SNIA PERSISTENT MEMORY + COMPUTATIONAL STORAGE SUMMIT 2022
VIRTUAL EVENT • MAY 24-25, 2022

Accelerating Business Critical Oracle Workloads with VMware Persistent Memory using Intel DC Optane PMM

Sudhir Balasubramanian - Sr Staff Solution Architect - VMware

Arvind Jagannath - Product Line Manager - Cloud Platform - VMware

How Intel® Optane™ PMem benefits VMware vSphere users

Delivering improved performance and increased VM density, affordably

Solution on a DIMM

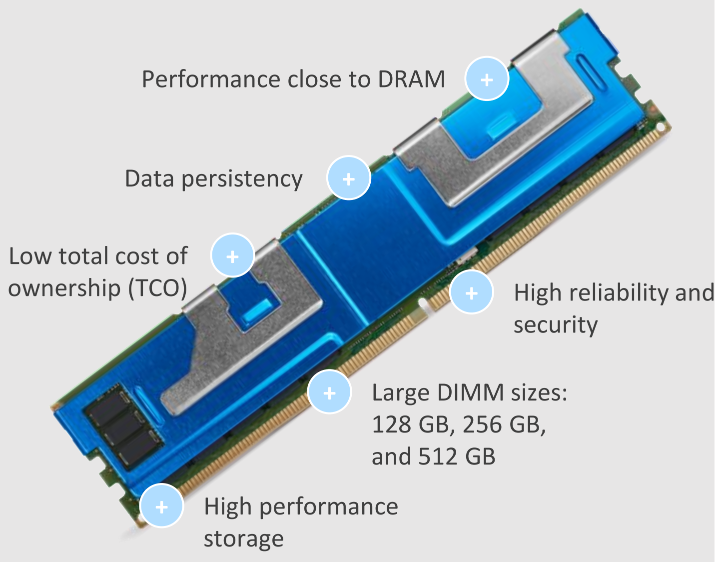

Memory Mode (MM)

Ideal for increasing memory capacity at lower cost than DRAM

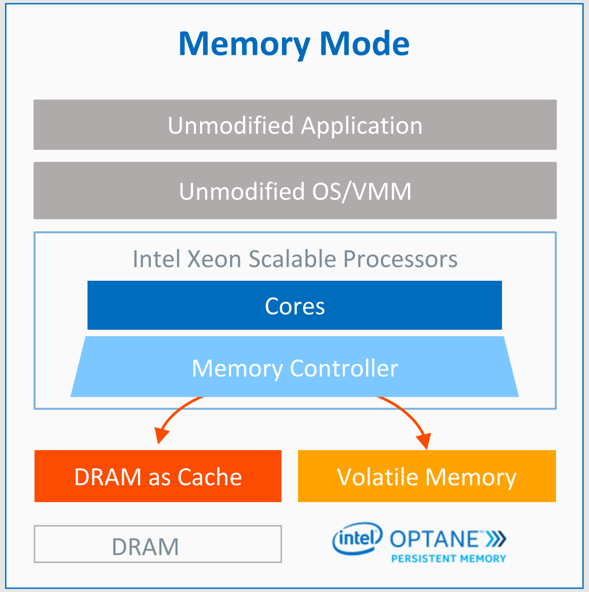

Memory Mode
(2nd level memory)
Volatile with DRAM caching

DRAM + App Direct Mode (AD)

For apps requiring data persistence and maximum capacity

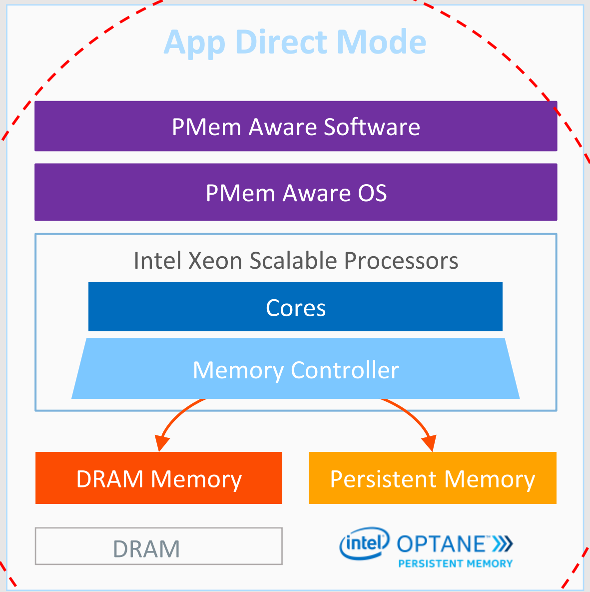

App Direct (AD)
Persistent programming model with no caching

App Direct Mode

- PMem & DRAM - independent memory resources under direct load/store control of the application
- PMem - byte-addressable mapped into system physical address space & directly accessible by applications

