

NVIDIA Corp. - (NVDA) | Investment Memo

December 2025 | Base Case PT: 238.53 (31% upside)

Summary

NVIDIA is the leading provider of accelerated computing hardware and software for AI and high-performance workloads. The stock has recently been re-rated for AI on the first phase of the larger AI infrastructure boom, but I believe the market still underestimates 1) how durable AI infrastructure demand will be, 2) how much pricing power NVIDIA's full-stack platform provides relative to competitors, 3) the extent to which supply, power and integration constrain the near term, and 4) the incremental demand from enterprises and verticals beyond hyperscalers.

My model values NVIDIA over a 12–24-month horizon anchored on forward P/E and EV/EBITDA multiples, which is cross-checked against a short-dated DCF. The weighted target price is \$238.53 vs a current price of \$182.31, which implies a 31% upside over the next 2 years.

- 1) Bull: \$246.54 (90% weight) – continued AI capex “arms race”, strong data-center growth, premium multiples
 - 2) Base: \$189.93 (5% weight) – strong but maturing AI cycle, modest multiple compression
 - 3) Bear: \$142.91 (5% weight) – AI digestion period, more aggressive pricing pressure
-

Investment Thesis

After a high growth past few years, many investors worry they're too late to NVIDIA and frame it as an “AI bubble” that will fade once hyperscaler spend normalizes and custom chips ramp. With long-dated DCF's extremely sensitive, the stock seeming trades on sentiment and headline multiples instead of detailed driver work. My view, informed from my field research conversations and the model is as follows:

- 1) **AI infrastructure demand is a multi-year program, not a one-off spike:** My field research (consultants, cloud architects, operators) consistently pointed to high conversion from pilots to production and quick follow-on orders once use cases show ROI, with clusters treated as long-lived infrastructure. Reflecting this, my model assumes above-consensus Data Center growth, taking revenue from roughly \$130B to ~\$400B over five years with a gradual slowdown, not a sharp drop.
 - 2) **NVIDIA's full-stack platform sustains premium margins and slows competitive share loss:** Contacts emphasized that CUDA, NVIDIA's software stack, and integrated systems/networking make it the default and sticky choice, with meaningful friction and risk in moving to alternatives. I translate this into structurally high gross margins and EBIT margins in the high-50s to low-60s and ~50% FCF margins in my base and bull cases, which supports using premium forward P/E and EV/EBITDA multiples versus broader semis and large-cap tech.
 - 3) **Physical bottlenecks, not weak demand, are the main near-term constraint—stretching the cycle:** Practitioners highlighted HBM/CoWoS capacity, data-center power/cooling, and integration timelines as the key brakes, saying projects are delayed in the queue rather than cancelled. In my model, this shows up as timing shifts in revenue and FCF (some years lighter/heavier) but a still-large five-year cash-flow opportunity, and it supports using a somewhat higher terminal growth rate than a typical mature semi in the DCF cross-check.
 - 4) **A second wave of enterprise and vertical demand is building behind hyperscaler:** Conversations in healthcare, financials, and industrials pointed to growing pipelines of “Phase 2/3” AI projects and internal “copycat” demand once one business unit proves value, often requiring hybrid or on-prem deployments where NVIDIA captures most of the stack. My model reflects this by assuming an increasing contribution from enterprise/vertical customers, helping sustain double-digit revenue growth and underpinning the premium—but still grounded—multiples that drive my ~\$238 scenario-weighted target price.
-

Business Description

NVIDIA generates revenue across four primary end markets (see revenue, geography, and customer tables). In FY2025, total revenue was \$130.5B, with data center making up 88% of the mix, gaming 9%, and all other segments a low-single-digit contribution.

- 1) **Data Center (AI systems + networking + software):** NVIDIA's core business is its data-center segment, which bundles AI systems (compute), networking hardware, and software/AI services into a full-stack infrastructure offering. In FY2025, data center generated \$115.2B, with roughly 89% of that coming from compute (GPUs/CPUs), underscoring that high-end AI accelerators are the economic engine of the company.
- 2) **Gaming & Professional Visualization:** This segment includes GeForce GPUs, GeForce NOW, RTX workstations, and the Omniverse platform sold to gamers, creators, and enterprises that need high-end graphics and visualization. Together, Gaming and ProViz contributed about 10% of FY2025 revenue, and remain important but more cyclical than data center.
- 3) **Automotive:** NVIDIA's automotive business centers on the DRIVE platform for ADAS/AV and Jetson modules for robotics/edge AI, sold to automakers, Tier-1s, and robotics developers. It accounts for ~1% of revenue today but represents a long-duration option on autonomous driving and robotics.

