STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

Riding the Long Tail of Optane's Comet

Emerging Memories, CXL, UCIe, and More

Jim Handy, Objective Analysis

Agenda

- Optane’s brief history
- Today’s Alternatives
- Optane’s Legacy
- CXL
- UCIe
- Future Thoughts

An Optane Timeline

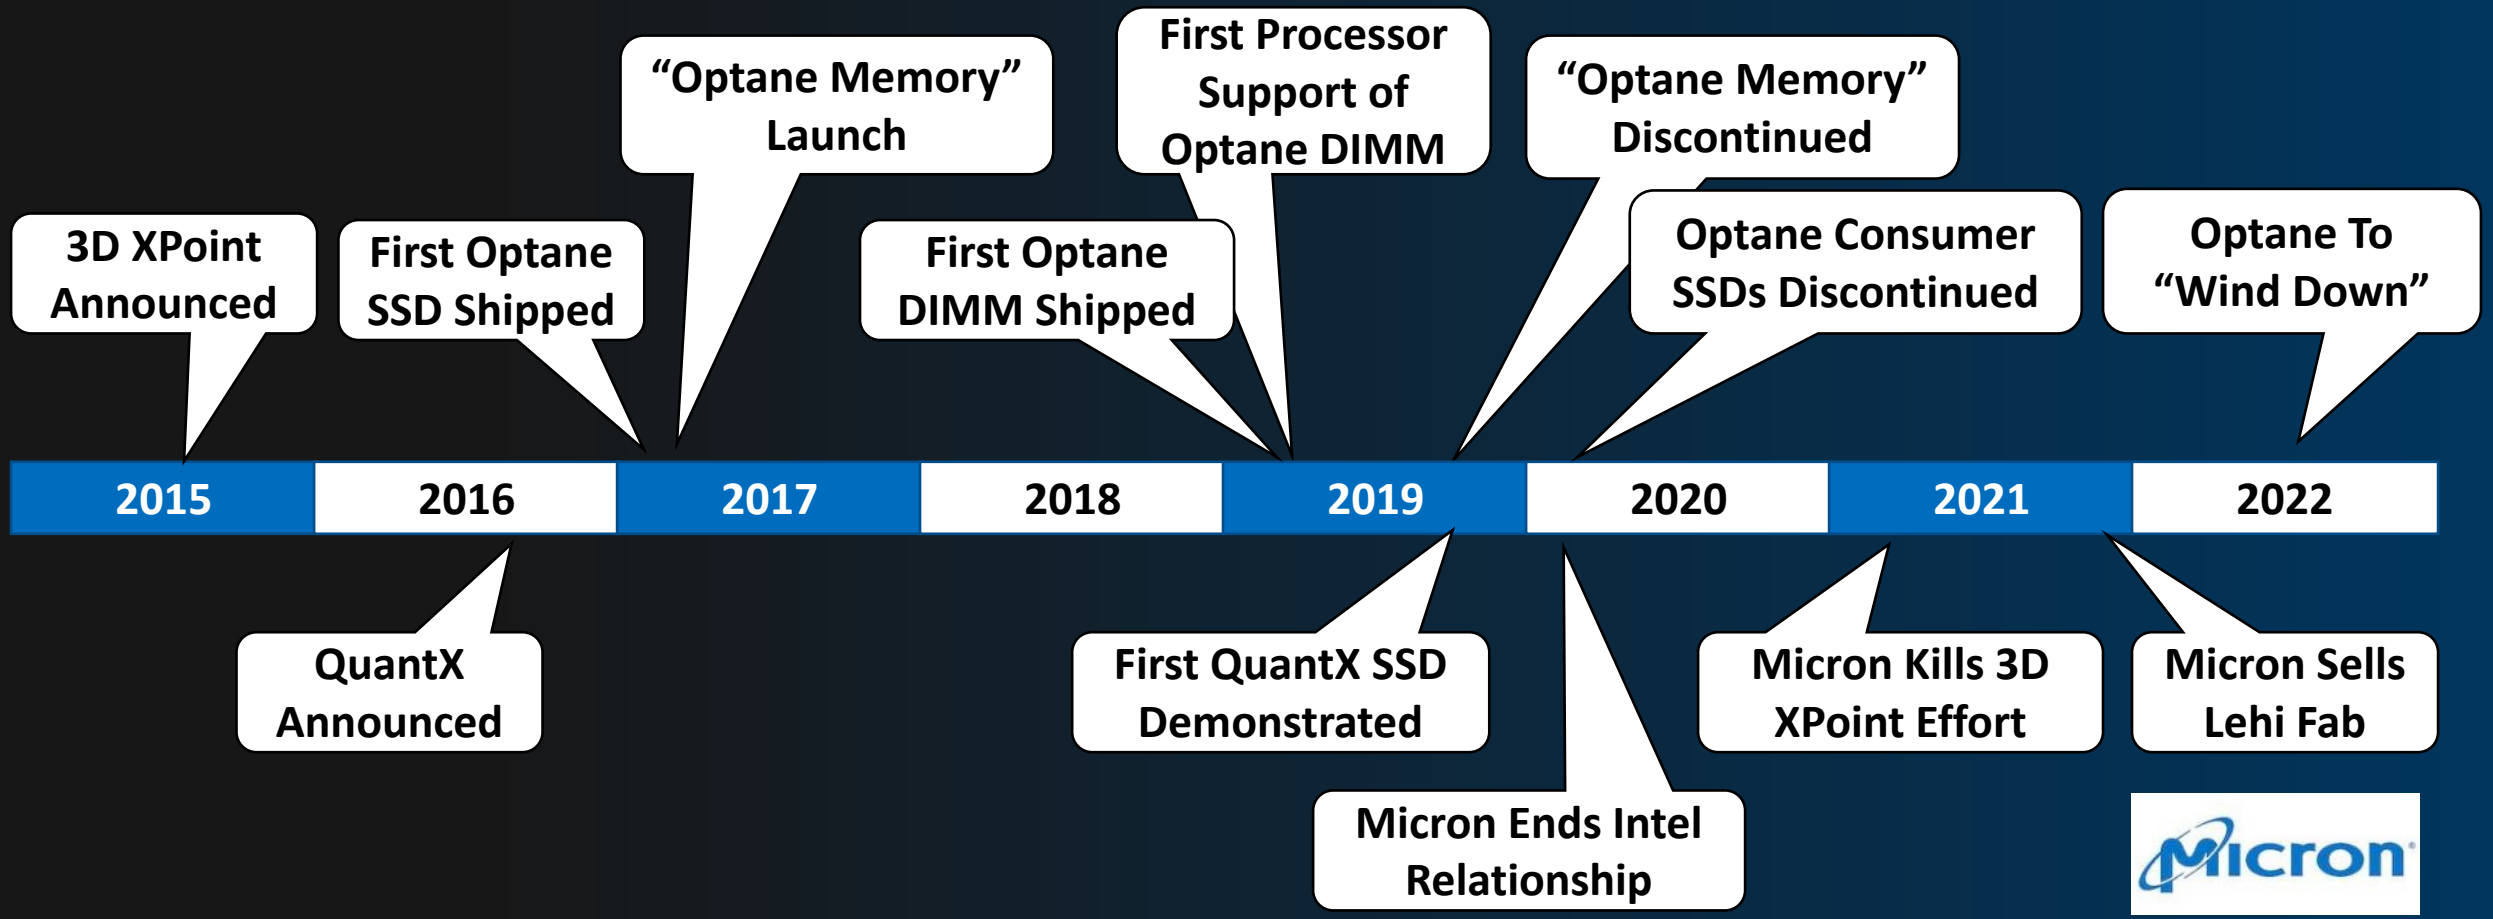

Today’s Alternatives

Alternatives to Optane

	Persistent?	Speed	Cost/DRAM	Issues
Optane	Yes	30%	50%	Winding Down
NVDIMM-N	Yes	100%	200%	Battery/Capacitor
MRAM DIMM	Yes	100%	1,000%	Compatibility
Fast SSD	Yes	0.1%	20%	Slow
Added DRAM	No	100%	100%+	Bus loading

Optane’s Still Around

- Current inventory fulfilling needs
- Ongoing low-level demand
- Support already in place

