



SUPERMICRO SOLUTIONS FOR THE MEMVERGE® MEMORY MACHINE™

Enterprise-class Big Memory Platforms Optimized for Expanding the Universe of Real-time apps.



CloudDC SuperServer SYS-120C-TN10NR



Ultra SuperServer SYS-120U-TNR



BigTwin® SuperServer SYS-220BT-HNTR

TABLE OF CONTENTS

- Executive Summary 1
- Overview of MemVerge Technology 2
- Supermicro – MemVerge Solution Advantages 2
- Joint Solution Architecture 3
- Use Cases 3
- Case Study - Big Memory Accelerates Single-Cell RNA Sequencing 4
- Solution Reference Architecture Configuration 5
- Conclusion 7

Executive Summary

According to IDC, real-time data will comprise almost 25% of all data by 2024. This Big Bang of real-time data is driving the expansion of real-time analytics and AI/ML applications into the mainstream. However, it’s not uncommon for data sets to exceed the size of system memory, creating the need to read and write data to a from storage (storage I/O), making real-time analytics almost impossible, and extending the resource required for large data analytics. Therefore, implementing data-driven, software-defined architectures that meet the demands of those data sets is critical.

Supermicro and the MemVerge Memory Machine solution offer integrated hardware and software that virtualizes DRAM and persistent memory so that all of the memory can be accessed without code changes and can scale out in clusters to provide the capacity needed by real-time analytics and AI/ML applications



SUPERMICRO

Supermicro (Nasdaq: SMCI), the leading innovator in high-performance, high-efficiency server and storage technology is a premier provider of advanced server Building Block Solutions® for Enterprise Data Center, Cloud Computing, Artificial Intelligence, and Edge Computing Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green®” initiative and provides customers with the most energy-efficient,

and deliver enterprise-class data services for high-availability. Data has become one of the most significant drivers of our economy, businesses, and IT infrastructures. Leveraging a powerful, software defined architecture to aggregate the performance and capacity of DRAM and persistent memory eliminates the roadblocks and enables innovative workflows—an optimized platform with the architecture, integration, support, and managed services to ensure the success of critical applications

Overview of MemVerge Technology

MemVerge® Memory Machine™ is an in-memory data management platform that creates a memory virtualization layer between applications and the underlying memory hardware. Its architecture assumes there are two tiers of memory, the first composed of DRAM and the second composed of DRAM or PMem. By intelligently managing the two tiers of memory, Memory Machine optimizes memory performance with reduced Total Cost of Ownership (TCO) of server infrastructure. In addition, it makes data manageable and provides higher availability with the unique ability to load, save, replicate, and recover terabytes of data in seconds, greatly accelerating the time to discovery.

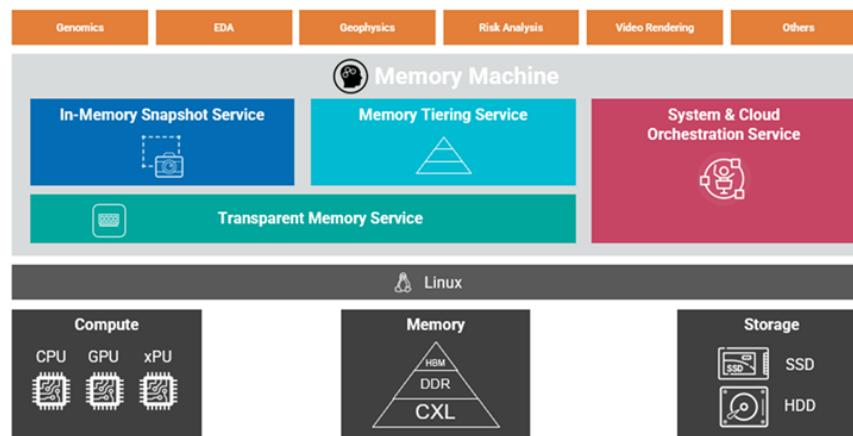


Figure 1 - MemVerge Memory Machine overview

Supermicro – MemVerge Solution Advantages

Supermicro and MemVerge have partnered to develop a best-in-class solution based on industry-leading SuperServers and Memory Machine™. The solution delivers an Enterprise-Class Big Memory platform to dramatically expand the in-memory computing power with compelling new capabilities, clear Total Cost of Ownership metrics, and data durability.

