



TECHNOLOGY BRIEF

## Increase Availability with Lightning Fast In-Memory Database Crash Recovery

## PROBLEM

# IMDB Crash Recovery is Slow

### “Slow” is The State of Art for IMDB Crash Recovery

Because the general availability of persistent memory occurred only recently, the vast majority of in-memory databases use DRAM as the storage medium. The advantage of DRAM is speedy access times in tens of nanoseconds. A big disadvantage is that all data is lost in the event of a power loss or crash of the application—making the recovery process lengthy—and flat-out unacceptable for mission-critical applications.

Typically, IMDB crash recovery involves copying each operation onto disk or SSD into a transaction log. Using this data, a crashed system can be restored by starting with a database configuration from the beginning of the day, then replaying the transaction log to catch up to the last saved state. For one financial services customer of MemVerge this process takes 3 hours to recover 500GB.

### The Blast Zone is Growing Wider and Deeper

The blast zone for IMDB crashes is growing wider. According to IDC, by 2021, 60-70% of the Global 2000 will have at least one mission-critical real-time workload. The blast zone grows deeper as memory density continues to increase. For example, a single server designed for in-memory databases can be configured with multiple terabytes of DRAM or persistent memory. One well known vendor offers a 2U server configuration with up to 3TB of DDR4-2933MHz, or 6TB of Intel® Optane™ DC Persistent Memory.

As the blast zone grows, so does the need to accelerate IMDB crash recovery.



3 hours to recover 500GB



6TB in 1 server

## SOLUTION

# Big Memory

### Defining Big Memory

Big Memory is a class of computing where the new normal is mission-critical applications and data living in byte-addressable, and much lower cost, persistent memory.

It has all the ingredients needed to handle the growth of IMDB blast zones by accelerating crash recovery. Big Memory can scale-out massively in a cluster and is protected by a new class of memory data services that provide snapshots, replication and lightning fast recovery.

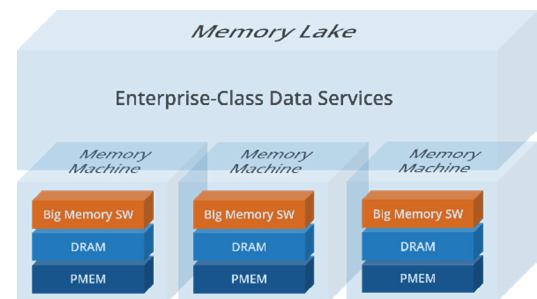
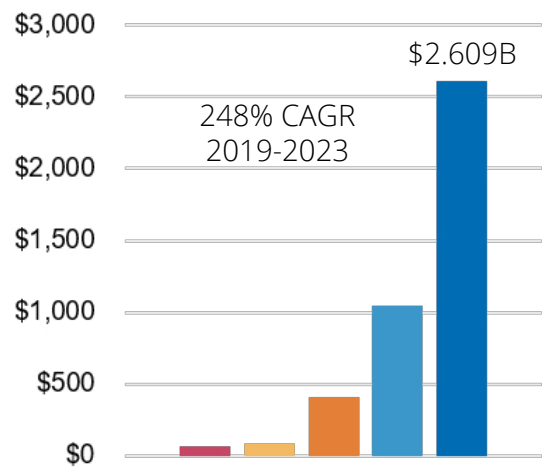
### The Foundation is Intel Optane DC Persistent Memory

The Big Memory market is only possible if lower cost persistent memory is pervasive. To that end, IDC forecasts revenue for persistent memory to grow at an explosive compound annual growth rate of 248% from 2019 to 2023.

### MemVerge Software is the Virtualization Layer

Wide deployment in business-critical tier-1 applications is only possible if a virtualization layer emerges to deliver HPC-class low latency and enterprise-class data protection. To that end, MemVerge pioneered Memory Machine™ software.

