

Scale Redis™ Memory Configurations for Larger Data Stores and Caching



Ultrastar DC ME200 Memory Extension Drive, NVMe™ 2.5-inch U.2 and AIC HH-HL form factors

TCO Highlights

- Better management of in-memory data stores through reduced server count
- Improve memory density
- Overcome limitations of DIMM slots per server for larger caching and scalable data stores
- Reduce overall data center footprint and power consumption
- Increase the number of instances per node within a Redis cluster for better CPU utilization

Product Features

- Up to 24TiB software-defined memory per 2U server (U.2)
- 99.9% latency < 1ms
- Works transparently with existing Linux® distributions and application stacks

Ultrastar® DC ME200 Memory Extension Drive with Redis Enables Scalable In-Memory Caching and Data Stores for a Better TCO

Redis in-memory data stores and caching engines improve application performance by storing or caching frequently accessed data items in main memory for faster data retrieval.

In order to achieve the highest performance, the entire dataset must be stored in-memory. Also, to future proof a growing dataset, or if the data being handled is larger than the available memory in a single server, Redis engines allow for scale-out configurations to multiple nodes using sharding.

The Redis scale-out approach requires careful balancing according to the amount of memory and processing power available on a per-node basis within the cluster.

Scalability and Need For High CPU: Memory Ratio

Redis can quickly become bottlenecked on memory if the message size is sufficiently small and CPU utilization is low. For example, modern server architectures are limited to 12 DIMMs per socket, and in the case of high-density servers, only 6-8 DIMMs per socket. Meanwhile, the processor core counts continue to increase with every generation, reducing the overall memory-to-core ratio even further.

As a result, the number of nodes for a large-scale cluster is determined by the amount of memory in each node, and not by the compute capacity per node. This is evident in typical deployments seeing CPU utilization on the nodes in the range of 10%-20% (for most common scenarios).

The result of introducing additional servers to house extra memory is unnecessarily high CAPEX and OPEX costs for the infrastructure, i.e., more hardware, more processors, more networking, more cooling, and a larger data center footprint.

Ultrastar DC ME200 Memory Extension Drive Overview

Ultrastar memory drive combines one or more custom NVMe drives, tuned for performance, with a software layer that expands system DRAM onto them. Unmodified Linux operating systems using this technology can address system memory up to eight times the capacity of the DRAM installed in a server with near-DRAM speeds. Memory-intensive Redis caching and memory stores can utilize this extra system RAM without any changes. For example, a 1U server with 256 GiB installed can make use of up to 2 TiB of Ultrastar memory.

Ultrastar Memory Drive Benchmarked Performance

The Redis benchmark uses high concurrency SET/GET operations of small (1kB) and large (100kB) messages, where the Redis server and client-load systems are connected over a 10GbE network connection.

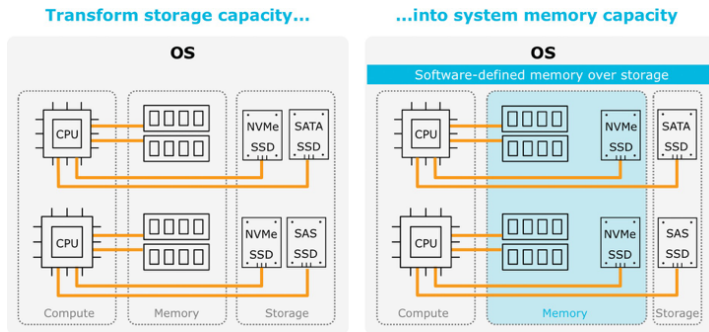

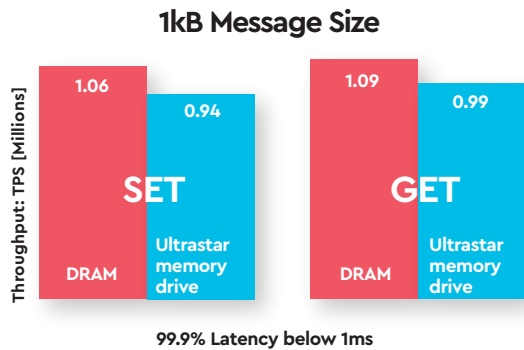


Figure 1 - The left diagram depicts the traditional compute-memory-storage architecture. The right diagram depicts how database instance configurations can take advantage of Ultrastar memory drives to augment server DRAM to create a virtualized memory pool to enable greater memory expansion.

