



Flash Controller Architecture for All Flash Arrays

Andy Walls

Distinguished Engineer, IBM Systems and
Technology Group

CTO, Flash Systems and Technology

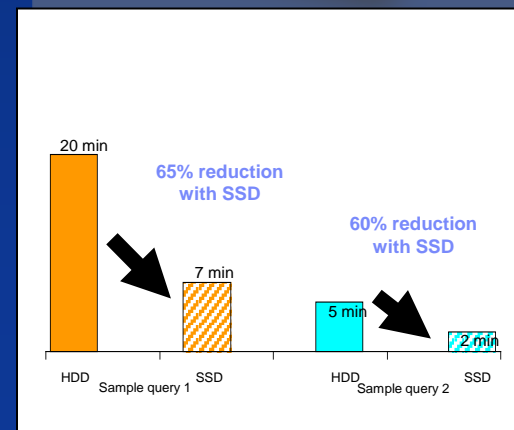



Once upon a Time,

A Disruption happened,

Solid State Drives in HDD form factors broke onto the scene!

And data centers saw that it was very good indeed!





SSDs Initiated changes to the industry

- SSDs plugged in where HDDs had and immediately gave lower response time
- Each SSD had an order of magnitude better Performance, and \$/IOPs and \$/W.
- PCIE Card formfactors provided reduced latency as well
- Infrastructure started to change
 - RAID Controllers
 - HBAs
 - Protocols (NVMe, SCSIe)
 - Software (Still largely to come)

HDD Form Factor SSDs

- Interchangeable with HDDs
 - Enabling performance and latency improvements in enterprise storage systems
 - Combine with Automatic caching/tiering software to make hybrid storage with optimized cost/performance.
- Gives ultimate flexibility to a storage system.
- Will generally be constrained to use HDD infrastructure



Big Data & Cloud Increasing IT Pressures

Big Data

- Businesses are challenged to create more insight and value out of more data
- Four dimensions of Big Data applications to be addressed: Volume, Velocity, Variety, Veracity
- IT infrastructures must not only be faster, but smarter



Cloud

- Businesses want to maximize infrastructure efficiency to become more agile
- Goal: standardize and consolidate applications on shared infrastructures
- IT infrastructures must be more scalable and reliable than ever before



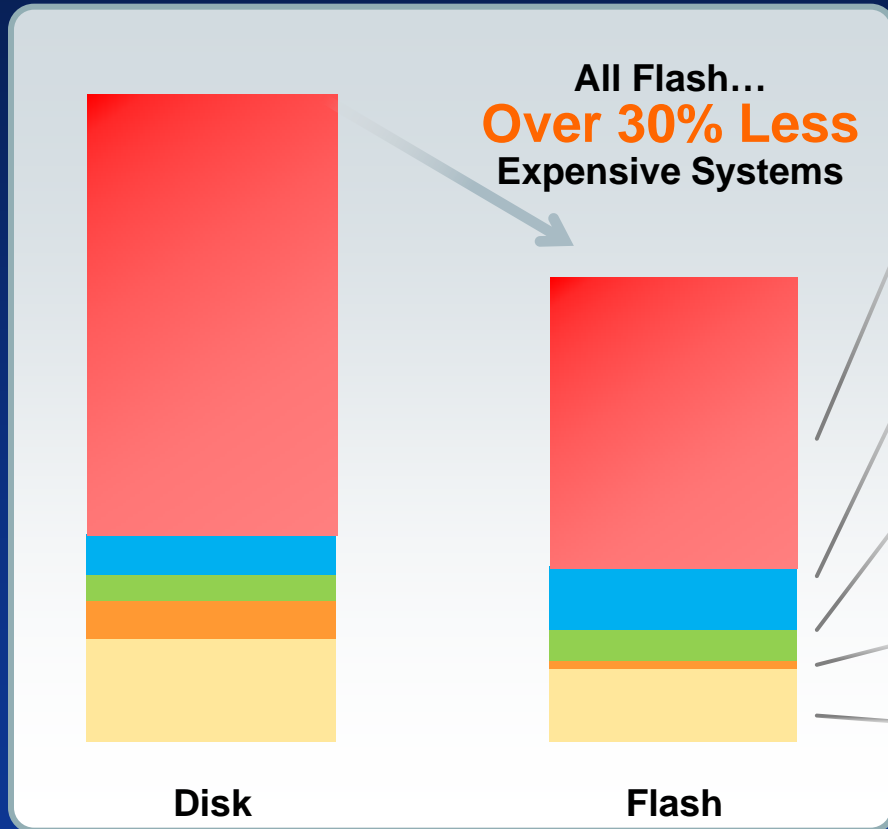
Emerging applications demand smarter, faster, scalable, & more reliable enterprise infrastructures



