Storage in the DIMM Socket

Adrian Proctor
Vice President, Marketing
Viking Technology

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the 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
  - The 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 the SNIA Education Committee.
- 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.

Abstract

- As data sets continue to grow, IT managers have begun seeking out new ways for memory technology to be deployed in the data center in order to take greater advantage of the performance and latency benefits.
- Non–Volatile DIMMs, or NVDIMMs, provide a persistent memory solution with the endurance and performance of DRAM coupled with the non–volatility associated with storage.
- This tutorial will provide a general overview of this emerging technology and how it plays in the data center.
- You will learn what an NVDIMM is, how it works, where it fits and why system architects should consider them for their next generation enterprise server and storage designs.

Agenda

- The Memory / Storage Problem (Latency)
- New Memory Technology Roadmap
- NVDIMMs – What they are
- NVDIMMs – How they work
- NVDIMMs – System Considerations
- NVDIMM Performance
- NVDIMM Ecosystem
- Summary

The Memory / Storage Problem: Latency

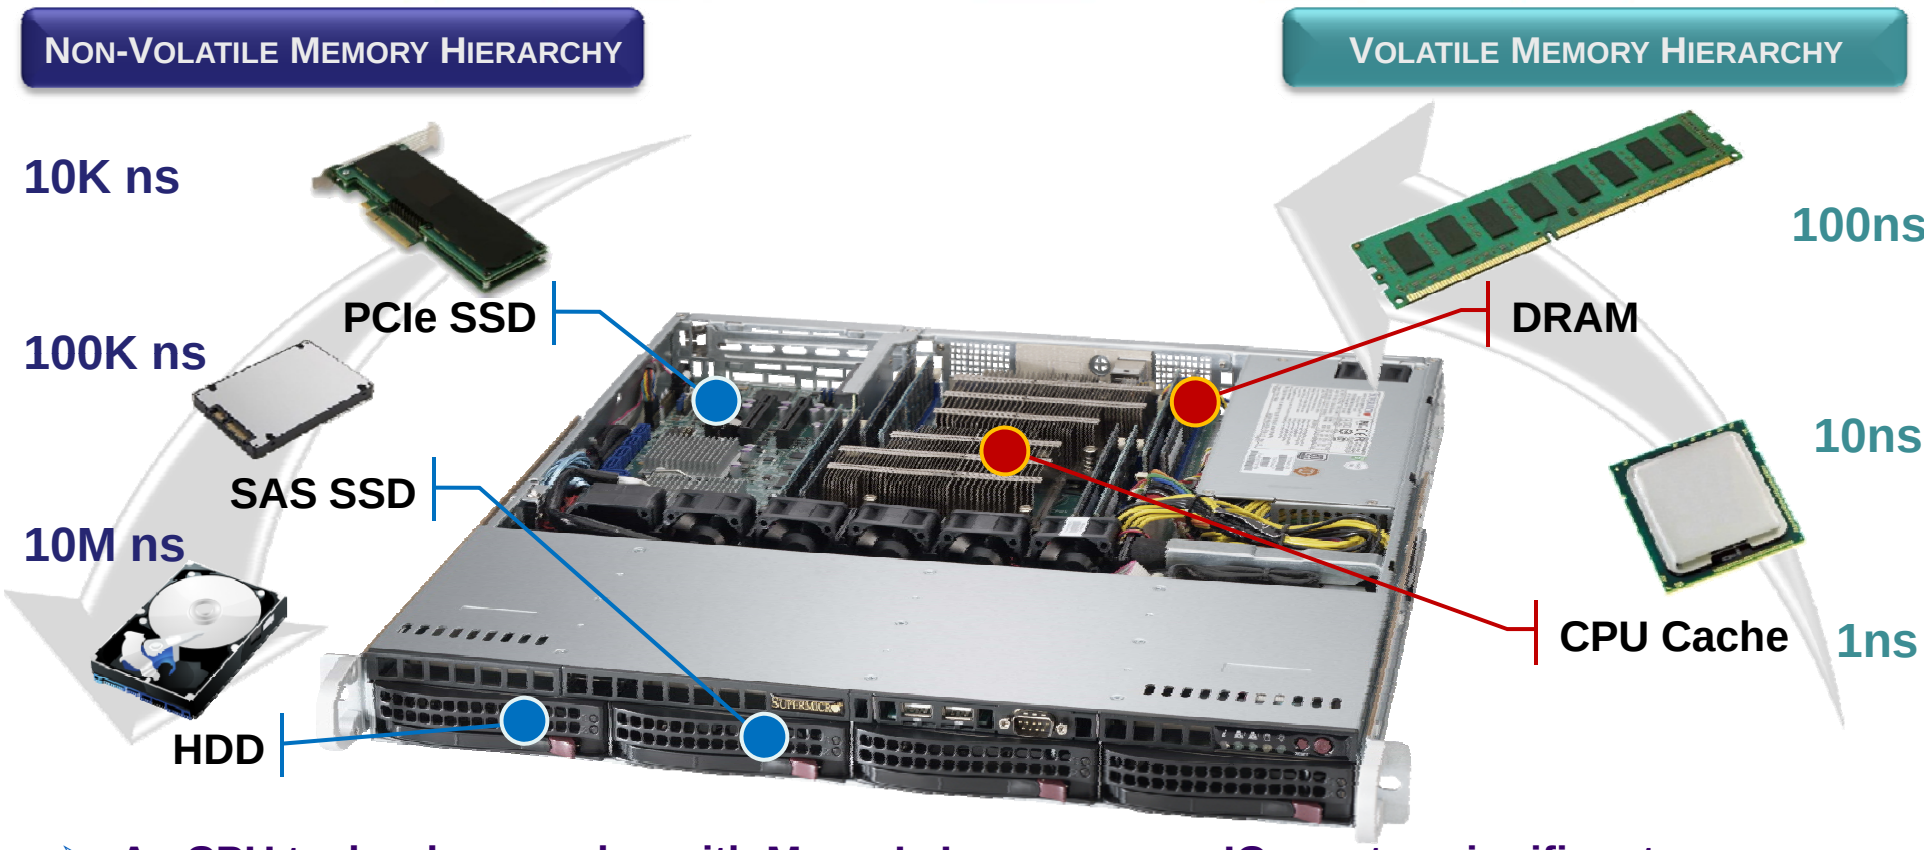

- As CPU technology scales with Moore's Law, memory IO creates significant performance bottlenecks
- The latency gap in memory / storage hierarchy needs to be bridged
- NVDIMM offers a solution today (Storage at DRAM Latency)

Memory/Storage Hierarchy

- Data-Intensive Applications Need Fast Access To Storage
- Large Performance Gap Between Main Memory And HDD
- SSDs Have Narrowed The Gap, But a Big Gap Still Exists
- Until an “SCM” becomes viable for mainstream adoption (2020?)

