Solid State Storage for the Enterprise

Woody Hutsell, Texas Memory Systems, Inc.
Brian McKean, LSI Corporation

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced without modification
  - The SNIA must be acknowledged as 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 nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
- The information presented herein represents the Author's personal opinion and current understanding of the 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.

© 2008 Storage Networking Industry Association. All Rights Reserved.

Abstract

Solid State Storage for the Enterprise

This presentation will discuss the types of solid state storage systems directed at enterprise. The presentation will look at DDR-RAM and Flash based storage devices and their technical differences. The presentation will provide some examples of the variance in the operational characteristics among different flash based SSD. The presentation will discuss the fit of various memory based storage devices with enterprise applications.

Learning Objectives

- **Understand types of solid state storage available for enterprise applications.**
- **Understand that not all Flash SSDs are the same.**
- **Discuss applications for solid state storage in the enterprise.**

© 2008 Storage Networking Industry Association. All Rights Reserved.

Agenda Topics

- Solid State Disk Overview
- Flash SSD in a HDD Form Factor

© 2008 Storage Networking Industry Association. All Rights Reserved.

Solid State Disk Overview

- What is a Solid State Disk?
- Flash and DDR Memory Characteristics
- Why Solid State Disk?
- SSD Killer Applications
- SSD and ILM

© 2008 Storage Networking Industry Association. All Rights Reserved.

What is a Solid State Disk?

- "A solid state drive (SSD) is a semiconductor-based block storage device that behaves as a virtual HDD and appears to the host device as a disk drive."
  - Source: IDC, Worldwide Solid State Drive 2008–2012 Forecast and Analysis: Entering the No-Spin Zone , Doc #212736, June 2008.

© 2008 Storage Networking Industry Association. All Rights Reserved.

What is a Solid State Disk?

Education SNIA

- Flash Memory Based
  - Same class of memory used in consumer electronics
  - Inherently non-volatile
  - Best known for ruggedness and good random read performance.
- DDR Memory Based
  - Same memory used in enterprise servers
  - Requires batteries and backup hard disks for non-volatility
  - Best known for outstanding performance and high cost.
- Cached Flash
  - Mix of DDR RAM and NAND Flash.

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash Memory Characteristics

- Two types of NAND Flash Memory
  - SLC – single layer
    - 100,000 writes per cell
    - Primarily used in industrial and military applications
    - Higher cost
    - 1.5 millisecond erase times; 200 microsecond write times; 25 microsecond read times
    - Maximum density – 16Gbit with 32Gbit on the way
    - SLC memory is best suited for the enterprise
  - MLC – multi layer
    - 10,000 writes per cell
    - Primarily used in consumer electronics
    - Lower cost
    - Half the performance of SLC
    - Maximum density – 32Gbit with 64Gbit on the way.
- Expect mixed SLC – MLC flash SSDs in the near future.

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash Memory Characteristics

- The process to write data to flash follows these steps:
  - Determine "block" to update. NAND flash is typically divided into 128KB blocks and further subdivided into 2KB pages.
  - Copy data from the existing "block" (if necessary)
  - Erase the "block" and reset cell to all "1's"
  - Rewrite the "block", only 0's can be written to a flash cell.

© 2008 Storage Networking Industry Association. All Rights Reserved.

DDR RAM Characteristics

- Unlimited writes per cell
- Primarily used as computer memory
- Higher cost and lower density than NAND flash
- Inherently volatile
- Highly reliable chip design
- 10-15 nanoseconds read and write times
- Maximum density 1Gbit (2Gbit is being sampled)
- DDR RAM is well suited for the enterprise.

© 2008 Storage Networking Industry Association. All Rights Reserved.

What is a Solid State Disk?

Education SNIA

1.8”	2.5”	3.5”	JBOD	Rack Mount
		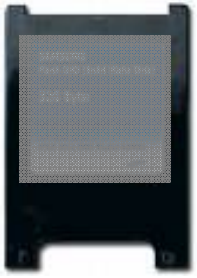		

Flash	Flash	DDR/Flash	DDR/Flash	DDR/Flash
SATA IDE	SATA IDE	SATA SCSI Fibre Channel	SCSI Fibre Channel SAS	SCSI Fibre Channel InfiniBand SAS

