STORAGE INDUSTRY
# SUMMIT

The Future of Computing:
The Convergence of Memory and Storage through Non-Volatile Memory (NVM)

JANUARY 28, 2014, SAN JOSE, CA

Tom Coughlin, Coughlin Associates
& Jim Handy, Objective Analysis

# Why Hybrid Storage Strategies Give the Best Bang for the Buck

SNIA
Solid State Storage Initiative

# Outline

- Different applications have different storage requirements
- Storage and memory trade-offs
- Hybrid hard disk drives
- Flash with magnetic tape
- Hybrid storage systems
- The evolution and impact of non-volatile memory
- Conclusions

# IOPS Survey: Application Types

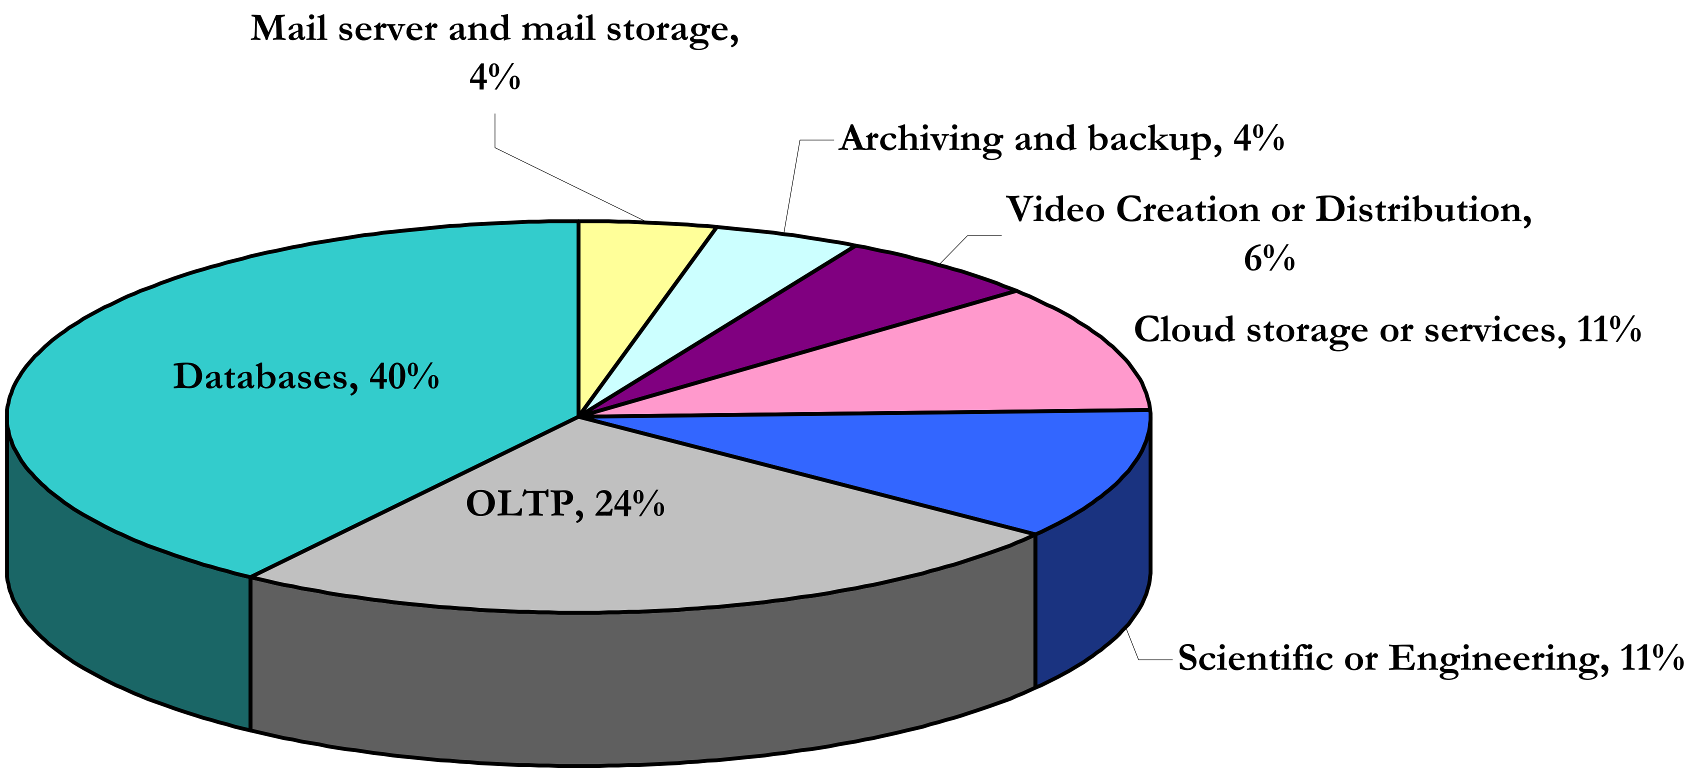


