


A close-up photograph of a blue printed circuit board (PCB) with various electronic components. A black microscope is positioned over the board, with its lens focused on a central chip. The background is dark and out of focus.

Nvidia Delivers Blowout Quarter, Guides Higher on AI Demand

Equity Insights
November 2025

LIGHTHOUSE
CANTON

Nvidia Delivers Blowout Quarter, Guides Higher on AI Demand

Executive Summary & Strategic Conclusion

NVIDIA delivered a standout performance in Q3 FY26, reporting \$57.0 billion in revenue, up 62% year-over-year, and GAAP EPS of \$1.30, both exceeding consensus expectations. The Data Center segment alone contributed \$51.2 billion, growing 66% year-over-year, and remains the primary engine of expansion as AI infrastructure demand continues to scale rapidly. Management's forward guidance — \$65 billion in revenue for Q4 FY26 — signals continued confidence in the durability of this demand cycle.

Source: TradingView

While the numbers are strong, it is the underlying structural dynamics that command investor attention. There are five critical themes emerging from this quarter that point to NVIDIA's evolving role as the foundational platform for the AI economy:

- **Unprecedented Demand Visibility:** With over \$500 billion in purchase commitments through 2026 and approximately \$350 billion yet to be fulfilled, NVIDIA's revenue base is being recalibrated upward. Supply chain metrics - inventory up 32% and supply commitments up 63% quarter-over-quarter, support this ramp.
- **Smooth Architectural Transitions:** The successful migration from Hopper to Blackwell, and now from GB200 to GB300, with GB300 already contributing two-thirds of Blackwell revenue validates NVIDIA's annual cadence. The early success of Rubin silicon from TSMC suggests no disruption ahead for the next-generation roadmap.
- **Software-Centric Longevity:** NVIDIA's CUDA platform extends the utility of its hardware across generations. Empirical evidence - continued full utilization of six-year-old A100 chips and \$2 billion in Hopper sales this quarter - undermines concerns around asset depreciation cycles in the cloud infrastructure space.
- **Networking Emerges as a Core Growth Driver:** With \$8.2 billion in quarterly networking revenue, split evenly between InfiniBand and Spectrum-X, NVIDIA is now the world's largest networking vendor. This further consolidates its position as an end-to-end AI systems provider, not merely a chip supplier.
- **Sustained Competitive Moat:** Benchmark leadership across both training and inference workloads, with Blackwell and Hopper outperforming rival chips on performance-per-watt and TCO. This cements NVIDIA's multi-year lead. Notably, Blackwell's 10x performance advantage in inference addresses a prior area of concern and reinforces its dominance across AI compute modalities.

Collectively, these trends not only explain NVIDIA's current outperformance but also suggest a structural re-rating of its long-term earnings power. The company is building a highly resilient, platform-wide moat across compute, interconnect, and software - positioning it at the center of the generative AI economy for years to come.

Key Financial Metrics:

- Revenue: \$57.0B (vs consensus ~\$56B)
- GAAP EPS: \$1.30 (vs consensus ~\$1.25)
- Gross Margin: 73.4% GAAP, 73.6% Non-GAAP
- Guidance(Q4 FY26):
- Revenue: \$65B $\pm$ 2% (above consensus ~\$62B)
- Gross Margin: ~75% Non-GAAP

Segment Highlights:**Data Center:**

Revenue: \$51.2B, up 25% q/q, 66% y/y

Driven by hyperscaler and enterprise adoption of the Blackwell GB 300

Networking revenue (Ethernet/InfiniBand): \$8.2B, +162% y/y

Management cites a \$500B+ revenue opportunity through 2026 across AI compute platforms

Gaming:

Revenue: \$4.3B, +30% y/y

Stable performance with potential cyclicalities

Professional Visualization:

Revenue: \$760M, +56% y/y

Continued adoption in AI-enhanced design workflows

Automotive & Robotics:

Revenue: \$592M, +32% y/y

Still a small contributor, but seen as a strategic long-term bet

1. Unpacking the \$500 Billion Backlog: A New Revenue Baseline Emerges

Nvidia previously disclosed that it had secured purchase commitments totaling approximately \$500 billion for calendar years 2025 and 2026. Of this, the company has already fulfilled \$150 billion, leaving a remaining backlog of \$350 billion over the next 14 months. Importantly, management noted that new commitments are continually being added, suggesting the ultimate realized revenue could exceed this figure.

This dynamic sets the stage for FY27 revenue to surpass current Street estimates, which stand at approximately \$304 billion, potentially triggering upward revisions in both revenue forecasts and valuation estimates.

The company's inventory position rose 32% quarter-over-quarter, and supply commitments climbed 63%, reinforcing NVIDIA's preparedness for an aggressive ramp in delivery volumes over the coming quarters.

2. Seamless Transition Between Product Generations Validates Annual Cadence

Concerns around potential demand fatigue due to NVIDIA's annual product release cadence — especially given the high cost of newer chips — have been effectively put to rest.

Management highlighted that the next-generation chips offer superior efficiency in terms of tokens per watt and total cost of operation (TCO), which continues to justify customer migration. The transition from Hopper to Blackwell, and now from Blackwell GB200 to GB300, has progressed smoothly. GB300 now accounts for two-thirds of all Blackwell revenue recognized this quarter.

