SNIA

Education

Solid State Storage Architectures

Jamon Bowen, Texas Memory Systems

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- 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

Solid State Storage Architectures

- The field of SSS has expanded into multiple classes of offerings (RAM, Enterprise Flash, and Consumer Flash) and ways to deploy the technology (form factor SSDs, PCIe, and rack mount). This presentation will quickly cover the features, relative costs, and performance of the different architectures as well as examples of the best fit applications. A current view of how SSS compares to HDD systems will be presented covering performance, price, and operational costs. Different methodologies of deploying SSS along side HDD storage will also be presented.

Overview

- Lots of SSS options available and lots of hype.
- There are a few fundamental differences between SSS deployment options.
- Determine the application architecture and requirements for deployment before evaluating vendors' offerings to save time.
- SSS is the future due to operational and performance cost savings.

Types of SSS

- RAM-based
  - Lowest latency media
  - Requires built-in batteries and non-volatile back-up storage
  - More expensive per GB
- Flash-based
  - Non-volatile chip, high density, low power
  - Chips require special handling for writes:
    - Wear leveling
    - Bad block replacement
  - Faster than disks, but at least 10x slower than RAM
  - Generally needs a lot of chips to get a lot of performance
  - Capable of disk density
  - Low power
  - MLC for mainly consumer apps

Types of Deployment Options

- Storage Centric
  - Typically form factor SSD
  - Designed to fit the requirements of storage system vendors.
- Server Centric
  - Designed to fit into the server
    - PCIe
    - Embedded on the motherboard
- Network Centric
  - Designed to attach to servers via a network
    - Appliance model

RAM SSS Attachment Options

- In the storage industry, RAM SSS deployed in the three attachment methods in the following ways:
- Storage Centric
  - SAN cache
- Server Centric
  - File system cache
  - RAM drive
- Network Centric
  - RAM SSS appliance

RAM SSS Appliance Applications

- A RAM SSS appliance provides the ability to add storage persistence to an application without sacrificing RAM performance and accommodating server clustering for availability.
- Typically involves clustered servers and mirrored RAM SSS
  - Database Logs
  - Financial applications
  - Real-time billing systems
- Typically small capacities

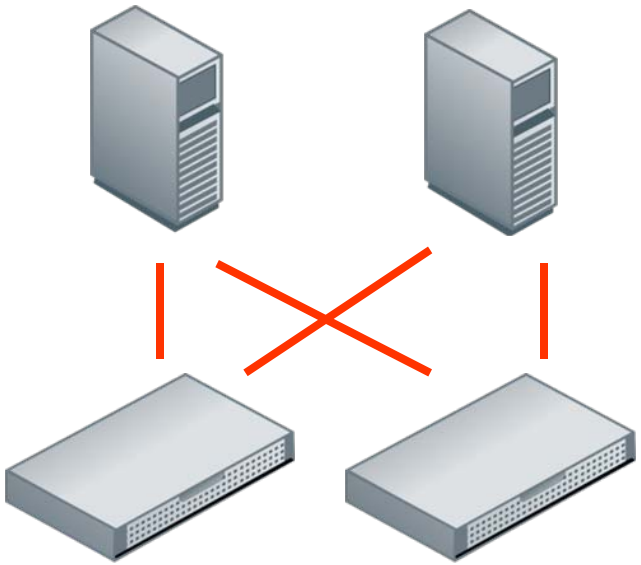

FLASH

- Dominates the discussion of SSS as it is a rapidly expanding market
- Factors driving growth:
  - Price per GB has been declining faster than other options (Disk or RAM)
  - Price per performance is extremely inexpensive
  - Dramatic space and power savings for persistent storage.
- Explosion in deployment options

Deployment Options Evaluation

- Typical vectors:
  - Performance
    - Latency
    - Bandwidth
    - IOPS
  - Cost
    - Minimum total cost
    - Cost per capacity
    - Cost per performance
    - Operational cost considerations
    - Management cost considerations
  - Features
    - *Determines the best fit Architectures*

Flash SSS Performance

- The media performance is a constant
- The Flash controller performance is highly variable from vendor to vendor, but isn't tied to any of the specific deployment options directly
- Are there any inherent performance differences?
  - YES! From a solution perspective
  - They can often be overstated

Performance

- Server Centric
  - Latency:
    - Bus Overhead Only
  - IOPS:
    - Single bus performance
  - Bandwidth:
    - Single bus performance

Performance (Cont.)

