



PERFORMANCE COMPARATIVE ANALYSIS OF MEMVERGE™ MEMORY MACHINE™ WITH REDIS WORKLOAD

Comparing server performance and cost

CONTENTS

Executive summary.....	2
Key takeaways.....	2
Configurations.....	3
Results.....	3
85M Keys Redis I/O Performance.....	3
300M Keys Redis Persistence (Snapshots to Storage and Memory).....	4
300M Keys Redis Recovery.....	5
HPE ProLiant DL380 Gen10.....	6
Bottom Line.....	6

EXECUTIVE SUMMARY

According to IDC, the growth of “real-time” data is leading the explosive growth of data and will comprise 25% of all data by 2024. The growth of real-time data is why applications incorporating AI/ML and real-time analytics are becoming the new normal. To meet real-time performance requirements, these apps process as much data as possible in-memory. The result is the cost, performance, and availability of volatile memory infrastructure is under stress.

In response, Big Memory Computing was invented to make memory abundant, persistent, and highly available. Big Memory Computing consists of DRAM and Intel® Optane™ persistent memory 100 series for HPE used with Big Memory software such as Memory Machine™ from MemVerge™ to virtualize the pool of memory and create a platform for data services to protect the data in-memory. Redis is an in-memory database but its data still needs to be saved or persisted and later restored from storage. This task becomes increasingly important and time-consuming as we work with larger data sets. This paper describes what we believe is the most efficient method of saving in-memory data.

With Memory Machine on Redis workloads, applications can allocate their in-memory data structures and runtime states transparently from a large pool of memory. DRAM and Intel Optane PMem 100 series for HPE are combined flexibly to provide both large capacity and DRAM-like memory access speed. With memory being managed by Memory Machine, applications can further leverage multiple enterprise in-memory data services to gain better performance and higher reliability. For instance, Memory Machine’s ZeroIO™ Snapshot allows instant capture of an application’s running state, which in turn enables instant application restart and recovery.

This paper informs and educates the sales field about the performance advantage of industry-standard servers using MemVerge Memory Machine software with a memory pool of DRAM and Intel Optane PMem 100 series for HPE. The workload results are from HPE internal lab testing. Using Intel Optane PMem 100 series for HPE with MemVerge on an HPE server, customers get added features and more memory at a lower cost per gigabyte of memory without sacrificing performance.

In a recent performance benchmark, HPE found significant performance advantages using Memory Machine software with Intel Optane PMem 100 series for HPE on Redis workloads.

Key takeaways

- **Memory Machine software and Intel Optane PMem 100 series for HPE can lower total TCO against a comparable DRAM configuration.**
- **Using MemVerge Memory Machine and Intel Optane PMem 100 series for HPE on HPE servers provides enterprise-class data protection for in-memory workloads:**
 - **ZeroIO™ snapshots are the industry’s first memory snapshots to persistent memory. Now, hundreds of gigabytes can be snapshot in 1-2 seconds, much faster than snapshots to storage, which are disruptive because they can take several minutes. Memory Machine can save data to persistent memory without cache up to 693X faster than saving DRAM to storage.**
 - **Auto-save takes advantage of non-disruptive ZeroIO snapshots. IT organizations are now encouraged to schedule frequent snapshots because they are non-disruptive.**
- **Memory Machine can restore data in seconds as the workload is already in persistent memory. Recovery using MemVerge Memory Machine is at least 8X faster.**

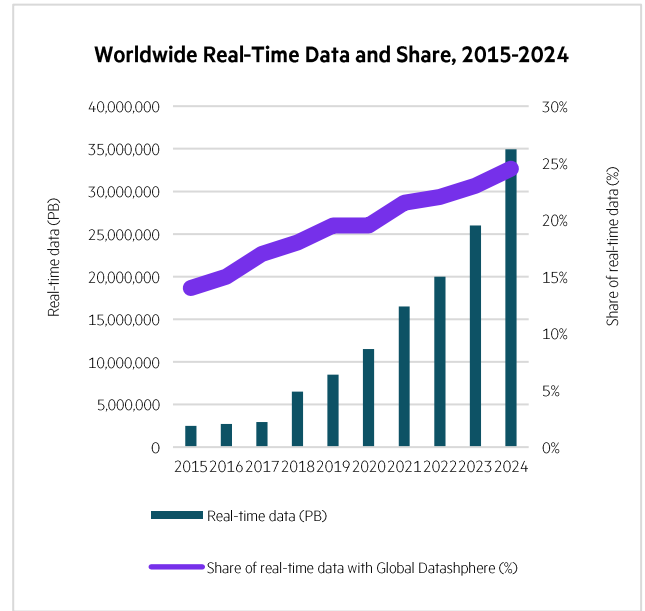


FIGURE 1. The growth of real-time data share.



