How SSDs Fit in Different Data Center Applications

Tahmid Rahman

Senior Technical Marketing Engineer

NVM Solutions Group

Agenda

- SSD market momentum and drivers
- Placement in server/storage applications
- Application specific requirements and workload characteristics
- Proof points with SSDs in transaction processing, IT, virtualization
- Call to action

Strong SSD Momentum in the Market

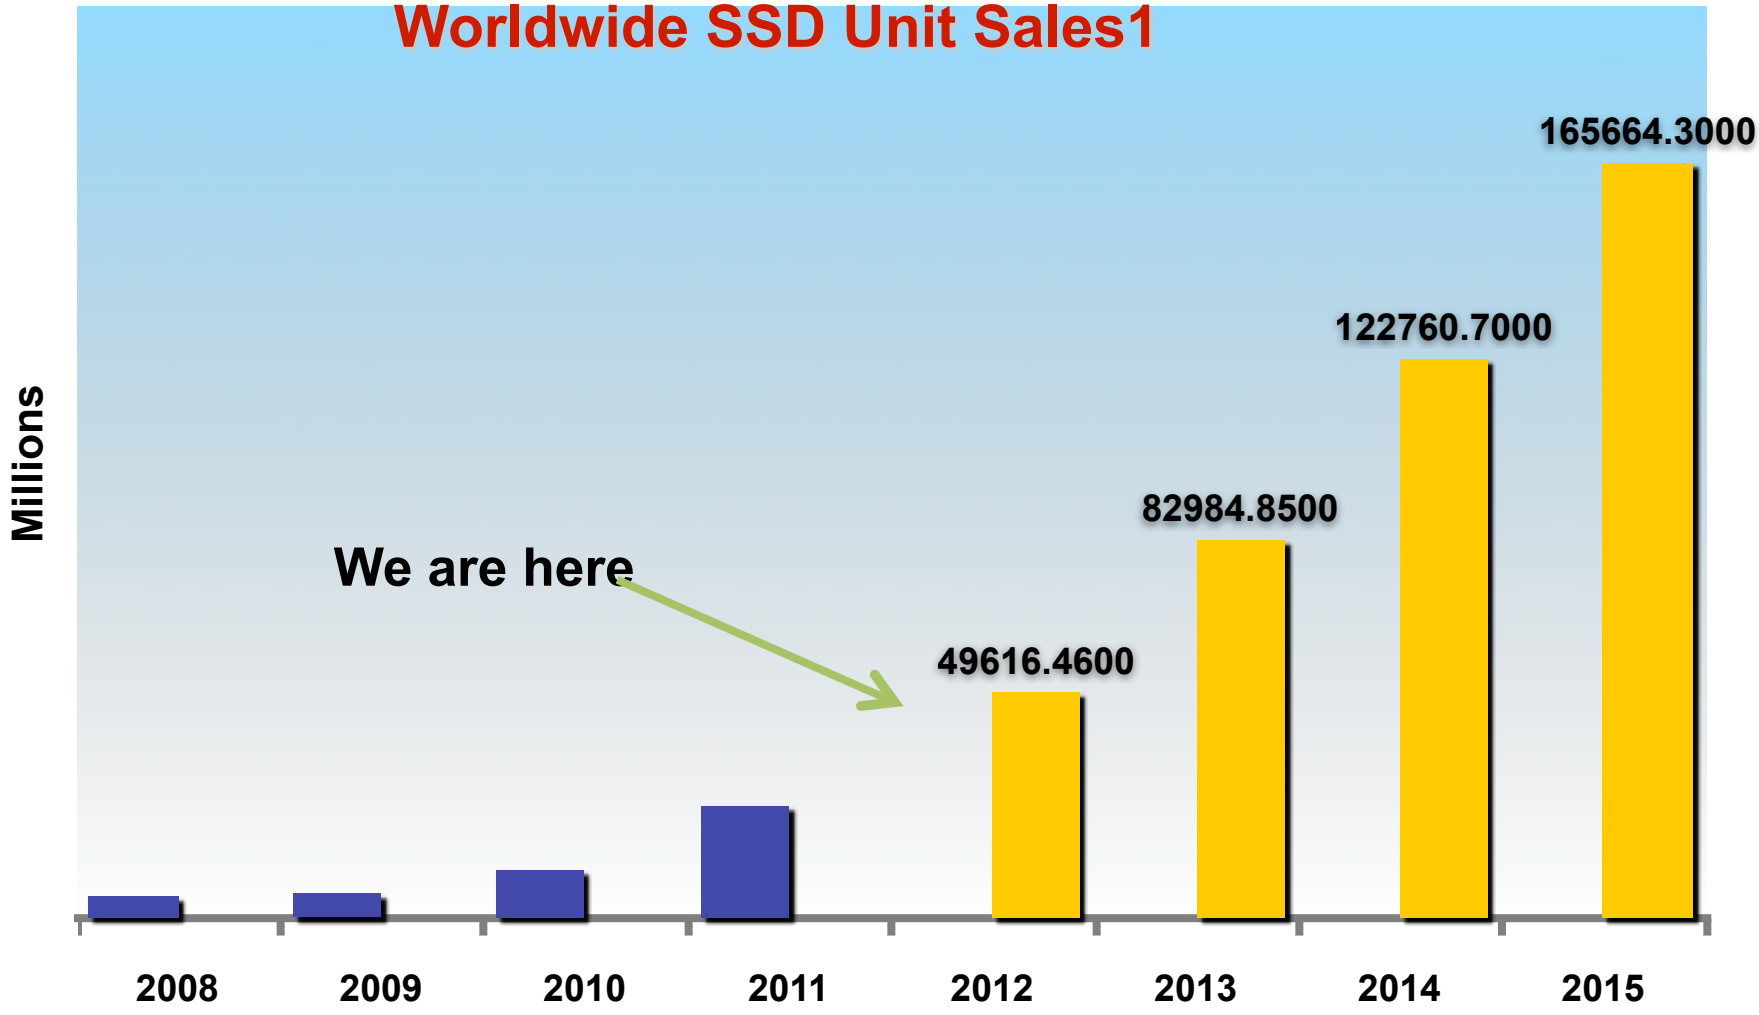

Strong SSD Momentum in the Market

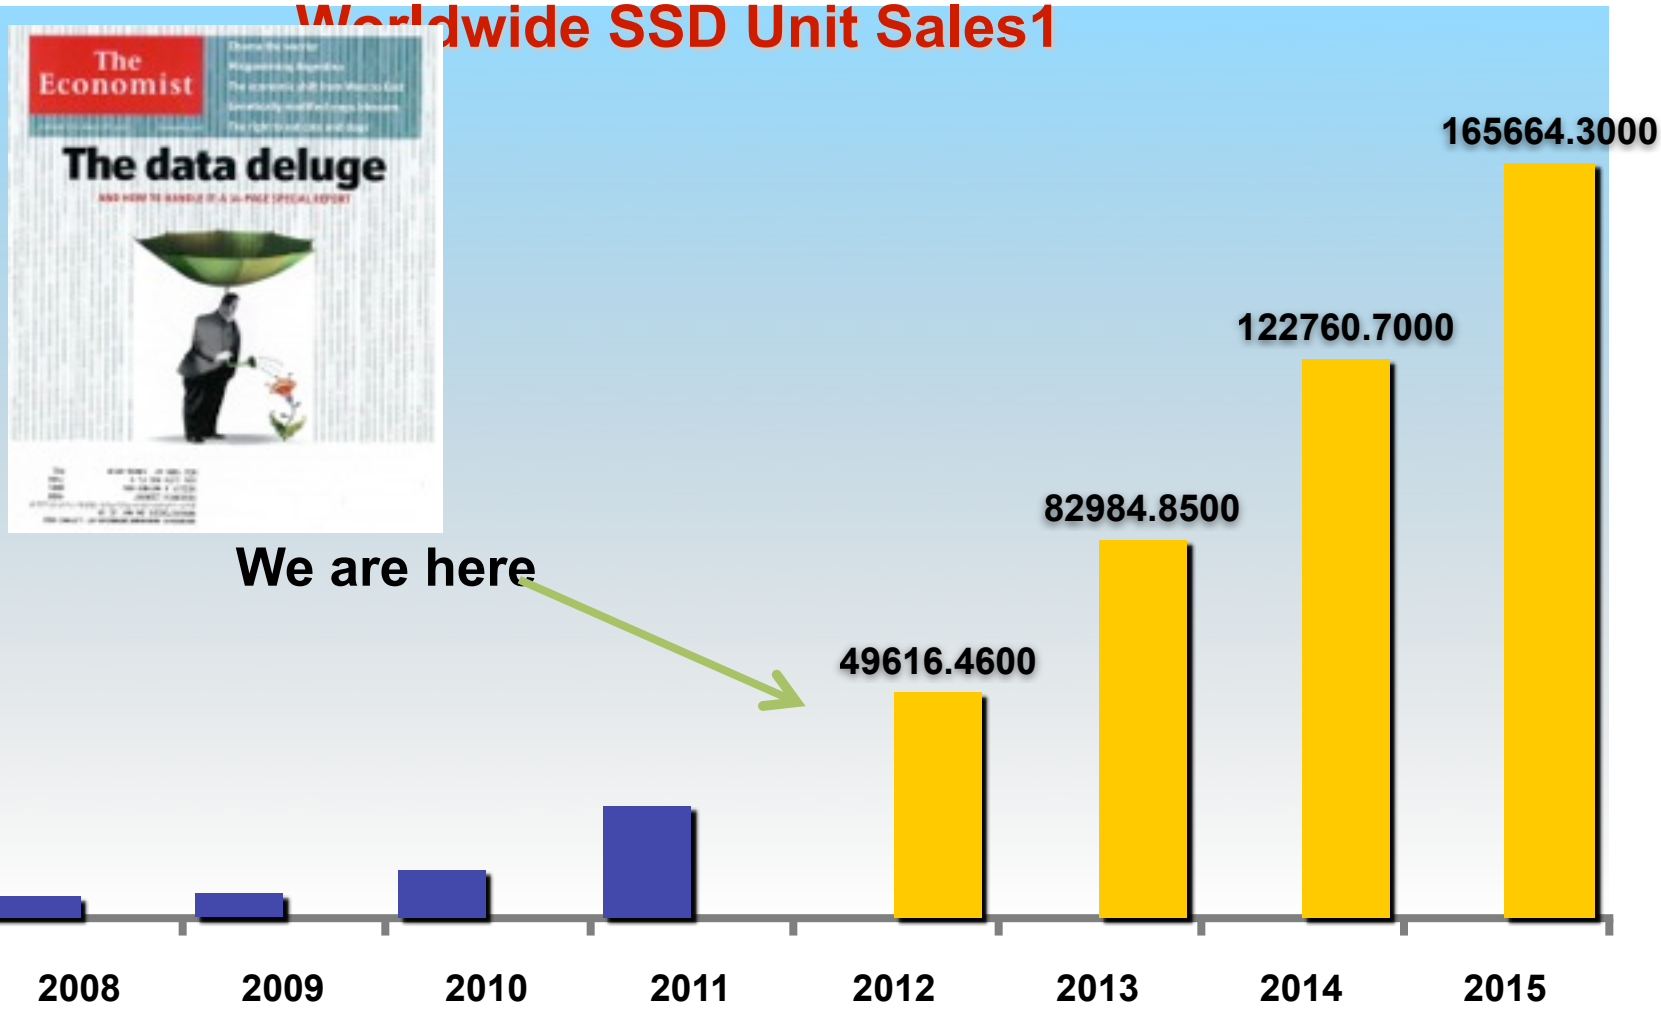

Strong SSD Momentum in the Market

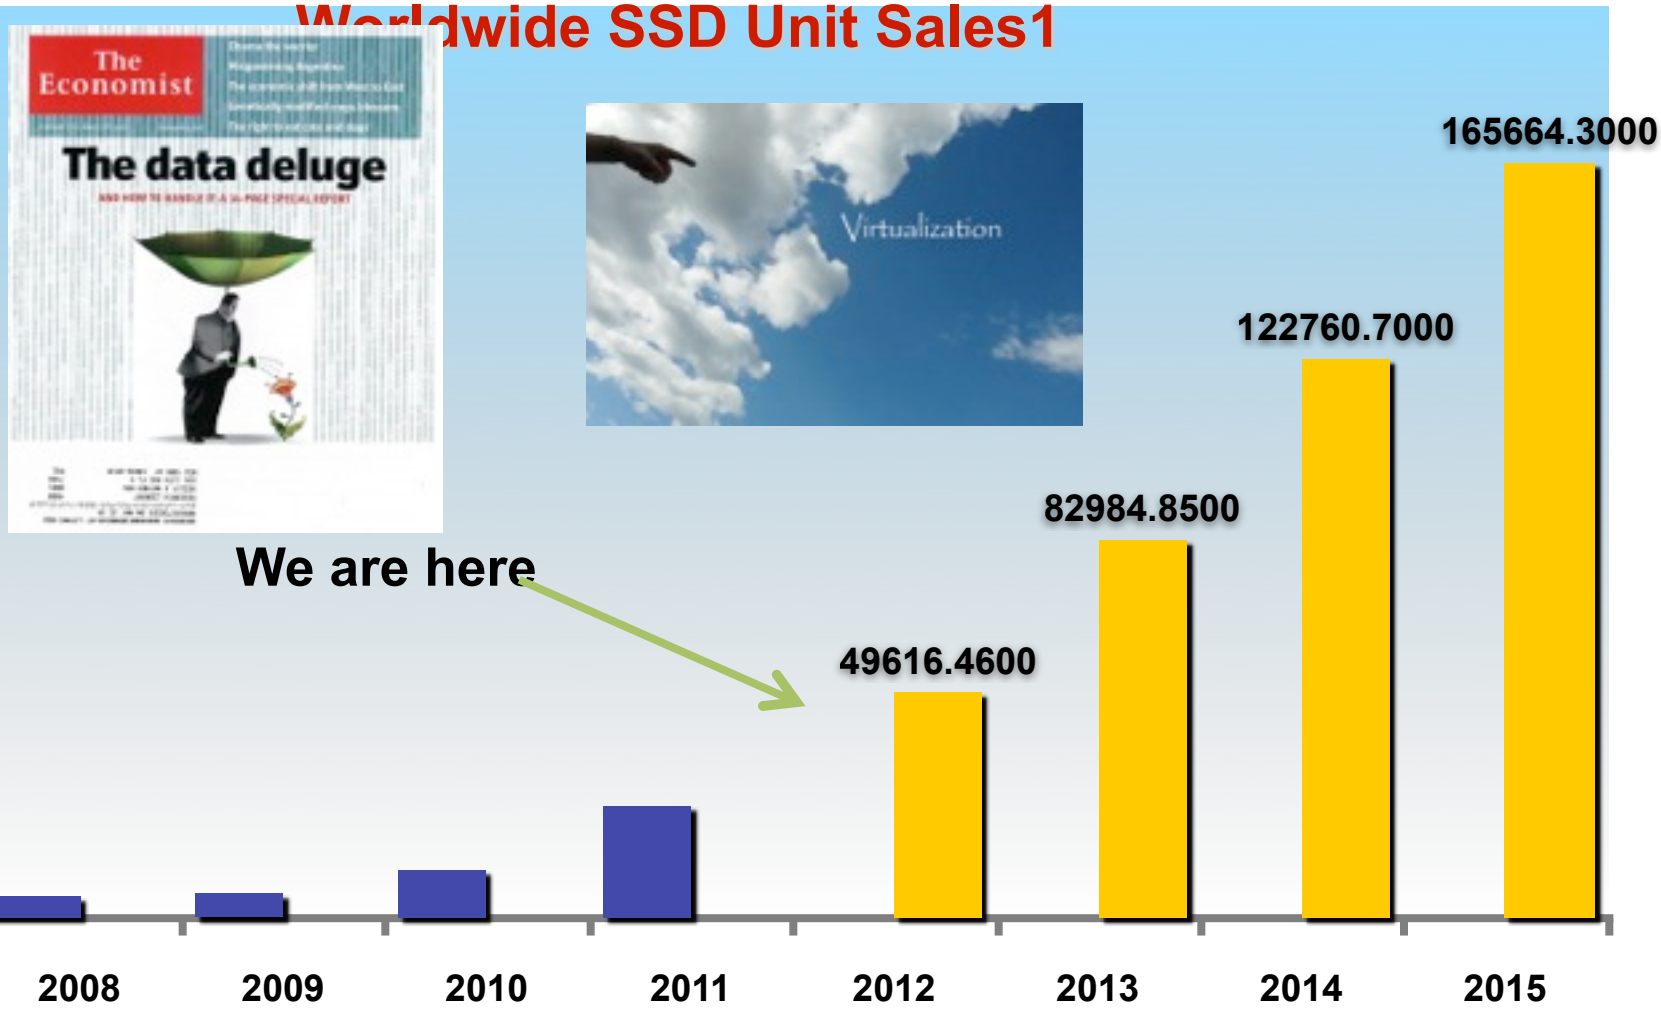

Strong SSD Momentum in the Market

Flash Memory SUMMIT

Storage Tiering

Worldwide SSD Unit Sales1

Millions

We are here

2008	2009	2010	2011	2012	2013	2014	2015
				49616.4600	82984.8500	122760.7000	165664.3000

Source: Gartner Q1'12

Drivers Behind Data Center Storage

- Architectural Changes –
  - Big data
  - Cloud
  - Software innovation for caching, tiering
- Server Side Innovations-
  - De-duplication, compression
