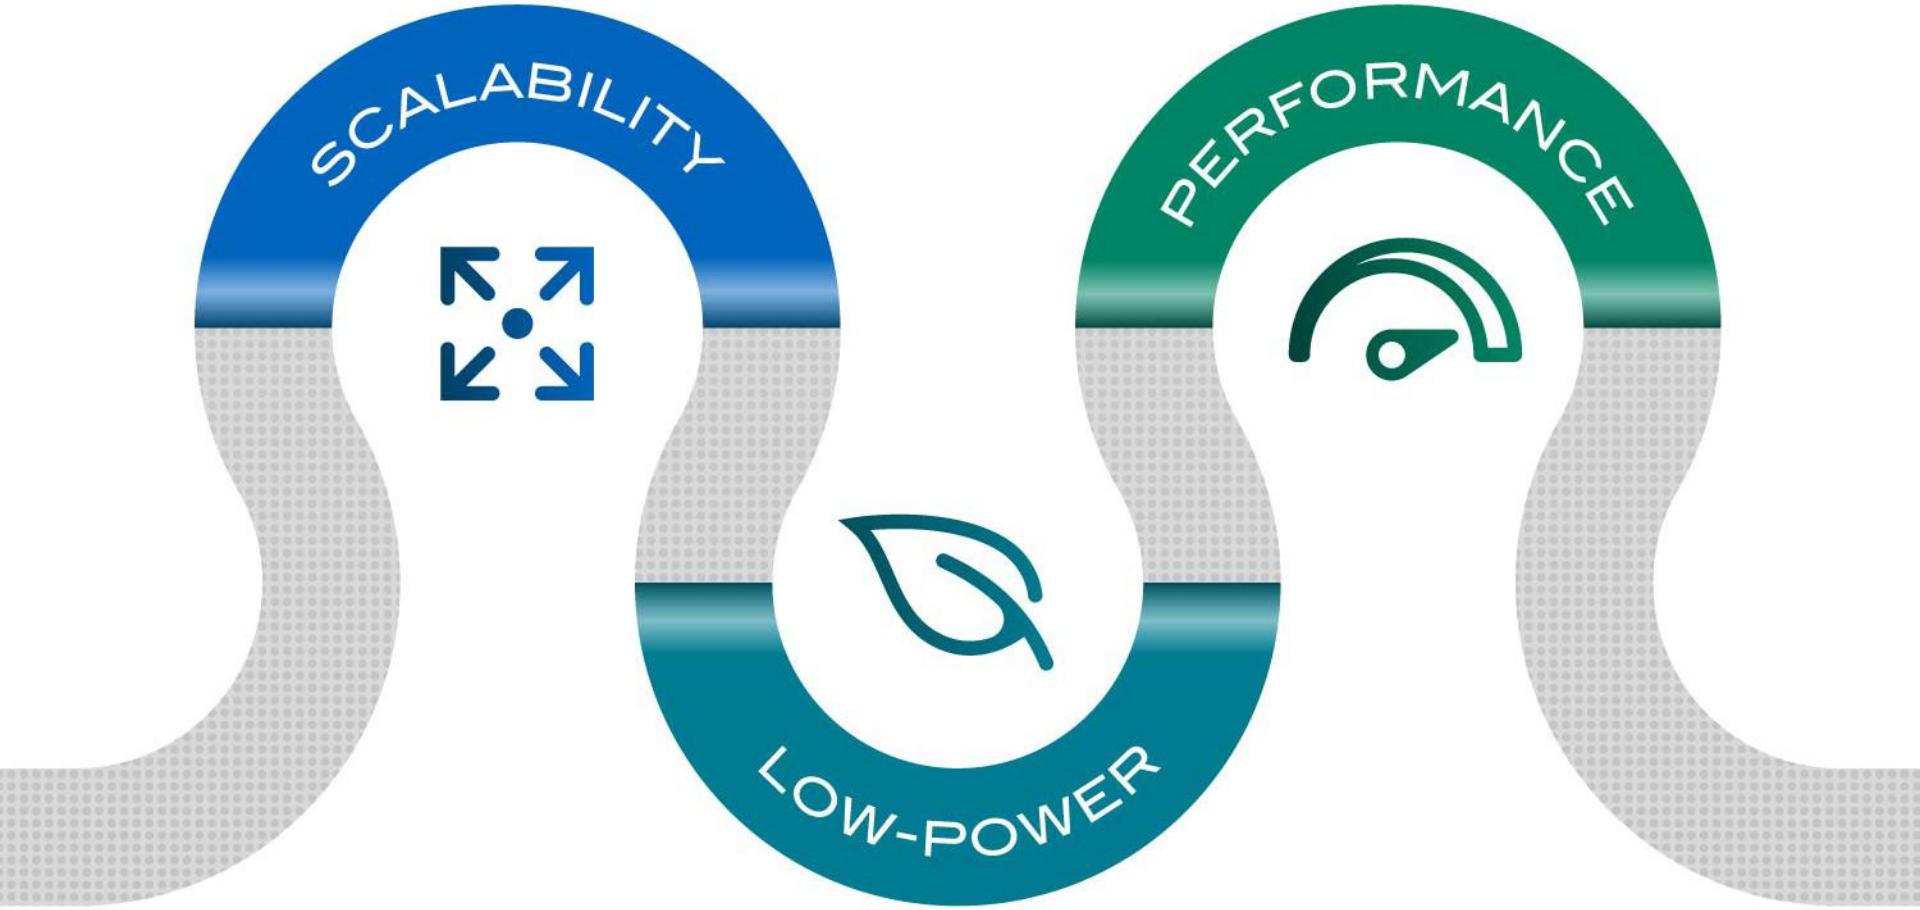

Single Controller 4/8TB SSD

Flash Memory Summit 2013
Santa Clara, CA

HakJune Oh, Director, Product Development
MOSAID Technologies Inc.