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Converging Trends

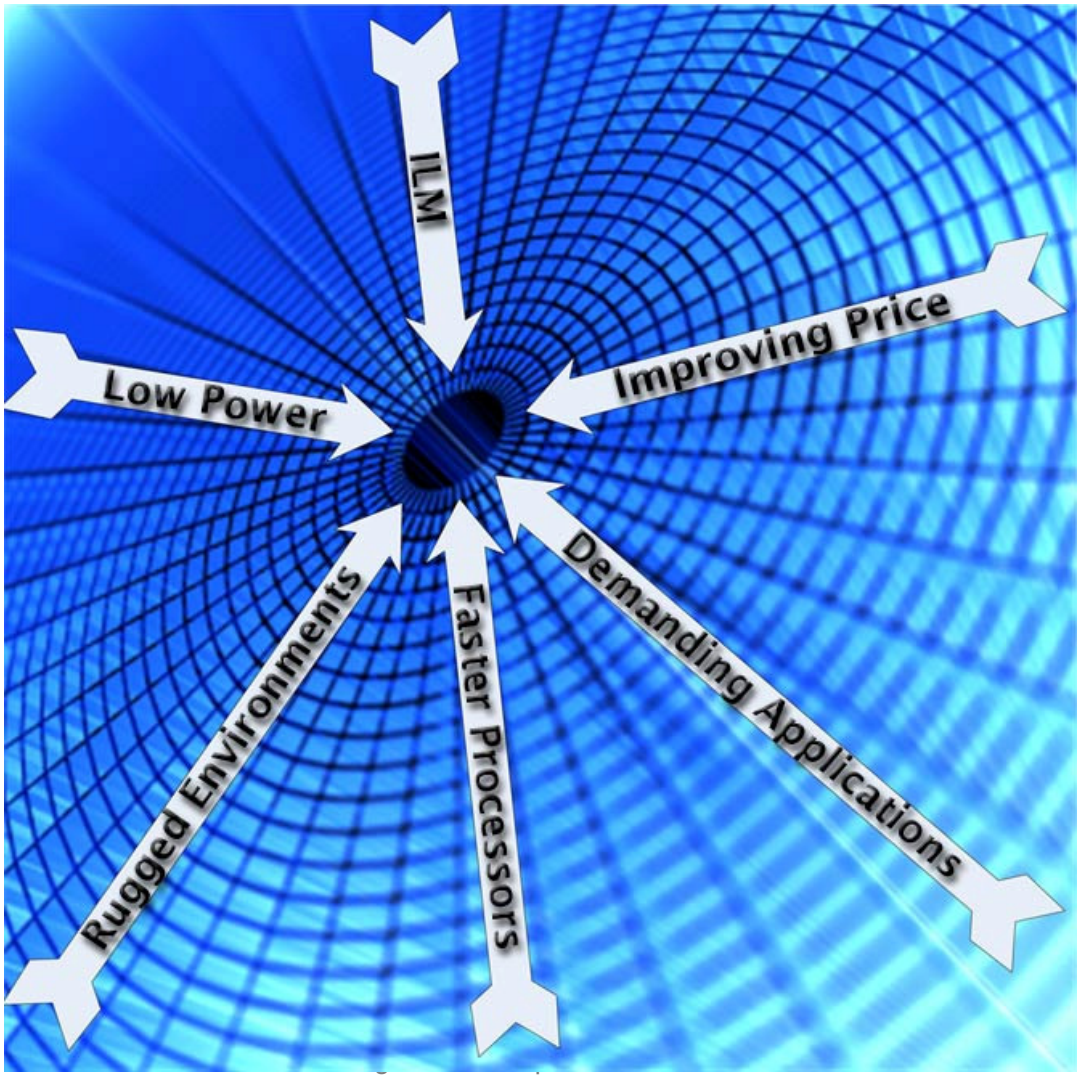

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Latency Matters

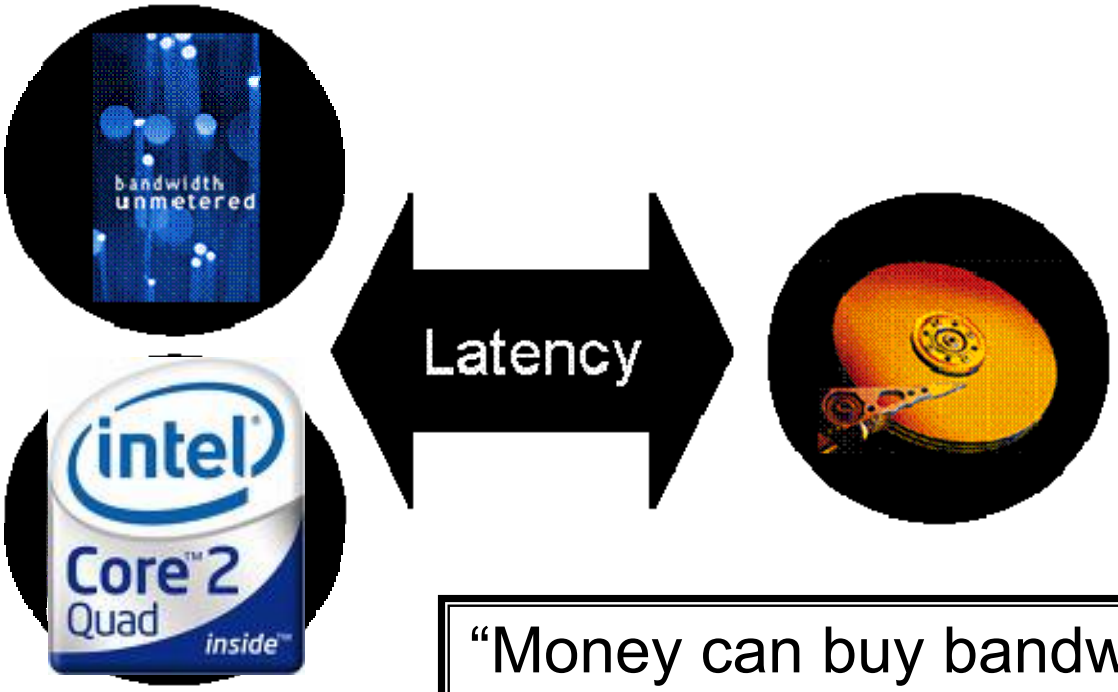

"Money can buy bandwidth, but latency is forever"

John R. Mashey, Chief Scientist SGI,
"Big Data and the Next Wave of InfraStress", USENIX, 1999

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Little Change in HDD RPM

- 1956 RAMAC
  - the first disk drive
  - 5 MB storage
  - 1,200 RPM
- 2007 SCSI Hard Drive
  - 400 GB storage
  - 15,000 RPM
- From 1956 to 2007:
  - 12.5 times increase in RPM
  - 80,000 times increase in capacity

