




Persistent Memory Advances

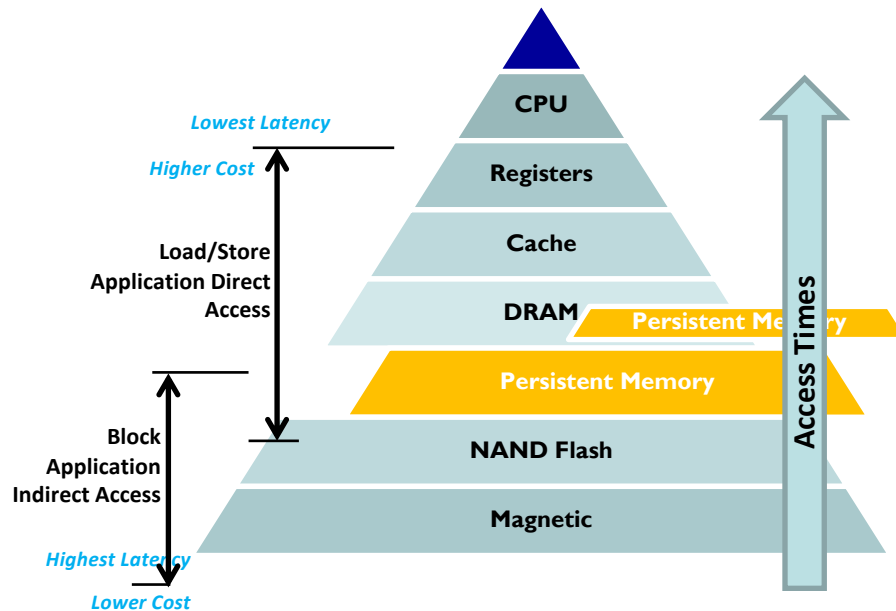
Arthur Sainio

*Persistent Memory and NVDIMM SIG Co-Chair, SNIA
Director, Product Marketing, SMART Modular Technologies*

- 
- A horizontal decorative bar spans the width of the slide, just below the header. It is composed of a series of colored rectangular segments in the following order from left to right: purple, light blue, light green, orange, grey, light blue, purple, grey, orange, grey, light blue, light green, and purple.
- **Why Persistent Memory is Important**
 - **Persistent Memory Applications**
 - **SNIA and Industry Alliance Efforts on Persistent Memory**

Why is Persistent Memory Important?

- Bridges the gap between DRAM and Flash
- Dramatically increases system performance
- Enables a fundamental change in computing architecture
- Apps, middleware and OSs are no longer bound by file system overhead in order to run persistent transactions




What It Is and Isn't and Why is it Important?



