

# Endless Memory



In search of the perfect analytics solution

# Endless Memory Team

## Moderator



**Eric Doyle**  
Director, Channels & Alliances  
MemVerge

## Big Memory Computing



**Mike Hoey**  
Director of Sales  
MemVerge

## Big Memory Use Cases



**Eddie Blanton**  
Director, Solutions Enablement, SSG  
Penguin Computing

# Stay until the end!





# Big Memory Technology

Mike Hoey

Director of Sales

MemVerge


[mike.hoey@memverge.com](mailto:mike.hoey@memverge.com)




# Digital Transformation Driving Big & Fast Data

## A Select Few Apps Today


**Financial Services**  
Digital Banking




**Life Sciences**  
Genomics




**Social Media**  
Customer Profiling



**Retail**  
Recommendation Engines



**Media & Entertainment**  
Animation/Visual Effects



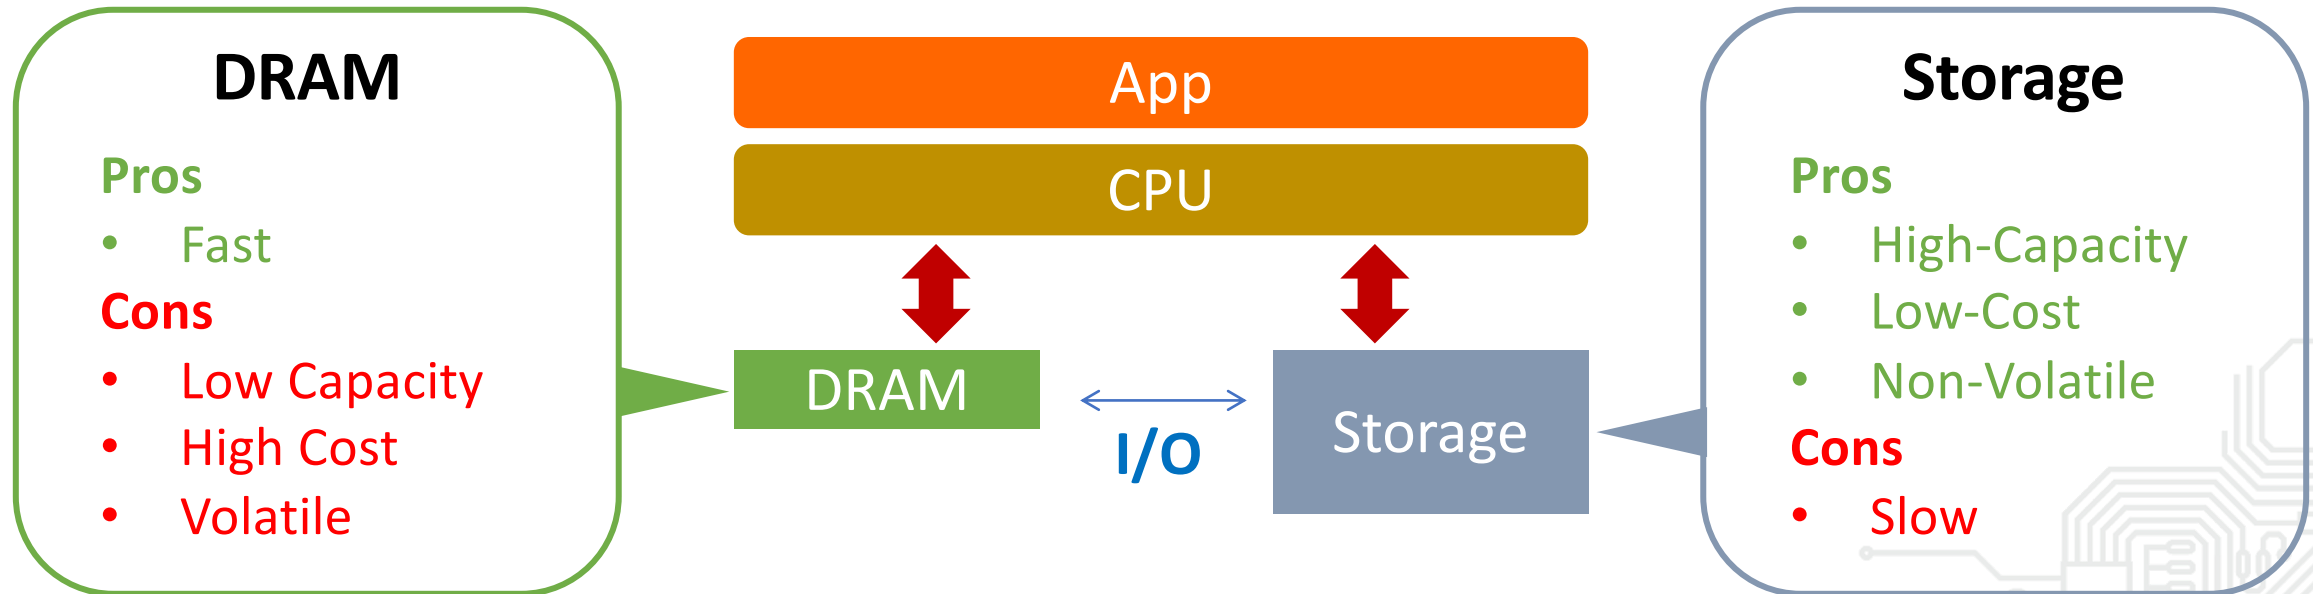
## Pervasive in the Future



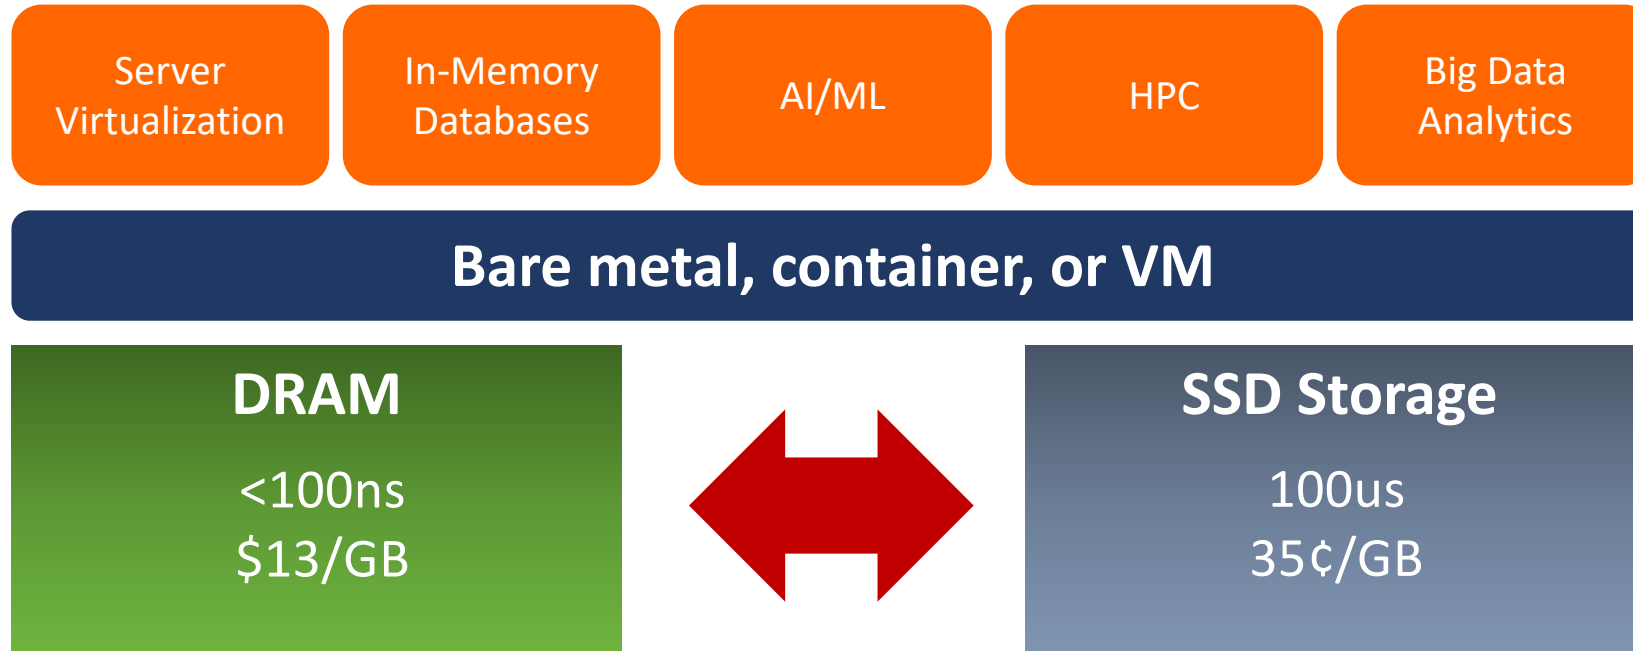
jupyter		PyTorch		R		AI/ML		TensorFlow		XGBoost			
databricks		SQL Server		redis		In-Memory Databases		hazelcast		MySQL		mongoDB	
RANCHER		RED HAT OPENSIFT		VMware Tanzu		Containers		AmazonEKS		Azure AKS		Google Kubernetes Engine	
vmware				Server Virtualization				KVM					
Alibaba.com		aws		Azure		Cloud & On-Premises		Google Cloud		Tencent 腾讯			

# The Memory-Storage Dichotomy: A Fundamental Bottleneck

## Apps Run in DRAM



# Brass Tacks...Root Cause



## IO to Storage kills performance

SSD storage media 1,000x slower than memory

Data serialization / deserialization overhead

Data transport overhead

# Real-Life Pain Points by the Numbers


## Financial Services Databases



**3hrs**

How long it takes to recover a fraud, risk analysis, or payment cluster from storage.

## Data Sciences Multi-Stage Analytics



**60%**

Of multi-stage pipeline throughput is repetitive loading and execution from storage.