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Data Access Times

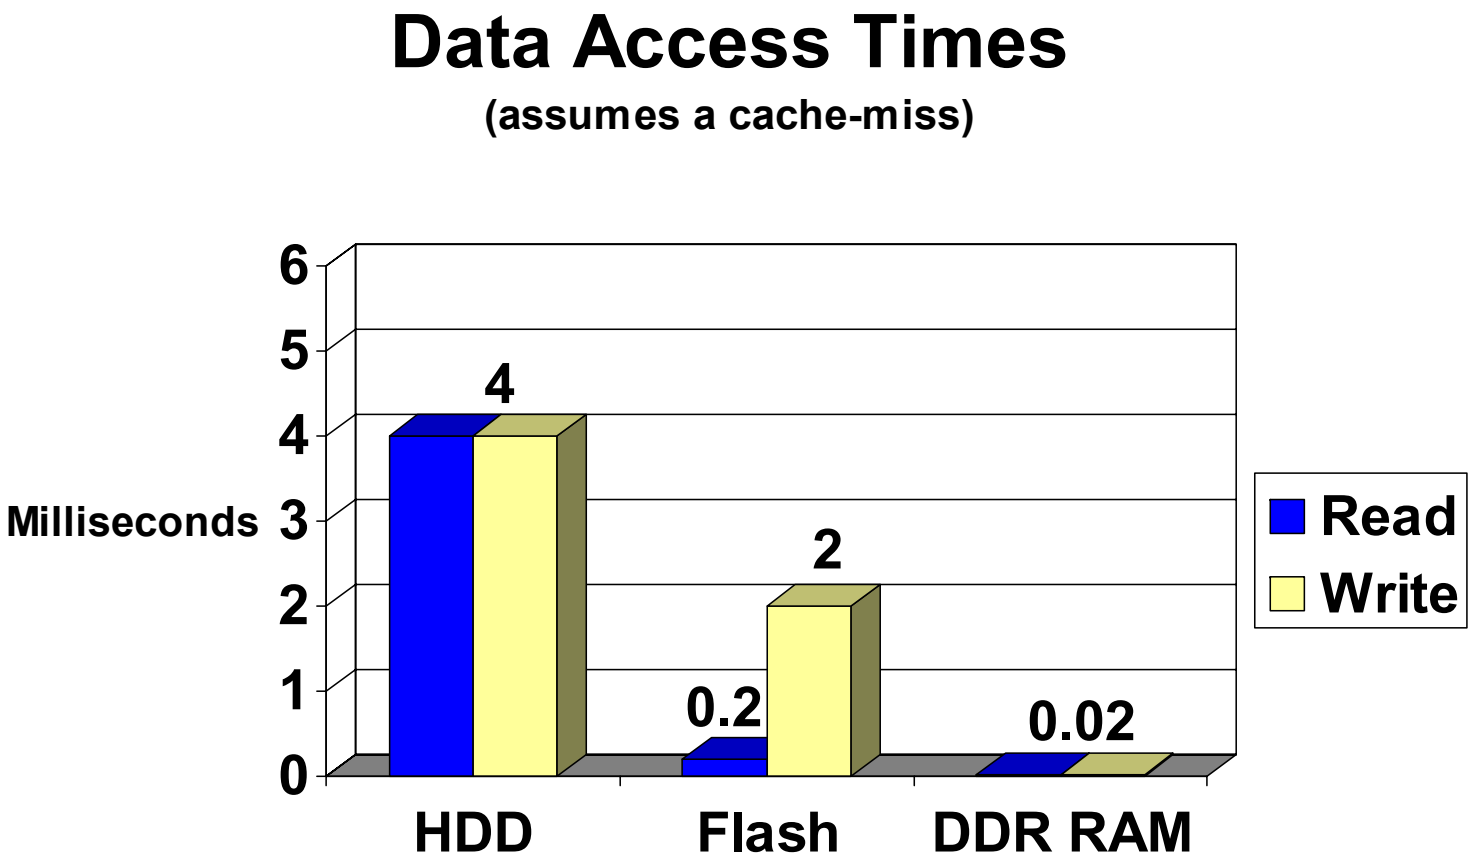

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Random I/O's

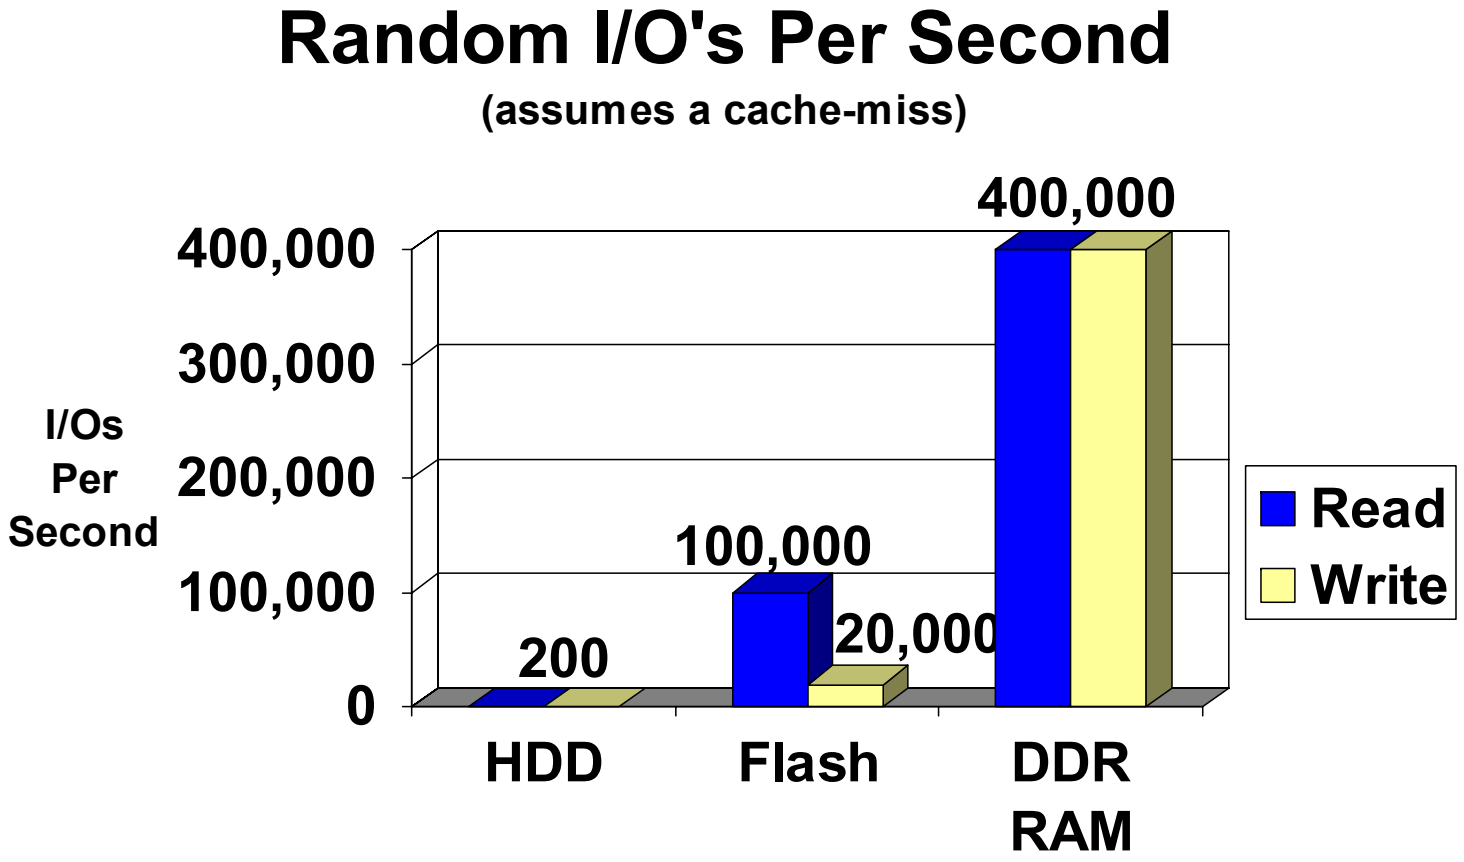

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Low Power

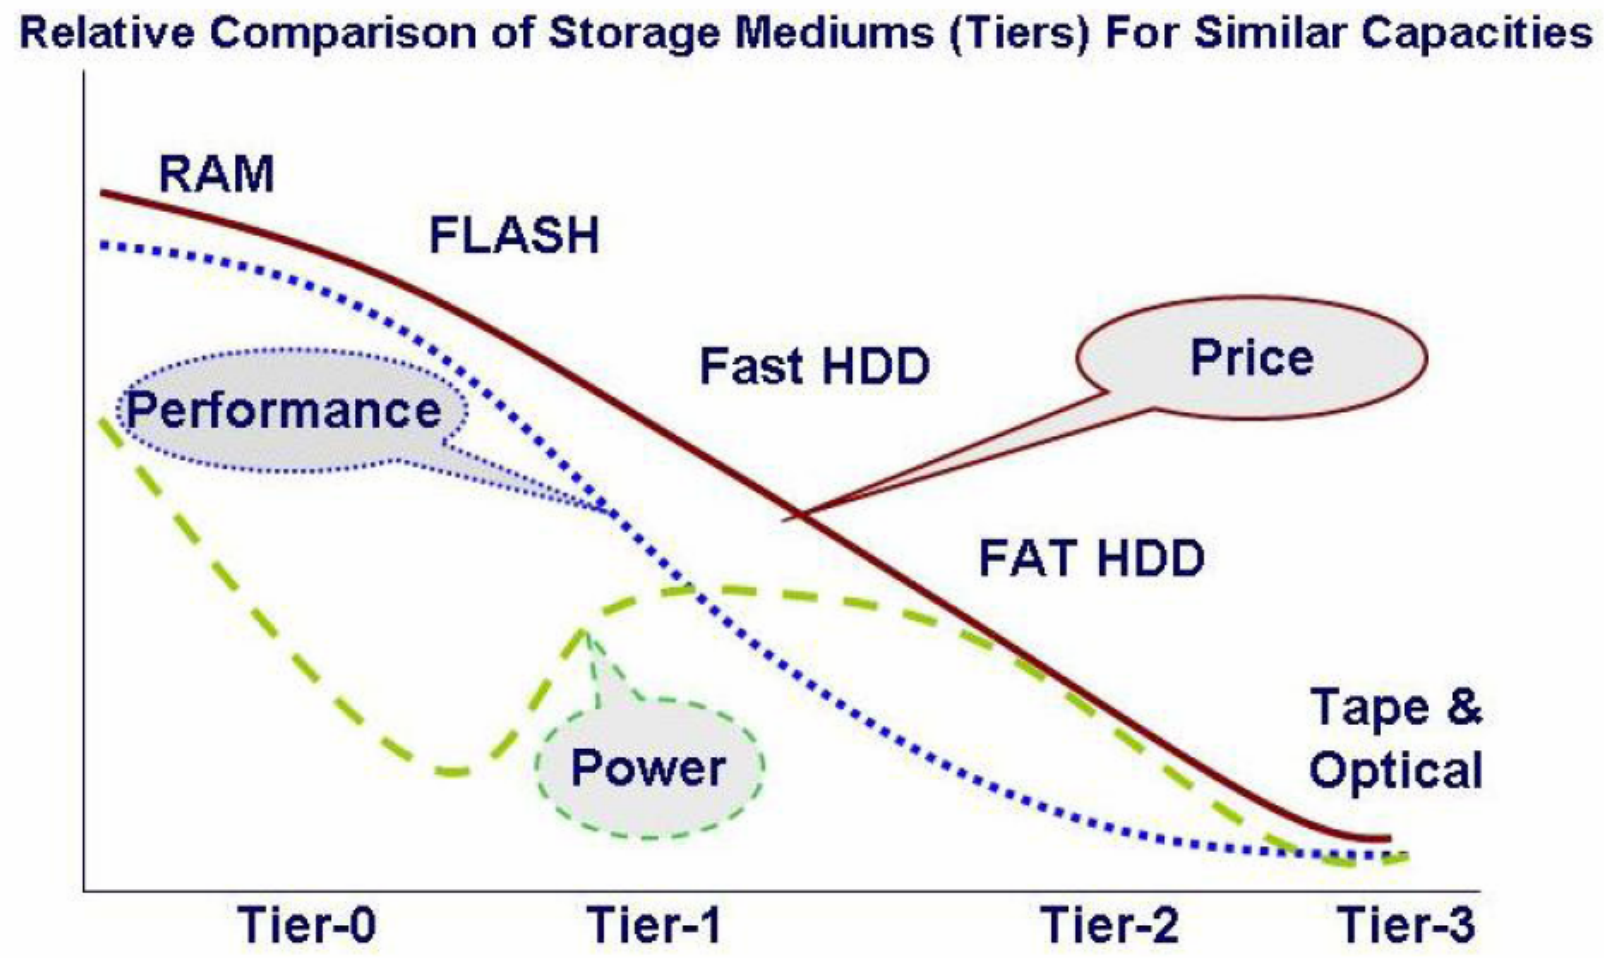

Source: Achieving Energy Efficiency using Flash SSD, 2008, Greg Shultz, Storage IO Group.

