SMART STORAGE™ SYSTEMS

Making NAND Better

Reaching the Final Latency Frontier

John Scaramuzzo

President SMART Storage Systems

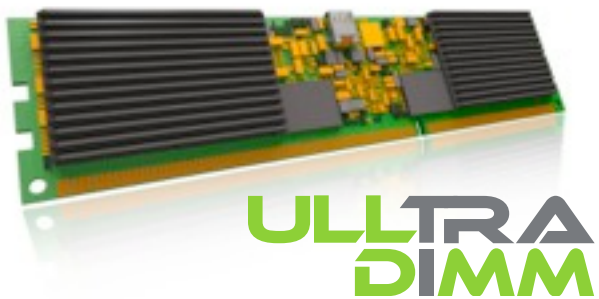

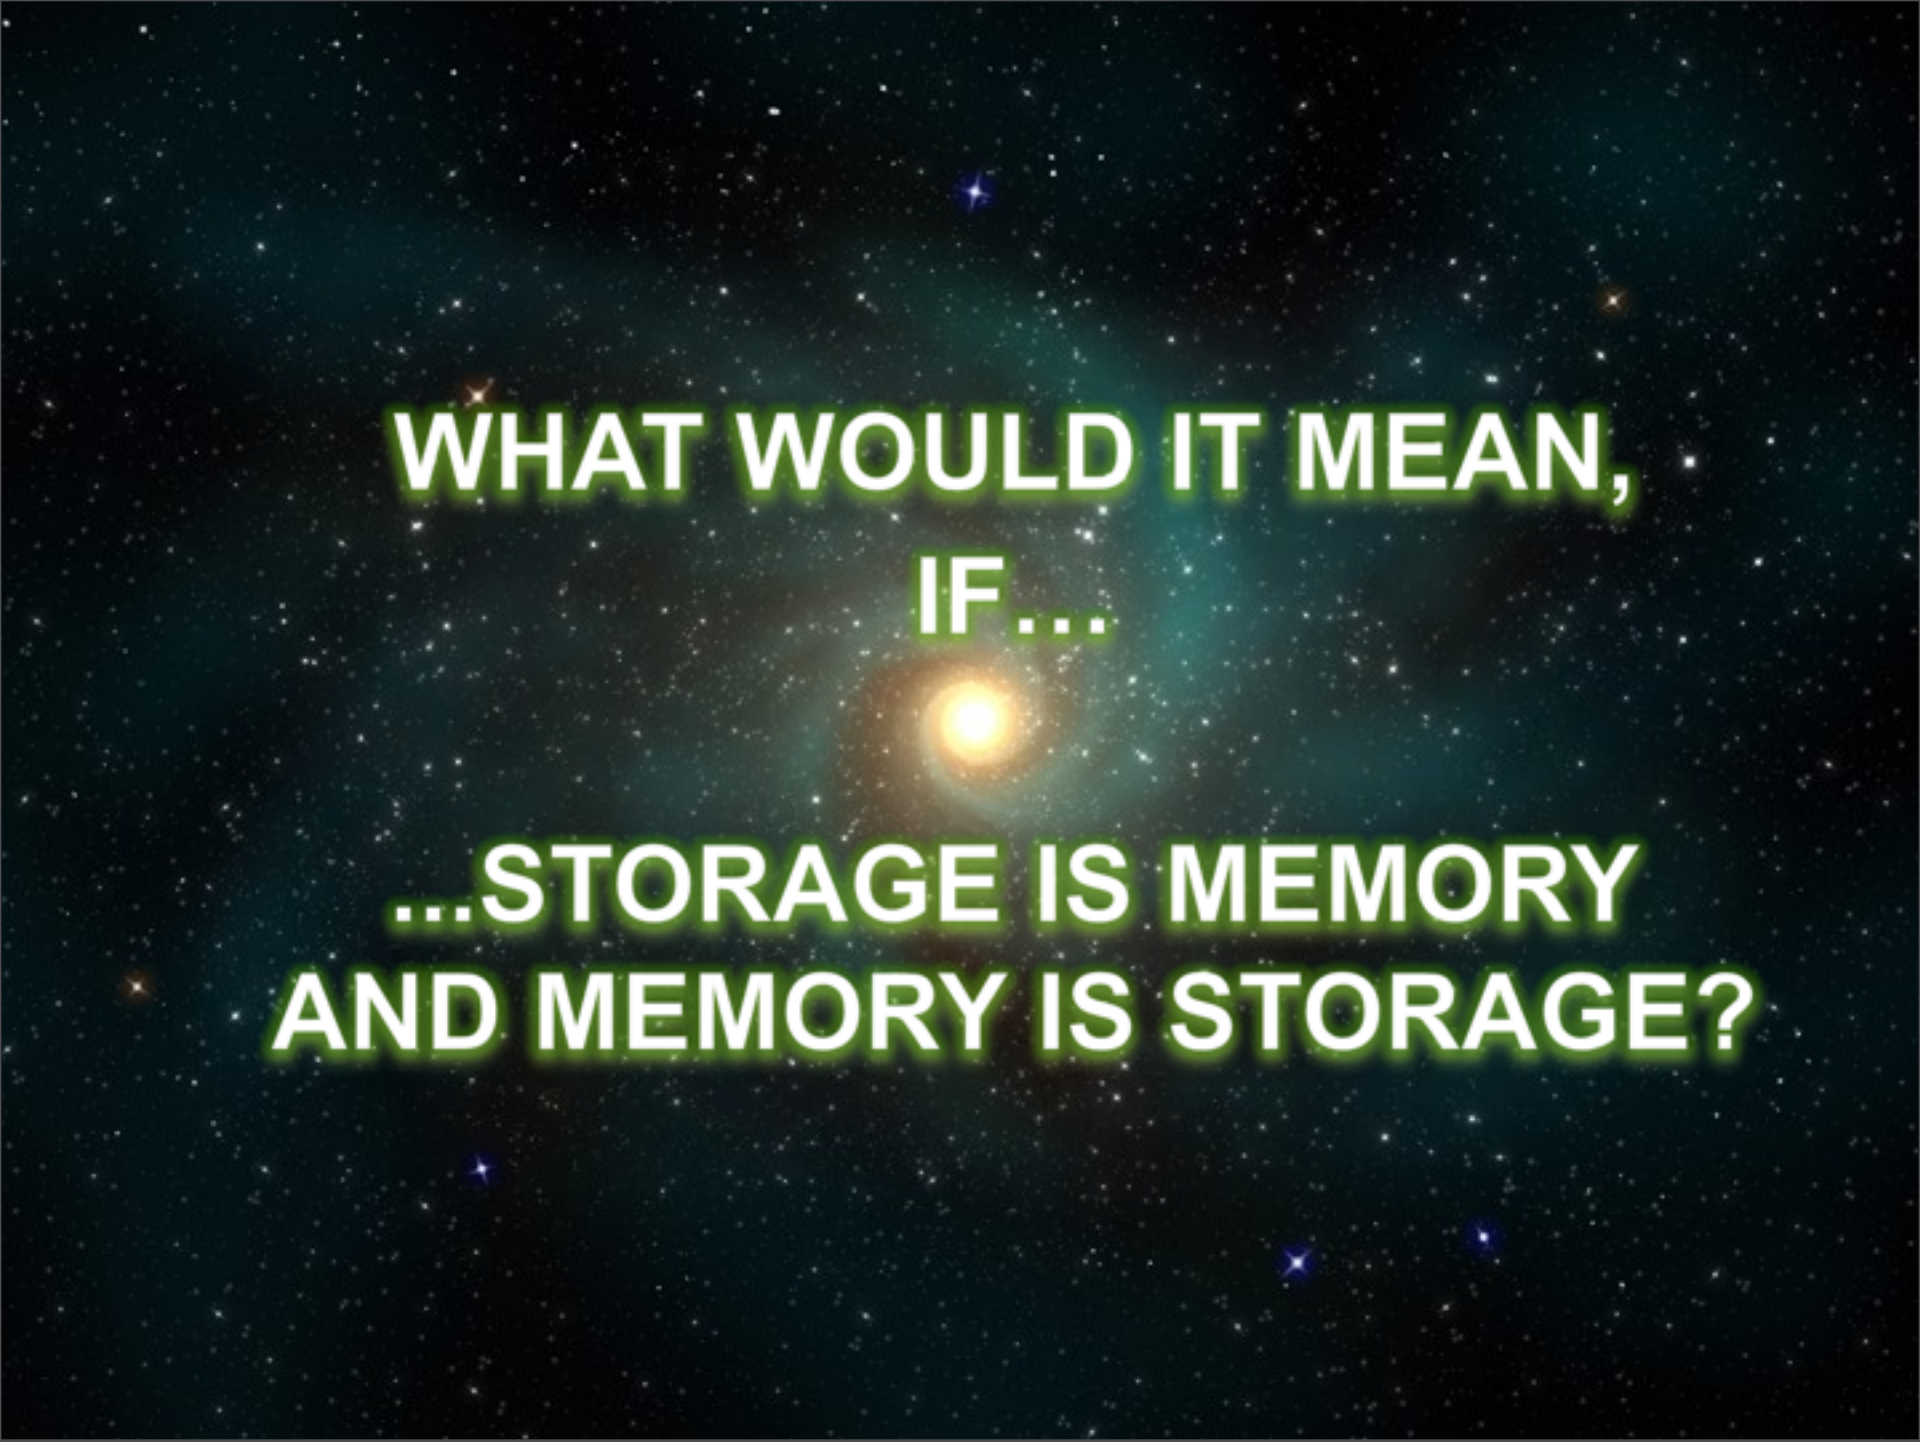
WHAT WOULD IT MEAN,
IF...
...STORAGE IS MEMORY
AND MEMORY IS STORAGE?

Our Expectations Are Changing

Source: PhoCus Wright

Our Expectations Are Changing

57% of **online consumers** will abandon a site after waiting **3 seconds** for a page to load

Source: PhoCus Wright

Our Expectations Are Changing

57% of **online consumers** will abandon a site after waiting **3 seconds** for a page to load

80% of these people **WILL NOT RETURN**

Source: PhoCus Wright

Our Businesses Are Changing

High Frequency Trading - 09:37am

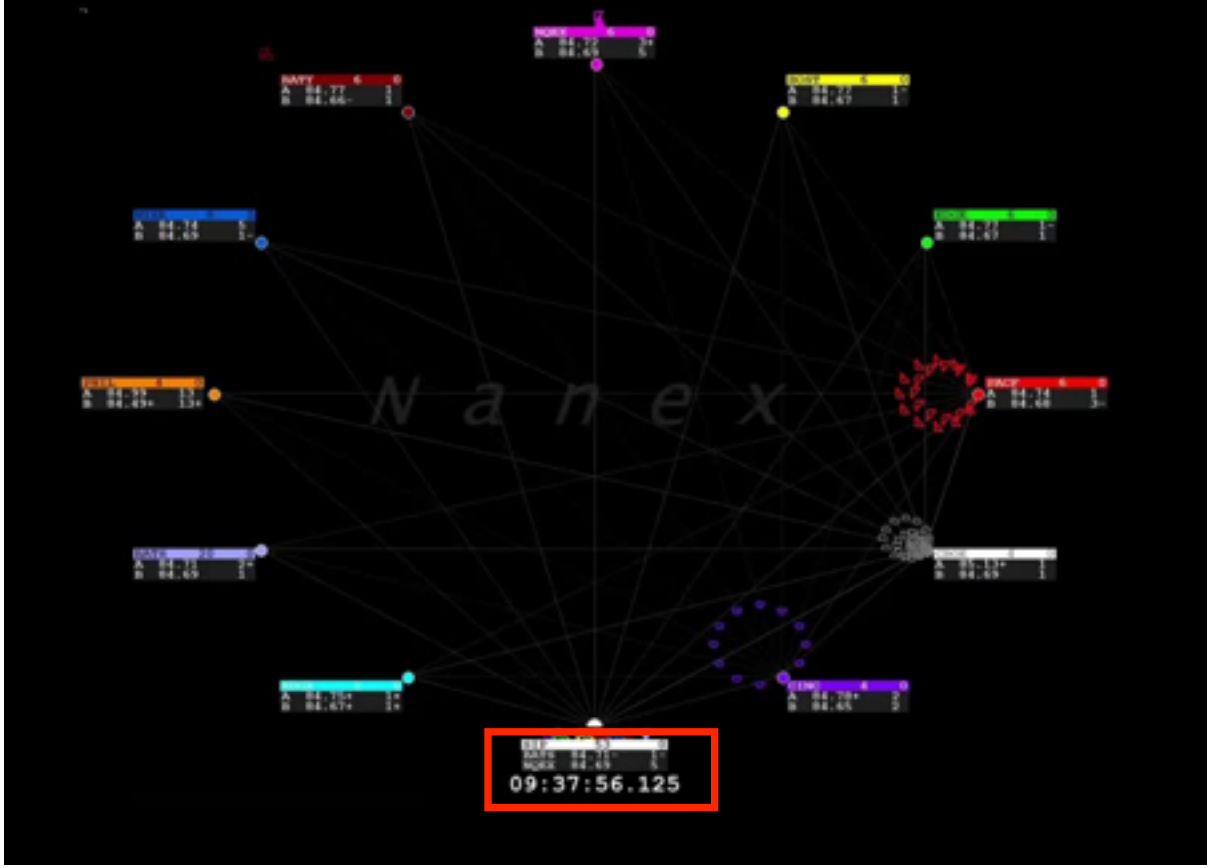

Our Businesses Are Changing

High Frequency Trading - 09:37am

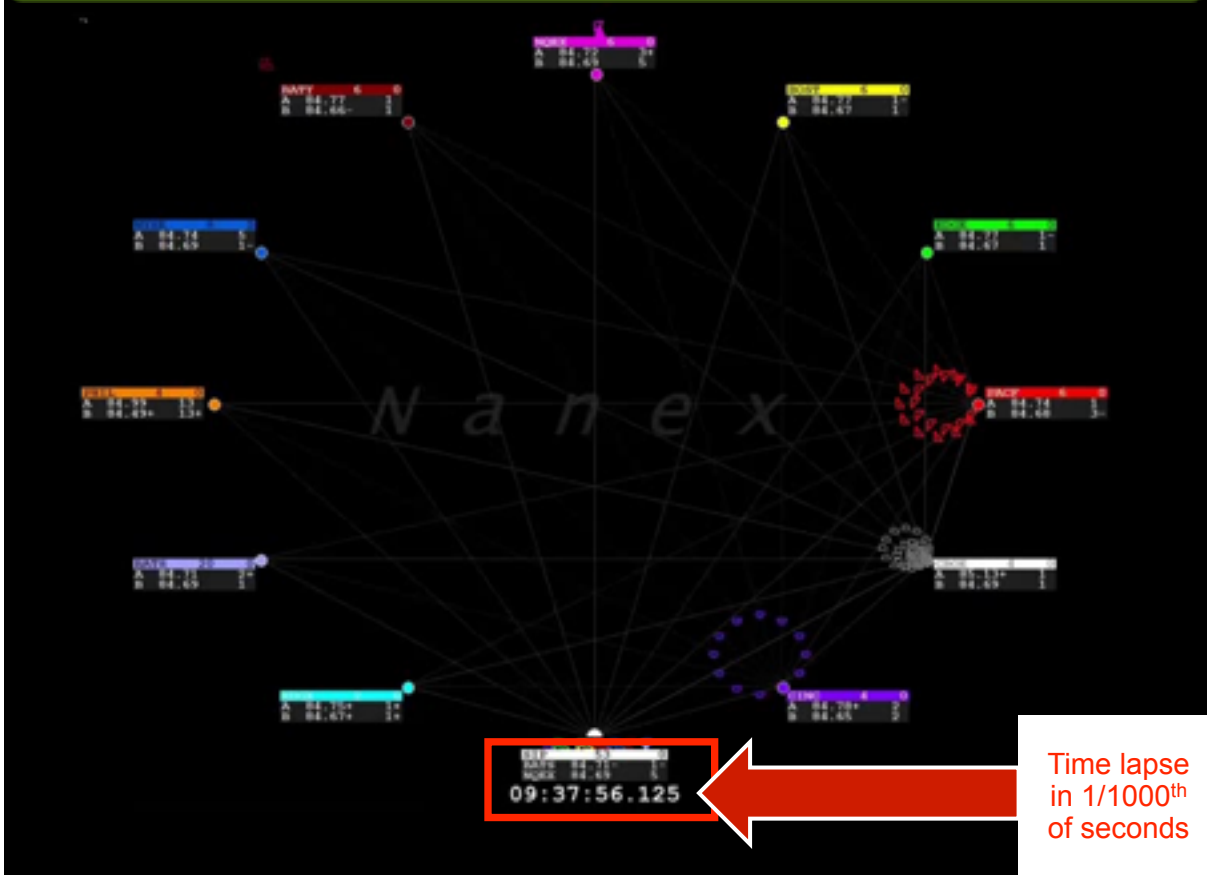

Our Businesses Are Changing

High Frequency Trading - 09:37am

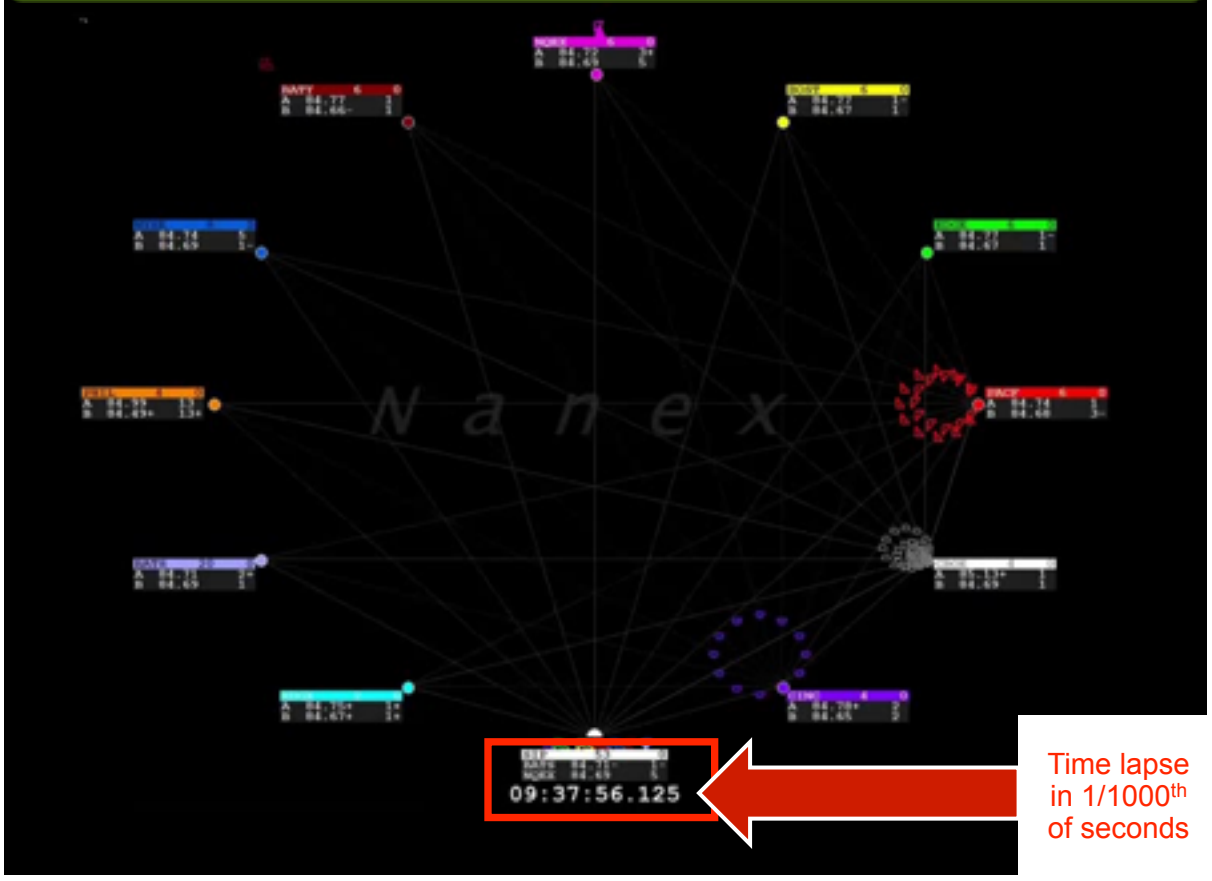

1ms improvement = $100M profit per year*

* Source: http://www.informationweek.com/wall-streets-quest-to-process-data-at-th/199200297

