

# SAS Enhancements: Additional Increases in Low-Latency Device Performance


August 17-19, 2010

Santa Clara Convention Center

Harry Mason, Director Industry Marketing - LSI Corp,  
President STA

Marty Czekalski, Sr. Staff Program Manager - Seagate  
Technology, Vice-President STA

- Founded 1995
- Member funded trade group
- Participation varies around 30 companies



# STA Market Successes

- Sustained SCSI in Standard High Volume Servers

- **Ultra160 (parallel SCSI)**
- **Ultra320 (parallel SCSI)**



- Effective Server Transition to SAS

- **3Gb/s SAS**
- **6Gb/s SAS**
- **12Gb/s SAS (in Process)**




- Drive Preference in External Storage Market

- Breadth of Supply
- Market Scale



- Host Connection into External Storage Markets

- Serial Architecture: Addressing; zoning, high density interconnects
- **SAS Advanced Connectivity Roadmap**



# STA Objectives

- “Best-In-Class” Block Storage I/F for Enterprise Devices
- Investment Protection For Storage OEMs
- Viable Long-Term Roadmaps
- Deployment of New Storage Technologies (TTM)
- Vendor Neutral Solutions
- Wide Variety of Device Types
  - HDDs
  - SSDs
  - HBA's/RAID
  - Scanners
  - Printers
  - OCR
  - Enclosures
  - Tape
  - Media Changers
  - Optical Drives
  - Switches
  - Converters
  - Bridges/Gateways
  - Object Based Storage



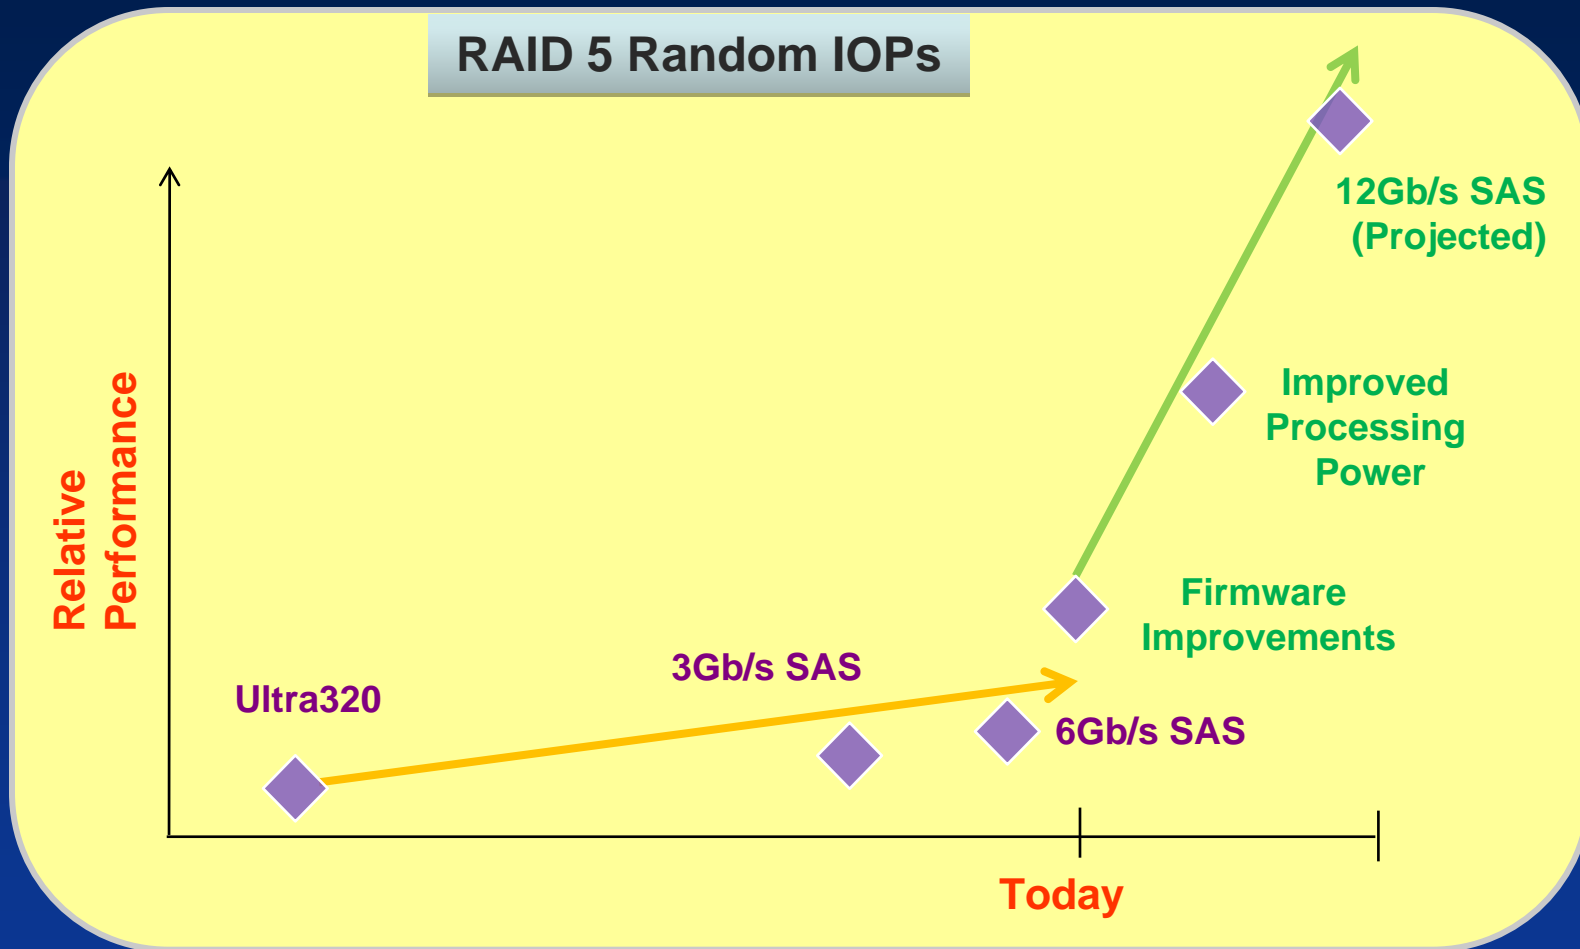
**Broad component offering = Broad range of solutions**

