Education

Solid State Storage:
Key to NextGen Enterprise & Cloud Storage

Anil Vasudeva, President & Chief Analyst, IMEX Research

Author: Anil Vasudeva, IMEX Research.com

SNIA Legal Notice

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

Source:©IMEX Research Industry Report 2009-12

Abstract

Solid State Storage: Key to NextGen Enterprise & Cloud Computing

Abstract

Computer architects dream of storage devices which can provide very high IOPs at minimal cost (IOPS/$/GB) using infinite cheap storage and instant access (low latency) for their applications/ workloads. Enterprise-Ready SSDs ©IMEX have started to fulfill that promise and being available as SATA and PCIe based Hybrid Storage products. But only the recent advent of advanced controllers and firmware which has allowed transparent mitigation of earlier issues related to reliability, endurance, data retention, performance, ease of management and quick integration using exiting storage interfaces. But the real killer software for their success in the enterprise have been the Automated Storage Tiering tools activated by monitoring the workload I/O access signatures and behavior over time and the ensuing smart migration of hot data to SSDs non-disruptively that has resulted in obtaining over 475% improvement in IOPS and 80% improvement in response time at peak loads

The presentation delineates the recently emerged technologies, storage characteristics (performance, cost, reliability and endurance etc.), and the applications that benefit the most from the use of SSDs in enterprise storage systems, workloads optimization using new generation of controllers and automated smart-tiering by specific vertical-industries as well as the economics of SSDs usage using real market data.

Learning Objectives

The session provides a clear illustrative the state-of-the-technology, storage characteristics (performance, cost, reliability and endurance etc.), overview of industry advances, specific applications that benefit the most from SSDs use, system implementation in enterprise storage systems allowing them to plan, implement and achieve stated benefits expected from using SSDs as a tiered storage specifically in OLTP/Database, Business Intelligence applications and cluster-based HPC workloads. The presentation illustrates real life case studies in SANs showing how optimally selected hierarchical storage systems encompassing SSDs and HDDs can achieve 65% lower TCO, 475% higher IOPS, 165% lower footprint while achieving a whopping 800% in $/IOPs in SAN and other tired storage systems, under different scenarios.

Source:©IMEX Research Industry Report 2009-12

Agenda – Solid State Storage

1. NextGen Data Center and Cloud Computing Infrastructure
2. Solid State Enabling New Systems Architecture
3. Improving Transaction Query Response Time and IOPS
4. Workload Characterization
5. Applications best utilizing Solid State Storage
6. New Intelligent Controllers: Heart of SS Storage Systems
7. Data Forensics and Tiered Mapping
8. Key Takeaways

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

IT Industry's Journey - Roadmap

Analytics – BI

Predictive Analytics - Unstructured Data

From Dashboards Visualization to Prediction Engines using Big Data.

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

Std. IT Infrastructure- Volume Economics HW/Syst SW

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

Source:: IMEX Research - Cloud Infrastructure Report ©2009-12

Data Centers & Cloud Infrastructure

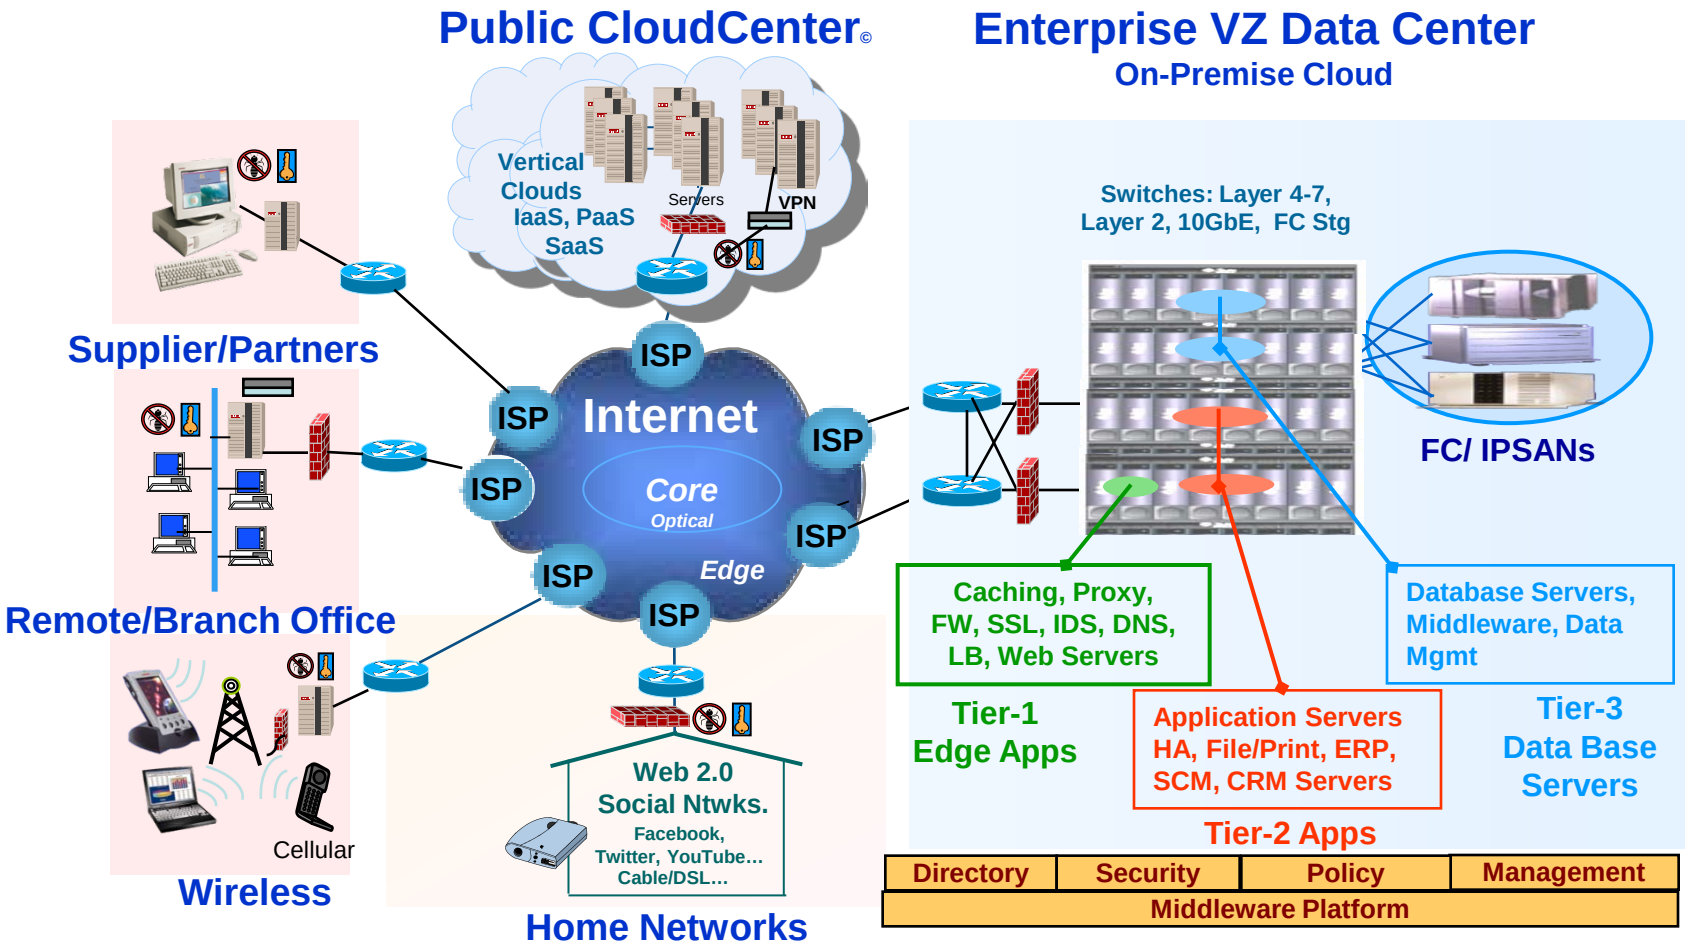

Request for data from a remote client to an enterprise data center crosses a myriad of systems and devices. Key is identifying bottlenecks & improving performance

Source:: IMEX Research - Cloud Infrastructure Report ©2009-11
Source:©IMEX Research Industry Report 2009-12

Market Segments by Apps/Workloads

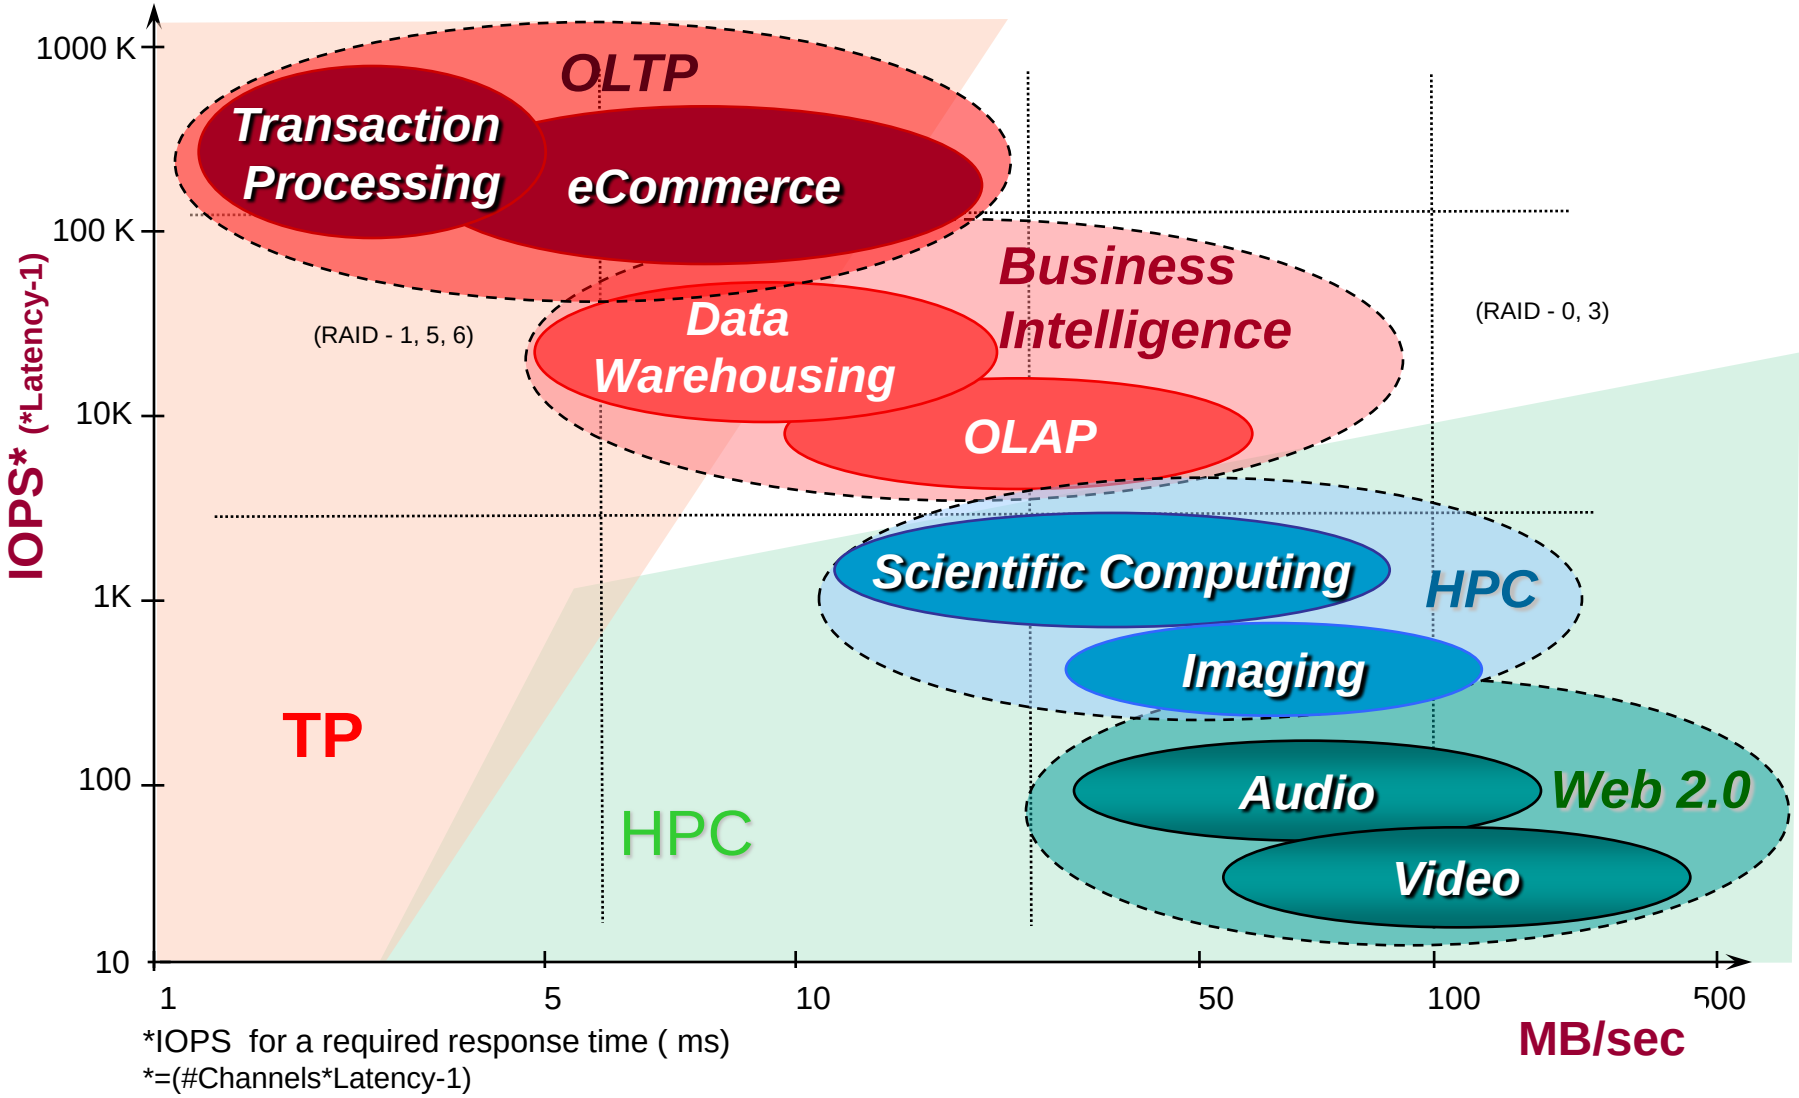

Source:: IMEX Research - Cloud Infrastructure Report ©2009-11
Source:©IMEX Research Industry Report 2009-12

