Storage Decisions CHICAGO JUNE 11-12, 2013

TechTarget

Solid State Technology – What's New?

Dennis Martin, President, Demartek

www.storagedecisions.com

Agenda: Solid State Technology – What's New?

- Demartek – About Us
- Solid-state storage overview
- Types of NAND flash (SLC, MLC, TLC)
- NAND flash endurance and performance
- SSD Capacity – devices and storage systems
- SSD Caching
- Sharing of SSD technology
- Future non-volatile memory technologies
- References

About Demartek Video

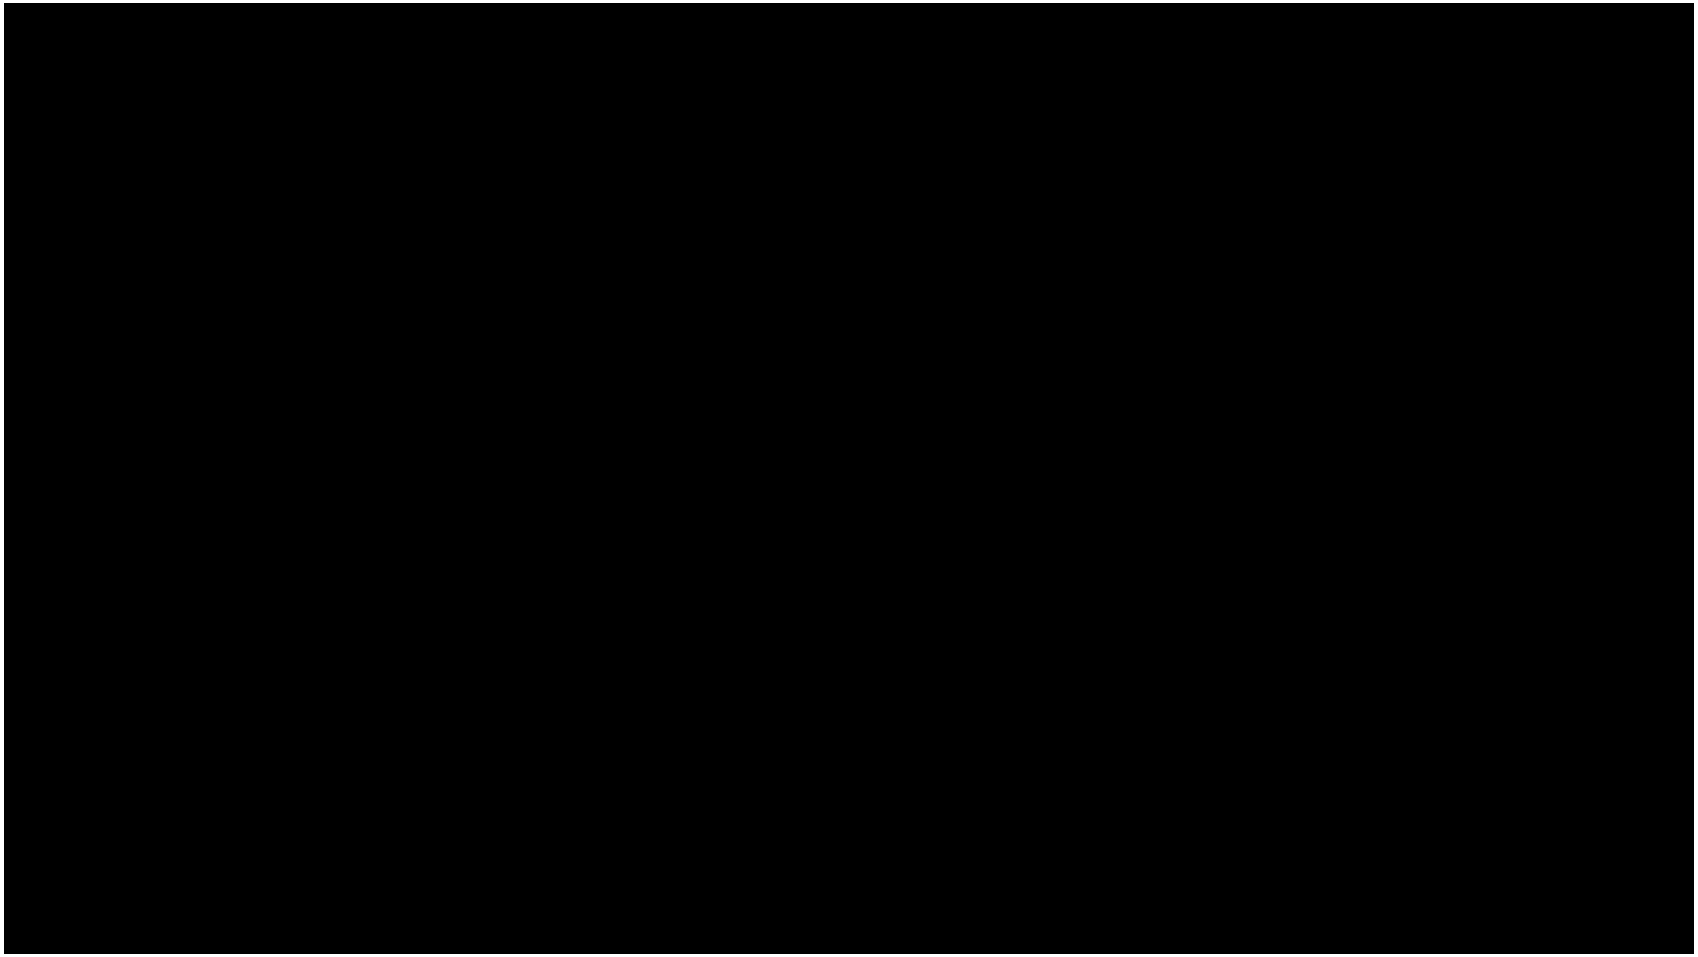

Demartek About Us video – http://www.youtube.com/demartek

Solid State Technology Overview

- Presents memory technology, such as DRAM or NAND flash, as storage media and appears as a disk drive to the operating system in most cases
  - Some motherboards allow dedicated SSD to act as a cache or other functions
- Very fast, no moving parts (no “seek time”)
- Variety of form factors
- Prices dropping
- Some SSDs use DRAM and NAND flash together
- Capacities doubling almost yearly

NAND Flash – What is it?

- A specific type of EEPROM
  - EEPROM – Electrically erasable programmable read-only memory
  - The underlying technology is a floating-gate transistor that holds a charge