# SAS Logical Abstraction Layer: Preservation and Innovation

- Preserves SCSI Command Set Across:
  - Successive Product Generations
  - Frequent Technology Perturbations
  - Multiple Vendors
- Lowers System Integration Cost/Time
- Investment Protection: S/W, Middleware, & Enhancements
- Delivers Enterprise Attributes
- Operates over numerous transport layers
  - ATA Packet Interface (ATAPI)
  - USB Mass Storage Class, UAS (USB Attached SCSI)
  - Firewire SBP-2
  - iSCSI (SCSI over TCP/IP)
  - Fibre Channel (FCP), FCoE
  - Parallel SCSI
  - Serial Attached SCSI (SAS)

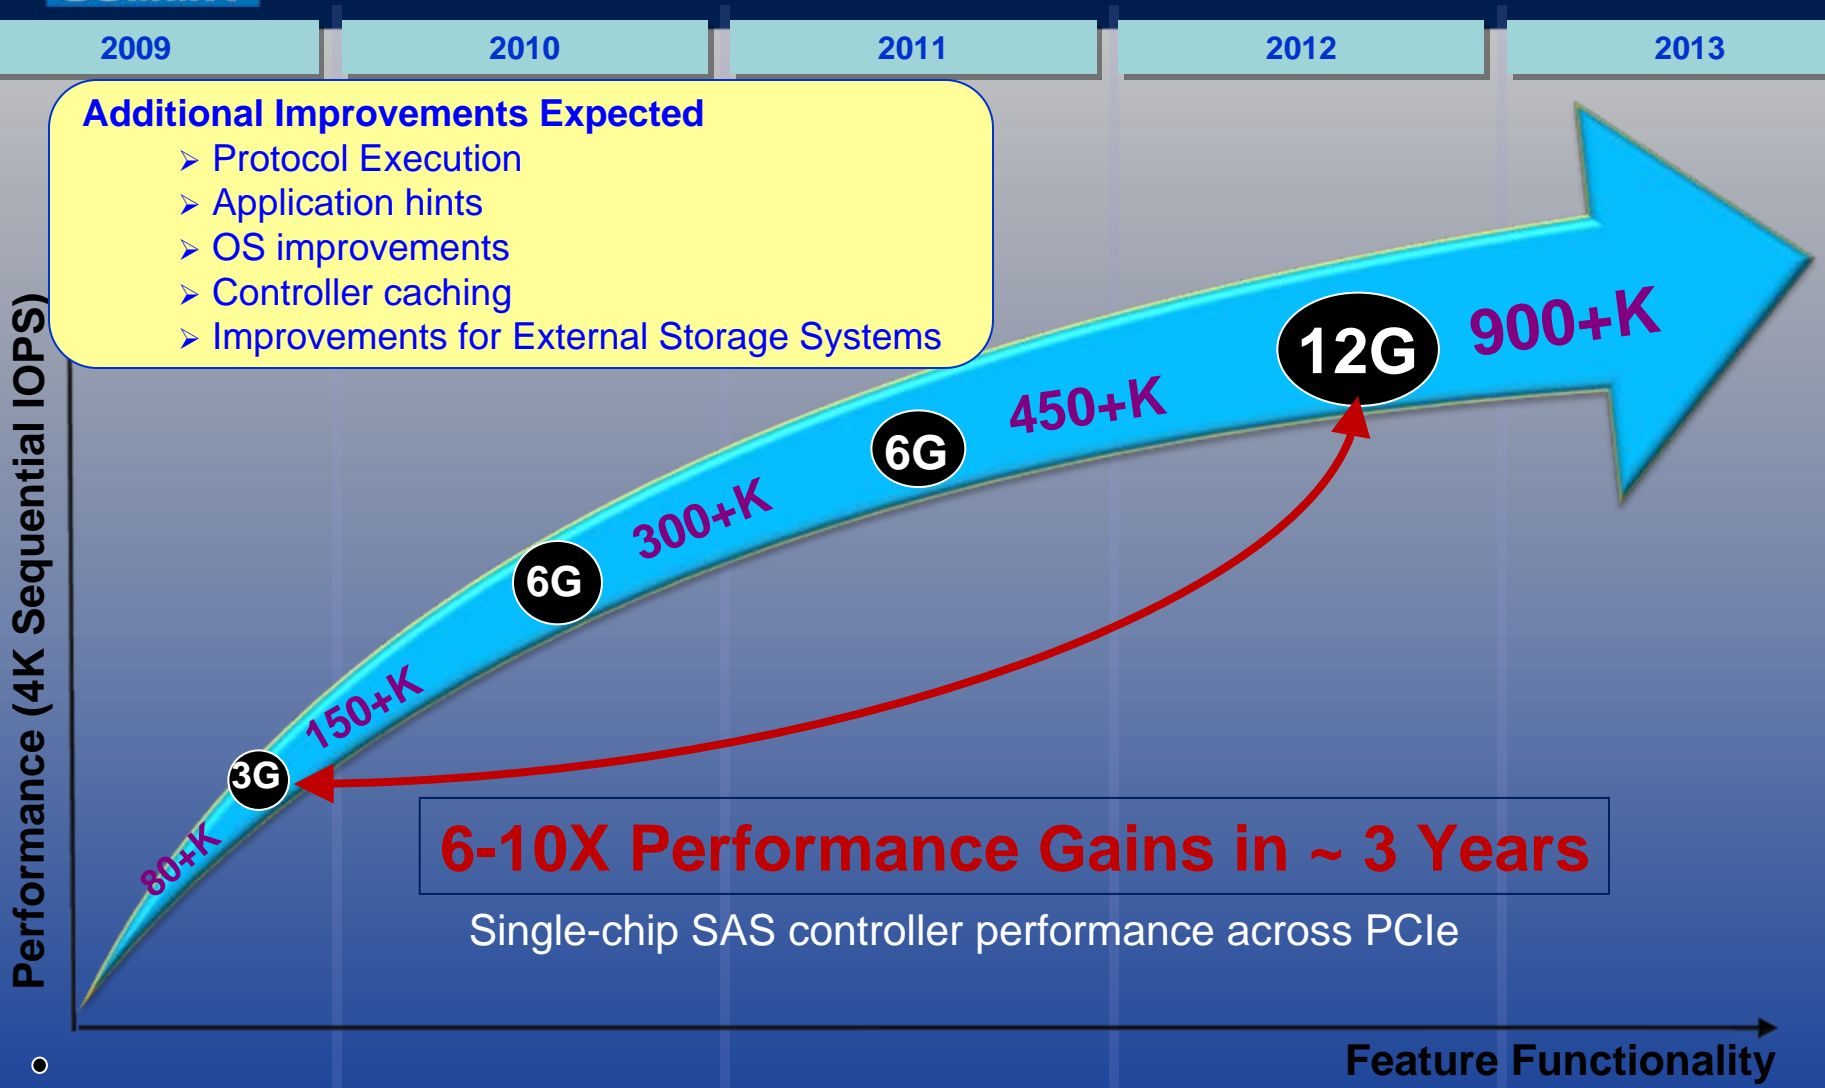


# SAS Performance Improvements



**Exceptional Gains without Exceptional Changes**

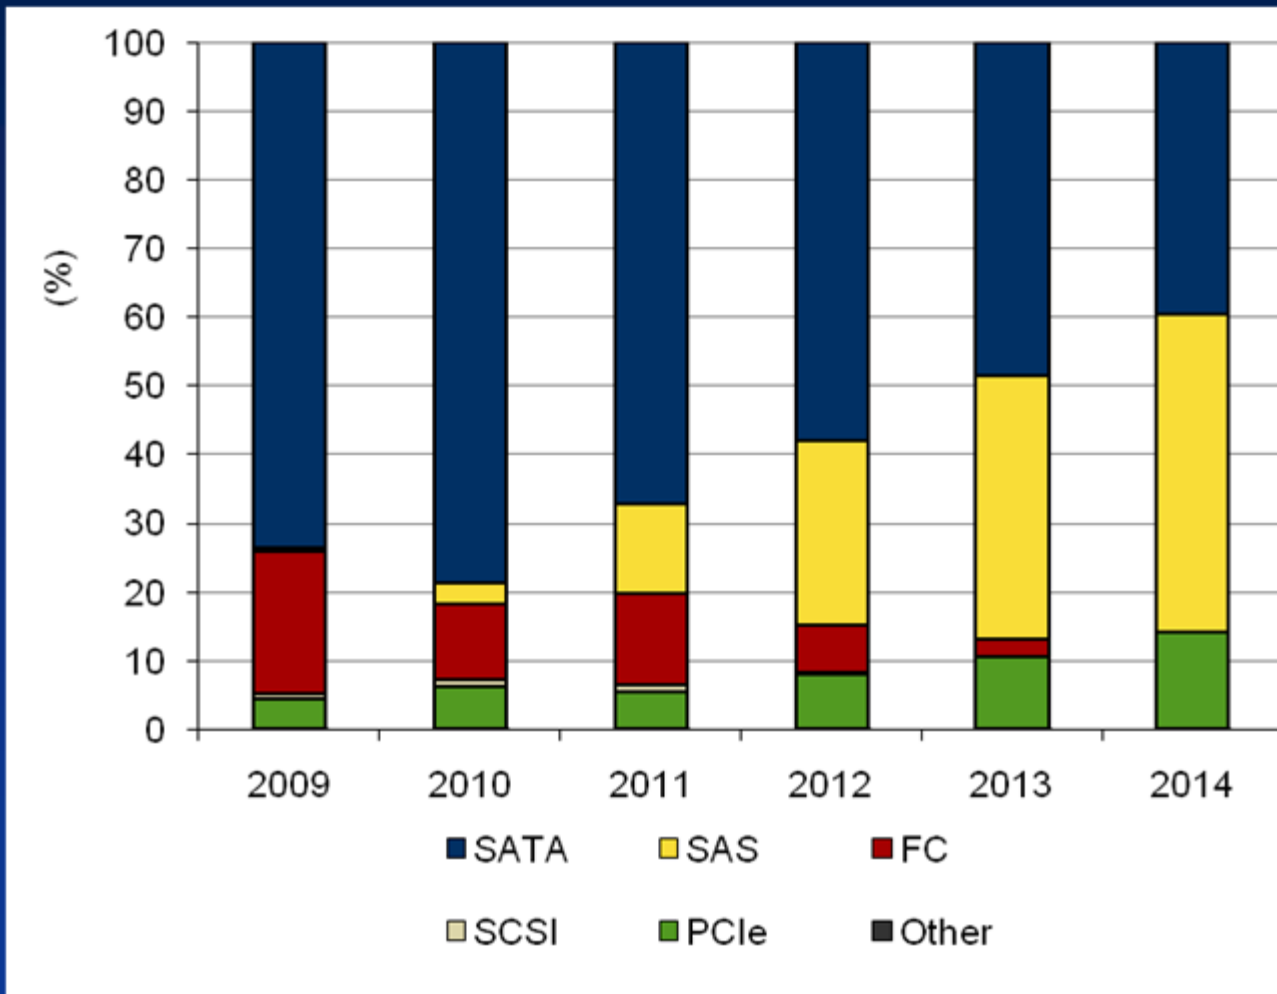
# SAS Projections – Existing Standards





# SAS/SATA SSD Market Share

Enterprise Shipments (Volume)




