



Flash Memory Summit

Memory Converged Infrastructure with Persistent Memory

Charles Fan
Co-Founder & CEO
MemVerge



Age of Machine-Generated Data Ushered in by AI, Big Data and IOT

Pain Points:

Storage IO is slow

DRAM is small and
expensive



Could there be a solution that makes storage faster and memory bigger?



Introducing Persistent Memory

Flash Memory Summit

- Intel delivered Optane DC Persistent Memory in Q2 2019
- Volatile Memory Mode
 - Bigger capacity than DRAM (3TB/socket in 2019, 6TB/socket in 2020)
- Block Storage Mode
 - Lower latency than NVMe SSDs (100-250ns)
- App Direct Mode
 - New Persistent Memory programming model





Flash Memory Summit

World's First Memory Converged Infrastructure



Distributed Memory Objects (DMO)


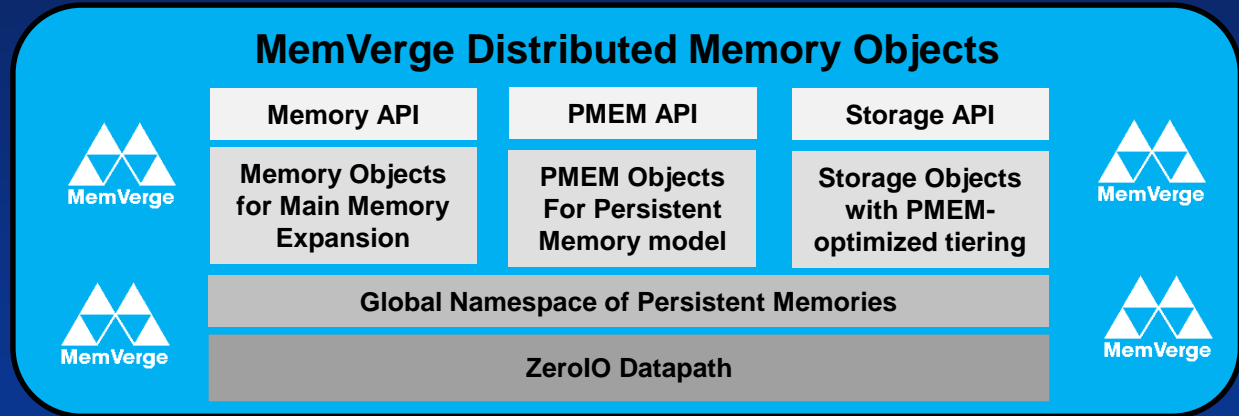

DMO is a Distributed Software System supporting both memory and storage APIs, served from a single pool of persistent memory

1. Fast Storage
2. Big Memory
3. No application programming model changes



Flash Memory Summit

Memory Converged Infrastructure Platform

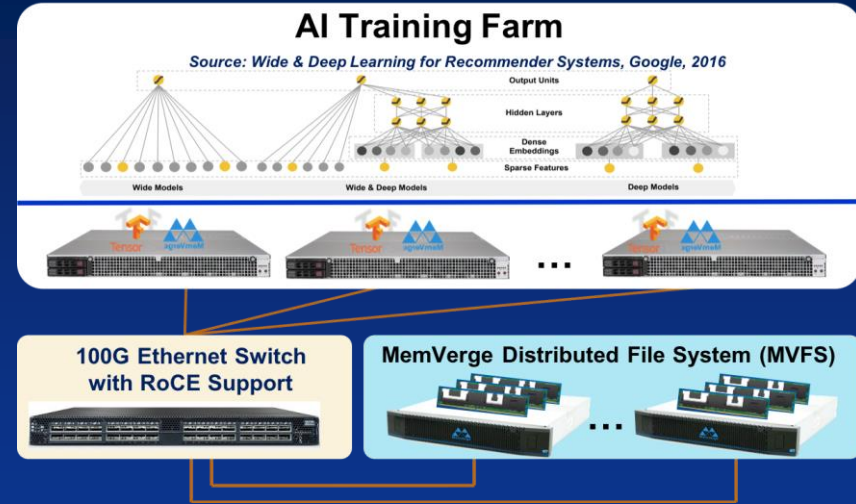




AI Training with Checkpointing

Flash Memory Summit

- 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

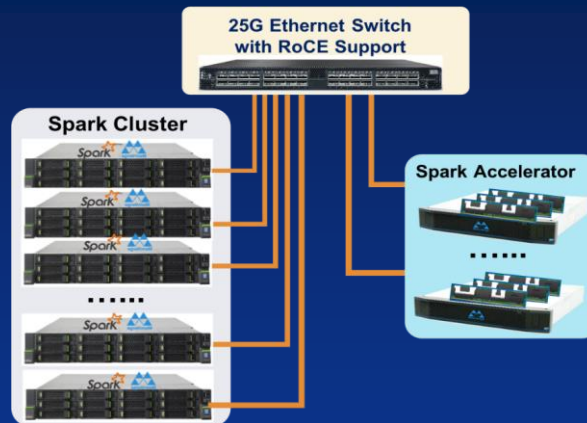
Instant
Checkpoint Recovery



Big Data Analytics with Spark

Flash Memory Summit

- 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



Flash Memory Summit

Memory Converged Infrastructure

1 μ s

Access Latency



MemVerge

up to **768TB**

Total Cluster
Memory

up to **10M+ IOPS**

Per Node



Flash Memory Summit

Thank You!