How Intel® Optane™ PMem Is Exposed inside VMware vSphere®

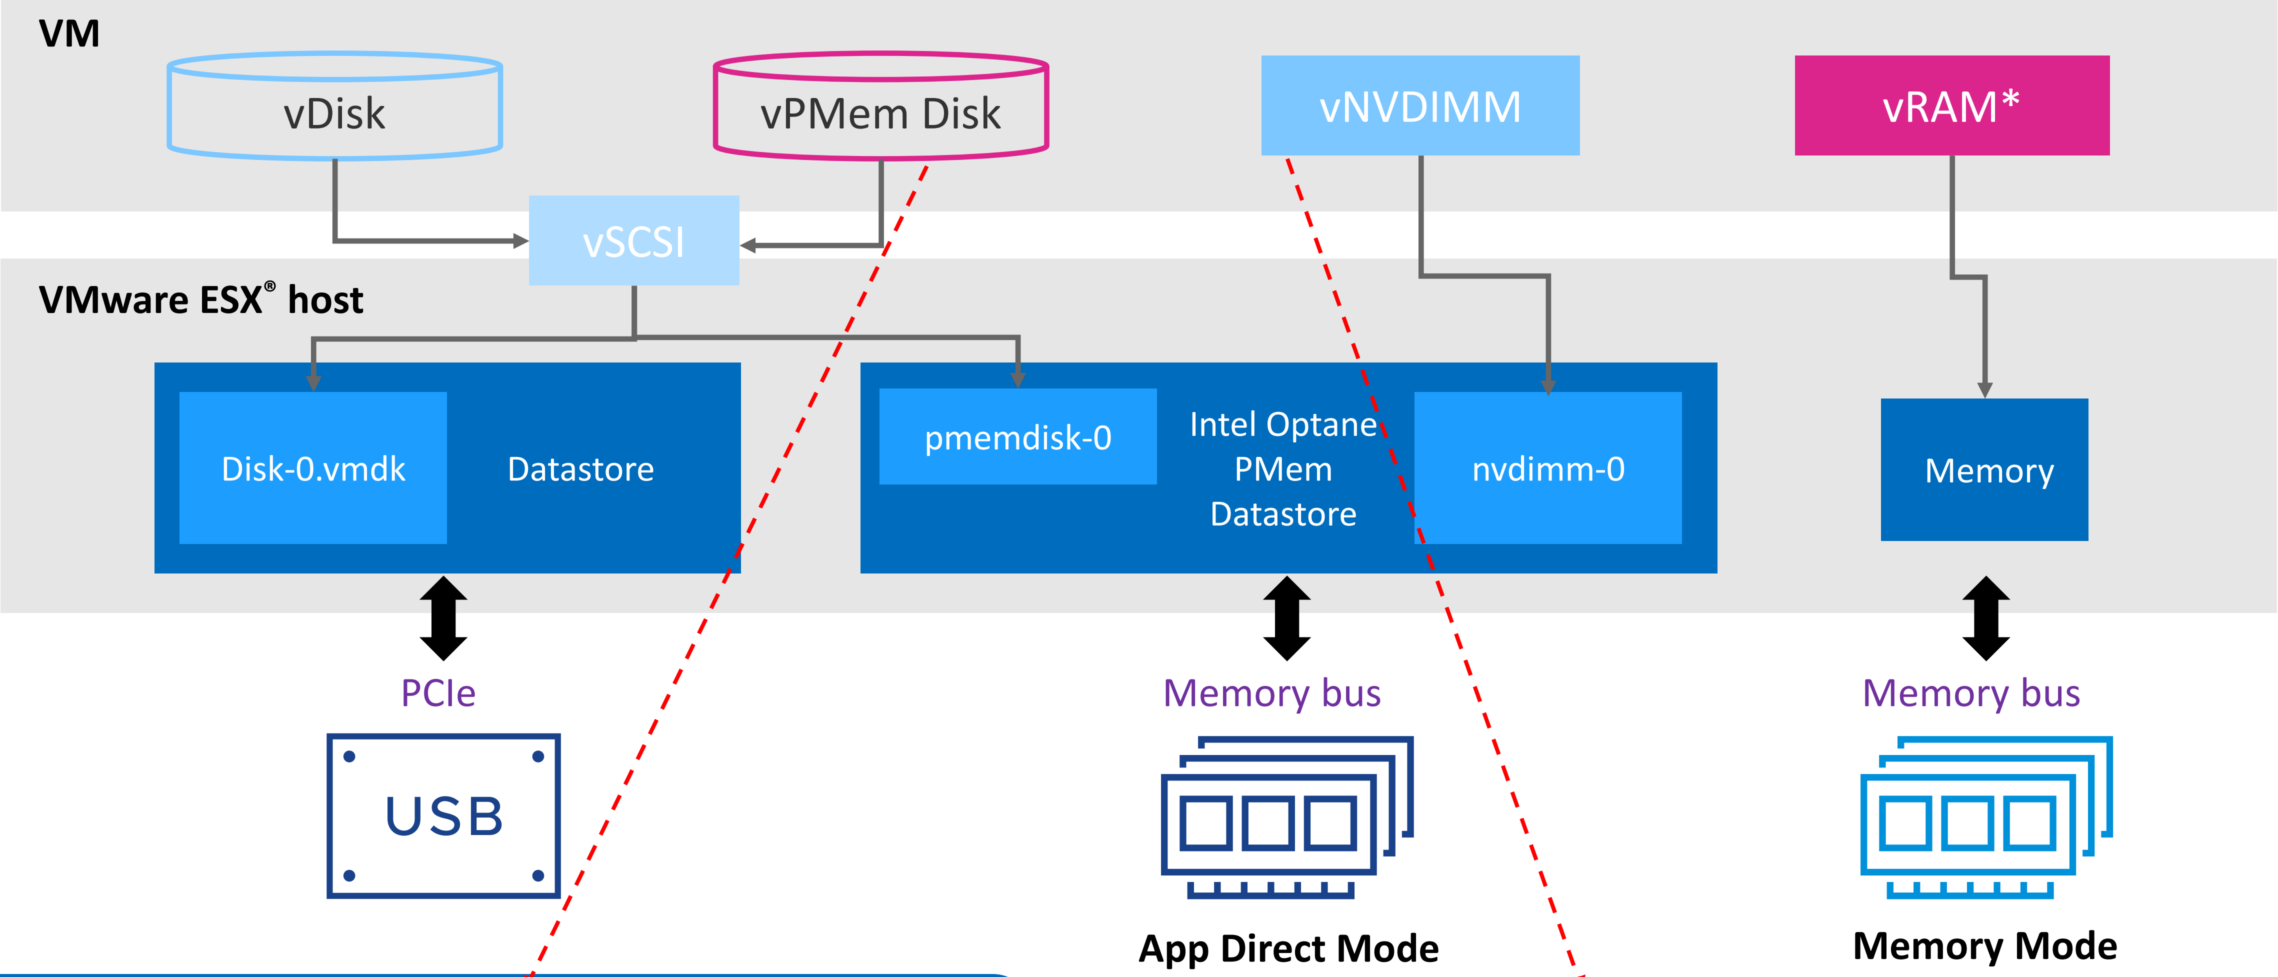

- Virtual PMem Disk (vPMEMDisk)
  - presents NVDIMM capacity as a host local datastore
  - vmdk on PMem datastore backed by PMem
  - GOS access regular SCSI/NVMe device (block device)
  - NO GOS changes required

Virtual PMem (vPMEM)
- exposed as byte addressable PMem
- exposes NVDIMM capacity to VM through a vNVDIMM device
- GOS directly use it as Block device / DAX mode