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Dropping Prices

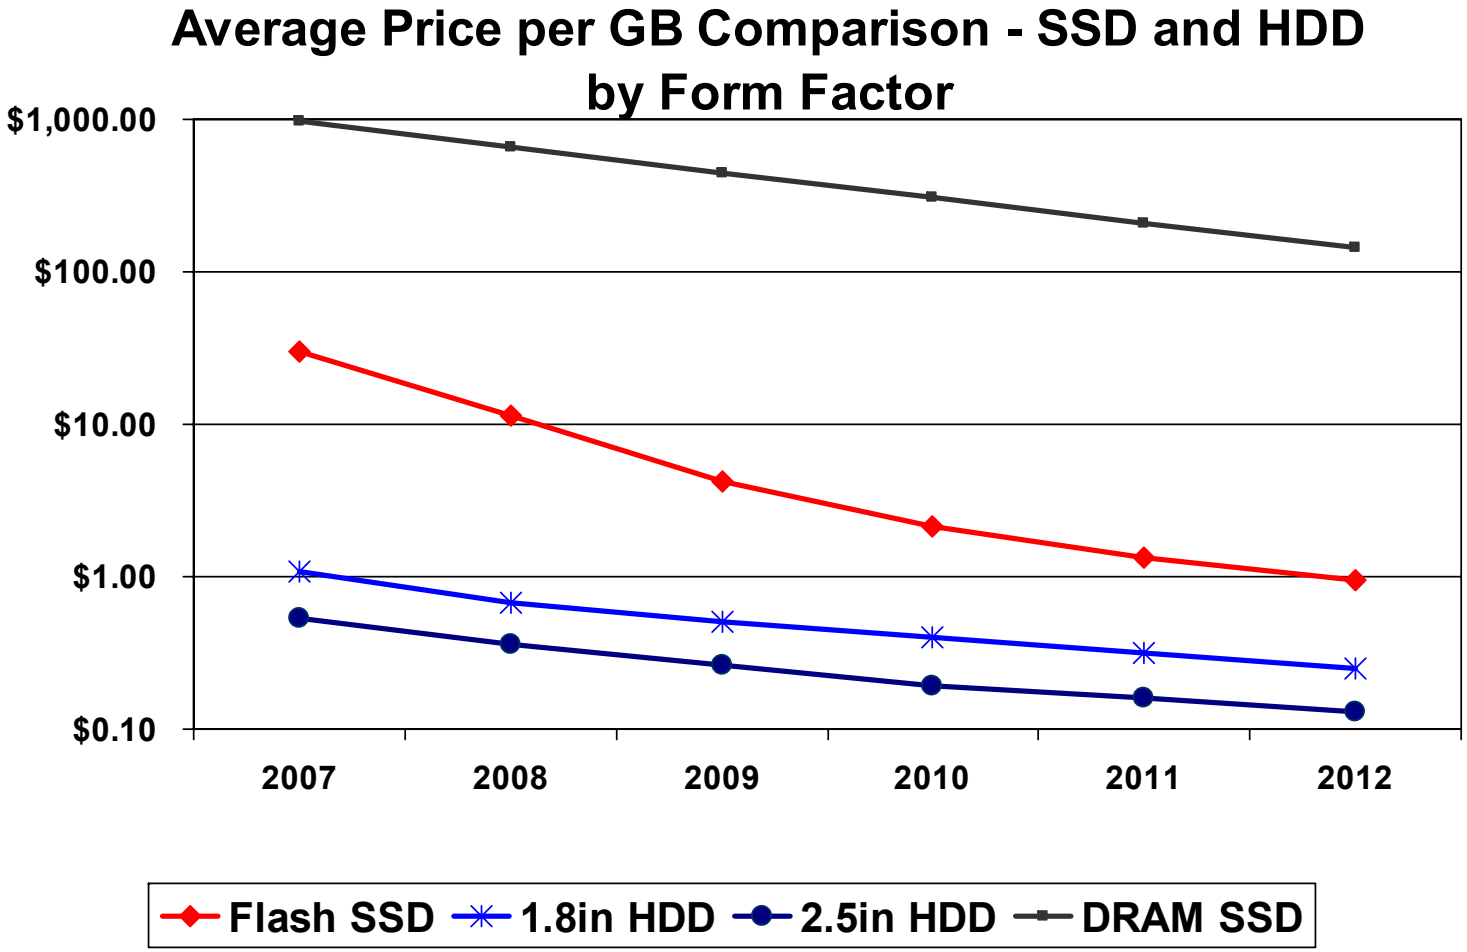

Note: DRAM SSD include additional components (I.e. batteries, etc)

Source: IDC, Worldwide Solid State Drive 2008–2012 Forecast and Analysis: Entering the No-Spin Zone , Doc #212736, June 2008.

© 2008 Storage Networking Industry Association. All Rights Reserved.

Why Solid State Disks?
Low Price for Performance

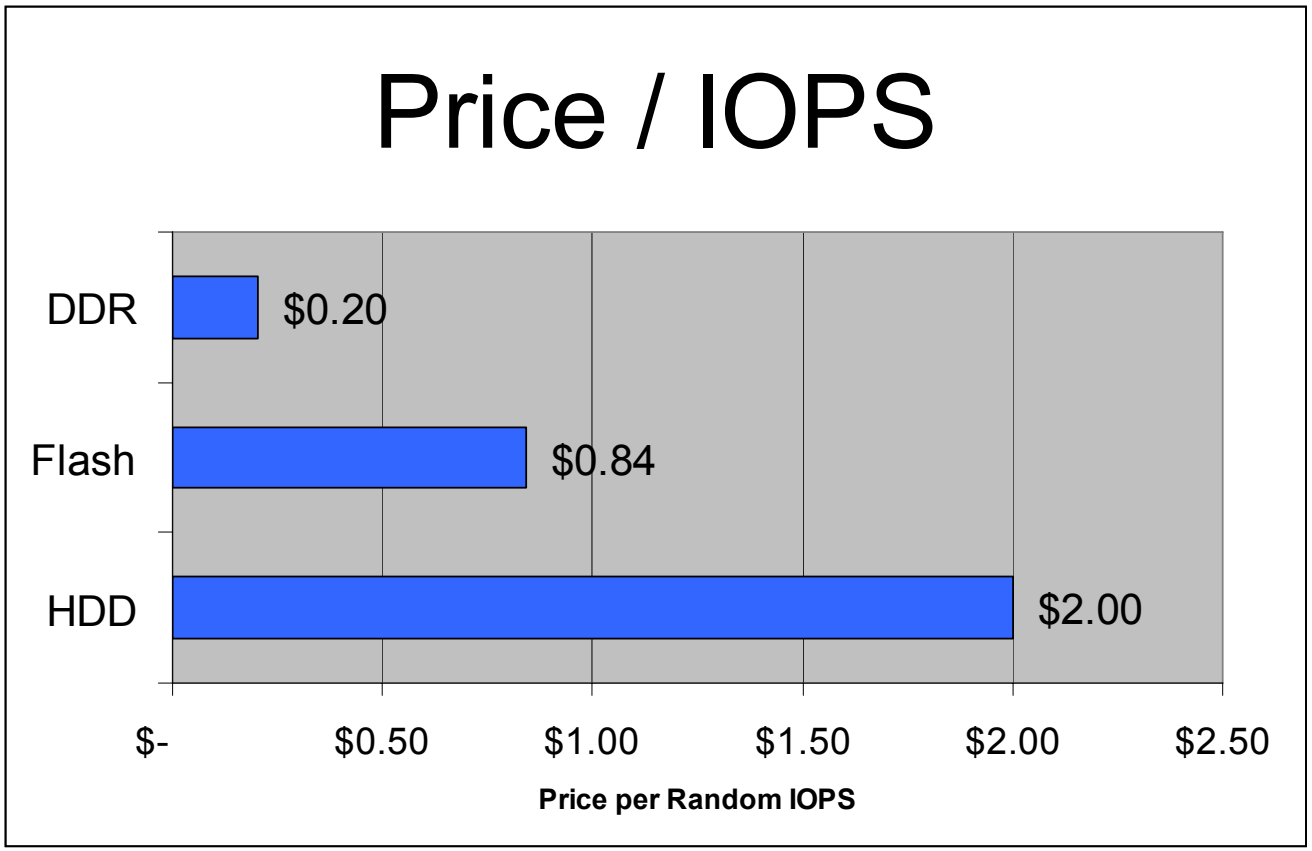

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD
Killer Applications

- Business notebook computer storage
  - Faster OS startup
  - Improved notebook read performance
  - Low power
  - Rugged

© 2008 Storage Networking Industry Association. All Rights Reserved.

Cached Flash RAID SSD
Killer Applications

Education SNIA

- Data warehousing
  - Terabytes of capacity needed as hot files are difficult to isolate to a small amount of data
  - Very read intensive applications
  - Cache will preserve write performance at levels the enterprise is accustomed to receiving
- Rendering
  - Small block random I/O application
  - Very read intensive application
  - Hot files are difficult to isolate

© 2008 Storage Networking Industry Association. All Rights Reserved.

