



Can DDR5 server memory deliver as promised?

Micron shares our real-world test results.

So, is DDR5 really what they say? There's been keen anticipation for the arrival of DDR5 server DRAM and the big performance gains it has promised.

The speed!

The doubled densities!

The vastly greater memory bandwidth!

Data center workloads have hit the limits of transferring data between devices and global memory, while bottlenecks slowed processing of today's high-volume workloads when fetching data from memory.

The question:


Do customers get breakthrough performance and impressive time to results with 4th Gen AMD EPYC™ Processors¹?

The answer:

Micron, one of the three global DRAM manufacturers, can weigh in here:

Yes,

because our tests performed in the real world² on DDR5 server DRAM show strong performance.



Data center memory made for the data center

Data centers today must handle more information at faster speeds and with greater efficiency. Industries around the globe are leveraging data for insight and innovation. To meet these demands, IT managers have been leveraging platforms that pack more processing cores into CPUs. Yet, counter to what they needed, the memory bandwidth per CPU core decreased. To address the shortfalls with DDR4, JEDEC announced the DDR5 spec in 2020 – the first DRAM memory spec engineered with a data-center-first focus. With the first platforms officially arrived from AMD and additional platforms on the horizon, data centers can now take advantage of the performance improvements this next-gen memory delivers.

Micron is a collaborator with JEDEC on DDR5 specs. Our enthusiastic YES is informed from the analysis we ran in partnership with AMD engineers on their testbed of processors supporting high-performant DDR5 DIMMs and fast PCIe® Gen 5 I/O. DDR5 makes processing data faster and more efficient. **The results that follow are actual performance numbers from tests of Micron® DDR5 Server DRAM and AMD EPYC™ 9004 Series Processors.**



Blazing fast DDR5 results on real DDR5 servers

This validation was performed by Micron's Data Center Workload Engineering (DCWE) team where tests compared the performance of Micron DDR4 versus DDR5 server memory running on the appropriate generations of AMD EPYC servers³, to isolate (to the extent possible) the DDR5 DRAM contribution to the workloads.

Micron's team chose representative workload groups that model typical performance challenges for high performance computing (HPC) applications: The STREAM Microbenchmark and three actual HPC scientific workloads. These results from real-world workloads indicated that with DDR5 memory, individual **HPC applications we ran performed an average of 2x faster**. Also, DDR5 enabled better parallelism on the compute node for these applications by progressively delivering more gains at increased core counts. Results like this in Figure 1 suggest that Micron[®] DDR5 Server DRAM and 9654 Series 4th Generation AMD EPYC[™] Processors are a potent combination for HPC.




Figure 1: Micron DDR5 doubles memory bandwidth with 4th-Gen AMD EPYC[™] processors using STREAM

1. Single socket 3rd Gen AMD EPYC CPU 7763 with Micron DDR4 3200 MHz system is capable of 189 GB/sec

2. Single socket 4th Gen AMD EPYC CPU 9654 with Micron DDR5 4800 MHz system is capable of 378 GB/sec

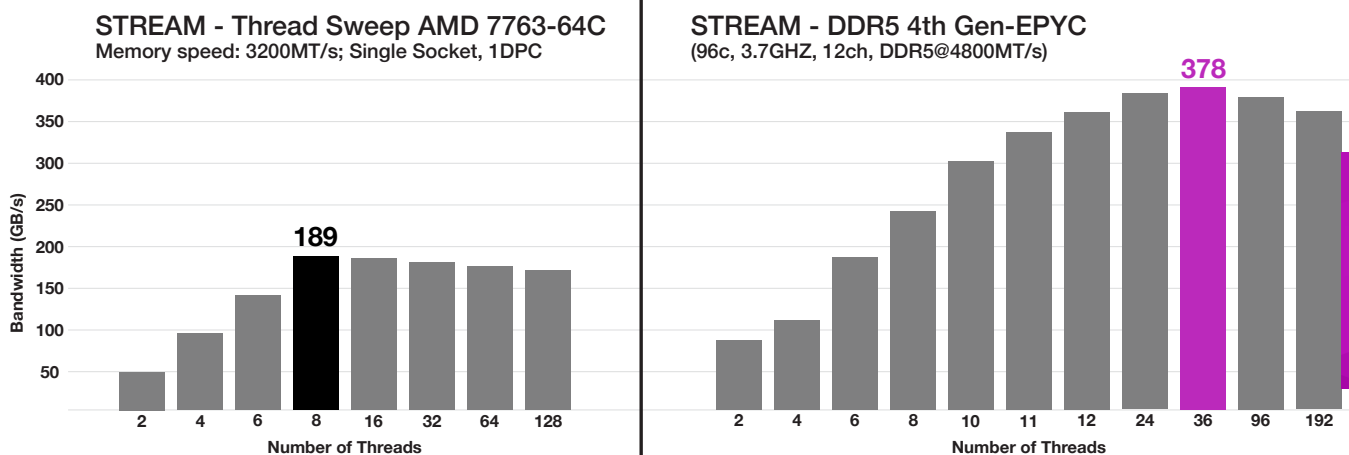
STREAM Memory Subsystem Microbenchmark

Often used to gauge data bottlenecks, STREAM is a simple, well-known benchmark that captures peak memory bandwidth for HPC systems. This noncomplex synthetic program tests simple vector kernels to measure sustainable memory bandwidth (in MB/s) versus the computation rate.

