


Application Crashes No Longer Have Starring Role

Intel® Optane™ Persistent Memory and MemVerge Memory Machine: Helping VFX and Animation artists recover from application crashes with amazing speed



Data is growing. By 2025, the world will create three times more data than it did between 2015 and 2020—and the demand for real-time responsiveness for these big data workloads is increasing exponentially, as well. It's as if the highway in your town needed to accommodate five times more traffic and everyone wanted to drive much faster than before.

Businesses and IT have had little time to prepare for the latest data explosion. Nowhere is this more apparent than in the media and entertainment industry where visual effects (VFX) and animation teams bring filmmakers' visions to life. According to TechRadar, creating an animated film such as *Rise of the Guardians* could take as many as 65 million hours of footage to end up with 90 minutes of movie-worthy content¹.

All of that takes processing power and memory—typically DRAM—to handle the data-centric apps that VFX artists and animators use. Loading scenes to a workstation can often take an artist or animator 10 minutes or longer, depending on the complexity and the application used. Of course, if there's a bottleneck loading the data or saving it—or if an application crashes—that time can extend to hours and productivity and timelines suffer.

“Both the capacity and velocity of data continues rising. Apps that are either very memory hungry or move large amounts of data are bottlenecked either by memory or the performance of storage I/O,” says Charles Fan, co-founder and CEO of MemVerge. “In the past year, we've seen an uptick of companies with either problems with the size of their memory, or I/O bottlenecks, or recovering from crashes.”

Companies can spend more money on DRAM, but the memory is volatile, expensive and does not provide the capacity needed. Persistent memory modules are able to hold a much higher capacity of data and are extremely stable—but persistent memory typically requires changes to an application to work.



Image courtesy of Stefan Ivanov

The New Star: MemVerge Memory Machine* Built with Intel® Optane™ Persistent Memory (PMem)

The MemVerge Memory Machine is the first in a new class of big memory software that virtualizes DRAM along with Intel Optane persistent memory so that memory can be accessed without code changes to applications. Compatible with existing applications, the Memory Machine looks like DRAM to apps—and utilizes 100% of DRAM and the persistent memory.

This gives media and entertainment companies like Chapeau Studios the capacity needed for intensive creative applications such as Maya*, it also saves them from application crashes by taking an in-memory snapshot of the work and eliminating I/O to or from memory and storage. The result gives artists the ability to recover terabytes of data from persistent memory in a few seconds, instead of taking minutes to hours from storage.

“We have created the world's first big memory software which delivers software-defined memory service to applications without requiring changes to them. The in-memory snapshot captures what is going on at any given time and then restores the data should the application stop. In particular, this is useful for movie editing where studios have thousands of artists working worldwide on editing, lighting, layout, animation, and special effects,” says Fan.

Fan continues, “The ability to roll back to a previous state—either because of a need by the artist or because of an application crash—is a big selling point. Without big memory and a snapshot, the artist would lose all the work. A crash is really painful to recover from.”


Transform an hour of downtime into seconds



Quickly recover data after a crash²

256 GB < **10** sec

Intel® Optane™ persistent memory and big memory software lets you recover from software crashes in mere seconds, without losing your work.




How do Intel Optane persistent memory and big memory work?

- Big memory software takes a snapshot of the artist's work.
- The snapshot is placed in Intel Optane persistent memory.
- Data in persistent memory retains its formatting and is closer to the CPU than data in SSD or disk storage.
- After a crash, hundreds of gigabytes of data can be restored within seconds.



Chapeau Studios: building bridges between creative and technology

Chapeau Studios is a tech-forward VFX and design collaborator that integrates traditional film techniques with technology, user experience, and design. The studio works with some of the largest tech and consumer brands to create everything from traditional media experiences to out-of-home advertising to media walls with extreme VFX to animations.



Founding Partner and Creative Lead, Ben Loomam, says that “Chapeau’s foundation is in VFX. We’re known for integrating photo-real CG into live action photography, amongst multiple other disciplines. It’s in our DNA to try and create efficiencies with technology.”

