Education

SAS & SATA Combine to Change the Storage Market

Marty Czekalski, Seagate Technology
Terry Gibbons, LSI Corporation

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- SAS and SATA Combine to Change the Storage Market
  - Serial Attached SCSI (SAS) allows systems to be built that accommodate large numbers of either SAS and/or SATA hard disk drives. This presentation, intended for OEM, System Builders and End-Users, describes the capabilities of the SAS interface, how it's designed to interoperate with SATA drives, and when combined, how these technologies can be combined to deliver some very compelling storage solutions.

SAS
Preserving the Past, Creating the Future

Education SNIA

Preserve Legacy SCSI

Preserves 25 Years of SCSI Middleware

Customer Choice

3.5" & 2.5" Form Factors

Plug Compatible

Multi-protocol

Usability

- Dual-ported
- Point-to-Point
- HDD, Tape, SSD
- Cost equal to SCSI

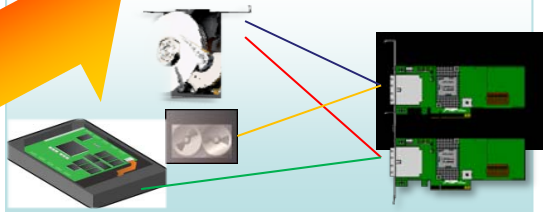

Serial Attached SCSI

Future Architected

- Protocol extends to new technologies
- Serial, Switchable
- SFF Connectors

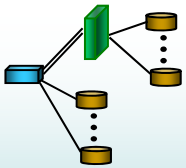

Scalable

1000's of Connections

Performance

Wide-Ports

Low Overhead

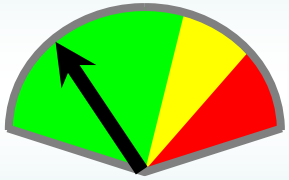

SAS Evolution
Supporting Key Storage Trends

Education SNIA

	Serial Attached SCSI	6Gb/s Serial Attached SCSI
T10 Specification	SAS-1 & SAS1.1	SAS-2
Distinguishing Features	➢Preserves legacy SCSI ➢SATA compatibility	**Expands SAS beyond traditional DAS Usage**
Storage Features Supported/Enabled	➢RAID 6 ➢Small Form Factor ➢HPC ➢High Capacity SAS Drives ➢Ultra320 SCSI replacement ➢Customer Choice ➢Blade servers	

Blind Mating and Hot Pluggable

- SATA Disk Connectors
  - Suitable for Both 3.5” and 2.5” Storage Devices
  - Includes data and voltage connections
  - Hot-Pluggable (staggered pins)
  - Blind Mating