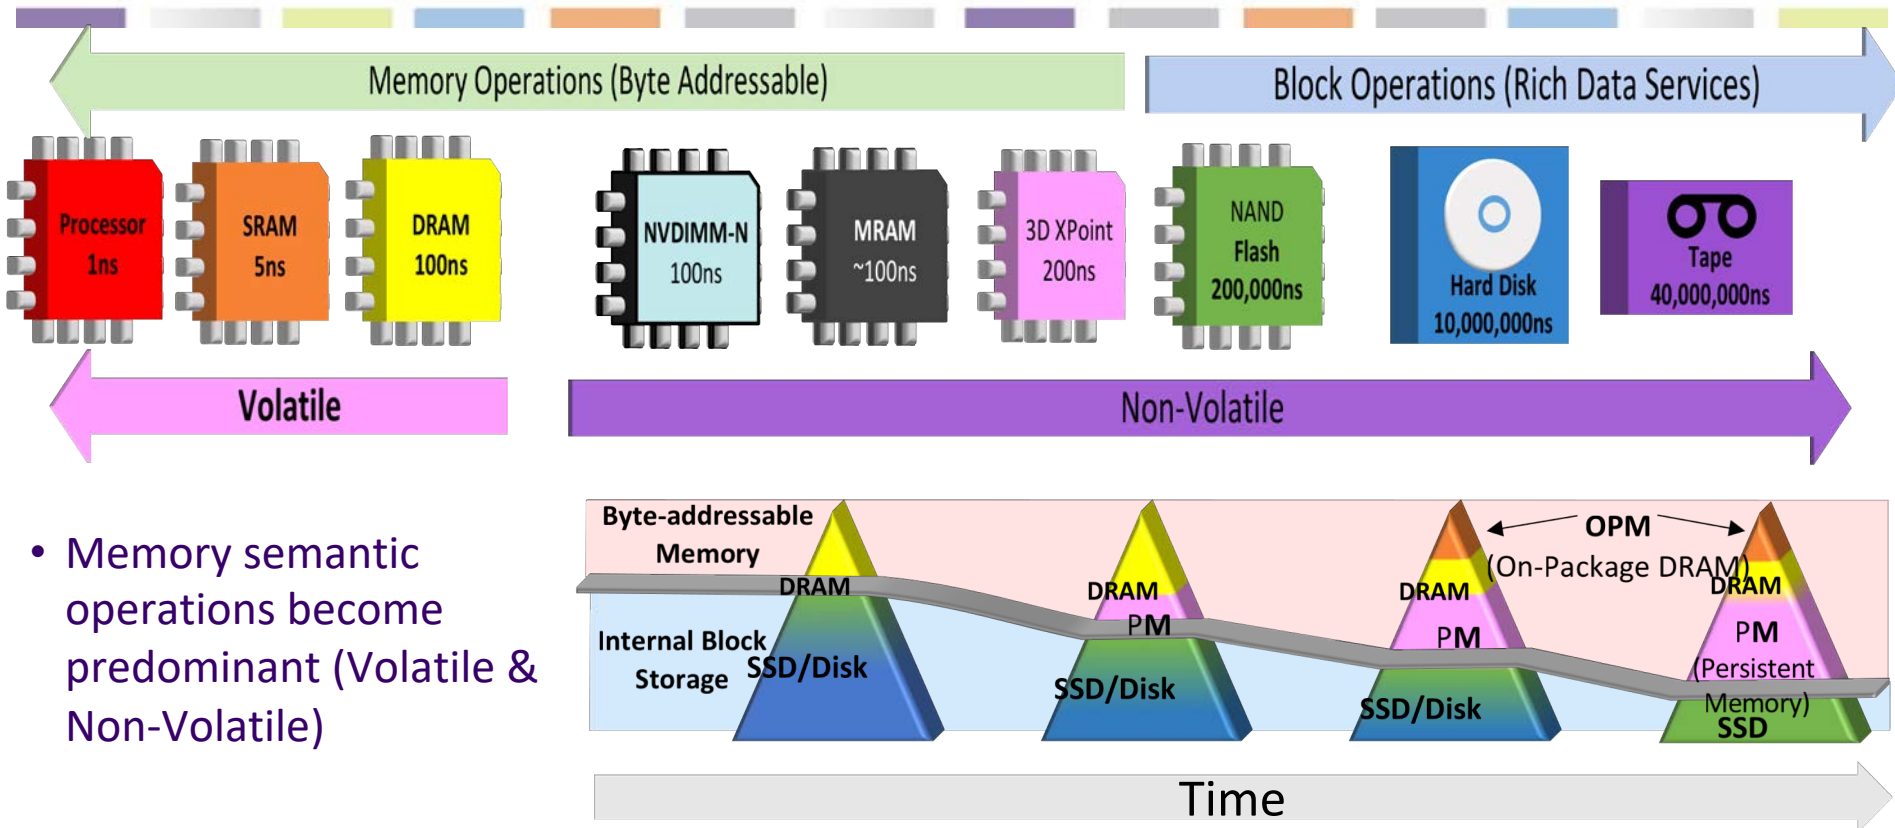
- **What is Persistent Memory?**
 - Non-Volatile
 - Byte Addressable
 - Low Latency <1 μ s
 - Densities greater than or equal to DRAM (for wide-scale adoption)
- **Why is persistent memory important?**
 - The vast majority of compute applications operate on 1, 2, 4, or 8 bytes at a time
 - For maximum performance memory technology must be directly addressable (load/store byte access)
 - With non-volatility, compute applications do not need serialization or to commit writes to a lower level (slower) memory tier
- **Non-Volatile Memory \neq Persistent Memory**
 - Non-Volatile memory is not necessarily Persistent Memory in the industry term usage
 - NAND flash is paged-based and has millisecond write speeds
 - In it's current form NAND Flash by itself is not PM

Storage vs. Memory



- Typical NUMA range: 0-200nS
- Typical context switch range: above 2-3μS

Memory and Storage are Converging



- Memory semantic operations become predominant (Volatile & Non-Volatile)