Big Data Analytics – Next Frontier in IT

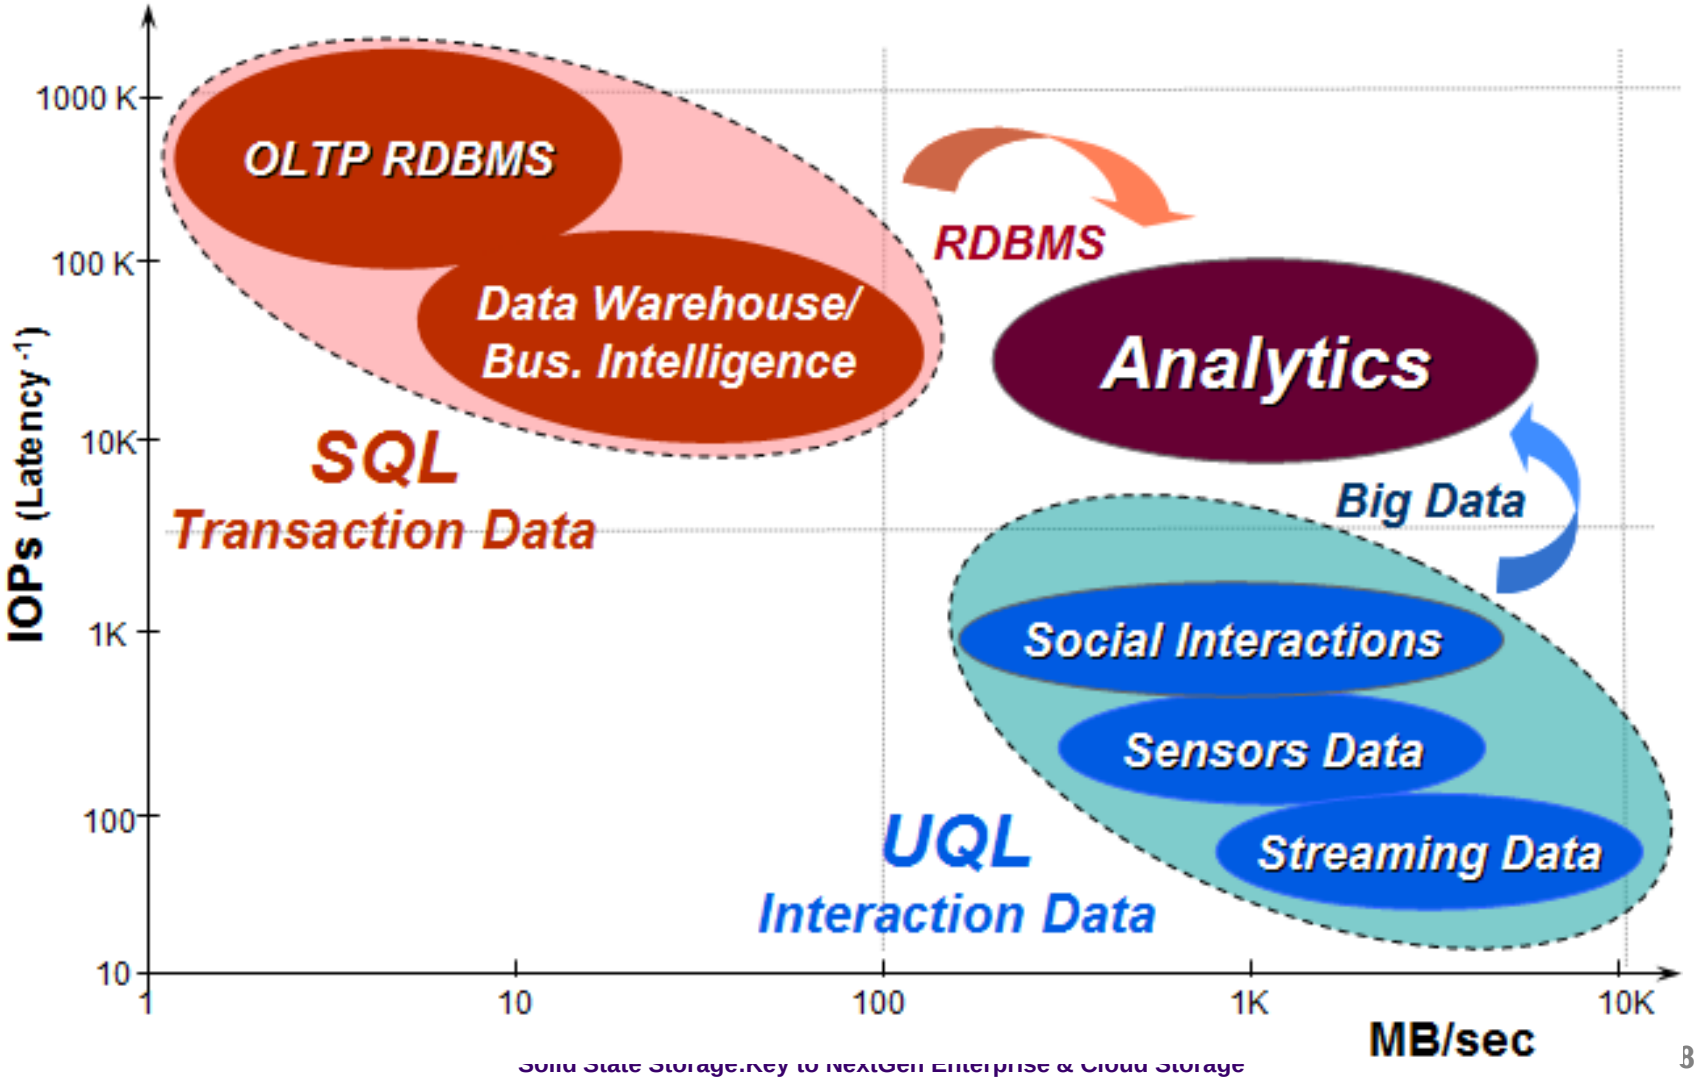

Source:©IMEX Research Industry Report 2009-12

Virtualized Cloud Infrastructure

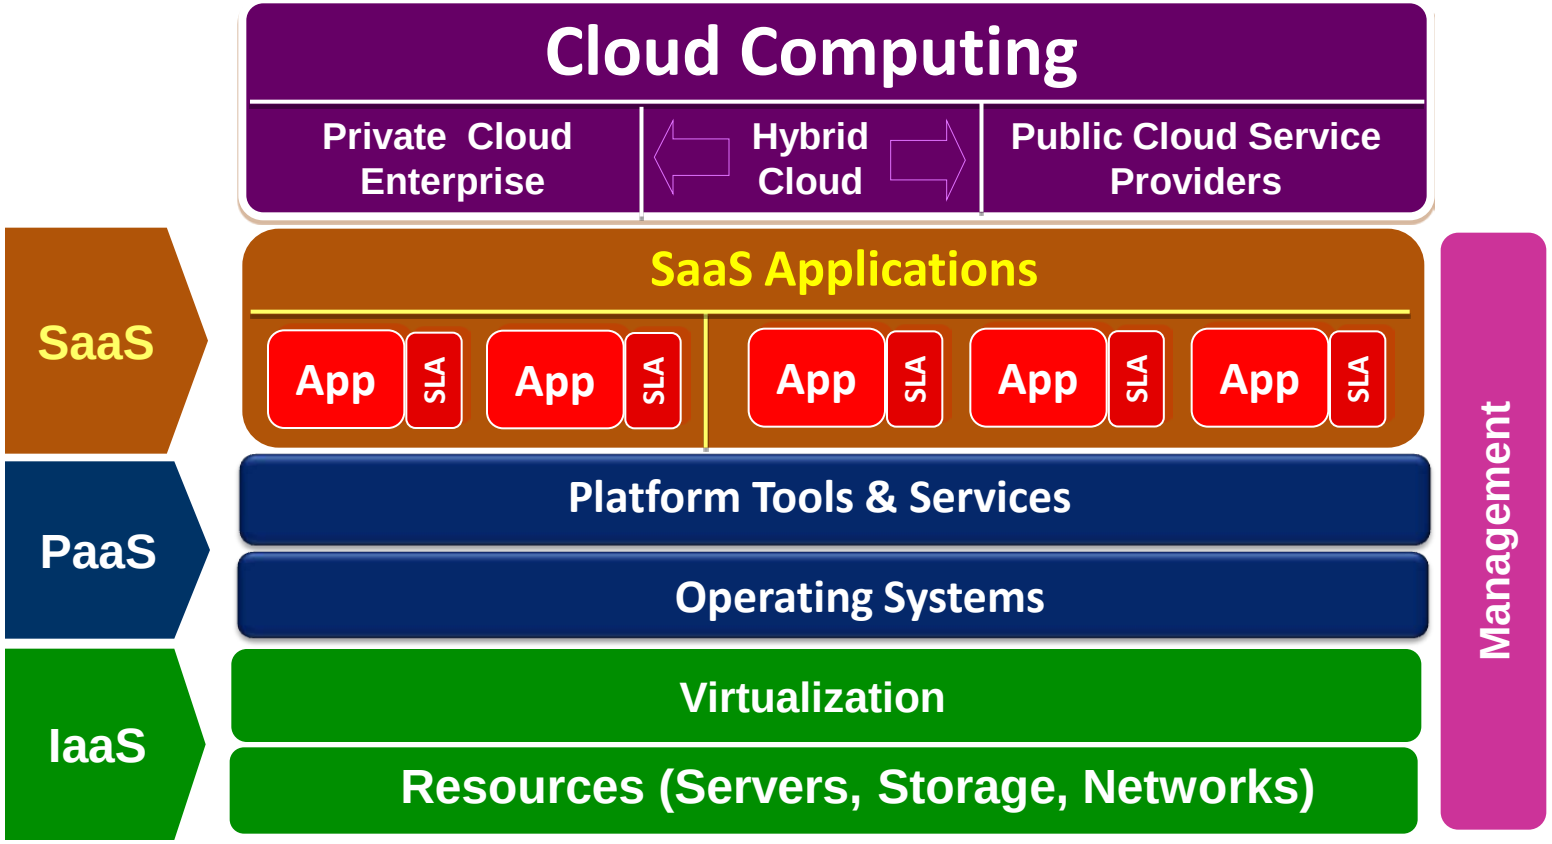

Application's SLA dictates the Resources Required to meet specific requirements of Availability, Performance, Cost, Security, Manageability etc.

