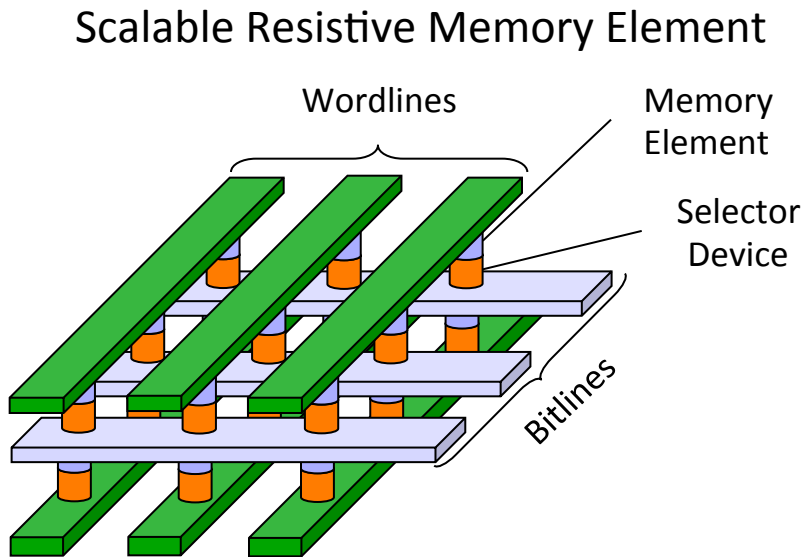


# Exposing Persistent Memory

# Next Generation Scalable NVM

## Resistive RAM NVM Options




Cross Point Array in Backend Layers  $\sim 4I^2$  Cell

Family	Defining Switching Characteristics
Phase Change Memory	Energy (heat) converts material between crystalline (conductive) and amorphous (resistive) <u>phases</u>
Magnetic Tunnel Junction (MTJ)	Switching of magnetic resistive layer by <u>spin-polarized electrons</u>
Electrochemical Cells (ECM)	Formation / dissolution of "nano-bridge" by <u>electrochemistry</u>
Binary Oxide Filament Cells	Reversible filament formation by <u>Oxidation-Reduction</u>
Interfacial Switching	<u>Oxygen vacancy drift</u> diffusion induced barrier modulation

**Many candidate next generation NVM technologies.  
Offer  $\sim 1000x$  speed-up over NAND, closer to DRAM**

# Exploiting Next Generation NVM

App to SSD IO Read Latency (QD=1, 4KB)



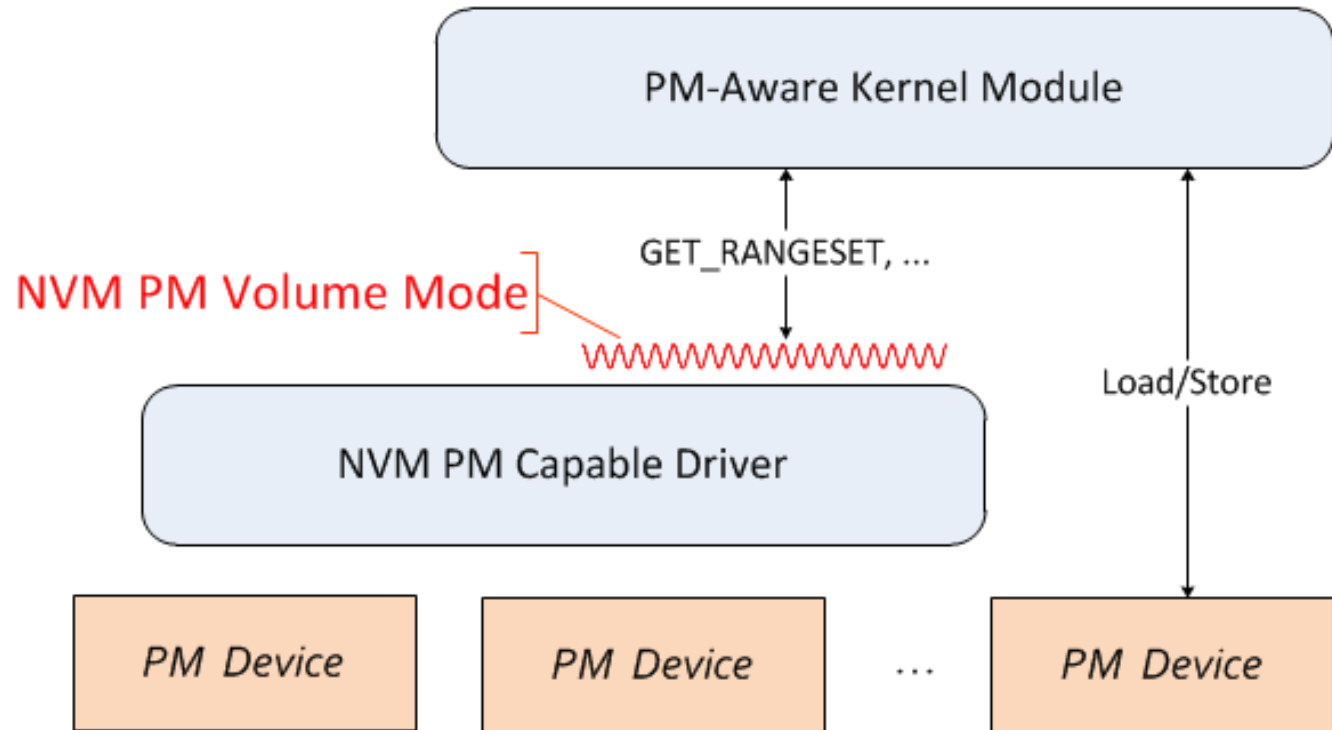
- With Next Generation NVM, the NVM is no longer the bottleneck
  - Need optimized platform storage interconnect
  - Need optimized software storage access methods

