SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2015

IMDB NDP Advances

Gil Russell, Principal Analyst

Semiscape™

CPU Centric vs. Data Centric Computing

- The nature of computing is no longer dominated by the execution of arithmetic and logic calculations but is in transition to the handling of large data volume and the cost of moving data to the locations where computations are performed
- In the computing-centric model data lives on disk, or even tape, and moves as needed to the CPU across a deep storage hierarchy - sufficient when computational aspects dominate data movement aspects
- In the data-centric model data lives in different storage levels within the hierarchy, with processing engines surrounding the data and operating on such data without moving it across the system

Semiscape

Data Access - Architecture Changes

- In-Memory Database (IMDB)
  - Entire Database Image Placed in DRAM (DIMMs)
    - Proximity to Memory Controller/Management Unit
    - Minimize Latency
    - Maximize Bandwidth
    - Preserve System Infrastructure Legacy
  - Lower system power
  - Data optimization may be required
  - Current size limited to 6 TB (64GB DIMMs)
  - Cost (till now) has been prohibitive

Semiscape

IMDB Market Motivation

- Market Drivers: Real Time Analytics and Transaction Processing (OLAP & OLTP)
  - "The reason for your IT to be on the Cloud"
- CAGR between 25% to 43%+ (SAP HANA 130% in 2013)
- Persistent Memory introduction
  - Maintain DB image durability
  - Lower system cost and lower power
- Equipped servers capable of "cognitive computing intelligence and predictive analytics" - next era of computing

