SNIA | Compute, Memory, and Storage

New Memories for Inference & General-Purpose Compute

Live Webinar

January 20, 2026

12:00 Noon PT/ 3:00 pm ET

Arthur Sainio,
SNIA PM SIG
Chair

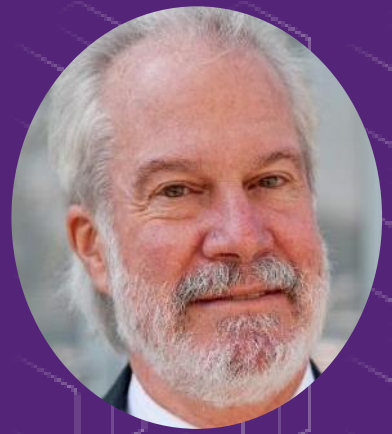

Jim Handy,
Objective
Analysis

Tom Coughlin,
Coughlin
Associates

Introductions

Arthur Sainio
Moderator
Co-Chair,
SNIA Persistent Memory
Special Interest Group

Tom Coughlin
Presenter
President, Coughlin Associates
Past President, IEEE

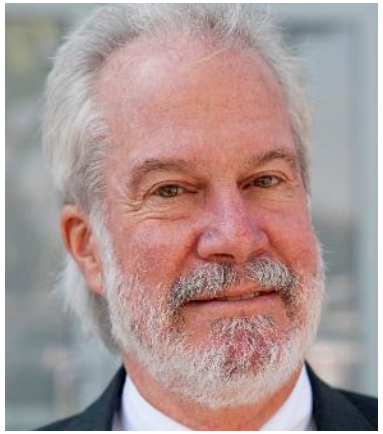

Jim Handy
Presenter
President,
Objective Analysis

SNIA | CMS A SNIA Community

The SNIA Community

200
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

SNIA | CMS A SNIA Community

SNIA Legal Notice

- The material contained in this presentation is copyrighted by SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

The industry leading member companies of the SNIA Compute, Memory, and Storage Community support the industry drive to combine processing with memory and storage, and to create new compute architectures and software to analyze and exploit the explosion of data creation over the next decade.

What We Do

Engage and Educate

- Computational Storage
- Persistent Memory
- PM and SSD Performance
- Emerging Memories
- Smart Data Accelerator
- Solid State Drives
- Solid State Systems
- SSD Form Factors

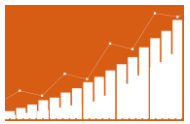

Accelerate Standards

- Computational Storage Architecture & Programming Model
- Computational Storage API
- NVM Programming Model
- Smart Data Accelerator Interface
- Solid State Storage Performance Test Specifications
- SSD Form Factor Specifications

How We Do It

Propel Technology Adoption

- Persistent Memory Programming Workshops
- SSD Form Factors Explained
- Computational Storage, Memory, and Solid State Drive Demonstrations at live and online technology events
- Interactive Webinars with Technology Industry Experts
- Videos on the SNIA Video YouTube Channel

SNIA | CMS A SNIA Community

Outline

- New Memory Technology Overview
- Why Now?
- Compute-In-Memory
- Persistence in Standard Architectures
- CXL: An Enabler
- Market sizing
- Q & A

SNIA | CMS
A SNIA Community

New Memory Technologies

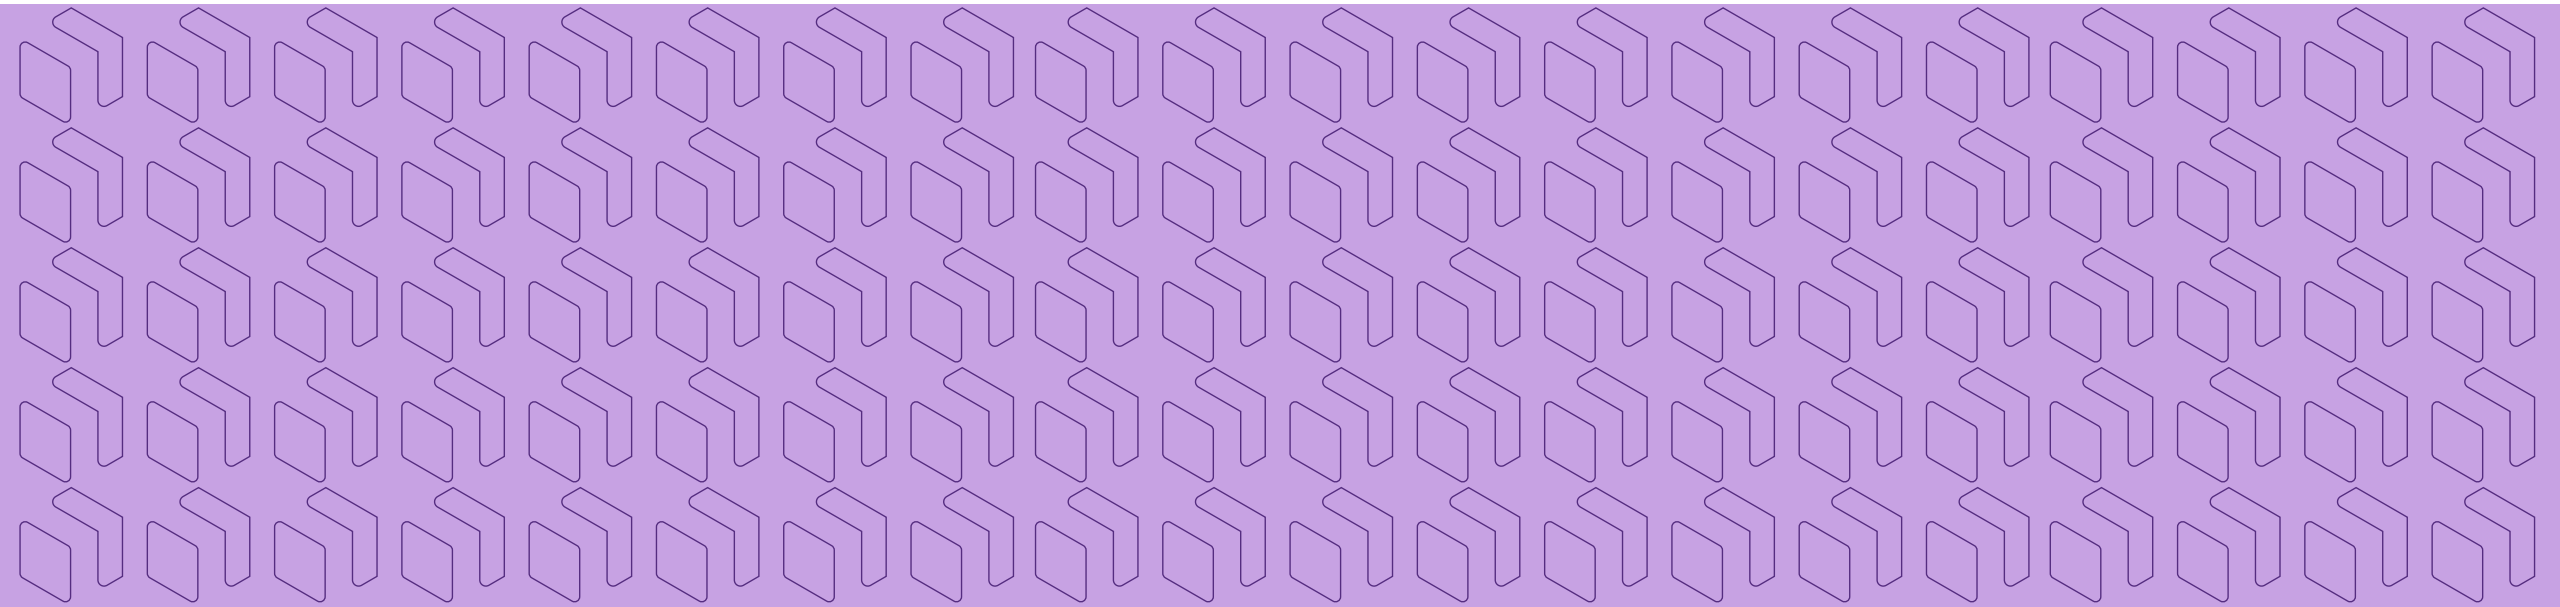

New Memory Technology Overview

- Magnetic: MRAM, STT, SoT, etc.
- Resistive: Conductive Bridge, Oxygen Vacancy
- Ferroelectric: FRAM, FeFET, Other
- Phase Change: PCM, 3D XPoint
- Other Types

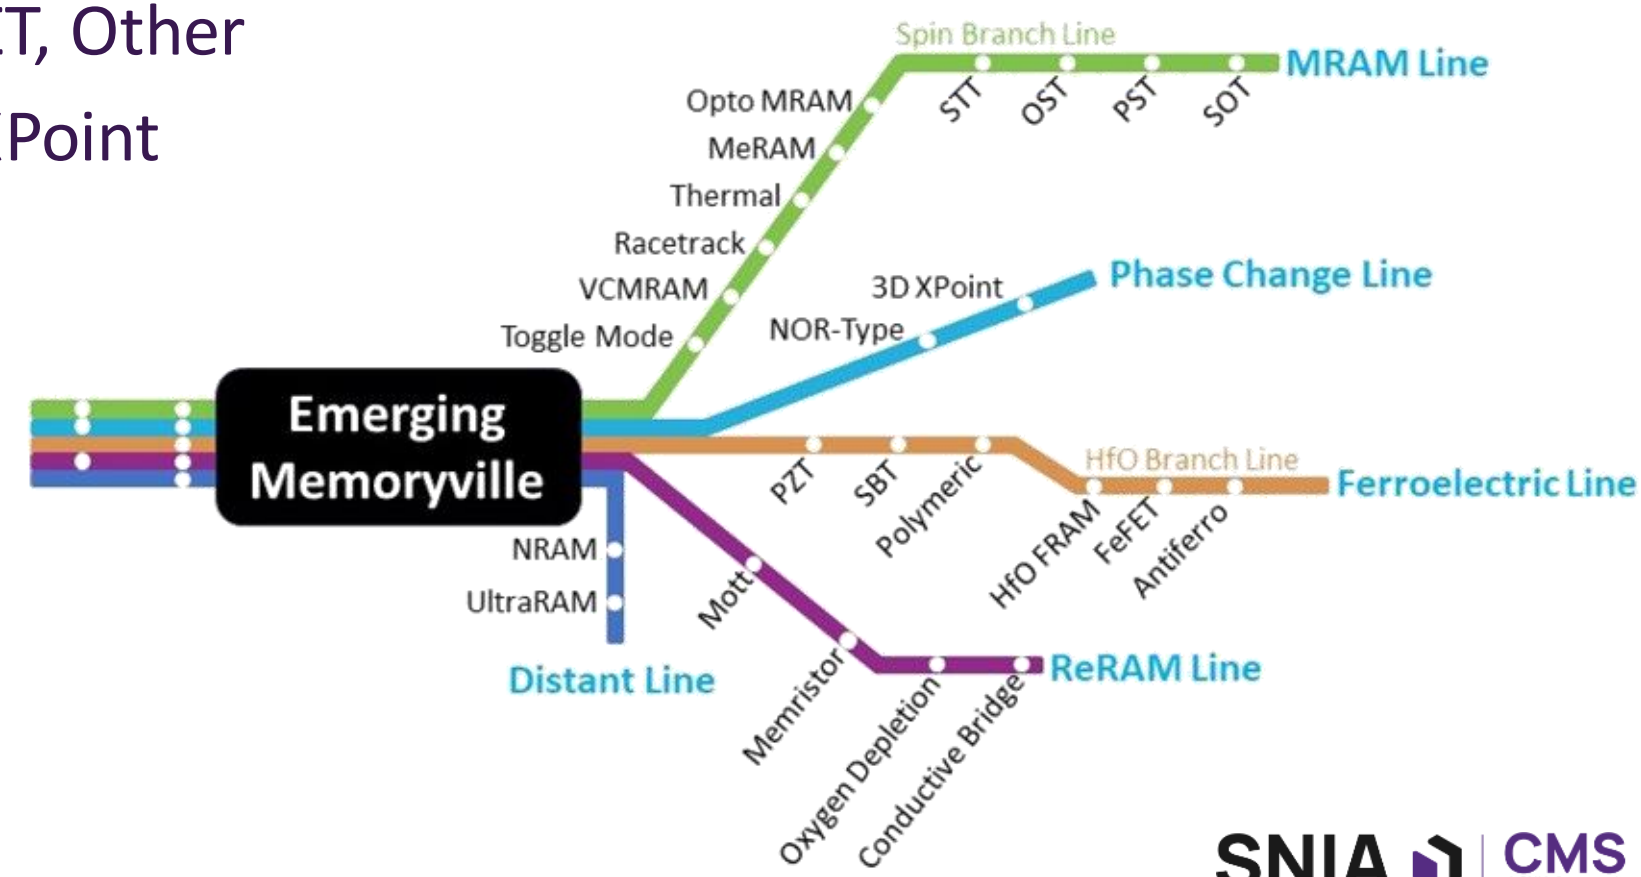

SNIA | CMS A SNIA Community

All New Memories Share Common Attributes

- 1-transistor cell
- Nonvolatile
- High-speed write
- Write-in-place
- Better scaling limits
- Improved endurance

MRAM

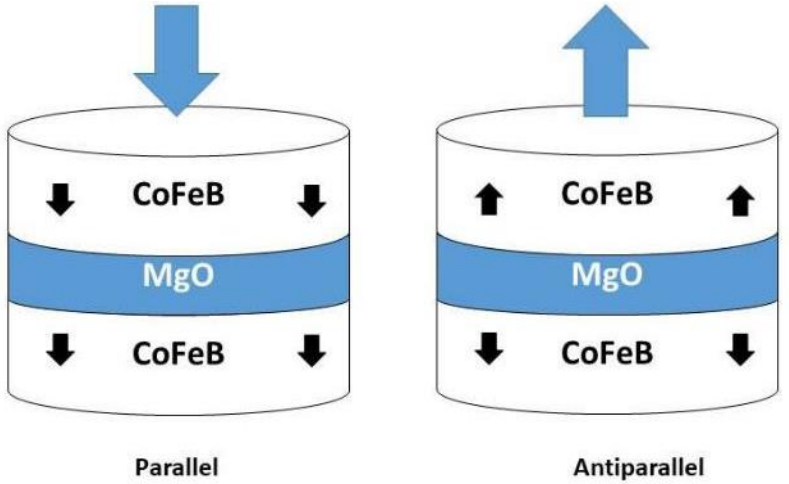

- In volume as embedded and discrete versions
  - Replacing NOR flash in embedded applications
- **Strengths**: Fast, low power, relatively small cell size, w range, radiation-tolerant
- **Weaknesses**: Requires new materials (well understood), large program currents, magnetic sensitivity, MLC difficult
- **Variants**: In-Plane Toggle, Perpendicular, STT, SOT, MeRAM…
- **Suppliers/Researchers**: Everspin, TSMC, Samsung, GlobalFoundries, Others
- **Outlook**: Well-established, broad support, in leading position

ReRAM

- In limited production as embedded memory
- **Strengths**: Can be made without new materials, radiation-tolerant,
- **Weaknesses**: Selector challenge,
- **Variants**: Oxygen Vacancy, Conductive Bridge, (Memristor)
- **Suppliers/Researchers**: Weebit, Fujitsu, Nuvoton, Renesas, TSMC, UMC, Winbond, Others
- **Outlook**: Has the potential to rise above others in the long run.

