Solid State Storage : Key to Efficient Tiered Storage Infrastructure

Anil Vasudeva
President & Chief Analyst
imex@imexresearch.com
408-268-0800

IMEX RESEARCH.COM

Agenda – Automated Tiered Storage

IMEX RESEARCH.COM

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. Data Forensics and Tiered Mapping
7. Selecting Automated Storage Tiering Software
8. Key Takeaways

Source: IMEX Research SSD Industry Report ©2011

IT Industry's Journey - Roadmap

IMEX RESEARCH.COM

SIVAC by IMEX Research

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)

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

Data Centers & Cloud Infrastructure

IMEX RESEARCH.COM

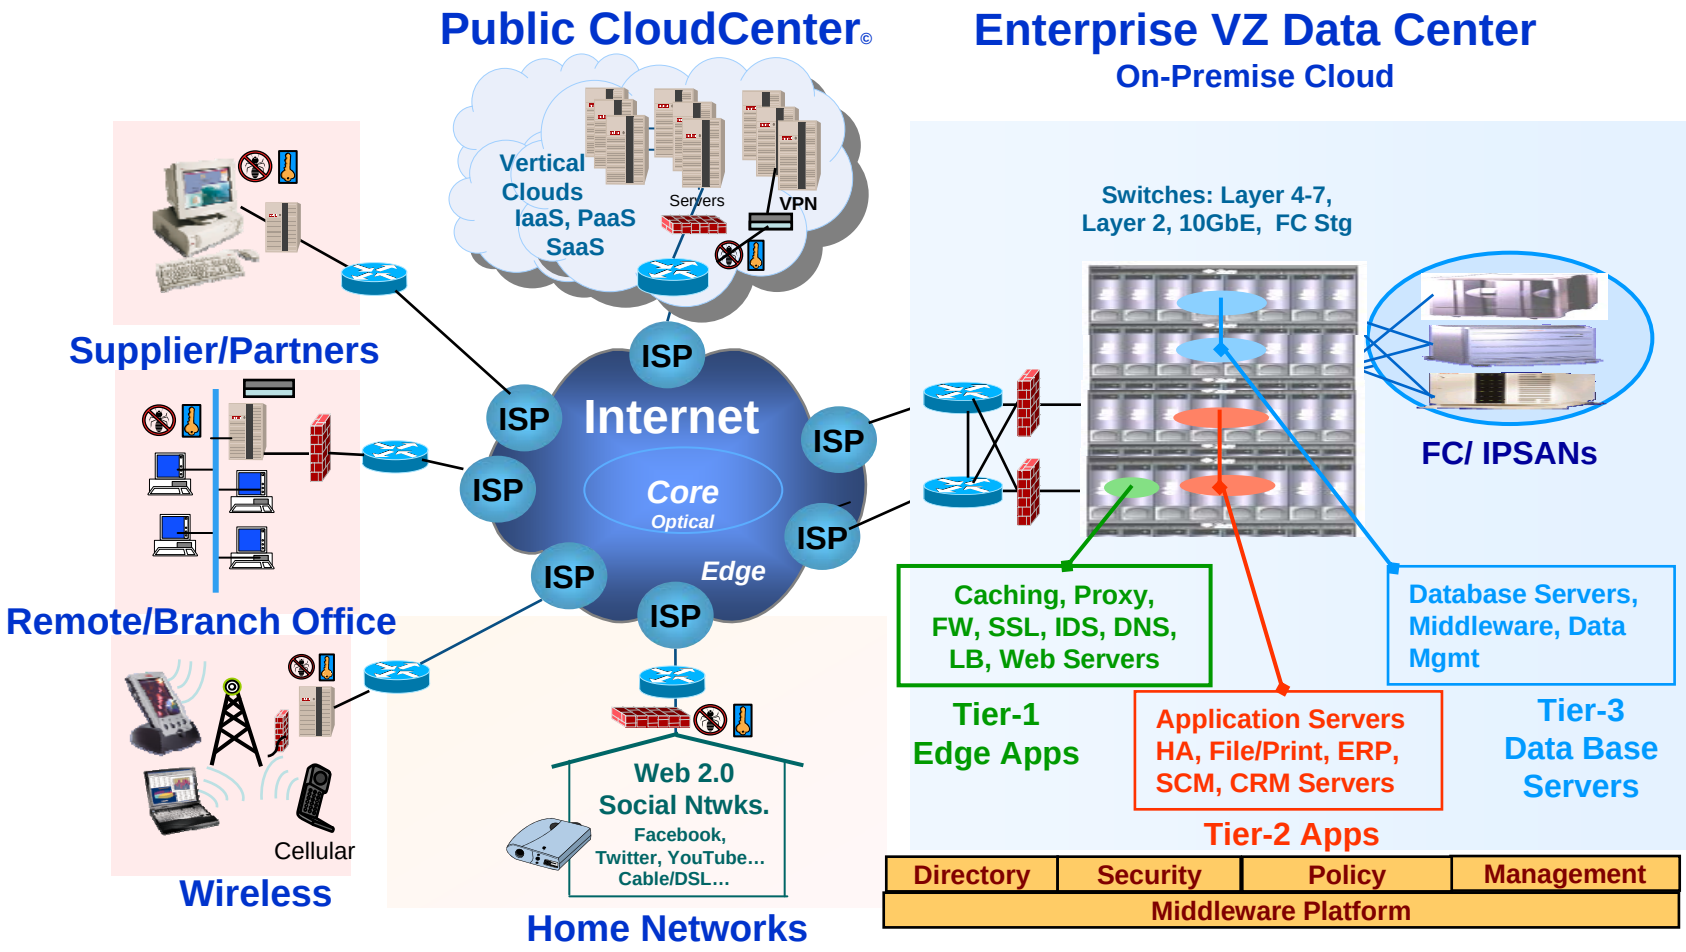

Request for data from a remote client to an enterprise data center crosses a myriad of systems and devices. Key is identifying bottlenecks & improving performance

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

Virtualized Cloud Infrastructure

IMEX RESEARCH.COM

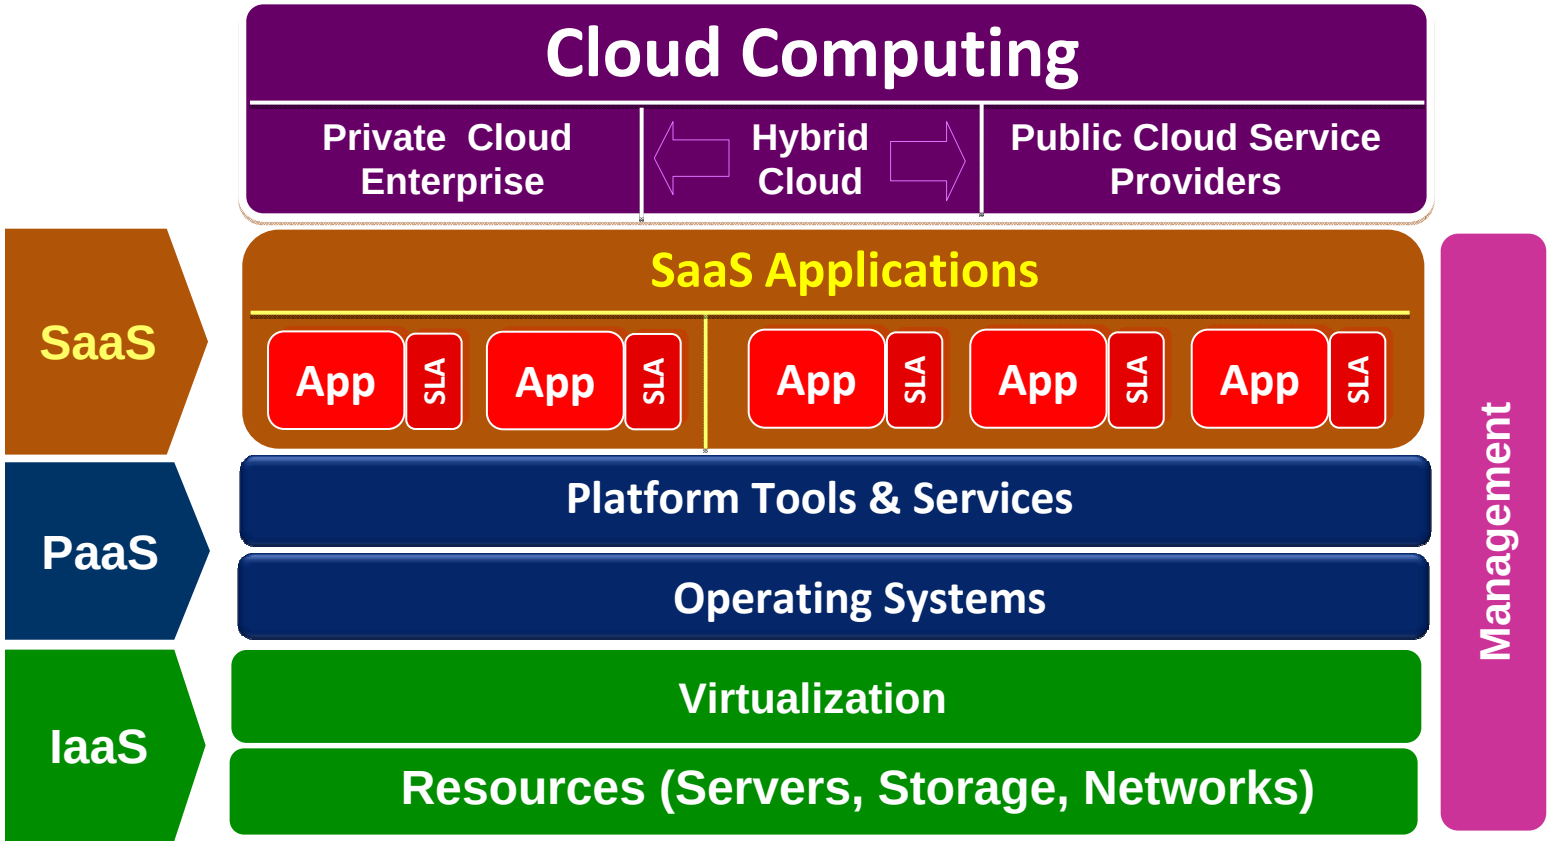

Application's SLA dictates the Resources Required to meet specific requirements of Availability, Performance, Cost, Security, Manageability etc.

