June 24, 2026

BostonPartners

Thinking about AI — An Investor's Framework

Duilio R. Ramallo, CFA | Senior Portfolio Manager

Utopia or Dystopia

What lies ahead?

"Any sufficiently advanced technology is indistinguishable from magic."

Arthur C. Clarke, 1973 revision of the "Hazards of Prophecy

Source: Google Nano Banana. Create a picture of the promise of AI moving from utopia to dystopia. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Google Nano Banana). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

Utopia or Dystopia

What lies ahead?

October 2024 → January 2026

Machines of Loving Grace

How AI Could Transform the World for the Better

The list of positive applications of powerful AI is extremely long (and includes robotics, manufacturing, energy, and much more, but I'm going to focus on a small number of areas that seem to me to have the greatest potential to directly improve the quality of human life. The five categories I am most excited about are:

1. Biology and physical health
2. Neuroscience and mental health
3. Economic development and poverty
4. Peace and governance
5. Work and meaning

- Dario Amodei, CEO and Co-Founder at Anthropic

The Adolescence of Technology

Confronting and Overcoming the Risks of Powerful AI

I think the best way to get a handle on the risks of AI is to ask the following question: suppose a literal "country of geniuses" were to materialize somewhere in the world in ~2027. Imagine, say, 50 million people, all of whom are much more capable than any Nobel Prize winner, statesman, or technologist. The analogy is not perfect, because these geniuses could have an extremely wide range of motivations and behavior, from completely pliant and obedient, to strange and alien in their motivations. But sticking with the analogy for now, suppose you were the national security advisor of a major state, responsible for assessing and responding to the situation. Imagine, further, that because AI systems can operate hundreds of times faster than humans, this "country" is operating with a time advantage relative to all other countries: for every cognitive action we can take, this country can take ten.

- Dario Amodei, CEO and Co-Founder at Anthropic

Hyperscalers Capital Spending

Over $5 trillion cumulative 2024-2030

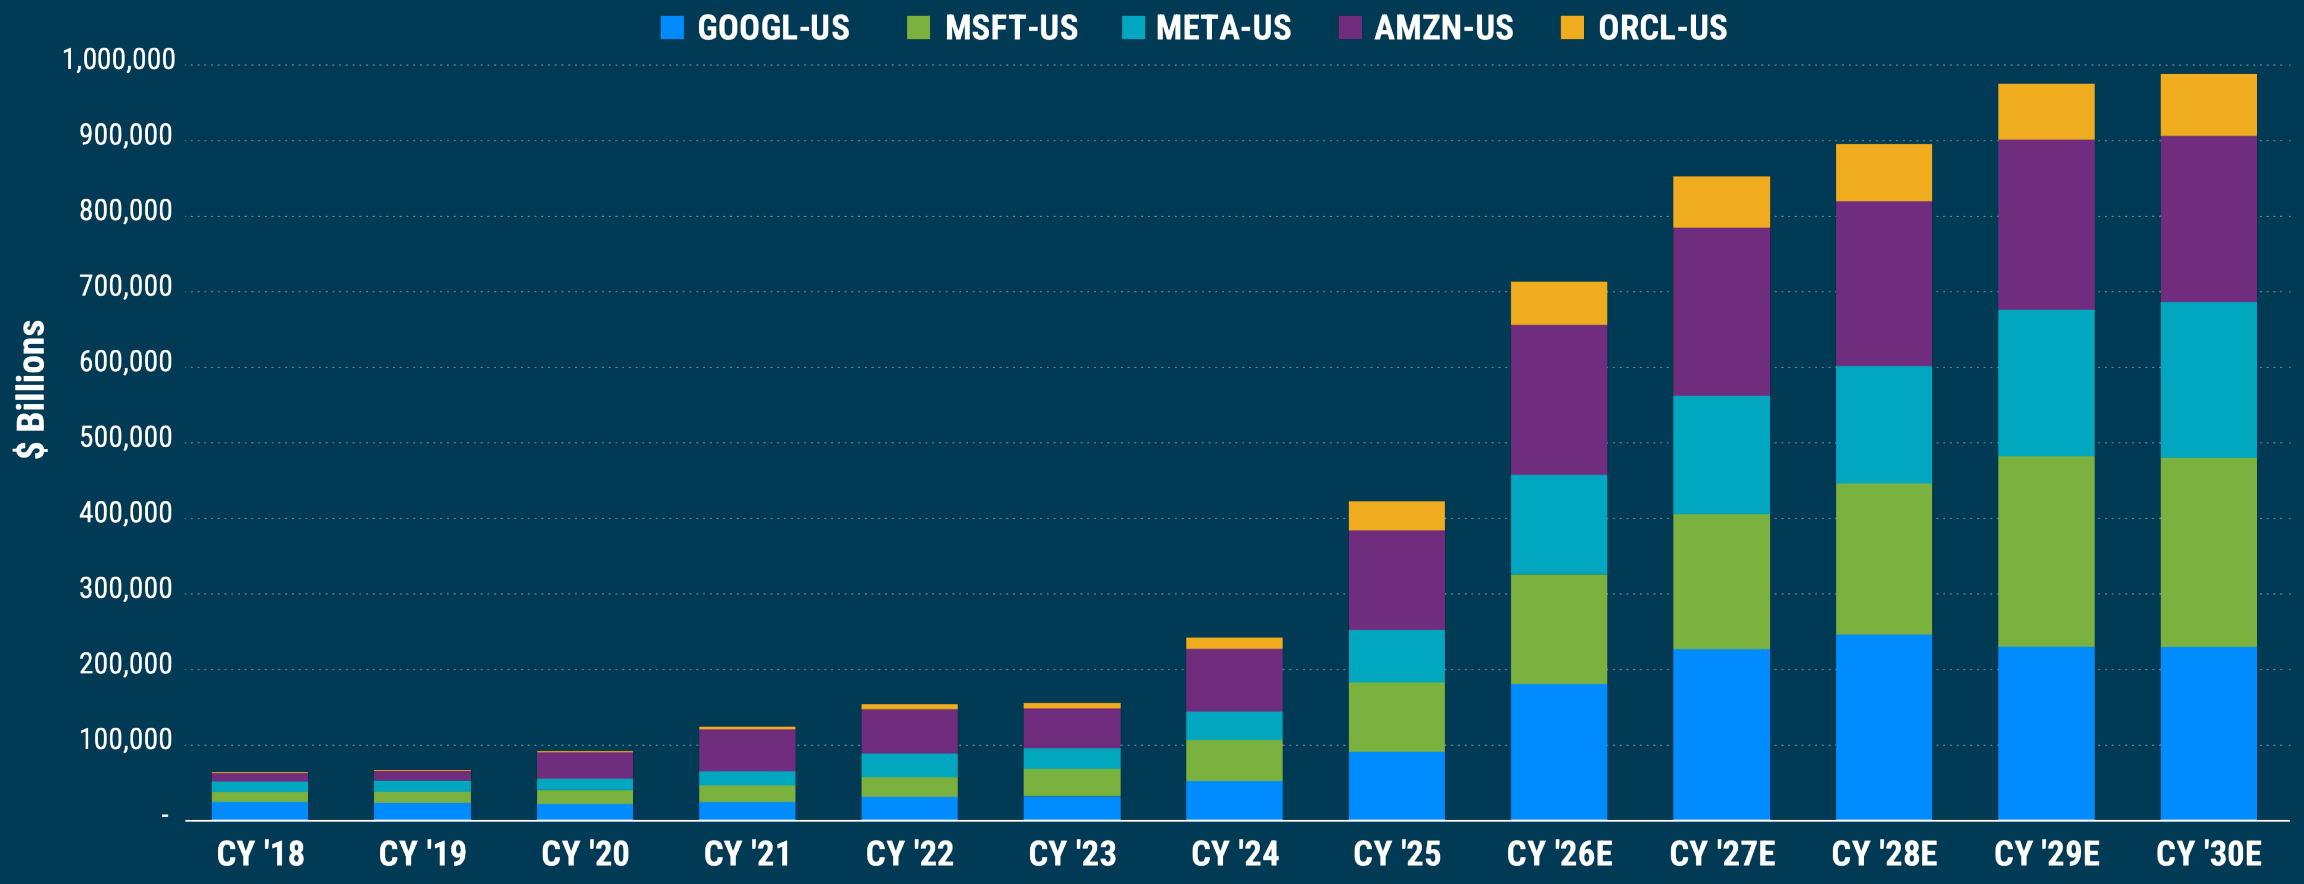

Data is as of May 15, 2026. Source: FactSet. Time period 2018-2030 is estimated spending. E: Estimated; Google (GOOGL), Microsoft (MSFT), Meta (META), Amazon (AMZN), Oracle (ORCL). Past performance is not an indication of future results. Please refer to the appendix for other important disclosures.

Hyperscalers and S&P 500 CapEx % (excluding financials)

Closing in on 50% of Total S&P 500 Spend by 2028

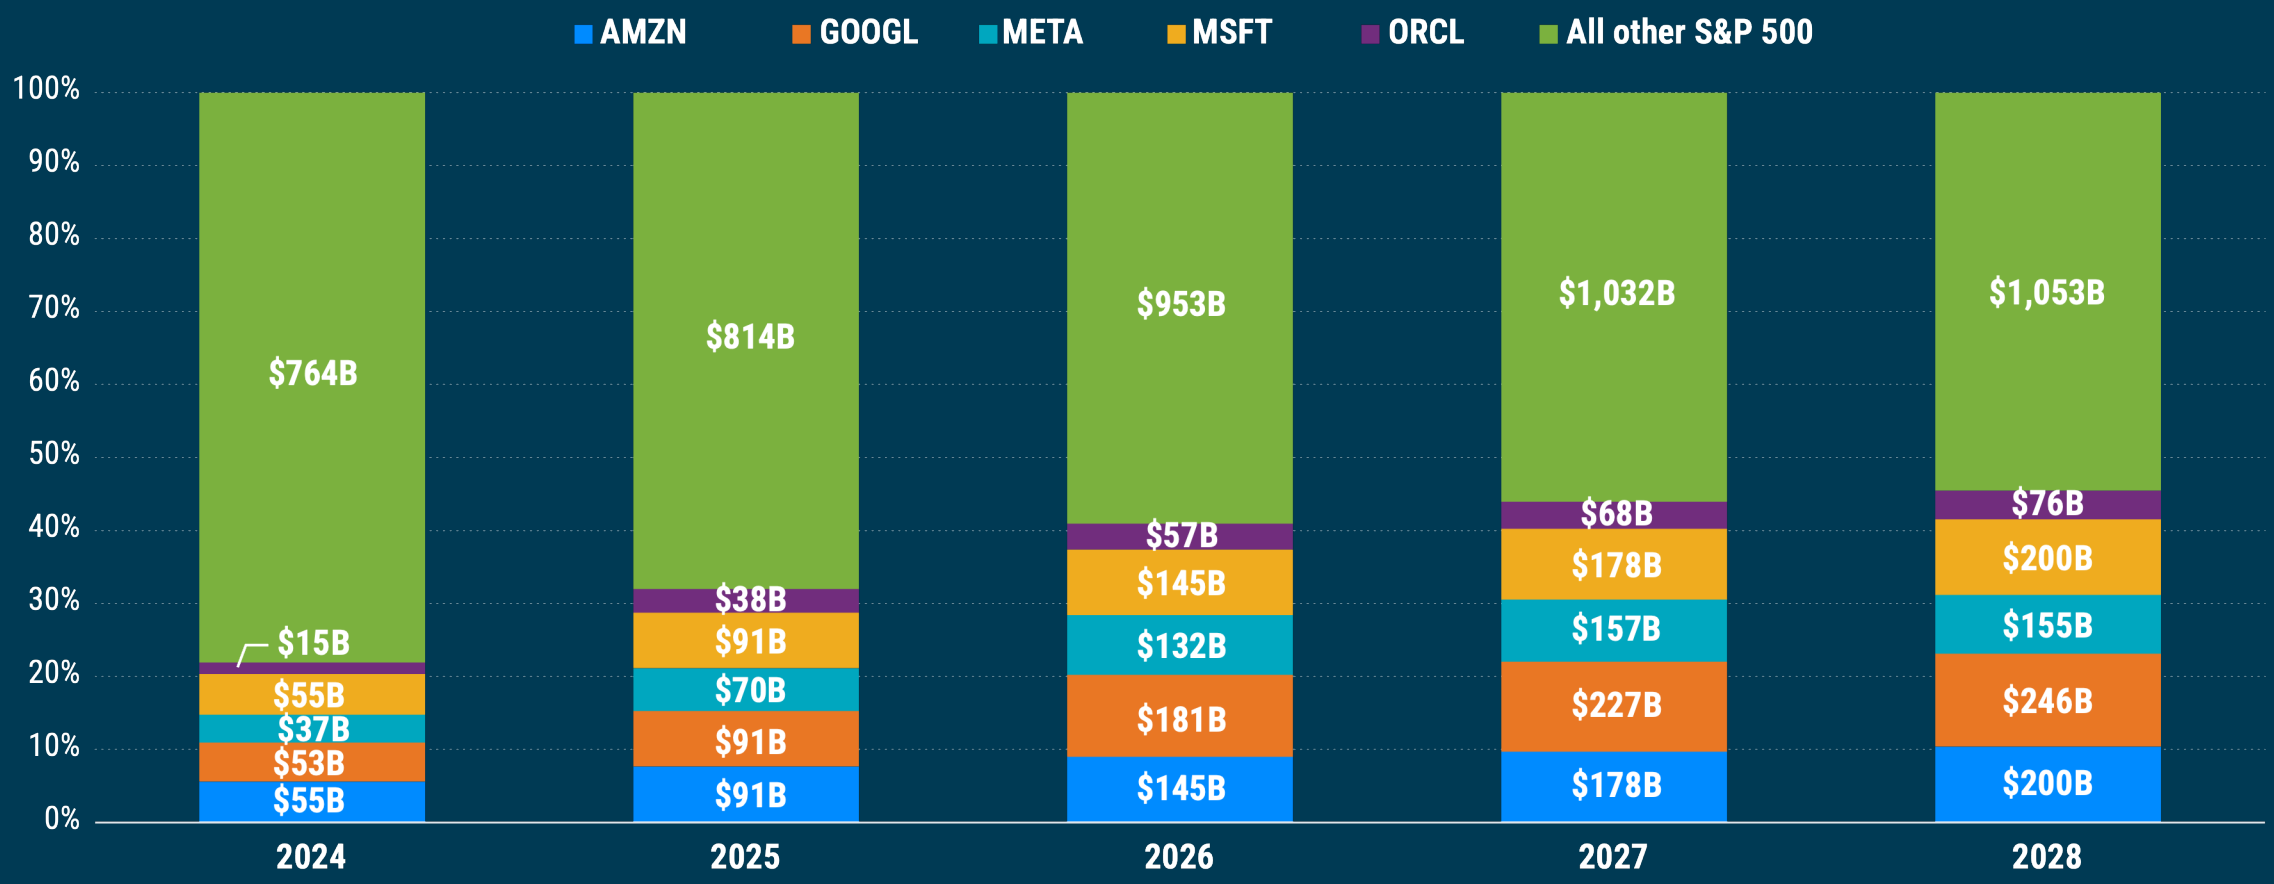

Data as of May 15, 2026. Source: FactSet. Amazon (AMZN), Alphabet (GOOGL), Meta Platforms (META), Microsoft (MSFT), Oracle (ORCL).
Past performance is not an indication of future results. Please refer to the appendix for other important disclosures.

