

# Emotigraph Analysis Report

Toshiya Hari: Good afternoon, everyone, and...

Generated: 2/26/2026 at 8:13:14 AM

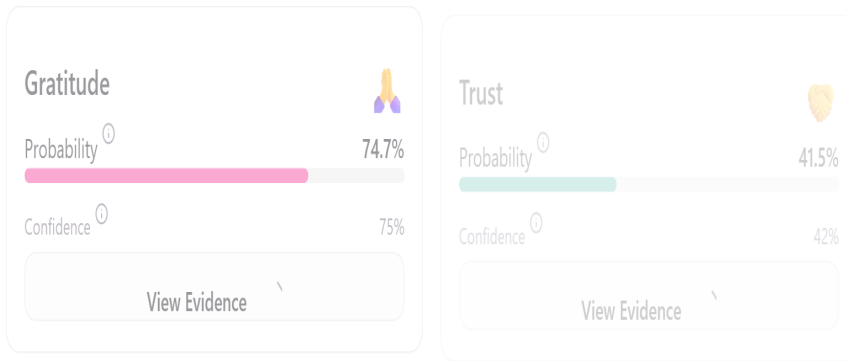
Analysis Date: 2/26/2026

Word Count: 6566 | Sentences: 423

## Overall Sentiment Dimensions




# Top Emotions




# Emotion Distribution

All detected emotions ranked by probability



# Emotion Spectrum

Distribution of emotions by intensity and category



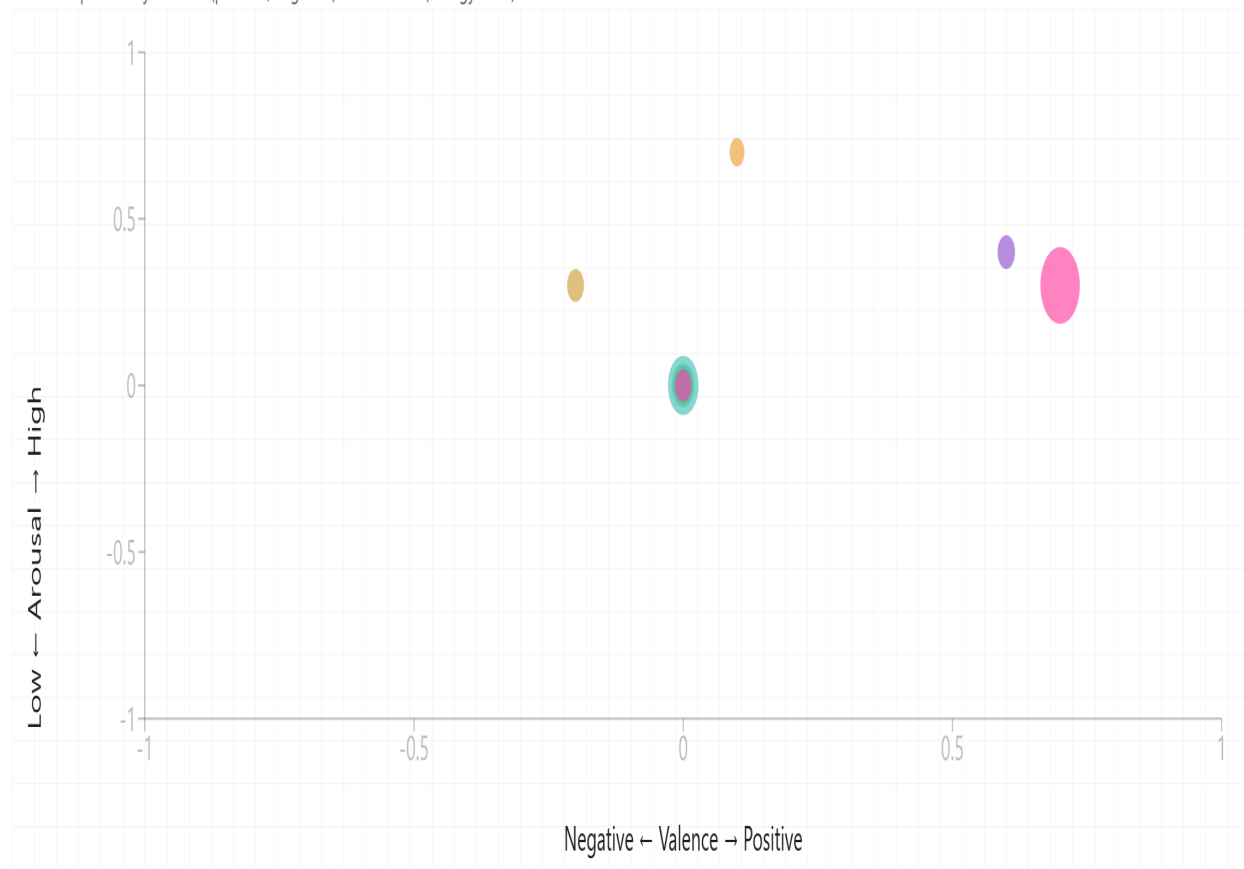
# Emotional Constellation

Emotions mapped by valence and arousal dimensions

# Emotional Landscape



Emotions plotted by valence (positive/negative) and arousal (energy level)



Low energy, Low energy,  
Negative: Depressed Positive: Calm

High energy, High energy,  
Negative: Tense Positive: Excited

# Emotional Flow Through Text

How emotions change throughout the document

## Emotional Arc

How emotions change throughout the document (6566 words)



# Color-Coded Emotional Text

Words underlined by their emotional associations

**Emotion Annotated Text**

Color Legend: Red (Anger), Blue (Sadness), Green (Happiness), Orange (Surprise), Purple (Disgust)

The following text is color-coded based on emotional associations. Words are underlined to indicate their emotional context. The text is a dense paragraph of placeholder text, where words are underlined and colored according to the legend above. The colors used include red, blue, green, orange, and purple, corresponding to the emotions listed in the legend.