Server to Storage IO Gap

IMEX RESEARCH.COM

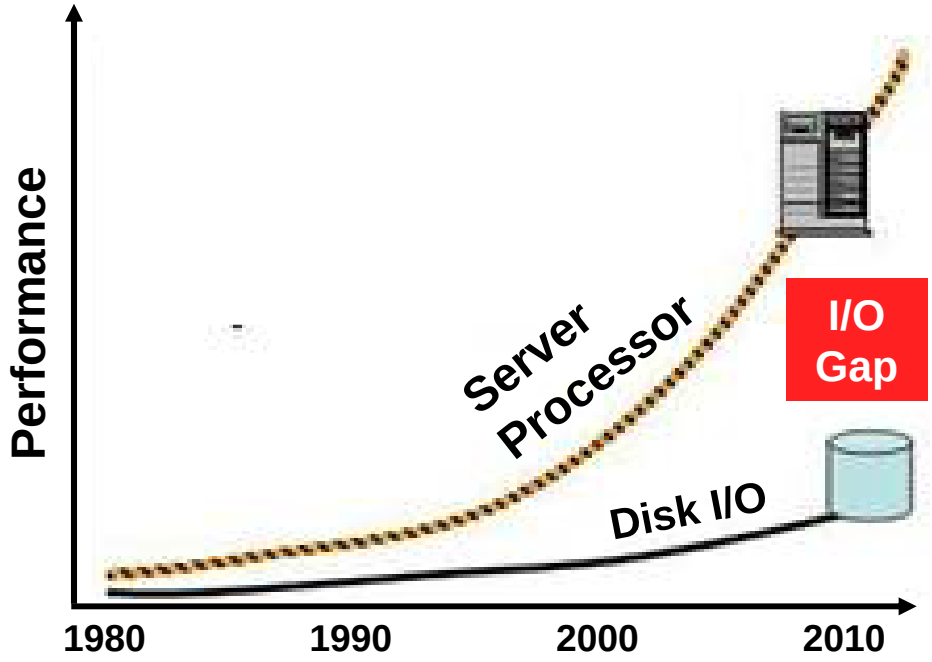

For Each Disk Operation:
Millions of CPU Operations or Hundreds of Thousands of Memory Operations can be done

