


Solid State Architectures in the Modern Data Center

Andy Walls
IBM Distinguished Engineer




SSDs Superior in Price Performance – HDDs Superior \$/GB

\$/IOP Compare between Enterprise HDDs and SSDs



- To get the same number of IOPs, SSDs are cheaper than HDDs
 - Especially considering infrastructure
- But capacity is much less
- RAID controllers to allow for all the IOPs and will continue to appear

Average Price per GB Comparison, Enterprise Flash SSD and Enterprise HDD, 2010-2015



- HDDs and SSDs \$/GB continuing to decline
- SSDs driven by
 - Geometry shrinks
 - More MLC usage
- HDDs driven by
 - Continuing TPI reductions
 - More usage of Slower and cheaper HDDs

Ref: IDC Worldwide 2010 – 2015 Solid State Drive Forecast and Analysis by Jeff Janukowicz

Problem and Vision statement

Problem:

Data center is facing the challenge to manage the explosion of information volume. The exponential growth of data implies higher cost on management, energy, floor space and capacity. Manual process to manage storage tiering cannot meet the demand anymore.

Solution: ***EASY Tier***

Enable customers to manage the explosion of information volume by removing the complexity of the data placement management in storage systems.

- Transparent to applications: automated and concurrent data migration
- Fine grain data movement: adaptive sub-disk level data migration
- Maximize performance, cost, energy, storage utilization
- Automated storage configuring and provisioning
- Alignment with business, application and system policy
- Consistent with thin technology
- Integrated solution with storage management Software

Client Challenges with SSD

- *High performance (IO/Sec) but high cost (\$/Gbyte)*

- SSDs are more expensive than magnetic disks (Although prices dropping)
- The challenge – use them for the data that really needs the speed in IOPs/Sec
- ***Do not make management more complex***
- ***Easy Tier makes the most use of the SSDs (\$/IOP) for hot data and places colder data on HDDs (\$/GB)***



Ways SSDs Can Accelerate Performance

- IOPs play to get more Transactions per second or to achieve similar transaction requirements with less cost.

- Higher Read Throughput or Scan Rate
 - Business Analytics
 - Data Warehousing

- Reduced and consistent Average Latency
 - Real time transaction processing with guaranteed response time
 - Requires reduced and manageable max latency

- Reduced batch windows to meet Service Level Agreements

SSD End Usage is Evolving

External Storage



- MLC adoption increasing
- Easy Tier increasing dramatically
 - Multi Level tiers including SSDs, 15K and Nearline HDDs
 - Feature improvements will further increase adoption
- Some pure SSD boxes
 - Fairly uniform distribution of data
 - Response time critical
- With optimized external storage like DS8800, latencies even in the SAN are good.

Direct Attach


- MLC adoption increasing
- Easy Tier will play here in many ways
 - Within networks controlled by a RAID Controller pair
 - Within the server across disparate SSDs directly attached via PCIE or SAS.
- DAS does offer some reductions in latency and powerful use models
- Workload Optimized Systems

Intelligent data placement EasyTier


Performance Improvement with minimal management



Automated Data Relocation





DB2 Brokerage Workload with Easy Tier: More Transactions with Less Response Time



Weighted Average RT (ms)				
Trade Activity	Lookup	Order	Update	Overall Transaction
DS8800 Base	9666	290	13332	1169
Easy Tier2	3441	198	5883	493
Benefit (%)	64.40%	31.72%	55.87%	57.83%

Flexible Configuration of Performance and Capacity with Auto-Rebalance



Before and After is Representative :
Data Movement Accelerated to Reduce Test Time

Conclusions

- Enterprise SSDs are starting to show significant growth
 - MLC is an important reason for that growth
- Clients need more than the underlying technology.
 - They need firmware and software to allow for tiering
 - They need applications which exploit it
- SSDs cheaper in terms of \$/IOPs
 - Expect more infrastructure which exposes those IOPs to the client
- Various applications will exploit different performance attributes
- We are likely to see radical changes to the datacenter as more clients go to pure SSD solutions
 - DAS is very important for SSDs
 - But sharing and clustering of those DAS solutions is essential
 - But the SAN will also stay relevant especially with Easy Tier