Page 6 of 23  
Generated by Emotigraph - sentiment.ws

# Emotion-Annotated Text

Original text with words color-coded by their emotional associations

## Color Legend:

- Gratitude
- Hope
- Admiration
- Anticipation
- Trust
- Satisfaction
- Confusion

Toshiya Hari: Good afternoon, everyone, and welcome to NVIDIA Corporation's conference call for 2026. With me today from NVIDIA Corporation are Jensen Huang, president and chief executive officer, and Colette Kress, executive vice president and chief financial officer. Our call is being webcast live on NVIDIA Corporation's website at [investors.nvidia.com](https://investors.nvidia.com). The content of today's call is NVIDIA Corporation's property. It cannot be reproduced or transcribed without our prior written consent. During this call, we may make forward-looking statements based on current expectations. These are subject to a number of significant risks and uncertainties, and our actual results may differ materially.

For a discussion of factors that could affect our future financial results and business, please refer to the disclosure in today's earnings release, our most recent Forms 10-K and 10-Q, and the reports that we may file on Form 8-K with the Securities and Exchange Commission. All our statements are made as of today, 02/25/2026, based on information currently available to us. Except as required by law, we assume no obligation to update any such statements. During this call, we will discuss non-GAAP financial measures. You can find a reconciliation of these non-GAAP financial measures to GAAP financial measures in our CFO commentary, which is posted on our website. With that, let me turn the call over to Colette.

Colette Kress: Thanks, Toshiya. We delivered another outstanding quarter, with record revenue, operating income, and free cash flow. Total revenue of \$68,000,000,000 was up 73% year over year, accelerating from Q3. Growth on a sequential basis was also a record as we added \$11,000,000,000 in Data Center revenue across a diverse and expanding set of customers including cloud providers, hyperscalers, AI model makers, enterprises, and sovereign nations. Demand for our Blackwell architecture, extreme codesign at data center scale, continues to strengthen as inference deployments grow in addition to training. The transition to accelerated computing and the infusion of AI across existing hyperscale workloads continue to fuel our growth.

Agentic and physical AI applications built on increasingly smarter and multimodal models are beginning to drive our financial performance.

On a full-year basis, Data Center generated revenue of \$194,000,000,000, up 68% year over year. We have now scaled our Data Center business by nearly 13x since the emergence of ChatGPT in fiscal 2023. We look ahead. We expect sequential revenue growth throughout calendar 2026, exceeding what was included in the \$500,000,000,000 Blackwell and Rubin revenue opportunity we shared last year. We believe we have inventory and supply commitments in place to address future demand, including shipments extending into calendar 2027. Every data center is power constrained. Customers make critical architectural decisions

based on performance per watt given these constraints and the need to maximize AI factory revenue.

SemiAnalysis declared NVIDIA Corporation inference king as recent results from InferenceX reinforced our inference leadership, with GB300 and NVL72 achieving up to 50x performance per watt and 35x lower cost per token, compared with offerings. And continuous optimization of CUDA software helped deliver up to 5x better performance on GV200 and NVL72, just within four months. NVIDIA Corporation produces the lowest cost per token, and data centers running on NVIDIA Corporation generate the highest revenues. Our pace of innovation, particularly at our scale, is unmatched.

Fueled by an annual R&D budget approaching \$20,000,000,000 and our ability to extreme codesign across compute and networking across chips, systems, algorithms, and software, we intend to deliver x-factor leaps in performance per watt every generation and extend our leadership position over the long term.

Q4 Data Center revenue of \$62,000,000,000 increased 75% year over year and 22% sequentially, driven primarily by sustained strength in Blackwell and the Blackwell Ultra ramp. With NVIDIA Corporation infrastructure in high demand, even Hopper and much of the six-year-old Ampere-based products are sold out in the cloud. Nearly a year has passed since the release of our Grace Blackwell GB200 NVL72 systems. Today nearly 9 gigawatts of infrastructure on Blackwell are deployed and consumed by the major cloud service providers, hyperscalers, AI model makers, and enterprises. Networking, a cornerstone of our data center-scale infrastructure offering, was a standout this quarter, generating \$11,000,000,000 in revenue, up more than 3.5x year over year.

Demand for our scale-up and scale-out technologies reached record levels, both growing double digits sequentially, driven by strong adoption of NVL72 scale-up switches as Grace Blackwell systems accounted for roughly two-thirds of Data Center revenue in the quarter. NVLink scale-up fabric has revolutionized computing and demonstrates the power of extreme codesign across all of the chips of the supercomputer and the full stack. In Q4, we announced that we will enable AWS with NVLink to integrate with their custom silicon. Momentum is strong with our Spectrum X Ethernet scale-up and scale-across networking as customers work to unify distributed data centers into integrated gigascale AI factories.

For the full year, our Networking business exceeded \$31,000,000,000 in revenue, up more than 10x compared to fiscal 2021, the year we acquired Mellanox.

