SNIA Education

Are SSDs Ready for Enterprise Storage Systems

Anil Vasudeva, President & Chief Analyst, **IMEX Research**

SNIA Legal Notice

Education SNIA

- 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

Education SNIA

- Are SSDs Ready for Enterprise Storage Systems

Computer architects dream of storage devices for their applications/workloads which can provide very high IOPs at minimal cost (IOPS/$/GB) and fast access (low latency). "Enterprise-Ready SSDs"©IMEX have started to fulfill that promise as they segment into SATA and PCIe based Storage products. A major factor for their quick adoption has been the advent of new controllers and firmware which have allowed them to transparently mitigate early issues related to reliability, endurance, data retention, performance, ease of management and interoperability with exiting storage interfaces. But their real success in enterprise adoption comes from Automated Storage Tiering activated by monitoring workload I/O access signatures and behavior over time and then non-disruptive migration of hot data to SSDs, resulting in over 200% improvement in IOPS and 80% improvement in response time at peak loads.

Learning Objectives:
- The presentation provides an overview of SSD technology, storage characteristics and applications that benefit the most from its usage. It also provides techniques for workloads optimization using automated smart-tiering and system implementation in enterprise storage systems together with economics of SSDs usage in real life.
- The presentation illustrates how optimally selected hybrid storage of SSDs and HDDs can achieve 65% lower TCO, 475% higher IOPS and 165% lower footprint while achieving a whopping 800% in $/IOPs in SANs and other storage systems under different scenarios.

Agenda

- IT DataCenter & Cloud Infrastructure Roadmap
- Storage Usage Patterns – Issues & Requirements
- NextGen SSDs for Enterprise Storage Systems
- Enterprise SSD Market/Product Segments by Interfaces
- SSD vs. HDDs vs. Hybrids - Price/Perf/Availability
- SLC vs. MLC SSDs – Technologies, Drivers & Challenges
- New Intelligent Controllers – Key for SSD Adoption
- AutoSmart Storage-Tiering Software – Usage & Impact
- Applications best suited for SSDs
- Key Takeaways

IT DataCenters & Cloud Infrastructure

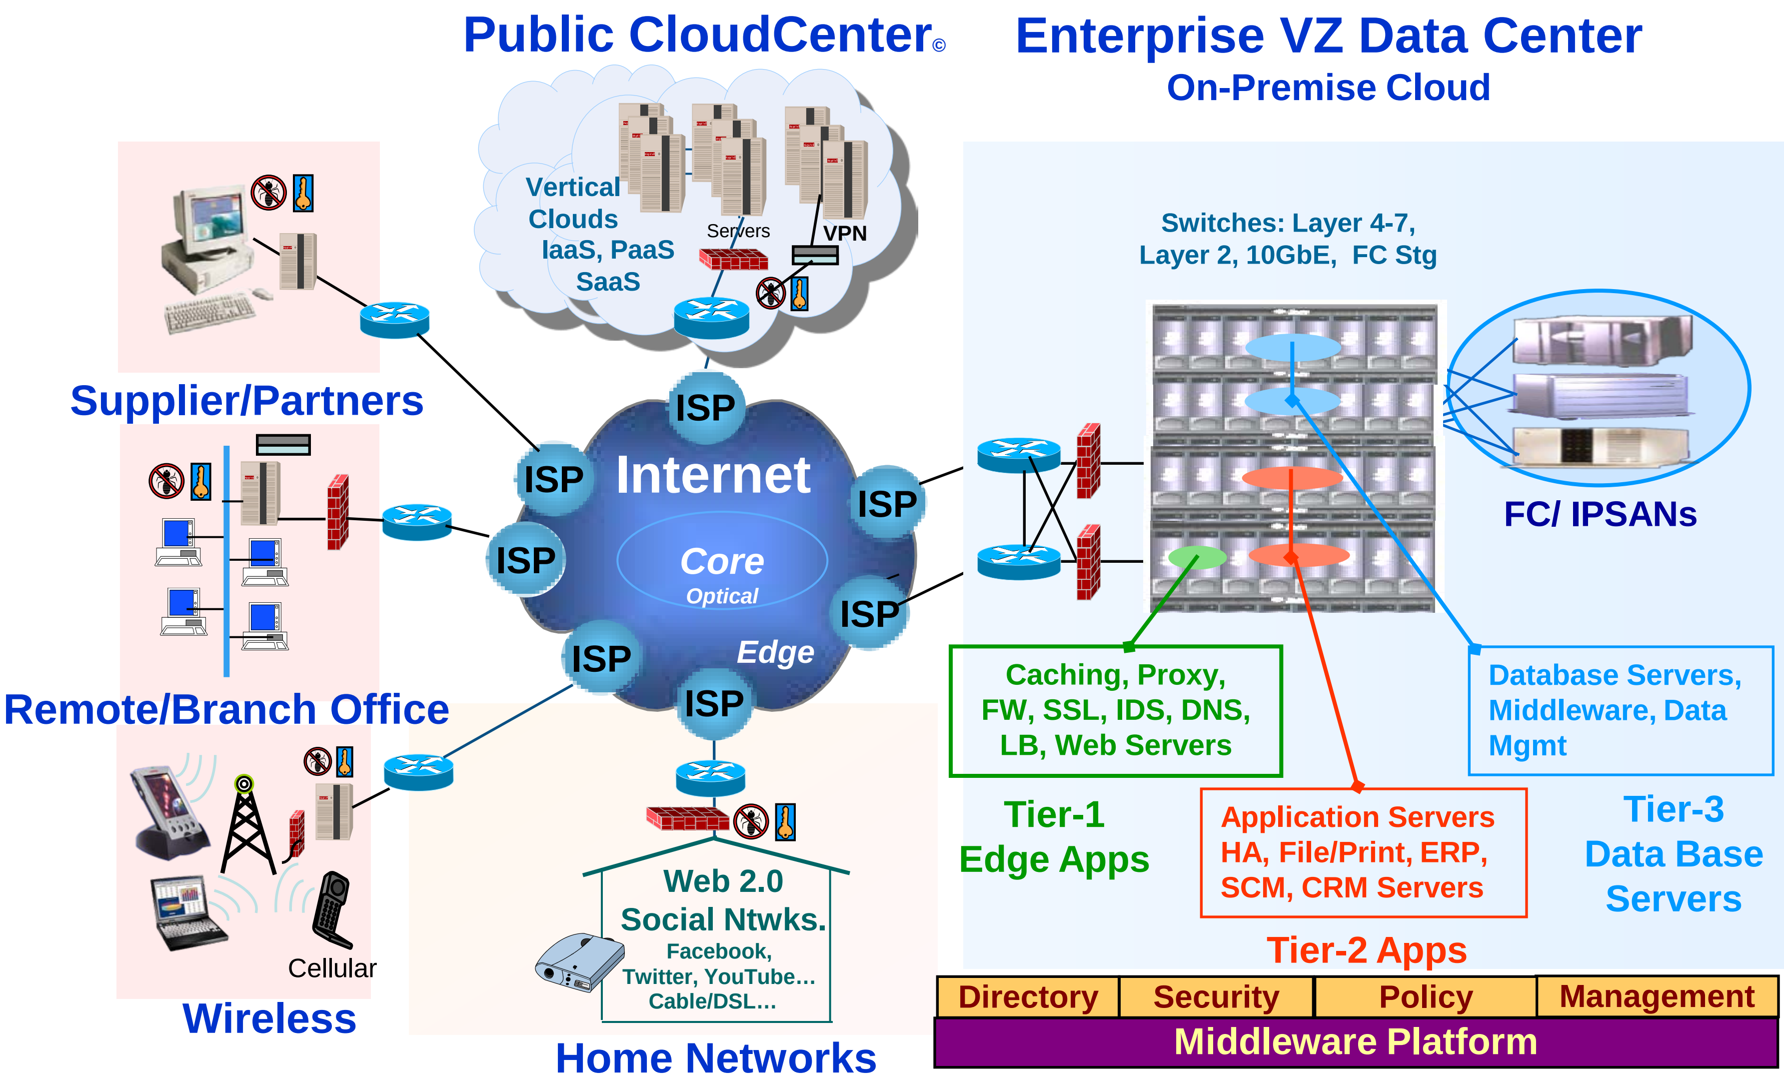

Data Source: IMEX Research
Cloud Infrastructure Report ©2009-11

IT Industry's Journey - Roadmap

Education SNIA

SIVAC

®IMEX

Cloudization

On-Premises > Private Clouds > Public Clouds

DC to Cloud-Aware Infrast. & Apps. Cascade migration to SPs/Public Clouds.

Automation

Automatically Maintains Application SLAs

(Self-Configuration, Self-Healing©IMEX, Self-Acctg. Charges etc)

Virtualization

Pools Resources. Provisions, Optimizes, Monitors

Shuffles Resources to optimize Delivery of various Business Services

Integration/Consolidation

Integrate Physical Infrast./Blades to meet CAPSIMS ®IMEX

Cost, Availability, Performance, Scalability, Inter-operability, Manageability & Security

Standardization

Standard IT Infrastructure- Volume Economics HW/Syst SW

(Servers, Storage, Networking Devices, System Software (OS, MW & Data Mgmt SW)

Data Source: IMEX Research
Cloud Infrastructure Report ©2009-11

Market Segments by Applications

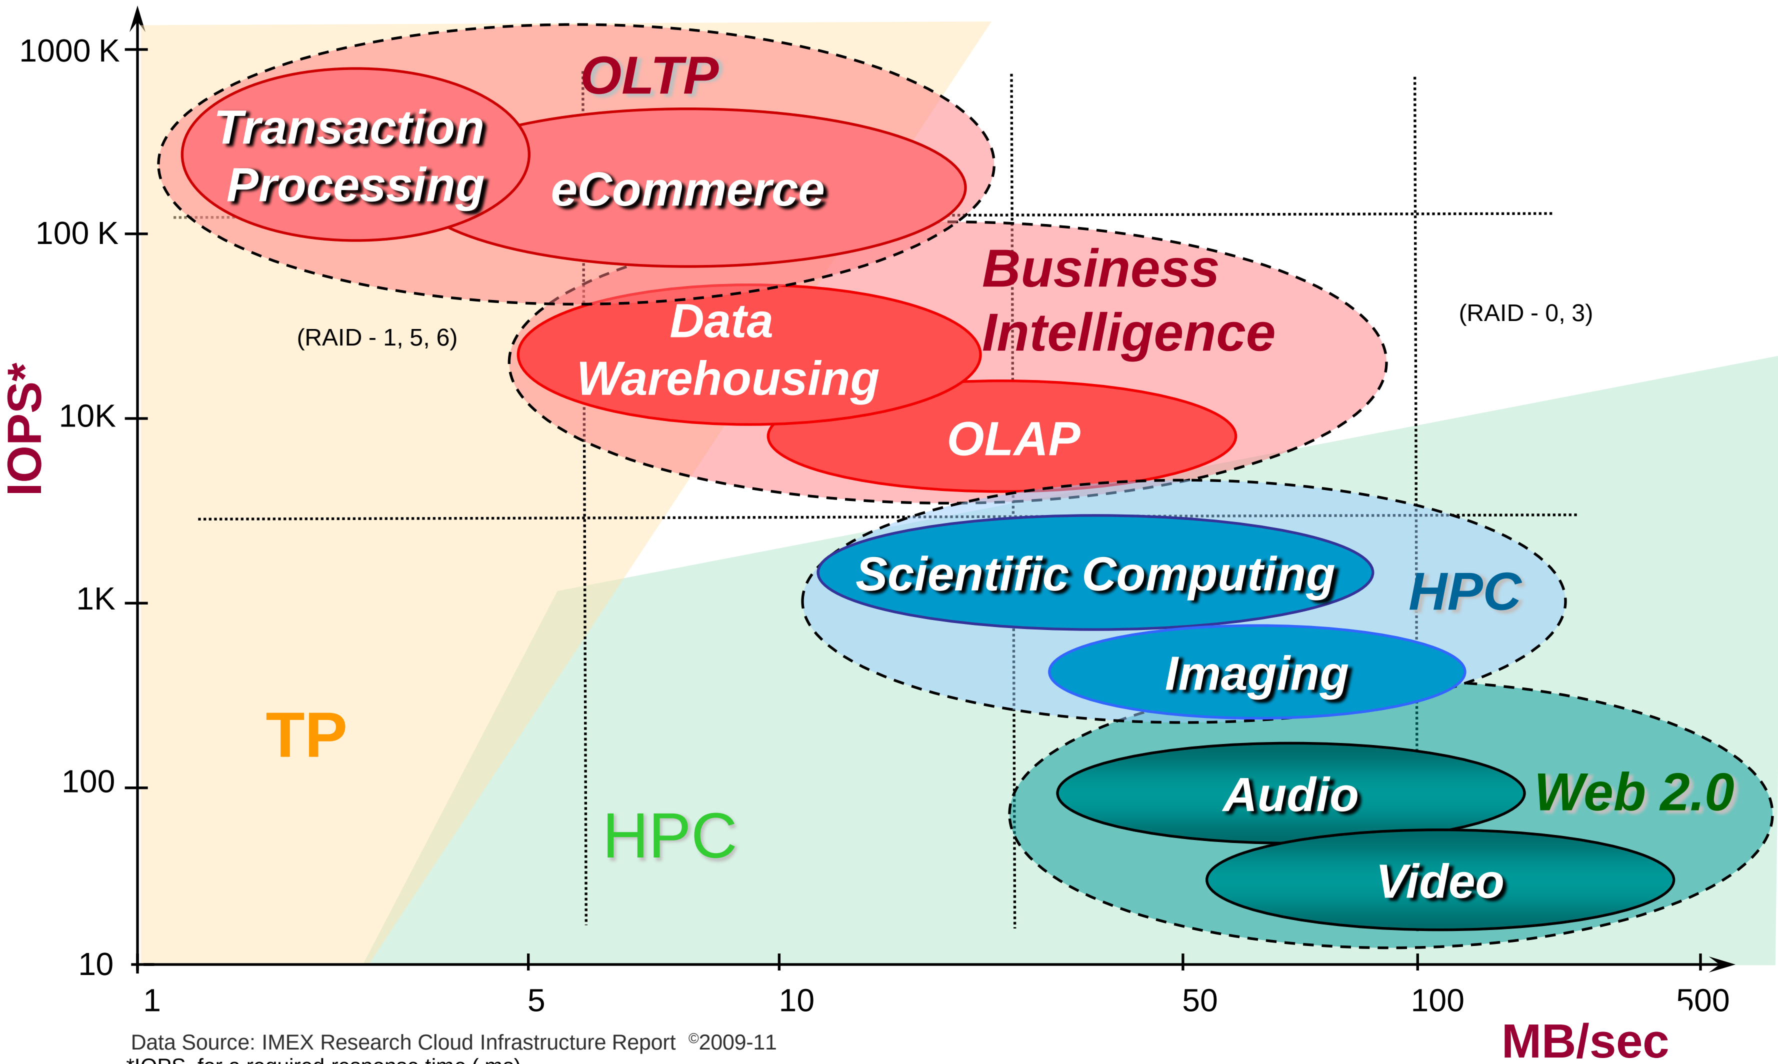

Data Source: IMEX Research Cloud Infrastructure Report ©2009-11
*IOPS for a required response time (ms)
*=(#Channels*Latency-1)