Historical Five-Year Capital Spend

Hyperscalers' relative spend historically high

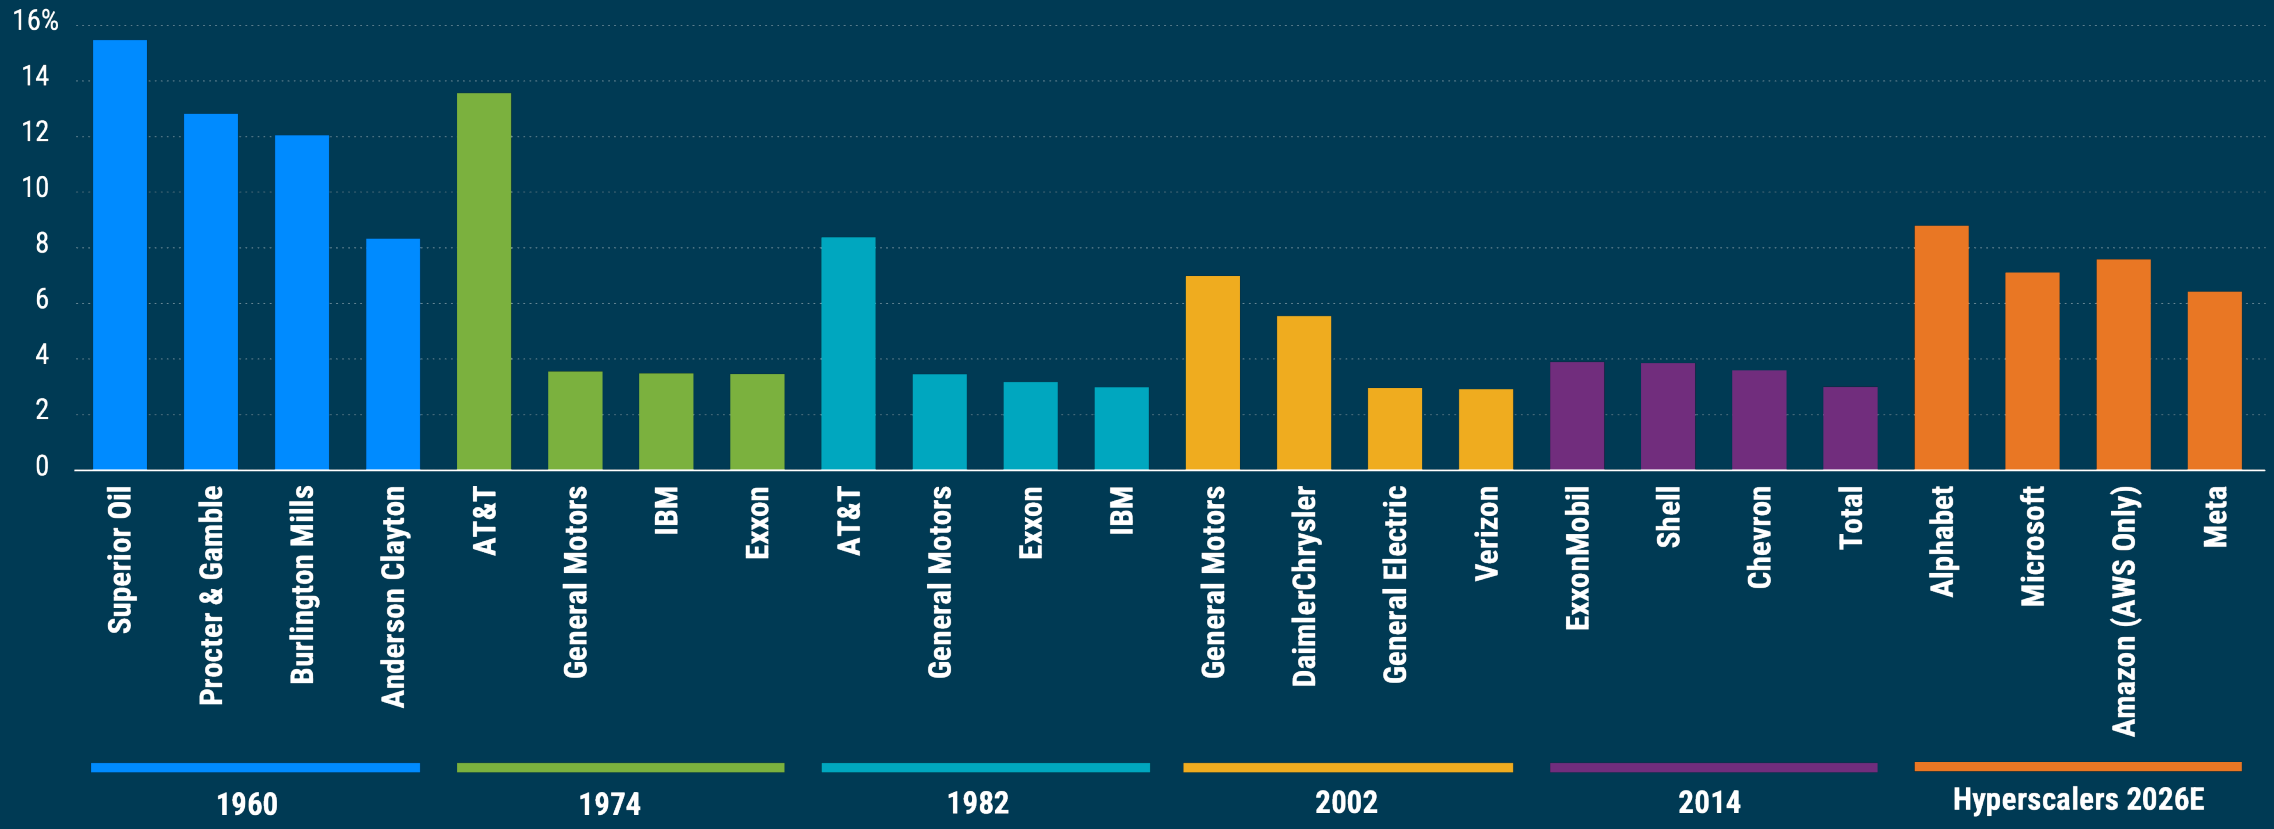

Data as of May 12, 2026. Source: Empirical Research Partners LLC.
Peaks in Capital Spending Concentration as of 1960, 1974, 1982, 2002, 2014 and 2026E (estimated).
Past performance is not an indication of future results. Please refer to the appendix for other important disclosures.

Costs to Build & Deploy are High

Estimates indicate $50 – $77 Billion

1 GW data center — NVL72 GB300 hardware cost breakdown

Blackwell Ultra NVL72 GB300 rack: $6.0M-$6.5M per rack | Sources: Tom's Hardware, Morgan Stanley, Bernstein Research, JLL (2026)

RACK HARDWARE SCENARIOS

SCENARIO	RACKS	PRICE/RACK	HARDWARE COST
Low estimate	5,500	$6.0M	~$33B
Mid estimate	6,000	$6.25M	~$38B
High estimate	7,000	$6.5M	~$46B
HARDWARE SUBTOTAL			**~$33B - $46B**

FULL COST SUMMARY

COST COMPONENT	ESTIMATE
Physical Infrastructure Shell, power, advanced liquid cooling	$15B - $25B
NVL72 GB300 rack hardware ~6,000 racks	$33B - $46B
Networking, storage, ancillary Switching fabric, NVMe, software	$3B - $6B
TOTAL CAPEX	**~$50B - $77B**

Key assumptions: ~120-140 kW IT load per rack, PUE ~1.2-1.3, ~770-830 MW usable IT capacity from 1 GW grid draw. **Benchmarks:** Nvidia Q2 FY26 earnings guidance $50B-$60B per GW; Bernstein modeled prior-gen GB200 at ~$35B per GW. GPUs alone represent ~39% of total project CapEx at this scale.

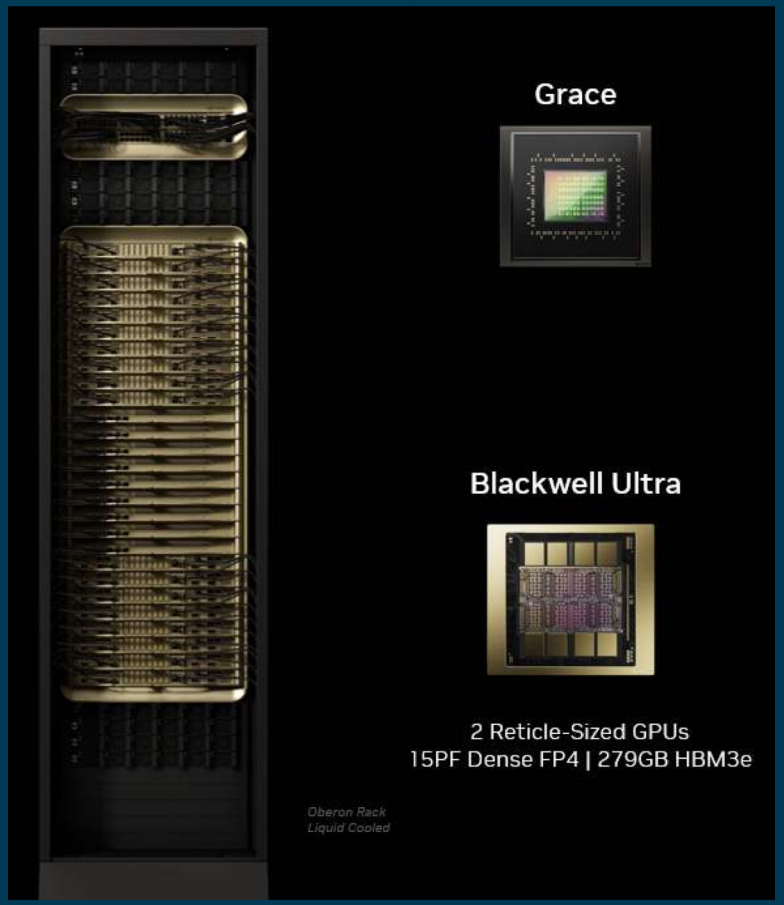

Data as of May 15, 2026. Source: Anthropic Claude. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Anthropic Claude). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

Costs, P&L Estimates of Building and Operating a Data Center

High revenue is required for an adequate ROE

1GW NVL72 GB300 Data Center @ 20% ROE	Cost	Useful Life (YRS)	D&A Op-X
GPUs	29	5	5.8
CPUs	1	5	0.2
Other Server Storage	3	5	0.6
Compute and Storage	33		6.6
Networking	4	5	0.8
Building, Power, Cooling	13	20	0.7
Total	**50**		8.1
Operating Costs			
Electricity ($0.15/Kwh)			1.3
Maintenance, Software & Other			1.2
Annual Cost/GW			10.6
Interest (50% Debt Financed @ 7%)			1.8
			12.3
20% ROE on 50% of Total Capital Investment	**5.0**		

INCOME STATEMENT	
Revenue	18.6
D&A	8.1
Operating Costs (electracy/maintenance)	2.5
EBIT	8.1
Interest	1.8
EBT	6.3
Tax @ 21%	1.3
Net Income	**5.0**
Net Margin (net Income/revenue)	27%
Revenue as a % of Total Investment ($50b/$18.6b)	**37%**
Operating Margin (EBIT/Revenue)	**43%**
Total Capital Investment/EBIT	16%
ROE	20%

1GW NVL72 GB300 Data Center @ 10% ROE	Cost	Useful Life	D&A Op-X
GPUs	29	5	5.8
CPUs	1	5	0.2
Other Server Storage	3	5	0.6
Compute and Storage	33		6.6
Networking	4	5	0.8
Building, Power, Cooling	13	20	0.7
Total	**50**		8.1
Operating Costs			
Electricity ($0.15/Kwh)			1.3
Maintenance, Software & Other			1.2
Annual Cost/GW			10.6
Interest (50% Debt Financed @ 7%)			1.8
			12.3
10% ROE on 50% of Total Capital Investment	**2.5**		

INCOME STATEMENT	
Revenue	15.5
D&A	8.1
Operating Costs (electracy/maintenance)	2.5
EBIT	5.0
Interest	1.8
EBT	3.2
Tax @ 21%	0.7
Net Income	**2.5**
Net Margin (net Income/revenue)	16%
Revenue as a % of Total Investment ($50b/$15.5b)	**31%**
Operating Margin	**32%**
Total Capital Investment/EBIT	10%
ROE	10%

Could leverage be higher, 4:1 debt to equity?

Blackstone Announces Joint Venture with Google to Create New TPU Cloud

- **$5 billion equity commitment**
- **500 MW capacity coming online 2027**
- **Potential value: $25 billion**
- **1 GW capacity could therefore be worth $50 billion**

Revenue of 31%-37% of Total Investment required for 10% and 20% ROE (50/50 debt/equity)

Data as of May 18, 2026. Source: Blackstone and Bloomberg.

Does the Opportunity justify Trillions in AI Spend?

What do we need to believe?

		20% ROE	10% ROE
	Revenue as a % of Capital Investment	**37%**	**31%**
	Estimated AI Capital Spending 2026-2030	4,400	4,400
✓	**Required Revenue**	**1,637**	**1,364**
	US population 2030 (est)	350	350
	European Union population 2030 (est)	460	460
	Total US & EU Population 2030 (est)	810	810
✓	**Revenue Required Per Capita/annum**	**$2,021**	**$1,684**
✓	**Revenue Required Per Capita/month**	**$168**	**$140**
	US GDP 2030 (est)	37,000	37,000
	European Union GDP 2030 (est)	28,000	28,000
	US + EU GDP 2030 (est)	65,000	65,000
✓	**AI Required Revenue % GDP**	**2.5%**	**2.1%**
	2026 Hyperscaler est Revenue	1,332	1,332
✓	**Required Revenue % of 2026 Revenue**	**123%**	**102%**
	Verizon 2026 est Revenue	117	117
	AT&T 2026 est Revenue	130	130
	T-Mobile 2026 est Revenue	95	95
	Verizon, AT&T and T-Mobile est 2026 Revenue	341	341
✓	**Required Revenue % of US 2026 Telco Revenue**	**480%**	**400%**
	Average annual salary of Software and Web Developers, Programmers, and Testers	136	136
	# of Employees: Required Revenue/Average Salary	12,035	10,029

Collette Kress, CFO Nvidia, 1Q27 Call, May 2026

"With analysts now forecasting Hyperscale CapEx to exceed $1 trillion in 2027 and agentic AI beginning to proliferate all industries, **AI infrastructure spending is on track to reach $3 trillion to $4 trillion annually by the end of this decade.**"

Can revenue support $3-$4 trillion of annual AI infrastructure Spend?

THIS IS JUST THE "FACTORY" LAYER

Data as of May 20, 2026. Source: FactSet; Google.