Persistent Memory Revenue Forecast 2019 – 2023 - IDC



[Download IDC Presentation Defining Big Memory](#)



## HOW IT WORKS

# Short Pause

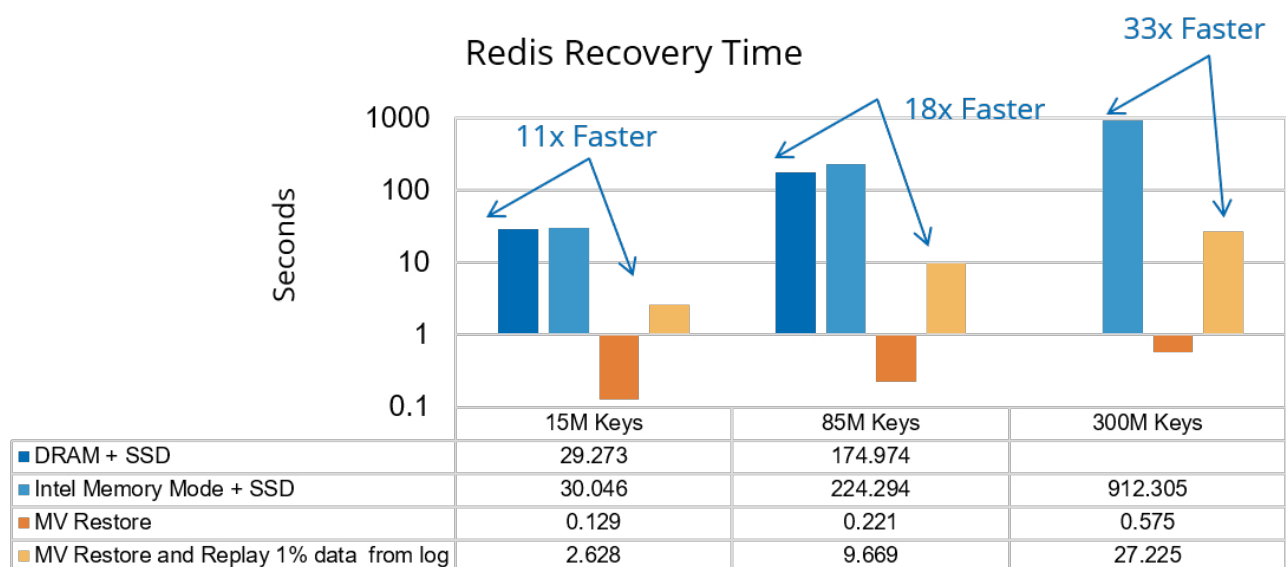
### Memory Machine Software Snapshot, Replication & Fast Recovery

MemVerge Memory Machine software virtualizes DRAM and persistent memory on single server and/or multiple servers in a cluster. It can then take non-disruptive snapshots of the data residing in memory while the application process keeps running. The snapshots can be taken on demand or scheduled at regular intervals. With these snapshots an application can be restored to a previous point in time to rapidly recover from a crash.

The snapshots can be retained on the server or moved to another server non-disruptively. This allows data written by a process on one server to be read by a process on a different server with very low latency. Based on this capability, a timeseries IMDB can be cloned efficiently to another server.

If an IMDB can tolerate a “short pause” of less than a second, Memory Machine can take snapshots as frequently as every minute. When the IMDB crashes, the database can be restored to the latest snapshot and then the transaction log can be replayed from the point in time when that latest snapshot was taken.

For the financial services customer mentioned in our problem statement, the time required to recover 500GB was slashed from 3 hours to 2 seconds. In testing performed by MemVerge, a Redis database can be recovered up to 33 times faster than recovery from storage.



## HOW IT WORKS

### No Pause

#### Memory Machine Software Clones DRAM to PMEM

If the IMDB cannot tolerate a short pause, the combination of snapshots and replication provides another solution.

The primary IMDB is run in DRAM with a clone IMDB running in PMEM. The transactions are replicated from the primary instance to the clone instance. Snapshots are scheduled on the clone instance without any impact to the performance of the primary IMDB.


When the primary IMDB crashes, a new primary instance can be brought up using the latest snapshot of the secondary IMDB. The transaction log can then be replayed to allow the new instance to catch up to the point where the original primary instance crashed. With this method, the crash recovery can happen within minutes or seconds, while the negative performance impact to the primary database instance is eliminated.

#### Summary

As the blast zone for IMDB crashes grows wider and deeper, Intel and MemVerge respond with Big Memory.

With Optane persistent memory and Memory Machine software, memory can safely scale-out to petabytes because it's now possible to recover from crashes in seconds.


The state-of-the-art for IMDB crash-recovery has been changed from "slow" to "fast".



# Increase Availability with Lightning Fast In-Memory Database Crash Recovery



It's Back to School Time



Watch ["The Art of Big and Fast Data"](#)

## Help us shape the future of Big Memory

- Join the Big Memory Community to receive news, tech tips and early access to new software releases.
- Attend Big Memory University and earn a certificate.
- Deploy a Memory Machine PoC to see the power of Big Memory for yourself.

Send me info