Company Overview

August 29, 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 markets, market opportunity, demand and growth drivers; our financial outlook; the benefits, impact, performance, features and availability of our products and technologies; the benefits, impact, features and timing of our collaborations or partnerships; third parties adopting our products and technologies; NVIDIA accelerated computing being broadly recognized as the way to advance computing as Moore's law ends and AI lifts off; accelerated computing being needed to tackle the most impactful opportunities of our time; AI driving a platform shift from general purpose to accelerated computing, and enabling new, never-before-possible applications; trillion dollars of installed global data center infrastructure transitioning to accelerated computing; broader enterprise adoption of AI and accelerated computing under way; AI and accelerated computing making possible the next big waves of autonomous machines and industrial digitalization; a rapidly growing universe of applications and industry innovation; the ability of developers to engage with NVIDIA through CUDA; AI augmenting creativity and productivity by orders of magnitude across industries ; generative AI as the most important computing platform of our generation; data centers becoming AI factories; large language models being one of today's most important advanced AI technologies, involving up to trillions of parameters that learn from text; full-stack and data center scale acceleration driving significant cost savings and workload scaling; the high ROI of high compute performance; NVIDIA powering the AI industrial revolution; the ability of developers to connect additional or third party services to the AI chatbot via cloud AI APIs; AI factories acting as trusted engines of generative AI; features of AI factories; nations using AI factories as sovereign national resources to process private datasets of companies, startups, universities and governments safely on shore to produce valuable insights; every important company running its own AI factories; NVIDIA generating recurring revenue from AI factories for their use of NVIDIA AI Enterprise, the operating system for enterprise AI, in addition to the up-front revenue opportunity from data center systems; our dividend program plan; our strategic investments; NVIDIA on track to source 100% renewable electricity for offices and data centers under operational control by end of FY25; and NVIDIA's on track to engage manufacturing suppliers comprising at least 67% of scope 3 category 1 GHG emissions with the goal of effecting supplier adoption of science-based targets by end of FY26 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 profit, non-GAAP gross margin, non-GAAP operating income, non-GAAP operating margin, and free cash flow. 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."

Headquarters: Santa Clara, CA

NVIDIA pioneered accelerated computing to help solve impactful challenges classical computers cannot. A quarter of a century in the making, NVIDIA accelerated computing is broadly recognized as the way to advance computing as Moore's law ends and AI lifts off.

NVIDIA's platform is installed in several hundred million computers, is available in every cloud and from every server maker, powers over 75% of the TOP500 supercomputers, and has over 5 million developers.

Headquarters: Santa Clara, CA
Headcount: ~29,600

NVIDIA's Accelerated Computing Platform

Full-stack innovation across silicon, systems and software

With nearly three decades of singular focus, NVIDIA is expert at accelerating software and scaling compute by a **Million-X**, going well beyond Moore's law

Accelerated computing requires **full-stack** innovation—optimizing across every layer of computing—from silicon and systems to software and algorithms, demanding deep understanding of the problem domain

Our full-stack platforms—NVIDIA AI and NVIDIA Omniverse—accelerate AI and industrial digitalization workloads

We accelerate workloads at **data center scale**, across thousands of compute nodes, treating the network and storage as part of the computing fabric

Our platform extends from the cloud and enterprise data centers to supercomputing centers, edge computing and PCs

What Is Accelerated Computing?

A full-stack approach: silicon, systems, software

Not just a superfast chip—accelerated computing is a full-stack combination of:

- Chip(s) with specialized processors
- Algorithms in acceleration libraries
- Domain experts to refactor applications

To speed-up compute-intensive parts of an application

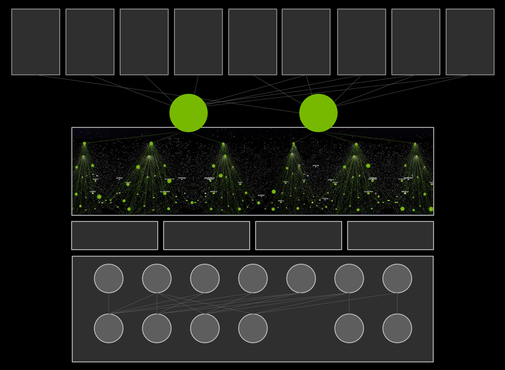

Amdahl's law:

The overall system speed-up (S) gained by optimizing a single part of a system by a factor (s) is limited by the proportion of execution time of that part (p).

$$S = \frac{1}{(1-p) + \frac{p}{s}}$$

For example:

- If 90% of the runtime can be accelerated by 100X, the application is sped up 9X
- If 99% of the runtime can be accelerated by 100X, the application is sped up 50X
- If 80% of the runtime can be accelerated by 500X, or even 1000X, the application is sped up 5X

Why Accelerated Computing?

Advancing computing in the post-Moore's Law era

Accelerated computing is needed to tackle the most impactful opportunities of our time—like AI, climate simulation, drug discovery, ray tracing, and robotics

NVIDIA is uniquely dedicated to accelerated computing —working top-to-bottom, refactoring applications and creating new algorithms, and bottom-to-top—inventing new specialized processors, like RT Core and Tensor Core

"It's the end of Moore's Law as we know it."

—John Hennessy Oct 23, 2018

"Moore's Law is dead."

