RESEARCH / EARNINGS AND VALUATION

NASDAQ: NVDA

NVIDIA Stock Q3 2026: Vera Rubin, a $108bn Guide and the Road to Fiscal 2028

BY QSI LABS RESEARCH, QSI LABS PUBLISHED 2 NOV 2026 UPDATED 2 NOV 2026

FIGURES AS OF Q2 FY2027 RESULTS AND Q3 2026 GUIDANCE; PRICES AS OF THE 30 OCT 2026 CLOSE

Our NVIDIA stock Q3 2026 review of NVIDIA Corporation (NASDAQ: NVDA) starts from a share price of $204.64, the 30 October 2026 close, about 16% below its 52-week high of $243.37 (Yahoo Finance). The business has never been bigger. In the quarter to 26 July 2026 NVIDIA booked $96.2bn of revenue, up 106% year on year, guided to $108.0bn for the quarter to 25 October, and told investors to expect roughly 70% revenue growth in fiscal 2028 (Q2 FY2027 press release; earnings call). The open questions are margins, memory costs, the Vera Rubin ramp and how much of the build-out NVIDIA is now financing itself.

The short answer. NVIDIA (NVDA) stock looks fully valued after Q3 2026. NVIDIA's latest quarter was the strongest in its history, and the Q3 2026 guide implies another $12bn of sequential growth. Our discounted cash flow model puts fair value at **$182.88 a share**, about 11% below the $204.64 close on 30 October 2026: the shares already price in growth somewhat ahead of our base case, and the next leg depends on Vera Rubin shipping on time with gross margins back at 72-73%.

> **About this analysis.** Written by QSI Labs Research, the equity research team at QSI Labs. Reported figures come from NVIDIA's SEC filings (Form 10-Q and Form 8-K), its CFO commentary and its Q2 FY2027 earnings call. Forecasts and the fair value come from the QSILabs NVIDIA discounted cash flow model, valuation date 30 October 2026. Figures are as of NVIDIA's latest reported quarter (Q2 FY2027, to 26 July 2026) and its guidance for the Q3 2026 quarter; prices are as of the 30 October 2026 Nasdaq close.

A note on timing. NVIDIA's fiscal year ends in late January, so the calendar Q3 2026 quarter is its fiscal Q3 FY2027 (27 July to 25 October 2026). Those results are due on 17 November

2026 (NVIDIA IR). Until then, the latest reported figures are fiscal Q2 FY2027, which we compare with the same quarter a year earlier, alongside the company's guidance for Q3 2026.

Key takeaways: NVIDIA stock Q3 2026

- **Fair value.** The QSILabs fair value for NVIDIA (NVDA) is $182.88 per share as of Q3 2026, against a $204.64 close on 30 October 2026, which puts the shares about 12% above fair value.
- **Latest results.** NVIDIA (NVDA) reported record revenue of $96.2bn for Q2 FY2027, the quarter to 26 July 2026, up 106% year on year, with Data Center revenue of $89.0bn and GAAP EPS of $2.46 (8-K Ex 99.1).
- **Q3 2026 guidance.** NVIDIA guided Q3 FY2027 revenue, the quarter to 25 October 2026, to $108.0bn plus or minus 2%, with a 74.0% gross margin and no Data Center compute revenue from China assumed; results are due on 17 November 2026 (CFO commentary; NVIDIA IR).
- **Growth and margins.** NVIDIA expects revenue growth of about 70% in fiscal 2028 (the year to January 2028), limited by supply, while memory costs push its gross margin to a 71-72% trough in Q4 FY2027 before 72-73% in FY2028 (Q2 FY2027 call).
- **Roadmap.** NVIDIA's Vera Rubin platform began production shipments in August 2026 and should be about 20% of Q3 FY2027 Data Center revenue; Rubin Ultra is dated for the second half of 2027 (Q2 FY2027 call; NVIDIA blog).

$96.2bn Q2 FY2027 revenue, +106% YoY 8-K EX 99.1	**$89.0bn** Data Center revenue, +117% YoY CFO COMMENTARY
$108.0bn Q3 FY2027 revenue guide, +/-2% CFO COMMENTARY	**~70%** FY2028 revenue growth outlook Q2 FY2027 CALL
$2.46 GAAP diluted EPS, from $1.08 8-K EX 99.1	**$182.88** QSILabs fair value per share QSILABS DCF

NVIDIA earnings: how did the latest quarter compare with last year?

NVIDIA key numbers: Q2 FY2027 (quarter to 26 Jul 2026) vs Q2 FY2026. Sources: Q2 FY2027 press release and CFO commentary (8-K, 26 Aug 2026).

METRIC	Q2 FY2027	Q2 FY2026	CHANGE
Revenue	$96,221m	$46,743m	+106%
Data Center revenue	$89,023m	$41,096m	+117%
Hyperscale	$48,710m	$24,168m	+102%
AI clouds, industrial and enterprise	$40,313m	$16,928m	+138%
Edge Computing revenue	$7,198m	$5,647m	+27%
GAAP gross margin	75.0%	72.4%	+2.6 pts
GAAP operating income	$63,734m	$28,440m	+124%
GAAP net income	$59,688m	$26,422m	+126%
GAAP diluted EPS	$2.46	$1.08	+128%
Non-GAAP diluted EPS	$2.22	$1.01	+120%
Operating cash flow	$24.1bn	$15.4bn	+56%

Every line roughly doubled. Revenue rose from $46.7bn to $96.2bn and GAAP gross margin widened by 2.6 points to 75.0%, as the mix moved to Blackwell Ultra (CFO commentary). GAAP operating expenses rose 55% to $8.4bn, driven by compute infrastructure for research and by compensation, yet operating income grew faster than revenue, at 124%.

Below the operating line, NVIDIA booked $7.8bn of net gains on equity securities, which is why GAAP EPS ($2.46) was higher than non-GAAP EPS ($2.22). The GAAP tax rate was 16.5% (CFO commentary). From the first quarter of FY2027, NVIDIA's non-GAAP figures no longer exclude stock-based compensation, and prior periods have been restated on the same basis (8-K Ex 99.1).

