

2026-03-22

## Executive Summary

NVIDIA is the dominant provider of AI compute infrastructure, benefiting from structurally increasing demand driven by hyperscaler investments and the growing importance of compute in the digital economy. Its full-stack offering, spanning GPUs, networking, systems, and software, creates high switching costs and reinforces its position as a key bottleneck in the AI value chain.

Despite strong financial performance, the share price has recently faced pressure due to investor uncertainty around the long-term economics of AI, the sustainability of current investment levels, and geopolitical risks. We view this as a short term dislocation rather than a deterioration in fundamentals, as AI infrastructure spending reflects a structural shift.

Based on our 5 year DCF analysis, we derive a target price of USD 244.64 per share, implying an upside of approximately 27% from the current share price. Given NVIDIA's dominant market position, strong pricing power, and durable long-term growth drivers, we initiate a BUY recommendation with a 12-month investment horizon.

## Key insights

**Accelerated Obsolescence and the 1- to 3-Year Refresh Cycle:** NVIDIA's aggressive shift to a one year product release cadence progressing rapidly from Hopper to Blackwell and the upcoming Vera Rubin architecture, which is driving a continuous cycle of accelerated hardware obsolescence. Because new generations like Blackwell offer up to 25 times better energy efficiency and staggering total cost of ownership improvements, older datacenter GPUs become economically non competitive for frontier model training within just 18 to 36 months. This dynamic forces major hyperscalers into an aggressive, recurring hardware replacement cycle, guaranteeing sustained long-term demand for NVIDIA's latest silicon rather than a one-time capacity buildout.

**Valuation Dislocation Driven by Short-Term Skepticism:** Despite NVIDIA consistently delivering exceptional revenue growth, record margins, and massive free cash flow, the share price has recently faced pressure due to short-term market skepticism regarding the long-term economics of AI and broader geopolitical risks. We view this disconnect as a sentiment-driven valuation dislocation rather than any deterioration in company fundamentals. Hyperscaler capital expenditures into AI infrastructure represent a structural, foundational investment phase that provides a tangible earnings anchor. As NVIDIA continues to execute and prove the sustainability of this demand, the gap between market perception and underlying fundamentals should close, driving a significant stock re-rating.

**Sovereign AI as a Structural Demand Layer:** Governments worldwide are recognizing AI infrastructure as a matter of national security. This has created a massive, entirely new customer base driving multi-year, politically motivated procurement that diversifies revenue away from corporate buyers.

## Analyst

<b>Alexander Gardos Ek</b>	Fund Analyst
<b>Gor Grigoryan</b>	Fund Analyst

## Market Data

Exchange	NASDAQ
Trading Currency	USD
Share Class	Common Stock
Shares OS	24,3 billion
Market Cap	4,4 TUSD
EV	4,3 TUSD

## Metrics & Drivers 2026 2027E 2028E

EV/EBITDA	32,5	14,49	14,99
P/E	36,8	21,5	16,1
EV/EBIT	61,7	60	60
EPS	4,9	6,4	11,0
CapEx/Revenue (%)	3	3	3


## Forecast, bUSD 2026 2027E 2028E

Revenue	215,94	367,09	477,22
- Growth Y-o-Y	65,47%	70,00%	30,00%
Gross Profit	153,46	275,32	357,92
- Gross Margin %	71,07%	75,00%	75,00%
EBIT	130,39	220,26	286,33
- EBIT Margin %	60,38%	60,00%	60,00%
Unlevered FCF	114,94	188,27	245,09

## Major Shareholders

The Vanguard Group, Inc.	9,33%
BlackRock, Inc.	7,98%
State Street Global Advisors, Inc.	4,08%
FMR LLC	3,65%
Jen-Hsun Huang	3,54%

## Stock Price Movement [Period] %



Beta (5Y)	2,38
52 Week Price Range	86,62 - 212,19
30D Avg Daily Volume	174 151 150
Free Float	96%



## Investment Thesis

### Accelerated Obsolescence and the 1- to 3-Year Refresh Cycle:

A critical, recurring growth driver for NVIDIA's Data Center segment is the rapid "accelerated obsolescence" of compute hardware, which is forcing hyperscalers into a continuous 1- to 3-year infrastructure refresh cycle. Historically, data center server platforms maintained a longer economic lifespan, but the unprecedented pace of AI scaling has drastically compressed this timeline.