Agenda

- **Introduction**
- **How to make a Big Capacity SSD**
- **8TB HLSSD with Single Controller**
- **HLSSD vs. Conventional SSD**
- **Summary**

HLNAND Introduction

- High Scalability – no practical limit on number of NAND die per channel with ring architecture – 16TB SSD in 3.5” form factor
- Excellent Throughput – point-to-point connection provides excellent signal integrity – 1600MB/s is achievable
- Low Power – single controller, low voltage IO, un-terminated bus, data truncation, hierarchical MCP – 50% lower power than conventional SSD
- Reduced System Cost – single controller, small footprint, reduced networking infrastructure - lower data center TCO

SSD Capacity

SSD Capacity =
(# of Channels) X (# of Ways/Channel)
X (NAND Die density)

- Flash controllers with 16- or 32- channels are too complex, too big and consumes high power
- # of ways / channel limited to 4 ~ 8 NAND dies because of RC loading effect in high speed DDR NANDs
- NAND die density is now 128Gb in MLC 1ynm

How to make high capacity SSD (Method 1)

- Make SSD controller with 16- or 32-channel
  - Chip size is too big → Low yield
  - High # of signal leads → Signal Integrity
  - Active and standby power consumption ↑
  - # of ways per channel is still limited to 4~8
  - Channel write bandwidth is limited to # of ways per channel

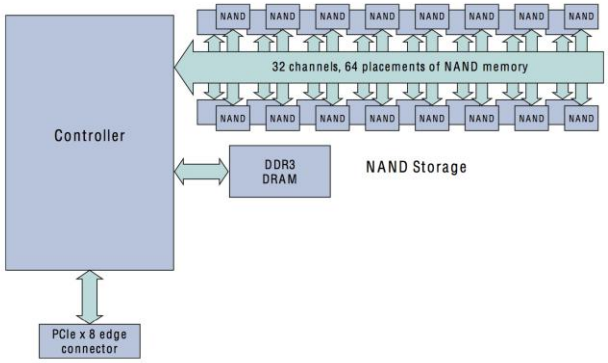

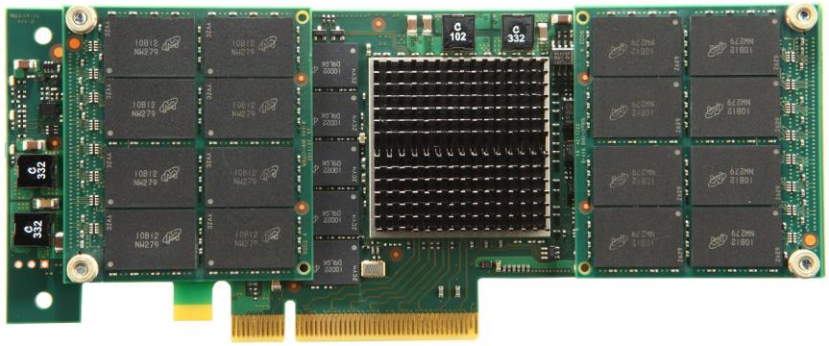

How to make high capacity SSD (Method 2)

- Make a Bridge-based SSD
  - PCIe-SAS/SATA RAID (RoC) + SAS/SATA SSDs
  - SAS/SATA protocol overhead
  - Medium IOPS with possible RoC bottleneck
  - Higher power and added RoC power

