

When compute became a receivable

Nvidia's second quarter was the largest in the company's history and the weakest cash quarter of the AI boom so far. Revenue rose 18% sequentially while operating cash flow fell 52%, because the money is now sitting in accounts receivable instead of the bank. At the same time, memory prices have turned the cost curve for frontier compute upward for the first time in this cycle, and Nvidia has pre-committed \$279 billion to hold its place in line. The AI buildout has quietly shifted from being funded out of cash flow to being funded out of balance sheets — and Nvidia's is now one of them.

AUTHOR

AT

Alicante Tech Vibes **Research Labs**

~8 MIN READ

·

ALL FIGURES SOURCED & DATED

\$96.2B Q2 FY27 REVENUE, +106% YOY · QUARTER ENDED 26 JUL 2026	~70% GUIDED FY2028 REVENUE GROWTH VS ~45% PRIOR CONSENSUS	60 days DAYS SALES OUTSTANDING, UP FROM 45 A QUARTER EARLIER
\$24.1B OPERATING CASH FLOW, DOWN 52% FROM \$50.3B IN Q1	\$279B SUPPLY & PURCHASE COMMITMENTS, FROM \$119B IN APRIL	71–72% GUIDED TROUGH GROSS MARGIN IN Q4 FY27, DOWN FROM 75.0%

I. The biggest quarter, and the thinnest cash

On 26 August 2026 Nvidia reported revenue of \$96.2 billion for the quarter ended 26 July 2026, up 106% year over year, with data centre revenue of \$89.0 billion and a 75.0% gross margin. Guidance for the current quarter is \$108.0 billion. Then, on the call, the company did something it had never done before: it guided a full year ahead, telling investors to expect roughly 70% revenue growth in fiscal 2028 — an implied \$690–700 billion against a sell-side consensus nearer \$570 billion. CFO Colette Kress framed the number as "constrained by supply", not demand.

The line that matters, though, is further down the cash flow statement. Nvidia's own CFO commentary puts operating cash flow at \$24.1 billion for the quarter. One quarter earlier, on \$81.6 billion of revenue, it was \$50.3 billion. **Nvidia booked 18% more revenue and collected less than half the cash.** Accounts receivable went from \$40.7 billion in April to \$63.1 billion in July; inventory went from \$25.8 billion to \$31.6 billion. Days sales outstanding — receivables divided by revenue, annualised — moved from about 45 days to 60. That is not a rounding difference. It is a change in what a sale is.

II. Two inversions arriving at once

The first inversion is financial. Nvidia is now extending working capital to its buyers. Kress attributed the 60-day DSO to extended terms for large investment-grade customers with shipments spanning multiple quarters, and reporting on the filing indicates terms running from 90 days out to a full year for some accounts, with five customers holding roughly 70% of the receivable balance. Separately, Nvidia disclosed financing platforms built with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR that are expected to raise more than \$500 billion, plus take-or-pay commitments on facility capacity to help neocloud operators reach project finance.

The second inversion is physical, and less discussed. Nvidia's gross margin is going down — 75.0% in Q2, 74.0% guided for Q3, a trough of 71–72% guided for the January quarter, settling at 72–73% next year once price increases land. The cause is memory. Kress described "extreme pricing conditions"; TrendForce had already forecast on 9 July 2026 that server DRAM contract prices would rise 13–18% quarter over quarter in 3Q26, with RDIMM bit supply growing only 15–20% year over year against faster server growth, and with long-term agreements shielding some hyperscalers so that everyone else absorbs the increase. On 2

June 2026 the same firm projected HBM contract prices rising by multiples in 2027.

For fifteen years the working assumption in every AI business case has been that the cost of a unit of compute falls. In this cycle the binding constraint has moved from logic to memory, and memory is a commodity market in shortage.

The cost of a frontier training or inference rack is rising, and Nvidia has chosen to eat roughly three points of margin before passing it on.

FIG. 1: NVIDIA TOTAL SUPPLY, MANUFACTURING & CAPACITY COMMITMENTS



BARS SCALED TO MAX (\$279B = 100%). SOURCE: NVIDIA 10-Q FILINGS AND Q2 FY27 CFO COMMENTARY, SEC, 27 AUG 2025 / 27 MAY 2026 / 26 AUG 2026.

FIG. 2: CASH CONVERSION — OPERATING CASH FLOW AS A SHARE OF REVENUE



BARS SCALED TO MAX (61.7% = 100%). COMPUTED FROM NVIDIA SEC FILINGS: \$50.3B / \$81.6B AND \$24.1B / \$96.2B.

III. What 60 days does and does not prove

The honest read is that a single quarter of stretched receivables is not evidence of anything fraudulent, and the loudest version of the bear case overreaches. Nvidia has faced this argument before — when DSO hit 53 days in late fiscal 2026, comparisons to Lucent-era vendor financing circulated, and the company responded that its DSO sat near long-run averages and that it does not use special purpose entities. The strongest counter-argument is comparative: suppliers selling into the same cloud customers — Arista, Celestica, Vertiv — routinely run 60–70 day DSO. On that benchmark Nvidia has moved from unusually fast collection to merely normal collection.

