AZ STOCKS
Stocks from A to Z

Nvidia

The company that built the infrastructure of AI, with software margins and a valuation that already prices in much of its success

DEEP DIVE

November 2026

azstocks.net

COMPANY PROFILE

Nvidia

THESIS

Nvidia has become the central company of AI infrastructure, but its own customers want to reduce their dependence on it. In this report I assess how defensible its position is and what the company is really worth.

THE COMPANY

Nvidia defines itself as an accelerated computing platform company. It does not sell only chips: it sells a complete ecosystem of hardware (GPUs, NVLink, InfiniBand), software (CUDA, TensorRT), frameworks and services.

That vertical integration is what makes the business so hard to replicate. Since FY2027 it presents its business in two large platforms: Data Center and Edge Computing.

TARGET PRICE
USD 268

UPSIDE / DOWNSIDE
+43.1%

SECTOR
Semiconductors

EXCHANGE
NASDAQ

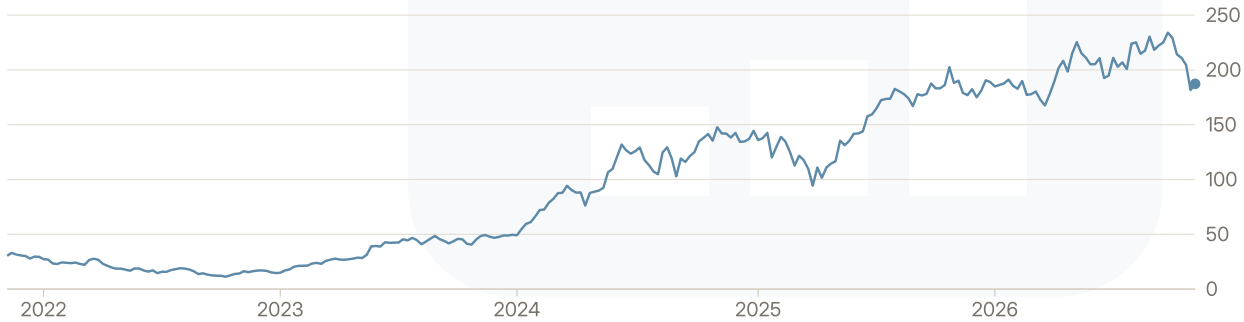
CURRENCY
USD

LAST UPDATED
November 2026

MARKET DATA AS OF 2026-11-06

PRICE	52-WEEK RANGE	5-YEAR CHANGE	BETA (5Y)
USD 187.31	164.27 – 243.37	+516%	2.33

WEEKLY SHARE PRICE, LAST 5 YEARS



Legal disclaimer

The information contained in this document is for educational and informational purposes only. It does not constitute an offer, recommendation, or financial, investment, legal or tax advice. The opinions expressed reflect a personal analysis and may change without notice.

Investing in financial markets involves risks, including the possible total loss of the capital invested. Each reader or investor is responsible for doing their own research and, if necessary, consulting a registered financial advisor before making investment decisions.

The author assumes no responsibility for losses or damages arising directly or indirectly from the use of the information presented here.

Contents

01	Introduction	05
02	Business model	06
03	Business quality	09
04	Management and capital allocation	15
05	Competitive position	17
06	The bulls vs. the bears	24
07	Valuation	25
08	Conclusion	28

SECTION 01

Introduction

Nvidia is today not only the most valuable semiconductor company in history but also the company with the largest market capitalization in the world. It may also be the company that best embodies the fundamental tension of this moment: the market does not necessarily assign it absurd multiples relative to its growth, but it does give it an enormous market cap and very demanding expectations, on the premise that artificial intelligence is not a fad but a transformation of the global economy, and that Nvidia will remain at the center of that transformation for much of the next decade.

That premise may be right. But it can also leave little room for error if ASICs capture a relevant share of inference, if China stays permanently blocked, if hyperscalers slow their CapEx or if a competitor manages to erode, even partially, the CUDA ecosystem.

Nvidia was founded in 1993 by Jensen Huang, Chris Malachowsky and Curtis Priem in Santa Clara, California, with the initial goal of revolutionizing graphics for video games. It trades on NASDAQ under the ticker **NVDA**. Today, with more than 36,000 employees, the company does not sell chips: it sells the computing infrastructure on which the AI era is being built.

Nvidia's stated mission is **"to accelerate computing for the benefit of humanity."** What began as a video game graphics business is today the backbone for training the world's most advanced language models, from GPT to Gemini to Claude. Practically all the generative AI you use today was trained on Nvidia hardware.

This analysis shows why Nvidia built one of the strongest moats in the technology industry, why that moat is being challenged from within by its own customers, and whether it can still be an attractive investment at the current valuation.

SECTION 02

Business model

Nvidia describes itself as an “accelerated computing platform” company. It does not sell only chips: it sells complete ecosystems that include hardware (GPUs, NVLink, InfiniBand), software (CUDA, cuDNN, TensorRT, Triton), frameworks (NeMo, cuOpt) and services (DGX Cloud, NIM microservices). The vertical integration of that stack is what makes the business extraordinarily hard to replicate.

Starting in FY2027, Nvidia began reorganizing how it presents its business to better reflect its current and future drivers. The company now talks about two large platforms: **Data Center** and **Edge Computing**. Within Data Center it separates Hyperscale from ACIE, which includes AI Clouds, Industrial and Enterprise. Edge Computing groups processing devices for physical AI and agentic AI, including PCs, consoles, workstations, robotics, automotive and AI-RAN. To keep historical comparability, this analysis still uses the traditional segments, but the reporting change is relevant: Nvidia no longer wants to be read as a GPU company but as a complete AI infrastructure platform.

> ***What is a GPU?*** *A GPU (Graphics Processing Unit) is a chip originally designed for video games. Unlike the CPU (a computer’s main processor, which does a few things very intelligently), the GPU has thousands of small cores doing millions of simple calculations at the same time. That massive parallelism turned out to be perfect for training artificial intelligence: training an AI model is basically doing the same mathematical calculation trillions of times. In 2012 researchers discovered that Nvidia GPUs could do that work 100 times faster than any alternative. From that moment, Nvidia stopped being a video game company.*

Segment 1: Data Center

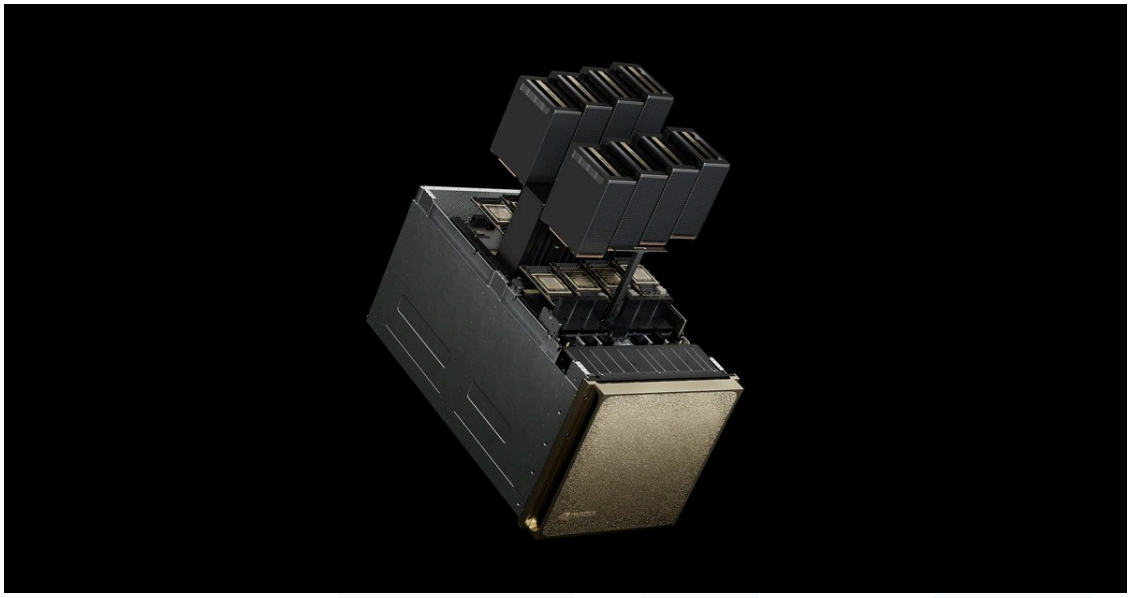
Data Center is the absolute heart of Nvidia. It generated $193.7B in FY2026, 89.7% of total revenue, and $62.3B in Q4 FY2026, +75% year over year. In Q1 FY2027 it accelerated again to $75.2B, and in Q2 FY2027 it reached **$89.0B, +117% year over year**. Even after a gigantic comparison base, Nvidia’s core business keeps growing at a speed with practically no precedent in semiconductor history. Nvidia’s Data Center alone already has more revenue than many of the world’s most relevant technology companies.

The flagship product of the current era is the **Blackwell** architecture, introduced in 2024 and scaling aggressively in 2025. The B200 chip offers up to 20 petaflops of inference performance and 4 petabytes per second of memory bandwidth with HBM3e. In plain terms:

- 20 petaflops = it can do 20 million billion mathematical operations per second. A calculator does one operation at a time; this chip does 20,000,000,000,000,000 at once. That power is what lets an AI model answer complex questions in seconds instead of hours.

- 4 petabytes per second of memory bandwidth = the speed at which the chip can read and write data in its memory. More bandwidth means the chip never waits; it is always processing.
- HBM3e is the type of memory that achieves that speed. HBM stands for *High Bandwidth Memory*, stacked in layers on the chip to minimize the distance the data travels.

DGX B200 systems, which integrate 8 B200 GPUs, sell for $300,000 to $400,000 each. Complete racks (GB200 NVL72) reach about $3 million per rack. And demand consistently exceeds supply.



Source: Nvidia.

The next step is **Vera Rubin**, the architecture that succeeds Blackwell and is already in full production. Vera Rubin should not be understood as a single chip but as a complete rack-scale computing platform. The Vera Rubin NVL72 system integrates 36 Vera CPUs and 72 Rubin GPUs, connected by NVLink 6, together with ConnectX-9 and BlueField-4. The idea is clear: Nvidia does not want to sell only the accelerator, but the whole “AI factory”.

