Hybrid Storage Performance Characteristics

Kirill Malkin

CTO, Starboard Storage Systems

Who is Starboard Storage?

- Designer and innovator of Hybrid Storage
- Innovative Multiprotocol Accelerated Storage Technology (MAST) architecture
- Unified and multiprotocol (block-based architecture with file capability)
- Solid state used as an Accelerator
- Hard disks managed in a single storage pool accommodating any size and type of drive.

What is Hybrid Storage?

- Storage systems that store data using more than one storage device technology
  - Typically HDD and SSD
- Key benefit is achieving two goals:
  - Get the performance of flash
  - Get the capacity at the cost of hard disks
    - and peace of mind
- Hybrid storage can be block, file or unified
  - Also a mix of access protocols is supported

Why would Hybrid Work?

- Real-life applications are subject to locality of reference phenomenon
- Locality is relatively small (5% of capacity) and can be tracked
  - In theory, only a small fraction of storage needs to be high-performance

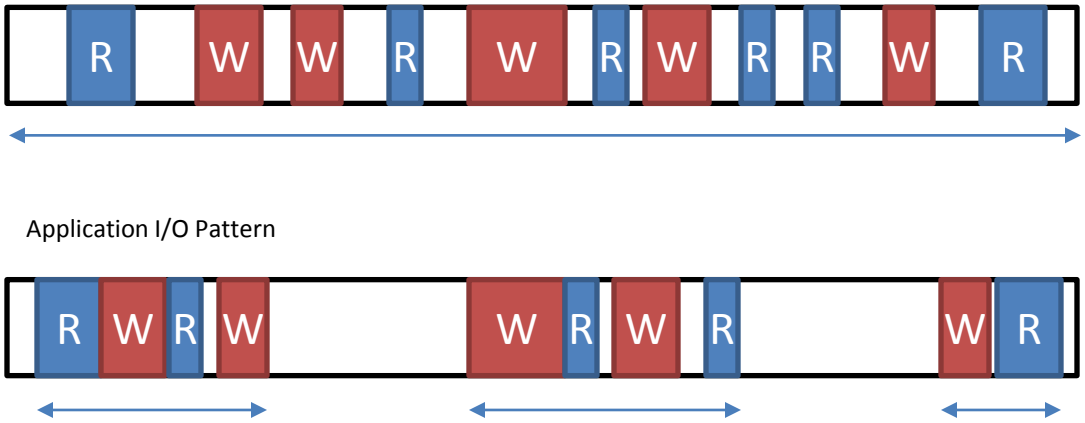

Spatial Locality

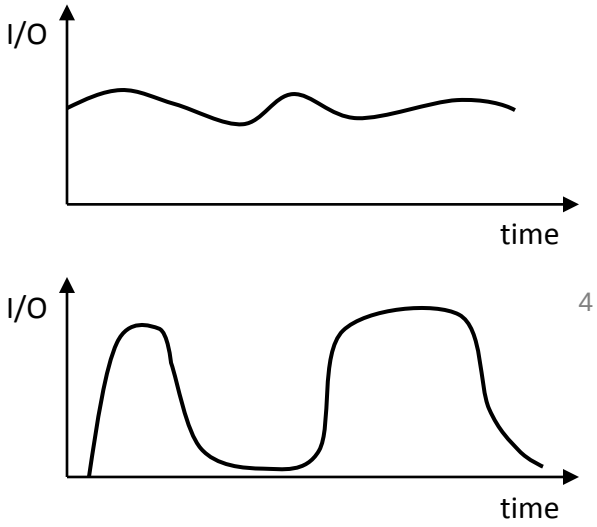

Temporal Locality

How Does it Work?

- Store on flash to boost performance when needed
- Transparent to workload
- Several hybrid storage architectures exist:
  - Solid state as a tier
  - Solid state as a cache
  - Solid state as accelerator

