



## **SAS Anchors Enterprise SSD Deployment**

- Harry Mason – president, SCSI Trade Association; director, Industry Marketing, LSI Corporation
- Marty Czekalski – vice president, SCSI Trade Association; Interface & Emerging Architecture program manager, Seagate

*Santa Clara Convention Center, August 9*

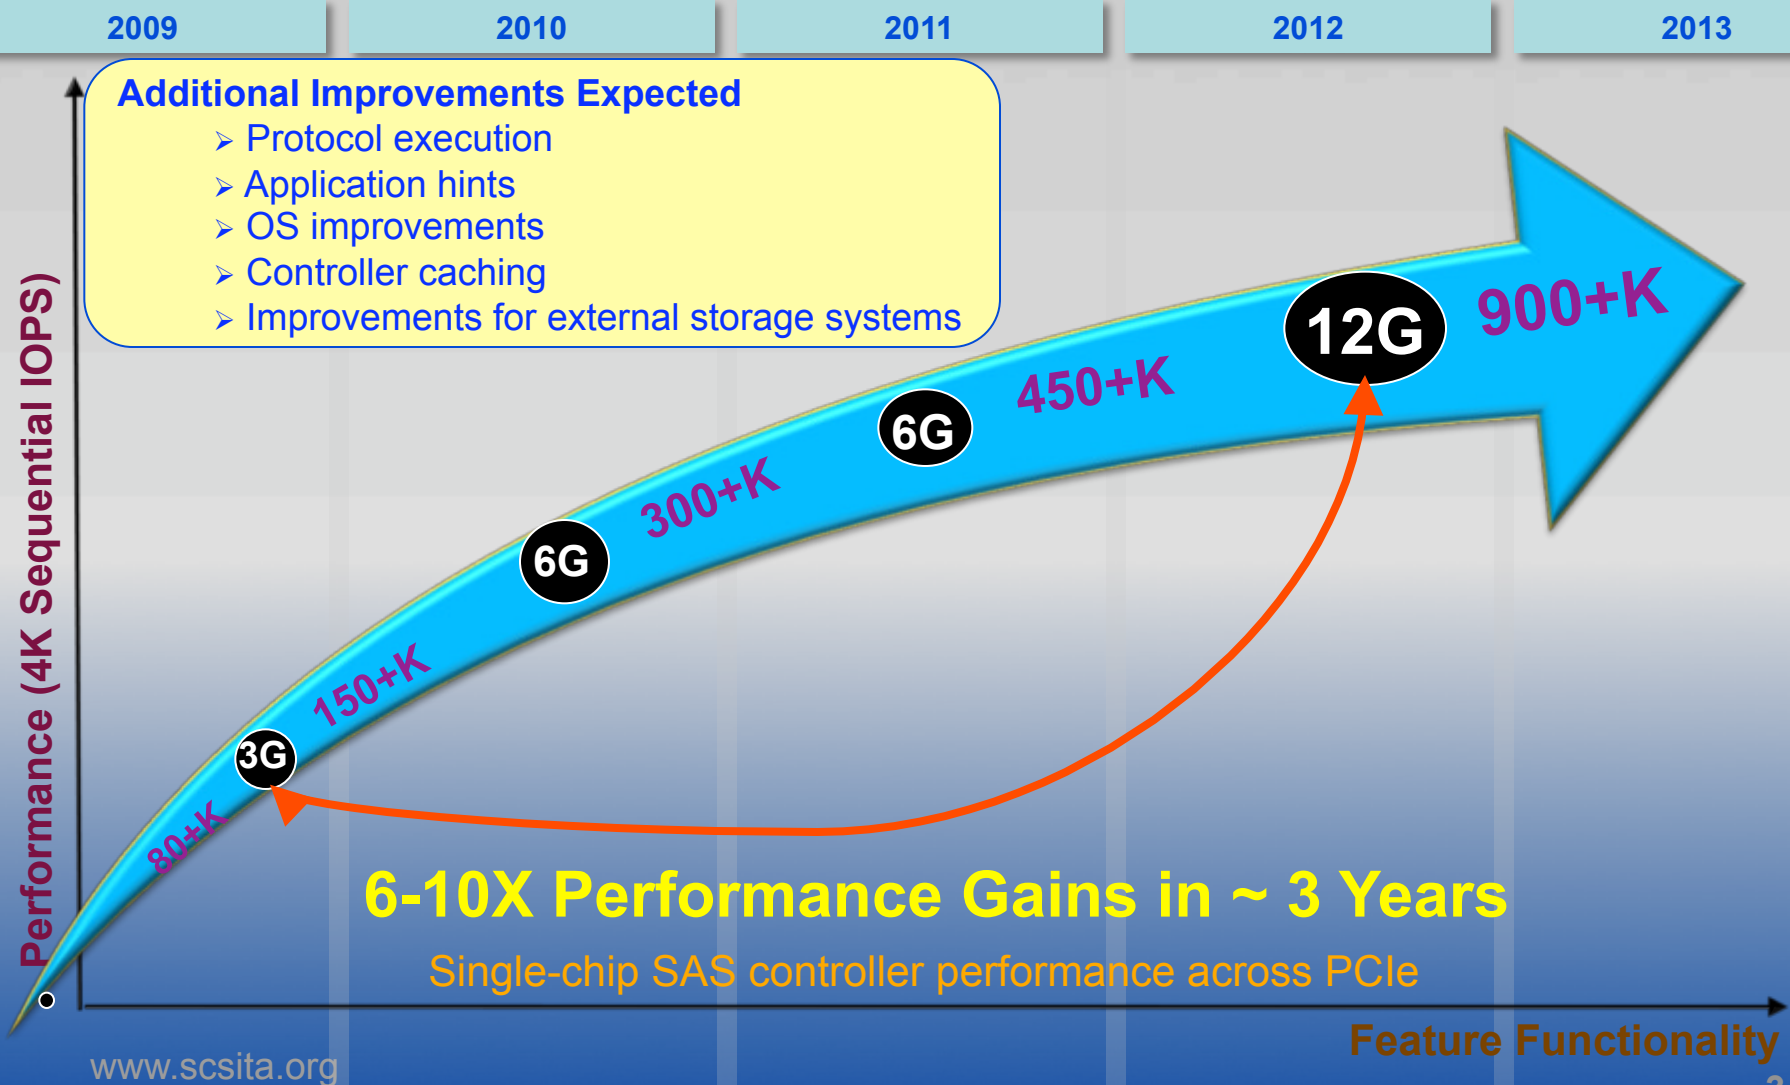


## SAS is the Enterprise




- Preserves Storage Investment – Logical SCSI
- Broad Open Industry Standards Support
- Dynamic Platform for Storage Innovation
- Enterprise Proven – RAS
- Depth and Breath of Infrastructure
- Class Drivers Personalize Non-HDD Devices
- Operates Across Numerous Transports
  - iSCSI, ATAPI, FC, Firewire, USB, PCIe

# SAS Projections – Existing Standards

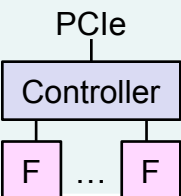
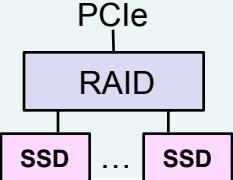


# Server-attached SSD Forecast

## Server-attached SSD Units

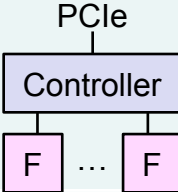
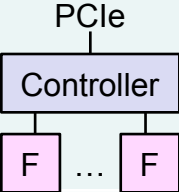


# Enterprise Interfaces: PCIe SSDs

	Native	Aggregator
Commands/Transport	 <p>Proprietary (FTL<sup>1</sup> in host/main memory)</p>	 <p>SCSI or SATA (Multiple SSDs &amp; controller on card)</p>
Committee	None	None
Standards Based	No	<b>Yes</b>
Performance with Flash	High	High
CPU Overhead	High	<b>Low</b>
Latency with short queue	Very Low	Low
Latency with deep queue	Moderate	<b>Low</b>
Use Case Extensibility	No	<b>Yes (RAID, HBA, etc)</b>
Maturity	Evolving	<b>Based on Proven Industry Architectures</b>
Enterprise feature set (PI, Security, Mgmt, etc.)	No	Depends on implementation