Our Computing Platforms Are Changing

SMART STORAGE SYSTEMS

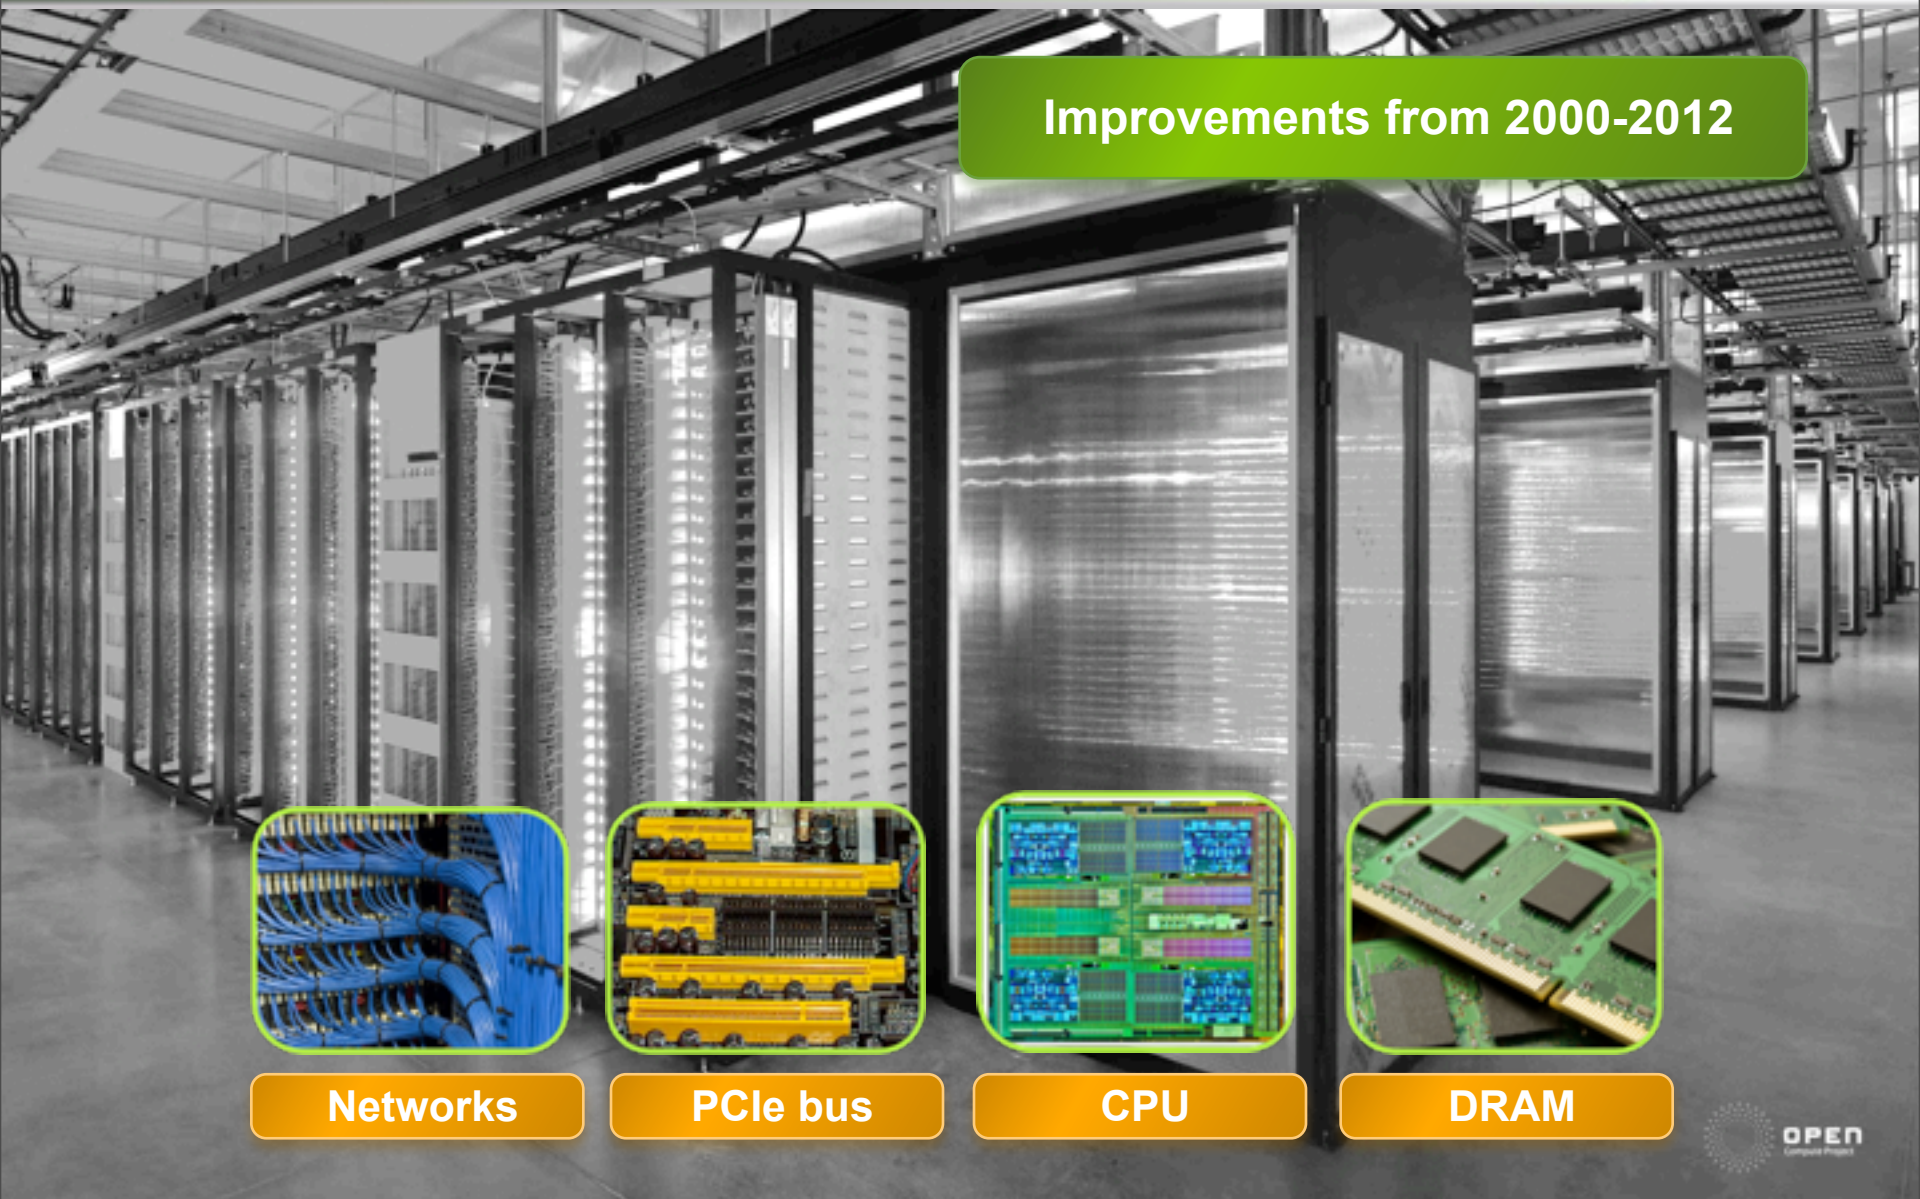

Our Computing Platforms Are Changing

SMART STORAGE SYSTEMS

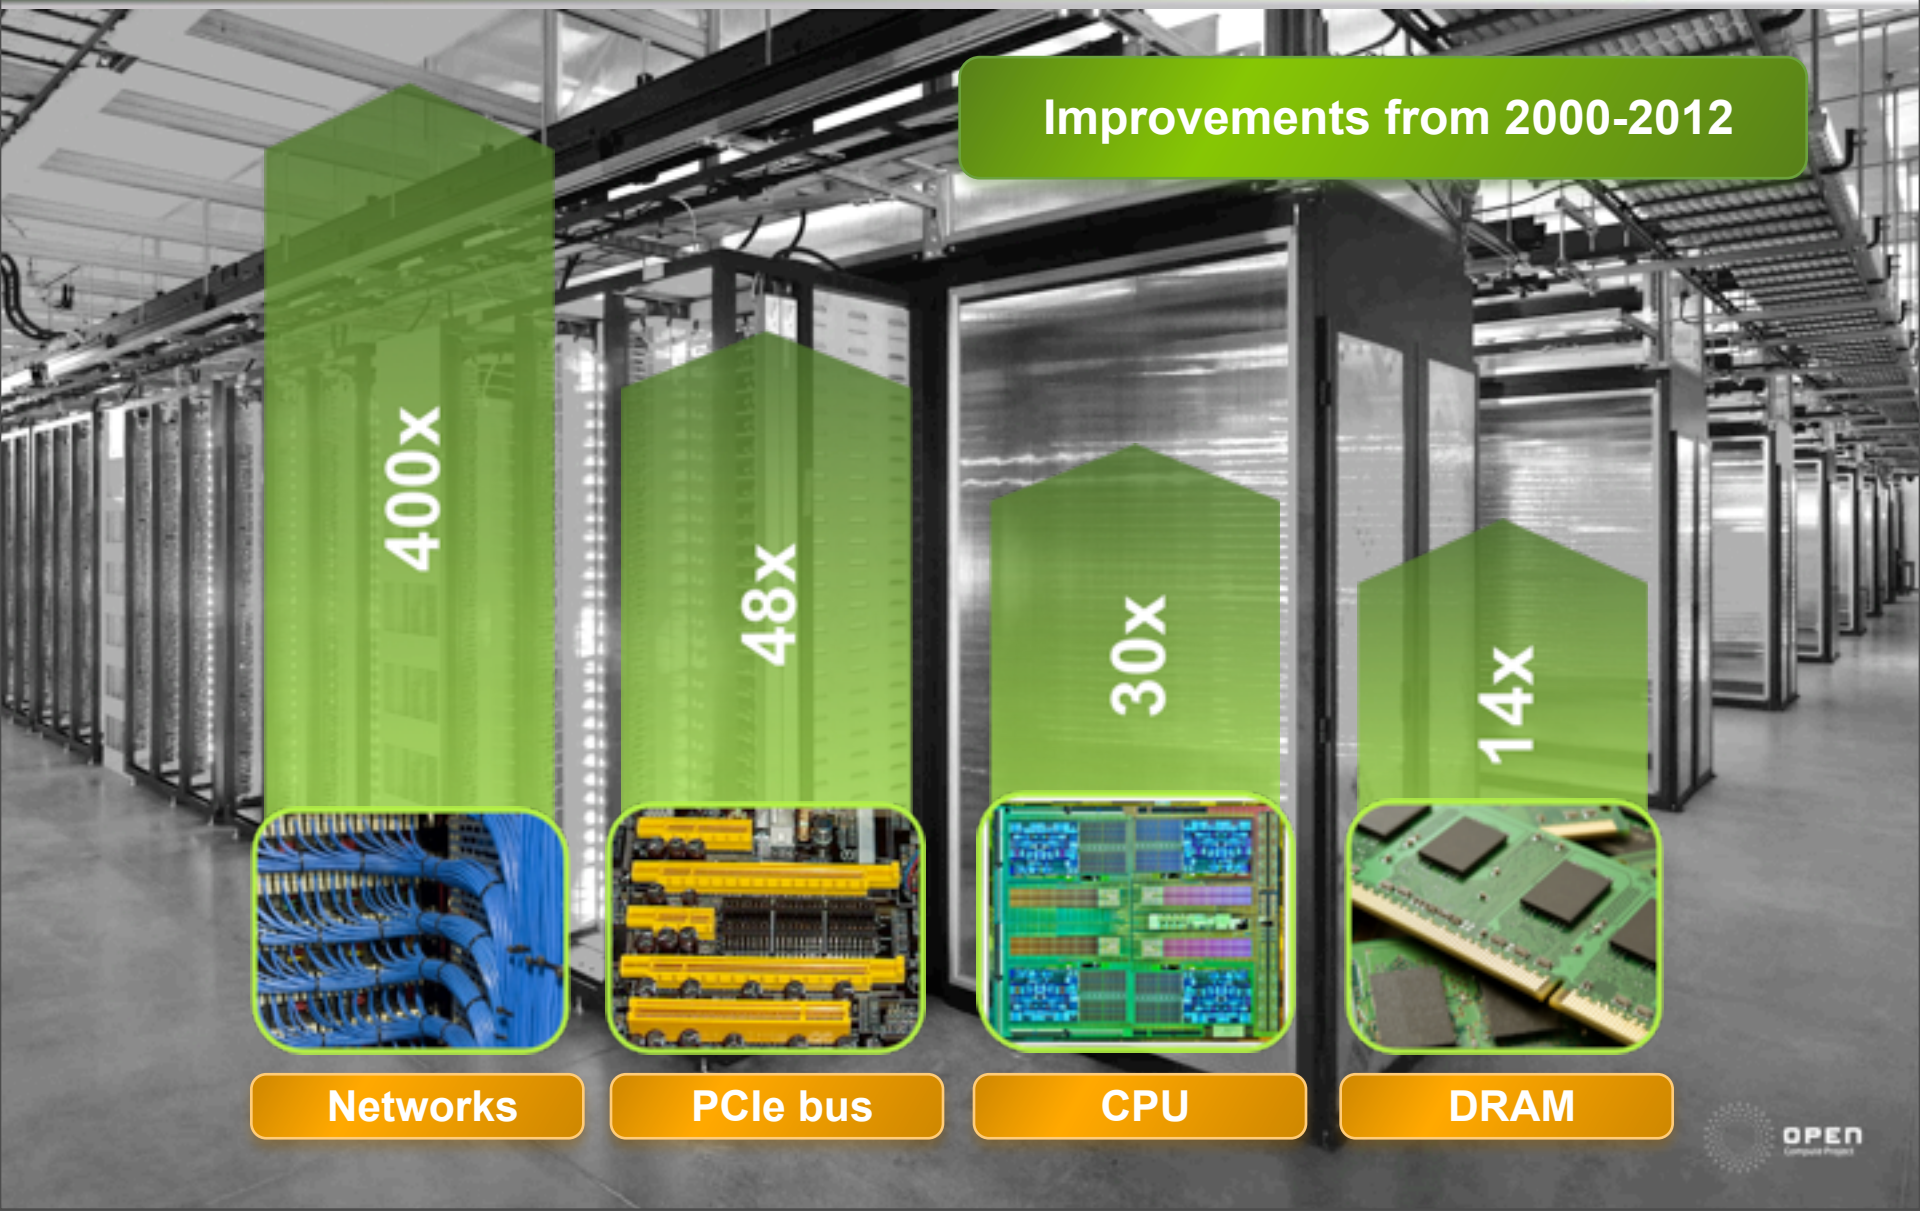

The I/O Bottleneck

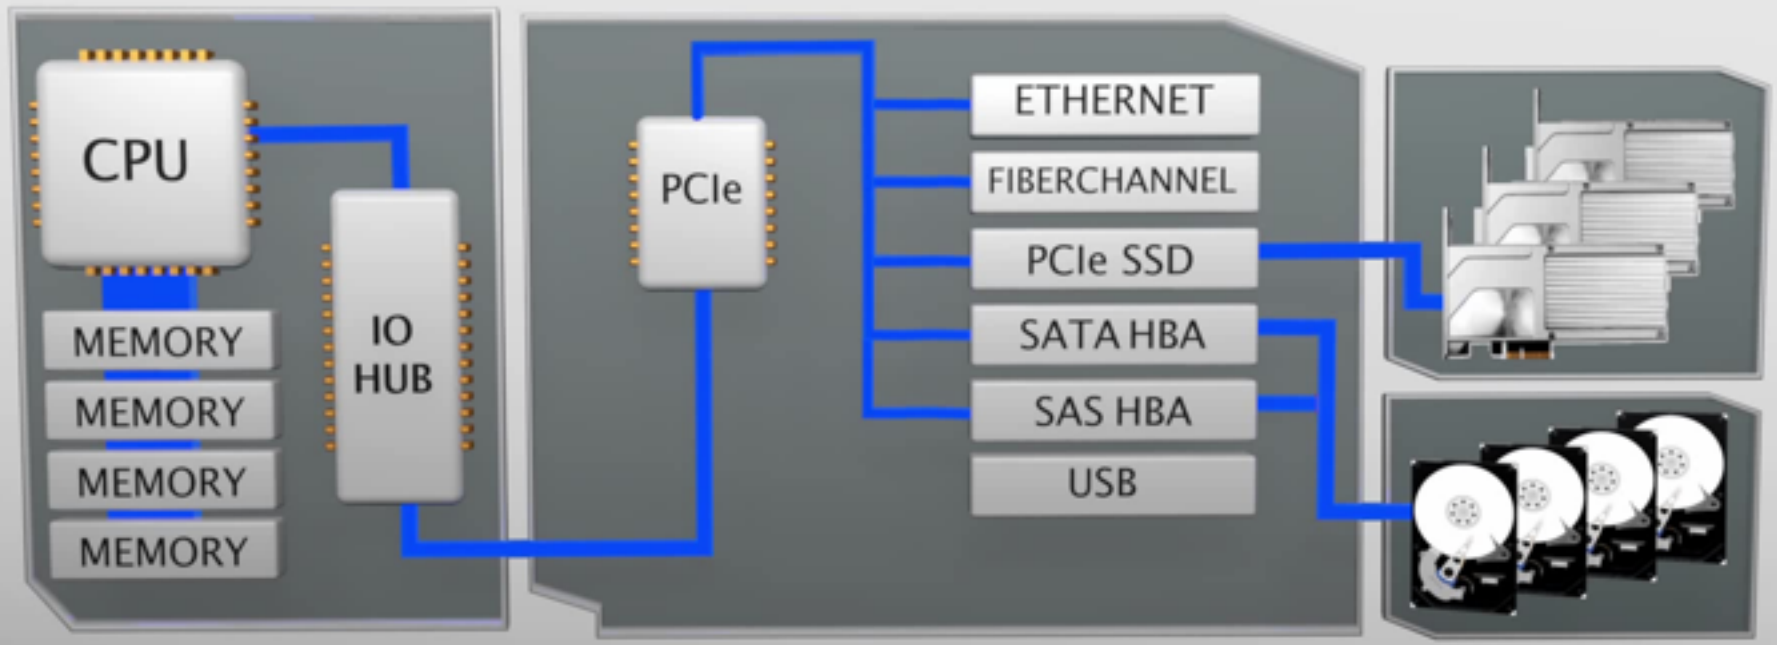

Connecting Flash to the Memory Bus

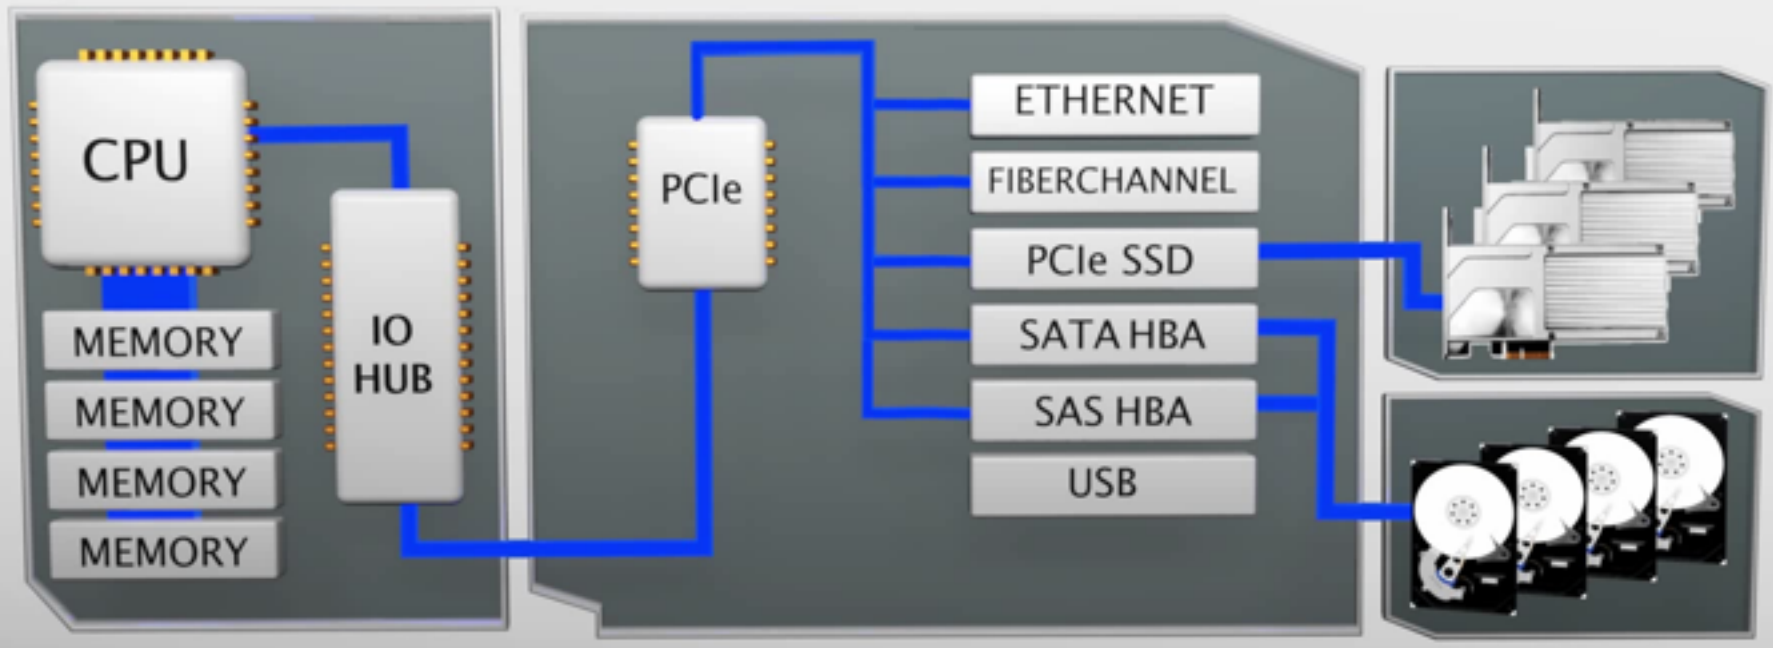

Connecting Flash to the Memory Bus

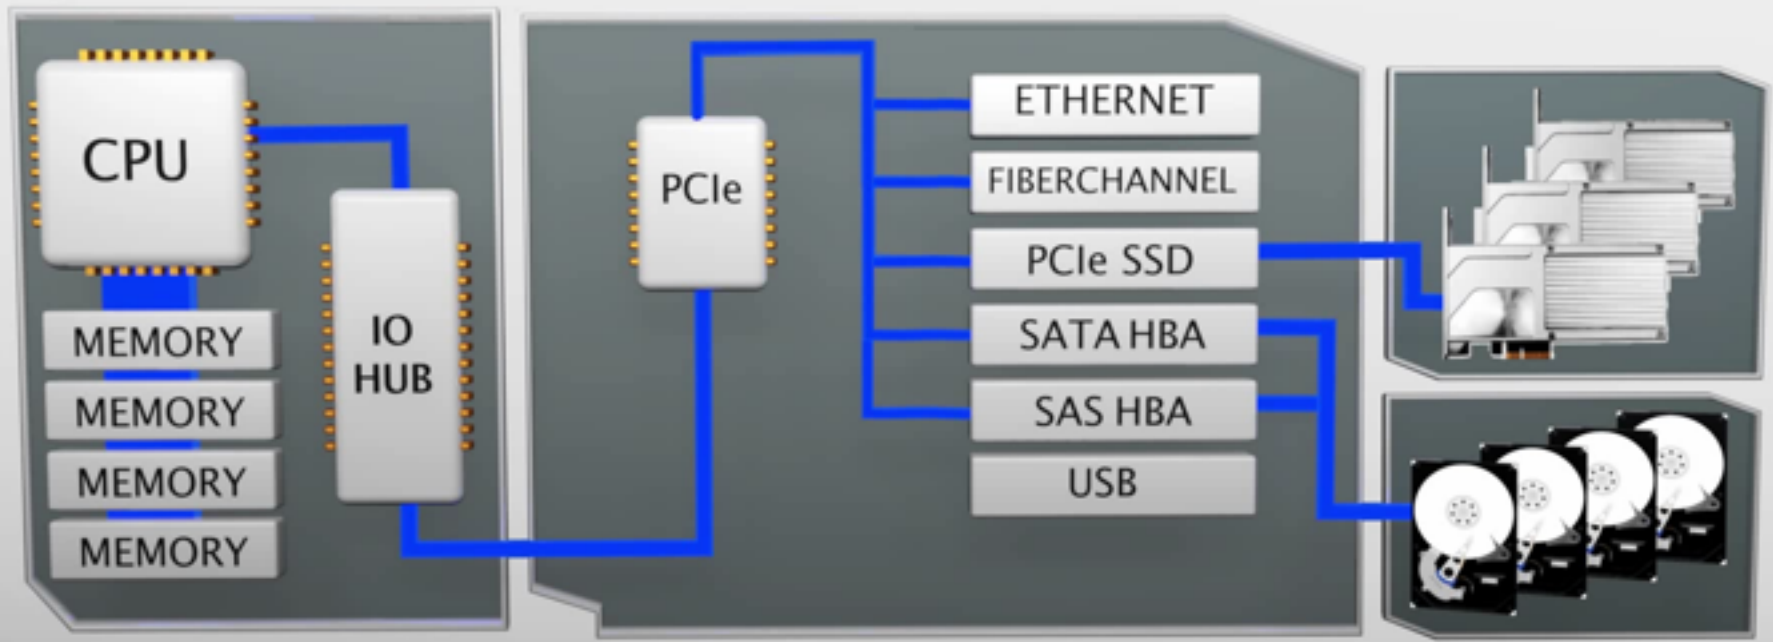

Connecting Flash to the Memory Bus eliminates contention and latency on the I/O hub