Oracle Workloads using vSphere Persistent Memory – Use Cases

Oracle Workloads using Persistent Memory – Challenges & Solution

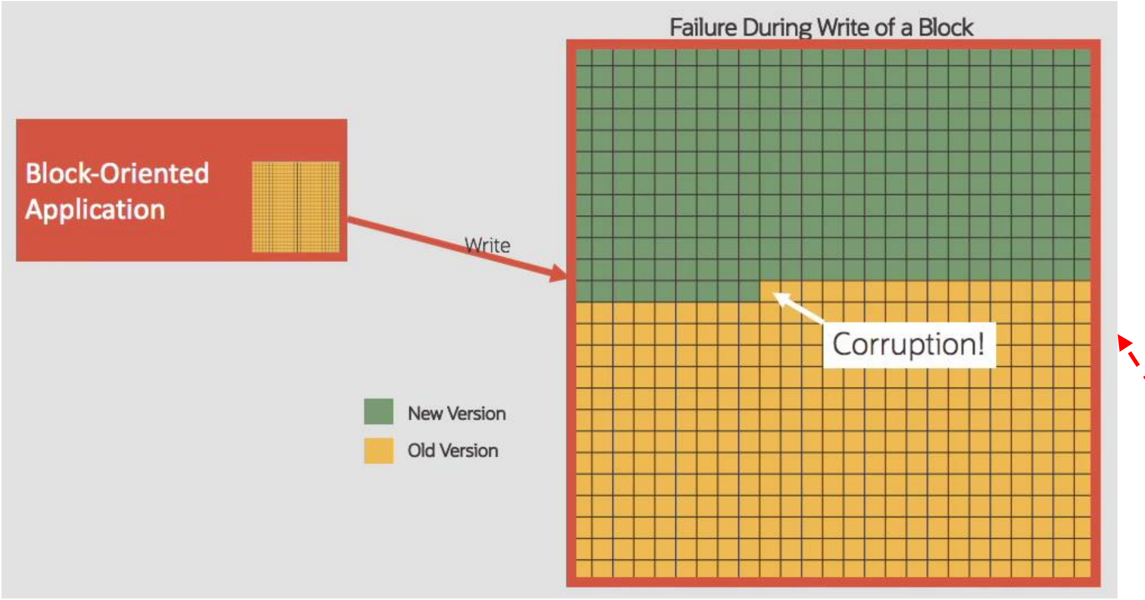

https://blogs.oracle.com/database/post/persistent-memory-primer

File Systems and Devices on Persistent Memory (PMEM) in Database Servers May Cause Database Corruption (Doc ID 2608116.1)

- Challenges with Native PMem
  - natively operates byte-by-byte - not block
    - data persistence is 8-bytes at a time
- Oracle DB's - based on BLOCK construct (2k-32k block size)
- Writing one 8k Oracle block = Writing 1024 x 8-Bytes PMEM chunks
- Power failure / other abnormal condition
  - can result in fractured / torn block portions of the block containing old data & other portions with new data
  - App level change needed to tolerate fracturing of blocks else block corruption

Intel Optane DC PMM in App Direct mode backed Oracle 21c Persistent Memory Filestore on VMware Platform

Use Case 1 –Accelerating Redo log files using Oracle 21c Persistent Memory Filestore on vSphere 7.x using Intel DC Optane PMM

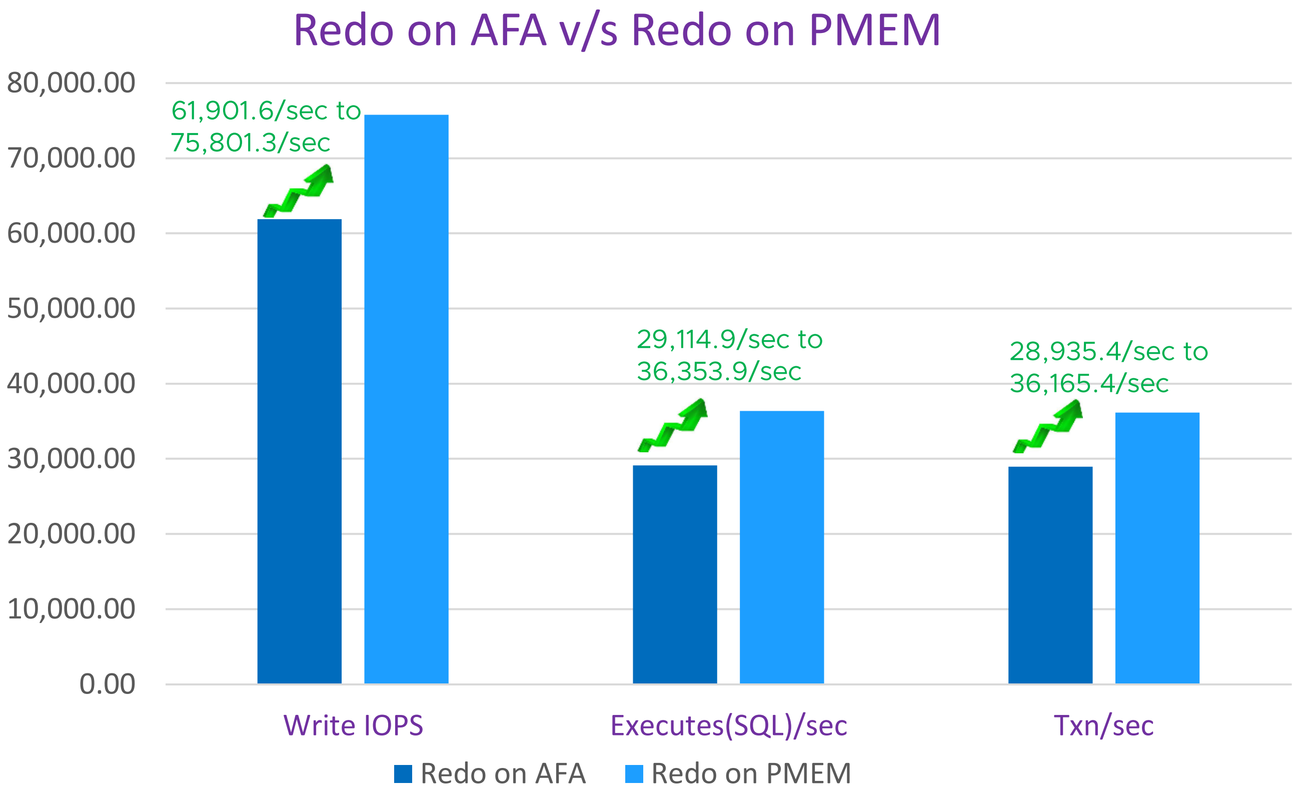

- VM on ESXi 7.0.2
  - 24 vCPU's, 256GB memory
  - OEL 8.5 UEK
  - Oracle 21.5 version with SGA = 96G, PGA=20G
  - Oracle Standalone DB with ASM & ASMLIB
    - Oracle on VMware Best Practices Followed

- Load Generator chosen as SLOB 2.5.4.0
  - UPDATE_PCT=100
  - SCALE=90G
  - WORK_UNIT=3
  - REDO_STRESS=HEAVY
    - Work Unit minimum size chosen to drive most amount of IO with heavy stress on redo to study performance improvement using PMEM

https://blogs.vmware.com/apps/2022/02/acclerate-oracle-redo-pmem-filestore-intel-dcpmm-vmware.html

Remember – Any performance data is a result of the combination of hardware configuration, software configuration, test methodology & test tool, workload profile used in the testing

Use Case 1 –Accelerating Redo log files using Oracle 21c Persistent Memory Filestore on vSphere 7.x using Intel DC Optane PMM - other metrics