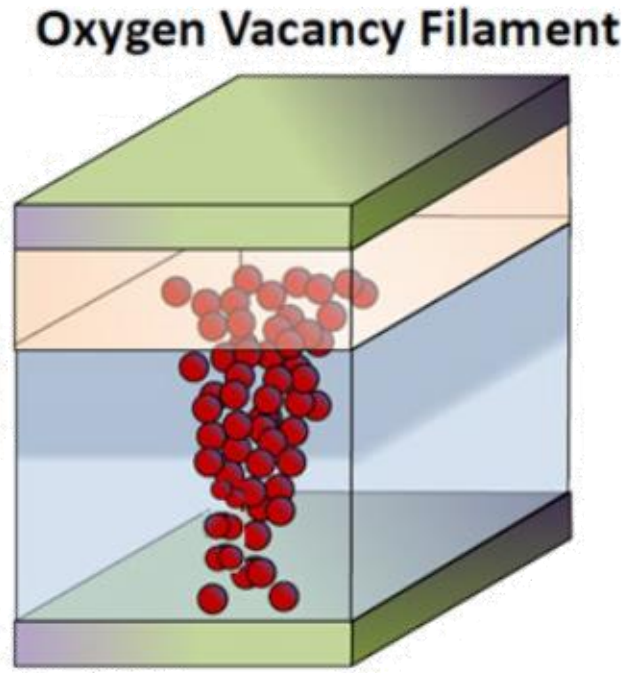

FRAM

- Very high production volume in some forms, but low wafer run rate
- **Strengths**: Extremely low write energy, high speed, 10^{15} endurance, could use volume HfO process
- **Weaknesses**: Current volume process doesn't shrink & attacks Silicon, new process struggling to reach production
- **Variants**: PZT, SBT, HfO
- **Suppliers/Researchers**: Infineon, TI, IBM, FMC, Others
- **Outlook**: A production process breakthrough could change the market

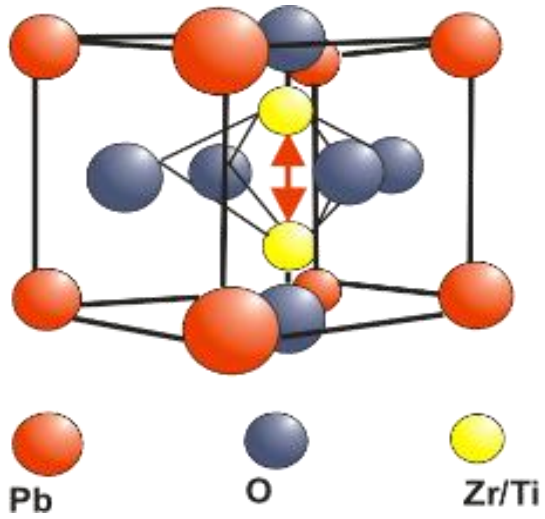

PCM

- One embedded processor and a failed high-volume product
- **Strengths**: Well-understood, matches selector material, small cell
- **Weaknesses**: High-energy write, challenging materials, thermal sensitivity
- **Variants**: Thermal vs. bulk change
- **Suppliers/Researchers**: Intel, Micron, Sandisk, IBM, others
- **Outlook**: No major commercialization efforts
  - Intel's learning not being used

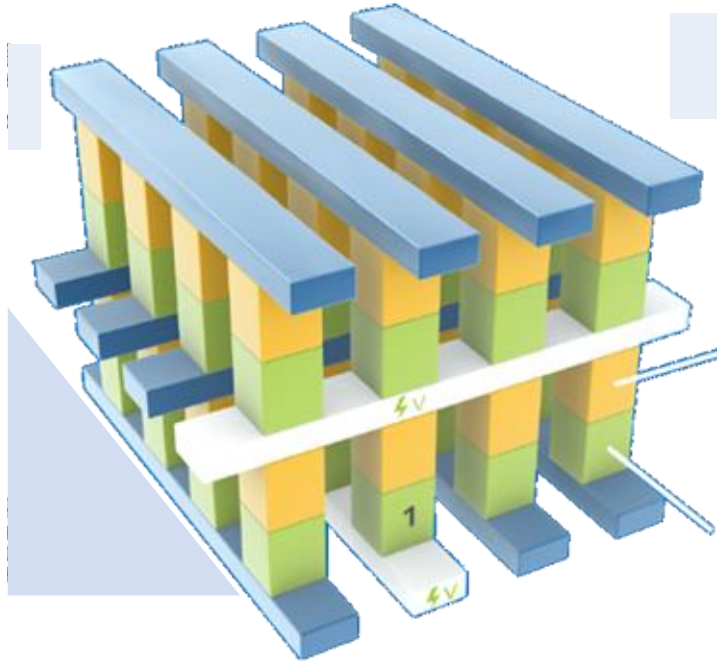

No Shortage of New Ideas

- Resonant Tunneling
- Mott Memories
- Nanotubes
- Quantum dots
- Memcapacitor
- Meminductor
- Polymeric
- Electrochemical
- Selector-Only Memory
- DNA

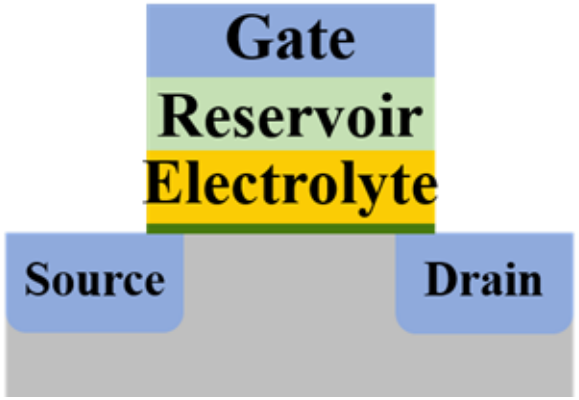

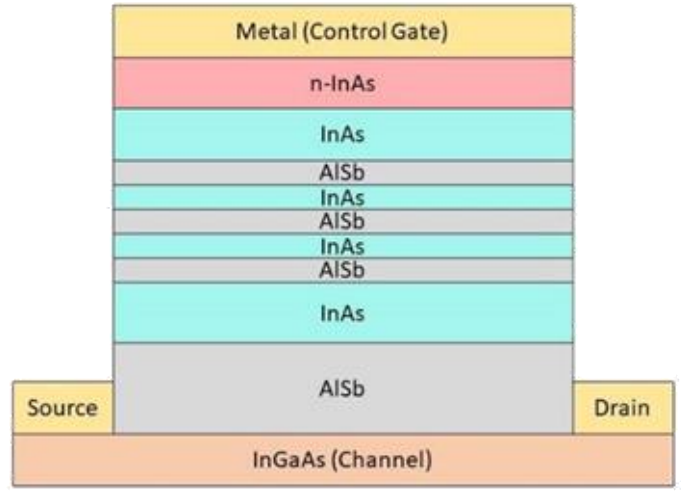

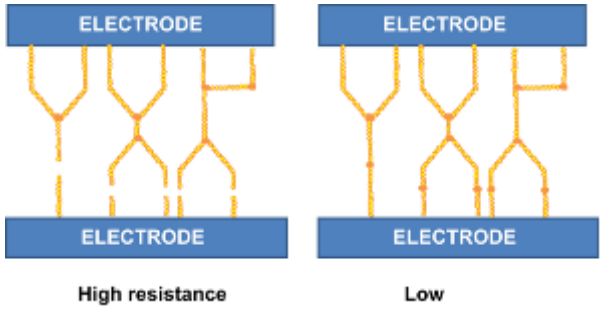

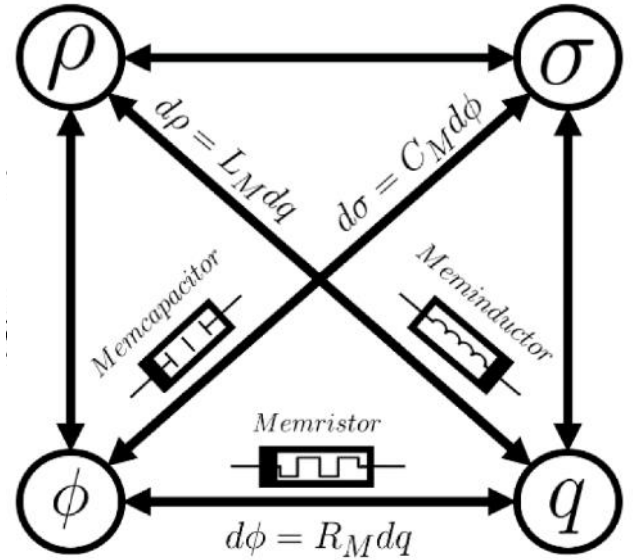

SNIA CMS A SNIA Community

SNIA | CMS
A SNIA Community

Why Now?

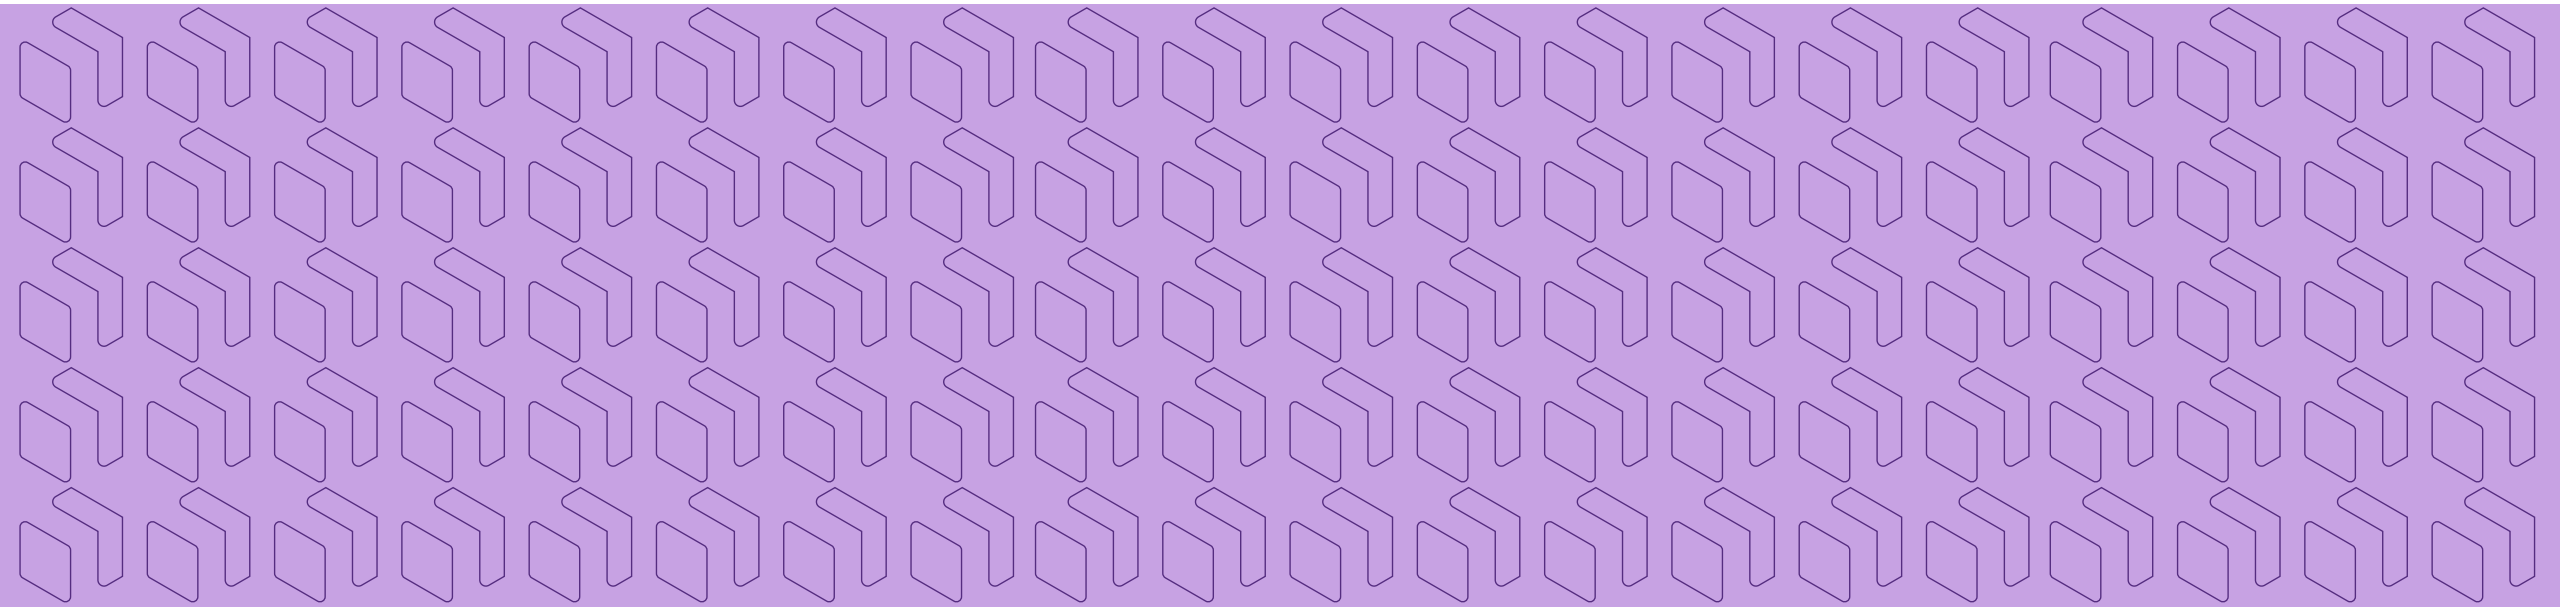

New Memory Growth Drivers

- End of NOR scaling
  - Embedded memories require a new technology after 28nm
  - Lack of NOR is limiting to MCUs, ASICs, etc.
- SRAM scaling slowdown
  - Lithographic scaling benefits are being limited by SRAM
  - This weakens cost improvements for MPUs, FPGAs, & other SoCs
- DRAM capacitor issues
  - DRAM capacitors went 3D in the 1990s. Few remaining shrink options
  - The largest memory market is facing trouble

