



Micron Innovates From the Data Center to the Edge With NVIDIA

March 18, 2025 at 4:24 PM EDT

Micron HBM3E 12H and LPDDR5X-based SOCAMM solutions designed to unlock full potential of AI platforms

SAN JOSE, Calif., March 18, 2025 (GLOBE NEWSWIRE) -- **GTC 2025** -- Secular growth of AI is built on the foundation of high-performance, high-bandwidth memory solutions. These high-performing memory solutions are critical to unlock the capabilities of GPUs and processors. Micron Technology, Inc. (Nasdaq: MU), today announced it is the world's first and only memory company shipping both HBM3E and SOCAMM (small outline compression attached memory module) products for AI servers in the data center. This extends Micron's industry leadership in designing and delivering low-power DDR (LPDDR) for data center applications.

Micron's SOCAMM, a modular LPDDR5X memory solution, was developed in collaboration with NVIDIA to support the [NVIDIA GB300 Grace™ Blackwell Ultra Superchip](#). The Micron HBM3E 12H 36GB is also designed into the NVIDIA HGX™ B300 NVL16 and GB300 NVL72 platforms, while the HBM3E 8H 24GB is available for the NVIDIA HGX B200 and [GB200 NVL72](#) platforms. The deployment of Micron HBM3E products in [NVIDIA Hopper](#) and [NVIDIA Blackwell](#) systems underscores Micron's critical role in accelerating AI workloads.

Think AI, think memory, think Micron

At GTC 2025, Micron will showcase its complete AI memory and storage portfolio to fuel AI from the data center to the edge, highlighting the deep alignment between Micron and its ecosystem partners. Micron's broad portfolio includes HBM3E 8H 24GB and HBM3E 12H 36GB, LPDDR5X SOCAMMs, GDDR7 and high-capacity DDR5 RDIMMs and MRDIMMs. Additionally, Micron offers an industry-leading portfolio of data center SSDs and automotive and industrial products such as UFS4.1, NVMe® SSDs and LPDDR5X, all of which are suited for edge compute applications.

"AI is driving a paradigm shift in computing, and memory is at the heart of this evolution. Micron's contributions to the NVIDIA Grace Blackwell platform yields significant performance and power-saving benefits for AI training and inference applications," said Raj Narasimhan, senior vice president and general manager of Micron's Compute and Networking Business Unit. "HBM and LP memory solutions help unlock improved computational capabilities for GPUs."

SOCAMM: a new standard for AI memory performance and efficiency

Micron's SOCAMM solution is now in volume production. The modular SOCAMM solution enables accelerated data processing, superior performance, unmatched power efficiency and enhanced serviceability to provide high-capacity memory for increasing AI workload requirements.

Micron SOCAMM is the world's fastest, smallest, lowest-power and highest capacity modular memory solution,¹ designed to meet the demands of AI servers and data-intensive applications. This new SOCAMM solution enables data centers to get the same compute capacity with better bandwidth, improved power consumption and scaling capabilities to provide infrastructure flexibility.

- **Fastest:** SOCAMMs provide over 2.5 times higher bandwidth at the same capacity when compared to RDIMMs, allowing faster access to larger training datasets and more complex models, as well as increasing throughput for inference workloads.²
- **Smallest:** At 14x90mm, the innovative SOCAMM form factor occupies one-third of the size of the industry-standard RDIMM form factor, enabling compact, efficient server design.³
- **Lowest power:** Leveraging LPDDR5X memory, SOCAMM products consume one-third the power compared to standard DDR5 RDIMMs, inflecting the power performance curve in AI architectures.⁴
- **Highest capacity:** SOCAMM solutions use four placements of 16-die stacks of LPDDR5X memory to enable a 128GB memory module, offering the highest capacity LPDDR5X memory solution, which is essential for advancements towards faster AI model training and increased concurrent users for inference workloads.
- **Optimized scalability and serviceability:** SOCAMM's modular design and innovative stacking technology improve serviceability and aid the design of liquid-cooled servers. The enhanced error correction feature in Micron's LPDDR5X with data center-focused test flows, provides an optimized memory solution designed for the data center.

