

Investment Thesis | NVIDIA.

Rating: | 4 Year IRR:
BUY! | **23.2%**

 InFinance |  SENNA
Organização Estudantil de Finanças | Insuper

Long NVIDIA: The time to own the Future is now

1. Proprietary, integrated ecosystem underpins Nvidia's edge

+85% Market Share
In GPUs

Hardware + Software + CUDA =
Competitive Advantage

2. Jensen's long-view, powered by Nvidia's human capital

Jensen owns ~3%
of NVIDIA

16 Years
Average executive tenure

3. Why buy NVIDIA: ecosystem + talent primed to beat expectations

Ecosystem + Talent =
Structural advantages in AI

Just the beginning of a
Promising market

Thesis numbers summary:

23.2%
2028 IRR

Entry Value ₂₀₂₅
US\$3.4 Trillions

29x P/E
Entry Multiple₁

28x P/E
Exit Multiple

Exit Value ₂₀₂₈
US\$5.7 Trillions

NVIDIA at a Glance

Founded in 1993 with a focus on gaming, NVIDIA is now the world's most valuable company, driven primarily by its leadership in data centers

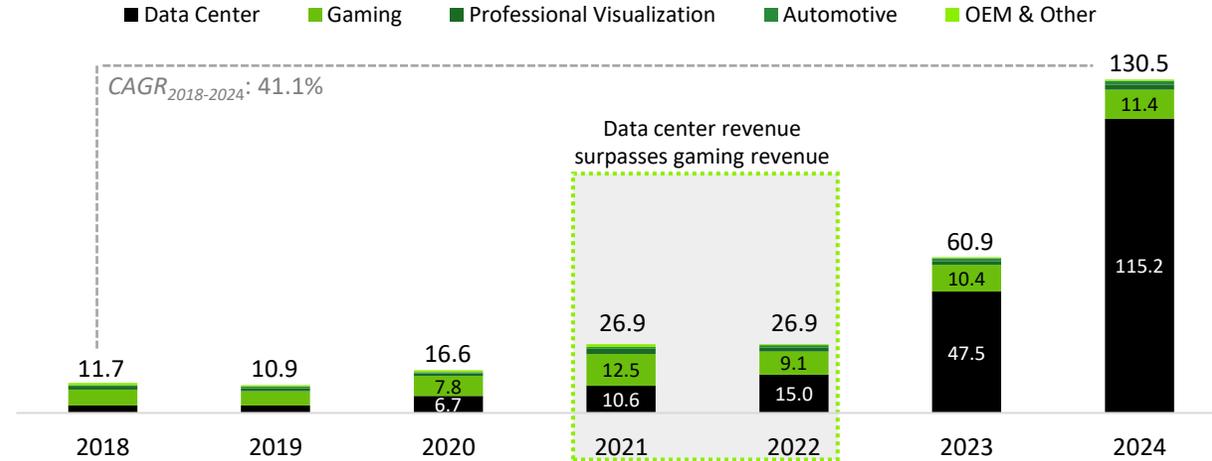
I. Once gaming-focused, NVIDIA's product line now underpins AI infrastructure

NVIDIA main products



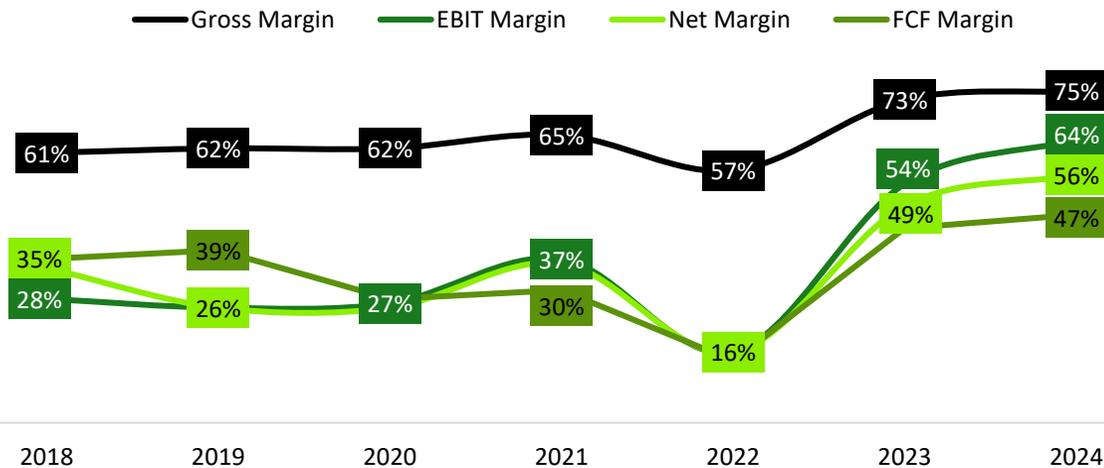
II. AI overtook gaming and data centers became NVIDIA's core business

NVIDIA Revenue by Segment (US\$ billion)



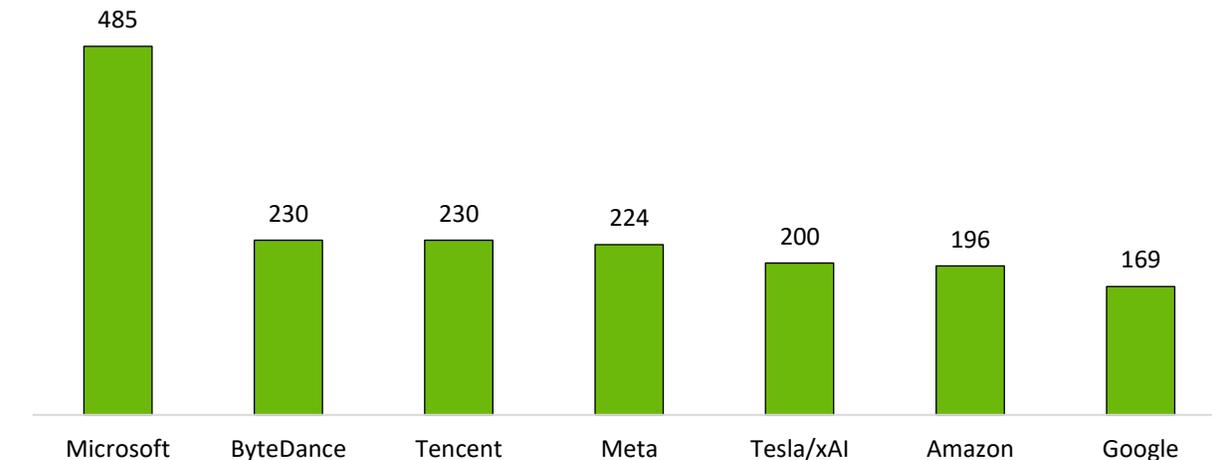
III. Nvidia tripled its margins since 2022, driven by AI and strong operating leverage

NVIDIA's margins (%)



IV. Top buyers like Microsoft made NVIDIA the core enabler of the AI era

Estimated shipments of Nvidia Hopper GPUs in 2024, by customer (in thousands)

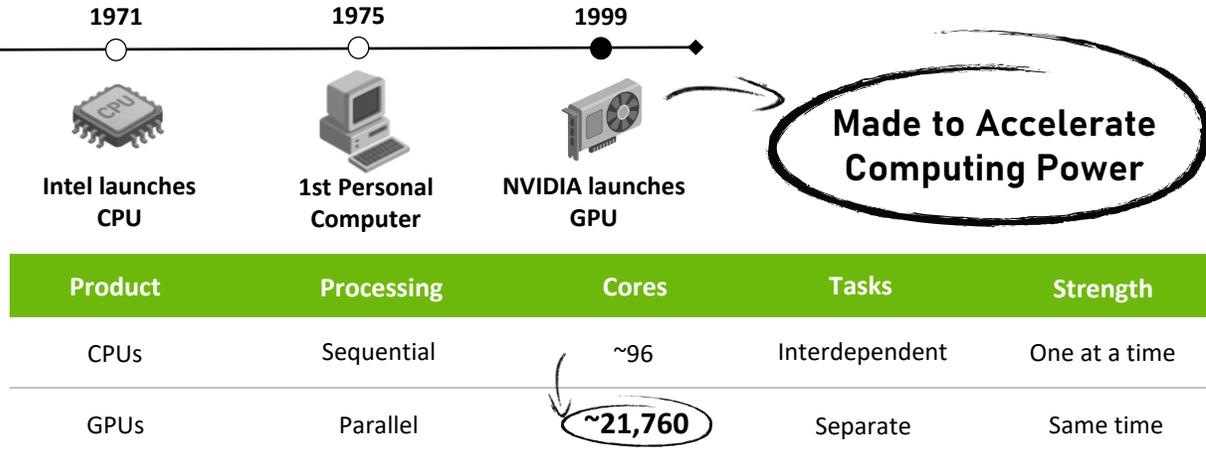


GPUs: NVIDIA's Main Product

Pioneered by NVIDIA, the GPU unlocked unprecedented leaps in data processing power across multiple applications

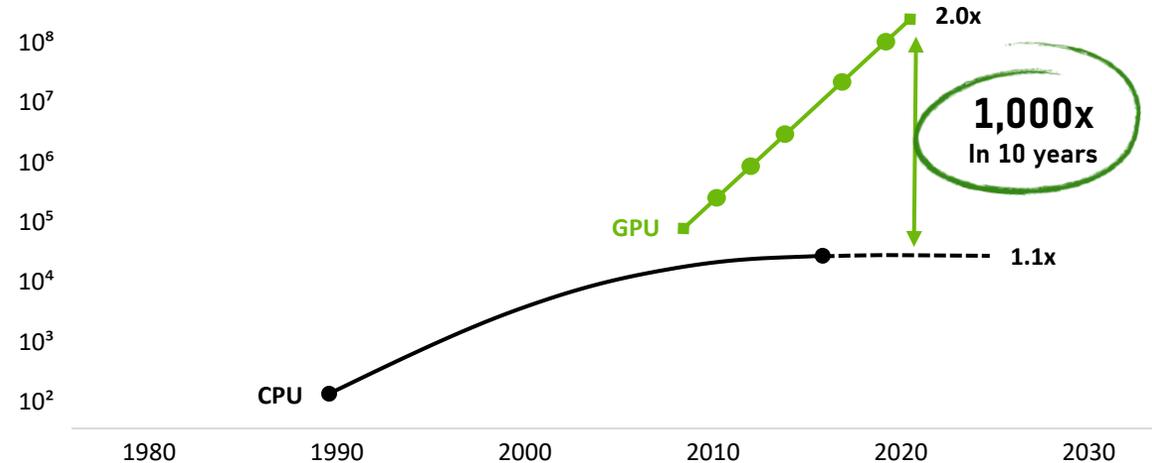
I. Built to ease CPU strain, GPUs specialized in rendering tasks through parallel processing

What is a GPU?



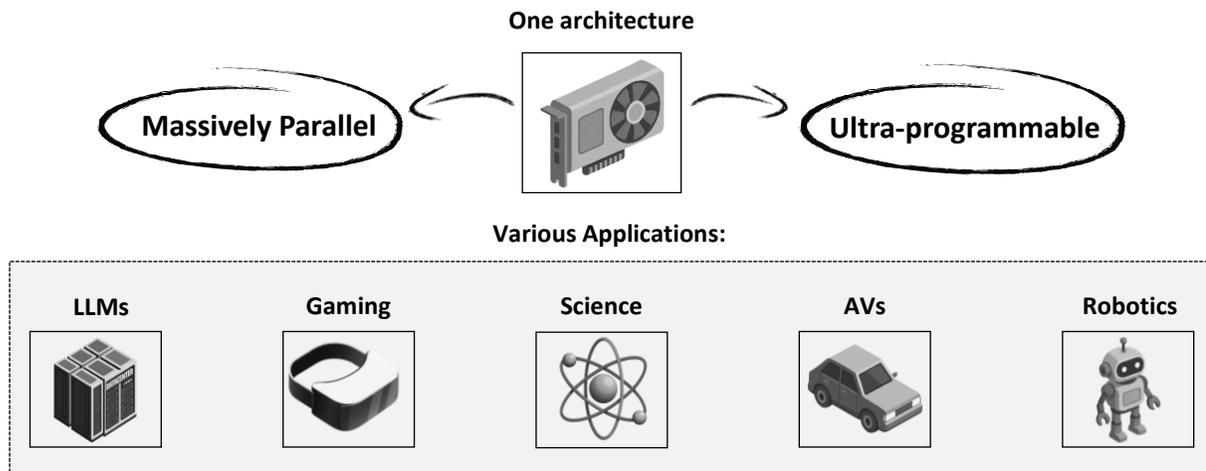
II. By breaking Moore's Law, GPUs disrupted the ceiling of computational power

CPU Vs. GPU in Trillions of Operations per Second (TOPS)



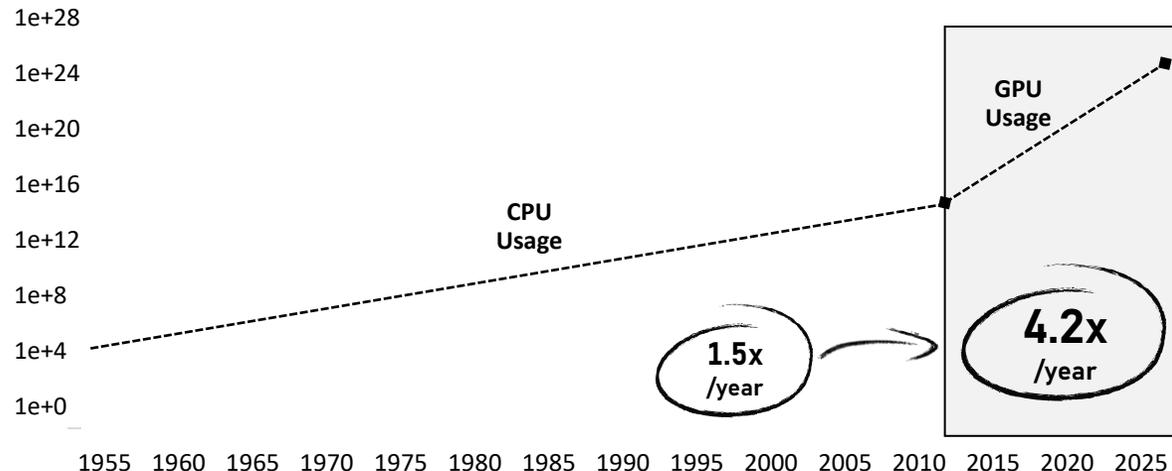
III. Programmability and high data throughput make GPUs ideal for heavy workloads

Flexibility through end markets



IV. GPUs powered extraordinary ML breakthroughs with parallelism and scalability

Training Compute of Notable Machine Learning Systems Over Time (FLOP)

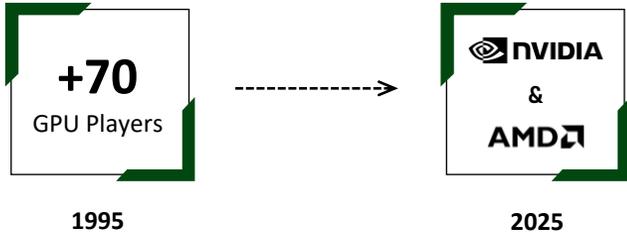


MOAT I: Ahead in GPU Technology

Early vision and execution secured GPU leadership, a position it's poised to strengthen in the future

I. Execution ensured NVIDIA's survival in a market where competitors disappeared

What about the competitors in this market?



Jon Y., Asianometry Founder

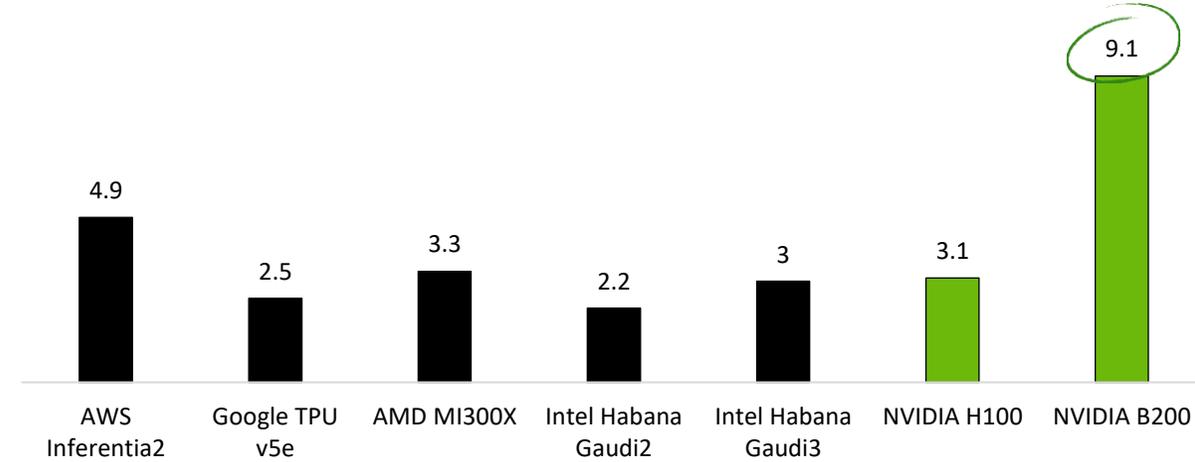


What has driven NVIDIA's GPU edge over AMD over time?

"The main reason is due to their **GPU-centric vision** and the **exceptional execution** capability of their R&D team"

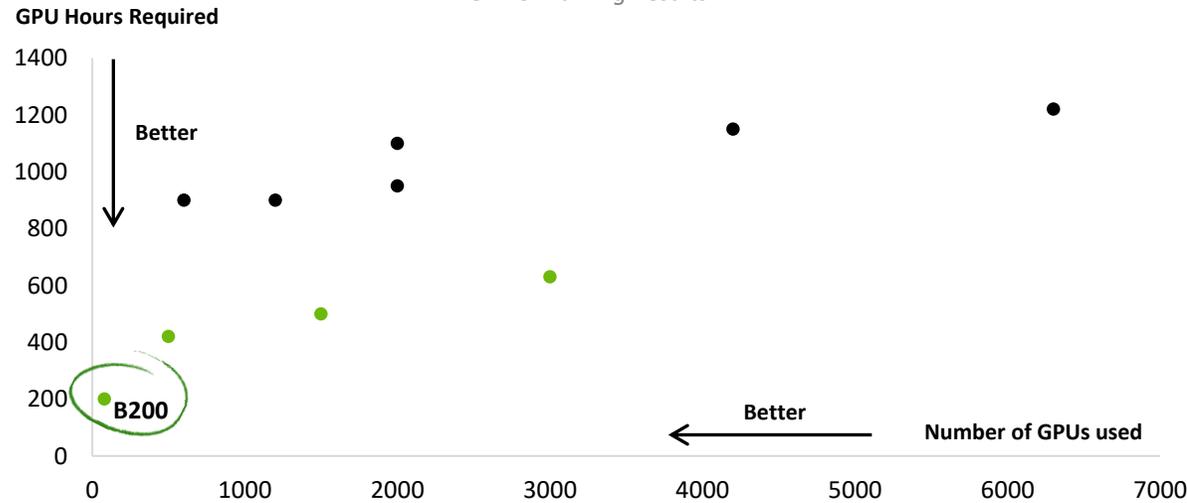
II. Delivers superior cost-benefit in its GPUs when compared to rivals

Chip-level cost-performance ratio (Performance/Cost)



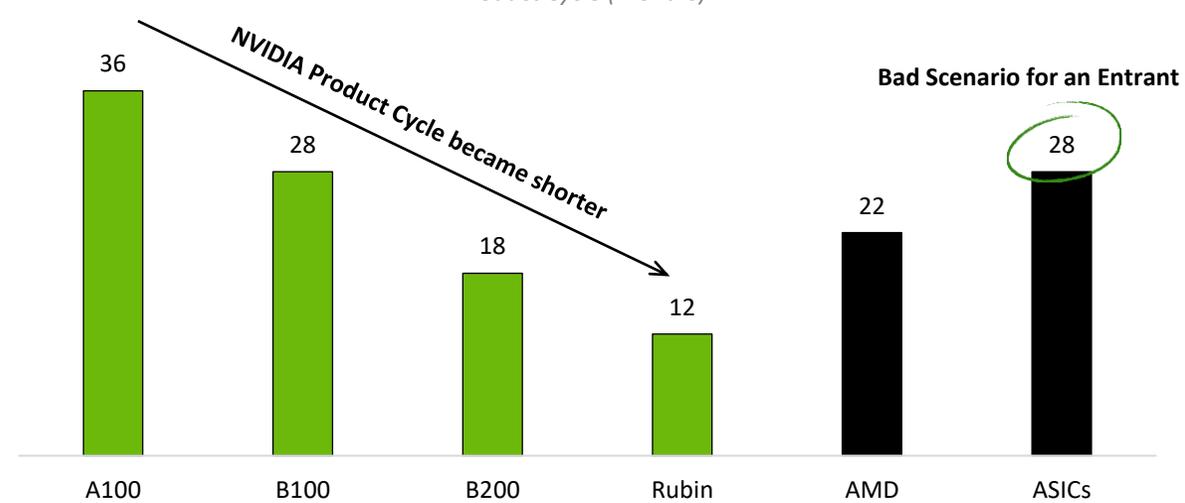
III. Unparalleled performance at the forefront of AI Training

GPT-3 Training Results



IV. Shortening product launch cycles, entrants face a suffocating competitive landscape

Product Cycle (Months)



NVIDIA Ecosystem Human Capital AI Market Valuation Scenarios

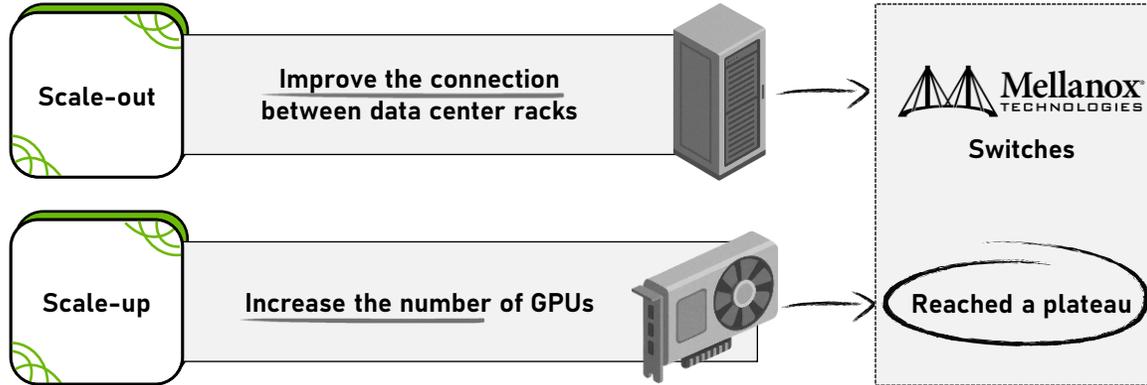
MOAT II: NVLink Scale Capabilities

With NVLink, NVIDIA introduced a new scale-up paradigm through ultra-fast GPU interconnections

I. To boost performance, simply increasing the number of GPUs eventually hits a limit

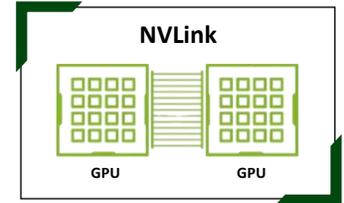
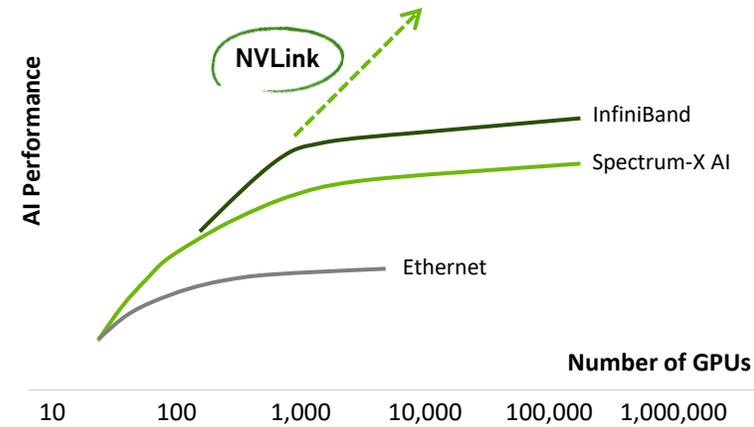
Why use so many GPUs?

Approaches to increase your computing power:



II. NVLink broke the scaling barrier—enabling limitless performance through added GPUs

AI Performance Across Different Products use



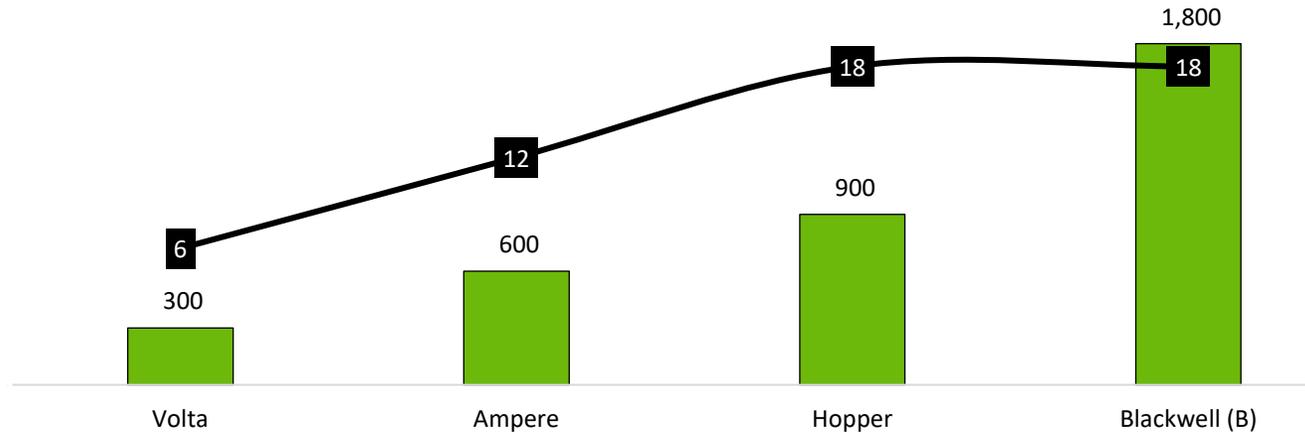
Faster Data Exchange

Limitless Performance Evolution

III. Accelerated GPU communication made it possible to unify 72 GPUs into a single compute entity—the next leap is linking full racks and turning them into an ultra GPU

Connectivity Performance

Transmission (GB/s) — Links/GPU



Present:



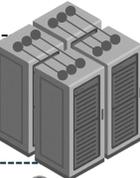
GB200
1 Grace GPU
2 Blackwell GPUs



GB200 NVL72
36x GB200 Superchips
72 GPUs + 36 CPUs

Next Step:

Connecting GPUs across entire racks



Future:

AI Fabric: A One-stop Shop

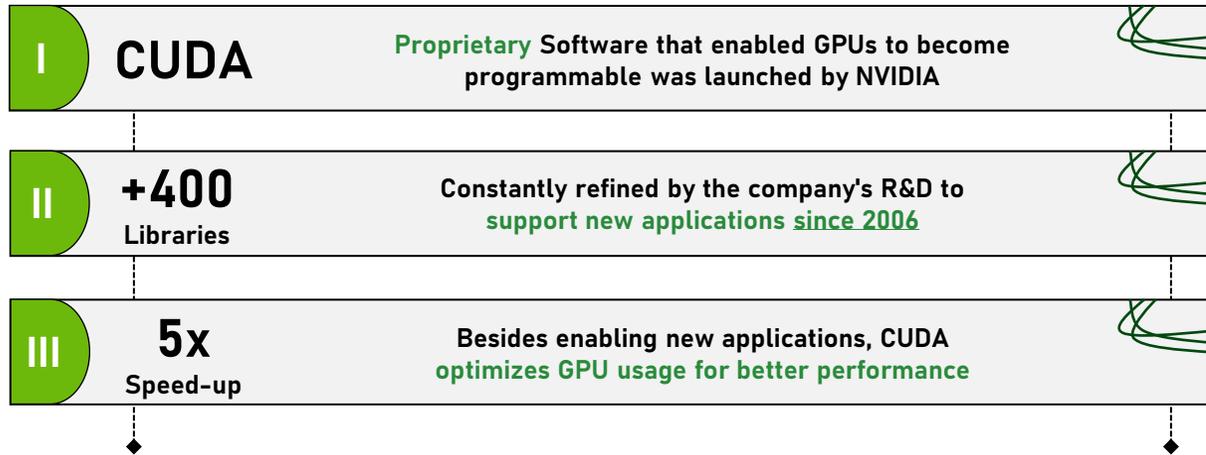


MOAT III: Software Switching Costs

NVIDIA's core moat, CUDA, provides industry-leading GPU performance and reinforces customer dependency

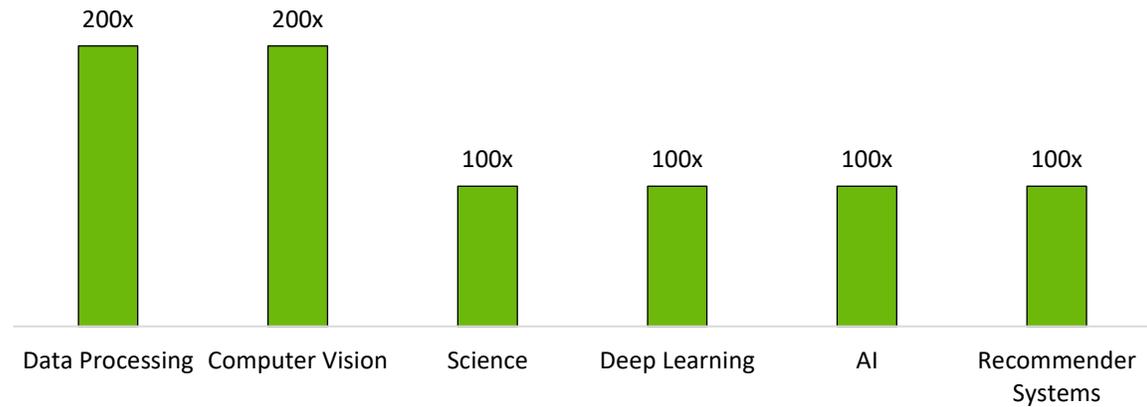
I. Launched in 2006, CUDA enabled GPU programmability and performance enhancement

What made the GPU so programmable?



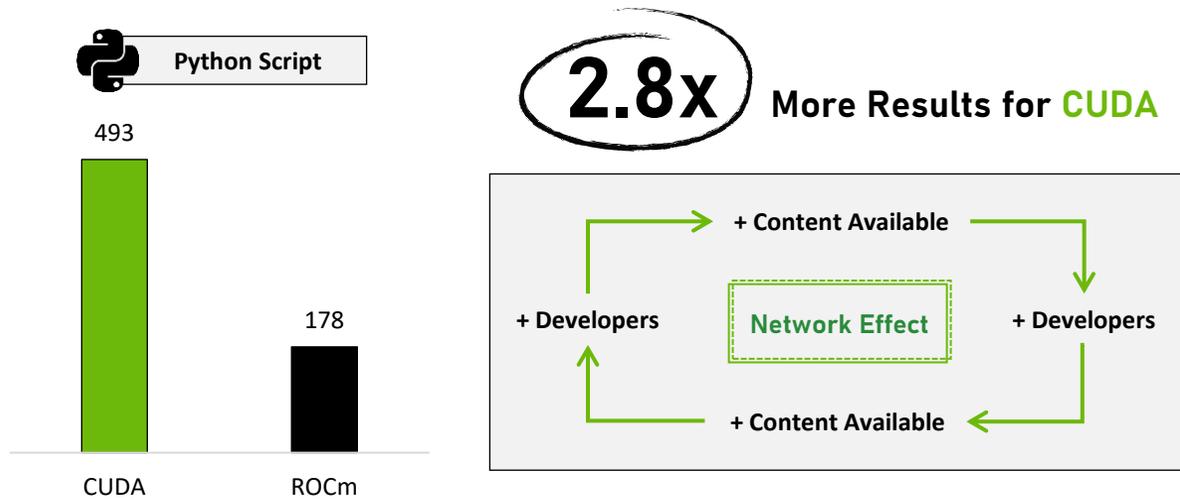
II. CUDA enhances GPU efficiency by enabling fine-grained programmability

Performance Enhancing (Speedup)



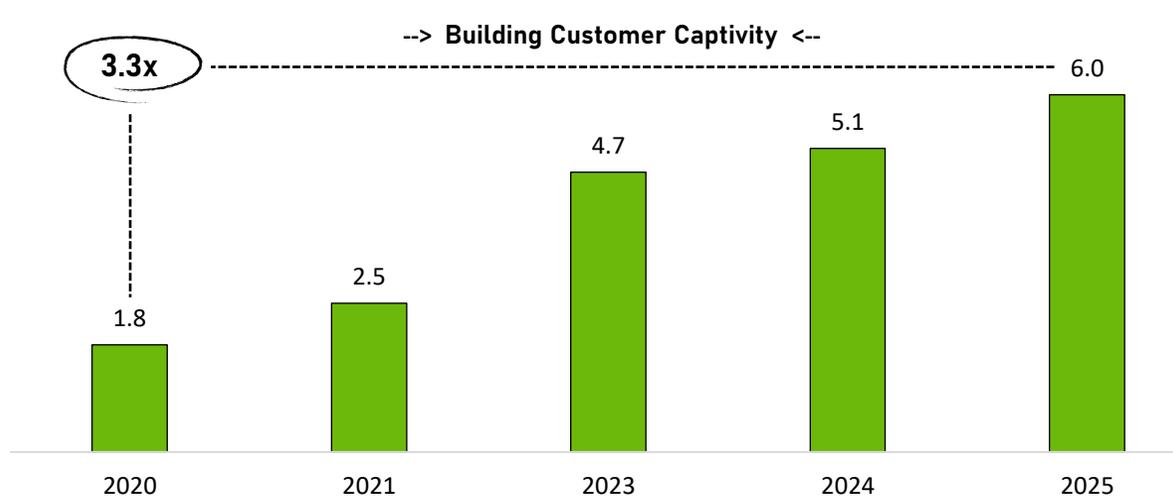
III. AMD's solution late 2016 launch gave CUDA a head start through the network effect

CUDA (NVIDIA) vs. ROCm (AMD): Measuring Content Presence on YouTube



IV. More developers, stronger lock-in—CUDA reinforces retention

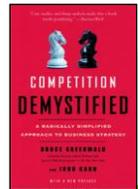
CUDA Developers (Million)



Competitive Advantage

With a unique blend of competitive strengths and world-class execution, NVIDIA is strongly positioned to lead into the future

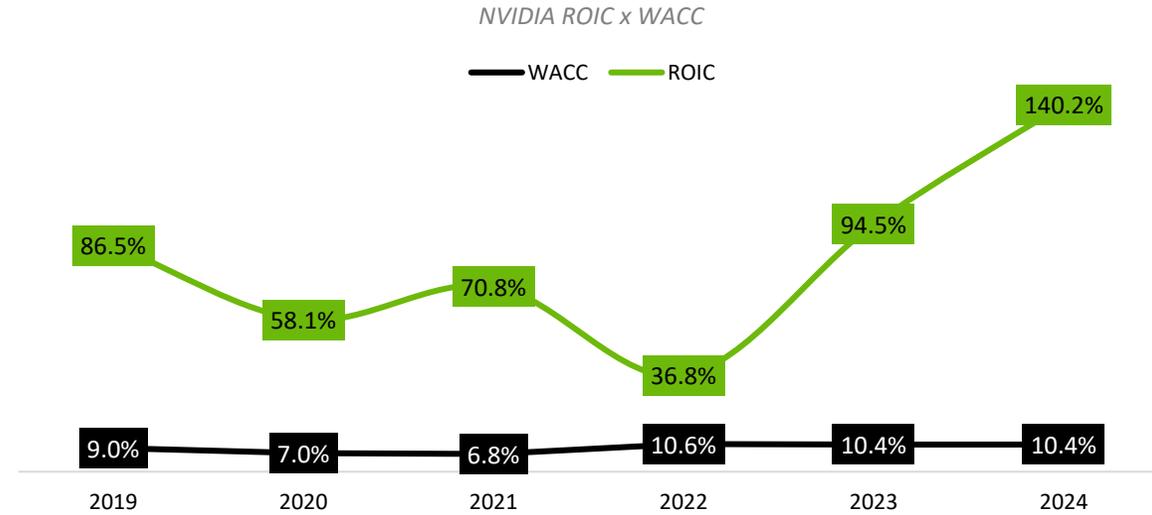
I. A winning business model that compounds over time



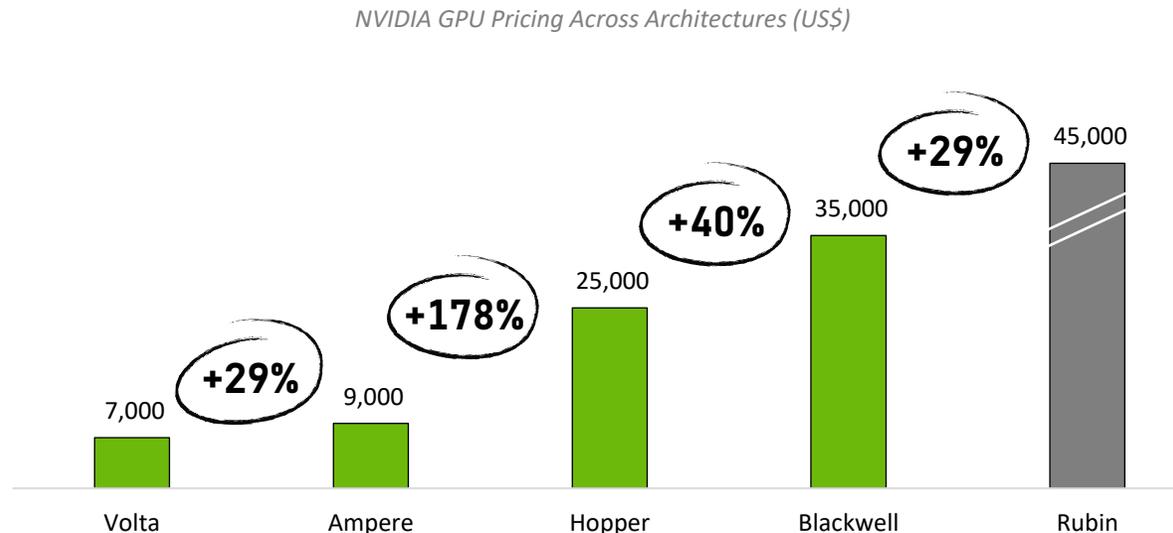
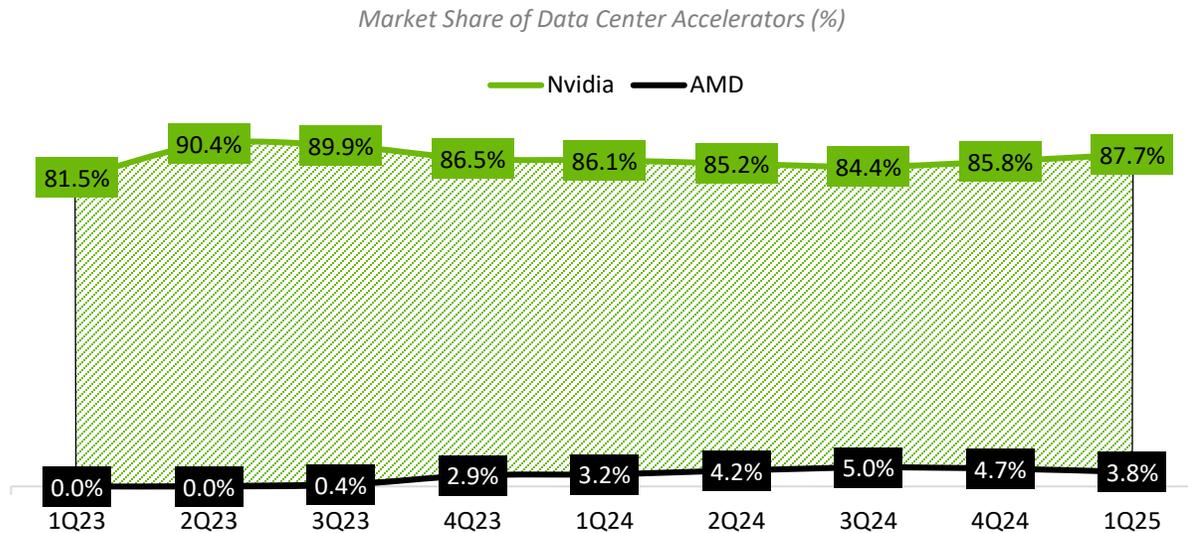
Competition Demystified, Bruce Greenwald

“The most powerful competitive advantages arise when customer captivity is combined with economies of scale.”

II. The underlying economics highlight the presence of significant entry barriers



III. Sustained market share highlights structural dominance, which underpins pricing power and signals strong switching barriers



Skilled Human Capital

While the market sees chips, NVIDIA invests where few look: In the people who make the future possible

I. Jensen's long-term vision is grounded in identifying unmet needs early, and deliberately building the infrastructure to meet them before others even recognize the gap



Jensen Huang, NVIDIA's Founder & CEO

"No matter how successful you were yesterday, if you don't innovate today, you become irrelevant. That's why at Nvidia, we wake up every morning as if we were running out of time"

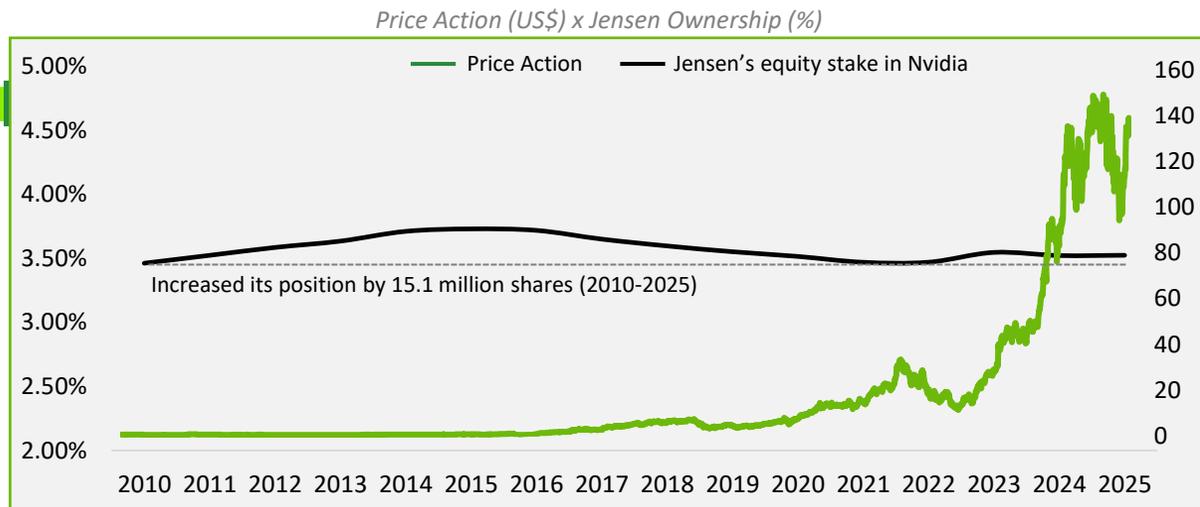


43% based on Short-Term Performance

Jensen's Remuneration

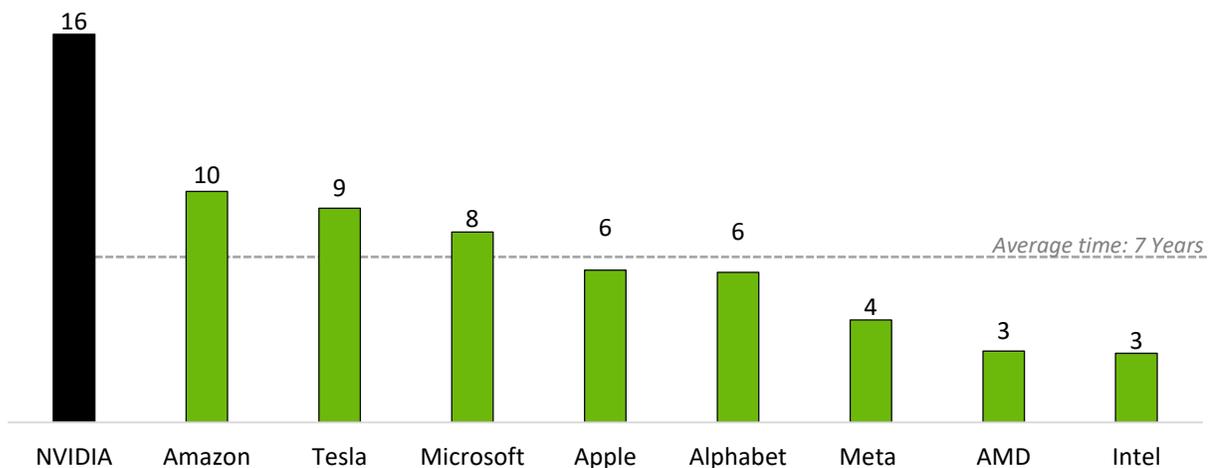
52% based on Long-Term Performance

5% is fixed



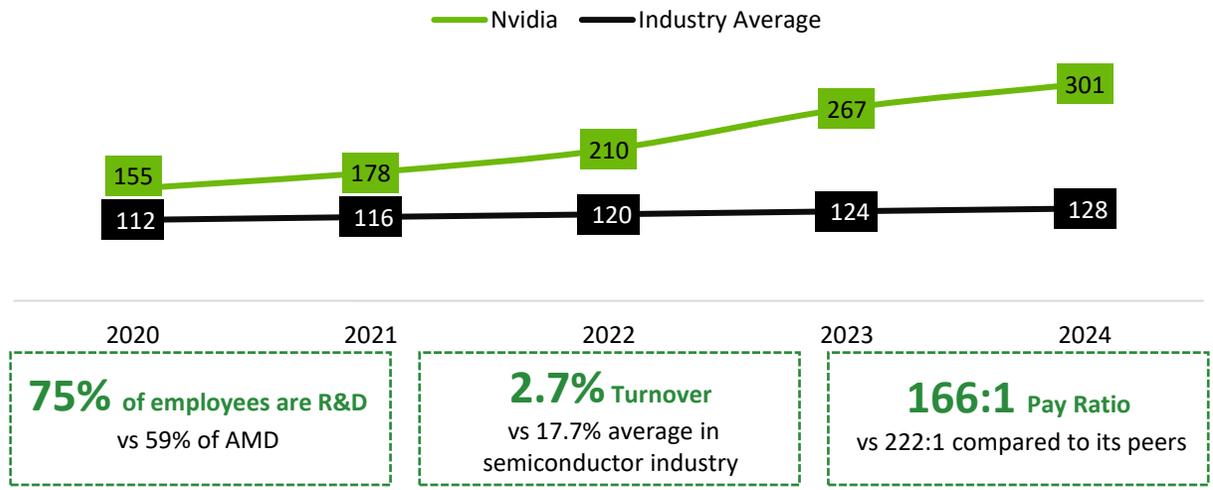
II. NVIDIA's long executive tenure quietly drives an execution edge rivals can't match

Average executive tenure (Years)



III. Beyond leadership, NVIDIA's skilled workforce makes execution its core strength

NVIDIA Employee Pay vs Semiconductor Industry Average (US\$ Thousand)

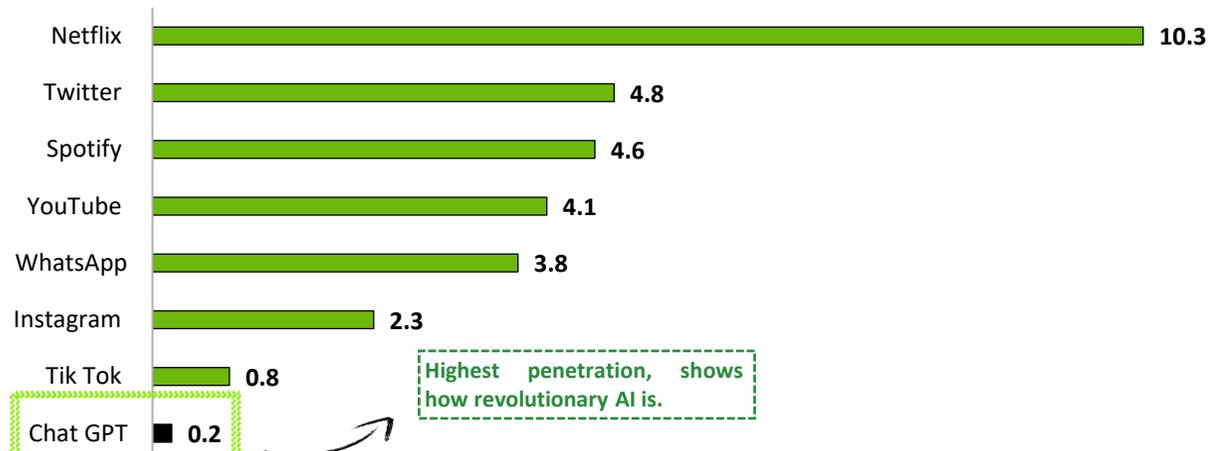


It's Still Early Days for AI

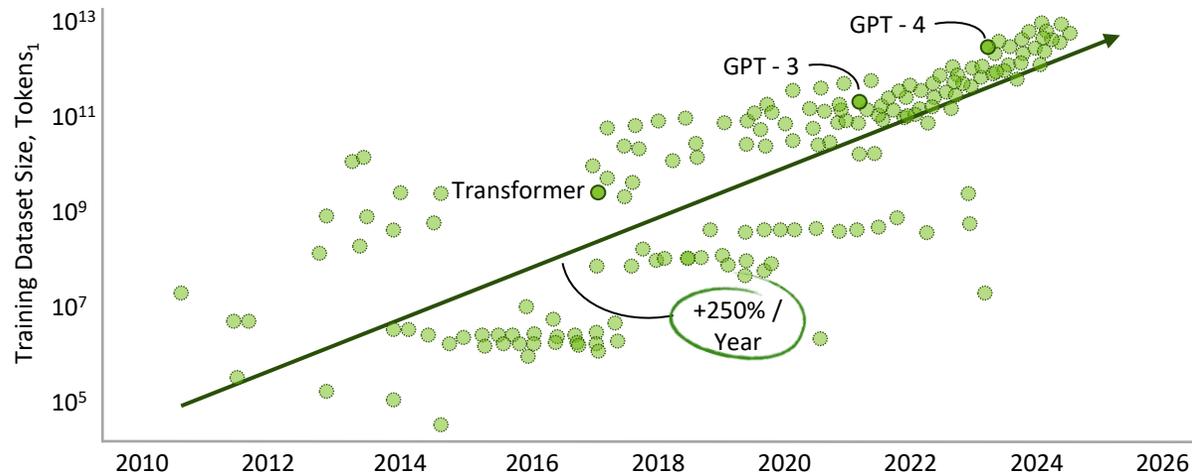
AI has already achieved unprecedented global penetration, and the biggest wave of growth is still ahead

I. AI is scaling faster than any technology in history, and it's only getting smarter, more capable, and more deeply embedded across every layer of the global economy

Years to Reach 100 MM Users

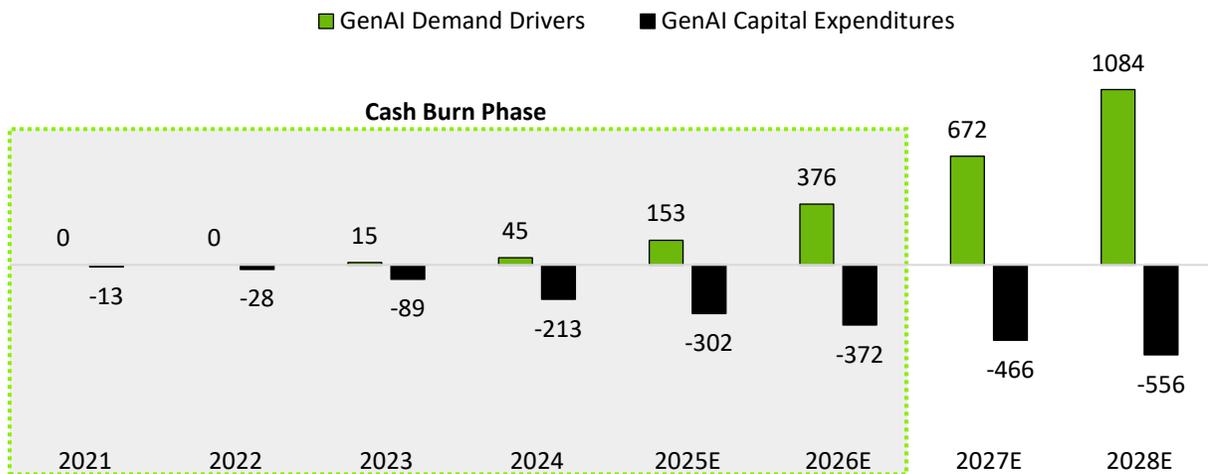


AI Model training Dataset Size by Model Release Year

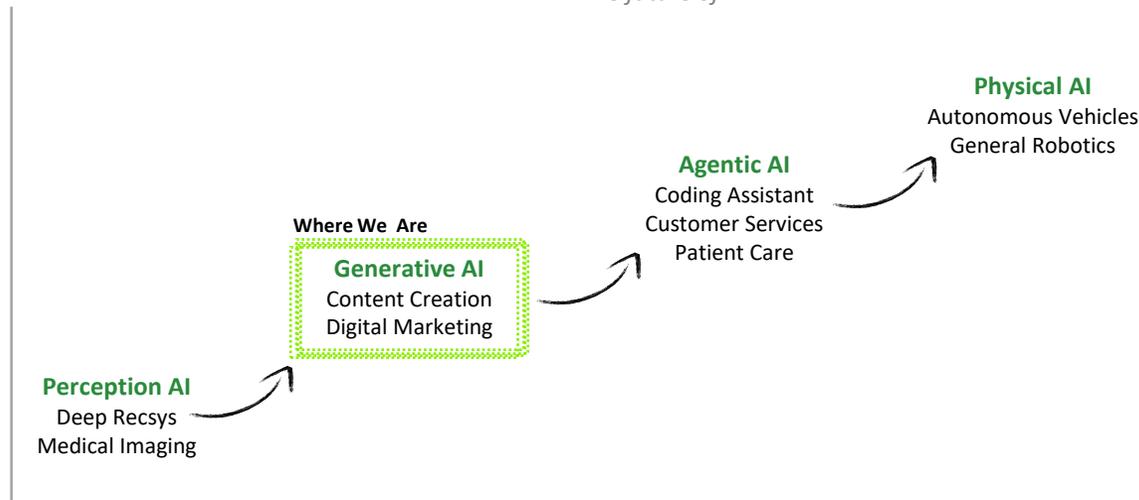


II. AI is still in its early stages, and today's investments are building the foundation for exponential future value creation across the global economy

GenAI Capex vs. Revenue (US\$ Billions)



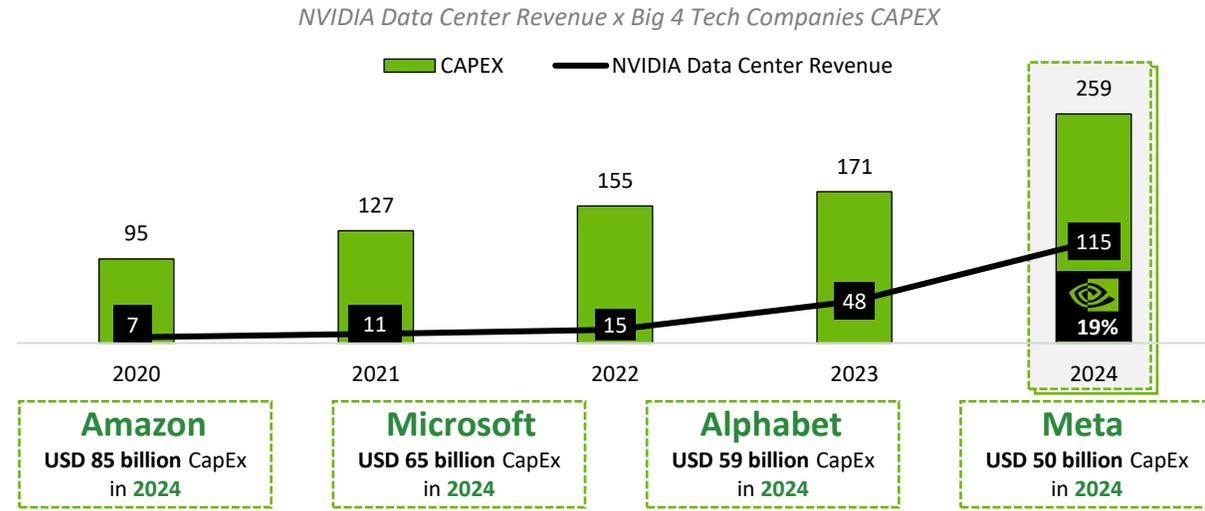
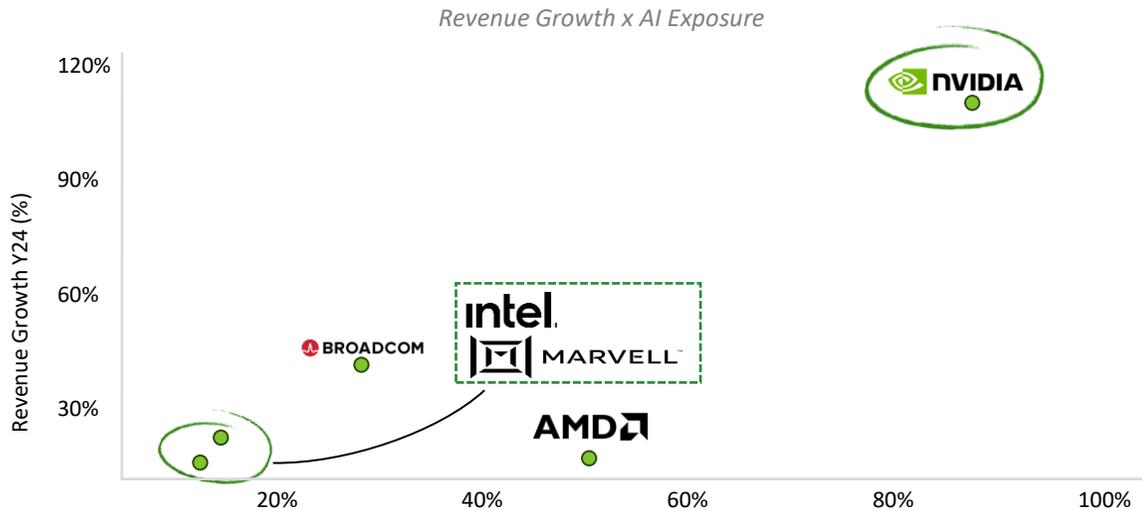
The future of AI



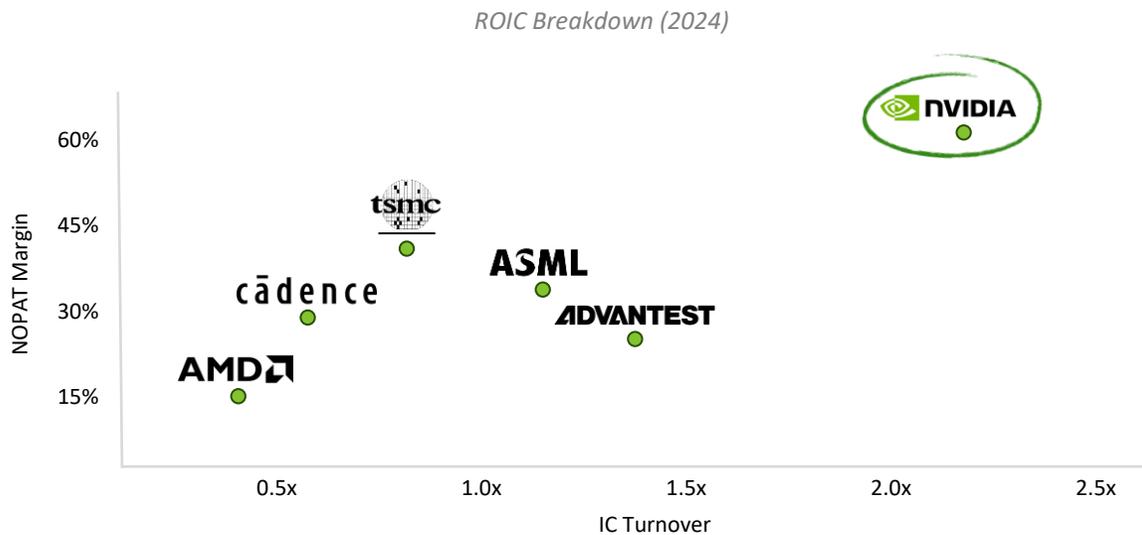
NVIDIA: Best Positioned to Capture the AI Future

No other company is as well positioned as NVIDIA to capture the AI boom and absorb hyperscaler CapEx

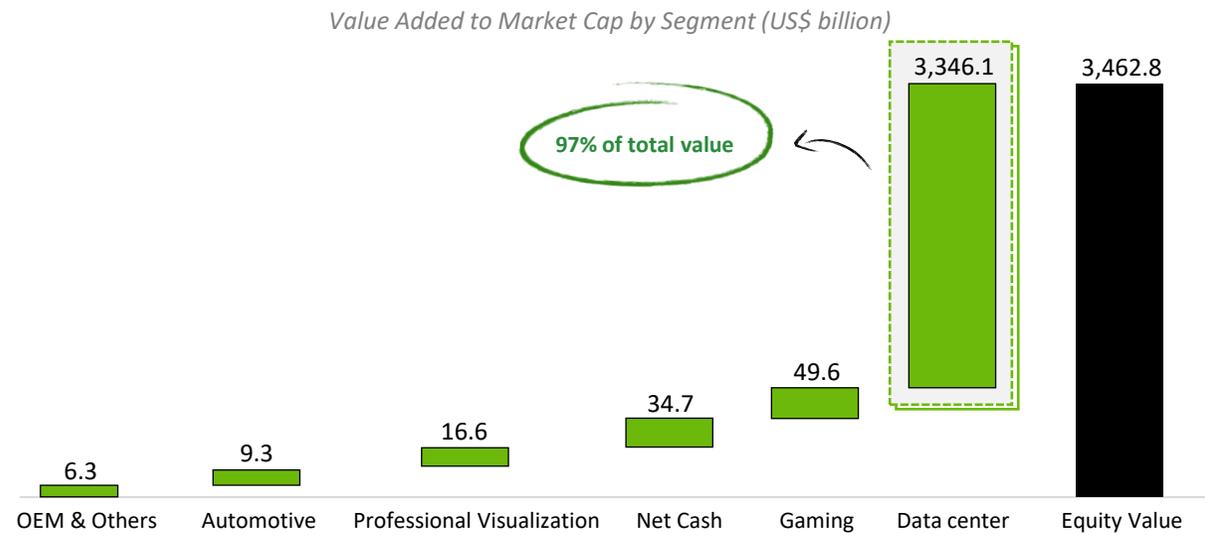
I. NVIDIA is the best-positioned company to capture the AI growth cycle, as evidenced by its strong revenue acceleration and increasing share of Big Techs' CapEx



II. No other company turns AI demand into economic value as effectively as NVIDIA



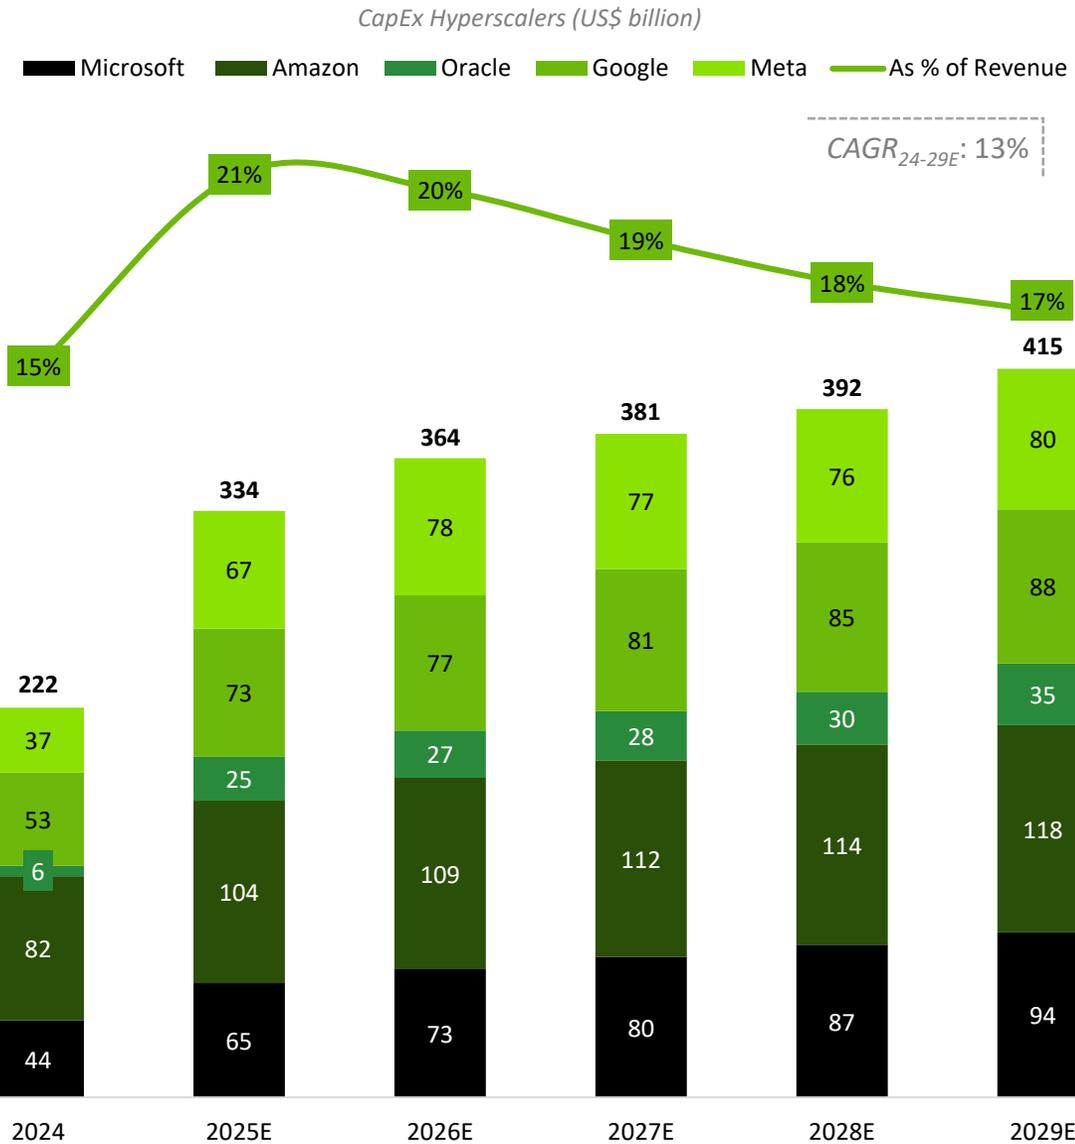
III. NVIDIA is priced as an AI company, data center drives nearly all of its value



Hyperscalers CapEx Forecast

AI is reshaping infrastructure priorities, hyperscalers are rapidly scaling CapEx to meet compute demand

I. Cloud and AI investments are accelerating, Hyperscalers CapEx to grow at 13% CAGR



II. Rising compute needs are triggering massive CapEx expansion across Hyperscalers

Why Are Hyperscalers Ramping Up CapEx?

Meta Zuckerberg indicated that to support LLaMA-5/6 and real-time inference, **Meta will have to triple its infrastructure**, pushing CapEx to even higher levels this decade

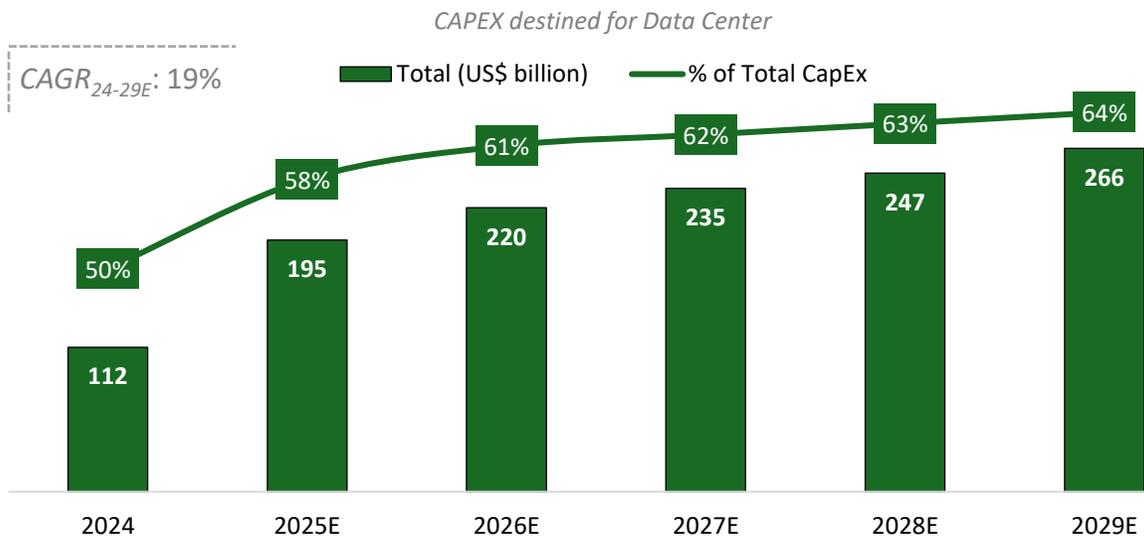
ORACLE Facing Microsoft–OpenAI’s “Stargate” super-cluster, Oracle must raise CapEx to keep OCI’s compute and latency competitive **while maintaining its sovereign-cloud edge**

Google CFO Anat Ashkenazi says Google’s capacity is tight, so the company will ramp up data-center **CapEx to close the AI demand gap**

amazon AWS’s multi-year **AI backlog already far outweighs its upcoming capital plan**, prompting CFO Brian Olsavsky to push for an unprecedented data-center expansion

Microsoft Demand for **Copilot and Azure AI pushes Microsoft to boost data-center CapEx**, with the “Stargate” super-cluster now setting its build scale

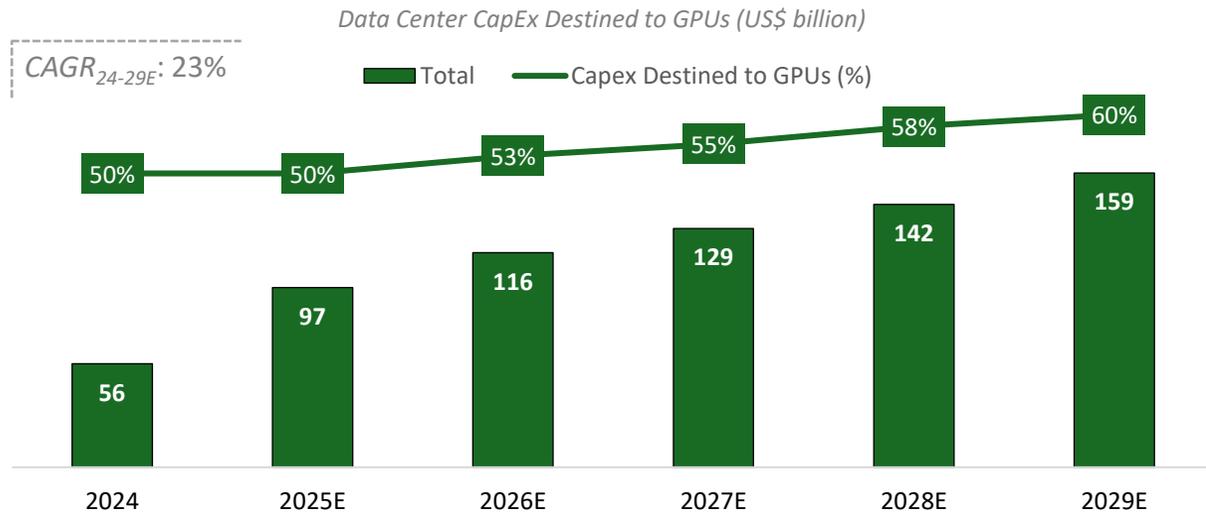
III. Most CapEx now flows into data centers, highlighting their role in Hyperscaler growth



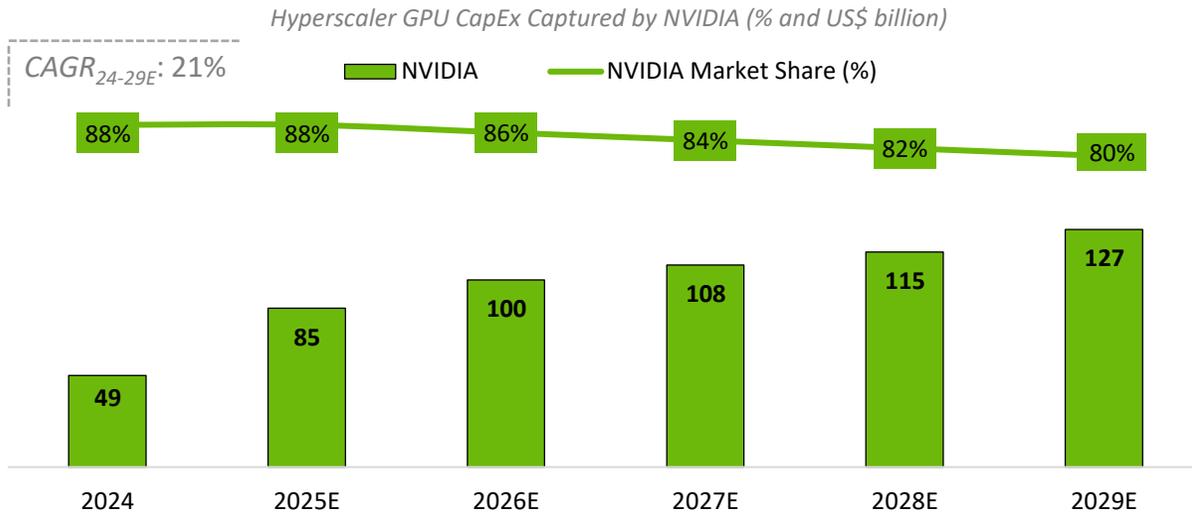
Data Center Revenue Build Up

NVIDIA is uniquely positioned to capitalize on rising AI accelerator demand, and our projections reflect that upside

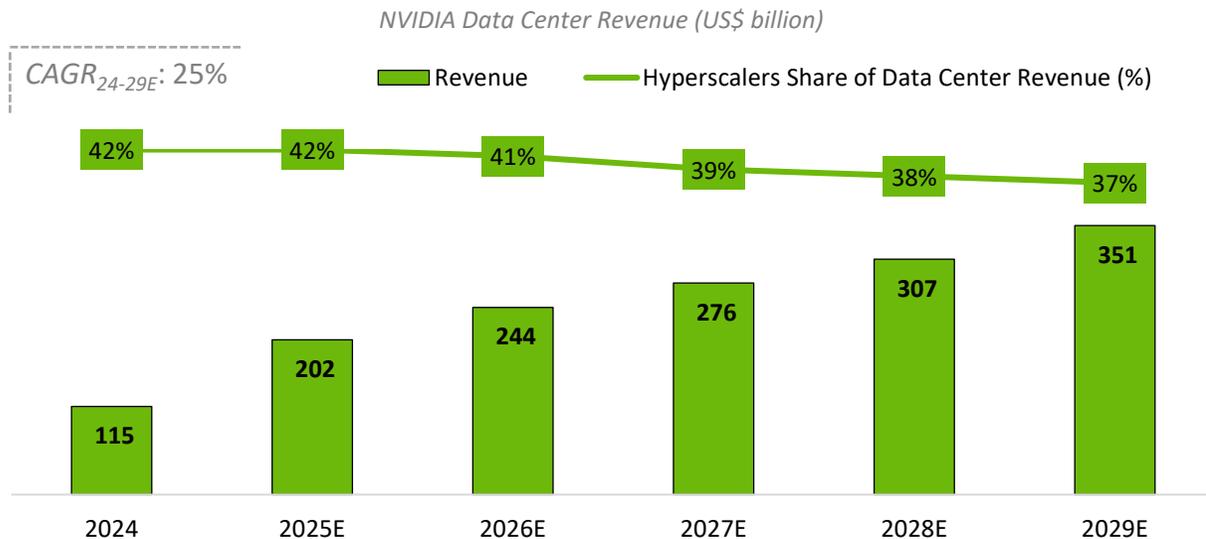
I. Hyperscalers are allocating more to GPUs, driven by the AI shift in data centers



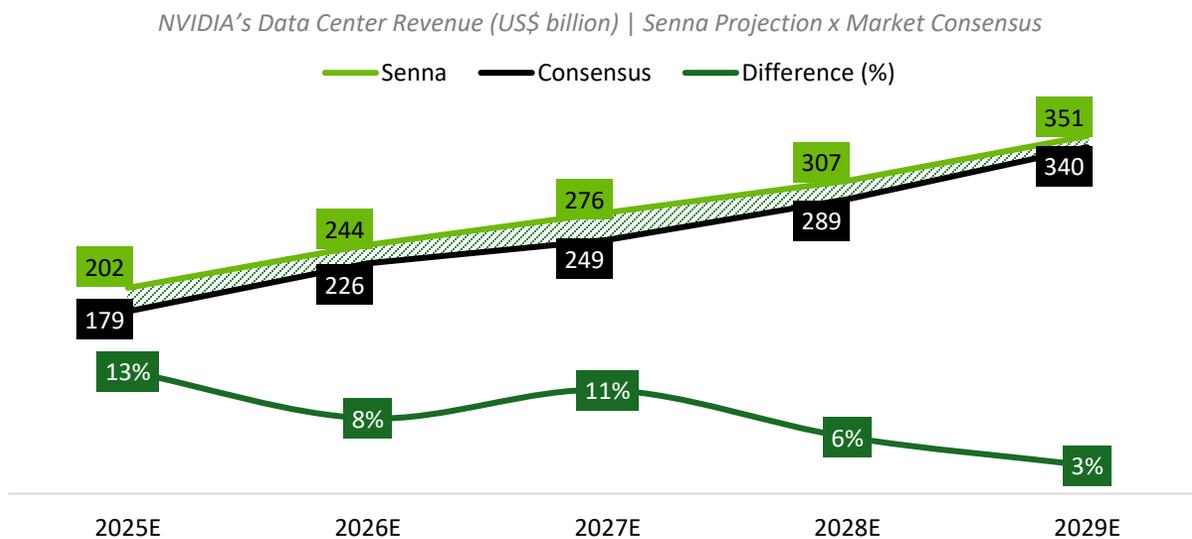
II. NVIDIA Leads Hyperscaler GPU CapEx, Even as Share Slightly Declines



III. As the hyperscalers' share declines, we backsolve to reach total revenue



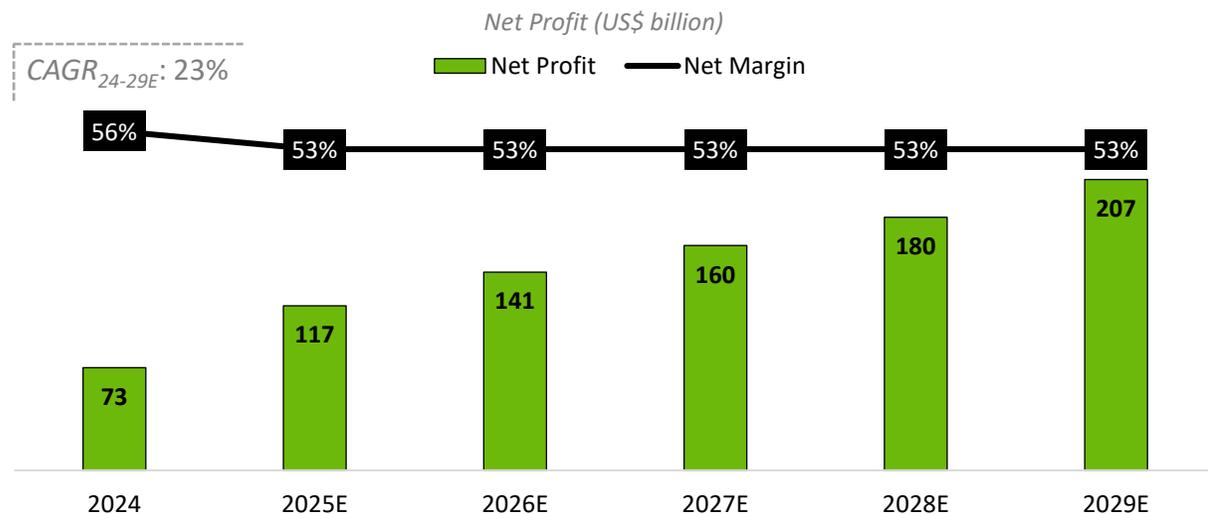
IV. Our projections exceed consensus, driven by accelerated AI infrastructure growth



Is our Desired IRR Viable? Yes!

