NVIDIA CORPORATION

Nasdaq: NVDA | AI accelerators, rack-scale systems, NVLink, InfiniBand and Ethernet networking | Santa Clara, California

Tommy Lau | 2 November 2026 | Initiation of coverage | Independent research

LONG	**Target price:** US$245.00	**Last price:** US$204.64 (Nasdaq close, 30 October 2026)	**Upside:** +20%

Nvidia already sells about twelve gigawatts of AI a year. Consensus needs half as much again next year; the price assumes the power stops growing after that.

Nvidia designs the chips, racks and networks inside AI data centres. My journal pieces argue that power, not chips, sets the pace of AI. Nvidia's own filings now agree: it lists energy among the things that move its demand, and since August it has guaranteed up to US$108.5bn of its customers' land, power and buildings. On its own yardstick, about US$41bn of revenue for each gigawatt of computers, its revenue this year equips about 12 gigawatts.

The FY28 consensus needs about 51% more gigawatts than that, which the grid and equipment queues on this site make hard. I assume 20% and then 15%, so my FY28 EPS of US$12.25 is 23% below consensus. Yet at 16 times, the price needs FY29 EPS of only US$12.79: it pays for one more year of growth and little after.

My call: LONG, low conviction, target US$245. My base case is worth US$244, up 19%. The bear case is US$114 (down 44%) and the bull case US$394 (up 93%). Conviction is low because the growth in powered gigawatts is my judgement, and no one publishes it.

Stock data

LAST PRICE
US$204.64, Nasdaq close 30 October 2026; 52-week range US$164.27 to 243.37 (intraday)

MARKET VALUE
About US$4.93 trillion on 24.1bn shares (21 Aug 2026)

NET CASH, 26 JUL 2026
US$23.2bn: cash and debt securities US$56.6bn, debt US$33.4bn

P/E, MY ESTIMATES
22x FY27E, 17x FY28E, 13x FY29E (FY27 ends about 31 Jan 2027)

CONSENSUS (STOCKANALYSIS.COM)
EPS US$9.31 FY27, US$15.91 FY28; average target US$328.95

Q3 FY27 GUIDE (26 AUG)
Revenue US$108.0bn, plus or minus 2%; gross margin 74.0%; no Data Center compute revenue from China assumed

GUARANTEES
Up to US$108.5bn for customers' land, power and shell, US$105bn of it for 4.25 GW in Ohio

SUPPLY COMMITMENTS
US$279bn at 26 Jul 2026, up from US$119bn a quarter earlier, mainly memory

NEXT RESULTS
Q3 FY27, 17 November 2026 (Nvidia release, 28 Oct 2026)

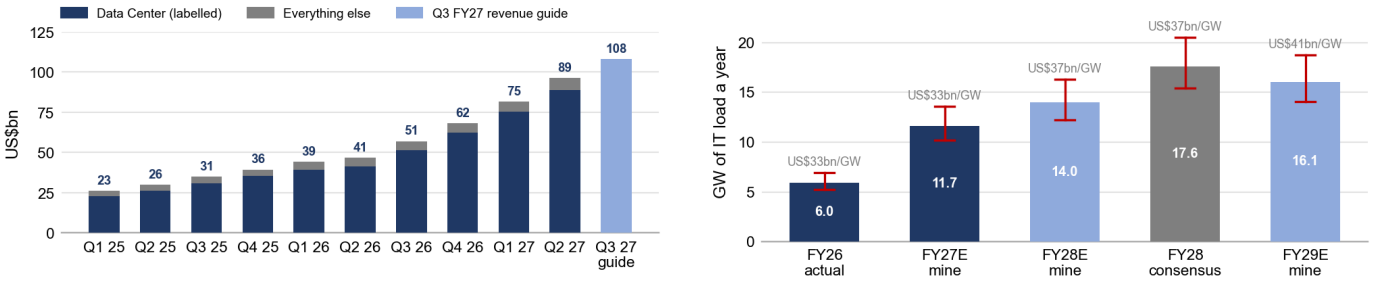
Key financials (non-GAAP from FY27; FY27 ends about 31 January 2027)

US$ billion unless stated	FY25A	FY26A	FY27E*	FY28E*	FY29E*
Data Center revenue	115.2	193.7	379.3	523.4	662.1
Edge Computing and other	15.3	22.2	28.6	33.0	37.0
Revenue	130.5	215.9	407.8	556.4	699.1
Revenue growth (%)	114	65	89	36	26
GW of IT load equipped (my yardstick)		6.0	11.7	14.0	16.1
Non-GAAP operating margin (%)			65.9%	63.0%	62.0%
EPS, US$ (GAAP to FY26)	2.94	4.90	9.32	12.25	15.26
Consensus EPS (US$)			9.31	15.91	n/v