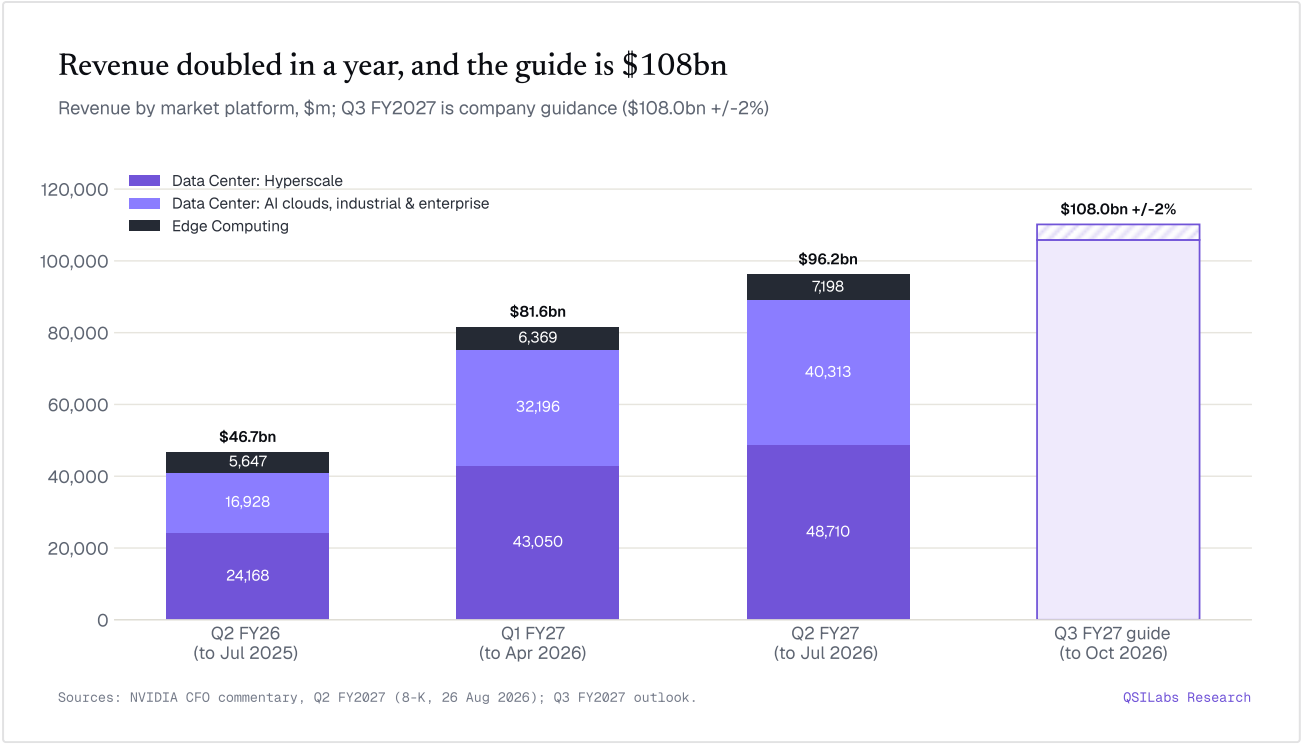


FIGURE 1 Revenue more than doubled in a year to $96.2bn. Hyperscale rose 102% and AI clouds, industrial and enterprise 138%. The Q3 FY2027 guide of $108bn would add another $12bn.

Concentration is still high. One direct customer accounted for 16% of revenue in the quarter, against two customers at 23% and 16% a year earlier (Form 10-Q). In the first half of FY2027, three direct customers were 16%, 15% and 13% of revenue.

Data Center, Hyperscale and Edge: where NVIDIA's growth came from

NVIDIA's growth came overwhelmingly from Data Center, which was $89.0bn of the $96.2bn quarter and rose 117% year on year (8-K Ex 99.1). NVIDIA now reports two market platforms: Data Center, split into Hyperscale and AI clouds, industrial and enterprise (ACIE), and Edge Computing. By reportable segment, Compute & Networking was $88.3bn (+114%) and Graphics $7.9bn (+46%) (CFO commentary).

Hyperscale: $48.7bn, up 102%

NVIDIA's Hyperscale revenue, from the largest cloud providers, more than doubled to $48.7bn and rose 13% on the quarter, on the strength of Blackwell Ultra. Management said the top five hyperscalers are expected to spend nearly $800bn on capex in 2026 and $1.3tn in 2027, against a cloud industry backlog of more than $2tn. AWS will deploy an additional 2 million NVIDIA GPUs from this quarter through Q2 FY2029, along with Vera CPUs (Q2

FY2027 call). Management expects Hyperscale growth to re-accelerate in Q4 and into FY2028 as Vera Rubin supply grows.

AI clouds, industrial and enterprise: $40.3bn, up 138%

AI clouds, industrial and enterprise (ACIE) revenue rose 138% year on year to $40.3bn, making it NVIDIA's fastest-growing market, and it rose 25% on the quarter. It covers neoclouds, sovereign AI, AI start-ups and enterprises, and management expects it to represent roughly half of Data Center revenue. Neocloud partners are expected to end the year with 8 GW of installed capacity, up from about 3 GW at the end of 2025, and the sovereign business, mostly through regional neoclouds, grew 35% sequentially and more than tripled year on year (Q2 FY2027 call). Most of the Q3 sequential growth is expected to come from ACIE.

Edge Computing: $7.2bn, up 27%

NVIDIA's Edge Computing revenue, the market platform that covers everything outside Data Center, rose 27% year on year to $7.2bn and 13% on the quarter, led by Blackwell workstations, partly offset by slower consumer PC sales as memory and system prices rose (CFO commentary). New products included RTX Spark for Windows PCs with Microsoft, DGX Station for Windows and the Cosmos 3 and Alpamayo 2 Super open models for physical AI and autonomous driving (8-K Ex 99.1).

Networking, CPUs and Groq

NVIDIA's networking, CPU and inference products are growing quickly alongside its GPUs. Networking revenue grew 18% sequentially, and Spectrum-X Ethernet grew 2.6 times year on year. Grace CPU revenue passed $5bn on a trailing twelve-month basis, the next-generation Vera CPU is in full production, and management's preliminary view is that CPU revenue more than doubles in FY2028. Groq 3 LPX, NVIDIA's first rack-scale LPU system for low-latency inference, is in full production, with volume shipments to early adopters starting with Nebius (Q2 FY2027 call). Shipments of Hopper products to China were less than 1% of Data Center revenue.