# Exposing PMem for Other Kernel Modules

- NVM Volumes
  - PM Capable
  - A list of physical ranges of NVM
- Operations
  - GET\_RANGESET
  - ...

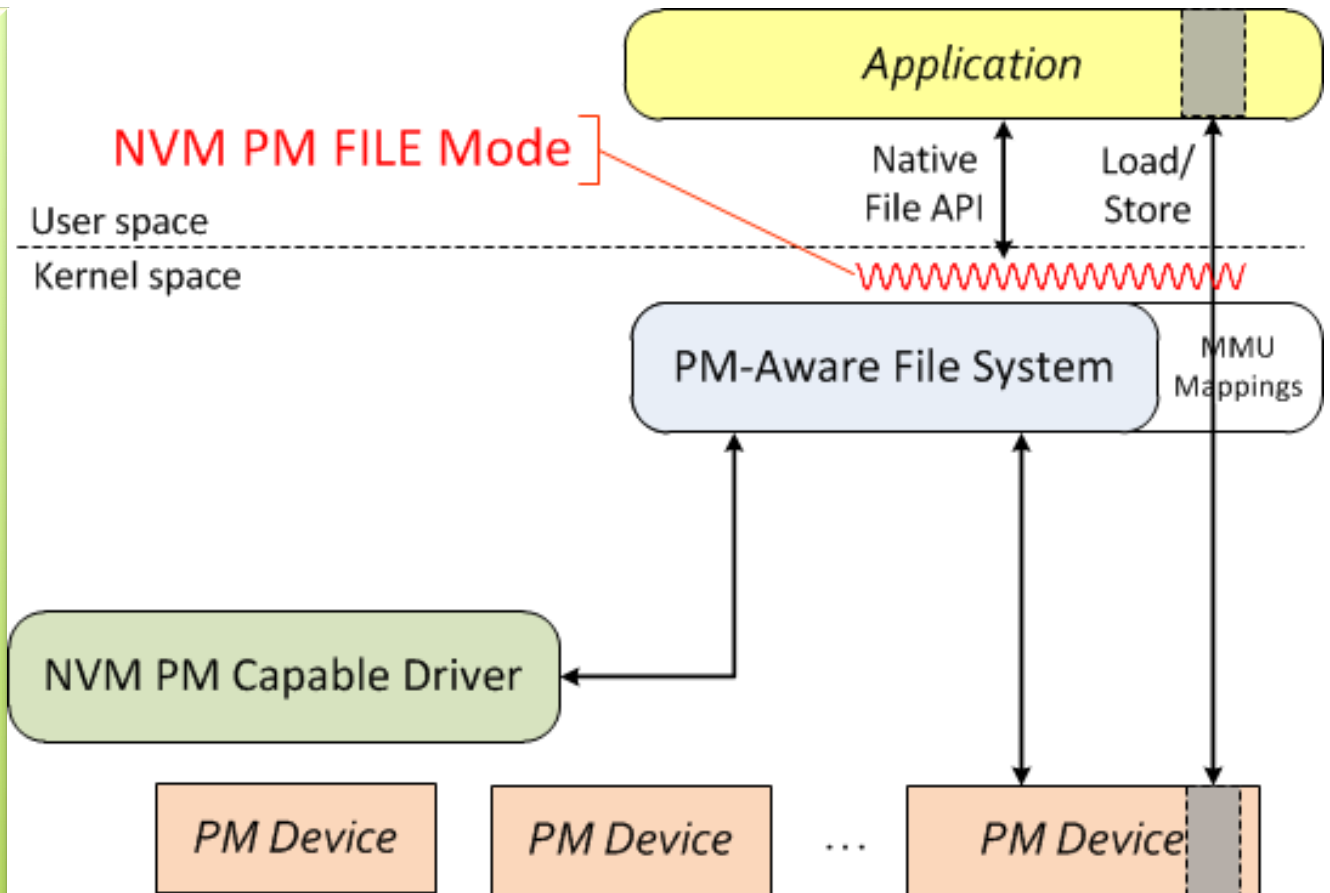
User space

Kernel space



# Exposing PMem for Applications

- NVM Files
  - PM Capable
  - Native file APIs and management
- Operations
  - Native open/close read/write
  - NVM.PM.FILE.MAP
  - ...