Storage Is Changing Too…

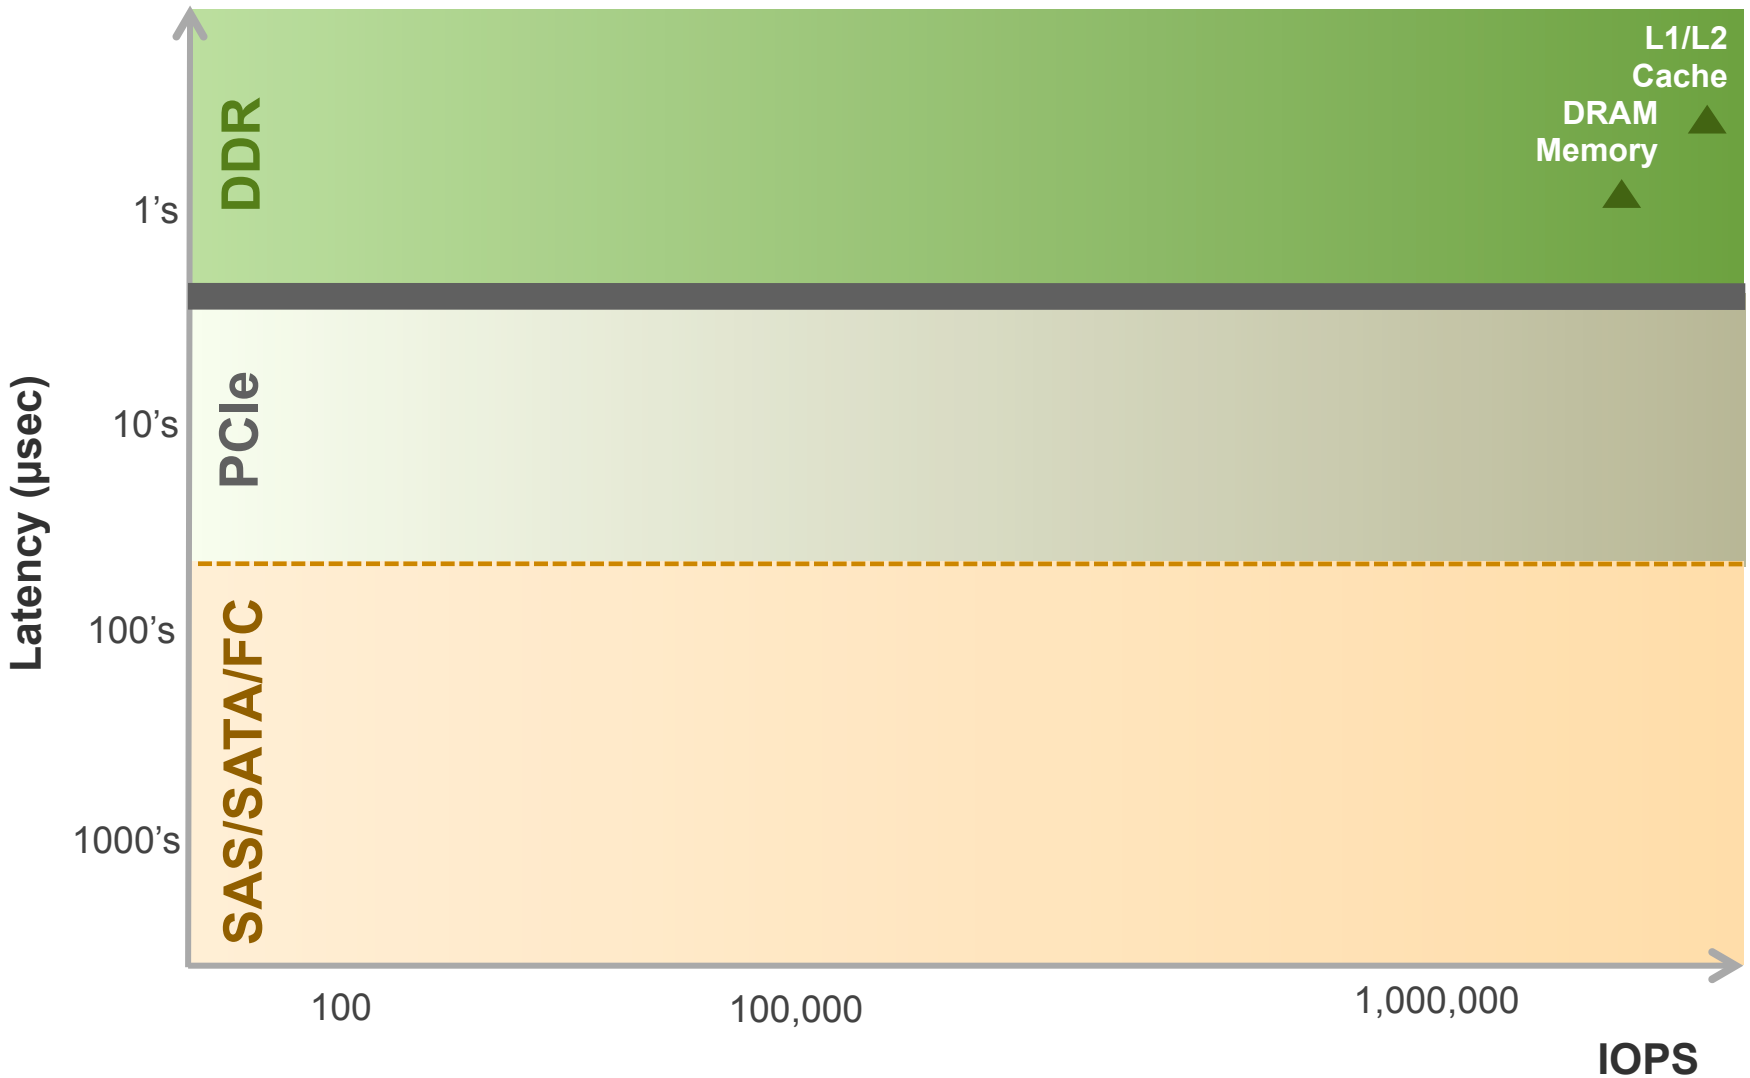

Storage Is Changing Too…

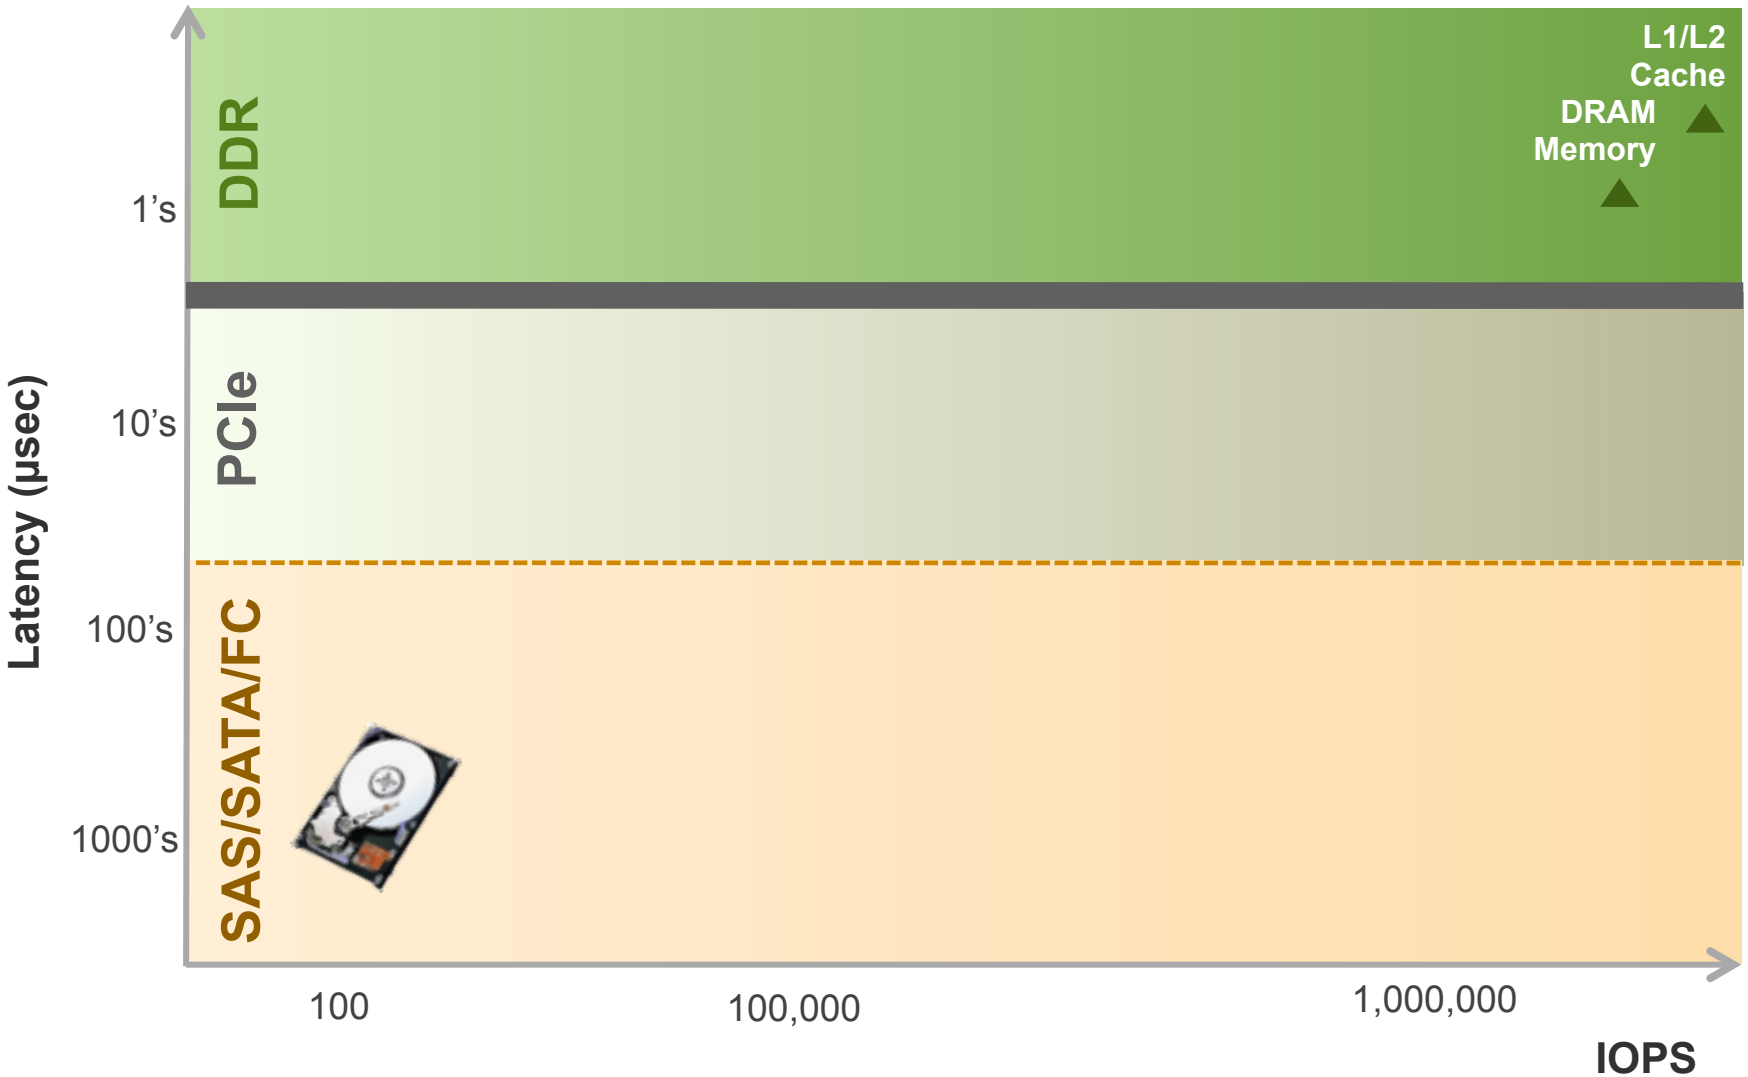

Storage Is Changing Too…

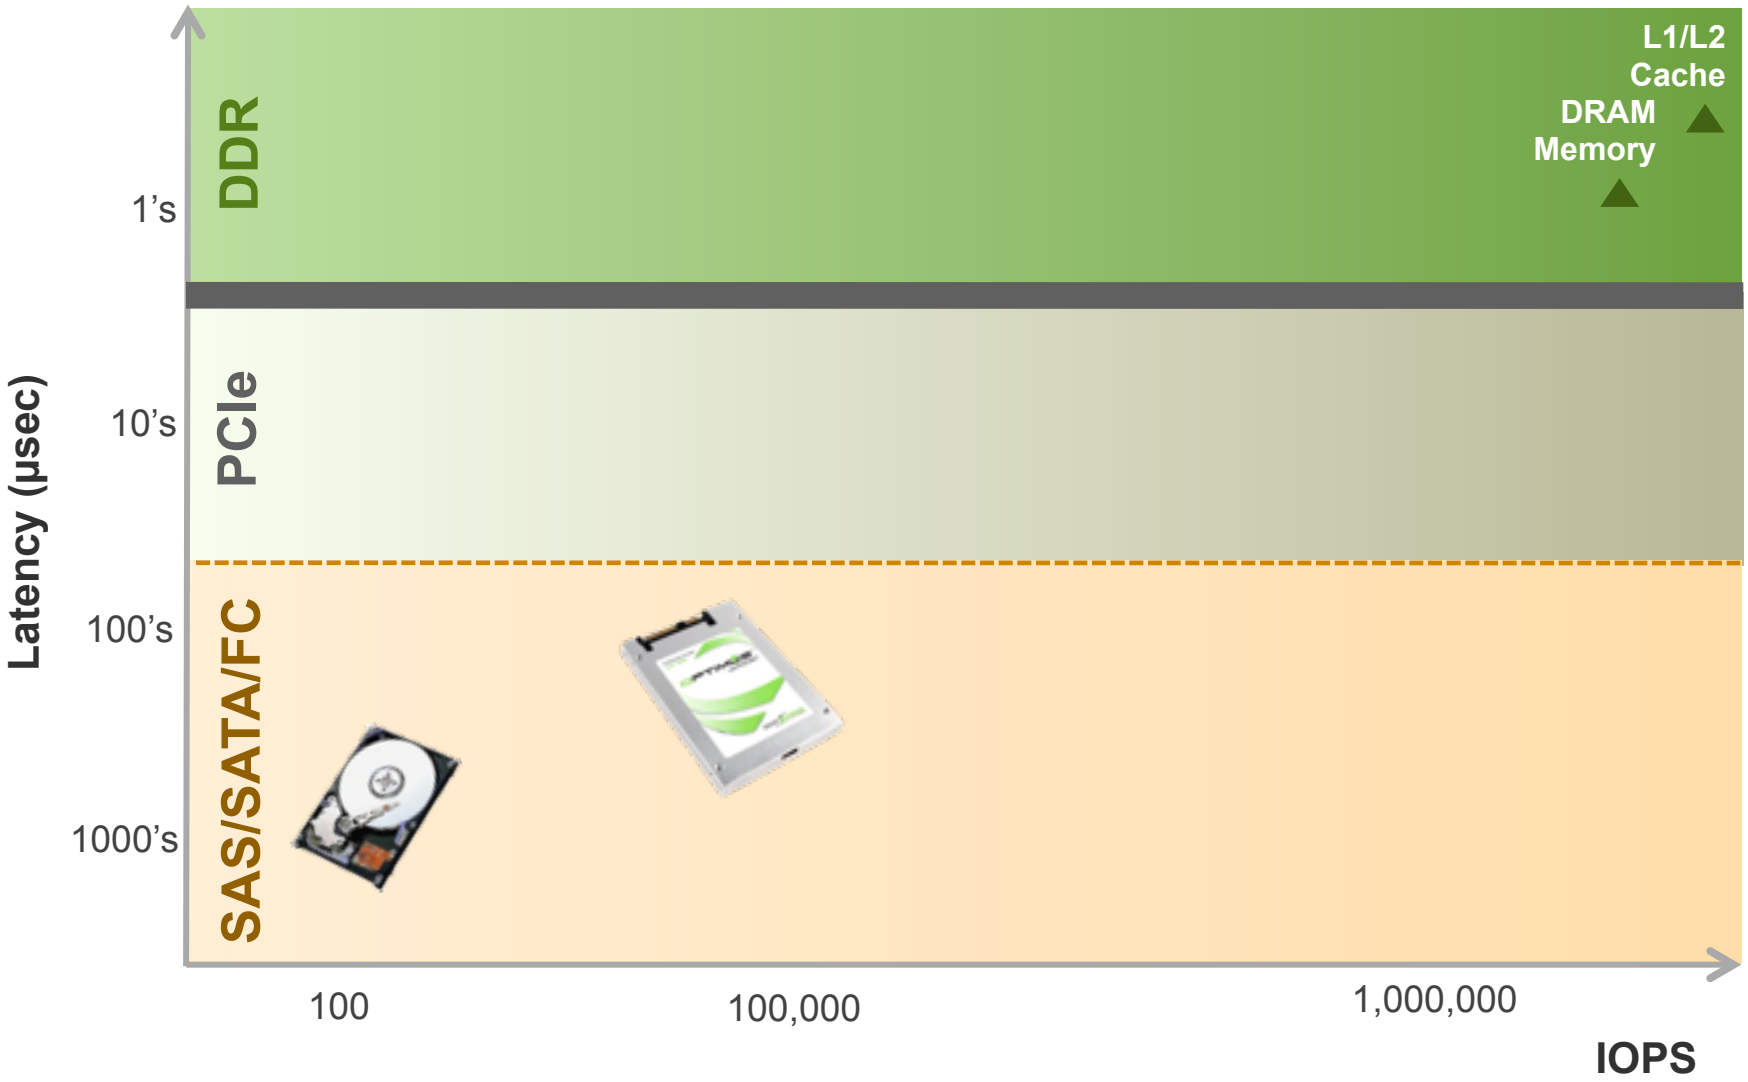

Storage Is Changing Too…

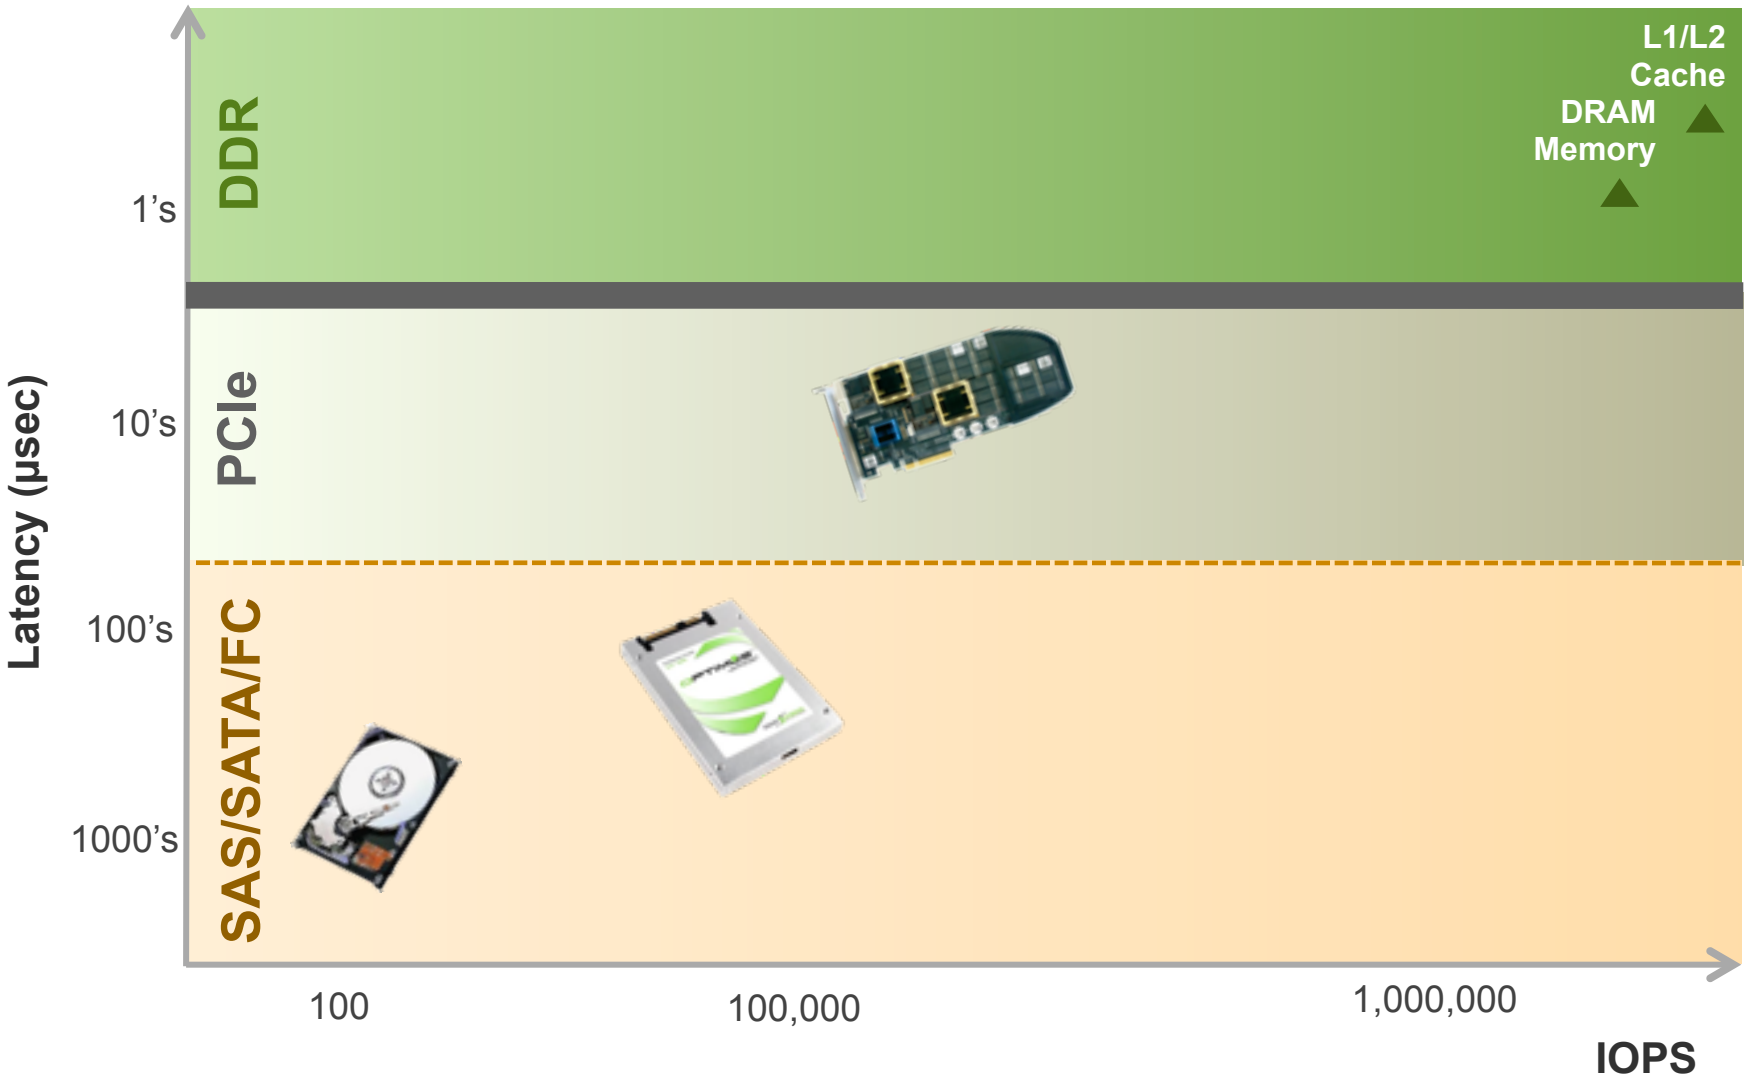

Storage Is Changing Too…

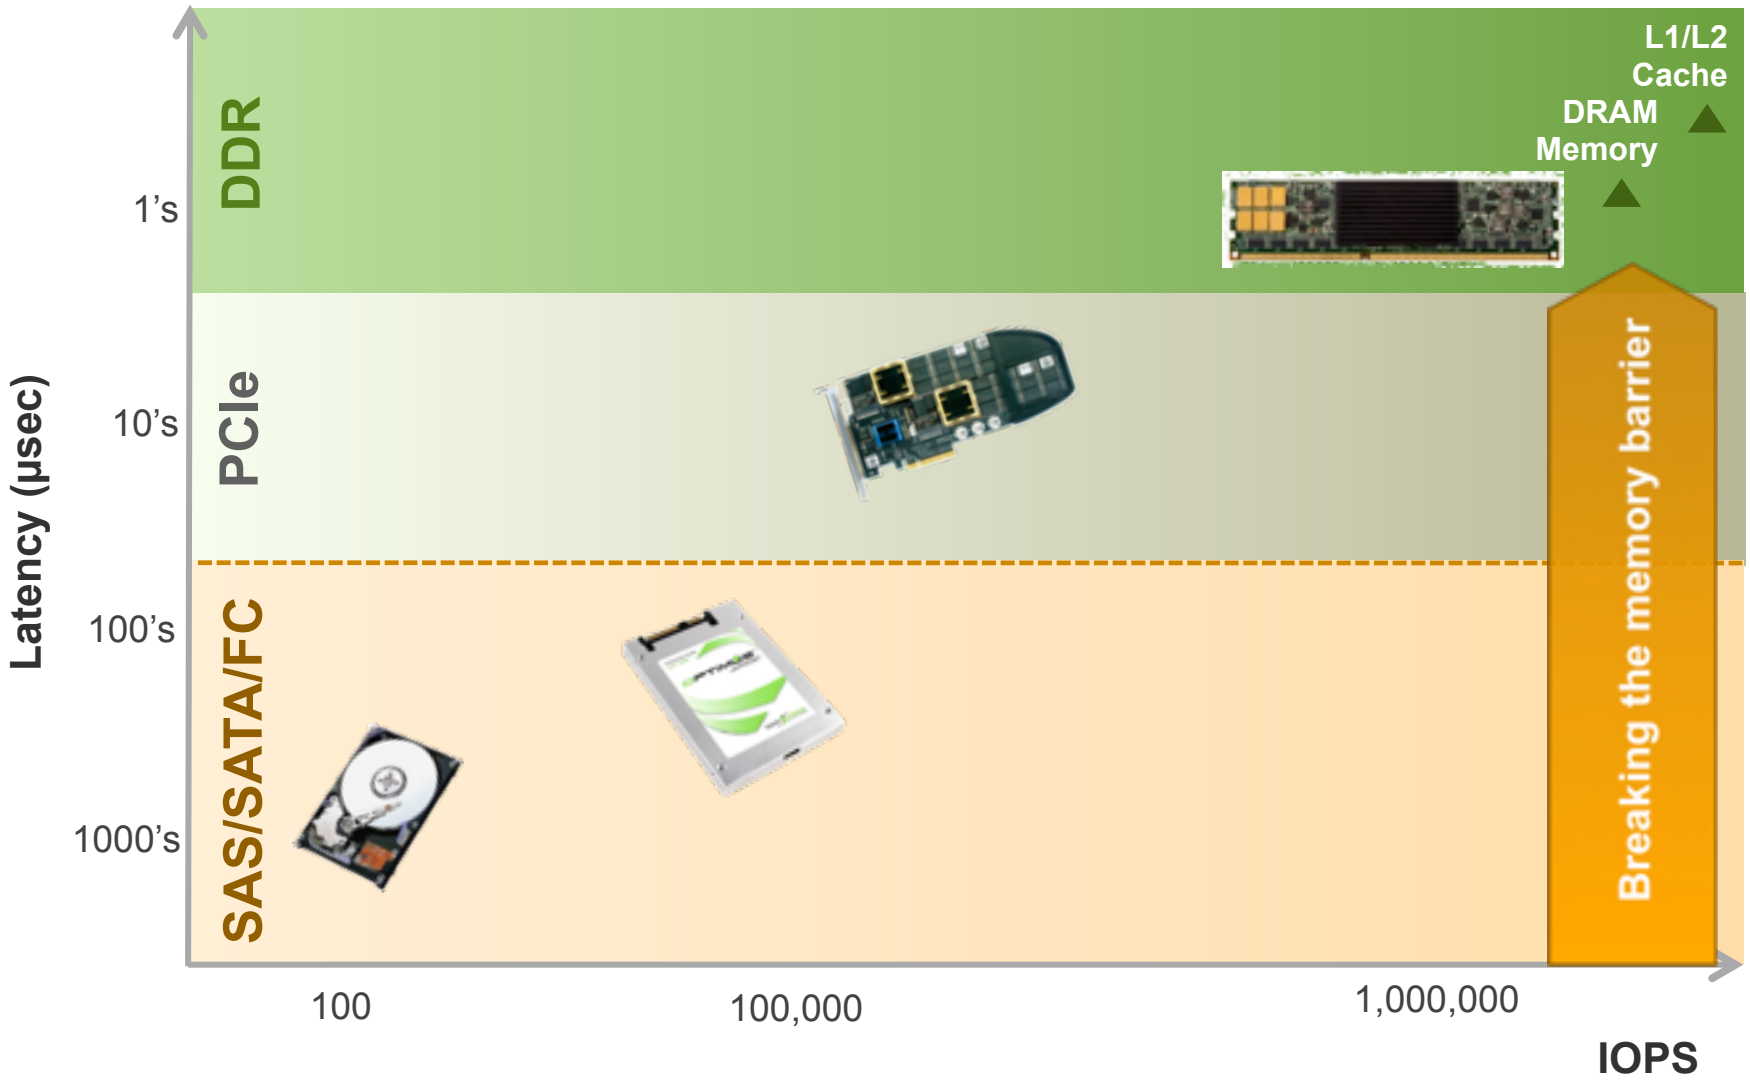

Scalable I/O Performance, Constant Latency

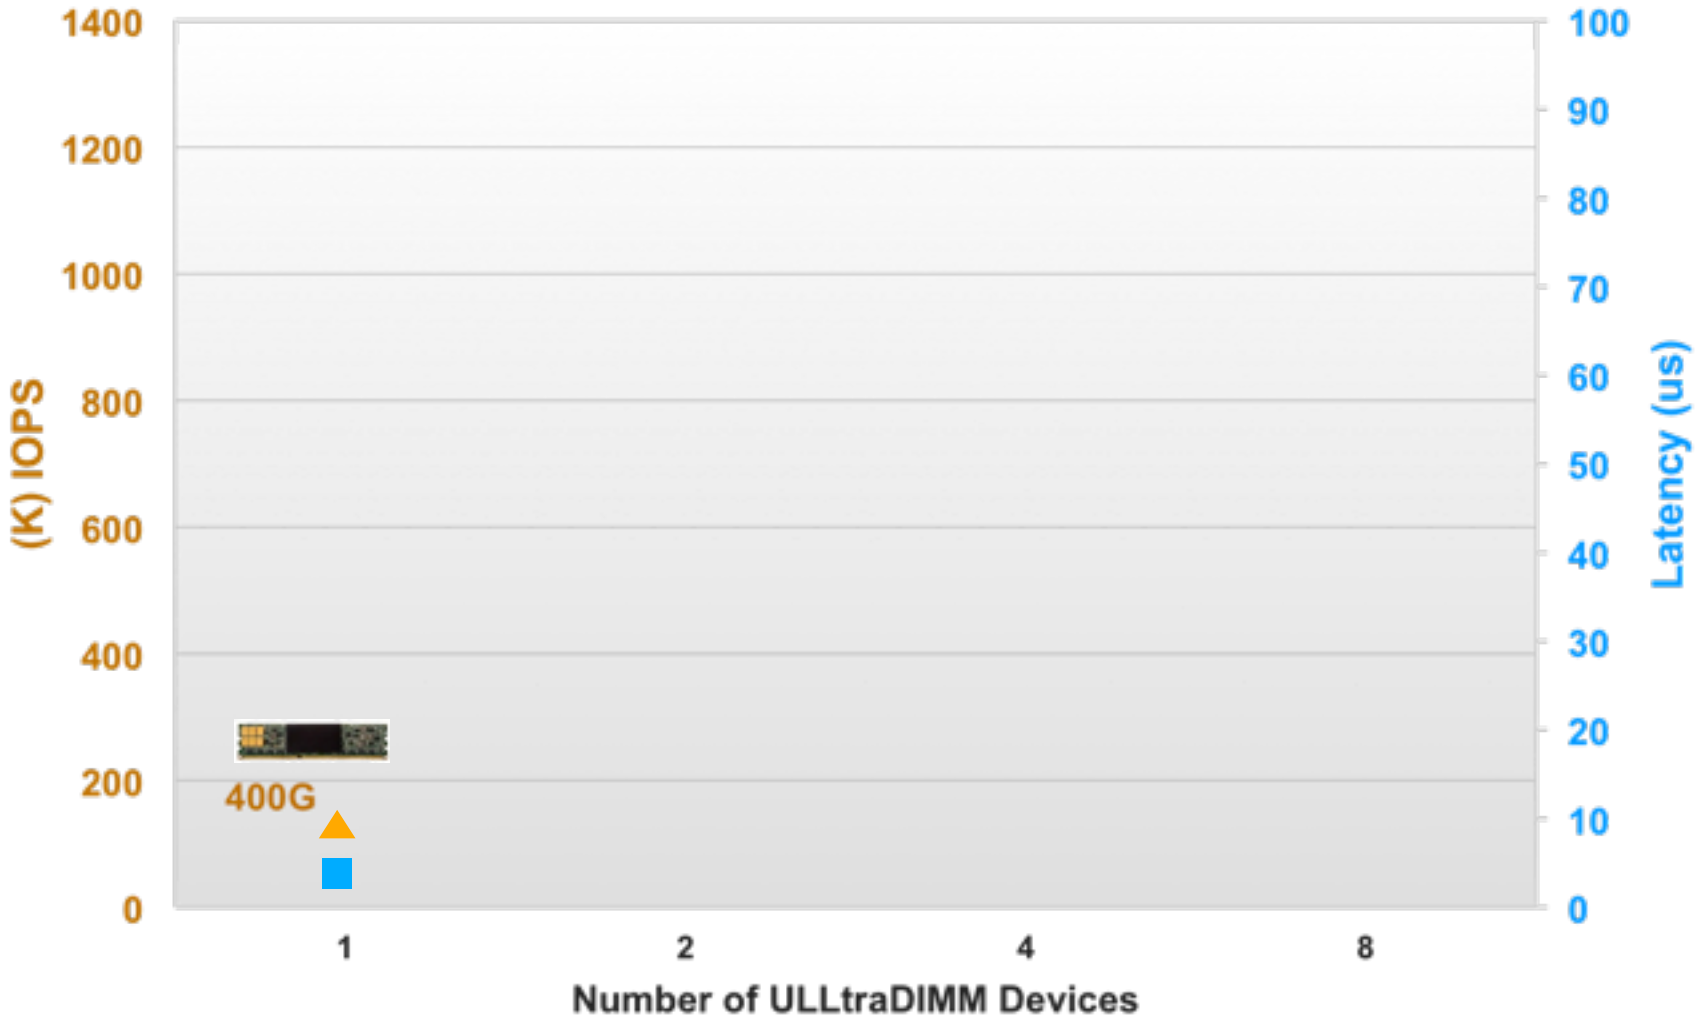

Scalable I/O Performance, Constant Latency

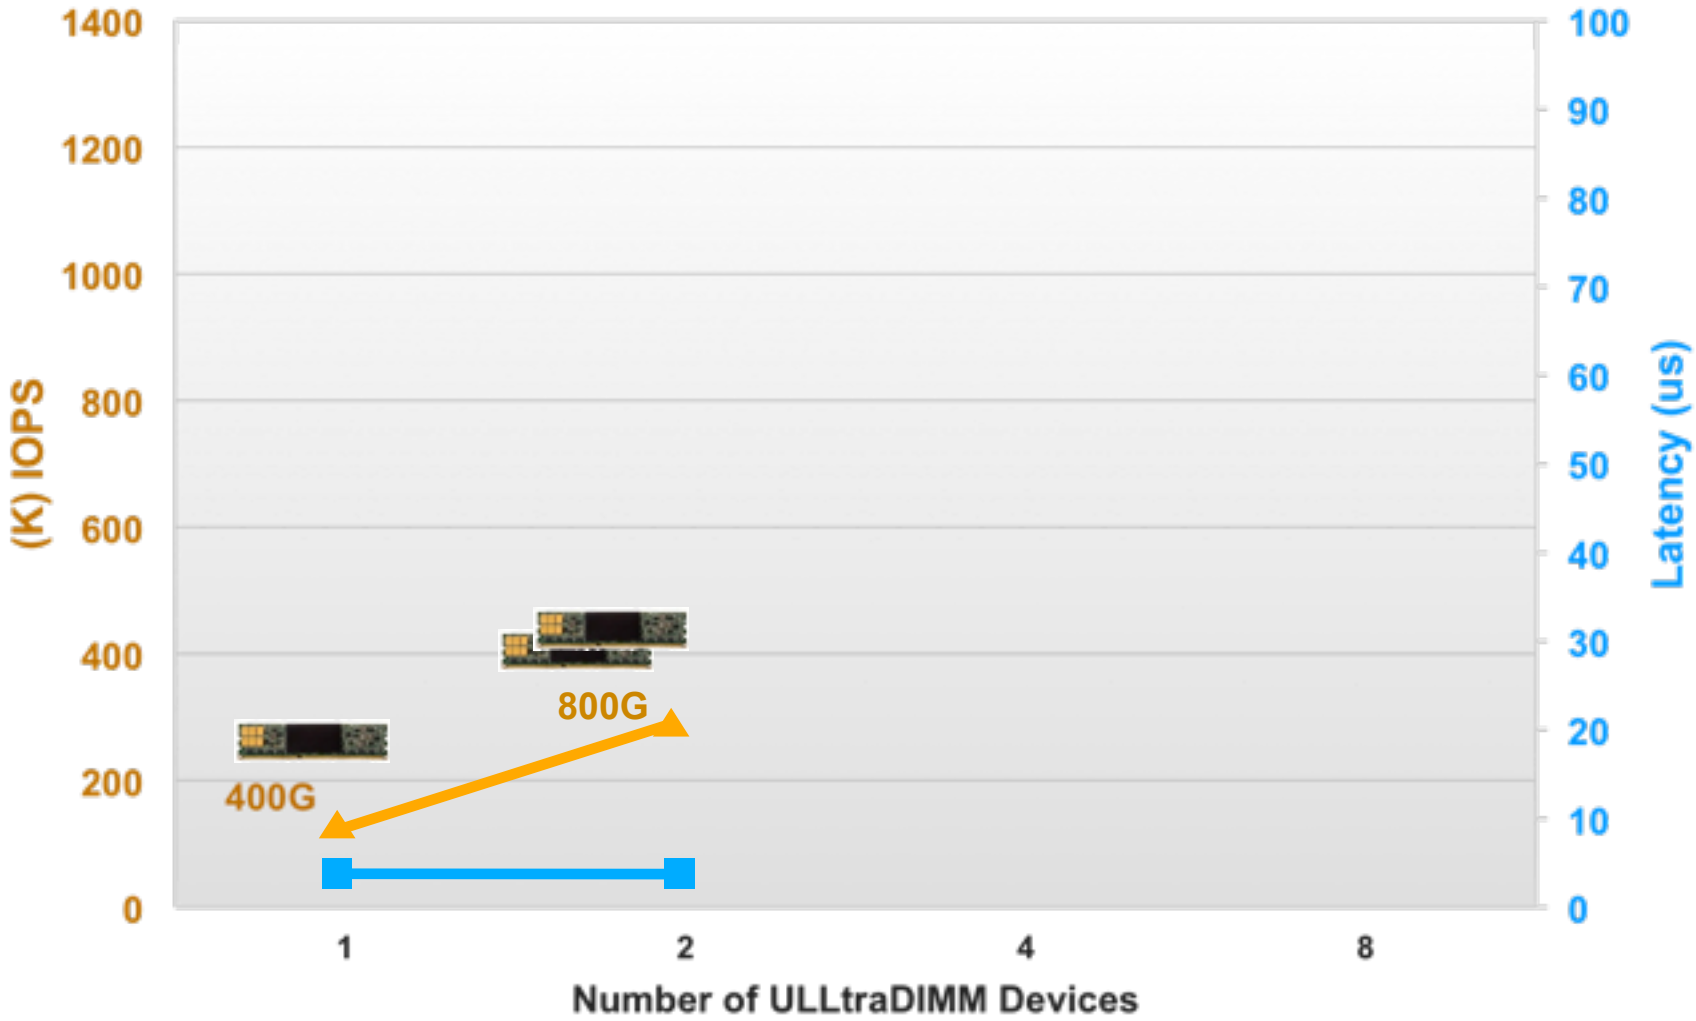

Scalable I/O Performance, Constant Latency