- Memory Virtualization Platform: Memory Machine virtualizes DRAM and Persistent Memory so that data can be accessed, tiered, scaled, and protected while in memory
- Software-Defined Memory Service: Compatible with existing applications. Provides access to persistent memory without changes to applications. The persistent memory looks like DRAM to applications.
- Tiered memory 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 snapshots & 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

Joint Solution Architecture

Supermicro MemVerge® Memory Machine™ Ready Server Platforms solution architecture Diagram:

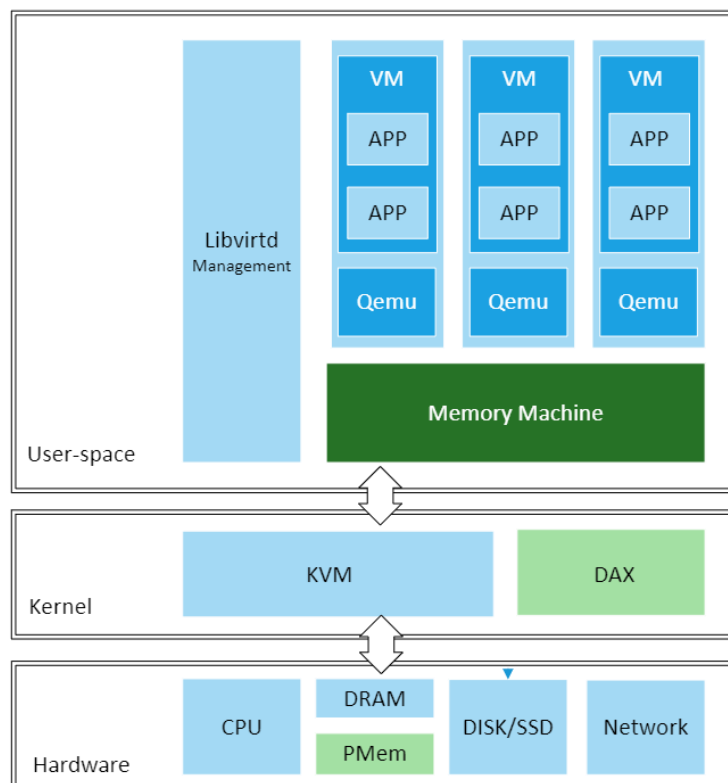


Figure 2 - Supermicro and MemVerge Joint Solution Architecture

MemVerge Use case:

When 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.

Some analytical pipelines take a long time to complete because there is no 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.

Some applications cannot run because there is insufficient 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 unavailable to the application.

Case Study - Big Memory Accelerates Single-Cell RNA Sequencing

Working with a leader in Single Cell RNA Sequencing and Analytics, Analytical Biosciences, the big memory data science workbench can dramatically increase the productivity of the genomic computation pipeline. Using snaps instead of storage IO, the data loading time for every stage is significantly reduced, up to 800X for some stages. It can save the overall execution time by at least 60% and delivers higher productivity by allowing data scientists to roll back and branch to perform what-if analysis.