- 4) **OEM & Other:** Legacy and miscellaneous OEM products now represent <0.5% of revenue and are non-core to the thesis.

In sum, NVIDIA is no longer a generic GPU vendor; it is the AI infrastructure choke point for a small set of hyperscaler “AI factories.” Data Center is the economic engine, bundling AI systems, networking, and CUDA software into a full-stack platform, while Gaming/ProViz remains important but more cyclical and Auto/OEM are long-dated options rather than core to the thesis. Revenue is anchored in U.S. and Asia cloud hubs, with the U.S. and key Asian data-center regions together representing the vast majority of FY2025 sales. Exhibits 1 and 2 reinforce this framing by showing NVIDIA’s dominant share across the AI stack and its concentrated dependence on a handful of hyperscaler customers whose AI capex decisions ultimately drive the story.

Exhibit 1

Sub-Segment	NVIDIA Approx. Share	Key Competitors (Approx. Share)
High-end data-center GPUs (AI accelerators for training + inference)	~90% (range 85–95%) of data-center GPU units/revenue	AMD + others ~5–15% combined (AMD Instinct line plus Google TPU, AWS Trainium/Inferentia, etc.)
AI data-center chips overall (all AI chips, not just GPUs)	~65% of AI data-center chip revenue	Intel ~22%, AMD ~11%, others ~2%
InfiniBand networking hardware (HPC / AI cluster ports)	~80–85% of InfiniBand ports shipped by NVIDIA Networking (ex-Mellanox)	Others ~15–20% (Intel, Oracle, Cisco, IBM, etc.)
AI back-end network protocols (inside AI clusters – InfiniBand vs Ethernet)	InfiniBand 80% of AI back-end networks	Ethernet ~20% (Broadcom, Cisco, Arista...)
GPU compute software ecosystems (developers / workloads)	CUDA = clear majority of GPU-compute workloads	AMD ROCm, Intel oneAPI/SYCL and others = minority share

Exhibit 2

Customer	% of FY2025 Revenue	% of Q2 FY2026 Revenue
Microsoft (Azure / OpenAI)	~12% FY2025	~23% Q2 FY2026
Amazon (AWS)	~11% FY2025	~16% Q2 FY2026
Google (GCP)	~11% FY2025	est. 10–12% Q2 FY2026
Meta (Facebook)	≥10% indirectly	part of top 3 DC customers
Other (Tencent, Baidu, Oracle Cloud, etc.)	<10% individually	Smaller shares

Moat Analysis

- 1) **CUDA Ecosystem:** NVIDIA’s CUDA platform houses millions of developers and over 500 specialized libraries built around it. This ecosystem makes it difficult for customers to switch because of massive costs. Developers and enterprises that have invested in CUDA tools and workflows are reluctant to move into alternatives, even if competing hardware is cheaper.
- 2) **Full-stack Integration:** NVIDIA controls the entire AI infrastructure stack—chips, interconnects, systems and software. The integration ensures that deployment is faster for companies and optimizes performance relative to piecemeal solutions. This provides customers with a time-to-value that competitors struggle to match.
- 3) **Networking Lock-in:** With InfiniBand and spectrum-x ethernet, NVIDIA also controls how GPUs communicate at scale. Once a customer has built GPU clusters around NVIDIA’s networking technology, it is very difficult to switch. Reinforcing lock-in at the system level.
- 4) **Pressure Points:** Rising competition from AMD’s MI300 GPUs, hyperscaler in house chips, and potential regulatory restrictions challenge NVIDIA’s moat. While currently none of them can match the breath of NVIDIA’s ecosystem, these threats can continue to erode pricing power and customer dependency.

