



Education

# **Solving the I/O Bottleneck in NextGen Data Centers & Cloud Computing**

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

- 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.**

## ◆ Solving the I/O Bottleneck in NextGen Data Centers & Cloud Computing

Virtualization brings tremendous advantages to data centers of all sizes – large and small. But the rapid proliferation of virtual machines (VMs) per physical server, in its wake, creates a highly randomized I/O problem, raising performance bottlenecks. To address this I/O issue, the IT industry and storage administrators in particular, have begun the adoption of a slew of newer technology solutions - ranging from faster and larger memories in cache, SSDs for higher IOPS cost effectively, high bandwidth (10GbE) networks using NPIVs for virtualized networks between VMs and shared storage systems along with embedded intelligence software to optimize various VM workloads – all in order to meet various SLA metrics of performance, availability, cost etc..

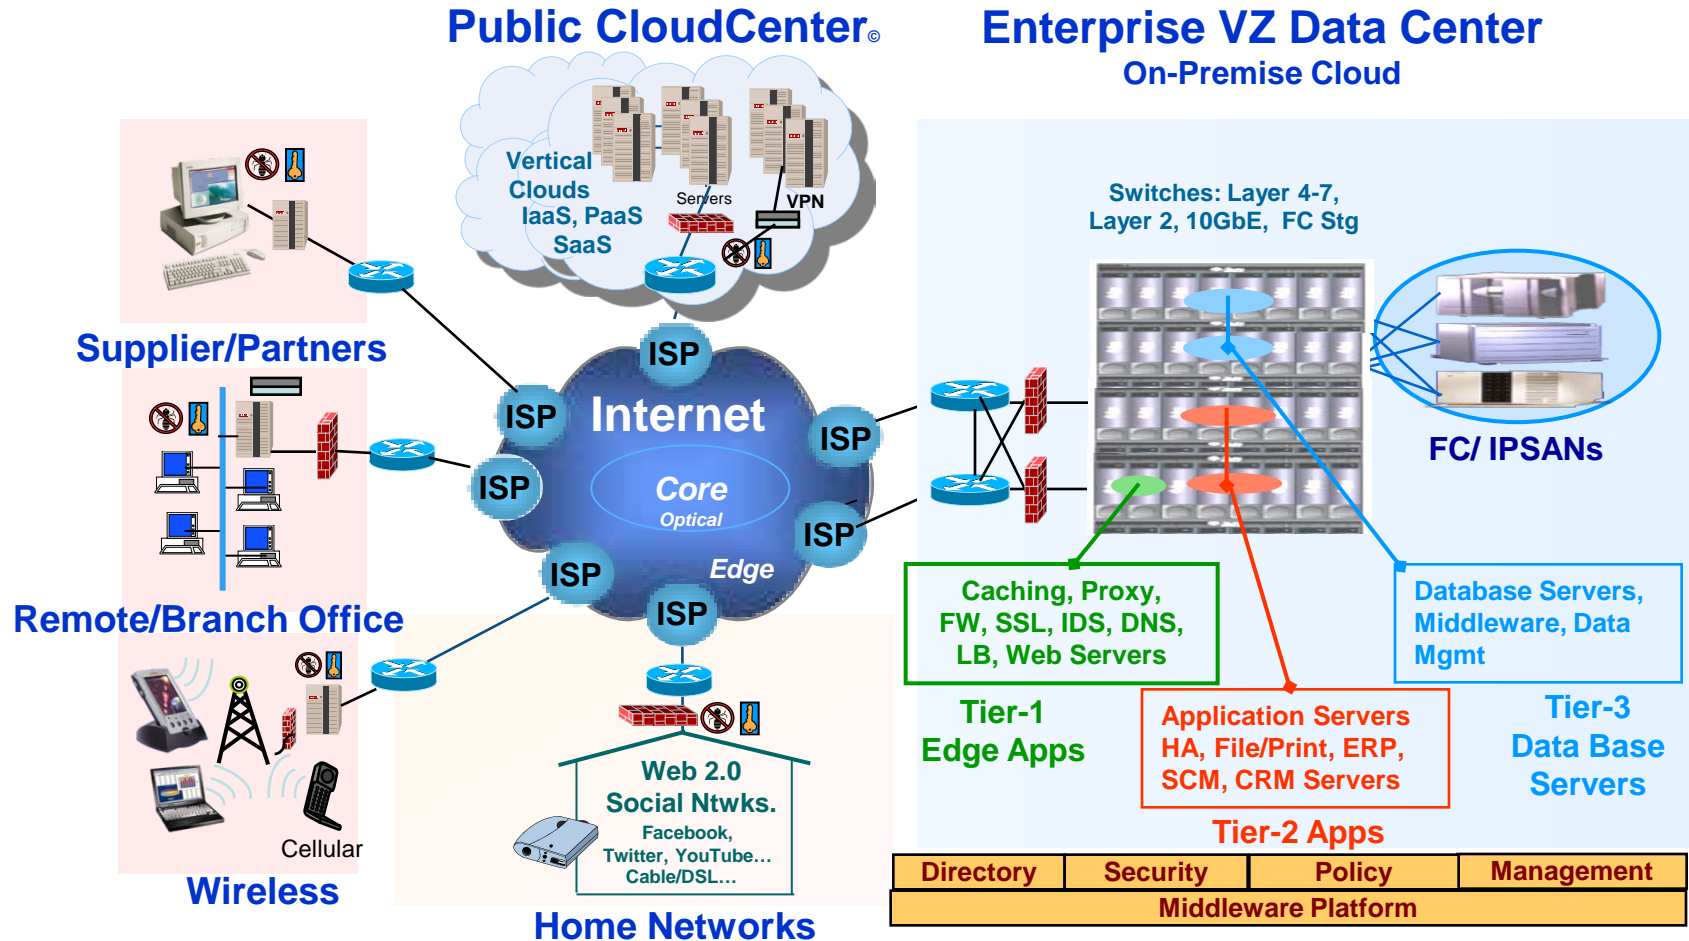
### **Learning Objectives:**

This presentation provides a simple and illustrative view of the impact of Server Virtualization on existing storage I/O solutions and best practices. It delineates the roles, capabilities and cost effectiveness of emerging technologies in mitigating the I/O bottlenecks so the IT infrastructure implementers can achieve their targeted performance under various workloads, from their storage systems in virtualized data centers.