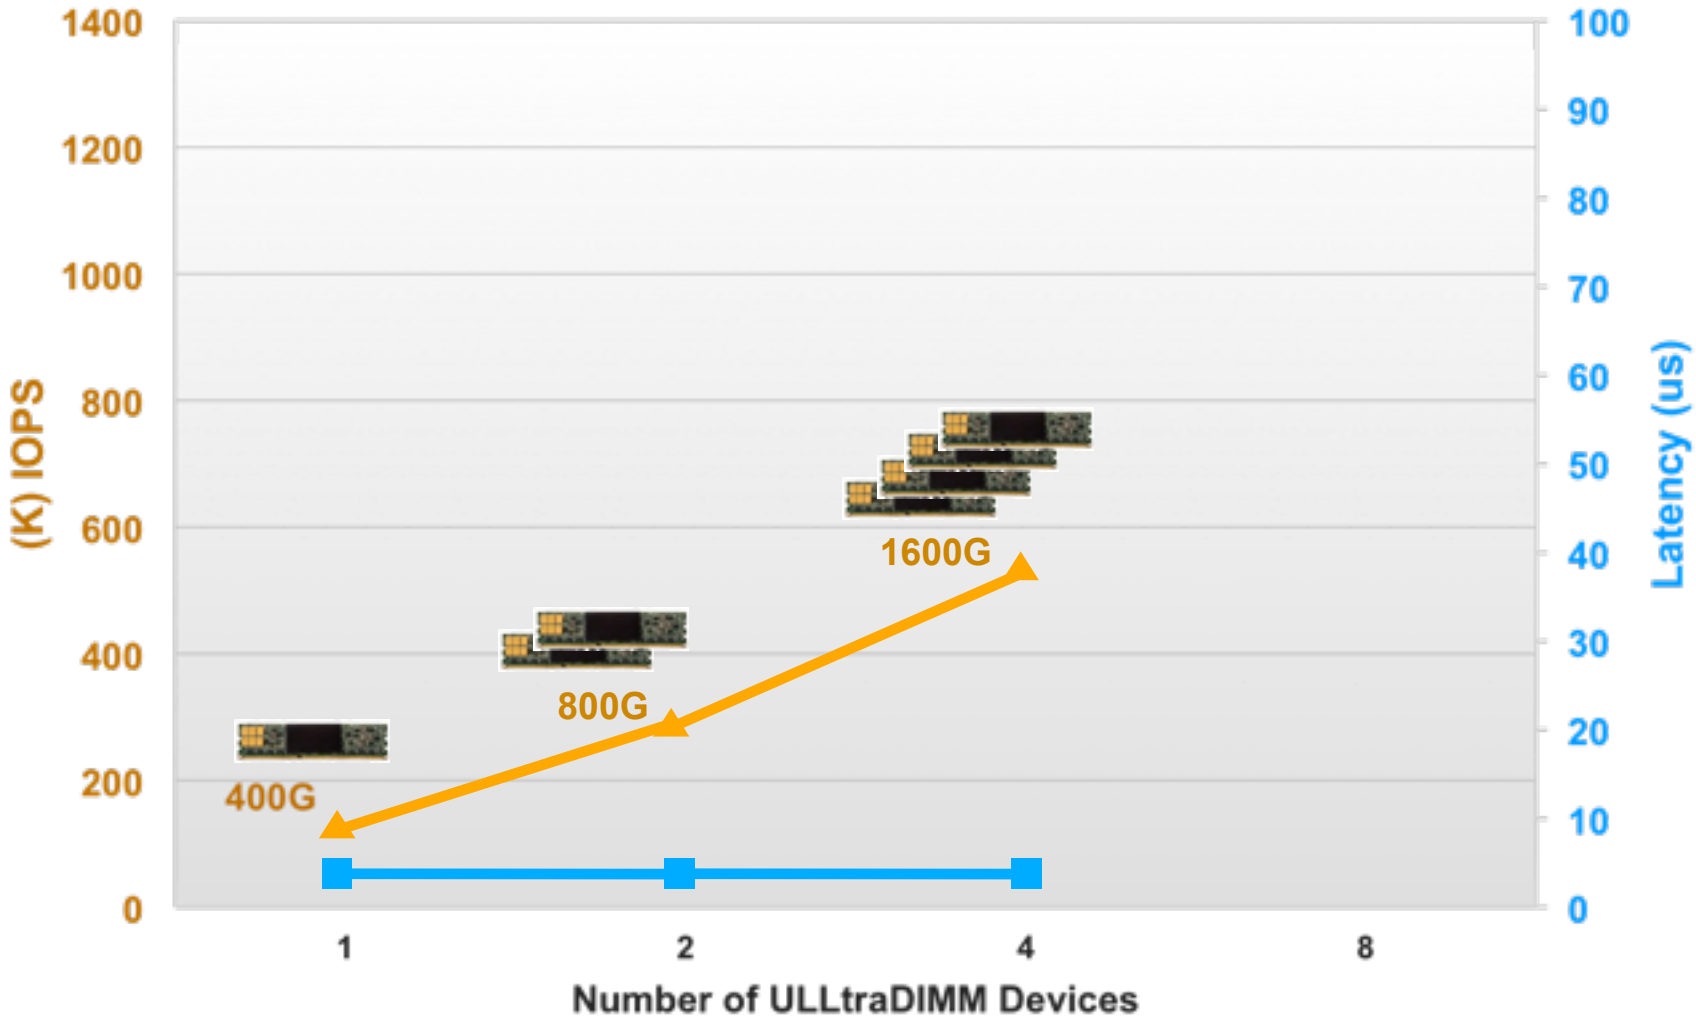

Scalable I/O Performance, Constant Latency

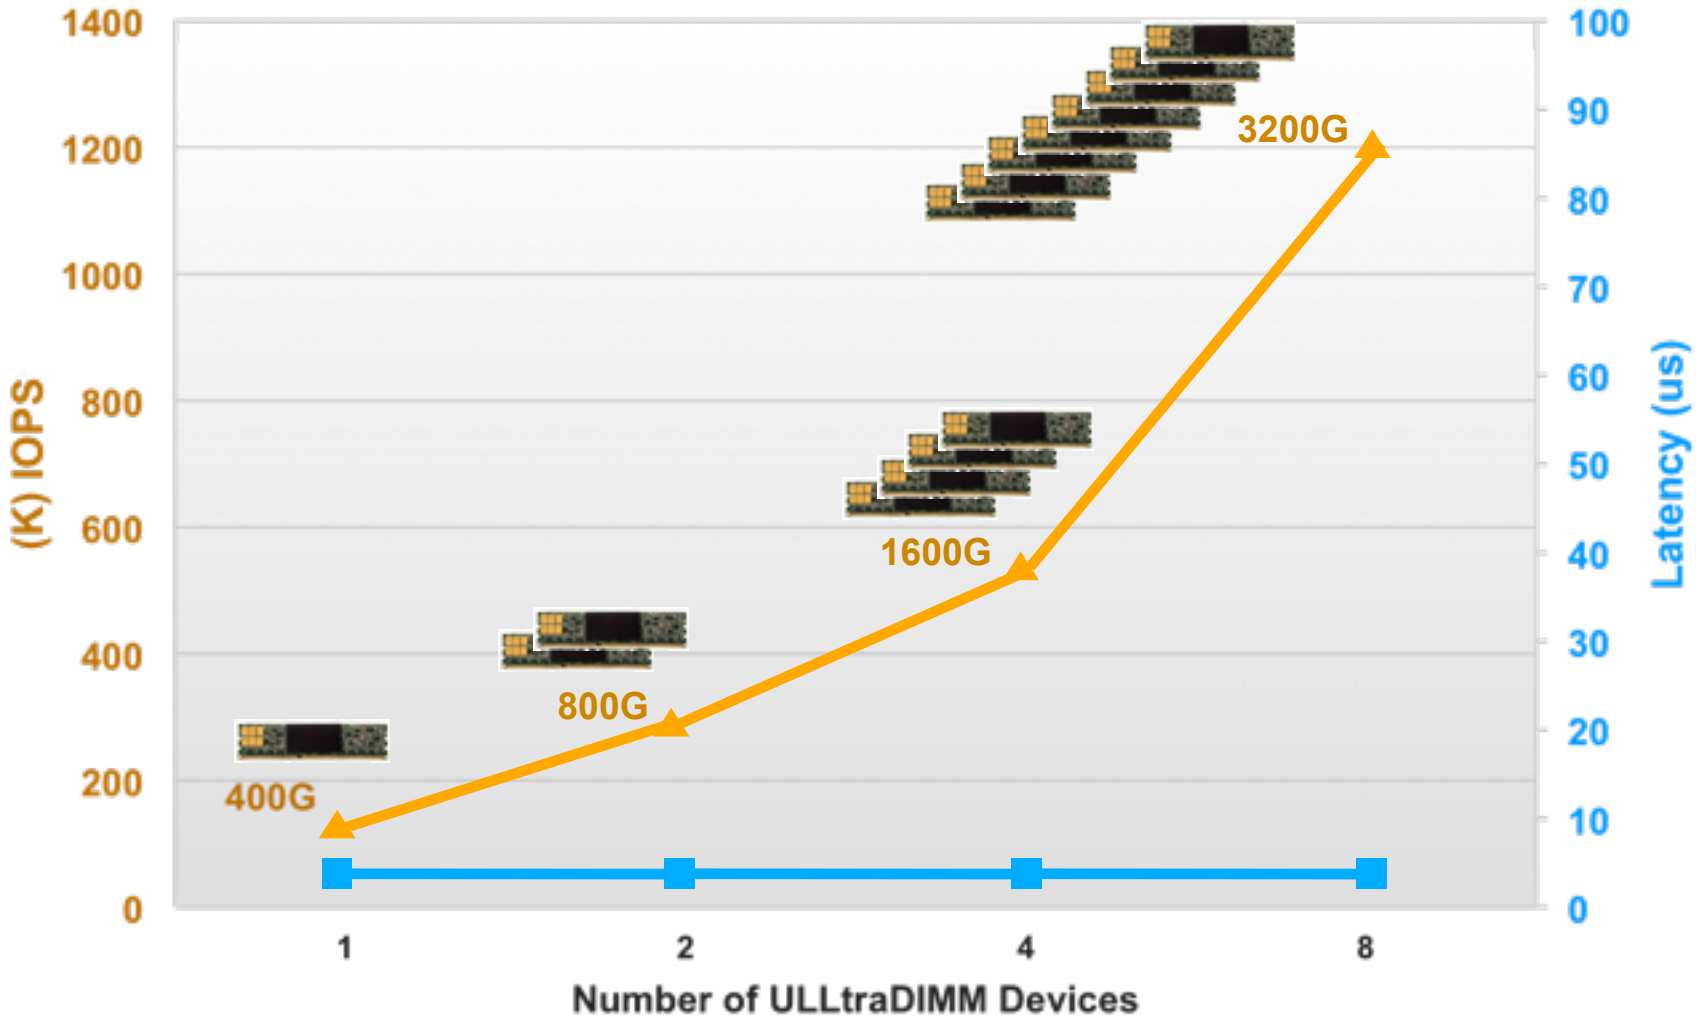

Performance alone isn't enough…

Key Storage Attributes

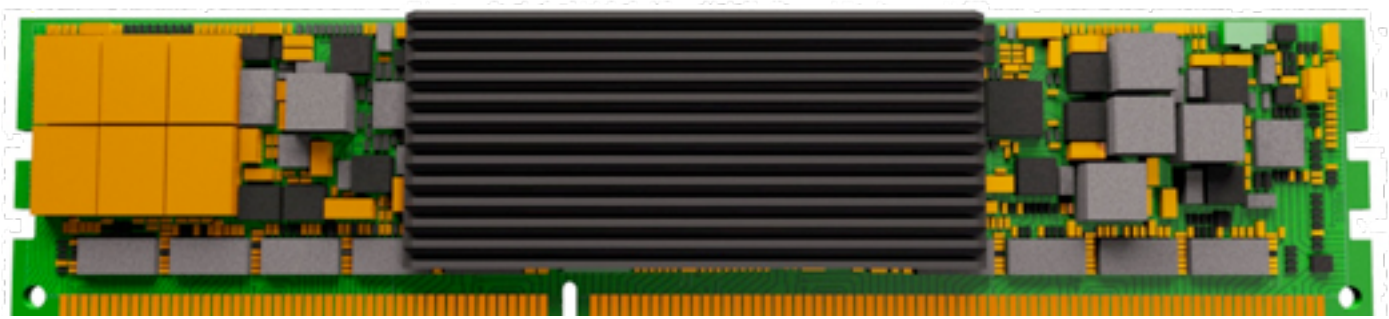

Performance alone isn't enough…

Key Storage Attributes

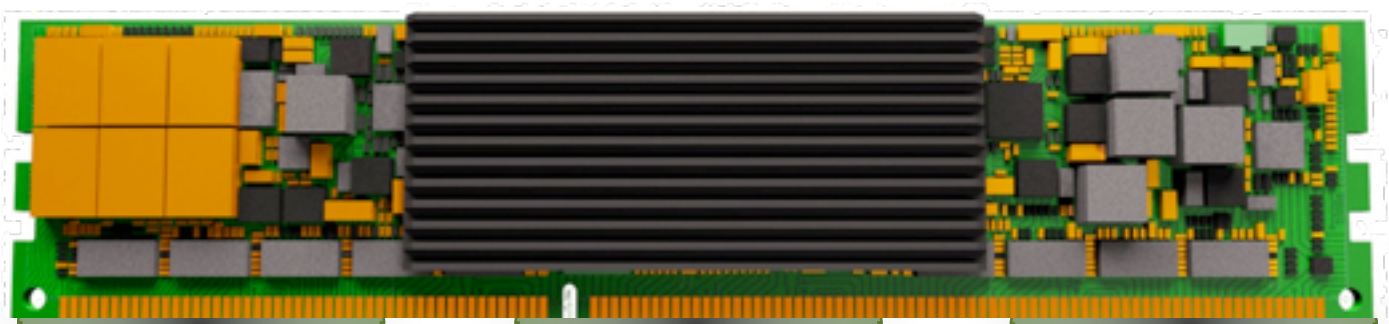

- Memory Channel Interface
- Enterprise Class Reliability, Endurance & Support
- Cost Effective

Utilizing Flash Over DDR3

Block Device Driver

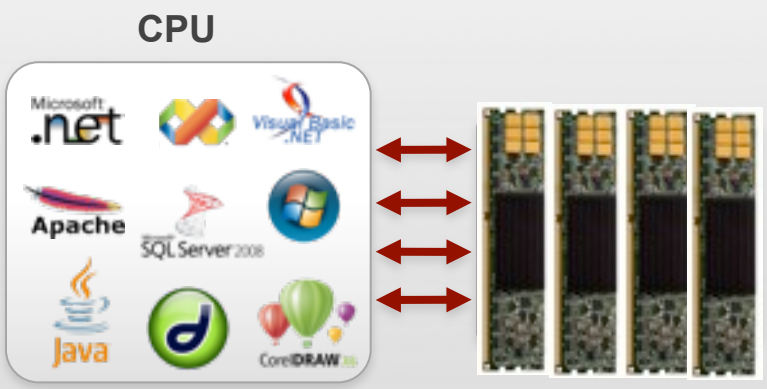

- Traditional storage I/O Device