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 do 100/140 IOPS

SSD Filling Price/Perf Gaps in Storage

IMEX RESEARCH.COM

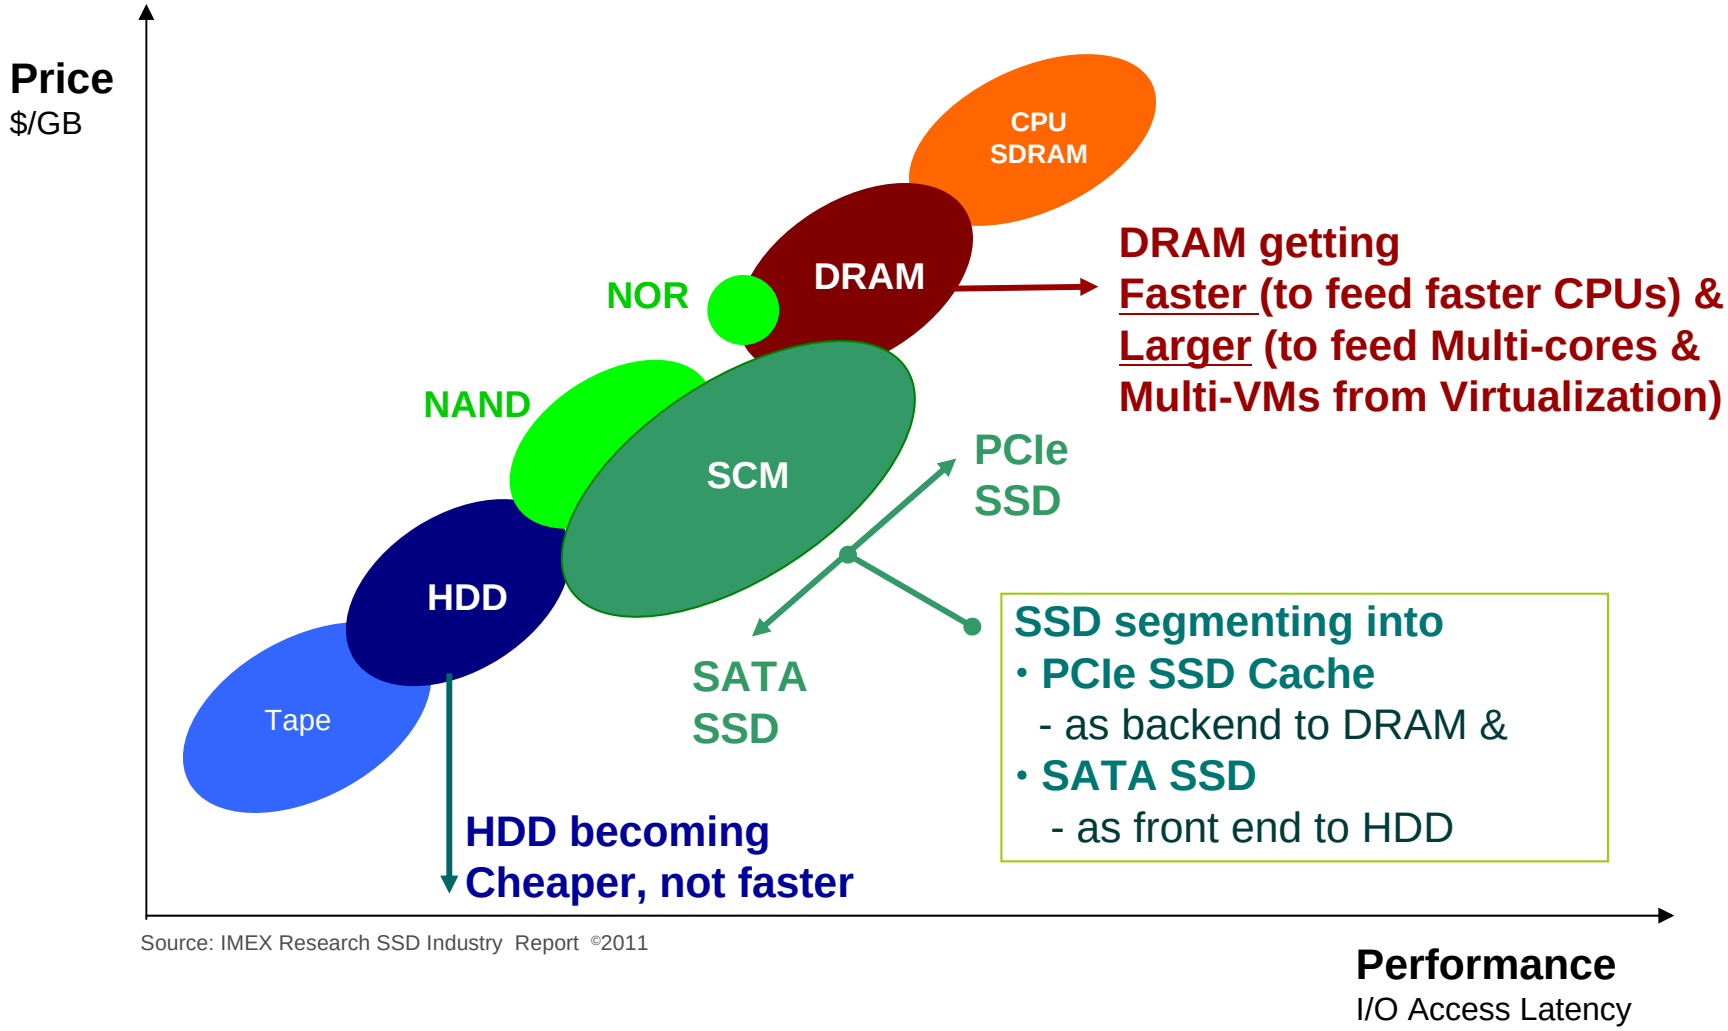

Source: IMEX Research SSD Industry Report ©2011

Source: IMEX Research SSD Industry Report ©2011

SSDs - Price Erosion & IOPS/GB

IMEX RESEARCH.COM

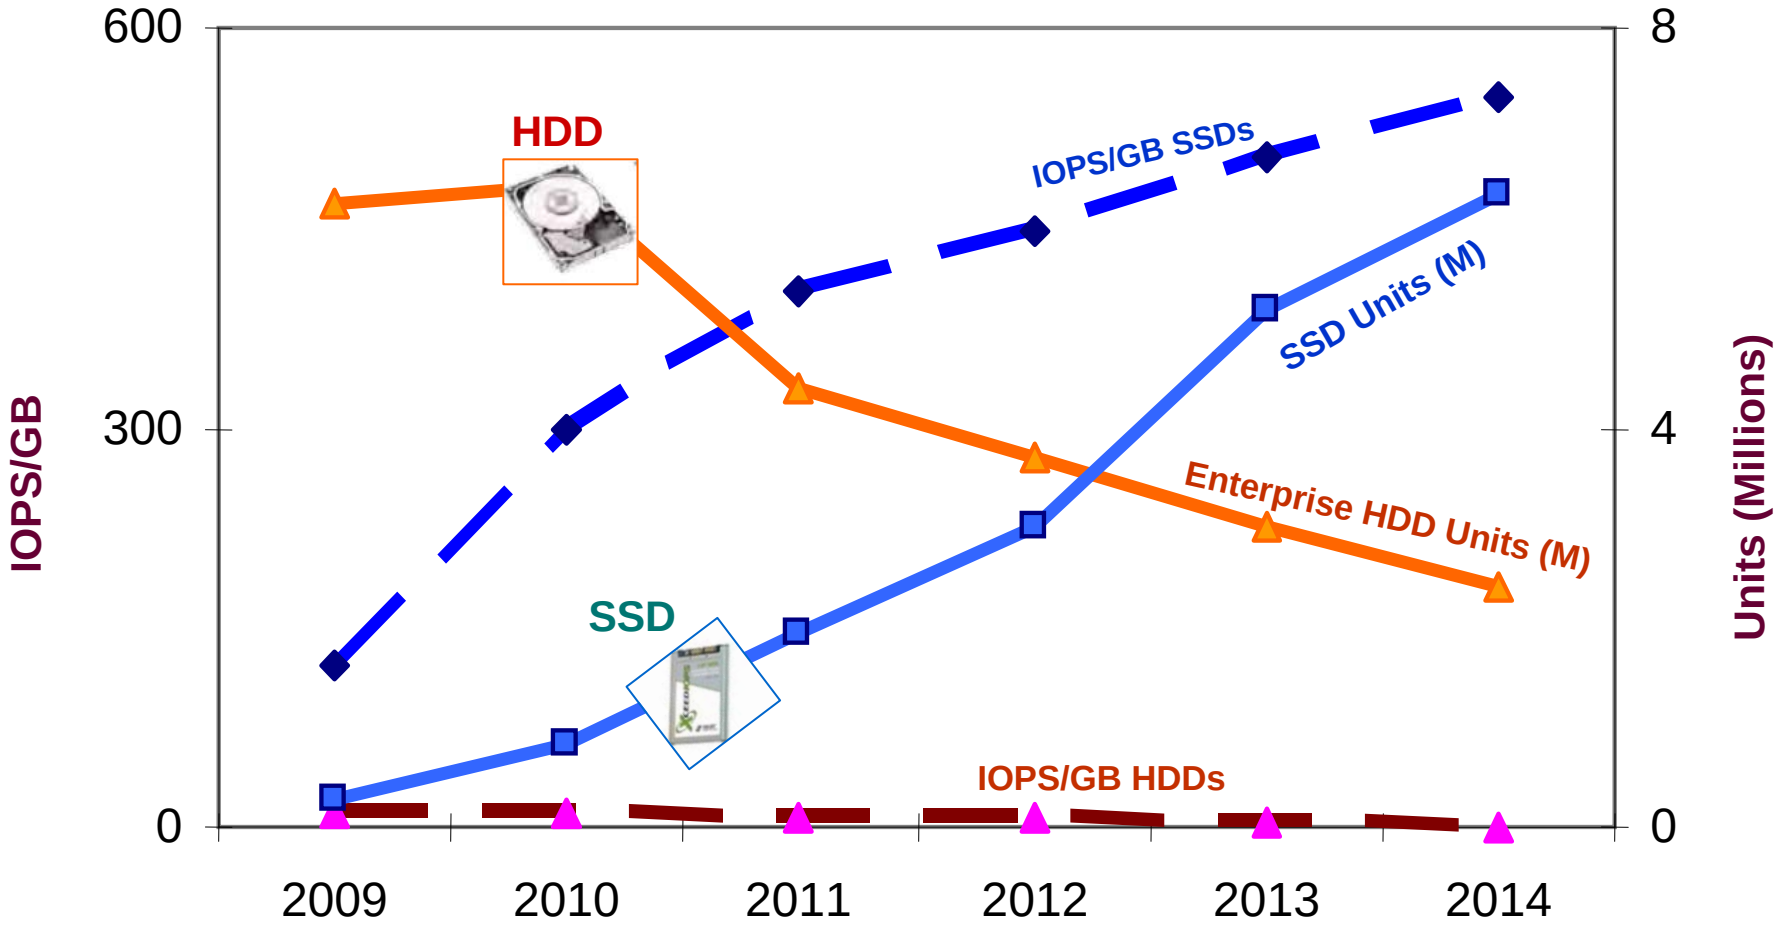

Key to Database performance are random IOPS. SSDs outshine HDD in IO price/performance – a major reason, besides better space and power, for their explosive growth.

Source: IMEX Research SSD Industry Report ©2011

Advantage Solid State Storage (vs. HDDs)

IMEX RESEARCH.COM

HDD	Parameter	SSD	Improvement SSD vs. HDD
	Concurrent Access		900 %
	Data Access Time ms		<1 %
	IOPS		475 %
	Read Speed		500%
1.0	MTBF (Million Hrs) *	2.1	110 %
<5%	Failure Rate (AFR%) **	<=3%	40 %
10^(-14)	UBER **	10^ (-16)	16 %
11.4 GB/W	Power Efficiency	570 GB/W	5,000 %
43.1 IOPS/W	Performance/Power	42,850 IOPS/W	100,000 %
6.8 Watts	Idling Power	0.5 Watts	93 %
10.1 Watts	Load Power	0.9 Watts	91 %
1.0 GB/in3	Storage Density	16 GB/in3	1600 %
4.2 IOPS/in3	Performance Density	1,250 IOPS/in3	30,000 %
	Shock/Vibration/Noise		800/1600%/30dBLess
	Weight		50 %
	Maintenance/Op.Time #		50 %

#Reduced -Booting Up, -Virus Scan, -Defrag, -RAID Build, -Patching, -Data Restoration

** JEDEC's Mfr's Required Specs

Source: IMEX Research SSD Industry Report ©2011

Improving DB Query Response Time

IMEX RESEARCH.COM

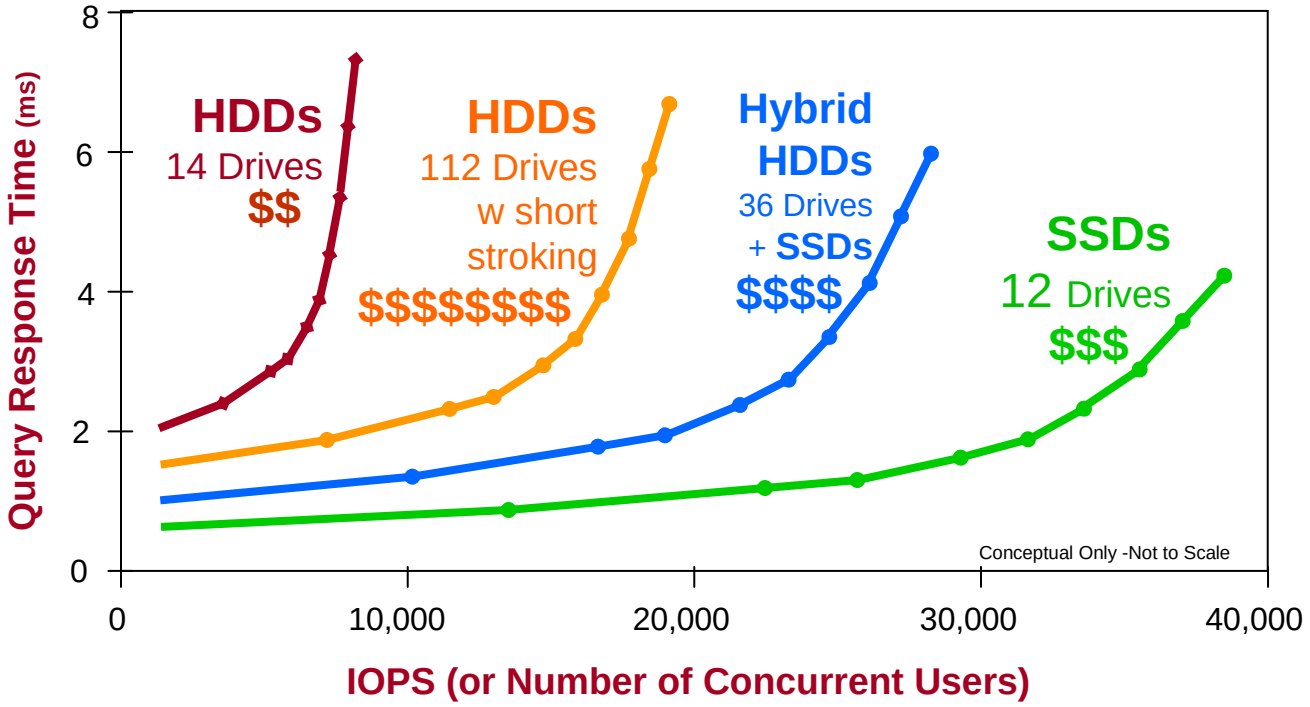

For a targeted query response time in DB & OLTP applications, many more concurrent users can be added cost-effectively when using SSDs or SSD + HDDs storage vs. adding more HDDs or short-stroking HDDs