To capitalize on this, NVIDIA has aggressively transitioned to a one-year product release cadence, advancing rapidly from Hopper (2022) to Blackwell (2024), and now to the upcoming Vera Rubin architecture (2026). Because each successive generation delivers staggering improvements in total cost of ownership (TCO) and energy efficiency such as Blackwell offering up to 25 times better efficiency than Hopper for specific workloads, older hardware becomes economically non competitive for frontier model training within just 18 to 36 months.

Industry insiders, including architects at major hyperscalers like Google, now project the effective service life of a datacenter GPU to be merely one to three years. Consequently, as the GPU with best performance naturally dominates the market, major customers are purchasing next generation systems not exclusively to expand their net compute capacity, but to aggressively upgrade and replace existing infrastructure, effectively guaranteeing a massive, recurring replacement cycle for NVIDIA's latest silicon.

### Valuation Dislocation Driven by Short-Term Skepticism:

Despite consistently delivering exceptional financial performance, characterized by rapid revenue growth, record-high gross margins, and sustained demand from hyperscalers, NVIDIA's share price has recently come under pressure. This reflects increasing investor uncertainty around the long-term economics of AI, including the value creation of AI applications, the sustainability of current investment levels, and how demand may evolve beyond the initial build-out phase. Geopolitical tensions and trade restrictions further add uncertainty around supply chains and end-market access.

We view this disconnect as a short-term sentiment-driven dislocation rather than a deterioration in fundamentals. Hyperscaler CAPEX, led by Microsoft, Amazon, and Alphabet, represents a structural reinvestment into AI infrastructure rather than a cyclical spike in demand. Unlike prior technology cycles, NVIDIA is already generating substantial and rapidly growing free cash flow, providing a tangible earnings anchor that supports valuation.

Over time, we expect this gap between market perception and underlying fundamentals to close. As NVIDIA continues to execute and demonstrate sustained earnings power, free cash flow generation and its dominant position within the AI ecosystem should drive a re-rating, with the share price converging toward its intrinsic value as short-term uncertainty subsides.

### Sovereign AI as a Structural Demand Layer:


Governments worldwide are increasingly treating AI infrastructure as a matter of national security and economic competitiveness. Nations across Europe, the Middle East, and Asia-Pacific are committing to building domestic AI capabilities, from data centers to foundation models, rather than relying solely on US hyperscalers. NVIDIA has emerged as the primary enabler of this trend, securing large-scale agreements with governments in Japan, France, India, and the UAE. This sovereign demand represents a structurally new and largely incremental customer base, independent of hyperscaler capex cycles. Unlike corporate customers, government procurement tends to be multi-year, politically driven, and less sensitive to short-term ROI concerns, providing NVIDIA with a more durable and diversified revenue base as the AI buildout enters its next phase.

### NVIDIA Product release cadence

Architecture	Launch Year	Approximate Cycle Gap
Ampere	2020	~30 Months
Hopper	2022	24 Months
Blackwell	2024	24 Months
Vera Rubin	2026	12 to 18 Months

Source: NVIDIA, Team analysis

### Free cash flow




Source: FiscalAI

### Forward P/E



Source: FiscalAI

### Earnings Per Share



Source: Yahoo Finance

## Company Overview

NVIDIA has evolved from a GPU-focused company into the leading provider of AI compute infrastructure, positioning itself at the center of the generative AI value chain. Leveraging its parallel processing architecture, NVIDIA delivers superior performance-per-watt for large-scale training and inference workloads, making its GPUs the preferred choice for hyperscalers and AI developers.