For a company of such high quality, we believe there is ample room to allocate capital and generate a satisfactory return

I. We are exiting at a 2029 forward multiple equivalent to pre-AI years. Additionally, we conducted a reverse P/E analysis, which reinforced our decision within a satisfactory IRR



Entry & Exit: Taking a Conservative Stance



II. With a multiple premium to the sector well below historical averages—especially when compared to peers of significantly lower quality—we believe this is a compelling entry point

Required P/E Analysis

Senna's Revenue CAGR: 24%

IRR	Revenue CAGR						
	15%	20%	25%	30%	35%	40%	45%
Min: 15%	34x	27x	22x	18x	15x	12x	10x
Base: 20%	39x	31x	25x	21x	17x	14x	12x
Bull: 30%	44x	35x	29x	23x	19x	16x	13x

Average Premium:

5y: 52%

10y: 46%

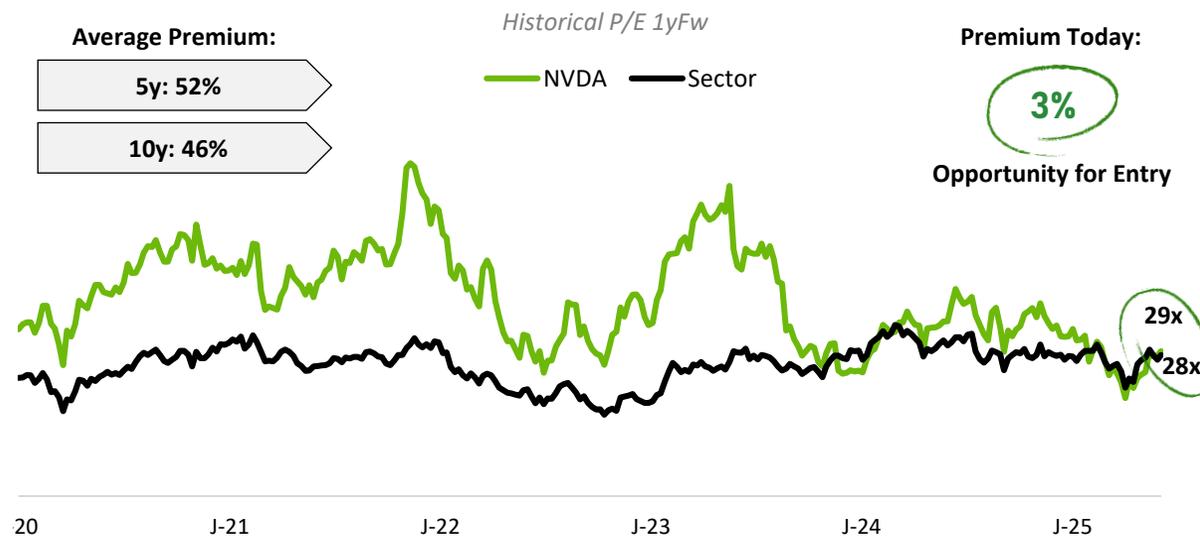
Historical P/E 1yFw

NVDA Sector

Premium Today:

3%

Opportunity for Entry



The P/E ratio has exceeded 25x in **83%** of the weeks over the **past 10 years**



Risks and Scenarios

From Opportunity to Threat: Mapping the IRR Upside and Stress-Testing the Core Risks

I. Mapping IRR Potential Across Bear, Base, and Bull AI Outlooks

IRR Scenarios Analysis

	Bear	Base	Bull
Marginal CapEx to Data Center (%)	70%	80%	90%
CapEx destined to Accelerators (%)	50%	60%	70%
NVIDIA GPU Market Share (%)	70%	80%	90%
Gross Margin	65%	75%	80%
Exit Multiple	15x	28x	35x
IRR	(14%)	23%	35%

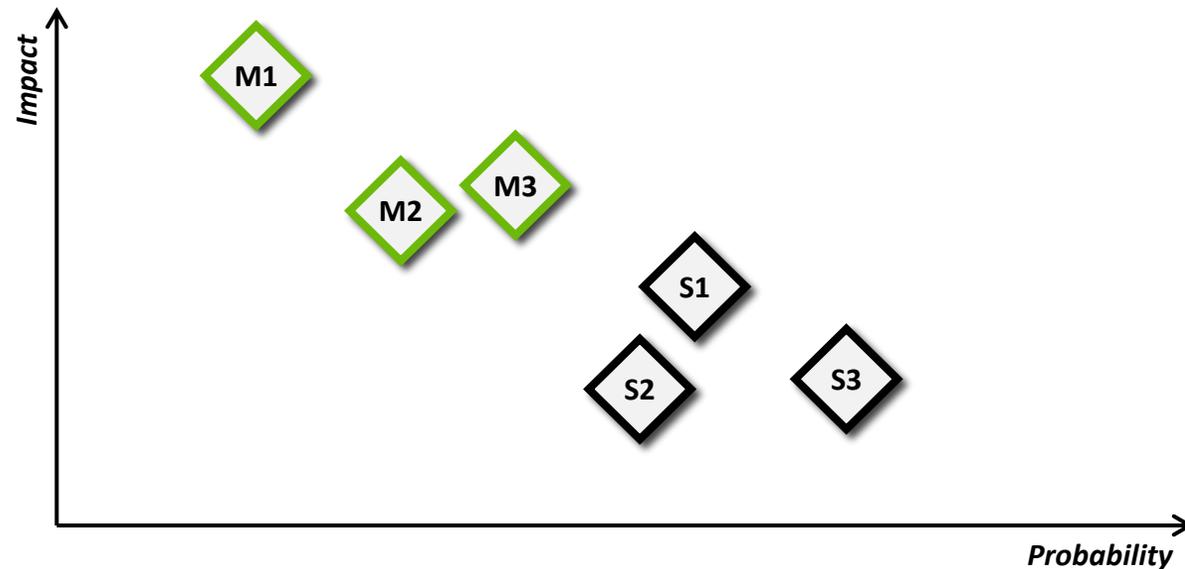
II. Stress Testing the Thesis: How Core Risks Could Break the Upside

IRR Scenario Analysis of Key Risks

ASICs CHIP	Scenario 1	Scenario 2	Scenario 3
NVIDIA Market Share (%)	55%	65%	75%
Gross Margin (%)	60%	65%	70%
IRR	3%	9%	15%
Slow Down in AI Investments	Scenario 1	Scenario 2	Scenario 3
Marginal CapEx to Data Center (%)	40%	50%	60%
CapEx destined to Accelerators (%)	35%	40%	45%
IRR	(4%)	(2%)	8%

III. Risk Matrix: Identifying the Most Impactful Threats to NVIDIA's Upside

Risk Matrix



Sector	Macro
S1 Asics Chip	M1 China Invades Taiwan
S2 Hardware Competition	M2 Economic Deceleration Risk
S3 DeepSeek V2 Efficient	M3 Slowdown in AI Investments



Case in a Nutshell

NVDA: We are Long!

NVIDIA's ecosystem creates a competitive advantage

ROIC 140%

Above WACC

+85% Market Share

In GPUs

Management and human capital as a key differentiator

16 Years

Average executive tenure

Jensen owns ~3%

of NVIDIA

Nvidia is well positioned to capture the AI market

US\$ 351 Billion

2029E Data Center Revenue

25% CAGR_{25E-29E}

In Data Center Revenue

Entry Multiple: 29x P/E

23.2%

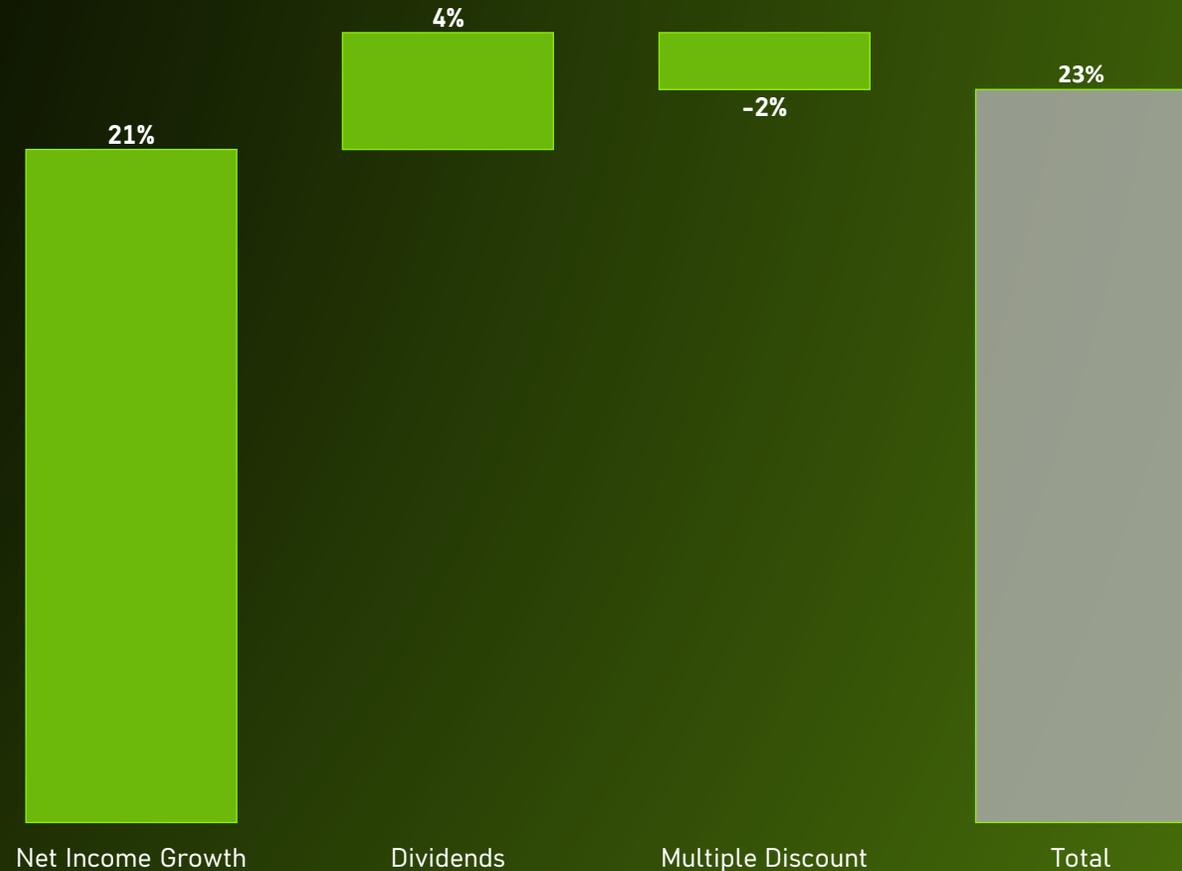
IRR

10.3%

Cost of Equity

Exit Multiple: 28x P/E

IRR Breakdown



Net Income Growth

Dividends

Multiple Discount

Total

Impact on our thesis

Key insights for thesis returns

Data center investments

New releases products

New market trends

Management



We are open for Q&A!



Presentation Index

Long NVIDIA

NVIDIA at a glance

Ecosystem

GPUs:NVIDIA's Main Product

MOAT I: Ahead in GPU Technology

MOAT II: NVLink Scale Capabilities

MOAT III: Software Switching Costs

Competitive Advantage

Management

Skilled Human Capital

Capture AI Market

It's Still Early Days for AI

NVIDIA: Best Positioned to to Capture the AI Future

Hyperscalers CapEx Forecast

Data Center Revenue Build-Up

Valuation

Is our Desired IRR Viable? Yes!

Risks and Scenarios

Case in a Nutshell

Our Conversations:

Call 1 – Guilherme Amaral | Kinea

Call 2 – Jon Y | Asianometry

Call 3 – José Moreira | PRAGMA

Call 4 – Gabriel Oliveira | Verde Asset

Call 5 – João Pedro Freitas | Mainú

Call 6 – Ilan Crohmal | Occam

Call 7 – Karina Fugita | GeoCapital

Call 8 – Caio Bessa | M Square

Call 9 – Lucas Dias | Aster Capital

Call 10 – Igor Fernandes | AZ Quest

Call 11 – Adriano Marques | Ascenty

Appendix Index

Data Center Investments

Stargate: The New Race for AI Infrastructure
Data center Demand

Multiple Rationale

- I. Rationale Behind Our Exit Multiple Assumption
- II. Rationale Behind Our Exit Multiple Assumption

Risks

DeepSeek Day
Jevons Paradox

- I. DeepSeek x GPT-4o. Who is Better?
- II. DeepSeek x GPT-4o. Who is Better?
- III. DeepSeek x GPT-4o. Who is Better?

And in the end, who was Better, DeepSeek or GPT-4o?

Can Energy be a risk for AI?

- I. Risk: Slowdown in AI Investments
- II. Risk: Slowdown in AI Investments

Value Generation With AI

Build-Up Sanity Check | Goldman Sachs

Build-Up Sanity Check | Dell'Oro

Build-Up Sanity Check | AMD

Build-Up Sanity Check | McKinsey

I. AI Bubble

II. AI Bubble

Key-man Risk

How Could This Investment Lose Money?

NVIDIA

I. Robotics

II. Robotics

Current Overview

ARM

Ampere Computing

What Happened in NVIDIA Margin in 2022

Value Added

Shareholder Structure

Performance Metrics & Goals

Inputs

Marginal CapEx

How Marginal CapEx Reflects in our Model?

GPU's Share

How GPU's Share reflects in our Model?

Hyperscalers % of Revenue

Ex-Hyperscalers % of Revenue

How Hyperscalers Share Reflects in our Model?

Competition

AMD at a Glance

ASICs at a Glance

Unconvincing Playbook

AI Rush Isn't Making It Any Easier

How Does This Translate into our Model?

What if they introduced better software than

CUDA?

Is NVIDIA the Best Investment in

Semiconductors?

ROIC Comparison

Geopolitical

Chinas as an Opportunity

China x Taiwan

Why Taiwan is so Important to NVIDIA? TSMC
A Powerful Alliance Between Chip Design and
Manufacturing Leaders
Trade War: Direct Impact on NVIDIA

DCF

Why not DCF?

FCFF Valuation

FCFE Valuation

DCF Scenarios

WACC

Perpetuity Growth

Senna x Consensus

Where we Disagree with the Consensus?

Supply Analysis

Gross Margin

EBIT Margin

Net Margin

Sequoia Paper

Model

NVIDIA's Revenue

NVIDIA Total Revenue 2010-2024

NVIDIA Data Center Revenue 2010-2024

NVIDIA Projected Total Revenue

NVIDIA Projected Data Center Revenue

NVIDIA Revenue per Segment (%)

NVIDIA Revenue per Segment (US\$)

NVIDIA Projected Margin

Phase 2 | Presentation

Data Center Investments

Stargate: The New Race for AI Infrastructure

With up to \$500B in investments, OpenAI, SoftBank, and Oracle are building the global backbone of generative intelligence

Abilene, Texas



Stargate is a **\$500B initiative by OpenAI, SoftBank, Oracle, and MGX to build AI-focused data centers** across the globe. It aims to secure large-scale compute for training advanced models and supporting national AI strategies, starting in Texas and the UAE

How Nvidia can win with Stargate?

1. Massive GPU sales

Stargate will require hundreds of thousands of high-end AI chips, and NVIDIA is the dominant supplier. Oracle, a key partner in Stargate, has already committed to purchasing **up to \$40 billion** worth of NVIDIA GPUs to power the initial U.S. data centers

2. Software and Licensing

Beyond hardware, NVIDIA earns recurring revenue from its AI frameworks, and proprietary tools, which are essential for training and deploying models at scale. **Stargate's reliance on these tools further locks in revenue**

Sam Altman, CEO at OpenAI

"We believe Stargate will become the foundational infrastructure for the future of artificial intelligence, not just to train more powerful models, but to ensure they are developed safely, aligned with human values, and accessible to the world's democracies"

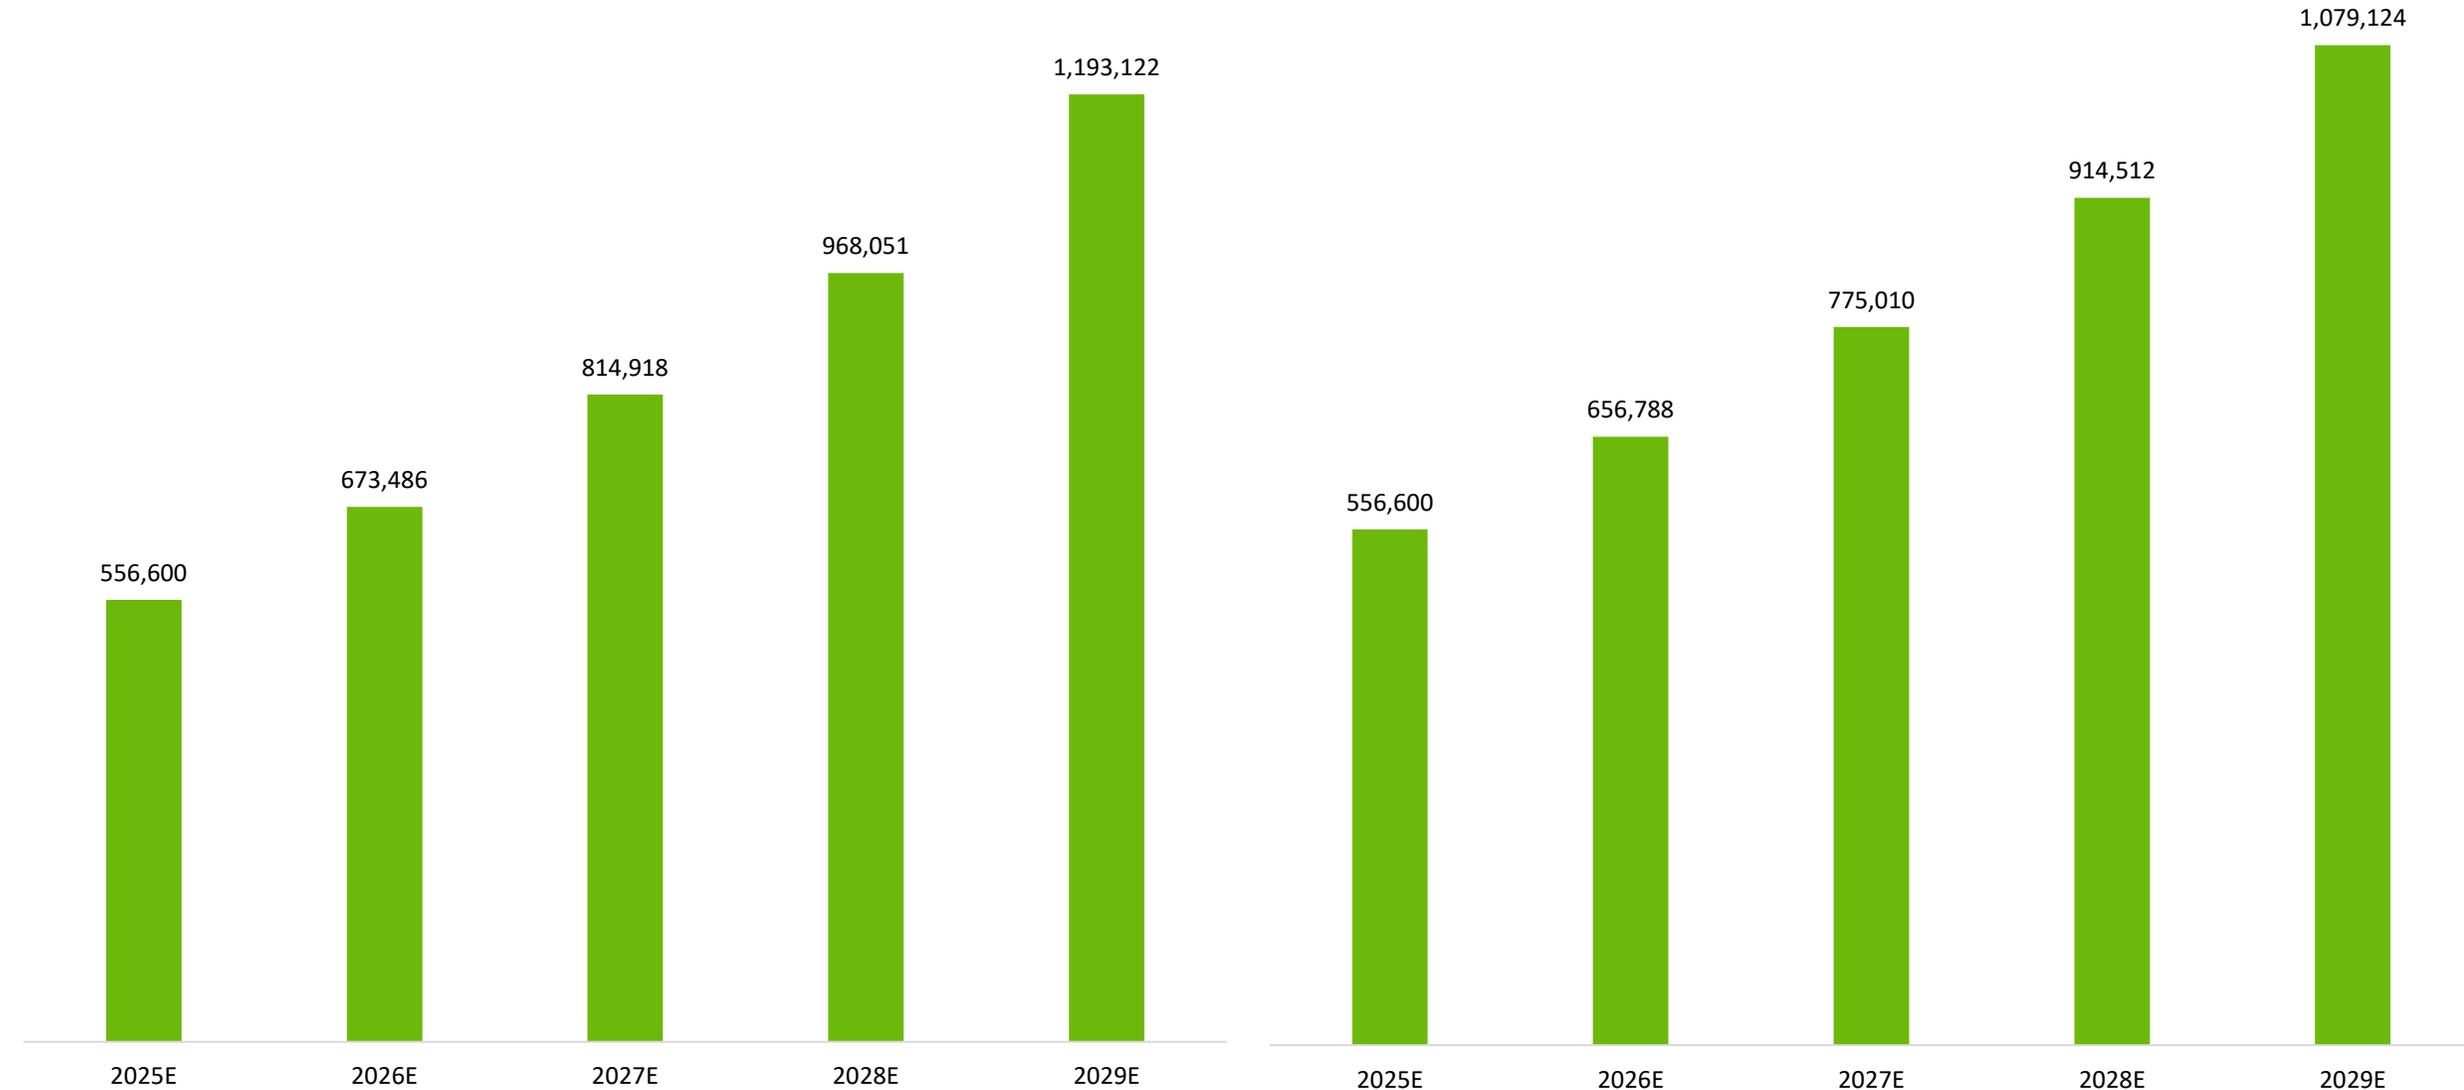


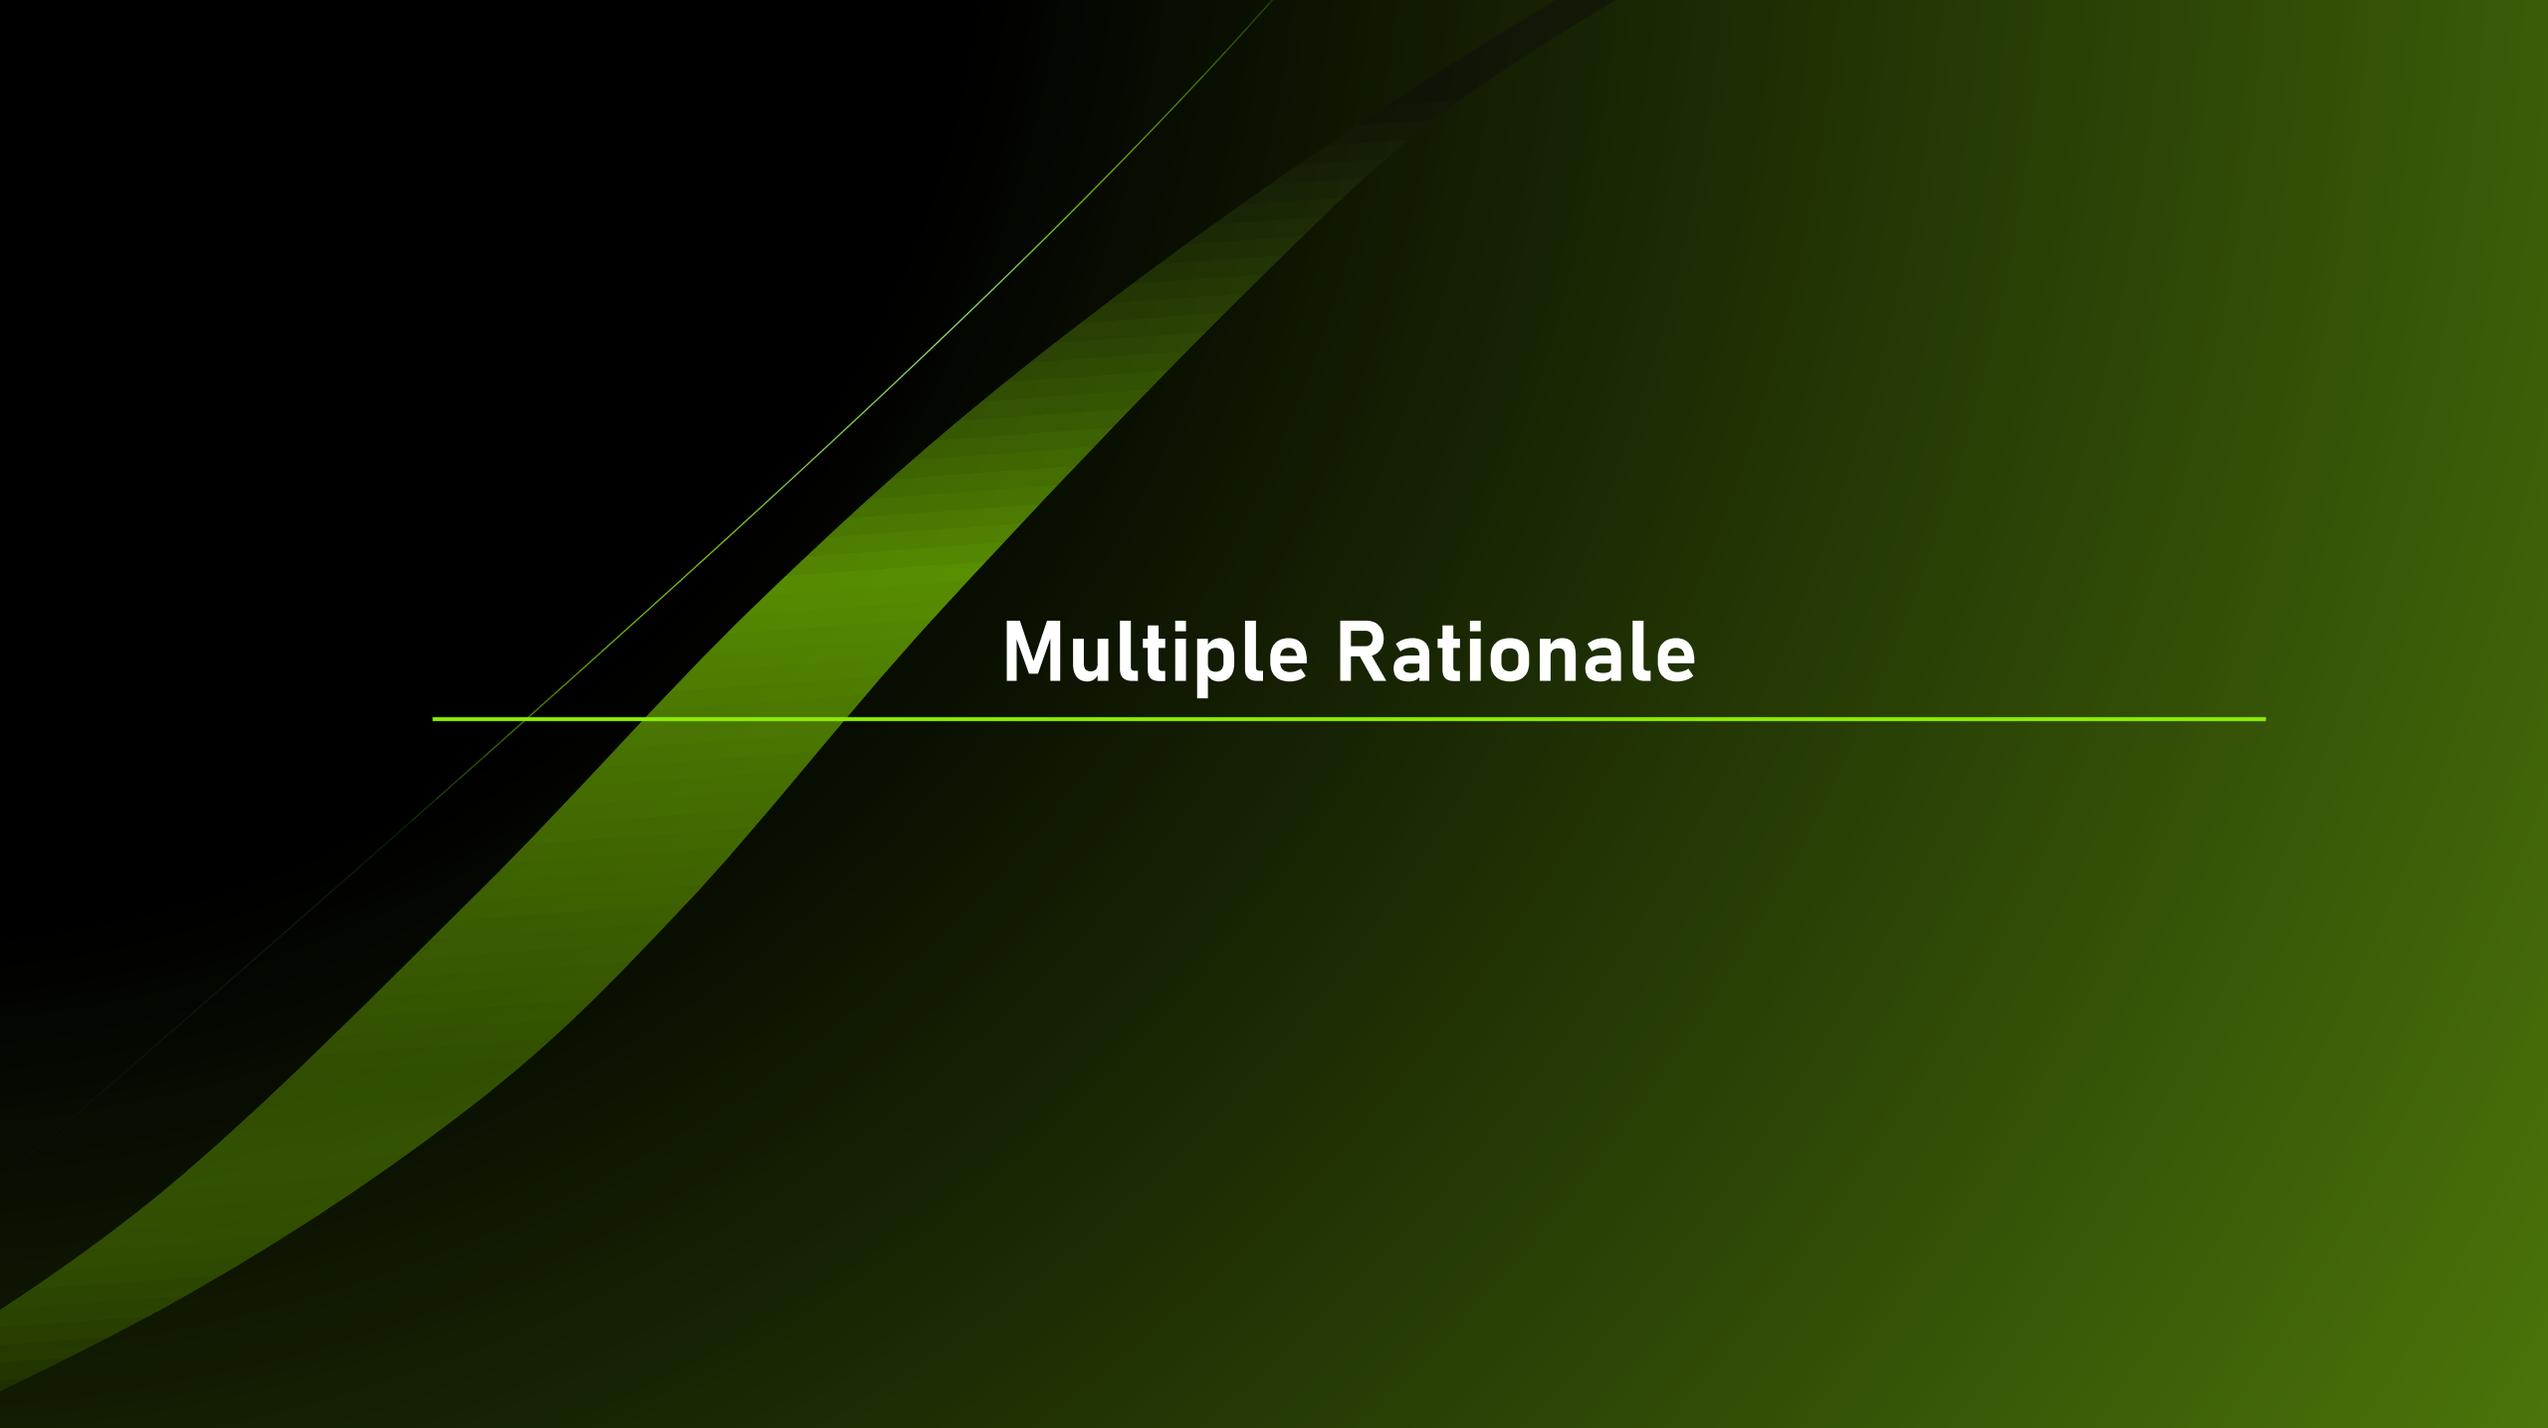
Data Center Demand

Market valuations increasingly hinge on expectations of sustained, AI-driven data center demand

Dell'oro projections

McKinsey Projections

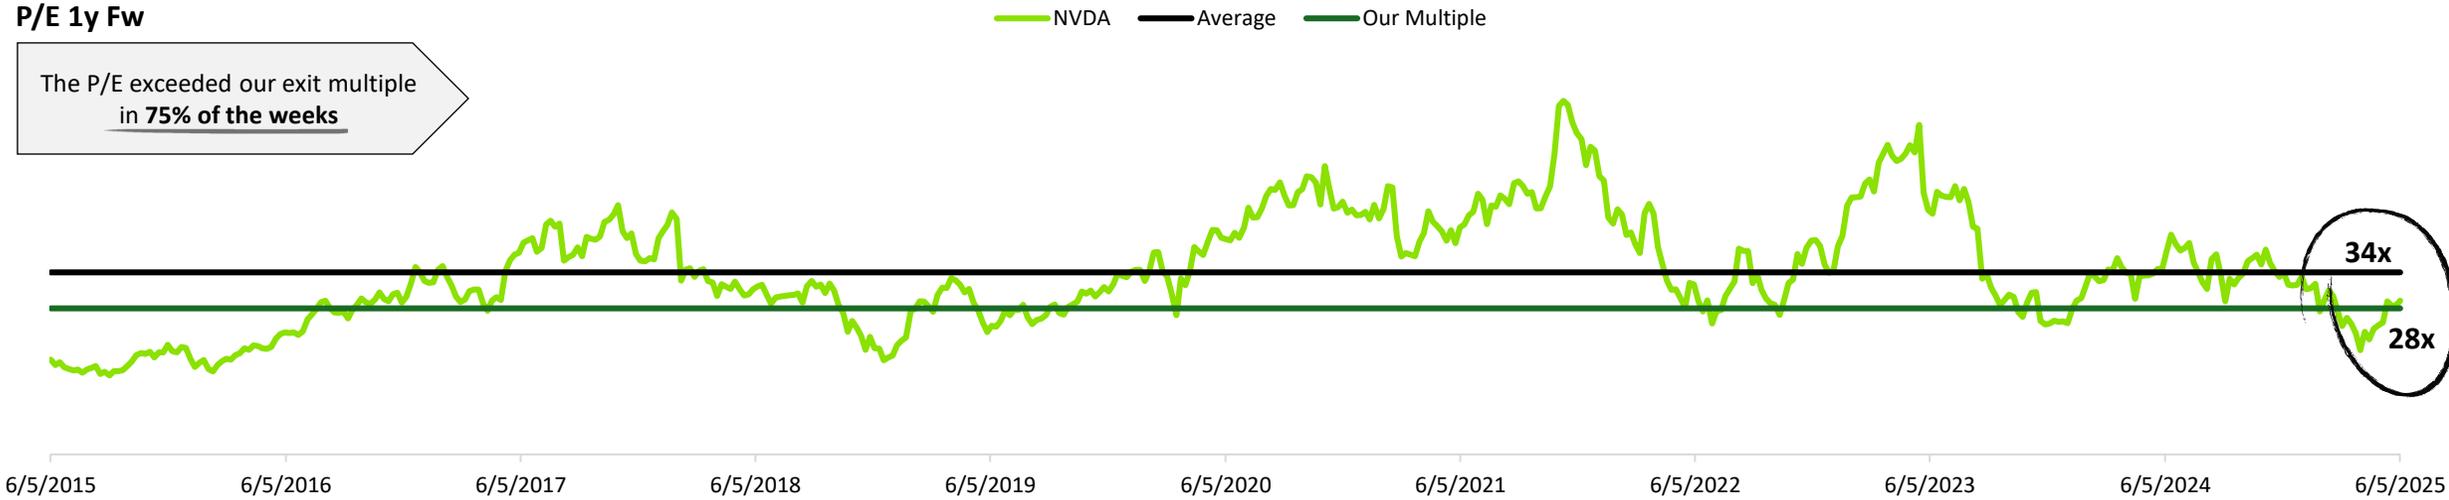




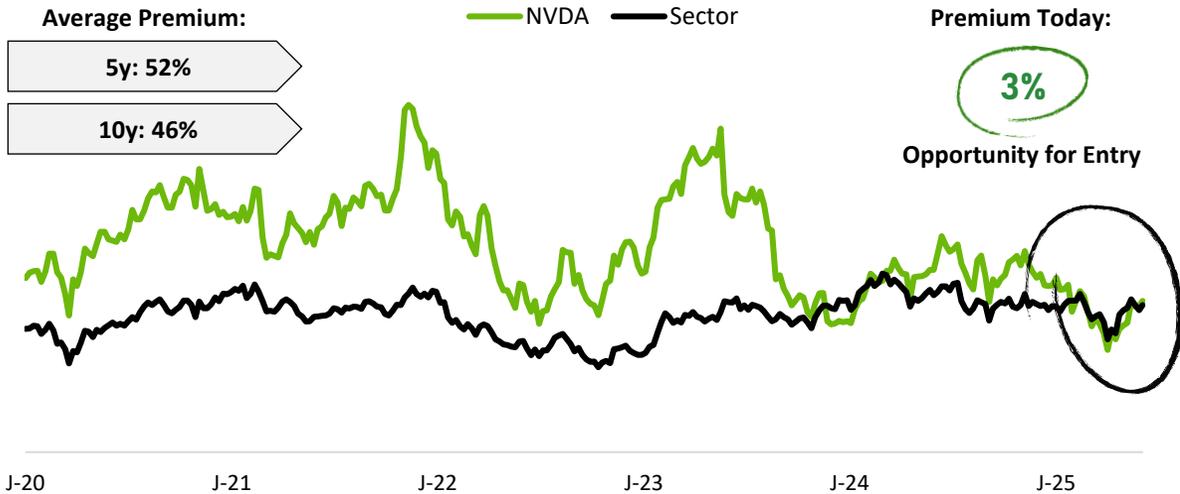
Multiple Rationale

I. Rationale Behind Our Exit Multiple Assumption

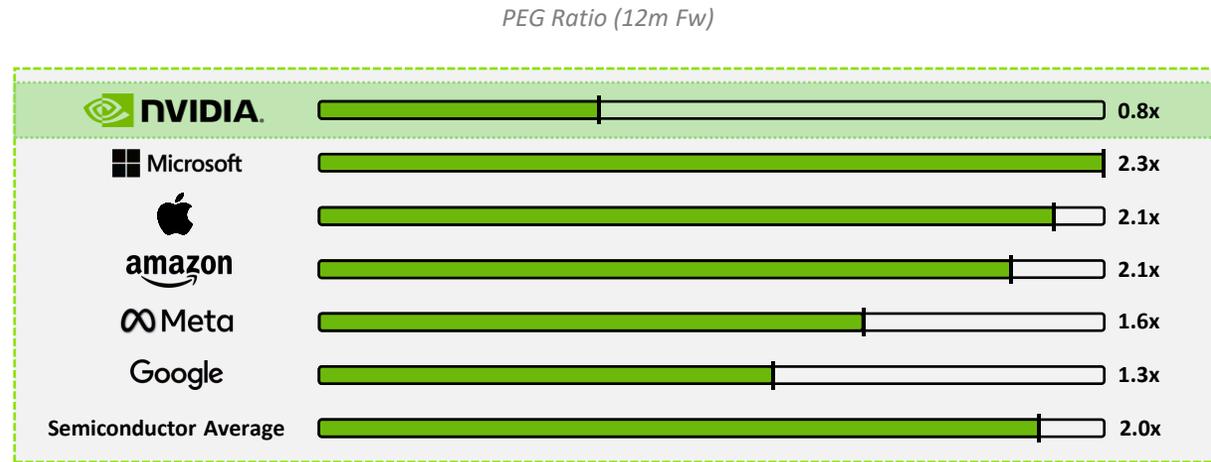
I. Historical data shows the current multiple is well below previous averages



II. The current premium to the sector is well below historical averages



III. NVIDIA's exceptional earnings growth justifies its P/E multiple



II. Rationale Behind Our Exit Multiple Assumption

Although we believe the AI boom still has much room to unfold, we used a conservative assumption for our multiple, based on a normalized moment for the company

10y	(2015~2025)							
Average	34.5							
5y	(2015~2020)	(2020~2025)						
Average	30.9	39.6						
4y	(2015~2018)	(2018~2021)	(2021~2024)					
Average	28.5	37.4	40.2					
3y	(2015~2017)	(2017~2019)	(2019~2021)	(2021~2023)	(2023~2025)			
Average	27.7	31.9	39.7	42.2	36.3			
2y	(2015~2016)	(2016~2017)	(2017~2018)	(2018~2019)	(2020~2021)	(2021~2022)	(2023~2024)	
Average	21.8	30.6	33.8	29.3	45.4	42.1	38.3	

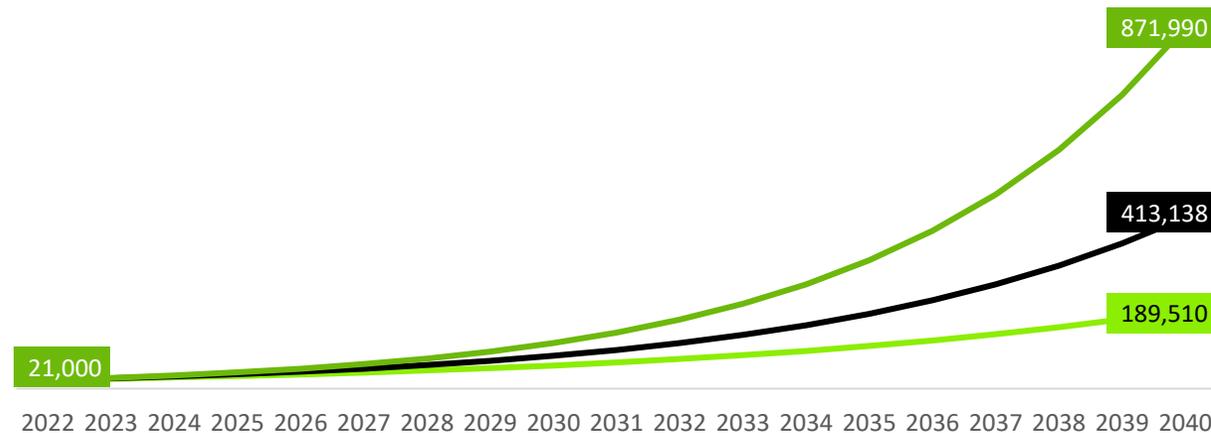
We chose a multiple from before the AI boom

Only in one period, among all observed, was the multiple lower than the one we selected

In addition to using a conservative multiple, we believe the company may still unlock new revenue streams — such as through Robotics — which could trigger a new wave of growth

Robotics Industry (US\$ million)

— Bear — Base — Bull



Monetization Opportunities through Robotics

NVIDIA is betting on robotics as its next growth engine, monetizing through Isaac (software licensing) and Omniverse + Cosmos (realistic and scalable simulation). **These platforms expand its reach beyond hardware, targeting a trillion-dollar market.**

Isaac Platform: A comprehensive framework for the development, simulation, and deployment of autonomous robots. Revenue comes from software licensing and integration.

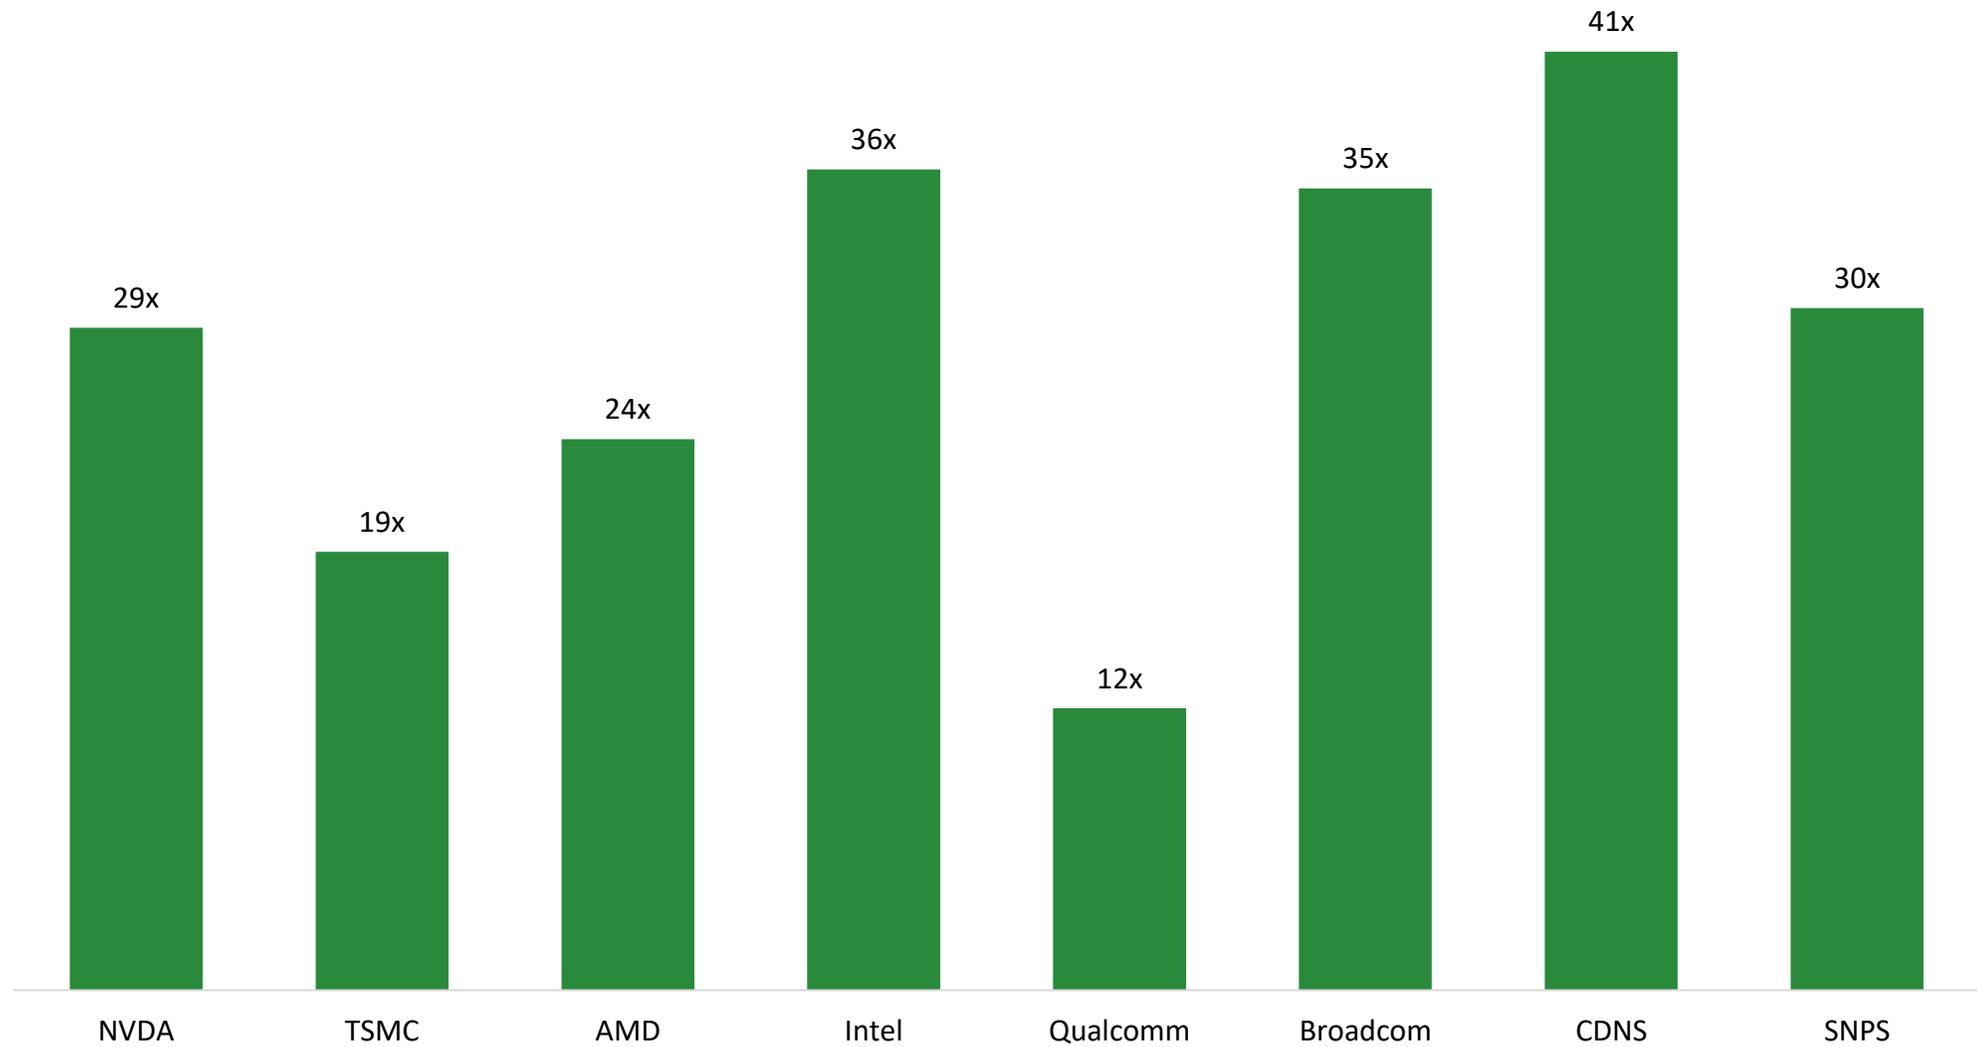
Simulation with Omniverse: High-fidelity environment simulation for virtual robot training, powered by a precise physics engine and integrated with the NVIDIA Cosmos platform for scalable deployment.

Jensen Hung, Nvidia's Founder & CEO

"I think this is likely to be the next multi-trillion-dollar industry."



Sector Multiples



The image features a dark green background with a lighter green curved band that starts from the bottom left and curves upwards towards the top right. A thin, light green horizontal line spans across the middle of the image, intersecting the curved band. The word "Risks" is written in white, bold, sans-serif font, centered horizontally and positioned just above the horizontal line.

Risks

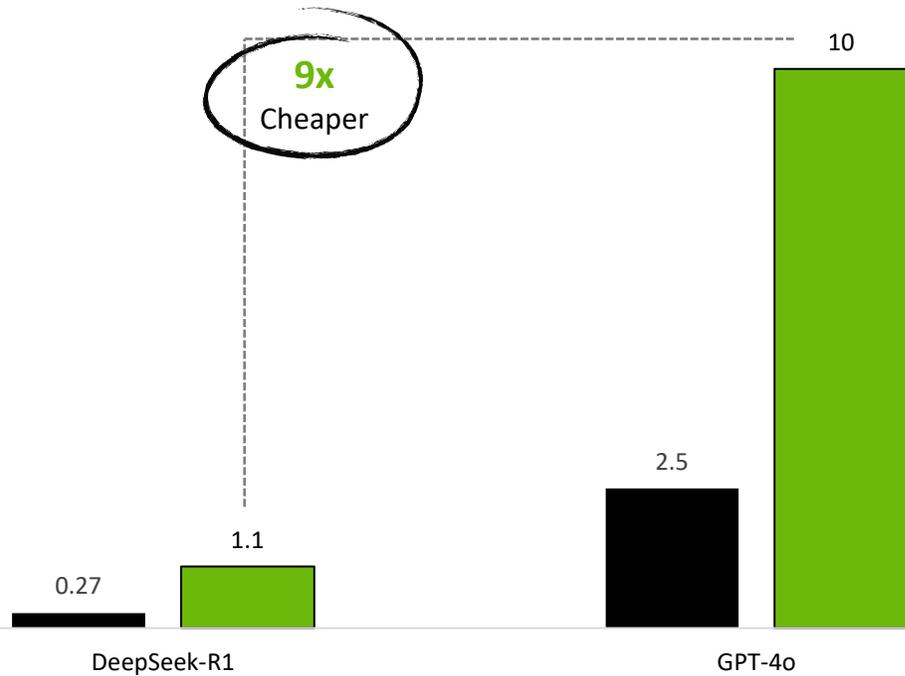
DeepSeek Day

DeepSeek's rapid ascent challenges incumbents, raising questions about model openness, security, and competitive stability

I. DeepSeek's 9x lower cost exposes Big Tech's monetization fragility

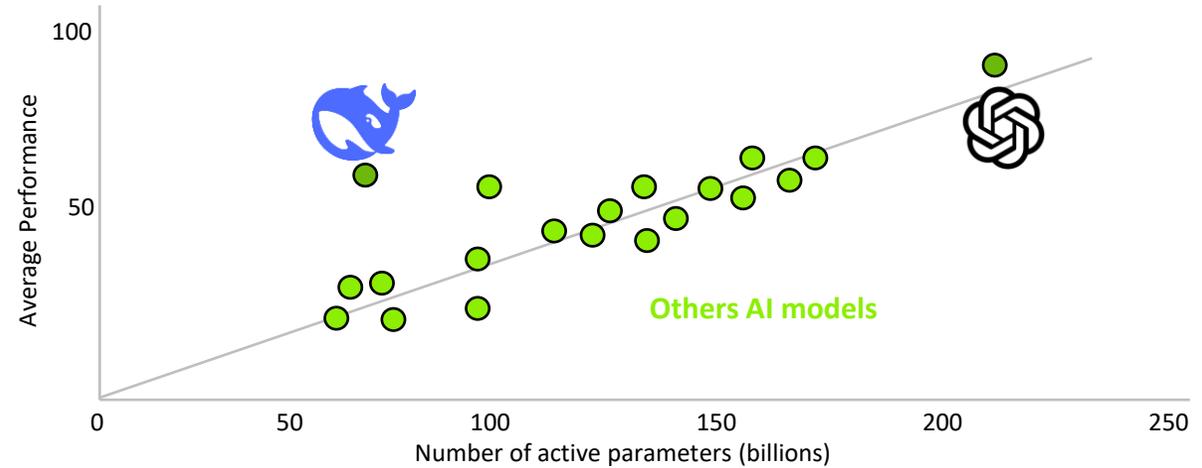
Price per million token (US\$)

■ Input Cost ■ Output Cost



II. DeepSeek breaks the parameter-performance tradeoff

Performance vs. Activated Parameters



III. DeepSeek fizzled due to hype, mistrust, and no global play

Inflated Benchmarks

Initial performance was overhyped, based on self-reported metrics with little third-party validation

Lack of Ecosystem

No strong APIs, enterprise integration, or developer community like OpenAI or Anthropic

Trust Issues

Concerns around alignment, safety, and transparency limited adoption outside China

No Global Strategy

Without a clear international presence or monetization plan, DeepSeek remains locally constrained

<https://www.the-sun.com/tech/13396198/nvidia-most-valuable-company-loses-billions>

THE Sun

CHIPS ARE DOWN:
World's most valuable company making AI chips loses \$600bn in biggest market loss EVER after China's DeepSeek launch



Jevons Paradox

DeepSeek V2 and NVIDIA's GPU efficiency lower inference costs, but in Jevons' logic, this only accelerates total GPU demand

I. Explaining the Jevons Paradox: Why Efficiency Can Drive More Consumption

William Javons Observation

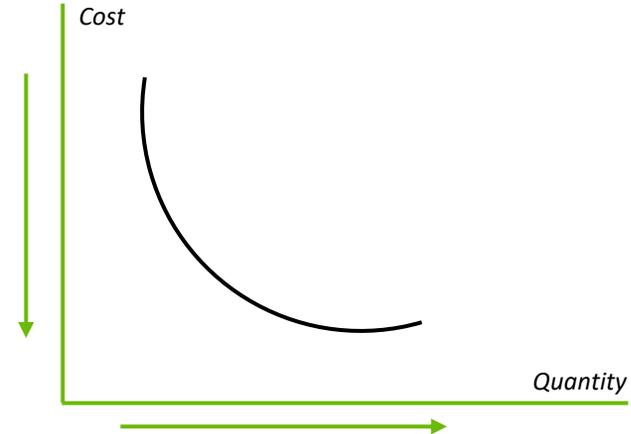
William Stanley Javons, 1865



Coal got more efficient

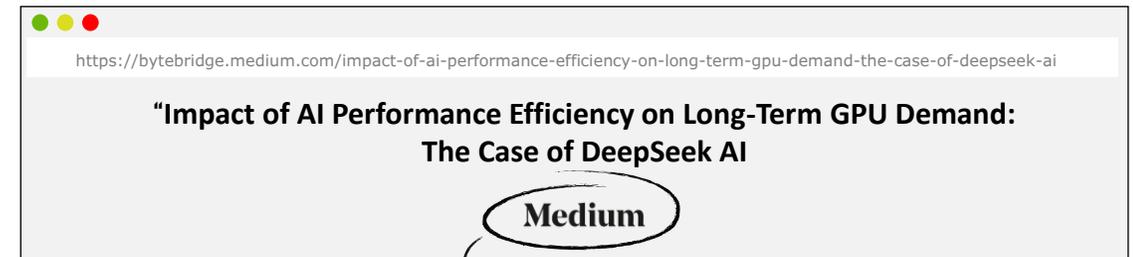
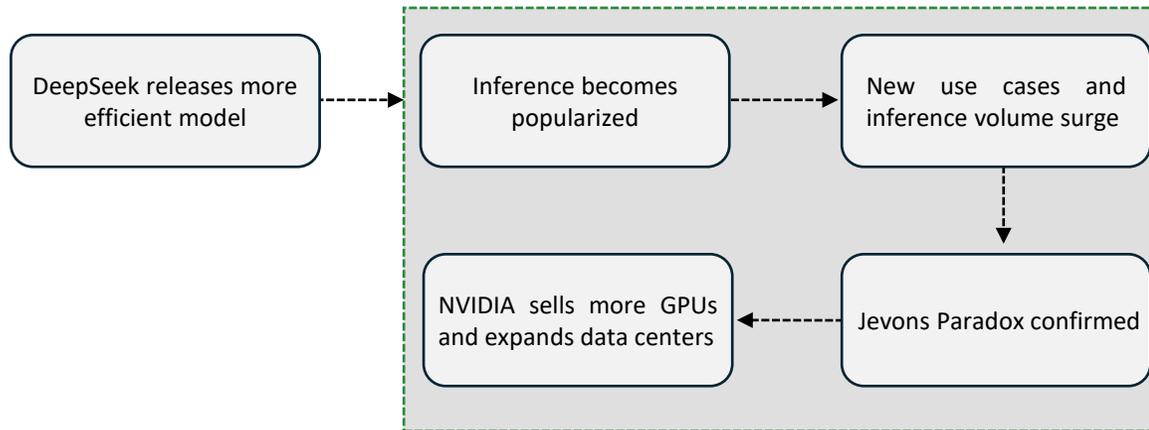
Jevons observed during the Industrial Revolution that, **after coal became more efficient to use, its total consumption increased sharply**. The lower cost and greater efficiency led to widespread adoption across factories, transport, and industry — accelerating overall demand

Javons Paradox



The **Jevons Paradox** shows that **increasing efficiency can lead to higher overall consumption**. As using the resource becomes cheaper and more accessible, demand grows — often surpassing the efficiency gains

II. DeepSeek V2's efficiency could lower inference costs, potentially expanding AI adoption and increasing GPU demand



"The adoption of energy-efficient AI models like DeepSeek AI is expected to **influence the long-term demand for GPUs in various industries**"

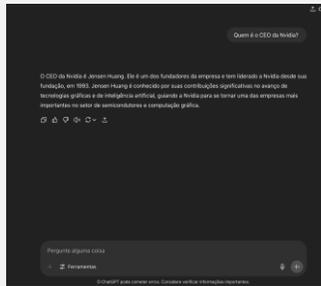


DeepSeek x GPT-4o. Who is Better?

We did this analysis in three steps. **Easy questions**, medium question and hard questions

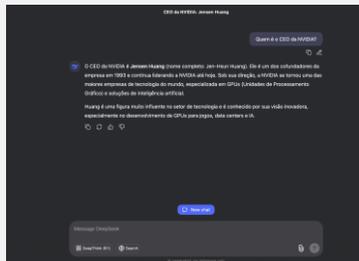
Question 1: Quem é o CEO da NVIDIA?

GPT



Time: 2 seconds and 36 thousandths

DeepSeek

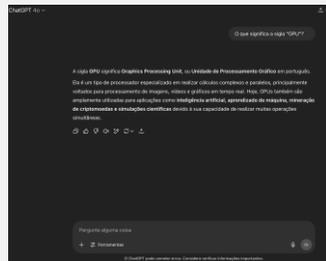


Time: 4 seconds and 85 thousandths

Time: GPT Better
Response: Equal

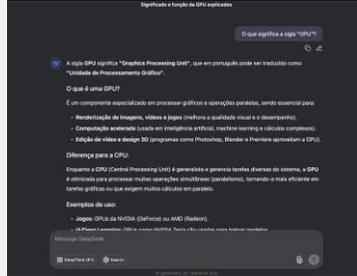
Question 2: O que significa a sigla "GPU"?

GPT



Time: 1 seconds and 58 thousandths

DeepSeek

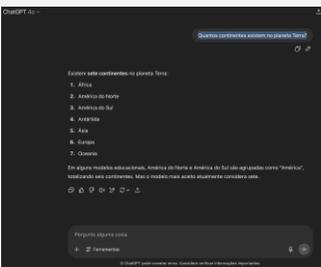


Time: 4 seconds and 76 thousandths

Time: GPT Better
Response: DeepSeek Better

Question 3: Quantos continentes existem no planeta Terra?

GPT



Time: 1 seconds and 44 thousandths

DeepSeek

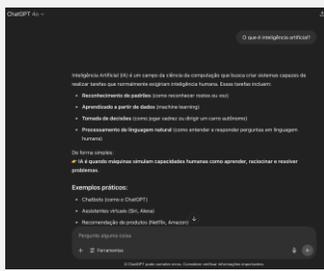


Time: 4 seconds and 85 thousandths

Time: GPT Better
Response: Equal

Question 4: O que é inteligência artificial?

GPT



Time: 1 seconds and 14 thousandths

DeepSeek



Time: 4 seconds and 81 thousandths

Time: GPT Better
Response: Equal

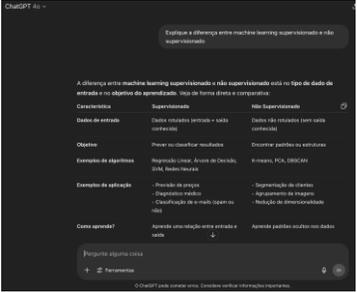


DeepSeek x GPT-4o. Who is Better?

We did this analysis in three steps. Easy questions, **medium question** and hard questions

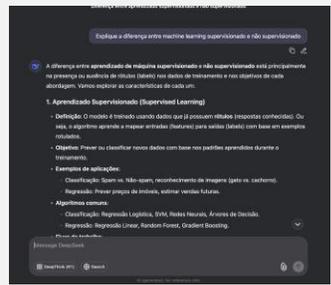
Question 1: Explique a diferença entre machine learning supervisionado e não supervisionado

GPT



Time: 1 seconds and 84 thousandths

DeepSeek

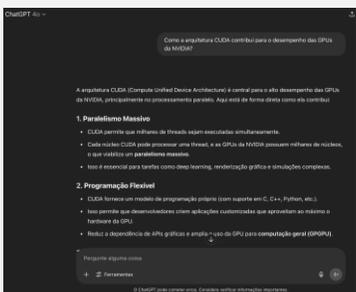


Time: 3 seconds and 91 thousandths

Time: GPT Better
Response: GPT Better

Question 2: Como a arquitetura CUDA contribui para o desempenho das GPUs da NVIDIA?

GPT



Time: 1 seconds and 66 thousandths

DeepSeek



Time: 4 seconds and 23 thousandths

Time: GPT Better
Response: DeepSeek Better

Question 3: Compare os modelos de linguagem GPT-4 e LLaMA-2 em termos de arquitetura

GPT



Time: 1 seconds and 38 thousandths

DeepSeek

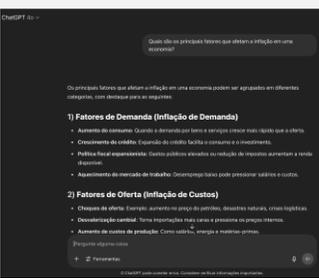


Time: 4 seconds and 67 thousandths

Time: GPT Better
Response: GPT Better

Question 4: Quais são os principais fatores que afetam a inflação em uma economia?

GPT



Time: 1 seconds and 68 thousandths

DeepSeek



Time: 4 seconds and 45 thousandths

Time: GPT Better
Response: Equal



DeepSeek x GPT-4o. Who is Better?

We did this analysis in three steps. Easy questions, medium question and **hard questions**

Question 1: Avalie criticamente os desafios éticos do uso de LLMs em ambientes corporativos

GPT



Time: 2 seconds and 56 thousandths

DeepSeek

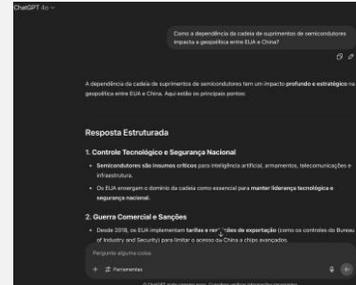


Time: 4 seconds and 99 thousandths

Time: GPT Better
Response: Equal

Question 2: Como a dependência da cadeia de suprimentos de semicondutores impacta a geopolítica entre EUA e China?

GPT



Time: 1 seconds and 39 thousandths

DeepSeek



Time: 5 seconds and 12 thousandths

Time: GPT Better
Response: GPT Better

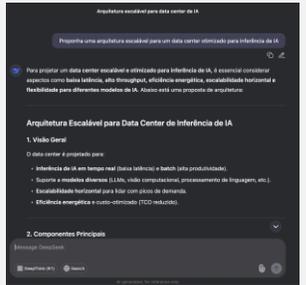
Question 3: Proponha uma arquitetura escalável para um data center otimizado para inferência de IA

GPT



Time: 1 seconds and 39 thousandths

DeepSeek



Time: 4 seconds and 85 thousandths

Time: GPT Better
Response: GPT Better

Question 4: Qual é a relação entre o fim da Lei de Moore e a ascensão do modelo de computação acelerada da NVIDIA?

GPT



Time: 1 seconds and 71 thousandths

DeepSeek



Time: 4 seconds and 57 thousandths

Time: GPT Better
Response: GPT Better

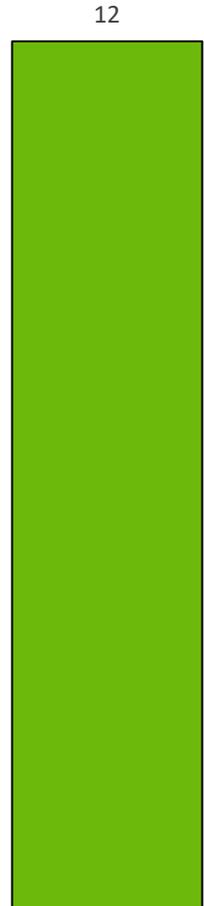


And in the end, who was better, DeepSeek or GPT-4o?

A Final Comparison Between DeepSeek and GPT-4o

GPT-4o proved to be dominant, delivering faster responses

Time



GPT - 4o

0

DeepSeek

GPT-4o and DeepSeek were similar in responses, but GPT-4o held the advantage

Response



GPT-4o

2

DeepSeek



Equal

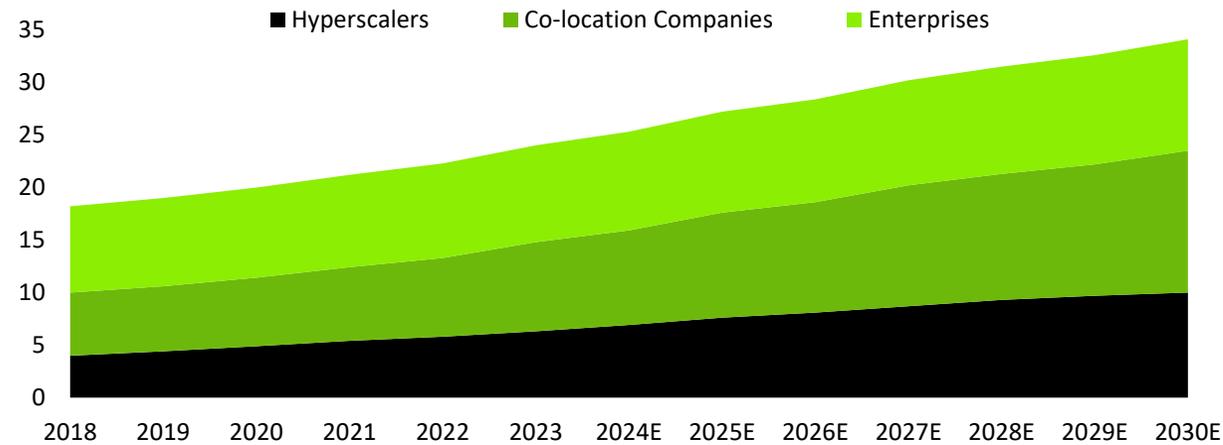


Can Energy be a Risk for AI?

Rising power demands may challenge scalability and sustainability of AI growth

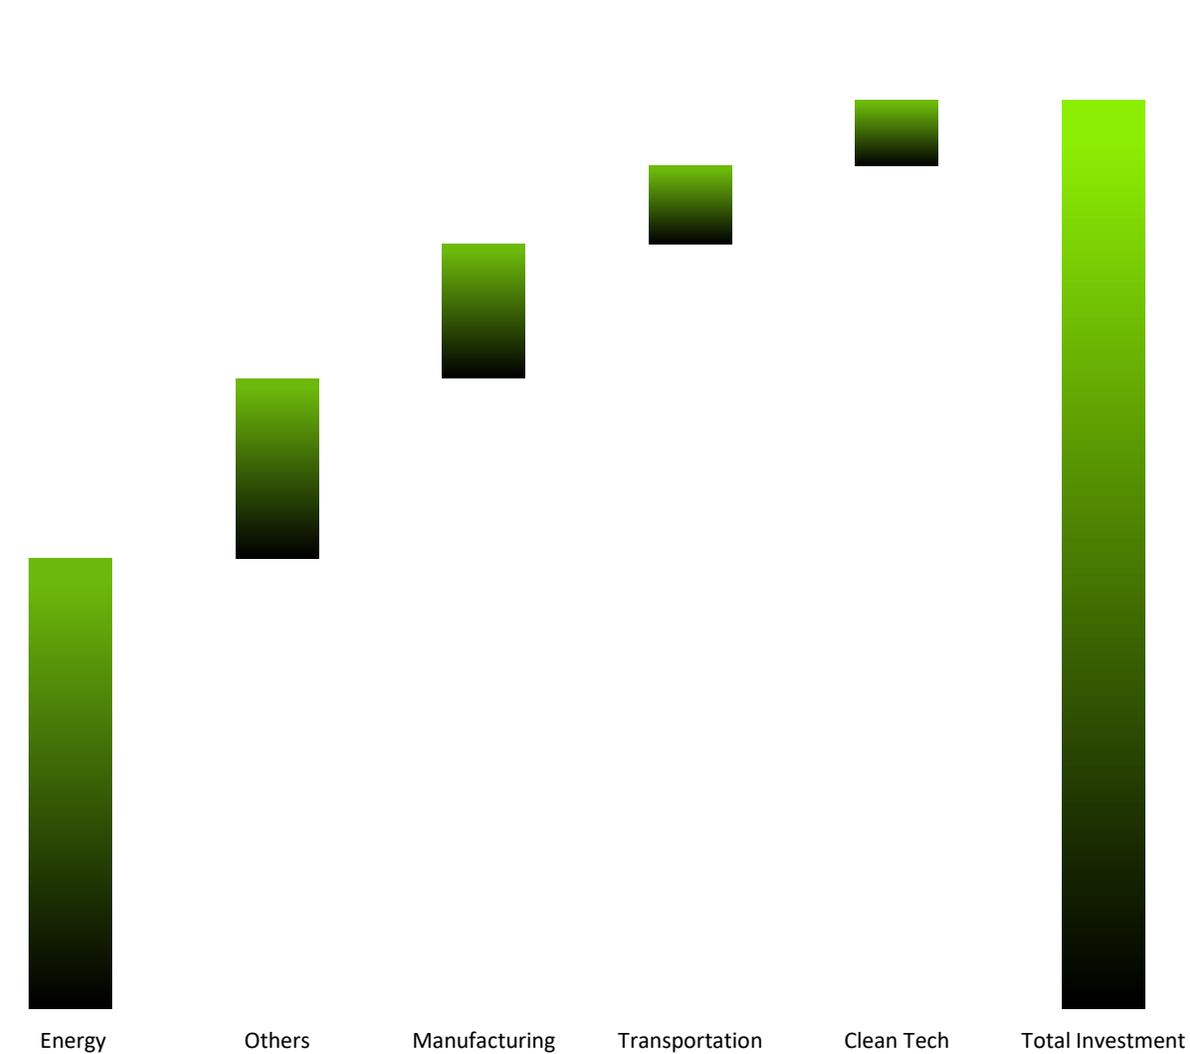
Data centers are booming, led by heavy U.S. tech investments

Data center power consumption (gigawatts), by providers, in USA



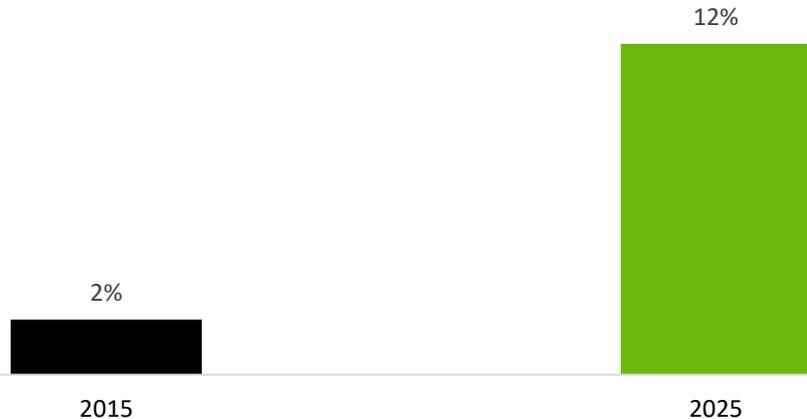
The government is investing to improve energy infrastructure

Total Government Investment in Power Grid (USD Bn)



U.S. data center energy use will grow 6x by 2025, driven by AI and cloud demand

Data Center as a % of USA energy

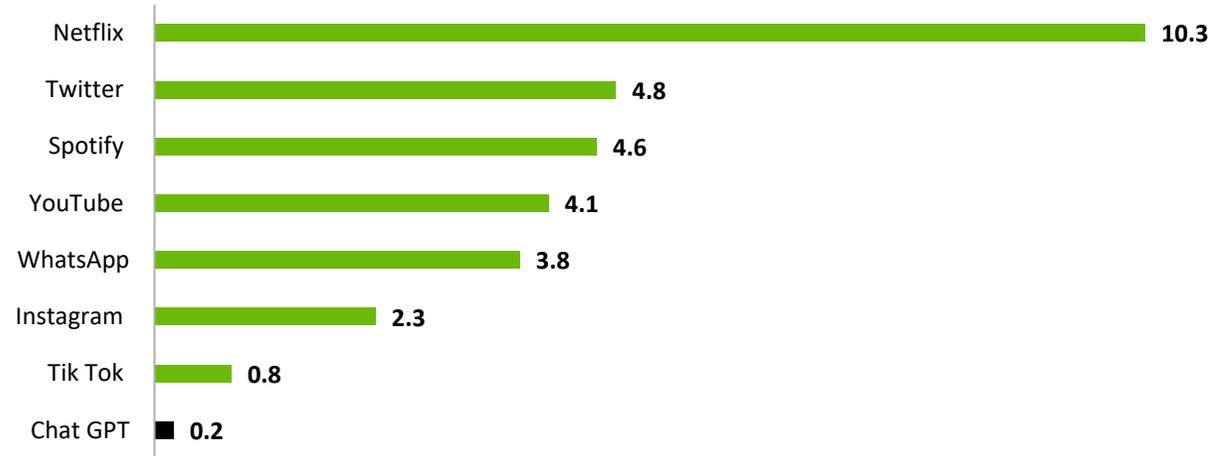


I. Risk: Slowdown in AI Investments

Inference Is the New Bottleneck, Reinforcing the Need for Continued AI Investment

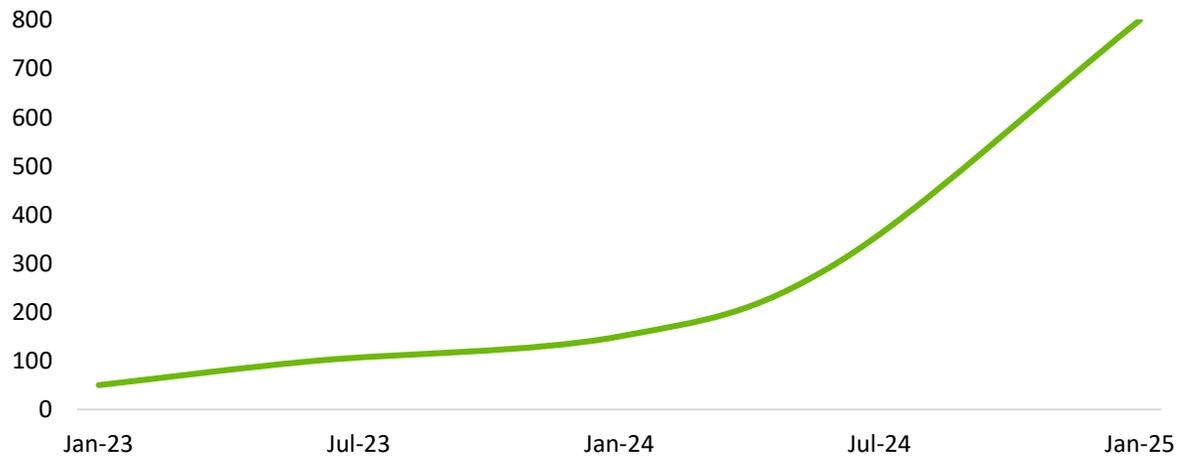
I. Generative AI Adoption Is Scaling Faster Than Any Prior Technology

Years to Reach 100 MM Users



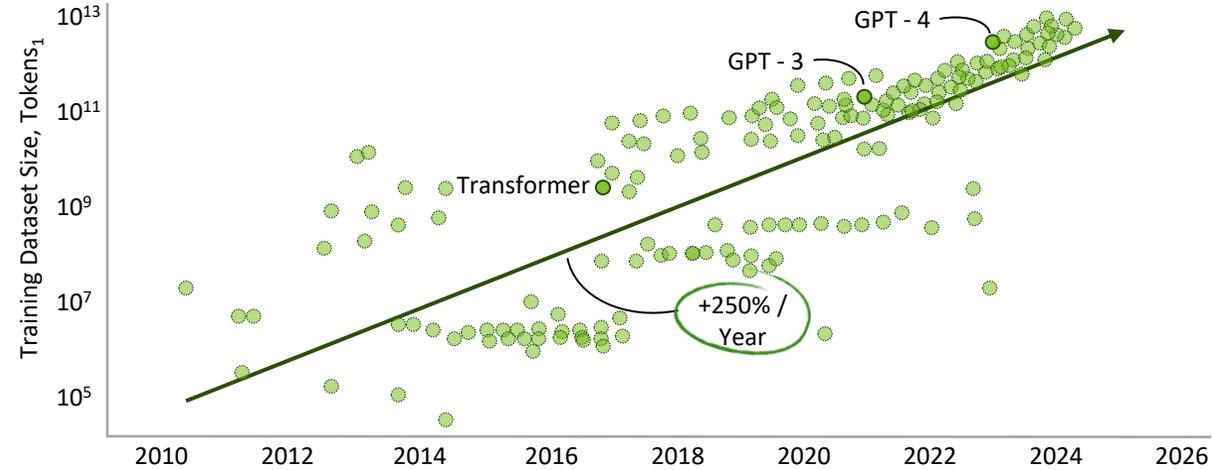
III. User Growth is Exploding — Fueling Persistent Inference Demand

Number of weekly users in Chat GPT (millions)



II. No Signs of Slowdown: AI Models Keep Getting Larger and Heavier

AI Model training Dataset Size by Model Release Year



IV. Why Inference Will Keep Growing — With or Without New Training Breakthroughs

Drivers Behind the Exponential Growth of AI Inference Demand



Jensen Hung, Nvidia's Founder & CEO

"With the rise of inference-based AI, compute requirements have already increased by 100x — and as reasoning models advance, inference will become one of the main drivers of AI infrastructure demand"



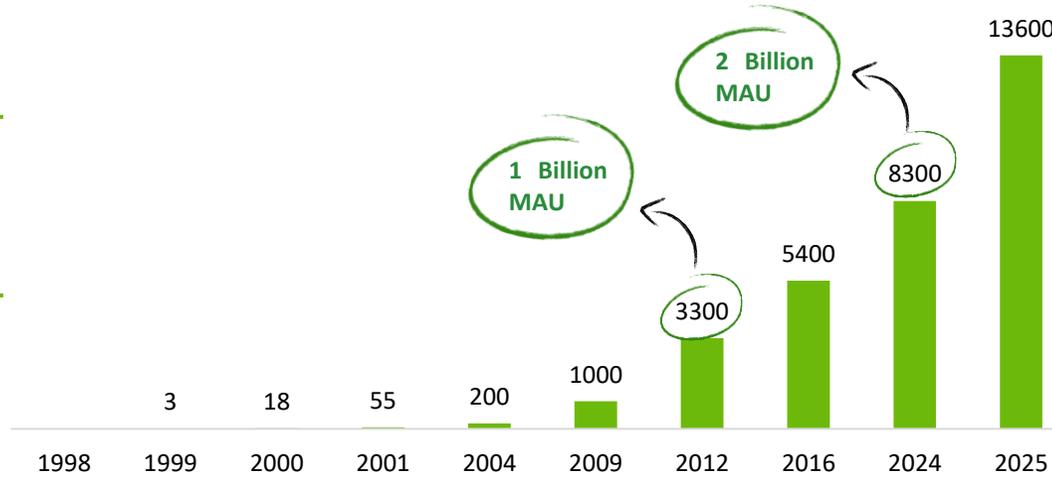
II. Risk: Slowdown in AI Investments

Adoption curves from past tech giants show ChatGPT's growth is still in its early stages

I. Google's Growth Trajectory Shows That ChatGPT's Expansion Is Likely Just Beginning

Million Users	2 Months	3 Year	13 Year
	1	10	1000
	100	800	

Google Daily Searches (million)

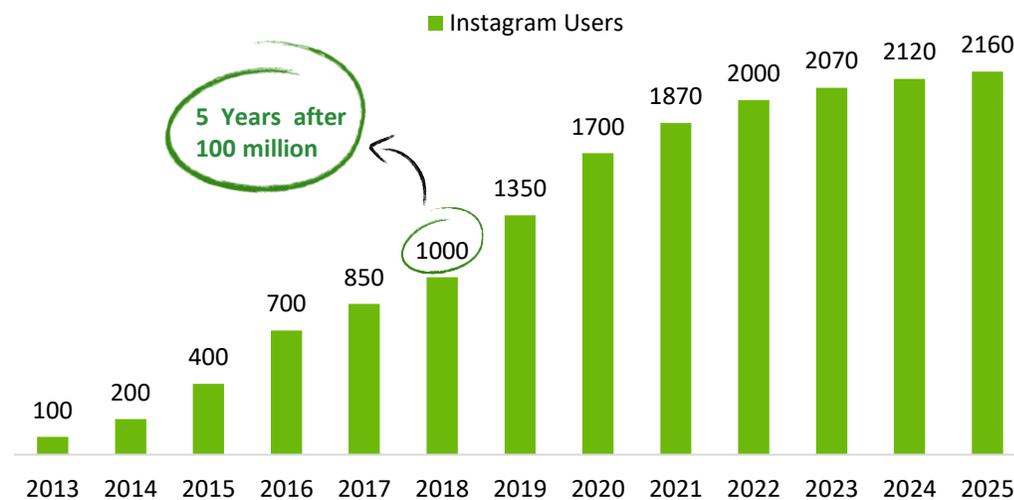


Google took over a decade to reach 1 billion users — but its real scale came afterward: daily searches grew from 3.3 billion in 2012 to 13.6 billion in 2025, more than 4x growth even with a mature user base. ChatGPT, by contrast, reached 100 million users in just 2 months and is projected to hit 800 million monthly active users in under three years. Despite its record-breaking adoption, we believe — just like with Google — the biggest growth is still ahead, both in user base and in daily engagement per user.