Source: IMEX Research SSD Industry Report ©2011

Buying SSD vs. Memory to Improve TPS

IMEX RESEARCH.COM

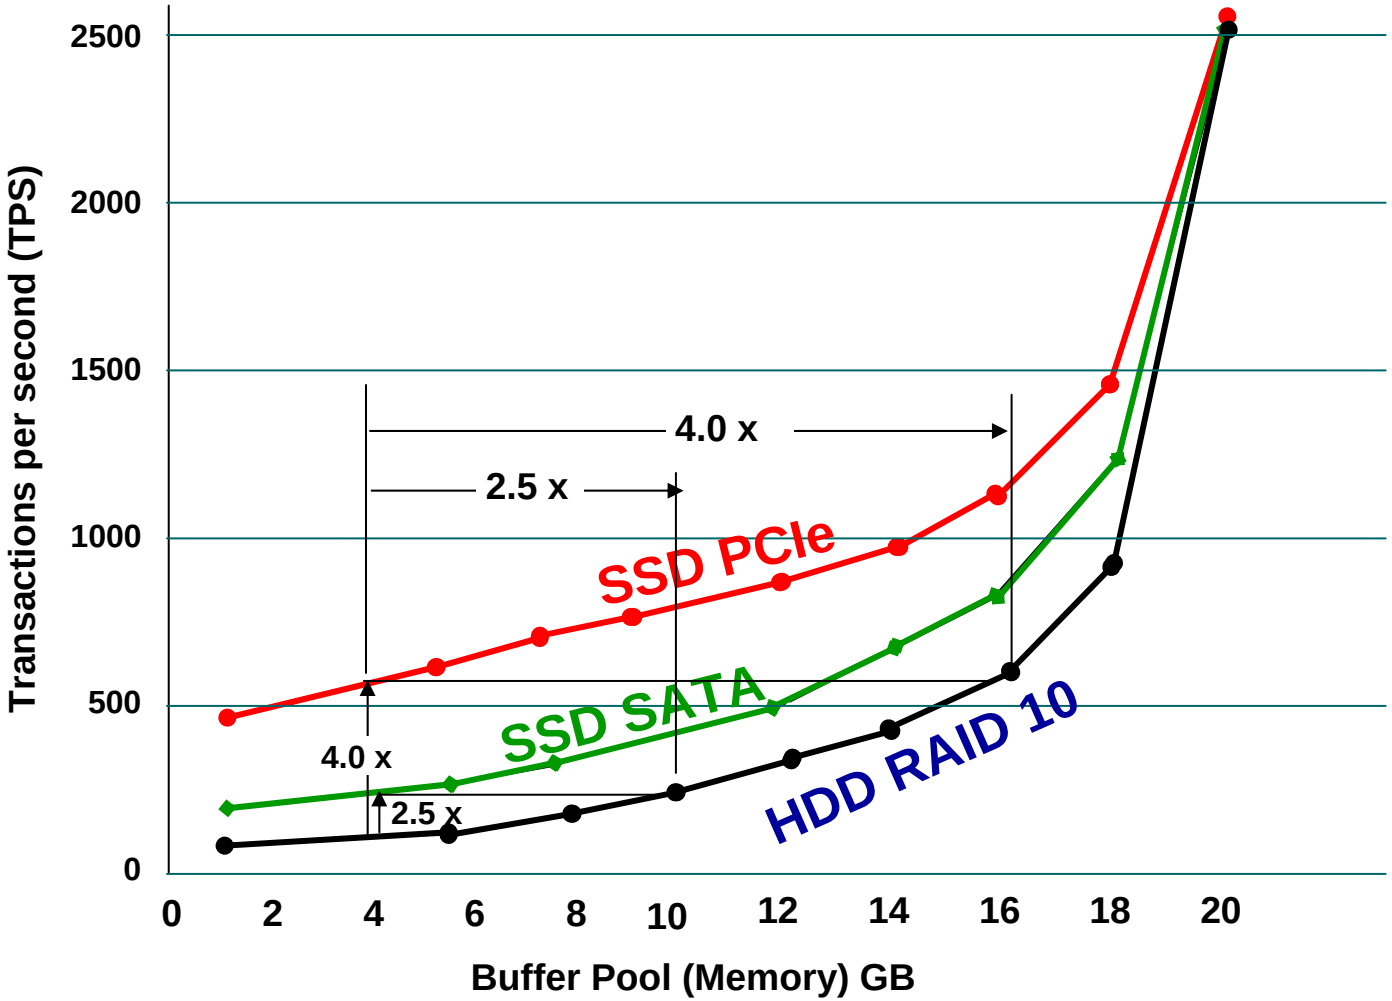

To achieve a certain TPS improvement, its cheaper to deploy SSD vs increased buffer memory (in GB costs) needed with using HDDs alone.