- Bits are erased and programmed in blocks
  - Process is known as the program-erase (P-E) cycle
  - Flash blocks are typically 4KB, some larger
- Non-volatile
- Quiet, low-power, low-weight, low-heat
- Types – SLC, MLC, TLC

Types of NAND Flash

- Single-level Cell (SLC) – 1 bit per cell
- Multi-level Cell (MLC) – 2 bits per cell
  - Consumer grade (cMLC)
  - Enterprise/Endurance grade (eMLC)
- Triple-level Cell (TLC or MLC-3) – 3 bits per cell

	SLC	MLC	TLC (MLC-3)
Bits per cell	1	2	3
Performance	Fastest		Slowest
Endurance	Longest		Shortest
Capacity	Smallest		Largest
Error Prob.	Lowest		Highest
Price per GB	Highest		Lowest
Applications	Enterprise	Enterprise / Consumer	Consumer

NAND Flash – Endurance

- Single-level cell (SLC)
  - SLC typical life of 100,000 write cycles
- Multi-level cell (MLC)
  - MLC typical life 10,000 or fewer write cycles
    - MLC-2: 3,000 – 10,000 write cycles
    - MLC-3: 300 – 3,000 write cycles
  - “Enterprise MLC” (eMLC): 20,000 – 30,000 write cycles
    - Based on MLC-2
    - Better name is probably “Endurance MLC”
- As die size decreases, endurance also decreases
  - This may be fine for consumer, but not enterprise applications