Semiscape

Lower Latency Opens Up Business Value

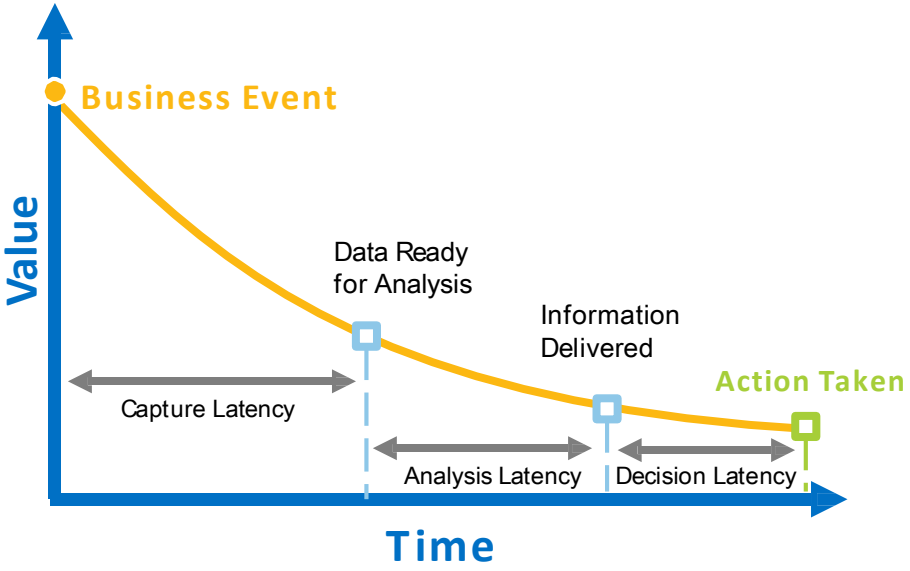

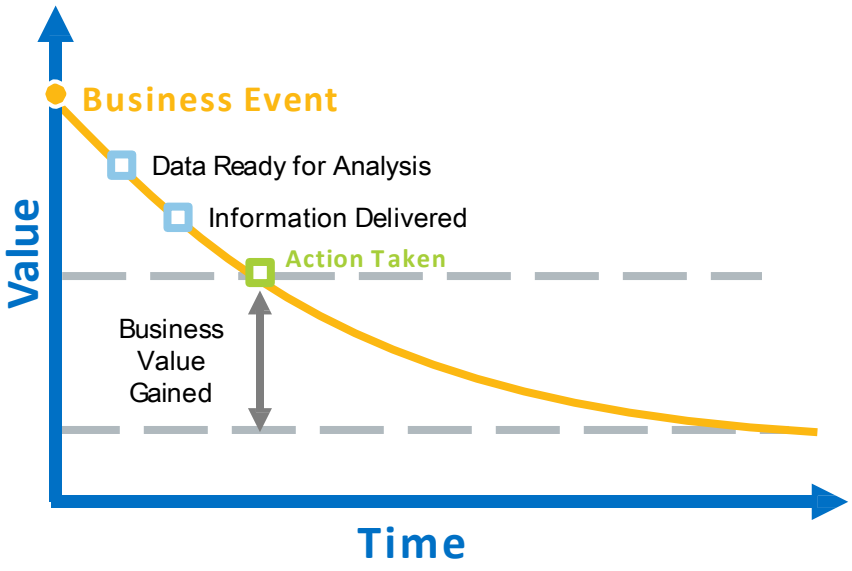

Source: Greg Matson Director of Strategic Planning, Non-Volatile Memory Solutions Group, Intel Corporation – IDF 15
http://jtonedm.com/2012/11/21/decision-latency-revisited