Source:©IMEX Research Industry Report 2009-12

Anatomy of Data Access

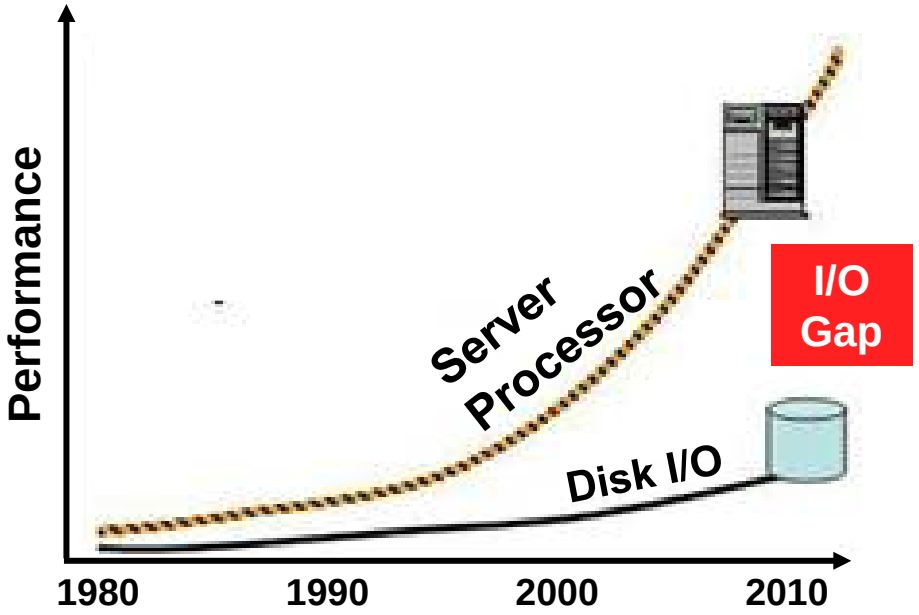

For the time it takes to do Each Disk Operation:

- Millions of CPU Operations can be done
- Hundreds of Thousands of Memory Operations can be accomplished

Anatomy of Data Access

Time taken by CPU, Memory, Network, Disk for a typical I/O Operation during a Data Access

L1 cache reference	0.5 ns
Branch mispredict	5 ns
L2 cache reference	7 ns
Mutex lock/unlock	25 ns
Main memory reference	100 ns
Compress 1K bytes with Zippy	3,000 ns
Send 2K bytes over 1 Gbps network	20,000 ns
Read 1 MB sequentially from memory	250,000 ns
Round trip within same datacenter	500,000 ns
Disk seek	10,000,000 ns
Read 1 MB sequentially from disk	20,000,000 ns
Send packet CA->Netherlands->CA	150,000,000 ns

A 7.2K/15k rpm HDD can only do 100/200 IOPS*

*** (typical/worst cases considered)**

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

Data Center Performance Bottlenecks

Education SNIA

Applications
Excessive Locking
Data Contention
I/O Delays/Errors

Network I/O
Network Congestion
Dropped packets
Data Retransmissions
Timeouts
Component Failures

Storage I/O Connect
Lack of Bandwidth
Overloaded PCIe Connect
Storage Device Contention

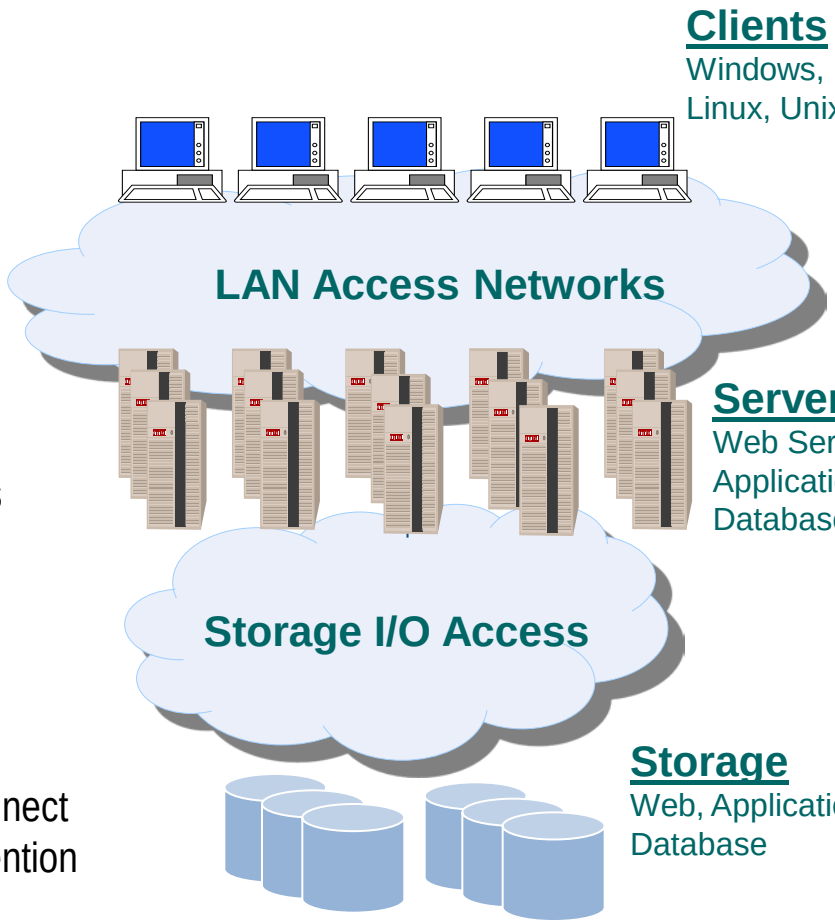

User Bottlenecks
Connectivity Timeouts,
Workload Surges

Server Bottlenecks
Lack of Srvr Power
IO Wait & Queuing CPU
Overhead I/O Timeouts

Device Bottlenecks
Device I/O Hotspots
Cache Flush
Lack of Storage Capacity

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

Automated Storage Tiering - Next Frontier in Storage Efficiency

Education SNIA

Data Protection	Storage Efficiency
Back Up/Archive/DR	Auto Tiered Storage
RAID – 0,1,5,6,10	Storage Virtualization
Virtual Tape	Thin Provisioning
Replication	Deduplication
	MAID

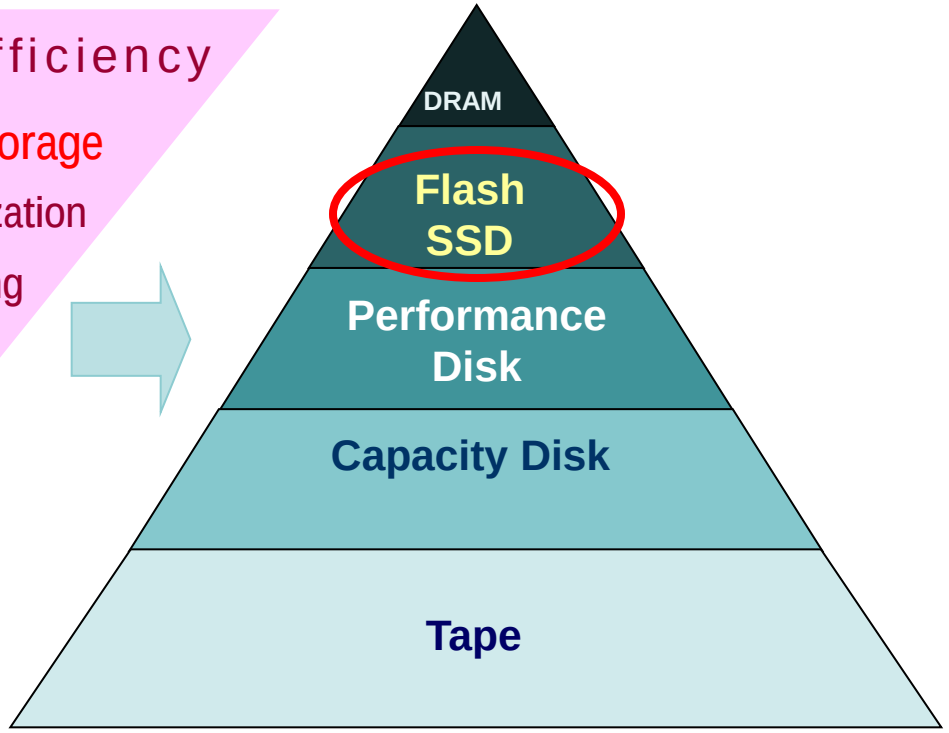

Auto Tiering System using SSDs

Data Class (Tiers 0,1,2,3)
Storage Media Type (Flash/Disk/Tape)
Policy Engines (Workload Mgmt)
Transparent Migration (Data Placement)
File Virtualization (Uninterrupted App.Opns.in Migration)