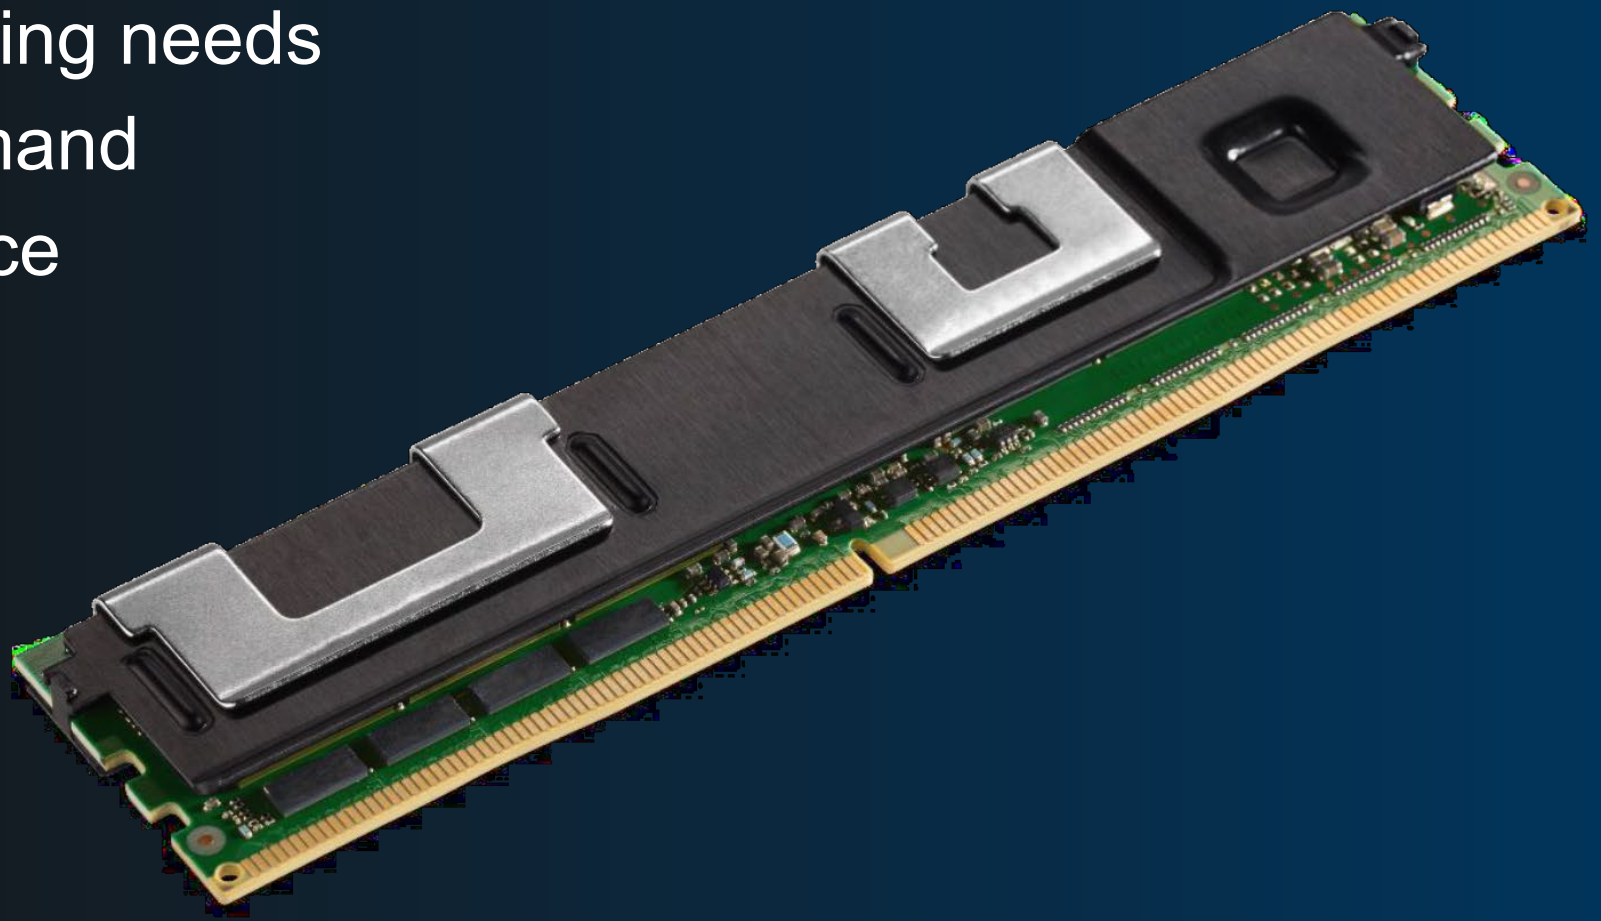

NVDIMM-N

- Faster than Optane
- >2X the cost of DRAM
- Requires back-up power source

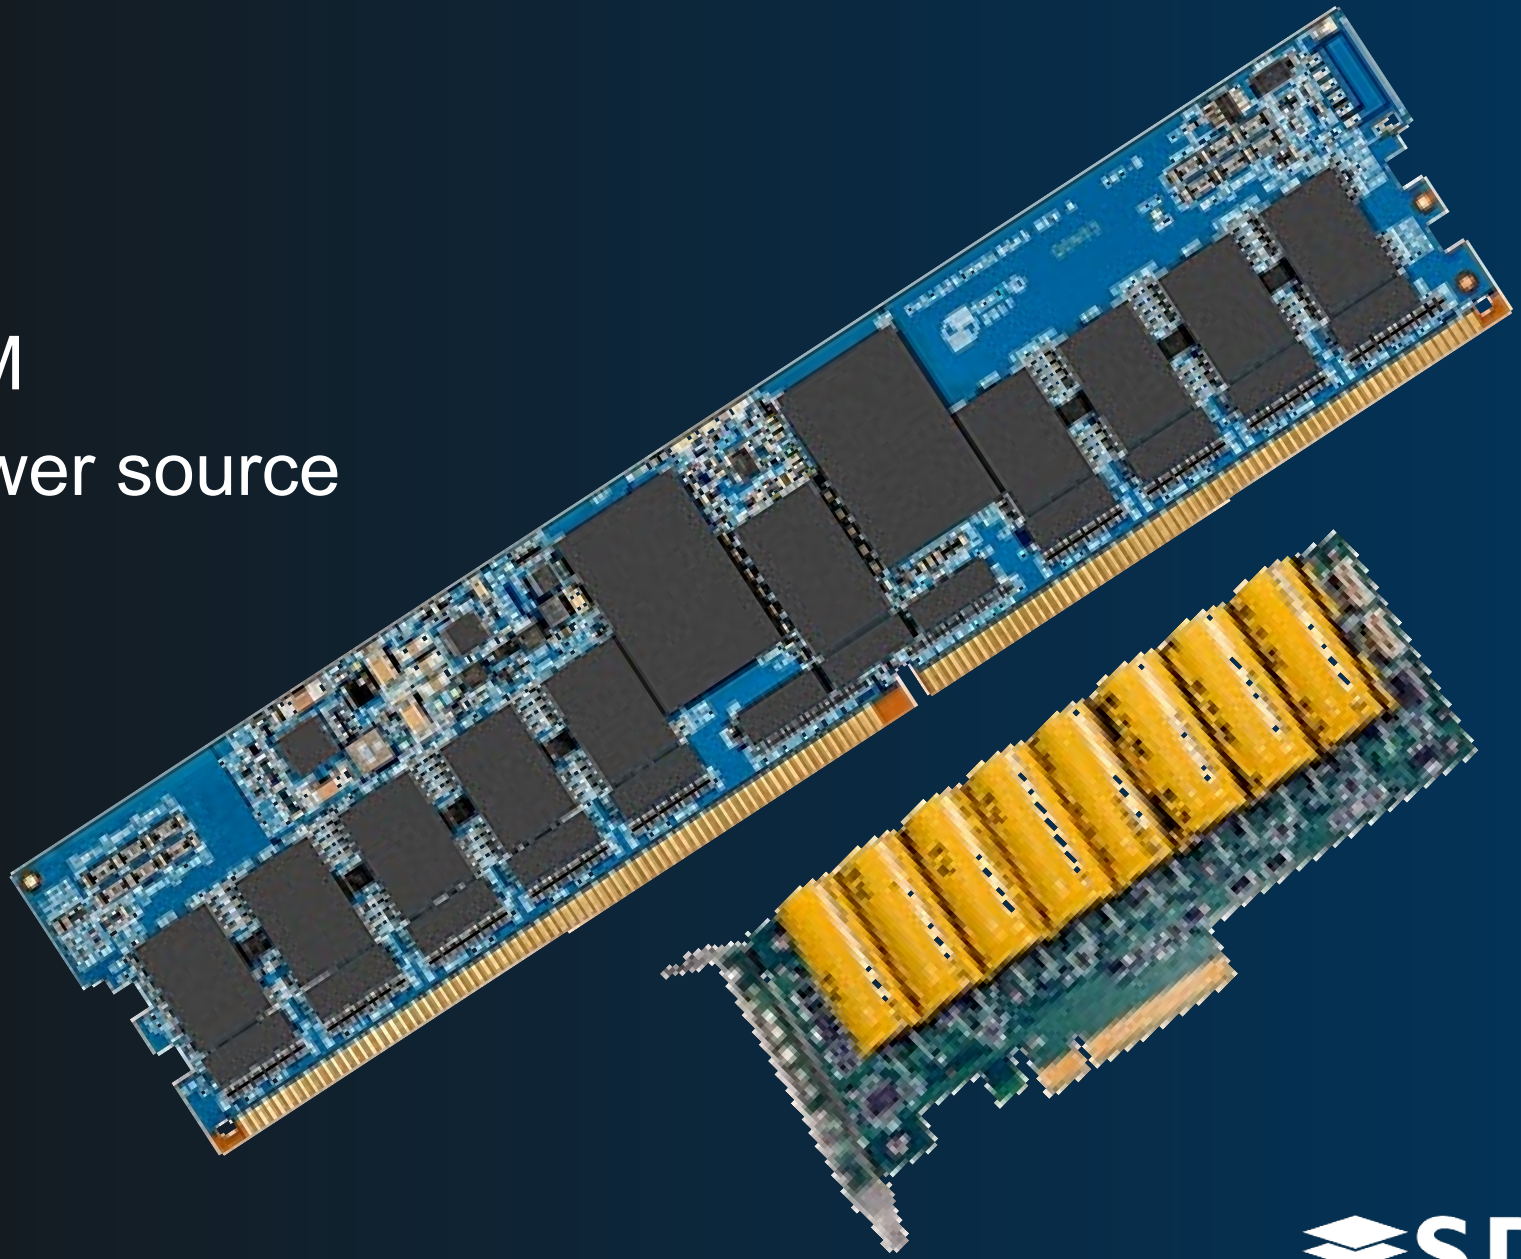

MRAM DIMM

- Production started in 2017
- DDR3 only
- Requires changes to processor
- >100X the cost of DRAM