## Cloud Infrastructure Compute Instances

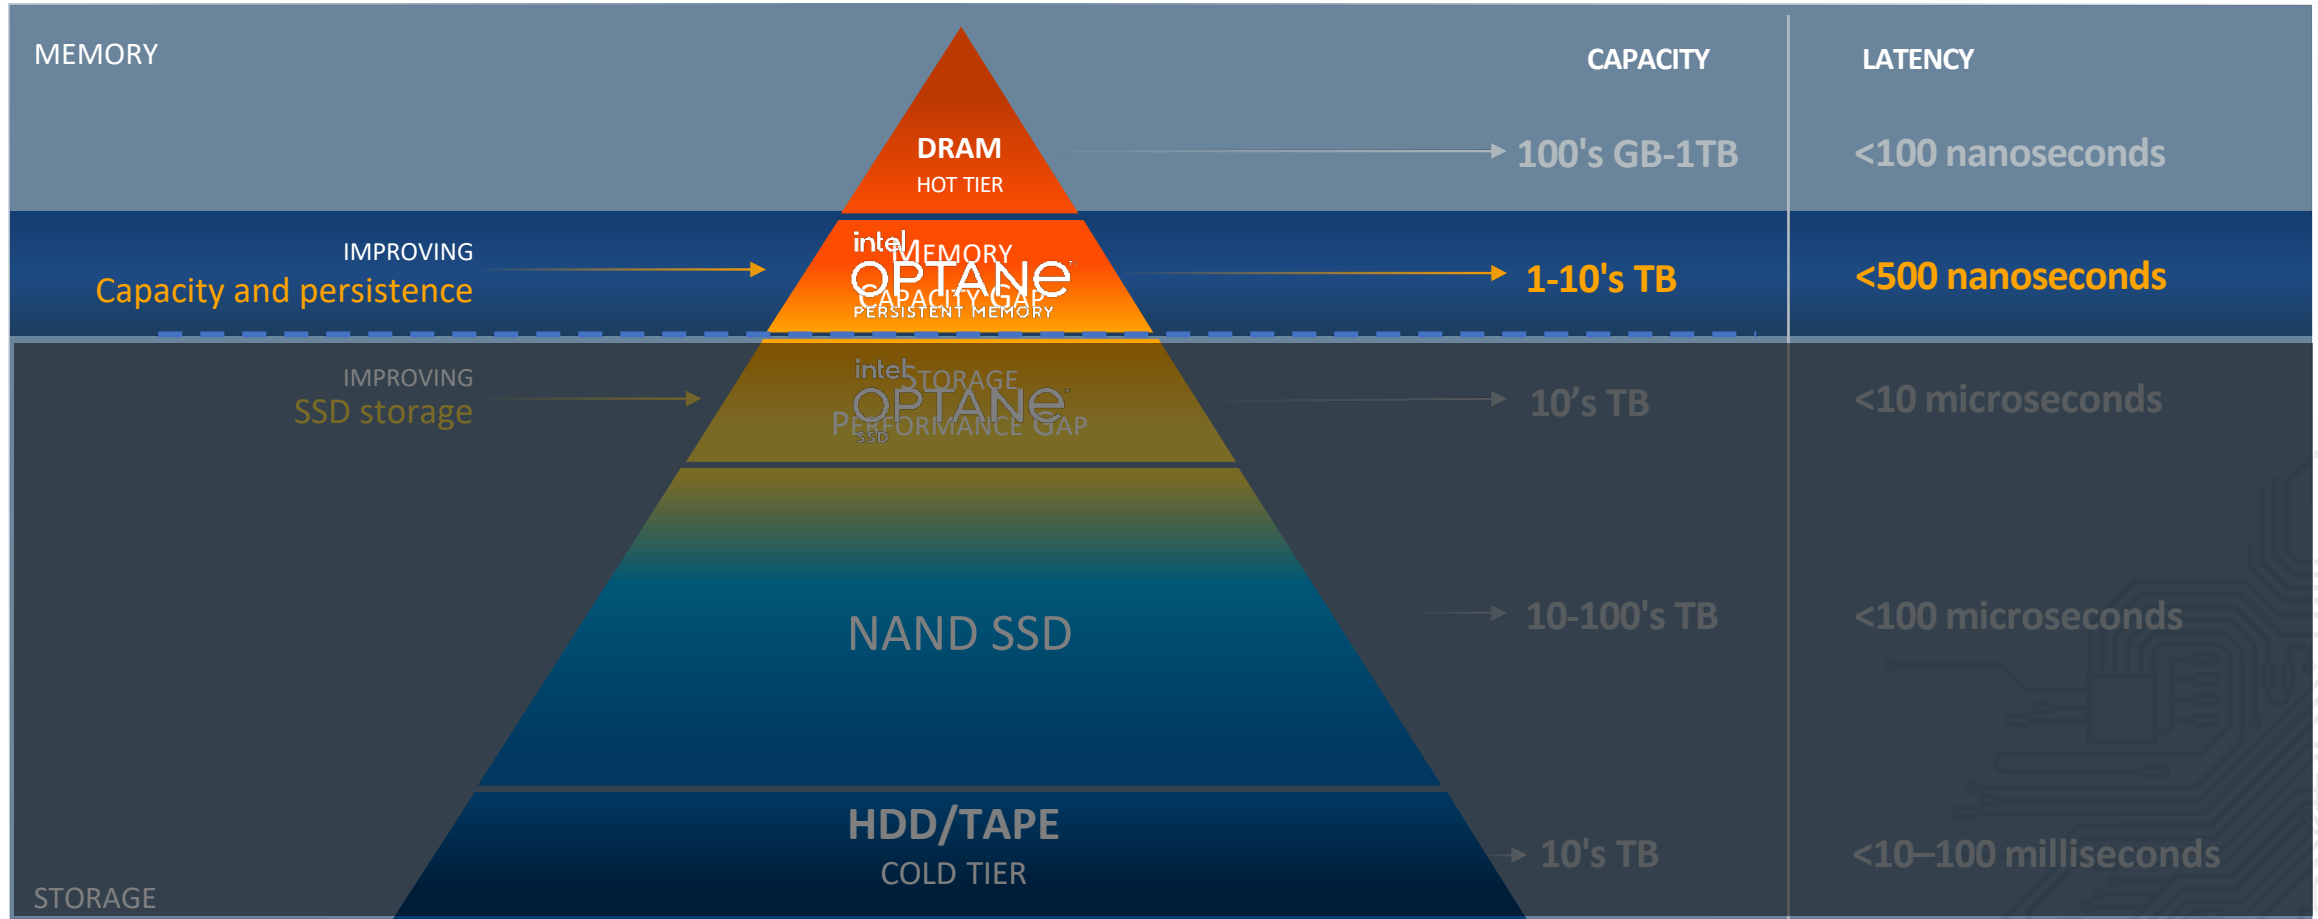


**\$500K**

Server overhead due to memory capacity limitations.