  - Thin-provisioning
  - Virtualization
- Interface transitions
  - SATA/SAS to PCIe
  - AHCI based to NVMe
- SSD endurance and performance grades
  - Endurance classes – high, medium, standard
  - Optimization for access – read intensive, write intensive, mixed workload
  - Different “out of the factory” spare area level

SSD Placement in Server/Storage Application

Usage	Applications	Compute (Servers)	External Storage
Cache  (Low, Deterministic Latency, $/IOP	**IPDC Web 2.0**	Persistent cache: (Block Cache, User space buffer cache	Persistent Cache: (e.g. OS block, metadata, de-dupe
Performance ($/IOP/GB) 	**IPDC web2.0 Volume**	Hot Application Data (Web, Database, Email, Search, Videos, IPDC etc)	Hot Application Data
Capacity ($/TB, Watt/TB) 	**Data Warehouse**	Luke-warm Application Data	Cold/Luke-warm
Boot ($/GB) 	**All Server Applicatio**	Local boot data (Operating System, Hypervisor, SWAP, VM, Application Image)	Local boot Data:

Highest Requirements for Data Center SSDs

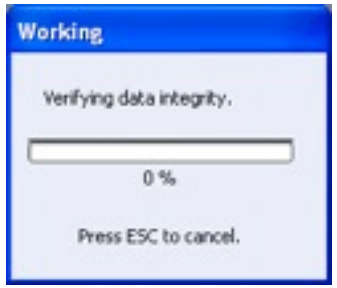

