SNIA

Education

Executive Overview and Current Topics in Solid State Storage

Rob Peglar
Xiotech Corporation

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.

Agenda

- Goal of this tutorial:
  Executive overview of Solid State Storage
  - Leading into more detailed Tutorials
- What is Solid State Storage?
- The economics of Solid State Storage
- A survey through various Solid State Storage designs
- Current topics, especially related to enterprise use

- Q&A

A Blast from the Past – August 1981

The Good Old Days?

Processor	8088
Speed	4.77 MHz
RAM	16KB
Storage	Cassette Tape, optionally 5.25" 160KB floppy drives
Expansion	5 expansion slots
Bus	Industry Standard Architecture (ISA)
Video	Initially CGA (320x200x16 color, 640x200x2 color) or monochrome (80x25 text only))
I/O	Parallel, Serial
OS	Basic 1 (ROM)
Killer App	VisiCalc

Fast Forward – to Today

- Today, we have CPUs which are 1,000x faster
  - Instead of MHz, we have GHz
  - Instead of one core, we have multi-core
- Today, we have RAM which is 1,000,000x larger
  - Instead of KB, we have GB
  - Some machines are approaching TB (!)
- Today, we have storage which is 1,000,000x deeper
  - Instead of MB, we have TB
- So what's the problem?

The Problem WAS – AND IS – I/O

- In a perfect world, I/O would not be necessary
  - 1st level store would hold everything, forever
- Access Density – IOPS/GB
  - Getting WORSE over time for rotating magnetic
  - Will it get worse over time for non-rotating SSD?
- Access Density Example (more is better):
  - August 1981 – 625 KB/s, 8.33 ms, 3,600 RPM, 20 IOPS
    - IOPS/GB = 20 / .001 = **20,000**
  - Today – 170,000 KB/s, 2.9 ms, 15,000 RPM, 250 IOPS
    - IOPS/GB = 250 / 300 = **0.833**

Solid State Storage – What is it?

- Storage devices constructed from chips instead of rotating platters or streaming tape
- Solid State Storage is not new
  - Dates back to the 1980's and mainframes
- Various form factors
  - HDD forms using HDD interfaces (e.g. FC, SATA)
  - Non-HDD forms such as PCI-Express cards
- Two basic technologies:
  - DRAM
    - Usually accompanied by battery/persistent media
  - NAND Flash
    - Other Flash types are becoming less viable

Solid State Storage – Why?

Images of HDD and Representative SSS

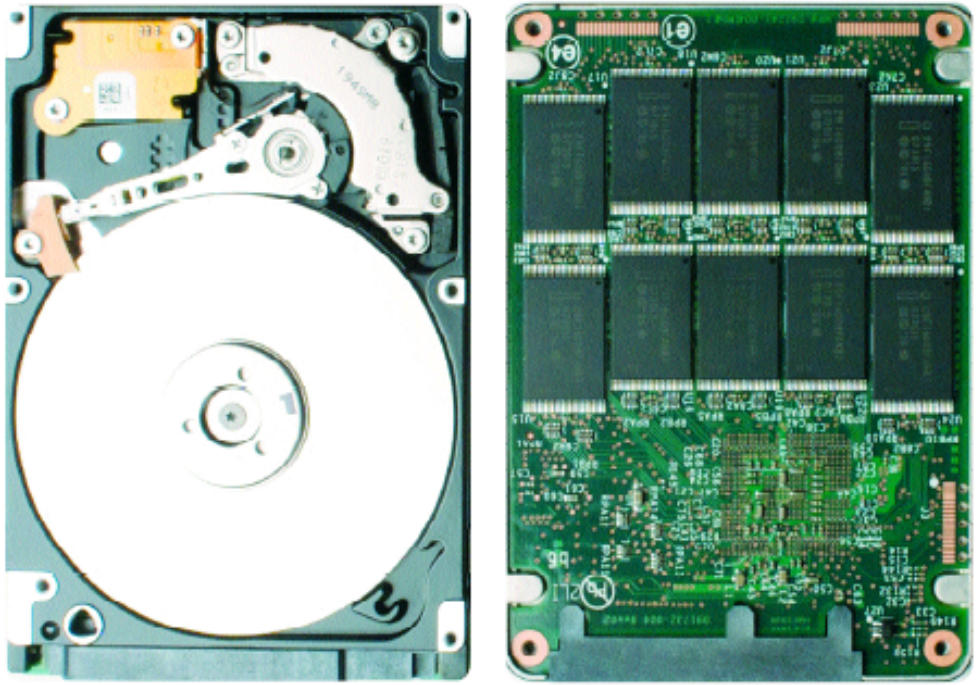

Photo provided by Intel Corp 2008

- Flash as Magnetic Disk Replacement
  - Write cycles, cost/GB, media lifetime, TCO
- DRAM as (controller) Cache Replacement
  - Cost/GB, TCO, expandability/flexibility
- Consider the Application Workloads
  - Transactional (random) versus Batch (sequential)
  - Block versus Files
  - Structured versus Unstructured

What functionality do users need?