Configurations

For this setup, an HPE ProLiant DL380 Gen10 server was configured with two processors and MemVerge Memory Machine using Redis Persistence workloads. Redis provides a different range of persistence options; for these benchmarks, Redis Persistence executed point-in-time snapshots of the database in three different intervals. Memory Machine DRAM/PMEM ratio is configurable. Three different sizes were used to showcase the different results of no DRAM to a larger amount of DRAM in the memory pool:

- MemVerge Memory Machine without cache
- MemVerge Memory Machine with 60GB of DRAM
- MemVerge Memory Machine with 120GB of DRAM

The baseline was Intel Memory Mode RDG enabled which captured a point-in-time snapshot (ms) of 20+ minutes.

RESULTS

85M Keys Redis I/O Performance

Application performance

MemVerge Memory Machine combines PMem + DRAM to create a large memory pool with DRAM-like performance and provides storage data services at a lower cost. The graphs below compare latency and operations per second of running large memory-intensive workloads on a DRAM-only system to MemVerge Memory Machine systems combining different ratios of (DRAM+PMem) in standard get/set performance tests.

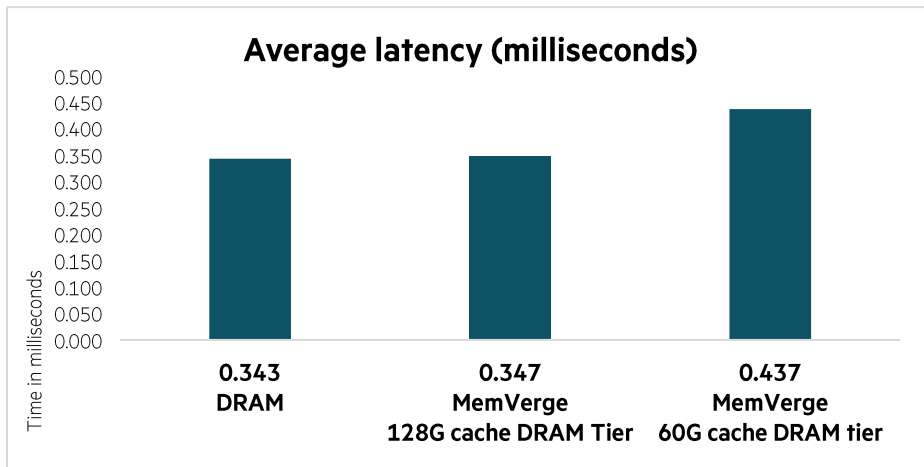


FIGURE 2. Comparing operations per second of large memory-intensive workloads on a DRAM-only system to MemVerge Memory Machine.

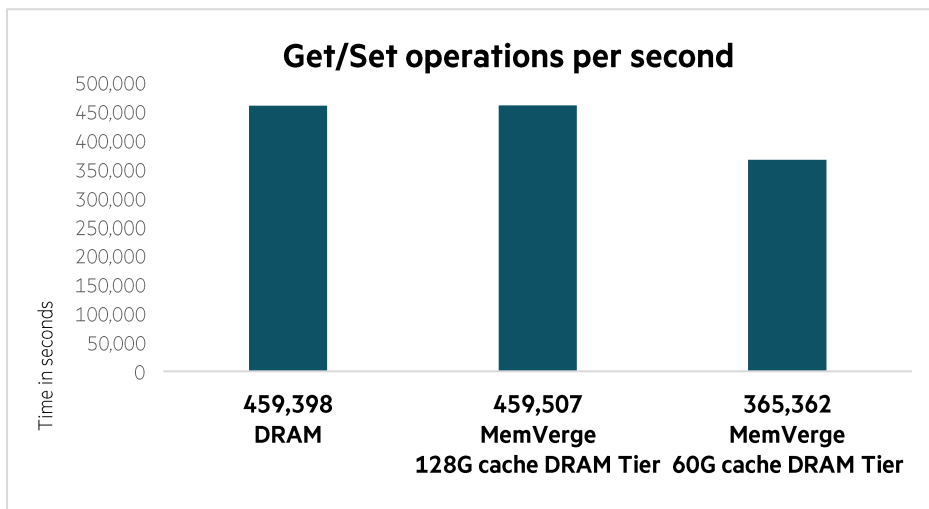


FIGURE 3. Comparing the latency of large memory-intensive workloads on a DRAM-only system to MemVerge Memory Machine.



