



Education

Solid State Storage in a Hard Disk Package

Brian McKean, LSI Corporation

SNIA Legal Notice

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

Solid State Storage in a Hard Disk Package.

- ▶ For those interested in considering flash based SSD as a replacement for Hard Disk Drives (HDD) in storage systems, this tutorial will provide some things to look at as part of that process. While using the same form factor and interfaces as HDD, SSD have a greater variation in features and performance. Without discussing specific vendors, this tutorial will show examples of these variations and discuss how they affect using the drives in storage systems.
- ▶ **Learning Objectives:**
 - ◆ Understand the variety of HDD form factor flash SSD available
 - ◆ Learn some of the ways SSD differ from HDD and from each other
 - ◆ Learn aspects of SSD features that may affect deployment

Agenda Topics

- Flash SSD
- Price
- Performance
- Capacity
- Power
- Wear Life
- System Design Considerations
- Summary

➤ This Tutorial

- ◆ SSD: SSS in Hard Disk Drive form factor
 - > With HDD interface
- ◆ Based on flash technology

➤ Not in this Tutorial

- ◆ Other solid state storage architectures
 - > PCIe card
 - > Rack Mount
 - > Memory Card
 - > USB Stick
- ◆ SSS based on DRAM

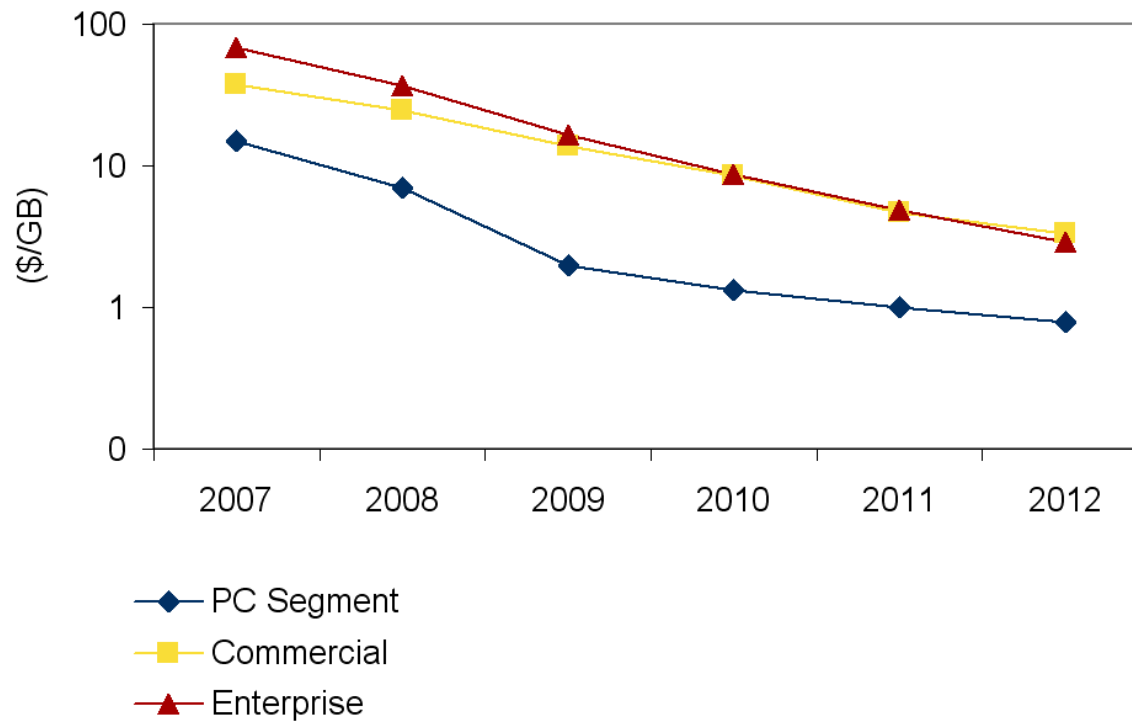
Flash SSD

As HDD replacement

- ◆ Finding the value proposition is not easy
 - ◆ SSD cost significantly more per GB than Enterprise HDD
 - ◆ The benefits of SSD are very dependent on workload characteristics
- ◆ It is \$\$/IOP
 - ◆ But finding that is not simple
 - ◆ And it's not the whole story

