Investor Presentation
Q1 FY26

June 2025

Certain matters in this presentation including, but not limited to, statements as to: 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 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 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

- Q1 FY26 Earnings Summary
- Key Announcements This Quarter
- Reconciliation of Non-GAAP to GAAP Financial Measures

Q1 FY26
Earnings Summary

Highlights

Strong quarter exceeding our outlook in a challenging operating environment

- Total revenue up 69% Y/Y to $44.1B, above outlook of $43.0B +/- 2%
- Data Center up 73% Y/Y to $39.1B
- Gaming up 42% Y/Y to $3.8B

Data Center revenue driven by the continued Blackwell ramp – fastest in our company's history

- Seeing a sharp jump in inference demand; Microsoft, Google, OpenAI all seeing a step-function leap in token generation
- AI factory deployments accelerating; nearly 100 NVIDIA-powered AI factories in flight this quarter, up 2x y/y; average number of GPUs powering each factory also doubling in the same period
- Have line of sight to AI factory projects requiring tens of gigawatts of NVIDIA AI infrastructure in the not-too-distant future

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

- Against robust demand backdrop, we greatly improved our supply and availability in Q1 and expect to continue these efforts in Q2
- Brought Blackwell architecture to mainstream gaming with the launch of GeForce RTX 5060 and 5060 Ti

Q1 FY26 Financial Summary

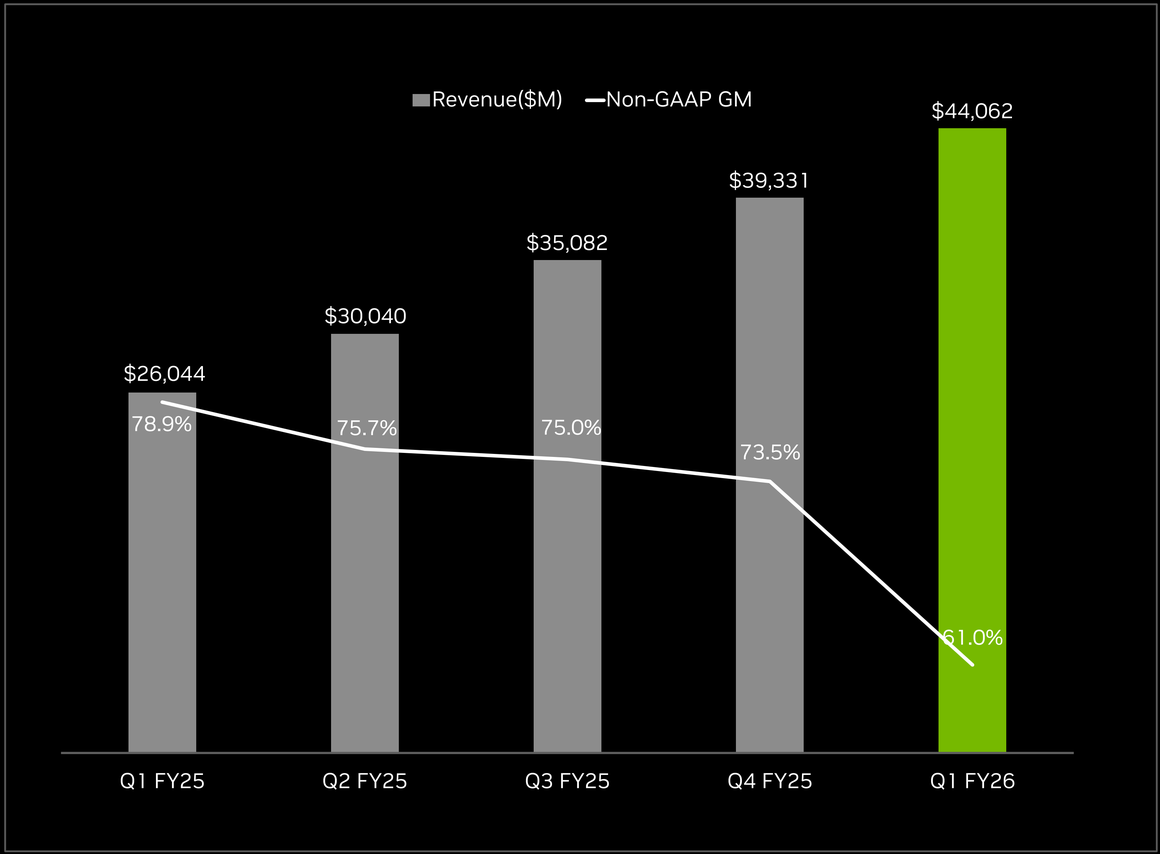

	GAAP			Non-GAAP		
	Q1 FY26*	Y/Y	Q/Q	Q1 FY26*	Y/Y	Q/Q
Revenue	$44,062	+69%	+12%	$44,062	+69%	+12%
Gross Margin*	60.5%	-17.9 pts	-12.5 pts	61.0%	-17.9 pts	-12.5 pts
Operating Income	$21,638	+28%	-10%	$23,275	+29%	-9%
Net Income	$18,775	+26%	-15%	$19,894	+31%	-10%
Diluted EPS	$0.76	+27%	-15%	$0.81	+33%	-9%
Cash Flow from Ops	$27,414	+79%	+65%	$27,414	+79%	+65%

All dollar figures are in millions other than EPS. Refer to Appendix for reconciliation of Non-GAAP measures.
**Q1 FY26 included a $4.5 billion charge associated with H20 excess inventory and purchase obligations*