# Re-architecting the Memory/Storage Hierarchy

Intel® Optane™ Persistent Memory (PMem) fills the DRAM -> NAND gap



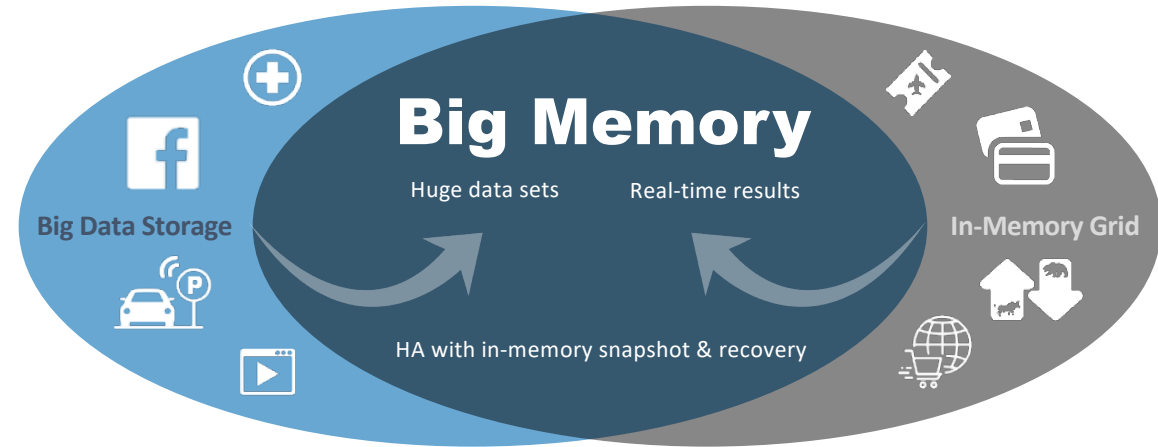
# Introducing MemVerge®

Led by the creators of VSAN  
& XtremIO

Market share leaders and  
product lines with \$ billions in  
annual revenue.



## Pioneering Big Memory Computing

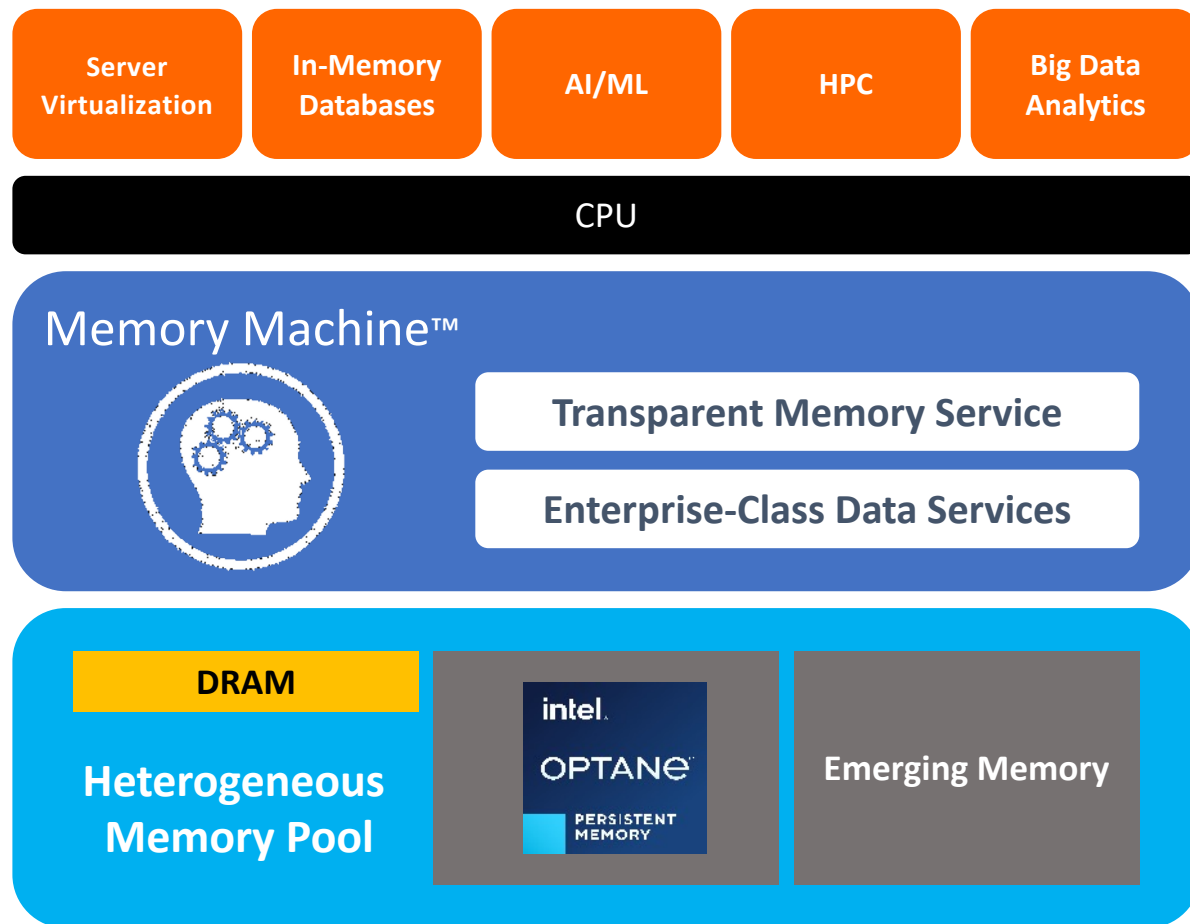


## Investors



# Memory Machine™

## *Big Memory without Compromises*



### Higher Performance

- Elimination of Storage I/O
- Solves Data-Greater-than-Memory (DGM) Problem

### Lower Cost

- Higher Consolidation Ratio
- DRAM-like Performance @ 40% less cost

### Enterprise-class features for HA & Productivity

- ZeroIO™ In-Memory Snapshot
- HA with AutoSave, Time Machine, Replication, and Instant Recovery
- Higher Productivity with IMDB and ML Cloning

