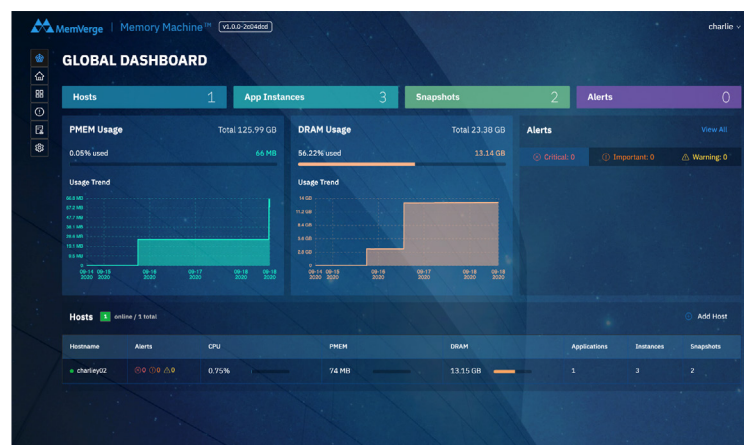


Memory Machine™

Enterprise-Class Big Memory Software

Key Features & Benefits

- Memory Virtualization Platform –** Memory Machine virtualizes DRAM and Persistent Memory so that data can be accessed, tiered, scaled, and protected in-memory.
- Software-Defined Memory Service –** Compatible with existing applications. Provides access to persistent memory without changes to applications. It looks like DRAM to your apps.
- Tier persistent memory and DRAM for maximum capacity and optimum performance –** Unlike approaches that cache PMEM data in DRAM, which lowers useable capacity, Memory Machine utilizes 100% of the DRAM and PMEM capacity. For performance, hot data is intelligently moved to the DRAM “fast tier” while warm data is placed in the PMEM “persistent tier.”
- Low-latency memory replication –** For apps like Pub/Sub, Memory Machine uses RDMA to publish high volumes of messages to remote subscribers with ultra-low latency.
- Recover TBs of data in seconds with ZeroIO™ memory snapshots –** To protect large memory blast zones, Memory Machine provides memory snapshot & replication for lightning-fast crash recovery that requires zero IO to storage.
- Clone databases in seconds –** Snapshots can be used for developers and other applications to clone terabyte database instances while sharing existing memory.



For the Expanding Universe of Real-Time Apps

According to IDC, real-time data was less than 5% of all data in 2015 but is projected to comprise almost 25% of all data by 2024. They project that by 2021, 60-70% of the Global 2000 will have at least one mission-critical real-time workload.

This Big Bang of real-time data is driving the expansion of real-time analytics and AI/ML applications into the mainstream. Another result is real-time applications are out-pacing the capacity, performance, and availability capabilities of in-memory infrastructure.

MemVerge responds to the need for a modern in-memory computing model with Memory Machine, the first in the new class of Big Memory software that virtualizes DRAM and persistent memory so that memory can be accessed without code changes, scale-out in clusters to provide the capacity needed by real-time analytics and AI/ML apps, and deliver enterprise-class data services for high-availability.

Memory Machine Standard Edition: The Best Way to Use Persistent Memory to Expand Memory

For businesses that need to address the cost of their growing memory footprint without compromising performance, Memory Machine Standard Edition provides applications with transparent access to persistent memory that emulates DRAM and is volatile like DRAM. Memory Machine also powers lower-cost persistent memory to perform like DRAM. Instead of modifying multiple applications, IT organizations can now efficiently deploy Big Memory across a data center with a single Memory Machine virtualization layer.

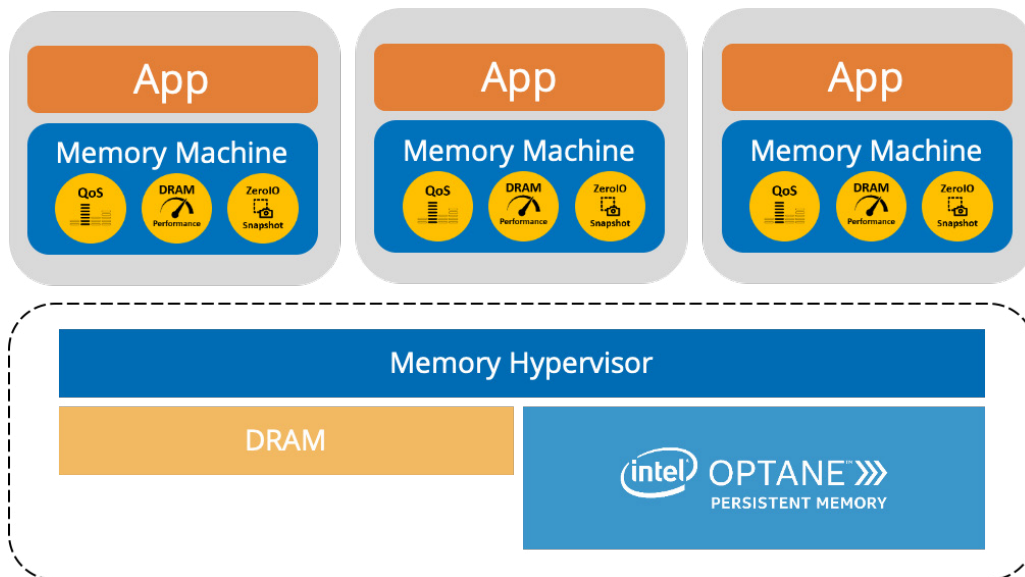
Memory Machine Advanced Edition: Enterprise-Class Memory Data Services

Memory Machine Advanced Edition includes all the capabilities of Memory Machine Standard Edition, plus transparent access to persistence via the industry's first enterprise-class data services for highly available memory.