Source: IMEX Research SSD Industry Report ©2011

Market Segments by Apps/Workloads

IMEX RESEARCH.COM

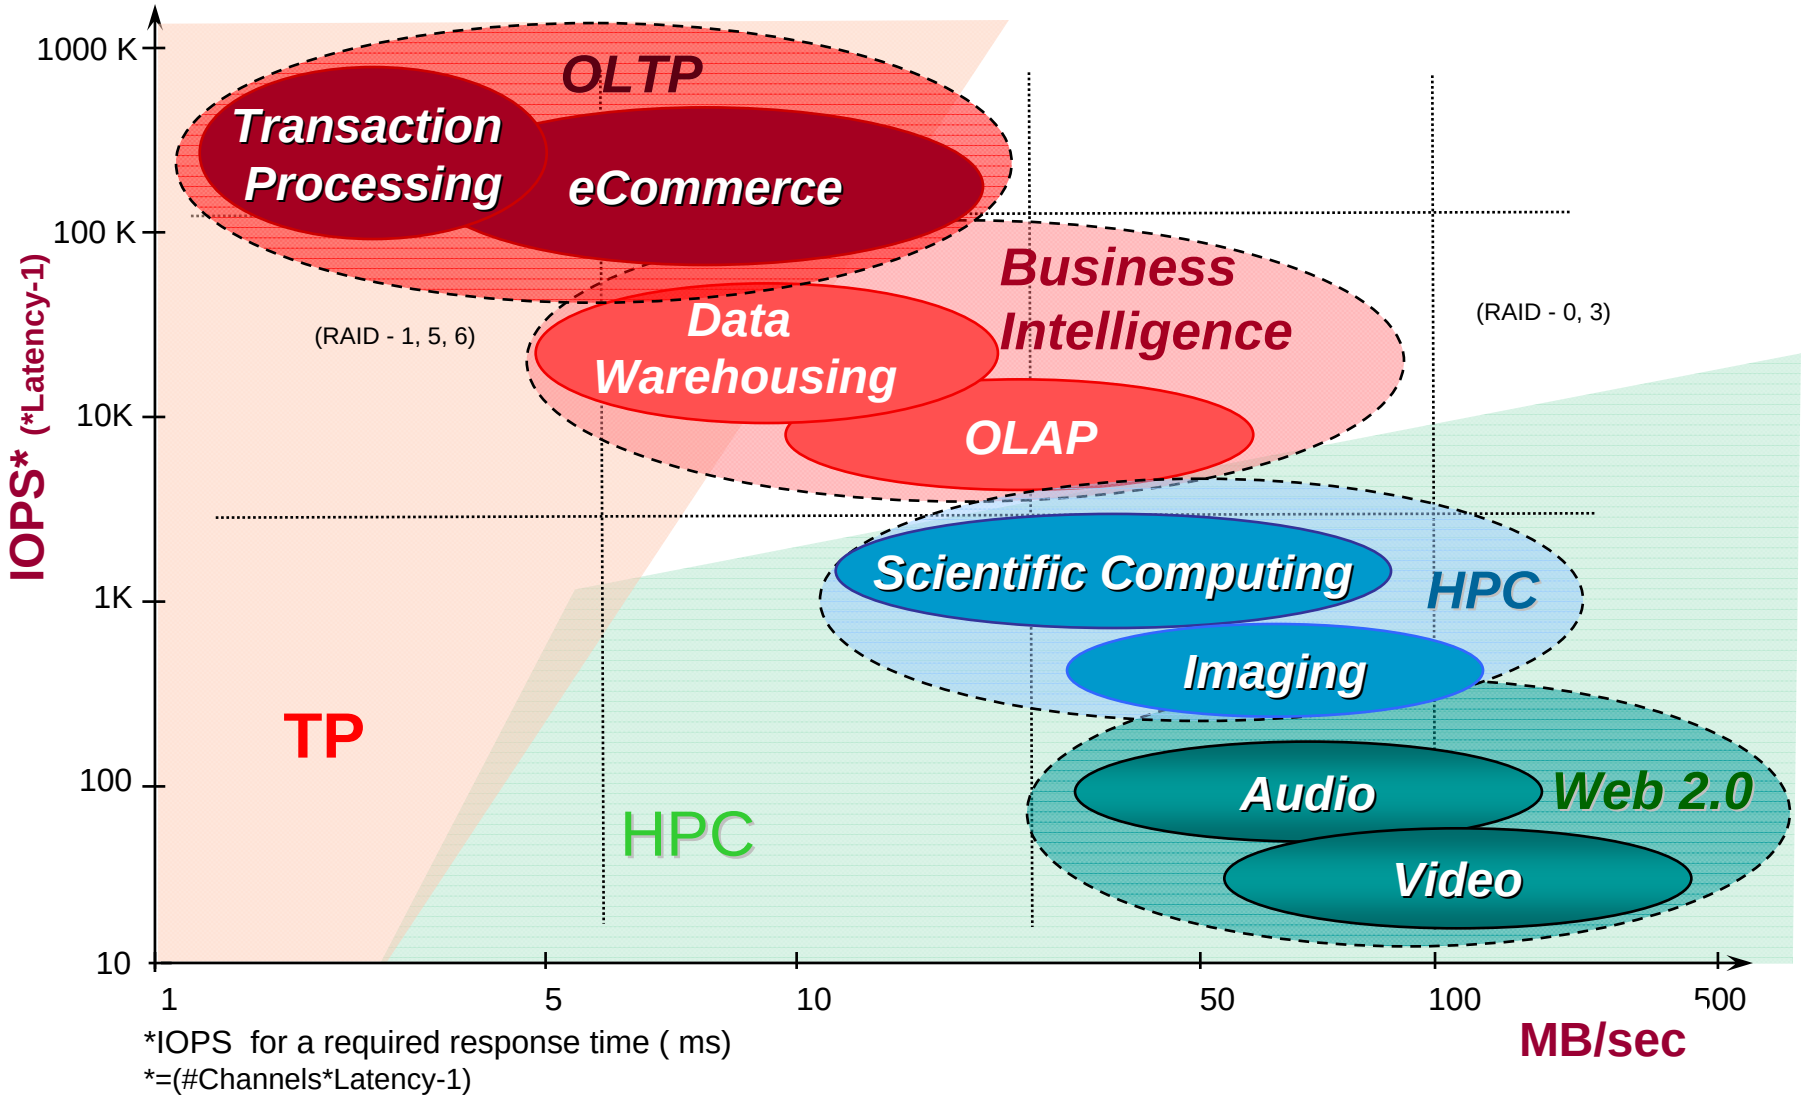

Source:: IMEX Research - Cloud Infrastructure Report ©2009-11

Data Storage Usage – In DCs & Cloud ©

IMEX RESEARCH.COM

I/O Access Frequency vs. Percent of Corporate Data

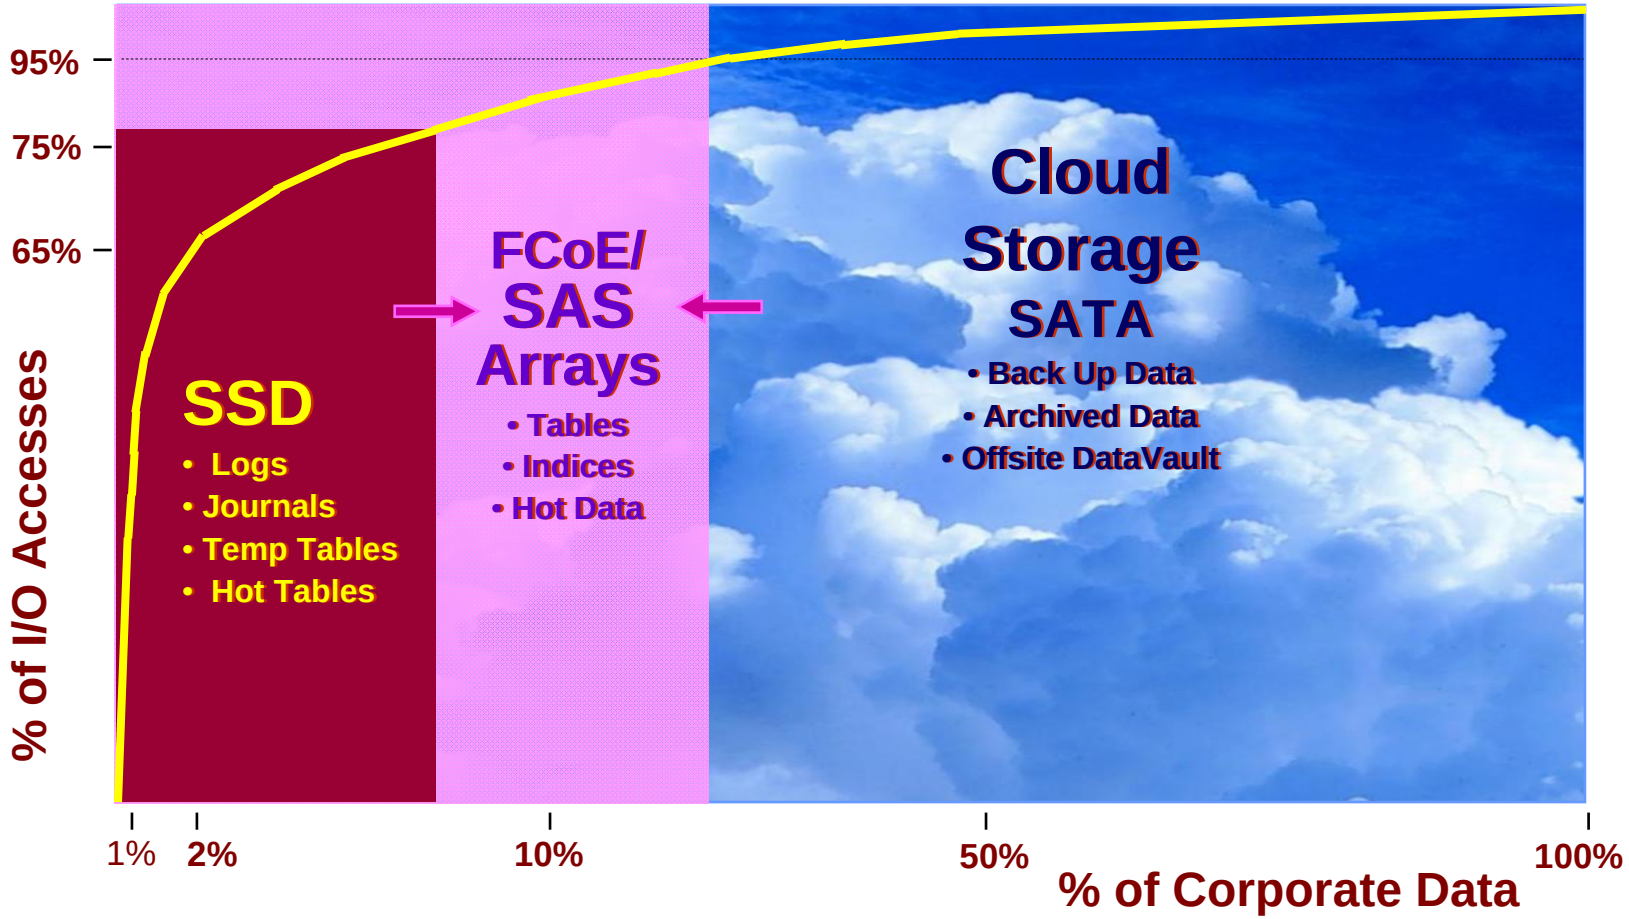

Source:: IMEX Research - Cloud Infrastructure Report ©2009-11

Applications Best Suited for SSDs

IMEX RESEARCH.COM

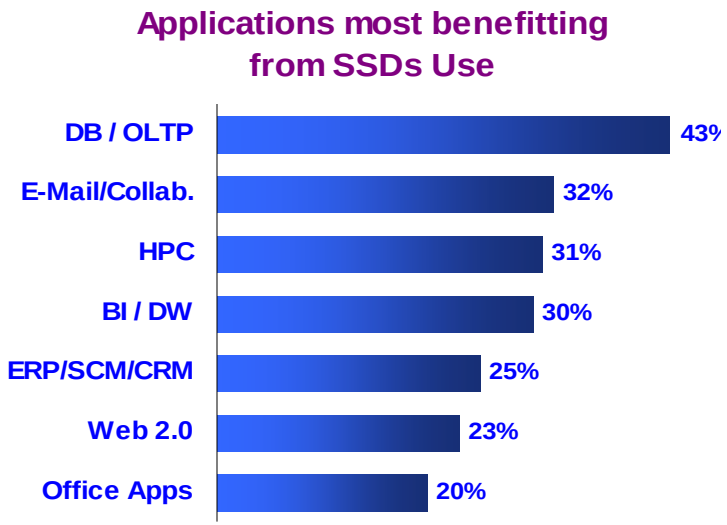

Apps and impact from SSD Usage

- **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

Applications best suited for SSDs: Workloads Characterization

IMEX RESEARCH.COM

Storage performance, management and costs are big issues in running Databases

- **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

Applications Best Suited for SSDs: Data Warehouse/BI

IMEX RESEARCH.COM

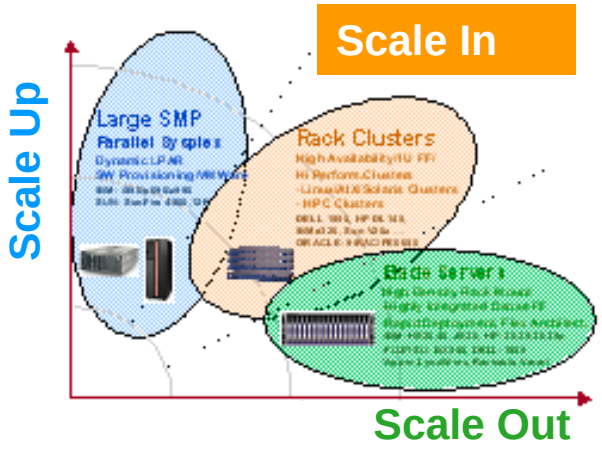

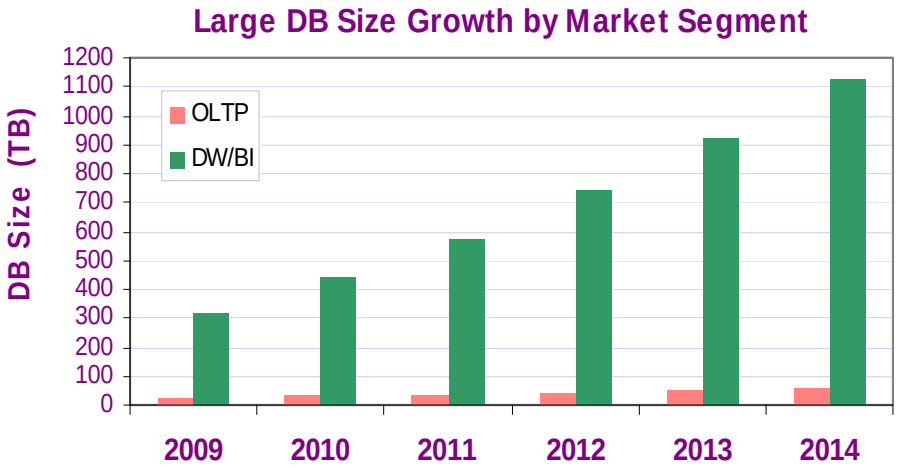

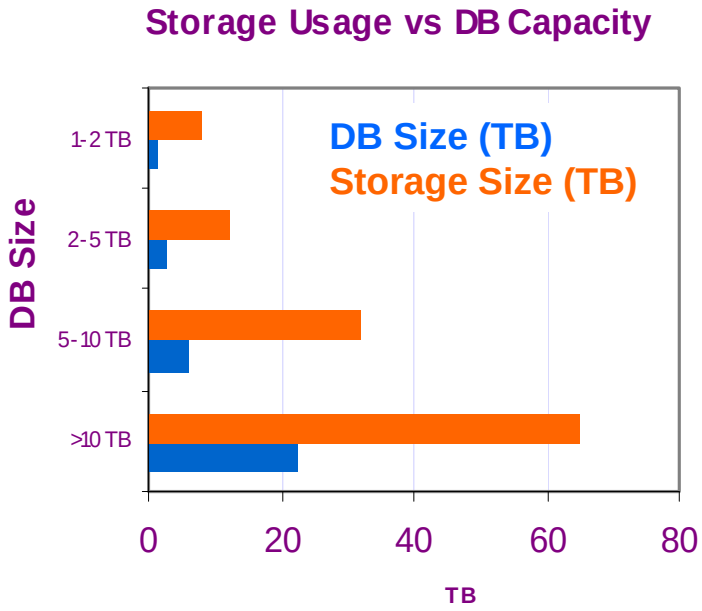

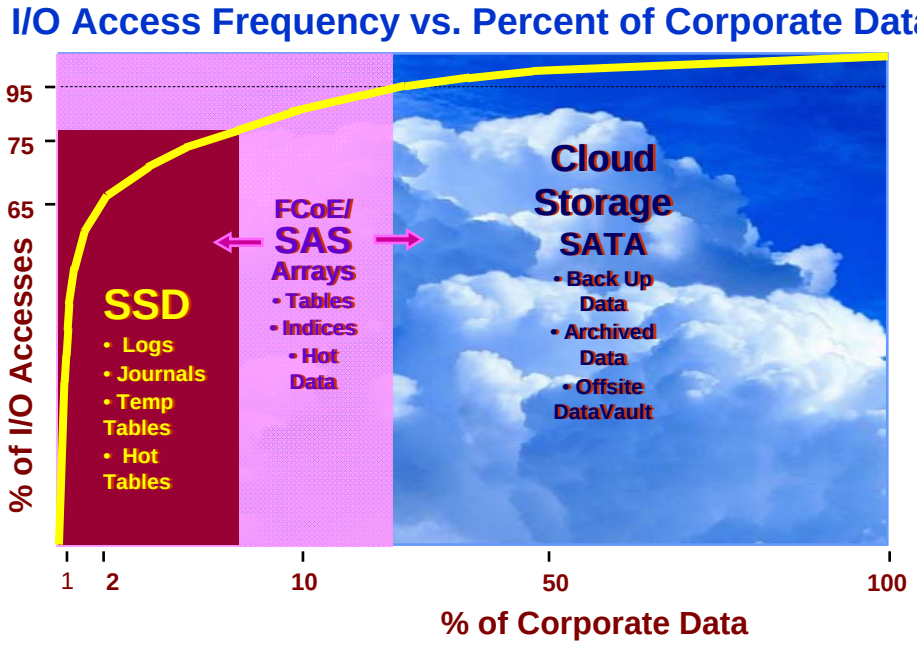

Data Source: IMEX Research
Cloud Infrastructure Report ©2009-11

Apps Best Suited for SSDs:
HPC/Web 2.0

IMEX RESEARCH.COM

VDI

Instant On Boot Ups
Rugged, Low Power

1GB/s, __ms

Commercial Visualization