How Many IOPS is Enough Report, T. Coughlin and J. Handy

# How Many IOPS Does This App Need?

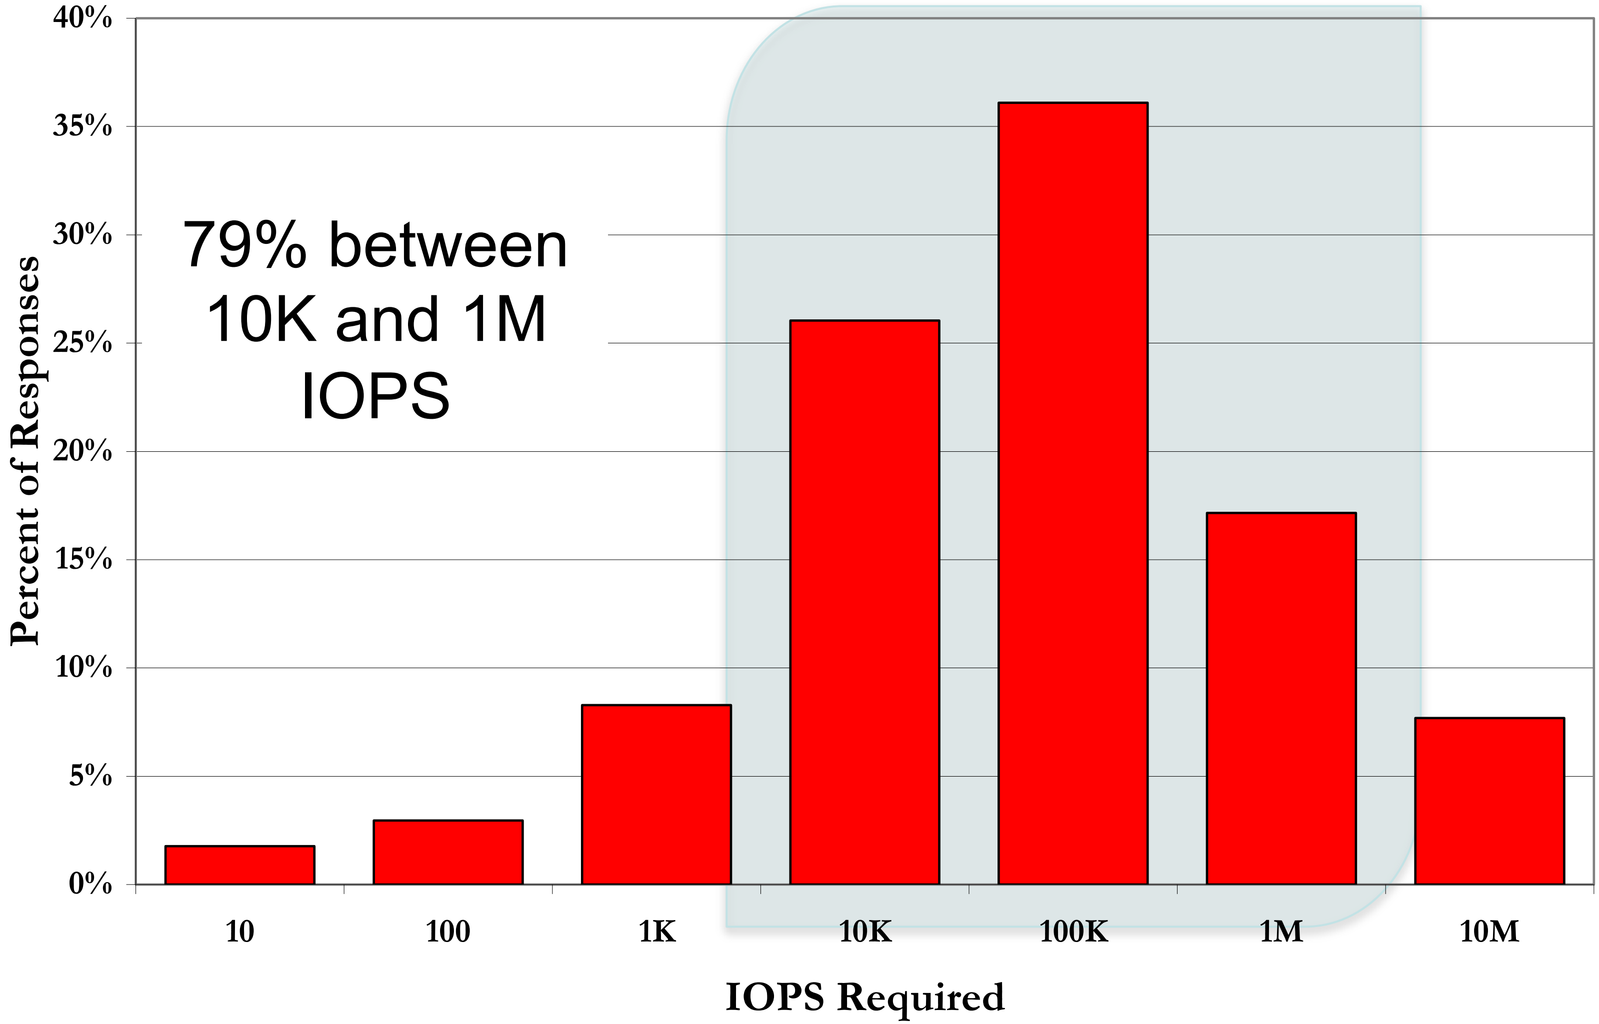

# How Much Storage Does the App Need?

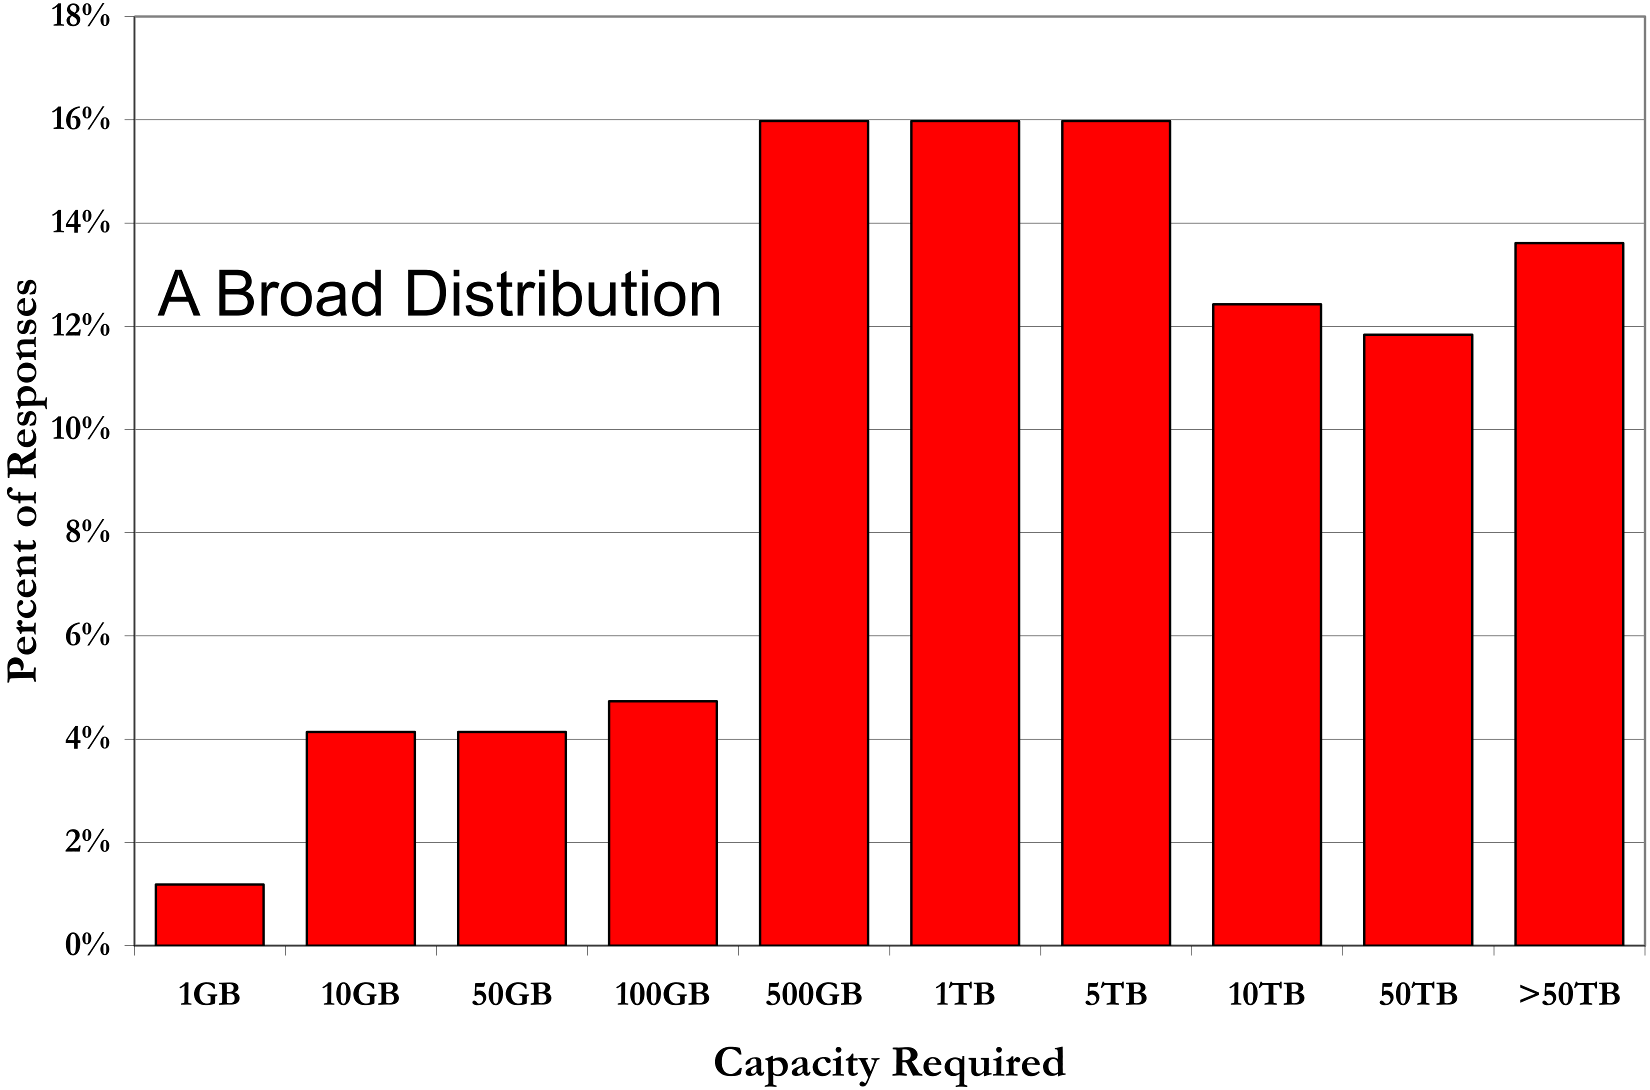

# Latency Requirement

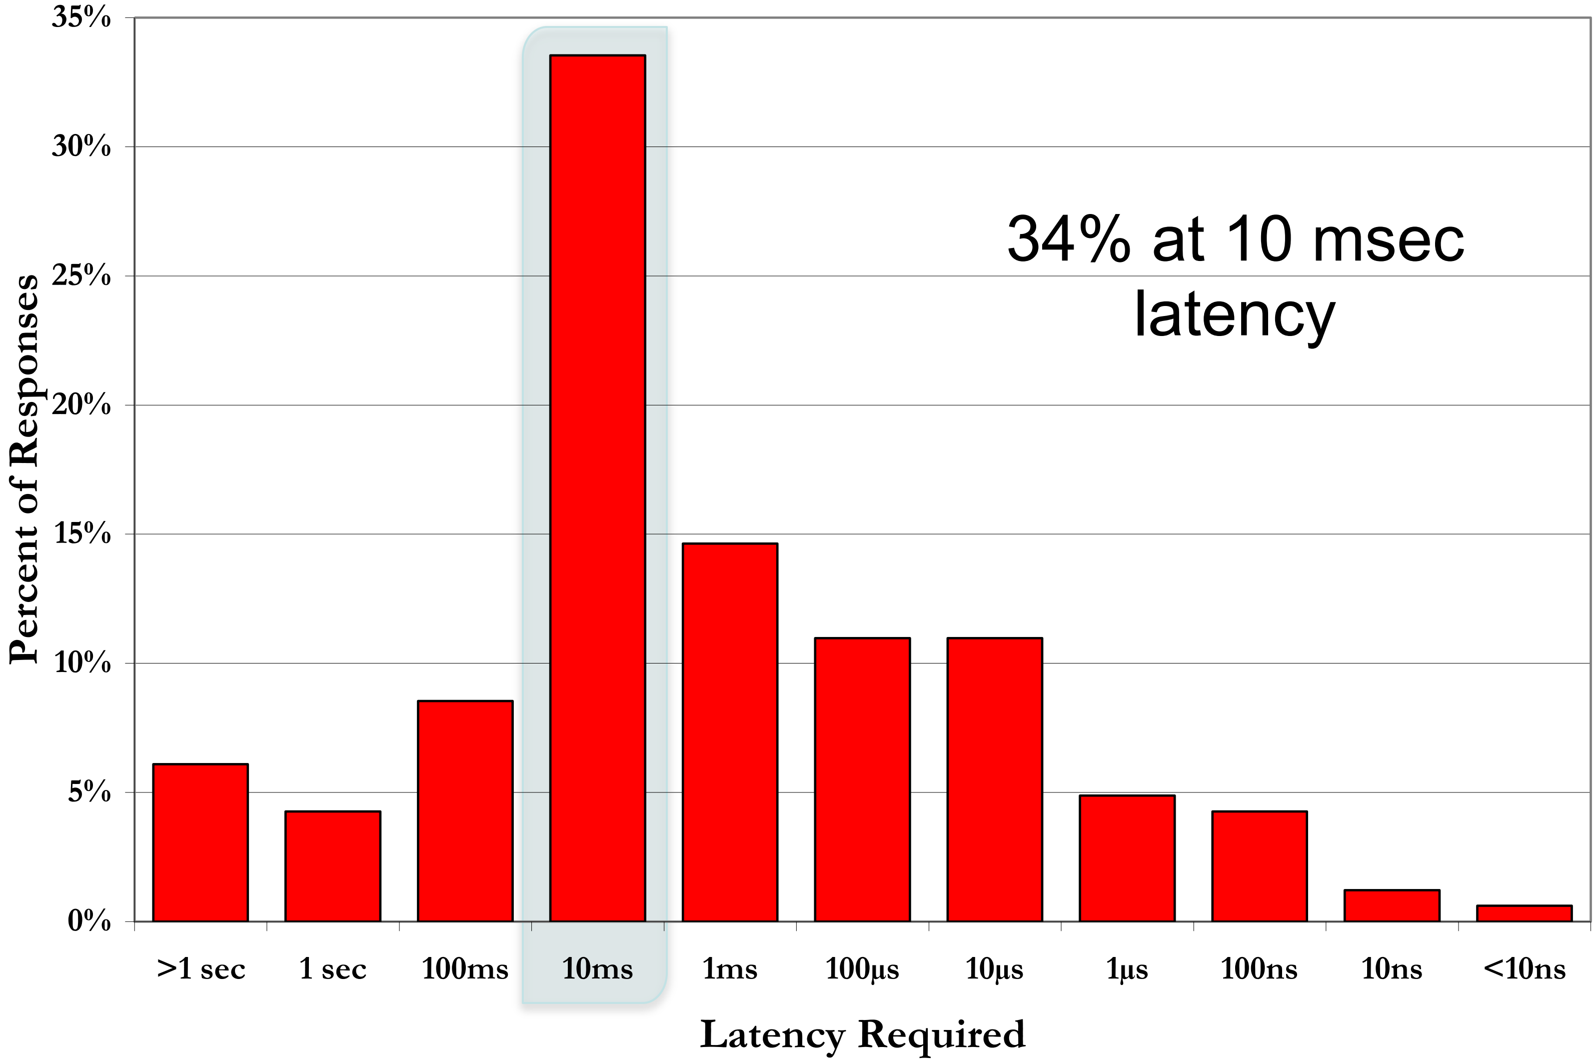

# We Need Multiple Storage Devices

- Designers must trade off price and performance
  - “What’s the least expensive combination that meets my needs?”
- Different storage devices combine to balance cost vs. capability
- This leads to interesting combinations of different memory and storage technologies in order to achieve design objectives
- NV memory makes distinction of storage and memory fuzzier

# Devices for Modern Storage

- Hard disk drives
  - Cold Storage Drives with SMR and He-filled drives promise 5-6 TB models by 2014
  - Hybrid HDDs, as thin as 5 mm
  - New interfaces—Kinetic from Seagate
- Magnetic data tape
  - LTO 5,6 with LTFS
  - Oracle 8.5 TB with LTFS
  - Object based tape
- Flash Memory
  - Sub-20 nm SSDs
  - 15-16 nm flash in 2014


# Storage/Memory Hierarchy

- Applications can have various performance and storage capacity requirements
- A trade-off between $/GB or $/IOPS depending upon the application
- Trade-offs also possible between memory and storage
- New solid state technology offer possibility of persistent memory (e.g. MRAM)

