SNIA PERSISTENT MEMORY + SUMMIT 2021
COMPUTATIONAL STORAGE

FROM DATACENTER TO EDGE : VIRTUAL EVENT
APRIL 21-22, 2021

# 4 Top Use Cases for Big Memory

## Today and Tomorrow

Charles Fan, CEO & Co-Founder, MemVerge

# The Future of Infrastructure



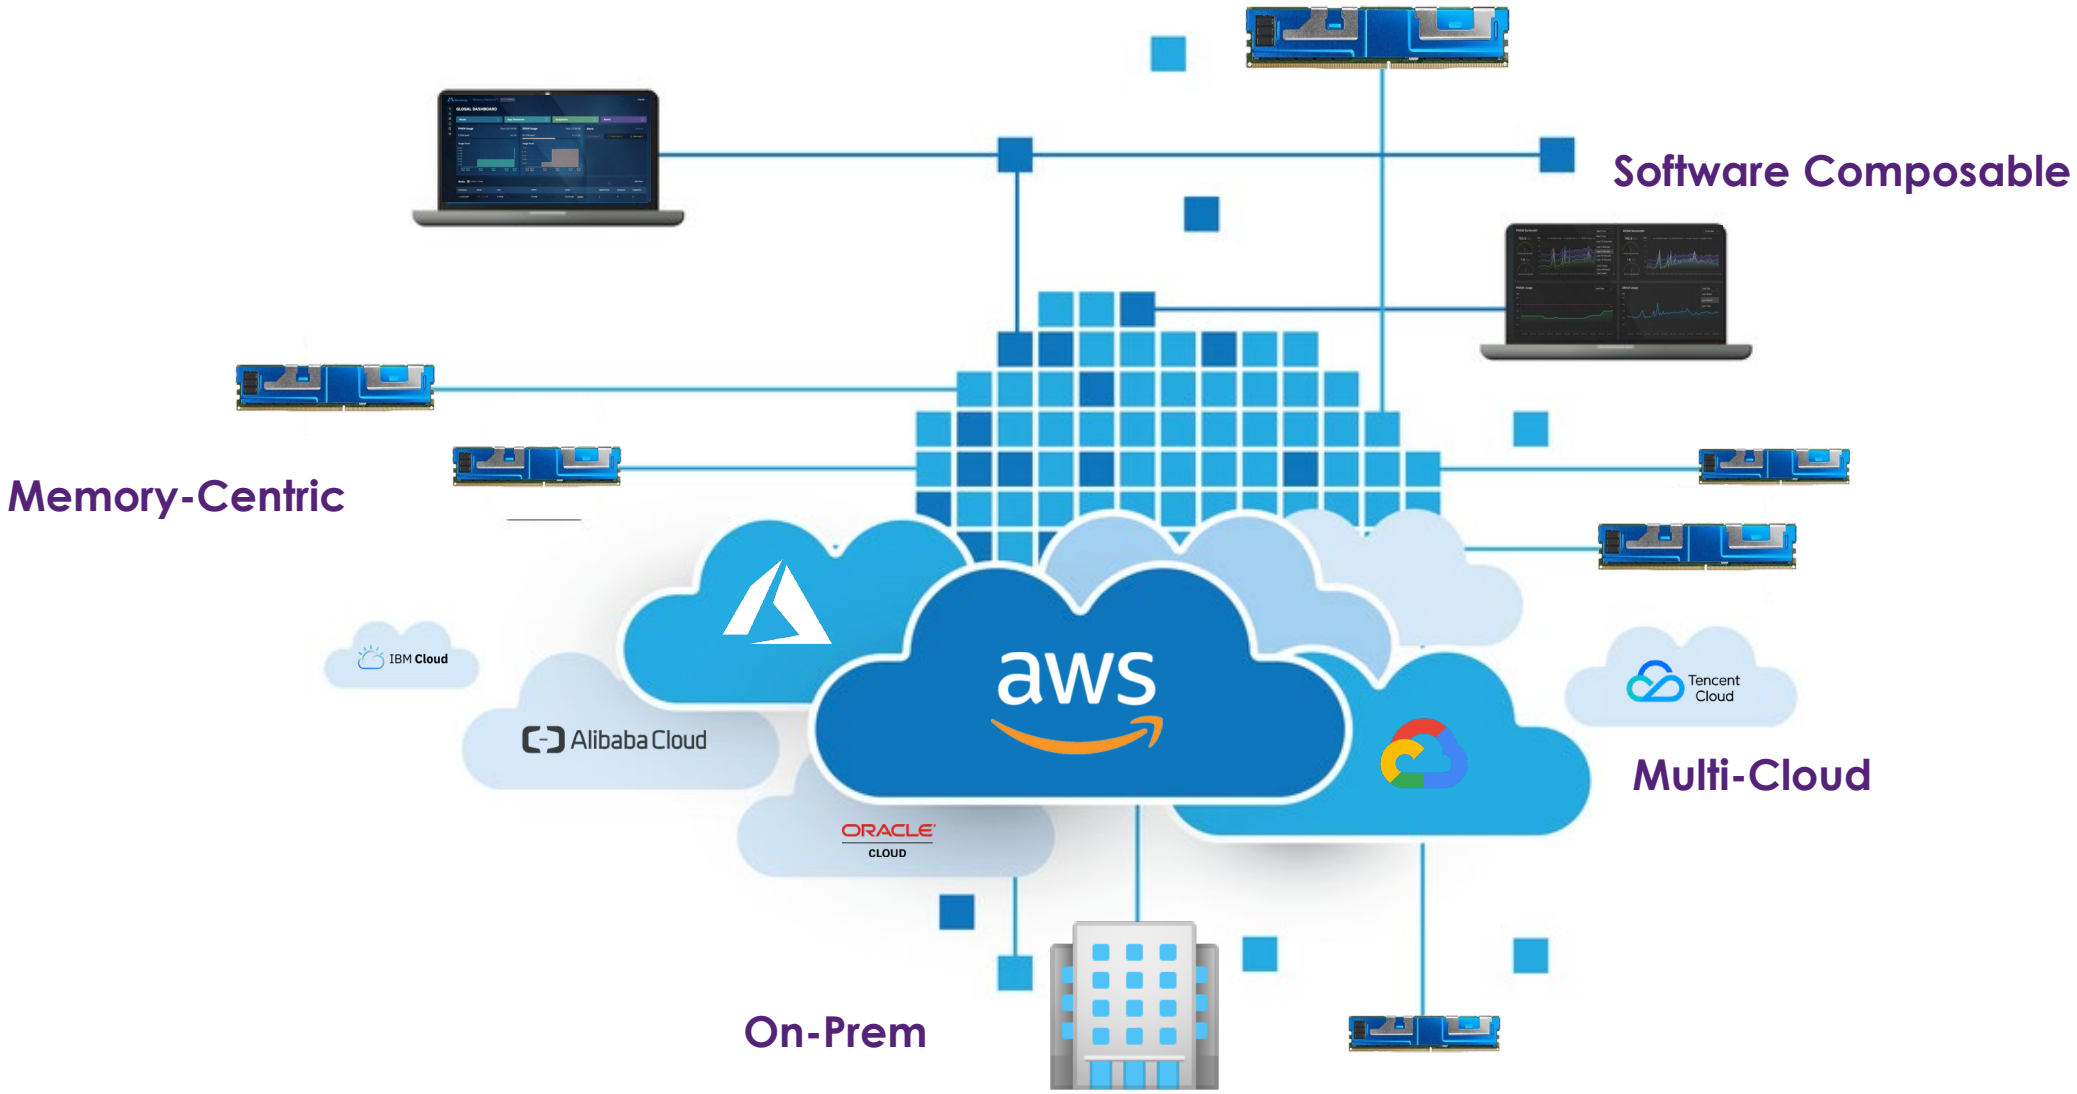

# Big Memory Computing for Big and Fast Data



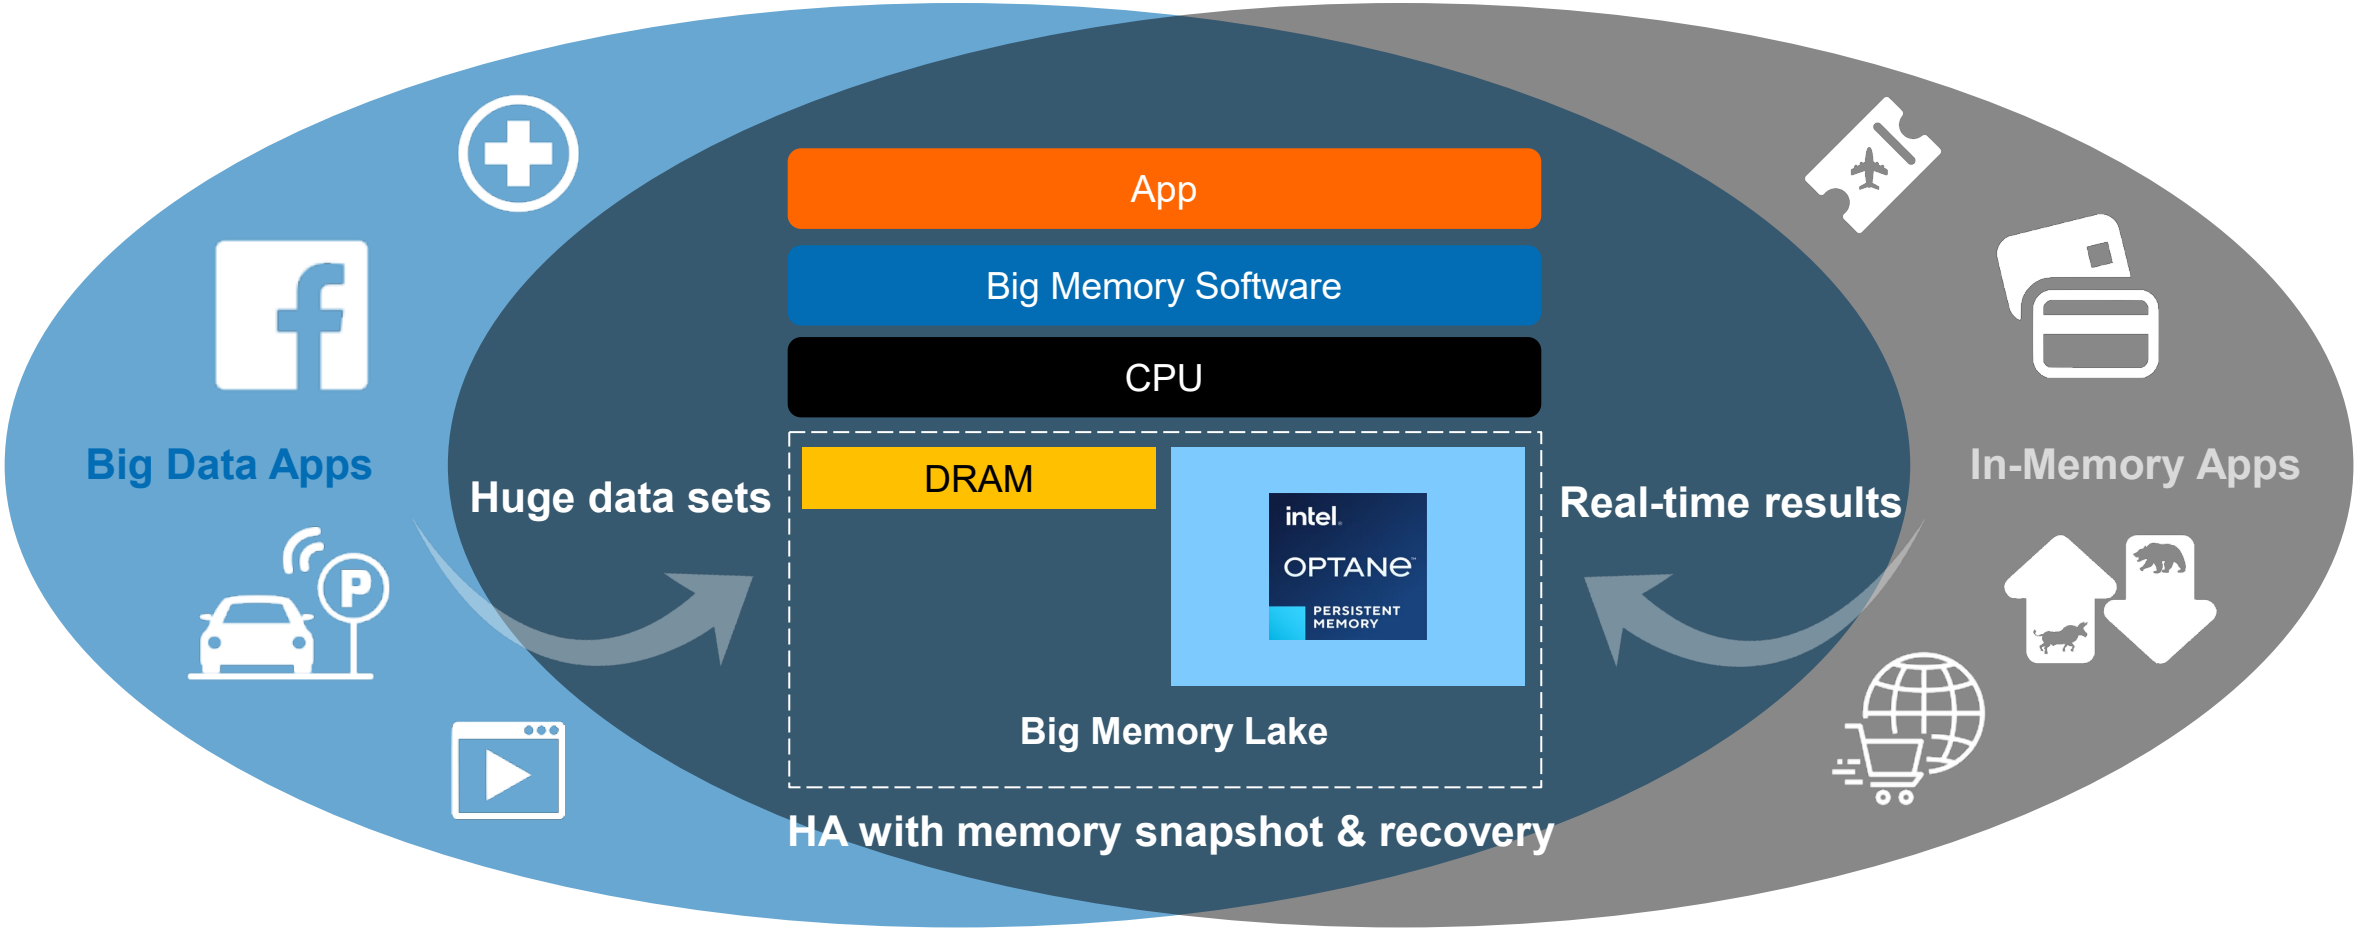

# Platforms & Apps Needing Lower Cost, Higher Capacity, and HA Memory



**Financial Services**
Digital Banking


**Health Sciences**
Genomics

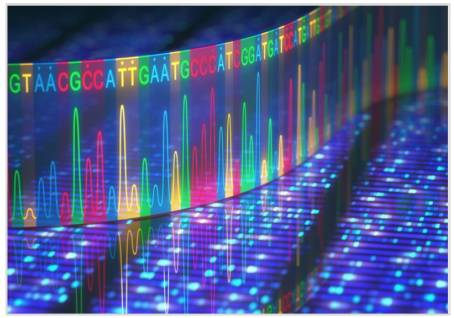

**Media & Entertainment**
Animation/Visual Effects

**Social Media**
Customer Profiling

**Retail**
Recommendation Engines


jupyter PyTorch R **AI/ML** TensorFlow **XGBoost**

databricks Microsoft SQL Server redis **In-Memory Databases** hazelcast MySQL mongoDB