Corporate DataCenter Storage Usage

Education SNIA

I/O Access Frequency vs. Percent of Corporate Data

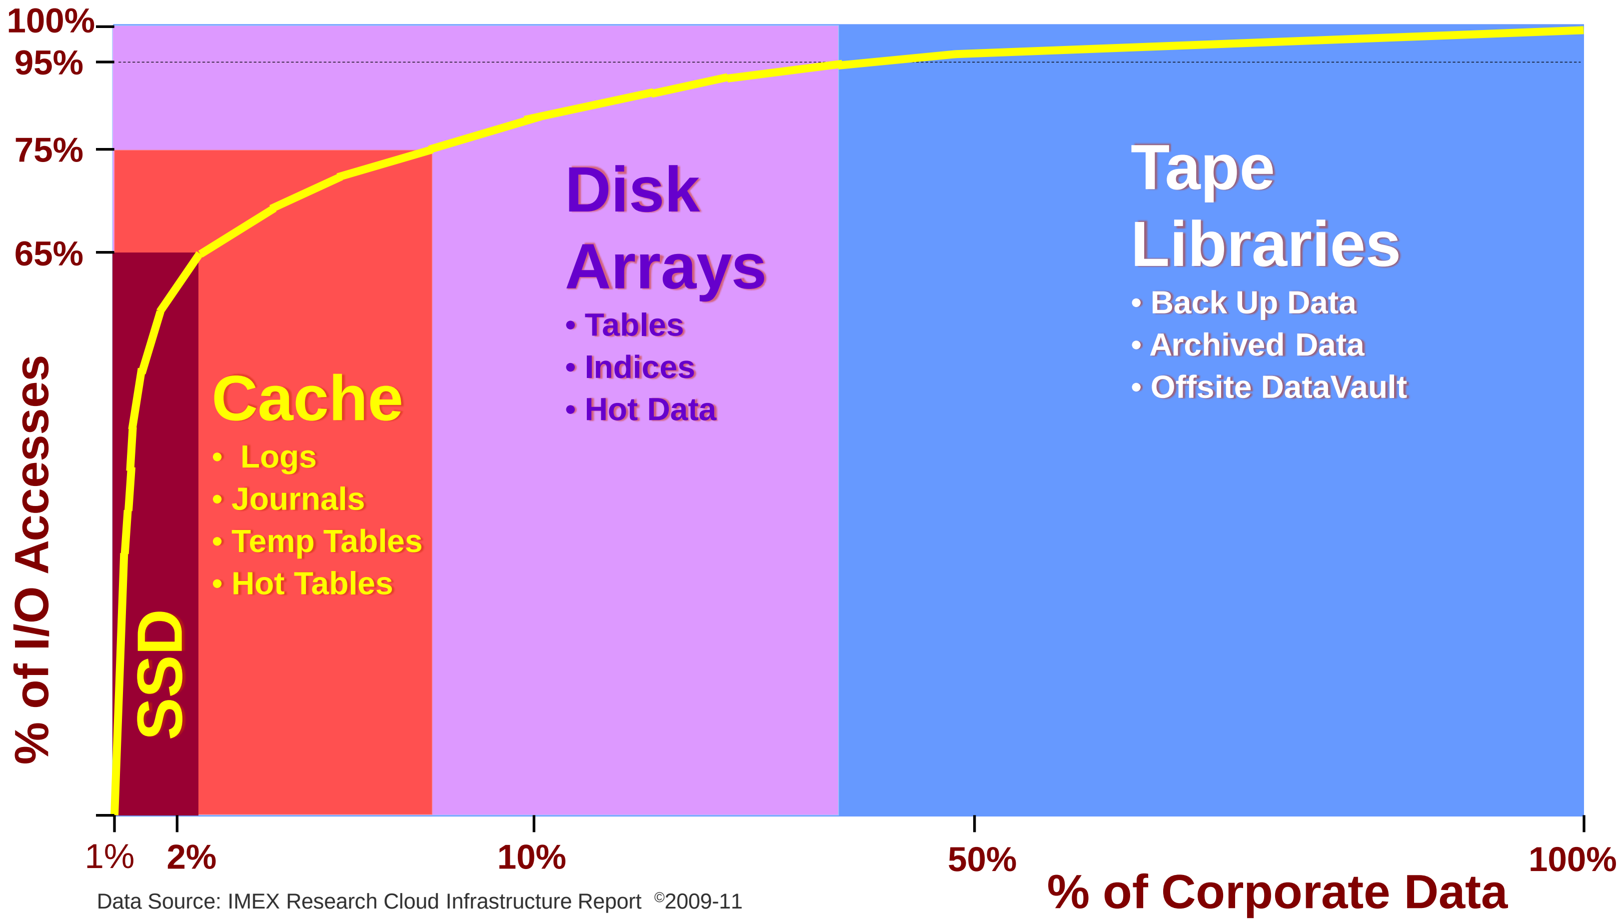

Data Source: IMEX Research Cloud Infrastructure Report ©2009-11

Cloud MegaDataCenter© Storage Usage

Education SNIA

I/O Access Frequency vs. Percent of Corporate Data

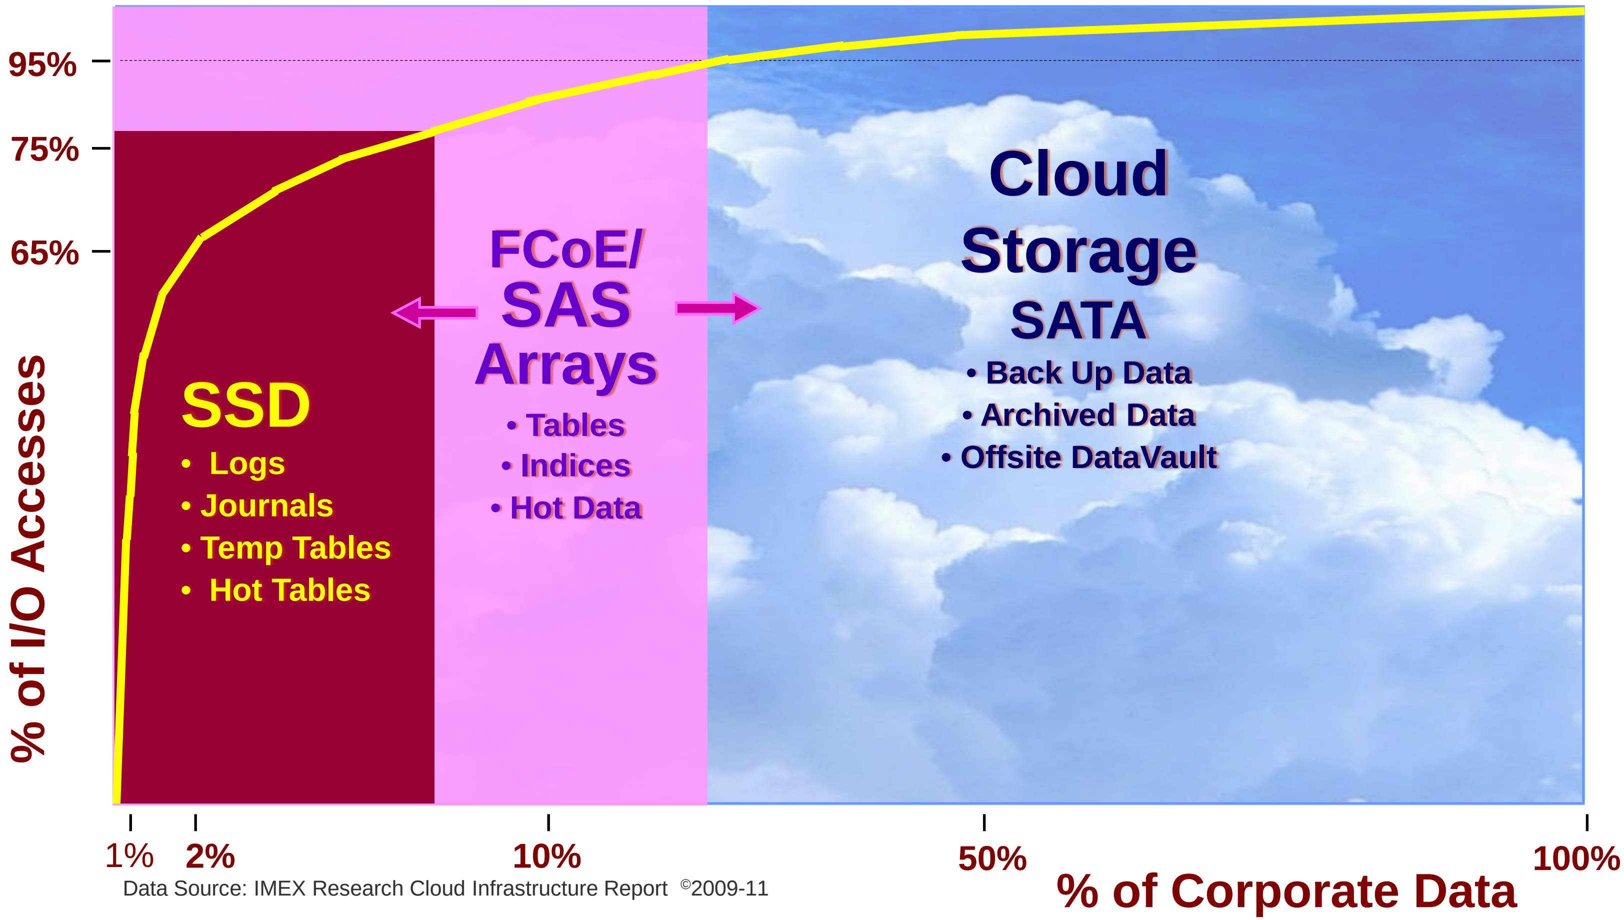

Data Source: IMEX Research Cloud Infrastructure Report ©2009-11

Data Storage Usage – Access & Longevity

Education SNIA

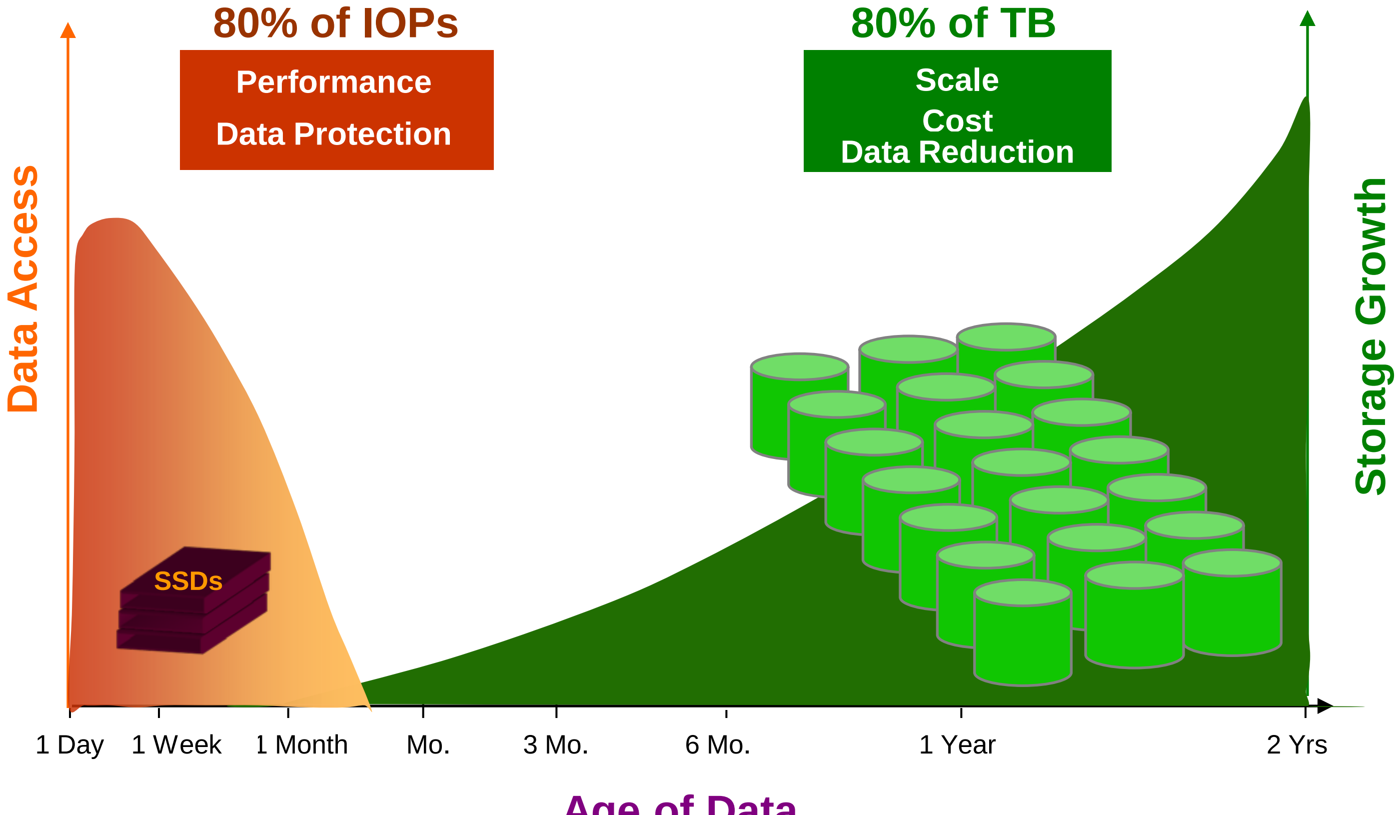

Data Source: IMEX Research Cloud Infrastructure Report ©2009-11

NAND: Enabling Consumer to Enterprise Mkts

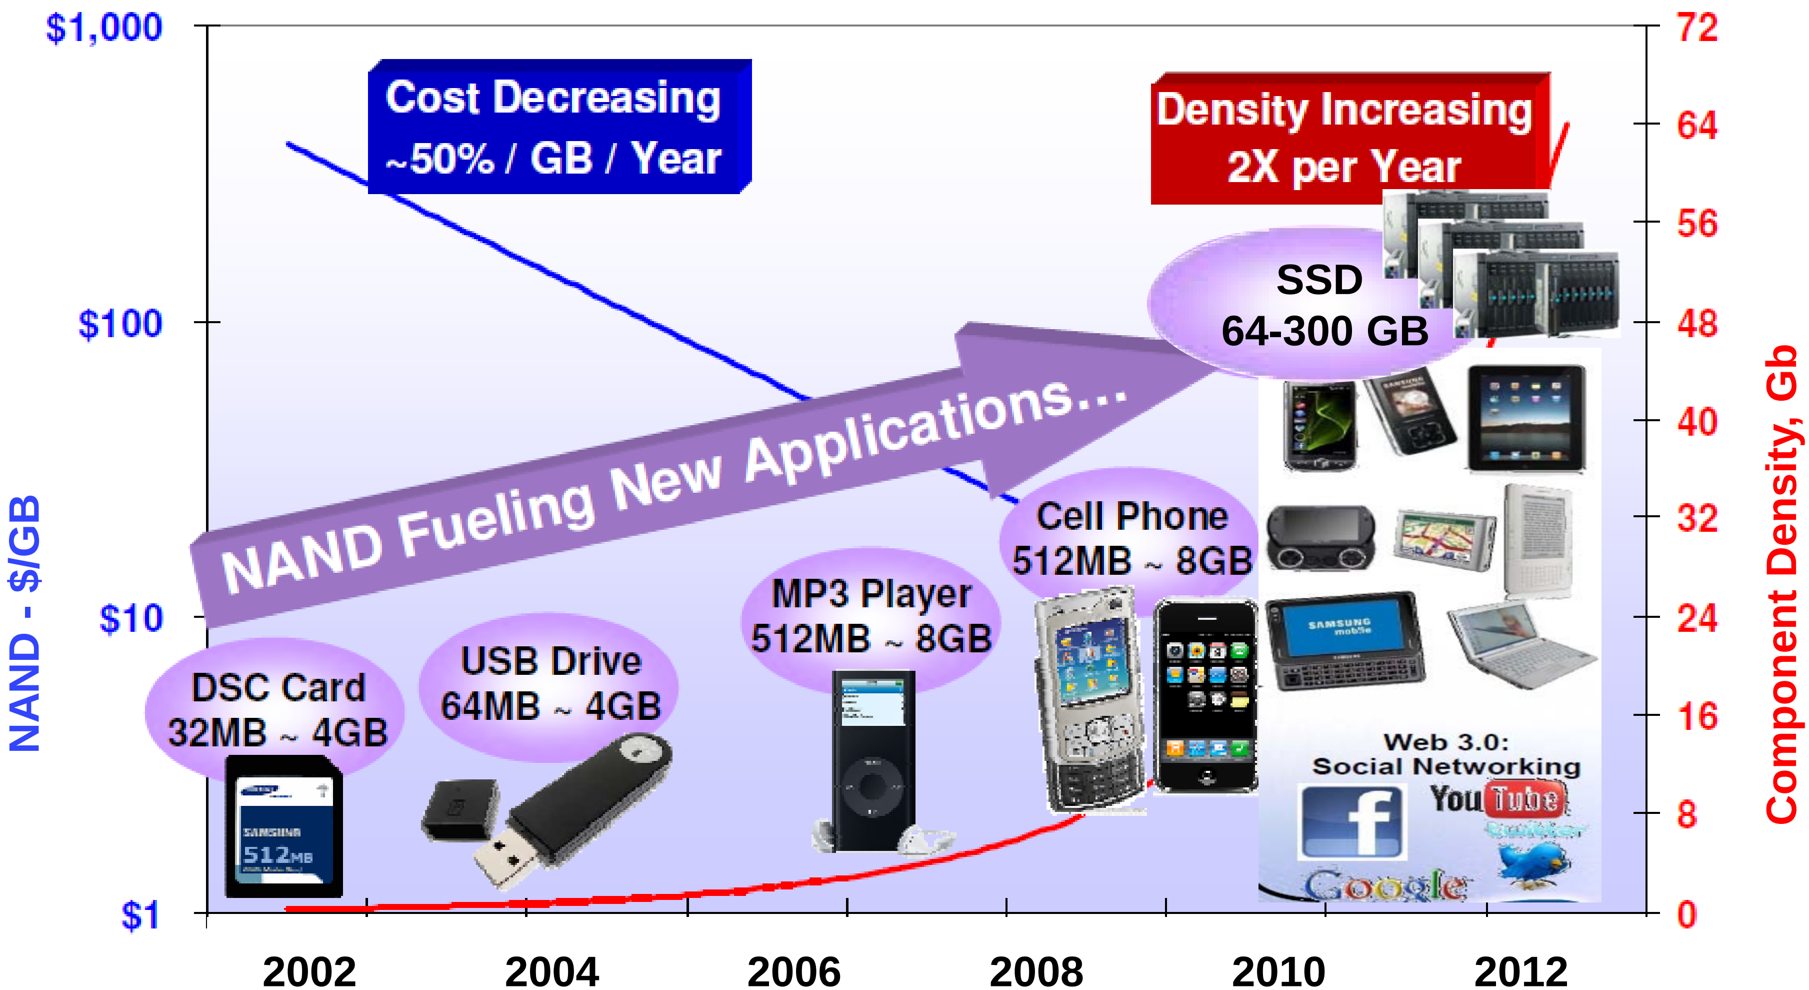

Source: Samsung & IMEX Research

Enterprise SSDs Trends - Cost

Education SNIA

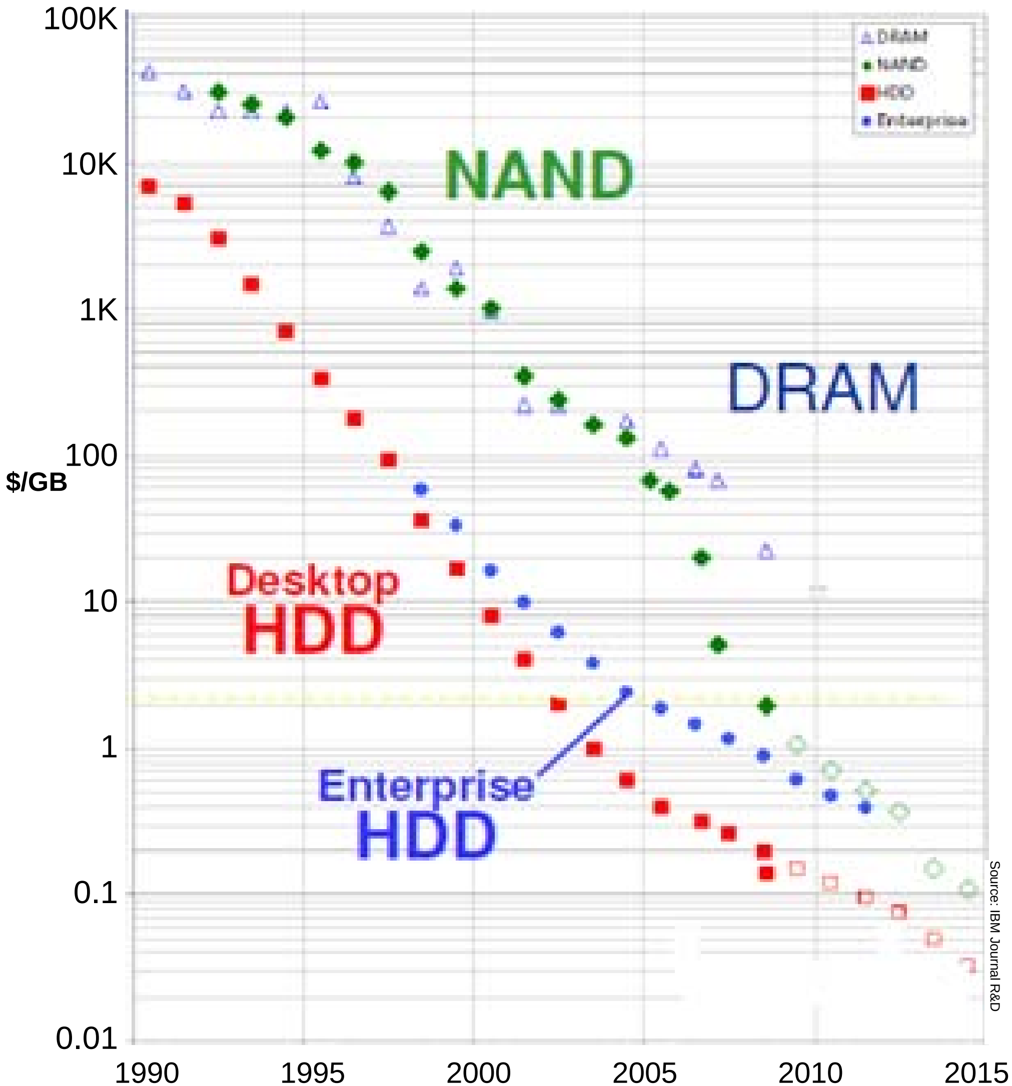

Price Erosion Trends

- Driven by an explosion in the use of cost-sensitive handheld mobile devices, MLC NAND has seen an explosive growth.