Multiple Persistent Memory technologies are nearing commercialization

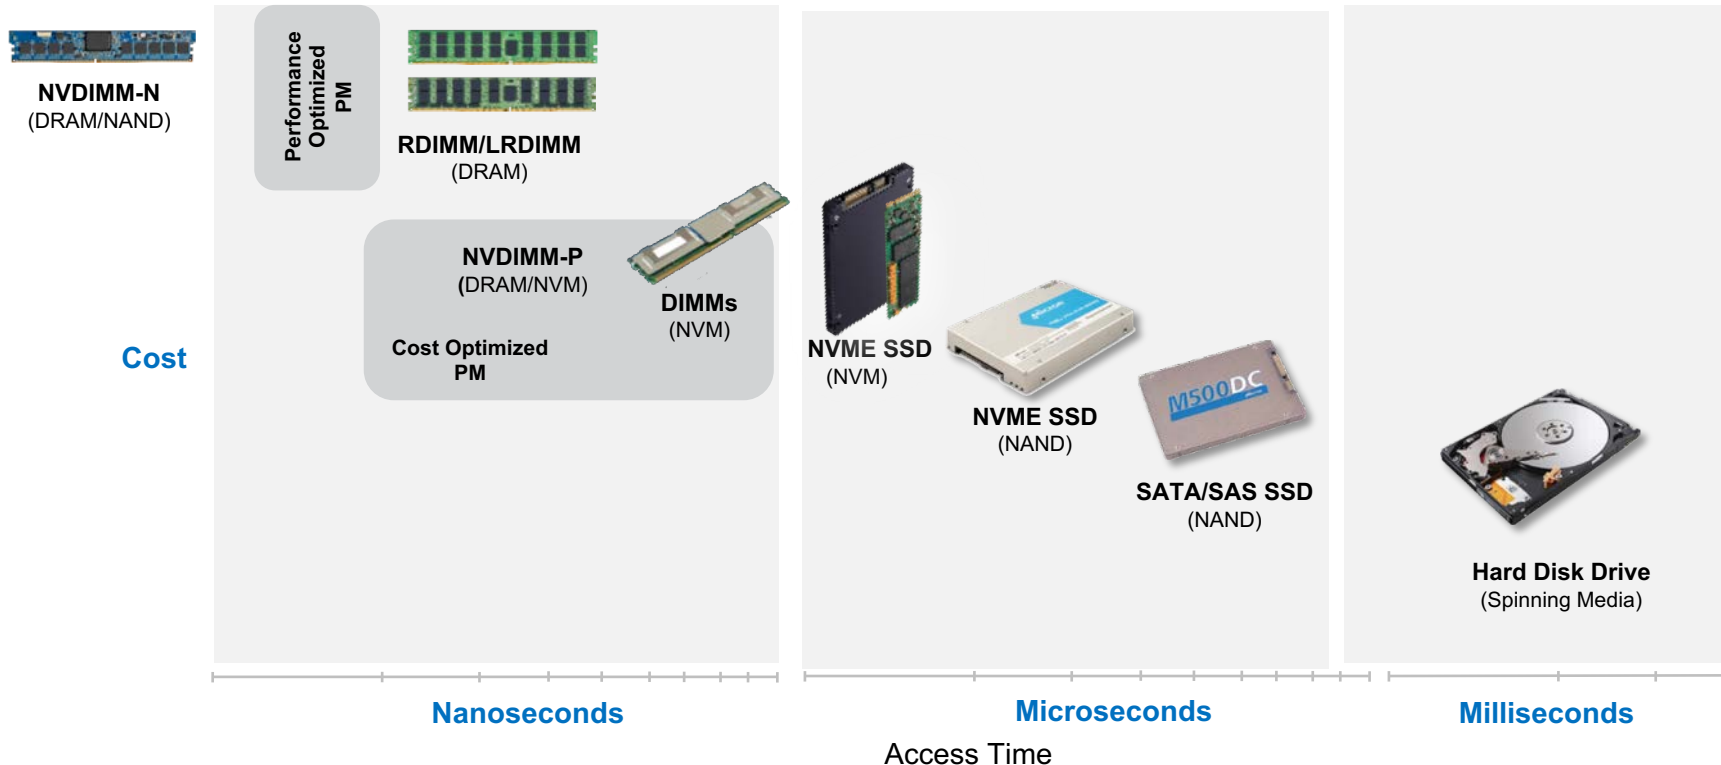
- **Phase Change (PCM):** a middle ground between DRAM and Flash
- **MRAM:** DRAM replacement? density past 8Gb, lower idle power
- **ReRAM:** Flash replacement? High density, better endurance
- **CNTRAM:** Carbon Nanotube based memory – another DRAM replacement?