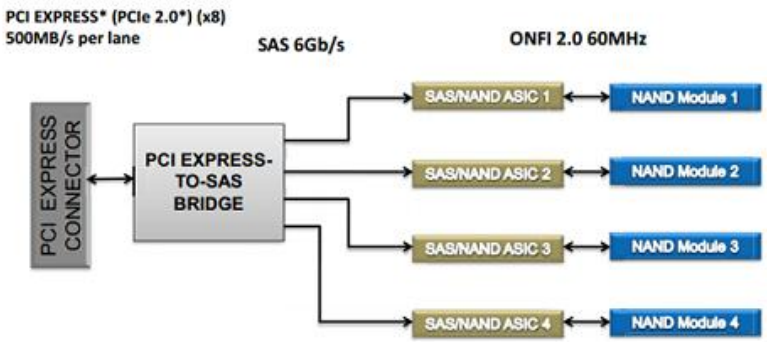

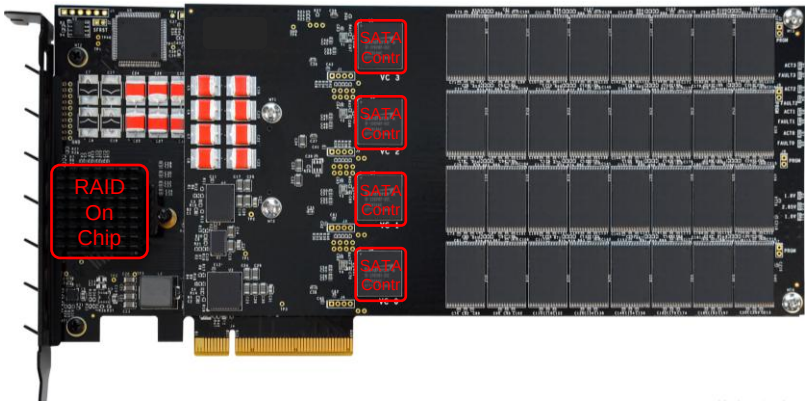

HLNAND Flash

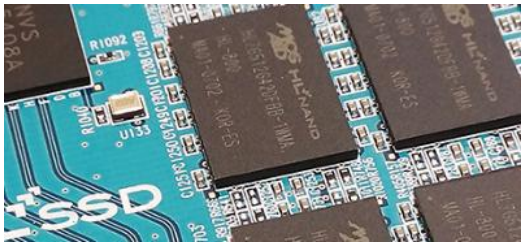

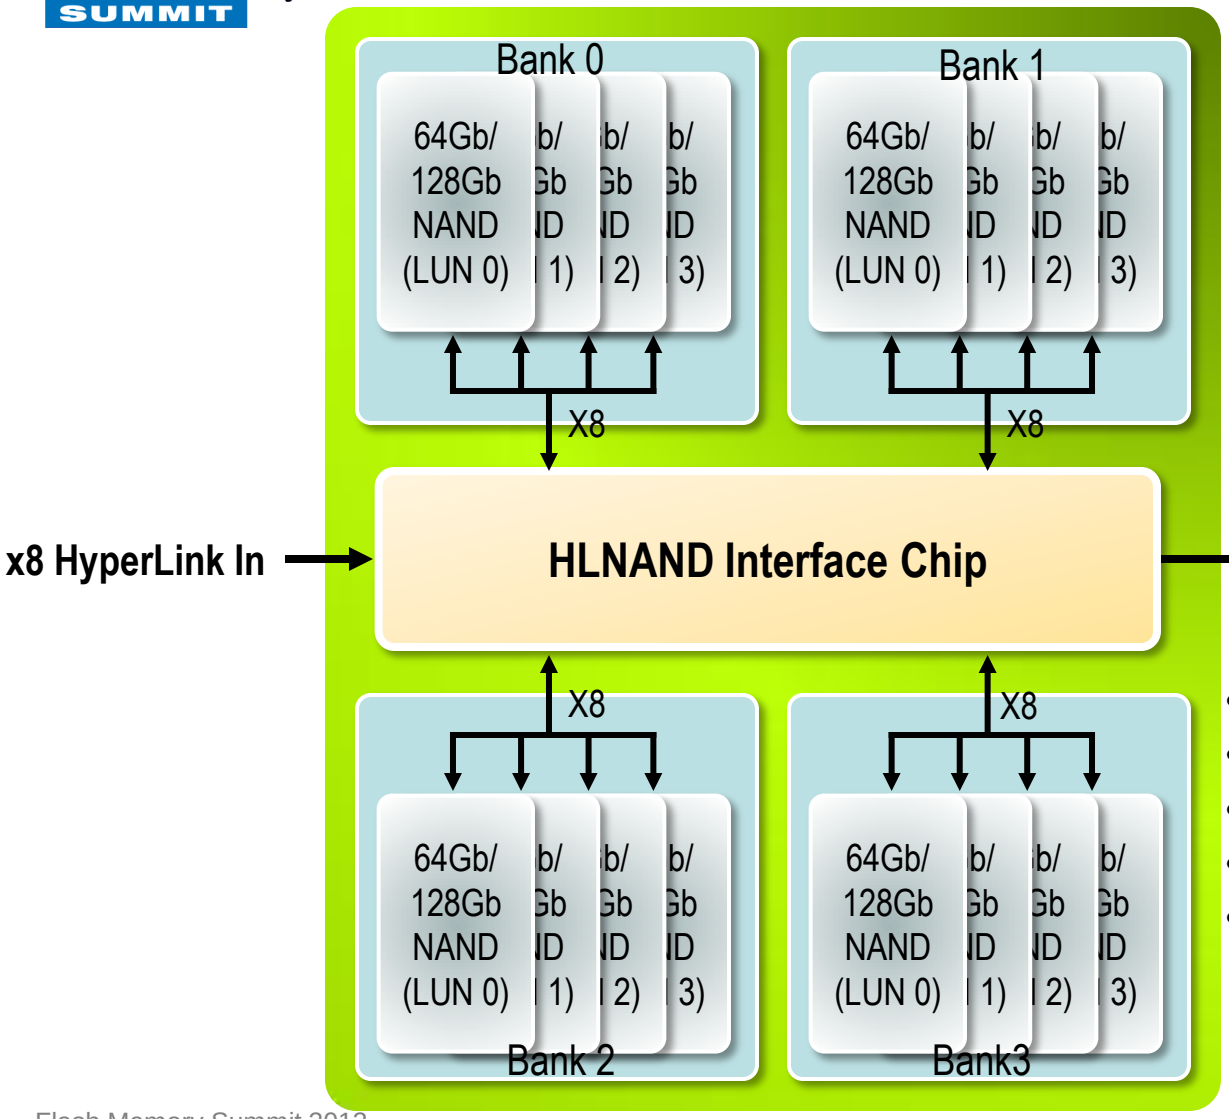

- DDR-533/667/800
- Up to 16-Die-Stack
- Up to 128Gb NAND Die
- Up to 2Tb / Package
- Support SLC, MLC, eMLC, TLC in 2x, 1x, 1y nm Process Technology

HLNAND Controller

- Dual-Core with Hardware Accelerator
- 28nm Process Technology
- SATA3 & PCIe Gen.2 X4, NVMe
- 8-CH HyperLink
- Up to 16TB Capacity
- Enhanced 2-Dimensional data randomization
- Enhanced BCH ECC up to 80bits/2KB
- Read refresh by bit error monitoring
- End-to-end data protection
- Full disk encryption with AES-128/256 ECB/CBC/CTR/XTS
- Smart power management for low and peak-power control
- Package: 17mm x 17mm FCFBGA

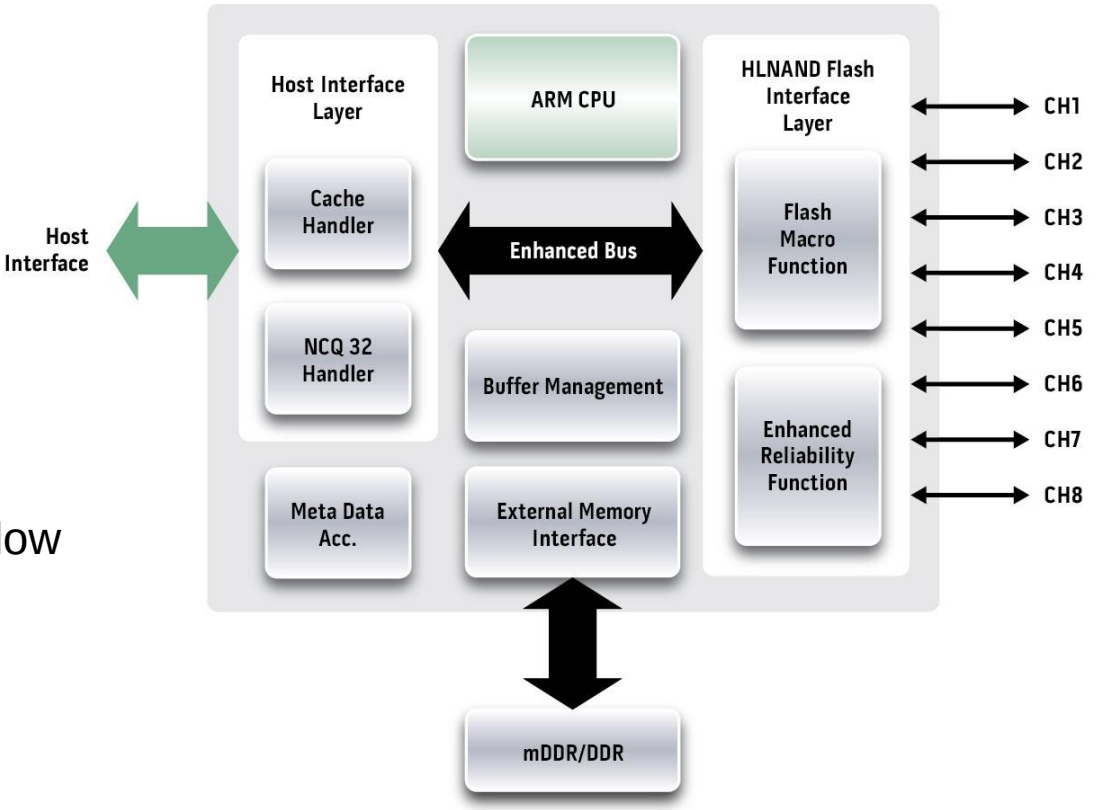

HLNAND Architecture

- Point-to-Point Daisy-chain Ring Topology Channel
- Synchronous DDR signaling (DDR-266 ~ DDR-1600)
- Packet Protocol Command & Addressing
- Up to 128 HLNAND MCPs* per single channel