- On enterprise side Clustered low cost servers used in multiple environments from DB to BI to HPC applications besides being driven by Cloud Service Providers are providing an overall growth of 107% cagr in Computing SSDs GB
- SSD units are forecasted to grow at 86% cagr during the 2010-14 time frame.

SSD Filling Price/Perf Gaps in Storage

Education SNIA

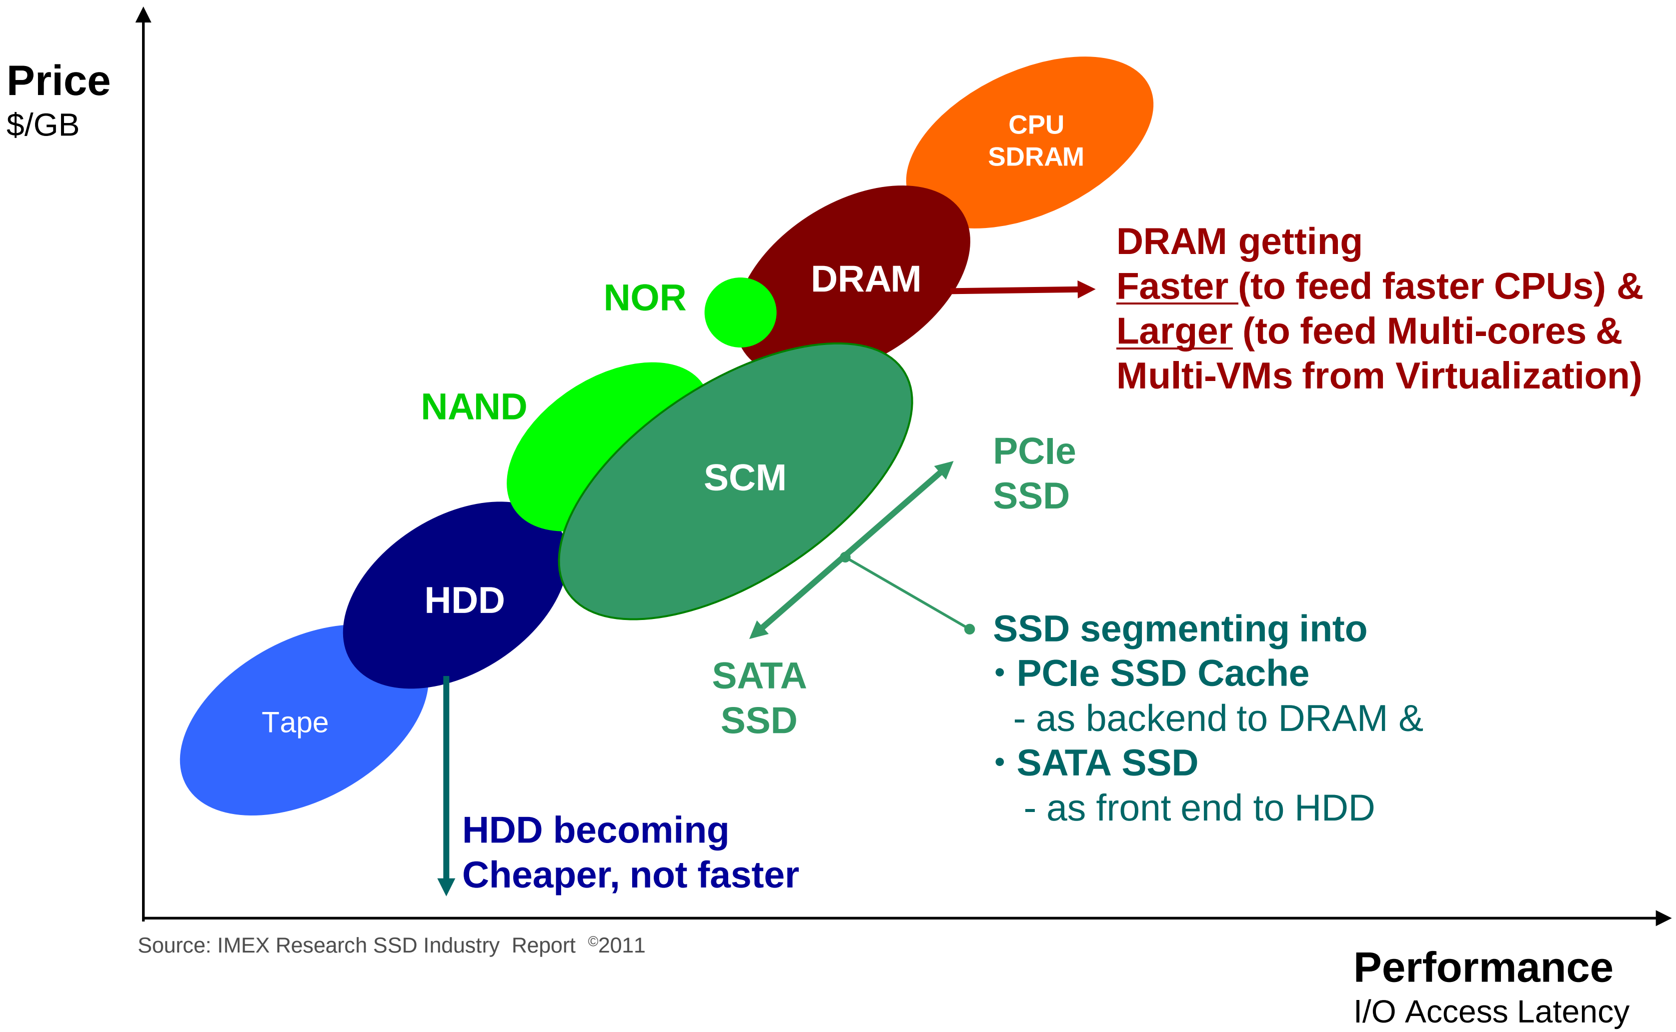

Source: IMEX Research SSD Industry Report ©2011

SCM – A new Storage Class Memory

SCM (Storage Class Memory)

Solid State Memory filling the gap between DRAMs & HDDs
Marketplace segmenting SCMs into SATA and PCIe based SSDs

Key Metrics Required of SCMs

Device - Capacity (GB), Cost ($/GB),

- **Performance** - Latency (Random/Block RW Access-ms);
  - Bandwidth W(R/W- GB/sec)
- **Data Integrity** - BER (Better than 1 in 10^17)
- **Reliability** - Write Endurance (No. of writes before death); Data Retention (Years); MTBF (millions of Hrs),
- **Environment** - Power Consumption (Watts); Volumetric Density (TB/cu.in.); Power On/Off Time (sec),
- **Resistance** - Shock/Vibration (g-force); Temp./Voltage Extremes 4-Corner (°C,V); Radiation (Rad)