IT organizations rely on data services for their disk and all-flash storage, especially snapshots and replication which enable crash recovery. Memory Machine's patented ZeroIO™ memory snapshot technology eliminates IO to storage. The result is the ability to snapshot and recover terabytes of data from persistent memory in a few seconds, instead of taking minutes to hours from storage. In addition, Memory Machine's Pub/Sub feature provides a lightning fast in-memory message bus that can be used to achieve high availability through data replication.

Software Defined Memory Architecture

Memory Machine is a Linux-based software subscription deployed on a single server or in a cluster. Once DRAM and persistent memory are virtualized, the Memory Machine hypervisor can guarantee performance to specific applications with QoS, set-up DRAM as a "fast" tier for persistent memory and provide enterprise-class memory data services based on ZeroIO™ memory-to-memory snapshots.



A Memory Machine can be configured for each app so the 3 major capabilities can be tailored to the specific needs of each app. In addition, 4 enterprise-class data services are available which are based on ZeroIO Snapshots: Time Travel, AutoSave, Thin Clone and App Migration.

Key Capabilities of Software Defined Memory



DRAM Performance

A Memory Machine breakthrough is the ability to make a virtualized pool of DRAM and Intel Optane Persistent Memory go DRAM-fast. The “DRAM Performance” feature allows organizations to tune the amount of DRAM and PMEM allocated to a single app or to multiple apps. Organizations can now maintain DRAM performance while enjoying the lower cost, higher capacity, and persistence of PMEM.



Quality of Service (QoS)

When multiple applications are sharing memory, non-critical “noisy neighbors” can disrupt the performance of business-critical applications. The Memory Machine QoS feature allows server admins to isolate applications and prioritize access for those business-critical applications.



ZeroIO™ Snapshots


Memory Machine features the world’s first memory-to-memory snapshot technology. The ability to snapshot data from DRAM to persistent memory transforms the volatile low-availability memory tier into a high availability tier. Traditional snapshots and recovery of large in-memory data sets to-and-from storage are disruptive because they take minutes to hours. Memory snapshots of large data sets are non-disruptive because they take as little as one second. Organizations can now achieve high availability by deploying frequent snapshots and slashing crash recovery time to seconds. Memory Machine ZeroIO Snapshots are also the foundation of the 4 data services below.

Enterprise-Class Memory Data Services



Time Travel

Memory Machine allows an application to roll back to any of the snapshots previously taken, therefore creating “time travel” for that application.



AutoSave

Memory Machine can be configured to automatically take snapshots at predetermined intervals to protect an app from the effects of a crash. After a crash, the app will automatically recover to the last snapshot taken.



Thin Clone

Clones of in-memory databases are frequently made so DevOps or other applications can have access to the database. The process can be complex, lengthy and requires additional memory. The Memory Machine Thin Clone service uses ZeroIO snapshot technology to create another instance of the last snapshot of the in-memory app in a couple of seconds, and without additional physical memory.



App Migration

After an in-memory snapshot is taken, this snapshot can be moved to another server, and a new application instance can be created from this snapshot.

Supported In-Memory Applications

Kx: kdb+	Time-series in-memory database
Autodesk: Maya	3D animation and rendering
Apache Software Foundation: Spark	In-memory big data analytics framework
Redis Labs: Redis	In-memory key-value database
Apache Software Foundation: TensorFlow	Machine learning framework
Open Source: PyTorch	Machine learning framework
Apache Software Foundation: Flink	Stream processing framework
Apache Software Foundation: XGBoost	Gradient Boosting framework
Hazelcast: Hazelcast	In-memory data grid
Apache Software Foundation: Cassandra	Wide column store, NoSQL database
Apache Software Foundation: RocksDB	Persistent key value database
Apache Software Foundation: HBase	Non-relational distributed database
GNU: MySQL	Relational database management system
BSD: Memcached	In-memory key-value database
Linux Kernel Community: KVM	Hypervisor
MongoDB, Inc.: MongoDB	NoSQL database
R Foundation: R	Environment for statistical computing and graphics
** MemVerge is continuously qualifying new apps. Support for apps is available upon request.	

Supported Environments

Hardware	Second Generation Intel® Xeon® Scalable Processors, with Optane DCPM memory
Operating Systems	CentOS: 8.1 with kernel 5.7.4; 8.0 with kernel 4.20; 7.8 with kernel 5.7.6; 7.7 with kernel 3.10.0-1062.12.1.el7.x86_64, 3.10.0-1062.9.1.el7.x86_64; 7.6 with stock kernel or kernel 4.20. RHEL: 8.1 or 8.0 with stock kernels.
Hypervisors	VMware ESXi running on above HW (vSphere 6.7, 7.0), QEMU-KVM
Containers	OpenShift, VMware Tanzu, and other Kubernetes based container platforms