Investor Presentation
Q2 FY25

September 3, 2024

Except for the historical information contained herein, certain matters in this presentation including, but not limited to, statements as to: our financial position; our business plans; our markets, market opportunity, demand and growth drivers; supply for our products; our financial outlook; the benefits, impact, performance, features, and availability of our products and technologies; third parties adopting or using our products or technologies and the benefits and impact thereof; the benefits and impact of our collaborations or partnerships with third parties, and the features and availability of their offerings; our customers continuing to accelerate their Hopper architecture purchases while gearing up to adopt Blackwell; demand coming from frontier model makers, consumer internet services, and tens of thousands of companies building gen AI applications; sovereign AI revenue reaching low double-digit billions this year; every PC with RTX capable of delivering up to1,300 AI TOPS; game developers continuing to embrace NVIDIA ACE; RTX being set to revolutionize consumer experiences with gen AI; Blackwell being widely sampling; timing of Blackwell production ramp and shipment; demand and supply for Blackwell; Hopper shipment; GeForce NOW library continuing to expand; NVIDIA MGS modular reference architecture in over 100 Blackwell systems, providing computer manufacturers with a reference architecture to quickly and cost-effectively build Blackwell design configurations; NVIDIA Spectrum-X Ethernet networking platform gaining wide adoption; CoreWeave, GMO Internet Group, Lambda, Scaleway, STPX Global, and Yotta bringing extreme networking performance to their AI infrastructures; plans announcing the launch of new Spectrum-X products every year, delivering increased bandwidth, ports, enhanced software feature sets, and programmability to drive leading AI Ethernet networking performance; NIM enabling enterprises to maximize their infrastructure investments; technology partners integrating NIM into their platforms; enterprises being able to deploy AI applications in production with NIM through the NVIDIA AI Enterprise software platform; NVIDIA ACE generative AI microservices enabling the creation and operation of digital humans or generative AI avatars; the next wave of AI being physical AI or AI that understands the laws of physics and NVIDIA Omniverse and NVIDIA Isaac being key platforms driving this transition; NVIDIA's RTX AI PCs being set to revolutionize consumer experiences with generative AI; major advancements to OpenUSD expanding adoption of the universal 3D data interchange framework to robotics, industrial design and engineering, and accelerating developers' abilities to build highly accurate virtual worlds; availability of the world's first generative AI models for OpenUSD; the integration of OpenUSD pipelines; new AI microservices allowing roboticists to enhance simulation workflows for generative physical AI; enterprises being able to continually refine their AI applications based on user feedback with NIM Agent Blueprints, creating a data-driven AI flywheel are forward-looking statements.

These forward-looking statements and any other forward-looking statements that go beyond historical facts that are made in this presentation are subject to risks and uncertainties that may cause actual results to differ materially. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences and demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems and other factors.

NVIDIA has based these forward-looking statements largely on its current expectations and projections about future events and trends that it believes may affect its financial condition, results of operations, business strategy, short-term and long-term business operations and objectives, and financial needs. These forward-looking statements are subject to a number of risks and uncertainties, and you should not rely upon the forward-looking statements as predictions of future events. The future events and trends discussed in this presentation may not occur and actual results could differ materially and adversely from those anticipated or implied in the forward-looking statements. Although NVIDIA believes that the expectations reflected in the forward-looking statements are reasonable, the company cannot guarantee that future results, levels of activity, performance, achievements or events and circumstances reflected in the forward-looking statements will occur. Except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances. For a complete discussion of factors that could materially affect our financial results and operations, please refer to the reports we file from time to time with the SEC, including our most recent Annual Report on Form 10-K, Quarterly Reports on Form 10-Q, and Current Reports on Form 8-K. Copies of reports we file with the SEC are posted on our website and are available from NVIDIA without charge.

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

- Q2 FY25 Earnings Summary
- Key Announcements This Quarter
- Reconciliation of Non-GAAP to GAAP Financial Measures

Q2 FY25
Earnings Summary

Highlights

Record quarter driven by strong Data Center growth

- Total revenue up 122% Y/Y to $30.0B, well above outlook of $28.0B +/- 2%
- Record Data Center up 154% Y/Y to $26.3B
- Gaming up 16% Y/Y to $2.9B

Data Center revenue driven by strong demand for NVIDIA Hopper GPU computing and networking platforms

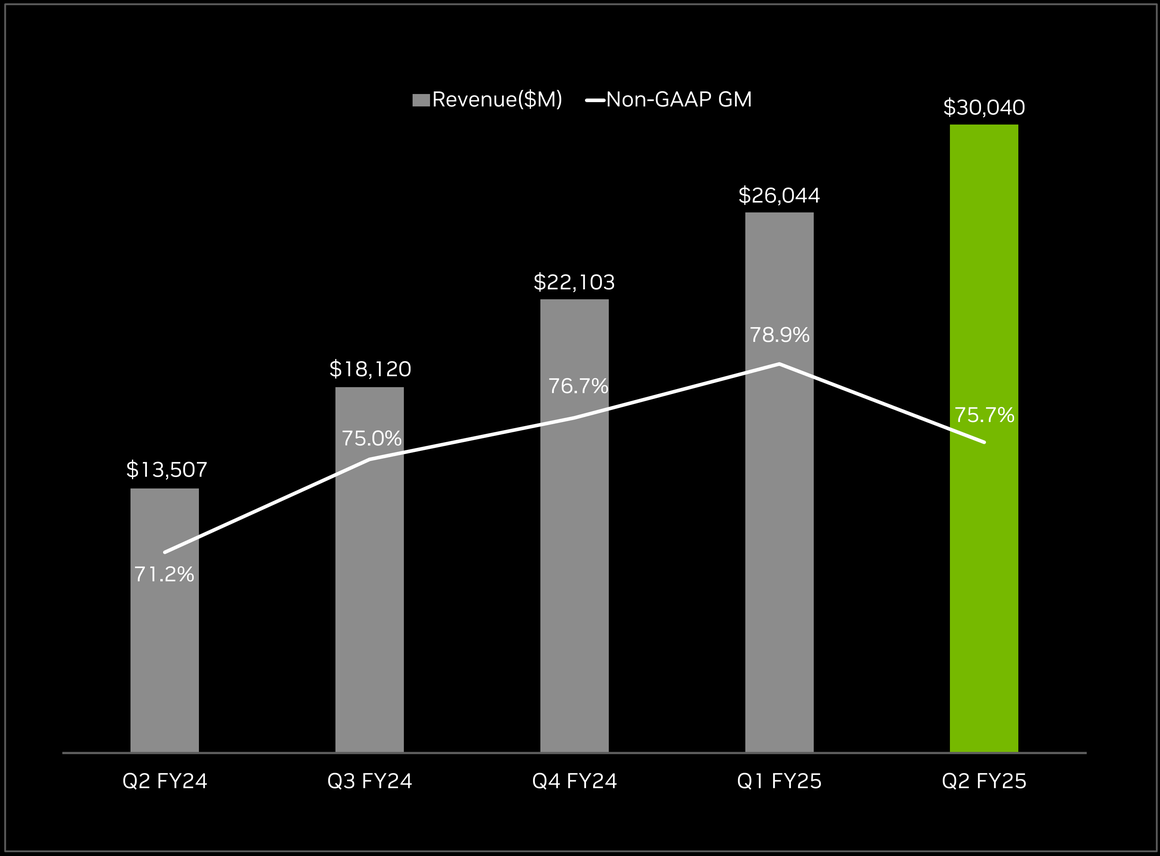
- Customers continue to accelerate their Hopper architecture purchases while gearing up to adopt Blackwell
- Demand coming from frontier model makers, consumer internet services, and tens of thousands of companies building gen AI applications
- Ethernet for AI revenue doubled Q/Q with hundreds of customers adopting our Ethernet offerings
- Sovereign AI revenue can reach low double-digit $ billions this year