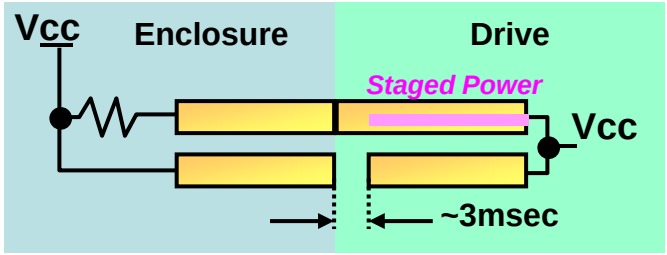

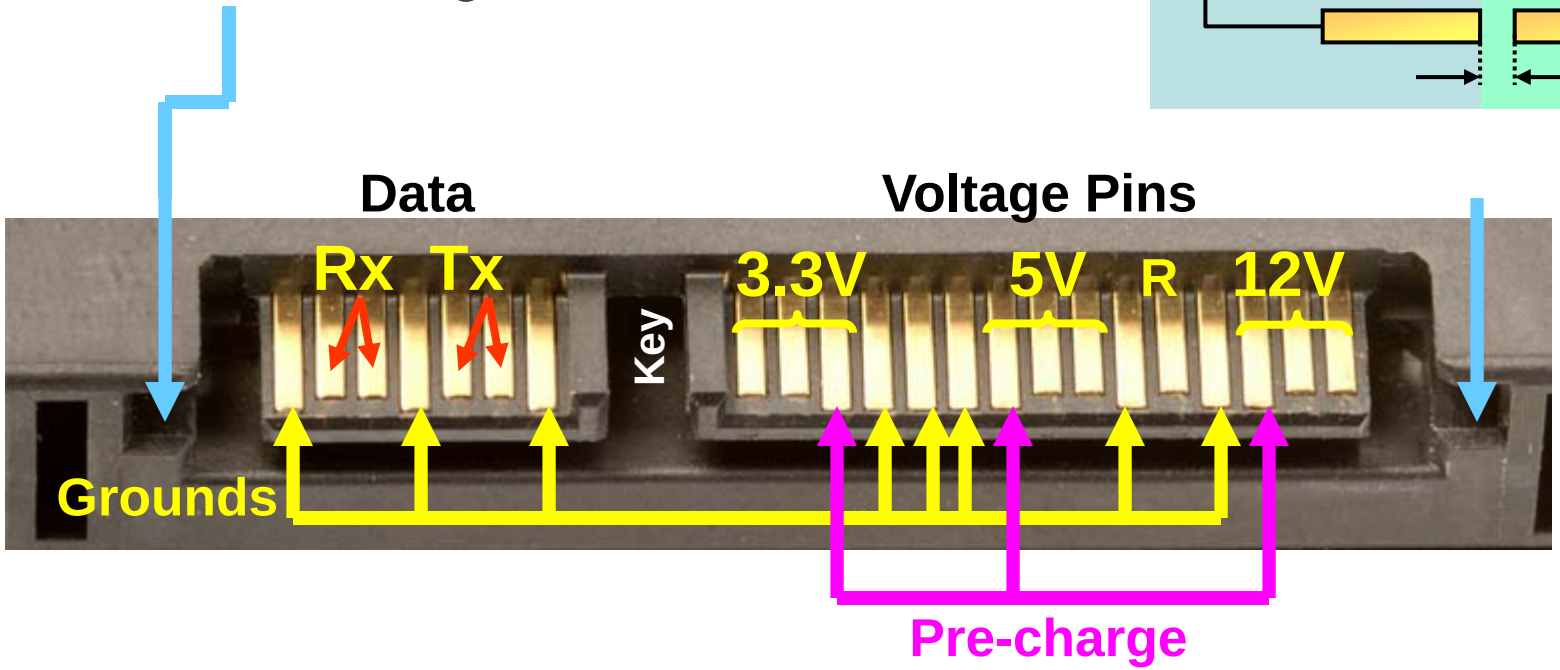

SAS/SATA Compatibility

Education SNIA

Disk Drive Connectors

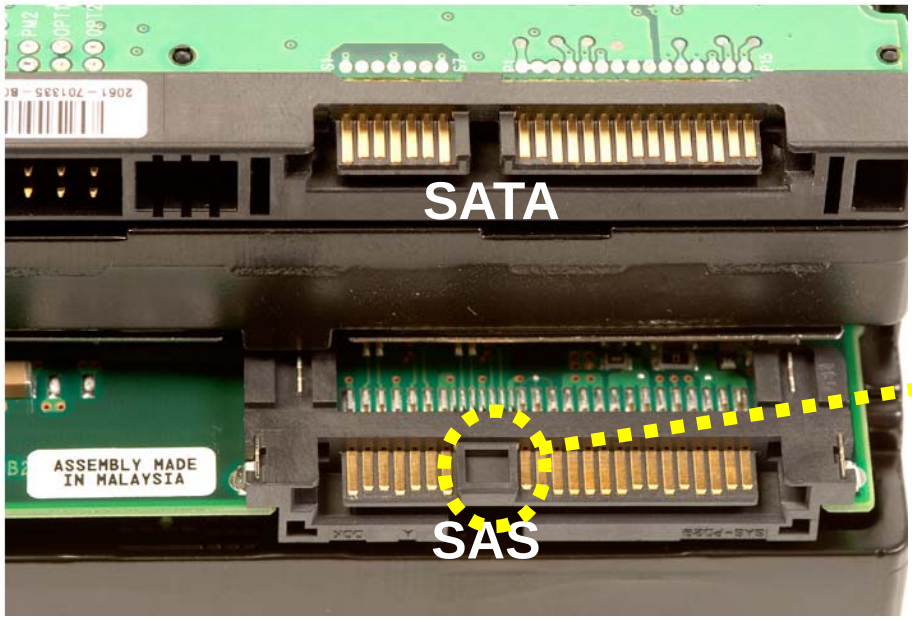

SAS Connector Flip Side

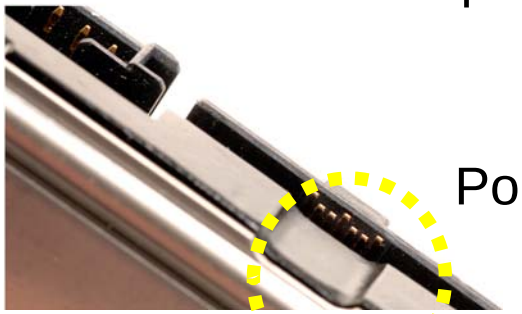

SAS Backplane Connector

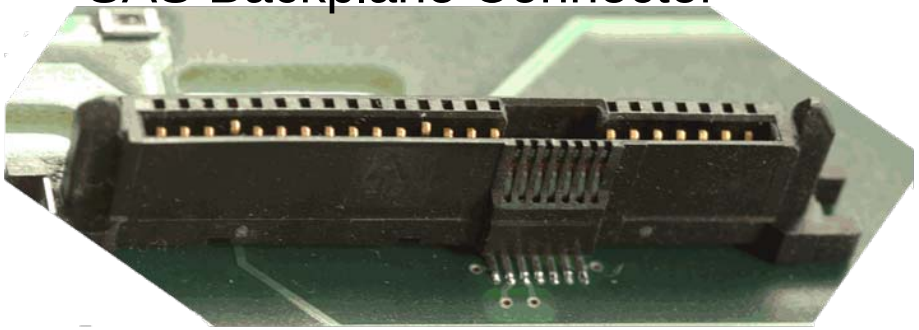

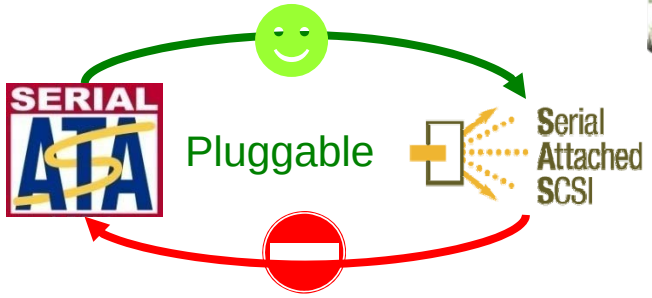

Accommodates both SAS & SATA Drives

World-wide HDD Shipments

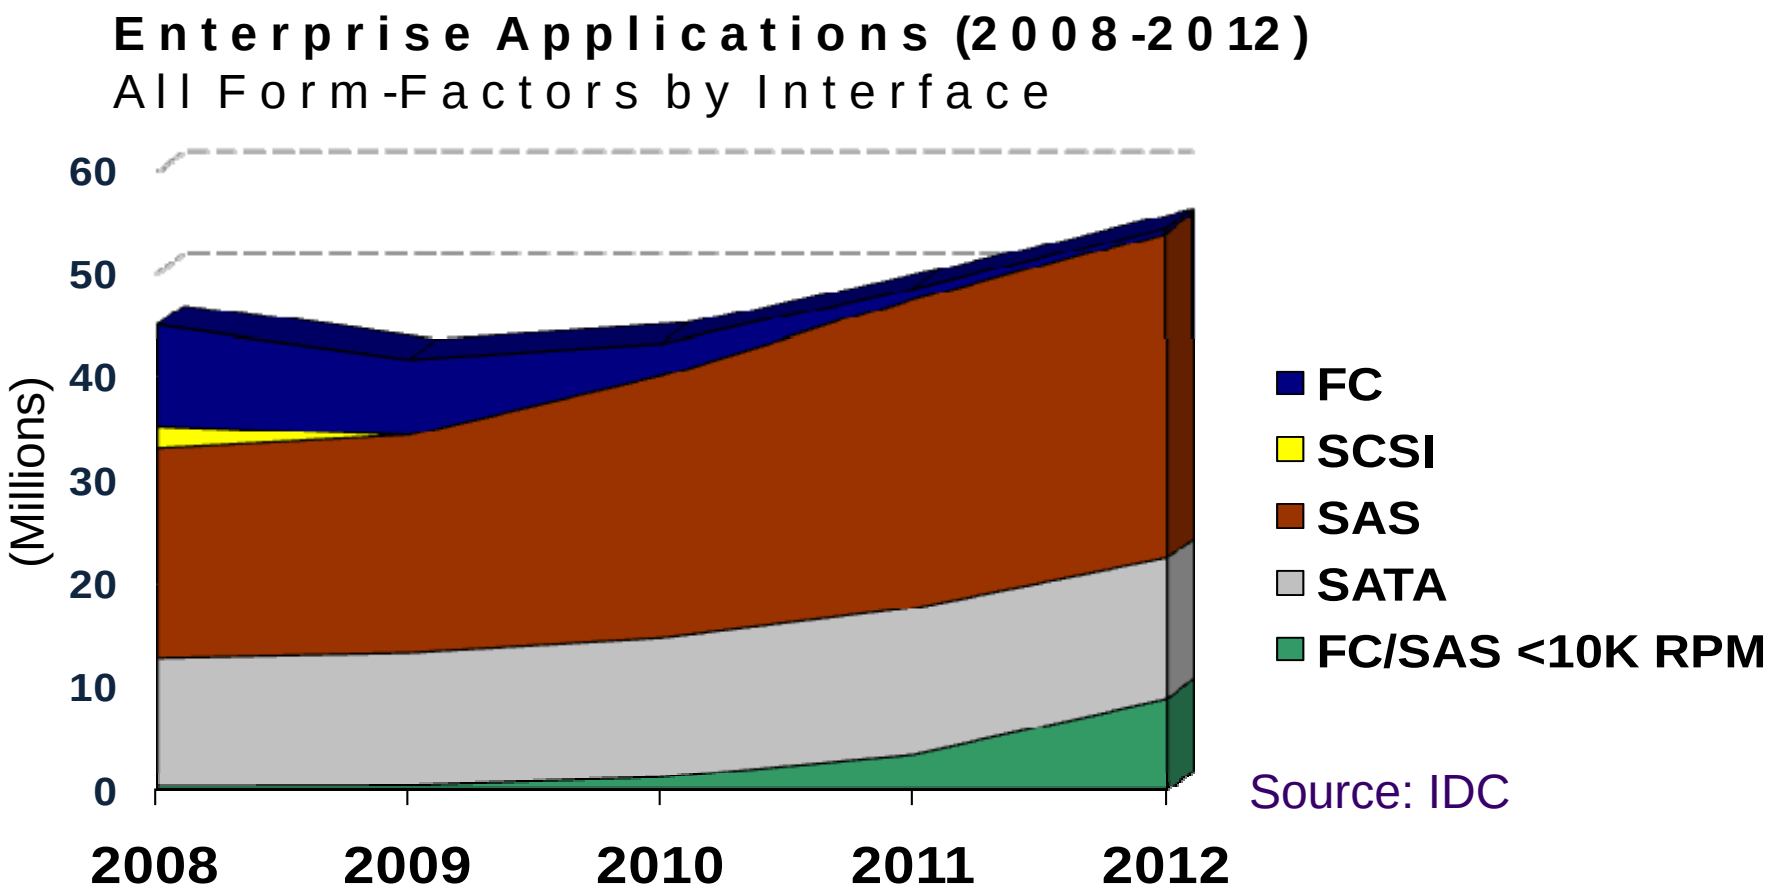

Source: IDC

A significant percentage of SATA in the enterprise will be deployed on SAS

Storage Growth & Innovation Drives Customer Choices

Prevailing wisdom: Storage Consolidation was accomplished through networked storage solutions (NAS and SAN)

Virtualization & Live Migration drive Networked Storage

Virtualization

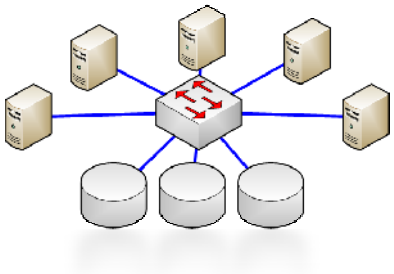

Networked Attached Storage

Direct Attached Storage

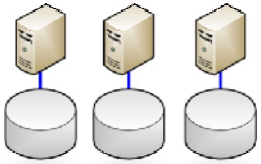

Application Intelligence

Application intelligence driving applications back to DAS to reduce cost and complexity

Applications Get Smarter

- Applications gaining intelligence
  - Applications growing new capabilities without requiring SAN storage model
  - High-availability clustering through application based local replication
  - Disaster recovery through application based remote replication
  - Greater performance through tight integration with storage subsystems
- Messaging Applications:
  - Reduced I/O load profile reduces per user spindle count requirements,
    - Reduce costs and/or
    - Enabling greater per server capabilities (user count/mailbox size)
  - Clustering techniques, like Cluster Continuous Replication (CCR) and Standby Continuous Replication (SCR), enable robust fault-tolerant deployments for DAS deployments

The Direct Attached Storage (DAS) Revival

- New DAS technologies are more powerful than ever
  - Serial Attached SCSI (SAS) and Serial Attached ATA (SATA) deliver a new level of performance and cost
  - Improved application architectures coupled with these interfaces, improves economics of storage
  - Reduced latency important for SSDs
- SAS based storage delivers high performance and high capacity storage for transactional deployments
  - SAS JBODs coupled with CCR and SRC capabilities, deliver a new level of price performance for messaging environments
  - SAS JBODs offer a breadth of storage devices to satisfy the IO processing requirements of databases and web services

Who Consumes Enterprise Drives?