Semiscape

Storage Hierarchy Tomorrow

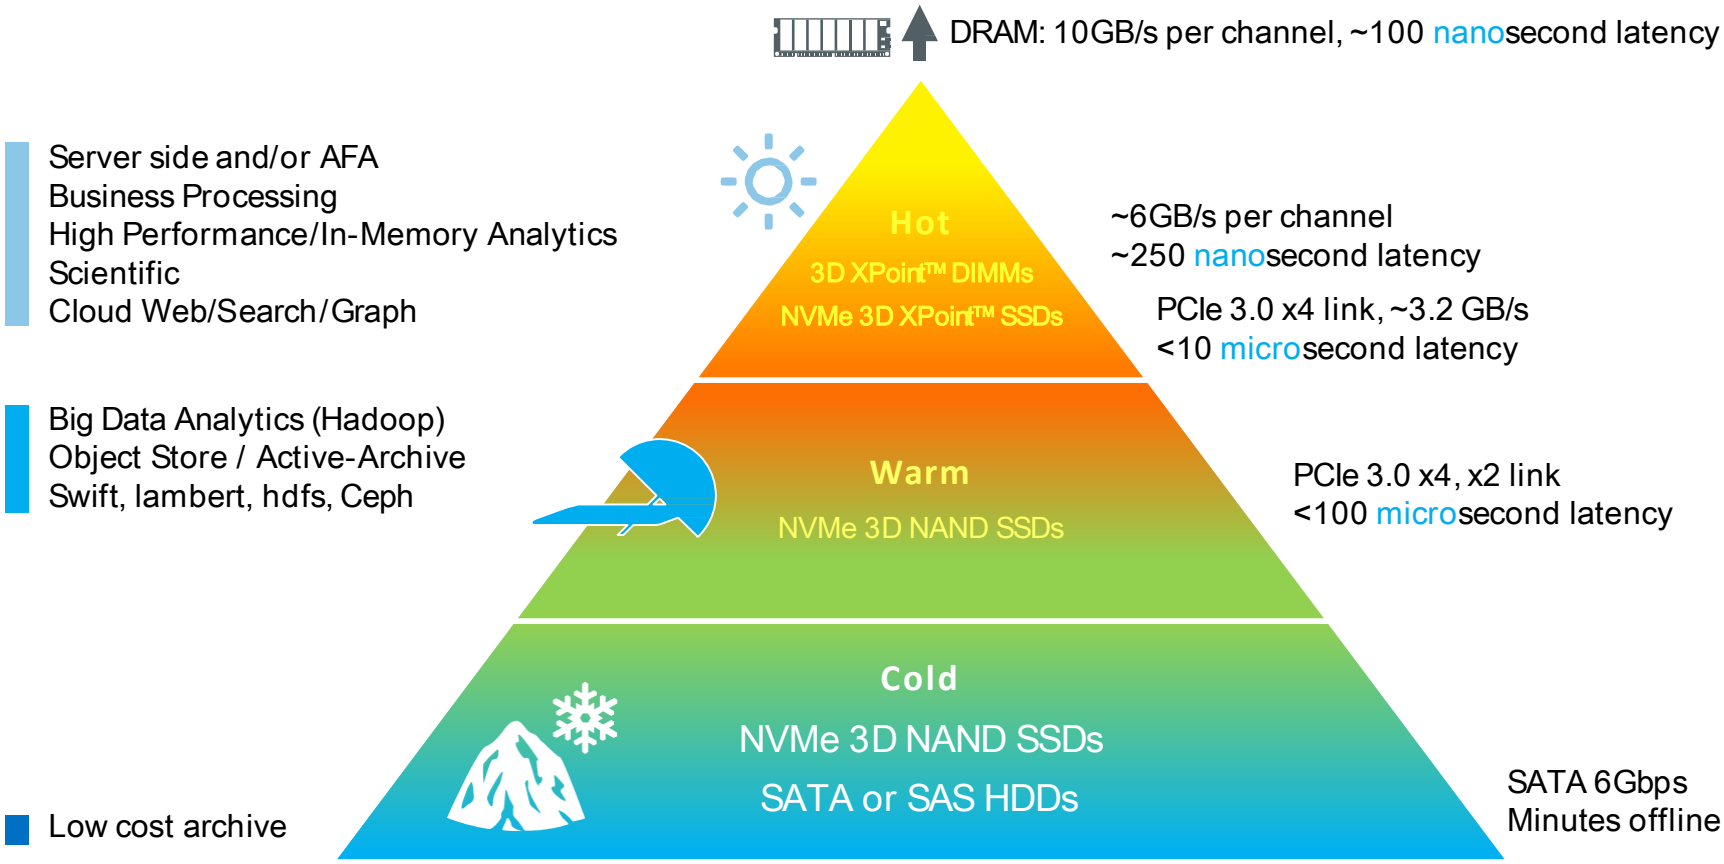

Comparisons between memory technologies based on in-market product specifications and internal Intel specifications.

Source: Greg Matson Director of Strategic Planning, Non-Volatile Memory Solutions Group, Intel Corporation – IDF 15

Semiscape

SNIA NVM Programing model

The Four Modes

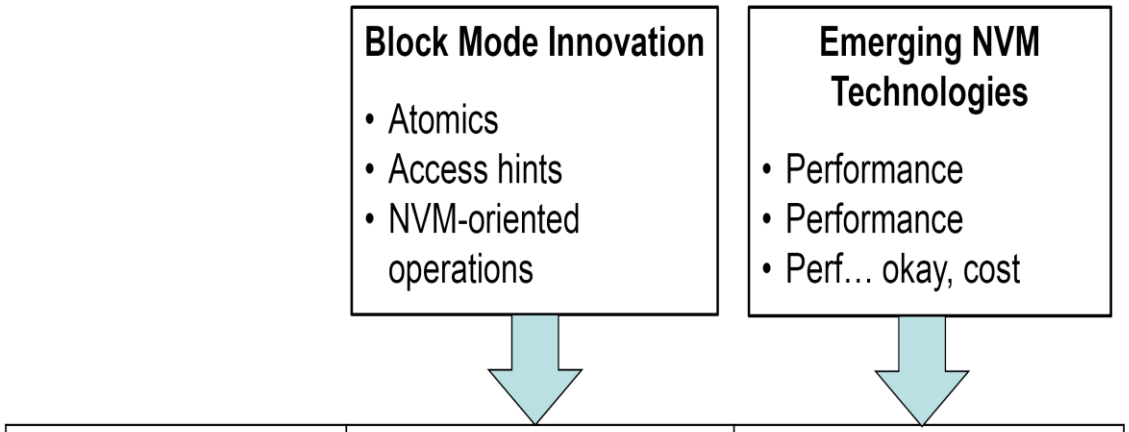

	Traditional	Persistent Memory
User View	NVM.FILE	NVM.PM.FILE
Kernel Protected	NVM.BLOCK	NVM.PM.VOLUME
Media Type	Disk Drive	Persistent Memory
NVDIMM	Disk-Like	Memory-Like