What the quarter does establish is directional and it is not ambiguous. Receivables absorbed roughly \$22 billion of cash in three months. Concentration rose rather than fell: five names now carry about 70% of the balance, against 56% a year earlier. And the terms are not incidental — Kress addressed the criticism head-on rather than denying the structure, saying the company knows some will call this circular financing and characterising the

exposure as low-risk, high-reward. That is an argument about credit quality, not an argument that credit is not being extended.

AI has reached its inflection point. Its tokens are productive and profitable. Now, compute is revenue.

JENSEN HUANG, NVIDIA CEO, IN THE Q2 FY2027 RESULTS RELEASE, 26 AUGUST 2026. IN THE SAME QUARTER, \$22B OF THAT REVENUE CONVERTED INTO RECEIVABLES RATHER THAN CASH.

IV. Who actually funds the next \$700 billion

Work the arithmetic forward with stated assumptions. Fiscal 2027 lands near \$406 billion if the \$108 billion Q3 guide holds and Q4 comes in around \$120 billion. Seventy per cent growth on that base is roughly \$690 billion of fiscal 2028 revenue. Nvidia's customers must therefore find something on the order of \$690 billion for Nvidia alone, before power, land, buildings, networking outside Nvidia's stack, and the memory premium they are also paying. Set against hyperscaler capital expenditure already running near \$600 billion across Alphabet, Amazon, Meta and Microsoft in calendar 2026, this is not a marginal increment. It is a second buildout stacked on the first.

The four hyperscalers can largely fund that from operating cash flow and investment-grade debt, and they have been issuing heavily to do it. The other half of Nvidia's business cannot. Kress put the enterprise, neocloud and sovereign segment at about \$40 billion in the quarter, growing 138% year over year, and Huang described non-hyperscale customers as roughly half the business growing at 100% annually. Those buyers — neoclouds especially — are thinly capitalised relative to the assets they are buying, which is precisely why Nvidia is arranging \$500 billion of third-party financing and offering take-or-pay backstops. **The fastest-growing half of the demand base is the half that needs someone else's balance sheet to buy.**

V. The second-order effect: token prices may stop falling

Most enterprise AI business cases written in 2025 and 2026 assume the price of inference keeps collapsing, so that workloads uneconomic today become

economic on a two-year view. Two things in this quarter cut against that. First, Nvidia has told the market it will raise prices on upcoming systems to recover memory costs. Second, memory shortage is a supply-side problem with multi-year lead times; TrendForce's own supply arithmetic has RDIMM bits growing 15–20% against faster demand, and HBM tightening into 2027.

If the delivered cost per unit of frontier compute is flat to rising through 2027, the effect is not that AI adoption stops — it is that the marginal workload does not clear. Batch-scale, low-value-per-token applications get deferred; high-value-per-task agentic work, where a dollar of inference replaces many dollars of labour, keeps clearing easily. That is a composition shift, and it is the mechanism by which a compute shortage becomes visible in the real economy: not as an AI slowdown, but as a narrowing of which AI use cases are worth running. For anyone modelling AI-driven margin expansion in a forecast period, the assumption to test is the cost-per-token glide path, not the adoption rate.

VI. The precedent, and where it breaks

The comparison bears reach for is telecoms in 1999–2001, when equipment vendors financed carriers to buy their own gear and the receivables turned out to be the earnings. The structural rhyme is real: vendor-extended terms, concentrated counterparties, capacity commitments made years ahead of the revenue they serve. The differences are also real, and this is a judgment rather than a fact. Nvidia's largest counterparties are among the most cash-generative companies in the world, not leveraged start-up carriers; the assets being bought are producing revenue today rather than lighting dark fibre for future demand; and Nvidia is funding commitments from a business generating tens of billions in quarterly profit, having still returned \$26 billion to shareholders in the quarter. The exposure that resembles 2000 is not the hyperscaler book. It is the neocloud and sovereign tail, and that tail is growing fastest.

VII. How it plays out: three scenarios

Timeframes run to end-2027. Probabilities are analytical judgment, not market-implied, and sum to 100%.

1 · Baseline: normalised, not resolved ~55% THROUGH 2027 · STEADY

Nvidia delivers \$600–700 billion in fiscal 2028. DSO settles in a 55–65 day band rather than reverting to 45, because extended terms are now a structural feature of selling to capacity-constrained buyers. Gross margin troughs near 72% and recovers as price increases land. Watch for: DSO stabilising rather than climbing in the November and February prints; free cash flow recovering above \$40 billion a quarter by mid-2027; no neocloud payment default reaching disclosure.

2 · Acceleration: the financing closes ~20% 2027 · FAST

The \$500 billion asset-manager platforms fund at scale, converting Nvidia's own receivable risk into third-party project debt, and agentic enterprise workloads absorb capacity as fast as it is built. Fiscal 2028 beats the 70% guide. Watch for: announced first closes and drawdowns on the Apollo/Blackstone/KKR vehicles; Nvidia's receivable growth decelerating below revenue growth; hyperscaler 2027 capex guides above \$700 billion in aggregate at the January and February earnings rounds.