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

SSD Filling Price/Perf Gaps in Storage

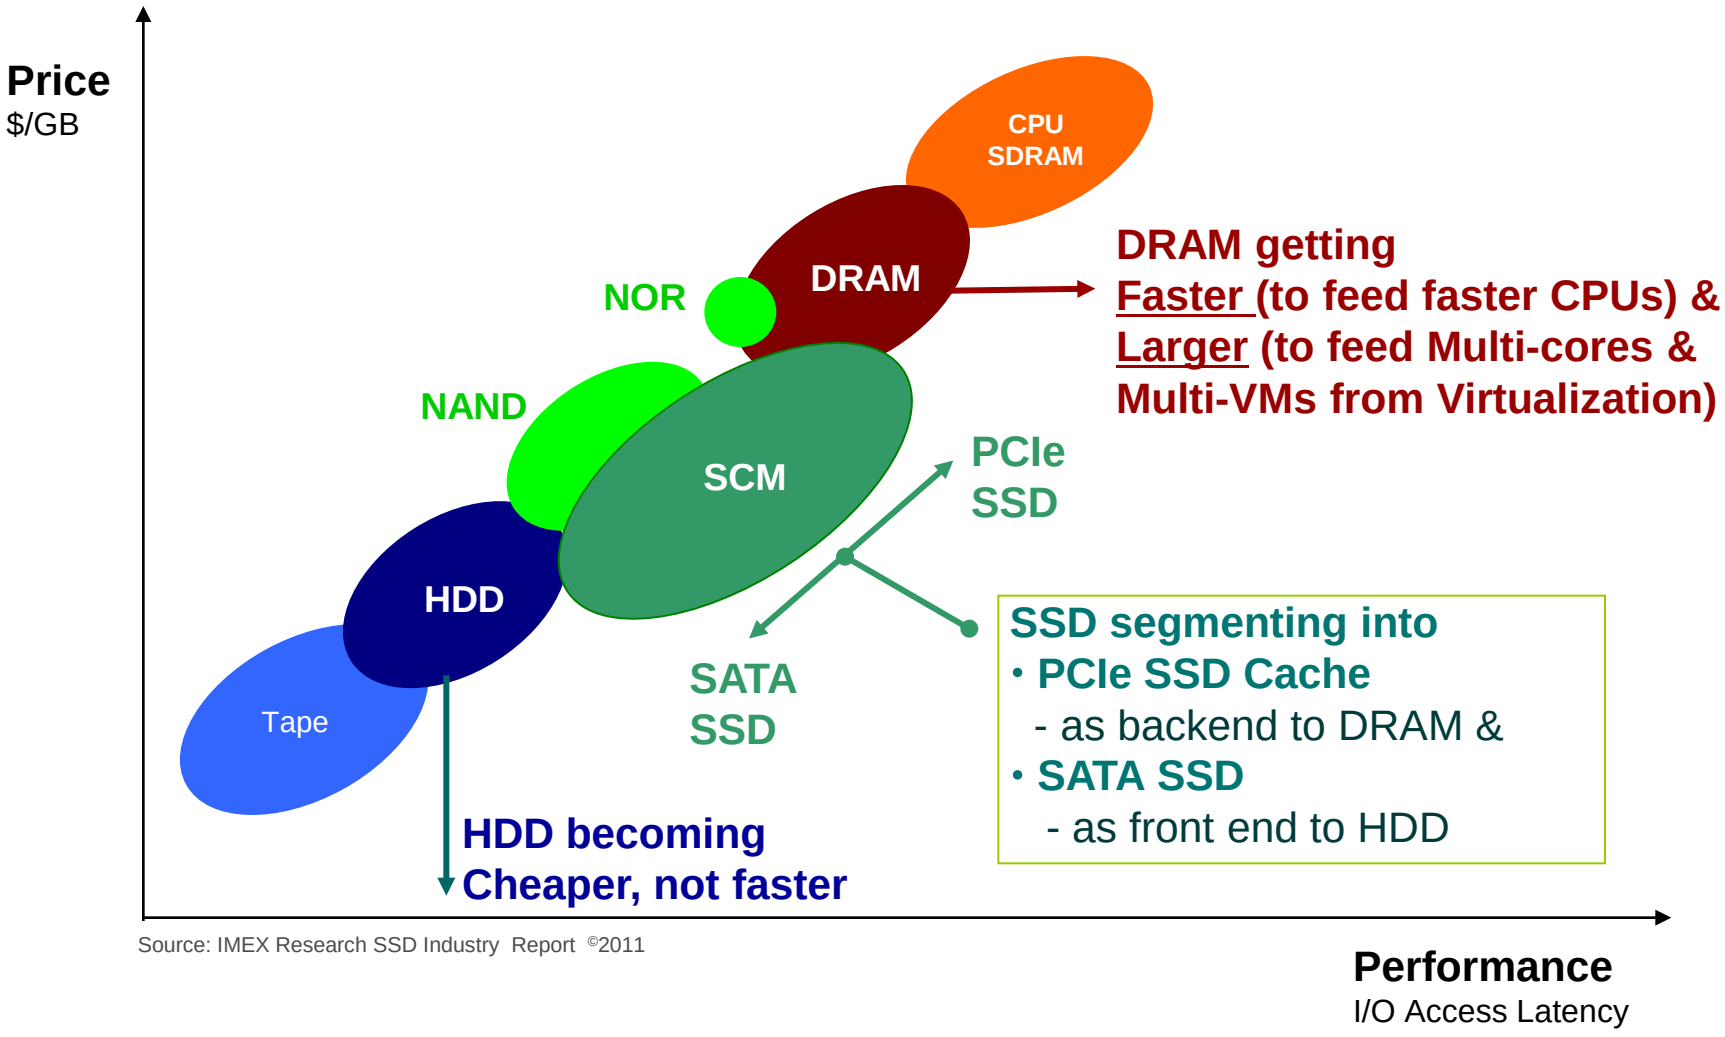

Source: IMEX Research SSD Industry Report ©2011

SSDs - Price Erosion & IOPS/GB

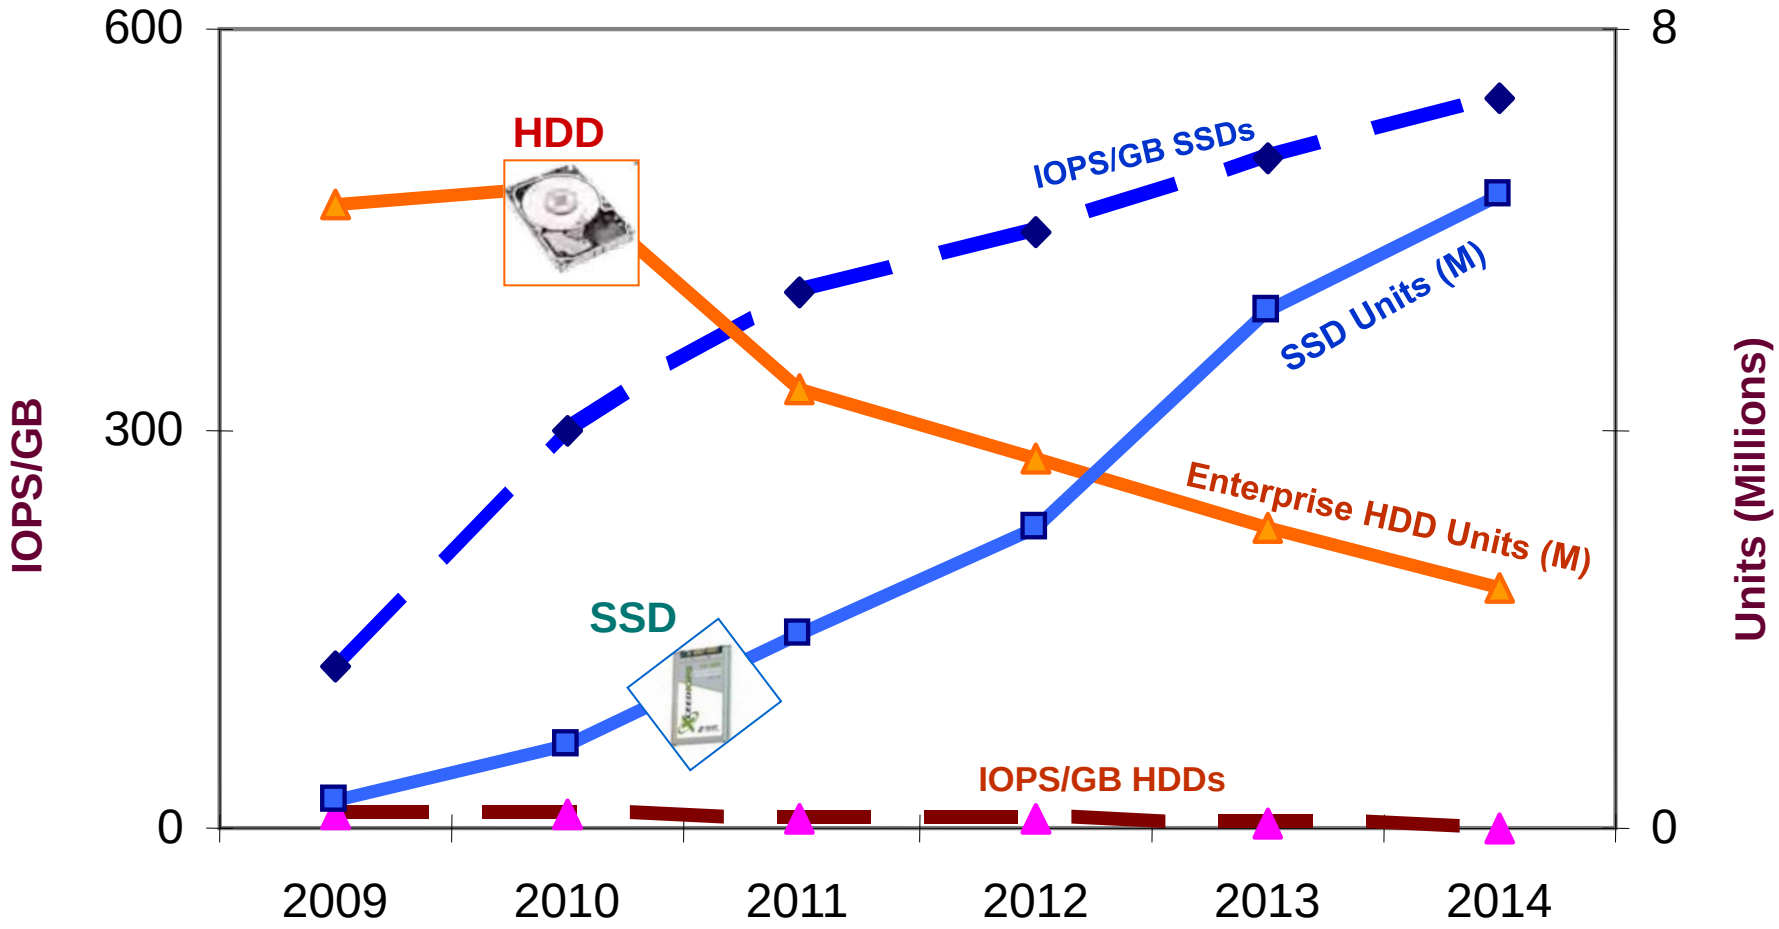

Note: 2U storage rack, • 2.5" HDD max cap = 400GB / 24 HDDs, de-stroked to 20%, • 2.5" SSD max cap = 800GB / 36 SSDs

Key to Database performance are random IOPS. SSDs outshine HDD in IO price/performance – a major reason, besides better space and power, for their explosive growth.

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

Improving Query Response Time for DB

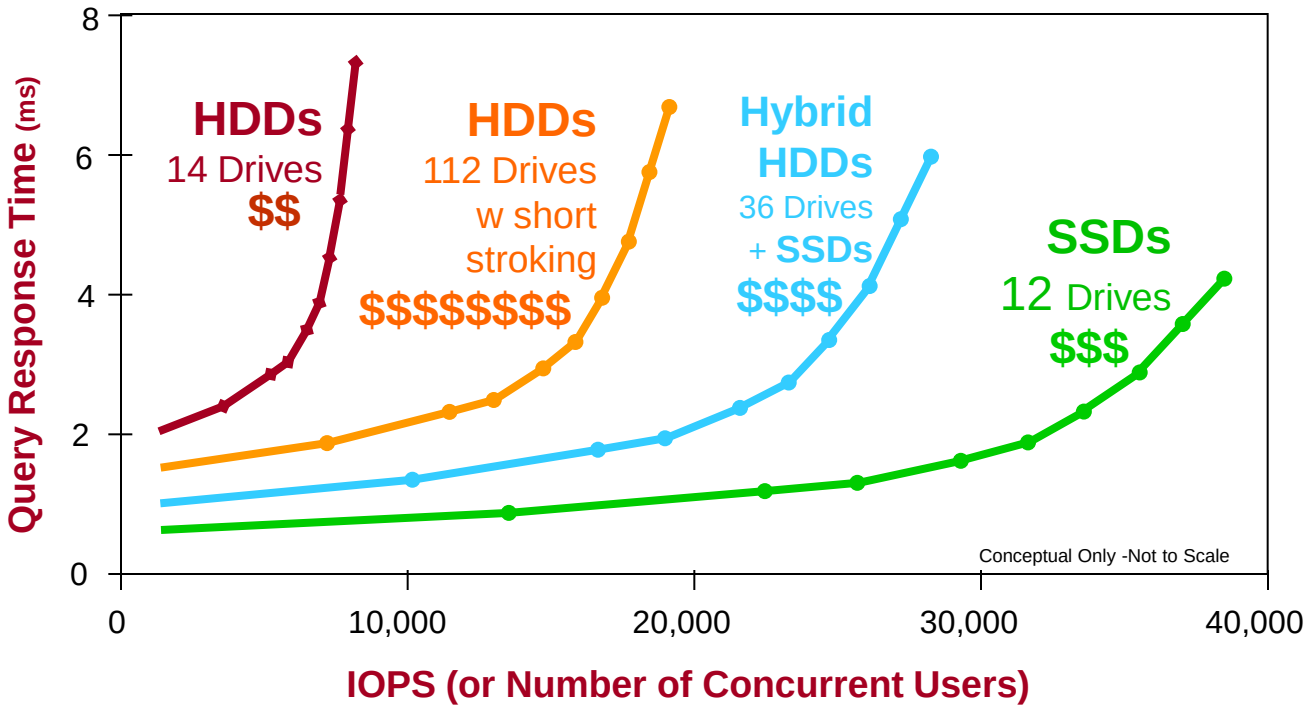

For a targeted query response time, many more concurrent users can be added cost-effectively when using SSDs or SSD + HDDs storage vs. forcing more HDDs or short-stroking HDDs – particularly when running DB & OLTP Apps

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

SCM: A New Storage Class Memory

Education SNIA

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 Storage Class Memories

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

- **Performance** - Latency (Random/Block RW Access-ms); Bandwidth BW(R/W- GB/sec)
- **Data Integrity** - BER (Better than 1 in 10^17)
- **Reliability** - Write Endurance (30K PE Cycles No. of writes before death);
- - Data Retention (5 Years); MTBF (2 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 (oC,V); Radiation (Rad)

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

Use SSDs where best price/perf achieved

Databases have different types of performance demands on storage

- Data access
- Index Reads (Read Intensive operations)
  - DB file sequential read
  - DB file scattered read
- Redo/Undo logs, Temp table space (Write intensive operations)
  - Log file parallel write
  - Log file sync
  - Log file single write
- Buffer Cache events (Write intensive operations)
  - Free buffer wait
  - Control file parallel write
  - Buffer busy waits