Margins: why NVIDIA's gross margin is heading to 71-72%

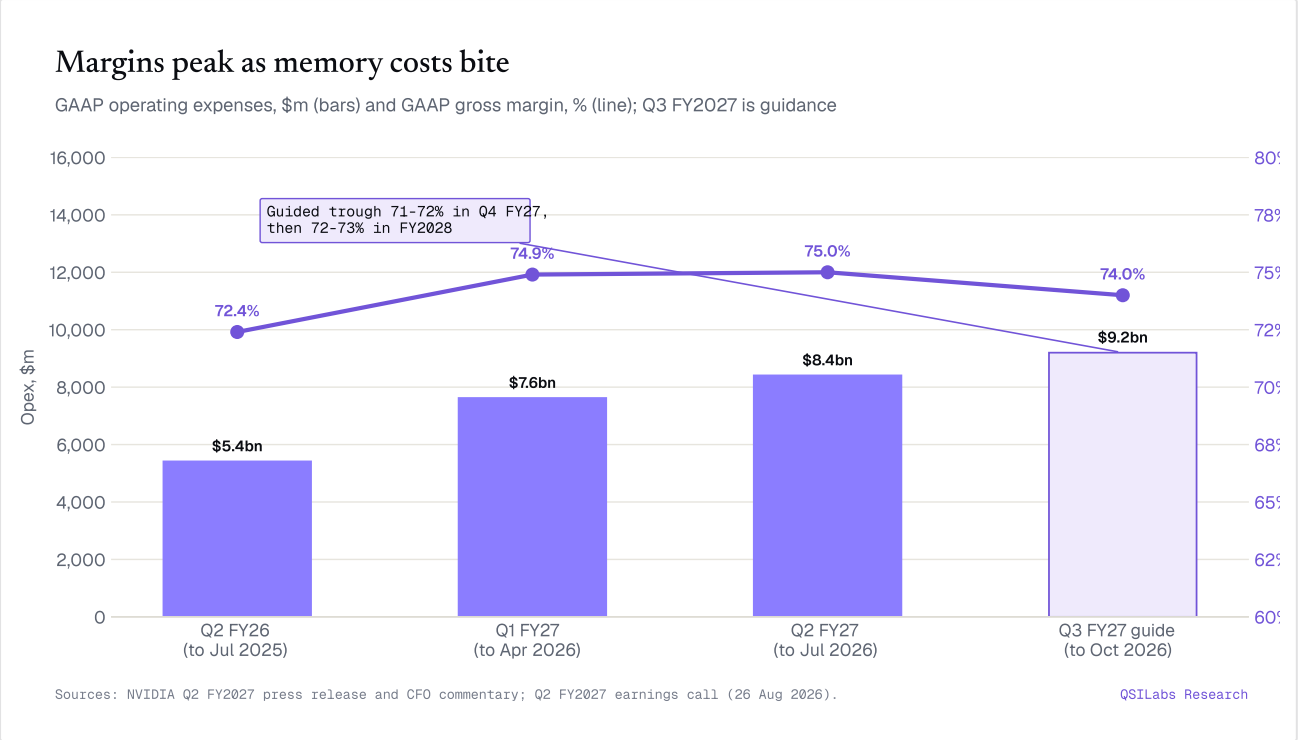


FIGURE 2 GAAP gross margin reached 75.0% in Q2 FY2027. The Q3 guide is 74.0%, with a trough of 71-72% expected in Q4 before 72-73% in FY2028. Opex is guided to $9.2bn in Q3.

NVIDIA's gross margin is falling because memory costs are rising, not because of pricing pressure. Management said memory prices have risen by more than it expected and are "headed even higher into next year", and reset its margin outlook: 74.0% plus or minus 50 basis points in Q3, a bottom of 71-72% in Q4, and 72-73% in FY2028 once price increases already agreed take effect in the first quarter (Q2 FY2027 call). Full-year operating expenses are now expected to grow in the low 50s per cent, driven by a broader product portfolio and heavier use of AI tools.

That is a small change in a very large business. A three-point fall in gross margin on more than $100bn of quarterly revenue is roughly $3bn a quarter of gross profit, but revenue is still growing by $10bn or more a quarter. In our model, the operating margin stays at about 66% from FY2028.

