



NVIDIA Non-Deal Roadshow

July, 2026




Certain matters in this presentation including, but not limited to, statements as to: expectations with respect to growth, demand, performance, benefits and availability of NVIDIA's products, services, and technologies, including Blackwell and Rubin, and related trends and drivers; NVIDIA's financial position and financial and business outlook; NVIDIA's future cash dividends and other returns to shareholders; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; third parties adopting our products and technologies, expectations with respect to technology developments, and related trends and drivers; projected market growth, opportunities 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 management's beliefs and assumptions and on information currently available to management and are subject to risks and uncertainties that could cause results to be materially different than expectations.

Important factors that could cause actual results to differ materially include: global economic and political conditions; NVIDIA's reliance on third parties to manufacture, assemble, package and test NVIDIA's products; the impact of technological development and competition; development of new products and technologies or enhancements to NVIDIA's existing products and technologies; market acceptance of NVIDIA's products or NVIDIA's partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of NVIDIA's products or technologies when integrated into systems; NVIDIA's 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 NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its Annual Report on Form 10-K and Quarterly Reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

Many of the products and features described herein remain in various stages and will be offered on a when-and-if-available basis. The statements within are not intended to be, and should not be interpreted as a commitment, promise, or legal obligation, and the development, release, and timing of any features or functionalities described for our products is subject to change and remains at the sole discretion of NVIDIA. NVIDIA will have no liability for failure to deliver or delay in the delivery of any of the products, features or functions set forth herein.

NVIDIA uses certain non-GAAP measures in this presentation including free cash flow. NVIDIA believes the presentation of its non-GAAP financial measures enhances investors' overall understanding of the company's historical financial performance. The presentation of the company's non-GAAP financial measures is not meant to be considered in isolation or as a substitute for the company's financial results prepared in accordance with GAAP, and the company's non-GAAP measures may be different from non-GAAP measures used by other companies. Further information relevant to the interpretation of non-GAAP financial measures, and reconciliations of these non-GAAP financial measures to the most comparable GAAP measures, may be found in the slide titled "Reconciliation of Non-GAAP to GAAP Financial Measures".



NVIDIA

NVIDIA's invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined computer graphics, revolutionized accelerated computing, and ignited the era of modern AI.

Today, NVIDIA is a foundational company, driving one of the most important computing platform shifts in decades as it builds the intelligence infrastructure layer for the global AI economy.


Every major technology era creates a new foundational layer. Electricity required power grids. The internet required networks, data centers, and cloud platforms. AI requires a new kind of infrastructure: AI factories that convert data and energy into intelligence – monetizable tokens.

As software, systems, robots, factories, vehicles, and enterprises become AI-powered, they will need infrastructure that is fast, trusted, scalable, and widely adopted. NVIDIA is the full-stack infrastructure platform for the age of AI: the company that turns chips, systems, software, developers, and data centers into AI factories.

Headquarters: Santa Clara, CA | Employees: 42,000+

NVIDIA is the Platform Layer of the AI Economy


NVIDIA started as the chip that made AI possible. Today, it is becoming the builder of AI infrastructure: a full-stack platform — a complete infrastructure spanning GPUs, CPUs, networking, systems, software, developer tools, and cloud-scale deployment.




Other Companies Sell Chips — NVIDIA Builds Full-Stack AI Factories

This is what makes NVIDIA different. Other chip companies sell components. NVIDIA builds complete gigawatt-scale AI factory systems: integrated, yet customizable, infrastructure that can be deployed across every major cloud, inside enterprise data centers, at the edge, and into robotics and physical AI systems.


Vera Rubin NVL72
Compute Tray




Vera Rubin NVL72
NVLink Switch Tray




Vera CPU Tray




Spectrum-6 SPX
Switch Tray



Groq 3 LPX Tray



Vera BlueField-4
STX Storage Tray



NVIDIA Wins Through Extreme Co-Design

NVIDIA wins through extreme co-design. It owns the critical layers of the AI factory — CPU, GPU, LPU, DPU, networking fabric, systems, software, libraries, model execution, and full system architecture — and engineers them into one integrated platform. This advantage compounds because NVIDIA is world-class across every major chip and system layer. Combined with a relentless one-year roadmap cadence, NVIDIA keeps raising the performance bar in a way a single-chip approach cannot sustain. That is why NVIDIA continues to lead across training, inference, and agent workloads: it optimizes the full AI factory, not just the GPU.




NVIDIA Creates Value for Every Part of the AI Ecosystem

NVIDIA creates value for every part of the AI ecosystem.


Developers get a high-performance platform with a global reach, enabling them to build once and deploy everywhere. CSPs get the AI platform their customers already demand, with the performance needed to drive utilization, revenue, and profitable growth. AI clouds get bankability: a trusted platform that helps them raise capital, win customers, and scale quickly. OEMs get a complete AI factory platform with global enterprise demand. Enterprises get control: a trusted AI factory they can own, secure, and operate without building the entire stack from scratch.