According to Nvidia, Vera Rubin NVL72 aims to train mixture-of-experts models (a technique that splits a large model into specialized sub-networks called “experts”) using a quarter of the GPUs needed with Blackwell, and to deliver up to 10x more inference throughput per watt, with a cost per token up to 10 times lower. This is key because the debate about Nvidia is no longer only about training but about the cost of inference. If each generation lowers the cost per token faster than its customers’ ASICs, Nvidia can keep a much larger share of the inference market than the bears expect.

Product cadence remains one of the most powerful mechanisms of the moat. Hopper launched in 2022, Blackwell in 2024 and Vera Rubin entered production in 2026. Each generation improves performance, efficiency and cost per token, forcing hyperscalers to keep investing so as not to fall behind. It is not a subscription in a contractual sense, but it works as technological recurrence.

Within Data Center there is a subcategory that deserves attention: **networking**. In Q1 FY2027, Data Center Networking reached $14.8B, up 199% year over year. InfiniBand, NVLink, Spectrum-X, ConnectX and BlueField are increasingly critical as clusters scale. At some point, the network that

connects the GPUs becomes as important as the GPUs themselves, and Nvidia also has a dominant position there.

Segment 2: Gaming

Gaming built Nvidia over 30 years and is today the number two segment with **$16B in FY2026 (+43% YoY)**. The **GeForce RTX** line, from entry-level models to the RTX 5090, dominates the consumer GPU market with a share that has historically exceeded 80% of the discrete market.

What makes gaming interesting today is the convergence between video games and artificial intelligence. GeForce cards include the same **Tensor Cores** (cores specialized for AI) as data center GPUs, and technologies like **DLSS 4** use AI to improve image quality: the game renders at lower resolution and a neural network "reconstructs" the image at high quality. AMD cannot replicate that differentiation without an equivalent software ecosystem.

But gaming matters strategically for something beyond sales: it produces **developers familiar with CUDA**. An engineering student who learns to program for games using Nvidia technology eventually brings those skills to the data center. Gaming is the funnel that feeds the ecosystem, and that ecosystem is the company's deepest moat.

Segment 3: Professional Visualization

Professional Visualization includes professional-grade GPUs for 3D design, engineering, simulation and visual effects. It generated **$3.2B in FY2026 (+70% YoY)**, driven by the adoption of Blackwell GPUs in professional notebooks.

Its most interesting bet is **Omniverse**, a software platform for creating "digital twins" of real physical objects. A factory can build an exact digital replica of its production line, simulate robots and workers, test configurations and optimize efficiency without moving a single physical machine. BMW, Amazon and dozens of industrial companies already use it. It is the kind of application that turns hardware sold into recurring software revenue.

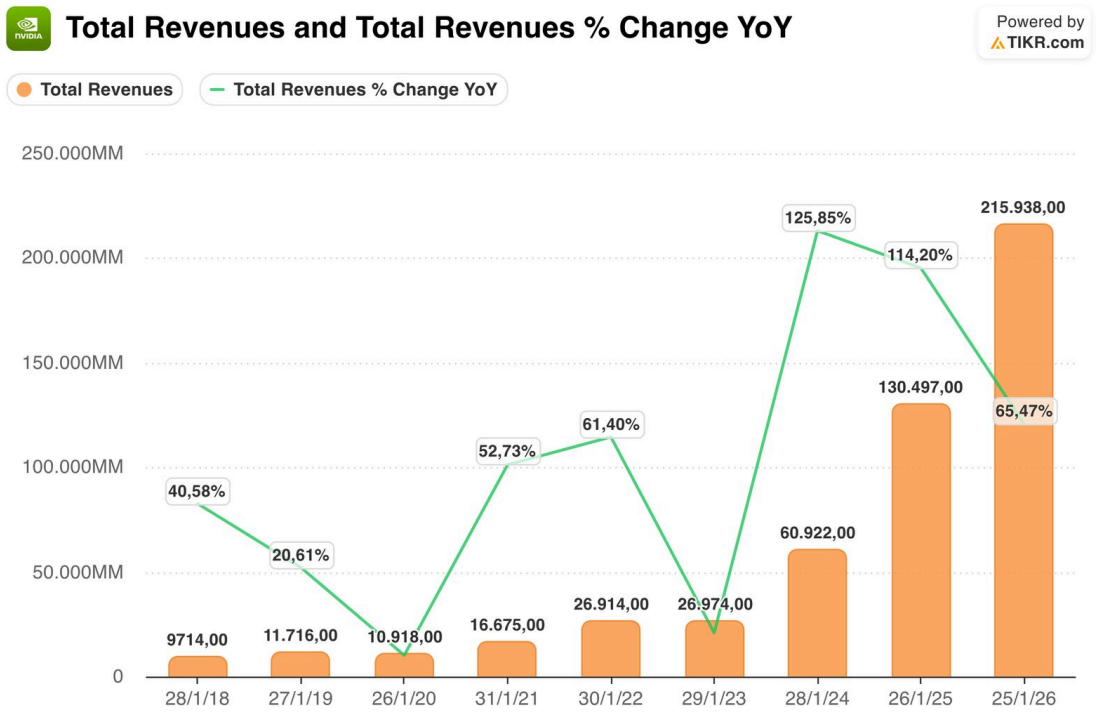
Segment 4: Automotive

Automotive is the smallest segment but one with high growth: **$2.3B in FY2026 (+39% YoY)**, driven by the adoption of the DRIVE Orin and DRIVE Thor platforms in production vehicles. Toyota, Mercedes-Benz, Volvo, BYD and dozens of manufacturers have integrated Nvidia chips into their driver assistance systems. If robotaxis scale as an industry, Nvidia has a position in every vehicle, regardless of the manufacturer.

SECTION 03

Business quality

Revenue



Source: Tikr.

Nvidia's revenue trajectory has no parallel in the history of any semiconductor company. In FY2018 the company generated $9.7B. In FY2023, $26.9B. The jump began in FY2024 with $60.9B (+126% YoY) and accelerated in FY2025 with $130.5B (+114% YoY). **In FY2026 revenue reached $215.9B, 65% growth on a base that was already a record.** In absolute terms, Nvidia added more than $85B of revenue in a single fiscal year.

To put those numbers in perspective: Intel, the semiconductor company that dominated computing for decades, generated $54B in 2024. Nvidia, in the fourth quarter of FY2026 alone, had $68B of revenue. A single Nvidia quarter is now larger than Intel's entire year.

The composition of growth also matters. Q4 FY2026's $68.1B came from $62.3B in Data Center (+75% YoY) and $6.5B in Gaming (+97% YoY). In Q1 FY2027 Nvidia beat that level again with $81.6B of revenue, and in Q2 it reached **$96.2B, +106% year over year**. For the third quarter it guided to $108B, which would imply a FY2027 of around $400 billion, almost double FY2026. Most striking, that guidance assumes no Data Center Compute revenue from China.

Latest results: second quarter of FY2027

On August 26 Nvidia reported results for the quarter ended July 2026, the latest available at the time of this analysis (third-quarter results come out at the end of November):

- **Revenue:** $96.2B, +18% sequentially and +106% year over year.
- **Data Center:** $89.0B, +117% year over year. Under the new reporting, the rest of the business (gaming, workstations, robotics, automotive) is grouped into **Edge Computing**, with $7.2B (+27%).
- **Gross margin:** 75.0%, up from 72.4% a year earlier.
- **Operating income:** $63.7B (+124%), a 66% margin. **Net income:** $59.7B.
- **Shareholder returns:** about $26 billion between buybacks and dividends in the quarter, with $99 billion still available under the buyback authorization.
- **Vera Rubin:** already in full production, with racks running at CoreWeave, Google Cloud, Microsoft Azure, Oracle Cloud and Nebius.

For the third quarter the company expects revenue of **$108 billion (±2%)** and a 74% gross margin, excluding Data Center Compute sales to China. Execution remains flawless. The stock, however, trades around $187 today, below its highs, amid market doubts about the sustainability of AI spending.

The history of Nvidia's revenue between FY2023 and FY2026 is also the history of the launch of ChatGPT: that November 2022 event unlocked a demand for AI computing capacity that no capacity planning model had contemplated. Hyperscalers (Microsoft, Google, Amazon, Meta) had to rush to buy GPUs so as not to fall behind. Nvidia was the only supplier with the hardware and software ecosystem ready to scale.

Vertical analysis

INCOME STATEMENT

Fiscal year Income Statement — % of revenue	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Total Revenues	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
Cost of Goods Sold	(41.2%)	(40.1%)	(38.8%)	(38.0%)	(36.7%)	(35.1%)	(43.1%)	(27.3%)	(25.0%)	(28.9%)	**(27.1%)**	**(31.9%)**
Gross Profit	**58.8%**	**59.9%**	**61.2%**	**62.0%**	**63.3%**	**64.9%**	**56.9%**	**72.7%**	**75.0%**	**71.1%**	**72.9%**	**68.1%**
Selling General & Admin Expenses	(9.6%)	(8.4%)	(8.5%)	(10.0%)	(11.5%)	(8.0%)	(9.0%)	(4.4%)	(2.7%)	(2.1%)	**(3.1%)**	**(5.2%)**
R&D Expenses	(21.2%)	(18.5%)	(20.3%)	(25.9%)	(23.5%)	(19.6%)	(27.2%)	(14.2%)	(9.9%)	(8.6%)	**(10.9%)**	**(15.9%)**
Total Operating Expenses	**(30.8%)**	**(26.9%)**	**(28.7%)**	**(35.9%)**	**(35.0%)**	**(27.6%)**	**(36.3%)**	**(18.6%)**	**(12.6%)**	**(10.7%)**	**(14.0%)**	**(21.1%)**
Operating Income	**28.0%**	**33.0%**	**32.5%**	**26.1%**	**28.3%**	**37.3%**	**20.7%**	**54.1%**	**62.4%**	**60.4%**	**59.0%**	**47.0%**
Interest Expense	(0.8%)	(0.6%)	(0.5%)	(0.5%)	(1.1%)	(0.9%)	(1.0%)	(0.4%)	(0.2%)	(0.1%)	**(0.2%)**	**(0.5%)**
EBT Incl. Unusual Items	**27.6%**	**32.9%**	**33.3%**	**27.2%**	**26.4%**	**36.9%**	**15.5%**	**55.5%**	**64.4%**	**65.5%**	**61.8%**	**47.6%**
Income Tax Expense	(3.5%)	(1.5%)	2.1%	(1.6%)	(0.5%)	(0.7%)	0.7%	(6.7%)	(8.5%)	(9.9%)	**(8.4%)**	**(5.0%)**
Net Income	**24.1%**	**31.4%**	**35.3%**	**25.6%**	**26.0%**	**36.2%**	**16.2%**	**48.8%**	**55.8%**	**55.6%**	**53.4%**	**42.5%**
EBITDA	30.7%	35.1%	34.7%	29.6%	34.9%	41.7%	26.4%	56.6%	63.8%	61.7%	**60.7%**	**50.0%**

Source: author's analysis based on Tikr data.