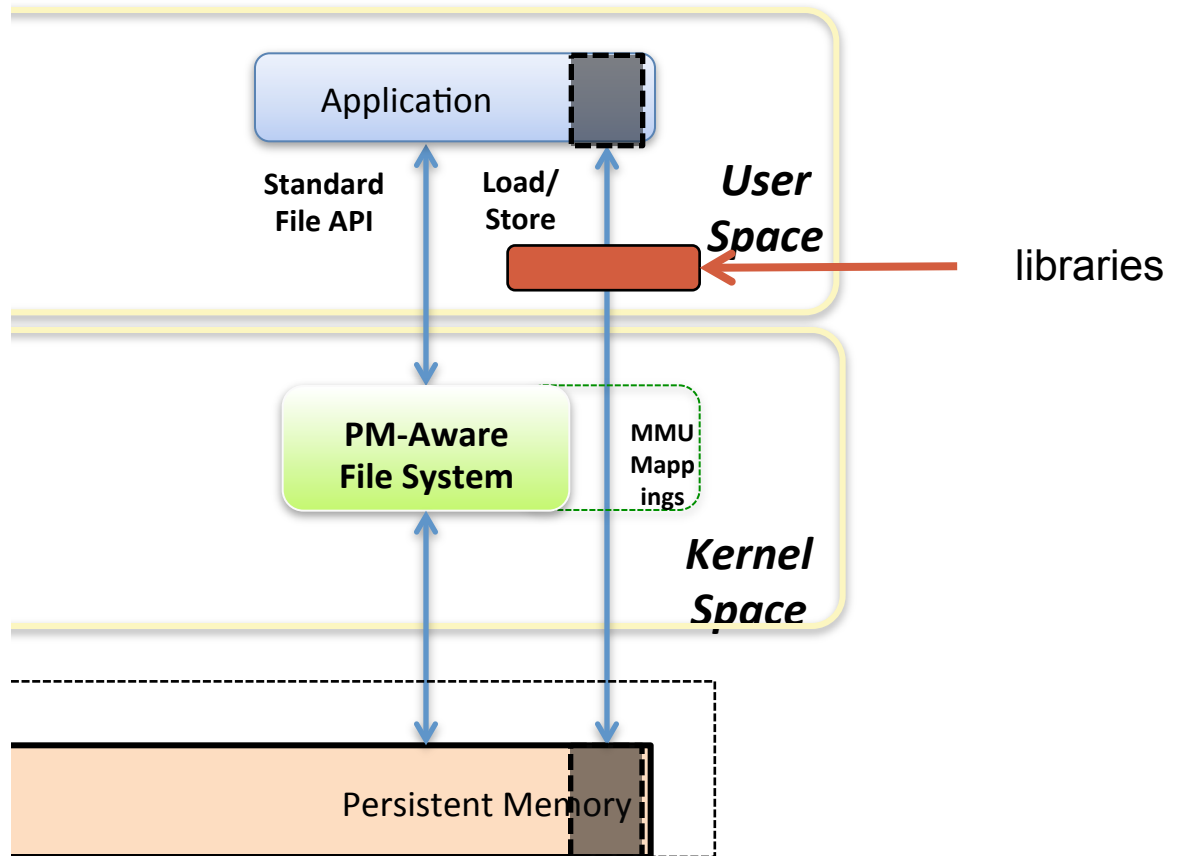
# Backup Slides

# The Value of Persistent Memory

- Data sets addressable with no DRAM footprint
  - At least, up to application if data copied to DRAM
- Typically DMA (and RDMA) to PM works as expected
  - RDMA directly to persistence – no buffer copy required!
- The “Warm Cache” effect
  - No time spend loading up memory
- Byte addressable