**No Application Changes needed !**



# Comparison of PMEM Deployment Modes

## Memory Mode – single pool of volatile memory

### Pros

- Higher VM density
- No application changes needed

### Cons

- DRAM capacity will be hidden
- Total memory capacity = PMEM capacity only
- PMEM becomes volatile
- Multi-VM performance issues
- Fixed DRAM-to-PMEM ratio for all VM's



## App Direct Mode – (2) distinct pools of memory

### Pros

- PMEM is persistent

### Cons

- Cannot be used to extend VM's memory
- Application changes needed



# Why MemVerge's Memory Machine ?

## *Big Memory without Compromises*

Capability	Intel Memory Mode	Intel App Direct Mode	Memory Machine™
Requires <b>changes to app</b>	No	Yes	No
<b>100% Capacity Utilization (DRAM + PMEM)</b>	No	Yes	Yes
Full <b>Persistence</b> in memory	No	Yes	Yes
<b>Auto Tiering/Tuning</b>	No	No	Yes
Modern memory allocator delivering <b>DRAM-like performance</b> for pool of DRAM + PMEM	No	No	Yes
<b>High availability</b> with local and cluster memory-to-memory snapshots and instant recovery	No	No	Yes
<b>Higher IT and business productivity</b> with cloning of IMDBs and ML stages, and on-prem/cloud snapshot mobility	No	No	Yes
<b>Lower TCO</b> no application rewrite, higher capacity utilization, and higher VM density / less servers with great performance	No	No	Yes

# Java C++ App

Lower TCO by consolidating servers and memory




## With DRAM DIMMs-Only

245 Servers  
x \$2,500  
\$612,500

768GB DRAM  
\$7.65/GB  
x 245 Servers  
\$1,439,424

Total  
\$2,051,924




## With Intel® Optane™ Persistent Memory and DRAM capacity @ 2:1

90 Servers  
x \$4,500  
\$405,000

768GB DRAM  
\$7.65/GB  
x 90 Servers  
\$528,768

1.5TB PMEM  
\$4.00/GB  
x 90 Servers  
\$552,960

Total  
\$1,486,728




★  
**Run time reduced**  
from: 12 -13 hours  
to: 3 - 4 hours

# Increase Productivity<sup>1</sup> of Data-Intensive Jobs


CASE STUDY

BioTech


Increase  
Productivity  
of Data-  
Intensive Jobs



## Reduce time to load data by 800x



## Reduce execution time by up to 25x or 60% across all stages



<sup>1</sup> Source – MemVerge. Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.



# Increase Availability with ZeroIO™ In-Memory Snapshots and Instant Recovery



ZeroIO™  
Snapshot Coordinator



Recover from storage



Recover from PMEM



*More memory → more instances per server and High Availability / fast recovery !*

# Increase Availability with ZeroIO™ In-Memory Snapshots / Auto Save and Restore



## Without Memory Machine: Manual, disruptive snapshots to storage

Manual snapshot



Snapshot takes seconds



Restore takes minutes



## With Memory Machine: Memory snapshots and instant recovery

Auto Save



Snapshot takes 1-2 sec



Restore takes 1-2 sec




★ *Reduced downtime & recovery, reduced overtime, and higher quality of work !*


# Lower TCO by Eliminating Noisy Neighbors to Increase VM Density



### Intel Memory Mode



### Memory Machine™




★ *Regained lost DRAM capacity, gained performance,  
and smoothed out Service Levels across VM's !*



# Big Memory is a Massive New Market

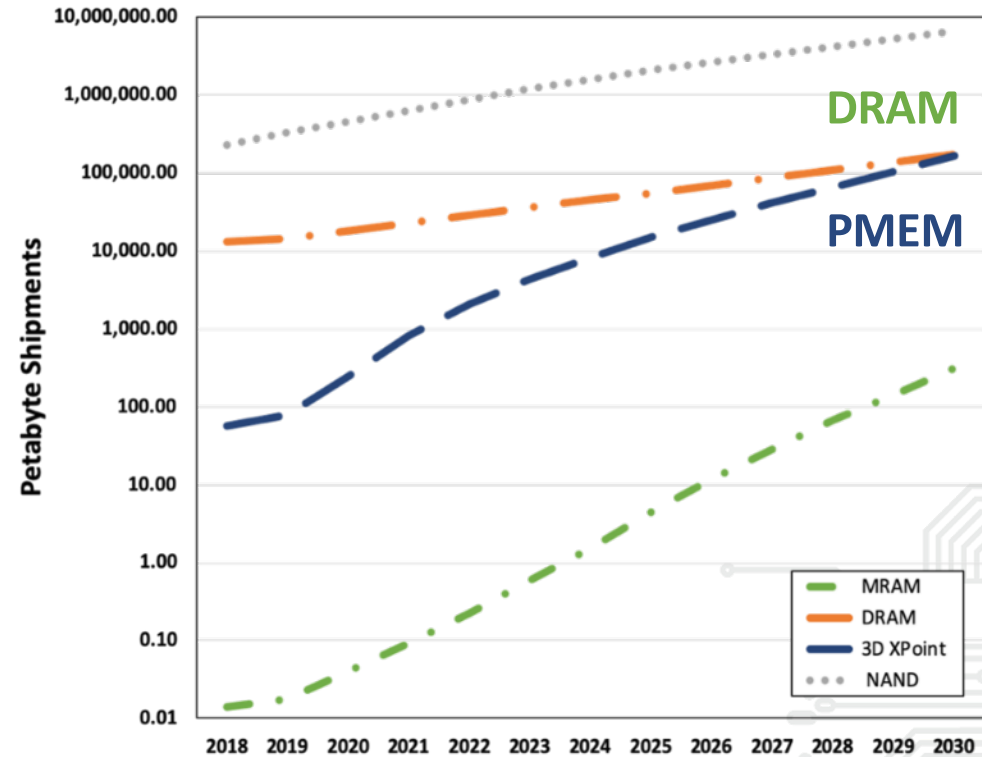
\$2.6B by 2023

Byte-Addressable Persistent Memory  
Revenue (\$M)