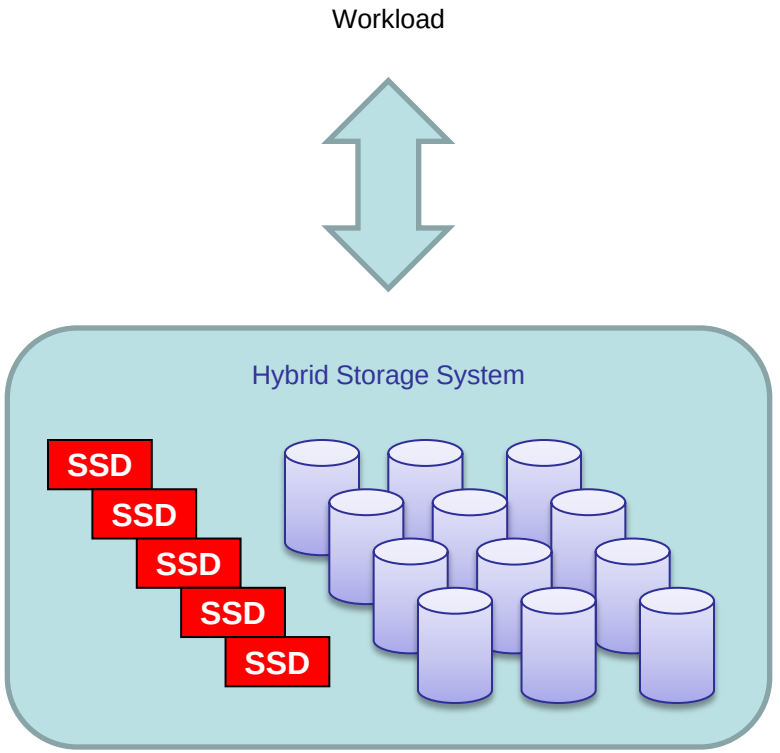

Solid State as a Tier

- Data segments statically placed to a storage tier
- Segments moved depending on policy & access patterns
- Pros:
  - Predictable
  - Simple
- Cons:
  - Wasteful, expensive
  - Inertia, overhead
  - Tough decisions

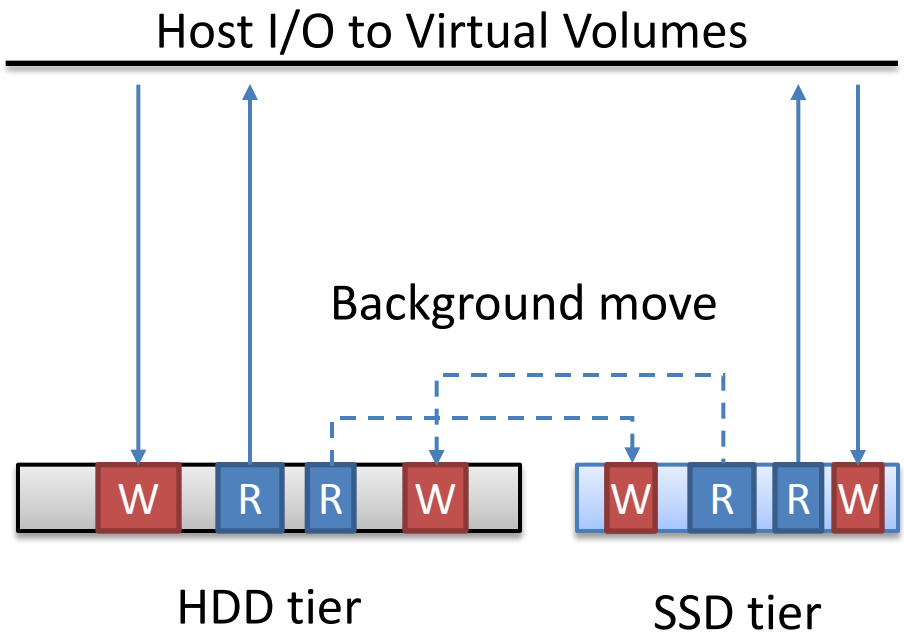

Solid State as a Cache

- Data segments are copied to SSD cache
- Segments evicted depending on policy & access patterns
- Pros:
  - Holistic
  - Simple
- Cons:
  - Unpredictable
  - Small, lost on reset