- Use as fast, ultra low latency, persistent storage
- Scalable I/O performance, consistent latency

Utilizing Flash Over DDR3

Block Device Driver

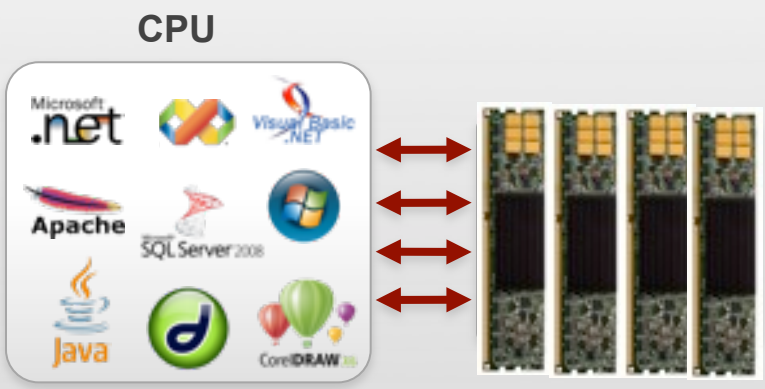

- Traditional storage I/O Device
- Use as fast, ultra low latency, persistent storage
- Scalable I/O performance, consistent latency

Memory Extension Driver

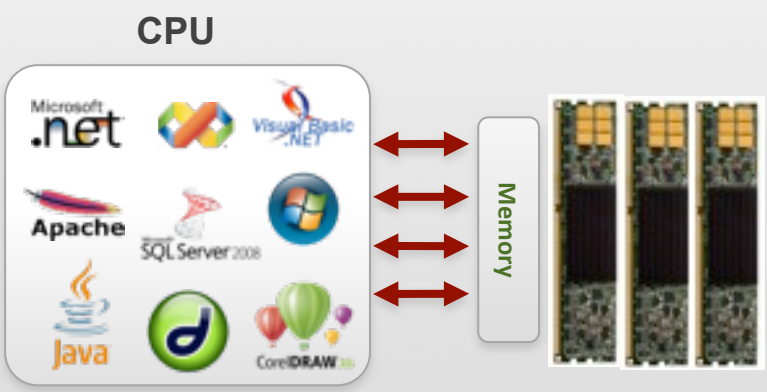

- Expands memory from tens of Gigabytes to Terabytes
- Ultra-fast paging at near DRAM performance, latency
- Allows large applications to run in memory space

