

# Big Memory Computing at the Turning of the Tide

by DCIG Lead Analyst Storage, Ken Clipperton



#### COMPANY

MemVerge 1525 McCarthy Blvd, Suite 218 Milpitas, CA 95035 +1.408.823.3902

#### memverge.com

#### INDUSTRY

In-Memory Computing Software-defined Memory

SOLUTION

Memory Machine

Persistent memories are now shipping in production volumes from multiple sources. At present, they are primarily being applied in embedded solutions and in a narrow set of high-value use cases. However, the time has come for enterprise technologists to begin planning their technology refreshes around a composable, datacentric, big memory architecture.

## **Dynamic Multi-dimensional Data Growth**

Organizations are capturing data from more sources and in a greater variety of file sizes than in the past. As resolutions increase, so do working set sizes. Consider that an 8K camera can generate up to 128 Gbps of data. That is nearly 1 TB per minute. Producing a feature film can easily require 2PB+ of storage capacity.

The expanded use of structured and especially unstructured data has created new dimensions of scalability that go beyond mere capacity. Organizations across many industries are seeking to use the data they capture in a new and varied ways. These include using Al-enabled video analytics to optimize retail operations or traffic flows in a city, and using in-depth analysis of high-resolution microscopy images to improve healthcare. (Figure 1.)

Many enterprises are now running workloads that until recently were the domain of academic and scientific high-performance computing (HPC) environments. In the process, these organizations are running into infrastructure performance limits on multiple fronts, whether compute, network, storage performance, or memory.



## **DRAM is Stretched to its Limits**

Although successive generations of the DDR specification have increased the performance of DRAM, these advances have not kept pace with advancements in CPUs.

Since 2012 CPU core counts have tripled, but bandwidth per CPU core has declined by 40%.<sup>1</sup>

DRAM capacities have also stagnated. The advances in DRAM capacity that have been achieved come at a high cost per bit, such as seen in HBM (High Bandwidth Memory.)

"Memories of all types are facing pressures as demands grow for greater capacity, lower cost, faster speeds, and lower power to handle the onslaught of new data being generated daily," Bryon Moyer, Semiconductor Engineering.<sup>2</sup> (Figure 2.)

#### Memory Trends: BW vs. CPU Core Count<sup>3</sup>







Figure 1



# The Decades-long Search for Next Generation Memory Key infrastructure providers such as Intel, IBM, Micron, and Samsung have long recognized that DRAM technology was reaching its limits in terms of both capacity

and bandwidth. Thus, these companies have for decades been exploring alternate memory technologies. These technologies include MRAM, FRAM, ReRAM and PCM. For a long time, these research efforts did not make it out of the lab and into general availability. That has changed.

# Storage Class Memory is Here in Production Volumes

Today multiple providers are shipping what is popularly called "storage class memory" in production volumes. These memories are coming from multiple major chip foundries including TSMC, Samsung, Global Foundries, and others.

The initial SCM leaders, in terms of shipped capacity, are 3D XPoint (Intel Optane) and MRAM. In 2021, MRAM shipments crossed the 100 TB threshold and 3D XPoint shipments are estimated at 1,000 petabytes.<sup>4</sup> (Figure 3.)



#### **Growth in New Memory Shipments**<sup>5</sup>

#### Figure 3

A key attribute of all storage-class memories is that they are persistent. Unlike DRAM, the data stored in SCM stays in memory even if the system is powered off. In that regard SCM is storage, like an SSD or thumb drive.

Another key attribute of SCM is the "M." It is a memory technology. CPUs and other data processing units can communicate with SCM using load/ store memory-semantics. These move data with minimal overhead and latency compared to disk-oriented protocols. Memory-semantics also enable efficient, fine-grained data access through byte-addressability.

Solution providers are currently using MRAM primarily in embedded devices, so we won't be loading up our PC's memory slots with MRAM anytime soon. However, providers are incorporating Intel Optane (typically 16 GB or 32 GB modules) into desktop PCs and laptops to accelerate initial boot times and launch times for individual programs.

## Intel Ice Lake (Finally) Delivers a Performance Bump

In the datacenter, Intel's Ice Lake datacenter CPU, announced in April 2021, offers up to 6 terabytes of system memory composed of DRAM plus Optane Persistent Memory (PMem). This expanded memory

capacity is teamed with up to 40 CPU cores, and up to 64 lanes of PCIe Gen4 per socket. This represents a 33% increase in memory capacity and a 250% increase in I/O bandwidth.