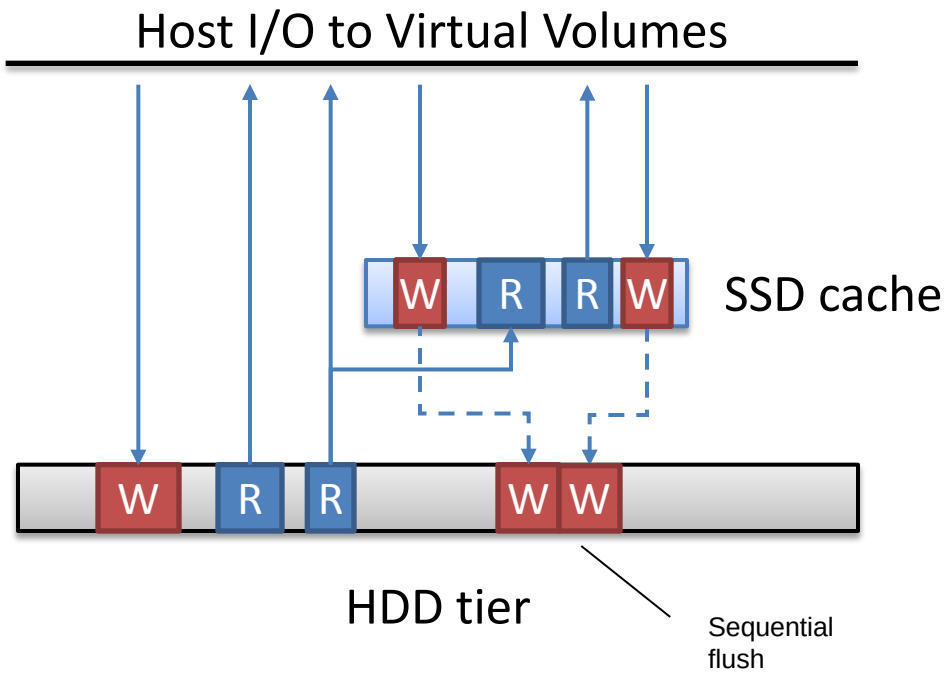

Solid State as a Cache - 2

- All data segments are compressed & copied to cheap SSD cache
- Redundant HDD tier is used as sequential log
- Reads mostly from SSD, rarely from HDD
- Pros:
  - Good performance initially
  - Inexpensive SSDs
- Cons:
  - Compressible workloads only
  - Requires HDD defrag

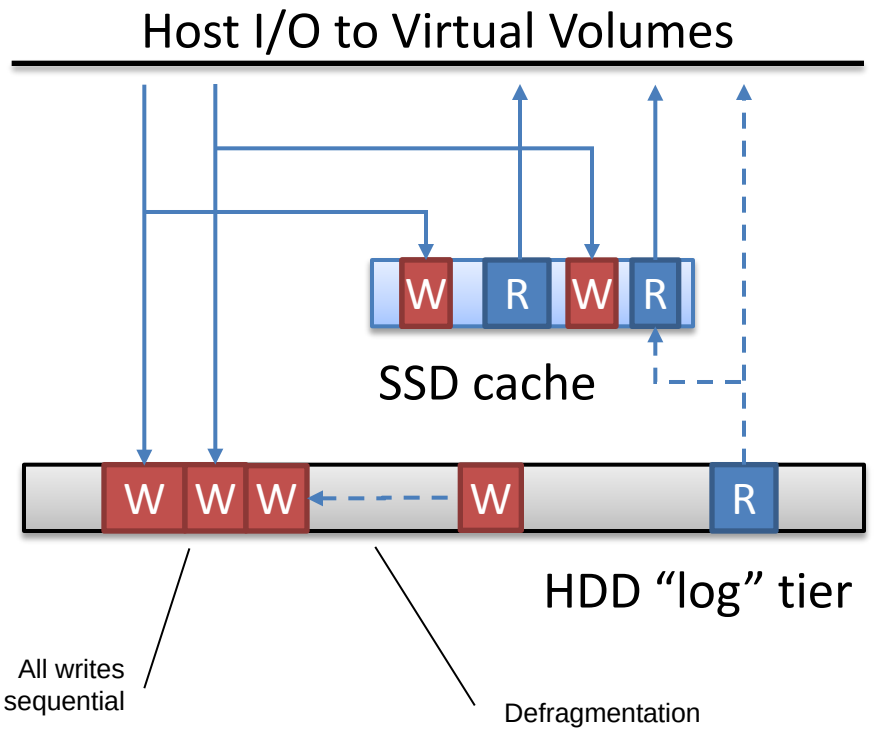

Solid State as Accelerator

- Data segments are copied or stored to multiple SSD tiers and to HDD tier
- Separate read & write channels
- Segments flushed or evicted depending on access patterns
- Pros:
  - Highly optimized
  - Mixed workloads
- Cons:
  - More complex