Our demand profile is broad, diverse, and expanding beyond just chatbots. First, there is a fundamental platform shift from classical machine learning to generative AI. Strong evidence of ROI as hyperscalers upgrade massive traditional workloads to generative AI, including search, ad generation, and content recommender systems, is encouraging our largest customers to accelerate their capital spending. For example, at Meta, advancements in their GEM model drove a 3.5x increase in ad clicks on Facebook and more than a 1% gain in conversions on Instagram, translating into meaningful revenue growth. With the same NVIDIA Corporation infrastructure, Meta Superintelligence Labs can train and deploy their frontier agentic AI systems. Frontier agentic systems have reached an inflection point.

Claude Code, Claude Cework, and OpenAI Codex have achieved useful intelligence. Adoption is skyrocketing, and tokens are profitable, driving extreme urgency to scale up compute. Compute directly translates to intelligence and revenue growth.

Analyst expectations for 2026 CapEx across the top five cloud providers and hyperscalers, who collectively account for a little over 50% of our Data Center revenue, are up nearly \$120,000,000,000 since the start of the year and approaching \$700,000,000,000. We continue to expect the transition of classic data center workloads to GPU-accelerated computing and the use of AI to enhance today's hyperscale workloads to contribute toward roughly half of our long-term opportunity. Every country will build and operate some parts of its AI infrastructure just like with electricity and the Internet today.

In fiscal year 2026, our Sovereign AI business more than tripled year over year to over \$30,000,000,000, driven primarily by customers based in Canada, France, the Netherlands, Singapore, and the UK. Over the long run, we expect our sovereign opportunity to grow at least in line with the AI infrastructure market, as countries spend on AI proportional to their GDP. While small amounts of H200 products for China-based customers were approved by the U.S. government, we have yet to generate any revenue, and we do not know whether any imports will be allowed into China.

Our competitors in China, bolstered by recent IPOs, are making progress and have the potential to disrupt the structure of the global AI industry over the long term. To sustain its leadership position in AI compute, America must engage every developer and be the platform of choice for every commercial business, including those in China. We will continue to engage with the U.S. and China governments and advocate for America's ability to compete around the world.

We unveiled the Rubin platform last month at CES, comprised of six new chips: the Vera CPU, Rubin GPU, NVLink 6 switch, ConnectX-9 SuperNIC, BlueField-4 DPU, and Spectrum-6 Ethernet switch. The platform will train MoE models with one-fourth the number of GPUs and reduce inference token costs by up to 10x compared to Blackwell. We shipped our first Vera Rubin samples to customers earlier this week, and we remain on track to commence production shipments in the second half of the year. Based on its modular, cable-free tray design, Rubin will deliver improved resiliency and serviceability relative to Blackwell. We expect every cloud model builder to deploy Vera Rubin.

Moving to Gaming. Gaming revenue of \$3,700,000,000 increased 47% year on year, driven by strong Blackwell demand and improved supply. GeForce RTX is the leading platform for PC gamers, creators, and developers. In Q4, we added several new technologies and advancements, including DLSS 4.5, which uses AI to bring game visuals to a new level, and 35% faster LLM inference across leading AI PC frameworks. Looking ahead, while end demand for our products remains strong and channel inventory levels are healthy, we expect supply constraints to be the headwind to Gaming in Q1 and beyond.

For Professional Visualization, it crossed the \$1,000,000,000 mark for the first time, with revenue of \$1,300,000,000, up 159% year over year and 74% sequentially. During the quarter, we launched the RTX Pro 5000 Blackwell workstation with 72 GB of fast memory for AI developers running LLMs and agentic

workflows. Automotive revenue of \$604,000,000 was up 6% year over year and was driven by robust demand for self-driving solutions. At CES, we introduced Alpamayo, the world's first open portfolio of reasoning, vision, language, action models, simulation blueprints, and datasets enabling vehicles that can think. The first passenger car featuring Alpamayo, built on NVIDIA Corporation Drive, will be on the road soon in the new Mercedes-Benz CLA.

Physical AI is here, having already contributed north of \$6,000,000,000 in NVIDIA Corporation revenue in fiscal year 2026. Robotaxi rides are growing exponentially, with commercial fleets from Waymo, Tesla, Uber, WeRide, and Zoox, and many others are expected to scale from thousands of vehicles in 2025 to millions over the next decade, creating a market poised to generate hundreds of billions of dollars of revenue. This expansion will demand orders of magnitude more compute, with every major OEM and service provider developing on NVIDIA Corporation's platform.

We continue to advance robotics development with the new NVIDIA Corporation Cosmos and Isaac GR00T open models and frameworks, and NVIDIA Corporation-powered robots and autonomous machines for leading companies, including Boston Dynamics, Caterpillar, FANUC Robotics, LG Electronics, and Neurorobotics. To accelerate industrial physical AI adoption, we also announced new expanding partnerships with Dassault Systèmes, Siemens, and Synopsys to bring NVIDIA Corporation AI Infrastructure, Omniverse Digital Twins, world models, and CUDA-X libraries to millions of researchers, designers, and engineers building the world's industries.

Let's move to the rest of the P&L. GAAP gross margin was 75% and non-GAAP gross margin was 75.2%, increasing sequentially as Blackwell continued to ramp. GAAP operating expenses were up 16% sequentially and up 21% on a non-GAAP basis related to new product introductions and compute and infrastructure costs. Non-GAAP effective tax rate for the fourth quarter was 15.4%, below our outlook for the quarter, primarily due to the impact of a one-time tax benefit. Inventory grew 8% quarter over quarter while purchase commitments also increased significantly as we have strategically secured inventory and capacity to meet demand beyond the next several quarters. This is further out in time than usual and reflects the longer demand visibility we have.