Data Center

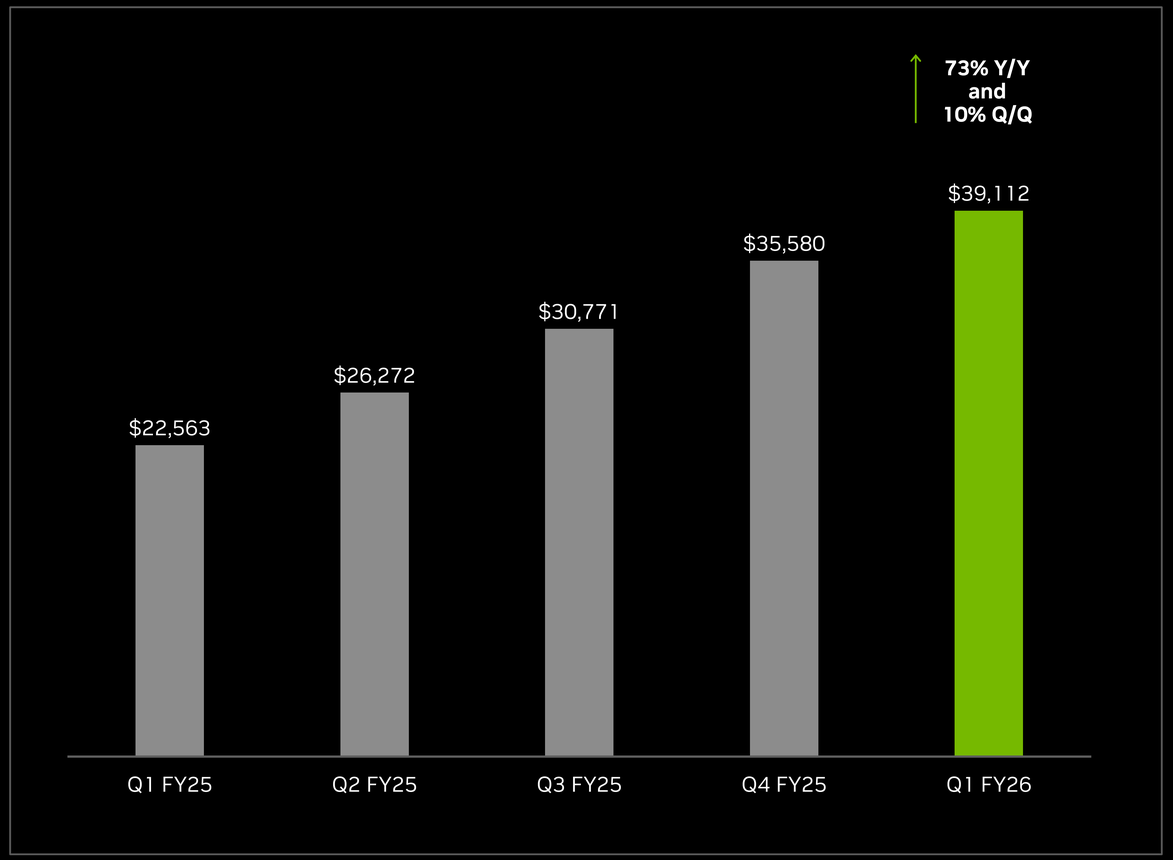

Revenue ($M)

Highlights

- Recognized $4.6B in H20 revenue; also recognized a $4.5 billion charge – write down of inventory and purchase obligations, less than anticipated; unable to ship $2.5B in H20 due to export controls
- Blackwell nearly 70% of DC compute revenue in Q1; Hopper transition nearly complete
- Major hyperscalers deploying nearly 1,000 NVL72 racks or 72,000 Blackwell GPUs per week; on track to further ramp output in Q2
- Sampling of GB300 systems began earlier in May at major CSPs; expect production shipments to commence later in Q2
- Networking grew 64% q/q
  - NVLink a new growth vector; Q1 shipments exceeded a billion dollars
  - Strong q/q & y/y Spectrum-X growth; now annualizing over $8B revenue
  - Spectrum-X adoption widespread: CoreWeave, Microsoft Azure, Oracle Cloud, and xAI; this quarter we added Google Cloud and Meta

Gaming

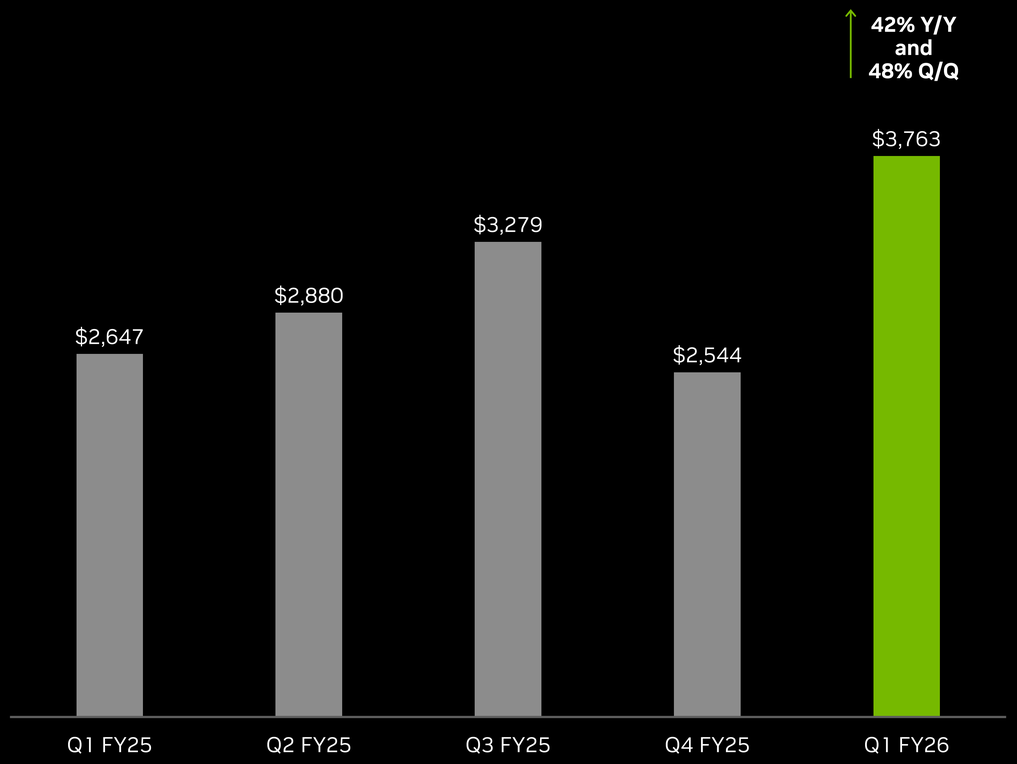

Revenue ($M)

Highlights

- Strong adoption by gamers, creatives, and AI enthusiasts have made Blackwell our fastest ramp ever
- Against robust demand, we greatly improved our supply and availability in Q1 and expect to continue these efforts in Q2
- With a 100M user installed base, GeForce represents the largest footprint for PC developers
- Added to AI PC laptop offerings, including models capable of running Microsoft's Copilot+
- In console gaming, the recently unveiled Nintendo Switch 2 leverages NVIDIA's neural rendering and AI technologies – including next gen custom RTX GPUs with DLSS technology

