



Education

# SAS: The Emerging Storage Fabric

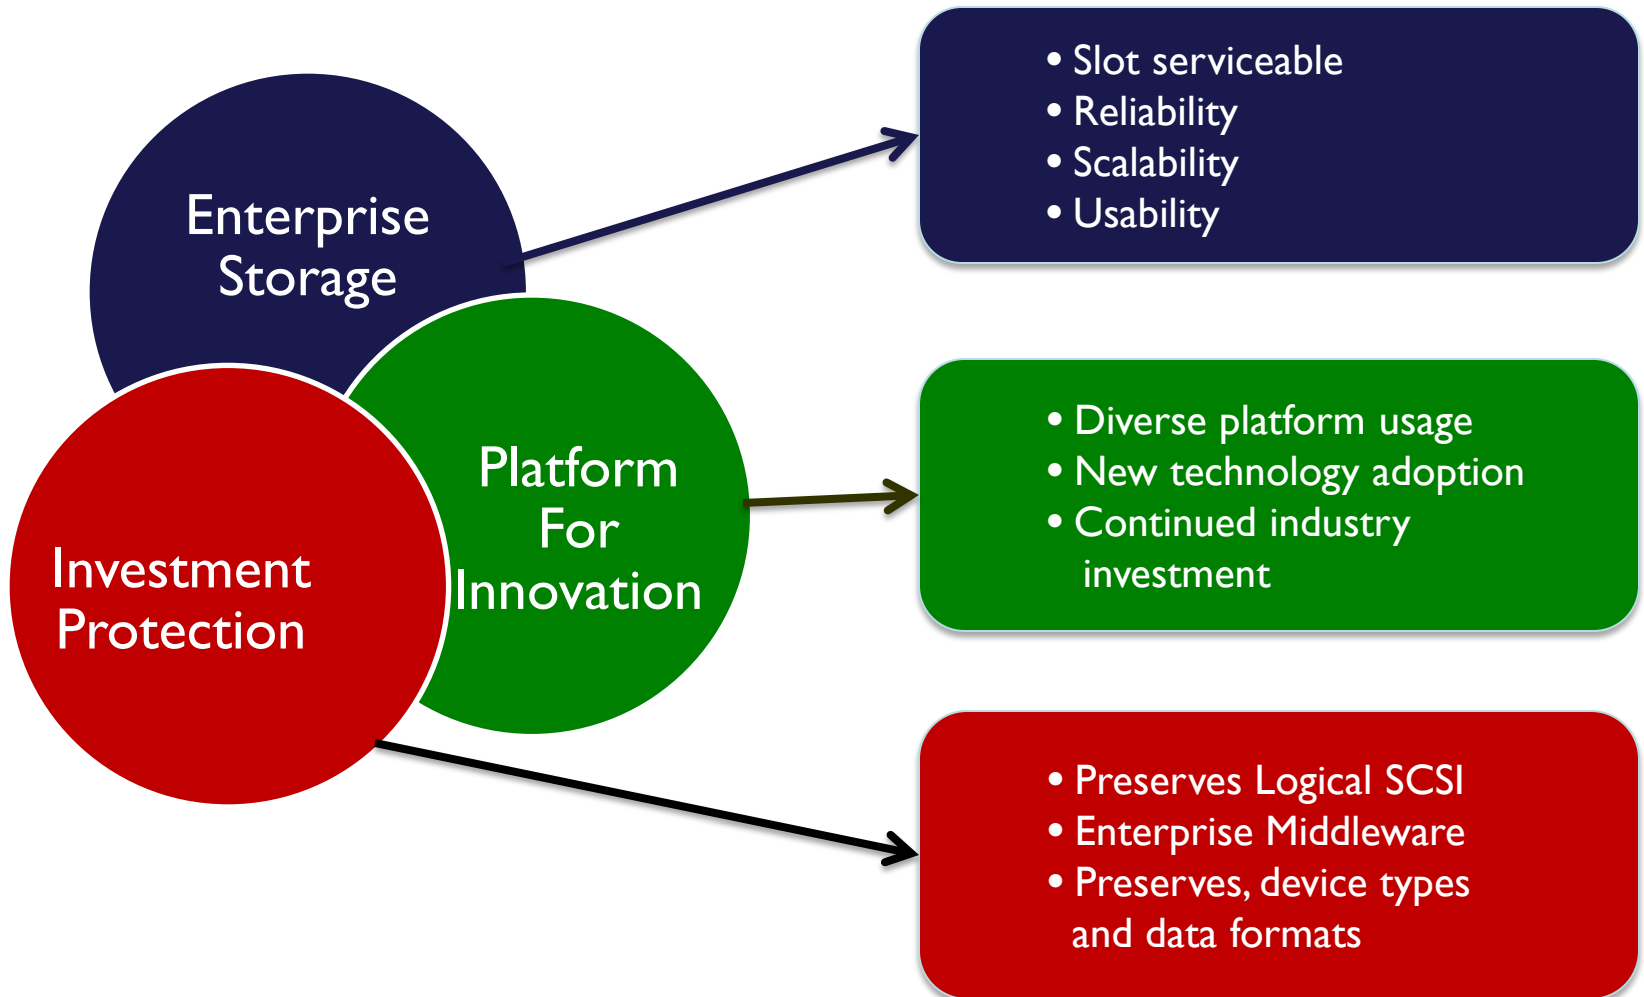
Harry Mason – President, SCSI Trade Association  
Marty Czekalski – Vice President, SCSI Trade Association

Author: Harry Mason – President, SCSI Trade Association  
Author: Marty Czekalski – Vice President, SCSI Trade Association

- 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.**


- SAS: The Emerging Storage Fabric
- SAS is the backbone of nearly every enterprise storage deployment, rapidly evolving, adding new features, enhanced capabilities and offering “no compromise” system performance. SAS not only excels as a device level interface, its versatility, reliability and scalability have made it the connectivity standard of choice for creating new Enterprise storage architectures.
- This presentation covers the advantages of using SAS as a device interface, and how its capabilities as a connectivity solution, are changing the way data centers are being deployed. Advantaging 12 Gb/s transfer rates, bandwidth aggregation, SAS Fabrics (including switches) active connections, and multi-function connectors (connectors that support SAS as well as PCIe Attached Storage devices) allows data center architects to create sustainable storage solutions that scale well into the future.

# SAS – Preservation and Innovation



# SAS & SATA Span the Storage Spectrum

## Direct Attach Storage



- **Controllers/ROCs/HBAs\***
- **Expanders**
- **SAS/SATA HDDs**
- **SAS/SATA SSDs**
- **Storage blades**

## SAS Fabrics



- **Expanders**
- **SAS switches**
- **Bridges**
- **Port multiplexers**

## External Storage



- **NAS/SAN heads**
- **Native SAS connect**
- **Controllers/ROCs/HBAs\***
- **Expanders**
- **SAS/SATA HDDs**
- **SAS/SATA SSDs**
- **SAS/SATA tape**

## HDD/SSD



- **SAS drives**
- **SATA drives**
- **Near-line SAS**
- **Drive carriers**
- **Drive controllers**

