



NVIDIA

Sustainability Report

Fiscal Year 2026



Table of Contents

Message From Our CEO

Introduction

FY26 Highlights	06
About NVIDIA	07
Sustainability Governance	08
Our Impact	09

Energy, Efficiency, and Climate

Product Energy Efficiency	11
Greenhouse Gas Emissions	12
Climate Risks and Opportunities	12
Energy and Environment	13

Our People

Recruitment	15
Benefits and Compensation	15
Pay and Promotion	16
Engagement and Retention	16
Learning and Development	17
Health and Safety	17
Community Resource Groups	18

Product Value Chain

Responsible Sourcing	19
Product Environmental Impact	21
Product Quality and Safety	22

Responsible Business

Our Code	23
Trustworthy AI	24
Human Rights	26

Sustainability Indicators

U.N. Sustainable Development Goals	39
------------------------------------	----

About This Report

Message From Our CEO

We are living in a historic moment. The era of AI has arrived — and with it, a new industrial revolution is underway.

AI is more than a model. It is infrastructure, like electricity and the internet — a foundation on which every industry, every scientific discipline, and every economy will be transformed. Every company will use it. Every country will build it.

AI is not weightless. It may feel digital, but it is deeply physical. Every token is produced by real systems: energy, chips, cooling, networking, software, supply chains, and people. As AI becomes essential infrastructure, we have a responsibility to build it efficiently, deploy it responsibly, and use it to solve important problems.

You can think of AI as a five-layer cake: energy, chips and systems, infrastructure, models, and applications. Energy is the foundation. Applications are where AI becomes useful — better healthcare, faster science, smarter infrastructure, climate resilience, productivity, and expanded human capability.

We are still early. The opportunity is enormous. So is the responsibility.

AI Begins with Energy

AI factories turn electricity into intelligence. Their output is bounded by the power available to them and by how efficiently that power is converted into computation. As AI expands from training to inference, and from chatbots to reasoning agents that plan, use tools, and act, demand for compute will continue to grow — and so will the energy demand.

This year, we announced partnerships with Emerald AI, EPRI, and other organizations in the energy sector to create flexible AI factories that can connect to the grid faster, generate intelligence at scale, and act as responsive grid assets — supporting the broader energy system, not just drawing from it. Powered by the NVIDIA Vera Rubin DSX AI Factory reference design, these AI factories can help unlock up to 100 gigawatts of capacity across the U.S. power system while supporting grid reliability.

The energy challenge is real. It is solvable. The path is engineering: better architecture, higher utilization, liquid cooling, grid-aware operations, renewable energy, and deep collaboration across computing, energy, and infrastructure.

Compute Efficiency is the North Star of AI

We must build AI factories where extreme performance and energy efficiency advance together. That requires designing the entire system — power, cooling, networking, software, and compute — with precision.

NVIDIA's greatest sustainability lever is architecture, and it starts with CUDA accelerated computing.

CUDA accelerated computing is a discontinuous step-function improvement in computation efficiency. By moving the right work from CPUs to GPUs, accelerated computing does more work with every watt of energy. Over two decades, CUDA has created a software ecosystem that compounds efficiency: developers create new algorithms, NVIDIA optimizes libraries, runtimes, compilers, and frameworks, and those improvements flow across the installed base. Customers do not just get performance on day one. They get continuous performance gains and energy efficiency over the life of the system.





The performance and efficiency of our computing platforms determine how much intelligence can be produced from every watt.

CUDA was the beginning. Today, NVIDIA architecture extends across the entire AI factory — GPUs, CPUs, LPUs, DPUs, networking, switching, storage, cooling, power, systems, and software. We do not optimize a single component. We optimize the system — chips, racks, networking, cooling, algorithms, software, and the AI factory around them. We do extreme co-design.

This year, the NVIDIA Vera Rubin platform marks a generational leap: seven breakthrough chips, five racks, one giant supercomputer built to power every phase of AI.

Vera Rubin NVL72 delivers up to 10x energy efficiency over the previous NVIDIA Blackwell architecture, enabling AI factories to generate significantly more intelligence per watt. Its warm-water, 45°C direct liquid cooling reduces cost, complexity, energy use, and water consumption.

Efficiency compounds. Better chips reduce energy per token. Better systems reduce data movement. Better cooling reduces waste. Better software increases utilization. Together, they make AI factories more productive, more economical, and more sustainable.

AI Expands Human Capacity – Creating New Jobs, Companies, and Industries

AI is infrastructure for the modern economy — transforming healthcare, energy, manufacturing, financial services, retail, transportation, and beyond.

The AI buildout is creating demand for electricians, plumbers, steelworkers, network technicians, installers, operators, and the people who will build and run the infrastructure of this new era.

At the same time, AI is expanding capacity across the knowledge economy. Agentic AI systems can reason, plan, use tools, and complete complex tasks. They are helping engineers write software, financial analysts study markets, customer service teams resolve issues, and clinicians reduce administrative burden.

This is not simply automation. It is capacity expansion.

Every job has a purpose, and that purpose is larger than its tasks. Coding is a task; innovation is the purpose. Reading scans is a task; caring for patients and diagnosing disease is the purpose. When AI helps with routine work, people can spend more time on judgment, creativity, communication, and care.

AI elevates work so people can pursue a higher purpose.

Lowe's is using AI-powered digital twins to improve operations across more than 1,750 stores. PepsiCo achieved a 20% throughput increase in initial deployments by simulating manufacturing and supply chain operations with physics-accurate digital twins built with Siemens and NVIDIA.

This is sustainability in practice: better planning, less waste, higher productivity, safer operations, and more useful work from the same resources.

Transforming Healthcare and Life Sciences With AI

AI is transforming medicine — from drug discovery and cancer detection to patient care in the most remote corners of the world.

Eli Lilly launched LillyPod, the world's most powerful AI factory owned by a pharmaceutical company, powered by an NVIDIA DGX SuperPOD with more than 1,000 NVIDIA Blackwell Ultra GPUs and delivering over 9,000 petaflops of AI performance. Assembled in just four months, LillyPod is accelerating genomics, drug discovery, and medical advancement at a scale never before possible.

NVIDIA Inception company Droplet Biosciences is using NVIDIA Parabricks to accelerate residual cancer detection, reducing analysis from 10 days to under five. NVIDIA Inception partner MedCognetics is bringing AI-powered breast cancer screening to rural India through mobile clinics, reaching thousands of patients, many of whom had never had a mammogram.

These are not abstract examples. They are faster science, earlier detection, broader access, and more capacity for the people delivering care.

Accelerating Climate Action With AI

Climate resilience is also a computing challenge.

Weather, oceans, coastlines, power grids, agriculture, and ecosystems are dynamic systems. They change continuously. They interact with one another. And the decisions that depend on them — where to build, when to prepare, how to protect communities, how to operate energy systems — are increasingly local and time-sensitive.

AI gives scientists, governments, and industries a new instrument to understand these systems faster, at higher resolution, and with more useful lead time.

NVIDIA Earth-2 is the world's first fully open, accelerated AI weather and climate platform.



The Earth-2 family of models makes weather AI accessible worldwide, from processing initial observation data to generating 15-day global forecasts or daily local storm forecasts. Earth-2 models can generate initial atmospheric conditions in seconds, work that once required hours on supercomputers, and produce forecasts up to 500x faster than traditional methods.

At UC Santa Cruz, researchers are using NVIDIA GPU-accelerated simulations to map coastal flooding and model the protective benefits of coral reefs and mangroves on a global scale.

AI can help us compress time, increase resolution, reduce uncertainty, and act sooner.

AI is a Catalyst to Upgrade Energy Grids and Invest in Sustainable Energy

AI demand and infrastructure investments are powerful market incentives to modernize the energy grid and invest in sustainable power generation. AI is the best opportunity in generations to invest in sustainable energy and finance new generation, storage, transmission, and grid software, pulling forward investment to modernize the grid.

We are investing in energy innovators across this future — from Emerald AI for grid flexibility and Redwood Materials for battery materials and storage, to companies working on longer-term clean power. AI will require more energy, and it can also accelerate the investment needed to build a more resilient and sustainable energy system.

Our Commitment to Sustainability and Community

We must also continue to improve how we operate.

For the second consecutive year, NVIDIA matched 100% of our global electricity consumption with clean energy. We continue to work to reduce emissions, improve efficiency across our operations, and design platforms like Vera Rubin so that powerful computing and sustainability advance together.

Our responsibility extends beyond our own footprint. Helping our communities has always been at the heart of NVIDIA's culture. In fiscal year 2026, NVIDIA's donated more than \$34 million in charitable giving and logged more than 122,000

volunteer hours to support more than 11,000 nonprofits in more than 70 countries.


The AI industrial revolution is not simply a technology story. It is a human story — about how we use AI to create jobs, cure diseases, accelerate climate solutions, strengthen communities, and expand what people can do.

AI is a tool. Like every great tool in history, what matters most is what we choose to build with it.

AI factories are here. A new industry is rising. It is a privilege to contribute to this moment — and a responsibility we carry with humility and determination.

We are making history together.

Jensen Huang,
Founder and CEO, NVIDIA



Introduction

FY26 Highlights



About NVIDIA

42,000

Employees globally

38

Countries NVIDIA
operates in globally

\$64M

Employee donations and company
matching through NVIDIA Foundation



Energy, Efficiency, and Climate

100%

Clean electricity purchased or
generated for our operations¹

35x

Inference per watt (TPS/MW) for
NVIDIA Vera Rubin + LPX compared
to NVIDIA Blackwell²

Top 9

Supercomputers on the
November 2025 Green500 list
are accelerated by NVIDIA



Our People

#1

On Glassdoor's Best Places
to Work in Tech & AI list

1,200

Employees participated in our
CRG mentorship program

3.7%

Overall turnover rate, compared to the
technology industry average of 21.9%



Product Value Chain

24%

Reduction of embodied carbon emissions
across large AI workloads between
NVIDIA HGX™ B200 and HGX H100

93%

Strategic suppliers audited
in the past two years

96%

NVIDIA GPU systems packaging was
recyclable materials by weight



About NVIDIA

NVIDIA pioneered accelerated computing to help solve the most challenging computational problems. NVIDIA is now a data-center-scale AI infrastructure company reshaping all industries.

Our technology stack includes the foundational NVIDIA® CUDA® development platform that runs on all NVIDIA GPUs, as well as hundreds of domain-specific software libraries, frameworks, algorithms, software development kits (SDKs), and application programming interfaces (APIs). This deep and broad software stack accelerates the performance and facilitates the deployment of NVIDIA Accelerated Computing for computationally intensive workloads, such as artificial intelligence, model training and inference, data analytics, scientific computing, robotics, and 3D graphics, with vertical-specific optimizations to address industries ranging from healthcare and telecom to automotive and manufacturing.

Introduced with the NVIDIA Blackwell architecture, our data-center-scale offerings feature extreme codesign where the infrastructure's chips, networking, systems, software, and algorithms are holistically architected and optimized to maximize performance and scale.


Hundreds of thousands of GPUs can be interconnected to function as a single giant computer. This type of data center architecture and scale are needed for the development and deployment of modern AI and accelerated computing applications.

The GPU was initially used to simulate human imagination, enabling the virtual worlds of video games and films. Today, it also simulates human intelligence, unlocking a deeper understanding of language, science, and the physical world. Its parallel processing capabilities, supported by tens of thousands of computing cores, are essential for deep learning algorithms. This form of AI, in which software writes itself by learning from large amounts of data, can serve as the brain of computers, robots, and self-driving cars that can perceive, understand, and reason about the world. GPU-powered AI solutions are being developed by thousands of enterprises to deliver services and products that would have been immensely difficult or even impossible with traditional coding. Examples include generative AI, which can create new content such as text, code, images, audio, video, molecule structures, and recommendation systems; and agentic AI, where systems of AI models work in concert to automatically complete a task.

