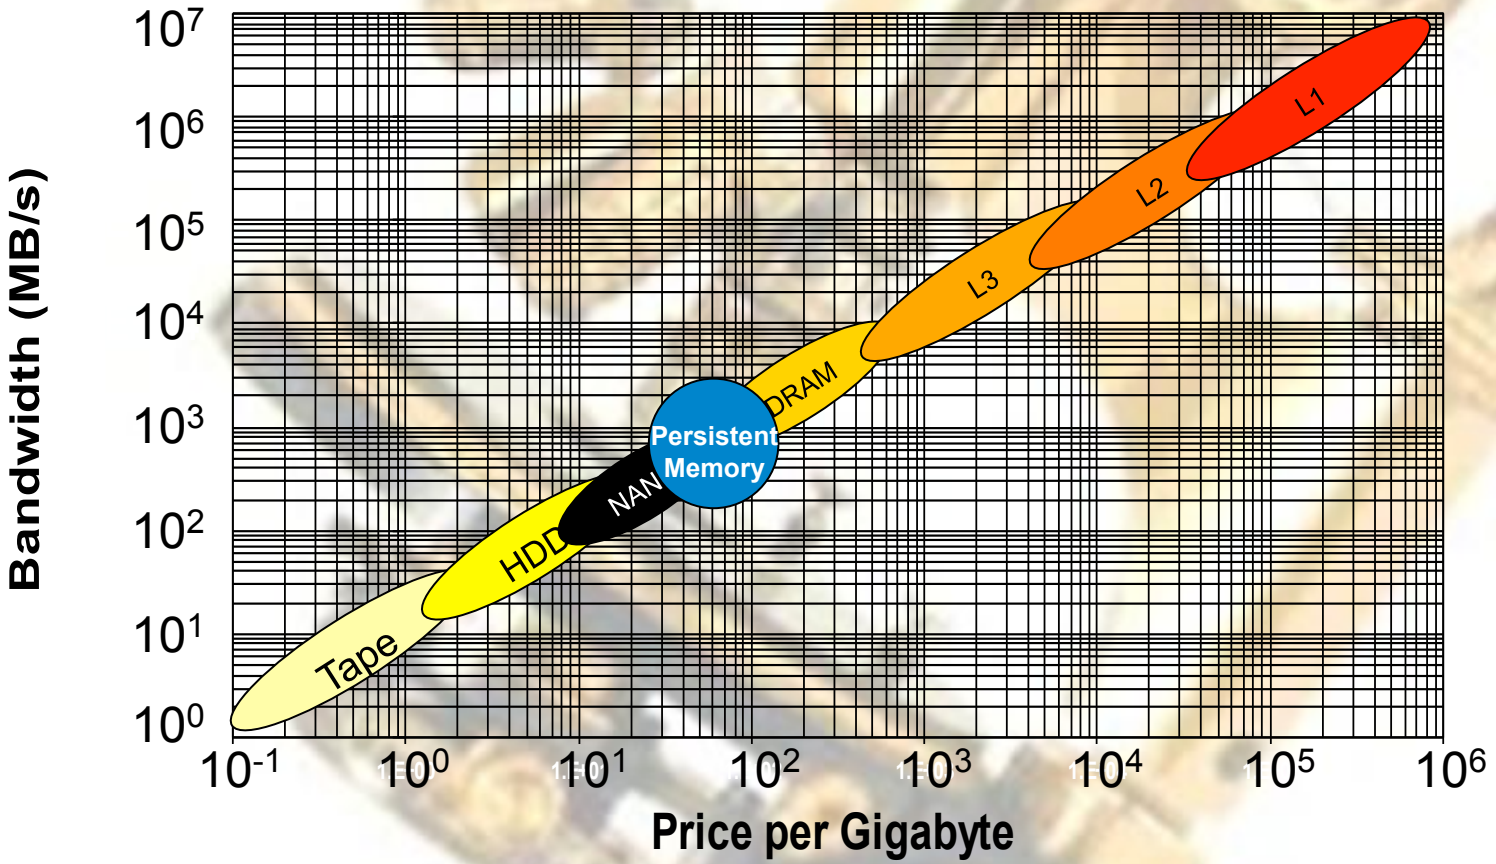

Cost/Performance MUST Fit The Hierarchy!

From: Solid State Drives in the Enterprise

Software Support

- Persistent Memory requires two levels of software support:
  - BIOS/Firmware
  - System Calls