NVIDIA cash flow, capex and the balance sheet

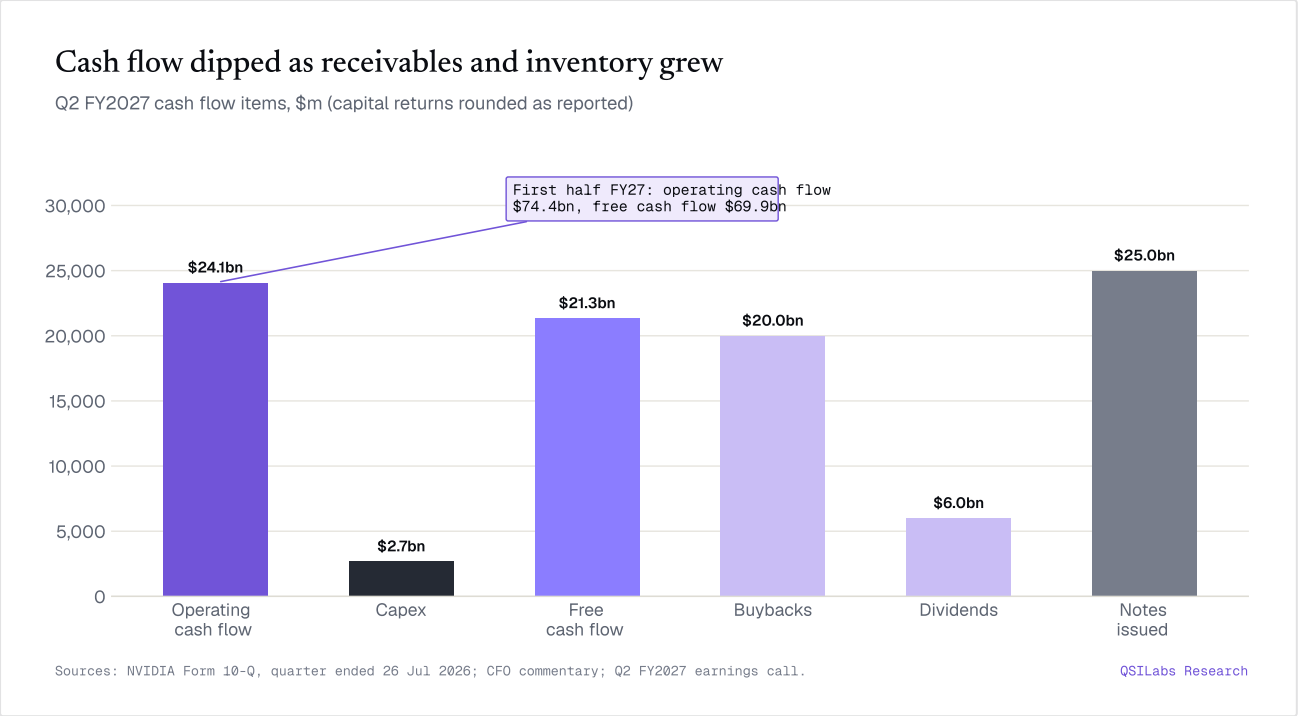


FIGURE 3 Operating cash flow was $24.1bn in Q2 FY2027, down from $50.3bn in Q1 as receivables and inventory grew. NVIDIA returned about $26bn to shareholders and raised $25bn of senior notes.

NVIDIA's operating cash flow was $24.1bn in Q2 FY2027, up from $15.4bn a year earlier but down from $50.3bn in Q1, because of higher working capital and cash taxes (CFO commentary). Capex was $2.7bn, leaving free cash flow of $21.3bn; for the first half, operating cash flow was $74.4bn and free cash flow $69.9bn (Form 10-Q).

Two working capital items stand out. Receivables were $63.1bn, with days sales outstanding up to 60 from 45, because of extended payment terms on large, multi-quarter agreements with investment-grade customers. Inventory rose to $31.6bn from $25.8bn ahead of Vera Rubin (CFO commentary).

NVIDIA returned a record $26bn in the quarter: $19.7bn of buybacks and $6.0bn of dividends at $0.25 a share (Form 10-Q). That is 60% of free cash flow year to date, against a plan of at least 50%, and $99.0bn of buyback authorisation remains (8-K Ex 99.1). It also issued $25.0bn of senior unsecured notes.

At 26 July, cash, cash equivalents and marketable debt securities were $56.6bn, against $33.4bn of short- and long-term debt. NVIDIA also held $42.8bn of marketable equity securities and $51.2bn of non-marketable equity investments, and total equity was $229.0bn (Form 10-Q).

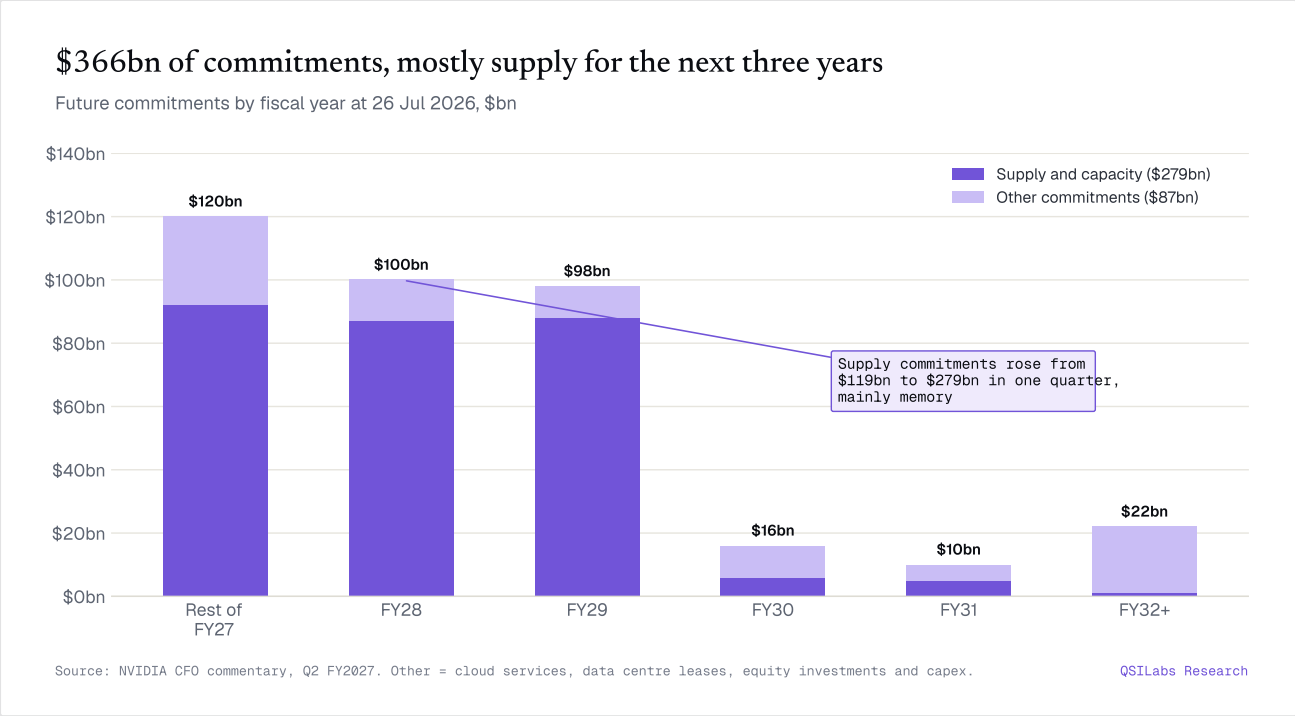


FIGURE 4 Commitments total $366bn, of which $279bn is supply and capacity, mostly memory. Nearly 90% falls in the rest of FY2027 to FY2029.

Commitments, guarantees and PORTS-Pike

NVIDIA's off-balance-sheet obligations grew sharply in the quarter. Supply and capacity commitments jumped from $119bn to $279bn in one quarter, mainly memory. Including cloud services, data centre leases, equity investments and capex, total commitments were $366bn, with another $56bn under AI cloud agreements and leases for third parties (CFO commentary).