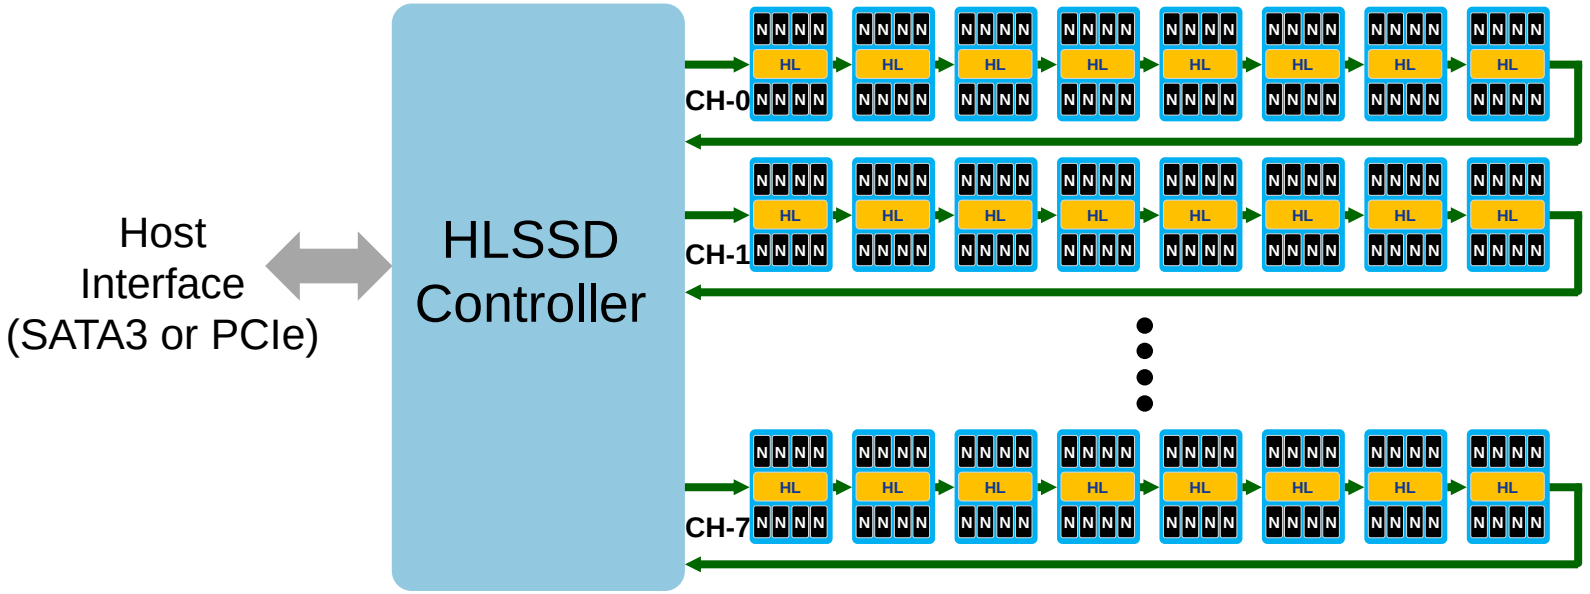

* Maximum available number of NANDs per channel = 2048 (based on 16-die-stack)

HLSSD

- Interface: Native SATA3 & PCIe Gen 2, 4 lanes (NVMe)
- Form Factor: 2.5”, 3.5”, SFF-8639, FHHL PCIe, FHFL PCIe
- Capacities: 256GB – 16TB
- Customized firmware for targeted use-cases (workload)
  - Up to 560MB/s R/W & up to 100K IOPS 4K Random R/W in SATA3
  - Up to 1.8GB/s R/W up to 320K IOPS 4K Random R/W in PCIe x4

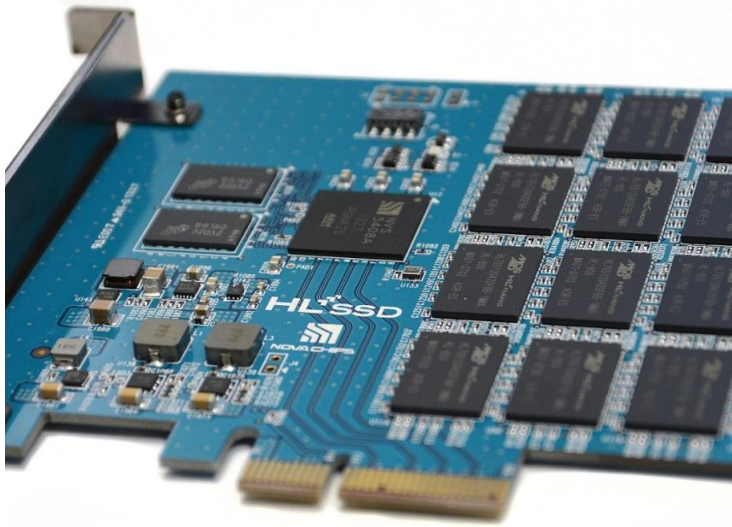

SFF-8639
Form Factor

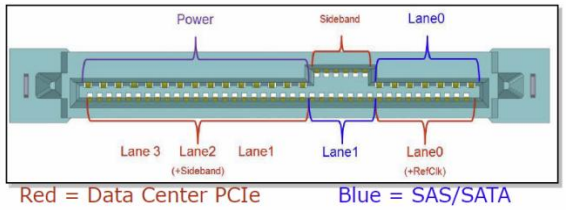

HLSSD Scalability

- 128 HLNAND MCPs (64 MCPs/side)
- 8 channels
- 16 MCPs per channel
- 8TB with 8-die stacked HLAND MCP (64Gb NAND die)
- 16TB with 8-die stacked HLAND MCP (128Gb NAND die)
- PCIe Gen 2.1, x4 lane, NVMe