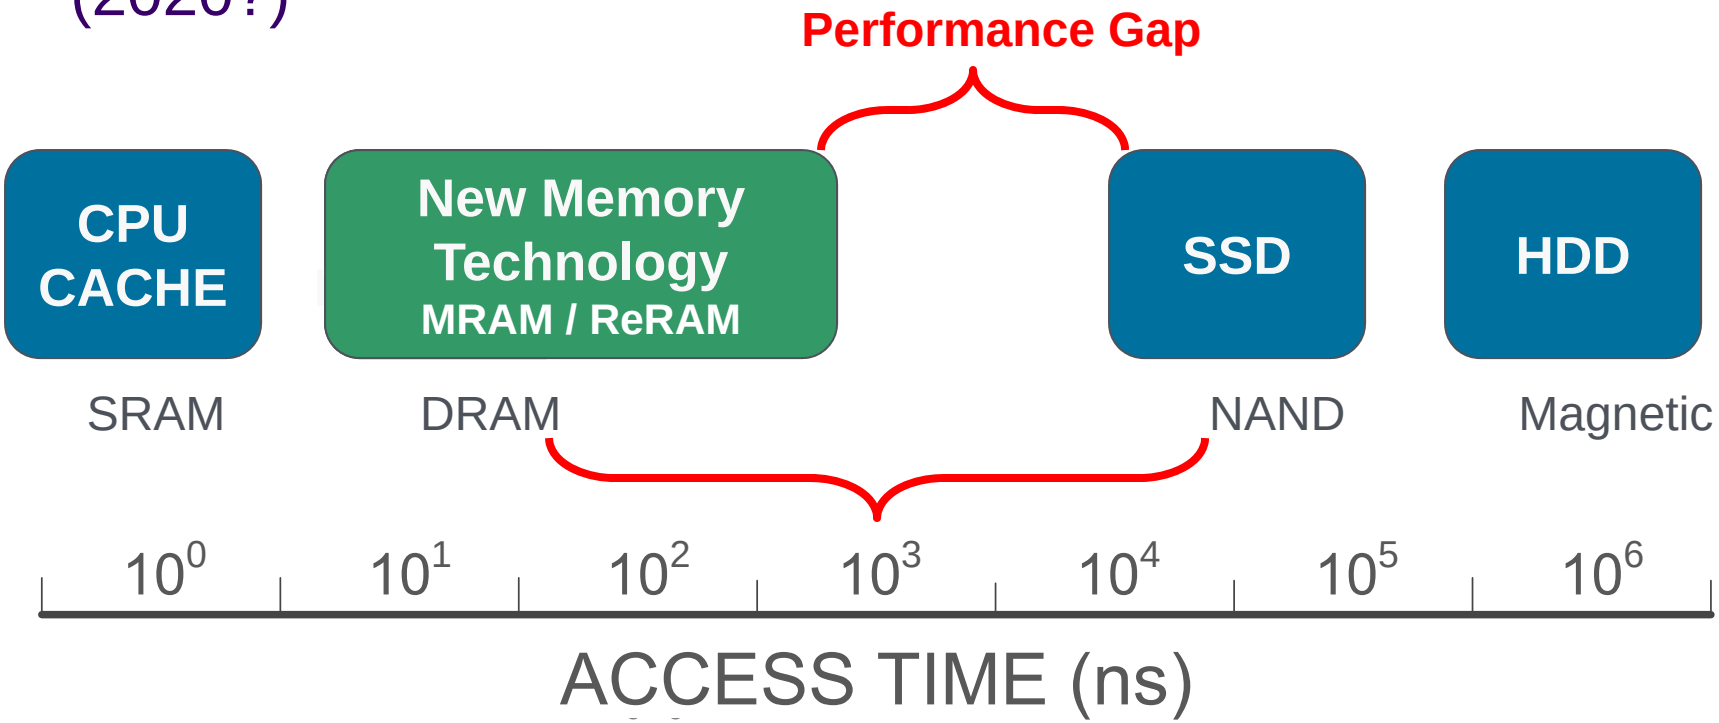

MAIN MEMORY ROADMAP

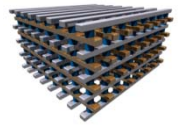

Volatile DRAM

Non-Volatile Memory

- MRAM
- PCM
- ReRAM

SCM

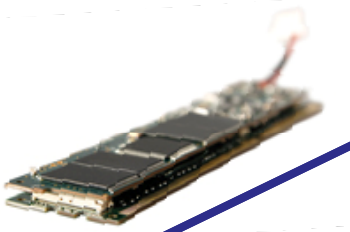

NVDIMM

DDR4

DDR3

- Scaling Problems
- MFG Concerns <2x nm

Relative Capability

2012 2014 2016 2020

WHAT THE INDUSTRY WANTS FROM MEMORY (THE HOLY GRAIL)

- Infinite Endurance
- Lowest Latency
- High Capacity
- Non-Volatile
- Low Power
- Scalability
- Low Cost

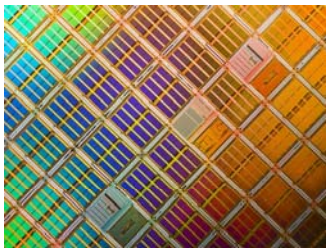

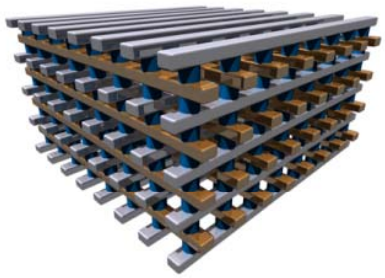

NVDIMM DOES MOST OF THIS TODAY…

NVDIMMs – What and Why

- Reside on the Memory Channel (DDR3/DDR4)
- Retain data in the event of an unexpected power loss
- Combines mature memory technologies (DRAM and Flash)
- Requires independent power source to ensure persistence
- Fits well with the NVM Programming Model (as precursor to SCM)
- Delivers new levels of storage performance
- Databases can run faster and recover more quickly
- Can enhance both SSD endurance and reliability

HOW IT WORKS

NORMAL OPERATION:

1. During normal operation, the NVDIMM appears like a standard DDR3 DRAM module.
   - » DRAM Latency (nanoseconds)
   - » DRAM Endurance (practically infinite)
   - » DRAM Bandwidth (12GB/s per NVDIMM)