SRAM No Longer Scales with Process

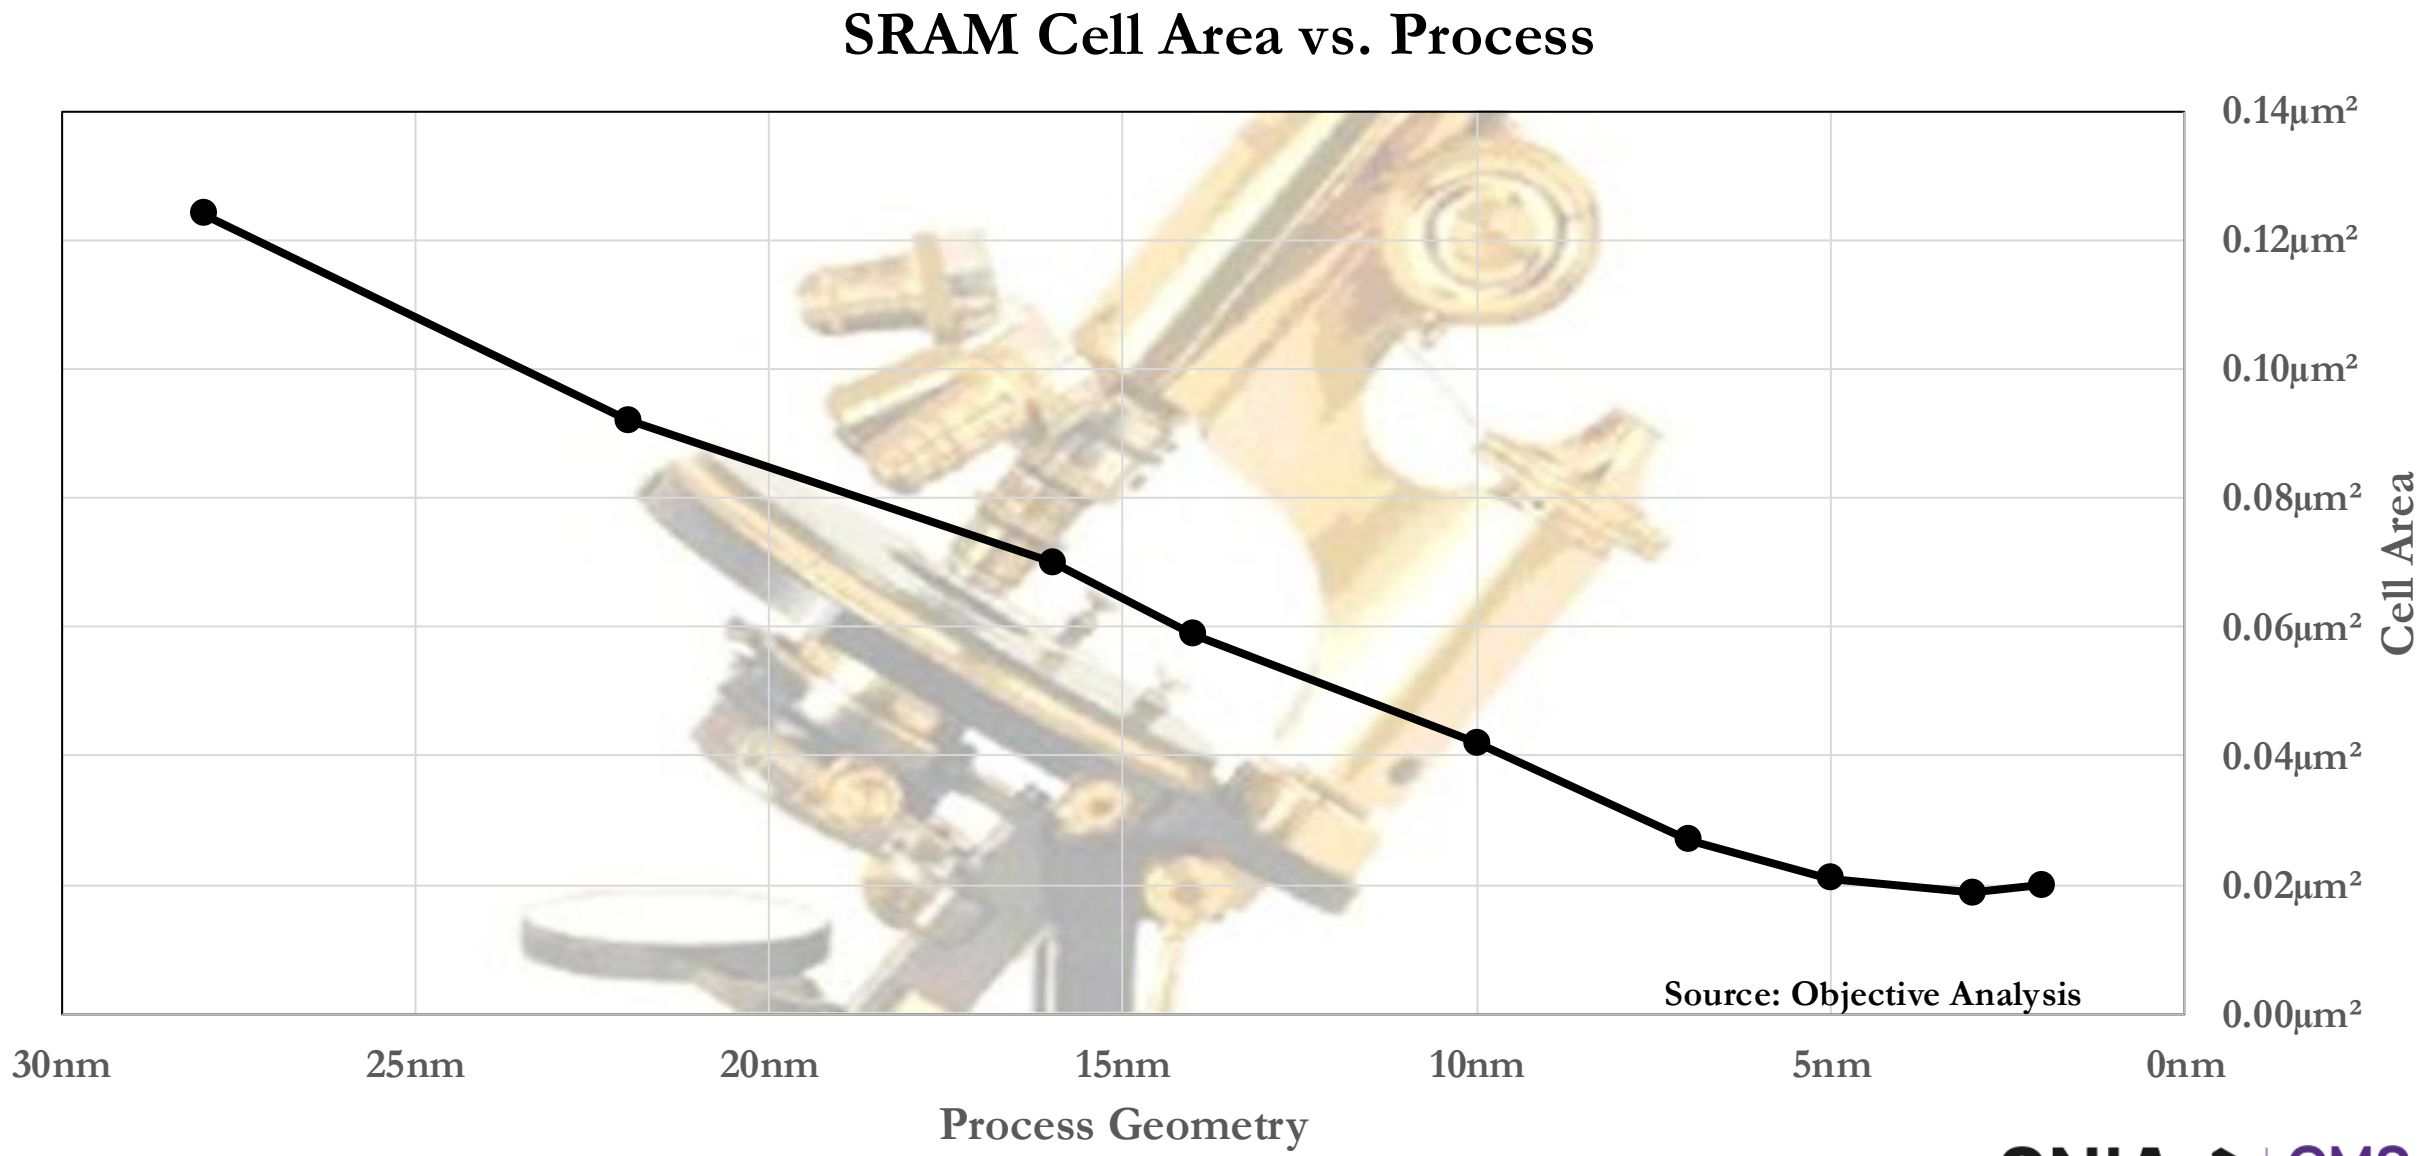

SNIA | CMS A SNIA Community

Why Chiplets Will Use New Memories

- Embedded NOR has stopped shrinking
- Embedded SRAM is running into trouble
- New Memories use different processes
  - These processes can shrink in step with logic shrinks
- Embedded SRAM is the most expensive SRAM on the planet!
  - But it's fast
    - But fast can be traded off against big
    - A big slow off-chip memory chiplet can be <u>really cheap</u>

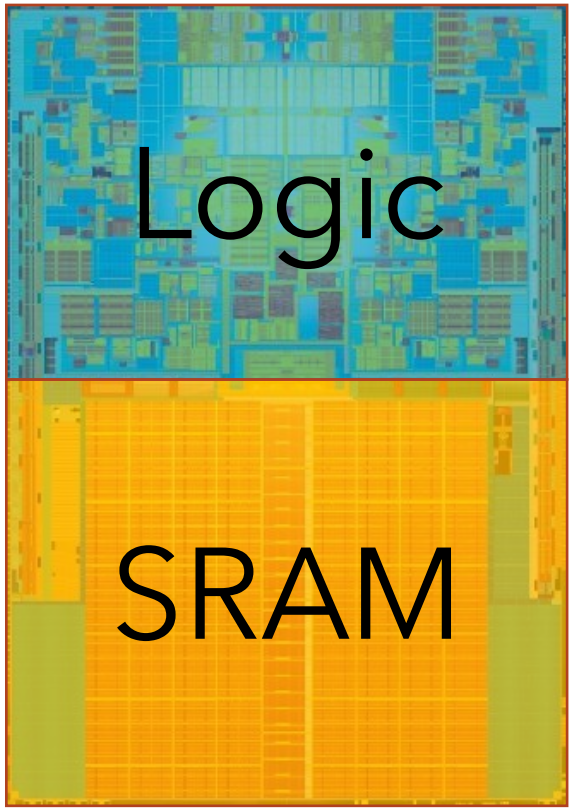

- **<u>Remove the SRAM to cut this chip's cost >50%!</u>**

SNIA | CMS
A SNIA Community

Compute-In-Memory

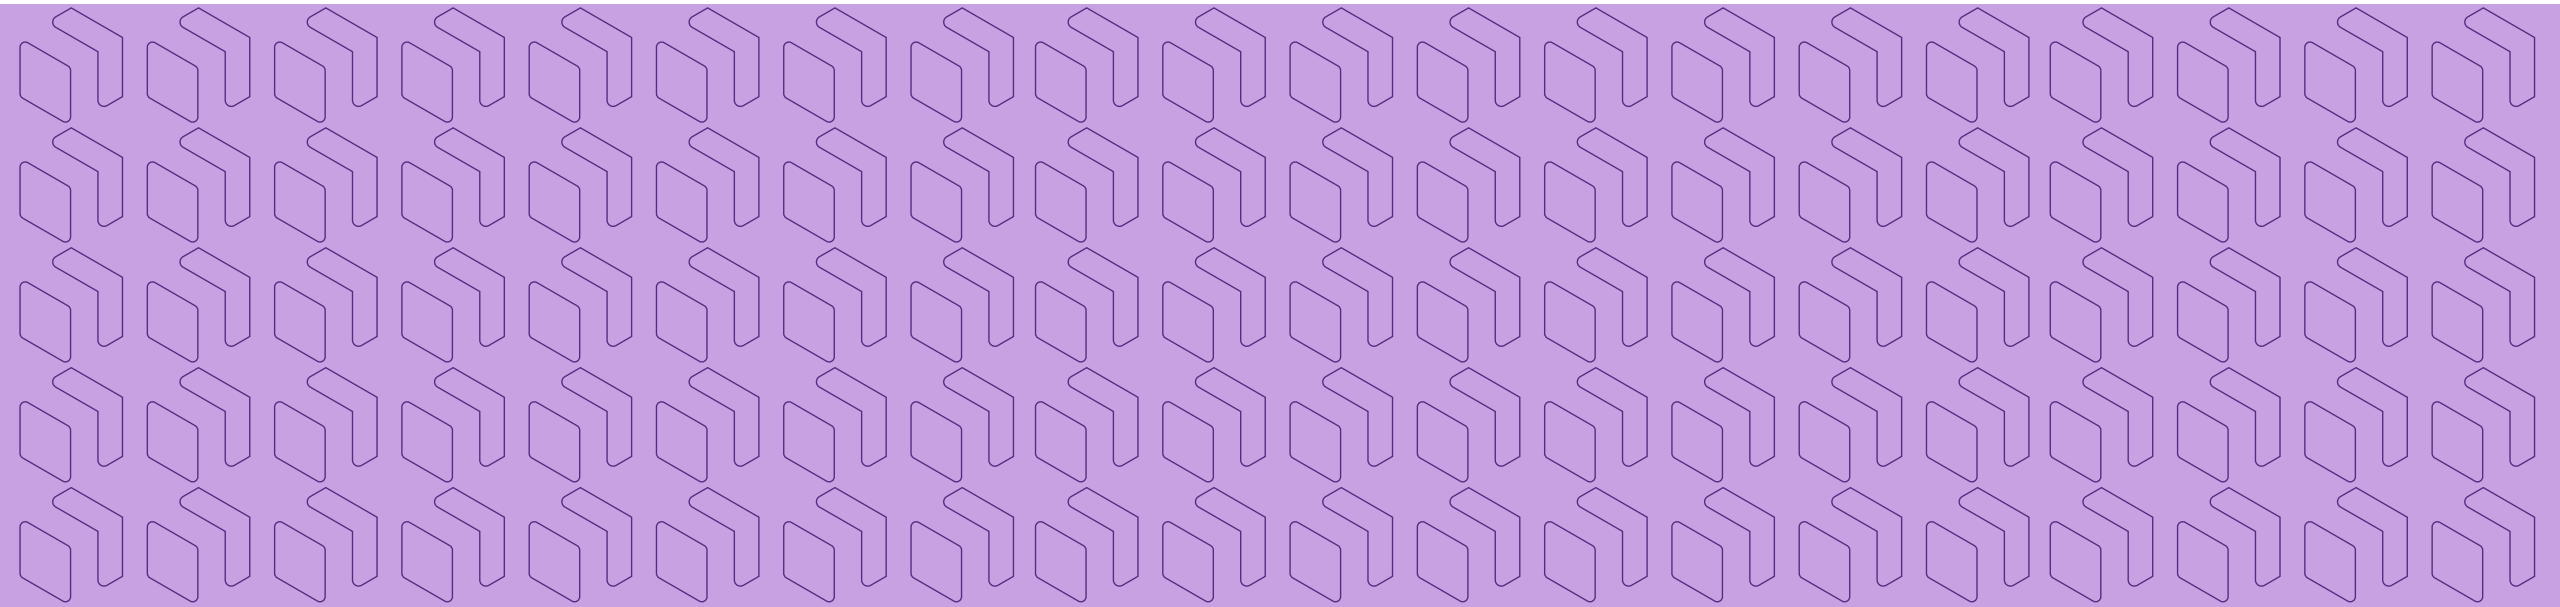

Compute-In-Memory Types

- Digital
  - Compute “In” Memory (at the bit)
    - TCAM is one example
  - Compute “Near” Memory (outside the array)
- Analog
  - Analog Neural Nets (ANN)
- Lots of names
  - Compute-In-Memory (CIM)
  - In-Memory Compute
  - Processor-In-Memory (PIM)

Guess What? Standard CPUs are CIM!