NAND Flash Performance

- IOPS
  - 10K – 250K reads per device
    - Enterprise HDDs – 100-200 IOPS
    - Desktop HDDs – < 100 IOPS
  - Writes are generally slower than reads
- Bandwidth
  - Up to 550 MB/s for 6 Gb/s SAS or SATA interface drives
    - Up to 1 GB/s for SAS wide-port drives
  - Up to 3.2 GB/s for PCIe cards
- Latencies
  - <1 ms for drives, depending on interface
  - Well below 1 ms for PCIe cards

NAND Flash Maximum Capacities Today

- Individual Devices
  - Drive form factor – 2 TB, 2.5-inch, 6 Gb/s SAS
    - Compare to capacities of 10K & 15K RPM HDDs
  - PCIe card – 10 TB
  - mSATA (mini PCIe) – 256 GB
  - SATADIMM stick – 480 GB
- Storage Arrays
  - All-flash, single-rack – 1 PB

SSD Form Factor – SCSI Express

- SCSI Express (2.5-inch PCIe)
  - Combines proven SCSI protocol with performance of PCIe bus
  - Products expected in 2014
  - Ideal for SSDs

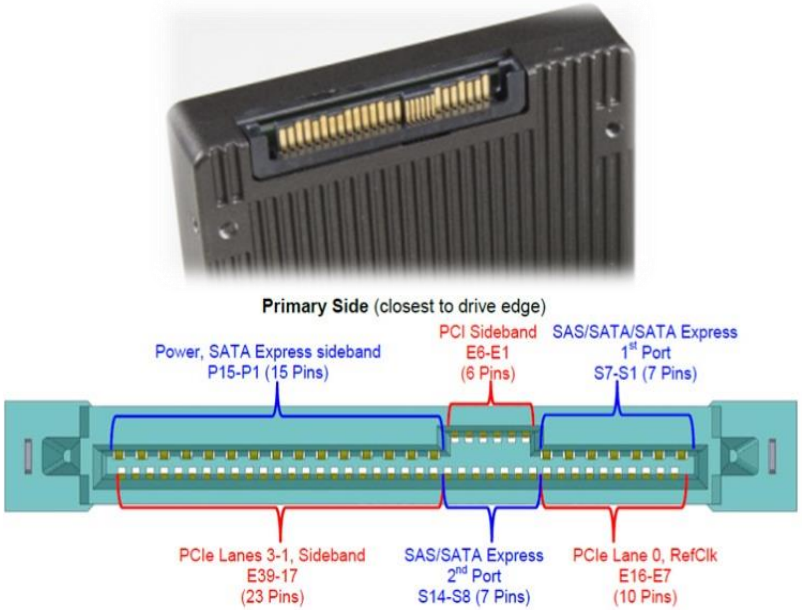

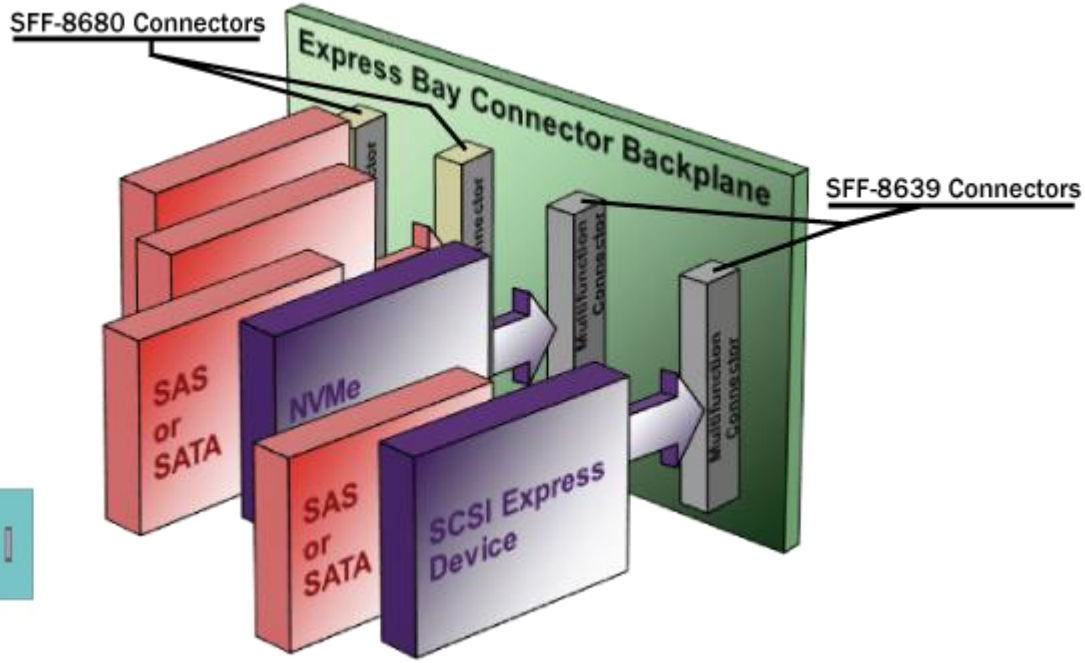

More information:

Demartek Storage Interface Comparison

SSD Form Factor – SATA Express

- SATA Express
  - SATA compatibility transported by the PCIe bus