NVIDIA has a platform strategy, bringing together hardware, systems, software, algorithms, libraries, AI models and training data sets, and services to create unique value for the markets we serve. Although the computing requirements of these end markets are diverse, we address them with a unified underlying programmable architecture that lets us support several multi-billion-dollar end markets with the same underlying technology, using a variety of software stacks developed internally or by third-party developers and partners. The large and growing number of developers and installed base across our platforms strengthens our ecosystem and increases the value of our platform for our customers.

Innovation is at our core. We have invested over \$76.7 billion in research and development since our inception, yielding inventions that are essential to modern computing. Our invention of the GPU in 1999 sparked the growth of the PC gaming market and redefined computer graphics. With our introduction of CUDA in 2006, we opened the parallel processing capabilities of our GPU to a broad range of compute-intensive applications, paving the way for the emergence of modern AI. In 2012, the AlexNet neural network, trained on NVIDIA GPUs, won the ImageNet computer image recognition competition, marking the

“Big Bang” moment of AI. We introduced our first Tensor Core GPU in 2017, built from the ground-up for the new era of AI, and our first autonomous driving system-on-chip (SoC) in 2018. Our acquisition of Mellanox in 2020 expanded our offerings to include networking, enabled our platforms to be data center scale, and led to the introduction of a new processor class—the data processing unit, or DPU.





Over the past five years, we have built full software stacks that run on top of our GPUs and CUDA to bring AI to the world’s largest industries, including NVIDIA DRIVE™ for autonomous driving, NVIDIA Clara™ for healthcare, NVIDIA Omniverse™ for physical AI applications, and NVIDIA AI Enterprise software—essentially an operating system for enterprise AI applications.

In 2023, we introduced our first data center CPU, NVIDIA Grace™, built for giant-scale AI and high-performance computing (HPC). In 2024, we launched the NVIDIA Blackwell architecture—connecting 36 Grace CPUs and 72 Blackwell GPUs in a data-center-scale, liquid-cooled design—for real-time trillion-parameter inference and training. In 2025, we launched and scaled the NVIDIA Blackwell Ultra platform, optimized for agentic, reasoning, and physical AI. Building on the architectural breakthroughs of NVIDIA Blackwell and leveraging NVIDIA Dynamo inference software, it delivers a significant increase in token throughput and reduction in cost per token compared to the NVIDIA Hopper™ generation.

More recently, in support of market development, we have accelerated the release cadence of our open AI model platforms, including NVIDIA Nemotron™ for agentic AI and NVIDIA Cosmos™ for physical AI.

With a strong engineering culture, we drive fast, yet harmonized, product and technology innovations in all dimensions of computing, including silicon, systems, networking, software, and algorithms.

NVIDIA’s sustainability reporting follows our fiscal calendar; FY26 data corresponds to the period January 27, 2025, to January 25, 2026.

Sustainability Governance

We assess corporate sustainability topics to identify our key social and environmental impacts, risks, and opportunities, using internal expertise, stakeholder expectations, and market trends in the process. Our double materiality assessment analyzed sustainability issues based on financial and impact materiality.

We determine the sustainability topics to include in our reporting based on our impacts, stakeholder expectations, and frameworks such as the Global Reporting Initiative (GRI), Sustainability Accounting Standards Board (SASB), the Task Force for Climate-Related Financial Disclosures (TCFD), and the United Nations Sustainable Development Goals (UN SDGs). These topics reflect our values, social and environmental contributions, and the matters we believe are important to our continued business success.

Board Oversight

NVIDIA’s Board of Directors (Board) oversees risk management at NVIDIA and delegates oversight of appropriate topics to its committees. The oversight responsibility of our Board and its committees is enabled by management reporting processes, including our enterprise risk management (ERM) process, which are designed to provide visibility to our Board about the identification, assessment, and management of critical risks and management’s risk mitigation strategies. For more information on enterprise risk oversight, please refer to NVIDIA’s annual [Proxy Statement](#).

The Board’s Nominating and Corporate Governance Committee (NCGC) reviews and discusses with management our governance of and reporting related to sustainability, including strategy, risks and opportunities, climate targets, and related programs and initiatives. NVIDIA’s Head of Sustainability updates the NCGC at least semiannually and the full Board annually on these topics.

The Board’s Audit Committee has direct oversight of our overall ERM and external auditors engaged for the limited assurance of select sustainability metrics.

Management Oversight

At the executive level, our Corporate Sustainability Steering Committee (CSSC) reviews and supports the effectiveness of our sustainability strategy and program. Our CSSC is composed of members of our executive leadership team, including Legal, Finance, Operations, and Human Resources. Feedback from the Board, the NCGC, and CSSC, along with specific input from our executive team, helps determine the focus and scope of our sustainability strategy and program.





Our Head of Sustainability reports to the CSSC while the Corporate Sustainability (CS) team manages the day-to-day implementation of our sustainability program, including monitoring targets, assessing risks and opportunities, and reporting.

Our Impact

Accelerated computing and artificial intelligence are enabling a renaissance of innovation across industries. By offering unprecedented levels of computing and intelligence, NVIDIA is helping address some of the world's most difficult environmental and social challenges—from climate resiliency and education accessibility to revolutionary drug discovery.

Energy and Climate

More organizations are discovering and utilizing AI's potential to address environmental challenges. Researchers are using accelerated computing to [map](#) coastal flooding and evaluate storm adaptation solutions. Customers are accelerating research and building energy-efficient AI factories [running](#) on renewable energy.

AI can support improved grid flexibility, energy availability, and reliability by optimizing grid infrastructure.

Today's electric systems are built to serve peak demand but are underutilized during most hours of the day. Power-flexible AI factories can help [unlock](#) up to 100 gigawatts of capacity across the U.S. power system by combining optimized infrastructure design with efficient use of existing assets and, where needed, new-build generation, while flexing during limited periods of grid stress to reduce the need for broader grid expansion to support reliability.

NVIDIA and Emerald AI are [working](#) with leading energy companies to pioneer flexible AI factories as grid assets. These next-generation AI factories will harness the new NVIDIA Vera Rubin DSX AI Factory reference design, which includes the DSX Flex software library for connecting AI factories to power-grid services. For accelerated deployment, the factories can use co-located energy generation and storage as bridge power for hybrid AI factories, then later harness these resources to flexibly supply the grid, accelerate AI factory interconnection, and support the broader power system. This approach helps bring AI capacity online faster while creating broader value for customers and communities.

AI also helps weather and climate prediction be more accessible than ever for scientists,

startups, developers, enterprises, and government agencies worldwide, offering the potential to improve climate resiliency. NVIDIA's [Earth-2](#) is the first [open](#), accelerated set of models and tools that enables developers to bring disparate weather and climate AI capabilities together. Accurate weather forecasting helps save lives and

protect environments—and is a cornerstone of decision-making in agriculture, energy, public health, and other industries. Earth-2 models are improving weather forecasting, advancing energy forecasting and resiliency of grid operations, and informing climate risk assessments.



Education

Researchers, educators, and students are at the forefront of innovation, pushing the boundaries of what's possible in academia and every industry. NVIDIA's cutting-edge technologies, available both on premises and in the cloud, are designed to support their most ambitious projects. University researchers are currently leveraging AI to advance Parkinson's disease detection and digital twins to improve emergency responses in urban areas and simulate autonomous vehicles for safe testing.

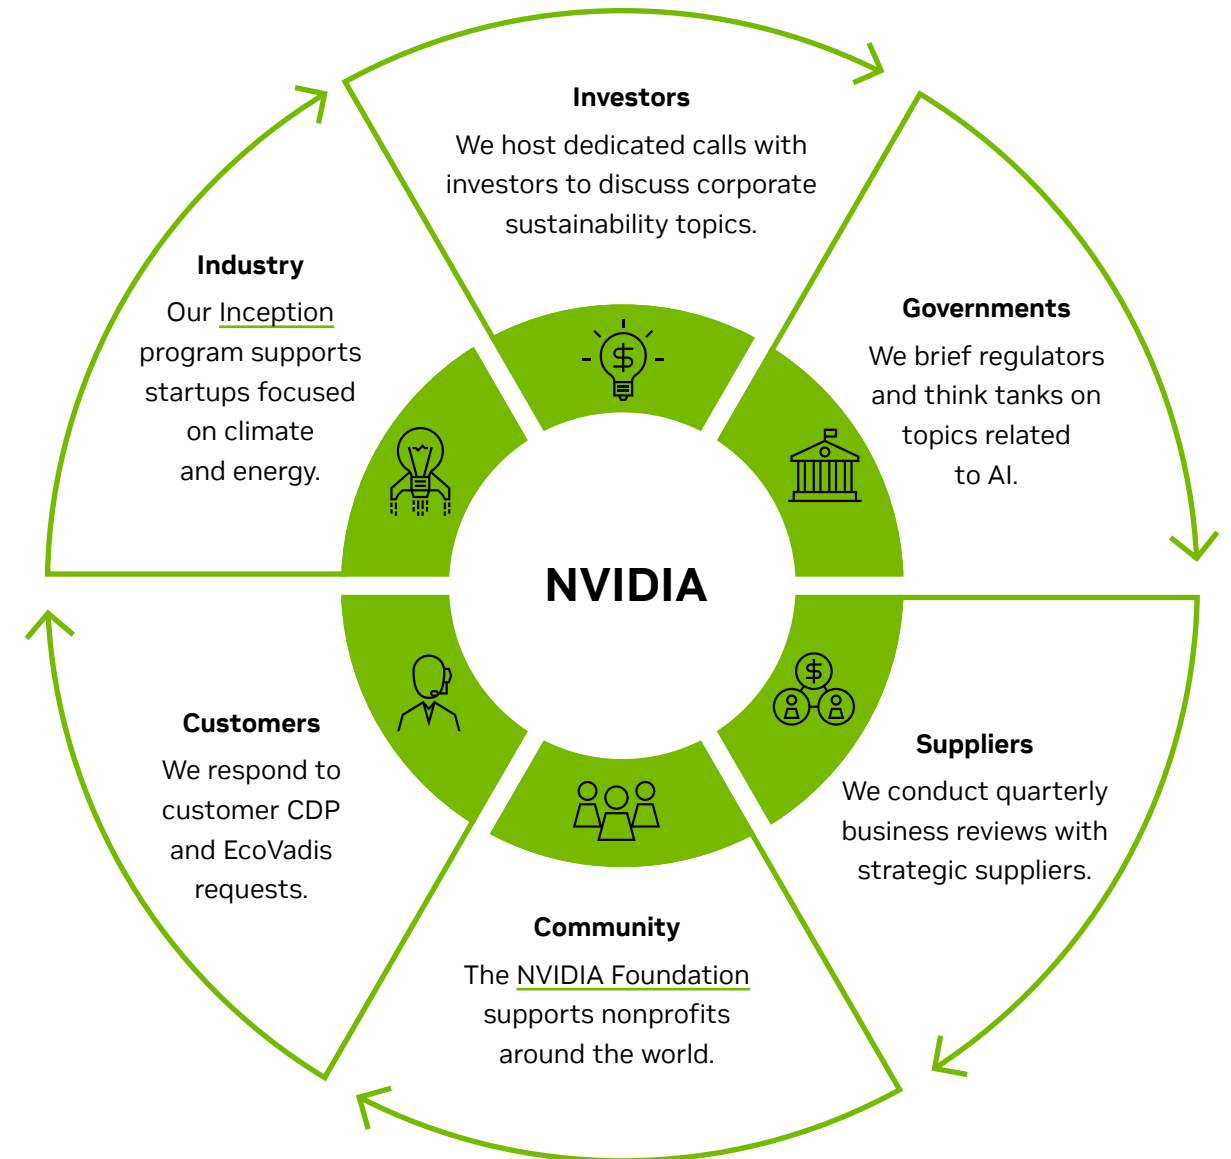
To empower early education, NVIDIA pledged \$25 million in support with AI education programs, partnering with two leading K-12 learning platforms to tailor the NVIDIA Deep Learning Institute (DLI) and NVIDIA Academy content to meet the instructional needs of U.S. K-12 classrooms.