Advantage: SSD vs. HDD in Enterprise Storage

Education SNIA

Manufacturer's Required Specs				
Endurance	Function Failures			UBER
Enterprise		<= 3%		<=10e-16
		SSD	**HDD**	Diff. %
MTBF	Million Hr	2.1	1.0	**-110%**
Failure Rate	Per Year	<<3%	4%	**-33%**

Reliability

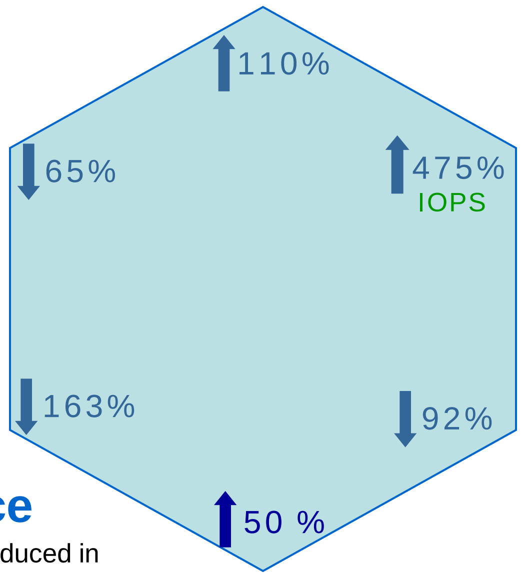

Environmental

Shock/Vibration Higher Resistance to -Shock 8x Better, Anti-Vibration 16x Better
Operating Temp Spec'd at 2x Wider Operating Temp Range
Noise 30dB Lower
Space Savings in Space $ at $/sq.ft Office Space For same IOPS, Fewer Frames, Switch Ports, Controllers, Cables, Power Supplies etc
Weight 50% Less Weight

Performance

RW Speed 5x Faster
Data Access Time <1%
Concurrent Access 900% Better
IOPS 475% Better

Power

92% Less Power, 38% Less Temp

		SSD	HDD	Svgs %
Idling	**Power** Watts	0.5	6.8	**93%**
	Temp Surf C	85	136	38%
Load	**Power** Watts	0.9	10.1	**91%**
	Temp Surf C	94	154	39%

Operating Maintenance

Maintenance & Operating Time Reduced in - Booting Up, - Virus Scan, - Defrag, - RAID Build, - Patching, - Data Restoration

TCO

50% of HDD

Source: IMEX Research SSD Industry Report 2011 ©IMEX 2009-11

Advantage: Enterprise SSDs vs. HDDs

HDD

SSD

	Parameter		Improvement SSD vs. HDD
1.0 GB/in3	Storage Density	16 GB/in3	1600 %
4.2 IOPS/in3	Performance Density	1,250 IOPS/in3	30,000 %
11.4 GB/W	Power Efficiency	570 GB/W	5,000 %
43.1 IOPS/W	Performance/Power	42,850 IOPS/W	100,000 %

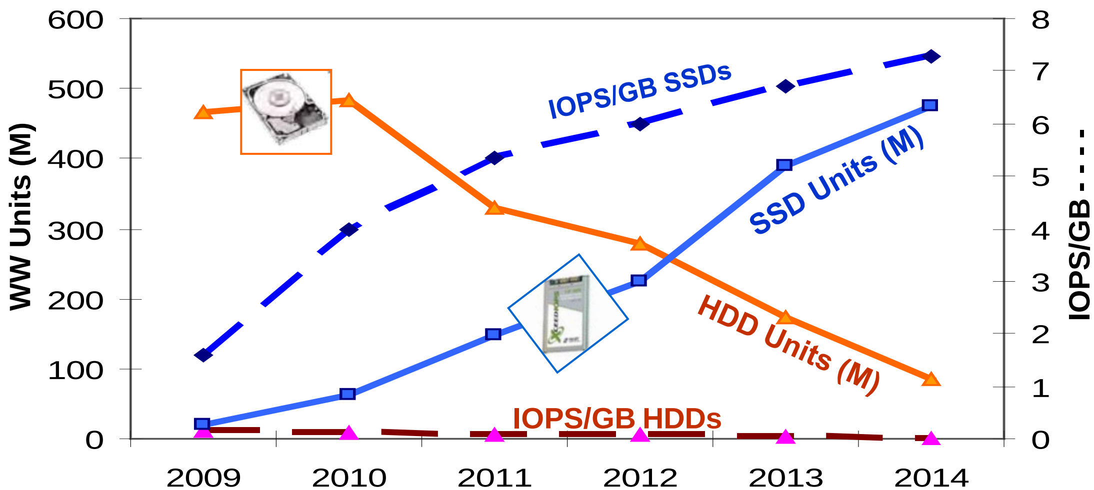

Note: 2U storage rack, • 2.5" HDD max cap = 400GB / 24 HDDs, de-stroked to 20%, • 2.5" SSD max cap = 800GB / 36 SSDs
Source: IMEX Research SSD Industry Report ©2011

Drivers & Challenges – MLC vs. SLC SSDs

Education SNIA

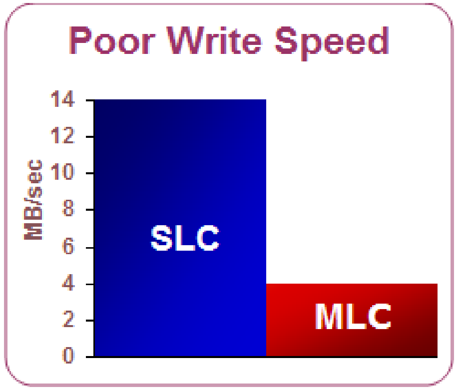

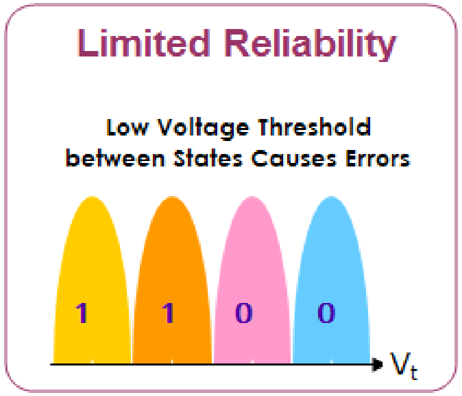

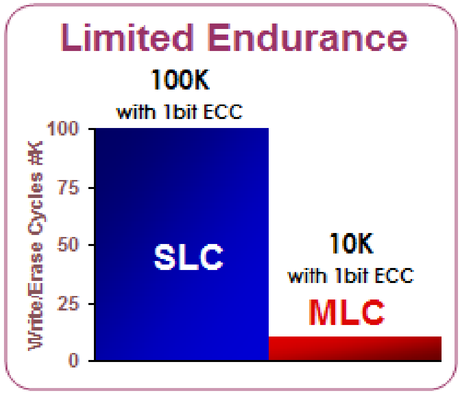

	Drivers	Challenges
Raw Media Reliability	•No moving parts •Predictable wear out •Post infant mortality catastrophic device failures rare	Higher density of MLC increases bit error rate High bit error rate increases with wear Program and Read Disturb Prevention, Partial Page Programming Data retention is poor at high temperature and wear
Media Performance	•Performance is excellent (vs.HDDs) •High performance/Watt (IOPS/Watt) •Low pin count: shared command / data bus, good balance	•NAND not really a random access device Block oriented; Slow effective write, erase/transfer/program) latency, Imbalanced R/W access speed •NAND Performance changes with wear, Some controllers do read/erase/modify/write, Others use inefficient garbage collection
Controller	•Transparently converts NAND Flash memory into storage device •Manages high bit error rate •Improves endurance to sustain a 5-year life cycle	•Interconnect •Number of NAND Flash Chips (Die); # of Buses (Real / Pipelined) •Data Protection (Int./Ext.RAID; DIF; ECC);Write Mitigation techniques •Effective Block (LBA; Sector) Size: Write Amplification •Garbage Collection (GC) Efficiency •Buffer Capacity & Management: Meta-data processing

Source: IMEX Research SSD Industry Report ©2011

MLC vs. SLC SSDs - Price Erosion

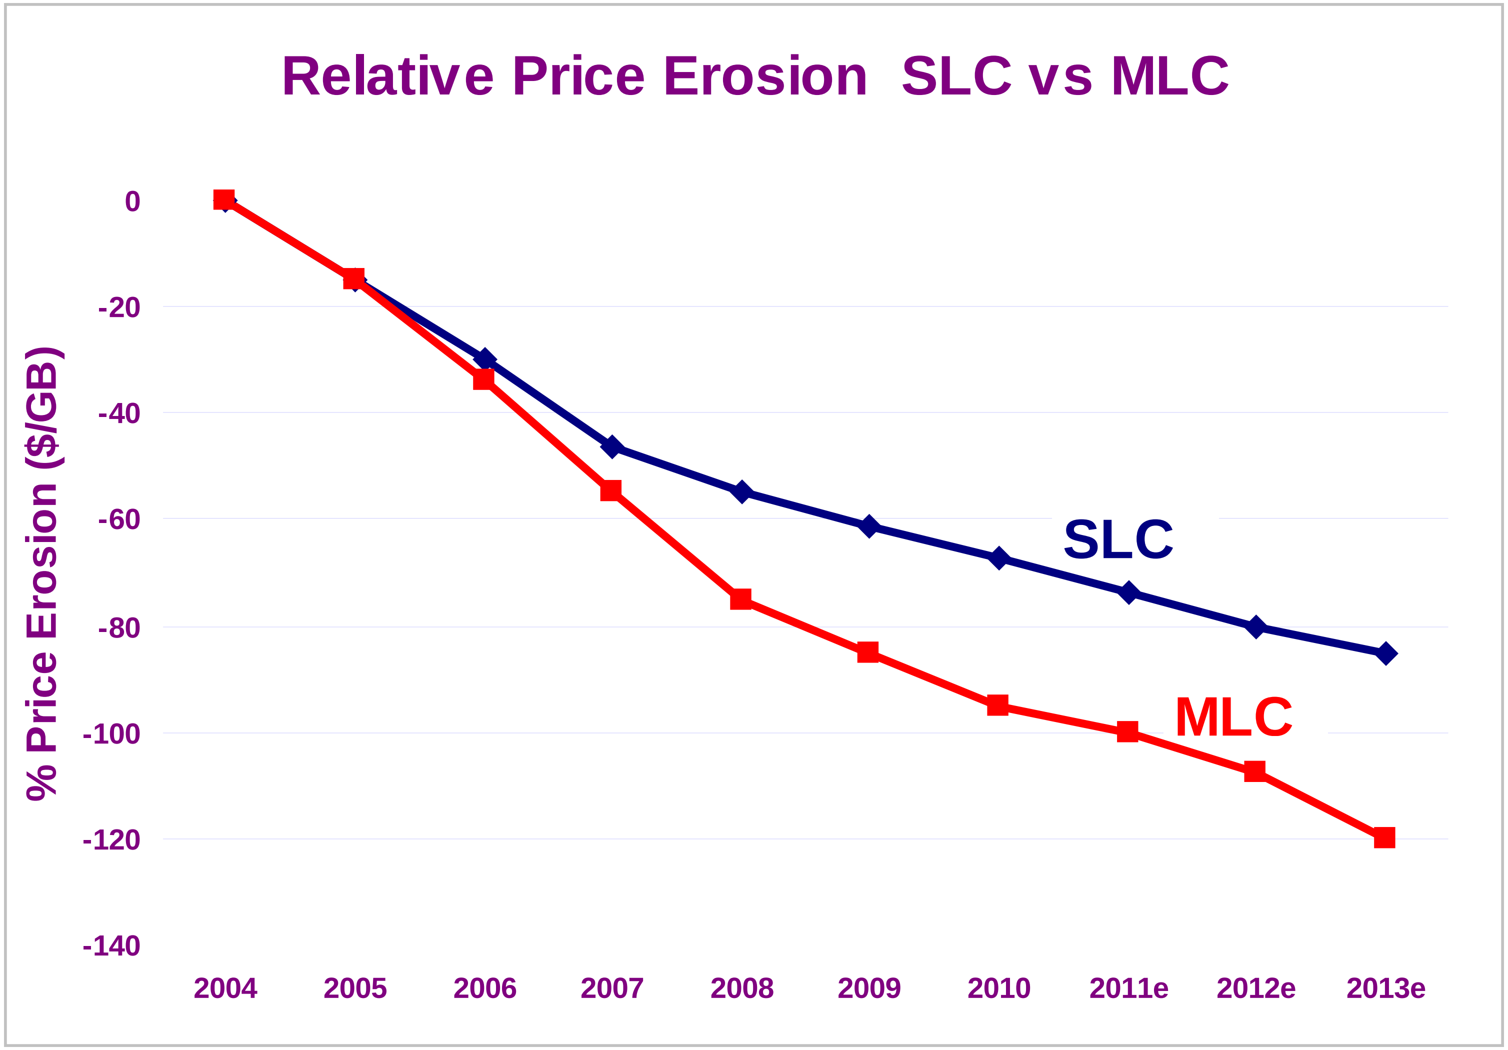

SSD Challenges & Solutions: Endurance/Wear-out

Reason for Endurance Limitation in SSDs

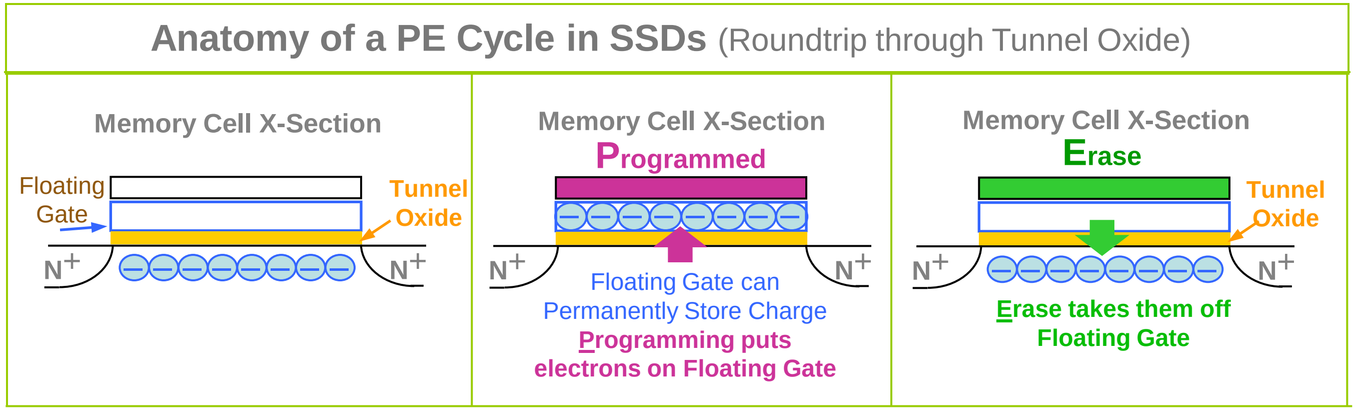

- Fundamentally – NAND Flash Memory Cell is an MOS Transistor with a Floating Gate that can permanently store charge
- Programming puts electrons in Floating Gate, Erase takes them off
- 1 Program/Erase (P/E) Cycle is a round trip by the electrons
- Electrons pass through Cell's Tunnel Oxide. Back & Forth round trips gradually damage the Tunnel Oxide over hundred thousands of trips (Program/Erase or PE cycles) resulting in Limited Endurance (or Wear-Out by PE cycles) in SSDs