All Flash Arrays

- Several possible advantages
 - Consistently low and just plain consistent latencies
 - Consistently high IOPs
 - Able to deal with dynamic workloads
 - Environmental savings
 - Can save software license costs
- Homogeneity is often simpler to manage from client point of view

Counter Intuitive – Flash can be Cheaper



- **38% Lower Software License Costs**
 - Fewer cores, lower maintenance
 - Database, infrastructure SW...
- **Higher Storage Utilization**
 - As much as 50%
 - Lower maintenance
 - Simplified management
- **17% Fewer Servers**
 - Fewer cores
 - Fewer network connections
 - Lower maintenance
- **74% Lower Environmental Costs**
 - Power / cooling, floor space
- **35% Lower Operational Support Costs**
 - Server / storage administration

Source: Wikibon, March 2013



All Flash Arrays can produce additional savings over arrays of SSDs.

- All Flash Systems can be designed differently
 - RAID which is geared toward Flash
 - Error recovery, performance recovery unique for Flash
 - Help Endurance at all levels of the system
 - Owning the entire data path gives more opportunities for latency reduction.
- Optimizing the above can yield additional cost savings.



Flash Controller comparison

Standalone SSD

- Single controller Ops and Throughput
- Must comply with standard protocols
- Each Drive responsible for emergency power off.
- Each drive only sees its data.
- Each drive responsible for Endurance

Flash Controller in All Flash Array

- Aggregate Ops and Throughput for system
- If both ends of wire owned then benefits can be derived.
- Emergency power source can be shared.
- Can be a sharing of data for optimization.
- Endurance responsibility becomes shared.





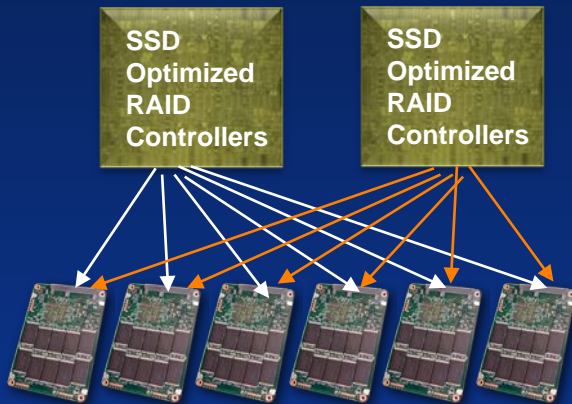
A Lake of storage vs. Pools of storage

Each Flash controller has visibility over its relatively small Flash domain

RAID Controller somewhat limited by standard interfaces.

Flash controller can be more integrated into overall system

RAID Controllers, management, flash controllers can be designed to work together.





Flash Controller for an All Flash Array Performance

- In an all flash array, the flash can be one large pool of Flash.
 - You can optimize either for capacity or performance
 - Controllers do not have to work autonomously.
 - Load balancing and hot spot detection
 - Data efficiency
 - High Access detection
- Protocols can be used which reduce overhead and increase parallelism.
- Potential to reduce



Endurance and Recovery

- If persistent data is stored on Flash – it must have no single points of failure or repair.
- Therefore, with standalone SSDs, the redundancy is above.
 - Software mirroring
 - RAID
- Opportunities for optimization in all flash arrays.
(Reducing overall storage efficiency)
- System level write amplification can be addressed
- Other tricks can be applied to help endurance over a larger pool of flash.



What Next?

- SSDs absolutely have a very important place
- All flash arrays can offer a single pool of Flash.
 - Optimize performance and Capacity
 - Tier between types of flash within the array
 - Work together for reduction of write amplification
 - Provide additional opportunities for endurance enhancements
 - Lowest latency
 - Consistent performance