Gaming end demand is strong and growing while channel inventory remains healthy

- Every PC with RTX is an AI PC, capable of delivering up to1,300 AI TOPS; now 200+ RTX AI laptop designs from leading PC manufacturers
- Game developers continue to embrace NVIDIA ACE, with Mecha BREAK the first announced game to use the technology
- RTX and gen AI set to revolutionize consumer experiences with 600 AI-powered applications & games, & an installed base of 100M devices

Announced an additional $50.0 billion to our share repurchase authorization

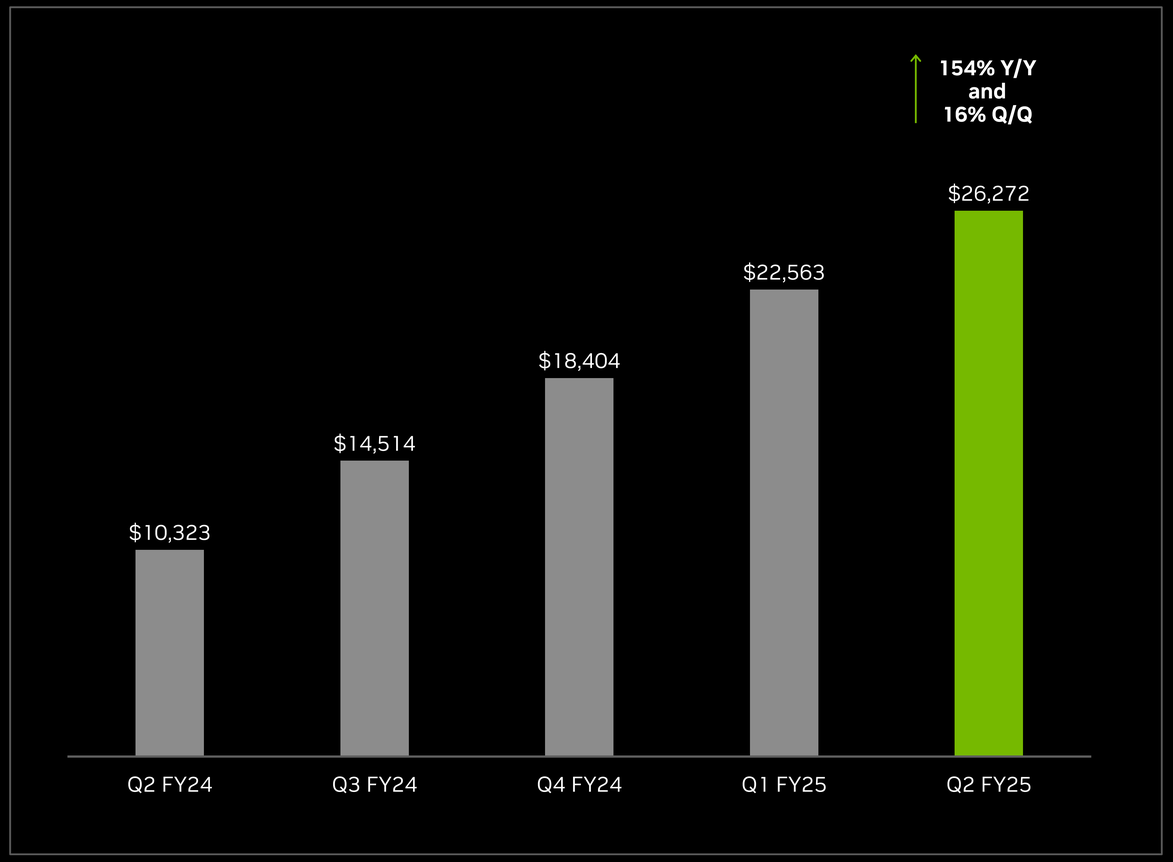
Q2 FY25 Financial Summary



	GAAP			Non-GAAP		
	Q2 FY25	Y/Y	Q/Q	Q2 FY25	Y/Y	Q/Q
Revenue	$30,040	+122%	+15%	$30,040	+122%	+15%
Gross Margin	75.1%	+5.0 pts	-3.3 pts	75.7%	+4.5 pts	-3.2 pts
Operating Income	$18,642	+174%	+10%	$19,937	+156%	+10%
Net Income	$16,599	+168%	+12%	$16,952	+152%	+11%
Diluted EPS	$0.67	+168%	+12%	$0.68	+152%	+11%
Cash Flow from Ops	$14,489	+128%	-6%	$14,489	+128%	-6%

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

Data Center

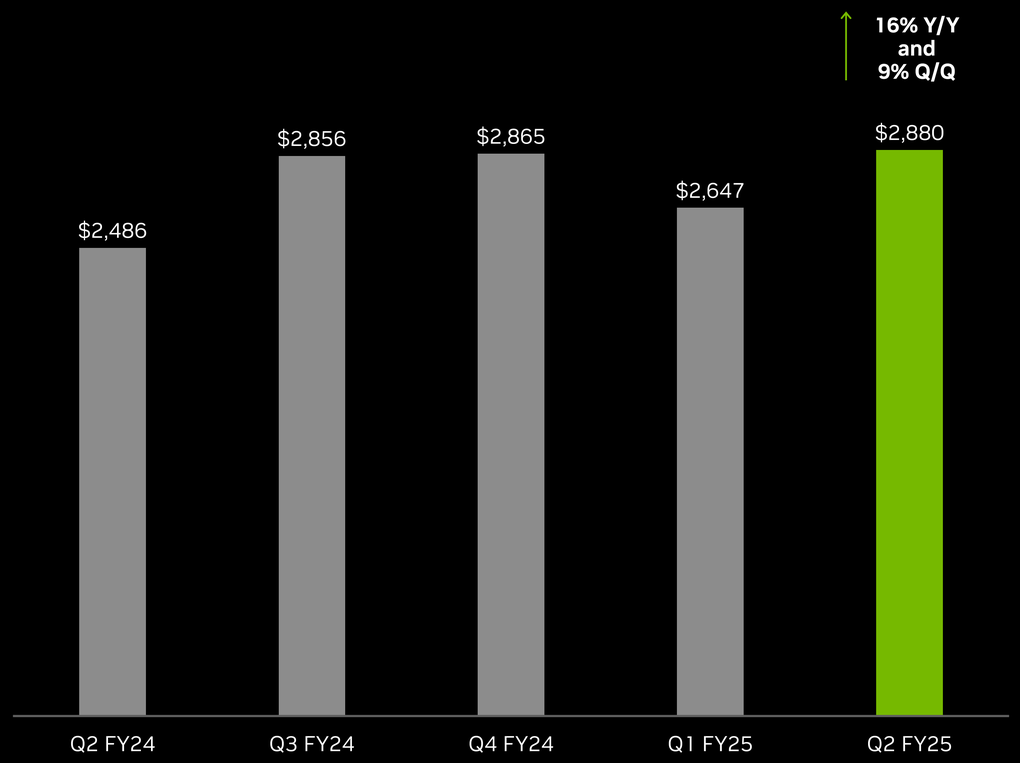


Revenue ($M)