SSD Challenges & Solutions: Endurance (Wear-Out)

Education SNIA

Challenge: Bad Block Mgmt

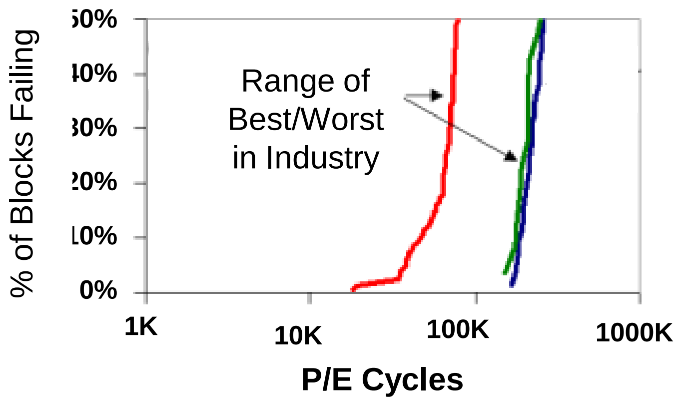

- The ability to erase slows down after a number of P/E Cycles.
- If NAND Memory block fails to erase, Controller is notified and another block from spares is used instead
- But there's no loss of data, so a failed NAND block does not pose a problem.
- Eventually devices will run out of spares
- The point where the % failing exceed number of spares is the most
  Basic Endurance Limit

Solution: Over Provisioning

- Over Provisioning by Increasing Spare blocks
  - Decreases user capacity but
  - Allows SSD to more efficiently complete random Writes
  - Improves Random Write Endurance and Performance
- Methods to Implement include:
  - Setting max LBA to limit visible drive capacity or
  - Create Smaller RAID Logical Drives or
  - Create Smaller Partitions

Promise:

- Depending on workload, endurance can vary
- Endurance should match usage needs of the system to minimize costs.
  - SSD used as cache for 10 HDDs. 2 PB writes of useful life will support this.(1.1 TB writes/day for 5 years.)

SSD Challenges & Solutions: Endurance (UBER)

Education SNIA

Challenge: Uncorrectable BER Mgmt

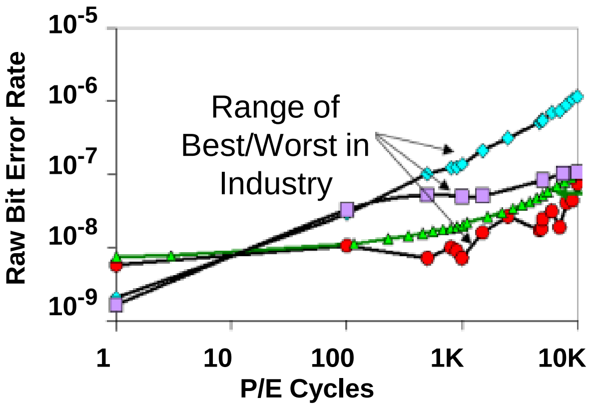

- A small of written bits gets flipped (similar to HDDs)
- This is Flash Media's **Raw Bit Error Rate (RBER)**
- ECC is used to correct/reduce this RBER
- RBER gradually increases with P/E cycles. Any bit error rate over ECC Correction capability is the **Uncorrected Bit Error Rate (UBER).** Reaching a UBER domain user data can become corrupted.
- UBER is kept low. JEDEC Spec is 1 in 10^{16} errors
- The point where UBER reaches this spec, is **Another Endurance Limit**

Solution: ECC

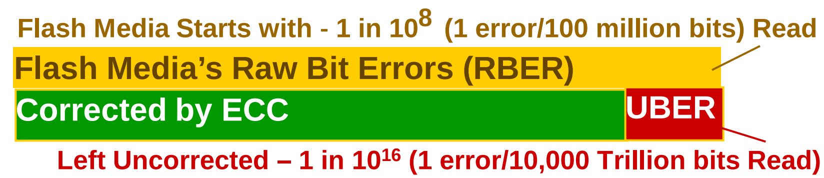

- Using modern ECC techniques based controllers, vendors are providing spec at 1 in 10^-17 UBER

SSD Challenges & Solutions: Data Retention

Education SNIA

Challenge: Data Retention

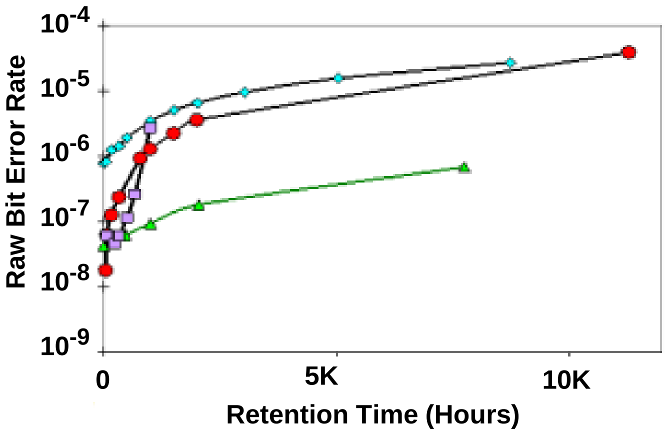

- After PE cycles, RBER increases with time. ECC corrects bit flips but only to a certain extent.
- So the industry lives with a required UBER and required Retention Time. This, in turn, determines the Safe PE cycles that device should be exercised to, prior to reaching the UBER and Retention time. This is also another endurance limit set by retention.

Solution: Data Retention Firmware

- Powered-On Firmware
  - To allow Higher Retention
- Balance out SSD Data Retention vs. Endurance
  - Lower Data Retention allows for higher endurance

SSD Challenges & Solutions: Functional Failure Defects

Education SNIA

Challenge: Electronic Component - Defects

Role of Defects in SSD Reliability

Wafer Process Defects	61%
Design Related & Test	10%
EOS/ESD	10%
Handling	9%
Process Errors	5%
Assembly & Test	5%

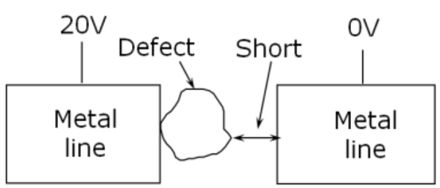

- All ICs have defects that cause failures. In Flash early life failures are caused by such defects.
- Defects can cause functional failures not just data loss. Most of NAND defect failures are caused by PE cycles, coming in from high PE voltages causing defects to short.
- The point where % failing from defects would reach unacceptable limits is another boundary for endurance.

Solution: Burn-Ins, Error Avoidance Algorithm

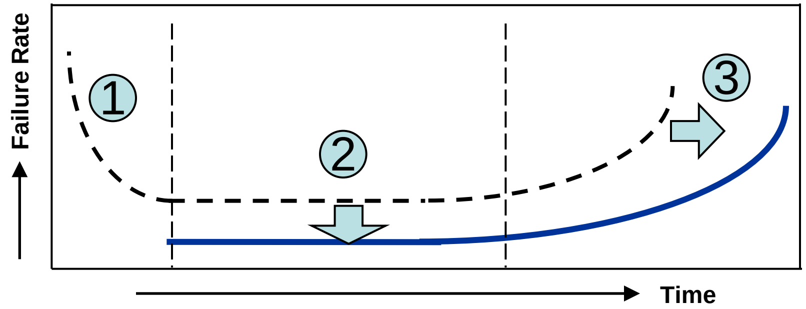

- Vigorous SSD Burn-In & Testing 1
  - Remove Infant Mortality

- Compute NAND 2
  - T_{read} to improve Read Disturbs
  - T_{PROG} to reduce Program Disturbs
- SSD Error Avoidance algorithms
  - ECC ASICS
- Wear Leveling to avoid Hot Spots
- Efficient Write Amplification Factor (WAF) 3
  - WAF=Data written to NAND /Data Written by Host to SSD
  - WAF dependent on (a) SSD FW algorithm built into SSD (b) Over Provisioning Amount (c) App Workload

SSD Challenges & Solutions: Industry Standard Testing

JEDEC Solution: Manufacturer Requirements

Class	Active Usage	Retention	Failures	UBER
	Power On	Power Off	FFR	
Client	8 Hrs/day (40°C)	1 yr. (40°C)	<=3%	<10^-15
Enterprise	24 Hrs/day (40°C)	3 mo.(40°C)	<=3%	<10^-16

JEDEC Solution: Specify Endurance, Verify Spec via EVT

Endurance spec is max TB written to SSD over which device meets spec
Rigorous verification of Spec using EVT (Endurance Verification Test)
JEDEC supplies the workload. Data continuously read and verified.
SSD must meet<3% fail, UBER <1 in 10^-16
EVT requires high/low temp stressing
EVT represents lifetime worth of Stress Test, so can be trusted
Accelerated Test (High Temp bake) and Unaccelerated Room Temp Retention Test required
Manufacturer provides 'gauge' informing user of % of endurance life used up

SSD Challenges & Solutions: Goals & Best Practices

Education SNIA

Goals & Best Practices

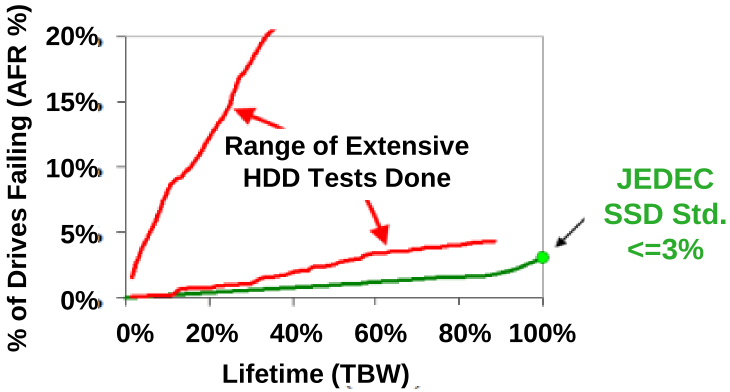

Afraid of SSD Adoption in your Enterprise?

Be aware of Tools & Best Practices… And you'll be OK !!

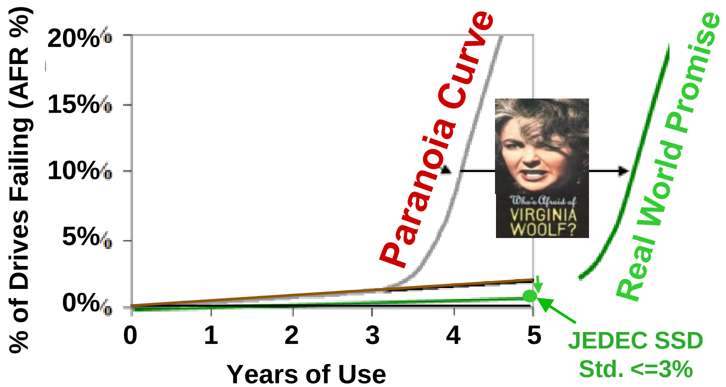

- All NAND will have finite Endurance Limits due to limitations imposed by:
  - Uncorrectable Bit Error Rates
  - Functional Failures
  - Data Retention Time
- Goal is to embody technologies to Improve Life (Years of Use)
  - Push Endurance Limit to the right beyond product life as required by SSD products
  - Push the defect rate down through Burn-Ins, Error Avoidance Algorithms and Practices. so the total <=3% defects and wear-outs issues combined
  - Target data errors to be < 1 in 10^16 for Enterprise SSDs for both TBW and Retentions specs.

Best Practices

- By leveraging Error Avoidance Algorithms, Verification Testing and Best Practices, so that total functional failure rate <=3% with defects and wear-outs issues combined
- In practice, endurance ratings are likely to be significantly higher than typical use, so data errors and failures will be even less.
- Capacity reduction can provide large increases in random performance and endurance.
- Select SSD based on confirmed EVT Ratings
- Use MLC within requirements of Endurance Limits

WW Enterprise SSD Market Opportunity

Education SNIA

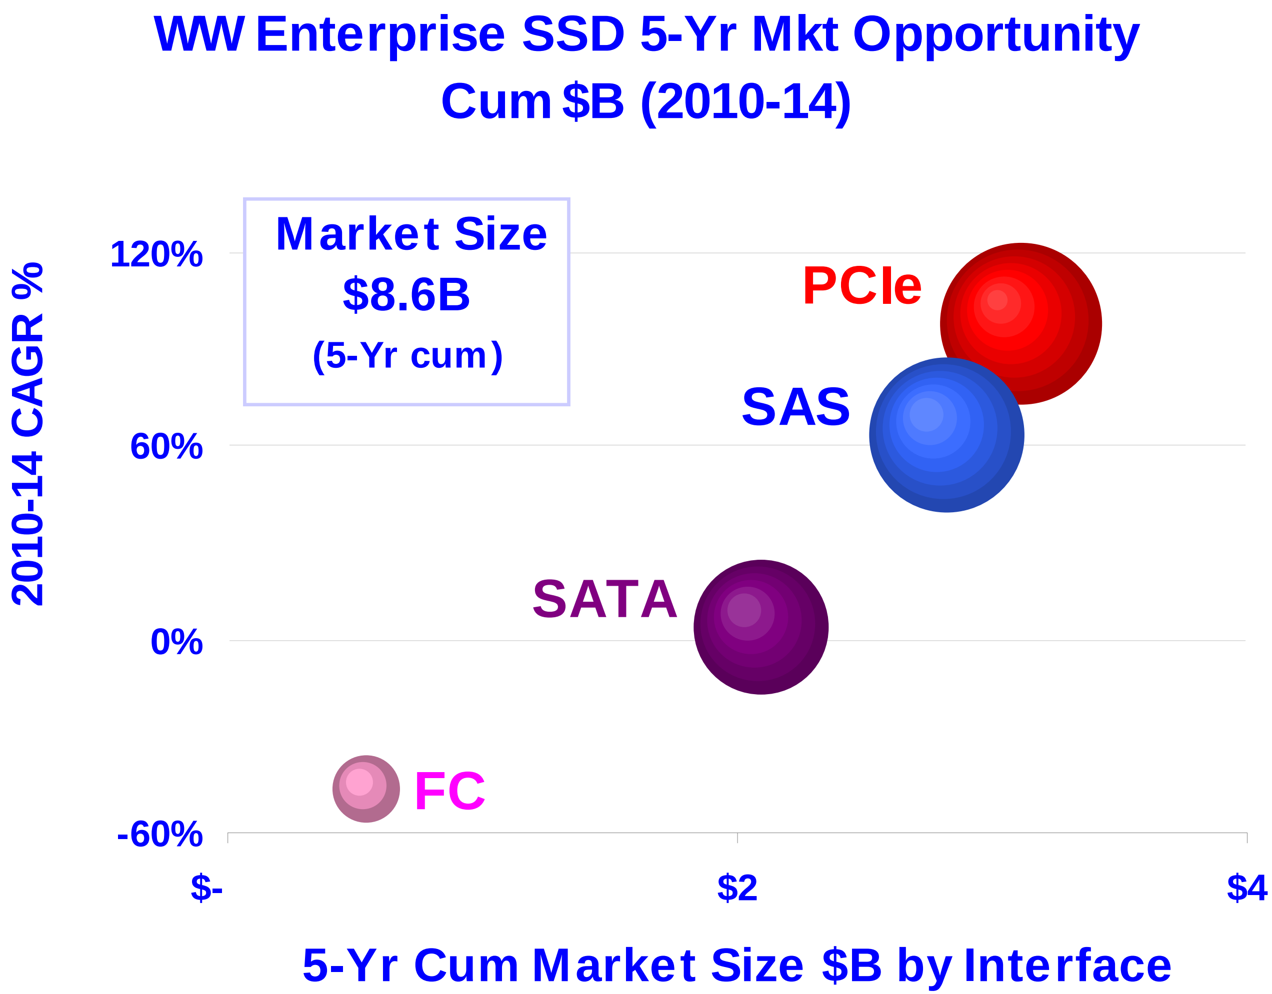

Source: IMEX Research SSD Industry Report ©2011

PCIe based SSD Storage

Education SNIA

PCIe based SSD Storage

- Target Market – Servers Storage
- SSD as backend storage to DRAM as the front end
- 36 PCIe Lanes Availability,
- 3/6 GB/s Performance (PCIe Gen2/3 x8),
- Low Latency in micro sec,
- Low Cost (via eliminating HBA cost)