While we expect tightness in the supply for our advanced architectures to persist, we remain confident in our ability to capitalize on the growth opportunity ahead with our scale, expansive supply chain, and long-standing partnerships continuing to serve us well. We generated free cash flow of \$35,000,000,000 in Q4 and \$97,000,000,000 in fiscal year 2026. For the year, we returned \$41,000,000,000, or 43% of free cash flow, to our shareholders in the form of share repurchases and dividends. We continue to invest in our technology and our ecosystem to cultivate market development, drive long-term growth, and ultimately yield total shareholder returns superior to the market or our peer group.

Importantly, we will continue to run a strategic and disciplined process as it relates to our investments, and we remain committed to returning capital to our shareholders.

Let me turn to the outlook for the first quarter. Starting this quarter, we will be including stock-based

compensation expense in our non-GAAP results. Stock-based compensation is a foundational component of our compensation program to attract and retain world-class talent. Let me first start with revenue. Total revenue is expected to be \$78,000,000,000, plus or minus 2%. We expect most of our growth to be driven by Data Center. Consistent with last quarter, we are not assuming any Data Center compute revenue from China in our outlook. GAAP and non-GAAP gross margins are expected to be 74.9% and 75%, respectively, plus or minus 50 basis points. For the full year, we continue to see gross margins in the mid-70s.

We will keep you updated on our progress as we prepare for the Vera Rubin transition. GAAP and non-GAAP operating expenses are expected to be approximately \$7,700,000,000 and \$7,500,000,000, respectively, including stock-based compensation expense of \$1,900,000,000. For the full fiscal year 2027, we expect GAAP and non-GAAP tax rates to be between 7% and 19%, excluding any discrete items and material changes to our tax environment. With that, let me turn the call over to Jensen. I think he has a few words for us.

Toshiya Hari: This quarter, we significantly deepened and expanded our partnerships with leading frontier model makers.

Jensen Huang: We recently celebrated OpenAI's launch of GPT-5.3 Codex, trained with and inferencing on Grace Blackwell and NVLink 72 systems. GPT-5.3 Codex can take on long-running tasks that involve research, tool use, and complex execution. 5.3 Codex is deployed broadly inside NVIDIA Corporation. Our engineers love it. We continue to work with OpenAI toward a partnership agreement and believe we are close. We are thrilled with our ongoing partnership with OpenAI, a once-in-a-generation company we have had the pleasure of partnering with since their first days. Meta Superintelligence Labs is scaling up at lightning speed. Last week, we announced that Meta is deploying millions of Blackwell and Rubin GPUs, NVIDIA Corporation CPUs, and Spectrum X Ethernet for training and inference.

This quarter, we announced a partnership with Anthropic and a \$10,000,000,000 investment in their company. Anthropic will train and inference on Grace Blackwell and Vera Rubin systems. Anthropic's Claude Cowork agent platform is revolutionary and has opened the floodgates for enterprise AI adoption. Between Claude Cowork and OpenAI, compute demand is skyrocketing, and the ChatGPT moment of agentic AI has arrived. With partnerships spanning Anthropic, Meta, OpenAI, and xAI, NVIDIA Corporation is deployed across every cloud, and with our ability to build full-stack AI infrastructure from the ground up or support them in the cloud, we are uniquely positioned to partner with frontier model builders at every stage: training, inference, and AI factory scale-out.

Finally, we recently entered into a non-licensing agreement with Groq for its low-latency inference technology and welcomed a team of brilliant engineers to NVIDIA Corporation. As we did with Mellanox, we will extend NVIDIA Corporation's architecture with Groq's innovations to enable new levels of AI infrastructure performance and value. We look forward to sharing more at GTC next month. K. Back to you.

Toshiya Hari: We will now open for questions. Operator, please poll for questions.

Sarah: At this time, I would like to remind everyone, in order to ask a question, press star then the number one on your telephone keypad. Your first question comes from Vivek Arya with Bank of America Securities.

Vivek Arya: Thanks for taking my question. I think you mentioned that you now have growth visibility into calendar 2027 also, and I think your purchase commitments kind of reflect that confidence. But, Jensen, I am curious. You know, when you look at your top cloud customers, cloud CapEx close to \$700,000,000,000 this year, many investors are concerned that it would be harder for this level to grow into next year, and for several of them, their cash flow generation capability is also getting compressed. So I know you are very confident about your road map, right, and your purchase commitments and whatnot. But how confident are you about your customers' ability to continue to grow their CapEx?

And if their CapEx does not grow, can NVIDIA Corporation still find a way to grow in that envelope? Thank you.

Jensen Huang: I am confident in their cash flow growing. And the reason for that is very simple. We have now seen the inflection of agentic AI and the usefulness of agents across the world and enterprises everywhere. You are seeing incredible compute demand because of it. This new world of AI, compute is revenues. Without compute, there is no way to generate tokens. Without tokens, there is no way to grow revenues. So in this new world of AI, compute equals revenues. And I am certain that at this point, with all of the—there is a modest amount of computers, you know, call it \$300,000,000,000 or \$400,000,000,000 worth of cash. Great. Thank you.