# Enterprise Interfaces: The Future of PCIe SSDs

	SOP/PQI <sup>1</sup>	NVMe <sup>2</sup>
Commands/Transport  <small> <sup>1</sup>SOP : SCSI over PCI Express  <sup>2</sup>NVMe : Non- Volatile Memory Express  <sup>3</sup>PCIe Queuing Interface  <sup>4</sup>INCITS : International Committee for Information Technology Standards           </small>	SOP/PQI <sup>3</sup> (FTL in controller) 	NVMe/NVMe (FTL in controller) 
Committee	<b>T10/INCITS<sup>4</sup></b>	Industry Working Group
Standards Based	<b>Yes (ANSI/ISO)</b>	No
Performance with Flash	Very High	Very High
CPU Overhead	Low	Low
Latency with short queue	Very Low	Very Low
Latency with deep queue	Low	Low
Use Case Extensibility	<b>Yes (RAID, HBA, etc.)</b>	No (NVM only)
Maturity	<b>Investment Protection</b>	TBD
Enterprise feature set (PI, Security, Mgmt, etc.)	<b>Full Support</b>	Limited



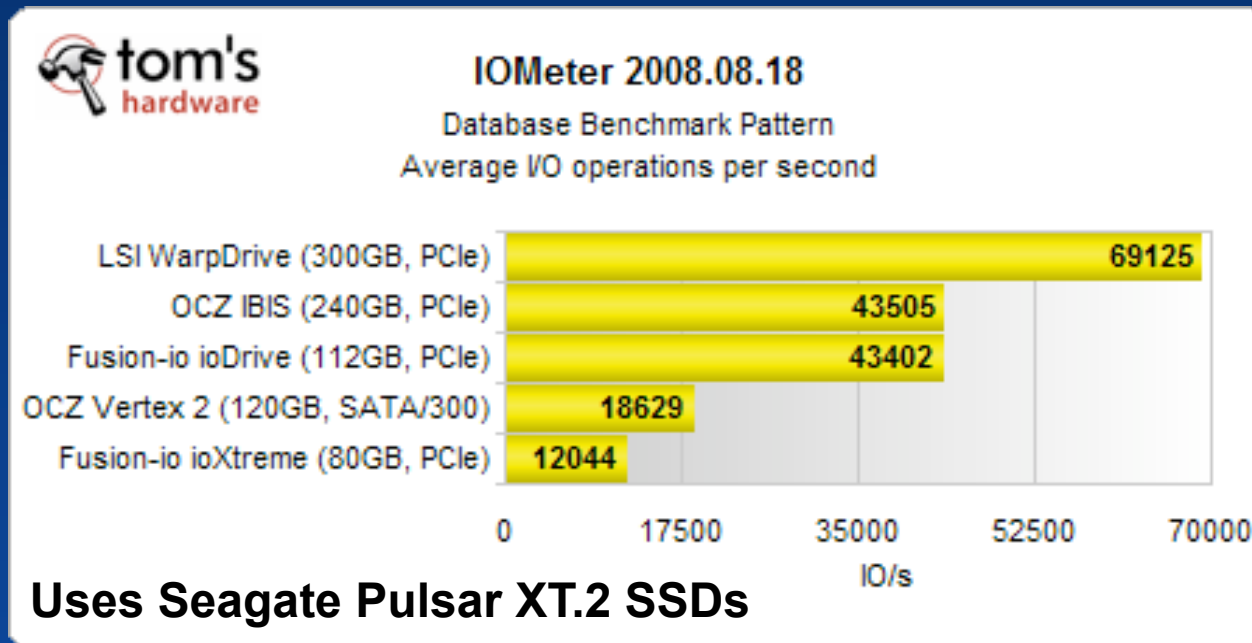
# Flash SSD Latency Facts

- Dominant contributor to SSD latency – Flash Devices
  - SLC access > 25us, MLC access > 50us, assuming no contention
  - Write/erase times are considerably longer, especially with enterprise versions, and getting longer with each flash generation
- Lager queue depths increase latency
  - Once a flash part begins access, other requests to that part must wait
  - Up to 8 flash die share bus access, causing each die to wait its turn
- Housekeeping activities add additional latency (address translation, garbage collection, wear leveling, etc.)