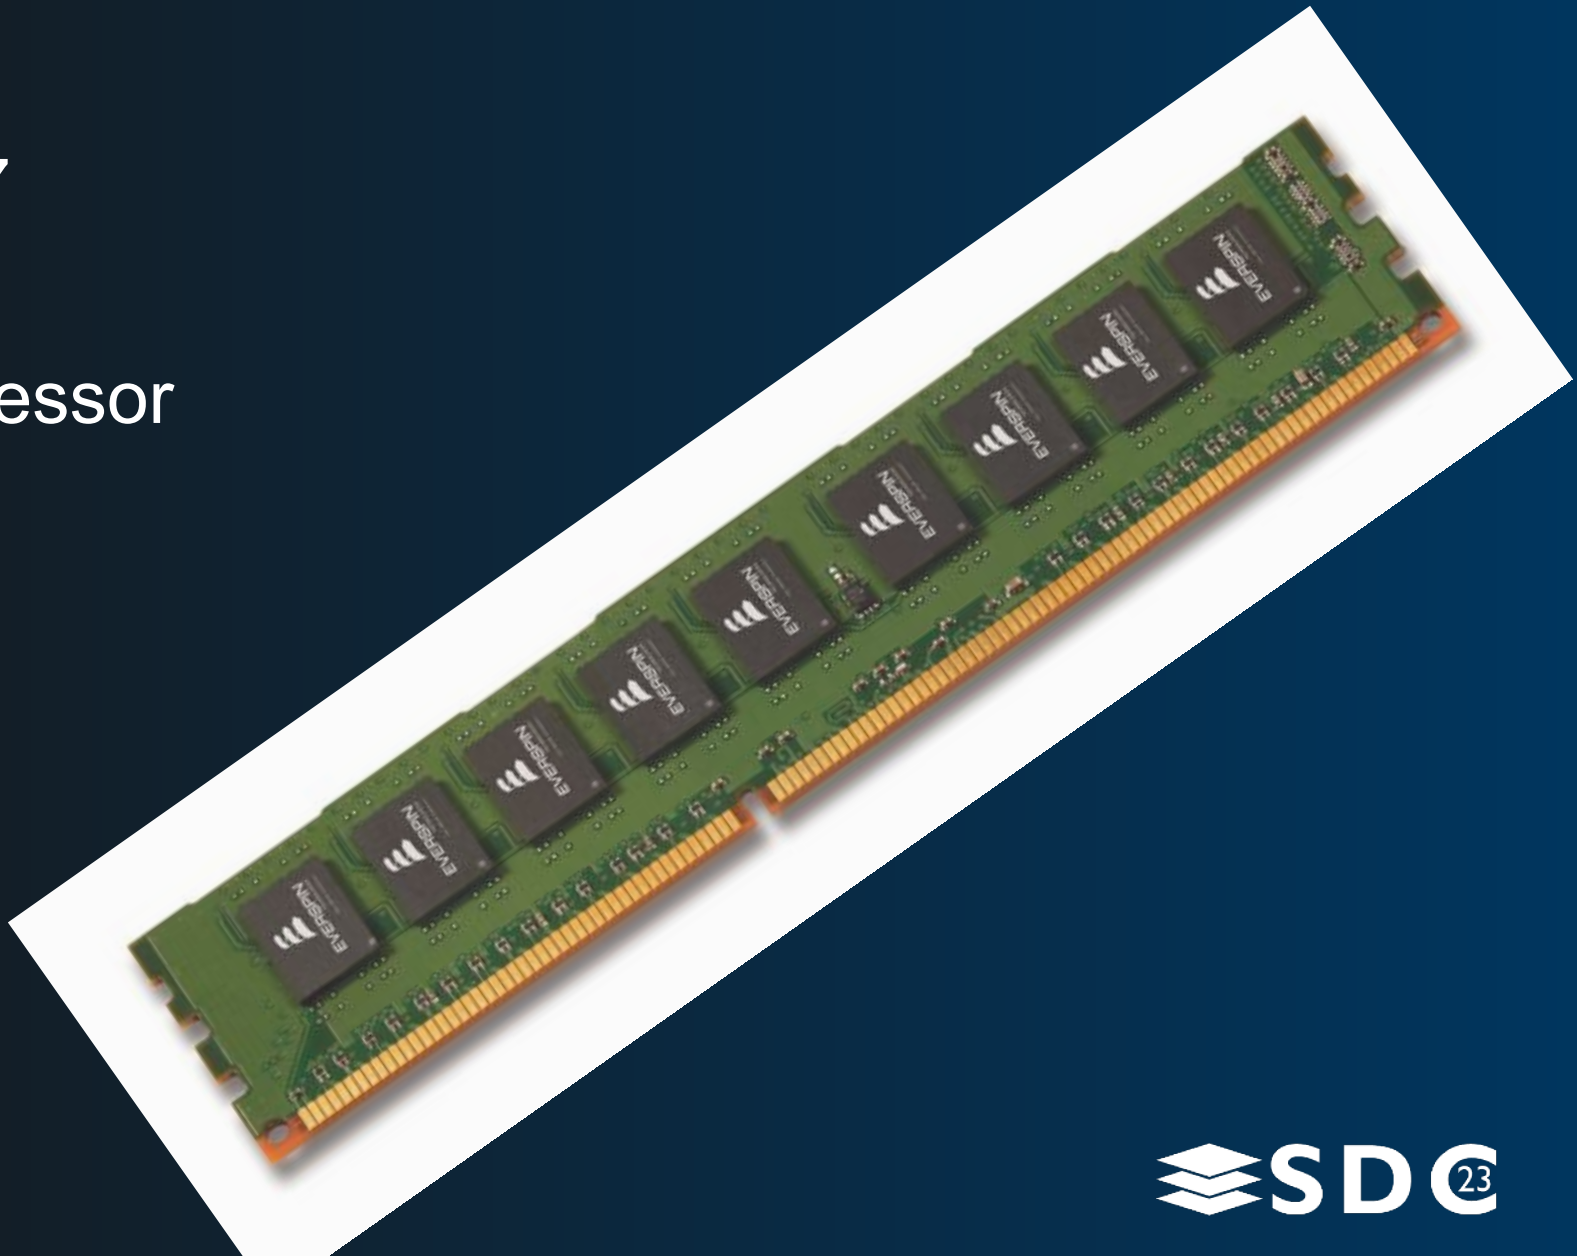

SDC 23

MRAM Is Already in the Enterprise

- IBM FlashCore Modules use MRAM instead of DRAM
  - Store translation tables
  - Buffers for write coalescing etc.
  - Easy way to protect data in flight
  - Fast path to persistence
- Consumer adoption is growing
  - Wearables, vehicles, health monitors, etc.
  - Drives growing wafer volume
  - Economies of scale will reduce prices

Fast SSD

- SSD? *Really*???
- Kioxia and Samsung both advocate this
  - Special NAND chip architectures
  - Uses SLC NAND
    - ~6X the price of TLC NAND
- Which performs better:
  - Fast & Small (DRAM) or
  - Slow & Big (NAND)?

Conundrum: Fast & Small, or Slow & Big?

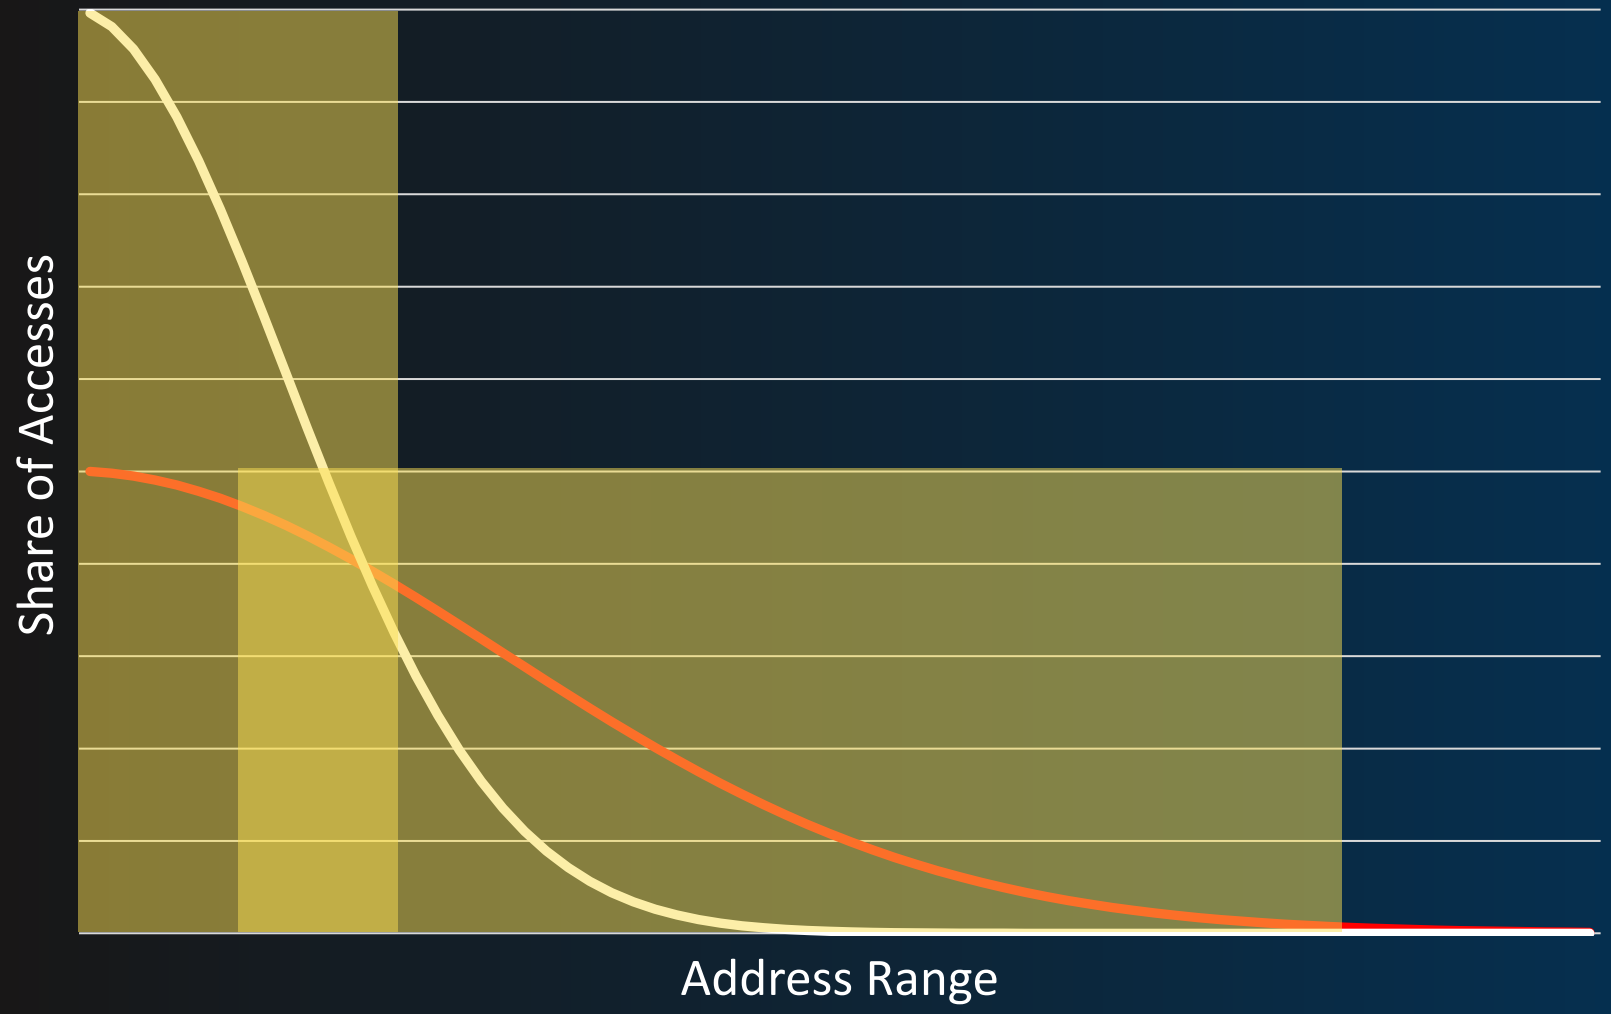

It’s Getting Harder To Add More DRAM

- “Fast & Small” includes large DRAM approaches
  - But large DRAMs increase loading, slowing the memory channel
  - Adding memory channels increases processor power & pin count
  - This is a thorny problem!
- IBM has been wrestling with this for years
  - POWER architecture uses buffered DIMMs with non-DDR interface
    - OpenCAPI led to the OMI interface
  - CXL is adding slower memory to the CXL channel
    - Disaggregated memory
    - Memory tiering
    - Will discuss this shortly

