Investor Presentation
Q2 FY26

September 2025

Certain matters in this presentation including, but not limited to, statements as to: the impact of H20 export licensing requirements; expectations with respect to growth, performance and benefits of NVIDIA's products, services, and technologies, including Blackwell, and related trends and drivers; expectations with respect to supply and demand for NVIDIA's products, services, and technologies, including Blackwell, and related matters including inventory, production and distribution; our financial position; projected market growth and trends, market opportunity, demand, and growth drivers; our financial and business outlook; our dividend program; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; third parties adopting our products and technologies; expectations with respect to technology developments, including Rubin, and related trends and drivers; expectations with respect to AI and related industries; and other statements that are not historical facts are forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, which are subject to the "safe harbor" created by those sections based on management's beliefs and assumptions and on information currently available to management and are subject to risks and uncertainties that could cause results to be materially different than expectations.

Important factors that could cause actual results to differ materially include: global economic and political conditions; NVIDIA's reliance on third parties to manufacture, assemble, package and test NVIDIA's products; the impact of technological development and competition; development of new products and technologies or enhancements to NVIDIA's existing product and technologies; market acceptance of NVIDIA's products or NVIDIA's partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of NVIDIA's products or technologies when integrated into systems; and changes in applicable laws and regulations, as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

Many of the products and features described herein remain in various stages and will be offered on a when-and-if-available basis. The statements within are not intended to be, and should not be interpreted as a commitment, promise, or legal obligation, and the development, release, and timing of any features or functionalities described for our products is subject to change and remains at the sole discretion of NVIDIA. NVIDIA will have no liability for failure to deliver or delay in the delivery of any of the products, features or functions set forth herein.

NVIDIA uses certain non-GAAP measures in this presentation including non-GAAP gross margin, non-GAAP gross margin excluding H20 related releases, non-GAAP operating income, non-GAAP net income, and non-GAAP diluted earnings per share. NVIDIA believes the presentation of its non-GAAP financial measures enhances investors' overall understanding of the company's historical financial performance. The presentation of the company's non-GAAP financial measures is not meant to be considered in isolation or as a substitute for the company's financial results prepared in accordance with GAAP, and the company's non-GAAP measures may be different from non-GAAP measures used by other companies. Further information relevant to the interpretation of non-GAAP financial measures, and reconciliations of these non-GAAP financial measures to the most comparable GAAP measures, may be found in the slide titled "Reconciliation of Non-GAAP to GAAP Financial Measures".

Content

- Q2 FY26 Earnings Summary
- Key Announcements This Quarter
- Reconciliation of Non-GAAP to GAAP Financial Measures

Q2 FY26
Earnings Summary

Highlights

Record quarter exceeded our outlook in what continues to be a dynamic external environment

- Total revenue up 56% Y/Y to $46.7B, above outlook of $45.0B +/- 2%
- Data Center up 56% Y/Y to $41.1B
- Gaming up 49% Y/Y to $4.3B

Record Data Center revenue driven by continued Blackwell ramp

- NVIDIA’s Blackwell platform reached record levels with 17% sequential growth
- Cloud service providers, neo clouds, enterprises and sovereigns are all contributing to growth
- We’re at the beginning of an industrial revolution; see $3 to $4 trillion in AI infrastructure spend by the end of the decade

Record Gaming revenue due to Blackwell adoption by gamers, creatives, and AI enthusiasts

- Strong sales continued and we increased supply availability
- For AI enthusiasts, on-device AI run best on RTX GPUs
- With the RTX platform stack, Windows developers can create AI applications designed to run on the world's largest AI PC user base