AI is also a powerful tool for accessibility. For example, NVIDIA partnered with the American Society for Deaf Children and Hello Monday to develop Signs, an AI-enabled platform that helps teach American Sign Language. NVIDIA is also collaborating with cities, states, and educational institutions to enhance AI education efforts and boost regional economic development.

Healthcare

NVIDIA technology enables healthcare institutions and researchers to harness the power of AI and HPC to improve patient care, accelerate research, and unlock medical breakthroughs. AI is driving a quantum leap forward in life sciences and healthcare, for drug discovery, healthcare robotics, medical imaging, genomics, and digital health.

For example, Lilly launched a new AI factory with over 1,000 NVIDIA Blackwell Ultra GPUs to help its teams make meaningful medical advancements faster, more accurately, and at unprecedented scale. With this foundation, Lilly teams can analyze genomes, explore billions of chemical possibilities, and apply AI across clinical development and manufacturing to design better trials, optimize production, and accelerate decision-making. Together, these capabilities enable faster, more precise, and more scalable creation and delivery of medicines. Lilly has announced plans for its new AI supercomputing infrastructure to run on 100% renewable electricity by 2030, using efficient liquid cooling and minimal incremental energy impact.



Our approach to stakeholder engagement.

Energy, Efficiency, and Climate

Product Energy Efficiency

Improving performance and energy efficiency is a principal goal in each step of our research, development, and design processes. We develop hardware, software, and networking technology to improve performance and energy efficiency, helping our customers solve problems and decrease emissions intensity.

Frontier AI models continue to increase in complexity and size as they enhance generative AI and enable scientific discoveries, and modern data centers and AI Factories require accelerated computing platforms to effectively execute these workloads.

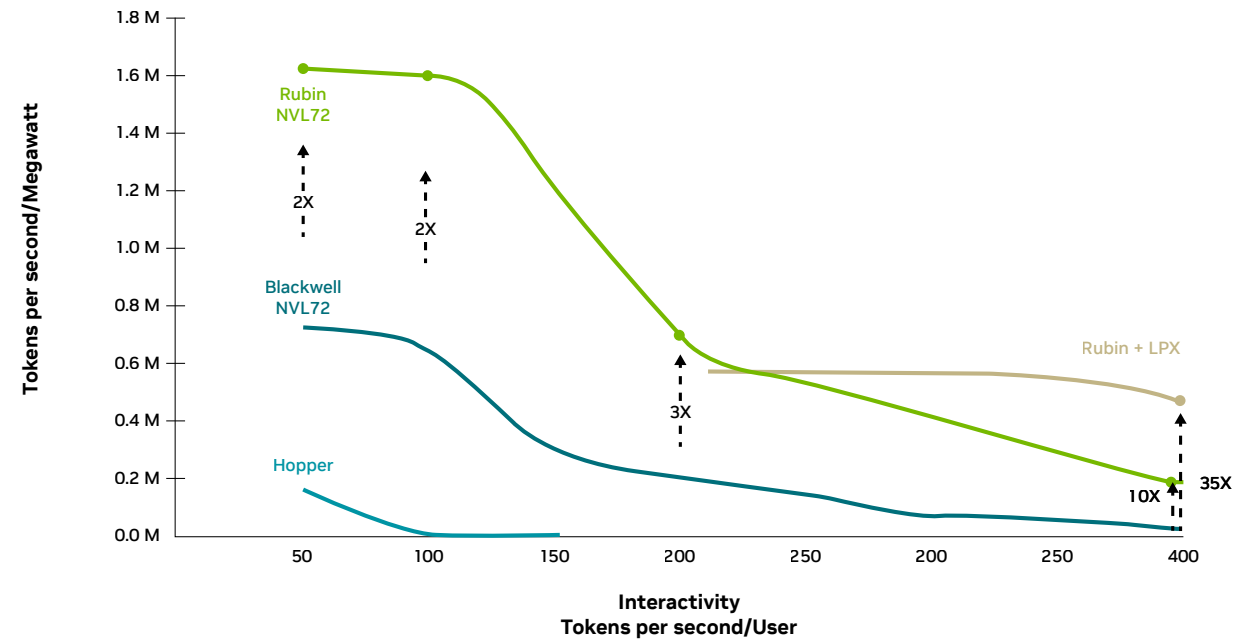
Accelerated computing is the use of specialized hardware to dramatically speed up computation, often using parallel processing to bundle frequently occurring tasks. It offloads demanding work that can bog down traditional processors that typically execute tasks in serial fashion and can offer higher performance and greater energy efficiency. Acceleration completes larger workloads more quickly, returning the computer to a low-power idle state when done and consuming significantly less energy overall for similar tasks when compared to traditional computing.

NVIDIA GPUs can run AI workloads that can't be completed with traditional computing. While our latest architecture, NVIDIA Rubin, has demonstrated up to 10x energy efficiency over the previous architecture,³ software updates continue to improve the performance of older products as well. For additional information about methodology behind comparisons in this section, please refer to Endnotes.

The NVIDIA computing stack is built for performance and energy-efficient compute:

- › NVIDIA Vera CPU: Enables software environments to run up to 50% faster with 2x the efficiency of traditional CPU infrastructure, unlocking agentic AI at scale.⁴
- › NVIDIA Groq 3 LPX + Vera Rubin NVL72: delivers 35x inference performance per watt for trillion-parameter models relative to Blackwell (GB200 NVL72).²
- › NVIDIA® BlueField® DPU: Offloads process-intensive tasks from the CPU, and our STX reference architecture provides up to 1.5x energy efficiency compared to the latest x86 based storage offering.⁵

Accelerated Computing is Energy Efficient Computing Large Parameter MOE



Simulation workload efficiency is equally important. The HPC and supercomputing community use the Green500 benchmark to measure high-precision (FP64) efficiency, and NVIDIA supercomputing systems top the leadership board, with nine of the top ten systems accelerated by NVIDIA technologies.

Top 9
Supercomputers on the November 2025 Green500 list are accelerated by NVIDIA.



Traditional data centers heavily rely on air cooling, which consumes significant energy to move and condition air. The NVIDIA Vera Rubin platform uses warm-water, single-phase direct liquid cooling (DLC) with a 45-degree Celsius supply temperature. Using liquid instead of air has many positive impacts inside and outside the data center, including higher performance per rack, reduced water consumption for cooling through the use of closed-loop heat rejection systems, and the ability to run data centers at higher ambient air temperatures, which further reduces energy consumption. This translates to significant cost, complexity, and energy savings compared to air cooling.

Rack and downtime inefficiencies can also impact the amount of grid power that actually reaches compute. Many operators size for worst-case peak rack power, leaving unused capacity usable for GPUs. Operating too close to thermal or electrical limits drives hardware faults, job restarts, and downtime. The NVIDIA Vera Rubin DSX AI Factory reference design [outlines](#) how to design, build, and operate the entire AI factory infrastructure stack. DSX MaxLPS helps ecosystem partners operate AI factories with up to 30% more GPUs within the same power envelope and higher throughput per watt. DSX Flex connects AI factories to power-grid services, enabling them to dynamically adjust power use and orchestrate

demand with hybrid onsite generation to save energy and maintain grid stability.

Greenhouse Gas Emissions

We assess our carbon footprint across our product lifecycle and climate risks, including regulations and market impacts.

To manage the greenhouse gas (GHG) emissions footprint of our data centers, labs, and offices, we strategically focus on working with key data center partners to optimize efficiency, managing our operations efficiently, and sourcing clean energy.

Last year, we met our goal to purchase or generate enough clean electricity to match 100% of our global electricity usage for sites under our operational control. In FY26, we continued to match 100% of our global electricity usage with clean electricity, reducing Scope 2 market-based emissions.¹ NVIDIA defines clean electricity as electricity generated from non-fossil fuel sources, including solar, wind, and other low-carbon technologies. The term refers to emissions at the point of generation and does not address full life cycle impacts.

We report Scope 1 and 2 GHG emissions in accordance with the GHG Protocol. We engaged an external third party to perform a limited assurance engagement over select

metrics presented for FY26. Find NVIDIA's FY26 management assertion and the Report of Independent Accountants [here](#). In FY26, we continued to mature our emissions calculation methodology for our Scope 3 Category 1 emissions, incorporating supplier-specific activity data for key components, replacing a full spend-based approach. This approach replaces generic industry average data with precise, supplier-specific information, enabling us to identify key decarbonization levers within our supply chain.

In FY26, NVIDIA achieved:

100%

Clean electricity purchased or generated, reducing our Scope 2 market-based emissions.

Science-Based Targets

Science-based targets are a useful tool for measurement that helps to inform work towards climate ambitions. NVIDIA recently adopted two emission reduction targets. Both targets were validated by the Science Based Targets initiative (SBTi) as consistent with a reduction in greenhouse gas emissions in line with climate science.

NVIDIA's goal is to, from a FY23 base year:

- › Reduce absolute Scope 1 and Scope 2 market-based emissions 50% by FY30
- › Reduce Scope 3 emissions intensity from use of sold GPU products by 75% per PFLOP by FY30

Our Scope 3 target is measured in emissions per petaFLOPS (PFLOPS), which represents one quadrillion (10¹⁵) floating-point operations per second. This is a common industry standard for measuring computational performance capacity. NVIDIA's ambition is to reduce energy and emissions per computation. We view this as NVIDIA's biggest opportunity to reduce emissions and to support global sustainability efforts.

As we work towards our goals, we will provide updates on our progress.

Climate Risks and Opportunities

In FY26, we performed a climate scenario analysis that took into account several climate-related scenarios across our global operations (offices, labs), upstream value chain (leased data centers, supplier manufacturing sites, warehouses), and downstream product impacts. In the analysis, we considered short- (0-1 year), medium- (1-5 years), and long-term



(over 5 years) time horizons. Climate models involve significant inherent uncertainty, including methodological limitations, data constraints, estimates, and projections.

Our analysis of climate risks and opportunities was subject to assumptions and limitations, with greater uncertainty over the long term.

As a result, we interpret scenario outputs as indicative rather than predictive.

The scenario analysis did not reveal any additional material climate-related financial risks to our business in the scenarios and aspects of our business examined.

Our scenario analysis considered a range of potential climate outcomes, including a low-emissions pathway aligned with the Paris Agreement and a high-emissions trajectory, spanning from 2025–2100. The analysis drew on established physical risk scenarios (RCP 2.6, 4.5, 6.0, and 8.5), water-specific projections using the WRI Aqueduct Water Risk Atlas, and transition risk scenarios (current policies, delayed transition, and net zero 2050).

NVIDIA does not currently anticipate any material financial impact to our business arising from climate-related risks. Please refer to our [2026 Form 10-K](#) filed with the SEC for information on the adverse impact climate change may have on our business.

The company may also still be subject to risks broadly applicable to the whole economy. While we have not presently identified material climate-related financial risks, potential impacts are considered in our business strategy.

Energy and Environment

We're committed to reducing our environmental impact by driving operational excellence. We identify and control environmental impacts and continuously improve our performance using a comprehensive environmental management system certified to [ISO 14001](#), which covers our Santa Clara, California headquarters buildings and Yokneam, Israel offices. Our dedicated Environmental, Health, and Safety team executes our environmental [policies](#) and practices, with actionable goals and metrics that are annually reviewed with leadership.

