

# Nvidia earnings AI boom meets margin test

Saxo Bank · 2026-08-27 · English original

---

## Key points:

Another beat-and-raise quarter: Revenue beat expectations by 4.2%, Q3 guidance came around 4% above consensus and management now expects approximately 70% growth in fiscal 2028 vs. current expectations of 45%, confirming that demand remains stronger than available supply.

Two company headwinds—and one macro hurdle: Memory costs are set to pull gross margin from 75% towards 71–72%, while free cash flow fell 56% quarter-on-quarter. Higher long-term bond yields add a separate valuation headwind for Nvidia and other long-duration AI stocks.

Positive for the AI ecosystem, but increasingly selective: The strongest read-through is for memory, networking, packaging and power infrastructure. Hyperscalers face a more mixed outcome as stronger AI demand comes with higher capex, weaker free cash flow and greater pressure to prove monetisation.

Nvidia delivered another beat-and-raise quarter, but the biggest positive was not the quarterly beat itself. It was management effectively telling investors that AI demand remains supply-constrained even at this scale—and guiding to approximately 70% revenue growth in fiscal 2028 vs. current expectations of 45%.

That is a powerful counter to the narrative that the AI capex cycle is already peaking. But at a market value above \$5 trillion, strong demand alone is not enough. Margins, cash conversion and the returns generated by Nvidia's customers will increasingly determine the next phase of the story.

Revenue gained 106% from a year earlier, while adjusted operating income increased 124%. Hyperscaler revenue also more than doubled, while gross margin was broadly in line with expectations at 75%.

The clearer disappointment was free cash flow. At \$21.3 billion, it was reportedly well below forecasts and down sharply from the previous quarter, although it remained 59% higher than a year ago.

## Demand was stronger—and broader—than expected

Data Centre revenue reached \$89 billion, representing more than 92% of Nvidia's total revenue.

Within Data Centre, hyperscale revenue rose 13% sequentially to \$48.7 billion. Revenue from AI clouds, industrial companies and enterprises increased a faster 25% to \$40.3 billion.

That distinction matters. Nvidia is no longer relying exclusively on a small group of US technology giants. AI-native companies, sovereign projects, industrial customers and enterprises are becoming increasingly important sources of demand.

It is still diversification within AI infrastructure rather than genuine business diversification. But it reduces the risk that Nvidia's growth depends entirely on the capex decisions of four or five hyperscalers.

## **The biggest bullish surprise came during the call**

Management expects revenue to grow by approximately 70% in fiscal 2028, which largely covers calendar 2027. Market expectations had reportedly been closer to 45%.

Nvidia also indicated that it can currently satisfy only around 70% of demand. In other words, the immediate constraint remains supply—not a shortage of customers.

That guidance was the main reason Nvidia shares moved more than 4% higher during extended trading after a more restrained initial reaction to the results.

It also changes the debate around peak growth. A \$108 billion quarterly revenue run rate no longer looks like the end of the cycle if Nvidia can deliver anything close to 70% growth next fiscal year.

## **Rubin execution risk has fallen, but not disappeared**

Nvidia commenced production shipments of Vera Rubin in early August, while inventory increased from \$25.8 billion to \$31.6 billion as the company prepared for the ramp.

This is reassuring because Nvidia is attempting another major architectural transition immediately after scaling Blackwell and Blackwell Ultra. So far, there is no sign of customers delaying spending while waiting for Rubin.

However, beginning production is not the same as proving a smooth, high-volume deployment. The next two quarters will test:

networking, power and cooling availability;

how quickly customers can install and utilise delivered systems.

The report reduces the risk of a product-transition slowdown, but the scale and complexity of Rubin mean execution remains important.

## **China is optionality, not part of the base case**

China contributed less than 1% of Nvidia's Data Centre revenue, while the company assumes no China Data Centre compute revenue in its Q3 forecast.

That means the \$108 billion guidance does not rely on regulatory relief. Any meaningful reopening would provide incremental upside.

The longer-term risk is that Nvidia's absence gives Chinese companies a protected market in which to develop domestic accelerators and software ecosystems. Even if restrictions eventually ease, some of that market could be difficult to recover.

For Asian markets, the near-term read-through is most supportive for Korea's memory suppliers and Taiwan's foundry, packaging and server supply chains. Nvidia's continued absence from China, meanwhile, strengthens the strategic case for domestic substitution.

## **Gross margins are the key earnings tension**

Nvidia delivered a 75% gross margin in Q2 but guided to 74% in Q3. Management subsequently indicated that margins could fall to around 71–72% in the January quarter before stabilising around 72–73% next fiscal year.

The biggest pressure is coming from expensive and scarce memory. Nvidia is planning price increases

on some upcoming systems, but there is a timing gap between higher component costs and the company's ability to recover them from customers.

The bullish interpretation is that margins are falling because AI demand is overwhelming the memory supply chain. Nvidia continues to have strong pricing power, while lower margins are more than offset by revenue growth.