	DAS deployments (Drive Interfaces by volume)	**SAN & NAS** deployments (Drive interfaces by volume)
2001	approximately **70%** (SCSI, P-ATA)	approximately **30%** (FC, SCSI, P-ATA)
2007	approximately **70%** (SAS, SCSI, SATA)	approximately **30%** (FC, SAS, SATA)
2013	TBD % (6Gb/s SAS, SATA)	TBD % (6Gb/s SAS, SATA)
Market Drivers	Application Intelligence, Price, Power, Performance DAS & S/W Improvements Flash	Virtualization Ease of Use Pervasive Networks

Source: Seagate(/Maxtor) Market Research

SAS's Unique Attributes

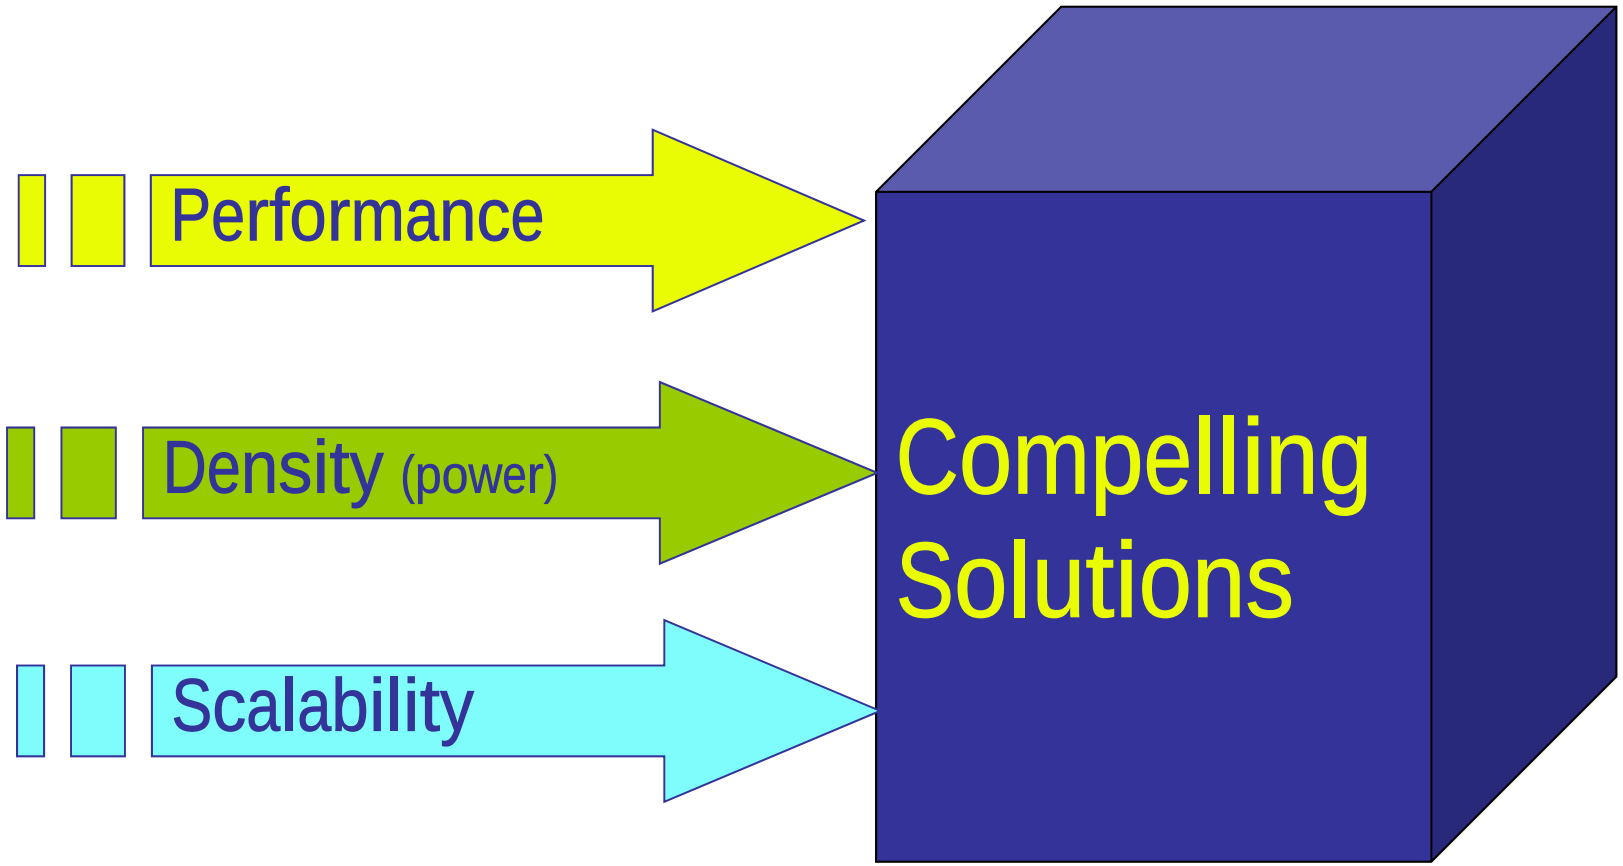

SAS: Bandwidth Aggregation

Education SNIA

Performance

- **Each SAS Link (Rx and Tx)**
  - 3Gb/s → 6Gb/s (full-duplex)
  - 6Gb/s → 12Gb/s (full-duplex)
- **Wide Ports**
  - Combine SAS links (6Gb/s SAS)
    - 2 ports – 24Gb/s (full duplex)
    - 4 ports – 48Gb/s (full-duplex)
- **Concurrency Brings Higher Performance**
  - Multiple concurrent I/O's

    (lots of drives operating concurrently)

Wide-Port (x4) yields 48Gb/s Full Duplex!

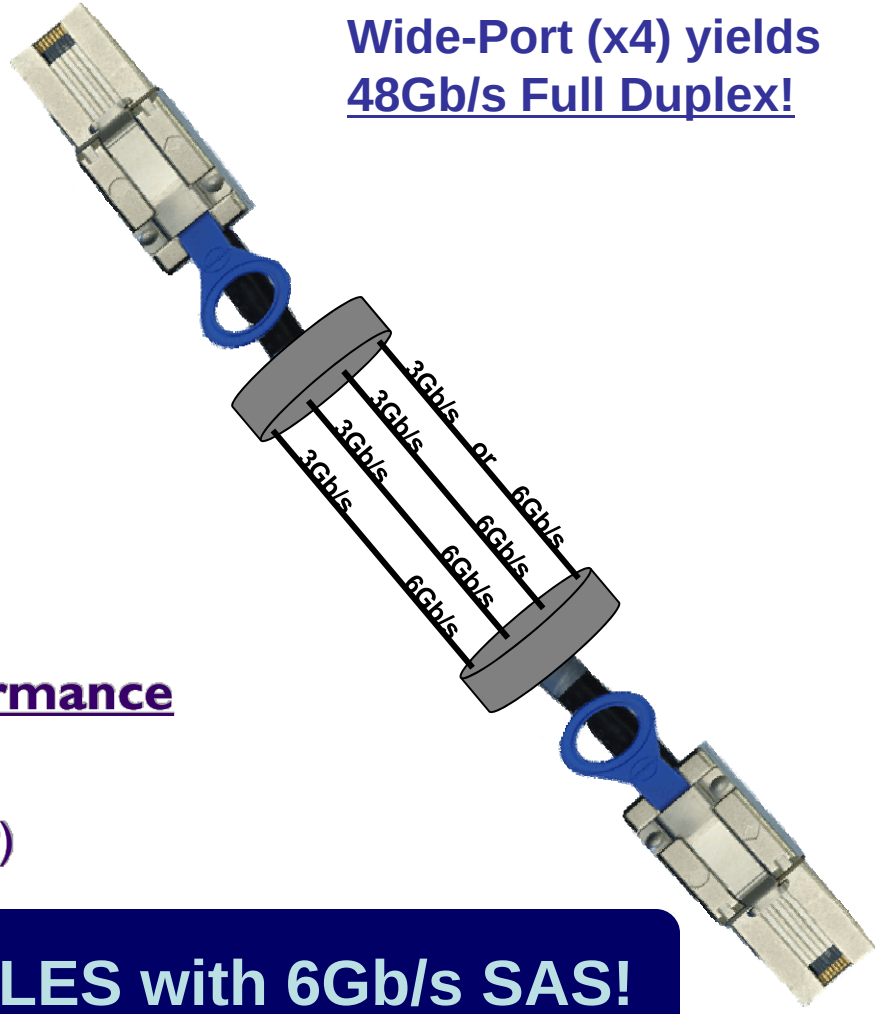

Theoretical Bandwidth DOUBLES with 6Gb/s SAS!