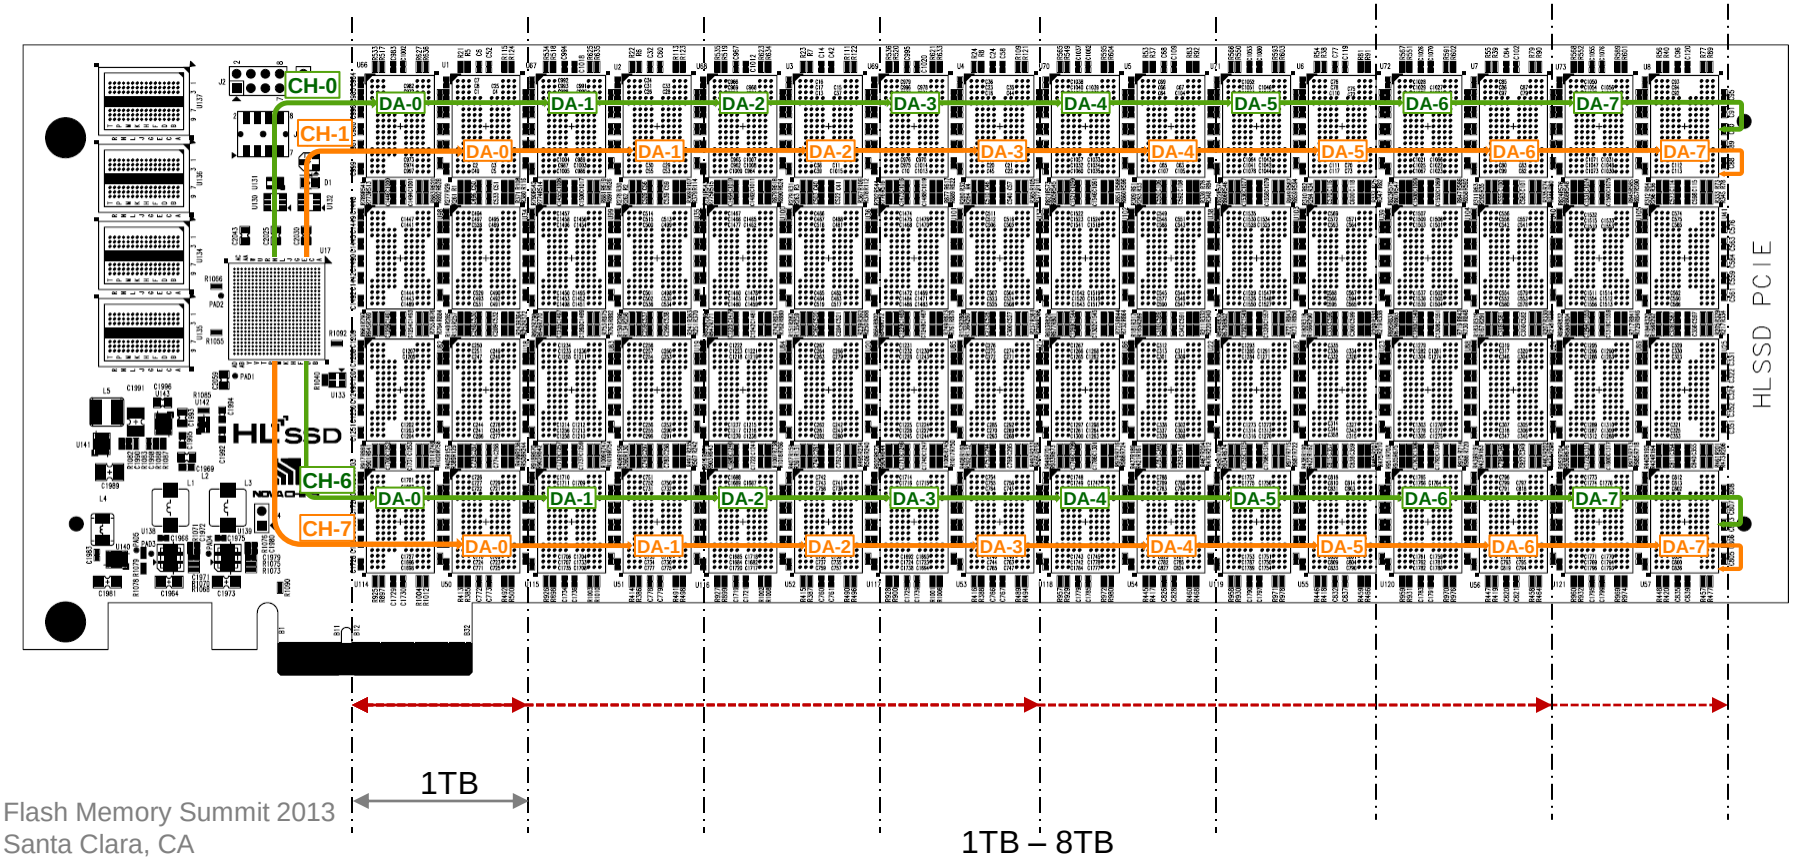

HLSSD Scalability (Cont'd)