Hyperscalers Control AI Compute

Over 60% controlled by 5 companies

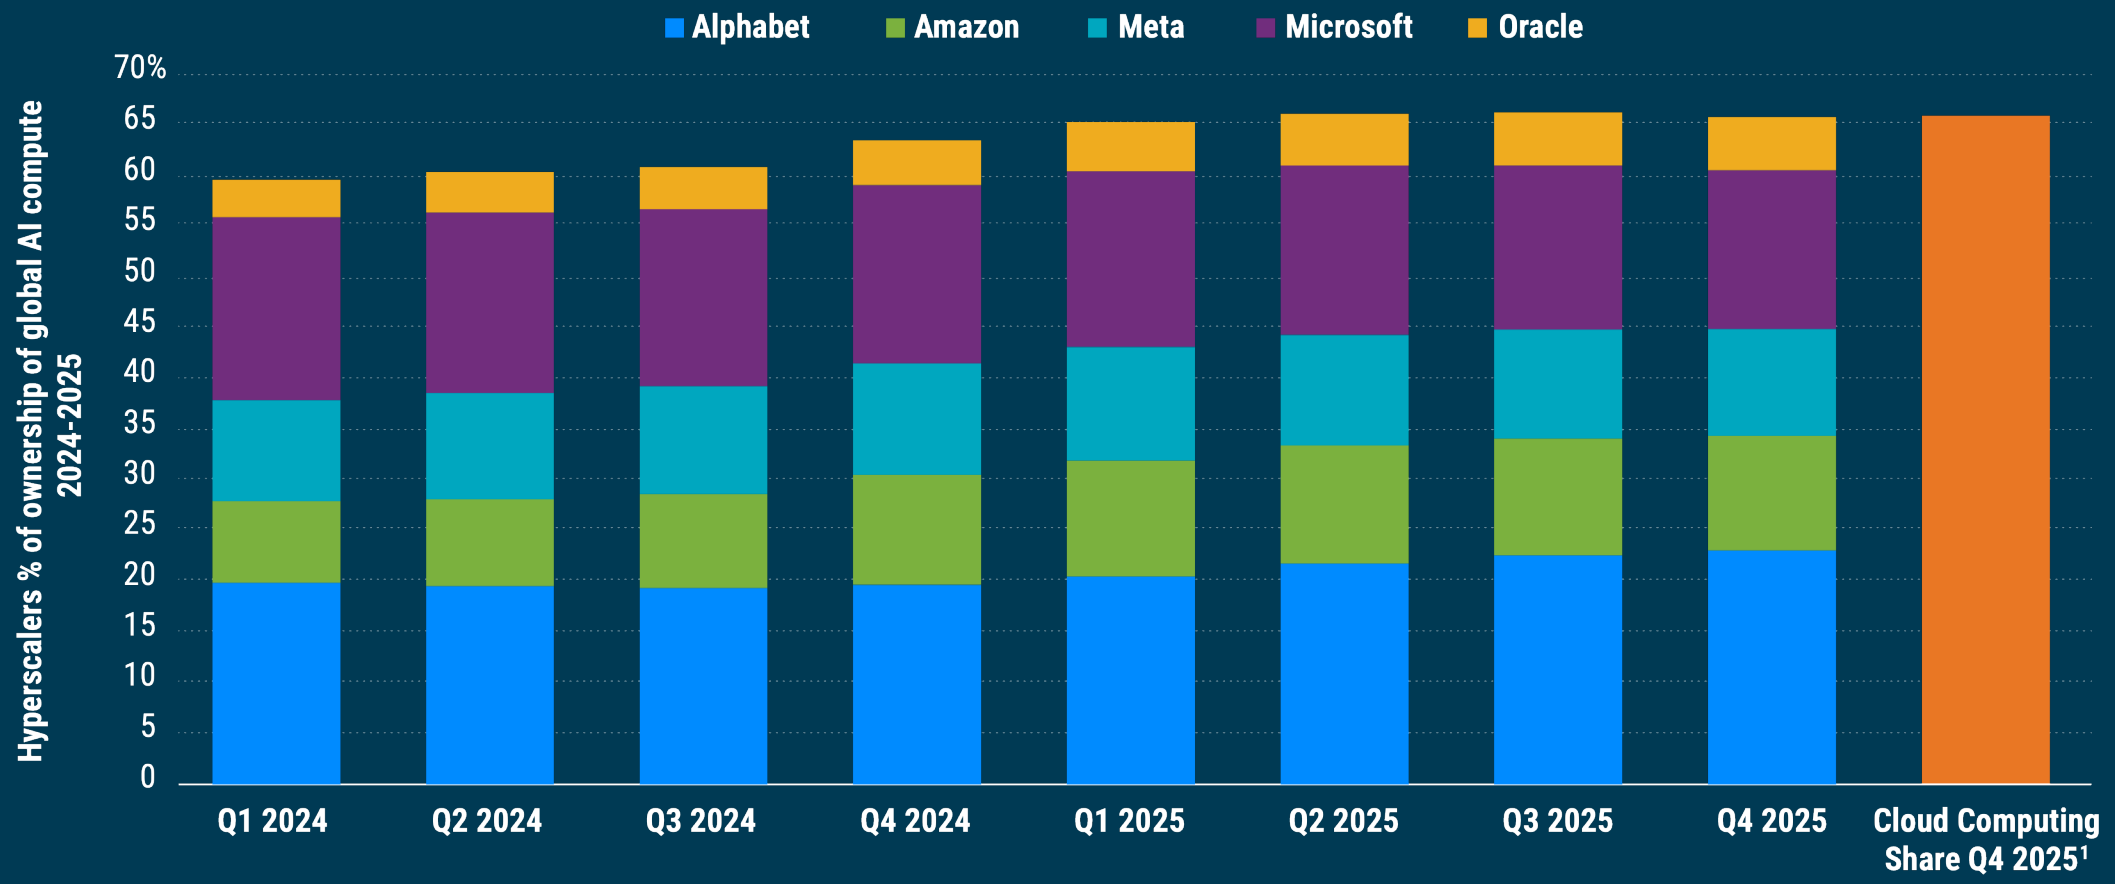

Data as of April 28, 2026. Source: EpochAI; Synergy Research Group; Empirical Research Partners LLC.
[1] Based on Amazon, Alphabet, Microsoft and Oracle; Meta does not offer cloud computing services.
Past performance is not an indication of future results. Please refer to the appendix for other important disclosures.

Risk of AI Compute Overbuild?

The dark fiber equivalent of the late 1990's?

The Rise And Fall Of The Fiber Giants

COMPANY	ROLE / INFRASTRUCTURE BUILT	FATE / CURRENT STATUS
Global Crossing	Built first privately financed undersea global fiber network (100k+ miles).	Declared bankruptcy in 2002 ($12.4B debt).
WorldCom	A major long-distance carrier dominating internet backbone.	Filed for bankruptcy in 2002 amid a record $11B accounting fraud.
Level 3 Communications	Built a massive national and international long-haul fiber network.	**Survivor.** Consolidated failed rivals.
Metromedia Fiber Network (MFN)	Focused high-capacity fiber within major cities	Filed for bankruptcy in 2002.
360networks	Built a transcontinental network and undersea cables.	Filed for bankruptcy in 2001.
Williams Communications (WilTel)	Laid fiber inside decommissioned natural gas pipelines.	Filed for bankruptcy in 2002.
Qwest Communications	Laid fiber along railroad rights-of-way (born from Southern Pacific Railroad).	**Survivor,** but faced accounting scandals. Acquired by CenturyLink **(Lumen)** in 2011.

Data as of May 15, 2026. Source: Google Gemini. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Google Gemini). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

Risk of AI Compute Overbuild?

Corning - will history rhyme?

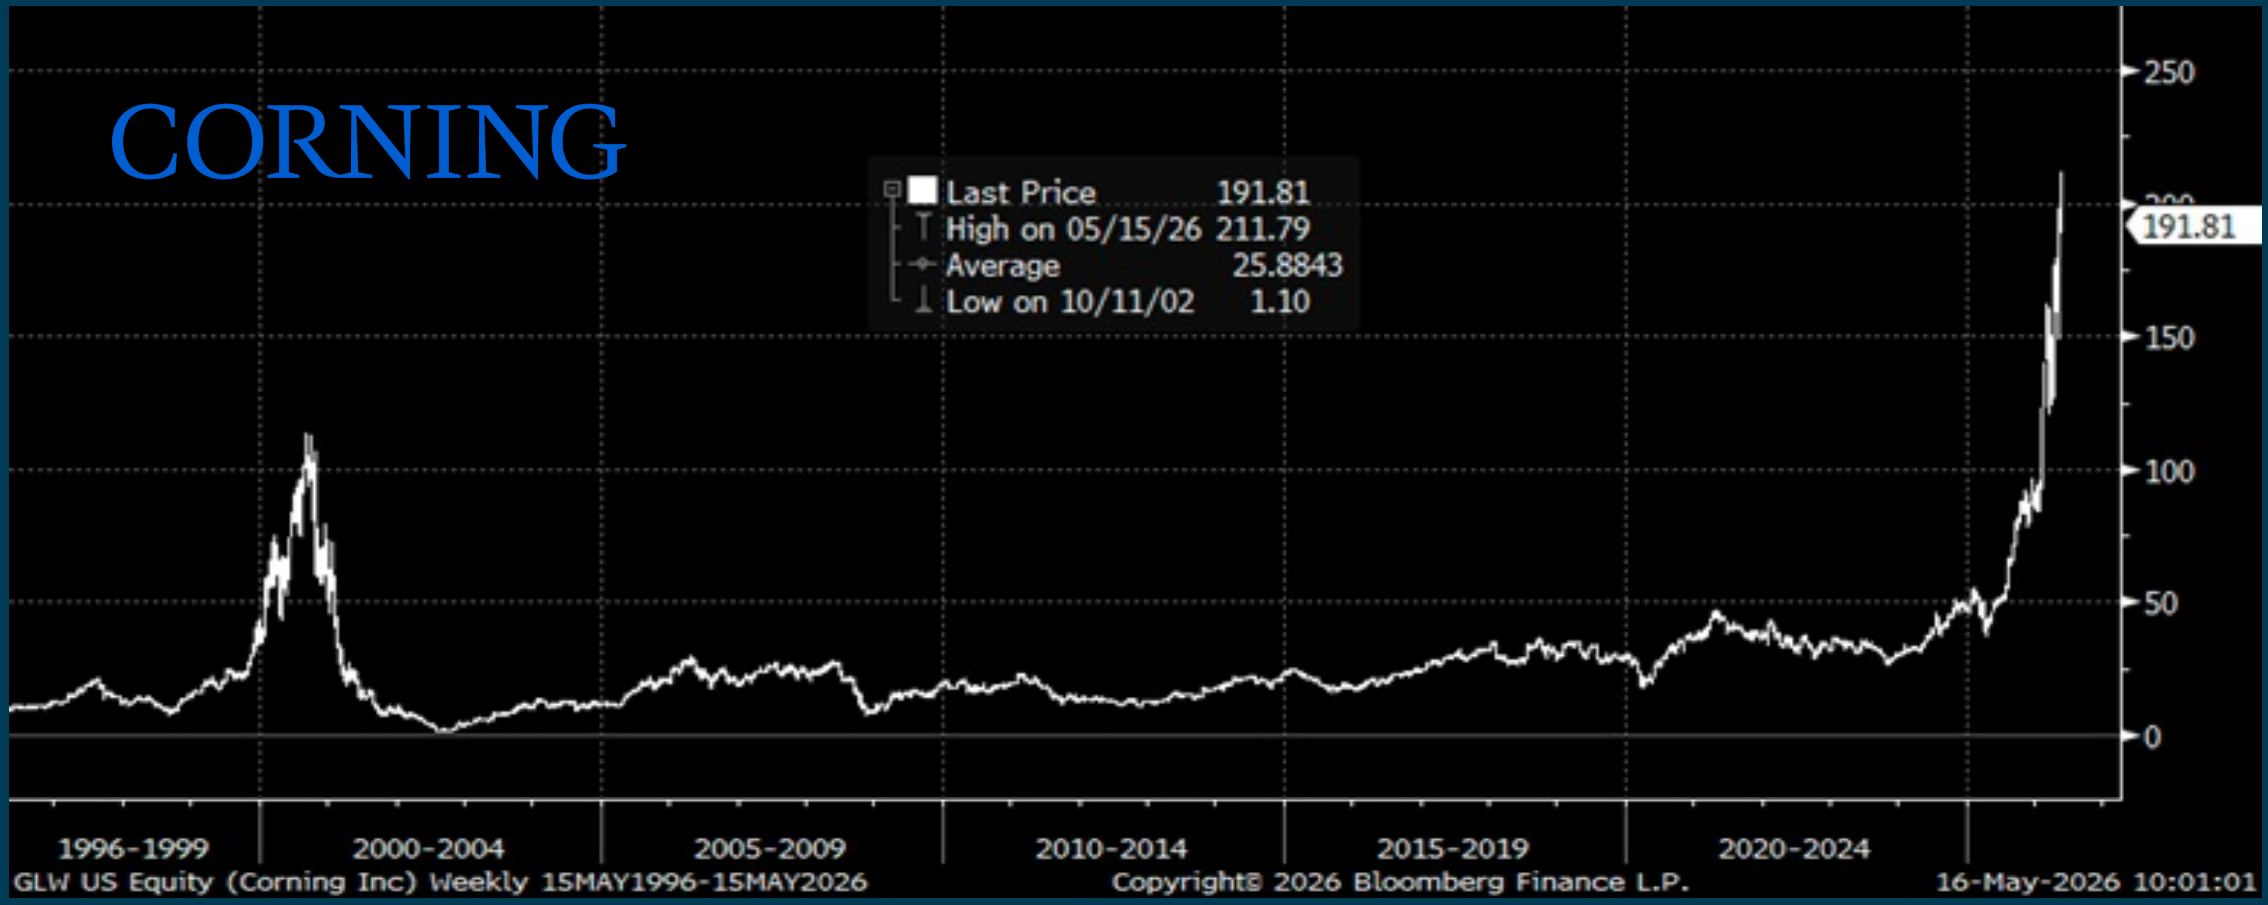

Data as of May 15, 2026. Source: Bloomberg.

US Wireless: Subscriber Connections & Annual Service Revenue

1995-2024

U.S. wireless subscriber connections, 1995-2024 (Millions)

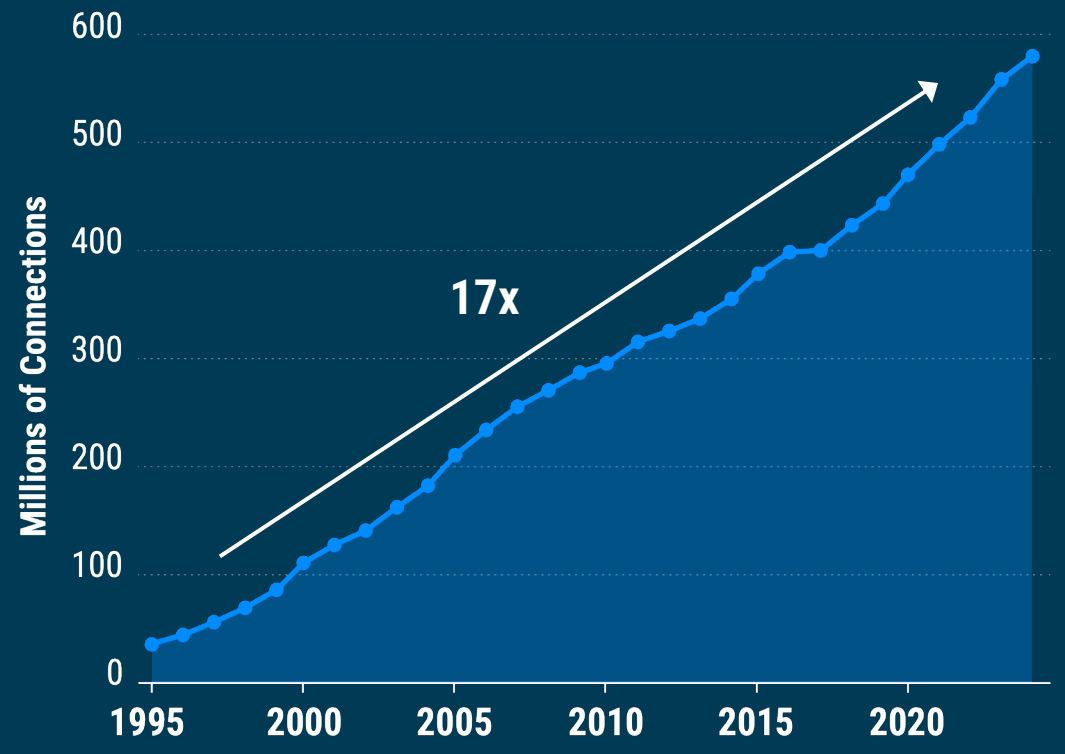

U.S. wireless annual service revenue, 1995-2024 (USD, nominal)