6Gb/s SAS: Significance to SSDs

Education SNIA

Performance

- Time to Market
  - Established Software
  - Low Integration Risk
  - Rapid System Qualification
  - Standards already exist
- Performance
  - High Bandwidth and IOPS
  - Years of Software Refinement
  - Advancements for Low Latency Storage
- Scalability/Serviceability
  - Logical Abstraction Layer
  - Infrastructure Supports Large Scale-outs
  - Field serviceable
- Technology Neutral - Market Resilience
  - Features Migrate seamlessly across OS's & Devices
  - Effective Platform for spurring innovation

Industry Compacts Storage

Education SNIA

Density

- SAS & Small Form Factor
  - SFF increases density
    - Scale & expand in same space
    - Ideal for Blade System
  - SFF Drives Consume 50% less Power than 3.5" counterparts
    - Lower heat loads & cooling costs at equivalent performance
      - Caution – Power per rack can still increase due to density increase
  - Higher Performance and Potentially Lower Cost
    - Multiple concurrent I/O's (lots of drives operating concurrently)
    - Lower TCO through common infrastructure
  - Additional Advantages
    - RAID 5 or 6 on a 1U server
    - SAS & SATA Drives in Common Drive Carrier

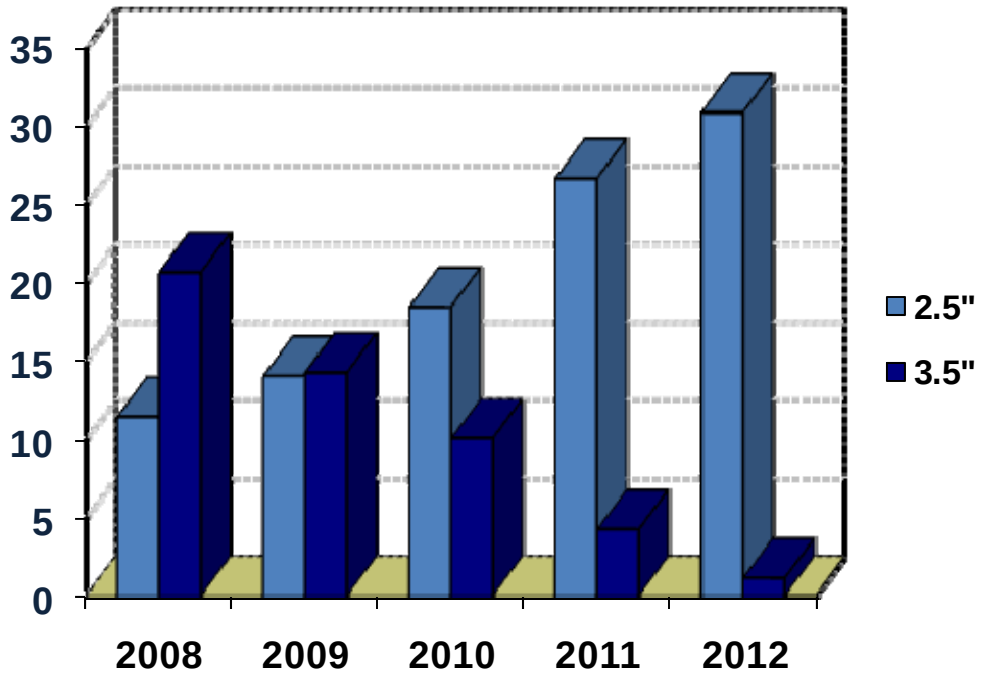

WW HDD Shipments (Millions of Units): Enterprise Applications by Form Factor Performance Optimized (FC, SCSI, SAS) Source: IDC

Connector Types

Density

InfiniBand Style

SAS 4X Connector

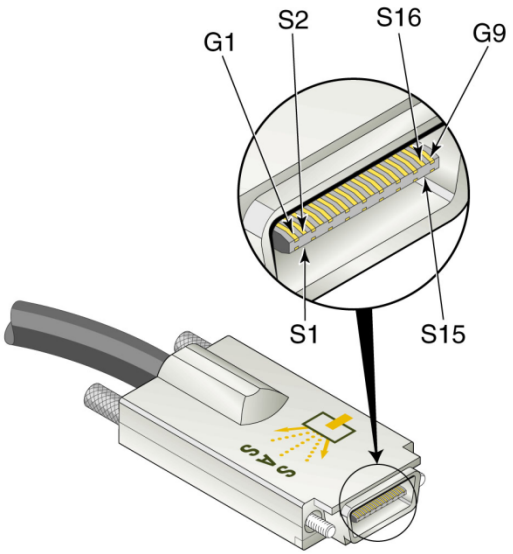

Mini SAS 4X Connector

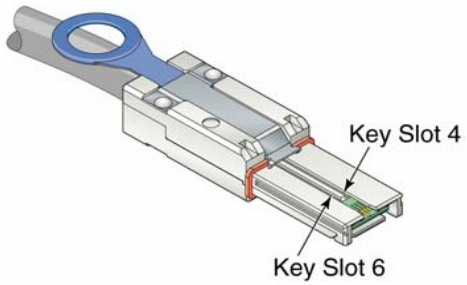

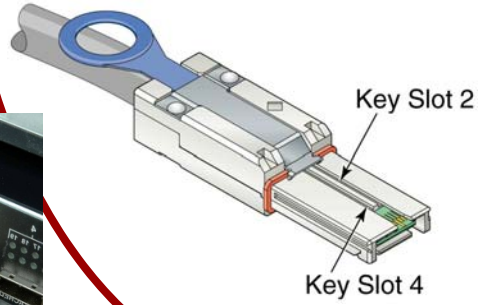

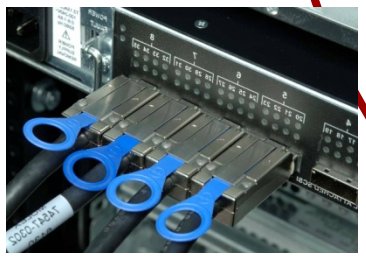

Preferred External Connection Scheme

Scaling outside the Server

Education SNIA

Density

Expand Your DAS Capabilities

and/or

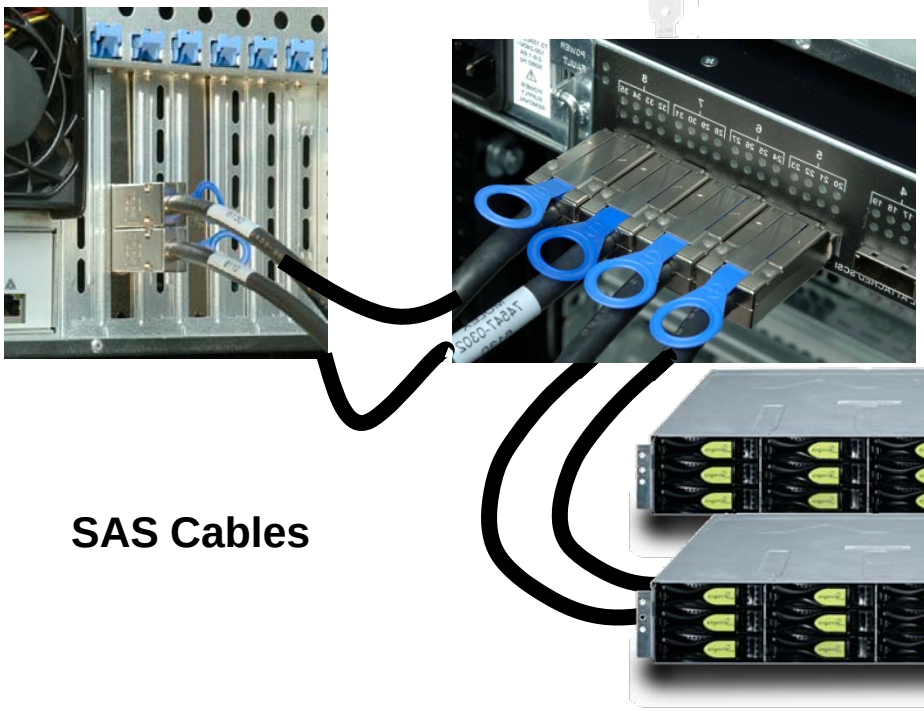

SAS Cables

Cascaded Enclosures

SAS: Improving Density and Airflow

Density

SATA/SAS Cable
VS Parallel

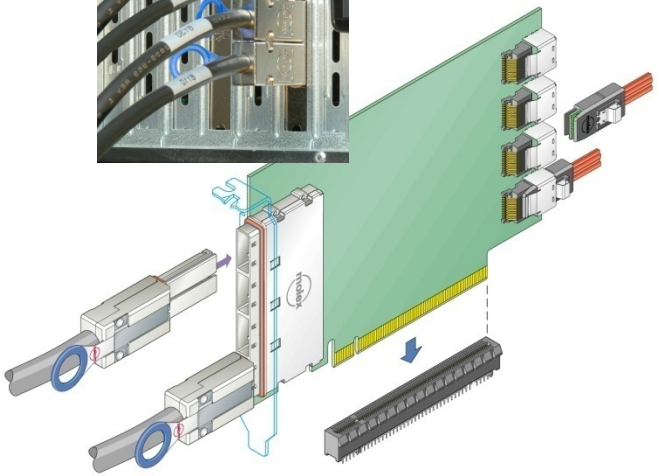

External & Internal
SAS 1.1 Mini SAS
(4x & 4i)