The Short Story: There Are Many Options

	Persistent?	Speed	Cost/DRAM	Issues
Optane	Yes	30%	50%	Winding Down
NVDIMM-N	Yes	100%	200%	Battery/Capacitor
MRAM DIMM	Yes	100%	1,000%	Compatibility
Fast SSD	Yes	0.1%	20%	Slow
Added DRAM	No	100%	100%+	Bus loading

Optane’s Legacy

Optane's Legacy: New Programming Models

SNIA NVM Programming Model is just the start

- Hierarchical memory tiers (HBM, DDR, CXL)
- Memory disaggregation is coming
  - Reduces "Stranded Memory"
- Models may move into the chiplet
  - Persistent cache (with an emerging memory)

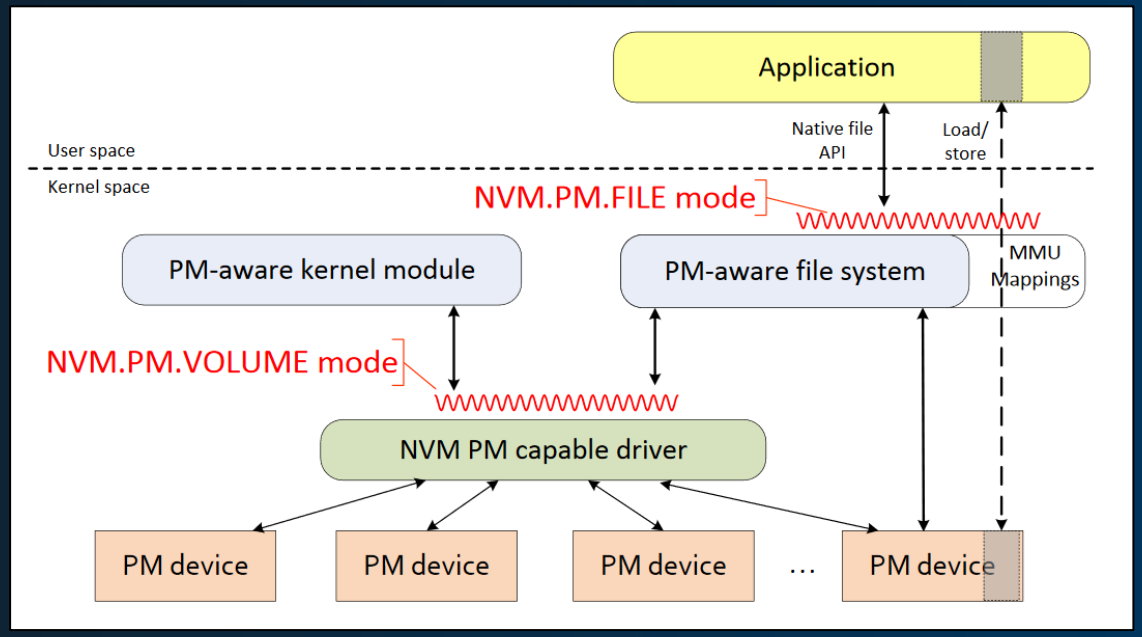

Optane's Legacy: A Fresh Look at Memory

OLD WAY

- All DRAM, all one speed
- Persistence is a storage thing
  - Slowed by context switches
- Memory is only put on the memory channel
- Only memory is put on the memory channel

NEW WAY

- Mixed memories, mixed speeds
- Persistence OK in memory
  - No context switches
- 4 channels: HBM, DDR, CXL, & UCIe
- Memory-Semantic SSDs on CXL