Flash SSD Price

Flash SSD pricing



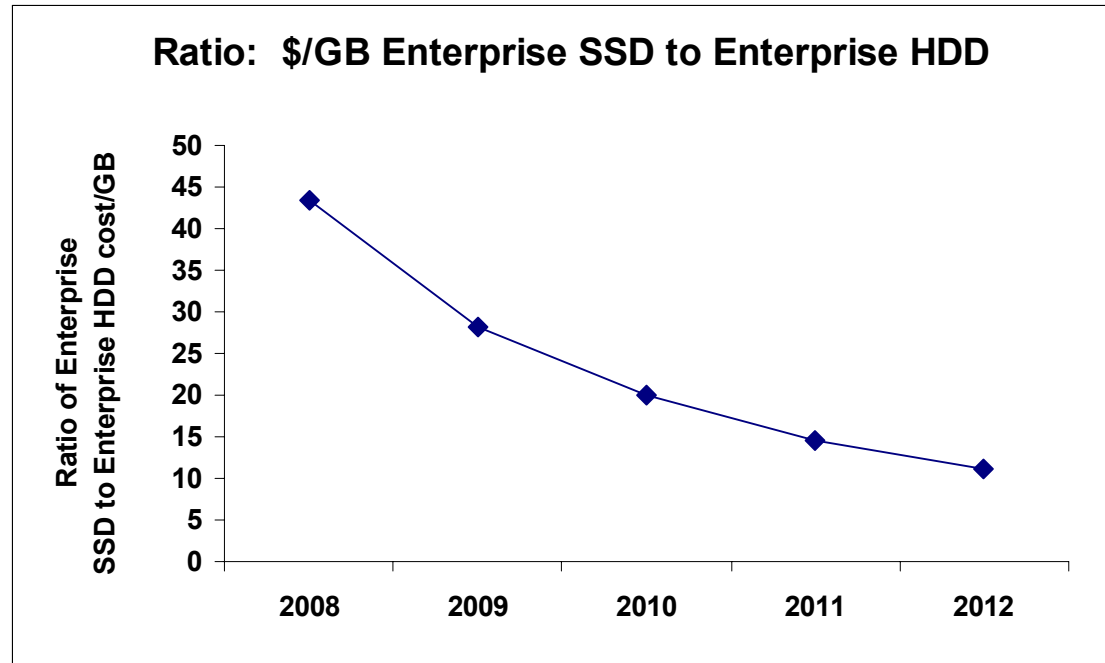
Source: Worldwide Solid State Drive 2009-2012 Forecast Update (IDC #216054, January 2009)

Flash SSD Price

- What do you get with Enterprise SSD vs. Consumer SSD?
 - ◆ Higher performance
 - › Especially higher write performance
 - ◆ Longer Life
 - › 5 years in a high IO environment vs. 3 years in a laptop
 - ◆ More attention to data integrity
 - › Better unrecoverable error rate
 - › Increased protection of the data path

Flash SSD Price

- Flash SSD pricing vs. HDD
 - ◆ Gap narrowing but not closing



Source: Worldwide Solid State Drive 2009-2012 Forecast Update (IDC #216054, January 2009) and Worldwide Hard Disk Drive 2009–2012 Forecast: Navigating the Transitions for Enterprise Applications (IDC #216394, Feb 2009) .

Flash SSD Performance

- HDD Performance today
 - ◆ Roughly correlated to rotational speed
- SSD Performance today
 - ◆ Better than HDD for small random reads
 - ◆ But if you add writes, performance:
 - › For some SSD may be slower than HDD
 - › Varies widely between drives
 - › Varies with work load
 - › Varies with previous access patterns