The vertical analysis of Nvidia's income statement reveals something extraordinary: a hardware company with software margins. **FY2026 gross margin was 71%**, comparable to companies like Microsoft or Salesforce, not to chipmakers that historically operate at 40-55%.

Gross margin was not always this high. In FY2019 and FY2020 it was 61-62%, already solid for the industry but more "normal". The jump came with the AI boom: in FY2024 margin rose to 72.7%, an unprecedented level for a company of its scale, driven by the extraordinary pricing power of being

the only supplier with the product and ecosystem the market urgently needed.

The slight compression to 71.1% in FY2026 (vs. 75.0% in FY2025) reflected the architecture transition, a greater weight of complete systems and one-off charges related to export restrictions. FY2027 results confirm that reading: second-quarter gross margin was 75.0% and third-quarter guidance is 74%, above the FY2026 annual margin. Gross margin remains one of the most important KPIs to monitor: if it falls persistently toward 65% or less, it could signal that AMD, ASICs or hyperscalers are capturing part of the bargaining power.

Nvidia's **operating margin** exceeded **60% in FY2026**, making it one of the most profitable companies relative to revenue in the whole technology industry. Selling and administrative expenses weighed only about 2-4% of sales over the last three years: Nvidia does not need an army of salespeople, because hyperscalers compete to get GPUs, not the other way around.

BALANCE SHEET

Fiscal year Balance Sheet — % of total assets	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Total Cash And Short Term Investments	**69.1%**	**63.2%**	**55.8%**	**62.9%**	**40.2%**	**48.0%**	**32.3%**	**39.5%**	**38.7%**	**30.2%**	**36.2%**	**37.8%**
Total Receivables	**8.4%**	**11.3%**	**10.7%**	**9.6%**	**8.4%**	**10.5%**	**9.3%**	**15.2%**	**20.7%**	**18.6%**	**18.2%**	**14.9%**
Inventory	8.1%	7.1%	11.8%	5.7%	6.3%	5.9%	12.5%	8.0%	9.0%	10.3%	**9.1%**	**9.2%**
Total Current Assets	**86.7%**	**82.3%**	**79.4%**	**79.1%**	**55.8%**	**65.2%**	**56.0%**	**67.5%**	**71.8%**	**60.7%**	**66.7%**	**64.3%**
Net Property Plant And Equipment	5.3%	8.9%	10.6%	13.2%	9.9%	8.2%	11.8%	8.0%	7.2%	6.4%	**7.2%**	**8.3%**
Long-term Investments	–	–	–	0.4%	0.5%	0.6%	0.7%	2.0%	3.0%	10.8%	**5.3%**	**3.4%**
Goodwill	6.3%	5.5%	4.6%	3.6%	14.6%	9.8%	10.6%	6.7%	4.6%	10.1%	**7.2%**	**8.4%**
Other Intangibles	1.1%	0.5%	0.3%	0.3%	9.5%	5.3%	4.1%	1.7%	0.7%	1.6%	**1.3%**	**2.7%**
Total Assets	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
Total Current Liabilities	**18.2%**	**10.3%**	**10.0%**	**10.3%**	**13.6%**	**9.8%**	**15.9%**	**16.2%**	**16.2%**	**15.6%**	**16.0%**	**14.7%**
Long-Term Debt	20.2%	17.7%	15.0%	11.5%	20.7%	24.8%	23.6%	12.9%	7.6%	3.6%	**8.0%**	**14.5%**
Capital Leases	–	–	–	3.2%	2.2%	1.7%	2.2%	1.7%	1.4%	1.2%	**1.4%**	**1.6%**
Total Liabilities	**41.4%**	**33.5%**	**29.7%**	**29.5%**	**41.3%**	**39.8%**	**46.3%**	**34.6%**	**28.9%**	**23.9%**	**29.2%**	**34.7%**
Total Equity	**58.6%**	**66.5%**	**70.3%**	**70.5%**	**58.7%**	**60.2%**	**53.7%**	**65.4%**	**71.1%**	**76.1%**	**70.8%**	**65.3%**
Total Debt	**28.3%**	**17.8%**	**15.0%**	**15.3%**	**26.4%**	**26.8%**	**29.2%**	**16.8%**	**8.9%**	**5.5%**	**10.4%**	**17.5%**
Net Debt	(40.8%)	(45.4%)	(40.9%)	(47.7%)	(13.8%)	(21.2%)	(3.1%)	(22.7%)	(29.8%)	(24.7%)	**(25.7%)**	**(20.3%)**

Source: author's analysis based on Tikr data.

Nvidia's balance sheet shows the transformation of a moderately indebted company with a conventional asset structure into a cash-generating machine with equity representing 76% of total assets.

Assets: the dominance of cash

The most striking feature over the decade is the absolute dominance of current assets, driven mainly by cash and short-term investments. In FY2017 cash and short-term investments were 69% of total assets. The share fluctuated over the years, hitting a relative low of 32% in FY2023, and then recovered. **In FY2026 cash and short-term investments represent 30% of total assets**, which in absolute terms equals tens of billions of dollars given the current size of the balance sheet.

Net property, plant and equipment tells a story consistent with the fabless model: **it never exceeded 13% of total assets in any year**, and in FY2026 it is just 6.4%. Nvidia designs chips but manufactures nothing, delegating industrial infrastructure to TSMC and its manufacturing partners. The result is a balance sheet that needs no physical anchors to generate extraordinary returns.

Inventory rose notably in recent years, from a historical 5-8% to **10.4% of assets in FY2026**, a direct reflection of the massive supply commitments with TSMC and HBM suppliers to guarantee future capacity. This is the asset item that most deserves monitoring: if demand slows suddenly, that inventory could become a significant charge, as already happened with the H20 episode in China.

Liabilities: structural deleveraging

In FY2017 total liabilities represented 41% of assets; in FY2026 the figure fell to **24%**. Long-term debt went from 20% of assets to just **3.6%**, reflecting both the deliberate repayment of obligations and the decision not to refinance once its own cash generation made it unnecessary.

Equity: silent accumulation

Equity went from 58.6% of assets in FY2017 to 76.1% in FY2026, driven fundamentally by retained earnings. Nvidia's balance sheet is that of a financially armored company: **little debt, abundant cash, dominant equity and minimal physical assets**. The remaining risks are not financial but strategic: the question is not whether Nvidia can pay its obligations, but whether it can defend its market position against its own customers' ASICs.

CASH FLOW

Fiscal year Cash Flow — % of revenue	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Net Income	**24.1%**	**31.4%**	**35.3%**	**25.6%**	**26.0%**	**36.2%**	**16.2%**	**48.8%**	**55.8%**	**55.6%**	**53.4%**	**42.5%**
Total Depreciation & Amortization	**2.7%**	**2.0%**	**2.2%**	**3.5%**	**6.6%**	**4.4%**	**5.7%**	**2.5%**	**1.4%**	**1.3%**	**1.7%**	**3.1%**
Stock-Based Compensation	3.6%	4.0%	4.8%	7.7%	8.4%	7.4%	10.0%	5.8%	3.6%	3.0%	**4.1%**	**6.0%**
Cash from Operations	**24.2%**	**36.1%**	**31.9%**	**43.6%**	**34.9%**	**33.8%**	**20.9%**	**46.1%**	**49.1%**	**47.6%**	**47.6%**	**39.5%**
Capital Expenditure	(2.5%)	(6.1%)	(5.1%)	(4.5%)	(6.8%)	(3.6%)	(6.8%)	(1.8%)	(2.5%)	(2.8%)	**(2.3%)**	**(3.5%)**
Cash Acquisitions	–	–	–	(0.0%)	(51.1%)	(1.0%)	(0.2%)	(0.1%)	(0.8%)	(0.7%)	**(0.5%)**	**(0.6%)**
Cash from Investing	**(11.5%)**	**13.2%**	**(35.0%)**	**56.3%**	**(118.0%)**	**(36.5%)**	**27.3%**	**(17.3%)**	**(15.6%)**	**(24.2%)**	**(19.1%)**	**(13.3%)**
Total Debt Issued	**28.8%**	–	–	–	**29.8%**	**18.5%**	–	–	–	–	–	**3.7%**
Total Debt Repaid	**(9.7%)**	**(8.4%)**	**(0.1%)**	–	–	**(3.7%)**	–	**(2.1%)**	**(1.0%)**	–	**(1.0%)**	**(1.3%)**
Repurchase of Common Stock	(13.2%)	(15.7%)	(22.3%)	(5.0%)	(5.6%)	(7.1%)	(42.7%)	(20.2%)	(31.1%)	(22.2%)	**(24.5%)**	**(24.7%)**
Common Dividends Paid	(3.8%)	(3.5%)	(3.2%)	(3.6%)	(2.4%)	(1.5%)	(1.5%)	(0.6%)	(0.6%)	(0.5%)	**(0.6%)**	**(0.9%)**
Cash from Financing	**4.2%**	**(26.2%)**	**(24.5%)**	**(7.3%)**	**22.8%**	**6.9%**	**(43.1%)**	**(22.4%)**	**(32.5%)**	**(22.4%)**	**(25.8%)**	**(22.7%)**
Free Cash Flow	21.6%	29.9%	26.8%	39.1%	28.1%	30.2%	14.1%	44.4%	46.6%	44.8%	**45.3%**	**36.0%**

Source: author's analysis based on Tikr data.

