



# Improving SAS for future generations

**LSI Presentation to the SCSI Trade Association**

**Timothy Hoglund  
LSI Fellow  
14-January-2008**

# Proposed Goals for SAS-3

- Forward and Backward Interoperability
  - Link Rate
    - 3 Physical Rate Generations: SAS-3 (12.0Gb), SAS-2 (6.0Gb), SAS-1.1 (3.0Gb)
    - 4 Connection Rate Generations (using rate matching for 1.5Gb)
  - Electrical
    - Backward
      - SAS-3 phys must function on SAS-2 interconnect
    - Forward
      - 6Gb SAS-2 phys must function on SAS-3 interconnect
  - Compatibility with future SATA specifications
- Performance
  - 2x bandwidth (1200MB/s)
  - Improve efficiency of the topology
  - Adopt enhanced method for bandwidth aggregation
- Improved cable reach
  - Active copper
  - Optical
- Power
  - Interface and Device power management mechanisms
- Robust
  - BER target  $10^{-15}$ 
    - Based on S-parameter based simulations, as per SAS-2

# Physical and Phy Layers

- Link Rate Alternatives for 1200MB/s bandwidth
  - 10Gb/s
    - Follow PCIe lead to double effective link rate
      - vs doubling of physical link rate
    - Encoding via 64b/66b or 64b/67b or scrambling with DC restoration
      - Longer run length increases ISI
      - Shift to different encoding scheme is 1 time fix only
    - Increased complexity
      - Have to replicate functionality of K characters
      - How to insure DC balance?
      - Clocking complications for non integrally related phy rates
      - Challenges for rate matching and multiplexing
  - 12Gb/s
    - Simpler compatibility model
    - 8b10b
    - Will require more advanced equalization techniques
  - LSI evaluating position
    - Starting to simulate channels at 12Gb/s

## LSI Recommendation

- 12Gbps with 8b10b encoding


## LSI 12Gb/s “Preview”

- Preliminary simulations against SAS-2 Channel Models
  - OIF-CEI-11G-LR compatible architecture at 12Gb
    - 3 TAP TX FIR
    - 8 TAP RX DFE
  - Voltage and timing noise sources
    - Inter-symbol interference (ISI)
    - Circuit noise (modeled as AWGN)
    - Random Jitter (RJ)
    - Duty-cycle distortion (DCD)
    - Crosstalk etc.
- Looked at horizontal and vertical eye margins for target BER of  $10^{-15}$
- Preliminary observations – “Not slam dunk, but shows promise”
  - Good performance achievable over most SAS-2 channels at 12Gb/s with significant horizontal and vertical margin
  - Crosstalk in some cases (HP19, HP20, HP21) becomes prohibitively large resulting in eye closure or reduced margin
  - Large insertion loss at Nyquist observed on HP24-26 channels makes them challenging to handle

## Desired cable/connector improvements for SAS

- Cheaper passive cables
- Longer cable length
  - Active cable
    - 20m minimum
    - Consider 50m-100m
- Cable recognition
  - I2C and serial EEPROM in cable
- Optical support

# SFF-8470, SFF-8088 & QSFP



## Concerns with QSFP

- Too long for a 2U short PCI card (MD2)
- InfiniBand-style pin ordering

# Link Layer

- Link Aggregation methods

- SAS multiplexing
- Store/forward or switched
- Virtual connections?
  - Priority
  - QoS

## SAS Multiplexing

- Expensive in controllers and expanders
- Does not play well with mixed speed devices
- Inflexible

## SAS Expander Buffering

- Greatest opportunity to maximize:
  - Topology efficiency and utilization
  - System performance
- Accommodate mixed speed devices, flexibly

## SAS has Virtual Phys

- Why not Virtual Connections?



- Topology efficiency and scaling
  - Arbitration/Fairness improvements?
    - Congestion mitigation
  - Preferential pathway allocation?

# Transport Layer

- Security/Authentication
  - What?
    - Authentication
      - Insure that the claimed sender is the actual sender
    - Integrity
      - Insure that no data has been created/modified/deleted between two devices
    - Confidentiality
      - Insure only intended recipients know what is being sent
  - How?
    - Don't re-invent, rather leverage from IPsec/FC-SP/SCSI-IKEv2
  - Why?
    - Data security threat models for storage will evolve beyond data-at-rest
    - SAS usage models likely to include “small fabric”
    - Devices within SAS domain can easily incorporate “snoop” port capability
    - Peer and competing interfaces will have solutions

## LSI Recommendation

- Pursue “SAS-SP” solution, perhaps address with separate SAS workgroup



# Application Layer

- Improved management protocol
  - SMP originally defined for simple expanders
  - Scope of SAS has grown
    - Includes more complex configuration and reporting
    - Amount of management data has grown significantly
  - Likely continued increase in both complexity and data set size
    - Authentication and Security

→ **Single frame, single connection SMP not best fit**

## LSI Recommendation

- Pursue standardization of either:
  - Enhanced SMP (disconnecting and multi-frame)
  - Encapsulation
    - SSP for in-band
    - ??? for out-of-band

## Moving Ahead

- Close on SAS-2 ASAP (Q108)
- Constrain Scope of SAS-2.X to improve SAS-2 connectivity
- Resource SAS-3 once SAS-2.X is forwarded, if not before (Q408)
  - Many potential enhancements
  - Additional complexity will require a significant effort to close
- Strongly recommend NOT breaking SAS into multiple specifications
  - Slows approval cycles
  - Encourages feature creep
- STA to provide SAS-3 Marketing requirements Q2-Q3 2008.

# Thank You