- Data Integrity
  - True End to End data protection
  - Power Loss Protection
  - Power loss cap self test
  - Protection of internal memory with ECC and parity

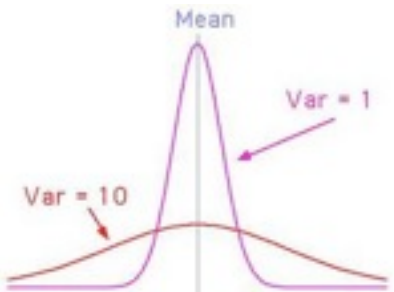

- Predictable Performance
  - IOPS variation needs to be within a narrow range
  - Latency outliers should be within a max value

- High Endurance Requirement
  - Two primary endurance evolving
    - Standard endurance 0.1-1 DWD
    - High endurance 10 DWD

Data Center Application Workload Characteristics

Applications	Transfer Size	% Random	% Read	Write. Endurance	Quality of service	
Media Streaming	64KB	Low	High	Med	Med	Sequential
Web-server Logging	8KB	Low	Low	Med	Med	
Search Engine	4KB/8KB/ 16KB	High	High	Low	High	
Video-On-Demand	128KB	High	High	Low	High	Random Read
Caching	512KB	High	High	Low	Med	
Decision Support	64KB	High	High	Low	High	
Content Delivery Network	16KB/32KB	High	Mixed	High	High	
Database OLTP (On Line Transaction Processing)	4KB/8KB	High	Mixed	High	High	Mixed Random