Nvidia's operating cash flow was never negative in any year of the period. **In FY2017 it represented 24.2% of revenue; in FY2026 it reached 47.6%**, almost double. Net income as a percentage of revenue rose from 24% to **55.6% in FY2026**, the highest in its history: Nvidia turns more than half of every dollar of revenue into net profit.

CapEx of 2.8% of revenue in FY2026 confirms the fabless nature of the business. For a company that generated 47% of its revenue in operating cash, spending less than 3% on CapEx produces a **free cash flow of 44.8% of revenue**, the highest of the series and one of the highest of any technology company of its scale.

Buybacks are the dominant use of cash: they represented 22% of revenue in FY2026, more than $48B. Shareholder returns became even more aggressive in FY2027: in the second quarter alone Nvidia returned about $26 billion between buybacks and dividends, and expanded its buyback authorization to $99 billion available. The message is clear: Nvidia generates so much cash that it can fund growth, secure supply chain capacity, invest in software and networking, and still return tens of billions of dollars per quarter to shareholders.

Financial ratios

- **Growth**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Growth												
Revenue growth		40.6%	20.6%	(6.8%)	52.7%	61.4%	0.2%	125.9%	114.2%	65.5%	**101.8%**	**73.4%**
EBIT growth		65.7%	18.5%	(25.2%)	65.9%	112.7%	(44.5%)	491.2%	147.0%	60.1%	**232.8%**	**153.3%**
Diluted EPS growth		100.0%	41.7%	(35.3%)	54.5%	129.4%	(56.4%)	600.0%	147.1%	66.7%	**271.2%**	**177.3%**
Change in diluted shares (dilution / buybacks)		(2.6%)	(1.1%)	(1.1%)	1.5%	1.0%	(1.1%)	(0.5%)	(0.5%)	(1.2%)	**(0.7%)**	**(0.5%)**

Source: author's analysis based on Tikr data.

Sales grew 126% in FY2024, 114% in FY2025 and 65% in FY2026, and are approaching 90% again in FY2027. EPS grew even faster because of operating leverage, and the share count falls slightly thanks to buybacks.

- **Profitability**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Profitability												
Gross margin	58.8%	59.9%	61.2%	62.0%	63.3%	64.9%	56.9%	72.7%	75.0%	71.1%	**72.9%**	**68.1%**
SG&A % revenue	9.6%	8.4%	8.5%	10.0%	11.5%	8.0%	9.0%	4.4%	2.7%	2.1%	**3.1%**	**5.2%**
R&D % revenue	21.2%	18.5%	20.3%	25.9%	23.5%	19.6%	27.2%	14.2%	9.9%	8.6%	**10.9%**	**15.9%**
EBIT margin	28.0%	33.0%	32.5%	26.1%	28.3%	37.3%	20.7%	54.1%	62.4%	60.4%	**59.0%**	**47.0%**
EBITDA margin (EBIT + D&A)	30.7%	35.1%	34.7%	29.6%	34.9%	41.7%	26.4%	56.6%	63.8%	61.7%	**60.7%**	**50.0%**
Net margin	24.1%	31.4%	35.3%	25.6%	26.0%	36.2%	16.2%	48.8%	55.8%	55.6%	**53.4%**	**42.5%**
Stock-based compensation % revenue	3.6%	4.0%	4.8%	7.7%	8.4%	7.4%	10.0%	5.8%	3.6%	3.0%	**4.1%**	**6.0%**
Effective tax rate	12.5%	4.7%	(6.3%)	5.9%	1.7%	1.9%	(4.5%)	12.0%	13.3%	15.1%	**13.5%**	**7.6%**

Source: author's analysis based on Tikr data.

Gross margin went from 56.9% in FY2023 to 75% in FY2025, with a slight dip to 71.1% in FY2026 that has already reversed in FY2027. EBIT margin exceeds 60% and net margin 55%: software margins in a hardware company. Overhead and R&D have been massively diluted (2.1% and 8.6% of sales) and stock-based compensation fell to 3%.

- **Cash flow and reinvestment**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Cash flow & reinvestment												
Cash from operations % revenue	24.2%	36.1%	31.9%	43.6%	34.9%	33.8%	20.9%	46.1%	49.1%	47.6%	**47.6%**	**39.5%**
FCF margin (TIKR FCF / revenue)	21.6%	29.9%	26.8%	39.1%	28.1%	30.2%	14.1%	44.4%	46.6%	44.8%	**45.3%**	**36.0%**
FCF conversion (FCF / net income)	89.8%	95.5%	75.9%	152.8%	108.4%	83.4%	87.2%	90.8%	83.5%	80.5%	**84.9%**	**85.1%**
Capex % revenue	2.5%	6.1%	5.1%	4.5%	6.8%	3.6%	6.8%	1.8%	2.5%	2.8%	**2.3%**	**3.5%**
D&A % revenue	2.7%	2.0%	2.2%	3.5%	6.6%	4.4%	5.7%	2.5%	1.4%	1.3%	**1.7%**	**3.1%**
Capex / D&A	0.9x	3.0x	2.3x	1.3x	1.0x	0.8x	1.2x	0.7x	1.7x	2.1x	**1.5x**	**1.3x**
D&A % beginning net PP&E		38.2%	26.3%	27.1%	47.9%	41.1%	42.8%	31.1%	35.4%	35.2%	**33.9%**	**37.1%**
Acquisitions % revenue	–	–	–	0.0%	51.1%	1.0%	0.2%	0.1%	0.8%	0.7%	**0.5%**	**0.6%**

Source: author's analysis based on Tikr data.

Free cash flow was 44.8% of sales in FY2026, with 80% conversion of net income. CapEx is only 2.8% of sales: the fabless model allows growth without large physical investments.

- **Working capital**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Working capital												
Receivables days (DSO)	43.6	47.5	44.4	55.4	53.2	63.1	51.8	59.9	64.5	65.0	**63.1**	**60.9**
Inventory days (DIO)	101.8	74.7	126.5	86.1	108.9	100.7	162.1	116.0	112.7	125.0	**117.9**	**123.3**
Payables days (DPO)	62.2	55.9	41.0	60.4	68.5	68.9	37.5	59.3	70.6	57.3	**62.4**	**58.7**
Cash conversion cycle (DSO + DIO − DPO)	83.2	66.3	129.8	81.1	93.6	94.8	176.4	116.6	106.7	132.7	**118.7**	**125.5**
Deferred revenue % revenue	1.3%	0.7%	1.2%	1.8%	2.7%	1.9%	2.1%	2.2%	1.4%	1.2%	**1.6%**	**1.8%**
Operating NWC % revenue	10.9%	10.6%	15.8%	10.1%	8.4%	12.5%	18.1%	14.7%	14.4%	15.6%	**14.9%**	**15.1%**

Source: author's analysis based on Tikr data.

This is the only area where Nvidia consumes meaningful cash: inventory rose to about 125 days and the cash conversion cycle to more than 130 days, because of supply commitments with TSMC and HBM memory suppliers. Operating working capital is around 15% of sales.

- **Returns on capital**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Returns on capital												
NOPAT = EBIT × (1 − normalized tax rate)	1,680	2,784	3,299	2,468	4,095	8,709	4,837	28,598	70,648	113,091	**70,779**	**45,177**
Invested capital = debt + equity − cash & ST inv.	1,747	2,363	3,908	3,950	12,929	17,235	20,836	28,050	46,099	106,149	**60,099**	**43,674**
ROIC = NOPAT / average invested capital		135.5%	105.2%	62.8%	48.5%	57.7%	25.4%	117.0%	190.6%	148.6%	**152.0%**	**107.9%**
ROIC incl. cash = NOPAT / average (debt + equity)		30.9%	31.7%	18.9%	20.8%	27.7%	13.3%	64.9%	98.6%	87.7%	**83.7%**	**58.4%**
Pre-tax ROIC = EBIT / average invested capital (reference)		156.2%	121.3%	72.4%	55.9%	66.6%	29.3%	134.9%	219.7%	171.3%	**175.3%**	**124.4%**
ROE = net income / average equity		46.1%	49.3%	26.0%	29.8%	44.8%	17.9%	91.5%	119.2%	101.5%	**104.0%**	**75.0%**
ROA = net income / average total assets		28.9%	33.8%	18.3%	18.8%	26.7%	10.2%	55.7%	82.2%	75.4%	**71.1%**	**50.0%**

Source: author's analysis based on Tikr data.

ROIC was 149% in FY2026 (191% in FY2025), and 88% including cash in invested capital. These levels have practically no comparison among companies of this size. ROE exceeds 100%.

- **Leverage and liquidity**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Leverage & liquidity												
Net debt (debt − cash & ST inv.)	(4,015)	(5,108)	(5,434)	(8,254)	(3,964)	(9,377)	(1,265)	(14,928)	(33,228)	(51,144)	**(33,100)**	**(21,988)**
Net debt / EBITDA	(1.9x)	(1.5x)	(1.3x)	(2.6x)	(0.7x)	(0.8x)	(0.2x)	(0.4x)	(0.4x)	(0.4x)	**(0.4x)**	**(0.4x)**
Interest coverage (EBIT / interest expense)	33.4x	52.6x	65.6x	54.7x	25.7x	42.5x	21.3x	128.3x	329.8x	503.4x	**320.5x**	**205.1x**
Debt / (debt + equity), book	32.6%	21.1%	17.5%	17.8%	31.0%	30.8%	35.2%	20.5%	11.2%	6.8%	**12.8%**	**20.9%**
Current ratio	4.77	8.03	7.94	7.67	4.09	6.65	3.52	4.17	4.44	3.91	**4.17**	**4.54**

Source: author's analysis based on Tikr data.

Nvidia had net cash of more than $51 billion at the end of FY2026, interest coverage of 500x and a current ratio of 3.9x. The balance sheet is not a risk in any reasonable scenario.