Highlights

- Compute revenue grew more than 2.5x and Networking revenue grew more than 2x from last year
- Key workloads driving growth include gen AI model training and inference; video, image, and text data pre and post processing with CUDA in AI workflows; synthetic data generation, AI powered recommender systems, SQL and vector data base processing
- NVIDIA H200 platform began ramping in Q2 shipping to large CSPs, consumer internet, and enterprise customers
- Blackwell is widely sampling, and production ramp is scheduled to begin in Q4 and continue into F2026; expect to ship several $B in Blackwell revenue in Q4
- Hopper shipments expected to increase in the 2H of fiscal 2025
- Hopper supply/availability have improved. Demand for Blackwell is well above supply, and we expect this to continue into next year
- Networking Y/Y growth was driven by InfiniBand and Ethernet for AI revenue, which includes Spectrum-X end-to-end ethernet platform

Gaming

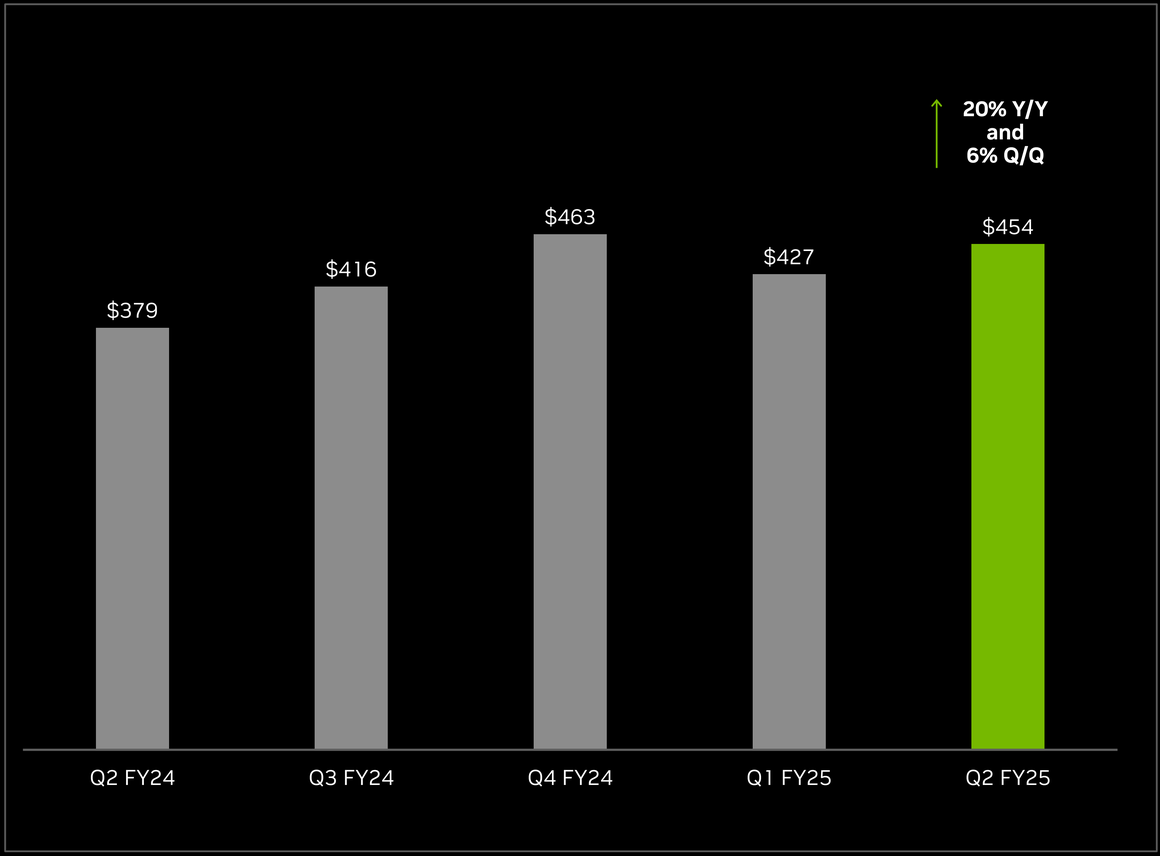


Revenue ($M)

Highlights

- Q/Q growth in console, notebook, and desktop revenue
- Every PC with RTX is an AI PC, and can deliver up to 1,300 AI TOPS; there are now 200+ RTX AI laptop designs from PC manufacturers
- Recently added RTX and DLSS titles include *Indiana Jones and the Great Circle*, *Dune Awakening* and *Dragon Age: The Veilguard*
- GeForce NOW library continues to expand with now over 2,000 titles, the most content of any cloud gaming service