Loomam isn’t kidding about needing efficiencies. One client—Facebook—reached out to Chapeau to create the social media giant’s “Friends Day” live action films. The films used Chapeau’s dynamic compositing process to create “live” video files, which were turned into millions of personalized videos. “Our first deliverable was 880 million films in 35 languages,” says Loomam.

Loomam, along with Mark Wright, Technology Manager for Chapeau, began using MemVerge Memory Machine on an HP Z8* workstation with Intel Optane persistent memory and Autodesk Maya (Maya) 3D graphics software—an industry-standard app for animators and VFX artists.

Wright talks more about the benefit of Memory Machine and Intel Optane persistent memory. “Let’s say an artist goes to lunch and someone needs their computer. They can use Memory Machine to take a snapshot of their work, go to lunch, and come back and Memory Machine puts them right where they left off—creatively speaking. It’s not just about opening up the scene, it’s about getting Maya back to where it was. There’s no trying to get back into the flow. You’re right there.”

Wright breaks down how the technology works for Chapeau’s artists. “As an artist, it functions similarly to autosave, by restoring your system to any snapshots you’ve previously taken. It gets rid of a big hurdle in allowing systems to


serve up scenes faster. The theoretical speed of Optane is ridiculous. Now the bottleneck is no longer storage speed.”

Loomam talks about the business benefits. “From a business point of view, we structure our bidding to account for system crashes and other erratic production issues”. If one artist has a hard crash and there’s a network domino effect, all the artists are affected. Suddenly that crash turns into 45 minutes per artist as you bring the network back online. Plus, there’s that human factor of not being able to get back into the flow. When you’re painting frames, for instance, you’re super-focused. After a crash, it’s hard to get back into the flow state.”

Both men say they are also looking at MemVerge’s Memory Machine and Intel Optane persistent memory not just as a way to recover from inevitable application crashes, but as a way to elevate creative work and push the boundaries.

“You can work differently because you have hardware like this,” says Wright. “Adding Intel Optane persistent memory into each artist workstation is not only affordable but it also adds the benefit of an additional insurance policy on delivery deadlines. It creates new ways of thinking about working on scenes and saving files. As MemVerge continues to refine their software to work on Windows* systems, I think you’ll see industry-wide adoption.”





Pushing the boundaries of what's possible

Hank Driskill, a technology and VFX visionary, has this to say about how the MemVerge Memory Machine with Intel® Optane™ PMem can help artists avoid losing their work and their creative flow.

“Scenes can sometimes take forty-five minutes to an hour to load and then, sometimes, saving can take several minutes so artists don't save very often. If a system crashes, then they lose that work and they lose their state of mind and the flow they were in. If they lose work, it can take half a day to recover and that can be catastrophic to the artist and the production timelines. Intel Optane PMem means an artist's work doesn't disappear if their system crashes.

“On every new film, artists are pushing the boundaries of what's possible. Artists gobble up new technology and then raise the bar on the art they're creating. This means that with Intel Optane PMem and MemVerge Memory Machine*, for instance, an artist might be able to create the same piece of art in half the time or a piece of art that is twice as good. I really do think this will be a big deal to the artists.”

It's not magic, it's Intel's persistence

Intel Optane PMem modules simply drop in alongside conventional DDR4 memory, allowing more data to remain close to the CPU with near-DRAM latency.

“People have been talking about persistent memory for 20 years,” says Charles Fan. “When I first heard about it, I thought it was magic, but I was also skeptical because no one had shipped it. Then Intel actually did.

“I think Optane is a revolutionary technology. I expect the product to continue to improve, the price to drop, and for Optane persistent memory to become an increasingly compelling technology. It's a wondrous technology and we are doing our part to make people adopt without having to change apps,” concludes Fan.

[Learn more about Intel Optane persistent memory: intel.com/optane](https://www.intel.com/optane)

[Learn more about MemVerge: memverge.com](https://www.memverge.com)

[Learn more about Chapeau Studios: chapeaustudios.com](https://www.chapeaustudios.com)



¹ <https://www.techradar.com/news/world-of-tech/inside-dreamworks-how-animated-movies-are-rendered-1127122#:~:text=Most%20films%20take%203%2D5,make%20more%20films%20per%20year.>

² Crash tests shows 256 gigabytes of data could be recovered from persistent memory in less than 10 seconds.