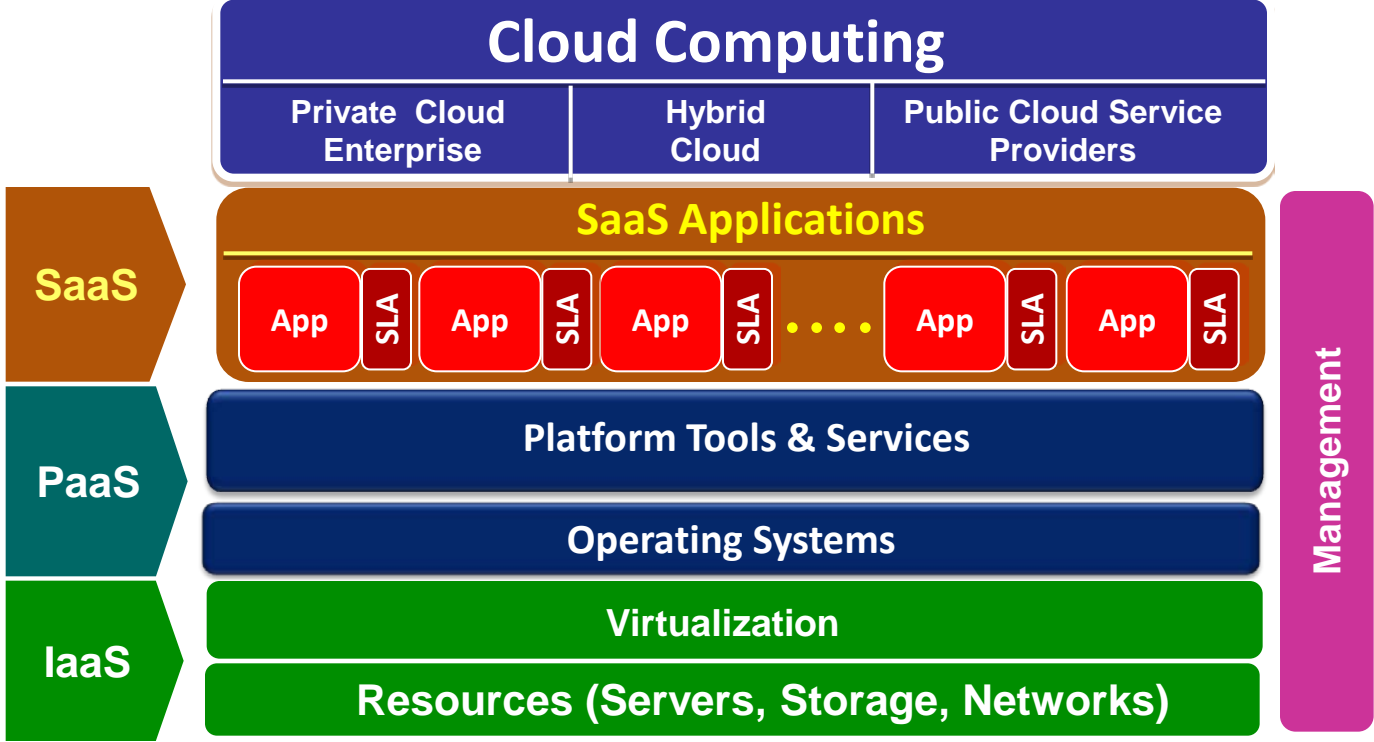
The audience will obtain a clear understanding of the factors that can create side-effects from Server Virtualization and become knowledgably better prepared to ask pertinent questions of their suppliers of IT Infrastructure Equipment, Storage Virtualization and Data Storage Management software as well as in their implementation to achieve targeted results in performance, availability, scalability, interoperability and data management in their virtualized data centers

- **Data Centers & Cloud Infrastructure**
  - ◆ Storage Usage in NGDC & Clouds
  - ◆ Performance Metrics by Workload
- **Data Center I/O Bottlenecks**
  - ◆ Anatomy of Data Access I/O Gaps
  - ◆ Role of SSD in Improving I/O Perf. Gap
- **Improving Query Response Time**
  - ◆ Apps Benefitting from Improved I/O
- **New Storage Economics**
  - ◆ SSD vs. HDD - IOPS/GB
- **I/O Forensics & Use in Auto Storage-Tiering**
- **Key Takeaways**