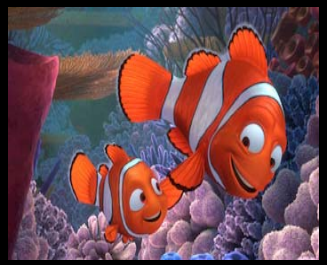

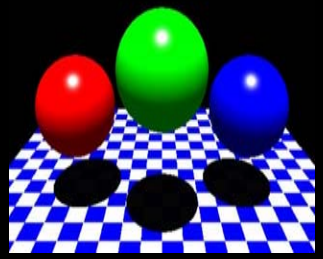

Rendering (Texture & Polygons)
Very Read Intensive, Small Block I/O

10 GB/s, __ms

Bioinformatics & Diagnostics

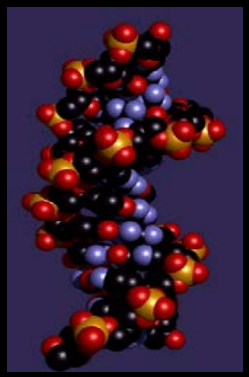

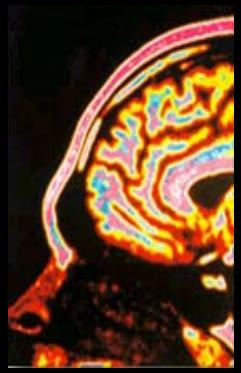

Decision Support Bus. Intelligence

Data Warehousing
Random IO, High OLTPM

1GB/s, __ms

Entertainment- VoD / U-Tube

Most Accessed Videos
Very Read Intensive

4 GB/s, __ms

Data: IMEX Research & Panasas

Automated Storage Tiering
- Next Frontier in Storage Efficiency

IMEX RESEARCH.COM

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

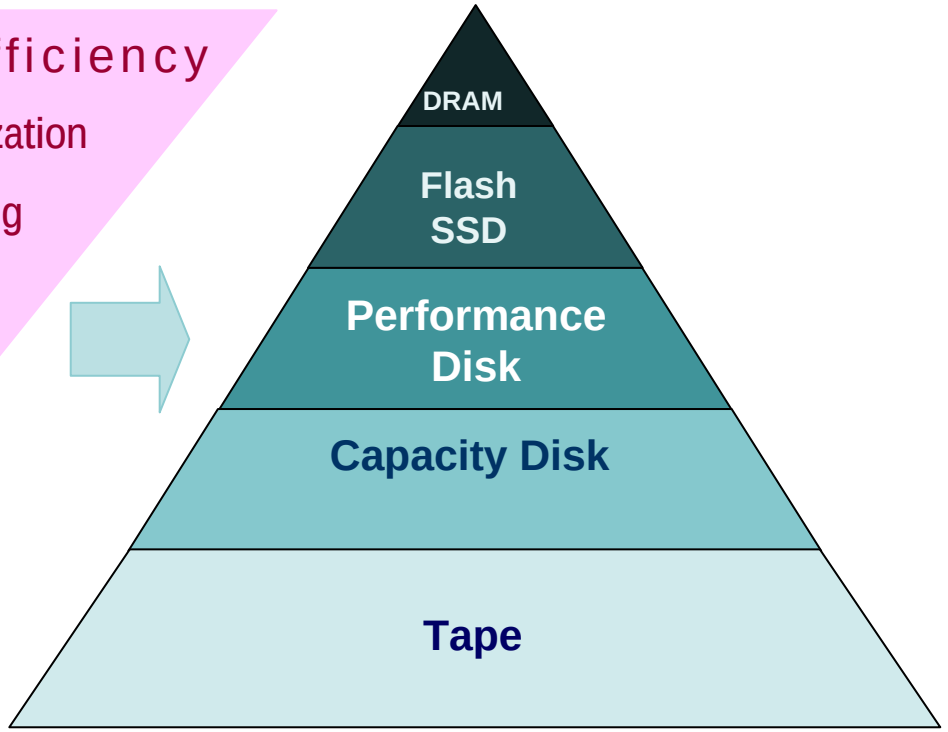

Auto Tiering

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

Automated Storage Tiering:
The Killer App for Enterprise SSDs

IMEX RESEARCH.COM

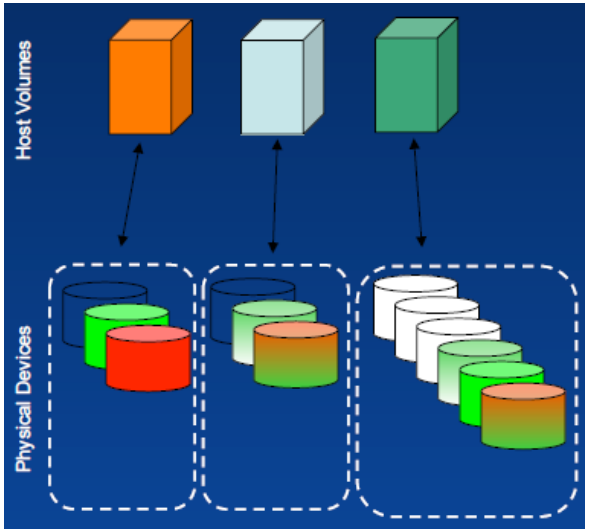

- **Traditional Disk Mapping**
  - Volumes have different characteristics. Applications need to place them on correct tiers of storage based on usage

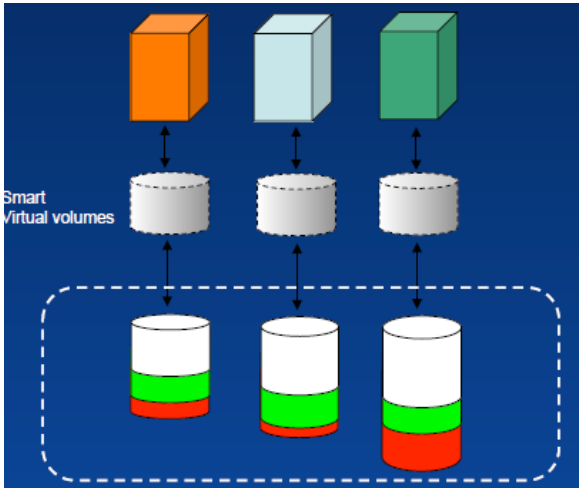

- **Smart Storage Mapping**
  - All volumes appear to be "logically" homogenous to apps. But data is placed at the right tier of storage based on its usage through smart data placement and migration