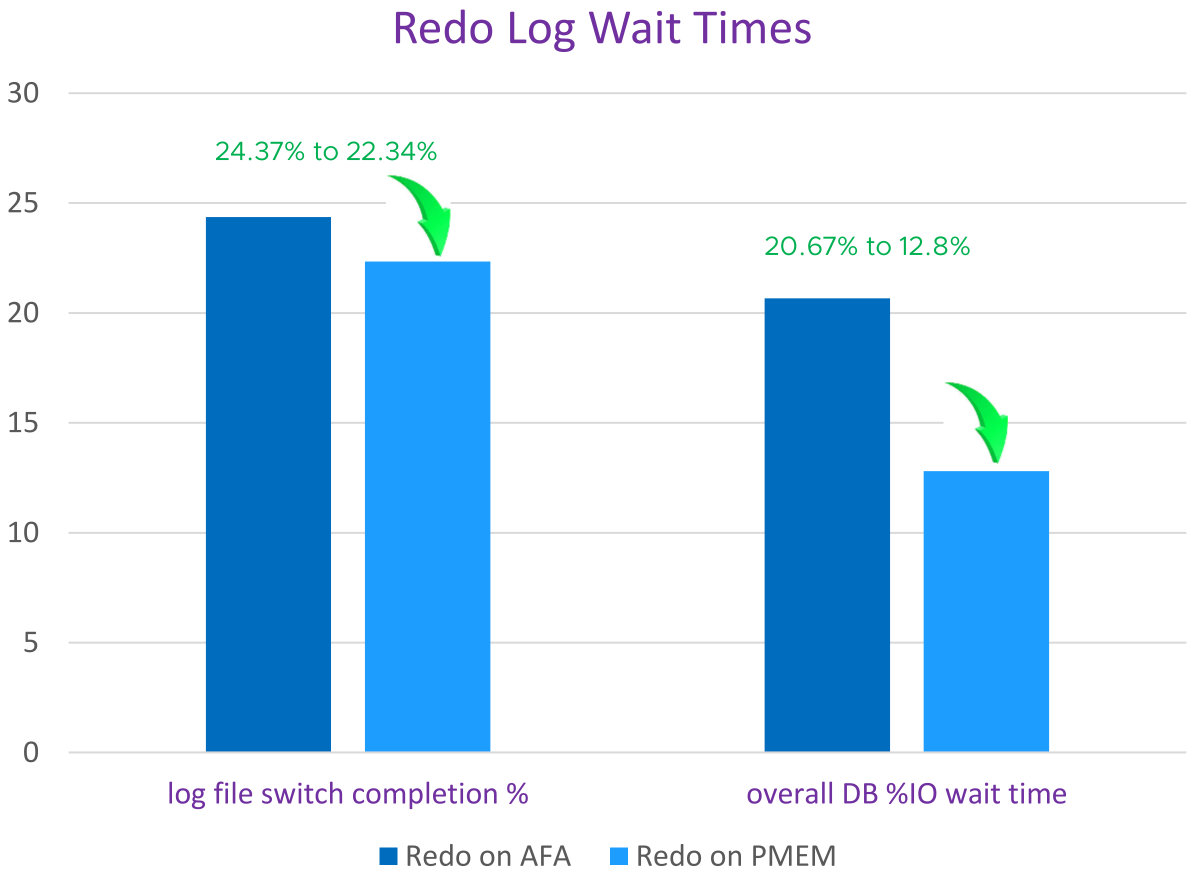

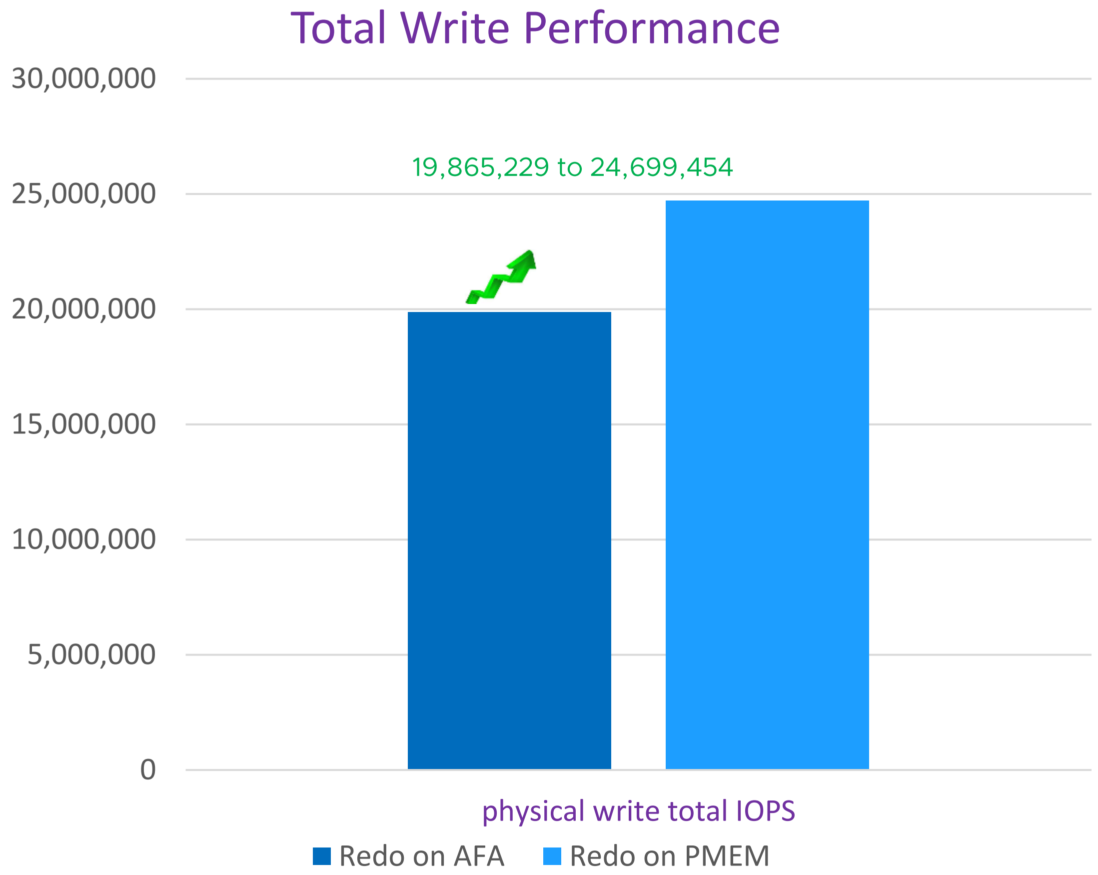

Redo Log on AFA v/s Redo Log on Oracle PMEM Filestore

- 'redo writes' - increased from 26,974 to 82,605 - indicative of IO load

https://blogs.vmware.com/apps/2022/02/acclerate-oracle-redo-pmem-filestore-intel-dcpmm-vmware.html

KEEP CALM 'CAUSE YOUR MILEAGE MAY VARY

Remember – Any performance data is a result of the combination of hardware configuration, software configuration, test methodology & test tool, workload profile used in the testing

Use Case 2 –Accelerating Oracle Flash Cache using Oracle 21c Persistent Memory Filestore on vSphere 7.x using Intel DC Optane PMM