Source: IDC



# Solid State Alternatives

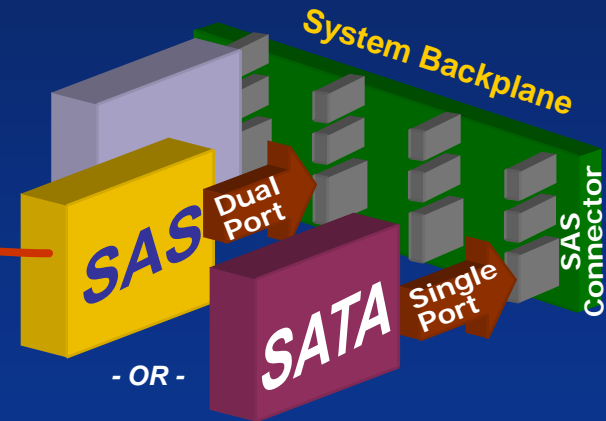
## Option A: Use Existing Drive slots

- High aggregate performance
- Moderate performance/slot (power limited)
- Hot Swap, Serviceability
- High Availability (two fault domains possible)



## Option B: PCIe Card

- High performance
- No Hot Swap, difficult servicing
- Availability issues (single fault domain)
  - Relies on application level replication



# Considerations for Low-Latency Storage Devices

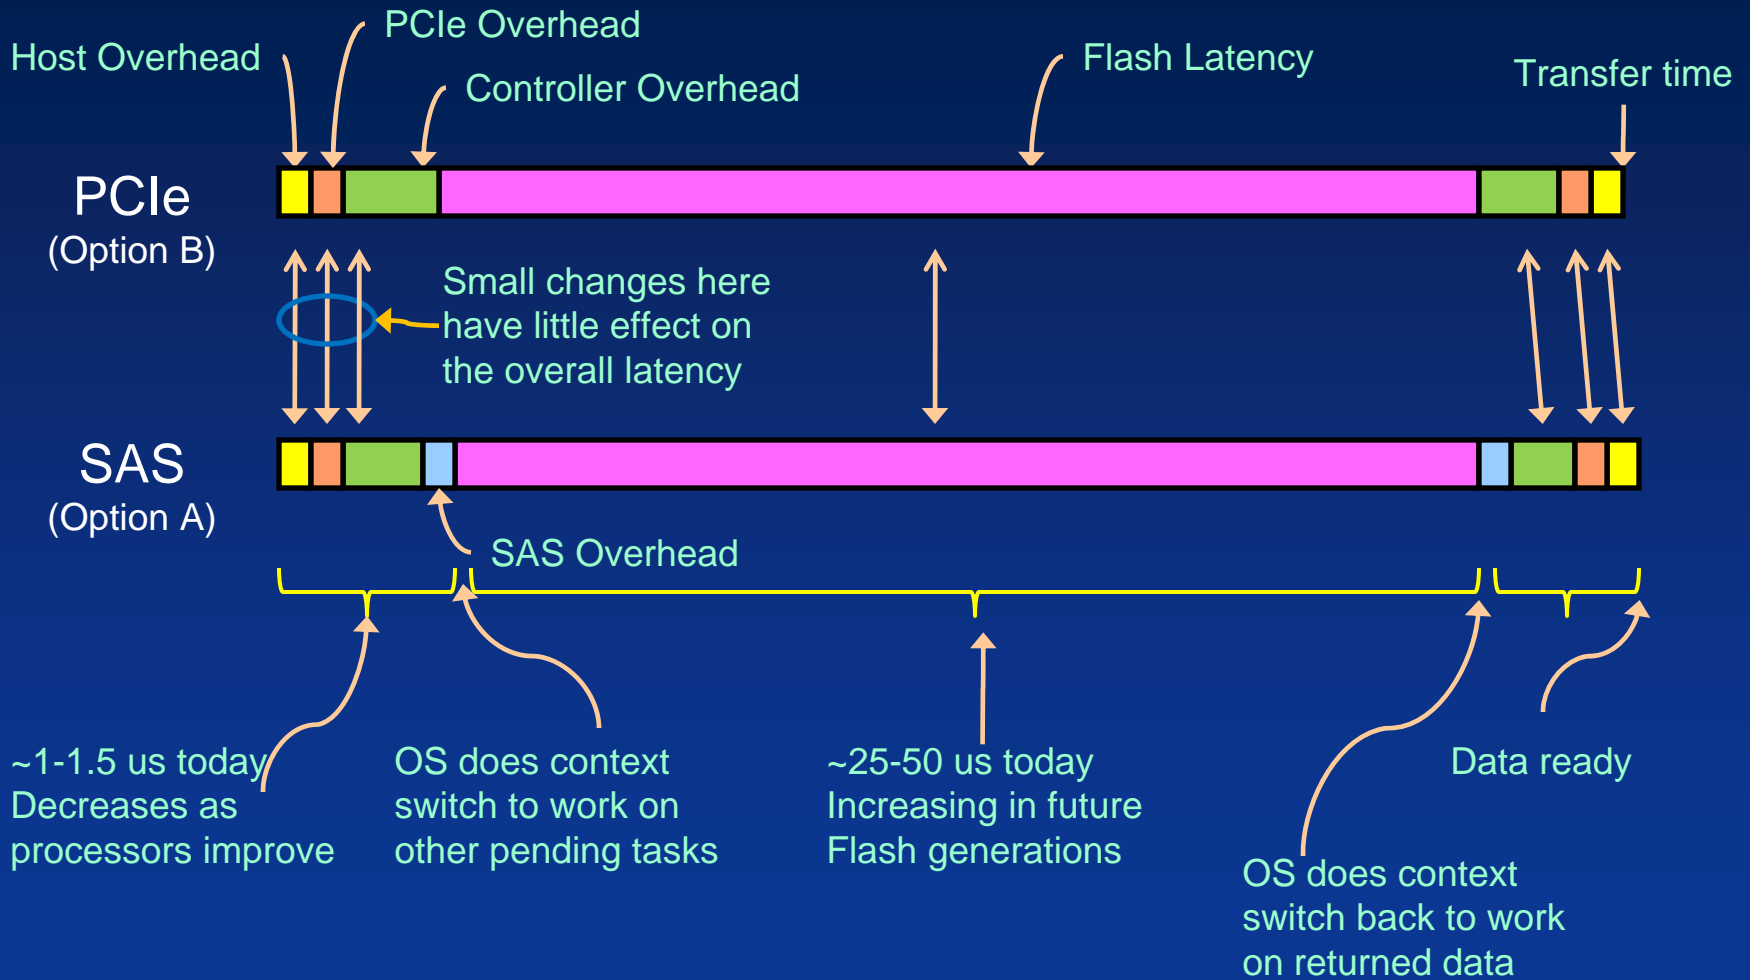
- Must Co-exist with SAS
- SSD Requirements
  - Non-persistent
    - Caching, I/O acceleration, etc.
    - Block Storage or Memory Model
  - Persistent - Enterprise Quality
    - Scales, hot plug, S/W, Robust Feature Set (zoning, encryption, etc.)
    - Multi-Vendor, technology transparent
    - SAS functionality is the Storage “Watermark”
- Performance Determination
  - Provisioned Power
  - Available bandwidth
  - Protocol latency