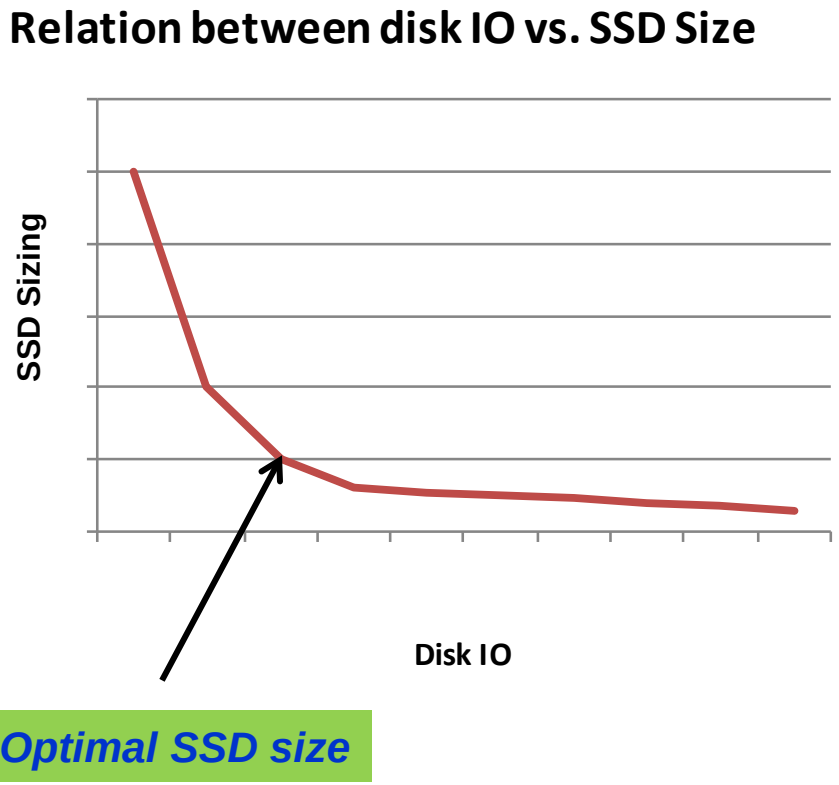

Source:©IMEX Research Industry Report 2009-12

Buying SSD vs. Memory to Improve TPS

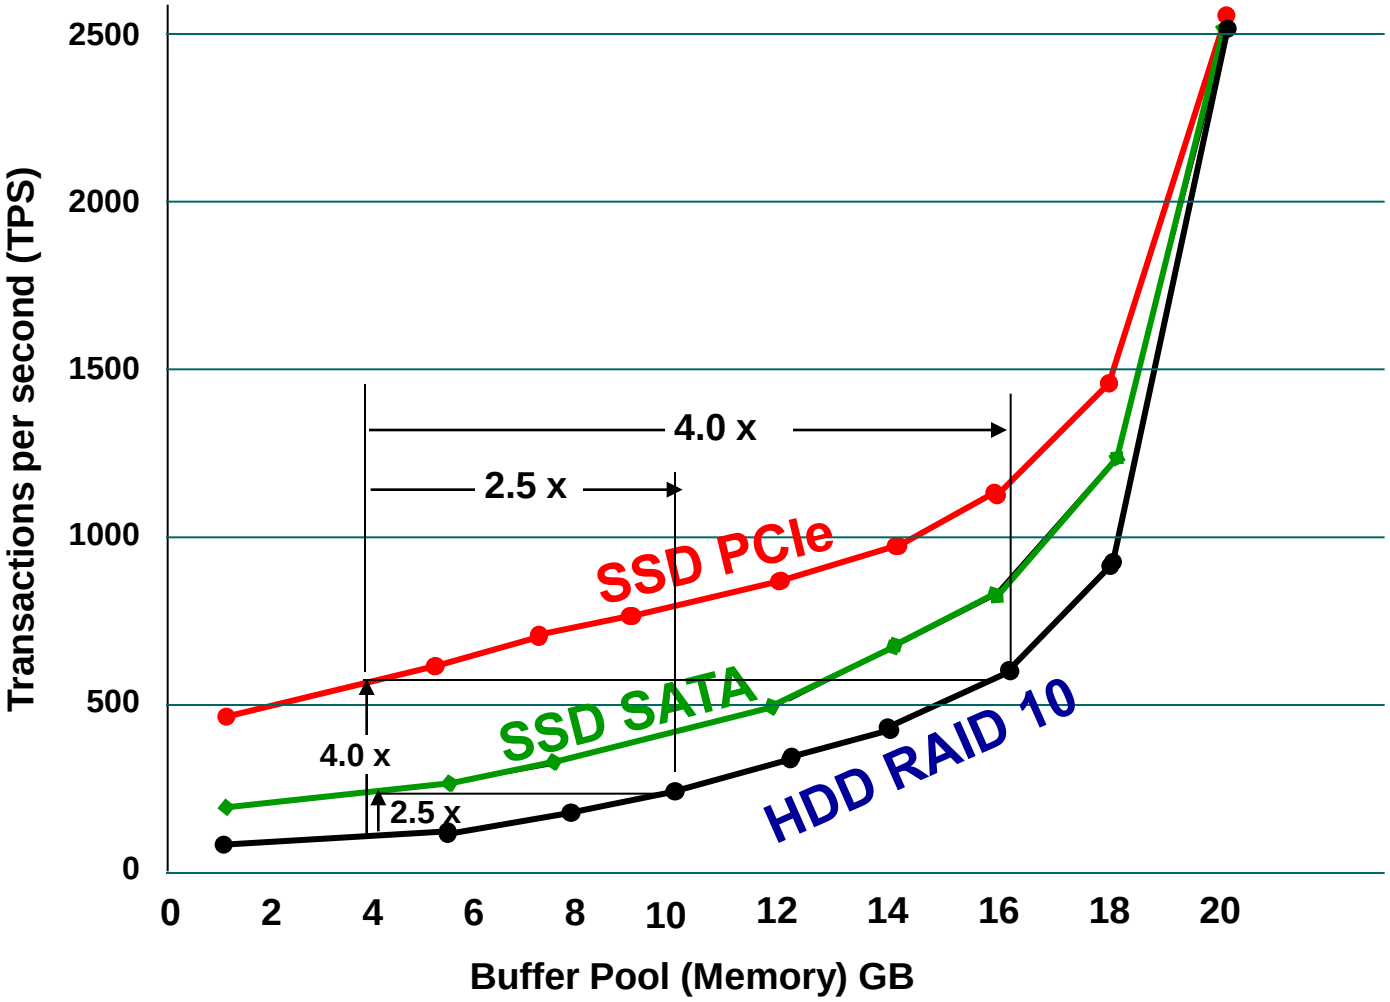

To achieve a certain TPS improvement, its cheaper to deploy SSD vs increased buffer memory (in GB costs) needed with using HDDs alone.

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

SSD Challenges & Solutions: Goals & Best Practices

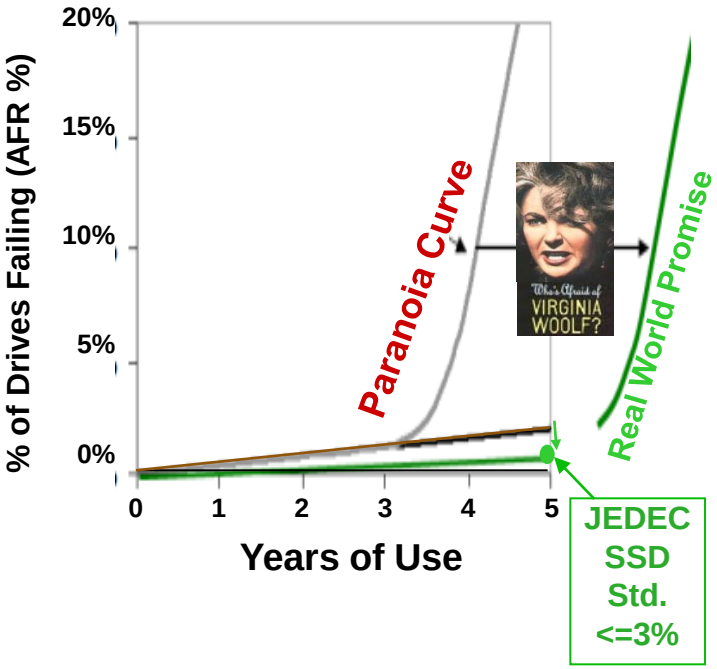

Concerned about SSD Adoption in your Enterprise ?

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

Best Practices

- By leveraging Error Avoidance Algorithms, and Best Practices of Verification Testing, to keep 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 Over-provisioning will provide large increases in random performance and endurance.
- Select SSD based on confirmed EVT Ratings
- Use MLC within requirements of Endurance Limits

Using Best-of-Breed Controllers to achieve <=3% AFR and JEDEC Endurance Verification Testing should allow Enterprise Capabile SSDs

Source: Intel IDF'10 & IMEX Research SSD Industry Report 2011 ©IMEX 2010-12

Source:©IMEX Research Industry Report 2009-12

Hybrid SSD Storage - Perf & TCO

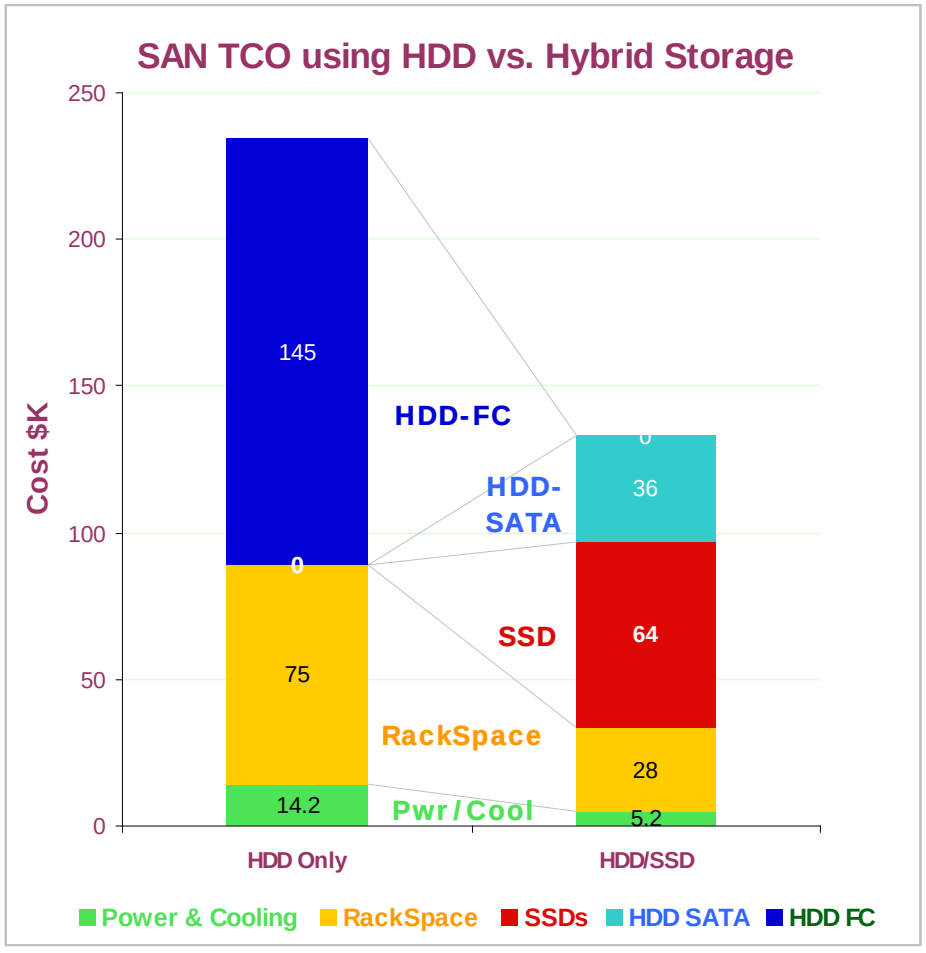

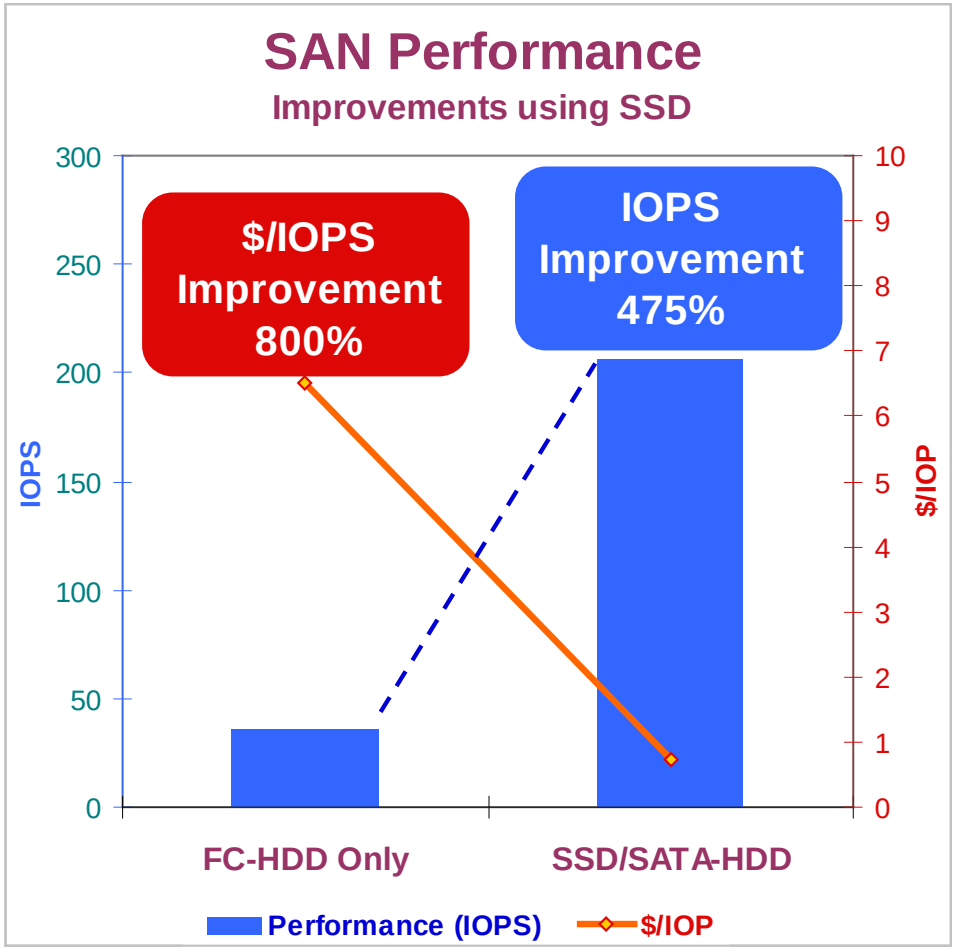

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

New Intelligent Controllers: SSD Storage Architecture

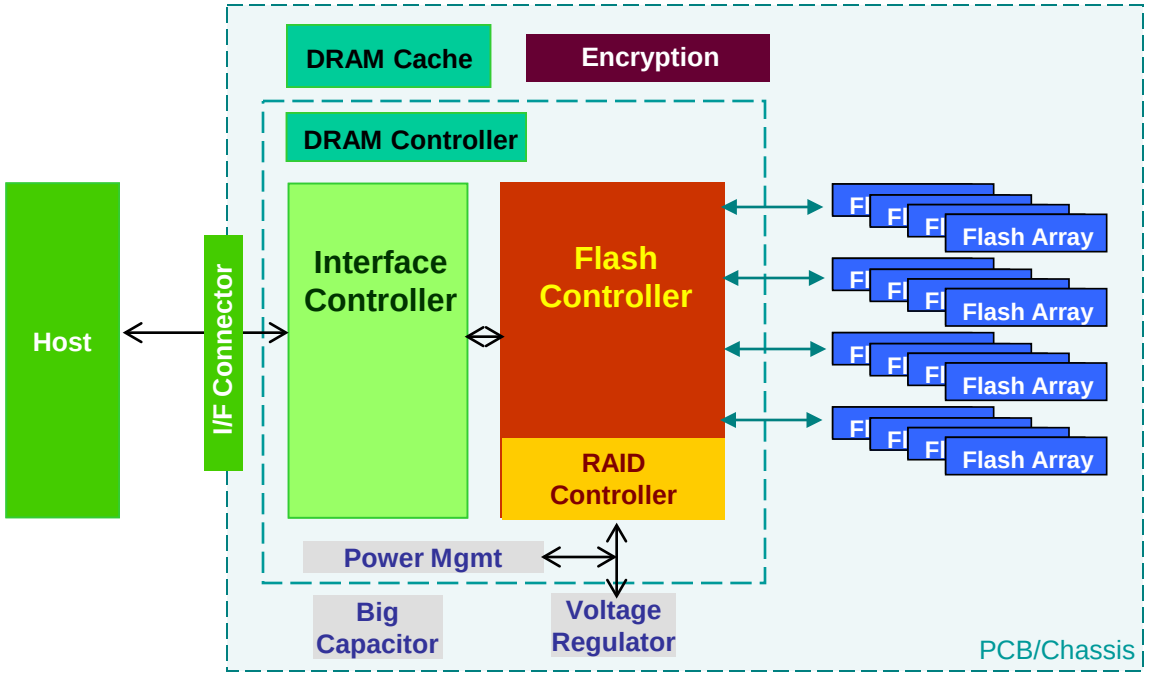

1	Interface Controller	Signaling Mgmt, Interpret WR/RD/Status Commands, Native Command Queuing, Move Data <-> Host
2	Flash Controller	Signaling Mgmt, Format, Interpret WR/RD/Status Commands for Flash Arrays, Move Data. Defect Mapping/Bad Block Mgmt, Wear Leveling, Physical<>Logical Translations, ECC…
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