Professional Visualization

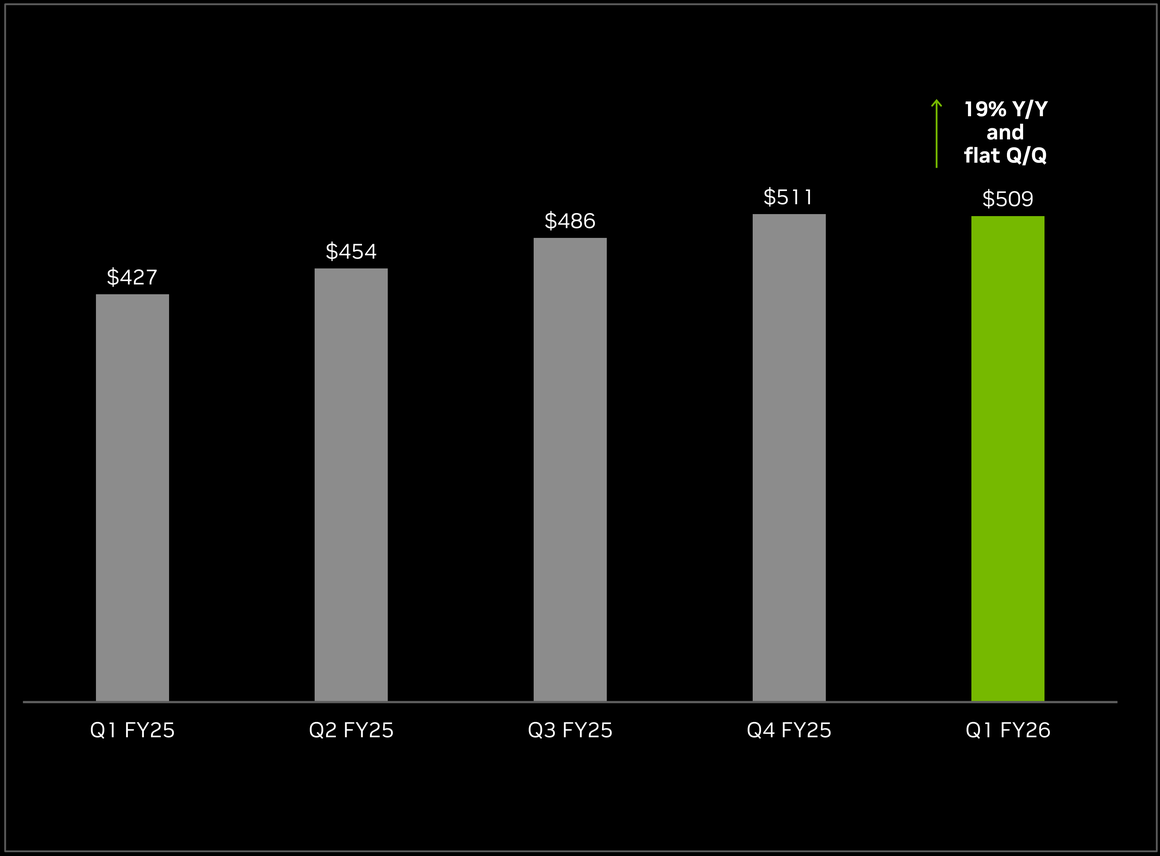

Revenue ($M)

Highlights

- Tariff-related uncertainty temporarily impacted Q1 shipments. End demand for our AI workstations is strong, and we expect sequential revenue growth to resume in Q2
- New DGX Spark will be available in calendar Q3 and DGX Station later this year
- Deepened Omniverse's integration and adoption into some of the world's leading software platforms including Databricks, SAP, and Schneider Electric
- Omniverse seeing great traction with tech manufacturing leaders including TSMC, Quanta, Foxconn and Pegatron

Automotive

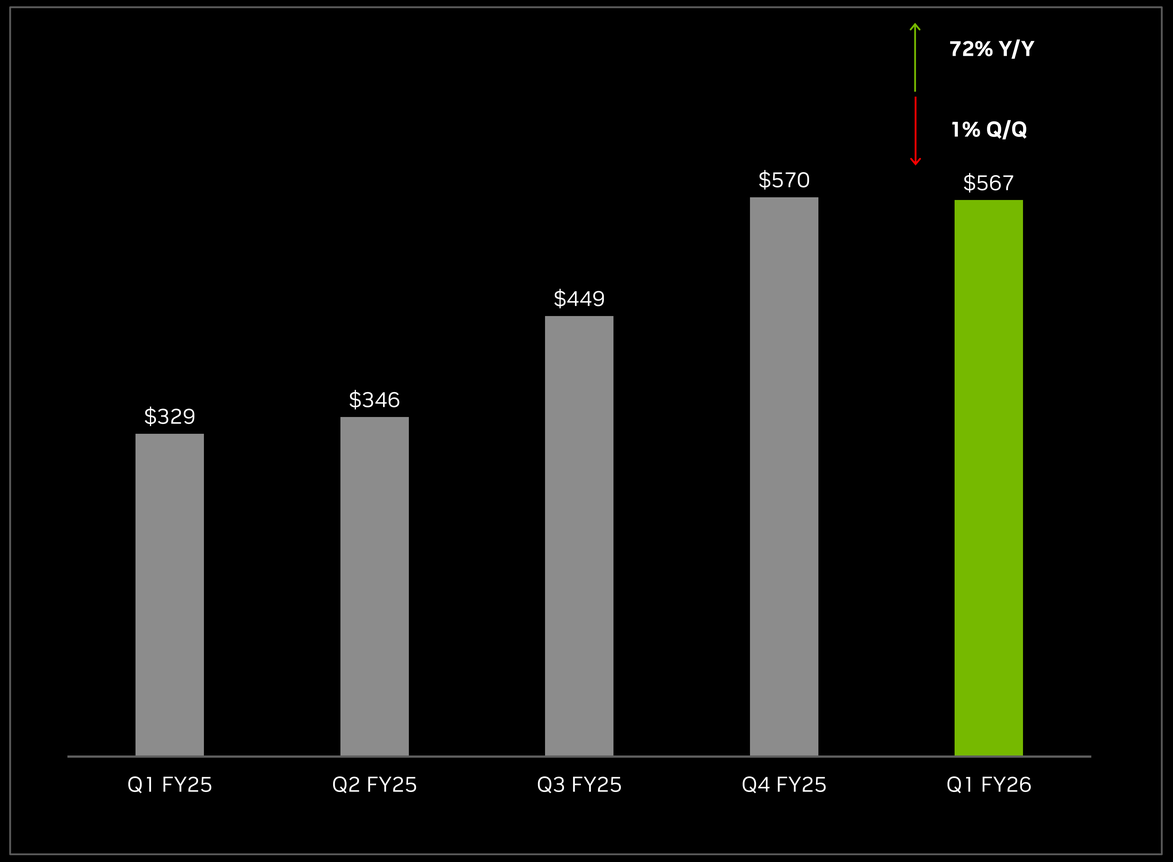

Revenue ($M)

Highlights

- Year-on-year growth was driven by the ramp in self-driving across a number of customers and robust end demand for NEVs
- Now in production with our full stack solution for Mercedes-Benz, starting with the new CLA, hitting roads in the next few months
- Billions of robots, hundreds of millions of autonomous vehicles, and hundreds of thousands of robotic factories and warehouses will be developed

