



Our Body of Work



Nothing makes me prouder than the incredible people who have made NVIDIA the company it is today. We want our company to be where they can do their life's work.

Together, we continue to drive advances in AI, HPC, gaming, creative design, autonomous vehicles, and robotics—some of the world's most impactful areas.

I want to thank NVIDIA employees, developers, partners, customers, and families for the amazing work you do. Exciting new frontiers lie ahead. Let's seek them out together.


A white handwritten signature of Jensen Huang, written in a cursive style.

Jensen Huang

Our Body of Work

NVIDIA pioneered accelerated computing to tackle challenges no one else can solve. Our work in AI and digital twins is transforming the world's largest industries and profoundly impacting society.


Photo by Jason O'Rear



Pioneering Accelerated Computing

Accelerated computing requires full-stack optimization, from chip architecture, systems, and acceleration libraries, to refactoring the applications. The global NVIDIA ecosystem spans 4.5 million developers, 40,000 companies, and over 3,300 applications.





Sparking the iPhone Moment of AI

The acceleration of deep learning ignited the big bang of AI. ChatGPT, a large language model powered by an NVIDIA DGX™ AI supercomputer, reached 100 million users in just two months. Its magical capabilities have captured the world's imagination. Generative AI is a new computing platform, like the PC, internet, and mobile-cloud.



What's the definition of a large language model?



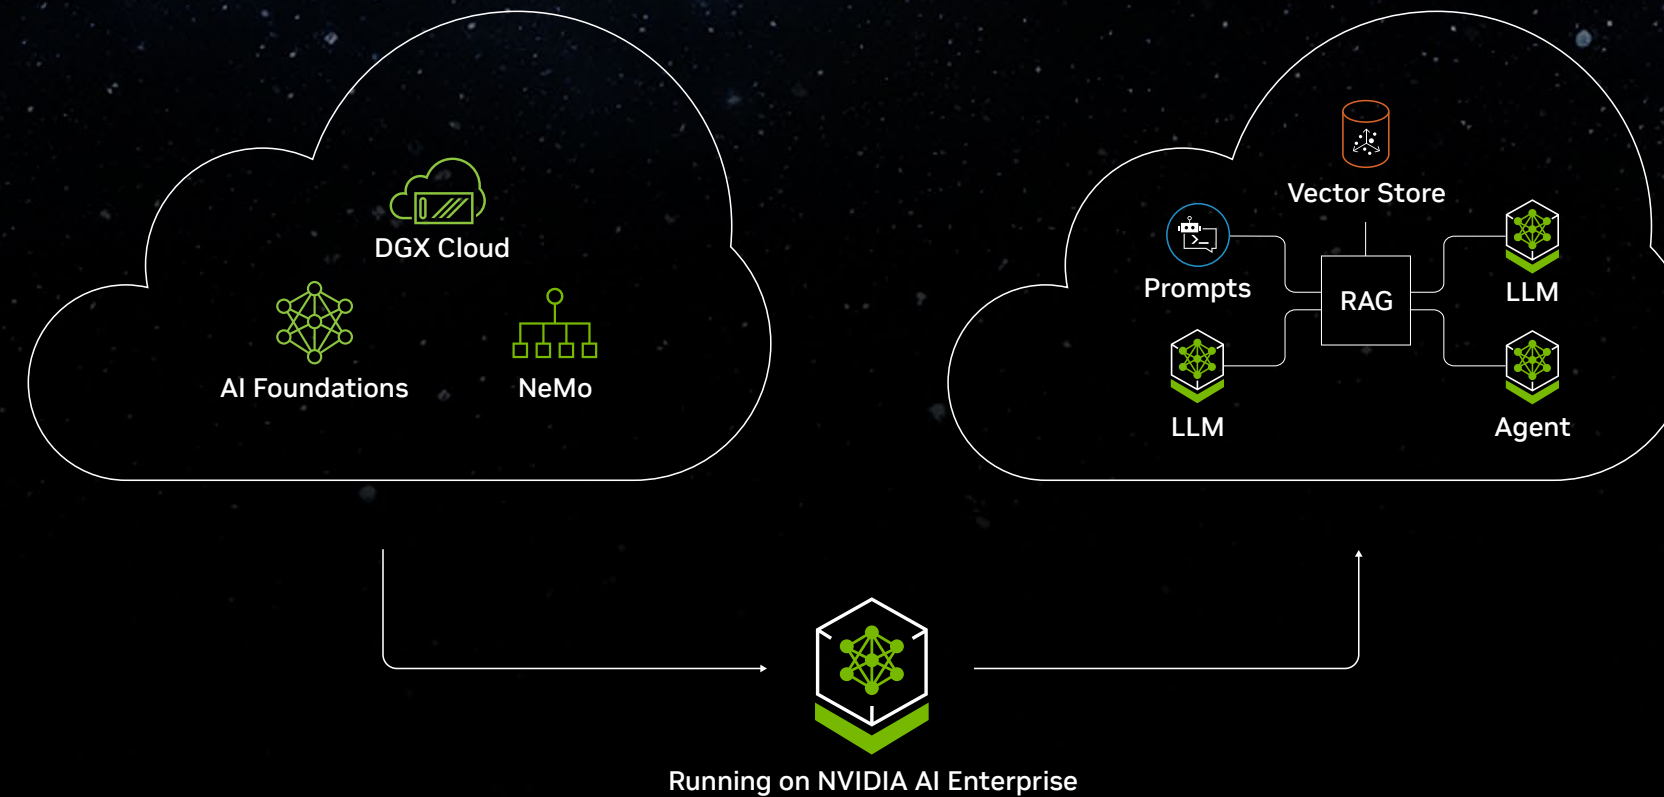
A large language model is a type of artificial intelligence system that has been trained on massive amounts of text data and can generate human-like language responses to input.