II. What took Instagram a decade, ChatGPT may achieve in just a few years

Million Users	2 Months	1 Year	3 Year
	1	10	100
	100	180	800

Instagram Monthly Active Users (million)



While most platforms take years to gain traction, ChatGPT reached 100 million users in just 2 months — a milestone that took Instagram years to approach. Three years in, ChatGPT has already hit 800 million users, signaling not just fast adoption, but a potential structural shift in how quickly transformative technologies scale. Looking at Instagram's trajectory, it took 5 years after reaching 100 million to hit 1 billion users, and over a decade to surpass 2.16 billion. If ChatGPT maintains even part of this trajectory, the implications for future scale are massive.



Value Generation with AI

Real-World Impact and What's Next: How Leading Companies Are Embedding AI at Scale

I. AI Is Already Delivering Tangible Business Value Across Industries

Real Value Generation seen on Companies

AI Use Case	Results
 AI-powered recommendation engine	+35% of revenue from personalized suggestions
 Generative AI for internal code generation	40% faster delivery of internal tools
 AI for demand forecasting and inventory management	+10% inventory efficiency, \$1B in savings
 AI-based predictive maintenance in equipment	30% reduction in unplanned downtime

II. Using AI Is No Longer Optional, It's a Standard in High-Performance Cultures

Case Study: Duolingo

Shopify Sent Internal Memo Making AI Usage Mandatory for All Employees

"Using AI effectively is now a fundamental expectation of everyone at Shopify." "Before asking for more headcount, teams must demonstrate why they cannot get things done using AI."

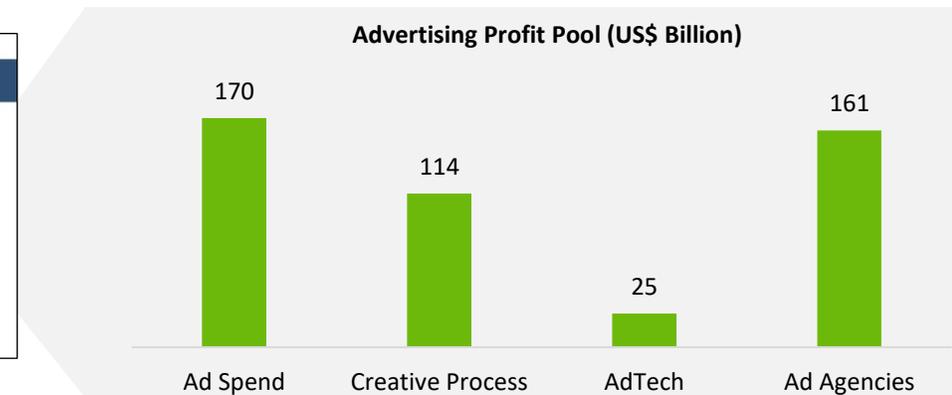
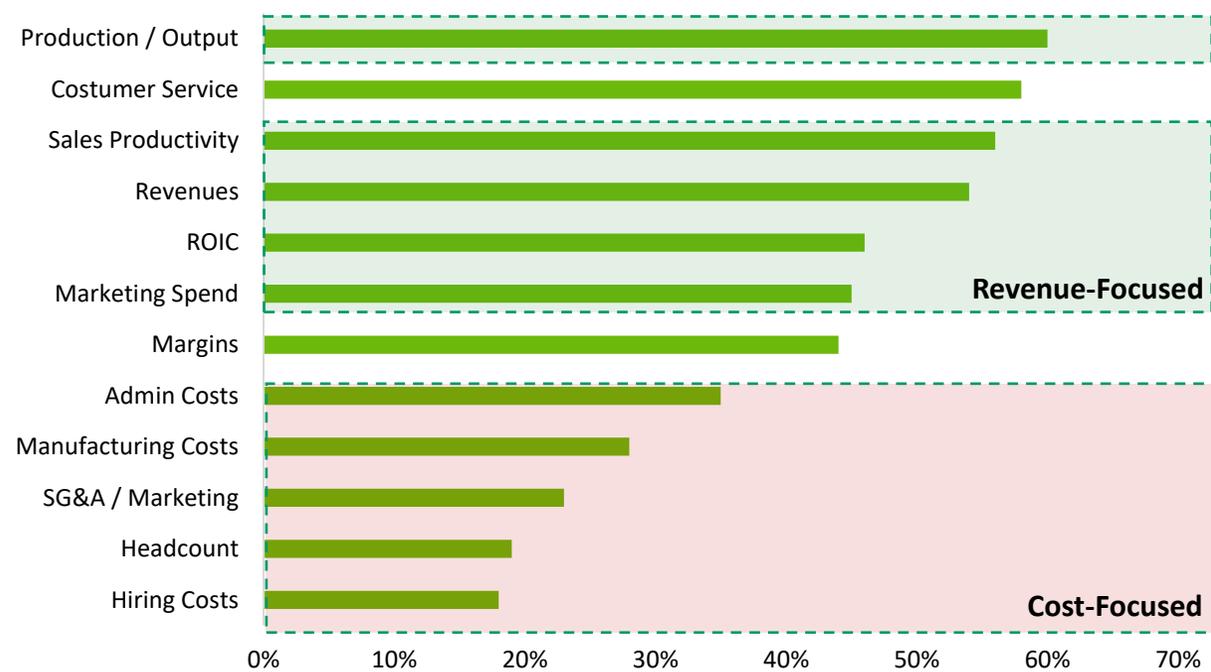
Case Study: Shopify

Duolingo's CEO Luis von Ahn issued an internal memo declaring that Duolingo is becoming an AI company

"Duolingo is going to be AI-first." "AI use will be part of what we evaluate in performance reviews." "We'd rather move with urgency and take small hits on quality than move slowly and miss the moment."

III. Global Enterprises Are Prioritizing GenAI to Transform Revenue and Cost Structures

GenAI Improvements Targeted for Global Enterprises over the next 2 years (2024)



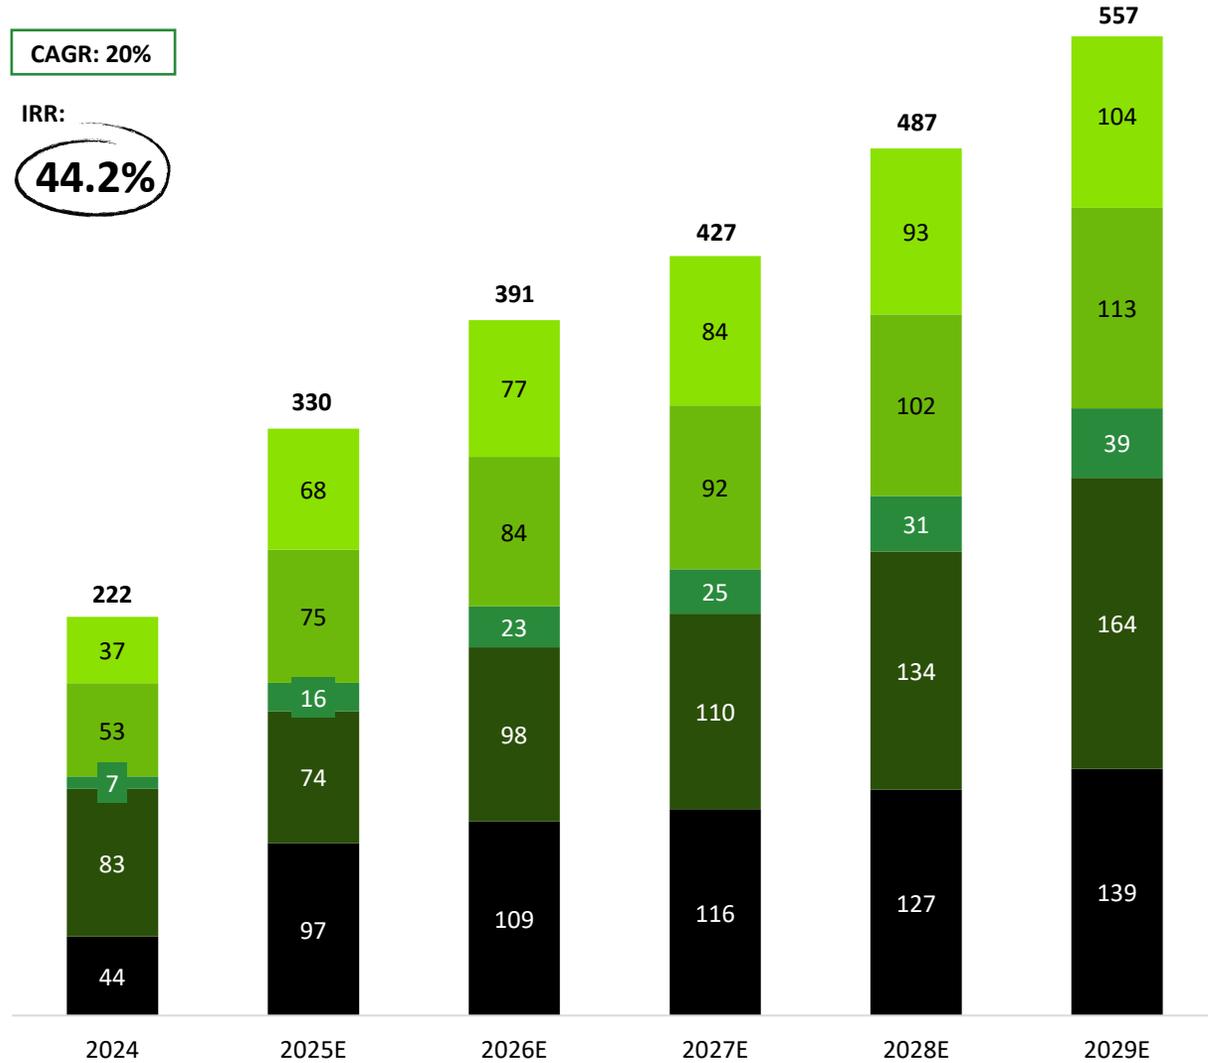
Build-Up Sanity Checks

Goldman Sachs

I. Goldman Sachs forecasts Hyperscalers Capex

Goldman Sachs Hyperscalers Data Center Capex (US\$ Billion)

■ Microsoft ■ Amazon ■ Oracle ■ Google ■ Meta



CAGR: 20%

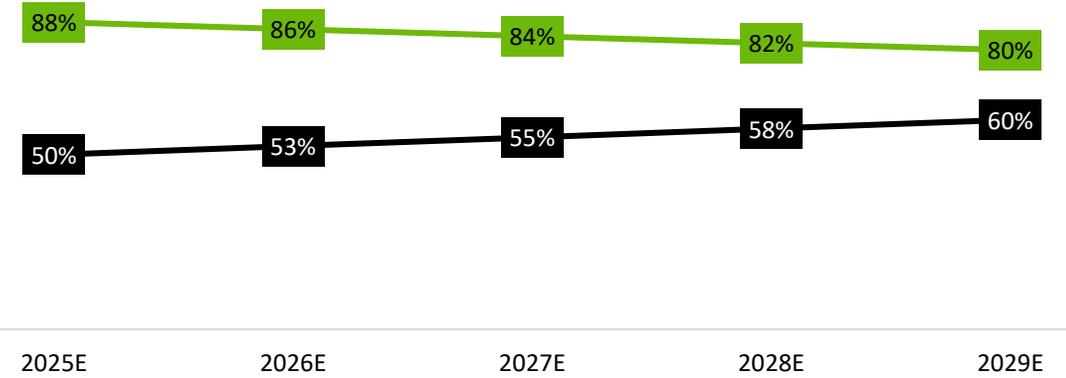
IRR:

44.2%

II. We project an expected NVIDIA Market Share to go from 88% in 2025 to 80% in 2029

Data Center Capex for GPUs and NVIDIA's Market Share (%)

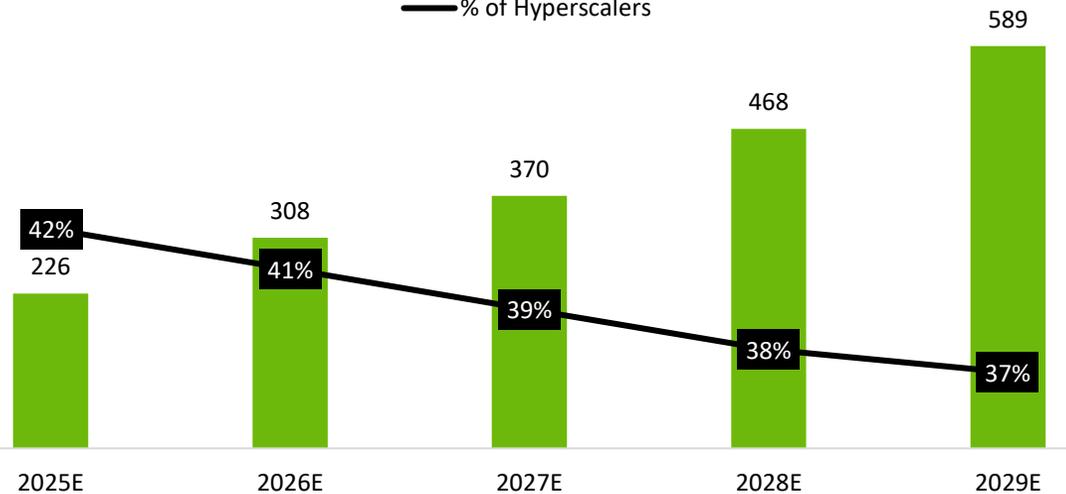
— % of Data Center Capex for GPUs — NVIDIA Market Share



III. NVIDIA Data Center Revenue in 2029 is 589 Billion

Nvidia's Data Center Revenue (US\$ Billion)

— % of Hyperscalers



Build-Up Sanity Checks

Dell'oro

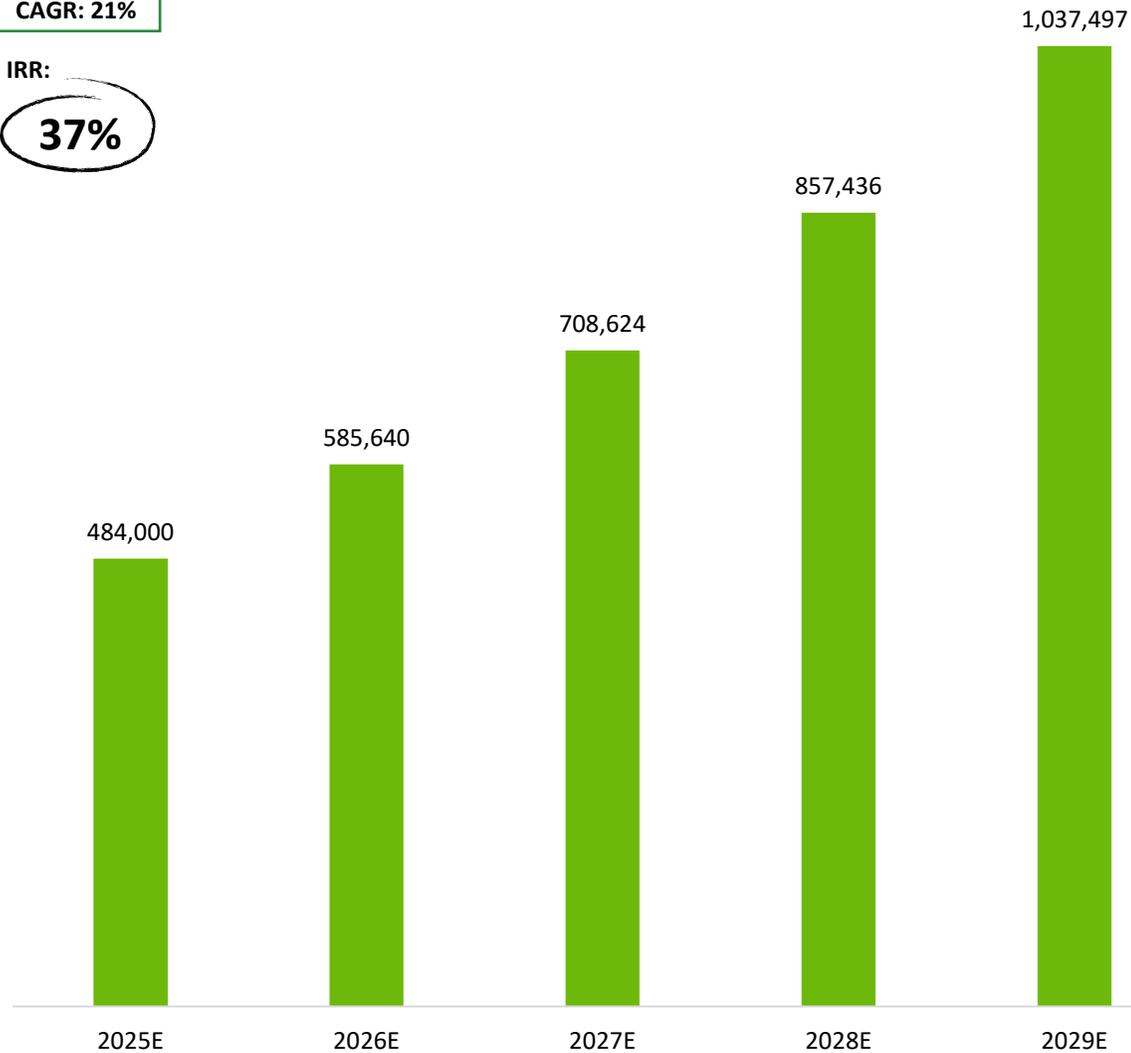
I. The Number Jensen Refers To: Global Data Center Capex

Dell'oro (Jensen) Global Data Center Capex

CAGR: 21%

IRR:

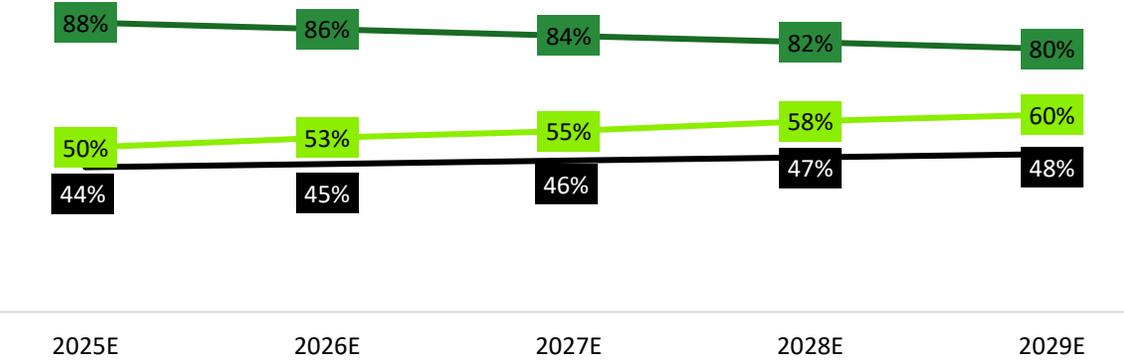
37%



II. Key Assumptions Behind Nvidia's Share of Capex

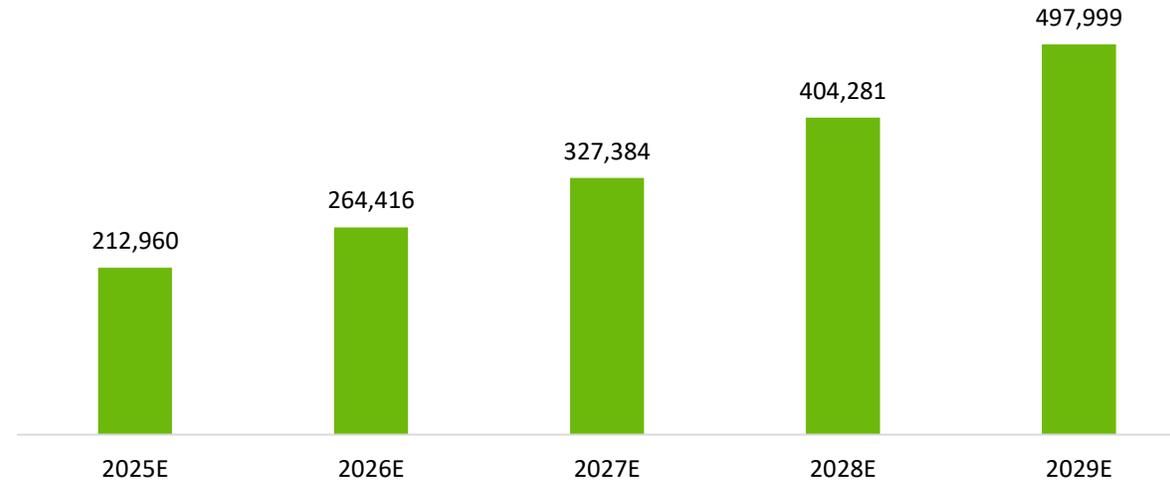
Assumptions

— % of Global Capex — % of Data Center Capex for GPUs — % of NVIDIA GPU Market Share



III. Projected Nvidia Data Center Revenue

Nvidia's Data Center Revenue

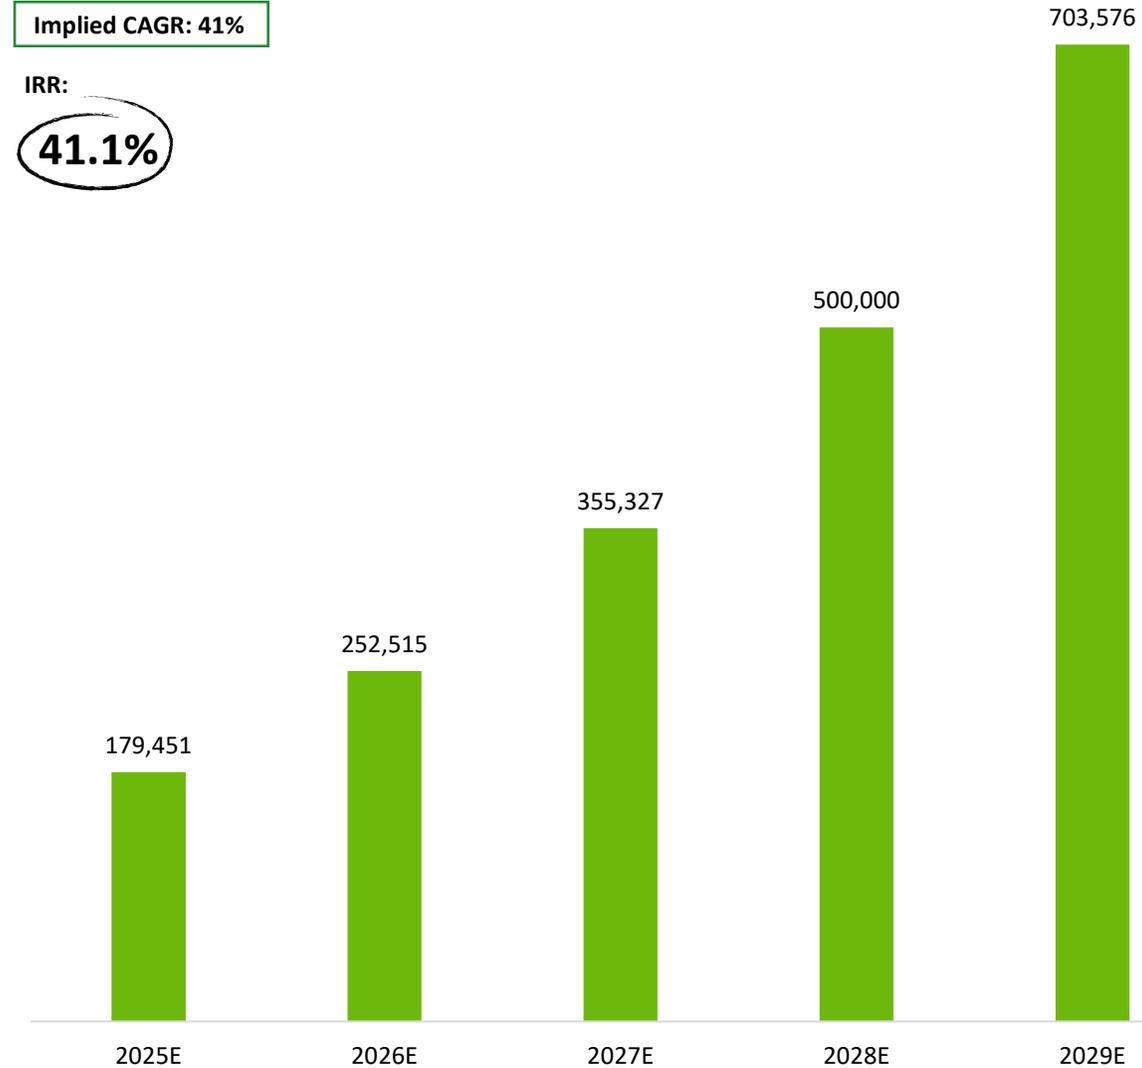


Build-Up Sanity Checks

AMD

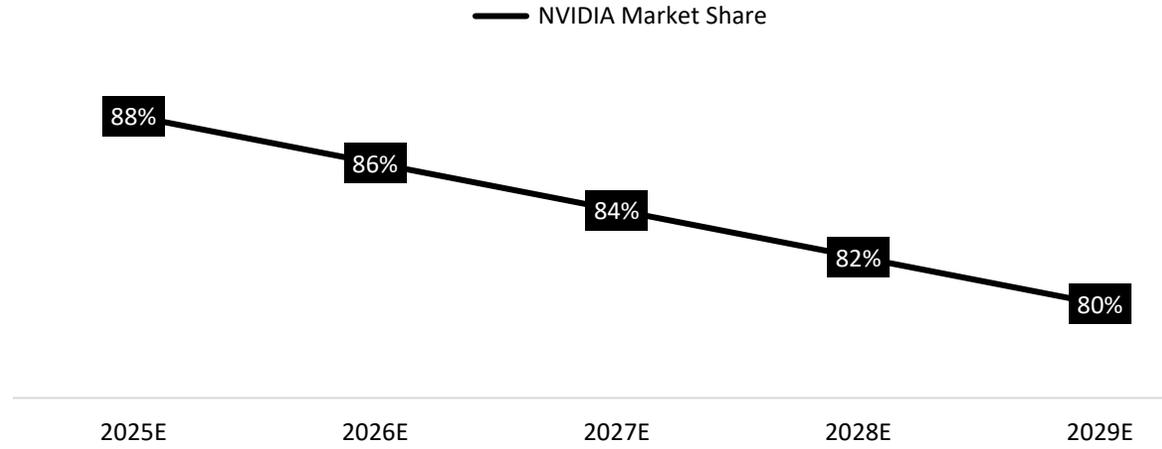
I. The Number Lisa Su Refers To: Total Accelerators Market Revenue

AMD (Lisa Su) Total Accelerators Market Revenue



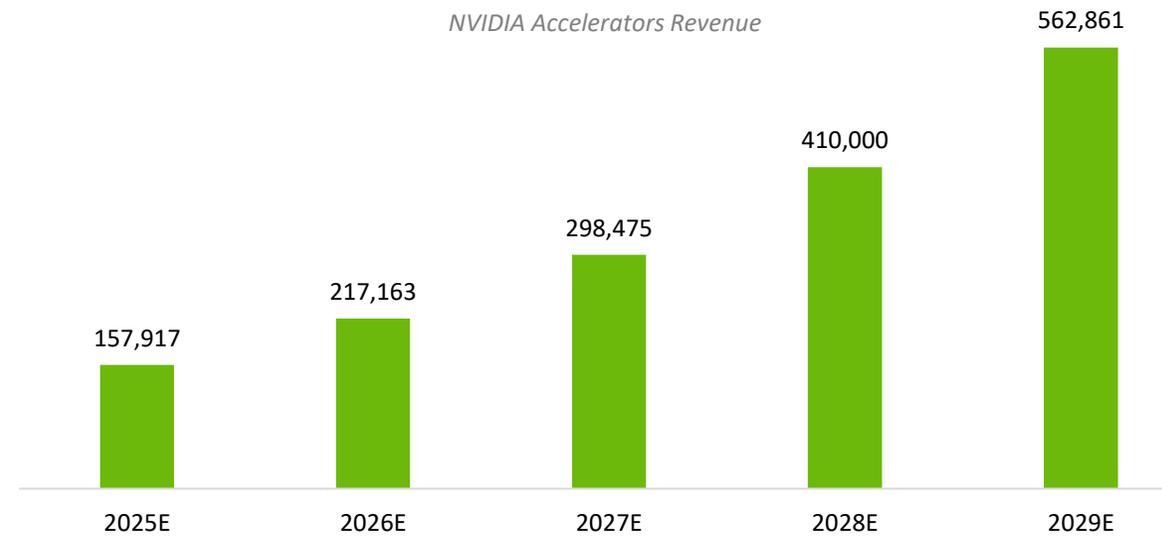
II. Key Assumptions Behind Nvidia's Share of Capex

Assumptions



III. Projected Nvidia Accelerators Revenue

NVIDIA Accelerators Revenue



Build-Up Sanity Checks

McKinsey

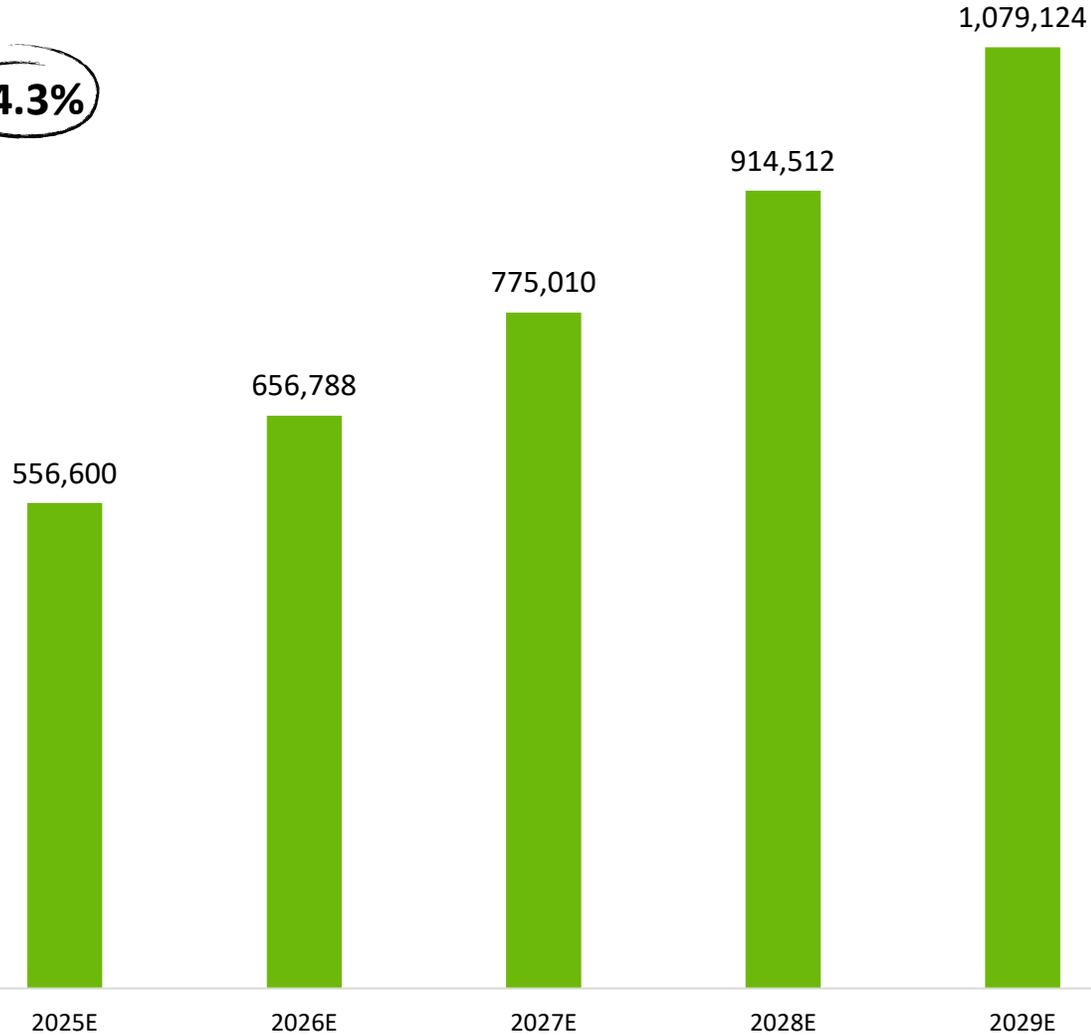
I. Data Center CapEx Driven by AI

McKinsey Global Data Center Capex Driven by AI (US\$ million)

CAGR: 18%

IRR:

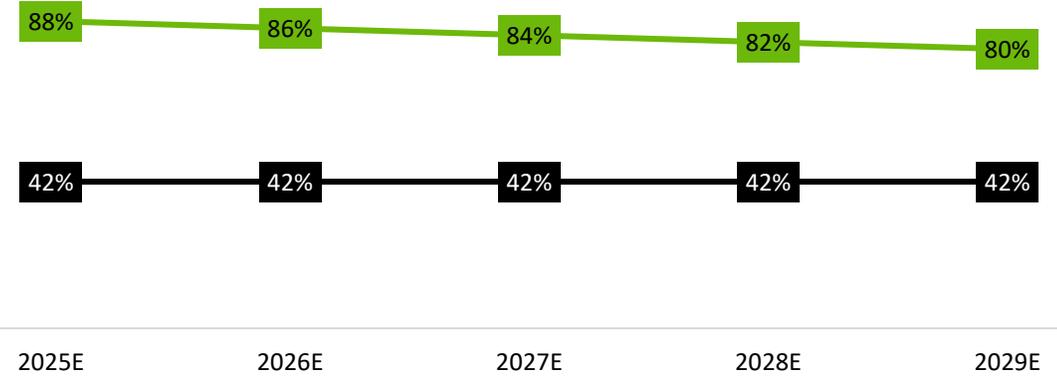
24.3%



II. We project a expected NVIDIA Market Share to go from 88% in 2025 to 80% in 2029

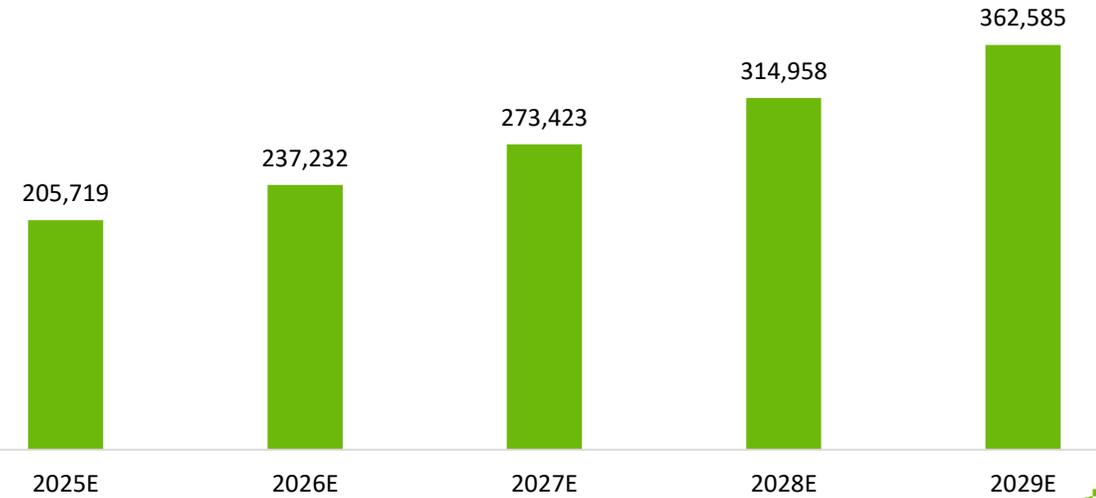
Data Center Capex for GPUs and NVIDIA's Market Share (%)

— % of Data Center Capex for GPUs — NVIDIA Market Share



III. NVIDIA Data Center Revenue in 2029 is 362.585

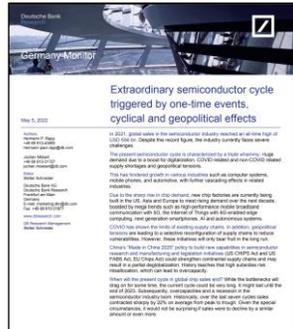
Nvidia's Data Center Revenue



I. AI Bubble

Using historical and economic factors, we conclude that NVIDIA is not in a bubble

I. Although it varies greatly, the indicated semiconductor cycle length is 27 months

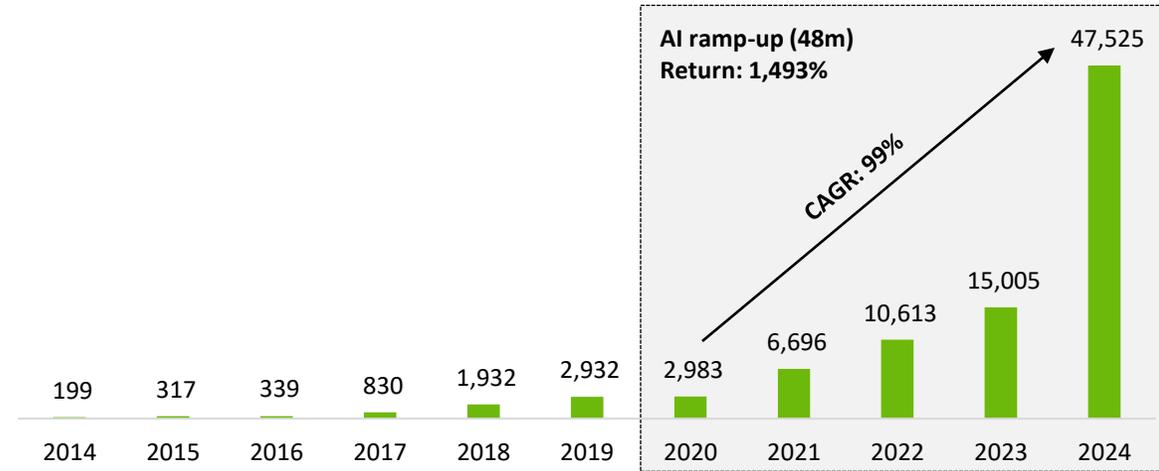


Semiconductor business cycles

Trough	Peak	Length in months	Sales in %	
			Trough-to-peak	Peak-to-trough
Oct 1998	Dec 2000	26.0	92.7	-46.3
Apr 2002	Mar 2005	35.0	90.0	-5.3
Jul 2005	Oct 2008	38.0	28.3	-39.0
Aug 2009	May 2011	21.0	85.7	-9.2
May 2013	Jan 2015	20.0	29.9	-9.3
Aug 2016	Nov 2018	27.0	62.7	-23.0
	Average	27.7	64.9	-22.0

II. With AI, the cycle has lasted 48 months with no signs of stopping

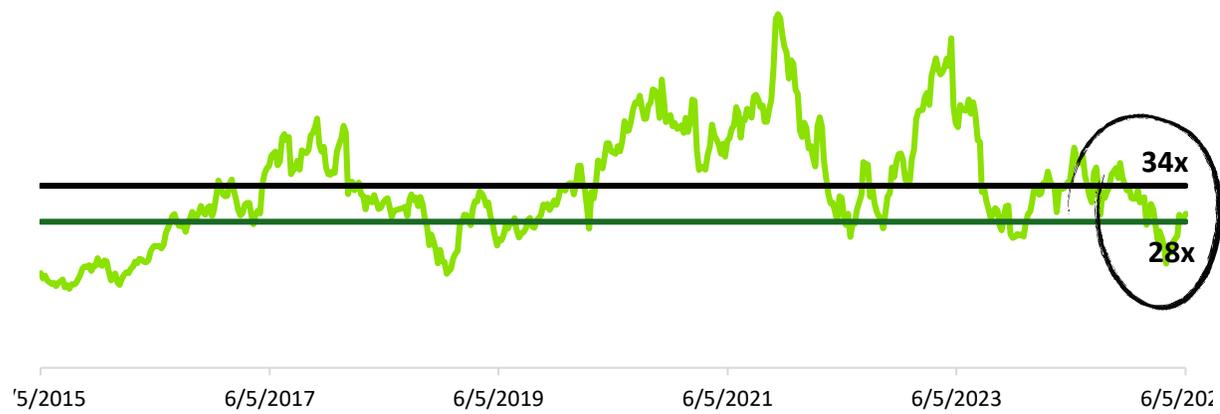
NVIDIA Data Center Revenue (US\$ Million)



III. Moreover, the company's multiples are below the historical average

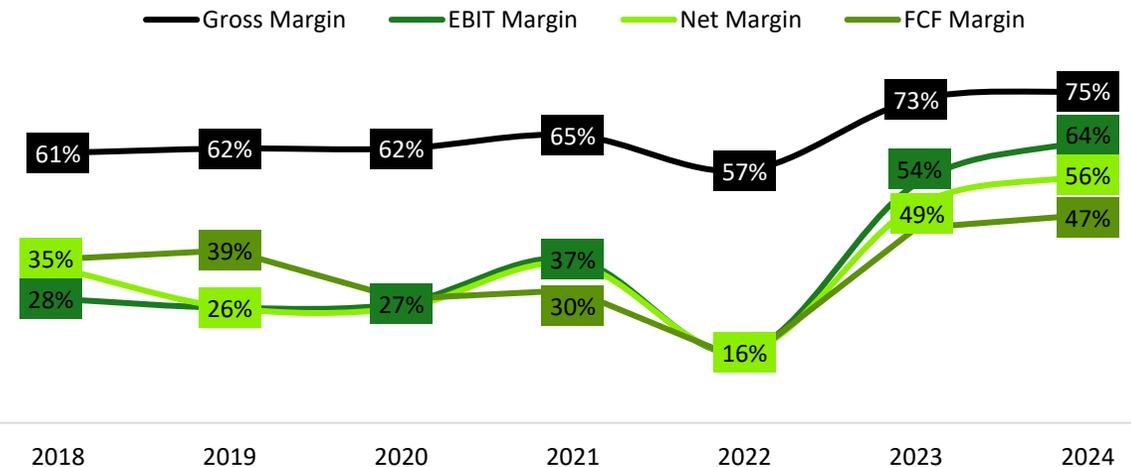
NVIDIA P/E 1yFw

NVDA



IV. Not only that, but it has also been generating cash

NVIDIA's margins (%)



II. AI Bubble

Using historical and economic factors, we conclude that NVIDIA is not in a bubble

I. We had the opportunity to attend Damodaran's lecture on his short position in AI

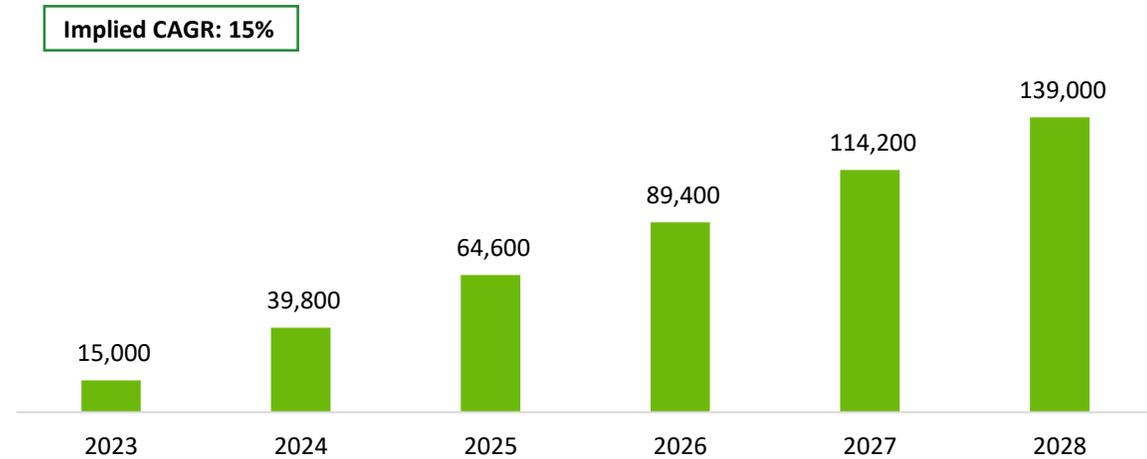


AI'S WINNERS, LOSERS AND WANNABES: BEYOND BUZZ WORDS!, April 2024

We had the opportunity to attend Aswath Damodaran's lecture on the AI bubble and his expectations for NVIDIA in 2024 at Insper

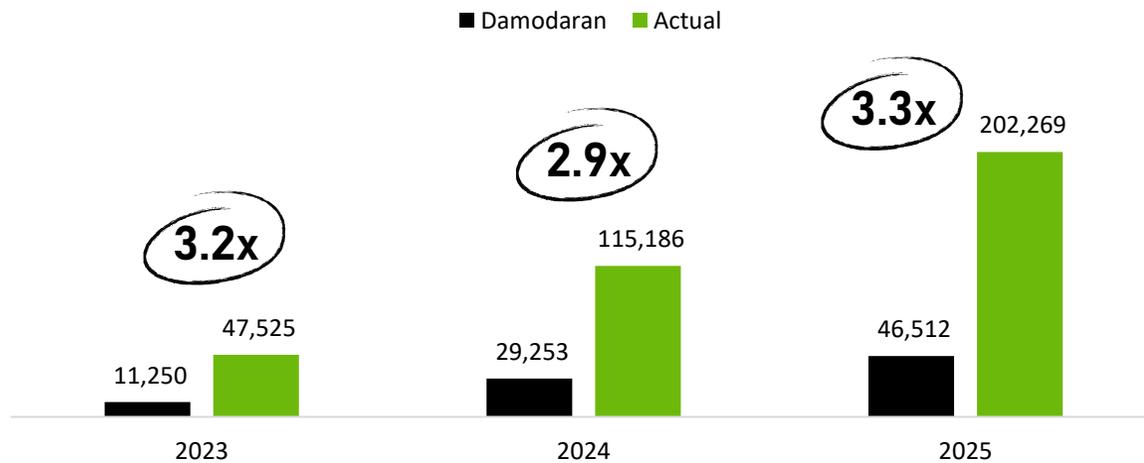
II. Damodaran's projections used a 15% CAGR for AI

NVIDIA Data Center Revenue (US\$ Million)



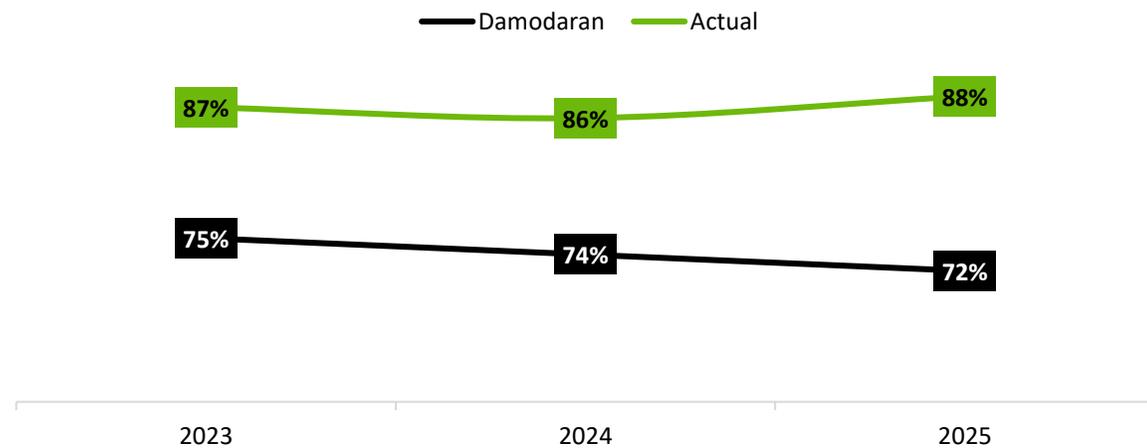
III. Technological disruption creates new growth paradigms, leading to massive errors

NVIDIA Data Center Revenue (US\$ Million)



IV. Moreover, NVIDIA's market dominance was underestimated

NVIDIA's Market Share (%)



Key-man Risk

Although led by a single man, we do not see this as a risk due to his continued and extreme willingness to remain with the company

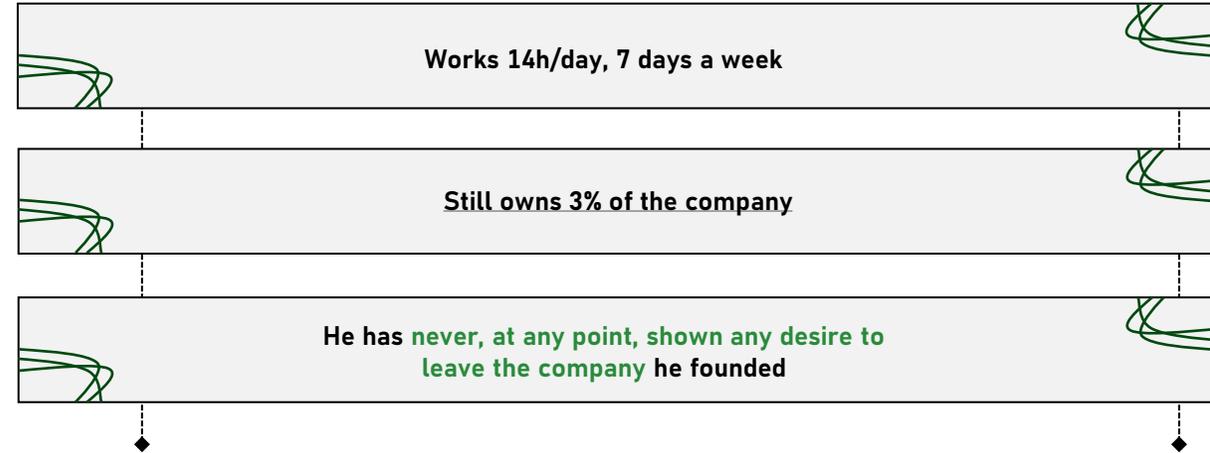
I. Despite being 62 years old, Jensen still upholds his culture of excellence and has never shown any sign of leaving the company



Jensen Huang, 62 years, NVIDIA CEO & Founder

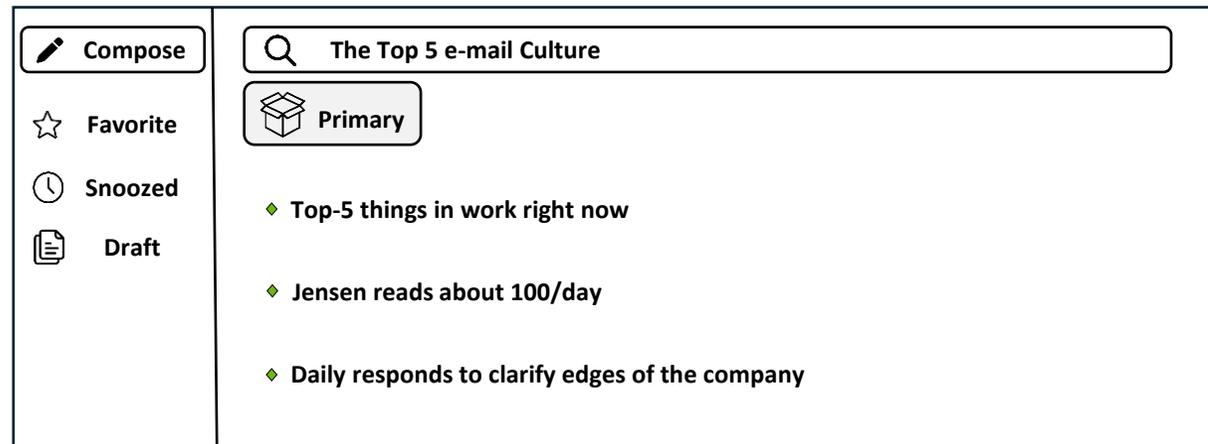
"I work from the moment I wake up to the moment I go to bed, and I work seven days a week. When I'm not working, I'm thinking about working"

"When I'm not working, I'm thinking about working, and when I'm working, I'm working. I sit through movies, but I don't remember them because I'm thinking about work"

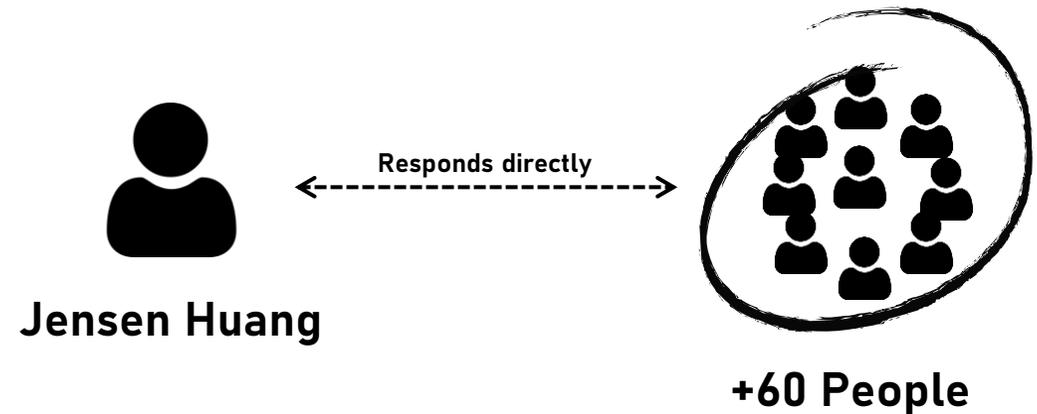


II. Jensen remains deeply involved through his "Top 5 email" culture and direct oversight of 60+ executives, showing his continued intensity at NVIDIA

Top-5 things culture



Organizational Structure



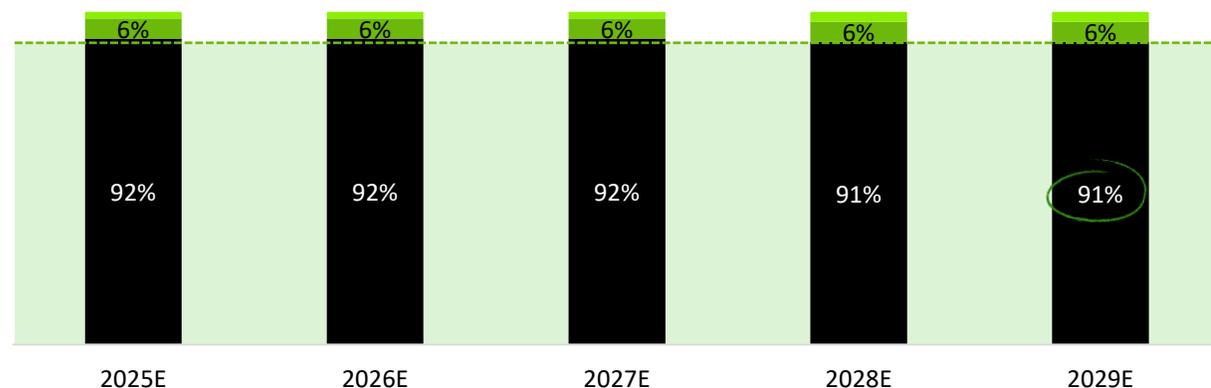
How Could This Investment Lose Money?

In a scenario where AI products are not monetized, we see a possibility of losing money

I. We are highly exposed to data center revenue, which is directly tied to AI development

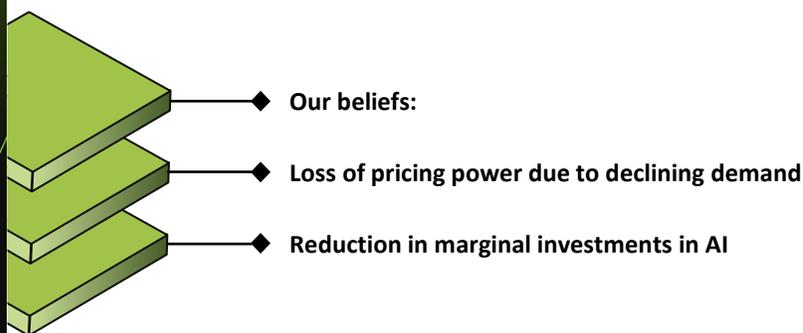
Projected revenue per segment

■ Data Center ■ Gaming ■ Others



III. If AI fails to create value, investments would drop and pricing power would decline

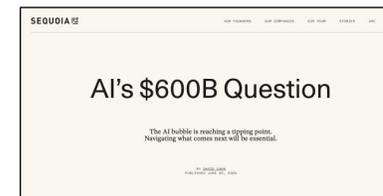
Consequences of not creating value



Therefore, we arrive at bear case assumptions by **zeroing out marginal investment in AI and reducing margins to pre-AI levels at 65%**

II. Applications may face a supply that exceeds demand

Sequoia's AI thesis



NVIDIA Data Center Run-Rate Revenue	179,308
Data Center Facility Build and Cost to Operate	50%
Implied Data Center AI Spend	358,616
Software Margin	50%
AI Revenue Required for Payback	717,232

Risk of AI Platform Saturation and Commoditization

If many companies offer similar AI platforms, it could trigger price wars, margin compression, and commoditization, making it harder to recover the heavy investments in GPUs and data centers

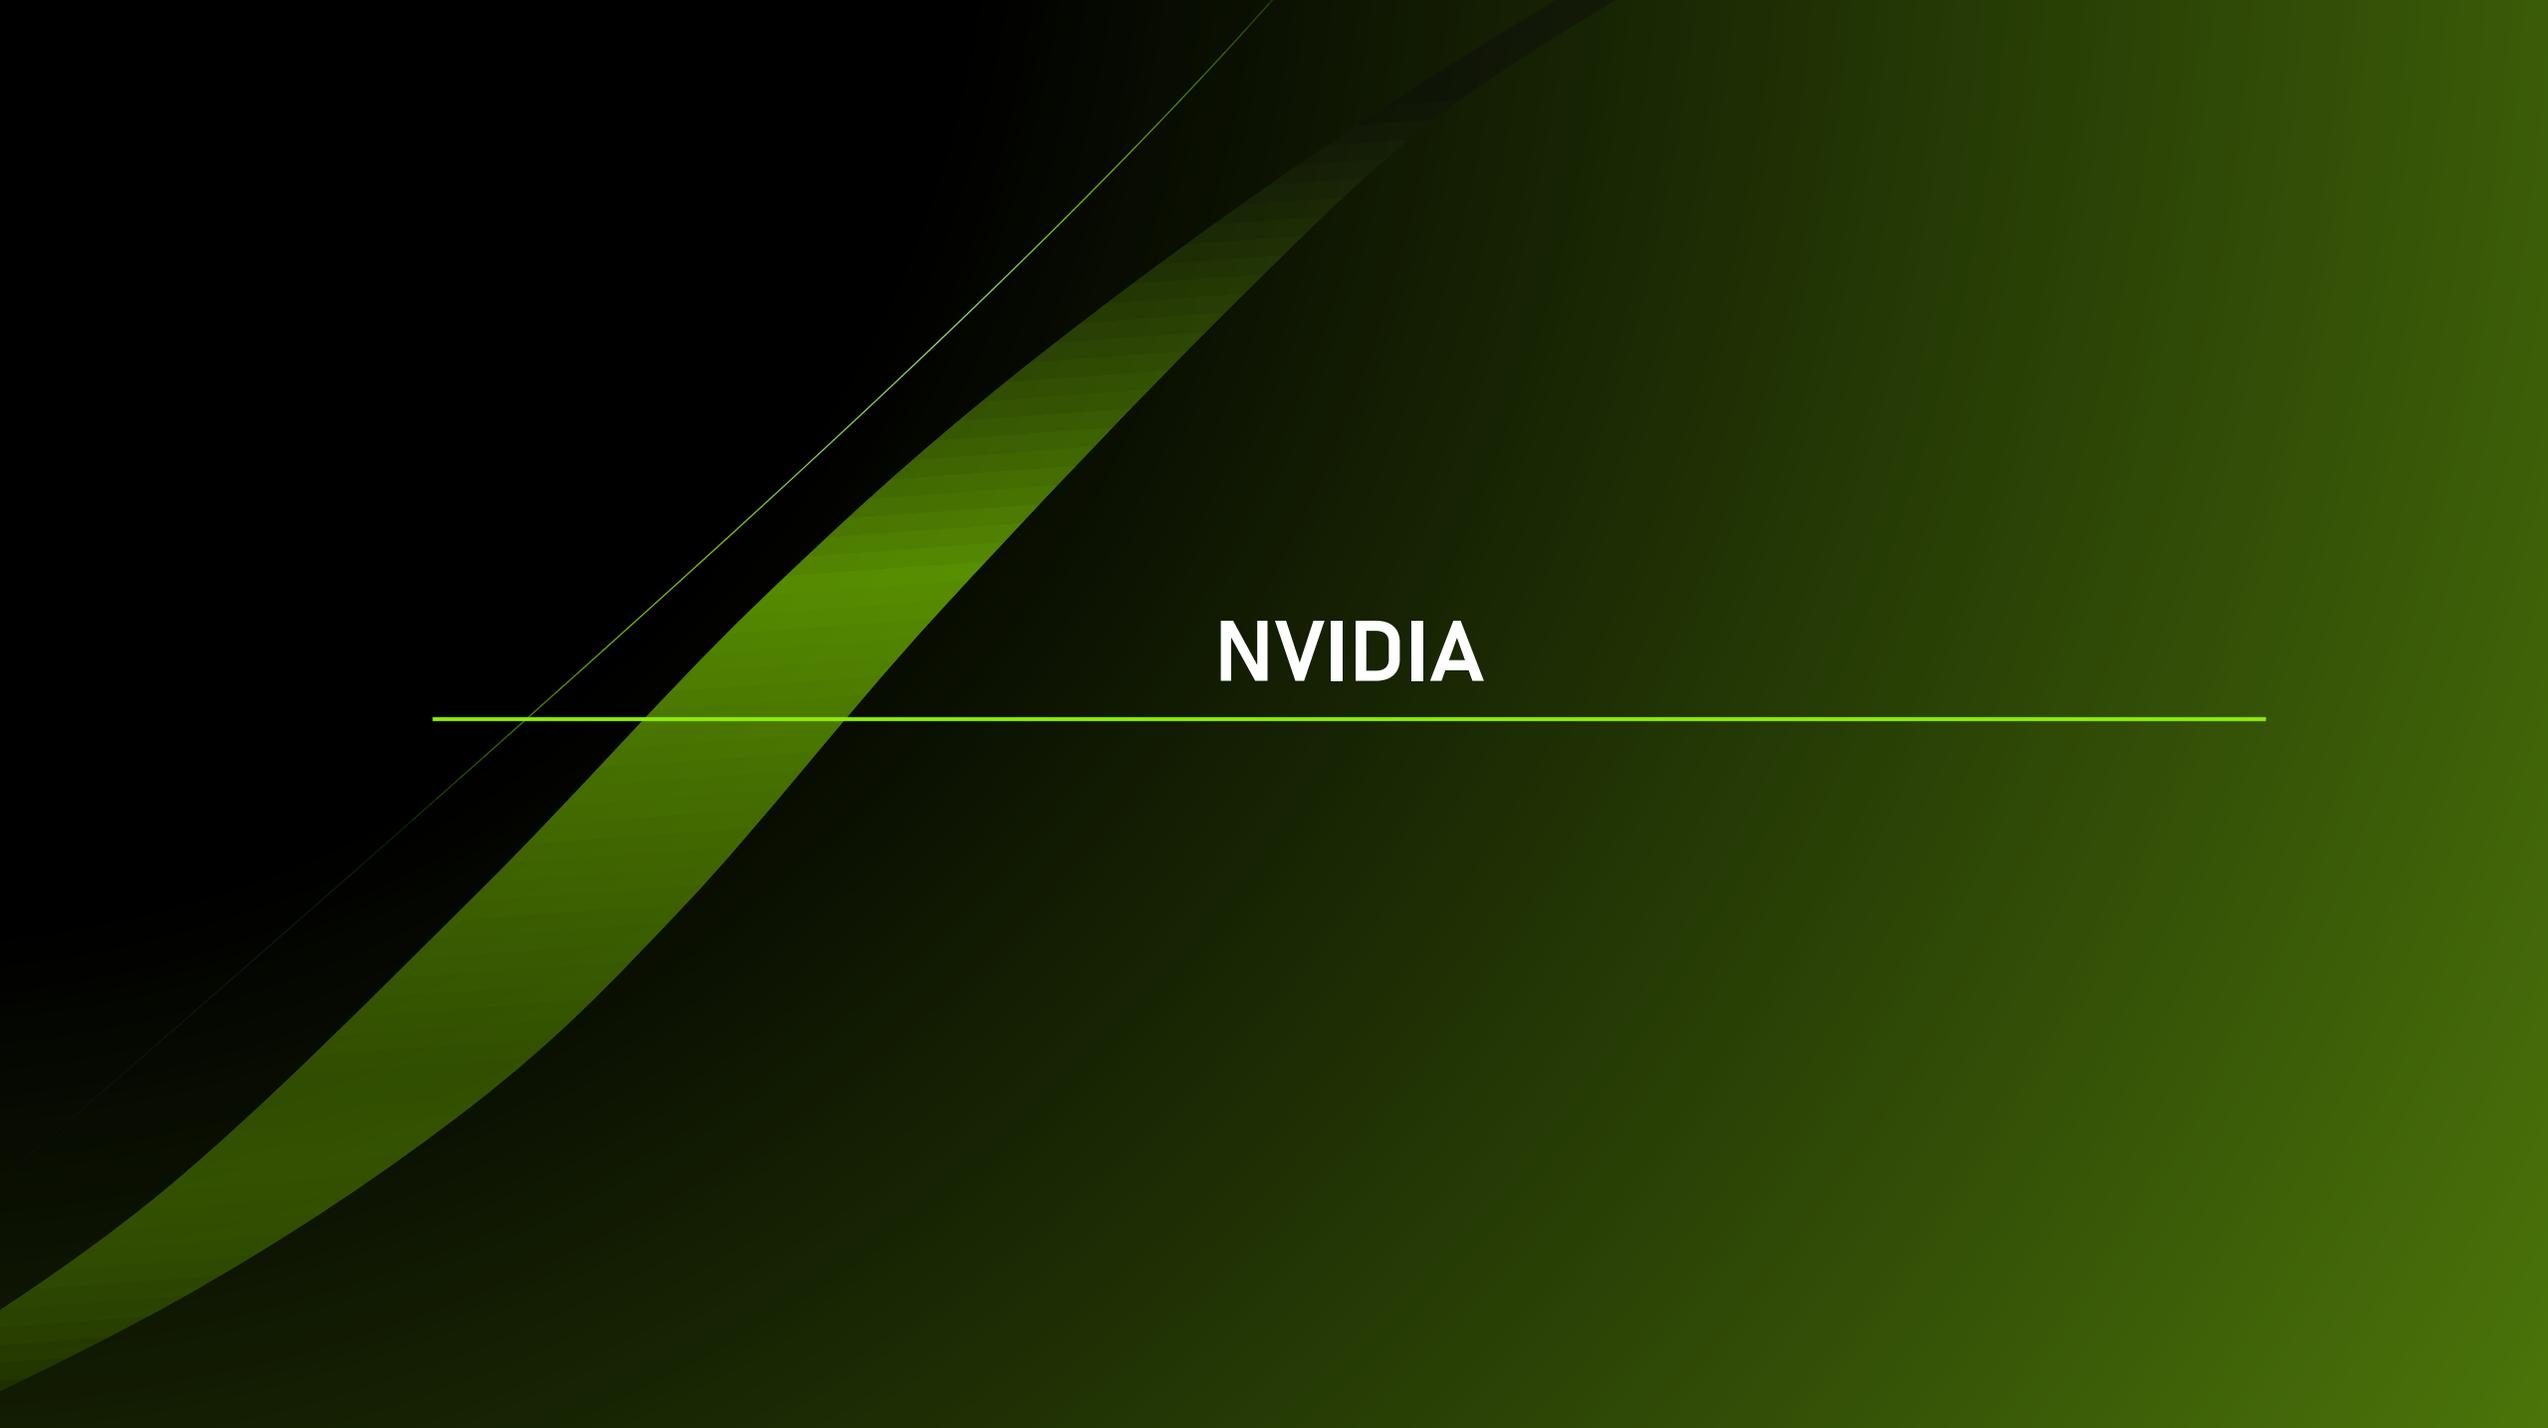
IV. We assume steady 2024 capex and lower gross margin going forward

Sensitivity Analysis

Gross Margin (%)	Marginal Capex (%)						
	0.0%	10.0%	20.0%	30.0%	40.0%	50.0%	60.0%
60.0%	-12.9%	-9.0%	-5.5%	-2.1%	1.0%	3.9%	6.7%
62.5%	-11.3%	-7.4%	-3.8%	-0.4%	2.8%	5.8%	8.6%
65.0%	-9.8%	-5.9%	-2.2%	1.3%	4.5%	7.6%	10.5%
67.5%	-8.4%	-4.4%	-0.6%	2.9%	6.2%	9.3%	12.3%
70.0%	-7.0%	-2.9%	0.9%	4.5%	7.8%	11.0%	14.0%
72.5%	-5.7%	-1.5%	2.4%	6.0%	9.4%	12.7%	15.7%
75.0%	-4.3%	-0.1%	3.9%	7.5%	11.0%	14.3%	17.4%

In a scenario where AI fails to generate enough value to drive marginal investments, we would see an **IRR of -10%**



The image features a dark green background with a large, curved, lighter green shape on the left side that resembles a stylized 'N'. A thin, light green horizontal line spans across the middle of the image, passing through the word 'NVIDIA'.

NVIDIA

Robotics I

A New Revenue Stream Emerging Beyond GPUs, Robotics Could Be NVIDIA's Next Growth Engine

I. Understanding Robotics — Applications Across Industries and the Measurable Economic Impact Already Being Delivered

What is robotics and What is used for?

Robotics enables the automation of physical tasks across industries — from manufacturing and logistics to healthcare and services — by combining AI with mechanical systems to perform repetitive or complex activities

- ➔ **Industrial automation:** assembly lines, logistics, semiconductor fabs
- ➔ **Warehouse & retail:** picking, sorting, restocking
- ➔ **Healthcare:** patient support, hospital delivery
- ➔ **Consumer services:** cleaning, elderly care, personal assistance
- ➔ **General-purpose humanoids:** designed to eventually perform everyday human tasks

Economic Impact

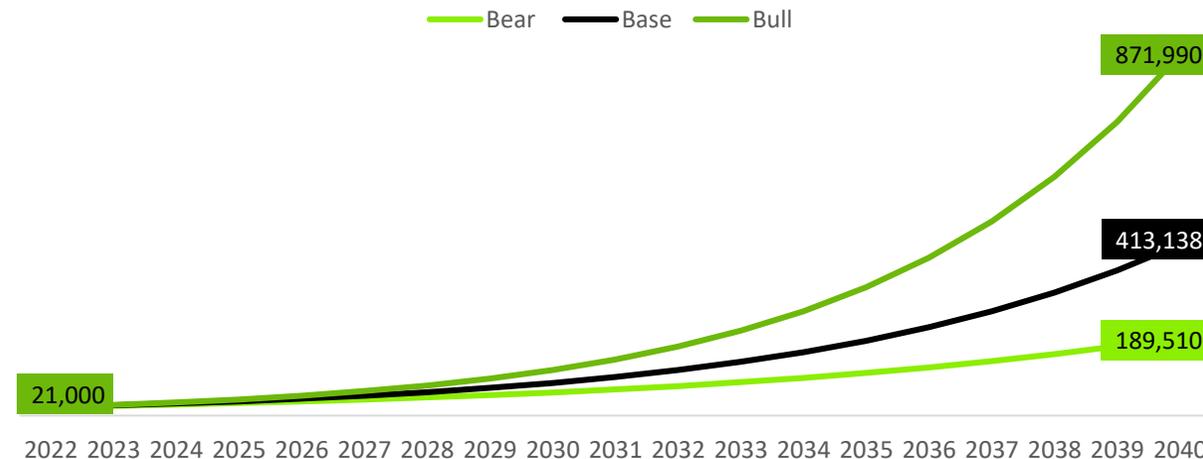
	-25% in operating costs	Up to US\$ 10B/year projected savings
	150→600 items/hour (×3.5)	30–40% reduction in labor costs
	+15% efficiency; -40% defects	+US\$ 10M/year on Cost Savings

II. New Companies and Capital Are Fueling What Could Become a Trillion-Dollar Robotics Market — With Strategic Investments Already Underway from Tech and Industry Leaders

New Companies and Investments

-  **Atlas – Boston Dynamics**
- Hyundai invested \$21 billion in partnership with Boston Dynamics, including the backbone for the production and deployment of robots like Atlas
-  **Figure 01 – Figure AI**
- NVIDIA invested approximately \$50 million in Figure AI, reinforcing its strategy to expand embedded AI in robotics
- Ambitious production target: up to 200,000 robots delivered by 2029
-  **Optimus - Tesla**
- Mass production of thousands of units is expected to begin in 2025, with plans to scale to 1 million units per year by the end of the decade
- “Optimus has the potential to be more significant than the car business long term” – Elon Musk

Robotics Industry (US\$ million)



Robotics II

A New Revenue Stream Emerging Beyond GPUs, Robotics Could Be NVIDIA's Next Growth Engine

I. Robotics Runs on Compute — And NVIDIA Provides the Performance and Platforms to Make It Real

Why Robots Need GPUs?

Modern robots must interpret complex sensor data, make real-time decisions, and interact with the physical world — all of which require high-throughput, low-latency AI inference

What Does GPU enable the Robot to do?