  - Speeds will be 8 Gb/s and 16 Gb/s
  - Specification in member review as of January 2013
  - Products expected in 2014

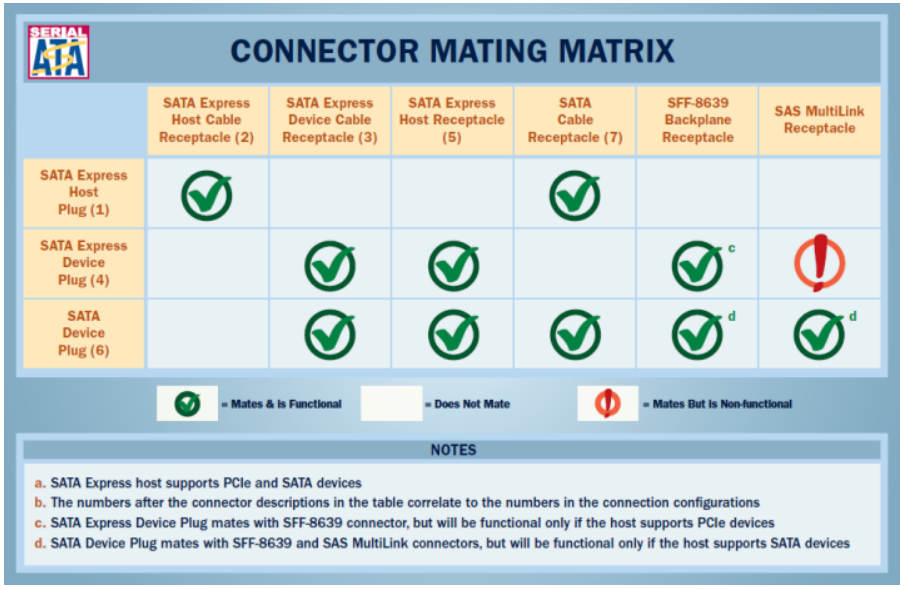

CONNECTOR MATING MATRIX

	SATA Express Host Cable Receptacle (2)	SATA Express Device Cable Receptacle (3)	SATA Express Host Receptacle (5)	SATA Cable Receptacle (7)	SFF-8639 Backplane Receptacle	SAS MultiLink Receptacle
SATA Express Host Plug (1)	✓			✓		
SATA Express Device Plug (4)		✓	✓		✓[c]	!
SATA Device Plug (6)		✓	✓	✓	✓[d]	✓[d]

✓ = Mates & is Functional = Does Not Mate ! = Mates But Is Non-functional

NOTES

a. SATA Express host supports PCIe and SATA devices
b. The numbers after the connector descriptions in the table correlate to the numbers in the connection configurations
c. SATA Express Device Plug mates with SFF-8639 connector, but will be functional only if the host supports PCIe devices
d. SATA Device Plug mates with SFF-8639 and SAS MultiLink connectors, but will be functional only if the host supports SATA devices

View larger version of this chart on the ***Demartek Storage Interface Comparison*** web page

NVM-Express (NVMe)