- **Shareholder returns and valuation**

Fiscal year	FY2017A	FY2018A	FY2019A	FY2020A	FY2021A	FY2022A	FY2023A	FY2024A	FY2025A	FY2026A	3Y avg	5Y avg
Shareholder returns & market valuation (TIKR, at fiscal year end)												
Dividends + buybacks % FCF	78.6%	64.0%	94.9%	22.0%	28.5%	28.3%	312.8%	47.0%	68.1%	50.7%	**55.3%**	**101.4%**
EV / EBITDA	26.6x	42.0x	22.6x	45.3x	54.8x	50.2x	70.2x	43.5x	41.6x	33.9x	**39.6x**	**47.9x**
EV / revenue	8.2x	14.7x	7.9x	13.4x	19.1x	20.9x	18.5x	24.6x	26.5x	20.9x	**24.0x**	**22.3x**
P/E	36.2x	48.4x	23.6x	54.8x	74.2x	58.6x	114.7x	50.7x	47.9x	38.0x	**45.5x**	**62.0x**
FCF yield (FCF / market cap)	2.5%	2.0%	3.2%	2.8%	1.5%	1.4%	0.8%	1.8%	1.7%	2.1%	**1.9%**	**1.6%**

Source: author's analysis based on Tikr data.

In FY2026 Nvidia returned about 50% of its free cash flow to shareholders, almost all through buybacks. At year end it traded at 34x EBITDA and 38x earnings, with a 2.1% FCF yield. With the stock's decline and the growth in FY2027 results, multiples on current earnings have come down significantly.

SECTION 04

Management and capital allocation

Jensen Huang, CEO and co-founder

Source: Nvidia.

Jensen is not simply the CEO of a successful company: **he is the architect of one of the most extraordinary strategic pivots in modern corporate history**, the man who took a video game company and turned it into the infrastructure on which global artificial intelligence is being built.

He was born on February 17, 1963 in Tainan, Taiwan. His childhood was marked by hardship: when he was 9, his parents sent him and his brother to live with an uncle in Tacoma, Washington. A mix-up landed him at the Oneida Baptist Institute in Kentucky, a religious reform school for troubled youth, where he cleaned bathrooms every day and was systematically bullied. That period forged a character that people who know him describe as resilient to the point of stubbornness.

He graduated in electrical engineering from Oregon State University in 1984 and completed his master's at Stanford. He worked at AMD and LSI Logic before founding Nvidia in April 1993 with Chris Malachowsky and Curtis Priem. They started at a Denny's, sketching the first ideas of Nvidia's architecture on a napkin.

The first big product was the **RIVA 128 in 1997**, and they almost did not make it. The previous chip had bet on an alternative method of processing triangles, and Microsoft decided DirectX would only support conventional triangles. Nvidia was weeks away from bankruptcy. Jensen redirected all resources to the RIVA 128 and launched it in record time. It was the first of many existential bets he won.

In 2006 he launched **CUDA** (Compute Unified Device Architecture), a software platform that allowed GPUs to be used for general-purpose computing beyond graphics. At the time almost nobody understood what it was for, and investors asked why the company was spending on something with no apparent market. Jensen used the term "zero-billion-dollar market" to describe bets worth making: so large that competitors did not bother to enter.

The turning point came in 2012 when **AlexNet**, the first modern deep learning model, developed by Geoffrey Hinton and his students in Toronto, won the ImageNet competition by a wide margin. It had been trained on Nvidia GPUs. Jensen immediately understood the implications, and from that moment Nvidia redirected the whole company toward AI.

In 2016 he personally delivered the first **DGX-1**, the first supercomputer designed specifically for AI, to a special recipient: OpenAI, then a small nonprofit in San Francisco.

What Jensen built over 30 years is one of the deepest technology platforms in history: CUDA, cuDNN, TensorRT, NeMo, Omniverse, DRIVE, the network of university researchers, the Deep Learning Institute courses and the partnerships with hyperscalers. **When the world needed AI infrastructure,**

Jensen had already been building it for 15 years before anyone asked for it. That is the essence of a founder-CEO with a long-term vision.

Jensen is 63 today and as intense as ever. His management style is known for being extremely demanding. He dislikes hierarchies: he prefers employees to send him short emails about what they are doing that week rather than formal reporting systems. His management maxim is to always operate at the "speed of light": calculate the minimum possible time to do something without constraints, and then work to get close to that ideal.

The question every investor must ask is: **what happens to Nvidia without Jensen Huang?** The honest answer is that nobody knows, because Jensen has no publicly identified successor. It is the largest key-person concentration risk of any company of its size.

Colette Kress, CFO

Colette Kress has been Nvidia's CFO since 2013, after serving as CFO at Cisco's business unit level and earlier at Texas Instruments. She is recognized as one of the strongest finance executives in the technology sector: her communication with investors is extraordinarily clear and precise, and she handles the company's volatility with remarkable consistency. The quality of Nvidia's quarterly CFO Commentary, which breaks down revenue by segment and geography in more detail than its peers, is largely her work.

SECTION 05

Competitive position

Industry analysis

The AI chip market is one of the fastest-growing markets today. The AI CapEx of the four big hyperscalers (Microsoft, Google, Amazon and Meta) exceeded $250B in 2025, and commitments for 2026 and 2027 are even larger. A significant fraction of that spending goes directly to Nvidia.

The global market for AI semiconductors was about $100B in 2024 and is projected to exceed $300B by 2028, with annual growth of 30-40% depending on the source and scenario. The most uncertain factor in those projections is the share that goes to GPUs versus custom ASICs.

The most important variable for understanding this market is the difference between **training and inference**. Training a model requires processing massive amounts of data to adjust billions of parameters: a workload that needs maximum flexibility and computing power, and Nvidia GPUs are by consensus the best tool available. Inference, by contrast, means running an already-trained model to answer questions: a more predictable, repetitive workload that is sensitive to cost per token. That is where ASICs (application-specific chips) have a natural advantage.

This distinction is key for Nvidia's future: if the AI market matures toward massive inference and ASICs capture a relevant part of it, Nvidia can keep its dominance in training and flexible workloads but lose part of the incremental growth in optimized inference. That is the most concrete risk scenario investors must consider.

The ASIC threat is no longer a distant possibility but a concrete dynamic. Anthropic expanded agreements with both Google/Broadcom and Amazon: with Google and Broadcom it secured access to multiple gigawatts of next-generation TPU capacity from 2027, and with Amazon it committed more than $100B over ten years for up to 5 GW of capacity on AWS, including Trainium. This confirms that the leading AI labs do not want to depend on a single accelerator supplier. Nvidia remains the most flexible and dominant platform, but production inference is starting to fragment between GPUs, TPUs and ASICs.

I do not think ASICs will "kill" Nvidia, but they can limit the size of the market Nvidia can really capture in the mature phase of AI. For training, research, new models and changing workloads, the flexibility of GPU + CUDA remains a huge advantage. For stable, repetitive inference at massive scale, ASICs can be more cost-efficient per token. The future is probably hybrid.

The CUDA moat: the hardest fortress to besiege

Before CUDA, GPUs had thousands of cores capable of parallel calculations, but programming them was extremely difficult and required very specific graphics knowledge. CUDA was the universal translator. It created a programming interface similar to C that let any programmer use those thousands of cores without knowing anything about graphics. Suddenly physicists, biologists, economists and engineers could use Nvidia GPUs for their own calculations, and AI researchers, who needed exactly that kind of massive parallelism, began building on that foundation.

The result is an ecosystem that today has more than **4 million developers** writing CUDA code. Every university that teaches machine learning teaches CUDA. The two most widely used AI frameworks in the world, PyTorch and TensorFlow, use CUDA internally.

On top of CUDA, Nvidia built an ecosystem of specialized libraries: **cuDNN** optimizes neural network operations, **TensorRT** optimizes trained models to respond faster in production, and **NCCL** coordinates communication between thousands of GPUs working in parallel. All are free, which is not generosity but strategy: every developer who uses them becomes more dependent on Nvidia hardware.

What makes the CUDA moat almost invulnerable is not the code itself but the ecosystem inertia accumulated over 20 years. Migrating all that production code to an alternative like ROCm (AMD's equivalent) is not a technical decision: it is a multi-year project with the risk of performance regression at every step, and customers spending billions on CapEx cannot afford that luxury.

