

Technology and products leadership for the AI era

October 21, 2026

Legal disclaimer

This presentation contains forward-looking statements regarding our industry, our strategic position, customer demand, investments and the goals for such investments, our business opportunities, our products and technology, including our product and technology roadmaps and expected capabilities and timing, and our financial and operating performance.

These forward-looking statements are subject to a number of risks and uncertainties that could cause actual results to differ materially. For a discussion of these risks and uncertainties, please refer to our filings with the Securities and Exchange Commission.

Although we believe that the expectations reflected in the forward-looking statements are reasonable, we cannot guarantee future results, levels of activity, performance, or achievements. We are under no duty to update any of the forward-looking statements, except as required by applicable law.

Scott DeBoer

President and Chief Technology
and Products Officer




Accelerating intelligence to enrich life for all

- AI has created unprecedented opportunity for memory and storage; Micron has executed and delivered
- Micron's technology leadership is foundational to delivering differentiated, industry-leading products
- AI accelerates Micron team members' innovation and execution
- Long-horizon R&D investments enhance Micron's future competitiveness
- Co-innovation with partners builds leadership solutions with Micron memory and storage at the core
- Micron is positioned to extend technology and product leadership for the next decade and beyond

Each wave of AI increases memory demand

AI related demand spans every layer of the memory and storage hierarchy



Illustrative for concept: not to scale
HBM capacity per GPU: NVIDIA product specifications (H100 SXM 80GB HBM3; Blackwell Ultra B300 288GB HBM3E). GB = gigabyte | GPU = graphics processing unit | HBM = high-bandwidth memory

Architecture evolving for AI workloads

Micron's portfolio spans the memory and storage hierarchy for the data center and beyond

Standard, custom and vertical HBM

Training and real-time inference at the highest bandwidth

High-capacity and low-power DRAM modules


Active model context

High-performance SSDs and AI-class NAND

KV cache and context offload, attention

High-capacity SSDs

Training data, checkpoints, archives at scale



Data Center Memory Hierarchy



Autonomous vehicles



Robotics and humanoids




AI PCs



Mobile

Technology enables industry firsts & best-in-class products




Claims of industry leadership, speed, and compactness apply at the time of announcement and are not ongoing claims.

ASIL = Automotive Safety Integrity Level | DDR5 = fifth-generation double data rate | G9 = ninth generation | GB = gigabyte | HBM4 = fourth-generation HBM | LPDDR5X = low-power double data rate 5X | MT/s = megatransfers per second
PCIe = peripheral component interconnect express | QLC = quad-level cell | RDIMM = registered dual in-line memory module | SOCAMM = small outline compression attached memory module
SOCAMM2 = second-generation SOCAMM, standardized in JEDEC | SSD = solid state drive | TB = terabyte | UFS = universal flash storage

Inference raises the bar for memory technology

Prompt

Your question and relevant context



Storage (SSDs)

Holds source information
Retrieved when needed



LLM inference

1 Prefill


Read the prompt

2 Keep


Save calculations to reuse

3 Decode

Read saved calculations per token



Saved calculations (KV cache)



Low-power DRAM modules


Energy-efficient working capacity*

HBM feeds the processor repeatedly



Answer

Response generated one token at a time



First token from prefill


Then one token
per decode step

Memory and storage innovation drive AI performance

AI = artificial intelligence | HBM = high-bandwidth memory | KV = key/value | LLM = large language model | SSD = solid-state drive


*Role depends on system architecture. Simplified flow. Storage is not read for every token. Bottlenecks vary.

Memory has never stopped scaling



>400x

more bits per DRAM wafer since 1995



At 1γ dimensions:
>1,500 capacitors
across the width of
a 50 μm hair

Source: Micron | GB = gigabyte | nm = nanometer | μm = micrometer

Micron’s sustained technology leadership

	2019	2020	2021	2022	2023	2024	2025	2026
DRAM 4 consecutive leadership nodes	1z First to volume production		1α Industry’s first shipment		1β Industry’s first shipment		1γ Industry’s first sample shipment	
HBM 3 consecutive generations performance and power leadership						HBM3E 8H volume production	HBM3E 12H shipments	HBM4 12H shipments HBM4E samples to key customers
NAND 3 consecutive leadership nodes		G7 Industry’s first shipment		G8 Industry’s first shipment		G9 Industry’s fastest; first G9-based SSD		