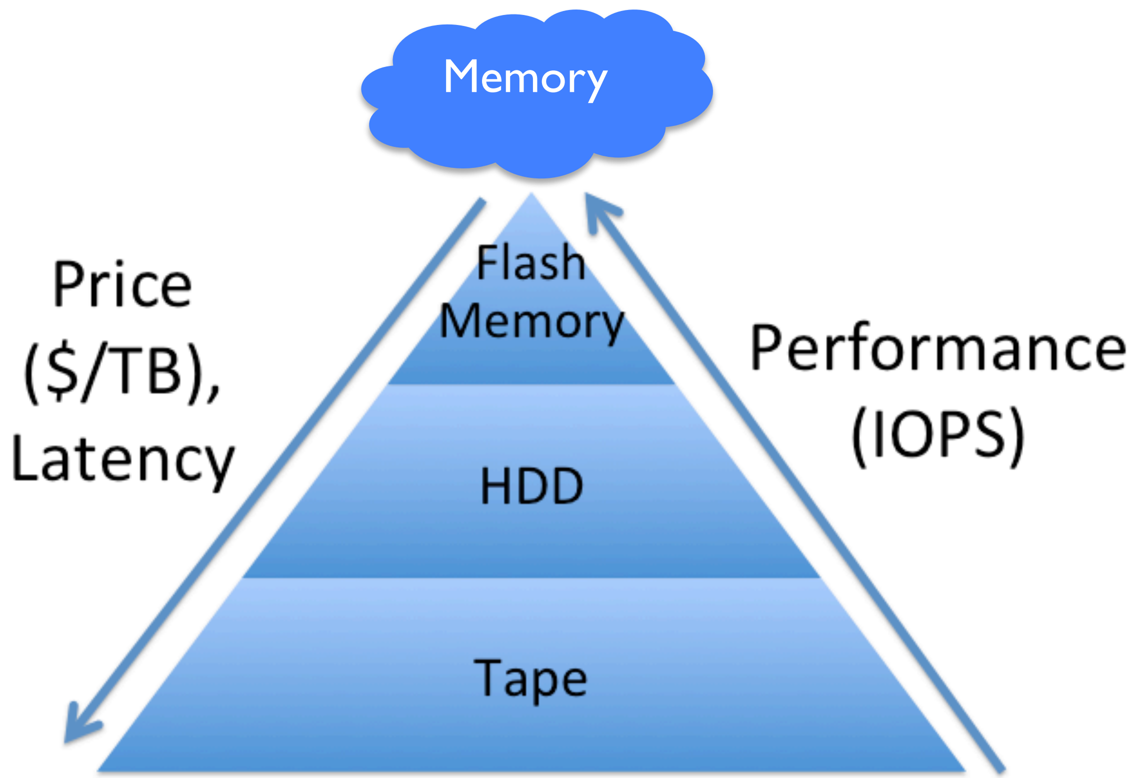

# Memory/Storage Price vs. Bandwidth

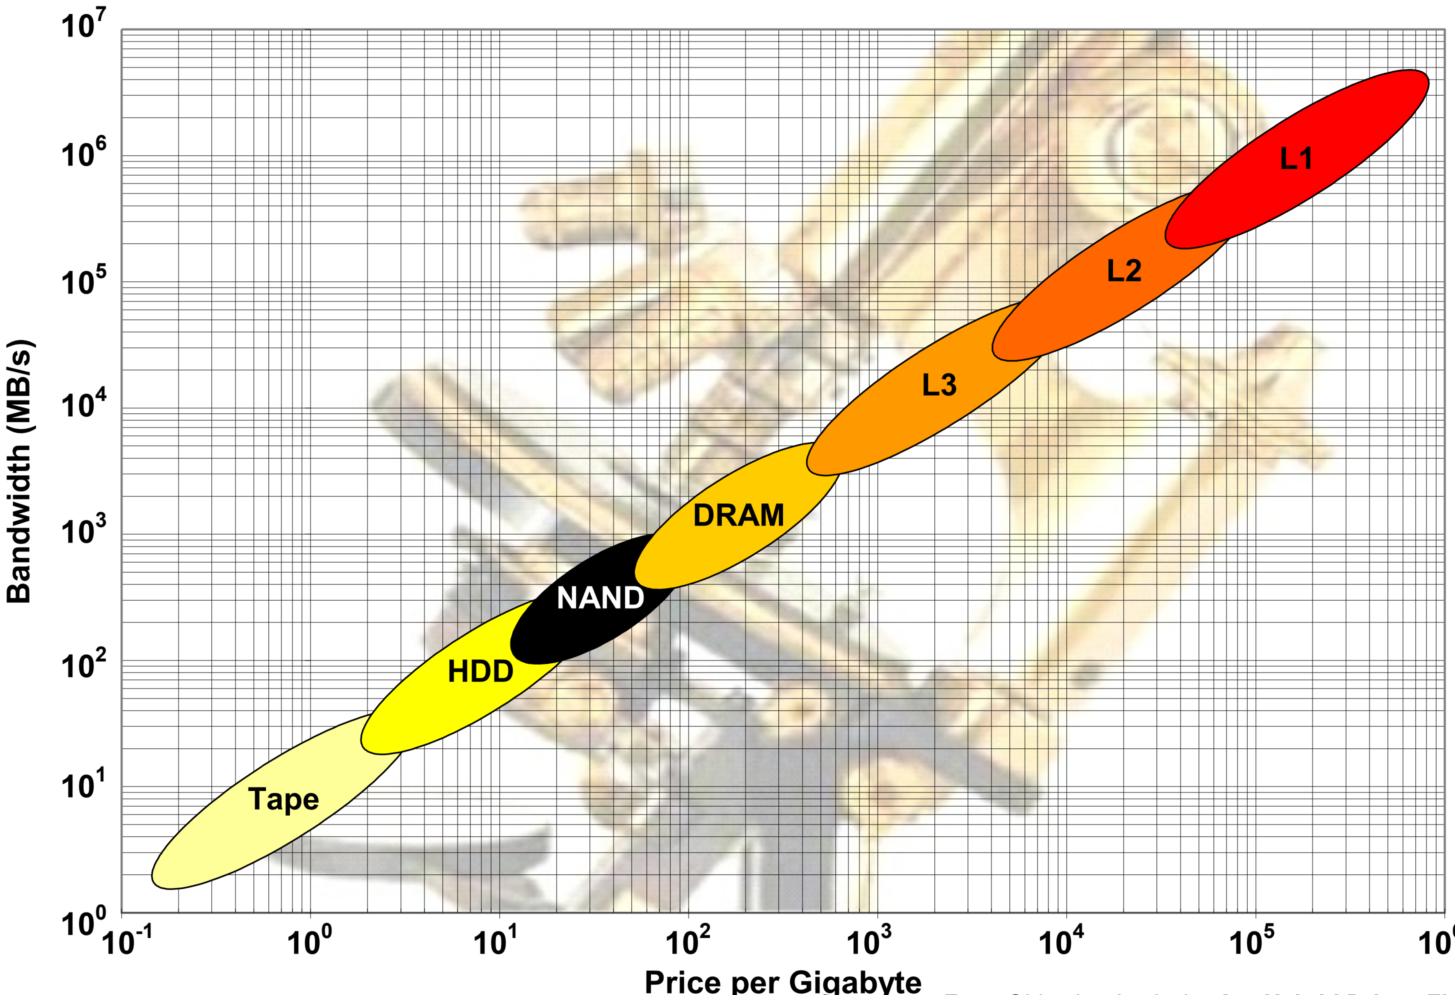


From Objective Analysis: ***Are Hybrid Drives Finally Coming of Age?***

# Hybrid HDDs (SSHDDs)

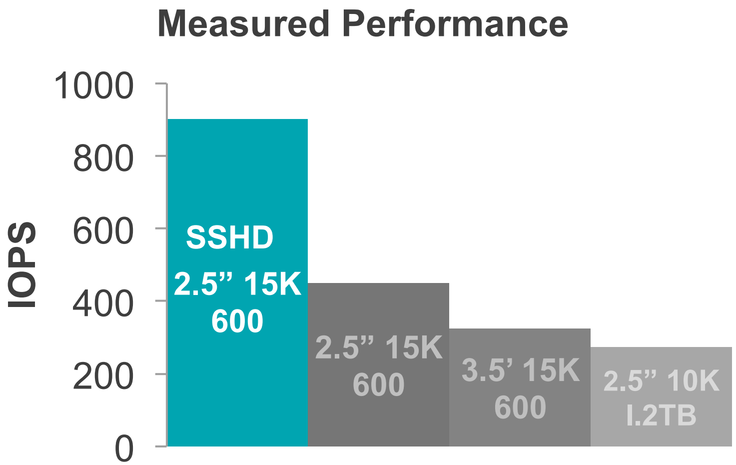


From Dave Anderson, Seagate, SV 2014

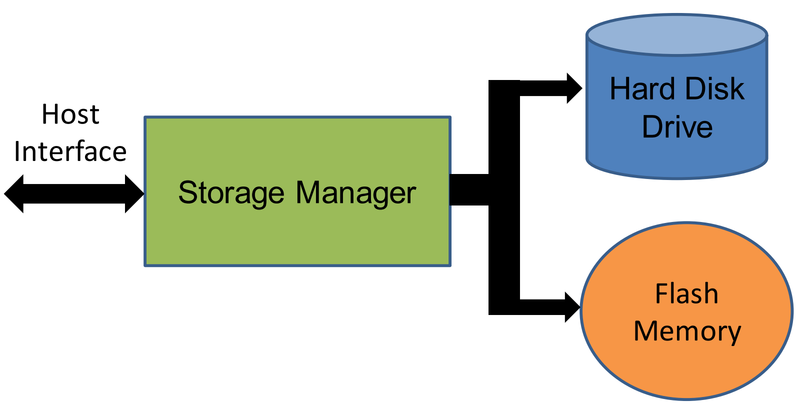


- Adds flash to HDDs
  - Performs like flash
  - Capacity of an HDDs
- Available from all HDD vendors
- Both client (2.5-inch for notebook market) and enterprise models