\* ROC = RAID on a Chip  
 HBA = Host Bus Adapter


# SAS Technology & Market Overview

## ➤ Basics of SAS Architecture & Deployment

# SAS Market Evolution Preserving the Past, Creating the Future

## Preserve Legacy SCSI

- 25 years of SCSI middleware




## Customer Choice

- 3.5" and 2.5" form factors
- Plug compatible
- Multi-protocol




## Usability

- Dual-ported
- Point-to-point
- Cost equal to SCSI




## Future Architected

- Protocol extends to new technologies
- Serial, switchable
- SFF connectors



## Scalable

- 1000s of connections




## Performance

- Wide ports
- Low overhead



# SAS/SATA Compatibility


## Disk Drive Connectors




## SAS Connector Flip Side



## SAS Backplane Connector

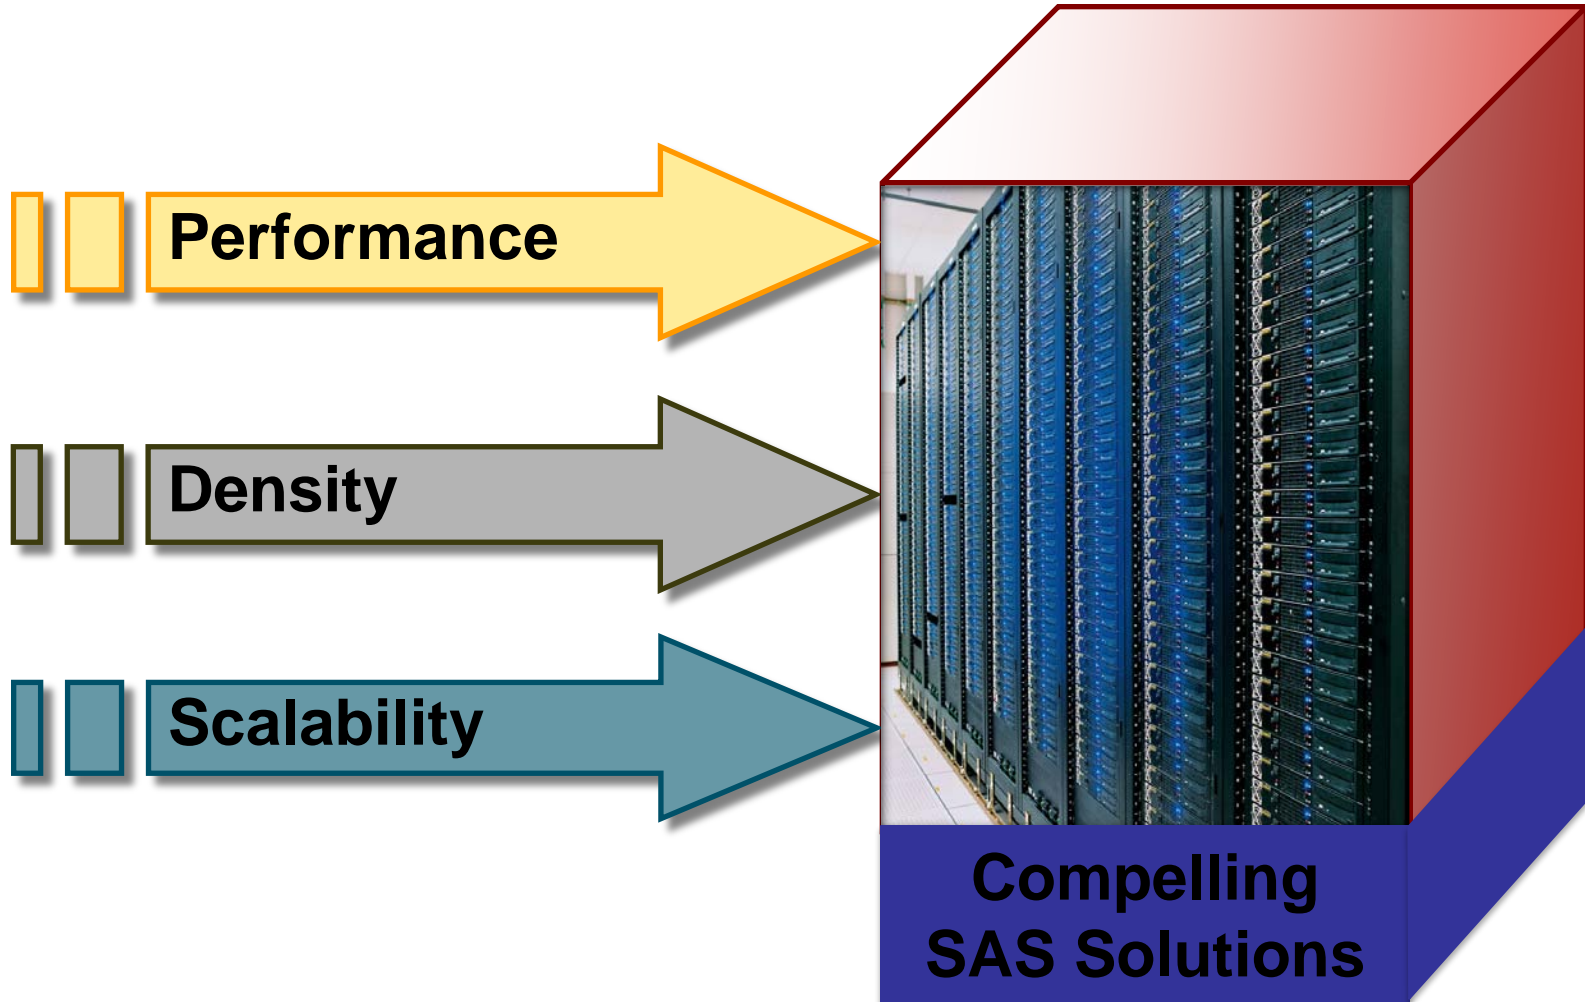


*Accommodates both  
SAS & SATA Drives*





# Technology Overview: Pulling it all together




# SAS: Bandwidth Aggregation

## Performance

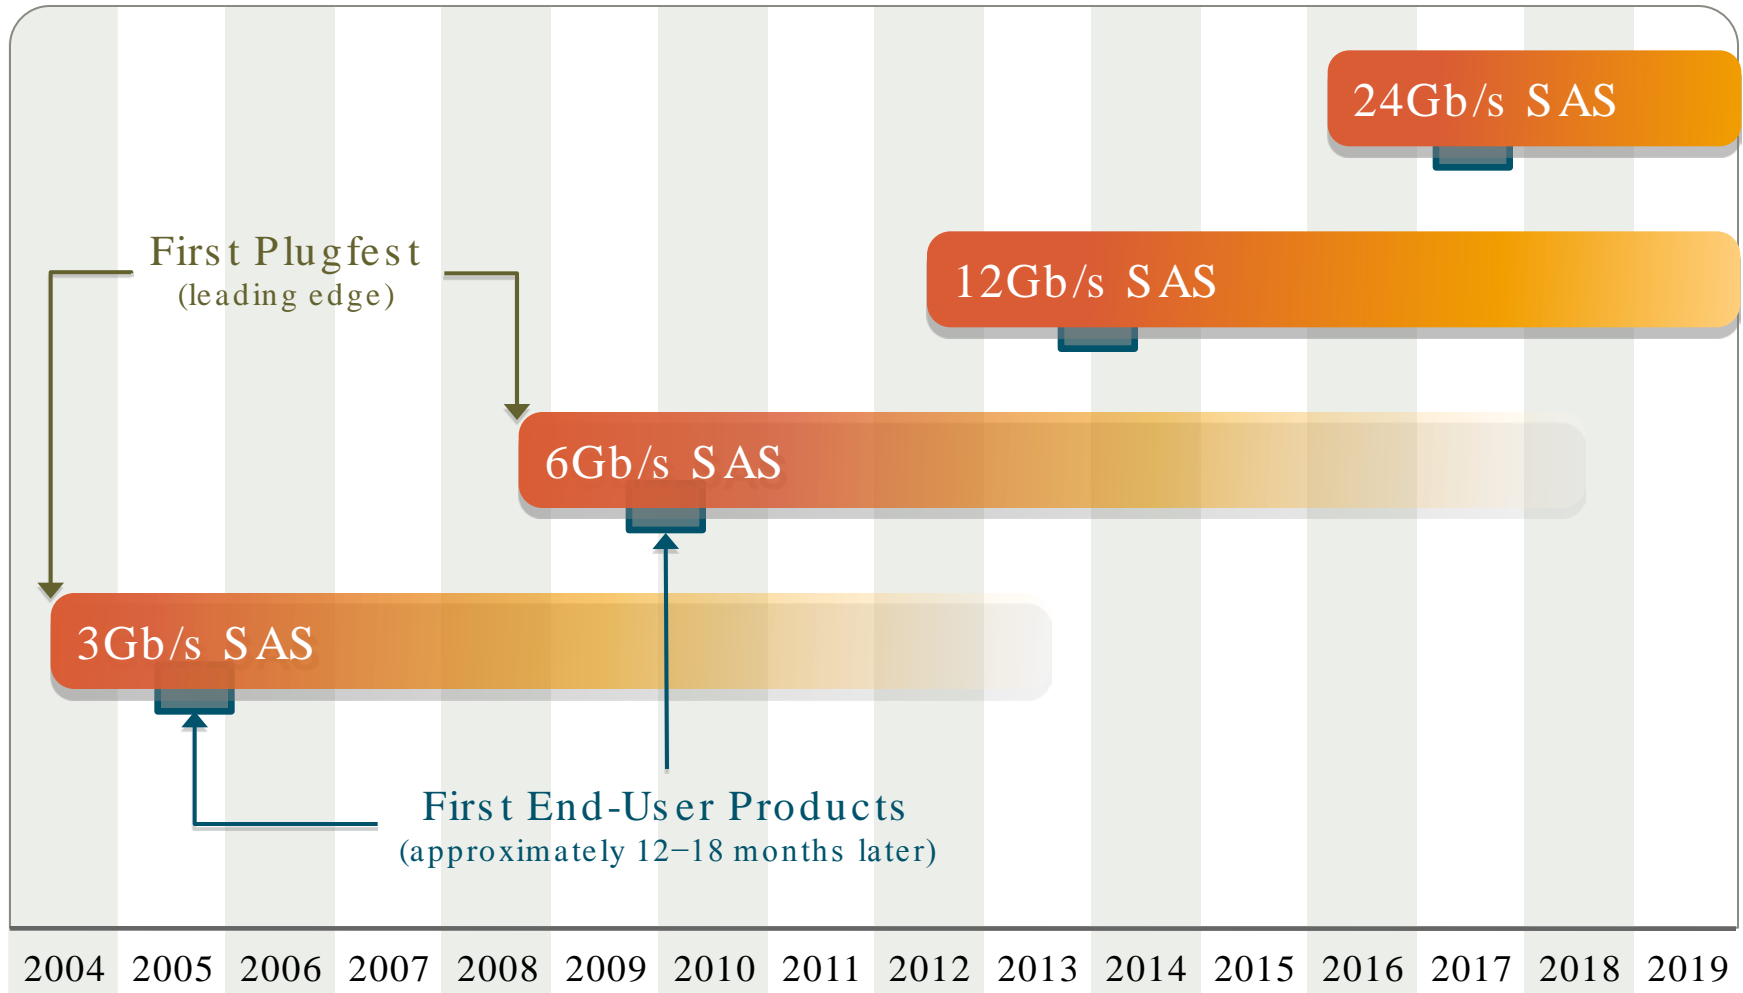
- Each SAS Link (Rx and Tx)
  - ◆ 6Gb/s > 12Gb/s (full-duplex)
  - ◆ 12Gb/s > 24Gb/s (full-duplex)
- Wide Ports
  - ◆ Combine SAS links (12Gb/s SAS)
    - > 2 ports > 48Gb/s (full duplex)
    - > 4 ports > 96Gb/s (full-duplex)
- Concurrency Brings Higher Performance Demands
  - ◆ Multiple concurrent I/O's  
(lots of drives operating concurrently)