Education SNIA

Application aspects of storage

- **Capacity**
  - Disk or Tape Size
  - Number of disks/channel
  - Number of tape devices
- **Performance**
  - Disk latency & seek time
  - Cache size & hit rate
  - Media rotation rate (RPM)
  - Responsiveness
- **Availability**
  - MTBF/MTTR (Rebuild time)
  - Path redundancy
  - Path bandwidth

Physical aspects of storage

- **Capacity**
  - Application requirements
  - Structured / unstructured
  - Growth potential
- **Performance**
  - Throughput / IOPS
  - Responsiveness
- **Availability**
  - Failure resistance
  - Recovery time/point
    - **RTO/RPO**
  - Simplification of change

Where does Solid State Fit?

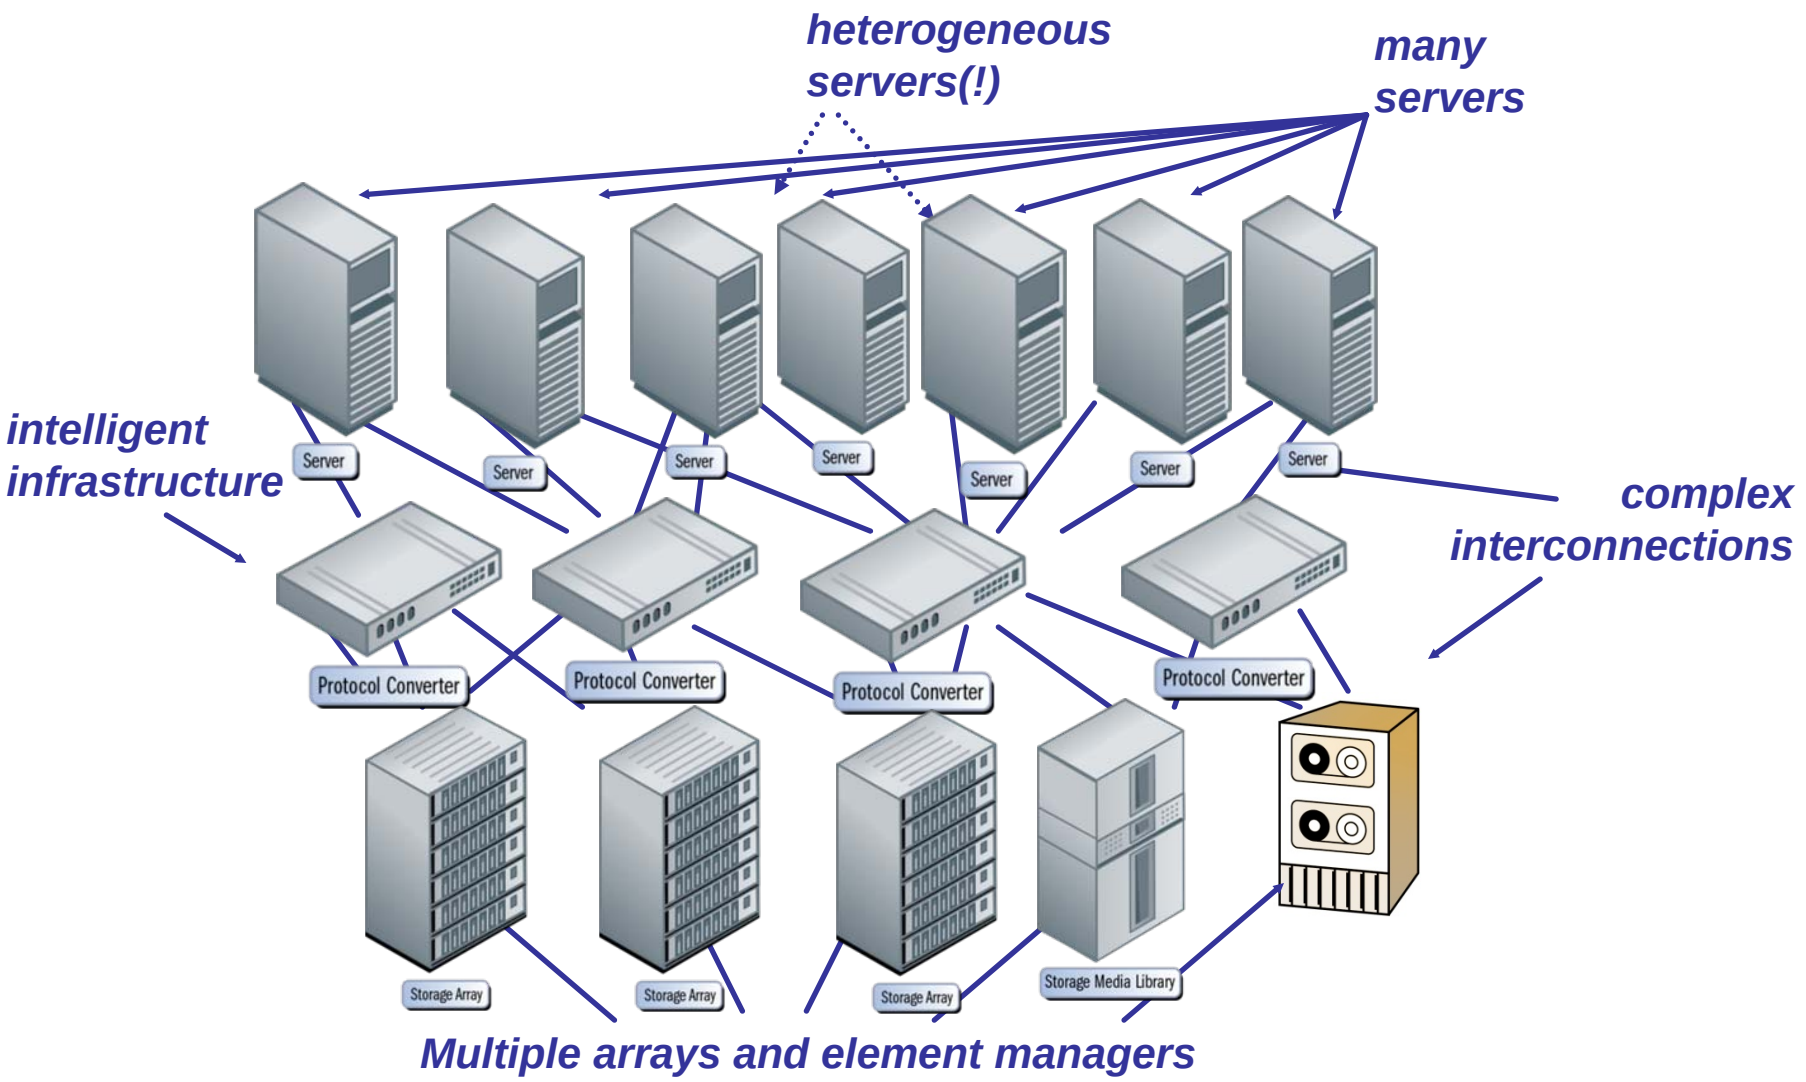

Application Requirements

- Applications don't want disks
  - They want space (more is better)
- Applications don't want IOPS
  - They want time (less is better)
- Applications do I/O because they *have* to
  - But they don't really *want* to
- The problem is not applications, it's application programmers and the OSes

Real-World Application Workloads

- Unstructured data
  - Unstructured data access is a poor fit for SSD
  - Exception – small, non-growing tagged files
  - OS images – boot-from-flash, page-to-DRAM
- Structured data
  - Structured data access is an excellent fit for SSD
  - Exception – large, growing table spaces
- Databases have key elements
  - Commit files – logs, redo, undo, tempDB
- **The key for senior management is to understand what workloads are critical to the business – ask the VPs of the business units**