NVDIMM-N

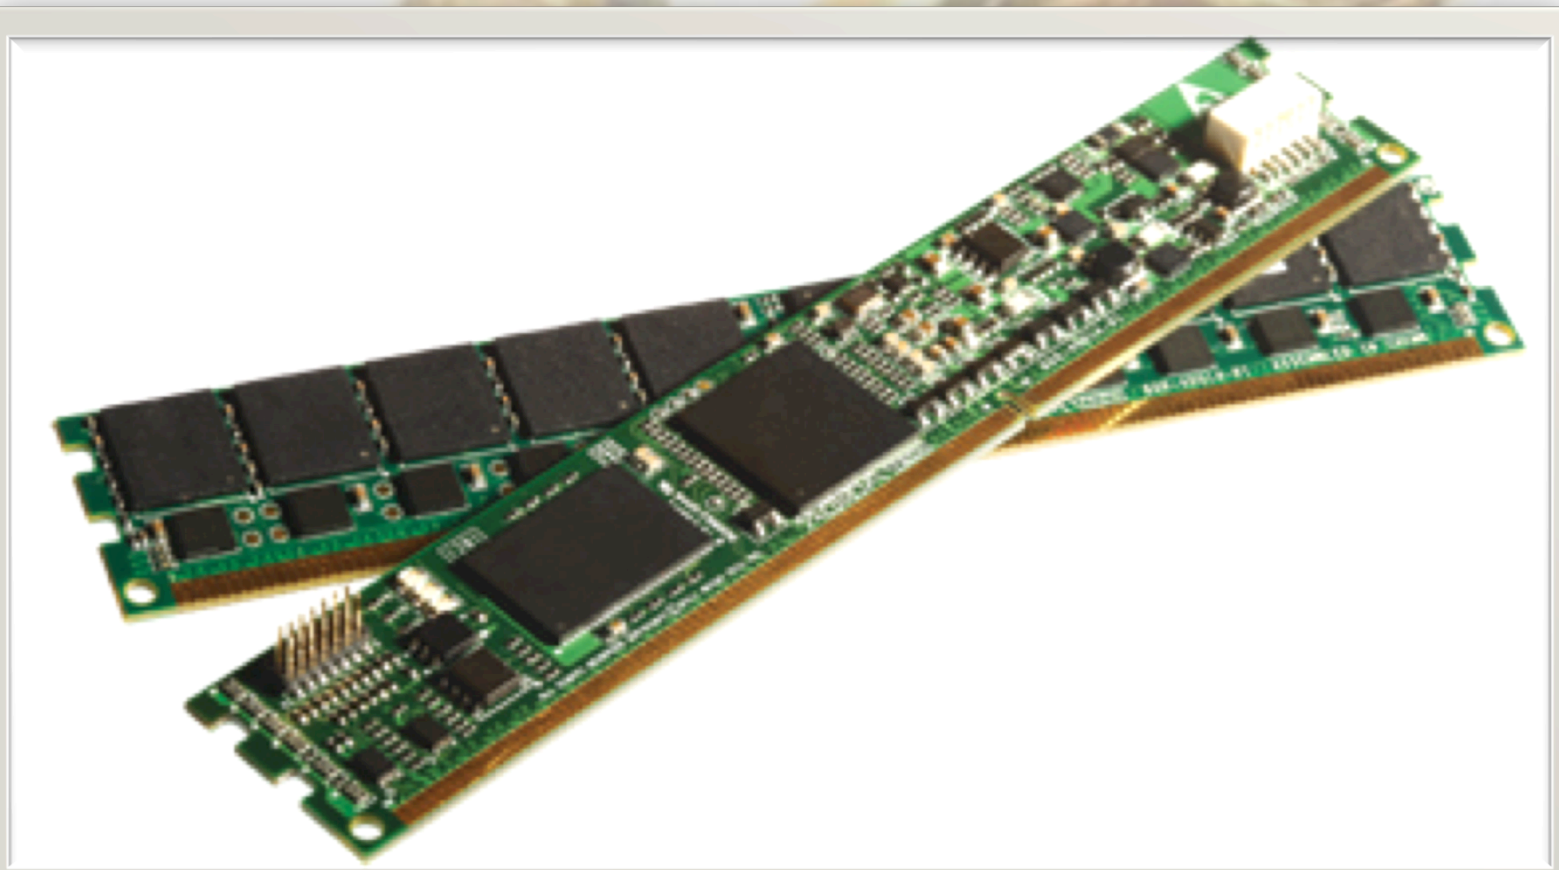

NVDIMM-N Operation

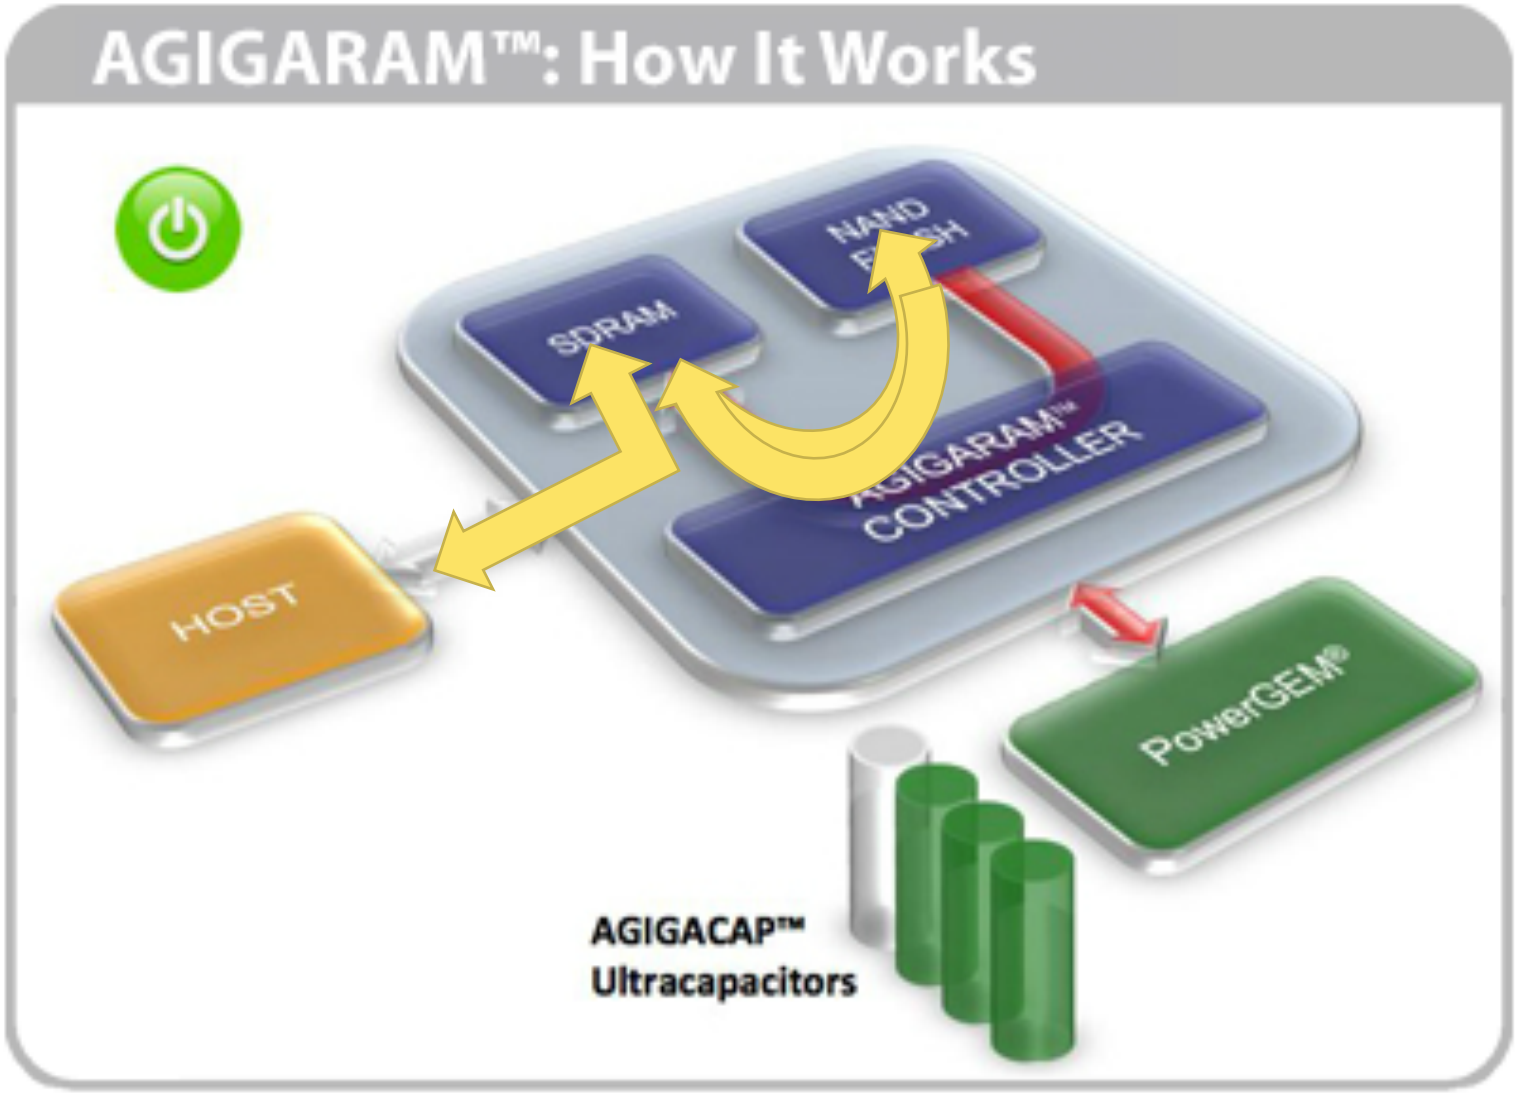

BIOS Support

Current BIOS

- DRAM corrupt at boot
- Power fail loses everything:
  - Memory
  - Registers
  - Cache

PM BIOS

- DRAM may boot with valid data
- Push DRAM to PM at power fail:
  - Memory
  - Registers
  - Cache

Different Kinds of NVDIMMs

	NVDIMM-N	NVDIMM-F	NVDIMM-P
Think of it as:	DRAM with Flash backup	SSD on the DRAM Bus	Emerging Memory, or Something Like It
Access Method	Byte or Block	Block	Likely Block (depends on what media type will be used)
Capacity Range	DRAM (tens of GB)	Flash (terabytes)	NVM (terabytes)
Latency	DRAM (tens of ns)	Flash (tens of µs)	NVM (hundreds of ns, TBD)
Other	Requires battery or capacitor		Specification in flux. Will support multiple media types

NVDIMM-F

Where Do NVDIMMs Fit?