Industry-leading HBM solutions

Micron continues its competitive lead in the AI industry by offering 50% increased capacity over the HBM3E 8H 24GB within the same cube form factor.⁵ Additionally, the HBM3E12H 36GB provides up to 20% lower power consumption compared to the competition's HBM3E 8H 24GB offering, while providing 50% higher HBM memory capacity.⁶

By continuing to deliver exceptional power and performance metrics, Micron aims to maintain its technology momentum as a leading AI memory solutions provider through the launch of HBM4. Micron's HBM4 solution is expected to boost performance by over 50% compared to HBM3E.⁷

Complete memory and storage solutions designed for AI from the data center to the edge

Micron also has a proven portfolio of storage products designed to meet the growing demands of AI workloads. Advancing storage technology in performance and power efficiency at the speed of light requires tight collaboration with ecosystem partners to ensure interoperability and a seamless customer experience. Micron delivers optimized SSDs for AI workloads such as: inference, training, data preparation, analytics and data lakes. Micron will be showcasing the following storage solutions at GTC:

- High-performance Micron 9550 NVMe and Micron 7450 NVMe SSDs included on the GB200 NVL72 recommended vendor list.
- Micron's PCIe Gen6 SSD, demonstrating over 27GB/s of bandwidth in successful interoperability testing with leading PCIe switch and retimer vendors, driving the industry to this new generation of flash storage.
- Storing more data in less space is essential to get the most out of AI data centers. The Micron 61.44TB 6550 ION NVMe SSD is the drive of choice for bleeding-edge AI cluster exascale storage solutions, by delivering over 44 petabytes of storage per rack,⁸ 14GB/s and 2 million IOPs per drive inside a 20-watt footprint.

As AI and generative AI expand and are integrated on-device at the edge, Micron is working closely with key ecosystem partners to deliver innovative solutions for AI for automotive, industrial and consumer. In addition to high performance requirements, these applications require enhanced quality, reliability and longevity requirements for application usage models.

- One example of this type of ecosystem collaboration is the integration of Micron LPDDR5X on the NVIDIA DRIVE AGX Orin platform. This combined solution provides increased processing performance and bandwidth while also reducing power consumption.
- By utilizing Micron's 1β (1-beta) DRAM node, LPDDR5X memory meets automotive and industrial requirements and offers higher speeds up to 9.6 Gbps and increased capacities from 32Gb to 128Gb to support higher bandwidth.
- Additionally, Micron LPDDR5X automotive products support operating environments from -40 degrees Celsius up to 125 degrees Celsius to provide a wide temperature range that meets automotive quality and standards.

Micron will exhibit its full data center memory and storage product portfolio at GTC, March 17 – 21, in **booth #541**.

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

Additional Resources:

- [SOCAMM web page](#)
- [HBM3E web page](#)
- [HBM3E 12H 36GB blog](#)
- [Micron SSDs](#)
- [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 Media Relations Contact

Kelly Sasso
 Micron Technology, Inc.
 +1 (208) 340-2410
ksasso@micron.com

¹ Calculations based on comparing one 64GB 128-bit bus SOCAMM to two 32GB 64-bit bus RDIMMs.

² Calculated using transfer speeds comparing 64GB 2R 8533MT/s SOCAMM and 64GB 2Rx4 6400MT/s RDIMMs.

³ Calculated area between one SOCAMM and one RDIMM.

⁴ Calculated based on power used in watts by one 128GB, 128-bit bus width SOCAMM compared to two 128GB, 128-bit bus width DDR5 RDIMMs.

⁵ Comparison based on HBM3E 36GB capacity versus HBM3E 24GB capacity when both are at the 12x10mm package size.

⁶ Based on internal calculations, and customer testing and feedback for Micron HBM3E versus the competition's HBM3E offerings.

⁷ Calculated bandwidth by comparing HBM4 and HBM3E specifications.

⁸ Assumes 20x 61.44TB E3.S SSDs in a 1U server with 20x E3.S slots available for storage and that 36 rack units are available for the servers in each rack.