Disc (Drive) Virtualization

Education SNIA

Physical disc drive

Disk Virtualization

LBA

000
001
002
003
004
005
006
.. nnn

Logical data layout

- Logical Block Addresses (LBA)
- 'Defect-Free'

Physical data layout

- C-H-S Addresses
- Media defects

Check out SNIA Tutorial:

Virtualization I – How, …?

Economics of Solid State Storage

- Memory densities 64x what they were 3 years ago
- $/GB is decreasing rapidly
  - But still nowhere near where rotating disk is $/GB
- Key metrics are now $/IOP, $/IOP/U, $/IOP/W
- SSS is 100-1000x faster access time
- Bandwidth:
  - Read bandwidth typically limited only by interface
  - Write bandwidth typically limited only by controller design
- Read/Random is typically the best use case (existing SSD)
  - Very consistent I/O response time

Check out SNIA Tutorial:
Economics of Solid State Storage: Perception and Reality

Business Impact of Solid State Storage

- Server CPUs today are multicore, GHz speeds
- Millions of CPU cycles during just one HDD seek
- SSS can potentially eliminate waste:
  - Server infrastructure – reduce CPU load/wait time
  - Storage infrastructure – reduce short-stroking HDDs
  - Applications – reduce I/O bounded-ness
  - Humans – reduce wait time, screen refreshes, queries
  - Data Centers – reduce power, cooling load, rackspace
- Reliability – a whole 'nother topic …
- SSS can potentially drive workloads more efficiently and thus save $$ and save time

Impact of Solid State Storage

- Use of internal techniques such as striping across multiple chips – parallel access – **saves time**