# The Performance Myth

- Option B (PCIe) is often perceived as being higher performance
- Option A, the SAS Reality:
  - Near equivalent access time latency (Limited Protocol Impact)
  - Excellent IOPs and BW today, with dramatic increases expected
  - Near “Lane for Link” bandwidth equivalency
  - Power dictates performance, not protocol
    - Number of active flash channels is limited by the power envelope
- Misconceptions regarding Option A performance are guided by:
  - Existing Power Envelope (9W vs 25W)
  - Single link vs Multi-lane Comparisons

**Multiple simultaneous transfers favor  
SAS for enterprise workloads**

# Latency Comparison



# Latency Comparison

## Option A: Flash-1(SAS)

- Currently shipping 6Gb/s SAS HBA
  - Standard SAS drivers
- 6 SATA SSDs (SLC)
  - HBA converts SCSI commands to SATA)
- Equal total capacity
- ~20W total power
- Preconditioned and in steady state

## Option B: Flash-2(PCIe)

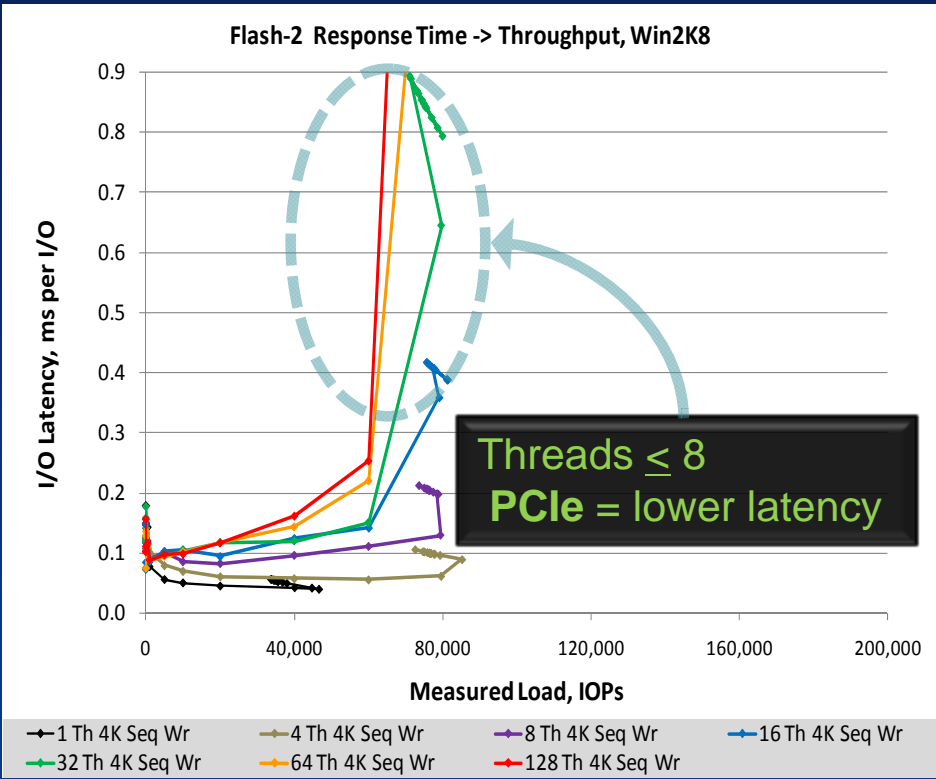
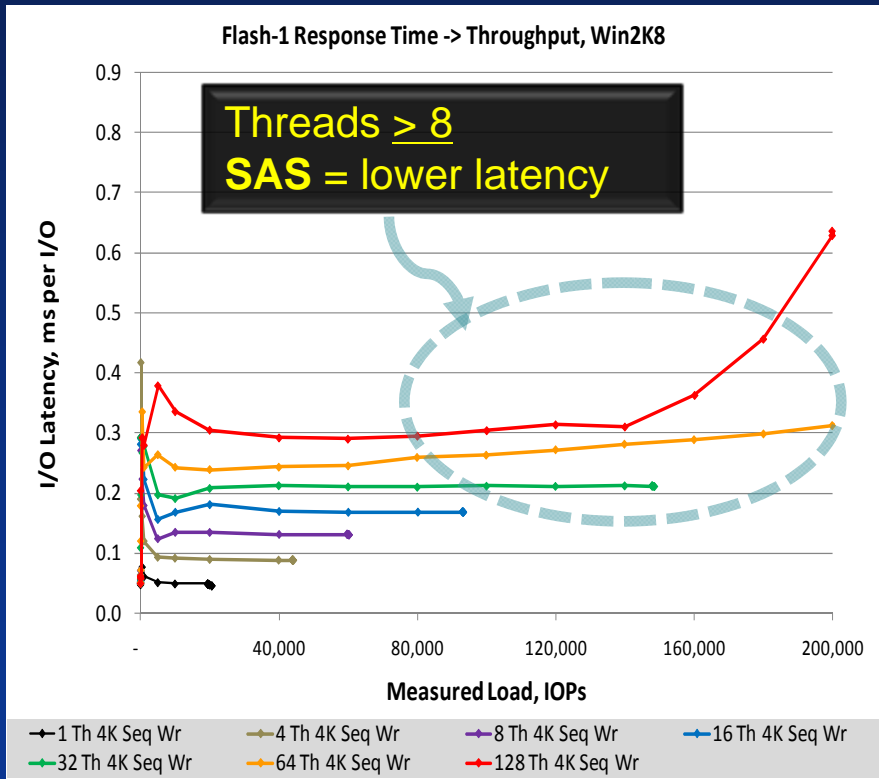
- Currently shipping native PCIe SSD (SLC)
  - Proprietary (non-SCSI) driver
- Equal total capacity
- ~20 W total power
- Preconditioned and in steady state