The most unusual item is a guarantee. In August NVIDIA agreed to provide credit support for the land, power and shell at SB Energy's PORTS-Pike campus in Ohio, which will host only NVIDIA infrastructure under 20-year leases to OpenAI. The guarantees are capped at $105bn, start in phases with the first expected in fiscal 2029, and fall as OpenAI pays its leases. NVIDIA says each generation of infrastructure at the site could represent about 1.5 million GPUs, or $150-200bn of NVIDIA revenue (CFO commentary; PORTS-Pike release). Total guarantee exposure is $108.5bn.

Management was open about the criticism: "we know some will call this circular financing". It said it has invested nearly $50bn in frontier AI labs, that it is setting up financing platforms with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR to raise more than $500bn of third-party capital, and that demand from labs it supports with its balance sheet should be roughly a quarter of its business next year (Q2 FY2027 call). We look at a similar question from the other side in our analysis of Anthropic's IPO filing and its compute obligations.

What is NVIDIA guiding for Q3 2026 and fiscal 2028?

NVIDIA is guiding to $108.0bn of revenue for Q3 2026 and about 70% revenue growth in fiscal 2028. For Q3 FY2027, NVIDIA guided to revenue of $108.0bn plus or minus 2%, GAAP and non-GAAP gross margins of 74.0% plus or minus 50 basis points, and GAAP operating expenses of about $9.2bn ($9.0bn non-GAAP). The full-year tax rate is expected at 16-18% (CFO commentary). At the midpoint, revenue would be up 12% on Q2.

For FY2028, which runs to January 2028, management gave a rare full-year view: revenue growth of approximately 70%, with supply the bottleneck at least through the end of that year. Jensen Huang said demand is "much greater than 70%" and that the company has supply for about 70% (Q2 FY2027 call).

As a top-down check, 70% growth on our FY2027 estimate of about $406bn implies roughly $690bn of revenue in FY2028 (QSILabs estimate). That is a little over half of the $1.3tn of 2027 capex management expects from the top five hyperscalers alone, before neoclouds, sovereigns and enterprises, which management says are about half of Data Center demand. The number is demanding, but consistent with what NVIDIA's customers say they will spend.

NVIDIA data-centre roadmap: Blackwell, Vera Rubin, Rubin Ultra and beyond

NVIDIA DATA-CENTRE ROADMAP AND 2026 MILESTONES

MAR 2025	GTC keynote dates Vera Rubin for 2H 2026 and Rubin Ultra for 2H 2027	NVIDIA blog
6 JUL 2026	SemiAnalysis reports the Kyber rack slipping to 2028; NVIDIA: "Our roadmap is intact"	CNBC; Tom's Hardware
AUG 2026	Vera Rubin production shipments begin; Groq 3 LPX in full production	Q2 FY2027 call
17 AUG 2026	PORTS-Pike: NVIDIA backs 4.25 GW of land, power and shell for OpenAI	8-K Ex 99.1
26 AUG 2026	Q2 FY2027 results; FY2028 growth outlook of about 70%	8-K; Q2 FY2027 call
2 SEP 2026	Agreement to acquire Hugging Face for about $11.9bn	8-K
21 OCT 2026	GTC Berlin: DSX POD cuts deployment from six months to 14 days	NVIDIA Newsroom
17 NOV 2026	Q3 FY2027 results (quarter ended 25 Oct 2026)	NVIDIA IR
EARLY 2027	First partner deliveries of DSX POD AI factories	NVIDIA Newsroom

NVIDIA's data-centre roadmap runs from Blackwell (shipping now) to Vera Rubin (shipping since August 2026), Rubin Ultra (second half of 2027) and the Kyber rack. The roadmap has one product rhythm: a new data-centre platform roughly every year, each sold as a full AI factory rather than a chip. At GTC 2025 the company dated Vera Rubin for the second half of 2026 and Rubin Ultra for the second half of 2027 (NVIDIA blog).

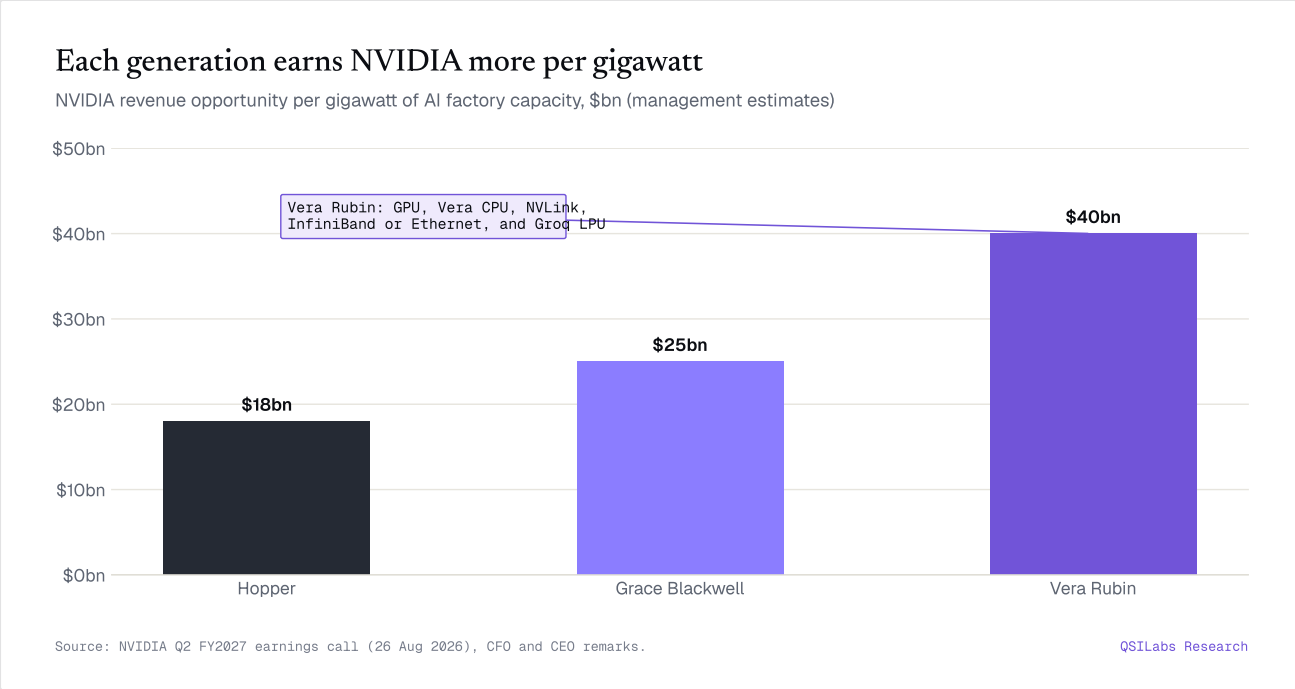


FIGURE 5 Management puts NVIDIA's revenue opportunity at about $18bn per gigawatt for Hopper, $25bn for Grace Blackwell and $40bn for Vera Rubin. More content per gigawatt means revenue can grow faster than power capacity.