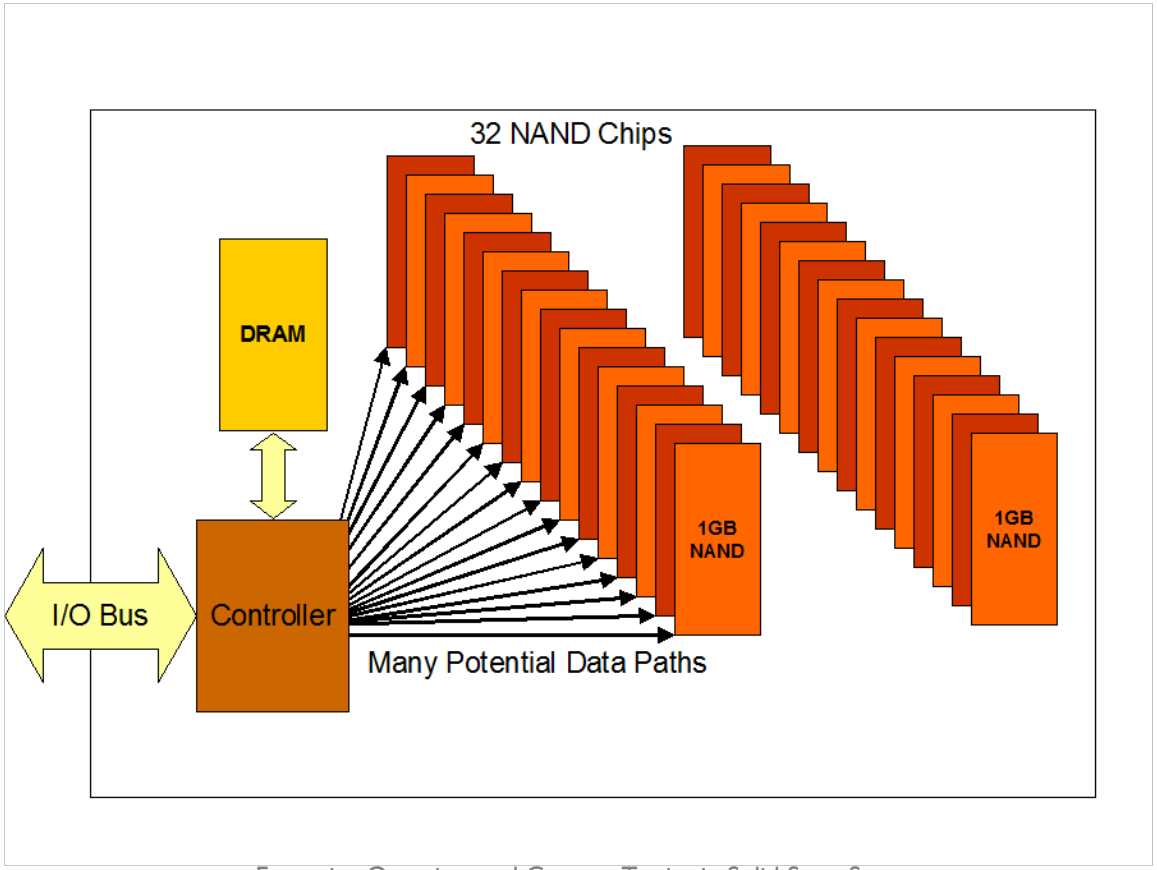

Tiered Storage and Solid State

- SAN
- NAS
- DAS

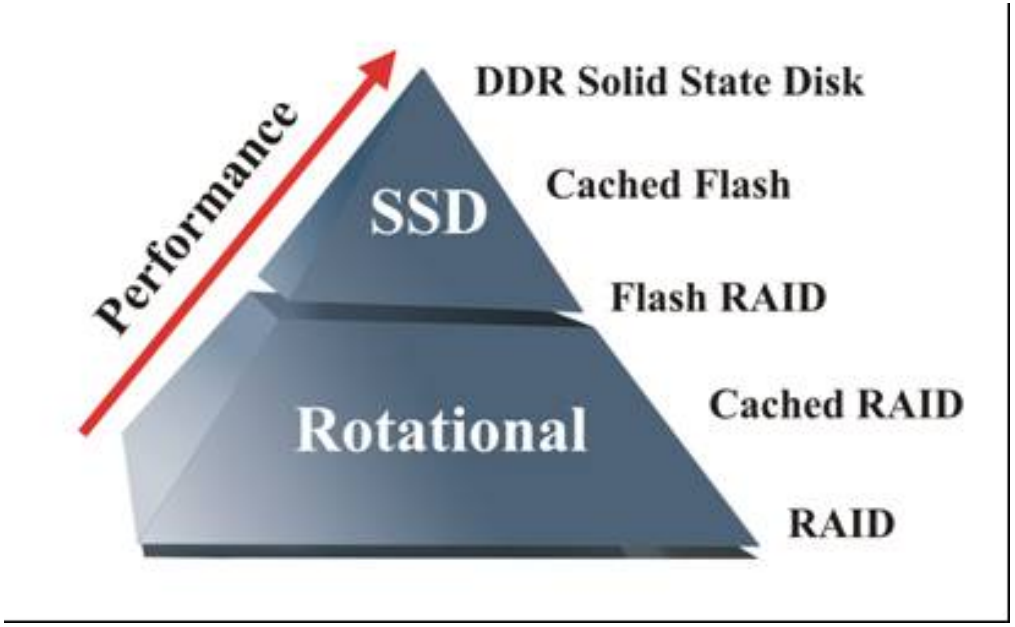

Architectural Approaches in SSS

- Arrays – take advantage of existing function
- Hybrid approach – use SSD and HDD in same array