Source:©IMEX Research Industry Report 2009-12

New Intelligent Controllers: Managing NAND Media in NexGen 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

HDD Media

Flash 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

Source:©IMEX Research Industry Report 2009-12

New Intelligent Controllers - Performance in Next Gen SSDs

Managing Factors Impacting 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** - "Burst" First On Board (FOB), Steady State post xPE Cycles

By using New Gen Controllers, performance of MLC SSDs starting to match performance of some SLC SSDs

Additional performance gains with interleaved memory banks, caching and other techniques

Source: ©IMEX Research Industry Report 2009-12

New Intelligent Controllers : Managing Endurance in NextGen SSDs

Education SNIA

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.

With advanced Errors Avoidance (ECC/Wear-Leveling etc) and capacity over-provisioning techniques, controllers successfully creating endurance for over 5-Year Product Life Cycle in Enterprise SSDs

Source:©IMEX Research Industry Report 2009-12

New Intelligent Controllers: Meeting Enterprise Requirements

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.

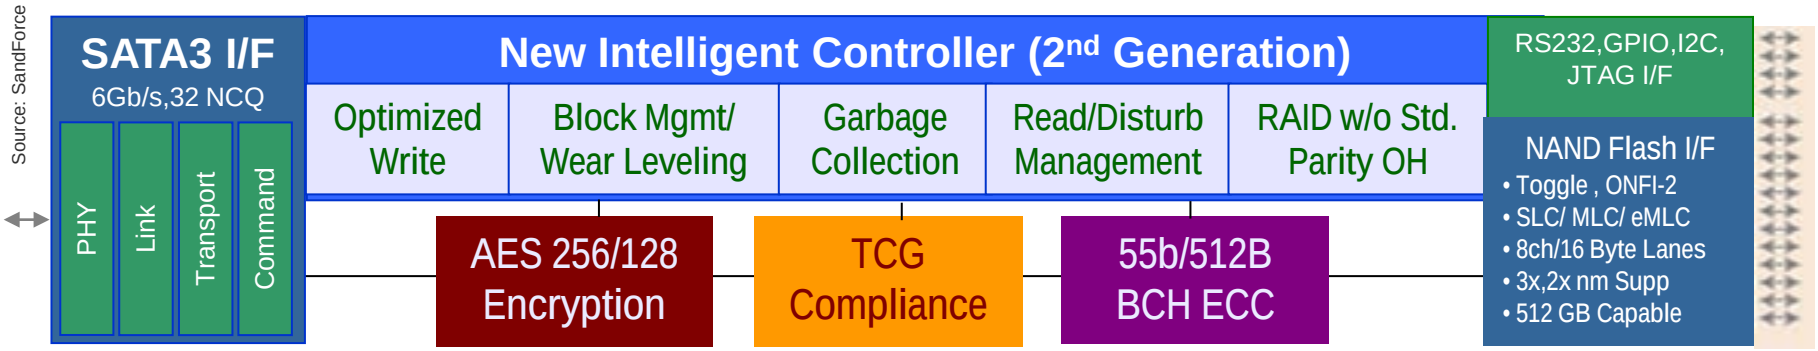

New Gen Controllers allow SSDs to meet Enterprise Class Availability/Performance/ over 5-Year Life/Scalability/ Auto-Configuration & Auto Data-Tiering

Source:©IMEX Research Industry Report 2009-12

Applications Best Suited for SSDs

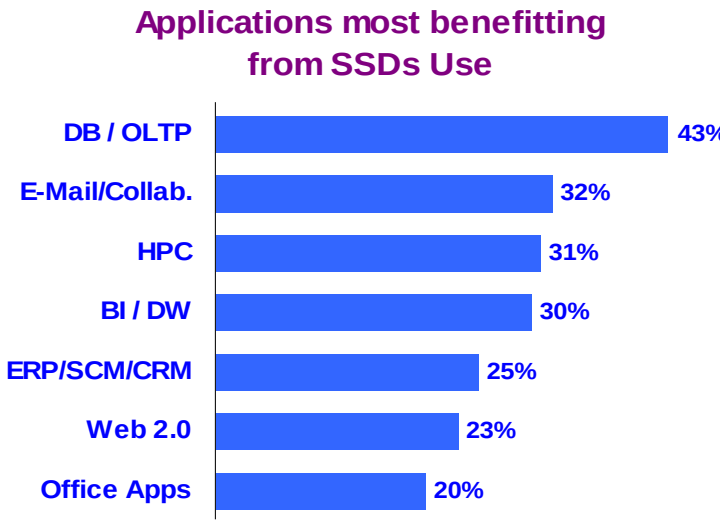

Apps and impact from SSD Usage

- **Databases**
  - Databases have key elements of commit files
  - logs, redo, undo, tempDB
- **Structured data vs. Unstructured Data**
  - Structured/SQL data access is an excellent fit for SSD
  - Exception–large, growing table spaces
  - 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

- **Financials/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

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

Apps best suited for SSDs: Workloads Characterization

- **Data Warehousing Workloads are I/O intensive**
  - Predominantly read based with low hit ratios on buffer pools
  - High concurrent sequential and random read levels
    - ✓ Sequential Reads requires high level of I/O Bandwidth (MB/sec)
    - ✓ Random Reads require high IOPS)
  - Write rates driven by life cycle management and sort operations
- **OLTP Workloads are strongly random I/O intensive**
  - Random I/O is more dominant
    - ✓ Read/write ratios of 80/20 are most common but can be 50/50
    - ✓ Can be difficult to build out test systems with sufficient I/O characteristics
- **Batch Workloads are more write intensive**
  - Sequential Writes requires high level of I/O Bandwidth (MB/sec)
- **Backup & Recovery times are critical for these workloads**
  - Backup operations drive high level of sequential IO
  - Recovery operation drives high levels of random I/O

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

Applications Best Suited for SSDs: Data Warehouse/BI

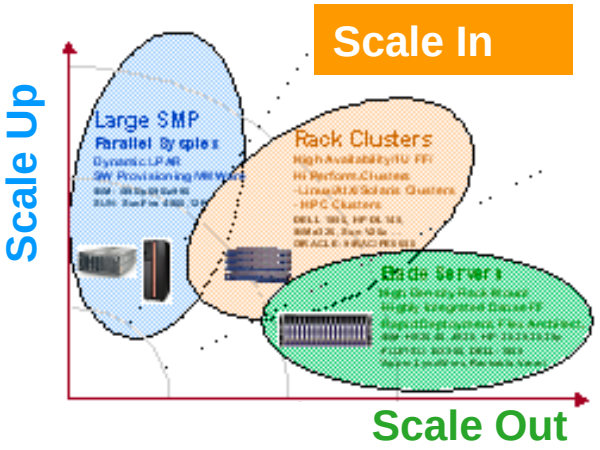

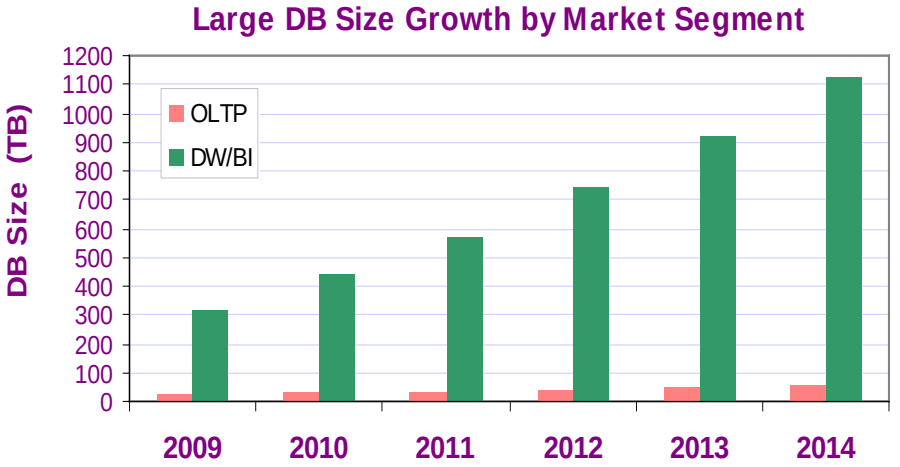

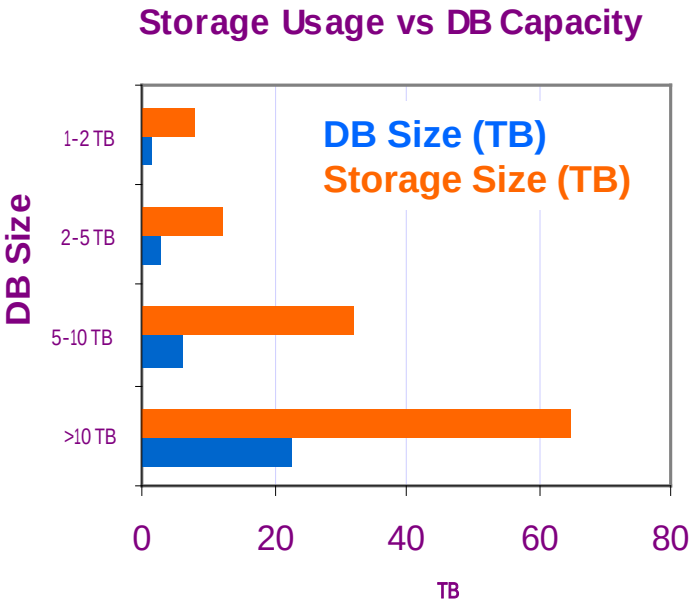

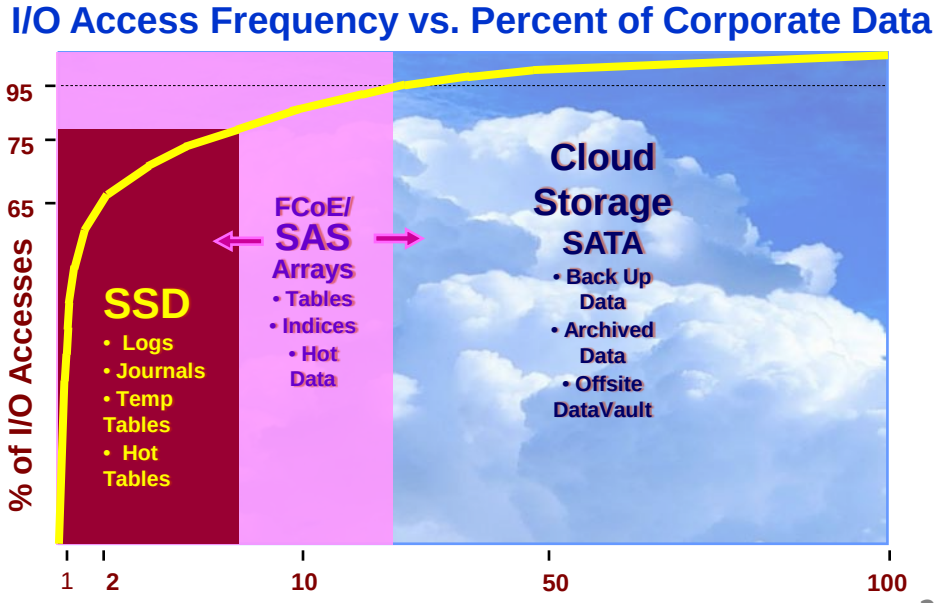

Data Source: IMEX Research
Cloud Infrastructure Report ©2009-11
Source:©IMEX Research Industry Report 2009-12