Main competitors

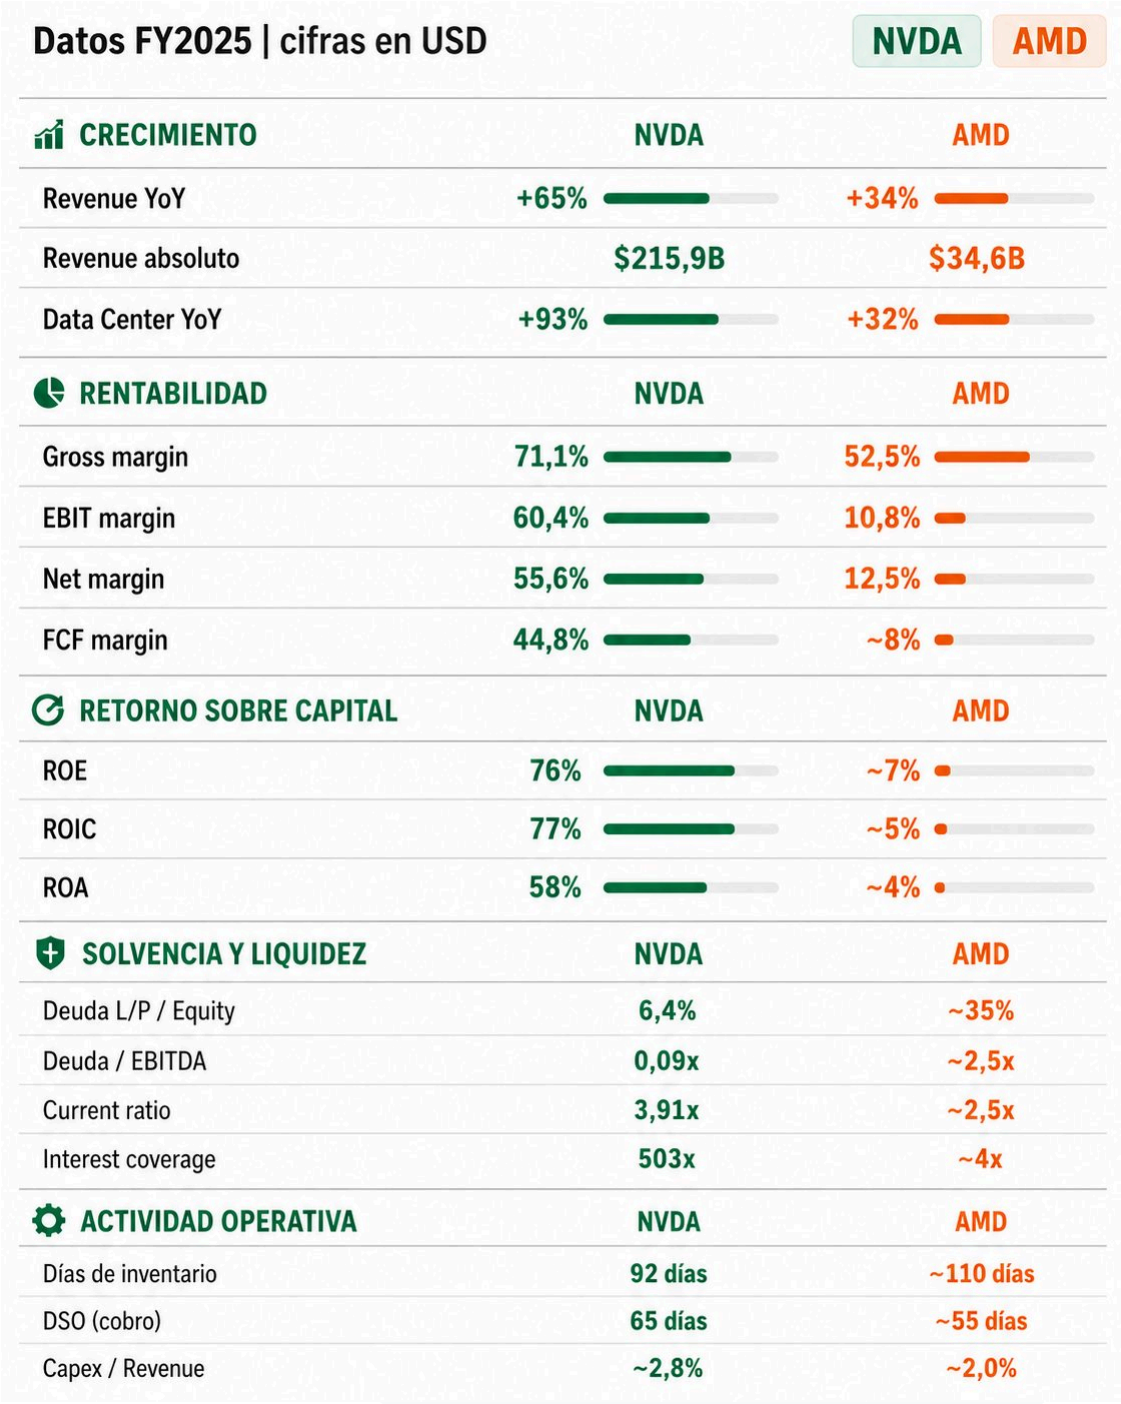
AMD (ADVANCED MICRO DEVICES)

AMD is Nvidia's most direct rival in AI GPUs and the only chipmaker offering a general-purpose alternative at scale. Its **Instinct MI300X** and MI350 are the first AMD chips that some hyperscalers are deploying as a partial alternative to Nvidia's H100/H200.

AMD's limitation is not the hardware, which in certain memory and bandwidth benchmarks is competitive, but ROCm, its equivalent of CUDA, which still has far less library and framework support. Customers who need certainty of results for billions of dollars of CapEx cannot afford to experiment with unproven alternatives. That gives Nvidia time to consolidate its advantage while AMD builds its ecosystem.

The MI450 ("Helios") planned for the second half of 2026 and the MI500 for 2027 are AMD's bet to close the gap. Lisa Su has shown with x86 that she can win back lost market share given enough time. In 2025 AMD had revenue of $34.6B (+34%), with records in Data Center and Client, and convinced the market it can scale its AI business, with OpenAI as the most significant milestone: the AI lab named AMD a preferred partner to deploy 6 GW of GPUs, the first 1 GW based on MI450. AMD also expanded its relationship with Meta, reinforcing the idea that hyperscalers want to fund a second source of accelerators to reduce their dependence on Nvidia.

NVIDIA vs AMD:



Datos FY2025 | cifras en USD	NVDA	AMD
CRECIMIENTO	**NVDA**	**AMD**
Revenue YoY	+65%	+34%
Revenue absoluto	$215,9B	$34,6B
Data Center YoY	+93%	+32%
RENTABILIDAD	**NVDA**	**AMD**
Gross margin	71,1%	52,5%
EBIT margin	60,4%	10,8%
Net margin	55,6%	12,5%
FCF margin	44,8%	~8%
RETORNO SOBRE CAPITAL	**NVDA**	**AMD**
ROE	76%	~7%
ROIC	77%	~5%
ROA	58%	~4%
SOLVENCIA Y LIQUIDEZ	**NVDA**	**AMD**
Deuda L/P / Equity	6,4%	~35%
Deuda / EBITDA	0,09x	~2,5x
Current ratio	3,91x	~2,5x
Interest coverage	503x	~4x
ACTIVIDAD OPERATIVA	**NVDA**	**AMD**
Días de inventario	92 días	~110 días
DSO (cobro)	65 días	~55 días
Capex / Revenue	~2,8%	~2,0%

Source: author's analysis based on Tikr data.

Nvidia grew **65% on a base of $130B of revenue**, adding more than $85B in a single fiscal year. AMD grew **34% on a base of $25B**. A single year of Nvidia's incremental growth equals two and a half times AMD's entire business.

In profitability the gap is **enormous**. Nvidia turns 71% of every dollar of revenue into gross profit, 60% into operating income and 45% into free cash flow. AMD operates with a 52% gross margin, an 11% EBIT margin and an FCF margin close to 8%. They are two business models with radically differ-

ent profitability structures, and the difference is not only scale but **pricing power**: Nvidia can charge what it wants because it has no real substitute, while AMD still partly competes on price. The return on capital comparison sums it up: Nvidia generates economic value at a speed without comparison in the industry, while AMD builds the foundations to compete in the next phase of the AI cycle.

GOOGLE (TPU IRONWOOD)

Google was the pioneer in AI ASICs. Its TPU has gone through several generations, and the latest, **TPU Ironwood (v7)**, offers performance comparable to Blackwell in certain workloads, with its main advantage being energy efficiency. Its limitation is that it is only available on Google Cloud. The most relevant data point in the TPU ecosystem is Anthropic, which expanded its relationship with Google and Broadcom with an agreement for multiple gigawatts of next-generation TPU capacity from 2027. ASICs are no longer internal lab projects: they are strategic infrastructure for one of the main foundation model developers in the world.

AMAZON WEB SERVICES (TRAINIUM)

AWS has two families of custom AI chips: **Trainium for training and Inferentia for inference**. Trainium3 was launched at the end of 2025, and Anthropic trains models on hundreds of thousands of Trainium2 chips. Anthropic signed a new agreement with Amazon for more than $100B over the next ten years. Hyperscalers do not necessarily want to replace Nvidia completely, but they do want a second source of compute to capture margin, reduce dependence and optimize specific workloads.

META, MICROSOFT, OPENAI AND THE ASICS OF THE FUTURE

Meta is developing **MTIA** for its recommendation and generative workloads. Microsoft has Maia for Azure and Copilot. OpenAI, with Broadcom as design partner, is launching its own chip. The pattern is universal: every hyperscaler that can afford the R&D is trying to reduce its dependence on Nvidia.

The bulls' argument is solid: ASICs are optimal for stable, predictable workloads, but GPUs are irreplaceable for research, prototyping and training new architectures. The bears' argument is also solid: if 75-80% of AI compute eventually is inference and ASICs capture 50-60% of that market, Nvidia's real TAM is a fraction of what the valuation implies.

The DeepSeek factor

January 27, 2025 was one of the most dramatic days in Nvidia's stock market history: the announcement of **DeepSeek-R1**, a Chinese model with capabilities similar to GPT-4 and a declared compute budget of $6M, wiped $589B off Nvidia's market capitalization in a single day, the largest one-day loss of value in stock market history.

The reality was more nuanced. DeepSeek used legitimate efficiency techniques, but the $6M figure excluded hardware and pre-training costs, and the model was developed on Nvidia GPUs. More importantly, **AI compute demand behaves like Jevons' paradox**: when efficiency improves and models become cheaper, the number of models and applications grows more than proportionally. The stock fully recovered in the following months, and Nvidia's results since then have confirmed that demand did not fall.

Export restrictions to China

The most concrete geopolitical risk for Nvidia is US export control policy toward China. In FY2022 China represented 26% of Nvidia's revenue; in FY2025 it was 13%, and part of that drop is explained by controls blocking the sale of the most advanced chips. In April 2025 the US restricted exports of the **H20**, the chip designed specifically for China, leading to a $4.5B charge for inventory and purchase commitments.

The risk has two readings. The pessimistic one: any further escalation can permanently eliminate a relevant part of sales and close a market of more than a billion people. The optimistic one: the demand China cannot meet with Nvidia is met with domestic alternatives like Huawei Ascend, which still lag in ecosystem and performance, preserving the Western advantage.

An important point is that Nvidia's FY2027 guidance assumed no Data Center Compute revenue from China. The company is modeling its growth without depending on the Chinese market for advanced accelerators, which reduces the risk that upcoming results depend on a regulatory reopening, but also confirms that a significant part of the Chinese TAM may remain out of reach for a long time.

SWOT analysis

STRENGTHS

- **An unmatched ecosystem moat:** CUDA, cuDNN, TensorRT, NeMo and millions of developers create switching barriers that take years to build and potentially decades to dismantle.
- **Extraordinary margins:** a 71-75% gross margin and an operating margin above 60% are software levels, reflecting genuine pricing power.
- **Leadership in the innovation cycle:** the Blackwell → Vera Rubin → Rubin Ultra roadmap sets an annual pace that keeps competitors permanently behind.
- **Growing vertical integration:** networking, software, systems, services and CPUs like Vera. Each additional layer makes substitution harder.
- **An AI factory platform, not just chips:** the more complex clusters become, the more value Nvidia captures outside the GPU.
- **Extraordinary FCF:** a business that generates cash faster than it can reinvest it productively.