SDC 23

New Thoughts on Context Switches

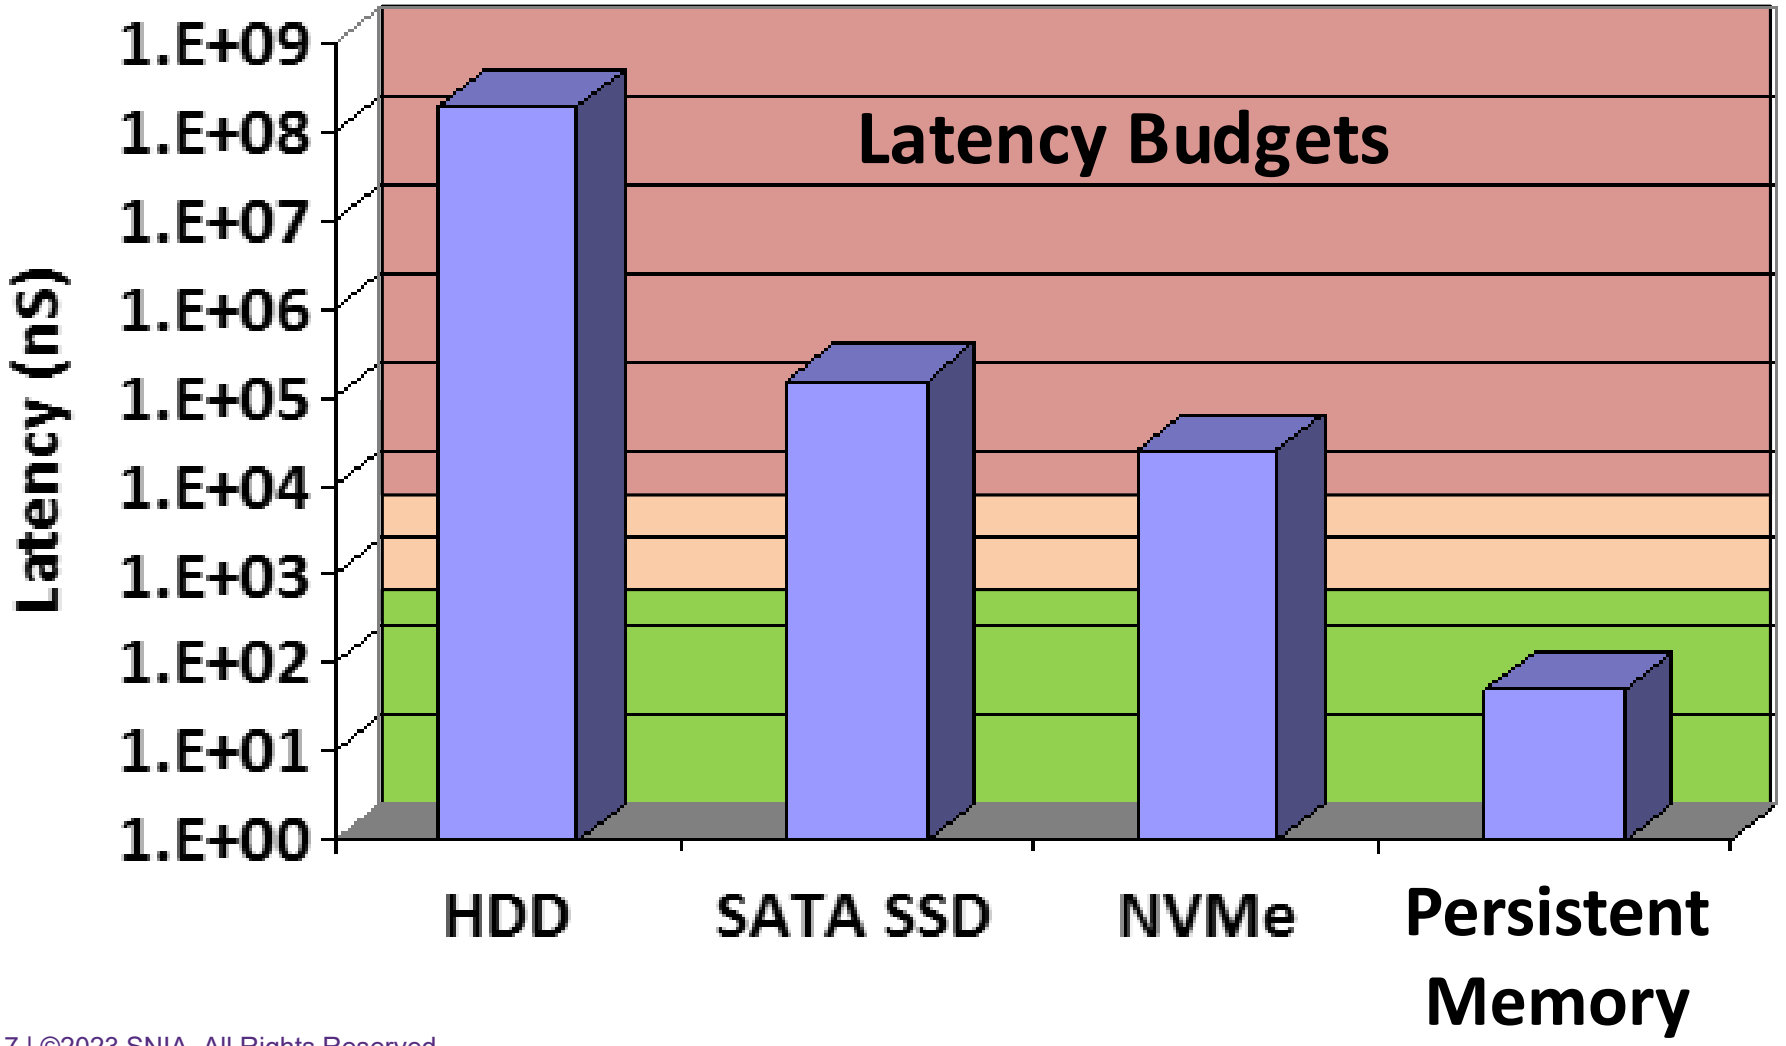

Context Switch

Polling used when a context switch would be too slow

CXL

DDR-T: Intel's Original Approach to Slower Memory

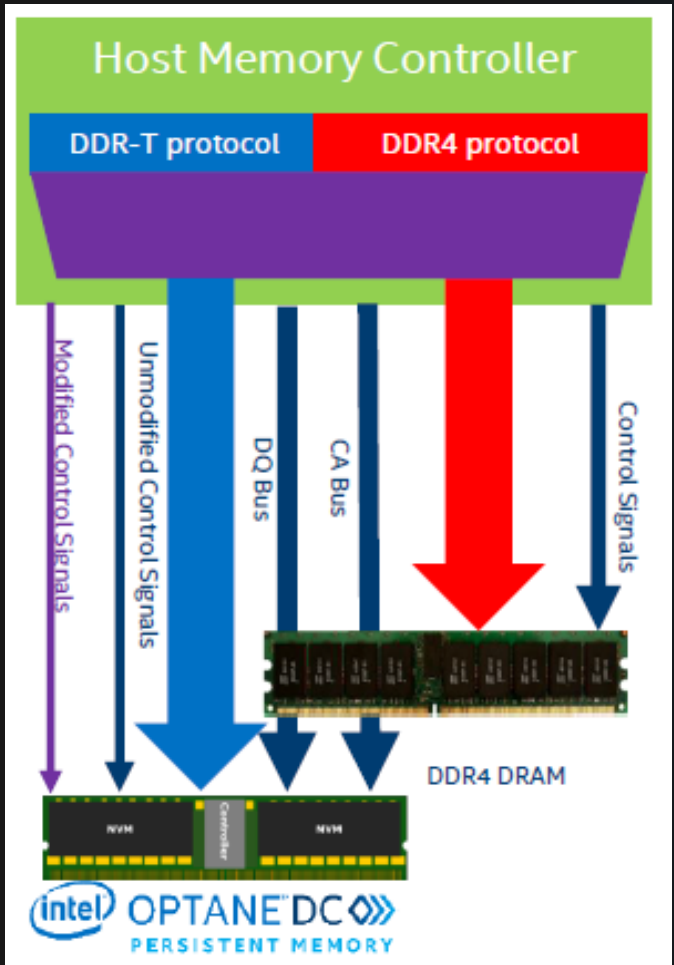

DDR-T for Optane

- Handles both fast & slow memory
  - Transactional protocol supports slow writes
- Based 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
- Migration from DDR4 to DDR5 a colossal headache!

SDC 23

CXL Solves Multiple Problems

- Removes processor's DDR limitation
  - A processor could use DDR4 or DDR5, but not both
  - CXL allows far memory to use any interface
    - With OMI near memory becomes similarly independent
- Supports memory disaggregation
  - No "Stranded Memory"
  - Memory pools can be dynamically allocated
  - Data sets can be moved from processor to processor
- Paves the way for UCIe