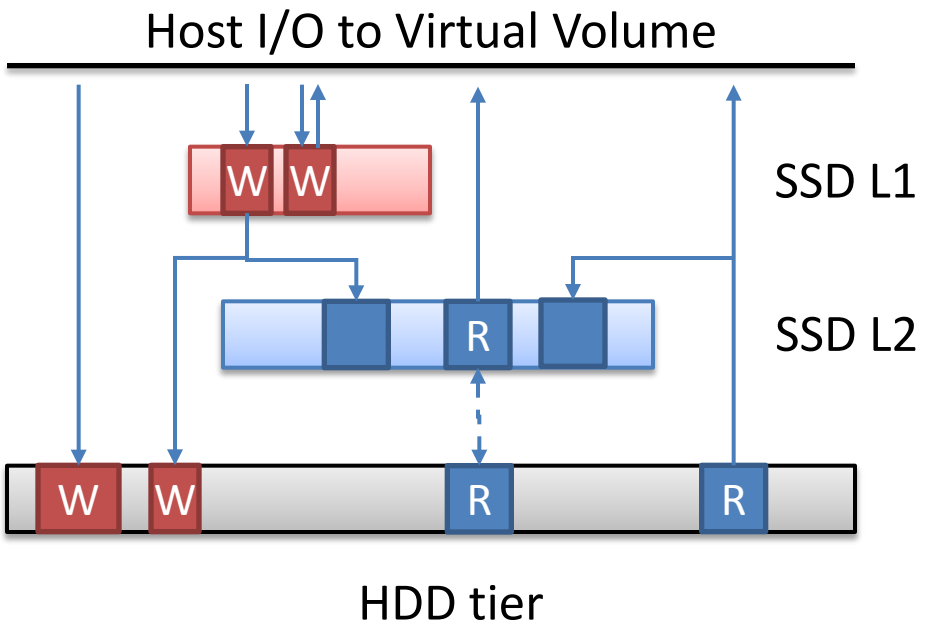

Hybrid Performance Optimization

- Data Reduction
  - Compression
    - Less written to SSD, less amplification
    - Most databases compress 2-4 times, some tests compress up to 25 times
    - HDD is less critical to compress though helpful
  - No-dupe
    - Zero copy snapshots & clones
    - Pointer-based EXTENDED COPY implementation
    - Thin provisioning & UNMAP
  - Dedupe
    - In-line & offline
    - Generally desired, may have performance impact

Performance Optimization - 2

- Metadata performance is key
  - Resides in memory, journaled to solid state
  - B+Tree designs recommended
- DRAM as compressed read cache
  - 100GB DRAM with compression means 200-400GB effective cache capacity
- Write-optimized flushes to HDD tier
- Read-optimized data placement into HDD tier
  - Eliminates need for defragmentation

STARBOARD
STORAGE SYSTEMS

HDD Tier Options

- Traditional hardware or software RAID
  - Group management complexities
  - Long rebuilds, spare management
  - Dedicated drives, lost capacity
  - Typically requires mirrored NVRAM for HA
- Dynamic disk pooling and data redundancy
  - Easy to manage & scale capacity
  - Metadata can be placed on SSD
  - Multiple concurrent redundancy levels
  - Accelerated architectures enable better data placement
  - Many tasks automated