➤ Random Read / Write Performance

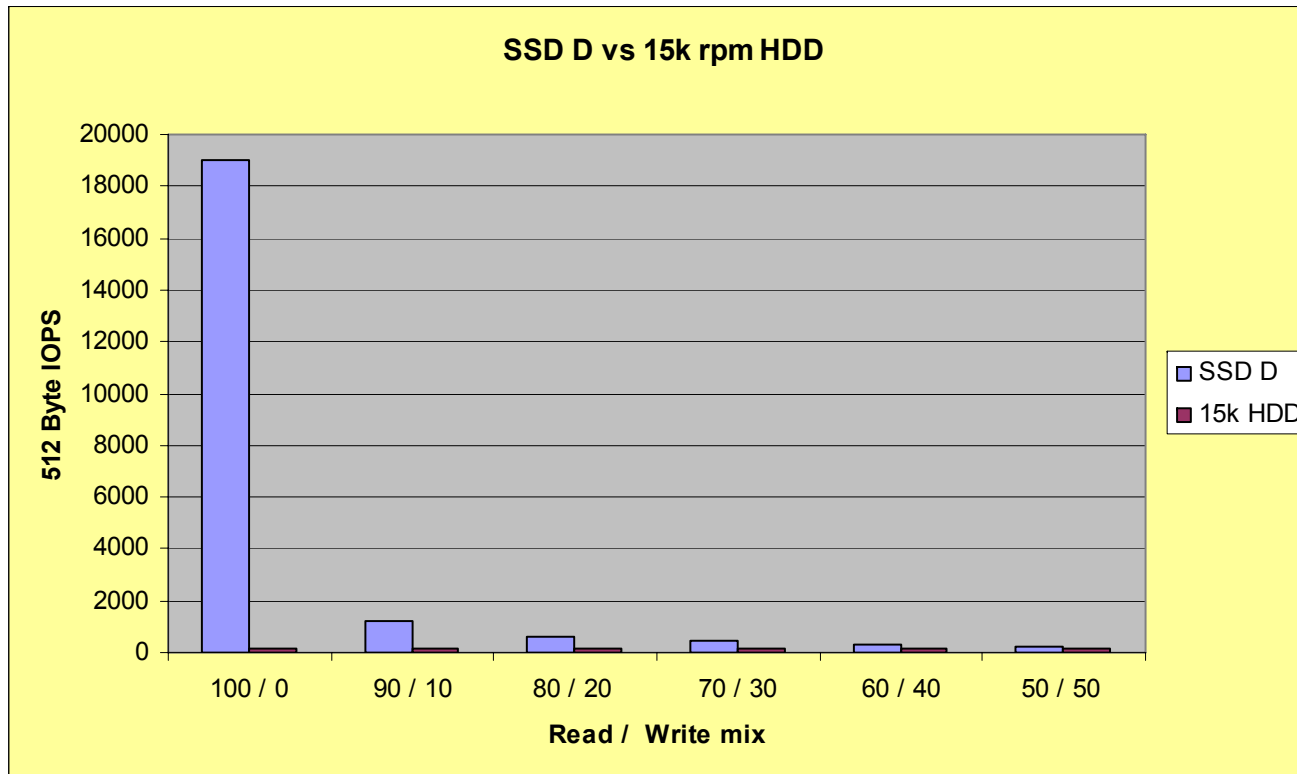
SSD Random 512 byte IOPs Performance		
	Read	Write
SSD A	45000	16000
SSD B	33000	3300
SSD C	10000	1500
SSD D	19000	130
15K rpm HDD	337	263
7.2K rpm HDD	90	79

➤ Sequential Read / Write Performance

SSD Sequential Performance MB/sec		
	Read	Write
SSD A	220	115
SSD B	250	170
SSD C	100	80
SSD D	100	80
15K rpm HDD	171	171
7.2K rpm HDD	105	105