Source: Industry Standard Benchmarks and Customer Engagement Data
Patterns will vary for unique customers

SSDs For Virtual Storage in the Cloud

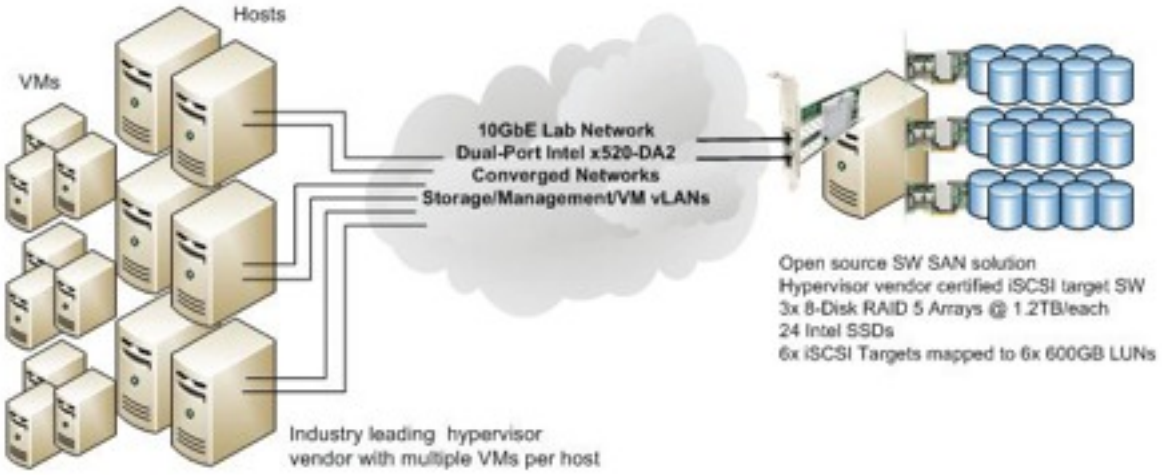

Challenges

- Reversed server to data store ratio (multiple VMs running on single array)
- Adding storage and cache is cost prohibitive

Solutions

- High Performance SSD 3x8 RAID5 meeting multiple VM random IOPS of ~100K w/ SW SAN solution

Impact

- Expanded performance at a lower cost >75% TCO reduction
  - 450 15K RPM HDDs vs. 24 Intel 710 SSDs
  - IT professional would spend $43K instead of $200K+

SSDs for Transaction Processing

- TPoX* (Transactional Processing over XML*) is an application-based benchmark that mimics a storage-bound online transaction procession over XML data for brokerage
- Intel® SSD 910 Series, reveals a replacement ratio of 1 to 180 with Standard Magnetic Drive Solutions
- 1 TB database can be compressed in one single PCIe card and meet the performance of 180 magnetic storage 15K RPM SAS drives

PCIe SSD Based Solution

Server:
Exercising Application Load

Database Storage:
1 PCIe SSD

HDD Based Solution

Server:
Exercising Application Load

4 x 4Gb/s

"Fiber Channel"

	PCIe SSD	Magnetic Storage 15K RPM
Number of drives	**1**	**180**
TPS score (steady-state run)	**13,516**	**13,742**
Latency (msec)	**5.93**	**7.15**
Drive Cost	**3,859**	**59,000**
Storage Subsystem Cost	**X**	**14,000**

Configuration: Intel® Xeon® Processor X5680 (3.33 GHz, 6.40 GT/s Intel® QPI) platforms with Intel® 7500 Chipset, 72GB (18x4GB), 800MHz DDR3 memory, SUSE SLES 11 SP1 operating system, DB2 9.7, and TPoX 2.0 using "M" factor scale (1 TB data size). Hitachi* HUS151P1 CLAR146 146GB SAS 15K RPM drives.

http://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/ssd-910-tpox-brief.pdf

Transaction Processing And Importance of Latency QoS

- Transaction processing requires dense IO (Higher IOPS/GB)
- Systems tune to have no “storage bottleneck”
- No Mercy for latency outliers and occasional drops of IOPS

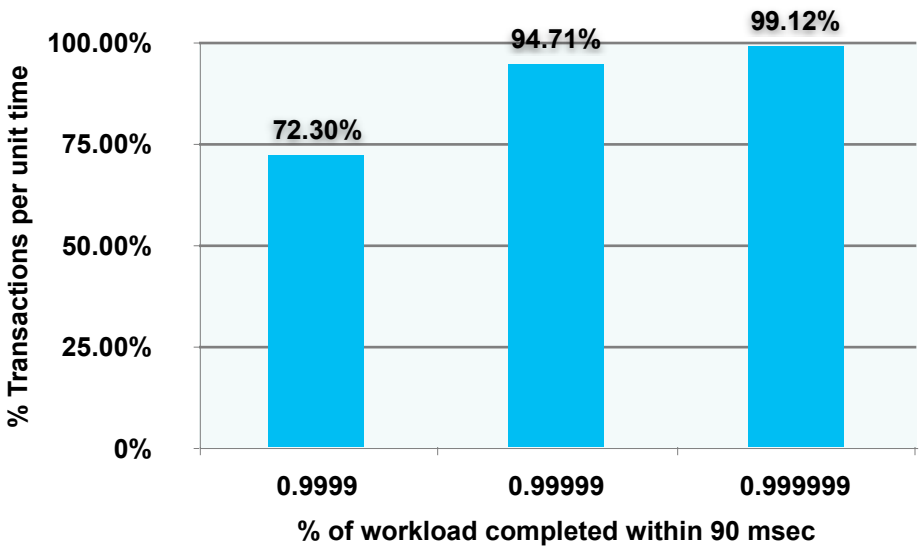

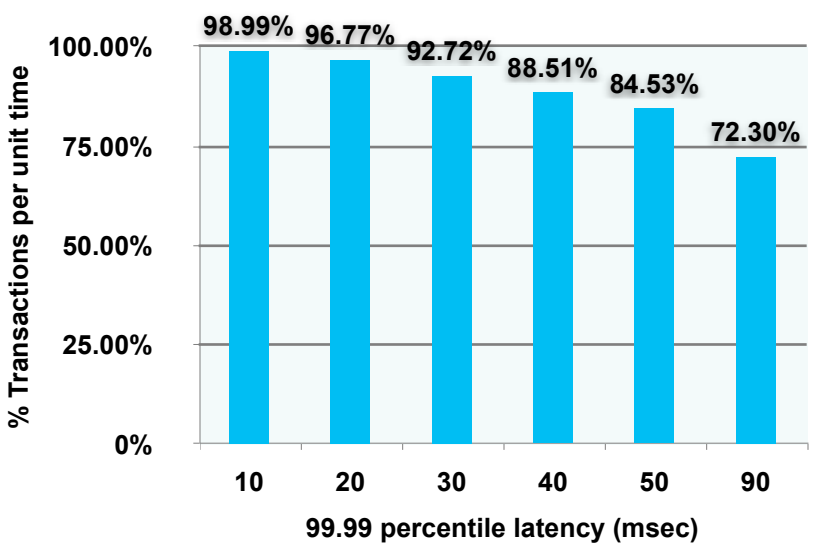

Acknowledgement: Terry Pricket, Jeff Smit, Intel SSG Lab

Transaction Processing And Importance of Latency QoS

- Transaction processing requires dense IO (Higher IOPS/GB)
- Systems tune to have no "storage bottleneck"
- No Mercy for latency outliers and occasional drops of IOPS