—Jensen Huang, GTC 2013

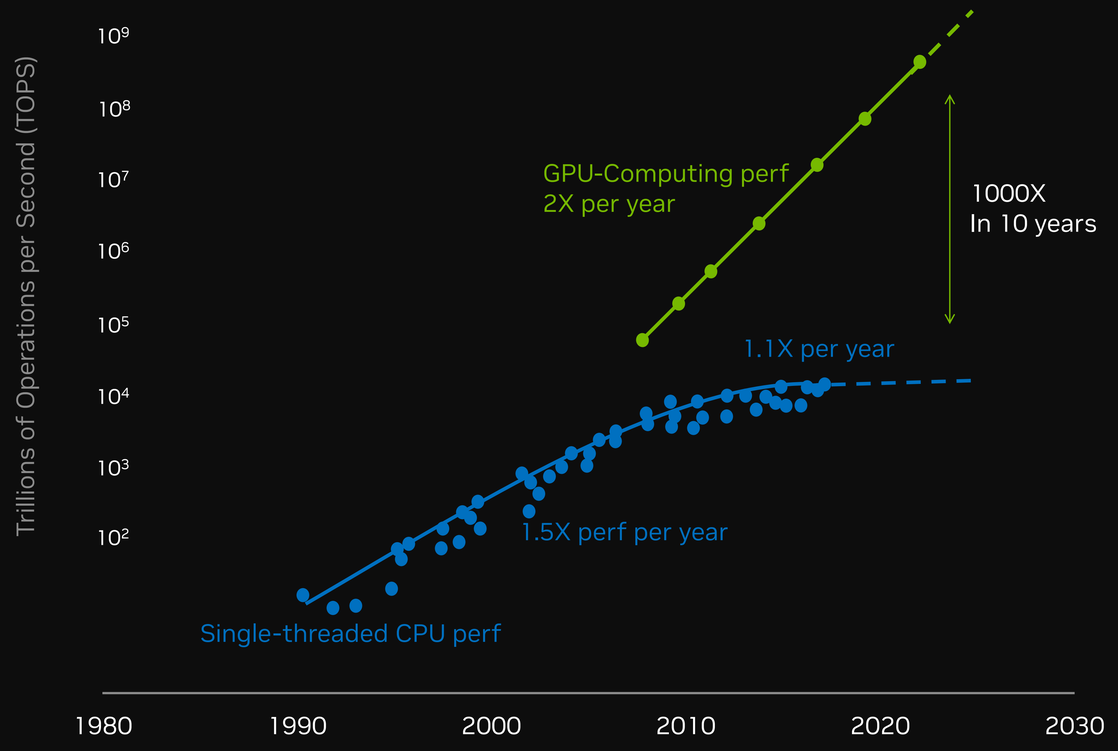

1000x AI Compute in 8 Years

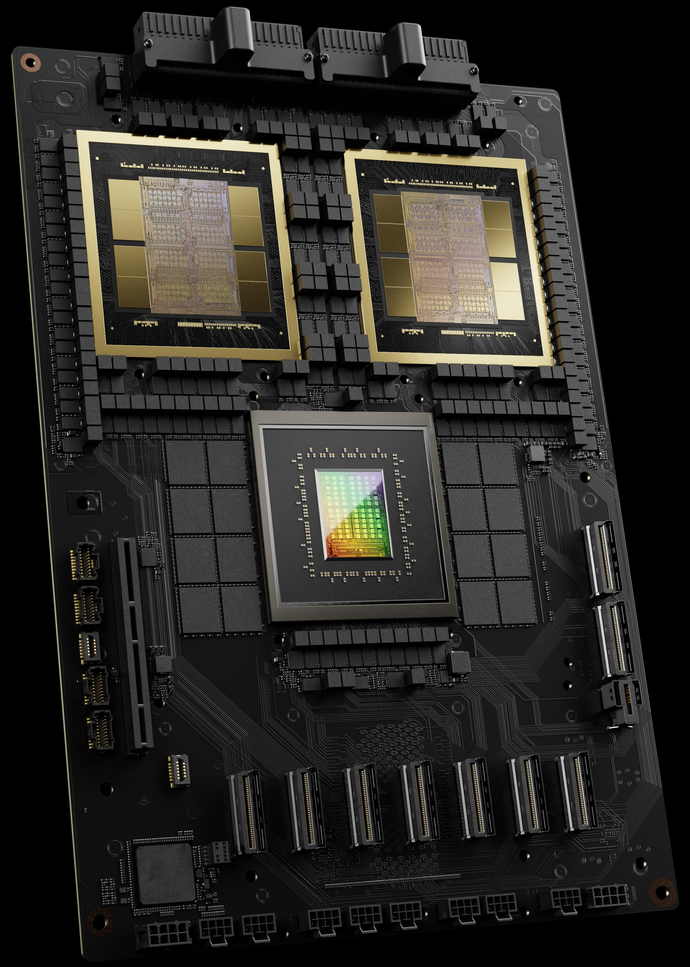

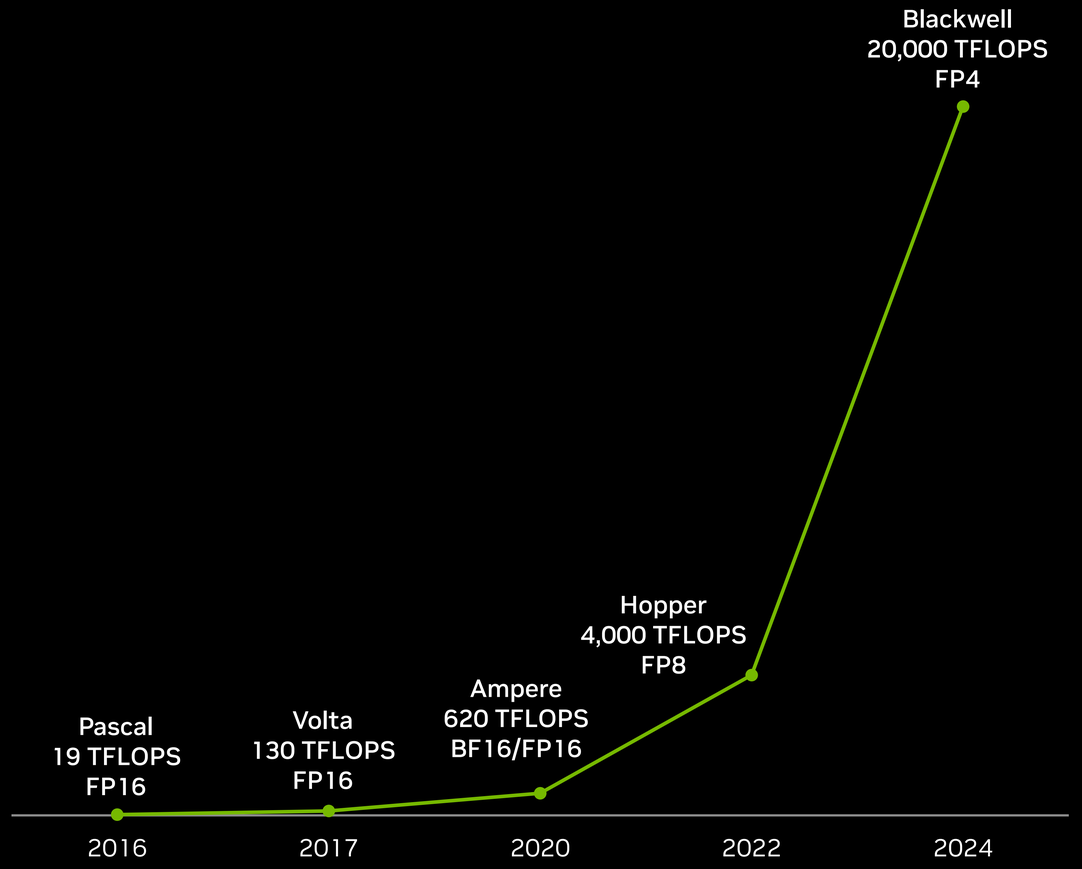

Waves of Adoption of Accelerated Computing

A generational computing platform shift

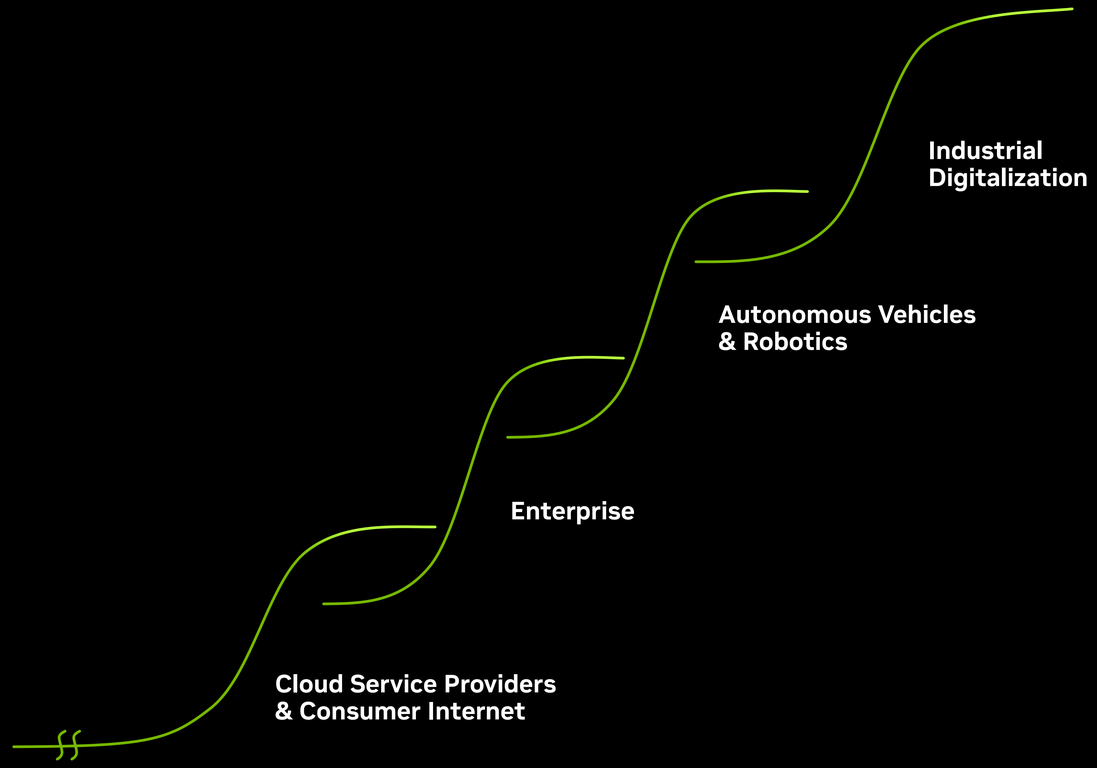

A new computing era has begun

Accelerated computing enabled the rise of AI, which is driving a platform shift from general purpose to accelerated computing, and enabling new, never-before-possible applications

The trillion dollars of installed global data center infrastructure will transition to accelerated computing to achieve better performance, energy-efficiency and cost by an order of magnitude

Hyperscale cloud service providers and consumer internet companies have been the early adopters of AI and accelerated computing, with broader enterprise adoption now under way

AI and accelerated computing will also make possible the next big waves—autonomous machines and industrial digitalization

NVIDIA Accelerated Computing for Every Wave

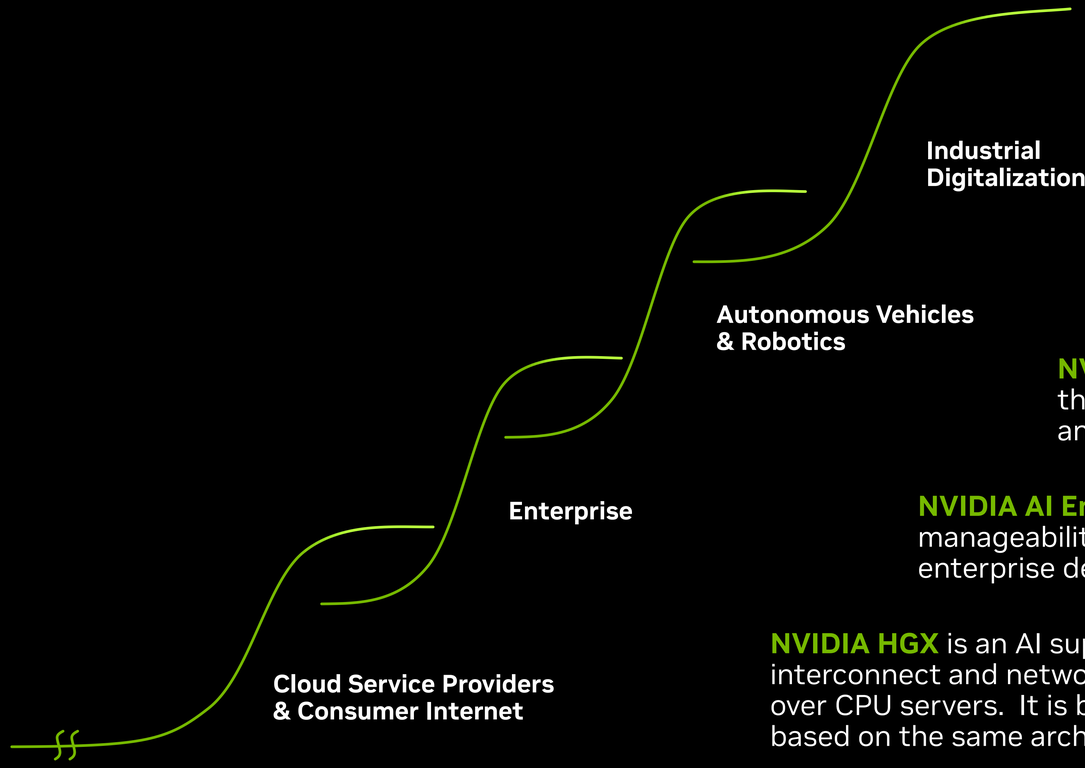

NVIDIA Omniverse is a software platform for designing, building, and operating 3D and virtual world simulations. It harnesses the power of NVIDIA graphics and AI technologies and runs on NVIDIA-powered data centers and workstations.

NVIDIA DRIVE is a full-stack platform for autonomous vehicles (AV) that includes hardware for in-car compute, such as the Orin system-on-chip, and the full AV and AI cockpit software stack.

NVIDIA DGX Cloud is a cloud service that allows enterprises immediate access to the infrastructure and software needed to train advanced models for generative AI and other groundbreaking applications.

NVIDIA AI Enterprise is the operating system of AI, with enterprise-grade security, stability, manageability and support. It is available on all major CSPs and server OEMs and supports enterprise deployment of AI in production.

NVIDIA HGX is an AI supercomputing platform purpose-built for AI. It includes 8 NVIDIA GPUs, as well as interconnect and networking technologies, delivering order-of-magnitude performance speed-ups for AI over CPU servers. It is broadly available from all major server OEMs/ODMs. **NVIDIA DGX**, an AI server based on the same architecture, along with NVIDIA AI software and support, is also available.

NVIDIA's Accelerated Computing Ecosystem

- The NVIDIA accelerated computing platform has attracted the largest ecosystem of developers, supporting a rapidly growing universe of applications and industry innovation
- Developers can engage with NVIDIA through CUDA—our parallel computing programming model introduced in 2006—or at higher layers of the stack, including libraries, pre-trained AI models, SDKs and other development tools

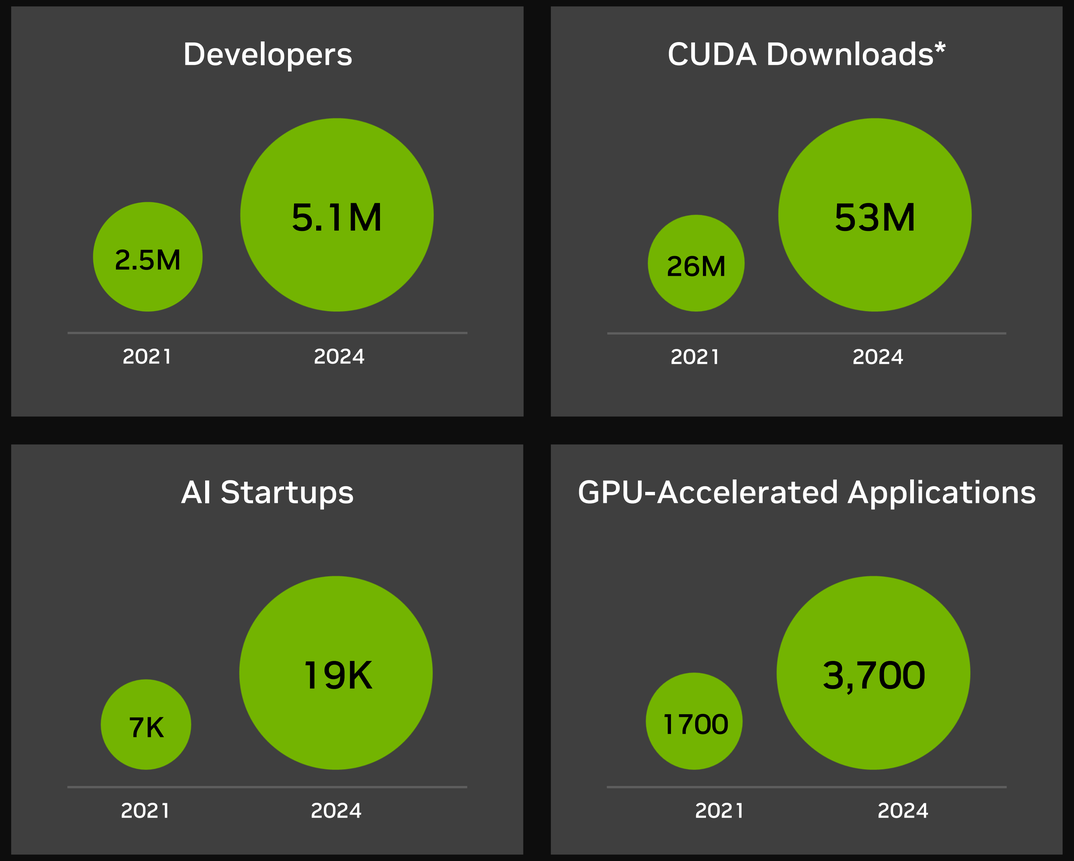

*Cumulative

NVIDIA’s Multi-Sided Platform and Flywheel

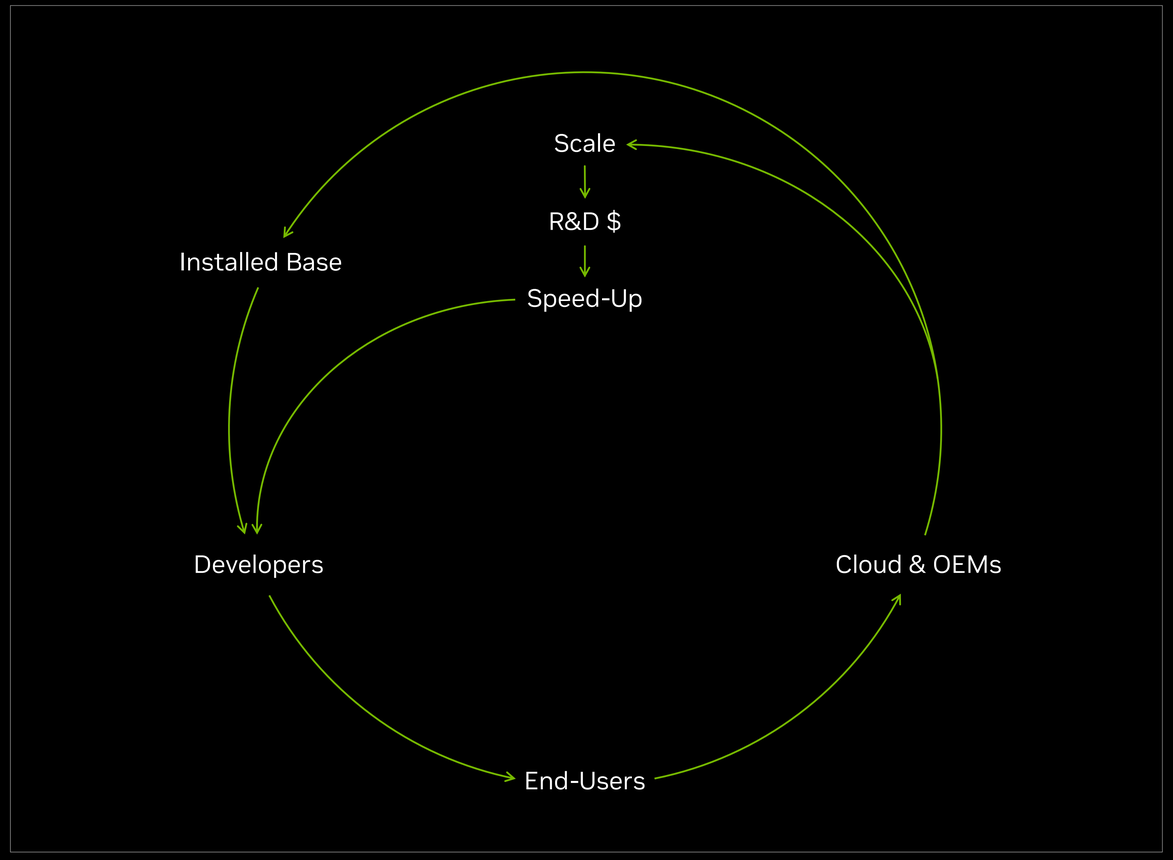

NVIDIA Accelerated Computing Virtuous Cycle

The virtuous cycle of NVIDIA’s accelerated computing starts with an installed base of several hundred million GPUs, all compatible with the CUDA programming model

- **For developers**—NVIDIA’s one architecture and large installed base give developer’s software the best performance and greatest reach
- **For end users**—NVIDIA is offered by virtually every computing provider and accelerates the most impactful applications from cloud to edge
- **For cloud providers and OEMs**—NVIDIA’s rich suite of Acceleration Platforms lets partners build one offering to address large markets including media & entertainment, healthcare, transportation, energy, financial services, manufacturing, retail, and more
- **For NVIDIA**—Deep engagement with developers, computing providers, and customers in diverse industries enables unmatched expertise, scale, and speed of innovation across the entire accelerated computing stack — propelling the flywheel

Huge ROI From AI Driving a Powerful New Investment Cycle

AI can augment creativity and productivity by orders of magnitude across industries

Knowledge workers will use copilots based on large language models to generate documents, answer questions, or summarize missed meetings, emails and chats—adding hours of productivity per week