To bring a more structured approach to managing energy efficiency at several of our key data center locations, we have an energy management system certified to ISO 50001. Our certification to the [ISO 50001](#) standard recognizes our efforts to properly measure and reduce energy consumption in our data centers.

Energy-Efficient Sites

Two of our headquarters buildings in Santa Clara and our Hyderabad, India campus are LEED Gold certified with solar panels to generate renewable electricity onsite. In FY26, we opened a new Shanghai, China campus that is LEED Gold certified. These buildings were designed to be energy and water efficient, built with sustainable raw materials and improved indoor environmental quality.

In FY26, we purchased or generated enough clean electricity to match 100% of our global electricity usage.¹ Our clean electricity sourcing includes on-site solar, long-term power purchase agreements, utility renewable electricity tariffs, and energy attribute certificates.

Waste Management

We aim to reduce the amount of waste we send to landfill through waste reduction, reuse, and recycling initiatives. In FY26, we worked on improving our waste data collection efforts by increasing our use of primary data and minimizing estimations.

For products we use for testing, R&D, and production purposes, we have programs in place to support internal reuse of equipment that hasn't reached the end of its useful life or financial depreciation life.



We properly dispose of IT assets, and used laptops are sold to a vendor for refurbishment and reuse, with a portion of the financial proceeds donated to the NVIDIA Foundation. For used equipment, we use a global specialist e-waste recycling vendor.

All our vendors follow strict certification guidelines and procedures to ensure proper tracking of the chain of custody, decommissioning, data erasure, and recycling of broken or unusable hardware.

Water Conservation

We use water in our direct operations for food service, landscaping, and sanitation. We take steps to use water efficiently, particularly in locations that are vulnerable to drought conditions. We also conduct an annual water risk assessment for all NVIDIA facilities and data centers, which helps identify which sites are located in water-stressed areas.

Our LEED Gold certified buildings incorporate a range of water-efficient measures. Our two LEED certified buildings at our Santa Clara headquarters include low-flow bathroom fixtures and the use of recycled water for toilet flushing, and the buildings are designed to achieve reductions in domestic water demand and potable water use for sewage conveyance when compared with buildings of standard design.



At our Santa Clara headquarters, our Voyager and Endeavor buildings are LEED Gold certified and connected by a three-acre park that's irrigated with reclaimed water and shaded by a trellis that houses solar panels.

We use reclaimed water in cooling towers and landscape irrigation systems and capture rainwater in bioswales. The three-acre park that connects the two buildings is planted with drought-tolerant trees, and the greenery is irrigated by reclaimed water.

Outside of Santa Clara, our Hyderabad, India, site has a water treatment plant that reuses the site's treated wastewater for landscape irrigation, and our Shanghai LEED certified buildings have low-flow pantry sinks and shower fixtures.

We're deploying closed-loop liquid cooling across our AI factories and leased data centers where applicable. By utilizing PG25 (25% propylene glycol with inhibitors) within an S45 (45°C warm-water) architecture, we capture nearly all heat directly from NVIDIA systems. This high-temperature design minimizes or eliminates the need for chillers and enables heat rejection via dry coolers and economizers. Dry coolers reject heat through facility water loops to ambient air, while economizers utilize two-phase refrigerant

systems to reject heat directly to ambient air, neither requiring water as part of the heat rejection process. This allows NVIDIA AI factory platforms to achieve industry-leading compute density per rack while operating efficiently across diverse climates with simplified, high-efficiency infrastructure and zero-water heat rejection capability. When fully implemented, this design can significantly reduce or eliminate water use for cooling compared to traditional evaporative systems.

Our People

Our employees are our greatest asset and play a key role in creating long-term value for our stakeholders. To execute our business strategy successfully, we focus on recruiting, developing, and retaining top global talent.

NVIDIA continues to be among America's best places to work as judged by employees. We were ranked #1 on Glassdoor's Best Places to Work in Tech & AI and #3 on Glassdoor's Best Places to Work overall for 2026.



Our employees also help to surface top talent, with over 40% of our new hires in FY26 coming from employee referrals. Additionally, we are focused on inclusive hiring, and we welcome employees of all backgrounds.

When recruiting new talent or developing our current employees, we strive to build a robust talent pipeline that includes top candidates from all backgrounds. We continue to focus on our employee referral program, talent-based hiring process, and partnerships with our community resource groups. We also partner with institutions and professional organizations around the world to reach top talent, attend conferences that serve many different populations, and host onsite events.

To strengthen our university pipeline, we invest in entry-level feeder programs around the world. This includes the [NVIDIA Ignite](#) program that prepares first- and second-year college students for an NVIDIA internship the following summer. We implement tools to help us identify a wide pool of top talent in the student population, through virtual and on-campus events.

Recruitment

As the demand for global technical talent continues to be high, we've grown our technical workforce and have been successful in attracting top talent to NVIDIA. We have attracted talent worldwide through our strong employer brand and differentiated hiring strategies for college, professional, and leadership talent. Within our workforce, over 80% have technical roles and more than half hold an advanced degree.

42,000
NVIDIAIANS

6,000
New in FY26

38
Countries



Benefits and Compensation

We offer comprehensive benefits to support our employees' and their families' physical health, well-being, and financial health. Programs include 401(k) programs in the U.S., statutory and supplemental pension programs outside the U.S., our employee stock purchase program, flexible work hours, and time-off policies. We evaluate our benefit offerings globally each year and aim to provide consistent support across the regions where we operate.

We offer tailored benefits based on the needs of our employees, including continuing support for parents—both new birth parents and those who wish to become parents. We provide employees with benefits such as reimbursement for eligible adoption, surrogacy, and fertility treatment expenses. Our parental leave program, available to all employees, enables parents to bond with their children. Non-birth parents can take up to 12 weeks of paid leave, and birth parents have up to 22 weeks of paid leave to ensure adequate recovery.



We partner with NVIDIA’s 10 different community resource groups to enhance support programs based on targeted needs, including military leave, student loan repayment, enhanced health insurance coverage for family members with developmental delays, and mental health counselor search tools.

We evaluate our benefit offerings annually to ensure employee needs are met and continuously seek feedback from employees to enhance our support.

Pay and Promotion

We strive to provide equitable compensation and opportunities for advancement to all employees and to achieve promotion parity based on a variety of considerations. We perform an annual review of peer compensation in the markets where we operate and track equity and parity in retention, promotions, and pay.

Each year, every employee’s performance is assessed, and managers provide feedback on goals and priorities.

Promotion readiness is reviewed, and compensation is adjusted appropriately. Employees are provided both cash and equity compensation.

Cash compensation adjustments are determined based on role, market compensation growth, and individual performance. Equity compensation is provided through restricted stock units (RSUs) that vest over time and provide value aligned to our stock price. We believe RSUs promote retention and align with stockholder interests.

Engagement and Retention

We want NVIDIA to be a place where people can grow their careers over their lifetime.

To ensure our continued success, we maintain a work environment where our employees are engaged and inspired.

To evaluate employee sentiment and engagement, we use pulse surveys, a suggestion box, and an anonymous third-party platform. Pulse surveys help us gain insight into employee experience and provide employee-generated ideas so we can take targeted action. This agile approach allows us to hear from employees more frequently, stay focused on specific areas, and react quickly to current events.

The suggestion box is an always-on, interactive tool where employees are able to ask our executive staff questions, submit ideas at any time, and vote on ideas about making our

company a better place to work. All questions are reviewed and receive a response.

We support a flexible work environment, which allows us to recruit the very best employees, regardless of where they live.

This flexibility supports diverse hiring and retention of talent, including working parents and other caregivers, and boosts employee engagement.

To give employees time to disconnect and recharge, we provide all NVIDIA’Ns eight “free days” per year, which are a direct result of employee feedback for structured time off.

Our employees tend to come and stay. NVIDIA’s overall turnover of 3.7% remains extremely low compared with the technology industry average of 21.9%

Pay Ratio	FY26	FY25	FY24
Women Men (Global)	99.5 100	99.4 100	99.5 100
Asian White (U.S.)	100.1 100	100.3 100	100.1 100
Black/African American White (U.S.)	100 100	100.6 100	101.5 100
Hispanic/Latino White (U.S.)	99.3 100	99.2 100	99.2 100

Turnover Rate	FY26	FY25	FY24
Overall turnover	3.7%	2.5%	2.7%
Men	3.7%	2.5%	2.7%
Women	3.6%	2.3%	2.8%
Asian (U.S.)	3.5%	2.0%	2.1%
Black/African American (U.S.)	6.3%	6.4%	3.1%
Hispanic/Latino (U.S.)	4.4%	2.7%	1.6%

Learning and Development

Being a lifetime learner is core to our culture. We have an extensive library of learning experiences that include online courses, live workshops, panel discussions, and speaker forums.

We constantly upgrade our learning offerings to ensure that our employees are exposed to the most current content and technologies available.

We offer [tuition reimbursement programs](#) to subsidize educational programs and advanced certifications.



We learn from our peers at our annual NTECH conference, where employees share their work with each other.

To support employee development, we provide opportunities to learn on the job through training courses, targeted development programs, mentoring and peer coaching, and ongoing feedback.

Our global Stride mentoring program offers one-on-one and group sessions, matching over 2,000 participants since its launch.

We implemented a career coaching service to provide one-on-one guidance to employees and encourage internal job mobility. Each year, we host a company-wide event that includes a series of development sessions for employees to build new skills, understand their strengths, and learn how to advance their careers at the company.

904,000+

Hours of learning logged by NVIDIA employees globally—approximately 21 hours per employee in FY26.

Health and Safety

Our Environmental, Health and Safety (EHS) team oversees workplace conditions for NVIDIANS globally, helping prevent injuries, protect well-being, and ensure our facilities not only meet but frequently exceed local safety requirements.

We promote safe work practices and maintain compliance with applicable health and safety legislation and internal policies in every region where we operate.

Our Integrated EHS Framework

To manage the complexities of a global workforce, our framework relies on a multi-layered approach:

- › Localized Oversight: A robust network of dedicated site safety officers who monitor real-time conditions.
- › Global Safety Programs: Standardized protocols covering essential topics such as ergonomics, ladder safety, vehicle safety, and chemical management.
- › Targeted Training: Hazard-specific online courses, including Hazardous Waste Management, Data Center Operations, and specialized Lab Safety modules.



With our increasing focus on advanced R&D, we have expanded our oversight to include specialized safety programs for robotics, laser safety, and radiation safety.

Hazard Identification and Risk Management

We proactively identify and assess workplace risks through routine safety inspections, external audits, and a transparent incident reporting system. Our Illness and Injury Prevention Program defines responsibilities, processes, and expectations for hazard assessment, incident reporting, and investigation. We encourage all employees to report hazards and near-misses and participate in safety committees and engagements, reinforcing a shared responsibility for safety. All incidents and subsequent investigation results are recorded in our Global Security Database, allowing us to implement corrective actions and continuously improve our safety performance across all global sites.

Community Resource Groups

We support 10 community resource groups (CRGs), which are open to all employees: Arabs & Allies Network, Asian Pacific Islanders & Allies, Black NVIDIAN Network, Early Career Network, Hispanic-Latino Network, NV Pride (LGBTQ+ employees and allies), NVIDIA Disabilities Network, Veterans@NVIDIA, South Asians & Allies, and Women in Technology.

We host topic-specific education events for CRGs relevant to their needs and feedback, including a session on our Inspire 365 initiative, which promotes employee giving and volunteering.