- Processor with on-chip SRAM
- SRAM with on-chip processor

- Very high speed internal data movement

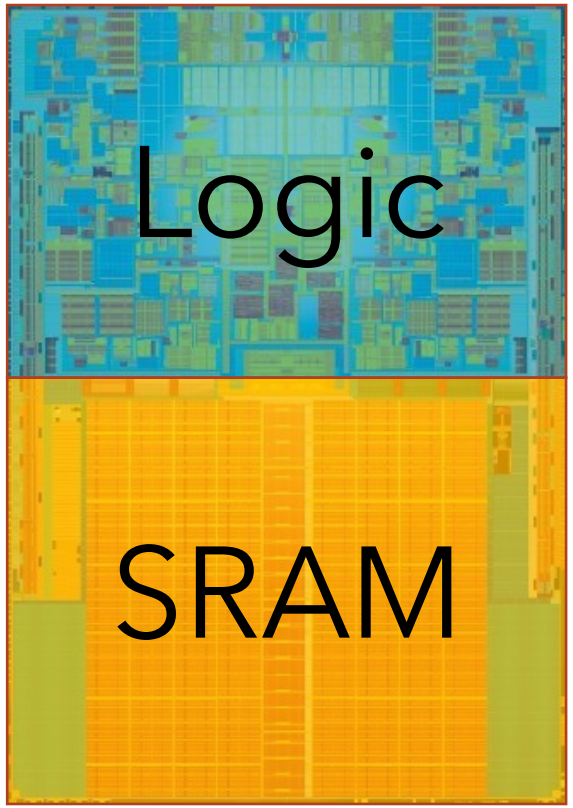

Analog Neural Networks

- Neural nets are great for endpoint applications
  - Low energy consumption
  - “Fast enough” performance
  - Low die complexity ► low cost
- There are drawbacks
  - No clear leadership or standards
  - Limited functionality
  - Very little software support
- Digital neural nets should stick around

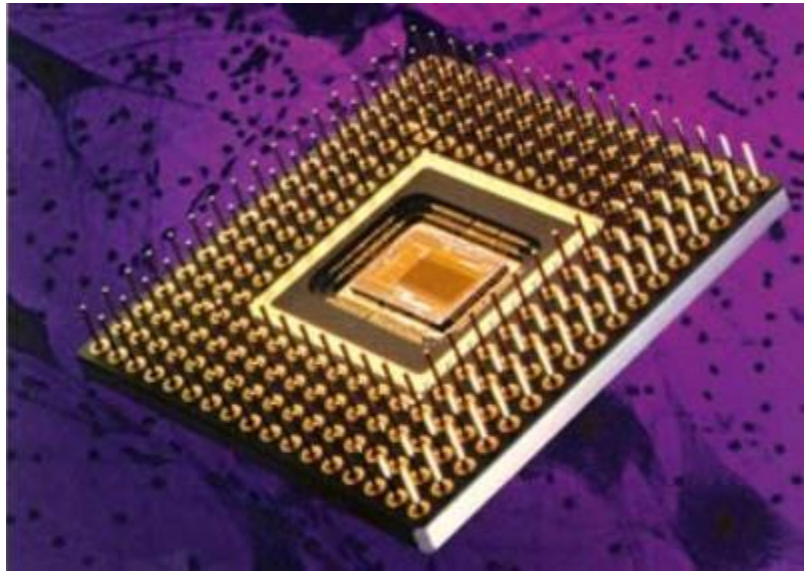

Persistence in Standard Architectures

Caches Are Headed Toward Persistence

- Use a new memory technology for cache
  - Smaller cell than SRAM
  - Continues to scale with process geometry
  - Slow and big caches can substitute for fast and small
- Put the slower caches onto chiplets
  - Cache chiplets can be made on a memory wafer
    - New memories are again a good option
  - Logic on foundry logic wafer
- Both of these solutions are nonvolatile/persistent
  - All new memory technologies scale
  - All new memory technologies are persistent

In the End, Nearly Everything Becomes Persistent

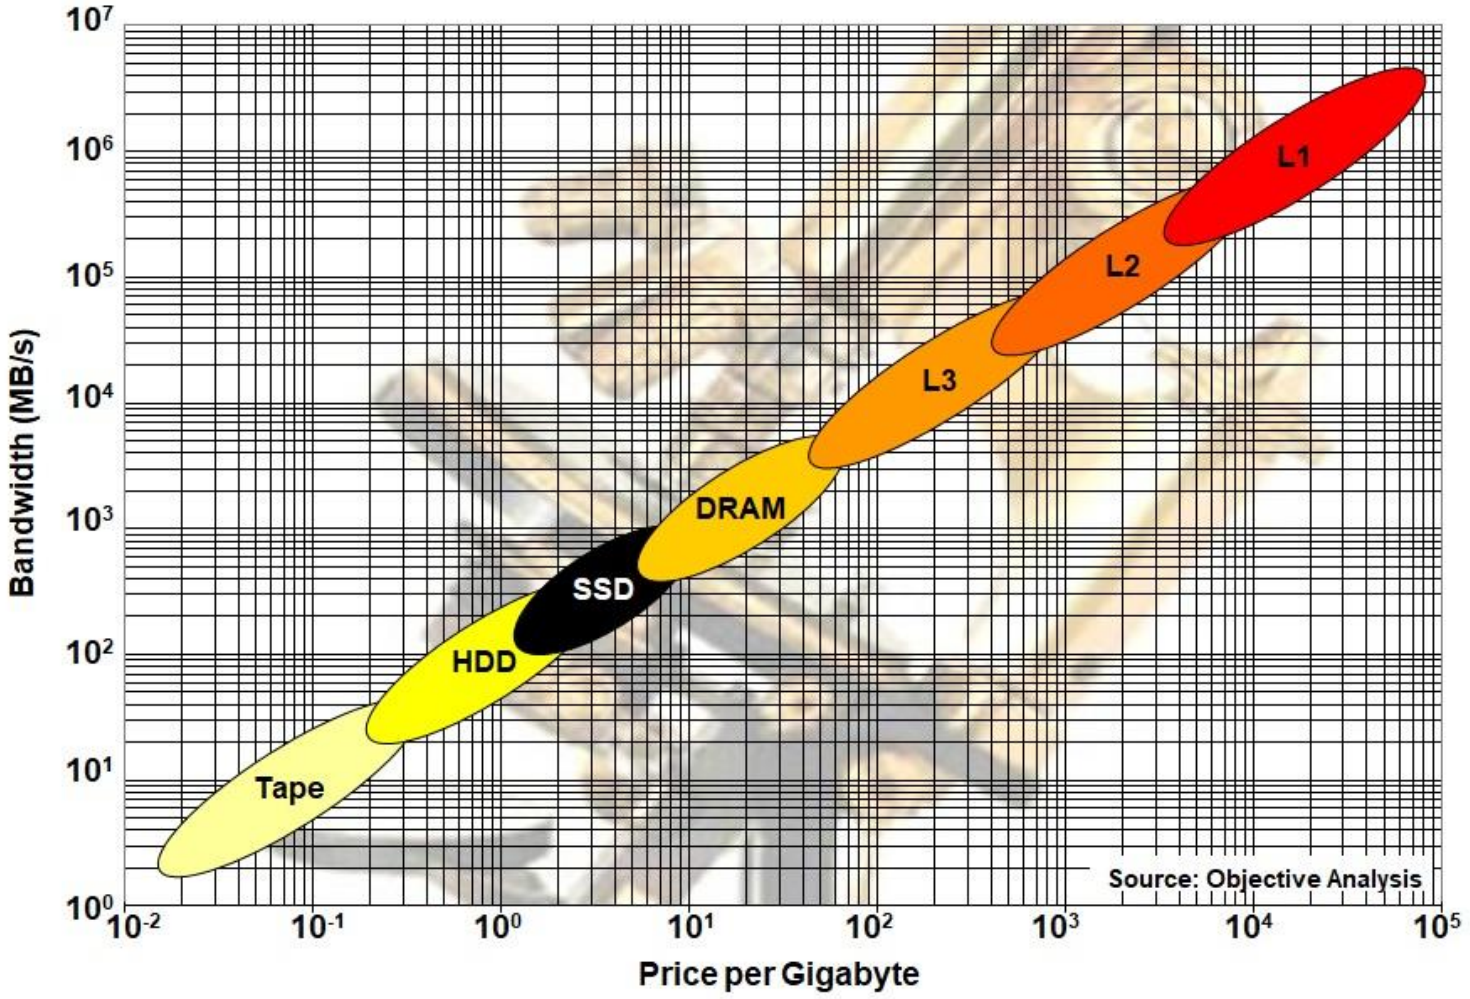

SNIA | CMS A SNIA Community

In the End, Nearly Everything Becomes Persistent

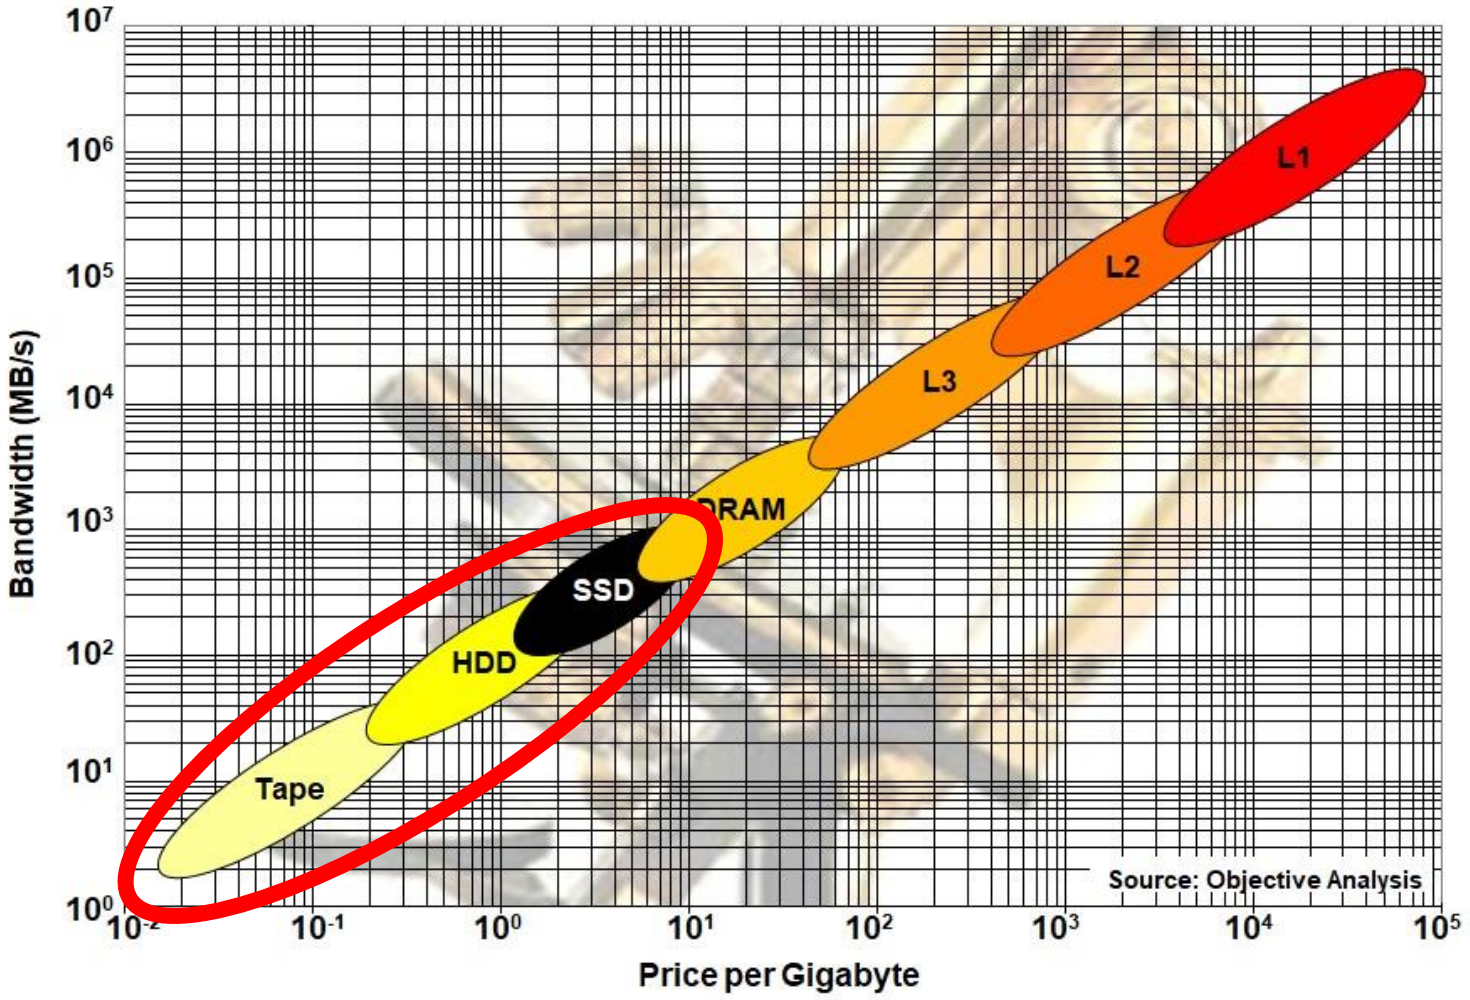

SNIA | CMS A SNIA Community

In the End, Nearly Everything Becomes Persistent

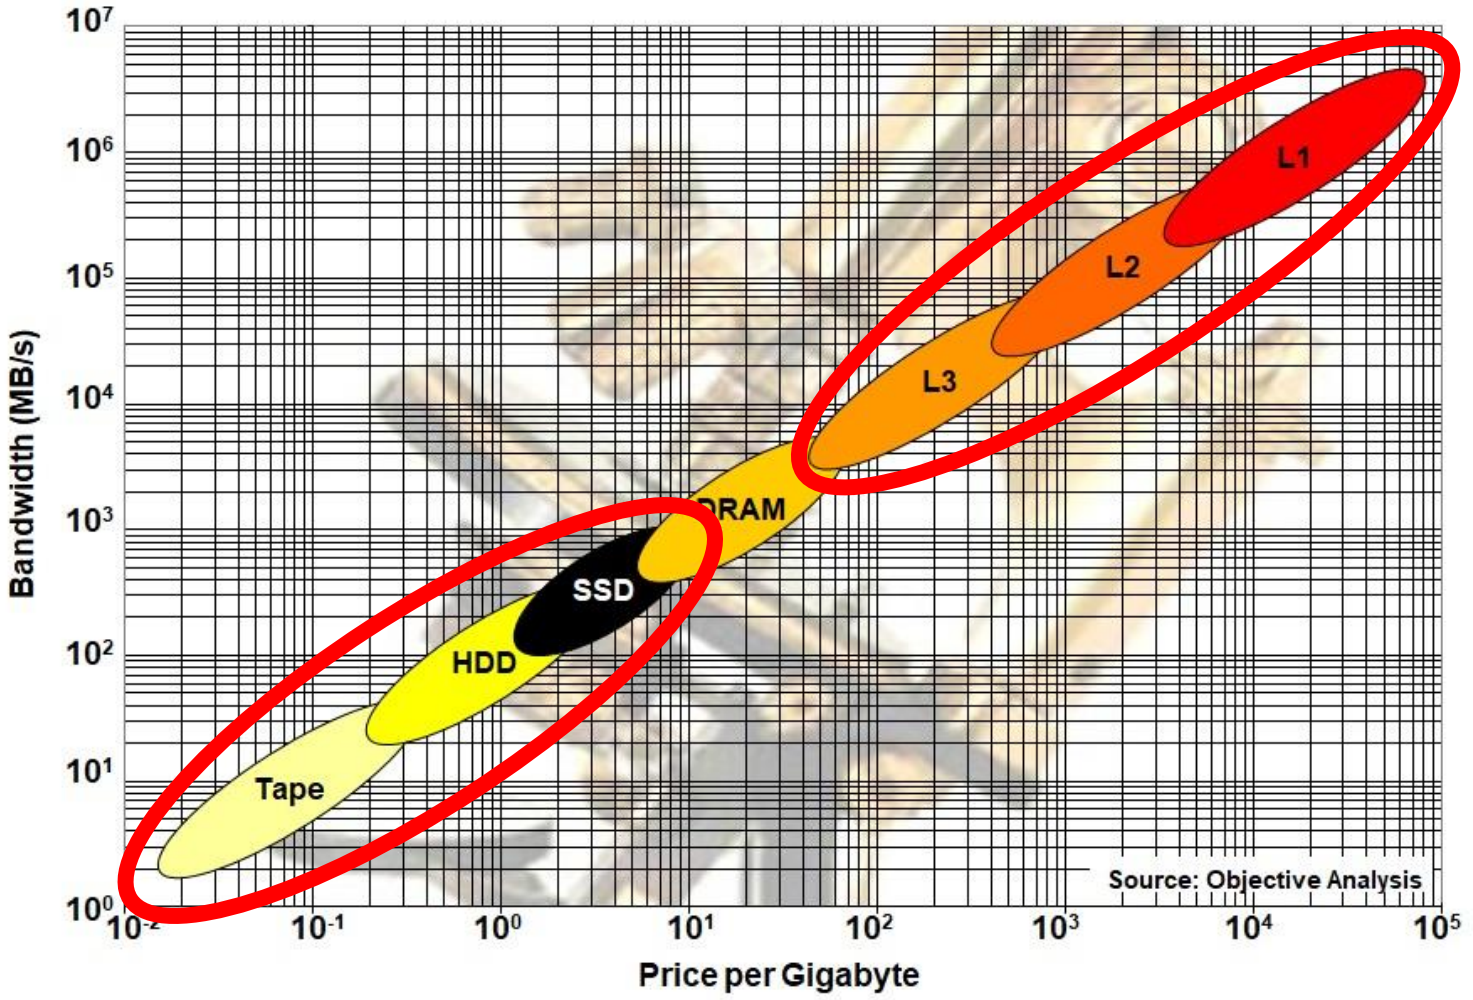

SNIA | CMS A SNIA Community

Enablers Are Falling Into Place

- SNIA Nonvolatile Memory Programming Model is here now
  - Limited application programs from Oracle & SAP
  - Designed to support future persistent architectures
- CXL supports varying memory speeds
  - Processors can communicate with DRAM, MRAM, ReRAM, etc.
- UCIe is based on CXL