Copilots specialized for fields such as software development, legal services or education can boost productivity by as much as 50%

Social media, search and e-commerce apps are using deep recommenders to offer more relevant content and ads to their customers, increasing engagement and monetization

Creators can generate stunning, photorealistic images with a single text prompt—compressing workflows that take days or weeks into minutes in industries from advertising to game development

Call center agents augmented with AI chatbots can dramatically increase productivity and customer satisfaction

Drug discovery, financial services, agriculture and food services and climate forecasting are seeing order-of-magnitude workflow acceleration from AI

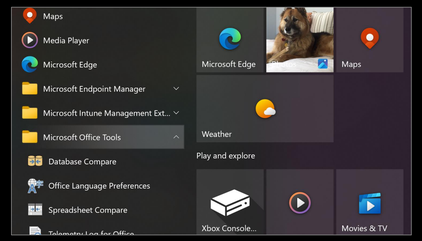

Office AI Copilots
Over 1B knowledge workers

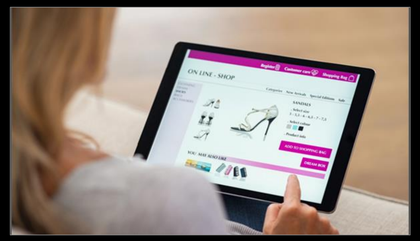

Search & Social Media
$700B in digital advertising annually

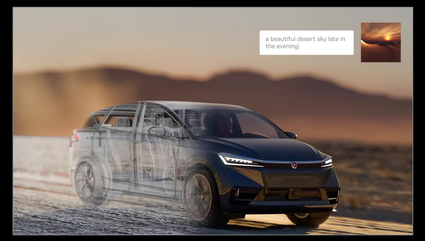

AI Content Creation
50M creators globally

Legal Services, Education
1M legal professionals in the US
9M educators in the US

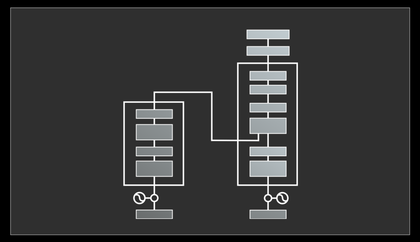

AI Software Development
30M software developers globally

Financial Services
678B annual credit card transactions

Customer Service with AI
15M call center agents globally

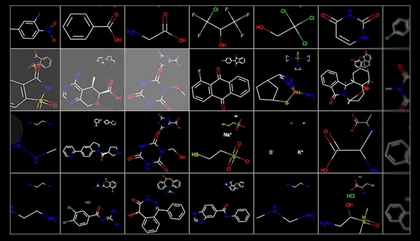

Drug Discovery
10^{18} molecules in chemical space
40 exabytes of genome data

Agri-Food | Climate
1B people in agri-food worldwide
Earth-2 for km-scale simulation

Source: Goldman Sachs, Cowen, Statista, Capital One, Wall Street Journal, Resource Watch, NVIDIA internal analysis

Generative AI

The most important computing platform of our generation

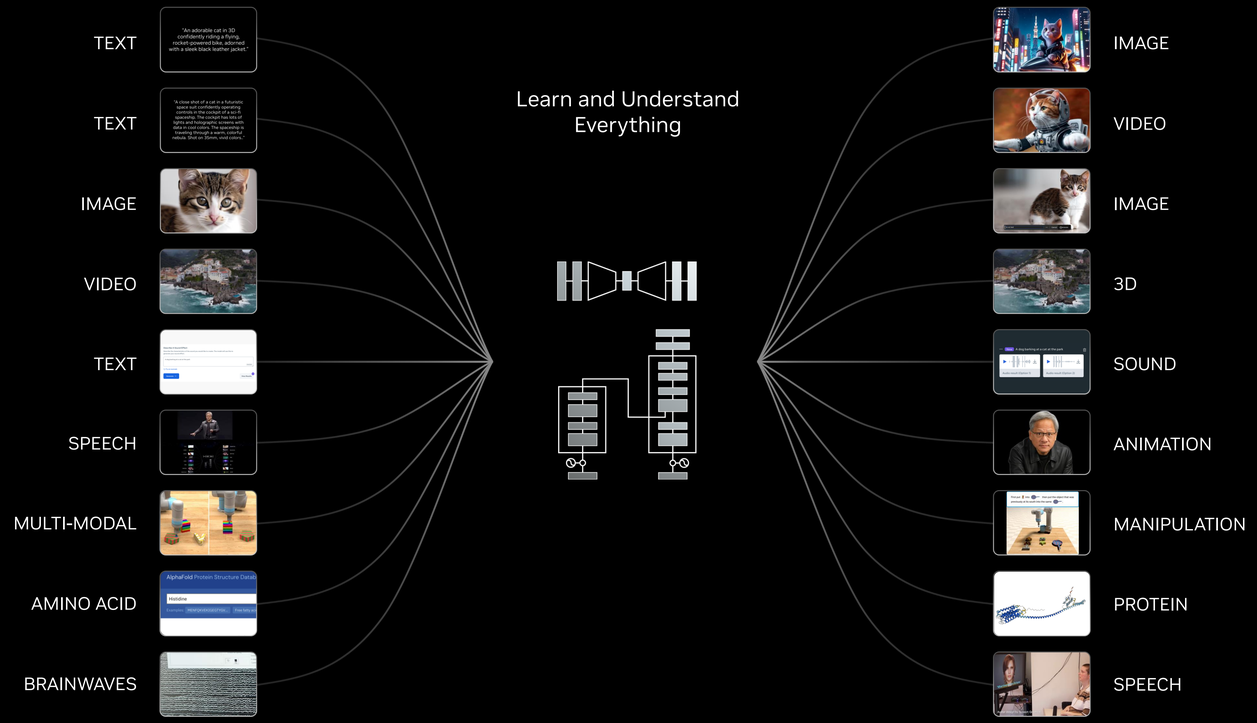

The era of generative AI has arrived, unlocking new opportunities for AI across many different applications

Generative AI is trained on large amounts of data to find patterns and relationships, learning the representation of almost anything with structure

It can then be prompted to generate text, images, video, code, or even proteins

For the very first time, computers can augment the human ability to generate information and create

1,600+ Generative AI companies are building on NVIDIA

Modern AI Is a Data Center Scale Computing Workload

Data centers are becoming AI factories: Data as input, intelligence as output

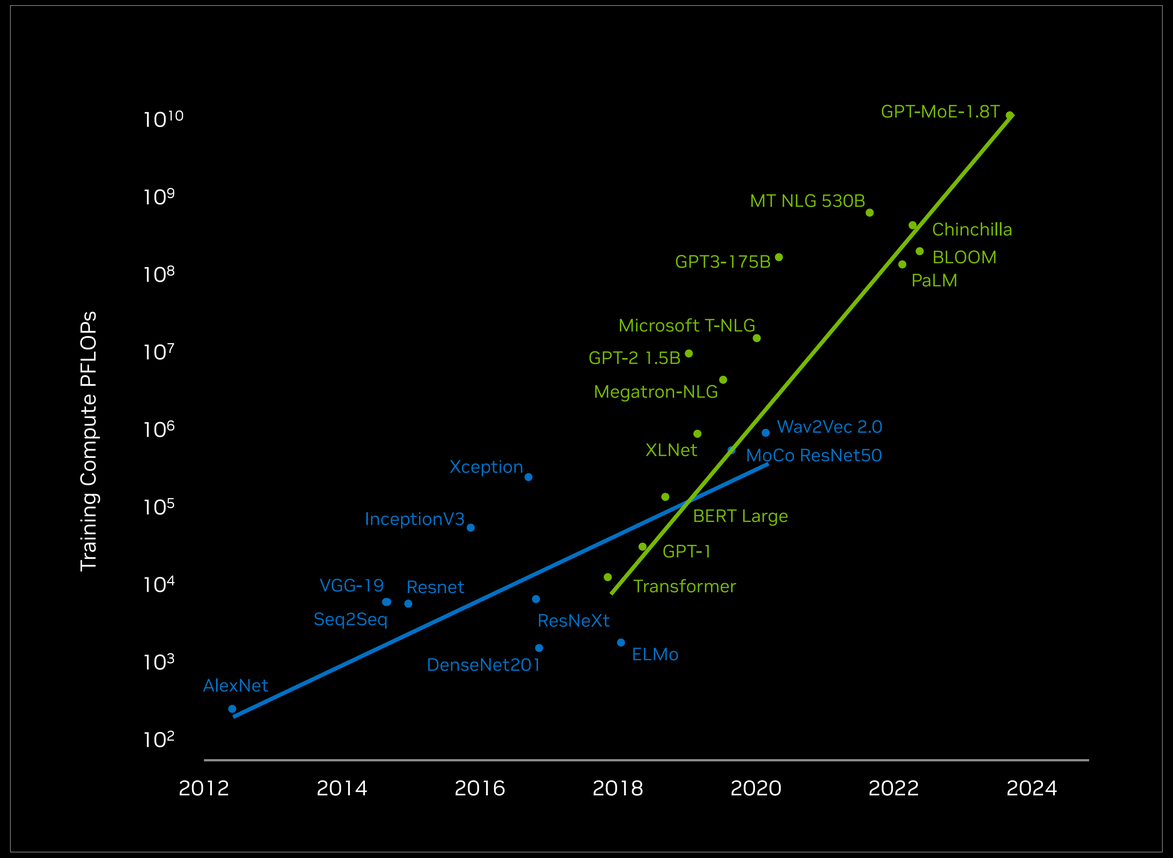

AI Training Computational Requirements

Large Language Models, based on the Transformer architecture, are one of today's most important advanced AI technologies, involving up to trillions of parameters that learn from text.

Developing them is an expensive, time-consuming process that demands deep technical expertise, distributed data center-scale infrastructure, and a full-stack accelerated computing approach.

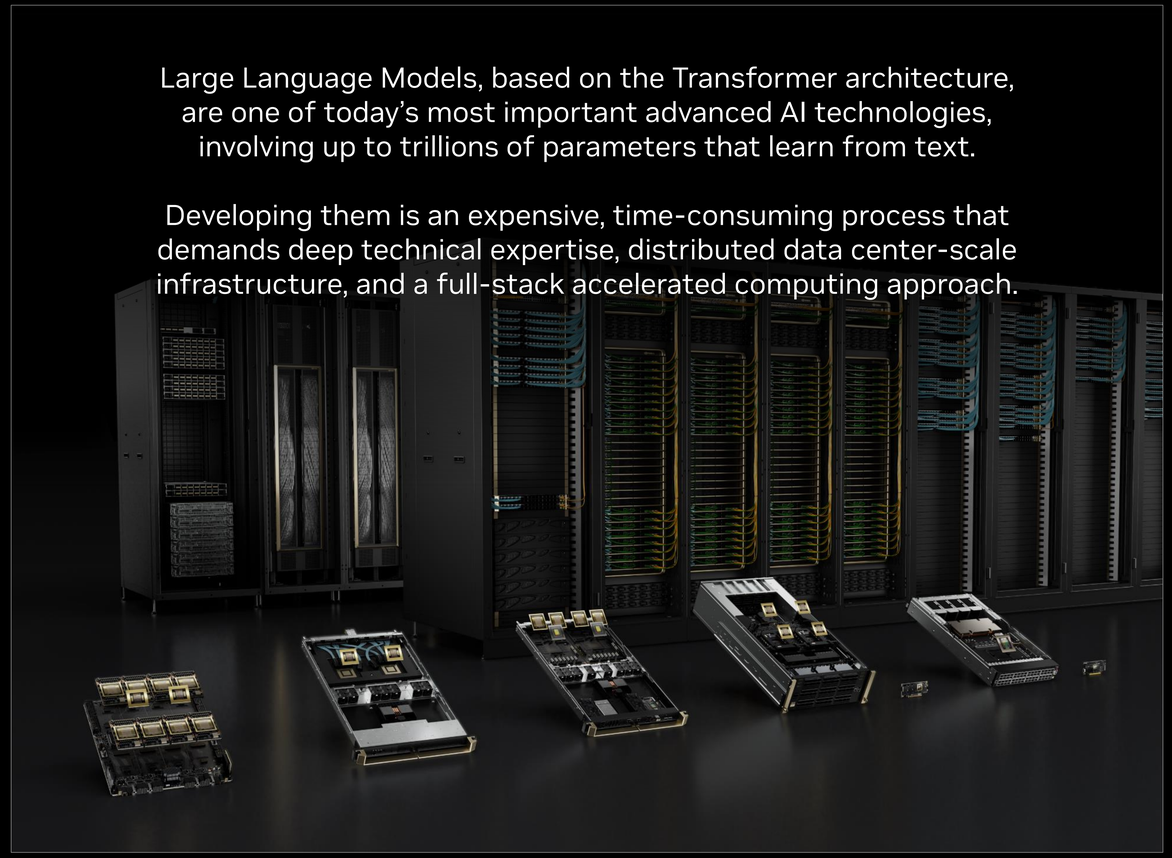

Fueling Giant-Scale AI Infrastructure

NVIDIA compute & networking GPU | DPU | Switch | CPU

NVIDIA

Full-Stack & Data Center Scale Acceleration

Drive significant cost savings and workload scaling

Classical Computing—960 CPU-only servers

Application

CPU server racks

Accelerated Computing—2 GPU servers

Application
Re-Engineered for Acceleration

CUDA-X Acceleration Libraries

Magnum IO

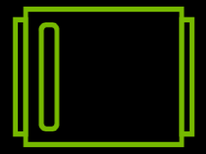

25X lower cost
84X better energy-efficiency

LLM Workload: Bert-Large Training and Inference | CPU Server: Dual-EYPC 7763 | GPU Server: Dual-EPYC 7763 + 8X H100 PCIe GPUs

The High ROI of High Compute Performance

$1 upfront investment in NVIDIA compute and networking can translate to $5 in CSP revenue over 4 years