# Latency Comparison: Flash-1(SAS) vs Flash-2(PCIe)

Latency dependent upon:

- Offered load (4K Seq Write)
- Queue or thread depth (multiple thread sizes)

Latency measured with  
VDBench under Win2K8



**SAS Concurrency favors enterprise workloads**

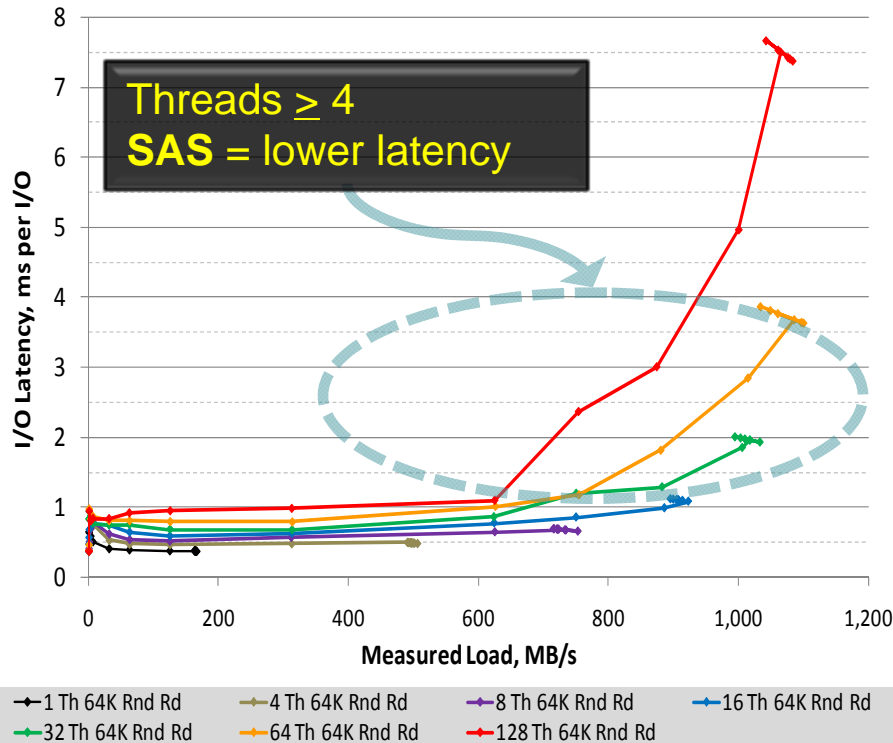
# Latency Comparison: Flash-1 (SAS) vs Flash-2 (PCIe)