Blackwell and Blackwell Ultra: today's revenue

NVIDIA's Blackwell architecture remains the "vast majority" of revenue, and Blackwell Ultra drove the Q2 increase in both Hyperscale revenue and gross margin (CFO commentary). Blackwell led every category in the MLPerf Training 6.0 benchmarks (8-K Ex 99.1).

Vera Rubin: the 2026-27 ramp

Vera Rubin is NVIDIA's next AI factory platform and is now shipping. It combines the Rubin GPU, the Vera CPU, NVLink, InfiniBand or Spectrum-X Ethernet and, for low-latency services, Groq LPUs. NVIDIA claims 30 times higher throughput per megawatt and 35 times lower cost per token than Grace Blackwell Ultra. Production shipments began in August, purchase orders have come from every major hyperscaler, AI cloud and system maker, and racks are running at CoreWeave, Google Cloud, Microsoft Azure, Oracle Cloud Infrastructure and Nebius (Q2 FY2027 call; 8-K Ex 99.1). Vera Rubin should be about 20% of Data Center revenue in Q3, which makes Q4 and FY2028 the real test of the ramp.

Rubin Ultra and Kyber: the timing question

NVIDIA still dates Rubin Ultra for the second half of 2027, but the timing of its Kyber rack is disputed. In July SemiAnalysis, a semiconductor research firm, reported that the Kyber rack, the 144-GPU system designed for Rubin Ultra, had slipped to 2028 because a key circuit

board is hard to manufacture. NVIDIA responded: "Our roadmap is intact" (CNBC; Tom's Hardware). The company has not given a new date. For investors, a later Kyber would matter mostly for competitive positioning in 2027-28, because Vera Rubin NVL72 racks, which NVIDIA says will make up "the vast majority of the world's data centers", are already shipping.

DSX POD: faster deployment

DSX POD is NVIDIA's prefabricated AI factory design. At GTC Berlin on 21 October NVIDIA introduced DSX POD, a reference design for Vera Rubin AI factories built by partners inside containers. Each pod holds seven Vera Rubin NVL72 racks, and NVIDIA says it cuts deployment at a prepared site from six months to as little as 14 days, with partner deliveries from early 2027 (NVIDIA Newsroom). It does not solve the slower problem of getting grid power to the site, but it shortens the time between a chip shipping and a customer earning revenue on it.

The demand outlook

Demand for NVIDIA's data-centre products still exceeds supply, according to management. The demand signals it cited are broad: global venture funding for AI of more than $400bn in the first half of 2026, nearly 20 AI companies with more than $1bn of annualised revenue, about 12 GW of existing and planned NVIDIA compute for OpenAI, and agentic AI using 15 to 100 times as much compute as a person prompting a model (Q2 FY2027 call). The planned $11.9bn acquisition of Hugging Face, expected to close in the first half of 2027, adds a leading platform for open models to NVIDIA's software stack (Form 8-K, 2 Sep 2026). Compute demand from AI developers is also the main driver in our SpaceX Q3 2026 analysis of xAI compute and Terafab.

Catalysts to watch

The next catalyst for NVIDIA (NVDA) stock is the Q3 FY2027 report on 17 November 2026.

- **17 November 2026: Q3 FY2027 results.** The guide is $108.0bn. The key lines are the Vera Rubin share of Data Center revenue, the Q4 guide and any update on the 71-72% margin trough.
- **Memory pricing.** The largest single swing factor in the FY2028 margin.
- **Rubin Ultra and Kyber dates.** Any confirmation or change of the second-half 2027 timing.

- **Supply.** Whether NVIDIA can close part of the gap between the 70% it can supply and the demand it says it sees.
- **Hugging Face.** Regulatory approval and closing in the first half of 2027.
- **China.** Any licence change. The outlook assumes no Data Center compute revenue from China.

