The Skinny

on

Memory Machine™

MemVerge

What is Memory Machine and what does it do?

For memory-intensive, long-running apps, big data has become unmanageable because it takes minutes-to-hours to load, save, replicate, and recover. MemVerge® Memory Machine™ is an in-memory data management platform that makes data manageable again with the unique ability to load, save, replicate, and recover terabytes of data in seconds.

Memory Machine intelligently allocates and manages system memory in data center servers. Awareness of memory locations, and how frequently they are accessed, is maintained so that Memory Machine can ensure optimal placement of data in memory.

Memory Machine in-memory data management software

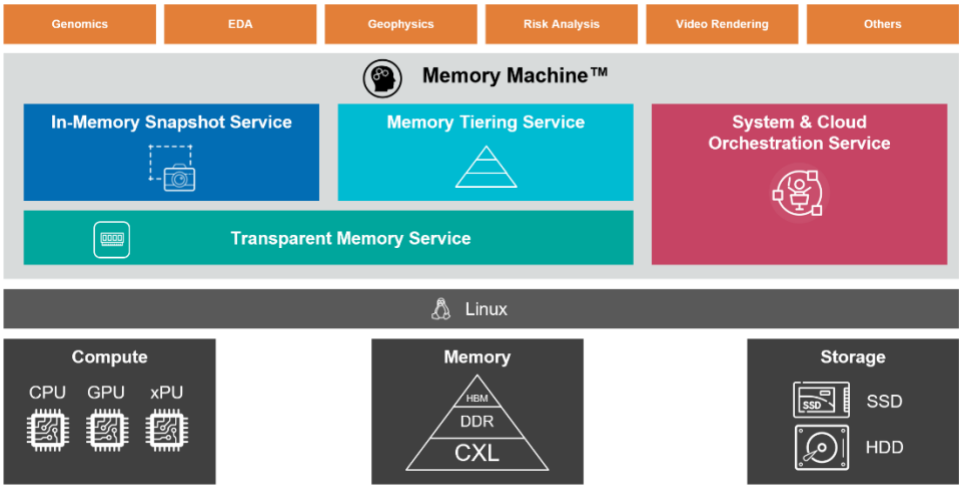

Metadata allows Memory Machine to capture ZeroIO™ In-Memory Snapshots, i.e., a complete moment-in-time picture of an application's state, including in-memory data and open files. In-memory snapshots are the foundation for in-memory data management including checkpoint-restore (for high availability), application cloning (for parallelizing execution pipelines), and in-memory storage (for reducing I/O overhead due to SSD access).

A few examples of snapshot-based in-memory data management

Memory Machine can solve challenges faced by customers today

Problem	Memory Machine solution
1. Customers have long-running, stateful applications that do not have robust checkpoint-restore functionality. If the application encounters a problem, the application must restart from the beginning.	Memory Machine takes ZeroIO™ In-memory Snapshots so that if an application instance encounters a problem, it can be restarted from any valid snapshot. Note: a list of supported applications is available.
2. Some applications take a long time to run to completion because of stages that incur time-consuming storage I/O — either loading data into memory or writing data to SSD.	Memory Machine keeps application snapshots in memory (although they can be written to SSD as well). A snapshot contains all the in-memory data that the application is using so that a snapshot can replace a write to an SSD-based file system. If the application instance needs to be rolled back to a previous stage (e.g., for ML parameter tuning), a snapshot can instantiate the application at that stage without a time-consuming I/O process. Note: the snapshot restarts the same application and snapshots cannot be accessed by other applications.
3. Some analytical pipelines take a long time to complete because there is not an easy way to parallelize some stages.	Memory Machine can create an AppCapsule — all the in-memory state and data associated with a running application instance (including open files residing on SSD). AppCapsule can be used to instantiate additional application instances. Consecutive pipeline stages that can be parallelized can be started simultaneously without requiring any storage I/O. Note: the degree of parallelization that can be achieved depends on the analytical pipeline and the dataset.

Customer problem	Memory Machine solution
4. Moving a stateful application instance from on-premises to the cloud (or vice versa), or from one cloud to another, renders the application unavailable for an unacceptably long period.	App Capsule can move an application instance from one compute platform (e.g., on-premises) to another (e.g., in the cloud), and restarted in seconds. Note: a list of the applications that can use AppCapsule is available. The target compute platform must use the same processor as the original compute platform.
5. Some applications cannot run because there is not enough memory (e.g., bioinformatics pipelines with large datasets). Other applications take a long time to execute because of limited memory — datasets must be processed in batches.	Intel® Optane™ Persistent Memory (PMEM) provides high-capacity memory but is not straightforward to use. Using Memory Machine, applications can access *all* the memory capacity (DRAM+PMEM) without requiring code changes. Note: when used in Intel Memory Mode, the DRAM portion is not available to the application.
6. Customers, especially Cloud Service Providers, want to improve financial metrics related to virtual servers by using less expensive, but higher capacity, memory.	With Memory Machine, some DRAM in a server can be replaced with less expensive PMEM without compromising performance, thereby reducing the average cost per GB. The extra memory can be used to host additional virtual machines on the server, improving total revenue and gross profit. Note: Financial performance varies and is not guaranteed by MemVerge.