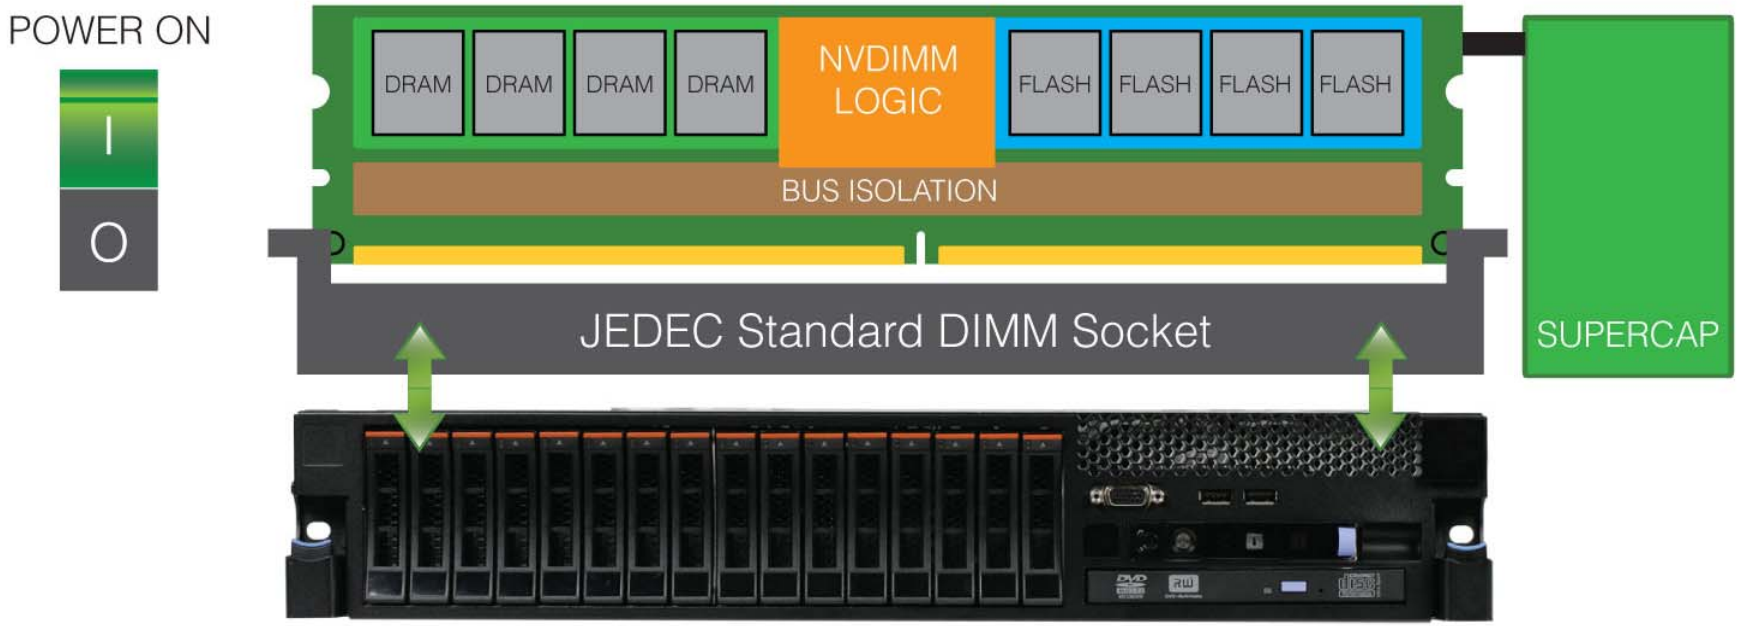

Shining Light on the DIMM Slot

HOW IT WORKS

POWER-FAIL - DRAM SAVE to FLASH:

1. **Power-Fail Event:** The NVDIMM isolated from the BUS. All data (incl. ECC) in the DRAM is SAVED to onboard Flash via the NVDIMM Logic. Super Capacitors provide hold up power to the module during this operation.
2. When the SAVE completes. The NVDIMM module is then shut down.

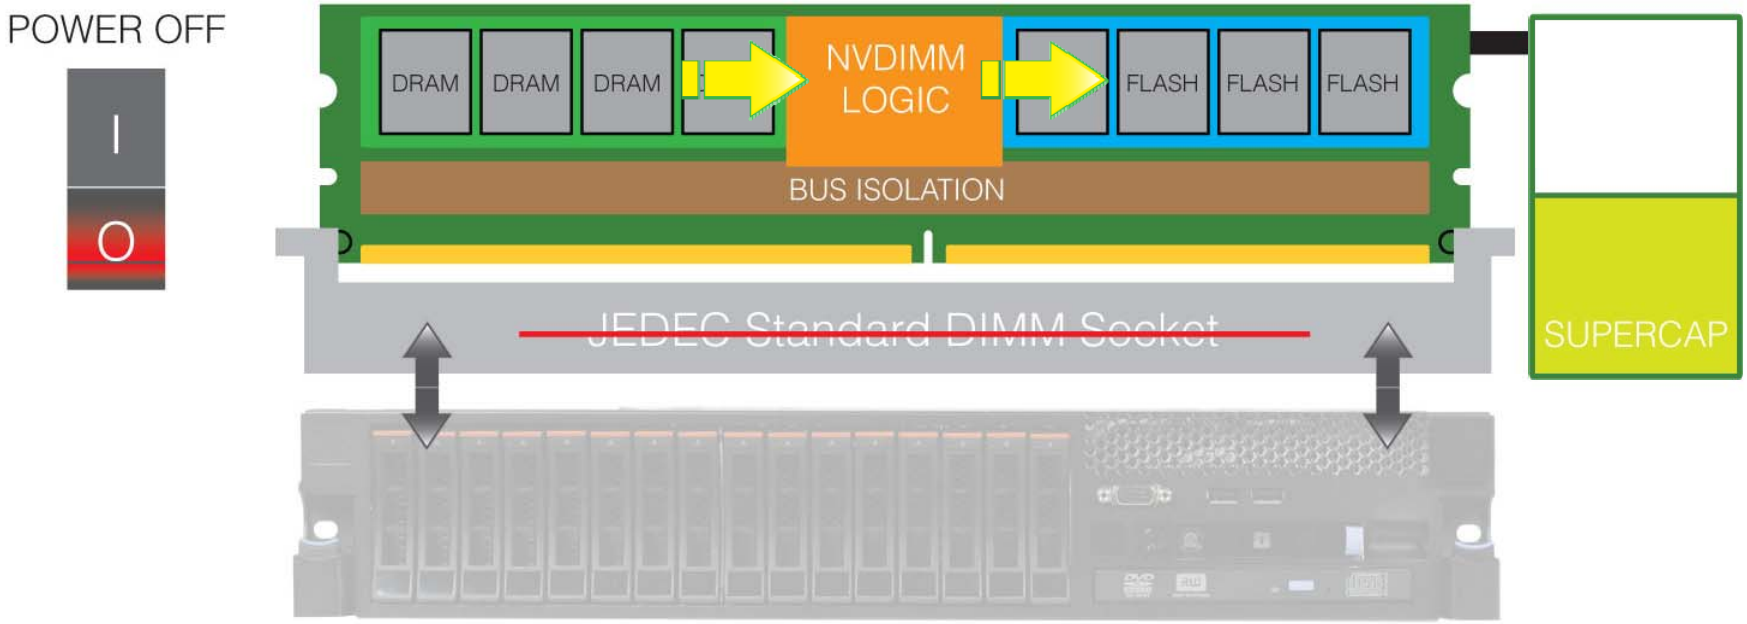

Shining Light on the DIMM Slot

HOW IT WORKS

POWER RESUME:

1. When power to the system is returned, Super Capacitors are re-charged & the data is RESTORED back from FLASH into the DRAM.

