

The logo for the Storage Networking Industry Association (SNIA), featuring the letters 'S', 'N', and 'I' stacked vertically in a bold, sans-serif font. A small square icon is positioned above the 'S'.

SNIA

PERSISTENT MEMORY

PM SUMMIT

JANUARY 24, 2018 | SAN JOSE, CA

# Workload Analysis and Acceleration Using NVDIMMs and Software

Brett Williams, Micron  
Per Brashers, Enmotus

# Infrastructure Now enabled for Persistent Memory

## Backup Triggers

SAVE\_N

RESET\_N

I2C Command

## Memory controller support

### Controller functions

BAEBI spec

### Driver support

Pmem, etc.

## OS support

Linux (native)

Windows Server 2016 (native)

## Debug tools

DSM, etc

## Motherboard support

Intel, Supermicro, Tyan, ASUS, Gigabyte, ASRock, etc.

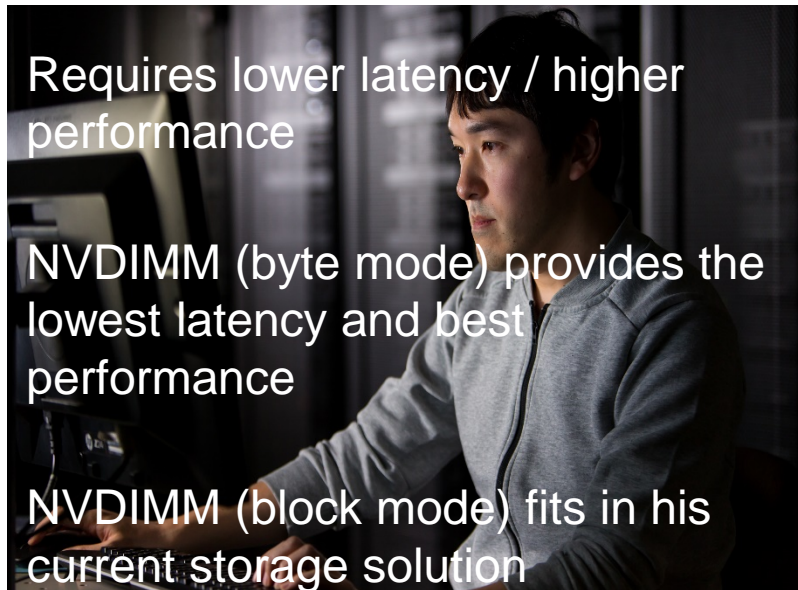
## System Support

HPE, Dell, Supermicro

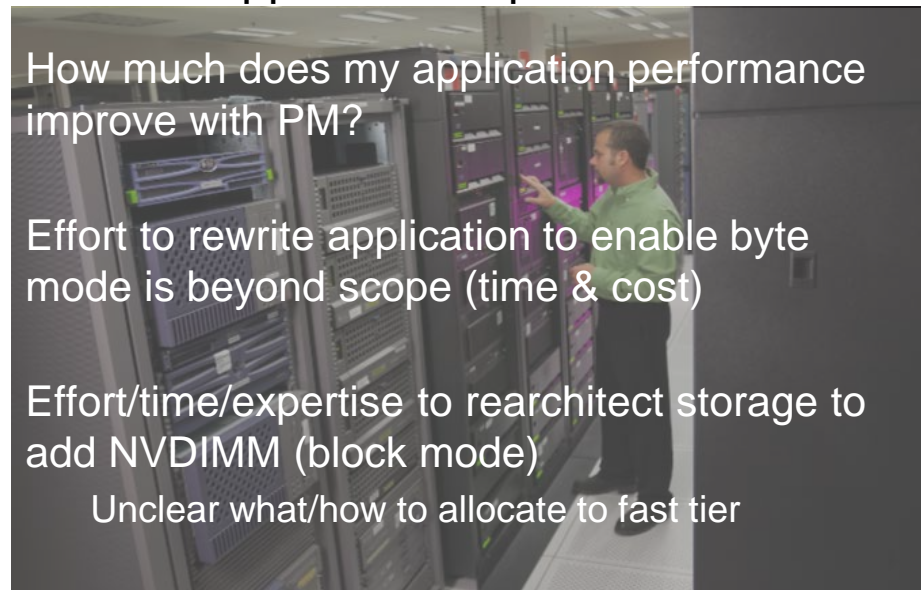


# Applications & Workloads are Not All Equal

## Customer



## Application impact




# 2018 – The year of “Workload Focus”

## ➤ Encourages tool capabilities

- ◆ Evaluate workloads – executing analytics
- ◆ Identify Hot, Cool, Cold data categories
- ◆ Automatic promotion of data to match workload use model

## ➤ Enables “Memory/Storage” Advisory

- ◆ Analyze workload for possible latency/performance improvements
- ◆ Optimize Memory/Storage for specific workload (capacity, performance, price)
- ◆ Incorporate workload learning into architecture development



The SNIA logo consists of a small square icon with a yellow top-left corner and a white bottom-right corner, followed by the letters 'SNIA' in a bold, purple, sans-serif font.