SNIA | CMS
A SNIA Community

CXL: An Enabler

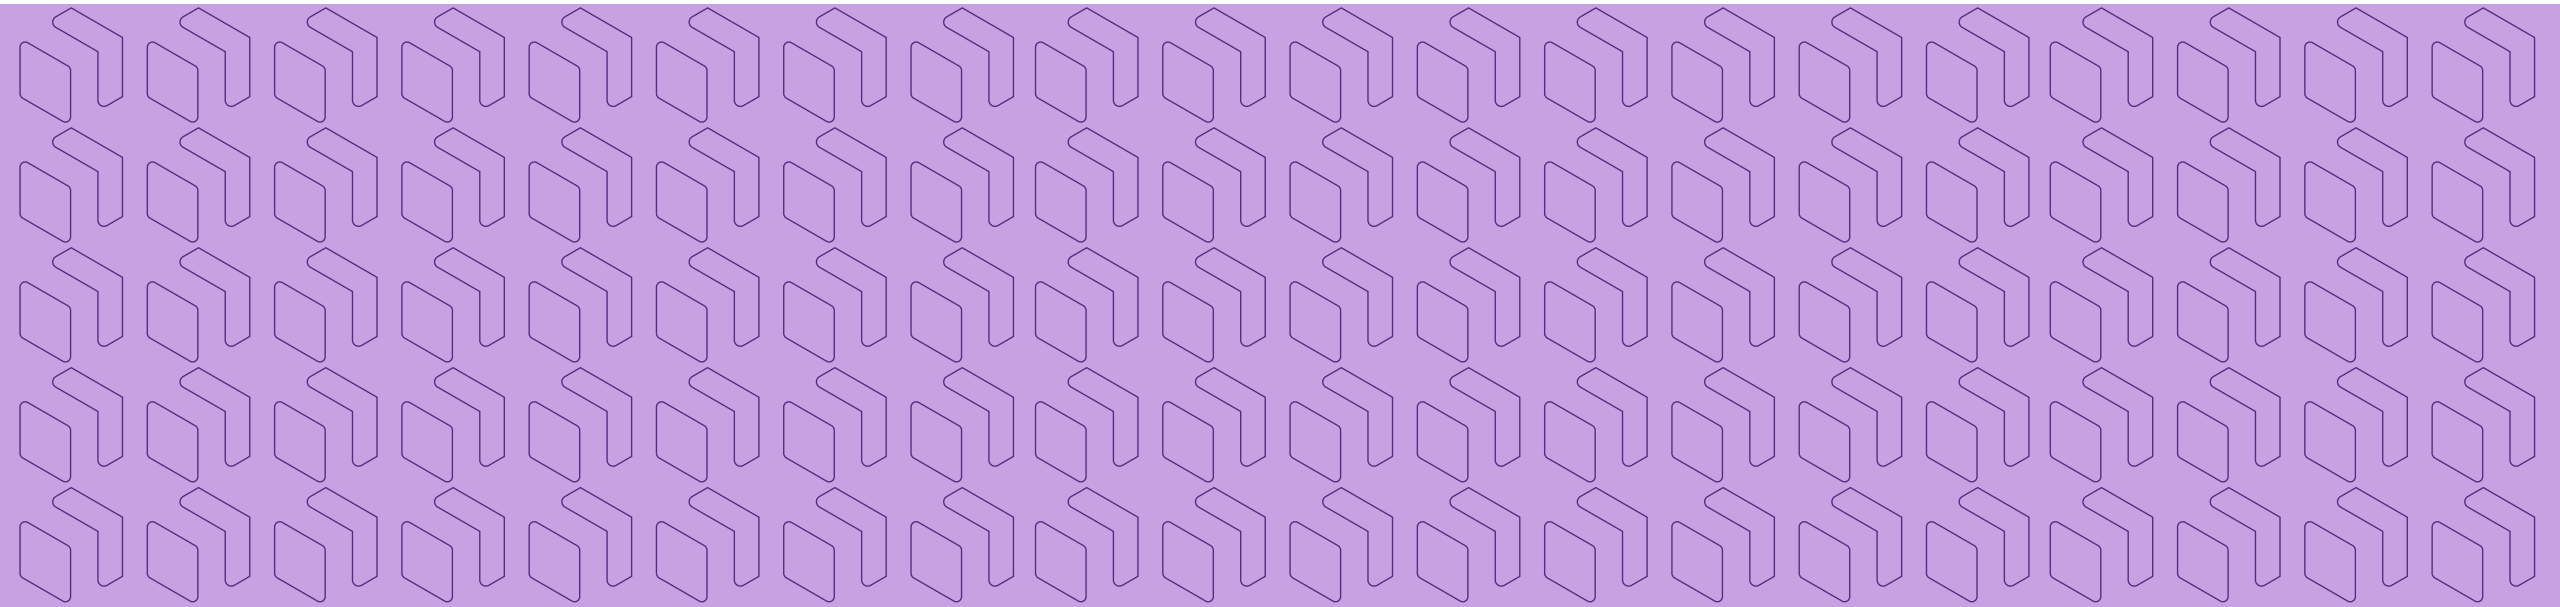

CXL Supports Any Memory, Volatile or Persistent

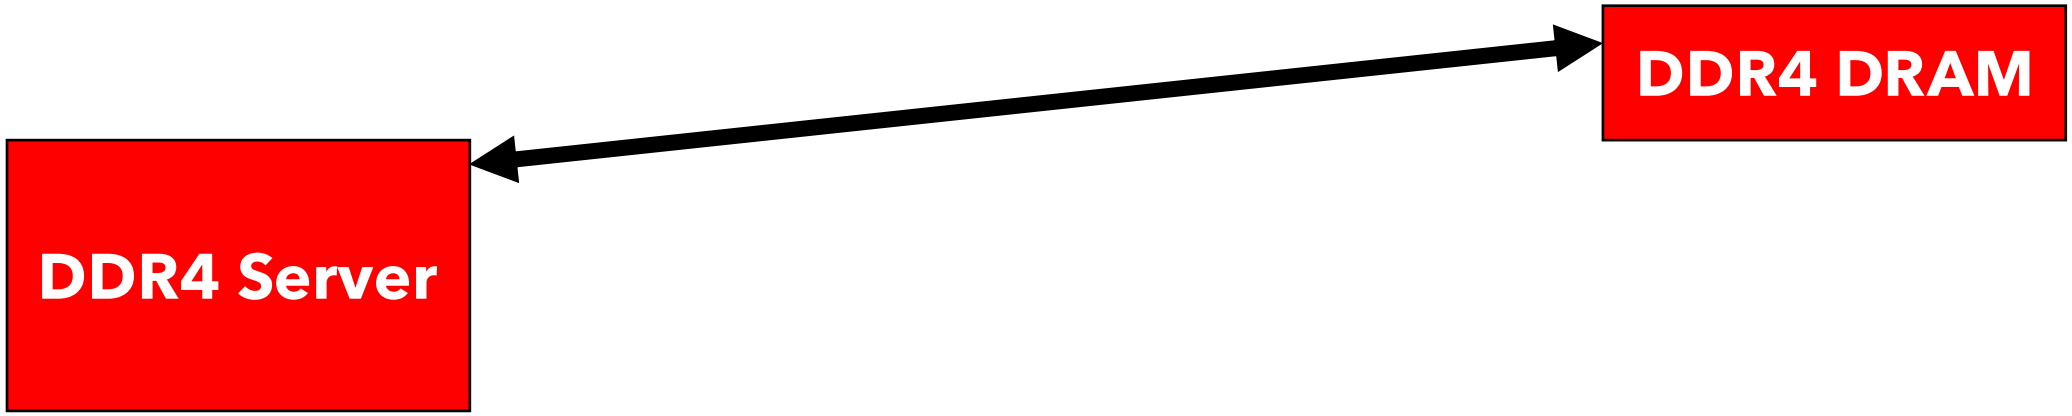

CXL Supports Any Memory, Volatile or Persistent

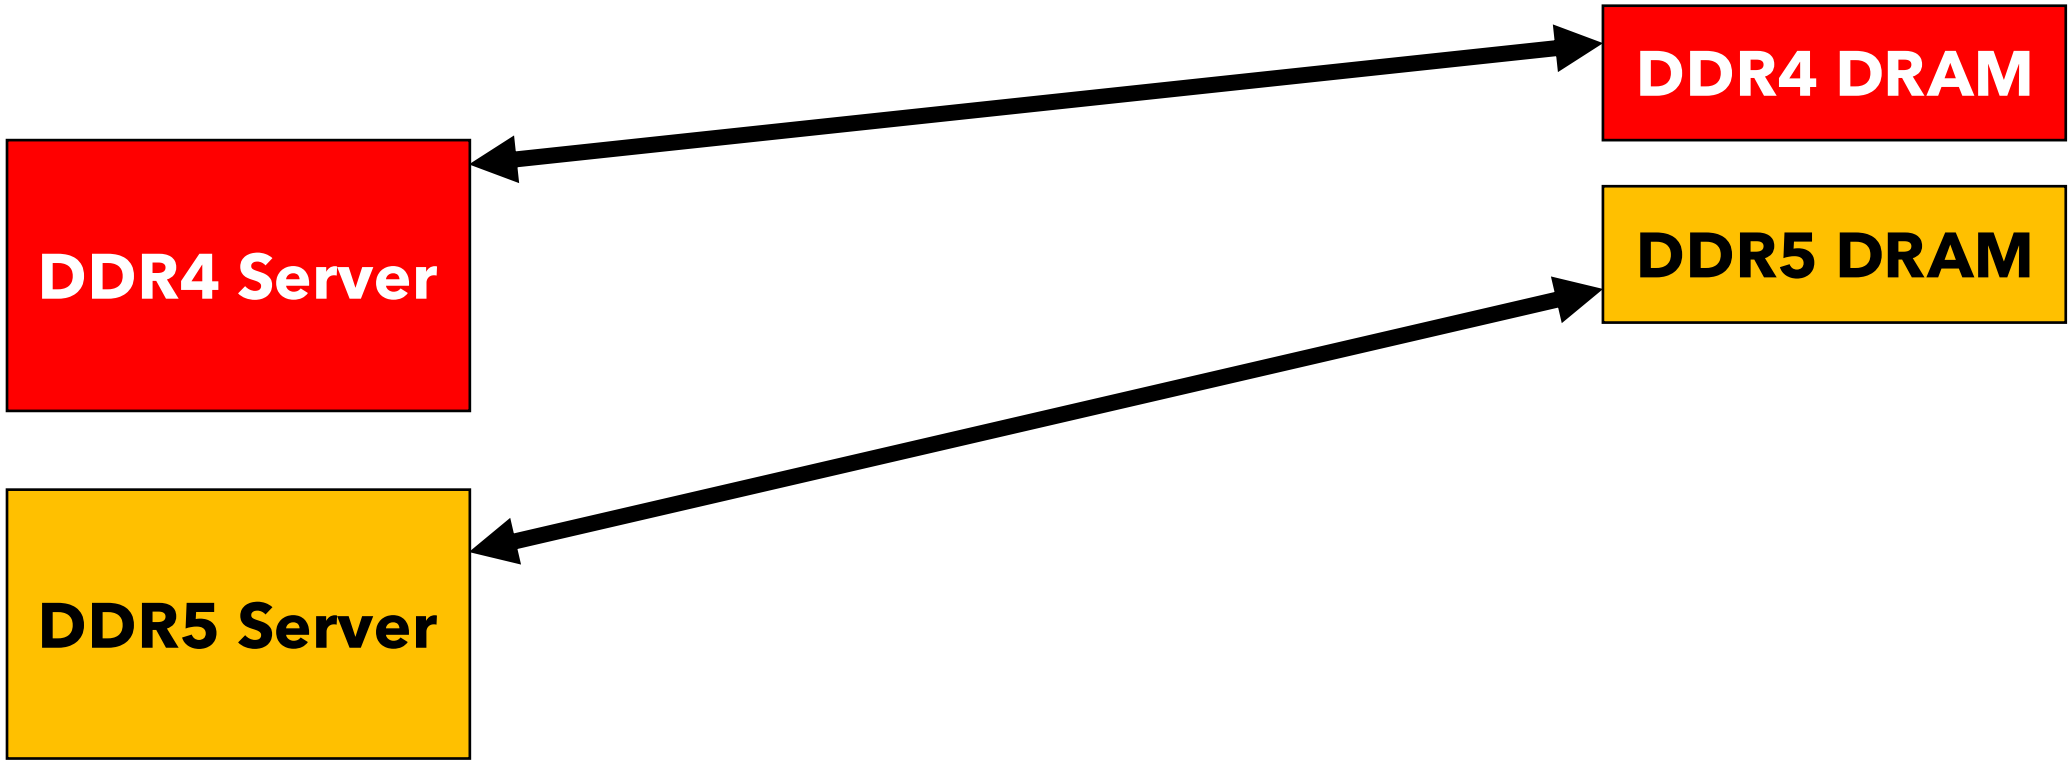

CXL Supports Any Memory, Volatile or Persistent

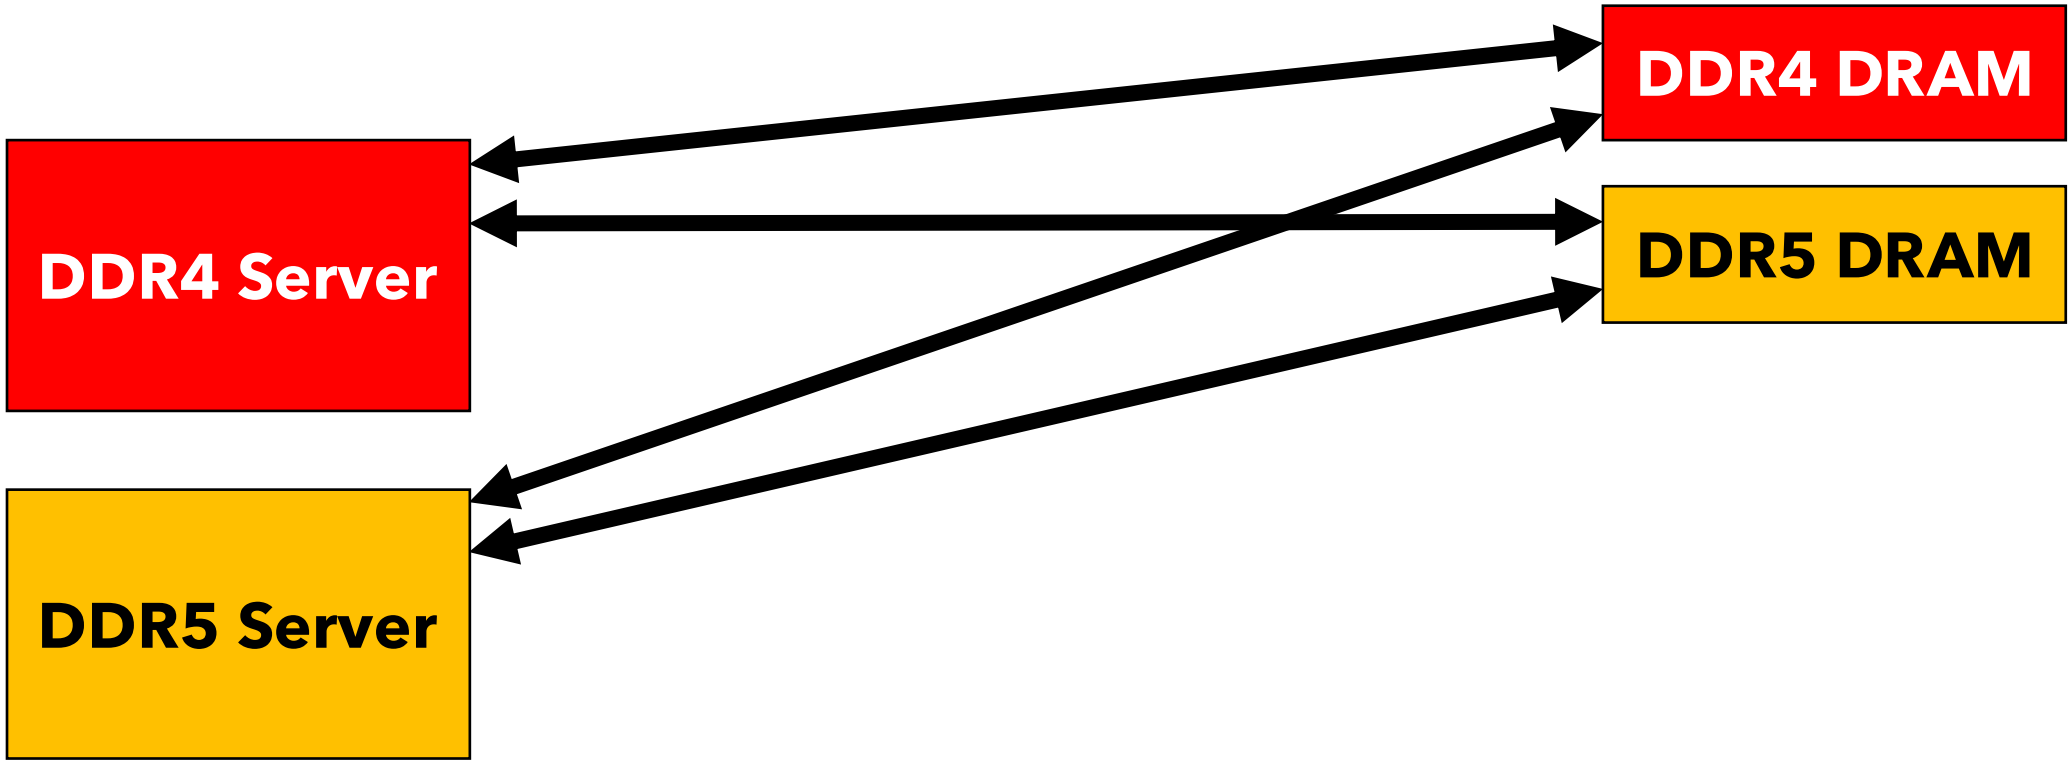

CXL Supports Any Memory, Volatile or Persistent

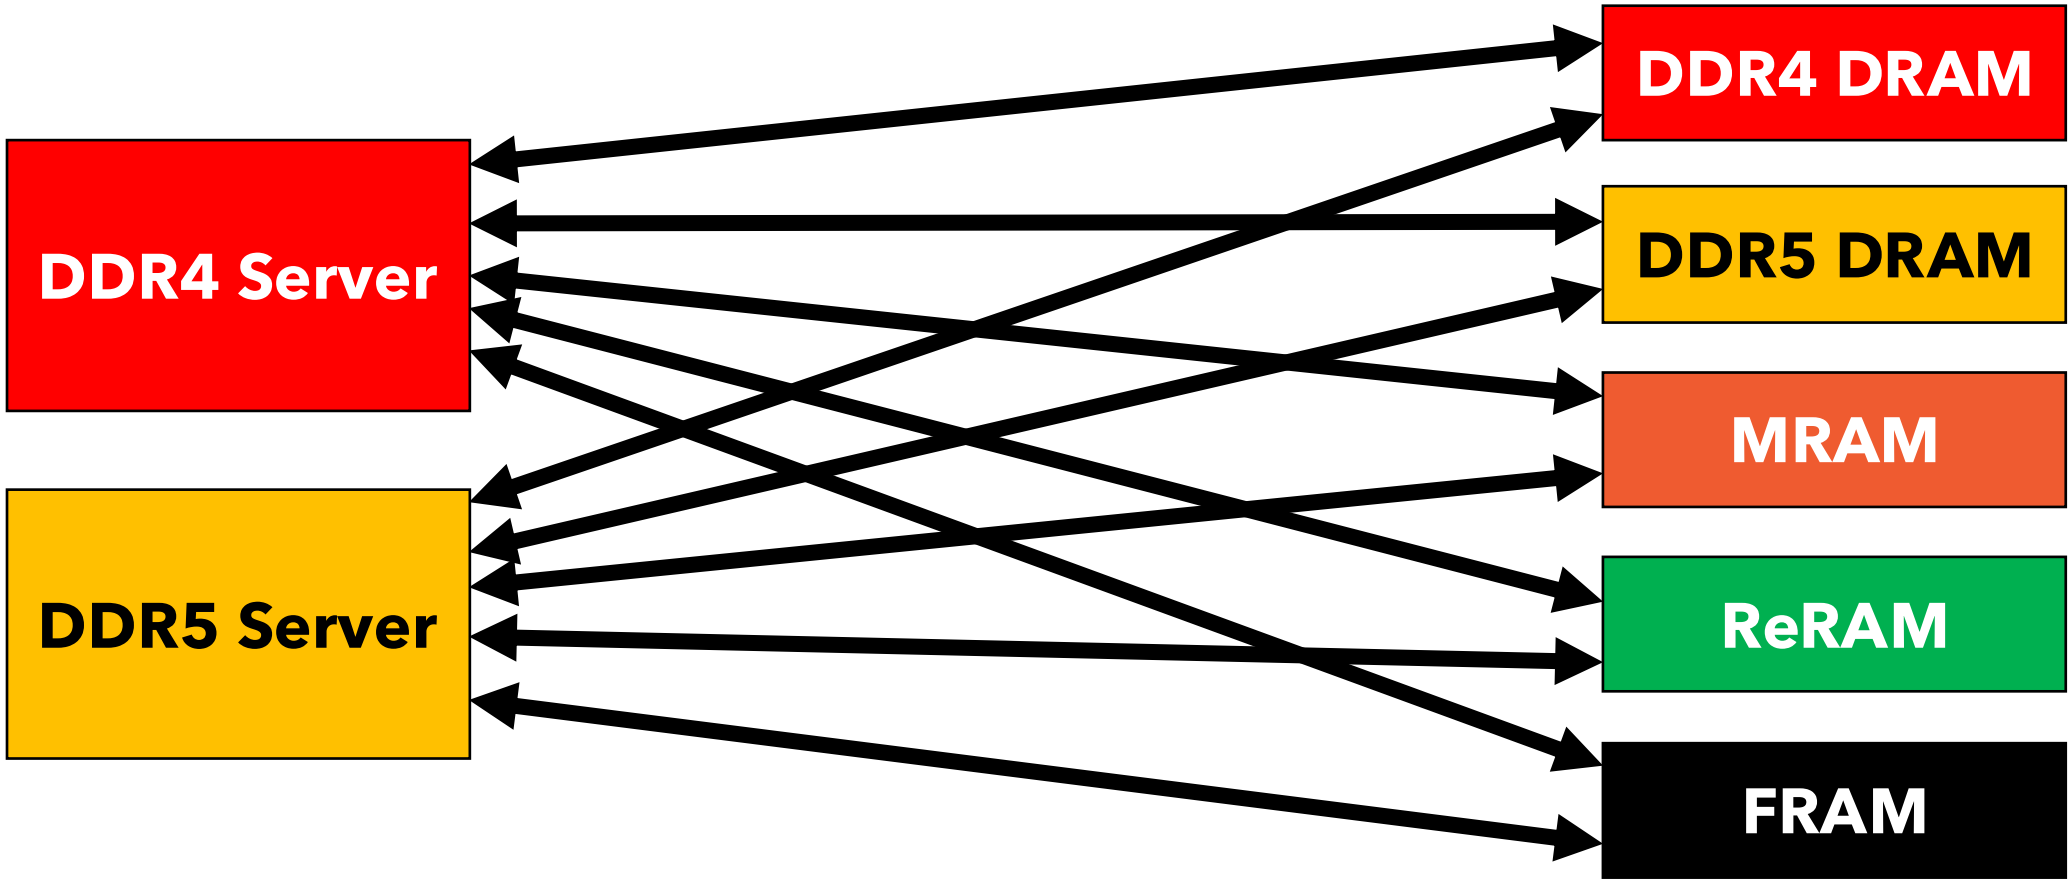

CXL Supports Any Memory, Volatile or Persistent

