

B U R K E W E A L T H M A N A G E M E N T

Delivering the Compute

Powering the AI Revolution

Focused Growth Portfolio | Delivering the Compute

N V D A · A S M L · M U

May 2026

The Thesis

Without compute, there is no way to generate tokens. Every dollar of AI revenue flowing through the economy — from chatbots to autonomous agents to enterprise copilots — depends on the chips, equipment, and memory that make computation possible.

The demand for AI compute is growing faster than supply. Hyperscalers are spending over \$600B in 2026 on infrastructure, yet capacity remains constrained. We have positioned the portfolio across the three companies that sit at the foundation of this buildout — the GPU designer, the lithography monopoly, and the memory supplier that makes it all work together.

Nvidia (NVDA)

GPU Design

The world's leading AI chip designer. Q1 revenue of \$81.6B (+85% YoY). ~\$1T in combined orders through 2027. Q2 guidance of \$91B beat consensus by \$4B.

ASML (ASML)

Lithography Equipment

The absolute monopoly in EUV lithography. Every advanced AI chip requires ASML machines. Zero competitors. 2026 guidance raised to €36-40B.


Micron (MU)

High-Bandwidth Memory

The only U.S.-based HBM supplier. Q2 revenue nearly tripled YoY to \$23.9B. Just crossed \$1 trillion market cap. HBM4 in volume production for Vera Rubin.

The Demand: Hyperscaler Capex Supercycle

Hyperscaler Capital Expenditure (Top 5)



Sources: Company reports, Bernstein

An unprecedented infrastructure buildout

The five largest hyperscalers are projected to spend over \$600B in 2026, with roughly 75% directed at AI infrastructure. Consensus capex estimates have been too low for two years running.

\$1.15T

Cumulative hyperscaler capex 2025-27

75%

Of 2026 capex directed at AI infrastructure

~\$1T

Combined Nvidia orders (Blackwell + Rubin thru 2027)

\$91B


Nvidia Q2 FY27 guide (beat \$87B consensus)

Every dollar of this capex flows through our three holdings: Nvidia designs the GPUs, ASML builds the machines that manufacture them, and Micron supplies the memory they require.

Nvidia: Generational Compute Gains

Each new GPU architecture delivers exponential improvements in AI inference efficiency — more tokens, less power, lower cost.

AI Inference: Relative Tokens per Watt by GPU Generation



Indexed to Hopper H100 = 1x | Sources: Nvidia, company presentations

Hopper → Blackwell

68x performance

Enabled frontier model training at scale. Cost per token dropped 87%.

Blackwell → Rubin

35x inference

Up to 10x lower cost per token. Vera Rubin in production, partner availability H2 2026.

Hopper → Rubin

900x total

Four years of compounding innovation. Annual cadence locked in through Feynman (2028).

Nvidia (NVDA)

The essential computing platform powering the AI revolution

KEY FACTS — Q1 FY2027 (May 20, 2026)

- Q1 Revenue: \$81.6B (+85% YoY, +20% QoQ) — record quarter
- Data Center Revenue: \$75.2B (+92% YoY)
- Non-GAAP EPS: \$1.87 vs. \$1.77 consensus
- Q2 FY2027 Guidance: \$91.0B (beat \$87B consensus by \$4B)
- Supply Commitments: \$119.0B (up from \$95.2B in Q4)
- Combined Order Backlog: ~\$1 trillion (Blackwell + Rubin thru 2027)
- \$80B additional share repurchase; dividend increased 25x
- Vera Rubin: 35x higher inference throughput vs Blackwell

WHY IT FITS THE THEME

Nvidia sits at the apex of the compute value chain. Every hyperscaler capex dollar ultimately flows through GPU procurement. The CUDA ecosystem creates competitive protection beyond hardware, and the ~\$1T legacy CPU-to-GPU transition represents a multi-year structural tailwind. The annual cadence — Blackwell, Rubin, Feynman — ensures the upgrade cycle remains alive.

INVESTMENT THESIS

Nvidia is the leading chip provider to the datacenter of the future. The \$1T legacy infrastructure transition, AI agent demand, and recommender system upgrades drive a convergence of demand that is structural, not cyclical. Combined orders of ~\$1 trillion through 2027 and Q2 guidance of \$91B provide extraordinary forward visibility.