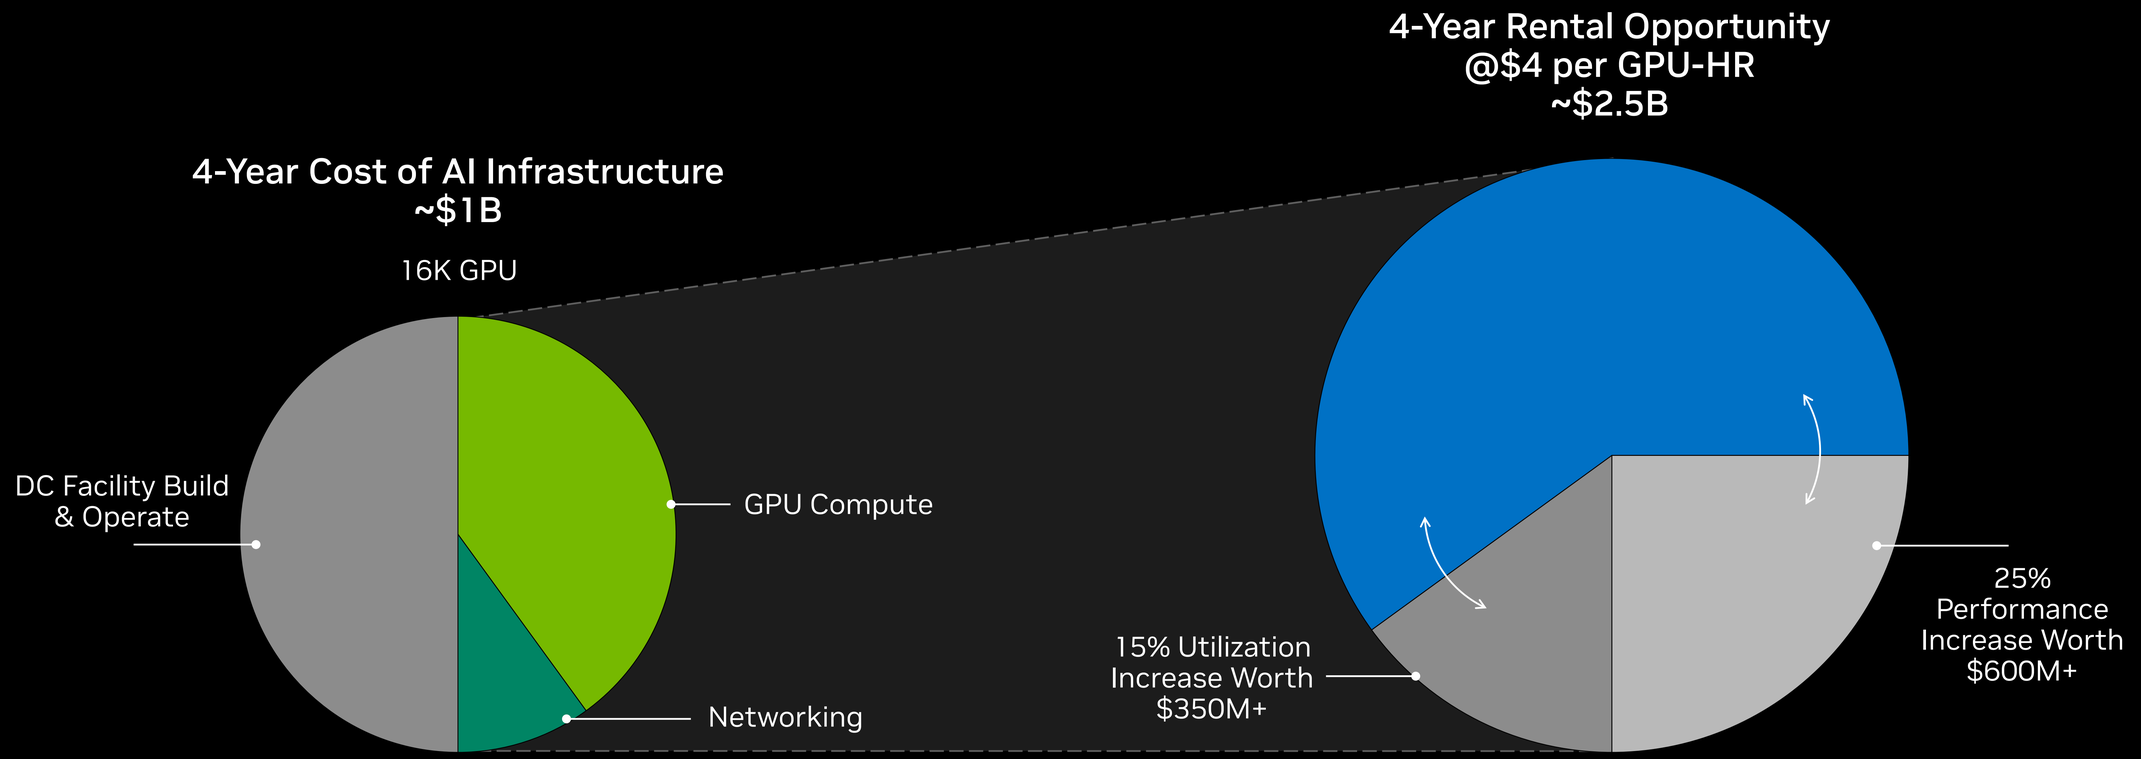

Illustrative example of NVIDIA GPU cost vs AI infrastructure total cost of ownership (TCO)

Powering the AI Industrial Revolution

Building and running enterprise Gen AI applications

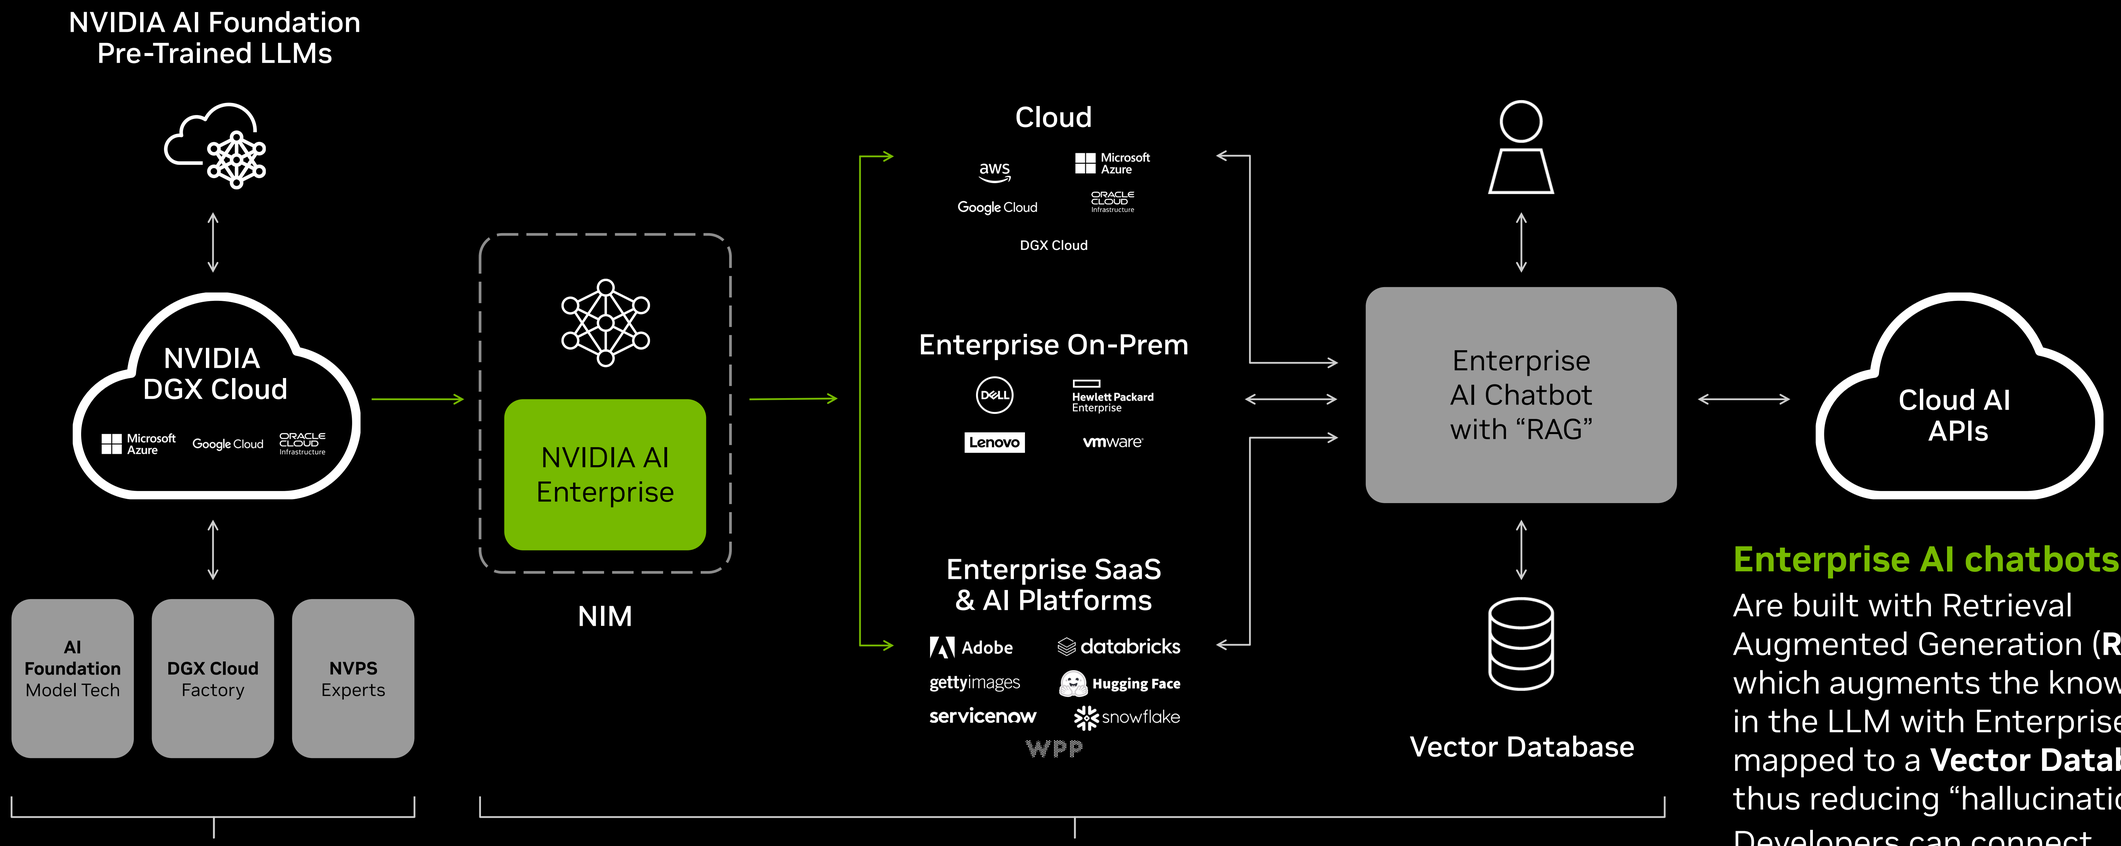

NVIDIA AI foundry service
for *building* Enterprise AI applications

NVIDIA AI enterprise ecosystem
for *running* Enterprise AI applications

Enterprise AI chatbots

Are built with Retrieval Augmented Generation (**RAG**), which augments the knowledge in the LLM with Enterprise data mapped to a **Vector Database**, thus reducing "hallucinations".

Developers can connect additional or 3rd party services to the AI chatbot via **cloud AI APIs**.

NVIDIA

The NVIDIA AI Foundry Model on DGX Cloud

For building enterprise AI applications

NVIDIA's "AI foundry" service leverages our AI infrastructure and expertise to build custom AI models for enterprise customers—analogous to a semiconductor foundry that uses its infrastructure and expertise to build custom chips for fabless customers.

An enterprise customer starts with an NVIDIA or 3rd party pre-trained AI model, available in **NVIDIA AI Foundations**. This model making service includes frameworks such as **NVIDIA NeMo** for custom LLMs and **NVIDIA Picasso** for custom generative AI for visual design.

With help from NVIDIA experts, the enterprise customer fine-tunes the model on their proprietary enterprise data and adds guardrails, using tools available in NVIDIA AI Foundations.

The fine-tuning and optimization is done on **NVIDIA DGX Cloud**, a cloud service that allows enterprises immediate access to NVIDIA AI infrastructure and software, hosted at partner cloud providers.

The enterprise customer ends up with a fully-trained and optimized AI model, fine-tuned on their proprietary enterprise data, that can be deployed anywhere—in the cloud or on-prem.

The NVIDIA AI Foundry model generates revenue based on per-node, per-month consumption of NVIDIA DGX Cloud.

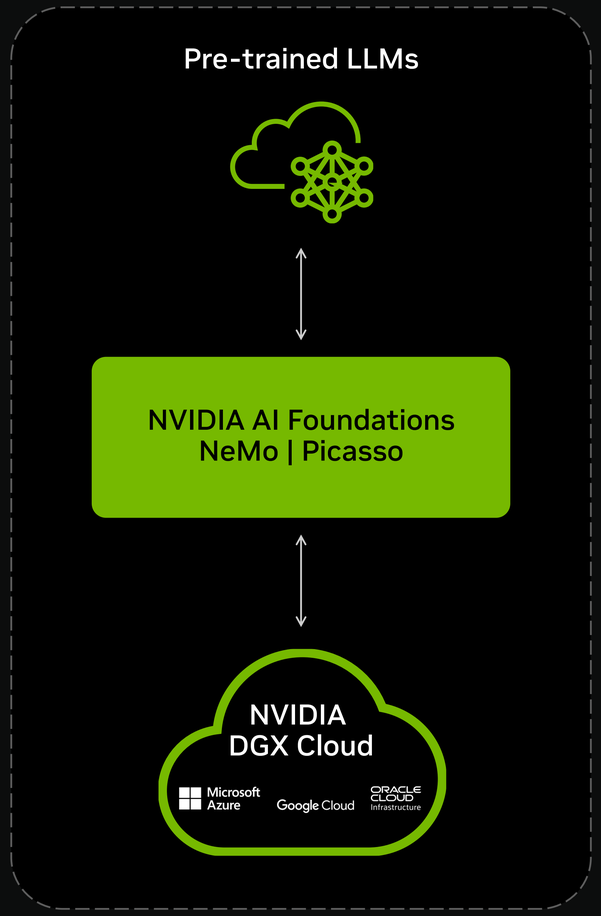

NVIDIA AI Foundry

AI Factories—A New Class of Data Centers

For running enterprise AI applications

AI factories are next-generation data centers that host advanced, full-stack accelerated computing platforms for the most computationally intensive tasks, where data comes in and intelligence comes out.

