Case for flash storage – how it can benefit your enterprise

Dejan Kocic

Hitachi Data Systems (HDS)

Case for flash storage - how it can benefit your enterprise

What is flash storage?

- SLC (single level cell)
- MLC (multi level cell)

and variations on each of the technologies

Differences between SLC and MLC flash:

	Access time	Density	Endurance	Operating environment
SLC	~ 0.3 ms	1 bit/cell	100,000+ PE cycles	Industrial
MLC	~ 0.7 ms	2+bit/cell	5,000-10,000 PE cycles	Commercial

Case for flash storage - how it can benefit your enterprise

- How flash storage compares to other types of storage

Technology	DRAM	Flash	HDD
Access time	**nanoseconds**	**microseconds**	**milliseconds**
Scale		1000 nanoseconds	1000 microseconds
Cost/GB	Very high	High	Moderately High

Case for flash storage - how it can benefit your enterprise

Deterrents for using flash

- Initially technology was not very stable
- Perception was that data can be more easily lost or corrupted as compared to traditional RAID protected HDDs
- **High Cost per GB**
- Not being able to justify purchasing of flash drives due to high cost/GB and unreliability concerns

Case for flash storage - how it can benefit your enterprise

What are the use cases for flash?

- Data can be classified as structured data or unstructured data

Data access can be categorized as:

- Sequential reads
- Sequential writes
- Random reads
- Random writes
- Combination of the above

Case for flash storage - how it can benefit your enterprise

What are the use cases for flash?

- Sequential reads are relatively fast due to prefetching and read-ahead algorithms used
- Sequential and random writes are always written to cache on storage arrays, then committed to disk
- Random reads are a problem with no good solution

Case for flash storage - how it can benefit your enterprise

What are the use cases for flash?

- Random reads involve high number of seek operations to position disk head at the specific place to be able to read or write data which takes more time than any other part of the disk read or write process
- Eliminating seek operations or reducing them would improve random read performance drastically

Case for flash storage - how it can benefit your enterprise

- What are the use cases for flash?
- Applications creating random reads:
  - Databases
    - Online transactional processing (OLTP)
    - Online analytical processing (OLAP)
    - Applications using databases (Data warehousing solutions, Microsoft SharePoint & similar apps)
  - Virtualization (VMware, Hyper-V, VDI solutions)
  - Different types of metadata
  - Operations involving large amounts of small files

Case for flash storage - how it can benefit your enterprise

- What are the use cases for flash?
  - Random read I/O is common performance problem for block and file (NAS) based storage
  - Several different solutions are available

Case for flash storage - how it can benefit your enterprise

Solutions

- Host attached flash (DAS)
- Network caching using flash
- Storage data tiering
- All flash storage array

Case for flash storage - how it can benefit your enterprise

- Host attached flash (DAS)
  - Easy to implement
  - Relatively cheap
  - Single point of failure
  - Doesn’t scale very fell
  - Limited to one host

Case for flash storage - how it can benefit your enterprise

Network caching using flash

- Solutions are available for IP and FC networks (SAN)
- Solutions for IP networks can cache data or metadata or both
- Solutions for IP networks support CIFS (SMB) and NFS protocols
- In both cases, IP network and FC caching reduces load on primary storage array
- IP and FC network caching can support multiple hosts, filers and storage arrays

Case for flash storage - how it can benefit your enterprise

Storage data tiering

- Since in many cases only 10%-20% of allocated and used storage is actively used storage, it makes sense to move rarely used or inactive storage to lower (cheaper) tiers of storage
- Storage tiers usually consist of 3 tiers: flash, 15K rpm or 10K rpm SAS drives and Near Line SAS (SATA) drives
- Data tiering allows users to gain better performance with relatively small amount of flash storage

All flash storage array

- When performance is the primary consideration, all flash storage array is probably the best solution
- Offers benefits of traditional storage arrays in terms of robustness and built in redundancies while being scalable and reliable
- Can provide millions of IOPS
- Usually expensive in the terms of cost/GB

Case for flash storage - how it can benefit your enterprise

- Solutions summary:

Solutions:	Supports more than one device	Single point of failure	File storage support	Block storage support	Scalable	Metadata only caching
Host attached flash (DAS)	NO	YES	YES	YES	NO	YES
Network caching	YES	NO	YES	YES	YES	YES
Storage data tiering	YES	NO	NO	YES	YES	NO
All flash storage array	YES	NO	NO	YES	YES	NO

Case for flash storage - how it can benefit your enterprise

- Solutions summary:
  - Depending on the workload and use case, one solution may be more appropriate than the others
  - It is good to know your data set, performance characteristics, application behavior… to be able to get most benefit from using flash storage
  - Know what problem you are trying to solve (e.g. performance on a host level, network level or storage level…)

Case for flash storage - how it can benefit your enterprise

True cost of storage

- Cost per GB is usually not the best measure of storage cost unless you need storage purely to meet capacity requirement
- In the most cases capacity requirement goes together with a performance requirement
- Often users purchase excess capacity to meet performance requirement
- This leads to increased usage of space in the datacenter, using more power, more cooling, higher maintenance cost, likely more manpower to manage it, etc.

Case for flash storage - how it can benefit your enterprise

True cost of storage

- Perception that HDD based storage solution will always be cheaper than flash storage solution may not always be true
- It may take several cabinets of HDD based storage to meet performance requirement while same performance requirement may be met with 8U or 12U all flash solution
- Difference in annual energy cost alone is huge, especially outside of USA where cost of electricity is much higher

Case for flash storage - how it can benefit your enterprise

- True cost of storage
  - Cost of extra space may be significant, especially in colocation facilities
  - Cost per I/O operation (IOPS) provides more accurate price of storage for use cases where performance is important consideration
  - Cost/IOPS usually comes out favorably or it is cheaper than traditional storage when total cost of solution is considered to include power, space, cooling, support cost, manpower cost, capacity licenses, etc.

Case for flash storage - how it can benefit your enterprise

Business justification for buying flash storage

- "It is fast" is usually not very good business justification
- Be prepared to describe the challenges in your environment and back it up with the data
- Describe how you think that flash solution can solve the problem and in the process also save some money to the company
- Flash storage may allow company to do things that were not possible to do with traditional storage, potentially creating new business opportunities

Case for flash storage - how it can benefit your enterprise

- Business justification for buying flash storage
  - In cases where timing is important and there are process dependencies, flash storage may increase availability of the environment, which may translate to more revenue
  - Being able to do more in less time may directly translate to increased revenue (e.g. high frequency stock trading)
  - Use information on previous slides to build a business case for flash storage

Case for flash storage - how it can benefit your enterprise

Conclusion

- Use of flash can remove bottlenecks in the environment
- It can speed up existing processes and allow business to do more in the same amount of time
- Caching using flash can extend life of legacy storage
- Storage data tiering can provide benefits of flash without having to buy large amounts of flash storage
- All flash arrays can provide millions of IOPS and sub-millisecond latency in small fraction of space as compared to traditional storage

Questions?