Professional Visualization

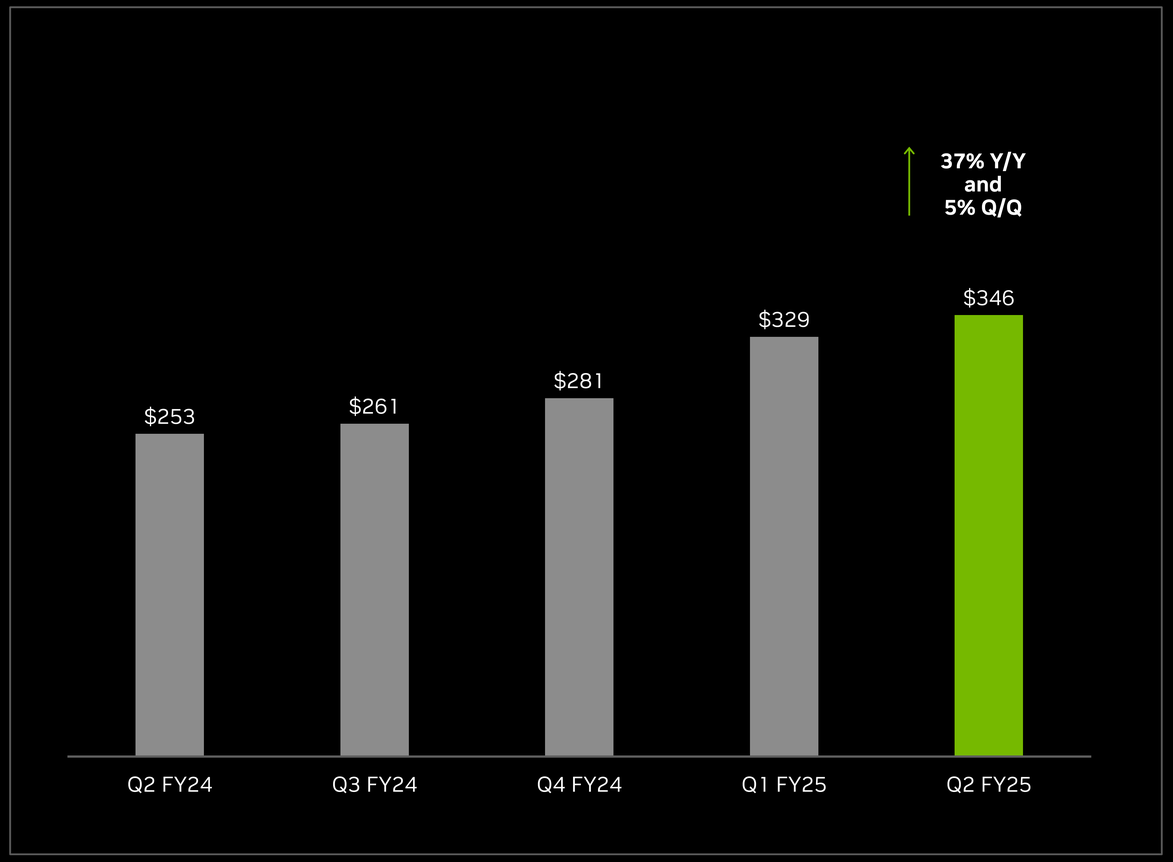


Revenue ($M)

Highlights

- Demand is being driven by AI and graphics use cases, including model fine-tuning and NVIDIA Omniverse-related workloads
- Automotive and manufacturing were among the key industry verticals driving growth
- Foxconn is using Omniverse to power digital twins of the physical plants that produce NVIDIA Blackwell systems
- Several large global enterprises, including Mercedes-Benz, signed multi-year contracts for NVIDIA Omniverse Cloud to build industrial digital twins of factories

Automotive

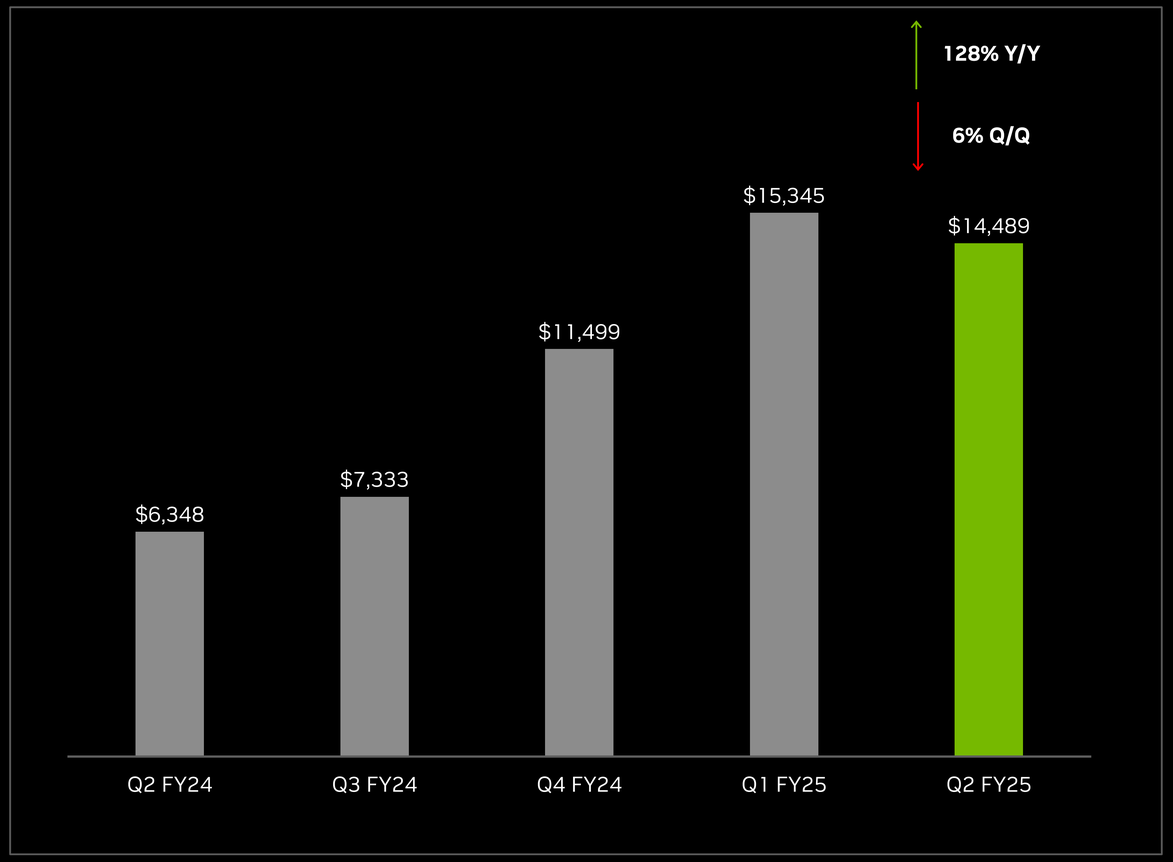


Revenue ($M)

Highlights

- Y/Y growth was driven by new customer ramps in self-driving platforms and increased demand for AI cockpit solutions
- At the Computer Vision and Pattern Recognition conference, NVIDIA won the Autonomous Grand Challenge in the ‘End-to-End Driving at Scale’ category, outperforming more than 400 entries worldwide
- Boston Dynamics, BYD Electronics, Figure, Intrinsic, Siemens, Skild AI, and Teradyne Robotics are using the NVIDIA Isaac robotics platform for autonomous robot arms, humanoids, and mobile robots

Sources & Uses of Cash



Cash Flow from Operations ($M)

Highlights

- Y/Y increase reflects higher revenue partially offset by higher cash taxes paid
- Q/Q decrease reflects cash taxes paid partially offset by higher revenue
- Utilized cash of $7.4B towards shareholder returns, including $7.2B in share repurchases and $246M in cash dividends
- Invested $1B in capex (includes principal payments on PP&E)
- Ended the quarter with $34.8B in gross cash and $8.5B in debt; $26.3B 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 FY25 Outlook