Sources & Uses of Cash

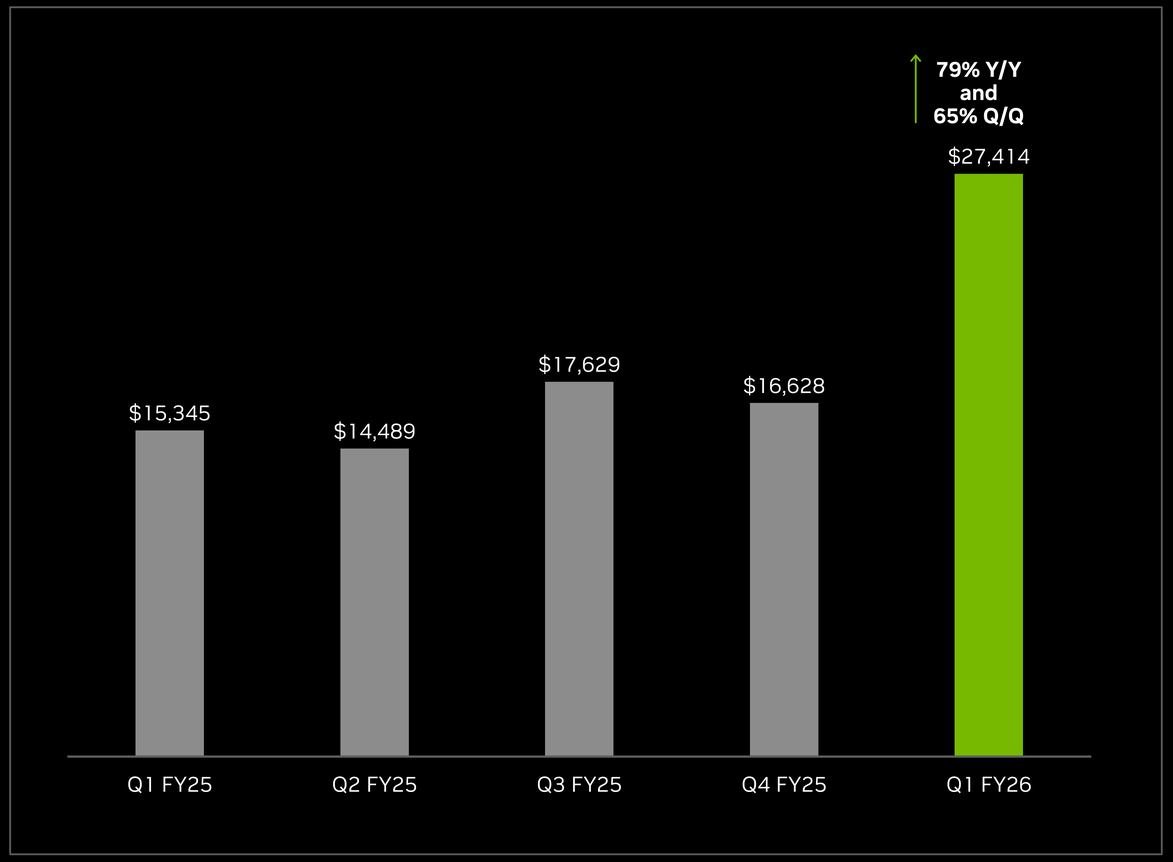

Cash Flow from Operations ($M)

Highlights

- Y/Y increase reflects higher revenue and timing of cash collections
- Q/Q increase was driven by higher revenue, timing of cash collections, and lower cash taxes
- Expect a substantial increase in cash taxes in Q2 related to estimated federal and state cash tax payments
- Utilized cash of $14.3B towards shareholder returns, including $14.1B in share repurchases and $244M in cash dividends
- Invested $1.3B in capex (includes principal payments on PP&E)
- Ended the quarter with $53.7B in gross cash and $8.5B in debt; $42.5B 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.

Q2 FY26 Outlook

Revenue	**$45.0 billion**, plus or minus 2% Expect modest sequential growth across all our platforms Continued ramp of Blackwell to be partially offset by a decline in China revenue Outlook reflects a loss in H20 revenue of ~$8.0 billion
Gross Margins	**71.8%** GAAP and **72.0%** non-GAAP, plus or minus 50 basis points Expect better Blackwell profitability to drive a modest sequential improvement Working toward achieving gross margins in the mid-70% range late this year
Operating Expense	Approximately **$5.7 billion** GAAP and **$4.0 billion** non-GAAP Expect full year FY26 operating expense growth to be in the mid-30% range
Other Income & Expense	Income of approximately **$450 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

NVIDIA Blackwell Ultra Engineered for the Era of AI Reasoning

- Built on the Blackwell architecture, Blackwell Ultra boosts training and test-time inference scaling for agentic and physical AI
- Configurations include:
  - **NVIDIA GB300 NVL72**; a rack-scale system delivering 1.5x more AI performance vs. NVIDIA GB200, and increases Blackwell's revenue opportunity by 50x for AI factories compared to Hopper. It connects 72 Blackwell Ultra GPUs + 36 Grace CPUs to act as a single GPU
  - **NVIDIA HGX B300 NVL16** – baseboard features 11x faster inference on large language models, 7x more compute and 4x more memory compared to Hopper
- AWS, Google Cloud, Microsoft Azure, Oracle Cloud Infrastructure and GPU cloud providers including CoreWeave, will be among the first to offer Blackwell Ultra instances.
- Sampling of GB300 systems began in May at the major CSPs. Production shipments expected to commence in Q2.

NVIDIA’s Spectrum-X and Quantum-X Silicon Photonics Switches Scale AI Factories to Millions of GPUs

- These switches fuse electronics and optics for unparalleled network performance
  - With 4x fewer lasers, they deliver 3.5x more power efficiency, 63x greater signal integrity, 10x better network resiliency, and 1.3x faster deployment than traditional solutions
- NVIDIA Spectrum-X Photonics switches configurations include:
  - 128 ports of 800Gb/s or 512 ports of 200Gb/s
  - 512 ports of 800Gb/s or 2,048 ports of 200Gb/s
- NVIDIA Quantum-X Photonics switches configurations include:
  - 144 ports of 800Gb/s on 200Gb/s SerDes using a liquid-cooled design
- NVIDIA Quantum-X Photonics InfiniBand switches are expected to be available later in 2025, with NVIDIA Spectrum-X Photonics Ethernet switches coming in 2026
- NVIDIA’s silicon photonics ecosystem includes TSMC, Browave, Coherent, Corning, Fabrinet, Foxconn, Lumentum, SENKO, SPIL, Sumitomo Electric Industries and TFC Communication

Spectrum-X Photonics
2nd Half 2026

512 Ports of 800G
2K x 200G

Quantum-X Photonics
2nd Half 2025

128 Ports of 800G
512 x 200G

144 Ports of 800G
576 x 200G

NVIDIA Dynamo Open-Source Inference Platform Drives Reasoning Inference Performance

- Dynamo, an AI inference serving software, maximizes data center token revenue by coordinating reasoning inference across thousands of GPUs processing millions of queries that output hundreds of millions of tokens
- It scales reasoning AI to deliver massive throughput gains – including 30x on DeepSeek R1 – faster response times, and lower model serving costs
- Backed by industry leaders including AWS, Cohere, CoreWeave, Dell, Google Cloud, Meta, Microsoft, NetApp, OpenAI, and Perplexity