\$25B by 2030

Emerging Memories  
Shipments (PB)

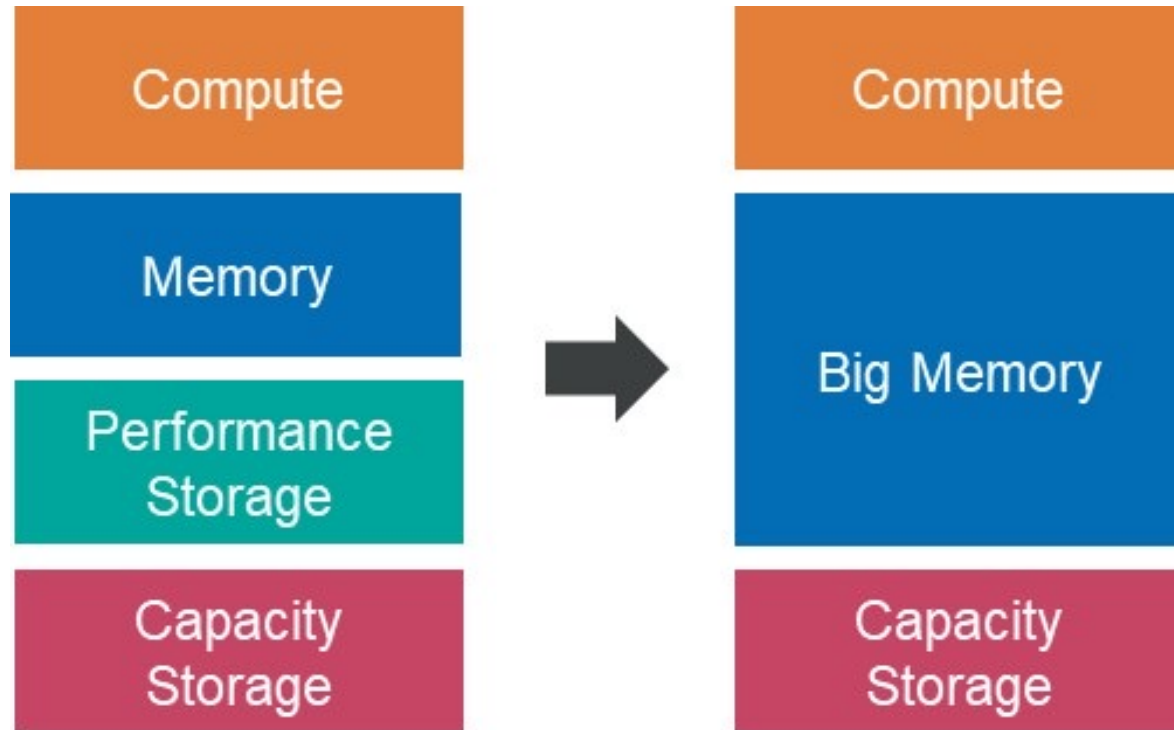


IDC: Digital Transformation Driving New "Big Memory" Requirements

Emerging Memories Find Their Direction: Objective Analysis and Coughlin Associates

# The Future is Memory Centric !

## Digital Transformation Driving New “Big Memory” Requirements



Source: IDC

## “Big Memory Computing”

- Enables the ability to run applications in memory for improved performance and efficiency
  - Leverages byte addressable memory media
- Includes enterprise-class data services to handle Tier-1 availability and management requirements
- Runs on a software-based memory virtualization layer on industry standard hardware without application modification
- The technology enabler for mission-critical real-time computing



# Live Data

Eddie Blanton

Director, Solutions Enablement, SSG

Penguin Computing

# LiveData - MemVerge Memory Machine

## Data Practice

**LiveData**  
Big Memory

LiveData with Memory Machine addresses provides a memory virtualization software layer that delivers software-defined memory services to the applications without application changes.

## Features

- Tier PMEM and DRAM
- Low-latency memory replication
- Virtualized platform for enterprise-class data services.
- Plug-compatible access with existing applications.
- Recover 100s of GBs in seconds with ZeroIO™ memory snapshots
- Clone databases in seconds