P/E at US$204.64 (x)	41.8	22.0	16.7	13.4

**FY27E is two reported quarters, the Q3 guide (US$108.0bn at 74.0% gross margin and US$9.0bn of costs) and my US$122bn Q4; FY28E and FY29E are my estimates. From Q1 FY27 Nvidia's non-GAAP measures include stock-based pay; FY25 and FY26 non-GAAP EPS on the old basis were US$2.99 and US$4.77, so history is shown on GAAP. Gigawatts are my translation at Nvidia's own PORTS yardstick (see the thesis). Consensus from stockanalysis.com (S&P Global data, updated 26 October 2026); FY28 is the page's 'next year' figure; n/v: FY29 consensus not verified. History from Nvidia's CFO commentaries.*

Two charts: the business, and the gigawatts it needs



Left: revenue by fiscal quarter, Data Center (labelled) and everything else, with the Q3 FY27 guide for total revenue (CFO commentaries). Right: Data Center revenue divided by Nvidia revenue per gigawatt of IT load, my arithmetic; the red bars span Nvidia's US$150bn to US$200bn for 4.25 GW at PORTS, set to FY29 and deflated by my assumed rises for earlier years. FY28 consensus is revenue less my Edge estimate. A yardstick, not a count.

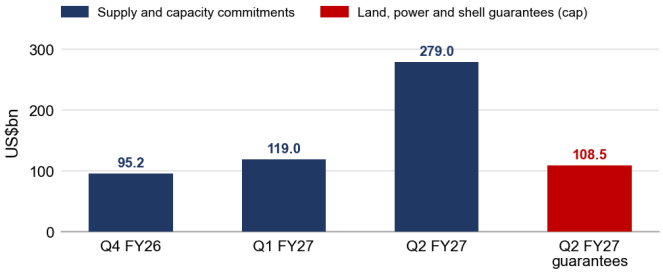
The thesis: the furniture maker waits for the house

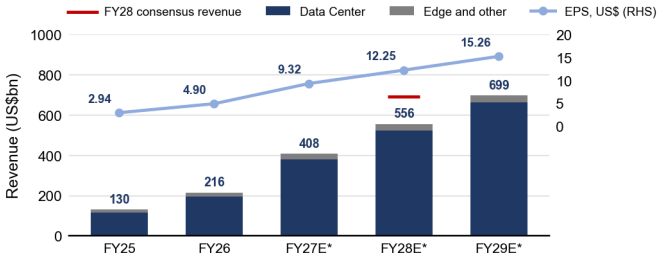
Go back to the housing estate in my piece on how AI capacity arrives. Houses are sold off plan, built, connected to the mains, handed over and lived in. Nvidia sells the furniture. For three years every house came with a removal van booked, so the furniture maker sold as fast as it could make sofas. Now the houses are the slow part, and the four claims this note rests on point to the same place: Nvidia has bought its inputs years ahead and sets its suppliers' designs, but it cannot buy a switched-on hall.

The chips now wait on the plug. Nvidia's 10-K for the year to 25 January 2026 lists whether customers can find enough data centre capacity "or energy" among the things that move its demand. Its CFO commentary of 26 August says securing land, power and shell "has become the next critical phase in the AI infrastructure buildout". On 17 August it guaranteed leases for about 4.25 GW of IT load at SB Energy's PORTS campus in Ohio, which OpenAI will rent and fill only with Nvidia systems. The guarantees are capped at US$105bn and take effect as halls become ready for service, expected from 2028. With US$3.5bn of smaller guarantees for AI clouds, Nvidia stands behind up to US$108.5bn. That is the stage my piece called the last one a builder controls, energised power, and the chip maker now pays to hurry it.

Nvidia's revenue, in gigawatts. Nvidia gave the yardstick: each generation at PORTS "could represent approximately 1.5 million NVIDIA GPUs, or approximately $150 billion to $200 billion in NVIDIA revenue". Over 4.25 GW that is US$35bn to US$47bn per gigawatt of computers, 9 to 14 times the US$3.25m to US$3.75m a megawatt Vertiv puts on its power and cooling gear. The halls are ready from 2028, so I set the US$41bn midpoint to FY29 and deflate it by my own assumed rises: US$37bn in FY28 and US$33bn in FY27. On my FY27 Data Center revenue of US$379bn that is about 11.7 GW of IT load in a year; on the undeflated PORTS range, 8 to 11 GW.