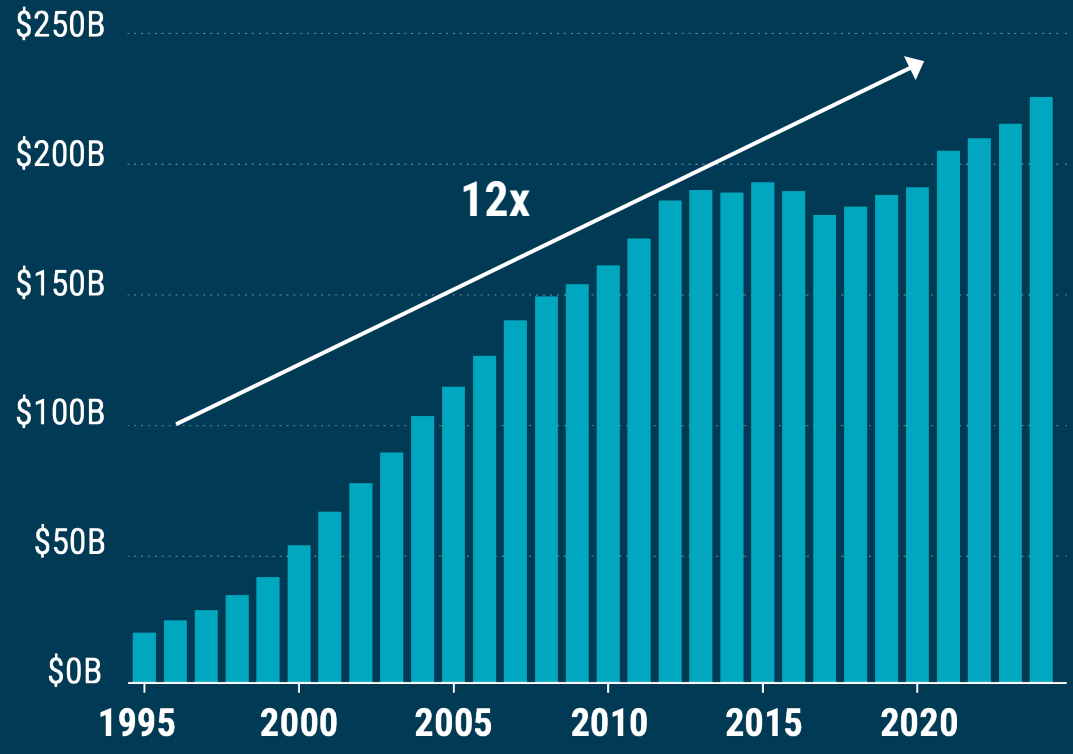

Data as of December 2024. Source: CTIA Annualized Wireless Industry Survey Results, year-end 1985-2024 (ctia Wireless Industry Indices Report).

US Wireless: Monthly ARPU & Cumulative Capital Investment

1995-2024

U.S. wireless monthly ARPU, 1995-2024 (USD, nominal)

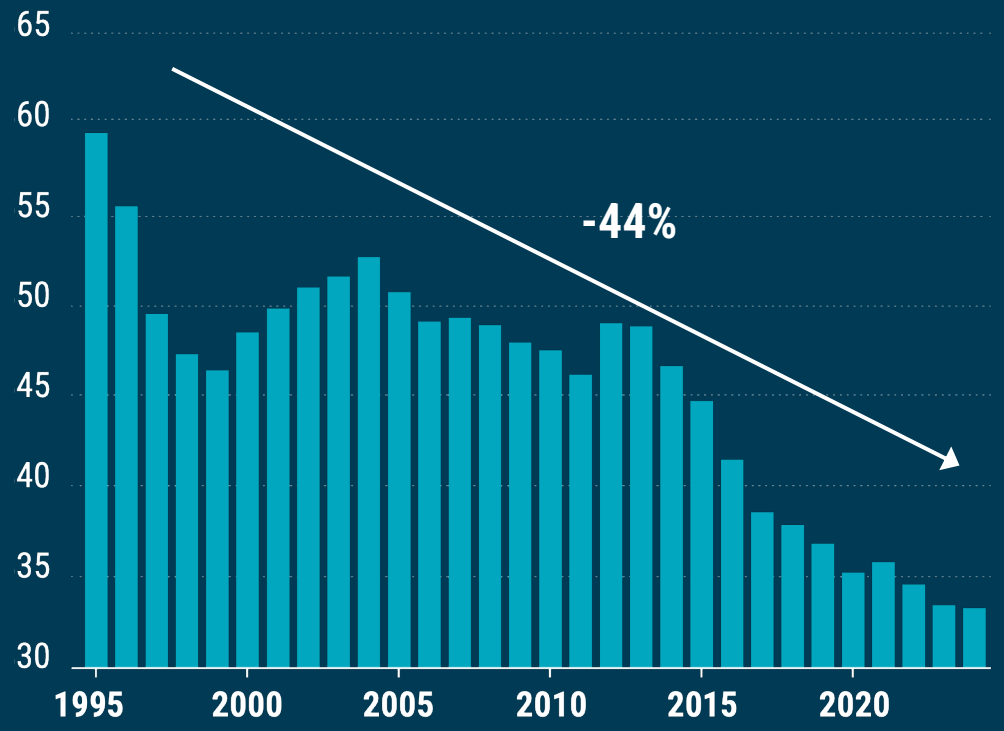

U.S. wireless capital investment, 1995-2024 (USD, nominal)

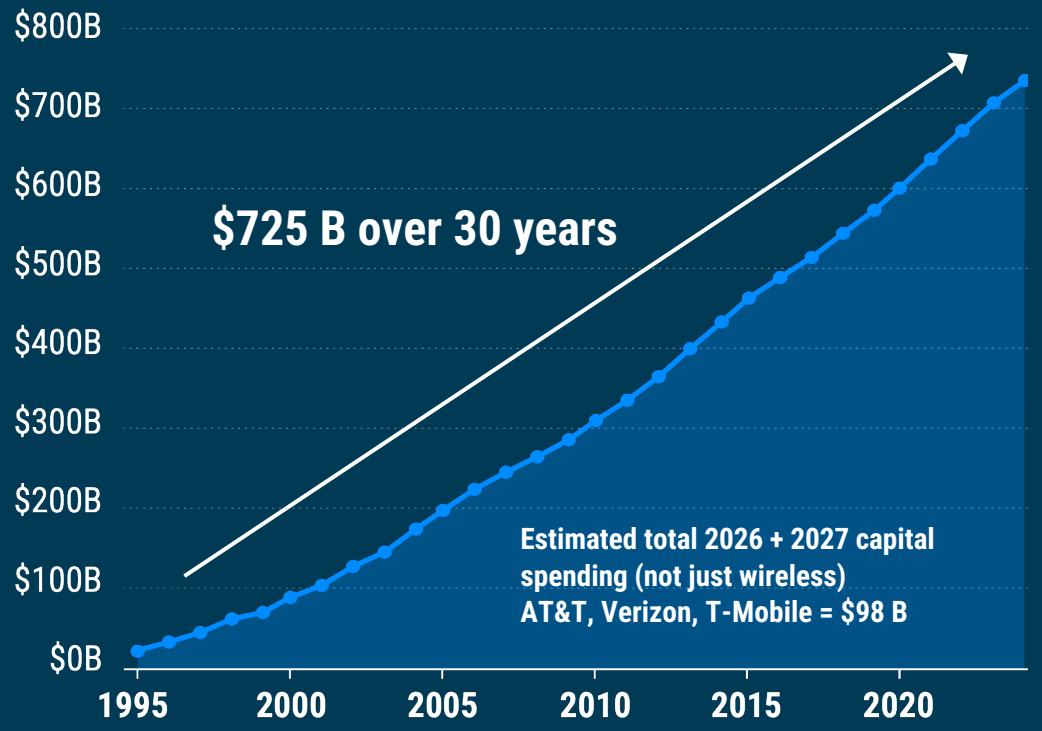

ARPU stands for **Average Revenue Per User**

Data as of December 2024. Source: CTIA Annualized Wireless Industry Survey Results, year-end 1985-2024 (ctia Wireless Industry Indices Report).

BostonPartners

U.S. Wireless: Data Usage

331x in 15-years

32T MB

increase in 2024, the largest year-over-year increase, ever

2x

Wireless data demand is on pace to double every two years.

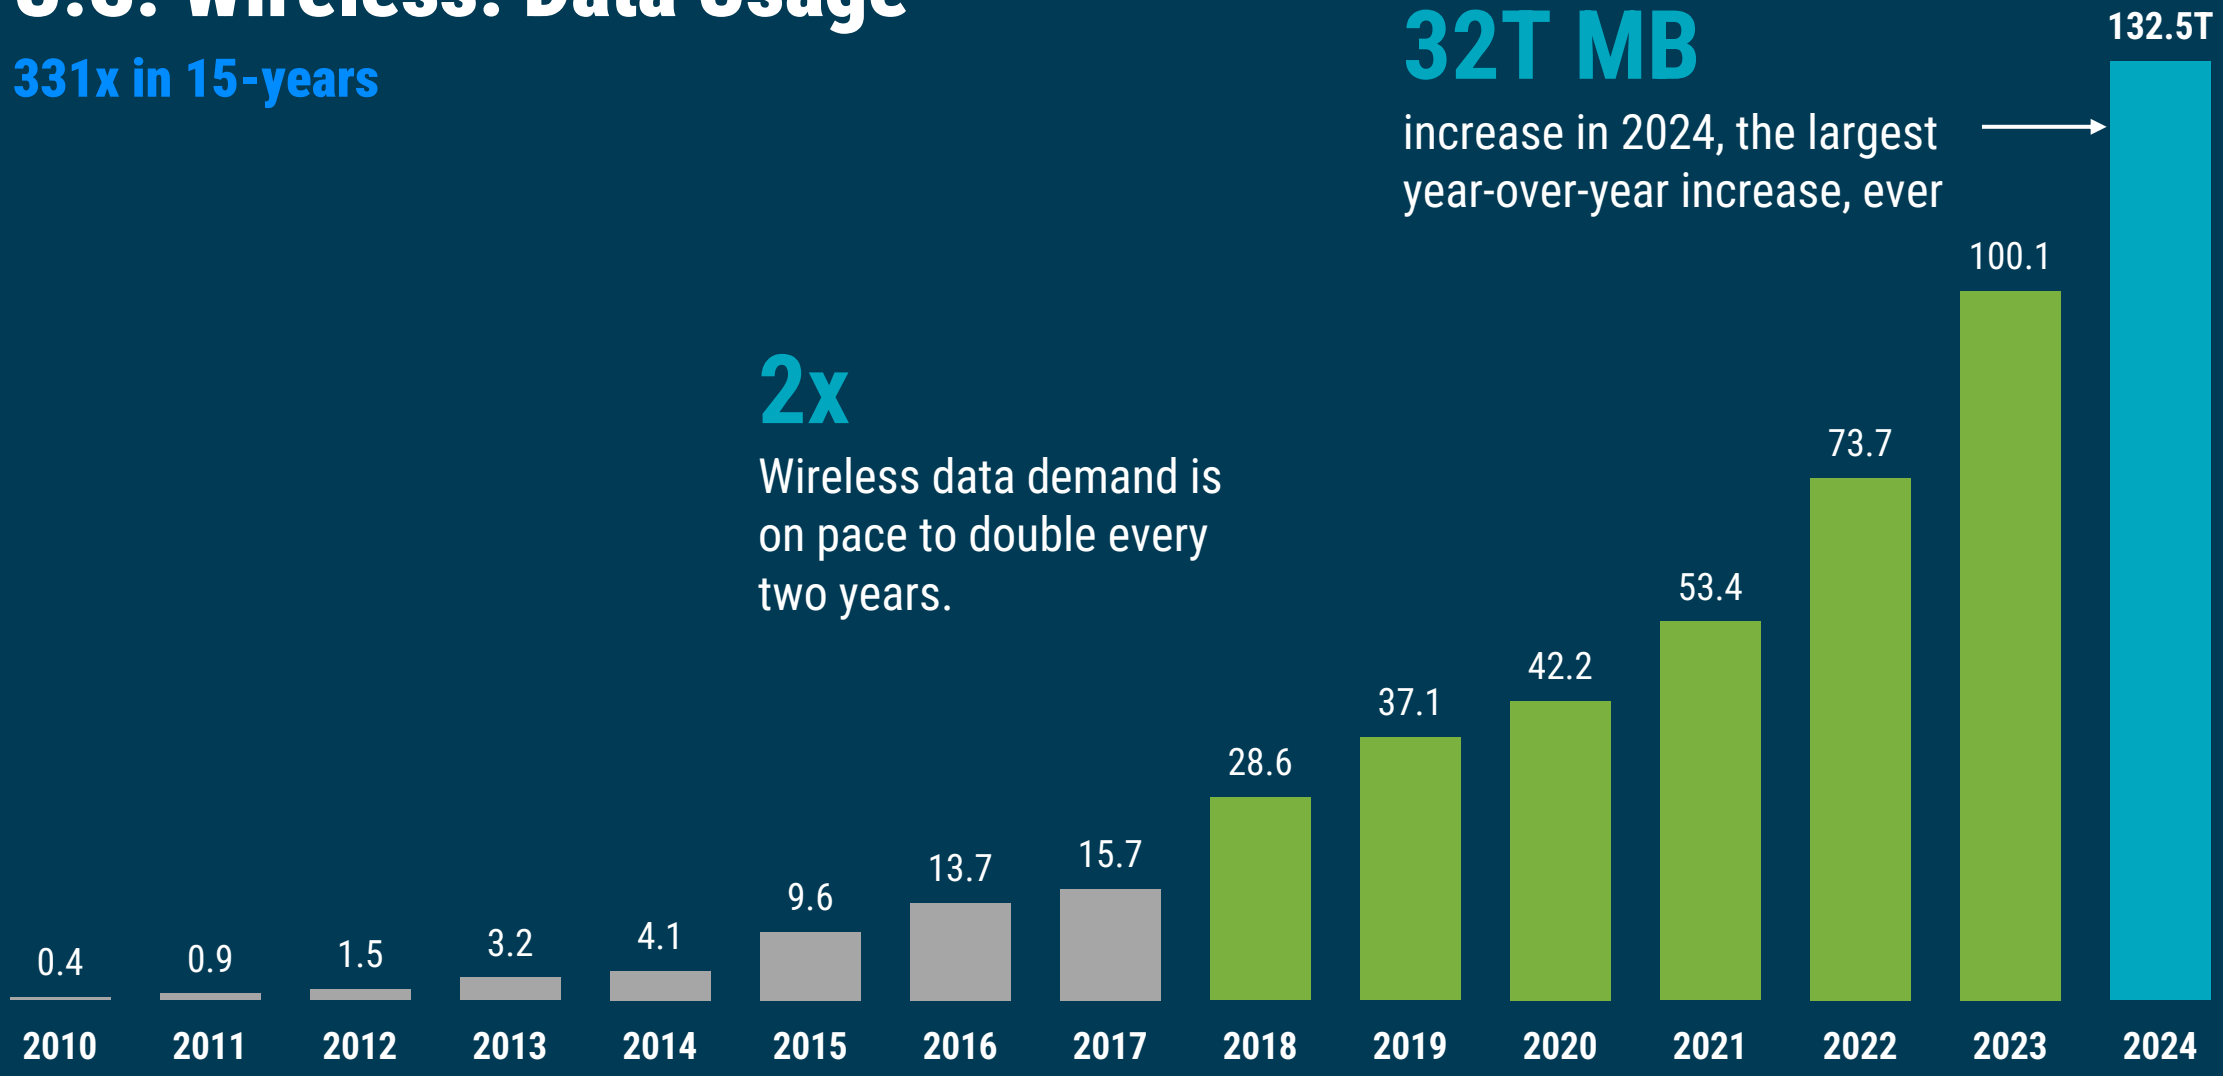

September 8, 2025. Source: ctia 2025 Annual Survey Highlights

BostonPartners

AI Token Usage

10,000x in 3-years?