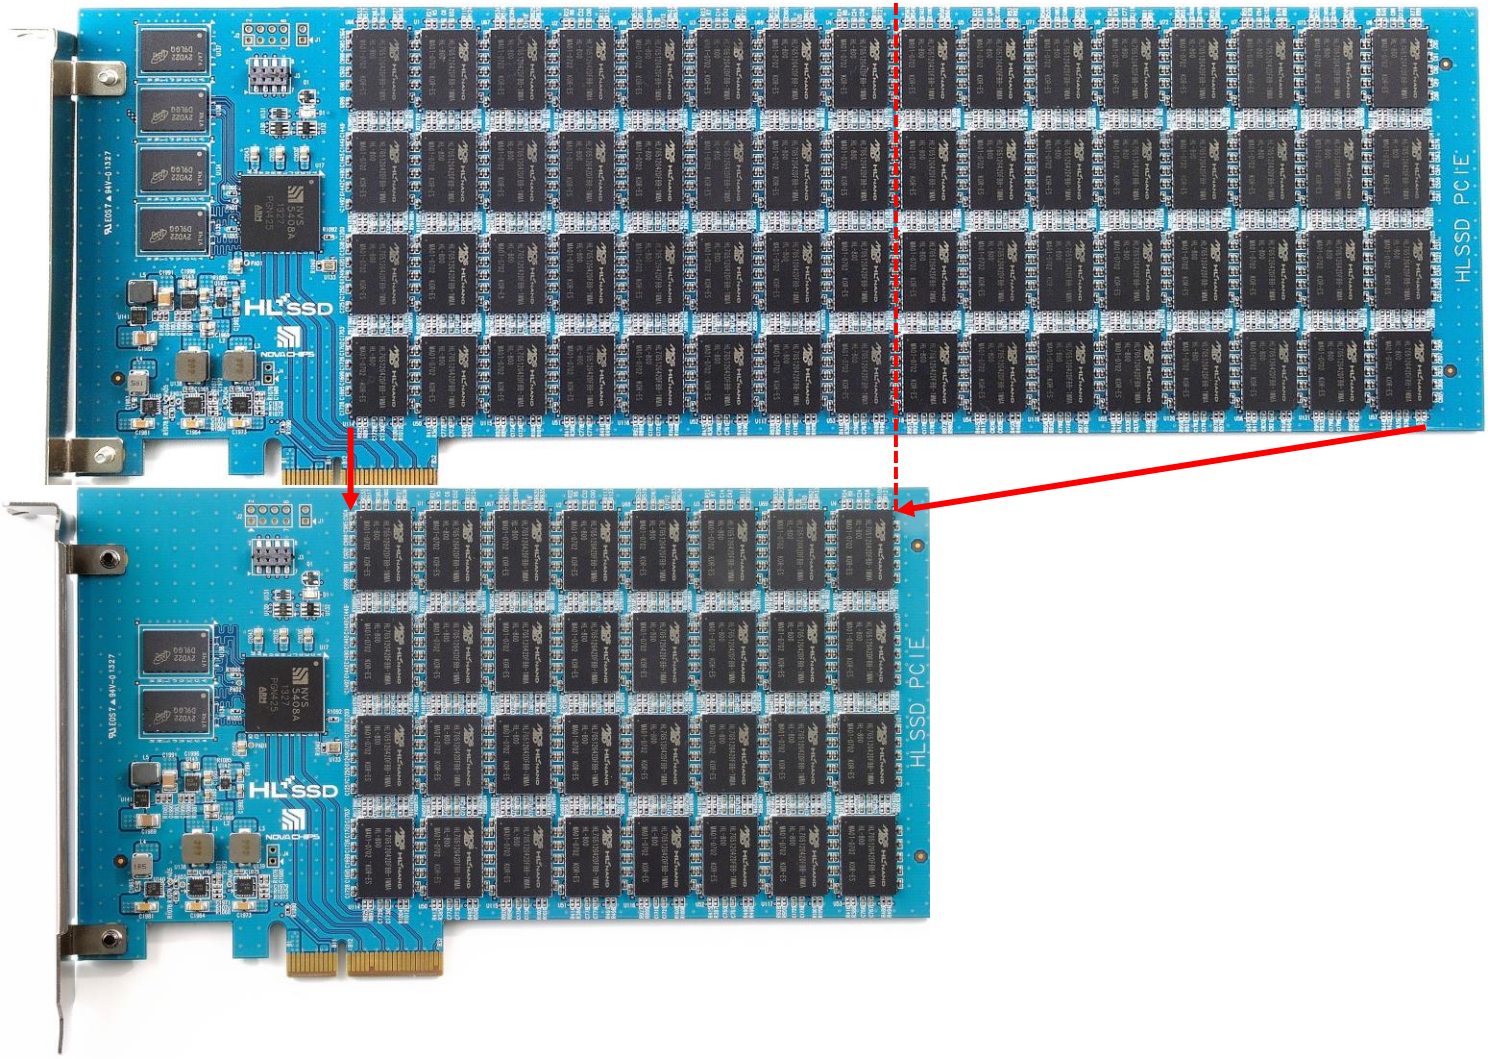

Power Comparison
(Bridge-based 8TB SSD)