Set that against what the site has measured. Texas has been asked for 474 GW of connections and had cleared 6.6 GW to the last step by 20 August. GE Vernova ships about 20 GW of gas turbines a year in all, with heavy turbines largely sold out to 2030. Nebius expects 5 GW contracted at the end of 2026 against at most 1 GW connected. The grid makers hold one to four years of orders. The FY28 consensus of US$692bn, less my US$33bn for edge devices, needs about US$659bn of Data Center revenue: with revenue per gigawatt up 15%, about 51% more gigawatts, roughly 17.6 against 11.7. I assume 20% and then 15%. That is faster than the makers add factories, because energised halls stand partly idle (Malaysia's data centres drew 1.26 of the 5.65 GW booked in June) and buyers add power of their own. It is my judgement, not a published figure, and the call turns on it.





Left: supply and capacity commitments at each quarter end, and the cap on land, power and shell guarantees at 26 July 2026 (CFO commentaries; Form 8-K of 17 August 2026). Right: revenue by fiscal year, Data Center and the rest, with EPS (GAAP to FY26, non-GAAP after) and the FY28 consensus revenue (red); FY27 to FY29 are my estimates.

The other three claims: Nvidia holds the rest of the chain. On memory, its commitments rose from US$119bn to US$279bn in one quarter, "primarily related to the procurement of memory", with US$87bn due in FY28. Its NVHBM design (26 August) puts Nvidia's controller in the base of each memory stack, "available from multiple memory providers": Nvidia draws the part and chooses the maker. On the rack, Nvidia sets the 800 volt design and dates its power rack to the second half of 2026 and its row power centre to 2027, the building step undated, so new racks can go into halls whose power rooms already exist. On the network, its Vera Rubin blueprints still specify pluggable switches; networking was US$14.8bn of US$75.2bn of Data Center revenue in the quarter to April, the last quarter with the split.

Each of the four claims holds on Nvidia's filings and releases.

Variant perception: what the market prices in, and where I differ

At US$204.64 Nvidia is worth about US$4.93 trillion, with US$23.2bn more cash and debt securities than debt. The shares closed at a high of US$239.24 on 6 October (intraday US$243.37) and are 14.5% below it. They trade at 22 times my FY27 EPS of US$9.32, which matches the US$9.31 consensus, and at 12.9 times the FY28 consensus. Twelve to thirteen times next year's consensus is low for a company still doubling its main business: the market does not believe the consensus will last.

Against consensus. Consensus needs the gigawatts equipped with Nvidia systems to rise about half in a year. My FY28 revenue is US$556bn, 20% below consensus, and EPS US$12.25, 23% below. **Against the price.** By October 2027 the market will price FY29, the year to January 2029. At 16 times, US$204.64 needs FY29 EPS of US$12.79. If gigawatts and revenue per gigawatt both stopped rising after FY28, my model gives US$12.24. The price assumes growth roughly stops after next year; the physical layer says it slows. Equipment makers are adding plants, Nvidia is paying to bring 4.25 GW forward, and each generation earns more per gigawatt.

The obvious objection is that chips already ship ahead of the power to run them. Receivables reached US$63.1bn in July, 60 days of sales against 45, on "extended payment terms on large, multi-quarter agreements"; the 10-Q says such terms run from 90 days up to a year for large data centre builds. If systems are reaching halls not yet energised, a pause follows. That is the bear case, and why the first test of the call is a quarter of falling Data Center revenue.

Valuation, with the working

I value Nvidia on EPS: it holds more cash than debt and its profit converts largely to cash. Revenue in my model is gigawatts equipped times revenue per gigawatt, plus edge devices. Base: gigawatts up 20% in FY28 and 15% in FY29, revenue per gigawatt up 15% and 10%, non-GAAP operating margin 63% and 62% as memory costs rise (the Q3 guide already steps gross margin down a point to 74%), tax 16.5% and 23.8bn shares. That gives FY29 EPS of US$15.26. I use 16 times, about today's 16.9 times forward multiple, rounded down: US$244.