- Direct user-mode access
  - No kernel code in data path
- Intel collaborating with the industry to ensure open standards and choice

# Building on the Basic PM Model

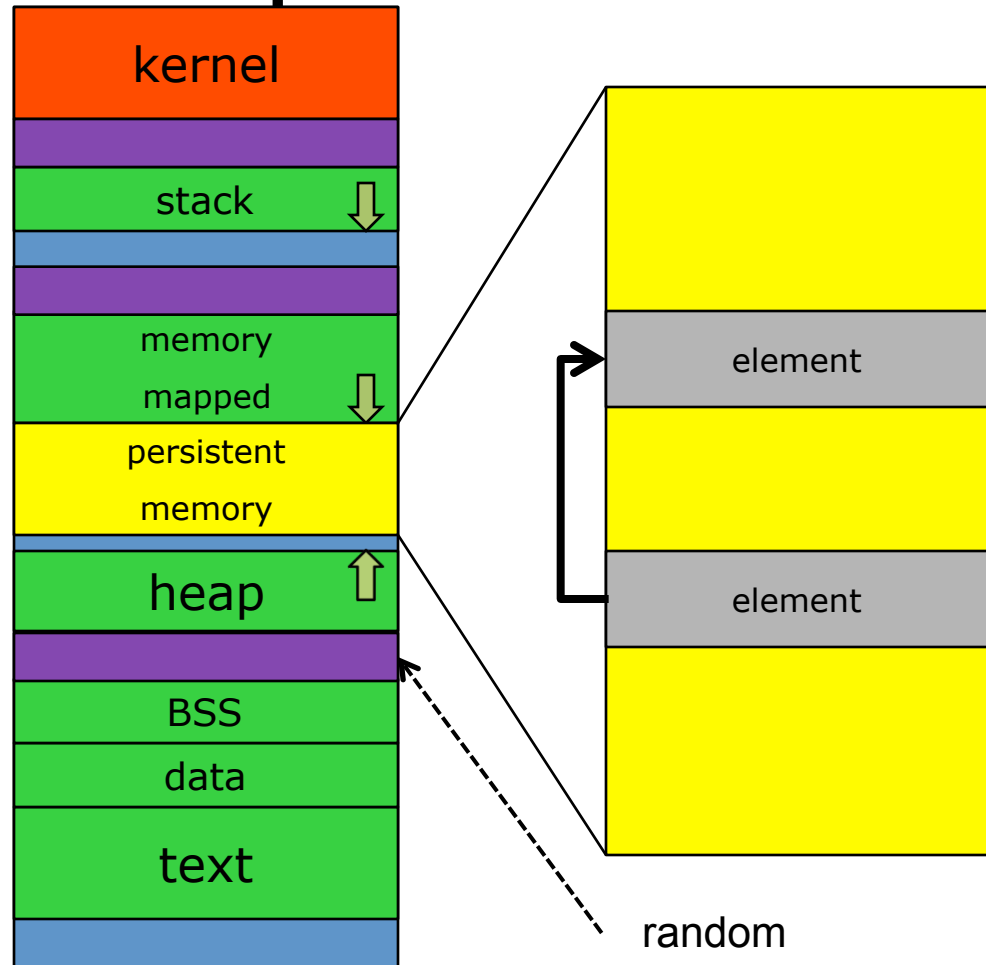
- NVM.PM.FILE programming model “surfaces” PM to application
- Still somewhat raw at that point
- Build on it with additional libraries
- **Eventually turn to language extensions**