## Latency dependent upon:

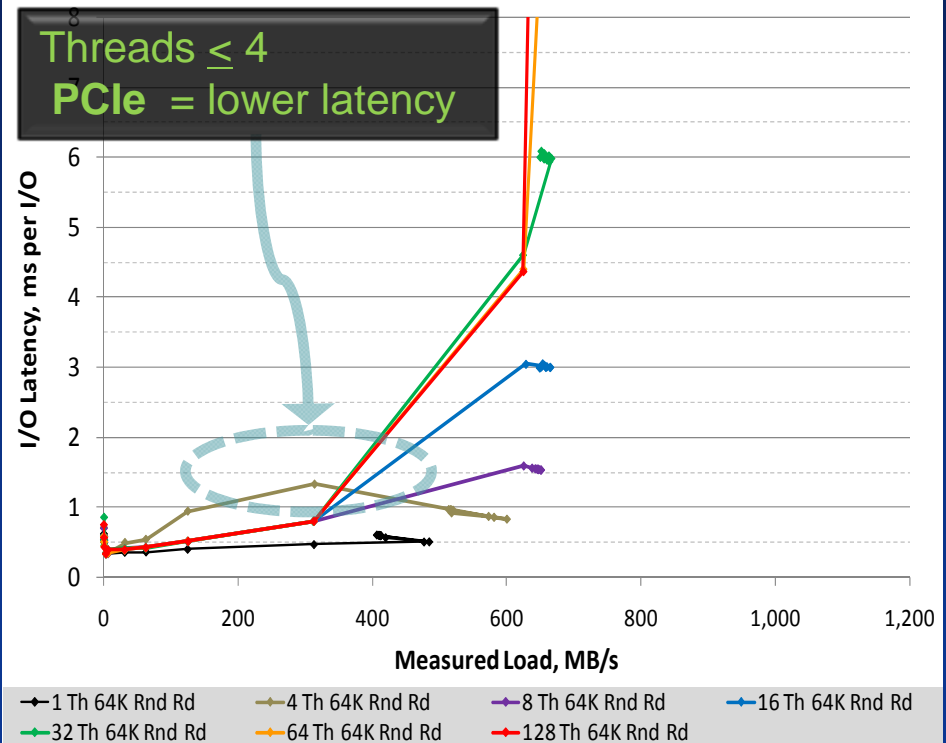
- Offered load (64K Random Read)
- Queue or thread depth (multiple thread sizes)

Latency measured with  
VDBench under Win2K8

Flash-1 ResponseTime -> Throughput, Win2K8



Flash-2 ResponseTime -> Throughput, Win2K8



**SAS Concurrency favors enterprise workloads**





# Key Comparisons Today

	PCIe	SAS
Links (BW)	X4 (4x500MB/s)	X1 (600MB/s)
Power Available	25W	9W (2.5")
Total Latency	>26 us	>26 us
Multi host protocol	No	Yes
High availability	No	Yes (Dual Port)
Scalability	Limited	Excellent
Robust proven protocol stack (security, encryption, end-to-end protection, zoning, distance, etc)	No	Yes
Hot Swap serviceable	No	Yes
Compatible with existing management SW	No	Yes

# MultiLink SAS Initiative

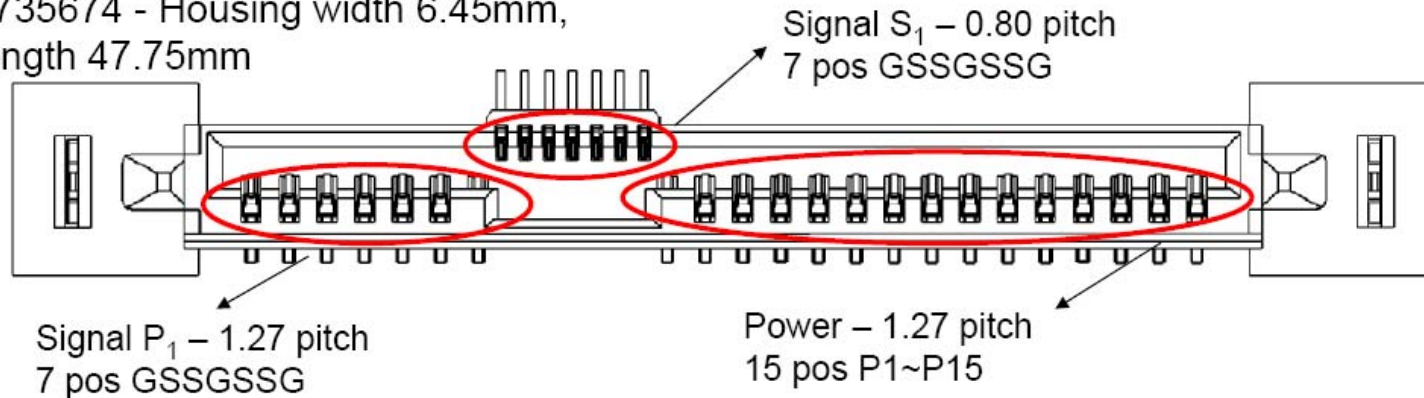
- STA Endorsed May 2010
- Defines new type of SAS Drive Slot
  - Increased Power/slot
  - Multiple SAS Links
- No Protocol Changes (Desired)
  - Work with existing infrastructure components
- Options for 12Gb/s and 24Gb/s link performance

# MultiLink SAS Objectives

- Externally-accessible backplane slot architecture for SAS drives
  - Additional Power provisioned for each MultiLink slot
  - Improved performance per slot for SSD drives
  - Optional for the system developer
- Backward Compatible Slot
  - Supports existing HDD & SDD drives
  - Extends existing backplane connectors to 4 SAS Links (2X2)
- Accommodates a variety of SSD form factors and configurations
- No protocol changes (desired)
  - Repackaging exercise
  - Power management & MultiLink SAS configuration, future enhancements