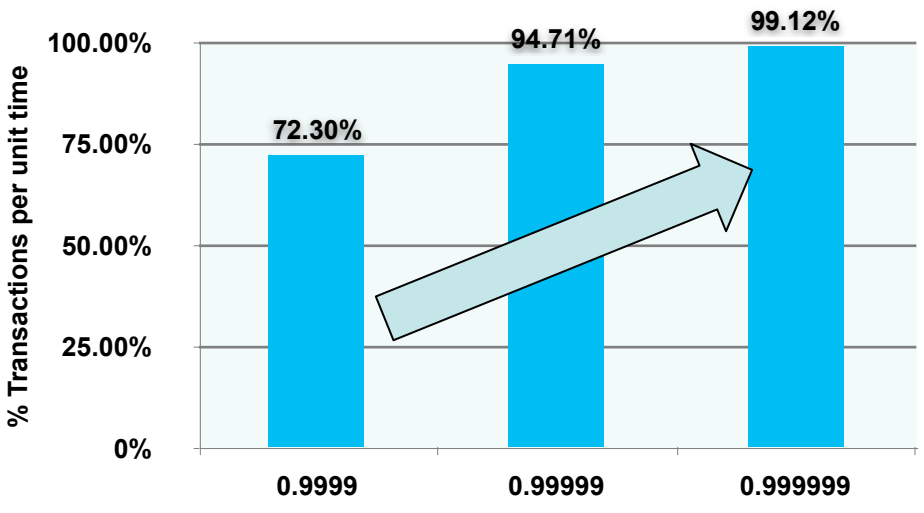

Gain transactions by 27% by moving to less frequent occurrence of outlier event (99.99% to 99.9999%)

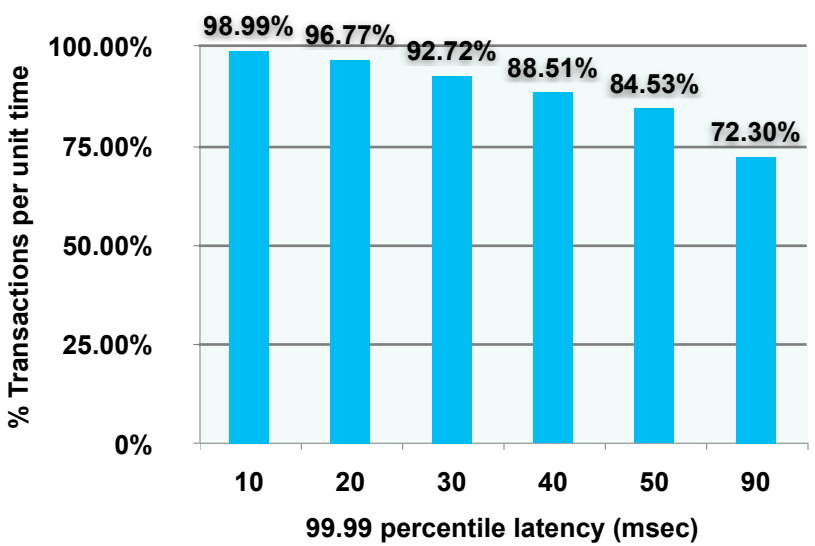

Acknowledgement: Terry Pricket, Jeff Smit, Intel SSG Lab

Transaction Processing And Importance of Latency QoS

- Transaction processing requires dense IO (Higher IOPS/GB)
- Systems tune to have no "storage bottleneck"
- No Mercy for latency outliers and occasional drops of IOPS

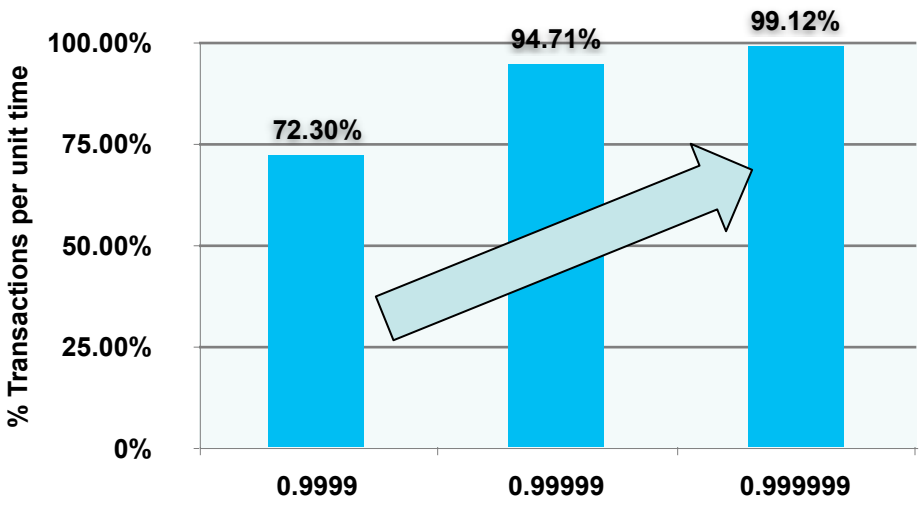

Gain transactions by 27% by moving to less frequent occurrence of outlier event (99.99% to 99.9999%)