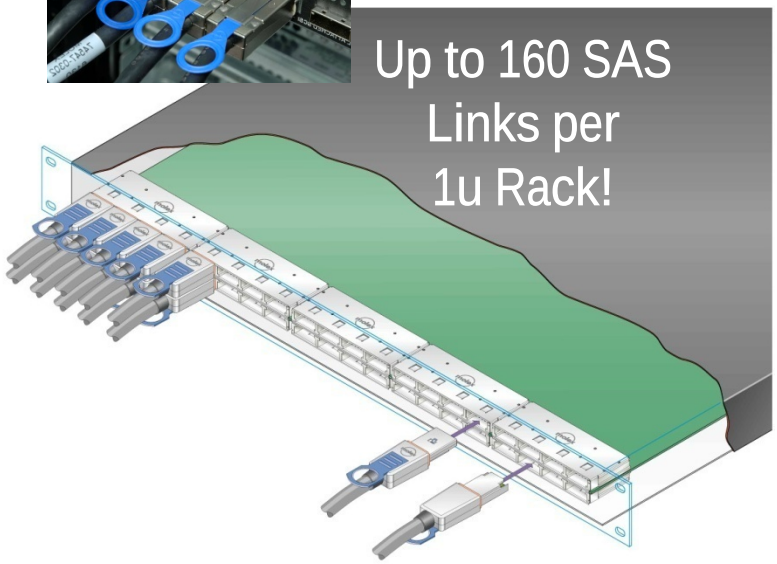

Highly Scalable Architecture
Connectivity and Bandwidth

Dramatically Improves Connectivity, Density & Airflow

Self Configuring Expanders Extend SAS

Scalability

- **1st Gen SAS Discovery**
  - Vendor Unique Zoning
  - Limited to 128
  - Discovery Executed by Initiator end-devices
    - Complicated large topologies
    - Problematic for zoning
- **2nd Gen SAS Discovery**
  - Standardized Zoning
  - Expanded to 1K
  - Discovery executed by SAS Expanders
    - Reduces time to discover large topologies
    - Eases burden on IO flow due to disk add/pull
    - Enables zoning of the topology