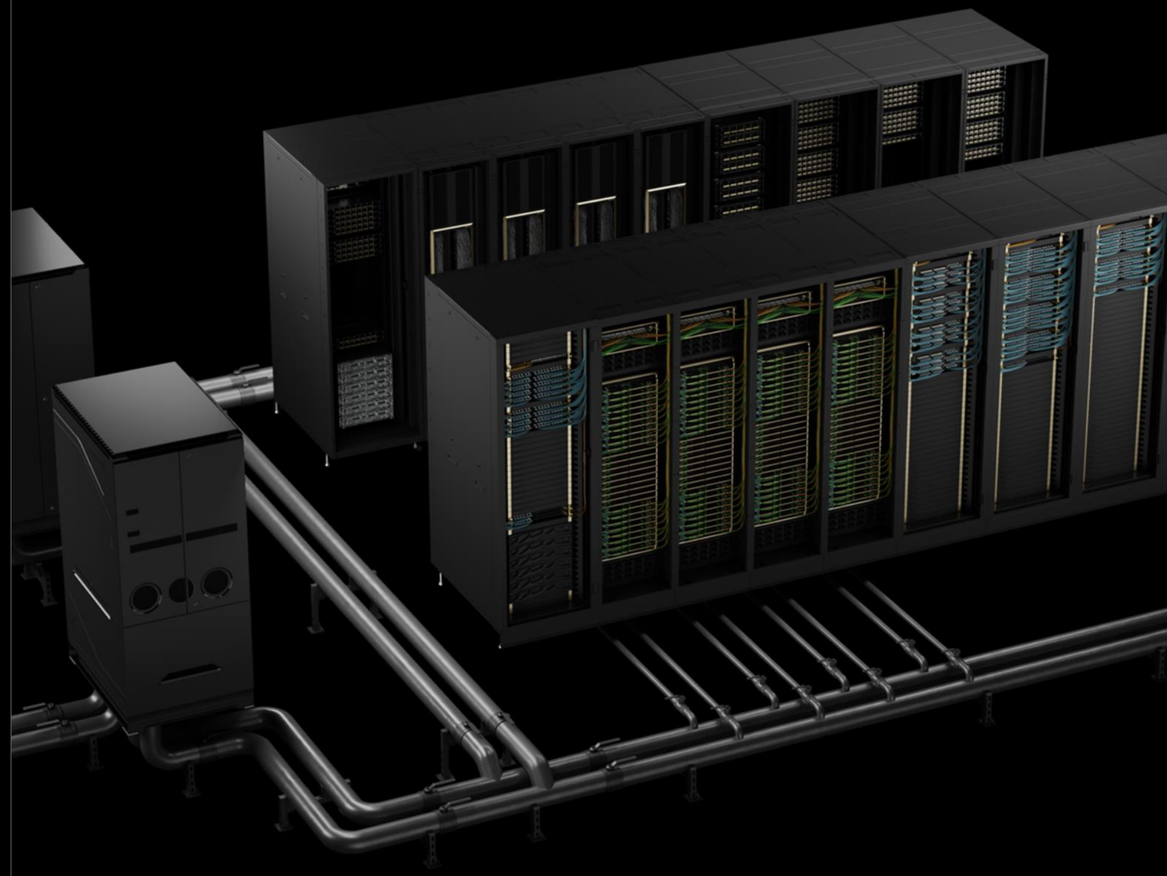
Revenue	**$32.5 billion**, plus or minus 2% Incorporates continued Hopper architecture growth and sampling of Blackwell products
Gross Margins	**74.4%** GAAP and **75.0%** non-GAAP, plus or minus 50 basis points For the full year, gross margins are expected to be in the mid-70% range
Operating Expense	Approximately **$4.3 billion** GAAP and **$3.0 billion** non-GAAP Full-year opex is expected to grow in the mid to upper 40% range
Other Income & Expense	Income of approximately **$350 million** for GAAP and non-GAAP Excluding gains and losses from non-affiliated investments and publicly-held equity securities
Tax Rate	**17.0%** GAAP and non-GAAP, plus or minus 1%, excluding discrete items

Refer to Appendix for reconciliation of Non-GAAP measures.

Key Announcements This Quarter

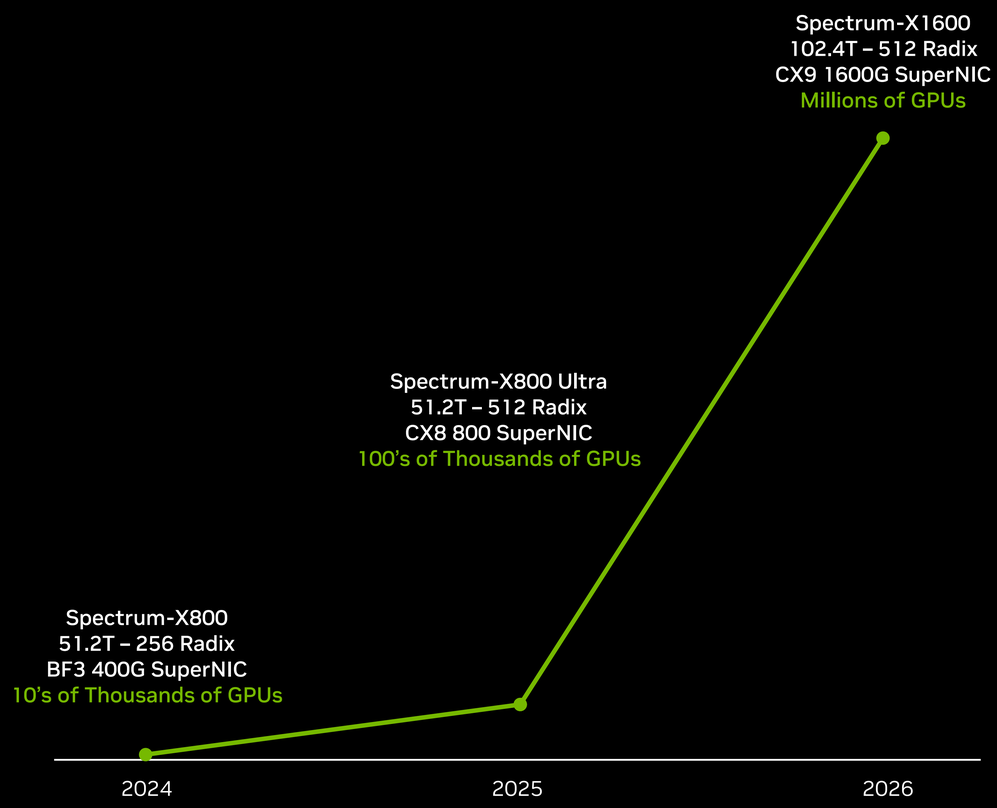
Computer Industry Joins NVIDIA to Build AI Factories and Data Centers for the Next Industrial Revolution

- NVIDIA and the world's top computer manufacturers unveiled an array of NVIDIA Blackwell architecture-powered systems and NVIDIA networking for building AI factories and data centers
- These include cloud, on-premises, embedded and edge AI systems ranging from single to multi-GPUs, x86- to Grace-based processors, and air- to liquid-cooling technology
- **NVIDIA MGX Modular Reference Architecture in Over 100 Blackwell Systems –** NVIDIA MGX now supports Blackwell products, including the new GB200 NVL2 for scale-out servers. NVIDIA MGX provides computer manufacturers with a reference architecture to quickly and cost-effectively build Blackwell system design configurations



