

The SNIA logo consists of a small square icon with a yellow-to-orange gradient, followed by the letters 'SNIA' in a bold, purple, sans-serif font.

SNIA

PERSISTENT MEMORY

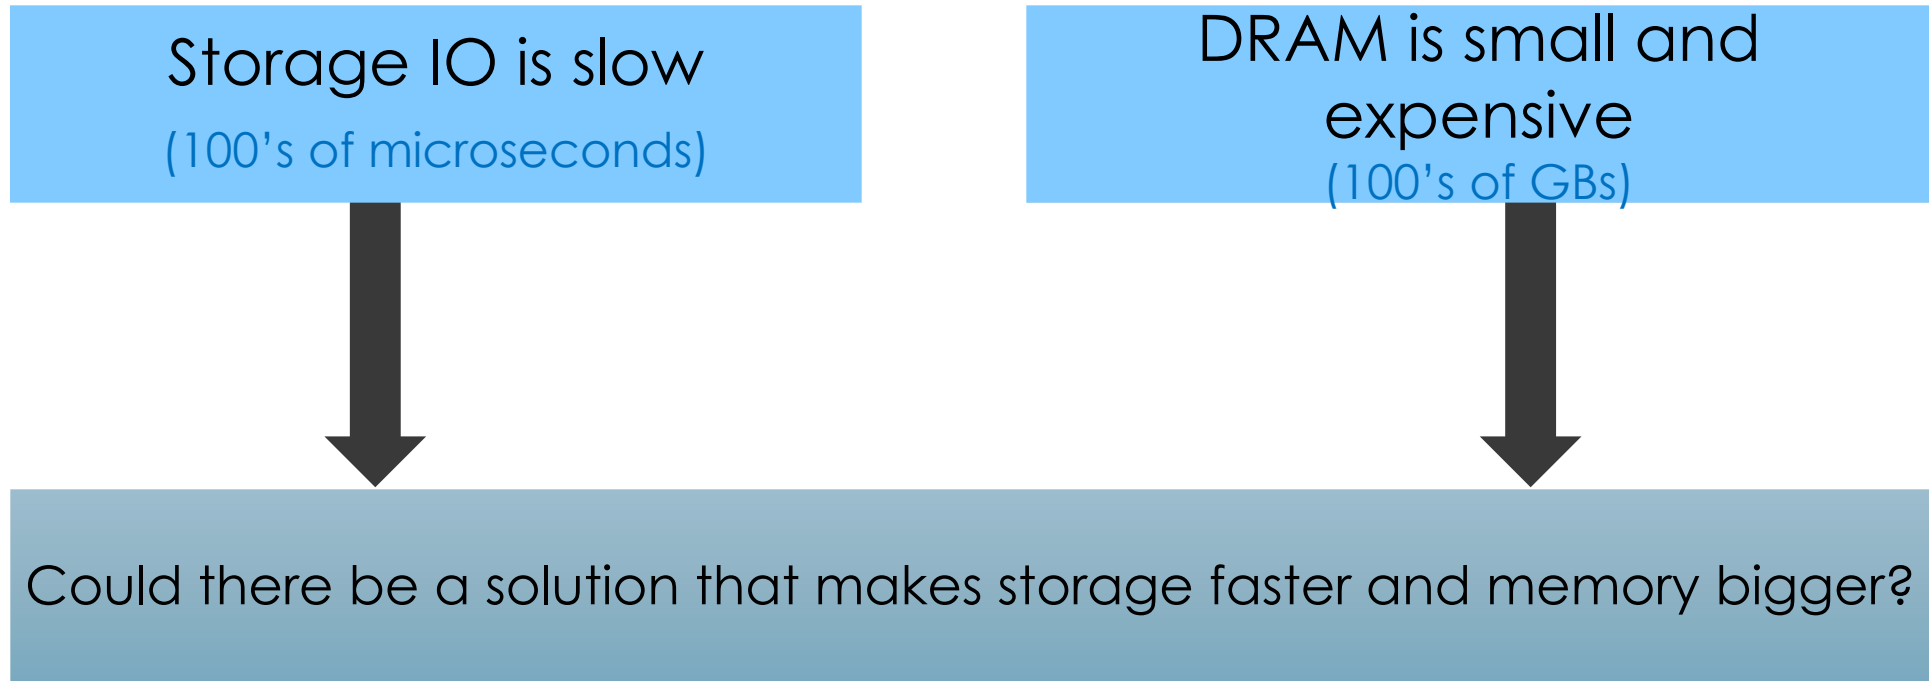
**PM** SUMMIT

JANUARY 23, 2020 | SANTA CLARA, CA

# Memory at Storage Scale Storage at Memory Speed

Charles Fan, Co-founder and CEO, MemVerge

# Big Data and Fast Data



Machine Learning, Big Data and IoT demand  
nanosecond speed + petabyte scale data  
infrastructure