RANCHER RED HAT OPENSHIFT VMware Tanzu **Containers** Amazon EKS Azure AKS Google Kubernetes Engine

vmware **Server Virtualization** KVM

Alibaba.com aws Azure **Cloud & On-Premises** Google Cloud Tencent 腾讯

# Memory Machine™
# Big Memory without Compromises



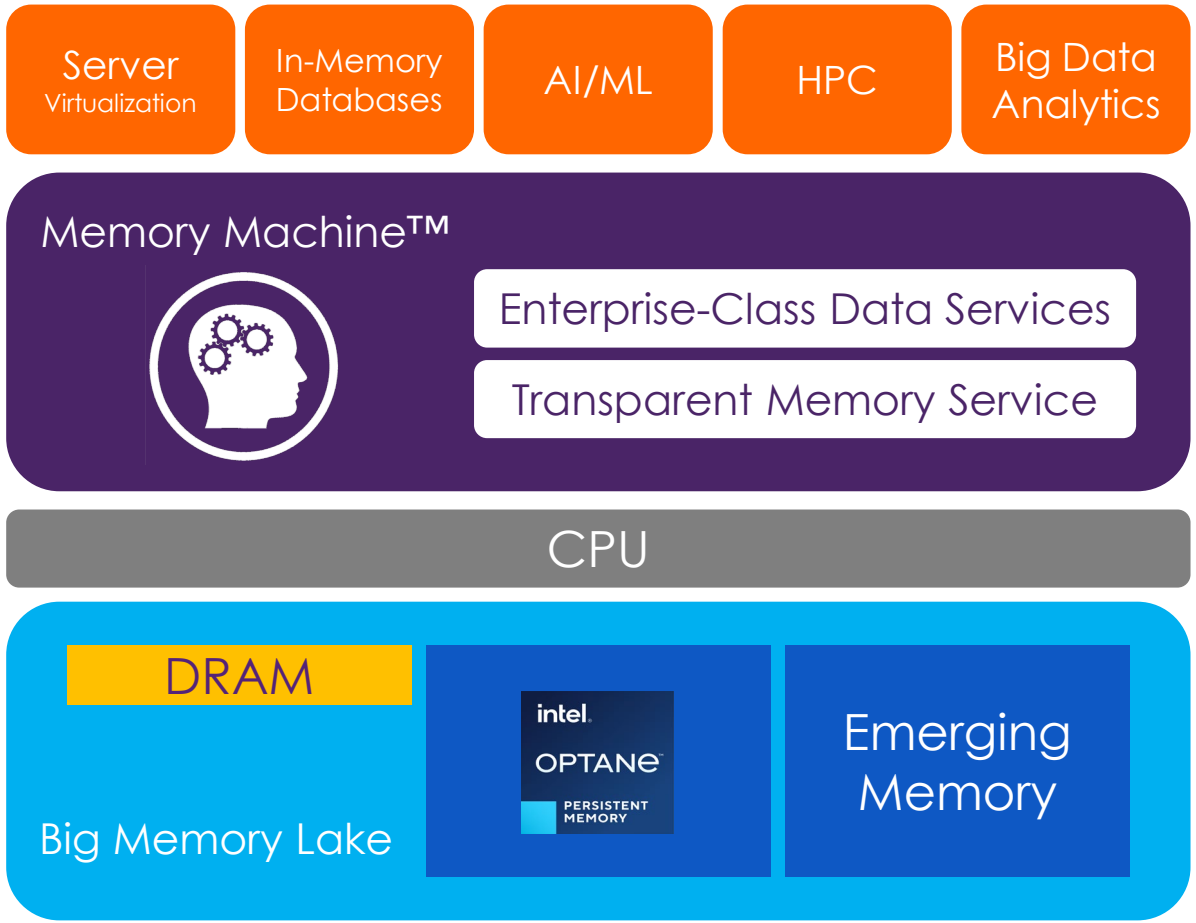


**New Memory Price/Performance**

- 30-40% memory cost savings
- Same DRAM-like performance

**Now Practical to Put all Data in Memory**

- Solving performance problems due to data-greater-than-memory (DGM)
- By eliminating IO to storage

**Unleashes New Class of Memory-based Data Services for Higher Availability and Productivity**

- ZeroIO™ in-memory snapshots
- HA with auto save, time machine, replication, and instant recovery
- Higher productivity with cloning of IMDBs and ML stages