What makes Memory Machine unique?

Transparent Memory Service

Although other companies have publicly recognized the importance of memory virtualization (VMware, for example), MemVerge is unique in having production-ready software available today.

The technology that powers Memory Machine is the Transparent Memory Service (TMS), which integrates different memory technologies into an abstraction layer that any application can access as a single, homogeneous, byte-addressable memory domain.

Memory Machine Transparent Memory Service

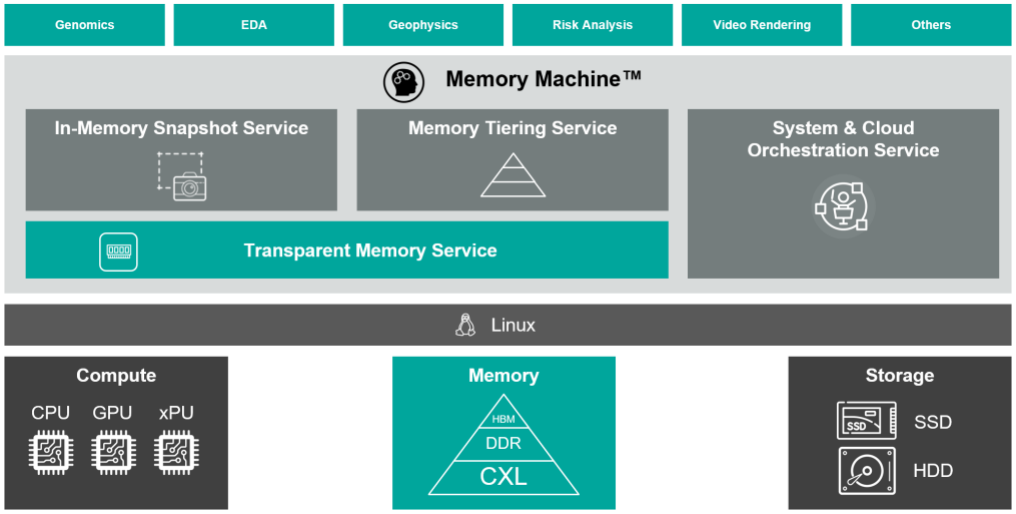

TMS is a form of memory virtualization, a capability that is increasing in importance as new memory types appear in the market – Intel Optane Persistent Memory today and CXL-based products in the next few years.

**The CXL interconnect will be the backbone,
and in-memory data management software the heart, of CXL fabrics of the future**

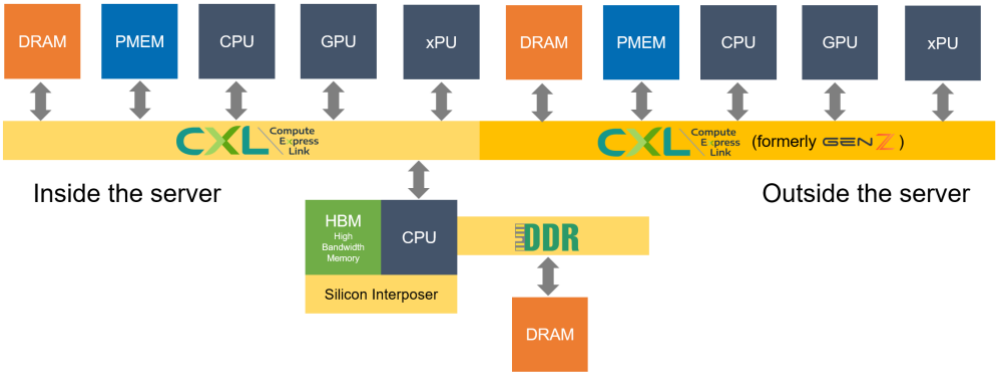

In-Memory Snapshot Service

Forms of checkpoint-restore (C/R) exist today, both open source and proprietary. Memory Machine is unique in providing C/R features without incurring storage I/O, which offers significant performance gains over traditional file system-based C/R approaches.