3 · Stall: one counterparty misses ~25% 2027 · ABRUPT

A levered neocloud or a sovereign programme cannot refinance, a receivable is written down or restructured, and credit spreads on AI-linked paper widen. Orders are deferred rather than cancelled, but the \$279 billion of commitments becomes an obligation against softening demand and margin falls below the guided trough. Watch for: any allowance for credit losses appearing in Nvidia's filings; secondary-market pricing on GPU-collateralised debt; a downward revision to the FY2028 growth frame — the guide Nvidia had never before given is also the one it can most visibly miss.

WHAT WOULD FALSIFY THIS THESIS

The thesis is that the AI buildout has shifted from cash-funded to credit-funded and that the compute cost curve has turned up. Both are testable within two quarters. If Nvidia's Q3 FY27 report shows DSO back near 45–50 days and operating cash flow above \$45 billion on \$108 billion of revenue, then this quarter was shipment timing, not a financing regime change, and the credit half of the argument fails. If gross margin holds at or above 74% into the January quarter instead of falling to the guided 71–72%, the memory-cost half fails too. The strongest standing counter-case is the comparables argument: at 60 days Nvidia collects no slower than Arista, Celestica or Vertiv selling to the same customers, and the company has explicitly denied using off-balance-sheet vehicles. A reader who weights that argument heavily should read this brief as a description of rising credit intensity rather than a warning of impairment.

WHAT TO WATCH, IN ORDER

1. **Nvidia Q3 FY2027 results, around late November 2026.** The single decisive datapoint: DSO, operating cash flow, and whether receivable concentration above 70% keeps rising. Check the 10-Q, not the press release.
2. **TrendForce 4Q26 server DRAM contract price prints, October 2026.** If the 13–18% quarterly pace continues, Nvidia's guided 71–72% margin trough is optimistic rather than conservative.
3. **Hyperscaler Q3 2026 earnings, late October 2026.** Capex guides and the mix of cash versus new debt issuance at Alphabet, Amazon, Meta and Microsoft show whether the top half of the demand base is still self-funding.
4. **First closes on the \$500 billion financing platforms, late 2026 into 2027.** Announced drawdowns with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs or KKR would move risk off Nvidia's receivable line; silence would mean it stays there.
5. **Any change to the zero-China assumption, from the Q3 FY27 guidance onward.** Nvidia is currently modelling no data centre compute revenue from China at all; a policy shift in either direction is a large, discrete swing factor sitting outside every number above.

SOURCES

1. Nvidia, "NVIDIA Announces Financial Results for Second Quarter Fiscal 2027," 26 Aug 2026 · globenewswire.com
2. Nvidia, CFO Commentary, Q2 FY2027, filed with the SEC, 26 Aug 2026 · sec.gov
3. Nvidia, Form 10-Q for the quarter ended 26 Apr 2026 · sec.gov
4. Nvidia, Form 10-Q for the quarter ended 27 Jul 2025 · nvidia.com
5. Fortune, "Nvidia unleashes 70% growth bombshell and defends against 'circular financing' doomsayers," 26 Aug 2026 · fortune.com
6. Yahoo Finance, "NVIDIA Q2 Earnings Call Highlights," 26 Aug 2026 · finance.yahoo.com
7. FXStreet, "60 days to pay: Nvidia is financing its own demand," 26 Aug 2026 · fxstreet.com
8. FXStreet, "Nvidia earnings: AI boom meets margin test," 27 Aug 2026 · fxstreet.com
9. TrendForce, "Long-Term Agreements Cap Price Increases; Server DRAM Contract Prices Expected to Rise 13-18% QoQ in 3Q26," 9 Jul 2026 · trendforce.com
10. TrendForce, "Tight DRAM Supply Gives Suppliers Greater Pricing Power in HBM," 2 Jun 2026 · trendforce.com
11. Barchart, "Why Is Everyone Worried About Nvidia's Days Sales Outstanding?" 2026 · barchart.com
12. CNBC, "Nvidia's 70% growth forecast puts it on track to become tech's No. 2 company by revenue," 26 Aug 2026 · cnbc.com

METHOD NOTE: FIGURES ARE ATTRIBUTED AND DATED INLINE; WHERE INDEPENDENT SOURCES DISAGREE, THE DISAGREEMENT IS REPORTED RATHER THAN AVERAGED. REPORTED SHARE-PRICE REACTION DIFFERS BY SOURCE AND BY TIMING — THE STOCK WAS QUOTED LOWER ON THE INITIAL PRINT AND HIGHER IN EXTENDED TRADING AFTER THE FISCAL 2028 GUIDE WAS GIVEN; NO PRICE MOVE IS RELIED ON IN THIS ANALYSIS. FISCAL 2028 REVENUE OF ROUGHLY \$690 BILLION IS AN ESTIMATE DERIVED FROM THE GUIDED ~70% GROWTH APPLIED TO AN ASSUMED FISCAL 2027 BASE NEAR \$406 BILLION, ITSELF ASSUMING THE \$108 BILLION Q3 GUIDE HOLDS AND Q4 LANDS NEAR \$120 BILLION.

Built by the [Alicante Tech Vibes](#) community, a tope.

2026-08-27