Q2 FY26 Financial Summary

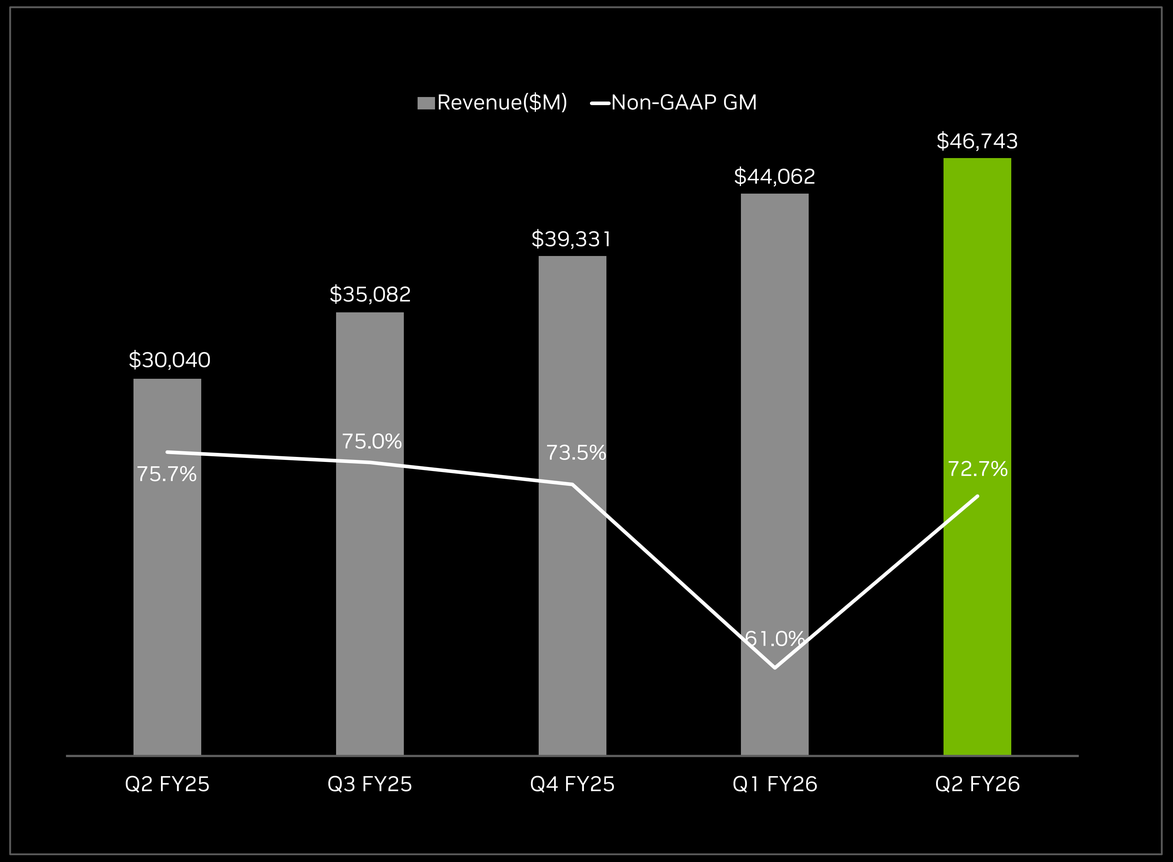

	GAAP			Non-GAAP		
	Q2 FY26	**Y/Y**	**Q/Q**	**Q2 FY26**	**Y/Y**	**Q/Q**
Revenue	$46,743	+56%	+6%	$46,743	+56%	+6%
Gross Margin (GM)	72.4%	-2.7 pts	+11.9 pts	72.7%	-3.0 pts	+11.7 pts
GM excluding H20 related charges/ releases, net				72.3%		+1.0 pt
Operating Income	$28,440	+53%	+31%	$30,165	+51%	+30%
Net Income	$26,422	+59%	+41%	$25,783	+52%	+30%
Diluted EPS	$1.08	+61%	+42%	$1.05	+54%	+30%
Cash Flow from Ops	$15,365	+6%	-44%	$15,365	+6%	-44%

All dollar figures are in millions other than EPS. Refer to Appendix for reconciliation of Non-GAAP measures.

Data Center

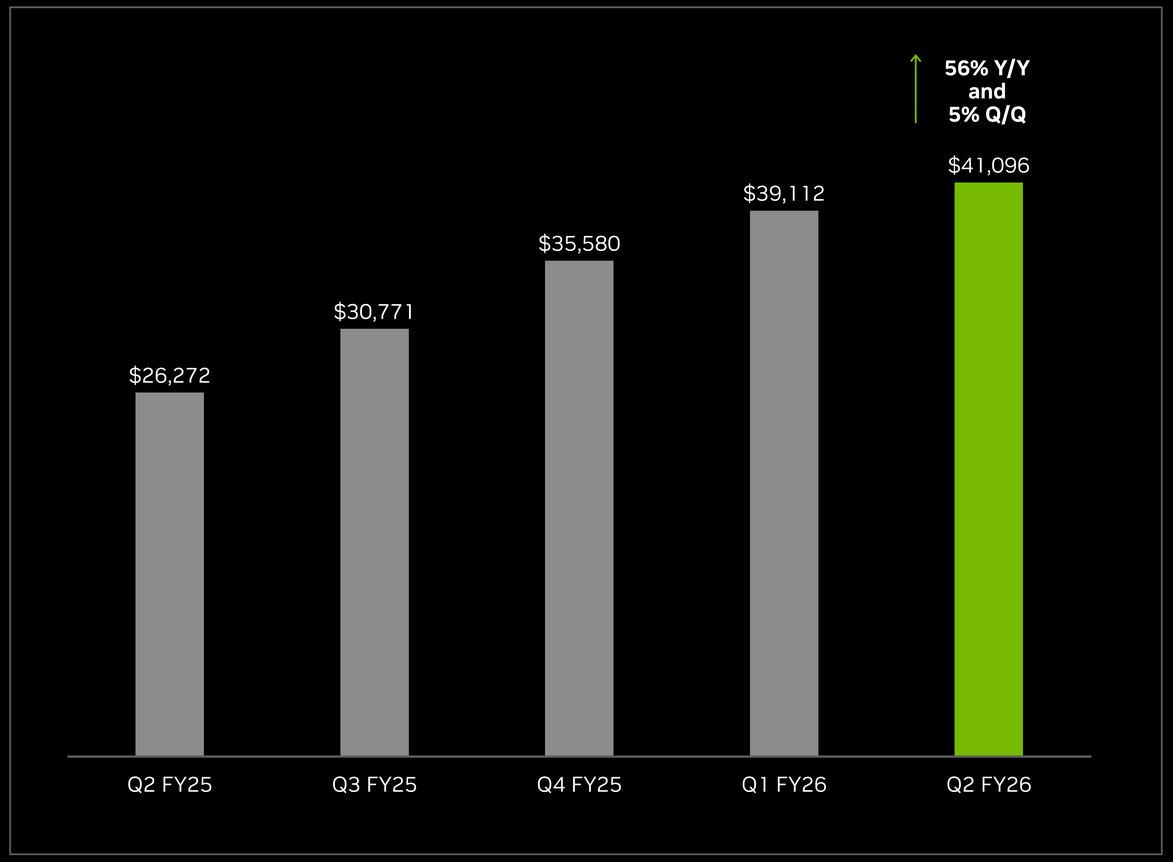

Revenue ($M)

Highlights

- The GB200 NVL system is seeing widespread adoption with deployments at CSPs and consumer internet customers
- New Blackwell Ultra platform had a strong quarter, generating tens of billions in revenue
- Current run rate back at full speed, producing approximately 1,000 racks per week; output expected to accelerate throughout Q3
- Expect widespread market availability of Blackwell Ultra in the second half of the year
- NVIDIA RTX PRO servers are in full production from the world's systems makers; poised to become a multi-billion-dollar product line
- On track to achieve over $20 billion in sovereign AI revenue this year, more than double that of last year
- Networking was a record, up 46% q/q and 98% y/y, with strong demand across Spectrum-X Ethernet, InfiniBand, and NVLink