These new data centers will act as trusted engines of generative AI.

Every important company will run its own AI factories to securely process its valuable proprietary data and turn it into monetizable tokens, encapsulating its knowledge, intelligence, and creativity.

Nations are using AI factories as sovereign national resources—processing private datasets of companies, startups, universities and governments safely on shore to produce valuable insights.

In addition to the up-front revenue opportunity from data center systems, NVIDIA can generate recurring revenue from AI factories with **NVIDIA AI Enterprise**, the operating system for enterprise AI.

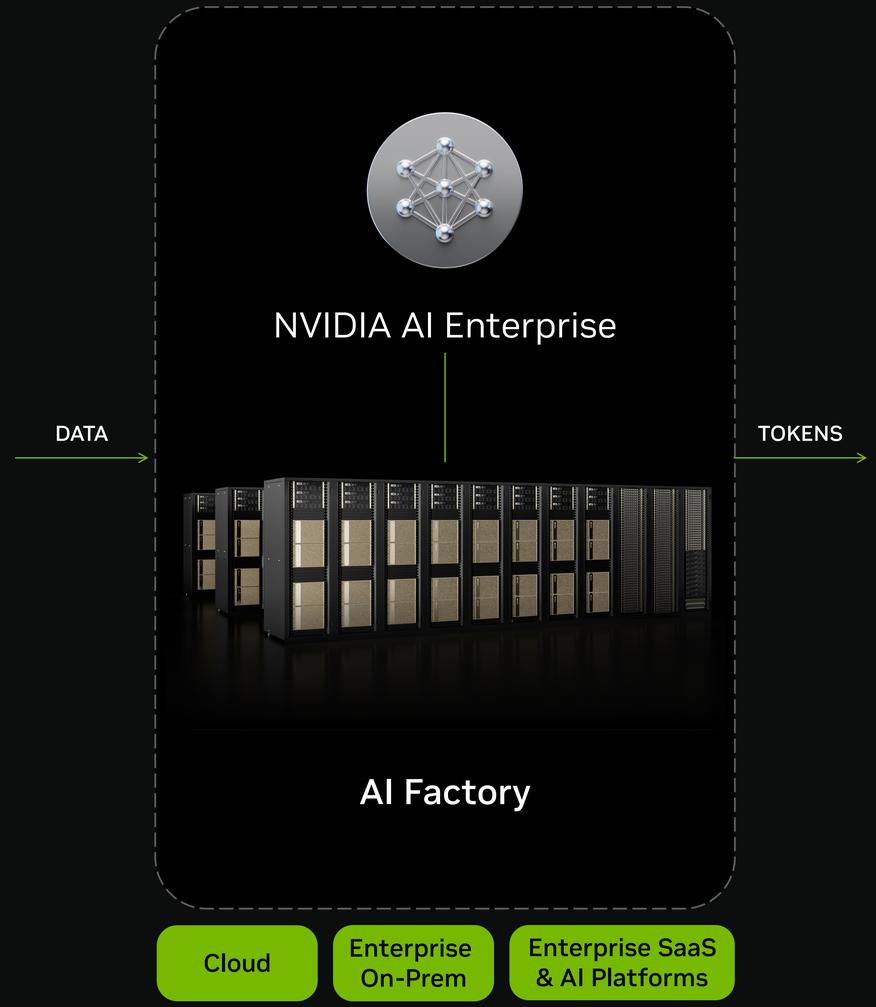

NVIDIA AI Enterprise

The operating system for enterprise AI

NVIDIA AI Enterprise

NVIDIA AI Enterprise is software for deploying and running AI with enterprise-grade security, API stability, manageability and support.

Cloud-native and available in every major cloud marketplace.

Certified to run on servers and workstations from all major OEMs.

Supported by all major global system integrators.

Integrated with and distributed by VMware.

AI Use Cases and Workflows

LLM

Speech AI

Recommenders

Cybersecurity

Medical Imaging

Video Analytics

Route Optimization

More

Run Anywhere

NVIDIA AI Enterprise

Azure | GCP | OCI | AWS

Consumption pricing per GPU-hour

NVIDIA Certified Server
Dell | HPE | Lenovo

Subscription pricing per GPU/year
(included with H100 PCIe/DGX)

NVIDIA AI Enterprise

Broad and deep ecosystem and distribution to reach every enterprise

GSI & Service Delivery

accenture

Booz | Allen | Hamilton

Capgemini

Deloitte.

Infosys

TATA CONSULTANCY SERVICES

wipro

AI Platforms

databricks

Hugging Face

snowflake

Software Platforms

gettyimages

servicenow

shutterstock

Adobe

WPP

Public Cloud Marketplaces

aws

Google Cloud

Microsoft Azure

ORACLE CLOUD Infrastructure

Private Cloud

vmware

Server OEMs

BOXX

cisco

DELL Technologies

HPE GreenLake

hp

Lenovo

SUPERMICRO

NVIDIA Inference Microservice (NIM)

Extending reach of the platform, connecting millions of developers to hundreds of millions of CUDA GPUs in the installed base

Industry Standard APIs
Text, Speech, Image, Video, 3D, Biology

Triton Inference Server
cuDF, CV-CUDA, DALI, NCCL, Post Processing Decoder

Cloud Native Stack
GPU Operator, Network Operator

Enterprise Management
Health Check, Identity, Metrics, Monitoring, Secrets Management

Kubernetes

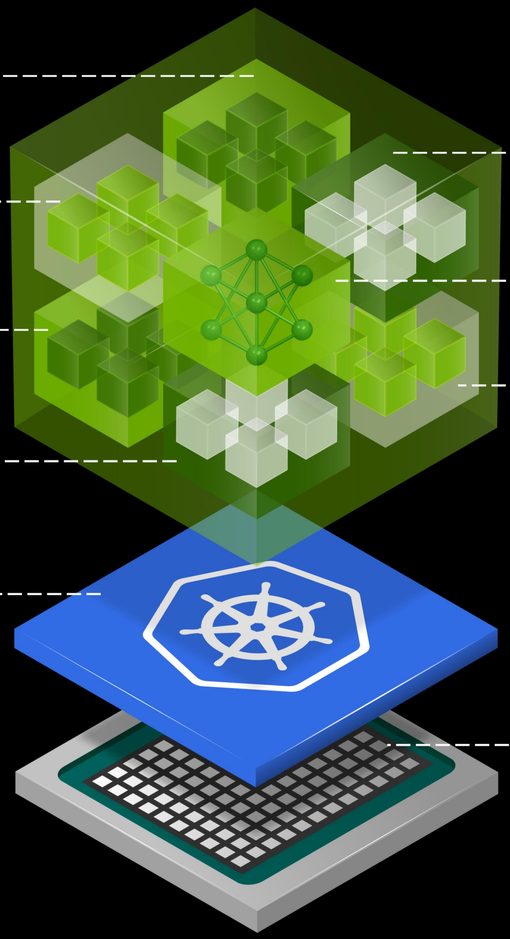

TensorRT LLM and Triton
cuBLAS , cuDNN, In-Flight Batching, Memory Optimization, FP8 Quantization

Optimized Model
Single GPU, Multi-GPU, Multi-Node

Customization Cache
P-Tuning, LORA, Model Weights

NVIDIA CUDA

Available Now as Part of NVIDIA AI Enterprise 5.0
$4,500/GPU/YEAR, $1/GPU/HOUR

NVIDIA

NVIDIA Go-to-Market Across Cloud and On-Premises

Reaching customers everywhere

NVIDIA

Partners

DGX Cloud

Google Cloud · Microsoft Azure · Oracle Cloud Infrastructure

NVIDIA AI Foundations - Cloud services for customizing and operating generative AI models

aws · Google Cloud · Microsoft Azure · Oracle Cloud Infrastructure

HGX

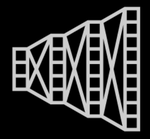

INFERENCE

Cloud

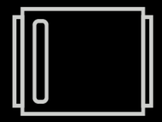

DGX

DELL Technologies · Hewlett Packard Enterprise · Lenovo

MGX

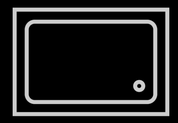

AGX

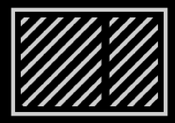

IGX

On-Prem

Driving Strong & Profitable Growth

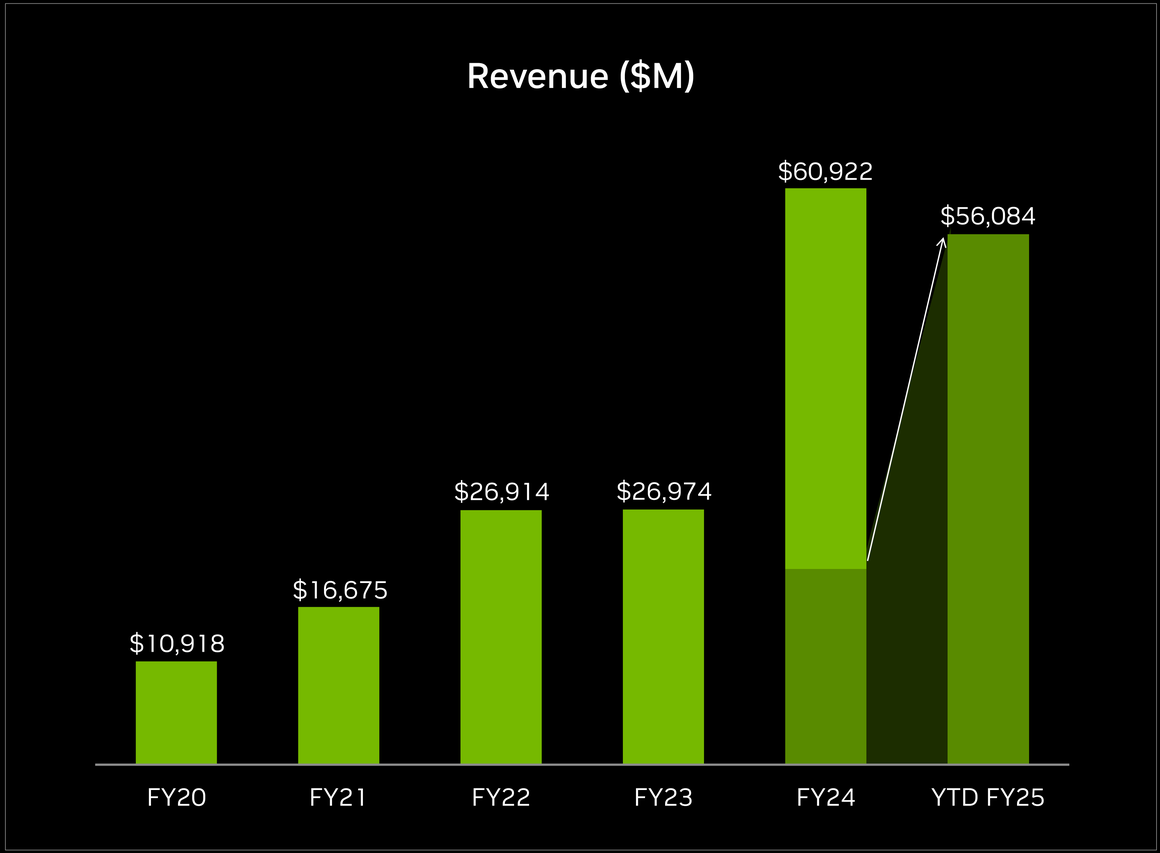

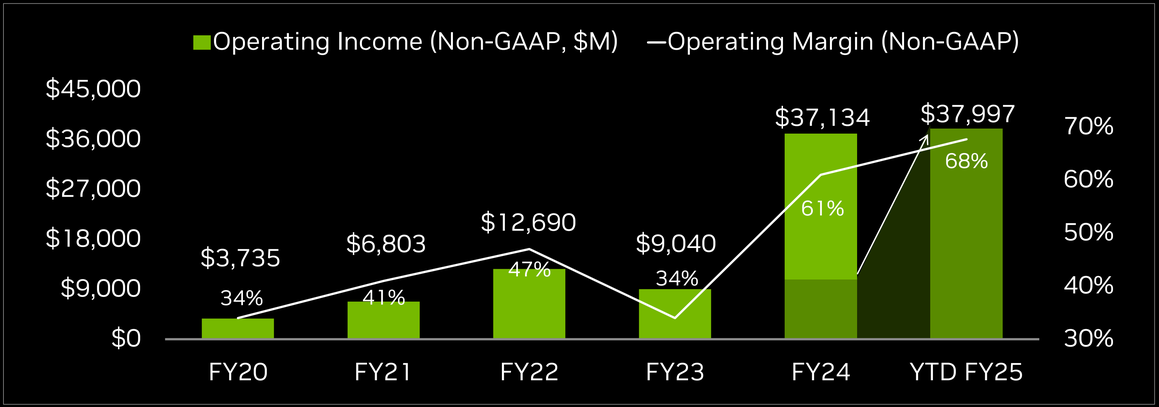

Fiscal year ends in January. Refer to Appendix for reconciliation of Non-GAAP measures. Operating margins rounded to the nearest percent.