➔ **See:** Computer Vision

➔ **Think:** AI models for reasoning and planning

➔ **Act:** Real-time control and movement

NVIDIA Robotics Products



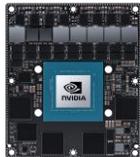
Jetson Orin Nano
US\$ 249
2023
Designed for entry-level Robots



Jetson AGX Orin
US\$ 1,500
2024
Designed for High performance Robots



Jetson Orin NX
US\$ 595
2023
Designed for Mid-tier robots



Jetson AGX Thor
~US\$ 3,000
Expected for 2025
Designed for Humanoids Robots

II. NVIDIA's Full-Stack Strategy to Monetize the Growing Robotics Market, Building a Developer-Led Ecosystem Like It Did with CUDA in AI

NVIDIA Platforms to monetize Robotics

Isaac Platform

NVIDIA Isaac is a modular platform for developing, simulating, and deploying autonomous robots, combining AI perception, planning, and control in one stack

Omniverse

Omniverse enables realistic, physics-based simulation of robots in virtual environments, accelerating training and testing before real-world deployment

Cosmos

Cosmos is a cloud platform for monitoring and managing robot fleets — allowing updates, orchestration, and real-time operations at scale

NVIDIA's Clients

Isaac Platform



Omniverse



Cosmos



Jensen Hung, Nvidia's Founder & CEO

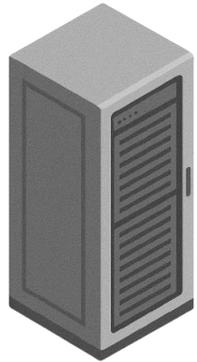
"I think this is likely to be the next multi-trillion-dollar industry"



Current Overview

NVIDIA is already delivering a product roadmap to establish itself as the AI fabric

I. Despite its explosive progress, the company continues to increase product performance exponentially, with an extreme leap in the Rubin architecture



Rubin Ultra NVL144

Second Half 2026

3.3x GB300 NVL72

13 TB/s HBM4
75TB Fast Memory

1.6x

260 TBs NVLink6
2x

Rubin Ultra NVL576

Second Half 2027

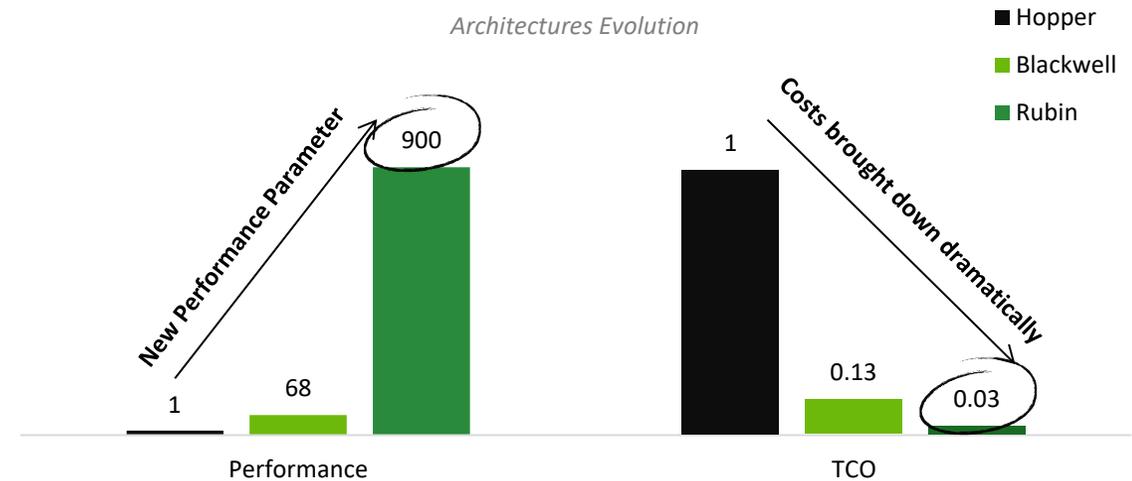
14x GB300 NVL72

4.6 PB/s HBM4e
365TB Fast Memory

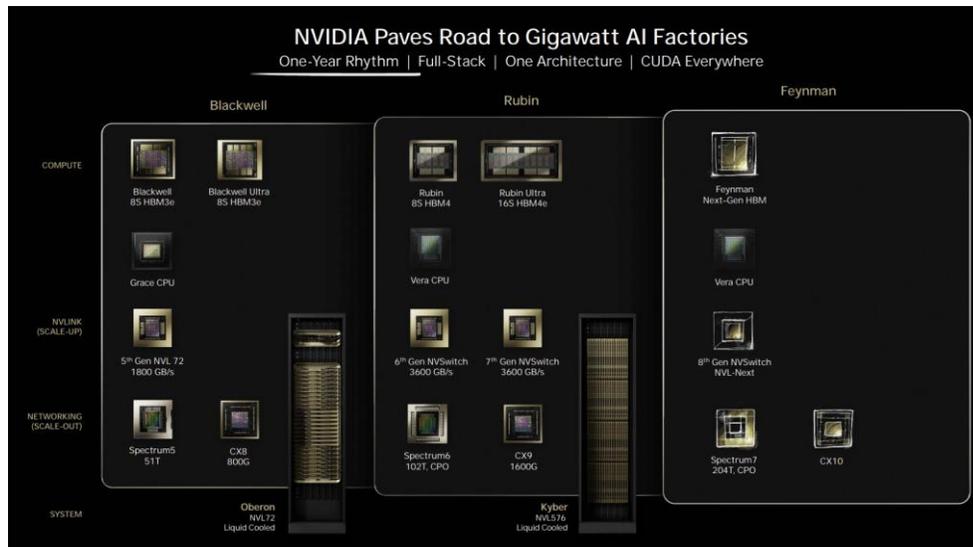
8x

1.5 PBs NVLink7
8x

Architectures Evolution



II. Having aligned its roadmap, the company is moving decisively across compute, networking, and software to cement its role as the cornerstone of the AI revolution



Present:



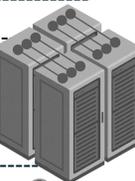
GB200
1 Grace GPU
2 Blackwell GPUs



GB200 NVL72
36x GB200 Superchips
72 GPUs + 36 CPUs

Next Step:

Connecting GPUs across entire racks



Future:

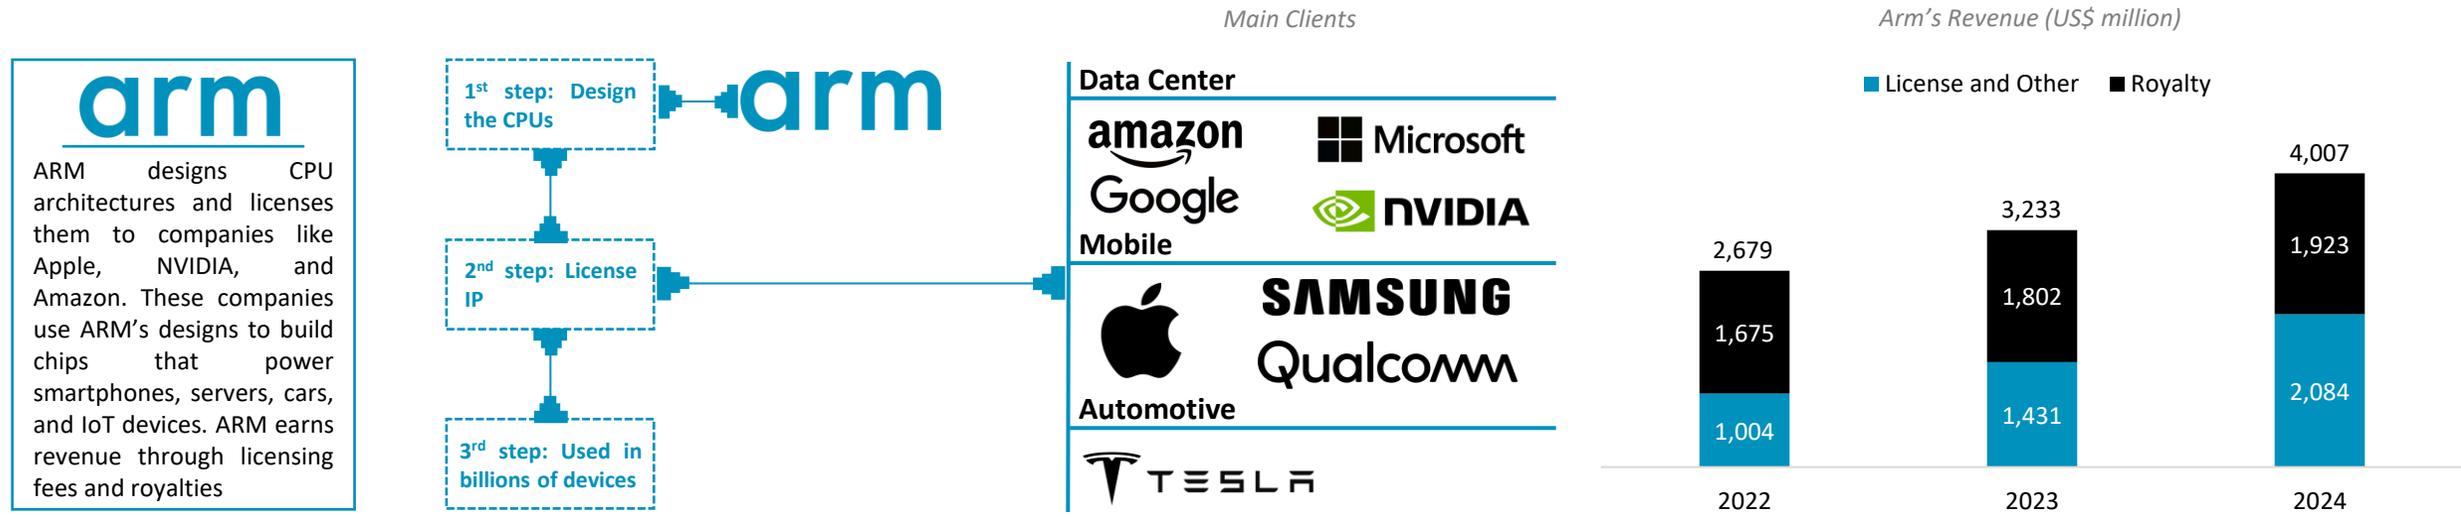
AI Fabric: A One-stop Shop



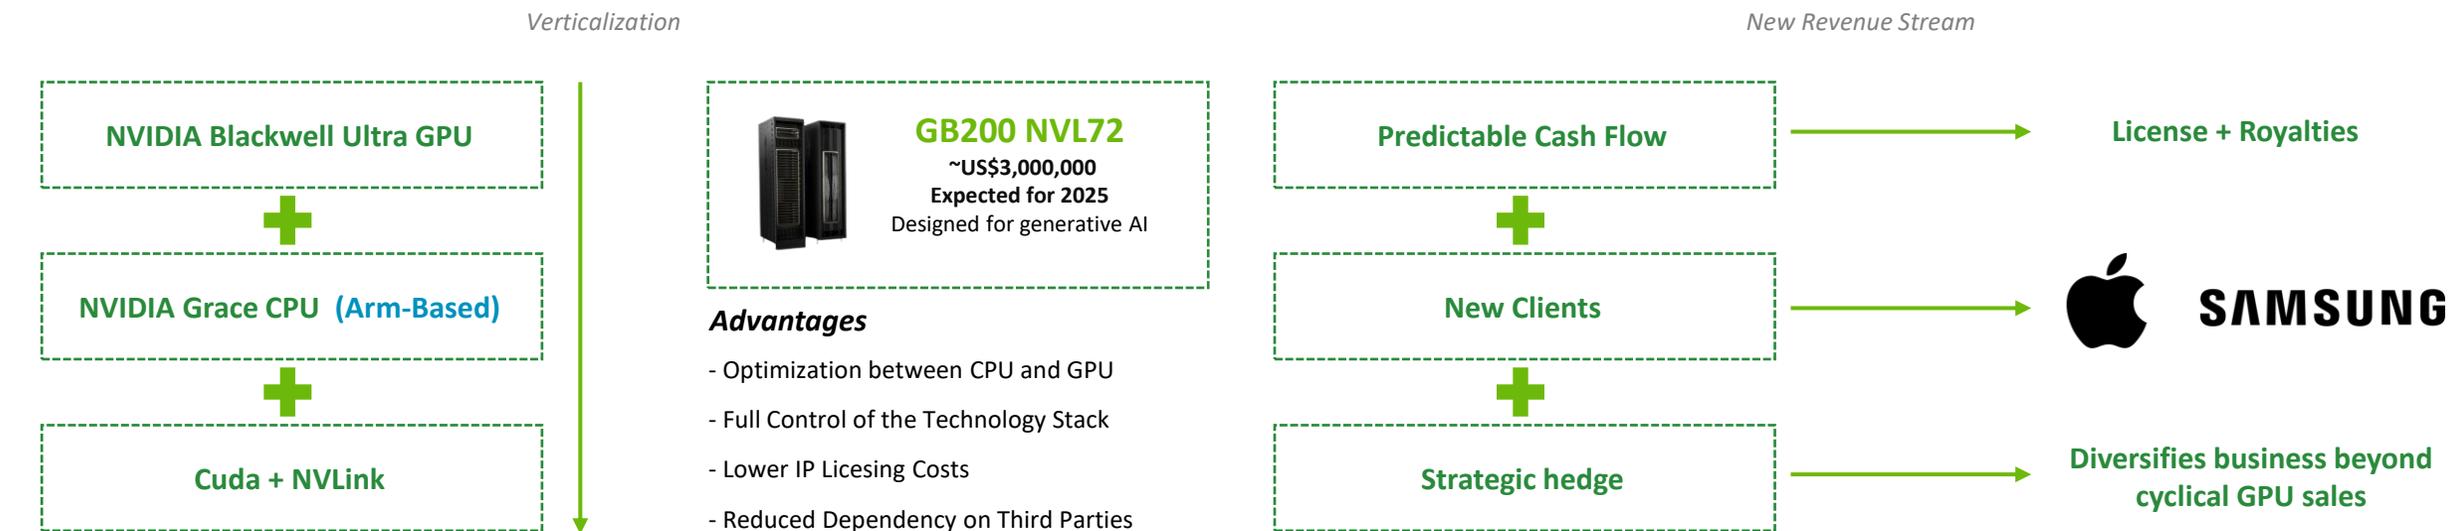
Arm

Acquiring ARM strengthens NVIDIA's positioning by uniting architectural control, product optimization, and recurring global revenue

I. ARM Designs the Brains of Modern Chips — Powering Billions of Devices and Generating Scalable, Recurring Revenue Through Licensing and Royalties



II. Acquiring ARM Strengthens NVIDIA's Full-Stack Technological Control — While Unlocking Recurring, Scalable Revenue Through Global Licensing and Royalties



Ampere Computing

Acquiring Ampere gives NVIDIA CPU control, enables full-stack AI optimization, and expands data center value capture

I. Ampere's Value Chain and Strategic Positioning in the Transition Toward Arm-Based Cloud Infrastructure

Overview

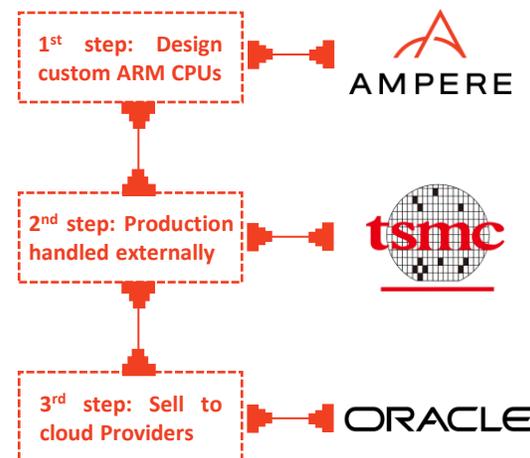
Value Chain



AMPERE

Ampere designs custom ARM-based CPUs optimized for cloud and data center workloads. Its chips are used by companies like Oracle and manufactured by partners like TSMC. Ampere earns revenue by selling these processors to cloud infrastructure providers

- ➔ **Founded:** 2017 by Renee James (ex-Intel President)
- ➔ **Focus:** Custom Arm-based CPUs for Hyperscalers
- ➔ **Business Model:** Fabless
- ➔ **Key Customers:** Oracle
- ➔ **Headquarters:** Santa Clara, California, USA

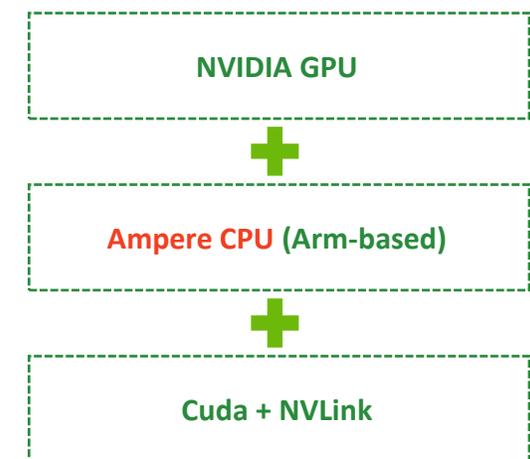
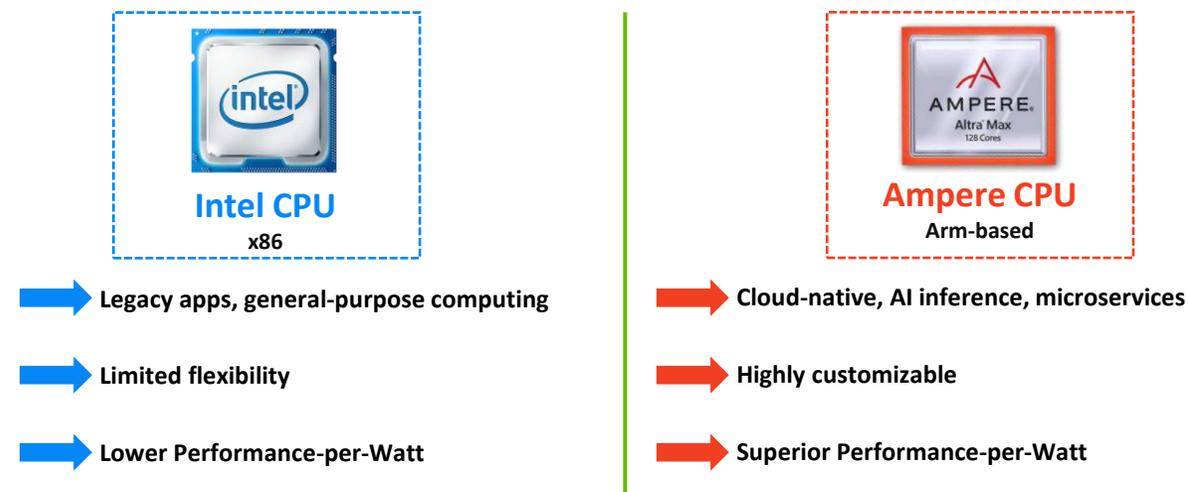



Ampere Ultra Max
Custom Arm-based CPU
Clients: Oracle, Equinix, HPE
2025
Designed for Hyperscalers and AI

II. Strategic Fit: How Ampere Fills NVIDIA's CPU Gap and Accelerates Its Vision for Full-Stack AI Infrastructure

CPUs Comparison

Verticalization

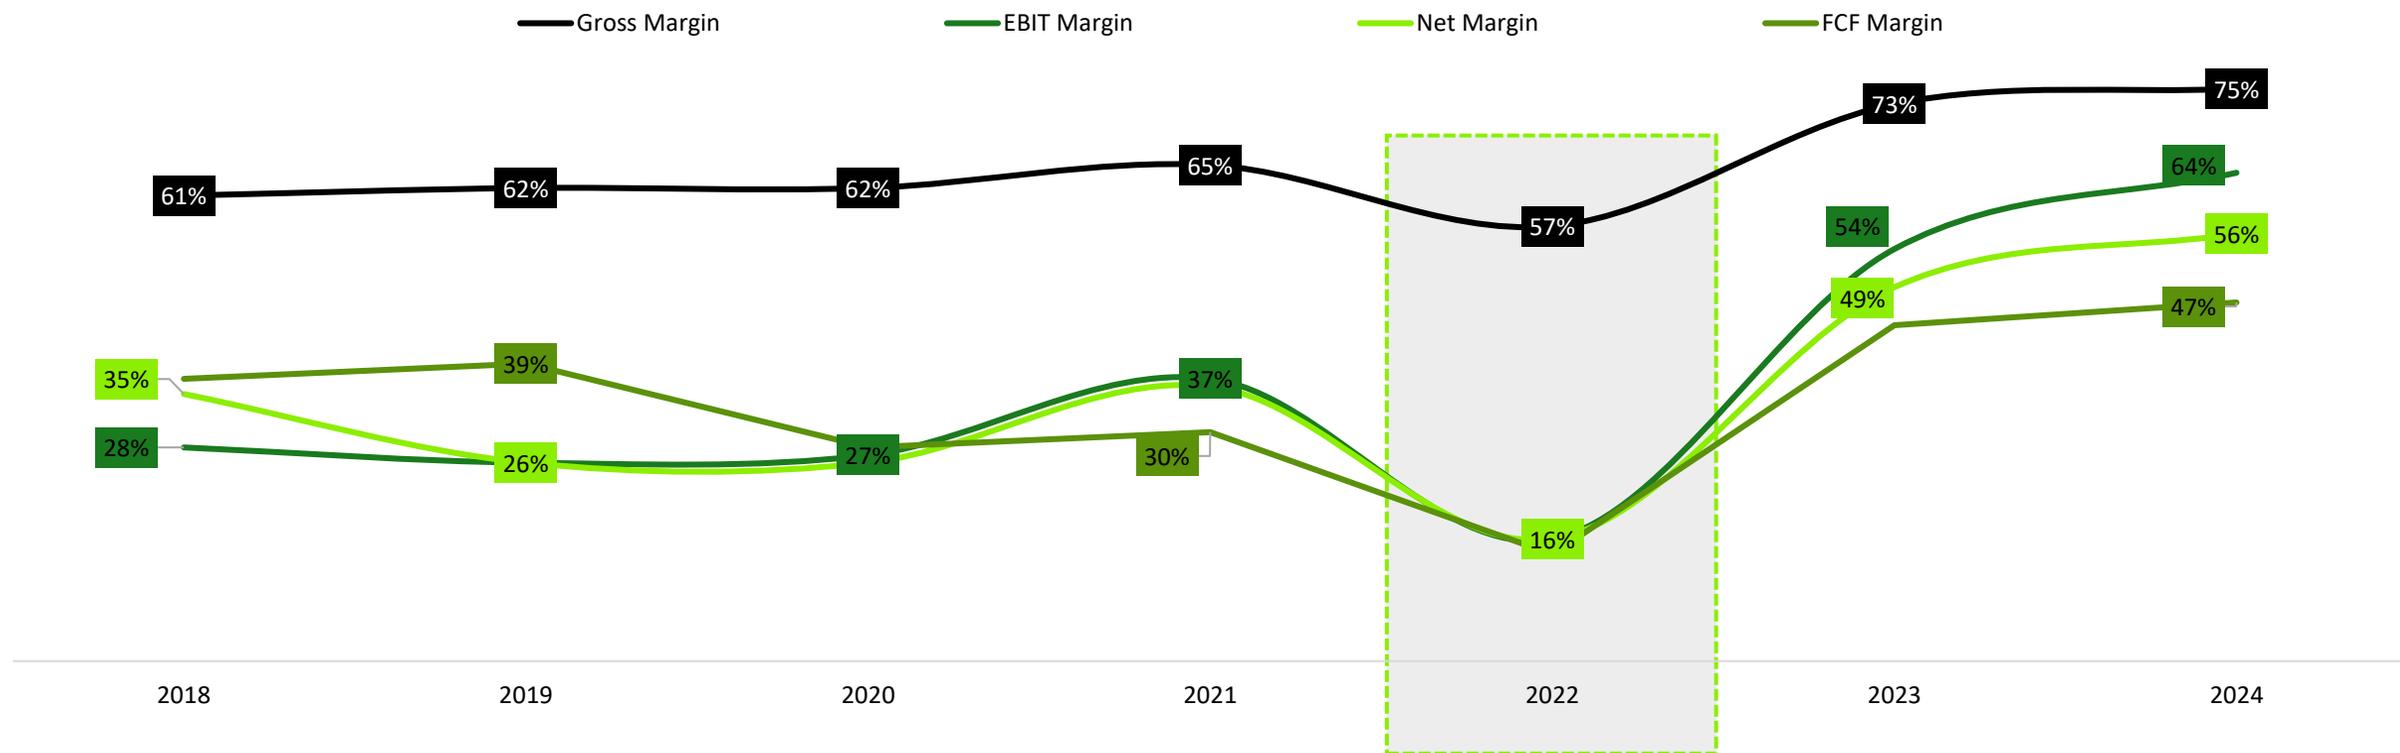


By combining NVIDIA GPUs, Ampere Arm-based CPUs, and its proprietary CUDA software + NVLink interconnect, NVIDIA is moving toward a fully vertically integrated AI computing platform

- Advantages**
- End-to-end performance gains
 - Less dependency on Intel/AMD
 - Greater value per system (hardware + software)
 - Faster rollout of AI/HPC/cloud chips
 - Tighter ecosystem with stronger competitive edge



What happened in Nvidia Margin in 2022



Gross margin for fiscal year 2023 declined from a year ago, driven by \$2.17 billion of inventory charges largely relating to excess supply of NVIDIA Ampere architecture Gaming and Data Center products as compared to the demand expectations for these products, particularly for the expected demand in China. The inventory charges were comprised of \$1.04 billion for inventory on hand and \$1.13 billion for inventory purchase obligations in excess of our demand expectations

Firsts restriction to China

NVIDIA 10-K 2023FY Q4



Value Added

How we conducted value added account

Unit: US\$ Milion

x Nvidia Price Today (04/06/2025)

Price today	139,99
Shares	24.400
Nvidia Market Cap	3.415.756
Total Debt	8.463
Total Cash	43.210
Net Debt	(34.747)

x Revenue by segment in 2024

Gaming	11.350
Professional visualization	1.878
Automotive	1.694
OEM & Others	389

x Value Added

Data center + Adjacenses Value Added	3.299.015
	96,6%
Data Center EV/Sales	28,34x

x EV/Sales Multiple

Gaming	4,38x
AMD	6,95x
Sony	1,80x
Professional visualization	8,88x
Autodesk	10,30x
Dassault Systems	7,44x
PTC	8,91x
Automotive	5,52x
Mobileye	7,35x
Qualcomm	3,69x
OEM & Others	16,21x
Broadcomm	22,50x
Marvell Tech	9,91x

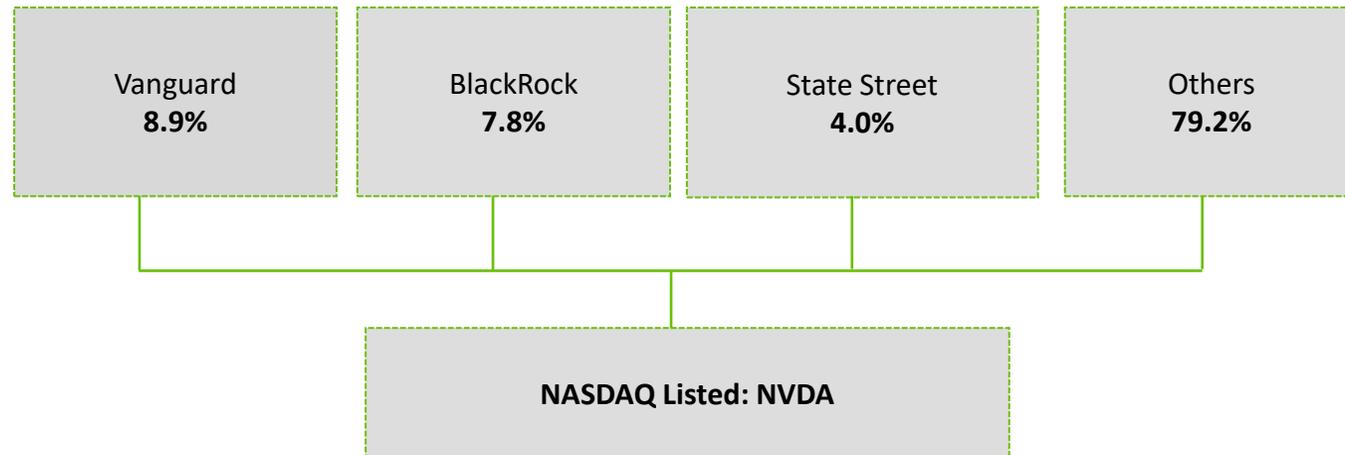
x Core Value - Net Debt

Gaming	49.656
Professional visualization	16.683
Automotive	9.351
OEM & Others	6.304
Core Value	81.994
Net Debt	(34.747)
Core - Net Debt	116.741



Shareholder structure

Shareholder Structure



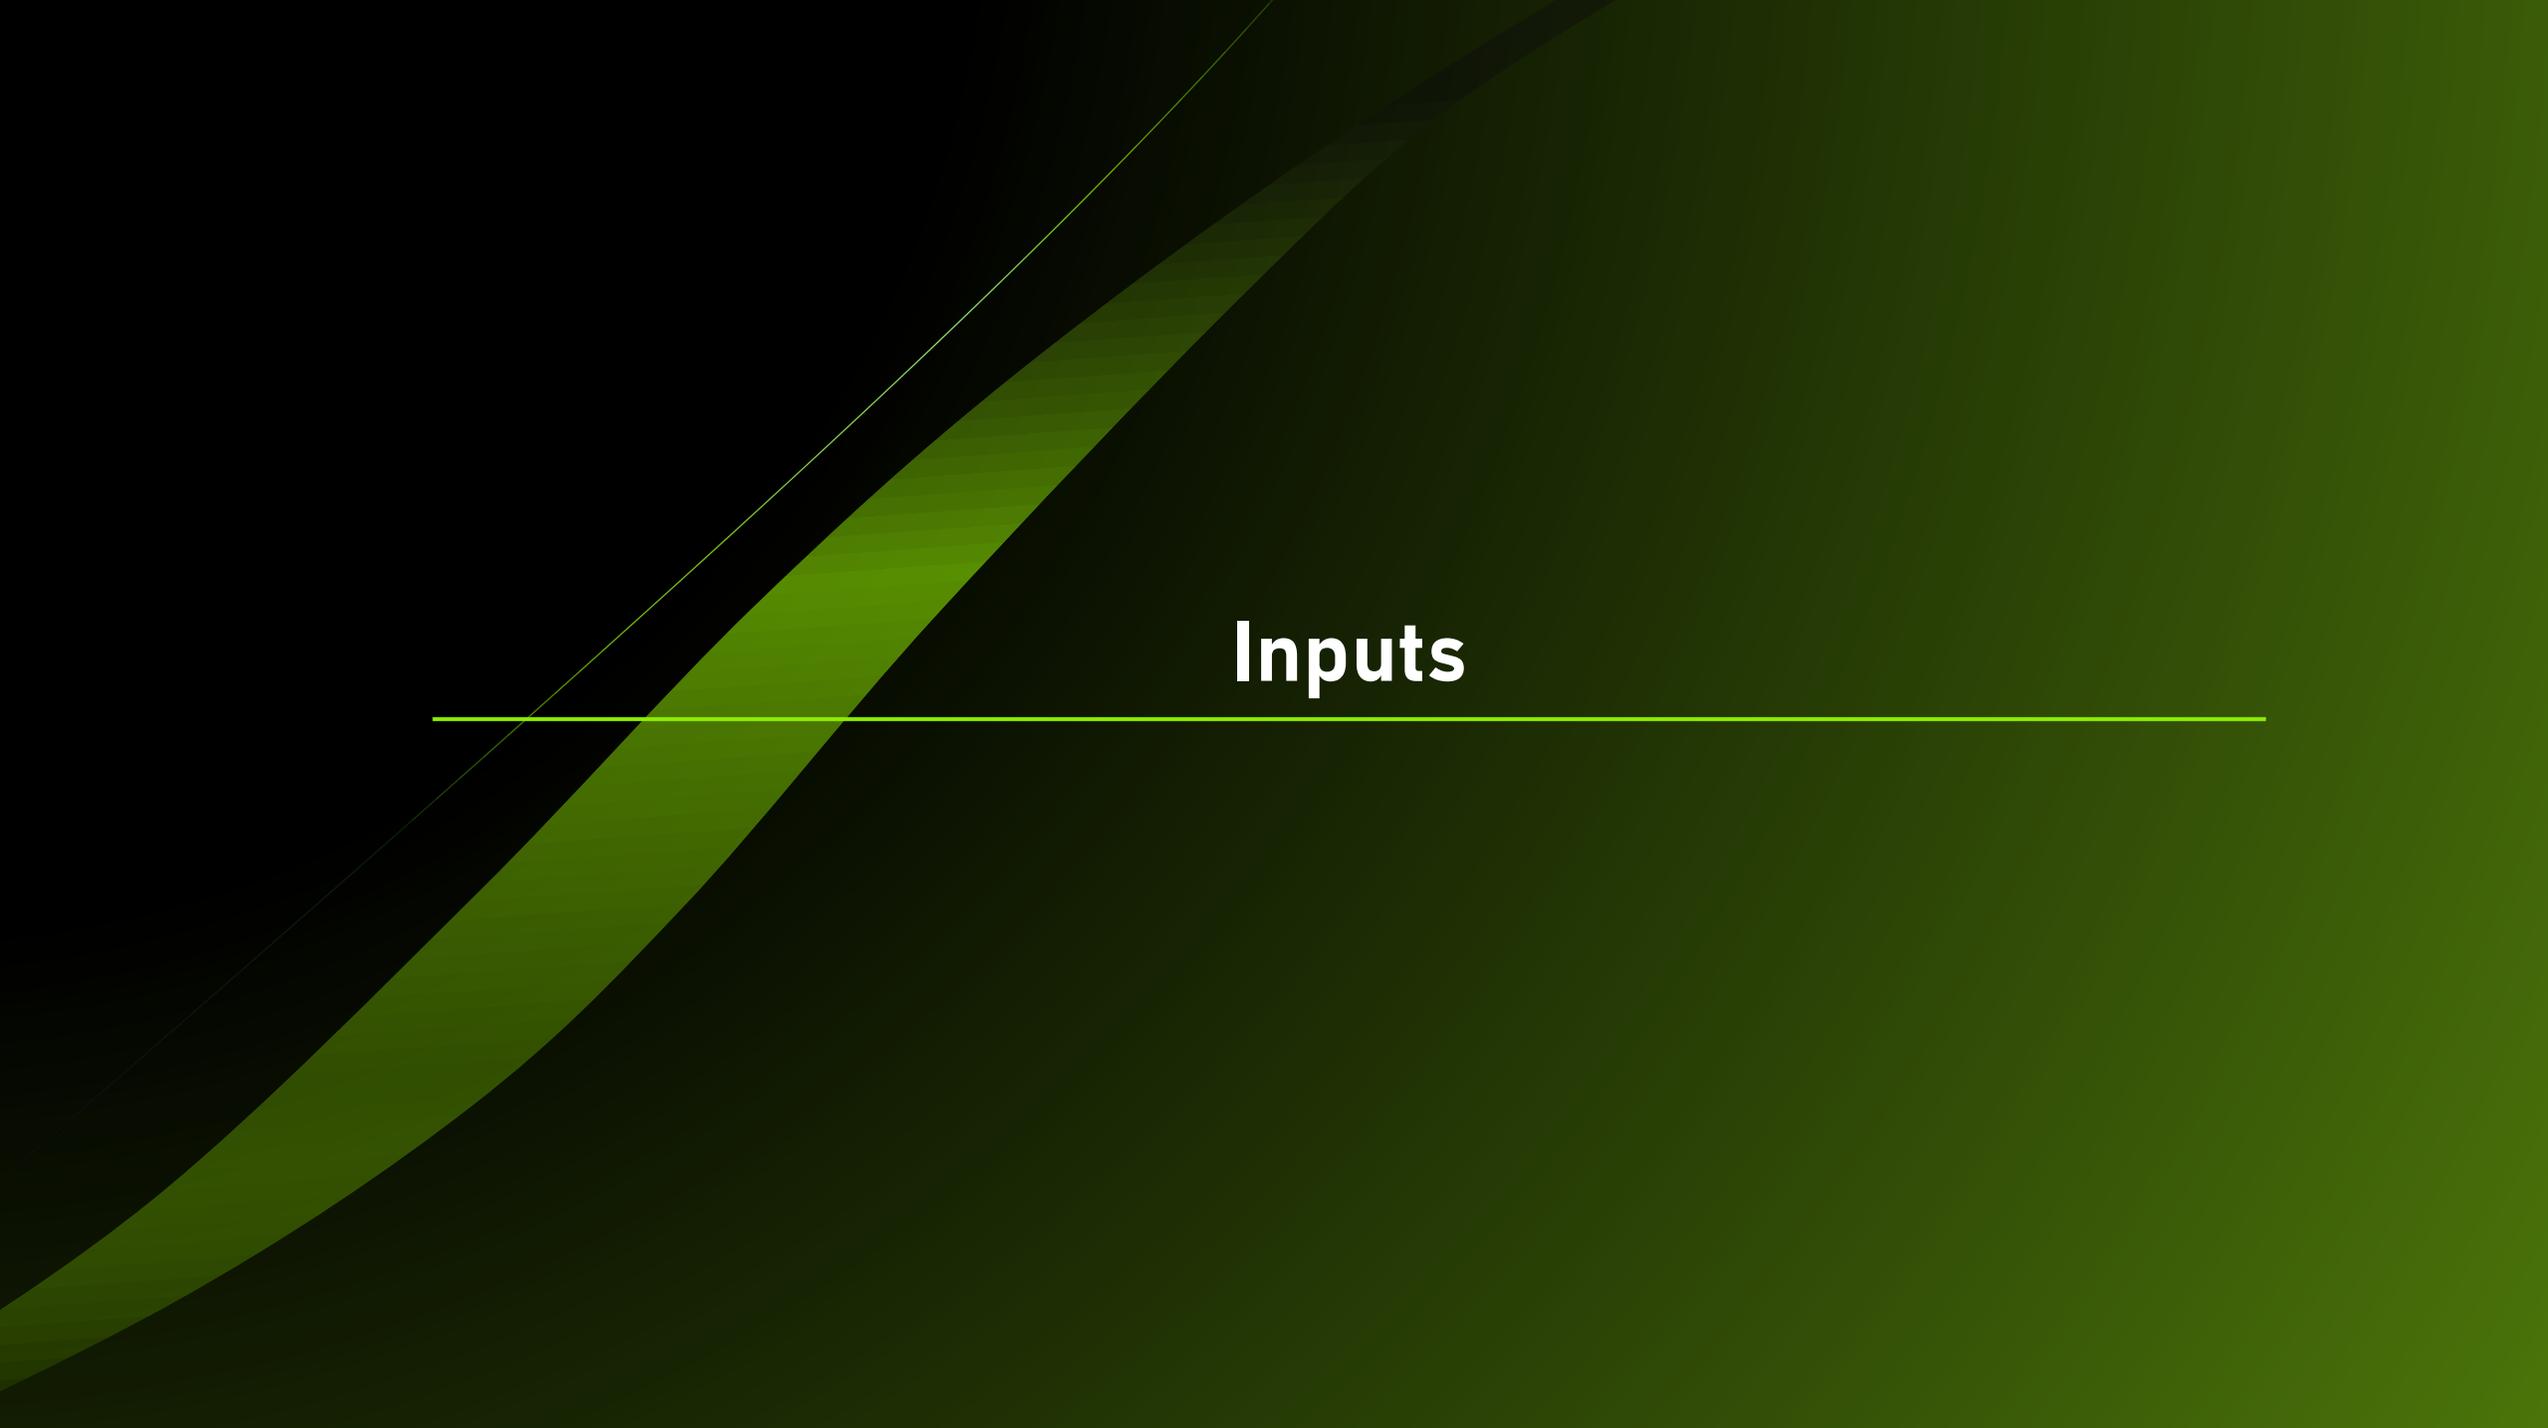
Performance Metrics & Goals

Fiscal 2025 performance metrics and goals for NEO pay were as set forth below:

PERFORMANCE METRICS			
	Variable Cash Plan	SY PSUs	MY PSUs
Metric	Revenue	Non-GAAP Operating Income	TSR relative to the S&P 500
Timeframe	1 year	1 year	3 years
CC's Rationale for Metric	Drives value, contributes to Company's long-term success Focuses on growth in new and existing markets Distinct, separate metric from Non-GAAP Operating Income	Drives value, contributes to Company's long-term success Reflects our annual revenue generation and effective operating expense management Distinct, separate metrics from revenue	Aligns directly with long-term shareholder value creation Provides comparison of our stock price performance, including dividends, against a capital market index in which we compete Relative performance goal accounts for macroeconomic factors impacting the market

PERFORMANCE GOALS						
	Variable Cash Plan		SY PSUs		MY PSUs	
	Fiscal 2025 Revenue	Payout as a % of Target Opportunity (1)	Fiscal 2025 Non-GAAP Operating Income (2)	Shares Eligible to Vest as a % of Target Opportunity (1)	Fiscal 2025 to 2027 3-Year Relative TSR (3)	Shares Eligible to Vest as a % of Target Opportunity (1)
Threshold	\$45.0 billion	50%	\$16.0 billion	50%	25th percentile	25%
Base Compensation Plan	\$90.0 billion	100%	\$56.0 billion	100%	50th percentile	100%
Stretch Compensation Plan .	\$110.0 billion	200%	\$72.0 billion	CEO 150% Other NEOs 200%	75th percentile	CEO 150% Other NEOs 200%



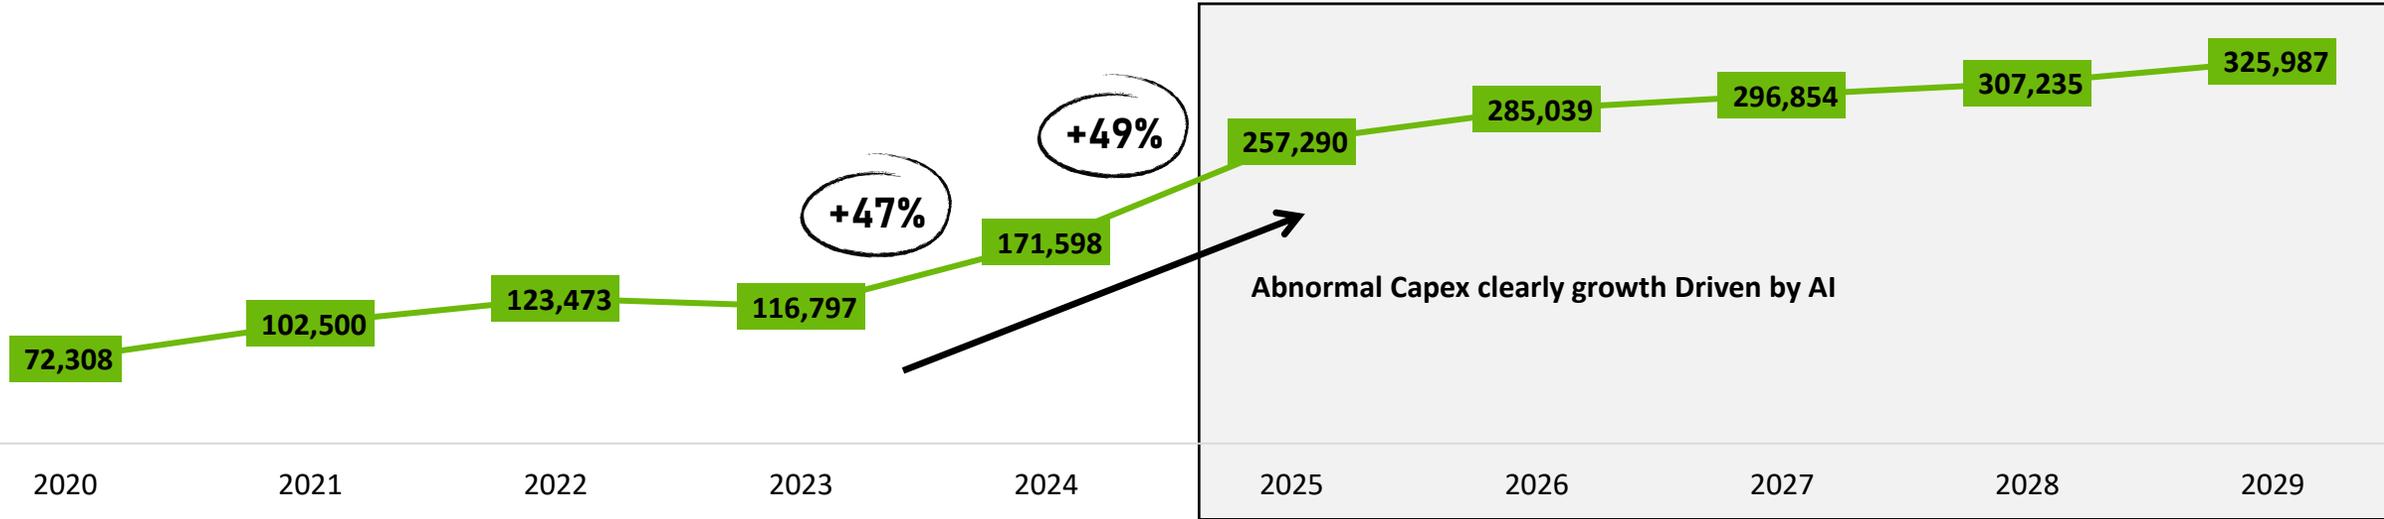
The image features a dark green background with a lighter green curved band that starts from the bottom left and curves upwards towards the top right. A thin, light green horizontal line spans across the middle of the image. The word "Inputs" is written in white, bold, sans-serif font, centered above the horizontal line.

Inputs

Marginal Capex

With the AI race underway, companies are allocating the majority of their marginal capex to Data Center's

I. The abnormal growth in CapEx signals an acceleration of investments in AI. This clearly indicates a strategic shift by big tech companies toward the new wave of innovation



II. Everyone is allocating incremental capital to AI, and NVIDIA is already seeing the impact in its revenue



Meta raised its 2025 CapEx guidance to US \$64 billion. CEO Mark Zuckerberg confirmed that US \$62 billion of that total will be directly devoted to AI infrastructure



Facing Microsoft–OpenAI's "Stargate" super-cluster, Oracle must raise CapEx to keep OCI's compute and latency competitive **while maintaining its sovereign-cloud edge**



Google expects to spend US \$75 billion in CapEx for 2025. CFO Anat Ashkenazi noted that **most of it will go to technical infrastructure, mainly servers and data centers**

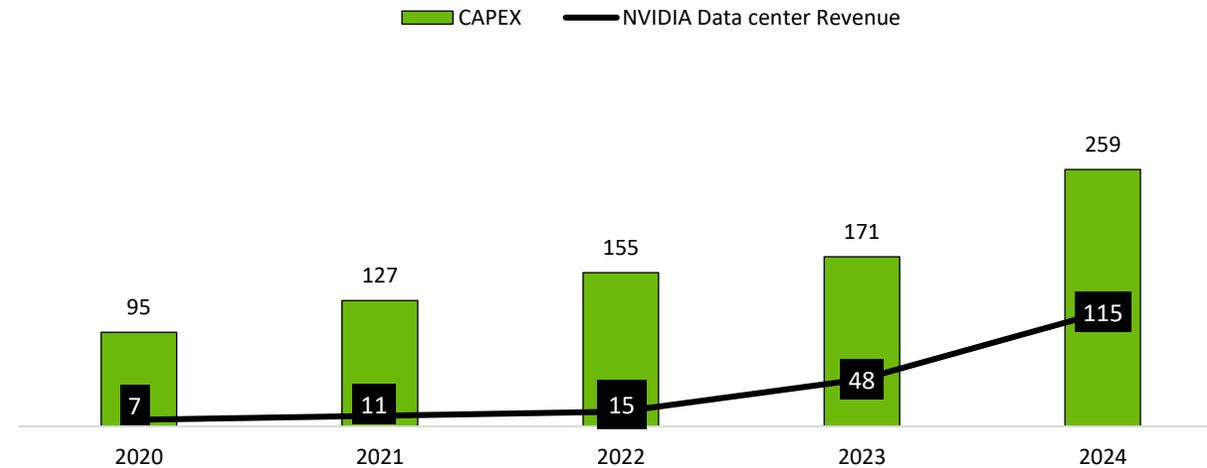


In February, Jassy said Amazon expects to spend more than **\$100 billion** in capex, the "vast majority" of which would go to AI infrastructure for Amazon Web Services



Microsoft plans to invest around US \$80 billion in data centers in FY25. CFO Amy Hood said capex will grow more slowly and include a higher share of short-lived assets

NVIDIA Data Center Revenue x Big 4 Tech Companies CAPEX

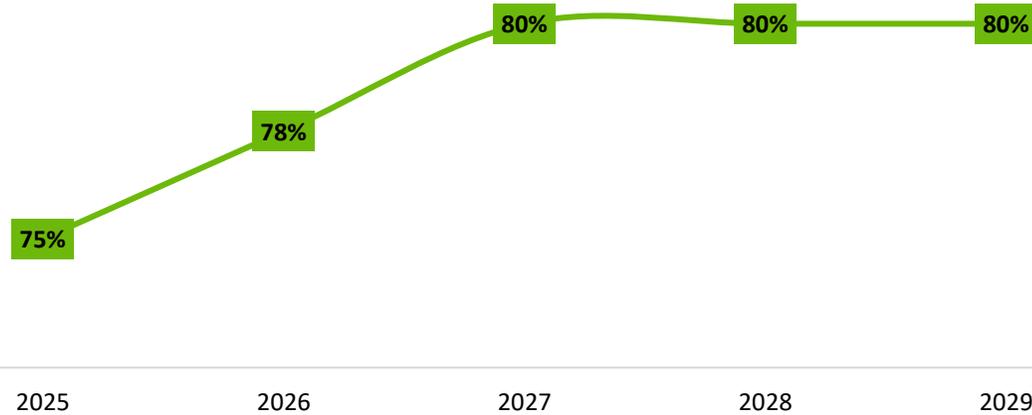


How Marginal Capex reflects in our Model?

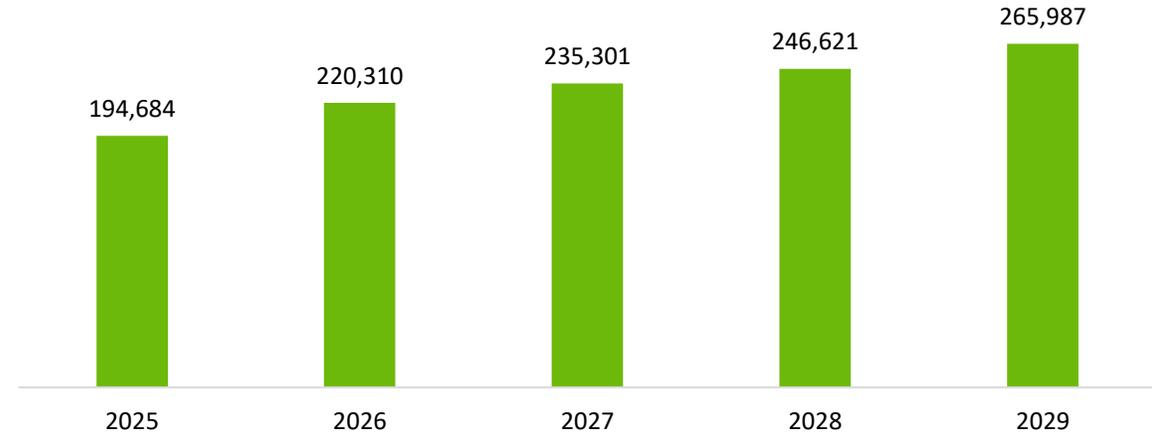
Data Center-focused Capex supports robust returns across a wide range of scenarios

I. For the following years, we assumed that the majority of marginal capex would be allocated to Data Centers, as indicated by the hyperscalers

% of Capex Destined for Data Center in Marginal Capex



Marginal Capex destined to Data center (US\$ Million)



II. When performing a sensitivity analysis, we found that in the vast majority of cases we still achieve a satisfactory IRR, which provides us with a margin of safety

IRR Sensitivity Analysis

Accelerators Capex (%)	Marginal Capex (%)						
	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%	95.0%
45.0%	9.3%	10.6%	11.8%	13.1%	14.3%	15.5%	16.7%
50.0%	12.8%	14.2%	15.5%	16.8%	18.1%	19.4%	20.6%
55.0%	16.2%	17.6%	19.0%	20.4%	21.7%	23.0%	24.3%
60.0%	19.4%	20.9%	22.3%	23.8%	25.1%	26.5%	27.8%
65.0%	22.5%	24.0%	25.5%	27.0%	28.4%	29.8%	31.2%
70.0%	25.4%	27.0%	28.5%	30.1%	31.6%	33.0%	34.5%
75.0%	28.2%	29.9%	31.5%	33.0%	34.6%	36.1%	37.6%

IRR exceeded 15% in 86% of the cases

IRR Sensitivity Analysis

Market Share (%)	Marginal Capex (%)						
	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%	95.0%
65.0%	11.3%	12.7%	14.0%	15.2%	16.5%	17.7%	18.9%
70.0%	13.9%	15.2%	16.6%	17.9%	19.2%	20.5%	21.7%
75.0%	16.3%	17.7%	19.1%	20.5%	21.8%	23.1%	24.4%
80.0%	18.7%	20.1%	21.5%	22.9%	24.3%	25.7%	27.0%
85.0%	20.9%	22.4%	23.9%	25.3%	26.7%	28.1%	29.5%
90.0%	23.1%	24.7%	26.2%	27.6%	29.1%	30.5%	31.9%
95.0%	25.3%	26.8%	28.4%	29.9%	31.3%	32.8%	34.2%

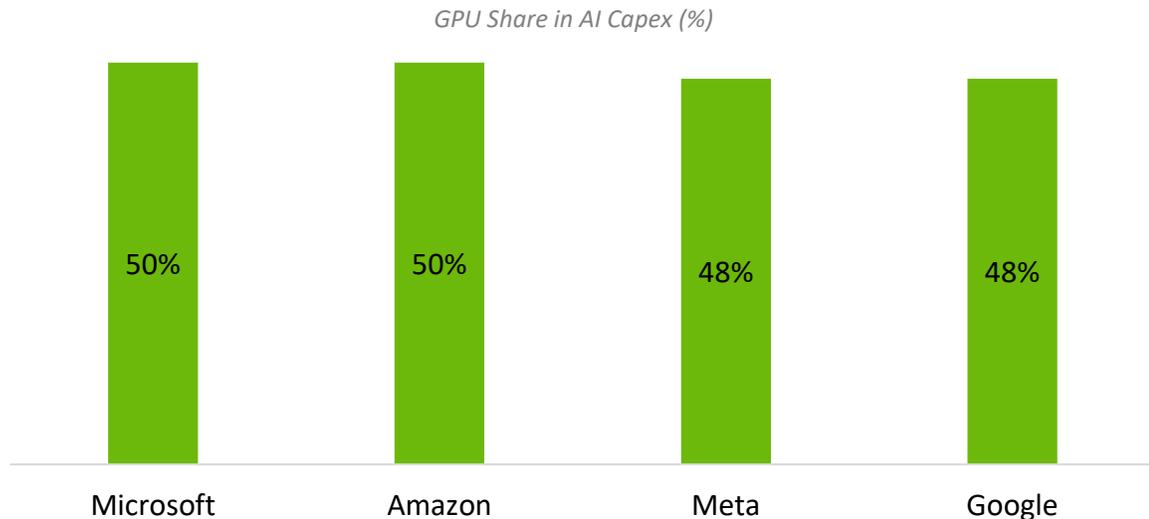
IRR exceeded 15% in 92% of the cases



GPUs Share

With the increased pricing power of GPUs, coupled with infrastructure stabilization, we believe in a rising share of overall spending

I. Today's GPU Share in AI Capex

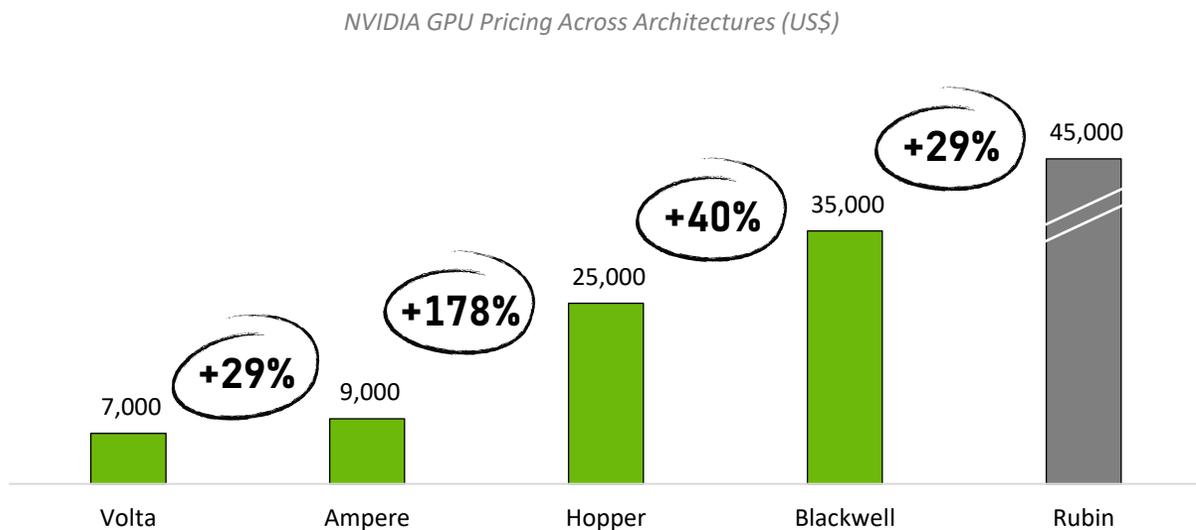


II. Unit Economics of a Data Center indicates that GPUs cost could get to 72%

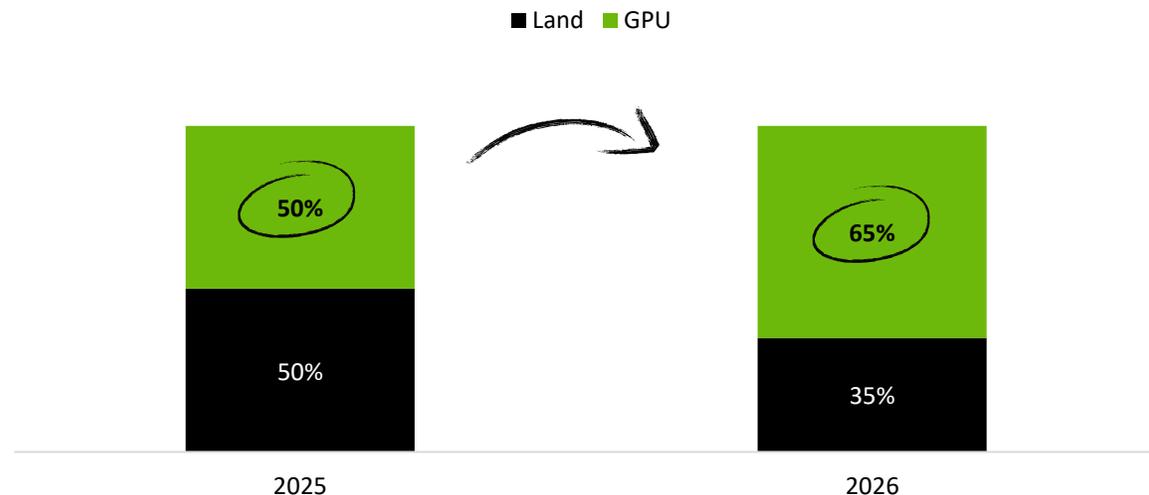
Hyperscaler capex split	Comment	% split
Power	Infra	25-30%
Building	Infra	10-15%
Cooling	Infra	10-15%
Networking	Chip	10-15%
Compute	Chip	40-45%

NVIDIA DGX H100 (Hardware Build Up)'	AI Server	%
CPU	5,200	
8 GPU + 4 NVSwitch Baseboard	195,000	72.4%
Memory	7,860	
Storage	3,456	
Others	57,494	
Total Cost	269,010	

III. But pricing power is likely to grow



IV. Microsoft's numbers already indicate a growth in the GPU share of the AI Capex mix

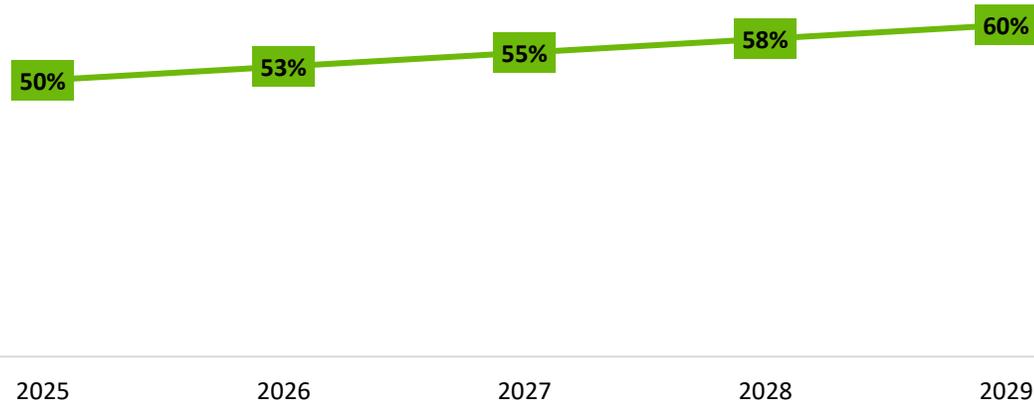


How GPU's Share reflects in our Model?

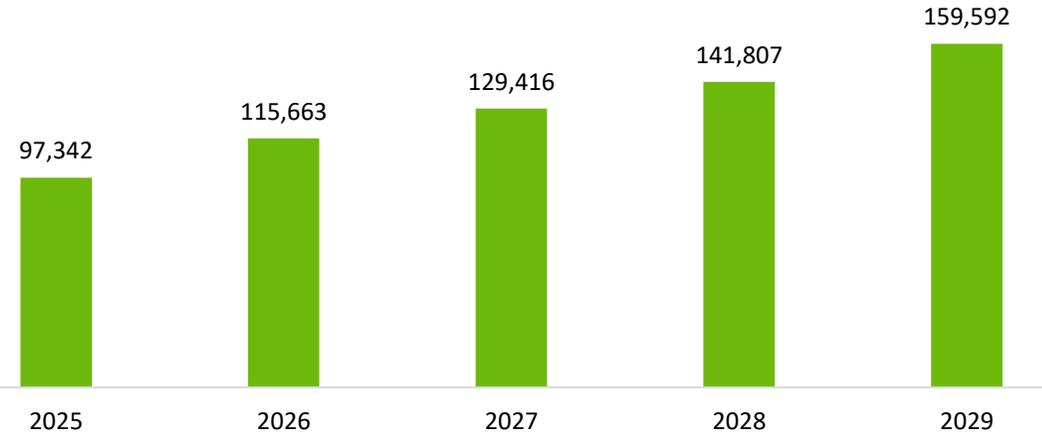
Rising GPU share supports investment case with solid IRR sensitivity

I. We increased the GPU share as a consequence of the companies' growing pricing power, as well as the need to replace them before the broader infrastructure

% of Hyperscalers Technology Capex (GPU Share)



Hyperscalers Accelerators TAM (US\$ Million)



II. When performing a sensitivity analysis, we found that in the vast majority of cases we still achieve a satisfactory IRR, which provides us with a margin of safety

IRR Sensitivity Analysis

Accelerators Capex (%)	Marginal Capex (%)						
	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%	95.0%
45.0%	9.3%	10.6%	11.8%	13.1%	14.3%	15.5%	16.7%
50.0%	12.8%	14.2%	15.5%	16.8%	18.1%	19.4%	20.6%
55.0%	16.2%	17.6%	19.0%	20.4%	21.7%	23.0%	24.3%
60.0%	19.4%	20.9%	22.3%	23.8%	25.1%	26.5%	27.8%
65.0%	22.5%	24.0%	25.5%	27.0%	28.4%	29.8%	31.2%
70.0%	25.4%	27.0%	28.5%	30.1%	31.6%	33.0%	34.5%
75.0%	28.2%	29.9%	31.5%	33.0%	34.6%	36.1%	37.6%

IRR exceeded 15% in 82% of the cases

IRR Sensitivity Analysis

Accelerators Capex (%)	Market Share NVIDIA						
	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%	95.0%
45.0%	6.0%	8.4%	10.6%	12.8%	14.9%	17.0%	19.0%
50.0%	9.4%	11.8%	14.2%	16.5%	18.7%	20.9%	23.0%
55.0%	12.6%	15.2%	17.6%	20.0%	22.4%	24.6%	26.8%
60.0%	15.6%	18.3%	20.9%	23.4%	25.8%	28.1%	30.4%
65.0%	18.5%	21.3%	24.0%	26.6%	29.1%	31.5%	33.9%
70.0%	21.3%	24.2%	27.0%	29.7%	32.2%	34.7%	37.2%
75.0%	24.0%	27.0%	29.8%	32.6%	35.3%	37.8%	40.3%

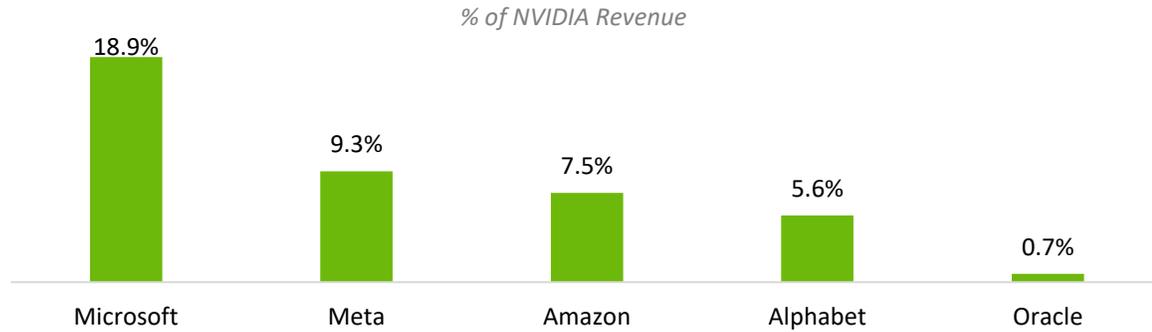
IRR exceeded 15% in 84% of the cases



Hyperscalers % of Revenue

Sovereign AI and corporate AI are the new sources of revenue diversification for NVIDIA

I. NVIDIA's Data Center revenue is composed of 42% from the world's major hyperscalers



Alexander Karp, Palantir CEO, **The Technological Republic**

The technologies we are creating — including innovative forms of AI that may challenge the country's current monopoly over global creative control — are themselves a product of a culture whose preservation and development we can no longer afford to neglect

III. Corporate AI is expected to grow, and NVIDIA is actively expanding in this space

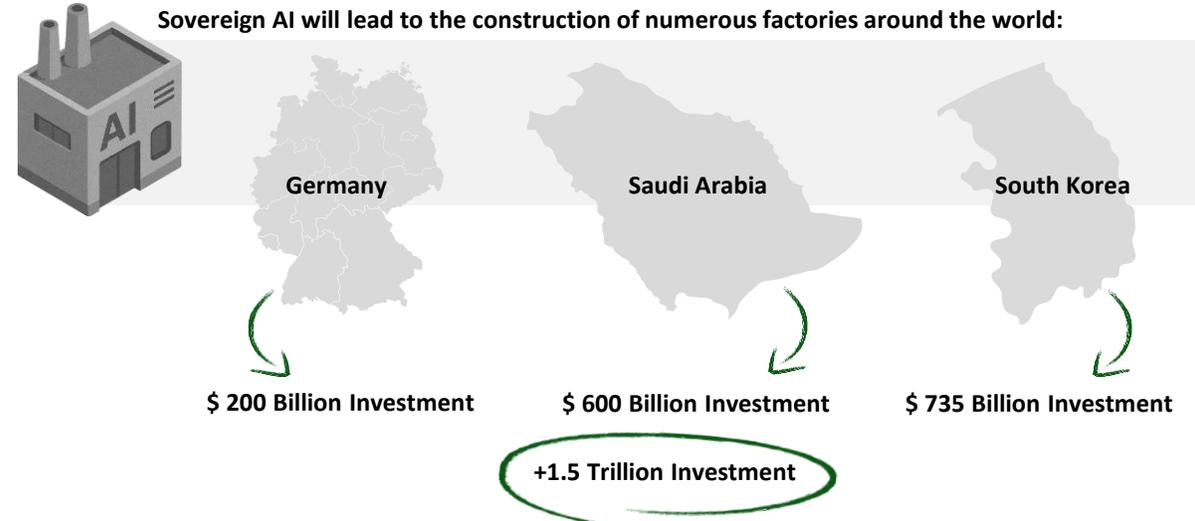
FINANCIAL TIMES

Nvidia seeks to build its business beyond Big Tech

Over recent months, Nvidia has also struck alliances with suppliers, including Cisco, Dell and HP, to help sell to enterprise customers, which manage their own corporate IT infrastructure instead of outsourcing to the cloud



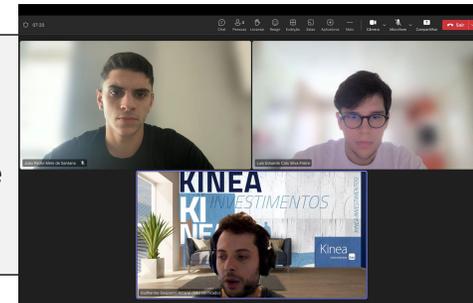
II. Sovereign AI will drive data center growth and increase demand for GPUs



IV. The share of hyperscalers in NVIDIA's revenue is expected to decrease by at least 2%/y

Guilherme Amaral, Kinea Global Analyst

“We expect at least a **2%/year compression** over time in the contribution of CSPs to NVIDIA's revenue. This is largely driven by the rise of sovereign AI demand and the company's ability to diversify its end markets”



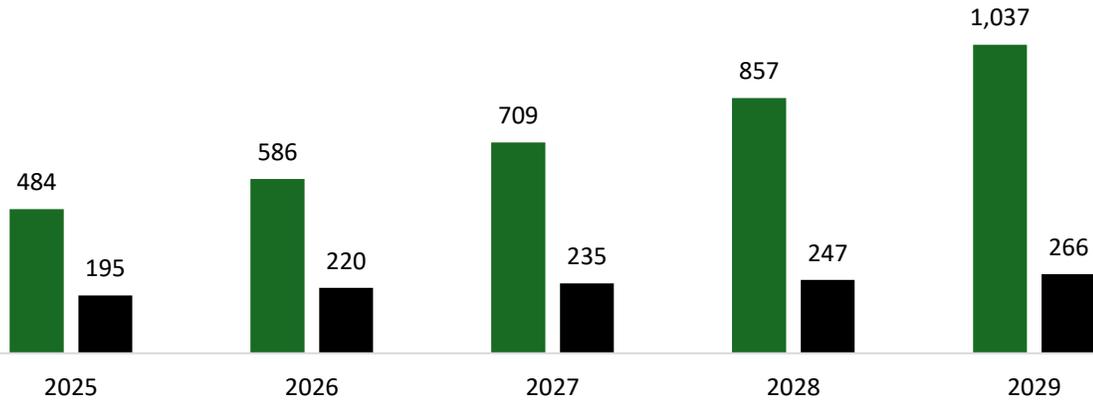
Ex-Hyperscalers % of Revenue

As other applications grow, hyperscalers become less relevant to total revenue

I. Hyperscalers' capex becomes relatively smaller as other applications grow

Data Center Capex

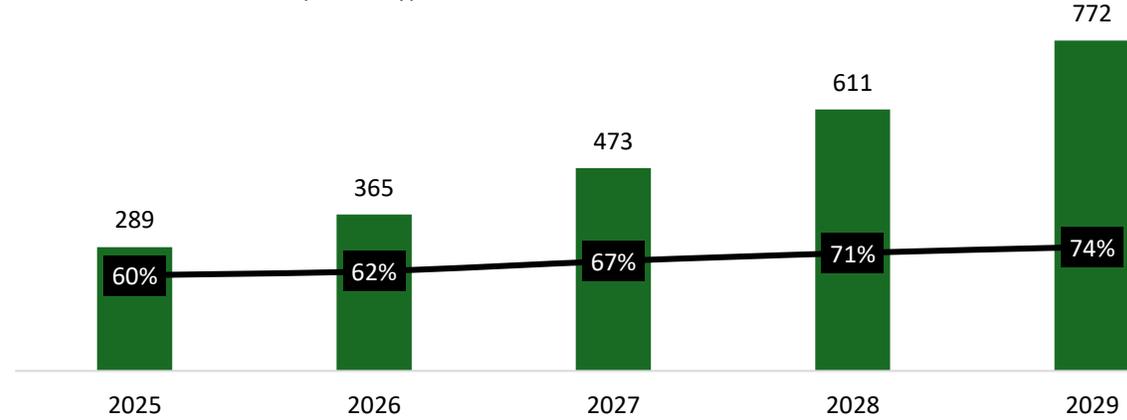
■ Dell'oro Total Capex ■ Hyperscalers Capex



II. As a result, a large portion of total capex shifts toward other players

DC Capex Ex-Hyperscalers

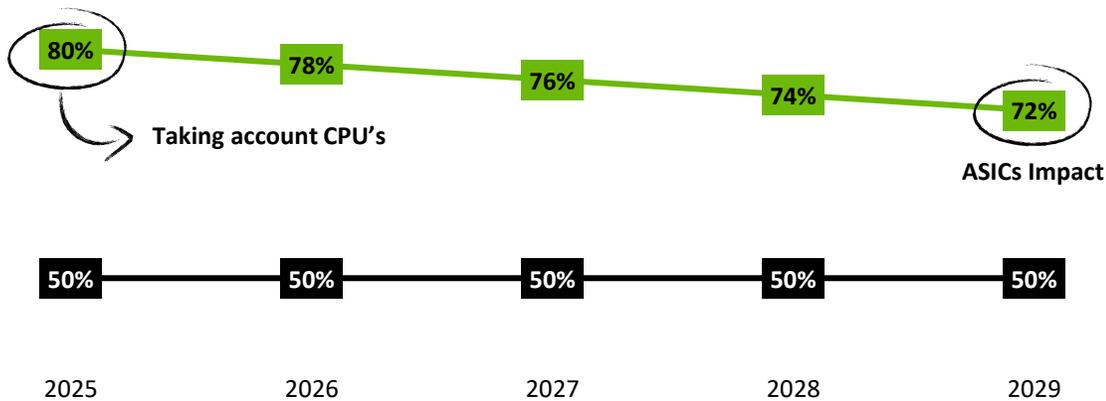
■ Capex Ex-Hyperscalers — % of Total Capex



III. Due to lower compute needs, NVIDIA holds less share in total Data Centers

Ex-Hyperscalers Input

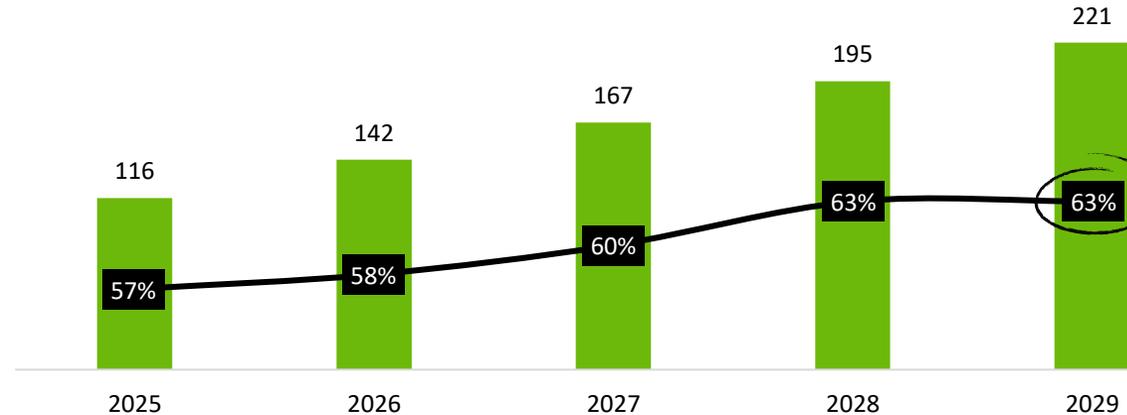
— % of Accelerators Capex — NVDA Market Share



IV. Still, the growth of sovereign and enterprise AI has a strong impact

Data Center Revenue (Ex-Hyperscalers)

■ NVDA DC Revenue (Ex-Hyperscalers) — % of NVDA DC Revenue

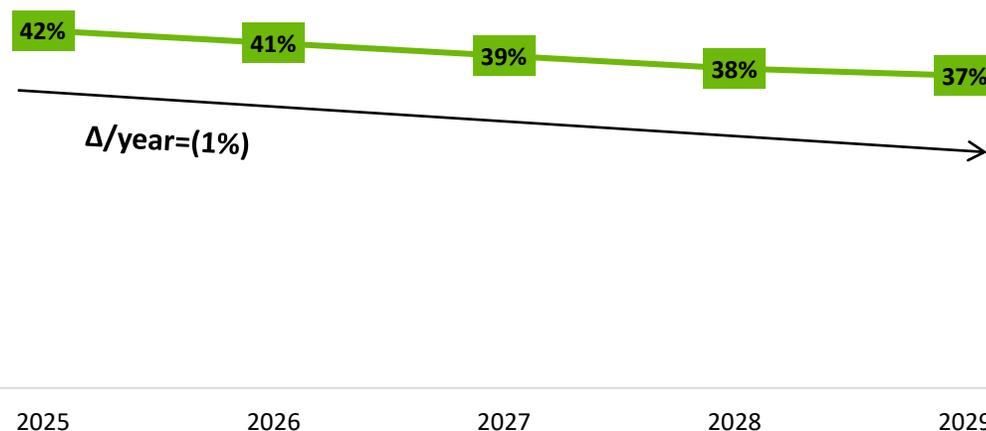


How Hyperscalers Share reflects in our Model?

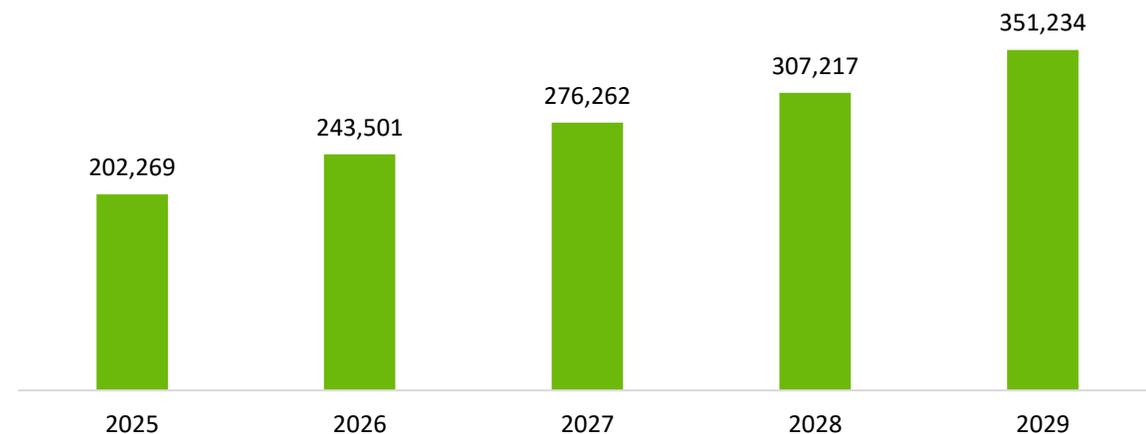
Declining hyperscaler's share supports investment case with solid IRR sensitivity

I. With the rise of sovereign AI and corporate AI, we conservatively reduce the hyperscalers' share of NVIDIA's revenue by 1 percentage point annually

Hyperscalers Share of NVIDIA Revenue (%)



NVIDIA Total Revenue (US\$ Million)



II. When performing a sensitivity analysis, we found that in the vast majority of cases we still achieve a satisfactory IRR, which provides us with a margin of safety

IRR Sensitivity Analysis

% of Hyperscalers	Market Share (%)						
	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%	95.0%
48.0%	6.3%	8.6%	10.9%	13.1%	15.2%	17.3%	19.3%
44.0%	9.0%	11.5%	13.9%	16.2%	18.4%	20.5%	22.6%
40.0%	12.2%	14.8%	17.3%	19.7%	22.0%	24.2%	26.4%
36.0%	15.9%	18.6%	21.2%	23.7%	26.1%	28.5%	30.7%
32.0%	20.2%	23.1%	25.8%	28.5%	31.0%	33.5%	35.9%
28.0%	25.4%	28.5%	31.4%	34.2%	36.9%	39.5%	42.0%
24.0%	31.8%	35.1%	38.2%	41.2%	44.1%	46.9%	49.6%

IRR exceeded 15% in 76% of the cases

Even if maintained at 40%, we would get a 20% IRR

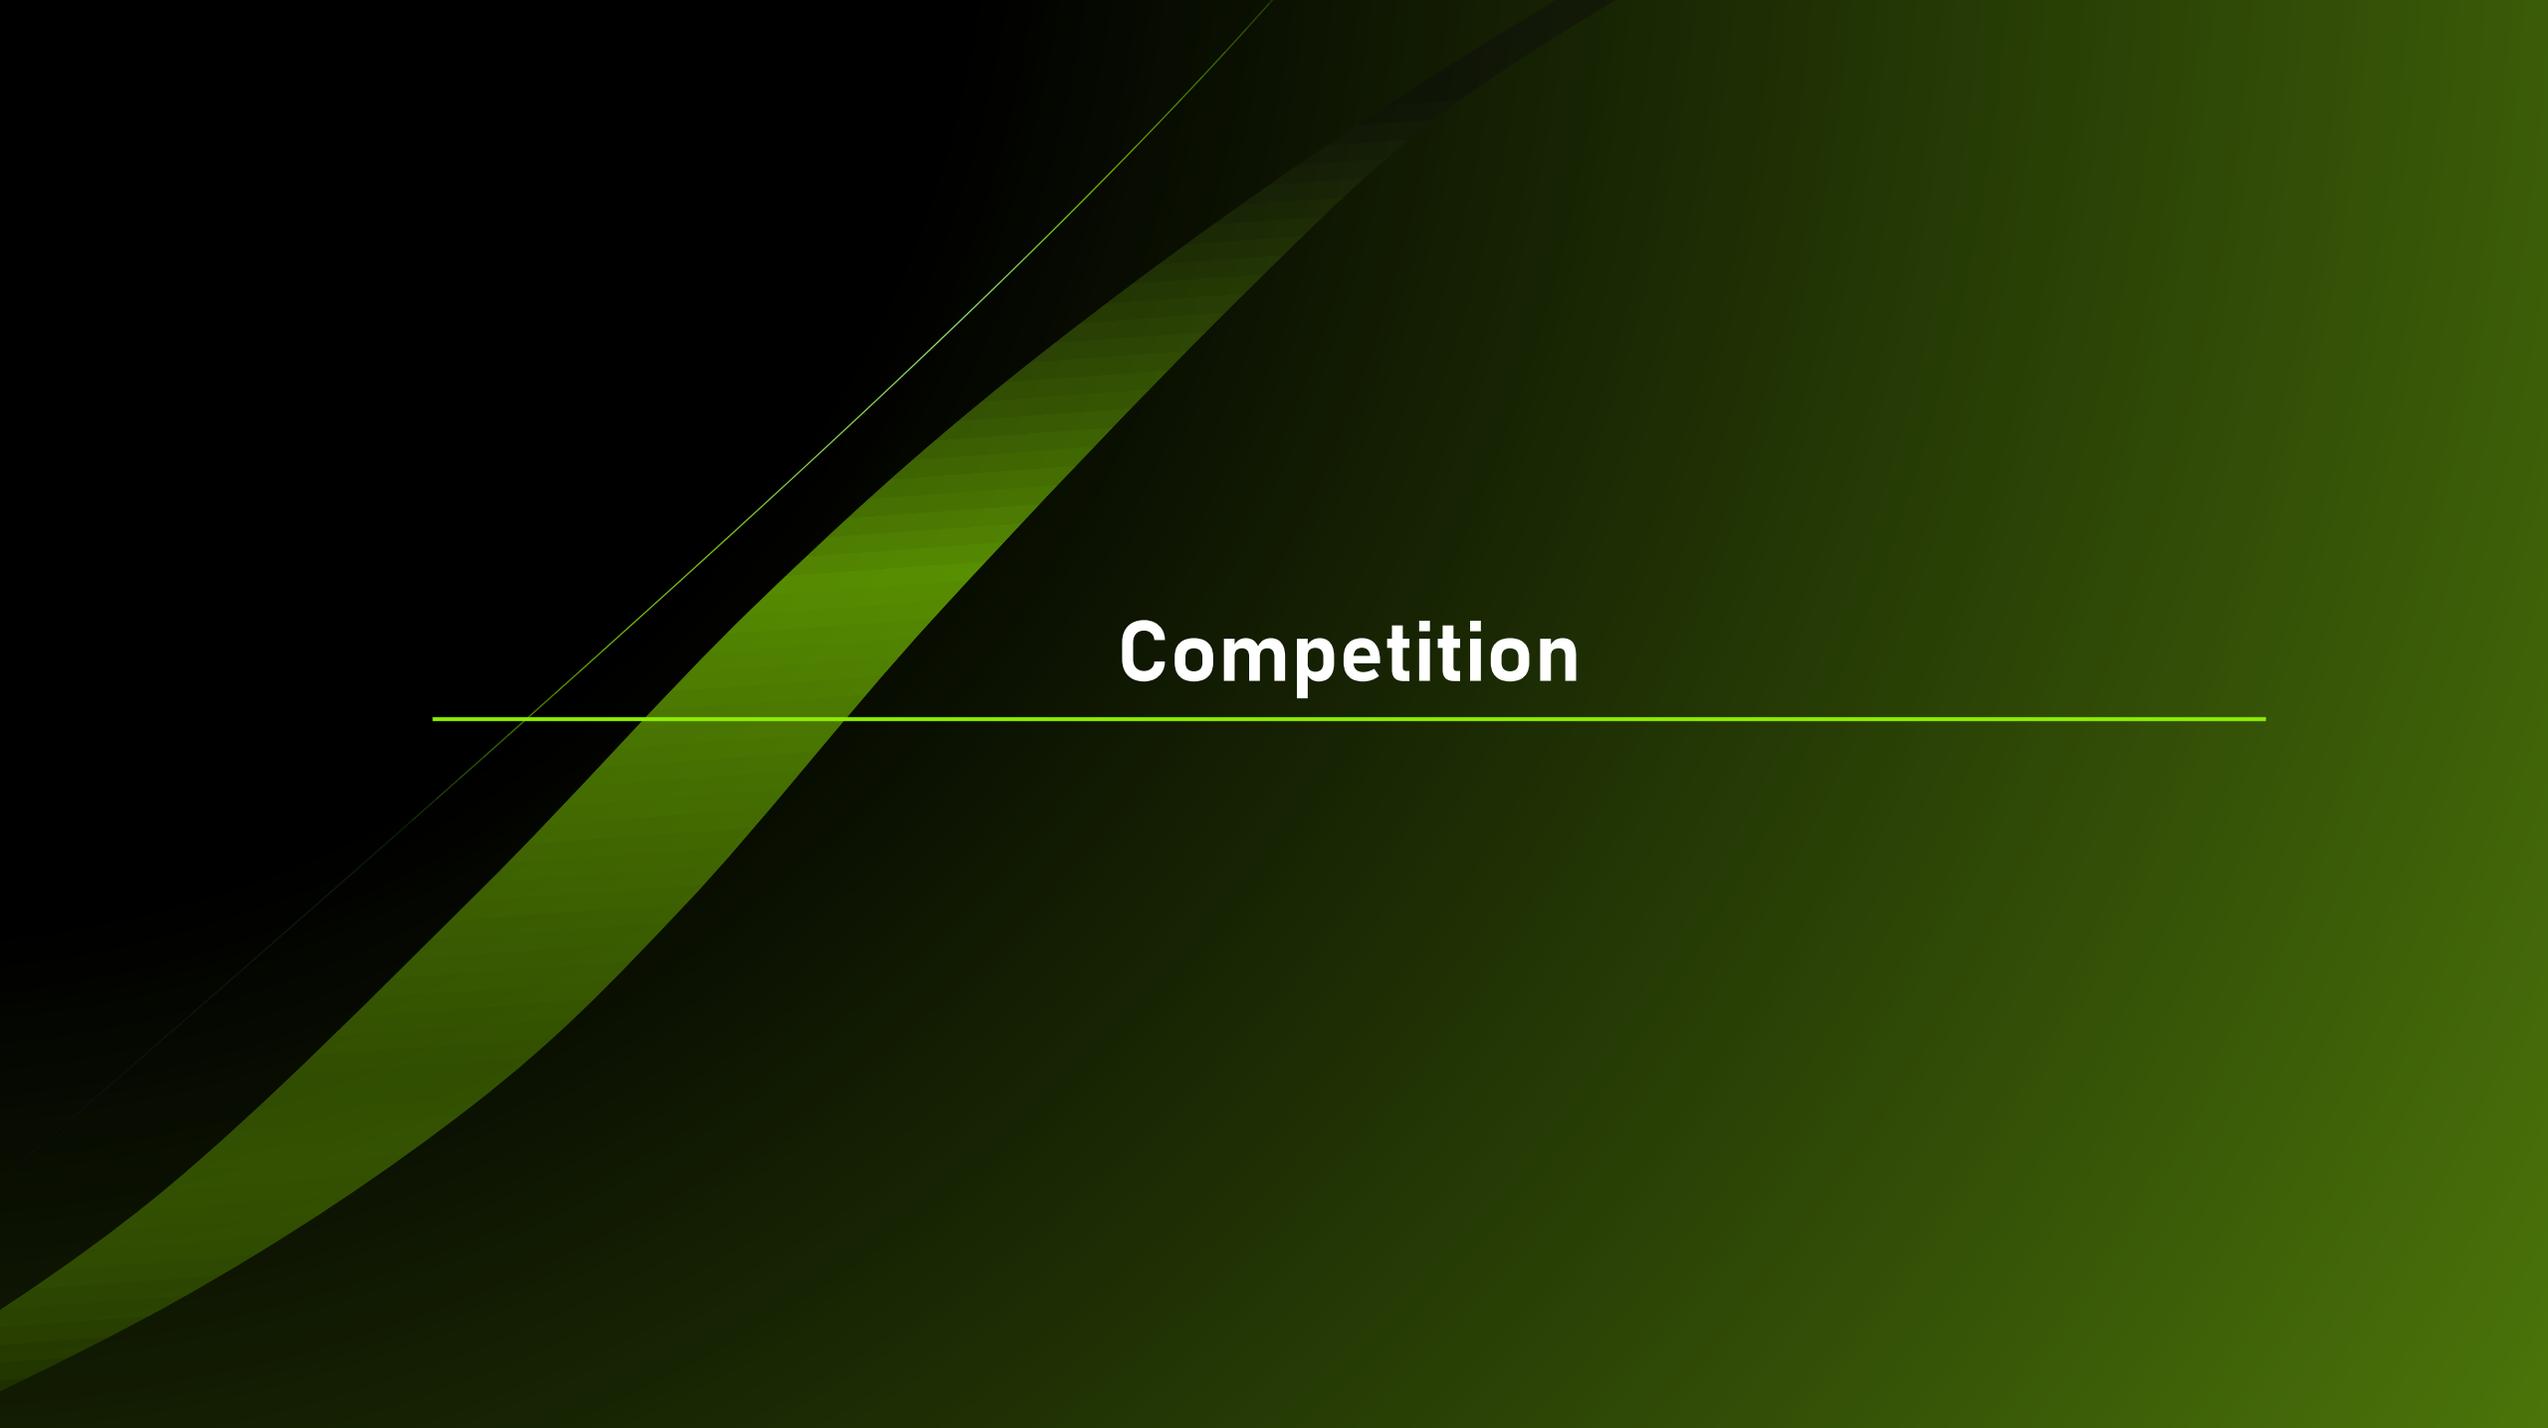
IRR Sensitivity Analysis

% of Hyperscalers	Marginal Capex (%)						
	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%	95.0%
48.0%	9.5%	10.8%	12.1%	13.4%	14.6%	15.8%	17.0%
44.0%	12.5%	13.8%	15.1%	16.5%	17.7%	19.0%	20.2%
40.0%	15.8%	17.2%	18.6%	20.0%	21.3%	22.6%	23.9%
36.0%	19.7%	21.2%	22.6%	24.1%	25.5%	26.8%	28.2%
32.0%	24.3%	25.8%	27.4%	28.9%	30.3%	31.8%	33.2%
28.0%	29.7%	31.4%	33.0%	34.6%	36.2%	37.7%	39.2%
24.0%	36.5%	38.3%	40.0%	41.7%	43.4%	45.0%	46.6%

IRR exceeded 15% in 77% of the cases

Even if maintained at 40%, we would get a 20% IRR





Competition

AMD at a Glance

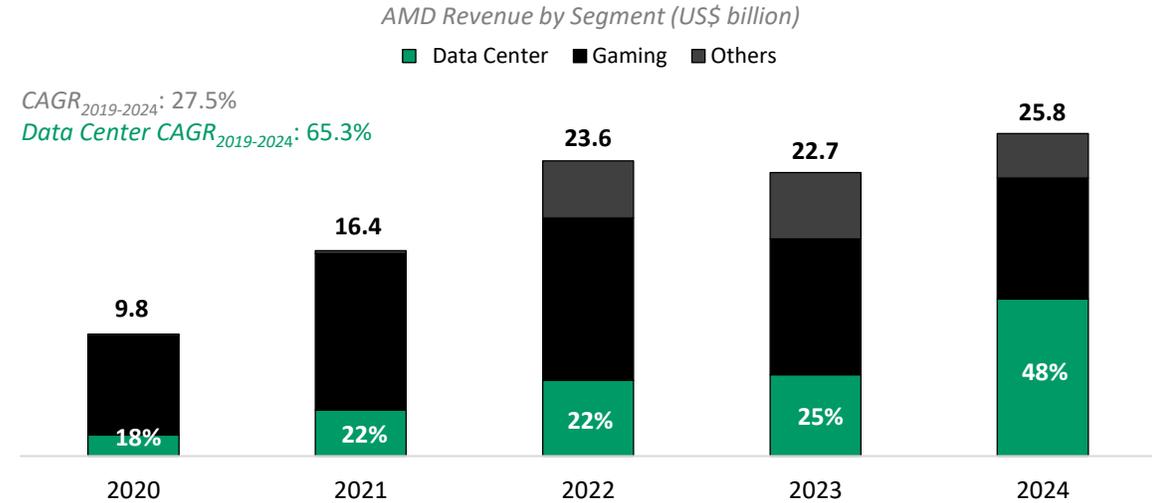
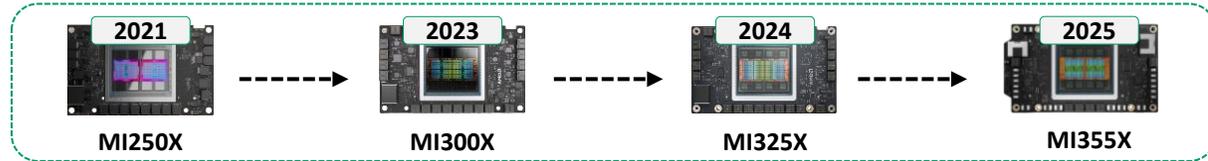
The underdog in the AI revolution is starting to show signs of activity

I. Originally a CPU player, AMD entered the Data Center GPU market late, but successfully established itself as the second-largest player



Continues to operate in the CPU market for gaming, but with an increasingly **strong focus on the data center** segment:

Data Center Roadmap:



II. Today, the company is attempting to replicate a similar strategy it used against Intel in the CPU market, aiming to capture some of the market share held by NVIDIA



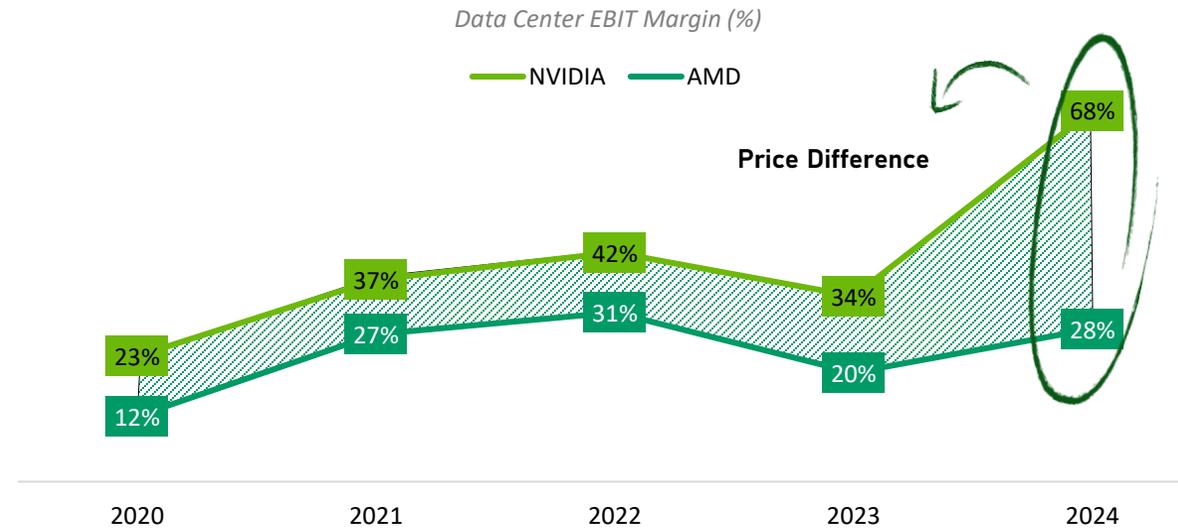
Dr. Lisa Su, AMD CEO

Jensen's cousin, AMD's CEO for over a decade, is now applying a **similar strategy against NVIDIA as she did with Intel**

TCO
Undercutting NVIDIA by pricing at approximately **30-40% lower**

ROCm
Made its ROCm software stack fully **open source**

UALink
Spearheading the UALink Consortium alongside companies like Intel and Broadcom to create an **open, vendor-neutral** alternative to NVIDIA's NVLink for GPU interconnects



ASICs at a Glance

As a way to reduce dependence on NVIDIA, ASICs have been gaining traction

I. ASIC chips can deliver lower costs due to their high application-specific design

What is an ASIC Chip?