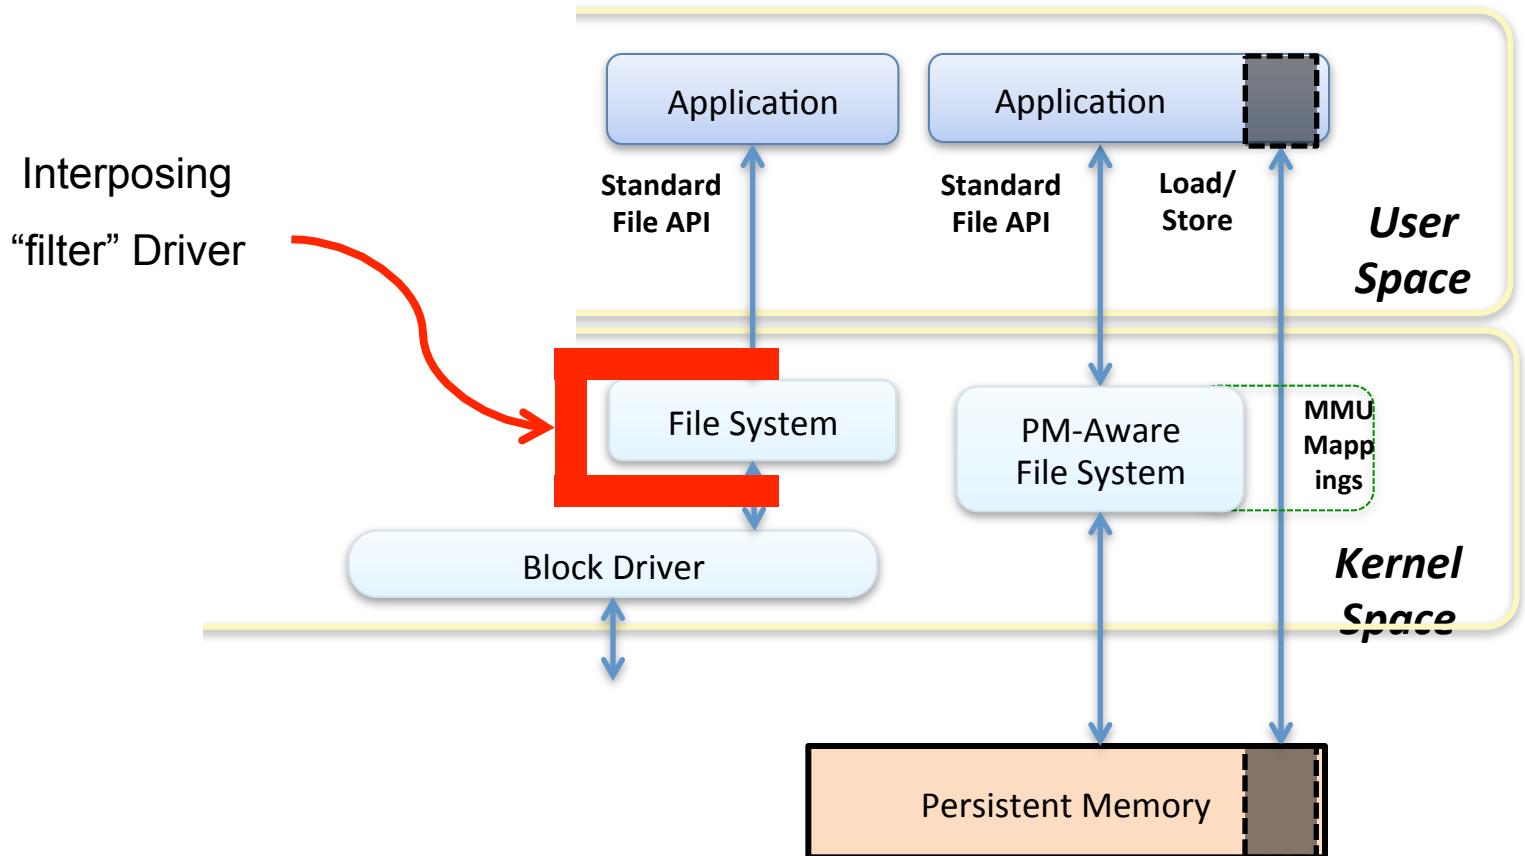


# Position Dependence

- Can pointers work across sessions?
- Will a PM file be mapped at the same address every time?



# The “C-Clamp”



# PM Interposition Points