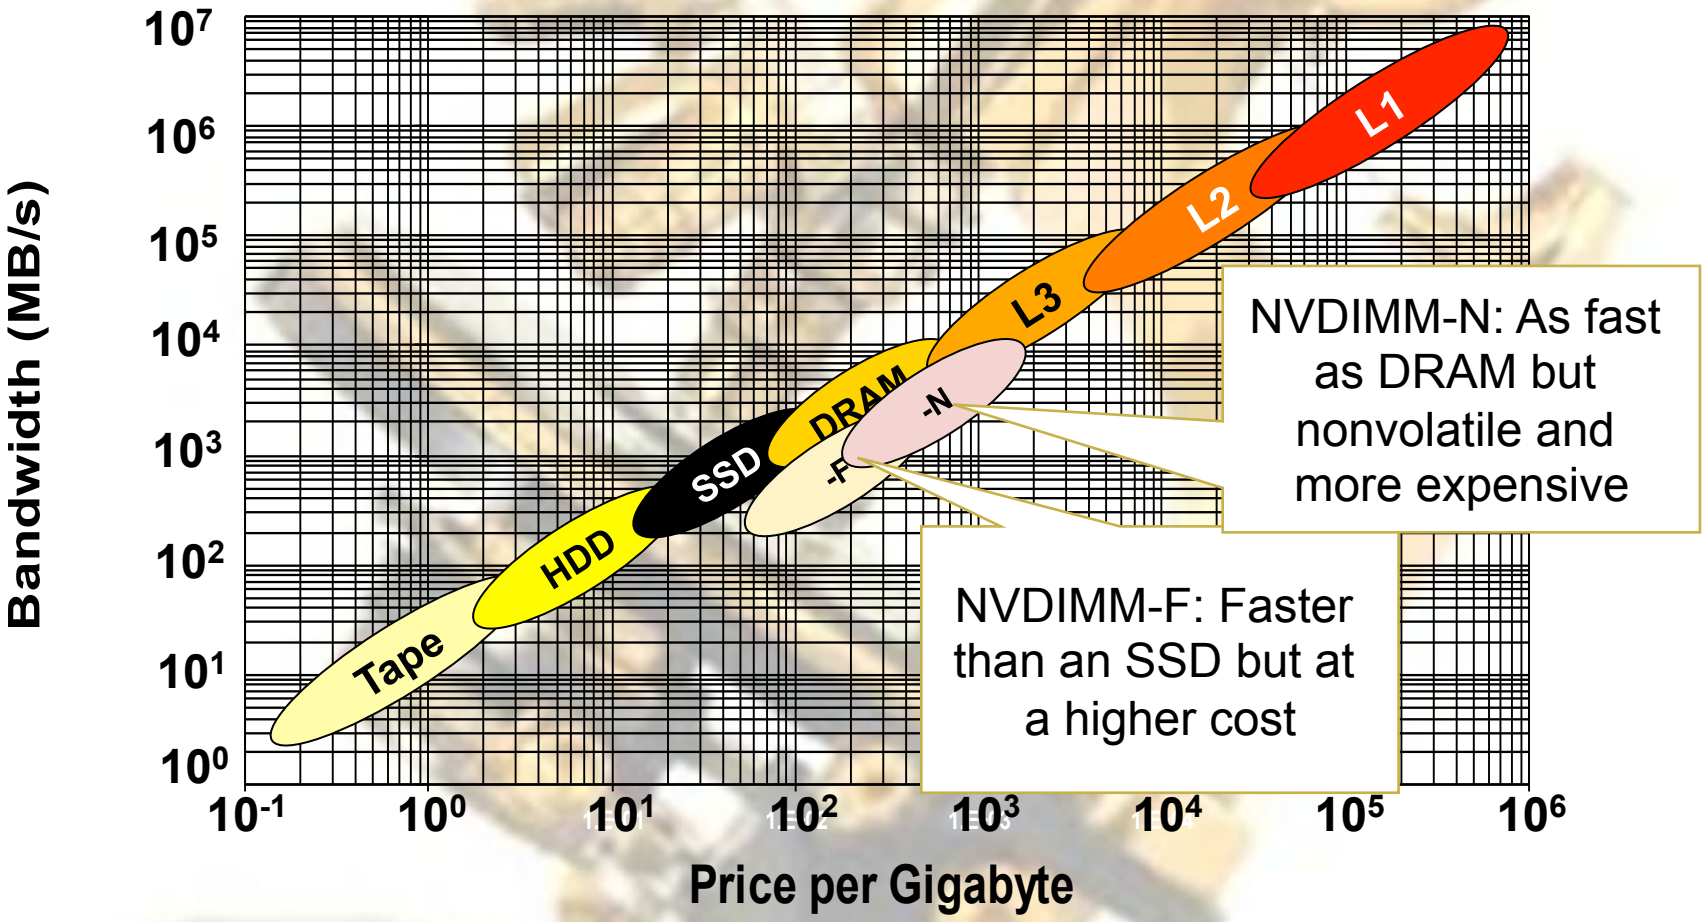

From: ***Solid State Drives in the Enterprise***

NVDIMM Unit Forecast

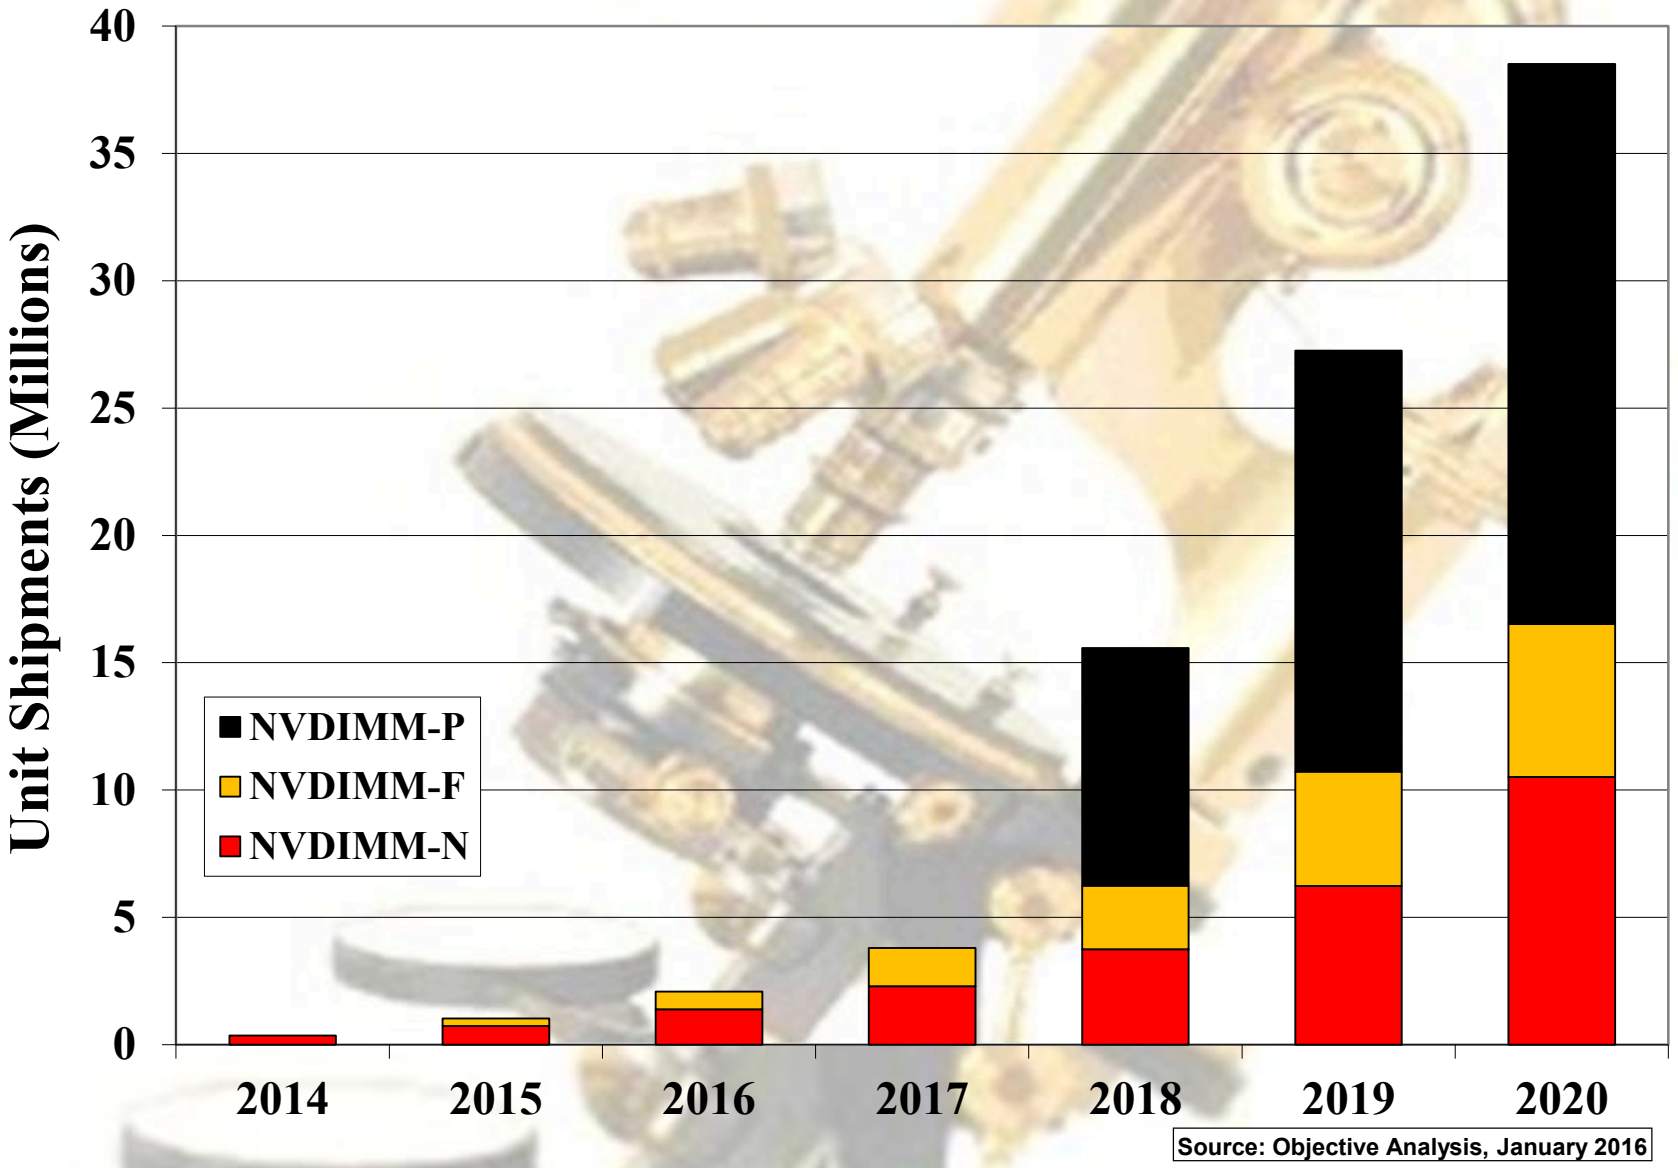

Operating System Support

- SNIA Persistent Memory Programming Model: https://www.snia.org/PM