AI Factories for All the World's Industries

In the future, every company will have AI factories. To help businesses easily deploy tailored generative AI applications to drive innovation and transformation across any industry, NVIDIA offers a custom AI model service built on world-class AI technology, AI factory, and model-making know-how.

This AI foundry service includes AI Foundation LLM models, NVIDIA DGX Cloud AI factories, and NVIDIA AI Enterprise acceleration runtime engines.

Through partnerships with Amazon, Google, Microsoft and Oracle, NVIDIA is bringing state-of-the-art AI capabilities within reach to thousands of organizations.




NVIDIA Accelerated Computing Is Sustainable Computing

Data centers are already about 1-2% of global electricity consumption. That consumption is expected to continue to grow. This continued growth is not sustainable.

If we switched accelerated computing workloads from CPU-only servers to GPU-accelerated systems worldwide, we estimate nearly 12 trillion watt-hours of energy savings a year, equivalent to the electricity requirements of nearly 1.7 million U.S. homes.


Acceleration is the best way to reclaim power and achieve sustainability and net-zero emissions.





NVIDIA Powers AI Factories

Data centers process mountains of continuous data to train and refine AI software. Companies are manufacturing intelligence, and their data centers are becoming giant AI factories. NVIDIA is the engine of the world's AI infrastructure.



Hopper— The Engine for the World's AI Infrastructure

The NVIDIA Hopper™ architecture is powering the next wave of AI data centers. The first Hopper-based GPU, the NVIDIA H100, comes packed with 80 billion transistors and delivers an order-of-magnitude performance leap over its predecessor.

NVIDIA DGX— Purpose-Built for the Unique Demands of AI

Our fourth-generation NVIDIA DGX system is the world's first AI platform to be built with the new H100 GPUs. Each DGX H100 provides 32 petaflops of AI performance at FP8 precision—6X more than the prior generation. The next-generation DGX SuperPOD™ will expand the frontiers of AI with the ability to run massive workloads with trillions of parameters.





Every Data Center Can Now Be a Generative AI Data Center

NVIDIA's inference platform provides one architecture for diverse AI workloads and maximum data center acceleration and elasticity: L4 GPUs for AI video; L40 GPUs for NVIDIA Omniverse and graphics rendering; H100 NVL for scaling out large language model inference; and Grace Hopper Superchips for recommender systems and vector databases.

NVIDIA Spectrum-X™ is an accelerated networking platform designed to improve the performance and efficiency of Ethernet-based AI clouds. Spectrum-X is supercharged by NVIDIA acceleration software and software development kits, allowing developers to build software-defined, cloud-native AI applications.

NVIDIA DGX GH200—A New Class of AI Supercomputer

The most efficient large memory supercomputer, the DGX GH200 enables the development of next-generation models for generative AI language applications, recommender systems, and data analytics workloads.