Sources: Micron press releases: 1z (Aug. 15, 2019), 1α (Jan. 26, 2021), 1β (Nov. 1, 2022), 1γ (Feb. 25, 2025), G7 (Nov. 9, 2020), G8 (July 26, 2022), G9 (July 30, 2024), HBM3E (Feb. 26, 2024).
 1z: first to mass-produce 16Gb DDR4 on 1z. 1γ: first to ship 1γ DDR5 samples. G9: first G9 TLC NAND shipped in an SSD. HBM3E power vs. competition at time of announcement.
 12H = 12 layers of DRAM die | 8H = 8 layers of DRAM die | G7 = seventh generation | G8 = eighth generation | G9 = ninth generation | HBM = high-bandwidth memory
 HBM3E = HBM third generation, extended | SSD = solid state drive

The pace of memory innovation

If a Porsche 911 had evolved at the same pace as memory between 1995 and 2026...



1995

270 hp · \$60,000 · 23 MPG



2026


2,700,000 hp · \$5.00 · 23,000 MPG

≈ the power of 10 aircraft carriers

Ratios applied to match each metric. Performance improvement is captured in the horsepower, cost improvement is captured in the price, and power efficiency is captured in the MPG.
hp = horsepower | MPG = miles per gallon


HBM4 vs. worldwide internet bandwidth

One HBM4 stack
2.8 TB/s



2026

~100 HBM4 stacks match
worldwide internet bandwidth



Worldwide internet bandwidth

≈ 282 TB/s

1 square = 1 stack (2.8 TB/s)

Memory shapes AI capacity, performance and cost per token

— AI workload

Context windows,
agents and
reasoning at scale

— Compute

FLOPS,
utilization and
stall time

— Memory

Bandwidth, capacity
and thermals

— Storage

Capacity and
cost per token

Memory has become the AI bottleneck
Memory capability and architecture must evolve

Accelerating intelligence to enrich life for all


- AI has created unprecedented opportunity for memory and storage; Micron has executed and delivered
- **Micron's technology leadership is foundational to delivering differentiated, industry leading products**
- AI accelerates Micron team members' innovation and execution
- Long-horizon R&D investments enhance Micron's future competitiveness
- Co-innovation with partners builds leadership solutions with Micron memory and storage at the core
- Micron is positioned to extend technology and product leadership for the next decade and beyond

Core technology leadership underpins future architectures

Memory design features, performance and power are the differentiators


DRAM

Planar DRAM



The DRAM cell is nearing atomic limits

3D DRAM



So, DRAM must now build up


NAND

Planar NAND



Scaled until the release of 3D NAND in 2015

3D NAND z-scaling



Same footprint, more bits with each layer


3D NAND xyz-scaling



Shrinking cells and increasing layers

DRAM roadmap: continuous scaling through innovation

Micron is extending 6F² leadership and developing 3D DRAM; not pursuing 4F² due to increased complexity/cost



Directional; no dates implied. Node names and phases subject to change.
EUV = extreme ultraviolet lithography | High-NA = high numerical aperture

1δ: the world's most advanced DRAM


Fifth consecutive node of DRAM industry leadership – ramping in the second half of 2027

Industry-leading technology manufactured in the U.S.

- Mid-20s percentage bit density improvement from 1γ; optimized cost
- Angstrom-level process control
- Equipment intelligence, new materials

Leadership in power, performance and cost

- Advanced HKMG CMOS
- High-performance products



Supports differentiated HBM5, vHBM, vEIM products

- Advanced TSV technology
- Thermal-aware design

Enables next-gen LPDDR, DDR and graphics modules

- Next-generation mobile memory
- High-density modules for AI workloads

DRAM CMOS logic roadmap and performance through 2030

Accelerated CMOS enabling future differentiated product requirements

>10x

CMOS logic performance

Poly SiON → 2nd Gen 3D CMOS

- Foundation of Micron's AI memory leadership
- Multiple nodes of power and performance leadership
- Advanced CMOS roadmap to meet AI bandwidth demands



Next-generation DRAM for AI

More bandwidth and capacity, with greater power efficiency

LP6 for data center and mobile

>35%

Improved power vs. LP5

High-reliability DDR6 for server

>2x

Bandwidth vs. DDR5

High-capacity RDIMM

1TB

Memory capacity


Performance improvements are targets. 1 TB RDIMM is a capability, not a product commitment.

