



NETWORKING
STORAGE

Everything You Wanted To Know About Storage But Were Too Proud To Ask: Part Taupe The Memory Pod

**Live Webcast
May 16, 2019
10:00 am PT**

Today's Presenters



Alan Bumgarner
Intel



Alex McDonald
NetApp



John Kim
SNIA NSF Chair
Mellanox

SNIA-at-a-Glance



185

industry leading
organizations



2,000

active contributing
members



50,000

IT end users & storage
pros worldwide

Learn more: snia.org/technical



SNIA Legal Notice

- ◆ The material contained in this presentation 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.
- ◆ 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.

- Memory in a Storage Context
- Background
 - ◆ A brief history
 - ◆ Why storage & memory have been seen as different
 - ◆ A sense of scale
- Persistent Memory
 - ◆ Programming PM (aka Non Volatile Memory)
 - ◆ The Hardware: NVDIMMs
 - ◆ OS and Application

A Little History - Memory

➤ Mechanical memories

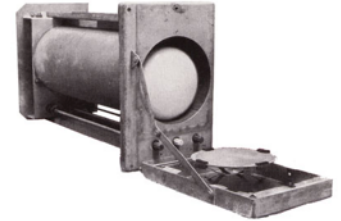
- ◆ Relays

➤ Williams–Kilburn tube (1946–47)

- ◆ Spots on phosphor

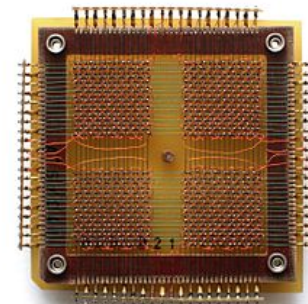
➤ Delay line memory (1947)

- ◆ Acoustic delay line that used mercury
- ◆ Alan Turing suggested using neat gin as he claimed it had similar acoustic properties



A Little History - Storage

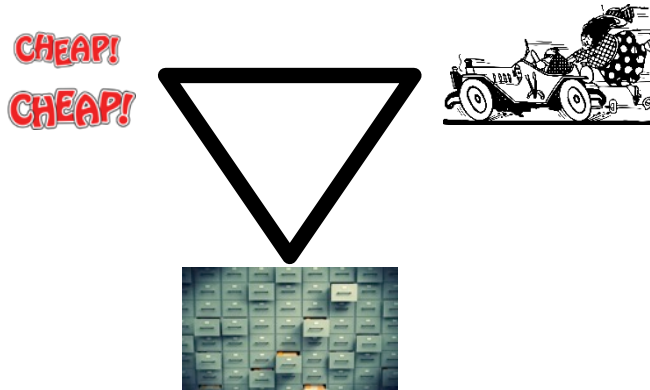
- Drum memory (1932)
- Magnetic core memory (1949)
 - ◆ Ferrite rings
 - ◆ “Thin film memory”
- Disks...
- And now ***persistent memory***