Memory Machine Snapshot Service creates the world's first in-memory snapshots. They are the foundation of in-memory data management which is needed by big memory apps to load, save, replicate, and recover terabytes of data in seconds.

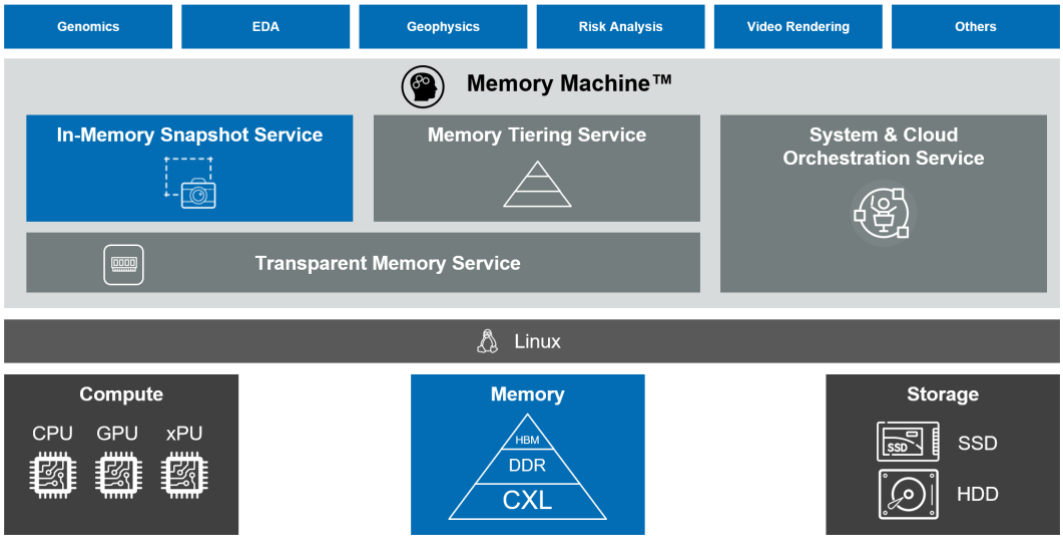

To take a ZeroIO In-Memory Snapshot, Memory Machine captures the entire in-memory state including execution state, open files and sockets, and associated data.

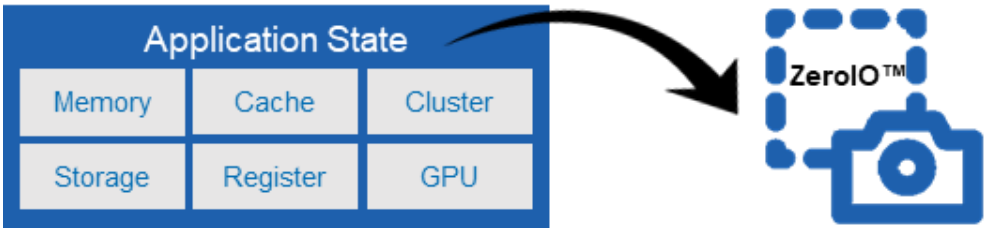

Compared to a C/R method that requires writing in-memory state to SSD, a ZeroIO snapshot pauses an application instance for a fraction of the time and does not incur the overhead of writing to disk. ZeroIO snapshots use a copy-on-write technique so that memory pages are not copied until their content is changed by an application instance. When a series of ZeroIO snapshots is taken, successive snapshots are incremental, that is, the change from the previous snapshot is captured rather than the entire in-memory state.

How Memory Machine works

Memory Machine is a user-space application that runs on popular Linux operating systems (Red Hat, Centos, and Ubuntu). The available memory capacity is virtualized, so that although an application sees a homogeneous memory domain, Memory Machine is managing memory in tiers. One tier — the performance tier — is always DRAM. The second tier — the capacity tier — is sometimes PMEM but it can be DRAM.

Memory Machine manages memory allocation in the two tiers by customizing components in common C libraries that are preloaded at run time. After allocating memory pages, Memory Machine monitors memory access so that "hot" pages can be moved into the performance tier and "cold" pages can be moved into the capacity tier. Additional optimizations are achieved through the intelligent use of Huge Pages. As a result, Memory Machine has complete visibility into how memory is being used, and thus, each application instance's complete in-memory state is available.