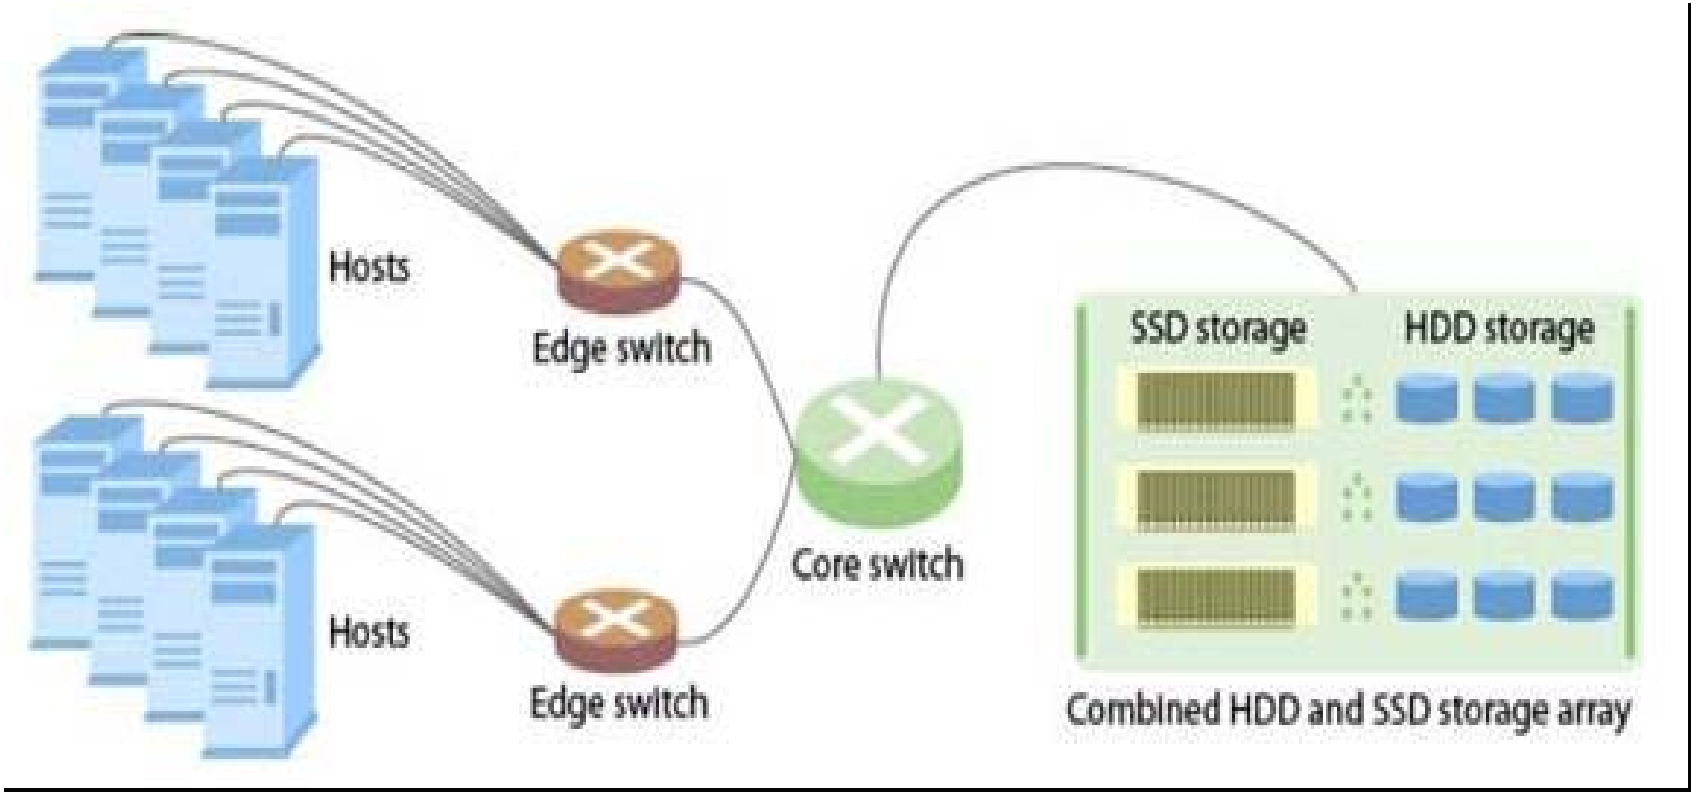

Architectural Approaches in SSS

- Arrays – take advantage of existing function
- Hybrid approach – use SSD and HDD in different arrays