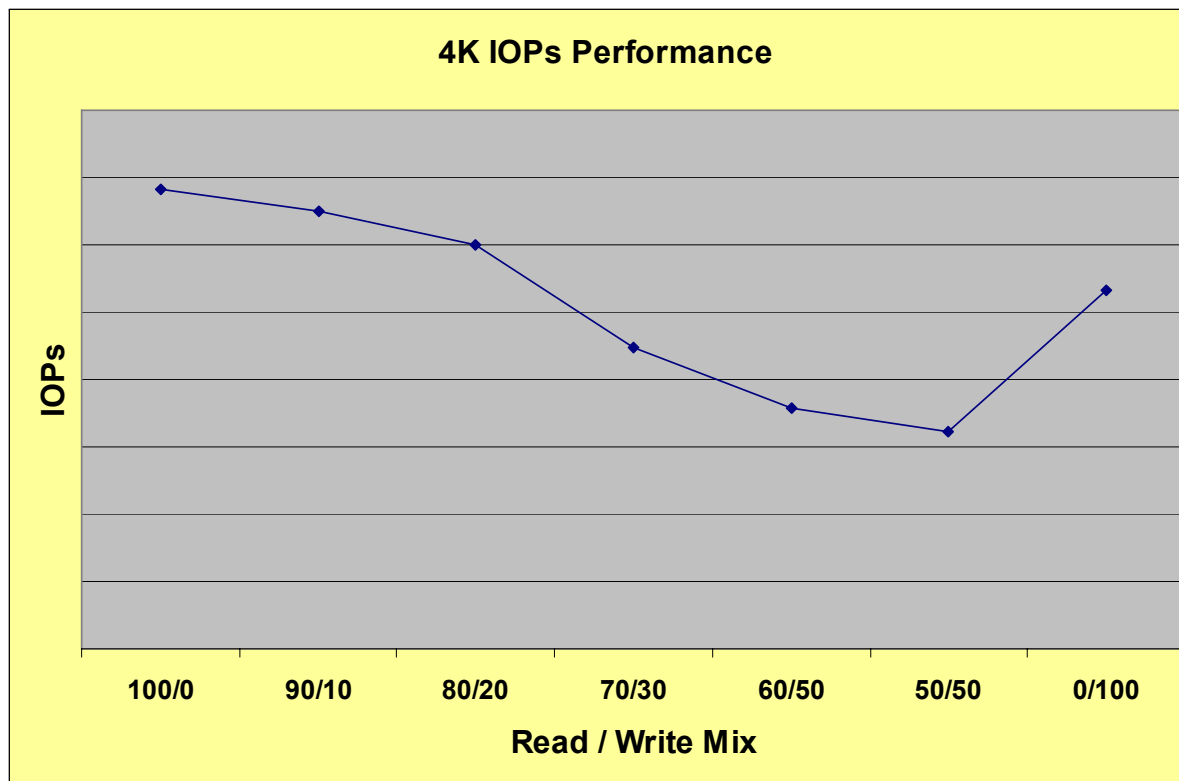
➤ SSD Performance Observations

- ◆ asymmetrical read and write performance can be dominated by time to write



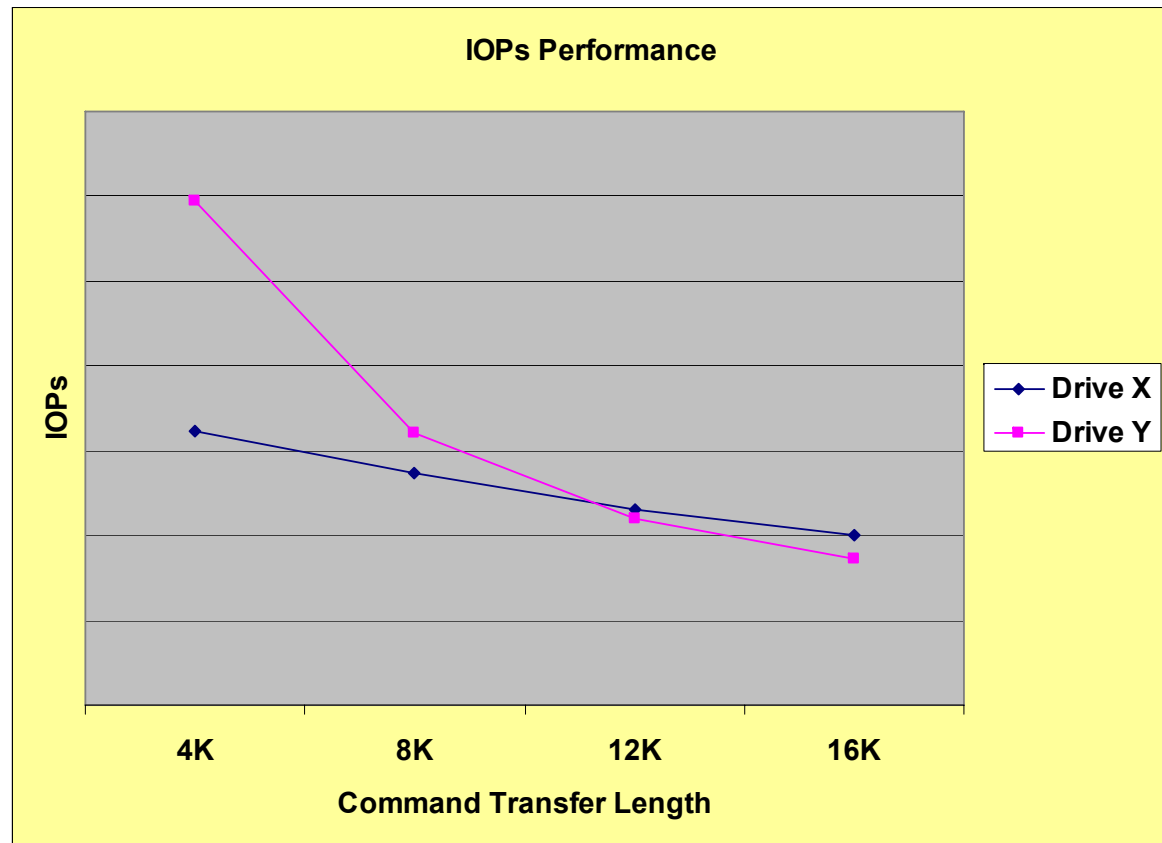
▶ SSD Performance Observations

- ◆ Mixing reads & write – for some SSD - slower than writes



➤ SSD Performance Observations

- ◆ Comparative performance between drives differs based on IO profile



➤ SSD Performance Observations

- ◆ Previous access patterns can affect current performance

512B 100% Random Write			
	After Seq Writes	After Random Writes	Drop
SSD B	2050	1470	28%
SSD C	2460	2380	3%

So how do you determine benefits?

- ◆ In addition to these interesting performance characteristics, there are other benchmarking challenges
 - ◆ An empty drive may perform much better than a full drive
 - ◆ Drives may have higher burst performance
 - › After sequential writes
 - › After idle time
 - ◆ Repeatability of performance measurements may be difficult
 - › Performance may vary depending on previous access pattern
- ◆ Know your work load! – and test to that.
- ◆ -or- sustained testing under real workload

Flash SSD Capacity

➤ Available Capacities

- ◆ 8 GB – 1 TB
- ◆ Most are 32GB – 256 GB

➤ Like HDD - Capacities vary versus advertised amount

- ◆ Actual useable bytes capacity for 64 GB drives as returned by drive

Vendor	Capacity
A	62500000000
B	62533296000