\$81.6B

Q1 Revenue

85%

YoY Growth

\$91B

Q2 Guide

~\$1T

Order Backlog


900x

Hopper→Rubin

ASML: The Irreplaceable Bottleneck

Every advanced AI chip in the world requires ASML's EUV lithography machines to manufacture. There are zero alternatives.

ASML System Backlog (€B)



Record demand, raising guidance

ASML reported Q1 2026 revenue of €8.8B (beat), gross margin of 53% (top of guidance), and raised full-year 2026 guidance to €36-40B from €34-39B. Order intake remains "very strong" as customers accelerate AI fab construction.

100%

Global EUV market share

€36-40B

2026 revenue guidance (raised)

\$8.0B

SK Hynix EUV order (Q1 2026)

€46.5B

Total performance obligations

ASML (ASML)

The monopoly provider of EUV lithography — essential equipment for every advanced AI chip

KEY FACTS

- Q1 2026 Revenue: €8.8B (beat) — Gross Margin 53.0%
- Q1 2026 EPS: €7.15 vs. €6.00 consensus
- 2025 Full-Year Revenue: €32.7B — Net Income €9.6B
- 2026 Guidance Raised: €36-40B (from €34-39B)
- EUV Market Position: 100% — absolute monopoly
- System Cost: \$200M-\$300M per EUV machine
- Performance Obligations: €46.5B (multi-year visibility)
- High-NA EUV entering early adoption, first systems shipped

WHY IT FITS THE THEME

ASML is the irreplaceable bottleneck in the semiconductor supply chain. Every advanced AI chip — whether from Nvidia, AMD, or custom hyperscaler silicon — requires ASML's EUV machines. Global decoupling is a net positive as Western governments build domestic foundries. The guidance raise to €36-40B confirms AI capex is accelerating, not plateauing.

INVESTMENT THESIS

ASML has rare long-term visibility in a cyclical industry. Its EUV monopoly is tied to structural semiconductor demand growth driven by AI. Global foundry investment is a multi-year tailwind as geopolitical decoupling accelerates chip reshoring. High-NA EUV extends the technology roadmap and deepens the competitive moat.

€8.8B

Q1 Revenue

100%

EUV Share

€36-40B

2026 Guide

53%

Q1 Gross Margin


€46.5B

Obligations


The Memory Imperative: HBM4 Powers Vera Rubin

As AI models grow to trillions of parameters, the bottleneck has shifted from raw compute power to memory bandwidth. Each new GPU generation demands exponentially more — and faster — memory.

HBM per GPU (GB) & Bandwidth



Global HBM Market (\$B)



Micron began high-volume HBM4 production in Q1 2026 for Vera Rubin. HBM4 delivers 2.3x bandwidth and 20% better power efficiency vs. HBM3E. First five-year strategic customer agreement signed with fixed volume commitments.

Micron Technology (MU)

The only U.S.-based HBM supplier — revenue nearly tripled, now a \$1 trillion company

KEY FACTS — Q2 FY2026 + May 2026 Update

- Q2 Revenue: \$23.86B — nearly tripled YoY (beat by \$3.6B)
- Q2 EPS: \$12.07 vs. \$9.33 consensus
- Q2 Gross Margin: 75% — operating profit \$16.1B
- Q3 Guidance: \$33.5B revenue, 81% gross margin, \$19.15 EPS
- Market Cap: Crossed \$1 trillion (May 26, 2026)
- UBS PT tripled to \$1,625 — reframed as "AI-native infrastructure giant"
- HBM4: Volume production for Vera Rubin; first 5-year customer agreement
- HBM4 48GB 16H samples shipped; HBM4E development underway
- Stock up 178% YTD; 30% dividend increase

WHY IT FITS THE THEME

AI datacenter demand has added a massive structural component to high-bandwidth memory. Each Rubin GPU requires 288GB of HBM4 across 8 stacks — a 50% increase over Blackwell. The "memory wall" is the new bottleneck: compute power without sufficient memory bandwidth is useless. Micron is the only U.S.-based supplier in this critical triopoly with ~24% market share.