Key Tension Points

- 1) **Supply Chain: HBM & Advanced Packaging (CoWoS):** At the supply-chain level, the key question centers on whether HBM and CoWoS stay a hard cap on NVIDIA’s data-center growth or ease enough to unlock another leg of upside. HBM capacity and CoWoS throughput control how many high-end, high-HBM GPUs NVIDIA can ship—so faster-than-expected ramps at SK hynix, Samsung, Micron, and TSMC would directly translate into more Blackwell-class systems and a richer margin mix than the market models today.

- 2) **Hyperscaler Capex Cadence & Buyer Power:** For hyperscalers, the issue concerns whether AI capex remains structurally high or slips into a sharper “digestion” as buyer power and custom silicon grow. Because a handful of “AI factory” customers (Microsoft, Amazon, Google, Meta) drive most data-center revenue, their AI budgets and discount expectations map almost one-for-one into NVIDIA’s units and pricing; my field work supports an elevated but lumpy, stair-step pattern rather than a near-term cliff.
- 3) **Power Availability & Cooling:** On the infrastructure side, the open question concerns whether power and cooling become a structural cap on AI build-outs or mostly a timing issue. AI factories are extremely power-dense, and grid, substation, permitting, and liquid-cooling bottlenecks can push projects out by quarters, so I treat power/cooling as phasing risk that adds volatility to quarterly growth instead of permanently destroying demand over my 18–24 month horizon.
- 4) **Competition: AMD & In-House Chips:** Competitively, the focus shifts to how fast AMD and custom cloud silicon can erode NVIDIA’s share and pricing before this AI cycle matures. Alternative accelerators will take pocket share in specific workloads, but based on ecosystem checks I still underwrite NVIDIA keeping a premium footprint at the high end because CUDA lock-in, time-to-train, and full-stack integration matter more than pure \$/FLOP for mission-critical models.
- 5) **Regulatory & Geopolitical Constraints (Export Controls, AI Policy)** On regulation, the question turns to whether export controls and AI policy merely reshape NVIDIA’s mix or materially cap TAM and margins. U.S. rules already limit high-end GPU sales into China and could tighten further, so I treat regulation as an asymmetric macro risk—a manageable headwind in my base case, but one of the clearest downside catalysts if a major new policy shock removes more high-margin demand than other regions can backfill.

Financial Analysis & Valuation

- 1) **Operating Model:** My operating model starts from NVIDIA’s segment income statement and builds a five-year view that treats data center as the economic engine of the company. In the active case, total revenue grows from roughly \$130B in FY2025 to ~\$400B by FY2030, with data center compounding faster and remaining the dominant share of the mix. Gross margin holds in the low-70s and EBIT margin in the low-60s as Blackwell ramps, networking attach increases, and opex grows slower than revenue. Free cash flow margins stay very high as capex and working capital remain modest. These assumptions are directly tied to my field work: I assume step-wise but elevated AI capex, treat HBM/CoWoS and power as phasing risk rather than permanent caps, and underwrite continued full-stack pricing power.
- 2) **DCF and Sensitivity:** Using this model, I run a FCFF DCF with a 9% WACC and 4% terminal growth and flex my key drivers (growth, margins, capex, NWC) to define bear, base, and bull cases. Based on these assumptions the DCF supports equity values of roughly \$143 (bear), \$190 (base), and \$247 (bull) per share. I then test how sensitive this is to macro and policy by running a two-way grid on WACC (7–11%) and terminal growth (2–6%) in Exhibit 6. The grid produces a wide range of outcomes (~\$140 to >\$1,000 per share), but the middle of the table (around 9% / 4%) lines up close to my bull-case value, which I use as an upper bound rather than a point estimate. The takeaway is that NVIDIA’s valuation is sensitive to discount-rate and long-run growth assumptions, but under a broad set of reasonable inputs the cash-flow profile can still support values well above today’s price.
- 3) **Peer & Implied Multiples (Exhibits 4-5):** To ground the DCF, I compare NVIDIA to three peer baskets in Exhibit 4: AI-focused semis (AVGO, AMD, MRVL), broader digital semis, and mega-cap platforms (MSFT, GOOGL, AMZN, META). Those groups generally trade around mid-single- to low-double-digit forward P/E and mid-teens EV/EBITDA on 2026–27 numbers; NVIDIA trades at a premium, which I think is justified by its faster growth, higher margins, and position as an AI infrastructure choke point. From the DCF I then back out justified multiples in each scenario (Exhibit 5). In the base case, NVIDIA settles at roughly low-20s P/E and high teens EV/EBITDA on Year-3 earnings, with the bear and bull cases bracketing that band. That keeps NVIDIA at a premium to traditional semis but within a reasonable range relative to other category-defining infrastructure and cloud platforms—consistent with my thesis that it should trade more like AI infra than a commodity chip name.
- 4) **Scenario-Weighted Target:** Finally, I convert the three scenarios into a 12–24 month price target. The bear, base, and bull outputs are: Bear: ~\$143 per share, Base: ~\$190 per share, and Bull: ~\$247 per share. Based on my field research (pilot conversion, follow-on order cadence, supply easing, and competitive dynamics), I view the bull path as most likely and assign probabilities of 5% / 5% / 90% to bear, base, and bull. This yields a scenario-weighted target price of ~\$238–\$239 per share, or a little over 30% upside from the current stock price. In other words, if NVIDIA can execute on the multi-year AI build-out I’ve modeled and avoid a major policy or execution shock, the stock still offers attractive, positively skewed risk/reward.