- Network Centric
  - Latency
    - Bus overhead + Bus to Network Adaptor Overhead + Network Adaptor to SSS overhead
  - IOPS
    - No inherent limitation, limit is solution dependent - maximum ability to aggregate connections.
  - Bandwidth
    - No inherent limitation, limit is solution dependent - maximum ability to aggregate connections.

Performance (Cont.)

- Storage Centric
  - Latency
    - Bus overhead + Bus to Network Adaptor + Network Adaptor to Storage Controller Bus + Storage Controller Processing Stack + Bus to back end storage Network Adaptor + Network Adaptor to internal SSD Bus
  - IOPS
    - No inherent limitation, limit is solution dependent - maximum ability to aggregate connections.
  - Bandwidth
    - No inherent limitation, limit is solution dependent - maximum ability to aggregate connections
- The performance is limited to whatever the storage controller can deliver
  - *Highly vendor specific*

Flash Performance Comparison

- Measuring the Latency impact of typical HBA
- Single Test System with both internal and external SSS

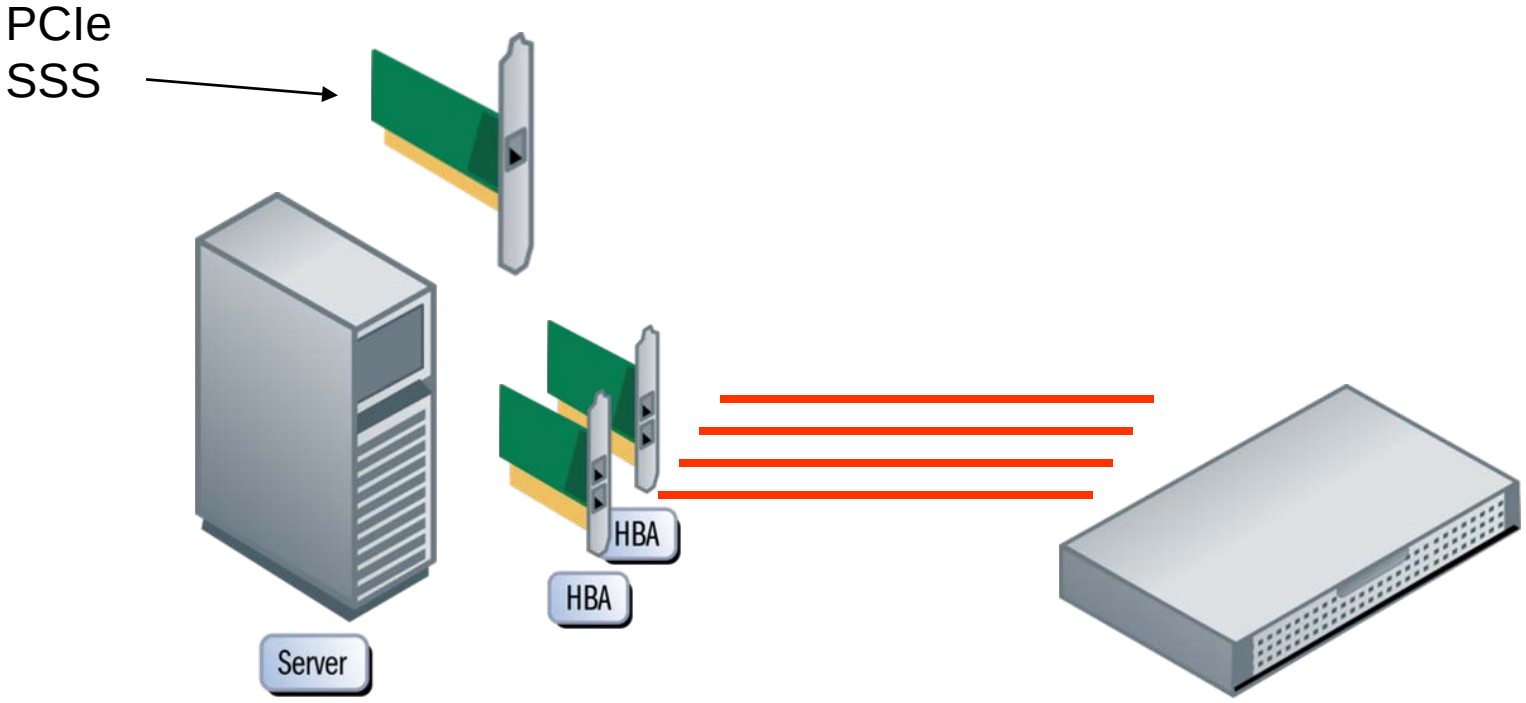

Response Time Comparison

- 4 KB Write Performance VS IOPS Load

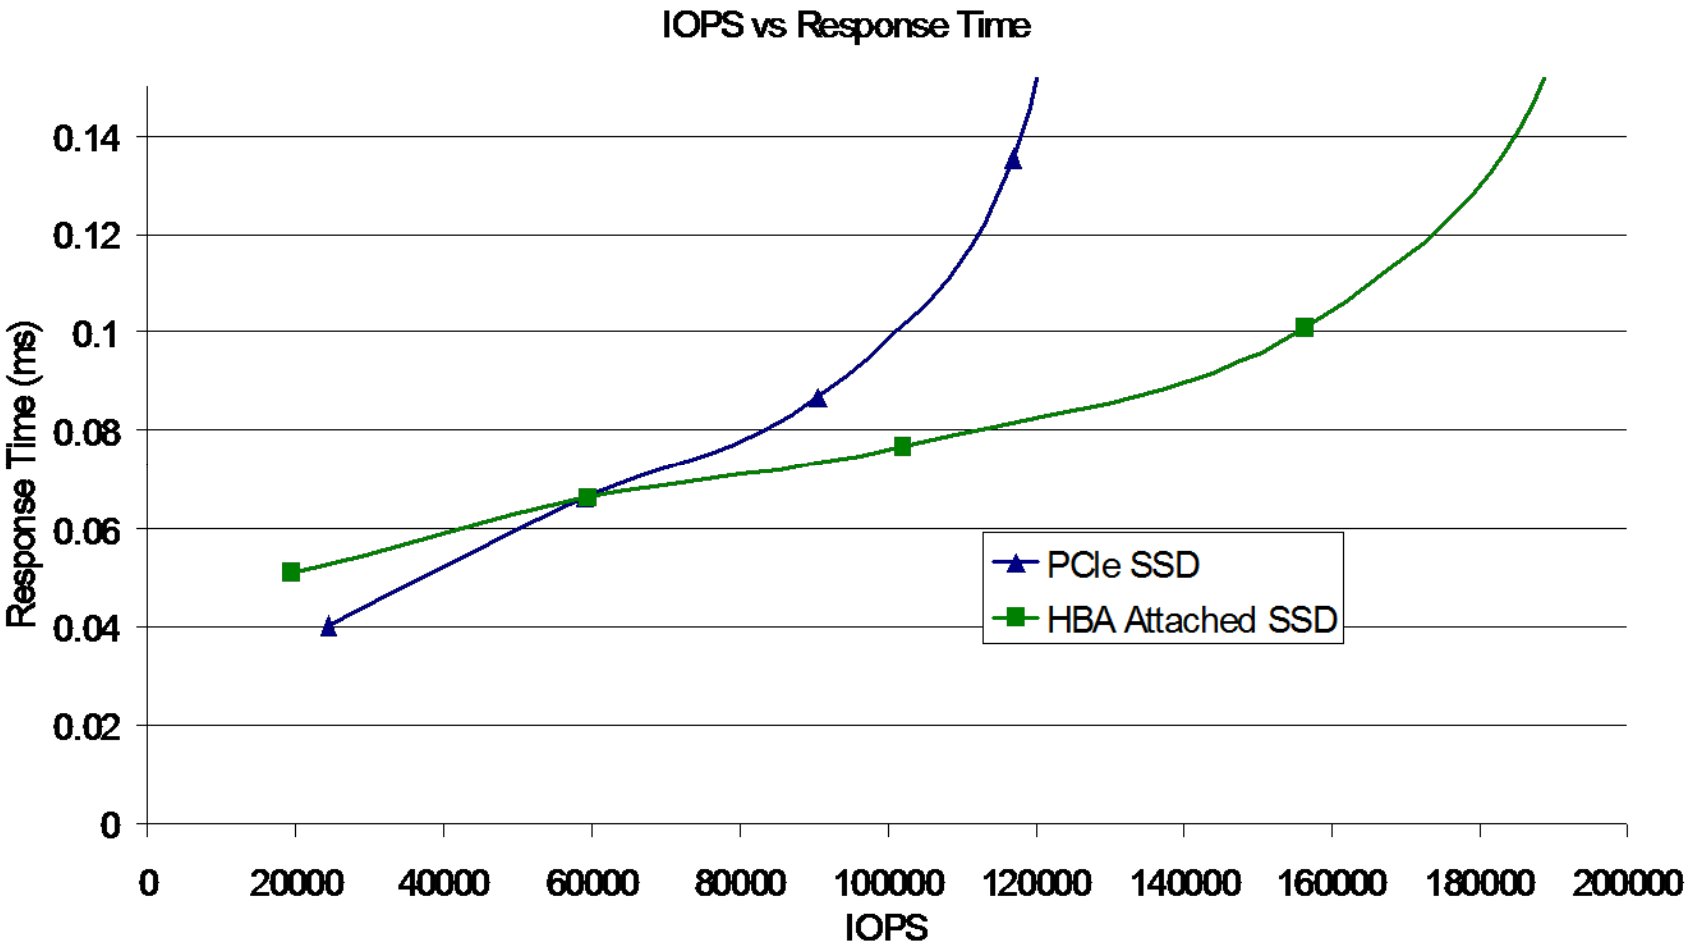

Cost

- Server Centric