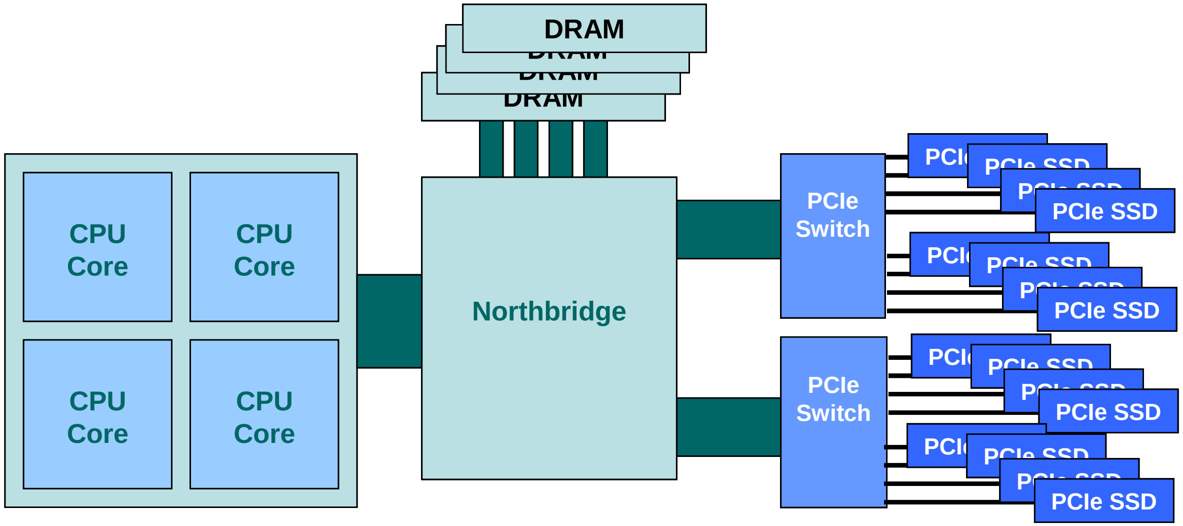

PCIe based SSD Storage

Usage for	Device	Storage Metric Target	Central Storage	Server Based
Caching		$/IOP, Latency *1	LBA Cache	LBA Cache
Performance	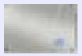	$/IOP/GB *2	Hot App Data	Hot App Data
Capacity	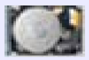	$/GB, Watts/GB *3	Cold/Lukewarm App Data	Lukewarm App Data

*1 – PCIe SSD performance enables new storage caching "IOPS Tier" as Application Managed Caching

*2 - PCIe SSDs of many flavors replace HDDs for High Performance Storage in some apps (e.g. Financial, DB etc)

*3 – HDDs best for Data at Rest as $/GB storage leader

PCIe SSD attributes of high IOPS, high Bandwidth, Low Latency and lower cost are a good match for Caching

Source: 2011

Hybrid SSD Storage

Education SNIA

Hybrid Storage – SAS or SATA SSD+HDD

- Target market – External Storage Systems
- Combines best features of SSDs - outstanding Read Performance (Latency, IOPs) and Throughput (MB/s) with extremely low cost of HDDs giving rise to a new class of storage - Hybrid Storage Devices
- SSD as Front End to HDD
- Controller emulates SSD as HDD
- Use of Adaptive Memory sends High IOPS requirements to SSD while capacity requiring Apps sent to HDD
- Simple Add on to SATA HDD Storage
- SAS 6Gb/sec announced by multi-vendors

Hybrid SSD Storage - Perf & TCO

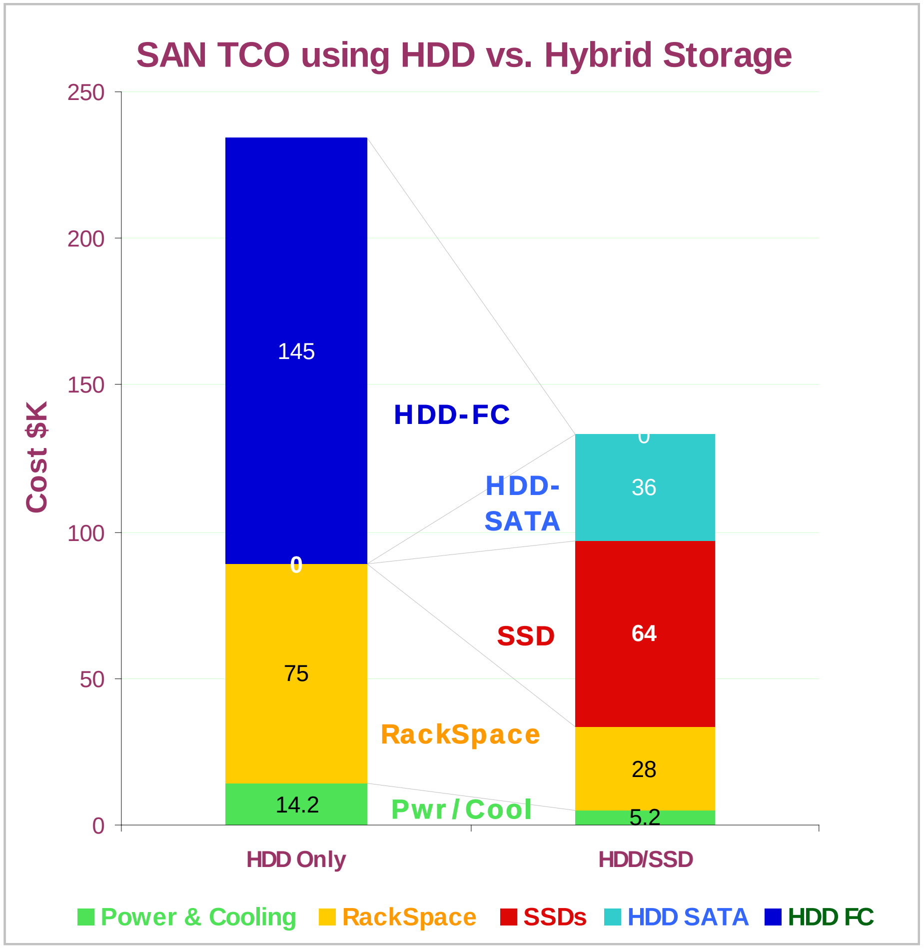

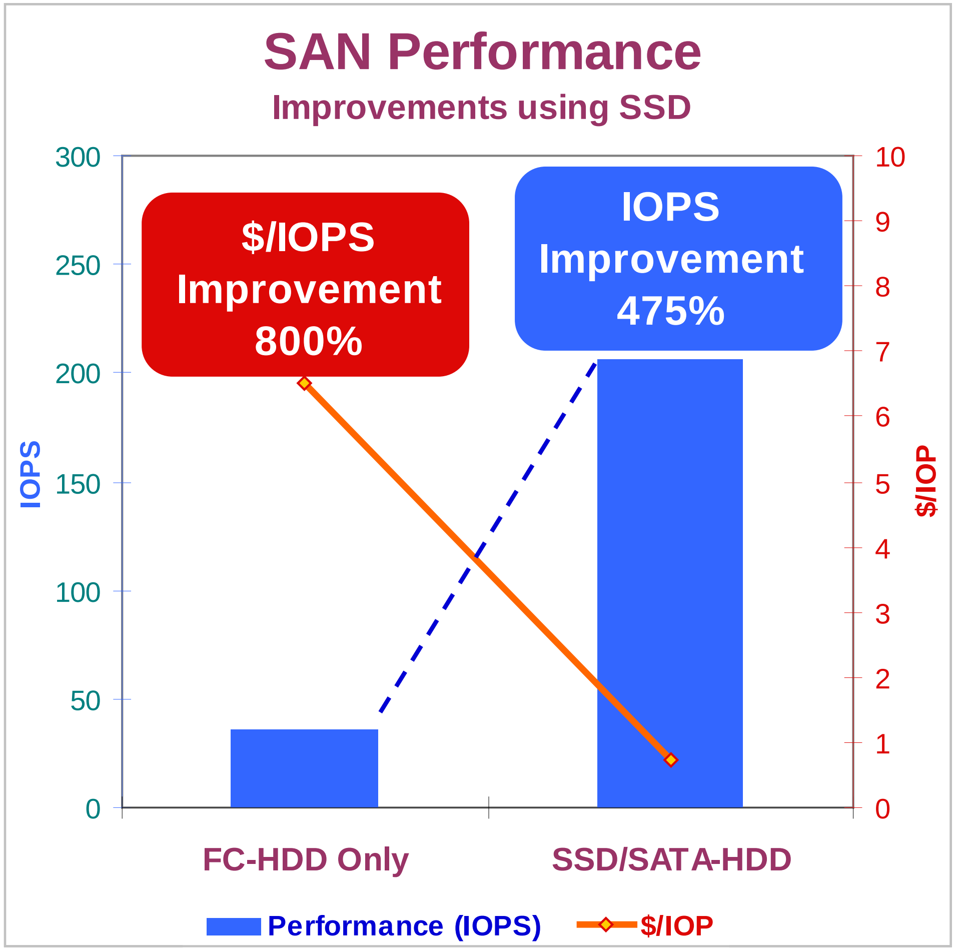

Source: IMEX Research SSD Industry Report ©2011

New Intelligent Controllers: SSD Storage Architecture

Education SNIA

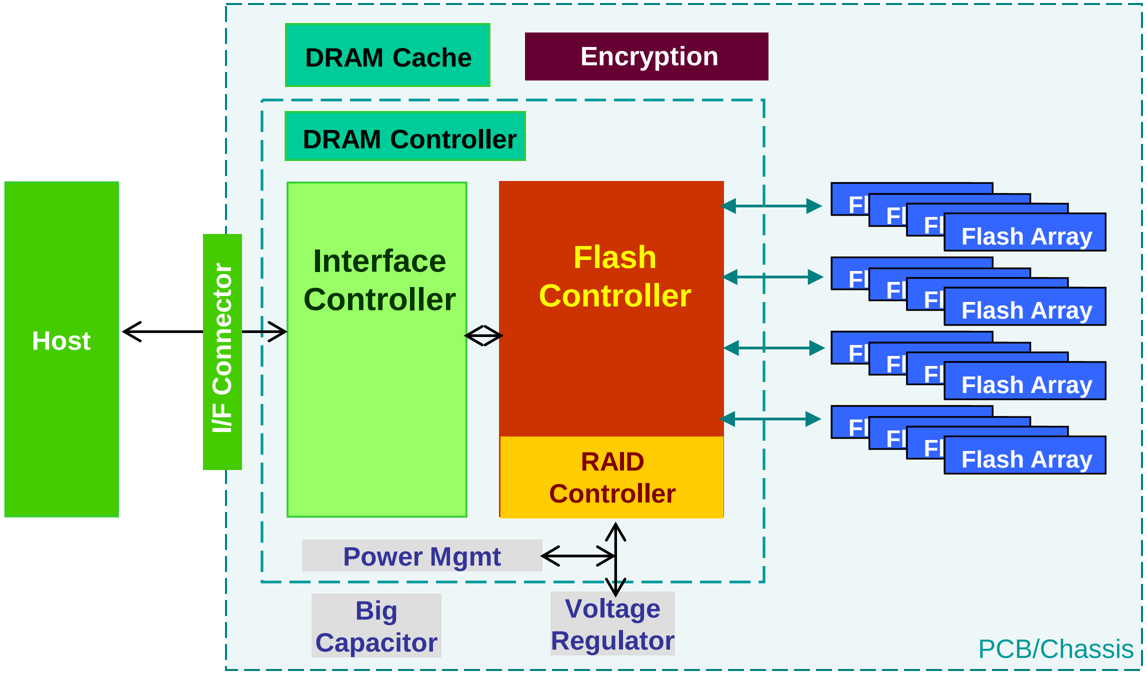

1	Interface Controller	Signaling Mgmt, Interpret WR/RD/Status Commands, Native Command Queuing, Move Data <-> Host
2	Flash Controller	Signaling Mgmt, Format, Interpret WR/RD/Status Commands for Flash Arrays, Move Data. Defect Mapping/Bad Block Mgmt, Wear Leveling, Physical<>Logical Translations, ECC…
3	RAID Controller	RAID Type & RD/WR/Parity Manipulation
4	Channels	Multiple Channel to Increase Speed between NAND Flash Arrays & Flash Controller
5	DRAM	Increase Performance using fast DRAM Cache Buffer
6	Power Failure	Power Failure Protection using Big Capacitor
7	Power Mgmt	Power/Performance Balancing, Sleep Mode Mgmt
8	Encryption	Security Schemes Implementation & Manipulation

New Intelligent Controllers: Managing NAND Media in NextGen SSDs

Education SNIA

Leveraging Long History of managing HDD's imperfect media & high error rates

- Characterizing the quality & capabilities of media
- Allocating data based on quality of media