Gaming

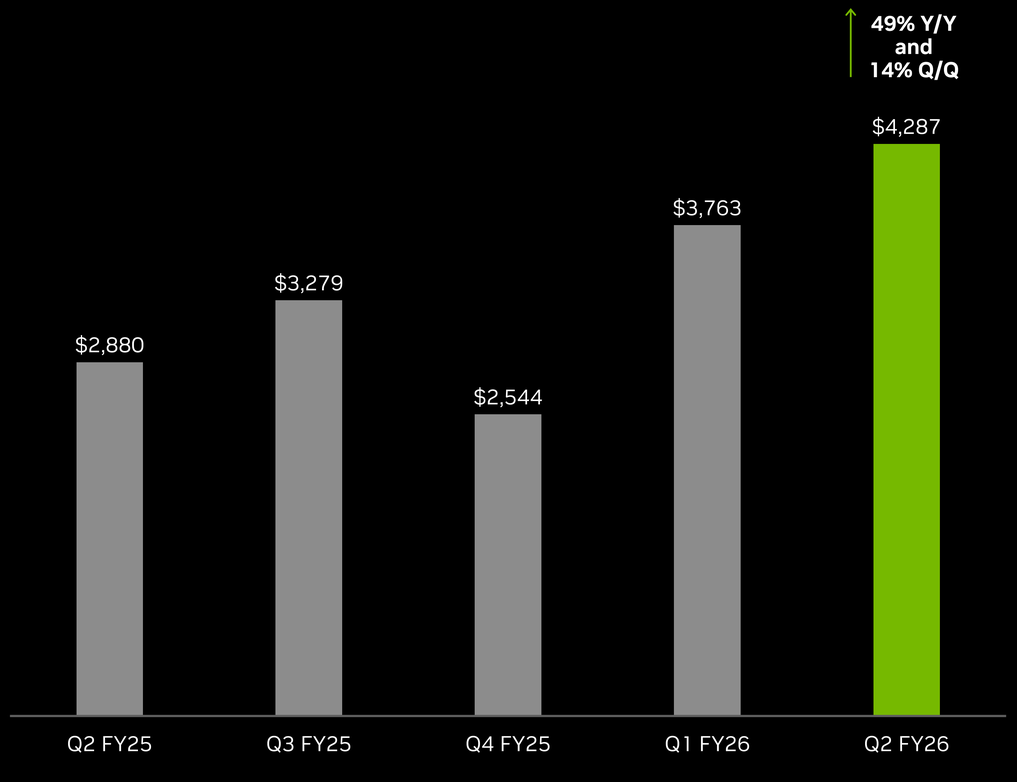

Revenue ($M)

Highlights

- Record gaming revenue driven by the ramp of Blackwell GeForce GPUs, as strong sales continued and supply availability increased
- Shipped GeForce RTX 5060 desktop GPUs this quarter
- Blackwell is coming to GeForce Now in September, the most significant upgrade, offering RTX 5080-class performance, minimal latency, and 5K resolution at 120 frames per second
- GeForce Now catalog doubling to over 4,500 titles, the largest library of any cloud gaming service
- Partnered with OpenAI to optimize their open-source gpt-oss models for high-quality, fast, and efficient inference on millions of RTX-enabled Windows devices

Professional Visualization

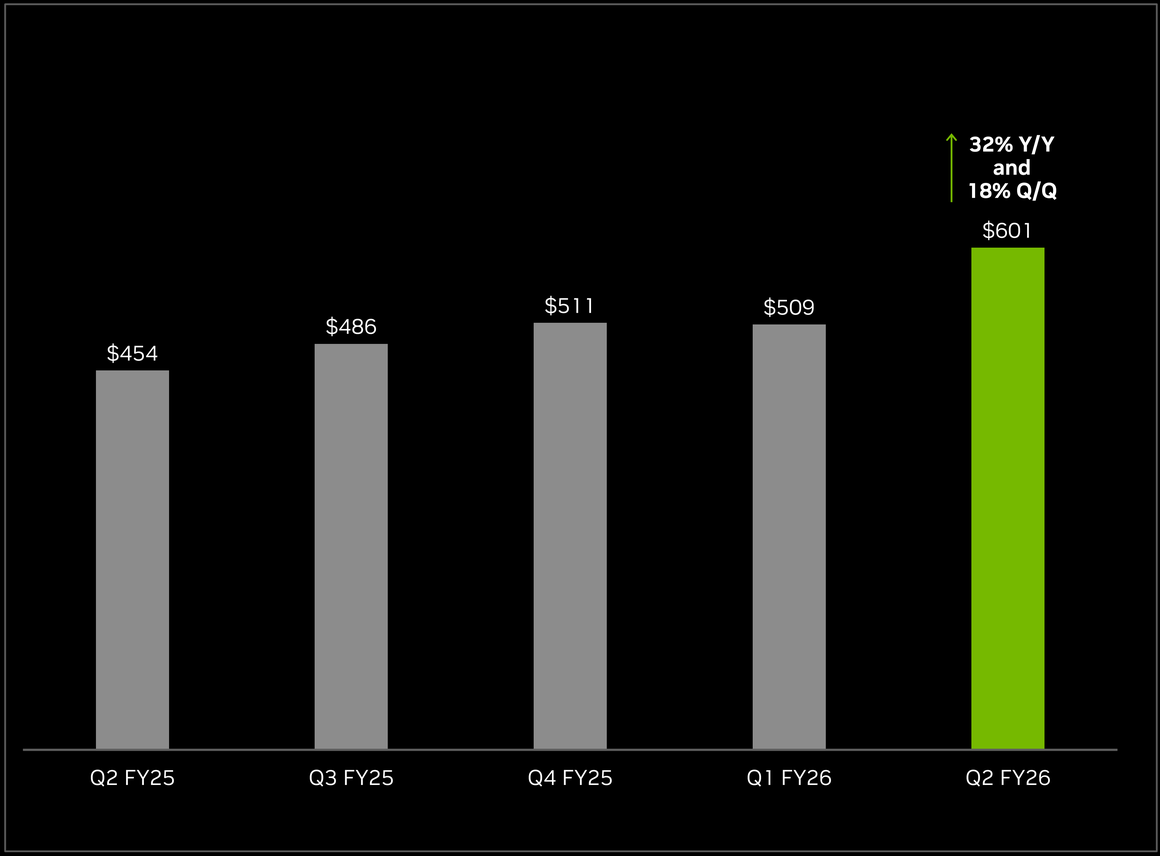

Revenue ($M)

Highlights

- Growth driven by the adoption of high-end RTX workstation GPUs for AI-powered workloads like design, simulation, and prototyping
- Key customers are leveraging NVIDIA solutions to accelerate their operations
- Activision Blizzard uses RTX workstations to enhance creative workflows, while robotics innovator Figure AI powers its humanoid robots with RTX embedded GPUs