C	61734912000
D	64408576000
E drive 1	65741760000
E drive 2	65721600000

64 GB = 64,000,000,000 Bytes

Flash SSD Power

- SSD Power per IOP advantage is huge
- SSDs also have power per GB advantage over enterprise HDDs

Power	Active Watts	Capacity GB	mWatt / Read IOP	mWatt / GB	Scaled W / Drive
SSD A	7.7	146	0.17	52.7	51%
SSD B	2.6	64	0.08	40.6	17%
SSD C	2.1	64	0.21	32.8	14%
SSD D	3.7	64	0.19	57.8	25%
15K rpm HDD	15.0	146	44.51	102.7	100%
7.2K rpm HDD	11.6	750	128.89	15.5	77%

Flash SSD Wear Life

- Factors affecting wear life
 - ◆ Flash Technology
 - › Single Level Cell (SLC) – 100,000 P/E cycles
 - › Multi-Level Cell (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

Flash SSD Wear Life

- But ... Do you really care a about MLC vs. SLC?
- The real question
 - ◆ For my application and workload what should I expect for the wear life of the SSD?
 - ◆ How is SSD wear out handled by the system?
- A typical spec ...
 - ◆ “>140 years @ 50GB write/day for 32GB SSD”
 - ◆ “1 petabyte of lifetime random writes”
 - ◆ “2 Million write erase cycles”

System Design Considerations

System Design Example

- ◆ **IO Workload**
 - 75/25 R/W Mix at 4K Transfer lengths
 - 16 commands queued
- ◆ SSD B IOPs = 6450
- ◆ FC HDD IOPs = 310

Prices

- ◆ 146 GB FC HDD - \$280 – price from web search (2/11/2009)
- ◆ 64 GB SSD B – \$805 price from web search (2/11/2009)