# Dual Drives—Combining HDDs and SSDs

- Separate HDDs and SSDs can also be used together to improve overall system performance
- The SSD can contain the operating system and important programs while the HDD contains user data
- Apple's Fusion Drive as well as WD's Dual Drive that contains a 120 GB SSD with a 1 TB HDD in a single package are examples of this
- Also many people have bought and installed separate SSDs and HDDs in their computers
- This is another way to achieve performance without sacrificing storage capacity

# How SSHDDs Work: 3 Memories Interacting

- **HDD Media: GBs to TBs**
  - Storage for all data
  - Primary copy of all cached data
- **DRAM: 128MB**
  - Buffers all Reads
  - Caches all Writes
- **NAND: 32GB**
  - Read Cache for active data
  - Non Volatile Cache for DRAM write cache
  - NVC size = DRAM Write Cache

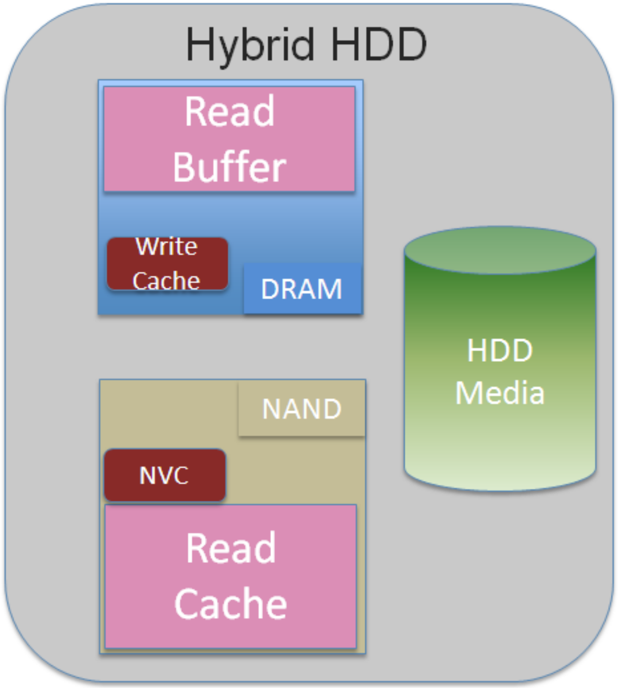


From Dave Anderson, Seagate, SV 2014

# Using Flash to Accelerate Reads

(Dave Anderson, Seagate at SV 2014)

1. First read from media thru DRAM
2. Next read from DRAM
3. Moved to cache if busy enough
4. Subsequent reads from Cache

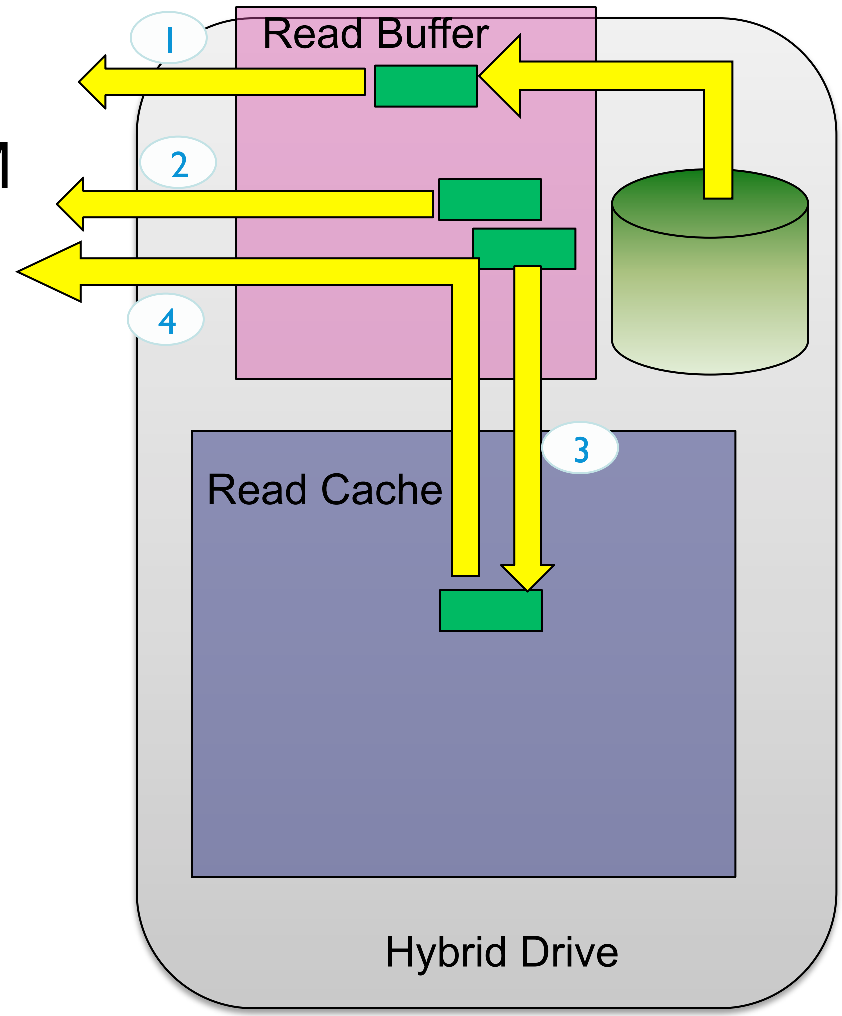






© 2014 Coughlin Associates & Objective Analysis

# Flash Accelerates and PROTECTS Writes

(Dave Anderson, Seagate at SV 2014)

1. System writes into DRAM write cache
2. Outstanding writes coalesced & written to media
3. If power is lost, back EMF powers writing cached data to NVC
4. On power up, NVC written to media

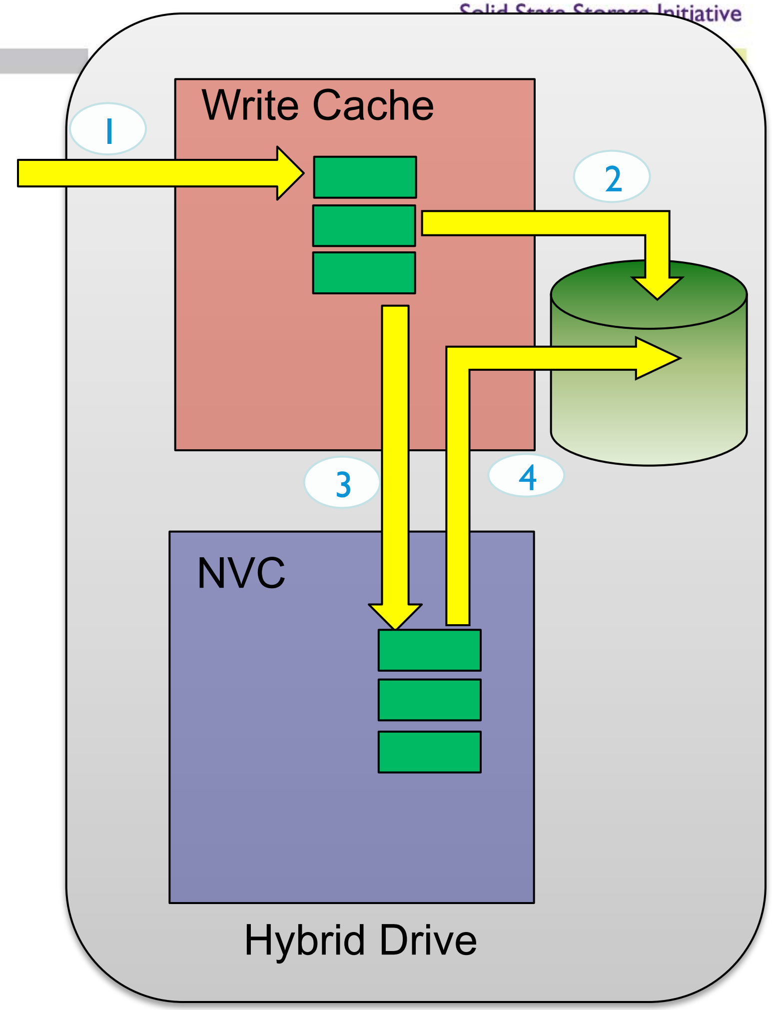


**All writes protected & preserved =**

**A new level of file system integrity!**

# SSDs and Tape


- Use of SSDs as front end to tape archive
- Running tape write and read speeds are very fast and flash memory is much better as a front end archive cache than HDDs
- Xendata has also had SSD based tape front end products for the last couple of years

# Hybrid Storage at the System Level

All-flash systems are finding roles either by themselves or more often combined with HDD arrays to achieve trade-offs between performance and capacity.