Widespread Adoption of NVIDIA Spectrum-X Ethernet for AI

- NVIDIA Spectrum-X Ethernet networking platform built for AI is now generally available from OEM & ODM partners; gaining wide adoption by Cloud Service Providers, GPU Cloud Providers and Enterprises
- CoreWeave, GMO Internet Group, Lambda, Scaleway, STPX Global and Yotta are among the first AI cloud service providers embracing it to bring extreme networking performance to their AI infrastructures
- Announced plans to launch new Spectrum-X products every year, delivering increased bandwidth, ports, enhanced software feature sets and programmability to drive leading AI Ethernet networking performance



SPECTRUM-X ROADMAP
ETHERNET FOR AI

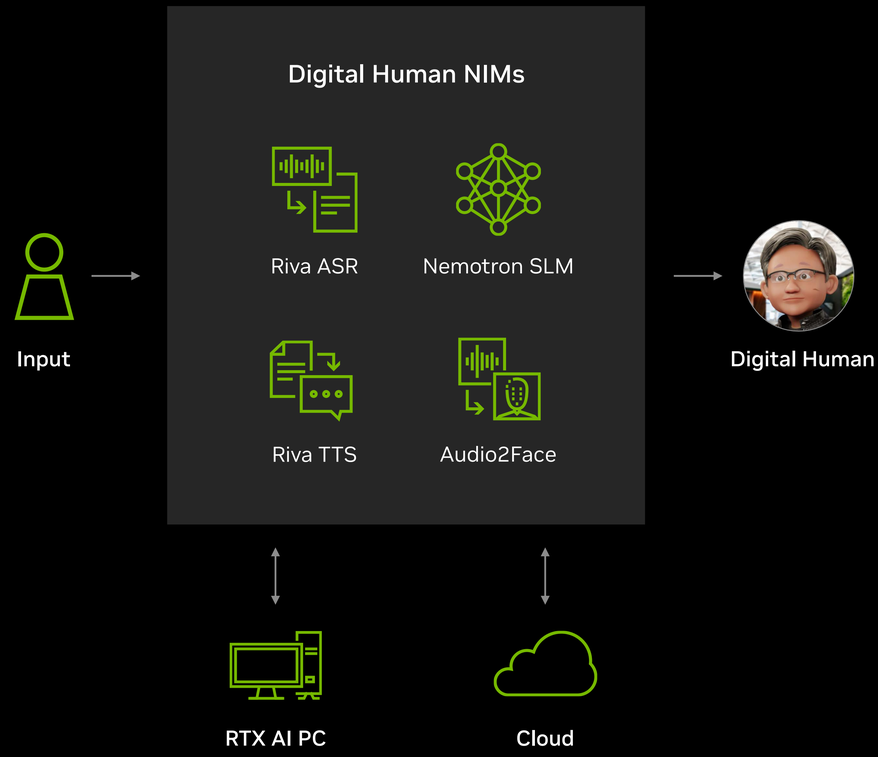
NVIDIA NIM Inference Microservices Now Generally Available, Adopted by 150+ Partners

- NVIDIA NIM are inference microservices that provide pre-trained models as optimized containers to deploy on clouds, data centers or workstations, enabling developers to quickly and easily build generative AI applications
- NIM enables enterprises to maximize their infrastructure investments. For example, running Meta Llama 3-8B in a NIM produces up to 3x more generative AI tokens on accelerated infrastructure than without NIM
- Many technology partners, including Cadence, Cloudera, Cohesity, DataStax, NetApp, Scale AI and Synopsys, are integrating NIM into their platforms to speed generative AI deployments for domain-specific applications, such as copilots, code assistants, and digital human avatars
- Industry leaders Foxconn, Pegatron, Amdocs, Lowe's, ServiceNow and Siemens are among the businesses using NIM for generative AI applications in manufacturing, healthcare, financial services, retail, customer service and more
- Enterprises can deploy AI applications in production with NIM through the NVIDIA AI Enterprise software platform



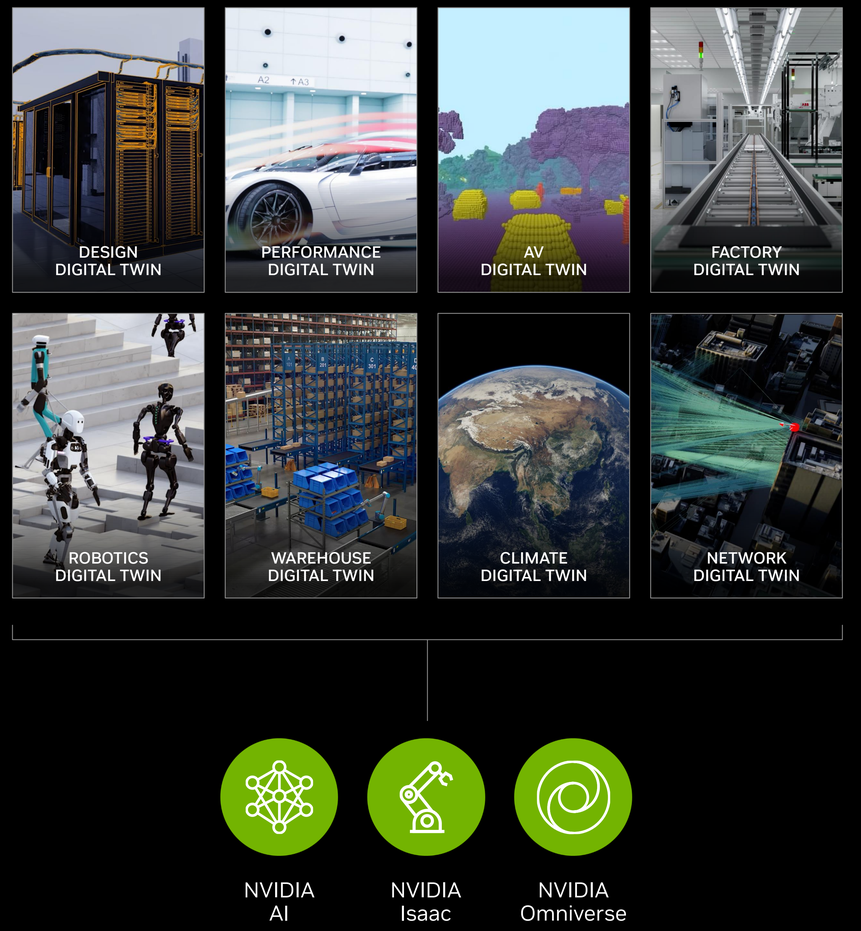
NVIDIA ACE Microservices for Future of Generative AI Avatars Now Available