Cached Flash RAID SSD
Killer Applications

- Video on demand
  - Storing frequently accessed movies requires 1TB plus capacity
  - Very read intensive applications
  - Less frequently accessed content still stored in HDD RAID
- Seismic processing
  - Large capacity of data (4TB plus), broadly and randomly accessed
  - Very read intensive application

© 2008 Storage Networking Industry Association. All Rights Reserved.

DDR RAM SSD
Killer Applications

- Write intensive OLTP environments
  - Heavy writes make these environments unsuitable for flash systems
  - Store transaction logs, temporary space, and undo segments on SSD
- Web transaction databases
  - Extreme concurrency
  - High read and write volumes
- Single-threaded applications
  - The ultimate latency sensitive application
  - Often found in the financial industry

© 2008 Storage Networking Industry Association. All Rights Reserved.

DDR RAM SSD
Killer Applications

Education SNIA

- Mission critical data warehousing
  - When company profitability or lives are on the line, nothing is faster than a DDR RAM based solution
  - Data Warehouses with simultaneous high ingest rates and high query rates

© 2008 Storage Networking Industry Association. All Rights Reserved.

SSD and ILM

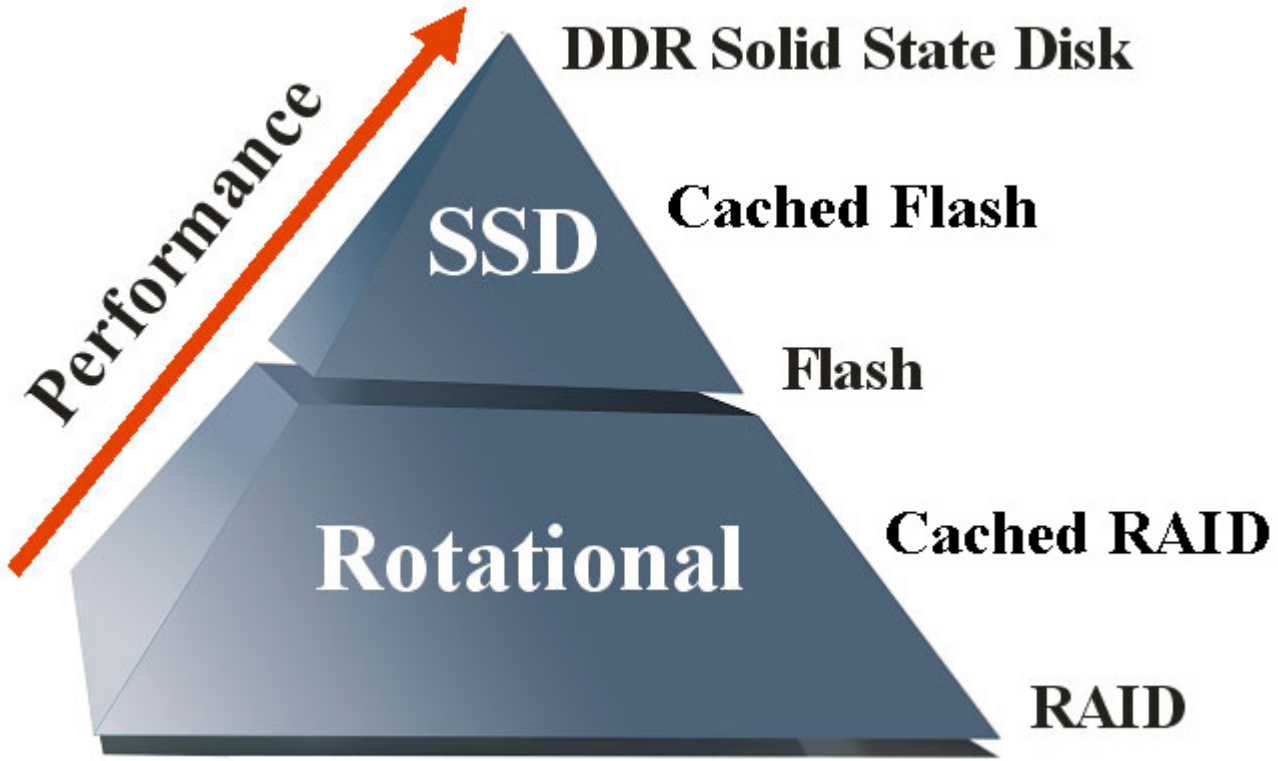

© 2008 Storage Networking Industry Association. All Rights Reserved.

