



## Micron in High-Volume Production of HBM4 Designed for NVIDIA Vera Rubin, PCIe Gen6 SSD and SOCAMM2

March 16, 2026 at 5:00 PM EDT

### News highlights:

- **HBM4 36GB 12H in high-volume production**, designed for NVIDIA® Vera Rubin — greater than 2.8 TB/s<sup>1</sup> and with 20% better power efficiency<sup>2</sup>
- **Industry's first PCIe® Gen6 SSD in high-volume production**<sup>3</sup> — the Micron 9650 data center SSD delivers up to two times the read performance of Gen5 at 100% higher performance per watt<sup>4</sup> and optimized for agentic AI workloads on NVIDIA BlueField-4 STX architecture
- **192GB SOCAMM2 in high-volume production** — expanding low-power, high-capacity memory for AI and HPC workloads on the NVIDIA Vera Rubin platform, part of a broad portfolio of SOCAMM2 products spanning 48GB to 256GB capacities

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

SAN JOSE, Calif., March 16, 2026 (GLOBE NEWSWIRE) -- **GTC 2026** — AI-optimized memory and storage have become strategic assets driving system performance to enable AI workloads and infrastructure to deliver real-world value. Micron Technology, Inc. (Nasdaq: MU), has begun volume shipment of its HBM4 36GB 12H in the first quarter of calendar year 2026 and is designed for NVIDIA Vera Rubin. With HBM4, Micron achieves over 11 Gb/s pin speeds,<sup>5</sup> enabling a bandwidth greater than 2.8 TB/s, representing a 2.3 times bandwidth and greater than 20% power efficiency<sup>2</sup> improvement over its HBM3E.

Looking towards further HBM cube capacity expansion, Micron has demonstrated advanced packaging capability of stacking 16 die of HBM by shipping samples of HBM4 48GB 16H to customers. This milestone delivers a 33% increase in capacity per HBM placement compared to the HBM4 36GB 12H offering.<sup>6</sup>

“The next era of AI will be defined by tightly integrated platforms developed through joint engineering innovations across the ecosystem. Our close collaboration with NVIDIA ensures that compute and memory are designed to scale together from day one,” said Sumit Sadana, executive vice president and chief business officer at Micron Technology. “At the heart of this is Micron’s HBM4, the engine of AI, delivering unprecedented bandwidth, capacity and power efficiency. With HBM4 36GB 12H, alongside the industry’s first SOCAMM2 and Gen6 SSD now in high-volume production, Micron’s memory and storage form a core foundation that unlocks the full potential of next-generation AI.”

Micron SOCAMM2 is designed for NVIDIA Vera Rubin NVL72 systems and standalone [NVIDIA Vera CPU](#) platforms, enabling up to 2TB of memory and 1.2 TB/s of bandwidth per CPU.

Micron is the first company to mass-produce a PCIe Gen6 data center SSD.<sup>3</sup> The Micron 9650 is optimized for energy efficiency and liquid-cooled environments, delivering high-speed, low-latency data access for AI training and inference workloads with the NVIDIA BlueField-4 STX reference architecture, supporting up to 28 GB/s sequential read throughput and 5.5 million random read IOPS. Micron 7600 and 9550 SSDs offer customers PCIe Gen5 SSDs, increasing their architectural design choices.

### Showcasing Micron innovation at NVIDIA GTC 2026

At NVIDIA GTC 2026, Micron will showcase how its advanced memory and storage portfolio enables end-to-end AI acceleration from the data center to the edge. Attendees can find detailed information at the **Micron booth, 1407**.

### Additional resources:

- [Micron AI data center solutions webpage](#)
- [Micron HBM4 image gallery](#)
- [Micron 9650 PCIe Gen6 image gallery](#)
- [From breakthrough demo to deployment path: SCADA on production-grade PCIe Gen6 hardware at NVIDIA GTC 2026](#)

### 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. 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](https://micron.com).

© 2026 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. NVIDIA and the NVIDIA logo are trademarks and/or registered

trademarks of NVIDIA Corporation in the U.S. and other countries. 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

**Micron Investor Relations Contact:**

Satya Kumar  
+1 (408) 450-6199  
satyakumar@micron.com

---

<sup>1</sup> Comparing the bandwidth of HBM4 to the previous generation HBM3E with the same capacity and stack height (36GB 12H).

<sup>2</sup> Comparison of HBM power efficiency: (HBM4 12H vs. HBM3E 12H at 9.2 Gbps; workload pattern 60/30/10/75) [Internal power calculator tool].

<sup>3</sup> On February 12, 2026, Micron [announced](#) mass production of the 9650 SSD. At the time of blog publication, no other Gen6 SSDs from the top five OEM data center SSDs suppliers by revenue had announced mass production, according to the Forward Insights analyst report *SSD Supplier Status Q4/25* (Nov 2025).

<sup>4</sup> For the Micron 9650 SSD performance results, see table 1 in the [Micron 9650 SSD Product Brief](#).

<sup>5</sup> Based on internal Micron testing and confidential customer test vehicle validation.

<sup>6</sup> The 48GB cube provides a 33% increase in capacity per HBM placement when comparing the capacity of HBM4 36GB 12H to HBM4 48GB 16H.