Automotive

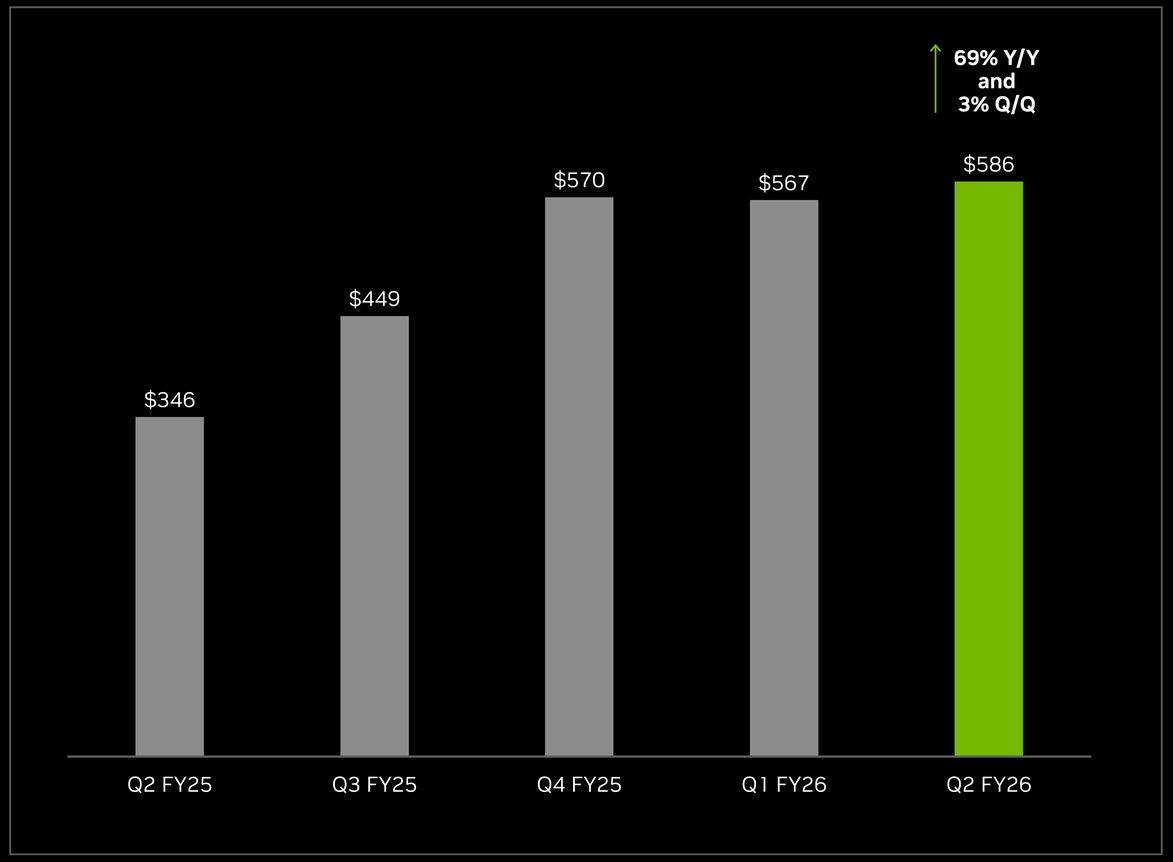

Revenue ($M)

Highlights

- Growth primarily driven by self-driving solutions
- Began shipments of NVIDIA Thor SoC, the successor to Orin
- NVIDIA's full-stack DRIVE AV software platform is now in production, opening up billions in new revenue opportunities for NVIDIA while improving vehicle safety and autonomy

Sources & Uses of Cash

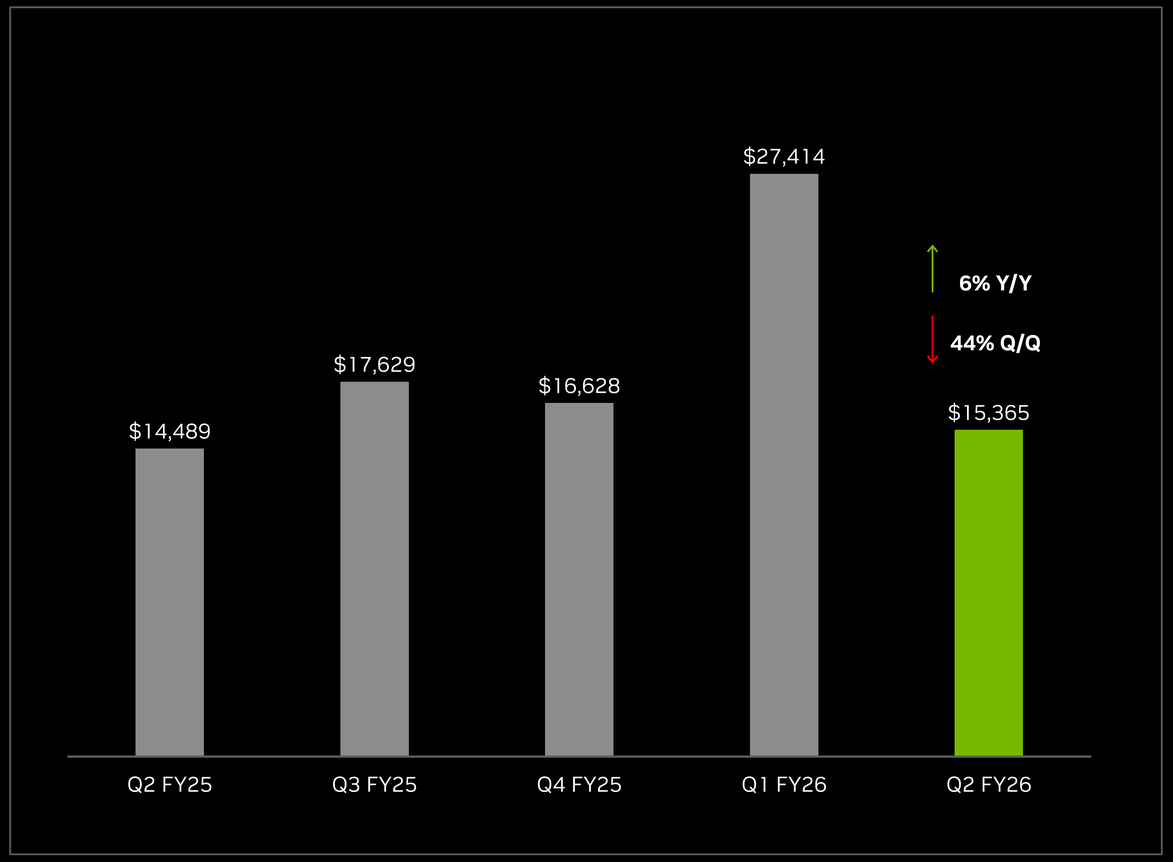

Cash Flow from Operations ($M)

Highlights

- Y/Y increase reflects growth in revenue, partially offset by an increase in working capital
- Q/Q decrease was mainly driven by $8.1B in taxes paid in the quarter
- Returned $10.0B to shareholders in the second quarter through $9.7B of share repurchases and $244M of cash dividends
- Invested $1.9B in capex (includes principal payments on PP&E)
- Ended the quarter with $56.8B in gross cash and $8.5B in debt; $48.3B in net cash

Gross cash is defined as cash/cash equivalents & marketable securities.
Net cash is defined as gross cash less debt.
Debt is defined as principal value of debt.