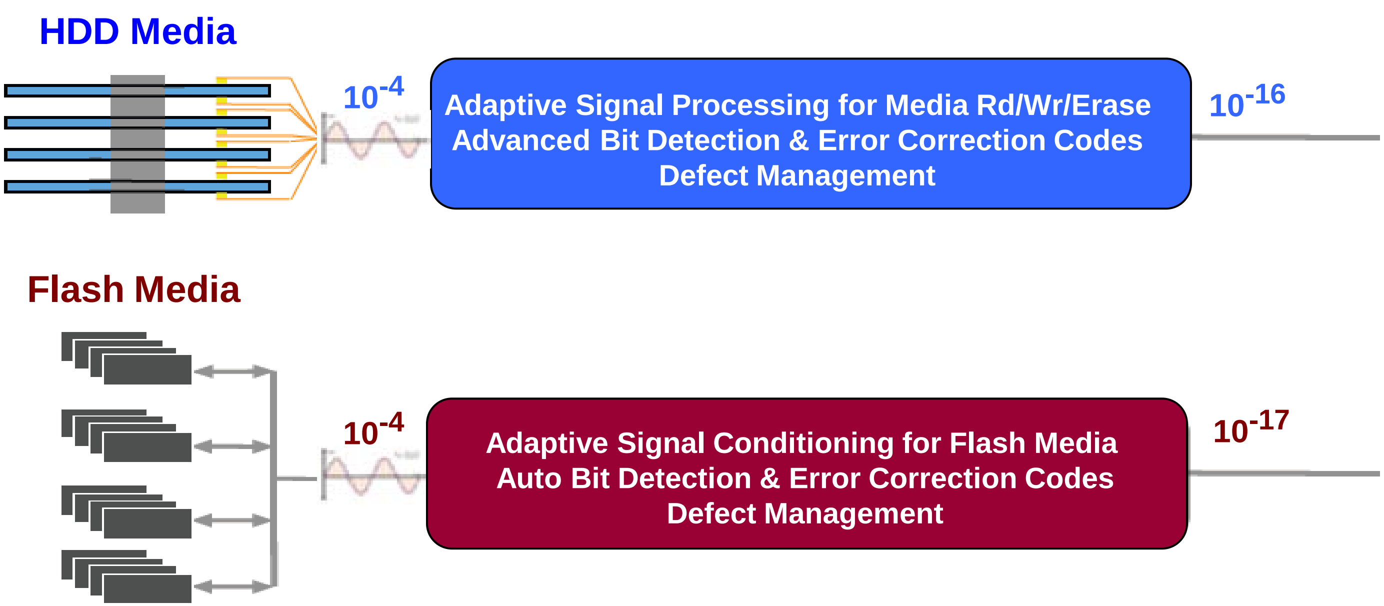

Leveraging Long History of managing HDD's imperfect media & high error rates

- Endurance for Long Life Cycle
- Reliability through RAID of Flash Elements
- Adaptive Digital Signal Processing Technology
- Dynamically adjust Read/Write characteristics of each chip
- Tune adjustments over life of media
- ECCs - PRML Deploying Enhanced Error Correction Codes

Source: IMEX Research SSD Industry Report ©2011

New Intelligent Controllers: Managing Enterprise Requirements

Education SNIA

Meet Enterprise Requirements

- **Always-On 24x7 Reliability** and performance supersede cost
- **Fast I/O Performance** required by business-critical applications and
- **5-Yr. Life Cycle Endurance** required by mission-critical applications in the enterprise.
- **Use State-of-the-Art** new sophisticated controllers and firmware technologies to run mission critical applications in the enterprise, using
  - Robust ECC, Internal RAID, Wear Leveling (To reduce hot spots), Spare Capacity, Write Amplification, Avoidance, Garbage Collection Efficiency, Wear Out Prediction Management etc.

Source: SandForce

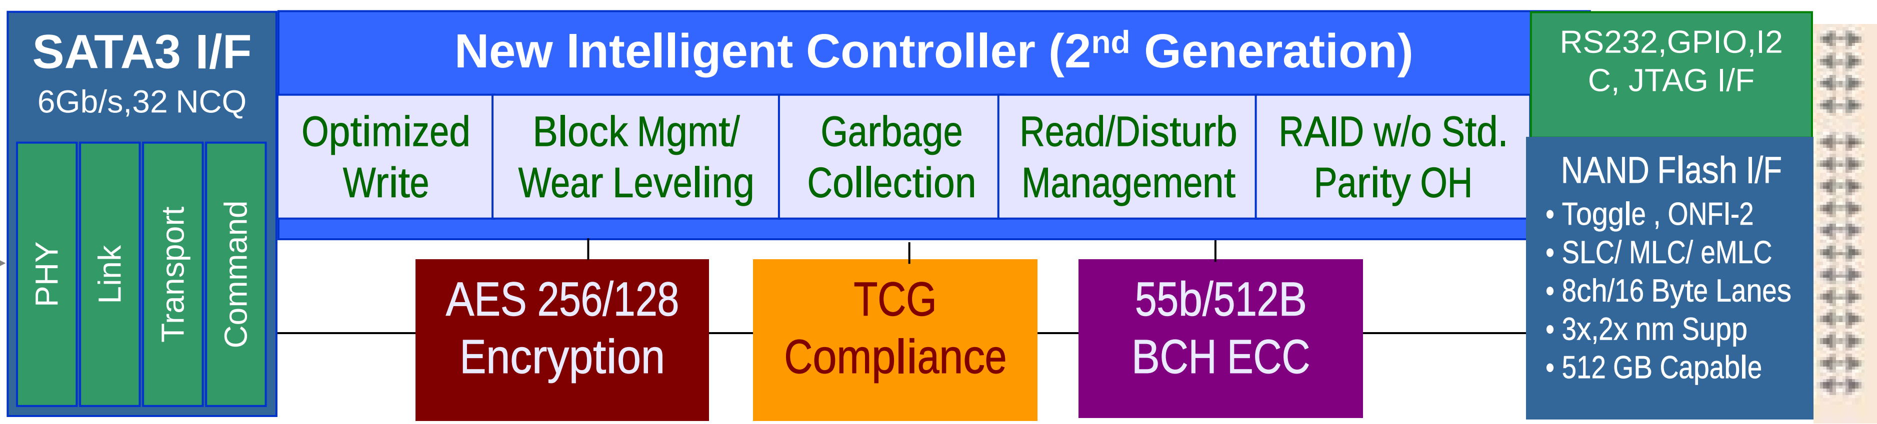

Managing Endurance

To overcome NAND's earlier endurance shortfalls due to limitation in write/erase cycles/block, intelligent controllers manage NAND SSDs using

- **ECC Techniques** – Correct and guard against bit failures, same as in HDDs
- **Wear Leveling Algorithms** – Writing data to evenly distributes it over all available cells to avoids a block of cells being overused and cause failures.
- **Over-provisioning Capacity** – Extra spare raw blocks are designed-in as headroom and included to replace those blocks that get overused or go bad. Additionally provide enough room for wear-leveling algorithms to enhance reliability of the device over its life-cycle.
- Typical SSD device's specified GB device will actually contain 20-25% extra raw capacity to meet these criterions.

Managing Reliability

Multiple techniques are being used to improve the reliability, such as:

In-Flight

Corruption upstream disk controllers, Corruption in SSD controller itself

Flush at power loss using large cap elements

At-Rest

ECC

Scanning & scrubbing

Redundancy

Meta-data

Error correcting memory

Data integrity field

New Intelligent Controllers
- Managing Performance in Next Gen SSDs

Education SNIA

Managing Performance / Key Metrics Impact

- Factors Impact Performance
- **Hardware -** CPU, Interface, Chipset ...
- **System SW** - OS, App, Drivers, Caches, SSD specific TRIM, Purge, …
- **Device -** Flash Generation, Parallelism, Caching Strategy, Wear-Leveling, Garbage Collection, Warranty Strategy…
- **Write History** - TBW, spares…)
- **Workload -** Random, Sequential, R/W Mix, Queues, Threads…
- **Pre-Conditioning -** Random, Sequential, Amount …
- **Performance -** Short "Burst" First On Board (FOB)

  - Steady State post xPE Cycles

Using interleaved memory banks, caching and other techniques being designed in modern controllers, the performance of MLC SSDs today started to match and even outshines

performance offered by some SLC SSDs

AutoSmart Storage-Tiering SW: Storage Mapping

Automated Storage Tiering Principles

- Continuously monitor and analyze data access on the tiers
- Automatically elevate hot data to "Hot Tiers" and demote cool data/volumes to "Lower Tiers. Allocate and relocate volumes on each tier based on use
- Reduces OPEX vs. managing SANs manually. All major Computer System manufacturers adopted it such as FAST, Easy Tier, Data Progression, Adaptive Optimization, Dynamic Tiering, Smart Pools...

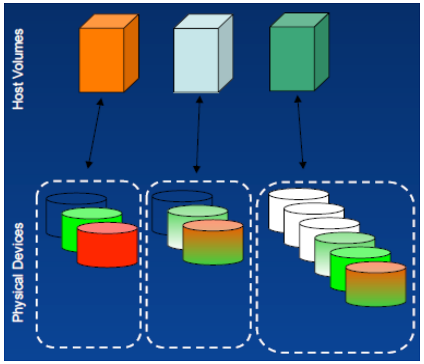

Traditional Disk Mapping

- Volumes have different characteristics. Applications need to place them on correct tiers of storage based on usage

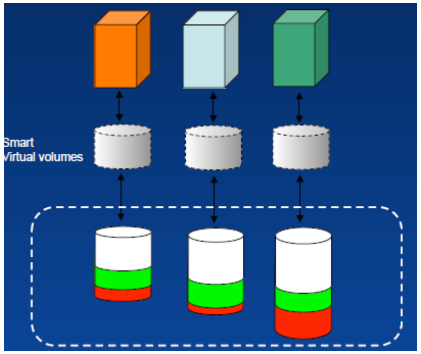

Smart Storage Mapping

- All volumes appear to be "logically" homogenous to apps. But <u>data is placed at the right tier</u> of storage based on its usage through smart data placement and migration

Source: IBM & IMEX Research SSD Industry Report ©2011

AutoSmart Storage-Tiering SW: Workload I/O Monitoring/Smart Migrations

Education SNIA

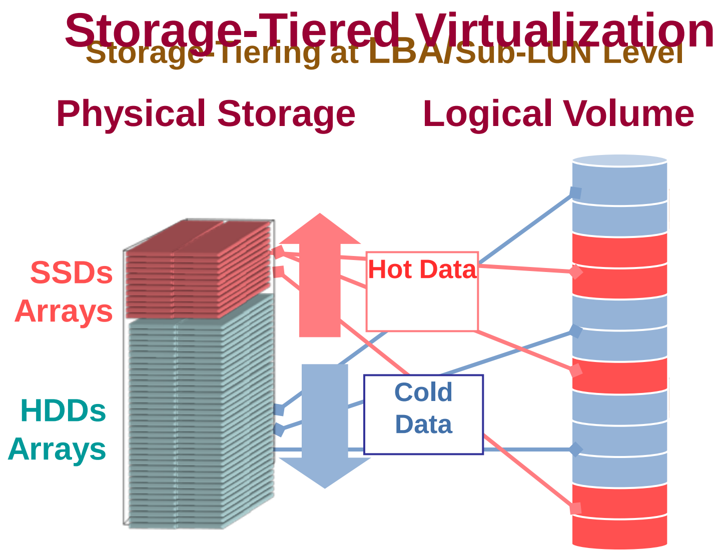

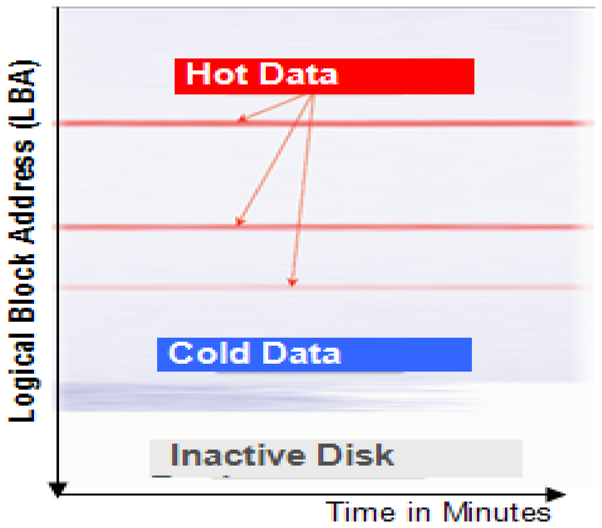

LBA Monitoring and Tiered Placement

- Every workload has unique I/O access signature
- Historical performance data for a LUN can identify performance skews & hot data regions by LBAs
- Using Smart Tiering identify hot LBA regions and non-disruptively migrate hot data from HDD to SSDs.
- Typically 4-8% of data becomes a candidate and when migrated to SSDs can provide response time reduction of ~65% at peak loads.

Source: IBM & IMEX Research SSD Industry Report 2011 ©IMEX 2010-11

AutoSmart Storage-Tiering SW: Enhancing Database Throughput

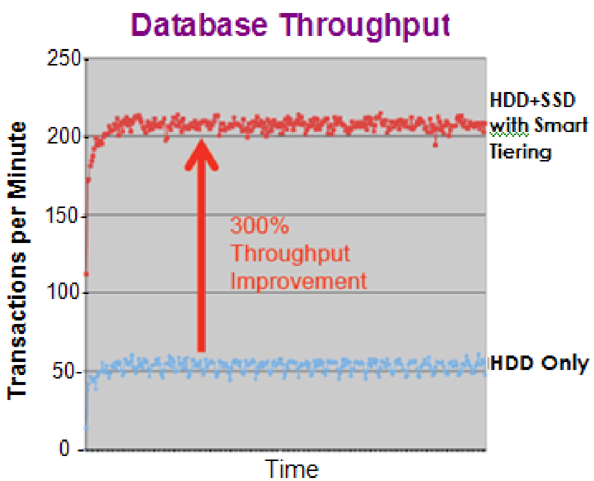

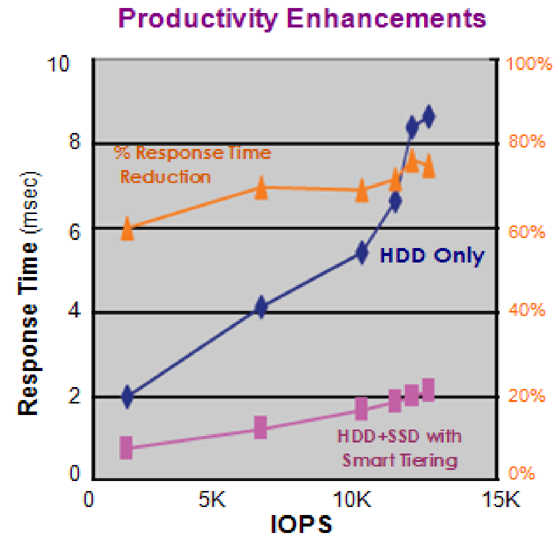

- **DB Throughput Optimization**
  - Every workload has unique I/O access signature and historical behavior
  - identify hot "database objects" and smartly placed in the right tier.
  - Scalable Throughput Improvement - 300%
  - Substantial IO Bound Transaction Response time Improvement - 45%-75%

- **Productivity (Response Time) Improvement**
  - Using automated reallocation of hot spot data (typically 5-10% of total data) to SSDs, performance improvements is achieved
  - Response time reduction of around 70+% or
  - Through put (IOPS) increase of 200% for any I/O intensive loads experienced by Time-Perishable Online Transactions like:♦ Airlines Reservations, ♦ Wall Street ♦ Investment Banking Stock Transactions ♦ Financial Institutions Hedge Funds etc. as well as Low Latency seeking HPC Clustered Systems etc.