- Scalable host controller interface designed for enterprise and client SSDs
  - Targeted at “high-frequency” storage applications
  - Goal is to streamline access to SSD devices that are directly connected to the PCIe bus, including the storage stack in the O.S.
  - Compatible with SCSI/SAS, but uses an underlying command set of six I/O commands for efficiency
- Enterprise samples expected in 2013, G.A. in 2014
- More information:
  - http://www.demartek.com/Demartek_Comments_IDF2012_and_NVMe.html
  - http://www.nvmexpress.org/

SSD Caching Basics

- Caching controller identifies any frequently accessed data (“hot data”) and automatically moves ***a copy*** of the hot data to SSD media
- SSD impact
  - Multiple applications can benefit from the SSD cache simultaneously
  - Performance improves over time, as cache is populated with data
    - This is known as “cache warm-up” or “cache ramp-up”
- Some caching solutions cache only the reads, others cache both reads and writes
- Overall HDD I/O load is reduced – Fewer I/Os

SSD Caching Architecture

- SSD caching can be added:
  - Server-side
  - In the network
  - In the storage system
- We have seen increased performance benefits by combining server-side with the others

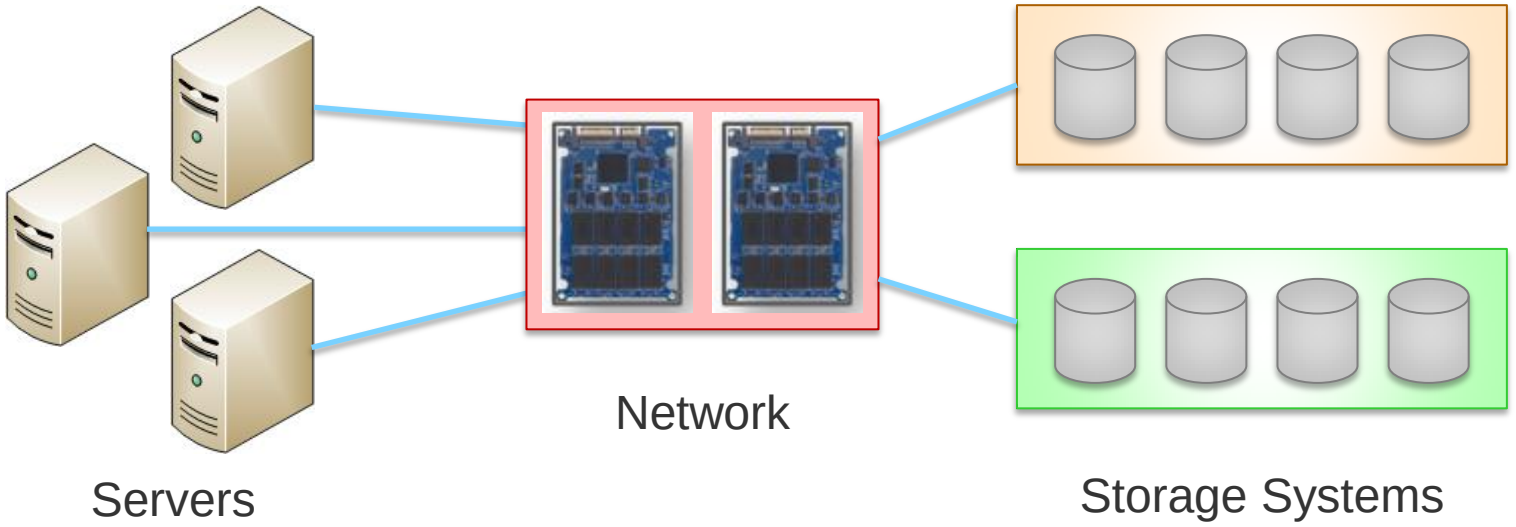

SSD Caching Workloads

- Caching Algorithms
  - Prefer random I/O, small to moderate block sizes
- Cache Friendly Workloads
  - Hot spots with repeated access
  - OLTP databases
  - Database indexes
  - File system table of contents (MFT, inodes, etc.)
- Cache Un-friendly Workloads
  - Data that is accessed approximately evenly and is larger than the cache