**All With No Changes to Your Application!**

# Top 4 Use Cases Today



Cloud Infrastructure

Databases

Genomics

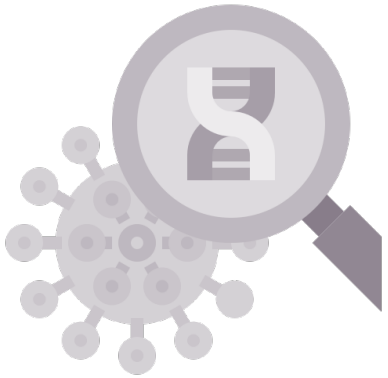

Animation & VFX

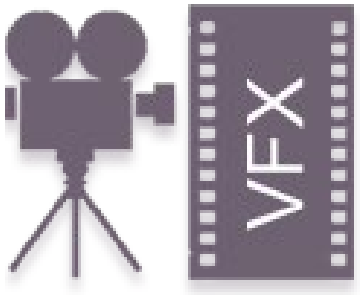

# Cloud Infrastructure



VM density is money

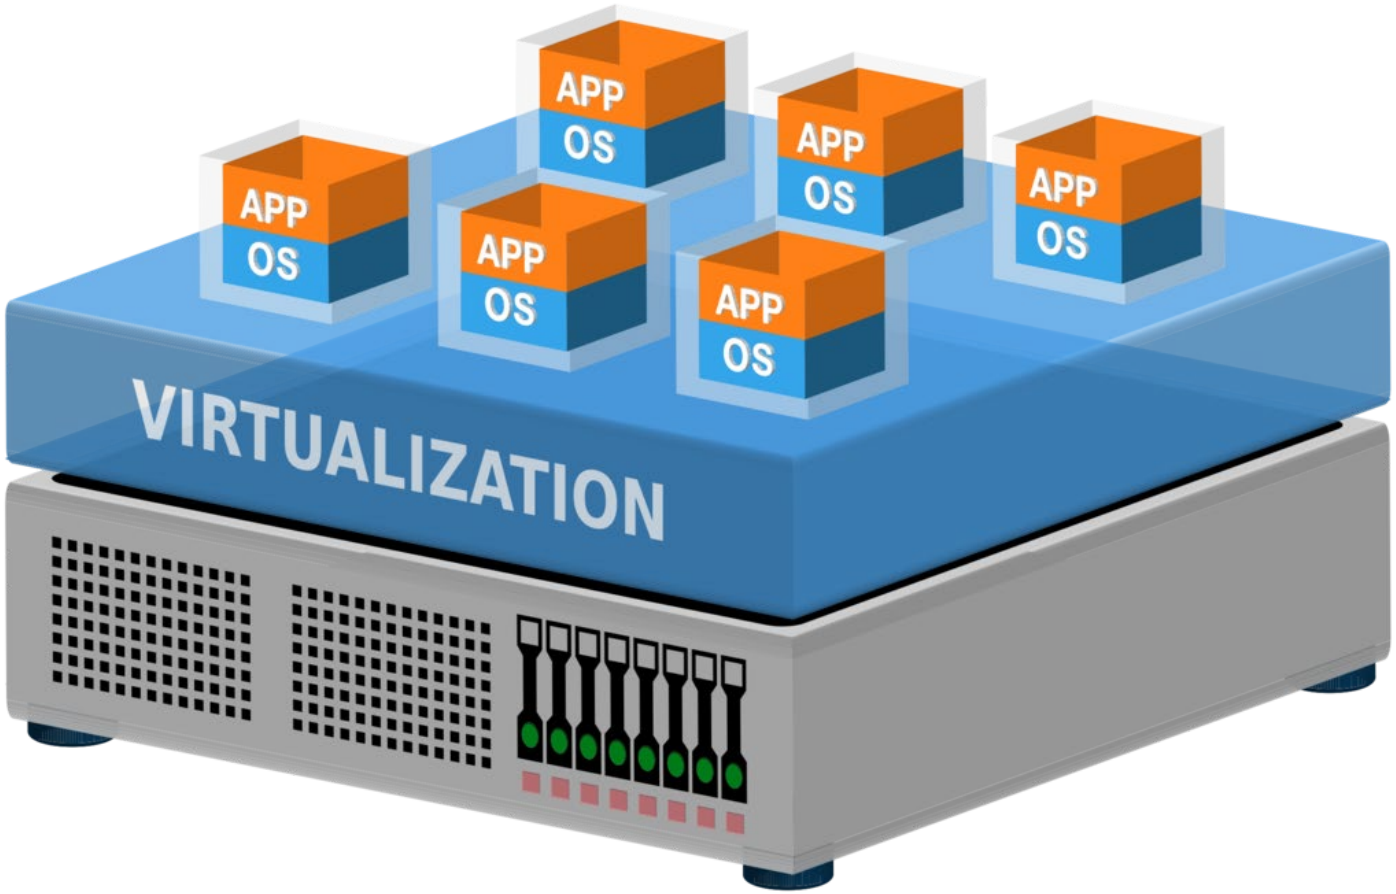

# Big Memory for Virtualized Servers



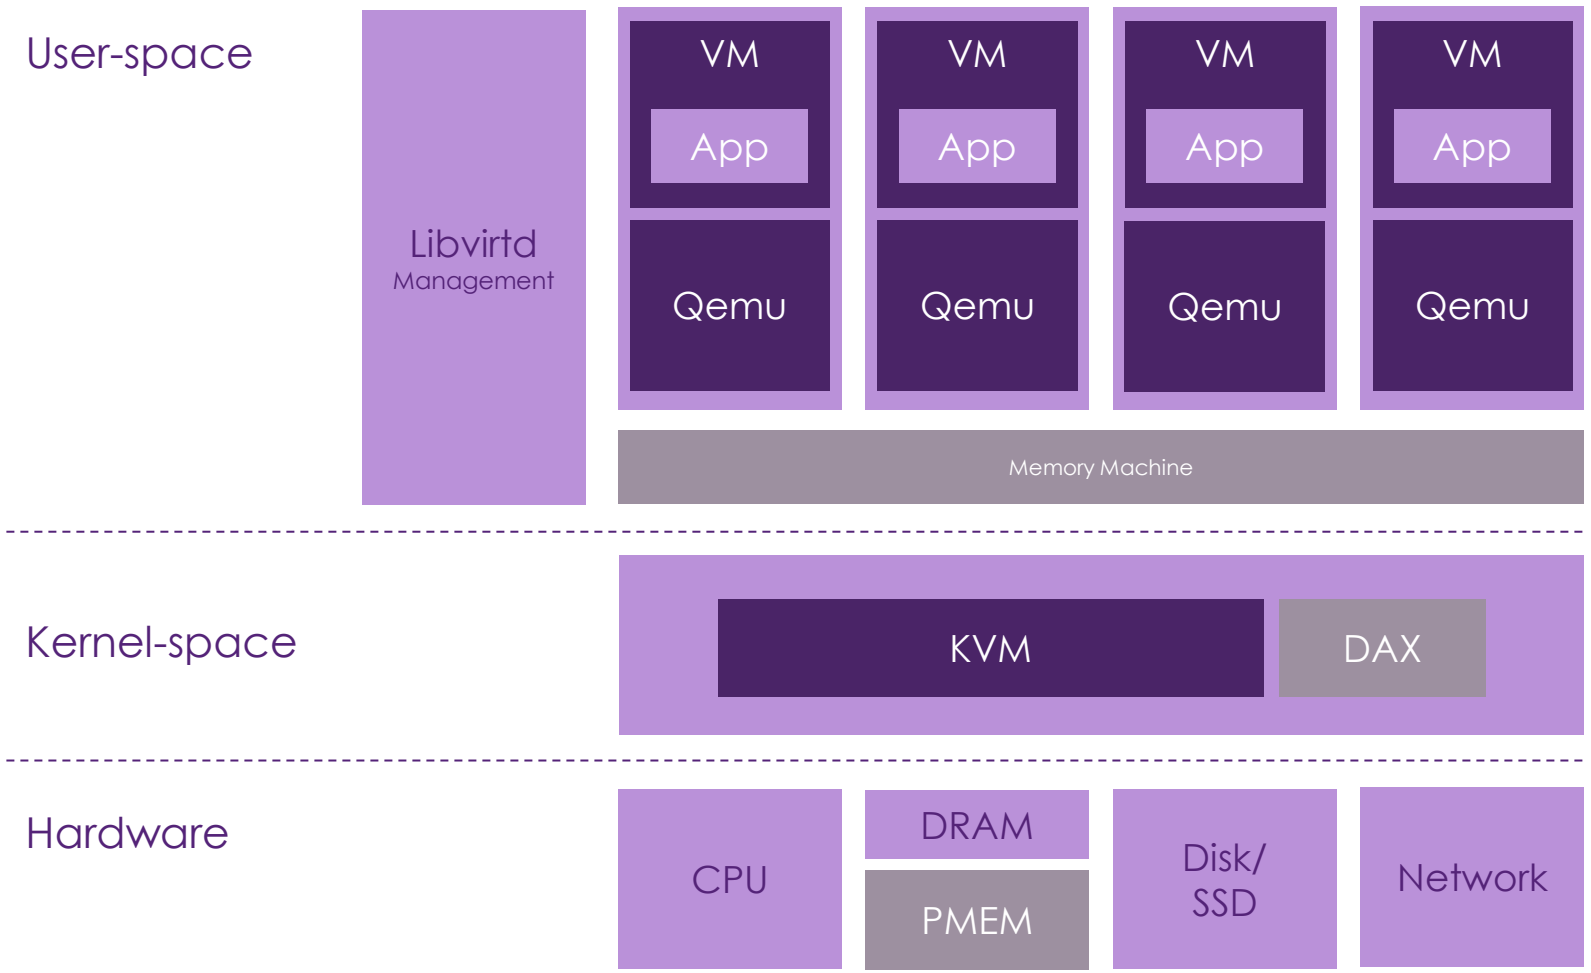


**Memory Virtualization**
KVM allocates memory from a tiered memory pool of software-defined memory (DRAM + PMEM)

**glibc compatible**
No application change, rewrite or recompile is needed

**Provisioning on a per-app basis:**
Maximum DRAM and PMEM
Different DRAM and PMEM ratios
Dynamic tuning of DRAM tier size for each VM

**Resource isolation on a per-app basis**
Allocate from different memory pools, avoiding noisy neighbours

**Monitoring and visualization**
of memory usage of multiple physical servers and each app

**High performance**
DRAM + PMEM pool at near-DRAM performance

# Cloud Server Consolidation & Lower TCO

## By Eliminating Noisy Neighbors & Increasing Memory Density



### Intel Memory Mode

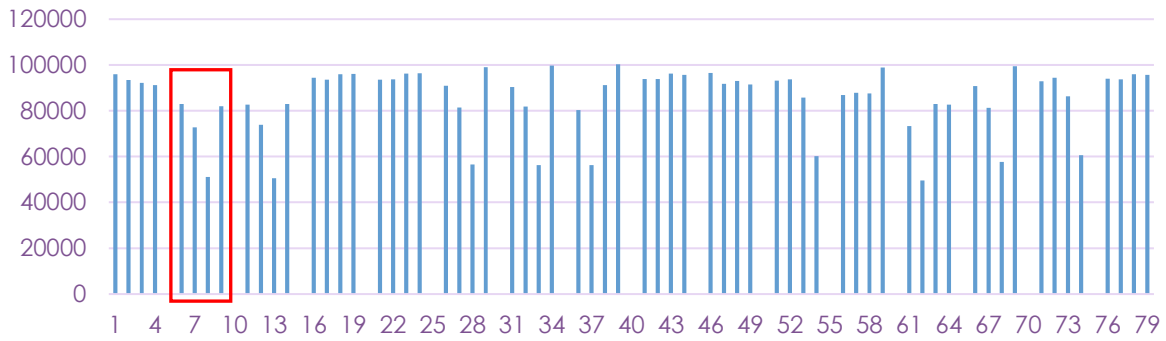


### Memory Machine

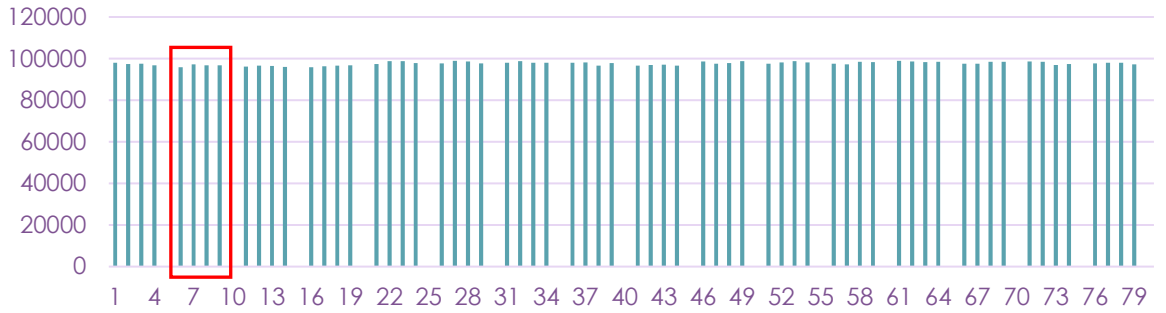


### With DRAM DIMMs-Only

| 245 Servers | 768GB DRAM | |
|---|---|---|
| x $2,500 | $7.65/GB | |
| $612,500 | x 245 Servers | Total |
| | $1,439,424 | $2,051,924 |

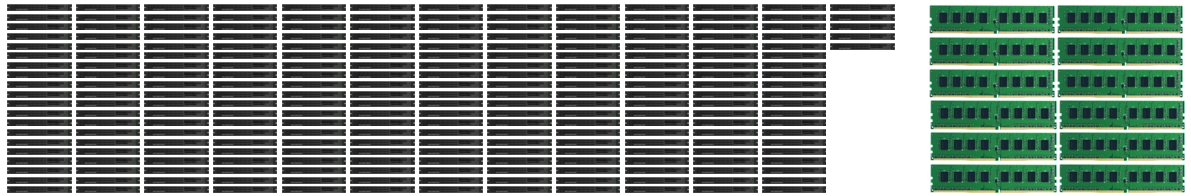

### With Intel® Optane™ Persistent Memory and DRAM capacity @ 2:1

| 90 Servers | 768GB DRAM | 1.5TB PMEM | |
|---|---|---|---|
| x $4,500 | $7.65/GB | $4.00/GB | |
| $405,000 | x 90 Servers | x 90 Servers | Total |
| | $528,768 | $552,960 | $1,486,728 |

