



Getting the Most from Your SAS Infrastructure

Marty Czekalski – Seagate Technology
President – SCSI Trade Association

SCSI Logical Abstraction Layer: A foundation for 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

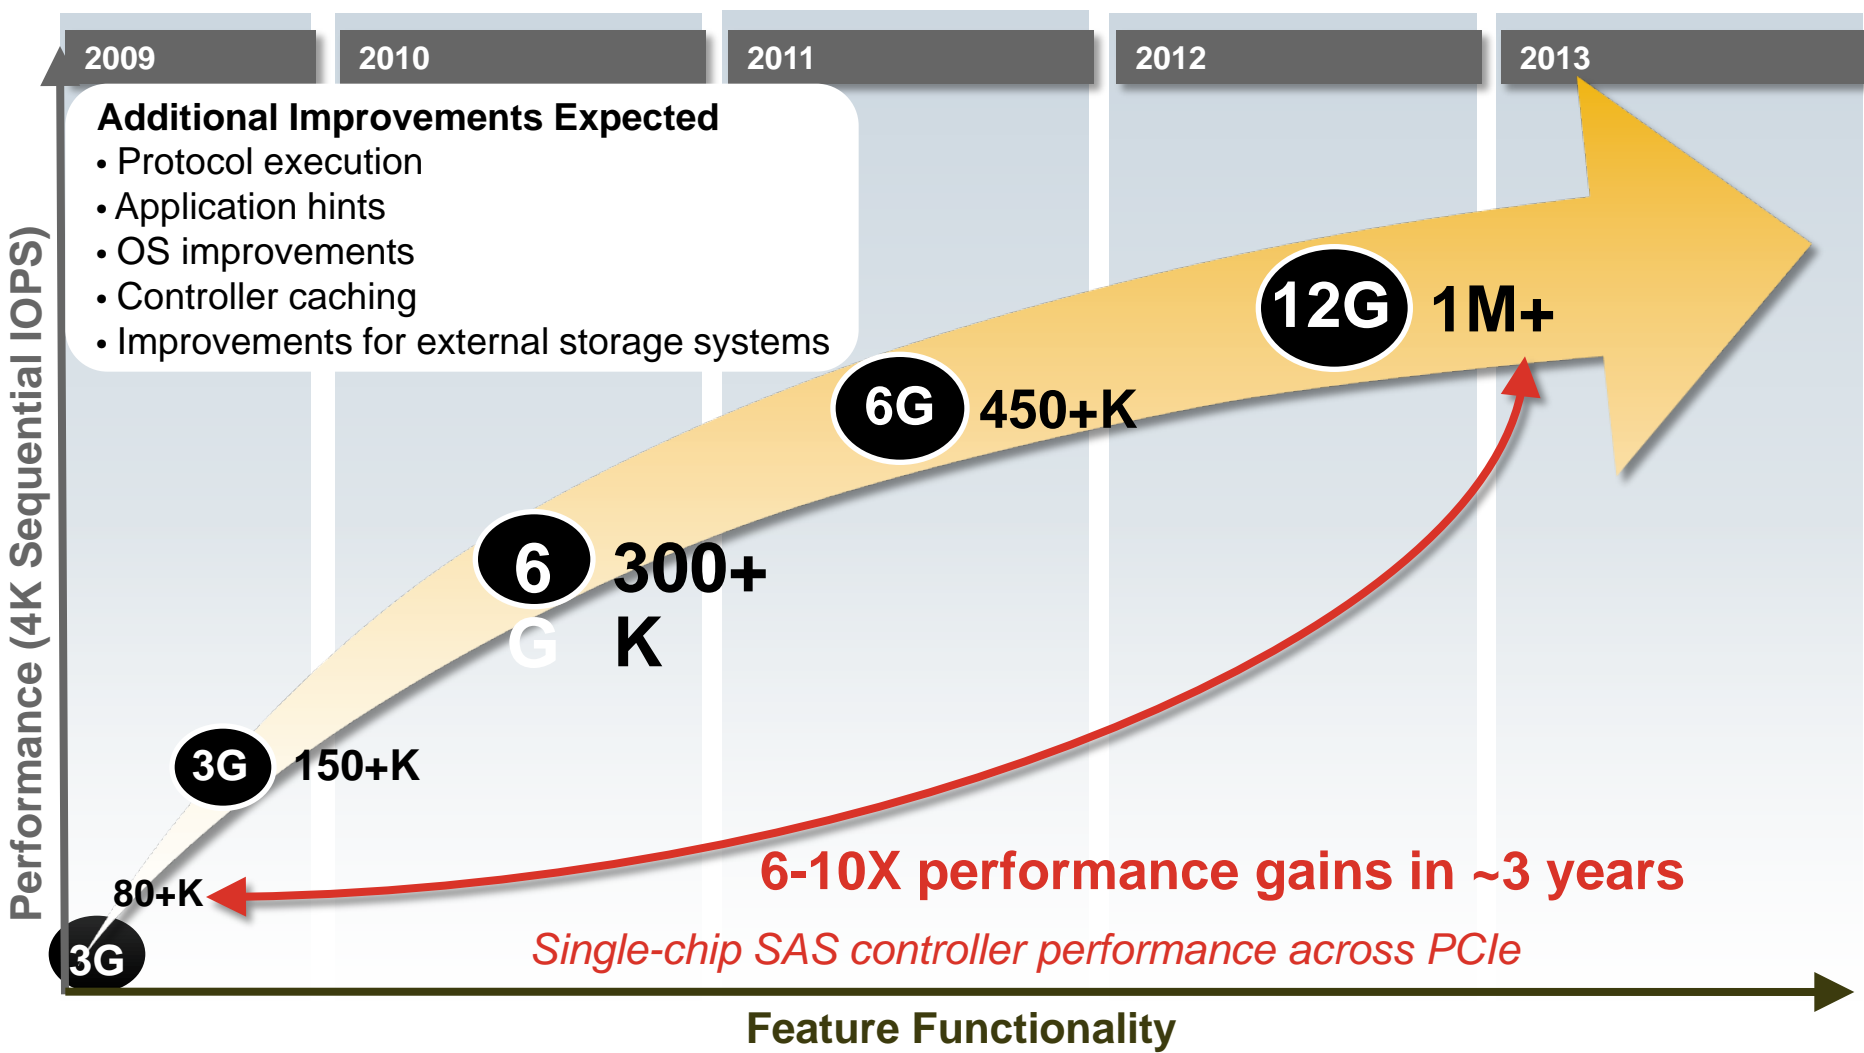
Operates over numerous transport layers

- ATA Packet Interface (ATAPI: ATA, SATA), USB Mass Storage Class, UAS (USB Attached SCSI, Memory Sticks), Firewire (1394), Infiniband, iSCSI, Fibre Channel (FCP, FCoE), Parallel SCSI, Serial Attached SCSI (SAS)
- SCSI Over PCIe (SOP, PQI)