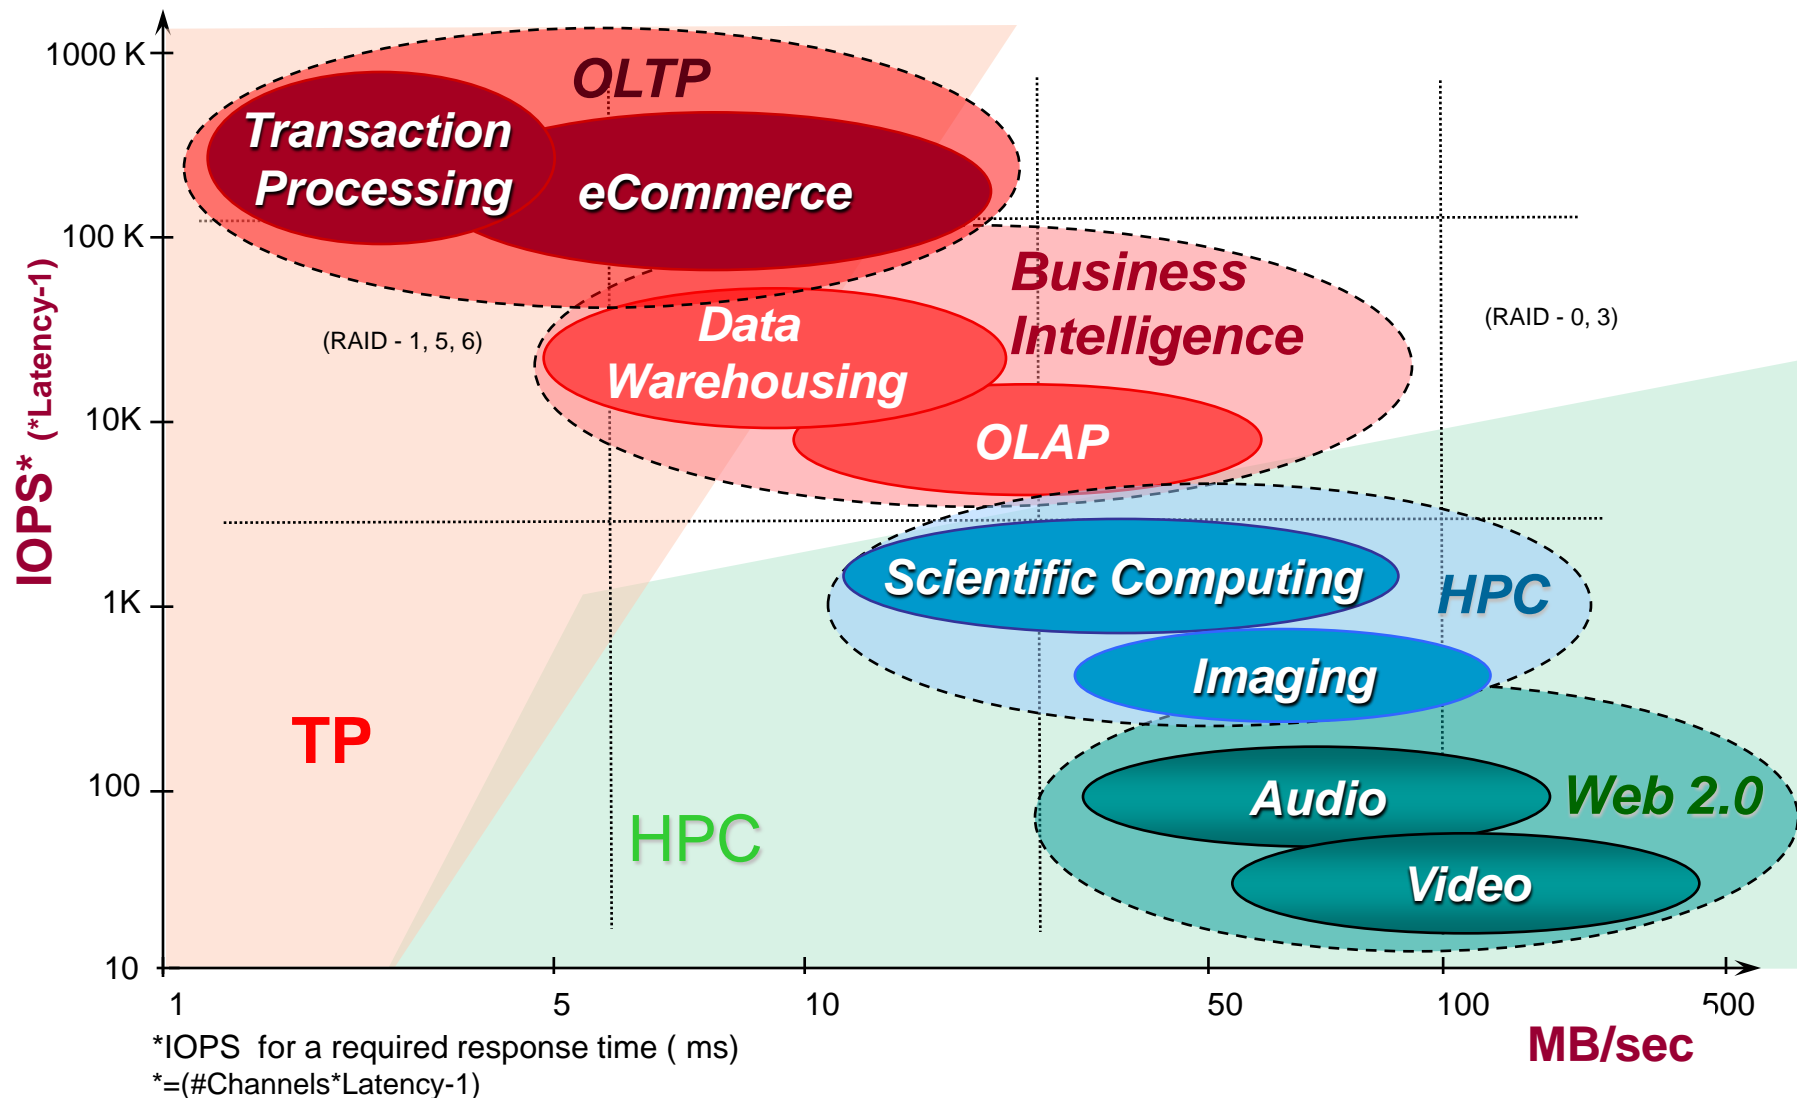


**A request for data from a remote client to an enterprise data center crosses a myriad of systems and devices. Key is understanding where potential I/O bottlenecks can happen**




**In Cloud Computing, Applications will carry SLAs as Metadata, dictating Availability, Performance, Cost, Security criterion and marshalling IaaS resources including Virtualized Tiered-Storage Elements (SSDs/HDDs/Tape/Data Placement & Mgmt. SW) required to mitigate any IO bottlenecks.**

# Performance Metrics by Workload



# 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 Zip	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 100/140 IOPS

# Data Center Performance Bottlenecks

## Applications

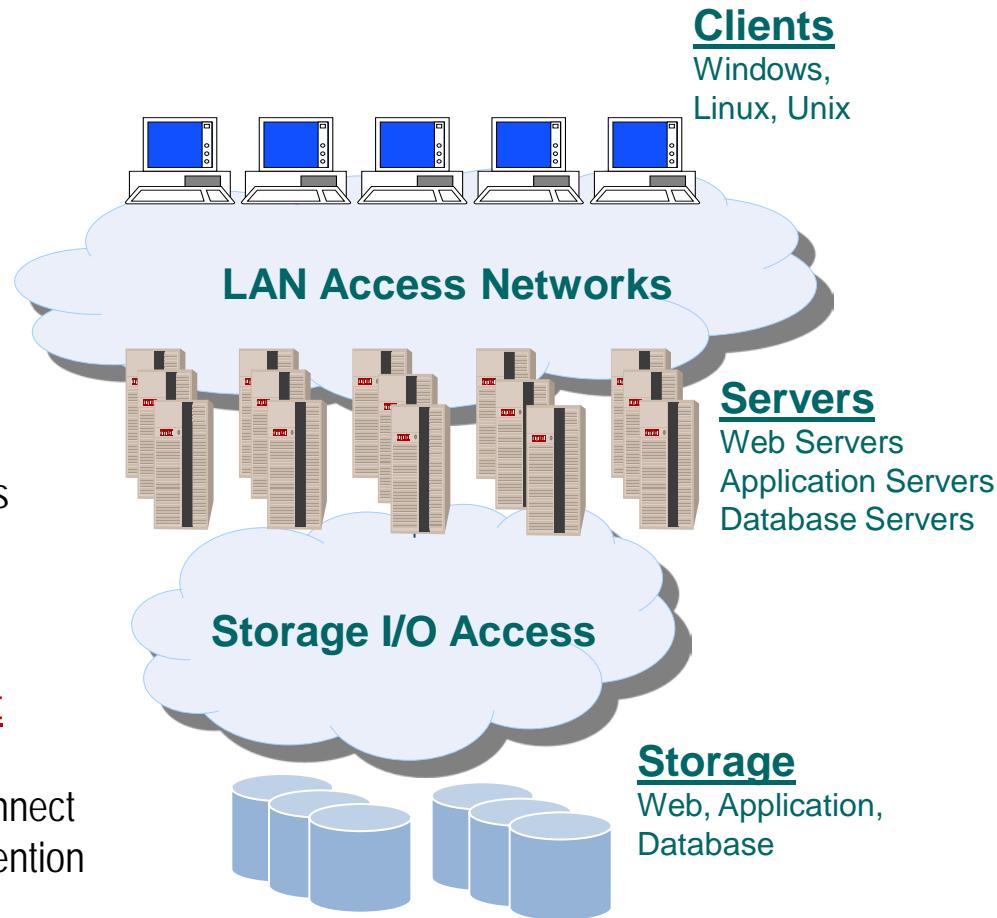
Excessive Locking  
Data Contention  
I/O Delays/Errors