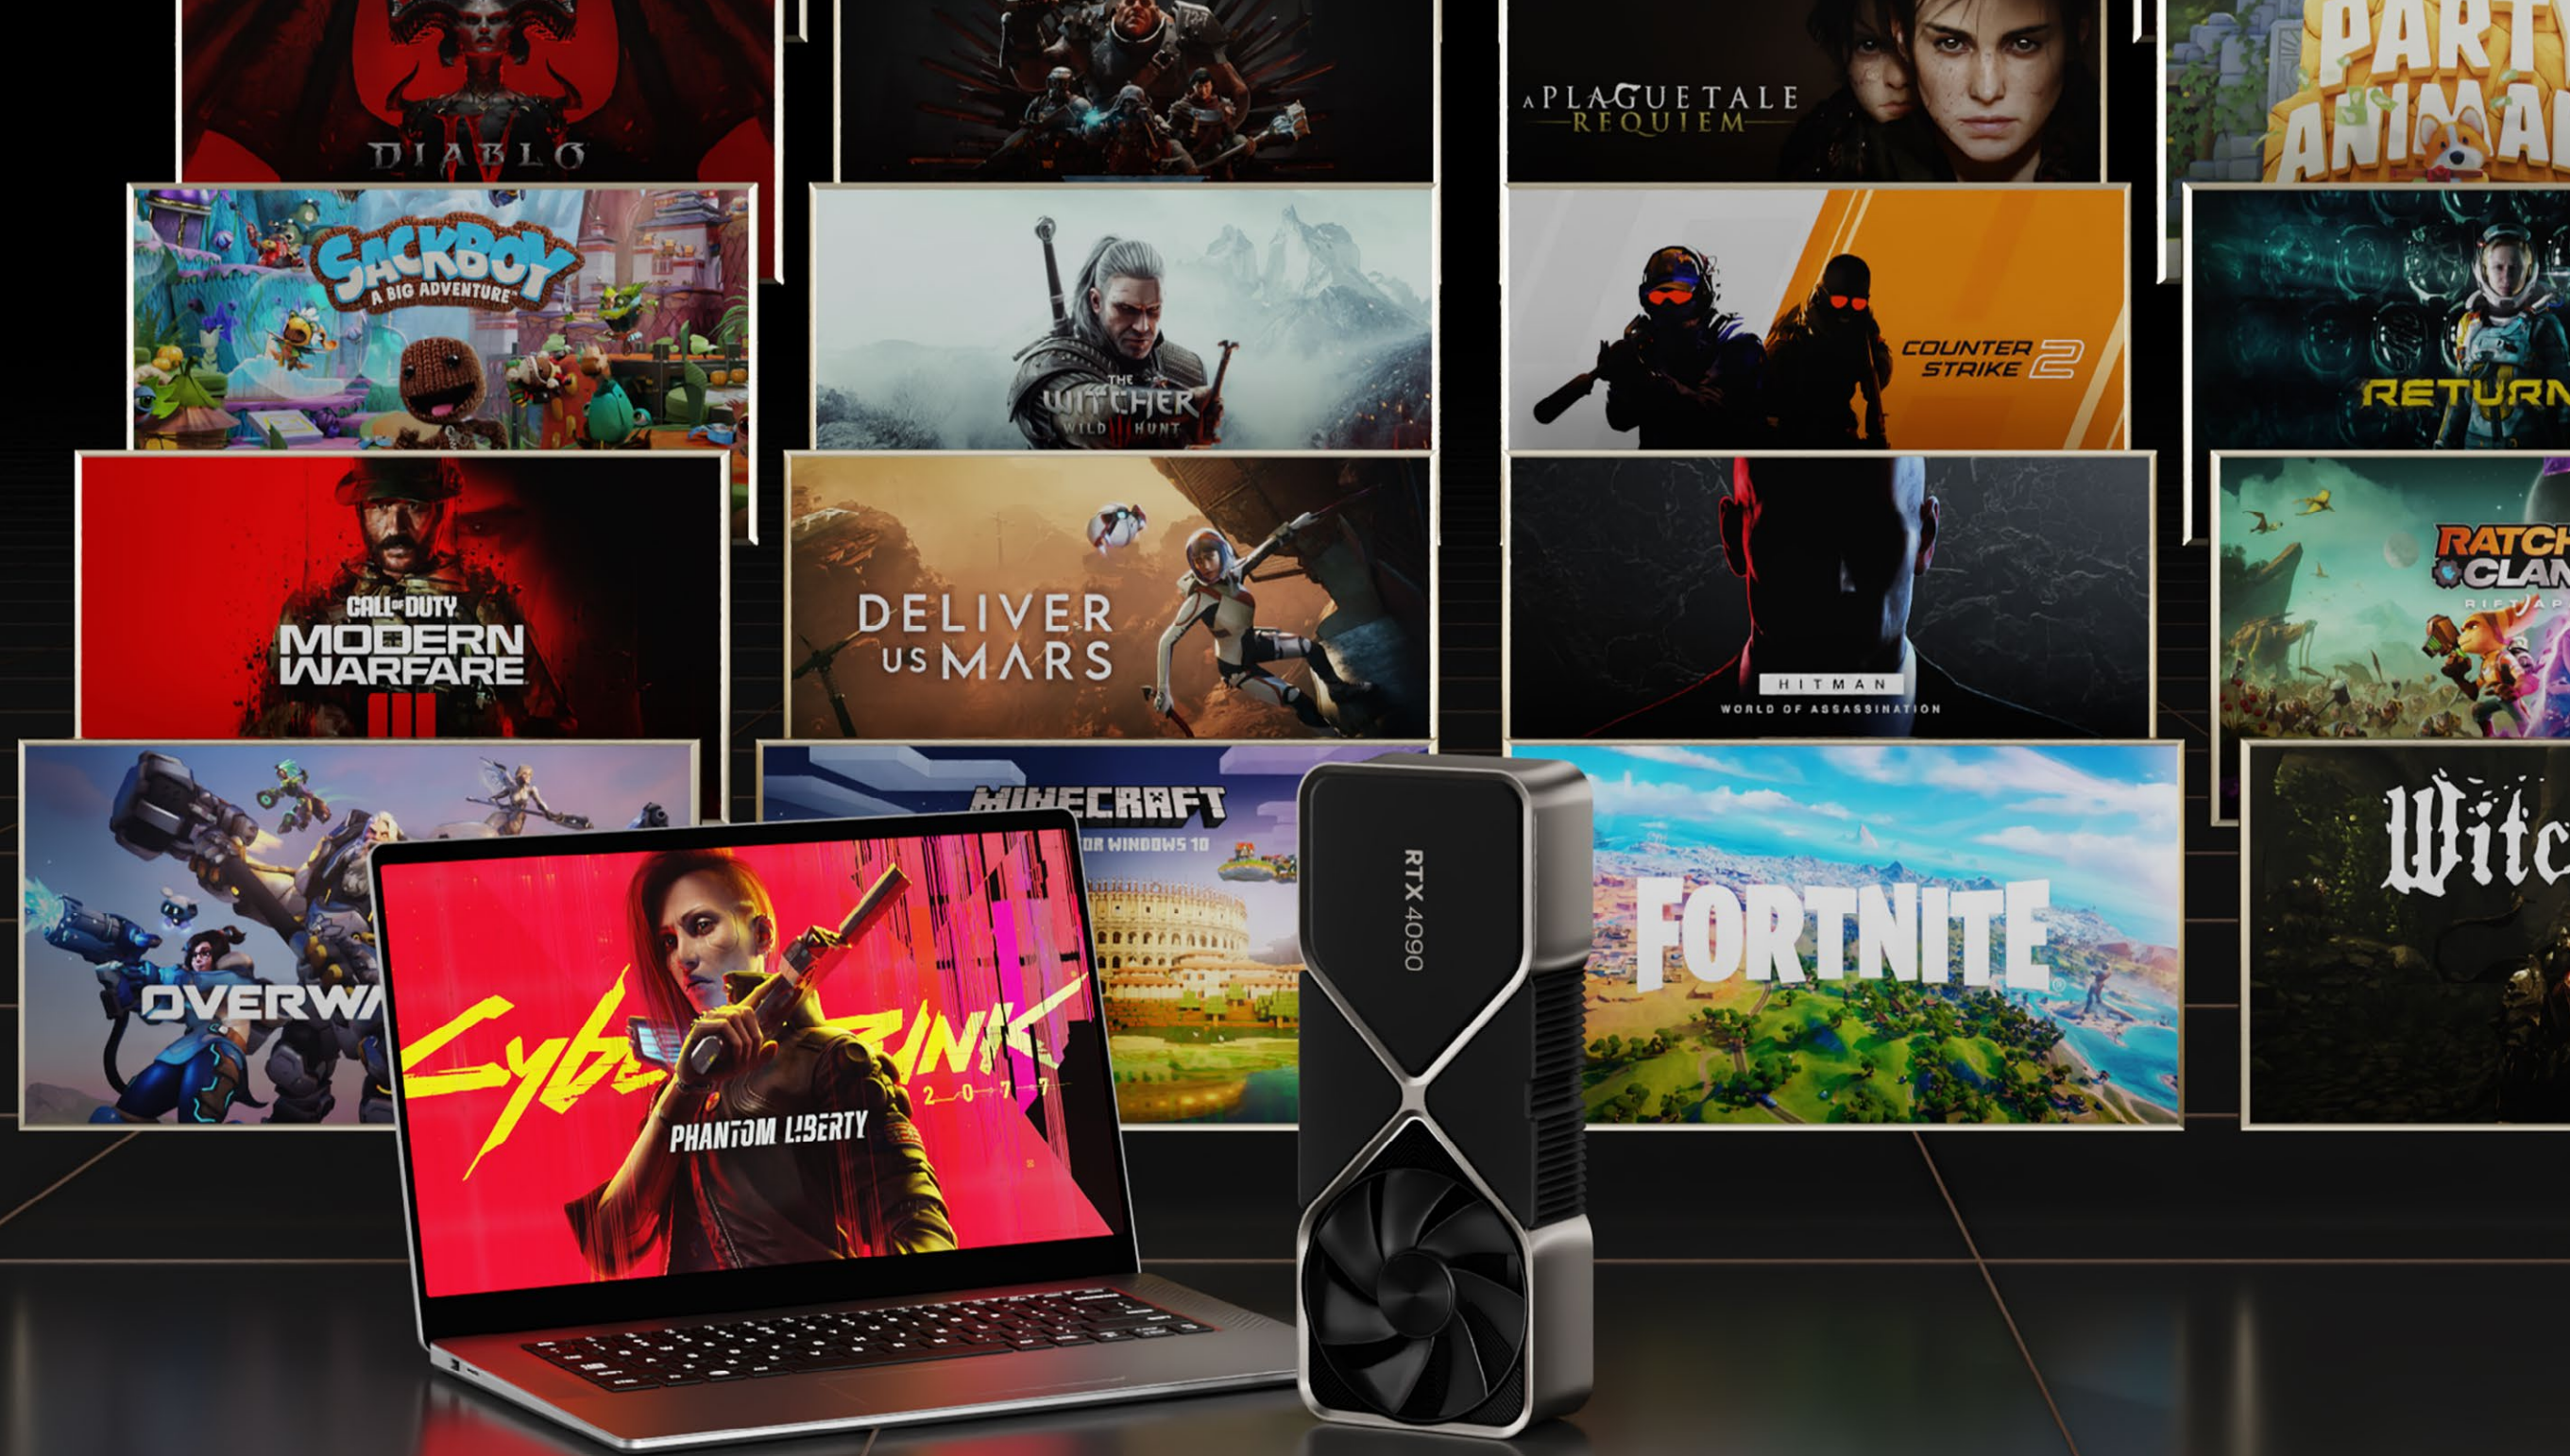
NVIDIA Reinvents Modern Graphics

We invented the programmable shading GPU nearly a quarter century ago, defining modern real-time computer graphics.

With NVIDIA RTX™ we have reinvented computer graphics again. This new rendering approach fuses rasterization and programmable shading with ray tracing and AI to make PC games look much more beautiful and realistic—almost cinematic.

NVIDIA RTX Resets Gaming

RTX is everywhere. More than 500 games and apps now use RTX to deliver stunning ray-traced graphics—including AAA blockbusters like Cyberpunk 2077, Fortnite, Minecraft, and more.



NVIDIA Cloud Gaming— Bringing RTX to Billions

With the power of NVIDIA® GeForce® GPUs in the cloud, GeForce NOW™ instantly transforms nearly any device into a powerful PC gaming machine. Any gamer can stream titles from the top digital game stores. Over 25 million members in 100+ countries now have access to more than 1,500 games.


And, recently, NVIDIA and Microsoft signed a 10-year deal to bring the Xbox PC game library to GeForce NOW.





NVIDIA Studio— Accelerated Computing Platform for Creators

Our industry-leading GPUs, paired with our exclusive driver technology and software, enhance creative apps with a level of performance and ability that is nothing short of inspiring. With NVIDIA Studio, creators are free to realize their most ambitious projects yet.




NVIDIA Supercharges Healthcare



From medical imaging to drug discovery, genomics to patient monitoring, researchers across life sciences are fusing traditional simulations and AI to solve the next grand challenges.

Generative AI Is Transforming the Pharmaceutical Industry

Drug discovery is a nearly \$2 trillion industry with \$250 billion dedicated to R&D. The industry is now jumping onto generative AI to discover disease targets, design novel molecules or protein-based drugs, and predict the behavior of medicines in the body.