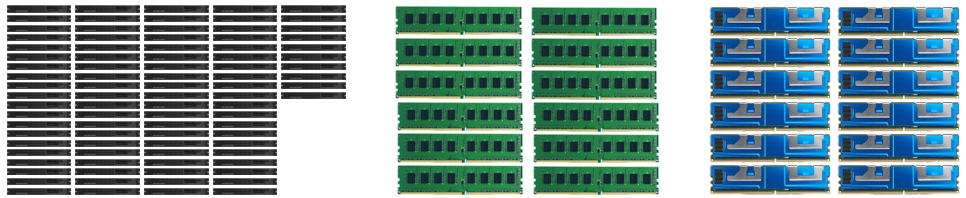

# Databases



Data frequently becomes greater than memory and IO to storage is

**slow**


Scaling DRAM to improve performance is

**expensive**

Big memory blast zone is big making recovery from storage slow and

**disruptive**

# 25% Greater Capacity Utilization



redis hazelcast®

Installed = 384GB DRAM + 1,536GB of PMEM = 1,920GB

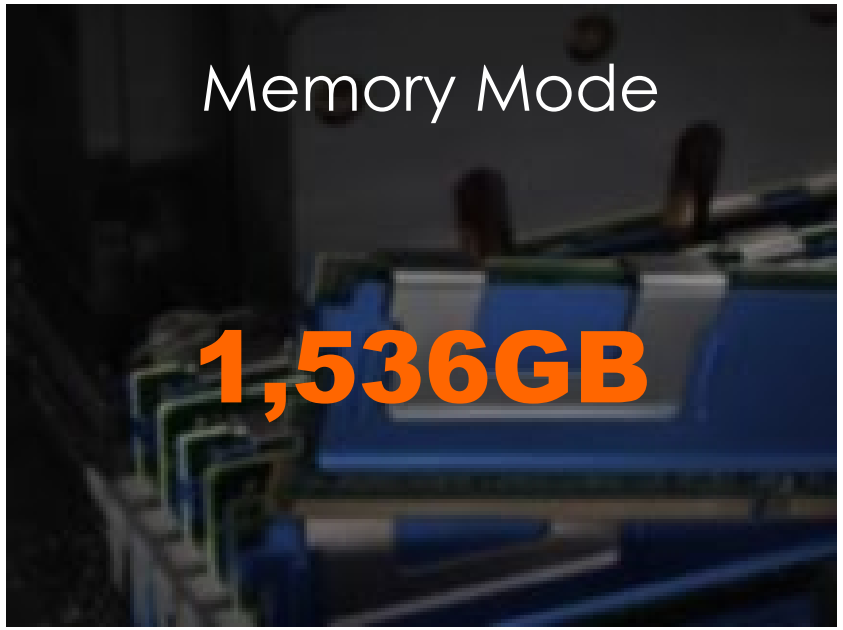


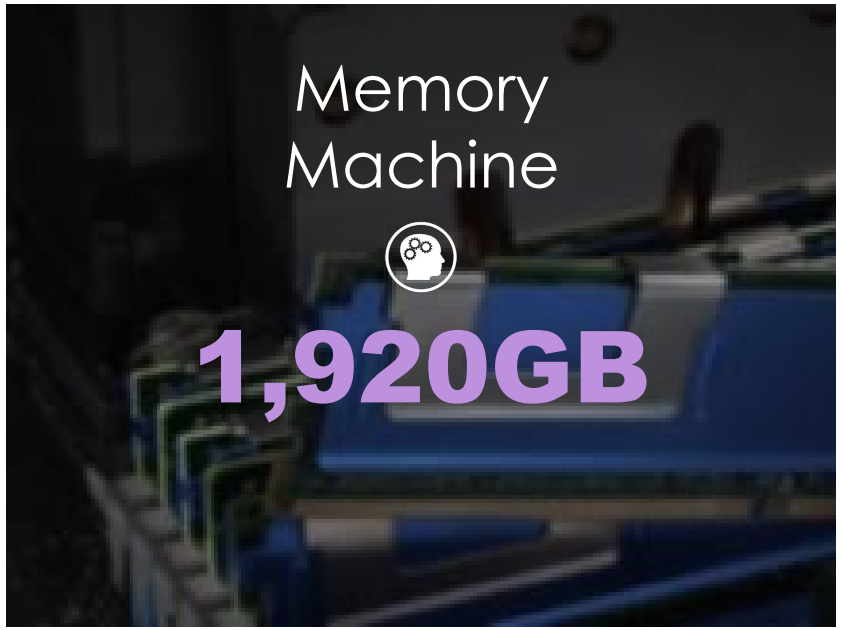

# More & Bigger Shards & Members/Node



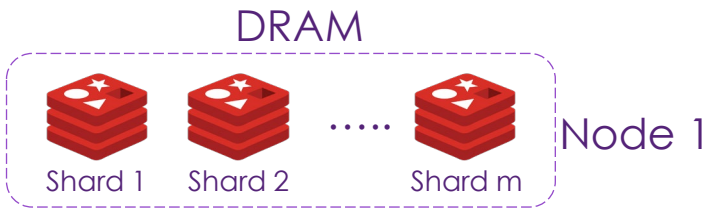



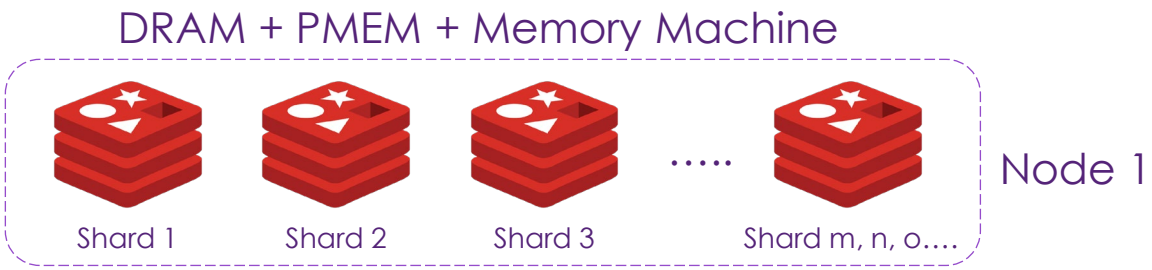


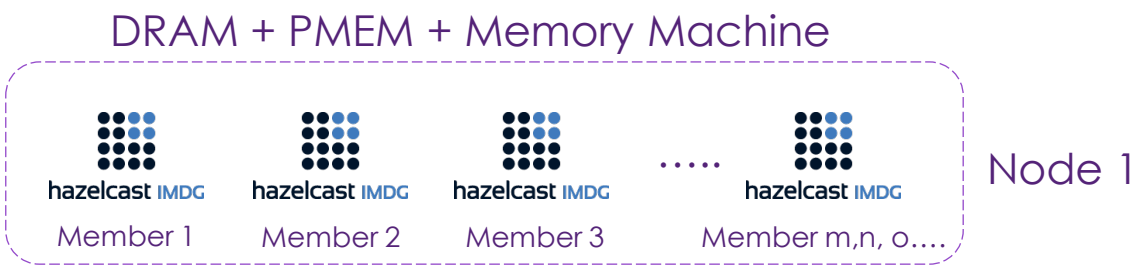

# No Memory Virtualization Overhead



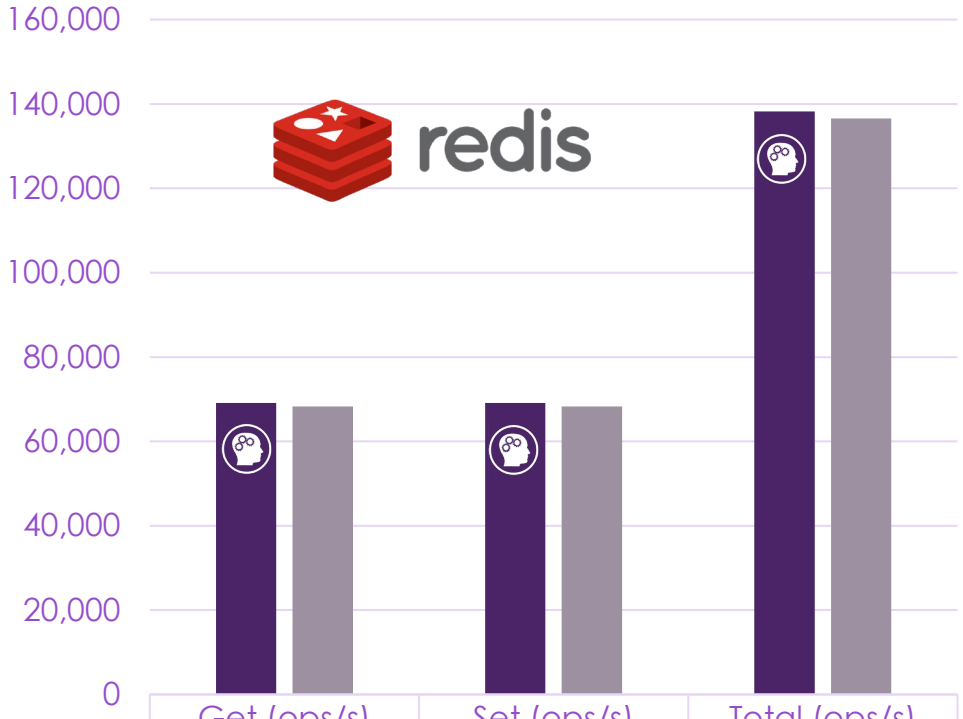


| | Get (ops/s) | Set (ops/s) | Total (ops/s) |
|---|---|---|---|
| DRAM & Memory Machine | 69,129 | 69,131 | 138,261 |
| DRAM & OS | 68,294 | 68,296 | 136,591 |

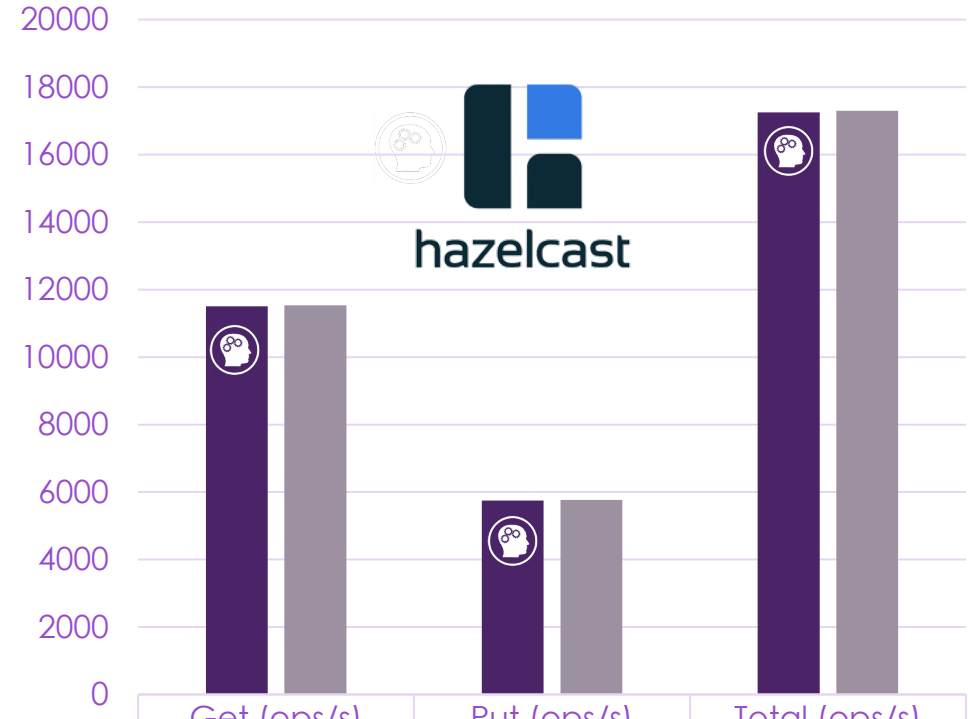


| | Get (ops/s) | Put (ops/s) | Total (ops/s) |
|---|---|---|---|
| DRAM & Memory Machine | 11500 | 5750 | 17250 |
| DRAM only | 11530 | 5770 | 17300 |

