

NVIDIA SAYS THE AI BOOM ISN'T OVER YET

NVDA +6.16 ▲
184.25

EPISODE
429



HOSTED BY
**DAMIEN
KLASSEN**

Disclaimer:

The information provided on this presentation is general in nature and does not constitute personal financial advice. The information has been prepared without taking into account your personal objectives, financial situation or needs. Because of this, you should consider the appropriateness of the information for your own objectives, financial situation and needs before acting on it.

Also, before you decide to invest in a financial product arranged by a representative of Nucleus Wealth Management Pty Ltd, ABN 54 614 386 266, corporate authorised representative of Nucleus Advice Pty Ltd AFSL 515796 (Nucleus Wealth or we or us), it is important that you read and consider the Product Disclosure Statement relating to the product before making any decision about whether to invest in it.

NVIDIA:

The Central Bank of AI?

Latest result, data-centre capex, profitability and the financing flywheel

Core question: when does the bottleneck shift from chips, power and money to actual end-user demand?

1. The result: extraordinary, and still supply constrained

Revenue

\$96.2bn

+106% YoY

Data Center

\$89.0bn

+117% YoY

Q3 guidance

\$108bn

±2%

Gross margin

75.0%

non-GAAP


Data Center split



Almost half of Data Center is now outside the traditional hyperscalers — and growing faster.

2. Data-centre capex: from chip scarcity to capital scarcity

NVIDIA estimate: top-five hyperscaler capex




But the spending is broadening beyond Big Tech:

- Neoclouds need debt, leases and backstops
- Power/site constraints push behind-the-meter solutions
- Compute density increases NVIDIA revenue per GW
- The boom is no longer just a cloud capex line item

The AI boom has moved from “can I get GPUs?” to “who can finance the next gigawatt?”

3. Physical capacity: power is the hard constraint




4. Profitability: who captures the revenue?

On a GB300 Cluster, Revenue Per Gigawatt Climbs 8× Up the AI Stack SpaceX Prices Between Neocloud Rental and Model-Layer Tokens

Annualized revenue per gigawatt by layer — \$B per GW per year. Multiples shown vs. neocloud 5-yr IaaS average (= 1.0×).

Sources: SemiAnalysis Tokenomics Model, SemiAnalysis Inference Simulator

sales@semianalysis.com



5. Data-centre profitability: EBITDA is not return on capital

CoreWeave Q2 2026

Revenue	\$2.575bn
Adjusted EBITDA	\$1.510bn
Adjusted EBITDA margin	59%
Adjusted operating margin	5%
Net loss margin	-24%

Hyperscaler cloud segments look far stronger because they wrap GPUs with storage, networking, databases, software and customer lock-in.

GPU rental can look fantastic before depreciation and interest.

Podcast question:

Is the boom earning high returns on capital — or only high cash gross margins before the capital bill arrives?

6. The tipping point: not GW, but profitable tokens

Useful capacity = tokens per second × acceptable latency × profitable price

Supply curve

More GW
+ more chips per rack
+ more tokens per MW
+ lower token cost

Demand curve


Agents
+ reasoning loops
+ tool calls
+ enterprise workflows
+ Jevons paradox

Watch list

- Rental prices fall below full cost
- Utilisation rolls over
- Residual GPU values collapse
- More generous vendor support needed
- Cloud revenue slows while capex remains high

7. The \$7tn question: who funds the infrastructure?

The Explosion of AI Debt Financing vs.
US Asset Backed Credit Markets (\$T, 2023-2029E)



Source visual: uploaded SemiAnalysis slide — Debt Issues?

8. NVIDIA financing: backstop today, revenue share tomorrow

Nvidia Backstop - Indicative Terms

Backstop Pricing



Year 1	\$/hr/GPU	\$3.68
Year 2	\$/hr/GPU	\$3.15
Year 3	\$/hr/GPU	\$2.62
Year 4	\$/hr/GPU	\$2.10
Year 5	\$/hr/GPU	\$1.57
Year 6	\$/hr/GPU	\$1.04
Average Backstop Price	\$/hr/GPU	\$2.36

Backstop Duration

Years 6

Nvidia Rev Share above Backstop

% ~40-60%

Mechanism

- Neocloud buys NVIDIA hardware
- NVIDIA backstops unsold capacity
- Bank lends against the cash-flow floor
- NVIDIA may capture 40–60% of upside in the SemiAnalysis illustration
- Economic form: NVIDIA writes a put option on compute demand

NVIDIA gets hardware margin first — then potentially a share of rental economics.

9. How big could the extra NVIDIA profit pool be?

Illustrative economics combining the uploaded pricing and backstop slides

1GW IaaS revenue ~\$12bn/year

Backstop average \$2.36/GPU-hour

Illustrative rental \$4.27/GPU-hour

Excess \$1.91/GPU-hour

NVIDIA share 40–60% of
excess

**Rough (non?) recurring revenue per
1GW: ~\$2.1bn–\$3.2bn/year**

**Six-year mechanical total:
~\$13bn–\$19bn**

Big caveat: this is an illustration, not disclosed NVIDIA guidance. Actual utilisation, contract floors, pricing, duration and revenue-share terms can differ materially.

10. NVIDIA's "second balance sheet"

Accounts receivable	\$63bn	Supply commitments	\$279bn
AI-cloud commitments	\$36bn	SB Energy / OpenAI guarantee	up to \$105bn
Equity investments	\$99bn	Third-party finance platform	>\$500bn
Future equity commitments	\$25bn	AI labs using balance sheet	~1/4 of next-year business

Do not add these numbers together — they are different types of exposure.

But together they show NVIDIA is no longer just selling into the AI ecosystem. It is financing, guaranteeing, investing in and coordinating it.

11. Conclusion: boom, bubble, or financed transition?

Bull case

- Demand still exceeds supply
- Cloud revenue growth is accelerating
- Agentic AI expands compute intensity
- NVIDIA captures more per GW
- Financing extends the runway

Bear case

- Supplier finance muddies demand signals
- Debt and leases build fragility
- GPU economics hit depreciation and interest
- Residual values could fall fast
- Backstops are put options

Balanced takeaway

We are not yet seeing clear evidence of excess AI capacity. But the industry now needs enormous financing support to maintain its growth rate. That does not prove the boom is ending; it means the consequences will be larger when it eventually does.

Final question: when does the bottleneck shift from chips, power and money to actual demand?

More from Nucleus Wealth



Find us on all major
podcast platforms:



Find out what your personalised
investment could look like in
10-20 minutes – commitment
free & at no cost.

Read insightful blog
articles:

nucleuswealth.com/blog

Nucleus Wealth on social
media:



Facebook.com/nucleuswealth



Linkedin.com/company/nucleuswealth



Instagram.com/nucleus_wealth



Twitter.com/nucleuswealth

*Have a guest or topic
suggestion for the show?
Leave a comment down below.*