

The logo for SNIA (Storage Networking Industry Association) features the letters 'S', 'N', and 'I' stacked vertically in a bold, sans-serif font. To the right of the letters is a small square icon divided into four quadrants, with the top-left quadrant being yellow and the other three being white.

# PERSISTENT MEMORY PM SUMMIT

JANUARY 24, 2018 | SAN JOSE, CA

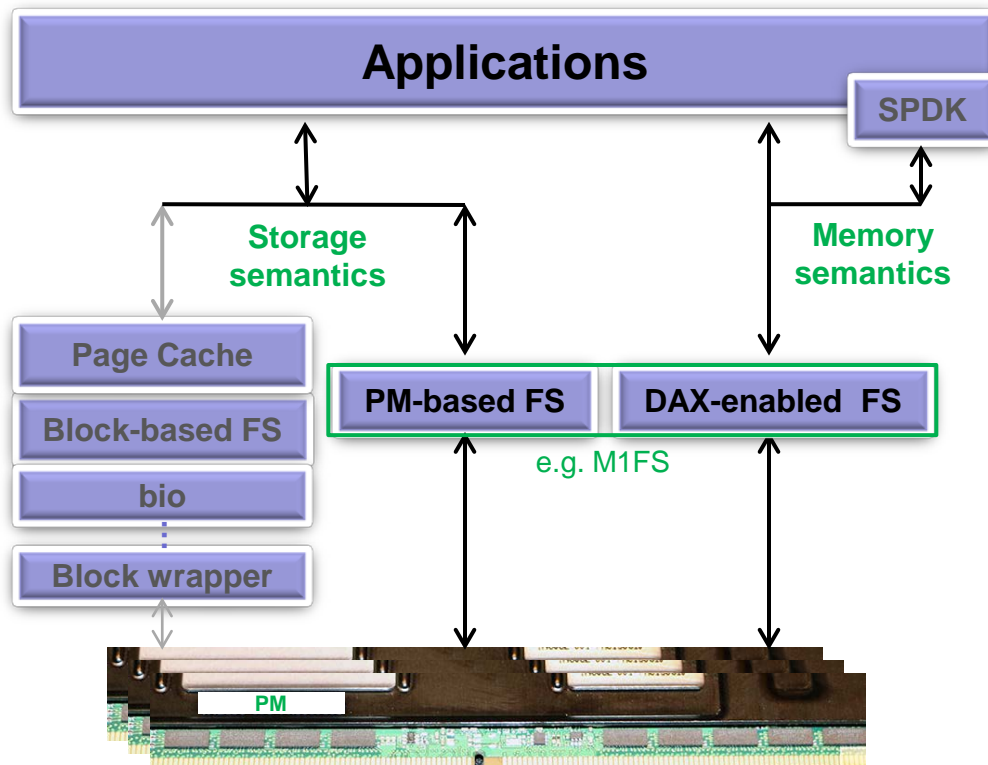
## ZUFS

### Simplifying the Development of PM-based File Systems

Dr. Amit Golander, NetApp

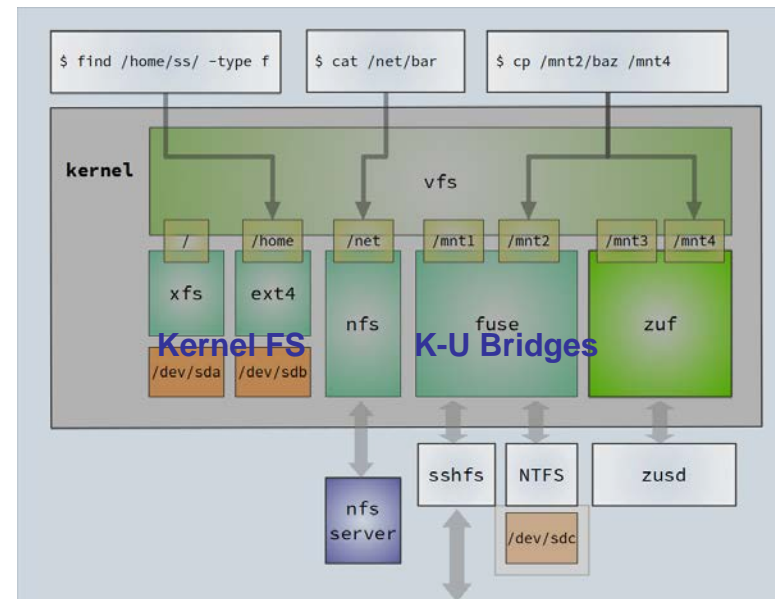