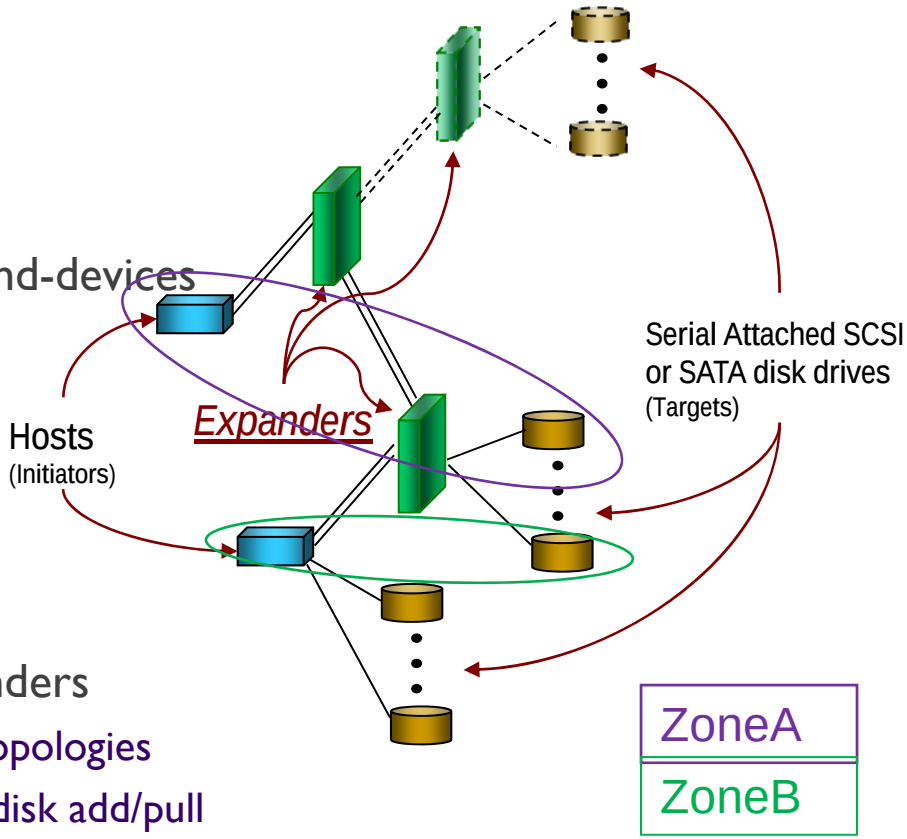

Enables Richer SAS Topologies

Scaling SAS Architecture

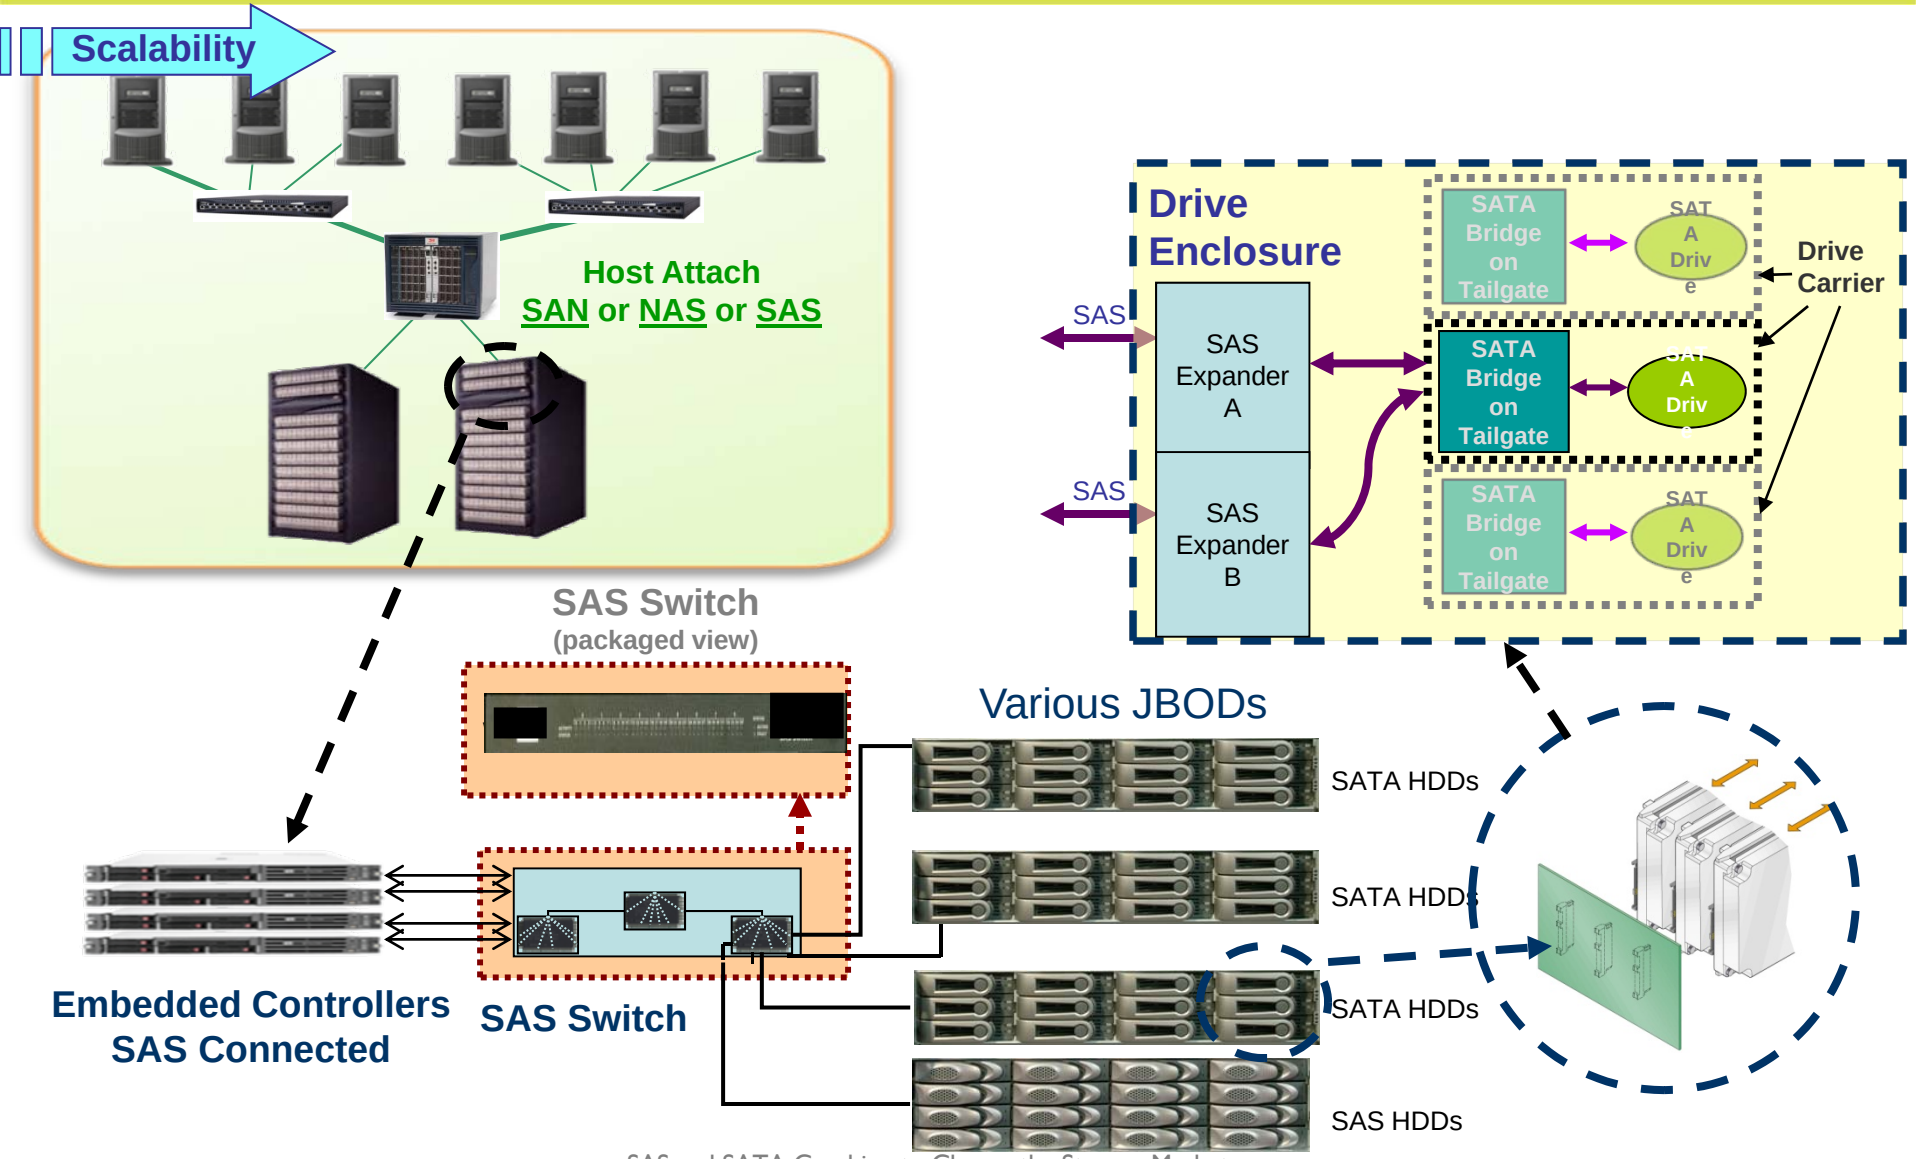

Cascading Enclosures

Education SNIA

Scalability

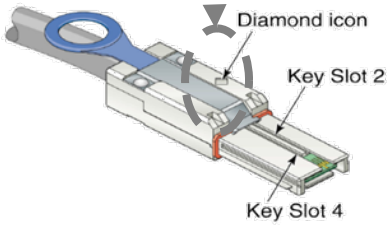

◆ = Cable Plug Connector to End Device or Enclosure OUT / Universal Port

● = Cable Plug Connector to End Device or Enclosure IN / Universal Port

●/◆ = Receptacle Connector for End Device or Enclosure Universal Port

Note: IN/OUT below may currently be substituted with Universal Port if Table-Table Routing is utilized