TABLE 2. 85M Keys Redis I/O performance with average latency and average ops.

Test type	Set P99	Get P99	Average latency	Average OPS
DRAM	0.55	0.55	0.343	459,398
Intel Memory Mode (192 GB cache)	0.56	0.56	0.344	458,085
MemVerge DRAM cache 60 GB	0.64	0.64	0.437	365,362
MemVerge DRAM cache 1200GB DRAM	0.56	0.56	0.347	459,507

300M Keys Redis Persistence (Snapshots to Storage and Memory)

320 GB data set

Redis “RDB Persistence” performs point-in-time snapshots of datasets at specified intervals. By default, Redis saves snapshots of the dataset on disk or SSD in a binary file called dump. Redis RDB can be configured to save the dataset every N seconds if there are at least M changes in the dataset, or you can manually call the SAVE or BGSAVE commands.

Testing by HPE shows the advantage of using Memory Machine to perform snapshots to memory instead of to SSD. When doing snapshots to PMem without DRAM Cache, Memory Machine disrupts the application for only 1.79 seconds. This is a big improvement over native Redis snapshots to SSD which disrupt the application for over 20 minutes.

The ability to do non-disruptive snapshots makes it possible for IT organizations to schedule frequent snapshots and dramatically reduce RPO.

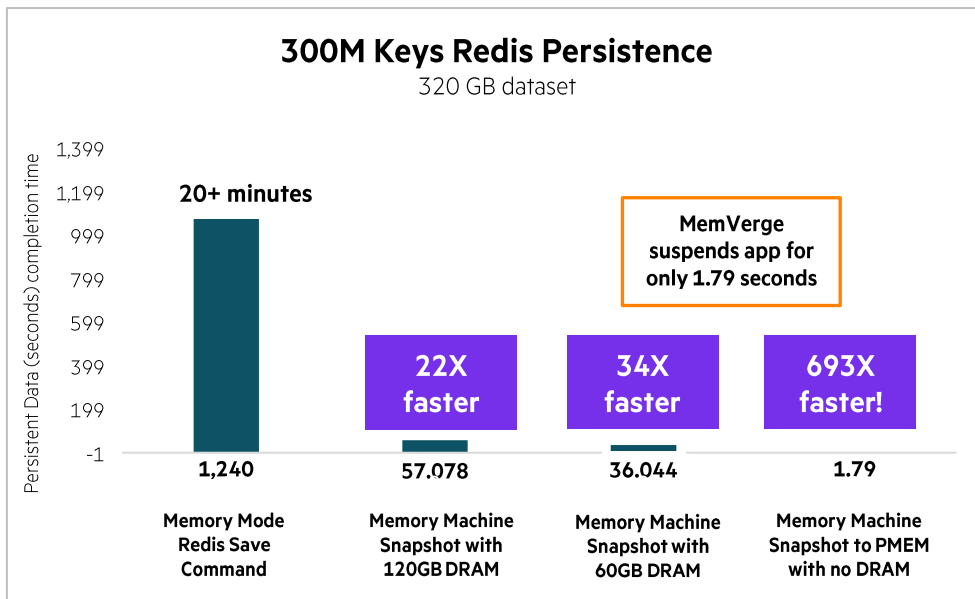


FIGURE 4. Results of 300M Keys Redis Persistence workloads.

TABLE 1. Comparison of 300M Keys Redis Persistence (320 GB database) with different MemVerge configurations

MM Snapshot	Snapshot total time (ms)
MemVerge without cache	1,790 (1.79 sec)
MemVerge with 60GB DRAM	36,044 (36 sec)
MemVerge with 120GB DRAM	57,078 (57 sec)



300M Keys Redis Recovery

In the past, restoring from snapshots meant restoring from a daily log because saving the data was very infrequent. Saving was infrequent because of how long and disruptive snapshots were. With Memory Machine and Intel Optane PMem 100 series for HPE, frequent snapshots are encouraged, which in turn reduces the replay (recovery) time. Best case, restoring a Redis database with 300M keys with 0% replay takes only 0.5 seconds, or 1,500X faster than restoring from storage. Restoring the same 300M keys with 5% replay is 8X faster with Memory Machine.