# PCIe Flash Storage

SNIA™
Solid State Storage Initiative


- PCIe is becoming a critical interface
- Note only are there PCIe storage devices but also next generation storage interfaces such as SAS will be based on PCIe
- Direct connection technology, Thunderbolt now support 20 Gbps raw data rates

# Non-Volatile Memory

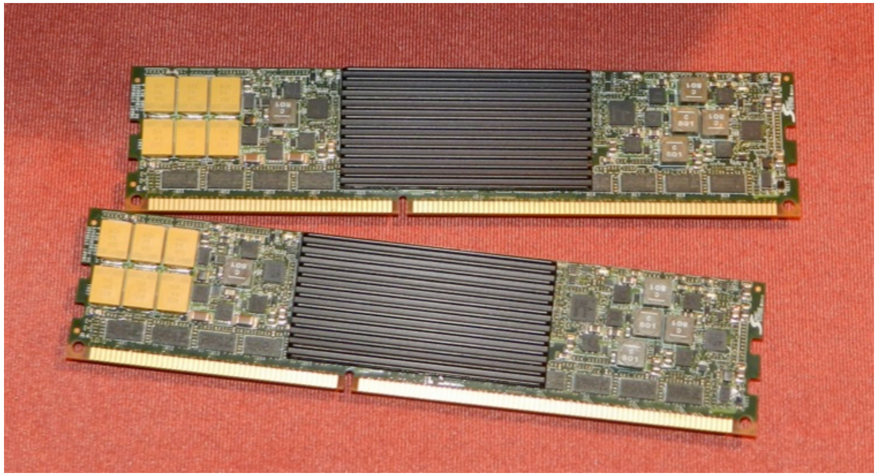


- NV Memory technology—e.g. ULLtraDIMM from SanDisk/ Diablo & SNIA Flash NVDIMM standard
  - Changing architectures of computers with NVMs,
  - Maybe even MRAM or ReRAM in future
- 3D Flash Memory 1TB capacities announced by Samsung

# NV-DIMM

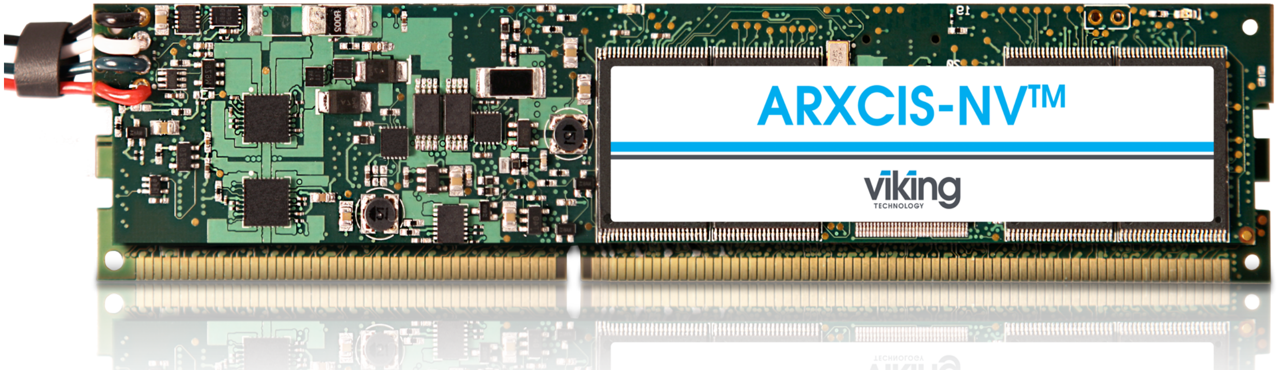


- A Non-Volatile DRAM Module (Fusion of both DRAM and NAND)
  - Resides on the DDR3/DDR4 channel. Standard Memory Interface.
  - Host sees the DRAM capacity (no direct access to the NAND)
- Leverages Beneficial Characteristics Of Each Memory Technology
  - Latency, Speed, endurance, and random byte addressability of DRAM
  - Non-volatility of NAND Flash
- Enables Main Memory Persistence
  - Data written to DRAM is preserved through system power loss
  - Main memory becomes non-volatile but operates at speed of DRAM