Q3 FY26 Outlook

Revenue	**$54.0 billion**, plus or minus 2% Does not assume any H20 shipments to China If geopolitical issues recede, could ship $2 to $5 billion dollars in H20 revenue in Q3
Gross Margins	**73.3%** GAAP and **73.5%** non-GAAP, plus or minus 50 basis points Continue to expect to exit the year with non-GAAP gross margins in the mid-70% range
Operating Expense	Approximately **$5.9 billion** GAAP and **$4.2 billion** non-GAAP Expect full year fiscal 2026 operating expense growth to be in the high-30% range
Other Income & Expense	Income of approximately **$500 million** for GAAP and non-GAAP Excluding gains and losses from non-marketable and publicly-held equity securities
Tax Rate	**16.5%** GAAP and non-GAAP, plus or minus 1%, excluding discrete items

Refer to Appendix for reconciliation of Non-GAAP measures.

Key Announcements This Quarter

Europe Builds Over 3,000 Exaflops of NVIDIA AI Infrastructure to Fuel Region's Next Industrial Transformation

- NVIDIA is working with European nations, and technology and industry leaders, to build NVIDIA Blackwell AI infrastructure that will strengthen digital sovereignty, support economic growth and position the continent as a leader in the AI industrial revolution
- In France, Mistral AI is building an end-to-end cloud platform powered by 18K NVIDIA Grace Blackwell systems in the first phase, with plans to expand across multiple sites in 2026
- In the U.K., NVIDIA Cloud Partners Nebius and Nscale announced the first phase of their AI infrastructure development plans to deploy 14K Blackwell GPUs to power new data centers
- In Germany, NVIDIA and partners are building the world's first industrial AI cloud for European manufacturers, powered by DGX B200 systems and RTX PRO Servers
- These deployments will enable the EU to securely develop, train and deploy agentic and physical AI applications

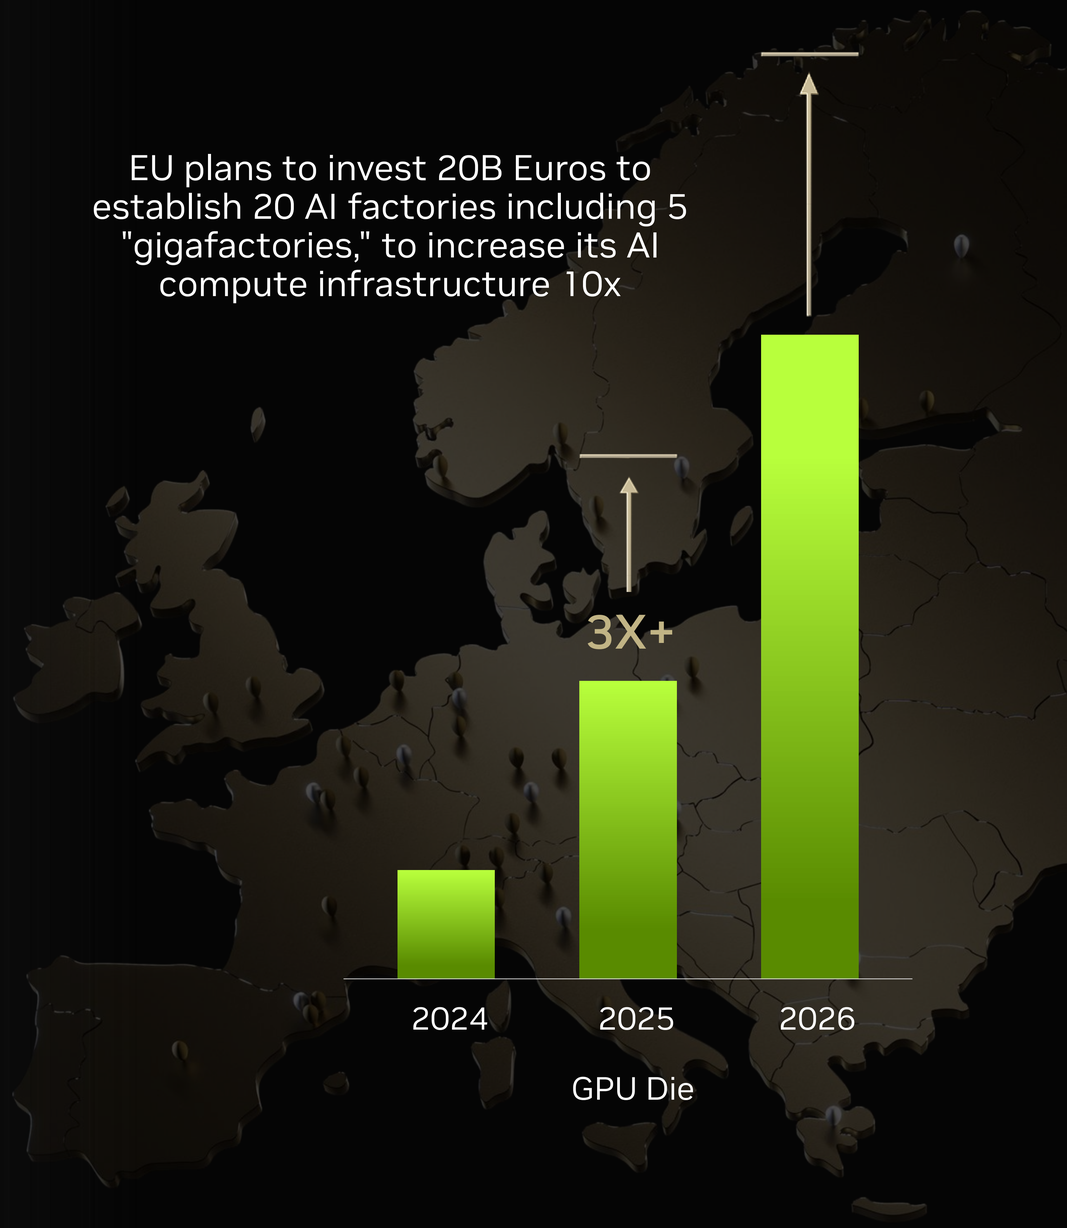

NVIDIA Partners With Europe Model Builders and Cloud Providers to Accelerate Region's Leap Into AI