Any Memory Talks to Any Server

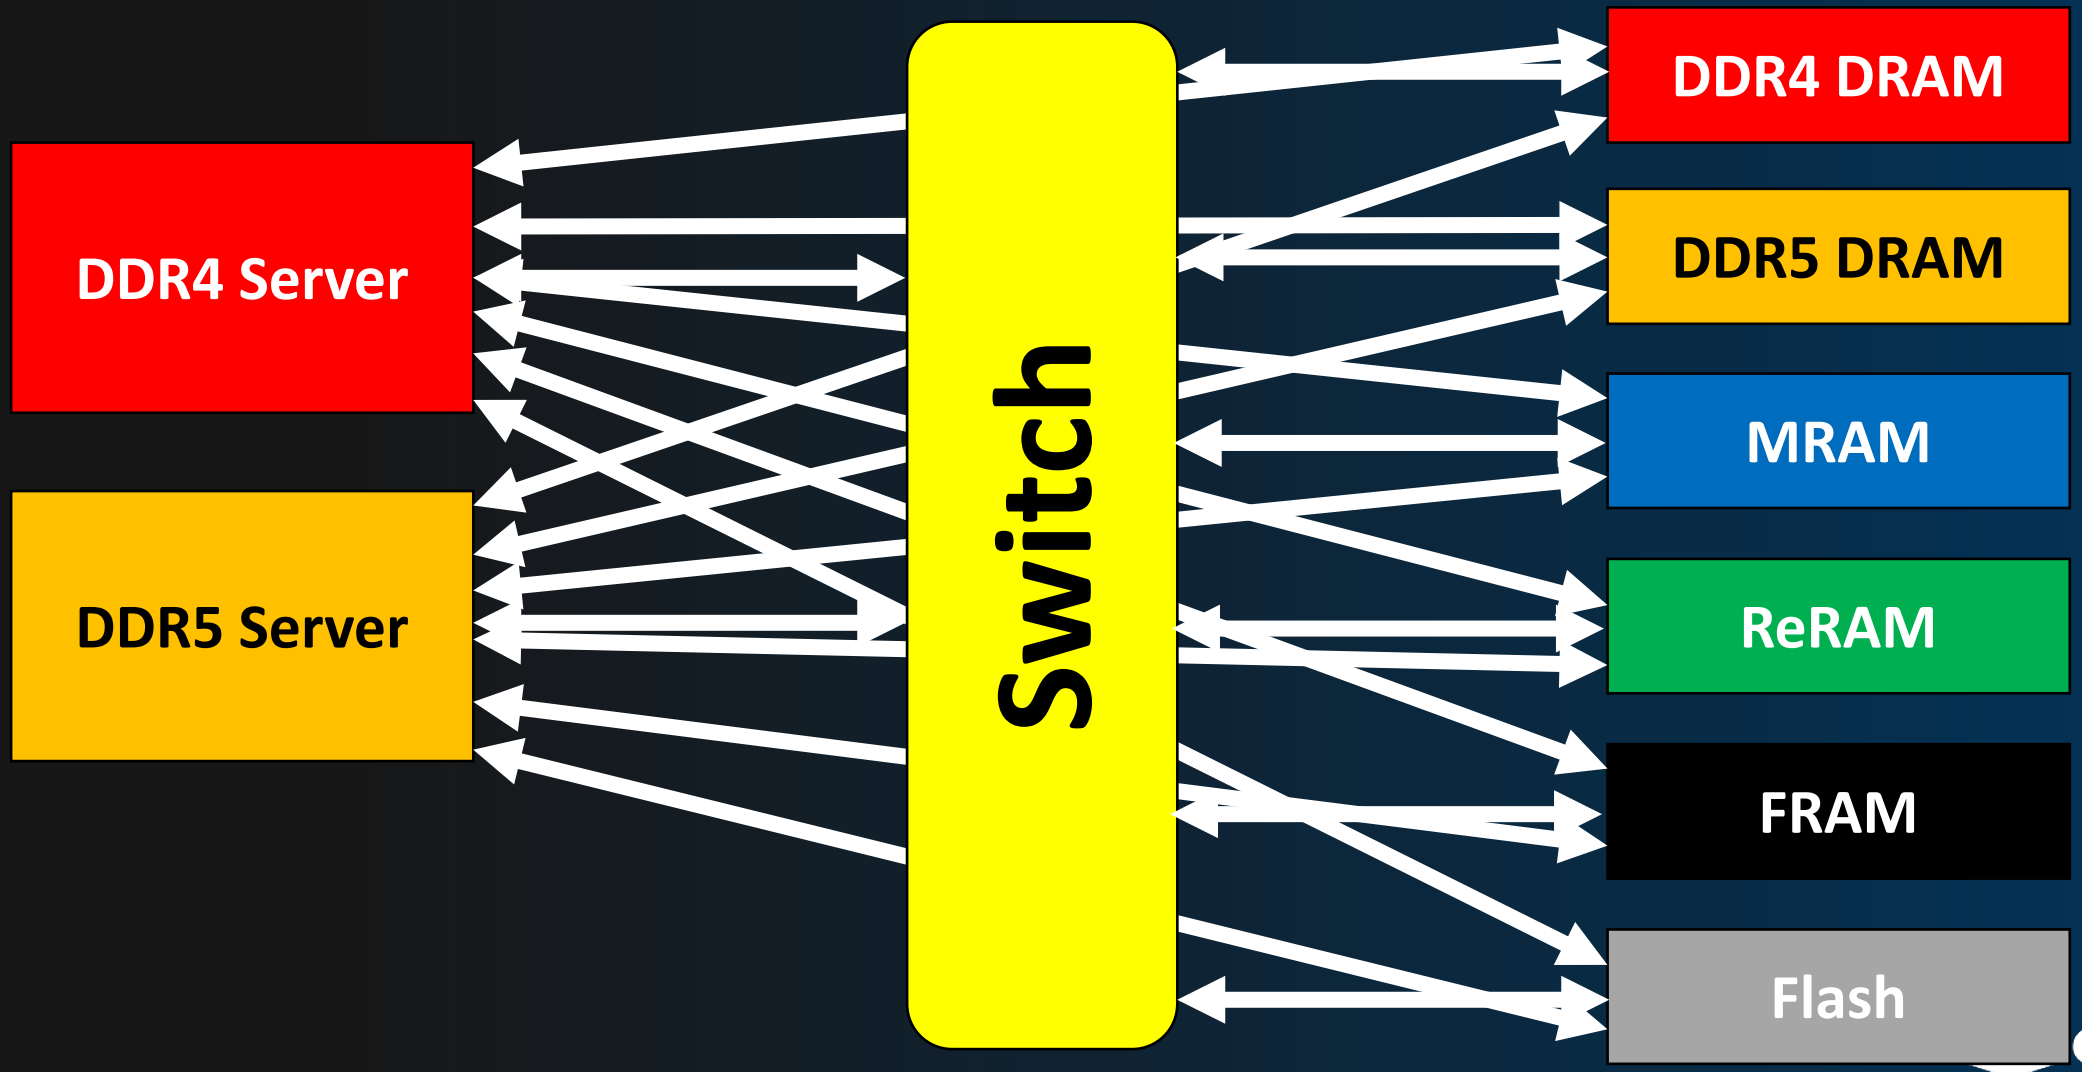

CXL 3.0 Supports Memory Fabrics

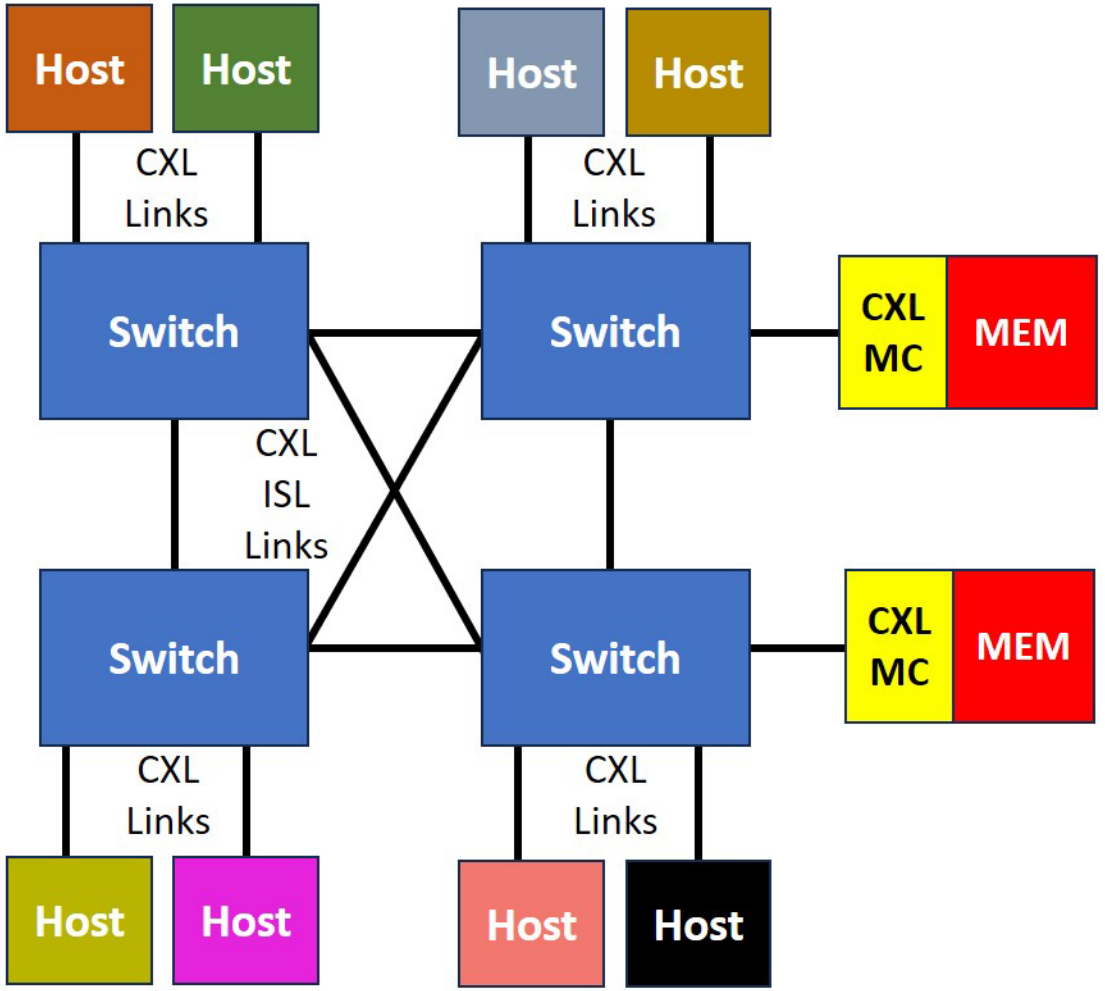

- Near Memory at CPU
- Far Memory on CXL
- CXL to support multiple Far Memory configurations
  - Large Memory
  - Memory Pools
  - Memory Sharing
    - Used for trading messages
  - Memory Fabrics
- No memory interface dependencies

UCIe

UCIe is CXL for Chiplets

A Standardized Chiplet Interface

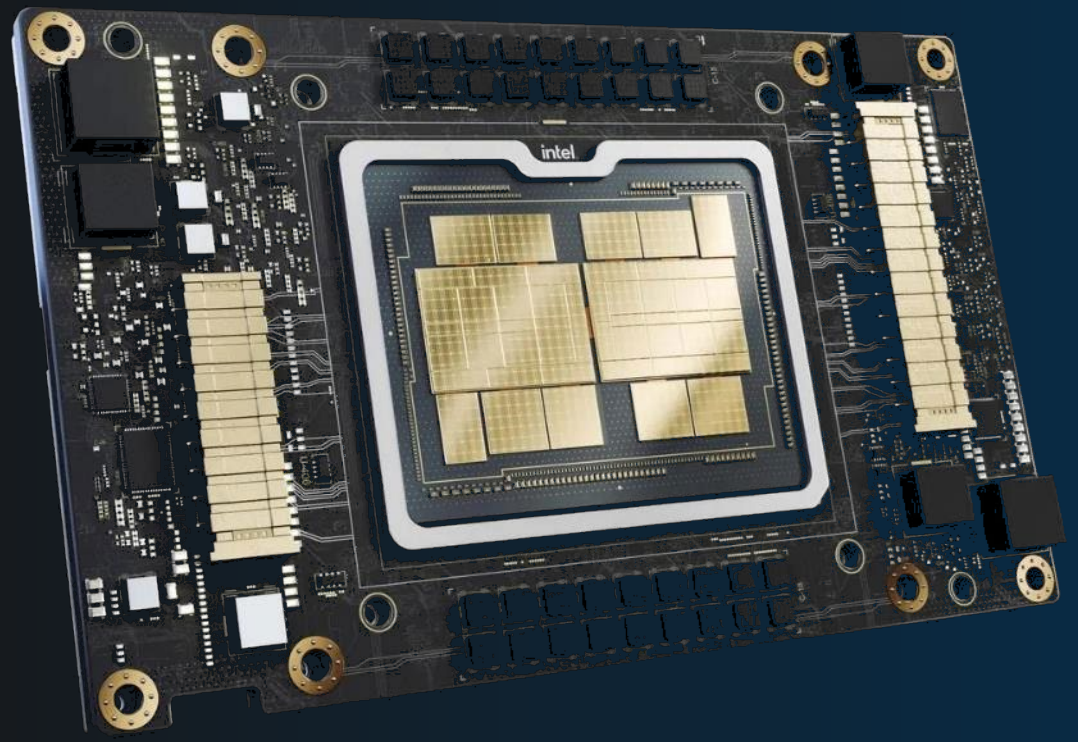

Supports multiple sources, and multiple customers

UCIe and Memories

- Mixed processes optimize cost/performance
  - Logic in a CMOS logic process
    - In logic SRAM & NOR flash are the only options for on-die memory
  - Memory chiplet in a memory process
    - DRAM, MRAM, ReRAM, FRAM, PCM…
  - Significant die area & cost reductions
- Commoditizes chiplets
  - One memory chiplet can be used by multiple logic companies
    - Increases volume, lowers costs
  - All vendors' parts equivalent
    - Vendors compete on price

SRAM Is No Longer Suited to CMOS Logic

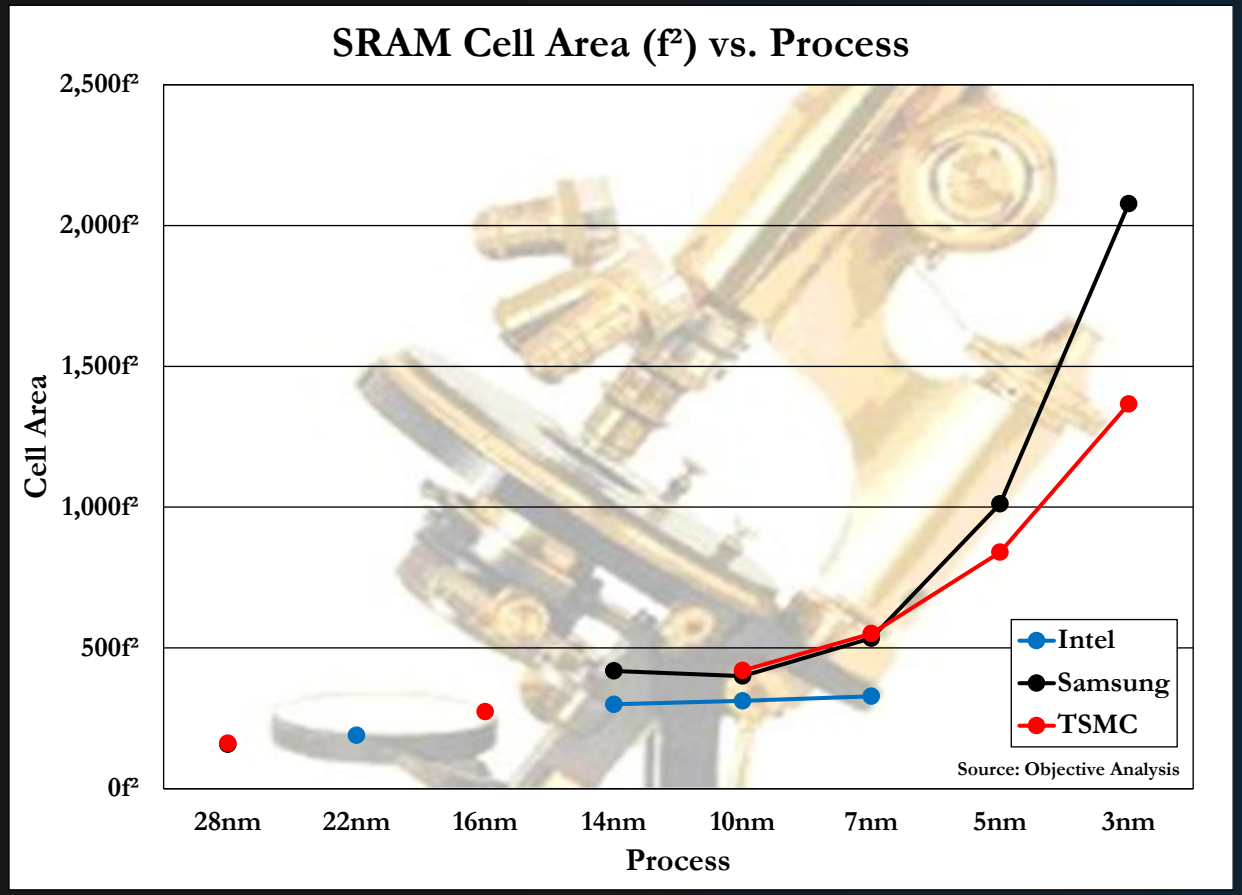

From: Emerging Memories Branch Out

- SRAM doesn't scale with logic process
  - Cost increases with smaller geometries
- Emerging memories can solve this problem
- Future caches will use emerging memories
  - Larger capacities
  - Cheaper
  - Persistent

SDC 23

Chiplet Memory Can Be Persistent

- Persistent code and data memory, and even caches
- Software will need to catch up
  - The SNIA NVM Programming Model is the basis for this
- Security concerns
  - What if the persistent cache chip falls into the wrong hands?
  - Should cache lines be erased when invalidated?
  - Should all memory communications and NVM data at rest be encrypted?