WEAKNESSES

- **Concentration in one segment:** around 90% of revenue comes from Data Center and a handful of customers. A slowdown in hyperscaler AI CapEx hits results directly.
- **Total dependence on TSMC:** all of Nvidia's advanced manufacturing goes through TSMC in Taiwan.
- **Key person:** Jensen Huang has no identified successor.
- **Exposure to CapEx cycles:** hyperscalers can moderate spending if their own AI businesses do not generate the expected return.

OPPORTUNITIES

- **Agentic AI:** the shift from models that answer questions to agents that execute tasks multiplies inference demand.
- **Physical AI (robots, autonomous vehicles):** DRIVE Thor and Jetson give Nvidia a position in every vehicle and robot if these industries scale.
- **Sovereign AI:** governments around the world want their own "AI factories", built on Nvidia hardware.
- **Omniverse and digital twins:** real-time industrial simulation is an emerging market where Nvidia has a technological edge.

THREATS

- **Migration to ASICs in inference:** the risk is not that Nvidia loses training, but that a material part of stable, massive inference migrates to TPUs, Trainium, Maia, MTIA or other ASICs.
- **Escalating export restrictions:** each additional round of controls can cut off another segment of the Chinese market.
- **AMD with a mature ecosystem:** if ROCm reaches parity with CUDA, price becomes the deciding factor.
- **An "AI bubble":** if AI monetization disappoints and hyperscalers cut CapEx, the adjustment in Nvidia's revenue could be violent given the expectations built into the valuation.

Porter's Five Forces

1. RIVALRY AMONG EXISTING COMPETITORS: LOW-MEDIUM

In the short term, effective rivalry in frontier AI GPUs is low: AMD is the only competitor with a general-purpose product, and its share of AI data center is estimated at 5-8%. But the dynamic is changing: AMD is accelerating its roadmap, hyperscaler ASICs are real competition in inference, and startups like Cerebras and Groq compete in niches. Rivalry is low today but could be medium in 2027-2028.

2. THREAT OF NEW ENTRANTS: LOW

Barriers to entry are extraordinarily high. Designing a competitive chip requires engineering teams of hundreds of people with GPU architecture experience, manufacturing capacity at TSMC or Samsung, and above all a software ecosystem. The only viable entry path is a hyperscaler's own ASIC, which is not a "new entrant" but a customer becoming a partial competitor.

3. BARGAINING POWER OF SUPPLIERS: MEDIUM

Nvidia's most critical supplier is TSMC, the only manufacturer able to produce advanced-node chips at scale. But the relationship is mutually dependent: Nvidia is one of TSMC's largest customers. SK Hynix and Samsung, as HBM suppliers, also have significant bargaining power, since HBM shortages are one of the bottlenecks in GPU production.

4. BARGAINING POWER OF CUSTOMERS: MEDIUM, BUT STILL CONTAINED

Customers still have limited power because demand for Blackwell and Vera Rubin systems exceeds supply. But bargaining power is rising: Microsoft, Google, Amazon, Meta, OpenAI and Anthropic have enough scale to fund their own alternatives or deals with AMD and Broadcom. As long as Nvidia offers the lowest total cost per token, the lowest execution risk and the best software, it will keep its pricing power. The indicator to watch is gross margin.

5. THREAT OF SUBSTITUTES: MEDIUM-HIGH

This is the most relevant factor for the long-term analysis. ASICs are functional substitutes for GPUs in inference, with better energy efficiency and lower cost per inference. They do not replace GPUs in training and research, but most production AI compute may end up being inference. The key is to distinguish between total and partial substitution: Nvidia will probably remain dominant in training, research and flexible workloads; the risk is that production inference fragments toward more efficient proprietary solutions.

SECTION 06

The bulls vs. the bears

What do the bulls say?

The first argument is **execution**: revenue doubling on a $200 billion base, 75% margins and guidance that keeps rising quarter after quarter. Nobody in semiconductor history has grown like this.

The second is the **ecosystem**: CUDA, networking and complete systems make Nvidia the supplier of the whole “AI factory”, not of a chip. Each generation (Blackwell, Vera Rubin) lowers the cost per token and forces customers to keep buying.

The third is **future demand**: agents, reasoning, robotics and sovereign AI could multiply the need for compute for many years.

What do the bears say?

The first argument is **cyclicality**: revenue depends on the CapEx of a handful of customers. If that spending eases, the drop could be abrupt, as in previous semiconductor cycles.

The second is **ASICs**: TPUs, Trainium, Maia, MTIA and OpenAI’s chips target precisely inference, which could be most of compute in the long run.

The third is **margins**: current levels reflect scarcity. When supply catches up with demand and AMD matures, it is reasonable to expect them to normalize.

SECTION 07

Valuation

I built a ten-year **discounted cash flow (DCF)** model with two methods for the terminal value: **perpetuity growth** and an **exit multiple on EBITDA**.

Model inputs and assumptions

2. Operating drivers by scenario (explicit forecast)

Driver	Δ vs Base	FY2024A	FY2025A	FY2026A	3Y avg	FY2027E	FY2028E	FY2029E	FY2030E	FY2031E	FY2032E	FY2033E	FY2034E	FY2035E	FY2036E
Forecast year #						*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
Revenue growth ▸ selected scenario		125.9%	114.2%	65.5%	**100.0%**	**90.0%**	**68.0%**	**30.0%**	**12.0%**	**9.0%**	**7.0%**	**6.0%**	**6.0%**	**5.0%**	**4.5%**
Best case	4.0%					94.0%	72.0%	34.0%	16.0%	13.0%	11.0%	10.0%	10.0%	9.0%	8.5%
Base case						90.0%	68.0%	30.0%	12.0%	9.0%	7.0%	6.0%	6.0%	5.0%	4.5%
Worst case	(7.0%)					83.0%	61.0%	23.0%	5.0%	2.0%	–	(1.0%)	(1.0%)	(2.0%)	(2.5%)
Gross margin ▸ selected scenario		72.7%	75.0%	71.1%	**72.9%**	**74.5%**	**74.0%**	**73.0%**	**71.0%**	**69.0%**	**68.0%**	**67.0%**	**66.0%**	**65.0%**	**64.5%**
Best case	1.0%					75.5%	75.0%	74.0%	72.0%	70.0%	69.0%	68.0%	67.0%	66.0%	65.5%
Base case						74.5%	74.0%	73.0%	71.0%	69.0%	68.0%	67.0%	66.0%	65.0%	64.5%
Worst case	(3.0%)					71.5%	71.0%	70.0%	68.0%	66.0%	65.0%	64.0%	63.0%	62.0%	61.5%
EBIT margin ▸ selected scenario		54.1%	62.4%	60.4%	**59.0%**	**65.5%**	**66.0%**	**66.4%**	**62.0%**	**58.0%**	**56.0%**	**54.0%**	**52.0%**	**51.0%**	**50.0%**
Best case	3.0%					68.5%	69.0%	69.4%	65.0%	61.0%	59.0%	57.0%	55.0%	54.0%	53.0%
Base case						65.5%	66.0%	66.4%	62.0%	58.0%	56.0%	54.0%	52.0%	51.0%	50.0%
Worst case	(8.0%)					57.5%	58.0%	58.4%	54.0%	50.0%	48.0%	46.0%	44.0%	43.0%	42.0%
Capex % of revenue ▸ selected scenario		1.8%	2.5%	2.8%	**2.3%**	**2.3%**	**2.3%**	**2.4%**	**2.5%**	**2.5%**	**2.5%**	**2.5%**	**2.5%**	**2.5%**	**2.5%**
Best case	(0.5%)					2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%	2.0%
Base case						2.3%	2.3%	2.4%	2.5%	2.5%	2.5%	2.5%	2.5%	2.5%	2.5%
Worst case	0.5%					2.8%	2.8%	2.9%	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%

Source: author's analysis.

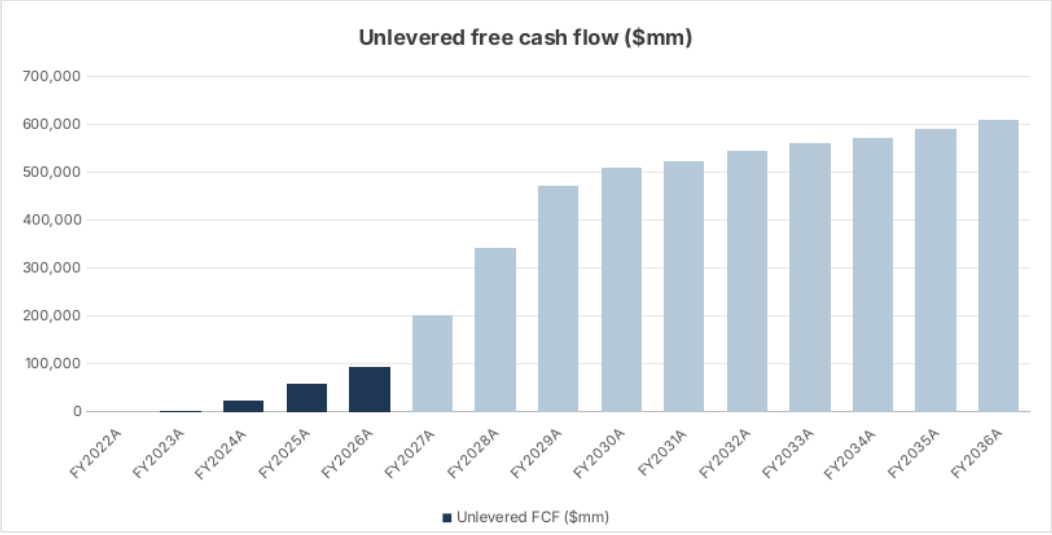
Revenue growth: in the Base Case sales grow 90% in FY2027 (around $410 billion, in line with the first two quarters and third-quarter guidance), 68% in FY2028 and 30% in FY2029. Growth then gradually slows to 4.5% in FY2036. Nvidia still has several years of very strong growth ahead, driven by Blackwell, Vera Rubin, networking and inference, but that rate inevitably compresses as the revenue base becomes gigantic.