  - Lowest minimum cost
  - Cost per capacity determined by chip used & vendor
    - Biggest drawback is fixed capacity vs capacity used
  - Cost per performance depends on controller design and number of chips
  - Operational cost consideration
    - Lowest operational costs
    - Highest management costs

Cost

- Network Centric
  - Minimum cost - tens of thousands range
  - Cost per capacity determined by chip used & vendor
  - Cost per performance: type of chip, number of chips, controller (maybe add spare area size).
  - Operational cost consideration
    - Low operational costs
    - Some management centralization

Cost

- Storage Centric
  - High minimum cost – dominated by storage controller cost.
  - Cost per capacity typically higher to cover OEM markup
  - Cost per performance highest
  - Operational cost consideration
    - Better operational cost than disk
    - Lowest management costs.

Features

- Server Centric
  - Minimum
  - Driver Installation requirement
- Network Centric
  - Shareable, server cluster support
  - Central management and monitoring.
- Storage Centric
  - Leverages Storage Controller Feature set
  - ***Primary benefit of this deployment method***

Architectural Deployment Models

- Scale UP Architecture
  - Monolithic design of servers and storage
- Scale OUT Architecture
  - Distributed design of servers and storage
- Mixed
  - Scaled out servers with scaled up storage
  - Scaled up servers with scaled out storage

Scale UP Architecture

- Monolithic Architecture
  - Single application image
  - Well suited to applications that require tightly correlated datasets and a shared state.
- External system design if:
  - High Availably Requirements
    - Active/active clustering - external system
  - Capacity
    - External system can scale to much higher capacities.
- Internal solution
  - Fits this model when a 3-4u server or workstation is used to provide a powerful single server solution.

Scale OUT Architecture

- Distributed modular architecture.
- Well suited to applications with easily partition-able, loosely correlated datasets.
- Internal
  - Deployed with single card per node.
  - Shared nothing cluster for availability.
- External
  - Many pairs of HA server pairs are scaled horizontally.
  - Mirrored external system per node.

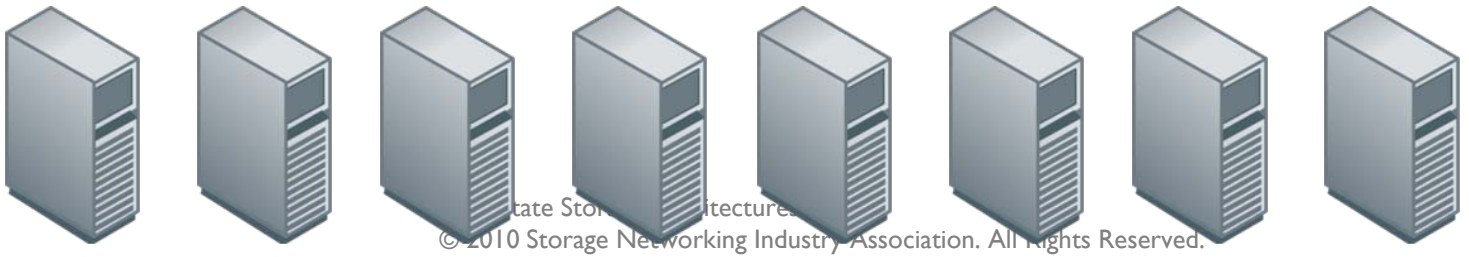

Scaled UP Storage

- Compute resources scale out
  - HPC setup
  - Many node Oracle RAC
- Storage scales up
  - Single shared namespace