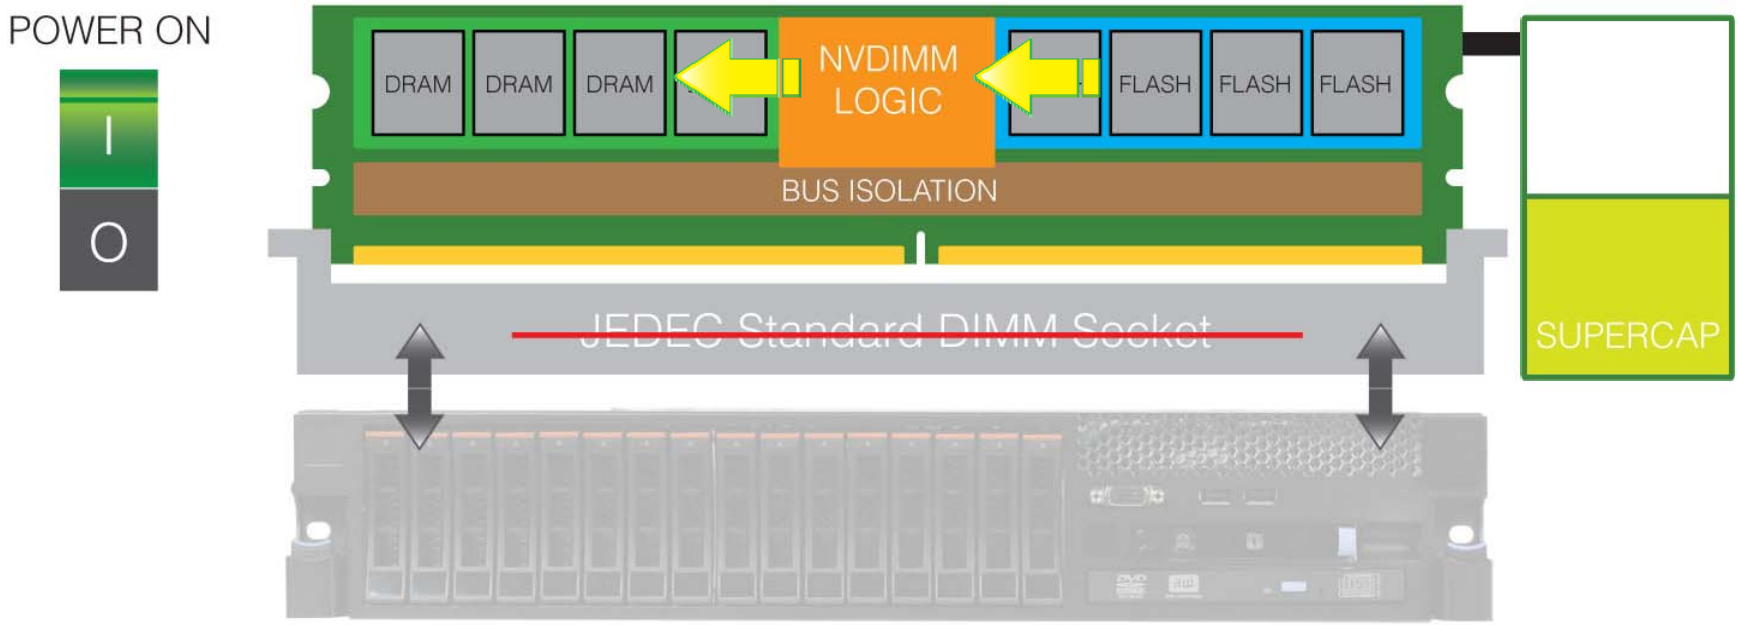

Shining Light on the DIMM Slot

HOW IT WORKS

NORMAL STATE:

1. Once all data is RESTORED back into the DRAM, the NVDIMM is ready for I/O transactions with the host system.
2. Host system finishes BOOT and normal NVDIMM operation continues.