Our results:

We showed **double the memory bandwidth** for a single socket DDR5 system with STREAM (Figure 2). Micron used the Alma 9 Linux Kernel 5.14 software stack and a STREAM array size of 2.5 billion for a footprint of 56GB. The BIOS settings had cache prefetchers enabled. The Micron DDR5-AMD single-socket system improved memory bandwidth from a peak of 189GB/s with the Micron DDR4-AMD system to 378GB/s, which is 2x more, due mostly to gains from a 50% increase in core count and memory channels³.

Figure 2: Memory bandwidth as core count increased in STREAM; DDR4 DRAM hit a peak of 189GB/s while DDR5 DRAM hit a peak of 378GB/s



Why it's useful:

Some high-end CPUs are so much faster than the computer memory systems that feed them, they may spend up to 95% of their rated peak speeds idle and waiting for cache to be satisfied. Testing on datasets bigger than the available cache on HPC systems, as this benchmark does, is a useful indicator of performance with very large, contiguous, long-vector memory-access applications. Examples include AI training and inference, data analytics, and HPC, including hyperscale cloud applications.

2x
more bandwidth

High performance computing workloads

Micron also benchmarked execution time improvements (a speed test) of three specific scientific computing applications. These are often large, data-intensive workloads that must be split into millions of operations that run in parallel and use terabytes of data. The three scientific computing benchmarks that Micron chose were:

1. Computational fluid dynamics with the OpenFOAM application used “Motorbike 240M cells” dataset
2. Weather Research and Forecasting (WRF) is a numerical weather prediction system, and we used the “CONUS 2.5km” dataset
3. CP2K is a quantum chemistry and solid-state physics software package for molecular dynamics that can perform atomistic simulations; we used the “H2O-dft-ls.NREP4” dataset

Our results:

Micron compared DDR5 vs. DDR4 on isolation tests that matched the number of cores used, channels per NUMA node and CPU frequency. Performance with DDR5 was up to 2x faster for the three HPC applications. For the computational fluid dynamics benchmark (OpenFOAM), using a data set for a motorbike turbulence simulation to calculate steady air flow around a motorcycle and rider, tests showed that it **runs 140% faster with DDR5** (Figure 3). Using the Continental United States (CONUS) at 2.5 km lateral resolution dataset with OpenMPI 4.1.1 to test WRF showed that it **runs 69% faster** with DDR5.




Figure 3: The OpenFOAM computational fluid dynamics benchmark with DDR5 vs DDR4 shows DDR5 is 140% faster

Why it's useful:

These *are* real-world applications. HPC workloads are notably memory-bandwidth hungry, and the lack of sufficient memory bandwidth often constrains their performance. These complex workloads are focused on solving some of mankind's most challenging problems, including weather and climate simulations, seismic modelling, chemical, physics and biological analysis.


These workloads also highlight the types of simulations, predictions and models of complex systems with large datasets that are often grouped with AI and other data analytics to support complex workflow analysis for both business and science, which help reduce development time and costs in a broad range of applications.

Now that DDR5 server platforms are here, are you ready?

DDR5 server DRAM is such a generational jump in memory that, of course, the testing and validation of products will be more intricate than that for DDR4. Micron helps you leverage DDR5 server memory and crush even the most challenging workloads. So, unleash your most data-intensive applications.

Micron led the way in enabling the ecosystem transition to DDR5 with the industry's only [DDR5 Technology Enablement Program](#). Micron® DDR5 Server DRAM launched at speeds of 4800MT/s to immediately deliver an 85% or higher increase in performance over DDR4⁴. Reach out to MicronCPG if you'd like to learn further testing details. Check back as more DDR5 server DRAM workload benchmarks are on the horizon once additional platforms have been announced.

For more info: microncpg.com/serverDDR5





SOURCES

1. AMD, the AMD arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc.
2. For the purpose of this article, "real-world results" means that all benchmark tests cited were performed on physical servers equipped with Micron® DDR5 Server DRAM and AMD EPYC™ 9004 Series Processors, as opposed to synthetic benchmarks or theoretical mathematical extrapolations.
3. STREAM Memory Subsystem Microbenchmark, OpenFOAM, WRF, and CP2K benchmark tests were performed from August through November 2022 by Micron engineering teams on DDR4-capable 64-core EPYC 7763 systems with DDR4-3200 MHz fully populated with 64 GB RDIMMs and DDR5-capable systems with AMD 96 core EPYC™ 9654 processors with DDR5-4800 MHz fully populated with 64 GB RDIMMs.
4. DDR5 launch data rate of 4800MT/s transfers 2x the data of the maximum standard DDR4 data rate of 3200MT/s, based on actual testing using 4th-Generation AMD EPYC processors. JEDEC's projected speeds of 8800MT/s for DDR5 will be 2.75x faster.