Growth in the Enterprise SSD Market

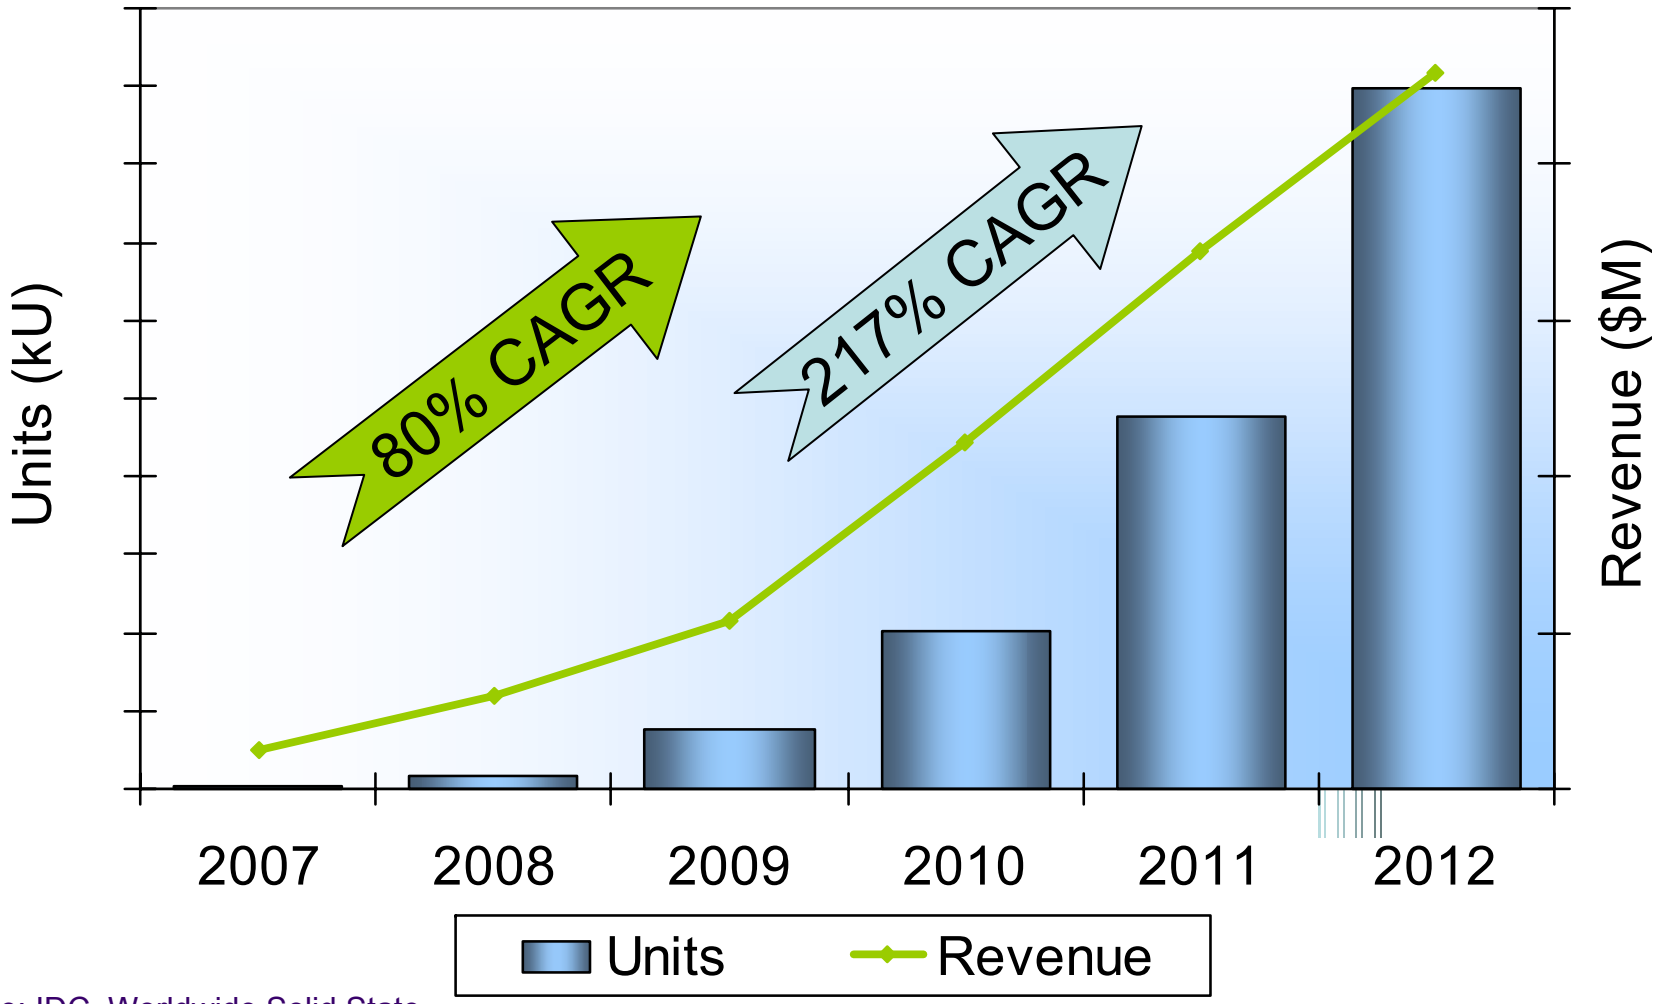

Source: IDC, Worldwide Solid State Drive 2008–2012 Forecast and Analysis: Entering the No-Spin Zone , Doc #212736, June 2008.

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD in an HDD Form Factor

Education SNIA

Moving from HDD to SSD

- Flash SSD Characteristics
- Market Segmentation
- Performance Comparisons
- Wear Life
- System Design Considerations

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD Characteristics

Education SNIA

- Consistent Characteristics
  - Small block random read performance is excellent versus HDD
  - Read / write performance is asymmetric
    - Writes slower than reads
  - Sequential performance today not much better than enterprise HDD
  - Power less than HDD
  - More expensive per capacity than HDD / less expensive per IOP

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD Market Segments

- Flash SSD pricing

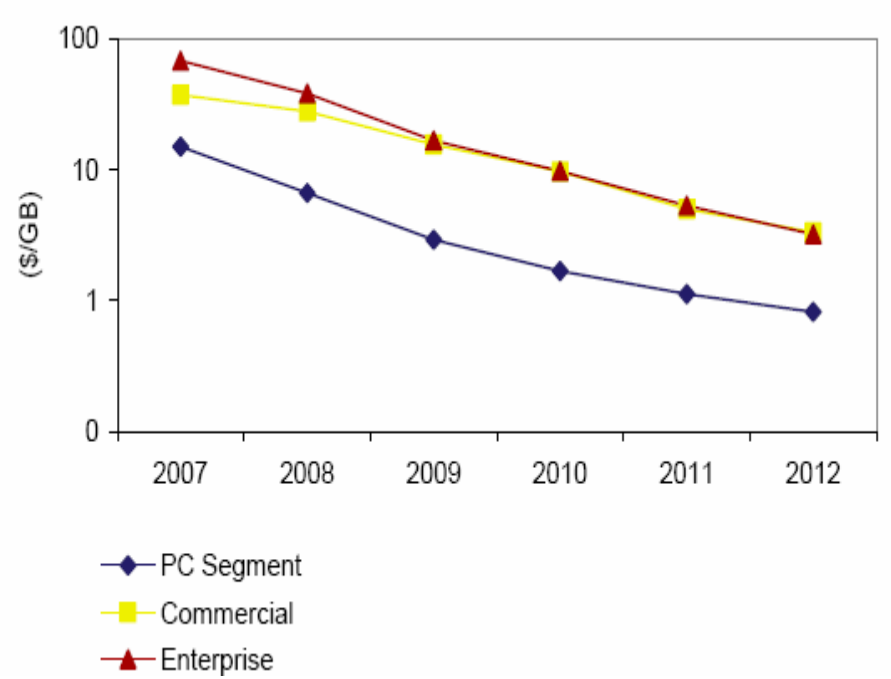

Source: IDC 2008: Worldwide Flash SSD Average Price per Gigabyte by Segment , 2007-2012

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

- HDD Performance
  - Roughly correlated to rotational speed
- SSD Performance
  - Varies widely between drives
  - Varies with work load in surprising ways
  - Varies with previous access patterns

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

SSD Random 512 byte IOPs Performance	Read	Write
SSD A	45000	16000
SSD B	19000	130
SSD C	7000	15
SSD D	6300	926
15K rpm HDD*	185	170
7.2K rpm HDD*	79	73
5.4K rpm HDD*	60	57

* calculated from data sheet seek time

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

Education SNIA

• Be careful about performance assumptions when dealing with asymmetrical read and write performance

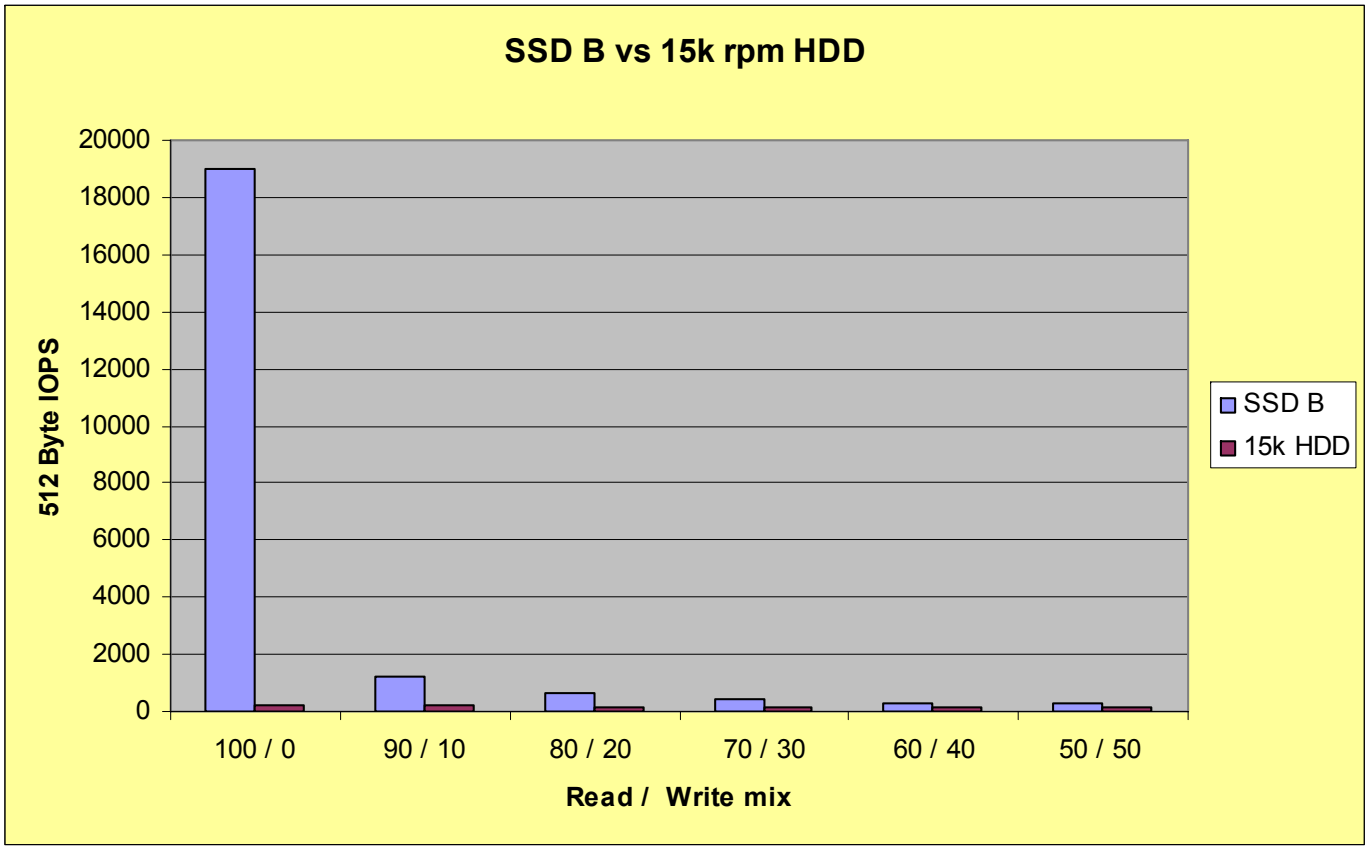

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

SSD Sequential Performance MB/sec	Read	Write
SSD A	220	115
SSD B	130	120
SSD C	57	38
SSD D	100	80
15K rpm HDD	171	171
7.2K rpm HDD	105	105
5.4K rpm HDD	61	61