## 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

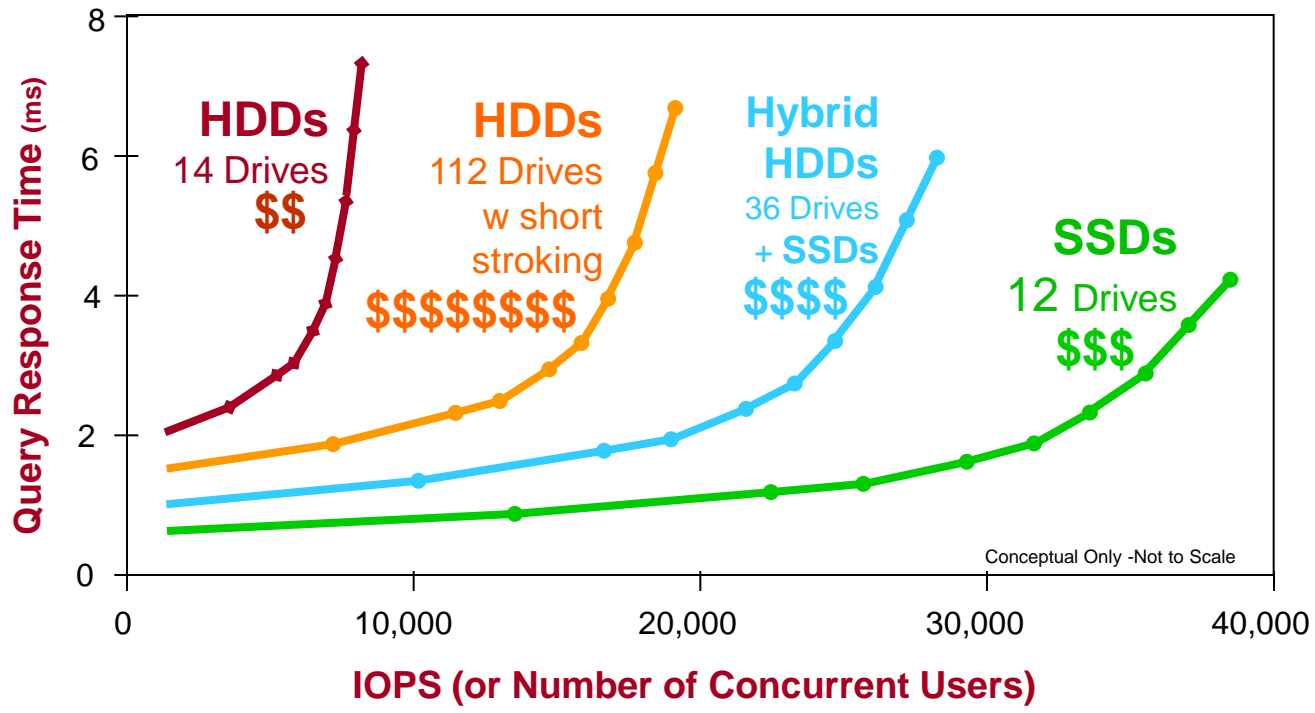
## Server Bottlenecks

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

## Device Bottlenecks

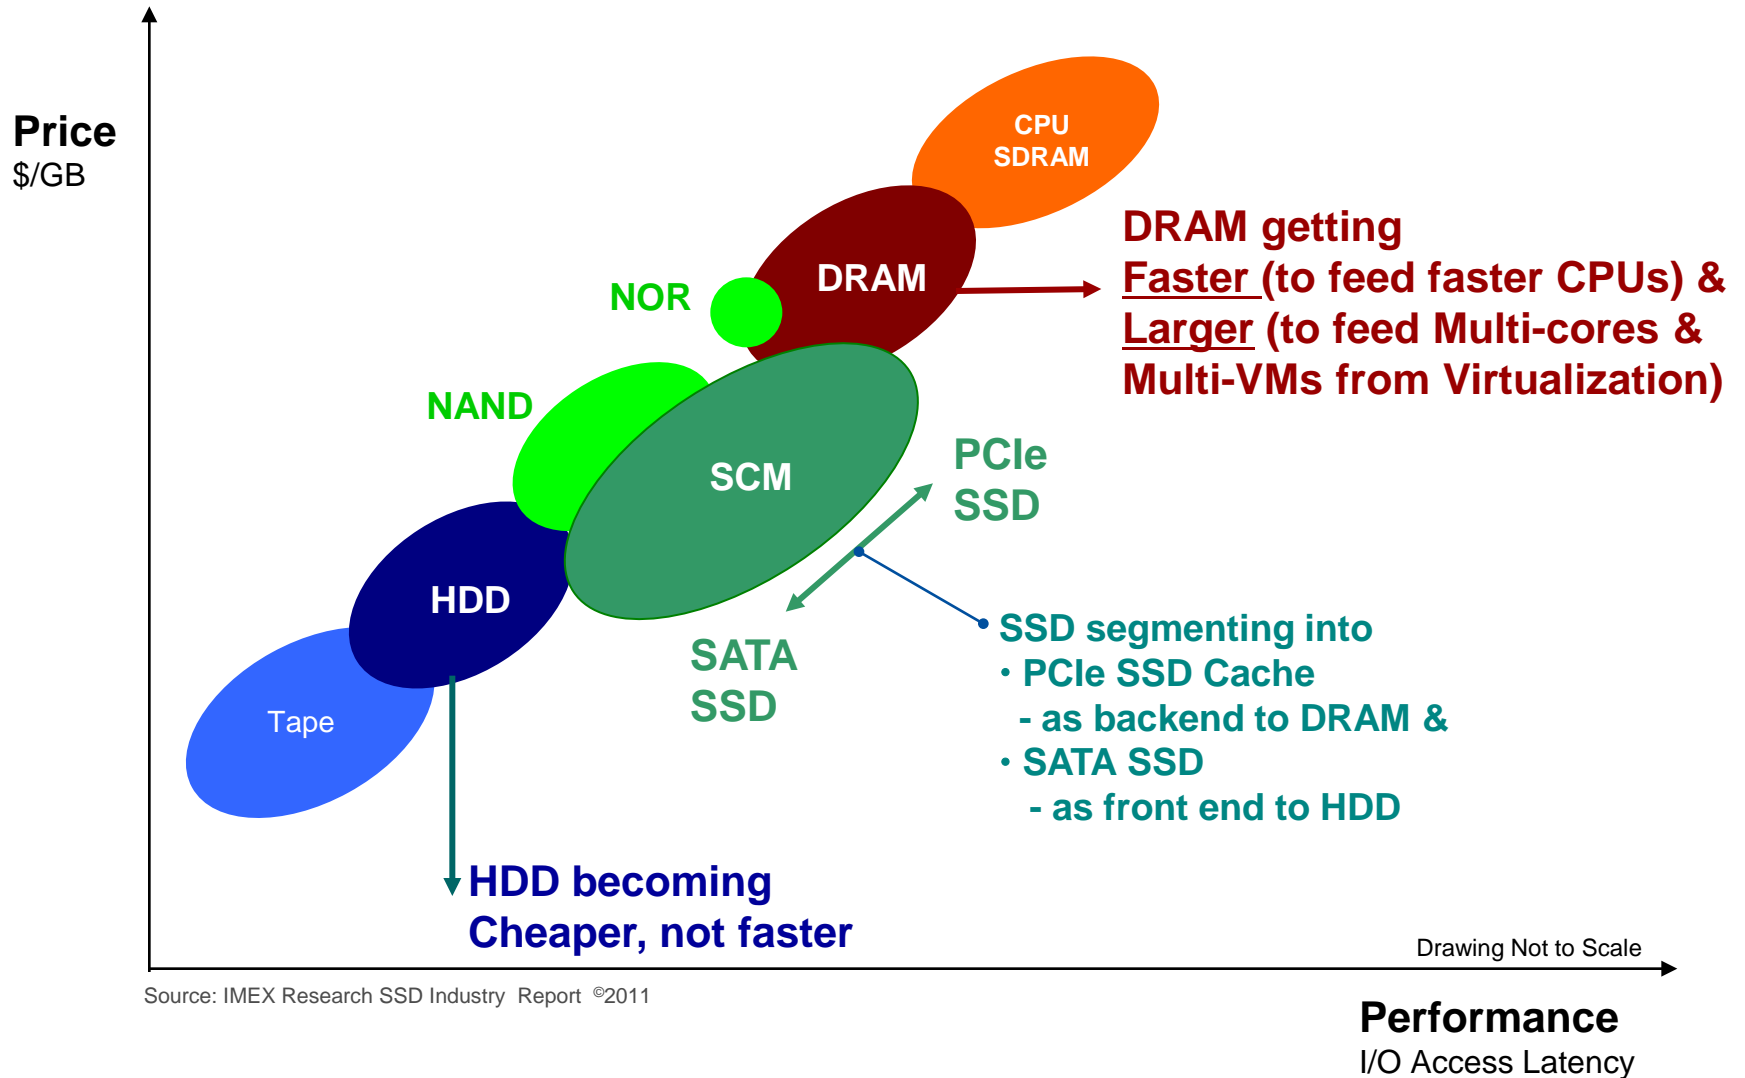
Device I/O Hotspots  
Cache Flush  
Lack of Storage Capacity