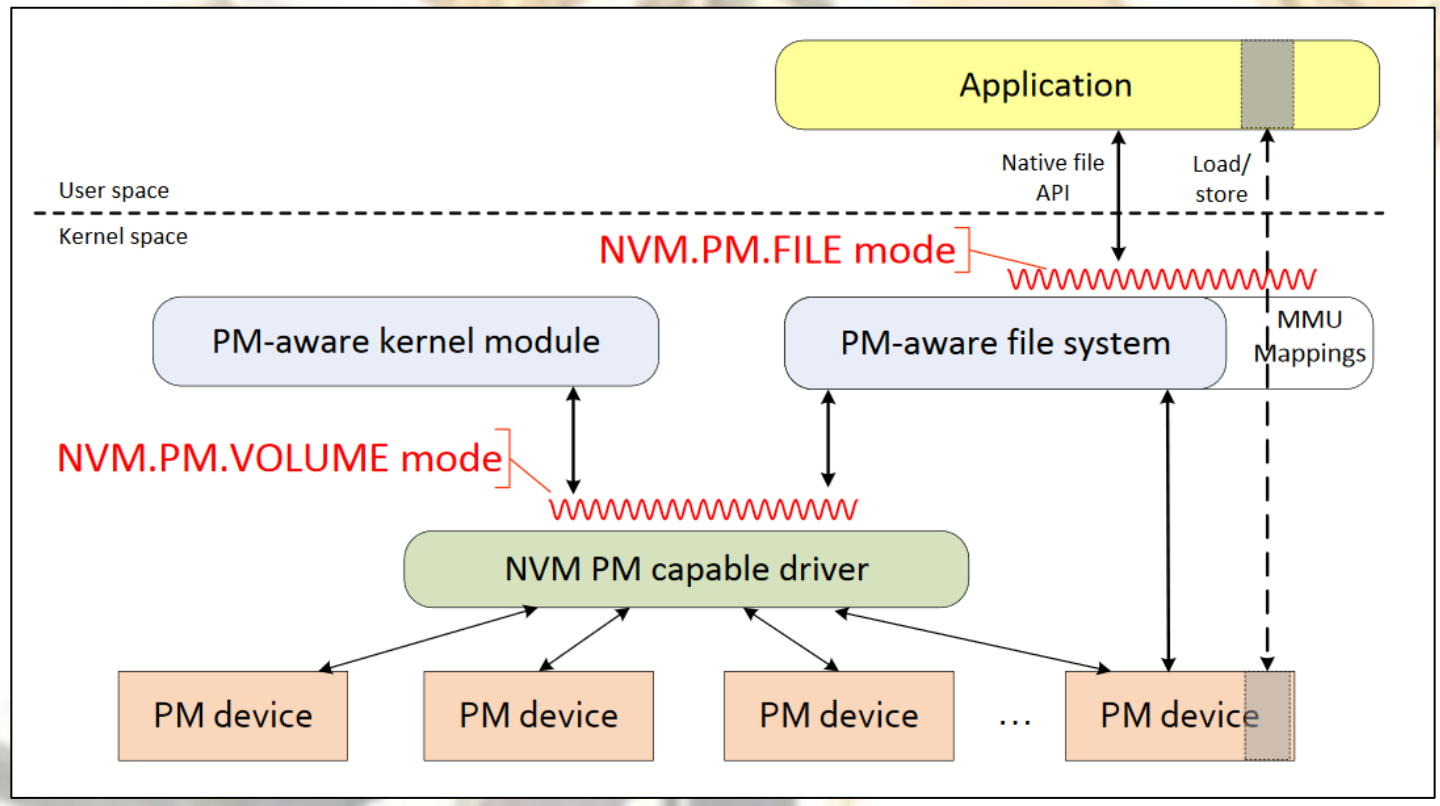

Processor Support

- Intel's new IA instructions
  - Cache & write buffer flush
  - NVDIMM-N "Commit": Copy DRAM to NAND
- Context switch issue?

Context Switches Become The Issue

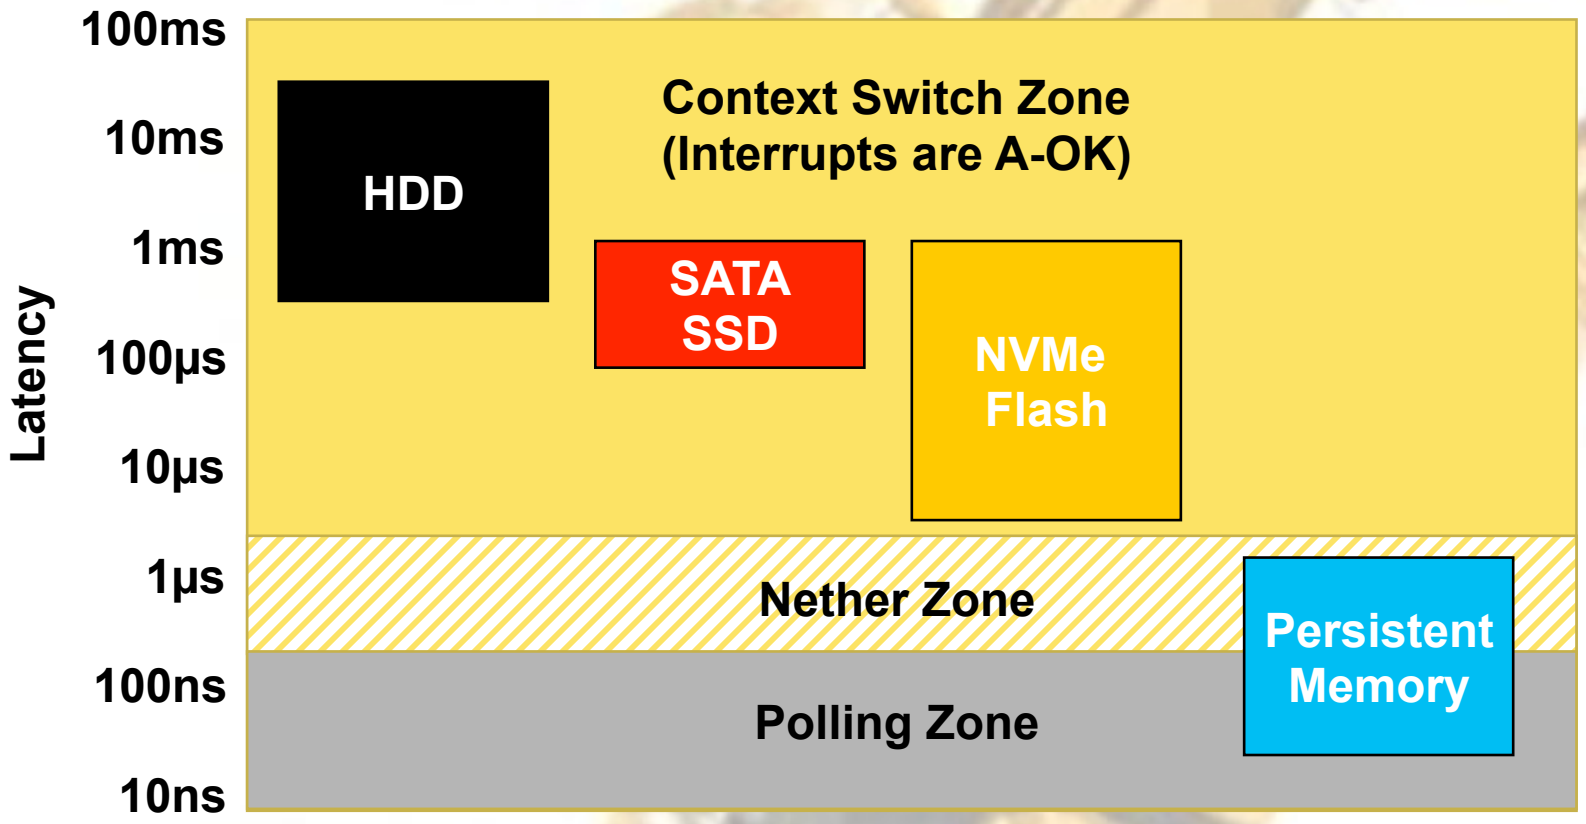

Making Sense of Persistent Memory

- Why is it needed?
  - Where does it fit?
- How will it impact other technologies?
  - Hardware
    - HDD, SSD, & DRAM
  - Software/Firmware
  - Processors
- What new chip technologies are coming?