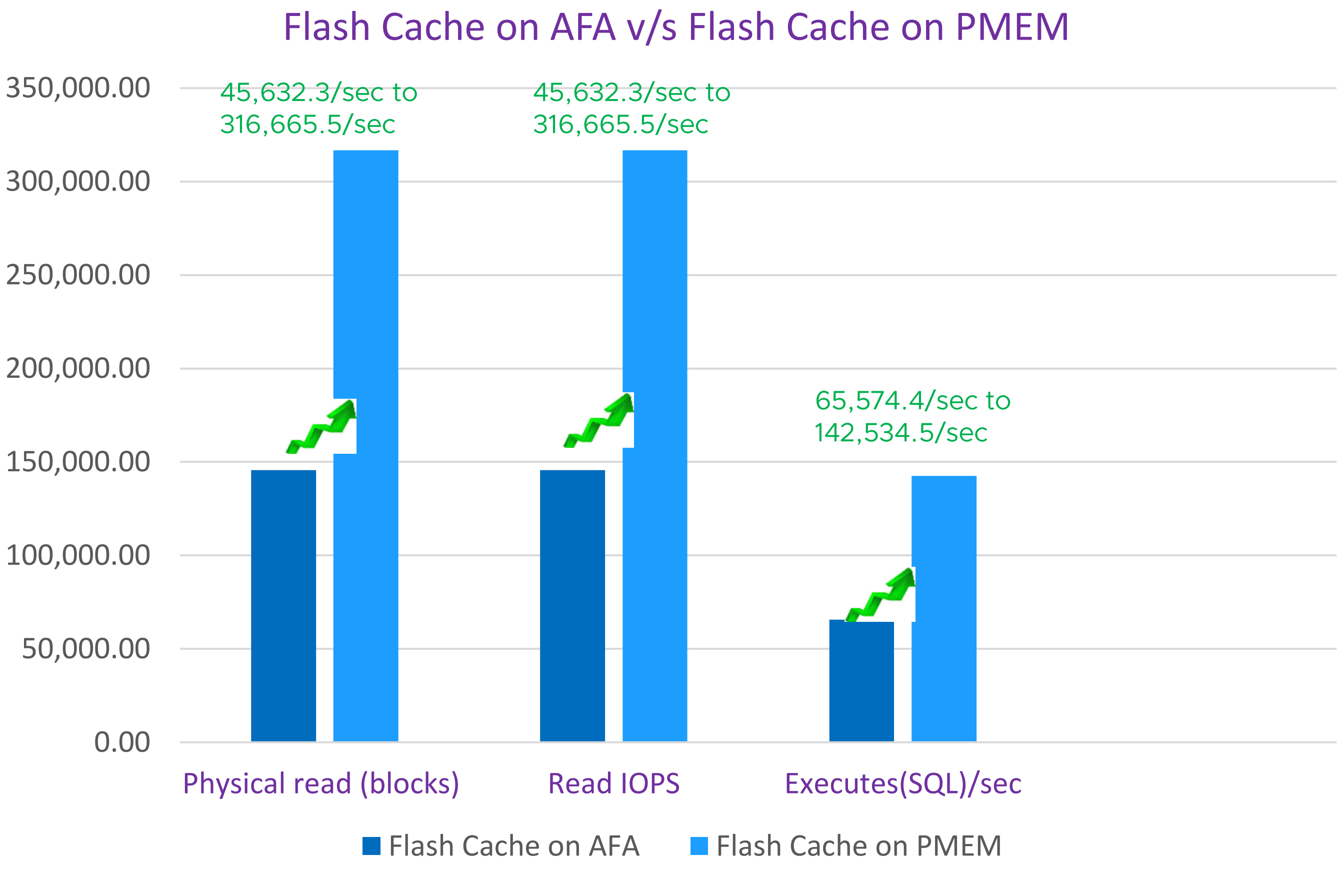

- VM on ESXi 7.0.2
  - 24 vCPU's, 256GB memory
  - OEL 8.5 UEK
  - Oracle 21.5 version with SGA = 128G, PGA=20G
  - DB_FLASH_CACHE_SIZE = 384G
  - Oracle Standalone DB with ASM & ASMLIB
    - Oracle on VMware Best Practices Followed

- Load Generator chosen as SLOB 2.5.4.0
  - UPDATE_PCT=0
  - SCALE=200G
  - WORK_UNIT=32
  - REDO_STRESS=LITE
    - update pct =0 as it is a heavy READ only test to study performance improvement of Flash Cache using PMEM

https://blogs.vmware.com/apps/2022/04/oracle-flashcache-pmem-filestore-intel-dcoptane-pmem-vmware.html

KEEP CALM 'CAUSE YOUR MILEAGE MAY VARY

Remember – Any performance data is a result of the combination of hardware configuration, software configuration, test methodology & test tool, workload profile used in the testing

Use Case 2 –Accelerating Oracle Flash Cache using Oracle 21c Persistent Memory Filestore on vSphere 7.x using Intel DC Optane PMM