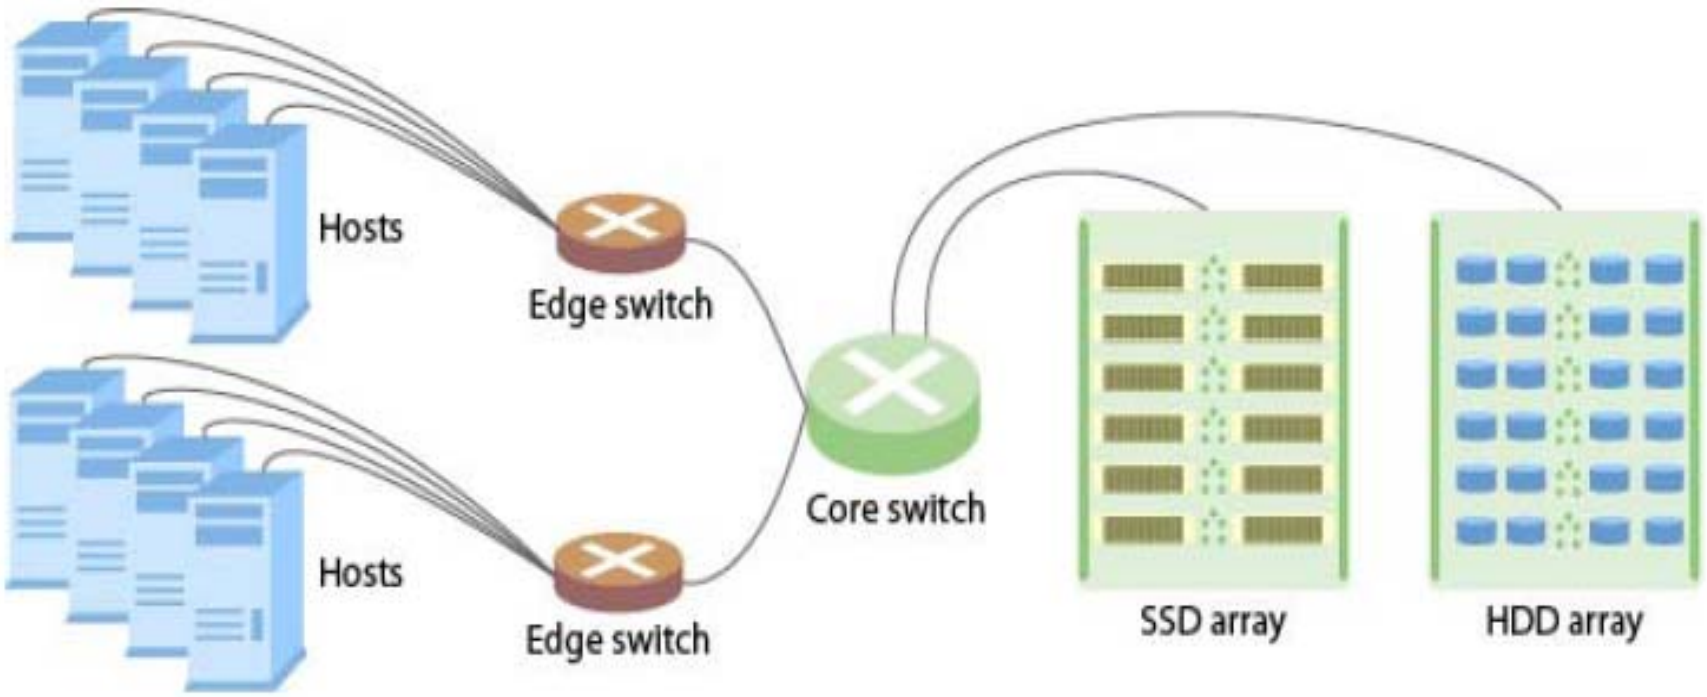

Architectural Approaches in SSS

- Servers – take advantage of existing busses/interconnect
- Brings SSS closer to the CPU/RAM, architecturally