Apr 2024	May 2025	Jul 2025	Q4 2025
9.7T	**480T**	**980T**	**~14T/day**
Google monthly tokens	Google monthly tokens	Google monthly tokens	Google via APIs

AI token usage: monthly volume, March 2023 - December 2025

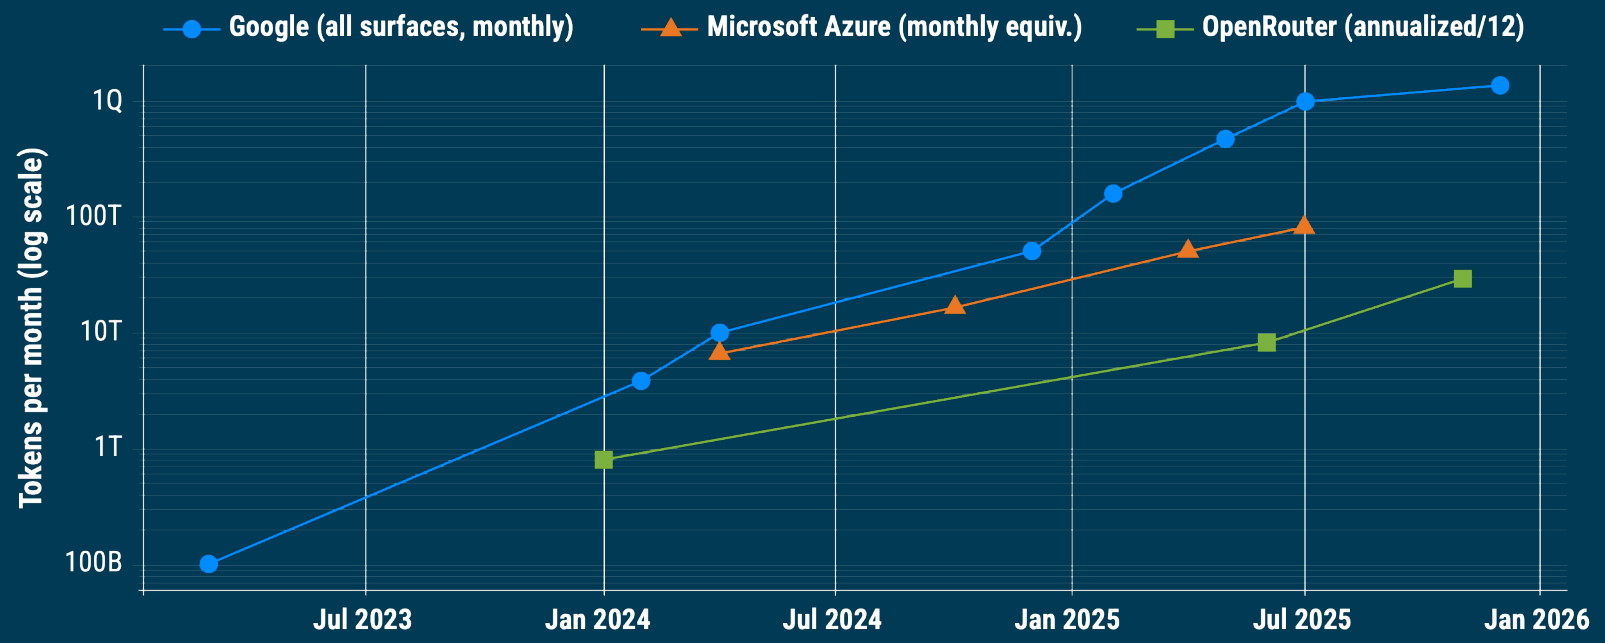

Key takeaways

- **Google Monthly Tokens up ~10,000x**
- **Average Prompt size increased 4x**
- **Completion length increased 3x**
- **>50% output tokens from programming**
- **1 quadrillion monthly tokens**

Data time period is March 2023 through December 2025. Source: Google I/O 2024 & 2025; Alphabet Q2/Q4 2025 earnings; Microsoft Q3/Q4 2025 earnings; OpenRouter State of AI (a16z, Dec. 2025); Interconnects.
Note: Y-axis is logarithmic. Series are not strictly applies-to-apples - providers count tokens across different surfaces and APIs.

Revenue has Followed

Open AI and Anthropic Annualized Recurring Revenue

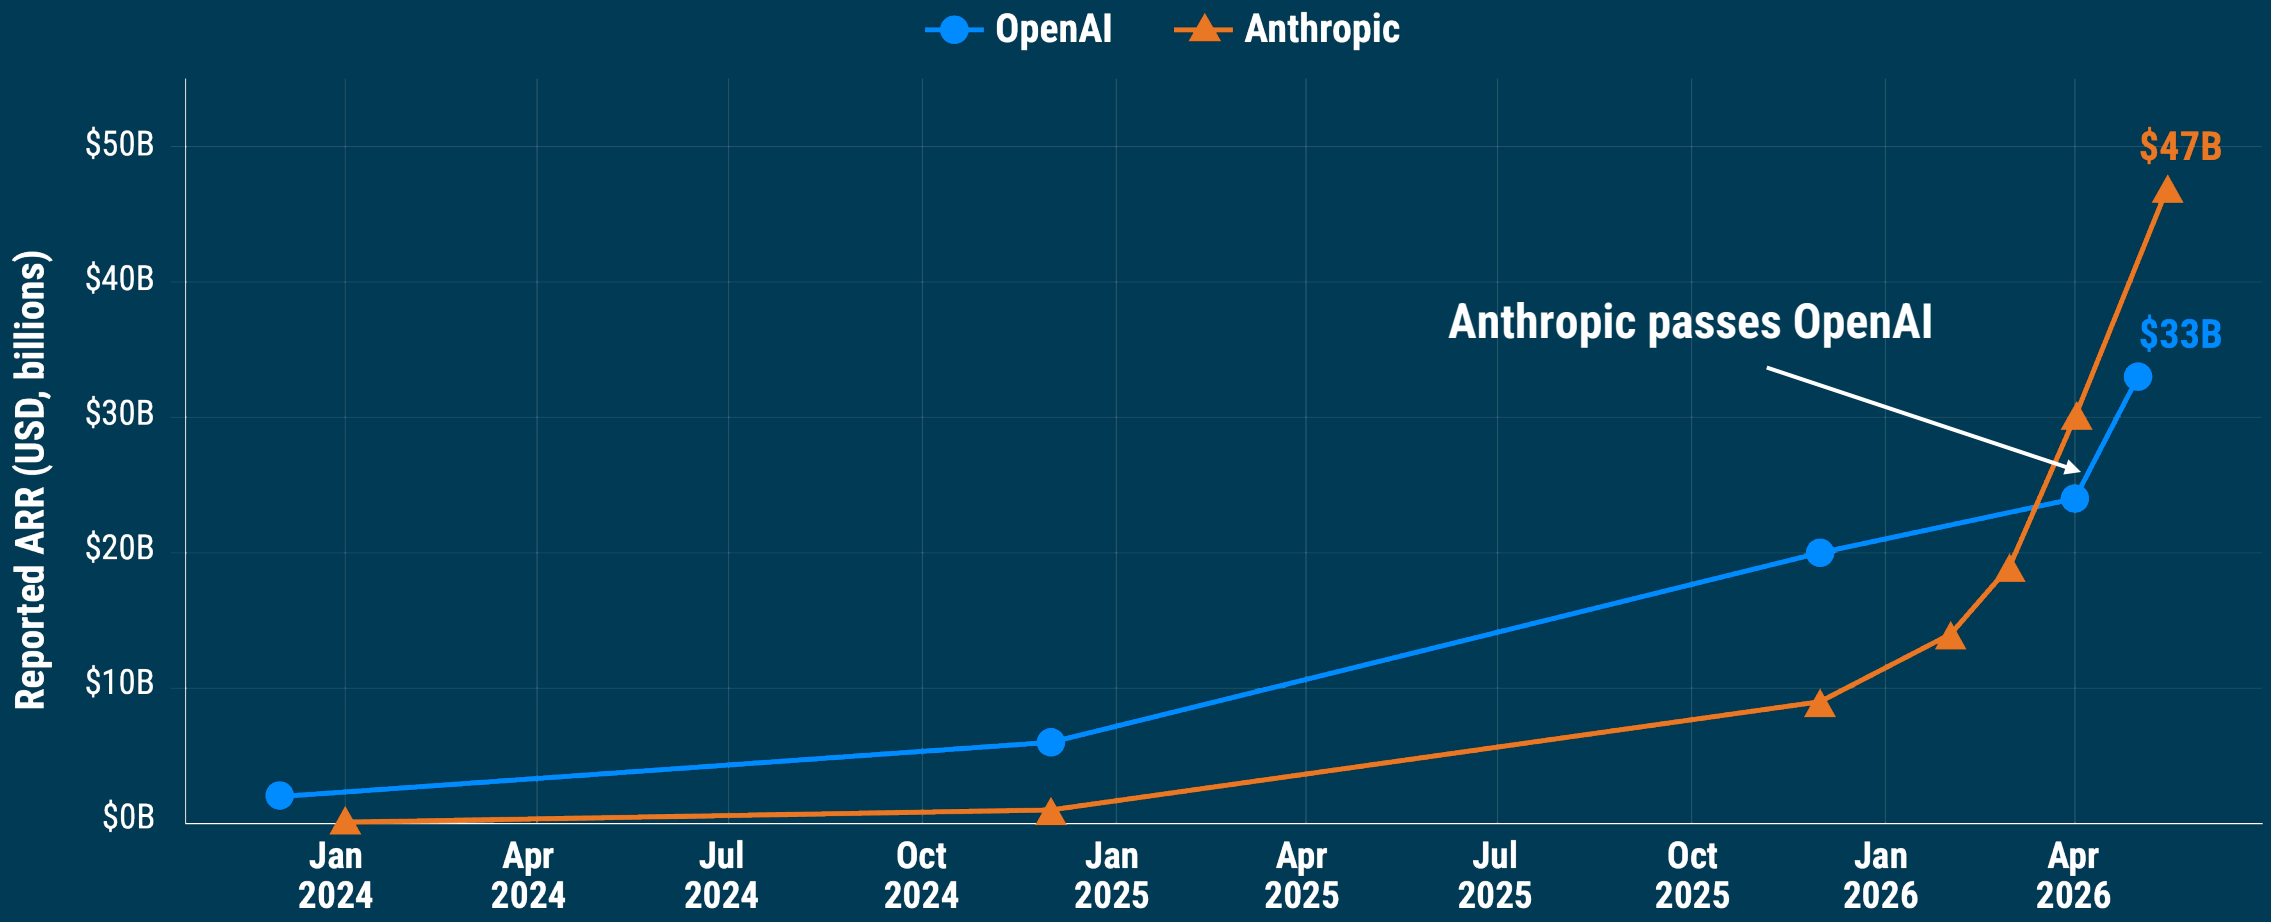

Data from November 2024 - May 2026. Source: Anthropic Claude. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Anthropic Claude). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

The Five Layer Cake

Jensen Huang Describes the Modern AI Ecosystem

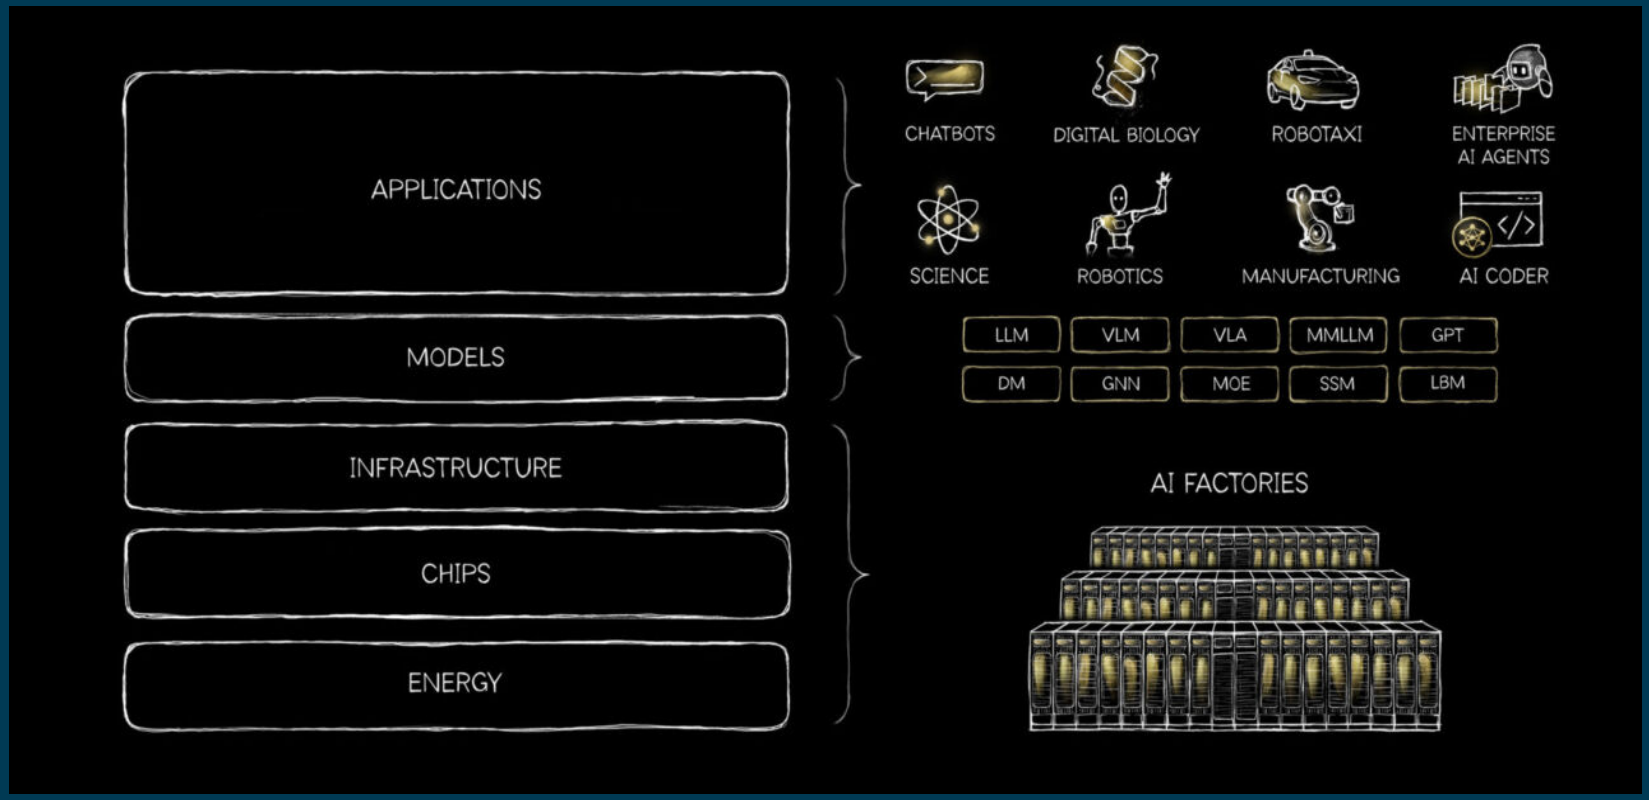

Where are the profit pools - chips, infrastructure, models, apps?

"...We have to be a product company because there's kind of a lot of reasons for lab to build products. And, you know, this actually wasn't clear early on. Like very early on in Anthropics history, this is before I joined, this was actually like an active debate. Should we even build products? Is this actually like a useful thing to do?"

Boris Cherny, Creator & Head of Claude Code @Anthropic, Big Technology Podcast (4:32)

Can there be any value at the application layer of the five-layer-cake? Or are applications simply a tool or an interface to access "models" (every kind of model, LLM, VLM, VLA, MMLLM..., etc.) and generate tokens? Is there any value in an AI coding agent if you can build your own? Is there any value in a co-working app/agent if you can build your own?