FMS 14
Jim Pappas, Intel

Eliminate File System Latency with Memory Mapped Files

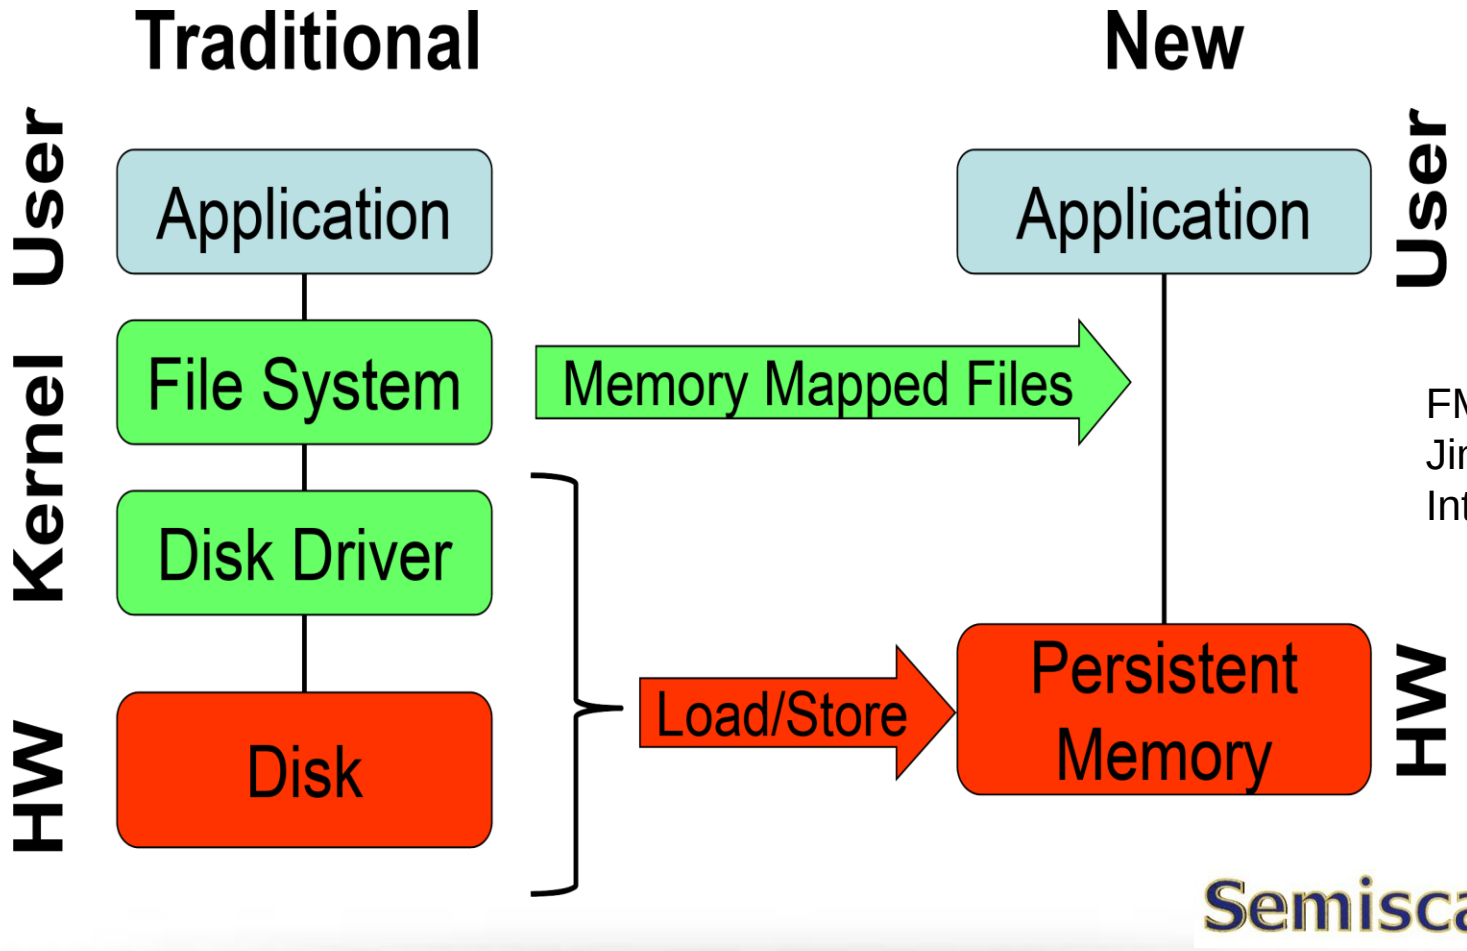

FMS 14
Jim Pappas, Intel

Semiscape

Application View of I/O Elimination

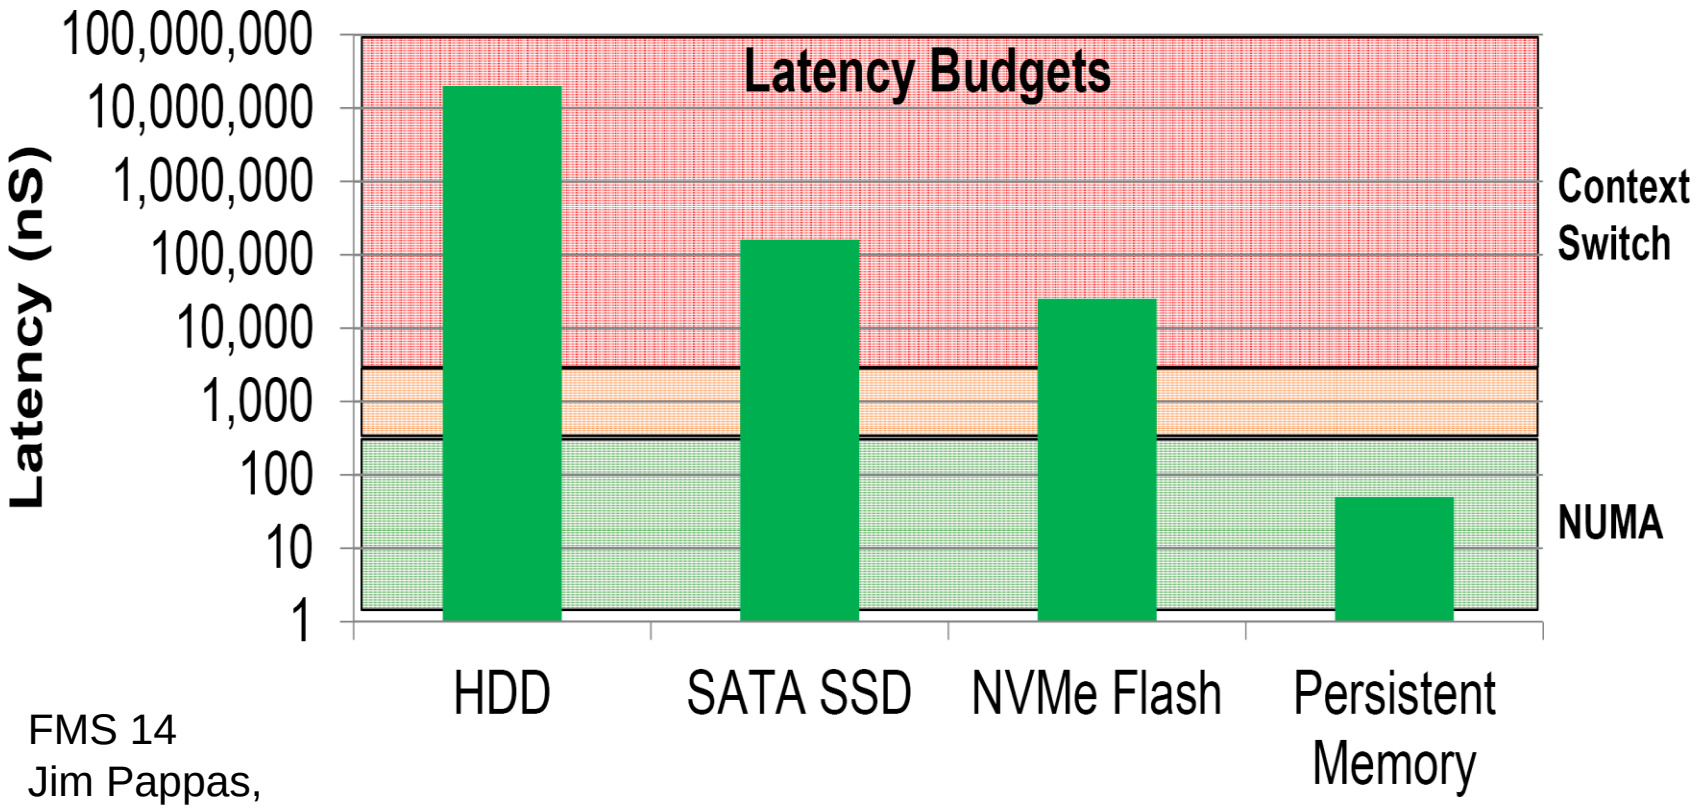