Source: IMEX Research SSD Industry Report ©2011

Automated Storage Tiering:
Workload I/O Monitoring & Migration

IMEX RESEARCH.COM

Storage-Tiered Virtualization

Storage-Tiering at LBA/Sub-LUN Level

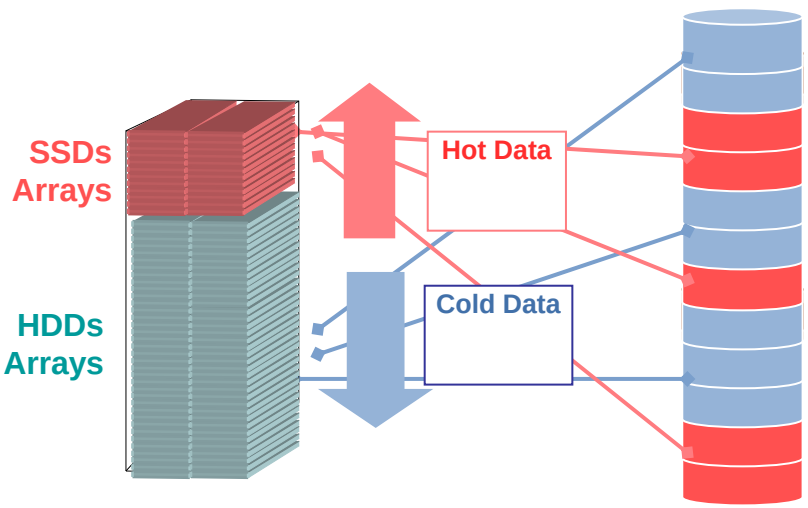

- **Automated Storage Tiering**
  - Continuously monitor and analyze data access on the tiers
  - Automatically elevate hot data to “Hot Tiers” and demote cool data/volumes to “Lower Tiers
  - Allocate and relocate volumes on each tier based on use
  - Automated Migration reduces OPEX to otherwise SANs managed manually

Source: IMEX Research SSD Industry Report ©2011

Automated Storage Tiering: Workload I/O Monitoring & Migration

IMEX RESEARCH.COM

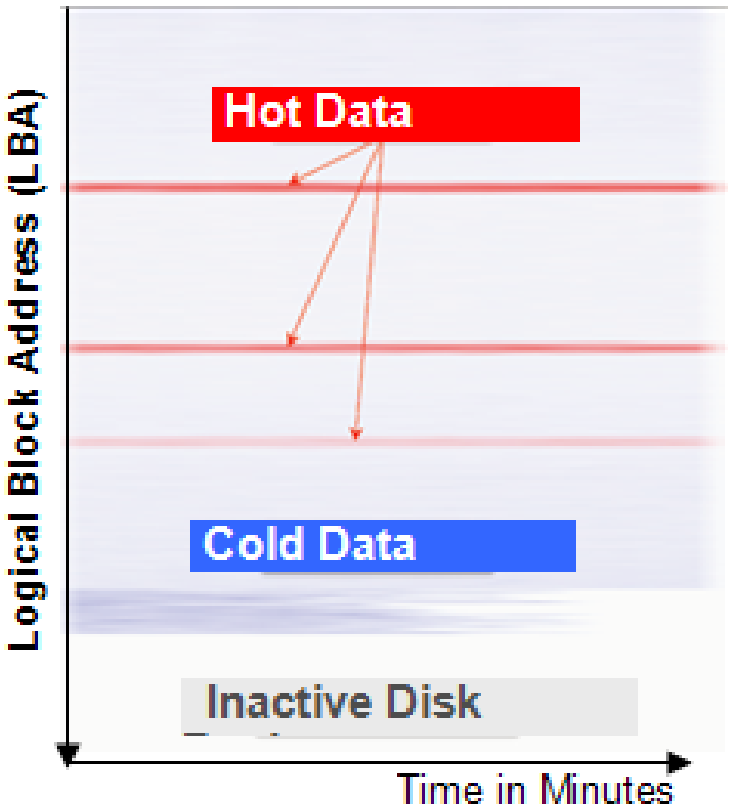

- **LBA Monitoring and Tiered Placement**
  - Every workload has unique I/O access signature
  - Historical performance data for a LUN can identify performance skews & hot data regions by LBAs
  - Using Smart Tiering identify hot LBA regions and non-disruptivelymigrate hot data from HDD to SSDs.
  -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

Source: IMEX Research SSD Industry Report ©2011

Automated Storage Tiering: Improving Response Time

IMEX RESEARCH.COM

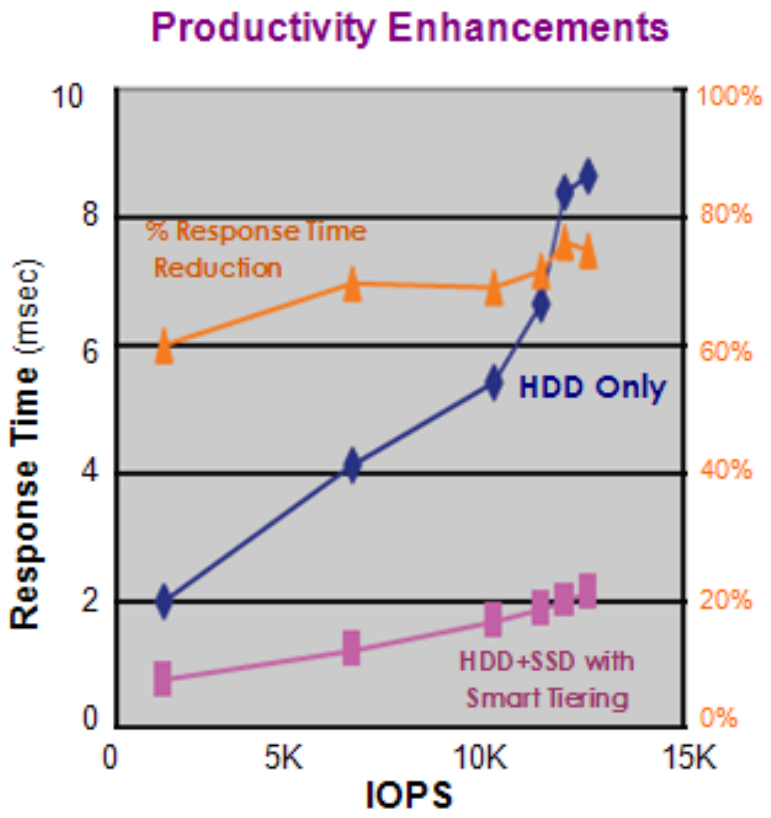

- **Productivity Improvements (Response Time)**
  - With automated reallocation of hot spot data (~ 5-10% of total data) to SSDs, performance improvements
    - **Response time reduction** ~70+% **IOPS increase of 200%** for any I/O intensive workloads in Time-Perishable OLTP markets: Airlines Reservations, 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

Automated Storage Tiering: Enhancing Database Throughput

IMEX RESEARCH.COM

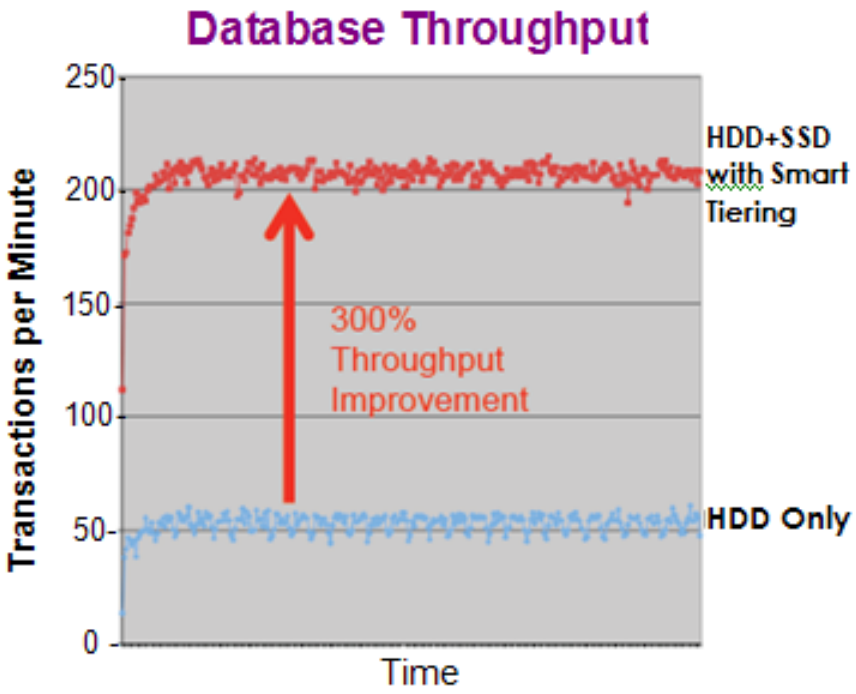

- **DB Throughput Optimization**
  - Every workload has unique I/O access signature and historical behavior
  - identify hot “database objects” and smartly placed in the right tier.
  - Scalable Throughput Improvement - 300%
  - Substantial IO Bound Transaction Response time Improvement - 45%-75%

Source: IMEX Research SSD Industry Report ©2011

Automated Storage Tiering:
Storage Tiering – Best Practices

IMEX RESEARCH.COM

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

SSS in DB Environments: Enabling new System Architectures

IMEX RESEARCH.COM

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**
- **System Test HW:**
  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

SSS in DB Environments:
DB Improvements with Flash SSDs

IMEX RESEARCH.COM

- **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

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

IMEX RESEARCH.COM

Goals & Implementation

- 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
- Use **Partitioning Software to Classify Data**
  - 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 Performance 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