And about some of the strategic investments and potentially OpenAI. Core with this model, but also partner on NVIDIA Corporation's platform. We are in every cloud. We are in every center. We are all over the world. Edge and dozens of AI natives are built on top of AI ecosystem. For language, or physical AI, or AI physics. All of these ecosystems to be built on top of NVIDIA Corporation. To invest into the ecosystem across the entire stack. Sure, today than it used to be. Partially, take down we have compute every aspect of that. And getting to AI models or deep. And as I mentioned before, scale up. Scale up. GPUs, nodes of and each of them.

That we do. The rack is really quite incredible. We also love the low work out, of course. And some people want to integrate that extends Ethernet with artificial intelligence way incredibly good at that. Our Spectrum X performance really shows it. Billion-dollar AI factory for your data center, that is NVIDIA Corporation's Networking business is really Boston Interface. Every time you cross an interface, you add latency, you add power unnecessarily. We are not allergic to Violet. We use Violets already. But we try to do so. And so when you look at the Grace Blackwell architecture and the Rubin architecture, use two-time radical limited ties into the bottom, and that reduces the amount architecture. Of the competitors.

If you look at our software advantage, but where software starts and architecture starts and ends is kind of hard to tell. It is—our software is effective because our architecture is so good. And so the CUDA architecture is unquestionably more effective, more efficient, I mean, per FLOP per watt. Generations of architectures that GPUs will all benefit. And so we will continue to do that, and it allows us to extend the useful life, allows us to

have innovation, flexibility, and velocity, which translates—and very importantly—performance customers.

And so what we will do with Groq is—you will come to see, come to see GTC—but what we will do is we will extend our architecture with Groq as an accel, extended NVIDIA Corporation's architecture with Mellanox.

Sarah: The next question comes from Stacy Rasgon.

Stacy Rasgon: Thanks for taking my questions.

Colette Kress: This is a very great architecture that hopes that just today quickly standing up, and have already planned on many different orders across the different customers to determine how beginning ramp will start in the second half. And with Fusion, in terms of the strong demand, too early for us to know at this time. We will get back to you as soon as we can.

Sarah: Your next question comes from Atif Malik with Citi.

Atif Malik: Thank you for taking my question. Jensen, can you talk about the importance of CUDA as more of the investment dollars in AI are coming from inference workloads?

Jensen Huang: No—the entire stack from TensorRT-LLM that we introduced, a most performant inference stack, optimizing it for NVL, requires us to discovery and invent new parallelization algorithms that sits on top of CUDA. Is to distribute the workload and the inferencing, take advantage of the aggregate. And so these systems, these agentic systems, are spawning off different agents, working with the team, then number of tokens that are being generated is really, really gone exponential. And so we need to inference at a much higher speed. And when you are inferencing at a much higher speed and each one of those tokens are dollarized, it directly translates into revenues. And so it equals revenues for our customer.

The data centers, which translates in a gigawatt directly to revenues. And so you could see that every CSP understands this now, every hyperscaler understands this. That CapEx translates to compute. Compute with the right architecture translates to maximizing revenues, and compute equals revenues. Without investing capacity today, without investing in compute, there cannot be revenue growth. And that, I think, everybody understands. Compute equals revenues. Architecture is incredibly important. It is more than strategic now. It directly affects their earnings. And choosing the right architecture, the one with the best performance per watt, is literally everything.

Sarah: Your next question comes from Ben Reitzes with Melius Research.

Ben Reitzes: Yeah. Hey. Thanks. First, let me say kudos on including this in non-GAAP. I think that is great. Just on gross margin—margins typically in the mid-70s long term—should we read into the visibility on supply being available in calendar 2027 that it is sustainable until then? And then, Jensen, what about after that? Are there innovations in memory consumption you can unveil that make us feel better about the ability to keep margins at that level for a long time?

Colette Kress: The single most important lever of our gross margins is actually delivering generational leads to our customers. That is the single most important thing. If we could deliver generationally, performance per watt that exceeds dramatically what Moore's Law can do, if we can deliver performance per dollar, dramatically more than the cost of our systems and the price of our systems, then can continue to sustain our gross margins. That is the simple, most important concept. Every—the reason why we are moving so fast is because, number one, the demand for tokens in the world as a result of the inflection points that we have gone through has now gone completely exponential. I think we are all seeing that.

To the point where even our six-year-old GPUs in the cloud are completely concerned that the pricing is coming up. And so we know that the amount of computation necessary, the amount of compute necessary, a modern way of doing software is growing exponentially. And so our strategy is to deliver an entire AI infrastructure every single year. This year, we introduced six new chips. Rubin next generation will do every single generation. We are committed to deliver many x-factors of performance per watt and performance per dollar. And that pace, our ability to do extreme codesign, allows us to deliver that value and that benefit to the customer.

That is the single most vital thing as it relates to our value delivery.

Toshiya Hari: Your next question comes from Antoine Chiketan with New Street Research.

Antoine Chiketan: Hi. Thanks for taking my questions. I would like to ask about space data centers, which some of your customers are considering. How feasible do you think that is and what kind of horizon? And what do the economics look like today, and how do you think that evolves over time?

Jensen Huang: The economics are poor today, but it is going to improve over time. As you know, the way that space works is radically different than how it works down here. There is an abundance of energy. But solar panels are large. But there is plenty of space in space. The energy dissipation, it is cold in space. However, there is no airflow, and so the only way to dissipate heat is through conduction. And the radiators that you need to create are fairly large. Liquid cooling is obviously out of the question because it is heavy and, you know, pre-resistant.

And so the methods that we use here on Earth are a little different than the way we would do it in space. But there are many different competing models that really want to be done in space. And so MPS is already the world's first GPU in space, Hopper is in space. And one of the best use cases of GPUs in space is imaging.