Today's Memories Are Limited

	SRAM	DRAM	ROM	EEPROM	NOR	NAND
Nonvolatile	No	No	Yes	Yes	Yes	Yes
Erasable	Yes	Yes	No	Yes	Yes	Yes
Programmable	Yes	Yes	Factory	Yes	Yes	Yes
Smallest Write	Byte	Byte	N/A	Byte	Byte	Page
Smallest Read	Byte	Page	Byte	Byte	Byte	Page
Read Speed	V Fast	Fast	Fast	Fast	Fast	Slow
Write Speed	V Fast	Fast	N/A	Slow	Slow	Slow
Active Power	High	Med	Med	Med	Med	Med
Sleep Power	V Low	High	Zero	Zero	Zero	Zero
Price/GB	High	Low	Low	High	Med	V Low
Applications	Small Fast	Main Memory	Stable Code Volume	Serial #, Trim	Code	Data

Emerging Memories Perform Better

	MRAM	ReRAM	FRAM	PCM	XPoint
Nonvolatile	Yes	Yes	Yes	Yes	Yes
Erasable	Yes	Yes	Yes	Yes	Yes
Programmable	Yes	Yes	Yes	Yes	Yes
Smallest Write	Byte	Byte	Byte	Byte	Byte
Smallest Read	Byte	Byte	Byte	Byte	Byte
Read Speed	Fast	Fast	Fast	Fast	Fast
Write Speed	Fast	Fast	Fast	Fast	Fast
Active Power	Low	Med	Low	High	High?
Sleep Power	Low	Low	Low	Low	Low
Price/GB	High	High	High	High	High?
Applications	Niche	TBD	Low Power	Obsolete	Main Memory

Memory Attributes By Technology

	DRAM	NAND	MTJ MRAM	ST-MRAM	FRAM	PCM	RRAM
Read Speed	10ns	10,000ns	35ns	<10ns	55ns	16ns	10ns
Write Speed	10ns	200,000ns	35ns	<10ns	55ns	30ns	10ns
Standby Current	45mA	<1mA	<1mA	<1mA	<1mA	<1mA	<1mA
Endurance	Infinite	10^4	10^{11}	10^{13}	10^{14}	10^6	10^6
Retention (years)	10^{-9}	>10	>20	>20	>10	>10	>10
Scaling Limit	10nm?	TBD	65nm	5nm	5nm	5nm	5nm
Cell Size	6-8f²	4f² effective with MLC	35-40f²	8-9f²	8-20f²	4f²	4f²
Select Device	Transistor	Transistor	Transistor	Transistor	Transistor	Diode	Diode
Max Bits/ Cell	1	4	2	4	1	4	TBD

But, Most Importantly, "What Does It Cost?"