DDR5 = fifth-generation double data rate | DDR6 = sixth-generation double data rate | LP6 = low-power double data rate 6 | LP5 = low-power double data rate 5 | RDIMM = registered dual in-line memory module | TB = terabyte

Building on HBM leadership to drive AI performance

Each architecture step shortens the memory–compute path to deliver more tokens per watt


Standard HBM



Industry baseline

Memory beside compute on an interposer


Custom HBM



Co-designed with customers

Base-die logic adds differentiated features


Vertical HBM (vHBM)



Memory stacked on compute

Shortest memory–compute path

Custom base die




Higher bandwidth per watt · More tokens per watt

Advanced packaging innovations enable future HBM


Next-generation memory-compute architectures leverage novel packaging innovations

Interconnect density




- TSV and μ bump scaling
- Hybrid copper bonding
- Metallurgy innovations

Stacking technology



- 16H μ bump and hybrid stack
- Bandwidth and thermal improvement
- 2.5D HBM + 3D vHBM integration


Thermal innovations



- Heat spreader over PHY
- Front and backside innovations
- Thermally aware design

Token efficiency scales with bandwidth, energy efficiency, capacity

Product roadmap expands to custom HBM, vHBM and vEIM




vHBM and vEIM enablers

Packaging

Hybrid bonding


Higher stack density and finer interconnect pitch



Thermals

Advanced thermals

Lower thermal resistance and better heat spreading as power rises



System co-design

JEDEC + custom + 3D

Base die, PHY, packaging and workload optimized together



Token-efficiency positions are illustrative. Product timing represents targeted high-volume manufacturing; timing and performance targets are directional and subject to change.
HBM = high-bandwidth memory | JEDEC = Joint Electron Device Engineering Council | PHY = physical layer | TB/s = terabytes per second | vEIM = vertical edge inference memory | vHBM = vertical HBM

Proven leadership extends with G11 generational leap

Each generation of NAND technology leadership builds on the last

Scale Today

G9 NAND

Industry leading technology supporting a broad portfolio

Leap Ahead

G11 NAND


- ~2x bit density vs. G9
- Up to 1.5x I/O speed
- >25% less data transfer energy per bit vs. G9

G11: the world's most advanced NAND

- Wafer-on-wafer bonding
- >400 layers and industry-leading cell pitch
- Volume ramp in 2H 2027

Beyond G11: redefining the limits

- Breakthrough high-mobility material
- Ultra dense cell architecture
- Energy-efficient devices and data paths



Micron AI-class NAND design optimized for AI workloads

AI-class NAND architected to accelerate emerging AI workloads and applications

FINE-GRAINED ACCESS

High-IOPS AI-class NAND

High-IOPS SLC NAND designs for GPU-initiated, small-block data access beyond system memory

Emerging use cases

Vector search

Graph analytics

AI context access

NEAR-COMPUTE BANDWIDTH

High-bandwidth AI-class NAND

TSV-stacked NAND with a logic base die, enables high-capacity tier close to the accelerator

Emerging use cases

Model weights

Long context

KV cache

More data closer to compute improves utilization, scalability and system efficiency

Micron future AI-optimized storage solutions

AI-class NAND enables purpose-built storage solutions wherever AI data lives

Data Center and AI infrastructure



Maximize AI throughput

Keep accelerators fed with massive data sets held closer to compute

Next-generation SSDs

High-performance, high-capacity (up to 512TB)


High-IOPS

25M–100M IOPS small-block access

AI acceleration

Optimized accelerator data pipelines

On-prem and AI PC



Power private, local AI

High-bandwidth, power-efficient storage for private and on-device inference

On-prem services

High-capacity, low-latency private AI

Workstations

Model loading, fine-tuning, creation

AI PC

Power-efficient local AI across devices

Automotive, robotics, mobile and wearables



Scale AI to the edge

Compact, low-power, high-bandwidth storage for AI in the physical world

Automotive

Intelligent cockpits and driver assistance

Humanoid robotics

Real-time perception and context retention

Mobile and wearables

Low-power on-device AI and immersive AR

Accelerating intelligence to enrich life for all

- AI has created unprecedented opportunity for memory and storage; Micron has executed and delivered
- Micron's technology leadership is foundational to delivering differentiated, industry leading products
- **AI accelerates Micron team members' innovation and execution**
- Long-horizon R&D investments enhance Micron's future competitiveness
- Co-innovation with partners builds leadership solutions with Micron memory and storage at the core
- Micron is positioned to extend technology and product leadership for the next decade and beyond