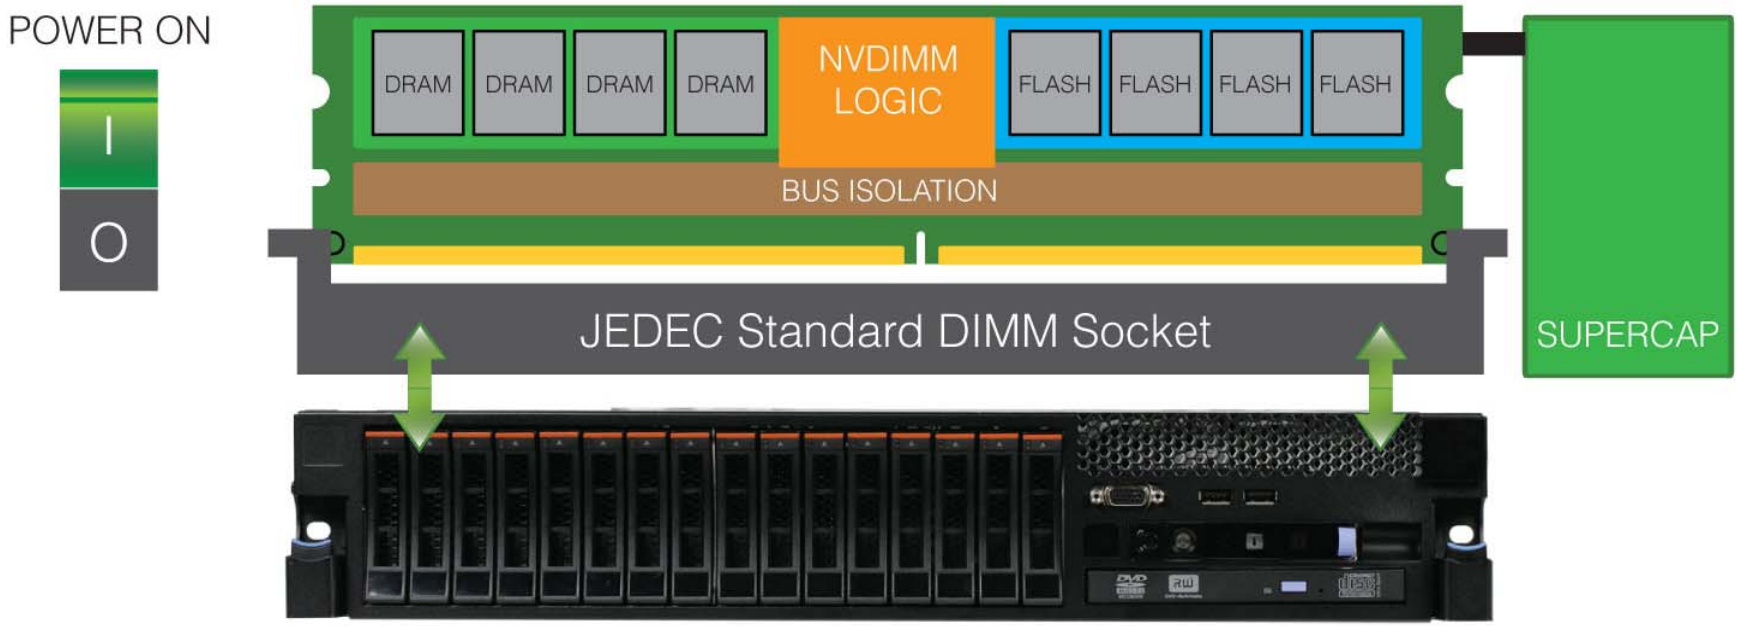

Shining Light on the DIMM Slot

System Considerations

The “Pieces of the Puzzle” that are required for NVDIMM Integration

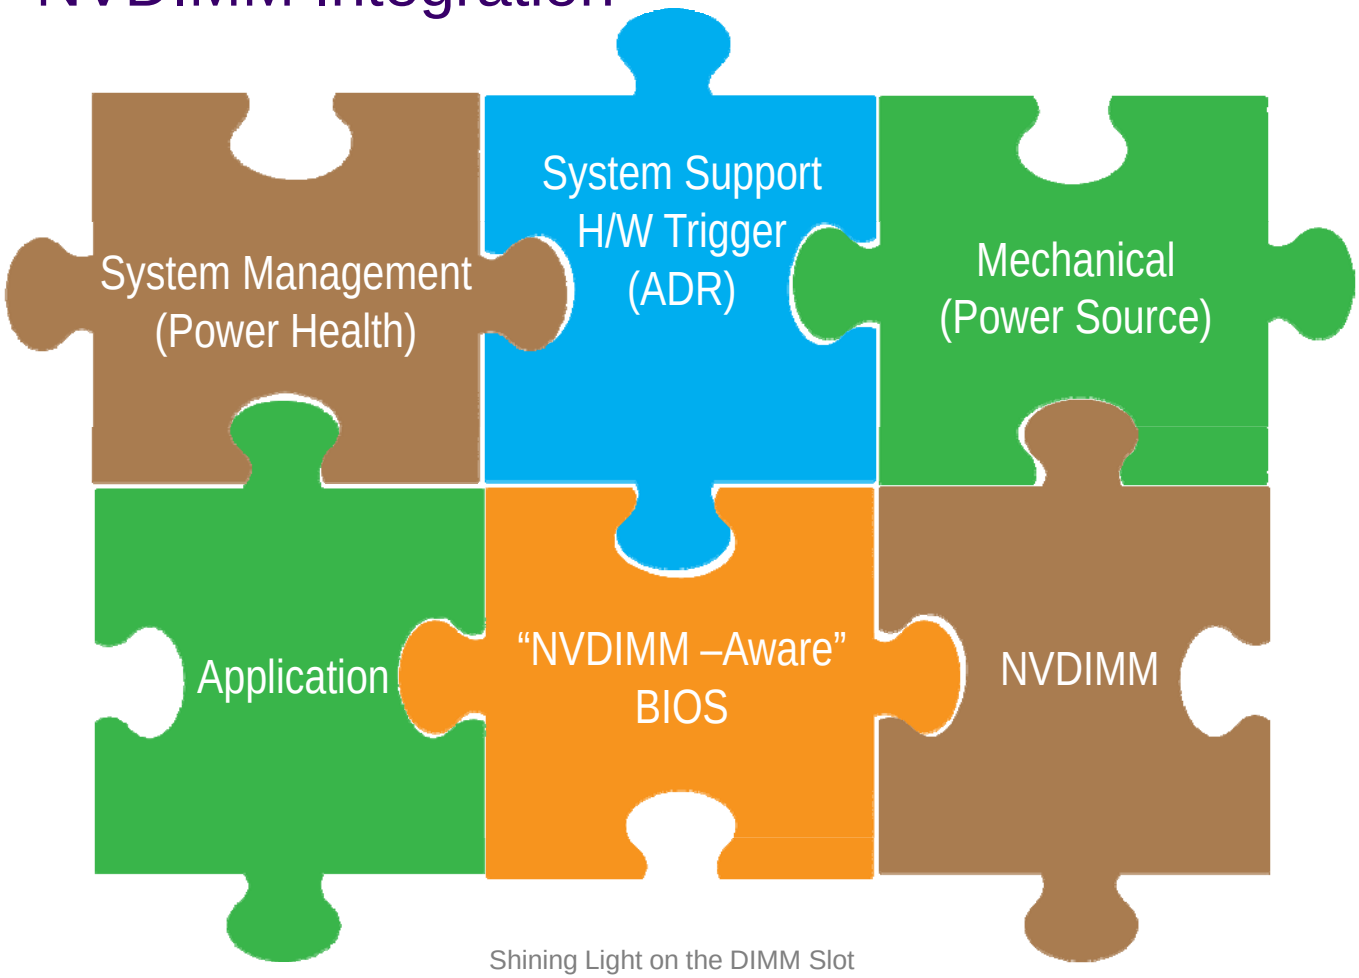

THE COST OF HIGH LATENCY

amazon.com

"…every 100ms of latency cost them 1% in sales"

Google

"…an extra 500ms in search page generation time dropped traffic by 20%"

TABB GROUP

"…a broker could lose $4M per millisecond if their electronic trading platform is 5ms behind the competition"