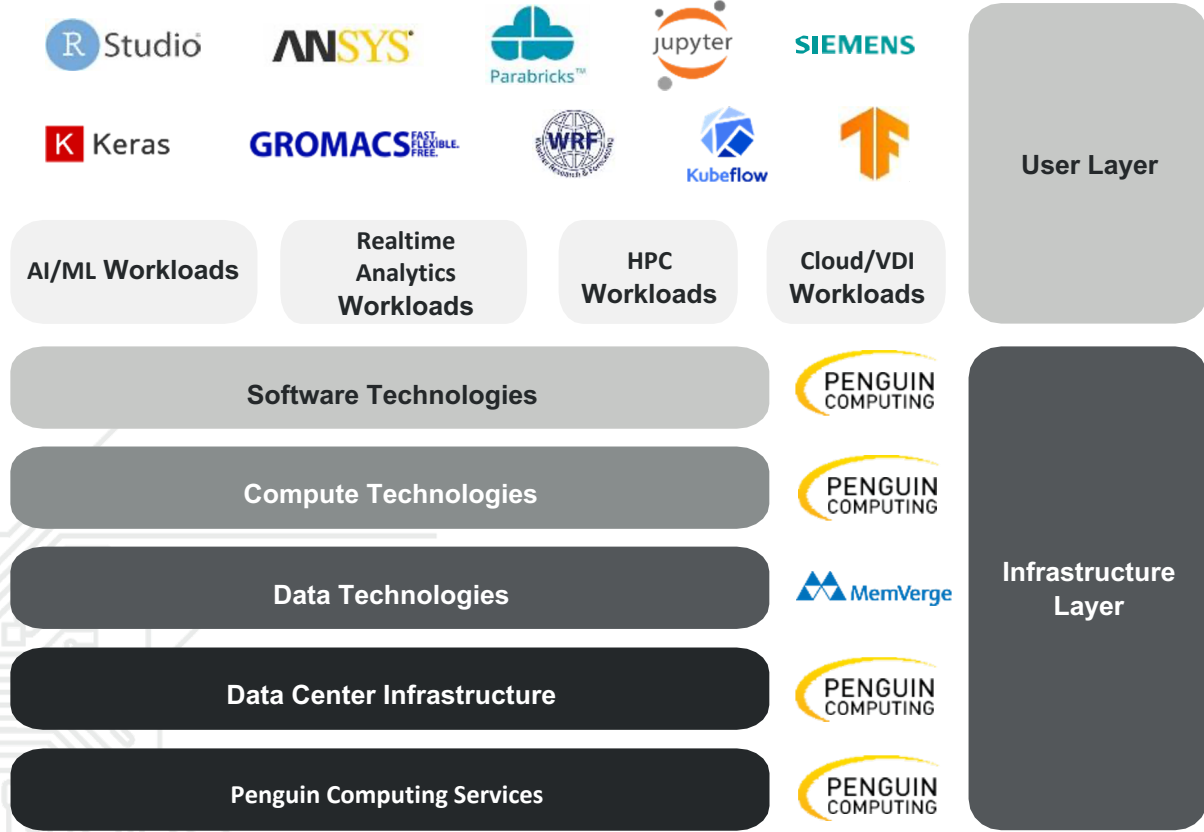
## Benefits

- Scale memory capacity and improve system performance.
- Maintain availability - recover in seconds, not hours.
- Enable agility with efficient clone deployment, and fast application rollbacks.
- Avoid application disruption or rewrite.
- Reduce latency.

# LiveData - MemVerge

## Penguin Computing LiveData Components

Data Practice





# LiveData Platform - Relion XE1112e

## Key Features

- 1U Form Factor
- Dual Intel Xeon Scalable Processor
- Up to 165W TDP with air cool
- Up to 1.5TB DDR4-2933MHz (12 x DIMM)
- Up to 6TB DDR4-2666MHz (12 x PMEM)
- 4 x 2.5" Hot Swappable SATA/SAS SSD
- 4 x 2.5" Hot Swappable SATA/SAS or NVMe SSD
- 1 x Dedicated BMC + 2x 10GBase-T
- 2 x PCIe Gen3 (x16) LP
- 1 x PCIe Gen3 (x16) OCP Mezzanine 2.0



# LiveData - MemVerge



## LiveData - MemVerge

### 32-Node Rack

#### *Memory*

- Up to 48 TB DDR4 DRAM
- Up to 192 TB Intel Optane PMEM
- Up to 240 TB Total Memory

#### *Compute*


- Up to 1792 Intel Scalable Cores
- Up to 3.9GHz Base Clock Speed

#### *Network*

- Up to 800 GB/s 1:1 FBB
- Ethernet, InfiniBand, Omni-path

#### *Power*

- 27.6 kW Maximum
- 22.1 kW Nominal





# LiveData - MemVerge



## LiveData - MemVerge

### 40-Node Rack

#### *Memory*

- Up to 60 TB DDR4 DRAM
- Up to 240 TB Intel Optane PMEM
- Up to 300 TB Total Memory

#### *Compute*


- Up to 2240 Intel Scalable Cores
- Up to 3.9GHz Base Clock Speed

#### *Network*

- Up to 1000 GB/s 1:1 FBB
- InfiniBand

#### *Power*

- 34.0 kW Maximum
- 27.2 kW Nominal





# Use Cases

# LiveData - 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



Real-Time Big Data  
Analytics  
in financial services,  
healthcare, retail, Federal  
Govt, and security.



AI/ML Analytics and  
Inferencing  
like fraud detection,  
image and video  
recognition, autonomous  
vehicles and healthcare.

# Big Memory for Financial Services

Financial Services

**LiveData**  
Big Memory



## Problem - IMDB Cluster takes 2 Hours to Restore

- Hazelcast in-memory database distributed across 40 servers.
- A cluster crash takes over 2 hours to completely restore
- Each server limited to 100GB of DRAM because scaling DRAM is extremely expensive
- Deploying lower cost PMEM sacrifices performance



# Big Memory for Financial Services

Financial Services

**LiveData**  
Big Memory



## Solution - IMDB Cluster Restored in Seconds


- MemVerge allows for Hazelcast cluster to be restored in seconds
- TCO - PMEM provides nearly twice the memory for the same cost
- 10x Server Consolidation
- MemVerge memory allocator technology allows PMEM to go as fast or faster than DRAM

# Big Memory for Financial Services

## Financial Services

### LiveData Big Memory

- DRAM = 48.3K Transactions Per Second
- PMEM + 16GB DRAM = 50K Tps (+3.5%)



# Big Memory for M&E

## Media & Entertainment

**LiveData**  
Big Memory



### Problem - 83,000 Hours of Creative Artist Time Lost Per Year


- "We need more RAM" is the top request from artists
- Creative artist time is the biggest cost item for any movie production!
- Creative artist's lose on average 1-2 hours of productivity per week
- 800 Artists x 1,600 hours per week = 83,200 hours per year



# Big Memory for M&E

Media & Entertainment

**LiveData**  
Big Memory



## Solution - Time Equivalent to 40 Creative Artists Restored

- Lower TCO: Workstations with double the Memory + MemVerge for the same cost or less
- “Autosave” improves productivity by potentially 40 Creative Engineers annually
- \$8M potential savings annually
- Higher quality movies & faster time to market