## The Need for Big Memory

As an example of the difference more memory can make, Intel tested start times for a 6TB SAP HANA database on the new system with PMem compared to a prior generation system that did not include PMem. Both systems had 720 TB of DRAM. The PMem-equipped system started in 4 minutes instead of 50 minutes, a 12.5x improvement. (Figure 4.)



#### Figure 4

This greater-than-10x improvement provides a proof point for the performance benefits that can be achieved when a working set fits entirely in memory instead of overflowing to the I/O subsystem. Unfortunately, some working sets are larger than 6TB. Whenever a working set exceeds memory capacity, the I/O subsystem comes into play. This introduces a significant drop in performance.

Even in the example cited above, Intel loaded the data into memory from an I/O device, in this case an NVMe SSD over the PCIe interface. Four (4) minutes is certainly better than 50 minutes, but if you are editing a movie that is long enough to disrupt the creative flow. That flow that may take more than an hour to recover. When persistent memory is available to a process and is bigger than the working set, recovery can be nearly instantaneous.

## Storage Class Memory is Necessary, but Insufficient

Storage class memory provides a much-needed boost in performance when integrated into a data center environment. By itself, however, it is not sufficient to deliver big memory that meets the performance demands of many real-time high-value applications, or to enable the next wave of workload consolidation.

## The Rise of Big Persistent Memory

Instead, we need a new memory infrastructure built around a new memory reality. The DRAM era of memory—volatile, scarce, and expensive—is drawing to a close. It is being replaced by the era of big memory—large, persistent, virtualized, and composable. The big persistent memory era is upon us.

Some aspects of this era are inherent to the technology. It is big (relative to DRAM) and persistent by design. It is just now beginning to open up a new set of opportunities to create value. For example, a narrow set of

high-value applications are already taking advantage of these attributes to accelerate discovery in the research center and reduce fraud in the financial industry.

#### Today: Data I/O. Tomorrow: Load/store.

Today, nearly all applications move data using traditional IO protocols. The adoption of SCM into infrastructures will enable a transition to much more efficient memory-semantics, yielding substantial benefits.

We are still in the early phases of a persistent-memory-enabled revolution in performance, cost and capacity that will change server, storage system, data center and software design over the next decade.

Nevertheless, multiple vendors are now shipping storage class memory as part of their enterprise storage systems. The revolution has begun.

## **A New Memory Infrastructure**

To make big memory broadly available to a wide range of standard workloads requires an ecosystem that goes beyond the storage class memory itself. Standards such as DDR5, PCIe 5.0, and NVDIMM-P supply the bandwidth and other attributes needed to take advantage of SCM. We also need memory virtualization, composable infrastructure standards, and compliant components to turn this big memory into a fabric-attached pool of memory that can be allocated and orchestrated much like containers are today. Now that SCM is actually available, the ecosystem is falling into place.

## Remain Calm, Look for Short Term Wins, Integrate Big Memory into Your Plans

We humans tend to overestimate short term impacts and underestimate long term impacts. Today, we are at the turning of the tide for big memory computing. The time has come for enterprise technologists to look for high-value tactical wins today, and to begin planning infrastructure refreshes around a composable, data-centric, big memory architecture.

#### Sources

- 1. CXL Consortium Memory Challenges and CXL Solutions [August 2020 Webinar] https://www.brighttalk.com/webcast/17915/425684
- 2. https://semiengineering.com/more-data-more-memory-scaling-issues/
- 3. De Dios & Associates. Compute Express Link<sup>™</sup> (CXL<sup>™</sup>): Memory Challenges and CXL Solutions. Compute Express Link<sup>™</sup> and CXL Consortium are trademarks of the Compute Express Link Consortium. https://www.brighttalk.com/webcast/17915/425684
- 4. snia-pm-cs-summit-Coughlin-Handy-Dynamic-Trends-2021-corr 5-10-21.pdf
- 5. SNIA Persistent Memory + Computational Storage Summit

#### **About DCIG**

The Data Center Intelligence Group (DCIG) empowers the IT industry with actionable analysis. DCIG analysts provide informed third-party analysis of various cloud, data protection, and data storage technologies. DCIG independently develops licensed content in the form of TOP 5 Reports and Solution Profiles. More information is available at **www.dcig.com**.



© 2021 DCIG, LLC. All rights reserved. This Executive White Paper is a product of DCIG, LLC. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. Product information was compiled from both publicly available and vendor-provided resources. While DCIG has attempted to verify that product information is correct and complete, feature support can change and is subject to interpretation. All features represent the opinion of DCIG. No negative inferences should be drawn against any product or vendor not included in this report. DCIG cannot be held responsible for any errors that may appear.