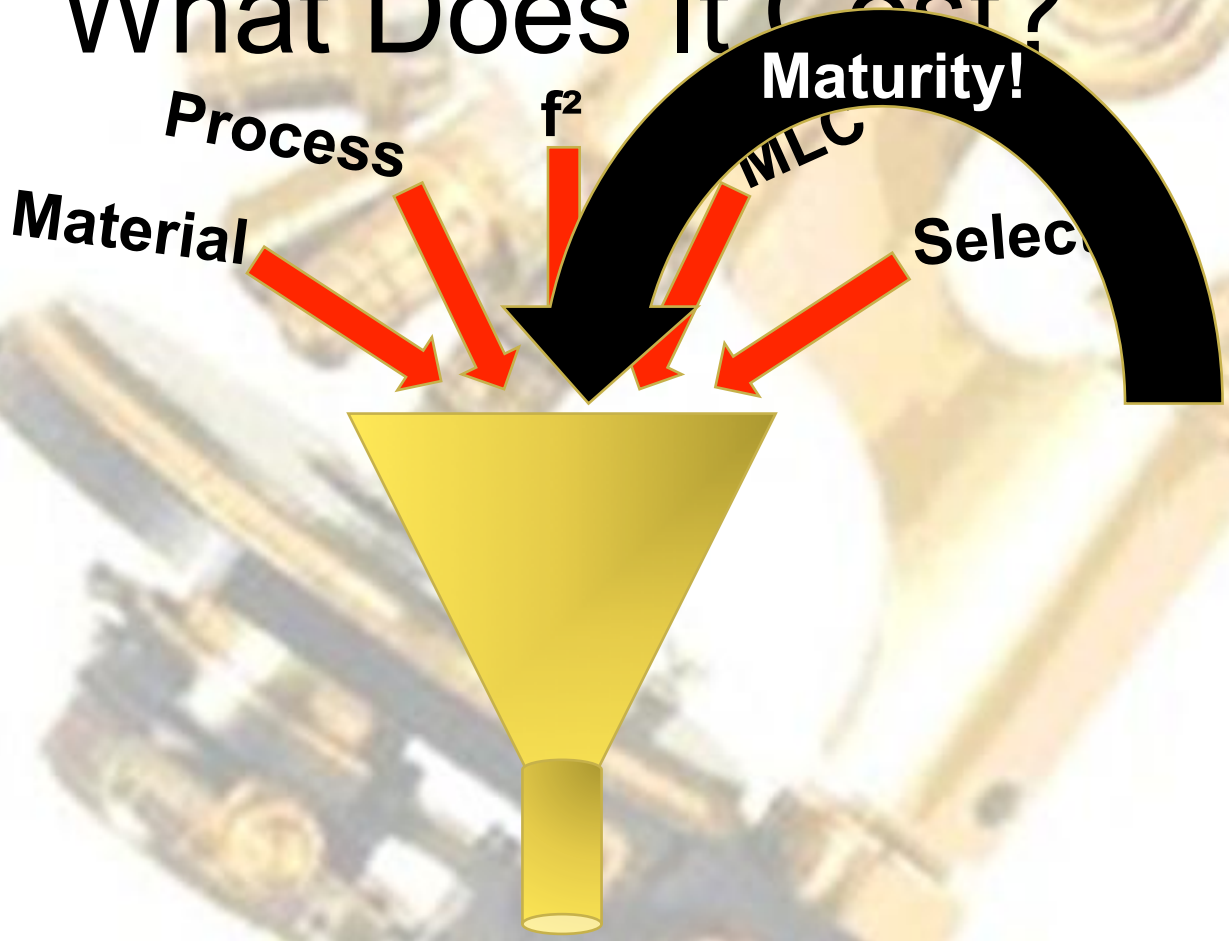

Summary

Memory/Storage Hierarchy

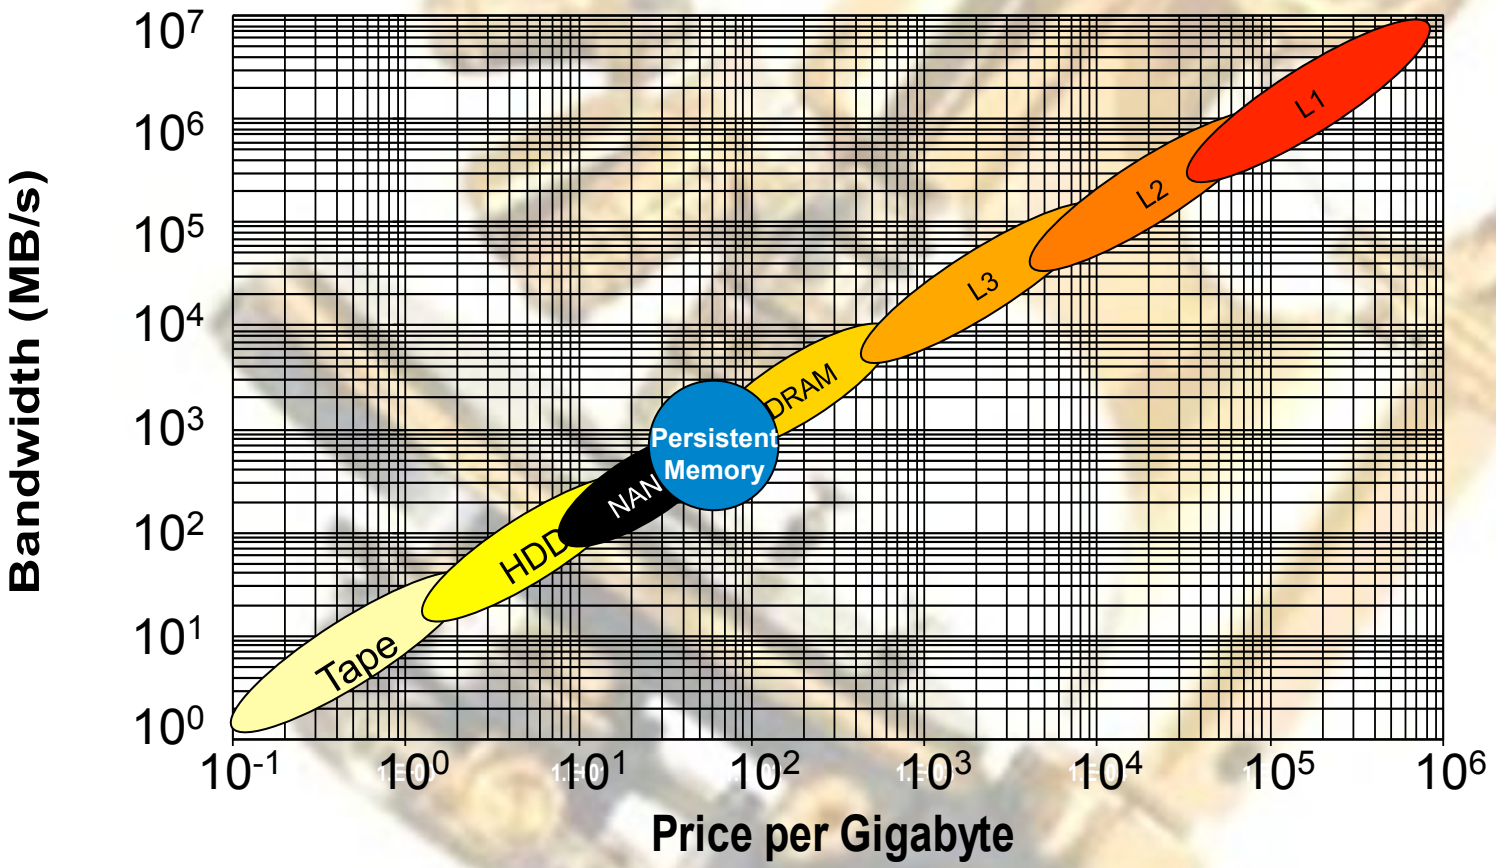

From: Solid State Drives in the Enterprise

Adding Layers Over Time

1960	1970	1980	1990	2000	2010	2020	2030
CPU	CPU	CPU	CPU	CPU	CPU	CPU	CPU
Core	Core	DRAM	Cache	Cache	Cache	Cache	Cache
Tape	HDD	HDD	DRAM	DRAM	DRAM	DRAM	HBM
	Tape	Floppy	HDD	HDD	Flash	SCM	SCM
			Floppy	CD-ROM	HDD	Flash	Flash
					Cloud	HDD	HDD
						Cloud	Cloud

Thank You!

Jim Handy

OBJECTIVE ANALYSIS – Semiconductor Market Research