Hybrid Architecture Example

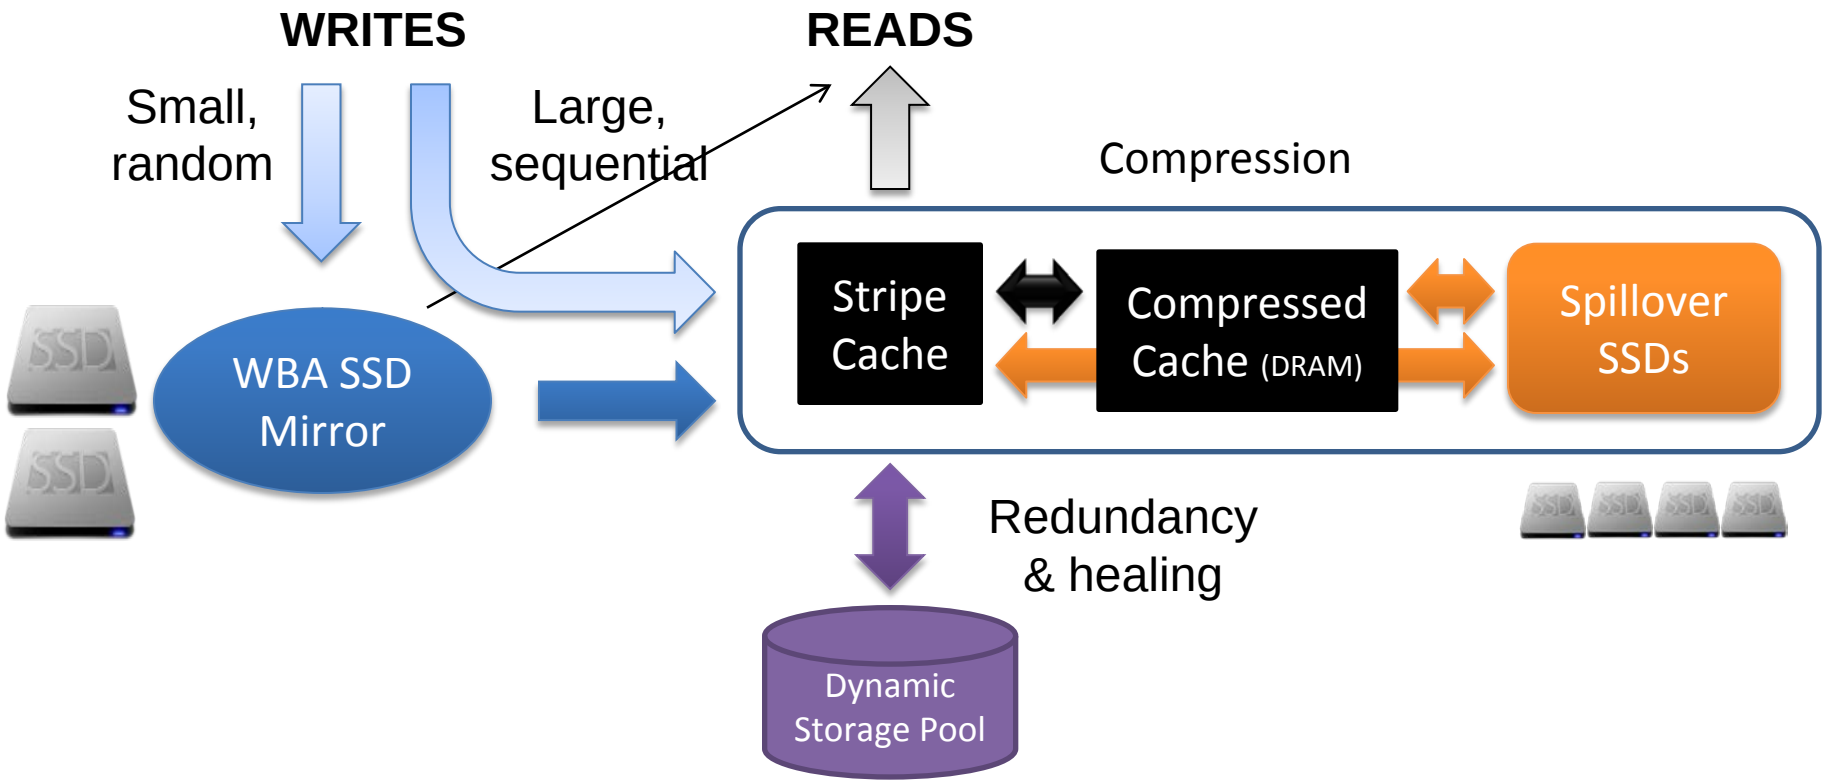

Hybrid Challenges

- Performance cliffs
  - Occur when cumulative capacity of reference exceeds high performance tiers
- Metadata growth
  - Huge pointer-based reference tables can hog memory and delay restart & failovers
- Testing
  - Generic random tools are misleading
  - Real-life (i.e. with locality of reference) workloads needed to demonstrate advantages of hybrid

Exchange Performance

Microsoft Exchange 2010 with 20,000 mailboxes

	Compellent Storage Center	NetApp FAS3220	Starboard AC4500
		 512GB Flash Cache	 2 x 200GB Write, 1TB Read Flash
Trans per sec	2632	3757	4010
Disks	60 x 15K 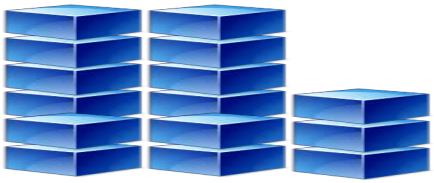	60 x 10K 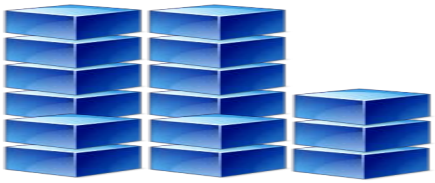	12 x 7200 RPM 