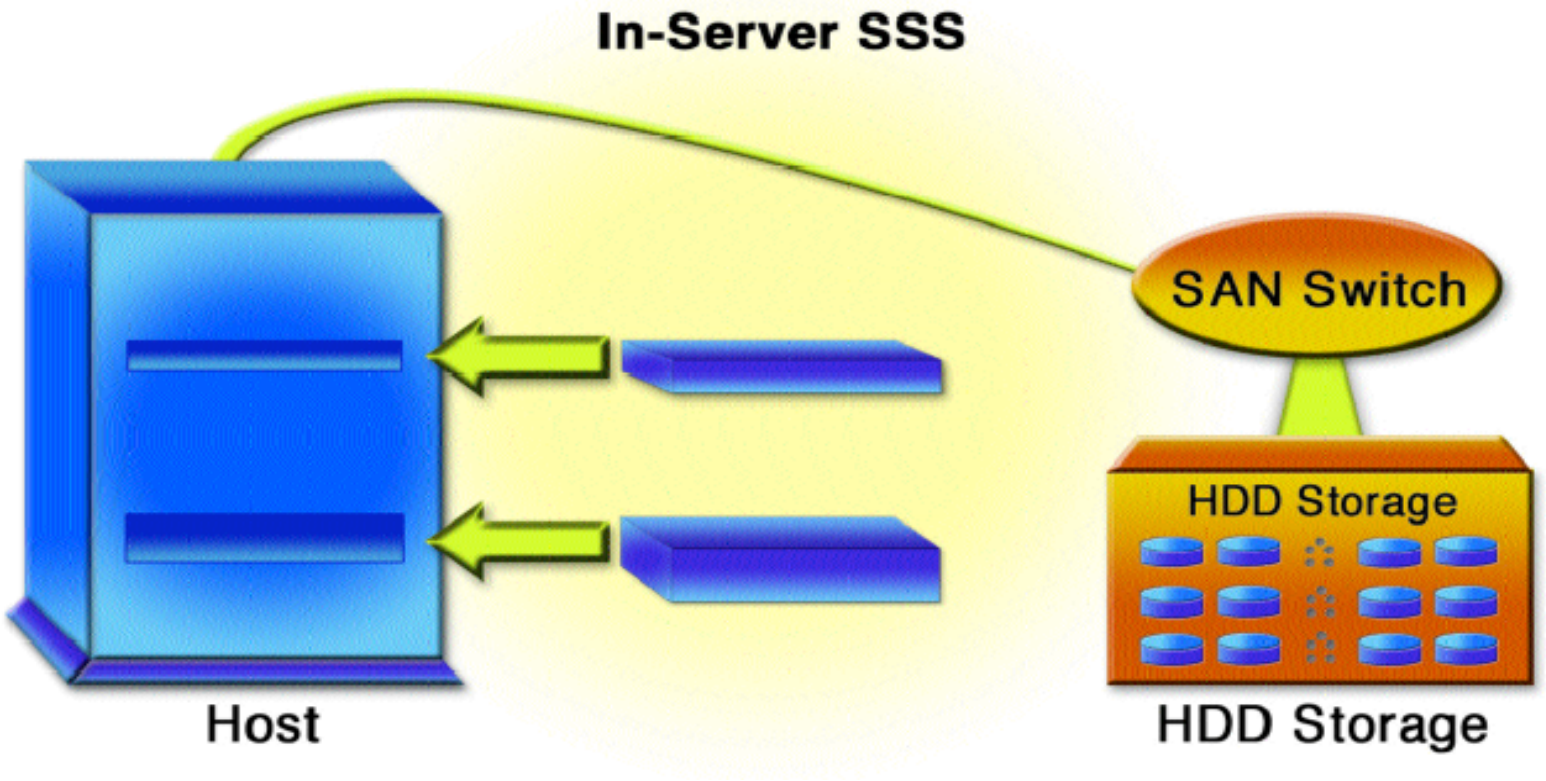

Architectural Approaches in SSS

- Standalone – take advantage of specifically isolating SSS
- Brings SSS to bear on a single server/compute job