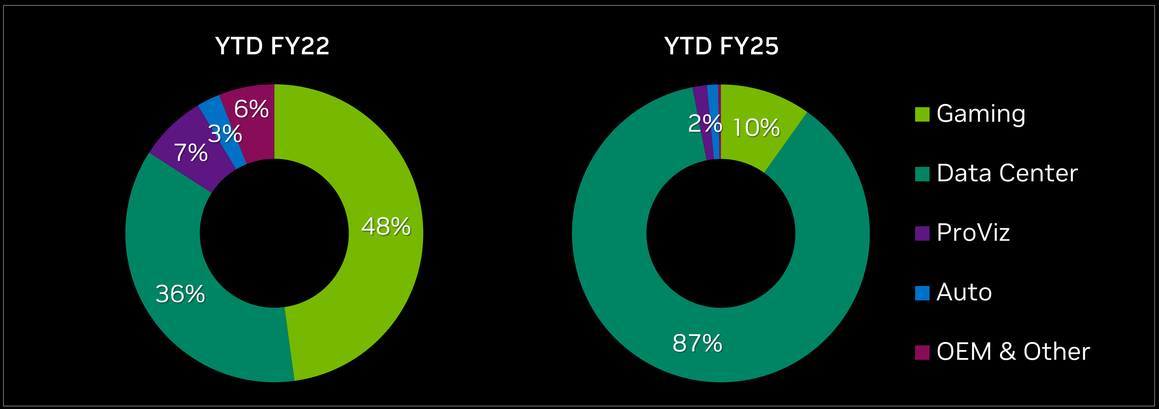

NVIDIA Gross Margins Reflect Value of Acceleration

Accelerated computing requires full-stack and data center-scale innovation across silicon, systems, algorithms and applications.

Significant expertise and effort are required, but application speed-ups can be incredible, resulting in dramatic cost and time-to-solution savings.

- For example, 2 NVIDIA HGX nodes with 16 NVIDIA H100 GPUs that cost $400K can replace 960 nodes of CPU servers that cost $10M for the same LLM workload.

NVIDIA chips carry the value of the full-stack, not just the chip.

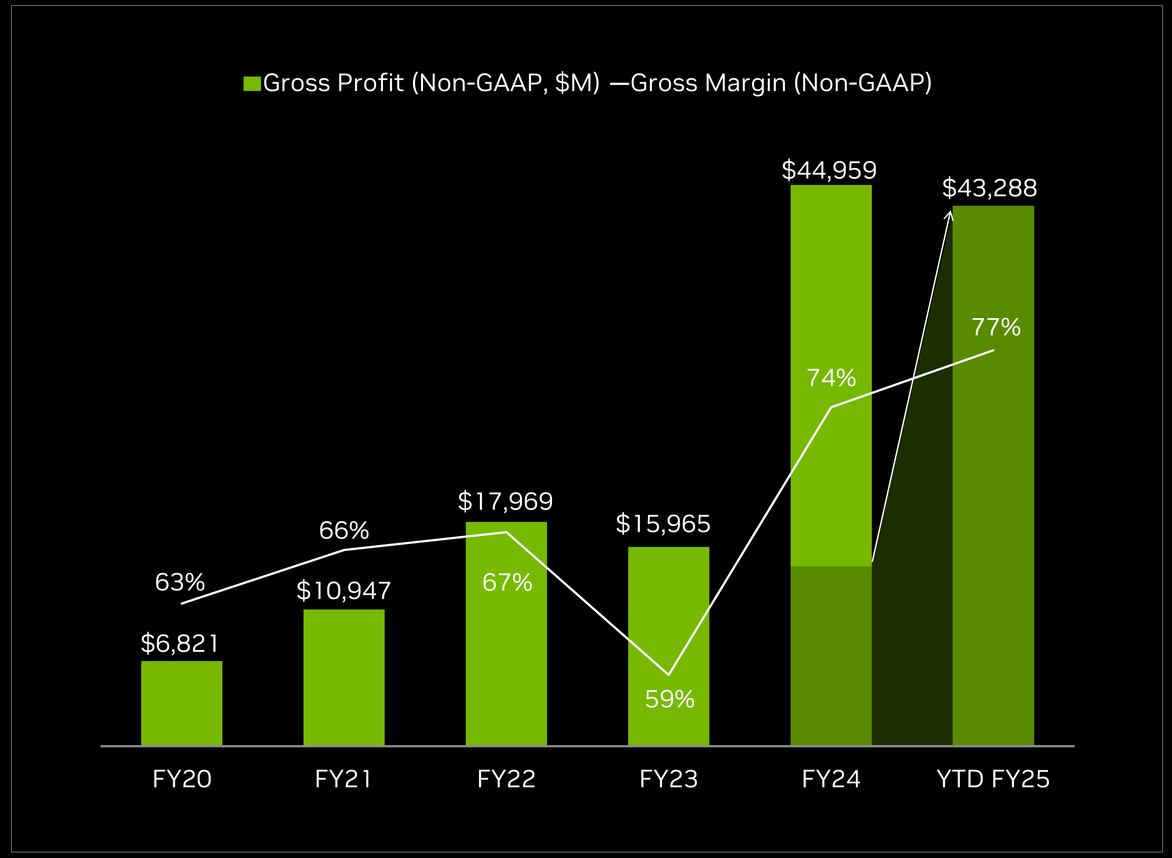

Cost comparison example based on latest available NVIDIA A100 GPU and Intel CPU inference results in the commercially available category of the MLPerf industry benchmark; includes related infrastructure costs such as networking.

Fiscal year ends in January. Refer to Appendix for reconciliation of Non-GAAP measures. Gross margins are rounded to the nearest percent.

Strong Cash Flow Generation

Free Cash Flow (Non-GAAP)

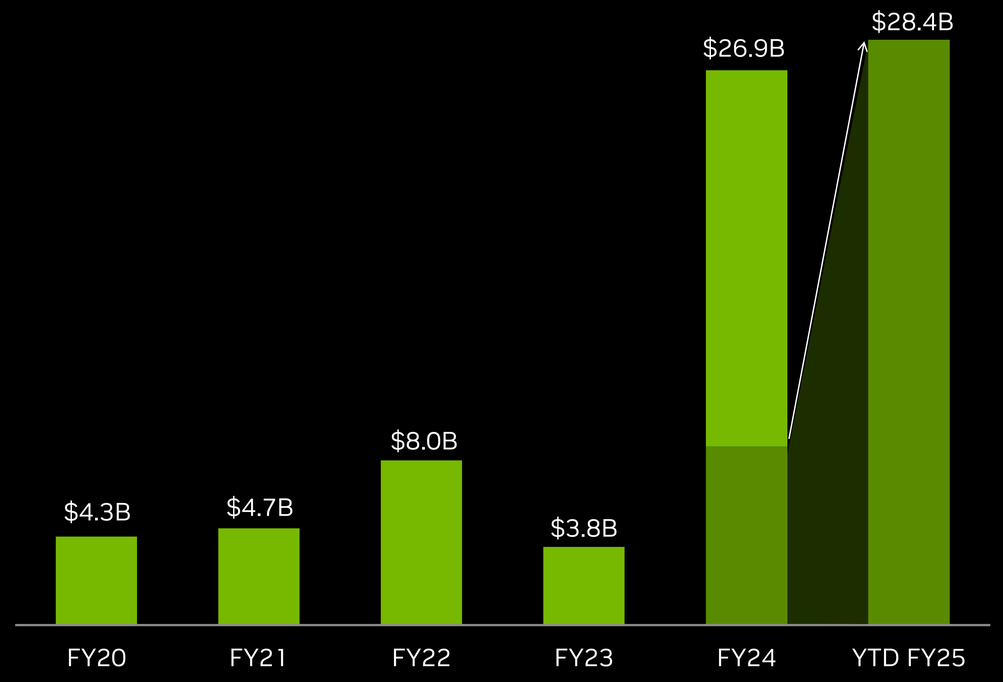

Fiscal year ends in January. Refer to Appendix for reconciliation of Non-GAAP measures.

Capital Allocation

Share Repurchase

Utilized $14.9B of cash for repurchases in H1 FY25

Additional $50B in stock repurchase authorization, adding to $7.5B which remained as of the end of Q2

Dividend

$344M in H1 FY25

Dividend increased by 150% in Q2 FY25

Plan to Maintain[1]

Strategic Investments

Growing Our Talent

Platform Reach & Ecosystem

[1] *Subject to continuing determination by our Board of Directors.*

Our Market Platforms at a Glance

Data Center

78% of FY24 Revenue

FY24 Revenue $47.5B
5-YR CAGR 75%

DGX/HGX/MGX/IGX systems

GPU | CPU | DPU | Networking
NVIDIA AI software

Gaming

17% of FY24 Revenue

FY24 Revenue $10.4B
5-YR CAGR 11%

GeForce GPUs for PC gaming

GeForce NOW cloud gaming

Professional Visualization

3% of FY24 Revenue

FY24 Revenue $1.6B
5-YR CAGR 7%

NVIDIA RTX GPUs
for workstations

Omniverse software

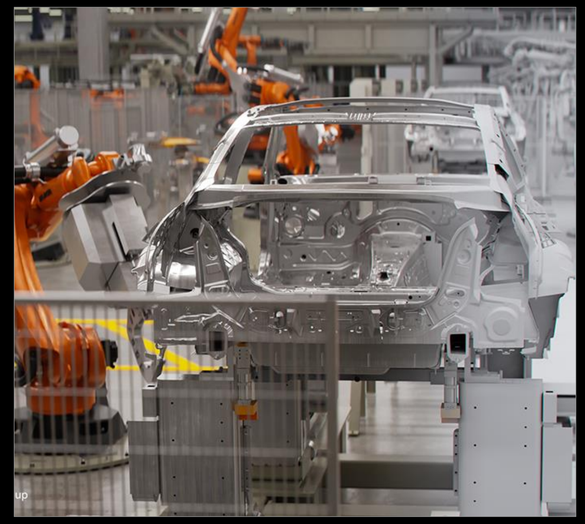

Automotive

2% of FY24 Revenue

FY24 Revenue $1.1B
5-YR CAGR 11%

DRIVE Hyperion sensor architecture
with AGX compute

DRIVE AV & IX full stack software
for ADAS, AV & AI cockpit

Data Center

The leading accelerated computing platform

Revenue ($M)

75% 5-YR CAGR
Through FY24

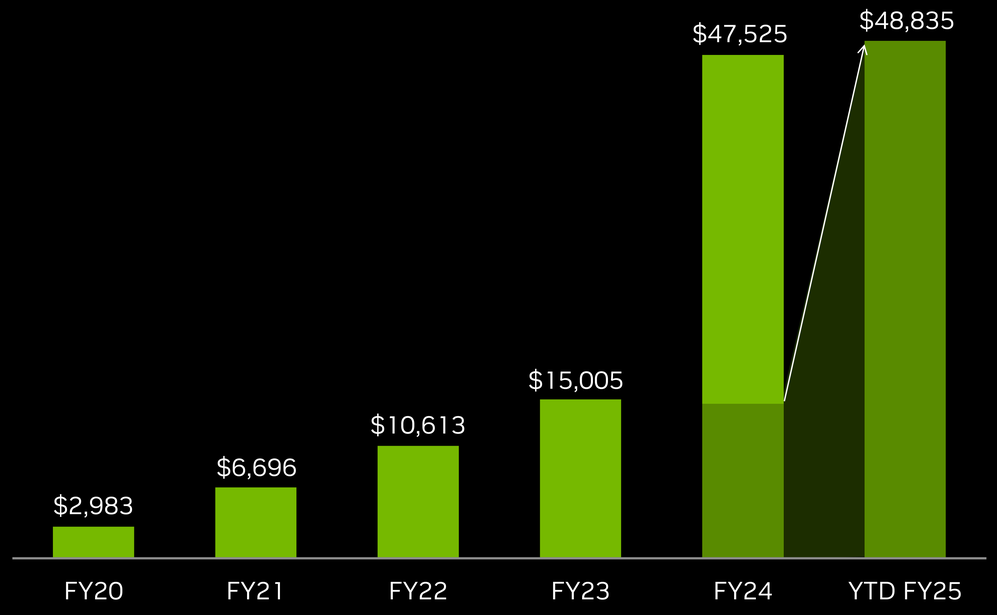

Leader in AI & HPC

#1 in AI training and inference

Used by all hyperscale and major cloud computing providers and over 40,000 companies

Powers over 75% of the TOP500 supercomputers

Growth Drivers

Broad data center platform transition from general-purpose to accelerated computing

Emergence of “AI factory” — optimized for refining data and training, inferencing, and generating AI

Broader and faster product launch cadence to meet a growing and diverse set of AI opportunities

DGX Cloud services and NVIDIA AI Enterprise software for building and running enterprise AI applications