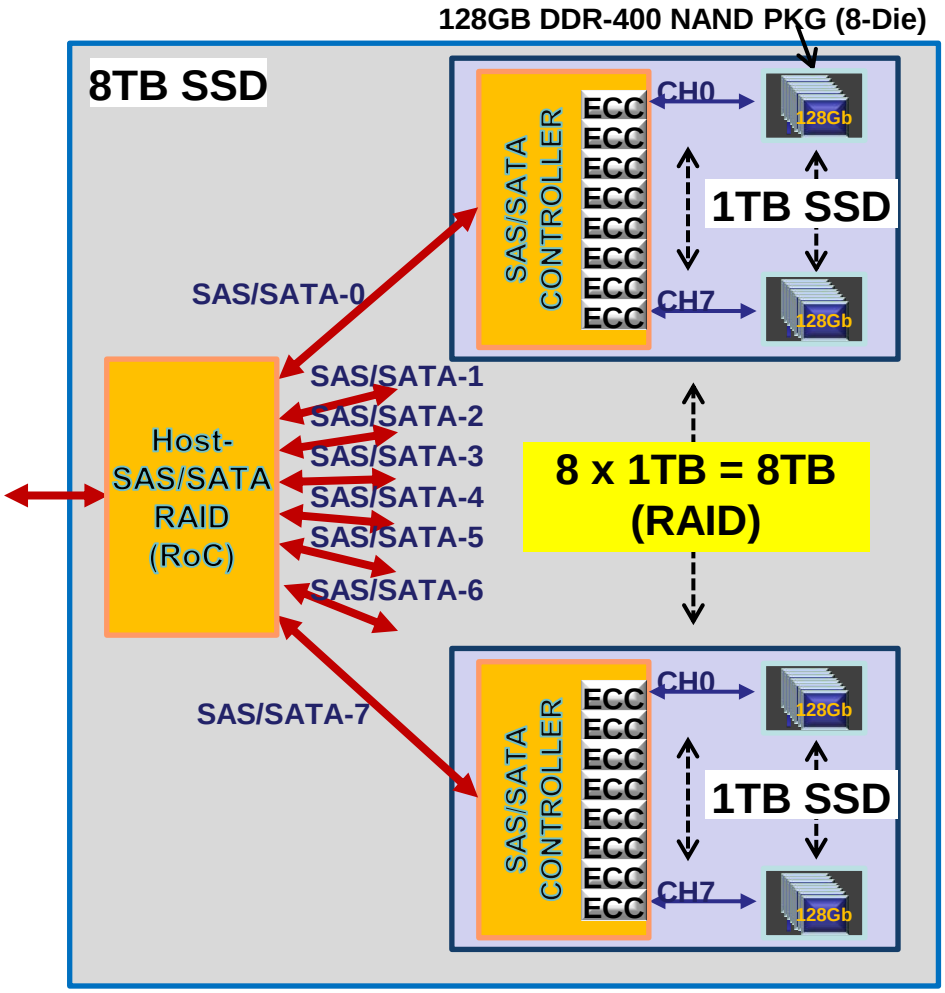

- 1 Host-SAS/SATA RAID (RoC) +
- 8 x 1TB SSD, each having;
  - 1 SAS/SATA Controller & 8 Channels, each having
    - 1 x 128GB PKG (8 NAND dies)
    - 400MB/s bandwidth
    - Vcc=3.3V Core, Vccq=1.8V
    - On-Die-Termination
    - 8 NAND die load
      - DQ pad – 4pF typ.
      - Input pad – 4pF typ.
    - PCB trace – 1pF
    - 24 Controller pins/Channel
      - 8 DQ
      - 8 CE#
      - 8 Controls (CLE,ALE,WE#,RE,RE#,DQS, DQS#,WP#)

Power Comparison
(Native 8TB HLSSD)

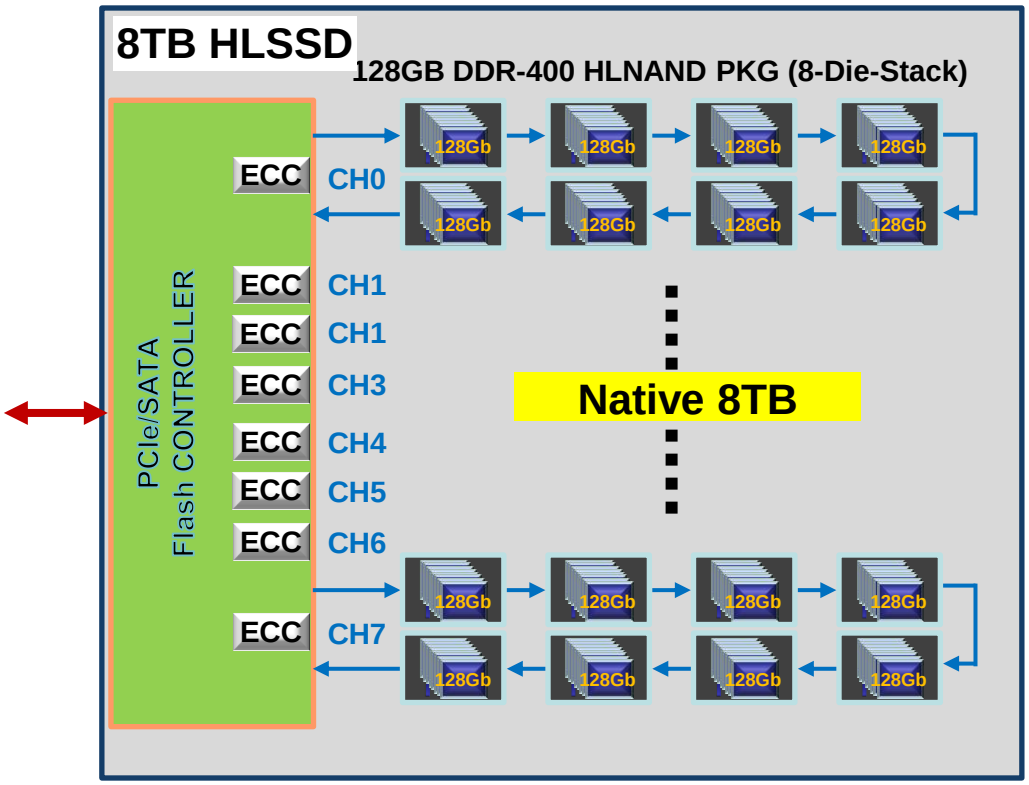

- 8 Channels
- Each channel having;
  - 8 x 128GB HLNAND MCP
  - 400MB/s bandwidth
  - 1.2V I/O
  - 24 controller pins/channel
    - 8 Data in
    - 8 Data out
    - 4 Clocks (diff. in/out)
    - 4 Controls
  - Bridge chip load
    - Input pad – 4pF
    - Output pad – 4pF
  - PCB trace – 1pF

8TB SSD Read Power

(SATA3 BW Saturation)

- HLSSD consumes 60% less power
- Conventional SSD consumes significant SATA channels power
- HLSSD achieves better power efficiency with reduced controller complexity