*Wide-Port (x4) yields  
96Gb/s Full Duplex!*



***Bandwidth DOUBLES with 12Gb/s SAS!***

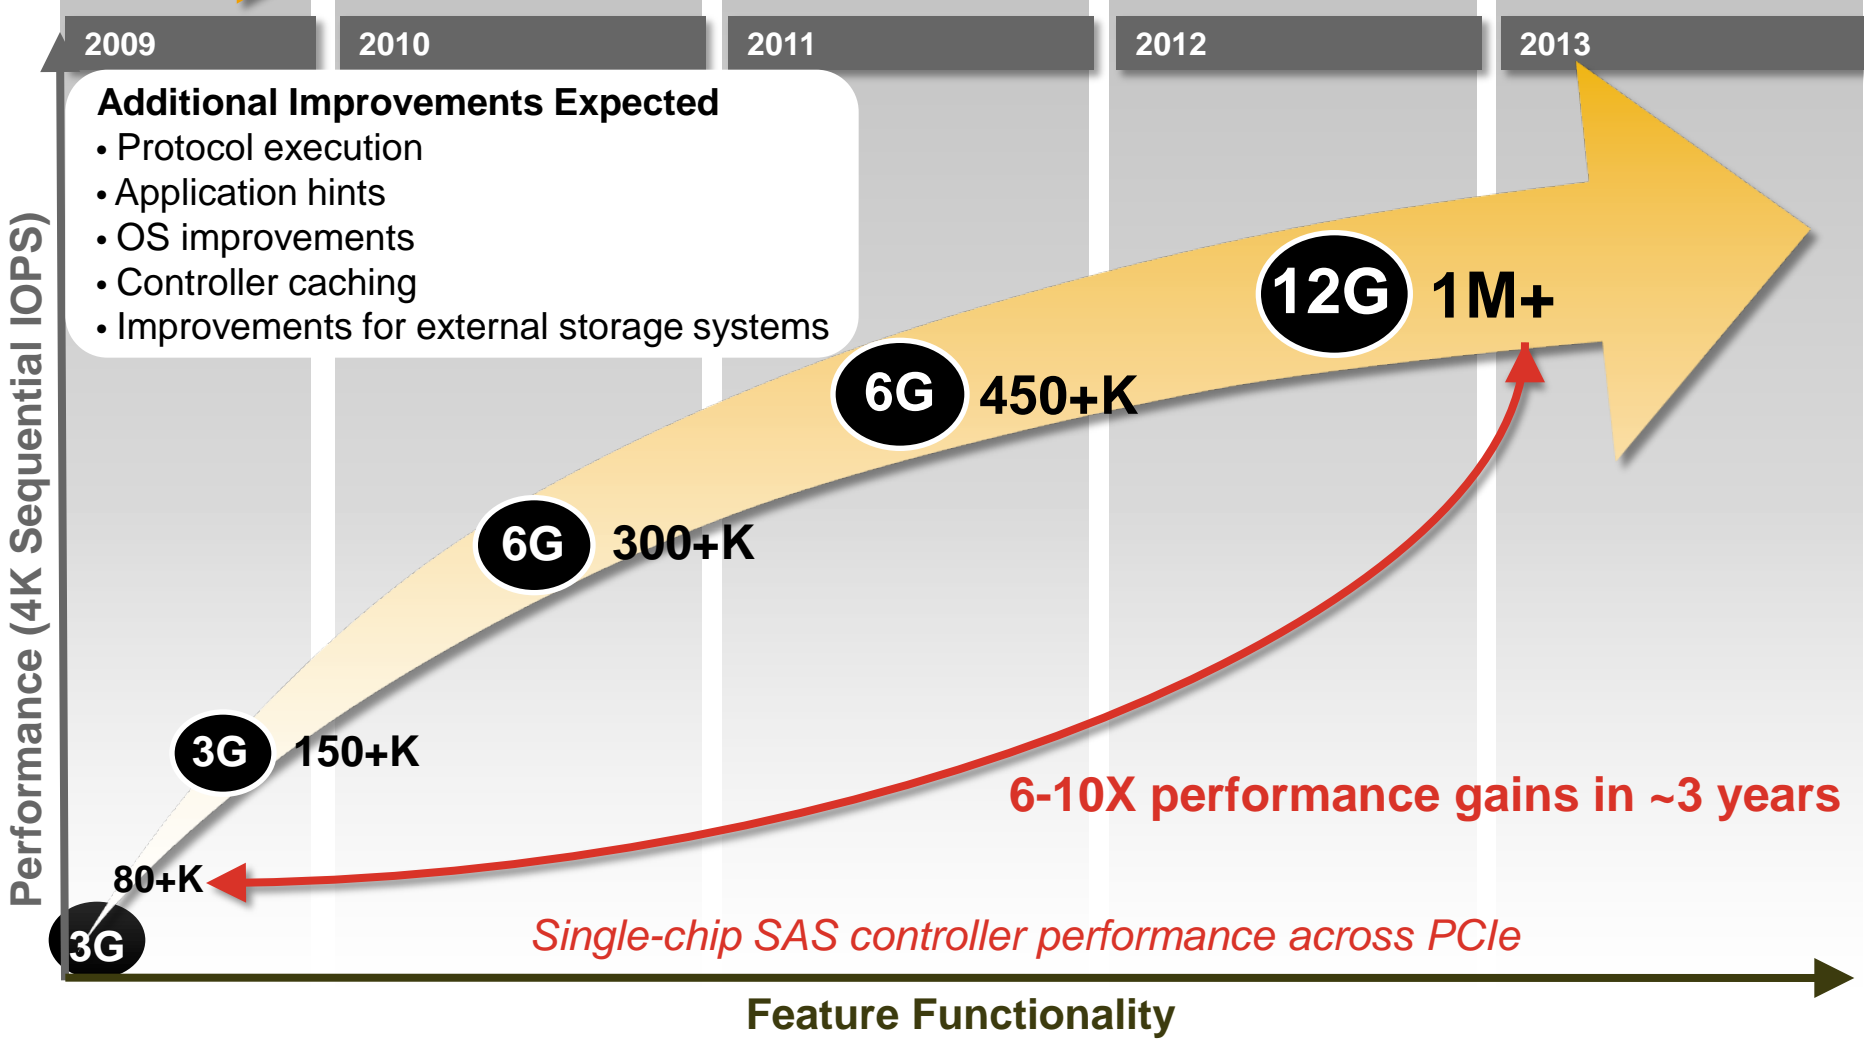
# SAS Performance Roadmap



*\*SAS Roadmap updated Nov 2010.*

# SAS Controller Projections


## - No SAS Changes




# Connector Types



## InfiniBand SAS 4X Connector



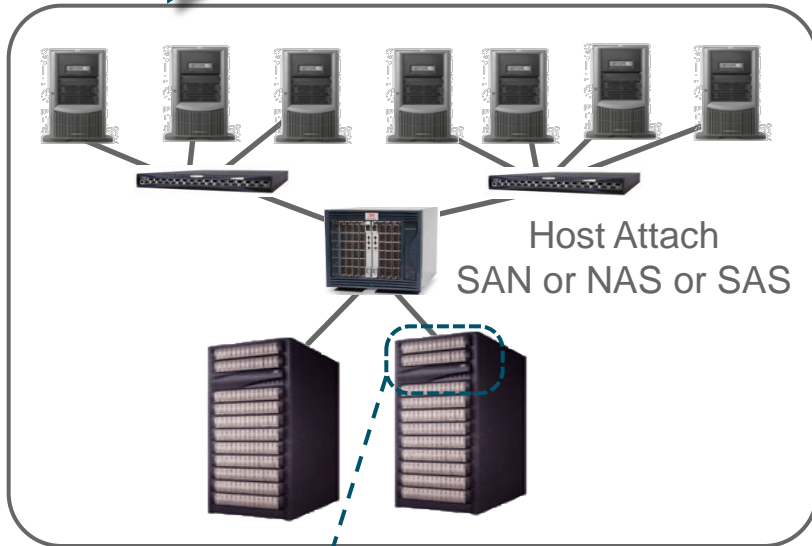
## External Mini-SAS SAS 4X Connector




***Preferred external connection scheme for 6Gb/s SAS***

# Scaling SAS Architecture


Scalability



SAS Switch (packaged view)




SAS Switch



Embedded Controllers  
SAS Connected

Various JBODs


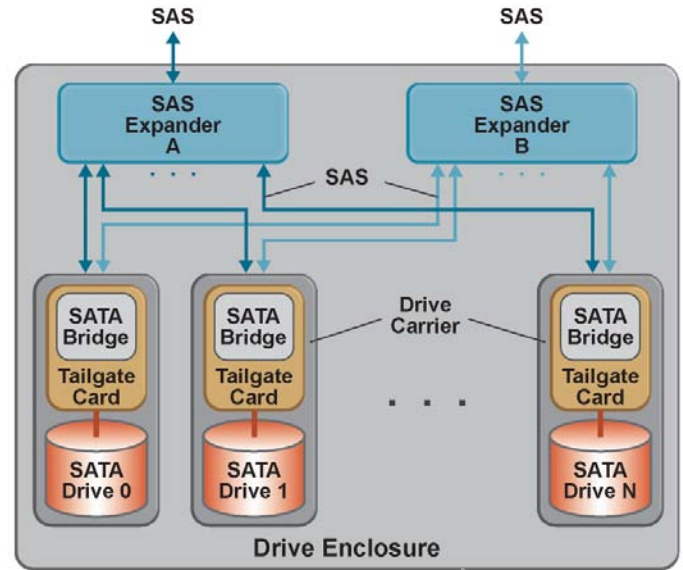


SATA HDDs

SATA HDDs

SATA HDDs

SAS HDDs

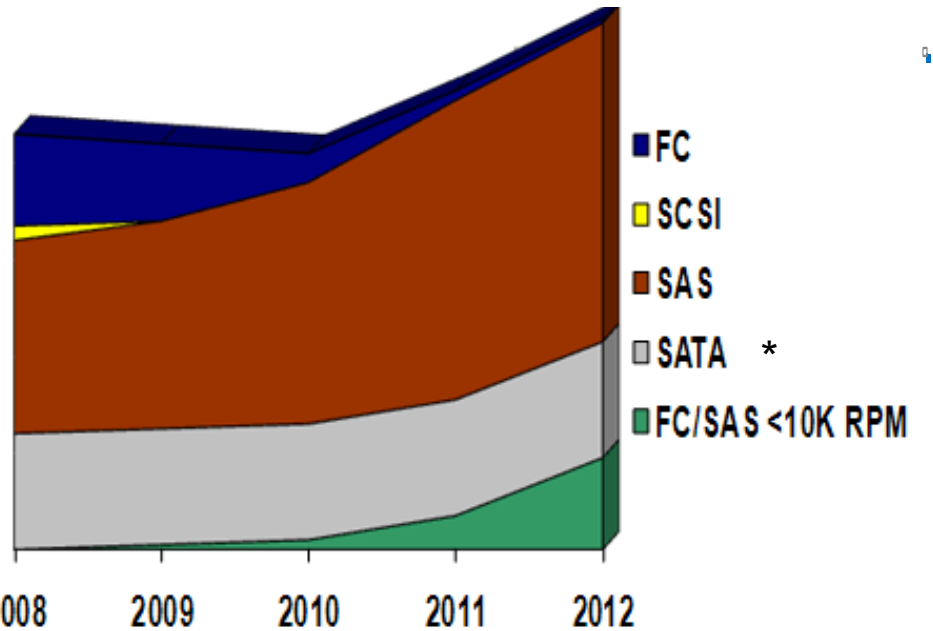


# Key Markets & Sample Deployments

➤ Primary Areas of SAS Usage

# Enterprise Drive Shipments


**Enterprise Applications (2008-2012)**  
 (all form-factors by interface)



\* Enterprise SATA drives often served by SAS Interconnects.