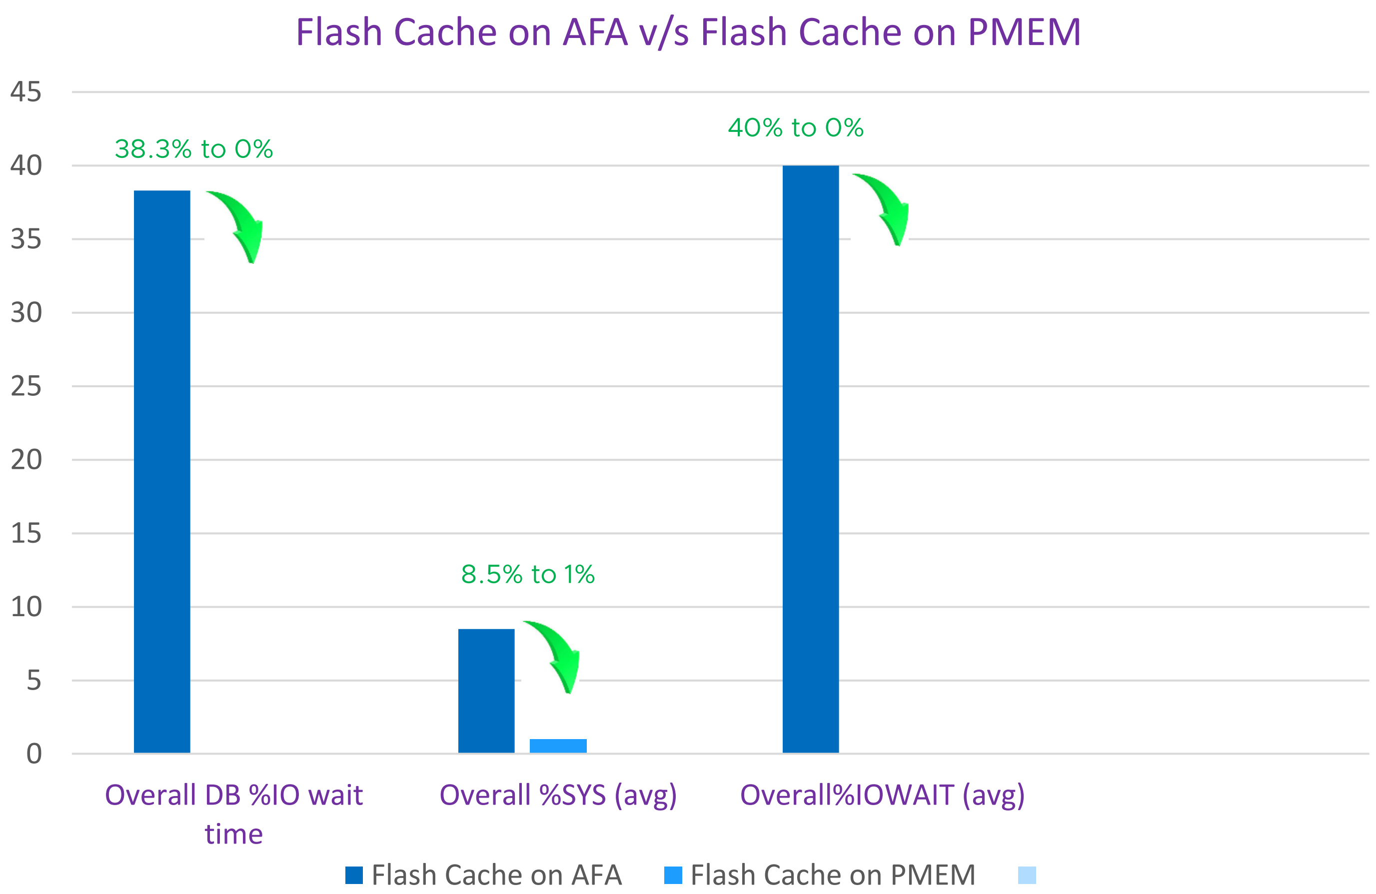

- Flash Cache on AFA v/s Flash Cache on Oracle PMEM Filestore
  - Overall database '%IO wait time' - reduced from 38.3% to 0%
  - OS CPU %sys - reduced from average 8.5 to average 1%
  - OS %iowait - reduced 40% average to 0%

https://blogs.vmware.com/apps/2022/04/oracle-flashcache-pmem-filestore-intel-dcoptane-pmem-vmware.html

KEEP CALM 'CAUSE YOUR MILEAGE MAY VARY

Remember – Any performance data is a result of the combination of hardware configuration, software configuration, test methodology & test tool, workload profile used in the testing

SNIA PERSISTENT MEMORY + COMPUTATIONAL STORAGE SUMMIT 2022

Project Capitola

The Solution: VMware Value Proposition

Addressing the growing need of in-memory computing!

WHAT CUSTOMERS HAVE

Inflexible Hardware

Over-Provisioned Memory

Project Capitola

WHAT CUSTOMERS WANT

Flexible Software

Defined Memory

Memory is the highest cost factor (50-80%)

Limited Inflexible Capacity

Failure prone

Software Defined Memory Tiering with uniform consumption model

Cluster level memory pooling, visibility and availability

Fully Integrated with VMware Stack

Reduce TCO of memory

Expand Capacity on demand

Faster failure recovery

The Solution: How Does it Work?