# Background

- ◆ **Plexistor (acquired by NetApp)**
  - ◆ PM-based FS pioneer since 2013
  - ◆ Contributing some of our IP
- ◆ **Our PM-based FS approach:**
  - ◆ Support legacy applications & Enable NPM (e.g. SPDK)
  - ◆ Feature rich
  - ◆ Integrate with NetApp product portfolio



# Kernel Vs. User Space FS Implementation

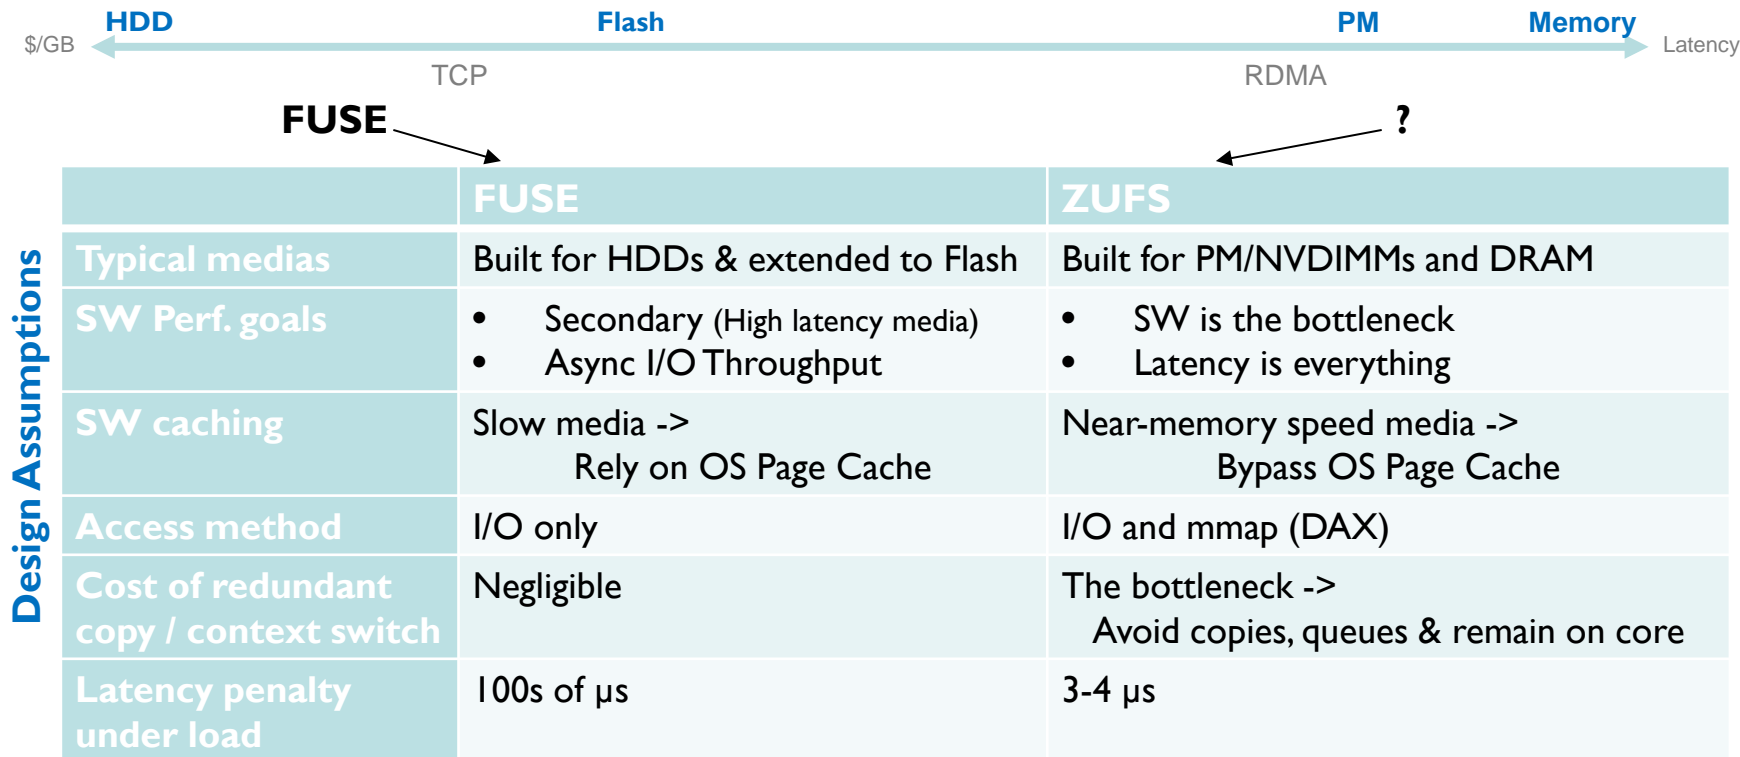
Kernel	User space
Fast (shortest path)	Portable
	Resilient (contained)
	Simpler to add functionality & Debug
	Fewer licensing restrictions