# Persistent Memory


- Intel Delivered Optane DC Persistent Memory based on 3D XPoint tech in Q2 2019.
  - Revenue projected to reach \$3.6B by 2023/2024
  - Additional major vendors to join the foray by 2022
- Storage at Memory Speed
  - 100-250ns R/W latency
- Memory at Storage Scale
  - 6TB/server
  - Persistent



**1000X**  
FASTER  
THAN NAND

**1000X**  
ENDURANCE  
OF NAND

**10X**  
DENSER  
THAN CONVENTIONAL MEMORY



**3D XPoint**

# World's First Memory Converged Infrastructure

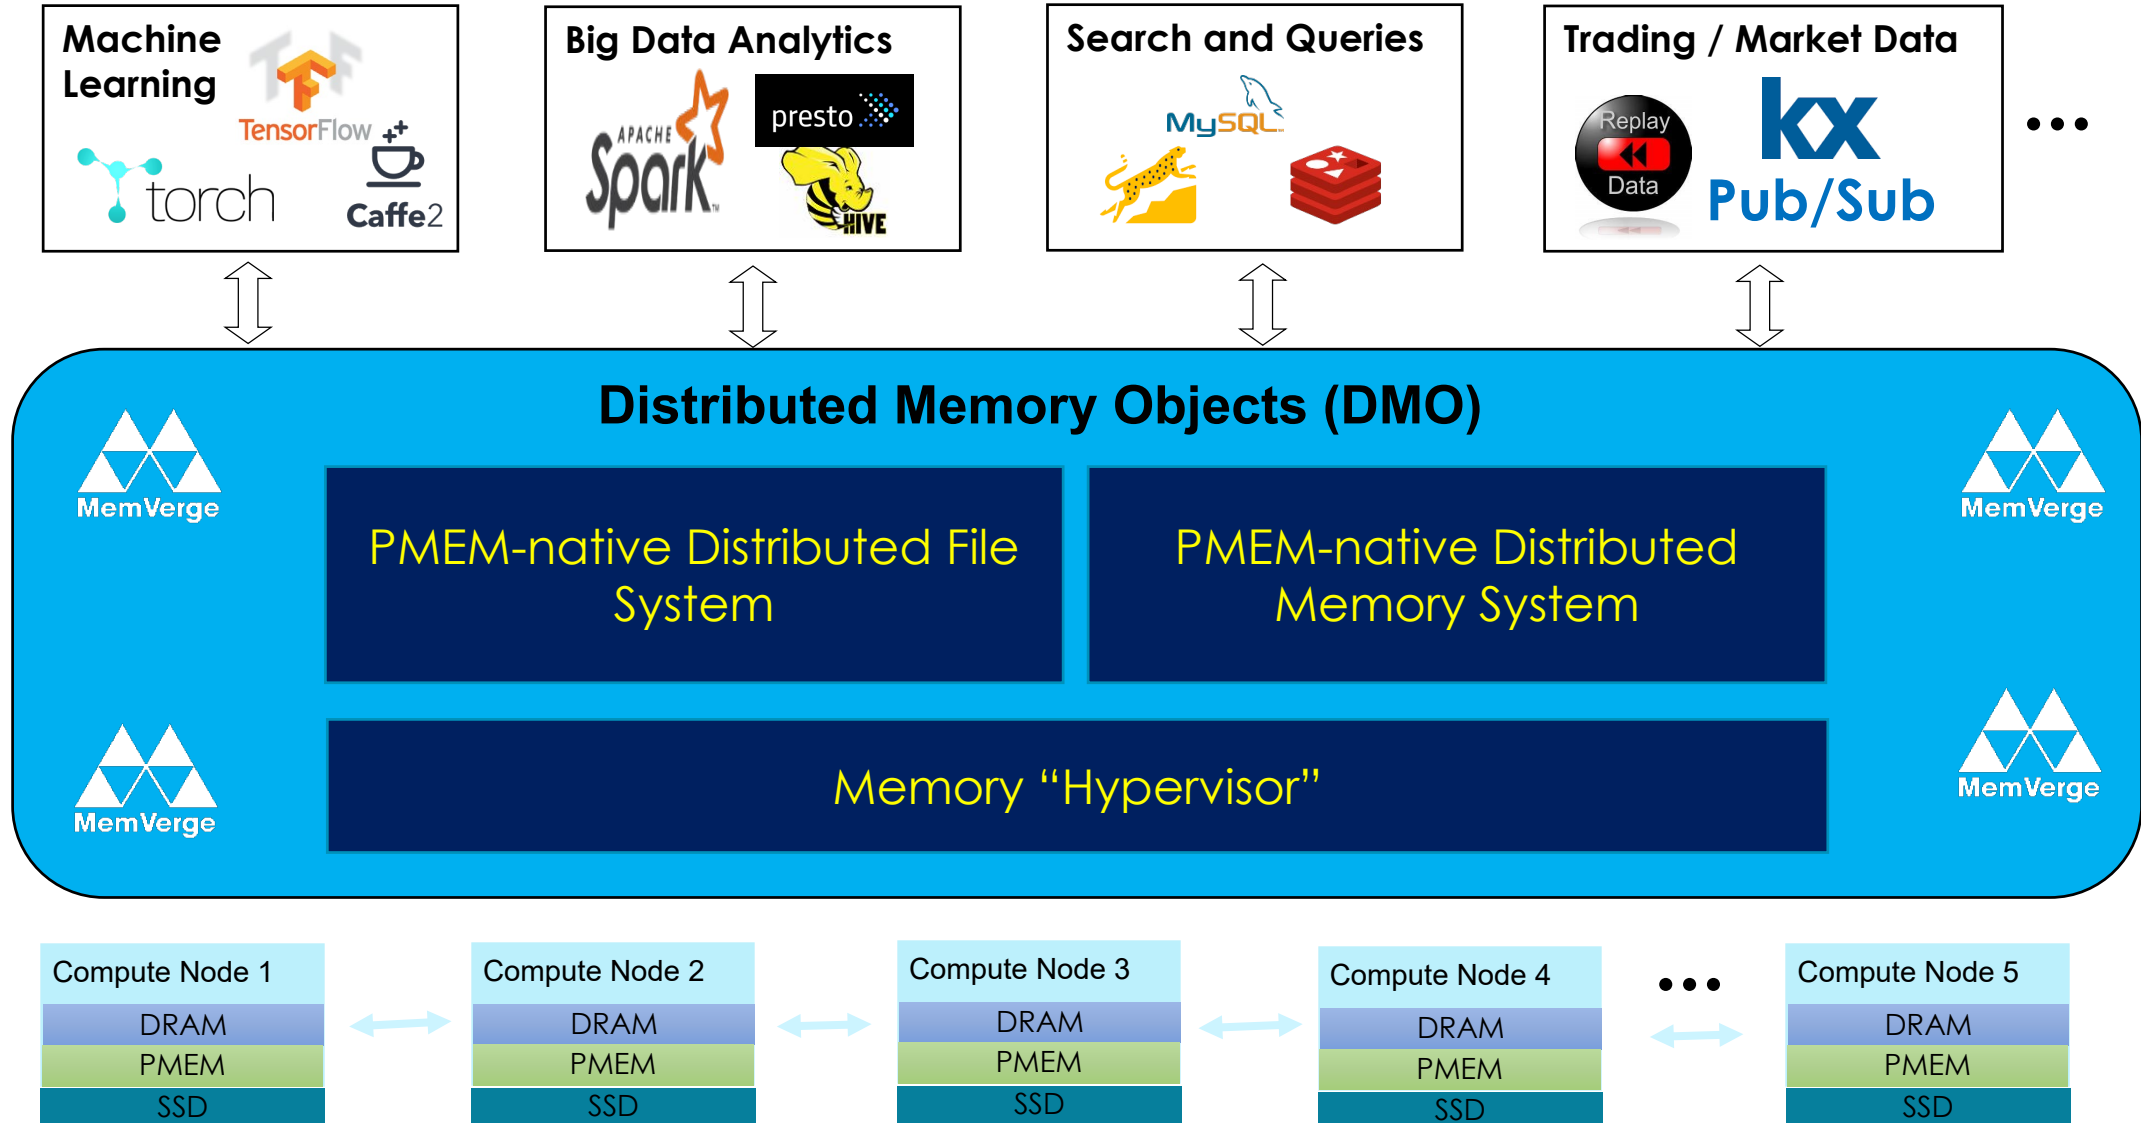
MemVerge software leverages Storage Class Memory technology to deliver larger memory and faster storage to applications without requiring application rewrites