Exhibit 3 – (Bull Case Shown)

Nvidia Operating Model								
Income Statement		Fiscal Year Ending January						
\$MM, unless otherwise noted	2023A	2024A	2025A	2026E	2027E	2028E	2029E	2030E
Data Center Revenue	15,005	47,525	115,186	195,816	283,933	411,704	514,629	607,263
% growth		216.7%	142.4%	70.0%	45.0%	45.0%	25.0%	18.0%
Gaming Revenue	9,067	10,447	11,350	17,025	19,579	22,320	24,998	27,498
% growth		15.2%	8.6%	50.0%	15.0%	14.0%	12.0%	10.0%
Professional Visualization Revenue	1,544	1,553	1,878	2,348	2,770	3,102	3,351	3,585
% growth		0.6%	20.9%	25.0%	18.0%	12.0%	8.0%	7.0%
Automotive Revenue	903	1,091	1,694	2,626	3,939	5,908	8,566	11,565
% growth		20.8%	55.3%	55.0%	50.0%	50.0%	45.0%	35.0%
OEM and Other Revenue	455	306	389	498	617	741	859	954
% growth		(32.7%)	27.1%	28.0%	24.0%	20.0%	16.0%	11.0%
Total Revenue	26,974	60,922	130,497	218,312	310,838	443,775	552,404	650,864
% growth		125.9%	114.2%	67.3%	42.4%	42.8%	24.5%	17.8%
Cost of Revenue	(11,618)	(16,621)	(32,639)	(61,127)	(77,710)	(106,506)	(129,815)	(149,699)
% margin	(43.1%)	(27.3%)	(25.0%)	(28.0%)	(25.0%)	(24.0%)	(23.5%)	(23.0%)
Gross Profit	15,356	44,301	97,858	157,185	233,129	337,269	422,589	501,166
% margin	56.9%	72.7%	75.0%	72.0%	75.0%	76.0%	76.5%	77.0%
Research and Development Costs	(7,339)	(8,675)	(12,914)	(17,902)	(21,759)	(30,177)	(37,563)	(44,259)
% margin	(27.2%)	(14.2%)	(9.9%)	(8.2%)	(7.0%)	(6.8%)	(6.8%)	(6.8%)
Sales, General and Administrative Costs	(2,440)	(2,654)	(3,491)	(4,366)	(4,973)	(6,213)	(7,734)	(8,461)
% margin	(9.0%)	(4.4%)	(2.7%)	(2.0%)	(1.6%)	(1.4%)	(1.4%)	(1.3%)
Acquisition Termination Costs	(1,353)	0	0	0	0	0	0	0
% margin	(5.0%)	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Operating Expenses	(11,132)	(11,329)	(16,405)	(22,268)	(26,732)	(36,390)	(45,297)	(52,720)
% margin	(41.3%)	(18.6%)	(12.6%)	(10.2%)	(8.6%)	(8.2%)	(8.2%)	(8.1%)
Operating Income (EBIT)	4,224	32,972	81,453	134,917	206,397	300,879	377,292	448,446
% margin	15.7%	54.1%	62.4%	61.8%	66.4%	67.8%	68.3%	68.9%
EBITDA	5,768	34,480	83,317	137,261	209,194	304,873	381,711	453,653
% margin	21.4%	56.6%	63.8%	62.9%	67.3%	68.7%	69.1%	69.7%