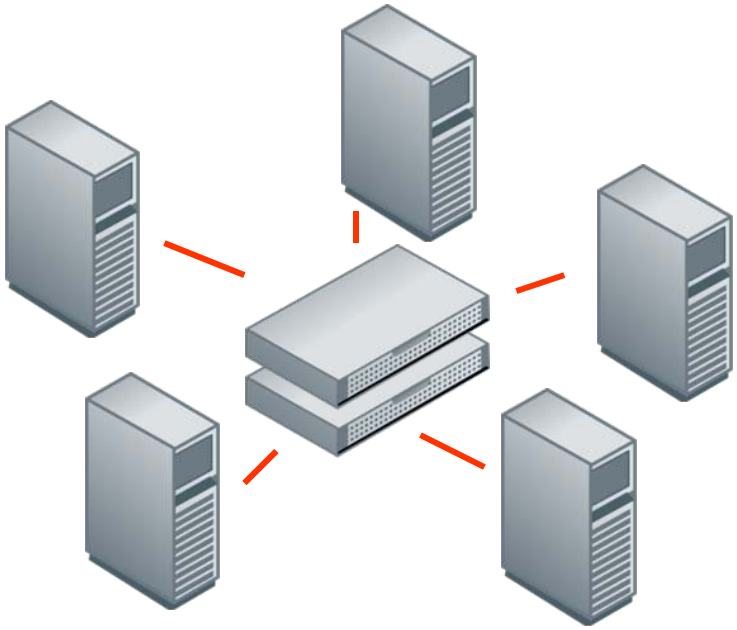

Scaled UP Storage

- Many different applications with different requirements
- Centralized data management primary benefit
  - Performance tradeoff
  - Some array features become more valuable the more applications in the same management framework (i.e. Dedupe)
  - Obvious SSS feature: Tiering

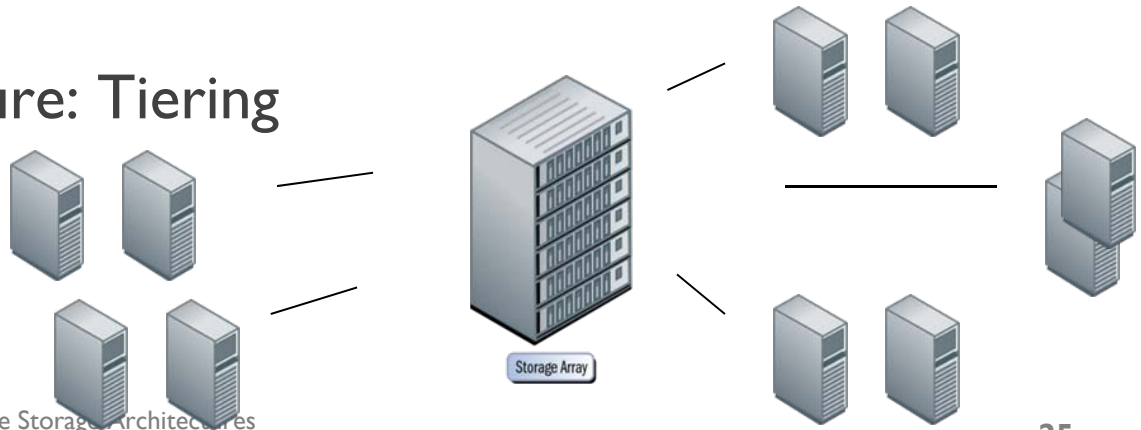

Why is SSS the future?

