NVIDIA CORPORATION

2026 ANNUAL MEETING OF STOCKHOLDERS QUESTIONS & ANSWERS

Below are questions we received prior to and during the NVIDIA Corporation 2026 Annual Meeting of Stockholders held on June 24, 2026. We have consolidated multiple questions and comments on the same or similar topic and provided a summary response, and we have limited each stockholder to one question. We have also removed questions and comments that presented general, economic, political, or other views that were not directly related to our business, the business of the meeting, or for matters of individual concern.

All responses, including any forward-looking statements, are made as of June 24, 2026, unless otherwise noted. We do not undertake, and expressly disclaim any duty or obligation, to update these responses, whether as a result of new information, new developments, or otherwise, except to the extent that disclosure may be required by law. Forward-looking statements are subject to many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in these forward-looking statements, as discussed in our SEC filings. Investors should review our most recent reports on Forms 8-K, 10-K, and 10-Q on the SEC's website at www.sec.gov for information regarding our business and factors affecting our business since June 24, 2026.

Thank you for your questions and comments and for your ongoing support of NVIDIA Corporation.

Q: How sustainable is the current AI infrastructure buildout? At what point could the principal drivers of your business mature and growth slow?

A: Our response to this question appeared at 34:55 in the 2026 Annual Meeting webcast.

Q: NVIDIA's GPUs now power the majority of the world's AI training infrastructure, but as inference workloads overtake training, how confident is NVIDIA that the GPU architecture remains the preferred platform for inference at scale?

A: Our response to this question appeared at 37:25 in the 2026 Annual Meeting webcast.

Q: Is supporting frontier systems worth the national security and reputational risk?

A: Our response to this question appeared at 39:45 in the 2026 Annual Meeting webcast.

Q: Is NVIDIA concerned that its products reach restricted or adversary end-users through diversion?

A: Our response to this question appeared at 42:30 in the 2026 Annual Meeting webcast.

Q: Why has NVIDIA offered opinions to the U.S. government regarding export controls? Is there a formal plan in place for selling data center compute to the China market?

A: Our response to this question appeared at 44:03 in the 2026 Annual Meeting webcast.

Q: Can you provide an update on your capital return policy? What should we expect in terms of dividends and share repurchases?

A: Our response to this question appeared at 45:52 in the 2026 Annual Meeting webcast.

Q: How is NVIDIA using AI internally?

A: NVIDIA uses AI across nearly every business function. Companywide there are already hundreds of digital workers working alongside our employees. Some examples include:

- Our hardware engineering teams use an NVIDIA-developed agentic platform called ChipNemo for chip design, verification testing, and more. Today, nearly 300 specialized AI agents are used by over 6,000 hardware engineers.

- 100% of our software engineers are using AI-powered code assistants to boost productivity and accelerate development cycles.

- We've built AI-powered search, summarization, and deduplication into our internal software bug tracking system, saving over 16,000 developer hours in just a few months.

- AI agents help our operations team run what-if scenarios, optimize supply-demand matching, and dynamically generate reports for better planning decisions. What-if scenarios that used to take 3-4 hours now run in 15 minutes.

We expect these digital workers to continue to grow in numbers - to thousands and eventually millions of AI Agents, supporting thousands of use cases. To enable this scale, we've built and deployed an internal AI Factory — a secure, on-prem platform that empowers teams to build, run, and iterate Agentic AI workflows. Internally, our inference hub now processes over 7T tokens/month across over 200+ agentic AI workflows that are in production.

Q: In certain communities in the US there seems to be growing backlash against data centers. What are you doing to increase citizen support? What benefits would you highlight?

A: While we acknowledge community concerns, transitioning computing workloads to NVIDIA's accelerated computing is actually the sustainable path forward. Our technology stack compresses massive AI workloads into dense systems, setting new standards for hardware efficiency while accelerating clean energy discovery and powering advanced climate simulations like Earth-2. Furthermore, we are establishing a blueprint for AI factories to be flexible, grid-responsive assets that can modulate their power demand, which helps communities avoid blackouts and can actually lower local utility costs by spreading fixed grid expenses.

