NVIDIA · Q2 FISCAL 2027

NVIDIA's AI boom is accelerating, but sustaining it now requires far bigger commitments to customers and suppliers

Demand and profitability remain exceptional. Cash collection and NVIDIA's expanding role in financing the AI ecosystem make the next stage more complicated.

Analysis based on NVIDIA's Earnings Release (8-K Exhibit 99.1), Exhibit 99.2, Form 10-Q and SEC XBRL companyfacts · Notes at the end

NVIDIA's quarter ended July 26, 2026, showed that the AI buildout is still accelerating, not merely coasting on last year's boom. Revenue rose 106% to $96.2 billion. Yet the cleanest reading is not the bottom line: net profit received a $7.8 billion lift from gains on equity securities. The more important tension lies elsewhere. Customers are buying at extraordinary speed, while NVIDIA is extending more generous payment terms, securing far more supply and tying its own capital to the ecosystem.[1]

THE BASICS

NVIDIA designs the processors, networking equipment and software used to build and run accelerated computing systems. Its largest market is now the data center, where cloud providers and other companies buy complete AI infrastructure. It also serves workstation, gaming and edge-computing markets. NVIDIA relies heavily on outside manufacturers and must reserve production capacity long before finished systems reach customers.

THE PROFIT GLOSS DOES NOT ERASE THE OPERATING RESULT

Investment gains made an already remarkable quarter look even better. Reported net income was $59.7 billion, but that included $7.8 billion of gains on equity securities. Those gains reflect changes in the value of NVIDIA's investments rather than the sale of more computing systems, so they should not be read as evidence that the core business suddenly became more productive.[1]

The operating business did not need that gloss. Profit from NVIDIA's day-to-day operations climbed from $28.4 billion to $63.7 billion. That increase came before the investment gains and shows that the underlying surge in demand translated into much more profit, not just more revenue.[1]

NVIDIA also kept more of each sales dollar after the cost of producing its systems. Gross profit was about 75 cents per dollar of revenue, compared with roughly 72 cents a year earlier. Management attributes the improvement to a more profitable mix of Blackwell Ultra infrastructure.[1]

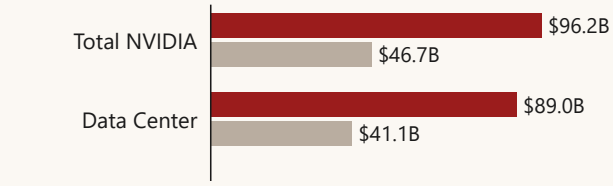
AI INFRASTRUCTURE DID ALMOST ALL THE WORK

The quarter was overwhelmingly a data-center story. Data Center revenue reached $89.0 billion, up from $41.1 billion a year earlier, and represented about 93% of total company revenue. Management says the rollout of Blackwell Ultra infrastructure drove the increase. This is extraordinary momentum, but it also leaves NVIDIA dependent on one investment cycle: the global construction of AI computing capacity.[2]

China was not an important contributor to that result. Shipments of Hopper data-center products to China were less than 1% of Data Center revenue. NVIDIA's next-quarter revenue outlook of $108 billion assumes no Data Center compute revenue from China, making demand elsewhere the relevant test of whether the expansion can continue.[2]

How completely is AI infrastructure driving the quarter?

Quarterly revenue, USD billions



Data Center is NVIDIA's market-platform view, not a separately reportable accounting segment. Source: SEC XBRL companyfacts and NVIDIA Exhibit 99.2

THE CUSTOMER BASE IS BROADENING WITHIN AI, NOT BEYOND IT

Growth within the data-center business is spreading beyond the largest cloud operators. Revenue from AI clouds, industrial companies and enterprise customers rose 25% from the previous quarter to $40.3 billion. Hyperscale revenue also grew strongly, but more slowly, at 13%. That widening customer mix matters because it suggests AI infrastructure spending is no longer confined to a few giant cloud platforms.[2]

Concentration has not disappeared. One direct customer accounted for 16% of quarterly revenue. Direct customers can include distributors and cloud providers rather than the organizations ultimately using the equipment, but the figure still shows how much business can flow through a small number of counterparties. Rapid expansion and customer concentration are coexisting, not canceling each other out.[3]

THE NEXT PRODUCT CYCLE IS ALREADY DEMANDING RESOURCES

NVIDIA is building inventory before the next generation arrives. Inventory increased to $31.6 billion from $25.8 billion in the previous quarter as the company prepared to introduce Vera Rubin. Management says Rubin is already in full production, but NVIDIA also acknowledges supply constraints as it prepares to ship Blackwell and Rubin systems at the same time.[2]

The larger commitment is to future supply. Contracted supply and capacity commitments jumped to $279 billion from $119 billion in one quarter. Including cloud agreements, leases, investments and planned capital spending, NVIDIA disclosed $366 billion of future commitments. These are not all immediate expenses, but they show how much capacity the company is reserving before the eventual level and timing of demand are fully known.[2]

SALES RACED AHEAD OF CASH COLLECTION

The quarter produced much less cash than the previous one even as sales continued to rise. Cash generated by operations fell from $50.3 billion to $24.1 billion. The cash left after investment followed the same pattern, as the chart shows. The issue was not a collapse in reported demand. More money was tied up in customer bills, inventory and tax payments.[2]

Large customers are also taking longer to pay. NVIDIA's average collection period increased to 60 days from 45 days after it granted extended payment terms on large, multi-quarter agreements with certain investment-grade customers. Management says these arrangements will continue to affect when operating cash arrives. That makes cash collection an important check on the quality of future growth.[2]