Drive	Cap (GB)	IOPs	Price / GB	Drive Cost	\$\$ / IOP
15K RPM HDD	146	310	\$1.92	\$280.00	\$0.90
SSD B	64	6450	\$12.58	\$805.00	\$0.12

System Design Considerations


- While \$\$/IOPs is key, the typical IT constraints are
 - ◆ Same performance, lower cost
 - ◆ Better performance, same cost

System Design Considerations

▶ Baseline System – 100 15K RPM HDD

Baseline

	HDD System
SSD B	0
15K RPM HDD	100
System IOPs	31000
Drive Cost	\$28,000.00
System Capacity (GB)	14600
SSD % Capacity	0%
Capacity % of base	100%
Drive Power Used (W)	1500



System Design Considerations

➤ System 1 – Same performance – all SSD

	HDD System	Same IOPs
SSD B	0	5
15K RPM HDD	100	0
System IOPs	31000	32250
Drive Cost	\$28,000.00	\$4,025.00
System Capacity (GB)	14600	320
SSD % Capacity	0%	100%
Capacity % of base	100%	2%
Drive Power Used (W)	1500	13

System 1




System Design Considerations

➤ System 2 – Same cost – all SSD

	HDD System	Same Cost
SSD B	0	34
15K RPM HDD	100	0
System IOPs	31000	219300
Drive Cost	\$28,000.00	\$27,370.00
System Capacity (GB)	14600	2176
SSD % Capacity	0%	100%
Capacity % of base	100%	15%
Drive Power Used (W)	1500	88.4

System 2





System Design Considerations

- System 3 – Same capacity – 20% of capacity is made faster using SSD / 80% capacity on HDD

System 3

	HDD System	Same Cap.
SSD B	0	46
15K RPM HDD	100	80
System IOPs	31000	321500
Drive Cost	\$28,000.00	\$59,430.00
System Capacity (GB)	14600	14624
SSD % Capacity	0%	20%
Capacity % of base	100%	100%
Drive Power Used (W)	1500	1319.6





System Design Considerations

- System 4 – Same capacity – 20% of capacity is made faster using SSD / 80% capacity on Near Line SATA 7.2K HDD

Drive	Cap (GB)	IOPs	Price / GB	Drive Cost	\$\$ / IOP
15K RPM HDD	146	310	\$1.92	\$280.00	\$0.90
SSD B	64	6450	\$12.58	\$805.00	\$0.12
7.2K SATA HDD	750	61	\$0.13	\$100.00	\$1.64

System 4



	HDD System	SSD + NL SATA
SSD	0	46
HDD	100	16
System IOPs	31000	297676
Drive Cost	\$28,000.00	\$38,630.00
System Capacity (GB)	14600	14944
SSD % Capacity	0%	20%
Capacity % of base	100%	102%
Drive Power (W)	1500	305

System Design Considerations

Summary comparison

	HDD System	Same IOPs	Same Cost	Same Cap.	Use SATA
SSD B	0	5	34	46	46
HDD	100	0	0	80	16
System IOPs	31000	104%	707%	1037%	960%
Drive Cost	\$28,000.00	14%	98%	212%	138%
System Capacity (GB)	14600	2%	15%	100%	102%
Drive Power (W)	1500	1%	6%	88%	20%

System Design Considerations

➤ Other potential SSD benefits for TCO

- ◆ Fewer drives
 - › More capacity efficient RAID levels (RAID 5 versus RAID 10)
- ◆ Potentially fewer enclosures
- ◆ Lower power per drive



**Check out SNIA Tutorial:
Economics of Solid State
Storage: Perception and
Reality**

Summary

- A variety of flash SSD are available today
- How do you know if they'll work for you?
 - ◆ Know your workload
 - ◆ Determine sustained SSD performance with workload
- Best opportunities:
 - ◆ Small random IOs
 - ◆ Short stroke HDDs
 - ◆ Less capacity efficient RAID levels
- SSD performance characteristics vary widely
 - ◆ Pick the SSD that best fits the workload
- The price gap is narrowing
 - ◆ If the cost/benefit doesn't work today, it may work tomorrow

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

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

SNIA Education Committee

**IDC
Jeff Janukowicz
Khaled Amer
Mark Nossokoff
Jonathan Thatcher**

**Woody Hutsell
Marius Tudor
Karen Huffstutter
Ross Zwisler
Rob Peglar**