AI factories built on NVIDIA's Vera Rubin platform can be assets to communities. AI factories create valuable jobs, from trades to high-tech roles. They also provide significant tax revenues that can support local schools, community services and infrastructure. Through extreme co-design, we have been working hard to minimize environmental impacts in our DSX platform. Our S45 warm liquid cooling system for Vera Rubin essentially eliminates water consumption for datacenter operation in many locations, and it can also reduce noise by eliminating chillers. Vera Rubin delivers up 10x higher inference throughput per watt than Blackwell, so the same electricity generates even more tokens and more value.

We are actively working to deliver these results and to share them with policymakers, media and the public. We coordinate with partners to highlight joint successes and the positive connection we see between AI and sustainability, and we are optimistic that the local benefits of AI will become increasingly clear as DSX and Vera Rubin are rolled out.

Q: Which Board committee is responsible for assessing export-control compliance and chip-diversion risk, AI safety, and catastrophic-risk scenarios?

A: NVIDIA’s Nominating and Corporate Governance Committee and the entire Board work with NVIDIA’s executive management to oversee export control compliance, chip diversion risk, AI safety, and catastrophic risk scenarios. These responsibilities, and our broader risk management framework, are described in detail in our SEC filings. These topics are reviewed regularly at the Board and committee level and are integrated into our overall enterprise risk management processes.

This Board level oversight reflects the strategic importance of AI, export controls, and safety to NVIDIA’s long-term performance and to the interests of our shareholders.

Q: When may the stock split again?

A: In fiscal 2025, we completed a 10 for 1 split of our stock. Our board reviews our capital return program regularly and would consider a stock split if it is in the best interest of our stockholders.

Q: What is the expected quarterly stock dividend for this year and next?

A: Our quarterly dividend currently is $0.25 per share. We plan to regularly grow our dividend at a modest pace over time.

Q: What is the feasibility of selling controlled products into China in the next 12-24 months? What is the company waiting on in order to execute sales?

A: Please see our earnings materials from May 20, 2026 for information on our sales to China.

Q: Who do you formally recognize as your biggest competitor? What is your plan to mitigate competitive threats?

A: Please refer to our public filings with the SEC, including our Annual Reports on Form 10-K and Quarterly Reports on Form 10-Q.

Q: Has the board thought to award NVIDIA stock to USA military shareholders (active duty and veterans)?

A: We currently have no such plans.

Q: Regarding the future of the business, what are the main risks management is focused on?

A: We're focused on executing against our product roadmap and scaling our company and ecosystem to address our large long-term growth opportunities.

Q: What are the prospects for NVIDIA's stock performance over the next two years, and what factors could influence its performance?

A: We do not predict our stock price performance. We are focused on creating long-term shareholder value through superior product innovation, market development efforts, strong execution, and shareholder capital returns.

Q: Does NVIDIA provide products to support Rivian vehicle production?

A: Rivian and NVIDIA have partnered to bring advanced AI-powered driving capabilities to Rivian vehicles. Rivian trains its AI models running in the car using NVIDIA GPUs in the cloud.

Rivian also uses NVIDIA RTX PRO GPUs for engineering, visualization, and XR-based vehicle design workflows. NVIDIA's full stack computing platform addresses the entire vehicle lifecycle including design, data center AI development, simulation, AI-powered manufacturing, and in-car runtime computer.

Q: Is NVIDIA communicating enough to shareholders?

A: NVIDIA communicates extensively with financial analysts, investors, and shareholders at events throughout the year including GTC conferences, earnings calls, industry and financial conferences, fireside chats, podcasts, press releases, blogs, non-deal roadshows, and virtual and in-person meetings. GTC keynotes are viewed by tens of millions of people. NVIDIA management and investor relations engage with thousands of financial analysts, investors, and shareholders every quarter. For additional information on NVIDIA's news and events please visit our Investor Relations and newsroom websites.

Q: How has AI impacted employee compensation?

A: AI doesn't impact employee compensation directly. At NVIDIA, managers play an important role in determining employee compensation based on performance. However, AI enhances employee productivity and performance which, over time, could lead to employee advancement and higher compensation.

Q: How do I go about selling NVIDIA stock?

A: Please contact your broker for additional information on trading NVIDIA shares.

Q: Any current M&A plans?

A: We frequently review potential acquisition opportunities. These opportunities can enhance our platform with key capabilities and technologies and add key talent and engineering resources.

Q: Why are you trying to sell chips to China?