Is the application layer simply an extension of the model layer? Is an AI coding agent just a really good coding model?

Data as of March 10, 2026. Source: Nvidia.

The Five Layer Cake

Jensen Huang Describes the Modern AI Ecosystem

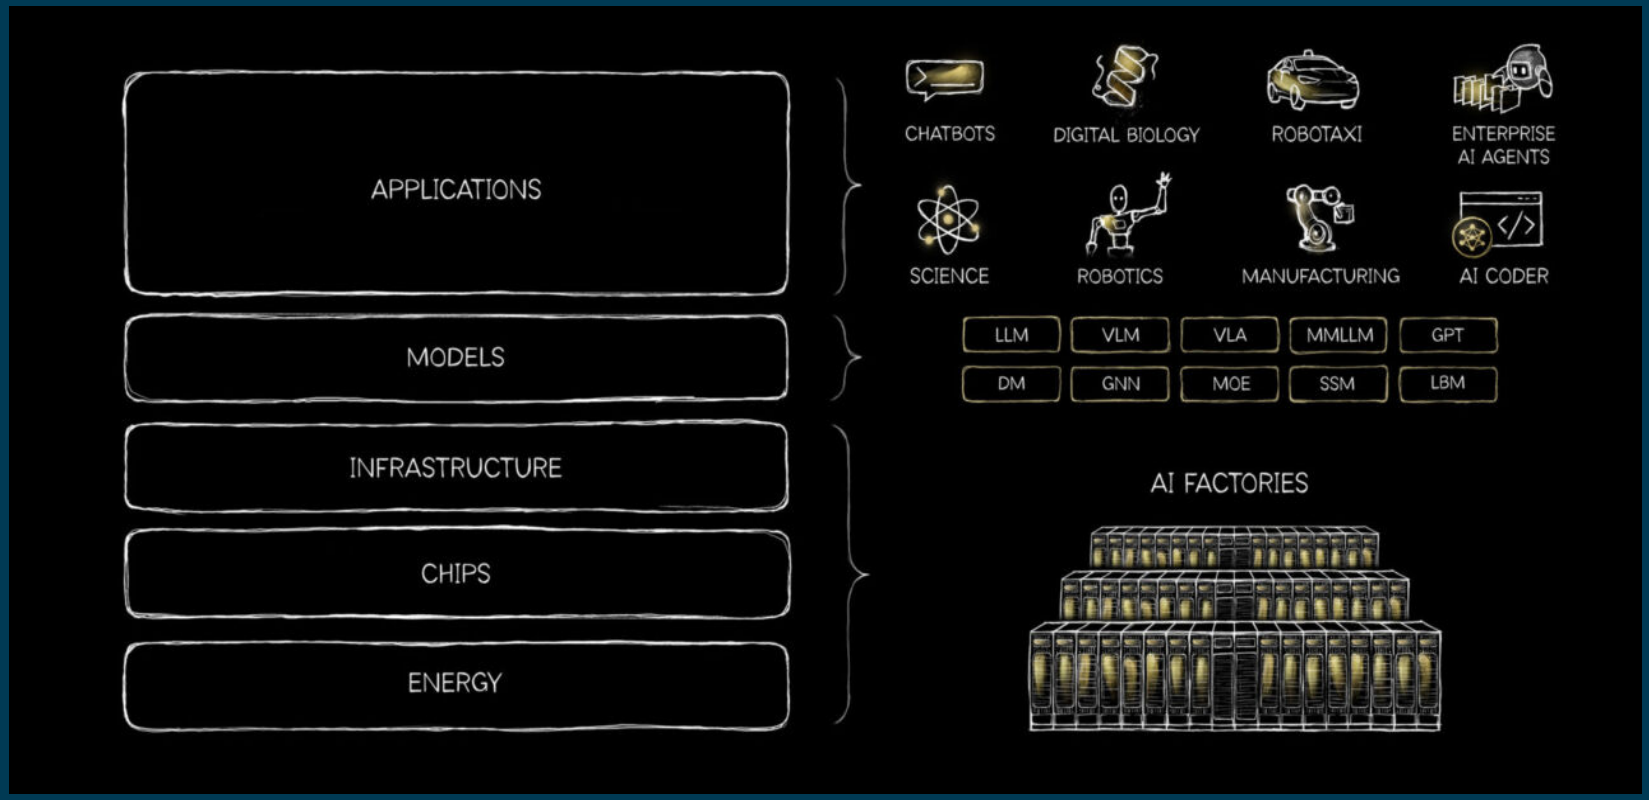

What business models will prevail? Subscription + Usage

"...there's no way to encapsulate all of that in a subscription product, right? So, I think the mental model I would have is AI is going to become a lot like Costco, where you pay for the membership and that gets you in the store. And that's actually the part of Costco's business that is, you know, the highest margin, and then you have everything you're buying in Costco, you know, you have confidence that there's like a max margin, right? And those are kind of like computer credits, right? And it's, you know, some people go to Costco and they just buy the hot dog. And then, you know, there's people who go and spend, you know, thousands of dollars every trip. And that depends on their needs."

Dmitry Shevelenko, Chief Business Officer @Perplexity, Big Technology Podcast (58:57)

Data as of March 10, 2026. Source: Nvidia.

Where are the Profit Pools in AI?

High Cost/Low Value: What needs to work?

April 2026

Open AI discontinues Sora, due to $15 million/day costs

AI use cases: cost vs. value

Cost = post-training + token intensity • Value = economic + societal

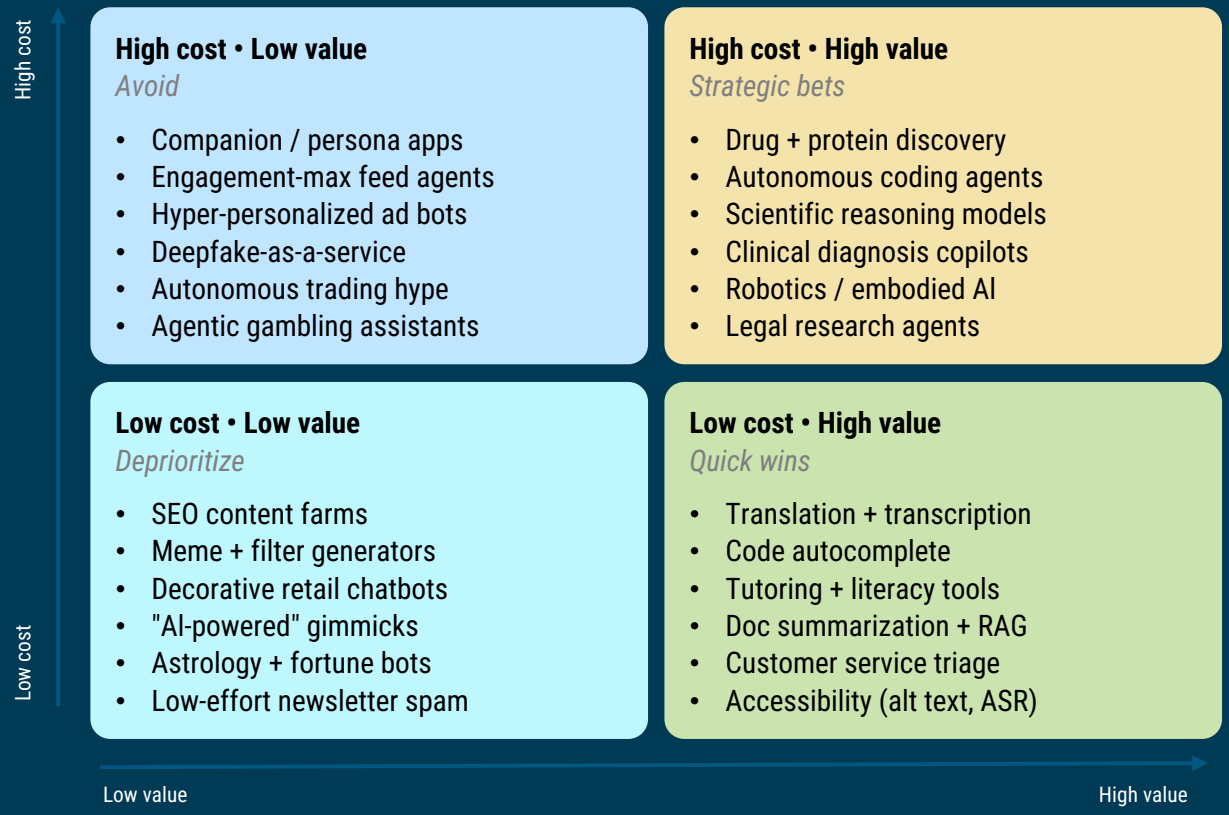

Data as of April 2026. Source: Anthropic Claude. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Anthropic Claude). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

Will the Model Market become Competitive

Hyperscalers are developing frontier models. What about open-source models?

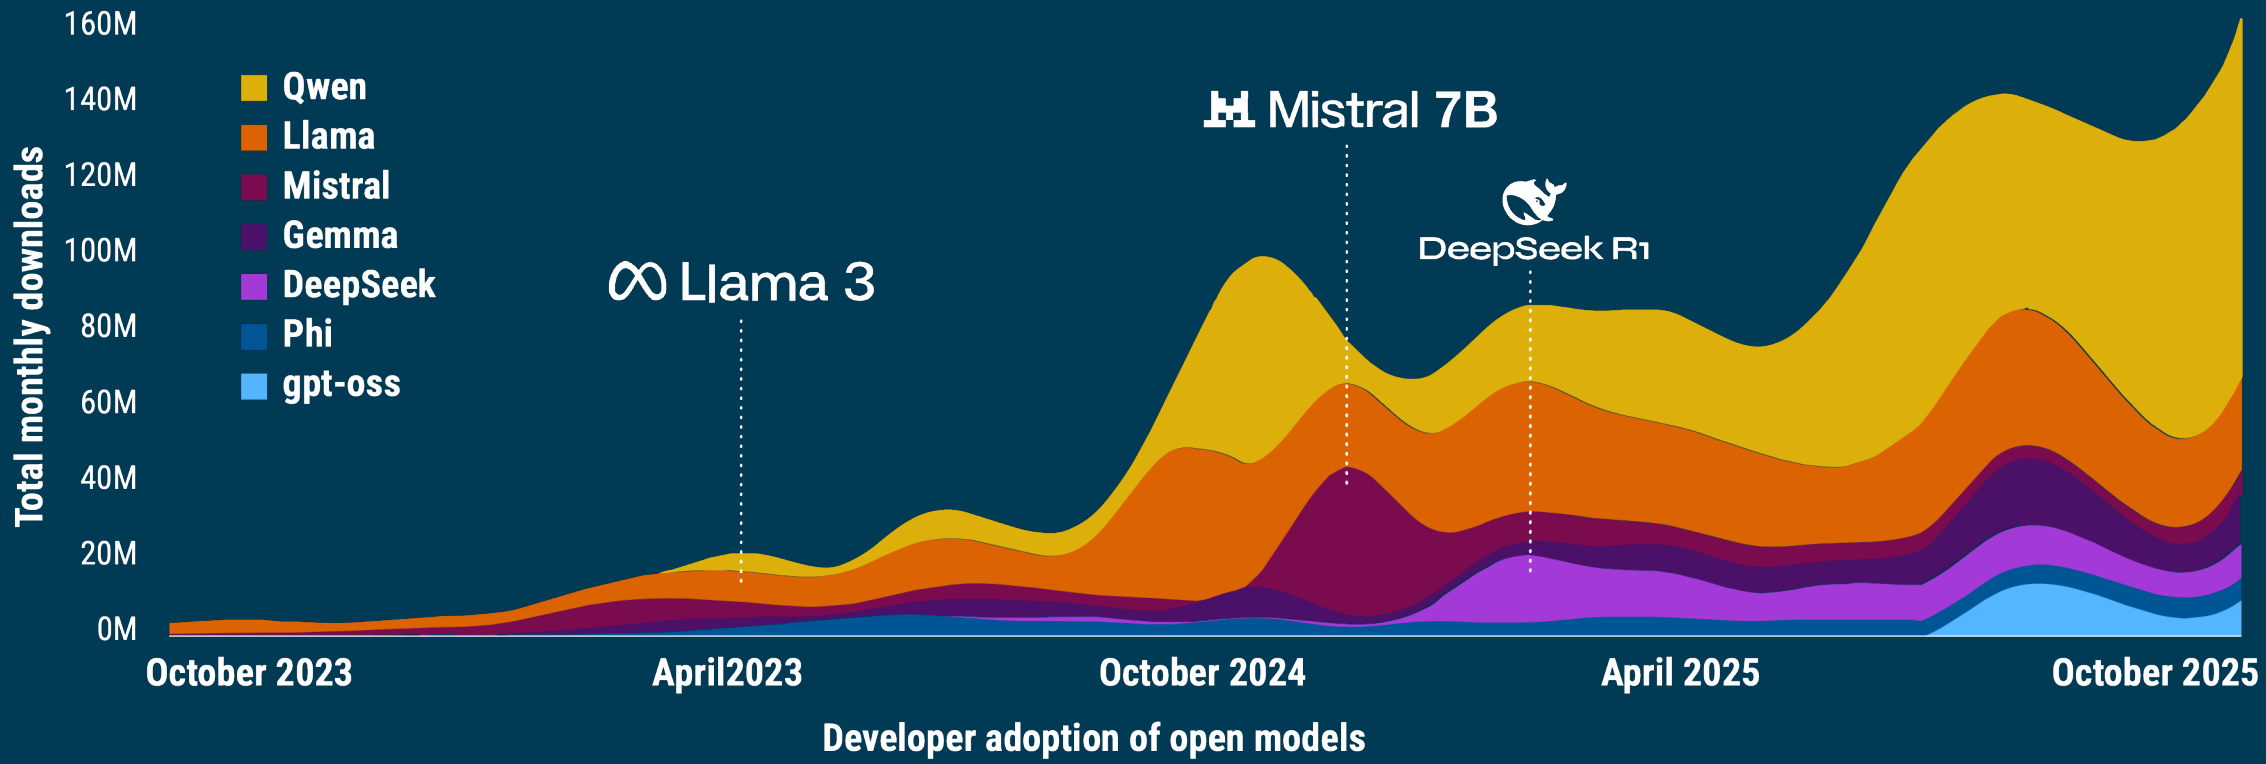

- Is there any reason to pay the frontier model developers when open-source models are sufficient?
- In the end, the most efficient model that leads to the least token utilization to complete a task should win.

Data as of October 2025. Source: Nvidia.

Coding Applications

Will the coding market commoditize?

There is a plethora of coding agents available, why pay for one?

AI Coding Agents — Comparison

CATEGORY	AGENT	VENDOR	INTERFACE	KEY STRENGTHS	BEST FOR
Terminal / CLI	Claude Code	Anthropic	CLI (terminal)	Strong customization via CLAUDE.md memory; excels at large-codebase reasoning and multi-file refactors.	Terminal-first developers wanting a flexible agent.
Terminal / CLI	Codex	OpenAI	CLI + cloud	High benchmark scores on agentic tasks; common head-to-head with Claude Code.	Complex refactors and autonomous execution.
Terminal / CLI	Gemini CLI	Google	CLI (open source)	Free Gemini 3.1 Pro via Google account; 1M token context; built-in Search grounding.	Individuals and OSS contributors wanting free access.
Terminal / CLI	OpenCode	Open source	CLI (open source)	Rapid growth (147K+ stars); Copilot auth partnership; runs fully offline with local models.	Security-conscious or offline workflows.
IDE-based	Cursor	Anysphere	AI-native IDE (VS Code fork)	Agent mode for multi-file edits; supports Composer, Claude, Gemini, and GPT backends.	Integrated IDE work with model flexibility.
IDE-based	GitHub Copilot	GitHub/Microsoft	IDE plugin (VS Code, JetBrains)	Mature inline completions; reliable pair-programming; now with agent mode.	Day-to-day coding and enterprise teams.
IDE-based	Cody	Sourcegraph	IDE plugin	Code graph for understanding huge repos and cross-file relationships.	Engineers in massive monorepos.
IDE-based	Junie	JetBrains	JetBrains IDEs	Autonomous planning, writing, refining, and testing with human-in-the-loop controls.	JetBrains users wanting native autonomy.
Browser / No-setup	Replit Agent	Replit	Browser IDE	End-to-end build (frontend, backend, DB, deploy); plans before coding.	Beginners and prototypers; zero local setup.
Browser / No-setup	Lovable	Lovable AI	Browser	Natural language to web apps; conversational iteration; one-click deploy.	Rapid web-app prototyping.
Browser / No-setup	v0	Vercel	Browser	Generates UI components/pages from prompts; tight Vercel integration.	Frontend prototypes and React components.
Browser /No-setup	Devin	Cognition	Cloud agent	Positioned as a fully autonomous AI software engineer for end-to-end tasks.	Hands-off delegation and longer-running work.