Allyship

Our allyship program is established and facilitated by employees to build a network of allies who model inclusive behaviors, amplify the work of others, and advocate for changes that lead to increased equity. The global program trains NVIDIANS to show up as allies for their peers, leading to a sense of belonging across the company. We also develop inclusive educational resources, made available to employees online.



Mentorship

We believe our mentorship programs cultivate a collaborative spirit throughout the company and nurture a culture of continuous learning and development. CRG participants experience mentoring in a one-on-one relationship or a group setting, reaching approximately 1,200 employees across the company. Specialized focuses include women, early career, and employees from traditionally underrepresented groups to ensure widespread readiness for future advancement.



Leadership

Our leadership development programs deliver tools, insights, and capabilities that managers need to excel in their roles. We also offer a program that educates all employees on the expectations of leaders at NVIDIA so they can prepare themselves to move into leadership. We provide managers with a resource guide that provides them with practical actions to create an inclusive employee experience.

Product Value Chain

We utilize a fabless and contracting manufacturing strategy, whereby we employ and partner with key suppliers for all phases of the manufacturing process, including wafer fabrication, assembly, testing, and packaging. We use the expertise of industry-leading suppliers that are certified by the International Organization for Standardization in such areas as fabrication, assembly, quality control and assurance, reliability, and testing.

We have expanded our supplier relationships to build redundancy and resilience in our operations to provide long-term manufacturing capacity aligned with growing customer demand. While currently our supply chain is mainly concentrated in Asia, we are expanding into the U.S. and Latin America. We utilize foundries, such as Taiwan Semiconductor Manufacturing Company Limited, or TSMC, and Samsung Electronics Co., Ltd., or Samsung, to produce our semiconductor wafers. We purchase memory from SK Hynix Inc., Micron Technology, Inc., and Samsung. We engage with independent subcontractors and contract manufacturers such as Hon Hai Precision Industry Co., Ltd., Wistron Corporation, and Fabrinet to perform assembly, testing, and packaging of our final products.

Responsible Sourcing

NVIDIA's supplier selection process incorporates social and environmental factors, including risk analysis on topics such as RBA code alignment, product compliance, and environmental initiatives.

We hold ourselves and our suppliers to high standards of behavior. We closely manage our supply chain to deliver innovative products that satisfy our customers' expectations in a socially and environmentally conscious manner.

Responsible Minerals

NVIDIA is committed to the responsible sourcing of minerals. We support, contribute to, and utilize industry-wide efforts to validate the source of minerals used in our products, aiming to ensure that they come from socially responsible sources and do not contribute to human conflict. Our goal is to use only conflict-free gold, tantalum, tungsten, and tin (3TG) in our products.

We evaluate smelter metrics and share an annual update in our [Conflict Minerals Report](#).



We have successfully removed non-compliant 3TG smelters and refiners from our supply chain to help us meet this goal. We received responses from 100% of our suppliers and component manufacturers with the requested information.

We also collect non-3TG smelter data for cobalt and mica from our strategic suppliers and report the initial results to customers. As the cobalt and mica programs through the RBA's Responsible Minerals Initiative (RMI)

mature, our goal is to continue to increase the percentage sourced from conflict-free refiners.

In FY26, we also implemented the updated Extended Minerals Reporting Template (EMRT) for graphite, lithium, nickel, and copper.

Our [Responsible Minerals Policy](#), sourcing goals, and steps for monitoring our supply chain for conflict minerals are framed around the Organization for Economic Cooperation and Development (OECD) Due Diligence Guidance for Responsible



Supply Chains of Minerals from Conflict-Affected and High-Risk Areas (CAHRAs). This provides detailed recommendations to help companies respect human rights and avoid contributing to conflict through their mineral purchasing decisions and practices.

Our conflict minerals due diligence program is designed to conform in all major respects with the framework recommended by the OECD. We've established a company management program endorsed by our Executive Vice President, Operations. We identify and assess the risks in our supply chain and have adopted a strategy and process for responding to the risks we've identified. As part of our diligence process, we review smelter data published by the Responsible Minerals Assurance Process (RMAP) of the RMI, as well as with the smelters that were audited by the London Bullion Market Association (LBMA) the Responsible Jewelry Council (RJC), and the Tungsten Industry—Conflict Minerals Council (TI-CMC). These groups carry out independent third-party audits of due diligence at identified points in the supply chain to verify smelter compliance status. Finally, NVIDIA makes various reports on our supply chain due diligence, including internally to management and externally through our [Form SD](#).

We're a member of the Public Private Alliance (PPA) for Responsible Minerals Trade and the RMI. The PPA provides funding and coordination support to organizations working within the Democratic Republic of Congo and adjacent countries to develop verifiable conflict-free supply chains, align due diligence programs and practices, encourage responsible sourcing from the region, promote transparency, and bolster in-region civil society and governmental capacity. We support these on-the-ground programs aimed at improving transparency for responsible sourcing and reducing human rights risks, including forced labor.

Additionally, we participate in various RMI work groups and align our program with the organization's tracking of additional minerals and materials, as well as CAHRAs. We monitor legal requirements to determine any future obligations regarding materials sourcing and high-risk regions of the world.

Strategic Suppliers

Strategic suppliers include those who produce or handle NVIDIA production materials, those for whom we closely manage quality requirements, those who design our branded products, and those who we're required to work with based on customer agreements.

For strategic suppliers, agreements are deployed and tracked through a Quarterly Business Review (QBR) or a Semi-annual Business Review (SBR) process to ensure they meet our expectations, which may influence our business allocation decisions. Five percent of the QBR score is based on environmental and social performance, and expectations vary by quarter. Each quarter, we assess these suppliers by product category and rank performance, and we perform an annual risk assessment based on RBA results, geography, and type of industry. We've assessed and adjusted business with suppliers who don't comply with minimum requirements.

We review all strategic suppliers against NVIDIA cybersecurity requirements and have performed a more in-depth assessment of higher-risk suppliers.

We looked at our suppliers' conformance to ISO 27001, ISO 28001, C-TPAT, and NIST standards and evaluated them for additional information-, product-, and physical-security expectations. If gaps are discovered, we expect an improvement plan to align with NVIDIA requirements. The cybersecurity risk rank determines supplier risk to be low, moderate, or high, and assessments or audits are conducted on a one- to three-year cycle based on our risk analysis.

Risk Assessment and Auditing

All our manufacturing suppliers are expected to comply with the RBA Code of Conduct and associated NVIDIA policies, including our Agreement for Manufacturer Environmental Compliance. Since 2016, all NVIDIA Master Service Agreements executed with suppliers require compliance with the RBA Code of Conduct.



Our assessment process involves using the RBA-Online system to evaluate existing and potential new suppliers against product compliance industry standards, social and environmental criteria, use of conflict minerals, the RBA Code of Conduct, and NVIDIA's Code of Conduct. All suppliers are expected to conduct an annual self-assessment questionnaire (SAQ), and we expect biennial Validated Assessment Program (VAP) onsite, third-party audits to validate the SAQ. We use the results of the assessment, which includes a spending analysis, to determine their overall risk.

To ensure consistency, suppliers can share VAP audits with customers. Any uncovered findings are also made available to those customers. Findings are aggregated anonymously by the RBA and disclosed in their public annual report.

Our goal is to audit 100% of our strategic suppliers every two years. In FY26, we reviewed VAP audits on 45% of our strategic suppliers, bringing total audits in the past two years to 93%. For all VAP audits conducted, we also engaged 100% of strategic suppliers on their corrective action plans, where applicable, and continue monitoring suppliers to ensure compliance.

Common audit findings include issues related to freely chosen employment, working hours, wages and benefits, emergency preparedness, and occupational safety. NVIDIA ensures strategic suppliers close all corrective actions in compliance with our Full RBA membership tier.

Product Environmental Impact

Product Carbon Footprints (PCFs) help us to understand the greenhouse gas emissions generated across our product lifecycle, especially those in manufacturing. In FY26, we [published](#) PCF summaries for two NVIDIA baseboard products. The PCFs were developed in conformance with the ISO 14067:2018 standard and verified by a third party. Both PCFs are a cradle-to-gate study including impacts from raw material extraction and refinement, material transport, component manufacturing, and assembly.

A comparative analysis of these PCF summaries demonstrates a 24% reduction of embodied carbon emissions across large workloads such as AI training and inference between generations. For a detailed analysis of how NVIDIA reduces both embodied and operational carbon emissions intensity, refer to our [blog](#).

We collaborate with our suppliers to collect product-specific data for high-impact components, improving data quality, relevance, and precision. Looking ahead, we intend to expand our product life cycle assessment portfolio, enhance our models with more accurate component data for high-volume products, and further integrate supplier-specific data to provide a comprehensive assessment of our products' environmental impact.

Manufacturing

Emissions are generated at every stage of our product lifecycle, including manufacturing within our supply chain. Since 2014, we've expected our key silicon manufacturing and systems contract manufacturing suppliers to report their annual energy and water usage, waste, greenhouse gas (GHG) emissions, and reduction goals and objectives through the RBA Environmental Survey or CDP. We also expect suppliers to have their GHG emissions verified by a third party. We use this supplier data to better understand our product manufacturing impact and allocate carbon emissions to our customers.

We maintain Full Material Disclosures for our chip designs and select system products, which demonstrates our compliance with restricted substances including halogenated flame retardants, lead, and mercury.



Information on materials we use is reported through various platforms, including Substances of Concern in Products (SCIP) and International Material Data System for Automotive.



Packaging

We continue to optimize the balance between ensuring that customers receive products in a protected, like-new condition while also minimizing our impact on the environment. We actively seek opportunities and implement appropriate solutions to reduce our usage of non-recyclable materials and packaging materials in general, while increasing the amount of recycled content in the materials we do use.

In FY26, our NVIDIA GPU systems packaging contained 96% recyclable materials by weight. NVIDIA continues to progress other sustainability initiatives, such as reducing our reliance on plastic foam, increasing applications for paper-based cushioning, replacing printed documentation with online versions when possible, working with our vendors to ensure they maintain FSC certification for their corrugated paper and expand certification at all new locations, and exploring certification for other paper packaging materials.

To meet compliance expectations, we also add material identification codes to our packaging parts for ease of disposal by end-users.

96%

Our NVIDIA GPU systems packaging contained 96% recyclable materials by weight in FY26.

Transport and Logistics

We plan, pack, and execute our raw material, work-in-progress, and finished goods shipments to optimize transport and logistics, which results in reduced fuel use.

This includes consolidating shipping volumes and weights and implementing a multimodal distribution system. An area of continued focus is to consolidate the number of pickups and shipments.

We expect carriers to report shipment data to support carbon emissions calculations and have sustainability initiatives, which are both covered in quarterly business reviews.

End-of-Life Management

NVIDIA-branded products are subject to electronic waste disposal requirements in our various global markets, and we take seriously our extended responsibility for these products.

In key regions such as the U.S. and Europe, we've established recycling programs in partnership with reputable third parties and provide information to consumers about how to recycle our branded products through labeling and our [website](#).

We help our customers [monetize](#) the residual value in their older NVIDIA DGX™ servers when refreshing to the latest technology, with the help of a remarketing partner. Older-generation DGX systems may be resold in the secondary market, and systems that are not resold are typically disassembled for materials recovery and recycling.

Product Quality and Safety

Safety is a principle that permeates our engineering culture. We incorporate it into every step of our product development process, from design and production to the end-user experience.

Our products undergo rigorous qualification and testing to ensure compliance with customer safety and reliability standards. With these standards in mind, we've established product safety technical committees to oversee safety throughout the product lifecycle.

Cross-functional teams manage all aspects of product quality. Customer quality engineering provides direct customer support for all quality-related issues, gathers information on customer failures and customer-observed failure rates, and processes returned material authorization.