Source: http://highscalability.com/

STORAGE: LATENCY & CAPACITY

Ecosystem performance gap between compute & storage

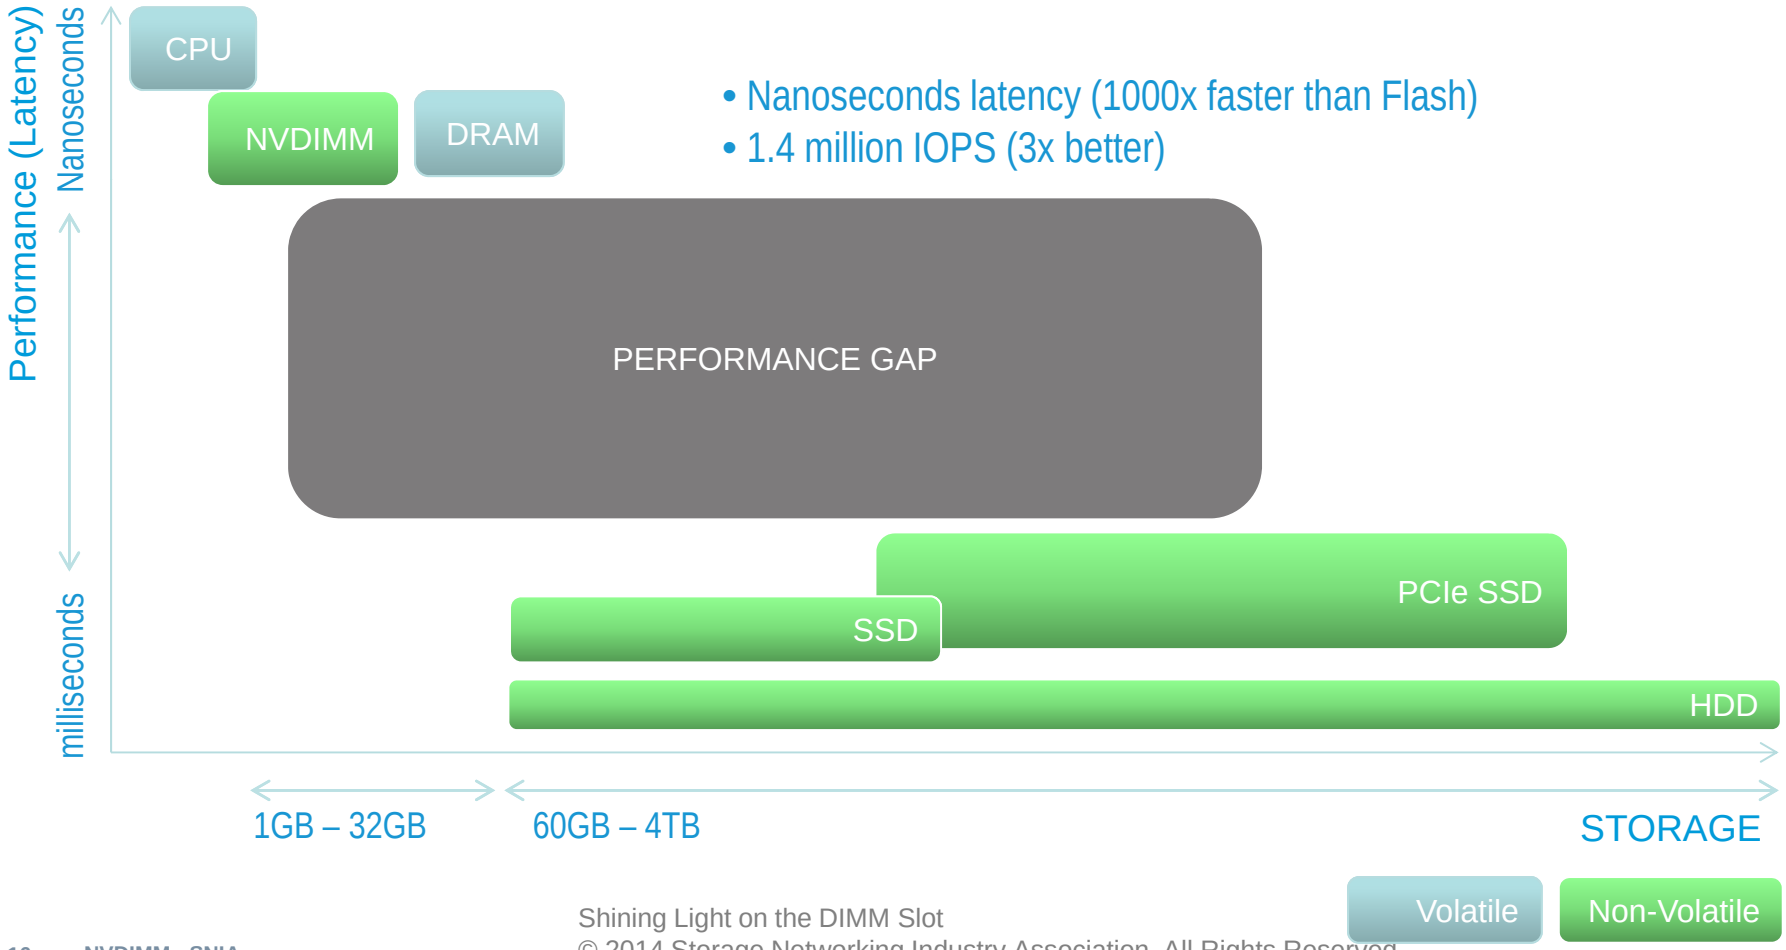

Example of NVDIMM Performance (BANDWIDTH – GB/sec)

Benchmark:

VDBENCH, Platform: Intel Sandybridge, Linux, Two DDR3-1333 NVDIMMs as interleaved pair (channel interleaving),

PRAMFS vs. SATA SSD as Linux block device

$ PER I/O : A NEW STORAGE METRIC ?