Furthermore, NVIDIA confirmed it has received test silicon for the upcoming Rubin platform from TSMC and expressed satisfaction with its performance — signaling a smooth ramp toward a H2 2026 launch. This is a marked improvement from the Blackwell development cycle, which faced issues such as overheating and the need for the chip's masking to be re-worked.

3. Extending Chip Lifespan Through Software Ecosystem Synergy

Nvidia emphasized that its unified CUDA software platform, which supports multiple hardware generations, significantly extends the useful life of its chips, reinforcing their TCO advantage for customers.

As evidence, the company pointed to its A100 GPUs, which remain at full utilization six years post-shipment, and noted \$2 billion in Hopper chip sales this quarter, underscoring the continued value of prior-generation products.

This longevity directly counters bearish narratives suggesting that cloud service providers are under-depreciating GPU assets, thereby overstating earnings. NVIDIA's commentary and data imply these assets retain meaningful utility well beyond typical depreciation cycles.

4. Networking: From Complementary to Core Business

Nvidia reported \$8.2 billion in networking revenue for the quarter, positioning it as the world's largest networking vendor. This figure was evenly split between InfiniBand (proprietary) and Spectrum-X Ethernet solutions, highlighting balanced momentum across both technologies.

Networking is becoming a core growth pillar, not merely an enabler. As AI model deployment scales, the demand for interconnected GPU clusters rises, boosting the need for high-bandwidth, low-latency networking solutions. This, in turn, reinforces adoption of CUDA, given its dependence on GPU-to-GPU communication efficiency.

NVIDIA is increasingly evolving into an end-to-end AI platform provider—offering compute, software, and connectivity—rather than just a chip supplier. While the overall networking market remains expansive enough to support multiple players, this development is likely to exert competitive pressure on vendors such as Arista Networks and other “blue-box” providers.

5. Competitive Position: Benchmarks Confirm Multi-Year Lead

Addressing competition, NVIDIA pointed to recent results from AI performance tests and benchmarks including MLPerf and Semi Analysis Inference MAX benchmark, where its GPUs achieved industry-best performance per watt and lowest TCO.

Crucially, both Blackwell and Hopper chips outperformed competitors, with Blackwell leading and Hopper coming in second—suggesting a multi-generational lead over rival architectures, including merchant GPUs from AMD and custom AI chips (ASICs/XPUs) from the likes of Broadcom.

Of particular note was Blackwell's 10x inference performance lead (over Hopper, which was 2nd best), which is strategically important as inference workloads are expected to dominate future AI deployments. The results decisively allay prior concerns about Nvidia's inference competitiveness, underscoring the company's advantage not just in training but also in real-world, production-scale AI applications.

LIGHTHOUSE CANTON

Singapore

16 Collyer Quay #11-02
Collyer Quay Centre
Singapore 049318

☎ +65 67130570

UAE

The Exchange Gate Village 11,
Unit 802 Dubai International
Financial Centre PO Box 507026
Dubai, UAE

☎ +971 45 861500

UK

24 Hanover Square,
London W1S 1JD

☎ +44 164 2843 487

India

Unit No-104A, Worldmark
2 Asset Delhi Aerocity,
New Delhi 110037

☎ +91 9650473961

Unit No 507/508, A Wing,
INS Tower, G Block, BKC,
Mumbai- 400051

☎ +91 9650473961

1st Floor, WeWork37,
Cunningham Cross Rd,
SRT Road Vasant Nagar,
Bengaluru-560001

☎ +91 9900096873

RK Swamy Centre, Hansa
Building, Door No:3, Thousand
Lights,
Chennai-600006

☎ +91 9650473961

Suite 502, Building 450, Central
Plaza, Genome Valley,
Shameerpet,
Hyderabad 500 078

☎ +91 9650473961

Unit No FF-10, FF Floor, Pragya
Accelerator, Block 15T GIFT CITY,
Gandhinagar
Gujrat-382355

☎ +91 9650473961

✉ info@lighthouse-canton.com

✉ service@lighthouse-canton.com

in Lighthouse Canton

DISCLAIMER

The contents of this document are confidential and are meant for the intended recipient only. If you are not the intended recipient, please delete all copies of this document and notify the sender immediately.

This document, provided as a general commentary, is for informational purposes only and is not to be construed as an offer to sell or solicit an offer to buy any financial instruments in any jurisdiction. This does not constitute any form of regulated financial advice, and your independent financial advisor should be consulted prior to taking any investment decision(s).

This document is based on information from Sources which are reliable but has not been independently verified by Lighthouse Canton Pte Ltd and its affiliate companies ("LC"). LC has taken the reasonable steps to verify the contents of this document and accept no liability for any loss arising from the use of any information contained herein. Please also note that past performances are not indicative of future performance.

Information contained herein are those of the author(s) and does not represent the views held by other parties. LC is also under no obligation to update you on any changes made to this document.

This document is prepared by Lighthouse Canton Pte Ltd and its affiliate company, Lighthouse Canton Capital (DIFC) Pte Ltd, which are regulated by Monetary Authority of Singapore ("MAS") and Dubai Financial Services Authority ("DFSA") respectively. MAS and DFSA has no responsibility for reviewing, verifying and approving the contents of this document and/or other associated documents. The contents of this document may not be reproduced or referenced, either in part or in full, without prior written permission from LC.

This document is confidential and is only intended for Accredited Investors and/or Professional Clients, as defined by MAS and DFSA.