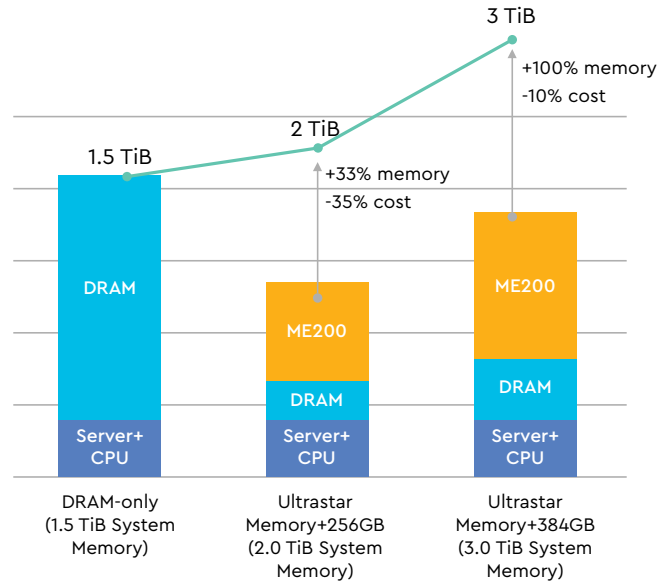
Ultrastar memory drives allow for multiple terabytes of main memory per server. However, since it is impossible to set up a comparable 6TB DRAM-only server, an apples-to-apples comparison was performed using a DRAM-only server of 768GB DDR4 vs. a server with 192GB DDR4 plus Ultrastar memory for a total of 768GB.

The graphs below indicate Ultrastar memory drive performance of Redis delivered at 86%-94% of DRAM-only performance with a dataset 4x larger than DRAM.



Cost Savings

The chart below compares two Ultrastar memory configurations to a baseline configuration of the same dual socket system with 1.5TB of DDR4.



- The first Ultrastar memory configuration shows that reducing DRAM from 1.5TiB to 256GiB and adding 2 Ultrastar memory devices reduces total system cost by 35% and provides applications with **33% more memory**.
- The second Ultrastar memory configuration shows that reducing DRAM from 1.5TiB to 384GiB and adding 2 Ultrastar memory devices, applications benefit from **100% more memory per node** while reducing the system cost by 10%.

The ability to have a server configuration with twice as much memory for a similar cost per server enables users to **reduce the number of servers in a Redis cluster by up to 50%**, while supporting the desired data-set size.

This has direct impact on CAPEX for the entire infrastructure, i.e., servers, networking, that is at least 50%, as well as data-center footprint and power consumption for further OPEX savings.

¹Memory capacity is indicated by GiB and TiB and based on binary values such that one gibibyte (GiB) is equal to 2³⁰ bytes and one tebibyte (TiB) is equal to 1024 GiB (2⁴⁰) bytes. Suggested expansion of 8x DRAM based on internal performance testing across a variety of industry standard benchmarks.

Western Digital.

5601 Great Oaks Parkway
San Jose, CA 95119, USA
US (Toll-Free): 800.801.4618
International: 408.717.6000

www.westerndigital.com

©2019 Western Digital Corporation or its affiliates. Produced 01/19, Rev. 03/19. All rights reserved. Western Digital, the Western Digital logo, and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Redis is a trademark of Redis Labs Ltd. Any rights therein are reserved to Redis Labs Ltd. Any use by Western Digital is for referential purposes only and does not indicate any sponsorship, endorsement or affiliation between Redis and Western Digital.

All other marks are the property of their respective owners. References to any Western Digital product does not mean it will be available in all countries.