# Improving Query Response Time in OLTP



**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**

# Role of SSD in Improving I/O Perf. Gap



**Performance**  
I/O Access Latency

## ➤ **SCM (Storage Class Memory)**

Solid State Memory filling the gap between DRAMs & HDDs

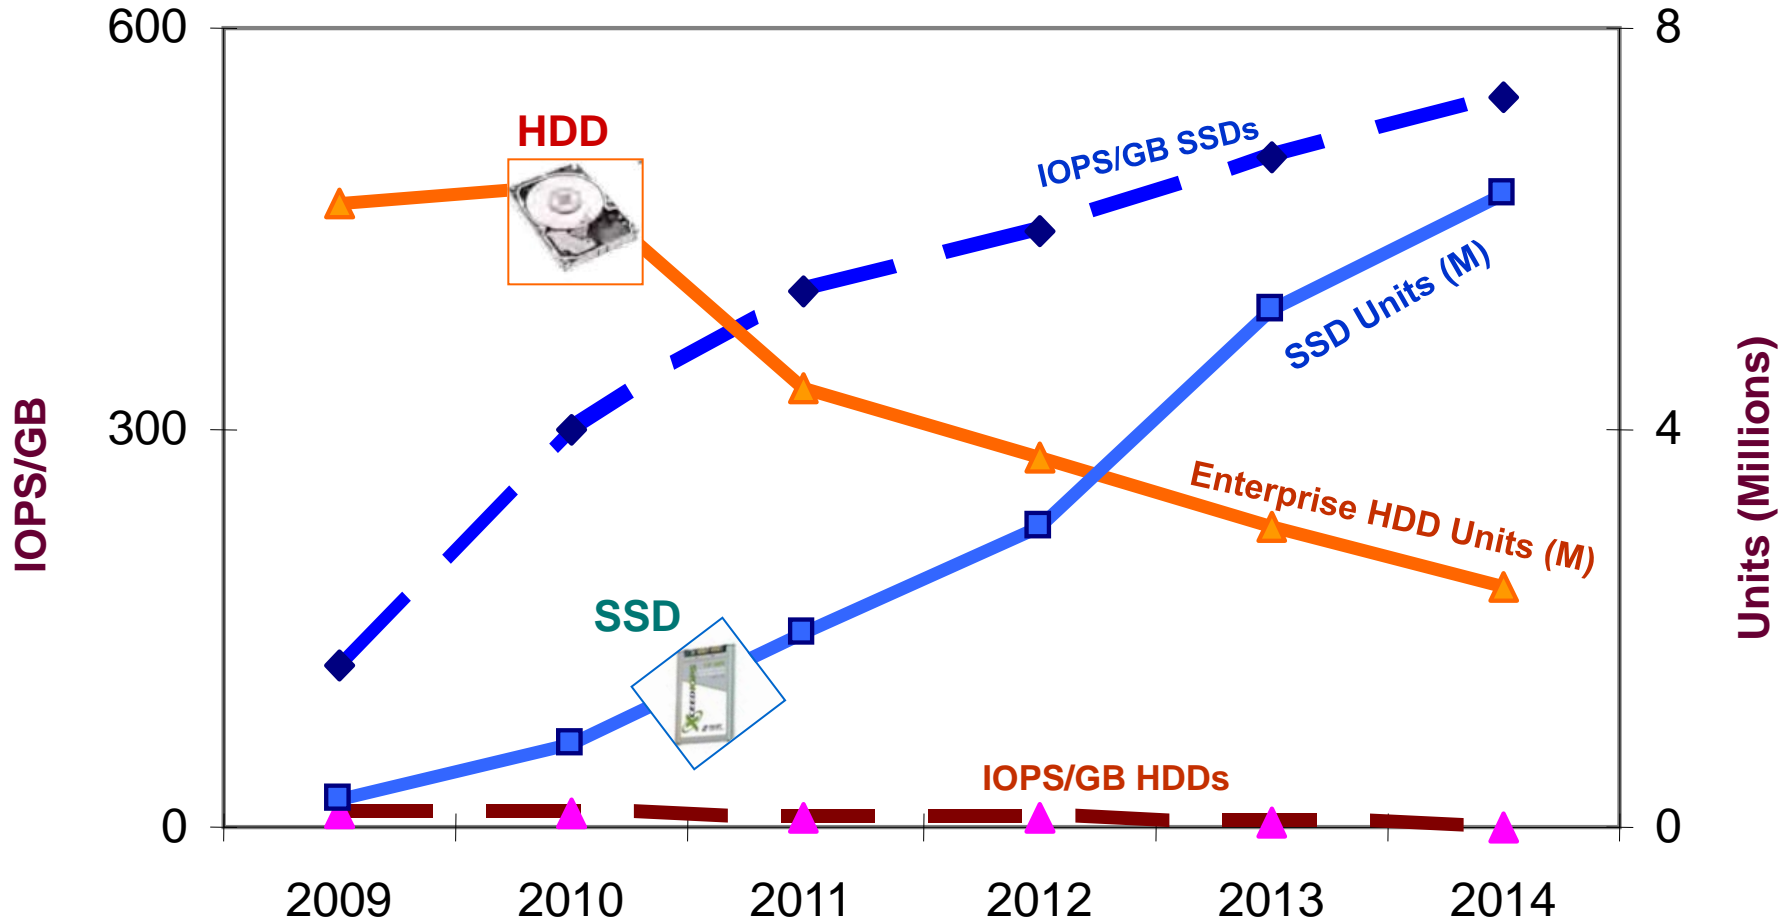
Marketplace segmenting SCMs into SATA and PCIe based SSDs