FMS 14
Jim Pappas, Intel

Example: Lenovo System x3850 X5

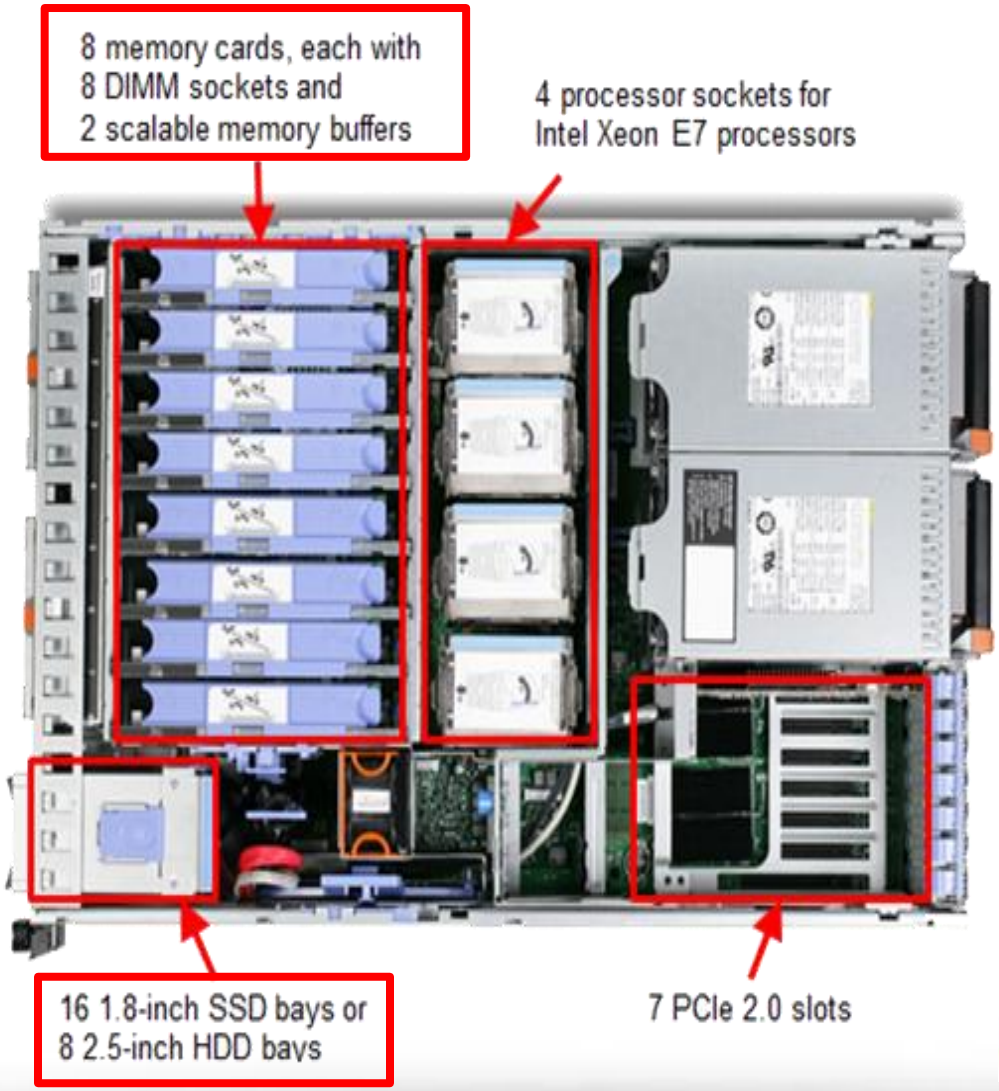

Semiscape

Intel DIMMs

Based on 3D XPoint™ Technology

- DDR4 electrical & physical co[illegible]ible
- Supported on next generation Intel® Xeon® platform
- Up to 4X system memory capacity, at significantly lower cost than DRAM
- Can deliver big memory benefits witho[illegible] modifications to OS or applications