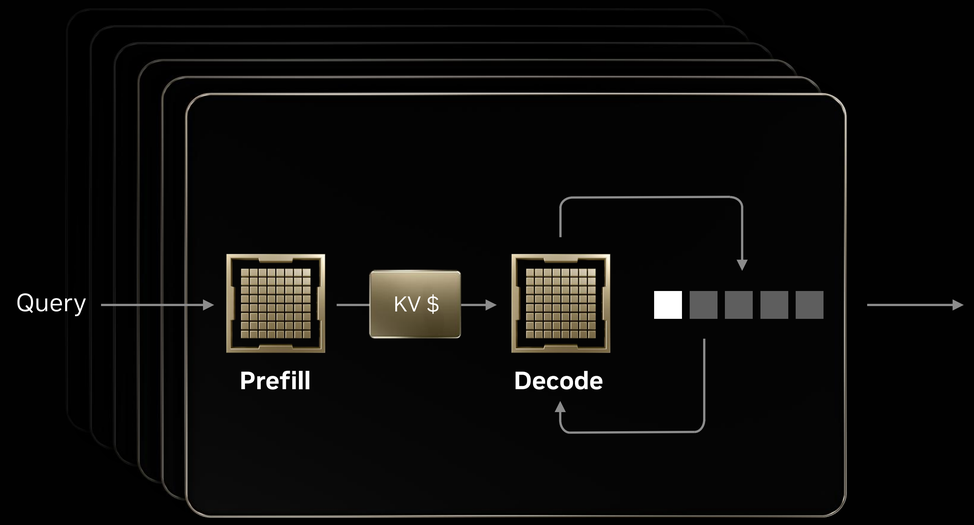

Disaggregated Inference

GPU Resource Allocation

KV Cache Routing

Communication Library (NIXL)

cohere Fireworks AI Meta Microsoft Azure perplexity

PyTorch SGL together.ai VAST vLLM

NVIDIA Turbocharges Agentic AI Development with Open Reasoning AI Models, Platforms, and Partnerships

- New open Llama Nemotron reasoning models accelerate agentic AI development for complex tasks. Post training by NVIDIA boosts model accuracy by up to 20% compared with the base model and optimizes inference speed by 5x compared with other leading open reasoning models
  - Accenture, Amdocs, Atlassian, Box, Cadence, CrowdStrike, Deloitte, IQVIA, Microsoft, SAP, and ServiceNow are pioneering AI agents with NVIDIA
- New NVIDIA AI Data Platform powered by NVIDIA Blackwell GPUs, NVIDIA BlueField DPUs, NVIDIA Spectrum-X networking, and NVIDIA Dynamo enables enterprise agentic workflows. AI Data Platform storage infrastructure with NVIDIA AI-Q Blueprint and NeMo accelerate data extraction and retrieval by up to 15x
  - Dell, HPE, IBM, NetApp, Nutanix, Pure Storage, and others will deliver these customized AI data platforms with NVIDIA
- Oracle and NVIDIA announce a landmark integration of Oracle Cloud Infrastructure and NVIDIA AI Enterprise, making 160+ AI tools and 100+ NIM microservices natively available on OCI

NVIDIA AI Data Platform

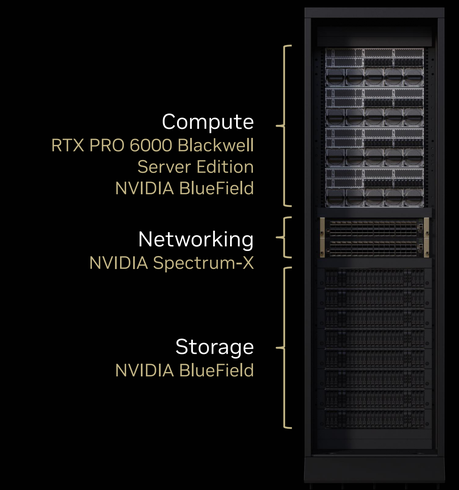

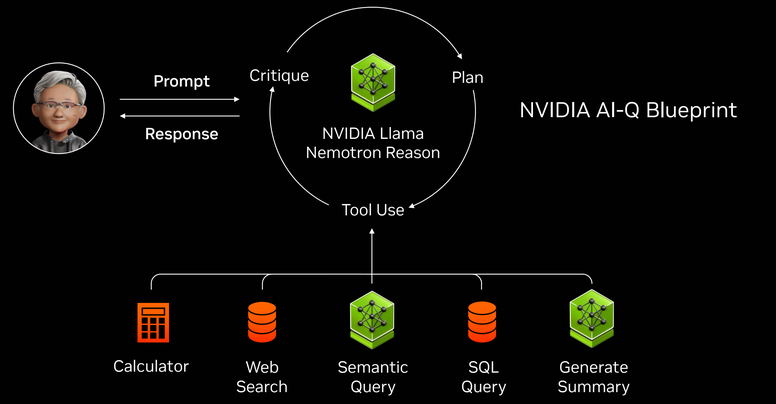

ddn · DELL Technologies · Hewlett Packard Enterprise · Hitachi Vantara · H3C · IBM · NetApp · NUTANIX · PURESTORAGE · VAST · WEKA

NVIDIA Blackwell RTX PRO Comes to Workstations and Servers for Designers, Developers, Data Scientists and Creatives

- We announced the NVIDIA RTX Pro Blackwell series - new desktop, laptop, and server GPUs, including the NVIDIA RTX PRO 6000 Blackwell Server Edition
  - They redefine professional workflows with AI inference, ray tracing, and neural rendering for AI, technical, creative, engineering, and design professionals
  - NVIDIA RTX PRO Blackwell GPUs feature up to 2x RT Core performance, 4,000 AI TOPS, FP4 precision and DLSS 4 multi-frame generation
- RTX PRO 6000 servers power NVIDIA AI and Omniverse applications for enterprises, and with a new NVIDIA Enterprise AI Factory validated design, are driving a shift in data centers from CPU-based systems to GPU-accelerated infrastructure
  - These servers are geared with new CX8 Switch-SuperNIC, providing the high-speed connection fabric to connect multiple GPUs to create an enterprise AI supercomputer
  - They offer enhanced performance and energy efficiency for virtually all enterprise workloads
- The NVIDIA Enterprise AI Factory validated design provides a blueprint for on-premises infrastructure, featuring RTX PRO Blackwell servers, NVIDIA Spectrum-X Ethernet, BlueField DPUs, NVIDIA-Certified Storage, and NVIDIA AI Enterprise software

New NVIDIA Foundation Models, Blueprints, Tools, and Partnerships to Advance Robotics With AI Pre-training, Post-training, and Reasoning

