



The \$1 Trillion AI Semiconductor Tectonic Shift

**Asia Pacific GSA Executive Forum
November 5th, 2025**

**Mario Morales, VP Head of Corporate Strategy and
Partnerships**

AMD 
together we advance_

Agenda

- Key Takeaways
- IT Spending Driven by AI
- Semiconductor and Capex Outlook
- Inference accelerates growth
- Enterprise AI moving from experiment to deployment
- Constraints lead to value capture for semiconductor ecosystem
- Recommendations and Q&A

Cautionary Statement


This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) such as worldwide IT spending trends, semiconductor and capital expenditure forecasts, Data Center Artificial Intelligence (AI) total addressable market (TAM), Edge AI TAM, AI usage, Enterprise AI forecast, Edge AI silicon forecast, foundry forecast, memory market forecast, expected AI benefits and opportunities, which are made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "projects" and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this presentation are based on current beliefs, assumptions and expectations, speak only as of the date of this presentation and involve risks and uncertainties that could cause actual results to differ materially from current expectations. Such statements are subject to certain known and unknown risks and uncertainties, many of which are difficult to predict and generally beyond AMD's control, that could cause actual results and other future events to differ materially from those expressed in, or implied or projected by, the forward-looking information and statements. Investors are urged to review in detail the risks and uncertainties in AMD's Securities and Exchange Commission filings, including but not limited to AMD's most recent reports on Forms 10-K and 10-Q.

AMD does not assume, and hereby disclaims, any obligation to update forward-looking statements made in this presentation, except as may be required by law.

Key Takeaways


- *Semiconductor industry to reach \$1 Trillion (Source: IDC Sept. 2025), two years earlier than consensus*
- AI is driving IT spending higher– Semi ecosystem is the primary beneficiary, as non-AI sectors like auto, consumer, and industrial continue prolonged correction
- Inference is set to accelerate as language models continue to get smarter, smaller, cheaper, and more open
- Enterprise adoption has been lagging so far but excitement building for the decade of AI agents and an overdue replacement in HW (AI enabled clients and smartphones)
- Investments have been concentrated in the DC market segment so far, Edge and Client are also set to benefit long term as AI inference and reasoning are adopted across devices and AI changes user interface (UI)
- Significant challenges remain including soaring costs, power shortages, and geopolitical uncertainty-strong and agile company leadership is a prerequisite for sustainable growth and competitive differentiation
- Entire semi ecosystem must scale and partner to realize the potential of AI

Worldwide IT Spending: AI Boom Offsetting Tariff Uncertainty



Source: IDC Worldwide Black Book (2025) growth in constant currency; excludes telecom spending and business/OT services


Today's IT spending epicenter is in the Data Center



Source: Worldwide Black Book, 2025 (growth in constant currency); Personal Device includes wearables and AR/VR viewers; Peripheral includes printers, MFPs and PC monitors; Software includes SaaS, PaaS and on-premise software


Semi and Capex Forecasts continue to be revised up

Total Semi Market Forecast



Source: IDC Semiconductors,
September 2025

Top 8 MDC Capex Wall Street Forecast




Source: Factset Consensus

Inference is set to accelerate


DC AI Accelerator TAM

Inference growing @ >80% CAGR



Source: AMD AAI (June'25)

Edge AI Silicon TAM




Source: IDC Sep 2025

- Reasoning models / Test Time Compute driving compute intensive inference
- AI use cases continue to proliferate as models become smarter, cheaper and more reliable

AI is getting better, cheaper and more open

Longer thinking => Intelligent Responses


ARC AGI 1 Benchmark % Accuracy



Cost per task (log scale)

Model prices falling ~10x per year


\$ per million tokens



— MMLU > 42 — MMLU > 83 — PlanBench > 35

Open Source is closing the gap

LM Arena Elo Scores




— #1 Proprietary Model — #1 OS Model

- AI usage is growing quickly: Google noted 980 trillion monthly tokens in July (2x growth in 2 months), Microsoft noted 500 trillion tokens for Microsoft Foundry API this year (up 7x vs. last year)

Enterprise AI – from experiment to deployment

Share of U.S. businesses with paid subscriptions to AI models, platforms, and tools



Source: Ramp Summer 2025 Business Spending report

Note: Ramp is a payment platform with 40K+ Enterprise customers

- AI today is primarily driven by consumer chatbot use cases, reportedly driving ~70% of OpenAI revenue despite only 5% of consumers being paying subscribers
- Enterprise adoption has been lagging – while many enterprises doing trials / experiments, at scale deployments are concentrated in select verticals / use cases (such as customer service and customer facing automation, call centers, coding)
- ROI is the key concern – AI needs to get smarter, cheaper, more reliable to unlock more use cases leveraging the enterprise data
- Substantial excitement for AI agents, with surveys indicating almost all enterprises expecting to launch AI agents at scale with positive ROI by 2027
- IDC expects Enterprise AI spending of \$1.5T in 2025 – 2028,