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

- SSD Performance Surprises
  - Mixing reads & write slower than writes

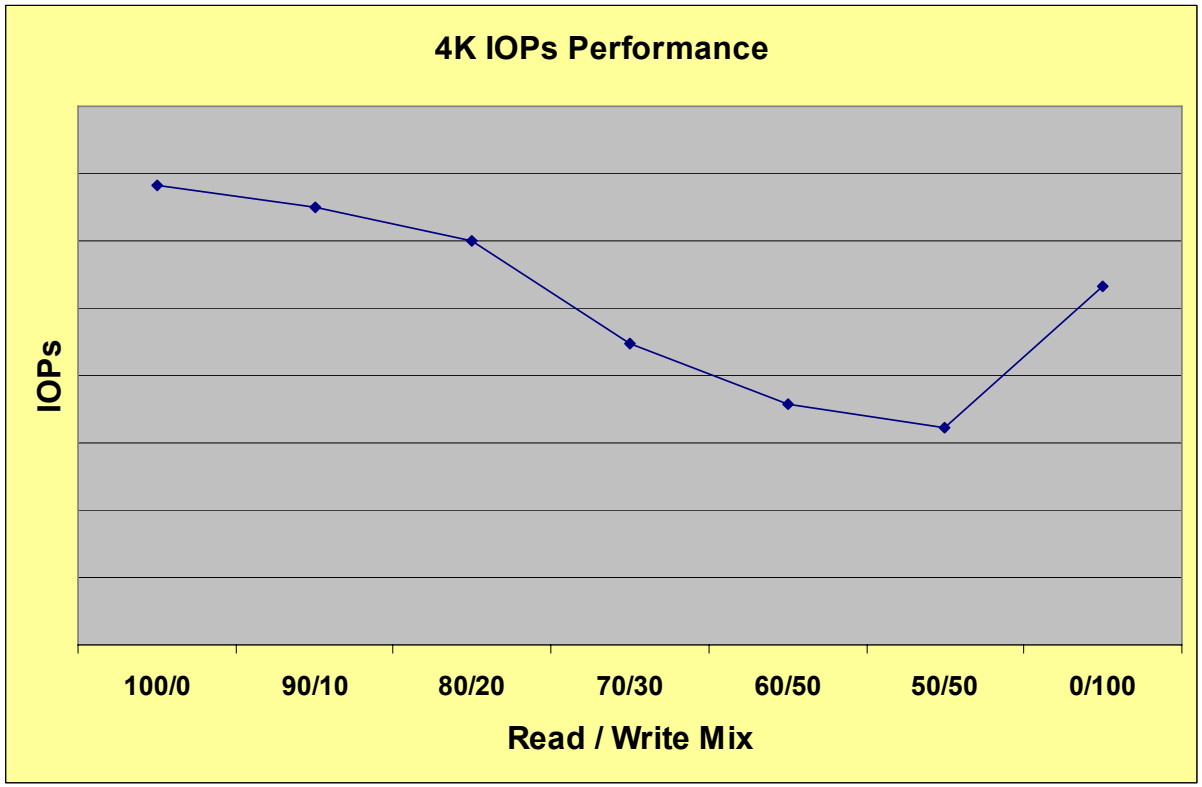

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

- SSD Performance Surprises
  - Comparative performance between drives differs based on IO profile

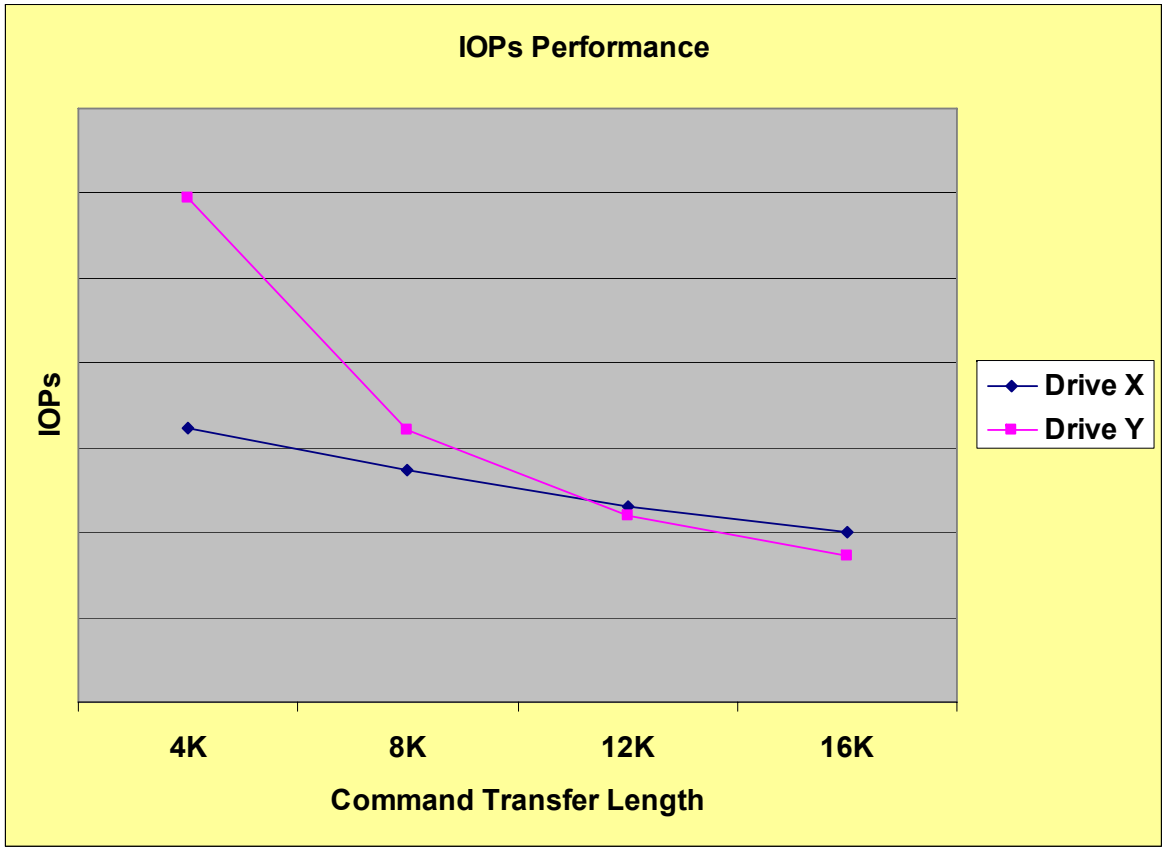

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD vs. HDD Performance

Education SNIA

- SSD Performance Surprises
  - Previous access patterns can dramatically affect current performance