- We launched NVIDIA Cosmos World Foundation Models (WFMs) including Cosmos Reason, Transfer, and Predict models and accompanying physical AI data tools, designed to revolutionize how AI understands and interacts with the physical world
- NVIDIA's Omniverse physical AI operating system is fueling industrial digitalization through expanded partnerships and new capabilities with major industrial software, service providers and leading companies integrating Omniverse and leveraging Omniverse Blueprints to enhance their industrial operations
  - Accenture, Ansys, Cadence, Databricks, Dematic, Hexagon, Omron, SAP, Schneider, Siemens connect Omniverse to leading software tools
  - Four new Blueprints enable robot-ready factories and large-scale synthetic data generation
  - Foxconn, GM, Hyundai Motor Group, KION Group, Mercedes-Benz, Pegatron and Schaeffler adopt Omniverse for Industrial AI
- GE HealthCare is using the new NVIDIA Isaac for Healthcare simulation platform to advance innovation in autonomous imaging
- GM will harness NVIDIA’s three core accelerated computing platforms for next-generation vehicles, factories, and robots

NVIDIA Omniverse

Cosmos

NVIDIA Powers Humanoid Robot Industry with Cloud-to-Robot Computing Platforms for Physical AI

- NVIDIA announced significant advancements in humanoid robot development, including the updated NVIDIA Isaac GR00T N1.5 foundation model for humanoid reasoning, the NVIDIA Isaac GR00T-Dreams blueprint for synthetic motion data generation, and Blackwell systems to accelerate humanoid development
- Leading robotics companies like Agility Robotics, Boston Dynamics, Fourier, General Robotics, and XPENG Robotics are adopting NVIDIA Isaac platform technologies to accelerate robotics development and deployment
- When developing the GR00T N1.5 model NVIDIA Research used the GR00T-Dreams blueprint to generate synthetic training data in just a few hours versus a manual teleoperation process of data collection that would take large teams of human demonstrators

Global Computer Makers Join NVIDIA to Launch AI-First DGX Personal Computing Systems

- Taiwan's leading system manufacturers are set to build NVIDIA DGX Spark and DGX Station systems, enabling a global ecosystem of developers, data scientists and researchers to prototype, fine-tune and inference AI models
- Equipped with the NVIDIA GB10 Grace Blackwell Superchip, the DGX Spark delivers up to 1 petaflop of AI compute. DGX Station features the NVIDIA GB300 Grace Blackwell Ultra Desktop Superchip and NVIDIA ConnectX-8 SuperNIC, offering up to 20 petaflops of AI performance
- Both systems run the NVIDIA DGX operating system and NVIDIA AI software stack to seamlessly deploy workloads to any accelerated data center or cloud infrastructure
- DGX Spark will be available from Acer, ASUS, Dell Technologies, GIGABYTE, HP, Lenovo and MSI, as well as global channel partners, starting in July. DGX Station is expected to be available from ASUS, Dell Technologies, GIGABYTE, HP and MSI later this year

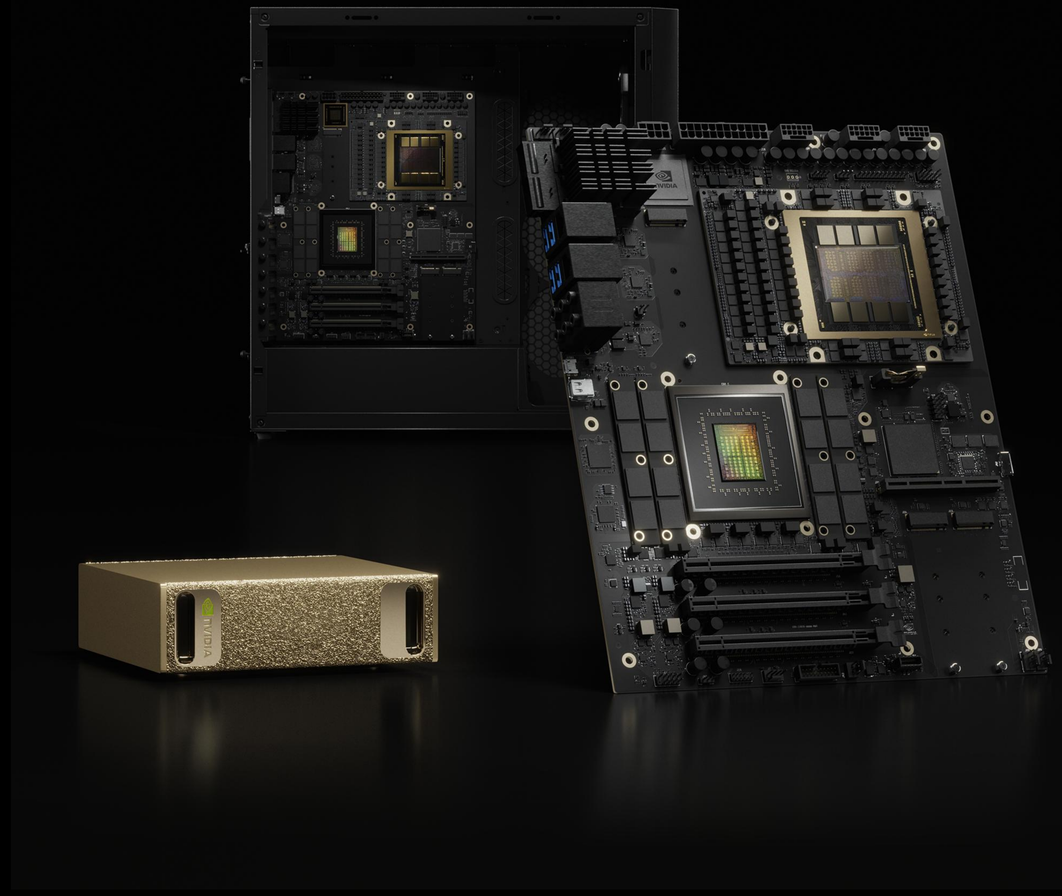

NVLink Fusion for Industry to Build Semi-Custom AI Infrastructure with NVIDIA Partner Ecosystem

- We unveiled NVIDIA NVLink Fusion – new NVLink chiplet silicon that lets industries build semi-custom AI infrastructure connected to NVIDIA's vast AI infrastructure ecosystem through NVLink, the world's most advanced and widely adopted computing fabric
- NVLink Fusion equips cloud providers with an easy path to scale out hybrid architecture AI factories to millions of NVIDIA GPUs and custom ASICs, unified with NVIDIA's NVLink rack-scale systems and NVIDIA networking platform featuring ConnectX-8 SuperNICs, Spectrum-X Ethernet and Quantum-X800 InfiniBand switches, with co-packaged optics available soon
- MediaTek, Marvell, Alchip Technologies, Astera Labs, Synopsys and Cadence are among the first to adopt NVLink Fusion, enabling custom silicon scale-up to meet the requirements of demanding AI workloads
- Using NVLink Fusion, Fujitsu and Qualcomm Technologies CPUs can be integrated with NVIDIA GPUs to build high-performance NVIDIA AI factories