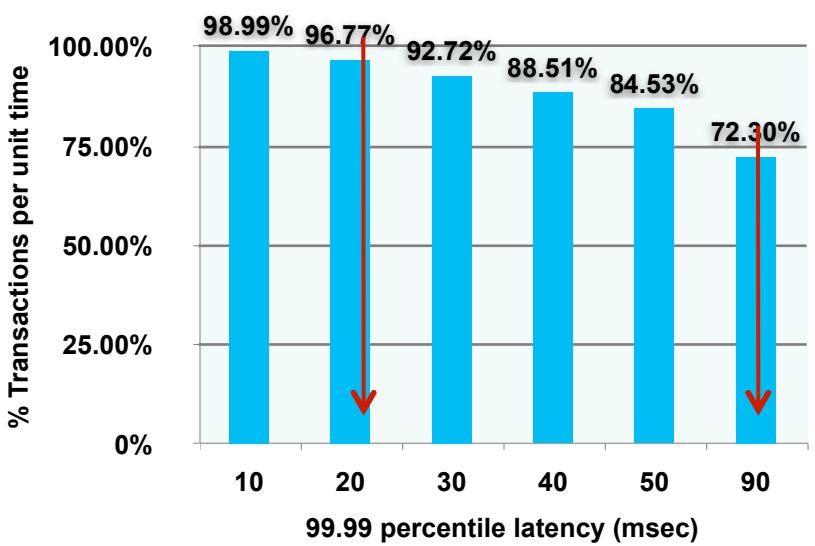

Acknowledgement: Terry Pricket, Jeff Smit, Intel SSG Lab

Transaction Processing And Importance of Latency QoS

- Transaction processing requires dense IO (Higher IOPS/GB)
- Systems tune to have no "storage bottleneck"
- No Mercy for latency outliers and occasional drops of IOPS

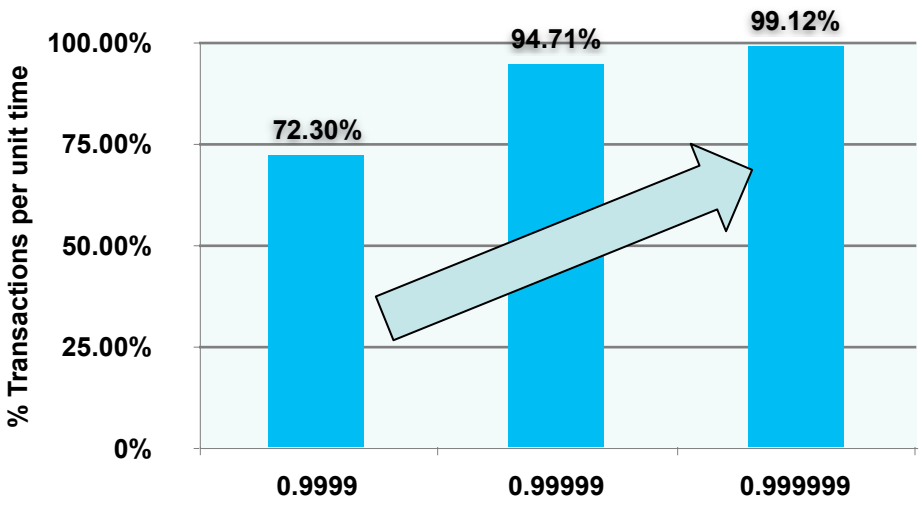

Gain transactions by 27% by moving to less frequent occurrence of outlier event (99.99% to 99.9999%)

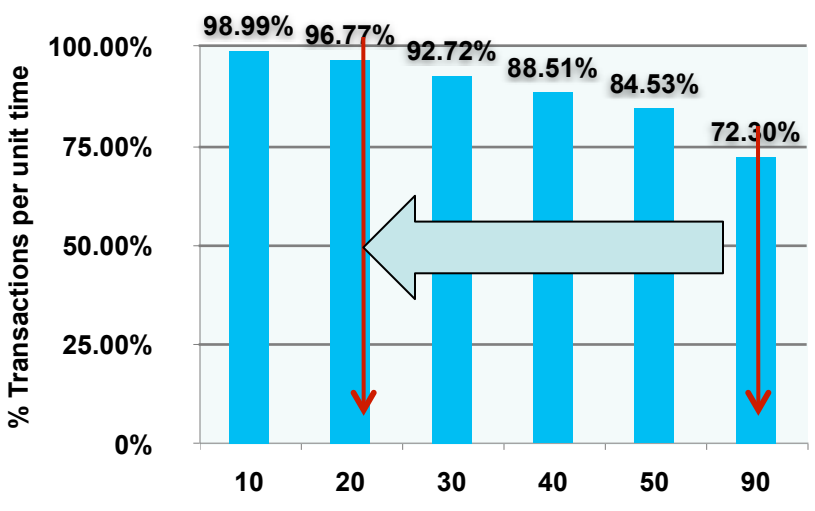

Gain transactions by 24% by moving the latency from 90 msec to 20 msec

Acknowledgement: Terry Pricket, Jeff Smit, Intel SSG Lab

SSDs for IT Management Services

- Automatic Updates for IT security patches
- Managing Design Simulation database
- Swap operation for over-flow memory
- Benchmarking and proof points

