




## Micron Delivers Industry's Highest Capacity SOCAMM2 for Low-Power DRAM in the AI Data Center

October 22, 2025 at 9:02 AM EDT

### 192GB SOCAMM2, built with LPDDR5X, extends Micron's leadership in power-efficient solutions for AI infrastructure

BOISE, Idaho, Oct. 22, 2025 (GLOBE NEWSWIRE) — In an era of unprecedented AI innovation and growth, the entire data center ecosystem is transforming toward more energy-efficient infrastructure to support sustainable growth. With memory playing an increasingly critical role in AI systems, low-power memory solutions have become central to this transformation. Micron Technology, Inc. (Nasdaq: MU) today announced customer sampling of 192GB SOCAMM2 (small outline compression attached memory modules) to enable broader adoption of low-power memory within AI data centers. SOCAMM2 extends the capabilities of Micron's [first-to-market LPDRAM SOCAMM](#), delivering 50% more capacity in the same compact footprint. The added capacity can significantly reduce time to first token (TTFT) by more than 80% in real-time inference workloads.<sup>1</sup> The 192GB SOCAMM2 uses Micron's most advanced 1-gamma DRAM process technology to deliver greater than 20% improvement in power efficiency,<sup>2</sup> further enabling power design optimization of large data center clusters. These savings become quite significant in full-rack AI installations, which can include greater than 50 terabytes of CPU-attached low-power DRAM main memory.<sup>3</sup> The modular design of SOCAMM2 improves serviceability and lays the groundwork for future capacity expansion.



Micron has begun customer sampling the 192GB SOCAMM2 to enable broader adoption of low-power memory within AI data centers. The 192GB SOCAMM2 uses Micron's most advanced 1-gamma DRAM process technology to deliver greater than 20% improvement in power efficiency, further enabling power design optimization of large data center clusters.

[A Media Snippet accompanying this announcement is available by clicking on this link.](#)

Building on a five-year collaboration with NVIDIA, Micron pioneered the use of low-power server memory in the data center. SOCAMM2 delivers LPDDR5X's inherent advantages — exceptionally low power consumption and high bandwidth — to the main memory of AI systems. Designed to meet the evolving demands of massive-context AI platforms, SOCAMM2 provides the high data throughput required for AI workloads while delivering new levels of energy efficiency, setting a new standard for AI training and inference systems. This combination of advantages will make SOCAMM2 a key memory solution for leading-edge AI platforms in the years ahead.

"As AI workloads become more complex and demanding, data center servers must achieve increased efficiency, delivering more tokens for every watt of power," said Raj Narasimhan, senior vice president and general manager of Micron's Cloud Memory Business Unit. "Micron's proven leadership in low-power DRAM ensures our SOCAMM2 modules provide the data throughput, energy efficiency, capacity and data center-class quality essential to powering the next generation of AI data center servers."

Through specialized design features and enhanced testing, Micron SOCAMM2 products transform low-power DRAM, initially designed for mobile phones, into data center-class solutions. Extensive experience in high-quality data center DDR memory helps Micron ensure that SOCAMM2 meets the stringent quality and reliability requirements of our data center customers.

SOCAMM2 improves power efficiency by more than two-thirds<sup>4</sup> compared with equivalent RDIMMs, while packing its performance into a module one-third the size<sup>5</sup>, optimizing data center footprint and maximizing capacity and bandwidth. SOCAMM's modular design and innovative stacking technology improves serviceability and aids the design of liquid-cooled servers.

Micron has been an active participant in the JEDEC SOCAMM2 specification definition and is working closely with industry partners to drive standards that will accelerate low-power adoption in AI data centers to improve power efficiency across the entire industry. SOCAMM2 customer samples are shipping now in capacities up to 192GB per module and speeds up to 9.6 Gbps, with high-volume production aligned to customer launch schedules.

#### Additional resources:

- [SOCAMM2 Technology Enablement Program webpage](#)
- [SOCAMM2 image gallery](#)

#### About Micron Technology, Inc.

Micron Technology, Inc., is an industry leader in innovative memory and storage solutions, transforming how the world uses information to enrich life *for all*. With a relentless focus on our customers, technology leadership and manufacturing and operational excellence, Micron delivers a rich portfolio of high-performance DRAM, NAND and NOR memory and storage products through our Micron® and Crucial® brands. Every day, the innovations that our people create fuel the data economy, enabling advances in artificial intelligence (AI) and compute-intensive applications that unleash opportunities — from the data center to the intelligent edge and across the client and mobile user experience. To learn more about Micron Technology, Inc. (Nasdaq: MU), visit [micron.com](#).

© 2025 Micron Technology, Inc. All rights reserved. Information, products, and/or specifications are subject to change without notice. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.

#### Micron Product and Technology Communications Contact:

Mengxi Liu Evensen  
+1 (408) 444-2276  
[productandtechnology@micron.com](mailto:productandtechnology@micron.com)

#### Micron Investor Relations Contact

Satya Kumar

<sup>1</sup> Performance improvement validated by Micron internal testing: Llama 3 70B model inference with OSL=128 on GH200 NVL2 (288GB HBM3E + 1TB LPDDR5x) using LMCACHE.

<sup>2</sup> Compared to Micron's previous generation LPDDR5X.

<sup>3</sup> Figure based on announced capacity of NVL144 rack systems.

<sup>4</sup> Calculated based on power used in watts by one 128GB, 128-bit bus width SOCAMM2 module compared to two 128GB, 128-bit bus width DDR5 RDIMMs.

<sup>5</sup> Calculation compares SOCAMM2 area (14 x 90mm) versus a standard server RDIMM.