Persistent Memory Technology Overview




Technology	FRAM	MRAM	ReRAM	PCM	3D XPoint	NRAM	NVDIMM-N
Density	4K-4Mb	256Mb	TBD	128Mb	128Gb	16Gb	32GB
Endurance	10^{12}	10^{12}	10^5	10^8	10^{12}	10^{11}	∞
Writes	Byte	Byte	Byte	Byte	Byte	Byte	Byte
Read Latency	70-100ns	70ns	25ns	20ns	100ns	100ns	40-180ns
Write Latency	70-100ns	70ns	12 μ s	65ns	500ns	100ns	40-180ns
Power	Low	Med-Low	Low	Med	Med	Low	High
Interface	DRAM	DDR3 DDR4	Flash-Like	Unique	Unique	DDR4	DDR3 DDR4
Availability	Limited	Prod'n	Alpha	Limited	Samp.	Samp.	Volume

Existing and Emerging Variations of Persistent Memory Products



The Role of the NVDIMM-N


- Paves the path for Persistent Memory DIMMs
 - ◆ Allows software development today
 - ◆ Gets the creative process started
 - ◆ Supports segment of PM application needs
- A vehicle to debug PM-based systems
 - ◆ Software ready when hardware ships



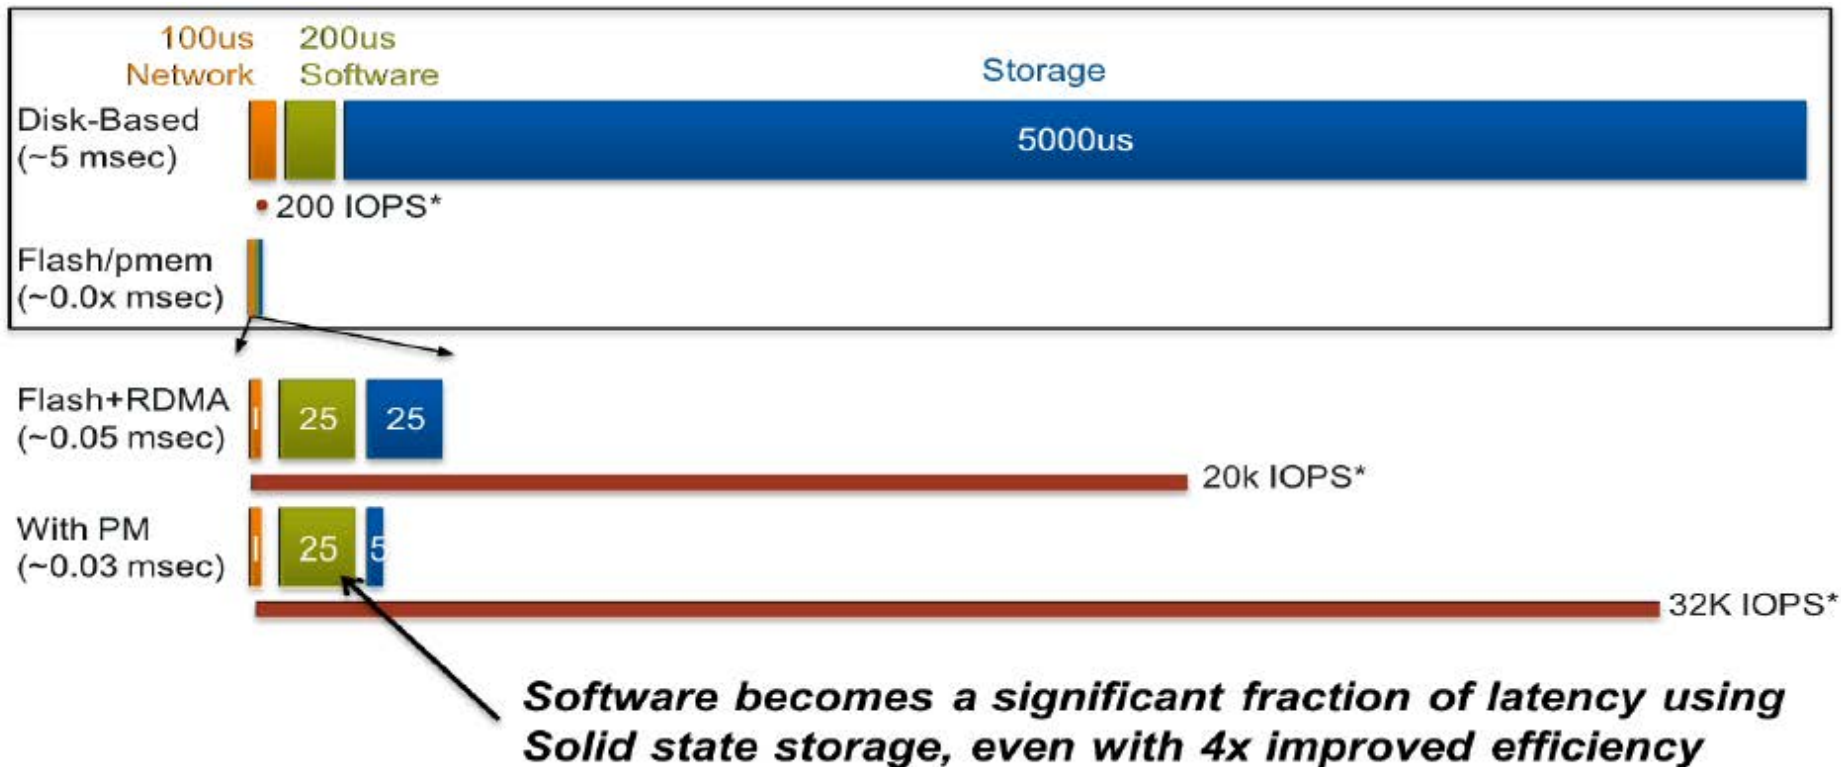
This Requires Standards!