Parallelism	Application	Architecture	Scalability
Fixed	Single-purpose	Hardware Only	Limited

Strength:



TCO

Lower TCO than NVIDIA GPUs due to their **simpler design and lower power consumption**. Their less complexity **enables reduced pricing** and operational costs

III. So far, Google's TPU stands as the only successful, first-mover alternative to GPUs

Case Study: Tensor Processor Unit

Google began its partnership with Broadcom in **2015** and has consistently launched new versions throughout **the past decade**

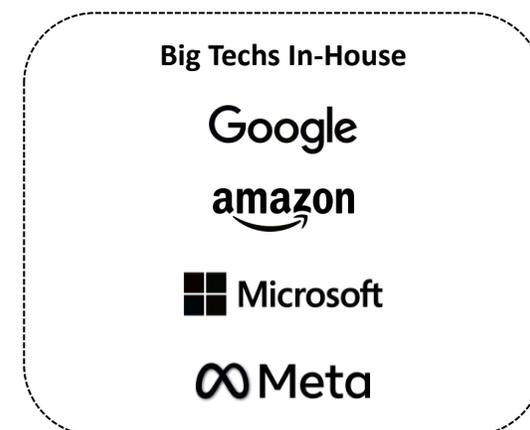
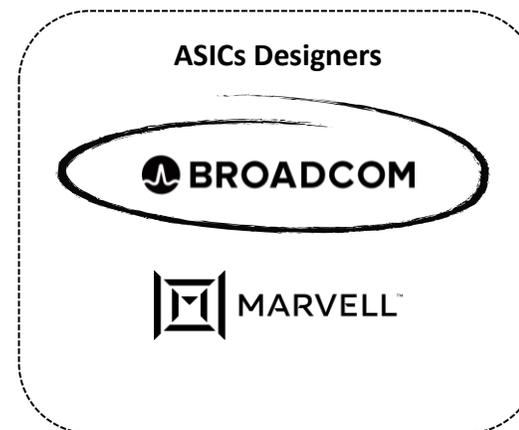
7,000\$ TPU x **35,000\$** GPU

Possible by the fact that TPUs cost approximately 3.7 times less to manufacture than GPUs



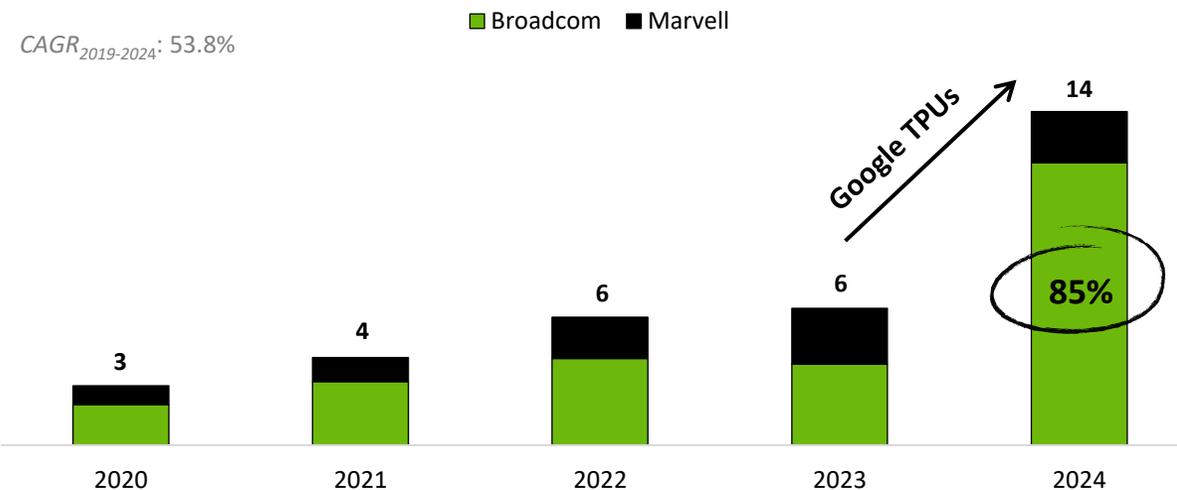
II. Big techs are partnering on chip design to strengthen their competitive edge

ASICs Companies



IV. As big tech projects grow, ASICs have been gaining traction

ASICs Designers Revenue (US\$ Billion)



Unconvincing Playbook

Despite ongoing efforts, it will be difficult to challenge the world's largest company in the short to medium term

I. Despite AMD's efforts to catch up to NVIDIA, its strategy has already shown weaknesses and is unlikely to significantly impact the market leader

Comparison of Resolved GitHub Issues: CUDA vs. ROCm



Why the Intel scenario is unlikely to happen again

Intel went through a staggering series of **execution missteps** between 2015 and 2021, with **repeated delays** and a **loss of strategic focus**



Insane Execution

+

Intel **failed to create meaningful switching costs**, so the winner was simply whoever delivered the better CPU (Same stack to all)



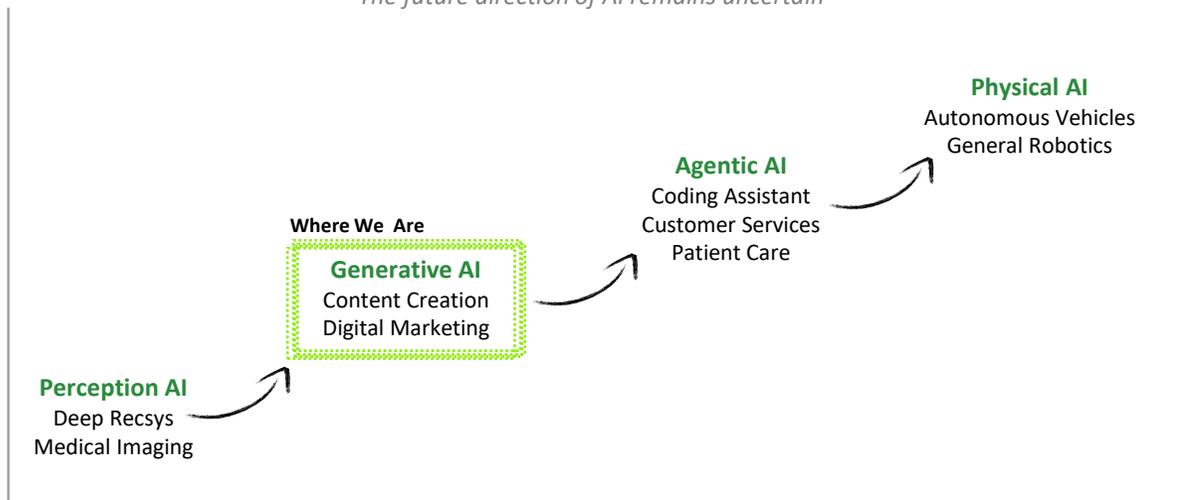
Full-Stack

UALink entered the race too far behind to realistically catch up:



II. In a world where the future of AI remains uncertain, the flexibility of GPUs proves essential — just as it has throughout their entire history

The future direction of AI remains uncertain



Cloud providers will not stick to a single architecture, as **their workloads constantly shift from one client to another**



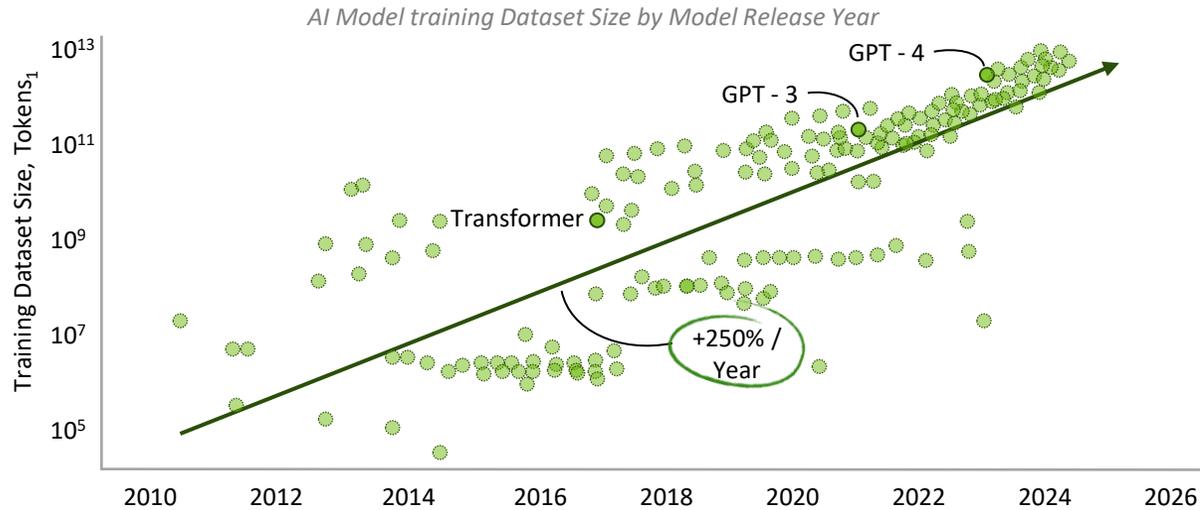
	Activity	Evolves	Ramp-up
GPU	Multi-purpose	5x in 3y	Plug-and-play
ASIC	One and only	No	Software Development



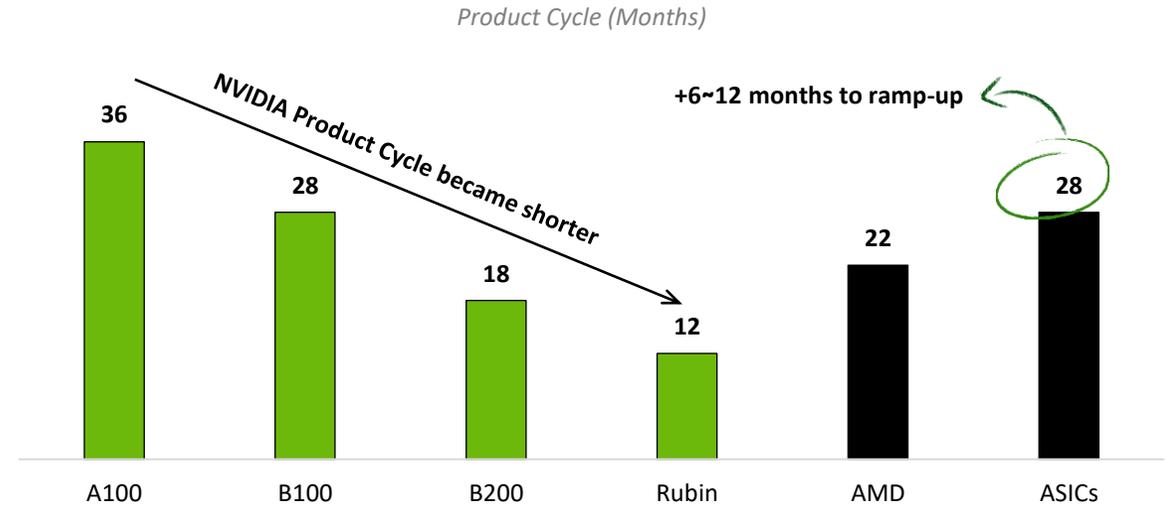
AI Rush Isn't Making It Any Easier

In a world where models evolve at lightning speed, everyone wants their solution delivered yesterday

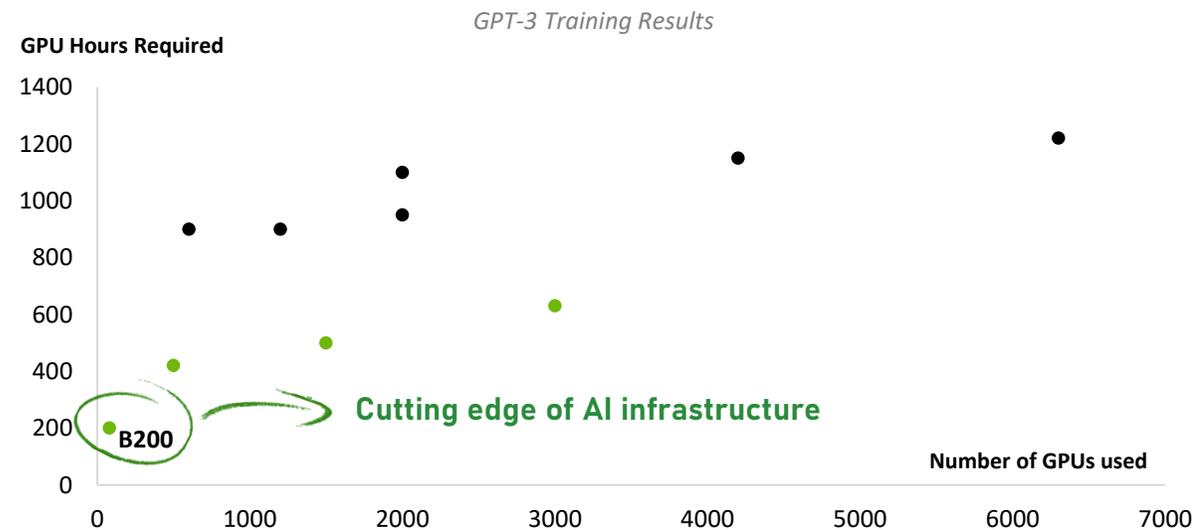
I. Models are advancing at an unimaginable pace



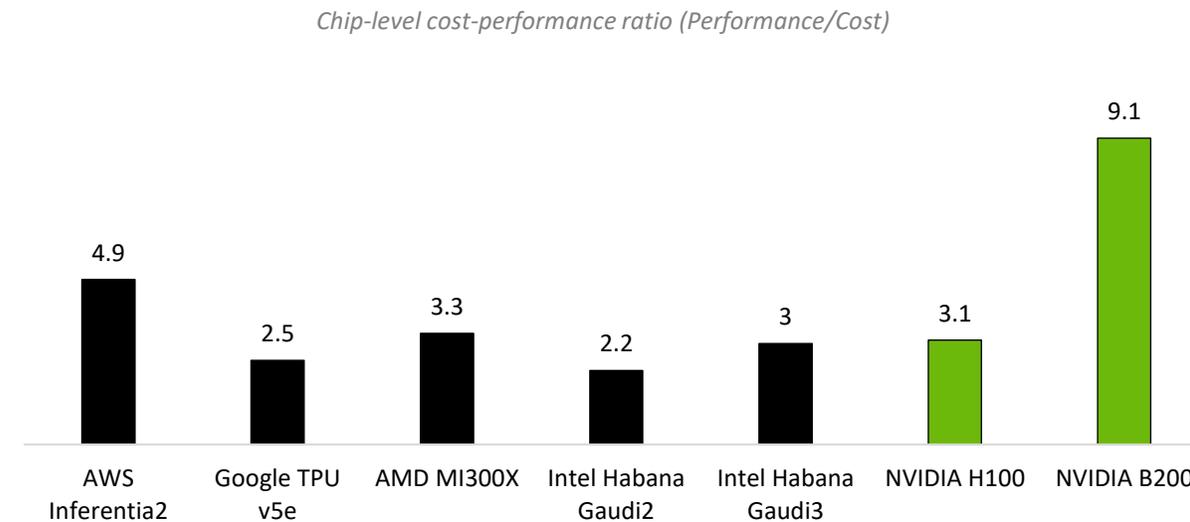
II. NVIDIA has been able to deliver new and improved products on very short timelines



III. Today, the company still delivers the best performance



IV. With every new product release, its cost-benefit ratio improves dramatically



How does this translate into our model?

Competitors are expected to serve more niche applications, while the lion's share of the market remains with NVIDIA

I. Even compared to the success case, GPUs are growing at an even faster pace

Even the best ASICs aren't growing as fast as GPUs

Morgan Stanley, Feb 25

“Specifically in the case of Google, we expect their purchases of NVIDIA to roughly double this year, while TPU will grow modestly. Some of that is Google's investment in enterprise cloud, but we are told that some of that is NVIDIA's very strong performance for LLM transformers even for some of the internal workloads”

II. Although unlikely to catch up, we expect them to find success in niche products

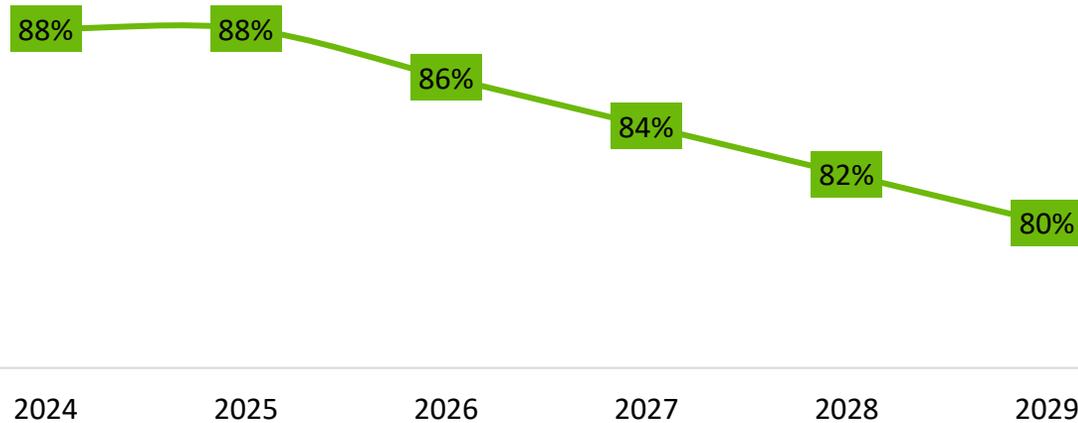
Despite that, we believe niche products will still capture some market share:

Matured Markets
Markets that mature over time can be overtaken by ASICs, just like Bitcoin mining was — especially when workloads become stable, predictable, and cost-sensitive

Corporate Clients
Not all GPU customers are hyperscalers or require ultra-high performance — AMD can take advantage of these markets by offering lower prices

III. Accordingly, our model reflects a gradual loss of market share over time, although we believe NVIDIA will remain highly dominant over the next five years

Senna's Market Share Projections (%)



Although we project a decline in share, we remain highly confident in NVIDIA's dominance over the coming years — **especially in the next five, which aligns with our planned exit**

Market Share (%)	Gross Margin (%)						
	60.0%	65.0%	70.0%	75.0%	80.0%	85.0%	90.0%
68.0%	6.8%	10.4%	13.8%	17.1%	20.2%	23.1%	25.9%
73.0%	9.0%	12.8%	16.3%	19.6%	22.8%	25.8%	28.7%
78.0%	11.3%	15.1%	18.7%	22.1%	25.3%	28.4%	31.3%
83.0%	13.4%	17.3%	20.9%	24.4%	27.7%	30.8%	33.8%
88.0%	15.4%	19.4%	23.1%	26.7%	30.0%	33.2%	36.3%
90.0%	16.2%	20.2%	24.0%	27.6%	30.9%	34.2%	37.3%
92.0%	17.0%	21.0%	24.8%	28.4%	31.8%	35.1%	38.2%

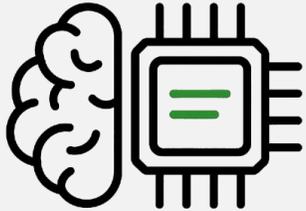
Even when stress-testing for extreme scenarios, we still find that **in most cases the investment makes sense with an exit in 2026**



What if better software than CUDA were introduced?

Can CUDA be change? We don't think so

1. Hardware Superiority



NVIDIA isn't just leading, it's setting the pace. Any challenger must deliver a clear performance leap; otherwise, users will simply keep using NVIDIA's proven platform. In this space, matching isn't enough, you have to outperform, and that bar keeps rising

2. Extreme Capital Intensity + Long Development Cycles



Building competitive hardware demands billions in R&D, manufacturing, and integration. But money alone isn't enough, development cycles take years, and by the time a rival product hits the market, NVIDIA will likely have launched its next generation. Time is a brutal adversary.

3. Transparency vs. Proprietary Innovation



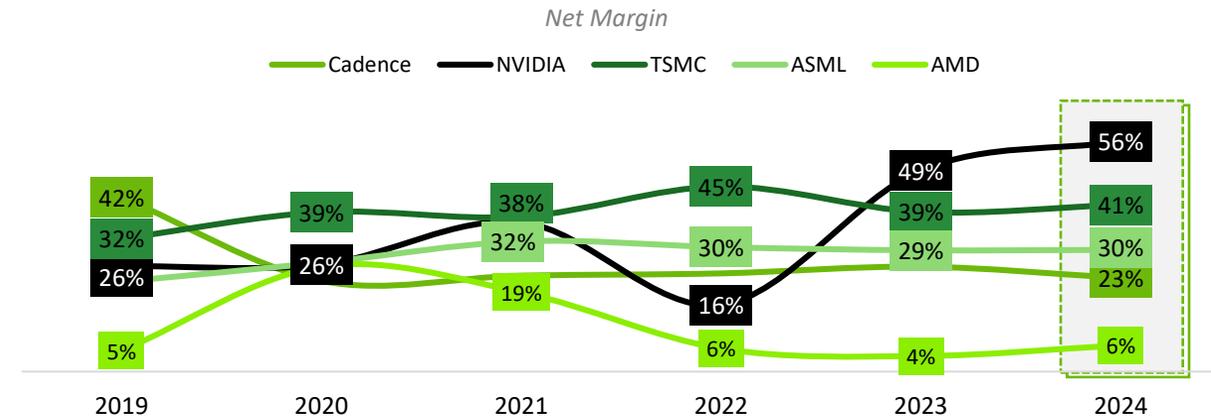
Challengers often go open-source to attract developers and build credibility, but this also exposes their roadmap and timing. Meanwhile, NVIDIA moves silently and strategically, supported by a proprietary ecosystem of innovative software libraries like CUDA, TensorRT, and cuDNN. Even if someone catches up, NVIDIA is already launching what comes next



Is Nvidia the Best Investment in Semiconductors? We Think so

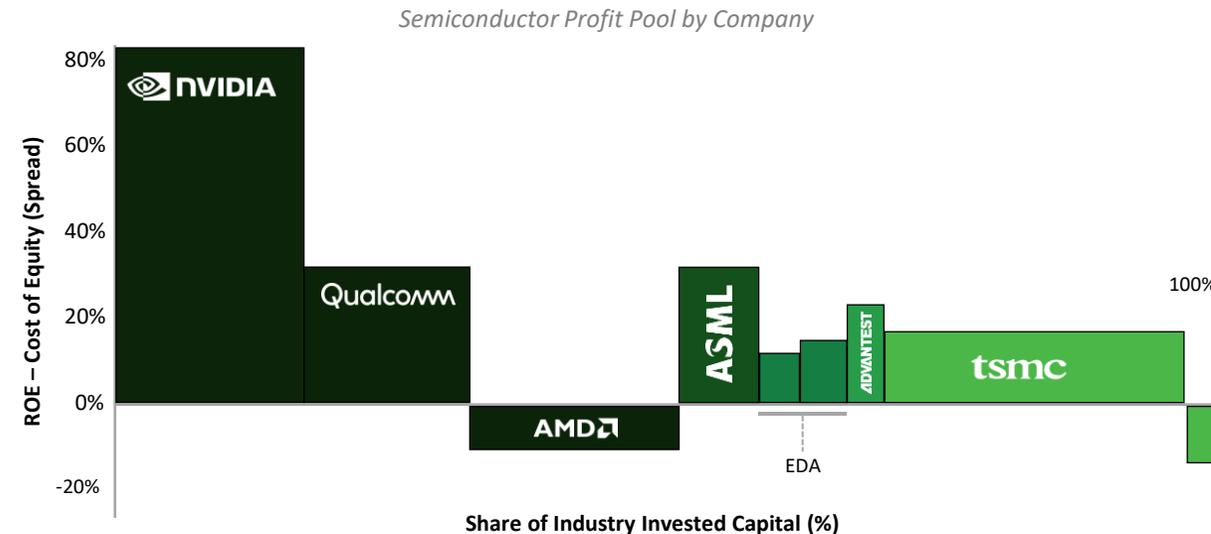
Why Nvidia Stands Out

Nvidia doubled its margin in 5 years while peers remained flat

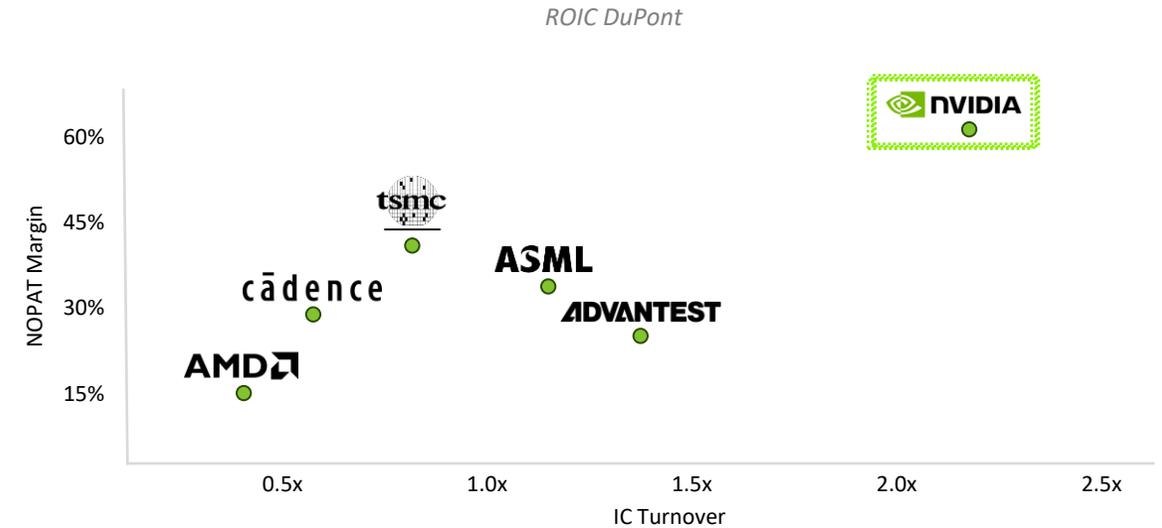


By repositioning its core business around Data Centers, Nvidia not only expanded its margins but also built a more defensible and scalable foundation, and given the sustained demand for AI infrastructure, this margin strength is likely to persist

Value creation led by NVIDIA



Nvidia exposes rivals' limits with unmatched scale, efficiency, and returns



Why we prefer NVIDIA than the others?

Nvidia scales with agility and high profitability, while TSMC relies on capital-intensive expansion, facing operational rigidity and diminishing returns

While TSMC excels in manufacturing efficiency, Nvidia captures greater value by delivering integrated solutions, resulting in consistently higher operating margins

Nvidia's model is more robust, combining control over design and software to act as a pricing maker, while competitors remain constrained by price pressure and external dependencies

Nvidia sits at the most strategic point of the semiconductor value chain, where architecture, software, and integration converge, enabling it to shape outcomes and capture outsized value



ROIC Comparison





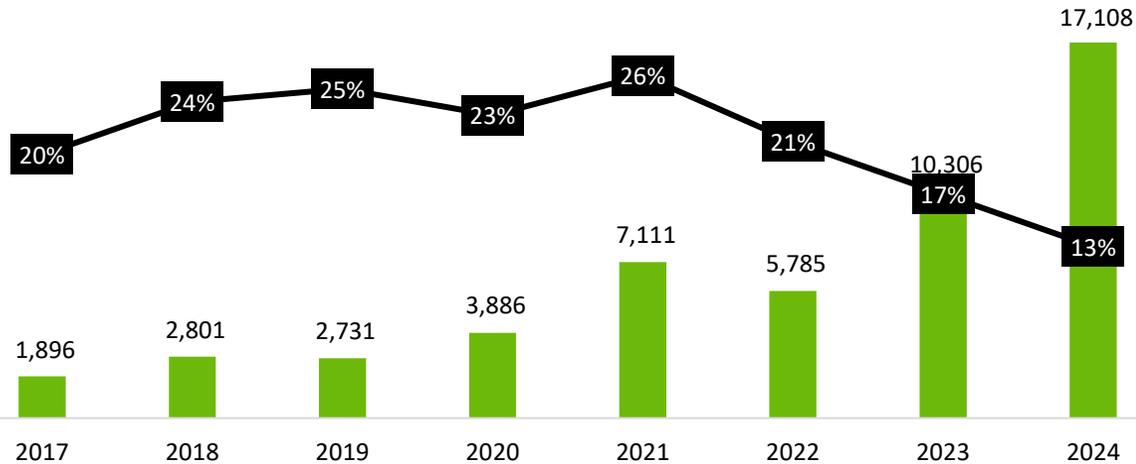
Geopolitical

China as an Opportunity

Once considered a risk, China now presents itself as an opportunity in the company's price action

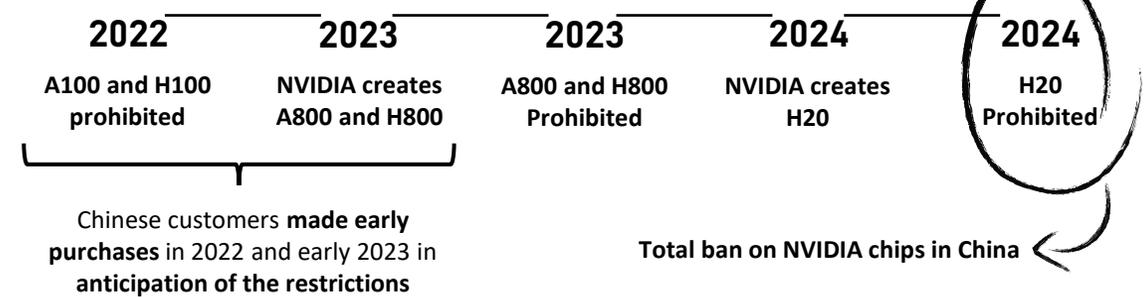
I. Historically, revenue from China has accounted for one-fifth of NVIDIA's total revenue

NVIDIA China Revenue (US\$)



II. The sanctions have always been followed by NVIDIA's swift countermeasures

Sanctions Roadmap



With the DeepSeek Day on January 27, 2025, the U.S. government became much more cautious and increased sanctions

III. After multiple sanctions, NVIDIA currently finds itself in a position where it exports no products to Chinese soil. As such, any positive developments now represent an opportunity

The market no longer expects any revenue from China

Morgan Stanley
BANK OF AMERICA
JPMorganChase

All major sell-side analysts covering NVIDIA have now reduced their **China revenue estimates to zero**

Nevertheless, any positive developments would be beneficial for the company

World^ Business^ Markets^ Legal^ Commentary^ REUTERS

EXCLUSIVE:

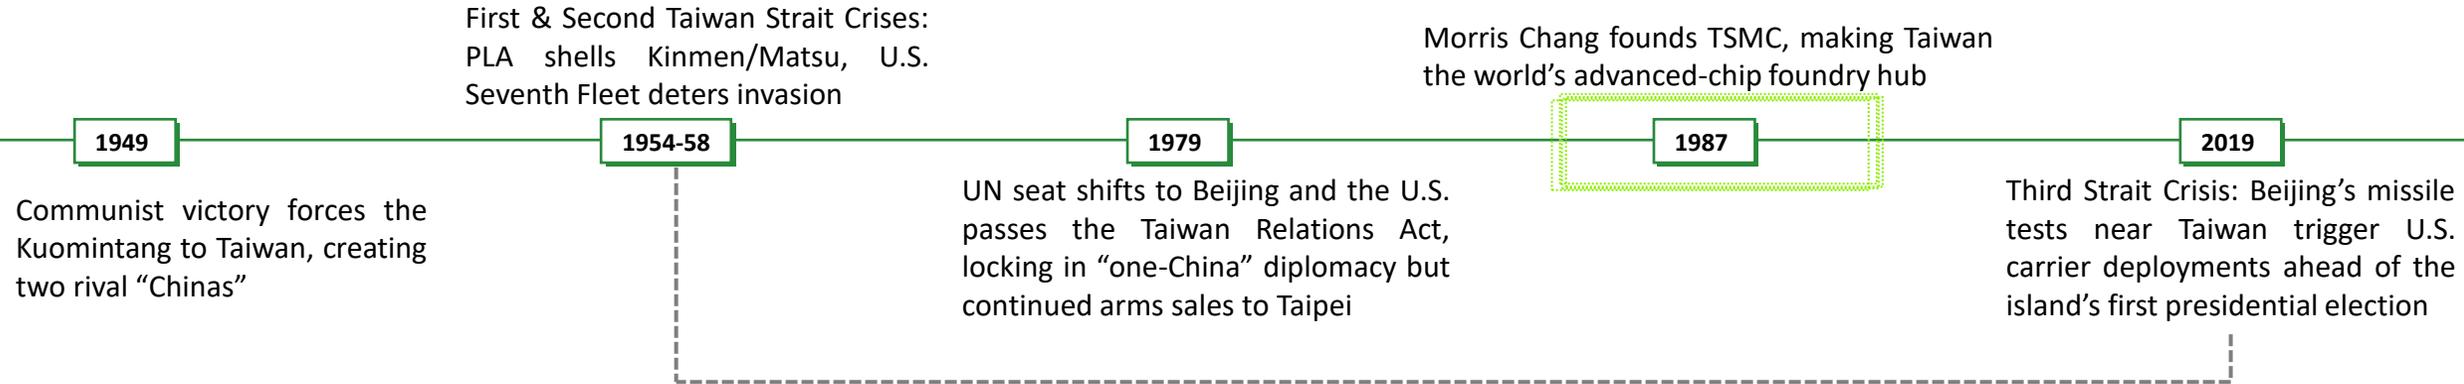
Nvidia to launch cheaper **Blackwell AI chip for China** after US export curbs, sources say

Now, any new developments within China are seen as opportunities for NVIDIA, no longer as risks



China x Taiwan

From shelling to silicon: a flashpoint still inches from war



Moments of near-war tension

II. What are the current tensions like?



<https://www.bloomberg.com/news/articles/2025-06-10/china-sends-two-aircraft-carriers-into-pacific-ocean-in-a-first>

"China Sends Two Aircraft Carriers Into Pacific Ocean in a First, marking an unprecedented show of naval reach and raising security concerns for Taipei, Tokyo and Washington"

Bloomberg



<https://braziljournal.com/guerra-de-trump-contra-a-china-abate-a-nvidia/>

"Beijing called its latest maneuvers "routine training" and urged Taipei to "stop hyping," while Taiwanese officials noted no live-fire activity and no further escalation despite a spike in PLA aircraft and ships"

**TAIPEI
TIMES**

III. What we can expect for the future? Just status quo

We believe the island of Taiwan will remain as it is today, and tensions will gradually subside



1. If they were to fail, a setback of this magnitude, with extremely high economic and military costs, could destabilize the Communist Party itself, a political risk that deters Beijing
2. CSIS wargames show China comes up short in most invasion scenarios
3. Crossing the strait and holding Taiwan would be extremely difficult, especially with U.S. and Japanese forces stepping in



Why Taiwan is so Important to NVIDIA? TSMC

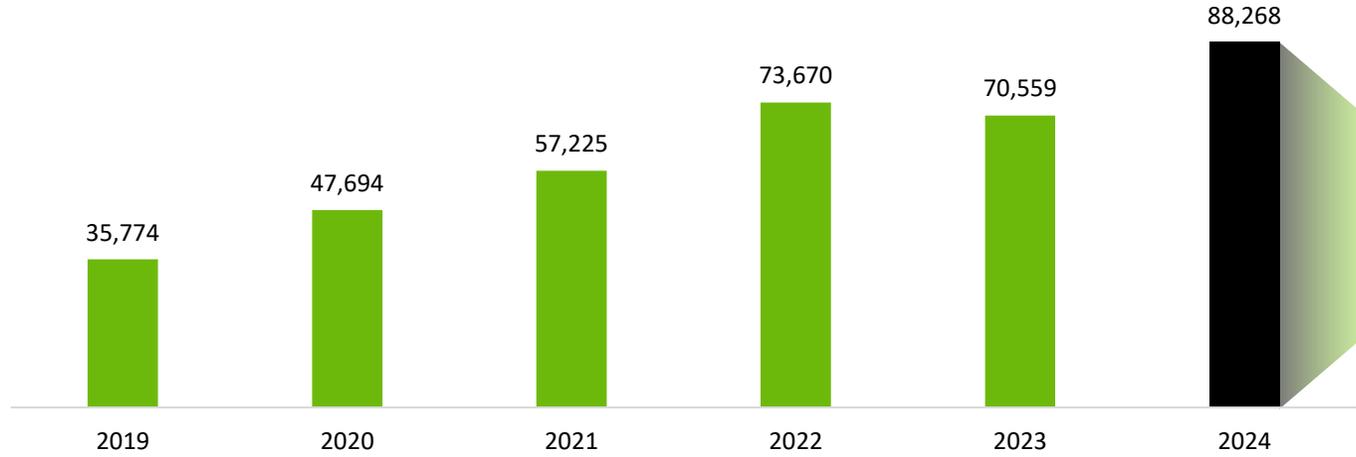
Because the world's most advanced chips still come from 100 miles off China's coast

I. TSMC's revenue boom is powered by advanced nodes, with 5nm and 3nm alone driving nearly half of all sales in 2024

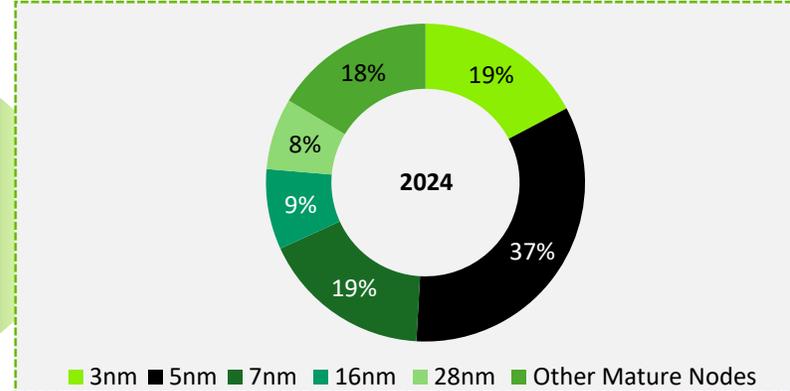
TSMC

Is the world's top chipmaker, producing advanced semiconductors (5nm, 3nm) for companies like Nvidia and Apple. Based in Taiwan, it supplies over 90% of the world's cutting-edge chips, making it vital to AI and global tech

TSMC Revenue (US\$ Billion)



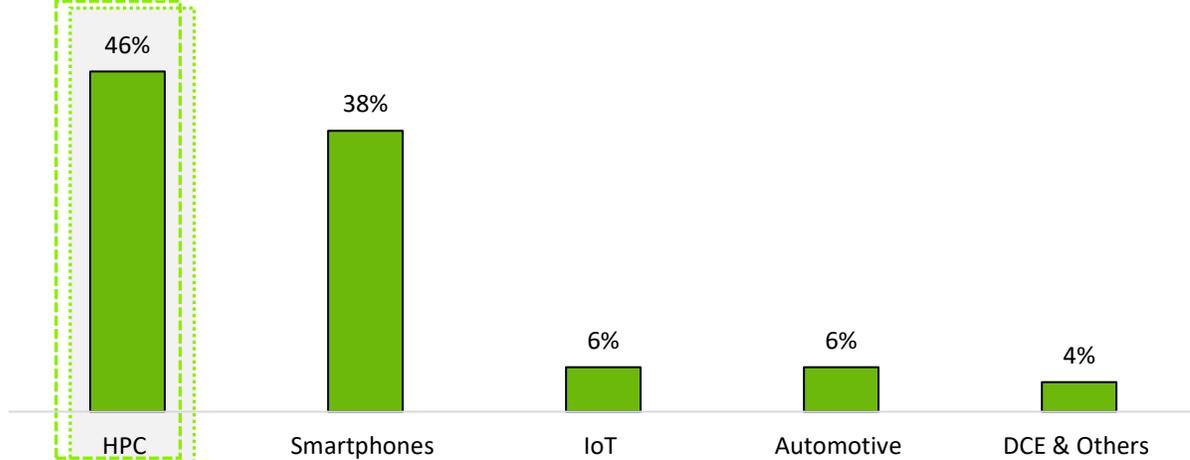
TSMC — wafer revenue by process node



II. TSMC's reliance on HPC makes Nvidia not just a client — but a vulnerability

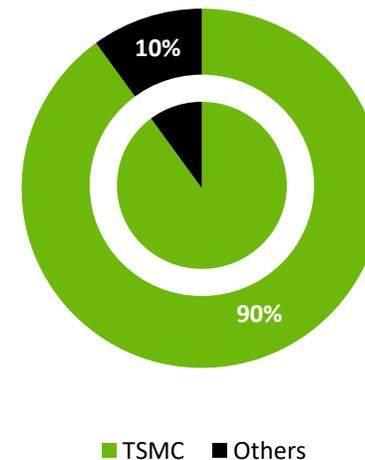
TSMC Revenue by Platform

NVIDIA Segment



III. A dominant monopoly in advanced semiconductor foundry

TSMC Market-Share in <7nm chips



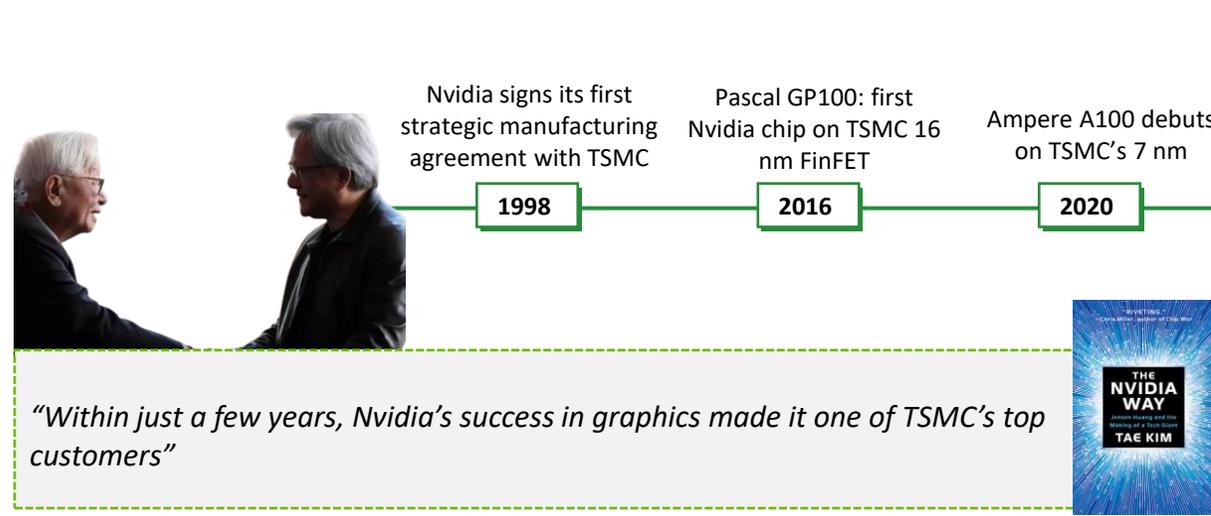
With 90% of <7nm chips coming from TSMC, the world has effectively outsourced cutting-edge computing to a single, vulnerable island



A Powerful Alliance Between Chip Design and Manufacturing Leaders

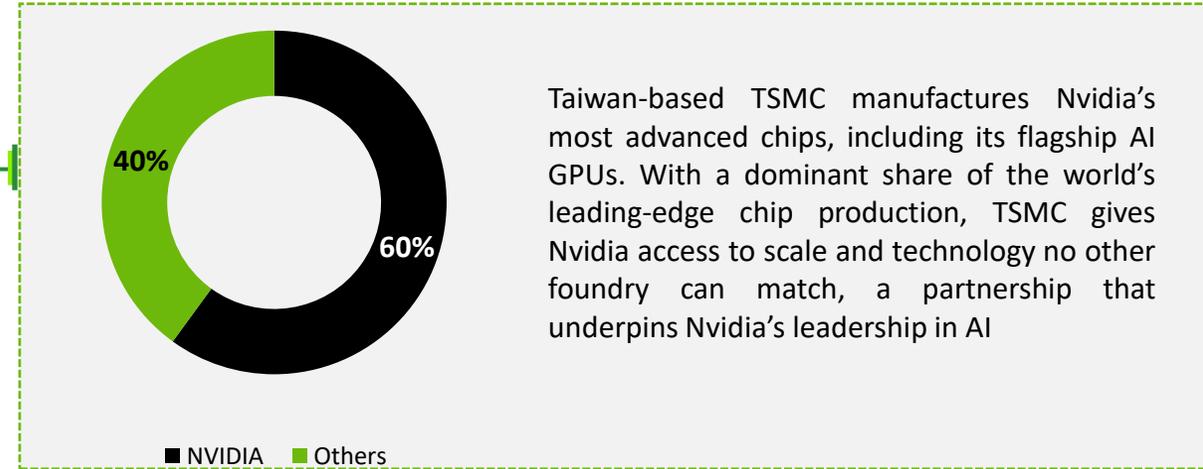
Nvidia and TSMC define the frontier of semiconductors

I. Nvidia's rise in AI is inseparable from its decades-long dependence on TSMC's most advanced nodes



"Within just a few years, Nvidia's success in graphics made it one of TSMC's top customers"

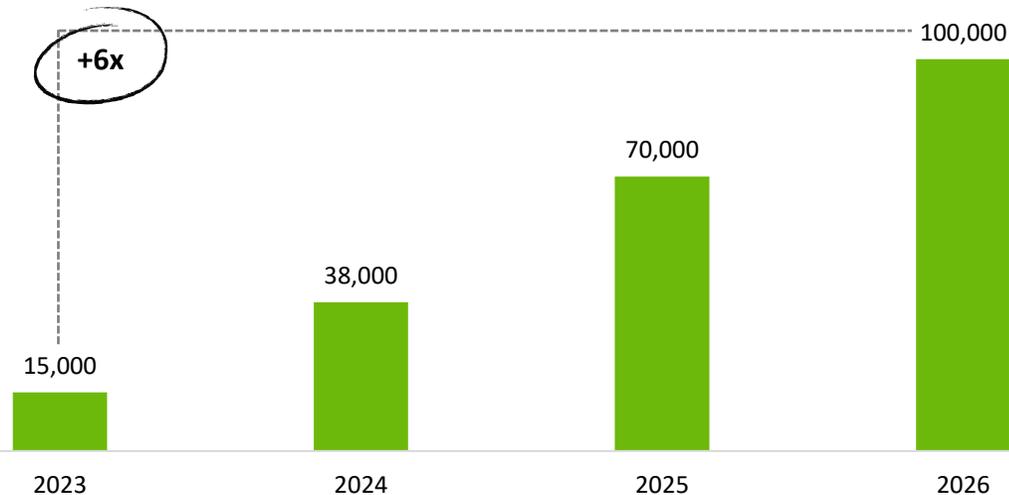
TSMC CoWoS Packaging Capacity (2025E)



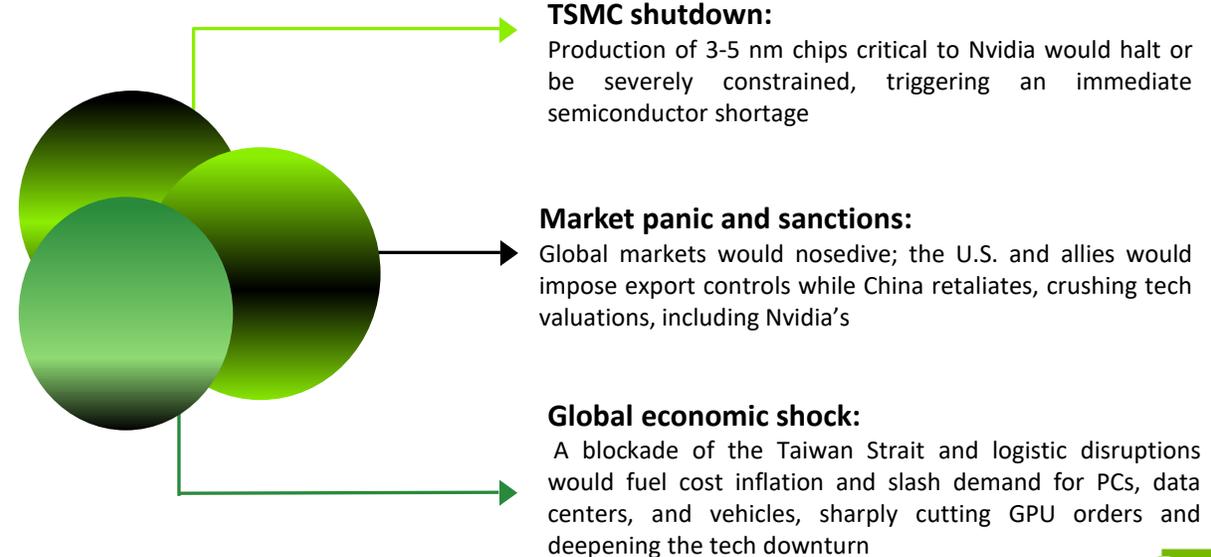
Taiwan-based TSMC manufactures Nvidia's most advanced chips, including its flagship AI GPUs. With a dominant share of the world's leading-edge chip production, TSMC gives Nvidia access to scale and technology no other foundry can match, a partnership that underpins Nvidia's leadership in AI

II. TSMC is scaling CoWoS fast to meet Nvidia's AI demand

CoWoS capacity (wafer-per-month, TSMC)



III. What would happen if China invaded Taiwan



Trade War: Direct Impact on Nvidia

How the tariff war between USA and China can affect Nvidia?



Donald Trump at "libertation day"

On 2 April 2025 President Trump signed Executive Order imposing a blanket 10 % tariff on almost all U.S. imports from 184 countries starting 5 April, while creating country-by-country "reciprocal" rates, many well above the baseline, to begin on 9 April

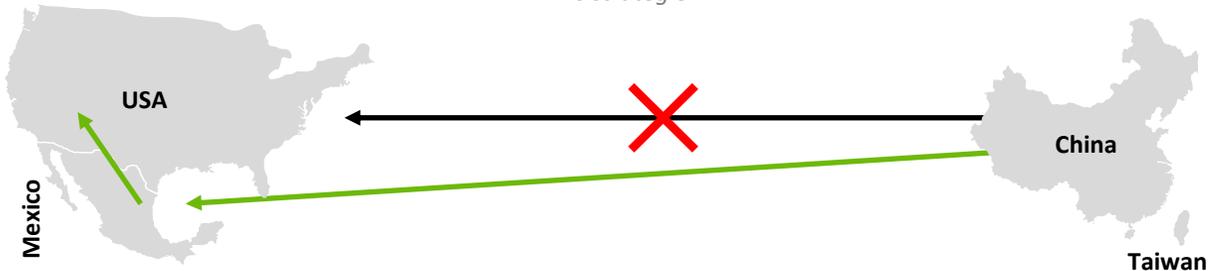
Pressure China to curb unfair practices, IP theft, subsidies, and forced tech transfers

Shield strategic U.S. industries such as steel, semiconductors, and clean energy under the banner of national security

Gain negotiating leverage, spotlighting America's large trade deficit with China

How this can affect Nvidia?

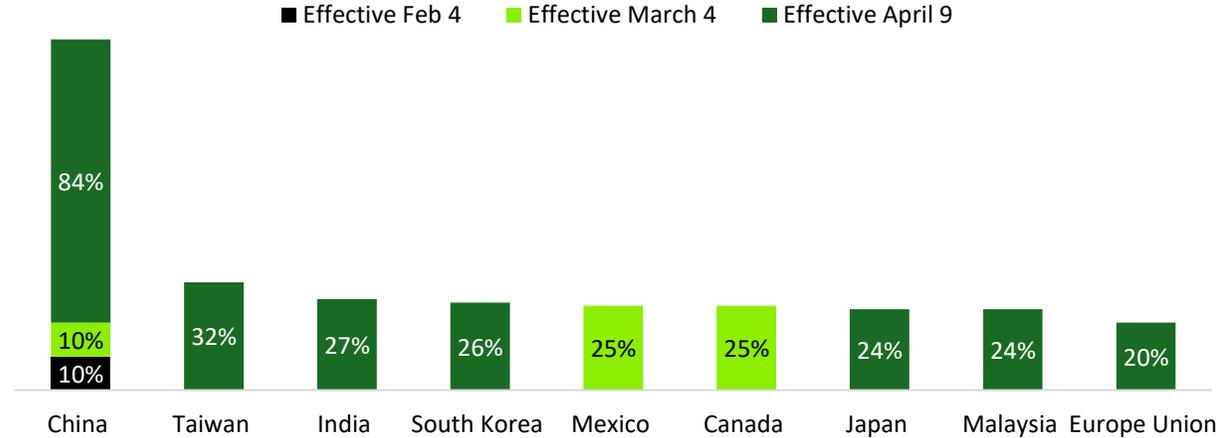
NVIDIA's strategie



Nvidia ships its finished chips, made in Taiwan and supplemented by China-sourced components, to **Foxconn plants in Mexico**, where PCB, memory, and firmware are integrated; this "substantial transformation" meets **USMCA rules of origin**, allowing the systems to **enter the United States tariff-free**

US tariff on China soars to 84 %, versus ~25 % on other partners

US import tariff rate on key trading partners



Anyway... NVIDIA has pricing power

Case Study about 2018's tariffs

In 2018, a **25 % tariff on China-made GPUs** raised Nvidia's costs. Nevertheless, it lifted card prices by about 10–15 %, **kept demand strong**, and closed the fiscal year with a gross margin above 60 %, underscoring its **robust pricing power**

Jensen Huang, Nvidia's Founder & CEO

"The impact of tariffs won't be meaningful"



The image features a dark green background with a lighter green curved band that starts from the bottom left and curves upwards towards the top right. A thin, light green horizontal line spans across the middle of the image. The text 'DCF' is centered in white, bold, sans-serif font, positioned above the horizontal line and within the green band.

DCF

Why not DCF?

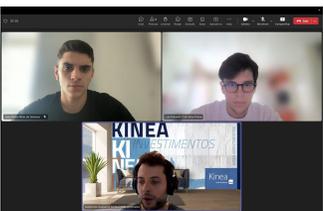
Why we believe it doesn't make sense to value Nvidia using a DCF

We had a call with some analysts, and a part of them said that running a DCF for Nvidia is just a theoretical exercise



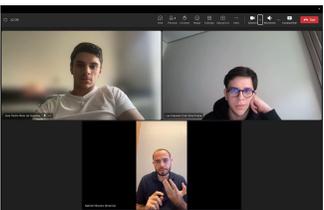
João Pedro Freitas – Mainú Capital

“With so much volatility and limited visibility into Nvidia’s future revenue streams, a DCF becomes more of a storytelling tool than a dependable valuation method. You’re basically modeling uncertainty on top of uncertainty”



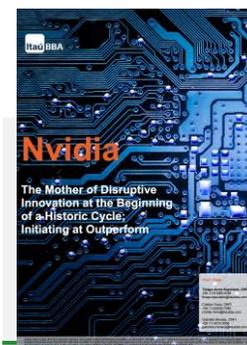
Guilherme Amaral – Kinea

“At Kinea, we don’t use DCF for tech. It just doesn’t make sense given the volatility and low visibility on long-term fundamentals. Multiples give us a cleaner read on what’s priced in”



Gabriel Oliveira – Verde Asset

“At Verde, we do run DCFs, but mostly to understand value boundaries and test assumptions. In the end, what really drives allocation is knowing what’s priced in today, and that comes much more from multiples and the narrative behind them”



Itaú Report about Nvidia

*“We initiate with an outperform rating, as we believe we are in the early innings of a strong cycle. Although **US investors typically don’t use a DCF**, we decided to go for it to value the company as we consider a long road for growth. Our TP is of USD 500”*

So, why we prefer use multiples...

The perpetuity fails to reflect the company’s real long-term growth potential, especially in fast-evolving sectors like semiconductors

Small changes in key assumptions can drastically alter the valuation, making the model unstable

The terminal value dominates the output, often accounting for the majority of the DCF, which reduces its reliability

In our model **70.5%** of Fair Value is **on perpetuity**



FCFF Valuation

Our DCF valuation for FCFF

\$mn	Current	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	Perpetuidade
Period		0,75	1,75	2,75	3,75	4,75	5,75	6,75	7,75	8,75	9,75	
Valuation FCFF												
EBIT		139.199	167.395	189.831	212.299	243.656	280.726	323.311	372.225	428.406	481.382	
(+) D&A		2.240	2.180	2.284	2.391	2.622	2.964	3.367	3.838	4.384	5.015	
(-) Taxes		(23.194)	(27.856)	(31.641)	(35.480)	(40.810)	(47.076)	(54.276)	(62.546)	(72.046)	(81.052)	
(+/-) Delta Working Capital		(6.841)	(6.822)	(5.428)	(5.436)	(7.586)	(8.707)	(9.993)	(11.469)	(13.163)	(7.125)	
(-) Capex		(2.775)	(3.337)	(3.784)	(4.232)	(4.857)	(5.574)	(6.398)	(7.343)	(8.427)	(9.672)	
FCFF		108.629	131.560	151.263	169.542	193.025	222.333	256.012	294.706	339.154	398.221	7.707.598
Present Value FCFF		100.843	110.600	115.159	116.890	120.516	125.709	131.086	136.653	142.416	151.433	3.077.541
Enterprise Value	4.328.847	Actual Price		140	Perp Growth		5,0%					
Net Debt	(34.747)	Shares Outstanding		24.400	WACC		10,4%					
Fair Value FCFF	4.363.594											
Fair Price FCFF	179											
Upside FCFF		27,7%										
		Current Market Cap 3.415.756,0										

Upside	Growth						
	3,5%	4,0%	4,5%	5%	5,5%	6,0%	6,5%
8,9%	40,6%	51,8%	65,6%	82,9%	105,2%	135,2%	177,5%
9,4%	27,2%	36,2%	47,0%	60,3%	77,0%	98,5%	127,4%
9,9%	15,9%	23,2%	31,9%	42,3%	55,2%	71,3%	92,1%
10,4%	6,2%	12,3%	19,4%	27,7%	37,8%	50,2%	65,8%
10,9%	-2,1%	3,0%	8,8%	15,7%	23,8%	33,5%	45,5%
11,4%	-9,3%	-5,0%	-0,1%	5,5%	12,1%	20,0%	29,4%
11,9%	-15,6%	-11,9%	-7,8%	-3,1%	2,4%	8,7%	16,3%

Upside	Gross Margin						
	60,0%	65,0%	70,0%	75,0%	80,0%	85,0%	90,0%
7,5%	7,5%	19,0%	30,5%	41,9%	53,4%	64,9%	76,4%
10,5%	3,9%	15,0%	26,1%	37,2%	48,3%	59,4%	70,5%
13,5%	0,3%	11,0%	21,7%	32,5%	43,2%	53,9%	64,7%
16,5%	-3,3%	7,0%	17,4%	27,7%	38,1%	48,4%	58,8%
19,5%	-6,9%	3,1%	13,1%	23,0%	33,0%	43,0%	53,0%
22,5%	-10,5%	-0,9%	8,7%	18,3%	27,9%	37,5%	47,1%
25,5%	-14,0%	-4,8%	4,4%	13,6%	22,8%	32,1%	41,3%



FCFE Valuation

Our DCF valuation for FCFE

\$mn	Current	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	Perpetuidade
Valuation FCFE												
FCFF		108.629	131.560	151.263	169.542	193.025	222.333	256.012	294.706	339.154	398.221	7.707.598
(+/-) Δ Debt		-	-	-	-	-	-	-	-	-	-	-
(-) Interest		(489)	(489)	(489)	(489)	(489)	(489)	(489)	(489)	(489)	(489)	(489)
FCFE		108.140	131.071	150.774	169.053	192.536	221.844	255.523	294.217	338.665	397.732	7.707.109
Present Value FCFE		100.379	110.165	114.747	116.497	120.139	125.342	130.725	136.293	142.055	151.061	2.927.218
Fair Value FCFE	4.174.623	Ke 10,4%										
Fair Price FCFE	171											
Upside FCFE	22,22%											

Upside	Growth						
	3,5%	4,0%	4,5%	5,0%	5,5%	6,0%	6,5%
8,9%	15,1%	21,3%	28,7%	37,4%	47,8%	60,6%	76,7%
9,4%	10,7%	16,8%	23,8%	32,1%	42,1%	54,3%	69,7%
9,9%	6,6%	12,4%	19,1%	27,0%	36,6%	48,3%	63,0%
10,4%	2,7%	8,2%	14,6%	22,2%	31,4%	42,6%	56,7%
10,9%	-1,1%	4,2%	10,3%	17,6%	26,4%	37,1%	50,6%
11,4%	-4,7%	0,4%	6,3%	13,2%	21,6%	31,9%	44,8%
11,9%	-8,1%	-3,3%	2,4%	9,0%	17,1%	26,9%	39,2%

Upside	Gross Margin						
	60,0%	65,0%	70,0%	75,0%	80,0%	85,0%	90,0%
7,5%	2,7%	13,7%	24,8%	35,9%	47,0%	58,1%	69,2%
10,5%	-0,8%	9,9%	20,6%	31,3%	42,1%	52,8%	63,5%
13,5%	-4,3%	6,1%	16,4%	26,8%	37,1%	47,5%	57,8%
16,5%	-7,8%	2,2%	12,2%	22,2%	32,2%	42,2%	52,2%
19,5%	-11,2%	-1,6%	8,0%	17,7%	27,3%	36,9%	46,6%
22,5%	-14,7%	-5,4%	3,9%	13,1%	22,4%	31,7%	40,9%
25,5%	-18,1%	-9,2%	-0,3%	8,6%	17,5%	26,4%	35,3%



DCF Scenarios

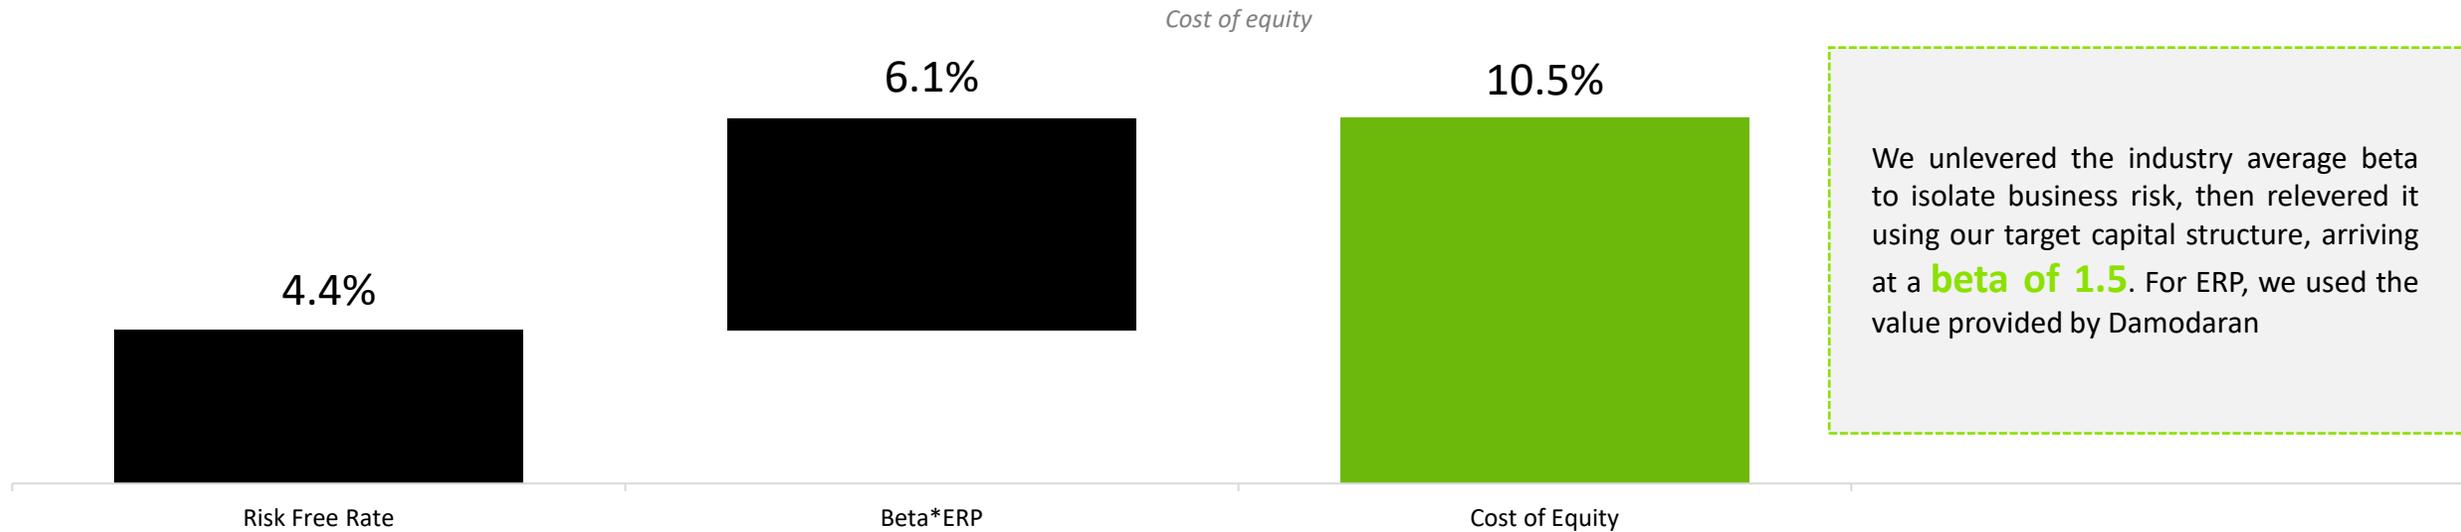
	Bear	Base	Bull
Marginal CapEx to Data Center (%)	70.0%	80.0%	90.0%
CapEx destined to Accelerators (%)	50.0%	60.0%	70.0%
NVIDIA GPU Market Share (%)	70.0%	80.0%	90.0%
Gross Margin	65%	75%	80%
Upside	(19.2%)	27.7%	72.8%



WACC

How we calculate our WACC

Cost of equity calculated based on CAPM



After to calculate the cost of equity, we calculated our WACC

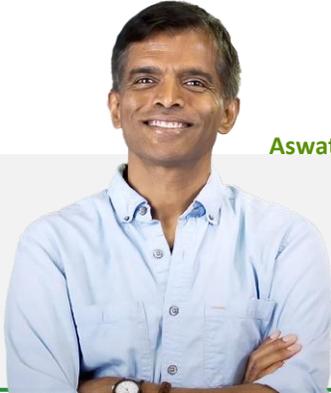


Perpetuity Growth

The reasons behind the growth rate we used in perpetuity

Perpetuity growth of 5.0% fits between risk-free rate and nominal growth

Aswath Insight



Aswath Damodaran – Professor at NYU

“One simple proxy for the nominal growth rate of the economy is the risk-free rate”

Risk free Rates and Nominal GDP Growth

- Risk free Rate = Expected Inflation + Expected Real Interest Rate**
- Nominal GDP Growth = Expected Inflation + Expected Real Growth**
- The real interest rate is what borrowers agree to return to lenders in real goods/services.
- The real growth rate in the economy measures the expected growth in the production of goods and services.

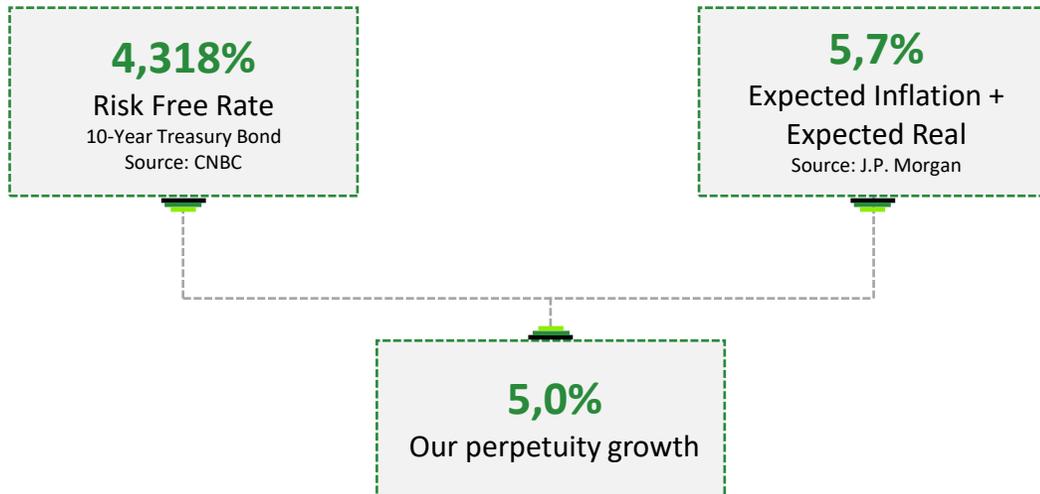
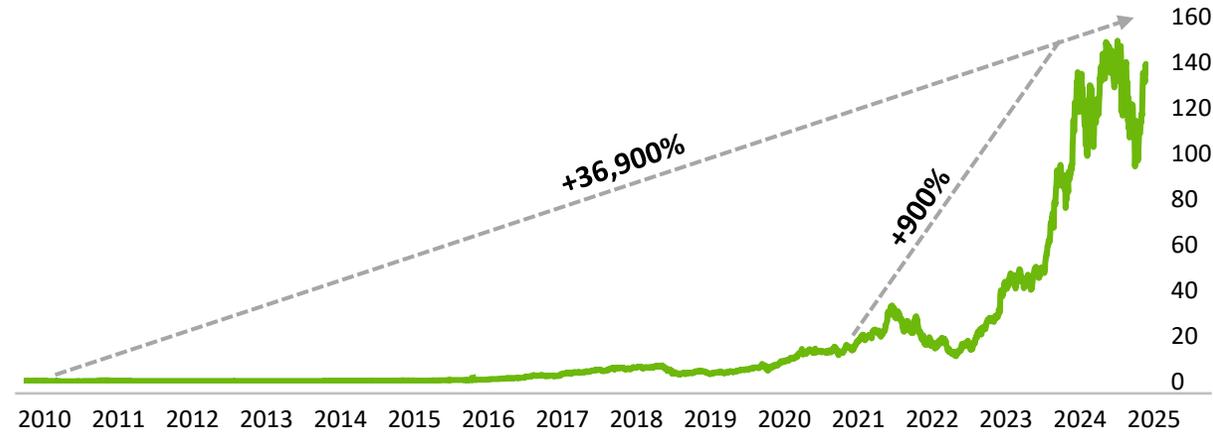
The argument for Risk free rate = Nominal GDP growth

- In the long term, the real growth rate cannot be lower than the real interest rate, since the growth in goods/services has to be enough to cover the promised rate.
- In the long term, the real growth rate can be higher than the real interest rate, to compensate risk taking. However, as economies mature, the difference should get smaller and since there will be growth companies in the economy, it is prudent to assume that the extra growth comes from these companies.

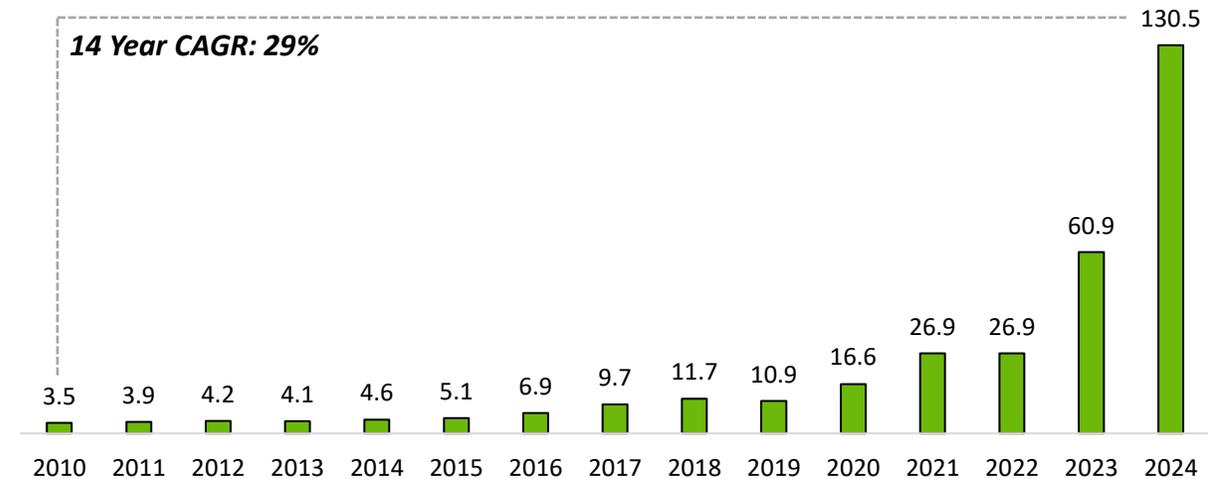
Period	10-Year T-Bond Rate	Inflation Rate	Real GDP Growth	Nominal GDP growth rate	Nominal GDP - T-Bond Rate
1964-2015	5.95%	3.41%	3.08%	6.47%	-0.52%
1964-1980	5.83%	4.49%	3.50%	7.98%	2.15%
1981-2008	6.88%	3.26%	3.04%	6.30%	-0.58%
2009-2015	2.57%	1.46%	1.47%	3.14%	0.57%

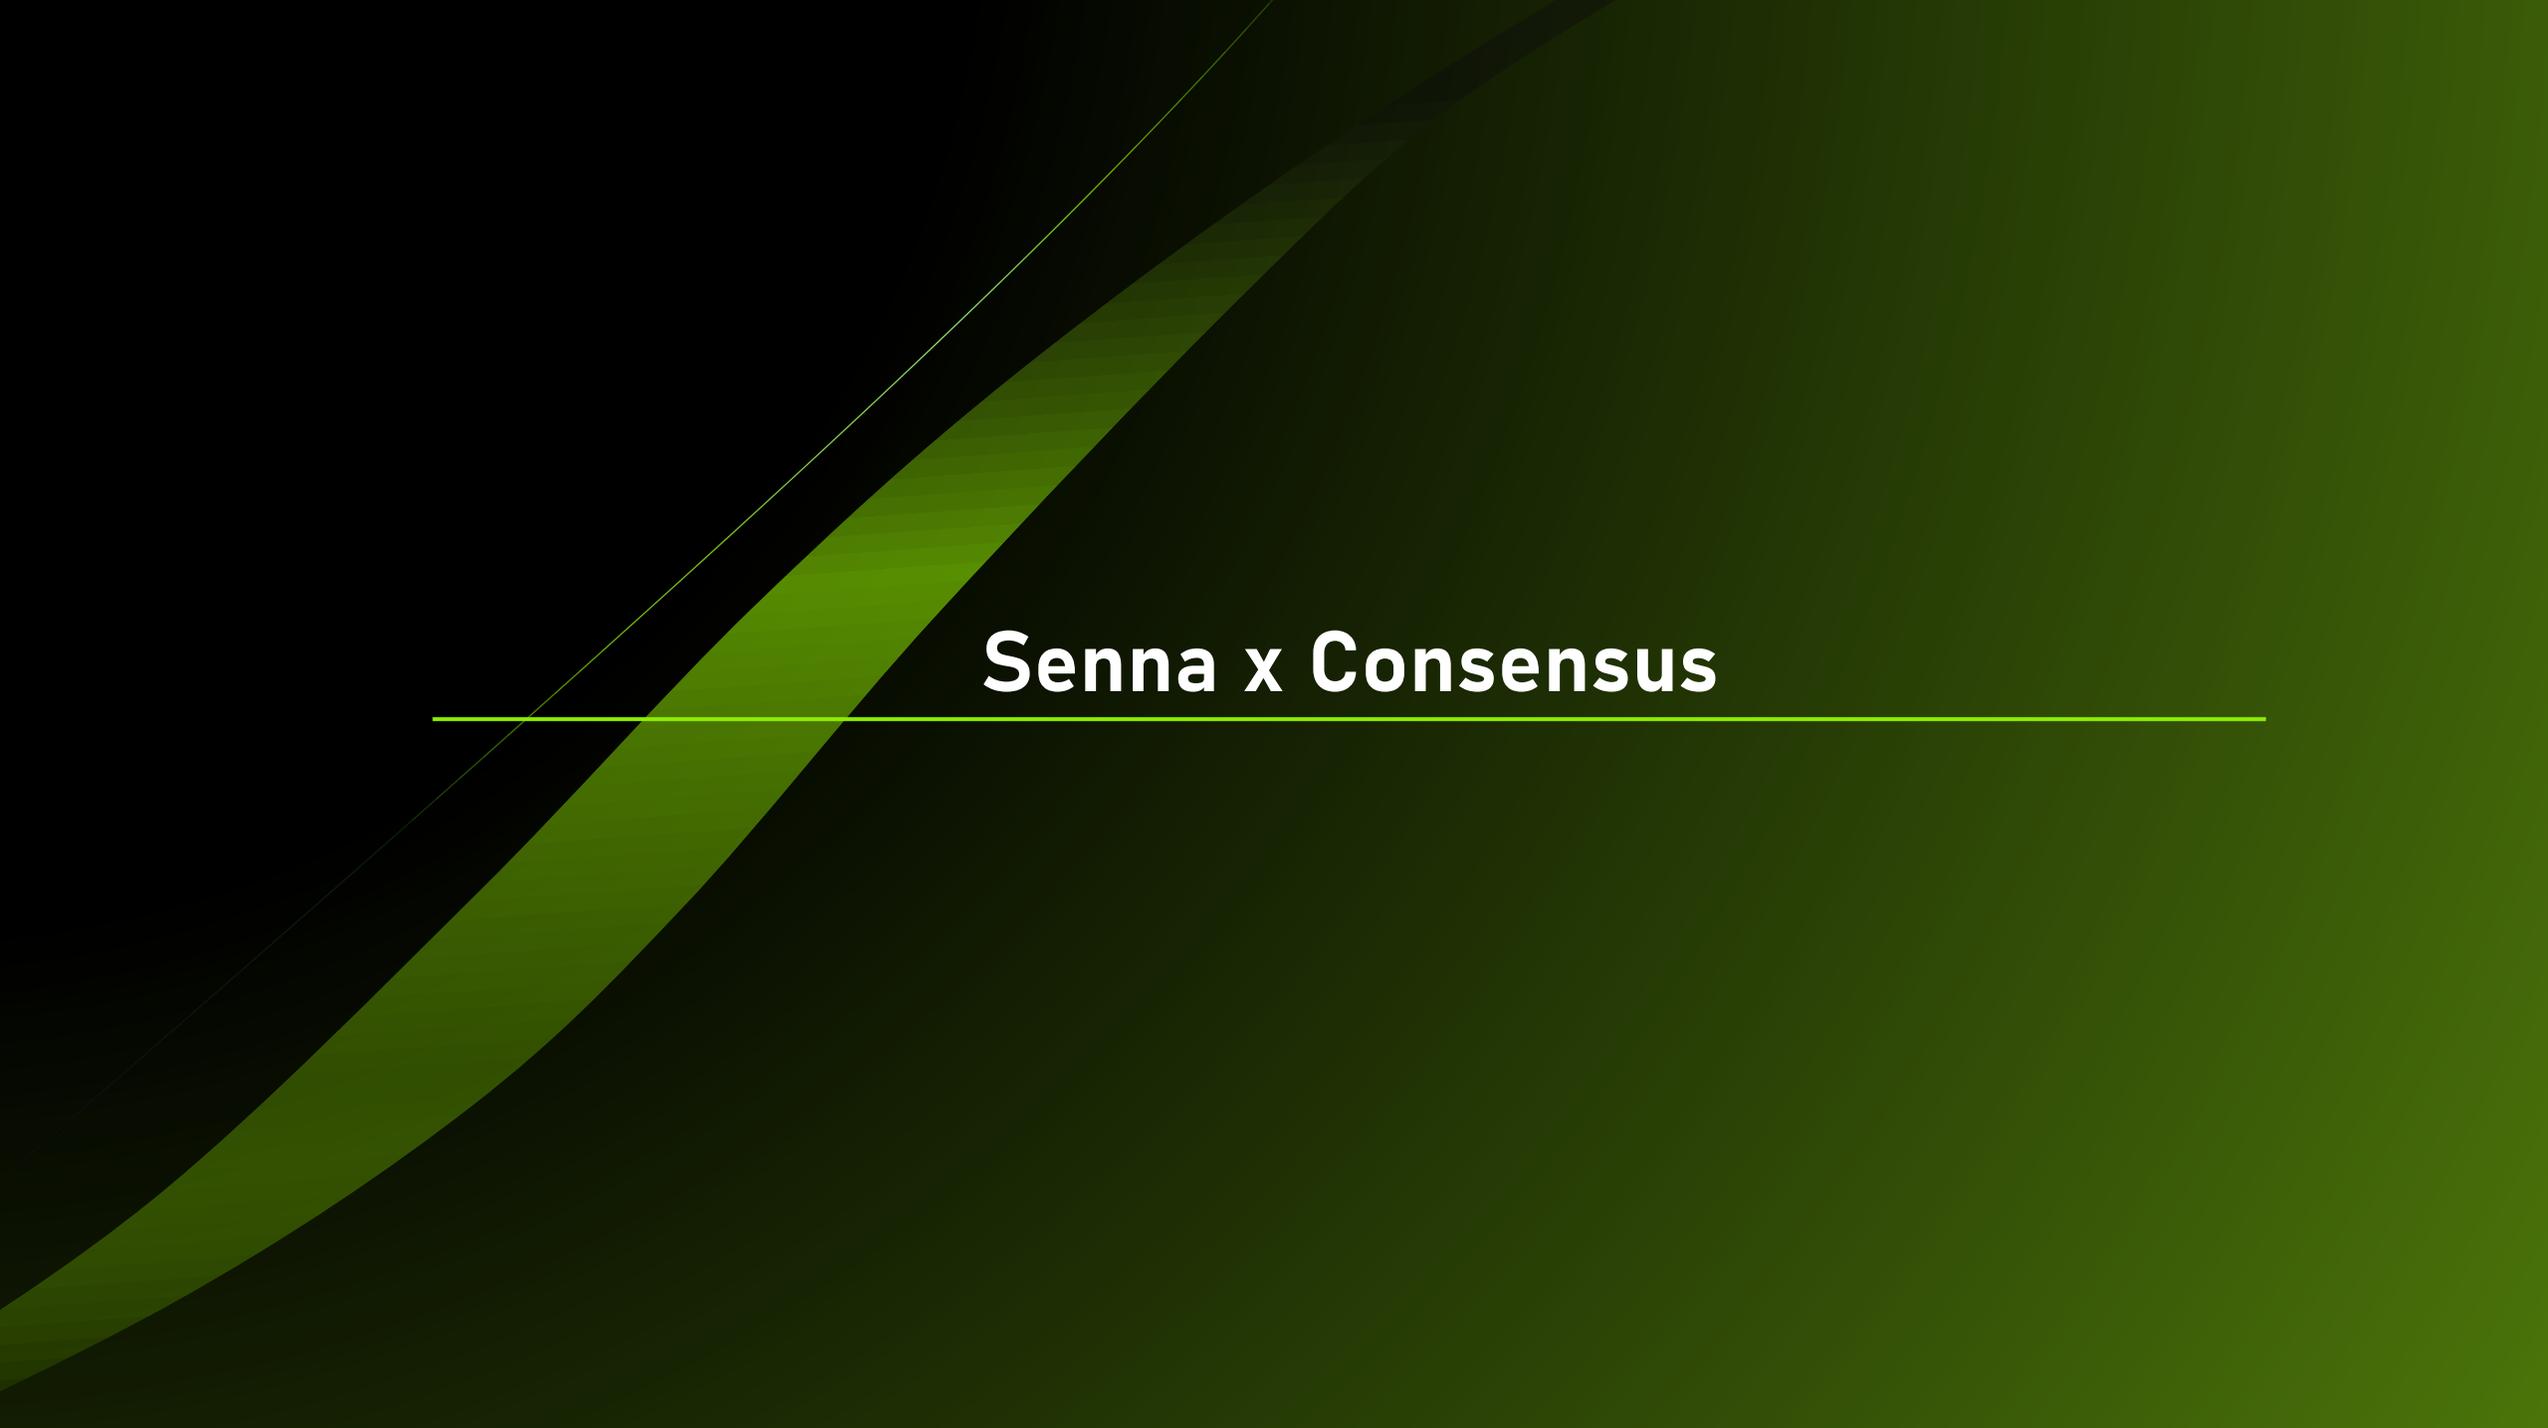
NVIDIA has shown steady growth, and we believe this will continue

NVDA's Price Action



NVIDIA's Revenue





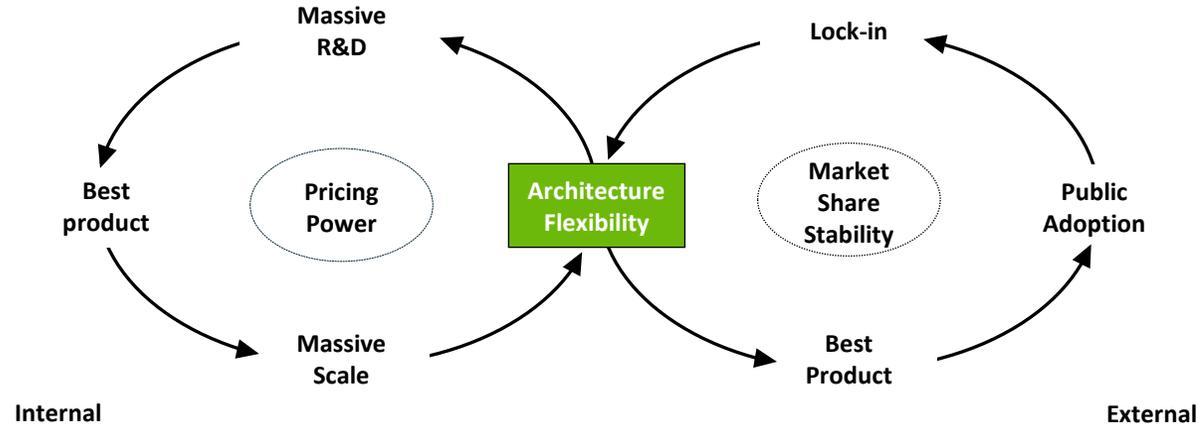
Senna x Consensus

Where we Disagree with the Consensus?