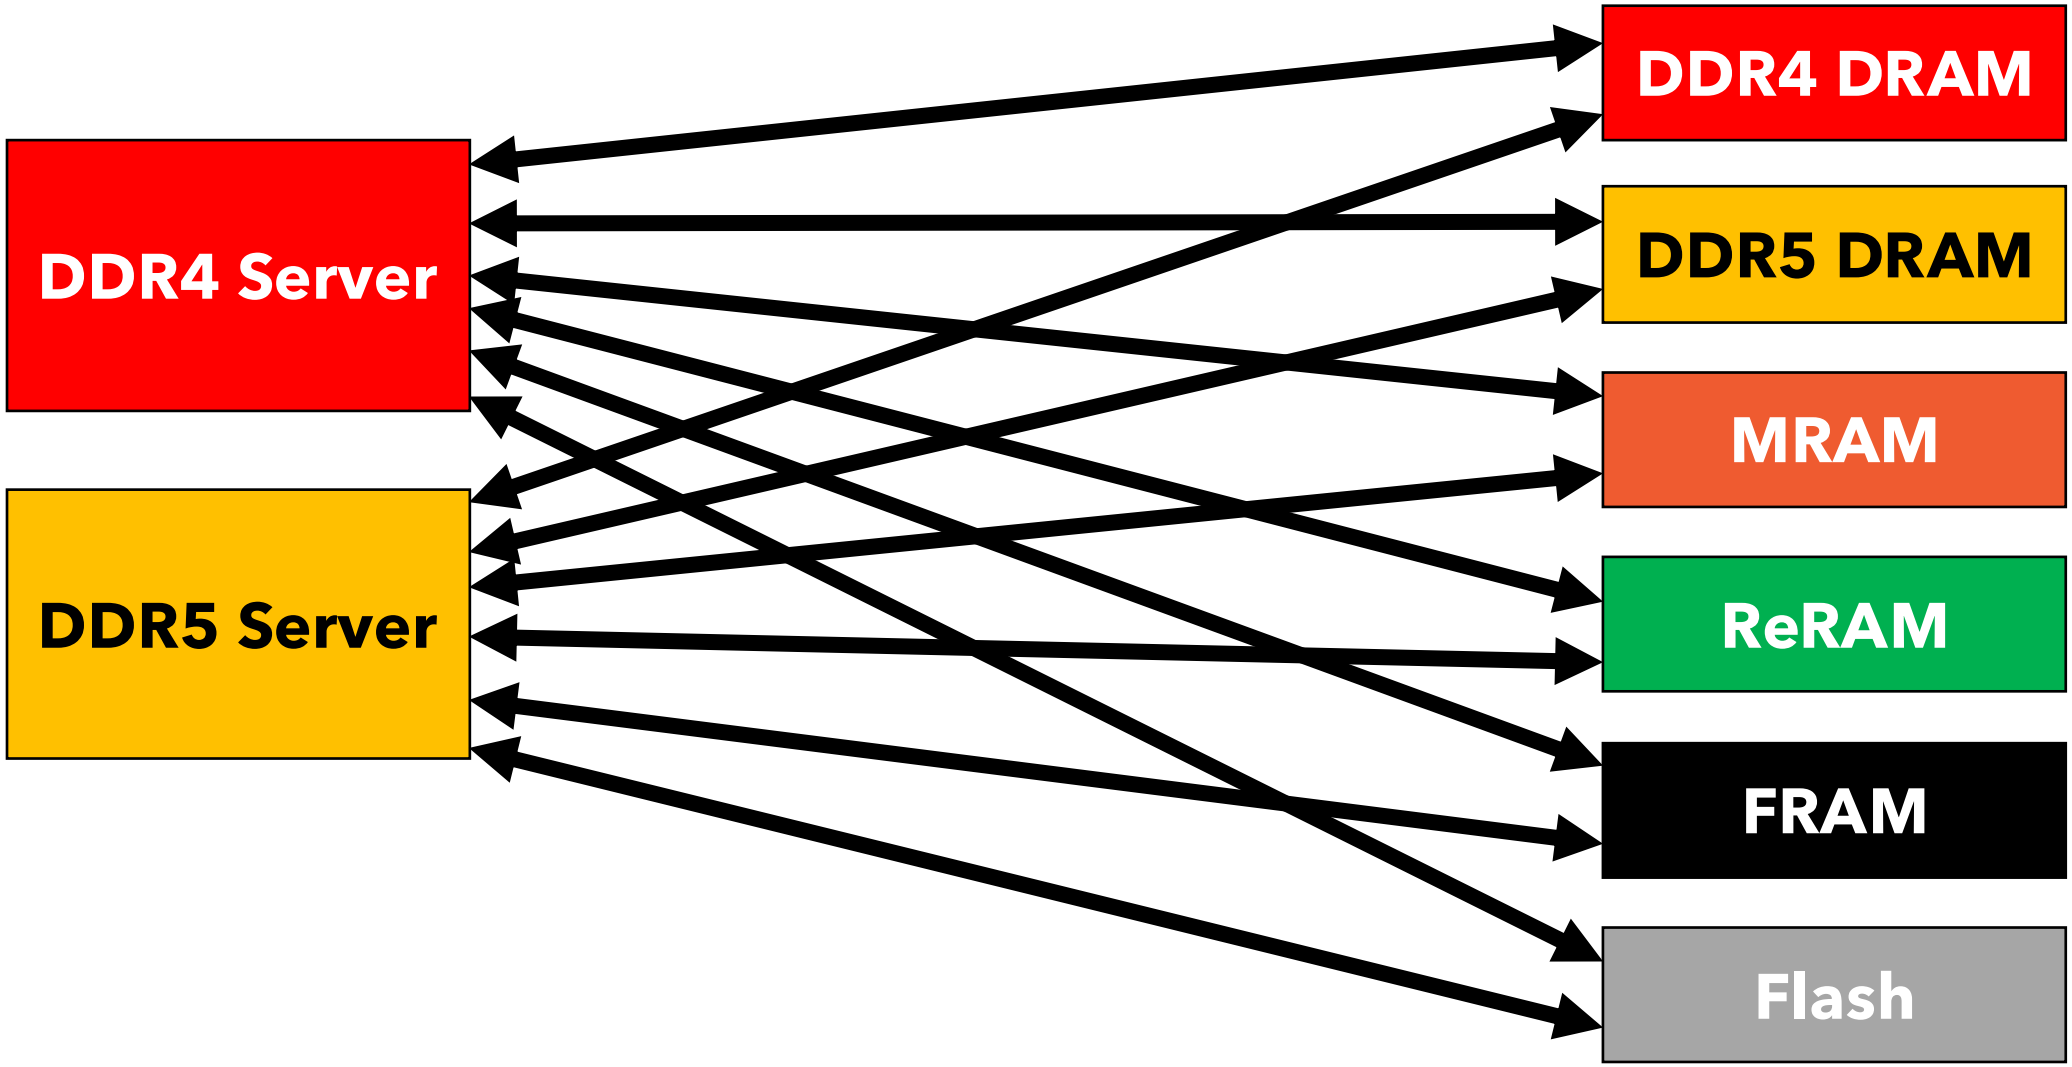

CXL Supports Any Memory, Volatile or Persistent

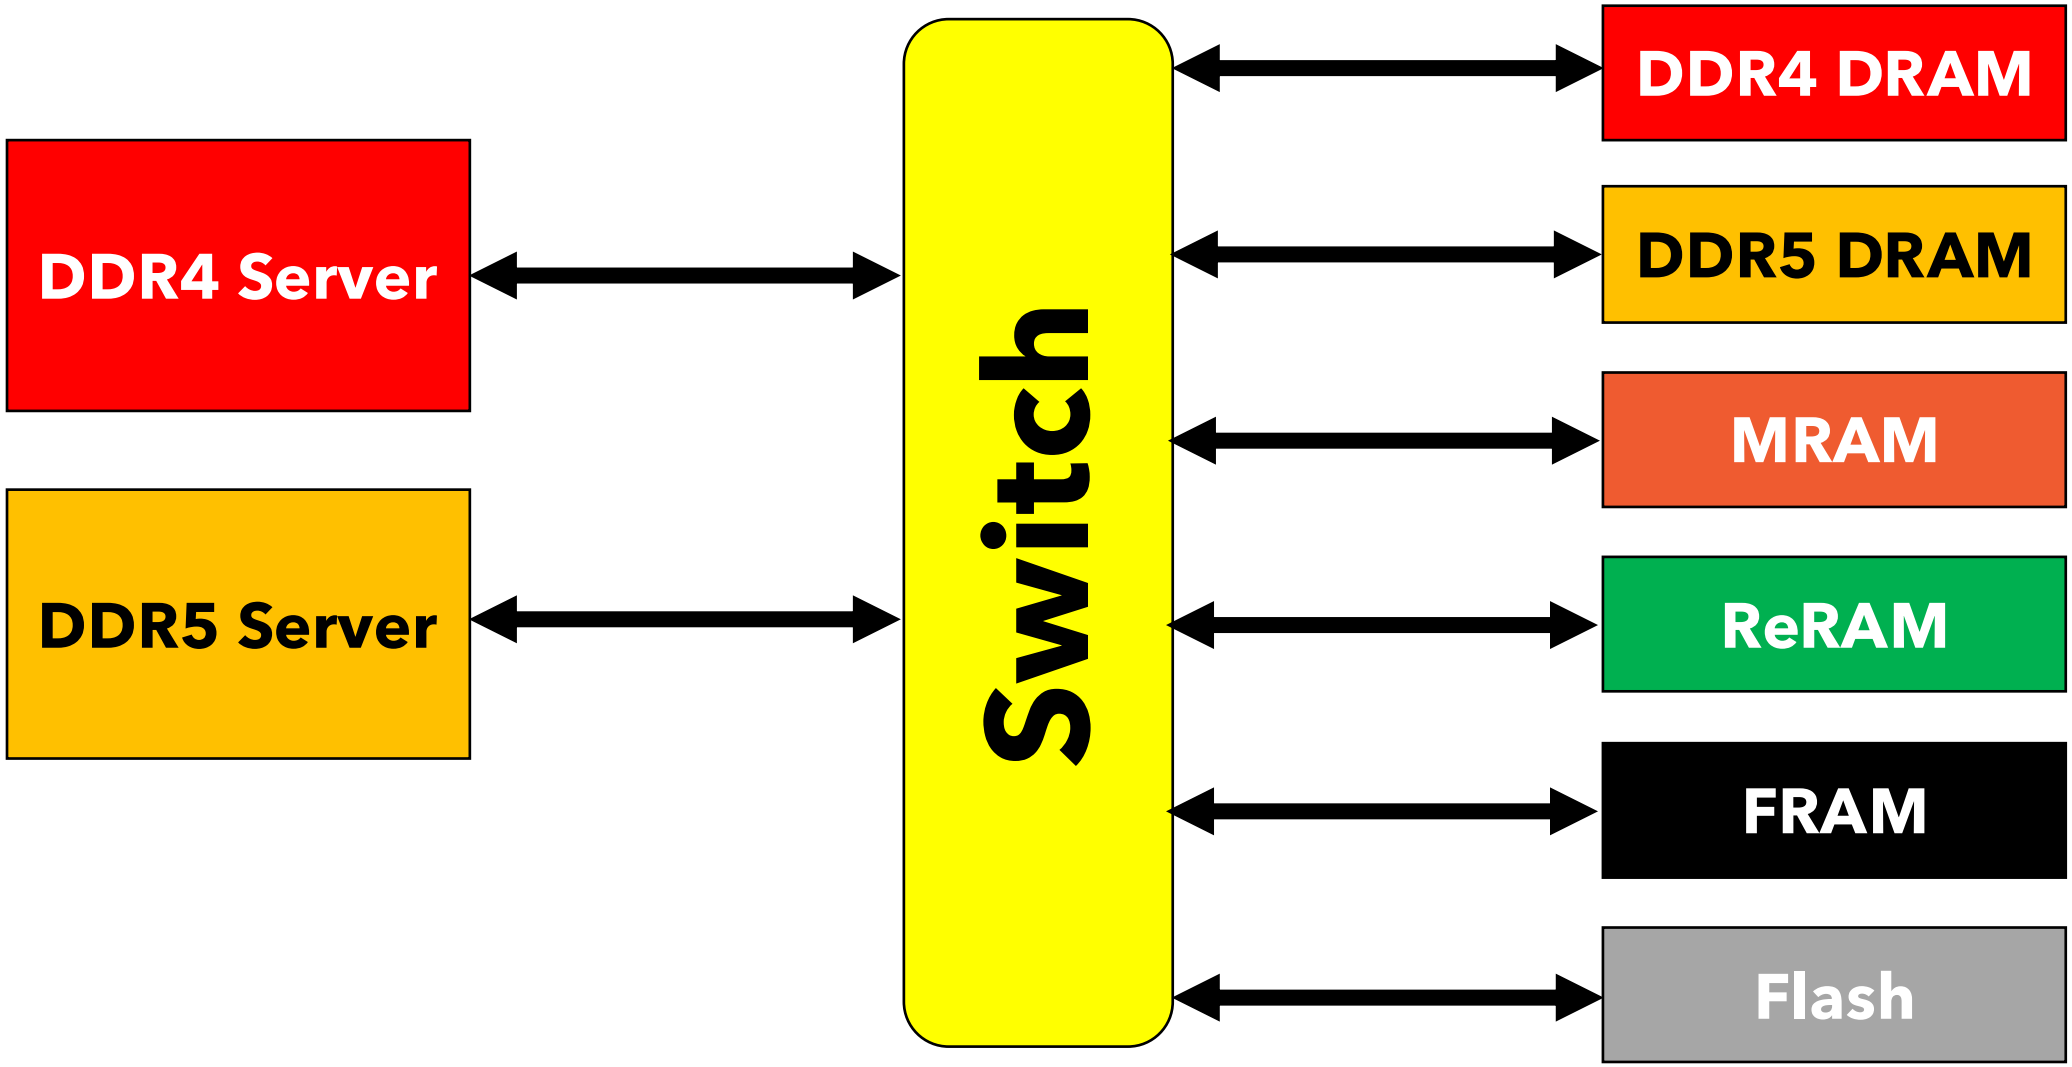

CXL Forecast

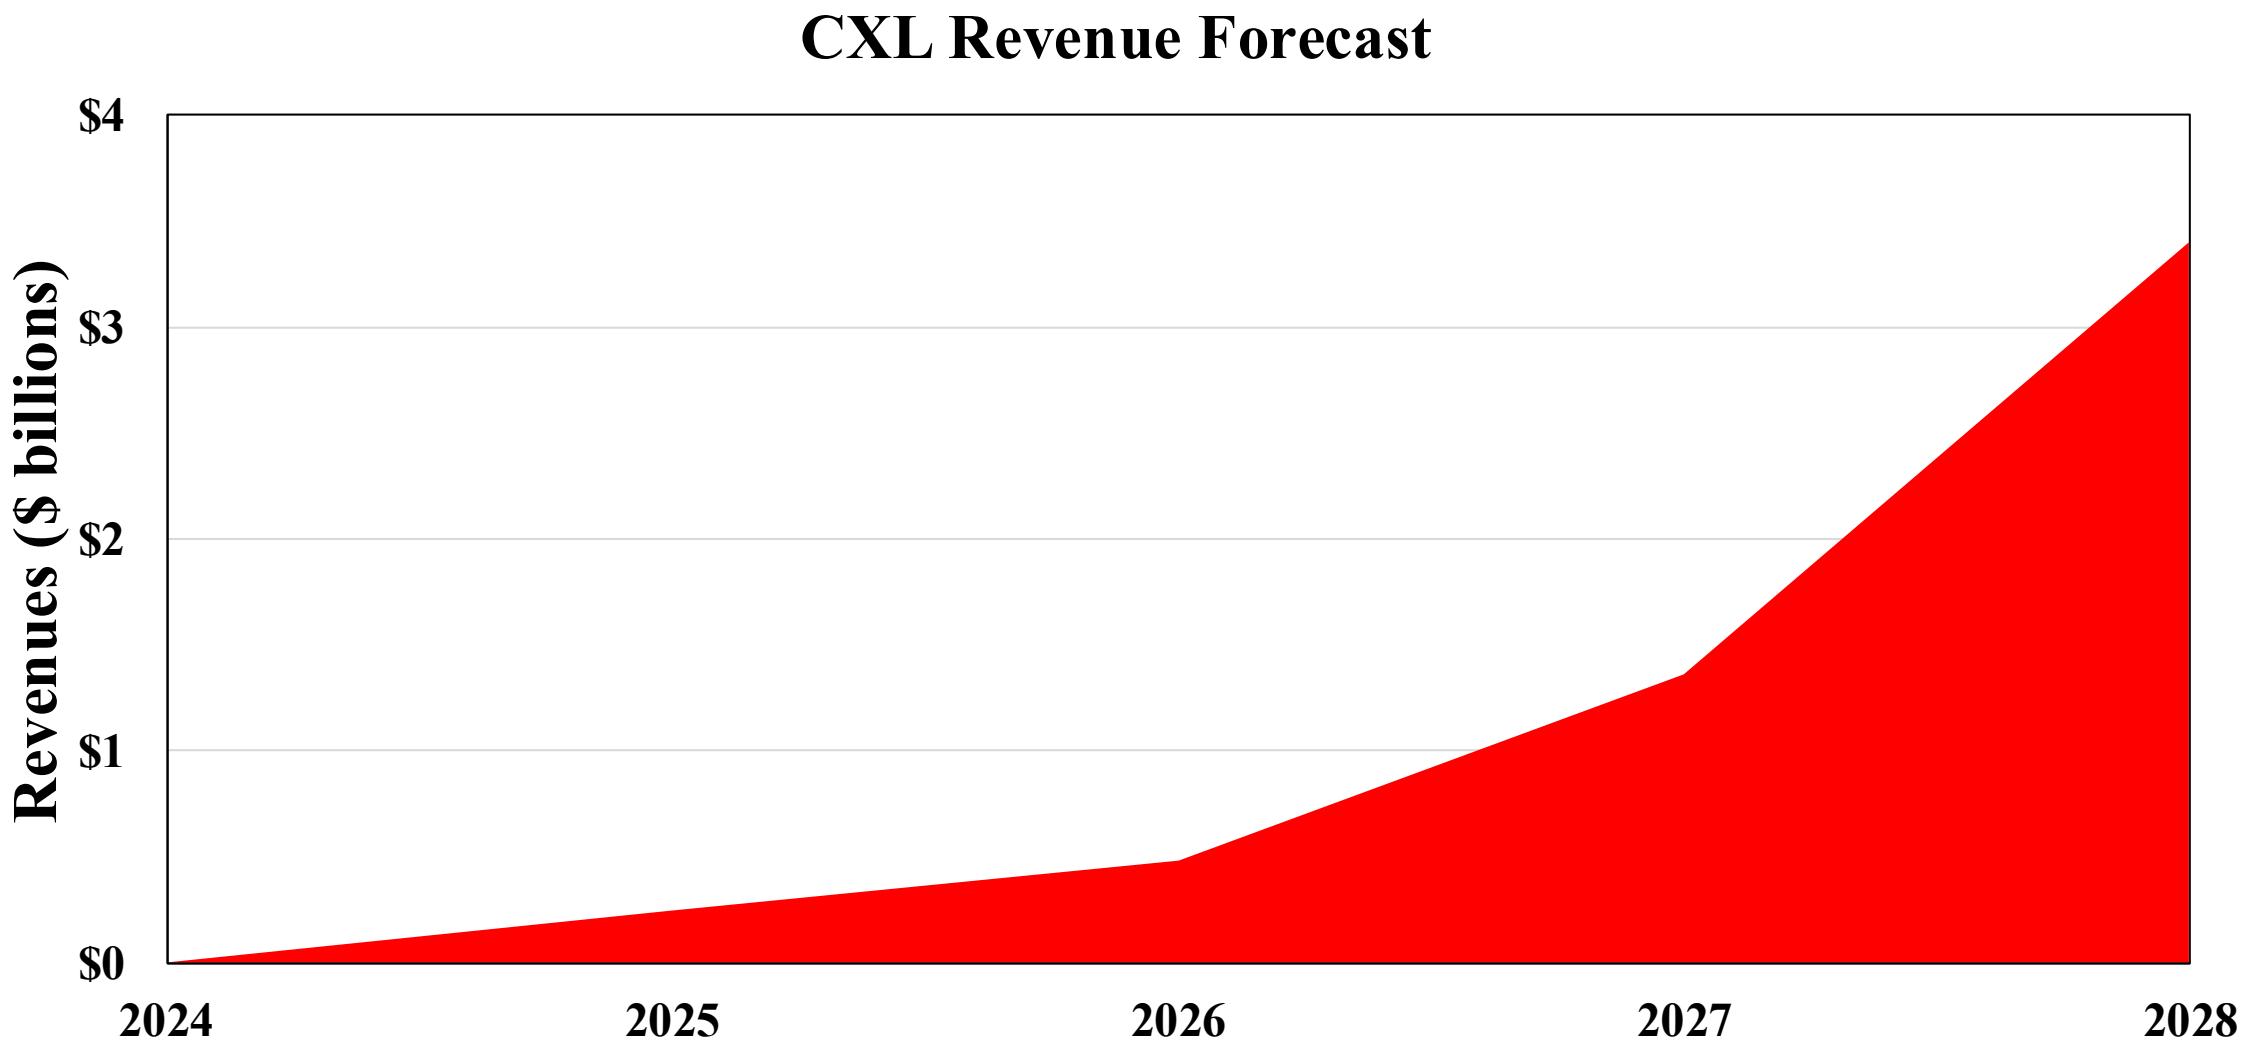

SNIA | CMS A SNIA Community

Long-Term Impact

- Re-thinking system architecture
  - Disaggregated memory
  - Processor arrays with mesh networks
  - Memory agnostic
- Better memory bandwidth & size vs. worse latency
  - Design-arounds will optimize for this

Report: *CXL Looks for the Perfect Home*