Note: Possible proposal to add Universal Cable Plug Connector

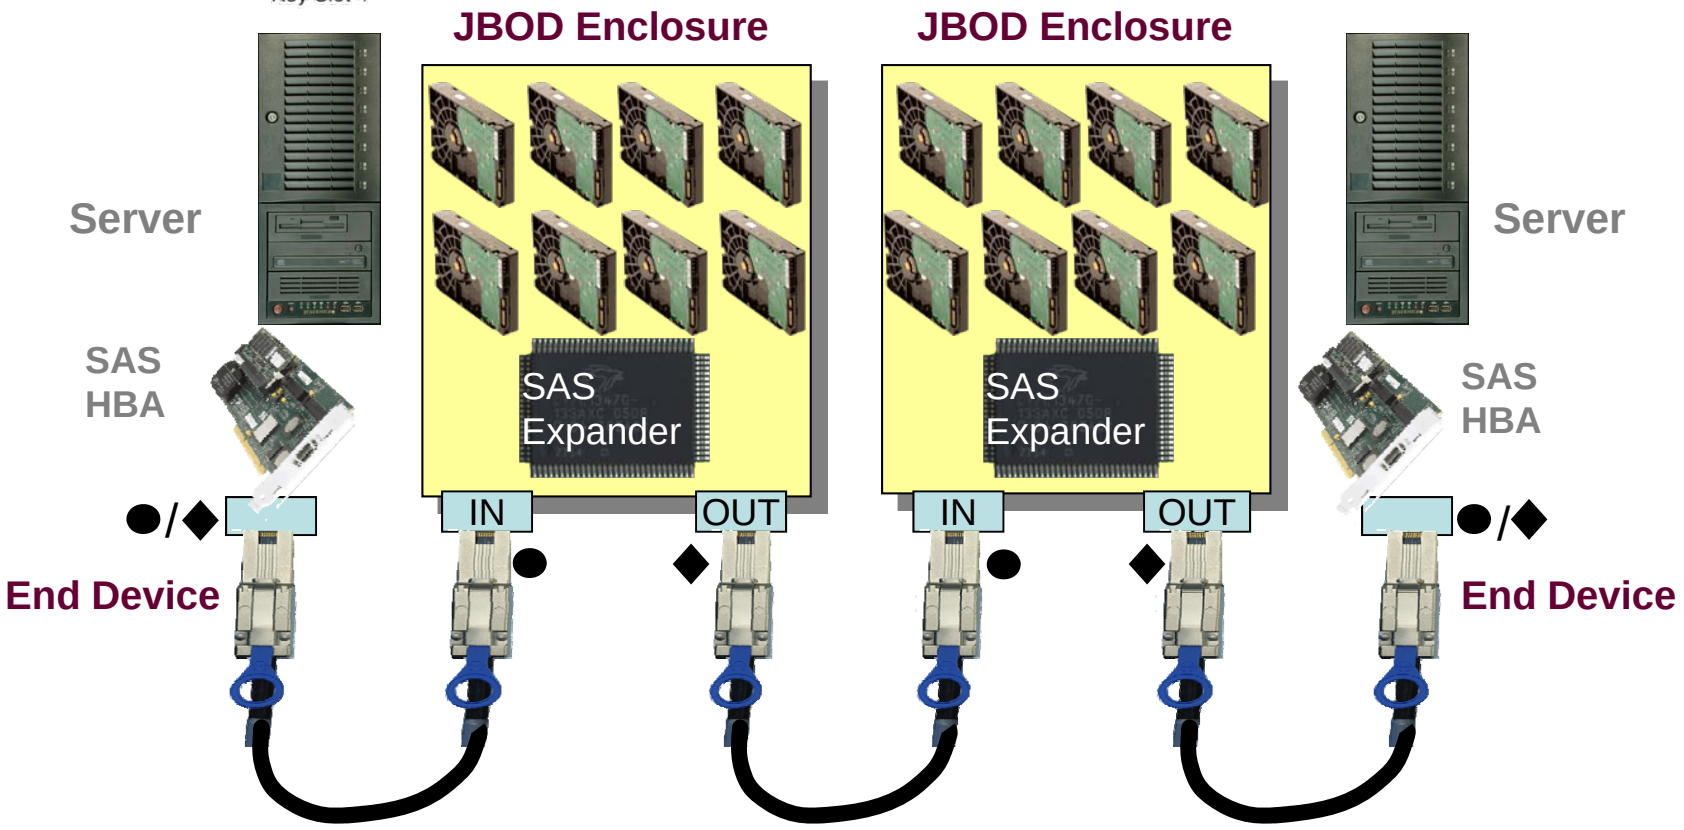

Switched SAS Applications

Solutions

Shareable DAS

RAID Array Drive Expansion

External Storage Controller to Drive Enclosure Expansion

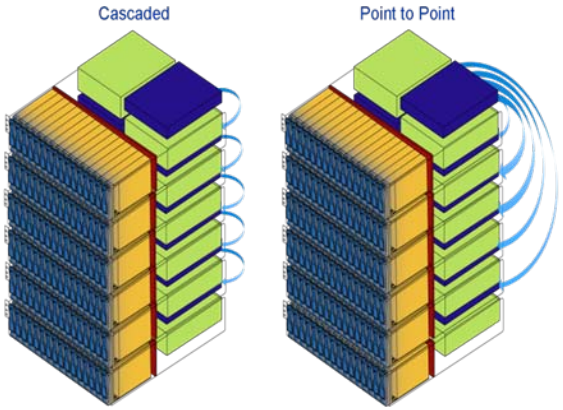

Blade Storage Mid-plane I/F

Lower Cost

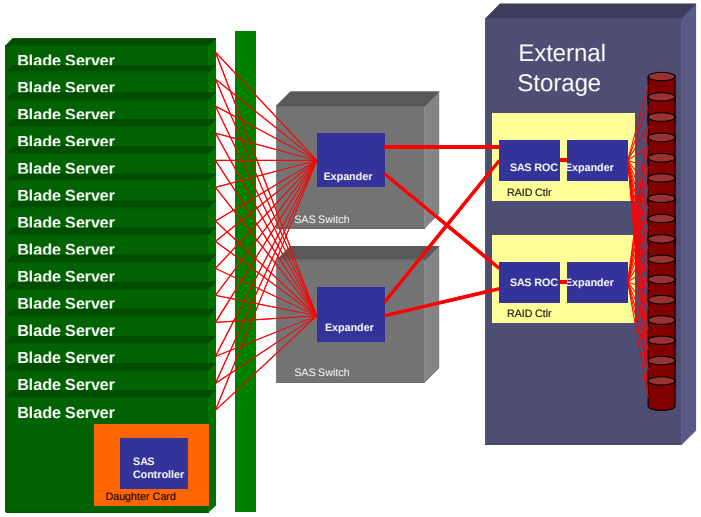

Higher Bandwidth

Benefits of Switched SAS

Solutions

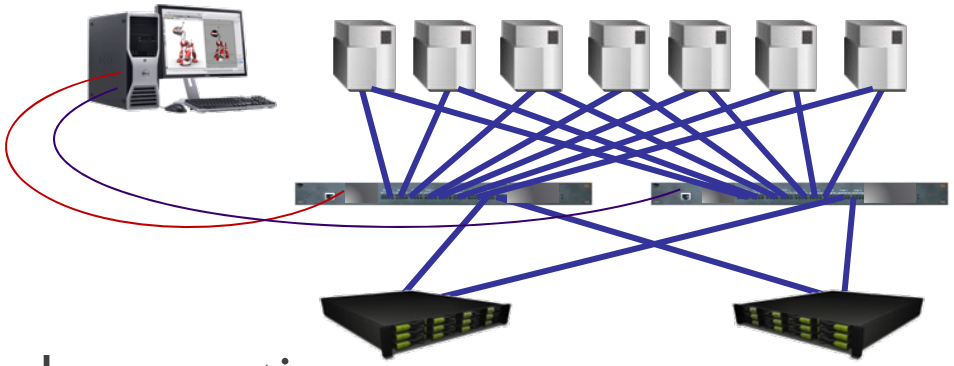

- Centralized Management
- Reliability
  - Improvement over cascaded connections
- Legacy protection/isolation
- Scalability
  - Thousands of storage devices
- Availability
  - Reduce latency and bottlenecks – improve fairness

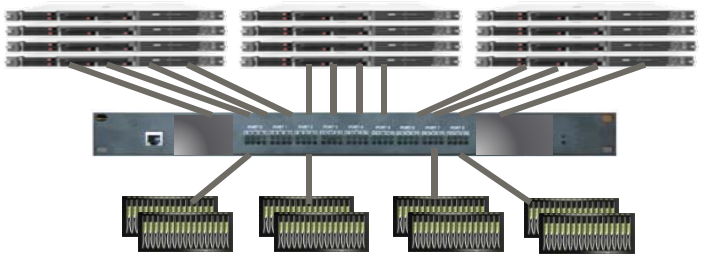

Scalable, Sharable DAS
SAS Value Proposition & Positioning

Education SNIA

Solutions

	Performance	Low Cost	Distance
Fibre Channel	X		X
GbE iSCSI		X	X
10 GbE iSCSI	X		X
SAS	**X**	**X**	

- Significant market segment underserved by SAN alternatives
  - Same room, modest scale (to few 10's of server & storage endpoints)
  - Compliments Application Intelligence and Application Messaging
  - Enhances Clustered Applications

SAS & SATA Span the Storage Spectrum

Education SNIA

Direct Attach Storage | **SAS Fabrics** | **External Storage** | **HDD/SDD**

SAS/SATA Connectivity Creates Storage Solutions

- Controllers/ROCs
- Expanders
- SAS/SATA HDDs
- SAS/SATA SSDs
- Storage Blades

- Expanders
- SAS Switches
- Bridges
- Port Multiplexers

- NAS/SAN Heads
- Native SAS Connect
- Controllers/ROCs
- Expanders
- SAS/SATA HDDs
- SAS/SATA SSDs
- SAS/SATA Tape

- SAS Drives
- SATA Drives
- Near-line SAS
- Drive Carriers
- Drive Controllers

Demonstrating SAS Infrastructure

Education SNIA

System Components

- One Quad-Core processor
- PCI-Express 2.0
- 3 – 6Gb/s SAS Controllers
- 30 Drives – (2.5" 6Gb/s 15K RPM, JBOD)
- Workstation motherboard using latest CPU & chipset
- Windows 2003

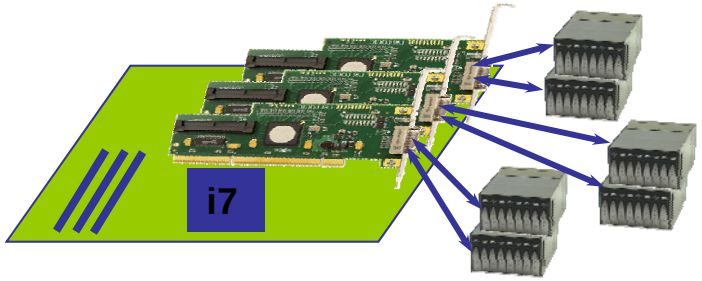

6Gb/s Serial Attached SCSI

Performance

- Iometer Throughput Benchmarks
  - 2KB sequential read or write
    - – **1,000,000 IOPs**
  - 256KB sequential read or write
    - – **6.5 GB/s**

Demonstrates the extremely high throughput available with standard high-volume components

SSD – Transaction Processing

System Components

- One Eight-Core processor / PCI-Express 2.0
- 6Gb/s SAS Controller
- 24 Drives – (2.5" 6Gb/s SAS 15K RPM, JBOD)
- Vs. 4 Drives – (3Gb/s **SATA SSD** – RAID 0)
- Windows 2003 64-bit

Performance

- Iometer Benchmark
- 4KB random workloads typical of Messaging, Databases, and Web Services
- 24 HDD vs. 4 SSD
  - IOPS =>11K vs. 68K (Random Write)

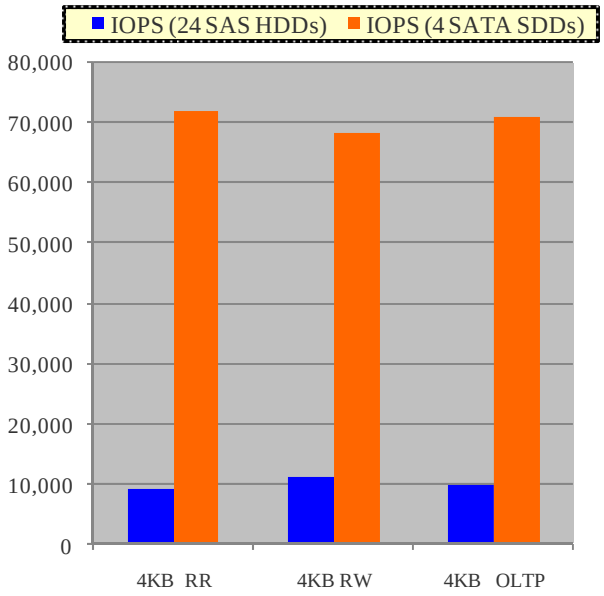

Reference

- SNIA SSSI TWG – Reference Test Platform Proposal – November 24, 2009

When SATA SSD is the choice – SAS/SATA controllers meet the IO processing requirements

What's coming?

- High Capacity SAS - Simplifies SAS Architecture
  - SAS interface on High Capacity Drives, 2 TB and higher
  - Architecturally more efficient, improves Enterprise System Integrity
- 6Gb/s SAS - Double transfer rate
  - 6Gb/s and 3Gb/s SATA compatible
  - Improved signaling and more efficient protocol
  - Data Protection Information (PI); ANSI T10/1799-D
- Data Integrity Extensions (DIX)
  - Proposed methods to increase efficiency of passing PI from OS to media
  - SNIA Architectural Model for Data Integrity (members only)
  - Oracle® Open Source Project: Linux Data Integrity Project
- Solid State Disks (SSD)
  - Random/Read performance improvements (10x to 100x)
  - Leverages existing middleware infrastructure
  - High demonstrated performance with further optimizations expected
- SAS Advanced Connectivity Roadmap

SAS Advanced Connectivity Objectives

- Drive market consistency
- Simplify cable and connector options
- Provide converged high-density connectivity
- Provide managed connectivity standards
- Provide active copper solution to 20m
- Provide optical solution to 100m