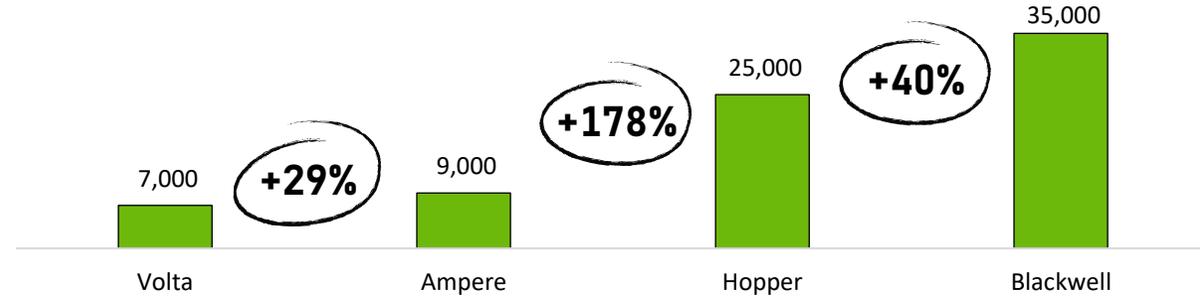
We diverge from consensus on gross margin, projecting 75% versus the consensus at 70%

The ecosystem is ruthless in locking in clients through switching costs and scale

NVIDIA Dynamics



II. Rare pricing power lets the company raise prices without pushback



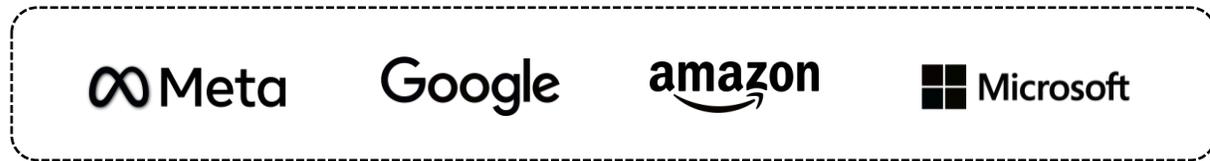
The Role of Switching Costs, Yale

" In software markets, switching costs are often particularly high due to incompatibility of data formats, user retraining, and the need to rewrite custom code. These costs can deter customers from switching even if alternative products are available and better. In markets with high switching costs, a monopolist can continue to charge high prices or maintain market dominance long after the competitive advantage that initially attracted customers has eroded "



III. Besides having pricing due to its ecosystem, it locks in clients through its contracts

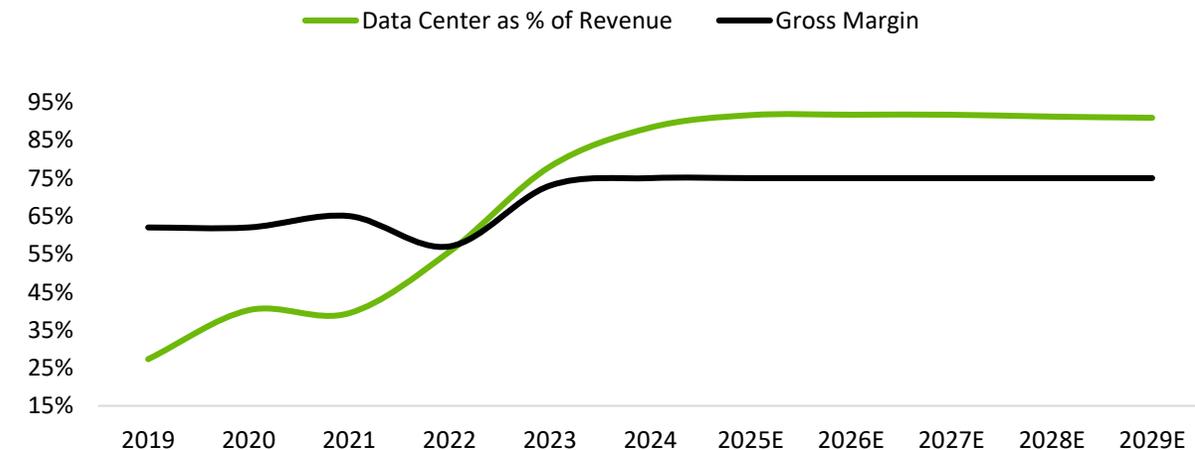
Multi-generation contracts



Clients must commit to purchasing future, unreleased GPU generations in order to access the current ones — agreeing to buy up to 3x the volume

IV. Higher Data Center share has driven NVIDIA's Gross Margin expansion

Data Center as % of Revenue x Gross Margin

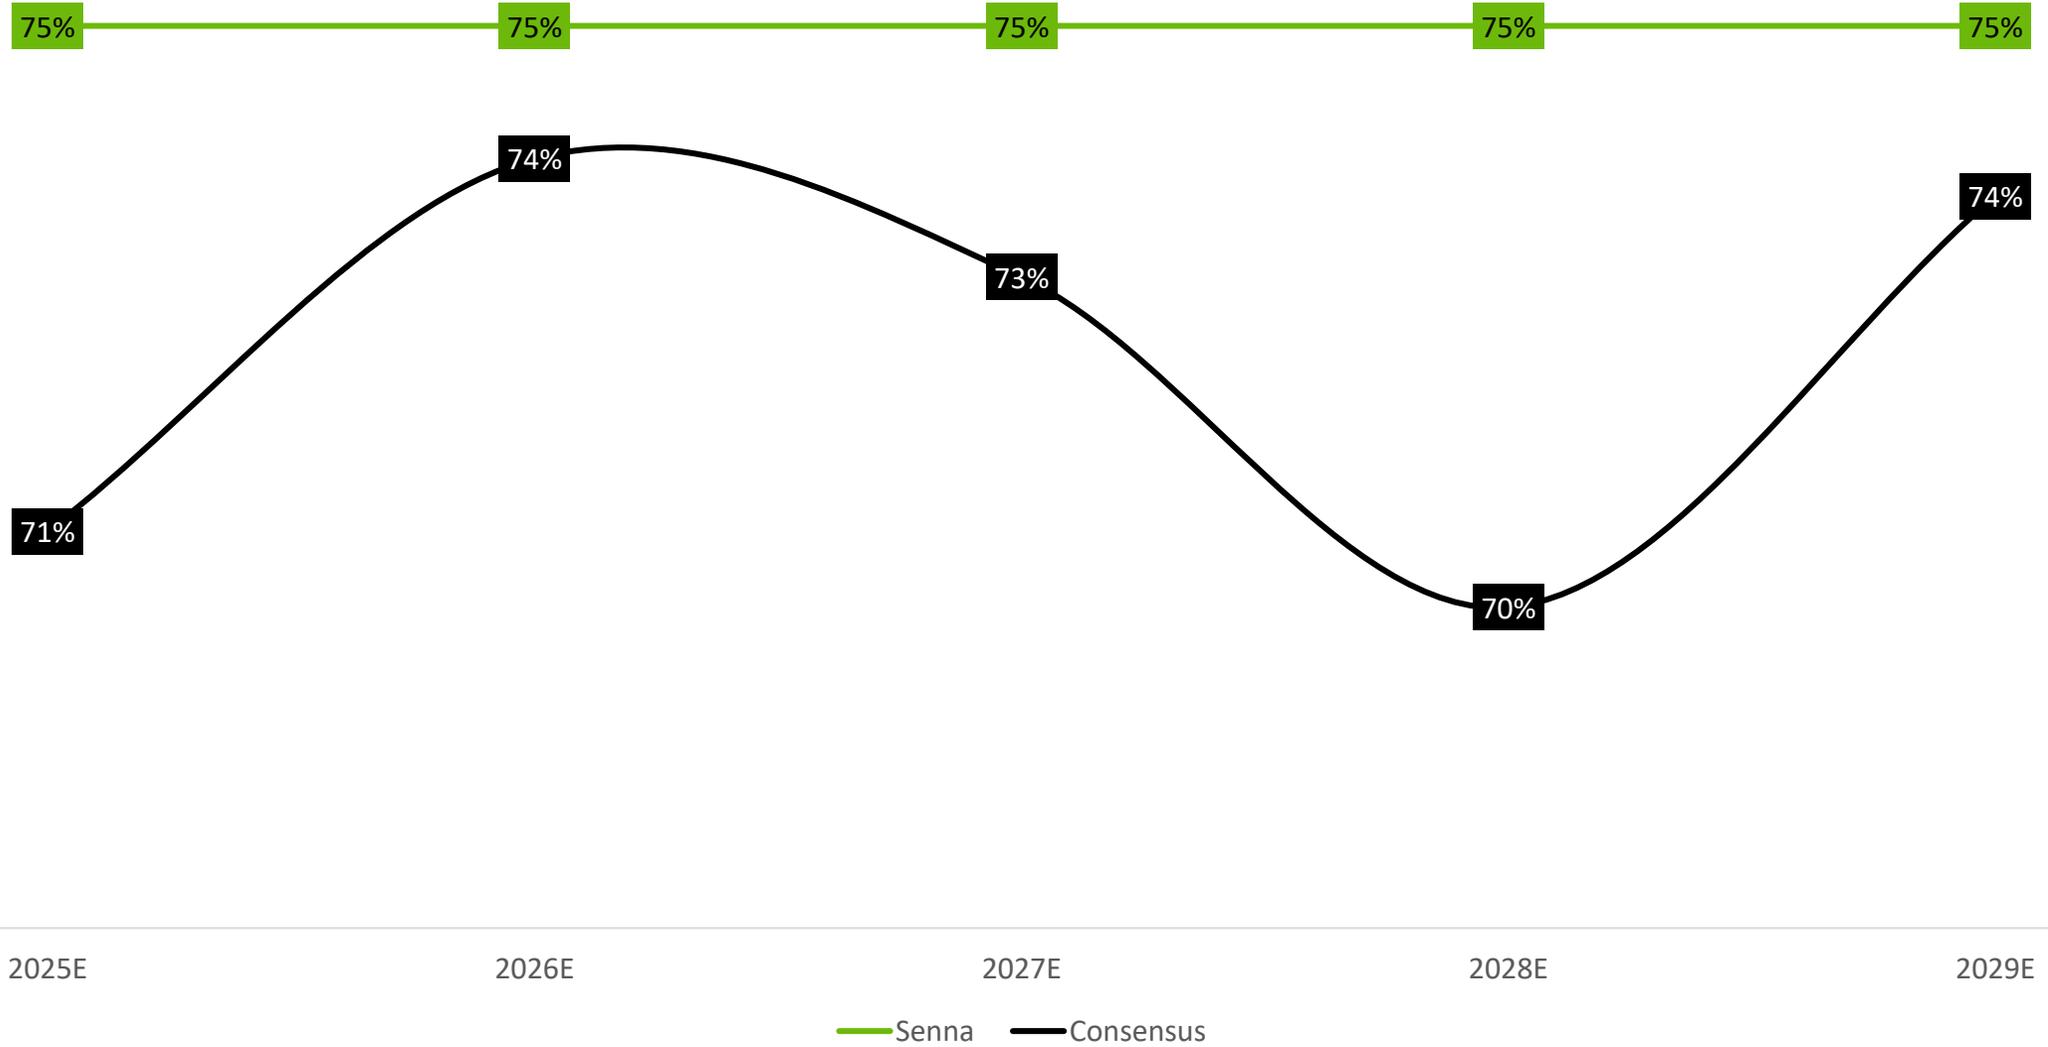


Supply Analysis

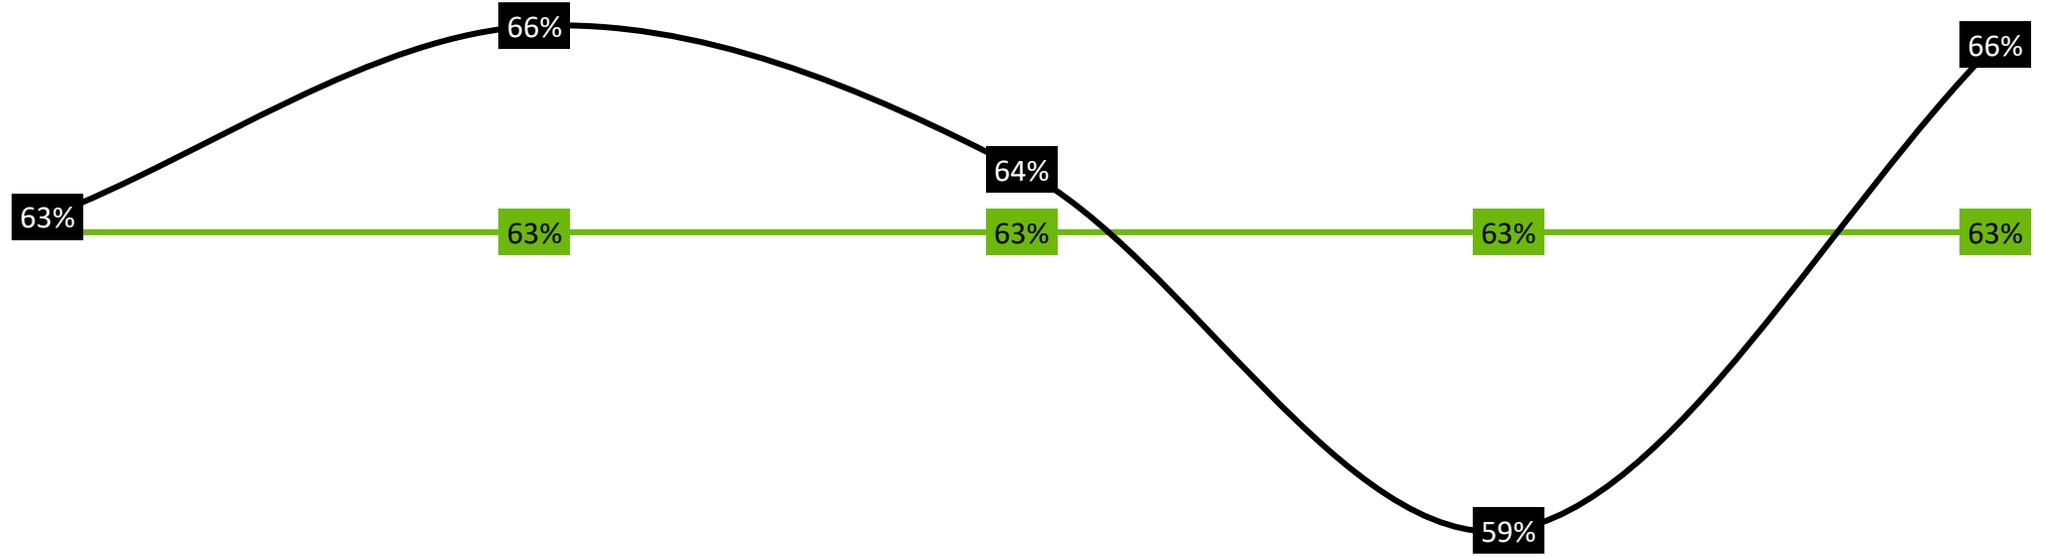
Supply Analysis	Q1	Q2	Q3	Q4
Wafers TSMC	120.000	156.000	184.000	210.000
Yield (%)	80%	80%	80%	80%
NVDA (%)	70%	70%	70%	70%
Wafers to NVIDIA	67.200	87.360	103.040	117.600
Others Production	30.000	30.000	30.000	30.000
Yield (%)	60%	60%	60%	60%
NVDA (%)	55%	55%	55%	55%
Wafers to NVIDIA	9.900	9.900	9.900	9.900
NVDA Total Wafers	77.100	97.260	112.940	127.500
Blackwell Wafers	52.428	68.082	90.352	121.125
% of Wafers to NVIDIA	68%	70%	80%	95%
Hopper Wafers	24.672	29.178	22.588	6.375
% of Wafers to NVIDIA	32%	30%	20%	5%
Blackwell Restriction	17	17	17	17
Hopper Restriction	29	29	29	29
NVIDIA Implied GPUs	1.606.764	2.003.556	2.191.036	2.244.000
Blackwell	891.276	1.157.394	1.535.984	2.059.125
Hopper	715.488	846.162	655.052	184.875
ASP Blackwell	31.000	31.000	31.000	31.000
ASP Hopper	20.000	18.000	15.000	15.000
Data Center Computing Revenue	41.939.316.000	51.110.130.000	57.441.284.000	66.606.000.000
		217.096.730.00		
NVDA Data Center Computing Revenue		0		



Gross Margin



EBIT Margin



2025E

2026E

2027E

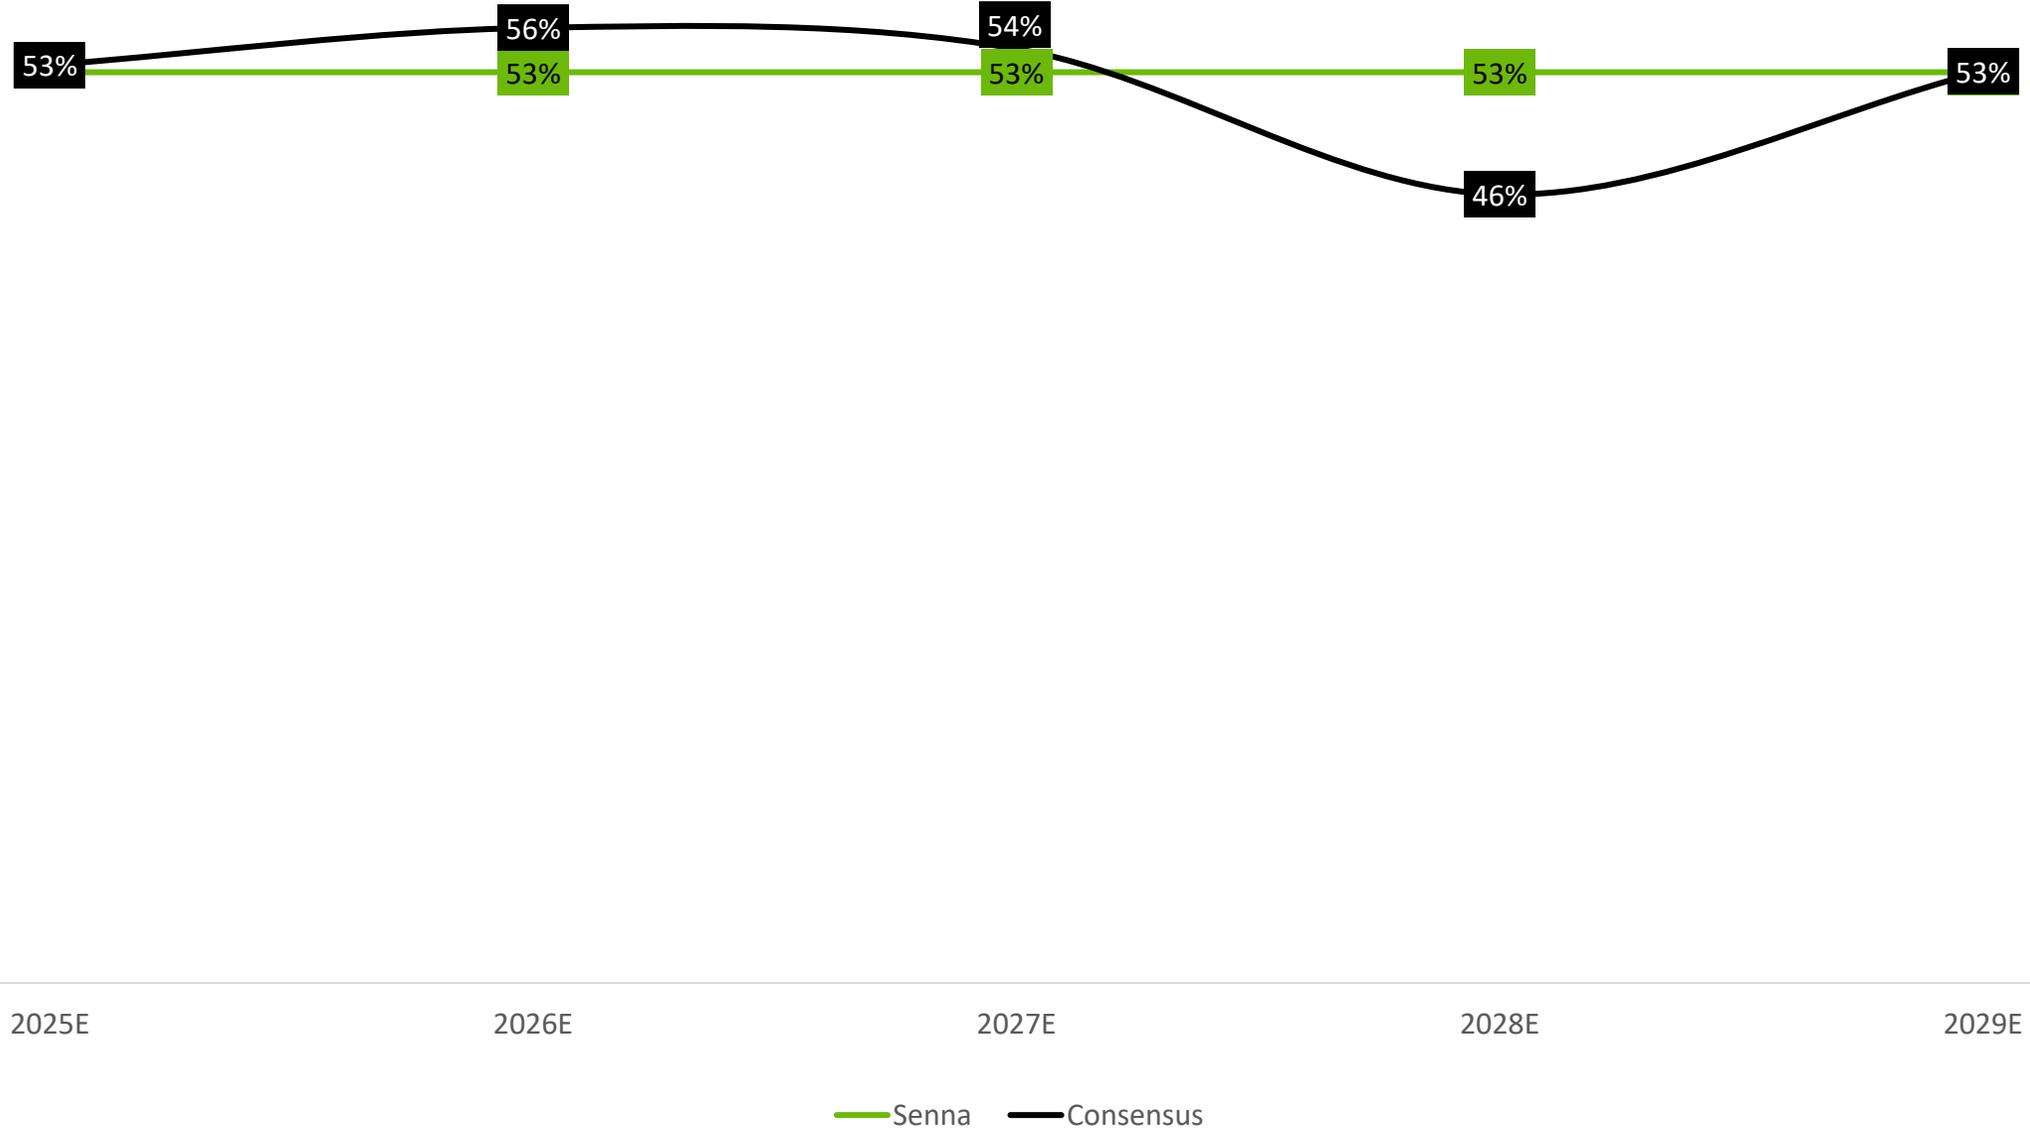
2028E

2029E

— Senna — Consensus



Net Margin



Sequoia Paper

Is AI a bubble?

I. The Sequoia paper questions if AI infrastructure investments are growing faster than revenue can justify

What we disagree about paper opinion?



Paper Overview

Sequoia Capital, in the article "AI's \$600B Question", questions whether AI infrastructure investments are growing faster than the market's revenue potential. They warn of a possible \$600 billion annual gap. Still, Sequoia is optimistic about founders focused on real end-user value. **The key message: be cautious, avoid the hype, and focus on sustainable growth**

1. NVIDIA's Ecosystem Creates Long-Term Value

Technologies like CUDA and NVLink build strong lock-in and protect GPU value over time

2. The AI Market Will Scale Faster Than Expected

Rapid AI adoption will unlock new uses and revenues, helping to close the gap

3. Infrastructure Is a Strategic Foundation, Not Excess

Current AI infrastructure investments are essential to support sustained market expansion

II. With our cost assumptions, the AI revenue required for payback drops significantly, suggesting the gap may be smaller than initially projected

	Q4 2023 ESTIMATE	Q4 2023 ACTUAL	Q1 2024 ACTUAL	Q4 2024 ESTIMATE
NVDA Data Center Run-Rate Revenue	\$50	\$74	\$90	\$150
Data Center Facility Build and Cost to Operate	50%	50%	50%	50%
Implied Data Center AI Spend	\$100	\$147	\$181	\$300
Software Margin	50%	50%	50%	50%
AI Revenue Required for Payback	\$200	\$294	\$363	\$600

Expected NVIDIA Revenue ÷ Data Center Costs = Implied AI Ramp Speed

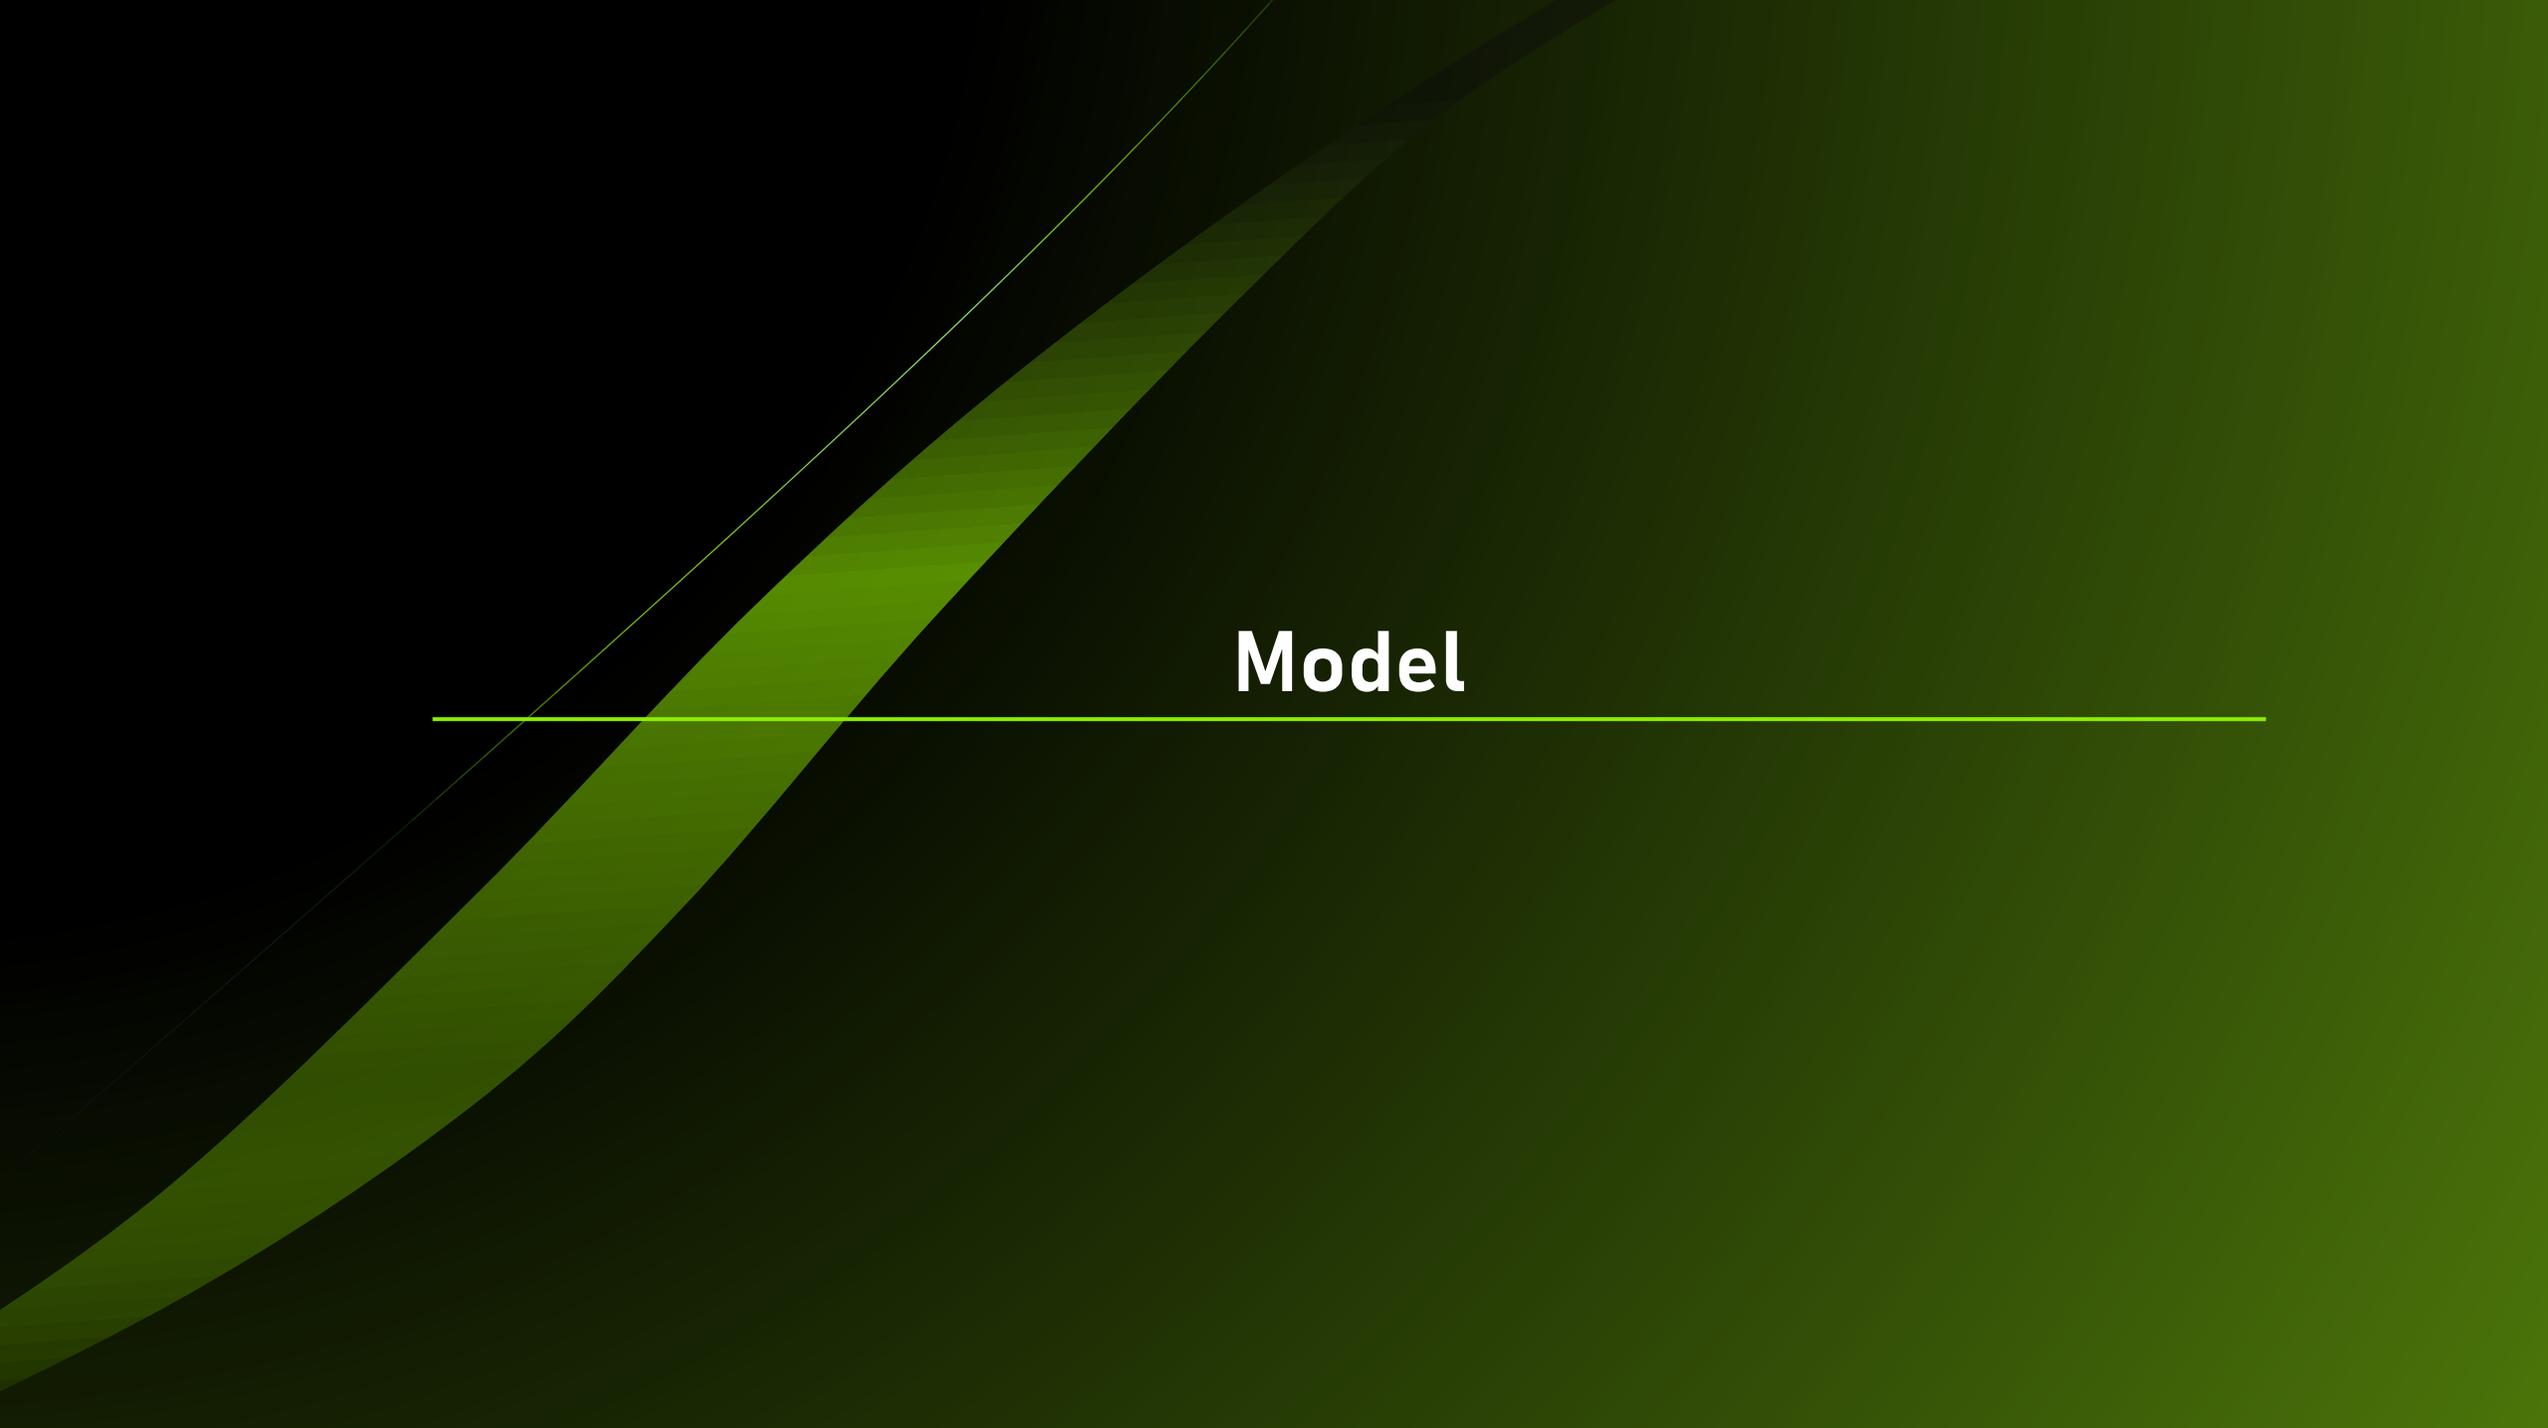
Implied AI Ramp Speed ÷ Software Margin = AI Required Revenue for Payback

US\$ Billion	Q4 2023 Actual	Q1 2024 Actual	Q4 2024 Estimate
NVIDIA Data Center Run-Rate Revenue	74	90	150
Data Center Facility Build and Cost to Operate	60%	60%	60%
Implied Data Center AI Sped	123	150	250
Software Margin	75%	75%	75%
AI Revenue Required For Payback	164	200	333

Data Center Facility Build and Cost to Operate: Using 60% better reflects current data center efficiency, with lower marginal costs, scale gains, and infrastructure reuse, making the assumption more realistic

Software Margin: A 75% software margin is justified by low marginal costs, premium pricing, and strong lock-in from proprietary ecosystems like CUDA



An abstract graphic featuring a dark green background. A thick, bright green curve starts from the bottom left and curves upwards towards the top right. A thin, light green horizontal line crosses the curve. The word "Model" is written in white, bold, sans-serif font, centered horizontally and positioned just above the horizontal line.

Model

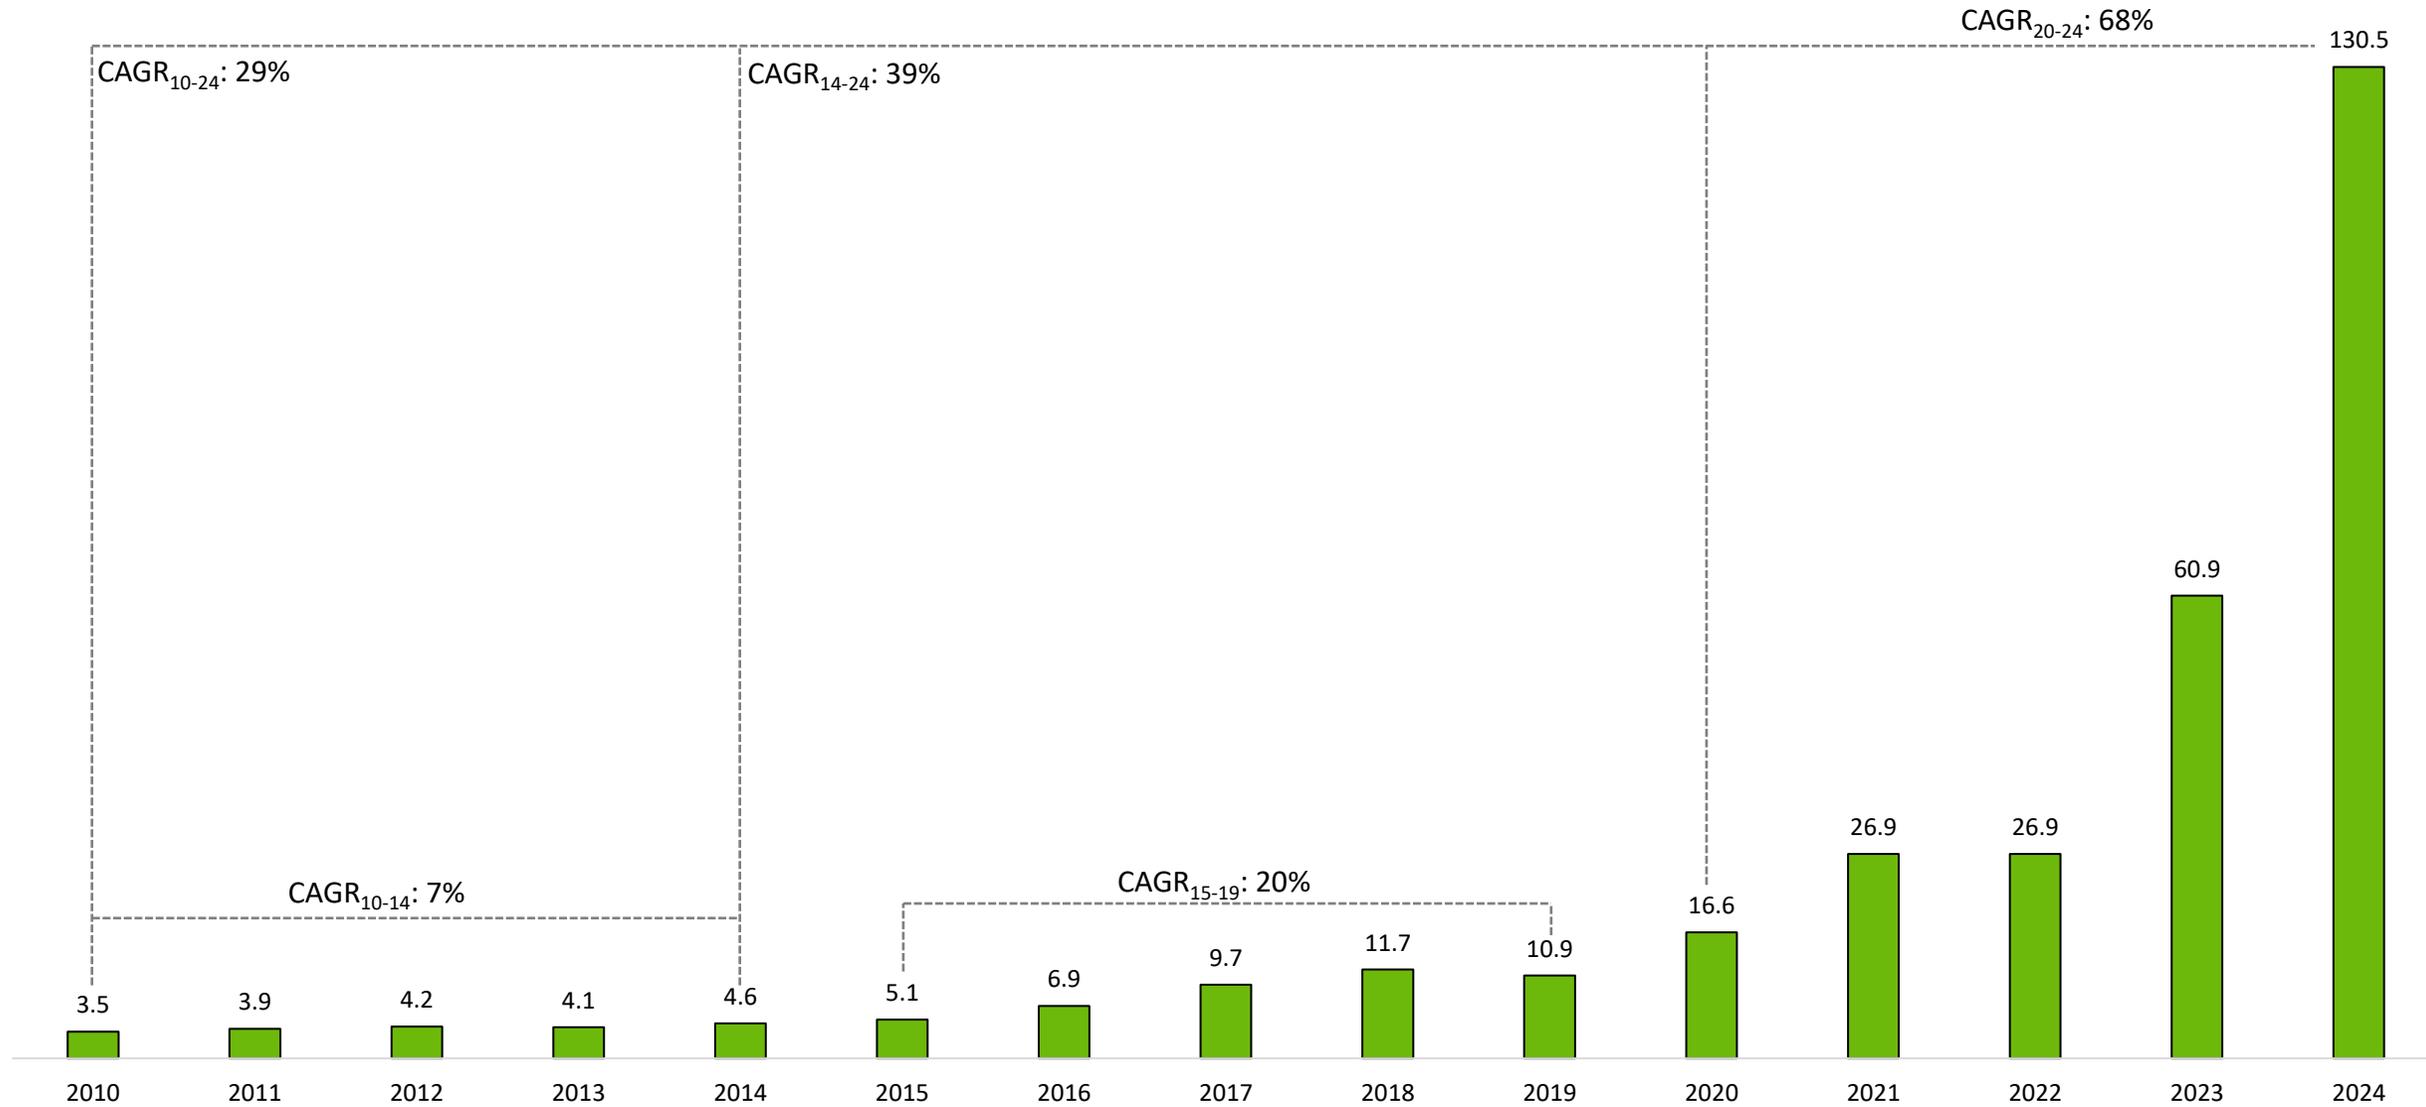
NVIDIA's Revenue

Nvidia's Data Center Revenue	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Data Center Revenue	2.983	6.696	10.613	15.005	47.525	115.186	202.269	243.501	276.262	307.217	351.234
YoY (%)		124%	58%	41%	217%	142%	76%	20%	13%	11%	14%
Compute	2.983	5.065	7.793	11.317	38.950	102.196	179.459	216.040	245.107	272.571	311.624
YoY (%)		70%	54%	45%	244%	162%	76%	20%	13%	11%	14%
% of Data Center Revenue	100%	76%	73%	75%	82%	89%	89%	89%	89%	89%	89%
Networking	-	1.631	2.820	3.688	8.575	12.990	22.811	27.461	31.155	34.646	39.610
YoY (%)			73%	31%	133%	51%	-	-	-	-	-
% of Compute Revenue		24%	27%	25%	18%	11%	11%	-	-	-	-
Nvidia's Gaming Revenue											
Gaming Revenue	5.518	7.759	12.462	9.067	10.447	11.350	13.717	16.021	17.241	19.785	22.705
YoY (%)		41%	61%	-27%	15%	9%	21%	17%	8%	15%	15%
PC Gaming	5.293	7.573	12.462	8.854	10.181	10.951	13.168	15.233	16.133	18.210	20.478
YoY (%)		43%	65%	-29%	15%	8%	20%	16%	6%	13%	12%
% of Gaming Revenue	96%	98%	100%	98%	97%	96%	96%	95%	94%	92%	90%
							20%	16%	6%	13%	12%
Console Gaming (Tegra)	225	186	-	213	266	399	548	788	1.108	1.575	2.227
YoY (%)		-17%	-100%	-	25%	50%	37%	44%	41%	42%	41%
% of Gaming Revenue	4%	2%	0%	2%	3%	4%	4%	5%	6%	8%	10%
							37%	44%	41%	42%	41%
Nvidia's Others Revenues											
Professional Visualization	1.212	1.053	2.111	1.544	1.553	1.878	1.972	2.090	2.237	2.415	2.633
YoY (%)		-13%	100%	-27%	1%	21%	5%	6%	7%	8%	9%
							5%	6%	7%	8%	9%
Automotive	700	536	566	903	1.091	1.694	2.460	3.523	4.975	6.925	9.502
YoY (%)		-23%	6%	60%	21%	55%	45%	43%	41%	39%	37%
							45%	43%	41%	39%	37%
OEM	505	631	1.162	455	306	389	401	413	425	438	451
YoY (%)		25%	84%	-61%	-33%	27%	3%	3%	3%	3%	3%
							3%	3%	3%	3%	3%



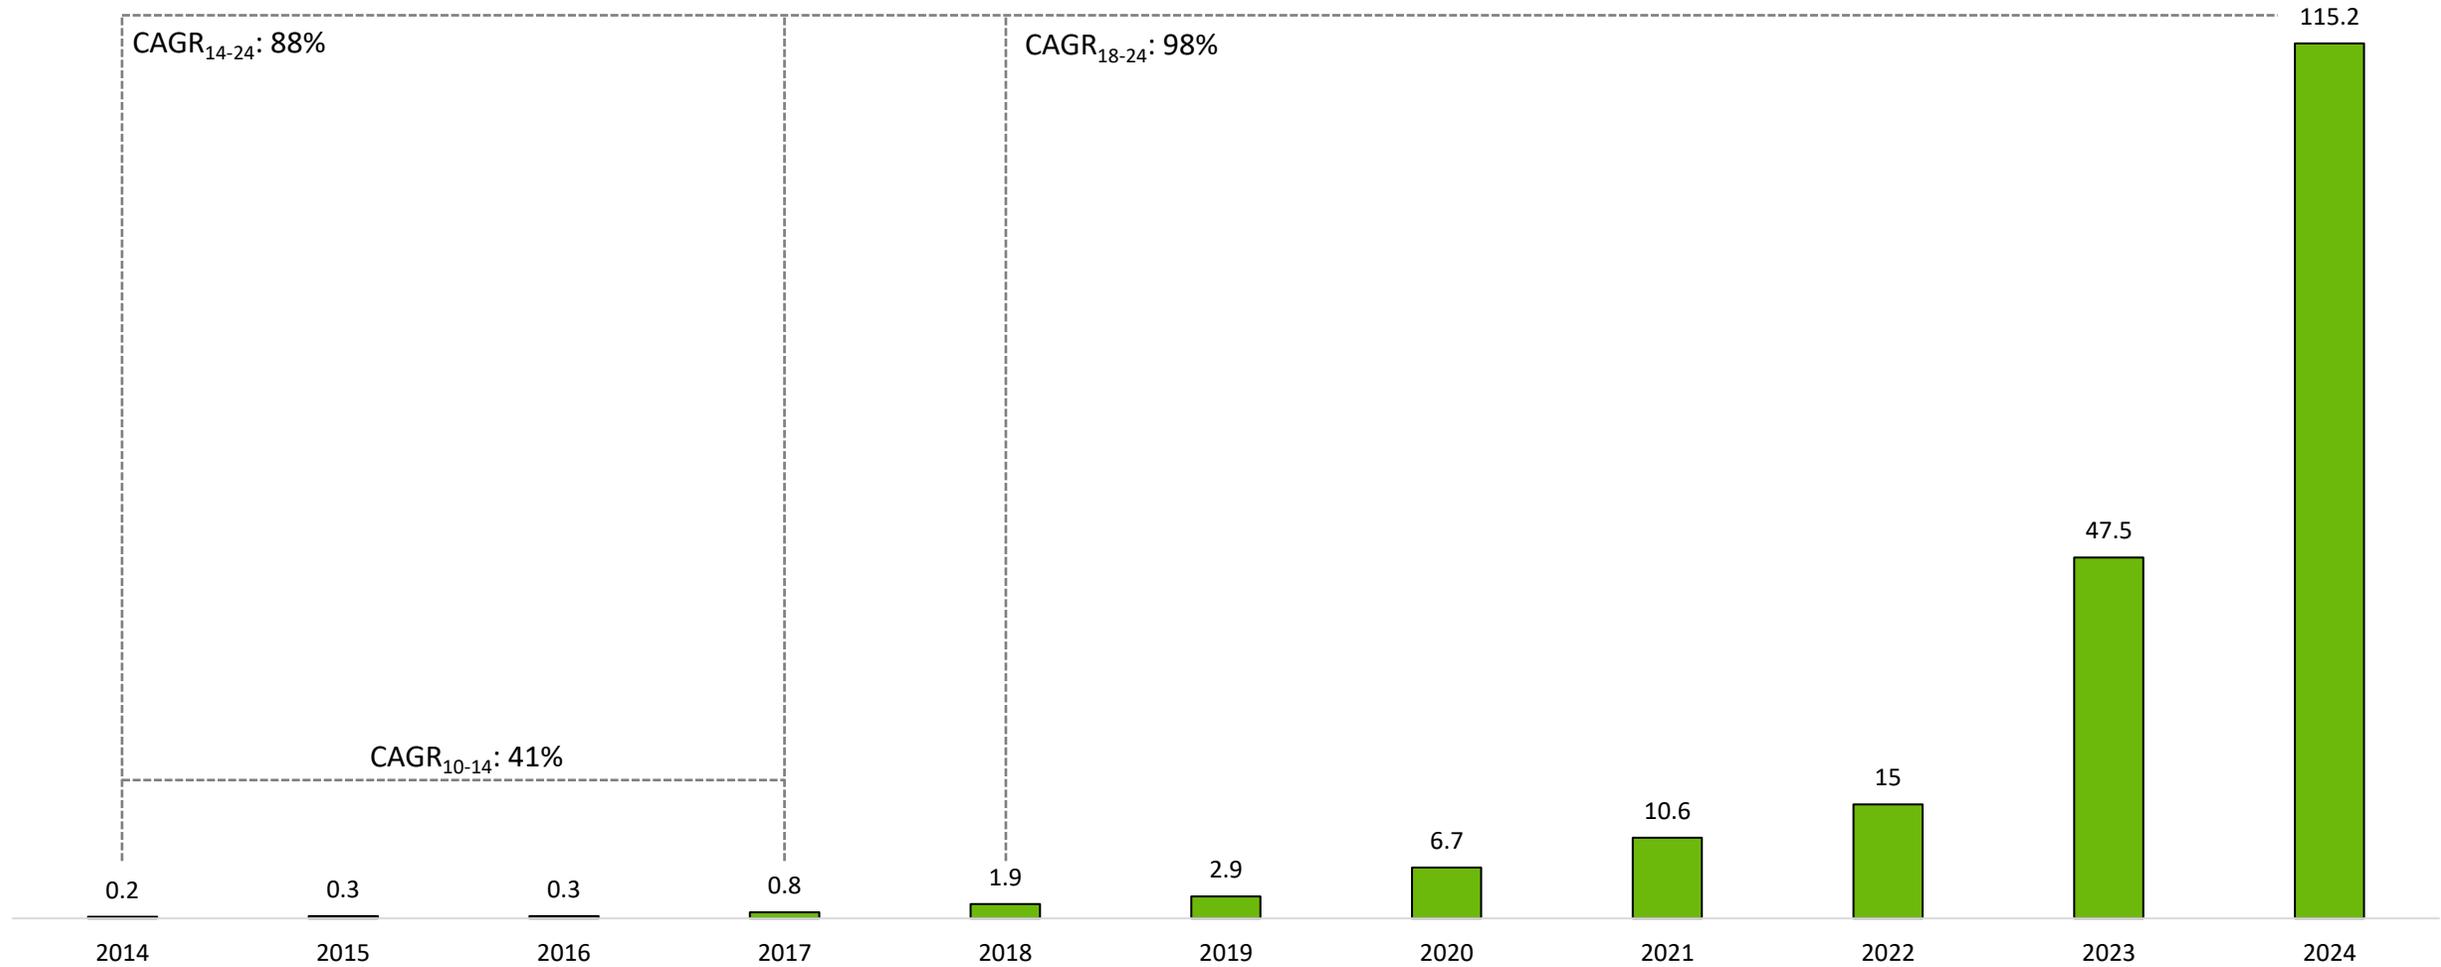
NVIDIA Total Revenue

NVIDIA's Total Revenue (US\$ Billion) 2010-2024



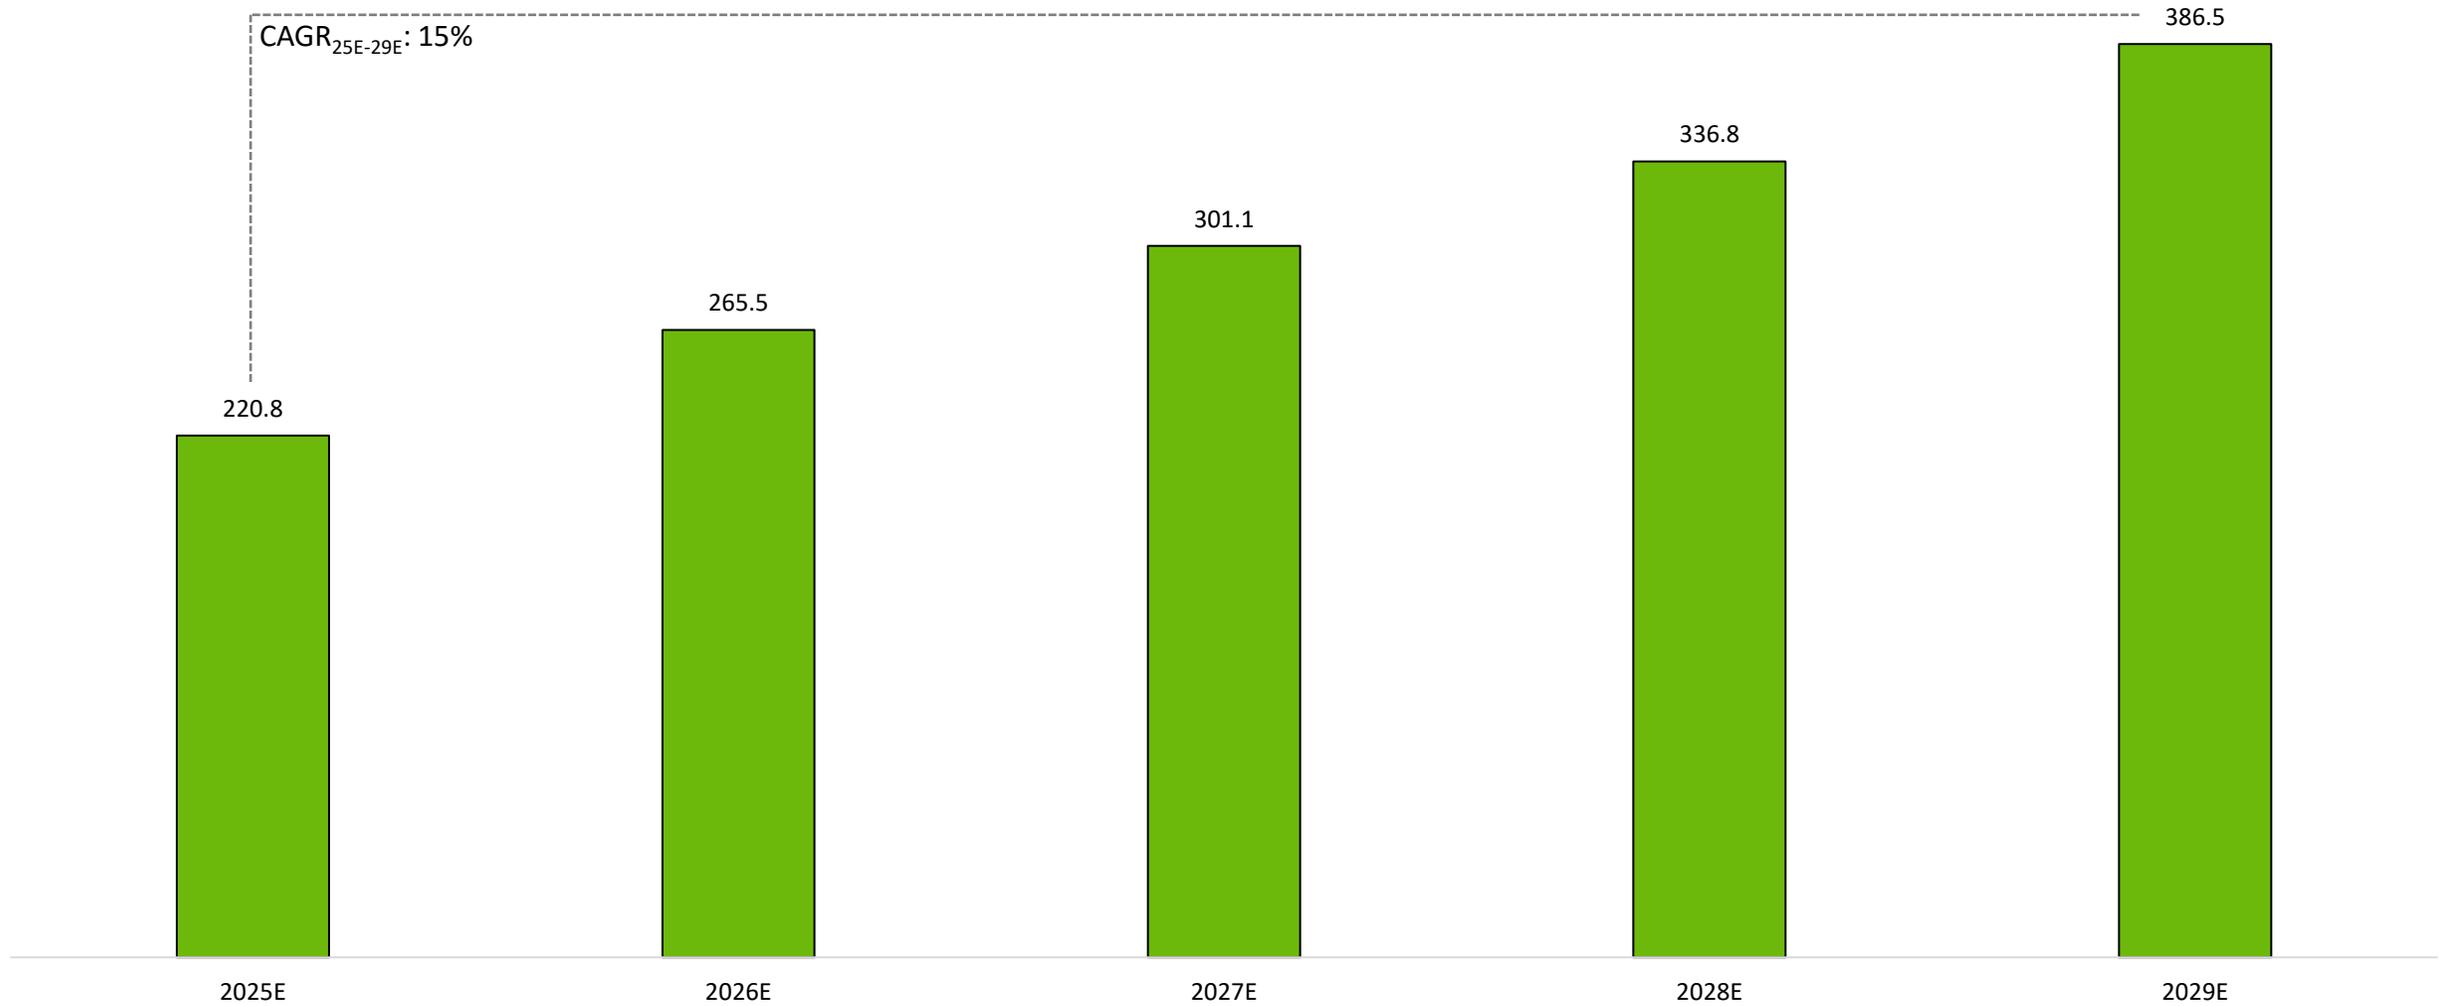
NVIDIA Data Center Revenue

NVIDIA's Total Revenue (US\$ Billion) 2010-2024



NVIDIA Projected Total Revenue

NVIDIA's Total Revenue (US\$ Billion) 2025E-2029E



NVIDIA Projected Data Center Revenue

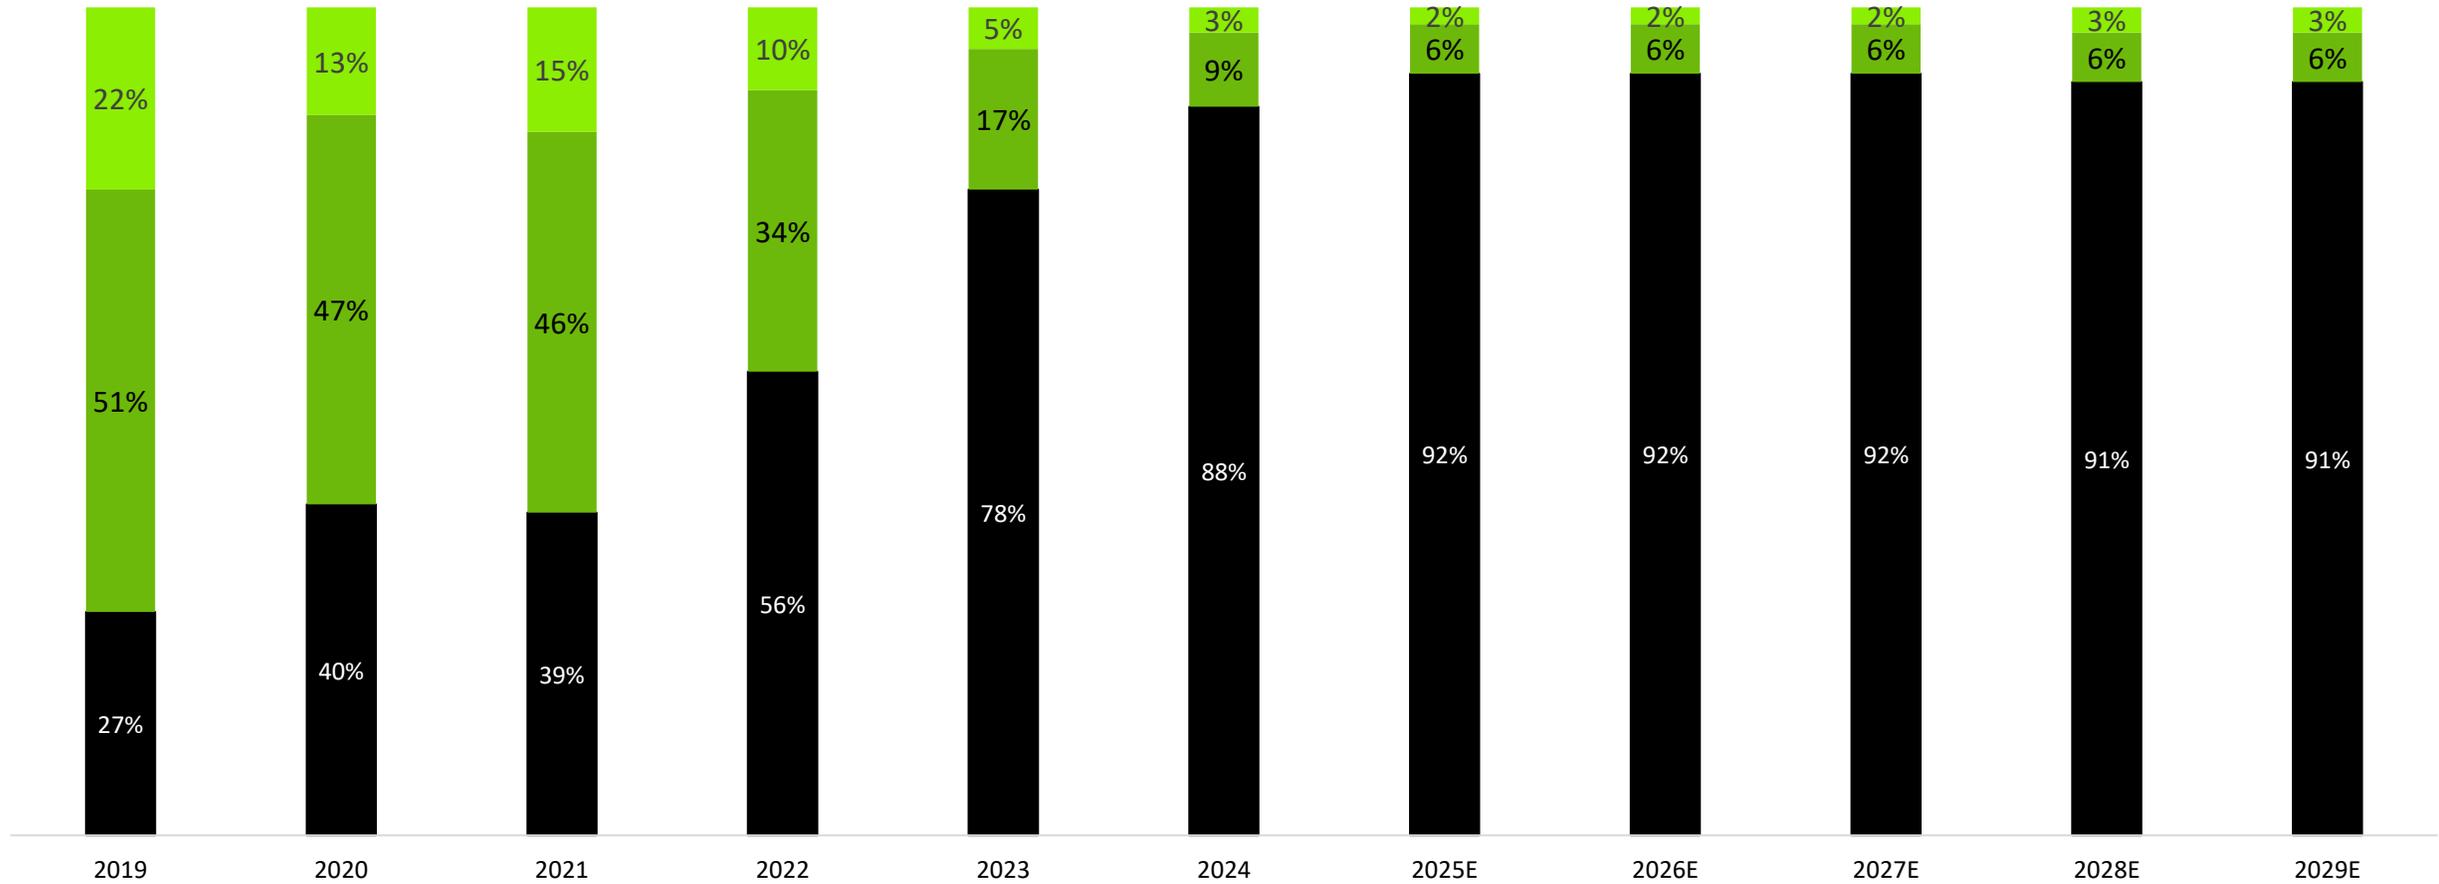
NVIDIA's Total Revenue (US\$ Billion) 2025E-2029E



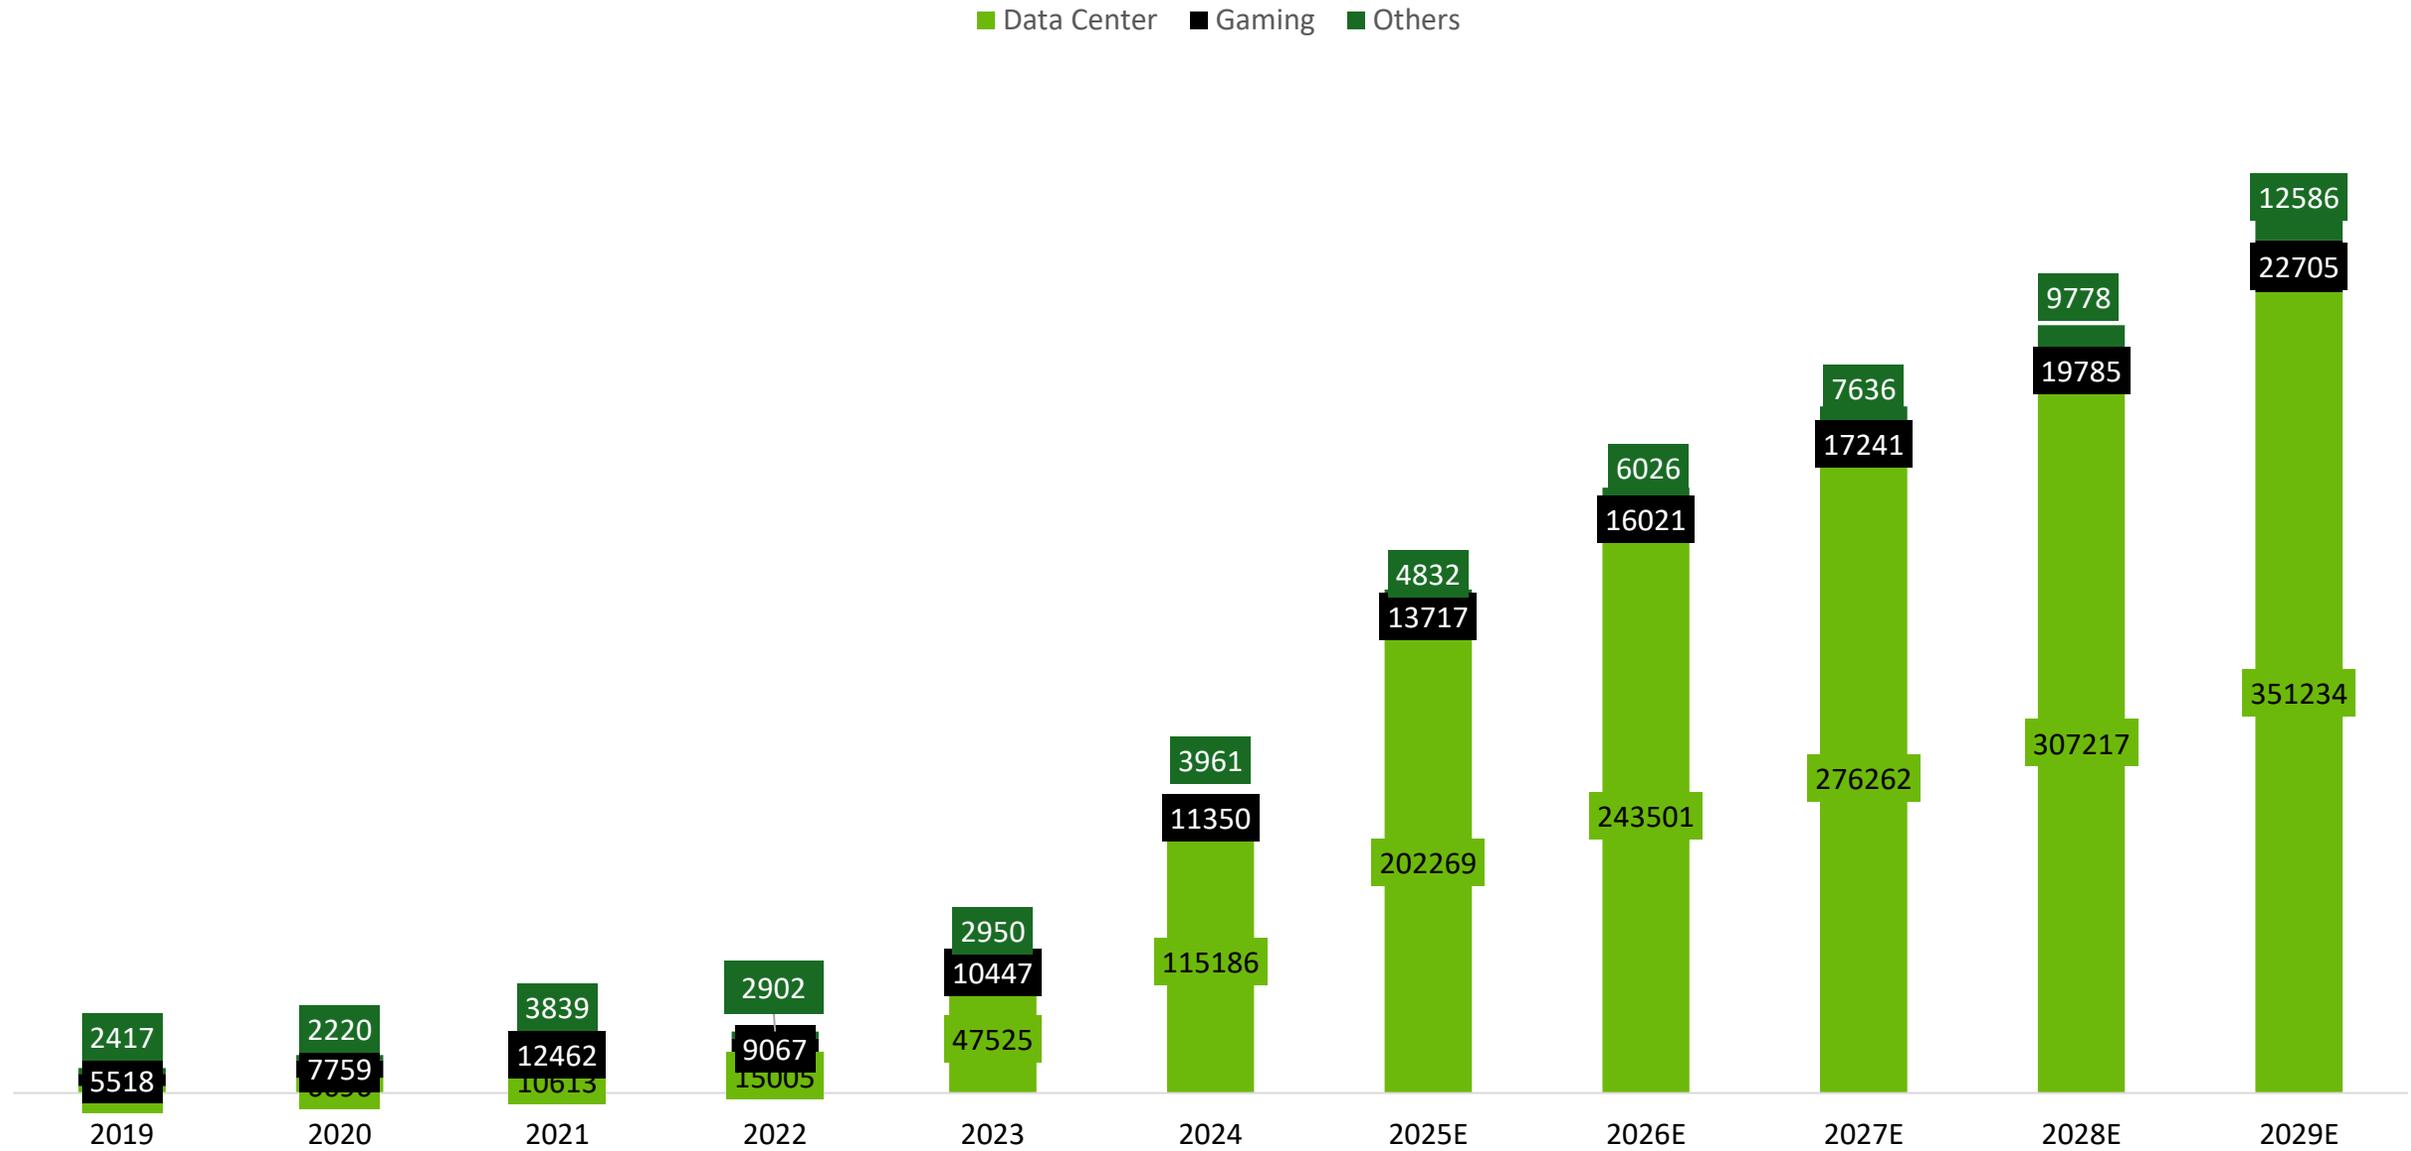
NVIDIA's Revenue per segment (%)