The less comfortable interpretation is that Nvidia's peak economics are being tested by rising system costs, greater complexity and more expensive inputs.

On quarterly revenue above \$100 billion, every percentage point of gross margin represents more than \$1 billion of gross profit.

## **Cash conversion was another weak spot**

The gap between earnings and cash flow was unusually wide. Adjusted net income reached \$54 billion, but free cash flow was only \$21.3 billion, implying cash conversion of less than 40%.

This largely reflected a heavy working-capital build:

Accounts receivable absorbed \$22.3 billion of cash.

Days sales outstanding rose from 45 to 60 days.

Cash taxes also weighed on operating cash flow.

Nvidia attributed the increase in receivables to extended payment terms on large, multi-quarter agreements with investment-grade customers. The inventory build reflects preparations for Rubin.

Some of this pressure may reverse as payments arrive and Rubin inventory is shipped, so the FCF decline does not necessarily signal weaker demand. However, it makes the quarter's earnings quality less clean and deserves monitoring as Nvidia extends customer terms and takes on larger financing commitments. Still, after such a strong run in the stock, cash conversion is now another metric investors cannot ignore.

Nvidia also returned around \$26 billion to shareholders during the quarter—more than the free cash flow it generated. The company can comfortably absorb this using its balance sheet and first-half cash generation, but the pace would be harder to sustain if weaker cash conversion persists.

GAAP net income of \$59.7 billion also included \$7.8 billion of gains on equity securities. Adjusted net income of \$54 billion is therefore the cleaner measure of underlying performance.

## **Nvidia is becoming an AI financier**

supply and capacity commitments of \$279 billion, up from \$119 billion last quarter, primarily related to memory;

approximately \$366 billion of total future commitments;

another \$56 billion of AI-cloud and third-party lease commitments;

maximum gross guarantees of \$108.5 billion, including as much as \$105 billion connected with the OpenAI and SB Energy data-centre development in Ohio;

partnerships intended to mobilise more than \$500 billion of third-party capital for AI infrastructure.

These arrangements can help Nvidia secure scarce supply, accelerate customer deployments and expand its addressable market.

But they also make the company's risk profile more complex. Investors increasingly need to consider customer credit quality, leases, guarantees, revenue-sharing agreements and Nvidia's equity investments—not just GPU shipments.

This does not automatically make the revenue circular. It does mean Nvidia is increasingly helping to create and finance the ecosystem into which it sells.

The weaker FCF result makes this development more relevant. Nvidia still has substantial financial flexibility, but expanding commitments alongside rising working-capital requirements could gradually reduce that cushion.

## **What Nvidia's results mean for the AI ecosystem**

**Memory:** Nvidia's margin guidance confirms that HBM and server-memory supply remains tight today, supporting pricing power for SK Hynix, Micron and other suppliers. But supply is catching up. The next stage of the memory trade depends on whether AI demand and rising memory content per accelerator can absorb additional capacity—not simply on shortages pushing prices higher.

**Hyperscalers and custom chips:** Nvidia's outlook gives another vote of confidence to hyperscaler capex. But the more expensive Nvidia GPUs and memory become, the stronger the incentive for companies to develop custom silicon. This does not undermine Google TPUs, Amazon Trainium or the Broadcom and Marvell ASIC opportunity. Nvidia can remain dominant in frontier AI while custom processors take a larger share of specific inference workloads.

**Networking, optics, power and cooling:** Rubin, Spectrum-6 networking and increasingly large AI factories suggest the bottleneck is moving beyond GPUs into memory bandwidth, data movement, electricity, cooling and data-centre capacity. These remain attractive second-order areas of the infrastructure buildout.

**AI clouds:** The 25% sequential growth in Nvidia's AI-cloud, industrial and enterprise segment supports selected providers such as CoreWeave and Nebius. But high leverage, customer concentration and dependence on external financing make this one of the riskier parts of the ecosystem.

**Foundry and packaging:** The Rubin ramp supports TSMC and advanced-packaging suppliers. Demand visibility remains strong, but capacity constraints, customer concentration and geopolitics cannot be ignored.

Overall, Nvidia's results validate the AI infrastructure trade—but they are not a signal to indiscriminately chase everything labelled AI. Rising component costs, higher yields and the need to prove returns will increasingly separate companies with pricing power and visible demand from the rest.

## **Bottom line**

Nvidia's report removes one major near-term concern around demand. The AI infrastructure cycle is not yet rolling over, Rubin is entering production and supply remains the binding constraint.

The next test is whether Nvidia can preserve enough of its extraordinary margins and restore stronger cash conversion as the business scales. Meanwhile, hyperscalers must show that rapidly rising AI expenditure can generate adequate returns—particularly when long-term yields offer investors an increasingly credible alternative.

For Nvidia and the wider AI ecosystem, the question is no longer whether growth exists. It is where the economics of that growth ultimately settle.

---

Source: <https://www.home.saxo/content/articles/equities/nvidia-earnings-ai-boom-meets-margin-test-27082026>