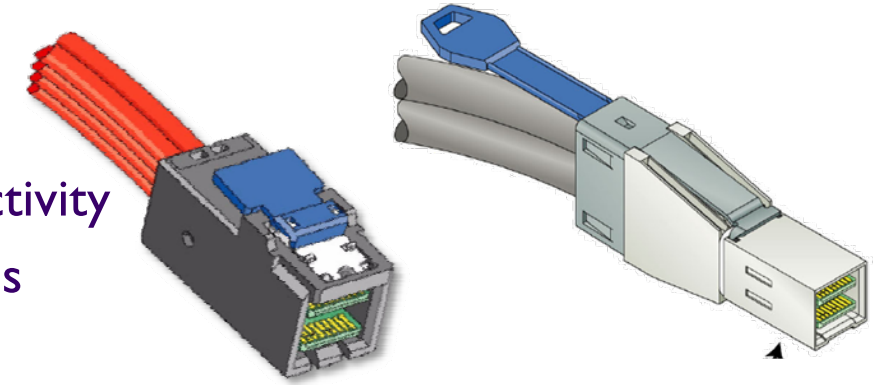

Internal similar to External

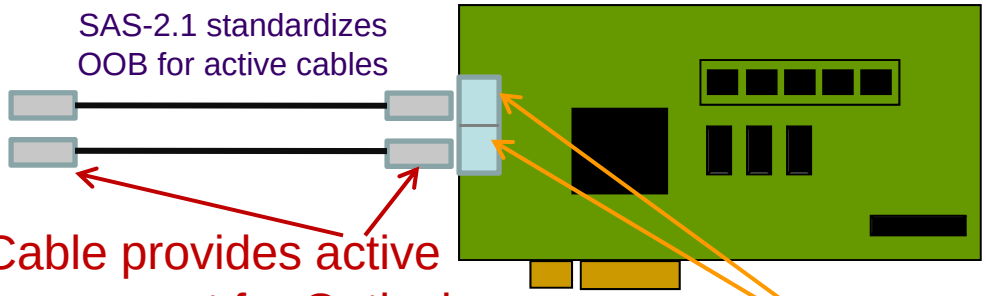

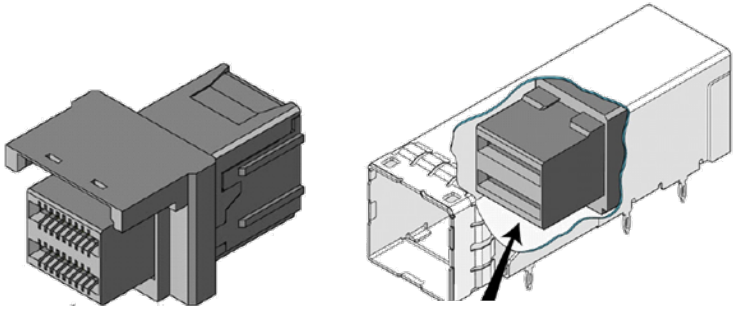

Passive, Active Copper, or Optical use same connector

Mini-SAS HD connectors courtesy:
Project T10/2125-D Revision 04
17 September 2009
American National Standard
Serial Attached SCSI - 2.1 (SAS-2.1)

- Support 6Gb/s SAS deployments
- Extensible to 12Gb/s SAS deployments

SAS Advanced Connectivity

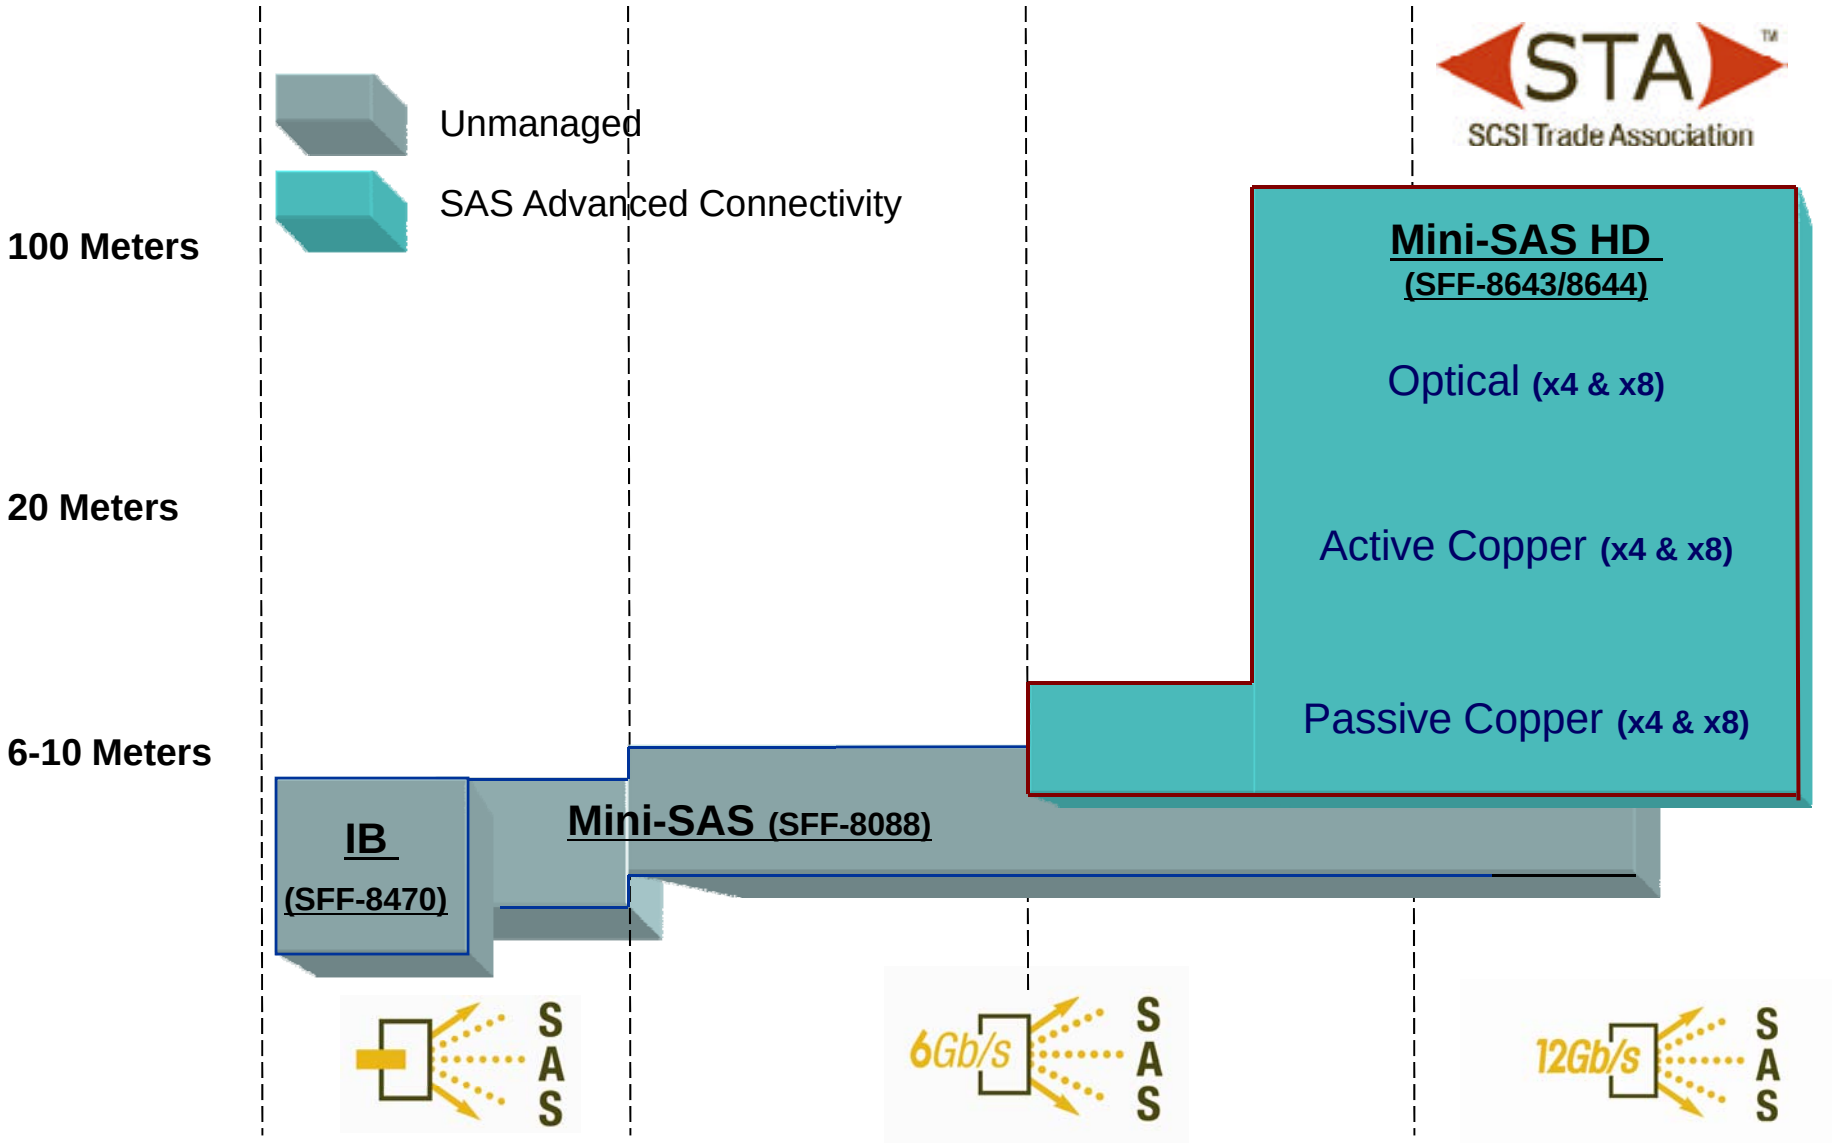

Additional Info Available at:

T10 (Serial Attached SCSI Spec development)
http://www.t10.org
SCSI Trade Association
http://www.scsita.org
Serial Storage Wire
http://www.serialstoragewire.com
SATA I/O
http://www.sata-io.org

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: trackstoragemgmt@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Marty Czekalski
Harry Mason
Jay Neer
Brandy Hammack
Benjamin Woo
Terry Gibbons
Nancy Clay
Rob Peglar
Wendy Betts