Creating a New Storage Interface

Creating a New Storage Interface

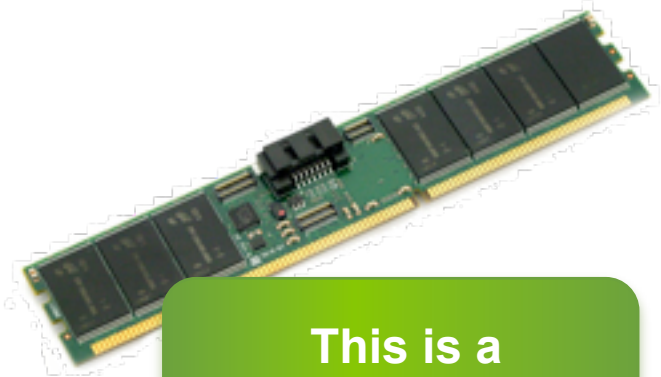

This is a
SATA DIMM

Creating a New Storage Interface

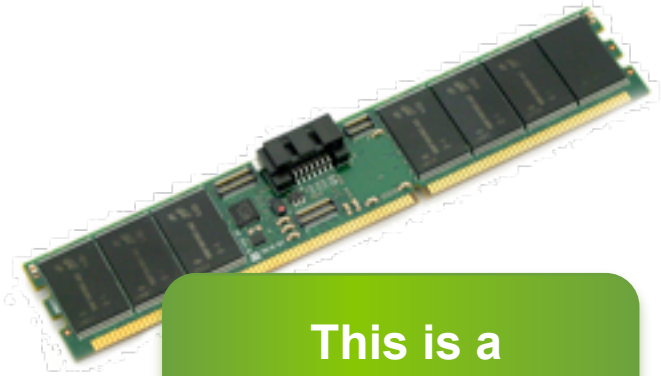

This is a SATA DIMM

This is DRAM with battery backup

Creating a New Storage Interface

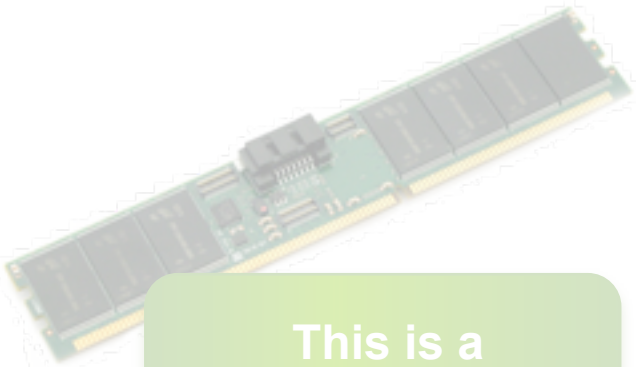

This is a SATA DIMM

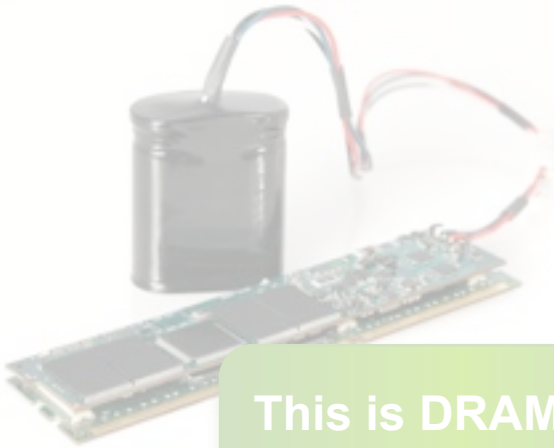

This is DRAM with battery backup

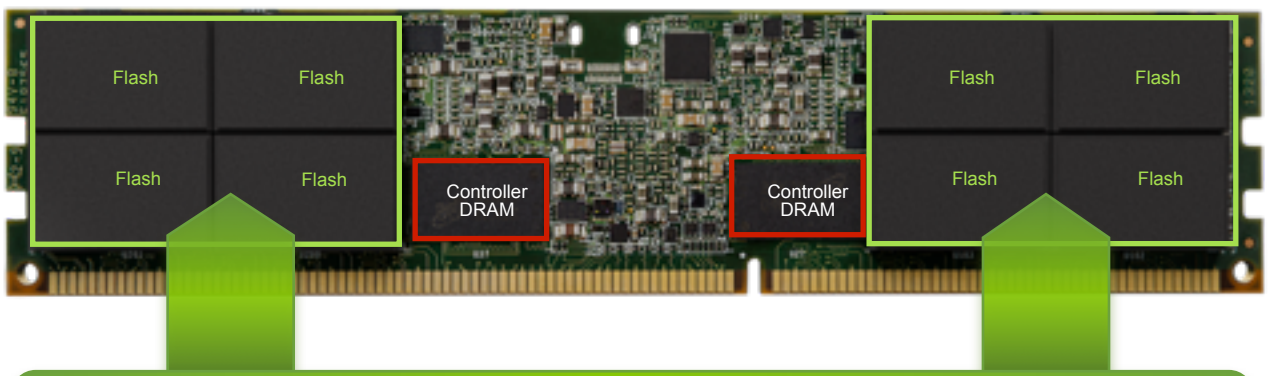

This is DDR3 directly to flash

Utilizing Existing Server Infrastructure

SMART STORAGE SYSTEMS

DDR3 interface & protocol

Scalable I/O & capacity

Utilizing Existing Server Infrastructure

DDR3 interface & protocol

Scalable I/O & capacity

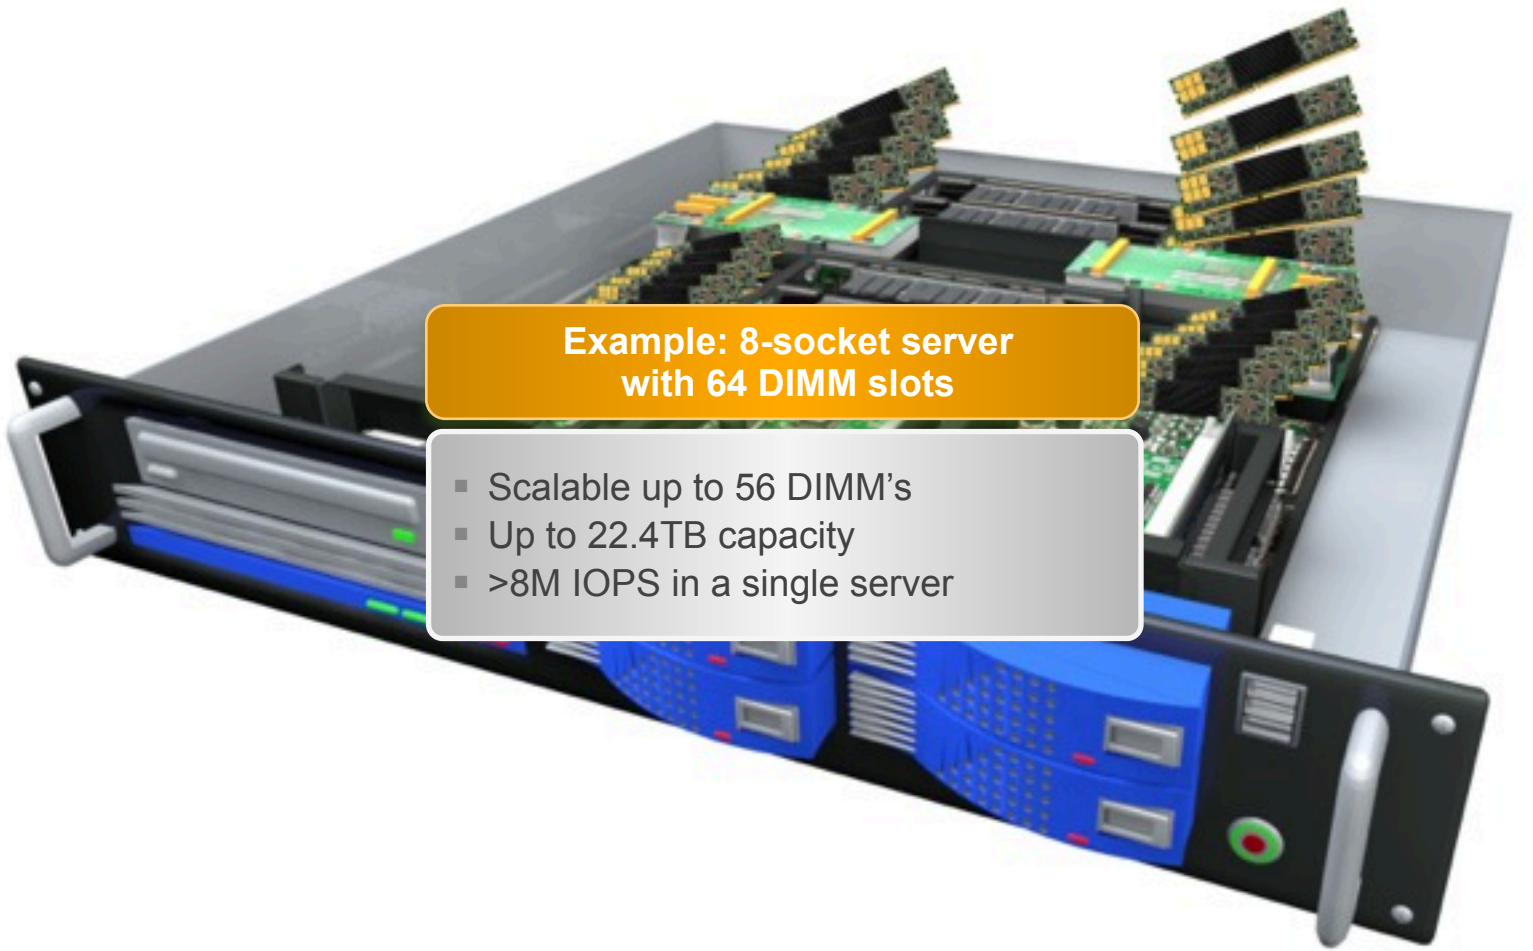

Example: 8-socket server with 64 DIMM slots

- Scalable up to 56 DIMM's
- Up to 22.4TB capacity
- >8M IOPS in a single server

Benefits To Applications

Benefits To Applications

Financial Services

- **Block device**
- Low, predictable latency
- Fast Interactive Data Analysis

Benefits To Applications

Financial Services

- **Block device**
- Low, predictable latency
- Fast Interactive Data Analysis

Database/ Cloud

- **Block device/ memory extension**
- Increase Transactions per Second
- Memcached consolidation

Benefits To Applications

Financial Services	Database/ Cloud	Virtualization
		
• **Block device** • Low, predictable latency • Fast Interactive Data Analysis	• **Block device/ memory extension** • Increase Transactions per Second • Memcached consolidation	• **Block device** • Increased VMs per Node • Faster response times per VM

Benefits To Applications

Financial Services	Database/ Cloud	Virtualization	Blade Server
			
• **Block device** • Low, predictable latency • Fast Interactive Data Analysis	• **Block device/ memory extension** • Increase Transactions per Second • Memcached consolidation	• **Block device** • Increased VMs per Node • Faster response times per VM	• **Block device** • Utilizes empty DIMM slots • Enables high density storage blades

Benefits To Applications

Financial Services	Database/ Cloud	Virtualization	Blade Server	In Memory Compute
				
• **Block device** • Low, predictable latency • Fast Interactive Data Analysis	• **Block device/ memory extension** • Increase Transactions per Second • Memcached consolidation	• **Block device** • Increased VMs per Node • Faster response times per VM	• **Block device** • Utilizes empty DIMM slots • Enables high density storage blades	• **Memory extension** • Reduce response times for analytics queries