## ➤ **Key Metrics Required of Storage Class Memories**

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

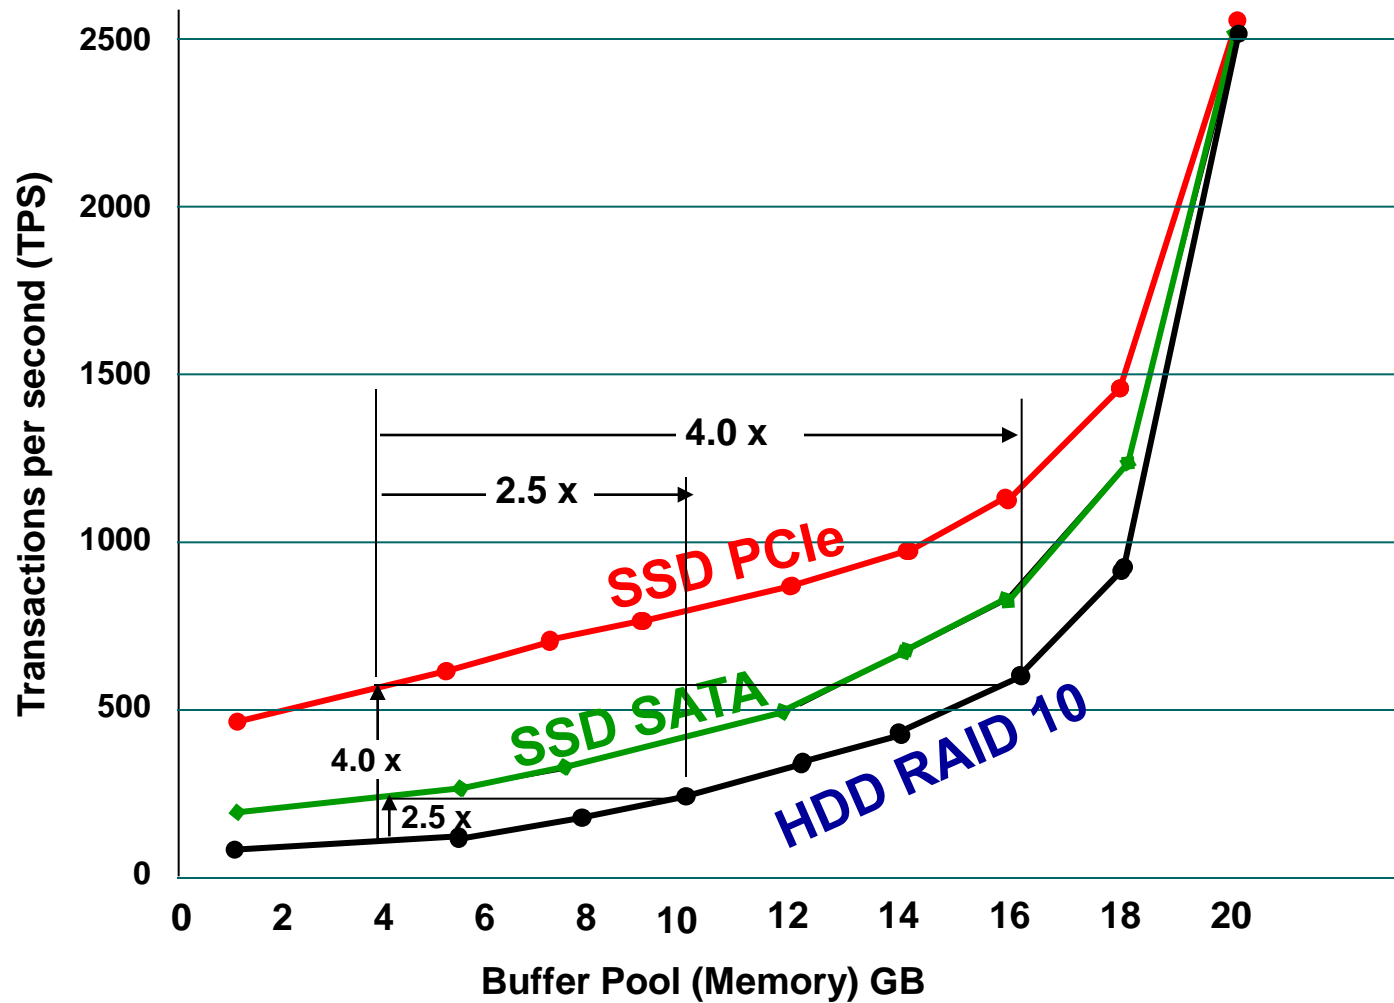
- **Performance** - Latency (Random/Block RW Access-ms); Bandwidth  
BW(R/W- GB/sec)
- **Data Integrity** - BER (Better than 1 in  $10^{17}$ )
- **Reliability** - Write Endurance (30K PE Cycles No. of writes before death);
  - - Data Retention (5 Years); MTBF (2 millions of Hrs),
- **Environment** -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