Data as of May 2026. Source: Faros AI; Zapier; MightyBot; Augment Code; Verdent AI; Anthropic Claude. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Anthropic Claude). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

Will the XPU market become competitive?

Hyperscalers are Developing CPU & GPU (XPU) Custom Silicon

HYPERSCALER	CUSTOM CPU	ACCELERATOR (GPU/ASIC)	STRATEGIC POSTURE
Google	Axion (Arm Neoverse); now used as host CPU for TPU systems.	TPU v7 "Ironwood" - 7 generations in; only program at true 10K+ chip cluster scale.	Most mature stack.
Amazon (AWS)	Graviton5 (192 cores, Arm Neoverse); >50% of new AWS CPU capacity for 3 years running.	Trainium3 (3nm, 144 GB HBM3E) for training; Inferentia for inference. ~$20B run rate.	Tight CPU + Nitro DPU + Trainium system integration.
Microsoft	Cobalt 200(132 cores, Arm Neoverse CSS V3).	Maia 200 - deploying in US Central data centers. Latest entrant.	Catching up.
Meta	No proprietary CPU - launch customer for Arm's turnkey AGI CPU platform.	MTIA v3 (mid-2026, Llama inference); MTIA v4 "Santa Barbara" (late 2026, first with HBM4).	Internal use only — no cloud sales.

- **All four use Arm-based CPUs (not x86), fabricate on TSMC 3nm, co-design AS/Cs with Broadcom, and aim custom accelerators primarily at the larger inference market — while Nvidia still wins frontier training.**
- **Anthropic is in talks with Microsoft to rent servers powered by Microsoft-designed AI server chips**

Data as of May 21, 2026. Source: Anthropic Claude; Bloomberg. The graphic presented herein was created with the assistance of a third-party generative artificial intelligence tool (Anthropic Claude). All outputs have been reviewed, edited, and validated by Boston Partners personnel to ensure accuracy, consistency with approved data sources, and compliance with applicable regulatory standards. Boston Partners is responsible for all final content.

Hyperscalers are Developing CPU & GPU (XPU) Custom Silicon

Is Google's TPU roadmap ahead of Nvidia's?

TPU v7 (Ironwood) vs. Blackwell & Rubin: what we know so far about efficiency.

Nvidia GPU & Google TPU efficiency (TFLOPs/US$k)

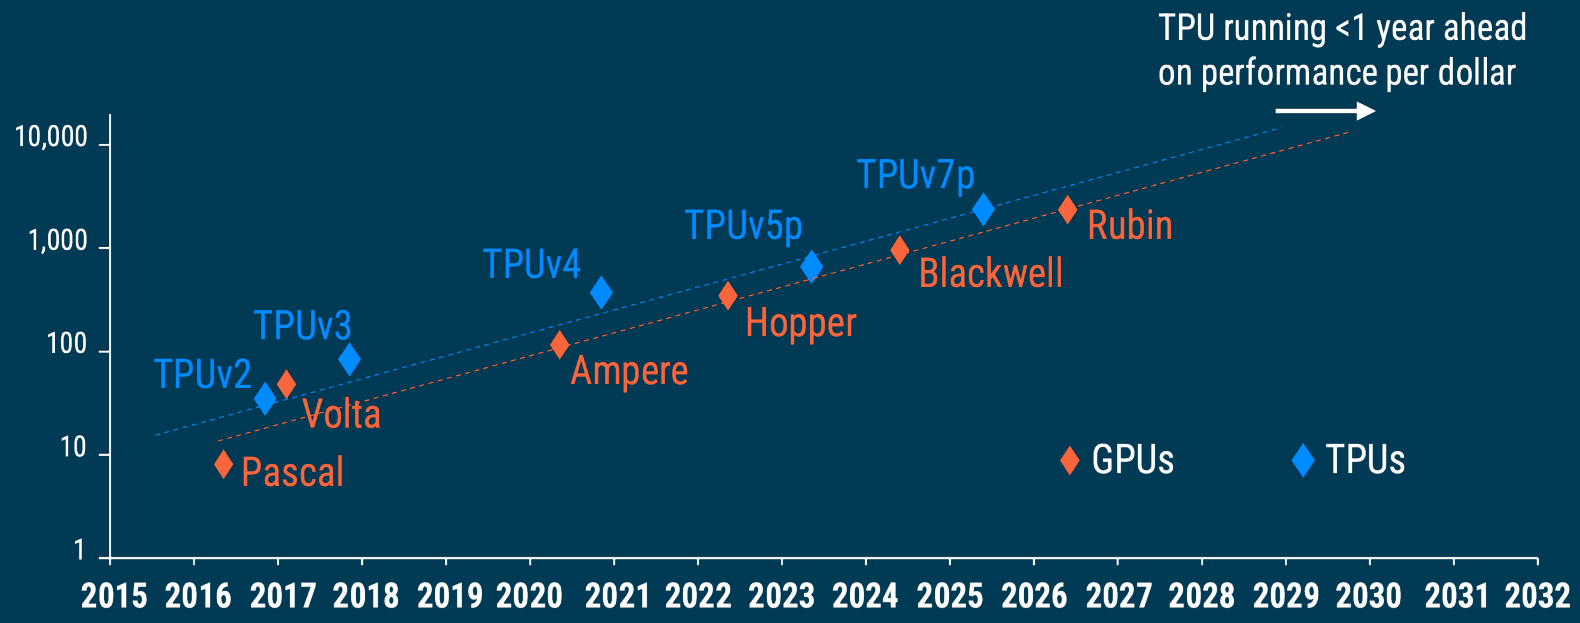

We expect roadmaps to run along parallel lines going forward, with TPU <1 year ahead.

Sundar Pichai, Alphabet CEO, Q1 2026 Earnings Call

- Our custom TPUs, Axion CPUs, and the latest Nvidia GPUs continue to form the industry's widest variety of compute options.
- Our TPUs continue our leadership in performance, cost and power efficiency for customers
- [As demand grows] we'll begin to deliver TPUs to a select group of customers in their own data centers' inner hardware configuration to expand our addressable market opportunity."

Data as of December 1, 2025. Source: Corporate reports and NewStreet Research estimates and analysis. 1 = peak TFLOPs (LHS) adjusted for TCO, sparsity, quantization and for higher ASIC FLOPs utilization (MFU).

Hyperscalers are Developing CPU & GPU (XPU) Custom Silicon

Amazon XPU expected to be competitive. Too late for the others?

Conclusion: everybody will try, but who will succeed & when remains uncertain

Google is there, Amazon is close, others still have to prove the can make it

	Training ASIC	Start of program	Start of XPU ramp'	Design partner	Probability of success	
Google	TPU	2013	2015	Broadcom	100%	Program in place and capturing 90% of AI and ML internal workloads at Google
Amazon	Trainium	2017	2019	Marvell	80%	Third generation in advanced development phase; planning to ramp hundreds of thousands of chips and a single cluster of 400k this year.
Microsoft	Maia	2022	2024	AMD / Marvell (from 2026)	30%	All these players have programs in place and will ramp chips. Whether they successfully catch up with the scale of GPU deployments is an open question. We expect a lot of hybrid deployments ranging from 20/80 to 80/20 splits between GPUs and ASICs
Meta	MTIA	2021	2023	Broadcom		
OpenAI	Nexus	2024?	2026?	Broadcom		
Apple	Baltra	2024?	2026?	Broadcom		
Tesla/xAI	Dojo	2020	2022	Marvell?		
ByteDance	?	2024?	2026?	Broadcom		

The race is open-ended. 50/50 share between ASICs and GPUs for leading clusters.

Data as of March 11, 2025. Source: Corporate reports and NewStreet Research estimates and analysis.
1- assume 2 years lead time after start of the program.

Next Generation Chips – GPU, CPU, LPX

35x "Better"/10x "factory" revenue per gigawatt?

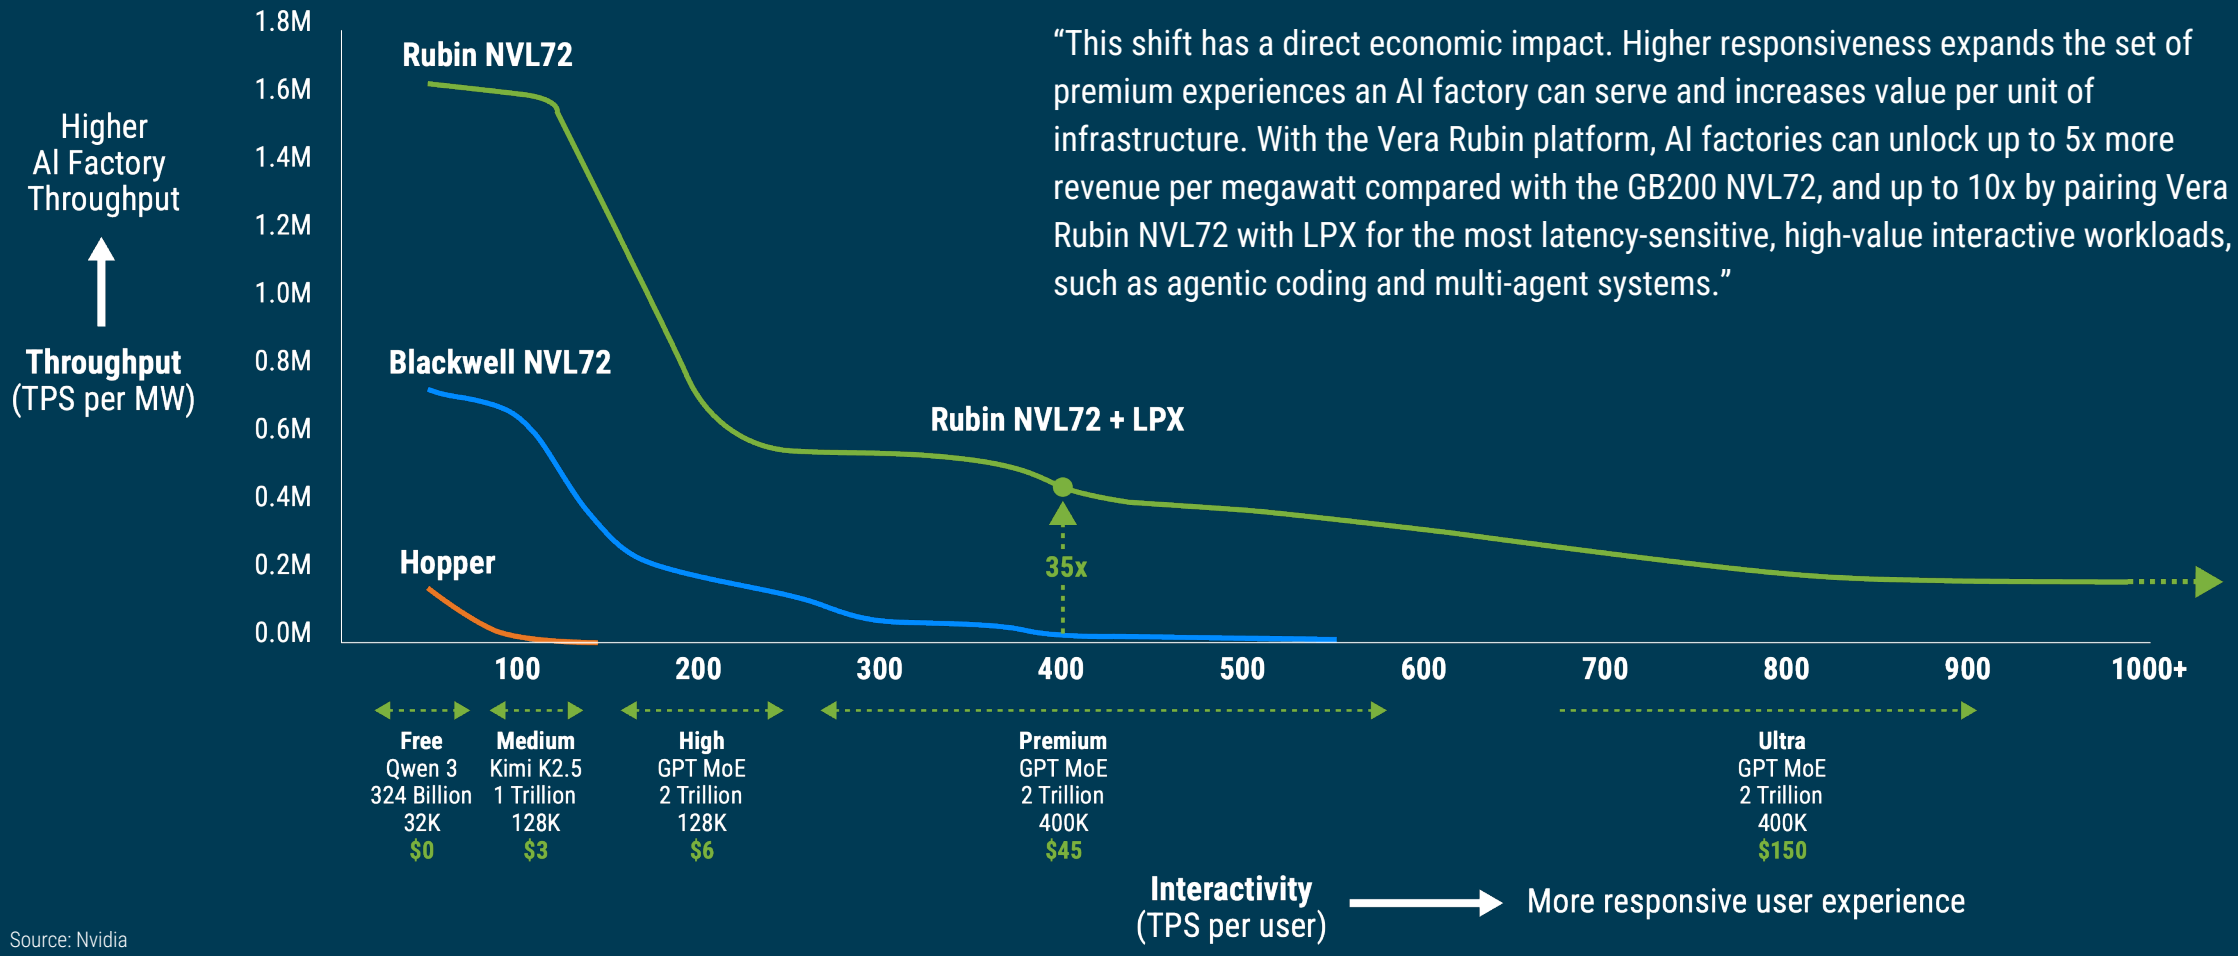

"This shift has a direct economic impact. Higher responsiveness expands the set of premium experiences an AI factory can serve and increases value per unit of infrastructure. With the Vera Rubin platform, AI factories can unlock up to 5x more revenue per megawatt compared with the GB200 NVL72, and up to 10x by pairing Vera Rubin NVL72 with LPX for the most latency-sensitive, high-value interactive workloads, such as agentic coding and multi-agent systems."