SPECsfs2008 (NFS)

	NetApp FAS3140	Hitachi NAS3080	Starboard AC-4500
op/sec	40,109	40,688	27,478
ART, ms	2.59	3.05	0.81
HDDs	80 x 15K	224 x 15K	28 x 15K
SSDs			2x200GB 800GB

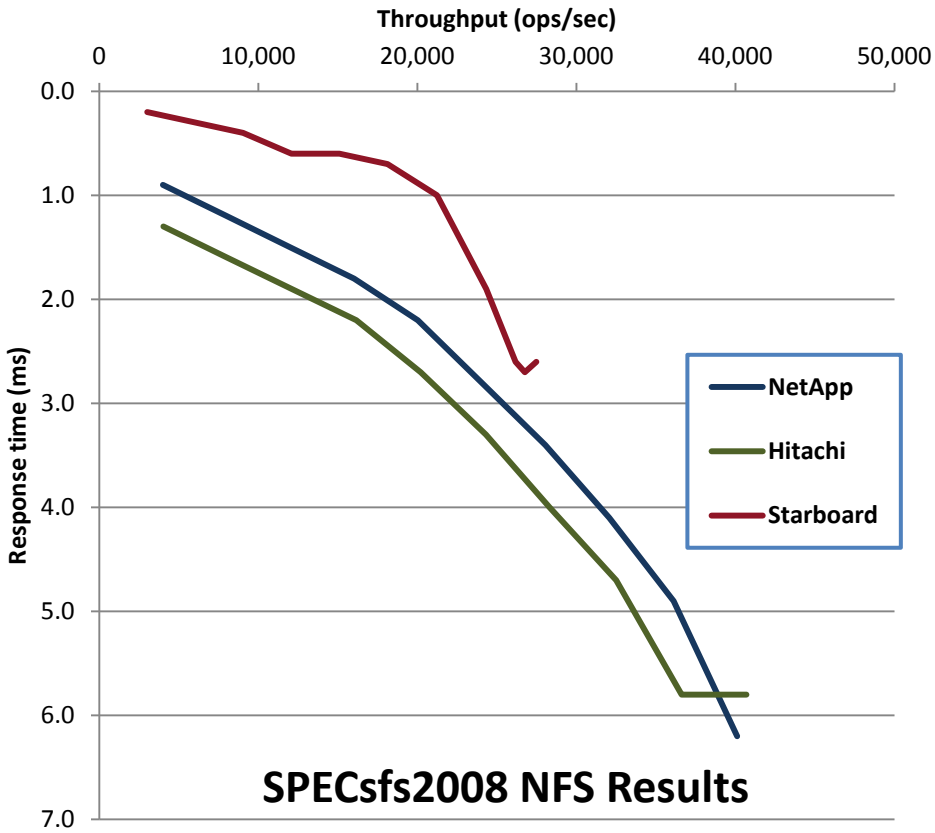

IOmeter Test Diagram

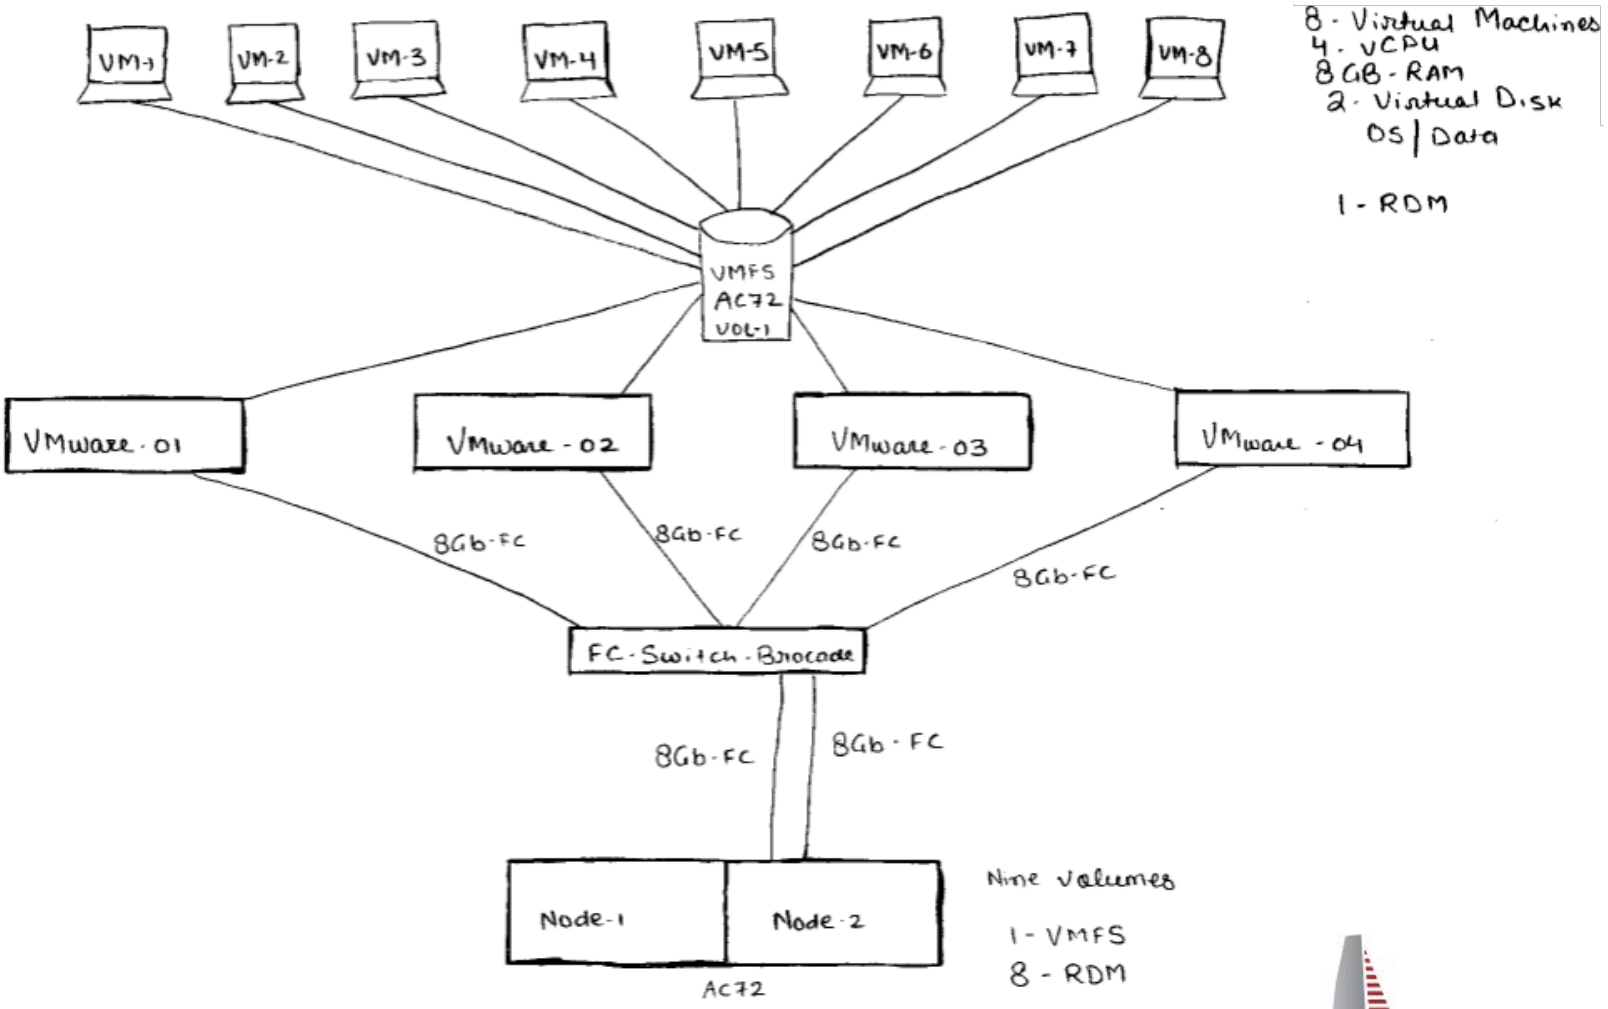

Cached Reads Performance