# NV-DIMM TECHNICAL COMPARISON

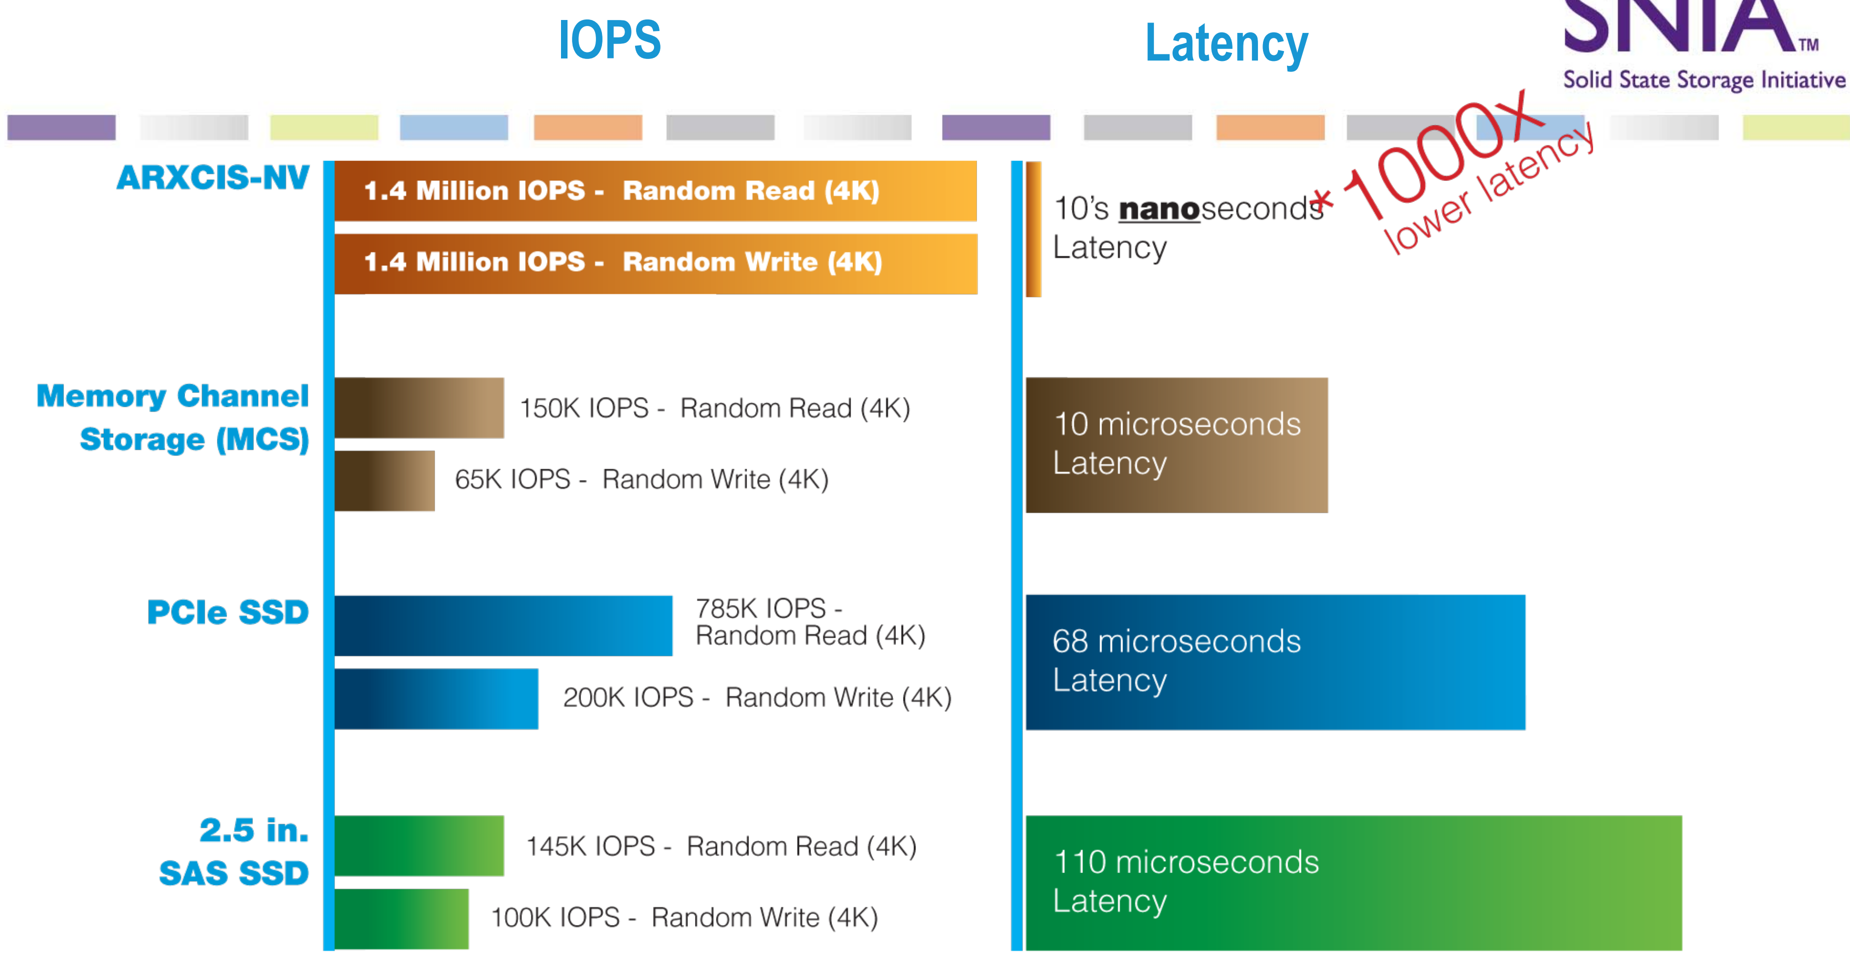


- 1000x speed up compared to NAND (latency)
- Rare – yet recently seen with HDD vs. SSD IOPS (~200 to 200k)

# Merging Storage and Memory

From 2014 Report on New Storage Technologies

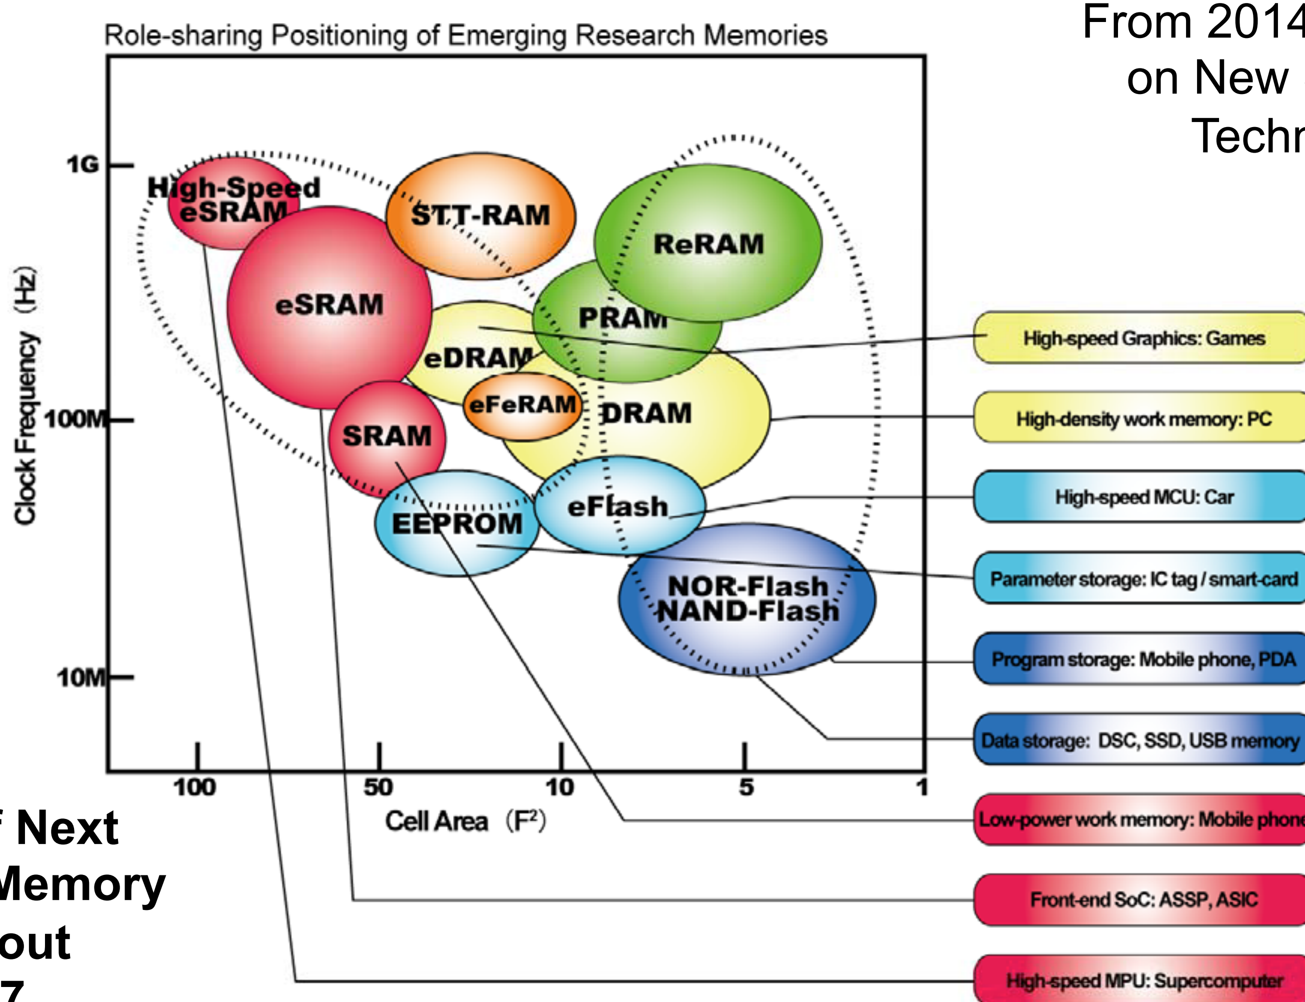


**Estimates of Next Generation Memory Market at about $91M by 2017**

# Next Generation Scalable NVM

(Jim Pappas, Intel at SV 2014)