Performance vs. Cost ($ per I/O) trade-off

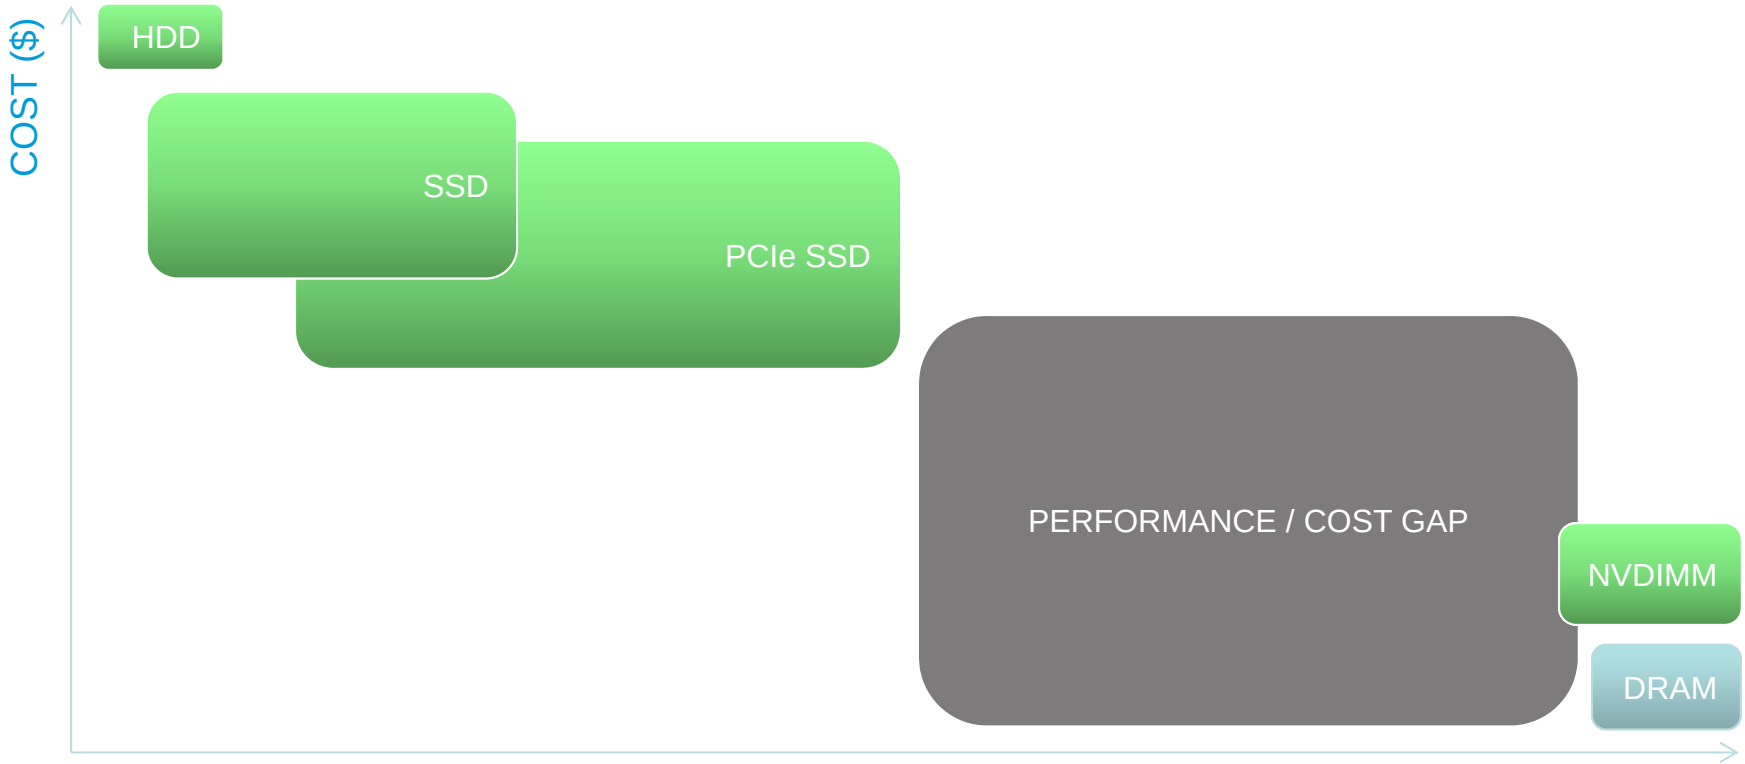

*Cost per PB written:
Best in Class SSD: $100.00 / PB versus. NVDIMM: $0.40 (250x cost savings)
Reference: http://www.vikingtechnology.com/uploads/NVDIMM_Technical_Comparison.pdf

NVDIMM ADOPTION

The flow of NVDIMM evolution and adoption

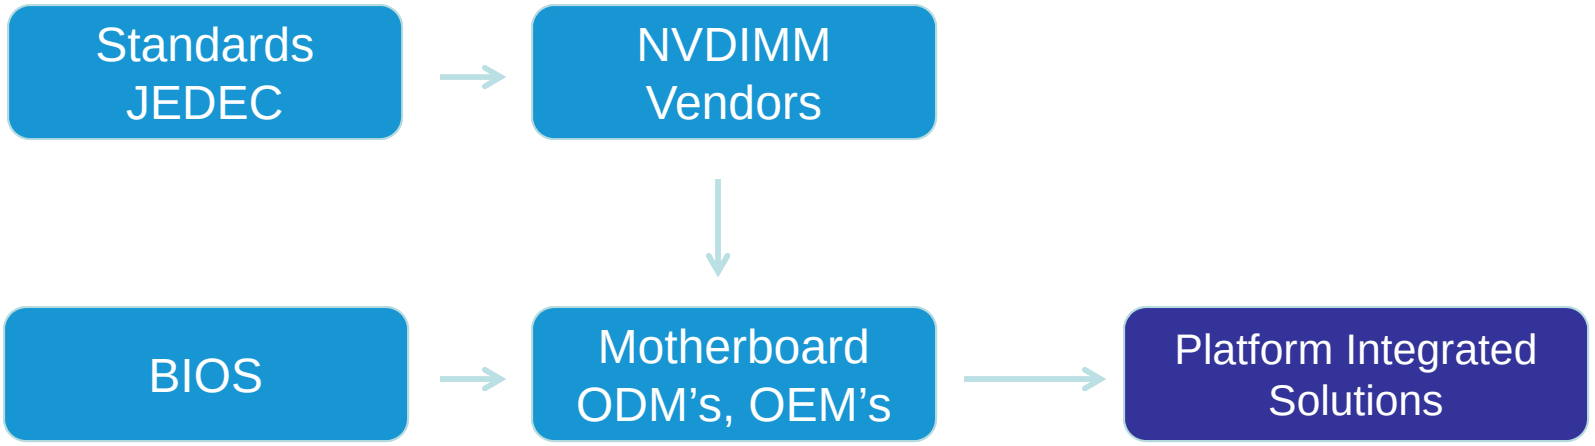

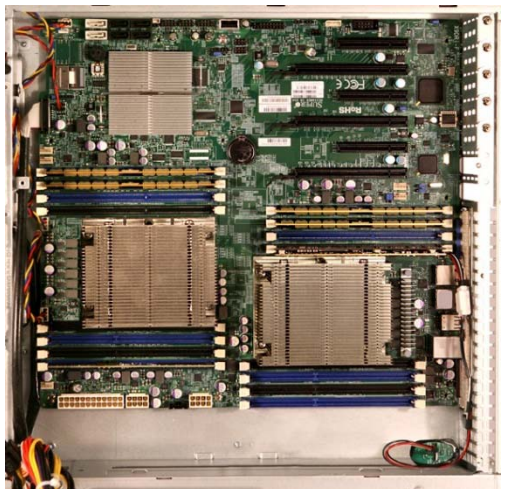

Shining Light on the DIMM Slot

Innovation and Alternate Memory Bus Device Options….