Scenarios: three cases, twelve months out

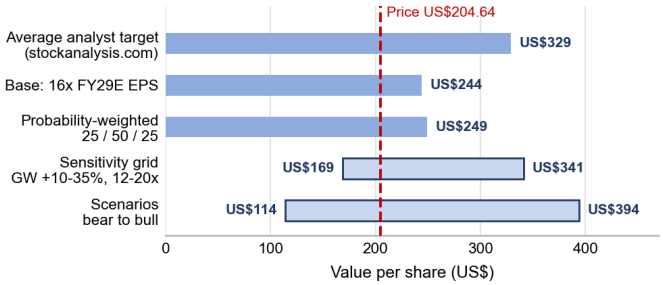
Case	What happens	GW FY28, FY29	US$/GW	FY29 EPS	Multiple	Value	Change	Weight
Bear	Chips outrun the power; growth stalls	+5%, flat	+10%, +5%	US$10.34	11x	US$114	down 44%	25%
Base	Power paces growth	+20%, +15%	+15%, +10%	US$15.26	16x	US$244	up 19%	50%
Bull	Consensus pace arrives	+50%, +30%	+15%, +10%	US$21.89	18x	US$394	up 93%	25%
Weighted	**25 / 50 / 25**					**about US$249**	**up 22%**	

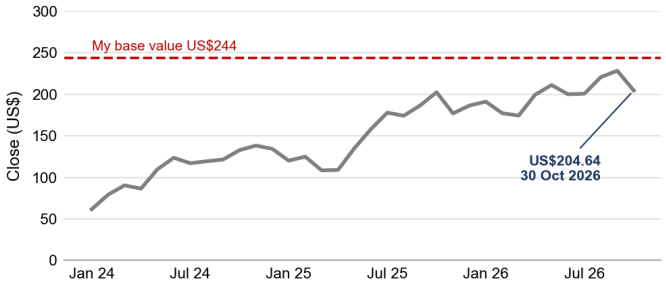
Operating margins: bear 60% and 59%, base 63% and 62%, bull 64%. The bull case gives FY28 EPS of US$15.35, close to the US$15.91 consensus. Gigawatts: bear 12.2, base 16.1, bull 22.7 GW in FY29. The bear case is not fanciful: the shares closed at US$149.43 on 6 January 2025 and at US$94.31 on 4 April 2025, 37% lower, and Nvidia's guarantees reach 4.7 times its net cash. The band: long when my base value is at least 15% above the price, short when at least 25% below, no call in between.

Sensitivity: FY28 gigawatt growth against the multiple

FY28 GW growth	12x	16x	20x
+10%	US$169	US$225	US$281
+20%	US$183	**US$244**	US$305
+35%	US$205	US$273	US$341

Seven of the nine cells sit above today's price, and five by more than 15%. The two below it both need the multiple to fall to 12 times, with growth of 20% or less; at 12 times even 35% growth only matches the price. FY29 growth is held at 15% and every other input at the base. Every figure is live in the accompanying model.





Left: value per share from the base case, the weighted cases, the sensitivity grid and the bear to bull range, against the US$204.64 price (red) and the stockanalysis.com average analyst target. Right: NVDA month-end close, January 2024 to the 30 October 2026 close (Nasdaq.com, retrieved 2 November 2026), against my base value.

Peer landscape

Company	Ticker	Forward P/E	What they make
Nvidia	Nasdaq: NVDA	16.9x	GPUs, rack-scale systems, NVLink, InfiniBand and Ethernet networking
Broadcom	Nasdaq: AVGO	18.7x	Custom AI chips and Ethernet switch chips
TSMC	NYSE: TSM	19.1x	Foundry and CoWoS packaging for Nvidia's chips
AMD	Nasdaq: AMD	45.2x	GPUs and CPUs for data centres
Marvell	Nasdaq: MRVL	34.8x	Custom AI chips and optical DSPs
Micron	Nasdaq: MU	5.0x	DRAM and HBM, one of Nvidia's three memory suppliers

Forward P/E from stockanalysis.com at the 30 October 2026 close, retrieved 2 November 2026. Five-peer median 19.1 times; Micron's multiple reflects peak memory earnings. Multiples only; no view on the peers' shares.

Catalysts

When	Event	What I watch
17 November 2026	Q3 FY27 results (quarter to 25 October)	Data Center against the US$108.0bn total guide; gross margin against 74%; receivables and payment terms
Around late February 2027 (expected)	Q4 FY27 results and 10-K	Supply commitments; guarantees; any new land, power and shell deals
March 2027	GTC	A date for the DC power block and NVHBM; Kyber; Feynman's network
From 2028	First PORTS halls ready for service	Whether 4.25 GW arrives on time, and the guarantees take effect

Risks: what would hurt the view

Chips outrun the power. If buyers hold systems they cannot switch on, orders pause until halls catch up; longer terms and higher receivables fit, though they prove nothing yet. This is the bear case, a 44% fall. **The guarantees.** Up to US$108.5bn, 4.7 times net cash; for Ohio, OpenAI has agreed to reimburse Nvidia for any payment, so the exposure is to OpenAI's ability to pay. Nvidia also has US$36bn of AI cloud agreements and US$20bn of leases it expects to pass to others. **Concentration and custom chips.** One direct customer took 16% of revenue in the July quarter, and every large buyer designs chips of its own.