Source: Nvidia

Downstream "Picks-and-Shovels"

It has paid to chase

COHR +381% 1-year return (as of May 22, 2026)

COHERENT

Coherent Corp Periodicity Quarterly Source Standard Currency USD

Measure Revenue

	2024	2025	2026	2027	2028	2029
Q1 Sep	1.05B	1.35B	1.58B	2.14B	2.85B	3.49B
Q2 Dec	1.13B	1.43B	1.69B	2.31B	3.03B	3.76B
Q3 Mar	1.21B	1.50B	1.81B	2.46B	3.10B	4.33B
Q4 Jun	1.31B	1.53B	1.99B	2.63B	3.26B	4.74B
Year	4.71B	5.81B	7.06B	9.55B	12.45B	16.93B
Cal Yr	5.31B	6.29B	8.25B	10.97B	13.62B	--

(Fiscal Period: Reported, Estimated)

YoY % Growth

	2024	2025	2026	2027	2028	2029
Q1 Sep	-22%	28%	17%	35%	33%	22%
Q2 Dec	-17%	27%	17%	37%	31%	24%
Q3 Mar	-3%	24%	21%	36%	26%	40%
Q4 Jun	9%	16%	30%	32%	24%	45%
Year	-9%	23%	21%	35%	30%	36%
Cal Yr	15%	19%	31%	33%	24%	--

PoP % Growth

LITE +1150% 1-year return (as of May 22, 2026)

LUMENTUM

Lumentum Holdings Inc Periodicity Quarterly Source Standard Currency USD

Measure Revenue

	2024	2025	2026	2027	2028	2029
Q1 Sep	317.60M	336.90M	533.80M	1.15B	1.87B	2.53B
Q2 Dec	366.80M	402.20M	665.50M	1.33B	2.09B	2.66B
Q3 Mar	366.50M	425.20M	808.40M	1.49B	2.21B	2.94B
Q4 Jun	308.30M	480.70M	989.65M	1.66B	2.41B	3.15B
Year	1.36B	1.65B	3.00B	5.55B	8.57B	11.64B
Cal Yr	1.41B	2.11B	4.28B	7.11B	9.82B	--

(Fiscal Period: Reported, Estimated)

YoY % Growth

	2024	2025	2026	2027	2028	2029
Q1 Sep	-37%	6%	58%	115%	63%	36%
Q2 Dec	-28%	10%	65%	101%	56%	28%
Q3 Mar	-4%	16%	90%	85%	48%	33%
Q4 Jun	-17%	56%	106%	68%	45%	31%
Year	-23%	21%	82%	85%	54%	36%
Cal Yr	-2%	49%	103%	66%	38%	--

PoP % Growth

MRVL +223% 1-year return (as of May 22, 2026)

MARVELL

Marvell Technology Inc Periodicity Quarterly Source Standard Currency USD

Measure Revenue

	2024	2025	2026	2027	2028	2029
Q1 Apr	1.32B	1.16B	1.90B	2.40B	3.22B	4.41B
Q2 Jul	1.34B	1.27B	2.01B	2.59B	3.50B	4.68B
Q3 Oct	1.42B	1.52B	2.07B	2.82B	3.84B	4.96B
Q4 Jan	1.43B	1.82B	2.22B	3.06B	4.29B	5.26B
Year	5.51B	5.77B	8.19B	10.87B	14.94B	19.23B
Cal Yr	5.64B	8.06B	10.59B	14.45B	18.99B	21.07B

(Fiscal Period: Reported, Estimated)

YoY % Growth

	2024	2025	2026	2027	2028	2029
Q1 Apr	-9%	-12%	63%	27%	34%	37%
Q2 Jul	-12%	-5%	58%	29%	35%	33%
Q3 Oct	-8%	7%	37%	36%	36%	29%
Q4 Jan	1%	27%	22%	38%	40%	23%
Year	-7%	5%	42%	33%	37%	29%
Cal Yr	2%	43%	31%	36%	31%	11%

PoP % Growth

MU +704% 1-year return (as of May 22, 2026)

micron

Micron Technology Inc Periodicity Quarterly Source Standard Currency USD

Measure Revenue

	2023	2024	2025	2026	2027	2028
Q1 Nov	4.09B	4.73B	8.71B	13.64B	44.06B	49.17B
Q2 Feb	3.69B	5.82B	8.05B	23.86B	45.48B	48.46B
Q3 May	3.75B	6.81B	9.30B	33.57B	47.46B	47.64B
Q4 Aug	4.01B	7.75B	11.32B	40.19B	48.89B	47.19B
Year	15.54B	25.11B	37.38B	110.26B	178.99B	173.46B
Cal Yr	16.89B	29.84B	47.58B	148.88B	192.00B	192.43B

(Fiscal Period: Reported, Estimated)

YoY % Growth

	2023	2024	2025	2026	2027	2028
Q1 Nov	-47%	16%	84%	57%	223%	12%
Q2 Feb	-53%	58%	38%	196%	91%	7%
Q3 May	-57%	82%	37%	261%	41%	0%
Q4 Aug	-40%	93%	46%	255%	22%	-3%
Year	-49%	62%	49%	195%	62%	-3%
Cal Yr	-35%	77%	59%	213%	29%	0%

PoP % Growth

Hyperscalers are projected to spend $1.4 trillion 2026-2027. With that level of capital flowing, it doesn't take much to move revenues of small downstream companies that help solve AI scaling bottlenecks: memory, storage, optical.

Data as of May 22, 2026. Source: Bloomberg.

Downstream “Picks-and-Shovels”

It has paid to chase

Est. Revenue Progression, Trailing 6 Months

(October 1, 2025 – May 22, 2026)

— 2026 — 2027 — 2028

Coherent Corp (COHR)

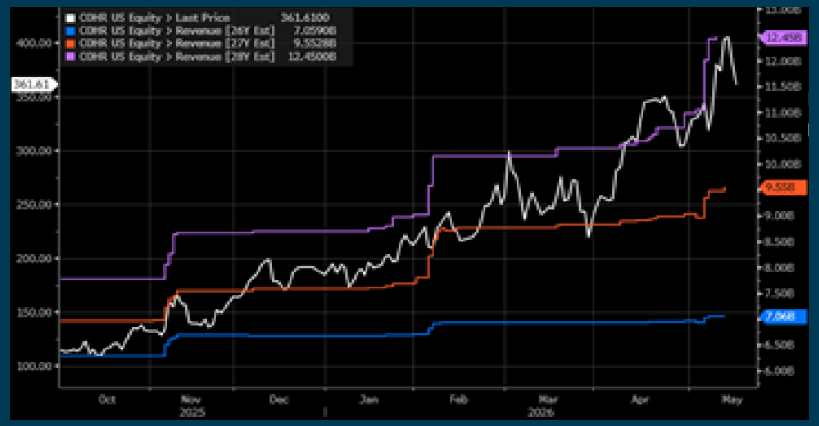

Lumentum (LITE)

Marvell (MRVL)

Micron (MU)

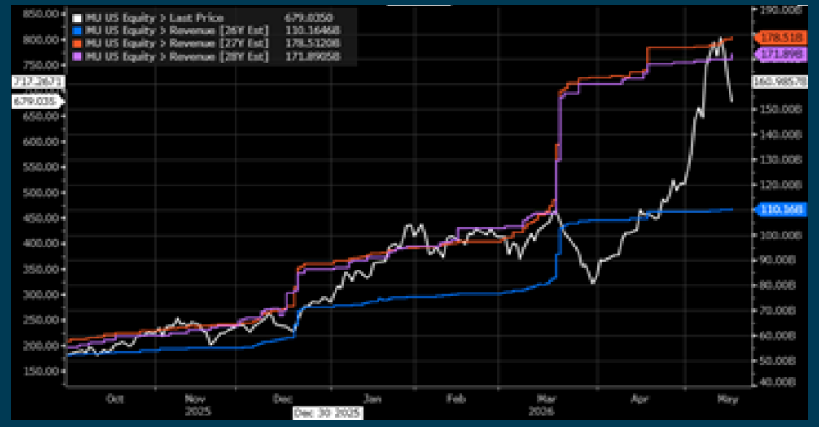

Hyperscaler Cloud Operating Margins

Will AI compute margins mirror cloud service operating margins?

Hyperscaler Cloud Service Operating Margins

amazon	Year Ended Dec. 31, 2023	Year Ended Dec. 31, 2024	Year Ended Dec. 31, 2025
AWS Operating Margin	27%	37%	35%
Microsoft	Year Ended June 30, 2023	Year Ended June 30, 2024	Year Ended June 30, 2025
Intelligent Cloud Operating Margin	39%	43%	42%
Google	Year Ended Dec. 31, 2023	Year Ended Dec. 31, 2024	Year Ended Dec. 31, 2025
Google Cloud Operating Margin	5%	14%	24%

- Cloud margins do not indicate cloud services are a commodity.
- If AI compute margins are in line with cloud service margins, ROI will be satisfactory to good.

Source: Amazon (AMZN) 2024 10-K; Microsoft (MSFT) 2024 10-K; Alphabet (GOOGL) 2025 10-K. ROI: Return on investment.

Micron Over the Years – the Boom and Bust of Memory Cycles

Revenue, Cost of Goods Sold and Gross Margin 1984-2026

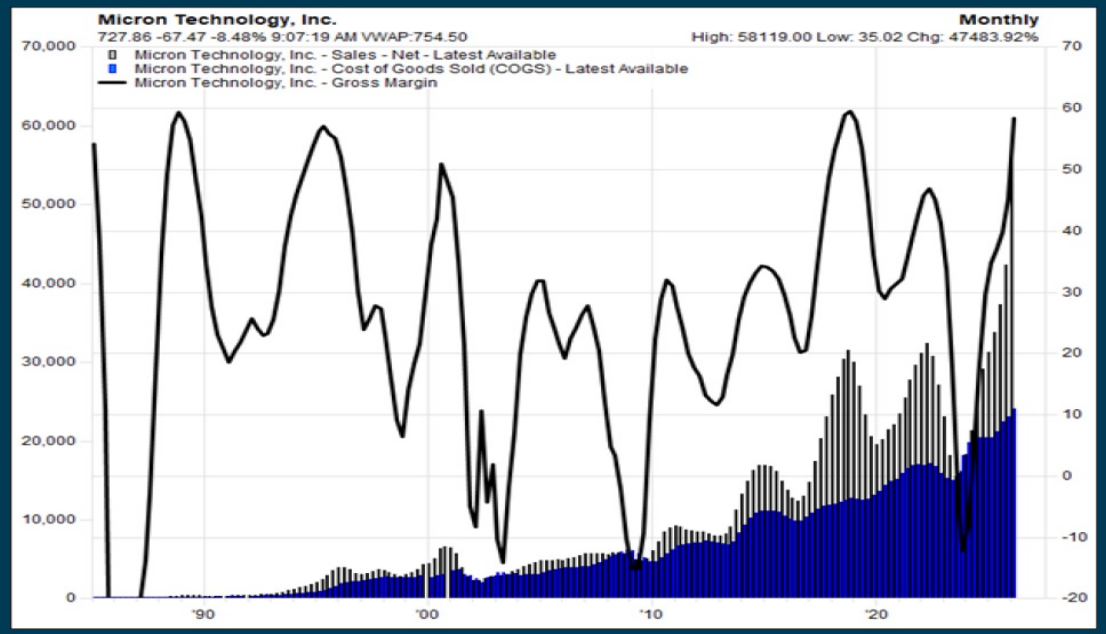

micron™

The 1980s
The rise of the personal computer (PC) created massive demand.

The 1990s
Windows 95 and the mainstreaming of the PC drove a massive "supercycle."

2017–2019
The cloud & server wave/smartphone upgrade super-cycle

2025–Present
The AI "supercycle."

Source: Factset

The "compute = revenue" thesis will not likely hold if AI continues to disrupt employment. Unemployment and AI success must be mutually exclusive; a 5-10 percentage point increase in unemploym

Thinking about AI – an Investor's Framework

Conclusions

- The "compute = revenue" thesis must hold, and there must be ROI through either products & services and/or productivity/cost avoidance.
    - If there's no ROI or productivity gain, there is no technology transition that really amounts to anything. *Aneel Bhusri, Workday Inc Co-Founder, CEO, 4Q26 Earnings Call 2/24/26*
    - How do you measure success for an agent? "…our kind of principle of the end state of some of these things is that everything should compress down to an outcome and a budget and that's probably about it." *Angela Jiang, Head of Product, Platform @ Anthropic, AI & I podcast (36:10)*
- Is AI spend realistic "$3 trillion to $4 trillion annually by the end of this decade"?
- The "compute = revenue" thesis will not likely hold if AI continues to disrupt employment. Unemployment and AI success must be mutually exclusive; a 5-10 percentage point increase in unemployment (recession levels) would not benefit AI.
- The lowest cost of compute will win, both for inference and training
- The models that solve for efficiency should ultimately win.
- Valuations are critical for long-term success. Momentum is critical for short-term success.
- It's early. Currently there are more questions than answers.
- Is it Dec. 1996, Greenspan's "irrational exuberance" or March 2000? Starting to feel a lot like the late 1990's, perhaps with more substance this go-around.
- There is enough evidence of AI's potential that some exposure is warranted, upstream, downstream, and selectively some areas that have been deemed disrupted, e.g., SaaS.

Thinking about AI – an Investor's Framework

Risks

- Policy Response slowing investment outright
  - "If the future of truth is scary, so is the future of jobs. China seems to understand that. Its courts are trying to balance widespread A.I. adoption with the unemployment it may bring about, my colleague Catie Edmondson reports. One precedent-setting decision recently said that an employee couldn't be fired after he'd been replaced by software." *Sam Sifton, NY Times, May 20, 2026*
  - Pope Leo XIV releases his first encyclical, entitled 'Magnifica humanitas: On Safeguarding the Human Person in the Time of Artificial Intelligence.' He appeals for the safeguarding of humanity, promotion of truth, dignity of work, social justice, and peace. *Vatican News, May 25, 2026*
- Policy Response taxing "AI"
- Supply chain shortages/disruptions delay buildout
- Power supply chain shortages
- The economy cannot absorb mass white-collar job losses without significant negative consequences.
- Timing of spend relative to usage – dark fiber example.
  - CEO Dario Amodei, "cone of uncertainty" data centers take one to two years to build, so companies are committing billions now for demand they can't yet verify. Buy too little and lose customers when you don't have enough capacity. Buy too much and revenue doesn't arrive on schedule, the math stops working.
- Token maximizing – will there be an ROI discussion at some point rationalizing usage and decreasing token demand?
  - Microsoft's Claude Code pilot burned through its full annual AI budget in months due to token-based billing, triggering a June 30 cancellation. The Verge, May 14, 2026
  - Uber Burns Its 2026 AI Budget In Four Months On Claude Code. Forbes, May 17, 2026
- Do we need more compute after "super intelligence"? When is enough, enough?

Thinking about AI – an Investor's Framework

Have we been here before?

A December 1951 advertisement for the IBM 604 Electronic Calculating Punch that was first produced in 1948.

Source: Thanks to Benedict Evans for pointing this out.; Wikimedia Commons (https://commons.wikimedia.org/wiki/File:IBM_150_Extra_Engineers_1951.jpg)

About Boston Partners

Boston Partners is a value equity manager with a distinctive approach to investing—one that combines attractive valuation characteristics with strong business fundamentals and positive business momentum in every portfolio. The consistent application of this approach over more than 30 years by an experienced and long-tenured team has created a proven record of performance across economic cycles, market capitalizations, and geographies.

8971994.2

BostonPartners

One Beacon Street | Boston, MA 02108 | tel: 617-832-8200 | bostonpartners.com