NVIDIA's Total Revenue (US\$ Billion) 2025E-2029E

■ Data Center ■ Gaming ■ Others

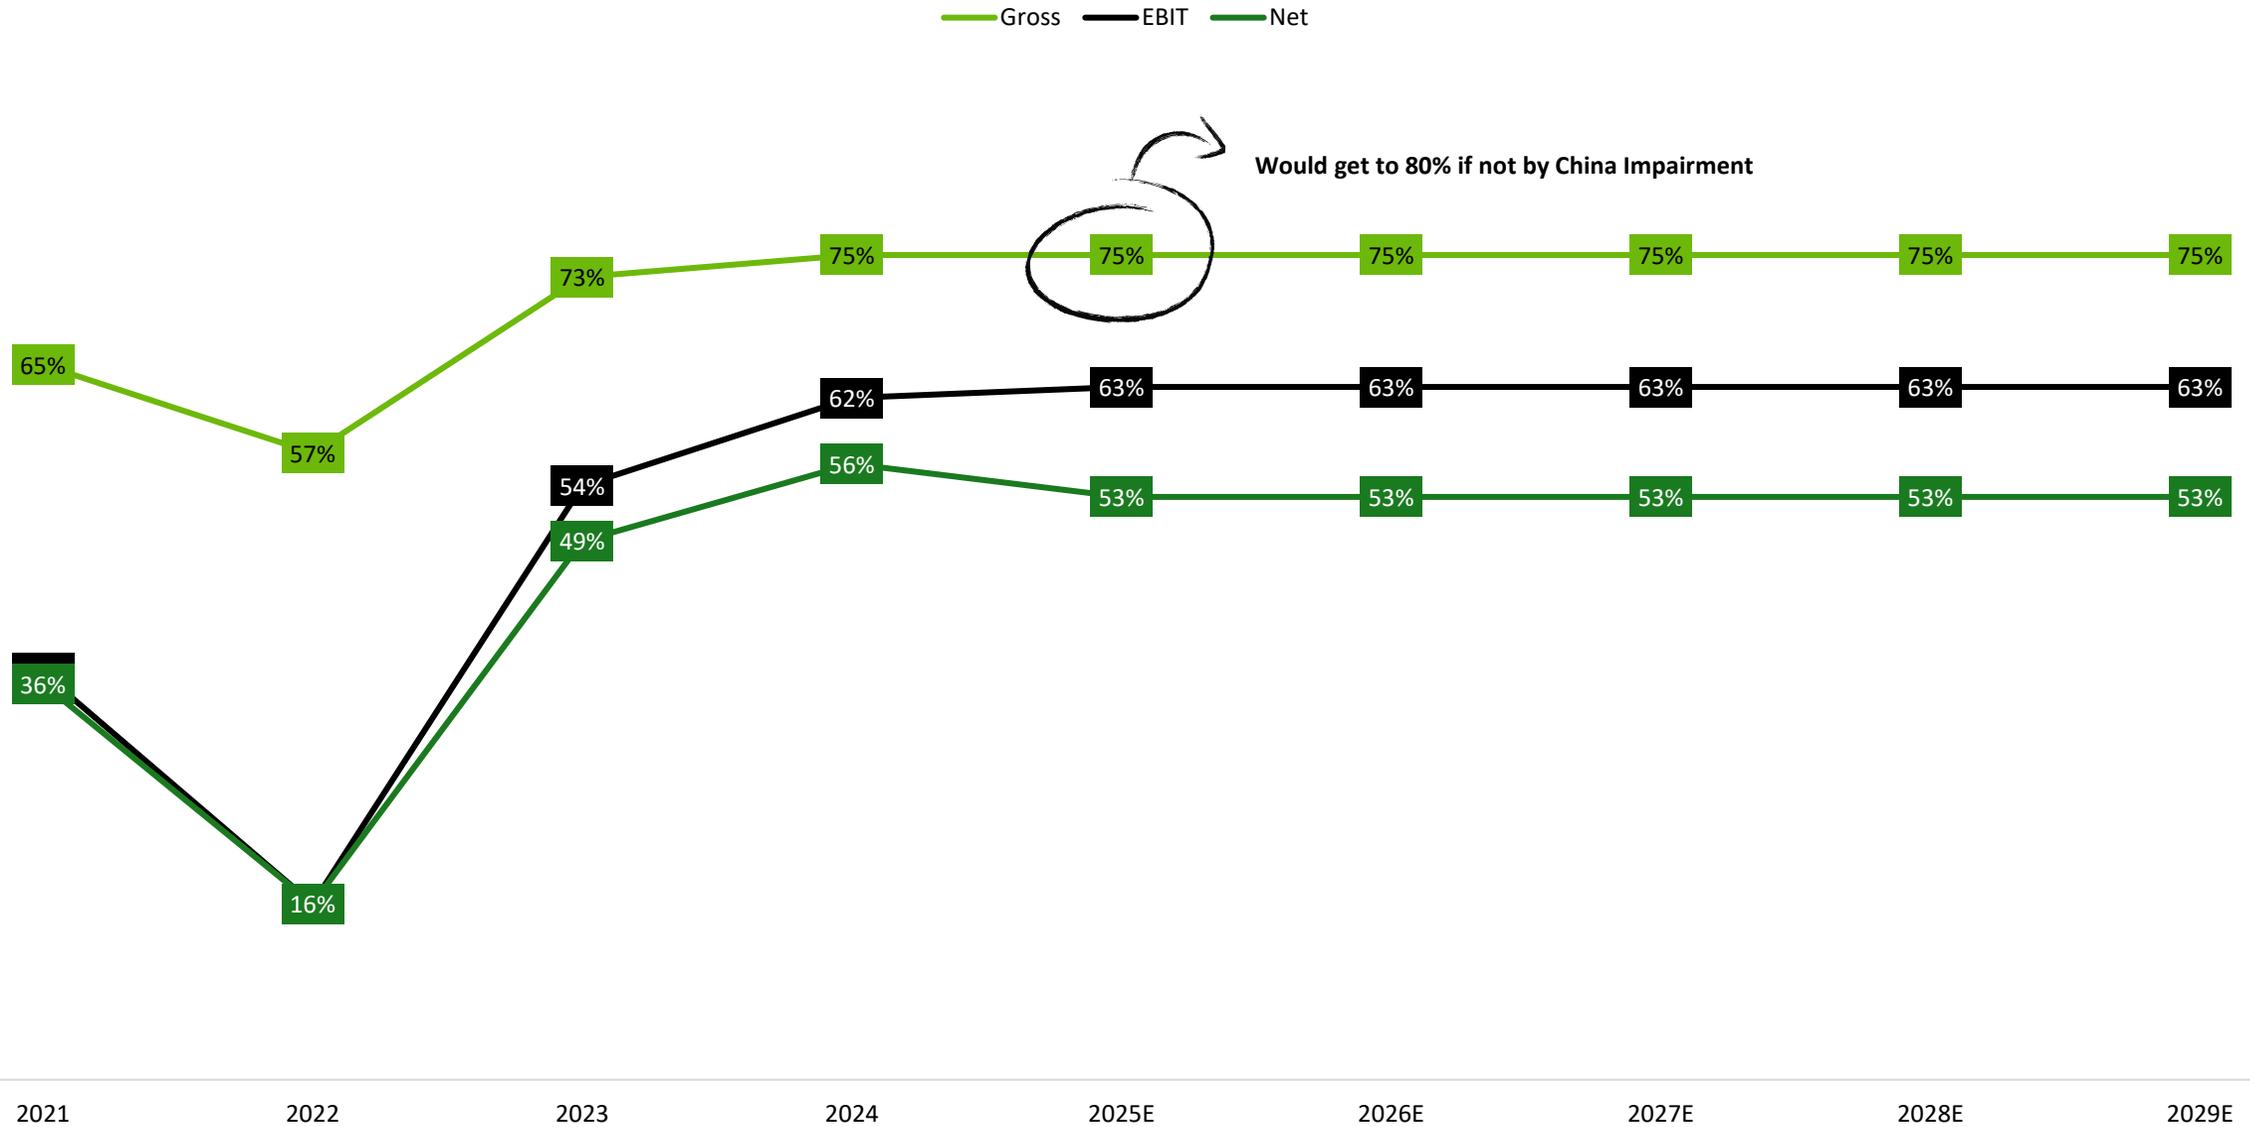


NVIDIA's Revenue per segment (US\$)

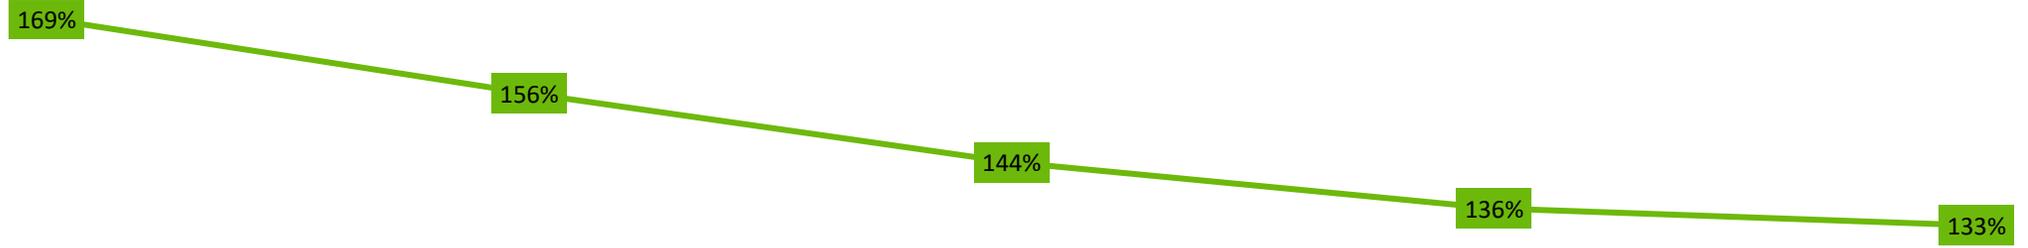


NVIDIA Projected Margins

NVIDIA's Margins 2025E-2029E



Projected ROIC



2025E

2026E

2027E

2028E

2029E



Sensitivity Analysis

We performed a sensitivity analysis to assess how changes in inputs would affect the project's IRR

		Exit Multiple						
		25x	26x	27x	28x	29x	30x	31x
Exit Year	IRR							
	2026	24,1%	28,9%	33,7%	38,5%	43,4%	48,2%	53,0%
	2027	19,8%	22,0%	24,3%	26,4%	28,6%	30,7%	32,8%
	2028	19,3%	20,8%	22,2%	23,2%	25,0%	26,4%	27,7%
	2029	19,2%	21,3%	22,3%	22,3%	23,4%	24,3%	25,3%



Revenue Build Up (TAM)

Unit: US\$ Milion	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
TAM Comparatives											
Dell'oro (Jensen) Global Data Center Capex				220.000	260.000	460.000	556.600	673.486	814.918	986.051	1.193.122
YoY (%)					18%	77%	21%				
Nvidia's Data Center Revenue	2.983	6.696	10.613	15.005	47.525	115.186	244.904	304.079	376.492	464.923	572.698
% of Global Capex				7%	18%	25%	44%	45%	46%	47%	48%
% of Data Center Capex for GPUs							50%	53%	55%	58%	60%
% of NVIDIA GPU Market Share							88%	86%	84%	82%	80%
McKinsey											
Data Ceter CAPEX driven by AI							556.600	656.788	775.010	914.512	1.079.124
% IT equipment (CPUs, GPUs, memory...)							60%				
% of IT equipment for GPUs							60%				
% of IT equipment for CPUs							10%				
% of NVIDIA Market Share							88%	86%	84%	82%	80%
CAGR Capex							18%				
Nvidia's Data Center Revenue							205.719	237.232	273.423	314.958	362.586
AMD (Lisa Su) Total Accelerators Market Revenue	2.983	5.065	11.487	17.360	45.446	127.528	179.451	252.515	355.327	500.000	703.576
NVIDIA	2.983	5.065	7.793	11.317	38.950	102.196					
AMD	-	-	3.694	6.043	6.496	12.579					
Market Share Assumption (%)						90%					
Implied CAGR (%)							41%	41%	41%	41%	41%
NVIDIA Accelerators Revenue							157.917	217.163	298.475	410.000	562.861
Market Share (%)							88%	88%	86%	84%	82%



Revenue Build Up (Hyperscale CapEx)

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Hyperscalers Capex											
Microsoft						22.239	37.829	44.039	50.698	56.221	62.379
YoY (%)											
Capex	13.925	15.441	20.622	23.886	28.107	44.477	65.265	72.607	80.052	86.955	94.652
(Δ) Capex for AI							15.591	21.801	28.460	33.983	40.140
							58%	61%	63%	65%	66%
Marginal Capex							20.788	28.130	35.575	42.478	50.175
Amazon						41.500	57.367	61.416	64.569	66.579	70.022
YoY (%)											
Capex	16.861	40.140	61.053	63.645	52.729	82.999	104.155	108.698	111.836	114.348	118.652
(Δ) Capex for AI							15.867	19.916	23.070	25.079	28.522
							55%	57%	58%	58%	59%
Marginal Capex							21.156	25.699	28.837	31.349	35.653
Oracle						3.433	17.012	19.237	20.263	21.627	26.029
YoY (%)											
Capex	1.660	1.564	2.135	4.511	8.695	6.866	24.971	27.258	27.903	29.608	35.111
(Δ) Capex for AI							13.579	15.804	16.830	18.194	22.596
							68%	71%	73%	73%	74%
Marginal Capex							18.105	20.392	21.037	22.742	28.245
Google						26.268	41.842	45.535	49.379	52.280	54.456
YoY (%)											
Capex	23.548	22.281	24.640	31.485	32.251	52.535	73.301	77.396	81.424	85.051	87.771
(Δ) Capex for AI							15.574	19.268	23.111	26.012	28.188
							57%	59%	61%	61%	62%
Marginal Capex							20.766	24.861	28.889	32.516	35.236
Meta						18.628	40.634	50.084	50.392	49.915	53.102
YoY (%)											
Capex	15.102	15.163	18.690	31.431	27.266	37.256	66.598	77.844	76.961	76.364	80.348
(Δ) Capex for AI							22.006	31.456	31.764	31.287	34.474
							61%	64%	65%	65%	66%
Marginal Capex							29.342	40.588	39.705	39.108	43.092

Bloomberg



Revenue Build Up (Premisses to Data Center)

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
% of Capex Destined for Data Center in Actual Capex						50%					
% of Capex Destined for Data Center in Marginal Capex							75%	78%	80%	80%	80%
Accelerators Capex						56.033	97.342	115.663	129.416	141.807	159.592
% of Hyperscalers Technology Capex						50%	50%	53%	55%	58%	60%
Capex of Hyperscalers to NVIDIA Revenue						49.309	85.661	99.470	108.709	116.282	127.674
% of NVIDIA Market Share						88%	88%	86%	84%	82%	80%
Top 5 Hyperscalers as 42.35%						116.433	202.269	243.501	276.262	307.217	351.234
Hyperscalers Share of NVIDIA Revenue (%)		-1%				42%	42%	41%	39%	38%	36%
Actual Data Center Revenue						115.186					
NVIDIA Data Center Revenue							202.269	243.501	276.262	307.217	351.234

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Nvidia's Data Center Revenue											
Data Center Revenue	2.983	6.696	10.613	15.005	47.525	115.186	202.269	243.501	276.262	307.217	351.234
YoY (%)		124%	58%	41%	217%	142%	76%	20%	13%	11%	14%
Compute	2.983	5.065	7.793	11.317	38.950	102.196	179.459	216.040	245.107	272.571	311.624
YoY (%)		70%	54%	45%	244%	162%	76%	20%	13%	11%	14%
% of Data Center Revenue	100%	76%	73%	75%	82%	89%	89%	89%	89%	89%	89%
Networking	-	1.631	2.820	3.688	8.575	12.990	22.811	27.461	31.155	34.646	39.610
YoY (%)			73%	31%	133%	51%	-	-	-	-	-
% of Compute Revenue		24%	27%	25%	18%	11%	11%	-	-	-	-



Revenue Build Up (Others)

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Nvidia's Gaming Revenue											
Gaming Revenue	5.518	7.759	12.462	9.067	10.447	11.350	13.717	16.021	17.241	19.785	22.705
YoY (%)		41%	61%	-27%	15%	9%	21%	17%	8%	15%	15%
PC Gaming	5.293	7.573	12.462	8.854	10.181	10.951	13.168	15.233	16.133	18.210	20.478
YoY (%)		43%	65%	-29%	15%	8%	20%	16%	6%	13%	12%
% of Gaming Revenue	96%	98%	100%	98%	97%	96%	20%	16%	6%	13%	12%
Console Gaming (Tegra)	225	186	-	213	266	399	548	788	1.108	1.575	2.227
YoY (%)		-17%	-100%	-	25%	50%	37%	44%	41%	42%	41%
% of Gaming Revenue	4%	2%	0%	2%	3%	4%	37%	44%	41%	42%	41%
Nvidia's Others Revenues											
Professional Visualization	1.212	1.053	2.111	1.544	1.553	1.878	1.972	2.090	2.237	2.415	2.633
YoY (%)		-13%	100%	-27%	1%	21%	5%	6%	7%	8%	9%
Automotive	700	536	566	903	1.091	1.694	2.460	3.523	4.975	6.925	9.502
YoY (%)		-23%	6%	60%	21%	55%	45%	43%	41%	39%	37%
OEM	505	631	1.162	455	306	389	401	413	425	438	451
YoY (%)		25%	84%	-61%	-33%	27%	3%	3%	3%	3%	3%



Income Statement Model

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Income Statement											
Revenue	10.918	16.675	26.914	26.974	60.922	130.497	220.818	265.548	301.139	336.781	386.525
<i>YoY (%)</i>		53%	61%	0%	126%	114%	69%	20%	13%	12%	15%
Cost of revenue	(4.150)	(6.279)	(9.439)	(11.618)	(16.621)	(32.639)	(55.121)	(66.287)	(75.171)	(84.068)	(96.485)
Gross profit	6.768	10.396	17.475	15.356	44.301	97.858	165.697	199.261	225.968	252.713	290.039
<i>Gross Margin</i>	62%	62%	65%	57%	73%	75%	75,0%	75,0%	75,0%	75,0%	75,0%
Amortized Intangible Assets	26	612	563	699	614	563	555	354	236	84	31
PP&E Amortizaion	355	486	611	844	894	892	1.685	1.826	2.048	2.307	2.591
Total D&A	381	1.098	1.174	1.543	1.508	1.455	2.240	2.180	2.284	2.391	2.622
EBITDA	3.227	5.630	11.215	5.767	34.480	82.908	141.439	169.575	192.115	214.689	246.278
<i>EBITDA Margin</i>	30%	34%	42%	21%	57%	64%	64%	64%	64%	64%	64%
Operating expenses	(3.922)	(5.864)	(7.434)	(11.132)	(11.329)	(16.405)	(26.498)	(31.866)	(36.137)	(40.414)	(46.383)
Research and development	2.829	3.924	5.268	7.339	8.675	12.914	20.978	25.227	28.608	31.994	36.720
<i>As a % of Revenue</i>	26%	24%	20%	27%	14%	10%	10%				
Sales, general and administrative	1.093	1.940	2.166	2.440	2.654	3.491	5.520	6.639	7.528	8.420	9.663
<i>As a % of Revenue</i>	10%	12%	8%	9%	4%	3%	3%				
Acquisition termination cost	-	-	-	1.353	-	-	-	-	-	-	-
<i>As a % of Revenue</i>	0%	0%	0%	5%	0%	0%	0%				
Operating income	2.846	4.532	10.041	4.224	32.972	81.453	139.199	167.395	189.831	212.299	243.656
<i>EBIT Margin</i>	26%	27%	37%	16%	54%	62%	63%	63%	63%	63%	63%
Interest income	178	57	29	267	866	1.786	1.862	1.921	2.419	3.220	4.166
Interest expense	(52)	(184)	(236)	(262)	(257)	(247)	(489)	(489)	(489)	(489)	(489)
Other, net	(2)	4	107	(48)	237	1.034	-	-	-	-	-
Income before income tax	2.970	4.409	9.941	4.181	33.818	84.026	140.572	168.827	191.761	215.030	247.333
<i>EBT Margin</i>	27%	26%	37%	16%	56%	64%	64%	64%	64%	64%	64%
Income tax expense	(174)	(77)	(189)	187	(4.058)	(11.146)	(23.194)	(27.856)	(31.641)	(35.480)	(40.810)
<i>Effective tax rate</i>	6%	2%	2%	4%	12%	13%	17%				
Net income	2.796	4.332	9.752	4.368	29.760	72.880	117.378	140.970	160.120	179.550	206.523
<i>Net Margin</i>	26%	26%	36%	16%	49%	56%	53%	53%	53%	53%	53%
<i>YoY (%)</i>		55%	125%	-55%	581%	145%	61%	20%	14%	12%	15%



Balance Sheet Model

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Balance Sheet											
Current assets	13.690	16.055	28.829	23.073	44.345	80.126	105.437	131.507	158.390	187.762	223.093
Cash and cash equivalents	10.896	847	1.990	3.389	7.280	8.589	22.712	39.803	59.540	81.756	107.101
Marketable securities	1	10.714	19.218	9.907	18.704	34.621	34.621	34.621	34.621	34.621	34.621
Accounts receivable, net	1.657	2.429	4.650	3.827	9.999	23.065	31.329	37.675	42.725	47.782	54.839
Inventories	979	1.826	2.605	5.159	5.282	10.080	13.003	15.637	17.733	19.832	22.761
Prepaid expenses and other current assets	157	239	366	791	3.080	3.771	3.771	3.771	3.771	3.771	3.771
Non-Current assets	3.625	12.736	15.358	18.109	21.383	31.475	32.009	33.167	34.666	36.508	38.743
Property and equipment, net	1.674	2.149	2.778	3.807	3.914	6.283	6.806	7.636	8.599	9.660	10.934
Operating lease assets	618	707	829	1.038	1.346	1.793	1.793	1.793	1.793	1.793	1.793
Goodwill	618	4.193	4.349	4.372	4.430	5.188	5.188	5.188	5.188	5.188	5.188
Intangible assets, net	49	2.737	2.339	1.676	1.112	807	819	1.146	1.683	2.463	3.423
Deferred income tax assets	548	806	1.222	3.396	6.081	10.979	10.979	10.979	10.979	10.979	10.979
Other assets	118	2.144	3.841	3.820	4.500	6.425	6.425	6.425	6.425	6.425	6.425
Total assets	17.315	28.791	44.187	41.182	65.728	111.601	137.446	164.674	193.057	224.269	261.836
Current liabilities:	1.784	3.925	4.335	6.563	10.631	18.047	22.393	24.552	26.270	27.990	30.390
Accounts payable	687	1.149	1.783	1.193	2.699	6.310	10.656	12.815	14.533	16.253	18.653
Accrued and other current liabilities	1.097	1.777	2.552	4.120	6.682	11.737	11.737	11.737	11.737	11.737	11.737
Short-term debt	-	999	-	1.250	1.250	-	-	-	-	-	-
Long-term liabilities	3.327	7.973	13.240	12.518	12.119	14.227	14.227	14.227	14.227	14.227	14.227
Long-term debt	1.991	5.964	10.946	9.703	8.459	8.463	8.463	8.463	8.463	8.463	8.463
Long-term operating lease liabilities	561	634	741	902	1.119	1.519	1.519	1.519	1.519	1.519	1.519
Other long-term liabilities	775	1.375	1.553	1.913	2.541	4.245	4.245	4.245	4.245	4.245	4.245
Total Liabilities	5.111	11.898	17.575	19.081	22.750	32.274	36.620	38.779	40.497	42.217	44.617
Total shareholders' equity	12.204	16.893	26.612	22.101	42.978	79.327	100.826	125.895	152.560	182.053	217.219
Preferred stock	-	-	-	-	-	-	-	-	-	-	-
Common stock	1	3	3	2	25	24	24	24	24	24	24
Additional paid-in capital	7.045	8.719	10.385	11.971	13.109	11.237	11.237	11.237	11.237	11.237	11.237
Accumulated other comprehensive income (loss)	1	19	(11)	(43)	27	28	28	28	28	28	28
Treasury stock, at cost	(9.814)	(10.756)	-	-	-	-	-	-	-	-	-
Retained earnings	14.971	18.908	16.235	10.171	29.817	68.038	89.537	114.606	141.271	170.764	205.930
Total liabilities and shareholders' equity	17.315	28.791	44.187	41.182	65.728	111.601	137.446	164.674	193.057	224.269	261.836
BS Check											
BS Check	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok



Cash Flow Model

Unit: US\$ Milion	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Cash Flow											
Net Income							117.378	140.970	160.120	179.550	206.523
(+)D&A							2.240	2.180	2.284	2.391	2.622
(+/-)Δ Working Capital							(6.841)	(6.822)	(5.428)	(5.436)	(7.586)
CFO							112.777	136.329	156.977	176.505	201.559
(-) PP&E CapEx							(2.208)	(2.655)	(3.011)	(3.368)	(3.865)
(-) Acquired intangibles Capex							(567)	(681)	(773)	(864)	(992)
CFI							(2.775)	(3.337)	(3.784)	(4.232)	(4.857)
(-) Principal Amortization							-	-	-	-	-
(-) SBC Buyback Compensation in Dividends							(15.846)	(19.056)	(21.610)	(24.168)	(27.737)
(-) Buyback Dividends							(80.033)	(96.845)	(111.845)	(125.890)	(143.620)
CFF							(95.879)	(115.901)	(133.455)	(150.057)	(171.358)
Cash BoP							8.589	22.712	39.803	59.540	81.756
Δ Cash							14.123	17.090	19.737	22.216	25.345
Cash EoP						8.589	22.712	39.803	59.540	81.756	107.101



Working Capital & Debt Model

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Working Capital											
Working Capital	1.949	3.106	5.472	7.793	12.582	26.835	33.676	40.498	45.925	51.361	58.947
Δ Working Capital		1.157	2.366	2.321	4.789	14.253	6.841	6.822	5.428	5.436	7.586
Days	365										
Working Capital - Assets	2.636	4.255	7.255	8.986	15.281	33.145	44.332	53.313	60.458	67.614	77.600
Accounts receivable	1.657	2.429	4.650	3.827	9.999	23.065	31.329	37.675	42.725	47.782	54.839
<i>As a days of revenue</i>	55	53	63	52	60	65	52				
Inventories	979	1.826	2.605	5.159	5.282	10.080	13.003	15.637	17.733	19.832	22.761
<i>As a days of cost of revenue</i>	86	106	101	162	116	113	86				
Working Capital - Liabilities	687	1.149	1.783	1.193	2.699	6.310	10.656	12.815	14.533	16.253	18.653
Accounts Payable	687	1.149	1.783	1.193	2.699	6.310	10.656	12.815	14.533	16.253	18.653
<i>As a days of cost of revenue</i>	60	67	69	37	59	71	71				
Debt											
Total Debt BOP							8.463	8.463	8.463	8.463	8.463
(+) New Debt							-	-	-	-	-
(-) Amortization							-	-	-	-	-
Total Debt EoP							8.463	8.463	8.463	8.463	8.463

We chose to maintain the same debt level, assuming the company is operating with an optimal capital structure. In this sense, we won't amortize the debt, because even if it were paid down in reality, the company would likely issue new debt



PP&E & Acquired Intangibles Model

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
PP&E											
PP&E	1.674	2.149	2.778	3.807	3.914	6.283	6.806	7.636	8.599	9.660	10.934
BoP		1.674	2.149	2.778	3.807	3.914	6.283	6.806	7.636	8.599	9.660
(+) Capex <i>As a % of revenue</i>		961 6%	1.240 5%	1.873 7%	1.001 2%	3.261 2%	2.208 1%	2.655	3.011	3.368	3.865
(-) Depreciation <i>As a % of PP&E</i>		486 29%	611 28%	844 30%	894 23%	892 23%	1.685 27%	1.826	2.048	2.307	2.591
EoP	1.674	2.149	2.778	3.807	3.914	6.283	6.806	7.636	8.599	9.660	10.934
Acquired Intangibles											
Acquired Intangibles	49	2.737	2.339	1.676	1.112	807	819	1.146	1.683	2.463	3.423
BoP		49	2.737	2.339	1.676	1.112	807	819	1.146	1.683	2.463
(+) Purchases of Intangible Assets and Investments <i>As a % of revenue</i>		3.300 20%	165 1%	36 0%	50 0%	258 0%	567 0,26%	681	773	864	992
(-) Amortization of Intangible Assets <i>As a % of PP&E</i>		612 1249%	563 21%	699 30%	614 37%	563 51%	555 69%	354 43%	236 21%	84 5%	31 1%
EoP	49	2.737	2.339	1.676	1.112	807	819	1.146	1.683	2.463	3.423



Others Model

Interest income, Interest expenses, SBC and Dividends

Unit: US\$ Million	2019	2020	2021	2022	2023	2024	2025E	2026E	2027E	2028E	2029E
Interest Income											
Interest Income							1.862	1.921	2.419	3.220	4.166
Cash BoP							43.210	57.333	74.424	94.161	116.377
SOFR							4,3%	3,4%	3,3%	3,4%	3,6%
Interest Expense											
Interest Expense							489	489	489	489	489
Debt BoP							8.463	8.463	8.463	8.463	8.463
Cost of Debt (%)							5,8%	5,8%	5,8%	5,8%	5,8%
Stock-based Compensation											
Stock-based Compensation	844	1.397	2.004	2.710	3.549	4.737	15.846	19.056	21.610	24.168	27.737
% of Revenue	8%	8%	7%	10%	6%	4%	7,2%				
Dividends											
Dividends							80.033	96.845	111.845	125.890	143.620
Cash Flow Before Dividends							94.156	113.936	131.583	148.106	168.965
Payout (%)							85,0%				
% of Net Income							68%	69%	70%	70%	70%
Share Buybacks as % of Net Income	14%	9%	4%	239%	33%	47%					
Share Buybacks + Dividends	(390)	(395)	(399)	(10.437)	(9.928)	(34.540)					
Net Income	2.796	4.332	9.752	4.368	29.760	72.880					

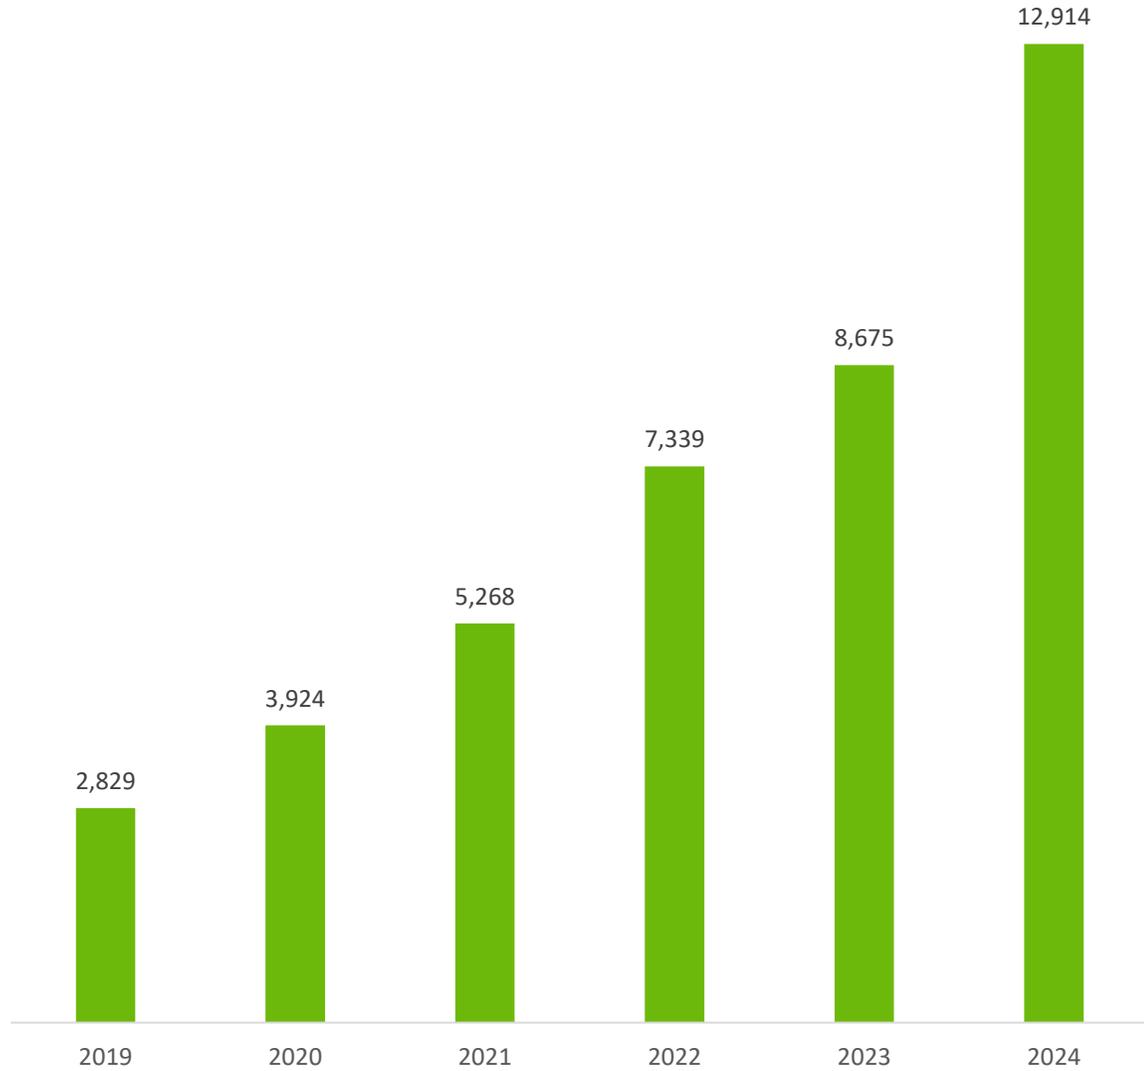
Given its already strong cash position, we chose to distribute dividends as a way to replace shareholder returns that would otherwise come from share buybacks



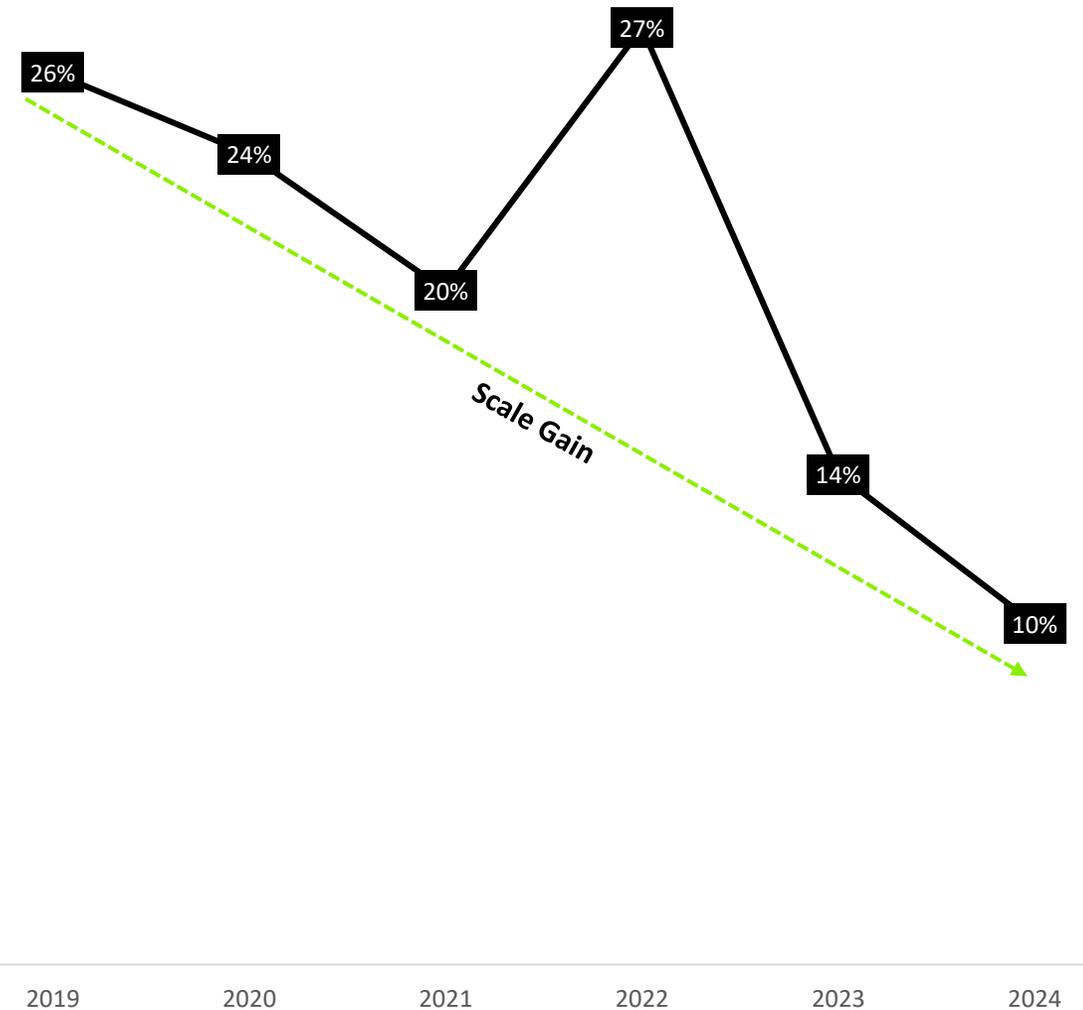
R&D Expenses

R&D dilution: scale is rising faster than innovation investment

R&D Expenses (USD millions)



R&D as a % of Revenue



ROIC

Unit: US\$ Milion		2019	2020	2021	2022	2023	2024
Amortization years		5					
R&D of the period							
	2019	2.829	3.924	5.268	7.339	8.675	12.914
	2020		566	566	566	566	566
	2021			785	785	785	785
	2022				1.054	1.054	1.054
	2023					1.468	1.468
	2024						1.735
R&D acumulation		2.829	6.187	10.105	15.039	19.842	27.149
Adjusted Nopat		5.343	7.752	13.693	9.568	33.242	76.986
Adjusted EBIT		5.675	7.890	13.958	9.159	37.775	88.760
EBIT		2.846	4.532	10.041	4.224	32.972	81.453
R&D		2.829	3.924	5.268	7.339	8.675	12.914
R&D Amortization		-	566	1.351	2.404	3.872	5.607
Effective Tax Rate		-6%	-2%	-2%	4%	-12%	-13%
Adjusted Invested capital		6.797	17.541	23.686	30.396	39.624	60.089
Adjusted ROIC		79%	44%	58%	31%	84%	128%
Adjusted ROIC (Ex-Goodwill)		86%	58%	71%	37%	94%	140%
NOPAT Margin							0,59
IC Turnover							2,17



Phase 2 | Presentation

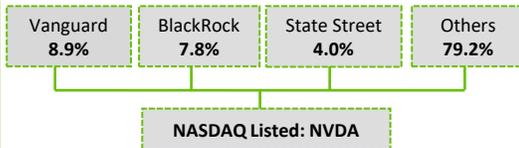
Nvidia in a Circuit

Leadership and Technology in Semiconductor Design



Nvidia is a global technology company known for creating powerful graphics processors and tools that help computers run faster and smarter. Founded in 1993, it started by making graphics cards for gaming but has since expanded into areas like artificial intelligence, self-driving cars, and cloud computing. Nvidia doesn't manufacture its own chips; instead, it designs them and relies on specialized factories, like TSMC, to produce the hardware.

Shareholder Structure



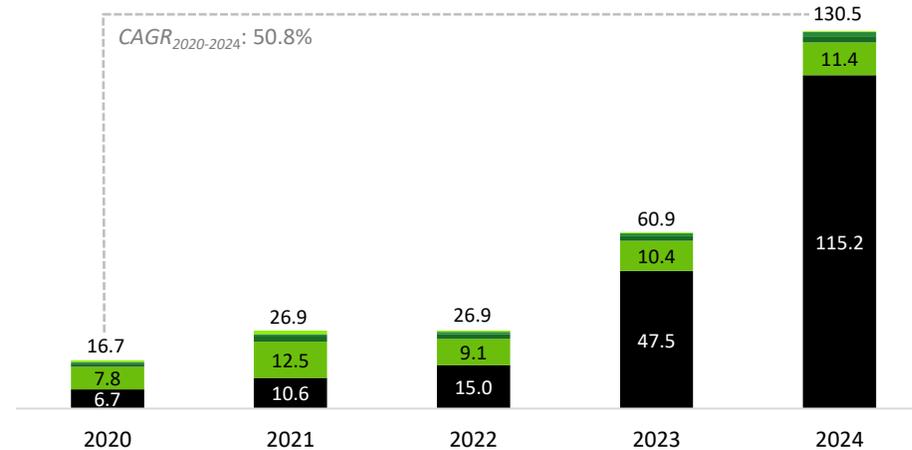
Key Indicators

Key Indicators	USD billions
Market Cap	3,415
Revenue	130.4
EBIT	81.4
Net Profit	72.8
Cash	8.5

Revenue by end-market (US\$ billion)

Historically, Nvidia's revenue was dominated by the gaming segment, but with the rise of artificial intelligence, the Data Center has gained prominence

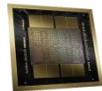
■ Data Center ■ Gaming ■ Professional Visualization ■ Automotive ■ OEM & Other



Nvidia main products

NVIDIA's high-end GPUs reflect a bold bet on AI dominance, extreme performance, and premium pricing, shaping an increasingly segmented computing market

Nvidia B200: US\$30,000 – US\$40,000



The NVIDIA B200 is a next-gen Blackwell GPU designed for generative AI, offering up to 1.4 PFLOPs, FP4 support, 1.8 TB/s NVLink bandwidth, and high energy efficiency for large-scale AI and HPC

Nvidia H200: ~ US\$30,000



The NVIDIA H200 is a Hopper-based GPU with 141 GB of HBM3e and 4.8 TB/s bandwidth, designed for generative AI and HPC, offering high performance and energy efficiency

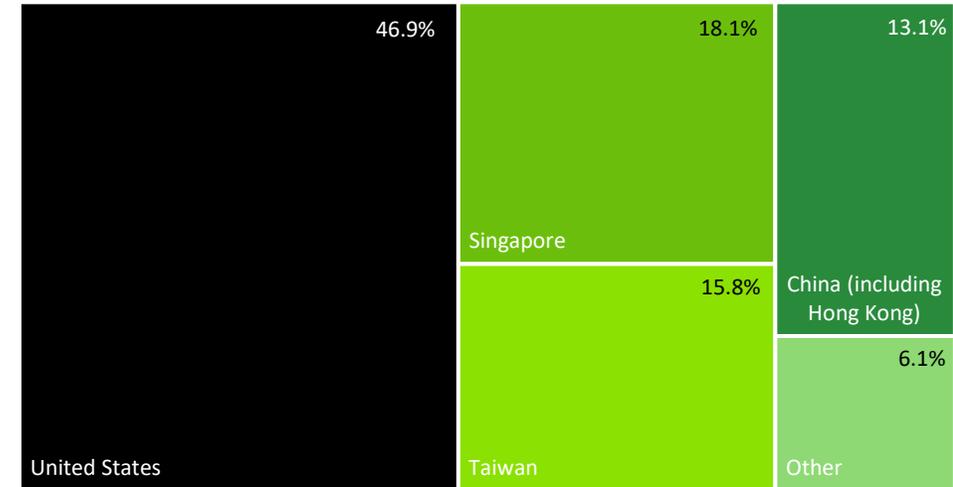
GeForce RTX 5090: US\$1,999



The GeForce RTX 5090 is NVIDIA's top consumer GPU, featuring Blackwell architecture, 21,760 CUDA cores, and 32 GB of GDDR7. It excels in gaming and AI with DLSS 4, ray tracing, and 8K support

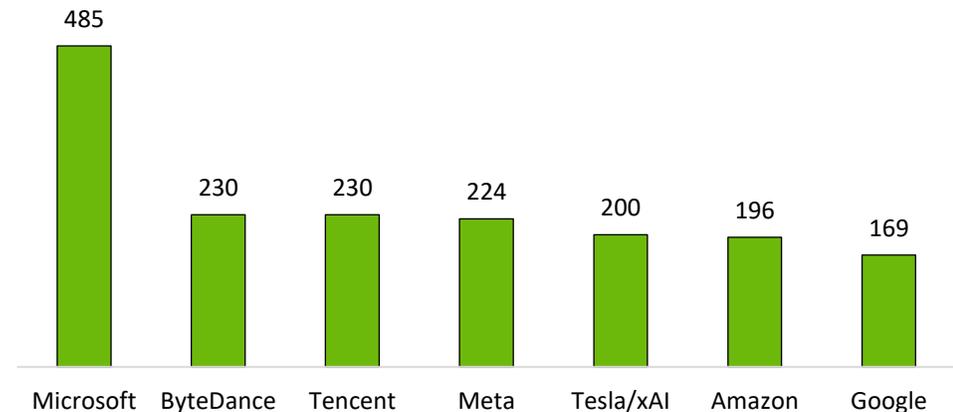
Geographic Revenue in 2024 (%)

With a strong presence in the U.S., NVIDIA also operates in Singapore and Taiwan but has been losing ground in China due to trade restrictions



Estimated shipments of Nvidia Hopper GPUs in 2024, by customer

In 2024, Microsoft purchased twice as many Nvidia Hopper GPUs as ByteDance and Tencent, leading AI investments and becoming Nvidia's top customer



Unlocking a Technological Revolution

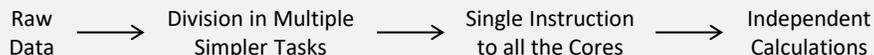
GPUs gained traction due to their scalable architecture, massive parallelism, and flexibility across applications



Why were GPUs so disruptive?

GPUs, with their **massively higher core counts**, enabled the **parallel execution of simpler, repetitive operations**—such as additions—across **large datasets**, offloading compute-bound workloads and allowing CPUs to focus on complex, interdependent tasks

Parallel Processing:



The simplicity and independence of GPU tasks make the **architecture inherently scalable**, as performance scales nearly linearly with core count

Massive Operation Scale:

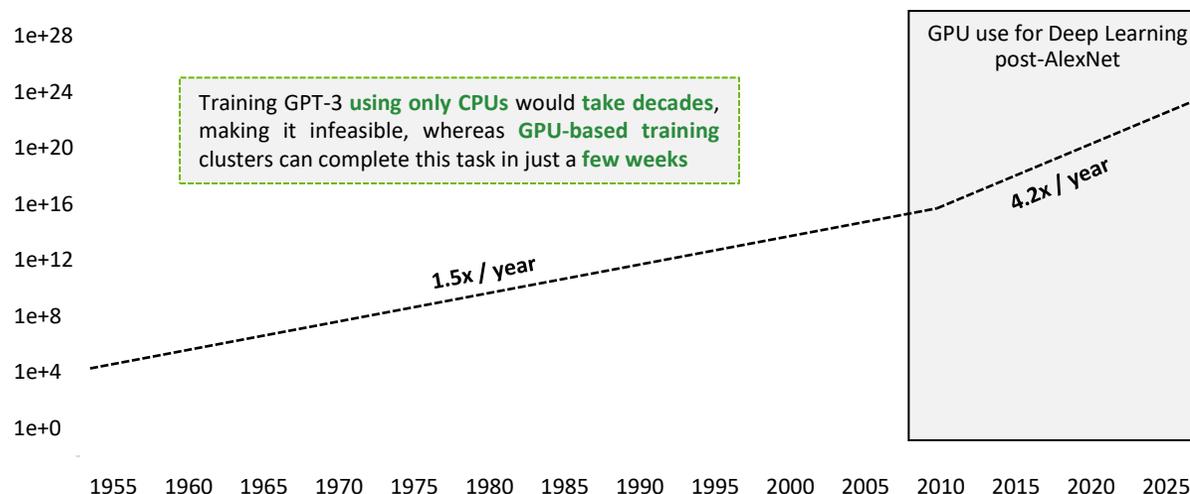
30 **trillion** Multiply-Add Operations per Second

$$FMA = a \times b + c$$

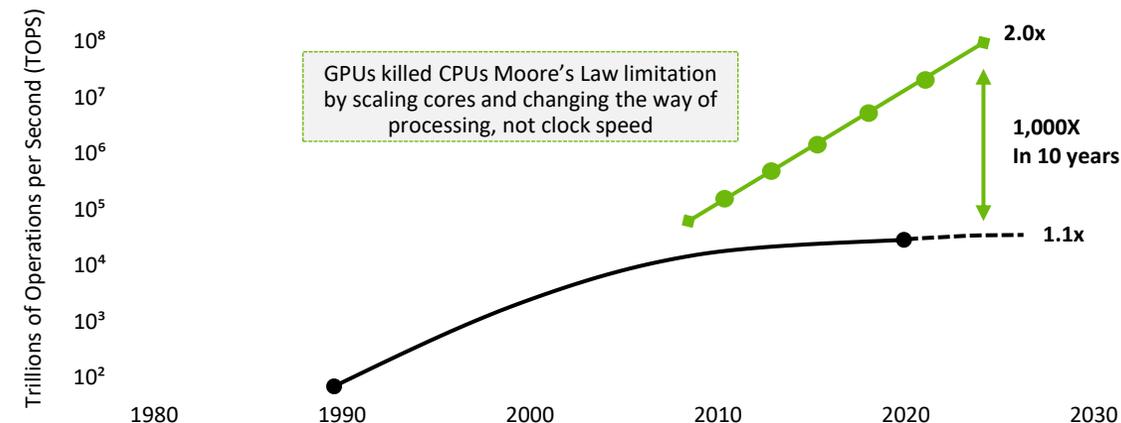
Product	Processing	Cores	Tasks	Strength
CPUs	Sequential	~96	Interdependent	Operating System
GPUs	Parallel	~21,760	Separate	Vectorized

Training Compute of Notable Machine Learning Systems Over Time (FLOP)

After the discovery of AlexNet, GPUs sparked the deep learning revolution



CPU performance advanced for decades under Moore's Law, driven by rising transistor counts and clock speeds. Over time, however, **gains plateaued due to physical constraints**—such as heat dissipation, quantum effects at nanoscales, and energy inefficiency. These limitations made **traditional CPU scaling unsustainable**, accelerating the shift toward alternative architectures like GPUs



Why Are GPUs remaining at the core of Technological Breakthroughs?

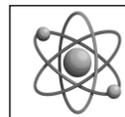
$$\text{Parallelism} + \text{General-purpose Architecture} = \text{Multiple end markets served}$$

Various Applications:



Data Center

Purpose-built to handle massive AI and scientific workloads, enabling high-throughput, low-latency compute at scale.



Professional Visualization

Real-time rendering and physics-accurate simulation, tailored for complex content creation. Formally entered this space with GPUs in the early 2000s, expanding rapidly from 2018 with Omniverse and RTX



Gaming

Engineered for high frame rates, and rich visual effects, ideal for interactive entertainment. Gaming has been its core since the late 1990s, evolving from basic graphics to AI-enhanced experiences

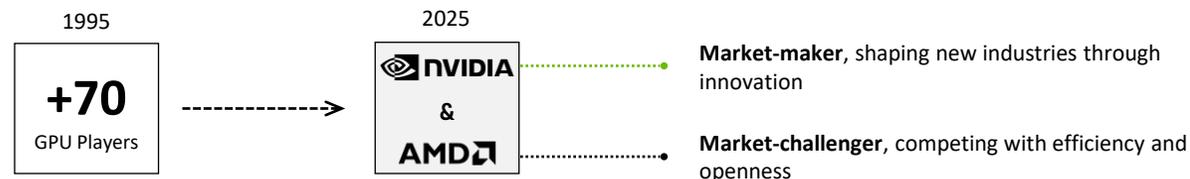


Three Steps Ahead: Unbothered by Competition

Consistently delivering superior results, reflecting deep expertise and the strength of the ecosystem built around GPUs

What about the competitors in this market?

Though rivals, NVIDIA focused early on building a GPU-centric ecosystem, while AMD spread across broader areas like CPUs

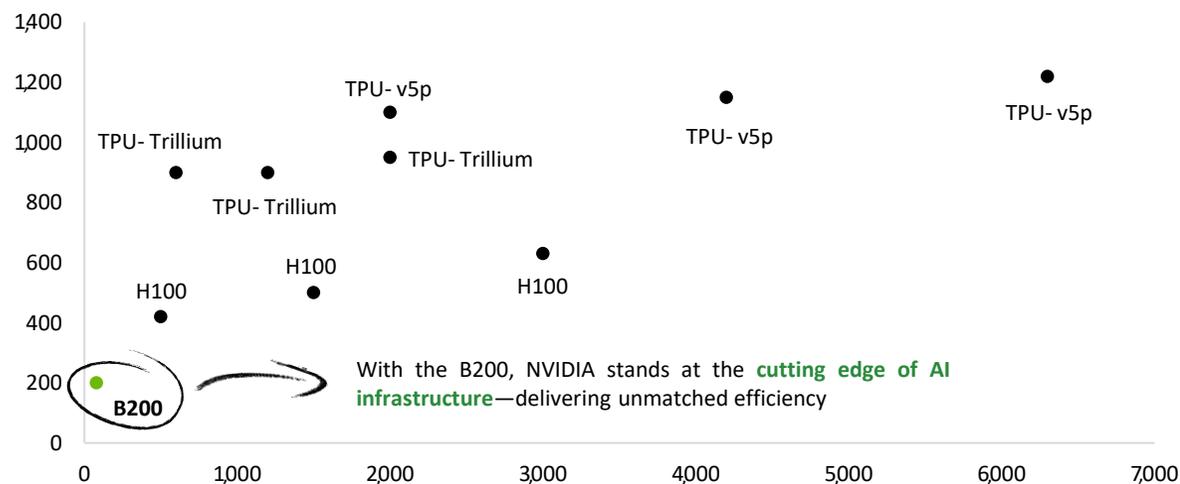


Despite AMD's efforts to enter the AI GPU market, there remains a significant performance gap in its flagship products — a gap NVIDIA has filled through its **ability to anticipate industry shifts**

Product	TT ₁	BW ₂	HBM Cap ₃	Strength
B200 (NVIDIA)	4.5 PFLOPS	8TB/s	192GB	AI Dominant
MI325X (AMD)	2.6 PFLOPS	6TB/s	256 GB	Abundant Memory

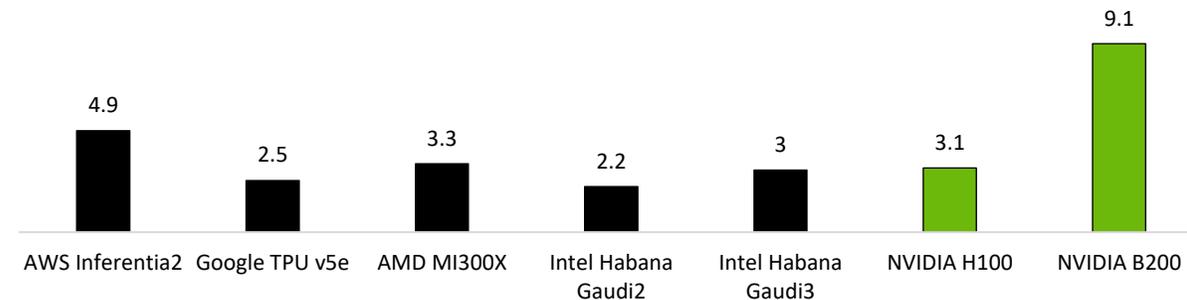
GPT-3 Training Results (GPU Hours Required x Number of GPUs Used)

Training LLM's becomes dramatically more efficient as each hardware generation unlocks more performance



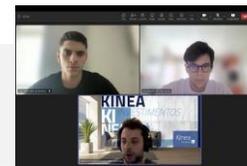
Chip level cost-performance ratio (Performance/Cost)

NVIDIA's B200 redefines cost-performance efficiency, nearly tripling competitors' metrics and highlighting how far rivals lag behind in delivering value at scale



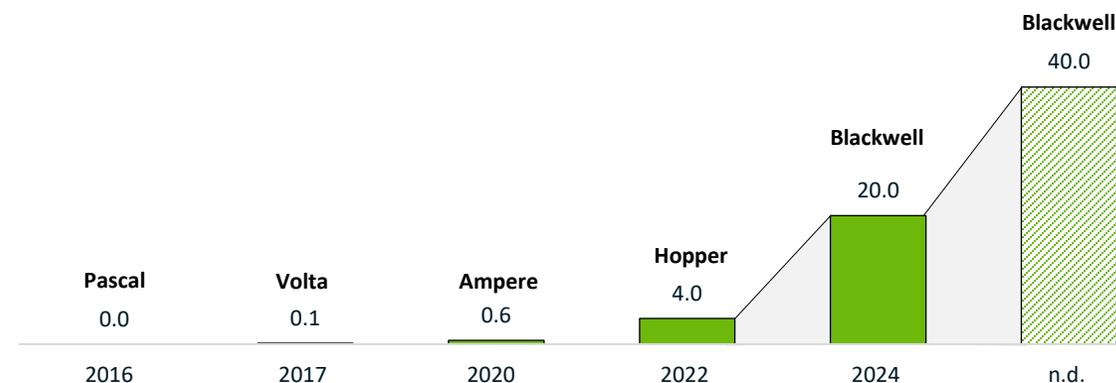
Guilherme Amaral, Kinea TMT Analyst

"NVIDIA has delivered higher performance by **capturing value-added layers** around the GPU business—not by full vertical integration, but by selectively controlling strategic components like CUDA and NVLink"



Evolution of NVIDIA GPU Processing Power (PFLOPs)

Unprecedented pace of performance scaling, with performance metrics results growing exponentially



NVIDIA consistently breaks through performance barriers, avoiding plateaus thanks to its unmatched pace of innovation. Through deep architectural redesigns and ecosystem control, each GPU generation brings exponential gains. This trajectory highlights NVIDIA's unique ability to reinvent and scale computing performance.



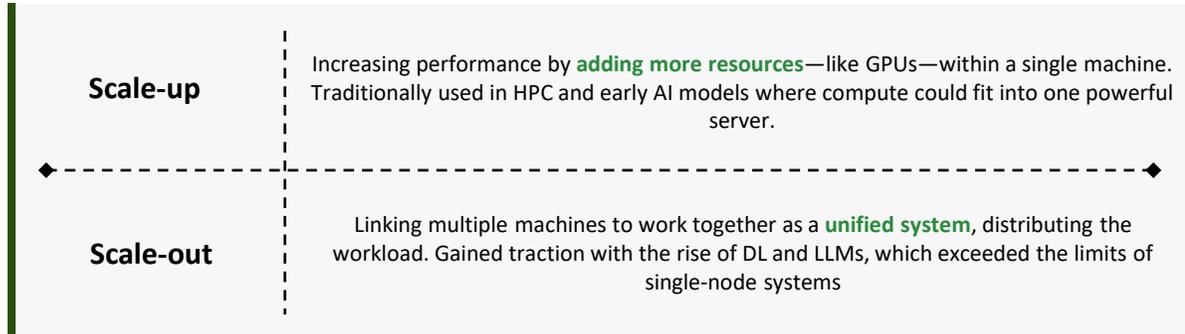
Scaling was Broken: NVLink Fixed It

NVLink unlocked the full potential of NVIDIA's GPU ecosystem, enabling seamless scalability beyond conventional limits

NVIDIA established its GPUs as top performers and scaled rapidly through sheer volume. However, to sustain performance gains as GPU clusters grew larger, new technologies were essential. This need for high-speed, low-latency communication across multiple GPUs led to the development of NVLink:

Why Data Centers have so many GPUs?

Due to GPUs parallelism capabilities it has two ways of improving its performance:



With the creation of NVLink, NVIDIA not only scaled the number of GPUs and boosted performance, but also unlocked **new product architectures** and **deeper integration across its ecosystem**—especially with CUDA, enabling seamless multi-GPU computing and more efficient software-hardware synergy

Expanding Possibilities for GPU Usage:



GB200
1 Grace GPU
2 Blackwell GPUs
Connected by NVLink-C2C
Unified Memory



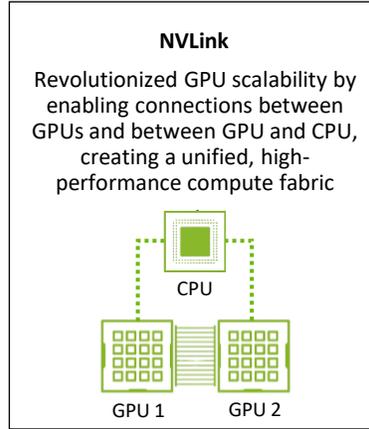
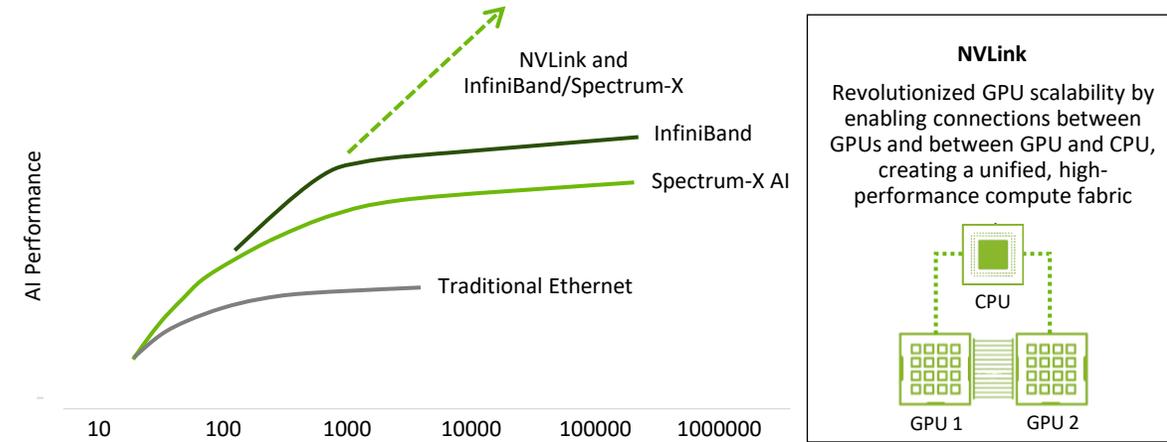
GB200 NVL72
36x GB200 Superchips
72 GPUs + 36 CPUs
Connected by NVLink Switch System
Neural ultra-speed

Jensen Huang, Nvidia's Founder & CEO



"NVLink wasn't just a technical innovation — it was the turning point that transformed NVIDIA from a chipmaker into a **builder of supercomputers**. By connecting GPUs with unprecedented bandwidth, we created a new class of superchips capable of powering the factories of the future"

Scale-up has been key to sustaining the growth of GPU performance



How is it performing?

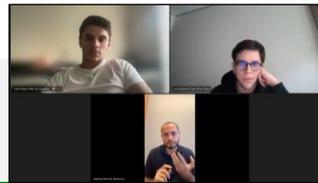
25% Reduction in Training Time

3X More throughput

NVIDIA's proprietary **NVLink** and **NVSwitch** technologies, reinforced by Mellanox's networking IP, provide ultra-low latency and high-bandwidth **GPU-to-GPU** and **node-to-node** interconnect. These capabilities enable efficient scaling of large AI and HPC workloads across thousands of GPUs

Gabriel Oliveira, Verde Global Equities Analyst

"NVIDIA has **NVLink and NVSwitch technologies**, which enable them to interconnect GPUs and nodes, delivering **faster performance than competitors**. A single NVLink can handle more data traffic than the entire global internet. **They make magic**"



CUDA: The Backbone of Nvidia's MOAT

Proprietary software has made it possible to program GPUs for specialized tasks, driving sustained performance gains over time

What made NVIDIA GPUs so flexible across End Markets?

With the launch of CUDA Software in 2006, GPUs evolved into more programmable machines, enabling the use of graphics cards across a wide range of applications and leveraging their parallel processing capabilities

Programmable GPUs = Flexibility for Applications + Performance Enhancing

Over nearly **two decades**, NVIDIA's software has been continuously refined and widely adopted by developers, creating a **high switching cost** due to deep-rooted academic training. This has fostered a strong developer base, reinforced by a **network effect**—more users mean more shared knowledge and institutional adoption

Proprietary Language (Only NVIDIA products)

+400 Libraries and +600 AI models

Optimization of GPUs constantly (even old ones)

Performance Enhancing

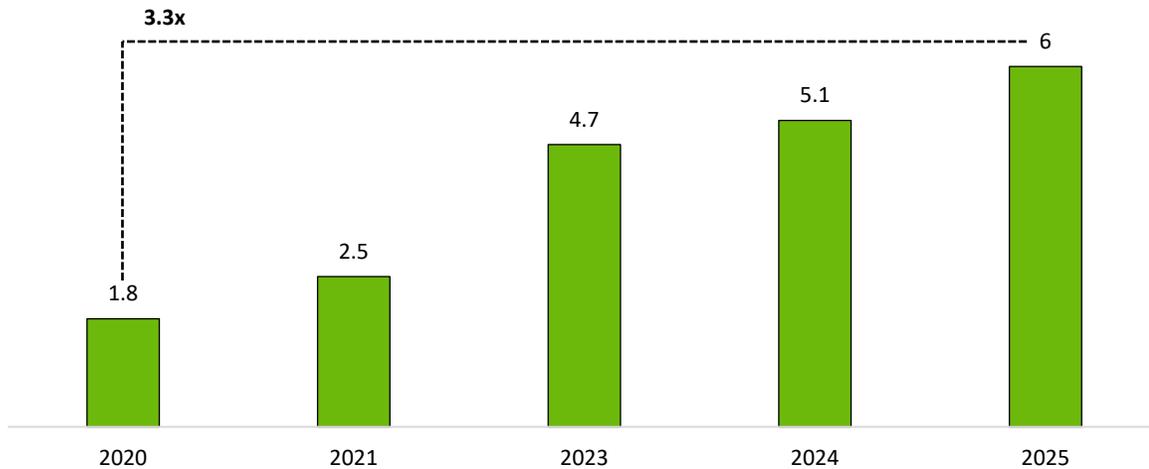
~200x Data Processing ~100x Deep Learning

~200x Computer Vision ~100x Agentic AI

~100x Science ~100x Recommender Systems

CUDA Developers (Million)

Beyond its accelerated growth, CUDA benefits from an already well-trained developer base, positioning it strongly for the years ahead



CUDA (NVIDIA) vs. ROCm (AMD): Measuring Content Presence on YouTube

We developed a Python script to quantify the volume of available content related to each software. Our findings indicate that CUDA is significantly more entrenched than ROCm, the AMD software

```
API_KEY = "AIzaSyAltVTXXGvd-bCcQajwYBvtLgwoEN0y-aI"

termos = ["CUDA", "ROCm"]

for termo in termos:
    print(f"\n== Buscando titulos contendo, {termo} (até 500 resultados) ==")
    titulos = buscar_todos_titulos_youtube(API_KEY, termo, max_por_termo=500)

    contagem =
    sum(1 for t in titulos
        if termo.lower() in t.lower())
    print(f"Total de titulos coletados {termo}: {len(titulos)}")
    print(f"Total de titulos que contém exatamente '{termo}': {contagem}\n")

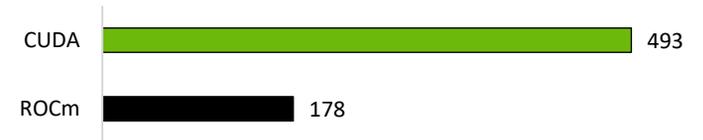
# (Opcional) Exibir os primeiros 20 titulos como amostra
print("Primeiros 20 titulos encontrados:")
for i, t in enumerate(titulos[:20], start=1):
    print(f"{i}. {t}")
```

Our results:

2.8x More Results for **CUDA**



The widespread availability of free content plays a crucial role in training new programmers with minimal effort



CUDA Downloads (Million)

The high number of CUDA downloads, combined with its non-transferable file formats, indicates a significant switching cost



The Role of Switching Costs, Yale

"In software markets, switching costs are often particularly high due to incompatibility of data formats, user retraining, and the need to rewrite custom code. These costs can deter customers from switching even if alternative products are available and better. In markets with high switching costs, a monopolist can continue to charge high prices or maintain market dominance long after the competitive advantage that initially attracted customers has eroded"



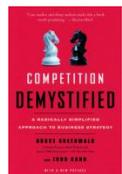
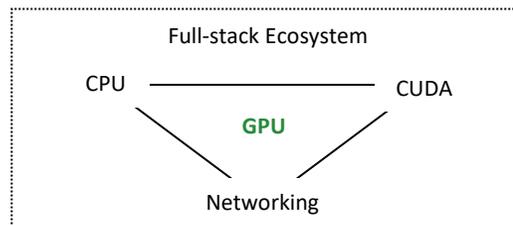
Stairway to Heaven: Winning Business Model? Check!

Building on deep technical moats, NVIDIA combines scale, lock-in, and execution to sustain market leadership

With CUDA as the backbone of its strategy, NVIDIA creates synergies across its entire product portfolio. This integration reinforces customer lock-in, resulting in market share stability and pricing power — a competitive edge continually reinforced by its unmatched architecture release cycle

Compiling Competitive Advantages:

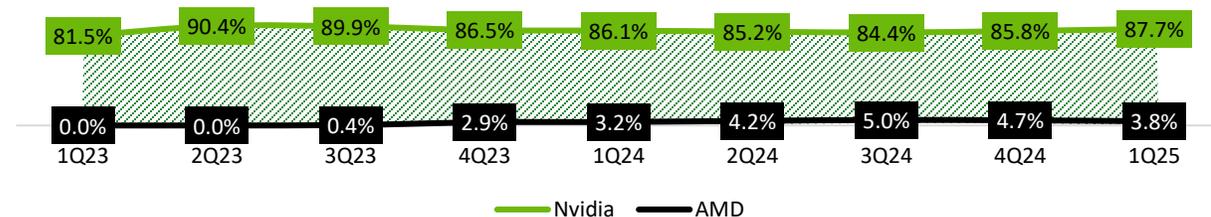
- ◆ Unmatched Technology Leadership
- ◆ Massive Scale Capability + General-purpose Architecture
- ◆ Switching Costs + Network Effect = Strong Customer Captivity



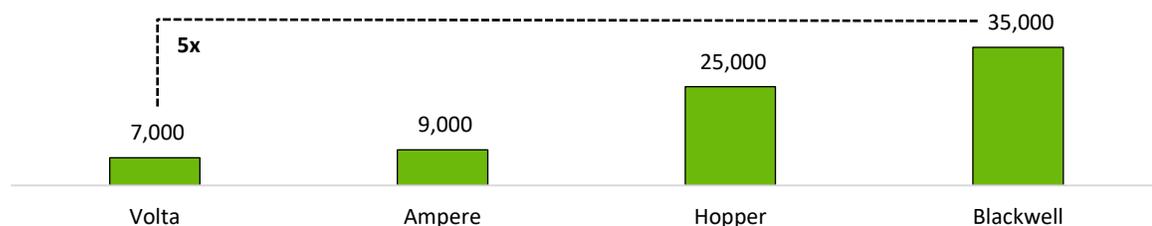
Competition Demystified, Bruce Greenwald

“The most powerful competitive advantages arise when customer captivity is combined with economies of scale. In such cases, a firm not only drives down its unit costs with volume but also makes it very hard for customers to leave, reinforcing its dominance over time.”

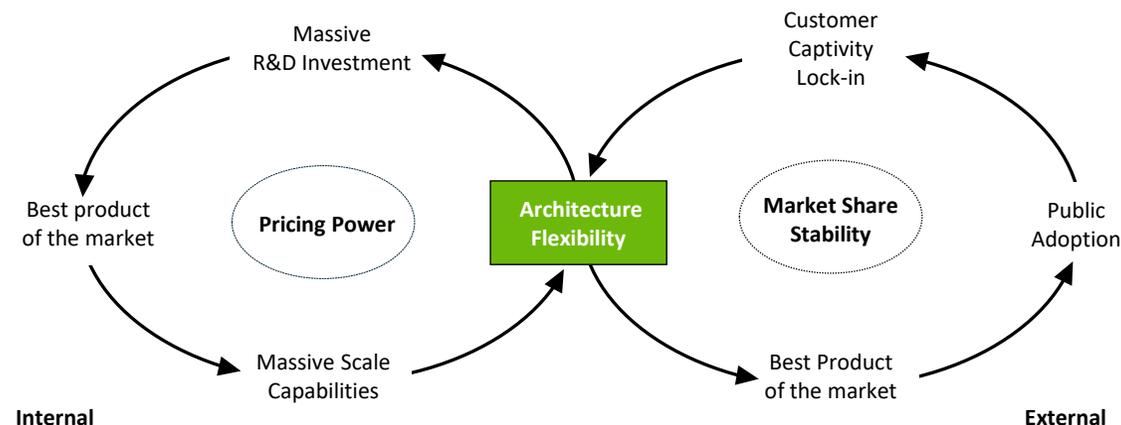
NVIDIA vs AMD Market Share of Data Center Accelerators(%)



NVIDIA GPU Pricing Power Across Architectures (US\$)



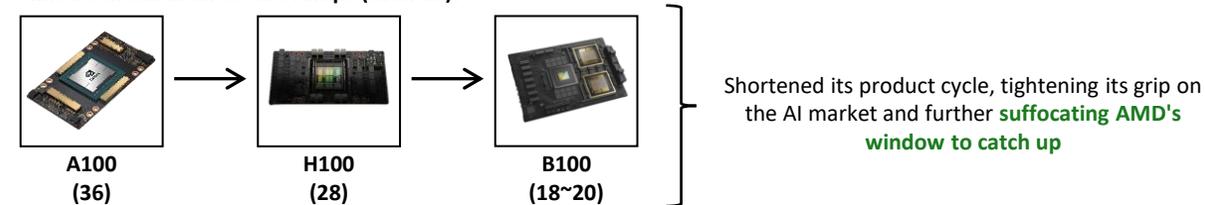
A self-reinforcing loop of winning outcomes:



Release cadence will make the difference larger and larger

By leveraging a unified architecture and reinforcing feedback loops across its ecosystem, NVIDIA has significantly accelerated its GPU launch cadence, widening its lead over competitors

Time between architecture ramps (Months):



As the inventor of the GPU, NVIDIA leverages unrivaled expertise, elite engineering talent, and visionary leadership under Jensen Huang. This deep-rooted strength enables faster architecture transitions and a tighter product cycle

Best Product

Barriers to Entry

Excellence in Execution



Envisioning the Future: Mission is the Boss

Vision backed by action: Jensen's decisions consistently anticipate where the market is heading, and get there first



Jensen Huang, Nvidia's Founder & CEO

"The technology industry doesn't reward the past, it only rewards the future. No matter how successful you were yesterday, if you don't innovate today, you become irrelevant. That's why at Nvidia, we wake up every morning as if we were running out of time"

Co-Founded Nvidia
Started Nvidia to focus on computing market

1993

TSMC Partnership
Early move to fabless model, ensuring scalable and advanced chip production

1998

The first modern GPU
GeForce 256 revolutionized graphics

1999

Created CUDA
Enabled GPUs for AI and scientific use

2006

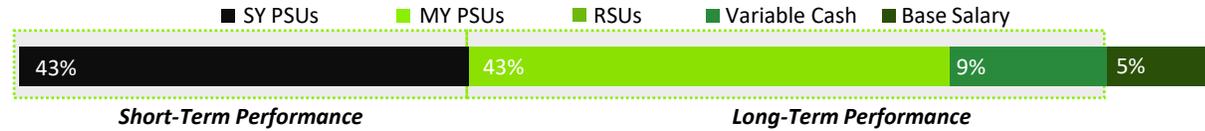
Mellanox Acquisition
Strengthened Nvidia's position in high-performance networking

2019

Launched Blackwell
Powered Nvidia's next-gen AI dominance globally

2024

CEO target pay mix 2024



Jensen's Visionary Decision Mentality:

Core Beliefs

+

Early Indicators of Future Success

=

Early Adopter of Disruptive End Markets

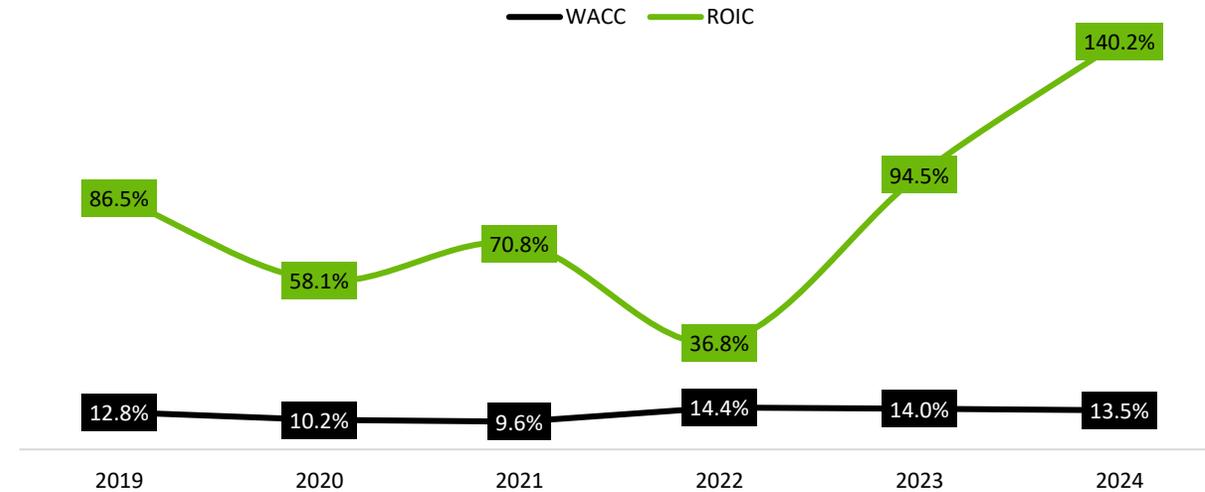
The Nvidia Way, Tae Kim

"Since Nvidia's founding, Jensen has insisted that all Nvidia employees work at the "Speed of Light". He wants their work to be constrained only by the laws of physics – not by internal politics or financial concerns [...] "Speed of the light gets you into the market faster and makes it really, really hard, if not impossible, for your competitors to do better", a former Nvidia executive said. "How fast can you do it, and why aren't you doing it that faster."