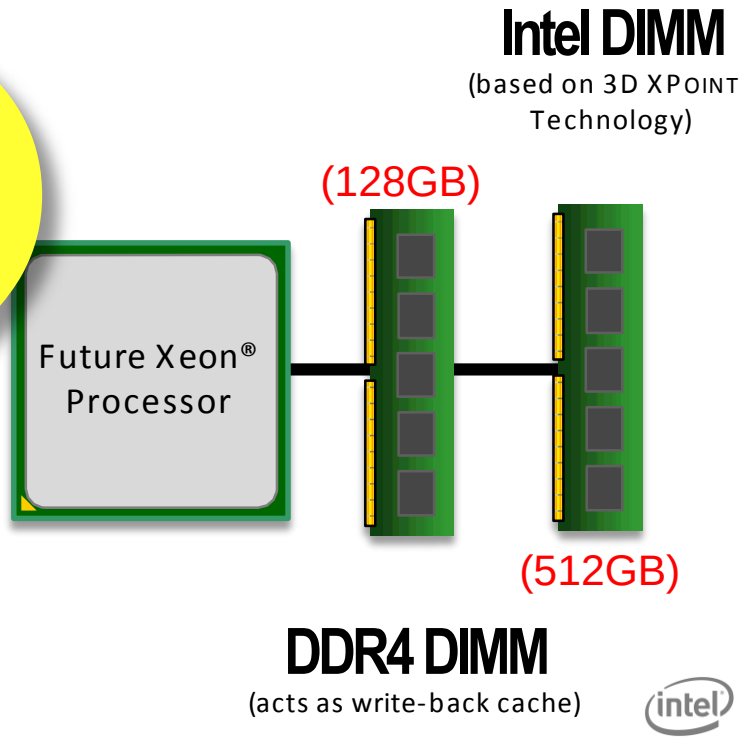

Source: "Intel® Non-Volatile Memory Inside, The Speed of Possibility Outside", Rob Crooke and Al Fazio, Slide 14, IDF 15.

Semiscape

Hypothetical IMDB DIMM Instantiation

IMDB Main Memory

- System Maximum: 64 TB
- Maximum DRAM portion enabled by:
  Introduction of 16Gb, 4 high TSV, DDR4, x16 chips,
  128GB Front Side LRDIMM
- Maximum Persistent Memory portion enabled by:
  introduction of 3D XPoint memory, 128Gb, 4 high TSV, x16 chips,
  1024GB Back Side LRDIMM
- DRAM Legacy support
- IMDB memory persistence support by XPoint Memory (64Gb per XPoint chip provisioning)
- No longer a need for “flush on fail” or recovery
- Power requirements significantly reduced