The desired missing link:  
 Near-memory speed Kernel-to-User bridge

# Why not extend FUSE to PM?

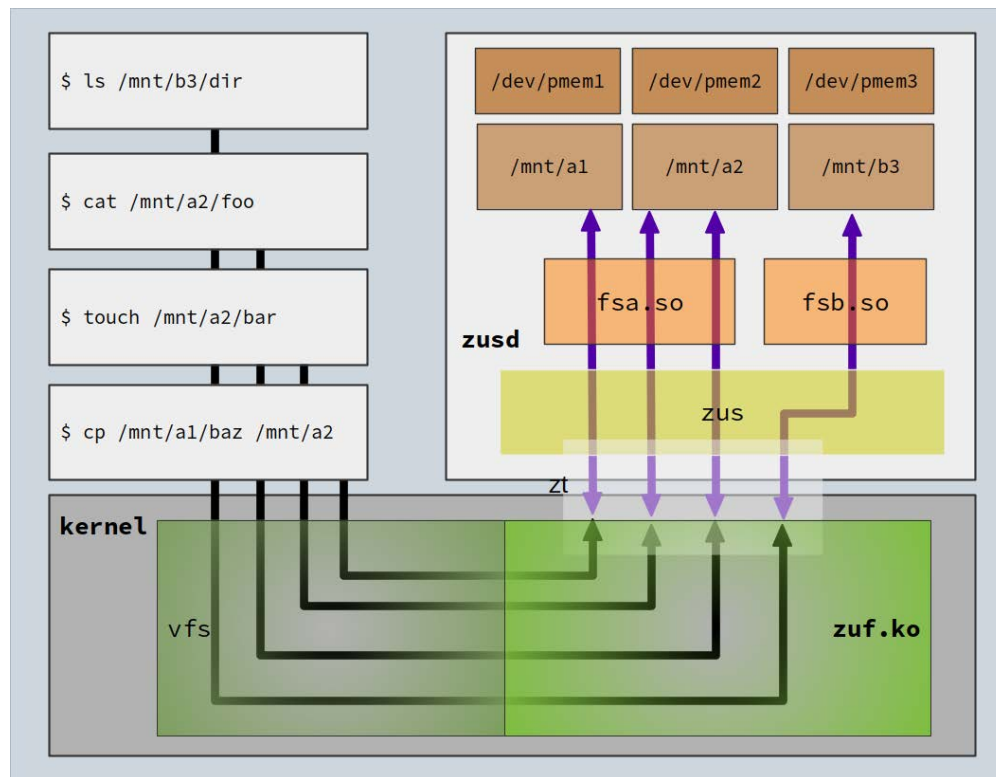
FUSE architecture is great for HDDs and ok(ish) for SSDs, but not suitable for PM