- NVIDIA is teaming with model builders and cloud providers across Europe and the Middle East to optimize sovereign LLMs, helping to accelerate enterprise AI adoption for the region's industries
- Model builders and AI consortiums are using NVIDIA Nemotron, to maximize model cost efficiency and accuracy for enterprise AI workloads, including agentic AI
- Model post-training and inference will run on AI infrastructure in Europe from NVIDIA Cloud Partners with DGX Cloud Lepton
- The open, sovereign models will provide a foundation for an integrated regional AI ecosystem that reflects local languages and culture

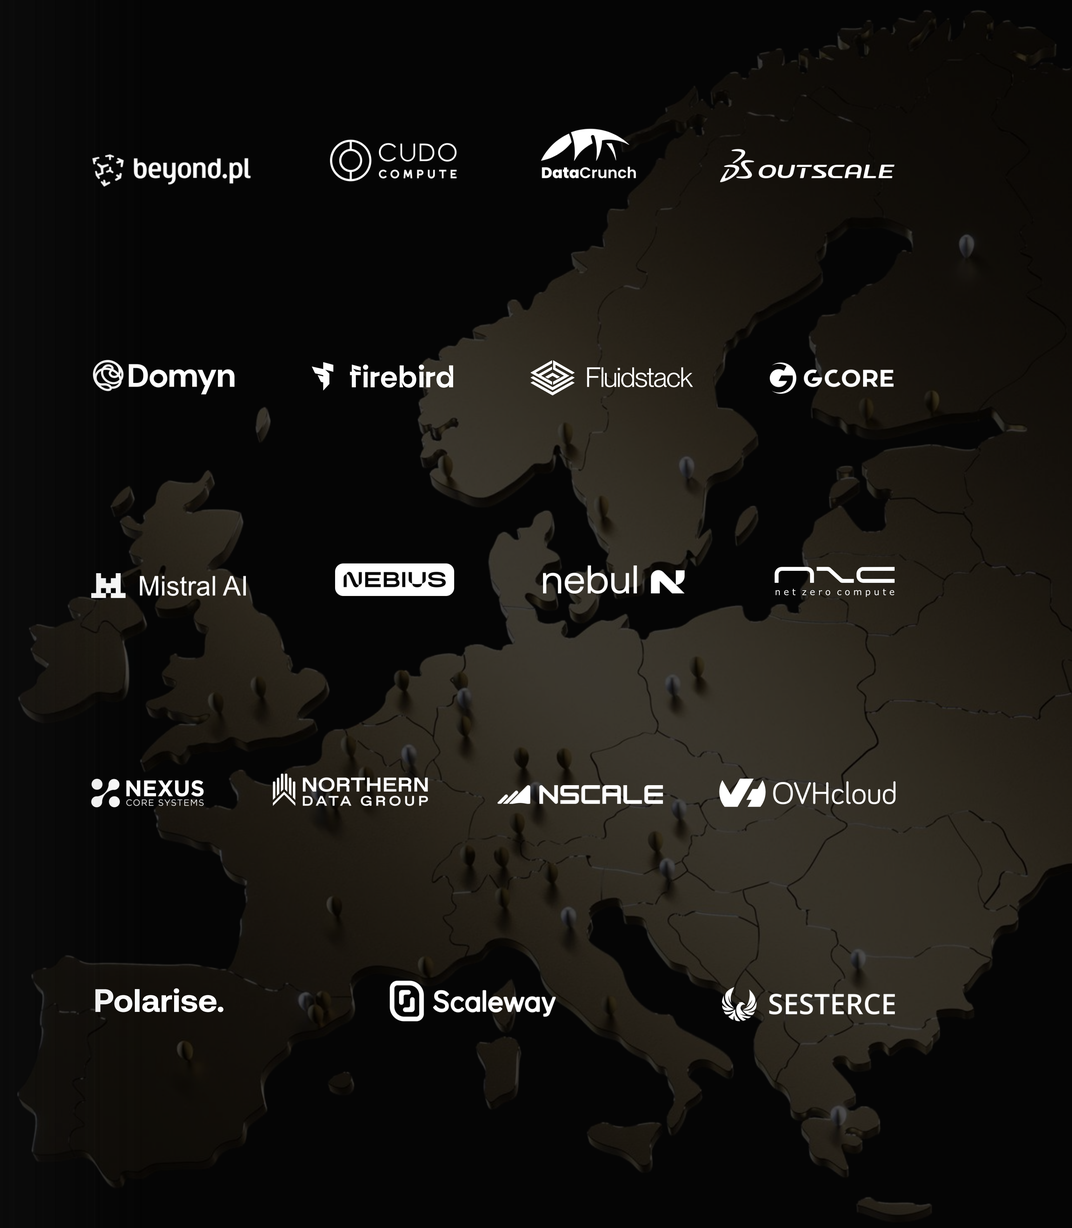

NVIDIA Partners With Novo Nordisk and DCAI to Advance Drug Discovery

- NVIDIA and Novo Nordisk are collaborating to accelerate drug discovery through innovative AI use cases
- Novo Nordisk will use the Gefion sovereign AI supercomputer – powered by NVIDIA DGX SuperPOD – for running drug discovery and agentic AI workloads
- And also:
  - NVIDIA BioNeMo for generative AI-powered drug discovery
  - NVIDIA NIM and NeMo for building customized agentic workflows
  - NVIDIA Omniverse to create physically accurate simulation environments for developing physical AI applications
- Will use NVIDIA BioNeMo, NIM, and NeMo to create customized generative AI models and AI agents that can be used for early research and clinical development
- NVIDIA Omniverse will be used to apply advanced simulation and physical AI technologies

DCAI

Siemens and NVIDIA Expand Partnership to Accelerate AI Capabilities in Manufacturing

SIEMENS | NVIDIA

- Siemens and NVIDIA announced an expansion of their partnership to accelerate the next era of industrial AI and digitalization and enable the factory of the future
- The partnership empowers industrial companies to leverage AI-powered technologies for next-gen factory automation
- Advanced AI agents will work seamlessly across the Siemens Industrial Copilot portfolio, executing entire AI-powered processes without human intervention
- Siemens' Operations Copilot brings generative AI to shopfloor operators; optimized to run on premises with NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs
- Siemens is also integrating NVIDIA BlueField DPUs for AI-driven cybersecurity