- Exhaustive report from Objective Analysis
- Covers all perspectives
  - Where CXL is useful, and where it isn't
  - Demand drivers for CXL DRAM modules
  - Opportunities outside of DRAM
  - Forecast (Revenues, units, ASP)
- Available for immediate download:
- Objective-Analysis.com/reports

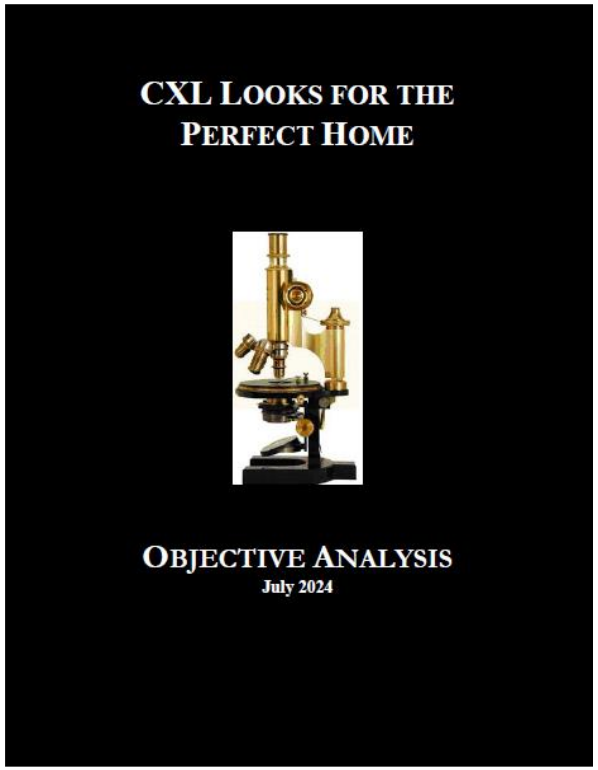

SNIA | CMS
A SNIA Community

Market Sizing

Market Factors

- Drivers: Why will it grow?
- Types: Embedded vs. Stand-Alone (Discrete) devices
- Economies of Scale
- One or Many Winners?

