SNIA Tutorial 1
A CASE FOR FLASH STORAGE – HOW TO CHOOSE FLASH STORAGE FOR YOUR APPLICATIONS

Dejan Kocic, NetApp

Flash Memory Storage 2018

Welcome to
SNIA Education Afternoon
at Flash Memory Summit 2018

Agenda

SNIA Global Education

1:00 pm – 1:50 pm	SNIA Tutorial 1 *A Case for Flash Storage* Dejan Kocic, NetApp
1:50 pm – 2:45 pm	SNIA Tutorial 2 *What if Programming and Networking Had a Storage Baby Pod?* John Kim, Mellanox Technologies and J Metz, Cisco Systems
2:45 pm – 3:00 pm	Break
3:00 pm – 3:50 pm	SNIA Tutorial 3 *Buffers, Queues, and Caches* John Kim, Mellanox Technologies and J Metz, Cisco Systems
4:00 pm – 5:00 pm	SNIA Tutorial 4 *Birds-of-a-Feather – Persistent Memory Futures* Jeff Chang, SNIA Persistent Memory and NVDIMM SIG Co-Chair

SNIA™ Global Education

170
industry leading organizations

2,500
active contributing members

50,000
IT end users & storage pros worldwide

Join SNIA at These Upcoming Events

SDC discount registration cards in FMS bags & at **SNIA booth 820**

Complimentary registration now open at **snia.org/pm-summit**

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

A Case for Flash Storage –

How to Choose Flash Storage for Your Applications

- Flash storage is becoming important factor for both consumers and the enterprises.
- There are many types of flash storage and it may be challenging to determine what type of flash is right for a specific application or a workload.
- This paper will examine different types of flash and their characteristics, suitability for different workloads, and how end users can get the most from the specific type of flash suitable for their needs.

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 128 layers) and variations on each of the technologies (Z-SSD/Z-NAND)
  - Persistent memory – 3D Xpoint – dual layer currently

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
3D Xpoint	7 µs	1 bit/cell	1,000,000+	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, does not apply to 3D Xpoint)

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

Storage controller, the way data is placed in NAND flash and write amplification are also related to wearing of flash cells. For 3D Xpoint durability, chalcogenide material used for storage is the key factor in the case of 3D Xpoint technology.

How Flash Storage Compares to Other Types of Storage

Technology	DRAM	3D Xpoint	Flash	HDD
Access time	Nanoseconds	7 µs	Microseconds	Milliseconds
Scale		1000 nanoseconds	1000 nanoseconds	1000 microseconds
Cost/GB	Very high	Cheaper than DRAM more expensive than Flash	High	Moderately High
Power consumption	Low	Higher than flash	~ 20% of HDDs	High
Heat generation	Low	Low to moderate	Low	High

Access Time Comparison

SNIA Global Education

	Nanoseconds (ns)	Microseconds (µs)	Milliseconds (ms)	If L1 Access is 1 second
L1 Cache Reference	0.5			1 sec
L2 Cache Reference	7			14 secs
DRAM Access	200			6 mins, 40 secs
Intel Octane 3D XPoint	7,000	7		3 hours, 53 mins, 20 secs
Micron 9100 NVMe PCIe SSD Write	30,000	30		16 hours, 40 mins
Mangstor NX NVMeF Array Write	30,000	30		16 hours, 40 mins
DSSD D5 NVMeF Array	100,000	100		2 days, 7 hours, 33 mins, 20 secs
Mangstor NX NVMeF Array Read	110,000	110		2 days, 13 hours, 6 mins, 40 secs
NVMe PCIe SSD Read	110,000	110		2 days, 13 hours, 6 mins, 40 secs
Micron 9100 NVMe PCIe SSD Read	120,000	120		2 days, 18 hours,m40 mins
Disk Seek	10,000,000	10,000	10	7 months, 10 days, 11 hours, 33 mins, 20 secs
DAS Disk Access	100,000,000	100,000	100	6 years, 4 months, 19 hours, 33 mins, 20 secs
SAN Array Access	200,000,000	200,000	200	9 years, 6 months, 2 days, 17 hours, 20 mins

What Are the Use Cases for Flash?

- Primary objective of flash storage is to reduce the latency
- Even with HDDs, 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
- For spinning disks, random reads are a problem with no good solution

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 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 (access latency) or reducing them would improve random read performance drastically

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, 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?

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

Solutions

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

Host Attached Flash (DAS or Persistent Memory)

- Easy to implement
- Cost dependent on solution (flash or PM)
- Single point of failure
- Doesn’t scale very well
- Limited to one host – scale out solutions exist with specialized software
- OS support needed (in the case of Persisent Memory)
- If very low latency is needed, this is the best solution

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
- With process of flash falling steadily, this is no longer an often used solution

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
- Traditionally storage filers usually consist of 3 tiers:
  - Flash
  - 10K rpm SAS drives and Near Line SAS or SATA drives
  - 3D Xpoint can be used as a DRAM cache
- 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 terms of cost/GB compared to HDD, but because of technologies like deduplication and compression price of effective storage is becoming comparable to cost of HDD storage

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 solutions 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 the most benefit using flash storage
- Know what problem you are trying to solve (e.g., performance on a host (application) level, network level, or storage level. . .)

Solutions Summary

- 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 training)

What Type of Flash to Choose

- Depending on the latency needs, endurance, cost and type of access, appropriate flash solution can be selected
- Below table summarizes characteristics of different types of flash and exact solution will depend on number of factors: application type, data access type, data access protocol, desired latency, scale out capabilities, availability…

	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
3D Xpoint	7 µs	1 bit/cell	1,000,000+	Consumer / Commercial

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?

Thank you!

Visit snia.org/education