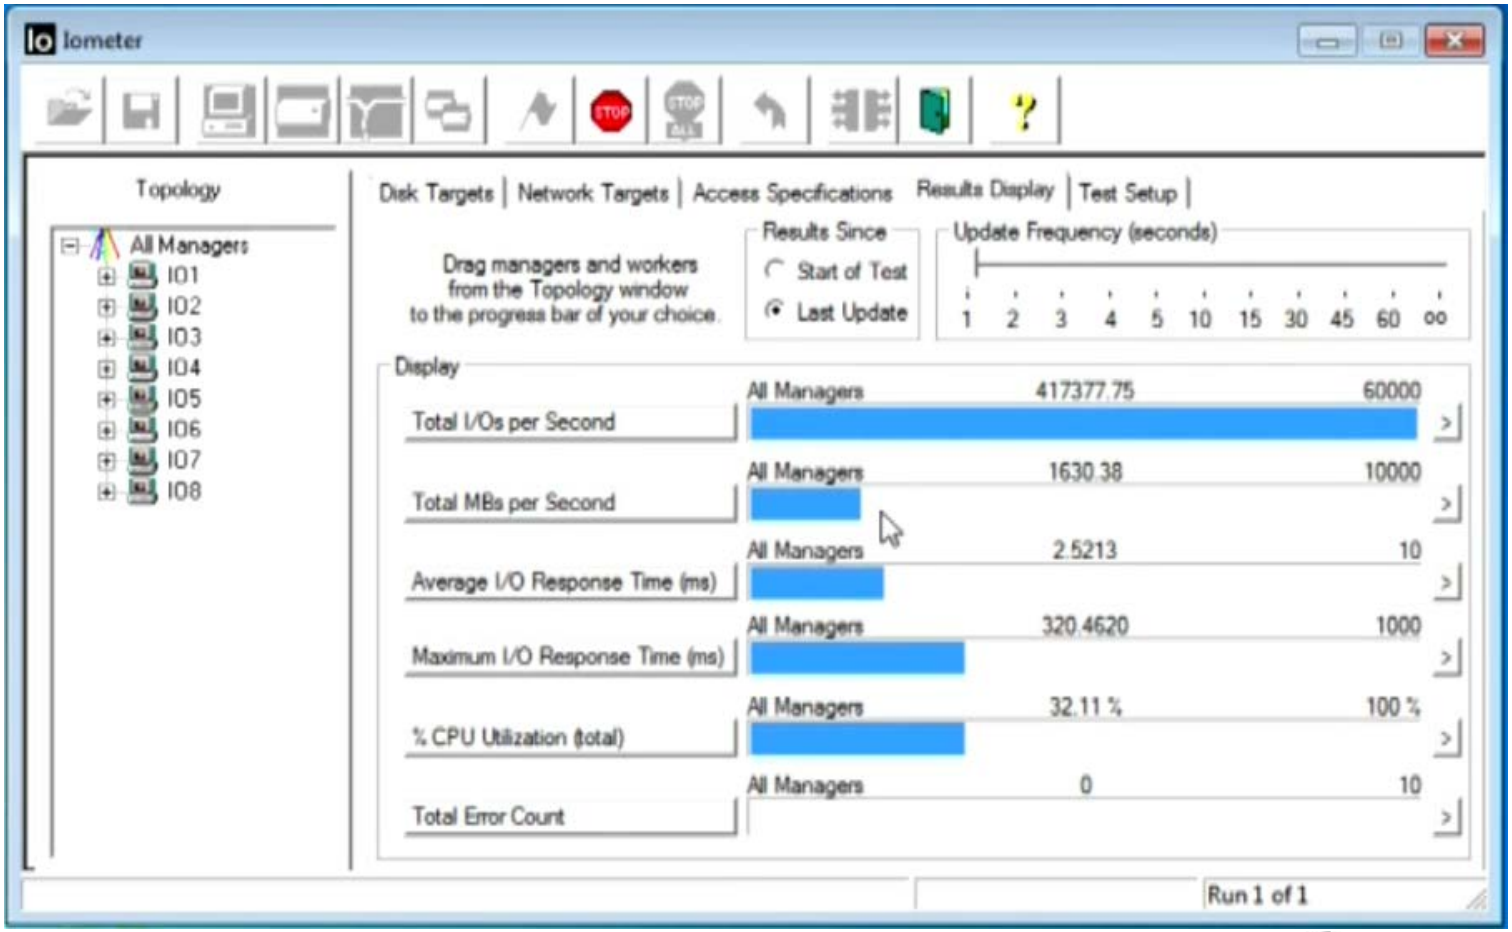

Flash Memory SUMMIT

Q & A

THANK YOU!

Visit us on Booth 911 in the Exhibit Hall

STARBOARD STORAGE SYSTEMS