# Big Memory for M&E

## Media & Entertainment

**LiveData**  
Big Memory

### Snapshot Time

- 99.66% Reduction

### Restore Time

- 99.98% Reduction

**Productivity**

**Artist Experience**

## Autosave

Auto Save

Snapshot  
takes 1-2 Sec

Restore  
takes 1-2 Sec



## Time Machine – Instant roll back / forward



# Big Memory for AI/ML Inferencing

AI/ML Inferencing

LiveData  
Big Memory



## Problem - Data Greater than Memory

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



# Big Memory for AI/ML Inferencing

AI/ML Inferencing

**LiveData**  
Big Memory

## Solution - Big Memory Lake

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



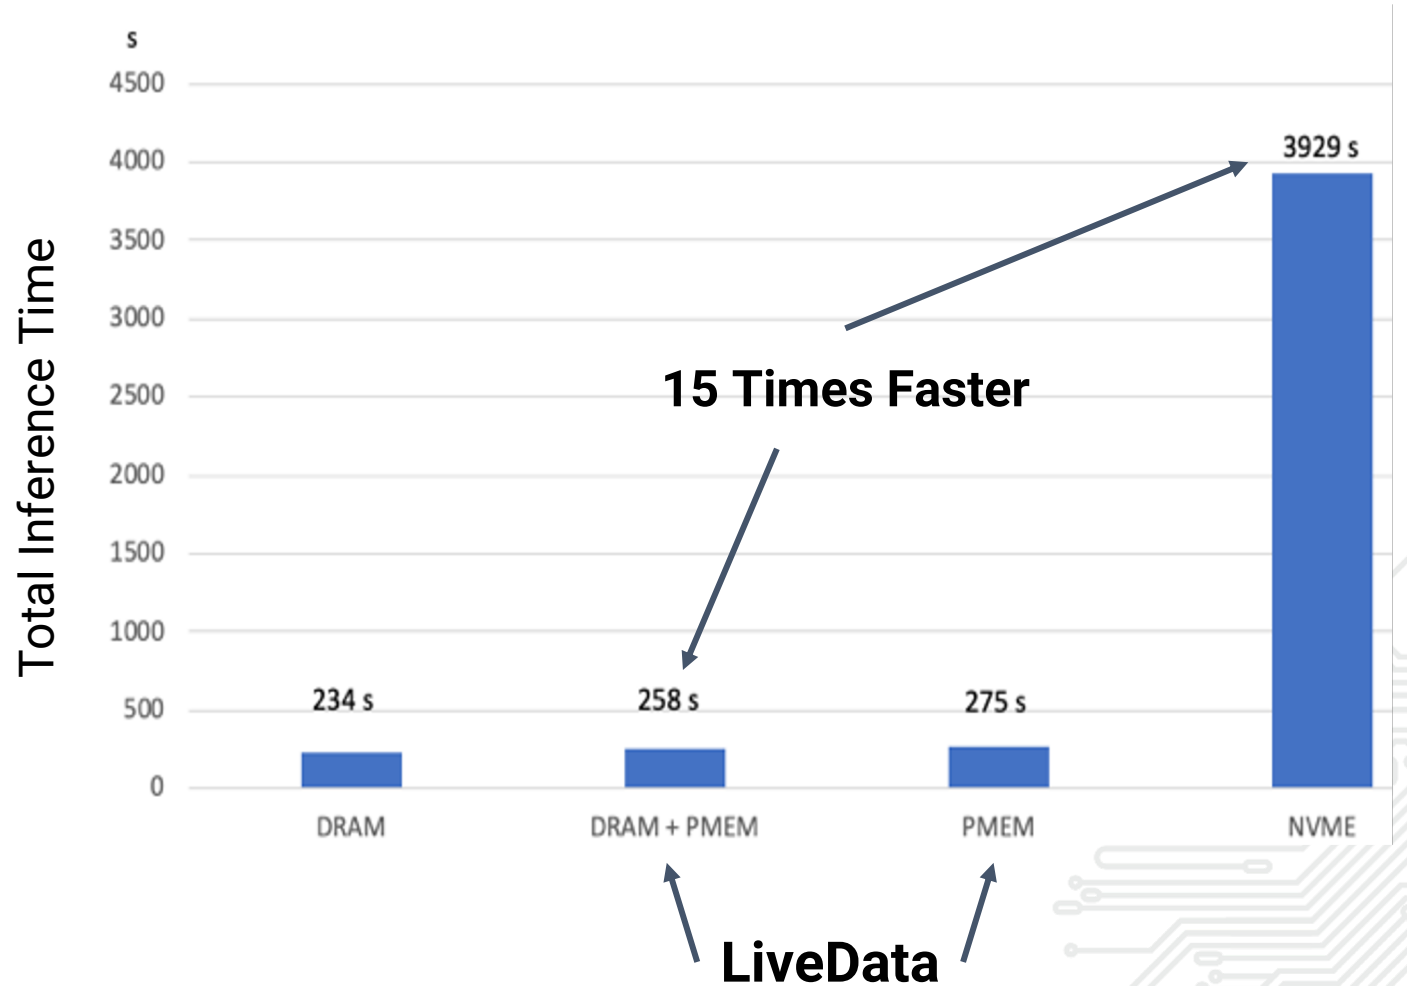
# Big Memory for AI/ML Inferencing

AI/ML Inferencing

LiveData  
Big Memory

## Inference Time


- DRAM+PMEM is within 10% of DRAM
- DRAM+PMEM is 15X Faster than NVMe






# Try it in your lab or ours

Server w/PMem



Memory Machine Software



Big Memory Labs



Systems Engineering Support





# Memory Machine™

For Big and Fast Data



# Q & A

# Contact Us

- **Penguin Computing**
  - Phone: 415-954-2800
  - [penguincomputing.com](http://penguincomputing.com)
- **MemVerge**
  - Phone: 408-823-3902
  - [memverge.com](http://memverge.com)



# Endless Memory



In search of the perfect analytics solution