Future Thoughts

Emerging Memory is Falling Into Place

- Leading-edge processes can’t use NOR
  - And SRAM is growing unattractive!
- Already some use in the enterprise
- Growing adoption in consumer applications
- Increased consumption will reduce prices
  - The economies of scale will accelerate emerging memory penetration
- Plus, they offer technical benefits
  - Fast
  - Very low power
  - Less messy than flash

Emerging Memories Are Around the Corner

MRAM

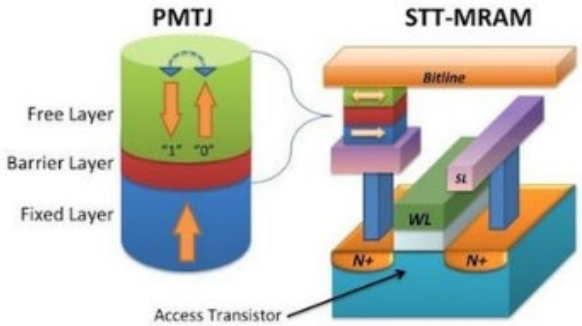

ReRAM

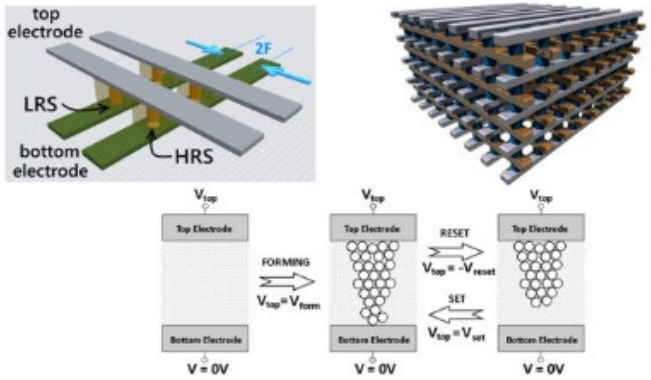

PCM

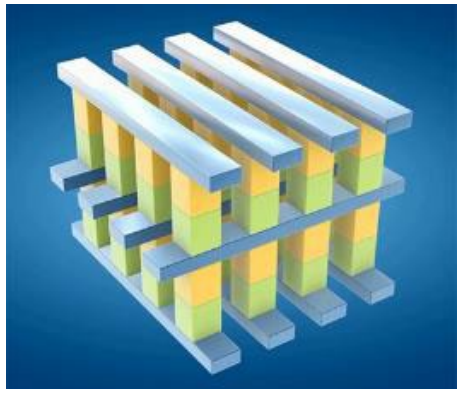

FRAM

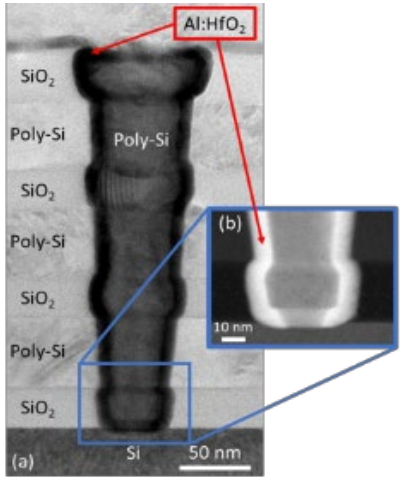

All New Memories Share Some Attributes

- Small single-element cell
  - Supports small/inexpensive die and 3D stacking
  - Promises to scale past DRAM & NAND flash
- Write in place
  - No “Block Erase”
  - More symmetrical read/write speeds
- Nonvolatile/Persistent
  - These can all be used as Persistent Memory: “PM”

The Future of Emerging Memories

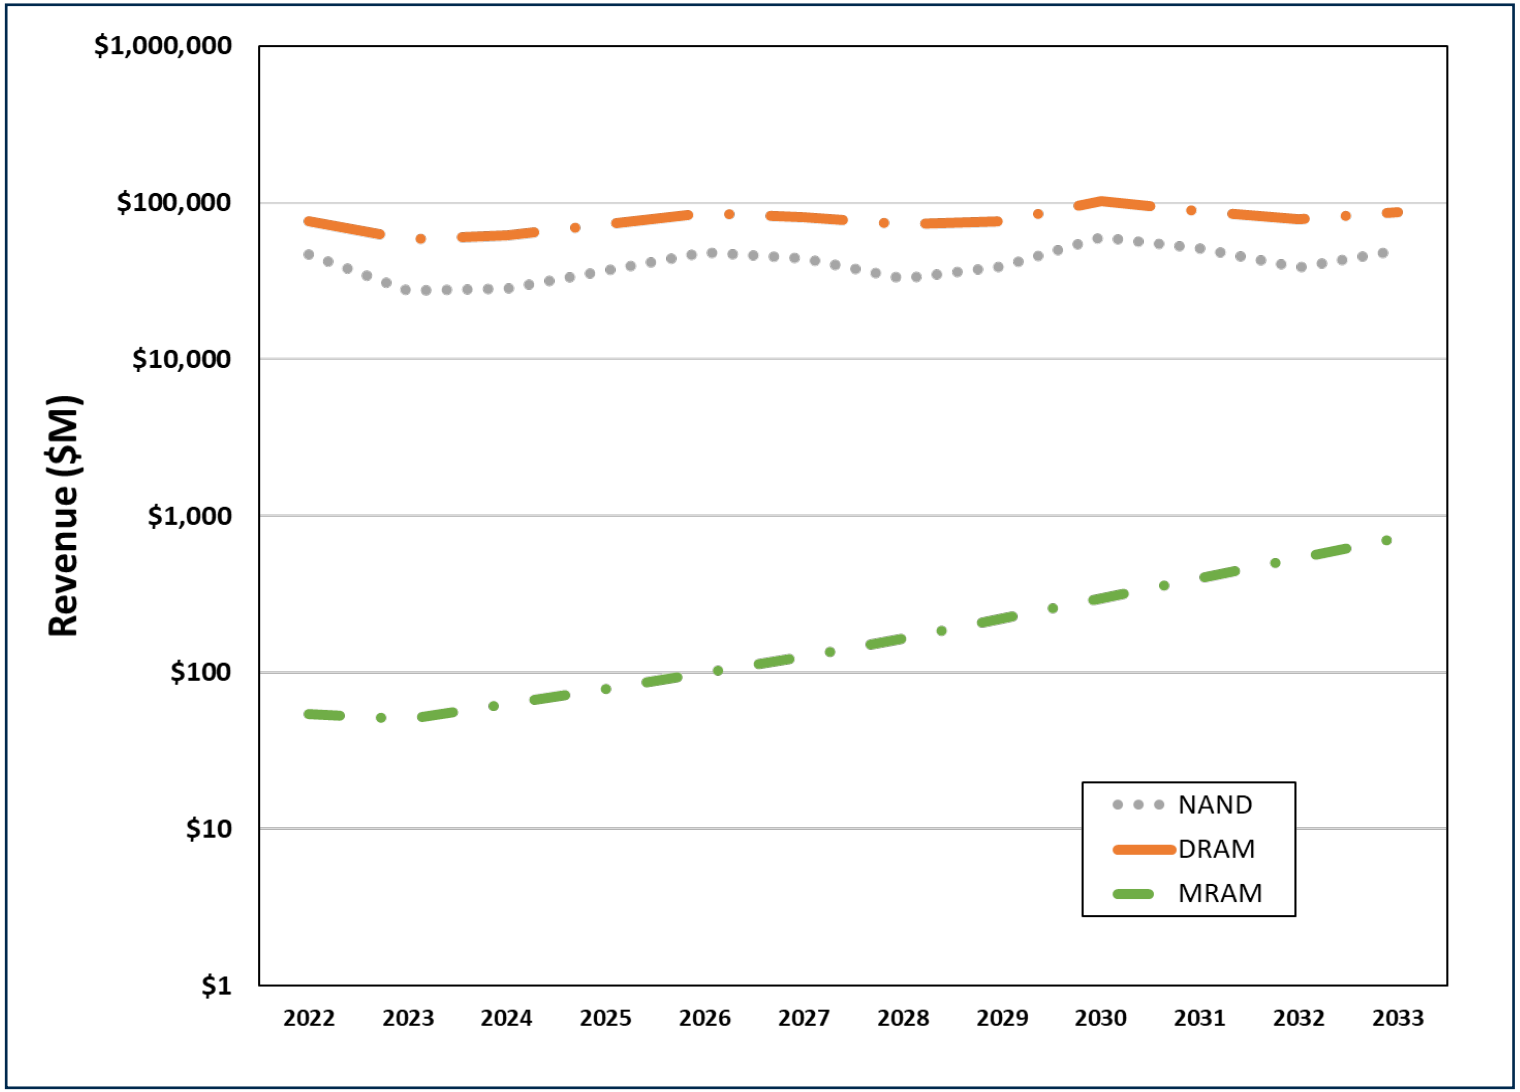

Emerging memory revenue forecast to grow significantly faster than DRAM or NAND flash

From: [Emerging Memories Branch Out]