Cached Storage - Performance Benefits

Intelligent Caching at the Storage Server Array

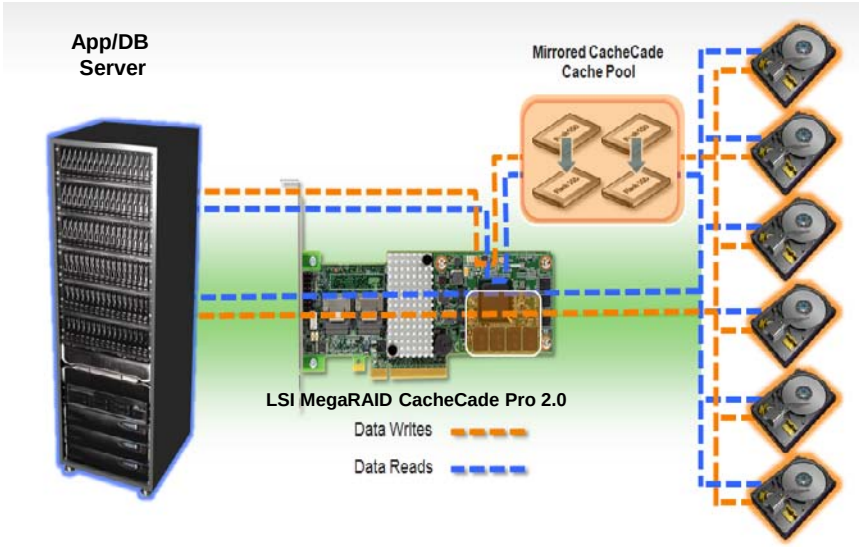

Application	Improvement over Cached vs.HDD only
Oracle OLTP Benchmarks	681%
SQL Server OLTP Benchmark	1251%
Neoload (Web Server Simulation	533%
SysBench (MySQL OLTP Server)	150%

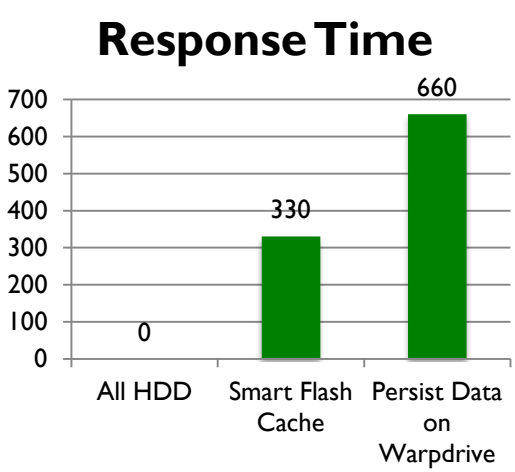

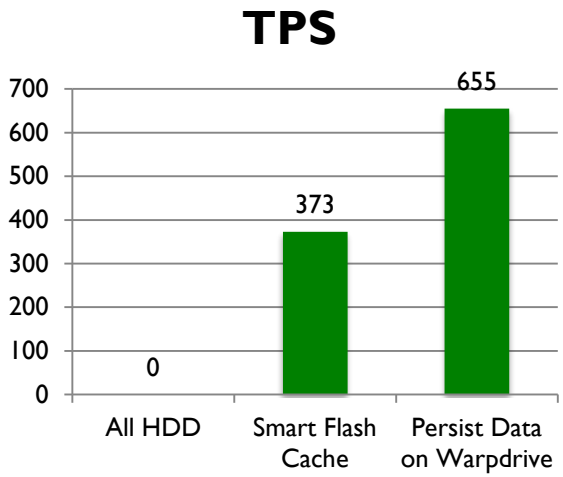

Source:©IMEX Research Industry Report 2009-12

New Storage Hierarchy in NGDC & Clouds

I/O Access Frequency vs. Percent of Corporate Data 2015

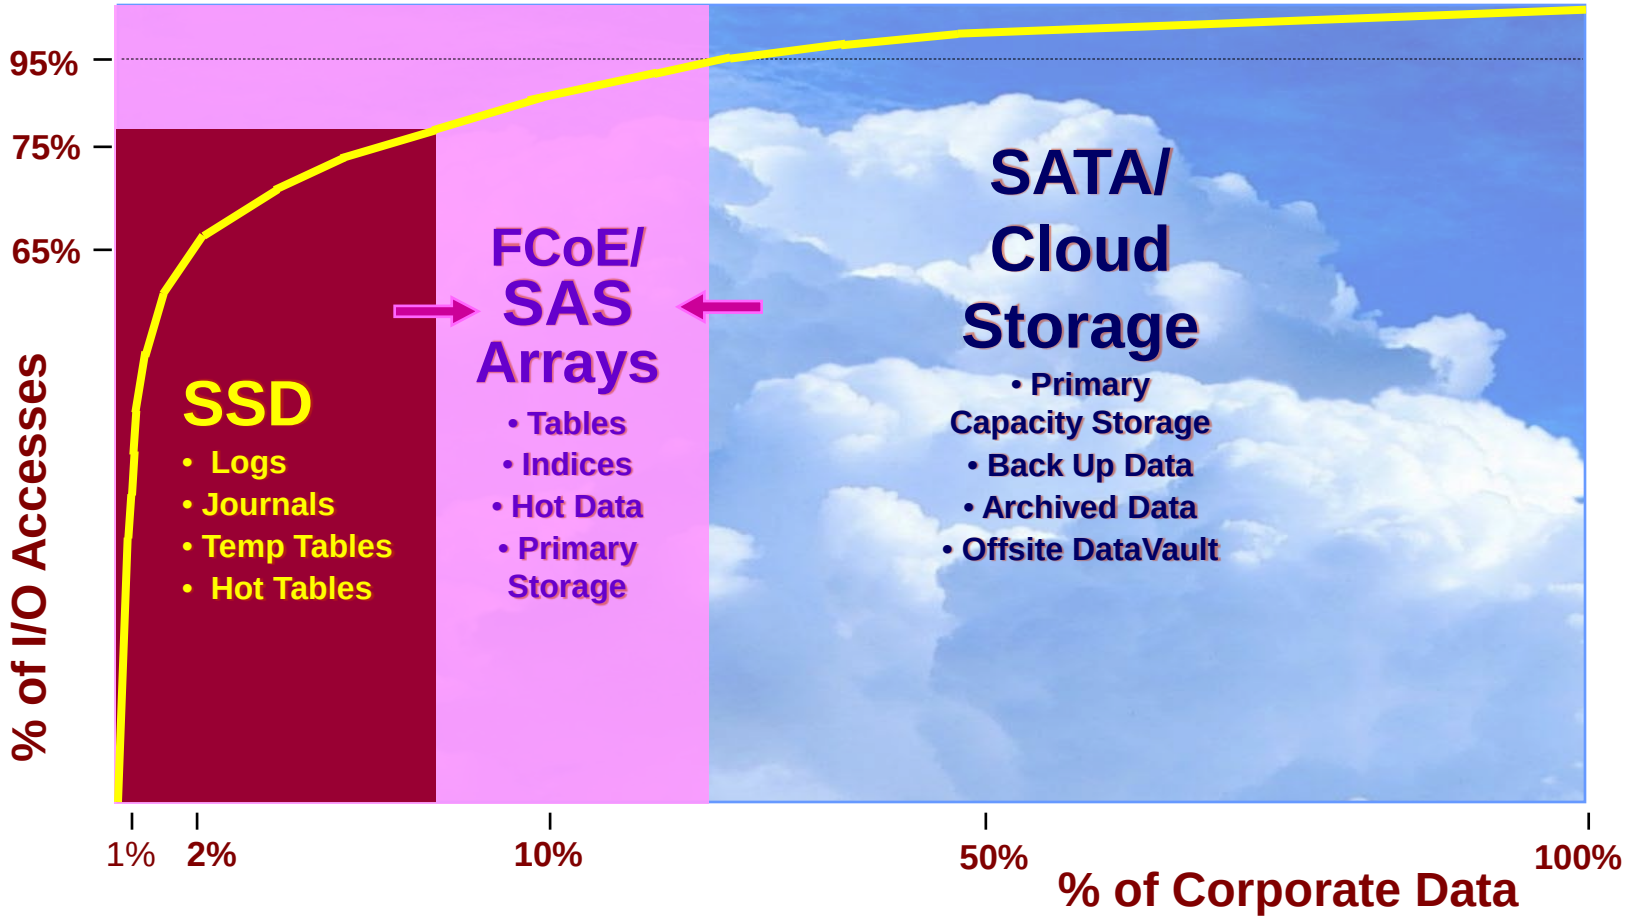

Data: IMEX Research - Cloud Infrastructure Report ©2009-11
Source:©IMEX Research Industry Report 2009-12

