Case for flash storage – How it can benefit your enterprise

Dejan Kocic

Netapp

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

Case for flash storage – How it can benefit your enterprise

This session will be interesting for technologists and the decision makers within an organization. Flash storage is becoming important factor for both consumers and in the enterprises and this session will explore good use cases for flash storage, benefits provided for those use cases and how to justify the purchase of flash storage.

The benefits of flash storage are many, but higher cost per GB of flash storage as compared to traditional storage is making it hard to justify buying the flash storage. What many storage customers don't necessarily realize is that price per megabyte or gigabyte of capacity is an outdated model and it does not represent true cost of storage. Many enterprise consumers of storage keep buying more and more hard drives, not to satisfy capacity requirement, but to satisfy performance requirement. When total cost of ownership is taken into consideration, flash storage becomes very competitive with much smaller footprint, cooling and power requirements. There are also certain workloads, which can benefit tremendously from using flash storage and those will be discussed in this presentation.

What is flash storage?

Flash storage

- Uses NAND flash memory (non-volatile) technology to store data. Primary technologies are:
  - SLC (single level cell)
  - MLC (multi (dual) level cell)
  - TLC (triple level cell)
  - 3D Flash (dual or triple level cell, currently up to 48 layers)

  and variations on each of the technologies

Differences between flash technologies:

	Random read	Density	Endurance (PE cycles)	Operating environment
SLC	25 µs	1 bit/cell	100,000+	Industrial
(e)MLC	50 µs	2 bits/cell	1,000-30,000	Commercial
TLC	100 µs	3 bits/cell	800-1,000	Consumer / Commercial
3D Flash	50 µs – 100 µs	2-3 bits/cell	MLC & TLC levels	Consumer / Commercial

What is flash storage?

- Program Erase (PE) Cycles – is sequence of writing, erasing and re-writing NAND flash. Each PE cycle is wearing out layer of silicon oxide (gate, only about 10nm thick)

	SLC	(e)MLC	TLC
Bits/cell	1	2	3
PE cycles	100,000	30,000	800 - 1000
Program time	300 µs	900 µs	900-1400 µs
Erase time	2 ms	3 ms	4-5 ms

Storage controller, the way data is placed in NAND flash and write amplification are also related to wearing of flash cells

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
Power consumption	Low	~ 20% of HDDs	High
Heat generation	Low	Low	High

Deterrents for using flash

- Initially technology was not very stable which is much improved now due to better controllers and algorithms used
- Perception was that data can be more easily lost or corrupted as compared to traditional RAID protected HDDs
- **High Cost per GB (which is rapidly declining)**
- Not being able to justify purchasing of flash drives due to high cost/GB and unreliability concerns

What are the use cases for flash?

- Data can be primarily classified as structured data or unstructured data, both types can benefit from flash storage

  Data access can be categorized as:

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

What are the use cases for flash? (Cont.)

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

What are the use cases for flash? (Cont.)

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

What are the use cases for flash? (Cont.)

- Applications creating random reads:
  - Databases
    - Online transactional processing (OLTP)
    - Online analytical processing (OLAP)
    - Applications using databases (Data warehousing solutions, Content Management Systems & similar apps)
  - Virtualization (Hypervisors & VDI solutions)
  - Different types of metadata
  - Operations involving large amounts of small files

What are the use cases for flash? (Cont.)

- Random read I/O is common performance problem for block (SAN) and file (NAS) based storage
- Several different solutions are available

Solutions

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

Host attached flash (DAS)

- Easy to implement
- Relatively cheap
- Single point of failure
- Doesn‘t scale very well
- Limited to one host

Network Caching Using Flash

- Solutions are available for IP (NAS) and FC (SAN) networks
- 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

Storage Data Tiering

- Since in most cases only 10%-20% of allocated 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, 10K rpm SAS drives and Near Line SAS or 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

Solutions Summary

Solutions:	Supports more than one host	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

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, workload… 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 (application) level, network level or storage level…)

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
- Users often 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.

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 just in annual energy cost alone is huge, especially outside of USA where cost of electricity is much higher

True Cost of Storage

- Cost of extra space may be significant, especially in colocation facilities
- Cost per I/O per second (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.

True Cost of Storage

- Example – requirement is 1,000,000 IOPS and 100TB of storage
  - Traditional 15K rpm 600GB SAS HDDs provide about 200 IOPS/drive so about 5000 drives are needed for performance requirement, then when 20% for parity and spare drives are added total is about 6000 drives or about 12-18 full cabinets (depending if full size or small factor drives were used (3.5” or 2.5”)) – 5000x 600GB drives is about 3PB of storage
  - Flash storage, 800GB enterprise drive can provide 2000* or more IOPS so about 500 drives are needed for performance requirement, then when 20% for parity and spare drives are added total is about 600 drives which can fit in about 2 and half cabinets

* Very conservative number which takes into consideration performance limit of the storage controller, SAS bus saturation, etc.

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
- Prepare the data and make proposal 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 and potentially creating new business opportunities

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 purchasing flash storage

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 a small fraction of space as compared to traditional disk based storage solutions

Questions?

Attribution & Feedback

The SNIA Education Committee thanks the following individuals for their contributions to this Tutorial.

Authorship History

Name/Date of Original Author here:
Dejan Kocic / 7/14/2014
Updates:
Dejan Kocic / 8/1/2014
Dejan Kocic / 6/23/2015

Additional Contributors

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***