- The reduced response time of SSS has been the historical reason for deployment
- Lower response time removes time that an application spends waiting on disk storage, accelerating the application.
- What if the latency of an SSS was the same as for an HDD?
  - Applications wouldn't run faster.
  - ***SSS would still be the future.***

HDD's Fundamental Flaw

- You have to buy the physical disk operating components for each disk that you buy (*Big difference compared to tape*)
- Multiple disks needed for random IO performance.

Platter cost per bit decreases rapidly every year. Single manufactured component.

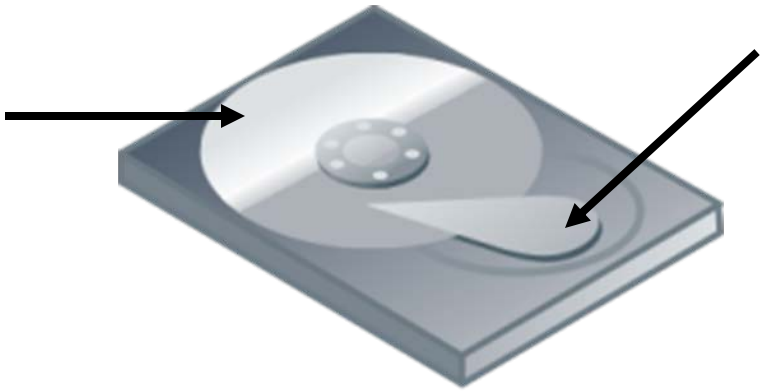

Many components in the head, armature, and motor. At low capacities the costs of these components is dominant.

Disc Drive

Performance Disks

- Much more expensive than consumer disks
- Fully burdened price includes the cost of the controller, often doubles price per GB
- Faster
- Less capacity
- Public price data available from the Storage Performance Council:
  - www.storageperformance.org