## SAS receptacle configuration

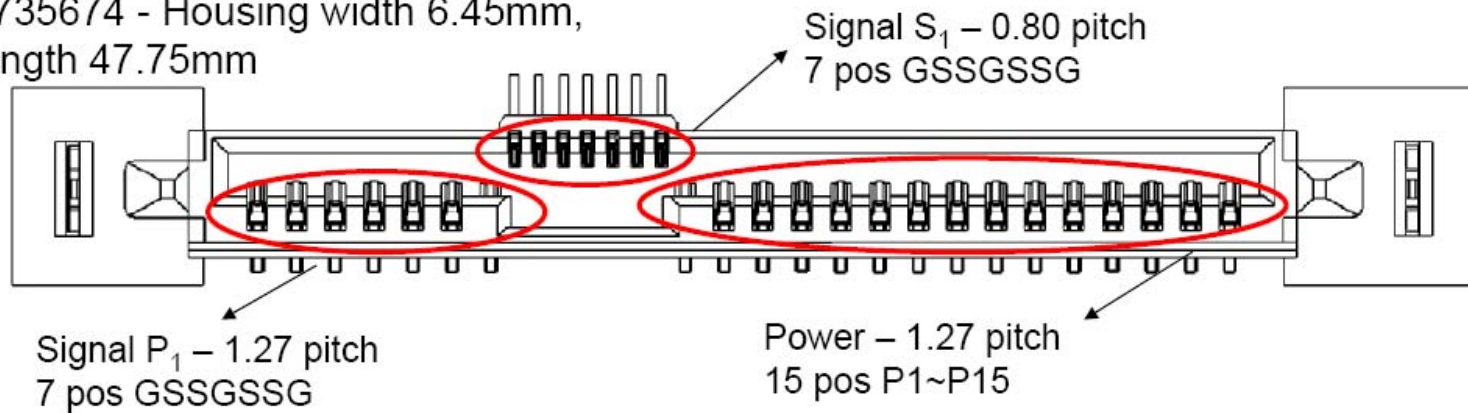
Current SAS receptacle configuration  
1735674 - Housing width 6.45mm,  
length 47.75mm




## SAS receptacle configuration

Current SAS receptacle configuration

1735674 - Housing width 6.45mm,  
length 47.75mm



Proposed SAS  
receptacle  
configuration



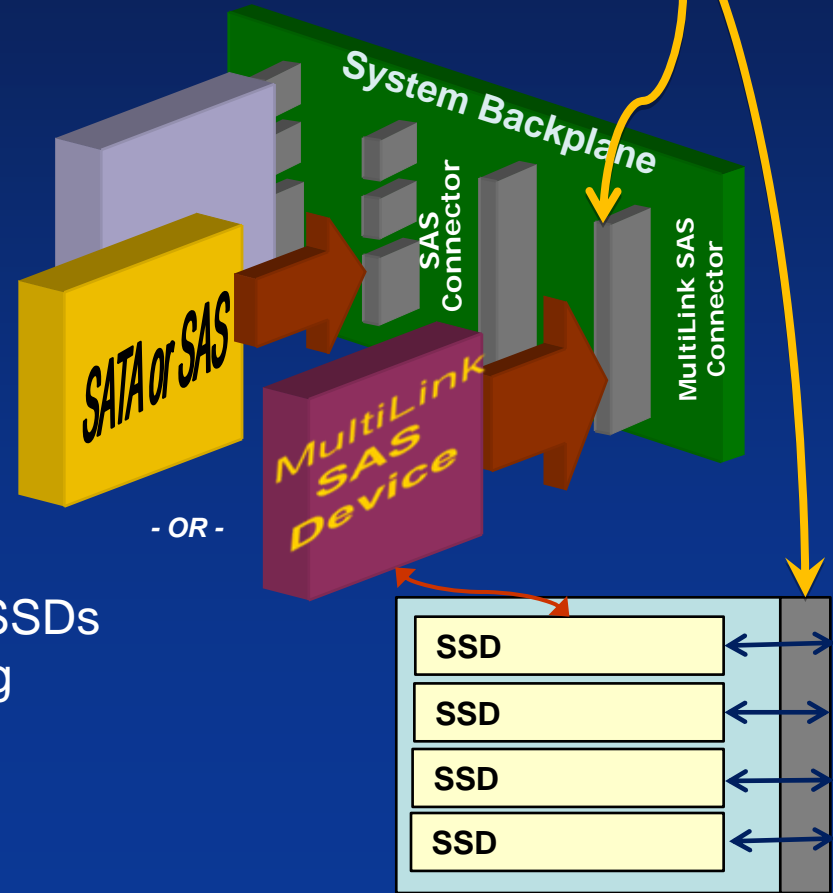
# MultiLink SAS Slot

MultiLink SAS Connector



## MultiLink SAS

- High Performance (20+W/slot)
- Hot Swap, Serviceability
- High Availability (2 fault domains possible)
- Low Implementation Risk:
  - Standard SAS drivers
  - Fully hardened protocol stack
  - Common Management Stack
- Low investment (repackaging)
- Flexible: Independent SSDs or wide-port SSDs
- Able to isolate tier-tier traffic w/o accessing system memory






# MultiLink SAS Roadmap: Backplane Slot Connector

Legacy Slot Compatible

New Slot





# MultiLink SAS timetable



**Kick-off**

**Validation and  
Proposals to  
T10 & SFF**

**Demos and  
prototypes**

**12Gb/s SAS  
Plugfest**

**OEM System Integration and Qualifications**

**Refreshed server platforms**

**New server  
platforms**

# SAS: New Enhancements Drive Performance Gains

- Power & Links/Lanes determine performance
- SAS concurrency favors enterprise workloads
- Mature eco-system: in place & multi-sourced
- Recent Dramatic Performance Improvements:
  - No changes to SAS standard
  - Gains from: Firmware, CPU power, Hardware assist, etc.
  - Further Gains: Application hints, OS streamlining, Protocol Execution, Controller caching, etc
- MultiLink SAS:
  - Added performance within a standard 15mm slot
  - An integral part of 12Gb/sec SAS deployment
- Sustainable time-to-market advantage (depth & breath of infrastructure)
- “Total solution” spans broad application space
- Any SSD/Flash storage solution MUST co-exist with SAS

# STA Member Companies