**eSSD Forecast by Interface**

Source: IDC, January 2011



**SAS is the Predominate Enterprise Drive Interface**



# Standard High Volume Servers


Solutions

## SAS in Servers – Features & Benefits:

- Internal storage scalable via expanders
- Robust 6Gb/s SAS performance
- Scalable I/O performance
- High reliability & redundancy
- Point-to-point serial - simplified cabling
- Compatible with SATA & SAS HDDs & SSDs




Rack Mount Servers  
 Tower Servers  
 Blade Servers




SAS: The Emerging Storage Fabric

## Solutions



## SAS External Drive Enclosures


- 1U to 8U Form Factor
- Supports SFF 2.5-inch & 3.5-inch SAS/SATA/SSD Drives
- Higher Capacity with LFF SAS & SATA
  - Higher Port Count/Rack
  - Lower Power/Drive
- Scalable Expander or Low Cost Non-expander base Chassis Solution
- Hot-swappable Drive Carriers, Cooling Fans & Power Supplies
- Single & Dual Expander Options Available



# Dual Path to External Storage

## Solutions

- Dual paths from host to a external storage enclosures as well as dual paths to individual (dual ported) SAS drives



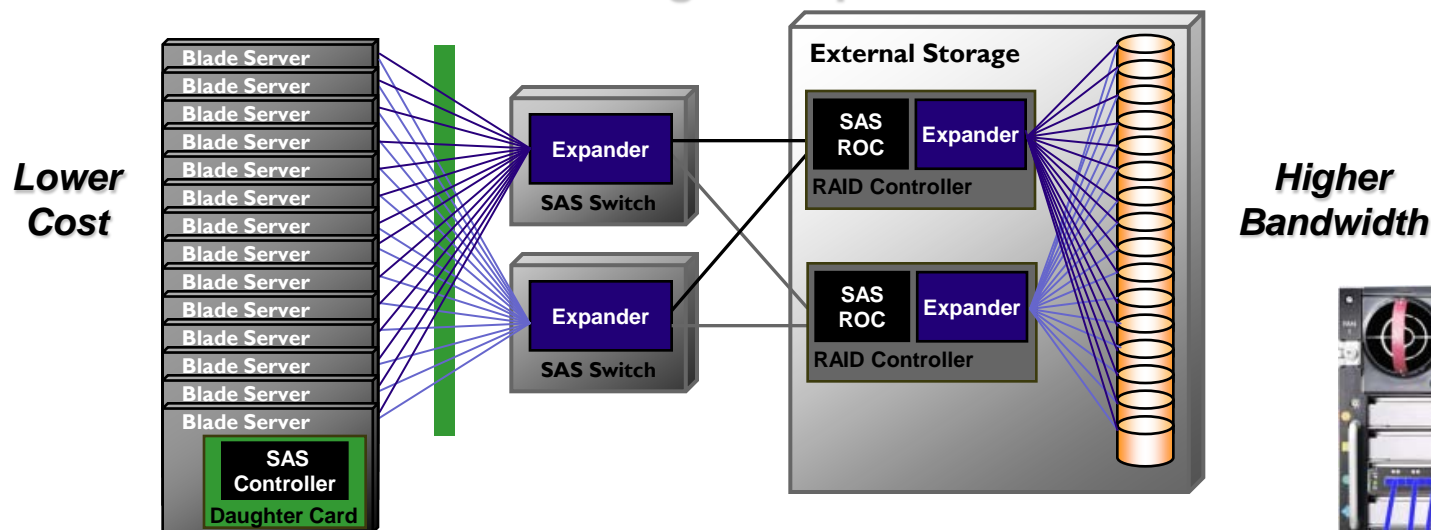
Dual path to each enclosure, and to individual dual ported drives

Dual Domain/Redundant Path: 4 enclosures behind 2 SAS ports on a single server

# External Storage - Blades


Solutions

## Blade Storage Mid-plane I/F



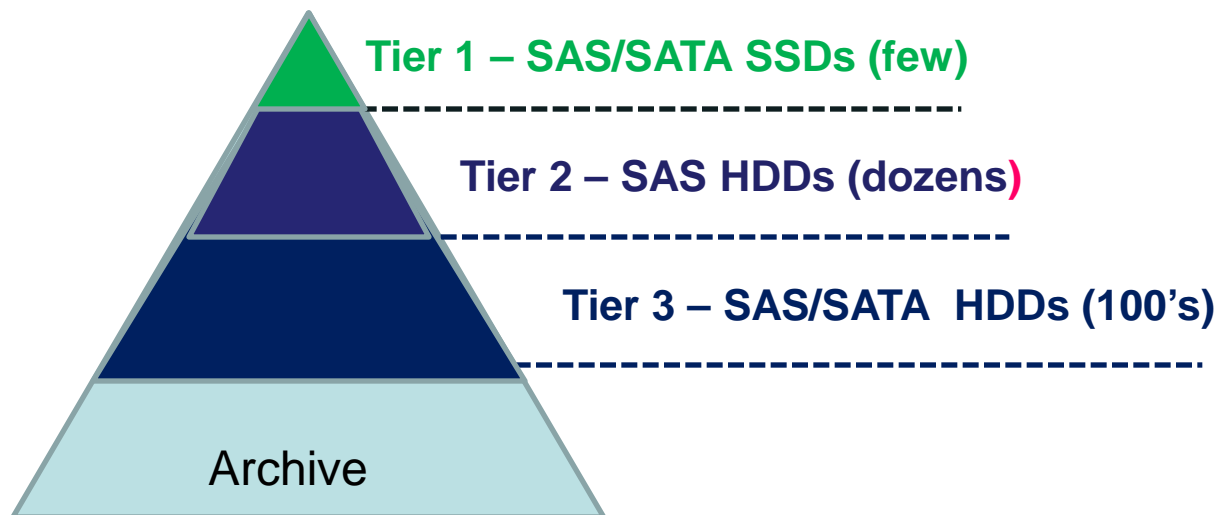
### Switch Benefits and Management:

- Central management
- Multiple servers to one or more storage JBODs
- Efficient scale out
- OS independent
- Direct Attached Storage (DAS) or Shared Storage
- Port based zoning
- Drive bay zoning



# 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



- Exploding capacity exceeding 60EB by 2016
- Content preservation & format conversion - Film, video tape, & other
- Performance demands - Shared databases, non-linear editing, & animation
- Content archiving - Grows to 60% of capacity demand in 2016
- Storage tiered (Tape, HDD, Optical, & Flash)
  - ◆ HDD arrays become the dominant mode for fixed content storage
  - ◆ Active archiving drives HDD growth
  - ◆ Flash memory use in cameras & content distribution
- Content distribution
  - ◆ Network (NAS) – Capacity Scaling
  - ◆ Direct attached storage (DAS) - Performance



# Extending SAS Innovation

➤ 12Gb/s SAS, MultiLink SAS™,  
Advanced Connectivity

# 12Gb/s SAS Industry Timelines

- T&M Equipment & Protocol Analyzers - early availability
- Internal & External Connectors - available or being tooled
- Early Component/HBA Testing - 4Q 2011
- HDDs and SSDs - 1H 2012
- STA Technology Showcase – 05/09/2012
- Industry Plugfest - 3Q 2012
- System Shipments - 2H 2013



## Mini-SAS HD External Connectors



Redesign of the current 6Gb/s connector to run 12Gb/s  
-6Gb/s – SFF-8644 → 12Gbps – SFF-8644

### Parts Availability for Right angle Receptacles

- First Pass Parts for eval – see below
- Low Volume Parts for First Pass Systems – see below
- Volume ramp – 1x1, 1x2, 1x4 Parts Shipping in High Volume

### Parts Availability for Cable Assemblies

- First Pass Plugs for eval – see below
- Low Volume Plugs for First Pass Systems – see below
- Volume ramp
  - The SAS 2.1 versions are shipping for 4x
  - Lower volumes shipping for 8x



#### Note:

- The technical content of SAS 3.0 is not final
- 12Gb/s components are based on currently available information.

# SSDs & MultiLink SAS™


➤ Increasing SAS Bandwidth

- ❖ Dominant contributor to SSD latency – Flash Components
  - ◆ SLC access > 25us, MLC access > 50us, assuming no access contention
- ❖ Lager 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 (fractions of a microsecond)**

## ➤ 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



SAS Connector  
(SFF-8482 & SFF-8680)



MultiLink SAS Connector  
(SFF-8630)



SAS: The Emerging Storage Fabric

© 2012 Storage Networking Industry Association. All Rights Reserved.