INVESTMENT THESIS

Massive AI datacenter capex is driving a prolonged positive memory cycle unlike any before. HBM pricing stability — with 60-70% of industry DDR5 volumes locked into long-term supply agreements — transforms historically hyper-cyclical memory into a durable business. UBS projects \$400B+ cumulative free cash flow through 2029. Q3 guidance of \$33.5B at 81% margins signals extraordinary momentum.

\$23.9B

Q2 Revenue

196%

YoY Growth

\$1T+

Market Cap

81%

Q3 GM Guide

\$33.5B

Q3 Rev Guide

Investment Outlook

Company	Ticker	Role	Latest Rev	YoY Growth	Key Catalyst	YTD
Nvidia	NVDA	GPU Design	\$81.6B (Q1)	85% YoY	\$91B Q2 guide; ~\$1T backlog	+15%
ASML	ASML	Lithography	€8.8B (Q1)	Beat est.	Guidance raised €36-40B	+53%
Micron	MU	HBM Memory	\$23.9B (Q2)	196% YoY	\$1T mkt cap; Q3 guide \$33.5B	+214%

OUTLOOK

These three companies represent the physical foundation of the AI revolution. No AI model trains, no agent reasons, no token generates without the chips Nvidia designs, the machines ASML builds to manufacture them, and the memory Micron supplies to feed them.

Since our last update, Nvidia reported record Q1 revenue of \$81.6B and guided Q2 to \$91B — both above consensus. ASML raised full-year guidance. And today, Micron crossed \$1 trillion in market cap after UBS tripled their price target, reframing the company as an AI-native infrastructure giant rather than a cyclical memory maker. The thesis is playing out.

All Three Companies Beat Most Recent Estimates | Multi-Year Demand Visibility Across the Board

YTD returns are as of 5/26/26. Returns reflect individual stock price performance and are not representative of any portfolio or account performance.

Portfolio Context

These three companies are the physical foundation upon which the entire AI economy is built.

Nvidia designs the GPUs. ASML builds the only machines that can manufacture them. Micron supplies the memory that makes them work. Together, they form an unbroken chain from silicon to intelligence.

The demand for AI compute is structural and growing. Hyperscalers are spending at record levels, GPU architectures are advancing on an annual cadence, and the memory requirements per chip are expanding with every generation. We believe the infrastructure buildout is in its early innings.

Without compute, there is no AI. Without these companies, there is no compute.

Important Disclosures

Burke Wealth Management, LLC ("BWM") is a registered investment adviser with the SEC. Registration does not imply a certain level of skill or training. This material is provided for informational purposes only and does not constitute investment advice, a recommendation, or an offer to buy or sell any security. The information contained herein should not be construed as personalized investment advice and should not be considered as a solicitation to buy or sell any security or to engage in a particular investment strategy.

Past performance is not indicative of future results. All investments involve risk, including the possible loss of principal. The year-to-date (YTD) returns presented represent individual stock price performance only and are not representative of any BWM portfolio, composite, or individual account performance. Actual client returns will differ based on timing of investments, fees, individual account holdings, and other factors. There is no guarantee that any investment strategy will achieve its objectives.

The securities identified and described herein do not represent all of the securities purchased, sold, or recommended for client accounts. It should not be assumed that any of the securities discussed were or will be profitable. References to specific securities are for illustrative and informational purposes only and are not intended as recommendations to purchase or sell such securities.

Forward-looking statements, estimates, and projections contained in this presentation are based on current expectations and assumptions that are subject to risks and uncertainties. Actual results may differ materially. BWM does not have any obligation to update or revise any forward-looking statements. For additional information about BWM, including its advisory fees and services, please refer to the firm's Form ADV Part 2A, which is available upon request or at adviserinfo.sec.gov.

Data Sources

Company Financials: Nvidia Q1 FY2027 8-K and CFO Commentary (SEC, May 20, 2026); Nvidia Q4 FY2026 8-K (SEC, Feb 25, 2026); ASML Q1 2026 Earnings Release (Apr 15, 2026); Micron Q2 FY2026 Earnings Release (Mar 18, 2026). All financial data sourced from company SEC filings and earnings press releases.

Market & Industry Data:

Hyperscaler capex estimates: Company reports (MSFT, GOOG, AMZN, META, ORCL).

Stock Prices & Returns:

YTD returns calculated using closing prices as of May 26, 2026. Individual stock returns do not reflect advisory fees, transaction costs, or taxes.