Margins: one of the most important points of the model. I start from a 74.5% gross margin in FY2027 and gradually normalize it to 64.5%. EBIT margin starts at 65.5%, stays around 66% in FY2028 and FY2029 and then converges toward 50%. The normalization reflects a greater weight of complete systems, potential price pressure, hyperscalers' own ASICs and AMD as a second source. Even so, terminal margins remain extraordinary.

CapEx: between 2.3% and 2.5% of sales, consistent with the fabless model.

What does each scenario mean?

- **Base Case:** demand for AI compute keeps growing for several years and Nvidia defends most of the market, with margins normalizing gradually.
- **Best Case:** agents and reasoning multiply inference demand and Nvidia defends its margins better. It adds 4 points of growth and 3 points of EBIT margin.
- **Worst Case:** the AI investment cycle normalizes sooner, ASICs win a relevant share of inference and margins compress sharply. It subtracts 7 points of growth and 8 points of EBIT margin.



Source: author's analysis.

Discount rate (WACC)

5. Discount rate (WACC)		Value
Risk-free rate		4.9%
Equity beta (levered)	x	1.30
Equity risk premium		4.2%
Country risk premium		-
Cost of equity		10.4%
Pre-tax cost of debt		4.5%
Tax rate (debt tax shield)		13.3%
After-tax cost of debt		3.9%
Market value of equity	$ mm	4,551,633
Total debt	$ mm	11,412
Debt / (Debt + Equity)		0.3%
Equity / (Debt + Equity)		99.7%
WACC		10.3%

Source: author's analysis.

I use a 4.9% risk-free rate, a 4.2% equity risk premium and a beta of 1.3. The five-year historical beta is higher (2.3), but it reflects the volatility of a stock that multiplied several times; for a company with net cash and this cash generation, a lower value seems more reasonable to me. The result is a **cost of equity of 10.4%**. Since it has practically no debt, the **resulting WACC is 10.3%**.

For the terminal value I use a long-run WACC of 10%, slightly lower than today's, because ten years from now Nvidia should be a more mature company with less volatile results.

To go from Enterprise Value to Equity Value I added net cash of $73.4 billion. The diluted share count is 24,300 million.

Model results

PERPETUITY METHOD

With 2.5% perpetual growth and a 25% return on new invested capital, the model gives an Enterprise Value of $6.15 trillion and an Equity Value of $6.22 trillion, equivalent to **$255.95 per share**.

NVIDIA Corporation

NVDA · Discounted Cash Flow · Perpetuity Growth Method

Valuation date 6 Nov 2026 · USD mm · Scenario: Base Case

AZ STOCKS Acciones de la A a la Z

INTRINSIC VALUE PER SHARE

$255.95

Current share price	$187.31
Upside / (downside)	+36.6%

VALUATION SUMMARY

Enterprise value	6,146,111
Equity value	6,219,506
PV of forecast cash flows	3,023,635
PV of terminal value	3,122,476
Terminal value % of EV	50.8%
Implied exit EV / EBITDA	10.1x

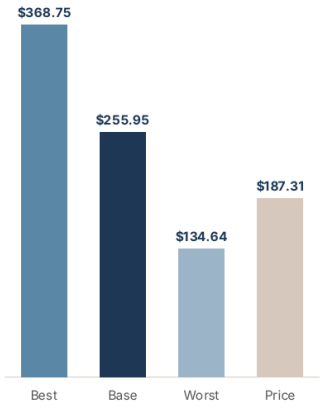
KEY ASSUMPTIONS

WACC (year 1)	10.3%
Terminal WACC	10.0%
Terminal growth (g)	2.5%
RONIC (terminal)	25.0%
Net debt & other claims	(73,395)
Diluted shares (mm)	24,300.0

SCENARIOS

Scenario	Terminal growth	Enterprise value	Equity value	Value per share	Upside
Best case	3.0%	8,887,242	8,960,637	**$368.75**	+96.9%
Base case	2.5%	6,146,111	6,219,506	**$255.95**	+36.6%
Worst case	2.0%	3,198,300	3,271,695	**$134.64**	(28.1%)
Current share price				*$187.31*	

VALUE PER SHARE VS. CURRENT PRICE



SENSITIVITY · VALUE PER SHARE

Terminal growth rate (g)

WACC / g	1.5%	2.0%	2.5%	3.0%	3.5%
9.3%	$278.53	$285.63	$293.82	$303.38	$314.67
9.8%	$261.17	$266.94	$273.53	$281.14	$290.02
10.3%	$245.87	$250.59	**$255.95**	$262.07	$269.13
10.8%	$232.28	$236.17	$240.55	$245.52	$251.19
11.3%	$220.14	$223.36	$226.97	$231.03	$235.62

Source: author's analysis.

In the Best Case the value is **$368.75** and in the Worst Case **$134.64**.

With a WACC half a point lower, the value rises to $274; with one half a point higher, it falls to $241.

EBITDA MULTIPLE METHOD

Applying a **12x** multiple to EBITDA in the last projected year of the Base Case (14x in the Best Case and 9x in the Worst Case), the model gives an Enterprise Value of $6.72 trillion and a value of **$279.69 per share**.

NVIDIA Corporation

NVDA · Discounted Cash Flow · Exit Multiple Method

Valuation date 6 Nov 2026 · USD mm · Scenario: Base Case

AZ STOCKS Acciones de la A a la Z

INTRINSIC VALUE PER SHARE

$279.69

Current share price	$187.31
Upside / (downside)	+49.3%

VALUATION SUMMARY

Enterprise value	6,722,962
Equity value	6,796,357
PV of forecast cash flows	3,023,635
PV of terminal value	3,699,327
Terminal value % of EV	55.0%
Implied perpetual growth	3.0%

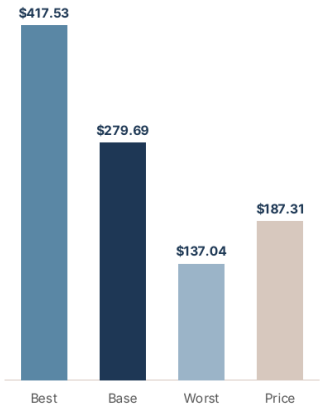
KEY ASSUMPTIONS

WACC (year 1)	10.3%
Terminal WACC	10.0%
Exit EV / EBITDA	12.0x
EBITDA, final forecast year	752,031
Net debt & other claims	(73,395)
Diluted shares (mm)	24,300.0

SCENARIOS

Scenario	Exit multiple	Enterprise value	Equity value	Value per share	Upside
Best case	14.0x	10,072,693	10,146,088	**$417.53**	+122.9%
Base case	12.0x	6,722,962	6,796,357	**$279.69**	+49.3%
Worst case	9.0x	3,256,593	3,329,988	**$137.04**	(26.8%)
Current share price				*$187.31*	

VALUE PER SHARE VS. CURRENT PRICE



SENSITIVITY · VALUE PER SHARE

Exit EV / EBITDA multiple

WACC / multiple	10.0x	11.0x	12.0x	13.0x	14.0x
9.3%	$270.54	$284.34	$298.14	$311.94	$325.74
9.8%	$262.26	$275.49	$288.72	$301.95	$315.18
10.3%	$254.31	$267.00	**$279.69**	$292.37	$305.06
10.8%	$246.69	$258.86	$271.03	$283.19	$295.36
11.3%	$239.38	$251.05	$262.72	$274.39	$286.06

Source: author's analysis.

I chose conservative multiples. They aim to value a more mature Nvidia at the end of the period: a very high-quality business, but one that should no longer be valued as a company in full hypergrowth. 12x implies 3.0% perpetual growth.

In the Best Case the value is **$417.53** and in the Worst Case **$137.04**.

Reading both methods together

Averaging both methods in the base case, I arrive at a **target price of $268 per share**, compared with a market price of $187.31. That is an upside of about 43%.

The two methods agree fairly well ($256 and $280), which makes the result more reliable. The range between scenarios, from about $135 to more than $400, shows that value depends mostly on how long Nvidia can sustain these margins.

SECTION 08

Conclusion

Nvidia is an extraordinary company. Not only because of its current numbers but because of the depth of the ecosystem it has built over more than two decades. CUDA, networking, NVLink, InfiniBand, Spectrum-X, complete systems like Blackwell and Vera Rubin, software libraries, developers, hyperscalers, AI labs and governments form a network that is extremely hard to replicate.

The company has achieved something very rare: turning a business originally associated with hardware into a platform with software margins. And in FY2027 it is proving it again: revenue doubling, 75% margins and guidance that keeps rising.

Even so, Nvidia being an exceptional company does not automatically make the stock a buy at any price. My model gives a **target price of $268 per share**, about 43% above the current price. That upside depends on the base case playing out, and it already assumes several more years of very strong growth.

The central point is that, to justify much higher prices, Nvidia has to execute very well for several more years: sustain Data Center growth, defend very high margins, get through the China restrictions, keep AMD from capturing too much share and, above all, show that its own customers' ASICs will not capture too large a part of inference.

That last point is, for me, the most important long-term risk. I do not think ASICs will "kill" Nvidia. But they can limit the size of the market the company can really capture in the mature stage of AI. Nvidia will probably remain dominant in training, research, new models and flexible workloads. The question is how much production inference will end up migrating to proprietary chips optimized for cost per token.

Even with those risks, Nvidia remains one of the highest-quality companies in the market: net cash, extraordinary margins, huge free cash flow generation and one of the deepest moats in the technology industry.

So my conclusion is not negative. Nvidia is still a company I would like to own over the long term. And at current prices my base case shows a reasonable margin of safety. For someone who wants to start a position, I would build it gradually, leaving room to add if the market offers a better entry point. For those who already own shares, I would not sell: I would hold the position and wait for better opportunities to add.

Alan



WEB
azstocks.net

SUBSTACK
@alaninvests

X
@alan_azstocks