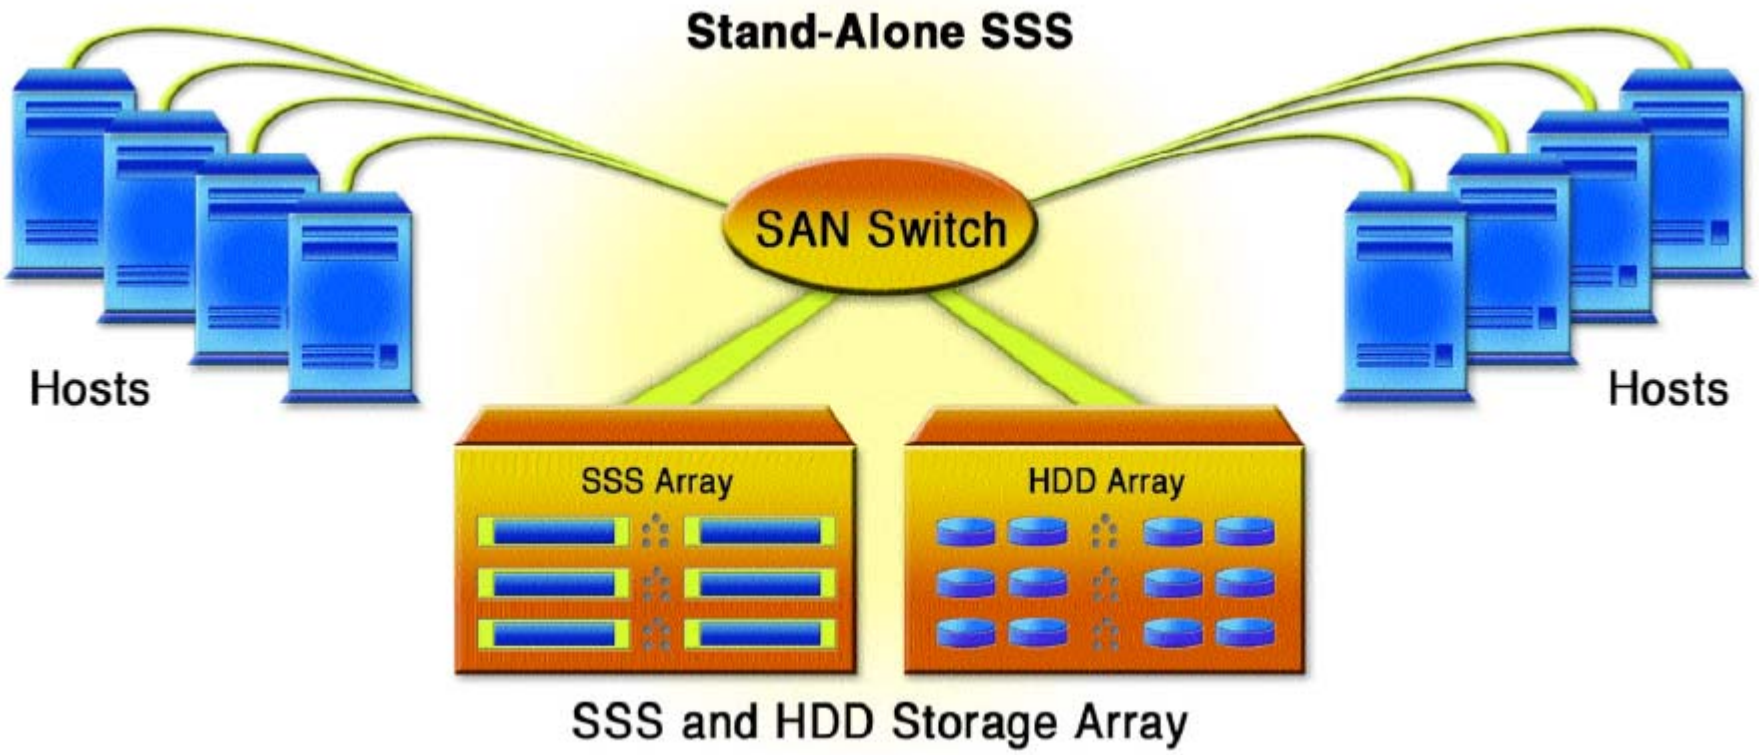

Final Thoughts on Solid State Storage

- What is the business impact on application workloads?
  - SSS may be better→ideal for virtualized workloads
  - SSS may be better for multiple workloads with heavy reads
    - E.g. two simultaneous sequential workloads = one large random
- What is the impact on server procurement?
  - Potential to entertain swap of RAM for SSS
  - Potential to use bus extension instead of channels
  - Potential to improve in-box reliability (no moving parts)
- Best practice – use all layers of the storage pyramid
- Application mix may change – layers remain the same

Check out SNIA Tutorial: SSS Performance

Final Thoughts on Solid State Storage

- What is the 'cost' of storage?
  - Acquisition Cost – very straightforward, a single figure
  - Operational Cost – not straightforward, multiple figures
  - Performance Cost – usually never calculated – but it should!
  - E.g. "if this job ran in 2 hours instead of 10, what's the $$$?"
- How to improve Performance Cost?
  - Replace slow components with fast(er) components
  - Add more components at the same or faster speed
    - VMs, CPU, RAM, slots, HBAs, channels, controllers, disks, …
- Best practice – use all layers of the storage pyramid
- Application mix may change – layers remain the same

Q&A / Feedback

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

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

Neal Ekker
Phil Mills
Jonathan Thatcher
Marius Tudor

SNIA Shared Storage Model *A Layered View*

IV. Application

III. File/record layer

IIIa. Database

IIIb. File system

II. Block aggregation

IIa. Host

IIb. Network

IIc. Device

I. Storage devices

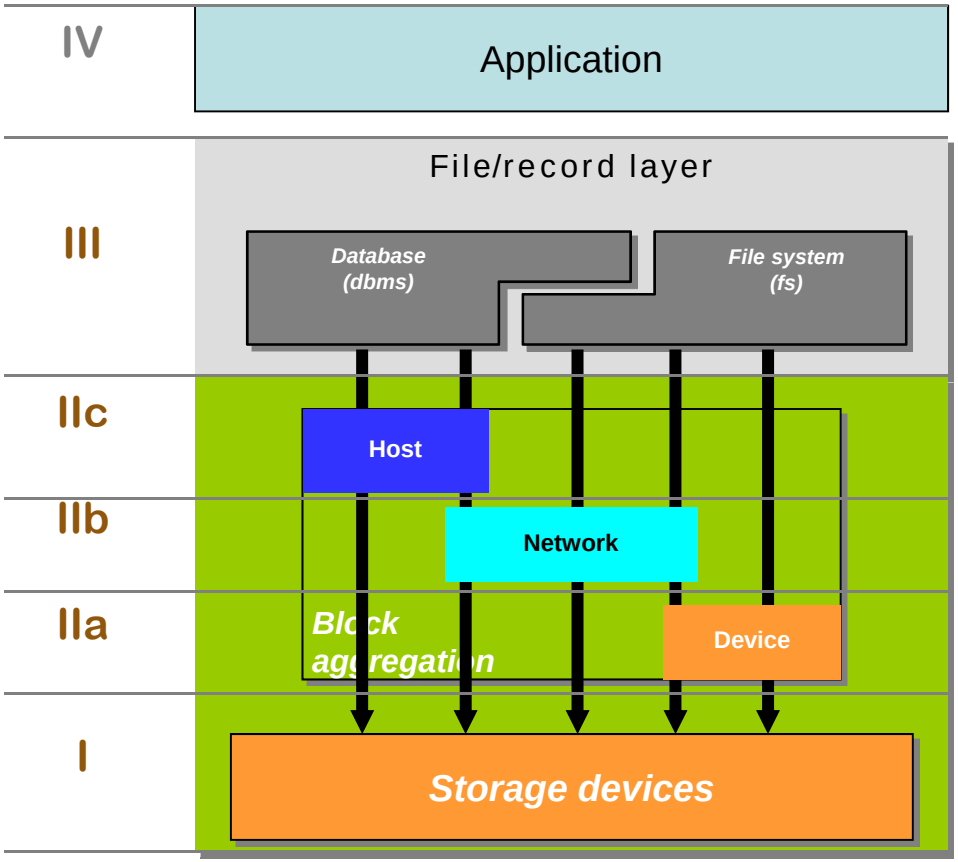

The SNIA Shared Storage Model uses the term “aggregation“ instead of “virtualization“