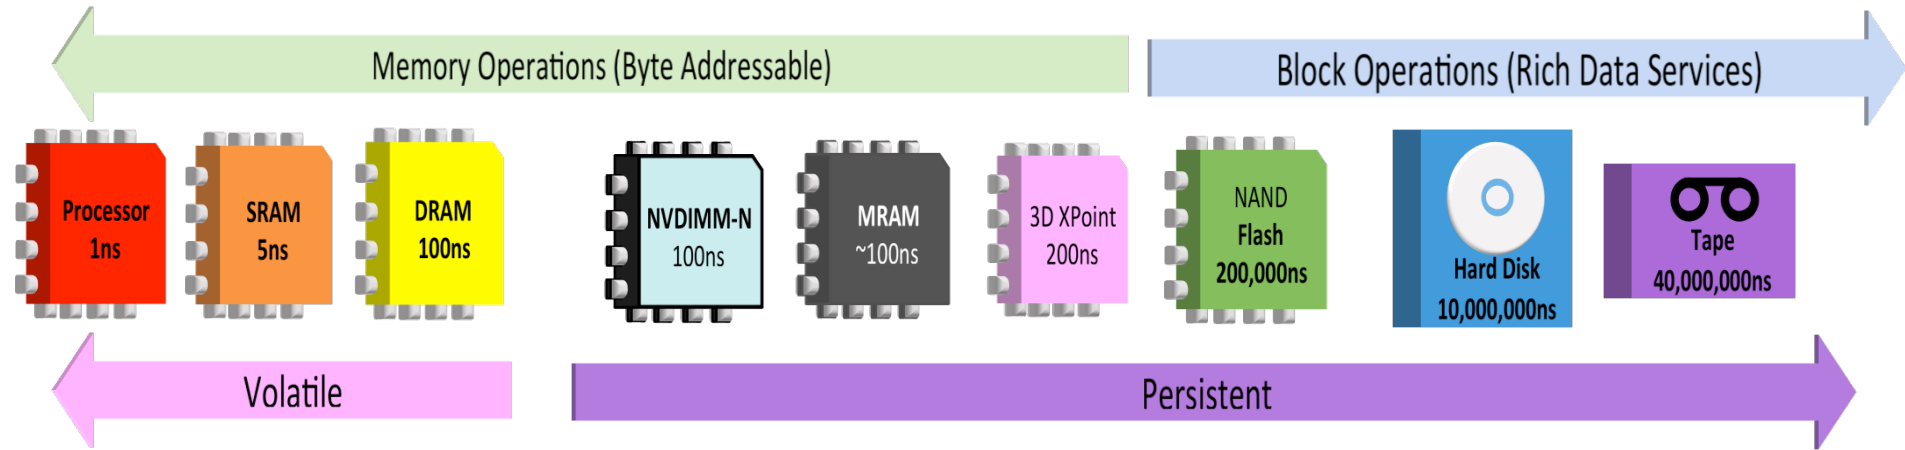
- Memory or storage?
 - ◆ Key attribute was not *persistance*, but *addressability*
- Driven by the available technologies
 - ◆ Cost, size, speed, persistence
- Programs deal with
 - ◆ Loads and stores for fine grain
 - ◆ Blocks of data for bulk
 - ◆ “Almost all programming can be viewed as an exercise in caching“, Terje Mathison

A Sense of Scale

- Size, speed and cost
- Classic “pick any two from three”



A Sense of Scale; Speed



➤ 1 ns = light travels approx. 30cm (1 foot)

Latency in Human Terms

➤ Memory Operation

- ◆ Getting an apple from the fridge (64B)

➤ Storage Operation

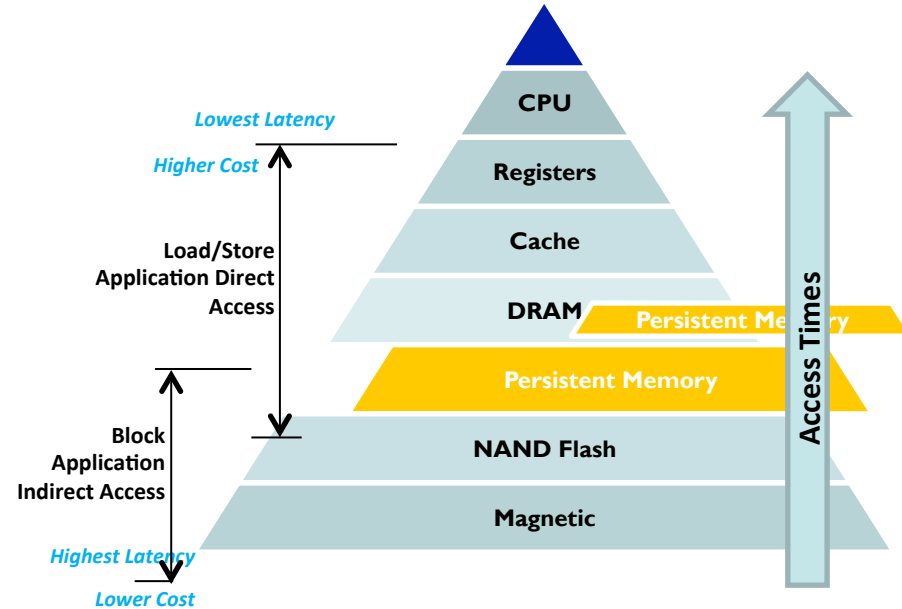
- ◆ Opening up an app
- ◆ Ordering from store
- ◆ Getting a box of apples shipped and delivered (4KB)

Time scales every developer should know.