That is NVIDIA's strategic advantage. It is not just a chip supplier. It is the platform that connects developers, clouds, infrastructure providers, OEMs, enterprises, and end users into one AI ecosystem.




Developers Want Performance and Reach That NVIDIA CUDA Provides


NVIDIA CUDA provides the performance and market reach developers want. CUDA-X libraries accelerate apps, enabling them to run orders of magnitude faster, more efficiently, and at lower cost than on CPUs. Over time the introduction of new libraries expands our platform's addressable set of applications and use cases, unlocking new markets and increasing our long-term growth opportunity. With 7,000+ apps, AI models, and libraries running on CUDA, NVIDIA addresses the largest market opportunity of any computing platform.




cuLitho
Computational
Lithography




cuDSS
cuSPARSE
cuFFT
AmgX
Computer Aided
Engineering




cuOpt
Decision
Optimization




Warp
Physics




cuDF
cuML
Data Science
and Processing



Megatron
Dynamo
NIXL
cuDNN
CUTLASS
Deep
Learning




cuEquivariance
cuTensor
Quantum
Chemistry



cuQuantum
CUDA-Q
Quantum
Computing




Earth-2
Weather
Analytics



MONAI
Medical
Imaging



Parabricks
Genomics




**Aerial
Sionna**
5G/6G
Signal Processing



cuPyNumeric
Numerical
Computing

Clouds Want Profitable Growth

AI factory lifetime revenue is the area under the curve. NVIDIA's scale, experience, and ability to execute extreme co-design enable earlier TTFT, better reliability, and higher tokens/watt. The CUDA ecosystem and installed base extend useful life.



Clouds Want Bankability

Reasons for NVIDIA GPU pricing durability

- Large and growing installed base across cloud, on-prem, edge
- Fungibility across AI/non-AI, training/inference
- CUDA software drives continuous performance enhancements
- Strong track record and existence of established market


A100 AI Cloud GPU pricing

- Trailing 3 months: +7%
- Trailing 6 months: +5%

H100 AI Cloud GPU pricing

- Trailing 3 months: +13%
- Trailing 6 months: +21%

AI Cloud GPU Pricing



Source: SemiAnalysis

Extreme Co-Design Repeatedly Delivers Performance Leadership in Training AND Inference

Training – MLPerf Training v6.0

- Won every benchmark and the only platform to submit on all tests
- Blackwell Ultra GB300 up to 1.6X faster than GB200


Inference – MLPerf Inference v6.0

- Highest-performing submissions across broadest range of models
- Only platform on every new benchmark
- GB300 NVL72 throughput up to 2.7X in a single round through software alone — slashing cost per million tokens
- Record reasoning throughput: 2.5M tokens/sec (DeepSeek-R1, offline)

Inference – SemiAnalysis InferenceX

- GB300 NVL72 delivers up to...
 - 50X higher throughput per megawatt
 - 35X lower cost per token vs. Hopper (agentic/low-latency workloads)
 - 1.5X lower cost per token vs. GB200 on long-context


MLPerf Training and Inference Wins




Extreme Co-Design Turns Performance Into Customer Economics

“GB NVL72 Inference King”

Tokens per Watt Drives Factory Revenue



Performance Drives Token Cost



DeepSeek R1 0528 · FP4 · 1K/1K · Source: SemiAnalysis InferenceX
Performance per watt determines the inference economics of a power-limited data center.
Through extreme co-design NVIDIA provides data center operators with the best inference performance and economics – the best ROI, token cost, and margins.



Agents Are the New AI Workload — NVIDIA Leads There Too

Artificial Analysis AgentPerf

- GB300 NVL72 runs up to 20X more agents per MW than H200
- Measured on DeepSeek V4 Pro (frontier MoE) using real coding-agent trajectories across 12+ languages
- Already in production with Baseten, DeepInfra and Together AI (e.g., Cursor on Blackwell)


CPU – Vera CPU, Phoronix Benchmarks

- 1.5X overall performance vs. latest-gen 128-core x86; 1.6X vs. prior-gen Grace
- Sustained 90% of peak memory bandwidth (highest of any CPU tested) and 4X memory bandwidth/core vs. x86

Models – Nemotron 3 Ultra, Artificial Analysis


- Lands in the most attractive quadrant on Artificial Analysis
- Frontier intelligence at ~5X faster output speed and ~30% lower cost to task completion vs. the best open models

NVIDIA GB300 NVL72 Delivers 20X Better Agentic Coding Performance



NVIDIA Ecosystem Flywheel Has Reached Escape Velocity

NVIDIA has reached a tipping point: a reinforcing flywheel where more developers build on NVIDIA, more models are optimized for NVIDIA, more workloads run on NVIDIA, more clouds deploy NVIDIA, and more enterprises choose NVIDIA. That demand, supported by NVIDIA's substantial free cash flow, gives the company the confidence to secure the supply required to build AI infrastructure at multi-gigawatt scale. At the same time, high utilization, broad customer demand, and longer infrastructure life make NVIDIA-based AI factories more profitable — attracting more capital and reinforcing demand. The result is a platform customers want, developers support, clouds deploy, and investors are willing to fund.