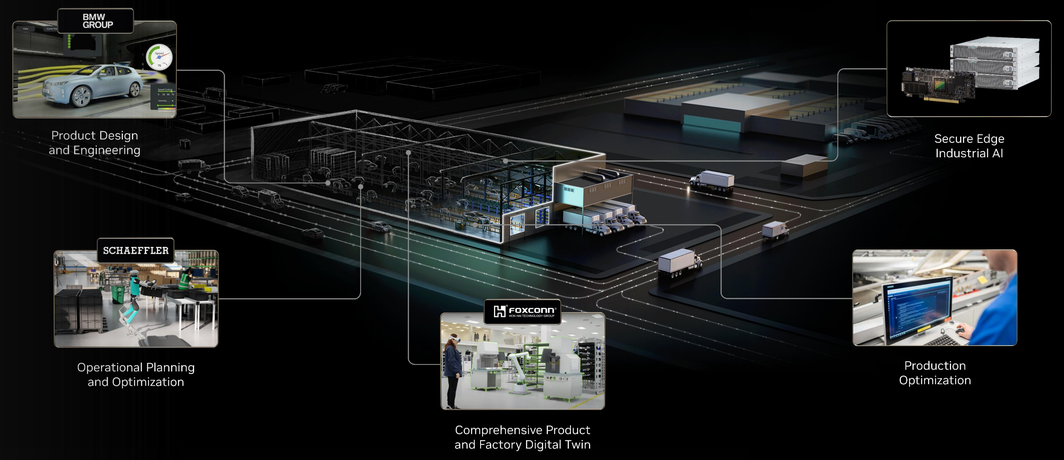

NVIDIA DGX Cloud Lepton Connects Europe's Developers to Global NVIDIA Compute Ecosystem

- NVIDIA is expanding NVIDIA DGX Cloud Lepton — an AI platform featuring a global compute marketplace that connects developers building agentic and physical AI applications — with GPUs now available from a growing network of cloud providers
- Mistral AI, Nebius, Nscale, Firebird, Fluidstack, Hydra Host, Scaleway and Together AI are now contributing NVIDIA Blackwell and other NVIDIA architecture GPUs to the marketplace
- AWS and Microsoft Azure will be the first large-scale cloud providers to participate in DGX Cloud Lepton, joining CoreWeave, Crusoe, Firmus, Foxconn, GMI Cloud, Lambda and Yotta Data Services
- Hugging Face is introducing Training Cluster as a Service, seamlessly connecting AI researchers and developers building foundation models with the NVIDIA compute ecosystem

NVIDIA Blackwell Delivers Breakthrough Performance in Latest MLPerf Training Results

- NVIDIA AI platform delivered the highest performance at scale on every benchmark and powered every result submitted on the benchmark's toughest large language model (LLM)-focused test: Llama 3.1 405B pretraining – first MLPerf training submissions using NVIDIA GB200 NVL72 systems
- NVIDIA was the only platform to submit results on every MLPerf Training v5.0 benchmark – underscoring its exceptional performance and versatility
- These performance leaps highlight advancements in Blackwell that enable faster and more cost-effective training
- NVIDIA is working with companies worldwide to build out AI factories – speeding the training and deployment of next-gen agentic AI applications that produce tokens and intelligence that can be applied to almost every industry and academic domain

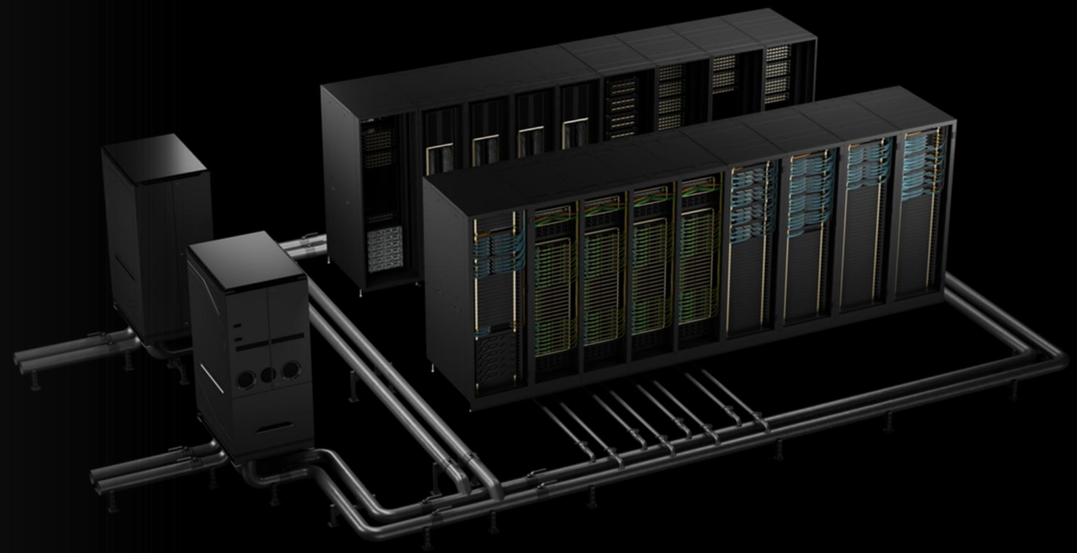

OpenAI and NVIDIA Propel AI Innovation With New Open Models Optimized for the World's Largest AI Inference Infrastructure

- OpenAI's two new open-weight AI reasoning models, oss-120b and gpt-oss-20b, bring cutting-edge AI development to developers, enthusiasts, enterprises, startups and governments everywhere
- This collaboration highlights the power of community-driven innovation and NVIDIA's foundational role in making AI accessible worldwide
- The models were trained on NVIDIA H100 GPUs and run inference best on the hundreds of millions of GPUs running the NVIDIA CUDA platform
- With software optimizations, the models offer optimal inference on NVIDIA GB200 NVL72 systems, achieving industry leading 1.5M tokens per second, driving massive efficiency for inference