Enterprise AI – barriers and challenges as the market shifts

- **Mckinsey:** >75% of enterprises have <20% of their workforce reskilled for AI; <10% have reskilled >50% of their workforce
- **BCG:** 93% are planning to invest in AI, but only 26% of enterprises have scaled AI / seeing ROI
- **KPMG:** Half of leaders are currently scaling their GenAI, but only a third anticipate being able to measure ROI, and none have reached that stage
- **IBM:** CEOs report that only 25% of AI initiatives have delivered expected ROI over the last few years, and only 16% have scaled enterprise wide
- **Accenture: (Sept 2025)** announce large reduction in employee workforce, as company shifts to AI driven services and retools to a staff that is augmented with AI agents to support customers
- **Gartner:** Demonstrating ROI is the number one barrier to scaling enterprise AI
- **Berkshire Hathaway (April 2025):** “We not yet made a conscious big-time effort to invest in AI in terms of pouring a lot of money into this opportunity. My guess is we will be in a state of readiness, and should that opportunity pop up we’ll be in a state where we’ll jump in promptly”
- **OpenAI** has launched an Enterprise Consulting offering to help customers deploy and scale AI which will compete against existing enterprise vendors and management consulting firms

Edge AI Silicon Revenue Forecast by Market Segment

5-year CAGR ('25-'30)	
Primary Client	11%
IoT/Endpoint	15%
Edge Infrastructure	32%

Edge Infrastructure IoT/Endpoint Primary Client




Source: IDC Embedded and Intelligent Systems, Semi AI Module 2025

AI growth also fuels Foundry and Memory Markets



Worldwide Foundry Market Forecast


Revenue USD



- Worldwide foundry market to grow by 21.8% this year, revised upward from our previous 18.7% forecast as TSMC's growth outlook increased from 25% to 32%.
- With strong AI demand serving as a catalyst, the worldwide foundry market CAGR is now expected to reach 15%, up from 14%, highlighting AI and 2nm technology hitting key inflection points of demand and supply

Source: IDC STSI Sep 2025

Memory Market Forecast; HBM CAGR Driving Investment and Growth



Source: IDC Semiconductors Sep 2025

AI Pendulum




Things to watch as AI evolves

Technology breakthroughs to maintain the current curve of innovation



AI Models – continue model improvement at the current trajectory



Technologies – improved fundamental technologies (leading edge nodes, lower power, optics, systems)



Specialization – models and AI agents specialized and optimized for a specific application or vertical market



AI Factories – ability for AI to scale up for mass deployment

Path to productivity and profitability



End User Adoption – ability of AI to increase productivity of end user cases



Business Model – players ability to monetize AI (e.g., traditional DC, telco, subscription SW, new models, ...)



AI Moves Outside DC – drive productivity in “real world” through physical implementations or embodied AI

The AI Pendulum: Growth constraints, costs, and value capture



External growth constraints

- **Power:**
 - Data center power demand set to double by 2030 (IEA)
 - US shortfall alone estimated at 15-45 GW (McKinsey, Morgan Stanley)
 - Power, not silicon, is becoming the primary bottleneck
- **Geopolitical uncertainty/Tech and Public Sector Converging:**
 - Tariffs & export controls disrupt supply chains
 - "On-shoring" increases build-out costs & complexity
 - Advocate for national and global initiatives integrating technology, policy, and economic strategy
 - Existing revenue exposure in China brings higher risk than loss of business due to export control for suppliers



Financial model imbalance


- **AI capex >> revenue:**
 - AI capex is tracking in the hundreds of billions of dollars annually
 - Model vendor revenue is much lower (in the tens of billions) but growing rapidly
 - Model builders like OAI become the next CSPs
 - OpenAI can be a single point of massive success...or
- **Model vendors are not profitable yet:**
 - Industry reports indicate leading model vendors are burning large amounts of cash and exploring new business models to finance the build-out
 - MDCs have called out that the risk of underinvesting >> risk of overinvesting, depreciation will impact margins



Semi winners moving up the stack

- **From chips → platforms**
 - Silicon + interconnect + full SW stack (SDK, libraries)
 - Reference systems & solution stacks
 - Recurring revenue layered on silicon
- **Margin migration**
 - Platform/solution mix drives blended gross margins toward software-like levels (~70%+)
- **Ecosystem moat**
 - Developer communities, ISV certifications, and portability raise switching costs as TAM expands
 - End Demand Spending (companies paying for AI services/models) needs to scale to balance investment

AI Mandate: Entire Semi ecosystem must partner to scale and realize AI opportunity



The "\$10K Chip, \$2 Part" Problem

- Mature-node analog, PMICs, and sensors are as critical as leading-edge wafers and accelerators
- A \$10,000 AI accelerator is useless if it's bottlenecked by the availability of a single \$2 component
- The new market bottlenecks to solve are memory, networking, interconnect, and advanced packaging
- Decades of under-investment in mature-node capacity (40nm+) created a structural supply risk and it exacerbates in the opposite direction as China concentrates production and investment in mature nodes driving down price and margins for the broader industry
- Advanced packaging, interconnects, and high-speed optics are important drivers of performance
- Scaling AI requires a "whole-ecosystem" approach: strategic partnerships and investment across logic, memory, packaging, and the entire mature-node supply chain