Disk Price/GB in Enterprise Array

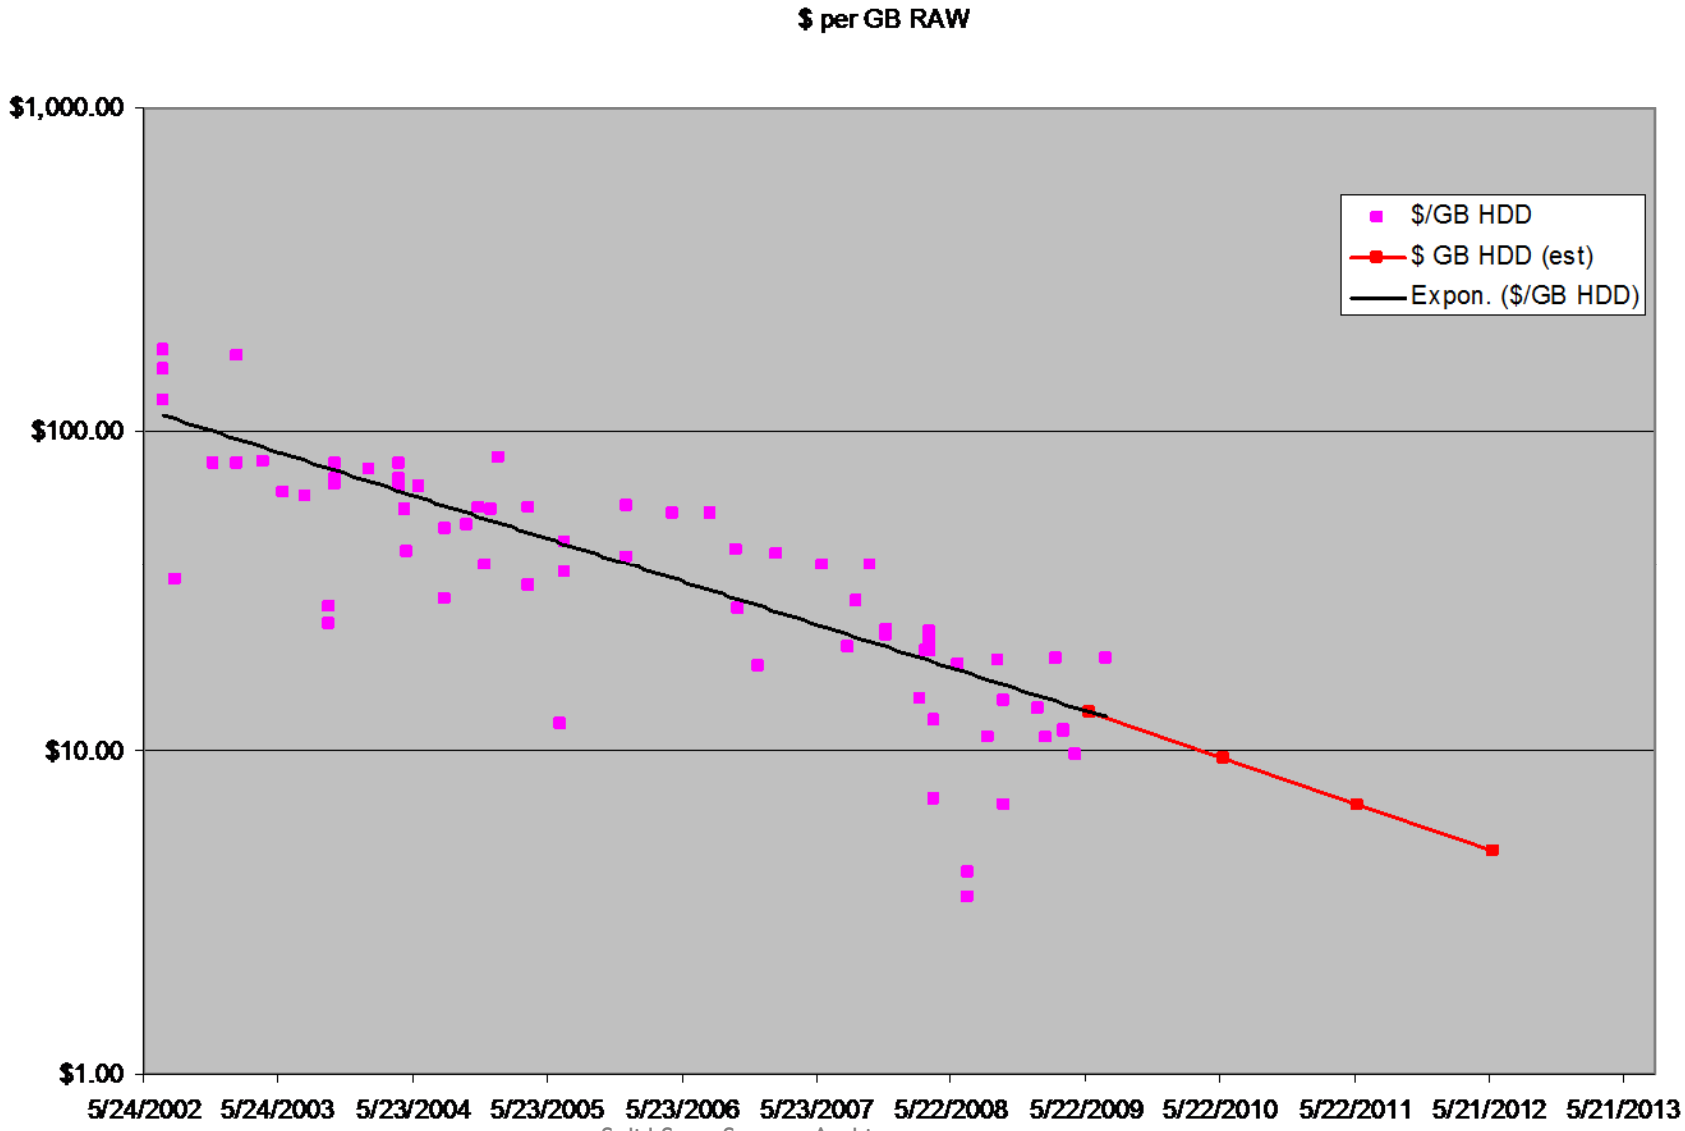

SSS

- Much lower cost for performance
- Much lower power and cooling requirements
- Much lower minimum cost
  - This is why net books mainly use SSS.
- As disks move from 3.5” to 2.5” the price burden of the fixed components gets worse!
- More and more SSS deployments will move from improving performance justifying the cost into a cost savings ROI calculation.

Summary

- Lots of SSS options available and lots of hype.
- There are a few fundamental differences between SSS deployment options.
- Determine the application architecture and requirements for deployment before evaluating vendors offerings to save time.
- SSS is the future due to operational and performance cost savings.

Q&A / Feedback

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

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

- SNIA Education Committee

Matt Key **Brian Groff**