# SSDs - Price Erosion & IOPS/GB



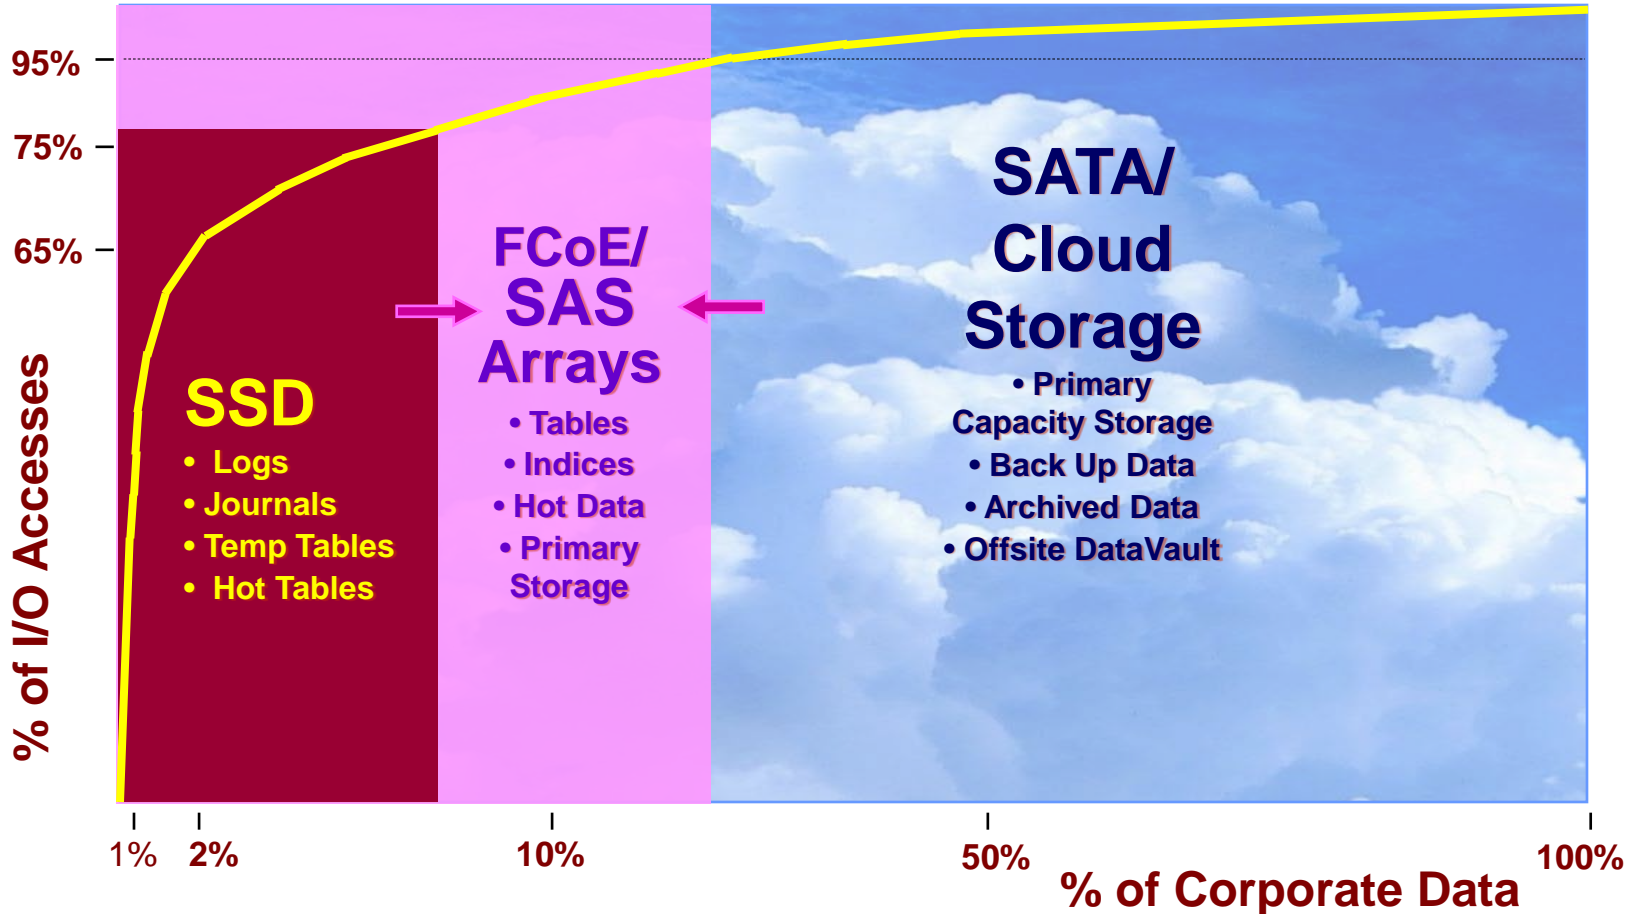
**Database Applications constitute dominant usage in IT . Key to DB performance are random IOPS. SSDs outshine HDD in IOPS price/performance – a major reason, besides better space and power, for their explosive growth.**