SNIA NVM Programming Technical Work Group (TWG)

Formed 06-11-2012



Persistent Memory – Challenging DRAM and Flash



* Max potential 1-thread random sector

- Both Linux and Microsoft joined the SNIA effort to help steer the direction of PM
 - Applications have direct Load/Store access to PM
 - End result: Both OS's are structured almost identically (ex: DAX)
 - No other unique drivers needed
- VMWare has also offered support for PM

Persistent Memory support in OS's ahead of volume adoption


- NVDIMM-N - the first HW available to run PM Applications.
 - Success stories emerge highlighting the application benefits of PM
 - Example: “Tail of Log” for SQL Server
 - Used in man All Flash Arrays
- Alliance formed between JEDEC and SNIA to effectively drive adoption of the technology

NVDIMM-N Prove the Benefits of PM

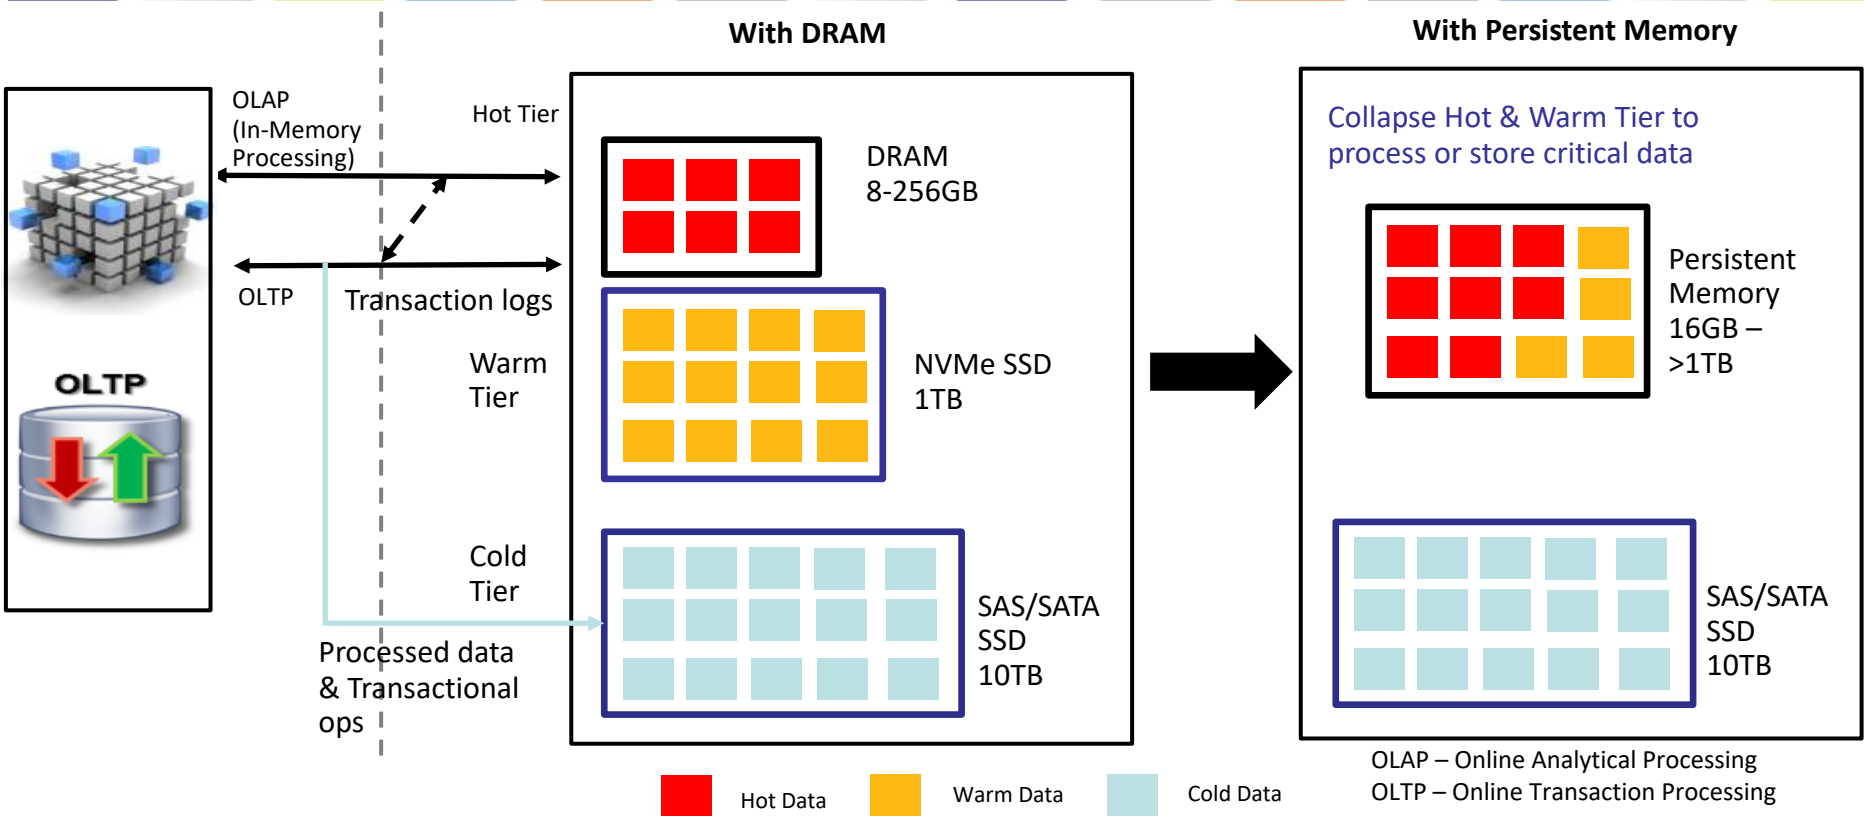


Persistent Memory Applications

What Applications will use PM?

- 
- A horizontal bar with a series of colored squares in purple, grey, yellow, blue, orange, and grey, spanning the width of the slide.
- Applications that have a large working set of data with a need for persistence
 - ◆ Using NVMe or standard SSDs add latency
 - ◆ Decreasing the latency to avoid disk access
 - In Memory Databases
 - ◆ Application driven data locality
 - ◆ Newer DB adaptations beginning to use PM
 - Productivity Improvements
 - ◆ Software infrastructure is enabled
 - ◆ Standard libraries are available