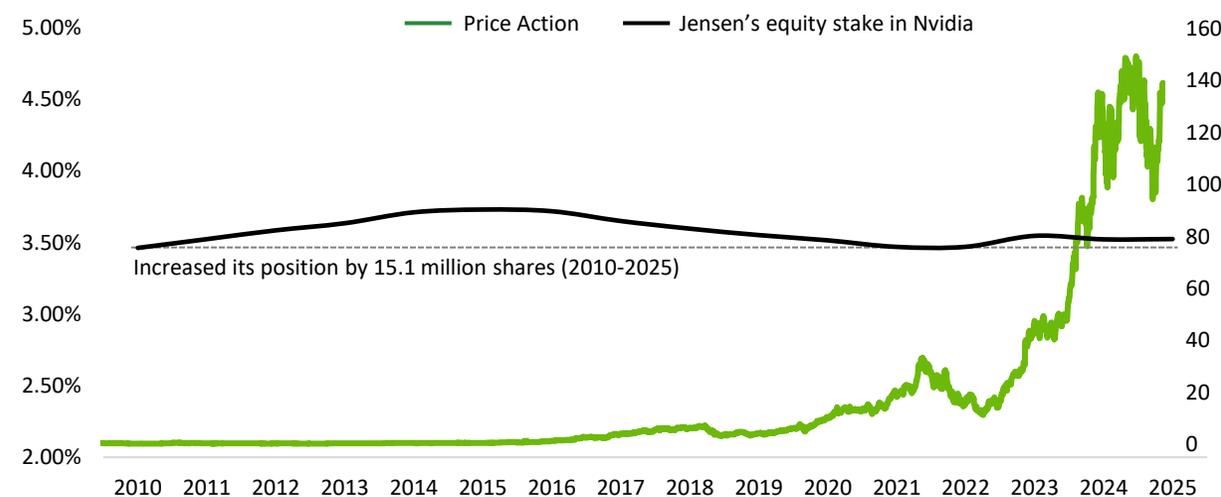
ROIC₁ x WACC (%)

Jensen's investment choices have proven highly efficient: Nvidia's ROIC comfortably exceeds its WACC, showing a return on capital well above the opportunity cost



Price Action (US\$) x Jensen Ownership (%)

Even after the stock has appreciated by more than 300,000% since its IPO, Jensen has maintained a substantial amount of his equity stake, and remains confident in Nvidia's long-term potential

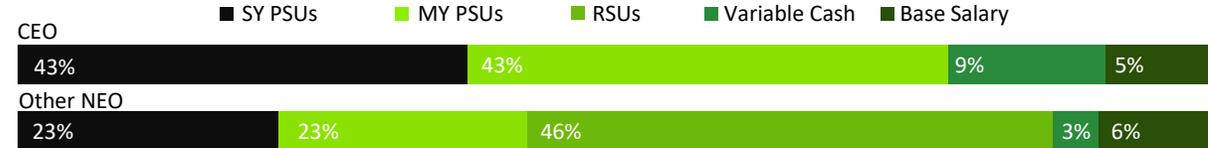


Where Execution Meets Consistency: Human capital advantage

How exceptional talent and a purpose-driven culture fuel technological leadership and innovation

CEO & Other NEO target pay mix 2024

Pay mix reflects a strong alignment between leadership incentives and the company's long-term, sustainable value creation goals



SY PSUs

Short-term focused, reward consistent achievement of key annual performance targets

MY PSUs

Long-term oriented, tied directly to multi-year goals, aligning with NVIDIA's strategic vision

RSUs

Long-term focused, encourage retention and align executives with shareholder value



Colette Kress
Nvidia CFO
Former CFO at Cisco
+13 Years
In-house



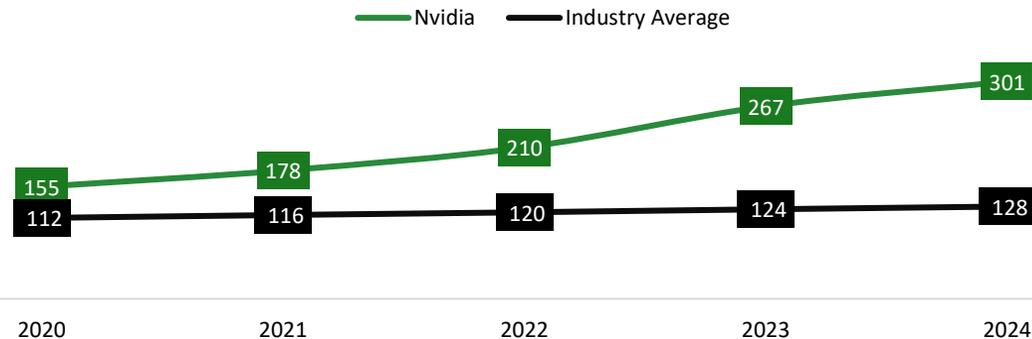
Bill Dally
CS₁ and SVP Research
Professor in Stanford
+16 Years
In-house



Ian Buck
VP at Data Center
Inventor of CUDA
+25 Years
In-house

NVIDIA Employee Pay vs Semiconductor Industry Average (US\$ Thousand)

But Nvidia is not built solely on strong executives. The company also relies on a highly specialized technical workforce, attracted by an above-average compensation policy



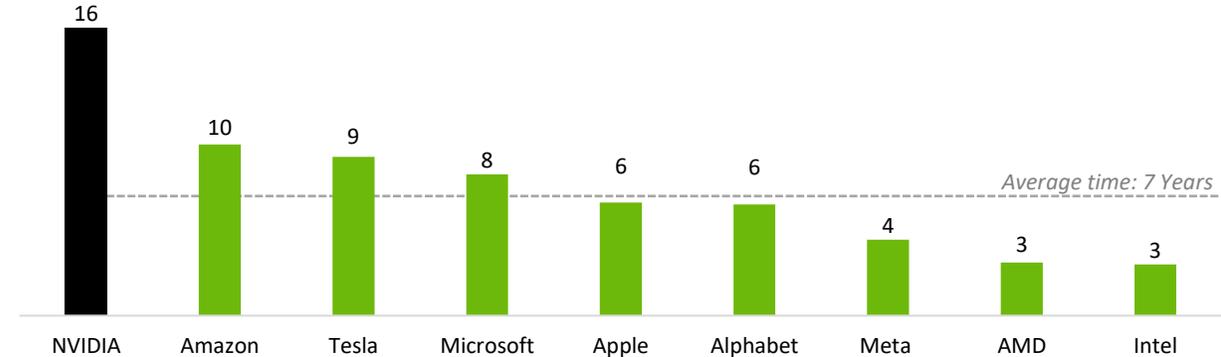
75% of employees are R&D
vs 59% of AMD

2.7% Turnover
vs 17.7% average in semiconductor industry

166:1 Pay Ratio₂
vs 222:1 compared to its peers₃

Average executive tenure (Years)

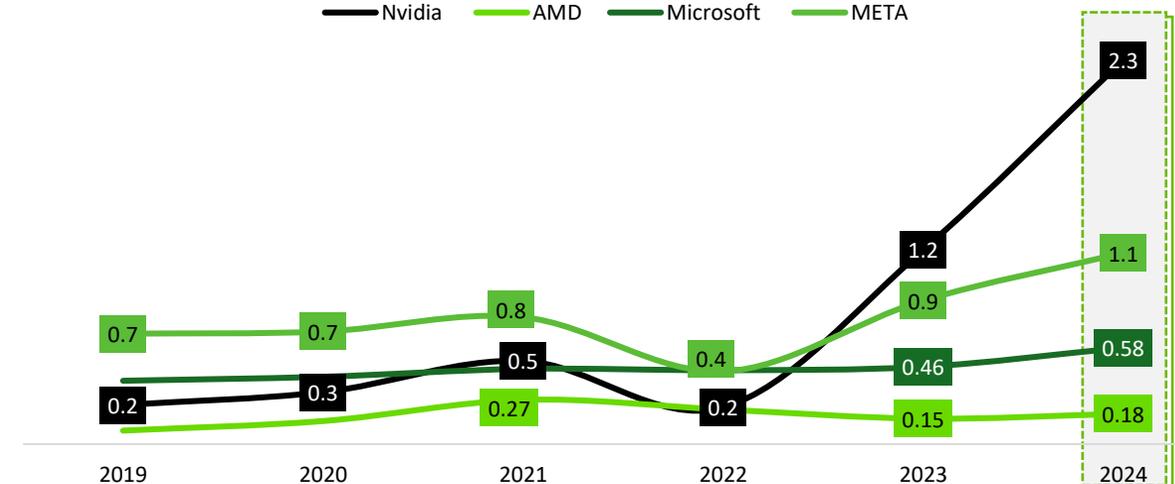
With an average executive tenure of 16 years, more than double the industry average, NVIDIA stands out for its leadership stability, enabling deeper expertise and long-term vision



With an average executive tenure of 16 years, **NVIDIA holds a clear competitive edge**. Long-standing leadership allows the company to move with greater strategic consistency, make better-informed decisions, and execute long-term plans more effectively than competitors like AMD and Intel, whose leadership turnover limits continuity and deep industry insight

EBITDA/Employee (US\$ millions)

NVIDIA's EBITDA per employee highlights a highly skilled and productive workforce, generating value at a scale far beyond industry peers



AI is the future and Nvidia is Best Positioned to Capture it

From talent to execution, NVIDIA's integrated model turns technical leadership into real-world market dominance

Complete Ecosystem

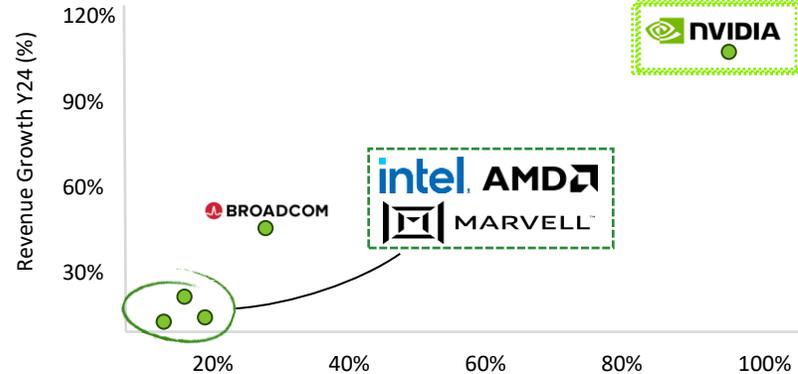
+

Specialized human capital

=

Positioning to capture future demand

Revenue growth (%) vs AI Exposure₁

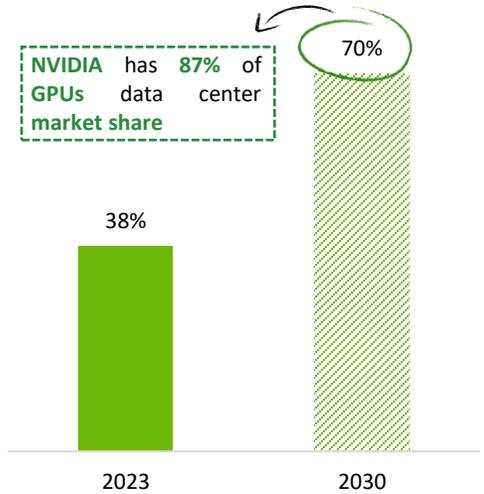


By building a robust and integrated ecosystem, guided by a well-defined long-term vision from its highly specialized talent and by positioning itself early in the AI market, Nvidia has established a strong competitive moat and, in practice, a monopoly in AI-focused GPUs, making it extremely difficult for new players to enter this segment

Demand for advanced-AI capacity (% of total data center capacity demand)

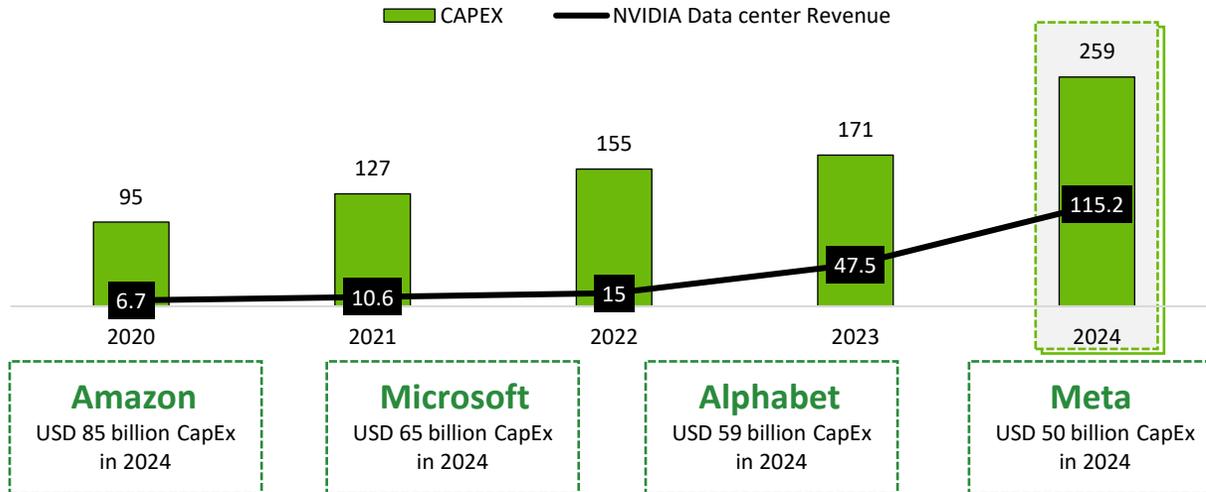
AI demand is accelerating and exposing the gap between those merely following the trend and those ready to lead it, like NVIDIA

Datacenter Type	AI	GPU Demand
HPC	X	High
Cloud / Hyperscaler	X	Very High
AI Accelerated	X	Very High
Colocation		Low
Enterprise		Low



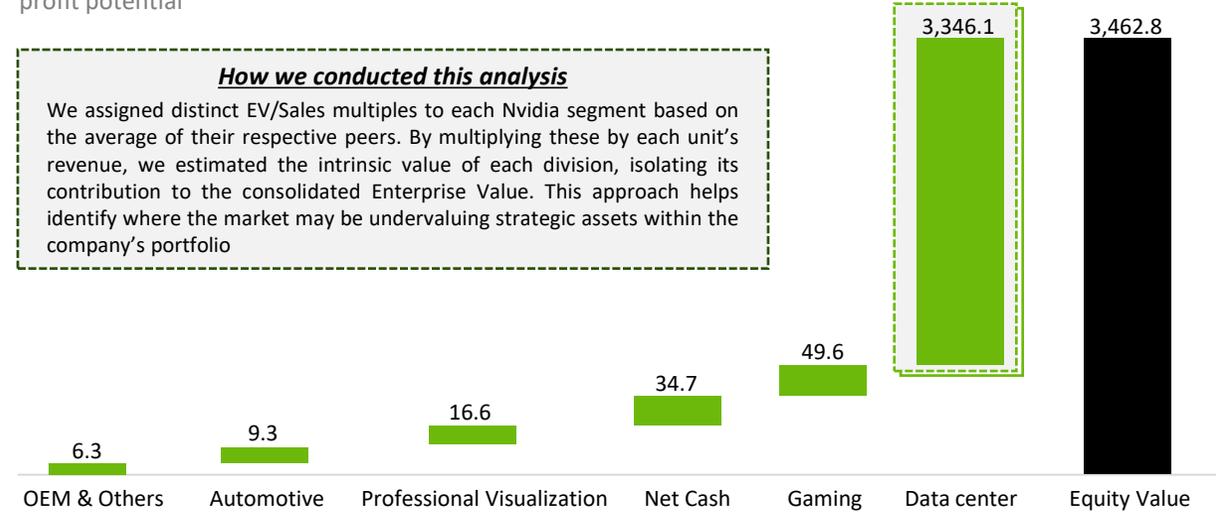
NVIDIA Data Center Revenue x Big 4 Tech Companies Capex

Big Techs are ramping up Capex to chase AI infrastructure, and NVIDIA is capturing the lion's share of this spending



Value Added to market cap by segment (US\$ million)

Roughly 97% of NVIDIA's market value stems from Data Center, a clear bet on its AI exposure and outsized profit potential



How we conducted this analysis

We assigned distinct EV/Sales multiples to each Nvidia segment based on the average of their respective peers. By multiplying these by each unit's revenue, we estimated the intrinsic value of each division, isolating its contribution to the consolidated Enterprise Value. This approach helps identify where the market may be undervaluing strategic assets within the company's portfolio



It's Still Early Days for AI

AI is rewriting the rules of value creation: driving productivity gains and unlocking new innovation and business models

Why is AI changing everything:

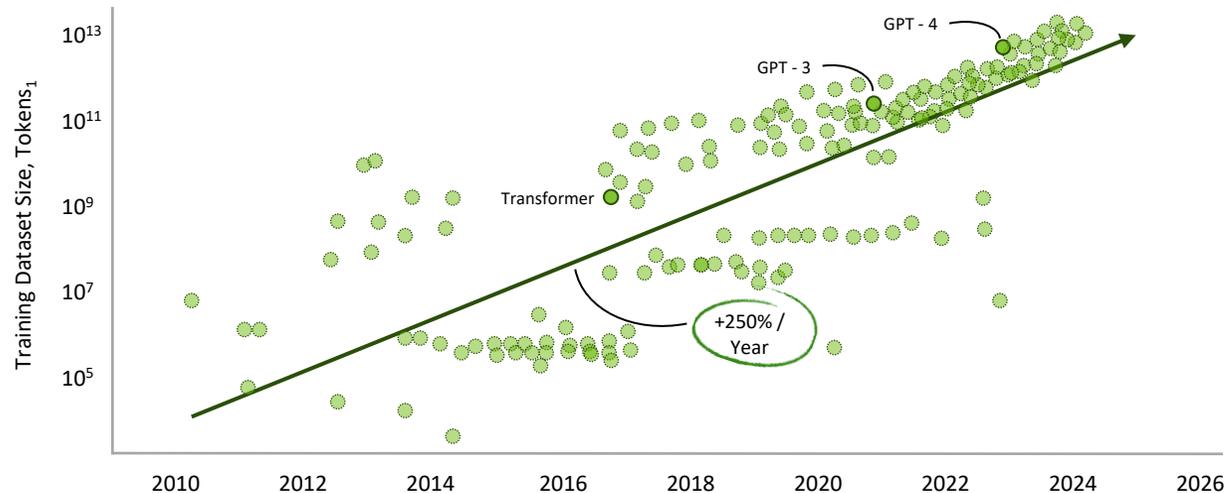
- **AI can learn and execute complex tasks** across multiple domains — something that used to require dozens of separate tools and systems
- **A simple and universal interface** drastically **lowers the barrier to AI adoption** — making it accessible to any employee, developer, or company
- **AI benefits** from massive **economies of scale and continuous improvement** — **the more it's used, the better** and more efficient it becomes

Why companies are going all in:

- **Massive productivity gains: AI cuts the cost of operational tasks across customer service, marketing, engineering, and legal**
- **AI powers entirely new products** — from copilots to intelligent search and diagnostics — **creating high-margin revenue opportunities** beyond cost reduction
- **Fear of falling behind** (the AI arms race): **Big tech and industry leaders know that whoever masters AI will dominate the next decade**

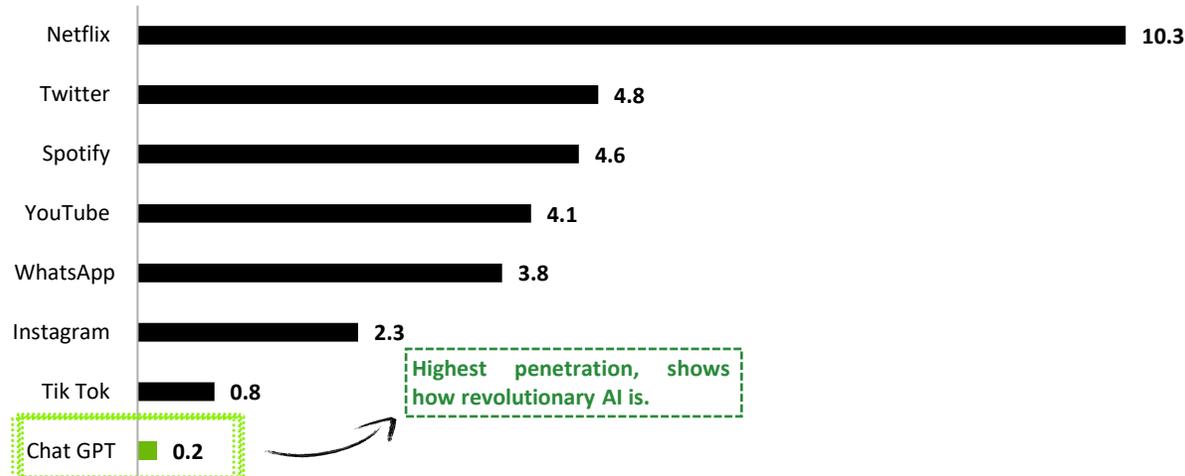
AI Model training Dataset Size by Model Release Year

The rapid rise of Generative AI has been fueled by an exponential increase in training data — with dataset sizes growing over 250% per year, enabling breakthroughs like GPT-3 and GPT-4



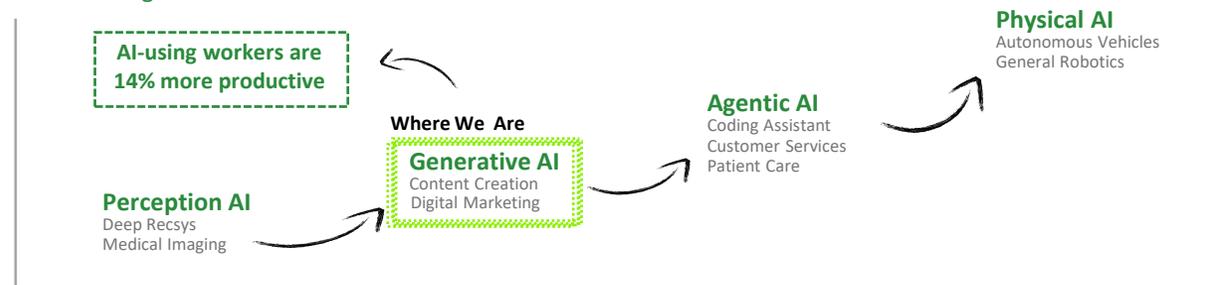
Years to Reach 100 MM Users

AI adoption shattered records, reaching 100 million users faster than any technology in history — a clear sign of the scale and speed of this new technological revolution



The Future of AI

We are still in the early stages of AI, and although there is still plenty of room to grow, Generative AI is already transforming the entire market



Andy Jassy, Amazon's CEO

"Generative AI is going to reinvent virtually every customer experience we know and enable altogether new ones about which we've only fantasized. ...Increasingly, you'll see AI change the norms in coding, search, shopping, personal assistants, primary care, cancer and drug research, biology, robotics, space, financial services, neighborhood networks, everything"



Capturing AI's Data Center Boom

How we see greater growth potential for NVIDIA and Why our view diverges from market consensus

Assumptions for NVIDIA's Data Center Revenue Forecast

We project faster growth in the data center GPU market than consensus expects, and believe NVIDIA is well positioned to maintain its leadership and capture the bulk of this upside

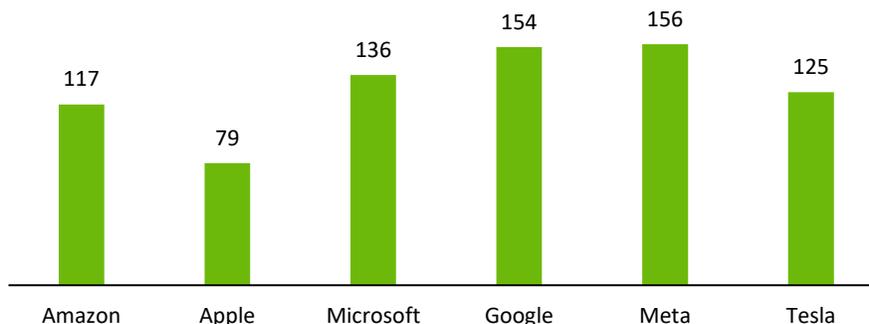
US\$ Million	2025E	2026E	2027E	2028E	2029E
Global Data Center Capex	556.600	673.486	814.918	986.051	1.193.122
Global Data Center Capex (%YOY)	21%	21%	21%	21%	21%
% of Data Center CAPEX for GPUs	45%	45%	45%	45%	45%
NVIDIA's Data Center Revenue	220.414	260.639	311.706	377.164	456.369
% Of Global CAPEX	40%	39%	38%	38%	38%
% of NVIDIA GPU Market Share	88%	86%	85%	85%	85%

Where We Diverge From the Market?

While the market is pricing in a deceleration in data center investments starting in 2028 due to concerns about overcapacity, we hold a different view. We believe the AI market is still in its early stages and will continue to expand significantly, driving sustained demand for high-performance infrastructure. The high volume of AI mentions in the latest MAG 7 earnings calls supports this thesis — indicating that major tech companies remain heavily focused on AI as a core growth driver and are likely to keep investing aggressively in data center capacity

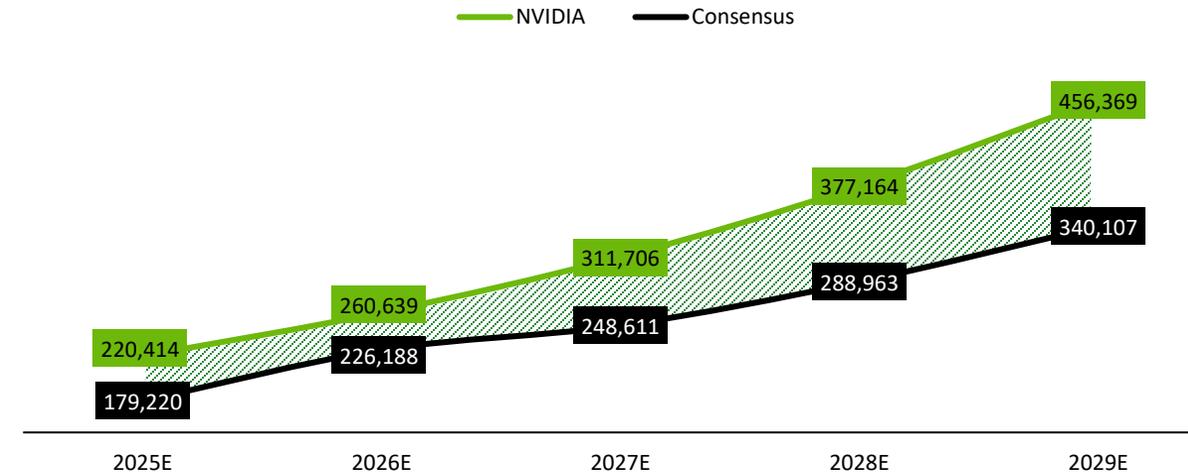
Mentions of AI in Latest Earnings Calls

```
import pdfplumber
palavras_alvo = ["ai", "artificial intelligence"]
def identificar_empresa(nome):
    nome_lower = nome.lower()
    if "apple" in nome_lower:
        return "Apple"
    elif "amazon" in nome_lower:
        return "Amazon"
    elif "microsoft" in nome_lower:
        return "Microsoft"
    elif "google" in nome_lower:
        return "Google"
    elif "meta" in nome_lower:
        return "Meta"
    elif "tesla" in nome_lower:
        return "Tesla"
    else:
        return "Desconhecido"
contagem_geral = {}
```



Data Center NVIDIA's Revenue x Consensus

We expect stronger data center growth than consensus, driven by AI, with the divergence becoming more pronounced from 2028 onward



Future Revenue Opportunities

Geopolitical tailwinds and robotics adoption could unlock incremental revenues — both currently excluded from our model but making a potential source of upside asymmetry



China Revenue

There is a possibility that NVIDIA is preparing a new U.S.-compliant version of its Blackwell chips for China. If approved, it could partially restore regional data center sales, potentially adding \$8–15B in annual revenue

When, Revenue Impact and Probability

- 2025
- 8-15 US\$ Billion
- Medium Probability



Robotics

We see a high probability of monetization through platforms like Omniverse, coupled with growing demand for GPUs powering robotics workloads — with potential for upside surprise

- 2027-?
- High Impact
- High Probability

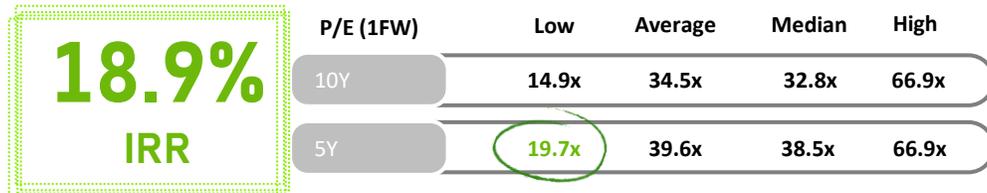


Is our Desired IRR Viable? Yes!

Our base case, grounded in conservative multiples, supports a feasible IRR

Entry & Exit: Taking a Conservative Stance — Exiting at the 5-Year Lowest

29x |————— 2028 —————| **20x**
Entry Multiple |—————| Exit Multiple



	2025	2026	2027	2028	2029
Net Income 1YF	144,327	171,250	206,368	248,763	0
(-) Entry Value	(3,415,756)	-	-	-	-
(+) Dividends	112,841	136,539	161,440	193,469	-
(+) Exit Value	-	-	-	4,975,256	-
(=) Cash Flow	(3,302,915)	136,539	161,440	5,168,725	0

Our Two-stage Fair-Multiple Model Signals a Favorable Entry Point:

	2026	2027	2028	2029
Fair Multiple	29.8	25.3	22.0	19.6
ROE	121%	121%	121%	121%
Growth	6%	6%	6%	6%
Discount Rate	11%	11%	11%	11%
Perpetuity Multiple	26.3	22.8	20.4	18.8
Growth Stage	3.4	2.5	1.6	0.8

Required P/E Analysis

5y CAGR: 421%

		Earnings CAGR						
		10%	15%	20%	25%	30%	35%	40%
IRR	Min: 15%	41x	32x	26x	21x	17x	14x	11x
	Base: 20%	47x	37x	29x	24x	19x	16x	13x
	Bull: 25%	53x	42x	34x	27x	22x	18x	15x

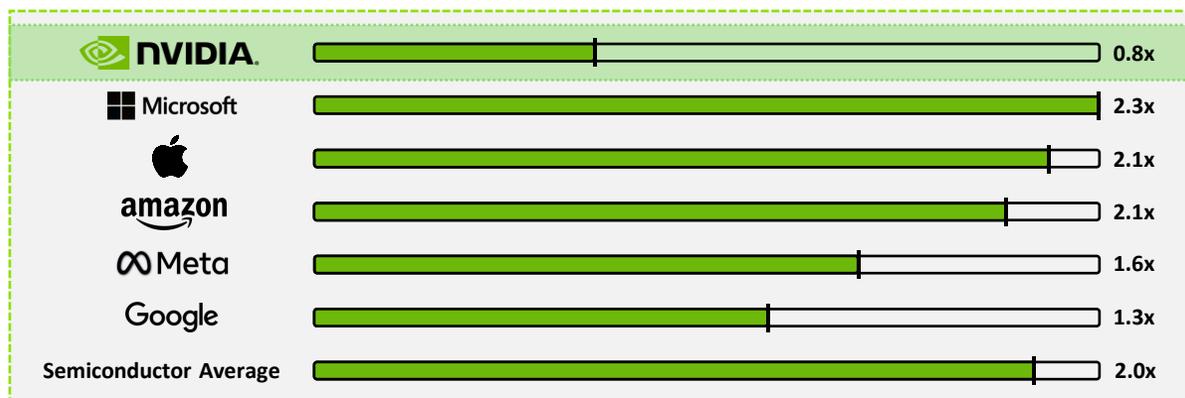
Required P/E Analysis

5y CAGR: 139%

		Revenue CAGR						
		15%	20%	25%	30%	35%	40%	45%
IRR	Min: 15%	35x	28x	23x	18x	15x	12x	10x
	Base: 20%	40x	32x	26x	20x	17x	14x	12x
	Bull: 25%	45x	36x	29x	24x	20x	16x	13x

PEG Ratio (12mFw):

NVIDIA's exceptional earnings growth justifies its P/E multiple



NVIDIA has consistently delivered EPS above consensus expectations over time, strengthening our conviction in its ability to consistently deliver:

93% of Quarter Results have **surprised** EPS Market Consensus

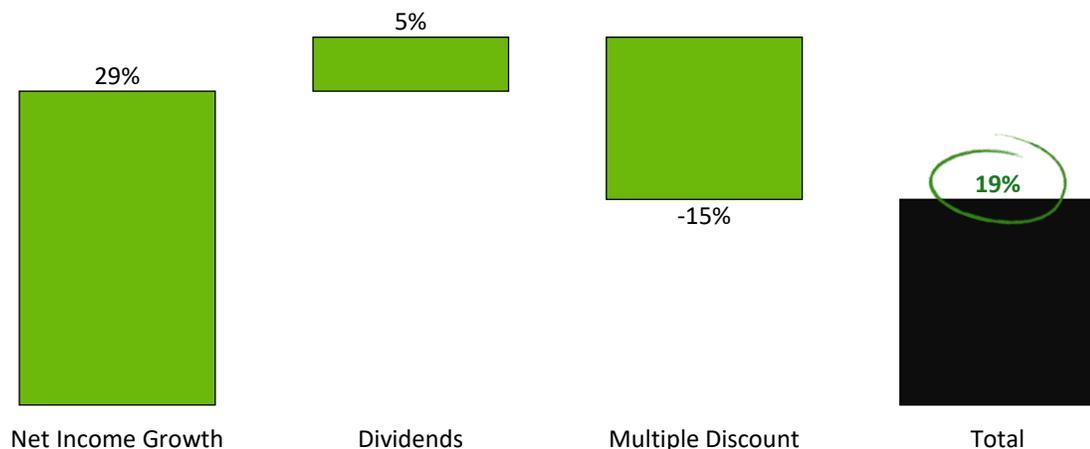


Thesis Scenarios & How much could we Lose?

Sensitivity Analysis of Our Base Case, and Downside Scenarios That Could Break the Thesis

What Drives Our NVIDIA IRR: A Decomposition of Return Drivers and Headwinds

IRR is largely a function of strong Net Income Growth — but held back by our cautious multiple



IRR Scenarios Analysis

While the upside potential remains compelling, disciplined monitoring of AI-driven capex is needed

	Bear	Base	Bull
Growth Global Data Center CapEx	18.0%	21.0%	25.0%
% of Data Center CapEx for GPUs	40.0%	45.0%	50.0%
% of NVIDIA GPU Market Share	75.0%	88.0%	90.0%
Gross Margin	65.0%	75.0%	80.0%
Exit Multiple	15x	20x	25x
IRR	(9.1%)	18.9%	43.0%

Where Our Thesis Could Break: Demand for AI Falls Short or Competitive Pressures Undermine Nvidia's GPU Leadership

IRR Impact Under Combined Downside Scenarios: Slower AI Infrastructure Investment and Structural Loss of Nvidia's GPU Market Share

ASICs Chips

These are **custom chips built for specific tasks**, and in some cases, they can outperform GPUs in efficiency and cost for AI workloads. **If hyperscalers like Google and Amazon successfully scale their own ASICs, Nvidia risks losing substantial GPU market share and the pricing power that underpins its high margins.** This could lead to a structural decline in profitability and a weakening of the competitive moat that currently supports its dominant position in AI infrastructure — with potential impact starting from 2028 onward.

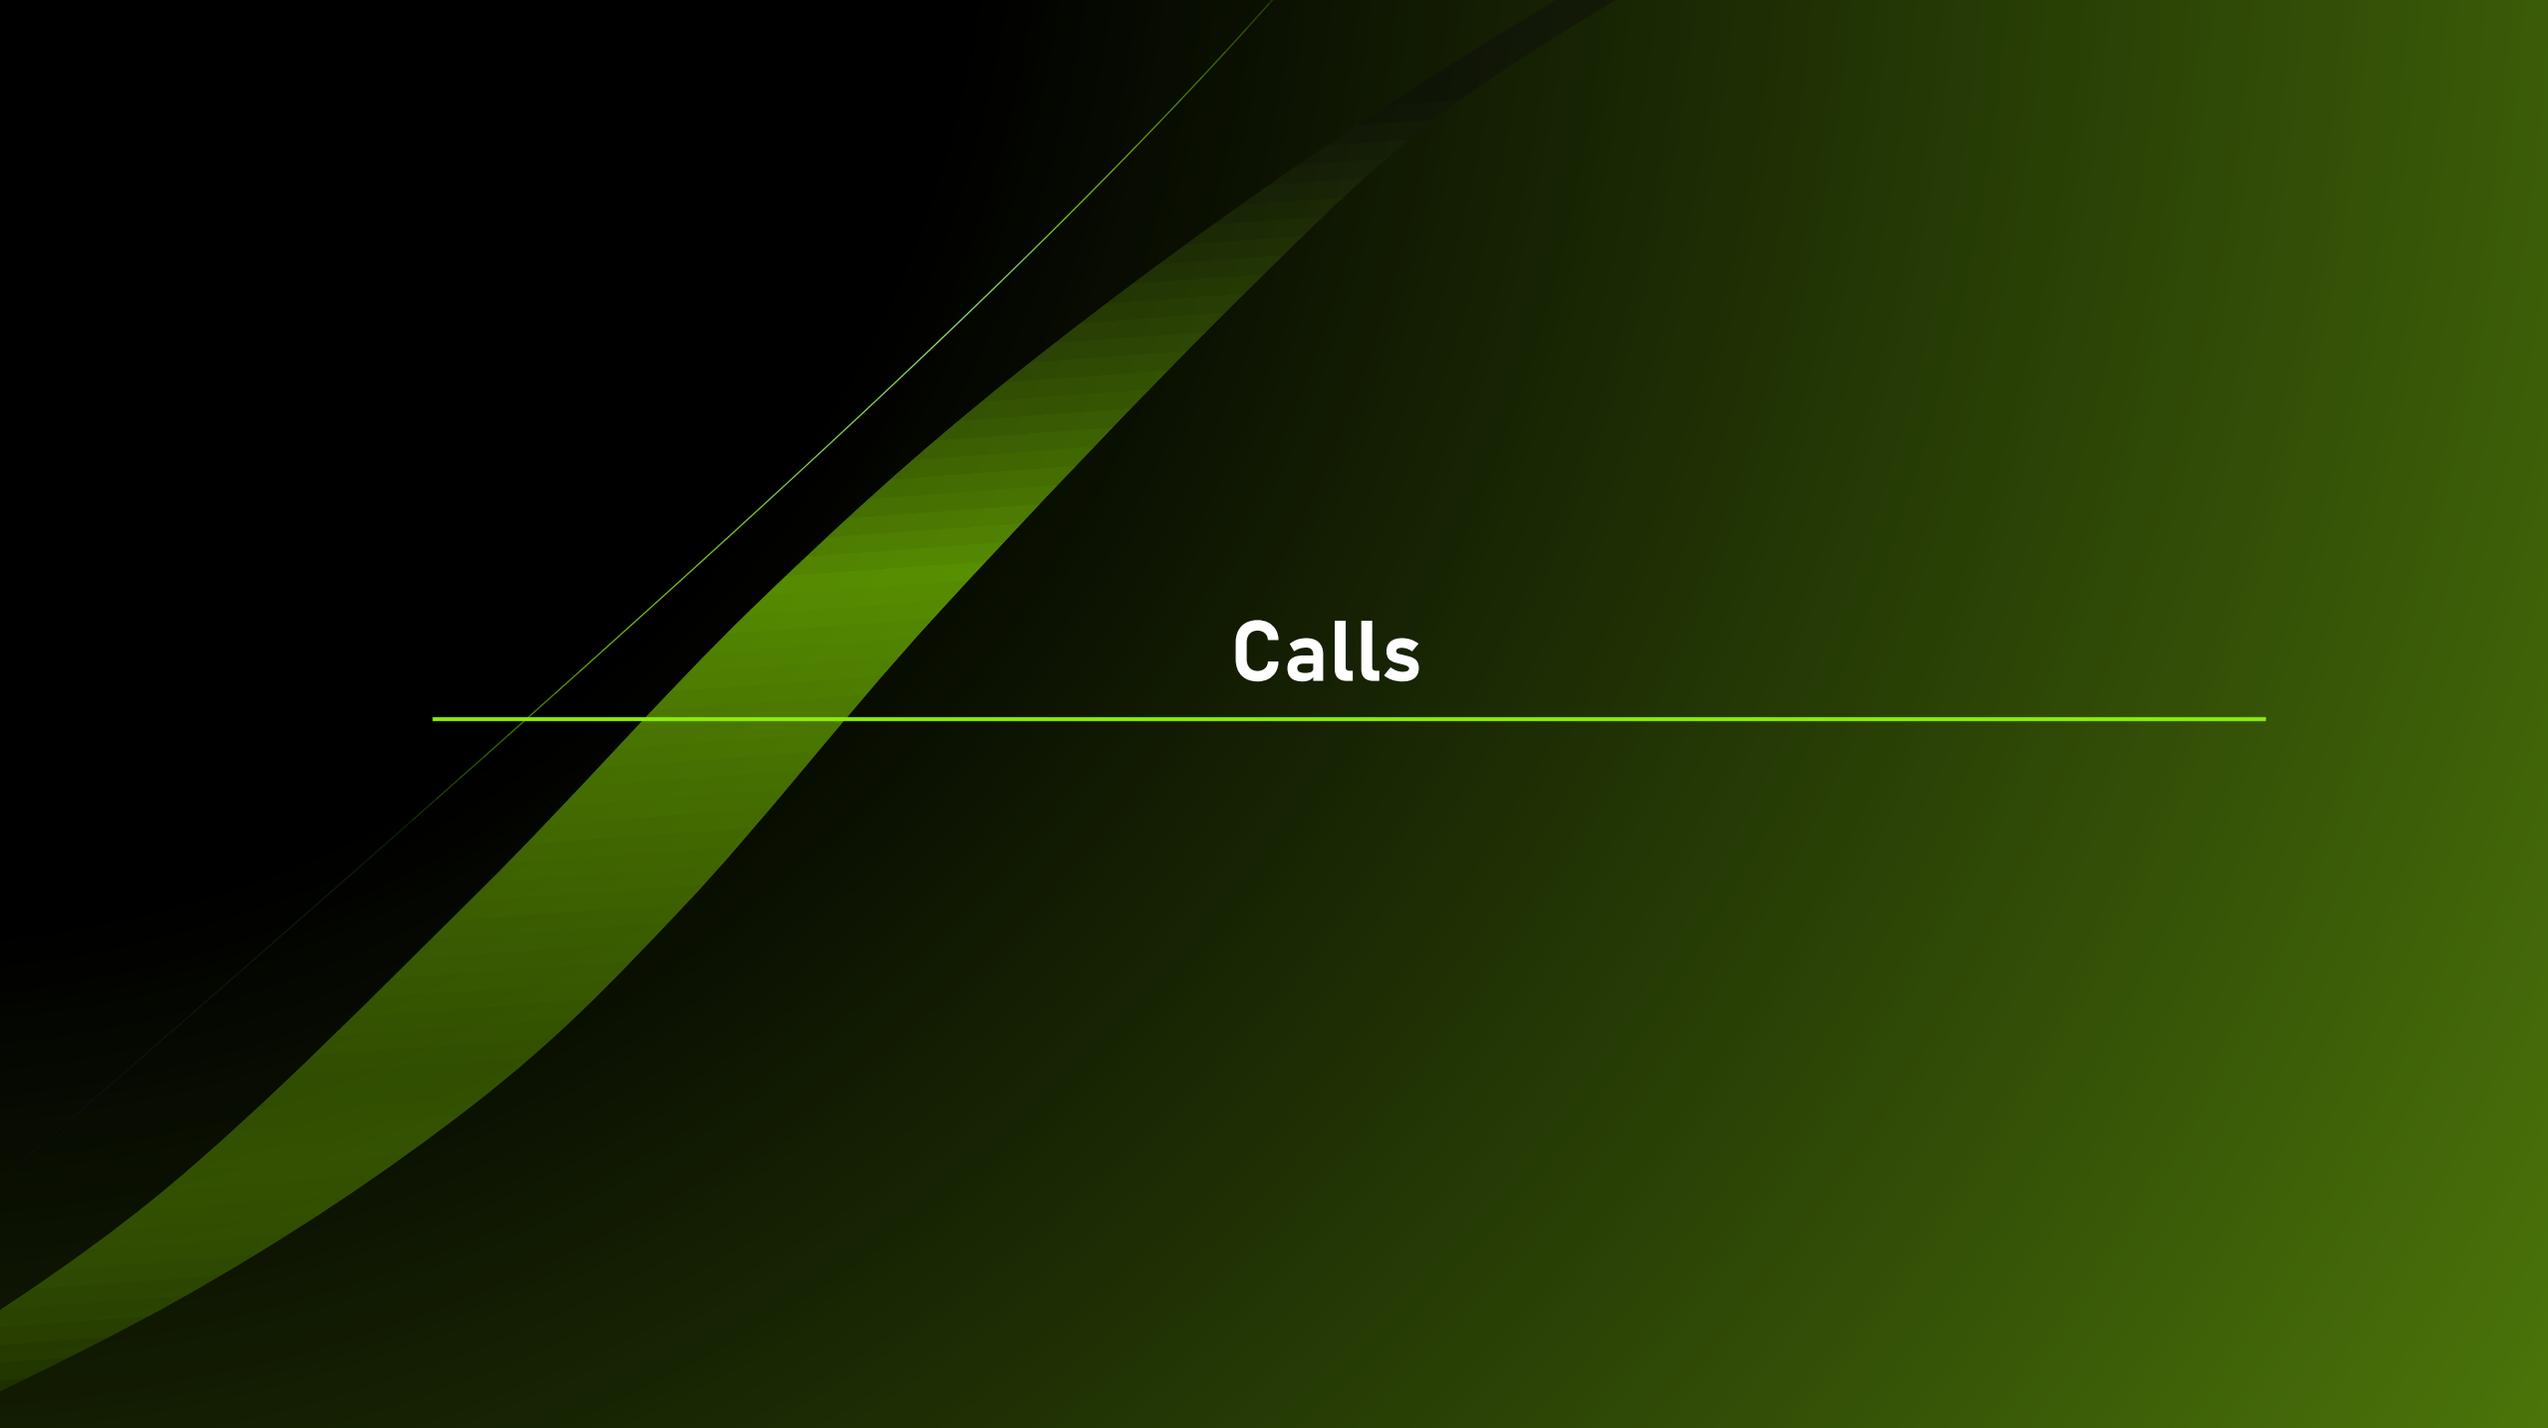
	Scenario 1	Scenario 2	Scenario 3
NVIDIA Market Share:	50%	60%	70%
Gross Margin:	60%	65%	70%
IRR	-7%	1%	9%

Slow Down in AI Investments

If **enterprises and hyperscalers begin to question the near-term returns from AI projects**, a pullback in spending could occur — especially after the recent wave of aggressive investment. This would not only lead to a **slower expansion of total data center infrastructure**, but also **reduce the percentage of CapEx allocated to GPUs.** Since Nvidia's growth is highly tied to the adoption of AI at scale, any hesitation or delay in AI monetization could directly pressure its revenue trajectory, lower utilization rates, and weaken its pricing leverage in the data center stack.

	Scenario 1	Scenario 2	Scenario 3
CAGR Capex DC:	15%	17%	19%
% of Accelerators Share:	30%	35%	40%
IRR	-4%	4%	11%

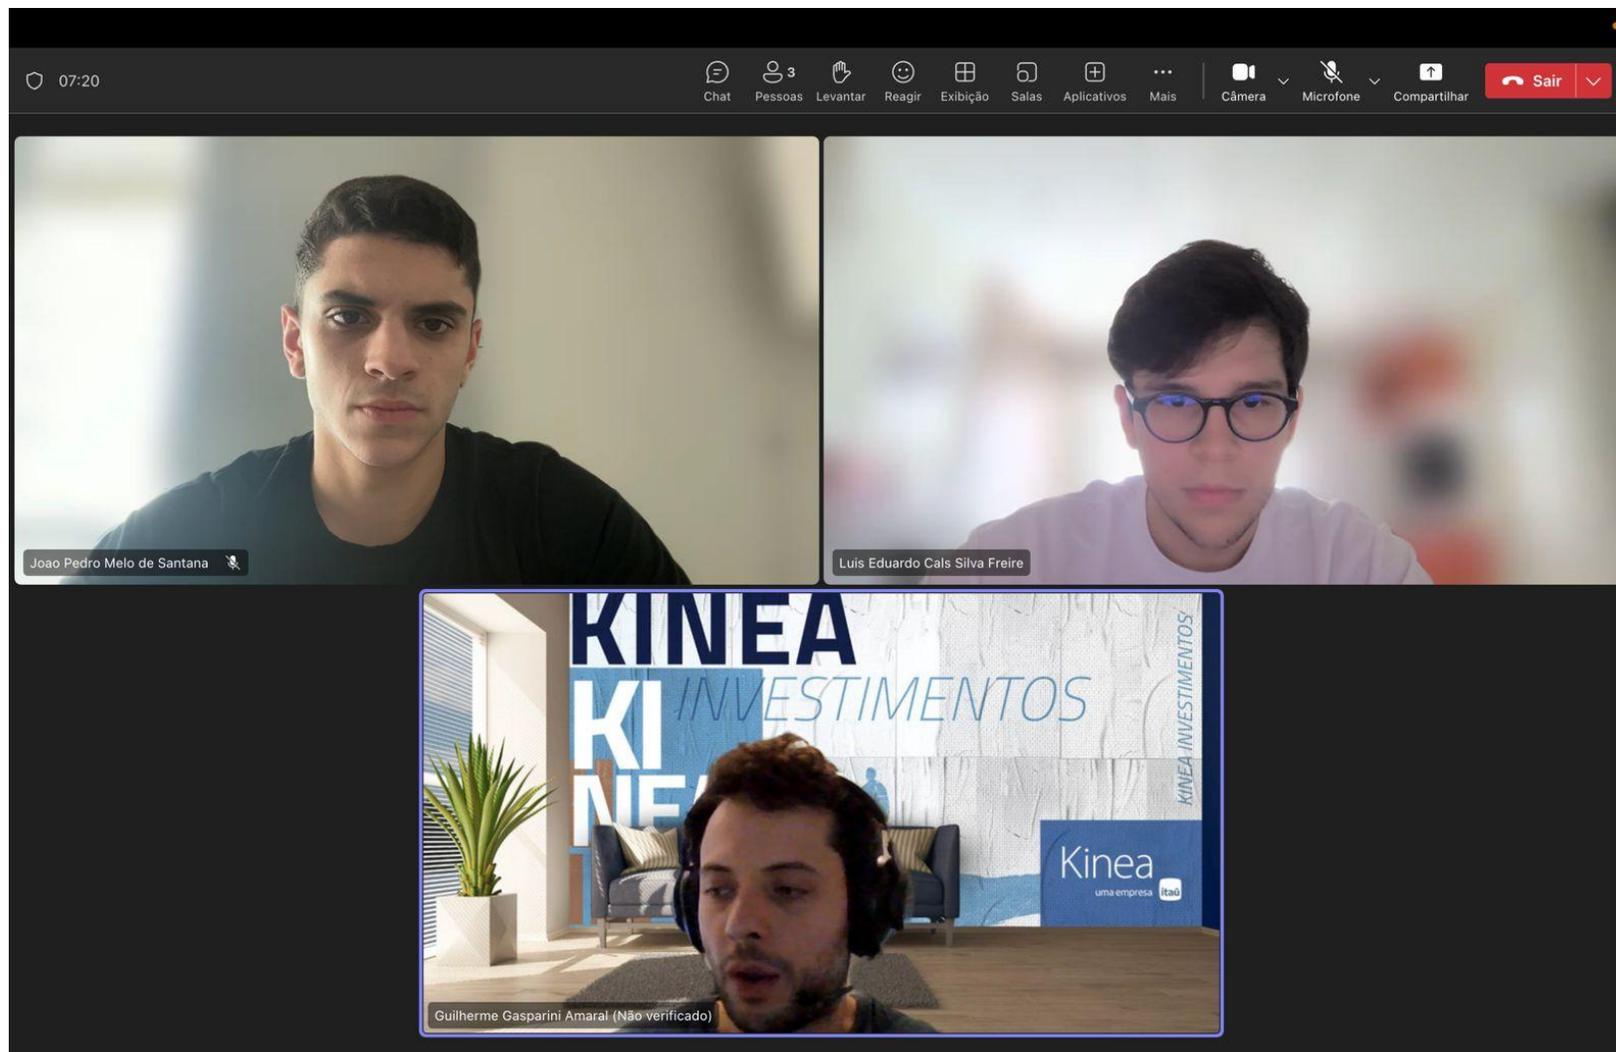




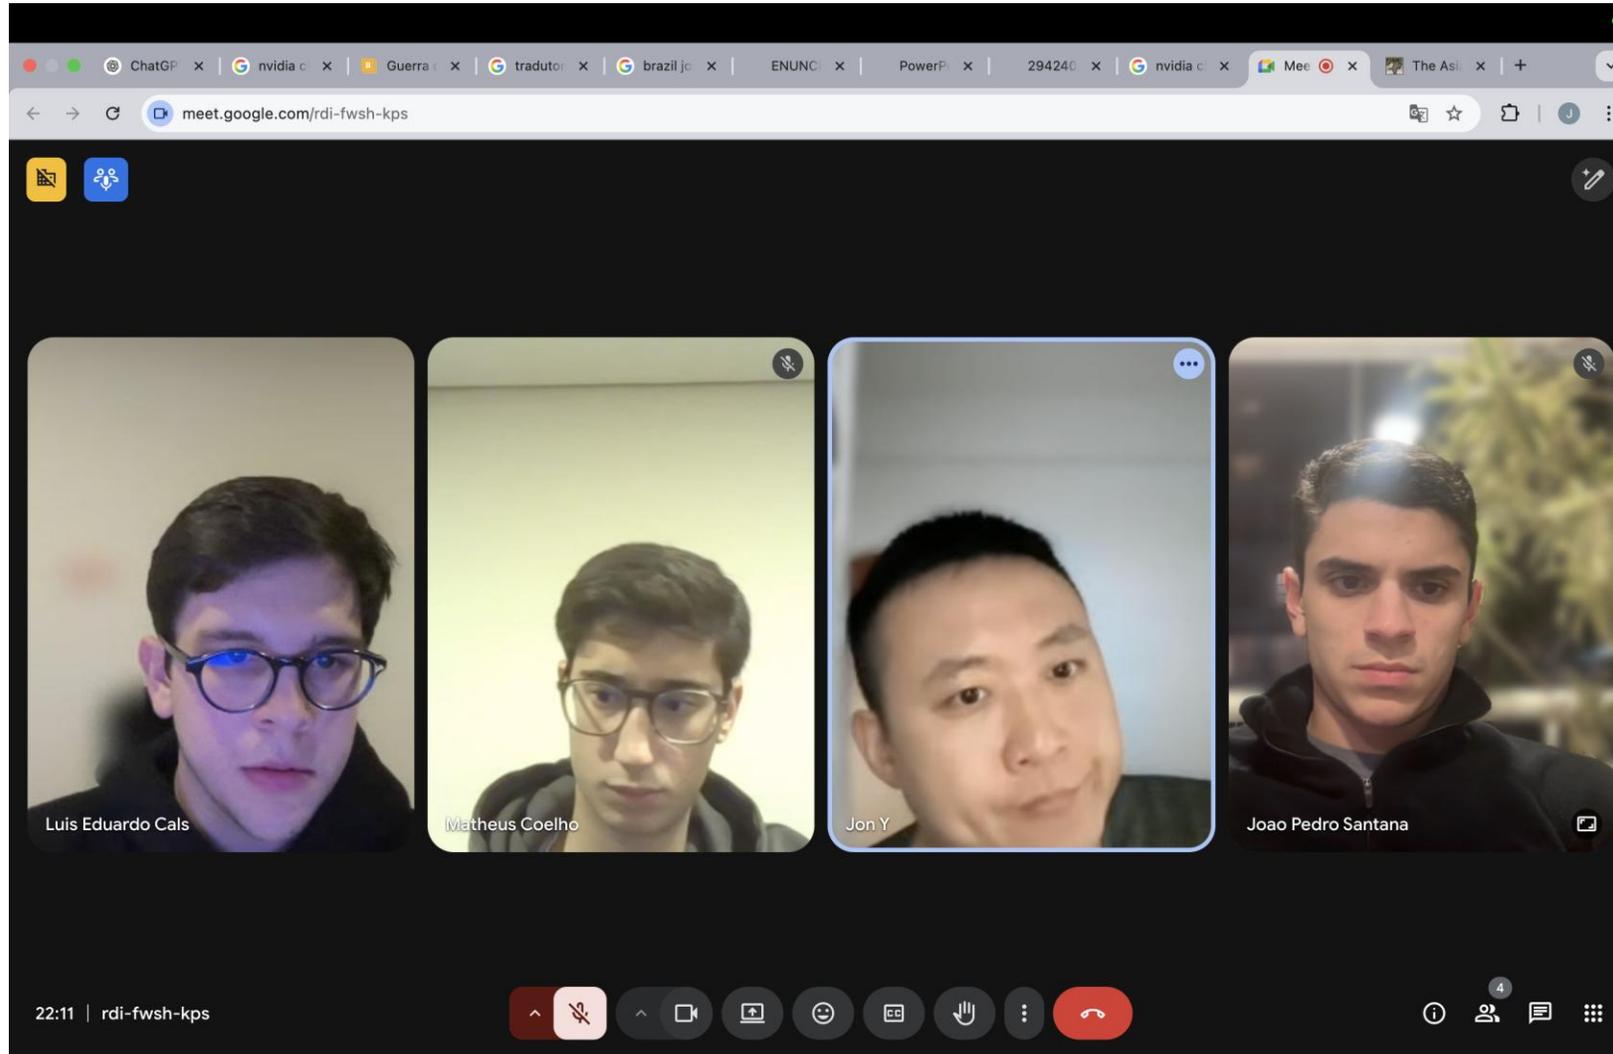
Calls

Call Guilherme Amaral – Kinea

Our conversation with Guilherme Amaral from Kinea helped us frame potential value triggers for NVIDIA over the coming months, adding depth to our timing and catalyst assessment

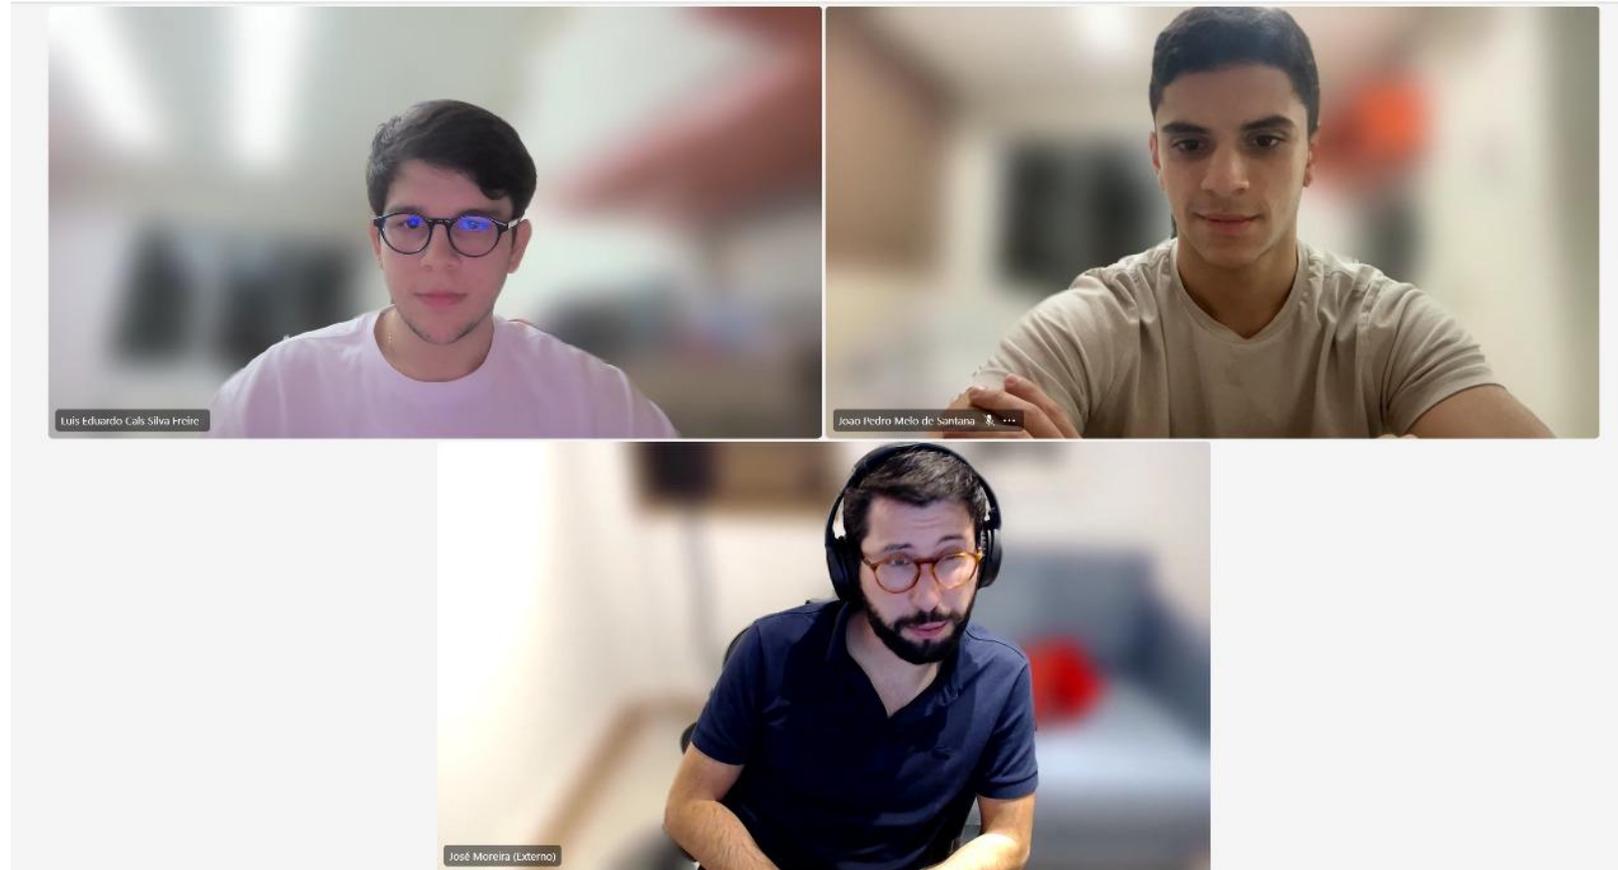


Call Jon Y – Asianometry



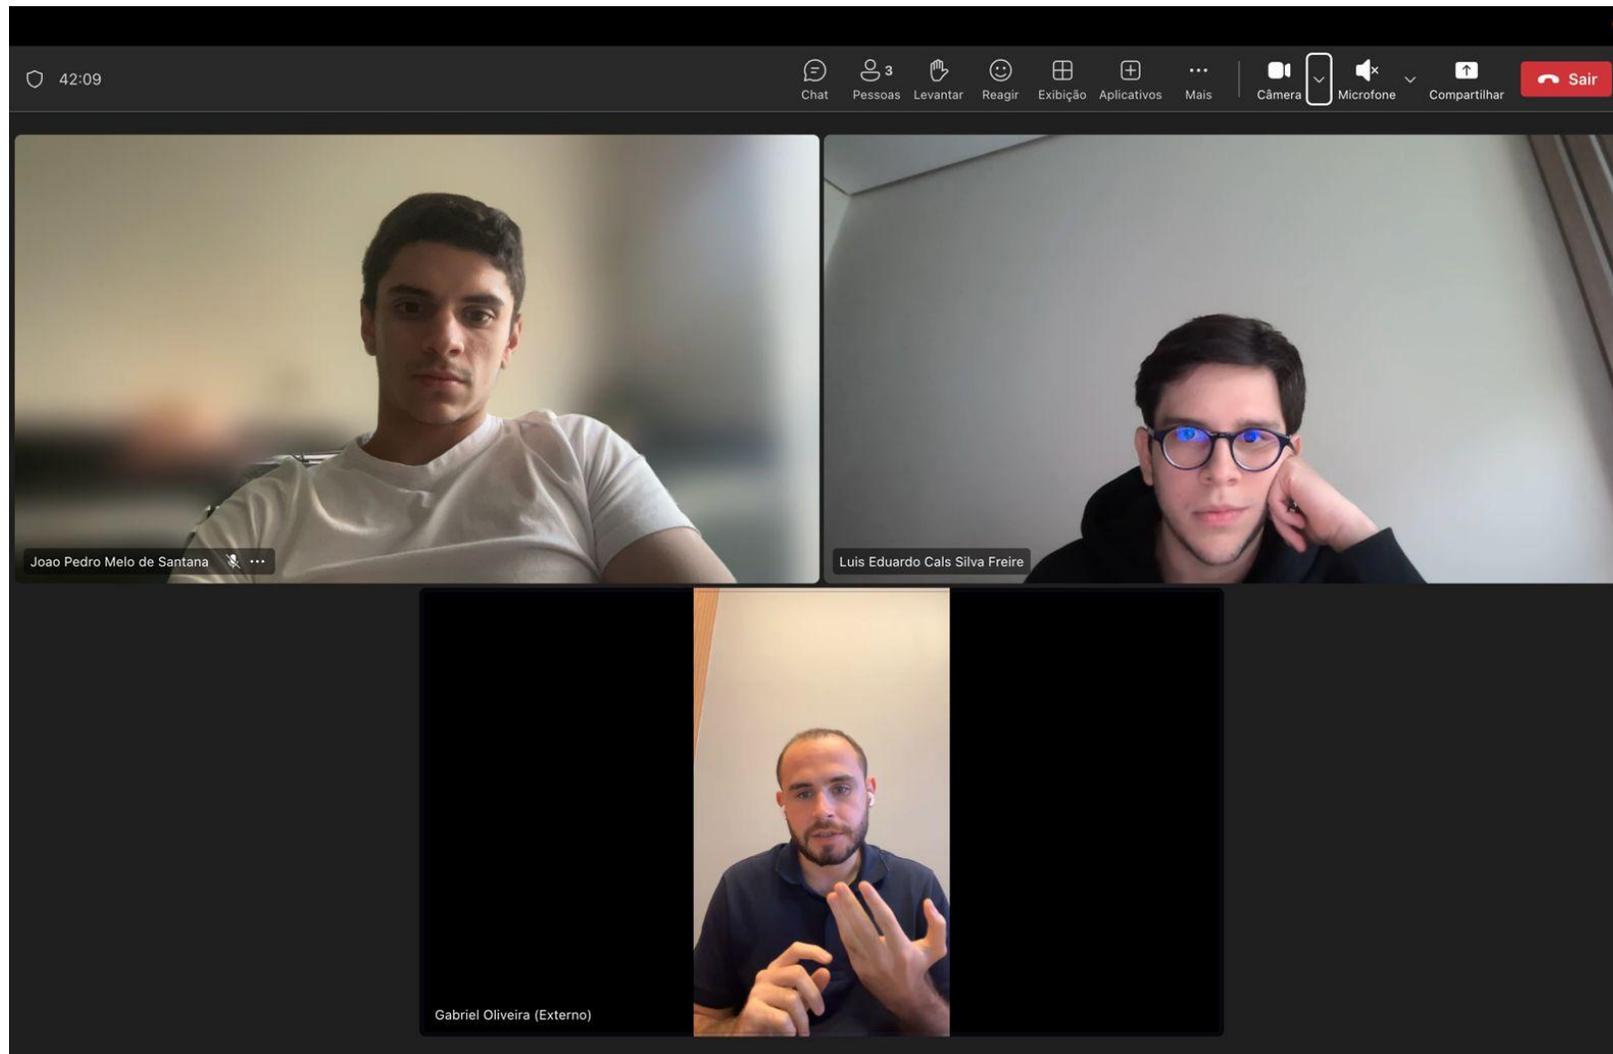
Call José Moreira – PRAGMA

Our conversation with José from PRAGMA helped us better understand the interdependencies across the semiconductor value chain

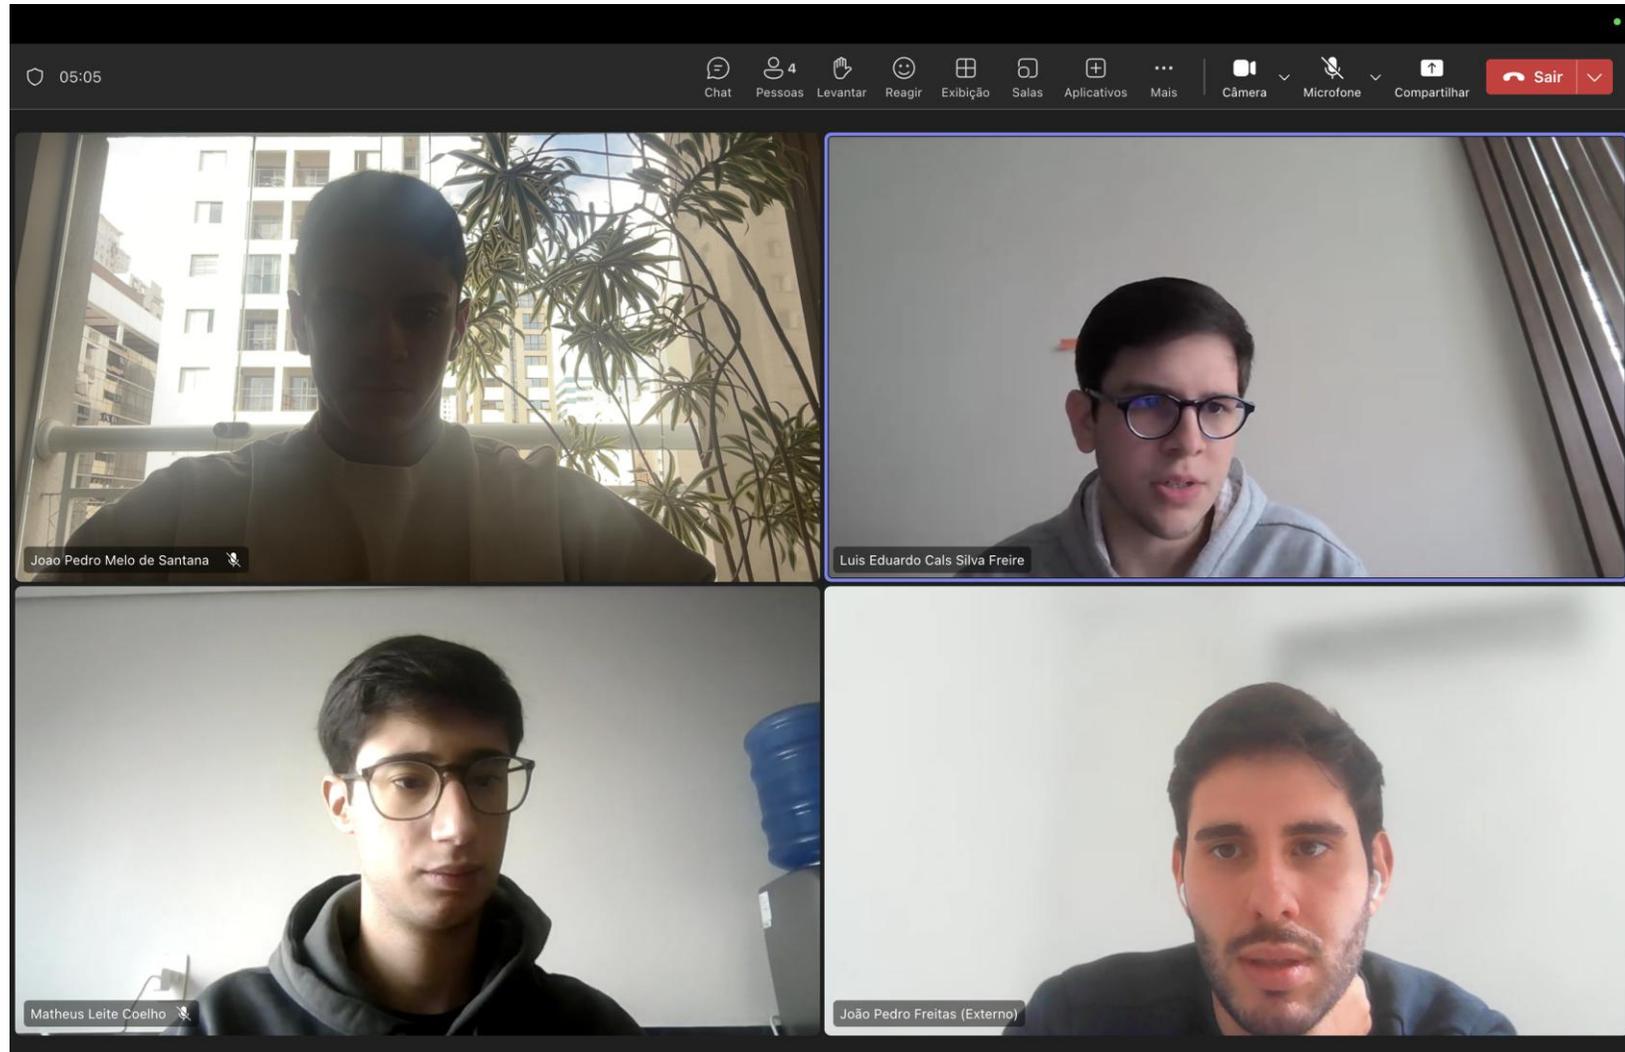


Call Gabriel Oliveira – Verde Asset

Our conversation with Gabriel Oliveira from Verde Asset provided valuable insights on NVIDIA's positioning and competitive dynamics, helping us strengthen our conviction on the company's long-term moat

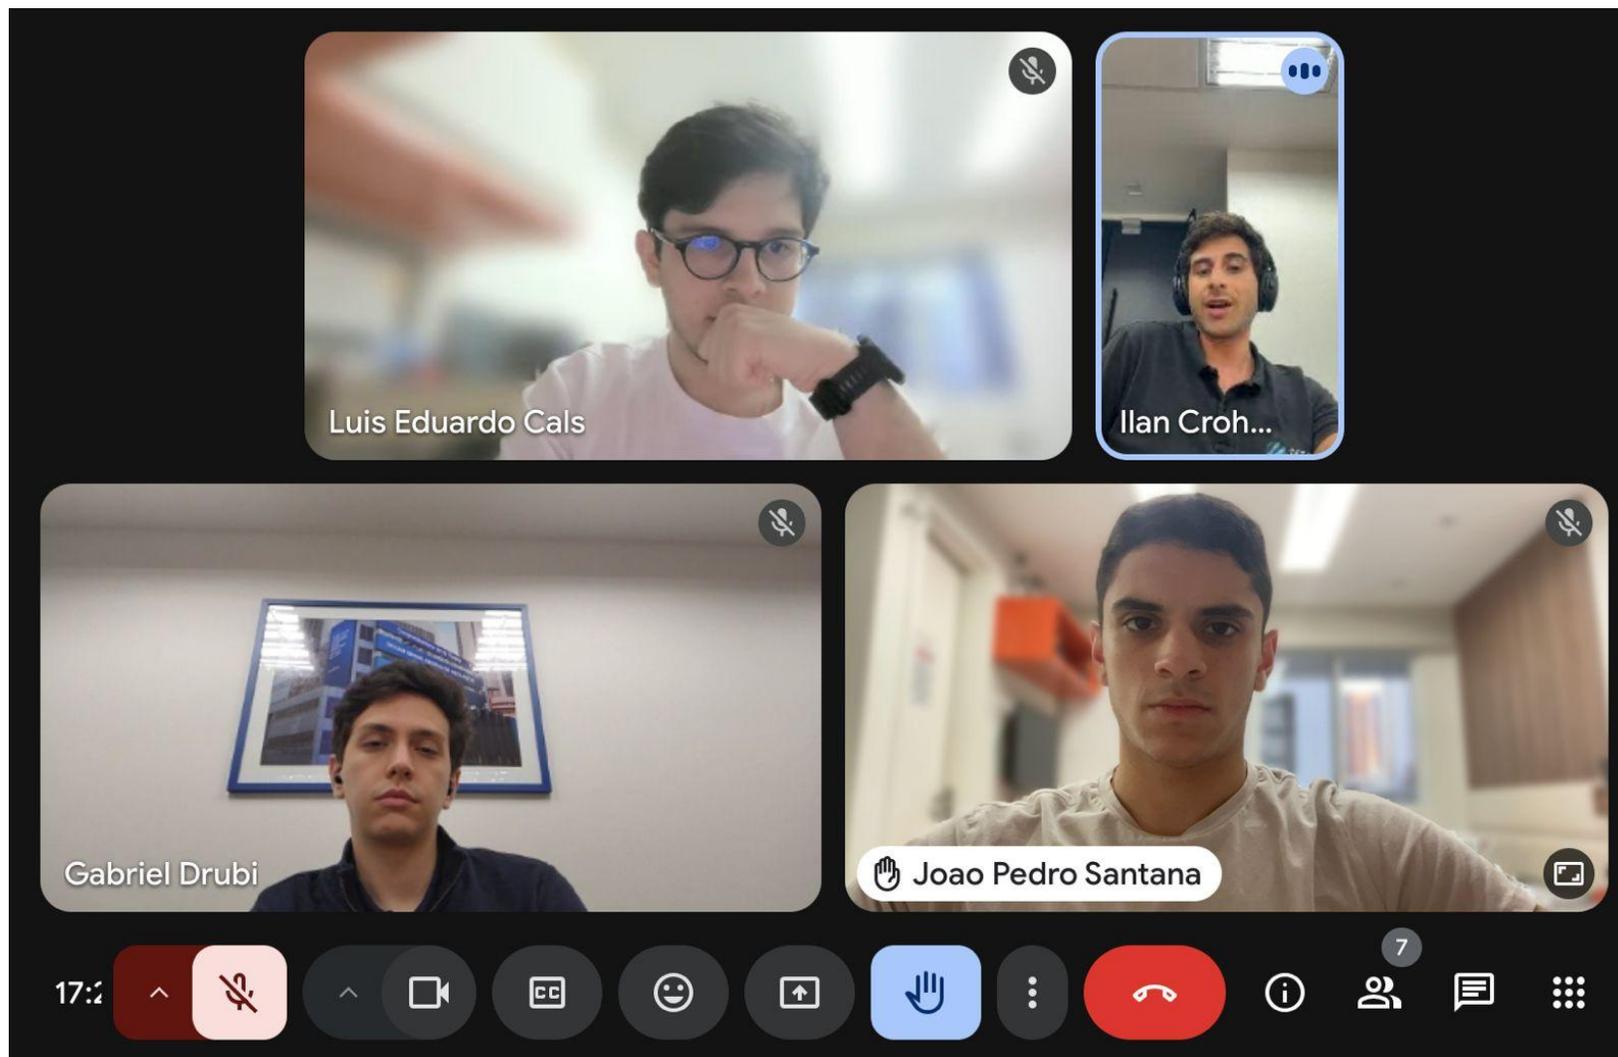


Call João Pedro Freitas – Mainú Capital



Call Ilan Crohmal – Occam

Our conversation with Ilan Crohmal from Occam helped us deepen our understanding of NVIDIA's strategic roadmap and its ability to sustain leadership in the AI compute space



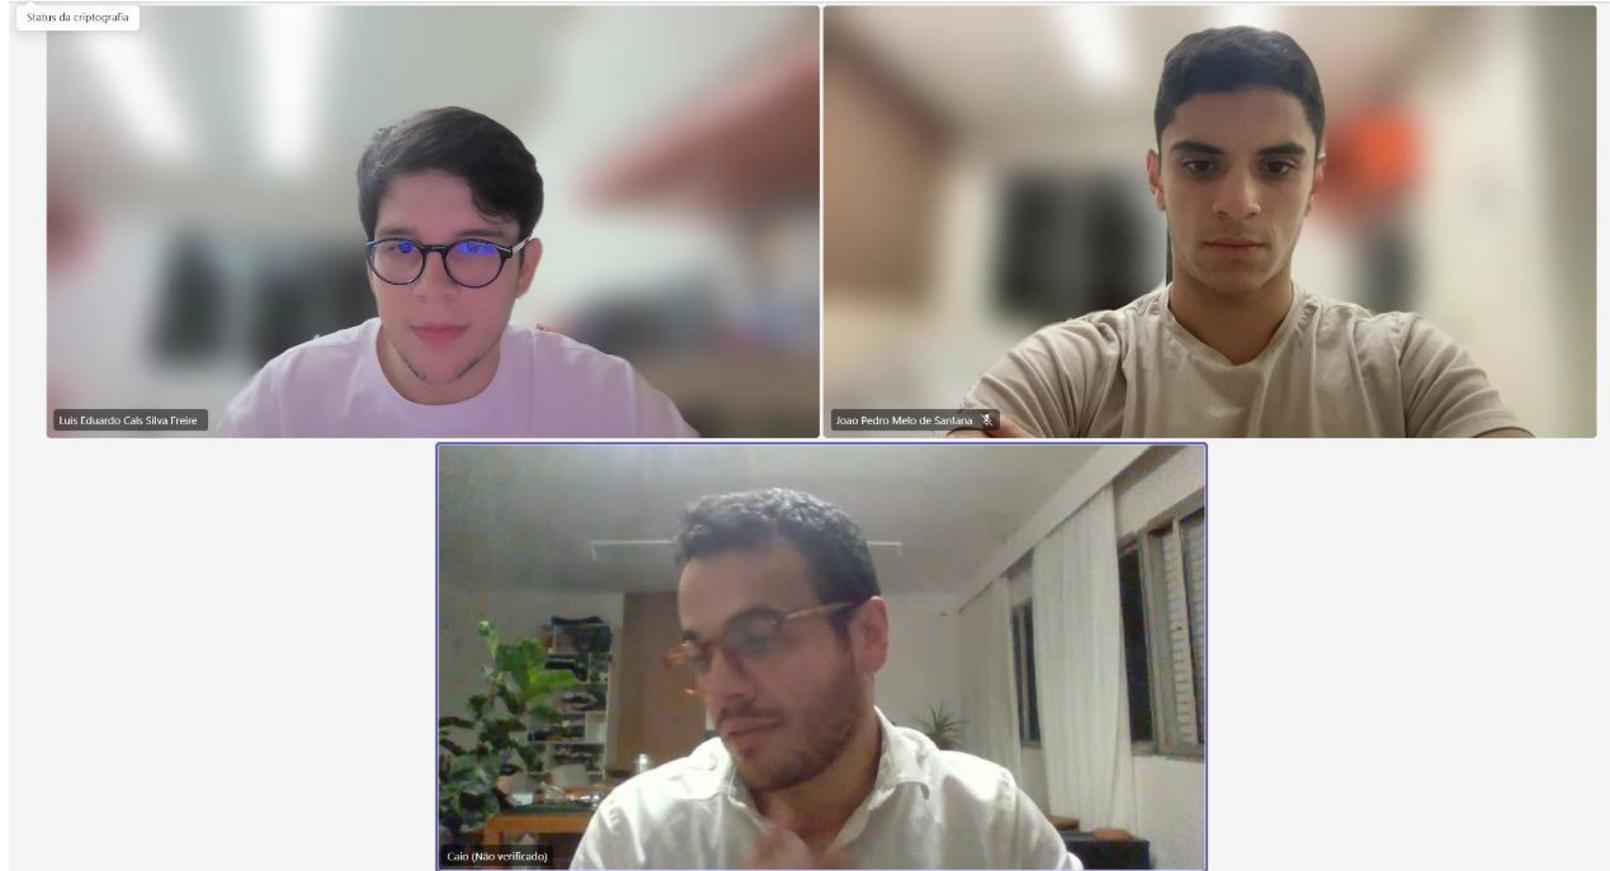
Call Karina Fugita – Geo Capital

Our conversation with Karina Fugita from GeoCapital reinforced the critical role of the semiconductor supply chain in enabling AI advancements, highlighting how structural investments in leading-edge infrastructure remain a key bottleneck and competitive moat



Call Caio Bessa – M Square

Our conversation with Caio Bessa helped us critically assess the key risks around our NVIDIA thesis and provided valuable perspective on how investors are currently positioning around the stock



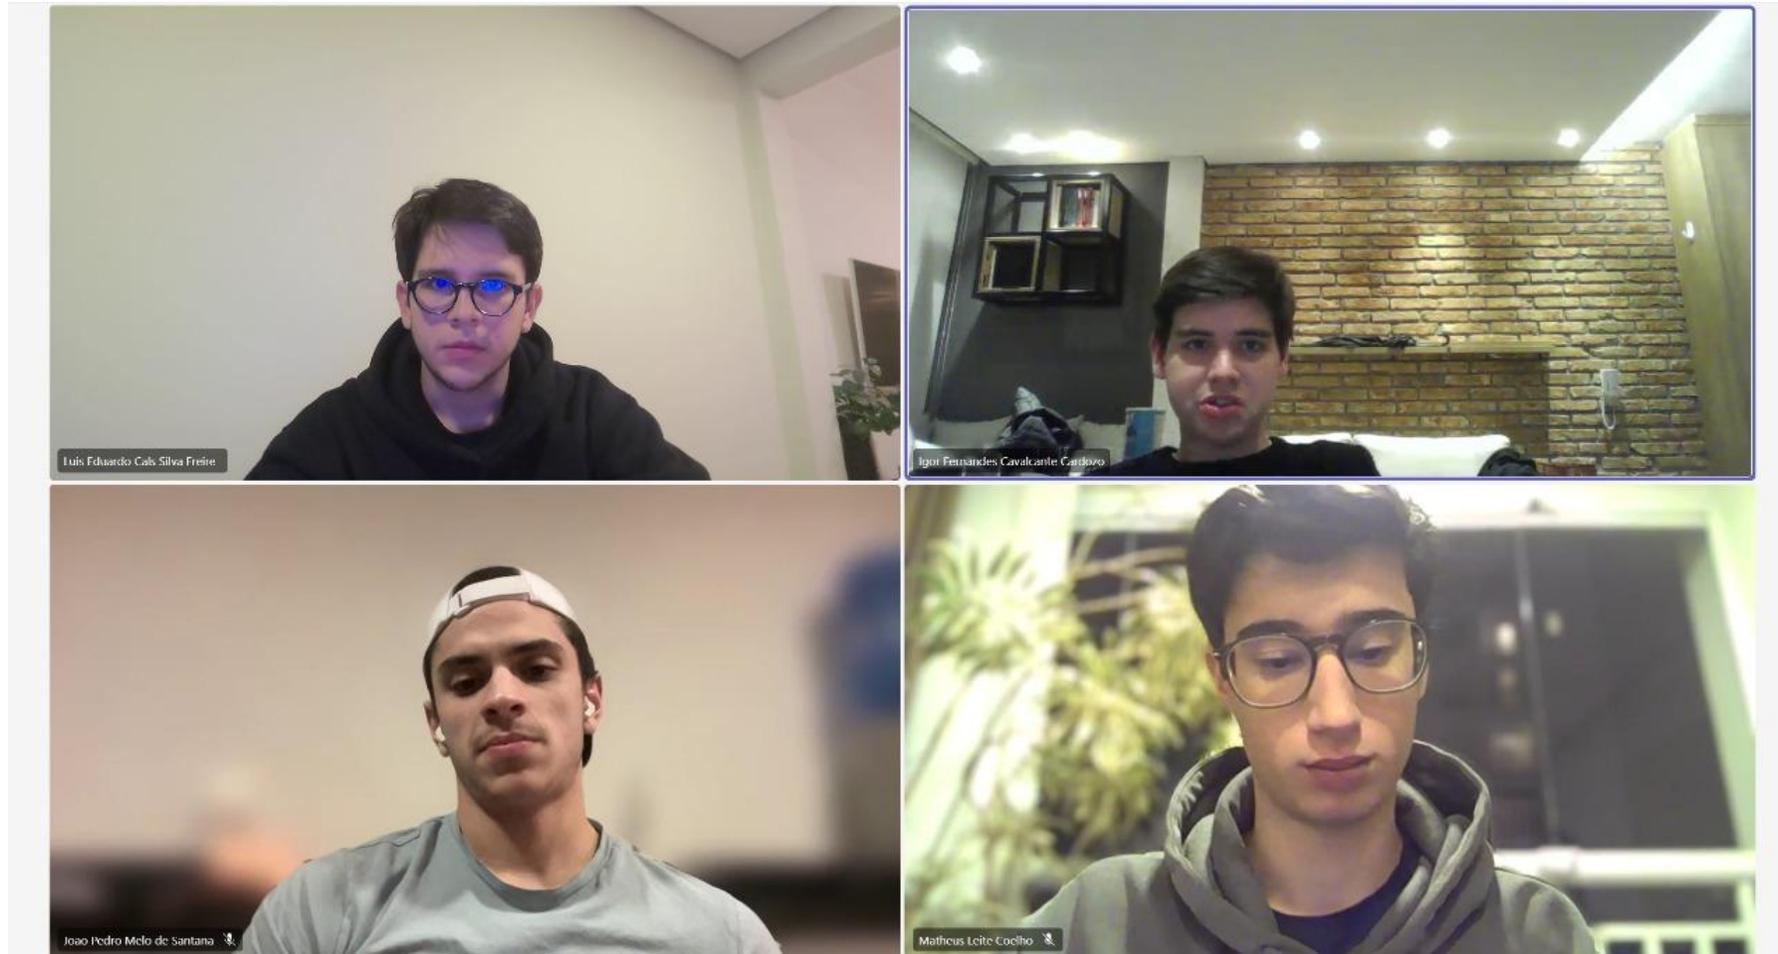
Call Lucas Dias – Aster Capital

Our conversation with Lucas from Aster Capital highlighted that AI demand remains a structural trend, with hyperscalers continuing to invest aggressively to support AI workloads and expanding use cases



Call Igor Fernandes – AZ Quest

Our conversation with Igor Fernandes from AZ Quest helped us refine our view on the company's management quality, highlighting the importance of strong execution capabilities and strategic vision in capturing the AI-driven growth opportunity



Call Adriano Marques – Ascenty

Our conversation with Adriano Marques, from Ascenty (the largest data center operator in Latin America), reinforced the strong and sustained demand for data center capacity, particularly driven by AI workloads and hyperscaler clients

