

# Case Study

Public Sector  
Embracing Digital Transformation



# Practical Use Cases of Optane Persistent Memory with MemVerge

Darren W Pulsipher & Charles Fan – June 7, 2021

**Persistent memory is changing the way people think about computing. Three areas are already taking massive leaps forward, thanks to MemVerge.**

Darren talks with Charles Fan, CEO of MemVerge, about use cases with their breakthrough technology in this episode. Their software utilizes Intel's Optane persistent memory innovatively, removing the bottleneck between memory and storage by combining them into a single pool of resources resulting in cost reduction, reliability, performance, and productivity.



[Youtube Channel](#)

**Podcast:** [SoundCloud](#)

**Keywords:** EDT, EDT51, Embracing digital, Optane Persistent Memory, PMEM, Genomic Sequencing, Large memory database, PMEM in the Cloud, MemVerge

MemVerge shipped the first version of their software in September 2020, and despite the pandemic, which slowed down educating customers about the new technology, a good pattern of use cases emerged.

## Cloud Service Providers Cost Reduction

The key metric for a cloud service operator is how many VMs they can deliver to their customers and at what cost. The size of the memory on the servers becomes the bottleneck of how many VMs they can allocate per server, limiting how low their price per VM can go.

MemVerge software with Optane delivers a larger amount of memory per server, allocating a larger number of VMs, therefore lowering the cost per VM and increasing the competitiveness of cloud service providers. The cost per VM could be three times cheaper.

## Reliability with Large Memory Data Bases

Financial customers such as stock exchanges, banks, and mutual funds use a lot of memory databases and in-memory applications. In addition to making their memory bigger so they can have more instances per server, MemVerge solves memory database availability challenges. If the data is not being persisted to storage all the time, but just sitting in memory, all the intraday data is lost if there is a crash. This is catastrophic. Even if you have logged all the transactions, you must replay the log to recover the database, which takes many minutes or even hours to recover.

MemVerge offers a new data service that has in-memory snapshot. It persists the database state to Optane, which is much faster than persisting onto storage. If there is a crash, you have the last snapshot captured on persistent memory and you can recover from that. The recovery only takes a minute or two, so it is a 60 to 100x improvement.

## Genomic Sequencing Reduction Through Memory Snapshot

In the area of genomics, MemVerge software in combination with Optane increases productivity exponentially. In a multistage data analytics workflow, bigger memory means more parallelism of the pipeline and processing so the whole process is faster. Snapshot is handy here, as well. If an organization is doing, for example, cancer or COVID research, and they need to do DNA or RNA sequences, they have to go through something like 50 stages of

processing. Each stage might take hours, and they need to take a checkpoint of the state of intermediate computation results for a few reasons: first, re-running or reproducing results, and second, comparing results if they modify some data. The checkpoints are saved onto storage and this takes from five to 30 minutes. In many cases, this can be more time consuming than the compute itself. So if one job takes 24 hours, they might use eight hours for compute and 16 hours just doing these IO jobs saving those intermediate states.

Rather than doing IO, MemVerge uses a snapshot after each stage and captures it on Optane persistent memory. Instead of 16 hours of IO, this process can take one minute. It's the new way of doing IO; you don't need to do the serialization or de-serialization to open a file, read, write, etc. All you have to do is take a snapshot.

Although this does take a lot of memory, with MemVerge, the memory is bigger than before, and it will continue to improve as Intel innovates. Two other features help with this issue. First, snapshots are taken periodically without creating full copies of the

memory state; they are only the change pages so the extra usage of memory is minimized. Second, MemVerge can keep up to 256 layers of snapshots to memory, but at the same time, you can export those snapshots off memory to storage servers or your own storage systems. This is done without interrupting or impacting your running application.

Essentially, you are creating a memory DVR because instead of only running your application forward, you can also run backward almost instantly. It's a new experience.

Genomics is just the first example of many workloads that could benefit from this technology.

Since MemVerge is a startup, they narrowly focus on the three areas of cloud service providers, financial large memory applications, genomics and related data science pipeline jobs, but these use cases all prove the power of the combination of Optane persistent memory and MemVerge software.



Intel® technologies may require enabled hardware, software, or service activation.  
No product or component can be absolutely secure.  
Your costs and results may vary.  
Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.  
©2021 Intel Corporation