SNIA

PERSISTENT MEMORY

**PM** SUMMIT

JANUARY 24, 2018 | SAN JOSE, CA

# Workload Analysis and Acceleration Using NVDIMMs and Software

Brett Williams, Micron

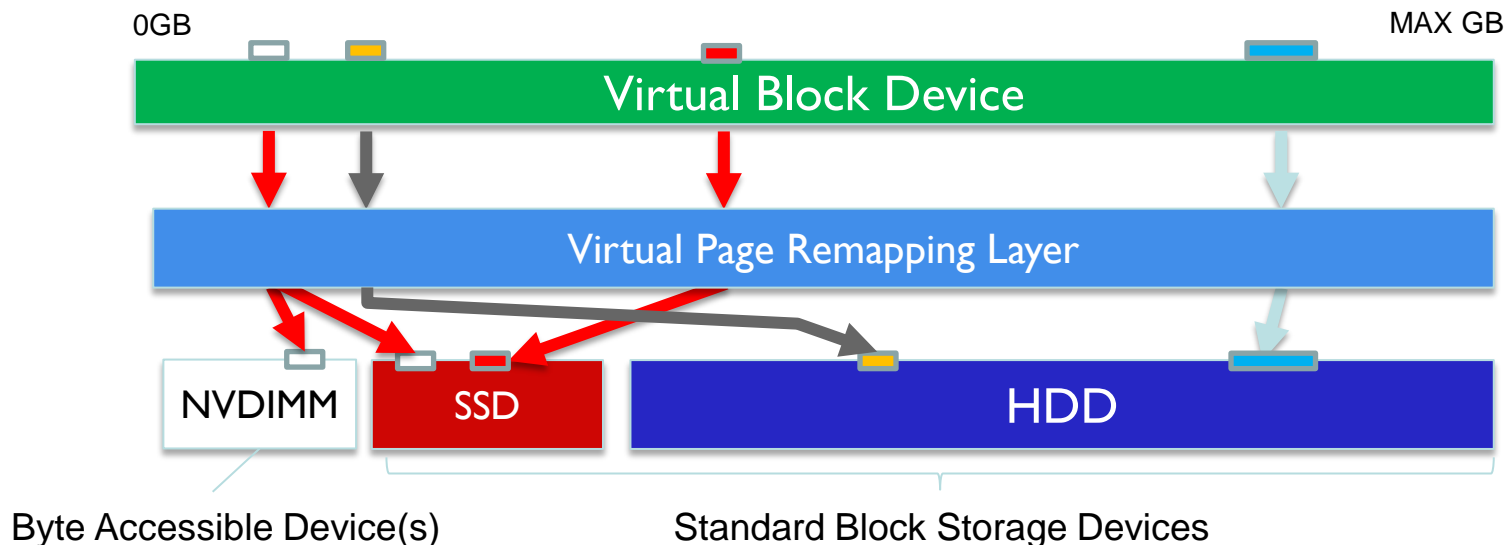
**Per Brashers, Enmotus**

# NVDIMM Adoption

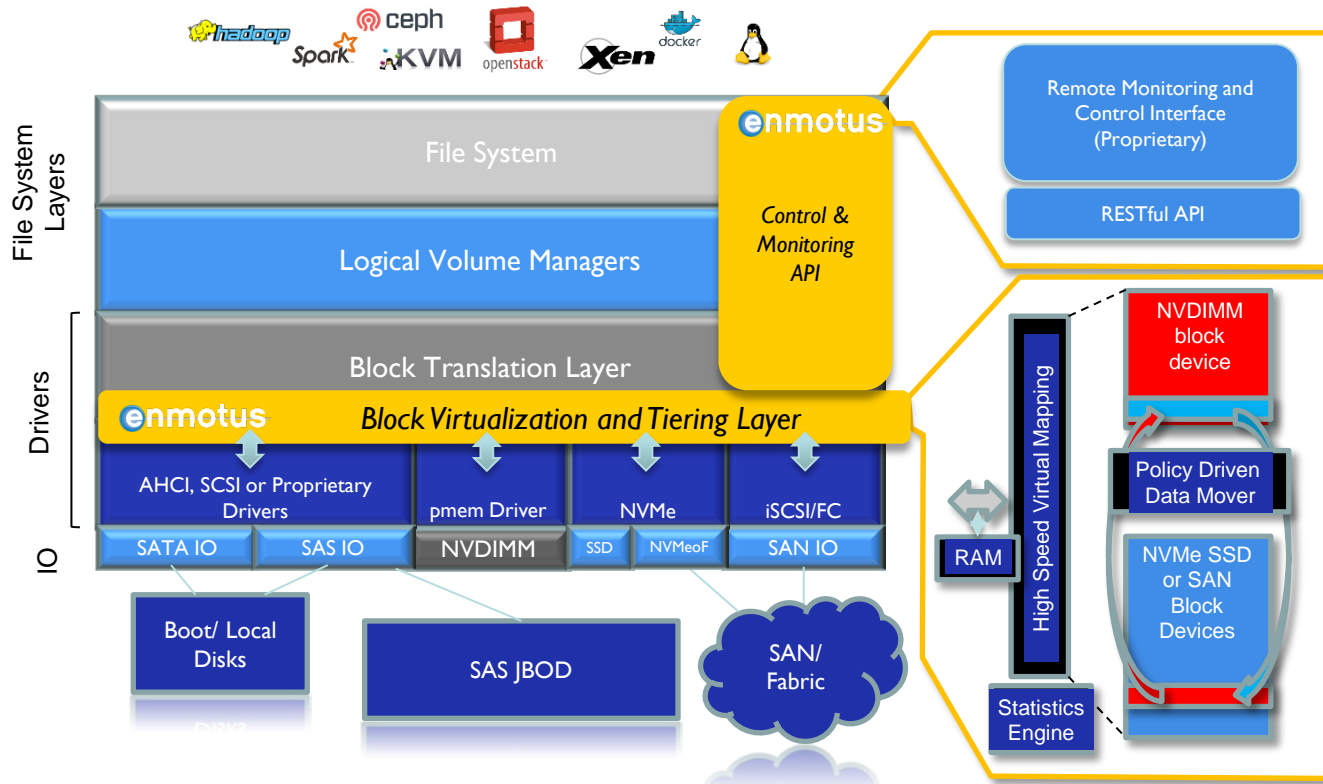
- ◆ Applications are not universally byte accessible aware .... yet
  - ◆ Will take many years before standard Linux and Windows can adopt NVDIMMs across the board
  - ◆ Not universally available to all software and application developers
  - ◆ Capacity limits applications from taking advantage w/o rearchitecting
  
- ◆ NVDIMM's are being deployed in specialized areas (e.g. embedded storage applications) but general server support is still limited to a few special SKUs
  - ◆ How far can adoption get if restricted in availability
  
- ◆ How can we accelerate adoption?
  - ◆ Bridge using block emulation – a language understood by ALL storage applications
  - ◆ Combine with multi-tiered architecture to create usable capacities

# Leveraging Performance Block IO Emulation-Tiering

- Objective: Bridge to existing applications using a language understood by all applications – block IO
- Leverage the built in block translation modes for Linux (pmem) and Windows
- Combine with performance tiering technology to solve the capacity issue



# Linux Example Stack (based on Enmotus stack)





# How Effective is NVDIMM Block Emulation/Tiering?



Per NVDIMM Performance Delta – Software Striping increased this to ~2M IOPS