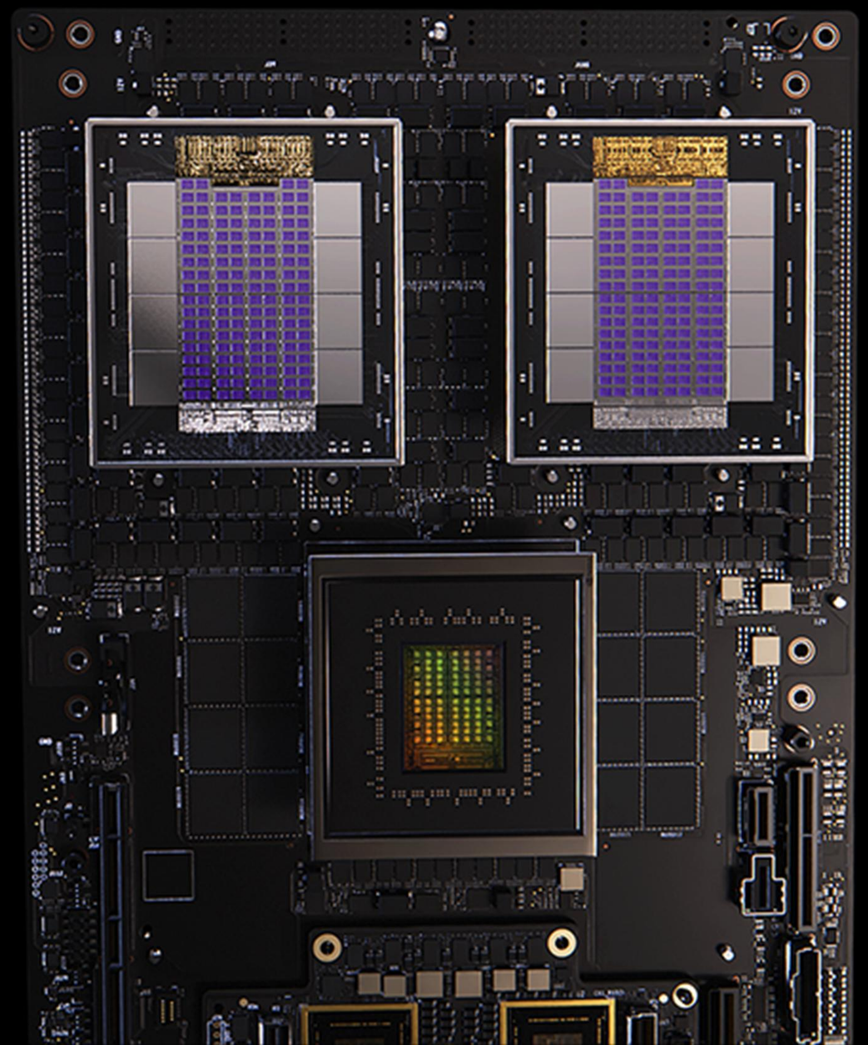

NVIDIA, Foxconn, and Taiwan Partner to Build Taiwan AI Cloud Service to Advance AI for Region's Science and Industry Technology

- NVIDIA and Foxconn Hon Hai Technology Group announced they are deepening their partnership and working with the Taiwan government to build an AI factory supercomputer that will deliver state-of-the-art NVIDIA infrastructure
- TSMC will be the leading customer. The Taiwan government will provide computing credits to spur AI in education, research, startups, and industry. This will strengthen Taiwan's position as a leader in science and technology and fuel innovation across Foxconn's three core pillars - smart cities, electric vehicles (EVs) and manufacturing
- The Foxconn AI factory will feature NVIDIA Blackwell Ultra systems, including the NVIDIA GB300 NVL72 rack-scale solution with NVIDIA NVLink, Quantum InfiniBand and Spectrum-X Ethernet networking
- Taiwan National Science and Technology Council to be among the first to harness the Foxconn AI supercomputer. The system will provide AI cloud computing resources to the Taiwan technology ecosystem, accelerating AI development and adoption

NVIDIA

FOXCONN®
HON HAI TECHNOLOGY GROUP

NVIDIA Blackwell Takes Pole Position in Latest MLPerf Inference Results

- Marked NVIDIA's first MLPerf submission using the rack-scale NVIDIA GB200 NVL72 system designed for AI reasoning
- NVIDIA Blackwell set new MLPerf inference records:
  - GB200 NVL72 system delivered up to 30x higher throughput on the Llama 3.1 405B benchmark vs. NVIDIA H200 NVL8. Only NVIDIA and its partners submitted results on this benchmark
  - NVIDIA DGX B200 system performance with eight Blackwell GPUs jumped 3x vs. eight H200 GPUs on the more challenging version of the Llama 2 70B benchmark
- Three years after launch, NVIDIA Hopper AI factory value continues increasing:
  - On the Llama 2 70B benchmark, H100 GPU throughput has increased by 1.5x vs. a year ago. The H200 GPU with larger and faster GPU memory, extends that increase to 1.6x
  - Hopper also ran every benchmark – Hopper can run a wide range of workloads and keep pace as models and usage scenarios grow more challenging
- 15 partners submitted stellar results on the NVIDIA platform – including CoreWeave, Dell, Google Cloud, Oracle Cloud Infrastructure – reflecting the reach of the NVIDIA platform

GB200 NVL72 Increases Token Throughput by 30x

New Llama 3.1 405B benchmark record

GB200 NVL72

72 Blackwell GPUs
36 Grace CPUs
13.4 TB HBM3e | 576 GB/s
130 TB/s NVLink

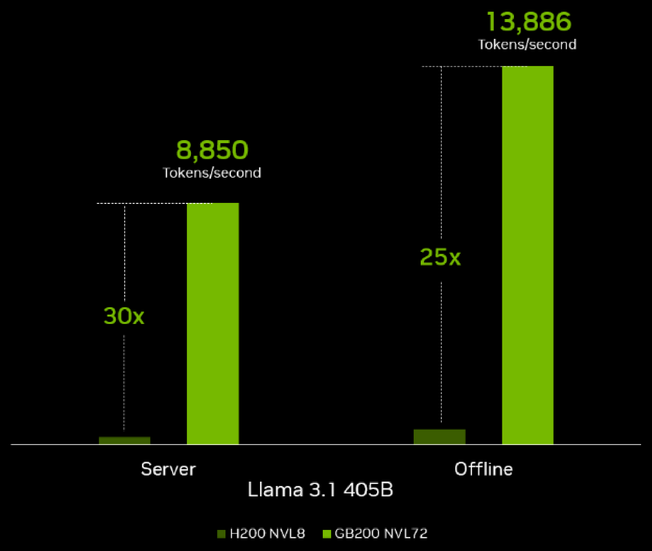

MLPerf Inference v5.0, Closed, Data Center. Results retrieved from www.mlcommons.org on April 2, 2025. Results retrieved from the following entries: 5.0-5.0-0058, 5.0-0060. The MLPerf name and logo are registered and unregistered trademarks of MLCommons Association in the United States and other countries. All rights reserved. Unauthorize use strictly prohibited. See www.mlcommons.org for more information