Phase-1: Local tiering with cluster support

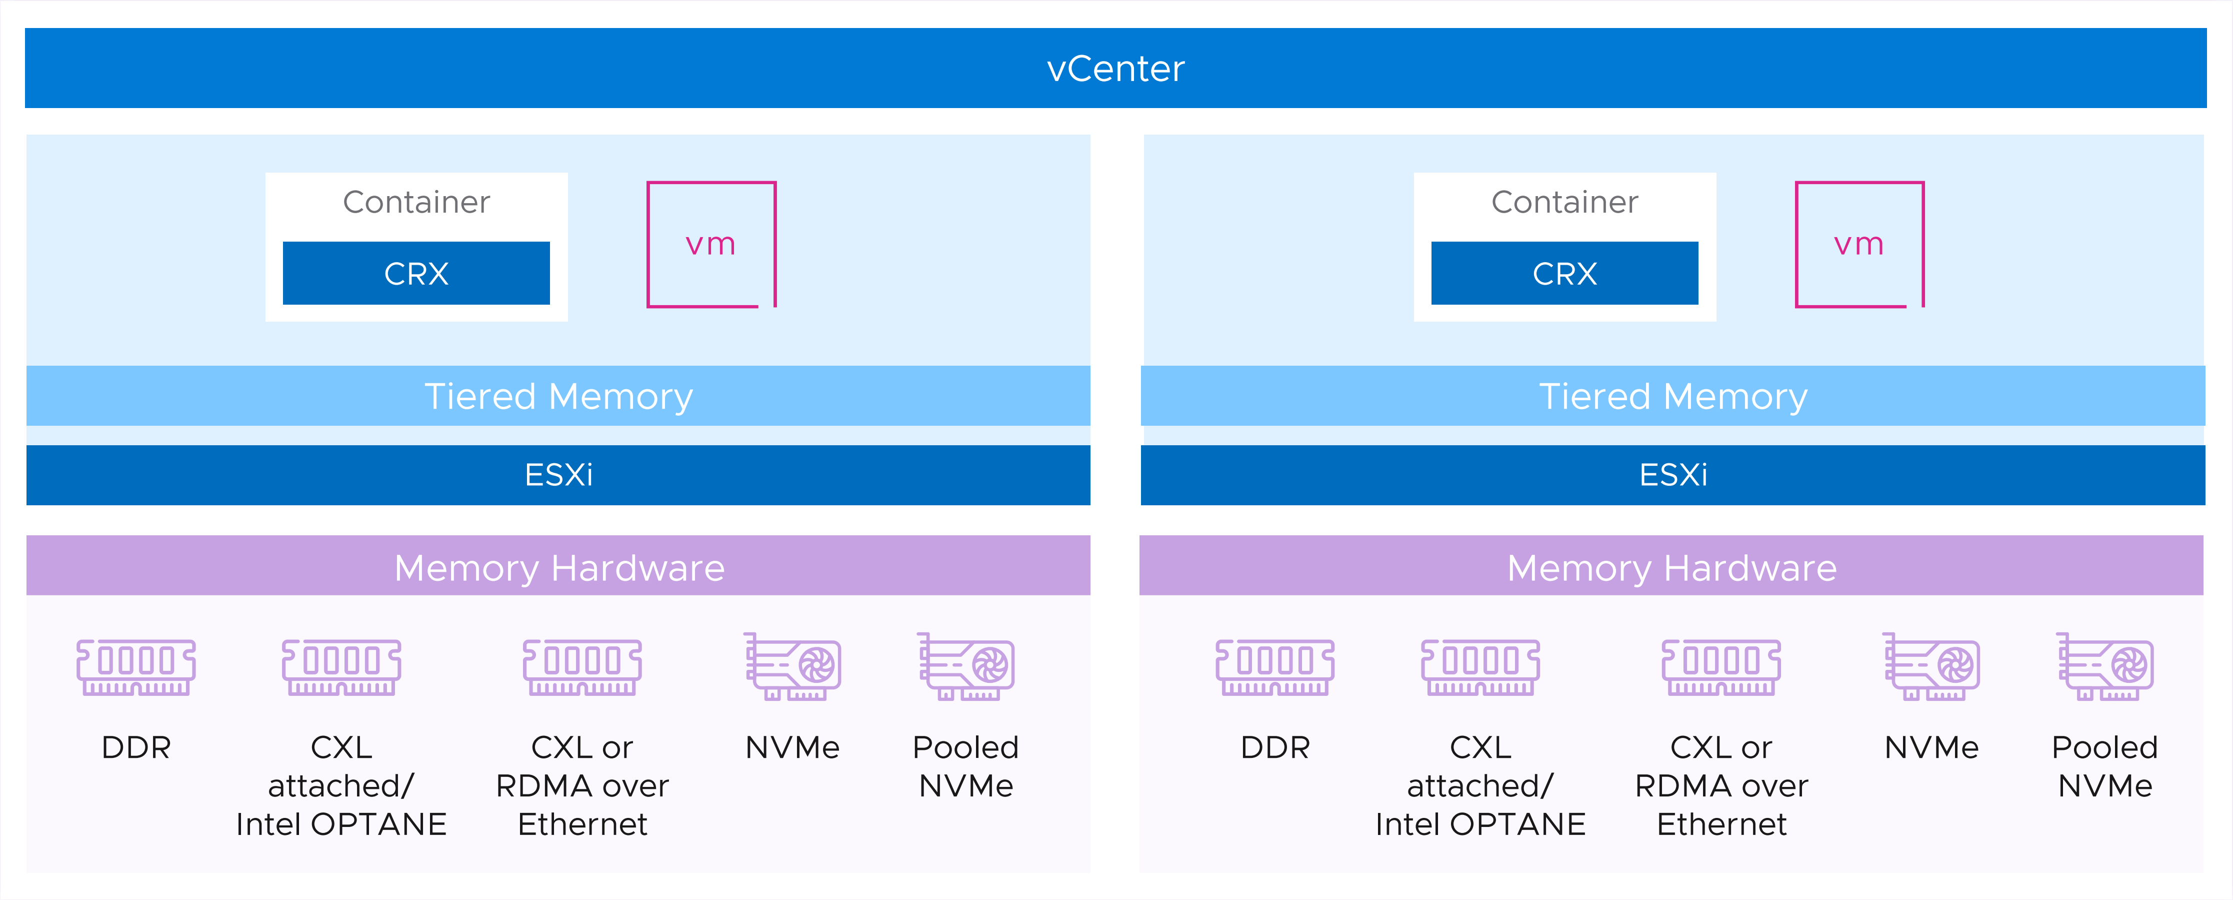

The Solution: How Does it Work?