- Demo at SuperComputing Nov 2017 on Dell 740XD with Micron NVDIMMs, Enmotus stack
- Multiple NVDIMM/pmem (block mode) tiered with NVMe drives
- ~2M IOPS achieved on fully tiered data using 6-8 NVDIMMs striped
- 18-20GB/s sustained rate on fully tiered data
- Utilizes NVMe SSD as the capacity portion and “slow” tier

# Real Time Block IO Analytics

- Emulation approach that works with today's apps allows better assessment capabilities
- Real time measurement and logging of IO statistics
- Available via log files or RESTful/JSON
- Identify and experiment
  - ◆ How much hot data there is instantaneously or over time
  - ◆ Can NVDIMM's help an application?
  - ◆ Trial NVDIMM insertion into running volume to assess if it can help

# Example Implementation – Real Time Tiering

-  NVMe SSDs  
(Slow Devices)
-  NVDIMM  
(Fast Devices)



IO Activity

Capacity Range



# Conclusions

- Block emulation and tiering for NVDIMMs offers significant performance gains with no change required to application or file system layers
  - ◆ Additional software layer introduces latency but still achieve 1M IOPs+
- Eases the transition for applications to full byte level access
- Allows adoption of NVDIMMs to occur more rapidly in supporting servers
- Combined with analytics tools, helps customers decide just how much NVDIMM will help their application rather than guessing

SNIA

PERSISTENT MEMORY

PM SUMMIT

JANUARY 24, 2018 | SAN JOSE, CA

Thank You!