### The Economic Moat: Software and Networking

NVIDIA's competitive advantage, or economic moat, is not merely a result of superior hardware. A critical component is CUDA (Compute Unified Device Architecture), NVIDIA's proprietary software platform. By providing a robust ecosystem for AI developers, NVIDIA has created high customer switching costs. That is, once a developer's workflow is built on CUDA, moving to a competitor's hardware becomes a costly and complex endeavor. Furthermore, NVIDIA has expanded its reach into high-performance networking, allowing customers to cluster thousands of GPUs together which is a near-requirement for modern large-scale AI training.

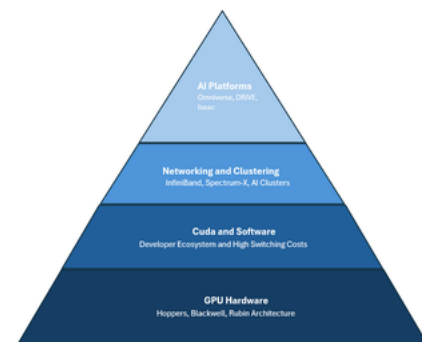
### Business Segments:

NVIDIA's operations are increasingly concentrated in its Data Center segment, which now accounts for approximately 90% of total revenue. This dominance is driven by the rapid adoption of the Blackwell and Rubin GPU architectures, complemented by high-performance networking solutions. While maintaining its global leadership in Gaming through the RTX 50-series, the company is strategically scaling its Automotive and Robotics divisions via the DRIVE and Isaac platforms to pioneer Physical AI. Additionally, NVIDIA continues to lead in industrial design and digital twin simulation through its Professional Visualization and Omniverse ecosystem.

### Business Model and Strategy:


NVIDIA operates a fabless model, outsourcing production to partners such as TSMC, enabling rapid scalability and structurally high gross margins (>70%). Despite rising competition from custom ASICs, NVIDIA's integrated hardware-software ecosystem and rapid innovation cycle continue to support its leadership in AI infrastructure.

### Economic Full Stack




Source: Team analysis

### Business Segments



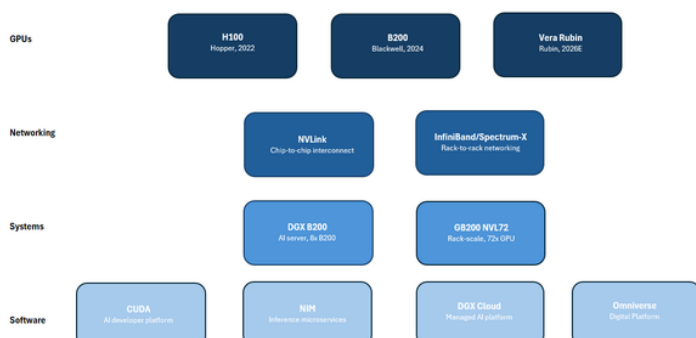
Source: NVIDIA FY2025 Annual Report

### Revenue and Gross Margins



Source: NVIDIA FY2025 Annual Report

### Data Center Products



Source: NVIDIA, Team analysis

## Market Overview

The global data center market is undergoing a fundamental structural transformation, shifting from general-purpose CPU-based architectures toward accelerated computing, primarily driven by GPUs. This transition is fueled by the exponential compute requirements of Large Language Models (LLMs) and generative AI applications, which demand significantly higher performance and scalability. Industry estimates suggest that the total server market could exceed \$1 trillion over the next decade, with growth largely driven by AI-related infrastructure investments. In the near term, demand is heavily concentrated among hyperscalers such as Microsoft, Alphabet, Amazon, and Meta, which are engaged in a strategic race to build out AI capabilities. Combined capital expenditures from these players surpassed \$380 billion in 2025 and are projected to approach \$600 billion in 2026 – a near-doubling in just two years, with a majority allocated to AI compute and networking. Given the competitive importance of AI leadership, underinvestment is widely perceived as a greater risk than overinvestment, supporting sustained demand for high-performance infrastructure.