A: US national security is paramount. To promote American leadership globally, it is important for the American technology stack to compete and win globally. Ceding the world's second largest AI market to competitors is antithetical to this goal.

Q: What is NVIDIA Investor Relations doing to attract and build stronger relationships with large, long-duration, value-oriented institutional investors?

A: NVIDIA management and investor relations engage with thousands of investors quarterly including those who are value oriented. On our recent earnings call, we announced a 25x increase to our quarterly dividend and authorized an additional $80 billion in share repurchases. Supported by our conviction in sustainable market growth and free cash flow generation, we plan to return to stockholders 50% or more of our free cash flow this year, next year, and beyond – growing both our share repurchases and dividend over time.

Q: How should shareholders reconcile management's highly positive public commentary about the potential equity upside of ecosystem partners with the investment case for continuing to hold NVIDIA stock?

A: AI, robotics, and accelerated computing are a once-in-a-generation growth opportunity measured in the trillions of dollars. We view this as an opportunity for NVIDIA as well as our ecosystem partners. Our confidence in our long-term stock price is underpinned by the long runway of growth ahead of us. We are enhancing shareholder returns by leveraging our strong margins, earnings growth, and cash flow generation to drive higher capital returns over time in the form of increasing stock repurchases and dividends.

Q: Has NVIDIA considered or is NVIDIA addressing the waste management industry? Could you please share more about whether NVIDIA's technologies can benefit this field?

A: Our accelerated computing and AI platforms address many industries, including waste management. For example, AMP Robotics uses NVIDIA's AI platform to automate the identification and sorting of materials in recycling facilities. These facilities require one-third fewer conveyor belts while increasing waste recovery rates from 75% at conventional plants to approximately 90%. This results in increased efficiency, reduced cost, improved material recovery, and reduced waste sent to landfills. Other NVIDIA technologies applicable to waste management use cases include our Jetson and Metropolis physical AI and edge AI platforms.

Q: Please advise the status of the class-action cryptocurrency-related lawsuit.

A: For the status of material litigation matters, please refer to our public filings with the SEC, including our Annual Reports on Form 10-K and Quarterly Reports on Form 10-Q.

Q: What is the growth path to $1T+ in annual revenue in 3 to 4 years? Will NVIDIA provide long term EPS guidance?

A: We provide an outlook for one quarter at a time. We have a multi-trillion dollar long term opportunity across our businesses powered by a diverse application and customer set including graphics, hyperscale, AI Cloud, enterprise, Sovereign AI, physical AI, and edge computing.

Q: What additional business opportunities, if any, did Jensen's trip to China with the President generate?

A: Jensen traveled to China to support the President and advocate for the right to compete globally, a requisite for American technology leadership.

Q: What is management's strategy to diversify the customer base?

A: Our business continues to diversify. Every major AI lab and thousands of enterprises across major industries including healthcare, financial services, and manufacturing are building and running on our platform. We continue to add major new inference customers including Anthropic and Apple. Our Sovereign AI revenue increased 3x last fiscal year and nearly forty countries representing $50 trillion in GDP are building NVIDIA-powered sovereign AI infrastructure. Our AI Cloud revenue for the fiscal first quarter of 2027 more than tripled year-over-year and the number of partner data centers exceeding 10MW nearly doubled, now surpassing 80 sites. Finally, over forty thousand AI startups are innovating on our platform including many physical AI and robotics companies who will further diversify our customer base as they ramp and scale.

Q: What is NVIDIA planning to do if the AI industry is not a long-term viable market?

A: AI is one of over 7,000 applications that run on NVIDIA's platform. Because Moore's Law has slowed, an increasing number of major applications outside of AI are moving to accelerated computing including data processing, computational lithography, and medical imaging. This alone can fuel demand for years to come. The inherent programmable nature of our GPU architecture combined with CUDA's development platform will enable NVIDIA to grow its addressable application set and long-term revenue opportunity over time. We estimate accelerating traditional applications will comprise roughly half of our long-term growth opportunity.

Q: Is NVIDIA working on analog compute solutions like those being done by Mythic.AI and Bittally Systems?

A: We have done work in this area and continually review external approaches. Our conclusion is that the approaches we've considered and the external approaches we've reviewed are not competitive with our best digital approaches when the entire system is considered.