## Resistive RAM NVM Options

Scalable Resistive Memory Element

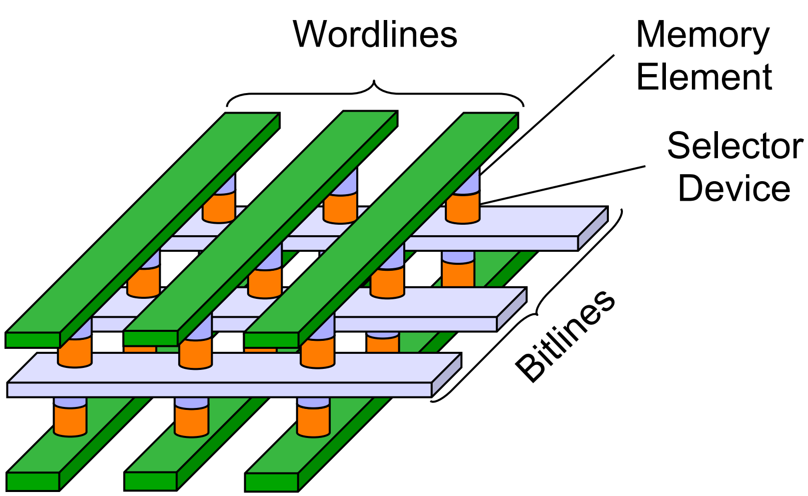


Cross Point Array in Backend Layers ~$4\lambda^2$ Cell

**~ 1,000x faster than NAND.**

| Family | Defining Switching Characteristics |
|---|---|
| Phase Change Memory (PCM) | Energy (heat) converts material between crystalline (conductive) and amorphous (resistive) phases |
| Magnetic Tunnel Junction (MTJ) | Switching of magnetic resistive layer by spin-polarized electrons |
| Electrochemical Cells (ECM) | Formation / dissolution of "nano-bridge" by electrochemistry |
| Binary Oxide Filament Cells | Reversible filament formation by Oxidation-Reduction |
| Interfacial Switching | Oxygen vacancy drift diffusion induced barrier modulation |

# Opportunities with Next Generation NVM

(Jim Pappas, Intel at SV 2014)

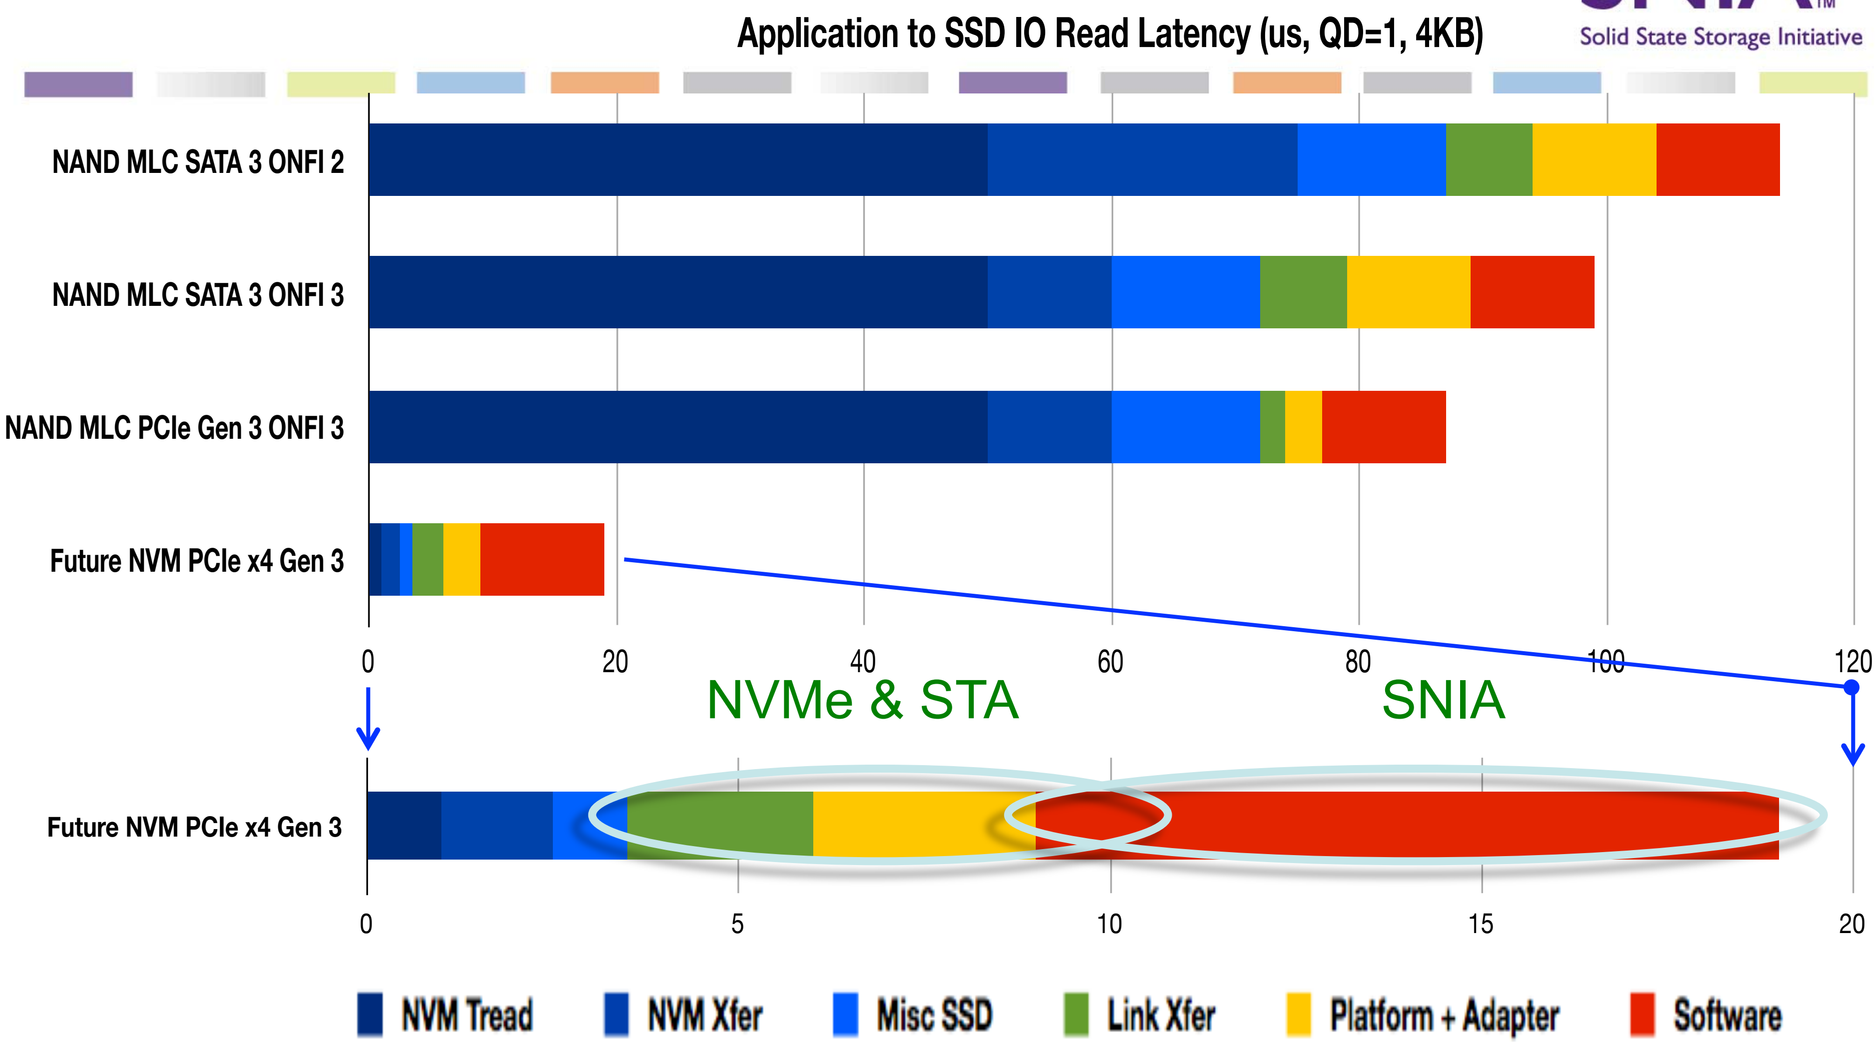


NVM Express/SCSI Express: Optimized storage interconnect & driver
SNIA NVM Programming TWG: Optimized system & application software

# Conclusions

- Different applications have different storage requirements
  - Drives different storage combinations to balance cost and performance
- Flash memory has enabled new hybrid architectures with other storage devices
  - Leads to hybrid HDDs and flash/tape solutions
- All-flash, combo flash, and HDD storage systems fill many important roles
- Non-volatile memory is changing the way computer systems are designed
  - Will be accelerated by new solid state non-volatile memory technology