OpenAI

NVIDIA

NVIDIA RTX PRO Servers are in Full Production From the World's Systems Makers

- Powered by NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs, these servers speed the shift from traditional CPU systems to accelerated computing platforms
- These air-cooled, PCIe-based systems integrate seamlessly into standard IT environments
- They run traditional enterprise IT applications as well as the most advanced agentic and physical AI applications
- Enterprises can achieve up to 45x better performance and 18x higher energy efficiency when running enterprise workloads
- Nearly 90 companies including many global leaders are already adopting RTX PRO servers
  - Hitachi for real-time simulation and digital twins
  - Lilly for drug discovery
  - Hyundai for factory design and AV validation
  - Disney for immersive storytelling

cisco

DELL Technologies

HPE

Lenovo

SUPERMICRO

Multi Workload Acceleration

vs. NVIDIA L40S

4X
FPS Real-Time Rendering

6X
Throughput LLM Inference

4X
Faster Synthetic Data gen

Data Center Consolidation

vs. CPU

45X
Performance Enterprise Workloads

Projected Performance. Subject to Change
- *Omniverse; Debrecen; Real Time Rendering – FPS. RTX PRO 6000 vs L40S, RT2+DLSS4 On vs DLSS3*
- *Llama3 70B Inference; 8K/256, 20 t/s/usr; RTX PRO 6000 FP4 vs L40S FP8*
- *NVIDIA Cosmos 7B; Text-Video Generation, 2.5s 720p Video; RTX PRO 6000 FP4 vs L40S FP8*
- *NVIDIA RTX PRO Server (2U Chassis, 2x RTX PRO 6000 Blackwell Server Edition GPU, 2x x86 CPU, 3kW) vs. 2x x86 CPU (2U Chassis, 1.1kW); Multi-Workload Relative Performance; Geomean of measured performance speedups for HPC Applications (FP32), Rendering (V-Ray), Spark RAPIDS, Encode Streams (1080p30, HEVC). Shown for representation only. Please contact partners for pricing.*

NVIDIA Opens Portals to World of Robotics With New Omniverse Libraries, Cosmos Physical AI Models and AI Computing Infrastructure

- New NVIDIA Omniverse libraries and Cosmos world foundation models (WFMs) accelerate the development and deployment of physical AI and robotics solutions
- Boston Dynamics, Disney Research, Figure AI, Hexagon, RAI Institute, Lightwheel and Skild AI are adopting Omniverse libraries, Isaac Sim and Isaac Lab to accelerate their AI robotics development, while Amazon Devices & Services is using them to power a new manufacturing solution
- Cosmos Reason is being used by Magna, and adopted by VAST Data, Milestone Systems and Linker Vision for autonomous delivery vehicles, automated traffic monitoring, visual inspection
- Lightwheel, Moon Surgical and Skild AI are using Cosmos Transfer to accelerate physical AI training

Images | Video | Lidar

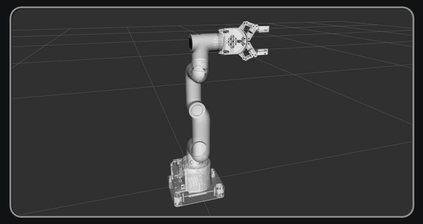

Robot Model
CAD, UDRF, MJCF

NVIDIA Isaac Sim

NVIDIA Blackwell-Powered Jetson Thor Now Available, the Ultimate Platform for Physical AI & Robotics

- These powerful new robotics computers are built for generative reasoning and designed to power millions of robots across diverse industries including manufacturing, logistics, agriculture, transportation, healthcare, and retail
- With over 2,000 teraflops of AI compute, it unlocks new capabilities such as simultaneous processing of multiple AI workstreams and real-time inference.
- Agility Robotics, Amazon Robotics, Boston Dynamics, Caterpillar, Figure, Hexagon, Medtronic and Meta are among the early Jetson Thor adopters
- Also available is the new NVIDIA DRIVE AGX Thor developer kit for safe, secure AV development
  - Runs increasingly complex AV systems equipped with reasoning vision language models, generative AI, and advanced sensor technologies
  - Leading automotive and AV companies are embracing Thor including Aurora, BYD, Continental, Desay SV, Nuro, PlusAI, Volvo Cars, Waabi, Wayve, and WeRide

7.5x Orin
AI Compute

3.5x Orin
Energy Efficiency

3.1x Orin
CPU Performance

10x Orin
I/O Throughput

NVIDIA

Reconciliation of Non-GAAP to GAAP Financial Measures

Reconciliation of Non-GAAP to GAAP Financial Measures

	Non-GAAP	Acquisition-Related and Other Costs (A)	Stock-Based Compensation (B)	H20 related releases	Other (C)	Tax Impact of Adjustments	GAAP
Q2 FY26							
Gross margin ($ in million)	$33,960	(49)	(58)		—	—	$33,853
	72.7%	(0.1)	(0.2)		—	—	72.4%
Gross margin excluding H20 related releases ($ in million)	$33,780	(49)	(58)	180	—	—	$33,853
	72.3%	(0.1)	(0.2)	0.4	—	—	72.4%
Operating income ($ in million)	$30,165	(86)	(1,624)		(15)	—	$28,440
Net income ($ in million)	$25,783	(86)	(1,624)		2,231	118	$26,422
Shares used in diluted per share calculation (millions)	24,532	—	—		—	—	24,532
Diluted EPS	$1.05	—	(0.06)		0.09	—	$1.08

A. *Consists of amortization of intangible assets, transaction costs, and certain compensation charges.*
B. *Stock-based compensation charge was allocated to cost of goods sold, research and development expense, and sales, general and administrative expense.*
C. *Other consists of legal settlement cost, net gain from non-marketable equity securities and publicly-held equity securities, net, and interest expense related to amortization of debt discount.*

Reconciliation of Non-GAAP to GAAP Financial Measures (contd.)

Gross Margin	Non-GAAP	Acquisition-Related and Other Costs (A)	Stock-Based Compensation (B)	GAAP
Q2 FY 2025	75.7%	(0.4)	(0.2)	75.1%
Q3 FY 2025	75.0%	(0.3)	(0.1)	74.6%
Q4 FY 2025	73.5%	(0.3)	(0.2)	73.0%
Q1 FY 2026	61.0%	(0.3)	(0.2)	60.5%

A. Consists of amortization of intangible assets.
B. Stock-based compensation charge was allocated to cost of goods sold.

Reconciliation of Non-GAAP to GAAP Financial Measures (contd.)

($ in Millions)	Q3 FY26 Outlook
Non-GAAP gross margin	73.5%
Impact of stock-based compensation expense, acquisition-related costs, and other costs	0.2%
GAAP gross margin	73.3%
Non-GAAP operating expenses	$4,200
Stock-based compensation expense, acquisition-related costs, and other costs	(1,700)
GAAP operating expenses	$5,900


NVIDIA