SCSI:
**The Most Widely Implemented Logical Storage
Protocol**

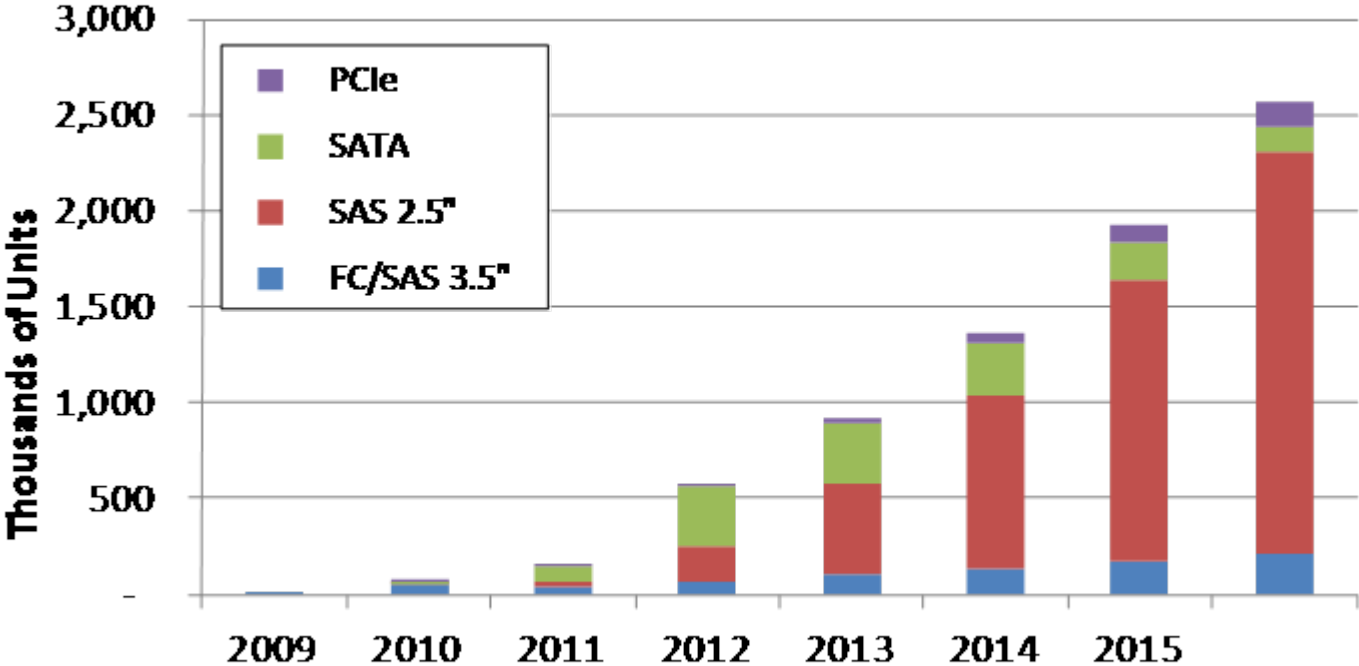
SAS Controller Projections

- No SAS Changes



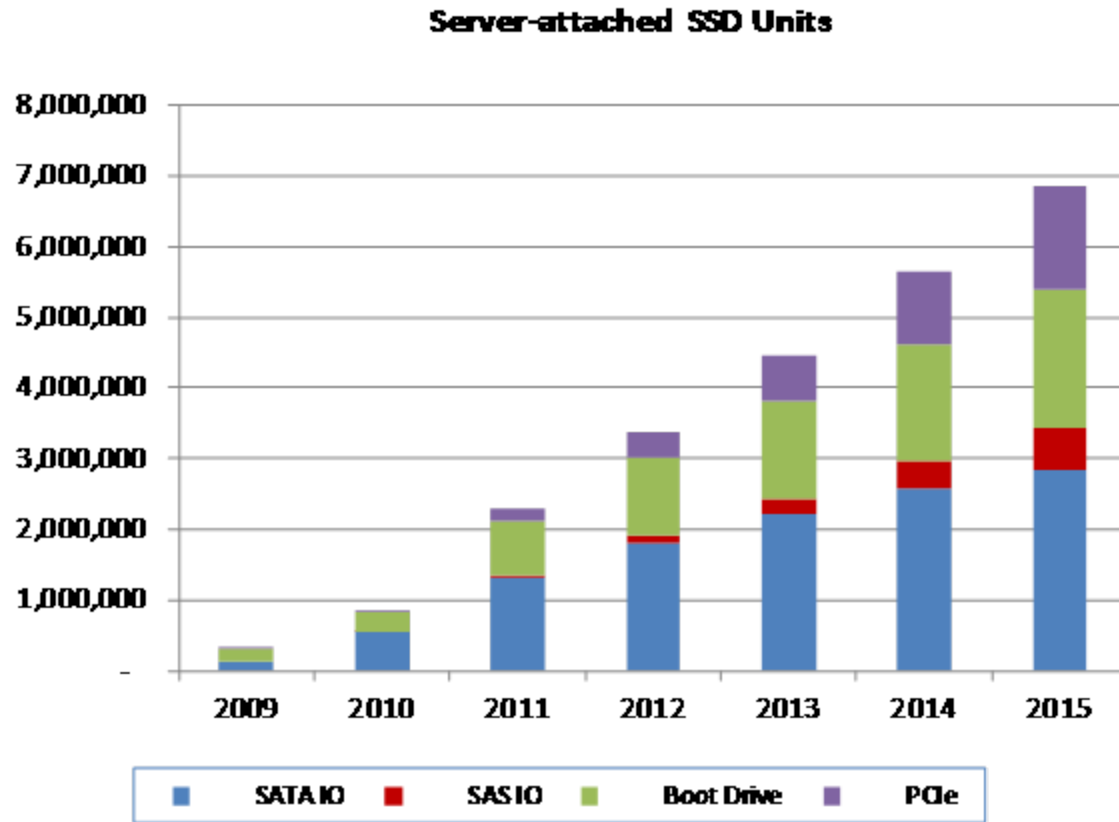
SAS is the preferred SSD Interface for Storage Systems

Storage-attached SSD Units



Forward-Insights 11-2011

Server Attached SSDs



Forward-Insights 11-2011

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
 - Can be SAS or PCIe attached
- 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
 - Increased slot power in Express Bay also usable by SAS devices
- Available bandwidth
 - Multi-link (96Gb/sec per backplane slot at 12Gb/sec)
- Protocol latency
 - HW improvements, NVM enhancements
 - SOP/PQI (SCSI Express)

Flash SSD Latency Facts

Dominant contributor to SSD latency – Flash Components

- SLC access > 25us, MLC access > 50us, assuming no access contention

Larger queue depths increase contention increasing latency

- Once a flash part starts its access, other requests to the same 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.)