AI makes Micron 2x faster: invention to volume production

Industry-leading use of AI across development teams; two examples included below

Data Center SSD controller

2x faster to silicon
tape-out with AI


Months to silicon tape-out



Data Center SSD product

>2x faster to silicon ready for
customer releases with AI

Months to silicon ready for customer releases



Accelerating intelligence to enrich life for all

- AI has created unprecedented opportunity for memory and storage; Micron has executed and delivered
- Micron's technology leadership is foundational to delivering differentiated, industry leading products
- AI accelerates Micron team members' innovation and execution
- **Long-horizon R&D investments enhance Micron's future competitiveness**
- **Co-innovation with partners builds leadership solutions with Micron memory and storage at the core**
- Micron is positioned to extend technology and product leadership for the next decade and beyond

Long-horizon R&D investments

Investing in capabilities for long-term leadership and a greater role in shaping the future

3-10 YEARS

**Micron
Ventures**

\$550M

total capital commitment

Investing across the AI stack to
strengthen ecosystem engagement

5-10 YEARS

**Albany NanoTech +
NY CREATES**

**Advanced
lithography**

Accelerating the high-NA EUV
and lithography roadmap

10+ YEARS

**Micron
Research Labs**

\$10B

planned over the next decade

First dedicated memory research hub in the
U.S., looking beyond a 10-year horizon

Building the next generation of memory
researchers and technology leaders

Separately, Micron has committed more than \$250B to manufacturing and R&D across the United States.


EUV = extreme ultraviolet lithography | High-NA = high numerical aperture | NY CREATES = New York Center for Research, Economic Advancement, Technology, Engineering and Science

Micron and ecosystem leaders create customer advantage

Solving customer challenges through memory and storage innovation, from startups to hyperscalers



+ micron™



>1.1 EB¹

in a single rack

Micron 6600 ION 245TB SSDs
deliver 15.8PB per 2U server


Customer challenge

AI data growth is outpacing finite rack space, power and cooling

Joint solution

- Engineering-level SSD and platform optimization preserves power and space for more compute.
- Delivers more than 2.1x throughput density vs. competitors.¹

Meta + micron™



LPDDR6

for data center systems

Performance • capacity
power-efficiency • reliability

Customer challenge

Power efficient scaling of memory bandwidth is a primary need for advanced compute and AI


What we are delivering

- Advancing LPDDR6 as data center-class memory architecture
- Meta platform and application expertise combined with Micron memory innovation for fleet-scale optimization.

¹ WEKA NeuralMesh data reduction, assuming 3:1. Results depend on configuration and workload. [Meta X Micron Joint Whitepaper: LPDDR6: A Scalable, Power-Efficient Memory for Data Centers](#)
EB = exabyte | LPDDR6 = low-power double data rate 6 | PB = petabyte | SSD = solid state drive | TB = terabyte

Co-innovation delivers scaled solutions to meet AI demand

One ecosystem, building leadership solutions together, with memory at the core



Accelerating intelligence to enrich life for all

- AI has created unprecedented opportunity for memory and storage; Micron has executed and delivered
- Micron's technology leadership is foundational to delivering differentiated, industry leading products
- AI accelerates Micron team members' innovation and execution
- Long-horizon R&D investments enhance Micron's future competitiveness
- Co-innovation with partners builds leadership solutions with Micron memory and storage at the core
- Micron is positioned to extend technology and product leadership for the next decade and beyond



© 2026 Micron Technology, Inc. All rights reserved. Information, products, and/or specifications are subject to change without notice. All information is provided on an "AS IS" basis without warranties of any kind. Statements regarding products, including statements regarding product features, availability, functionality, or compatibility, are provided for informational purposes only and do not modify the warranty, if any, applicable to any product. Drawings may not be to scale. Micron, the Micron logo, the M logo, Intelligence Accelerated™, and other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.