As part of the continuous improvement process, recommended design enhancements are brought to the engineering teams for incorporation into next-generation products. We monitor our quality-management system through internal audits and an annual third-party [ISO 9001](#) assessment.



Responsible Business

NVIDIA's Core Values



Innovation

Dream big, start small.
Take risks, learn fast.



Intellectual Honesty

Seek truth, learn
from mistakes,
share learnings.



Speed & Agility

Learn, adapt,
shape the world.



Excellence & Determination

Maintain the highest
standards.



One Team

Do what's best for
the company.

Our Code

Our Code of Conduct outlines our core values and establishes the expectations we have about how we conduct business. It guides our actions in our professional relationships with fellow employees, as well as our customers, partners, competitors, vendors, government regulators, shareholders, and the community at large. Our code applies to all employees and board members, and we expect all third parties we do business with, including consultants, contractors, and other service providers, to act in a manner consistent with it.

Our Code is translated into 13 different languages, which are available on our [website](#).

Ethics and Cybersecurity

We strive to achieve the highest standards of ethical conduct in all our business dealings and are committed to promoting a culture of integrity. NVIDIA employees receive Code of Conduct training upon hire and then every two years thereafter. This training covers topics like ethics, preventing workplace

harassment, and sustainability. Employees who have frequent contact with customers, partners, and suppliers (such as those in sales, finance, and procurement) complete additional global anti-bribery and anti-corruption training. In FY26, our employee training completion rates were over 95% for cybersecurity, 97% for our Code of Conduct, and 99% for anti-bribery and anti-corruption.

In addition to employee cybersecurity trainings, our Santa Clara headquarters is certified to [ISO 27001](#). For more information on our approach to cybersecurity, see pp. 32 of our [2026 Form 10-K](#) filed with the SEC.

Investigations Process

Our [Code of Conduct](#) reinforces our culture of integrity by setting clear expectations for employees and the third parties we do business with to act ethically and feel empowered to voice concerns without fear of retaliation.

Anyone can confidentially and anonymously submit a complaint about any matter using our [Speak Up](#) corporate hotline, which is hosted by an independent third party. Such

concerns may involve matters related to accounting, internal controls, auditing, the Code of Conduct, harassment, conflict of interest, or other potential violations of company policies, laws, or regulations.

Employees are encouraged to report suspected Code or other policy violations to their manager, a human resources representative, the legal department, and/or the anonymous hotline. We use an external third party for our Speak Up hotline to comply with applicable regulations and to give individuals confidence that their reports will remain anonymous if they prefer to report anonymously. We have a strict anti-retaliation policy that prohibits retaliation against anyone who, in good faith, reports a concern.

Our Compliance Committee is composed of our CFO, EVP of Operations, and SVP of Human Resources. NVIDIA's Compliance team ensures communication of and adherence to all NVIDIA policies; establishes procedures for the receipt, retention, and treatment of complaints; and promptly and thoroughly conducts investigations in partnership with relevant organizations within NVIDIA. The NCGC oversees the Compliance Committee



and periodically reviews our policies, practices, and investigation procedures in connection with our compliance program.

The Compliance Committee periodically reviews and assesses the effectiveness of its procedures and adjusts as necessary. Our third-party auditor reviews cases quarterly and process updates annually.

Employees who wish to understand NVIDIA's investigations process can contact the Compliance team or access an online training course.

Trustworthy AI

We endeavor to deliver AI models that:

- › Comply with relevant data protection and privacy laws,
- › Perform safely and as intended,
- › Communicate transparently about a model's design, capabilities, and limitations, and
- › Minimize unwanted bias and give all groups equal opportunity to benefit from AI.

Our [Trustworthy AI principles](#), which we share with our employees, customers, and partners, reflect our core values and [Code of Conduct](#).

Scope and Governance

To build [Trustworthy AI](#) (TAI), we develop practices, methodologies, and tools for datasets, machine learning frameworks, AI model development, and software development and testing. We build tools and processes that enable us, our customers and partners, and the broader ecosystem to build safer AI models and tools that are trustworthy.

NVIDIA's TAI efforts are led by our Head of AI & Legal Ethics and supported by dedicated product and risk teams. These teams are structured to work cross-functionally with engineering, product, product security, safety, and infrastructure teams to ensure our AI aligns to our principles. The teams also participate in external AI safety standards development efforts to contribute industry best practices.

Our legal team works with AI safety teams to monitor global regulations and relevant standards to ensure we maintain compliance. The internal AI Ethics Committee advises on generative AI development with members from our engineering, business, research, public policy, and legal organizations. In FY26, we continued to facilitate internal training opportunities and special TAI volunteer task forces to cross-train teams on how to apply TAI principles to product development and releases.



Privacy

AI should comply with privacy laws and regulations, and meet societal norms for personal data and information privacy.



Transparency

Make AI technology understandable to people. Explain, in non-technical language, how an AI system arrived at its output.



Safety and Security

Ensure that AI systems perform as intended and avoid unintended harm and malicious threats.



Nondiscrimination

Minimize bias in our AI systems and give all groups an equal opportunity to benefit from AI.

Standards and Frameworks

We look to global frameworks such as the EU's High-Level Expert Group on Artificial Intelligence and Singapore's Model Artificial Intelligence Governance Framework to inform our software development processes, from concept to performance validation to decommission. We participate in industry-driven initiatives like [ML Commons](#), as well as standards bodies like the IEEE. NVIDIA is also a member of the National Institute of Standards and Technology (NIST) AI

Safety Consortium and the [Content Authenticity Initiative](#), which are focused on the development and use of secure and trustworthy AI systems. In FY26, NVIDIA published several general-purpose AI models that comply with the EU AI Act's transparency requirements, including comprehensive technical documentation, model cards detailing capabilities and limitations, and required disclosures about training data.



AI Transparency

Open technologies—made available to developers and businesses to adopt, modify, and innovate with—have been part of every major technology shift, from the birth of the internet to the early days of cloud computing. AI transparency means that its adopters can understand how their models work and trust the results they provide.

NVIDIA Nemotron is a [collection](#) of open source AI technologies designed for efficient AI development at every stage and is part of NVIDIA’s wider efforts to provide open, transparent, and adaptable AI platforms for developers, industry leaders, and AI infrastructure builders across the private and public sectors.

[Model cards](#) detail how machine learning (ML) models work, ensuring transparency

and clear communication about how NVIDIA technology was built and should be used by developers working with our AI models. Information includes the ML model’s metadata, including the datasets that the model was trained on, performance measures, development and evaluation methodology, and technical limitations.

We continuously enhance our model cards to make them more understandable to consumers, investors, and policy makers. Recent enhancements have focused on improved readability with clearer language, increased detail around ethical considerations, and greater accessibility.

We [open source](#) transparency card documentation for blueprints, datasets, containers, and systems for industry to adopt and continuously update content and format. Teams are now using automation to help drive model card completion and compliance.

AI Risk Management

NVIDIA aims to reduce the risk of harm from deployment of AI models or systems. Model risk can come from many sources across the AI supply chain and within the deployment environment, including training

datasets, malicious attacks, or failure to comply with laws and regulations.

Our risk management guidance, based on the [NIST Risk Management Framework](#), is integrated into our AI product lifecycle process, and outlines the elements of development, validation, audit, and documentation to all employees working on models. This approach provides internal

development teams with guidance to achieve innovative solutions aligned to our Trustworthy AI principles.

NVIDIA is committed to providing open source tools for trustworthy AI model development and integration.



Full Technology Stack

NVIDIA offers an open, full-stack platform with powerful compute infrastructure, advanced AI and high-speed networking.



Industry Collaboration

NVIDIA works with industry leaders on open AI technologies that are reshaping how businesses operate.



Community Enablement

NVIDIA publishes models and training data openly, averaging nearly one new model or dataset daily on Hugging Face.



Safety and Trust

NVIDIA is committed to trustworthy AI usage and provides transparency for each stage of AI development and deployment.

1,000+

NVIDIA models across all domains meet the Model Card++ standard.



We've designed and continue to improve tools for internal developers to create TAI models:

- › Integrated third-party tools to conduct independent testing of high-risk models.
- › Collaborated with an external law firm to automate the identification of legal obligations when models are registered in our internal inventory.
- › Created a dashboard that measures a model's completion of essential steps prior to publication, including a current performance score and target goal.
- › Implemented model requirement and dataset approval processes that must be met prior to commercial and noncommercial releases.
- › Released model cards that accompany our published models, providing guidance on a model's uses and limitations, including ethical considerations such as explainability and privacy.

AI Safety

As LLMs become more widespread, the content safety risks associated with their use also increase. Deploying safeguards across various points in a model's architecture ensures that if one layer is compromised, others remain effective.

This approach enhances resilience against potential threats by providing redundant protective measures. Guardrails can be [implemented](#) at various locations in the model architecture.

To [evaluate](#) the safety of our models, we employ the Nemotron Content Safety Guard Model—a high-quality content safety evaluation dataset, LLM-based content safety classifier model, and evaluation benchmark. The model is backed by a broad content safety risk taxonomy that covers critical risks in human-LLM interactions, such as content including hate/identity hate, harassment, or violence. Our [NVIDIA NeMo™ Evaluator](#) provides a microservice to assess AI models across academic and custom benchmarks. Our [NeMo Guardrails](#) acts as a safety layer on top of language models to enforce predefined rules and policies during inference and offers an interface to run an internal red teaming to evaluate guardrails configurations.

Alongside content safety, we assess generative models for security weaknesses. [Garak](#)—Generative AI Red-teaming & Assessment Kit—is a vulnerability scanner for LLMs. It identifies a broad range of security weaknesses and unwanted behaviors in language model-based technology by scanning a model or dialog system to quickly discover

where it's working well and where it may be vulnerable to attack. Garak provides full reporting and informs red teaming activities.

In addition to automated safety evaluation with systems like the Nemotron family of products and Garak, human content red teaming further [evaluates](#) models using human interaction and response assessment to identify vulnerabilities and behaviors missed elsewhere.

Human Rights

The NCGC has oversight of policies and practices in connection with human rights and provides periodic updates to the Board. At the management level, our CSSC is responsible for reviewing and ensuring the effectiveness of our human rights program. Through this governance body and related executive engagement, management undertakes all actions it deems reasonable and necessary to ensure compliance with our [Human Rights Policy](#).

We regularly assess human rights-related risks and potential impacts, review our policies and management processes, and seek input from stakeholders on our approach. In FY25, NVIDIA conducted a human rights saliency assessment of potential human

rights impacts across our value chain. The saliency assessment identified five salient human rights risk areas for NVIDIA: responsible minerals sourcing; working conditions in the supply chain; clean, healthy, sustainable environment; responsible product development; and responsible product use. These assessment results guided updates to NVIDIA's human rights and responsible minerals policies in FY26 and will inform future due diligence and initiatives.

All NVIDIA employees are required to complete training on our code, which includes our commitment and guidelines to respect human rights. Additionally, certain employees take RBA e-Learning Academy courses and are engaged in RBA workgroups relevant to our supply chain operations.

Employees, suppliers, and other external stakeholders can confidentially and anonymously report a concern about human rights using our [Speak-Up line](#). For more information on our Speak-Up line, see our [Investigations Process](#) section.

NVIDIA has been a signatory of the UN Global Compact since June 2022.

Upstream Human Rights

We take the issue of forced labor very seriously and conduct regular audits to ensure all workers in our supply chain are treated with respect and dignity. We're full members of the Responsible Business Alliance (RBA) and work with them and similar industry groups to prevent and eliminate forced labor from global supply chains. Through the RBA, we also work with expert organizations, NGOs, and academia to advance our program and standards, and we support efforts at an industry level to identify and remediate actual and potential cases of labor and human rights abuse.