# Restore Entire Cluster in Seconds



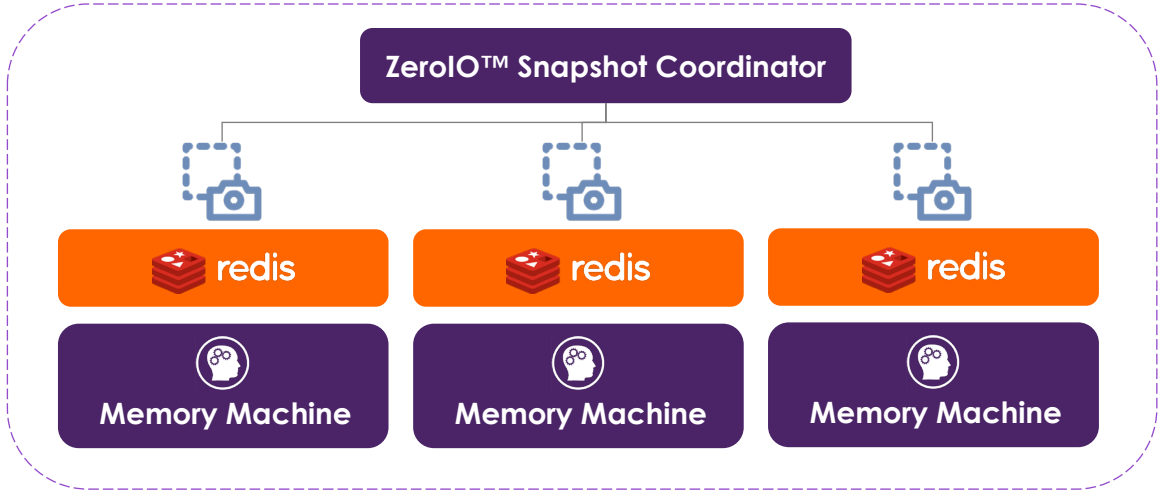


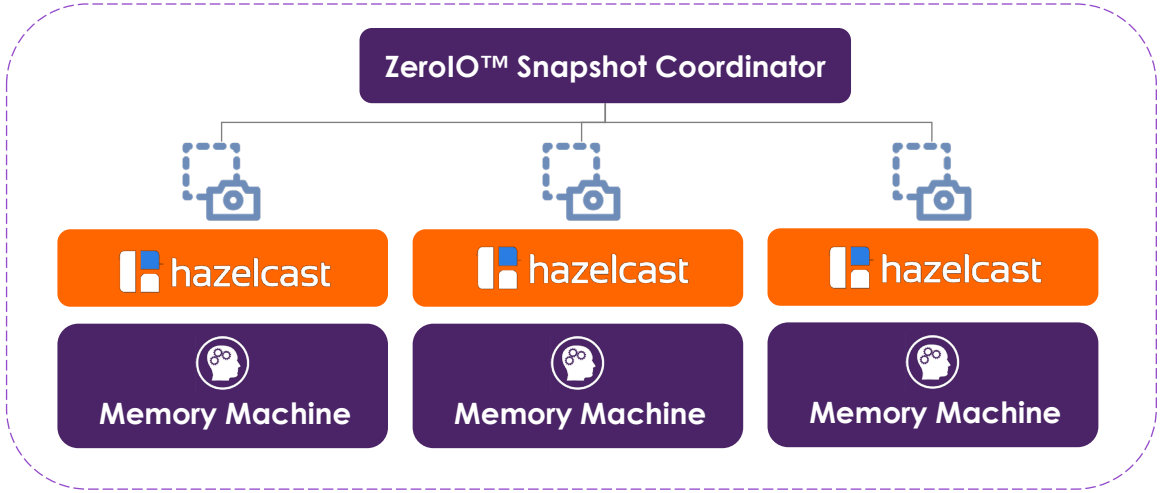


Recover from storage


Recover from PMEM

# Genomics



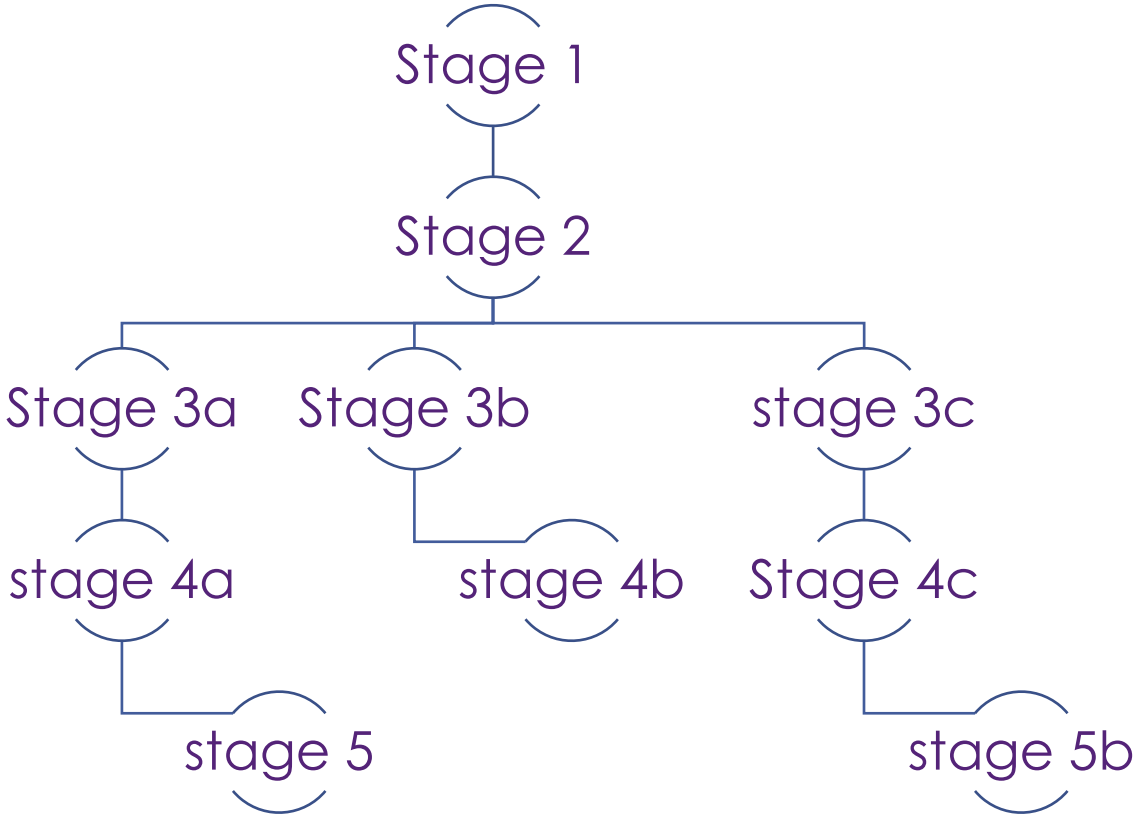


## Workload Attributes

- Large datasets
- Multi-stage pipeline
- Requires frequent checkpoints of intermediary stage results
- Frequent Rollbacks to tune parameters
- Branching to support what-if analyses

## Pain Points

- Checkpoint to disk and rollback extremely time-consuming
- Data loss risk
- Computation memory intensive

# Big Memory Solution



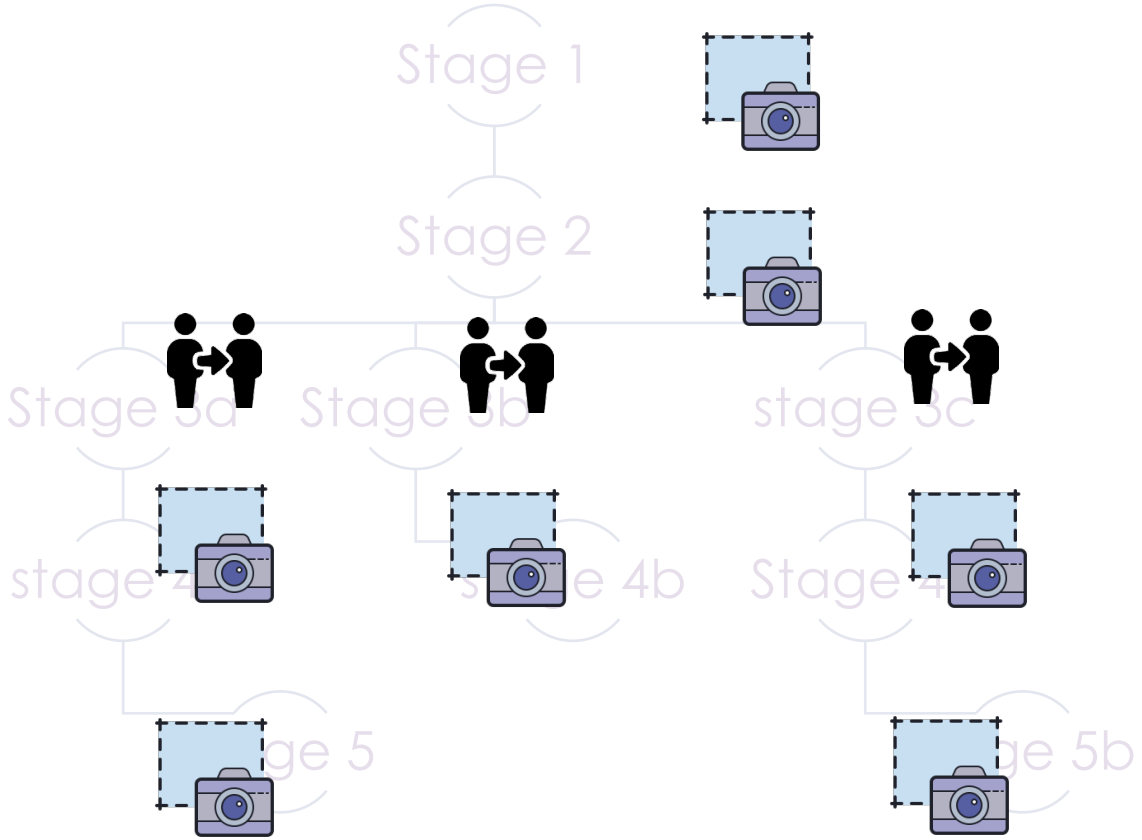


## Memory Snapshots

- ZeroIO™ snapshot: zero disk I/O
- Instant rollbacks
- Protects against data loss