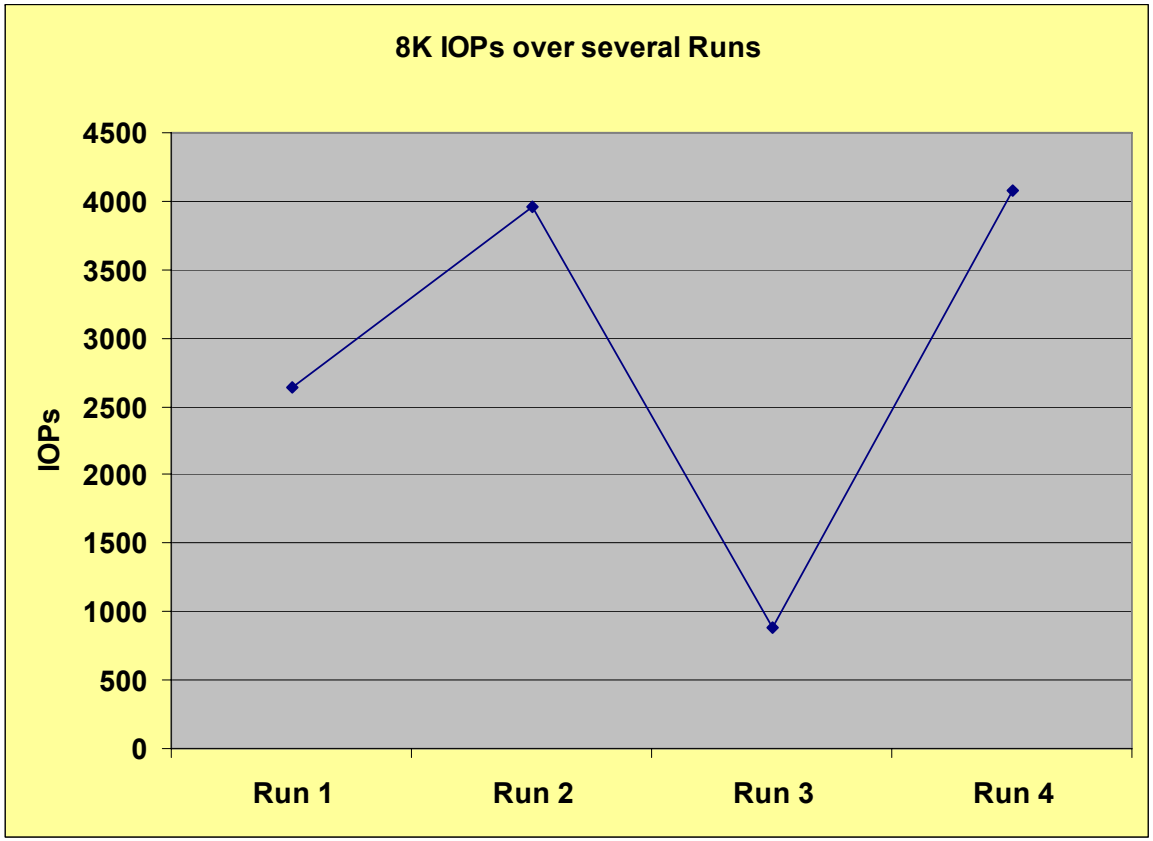

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD Wear Life

- Factors affecting wear life
  - Flash Technology
    - SLC – 100,000 P/E cycle
    - MLC – 10,000 P/E cycles
  - Controller Design
    - Average Flash Writes Per Host Write
    - Efficiency of Wear Leveling
    - SSD Rated Capacity
    - SSD Flash Capacity Above Rated Capacity
  - Use
    - Write rate at IO profiles
    - Duty Cycle

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD Wear Life

© 2008 Storage Networking Industry Association. All Rights Reserved.

Flash SSD Wear Life

- Wear Leveling Efficiency: 1.1
- Write Efficiency: 1.1
- Over Provisioning: 30%
- Duty Cycle: 80%
- R/W: 80/20
- User Capacity: 64GB

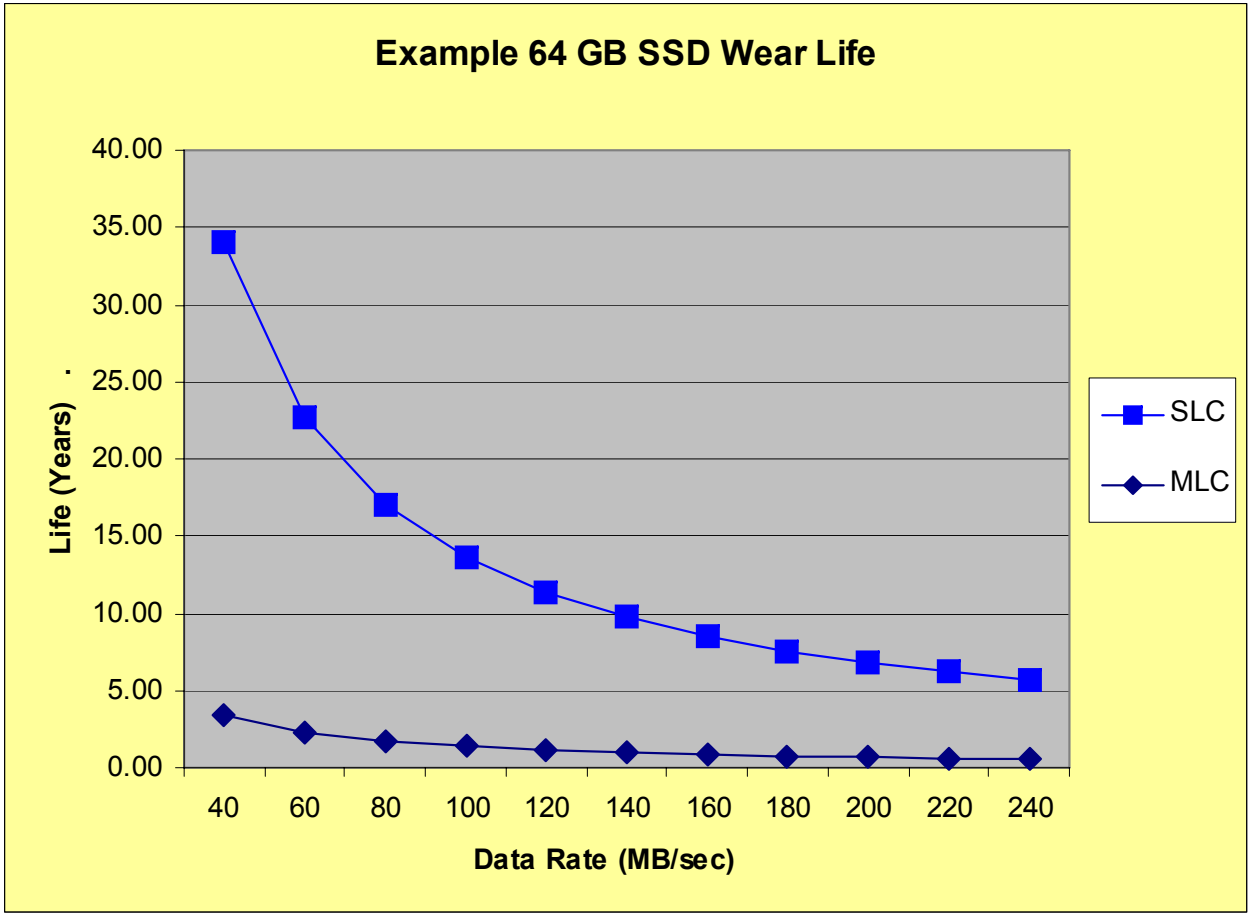

© 2008 Storage Networking Industry Association. All Rights Reserved.

System Design Considerations

- System Design Example
  - IO Workload
    - 75/25 R/W Mix at 4K Transfer lengths
    - 16 commands queued
  - SSD A IOPs = 16000
  - FC HDD IOPs = 310
- Prices
  - 73 GB FC HDD - $250 – price from web search
  - 64 Gb SSD A - $60/GB from IDC chart for enterprise drive price

Drive	Cap (GB)	IOPs	Price / GB	Drive Cost	$$ / IOP
15K RPM HDD	73	310	$5.00	$250.00	$0.81
SSD A	64	16000	$60.00	$3,840.00	$0.24

© 2008 Storage Networking Industry Association. All Rights Reserved.

System Design Considerations

- System 1 – Match IOPs
- System 2 – Match Price
- System 3 – Match Capacity with 20% of Capacity Fast

	HDD System	Match IOPs	Match Price	20% SSD
# SSD	0	2	6	23
# HDD	100	0	0	80
System IOPs @ profile	31000	32000	96000	392800
Drive Cost	$25,000.00	$7,680.00	$23,040.00	$108,320.00
System Capacity (GB)	7300	128	384	7312
SSD % Original Capacity	0.0%	1.8%	5.3%	20.1%
Drive Power Used (W)	1600	16	48	1464
% Power Used	100.0%	1.0%	3.0%	91.5%

© 2008 Storage Networking Industry Association. All Rights Reserved.

Summary

- SSD have a variety of technologies and form factors
- Lower SSD pricing has expanded the SSD market potential for higher performance storage tiers
- SSD performance characteristics vary widely
- Pick the SSD to match the application

© 2008 Storage Networking Industry Association. All Rights Reserved.

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: trackstorage@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

SNIA Education Committee

Jamon Bowen
IDC
Rob Peglar
Marius Tudor
Nathalie Gonzalez
Jeff Janukowicz
Brandy Barton
Greg Shultz

© 2008 Storage Networking Industry Association. All Rights Reserved.