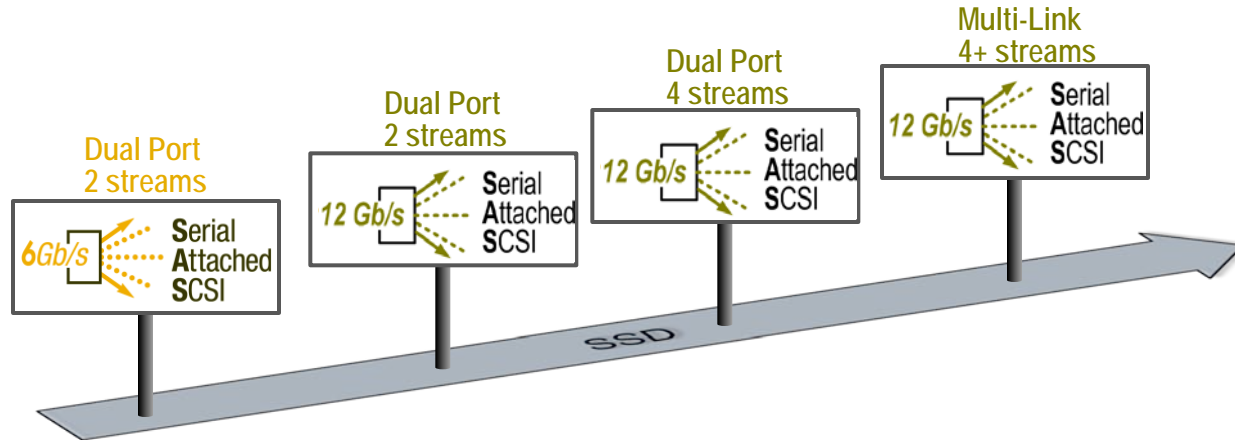
# Storage Interface Comparison

	<b>SATA</b>	<b>SAS</b>		<b>PCIe</b>		
	SATA	SAS	Multilink SAS	SOP/PQI	NVM Express	Proprietary
Drive Form Factors	1.8", 2.5", 3.5"	2.5", 3.5"	2.5"	2.5"	2.5", Card	Card
No of Ports / Lanes	1	1, 2	1, 2, 4	1, 2, 4	1, 2, 4 (8 on card)	1, 2, 4, 8
Command Set / Queuing Interface	ATA / SATA-IO	SCSI / SAS	SCSI / SAS	SCSI / SOP / PQI	NVM Express	Vendor-Specific
Transfer Rate	6Gb/s	12Gb/s	12Gb/s	8 Gb/s	8 Gb/s	8Gb/s
Drive Connector	SATA-IO	SFF-8680	SFF-9630 SFF-8639	SFF-8639	SFF-8639 (2.5"), CEM (Edge-Card)	CEM (Edge-Card)
Express Bay Compatible?	Yes (2.5")	Yes (2.5")	Yes (2.5")	Yes (2.5")	Yes (2.5")	N/A
Drive Power	9W Typical (2.5")	9W Typical (2.5")	Up to 25W	Up to 25W	Up to 25W	Vendor Specific
Max Bandwidth	0.6 GB/s	4.8 GB/s (x2)	9.6 GB/s (x4)	8 GB/s (x4)	8 GB/s (x4)	16GB/s (x8)

## System and Use Case Considerations

Host Driver Stack (Storage Controller / Direct Drives)	AHCI	IHV	IHV	Common Driver Possible (SOP/PQI)	Common Driver Possible (NVM Express)	IHV
Surprise Removal / Insertion ('Hot Plug')	Yes	Yes	Yes	Future Enhancement	Future Enhancement	Vendor Specific

# Innovating for Bandwidth Growth



**Today:** Leverage existing 6Gb SAS infrastructure for rapid deployment

**Tomorrow:** Enhance the SAS value proposition to accommodate SSD unique requirements:

### SAS Roadmap Progression:

- A) 6Gb SAS dual port, two streams
- B) 12Gb SAS dual port, two streams
- C) 12Gb SAS dual port, four streams
- D) 12Gb SAS multi-link, four+ streams

**SSD unique commands:** Trim, Unmap, others...

Benefits of SAS & proposed Multi-link enhancements SAS	
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>


# Advanced Connectivity

➤ Taking SAS deeper into the  
Enterprise




# SAS Advanced Connectivity Objectives


- Drive market consistency
- Simplify cable & connector options
- 2X density improvement
- Provide converged high-density connectivity
- Provide managed connectivity standards
- Provide active copper solution to 20m
- Provide optical solution to 100m
- Support 6Gb/s SAS deployments
- Extensible to 12Gb/s SAS deployments