Exhibit 4

Company	Headquarters	Current Trading Value			Price / Earnings			Enterprise Value / EBITDA		
		Share Price	Market Cap	Enterprise Value	NTM	CY 2026E	CY 2027E	NTM	CY 2026E	CY 2027E
AI Focused Semiconductor Peers										
Broadcom Inc.	United States	309.09	1,453,796	1,511,606	20.0x	18.8x	16.2x	30.3x	28.3x	24.7x
Advanced Micro Devices, Inc.	United States	184.42	299,285	297,304	7.9x	7.5x	6.4x	32.3x	28.6x	22.2x
Marvell Technology, Inc.	United States	79.32	68,390	72,016	7.8x	7.3x	6.4x	19.5x	18.4x	16.2x
Mean			607,157	626,975	11.9x	11.2x	9.7x	27.4x	25.1x	21.0x
Median			299,285	297,304	7.9x	7.5x	6.4x	30.3x	28.3x	22.2x
Digital Semiconductor Peers										
TSMC	Taiwan	40.09087	1,039,661	984,332	7.3x	7.0x	5.9x	10.8x	10.3x	8.9x
ASML Holding N.V.	Netherlands	750.4686	290,663	286,504	7.4x	7.3x	6.4x	20.1x	19.9x	16.7x
QUALCOMM Incorporated	United States	156.59	168,738	173,515	3.9x	3.9x	3.8x	10.8x	11.0x	10.1x
Texas Instruments Incorporated	United States	193.29	175,727	184,411	9.8x	9.5x	8.7x	20.3x	19.3x	16.9x
Arm Holdings plc	United Kingdom	141.6	149,954	147,444	28.3x	26.3x	23.1x	54.4x	48.2x	38.6x
Intel Corporation	United States	22.22	97,257	134,209	2.5x	2.5x	2.4x	8.4x	7.6x	5.8x
Analog Devices, Inc.	United States	237.63	117,923	122,766	10.5x	10.2x	9.2x	23.4x	22.4x	18.7x
Mean			291,418	290,455	10.0x	9.5x	8.5x	21.2x	19.8x	16.5x
Median			168,738	173,515	7.4x	7.3x	6.4x	20.1x	19.3x	16.7x
Leading Tech Companies										
Microsoft Corporation	United States	520.58	3,869,558	3,887,177	11.8x	11.2x	9.8x	20.2x	18.9x	16.4x
Alphabet Inc.	United States	203.03	2,448,314	2,394,834	5.6x	5.5x	4.9x	12.2x	11.7x	10.2x
Amazon.com, Inc.	United States	224.56	2,394,913	2,461,303	3.3x	3.2x	2.9x	13.1x	12.4x	10.6x
Meta Platforms, Inc.	United States	780.08	1,959,671	1,962,160	9.0x	8.6x	7.6x	14.8x	14.1x	12.2x
Mean			2,668,114	2,676,369	7.4x	7.1x	6.3x	15.1x	14.3x	12.3x
Median			2,267,633	2,272,766	6.0x	5.7x	5.1x	13.4x	12.7x	11.0x
NVIDIA										
NVIDIA	United States	\$181.59	4,481,096	4,432,665						
Base		\$189.93		4,650,062	63.2x	44.3x	31.3x	55.2x	37.0x	26.1x
Bear		\$142.91		3,498,820	47.6x	37.1x	28.3x	41.4x	30.9x	23.6x
Bull		\$246.54		6,035,930	82.1x	52.2x	34.3x	71.8x	43.6x	28.6x