We expect our suppliers to maintain progressive employment, health, safety, and ethical practices that meet or exceed applicable laws, the [RBA Code of Conduct](#), our [Code of Conduct](#), and our [Human Rights Policy](#). We also encourage suppliers to use the RBA Code of Conduct as a platform to go above and beyond compliance. We participate in RBA's Responsible Labor Initiative (RLI) and monitor our supply chain through VAP audits, which cover freely chosen employment, child labor, and freedom of association.

In FY26, we reviewed VAP audits on 45% of our strategic suppliers, bringing total audits in the past two years to 93%.

If we uncover findings during these audits, we work directly with suppliers to implement any corrective actions. Accordingly, we worked with relevant suppliers to track and remediate working hours, to address and comply with applicable wages and benefits payments, and to address and comply with zero hiring fees.

Product Due Diligence

We believe AI can and will enhance human welfare while respecting human rights.

NVIDIA is accelerating the AI revolution by creating platforms and computing tools that help developers, researchers, and data scientists innovate.

Our products are customizable and integrable across workflows and platforms. NVIDIA promotes responsible use of our products through responsible design and the development of tools customers can use to enhance product safety. We also support product development aligned with our principles and the mitigation of potential risks, including those related to human rights. For more on our efforts in delivering safe and ethical AI models, see our [Trustworthy AI](#) section.



Human Rights Principles

We support and respect internationally recognized human rights and principles, such as those in the:

- › [United Nations Global Compact](#)
- › [United Nations Guiding Principles](#)
- › [Universal Declaration of Human Rights](#)
- › [International Covenant on Civil and Political Rights](#)
- › [International Covenant on Economic, Social and Cultural Rights](#)
- › [Core Conventions of the International Labour Organization](#)
- › [ILO Declaration on Fundamental Principles and Rights at Work](#)



Sustainability Indicators

NVIDIA is committed to transparency, and the following tables provide our disclosure on key sustainability metrics. We report metrics relevant to our company and sector using the following frameworks and standards: Sustainability Accounting Standards Board (SASB) for the Technology and Communications sector and Semiconductor industry, and Global Reporting Initiative (GRI). We reference indicators that align with the United Nations Global Compact (UNGC) Communication on Progress (CoP). NVIDIA reports with reference to the 2021 GRI standards.

About NVIDIA

Metric	FY26	Reference Indicator
Organizational details	2026 10-K, pp. 4-5	GRI 2-1
Activities, value chain and other business relationships	2026 10-K, pp. 5-7 Product Value Chain	GRI 2-6
Governance structure and composition	Sustainability Governance Committee Composition 2026 Proxy Statement, pp. 14-28	GRI 2-9 UNGC G1, G3
Nomination and selection of the highest governance body	Corporate Governance Policies of the Board of Directors	GRI 2-10
Chair of the highest governance body	Board of Directors Committee Composition 2026 Proxy Statement, pp. 23-24	GRI 2-11
Collective knowledge of the highest governance body	2026 Proxy Statement, pp. 17-22	GRI 2-17
Evaluation of the performance of the highest governance body	2026 Proxy Statement, pp. 28-31	GRI 2-18
Remuneration policies and process to determine remuneration	2026 Proxy Statement, pp. 32-33, 38-53	GRI 2-19, 2-20
Annual total compensation ratio	2026 Proxy Statement, pp. 54-55	GRI 2-21
Statement on sustainable development strategy	Message From Our CEO 2026 GTC Keynote	GRI 2-22 UNGC G1



About NVIDIA

Metric	FY26	Reference Indicator
Membership associations	Public Policy Engagement	GRI 2-28
Approach to stakeholder engagement	Our People Product Value Chain	GRI 2-29
Compliance with laws and regulations	We consider significant fines that are required to be disclosed in the Company's SEC filings. See our 2026 Form 10-K filed with the SEC, pp. 9-11.	GRI 2-27

Sustainability Governance

Entities included in the organization's sustainability reporting	Our sustainability reporting covers the same entities as our financial reporting. See our 2026 Form 10-K filed with the SEC, pp. 88.	GRI 2-2
Reporting period, frequency and contact point	This report covers our fiscal year ended January 25, 2026 (FY26). We publish our Sustainability report annually, previous reports can be found in the NVIDIA Report Archive . We can be reached at NVIDIACorporateSustainability@nvidia.com	GRI 2-3
Approach to sustainability governance	Sustainability Governance	GRI 2-12, 2-13, 2-14 UNGC G1, G3

Societal Impact

Operations with local community engagement, impact assessments, and development programs	FY26 Foundation Annual Report	GRI 413-1
Political contributions	Public Policy Engagement	GRI 415-1



Energy, Efficiency, and Climate

Metric	FY26	FY25	FY24	Reference Indicator
GHG Emissions (MT CO₂e)*				
Scope 1	9,822	9,047	8,887	GRI 305-1 UNGC E5 SASB TC-SC-110a.1
Scope 2, market-based	568	0	40,555	GRI 305-2 UNGC E5
Scope 1 and 2, market-based	10,390	9,047	49,442	
Scope 2, location-based	308,891	228,378	178,087	GRI 305-2 UNGC E5
Scope 3	10,700,940	6,912,577	3,638,432	GRI 305-3 UNGC E6, E6.1
Category 1: Purchased goods and services	9,301,735	6,036,105	3,216,144	
Category 2: Capital goods	954,865	570,175	200,483	
Category 3: Fuel- and energy-related activities	103,390	75,035	61,590	
Category 4: Upstream transportation and distribution	137,747	78,199	72,562	
Category 5: Waste generated in operations	2,285	1,416	1,571	
Category 6: Business travel	52,882	36,032	17,132	
Category 7: Employee commuting	48,609	45,255	23,019	
Category 8: Upstream leased assets	99,427	70,360	45,931	
Scope 1 and 2 emissions by GHG type	Refer to our Report of Independent Accountants for Scope 1, Scope 2 market-based, and Scope 2 location-based FY26 emissions (mtCO ₂ e and metric tons) by relevant GHG type.			GRI 305-1, 305-2
External assurance	Report of Independent Accountants for select metrics for FY26.			GRI 2-5 UNGC G13

* Due to improved data collection and calculation methodology, some of the reported values in our prior-year GHG emissions no longer directly match the associated Report of Independent Accountants or data tables from those years. These values are not subject to assurance.



Energy, Efficiency, and Climate

Metric	FY26	FY25	FY24	Reference Indicator
GHG emissions intensity (Scope 1 and 2 mtCO ₂ e/\$M revenue)	0.05	0.07	0.81	GRI 305-4
Reduction of GHG emissions	Energy, Efficiency, and Climate			GRI 305-5 SASB TC-SC-110a.2
Energy (MWh)*				
Energy consumption	1,053,479	815,864	593,953	GRI 302-1
Non-renewable electricity	0	0	134,287	
Renewable electricity purchased	1,004,842	779,316	419,980	
Renewable electricity generated, onsite solar	622	818	599	
Fuel	46,258	35,729	39,087	
Purchased heating	862			
Purchased cooling	895			
Renewable electricity percentage (%)	100%	100%	76%	
Energy intensity (Energy consumption MWh/\$M revenue)	4.9	6.3	9.7	GRI 302-3
Reduction of energy consumption	Energy-Efficient Sites			GRI 302-4
Reductions in energy requirements of products and services	Product Energy Efficiency			GRI 302-5 SASB TC-SC-410a.2
Energy management in manufacturing	We are a fabless semiconductor company and do not have our own manufacturing facilities. See Product Environmental Impact for how we work with suppliers to track energy usage.			SASB TC-SC-130a.1

* Due to improved data collection and calculation methodology, some reported values in prior-year energy data no longer directly match the data tables from those years.



Energy, Efficiency, and Climate

Metric	FY26	FY25	FY24	Reference Indicator
Waste (metric tons)*				
Waste generated	5,257	3,299	2,346	GRI 306-3
Waste recycled/composted	4,358	2,759	1,659	GRI 306-4
Waste landfilled	899	539	687	GRI 306-5
Landfill diversion rate (%)	83%	84%	71%	GRI 306-4 UNGC E16
Management of waste-related impacts	Waste Management			GRI 306-1, 306-2
Amount of hazardous waste from manufacturing	We are a fabless semiconductor company and do not have our own manufacturing facilities. See Product Environmental Impact for how we work with suppliers to track energy usage.			SASB TC-SC-150a.1
Water (m³)**				
Water withdrawal	451,297	409,814	382,636	GRI 303-3 UNGC E12 SASB TC-SC-140a.1
Water consumption	105,531	141,772	134,219	GRI 303-5 UNGC E12 SASB TC-SC-140a.1
Water discharge	345,766	268,043	248,417	GRI 303-4
Percentage of water withdrawn in regions with High or Extremely High Baseline Water Stress	We analyze withdrawal from water-stressed areas as part of our annual water risk assessment using the WRI Aqueduct tool. In WRI Aqueduct, we define water stressed areas as high or extremely high baseline water stress (greater or equal to 3). The WRI assessment resulted in 23% of our FY25 facilities water withdrawals as located in water stressed areas.			GRI 303-3 UNGC E12 SASB TC-SC-140a.1
Interactions with water as a shared resource	Water Conservation			GRI 303-1

* Waste data for FY26 expanded to include our India sites; prior-year waste is reported only for our Santa Clara, CA headquarters location.

** Due to changes in data collection methods, FY26 water values are not comparable to prior-year data.



Our People

Metric	FY26	Reference Indicator
Employees	2026 10-K, pp. 11	GRI 2-7
Benefits and promotion of worker health	Benefits and Compensation Benefits at NVIDIA	GRI 401-2, 401-3, 403-6 UNGC L4
Programs for upgrading employee skills and average hours of training	Learning and Development	GRI 404-1, 404-2
Percentage of employees receiving regular performance and career development reviews	Pay and Promotion	GRI 404-3
Ratio of basic salary and remuneration of women to men	Pay and Promotion	GRI 405-2 UNGC L6

Employee Health and Safety

Metric	FY26	FY25	FY24	Reference Indicator
Lost-time case rate (U.S.)*	0.01	0.01	0.01	GRI 403-9, 403-10 UNGC L7
OSHA recordable incident rate (U.S.)*	0.07	0.08	0.11	
Fatalities (global)	None	None	None	GRI 403-9
Occupational health and safety	Health and Safety			GRI 403-1, 403-4, 403-5, 403-7, 403-8 SASB TC-SC-320a.1
Total amount of monetary losses as a result of legal proceedings associated with employee health and safety violations	We disclose information on legal proceedings in our quarterly reports on Form 10-Q and our Annual Reports on Form 10-K .			SASB TC-SC-320a.2

* The Lost Time Case Rate is calculated by multiplying the number of incidents that were lost time cases by 200,000 and dividing by the number of labor hours at the company.
The OSHA recordable incident rate calculated by multiplying the number of recordable cases by 200,000 and dividing by the number of labor hours at the company.



Workforce Data

Metric	FY26	FY25	FY24	Reference Indicator
Region*				
Americas	50.5%	49.7%	49.8%	GRI 2-7
APAC	16.4%	16.9%	16.8%	
EMEA	20.5%	20.2%	19.8%	
India	12.6%	13.3%	13.6%	
Employee Type*				
Executive	0.1%	0.1%	0.1%	GRI 405-1
Management	15.1%	15.7%	16.7%	
Regular employees	84.9%	84.2%	83.2%	
Age*				
20-30 years	21.7%	22.3%	22.0%	GRI 405-1
31-50 years	64.5%	63.7%	63.4%	
51+ years	13.8%	14.0%	14.6%	
Gender*				
Women	21.6%	21.2%	19.7%	GRI 405-1
Men	77.3%	77.5%	79.4%	
No data or not disclosed	1.1%	1.2%	0.9%	

* Numbers may not sum to 100% due to rounding effects.