The NVDIMM-SIG is in the process of developing taxonomy to help the industry distinguish product categories

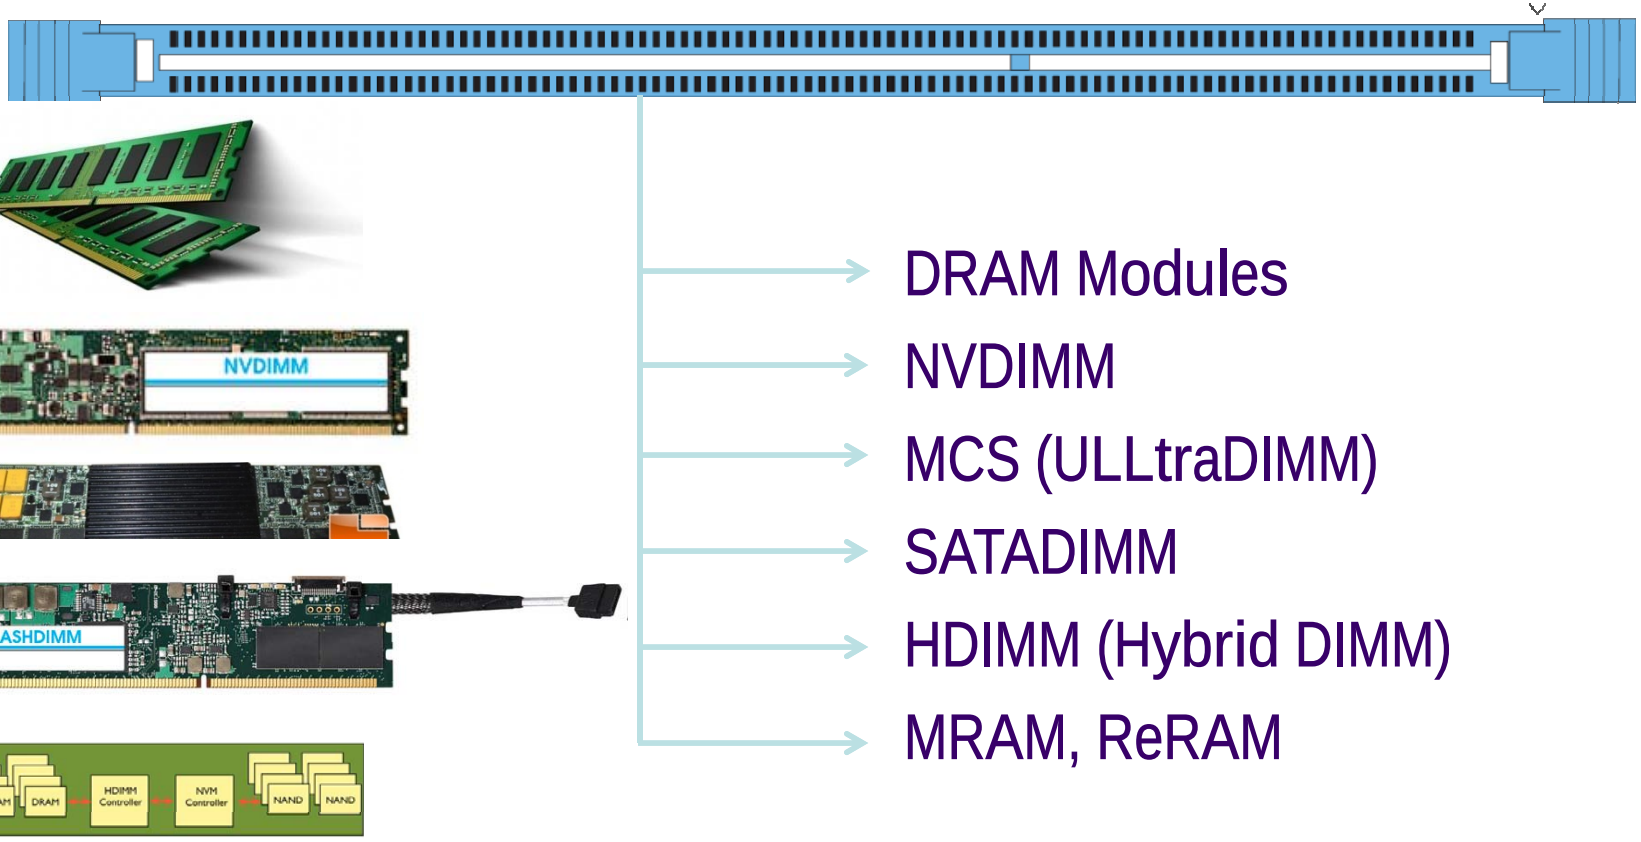

- DRAM Modules
- NVDIMM
- MCS (ULLtraDIMM)
- SATADIMM
- HDIMM (Hybrid DIMM)
- MRAM, ReRAM

Options – Flexibility – PRO's & CON's

- Flash is cheaper than DRAM $/GB
- NVDIMM has 1000x lower latency than Flash
- DRAM has practically infinite endurance

- Hyperscale want "Dense & Cheap" (WORM)
- Financial want low & predictable latency
- Storage wants higher I/O performance & increased data security
- No individual "BEST" choice – There are OPTIONS….

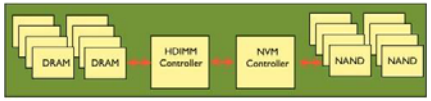

The Answer – Of Course…Is..

Standard Servers become Highly Flexible

There is a solution for whatever the Application Demands

- Highest Performing Storage - NVDIMM
- High Capacity Flash – PCIe SSD
- Lower Latency SSD – ULLtraDIMM

NVDIMM SIG

- NVDIMM Special Interest Group (SIG) formed January 2014
  - Organized under the SNIA Solid State Storage Initiative to help:
    - Accelerate awareness and adoption of NVDIMMs
  - Vendors collaborate to broaden industry support and knowledge
- SNIA's history of developing standards and providing education:
  - The NVM Programming Model Technical Working Group
  - Ideal venue for NVDIMM SIG support
- NVDIMM SIG will:
  - Educate on how system vendors can design in NVDIMM
  - Communicate industry standards as they evolve
  - Develop market understanding of NVDIMM technology
  - Communicate how new programming models help deliver value

Attribution & Feedback

The SNIA Education Committee would like to thank the following individuals for their contributions to this Tutorial.

Authorship History

Original Author:
Adrian Proctor 3/2014

Additional Contributors

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***