Operation	Latency	In human terms
L1 Cache hit	1 ns	A blink of an eye (~20 ms)
L2 Cache hit	3 ns	Noticeable flicker
L3 Cache hit	10 – 20 ns	Time to say “A”
Main memory	70 – 100 ns	Time to say a ten word sentence
Signal down a 200m fibre cable	1 µsec	One slide (speaking quickly)
SSD access	5 – 25 µsec	Time to reheat a meal (3 mins)
HDD access	8 msec	Time to flight around the world. (1.8 days)
Network packet from Germany to the USA	45 msec	Waiting for a 7 working day delivery

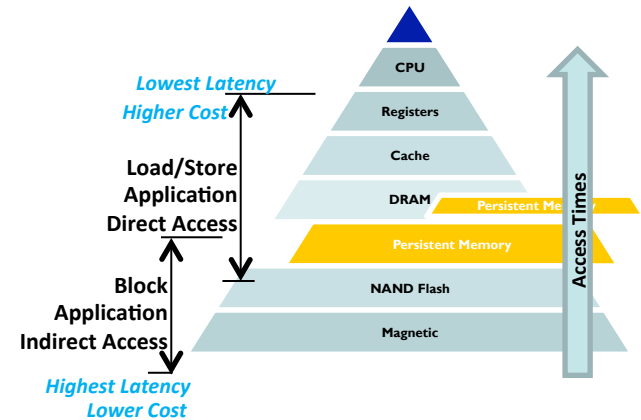
A Sense of Scale; Size & Cost

- At the top level, a few 1000s of bytes
- At the bottom, petabytes or more
- Each level represents a factor of approx 10^3
- Costs follow



A New Memory Paradigm

- Like memory (byte addressable) but like storage (persistent)
- A new paradigm
- “Almost all programming can be viewed as an exercise in caching“, Terje Mathison
 - ◆ Data used to move through the tiers from slow, big, cheap and persistent to fast, small, expensive and volatile
 - ◆ Hard boundary between the two, and a missing step
 - ◆ PM is a new tier that’s blending memory and storage



Persistent Memory (PM)

is a type of Non-Volatile Memory (NVM)

- **Disk-like non-volatile memory**
 - ◆ Persistent RAM disk
 - ◆ Appears as disk drives to applications
 - ◆ Accessed as traditional array of blocks
- **Memory-like non-volatile memory (PM)**
 - ◆ Appears as memory to applications
 - ◆ Applications store data directly in byte-addressable memory
 - ◆ No IO or even DMA is required

➤ Memory in a Storage Context

➤ Background

- ◆ A brief history
- ◆ Why storage & memory have been seen as different
- ◆ A sense of scale

➤ Persistent Memory

- ◆ Programming PM (aka Non Volatile Memory)
- ◆ The Hardware: NVDIMMs
- ◆ OS and Application

SNIA NVM Programming Model

- Version 1.2 approved by SNIA in June 2017
 - ◆ http://www.snia.org/tech_activities/standards/curr_standards/npm
- Expose new block and file features to applications
 - ◆ Atomicity capability and granularity
 - ◆ Thin provisioning management
- Use of memory mapped files for persistent memory
 - ◆ Existing abstraction that can act as a bridge
 - ◆ Limits the scope of application re-invention
 - ◆ Open source implementations available
- Programming Model, not API
 - ◆ Described in terms of attributes, actions and use cases
 - ◆ Implementations map actions and attributes to API's

NVMP TWG Work
continues on High
Availability Use
Cases and Practical
Implementations