Phase-2: Cluster wide pooling

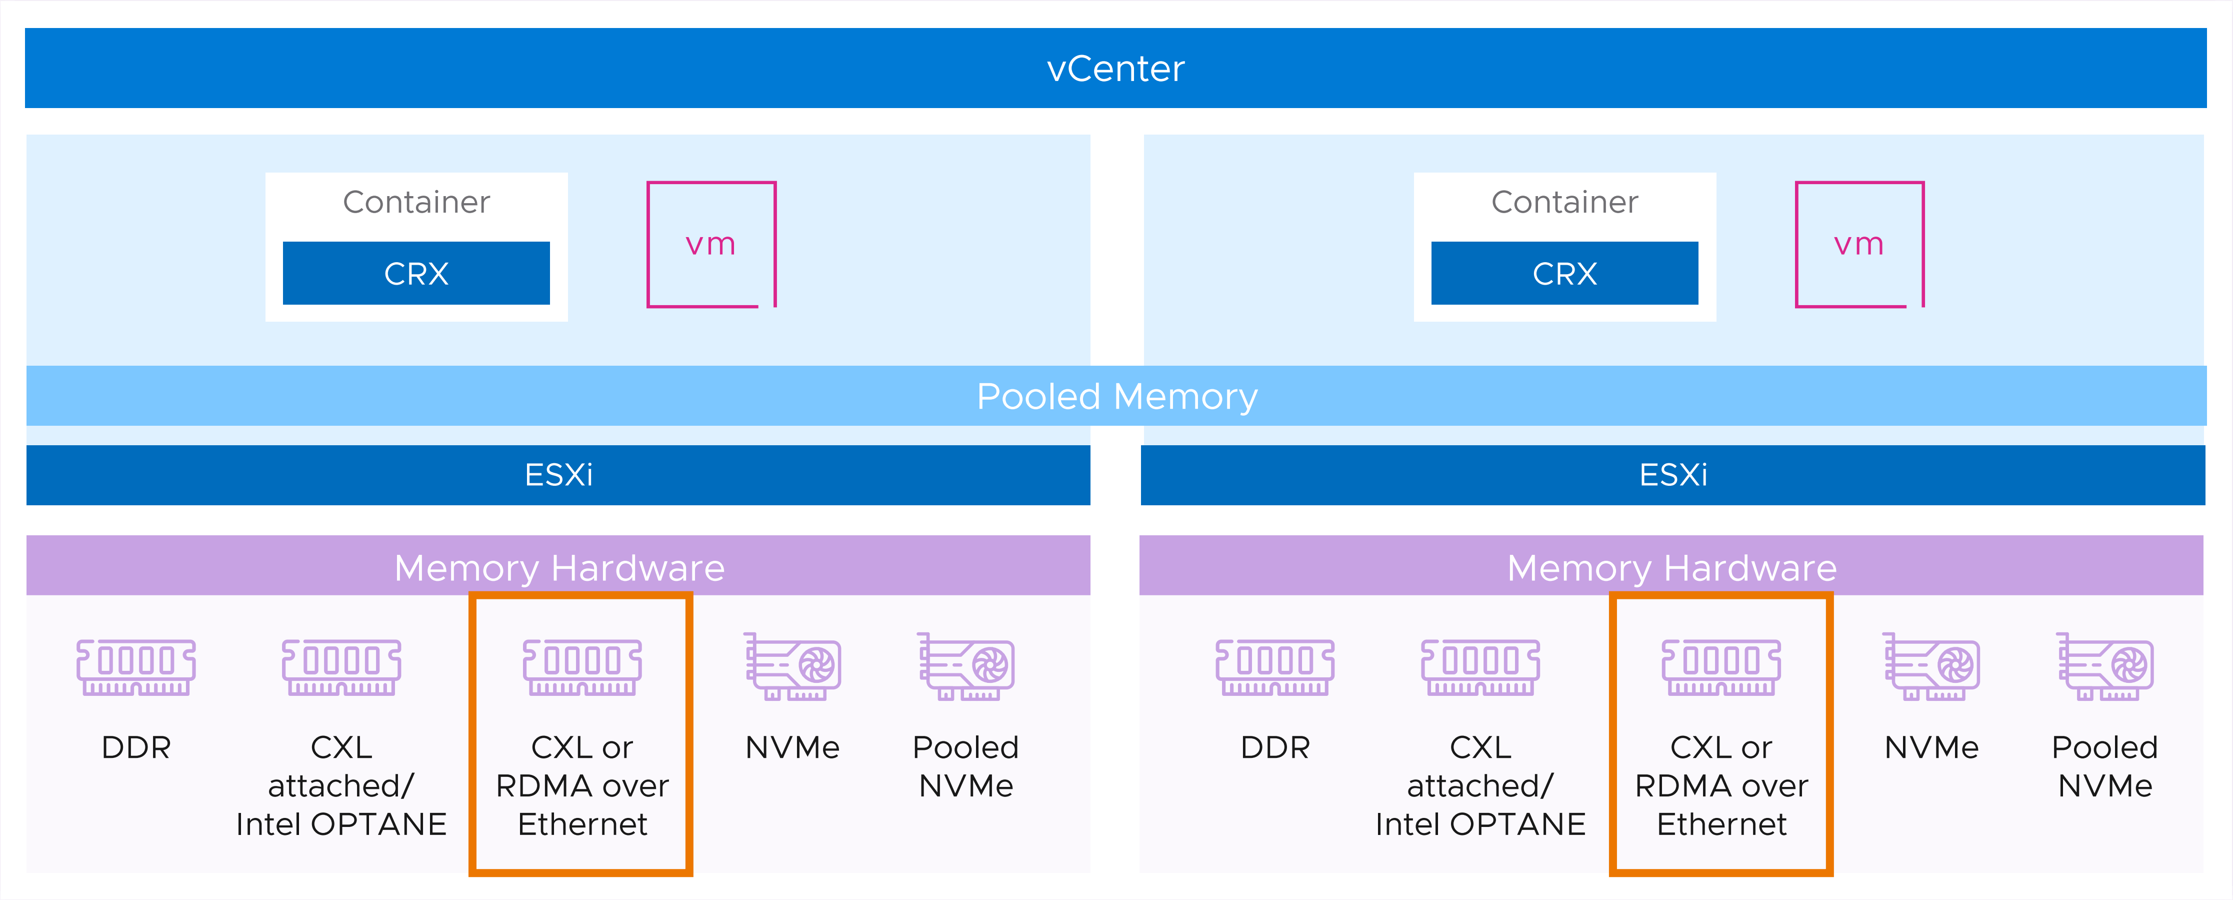

Project Capitola Enables Transparent Tiering

Built in vSphere, requires no modifications to applications or Operating Systems

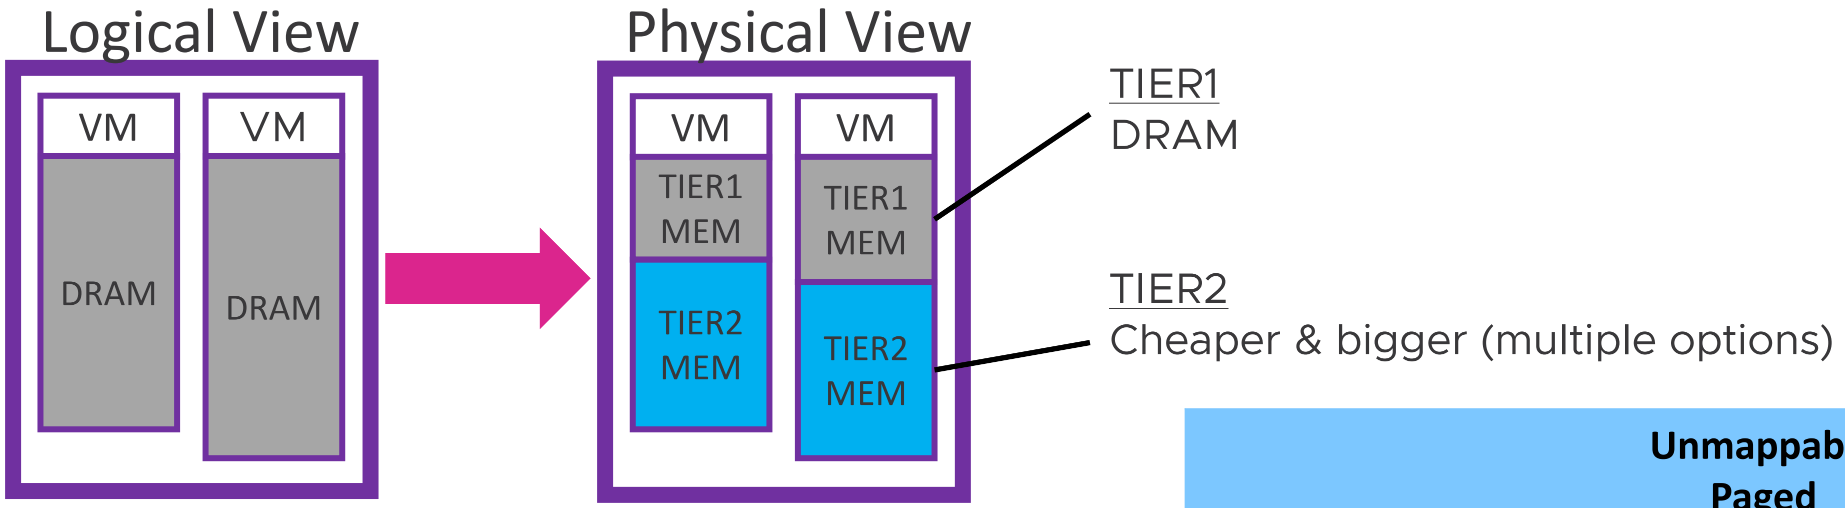

	Unmappable Paged	Mappable Byte-Addressable
Locally attached	Fast NVMe SSD	High-density mem
Pooled	DRAM High-density mem Fast NVMe SSD	DRAM High-density mem

- ✓ vSphere decides what tiers to use and when
- ✓ Memory available to the host is sum of all memory tiers

Key Takeaways

Key Takeaways

Larger datasets and real-time analytics requiring more **performance** are driving the trend for **larger memory**, and need to address capacity and density requirements

Memory is costly, hard to scale and manage, requiring new software -based solutions

Applications are also becoming more real-time, and benefit from being in memory. Mission critical applications (Oracle, SQL, SAP HANA) also benefit from such innovations

VMware's Project Capitola is bringing software defined, scalable memory solutions to address customers and ISV challenges for today and the future!

Oracle on VMware Collateral – One Stop Shop

- All Oracle on vSphere white papers including Oracle on VMware Hybrid Multi-Clouds (vSphere / vSAN / vVols / VMware Clouds) Best practices, Deployment guides, Workload characterization guide can all be found in the url below

Oracle on VMware Collateral – One Stop Shop

https://blogs.vmware.com/apps/2017/01/oracle-vmware-collateral-one-stop-shop.html

SNIA PERSISTENT MEMORY + COMPUTATIONAL STORAGE SUMMIT 2022
VIRTUAL EVENT • MAY 24-25, 2022

Please take a moment to rate this session.

Your feedback is important to us.