Example: High Frequency Trading

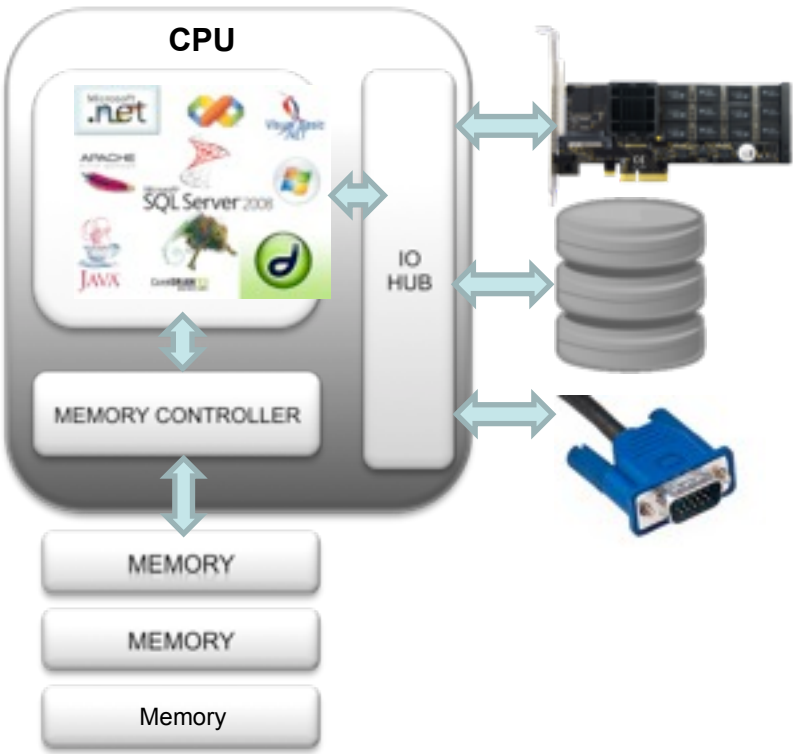

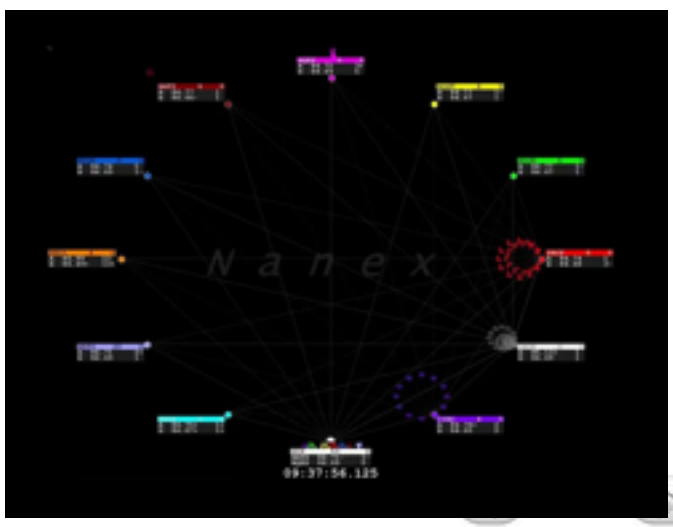

Example: High Frequency Trading

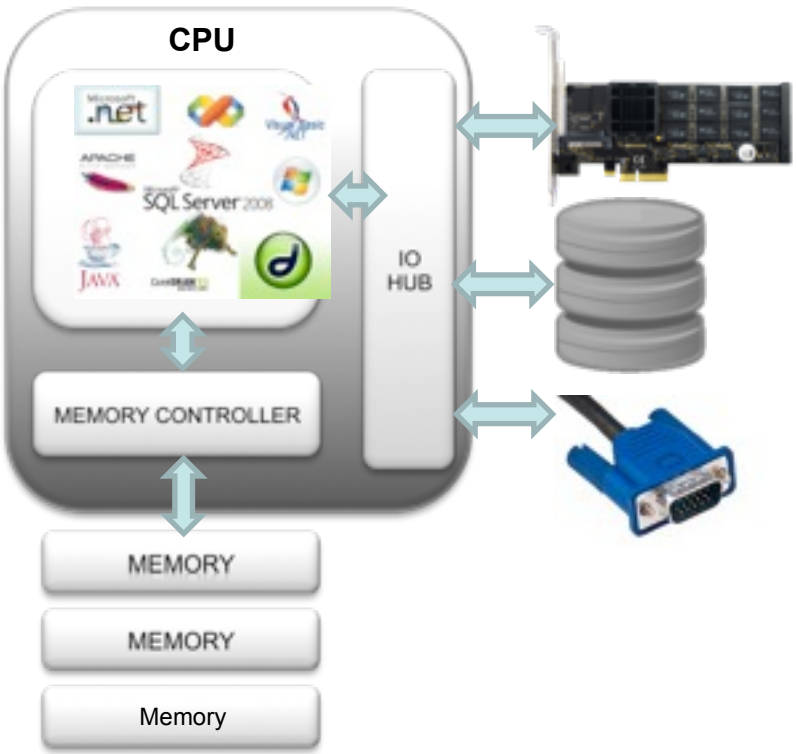

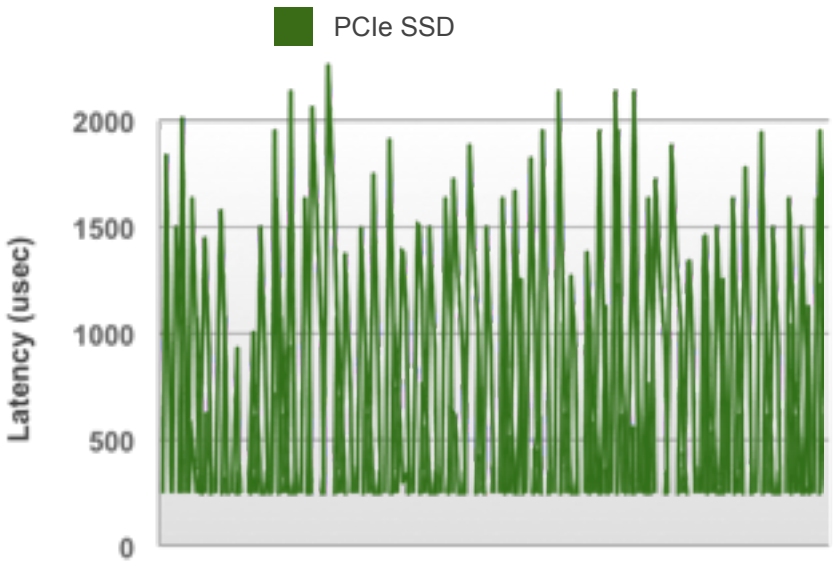

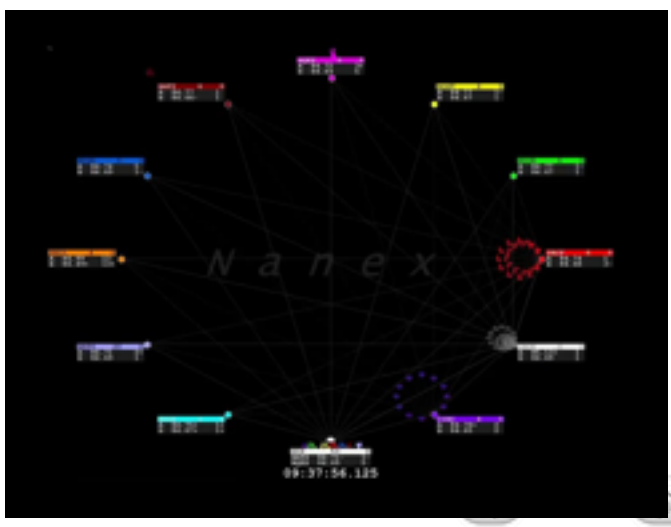

Example: High Frequency Trading

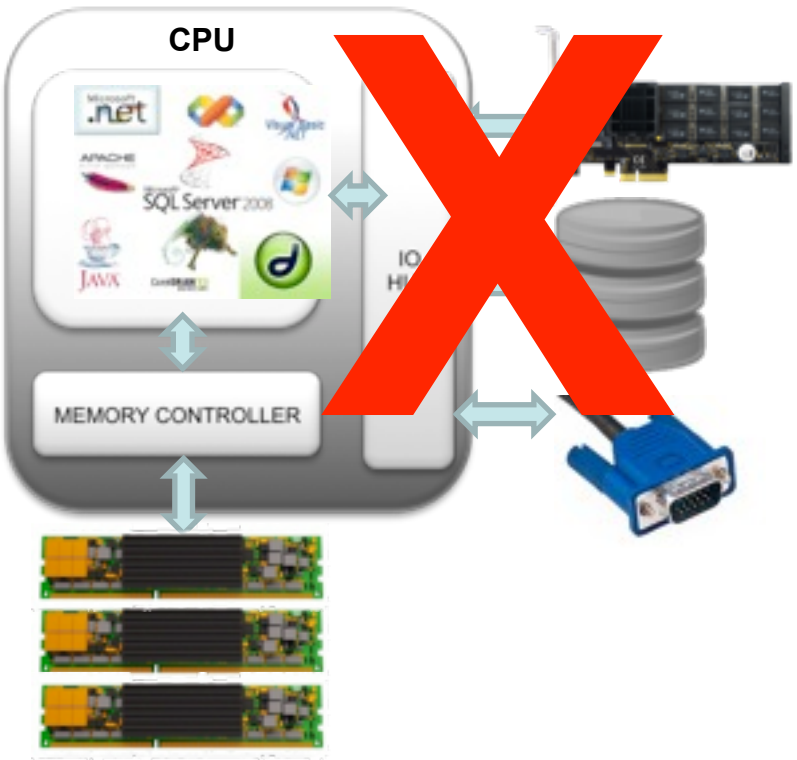

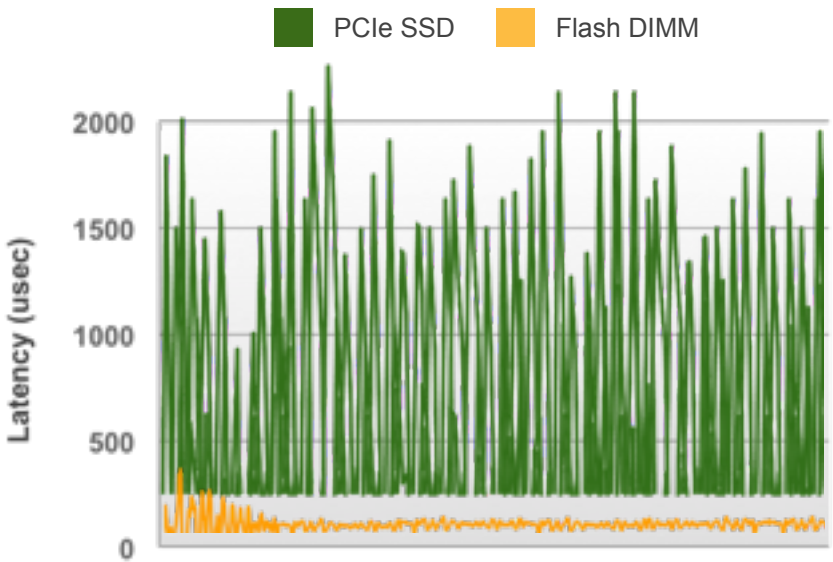

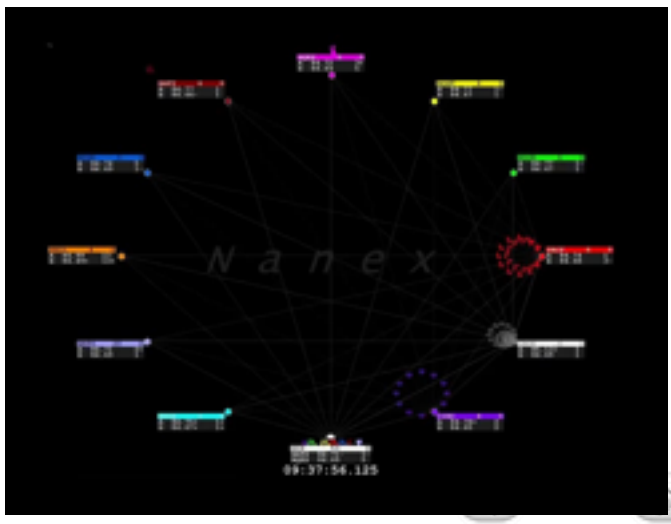

Example: High Frequency Trading

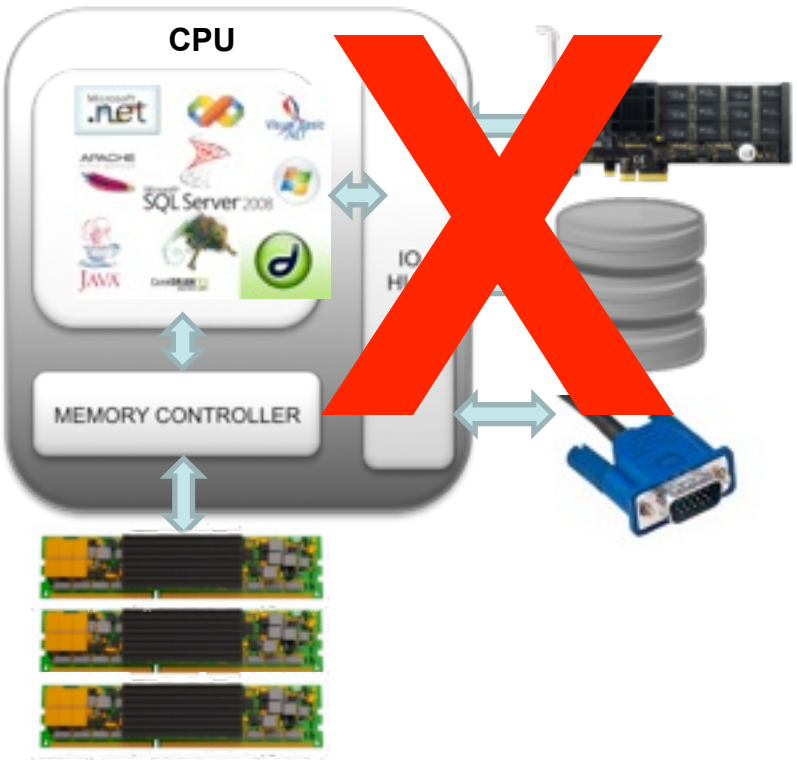

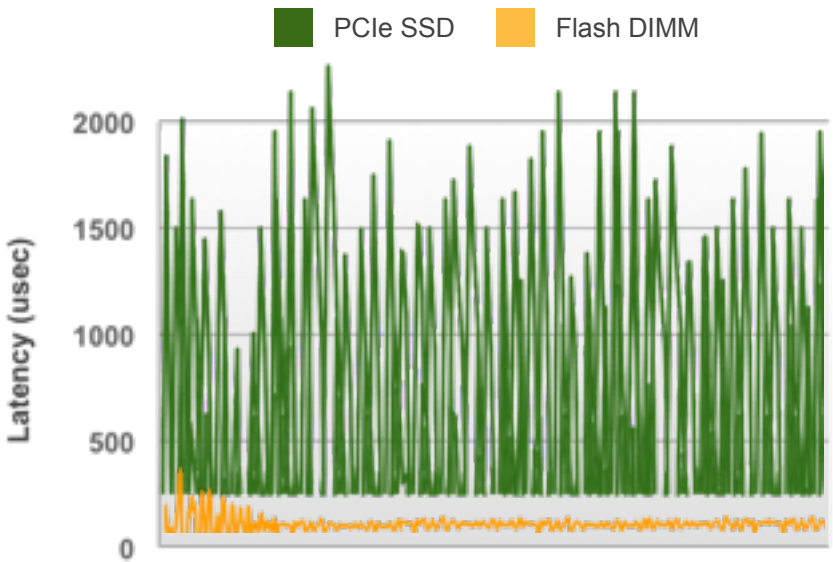

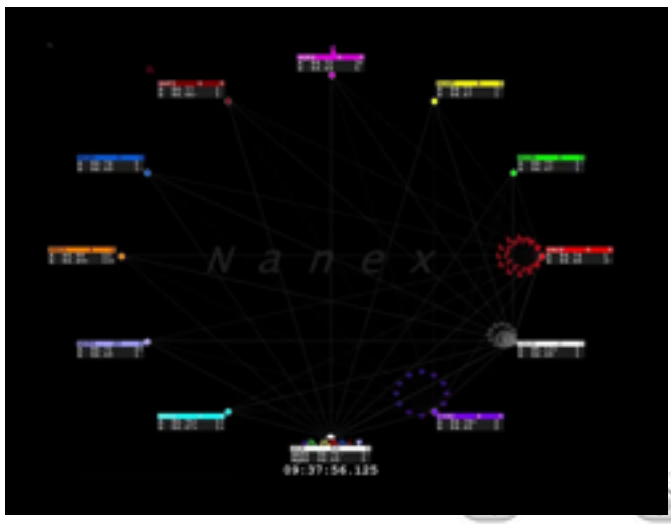

Predictable, deterministic latency will win the trade

Example: Blade Servers

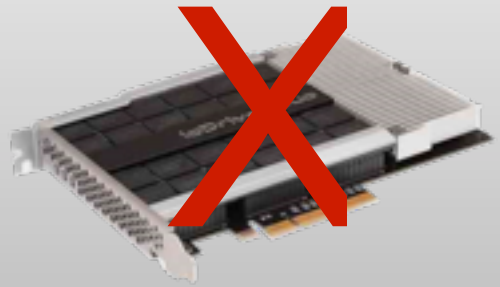

- Requires dedicated blade server slot
- Decreases server efficiency
- Custom Form Factor

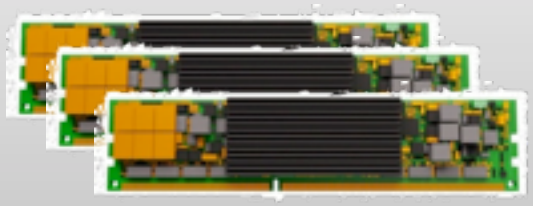

- Enables high density, high performance storage
- Block and memory expansion storage
- Improves blade server efficiency

Example: Blade Servers

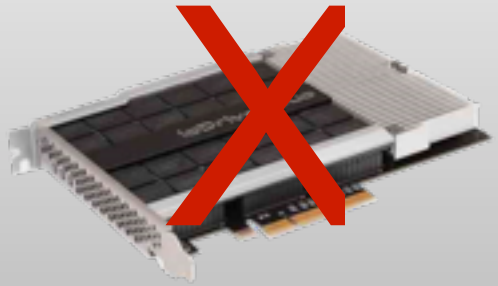

- Requires dedicated blade server slot
- Decreases server efficiency
- Custom Form Factor

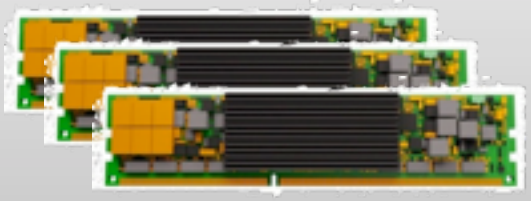

- Enables high density, high performance storage
- Block and memory expansion storage
- Improves blade server efficiency

Example: In Memory Compute

Example: In Memory Compute

SMART STORAGE SYSTEMS

Events Stream Benchmark

- Streams recorded data of a real event into SAP HANA database
- 500 inserts per transaction, 60 concurrent connections
- 15 minutes for storage

Example: In Memory Compute

SMART STORAGE SYSTEMS

Events Stream Benchmark

- Streams recorded data of a real event into SAP HANA database