Managing memory allocation in a 2-tier environment

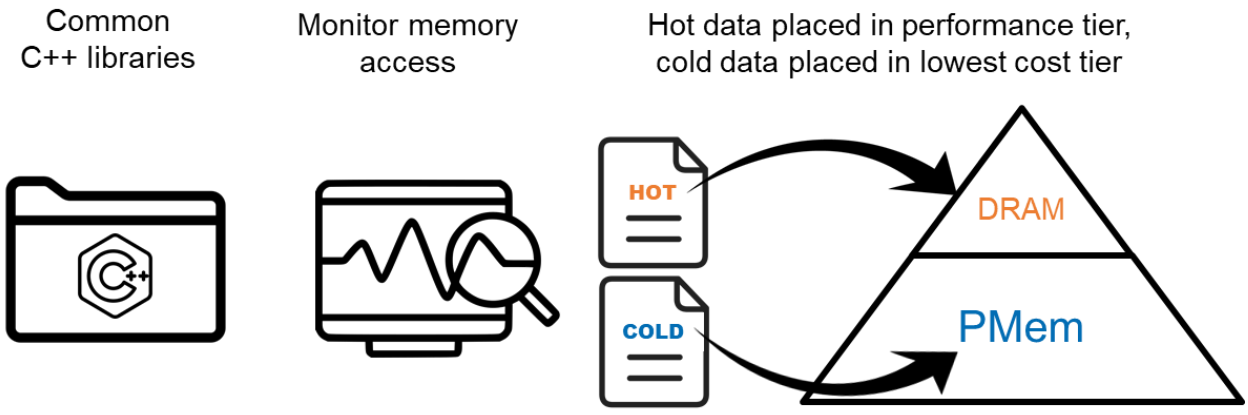

ZeroIO In-memory Snapshots begin by momentarily freezing the application instance so that data in the CPU caches can be flushed to their memory locations. The entire process tree is traversed to identify the associated memory locations and open files. The metadata is stored in the capacity memory tier, which is also non-volatile storage if the tier is composed of PMEM. If it is not PMEM and non-volatile storage is required, the metadata can be written to an SSD-based file system.

Use Cases

Memory Machine has been successfully deployed in production environments including:

- Bioinformatics
- Electronic Design Automations (EDA)
- Financial Services
- Hyperscale cloud service providers

Customer Success Stories

Analytical Biosciences (ABio) is a leader in providing single-cell RNA sequencing analysis services. Using Intel Optane Persistent Memory and MemVerge Memory Machine software, ABio was able to load datasets up to 800 times faster and to reduce execution times by a factor of 25 for some computational stages, compared to a traditional DRAM-based solution.

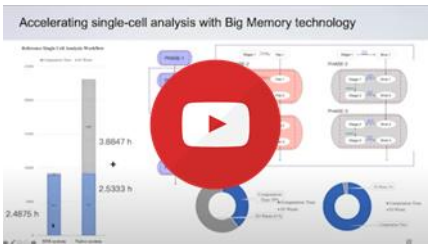

Dr. Chris Kang of ABIO explains how Memory Machine eliminated IO and slashed time-to-discovery by 58%.

The Huck Institutes of the Life Sciences at Penn State selected MemVerge Memory Machine to speed up genomic analysis of plant DNA in research to develop sustainable plant products. By deploying a pool of software-defined DRAM and PMEM, the life sciences institute has been able to speed up their analysis pipelines and run workloads that it has not been able to process before.

The Translational Genomics Research Institute (TGen), a non-profit research organization, deployed Memory Machine and PMEM to improve their time-to-discovery. Total pipeline run time was reduced from 6 ½ hours to a little over 4 hours (a 36% reduction). This was made possible by faster data load times, faster code execution with all data in-memory, and by using in-memory snapshots for parallel code execution.

TGen bioinformatician Austin Gutierrez describes how Memory Machine transformed their pipeline from single-threaded to multi-threaded and reduced time-to-discovery by 36%.

Hazelcast is the leading in-memory computing platform that enables organizations to leverage a resilient and elastic memory resource for data at rest and in motion. The Hazelcast platform combined with MemVerge memory virtualization and in-memory snapshots enables high-capacity memory and fault tolerance for applications in Financial Services such as credit card processing and risk analysis for Basel III Compliance.

Dale Kim of Hazelcast talks about how Memory Machine enriched fraud detection for credit card processing.

Learn more about Memory Machine

If you have a workload that would benefit from running on a Memory Machine-managed memory platform, contact one of our Solution Architects for a free, no obligation consultation.

Resources

MemVerge Website

Video: Big Memory and MemVerge Overview

Big Memory Cloud Website

IDC Report: Digital Transformation Driving New Big Memory Requirements

International Biopharmaceutical Industry article: Accelerating Single-Cell Genomic Sequencing with Big Memory

Solution Brief: Big Memory Accelerates Single-Cell RNA Sequencing

Demo Video: Memory Machine Cloud Fault Tolerance Service

TGen presentation useR! 2021

Analytical Biosciences presentation at HPC User Forum