To be able to image at extremely high resolutions using, of course, optics and artificial intelligence, to be able to do that computation of reprojection of different angles and be able to up-res and do noise reduction and just be able to see—be able to image at very high resolutions, extremely large scales, and very, very fast. It is hard to do that by sending petabytes and petabytes of imaging data back here on Earth and doing that work. It is easier just to do it out in space. And then ignore all of the data collected and processed until you see something interesting. And so artificial intelligence in space will have very good, very interesting applications.

Toshiya Hari: Your next question comes from Mark Zapakos with Evercore ISI.

Mark Zapakos: Hi. Thanks for taking my question. I want to pick up on the comment you made in the script. I believe, Colette, you said that hyperscalers were over 50% of revenues, but growth was led by the rest of your Data Center customers. Just to clarify, does that imply your non-hyperscale customers grew faster? And if so, are they doing different things than the hyperscalers, or the same things on a different scale? Do you expect a point where non-hyperscalers become a larger part of your business?

Jensen Huang: Our top five, as we articulated as being our CSPs or hypersigners, about 50%. Types of companies that we are working with. And it goes through our AI model makers, providers on our platform, and now we also have an extreme diversity of customers that we are seeing all the way across the world. And this will really benefit seeing that diversity in the ecosystem.

Colette Kress: That is in every cloud available at the edge. And we are now cultivating telecommunications. Also be a computing platform. Is a foregone conclusion, but somebody has to go and invent the technologies to make that possible. We created the tech—we created a platform called Ariel to go do that. Just by every single self-driving car. Our ability—CUDA's ability—to have the benefit of the performance of specialized processors on the one hand, with the Tensor Cores inside our GPUs, on the other hand, the flexibility of CUDA allows us to solve language problems, computer vision problems, robotics problems, biology problems, physics problems, and just about all kinds of AI and all kinds of computation algorithms.

And so the diversity of our customer base is one of the greatest strengths that we have. The second thing, of course, is without our own ecosystem, even if our processor was programmable, if we did not cultivate our ecosystem—talking about some of the things we were doing today, investing in our future ecosystem and our ecosystem—without our ecosystem, it is hard for us to go beyond what design wins we capture for somebody else's ecosystem. And so we could grow and expand our ecosystem very—because of the platform that we created.

And then lastly, of the things that is really important is the partnerships that we have with OpenAI and Anthropic with xAI, with Meta, now makes and enforces test by every single open source in the world. There is one and a half million AI models on Hugging Face, all of it runs on NVIDIA Corporation CUDA. And so an open source in totality probably represents the largest—second largest—model in the world, OpenAI is the largest, second largest probably all the collection of all the open sources. So NVIDIA Corporation's ability to run all of that makes our platform super fungible, super easy to use, and really safe to invest into.

And so that creates the diversity of customers and a diversity of the platforms and available in every single country and because, you know, we support the whole world's ecosystem.

Toshiya Hari: Your next question comes from Aaron Rakers with Wells Fargo.

Aaron Rakers: Yes. Thanks for taking the question. Sticking with the platform and extreme codesign, there

has been news about NVIDIA Corporation's ability or push to bring Vera CPUs to market on a standalone solution basis. Jensen, what is the importance of your place in architecture evolution as we move forward? Is this being driven more by the proliferation or the heterogeneity of workloads? How do you see that evolving for NVIDIA Corporation, particularly on a standalone CPU?

Jensen Huang: Thank you. No. Thanks. I will tell you some more about it at GTC. But at the highest level, we made fundamentally different architecture decisions about our CPUs compared to rest of the world CPUs. But it is the only ports LPDDR5. It is designed to be focused on very high data processing capabilities. The reason for that is because most of the computing problems that we are interested in are data-driven. Artificial intelligence is one. And the single-threaded performance in this ratio with bandwidth is just off the charts. And we made those architectural decisions because in the entire phase, the different phases of AI, from data processing—before you even do training—you have to do data processing.

So you have data processing, pre-training, and then post-training. Now the AIs are learning how to use tools. The usage of tools, many of those tools, run in CPU-only environments or they run in CPU- or GPU-accelerated environments. And Vera was designed to be an excellent CPU for post-training. And so some of the use cases in the entire pipeline of artificial intelligence includes using a lot of CPUs. You know, we love CPUs as well as GPUs. And when you accelerate the algorithms to the limit, as we have, Amdahl's Law would suggest that you need really, really fast single-threaded CPUs. And that is the reason why we built Grace to be extraordinarily great at single-threaded performance.

And Vera is off the charts better than that.

Tim Arcuri: Thanks a lot. I was wondering if you can talk about the deployment of capital. I know that you really jacked up the purchase commits, but it sounds like maybe you are over the hump on this and you are going to probably generate about \$100,000,000,000 in cash this year. Pretty much no matter how good the results have been, the stock has not really gone up much. So I would think that you probably feel like this is a pretty good price to be buying back a bunch of it here. Why not put a big stake in the ground and just have a huge share repo here? Thanks.

Colette Kress: Thanks for the question. We look at our capital return very, very carefully. And we do believe that one of the most important things that we can do is really supporting the extreme oak ecosystem that is in front of us. That stems from everywhere, from our suppliers and the work that we need to do to assure that we can have the supply that is needed and help them from a capacity. All the way that we are in terms of the early developers of the AI solutions that will be on our platform. So we will continue to make this a very important part of our process. And strategic investments.