**Internal similar to External**

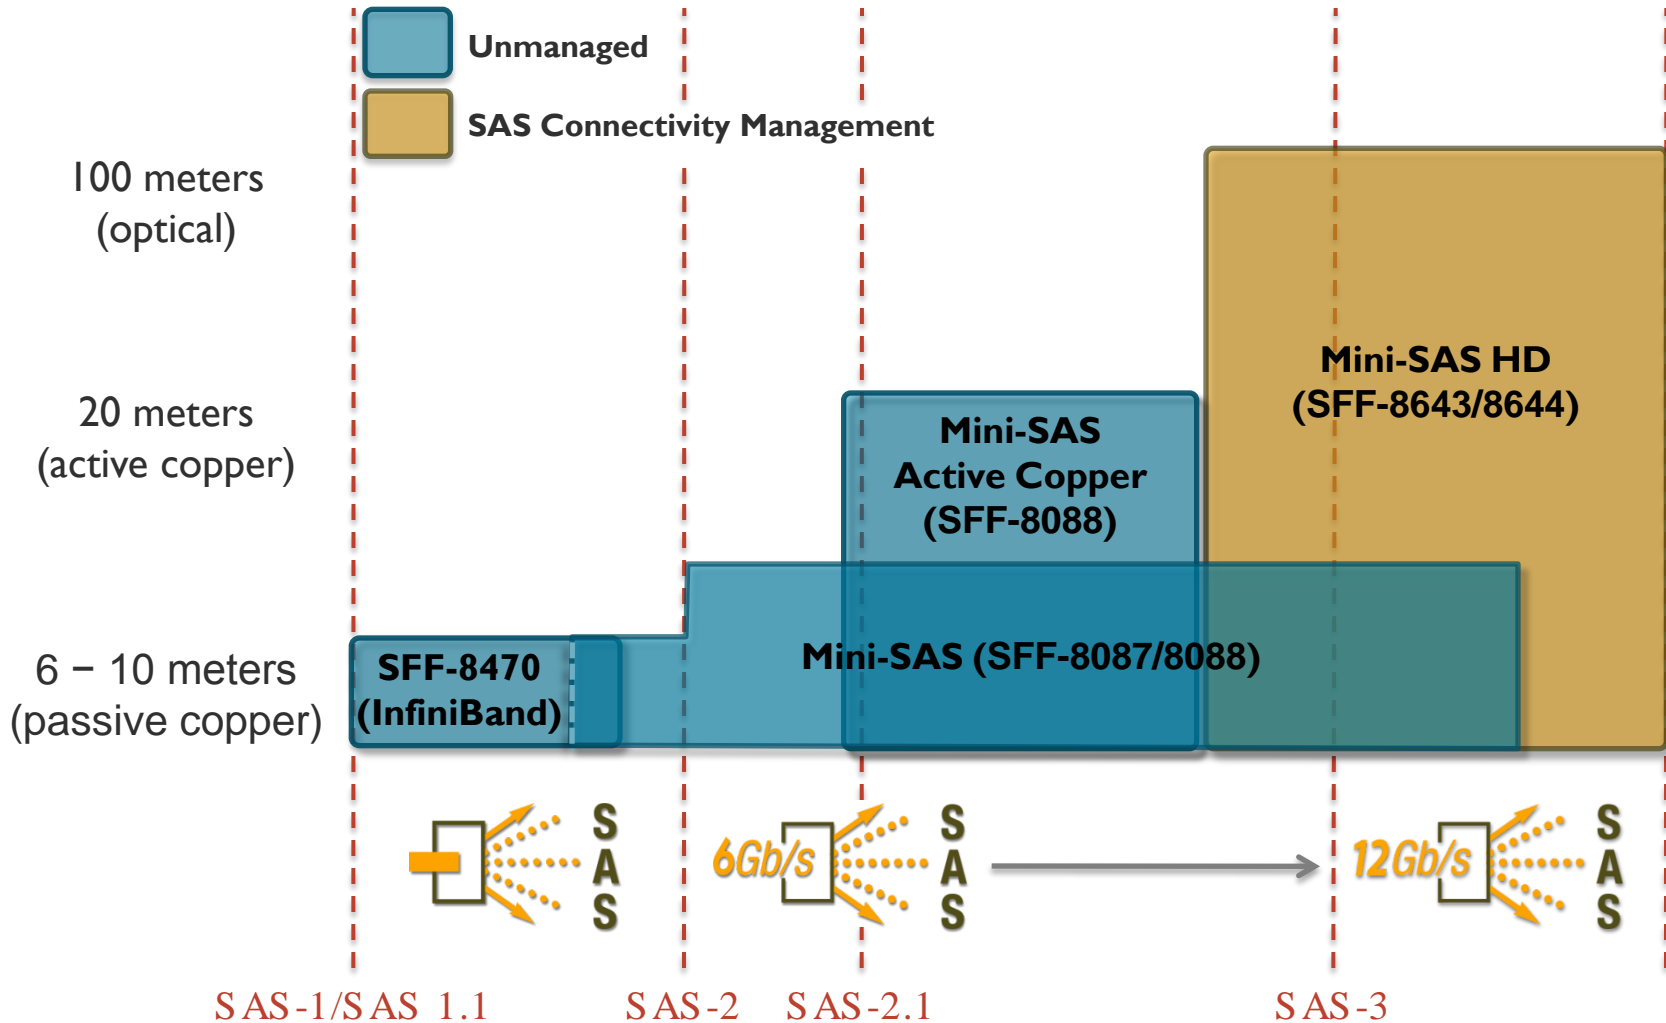


**Passive, Active Copper, or Optical use same connector**




Mini-SAS HD connectors courtesy:  
Project T10/2125-D Revision 04  
17 September 2009, ANSI SAS-2.1

# SAS Advanced Connectivity Roadmap – 2011




# SAS Price Advantage

**Based on switched connections to eight servers**




Configuration Cost

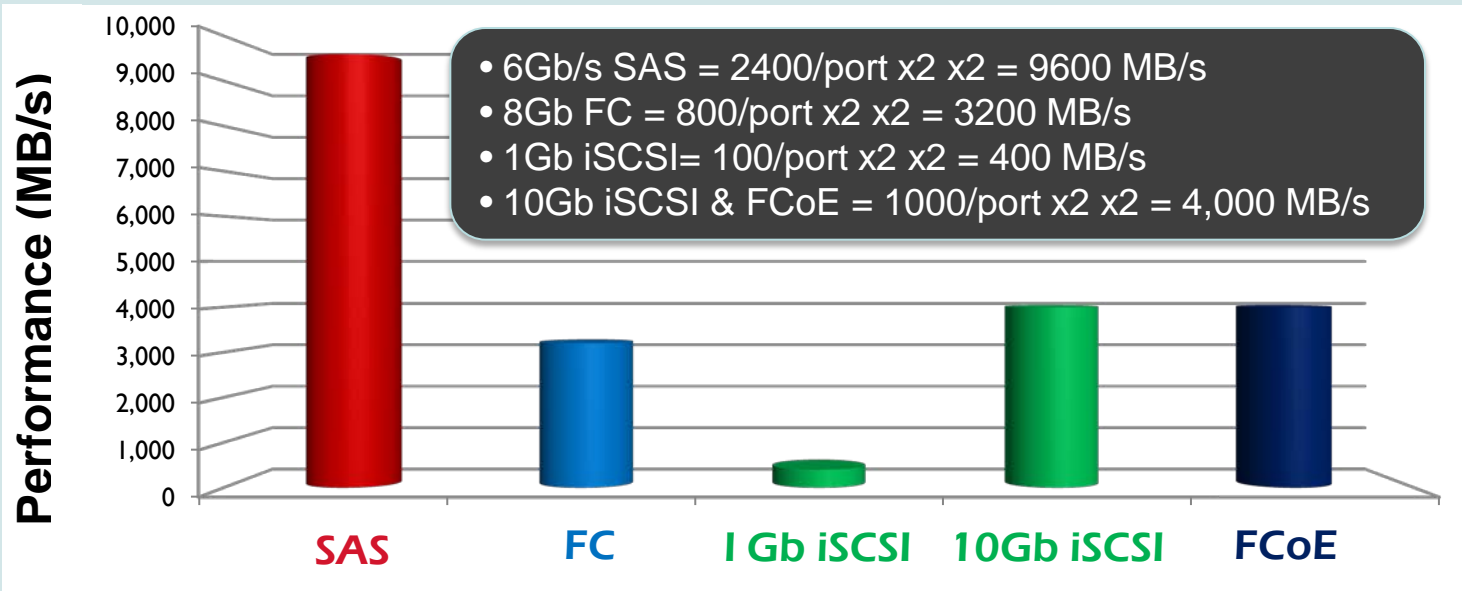


# SAS Performance Edge

*Based on switched connections to eight servers*



## Raw Performance




\*Performance is theoretical performance based on two external storage systems, dual controllers, with one interface host port per controller

SAS: The Emerging Storage Fabric

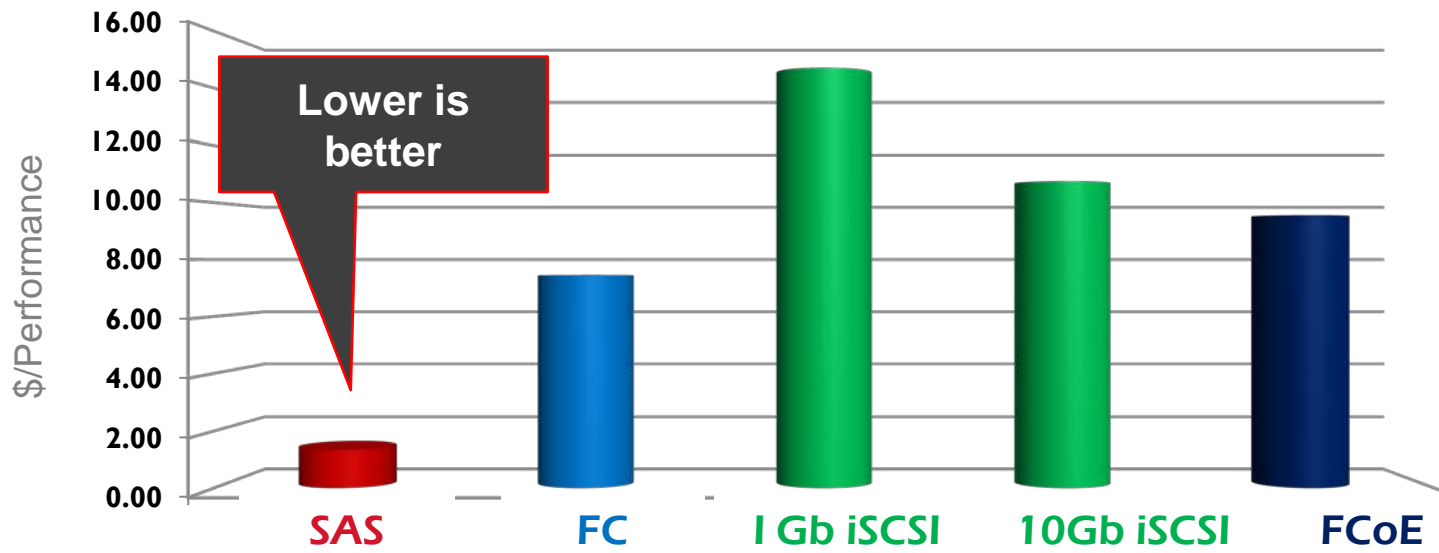
© 2012 Storage Networking Industry Association. All Rights Reserved.