### Industry Structure and NVIDIA's Strategic Position.

Within this rapidly expanding market, NVIDIA has established itself as the dominant supplier of high-end AI compute, holding an estimated AI accelerated GPU market share exceeding 80%. This positions the company as a critical bottleneck in the AI value chain, as hyperscalers rely heavily on NVIDIA's hardware, including its Hopper and Blackwell architectures, to scale their data center capabilities.

### Competitive Landscape:


Despite its dominant position, NVIDIA faces emerging competition from both traditional semiconductor peers and hyperscalers developing in-house solutions. AMD has positioned itself as the most credible merchant silicon competitor, with its Instinct series offering competitive performance in certain workloads. At the same time, hyperscalers are investing heavily in proprietary ASICs, such as Google's TPU and AWS's Trainium, aiming to reduce dependence on third-party suppliers and lower total cost of ownership, particularly for inference. However, these competitive threats are mitigated by NVIDIA's substantial competitive advantages. The company's CUDA software ecosystem has created significant switching costs, with millions of developers building AI workflows tightly integrated with NVIDIA hardware. In addition, NVIDIA benefits from a steep experience curve and a rapid innovation cadence, consistently launching new architectures ahead of competitors.

### AI accelerator Market Forecast (Billions)




Source: Northwise article

### Hyperscalers Capex (millions)



Source: FiscalAI

### AI Accelerated GPU Market Share 2025



Source: Northwise article

### Total cost of ownership and Cost per TFLOP comparison

Total cost of ownership Nvidia vs AMD		
Cost	Nvidia B200	AMD MI355X
Sell price	40000	25000
TDP (kW)	1	1,4
Electricity cost per 3 years	2102,4	2943,36
Software & engineering costs	Low	High
<b>TCO</b>	<b>42102,4</b>	<b>27943,36</b>
TFLOPS	9000	10000
<b>Cost per TFLOP</b>	<b>4,68</b>	<b>2,79</b>

Source: Silicon Analysts, Team Analysis

At first glance, AMD's MI355X appears cheaper with a lower acquisition price and higher raw TFLOP count. However, the comparison understates NVIDIA's true advantage. CUDA's native integration across all major AI frameworks eliminates the significant engineering costs required to optimize workloads on AMD's ROCm platform. When software utilization rates are factored in, NVIDIA's fully-loaded cost per effective TFLOP is competitive and in many enterprise deployments superior.

## Valuation & Financial Analysis

The base case models a massive 70% revenue spike in 2027E, which then stabilizes into 30%-40% annual growth through 2031E. It assumes NVIDIA maintains extreme pricing power, projecting flat 75% gross margins and 60% operating (EBIT) margins throughout the entire period. To sustain this growth, the models steady capital expenditures (CapEx) at 3.00% of revenue and Depreciation & Amortization at 1.50% Future cash flows are discounted at a rate of 10.45%. Because NVIDIA is virtually debt-free in this model (debt is only 0.19% of the capital structure), this rate is driven almost entirely by the 10.46% Cost of Equity (calculated using a 4.15% risk-free rate and a 1.44 Beta). The model applies a 2.00% long-term terminal growth rate for the period beyond 2031, which accounts for the vast majority of the company's total \$5.94 trillion Enterprise Value. Dividing the final equity value by 24.3 billion shares outstanding yields a base-case intrinsic value of \$244.64 per share. A sensitivity analysis stress-tests this \$244.64 price target against changes in the economy. It shows the implied share price could swing anywhere from \$172.20 (high WACC/low growth bear case) to \$419.00 (low WACC/high growth bull case).