But of course, we are still repurchasing our stock. We are still with our dividend as well. And we will continue to find the right unique opportunities within the year for doing those different purchases.

Sarah: Your final question comes from Jim Schneider with Goldman Sachs.

Jim Schneider: Thank you for taking my question. Jensen, you have previously outlined the potential to get to \$3 to \$4 trillion of data center CapEx by 2030, which implies a potential acceleration in growth rates, which you have guided to at least this next quarter. What are some of the key application areas most likely to drive that inflection? Is that physical AI, agentic, or something else? And do you still feel good about that \$3 to \$4 trillion envelope?

Jensen Huang: Yeah. Let us back that up and just reason through it from a few different ways. So the first way is on first principles. The way that software is done in the future using AI is token-driven. And I think everybody talks about tokenomics and talks about data centers generating tokens, and inference is about generating tokens, and we generate tokens. You know, we are just talking about tokens. How NVIDIA Corporation's NVLink 72 enabled us to generate tokens at 50 times better performance per unit energy than the previous generation. And so token generation is at the center of almost everything that relates to software in the future and relates to computing.

If you look at the way we use computing in the past, however, the amount of computation demand for software in the past is a tiny fraction of what is necessary in the future. And AI is here. AI is not going to go back. AI is only going to get better from here. And so if you think about it, you said, okay, well, the world was investing about \$300,000,000,000 to \$400,000,000,000 a year in classical computing. And now AI is here. And the amount of computation necessary is a thousand times higher than the way we used to do computing. The computing demand is just a lot higher.

And so if we continue to believe there is value in it, and we will talk about that in a second, then the world will invest to produce that token. And so the amount of token generation capability that the world needs is a lot more than \$700,000,000,000. And I am fairly confident that we are going to continue to generate tokens. We are going to continue to invest in compute capacity from this point out. And fundamentally because every single company depends on software, every software will depend on AI. And so every company will produce tokens. And that is the reason why I call them AI factories.

And whether you are a company in the cloud data centers, you have AI factories to generate tokens for your revenues. If you are an enterprise software company, you are going to generate tokens for the systems that are on top of your tools. If you are a robotics factory—and self-driving cars, first indication of that—you have huge supercomputers, which are basically AI factories, to generate tokens that goes into your cars, that becomes its AI. And then you also have to put computers inside the cars to continuously generate tokens. And so we are fairly sure now that this is the future of computing.

Now why is it so certain that this is the future of computing? And the reason for that is because the way we used to do software was pre-recorded. Everything was captured a priori. We pre-compiled the software. We pre-write the content. We pre-record the videos. But now everything is generative in real time. And when it is generated in real time, it can take into context of the person, the situation, the query, and the intentions could all be taken into consideration to generate the outcome of this new software called—we call AI—agentic AI. And so the amount of computation necessary is far, far greater than pre-recorded.

You know, just as a computer has a lot more computation capability than a DVD player that was pre-recorded, artificial intelligence needs a lot more computing capability than the way we used to do software in the past.

Now, the question about computation, about sustainability at the first level, is just at the computer science level. This is the way computing is going to be done now. From an industrial level, because as all of our companies, in the final analysis, are powered by software and the cloud companies are powered by software. And if the new software requires tokens to be generated and the tokens are monetized, then it stands to reason that their data center build-out directly drives their revenues. And so compute drives revenues. And I think they all understand that. I think people are increasingly starting to understand that as well.

And then lastly, you know, the benefits that AI produces for the world ultimately has to generate revenues. And we are seeing—right in front—being developed as we stand here, agentic AI has turned an inflection point. And it literally happened in the last couple, two, three months. Of course, inside the industry, we have been seeing it for a while—you know, probably six months or so. But the world is now awakened to the agentic AI inflection. The agents are super smart. They are solving real problems.

Coding is obviously supported by agentic systems now, and all of our coders here at NVIDIA Corporation are using systems—either Claude Code or OpenAI Codex—enormously, and oftentimes both, and Cursor, oftentimes all three, depends on the use case. But they have agents and co-designed partners, engineering partners, to help them solve problems. And you could see their revenue skyrocketing. You know, these companies, in the case of Anthropic, I think their revenues ten in a year. And they are severely capacity constrained because demand is just incredible. And the token demand is incredible. The token generation rate is growing exponentially. And the same thing with, of course, OpenAI. Their demand is incredible.

And so the more compute that they can stand online, bring online, the faster their revenues will grow. And that goes back to the comment that I was saying, that inference is revenues, that compute equals revenues. Now, in this new world. And in a lot of ways, that is the reason why we say it is a new industrial revolution. There are new factories, new infrastructure being built, and this new way of doing computing is not going to go back.

And so to the extent that we believe that producing tokens is going to be the future of computing, which I believe, and I think largely the industry believes, then we are going to be building out this capacity from this point forward and continue to expand from here.

Now, the thing that is the wave that we are seeing now is the agentic AI inflection, and the next inflection beyond that is physical AI, where we take AI and these agentic systems into the physical applications such as manufacturing, such as robotics. And so that is a giant opportunity ahead.

Toshiya Hari: Okay. This concludes the question-and-answer session. In closing, please note Jensen will be participating in a fireside chat at the Morgan Stanley TMT Conference in San Francisco on March 4. He will also be giving a keynote at GTC in San Jose on March 16. Our earnings call to discuss the results of our first

quarter of fiscal 2027 is scheduled for May 20. Thank you for joining us today. Operator, please go ahead and close the call.

Sarah: Thank you. This concludes today's conference call. You may now disconnect.