Differences in protocols & interconnects have negligible effects on latency, especially at the application level (a few microseconds)

SAS Connects the Tiers

Managed - More scale, more cables, greater need for management

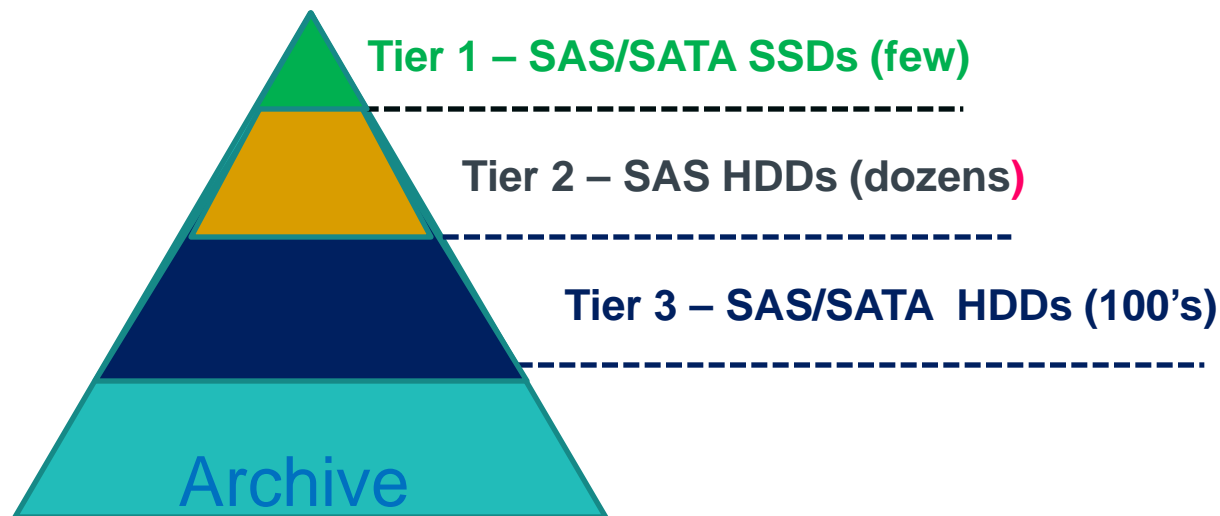
Distance - Active copper (20m), optical (100m)