## DCF & Sensitivity Analysis

DCF USD bn	2026	2027E	2028E	2029E	2030E	2031E
Revenue growth - %	-	70,00%	30,00%	35,00%	40,00%	35,00%
EBIT - %	60,38%	60,00%	60,00%	60,00%	60,00%	60,00%
<b>EBIT</b>	<b>130 387,00</b>	<b>220 256,76</b>	<b>286 333,79</b>	<b>386 550,61</b>	<b>541 170,86</b>	<b>730 580,66</b>
D&A	2 843,00	5 506,42	7 158,34	9 663,77	13 529,27	18 264,52
<b>EBITDA</b>	<b>133 230,00</b>	<b>225 763,18</b>	<b>293 492,13</b>	<b>396 214,38</b>	<b>554 700,13</b>	<b>748 845,18</b>
Tax	109525,08	185015,68	240520,38	324702,52	454583,52	613687,75
Tax rate	16%	16%	16%	16%	16%	16%
CapEx	6 042,00	6 223,26	6 409,96	6 602,26	6 800,32	7 004,33
Change in OWC	--	(8 618,00)	(3 967,40)	(3 821,36)	(4 164,22)	(4 551,55)
<b>Unlevered FCF</b>	<b>--</b>	<b>188266,24</b>	<b>245090,13</b>	<b>331928,25</b>	<b>465864,02</b>	<b>629938,65</b>
<b>Present value of FCF</b>	<b>--</b>	<b>170458,57</b>	<b>200917,89</b>	<b>246367,49</b>	<b>313072,38</b>	<b>383292,41</b>

WACC	
Cost of Debt	3,06%
Cost av Equity	10,46%
Tax Rate	16,00%
Levered Beta	1,44
Market Risk Premium	4,38%
Risk Free Rate	4,15%
Long-Term Growth Rate	2,00%
<b>WACC</b>	<b>10,45%</b>

DCF Valuation	
Sum of FCF	1314108,74
TV	4628412,58
<b>EV</b>	<b>5942521,32</b>
(-) Debt	8468,00
(+) Cash	10605,00
<b>Equity Value</b>	<b>5944658,32</b>
Shares outstanding	24300,00
<b>Equity Value per share</b>	<b>244,64</b>

## Sensitivity table

WACC	Terminal value						
	0,50%	1,00%	1,50%	2,00%	2,50%	3,00%	3,50%
8,45%	276,1	291,9	310,0	330,9	355,3	384,2	419,0
9,45%	241,2	253,1	266,6	281,8	299,3	319,4	343,0
10,45%	213,5	222,7	233,0	<b>244,5</b>	257,5	272,2	289,0
11,45%	190,9	198,2	206,3	215,3	225,2	236,3	248,8
12,45%	172,2	178,1	184,6	191,7	199,5	208,1	217,7

### Peer valuation and Financial Metrics

A hybrid peer group was selected to capture NVIDIA's dual market position. Traditional semiconductor competitors like AMD, Intel, and Broadcom establish a baseline for pure play silicon design valuations. Conversely, major hyperscaler customers like Alphabet, Amazon, and Meta represent the primary drivers of the AI infrastructure buildout. Blending these two groups ensures the valuation reflects both standard hardware dynamics and the massive capital expenditure premium of the AI ecosystem. The data clearly justifies NVIDIA's premium over this hybrid group, highlighting its exceptional 65.47 percent revenue growth and 113.80 percent Return on Invested Capital.


The valuation utilizes three Last Twelve Months trading multiples to provide a comprehensive view of the company's financial health. Enterprise Value to EBITDA serves as the professional standard, measuring pure operational cash flow by stripping away differences in capital structure, taxes, and depreciation. The Price-to-Earnings ratio is included to gauge broader retail market sentiment regarding bottom-line net income, despite potential distortions from debt or tax environments. Finally, the Enterprise Value to Sales multiple measures the company's value relative to its raw top-line revenue. This is particularly valuable for evaluating high-growth tech peers, as it normalizes comparisons across firms with varying profit margins or heavy research and development spending.