Useful AI Has Arrived


GitHub shows software coding surging in early 2026: 1.4B commits (nearly 3X), 90M pull requests merged, 20M new repos/month.

~30–40M software developers (~\$3T in annual salaries) are now producing ~3X the output — effectively ~\$9T of productivity — and more engineers are being hired, not fewer.

Useful AI means tokens are now profitable — and this incentivizes data center operators to expand compute.




Token Generation is Growing Exponentially



Source: OpenRouter

NVIDIA Powers Global AI Economy

Physical AI extends AI into the physical economy — manufacturing, robotics, autonomous vehicles and heavy industry. Addresses a \$50 trillion market and represents a multi-trillion-dollar long-term opportunity for NVIDIA. Only NVIDIA delivers all three computers necessary to build physical AI — Train (DGX), Simulate (Omniverse + Cosmos), and Run (Jetson Thor).



Omniverse

Libraries and microservices for developing Physical AI applications

Cosmos 3

Open world foundation model

Isaac GROOT

Open humanoid-robot reference platform

Alpamayo

The world's first reasoning autonomous vehicle model


NVIDIA is No Longer Just a GPU Story

Our demand is multi-industry and diversifying — AI Clouds, Industrial and Enterprise (ACIE) account for ~50% of Data Center revenue. ACIE — where NVIDIA is uniquely positioned to offer full-stack AI infrastructure — is expected to outgrow Hyperscale over the long run. We have visibility to more than \$1 trillion in cumulative Blackwell and Rubin revenue from CY2025 through CY2027. Growth opportunities spanning standalone CPUs, LPUs and storage are incremental to the \$1 trillion+.

	FY 2024	FY 2025	FY 2026	2027 (Q1)
Datacenter	\$48	\$115	\$194	\$75
Hyperscale	25	54	106	38
AI Clouds, Industrial Enterprise	23	61	88	37
Edge Computing	13	15	22	6
TOTAL	\$61	\$130	\$216	\$82


(\$ in Billions)

Blackwell + Rubin Opportunity




NVIDIA's Demand is Increasingly Diverse

Customer Diversity Enabled by Full-Stack AI Factory




Customer Diversity Enabled by CUDA Full-Stack Platform




Vera CPU — A New Pillar of Growth



1.8X Agentic Sandbox Performance vs. x86



Plan to Return 50% or More of Free Cash Flow



Cash Flow Strength and Capital Return

Cash Flow Strength

- Record operating cash flow, free cash flow, and cash balances — scaling rapidly, funding our strategic investments and capital return
- Operating cash flow increased from \$64B in FY25 to \$103B in FY26, and reached \$50B in Q1 FY27 alone
- Free cash flow grew from \$61B in FY25 to \$97B in FY26, and increased to \$49B in Q1 FY27 alone
- Cash and marketable securities stood at \$81B at the end of Q1 FY27, up from \$54B a year ago

Capital Return

- **Share Repurchases:** Announced a new \$80B share repurchase authorization, on top of ~\$39B remaining as of Q1 FY27
 - Returned a record ~\$20B to our shareholders in Q1 FY27, while executing strategic investments
- **Dividend Increase:** Increased quarterly dividend from \$0.01 to \$0.25 per share
- Plan to return 50% or more of free cash flow to shareholders — this year, next year, and beyond
- Combination of regular, modest dividend increases and expanding stock repurchases

Key Takeaways

- NVIDIA has evolved from a GPU to data center scale AI infrastructure provider
- Through extreme co-design, NVIDIA AI infrastructure delivers the highest token throughput and lowest token cost
- Architectural decisions are critical given ongoing power constraints – this favors NVIDIA
- AI represents the largest computing platform shift ever
- AI is now useful and tokens are profitable – driving unprecedented demand for AI infrastructure
- See more than \$1 trillion in cumulative Blackwell and Rubin revenue from CY2025 through CY2027
- Our demand is multi-industry and diversifying — non-Hyperscale is ~50% of Data Center revenue
- Physical AI is just around the corner — NVIDIA is uniquely positioned to offer the full stack
- Investing in R&D and our ecosystem is a top priority
- Plan to return 50% or more of FCF to shareholders — combination of regular/modest dividend increases and expanding stock repurchases

Reconciliation of Non-GAAP to GAAP Financial Measures

	Free Cash Flow	Purchases Related to Property and Equipment and Intangible Assets	Principal Payments on Property and Equipment and Intangible Assets	Net Cash Provided by Operating Activities
Q1FY27	\$48,554	(1,757)	(33)	\$50,344
FY26	\$96,575	(6,042)	(101)	\$102,718
FY25	\$60,724	(3,236)	(129)	\$64,089
FY24	\$26,947	(1,069)	(74)	\$28,090

(\$ in Millions)