Enterprise Patching and Security Compliance Performance Comparison

With 15K RPM Drive

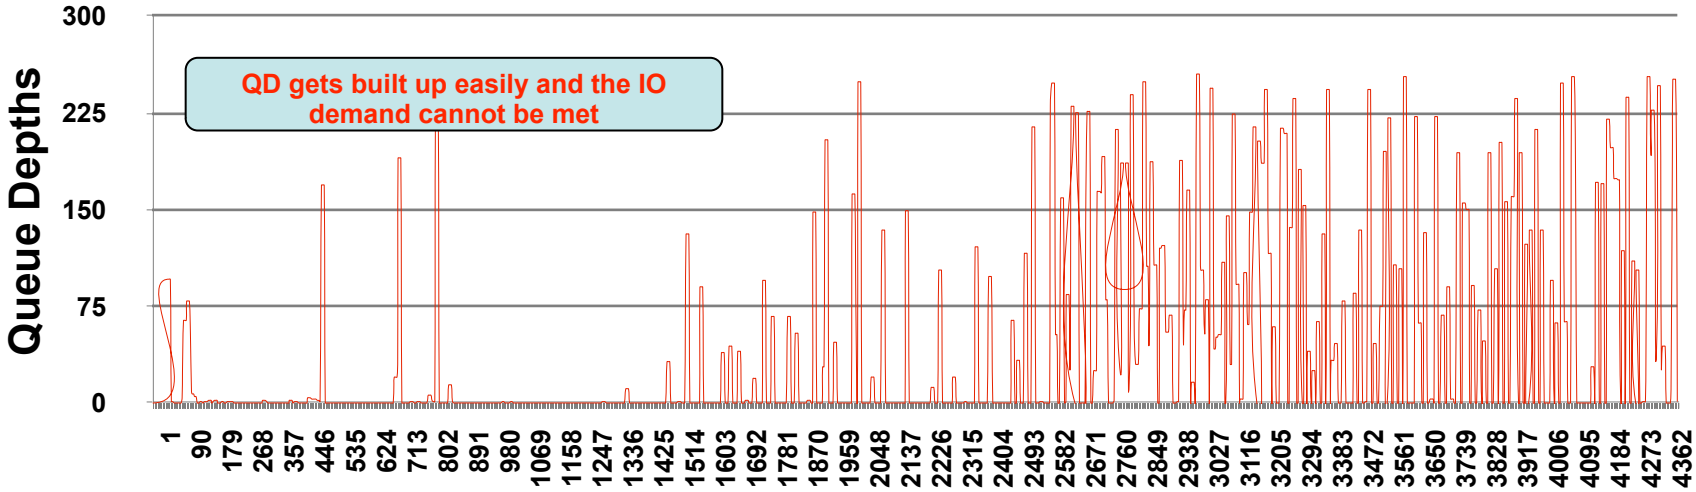

With SSDs deployed

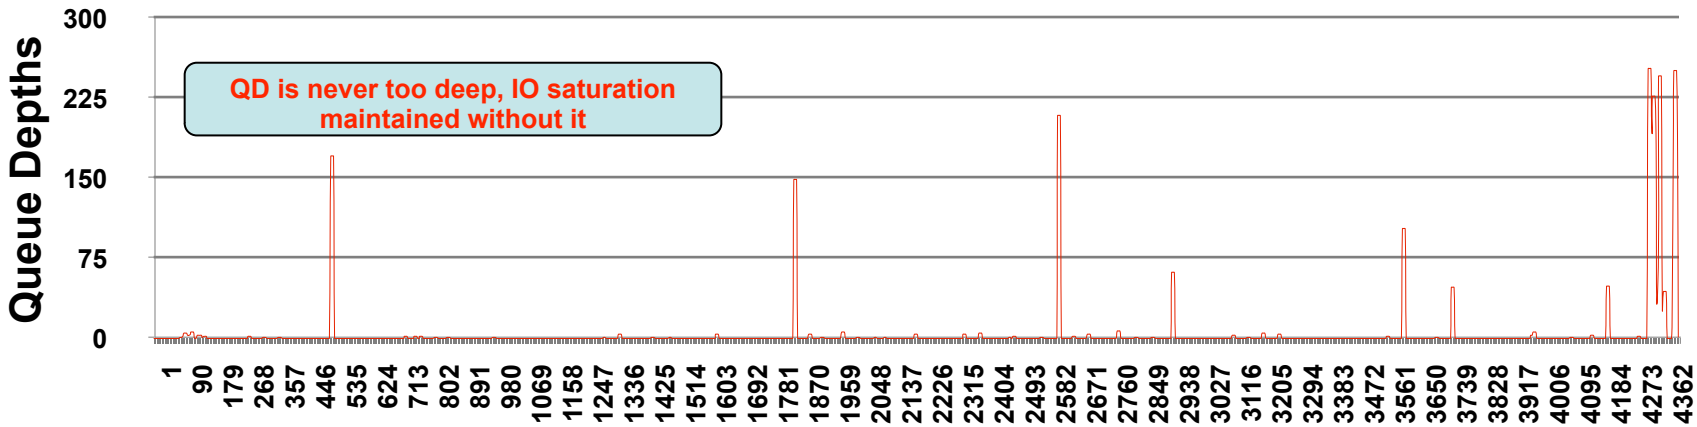

Acknowledgement: Christian Black, Intel IT Architect

Are All SSDs Ready for Transaction Processing?

Are All SSDs Ready for Transaction Processing?

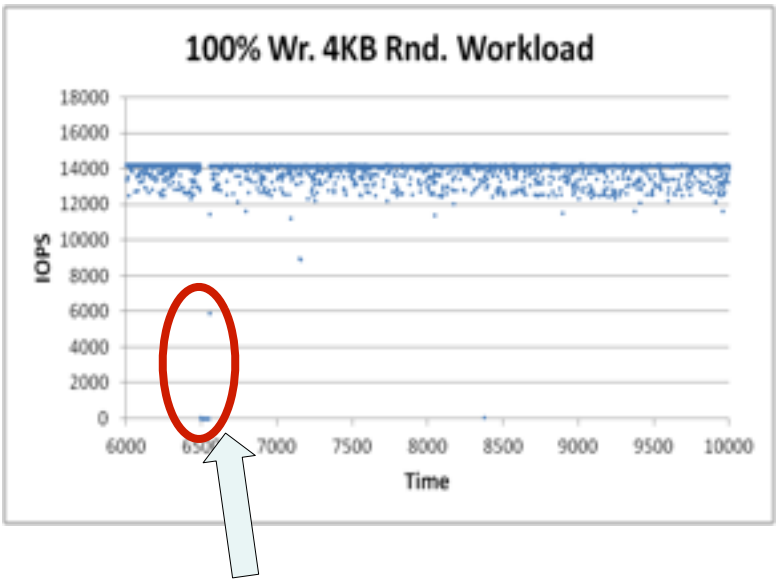

"Zero" IOPS!

Are All SSDs Ready for Transaction Processing?

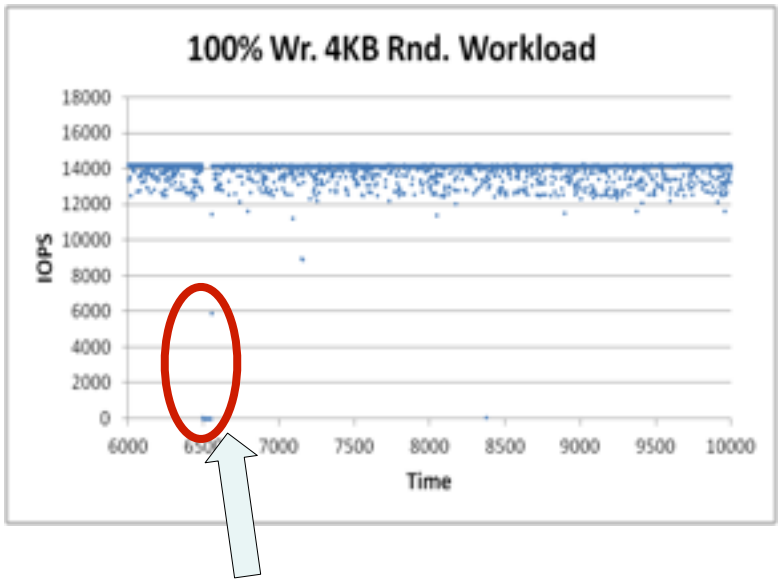

"Zero" IOPS!

← **RAID Array stalls and timeouts**

← **Higher drive counts to meet IO needs**

Are All SSDs Ready for Transaction Processing?