To calculate the implied intrinsic value, the average and median multiples across the six peers are synthesized and applied directly to NVIDIA's financial data. For example, a median peer EV/EBITDA of 21.90 or an average P/E of 44.40 is multiplied by NVIDIA's standalone results of nearly 216 billion dollars in revenue and over 133 billion dollars in EBITDA. After establishing the Enterprise Value, the model adjusts for NVIDIA's specific net debt and divides the total by its 24.3 billion outstanding shares. This yields a spectrum of implied share prices, such as 147 dollars under the average EV/EBITDA scenario or 142 dollars under the median Price-to-Earnings scenario, demonstrating what NVIDIA would be worth if valued exactly as an average peer within this elite cohort.

### Return on invested capital

	ROIC
<b>NVIDIA</b>	113,80%
<b>AMD</b>	6,60%
<b>Intel</b>	0,00%
<b>Broadcom</b>	17,90%
<b>Alphabet</b>	27,60%
<b>Amazon</b>	11,60%
<b>Meta</b>	24,50%

### Football Field Graph



### Different implied share prices

Current share price	178,56
Implied Share price (Average P/E)	218,892
Implied Share price (Average EV/EBITDA)	147,30
Implied Share price (Median P/E)	142,477
Implied Share price (Median EV/EBITDA)	72,98

Company name	Market cap (USD bn)	Enterprise value (USD bn)	EBITDA margin	ROA	Revenue Growth	P/E	EV/EBITDA
<b>NVIDIA</b>	4383	4500	62%	85%	66%	36,80	32,50
<b>AMD</b>	325	300	20%	11%	34%	76,40	47,20
<b>Intel</b>	225	150	24%	4%	-1%	-750,50	19,50
<b>Broadcom</b>	1496	1700	55%	12%	25%	61,60	41,40
<b>Alphabet</b>	3715	4000	37%	17%	15%	28,50	24,30
<b>Amazon</b>	2253	2600	20%	7%	12%	29,30	15,80
<b>Meta</b>	1557	1800	51%	25%	22%	26,20	15,30



## Valuation & Financial Analysis

### **DISCLAIMER:**

KTH Finance Society ("KTHFS") provides the following analyses, documents, and information solely for informational and educational purposes and general distribution. It is not intended to be treated as financial advice in any manner. The content presented in these analyses is based on sources, data, and individuals that KTHFS deems reliable. However, KTHFS cannot assure the accuracy or completeness of such information. Predictions and forward-looking statements included in this analysis are speculative and inherently uncertain. Users should exercise caution when basing decisions on this information. Any investment decisions influenced by the information provided by KTHFS, its employees, or affiliates should be considered the sole responsibility of the investor. The materials offered by KTHFS are meant to complement, not replace, other resources in the decision-making process for various investment types. Investors are encouraged to seek additional data and consult with a professional financial advisor before making any investment decisions. KTHFS expressly disclaims any liability for losses or damages of any kind arising from the use of its analyses, documents, and any other provided information. This research paper may also contain inaccuracies due to incorrect data sources or human error.

### **CONFLICTS OF INTEREST AND IMPARTIALITY:**

To maintain its independence and integrity, KTHFS has implemented strict compliance guidelines for its analysts. Furthermore, all analysts are bound by an agreement that mandates the disclosure of any potential conflicts of interest. These measures are in place to align with COMMISSION DELEGATED REGULATION (EU) 2016/958 of 9 March 2016, which supplements Regulation (EU) No 596/2014 of the European Parliament and of the Council. This regulation pertains to the regulatory technical standards for the objective presentation of investment recommendations or other information suggesting an investment strategy, as well as for the disclosure of specific interests or potential conflicts of interest. The author of this paper may hold an open position in the underlying instrument, but have no intentions to gain leverage of this papers potential perception by the markets by trading the instruments in a way thats aligned with any potential profitable outcomes.