IO Bottleneck Mitigation in Virtualized Servers

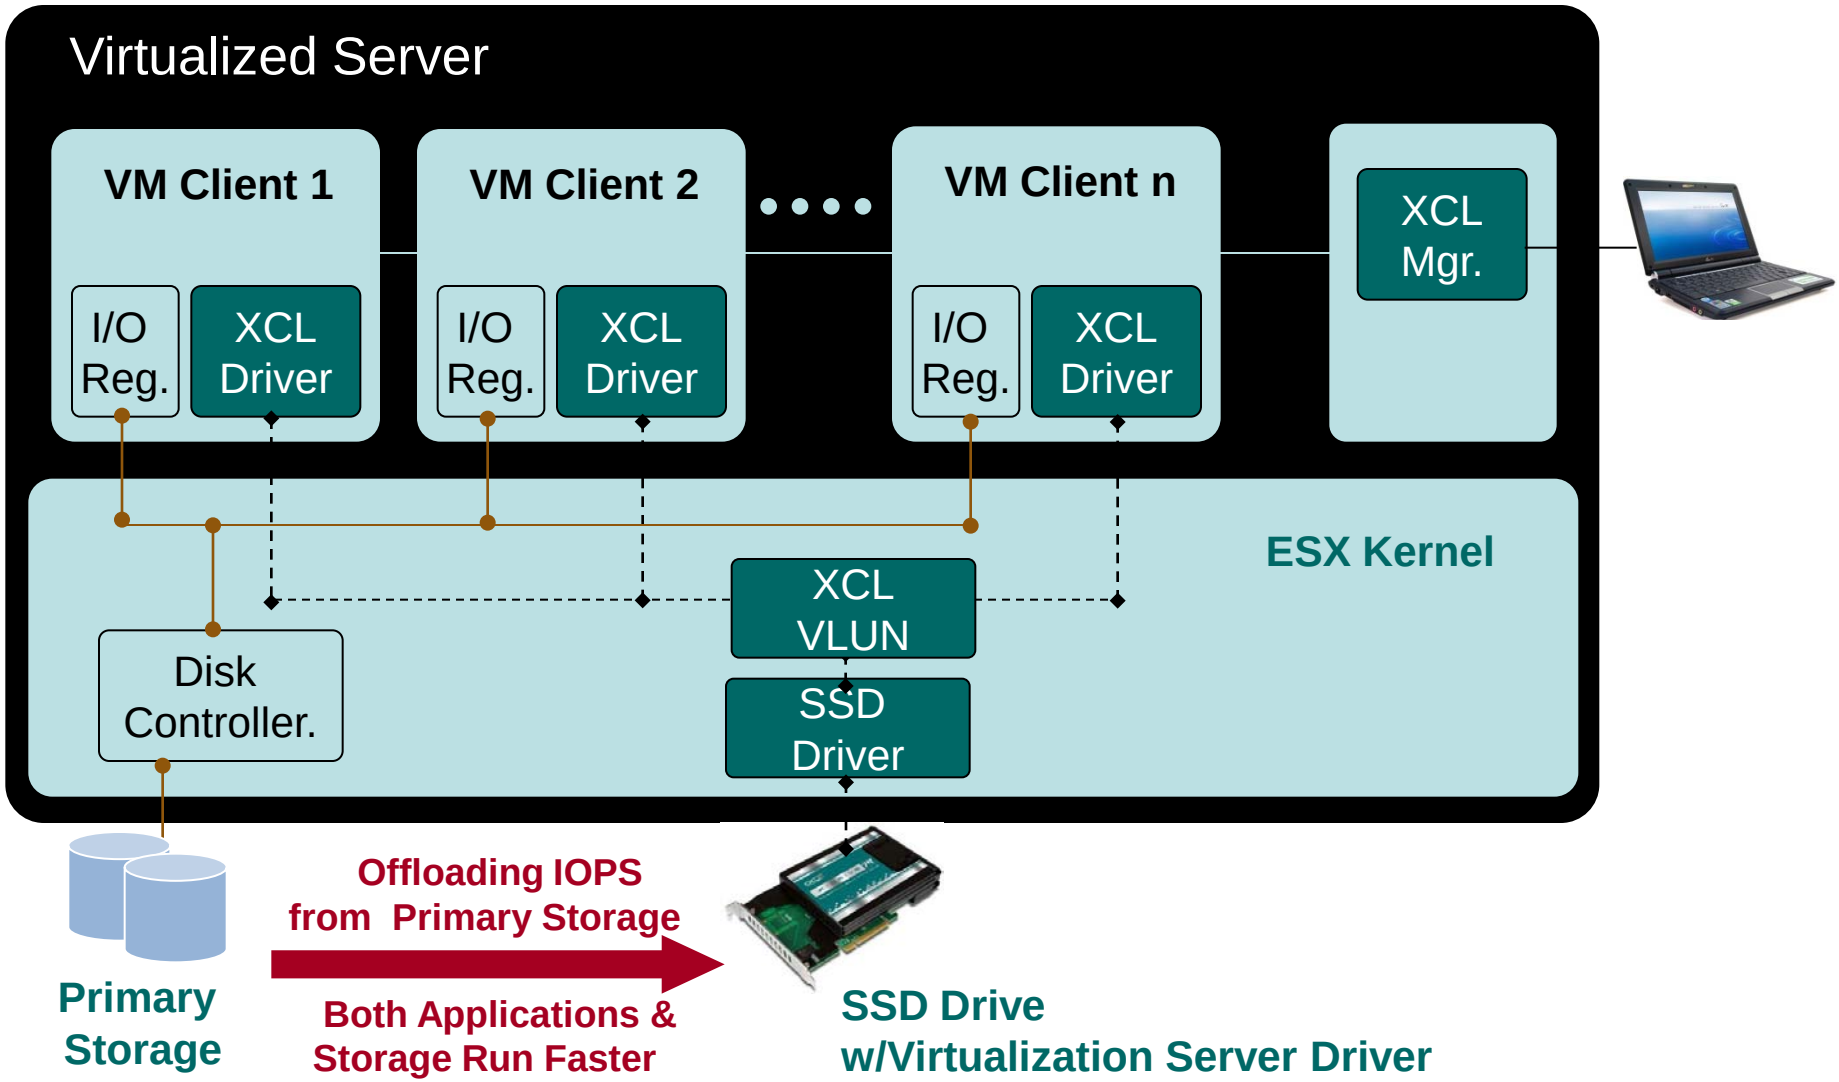

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

I/O Forensics for Auto Storage-Tiering

Storage-Tiered Virtualization

Storage-Tiering at LBA/Sub-LUN Level

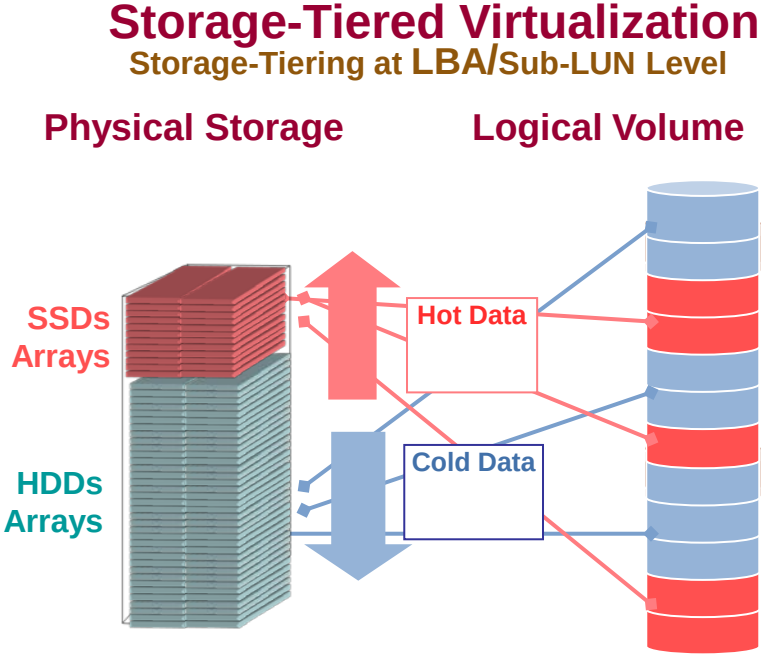

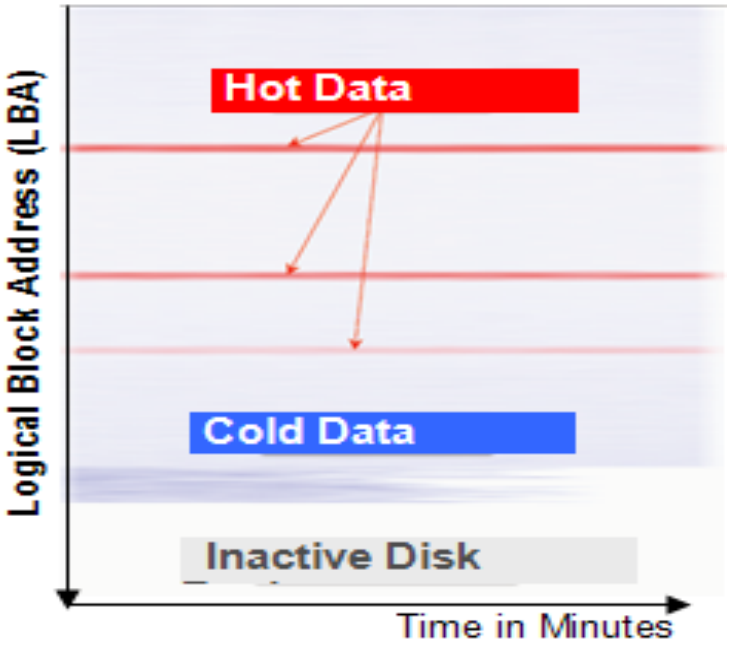

LBA Monitoring and Tiered Placement

- Every workload has unique I/O access signature
- Historical performance data for a LUN can identify performance skews & hot data regions by Logical Block Addresses

IBM & IMEX Research SSD Industry Report ©IMEX 2011
Source:©IMEX Research Industry Report 2009-12

Automated Storage Tiering: Improving Response Time

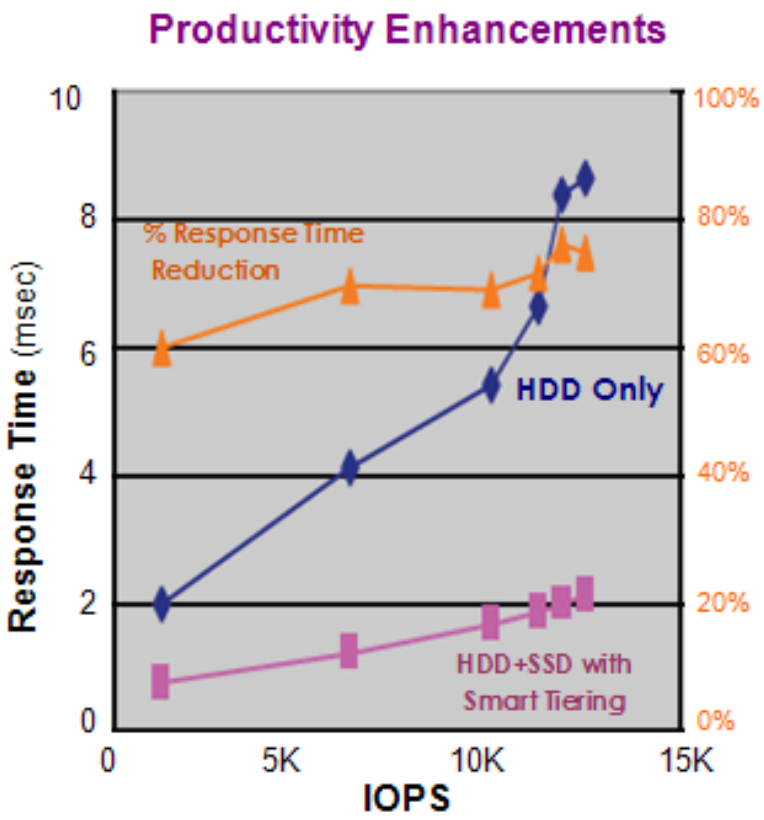

Productivity Improvements (Response Time)

- With automated reallocation of hot spot data (~ 5-10% of total data) to SSDs, performance improvements
  - **Response time reduction ~70+% IOPS increase of 200%** for any I/O intensive workloads in Time-Perishable OLTP markets: Airlines Reservations, Wall Street Investment Banking Stock Transactions, Financial Institutions Hedge Funds etc.
  - **Performance boost** in Low Latency seeking Systems (High Perf. Clustered Systems)

Source: IMEX Research SSD Industry Report ©2011

Source:©IMEX Research Industry Report 2009-12

SSS in DB Environments: DB Improvements with Flash SSDs

- Storage management, performance and cost - a big issue in DBs
- SSDs enable super IO performance and cost reduction in DBs

Improve Responsiveness	**Improve** • Insert/Update/Delete Performance • Random Read I/O Performance • Query Response Time • Sort Performance • Batch Performance too.
Reduce Costs	**Reduce** • DRAM size for Buffer Pools used to cache data on SSDs • Power/Cooling Space for housing Databases
New Added Benefits	**Reduce** • Database Recovery performance • I/O performance impact by Flash Copy • Skill levels required for DB tuning & monitoring

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

SSS in DB Environments: Best Practices for DB/DW/BI Apps

Goals

- Establish **Goals for SLAs** (Performance/Cost/Availability), BC/DR (RPO/RTO) & Compliance
- **Increase Performance for DB, Data Warehousing, OLTP Apps**:
  - Random I/O > 20x , Sequential I/O Bandwidth > 5x
  - Remove Stale data from Production Resources to improve performance
- **Classify Data -** Use **Partitioning Software**
- By Frequency of Access (Recent Usage) and
  - Capacity (by percent of total Data) using general guidelines as:
  - Hyperactive (1%), Active (5%), Less Active (20%), Historical (74%)

Implementation

- **Optimize Tiering** by Classifying Hot & Cold Data
  - Improve Query Performance by reducing number of I/Os
  - Reduce number of Disks Needed by 25-50% using advance compression software achieving 2-4x compression
- **Match Data Classification vs.Tiered Devices** accordingly
  - Flash, High Perf Disk, Low Cost Capacity Disk, Online Lowest Cost Archival Disk/Tape
- **Balance Cost vs. Performance** of Flash
  - More Data in Flash > Higher Cache Hit Ratio > Improved Data Performance
- **Create and Auto-Manage Tiering** (Monitoring, Migrations, Placements) without manual intervention

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

Storage Tiering - Best Practices Highlights

- **SSD-PCIe perform better** than SATA SSDs
  - Use Nehalem Class CPUs especially when using PCIe SSDs
- Put **Random Access Files on SSDs** (Index, Tables, Table Spaces)
  - Keep ample SSD Reserved Space to avoid massive SSD write deterioration
- Put **Sequentially Written Files on HDDs** since
  - HDDs better at Sequential Writes compared to SSDs
  - Removes SSD Write performance bottle necks
  - Increases SSD life
  - Archive Less Active Tables/Records to HDDs
- Leverage **Auto-Tiering Storage SW** to balance between SSDs and HDDs
  - Heat Mapping with Tier Managed Extent Pools
  - Workload Hot Spot Analysis
  - Smart Data Migration & Placement
  - Continuous Workload Monitoring
- Use Faster Networks (10GbE vs 1GbE) to avoid saturating DRBD
- Target **Price/Performance Economic Benefits of 150-800%**

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

SSD class memories fundamentally changing Computing Systems Architectures

Using SSDs, a leading Computer Systems company achieved:

- **Sustained 1 million IOPS with random 4K size**
- **70%RD/30%WR with Queue depth of 16**
- **System Latency 720 us Average**
- **Floor Space Less than 25% vs. Std. Disk Storage System**
- **Energy Used only 55% vs Std. System**
- **Comparable Total Cost New system vs. Standard System**

Note: System Test HW Used:
Host - 2 Servers (26 cores, 28 GB Memory), Cluster - 14 Storage Controller Nodes,
Storage - 31 Storage Arrays with 41 PCIe SSDs 160GB Each

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

Tiered Storage – Industry Ecosystem

Industry Status

- Integrated Storage Tiering Products offered by over top 10 Storage Vendors
- **Major Storage Vendors**
  - Automated Volume Level Tiering (SSD & HDD)
- **New Storage Start Ups**
  – Integrated Flash Caching & Block Level Tiering
- **Cloud Vendors**
  – Adding Shared Cloud for Lowest Cost Backup & Restore Storage

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

Key Takeaways

- **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
- **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**
  - **Data Forensics 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.Non-disruptively migrate hot data from HDD to SSDs.
- **Optimize Infrastructure to meet needs of Applications/SLA**
  - **Performance Economics/Benefits**
  - Typically 4-8% of data becomes a candidate and when migrated to SSDs can provide response time reduction of ~65% at peak loads

Source: IMEX Research SSD Industry Report ©2011
Source:©IMEX Research Industry Report 2009-12

Q&A / Feedback

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

Many thanks to the following individuals for their industry vision and leadership in identifying the I/O Bottlenecks
Late Jim Gray

Author: Anil Vasudeva

Send any questions or comments on this presentation to SNIA: tracktutorials@snia.org

Source:©IMEX Research Industry Report 2009-12