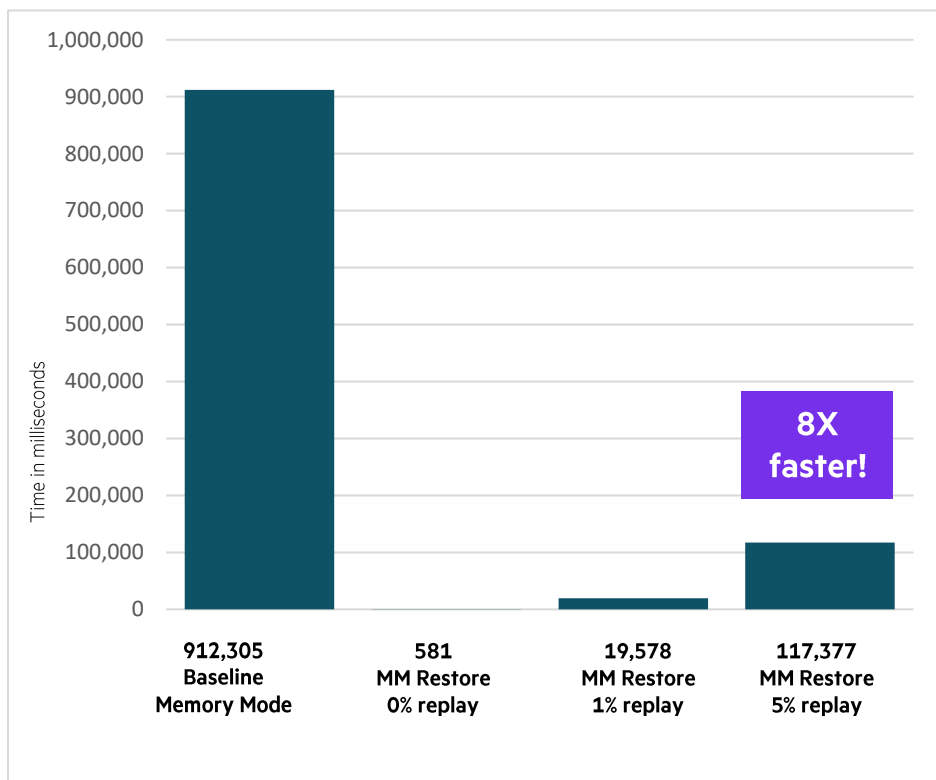


FIGURE 4. Results of 300M Keys Redis Recovery workloads.

TABLE 3. Comparison of 300M Keys Redis Recovery with different Memory Machine configurations.

Performance details	Restore time
DB size	300M Keys
DRAM	Wasn't possible. Exceeded DRAM capacity
Intel memory mode	912,305 (15 min)
MemVerge Restore 0% data replay	581 (.5 sec)
MemVerge Restore 1% data replay	19,758 (19.7 sec)
MemVerge Restore 5% data replay	117,377 (1.95 min)

NOTE

ZeroIO snapshots are extremely fast. If you take 20 times as many snapshots, you will replay 1/20 of log transactions (5%).



HPE PROLIANT DL380 GEN10

The HPE ProLiant DL380 Gen10 server delivers the latest in security, performance and expandability, backed by a comprehensive warranty. Standardize on the industry's most trusted compute platform. The HPE ProLiant DL380 Gen10 server is securely designed to reduce costs and complexity, featuring the First and Second Generation Intel® Xeon® Processor Scalable Family with up to a 60% performance gain¹ and 27% increase in cores², plus the HPE 2933 MT/s DDR4 SmartMemory supporting 3.0 TB. It supports 12 Gb/s SAS, and up to 20 NVMe drive plus a broad range of compute options. HPE Persistent Memory offers unprecedented levels of performance for databases and analytic workloads. Run everything from the most basic to mission-critical applications and deploy with confidence.



BOTTOM LINE

The enterprise application world is being transformed into a *real-time* application world dependent on lower-cost, high-performance, and highly available Big Memory. The bottom line of this report is that HPE ProLiant DL380 Gen10 servers equipped with Intel Optane PMem 100 series for HPE and Memory Machine software bring for the first time, enterprise-class data protection to Redis in-memory databases. With MemVerge, customers get added features and more memory at a lower cost without sacrificing performance.

LEARN MORE AT

[HPE ProLiant DL380 Gen10](#)

[Intel Optane PMem 100 series for HPE](#)

[MemVerge Memory Machine](#)

© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

MemVerge and Memory Machine are registered trademarks of MemVerge, Inc. Intel, the Intel logo, Xeon, and Optane are trademarks of Intel Corporation or its subsidiaries. All third-party marks are property of their respective owners.