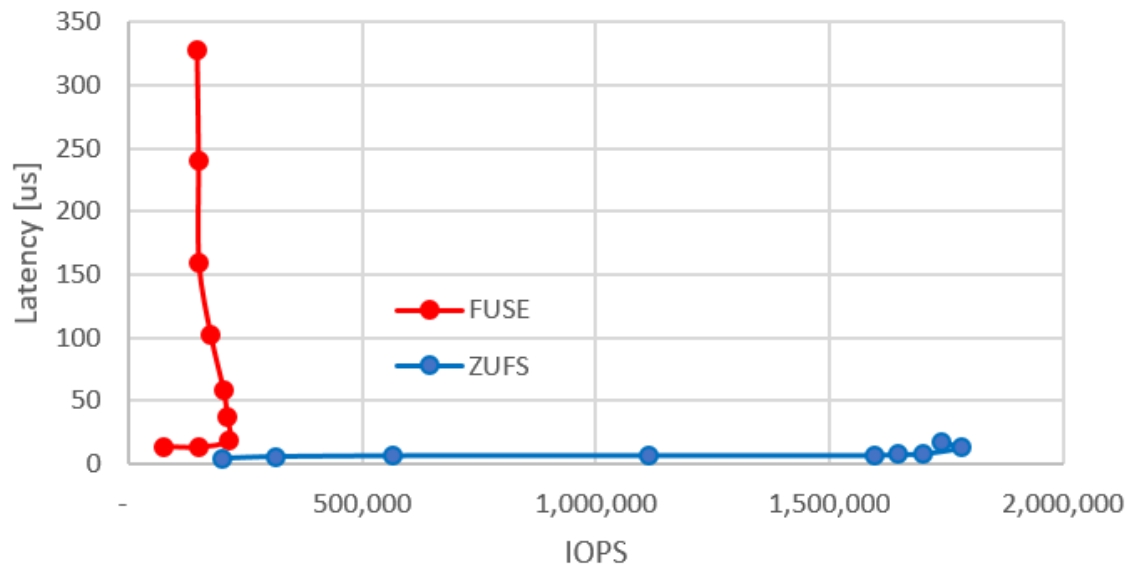
# ZUFS Features & Architecture

- ◆ **Low latency & Efficient**
  - ◆ Core & L1 cache affinity
  - ◆ Zero data copy
- ◆ **Manages devices**
  - ◆ Optimal pmem access
  - ◆ NUMA aware
  - ◆ Data mover to lower tier devices
- ◆ **Page table mapping supports I/O & DAX semantics**
- ◆ **Misc**
  - ◆ Async hook available
  - ◆ System service



# Preliminary Results (for PM)

## FUSE Vs ZUFS Penalty (PM, DirectIO)



- Measured on
  - Dual socket, XEON 2650v4 (48HT)
  - DRAM-backed PMEM type
- Random 4KB DirectIO write access

- ❖ ZUFS is a Kernel-to-User bridge designed for PM
  - ◆ Enables NetApp solutions
- ❖ Being contributed upstream
  - ◆ Hope to accelerate PM adoption and innovation
  - ◆ Link to Github (TBD)
- ❖ You're **welcome to use, review and contribute code**
  - ◆ [zufs@netapp.com](mailto:zufs@netapp.com)

Thank you