- NVIDIA ACE generative AI microservices, which help enable the creation and operation of digital humans or generative AI avatars, are now generally available for cloud, and in early access for RTX AI PCs
- NVIDIA ACE digital human technologies now generally available include:
  - NVIDIA Riva ASR, TTS and NMT — for automatic speech recognition, text-to-speech conversion and translation
  - NVIDIA Nemotron LLM — for language understanding and contextual response generation
  - NVIDIA Audio2Face — for realistic facial animation based on audio tracks
  - NVIDIA Omniverse RTX — for real-time, path-traced realistic skin and hair
- Newly announced technologies include:
  - NVIDIA Audio2Gesture — for generating body gestures based on audio tracks
  - NVIDIA Nemotron-3 4.5B — a new small language model for on-device RTX AI PC inference
- ACE technologies are in use by companies across customer service, gaming and healthcare, including Dell Technologies, ServiceNow, Aww Inc., Inventec, and Perfect World Games



Physical AI Is the Next Wave, Driving the Digitalization of $50T of Heavy Industries

- The next wave of AI is physical AI – or AI that understands the laws of physics – which is essential for robotics and industrial digitalization. NVIDIA Omniverse and NVIDIA Isaac are key platforms driving this transition
- Major Taiwanese electronics makers, including Delta Electronics, Foxconn, Pegatron and Wistron, are using NVIDIA technologies such as Omniverse, Metropolis, and Isaac to automate their factories
- Robotics leaders are adopting the NVIDIA Isaac robotics platform for the research, development and production of the next generation of AI-enabled autonomous machines such as robot arms, humanoids, and mobile robots
  - This includes BYD Electronics, Siemens, Teradyne Robotics and Intrinsic, an Alphabet Company
- NVIDIA AI Enterprise-IGX with Holoscan now generally available for Industrial Edge AI, providing an enterprise-grade platform for the medical, industrial and scientific computing sectors to develop and deploy edge AI solutions faster



NVIDIA Brings AI Assistants to Life with RTX AI PCs

- NVIDIA RTX AI PCs, with over 200 RTX AI laptops and more than 600 AI-powered apps and games, are set to revolutionize consumer experiences with generative AI
- New laptops from our partners feature up to GeForce RTX 4070 GPUs with Windows 11 AI PC capabilities and will receive a free update to Copilot+ PC experiences when available
- The RTX AI Toolkit is a new suite of tools and software development kits that aid developers in optimizing and deploying large generative AI models on Windows PCs
- Microsoft and NVIDIA are collaborating to help developers bring new generative AI capabilities to their Windows native and web apps with easy API access to RTX-accelerated small language models that run on-device as part of Windows Copilot Runtime

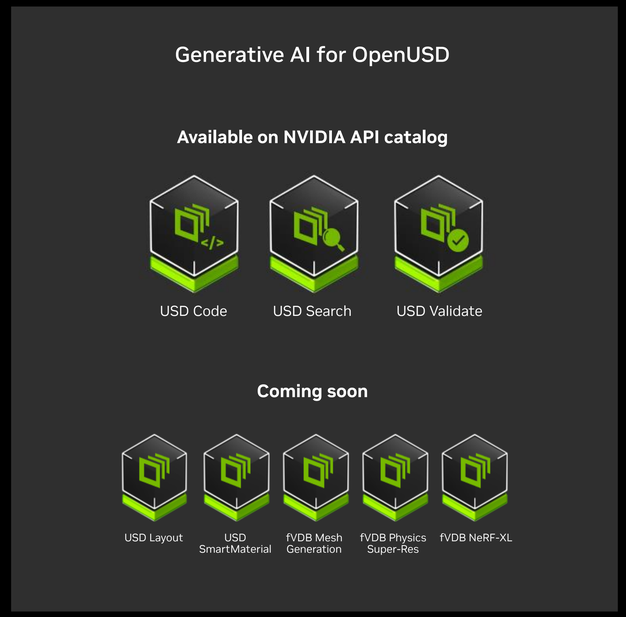


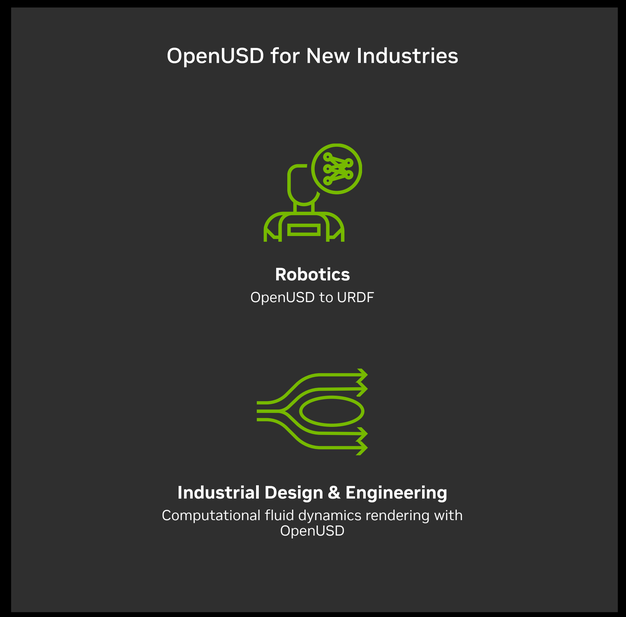
ANNOUNCING NEW RTX AI PCs

Now Over 200 RTX AI Laptops | Up to 700 AI TOPS | 7X Generative AI

NVIDIA Announces Generative AI Models and NIM Microservices for OpenUSD Language , Geometry, Physics and Materials

- Announced major advancements to Universal Scene Description (OpenUSD), that will expand adoption of the universal 3D data interchange framework to robotics, industrial design and engineering, and accelerate developers' abilities to build highly accurate virtual worlds
- The world's first generative AI models for OpenUSD, developed by NVIDIA, will be available as NVIDIA NIM microservices, enabling developers to incorporate gen AI copilots/agents into USD workflows
- NIM microservices and Omniverse are being used by Foxconn to create a digital twin of a factory under development
- WPP is implementing USD NIM microservices in its generative AI-enabled content creation pipeline for customers such as The Coca-Cola Company
- A series of new USD connectors opens OpenUSD to more industries. Siemens will integrate OpenUSD pipelines with its portfolio of simulation technologies
- NVIDIA also released a connector from Unified Robotics Description Format to OpenUSD, letting roboticists seamlessly bring their robot data across applications