Is NVIDIA stock undervalued after Q3 2026? Our fair value

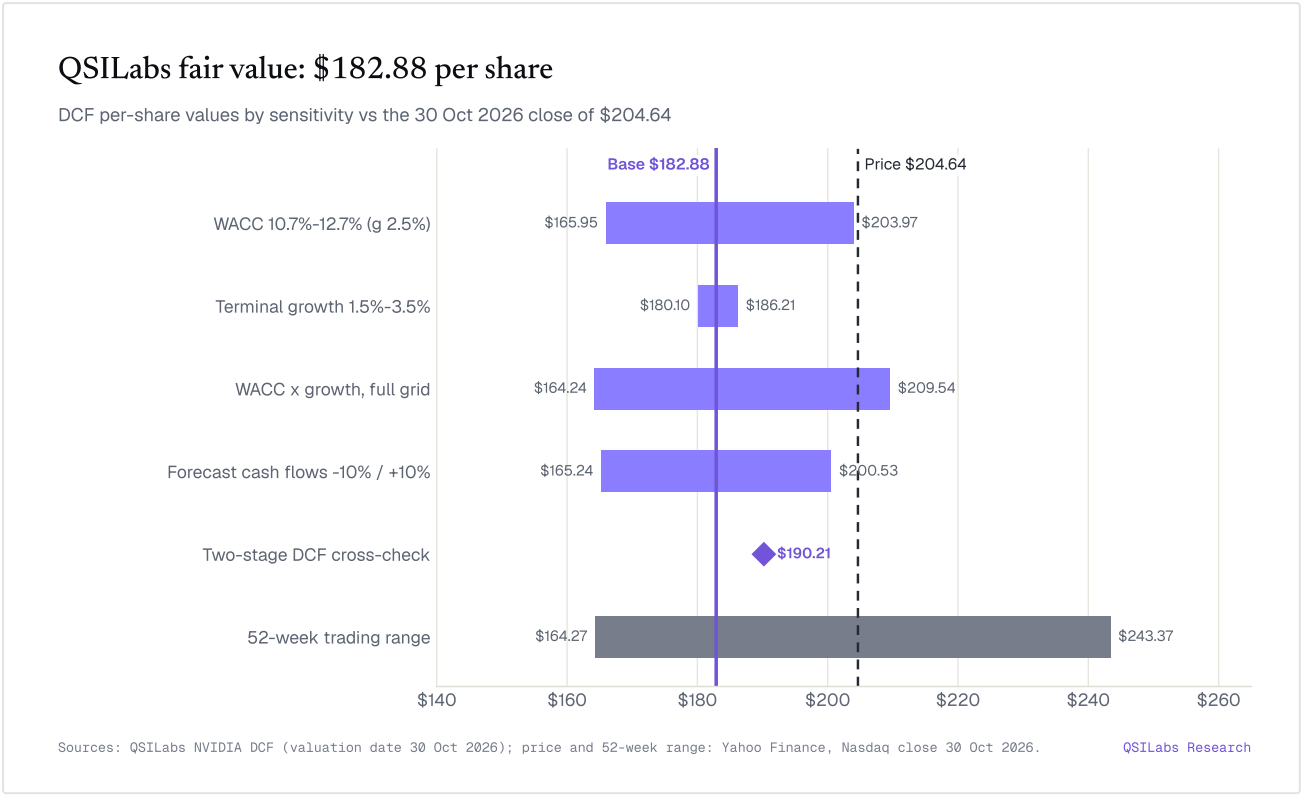


FIGURE 6 Our base case is $182.88 a share, against a $204.64 close. On single assumptions the range runs from about $164 to $210. Only the lowest discount rates, or cash flows about 12% above our forecast, reach today's price.

Our fair value for NVIDIA stock is **$182.88 a share**, about 11% below the 30 October close of $204.64. That makes the shares modestly overvalued on our numbers rather than cheap, despite the growth. The gap is not about the next year: we use management's FY2028 outlook of about 70% growth. It is about what comes after, and how much of the build-out repeats.

The market price implies either a discount rate of about 10.7% on our forecasts, or cash flows about 12% higher than ours in every year (QSILabs estimates). Both are plausible, which is

why we see the shares as fully valued rather than expensive.

Method

We value NVIDIA (NVDA) at $182.88 a share with a ten-year discounted cash flow model covering FY2027 to FY2036, valued at 30 October 2026.

- **Revenue.** FY2027 is the reported first half plus the $108bn Q3 guide and a QSILabs Q4 estimate of $120bn, about $406bn in total. FY2028 grows about 70% to $690bn, in line with management's outlook. We then model slower growth: about 20% in FY2029, 8% in FY2030, a digestion year in FY2031 with revenue down about 5%, and mid-single-digit growth to about $1.05tn in FY2036.
- **Margins.** The FY2027 operating margin is about 65%, reflecting the guided Q4 gross margin trough. From FY2028 we hold the operating margin at about 66%, consistent with a 72-73% gross margin less operating costs.
- **Cash flow.** Tax at the guided 16-18% range, and capex rising from $15bn in FY2028 to about $100bn by FY2036 as the business scales.
- **Discount rate.** A WACC of 11.7%, using a 5.12% risk-free rate, and terminal growth of 2.5%. The terminal value is 45% of enterprise value.
- **From enterprise value to equity.** Enterprise value is $4.28tn. Cash, marketable and non-marketable investments and other balance sheet items, net of debt, add $157bn, for equity value of $4.44tn, or $182.88 across 24,285m diluted shares. A two-stage version of the same model gives $190.21.

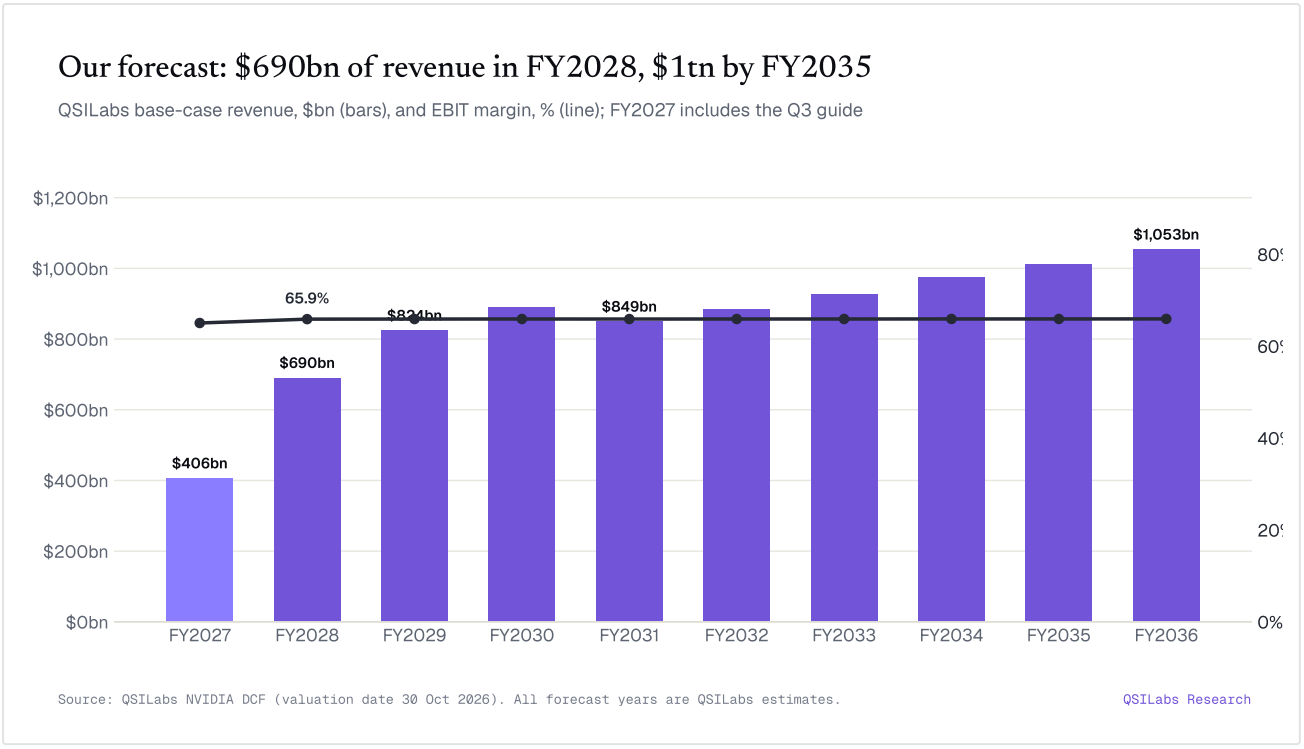


FIGURE 7 In our base case revenue reaches about $690bn in FY2028 and passes $1tn in FY2035. The operating margin holds near 66% from FY2028.