SSD Caching Performance – Effect on Application

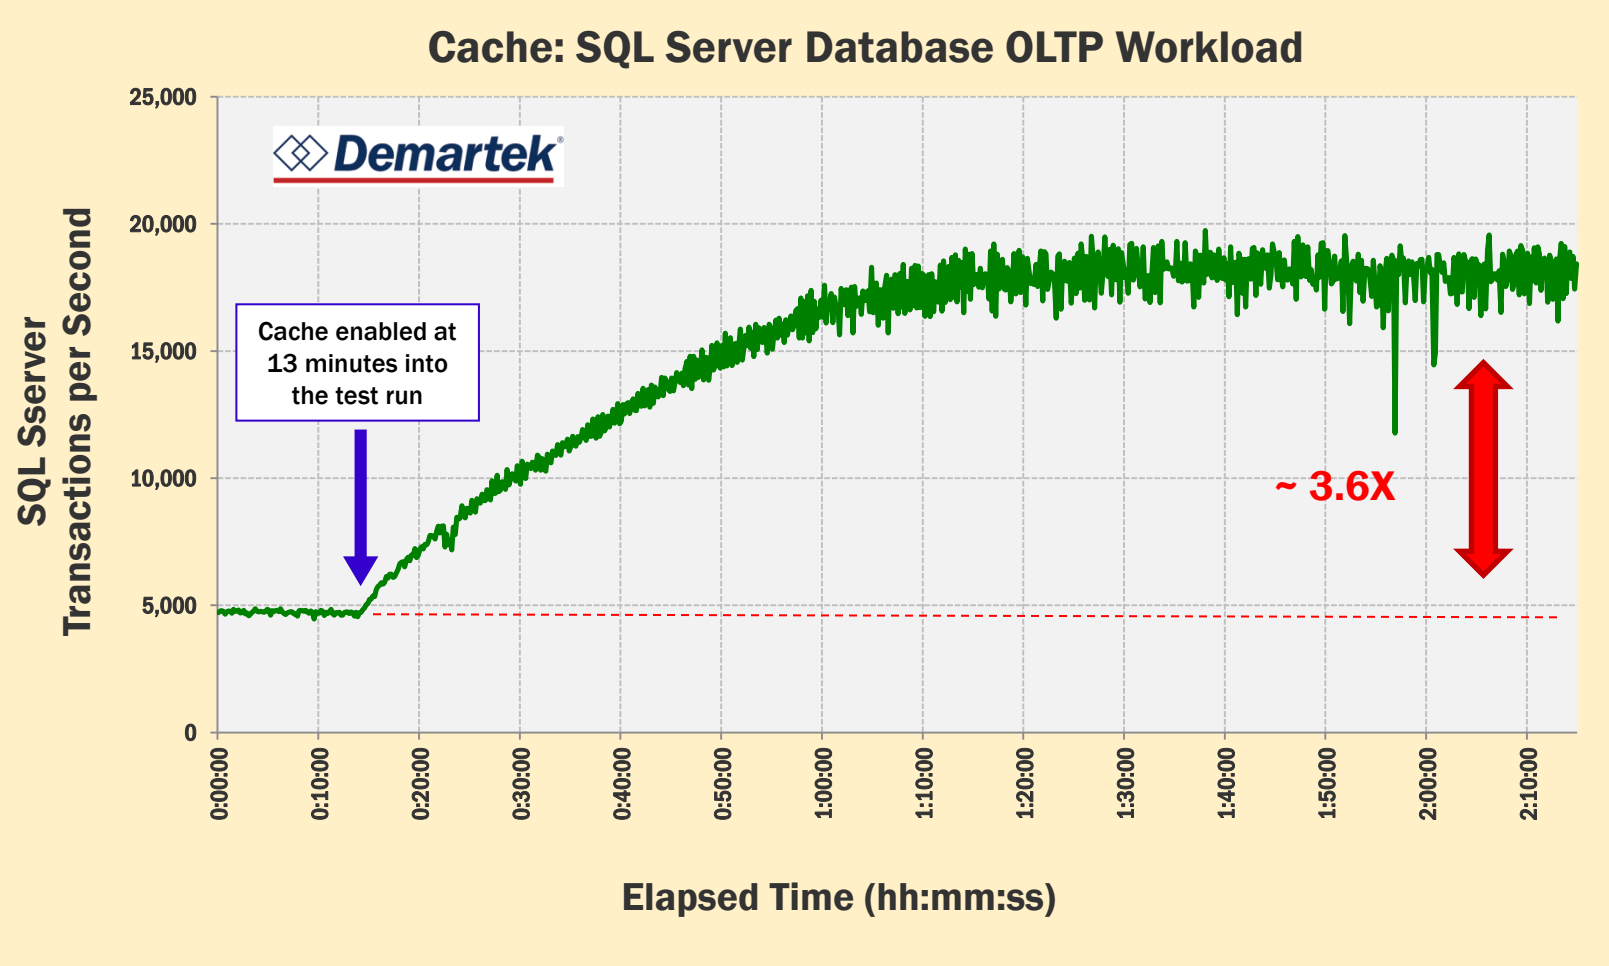

Source: http://www.demartek.com/Demartek_EMC_VFCache_Evaluation_2012-02.html

SSD Caching Performance – Effect on HDDs

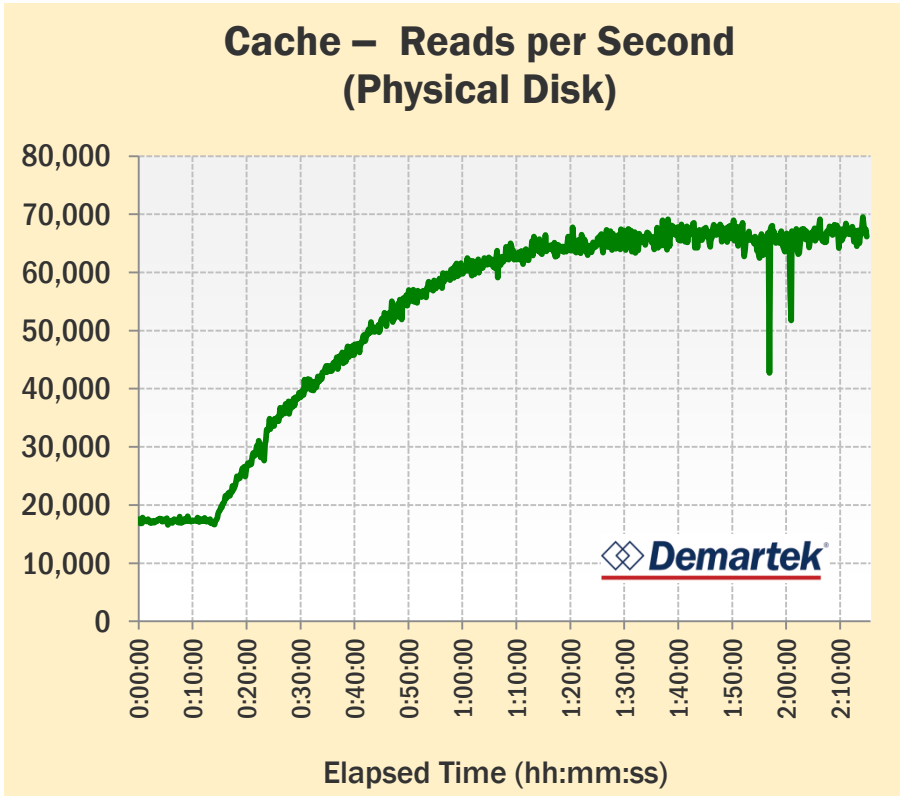

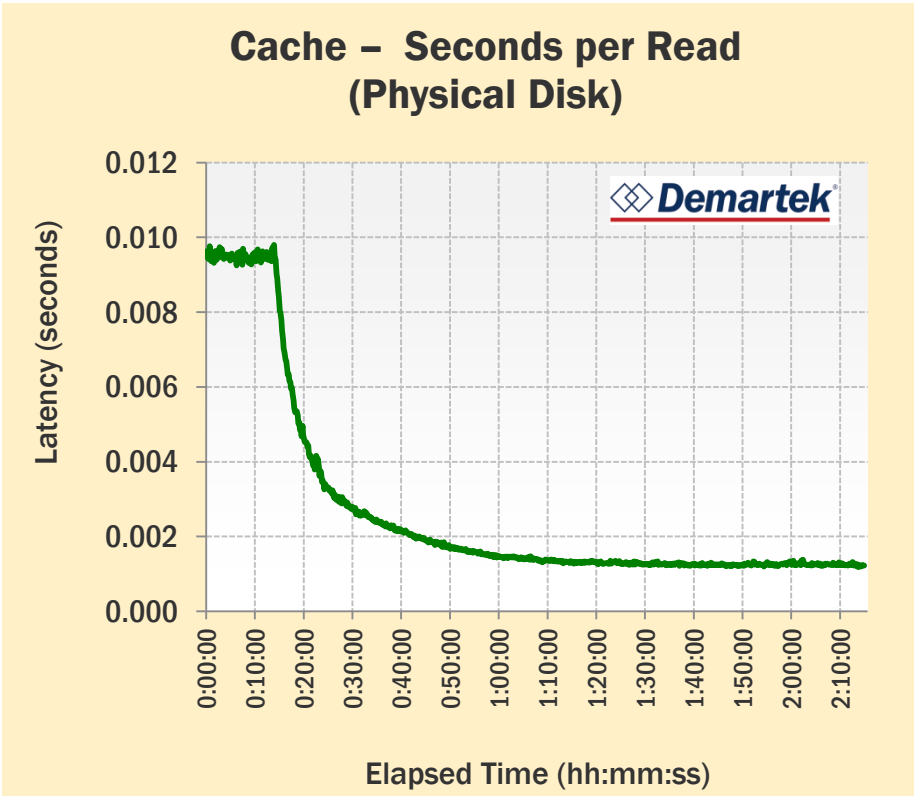

Source: http://www.demartek.com/Demartek_EMC_VFCache_Evaluation_2012-02.html

Future Non-Volatile Memory Technologies

- Today
  - NVDIMM (DDR3) – available today, same speed and capacities as DRAM, but more expensive
- Futures – near term
  - Phase Change Memory (PCM or PRAM) – probably closest to commercial viability, some shipping now
- Futures – moderate to long term
  - Ferroelectric (FeRAM)
  - Magnetic RAM (MRAM) – includes “Racetrack” & “Spin-Torque”
  - Resistive RAM (RRAM) – includes “Memristors”