New Report: Emerging Memories Branch Out

OBJECTIVE ANALYSIS
Semiconductor Market Research

Coughlin Associates

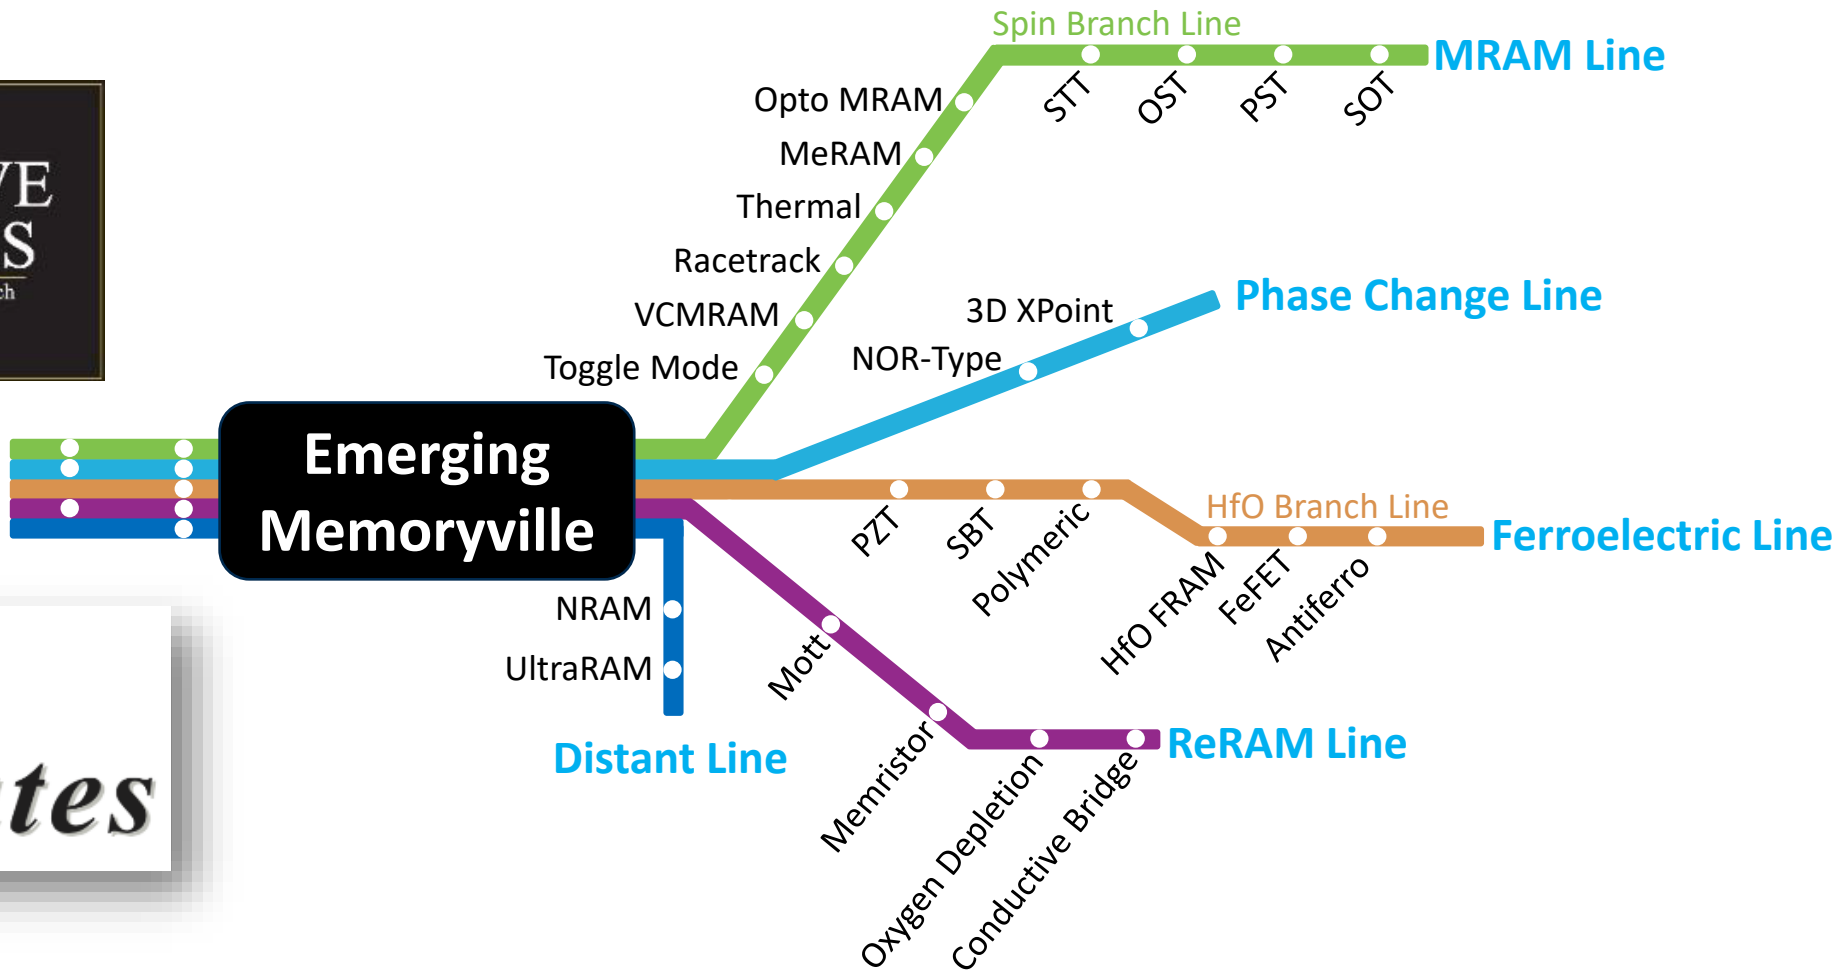

Now Available!

https://Objective-Analysis.com/reports/#Emerging

http://www.TomCoughlin.com/techpapers.htm

SDC 23

Summary

- Optane's short life founded a great legacy
  - New computing architectures and programming models
  - Many alternatives for the Optane user
- CXL has opened the door to new memory architectures
  - Processors no longer tied down to one interface, one memory type
- UCIe makes CXL's strengths available to chiplets
  - Chiplets are the path to future processors
- Emerging memories are poised to solve tomorrow's problems

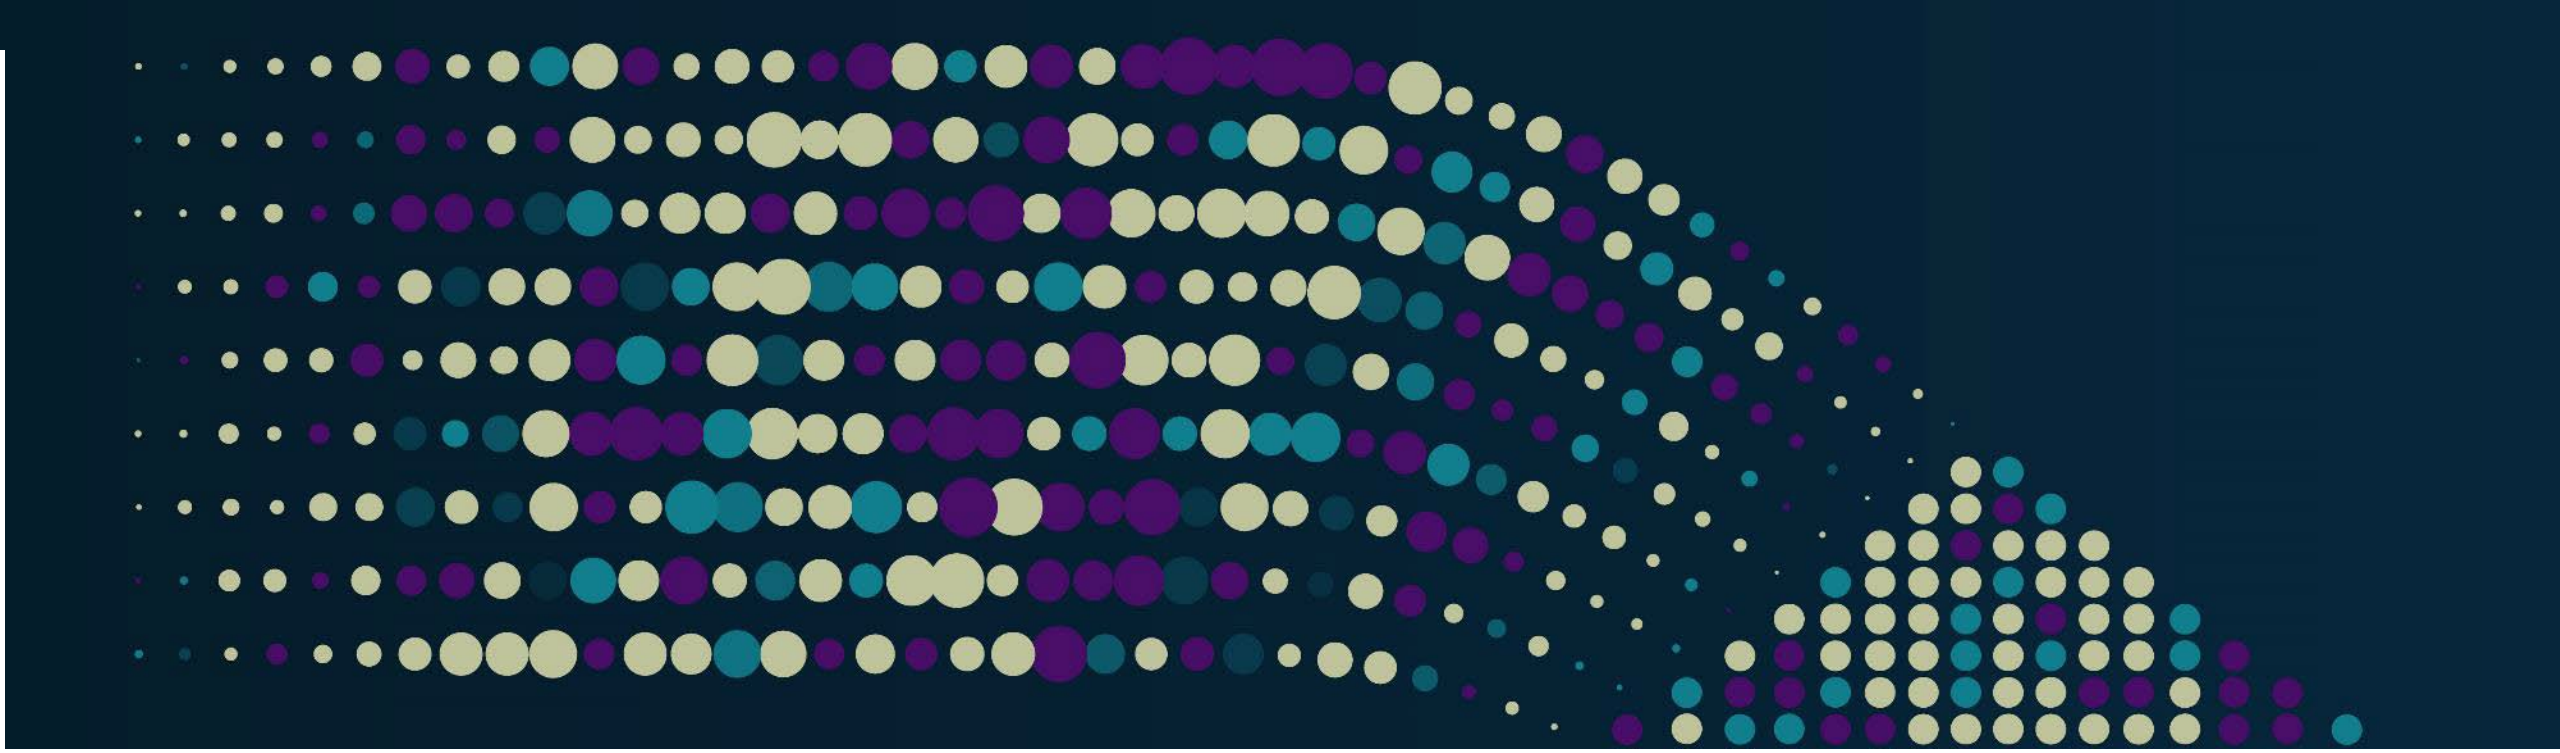

Please take a moment to rate this session.

Your feedback is important to us.