NVIDIA Blackwell Platform

HGX B100

NVLINK Switch

GB200 Superchip
Compute Node

Quantum X800 Switch
ConnectX-8 SuperNIC

Spectrum X800 Switch
BlueField-3 SuperNIC

Gaming

GeForce—the world's largest gaming platform

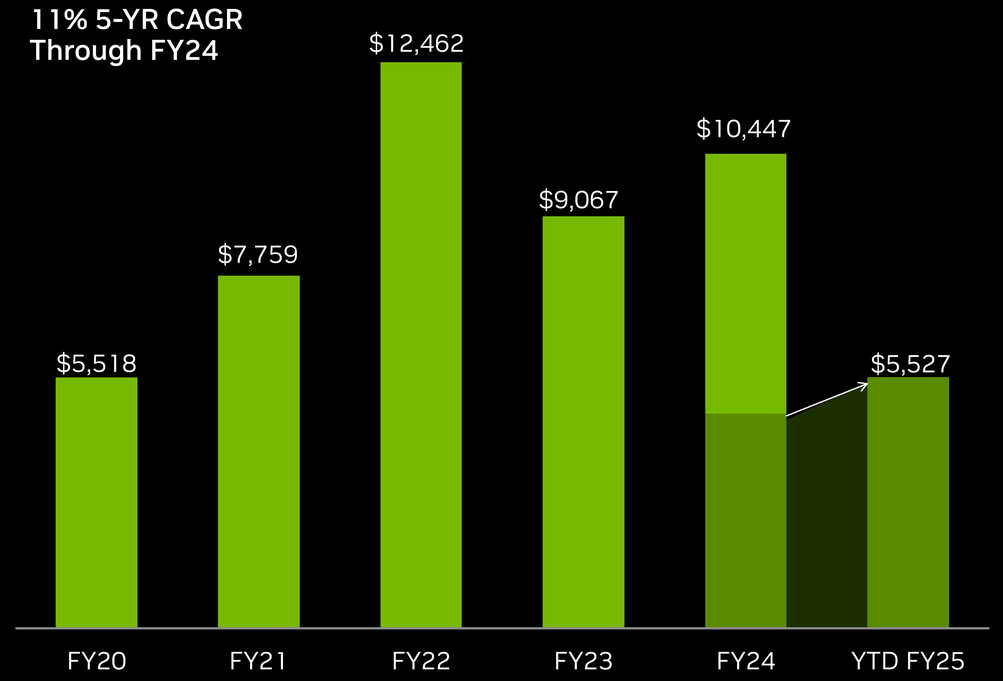

Leader in PC Gaming

Strong #1 market position

15 of the top 15 most popular GPUs on Steam

Leading performance & innovation

200M+ gamers on GeForce

Growth Drivers

Rising adoption of NVIDIA RTX in games

Expanding universe of gamers & creators

Gaming laptops & Gen AI on PCs

GeForce NOW Cloud gaming

GeForce Extends Growth, Large Upgrade Opportunity

GeForce Gaming Revenue

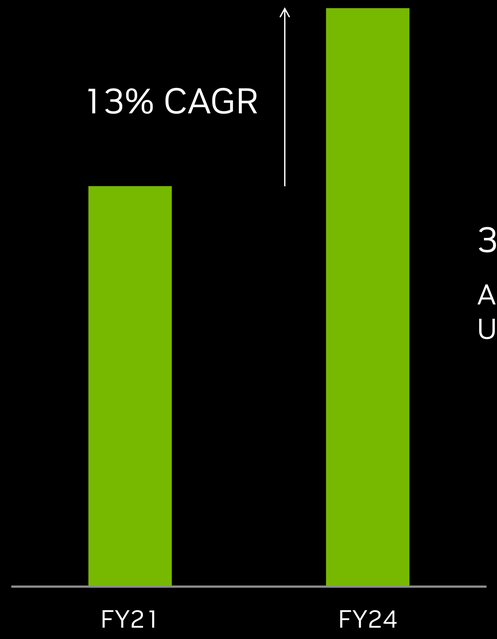

3YR CAGR

ASP 16%
Units -2%

Installed Base

53% RTX

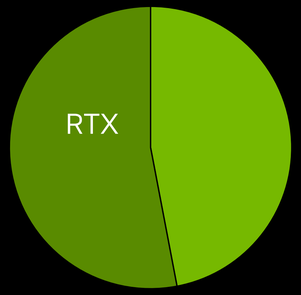

25% RTX3060+ Performance

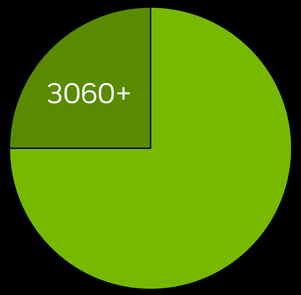

$699+ Cumulative Sell-Through $

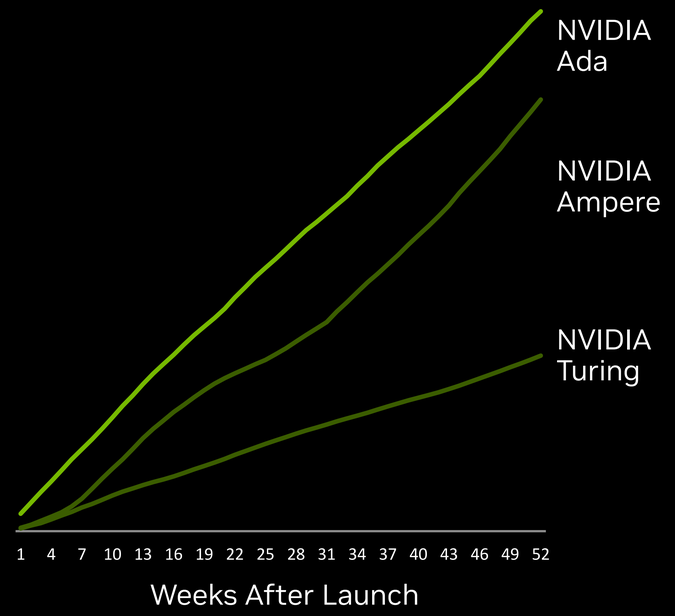

Installed Base Needs Upgrade

Source: NVIDIA estimates

Professional Visualization

Workstation graphics

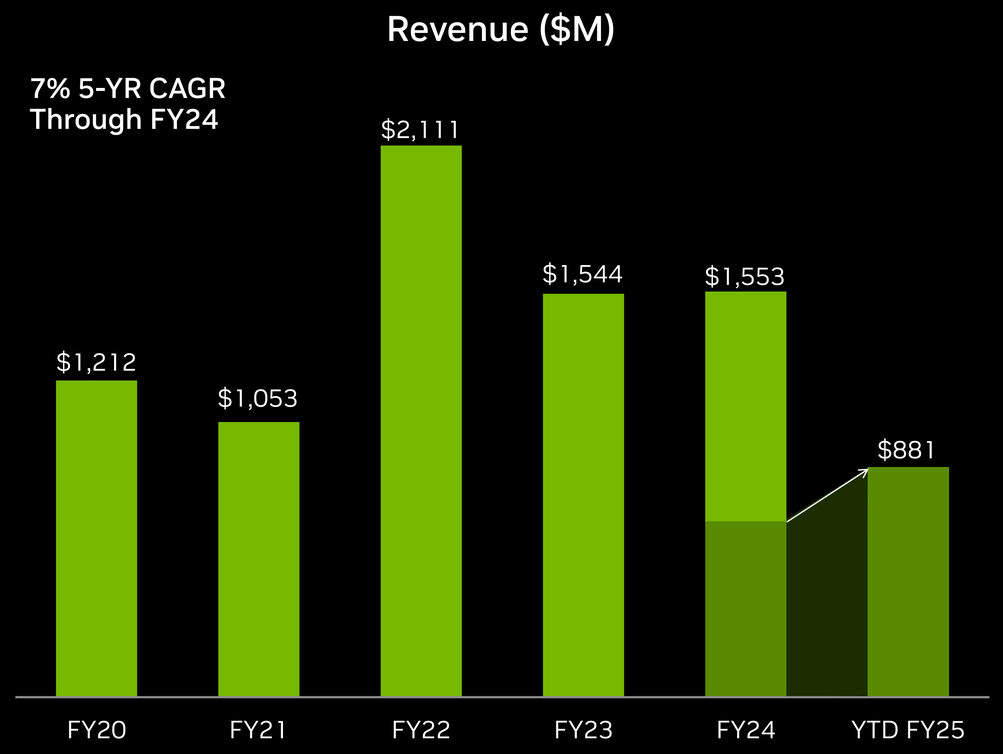

Leader in Workstation Graphics

95%+ market share in graphics for workstations

45M Designers and Creators

Strong software ecosystem with over 100 RTX accelerated and supported applications

Growth Drivers

Gen AI adoption across design and creative industries

Enterprise AI development, model fine tuning, cross-industry

Ray tracing revolutionizing design and content creation

Expanding universe of designers and creators

Omniverse for digital twins and collaborative 3D design

Hybrid work environments

Automotive

Autonomous vehicle and AI cockpit

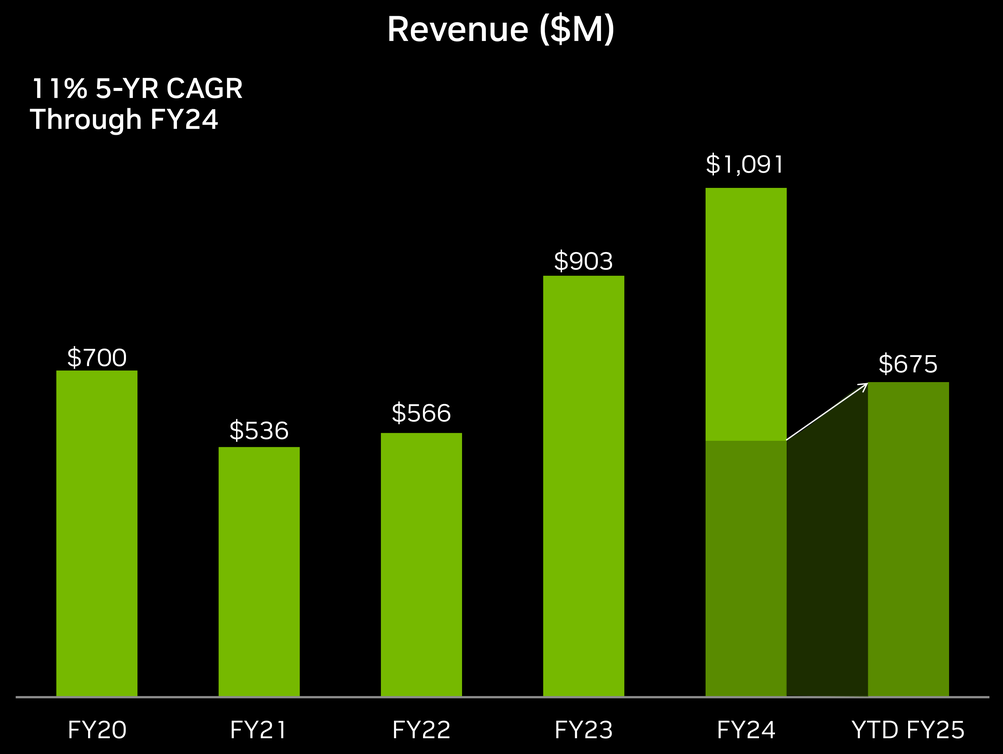

Leader in Autonomous Driving

NVIDIA DRIVE is our end-to-end Autonomous Vehicle (AV) and AI Cockpit platform featuring a full software stack and is powered by NVIDIA (systems-on-a-chip) SoCs in the vehicle

DRIVE Orin SoC ramp began in FY23

Next-generation DRIVE Thor SoC ramp to begin in FY26

Over 40 customers including 20 of top 30 EV makers, 7 of top 10 truck makers, 8 of top 10 robotaxi makers

Growth Drivers

Adoption of centralized car computing and software-defined vehicle architectures

AV software and services:

Mercedes-Benz

Jaguar Land Rover

$1 Trillion Long-Term Available Market Opportunity

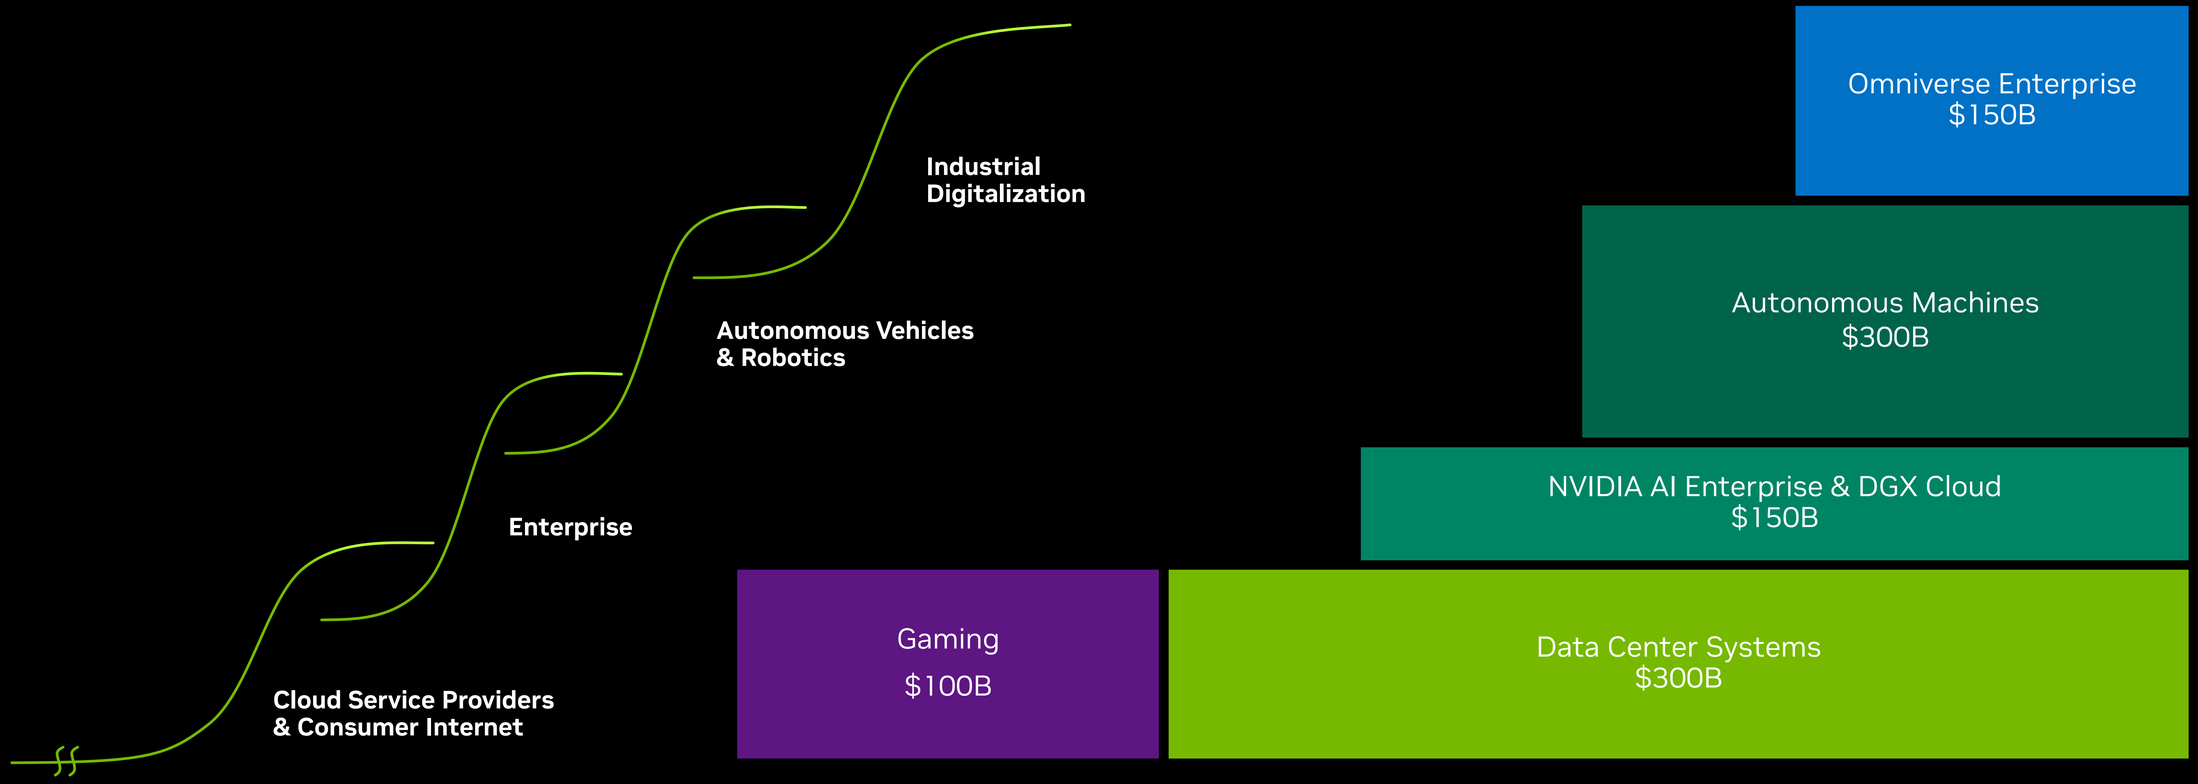

Financials

Annual Cash & Cash Flow Metrics

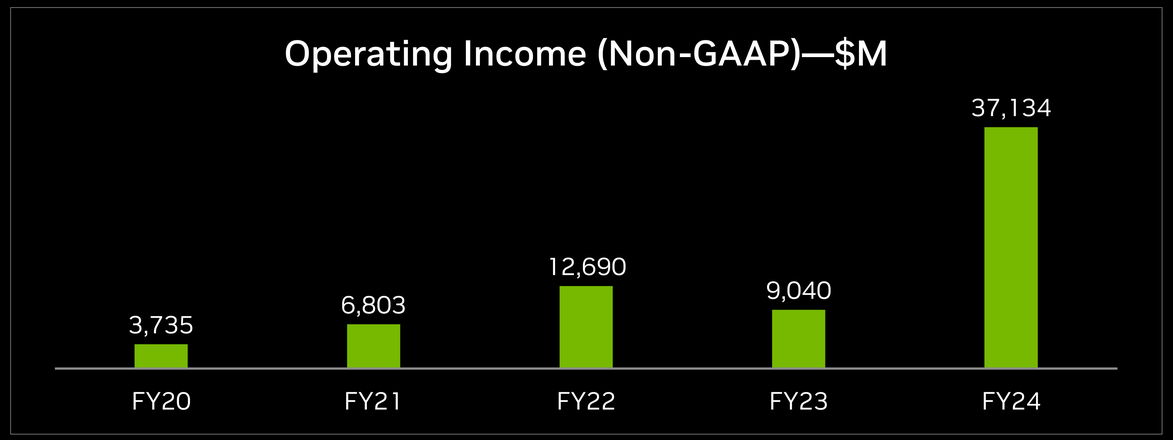

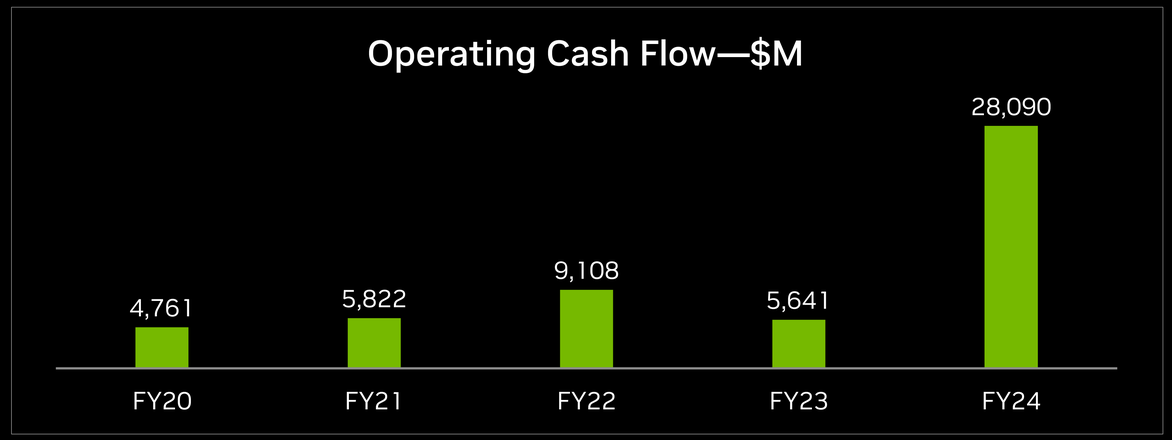

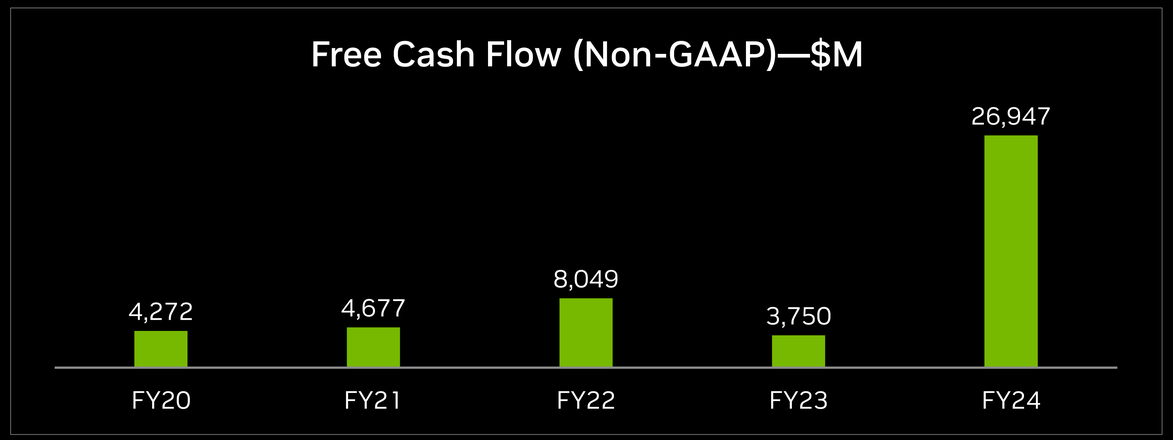

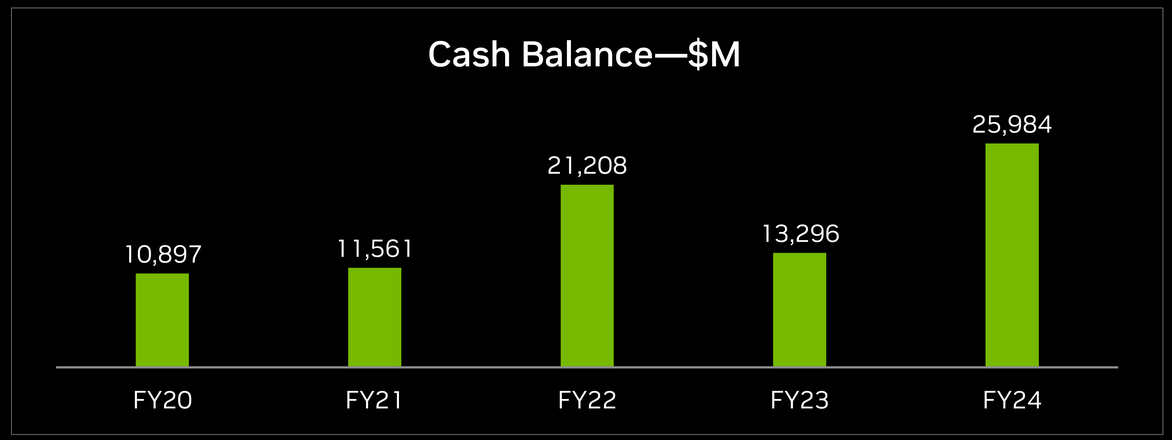

Cash balance is defined as cash and cash equivalents plus marketable securities
Refer to Appendix for reconciliation of non-GAAP measures

Corporate Sustainability

Environmentally Conscious

NVIDIA Blackwell GPUs are as much as 20X more energy efficient than CPUs for certain AI and HPC workloads

On track to source 100% renewable electricity for offices and data centers under operational control by end of FY25

On track to engage manufacturing suppliers comprising at least 67% of scope 3 category 1 GHG emissions with the goal of effecting supplier adoption of science-based targets by end of FY26

A Place For People To Do Their Life's Work

"Best Places to Work"
GLASSDOOR

"America's Most Sustainable Companies"
BARRON'S

"America's 100 Best Companies to Work For"
FORTUNE

"America's Most Responsible Companies"
NEWSWEEK

Management

Fast Company Magazine's World's 50 Most Innovative Companies

Fortune's World's Most Admired Companies

Time Magazine's 100 Most Influential Companies

Wall Street Journal's Management Top 250

Corporate Governance

50% of Board is Gender, Racially, or Ethnically Diverse

92% of Directors are independent

Reconciliation of Non-GAAP to GAAP Financial Measures

Reconciliation of Non-GAAP to GAAP Financial Measures

Gross Margin ($ in Millions & Margin Percentage)	Non-GAAP	Acquisition-Related and Other Costs (A)	Stock-Based Compensation (B)	Other (C)	GAAP
FY 2020	$6,821	—	(39)	(14)	$6,768
	62.5%	—	(0.4)	(0.1)	62.0%
FY 2021	$10,947	(425)	(88)	(38)	$10,396
	65.6%	(2.6)	(0.5)	(0.2)	62.3%
FY 2022	$17,969	(344)	(141)	(9)	$17,475
	66.8%	(1.4)	(0.5)	—	64.9%
FY 2023	$15,965	(455)	(138)	(16)	$15,356
	59.2%	(1.7)	(0.5)	(0.1)	56.9%
FY 2024	$44,959	(477)	(141)	(40)	$44,301
	73.8%	(0.8)	(0.2)	(0.1)	72.7%
YTD Q2'24	$14,417	(239)	(58)	(10)	$14,110
	69.7%	(1.2)	(0.3)	—	68.2%
YTD Q2'25	$43,288	(238)	(75)	4	$42,979
	77.2%	(0.5)	(0.1)	—	76.6%

A. Consists of amortization of intangible assets and inventory step-up
B. Stock-based compensation charge was allocated to cost of goods sold
C. Other consists of IP-related costs and assets held for sale related adjustments

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

Operating Income and Margin ($ in Millions & Margin Percentage)	Non-GAAP	Acquisition Termination Cost	Acquisition-Related and Other Costs (A)	Stock-Based Compensation (B)	Other (C)	GAAP
FY 2020	$3,735	—	(31)	(844)	(14)	$2,846
	34.2%	—	(0.3)	(7.7)	(0.1)	26.1%
FY 2021	$6,803	—	(836)	(1,397)	(38)	$4,532
	40.8%	—	(5.0)	(8.4)	(0.2)	27.2%
FY 2022	$12,690	—	(636)	(2,004)	(9)	$10,041
	47.2%	—	(2.5)	(7.4)	—	37.3%
FY 2023	$9,040	(1,353)	(674)	(2,710)	(79)	$4,224
	33.5%	(5.0)	(2.5)	(10.0)	(0.3)	15.7%
FY 2024	$37,134	—	(583)	(3,549)	(30)	$32,972
	61.0%	—	(1.0)	(5.8)	(0.1)	54.1%
YTD Q2'24	$10,828	—	(311)	(1,576)	—	$8,941
	52.3%	—	(1.5)	(7.6)	—	43.2%
YTD Q2'25	$37,997	—	(286)	(2,164)	4	$35,551
	67.8%	—	(0.5)	(3.9)	—	63.4%

A. Consists of amortization of acquisition-related intangible assets, inventory step-up, transaction costs, compensation charges, and other costs
B. Stock-based compensation charge was allocated to cost of goods sold, research and development expense, and sales, general and administrative expense
C. Comprises of legal settlement cost, contributions, restructuring costs and assets held for sale related adjustments

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

($ in Millions)	Free Cash Flow	Purchases Related to Property and Equipment and Intangible Assets	Principal Payments on Property and Equipment and Intangible Assets	Net Cash Provided by Operating Activities
FY 2020	$4,272	489	—	$4,761
FY 2021	$4,677	1,128	17	$5,822
FY 2022	$8,049	976	83	$9,108
FY 2023	$3,750	1,833	58	$5,641
FY 2024	$26,947	1,069	74	$28,090
YTD Q2’24	$8,691	537	31	$9,259
YTD Q2’25	$28,418	1,346	69	$29,833


NVIDIA