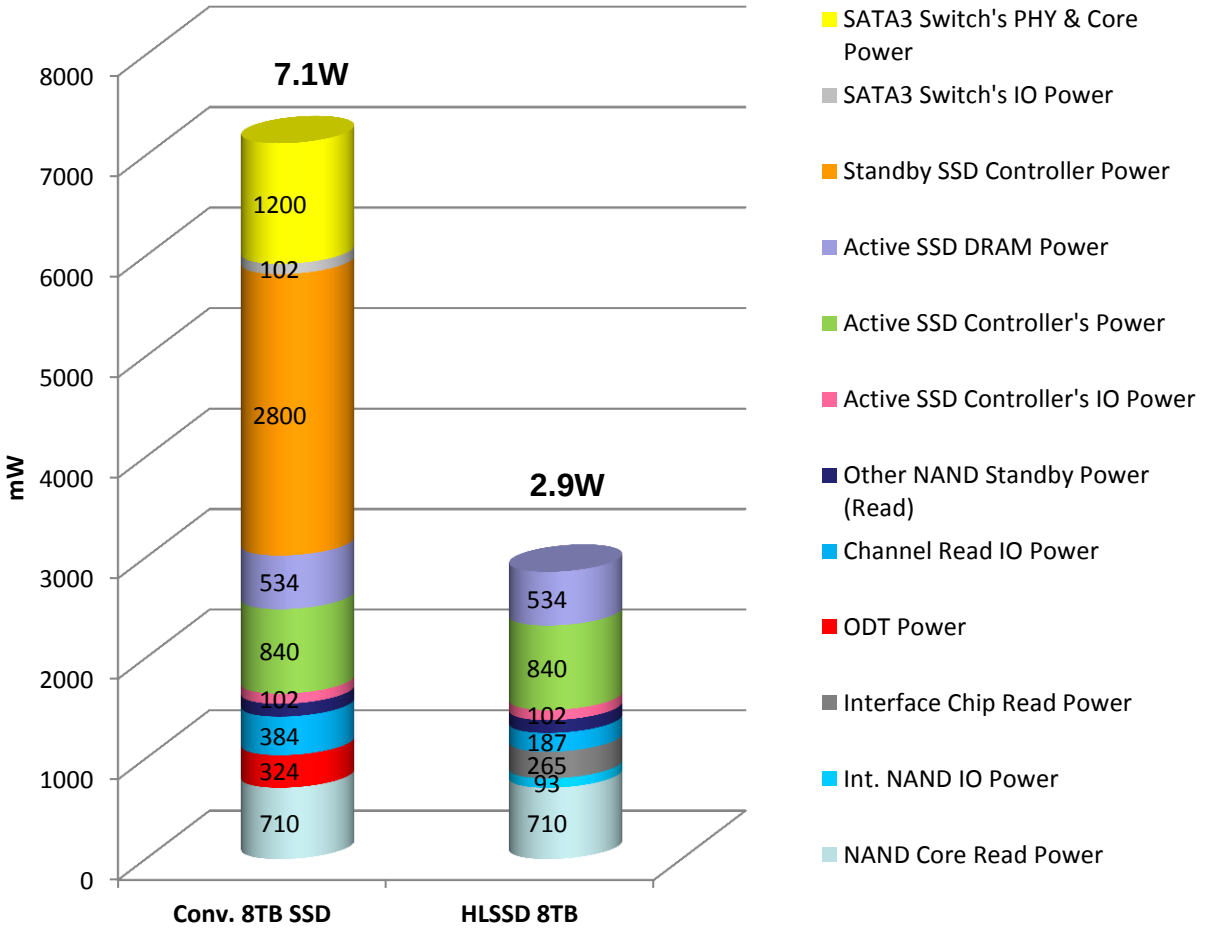

* Source: 2012 Flash Memory Summit, C. Lee and MOSAID's Internal Data

8TB SSD Write Power

(SATA3 BW Saturation)

- HLSSD consumes 40% less power
- Conventional SSD consumes significant SATA channels power
- NAND's core program operation dominates overall power

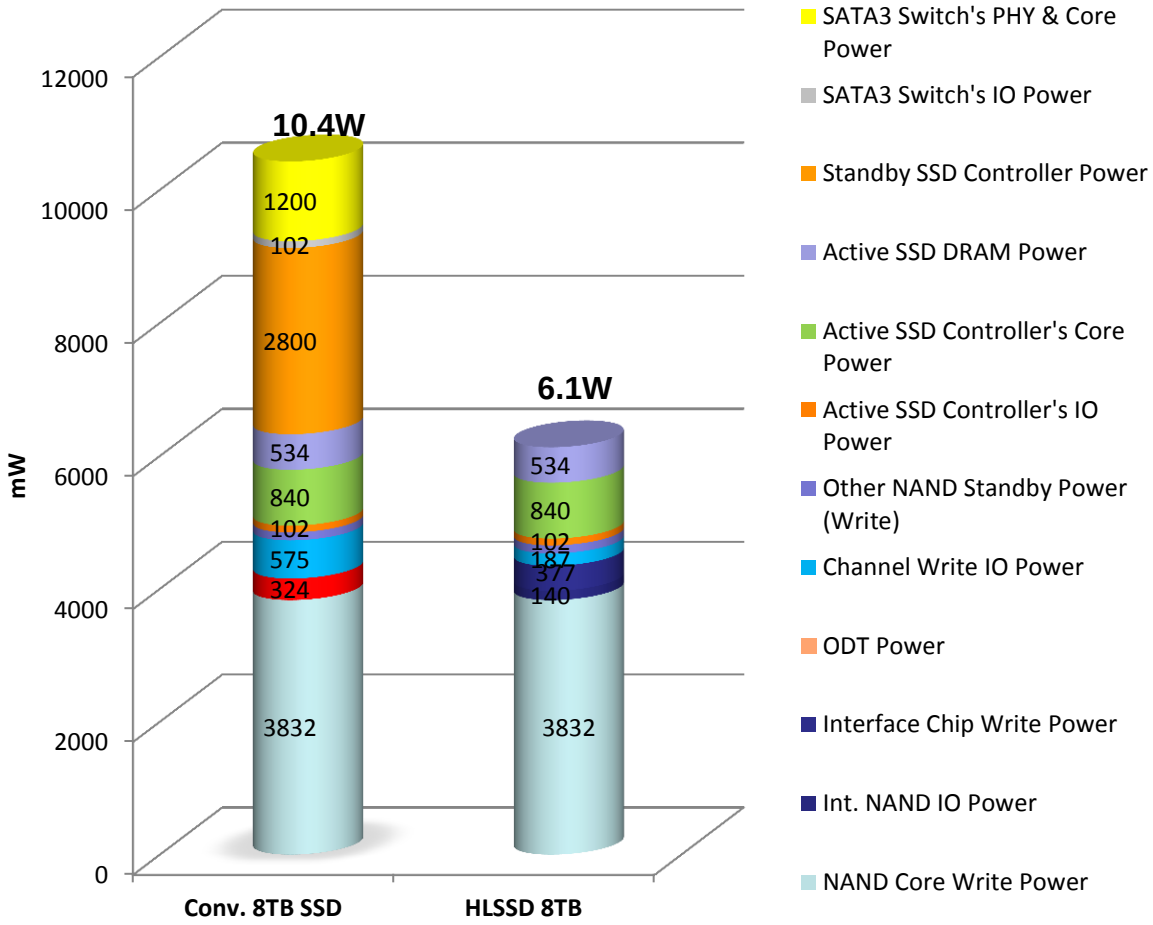

* Source: 2012 Flash Memory Summit, C. Lee and MOSAID's Internal Data

Summary

- Highly scalable HLNAND technology delivers affordable Terabyte-class SSDs
- Native PCIe controller based HLSSDs deliver the best $/IOPS, IOPS/W and IOPS/GB
- Provide fast and large storage pool to minimize the number of storage tiers – full flash array or no need for more than two “tiers” of storage

See HLSSD at Booth 714