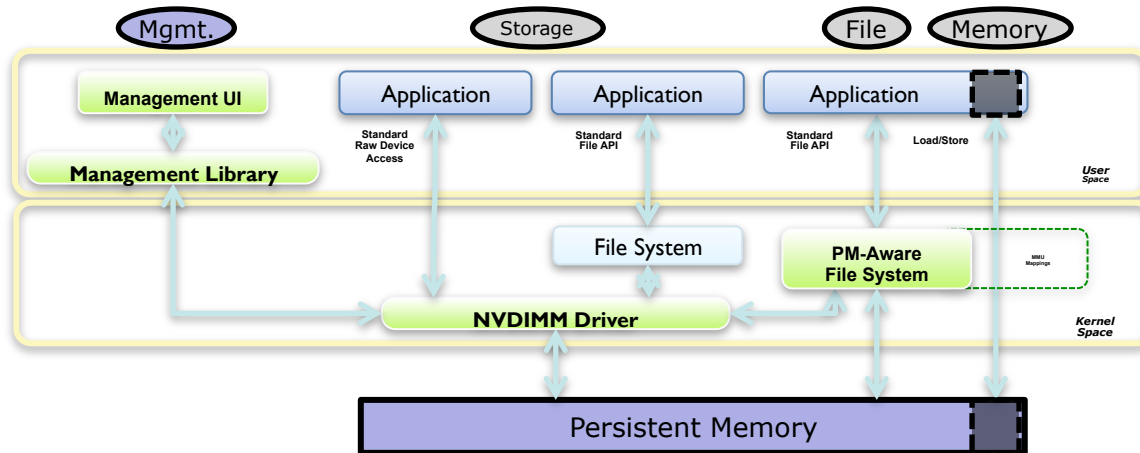
➤ Persistent memory...

- ◆ Allows load/store access like memory
- ◆ Is persistent like storage
- ◆ Exposed to applications using SNIA NVMP TWG model

➤ What isn't persistent memory:

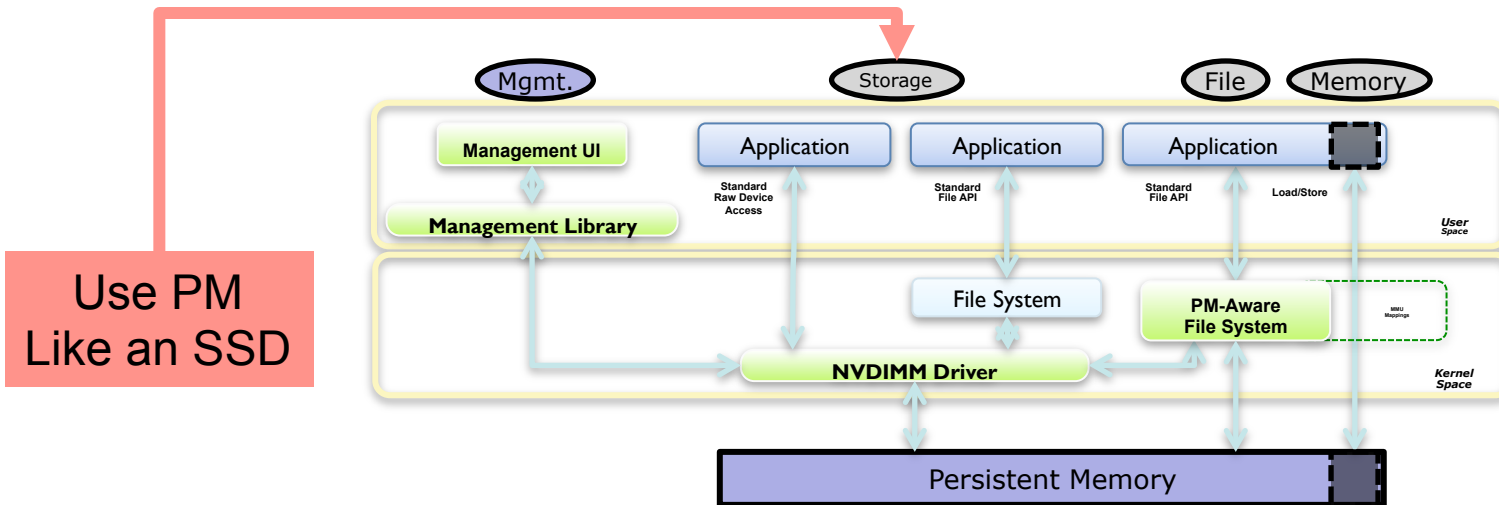
- ◆ Something that can only speak blocks (like a disk/SSD)
- ◆ Something that is too slow for load/store access
 - › SNIA TWG's language:
 - › Would reasonably stall the CPU waiting for a load to complete

- The programming model includes the storage APIs!