Memory and margin. US$279bn of commitments, mostly memory, lock in costs as prices rise; a gross margin below 70% would mean Nvidia cannot pass them on. **China.** Nvidia's guides assume no Data Center compute revenue from China; any licence is upside I do not count.

What would change my mind

Out of the long: a price above about US$212 with my numbers unchanged, which leaves the base less than 15% above it; or evidence that powered gigawatts for Nvidia grow 10% or less while the multiple falls towards 12 times. **Into more conviction:** a grid operator or equipment maker showing the queues shortening, or the first PORTS halls ready early, would move the base towards the bull case.

The call is wrong if Data Center revenue falls from one quarter to the next in any quarter Nvidia reports before October 2027, or its non-GAAP gross margin falls below 70% in any of those quarters, or Nvidia makes a payment under any of its land, power and shell guarantees.

Conclusion

The pieces this note rests on argue that power, not chips, sets the pace of AI. Nvidia's filings now agree: it lists energy among its demand risks and guarantees its customers' land, power and buildings. On its own PORTS yardstick its revenue already equips about 12 gigawatts a year, and consensus needs half as much again next year. I do not believe that pace, and my FY28 EPS is 23% below consensus. But the price goes further and assumes growth ends after next year. My call: LONG, low conviction, target US$245; base value US$244, range US$114 to US$394.

Appendix: forecast assumptions

Assumption	FY27E	FY28E	FY29E	Basis
Revenue, US$bn	407.8	556.4	699.1	FY27: two quarters, the Q3 guide and my Q4 of US$122bn
Growth in GW equipped		20%	15%	Mine: the pace of the queues, plus idle energised halls and on-site power
Rise in revenue per GW		15%	10%	Mine: Vera Rubin began production shipments in Q3 FY27 (10-Q)
Revenue per GW, US$bn	32.6	37.4	41.2	PORTS midpoint set to FY29; earlier years deflated (mine)
Edge Computing, US$bn	28.6	33	37	Q1 and Q2 reported; mine after
Non-GAAP operating margin	65.9%	63%	62%	Q2 66.5%; Q3 guide implies 65.7%; memory costs (mine)
Tax rate (non-GAAP)	16.5%	16.5%	16.5%	Guide 16% to 18% for FY27
Diluted shares, bn	24.2 (H2)	24.0	23.8	Buybacks net of dilution (mine)
Multiple on FY29 EPS			16x	Forward P/E 16.9x, rounded down

Full workings, with live formulas, in the accompanying model, 2026-11-02_Nvidia_Model.xlsx.

Sources

Nvidia CFO commentaries for Q3 FY25 to Q2 FY27, with comparative columns, for revenue, Data Center and sub-market revenue, non-GAAP margins and EPS, commitments, guarantees, balance sheet and the Q3 FY27 outlook. Nvidia Form 10-K for the year to 25 January 2026 (filed 25 February 2026): memory suppliers and demand risk factors. Form 10-Q for the quarter to 26 July 2026 (filed 26 August 2026): balance sheet, shares outstanding, payment terms, customer concentration, Vera Rubin production shipments and the SB Energy guarantees. Form 8-K of 17 August 2026 (PORTS campus). Nvidia release of 28 October 2026 (results date). Nvidia blogs of 11 August 2026 (800 VDC) and 26 August 2026 (NVHBM) and the DGX SuperPOD Vera Rubin reference architectures (15 September 2026), as cited in the journal pieces, which also give the ERCOT, GE Vernova, Nebius, Tenaga and equipment backlog figures. Share prices from Nasdaq.com historical data; consensus, average target and multiples from stockanalysis.com; all retrieved 2 November 2026. Gigawatt translations and estimates for FY27 to FY29 are the author's own.

About this note

I wrote this note and built the accompanying model independently. It is not a product of any bank or broker. It tests claims from my journal pieces, chiefly thephysicallayer.fyi/journal/contracted-energised-running, with memory-not-maths, powering-the-rack and ten-thousand-chips-one-thought; all notes and their track record are at thephysicallayer.fyi/research. Figures from company filings are labelled as such; derived figures are marked; the forecasts and view are my own estimates and my own view. I hold no position in Nvidia. Personal research, not investment advice.

Tommy Lau | The Physical Layer | thephysicallayer.fyi