Persistent Memory: Evolution of In-Memory Apps




Early PM System Support and Applications



A sample of companies publicly showing PM support

Persistent Memory Adds Value Across Diverse Applications

SNIA®



Relational Database

Scale-out Storage

Virtual Desktop Infrastructure

Big Data

In Memory Database

Middleware

HPC

MSFT SQL

MySQL

Maria DB

Oracle

Vmware
VSAN

MSFT Azure

Store Virtual

Vmware VDI

Citrix HDI

Mongo DB

Cloudera

HortonWorks

Hadoop

Cassandra

MSFT SQL Hadoop

SAP HANA

MSFT SQL Hekaton

XAP Gigaspace

Journaling, Transaction logs

Java

.NET

Optimized abstraction

HPC

Check point acceleration

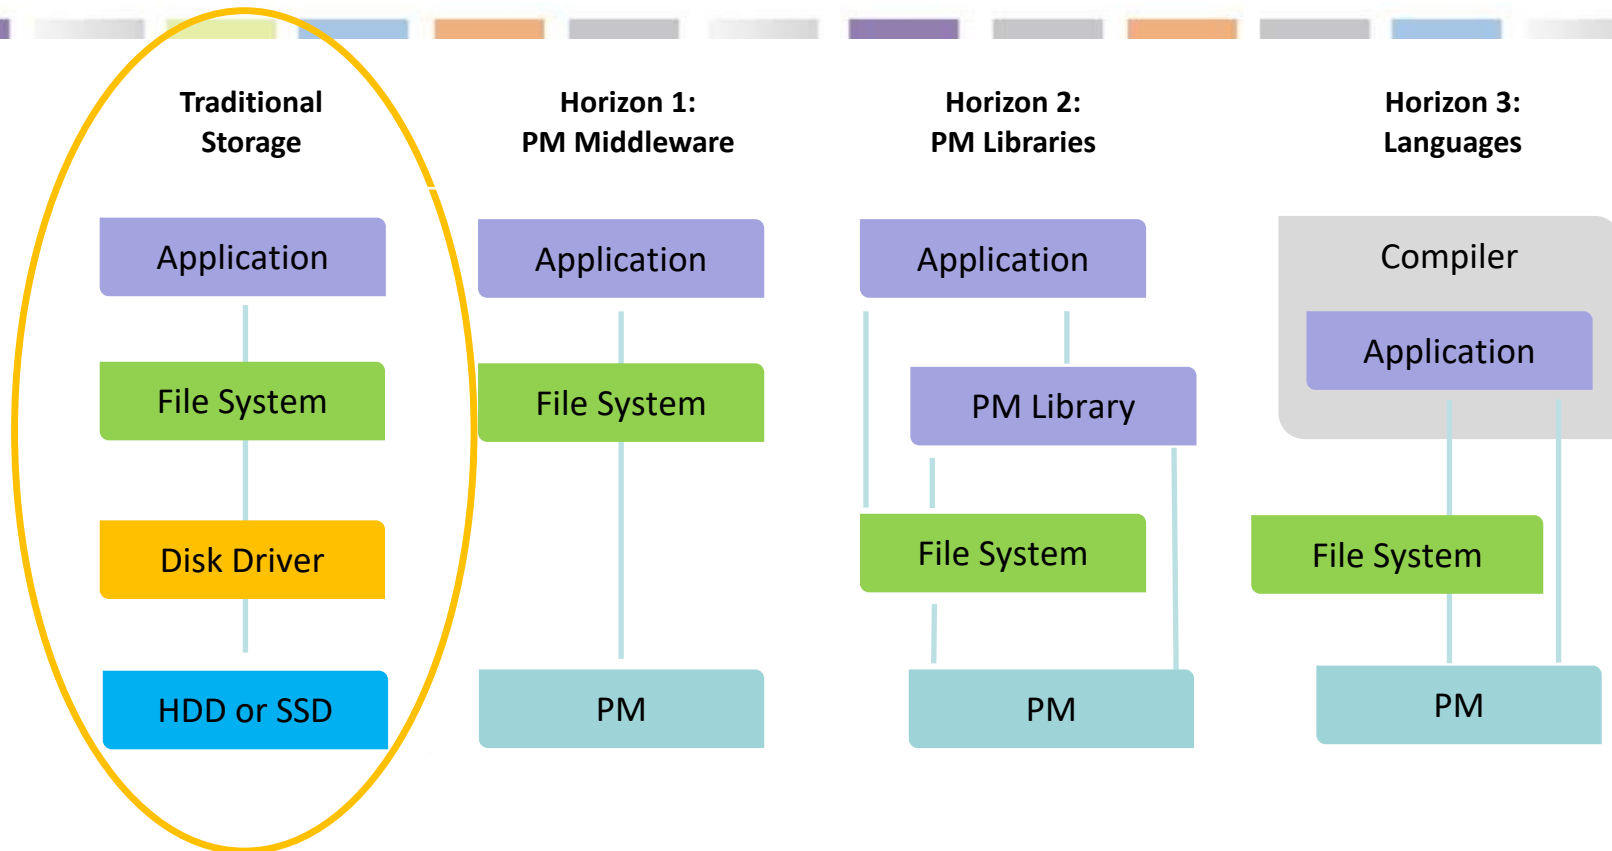
Log acceleration:
write combining and
caching