Exhibit 5

Implied Multiples						
Base						
Justified Mult	Year 0	Year 1	Year 2	Year 3	Implied Equity Value	\$ 4,650,062.42
P/E	63.2x	44.3x	31.3x	22.3x	Diluted shares outstanding	24,483.00
Price/Sales	35.6x	22.9x	17.2x	12.5x	Implied Equity Value per Share	\$ 189.93
EV/Sales	35.3x	22.6x	17.0x	12.3x		
EV/EBITDA	55.2x	37.0x	26.1x	18.5x		
EV/EBIT	55.2x	37.0x	26.1x	18.5x		
EV/NOPAT	63.7x	44.3x	31.2x	22.1x		
EV/FCFF	64.0x	45.3x	32.9x	22.3x		
Bear						
Justified Mult	Year 0	Year 1	Year 2	Year 3	Implied Equity Value	\$ 3,498,819.54
P/E	47.6x	37.1x	28.3x	21.6x	Diluted shares outstanding	24,483.00
Price/Sales	26.8x	18.2x	14.7x	11.5x	Implied Equity Value per Share	\$ 142.91
EV/Sales	26.4x	17.9x	14.5x	11.3x		
EV/EBITDA	41.4x	30.9x	23.6x	18.0x		
EV/EBIT	41.4x	30.9x	23.6x	18.0x		
EV/NOPAT	47.7x	37.1x	28.2x	21.5x		
EV/FCFF	48.0x	36.9x	30.3x	21.8x		
Bull						
Justified Mult	Year 0	Year 1	Year 2	Year 3	Implied Equity Value	\$ 6,035,929.71
P/E	82.1x	52.2x	34.3x	23.6x	Diluted shares outstanding	24,483.00
Price/Sales	46.3x	27.6x	19.4x	13.6x	Implied Equity Value per Share	\$ 246.54
EV/Sales	45.9x	27.4x	19.3x	13.5x		
EV/EBITDA	71.8x	43.6x	28.6x	19.6x		
EV/EBIT	71.8x	43.6x	28.6x	19.6x		
EV/NOPAT	82.8x	52.3x	34.3x	23.5x		
EV/FCFF	83.3x	54.4x	35.7x	23.7x		

Exhibit 6

	7.0%	8.0%	9.0%	10.0%	11.0%	Growth Rate
2.0%	\$265	\$218	\$184	\$159	\$140	
3.0%	\$323	\$255	\$210	\$178	\$154	
4.0%	\$421	\$312	\$247	\$203	\$172	
5.0%	\$616	\$406	\$301	\$238	\$196	
6.0%	\$1,202	\$594	\$392	\$290	\$230	
WACC						

Risks

- 1) Export controls & China exposure: Further tightening of U.S. export rules could restrict high-end GPU sales to China or push NVIDIA into lower-ASP parts, cutting into a key high-margin region and likely trimming ~10–20% of upside by pulling my bull case toward the base case.
Monitoring: NVIDIA China commentary, U.S. policy moves, and China vs. non-China data-center mix.
- 2) Supply chain shocks (HBM, CoWoS, substrates): Delays in HBM adds or CoWoS could push out Blackwell ramps and cap revenue, turning what I treat as “phasing risk” into a more material hit to the next 1–2 years of growth and delaying my valuation timeline by up to two years.
Monitoring: SK hynix/Samsung/Micron/TSMC commentary, CoWoS capacity adds, yields, and lead times.
- 3) Power & grid constraints: Limited power and cooling at major data-center hubs could delay large “AI factory” projects and make results lumpier than I model, potentially suppressing near-term revenue/FCF by high single digits vs. base
Monitoring: Utility interconnect queues, on-site power announcements, and permitting timelines in key regions.
- 4) Competition from AMD & in-house chips: Faster-than-expected adoption of AMD accelerators or custom ASICs for key workloads could compress NVIDIA’s unit share and pricing, taking 5–10 pts off data-center growth and pulling my premium EV/EBITDA and P/E multiples back toward the high-teens.
Monitoring: Hyperscaler disclosures on accelerator mix, AMD data-center revenue, and customer case studies
- 5) Management execution risk: Mis-steps on the roadmap (e.g., Blackwell delays), weak software execution, or poor capital allocation would directly undermine the thesis that NVIDIA can sustain high growth and margins and could take 20–30% off my target price via lower earnings and multiple compression.
Monitoring / diligence: Questions on supply commitments, software KPIs and adoption, export-control planning, and capital allocation between supplier prepayments, R&D, and buybacks.