Source: IBM & IMEX Research SSD Industry Report 2011 ©IMEX 2010-11

Applications Best Suited for SSDs

Education SNIA

Applications most benefitting from SSDs Use

Database/OLTP	43%
E-Mail/Collabor.	32%
HPC	31%
BI/DW	30%
ERP/SCM/CRM	25%
Web 2.0	23%
Office Apps	20%

Source: IMEX Research SSD Industry Report ©2011

Apps and impact from SSD Usage

- **Databases**
  - Databases have key elements of commit files – logs, redo, undo, tempDB
- **Structured data**
  - Structured data access is an excellent fit for SSD
  - Exception–large, growing table spaces
- **Unstructured data**
  - Unstructured data access is a poor fit for SSD
  - Exception – small, non-growing, tagged files
- **OS images**
  - boot-from-flash, page-to-DRAM

Typical Cases - Impact on Applications

- **Financial Credit Card/ATM Transactions**
  - **Improvements:** Batch Window 22%, App Response Time 50%, App I/O Rate 50%
- **Messaging Applications**
  - Cost Savings: 200+ FC HDDS into only 16 SSDs

Apps Best Suited for SSDs:
OLTP to Improve Query Response Time

Education SNIA

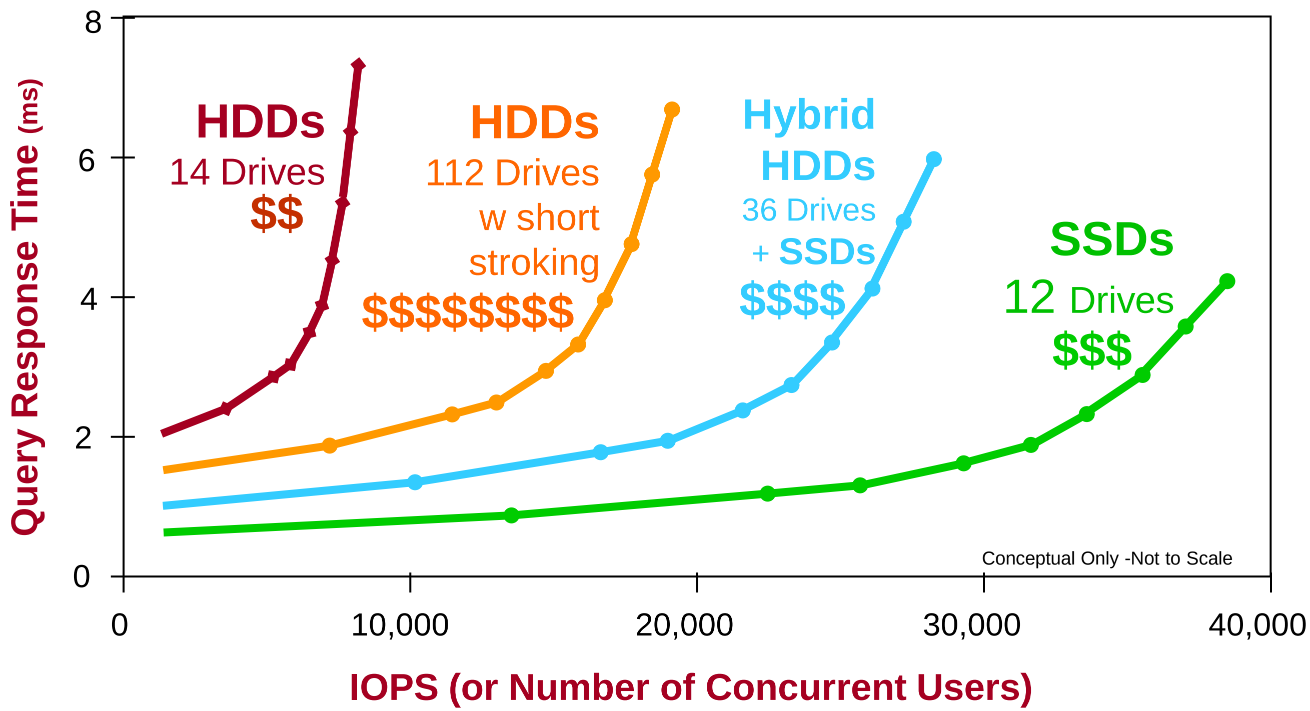

- **Improving Query Response Time**
  - Cost effective way to improve Query response time for a given number of users or servicing an increased number of users at a given response time is best served with use of SSDs or Hybrid (SSD + HDDs) approach, particularly for Database and Online Transaction Applications

Source: IMEX Research SSD Industry Report ©2011

Apps Best Suited for SSDs:
DB in Memory for Data Warehouse/BI

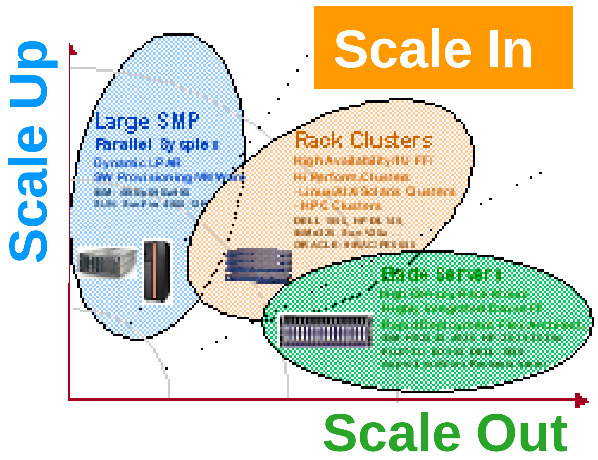

Large DB Size Growth by Market Segment

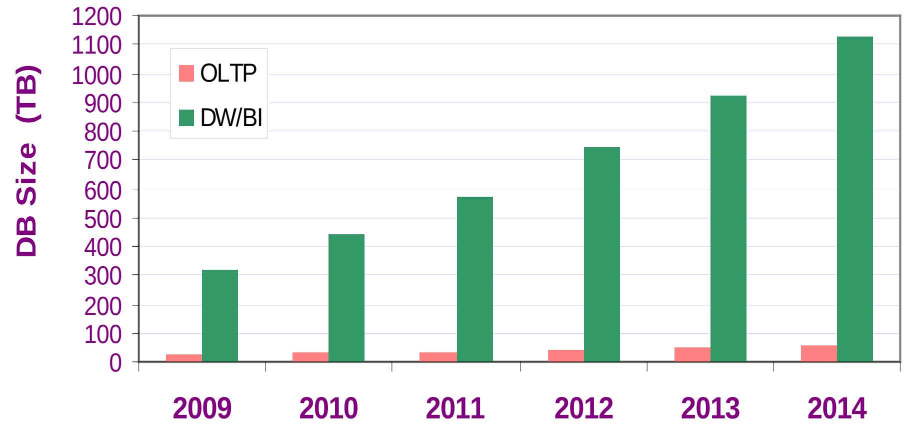

Storage Usage vs DB Capacity

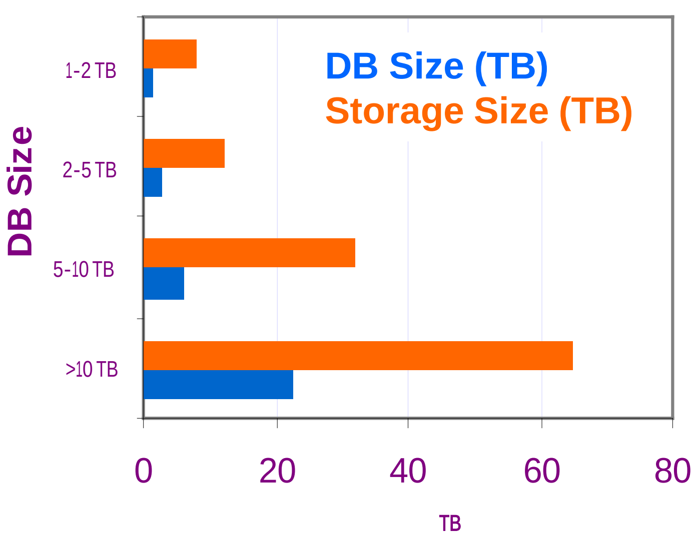

Corporate Data Usage
I/O Access Frequency vs. Percent of Data

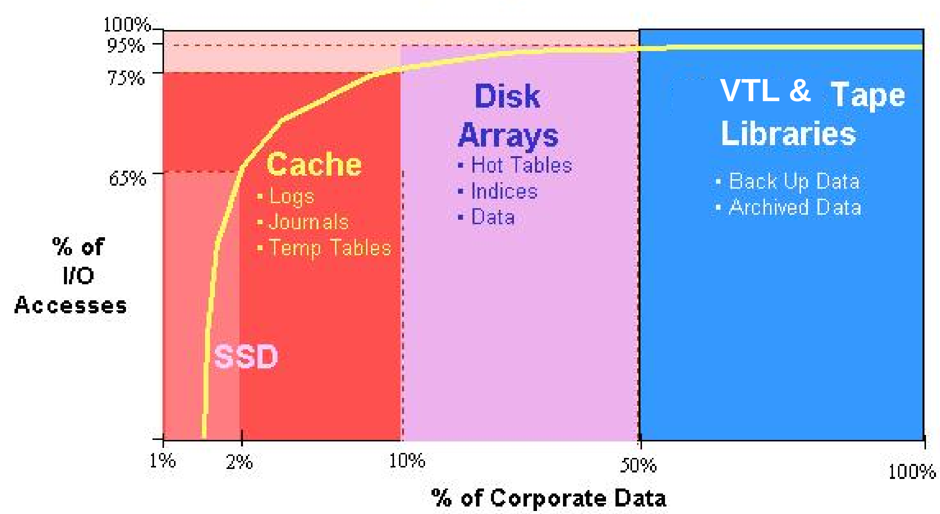

Data Source: IMEX Research
Cloud Infrastructure Report ©2009-11

Apps Best Suited for SSDs: Collaboration using Virtual Desktops

Education SNIA

Mitigating Boot Storms

- Boot Storm created by simultaneous Logins by users at start of office day
- Over provisioning SAN capacity just for short morning burst expensive, while sitting almost idle rest of the day
- Three potential solutions with pros and cons include:
- (1) DAS Storage, (2) Virtual Desktop Images on SSD (3) SS Cache to accelerate SAN,

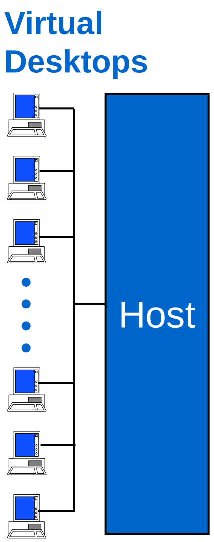

Solution	Pros	Con
New DAS Storage	- Popular with Desktop SW Vendors - Lowered Cost	- Additional Cost for Dedicated Storage - Wasted existing SAN Storage - Data Protection & Mgmt Challenges
SSD	- SSD Ideal for read intensive app - Instant-On/Immediate Boot Response – Images Stored with High Efficiency	- Using most expensive storage - High Speed Needed just for an hour - Not simple shoe-in w existing storage
SAN w SSD Accelerator	- Possibly best way to solve problem - Small SSD optimized for image store - No change to existing Data Protection	- Not feasible without existing SAN - SSD in SAN Integration still a challenge

Providing a perfect balance of access and storage is achieved through Integrating SATA HDDs with SSDs and using Automatic Tiering Solutions

Source: IMEX Research SSD Industry Report ©2011

Apps Best Suited for SSDs:
HPC/Web 2.0

Education SNIA

Smart Mobile Devices	Commercial Visualization	Bioinformatics & Diagnostics	Decision Support Bus. Intelligence	Entertainment-VoD / U-Tube


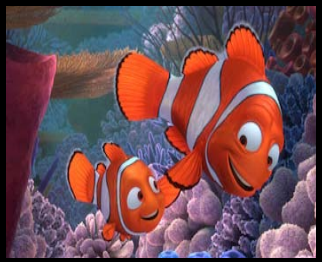
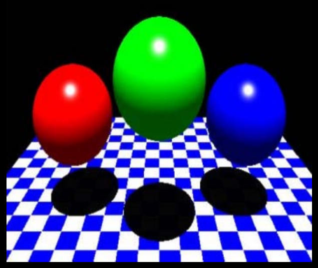
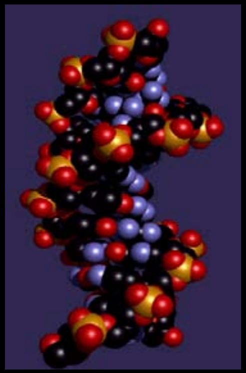
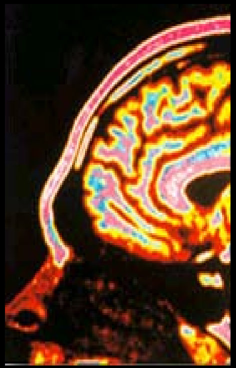


Data: IMEX Research & Panasas

Instant On Boot Ups Rugged, Low Power 1GB/s, __ms	Rendering (Texture & Polygons) Very Read Intensive, Small Block I/O 10 GB/s, __ms	Data Warehousing Random IO, High OLTPM 1GB/s, __ms	Most Accessed Videos Very Read Intensive 4 GB/s, __ms

Key Takeaways

- **Optimize Infrastructure to meet needs of Applications/SLA**
- **Solid State Storage creating a paradigm shift in Storage Industry**
  - Leverage the opportunity to optimize your computing infrastructure with SSD adoption after making a due diligence in selection of vendors/products, industry testing and interoperability
- **Enterprise SSD Market Segments: PCIe vs. SAS/SATA**
  - 5-Year cum Market $8.6B Segments by Revenues: 36% PCIe, 33% SAS, 24% SATA, 6% FC based SSDs
- **Understand Drivers and Challenges of SSDs for Enterprise Use**
- **Intelligent Controllers key to adoption & success of SSDs**
  - Mitigate Endurance, Wear-Out, Life issues
- **Optimize Transactions for Query Response Time vs. # of Users**
  - Improving Query Response time for a given number of users (IOPs) or Serving more users (IOPS) for a given query response time
- **Select Automated Storage Tiering Software**

Q&A / Feedback

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

Many thanks to the following individuals for their industry vision in this tutorial.

N. Mielke

Jim Elliott

John Scaramuzzo

Michael Raam

Manouch Moshayedi

Dean Klein

Joseph White

and

SNIA Education Committee