## Cloning

- Automatic Dedupe via copy-on-write
- Supports multi-branch what-if analyses

## Memory Capacity

- Virtualize physical memory types
- Optimized performance

# Time to Restore a Compute Stage



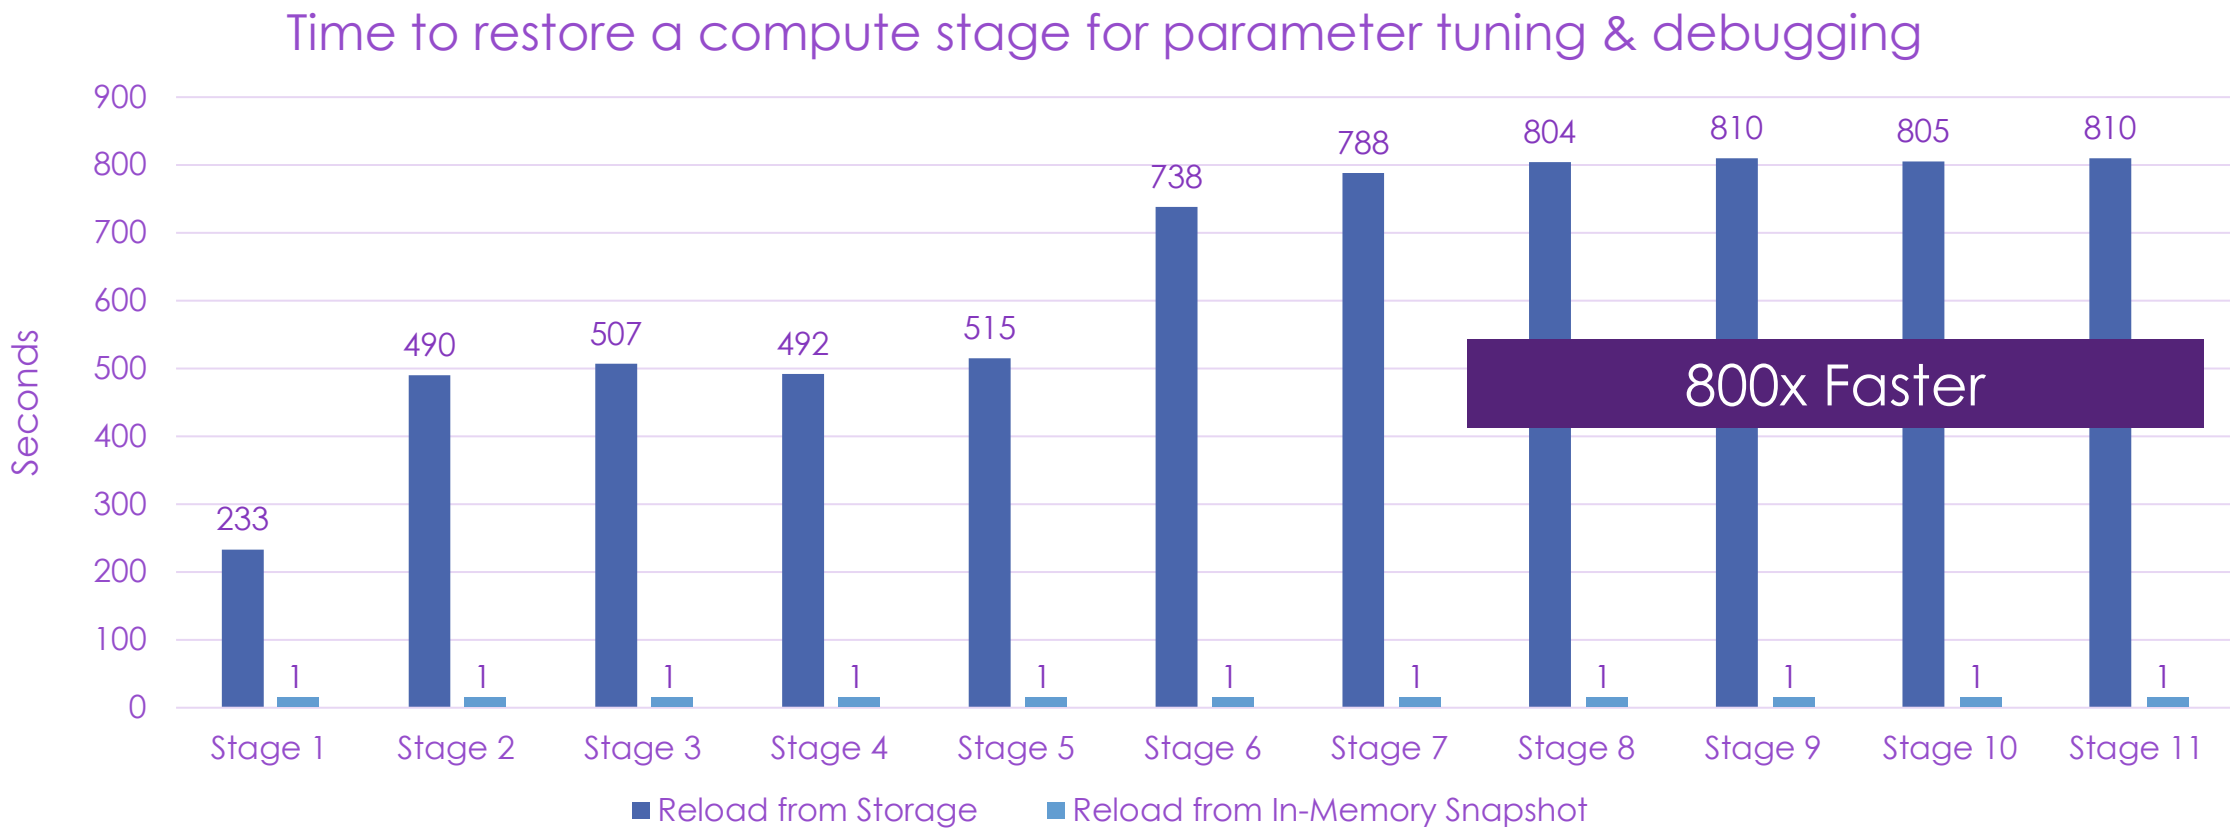

# Execution Time



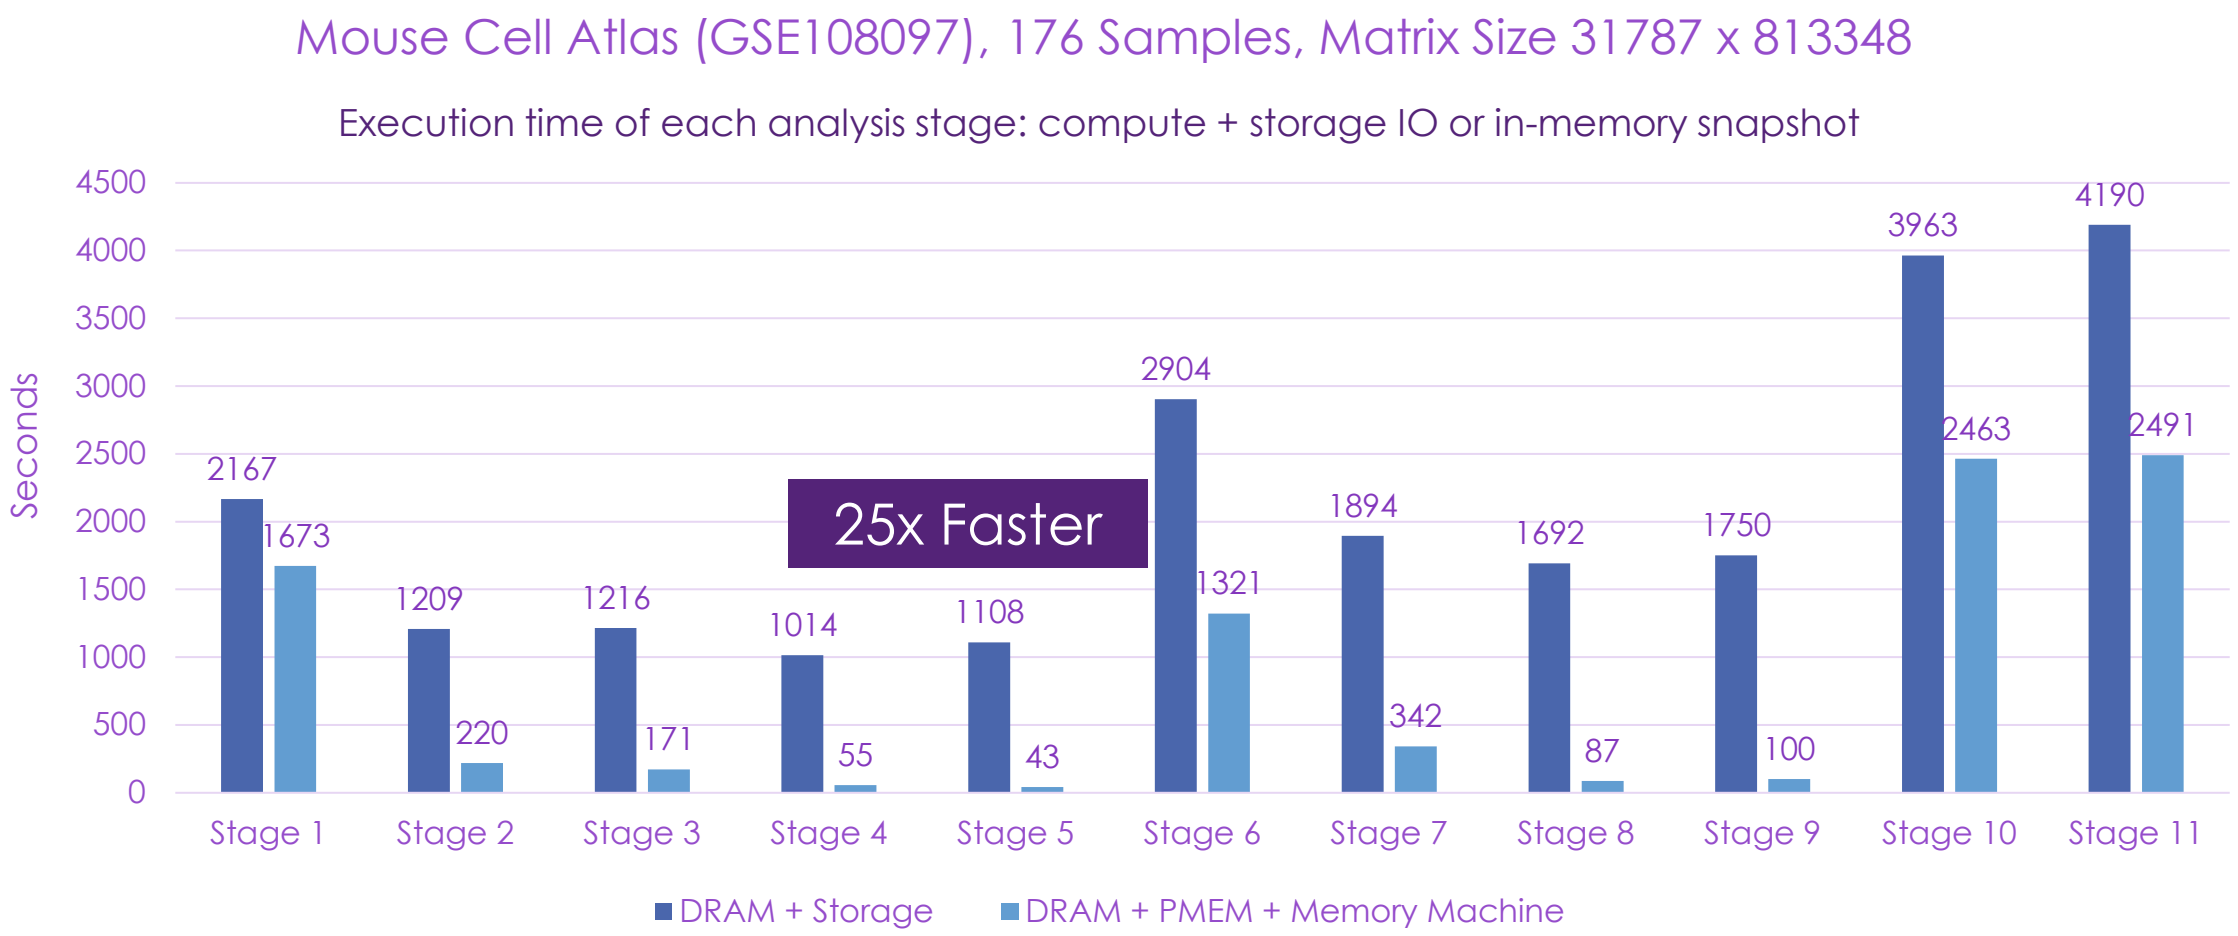

# Time Savings



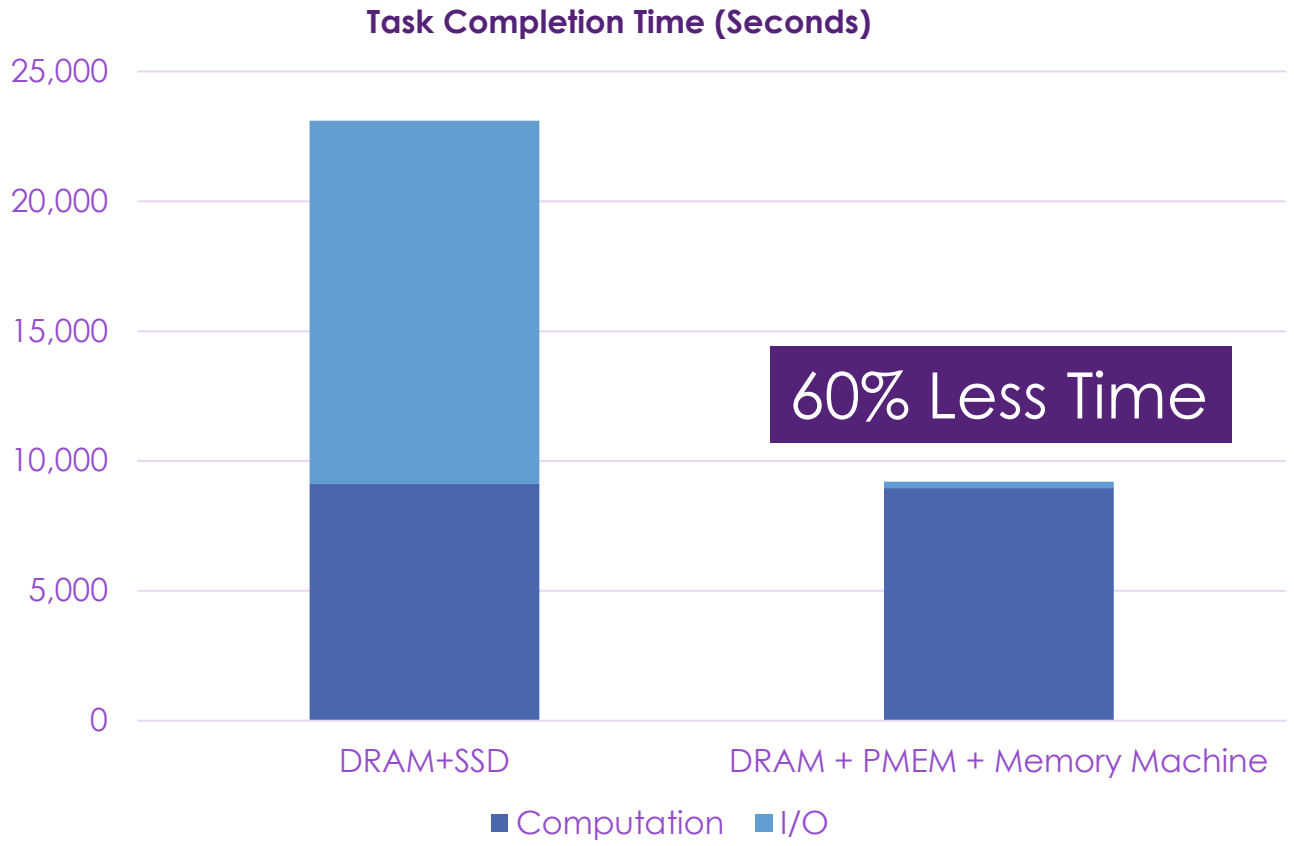


95% Less IO

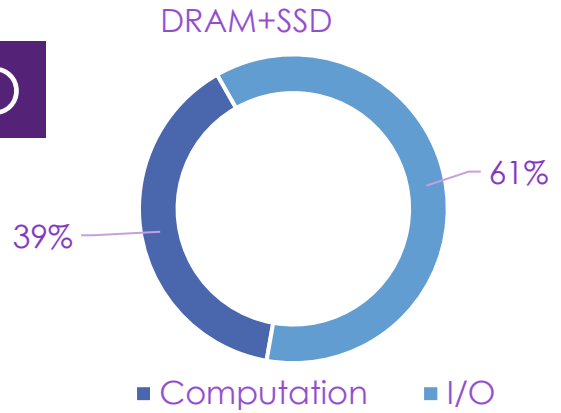


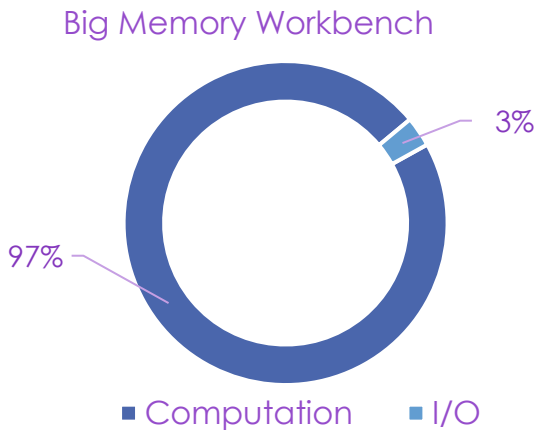

# Animation & VFX



- Bleeding edge visual effects provided by fragile plug-ins
- Frequent crashes
- Artists out of the zone after 30 seconds

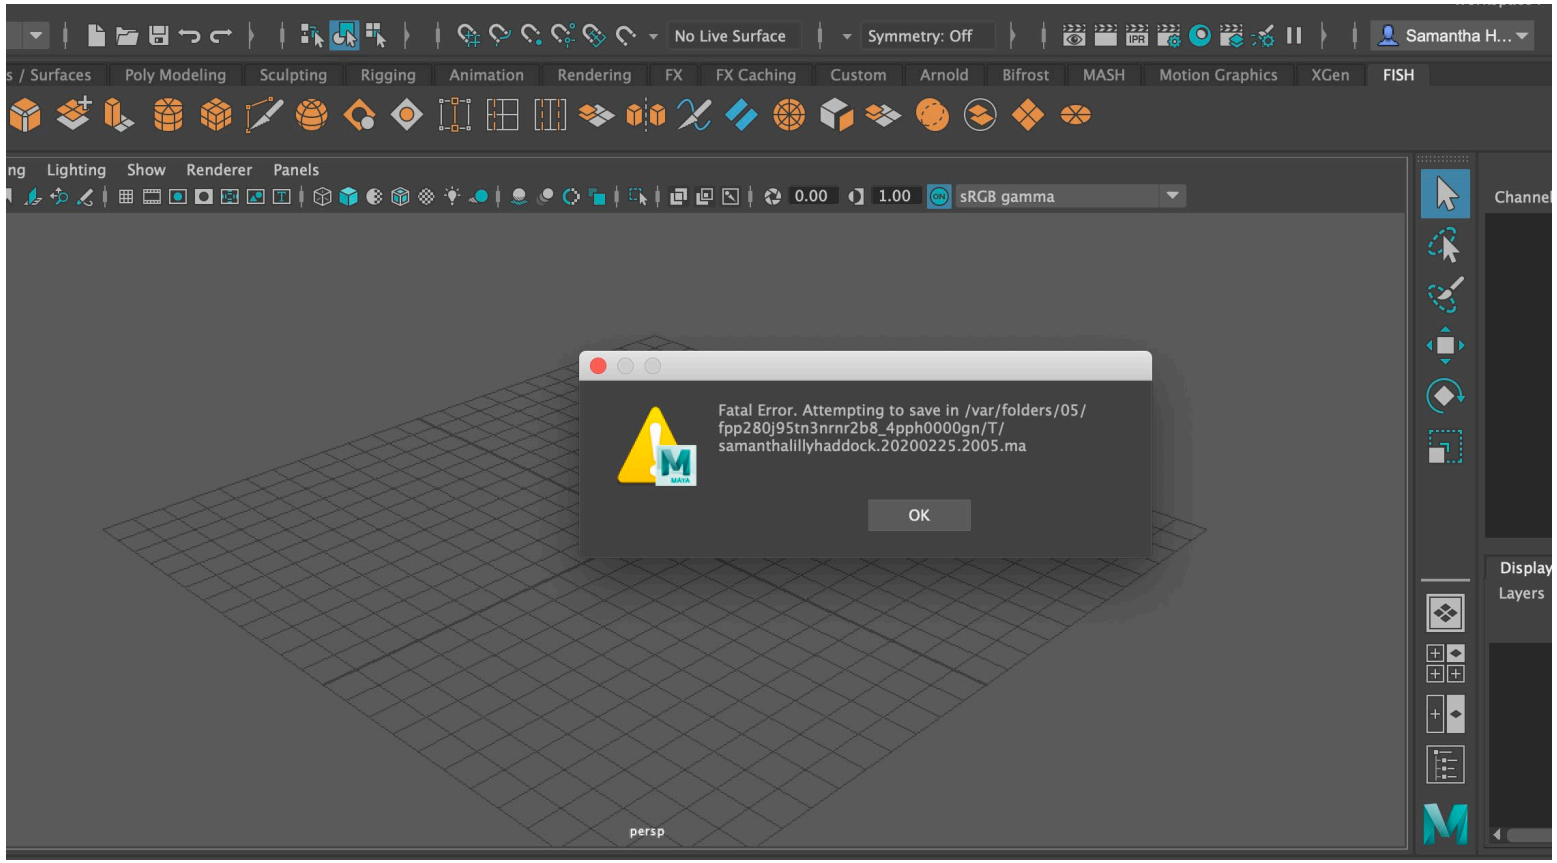

# Animation & VFX



## Keeping artists in the zone

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

Manual Snapshot

Restore


**With Memory Machine: Memory snapshots and instant recovery**

Auto Save


Restore


"Initially, we opened a poly-dense scene in Maya and it took two-and-a-half minutes. Then, we opened a scene from a snapshot we'd taken with Memory Machine and it took eight seconds. In addition to opening exponentially faster, another benefit of the Memory Machine snapshot is that it gets an artist right to the spot in the application where they were when they created a snapshot, there's no need to repopulate the entire application." - Mark Wright, Technology Manager for Chapeau Studios

# Tomorrow: Big Memory Unleashes Composable Memory Infrastructure



"…the ultimate vision of composable infrastructure includes a comprehensive range of disaggregated resources, including multiple processor, memory, cache and storage types. **Today, composable infrastructure is held back by a lack of technology to disaggregate DRAM from processors, industry-standard configurable fabrics** and cross-vendor APIs." – *Gartner: Understand the Hype, Hope and Reality of Composable Infrastructure*

