### Corporate Presentation

2020



### Introducing MemVerge

Founded by:

#### Shuki Bruck

XtremIO co-founder and Caltech professor



#### **Charles Fan**

VMware storage BU leader and creator of VSAN

#### Yue Li

Caltech post-doc and top researcher on non-volatile memory







### Introducing MemVerge

World-class team in Silicon Valley assembled from:



### DELLEMC





## **Digital Transformation is Widespread**

#### DX Maturity Distribution

IDC MaturityScape Benchmark: Future Enterprise — Maturity Distribution Across the Stages



91.1% of enterprises undergoing DX in the next three years More data-centric business models will drive AI/MI -infused analytics Performance and availability implications for enterprise storage Market evolution will drive demand for persistent memory technologies



#### **Realtime Workloads are on the Rise**



#### Worldwide Real-Time Data and Share, 2015-2024

Worldwide, data is growing at a 26.0% CAGR, and in 2024 there will be 143 zettabytes of data created

By 2021, 60-70% of the Global 2000 will have at least one mission-critical real-time workload



## PM Revenue Forecast, 2019 - 2023

Byte-Addressable Only



REVENUE (\$M)

MemVerge

IDC: Digital Transformation Driving New "Big Memory" Requirements

### Long Term Forecast

Emerging memories are well on their way to reach \$36 billion of combined revenues by 2030.

3D XPoint memory's sub-DRAM prices are expected to drive revenues to over \$25 billion by 2030



Emerging Memories Find Their Direction: Objective Analysis and Coughlin Associates



### **Big Memory Definition**



- Enables the ability to run applications in memory for improved performance and efficiency
  - Leverages byte addressable persistent
     memory media
- Includes enterprise-class data services to handle tier 1 availability and management requirements
- Runs as a software-based memory virtualization layer on industry standard hardware
- The technology enabler for missioncritical real-time computing

#### Our BIG MEMORY vision All applications live in memory

MemVerge

# Our Mission Open the door to Big Memory

A world of abundance, persistence and high availability





### Intel<sup>®</sup> optane<sup>™</sup> dc persistent memory Revolutionizing memory

### Why Not All Apps Can Run in Memory yet...

Not plug-and-play

App rewrite needed

Can't share memory

Siloed in servers

**No Data Services** 

Crash recovery is slow









### MemVerge Memory Machine<sup>™</sup>

**Software** Subscription



Virtualizes DRAM & PMEM



**Low Latency** PMEM over RDMA



Plug Compatible No re-writes



#### Memory Data Services Snapshot, Replication, Tiering



MemVerge

### **Big Memory = Optane + Memory Machine**



### Single-Node Memory Machine<sup>™</sup> Implementation





### Multi-Node Memory Machine<sup>™</sup> Implementation



### Financial Services are Killer Apps for Big Memory

Microseconds matter

In-Memory databases used

Data > memory More capacity needed

#### Growing blast zone

Memory data services needed









### **Uses Cases: Real-Time Workloads**

According to IDC, by 2021, 60-70% of the Global 2000 organizations will have at least one mission-critical real-time workload. Below are just a few examples of use cases that are implementing Big Memory now.



Latency-sensitive transactional workloads such as high-frequency trading



Real-time big data analytics in financial services, healthcare, and retail



AI/ML analytics and inferencing like fraud detection and smart security



### **Real Time Market Data Pub/Sub**

#### Objectives

- Market data event stream published to multiple subscribers with the lowest latency
- 2. Achieve fairness between the subscriber processes
- 3. Persist the event stream without incurring significant performance penalties





### **Real Time Market Data Pub/Sub**

#### Solution

- Memory Machine<sup>™</sup> software writes market data event stream to an inmemory bus
- 2. Background process commits the event stream to Persistent Memory synchronously or asynchronously
- 3. The event stream is replicated over RDMA to memory of other servers
- 4. Subscriber processes across all servers read the event stream with low latency



#### **Real Time Market Data Pub/Sub**

#### **Results with 210 Subscribers**



MemVerge

#### **Real Time Market Data Pub/Sub**

Results

0.5 uS

Avg. latency local host

Avg. latency remote host

**3 uS** 



99.9% tail latency



### **In-Memory Database Cloning & Crash Recovery**

#### Problem

- Need to run analytics, reporting or dev/test but concerned about taking performance hit on Primary instance
- 2. Application takes a long time to restart after crash or planned shutdown





### **In-Memory Database Cloning & Crash Recovery**

#### Solution

- Memory Machine takes instant snapshot, as frequently as every 1 minute
- 2. In-Memory Cloning easily creates a read replica of the primary instance using snapshot plus log replay
- 3. Fast restart from the database crash using snapshot plus log replay



#### **In-Memory Database Cloning & Crash Recovery**

Results

#### **Every Minute**

Fine grained Snapshots

Fast

Clone and Crash Recovery

Zero

Performance hit



### **Big Memory AI/ML Training and Inference**

#### Problem

- 1. When data is greater than the size of DRAM, AI/ML performance slows down dramatically
- 2. Memory-intensive inference jobs take a long time to load and restart





### **Big Memory AI/ML Training and Inference**

#### Solution

- Create big memory lakes consisting of DRAM and PMEM to provide capacity needed for all data including models and embeddings
- 2. Fast data recovery and restart by using in-memory data snapshot



### **Big Memory AI/ML Training and Inference**

#### Test results: 1000 libs (1 million records) case



**MemVerge** 



28

#### **Big Memory AI/ML Training and Inference**

Results

Up to 4x

Transactions per second

**Up to 100x** 

Lower latency

50%

Cost savings vs DRAM



### MemVerge Vision for Big Memory Industry By 2025...

- 1. Persistent Memory will be mainstream and data Infrastructure will be memory-centric.
- 2. Big Memory, consisting of PMEM and DRAM, will achieve petabyte-scale over clusters of servers interconnected by next-gen memory fabrics
- 3. Big Memory software will be needed to offer data services in memory, and every application will be run in-memory.

### **MemVerge Vision for Big Memory Industry**





## What happens in memory stays in memory...