NVIDIA to Manufacture American-Made AI Supercomputers in US for First Time

- NVIDIA is working with its manufacturing partners to design and build factories that, for the first time, will produce NVIDIA AI supercomputers entirely in the US
- With partners, the company has commissioned more than a million square feet of manufacturing space to build and test NVIDIA Blackwell chips in Arizona and AI supercomputers in Texas
- Within the next four years, NVIDIA plans to produce up to half a trillion dollars of AI infrastructure in the United States through partnerships with TSMC, Foxconn, Wistron, Amkor and SPIL
- Manufacturing NVIDIA AI chips and supercomputers for American AI factories is expected to create hundreds of thousands of jobs and drive trillions of dollars in economic security over the coming decades

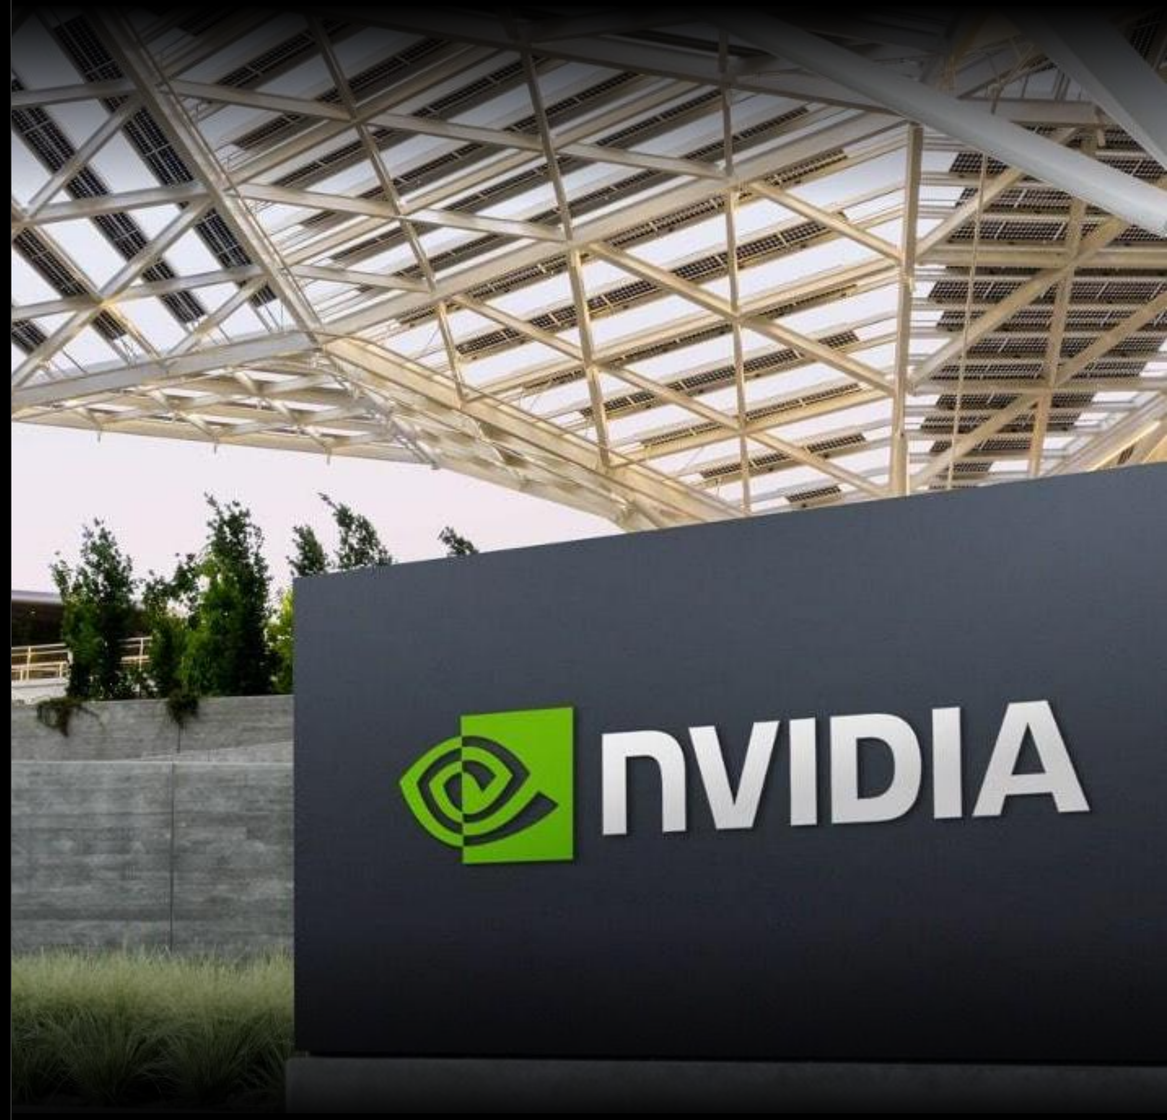

Reconciliation of Non-GAAP to GAAP Financial Measures

Reconciliation of Non-GAAP to GAAP Financial Measures

	Non-GAAP	Acquisition-Related and Other Costs (A)	Stock-Based Compensation (B)	Other (C)	Tax Impact of Adjustments	GAAP
Q1 FY26*						
Gross margin ($ in million)	$26,858	(123)	(64)	(3)	—	$26,668
	61.0%	(0.3)	(0.2)	—	—	60.5%
Operating income ($ in million)	$23,275	(160)	(1,474)	(3)	—	$21,638
Net income ($ in million)	$19,894	(160)	(1,474)	(179)	694	$18,775
Shares used in diluted per share calculation (millions)	24,611	—	—	—	—	24,611
Diluted EPS	$0.81	(0.01)	(0.06)	(0.01)	0.03	$0.76

**Q1 FY26 included a $4.5 billion charge associated with H20 excess inventory and purchase obligations*
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, losses 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
Q1 FY 2025	78.9%	(0.4)	(0.1)	78.4%
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%

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)	Q2 FY26 Outlook
Non-GAAP gross margin	72.0%
Impact of stock-based compensation expense, acquisition-related costs, and other costs	(0.2%)
GAAP gross margin	71.8%
Non-GAAP operating expenses	$4,000
Stock-based compensation expense, acquisition-related costs, and other costs	1,700
GAAP operating expenses	$5,700


NVIDIA