  - Conductive Metal Oxide (CMOx)
  - Solid Electrolyte

Future NV Memory Technologies Commentary

- The technologies listed on the previous page are interesting from a science and physics perspective. However, the key to their commercial viability hinges heavily on the cost to produce large quantities.
- The largest quantities of NAND flash today are produced for the consumer market, such as cell phones and tablet computers. The enterprise market for SSDs generally has a somewhat lower priority from the high-volume producers, simply because the volumes are not as large for enterprise products compared to consumer products.
- The next thing (PCM, possibly) will only become commercially viable if the manufacturers can get significantly better costs than NAND flash for equivalent features and capacities.

Demartek References

- Demartek SSD Zone
  - www.demartek.com/SSD
- Demartek SSD Deployment Guide
  - www.demartek.com/Demartek_SSD_Deployment_Guide.html
- Demartek Commentary – Horses, Buggies & SSDs
  - www.demartek.com/Demartek_Horses_Buggies_SSDs_Commentary.html
- Demartek Free Monthly Newsletter
  - www.demartek.com/Newsletter

Thank You!

Dennis Martin, President
dennis@demartek.com
www.linkedin.com/in/dennismartin

Demartek

(303) 940-7575
www.demartek.com
http://twitter.com/Demartek
www.youtube.com/Demartek
Skype: Demartek

To learn more about Demartek:

- Download the Aurasma App (Android/iPhone)
- Search and follow "Demartek"
- View image below with viewfinder.

*also on the back of Dennis' business card

Powered by: AURASMA