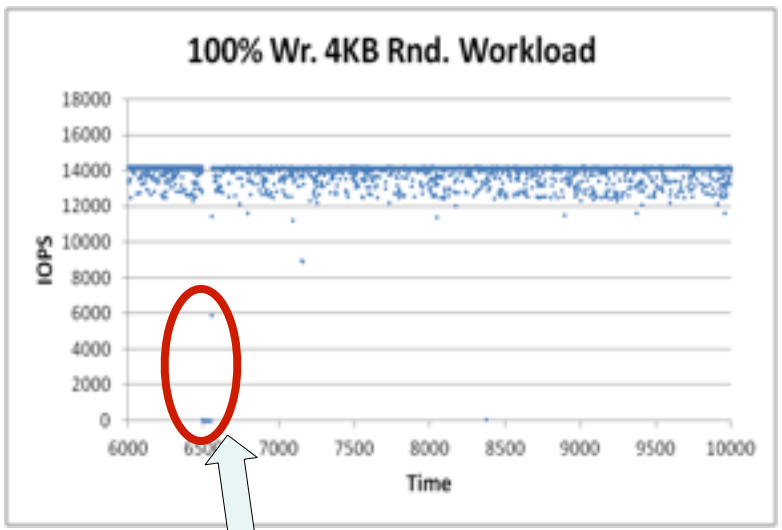

"Zero" IOPS!

← **RAID Array stalls and timeouts**

← **Higher drive counts to meet IO needs**

1 sec max latency!

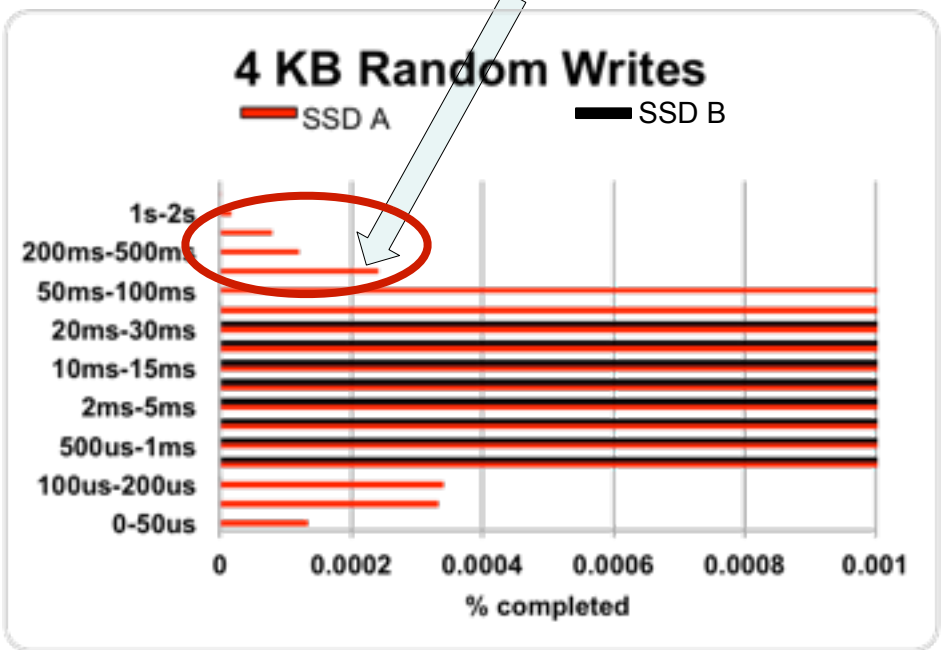

Are All SSDs Ready for Transaction Processing?

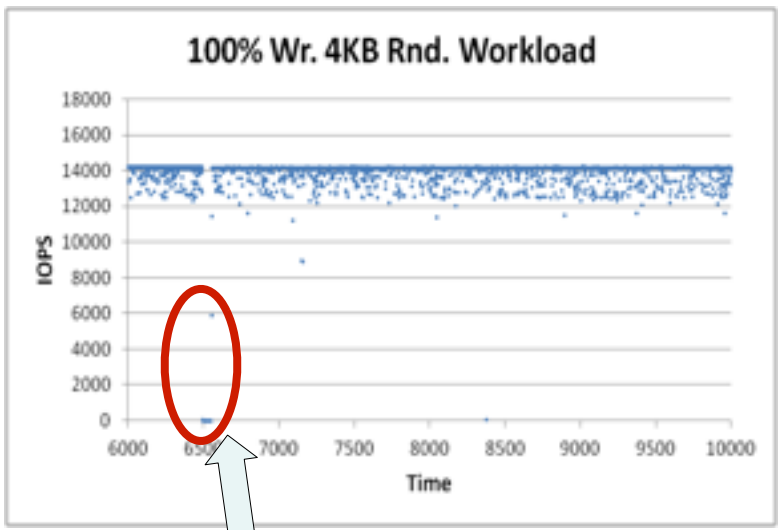

"Zero" IOPS!

← RAID Array stalls and timeouts

← Higher drive counts to meet IO needs

1 sec max latency!

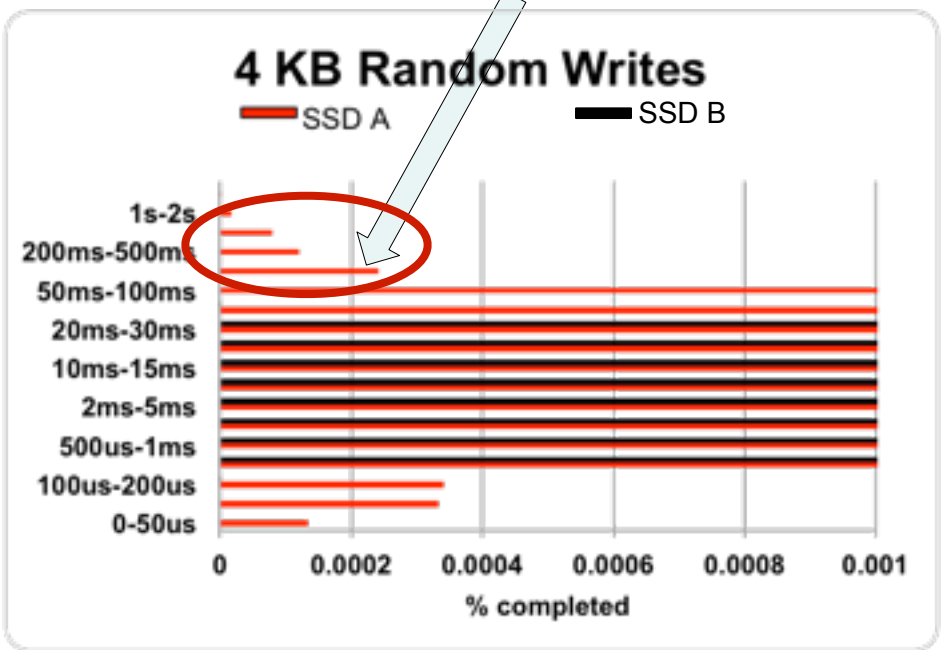

Negative SLA impacts →

Catastrophic for certain applications →

Call to Action

- Ample opportunity for SSD proliferation within data center
- Innovate around applications needs
- Use faster interface and technology to unleash NAND backend bandwidth