# Price-to-Performance Comparison

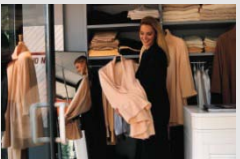
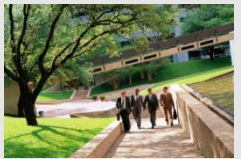

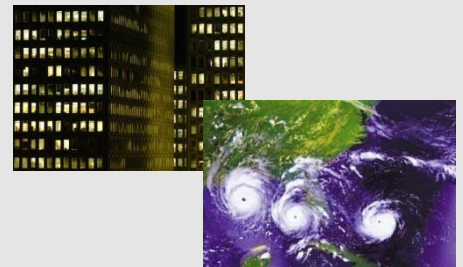
**Based on switched connections to eight servers**



Price/Performance Ratio by Interface

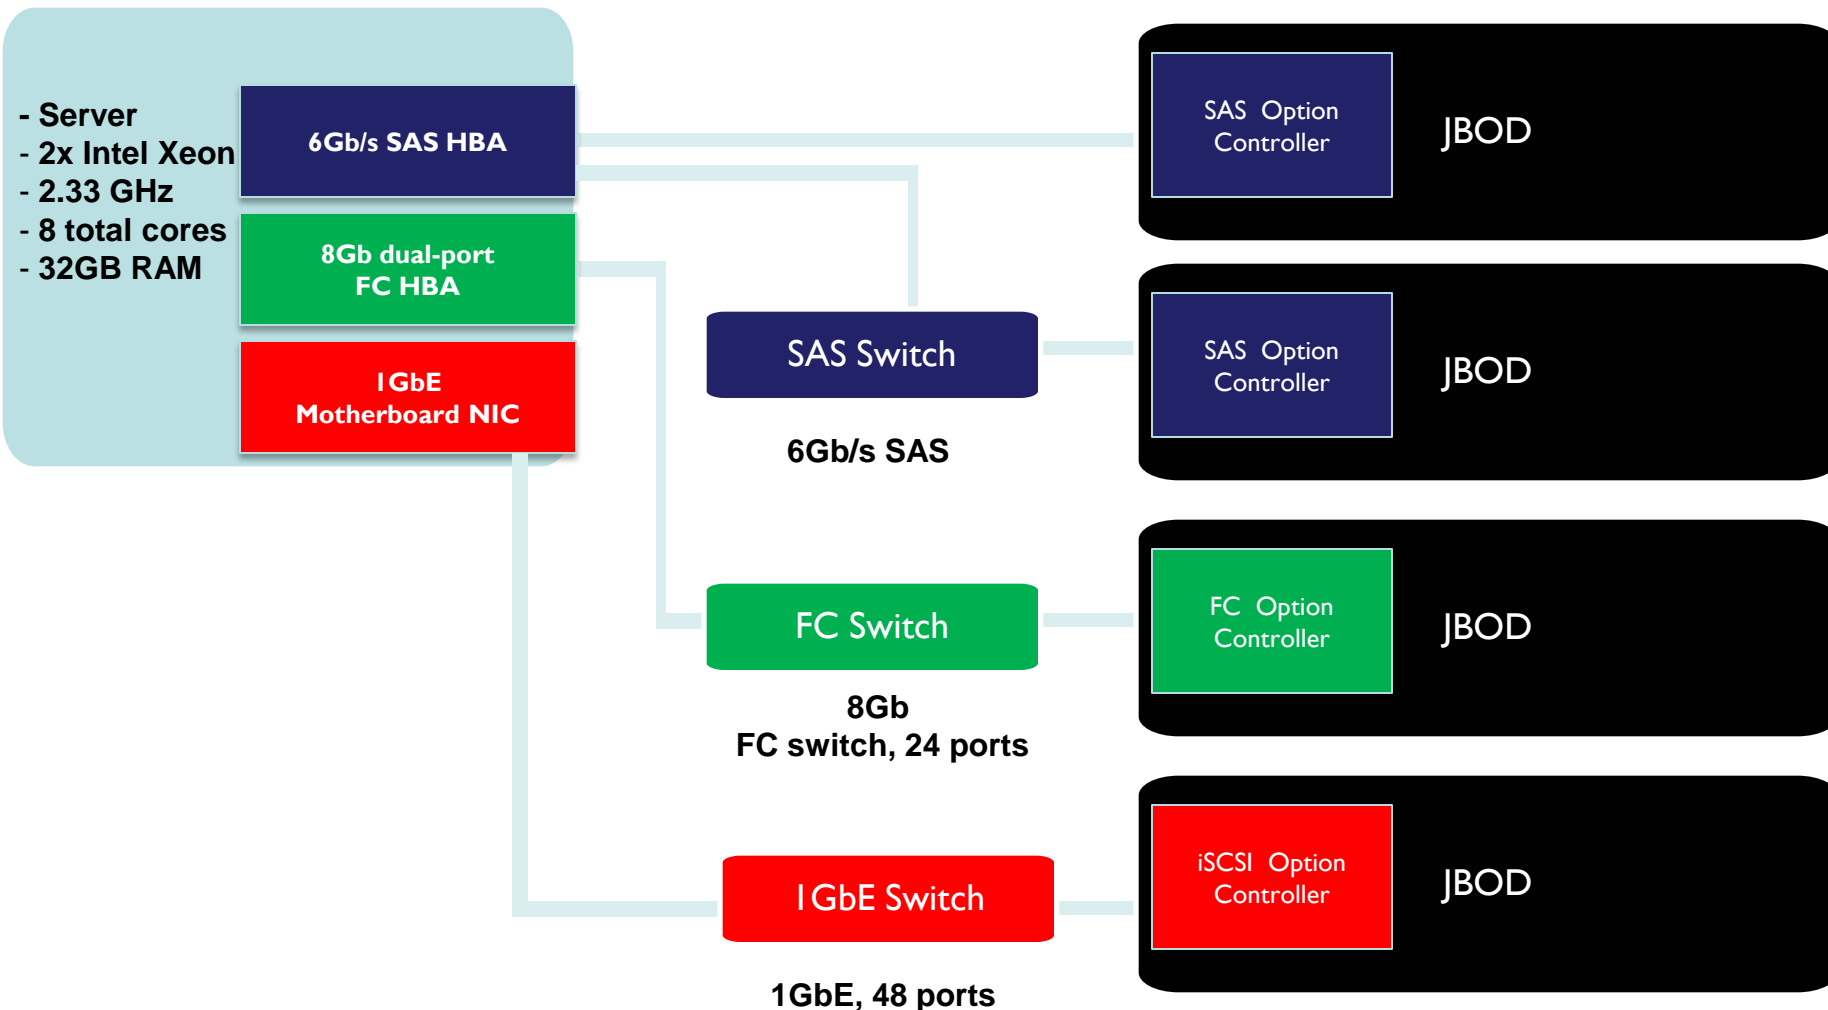


# Connection Distances Comparison

Meters	10	100	1000	10,000	Unlimited
					
	<b>SMB Remote site/ Campus</b>		<b>Large Campus Data Centers</b>	<b>Enterprise Multi-floor infrastructure Many-to-many</b>	<b>Global Enterprise Disaster Recovery</b>
<b>SAS</b>	Passive	Active	Casadable links have been demonstrated several hundred meters		
<b>FC</b>	[Blue bar representing distance range]				
<b>FCoE*</b>	Not IP Routable				
<b>iSCSI*</b>	IP Routable				


- As distance increases, performance can be negatively impacted
- This slide represents a generalization, in certain cases distances can be extended with multiple switches, tunnels, buffering/repeaters, etc.

# SQLIO Switched I/O Example



# SQLIO Performance Comparison


## SQLIO - IOPS





# SQLIO – Bandwidth Comparison

## SQLIO - MBPS



- Logical SCSI lives across the 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

- Please send any questions or comments on this presentation to SNIA: [tracktutorials@snia.org](mailto:tracktutorials@snia.org)

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

**- SNIA Education Committee**

**Chris Lionetti  
Nancy Clay**

**STA Members  
Dennis Martin**