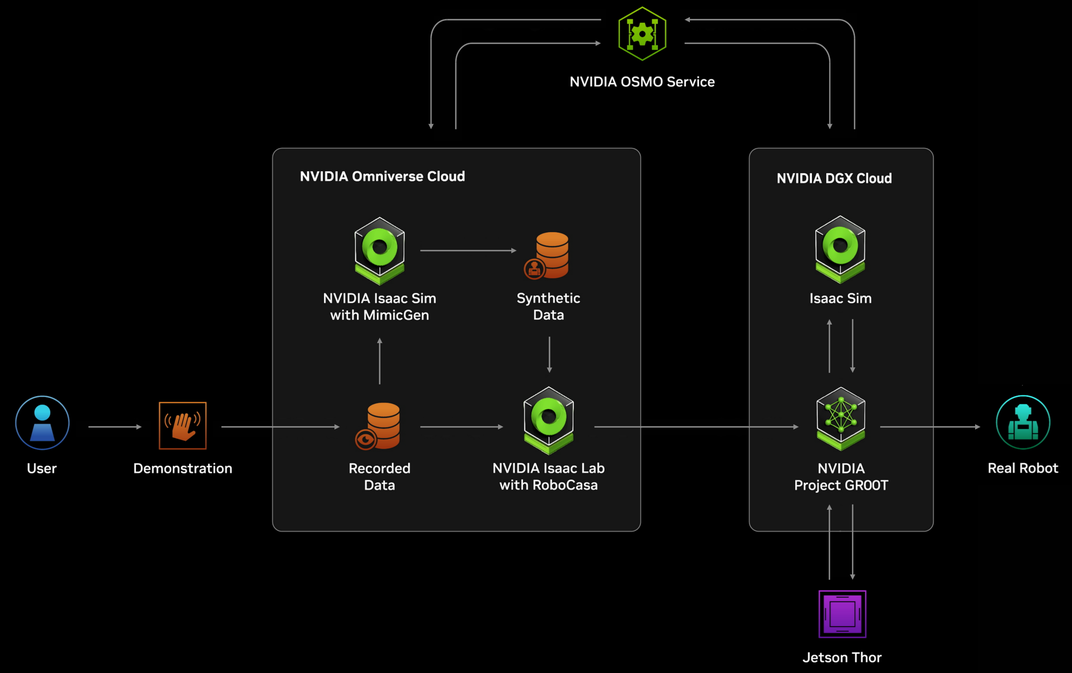




NVIDIA Accelerates Humanoid Robotics Development with New Offerings

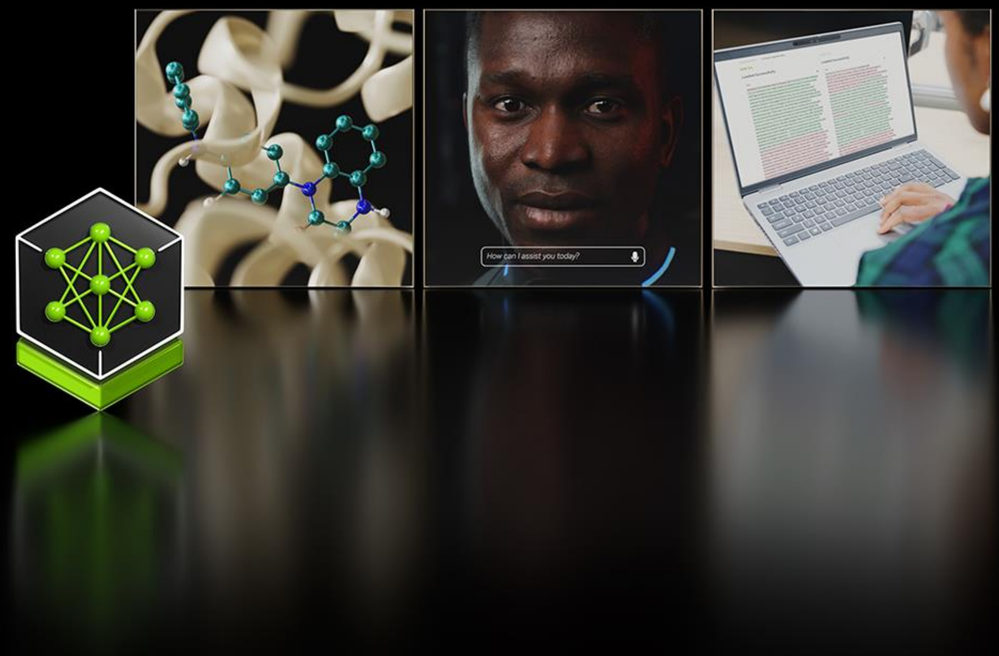
- Announced new offerings to accelerate humanoid development at a global scale including:
  - New NVIDIA NIM microservices and frameworks for robot simulation and learning
  - NVIDIA OSMO orchestration service for running multi-stage robotics workloads, cutting deployment and development cycle times from months to under a week
  - An AI- & simulation-enabled teleoperation workflow that allows developers to train robots using small amounts of human demonstration data

- The new AI microservices will allow roboticists to enhance simulation workflows for generative physical AI
  - MimicGen NIM generates synthetic motion data based on recorded teleoperated data from spatial computing devices
  - Robocasa NIM generates robot tasks and simulation-ready environments in OpenUSD



NVIDIA and Global Partners Launch NIM Agent Blueprints for Enterprises to Make Their Own AI

- NVIDIA NIM Agent Blueprint is a catalog of pretrained, customizable AI workflows that equip millions of enterprise developers with a full suite of software for building and deploying gen AI applications
- The catalog of customizable workflows speeds deployments of core gen AI use cases, starting with customer service, drug discovery and data extraction for PDFs, with more to come
- With NIM Agent Blueprints, enterprises can continually refine their AI applications based on user feedback, creating a data-driven AI flywheel
- Accenture, Deloitte, SoftServe and World Wide Technology are bringing NVIDIA NIM Agent Blueprints to enterprises worldwide. Cisco, Dell Technologies, Hewlett Packard Enterprise and Lenovo are offering NVIDIA-accelerated infrastructure and solutions to speed NIM Agent Blueprints deployments



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
Q2 FY25						
Gross margin ($ in million)	$22,729	(118)	(40)	3	—	$22,574
	75.7%	(0.4)	(0.2)	—	—	75.1%
Operating income ($ in million)	$19,937	(144)	(1,154)	3	—	$18,642
Net income ($ in million)	$16,952	(144)	(1,154)	195	750	$16,599
Shares used in diluted per share calculation (millions)	24,848	—	—	—	—	24,848
Diluted EPS	$0.68	—	—	—	—	$0.67

A. Consists of amortization of intangible assets and transaction costs.
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 IP-related costs, gains from non-affiliated investments 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)	Other (C)	GAAP
Q2 FY 2024	71.2%	(0.9)	(0.2)	—	70.1%
Q3 FY 2024	75.0%	(0.7)	(0.2)	(0.1)	74.0%
Q4 FY 2024	76.7%	(0.5)	(0.2)	—	76.0%
Q1 FY 2025	78.9%	(0.4)	(0.1)	—	78.4%

A. Consists of amortization of intangible assets.
B. Stock-based compensation charge was allocated to cost of goods sold.
C. Other consists of IP-related costs.

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

($ in Millions)	Q3 FY25 Outlook
Non-GAAP gross margin	75.0%
Impact of stock-based compensation expense, acquisition-related costs, and other costs	(0.6%)
GAAP gross margin	74.4%
Non-GAAP operating expenses	$3,000
Impact of stock-based compensation expense, acquisition-related costs, and other costs	1,250
GAAP operating expenses	$4,250


NVIDIA