Workforce Data

Metric	FY26	FY25	FY24	Reference Indicator
Positions Held By Women*				
In technical roles	17.4%	16.8%	15.2%	GRI 405-1
Managers	19.4%	18.7%	17.7%	
Leaders	14.8%	13.5%	13.1%	
Executive Officers	40.0%	40.0%	40.0%	
Board of Directors**	30.0%	38.5%	33.3%	
Race/Ethnicity (U.S.)*				
American Indian or Alaska Native	0.1%	0.1%	0.1%	GRI 405-1
Asian	52.6%	51.6%	51.0%	
Black or African American	2.0%	2.2%	2.1%	
Hispanic or Latino	4.1%	4.1%	3.9%	
Native Hawaiian or Other Pacific Islander	0.2%	0.2%	0.3%	
Two or More Races	2.0%	1.8%	1.5%	
White	33.2%	34.1%	35.2%	
No data or not disclosed	5.7%	5.8%	5.9%	
Employee Self-Identification (U.S.)				
Disability	4.7%	4.2%	3.2%	GRI 405-1
Veteran	1.5%	1.4%	1.5%	
New employee hires and employee turnover	Recruitment Engagement and Retention			GRI 401-1

* Numbers may not sum to 100% due to rounding effects.

** Board of Directors data as of May 12, 2026.



Product Value Chain

Metric	FY26	Reference Indicator
New suppliers that were screened using environmental or social criteria	Responsible Sourcing Human Rights	GRI 308-1, 414-1
Negative impacts in the supply chain and actions taken	Responsible Sourcing Human Rights NVIDIA Responsible Minerals Policy	GRI 308-2, 414-2 SASB TC-SC-440a.1
Operations and suppliers in which the right to freedom of association and collective bargaining may be at risk	Risk Assessment and Auditing Human Rights Policy	GRI 407-1 UNGC L3, L4, L5
Operations and suppliers at significant risk for incidents of child labor	Human Rights Risk Assessment and Auditing Human Rights Policy	GRI 408-1 UNGC L3, L4, L5
Operations and suppliers at significant risk for incidents of forced or compulsory labor	Human Rights Risk Assessment and Auditing Human Rights Policy	GRI 409-1 UNGC L3, L4, L5
Percentage of production from owned facilities	We are a fabless semiconductor company and do not have our own manufacturing facilities.	SASB TC-SC-000.B







Responsible Business

Metric	FY26	Reference Indicator
Conflicts of interest	Corporate Governance Policies of the Board of Directors 2026 Proxy Statement, pp. 40	GRI 2-15
Communication of critical concerns	Investigations Process	GRI 2-16
Policy commitments	Company Policies	GRI 2-23 UNGC G2, L2, E1
Embedding policy commitments	Our Code Code of Conduct	GRI 2-24
Mechanisms for seeking advice and raising concerns	Investigations Process	GRI 2-25, 2-26 UNGC G6
Total amount of monetary losses as a result of legal proceedings associated with anti-competitive behavior regulations	We disclose information on legal proceedings in our quarterly reports on Form 10-Q and our Annual Reports on Form 10-K .	SASB TC-SC-520a.1
Anti-corruption policies and procedures	Anti-Corruption Policy	GRI 205-1, 205-2 UNGC AC2, G2



U.N. Sustainable Development Goals

We continue to align our business activities to the UN Sustainable Development Goals. Here are a few ways we contributed to the SDGs in FY26.

SDG	Our Support	Our Impact
	<p>Our employees' well-being, physical, emotional, and financial wellness is a top priority, and we aim to support them by offering a suite of services where people can choose what works best for them. We partner with NVIDIA's 10 different community resource groups to enhance support programs based on targeted needs, including military leave, student loan repayment, enhanced health insurance coverage for family members with developmental delays, and mental health counselor search tools.</p>	<p>NVIDIA technology enables healthcare institutions and researchers to harness the power of AI and HPC to improve patient care, accelerate research, and unlock medical breakthroughs. For example, Lilly launched a new AI factory with over 1,000 NVIDIA Blackwell Ultra GPUs to help its teams make meaningful medical advancements faster, more accurately and at unprecedented scale. With this foundation, Lilly teams can analyze genomes, explore billions of chemical possibilities, and apply AI across clinical development and manufacturing to design better trials, optimize production, and accelerate decision-making.</p>
	<p>To strengthen our university pipeline, we invest in entry-level feeder programs around the world. This includes the NVIDIA Ignite program that prepares first- and second-year college students for an NVIDIA internship the following summer. We implement tools to help us identify a wide pool of top talent in the student population, through virtual and on-campus events.</p>	<p>NVIDIA's cutting-edge technologies, available both on premises and in the cloud, are designed to support higher education researchers in their most ambitious projects.</p> <p>To empower early education, NVIDIA pledged \$25 million in support with AI education programs, partnering with two leading K-12 learning platforms to tailor the NVIDIA Deep Learning Institute (DLI) and NVIDIA Academy content to meet the instructional needs of U.S. K-12 classrooms.</p>
	<p>Improving performance and energy efficiency is a principal goal in each step of our research, development, and design processes. We develop hardware, software, and networking technology to improve performance and energy efficiency, helping our customers solve problems and decrease emissions intensity. The NVIDIA computing stack is built for performance and energy-efficient compute—for example, our NVIDIA Vera CPU enables software environments to run up to 50% faster with 2x the efficiency of traditional CPU infrastructure, unlocking agentic AI at scale.⁴</p>	<p>AI can support improved grid flexibility, energy availability, and reliability by optimizing grid infrastructure. Power-flexible AI factories can help unlock up to 100 gigawatts of capacity across the U.S. power system by combining optimized infrastructure design with efficient use of existing assets and, where needed, new-build generation, while flexing during limited periods of grid stress to reduce the need for broader grid expansion to support reliability. These next-generation AI factories will harness the new NVIDIA Vera Rubin DSX AI Factory reference design, which includes the DSX Flex software library for connecting AI factories to power-grid services.</p>
	<p>To manage the greenhouse gas (GHG) emissions footprint of our data centers, labs, and offices, we strategically focus on working with key data center partners to optimize efficiency, managing our operations efficiently, and sourcing clean energy. Last year, we met our goal to purchase or generate enough clean electricity to match 100% of our global electricity usage for sites under our operational control. In FY26, we continued to match 100% of our global electricity usage with clean electricity, reducing Scope 2 market-based emissions.¹</p>	<p>AI helps weather and climate prediction be more accessible than ever for scientists, startups, developers, enterprises, and government agencies worldwide, offering the potential to improve climate resiliency. NVIDIA's Earth-2 is the first open, accelerated set of models and tools that enables developers to bring disparate weather and climate AI capabilities together. Accurate weather forecasting helps save lives and protect environments—and is a cornerstone of decision-making in agriculture, energy, public health and other industries.</p>

About This Report

We welcome feedback on this report and our performance. Send comments and suggestions to NVIDIACorporateSustainability@nvidia.com.

The discussion of topics included in this report and our other corporate responsibility disclosures should not be read as implying that such topics are “material” in the context of the U.S. federal securities laws, Delaware General Corporation Law, or any other regulatory framework, even where we use words such as “material” or “materiality.” Our approach to sustainability disclosures is informed by reporting frameworks, such as the GRI, that involve broader definitions of materiality than used for purposes of our compliance with SEC disclosure obligations. As a result, “materiality” for purposes of our sustainability reporting includes impacts on communities, the environment, and stakeholders such as employees, customers, and suppliers, and the inclusion of topics in our sustainability reporting, even when described as “material,” does not indicate that such topics are material to the Company’s business, operations, or financial condition.

NVIDIA’s reporting follows our fiscal calendar. FY26 data corresponds to the period from January 27, 2025 to January 25, 2026. The information contained in this report is accurate as of approximately June 12, 2026, unless stated otherwise. The information is subject to change, and NVIDIA will not necessarily disclose such changes. The information may be updated, amended, supplemented, or otherwise altered by subsequent reports or filings by NVIDIA.

Certain statements included or incorporated by reference in this report including, but not limited to, statements as to: expectations with respect to our growth and market opportunities and related trends and drivers; expectations with respect to the performance, impact, and benefits of our products, services and technologies and related trends and drivers; our strategies, priorities, goals, and objectives; expectations with respect to technology developments; projected market growth and trends; 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 our current expectations, estimates, and projections about our industry, our management’s beliefs and assumptions, and information currently available to management and are subject to risks and uncertainties that could cause results to be materially different than expectations.

These forward-looking statements are not guarantees of future performance and speak only as of the date hereof. Actual events may differ materially, perhaps adversely. In particular, goals and targets described in this report are aspirational and not guarantees or promises that the goals or targets will be met. Except as required by law, we disclaim any obligation to update these forward-looking statements to reflect future events or circumstances.

Important factors that could cause actual results to differ materially include: global economic and political 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 products and technologies; market acceptance of our products or our partners’ products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; our ability to realize the potential benefits of business investments or acquisitions; and changes in applicable laws and regulations, as well as other factors detailed from time to time in the most recent reports we file with the Securities and Exchange Commission including, but not limited to, our annual report on Form 10-K and quarterly reports on Form 10-Q.

Historical and forward-looking statements contained in this report may be based on standards for measuring progress that are still developing, internal controls and processes that continue to evolve, and assumptions that are subject to change in the future. Such historical and forward-looking statements may be subject to measurement uncertainties resulting from limitations inherent in the nature and the methods used for determining such information and may include estimates or approximations. We believe that such estimates are appropriate and reasonable; however, due to inherent uncertainties in making estimates and assumptions, actual results could differ from the original estimates. The precision of different measurement techniques may also vary. As we improve our methodologies and as new information becomes available, we may continue to revise our estimates and assumptions.

This report also includes certain information, including with respect to emissions factors, that is obtained from published sources or third parties. The accuracy and completeness of such information are not guaranteed. Such information is subject to assumptions, estimates, and other uncertainties, and we have not independently verified this information.

Endnotes

1. NVIDIA defines clean electricity as electricity generated from non-fossil fuel sources, including solar, wind, and other low-carbon technologies. The term refers to emissions at the point of generation and does not address full life cycle impacts. More information on this calculation can be found in the [Report of Independent Accountants](#).
2. Methodology: Energy efficiency was estimated using DLSim analytical projections, reported as modeled performance per megawatt for inference for NVIDIA Vera Rubin NVL72 (NVFP4) + Groq 3 LPX (FP8) and GB200 NVL72 (NVFP4) running the DLSIM model (GPT MoE 2 trillion, 400K) using common modeling assumptions. Results are based on NVIDIA internal simulations, may differ from measured silicon results and other deployments, and will not be updated over time.
3. Methodology: Energy efficiency was estimated using DLSim analytical projections, reported as modeled performance per megawatt for inference for NVIDIA Vera Rubin NVL72 and GB200 NVL72 running Kimi-K2 Thinking (32K ISL / 8K OSL) at NVFP4 precision using common modeling assumptions. Results are based on NVIDIA internal simulations, may differ from measured silicon results and other deployments, and will not be updated over time.
4. Methodology: Energy efficiency was estimated using industry-standard, complex software benchmarks including compiler workloads (LLVM/Clang, GCC) and data analytics (DuckDB). Comparisons were made between NVIDIA Vera stand-alone racks and x86-based racks running the same workloads. Results will not be updated over time.
5. Methodology: Energy efficiency was estimated using partner projections for performance per watt (Gbps/W/GPU) for NVIDIA Vera with CX-9 vs. the latest x86-based storage offering. Results will not be updated over time.