**Protocols & interconnects have negligible effects on latency.**

## Aggregator Example - SAS SSDs

“... the I/O capabilities of its LSISAS2008 controller, as it dominates the I/O test patterns for database, file server, Web server, and workstation scenarios.”




Selected by ICC for National Center for Supercomputing Applications (NCSA) “Dark Energy” project

[http://www.lsi.com/downloads/Public/Direct%20Assets/LSI/casestudy\\_ICC-NCSA.pdf](http://www.lsi.com/downloads/Public/Direct%20Assets/LSI/casestudy_ICC-NCSA.pdf)


# MultiLink SAS™ Slot

## ■ MultiLink SAS™

- High performance (20+W per 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 Location

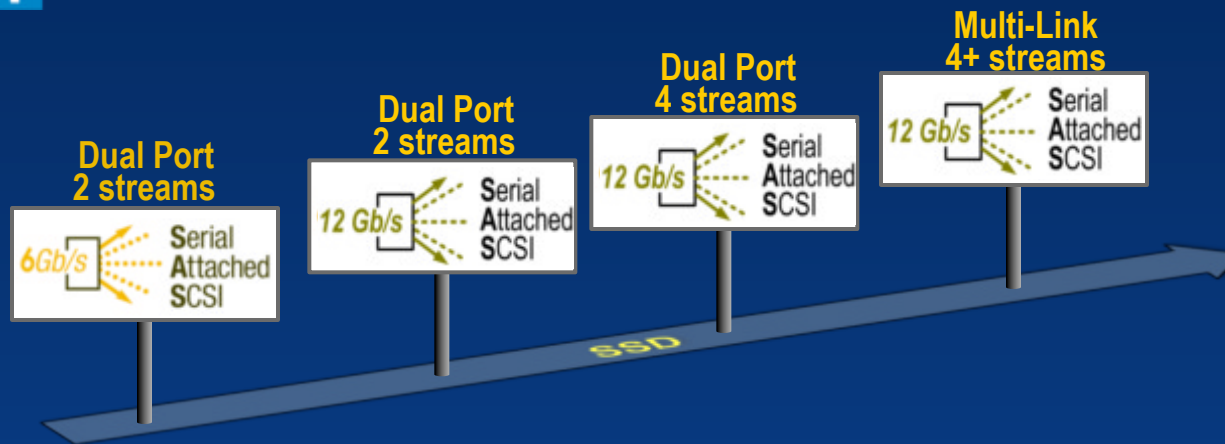




## Innate SAS Advantages

- Persistent & non-persistent usage models
- Concurrency – Performs best with large queues and high work loads
- Low CPU Utilization – Supports multi-tenant cloud service platforms
- Flexible Deployment Model – Cabled, PCIe, SAS slots, motherboard, etc.
- Solid legacy, rich infrastructure, & viable roadmap
- “No compromise” performance

# Innovating for Bandwidth Growth



- Today: Leverage 6Gb/s SAS infrastructure for rapid deployment
- Tomorrow: Enhance the SAS to support SSD unique requirements:
  - SAS Roadmap Progression:
    - A) 6Gb/s SAS dual port, 2 streams
    - B) 12Gb/s SAS dual port, 2 streams
    - C) 12Gb/s SAS dual port, 4 streams
    - D) 12Gb/s SAS multi-link, 4+ streams
  - SSD unique commands: Trim, Unmap, others...

Benefits of SAS & MultiLink SAS™ Enhancements	
Multiple Links (BW)	X4 (4x600MB/s)
Power Available	25W (2.5")
Total Latency	>26 us
Multi host protocol	<b>Yes</b>
High availability	<b>Yes (Dual Port)</b>
Scalability	<b>Excellent</b>
Robust proven protocol stack	<b>Yes</b>
Hot Swap serviceable	<b>Yes</b>
Compatible with existing management SW	<b>Yes</b>