PMEM-native Distributed  
File System

PMEM-native Distributed  
Memory System

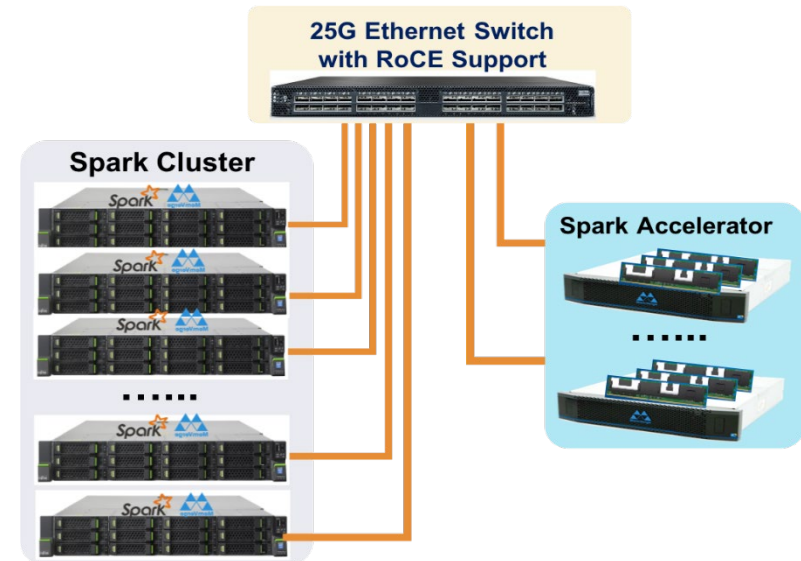
Memory "Hypervisor"

# Architecture and Use Cases



# Use Case: Big Data Analytics with Spark

- Problem
  - Spark SQL Out of DRAM
  - Disk I/O too slow
  - Data spill degrades performance
  - Local SSDs wear out by frequent intermediate data writes
- Solution
  - Adding MemVerge DMO to the Spark cluster accelerates the entire cluster
  - Moving intermediate state off Spark Elastic Computing nodes increased the cloud elasticity of the solution



**5X**

**Terasort Speed**

**7X**

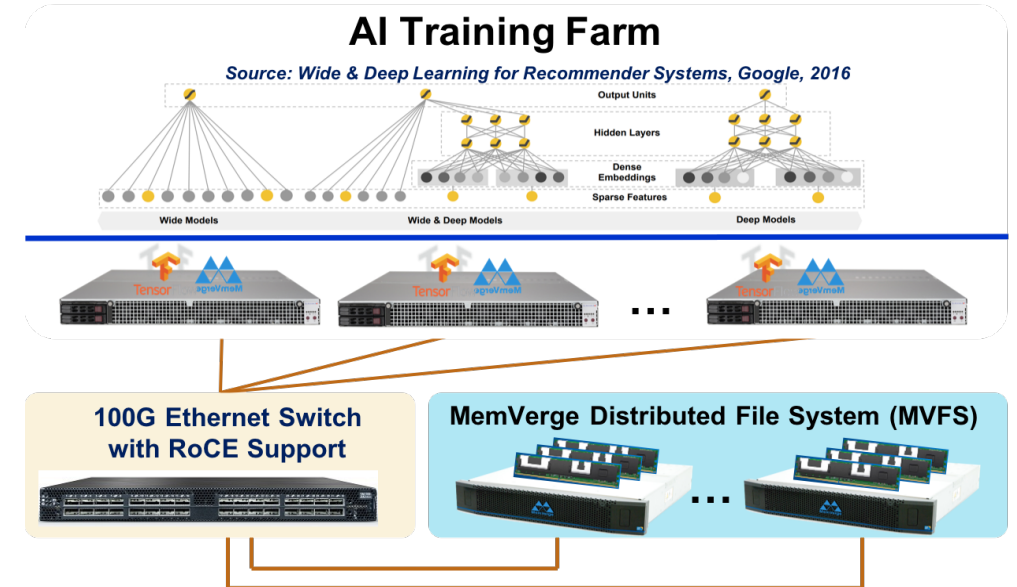
**RDD Caching Speed**

**100%**

**Cloud Elasticity**

# Use Case: AI Training with Checkpointing

- Problem
  - Model training takes a long time to complete for large datasets
  - Failure recovery is painful without frequent checkpointing
  - Data preprocessing and importing can take a long time
  - Delayed model deployment
- 
- Solution
  - MemVerge DMO, powered by Optane DC persistent memory, improves checkpointing speed and data loading speed




up to **6X**  
Training Speed

up to **350X**  
Data Import Speed

**Instant**  
Checkpoint Recovery

# Use Case: High Performance Financial Systems




- High Performance Market Data Service
  - Enables high bandwidth and low latency market data replay via large memory-mapped files
- Pub/Sub Service
  - Next-gen Pub/Sub architecture with RDMA-interconnected Persistent Memory
- KDB Acceleration
  - Improve the performance of existing deployments of KDB



# MemVerge Vision

- By 2025, Persistent Memory will be mainstream. Data Infrastructure will be memory-centric.
- Performance-tier storage will be replaced by Memory Converged Infrastructure (MCI) co-located with compute. Memory capacity will be expanded by the same MCI layer.
- MemVerge aspires to be a leader in MCI.



# Thank you

Please visit [www.snia.org/pmsummit](http://www.snia.org/pmsummit) for presentations