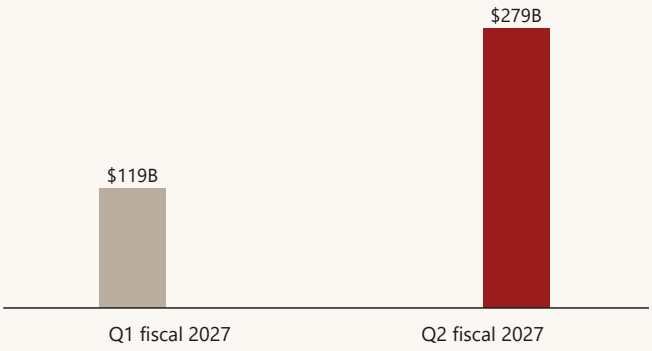
> *"Vera Rubin, now in full production, was built to power exactly this moment."*
>
> Jensen Huang, founder and CEO · Earnings Release

ONE KEY IDEA

Reserving supply early can protect NVIDIA from shortages and support faster growth. It also raises the cost of being wrong if demand, product timing or customer financing changes.

How much future supply has NVIDIA locked in?

Supply and capacity commitments, USD billions



These are contractual commitments extending over future periods, not spending incurred during the quarter. Source: NVIDIA Exhibit 99.2

NVIDIA IS BECOMING A BACKER OF THE ECOSYSTEM IT SUPPLIES

NVIDIA's relationship with emerging AI clouds increasingly goes beyond selling equipment. At quarter end, it held $99 billion of equity investments and had $36 billion of commitments with selected AI cloud partners. Under some agreements, NVIDIA sells infrastructure upfront, commits to use capacity itself and may later receive a share of the cloud provider's revenue. That can help new customers expand, but it also links NVIDIA more closely to their success.[3]

The company moved further in that direction after the quarter. It agreed to provide conditional support for land, power and building infrastructure at an SB Energy campus, with guarantee obligations capped at $105 billion. The guarantee is not a current cash payment: it becomes effective in phases as conditions are met, with the first phase expected in fiscal 2029. Even so, its scale shows how far NVIDIA is prepared to go to unlock future data-center construction.[3]

A SUPPLIER WITH SKIN IN THE GAME

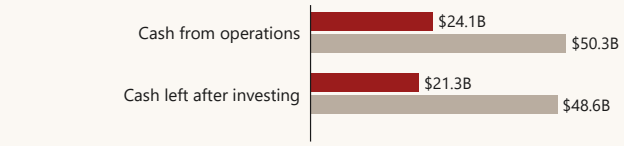
NVIDIA is no longer only waiting for customers to finance data centers and place equipment orders. It is investing in providers, reserving their capacity and, in selected cases, supporting the infrastructure beneath them.

THE NEXT QUARTER TESTS GROWTH AND DISCIPLINE TOGETHER

Management expects third-quarter revenue of $108 billion without assuming any China Data Center compute sales. At the same time, it expects gross profit to ease to about 74 cents per sales dollar from roughly 75 cents this quarter. The outlook therefore points to continued expansion, but not to endlessly improving economics. The question is whether NVIDIA can deliver the next product cycle while keeping supply, customer financing and cash collection under control.[2]

Why does the cash story look weaker than sales?

Quarterly cash generation, USD billions



"Cash left after investing" is NVIDIA's free cash flow measure, which subtracts purchases and principal payments for property, equipment and intangible assets. Source: NVIDIA Earnings Release and Exhibit 99.2

WHAT IT ALL MEANS · EARNINGS EXPLAINED INTERPRETATION

The quarter was better than the cash figures suggest and less clean than the net-profit headline implies: demand is exceptional, but NVIDIA is using more of its own capital and contractual capacity to keep the AI buildout moving.

THE ENGINE

Blackwell-led data-center demand remains the overwhelming source of growth, with customers outside the largest cloud platforms contributing more.

THE STRAIN

Customers paid more slowly while inventory and supply commitments rose ahead of the Rubin cycle.

THE PROFIT QUALITY

Investment gains flattered net income, but the operating result and stronger profit on each sales dollar show that the core improvement is real.

THE STRATEGIC SHIFT

NVIDIA is becoming an investor, capacity buyer and infrastructure backer for some of the same companies that purchase its systems.

Three questions for next quarter

1 Can NVIDIA reach its next-quarter revenue plan without China Data Center compute sales?

That would show whether demand in the rest of the world remains deep enough to sustain the current pace.

2 Do customers start paying faster?

A shorter collection period would help confirm that reported growth is turning into cash rather than remaining tied up in long contracts.

3 Does the Rubin rollout justify the larger inventory and supply commitments?

Successful shipments would validate NVIDIA's early capacity reservations; delays or new constraints would make those commitments more burdensome.

1 SEC XBRL companyfacts and NVIDIA's Form 10-Q provide the consolidated GAAP figures for the quarter ended July 26, 2026. NVIDIA Exhibit 99.2 reports the equity-security gain and management's explanation of gross margin. Dollar amounts in the article are rounded.

2 NVIDIA Exhibit 99.2 provides market-platform revenue, customer payment timing, inventory, cash flow and contractual commitments. The Earnings Release provides third-quarter guidance. NVIDIA recast prior Hyperscale and AI Clouds, Industrial, & Enterprise revenue after reclassifying one customer.

3 NVIDIA's Form 10-Q provides customer concentration, equity investments, AI cloud arrangements and related risks. It also describes the SB Energy guarantee signed after quarter end. The guarantee is conditional and phased; its cap is not a current expense or cash payment.

4 Calculations use reported pack values. Revenue growth compares the current quarter with the year-earlier quarter. Gross margin is gross profit divided by revenue. Data Center's share is market-platform revenue divided by consolidated revenue.