Figure 3 - Example of Restore Time Comparison

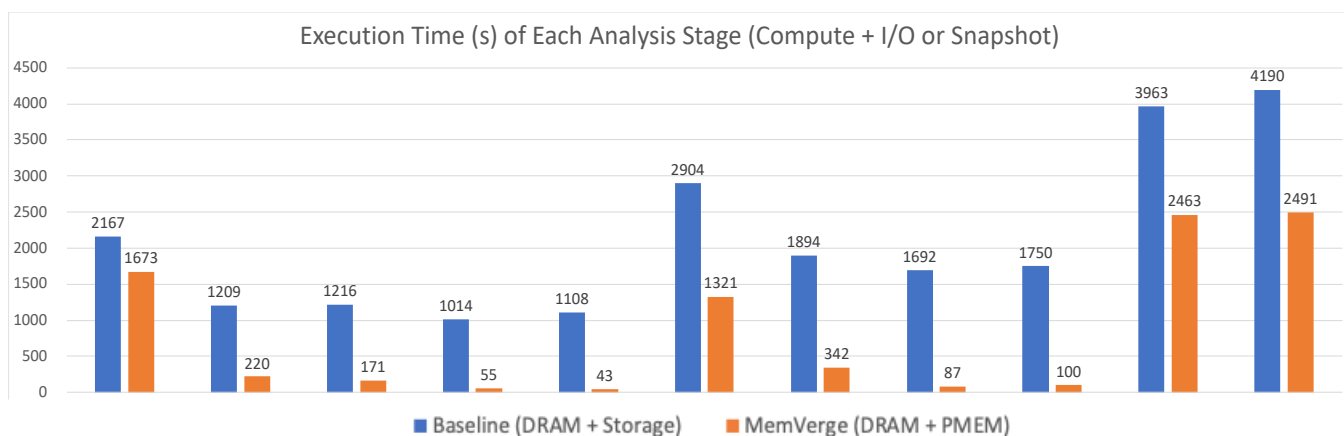


Figure 4 - Example of Execution Time comparison

Solution Reference Architecture Configuration

Supermicro MemVerge® Memory Machine™ reference architecture comes with three options for different use cases, two with rackmount server and one with a multimode system. Each option has two different configurations based on the capacity of DRAM and PMEM. As a result, MemVerge® Memory Machine™ Ready Server Platforms enable users to benefit from a two-tier memory hierarchy management to achieve improved performance when running any memory-intensive application on Linux.



Model SKU	Entry	Entry	High-End	High-End	Multi-node	Multi-node
	MV-EN-1000	MV-EN-2000	MV-ADV-4500	MV-ADV-5000	MV-MUL-1000	MV-MUL-1500
Form Factor	1U1N	1U1N	1U1N	1U1N	2U4N	2U4N
Server	SYS-120C-TN10R	SYS-120C-TN10R	SYS-120U-TNR	SYS-120U-TNR	SYS-220BT-HNTR	SYS-220BT-HNTR
CPU	2x ICX 4314	2x ICX 4314	2x ICX 5320	2x ICX 5320	2x ICX 4314	2x ICX 4314
Memory	128G (16G x8)	256G (32G x8)	512G (32G x16)	1024G (64G x16)	128G (16G x8)	256G (32G x8)
Optane PMEM	1024G (128G x8)	2048G (256G x8)	4096G (256G x16)	4096G (256G x16)	1024G (128G x8)	1024G (128G x8)
Boot Drive	2x 512G M.2	2x 512G M.2	2x 512G M.2	2x 512G M.2	2x 512G M.2	2x 512G M.2
NVMe	5.7T	5.7T	5.7T	5.7T	7.68T	7.68T
NIC Card	Mellanox® CX-6 DP	Mellanox® CX-6 DP	Mellanox® CX-6 DP	Mellanox® CX-6 DP	Mellanox® CX-6 DP	Mellanox® CX-6 DP

Figure 5 - Reference Architecture Configurations

MemVerge Memory Machine Software:

MemVerge has developed two versions of software for providing different levels of capabilities:

Memory Machine Standard Edition

The 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:

- Provides access to a pool of DRAM and (volatile) PMem without changes to apps
- Delivers higher performance and capacity than Intel Memory Mode
- More composable than Intel Memory Mode with visibility and control on a per-app basis
- Runs on bare metal, in VMs, and Containers

Memory Machine Advanced Edition

The 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:

- The suite of services powered by in-memory snapshots
- ZeroIO™ in-memory snapshots capture memory state in seconds
- Built on industry-standard DAX APIs for accessing persistent memory
- Supermicro systems configured in as a solution reference architecture

Supermicro MemVerge® Memory Machine™ Ready Server Platforms:

Entry - To provide baseline performance benefit brought from PMem and best cost-performance ratio out of three models (sizes)

CloudDC SuperServer SYS-120C-TN10R



- Dual sockets P+ (LGA-4189) 3rd Gen Intel® Xeon® Scalable Processors
- Intel® C621A Chipset
- 16 DIMMs up to 6TB 3DS ECC DDR4-3200: LRDIMM/RDIMM
- Intel® Optane™ Persistent Memory 200 series
- 2 PCI-E 4.0 x16 FHHL; 2 PCI-E 4.0 x16 AIOM (OCP 3.0); 2 PCI-E 3.0 x2 NVMe M.2
- Dual AIOM with NCSI (OCP 3.0) for networking, 1 dedicated IPMI LAN
- 10x 2.5" hot-swap hybrid NVMe/SATA/SAS drive bays
- 6 counter-rotating 4 cm PWM fans with optimal fan speed control, 2 air shrouds
- 860W redundant Platinum level 100-240Vac and 200-240 Vdc power supplies
- 1 VGA, 1 COM, 2 USB 3.0 (rear), 2 USB 2.0 (front)

High-End - To provide the most memory capacity of three models (sizes).

Ultra SuperServer SYS-120U-TNR



- Dual Socket P+ (LGA-4189) 3rd Gen Intel® Xeon® Scalable Processors
- Intel® C621A Chipset
- 32 DIMM Slots up to 8TB ECC DDR4-3200 LRDIMM/RDIMM;
- Intel® Optane™ Persistent Memory 200 series
- 1 Internal PCI-E 4.0 x16; 1 PCI-E 4.0 x16 (LP); 2 PCI-E 4.0 x16 (FH, 10.5"L)
- Flexible networking options
- 12x2.5" hot-swap hybrid NVMe/SATA/SAS drive bays

- 8 heavy duty fans with optimal fan speed control
- 1200W redundant Titanium level power supplies

Multi-Node - To provide baseline performance benefit from PMEM and higher density of the compute cluster. This is the best space efficiency out of the three models (sizes).

BigTwin SuperServer SYS-220BT-HNTR



- Four hot-pluggable systems (nodes) in a 2U form factor. Each node:
- Dual Socket P+ (LGA-4189) 3rd Gen Intel® Xeon® Scalable Processors
- Intel® C621A Chipset
- 16 DIMM Slots up to 6TB ECC DDR4-3200 LRDIMM/RDIMM;
- Intel® Optane™ Persistent Memory 200 series
- 2 PCI-E 4.0 x16 (LP) slot; Internal PCI-E 3.0 x8 for 2 M.2 NVMe/SATA support onboard;
- Optional NVMe Boot Controller via AOC-SMG3-2M2-B
- Network connectivity via AIOM (OCP 3.0 compliant)
- 6 Hot-swap 2.5" drive bays; 6 PCI-E 4.0 NVMe/SATA 2.5" drives
- 4 cooling fans per 2U enclosure, 16.5K RPM; Shared Cooling Design; Liquid Cooling Support
- 2600W Redundant Power Supplies Titanium Level (96%+); Shared Power Design

Conclusion and Resources

Real-time data is driving the expansion of real-time analytics and AI/ML applications into the mainstream and extending the need for large data analytics. Therefore, implementing data-driven, software-defined architectures that meet the demands of those data sets is critical.

Supermicro and MemVerge offer a solution based on the MemVerge Memory Machine with integrated hardware and software that unlocks Intel Optane PMem Performance, Capacity, Availability & Productivity. By virtualizing DRAM and persistent memory, the 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. Success will be assured with this optimized platform and the architecture, integration, support, and managed services it provides.

For more information, please visit Supermicro or MemVerge at: www.supermicro.com