NVIDIA BioNeMo provides state-of-the-art generative AI models for drug discovery, available from the cloud. And AI-powered medical devices can help clinicians detect and measure anomalies, up-level surgical skills, enhance image quality, and optimize workflows.





NVIDIA Drives Industrial Digitalization

NVIDIA Omniverse, our platform for industrial digitalization, builds virtual representations of physical things and assets—creating digital twins and connecting digital and physical worlds. Omniverse enables industries grounded in physical processes to become software-defined, realize unified digitalization, and connect large, highly skilled teams.




NVIDIA
Omniverse

USD COMPOSER




Design
Digital Twin

MODULUS




Performance
Digital Twin

DRIVE



AV
Digital Twin

METROPOLIS



Factory
Digital Twin

ISAAC



Robotics
Digital Twin

METROPOLIS



Warehouse
Digital Twin

Connecting Our Physical and Digital Worlds


As AI makes the leap to heavy industry, it needs to understand how to automate, design, navigate, and build based on the physics of our world. Digital twins via NVIDIA Omniverse enable AI to learn in a digital format. Optimizing changes virtually before deploying them reduces costs and speeds deployment.

NVIDIA DRIVE—Full Stack Autonomous Driving Platform


The NVIDIA DRIVE® family of products for autonomous vehicle development covers everything from the car to the data center.



DGX Data Center



Hyperion™ 8 Development Car



Omniverse Digital Twin



NVIDIA DRIVE Sim Turbocharges Developer Productivity to Get Self-Driving Cars on the Road

With NVIDIA DRIVE Sim™, features such as road elevation, road markings, islands, traffic signals, signs, and vertical posts are replicated at centimeter-level accuracy. Autonomous vehicles can drive millions of miles in a wide range of simulated scenarios so they hit the road running, safely.

Building the World's Most Advanced, Software-Defined Vehicles

This year, every next-generation Mercedes-Benz vehicle will include this first-of-its-kind software-defined computing architecture that includes the most powerful computer, system software, and applications for consumers. This marks the turning point of traditional vehicles becoming high-performance, updateable computing devices.






NVIDIA Omniverse Powers the Robotics Revolution

From smart automation in manufacturing to last-mile delivery, robots are becoming more ubiquitous in everyday life. The Isaac module in Omniverse is our platform for accelerating and enhancing robotics—from development to simulation to deployment.

Amazon Robotics Builds Digital Twins of Warehouses in NVIDIA Omniverse

Amazon has over 200 robotics facilities that handle millions of packages each day. Using NVIDIA Omniverse™ Enterprise and Isaac Sim™, Amazon Robotics is building AI-enabled digital twins of its warehouses to better optimize warehouse design and flow, and train more intelligent robotic solutions.





BMW Blends Reality and Virtual Worlds to Build Factory of the Future

BMW Group is using NVIDIA Omniverse to build a fully functioning factory digital twin before building it in the real world. Using NVIDIA AI and Omniverse has saved 20% on its factory fleet orchestration and planning.

An aerial photograph of a river delta, showing a large body of water branching into several smaller channels. The water is a deep blue, and the surrounding land is a lighter, sandy brown. A white rectangular text box is overlaid on the right side of the image.

NVIDIA's Earth-2 Initiative Aims to Accelerate Climate Research


Our Earth-2 initiative will be a digital twin of Earth. Simulating the world will help predict the complex multi-physics of Earth's atmosphere, land, sea, and ice caps at sufficiently high resolution. This will enable us to better predict the regional impacts of human actions over decades.

NVIDIA Is a Learning Machine

NVIDIA is united by a unique culture—the operating system of our company. We dream big, take risks, and learn from our mistakes together. Speed is key to our success. Craftsmanship is a passion. There are no org charts—the mission is the boss.

These beliefs inform everything we do, from designing amazing products to building one of the world's great companies—a place where people can do their life's work.





We're One Team Tackling Challenges No One Else Can Solve

NVIDIA employees are dedicated to building technology that moves humanity forward and to supporting the communities in which they work and live.

We've been recognized as a top company in social responsibility, and our employees are passionate donors to hundreds of charities around the globe.



**“Best Places to
Work in 2023”**

Glassdoor

**“100 Best Companies
to Work For”**

Fortune

**“Most Innovative
Companies”**

Fast Company

**“World’s Best
CEOs”**

Barron’s

**“World’s Best
Performing CEO”**

Harvard Business Review

**“50 Smartest
Companies”**

MIT Tech Review





Photo by Jason O'Rear