Economies of Scale

- A smaller die size doesn't guarantee a cheaper part
- Wafer volume is critical
- Pricing determines sales volume
- Sales volume determines production cost
- It's a chicken-and-egg problem

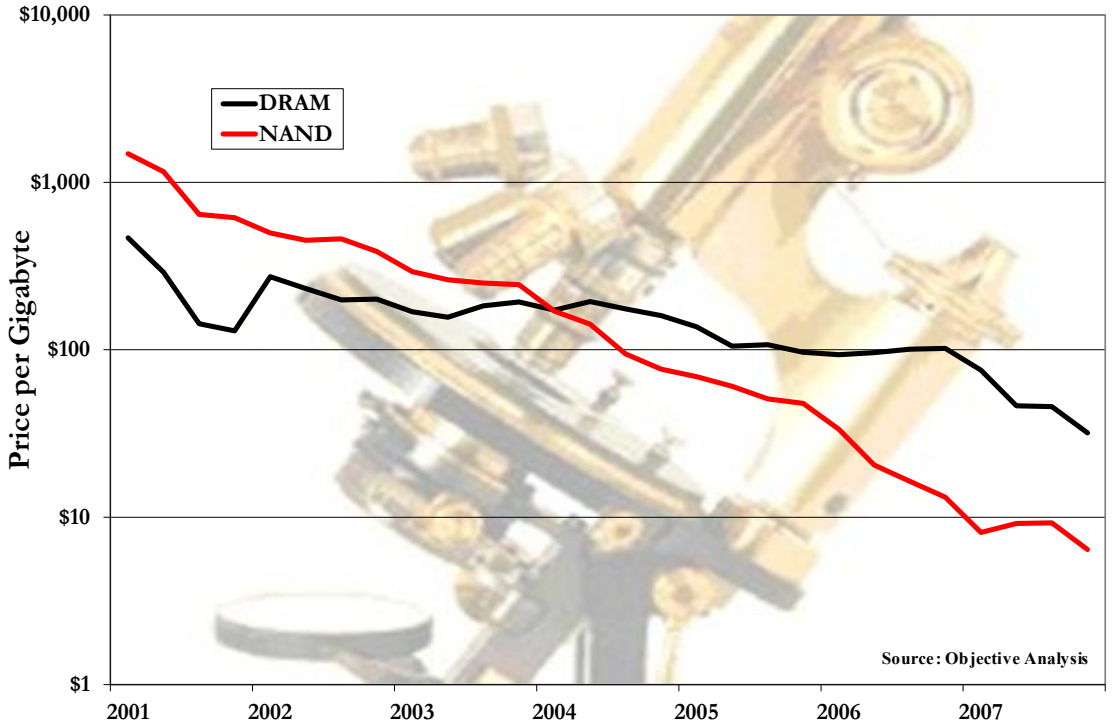

SNIA | CMS A SNIA Community

One or Many Winners?

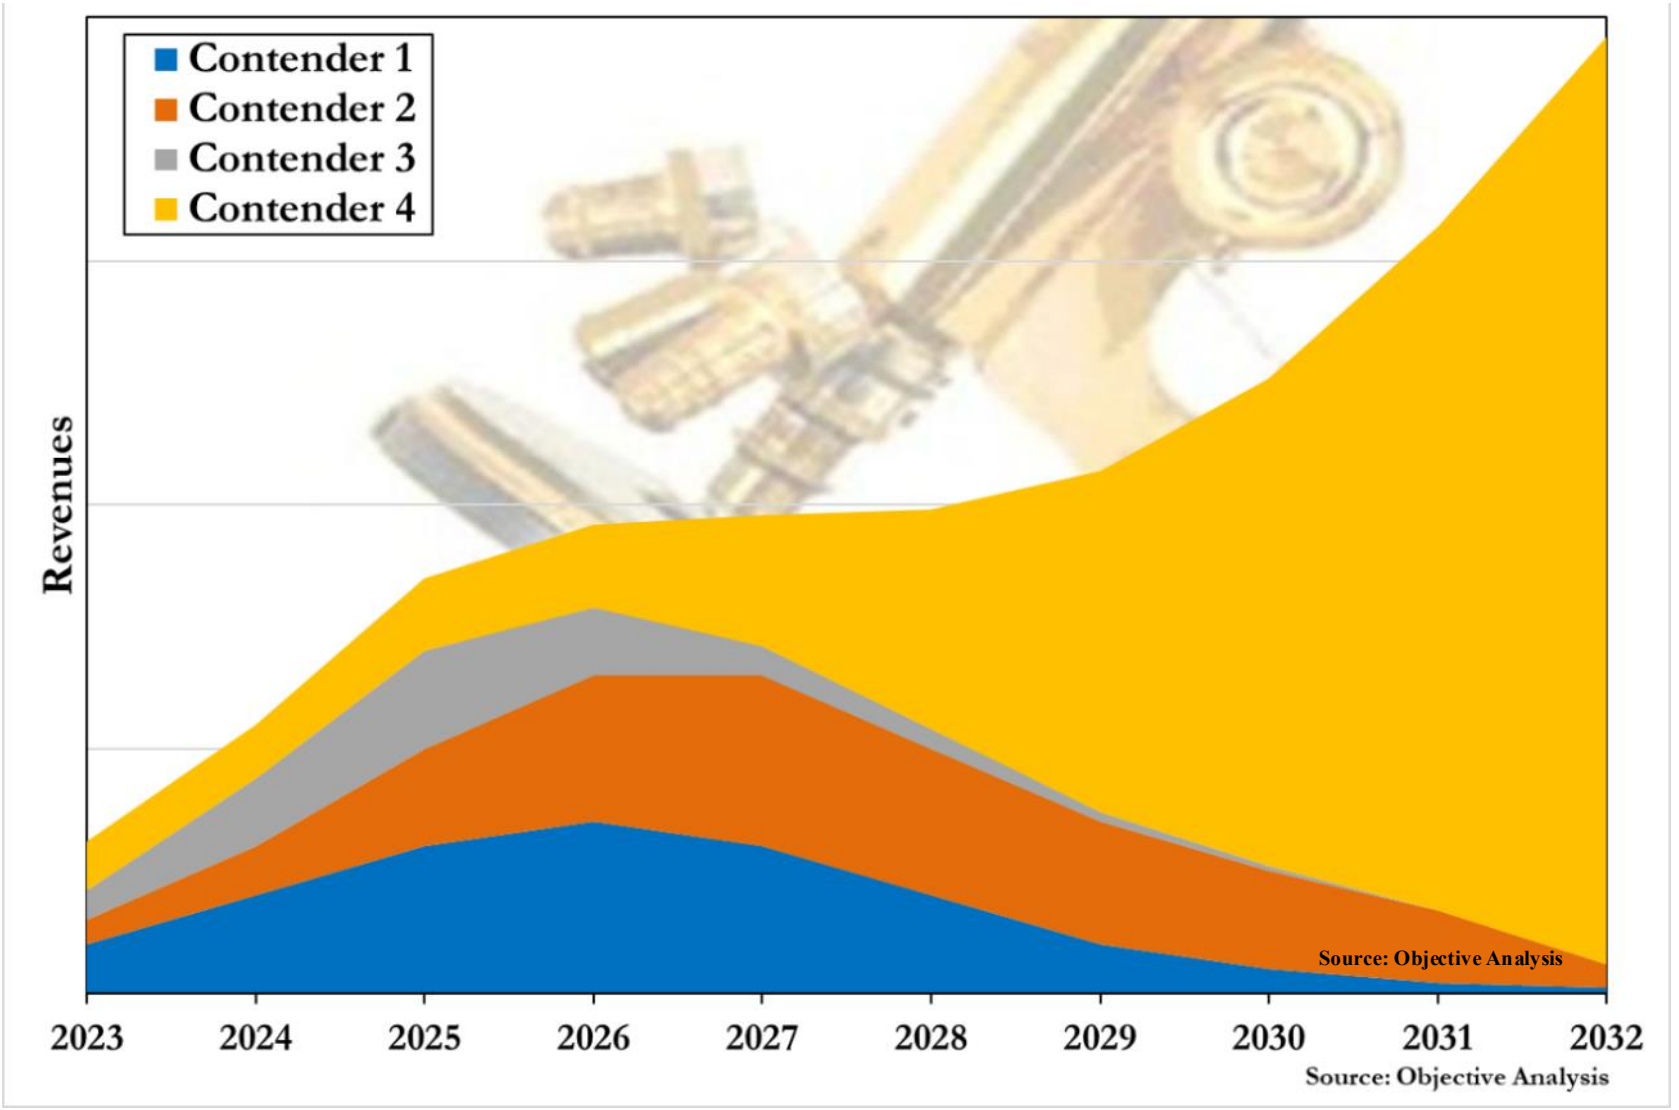

SNIA | CMS A SNIA Community

Forecast Assumptions

- Embedded memories drive early volumes
- Wafer volume drives economies of scale
  - Opens new markets for stand-alone chips
- Low price always wins out over better features
  - NOR/SRAM cell phones vs. NAND/DRAM
- We take no position on which technology will win out
  - Our only expectation is that only one will succeed in a big way

Rapid New Memory Terabyte Growth

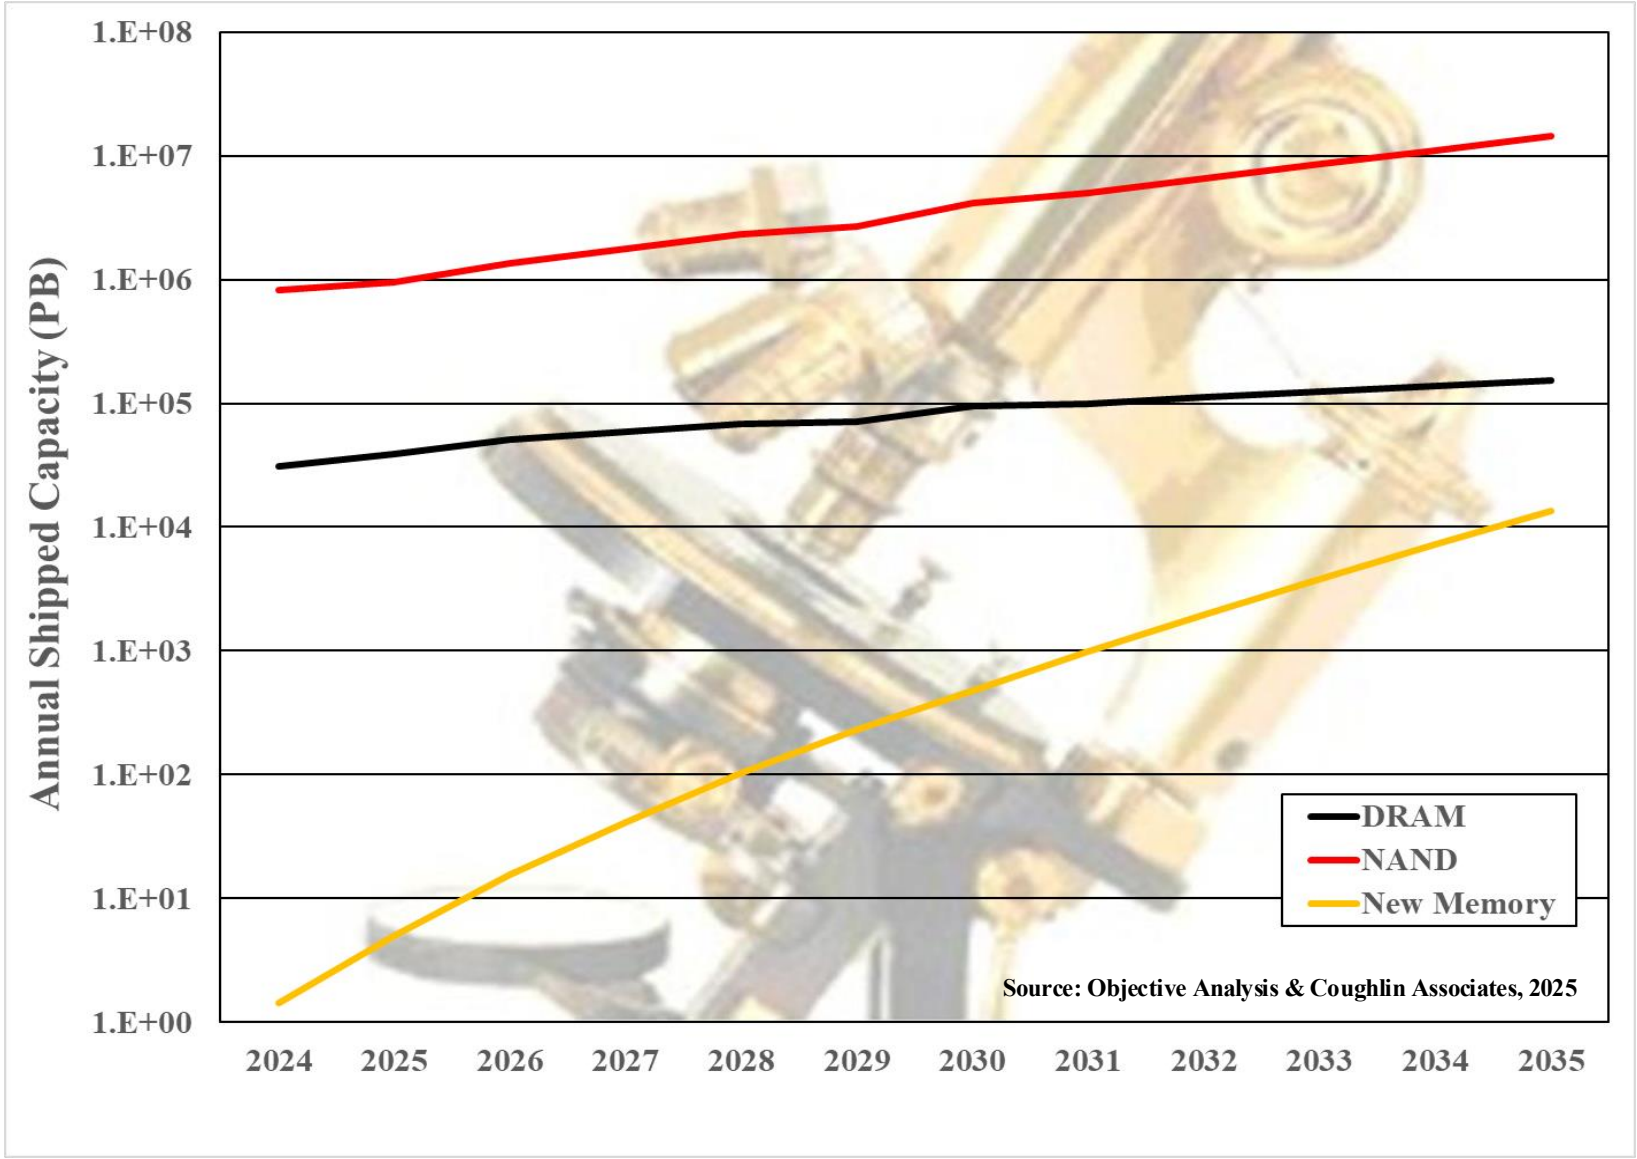

SNIA | CMS A SNIA Community

Report: *New Memory – Not Just for AI*

- Exhaustive examination of new memory technologies
  - 326 pages
  - 219 Figures
  - 36 Tables
  - 175 company profiles
- Technologies, forecasts, tool use, business issues, situation analysis

Order now for immediate download!

https://Objective-Analysis.com/reports/#Emerging
https://TomCoughlin.com/product/emerging-memory-report/

SNIA | CMS A SNIA Community

SNIA | CMS
A SNIA Community

Q&A