Performance - 6Gb/s & 12Gb/s SAS

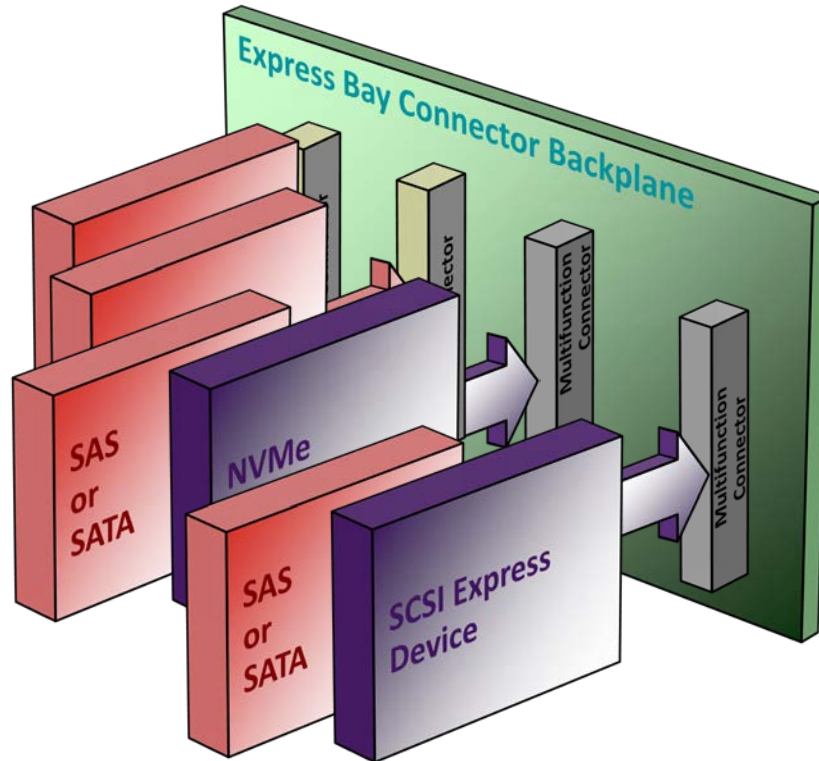
Density - More “beachfront,” More ports

Consistency - Standard method to scale distance & management

Converged - One solution for external active & passive



Express Bay





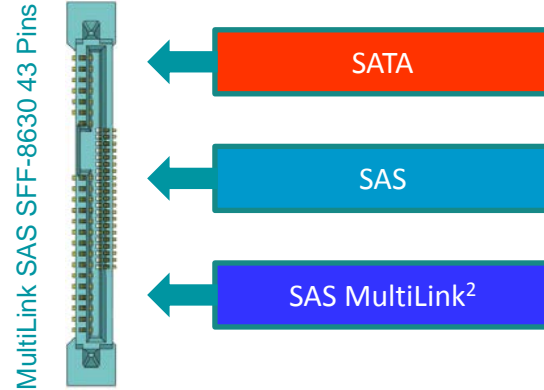

Express Bay

- Up to 25 Watts
- SFF-8639 connector
- PCI-SIG electrical work

Objectives

- Preserve the enterprise storage experience for PCI Express storage
- Meet SSD performance demands
- Serviceable, hot-pluggable Express Bay opens up new possibilities...

Connector Compatibility



- ¹ Max two links operate
- ² Four links operational
- ³ Two or four links operation depending on host provisioning

SCSI/SAS – Looking to the Future

- SCSI Express
- 12Gb/s, 24Gb/s SAS
- Extended Copy Feature
- Power Limit Control - up to 25W SAS devices
- Atomic Writes
- Hinting & other NVM features

SCSI Express



Proven SCSI protocol combined with PCIe creating an industry standard path to PCIe-based storage

- Enterprise storage for PCIe based storage devices
- Increased performance through lower latency
- Coexistence with SAS via Express Bay and common command set
- Unified management and programming interface

STA Member Companies



SCSI Express Components



Existing industry initiatives delivering enterprise storage using PCI Express

Technology

SCSI
SCSI Over PCIe (SOP)

PCIe Queuing Interface (PQI)

Express Bay connector (SFF 8639)

PCI Express

Description

The storage command set
Packages SCSI for a PQI queuing layer

Flexible, high-performance queuing layer

Accommodates PCIe, SAS, and SATA drives

Leading server I/O interconnect

12Gb/sec, 24Gb/sec SAS



12Gb/s SAS

- Standard in final stages of completion (SPL and SAS-3)
- 12Gb/s STA Plugfest in August 2012
- Doubles the throughput while maintaining the same distance use cases
 - Copper, Active and Optical cable options
- Expect Controller capabilities >1 million IOPs
- Volume ramp for end users 2H2013

24Gb/s SAS

- Connector proposal in the works
- Encoding scheme may change



Power Limit Control -up to 25W SAS devices

Allows system to decide how much power a device is allowed to consume. Both less or more than nominal.

For SSDs power = # of active die = performance

SAS SSDs will be able to take advantage of extra power provisioning above the nominal 9W limit in typical 2.5" slots.

- Example: SAS drive in a SAS MultiLink or Express bay (SFF-8630, SFF-8639) could be allowed to consume up to 25W



Atomic Writes



All or nothing written capability across multiple commands

- For single commands and across non contiguous LBA ranges

Benefits:

- Simplifies resilient system designs
 - Database, file system, etc.
- Improves system performance in these applications

Hinting & Other NVM Features

Pass “hints” to devices to make operations more efficient and increase performance

- Targeted at SSDs and hybrid drives, but also useful for HDDs

Direct attached devices don't need to continually OPEN and CLOSE connections

- Can be implemented within the existing standard
- Reduces latency on both SSDs and HDDs

NVM features and programming interfaces

- Leverages ongoing work in SNIA and T10



Summary

Logical SCSI lives across the Entire Storage Spectrum

Proven Enterprise attributes & growing footprint

Flexible architecture = Platform for Innovation

- Enterprise reliability, improved connectivity
- Low latency device performance, capacity scaling
- Multitude of other storage dimensions

A Multi-generational Investment

- Enterprise-hardened Middleware
- Preserves backward compatibility
- Roadmaps continue to preserve legacy investments

SAS, SCSI, and SCSI Express: Extending the Platform of Storage Innovation