# Improving TPS: Saving Memory with SSDs Use

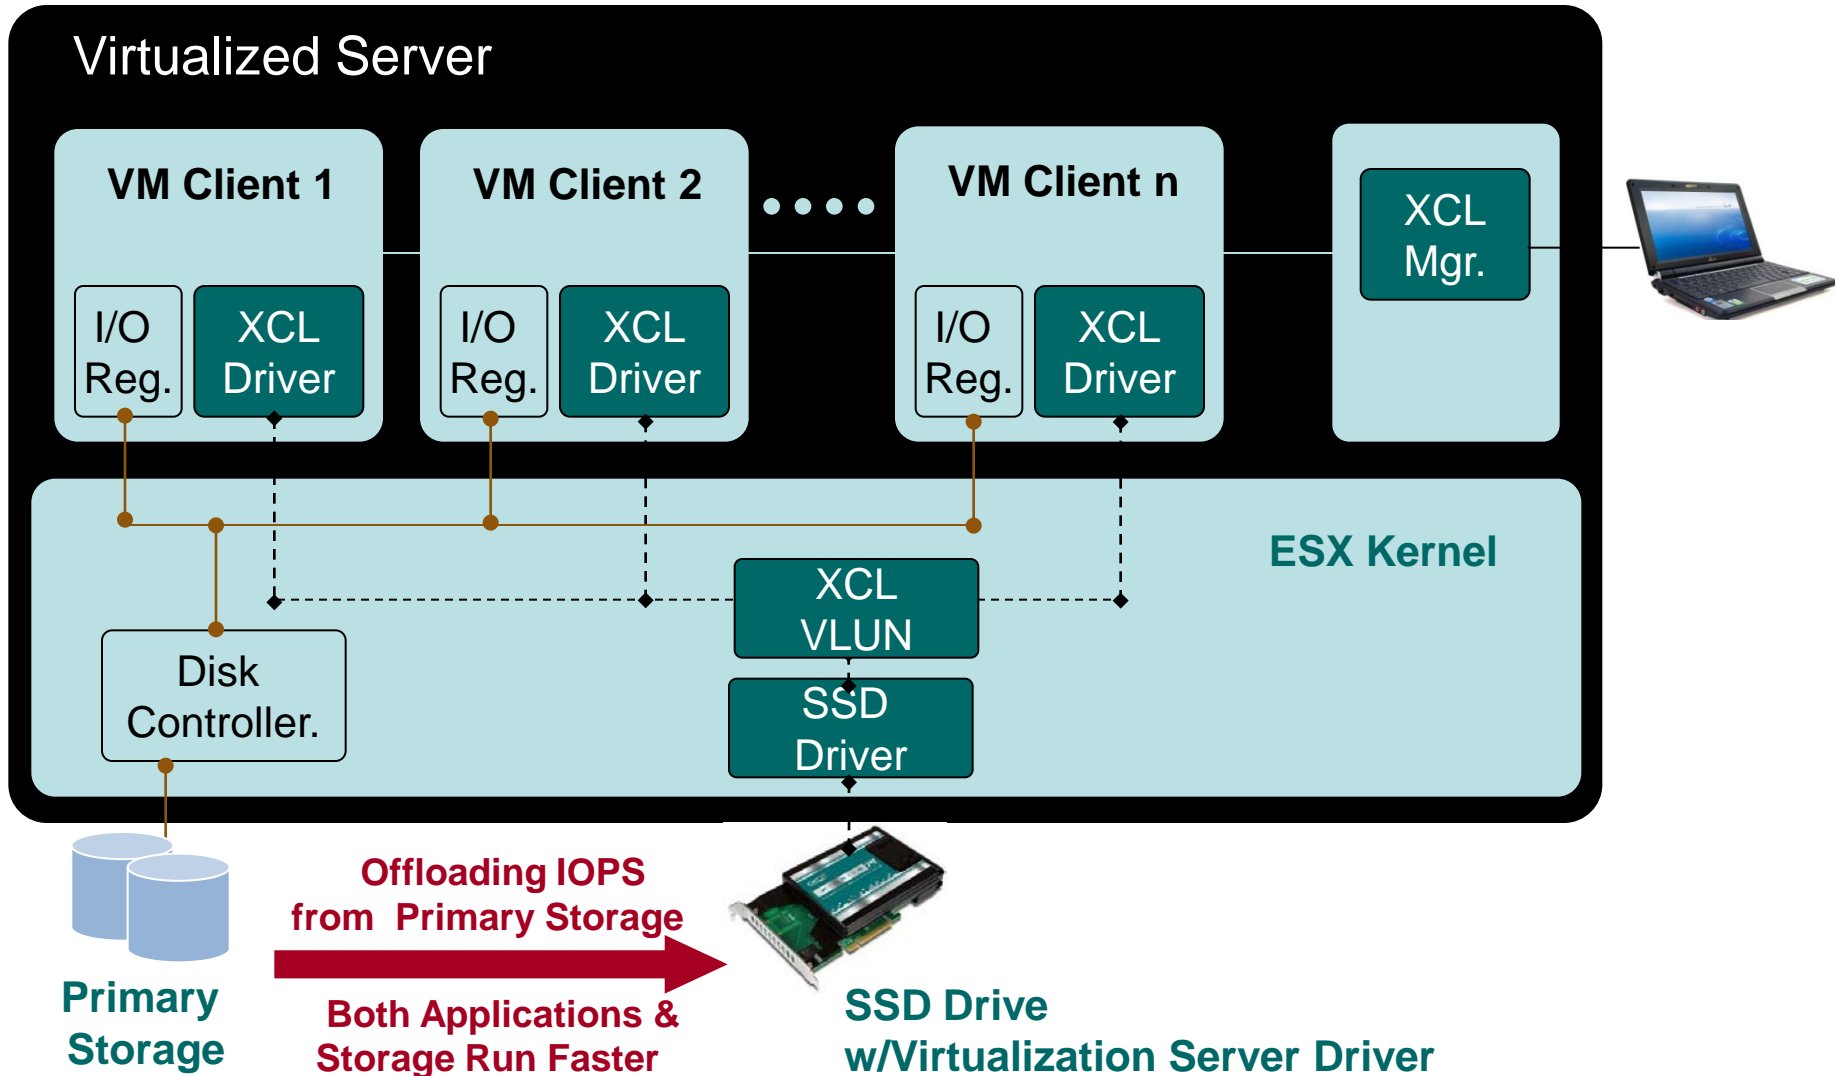


**To achieve a certain TPS improvement, its cheaper to deploy SSD vs increased buffer memory (in GB costs) needed with using HDDs alone.**

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





# IO Bottleneck Mitigation in Virtualized Servers



## 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

# Apps Benefitting from Improved I/O

## VDI






## 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

## ➤ Solving I/O Problems

- ◆ I/O Bottlenecks occur at multiple places in the Compute Stack, the largest being at Storage I/O
- ◆ SSD comes out cheaper/IOP for IO Intensive Apps
- ◆ To get of Reads – Improve Indexing, archive out old data
- ◆ Minimize the impact of writes – Get rid of temp tables/filesorts on slow disks.
- ◆ Compress big varchar/text/blobs

## ➤ 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
- Use Smart Tiering to 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

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

Late Jim Gray, Research Fellow, Microsoft  
Joshua Tseng  
Joseph White  
and the  
SNIA Education Committee

**Author: Anil Vasudeva, President & Chief Analyst, IMEX Research.com**

- Please send any questions or comments on this presentation to SNIA: [tracksolidstate@snia.org](mailto:tracksolidstate@snia.org)