- 500 inserts per transaction, 60 concurrent connections
- 15 minutes for storage

FlashDIMM: 495,397 inserts

Example: In Memory Compute

Events Stream Benchmark

- Streams recorded data of a real event into SAP HANA database
- 500 inserts per transaction, 60 concurrent connections
- 15 minutes for storage

FlashDIMM: 495,397 inserts

48%

PCIe SSD: 346,894 inserts

Example: In Memory Compute

Events Stream Benchmark

- Streams recorded data of a real event into SAP HANA database
- 500 inserts per transaction, 60 concurrent connections
- 15 minutes for storage

RAM Disk: 518,114 inserts

-5%

FlashDIMM: 495,397 inserts

48%

PCIe SSD: 346,894 inserts

Example: Virtual Desktop Infrastructure

Deploying 12,500 Virtual Desktops

Example: Virtual Desktop Infrastructure

Deploying 12,500 Virtual Desktops

75 Virtual Desktops

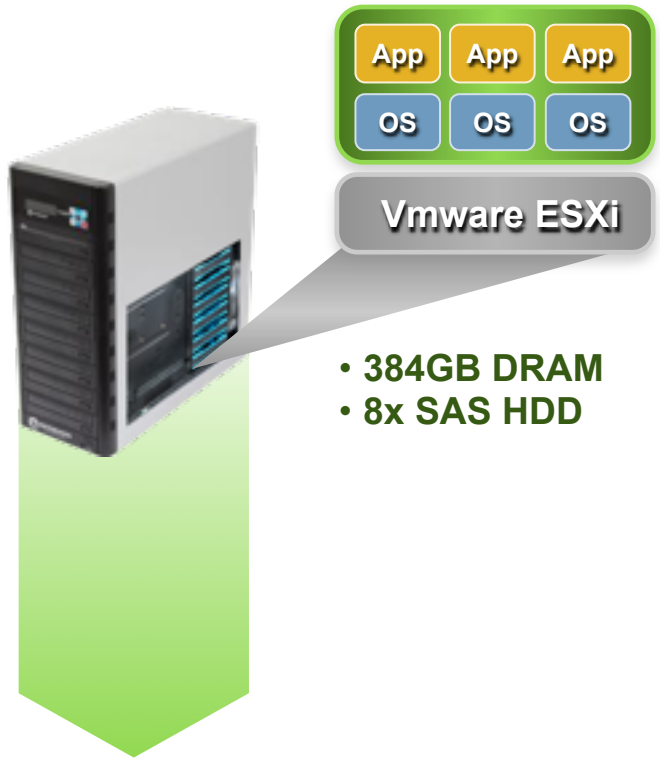

- 384GB DRAM
- 8x SAS HDD

167
Physical Servers

Example: Virtual Desktop Infrastructure

Deploying 12,500 Virtual Desktops

75 Virtual Desktops

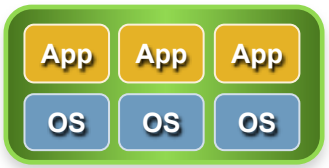

- 384GB DRAM
- 8x SAS HDD

167
Physical Servers

125 Virtual Desktops

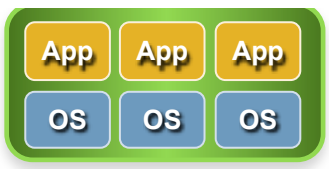

- 128GB DRAM
- 8x Flash DIMMs

100
Physical Servers

Example: Virtual Desktop Infrastructure

Deploying 12,500 Virtual Desktops

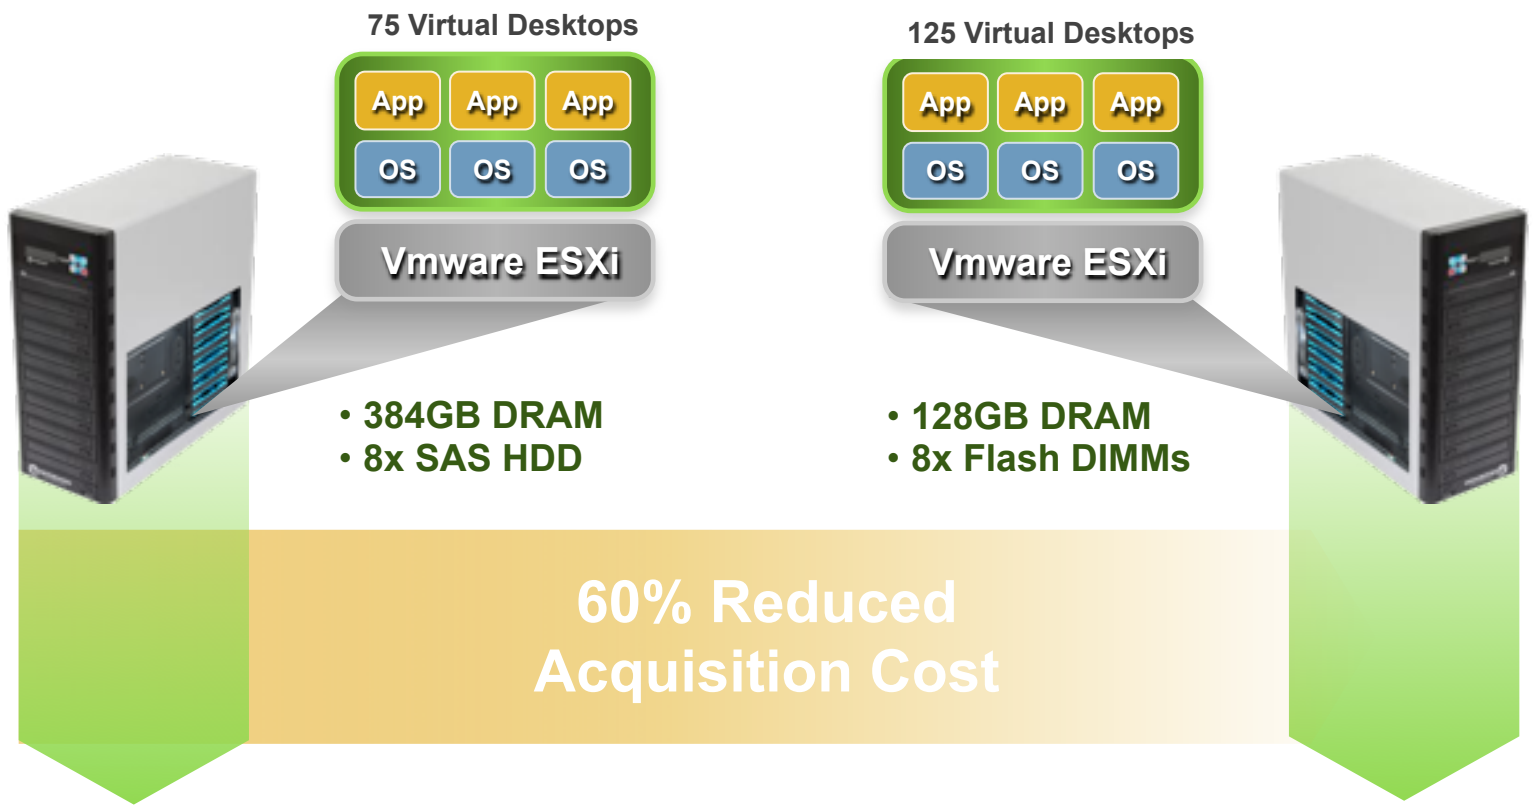

- 384GB DRAM
- 8x SAS HDD

- 128GB DRAM
- 8x Flash DIMMs

60% Reduced Acquisition Cost

167 Physical Servers

100 Physical Servers

Conclusion

Conclusion

Reaching the Final Latency Frontier

Leveraging DRAM interface

Enhancing Flash Endurance & Reliability

Storage IS memory, Memory is Storage

SMART STORAGE SYSTEMS™

Making NAND Better

Thank You