Tiering, caching, write
buffering, meta-data
storage


Higher VM
consolidation

Higher performance

Application Horizons - Today



Application Horizons – Ultimate Goal



NVDIMM-N and 3DXPoint Applications



- Many NAND flash storage array vendors are using NVDIMM-N modules for write acceleration and commit logging
 - These applications do not require a density higher than multiple GBs so they are well-suited for NVDIMM-N
- 3DXPoint is well-suited for PM applications like In-Memory databases that need 100's of GBs to TBs of persistent memory that is used in combination with DRAM

Example: Need for In Memory Persistent Database



DreamWorks



- 600TB's of data in one film
- Many small items in a large working set
- Substantial re-use and repeat file I/O
- Expensive to compute and convert
- Distributed clients doing similar things
- Writes are immutable; lockless updates

Goal with PM

- NVDIMMs in each workstation and server
- Accelerate local workflows
- Cluster of Persistent Memory servers
- Software stack that provides RPM-as-a-Service
- A way for apps to persist things and reduce trips through the storage stack
- A way for apps to find and get things
- That behaves like named shared memory

Example: Using Persistent Memory to Accelerate HCI Storage Performance



Differentiated value with Persistent Memory in HCI storage tier

Create a new persistent memory tier for metadata (benefits ALL apps)

1. Read-modify-write with persistent memory as byte addressable is 100X faster than block storage
2. Faster metadata access for dedup, checksum etc results in reduced CPU utilization and higher IOPS for all apps
3. Faster reboots due to persistence of metadata in persistent memory (save time for not having to rebuild metadata from logs)




Example: WDC IntelliFlash Write Cache



- Separate logging for incoming writes
- Write is acknowledged after persisting to the write cache
- Coalesced I/O is flushed to drives after dedupe and compression
- Uses high performance media as the latency is crucial for many applications like DBT and OLTP
- Best fit for NVDIMM



Example Results – Latency Comparison



- All flash array with 24TB capacity
- iSCSI protocol
- fio with 4 clients and 8 LUNs



SNIA and Alliance Efforts on Persistent Memory

Who is SNIA?



170
industry leading
organizations



2,500
active contributing
members




50,000
IT end users & storage
pros worldwide

What SNIA Does



- Leads the storage industry worldwide in developing and promoting:
 - Standards development and adoption of open source software
 - Interoperability assurance – plugfests and conformance testing
 - Technology acceleration – SIGs and collaborations
 - Global vendor neutral education – certification, webcasts, white papers



SNIA Efforts on Persistent Memory



- **SNIA Technical Council & Technical Working Groups (TWGs)**
 - ◆ Non-Volatile Memory Programming TWG
 - ◆ This is the body developing the NVM Programming Model
- **SNIA Standards (aka Technical Positions), Software & White Papers**
 - ◆ NVM Programming Model v1.2 (June 2017) – Technical Position
 - ◆ PM Remote Access for High Availability v1 (February 2016) –White Paper
 - ◆ PM Atomics and Transactions v1r1 (January 2017) – White Paper
- **SNIA Solid State Storage Initiative (SSSI)**
 - ◆ One of many Forums and Initiatives within SNIA
 - ◆ SSSI sponsors the Persistent Memory and NVDIMM Special Interest Group (SIG)
 - ◆ Deliverables: PM Summit, webcasts, videos, presentations, tutorials

- JEDEC standards address architectural, electrical, test, and SPD issues relating to memory design and manufacturing for commercial applications
- JC-45.6 subcommittee:
 - Hybrid Modules
 - This is the subcommittee that governs NVDIMM work
 - Most recent – JESD248A – DDR4 NVDIMM-N (March 2018)
 - Governs all behavior of module including backup/restore

Persistent Memory and NVDIMM SIG Charter



- To accelerate the awareness and adoption of Persistent Memories and NVDIMMs for computing architectures
- The Persistent Memory and NVDIMM SIG will:
 - ♦ Educate on the types, benefits, value, and integration of Persistent Memories
 - ♦ Communicate usage of the NVM Programming Model developed to simplify system integration of current and future PM technologies
 - ♦ ***Influence and collaborate with middleware and application vendors to support Persistent Memories***
 - ♦ Develop user perspective case studies, best practices, and vertical industry requirements
 - ♦ Coordinate with industry standards groups and promote industry standards related to PM and NVDIMM
 - ♦ Synchronize and communicate a common Persistent Memory taxonomy

- **SNIA Persistent Memory Application Enabling**
 - A program is being formed to enable the application development community to build Persistent Memory applications
 - Will launch in early 2019
 - ***Looking for new members/contributors***

Goal: Accelerate Development of PM Applications



Thank You!