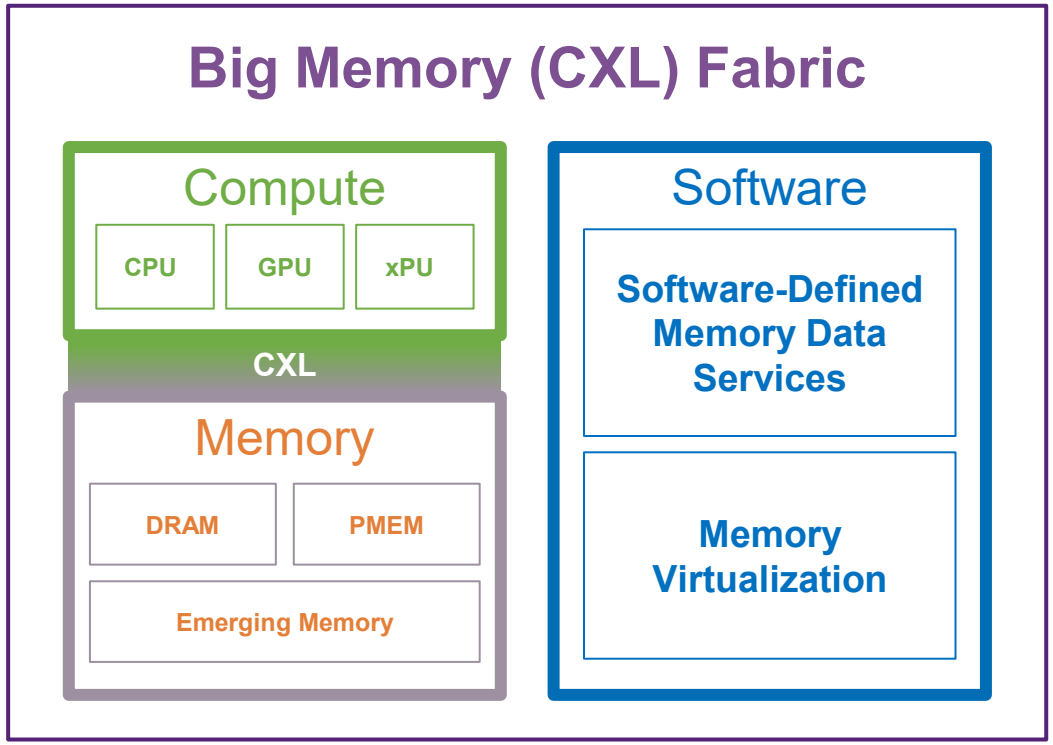

# Memory Tier Leverages New Technology Without Changes to Apps



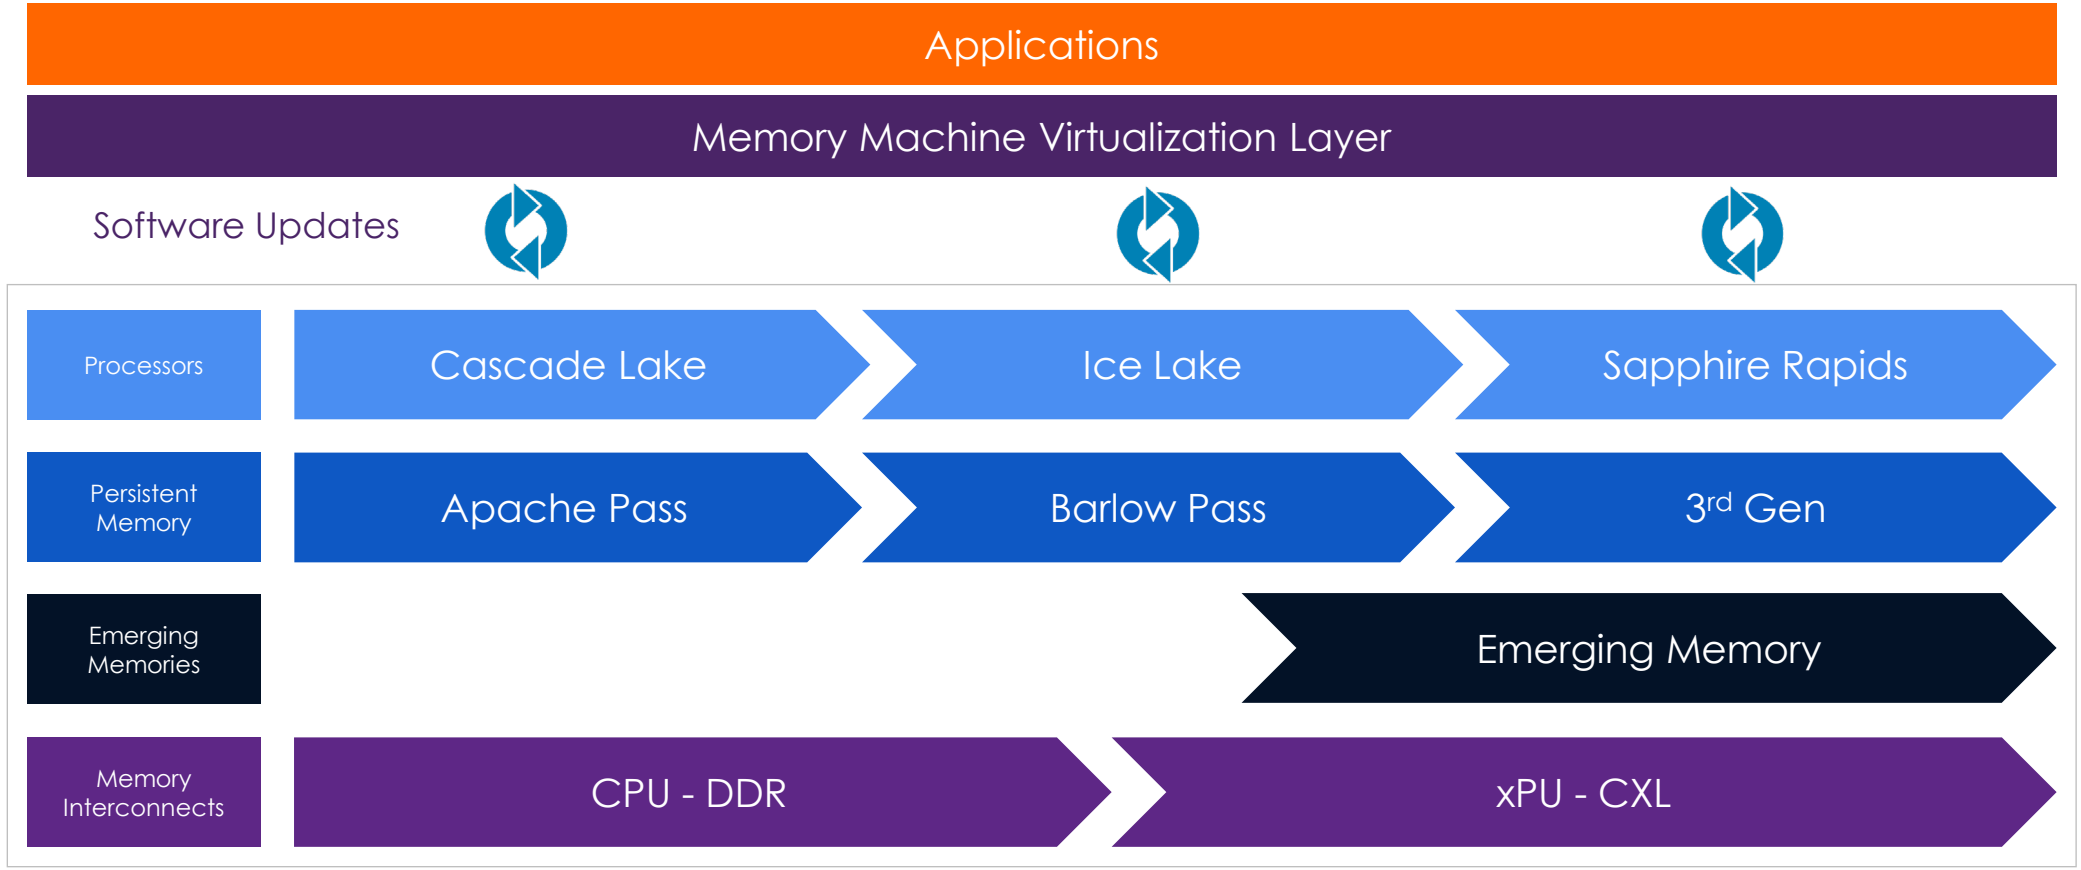

# The New Memory Pyramid



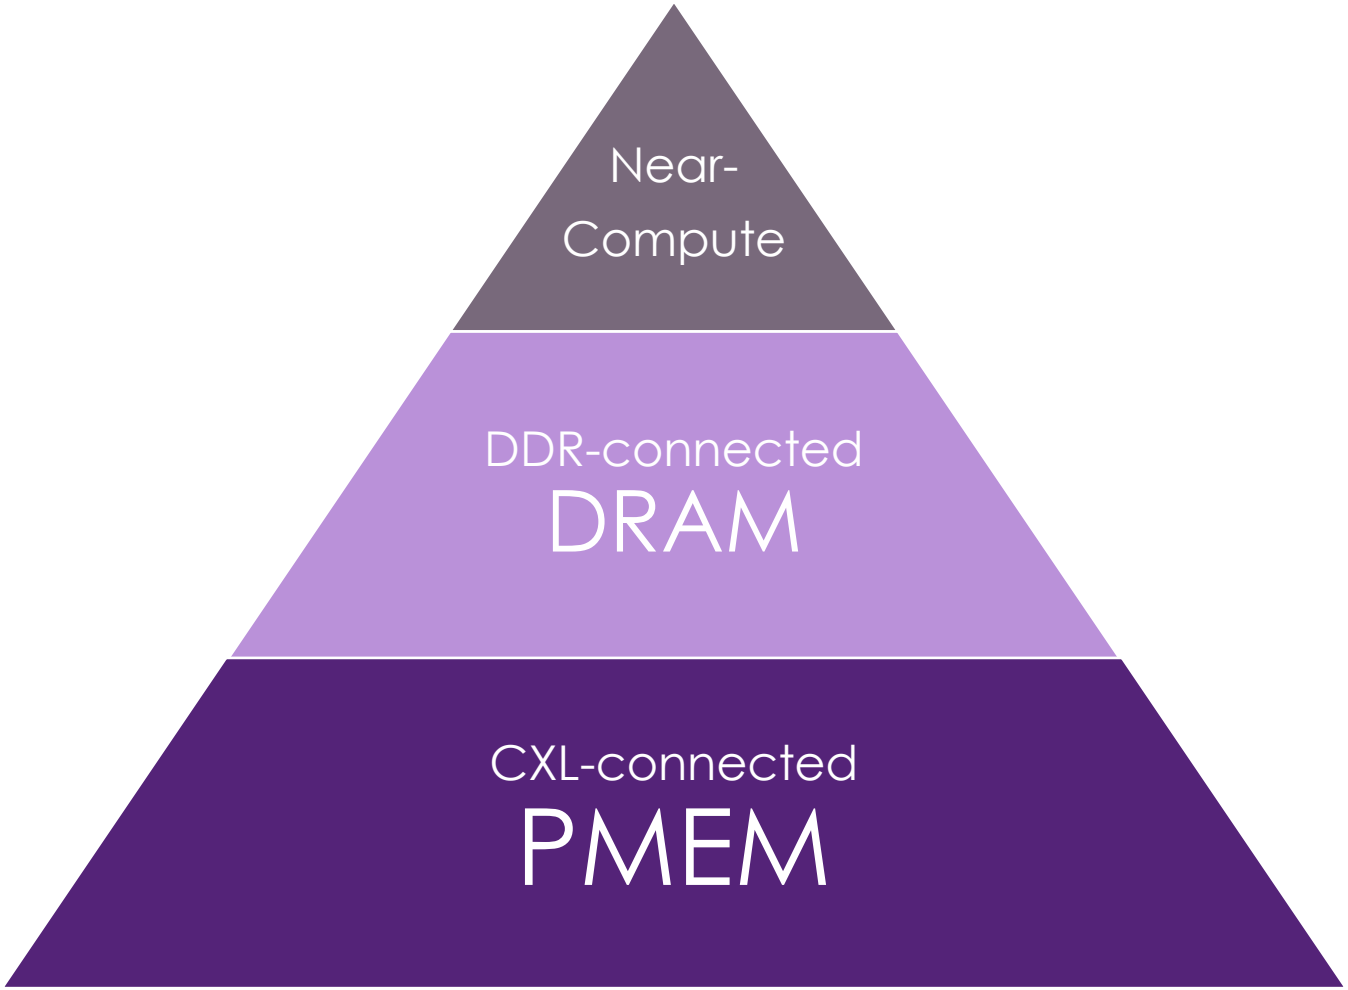

# Contact Us



Info@memverge.com

# Thank you

Please visit www.snia.org/pm-summit for presentations