Semiscape

Near Data Acceleration (NDA) Posit

Semiscape

Hypothetical IMDB DIMM Instantiation

128GB DDR4 Front Side

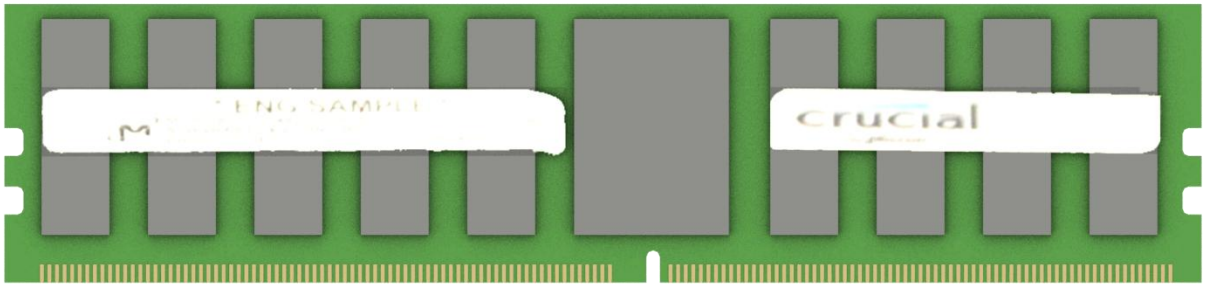

1024GB XPoint Back Side

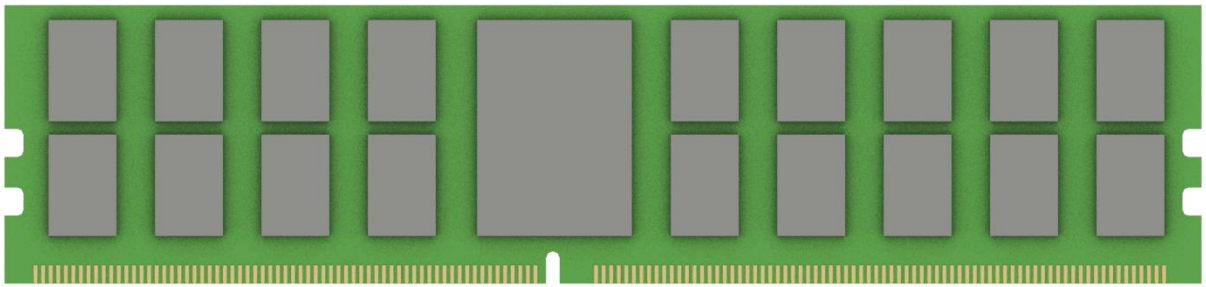

Semiscape

Near Data Processing (NDP)
A Simple Case

- NDP enables massively parallel search (query) functions (binary, ternary)
- Through Silicon Via (TSV) allows wide parallel DRAM &/OR XPoint data to connect to the logic interposer below stacked devices:
  - Consists of a simple FPGA search engine on the logic chip
  - Programmable application specific execution
  - Compatible with commodity DRAM & XPoint
  - System compatible with MapReduce extension
- Reduces system search power: ~ 60%
- Peak search: 19.2 GB/sec/DIMM
- Scales linearly with memory size

Semiscape

Near Data Processing (NDP)

- Hybrid Memory Cube (HMC) is an attractive candidate for NDP
  - ✓ Placing XPoint on an HMC stack with DRAM under consideration
  - ✓ XPoint dramatically increases density - increased latency and reduced bandwidth is the drawback
  - ✓ The idea of a 3D “hive’ matrix of “Smart Memory Cubes” [SMC] “hives” interconnected with photon pipes is under discussion/development
  - ✓ Persistence is not part of the HMC 2.0 Spec.

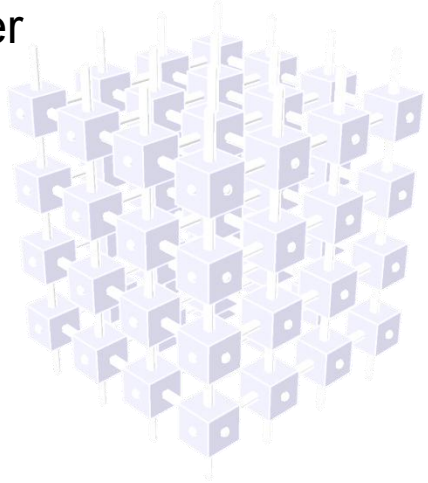

Semiscape

Near Data Processing (NDP) Posit

- Legacy compatibility is required for each successive DRAM generation – this may be ending
  - Silicon Photonics will replace copper but is not close to being cost effective
  - HPC budgets may cover high costs but the cost needs to come down an order of magnitude before widespread adoption begins
  - Introduction of XPoint memory opens a new episode if not an entirely new segue in computing
  - DRAM “generational step” performance has been decreasing
    - DDR4 may be the last generation of capacitor based storage devices

Semiscape

Summary

- IMDB & NDP is a practical near term solution for managing Big Data
- Moving from CPU Centric to Data Centric will rearrange the entire computing technology landscape
- Effect of Storage Class Memory on NAND-Flash SSDs:
  - ✓ Little or no affect in the near term
  - ✓ 2nd generation Storage Class Memory will begin to affect the NAND-Flash SSD market starting in 2017

Thank you for coming!

Semiscape