SSS in DB Environments: Enabling new Economics

IMEX RESEARCH.COM

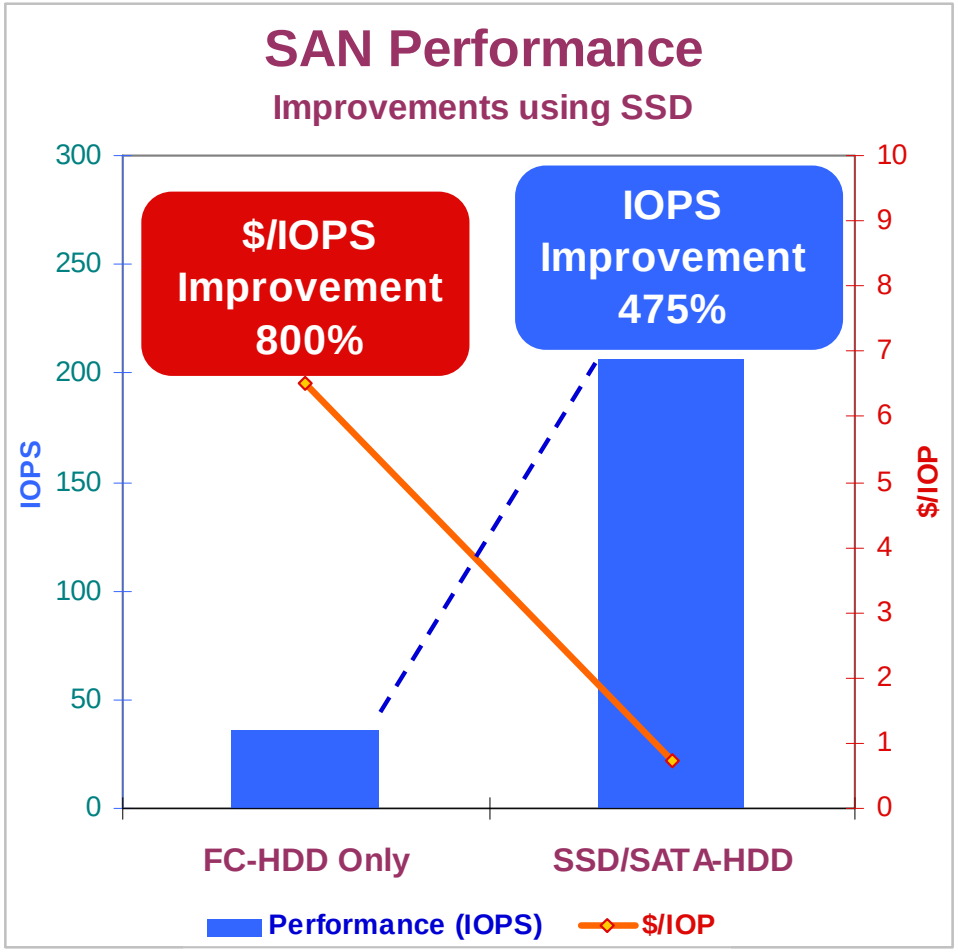

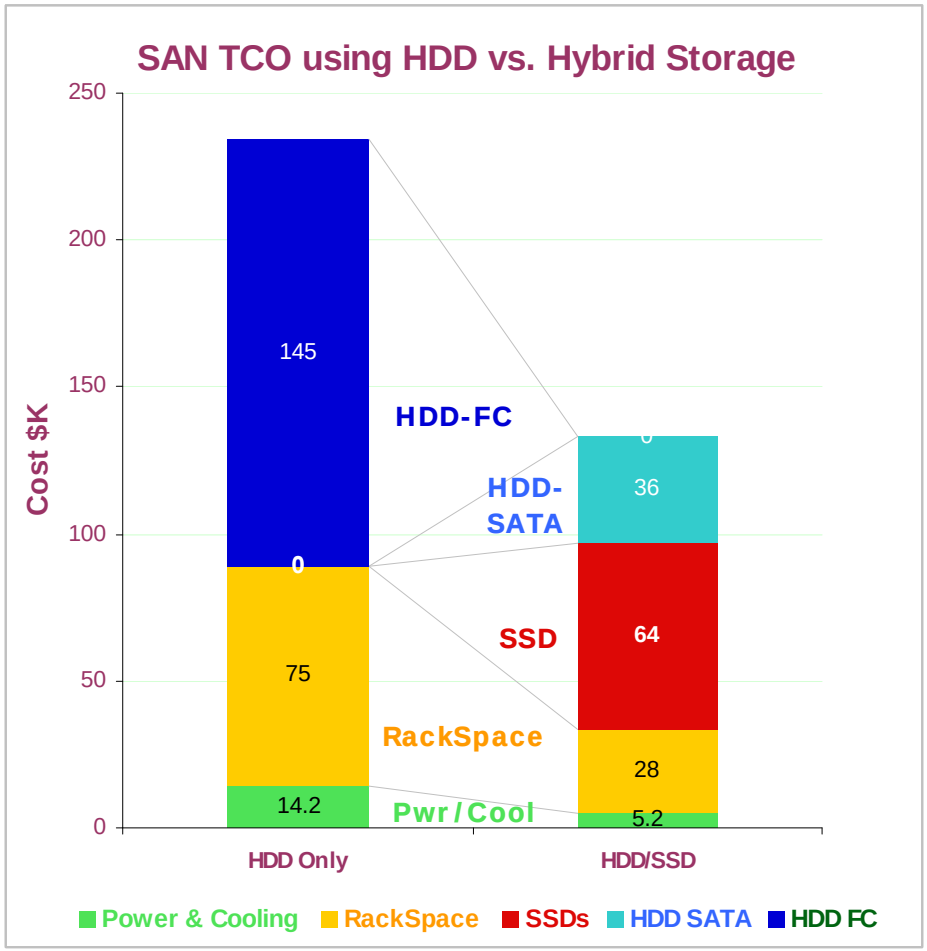

Source: IMEX Research SSD Industry Report ©2011

Tiered Storage – Industry Ecosystem

IMEX RESEARCH.COM

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

SSD Challenges & Solutions: Goals & Best Practices

IMEX RESEARCH.COM

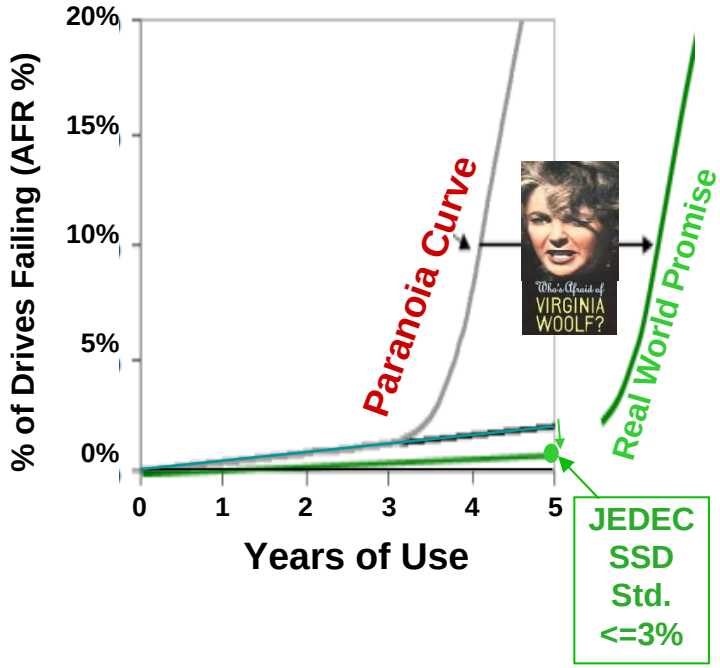

Concerned about SSD Adoption in your Enterprise ?

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 2010-11

New Intelligent Controllers: Meeting Enterprise Requirements

IMEX RESEARCH.COM

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.

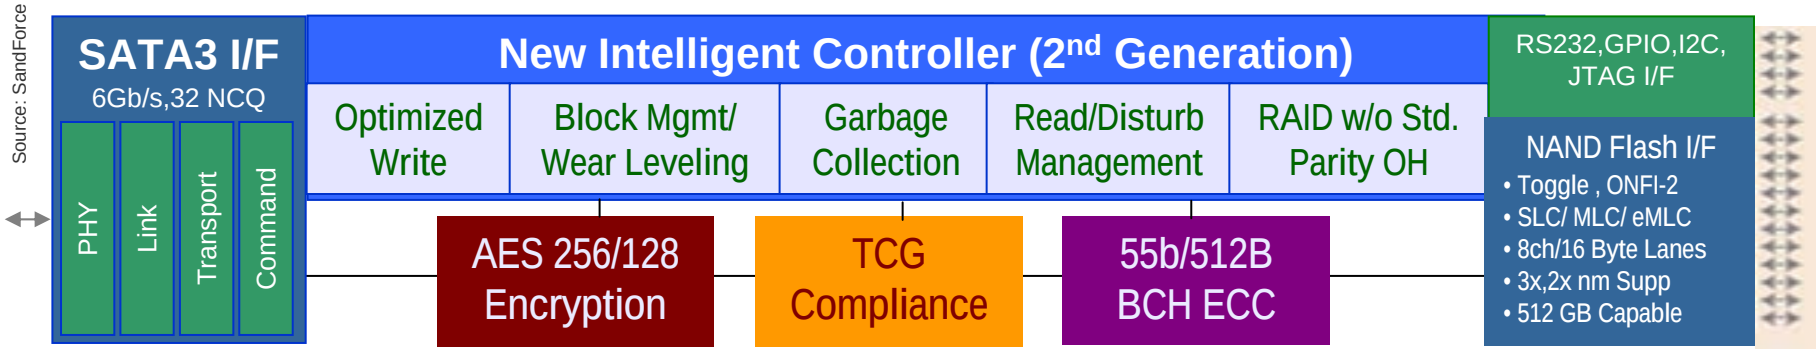

New Gen Controllers allow SSDs to meet Enterprise Class Availability/Performance/ over 5-Year Life/Scalability/ Auto-Configuration & Auto Data-Tiering

Source: IMEX Research SSD Industry Report ©2011

Key Takeaways

IMEX RESEARCH.COM

- **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