Note: In the digital version, underlined words indicate emotional associations

# Emotional Landscape

Emotions mapped by valence (positive/negative) and arousal (energy level)

<b>Gratitude</b>	74.7% probability
<b>Trust</b>	41.5% probability
<b>Hope</b>	19.4% probability
<b>Satisfaction</b>	14.0% probability
<b>Admiration</b>	7.9% probability
<b>Confusion</b>	7.0% probability
<b>Anticipation</b>	6.7% probability
<b>Surprise</b>	3.4% probability

# Emotional Arc

How emotions change throughout the document (6566 words)

The document shows varying emotional patterns across sections, with primary emotions shifting as the content progresses.

# Supporting Evidence

Detailed breakdown of what contributed to each emotion

## Gratitude - 74.7%

### Key Terms:

- "thanks" (appears 11x, NaN% relevance)
- "pleasure" (appears 1x, NaN% relevance)
- "partnership" (appears 3x, NaN% relevance)
- "love" (appears 3x, NaN% relevance)

### Example Quotes:

*"Thanks, Toshiya."*

Key terms: Thanks

"We are thrilled with our ongoing partnership with OpenAI, a once-in-a-generation company we have had the pleasure of partnering with since their first days."

Key terms: partnership, pleasure, partnering

"Our engineers love it."

Key terms: engineers, love

### Related Themes:

- Appreciation for Collaboration &&&&&
- Politeness &&&&&
- Acknowledgement of Contribution &&&&&

Coverage: 18 words (5.0% of text)

## Trust - 41.5%

### Key Terms:

- "confidence" (appears 1x, NaN% relevance)
- "commitments" (appears 4x, NaN% relevance)
- "ability" (appears 11x, NaN% relevance)
- "growth" (appears 12x, NaN% relevance)
- "leadership" (appears 3x, NaN% relevance)

### Example Quotes:

*"I am confident in their cash flow growing."*

Key terms: cash, flow, growing

"we remain confident in our ability to capitalize on the growth opportunity ahead with our scale, expansive supply chain, and long-standing partnerships continuing to serve us well."

Key terms: ability, growth, opportunity, partnerships

"We believe we have inventory and supply commitments in place to address future demand, including shipments extending into calendar 2027."

Key terms: inventory, supply, commitments, demand, shipments

### Related Themes:

- Financial Confidence &&&&
- Future Growth Belief &&&&
- Supply Chain Reliability &&&&
- Commitment to Shareholders &&&&

Coverage: 31 words (30.0% of text)

## Hope - 19.4%

### Key Terms:

- "expect" (appears 9x, NaN% relevance)
- "growth" (appears 12x, NaN% relevance)
- "opportunity" (appears 5x, NaN% relevance)
- "future" (appears 9x, NaN% relevance)
- "progress" (appears 2x, NaN% relevance)

### Example Quotes:

*"We look ahead. We expect sequential revenue growth throughout calendar 2026, exceeding what was included in the \$500,000,000 Blackwell and Rubin revenue opportunity we shared last year."*

Key terms: growth, revenue, opportunity

"Over the long run, we expect our sovereign opportunity to grow at least in line with the AI infrastructure market, as countries spend on AI proportional to their GDP."

Key terms: run, opportunity, grow, market

"that is a giant opportunity ahead."

Key terms: opportunity

### Related Themes:

- Future Revenue Growth &&&&
- Market Expansion &&&&
- Technological Advancement &&&&
- Long-term Opportunity &&&&

Coverage: 37 words (60.0% of text)

## Satisfaction - 14.0%

### Key Terms:

- "kudos" (appears 1x, NaN% relevance)
- "record" (appears 4x, NaN% relevance)
- "growth" (appears 12x, NaN% relevance)
- "innovation" (appears 2x, NaN% relevance)
- "leadership" (appears 3x, NaN% relevance)

### Example Quotes:

"We delivered another outstanding quarter, with record revenue, operating income, and free cash flow."

Key terms: quarter, record, revenue, income, cash, flow

"First, let me say kudos on including this in non-GAAP. I think that is great."

Key terms: kudos

"Our pace of innovation, particularly at our scale, is unmatched."

Key terms: pace, innovation, scale

### Related Themes:

- Financial Performance &&&&
- Record Breaking Results &&&&
- Market Leadership &&&&
- Praise from Analysts &&&&

Coverage: 22 words (40.0% of text)

## Admiration - 7.9%

### Key Terms:

- "king" (appears 1x, NaN% relevance)
- "innovation" (appears 2x, NaN% relevance)
- "celebrated" (appears 1x, NaN% relevance)
- "launch" (appears 1x, NaN% relevance)
- "partnership" (appears 3x, NaN% relevance)

### Example Quotes:

"We are thrilled with our ongoing partnership with OpenAI, a once-in-a-generation company we have had the pleasure of partnering with since their first days."

Key terms: partnership, company, pleasure, partnering

"Anthropic's Claude Cwork agent platform is revolutionary and has opened the floodgates for enterprise AI adoption."

Key terms: agent, platform, floodgates, adoption

"SemiAnalysis declared NVIDIA Corporation inference king as recent results from InferenceX reinforced our inference leadership, with GB300 and NVL72 achieving up to 50x performance per watt and 35x lower cost per token, compared with offerings."

Key terms: inference, king, results, leadership

### Related Themes:

- Praise for Partners &&&&
- Technological Superiority &&&&
- Industry Recognition &&&&
- Praise for Talent &&&&

Coverage: 8 words (20.0% of text)