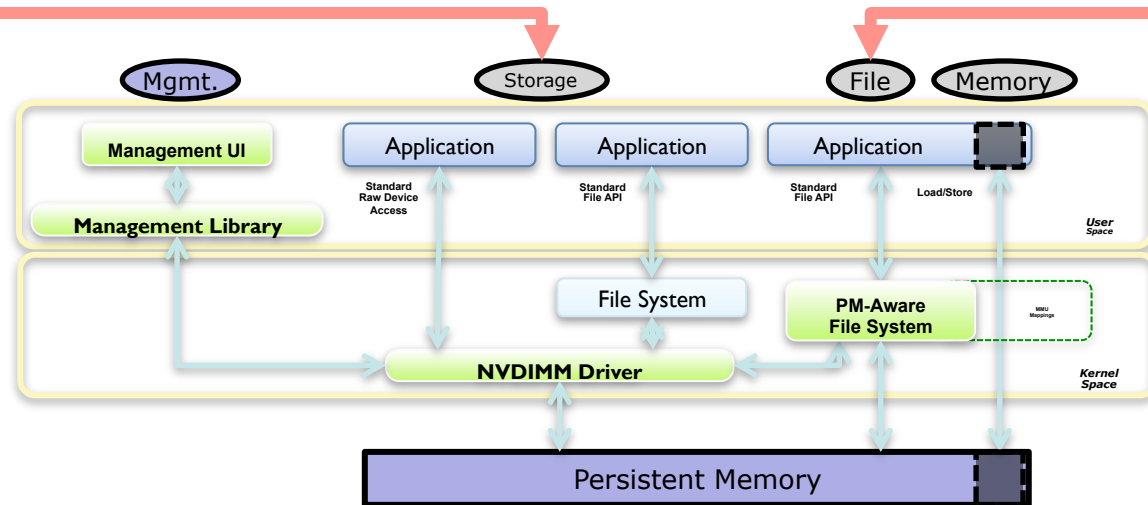
Often Forgotten: Storage Access

- The programming model includes the storage APIs!



Often Forgotten: DAX Access

- The programming model includes the storage APIs!



Use PM
Like an SSD

Use PM
Like an SSD
(no page cache)

“DAX”

No Application Modification

- **Using PM as a fast SSD**
 - ◆ Storage APIs work as expected
 - ◆ Memory-mapping files will page them into DRAM
- **Using PM as DAX**
 - ◆ Storage APIs work as expected
 - ◆ No paging (DAX stands for “Direct Access”)
- **Using PM as volatile capacity**
 - ◆ Just big main memory
 - ◆ Vendor-specific feature

NVDIMM Applications

- In-Memory Database: Journaling, reduced recovery time, Ex-large tables
- Traditional Database: Log acceleration by write combining and caching
- Enterprise Storage: Tiering, caching, write buffering and meta data storage
- Virtualization: Higher VM consolidation with greater memory density
- High-Performance Computing: Check point acceleration and/or elimination



- **Memory and Storage Differ by Access Model**
 - ◆ Speed, scale, cost
- **Persistent Memory Offers a New Solution**
 - ◆ Can be treated as memory or storage
- **PM Is Supported Today**
 - ◆ SNIA PM Programming Model
 - ◆ Support in Linux and Windows
 - ◆ Can use with or without application modifications

- This is our 10th “Too Proud To Ask” webcast in this series
- 9 More “Everything You Wanted To Know About Storage But Were Too Proud To Ask” on-demand at:
 - ◆ <https://www.snia.org/forums/nsf/knowledge/webcasts-topics>
- Next Live Webcast:
- Intro to Incast, Head of Line Blocking and Congestion Management
 - ◆ June 18, 2019, 10:00 am PT
 - ◆ Register at: <https://www.brighttalk.com/webcast/663/356343>

After This Webcast

- Please rate this webcast and provide us with feedback
- This webcast and a PDF of the slides will be posted to the SNIA Networking Storage Forum (NSF) website and available on-demand at www.snia.org/forums/nsf/knowledge/webcasts
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-NSF blog: sniansfblog.org
- Follow us on Twitter @SNIANSF

Thank You