FIGURE 8 Each half-point on the discount rate moves fair value by about $8-11 a share. Terminal growth matters less: 1.5% to 3.5% moves it by about $6.

NVIDIA stock risks

The main risks to NVIDIA (NVDA) stock are a slowdown in AI capex, rising memory costs and NVIDIA's growing financial exposure to its own customers.

- **A digestion cycle.** Our model already includes one down year. A deeper or earlier pause in hyperscaler capex would hit revenue and margins at the same time.
- **Memory and component costs.** Management has already cut its margin outlook once because of memory prices, and says they are still rising.
- **Balance-sheet exposure to customers.** $108.5bn of guarantees, nearly $50bn invested in AI labs and $366bn of commitments tie NVIDIA's results more closely to a small number of fast-growing, not always profitable, customers.
- **Concentration.** One direct customer was 16% of revenue in the quarter (Form 10-Q).
- **Custom chips.** Hyperscalers and AI labs, including OpenAI, are designing their own accelerators (Q2 FY2027 call).
- **Roadmap execution.** A delayed Kyber rack would leave NVIDIA with 72-GPU systems while competitors offer larger domains.
- **China and export rules.** The outlook excludes Data Center compute revenue from China; rules could tighten further, including for open models on Hugging Face (Form 8-K, 2 Sep 2026).
- **Working capital.** Longer payment terms took receivables to 60 days, and cash flow fell by half in a quarter when revenue grew 18%.

FAQ

What is NVIDIA's fair value after Q3 2026?

The QSILabs fair value for NVIDIA (NVDA) is $182.88 per share as of Q3 2026, about 11% below the $204.64 close on 30 October 2026. It comes from a ten-year discounted cash flow model; on single assumptions the range is about $164 to $210, and a two-stage version of the model gives $190.21.

Is NVIDIA stock overvalued or undervalued right now?

On QSILabs estimates NVIDIA (NVDA) stock is modestly overvalued, trading about 12% above our $182.88 fair value at the 30 October 2026 close of $204.64. The shares would be fairly valued at a discount rate of about 10.7%, or with cash flows about 12% above our forecasts.

How did NVIDIA do in its latest earnings report?

NVIDIA reported record revenue of $96.2bn for Q2 FY2027, the quarter to 26 July 2026, up 106% year on year. Data Center revenue was $89.0bn, up 117%, and GAAP net income was $59.7bn, or $2.46 per diluted share against $1.08 a year earlier (8-K Ex 99.1).

When is NVIDIA's Q3 2026 earnings date?

NVIDIA reports its Q3 FY2027 results, covering the calendar Q3 2026 quarter to 25 October 2026, on 17 November 2026 (NVIDIA IR).

What revenue is NVIDIA expecting for Q3 2026?

NVIDIA expects Q3 FY2027 revenue of $108.0bn, plus or minus 2%. It guided gross margins to 74.0% plus or minus 50 basis points and GAAP operating expenses to about $9.2bn, and assumed no Data Center compute revenue from China (CFO commentary).

What is NVIDIA Vera Rubin and when does it ship?

Vera Rubin is NVIDIA's next AI factory platform, combining the Rubin GPU, the Vera CPU, NVLink and networking, and it began production shipments in August 2026. Management expects it to be about 20% of Data Center revenue in Q3 FY2027 (Q2 FY2027 call).

Is NVIDIA's Rubin Ultra delayed?

NVIDIA has not changed its second-half 2027 date for Rubin Ultra. In July 2026 SemiAnalysis reported that the Kyber rack for Rubin Ultra had slipped to 2028; NVIDIA responded "Our roadmap is intact" (CNBC; Tom's Hardware).

Why is NVIDIA's gross margin going down?

NVIDIA's gross margin is going down mainly because memory prices have risen more than the company expected. It guided gross margin to 74.0% in Q3 FY2027 and 71-72% in Q4, before 72-73% in FY2028 as agreed price increases take effect (Q2 FY2027 call).

Sources

1. NVIDIA, Q2 FY2027 earnings press release (Form 8-K, Exhibit 99.1, 26 Aug 2026)
2. NVIDIA, CFO commentary on Q2 FY2027 results (Form 8-K, Exhibit 99.2, 26 Aug 2026)
3. NVIDIA, Form 10-Q for the quarter ended 26 Jul 2026
4. NVIDIA, Form 8-K: agreement to acquire Hugging Face (2 Sep 2026)
5. NVIDIA, press release: NVIDIA guarantees SB Energy's PORTS-Pike campus (Form 8-K, Exhibit 99.1, 17 Aug 2026)
6. NVIDIA Investor Relations: Q2 FY2027 financial results
7. NVIDIA Q2 FY2027 earnings call transcript (26 Aug 2026)
8. NVIDIA Investor Relations: conference call for Q3 FY2027 results (17 Nov 2026)
9. NVIDIA Newsroom: NVIDIA introduces DSX POD AI factories (GTC Berlin, 21 Oct 2026)
10. NVIDIA blog: GTC 2025 keynote live updates (roadmap: Vera Rubin 2H 2026, Rubin Ultra 2H 2027)
11. CNBC: Nvidia's Kyber rack system delayed to 2028, per SemiAnalysis (6 Jul 2026)
12. Tom's Hardware: Nvidia says 'our roadmap is intact' (Jul 2026)
13. Yahoo Finance: NVDA price history and 30 Oct 2026 close
14. QSILabs NVIDIA discounted cash flow model (internal, valuation date 30 Oct 2026)

Disclaimer

This article is for information only and is not investment advice. Forecasts and fair values are QSILabs estimates and may be wrong. Company names and trademarks belong to their owners and are used solely for identification. Such use does not imply any affiliation with, sponsorship by, or endorsement of, QSI Labs. For more of our work, see the QSI Labs research hub.