NVIDIA Corporation

April 10, 2026

Disclaimer

The analyses and conclusions contained in this presentation are based on publicly available information. There may be confidential or otherwise nonpublic information in the possession of the companies discussed in this presentation that could lead these companies and others to disagree with these analyses, conclusions and opinions.

This presentation may include forward-looking statements, estimates, and projections which reflect various assumptions that may not be accurate.

The content expresses the views of the author as of the time of writing and such views are subject to change. This presentation and the information contained herein is not investment advice or a recommendation or solicitation to buy or sell any securities. Past performance is not indicative of future results. All investments involve risk, including loss of principal.

Overview

Ticker: NVDA

Stock Price:

$188.63

- Accelerated-computing company designing GPUs, CPUs, networking, systems, and software for AI, data-center computing, graphics, simulation, and autonomous systems
  - Founded 1993; IPO 1999
- FY2026 revenue: $215.9B (+65.5% YoY)
- Market cap: ~$4.58 trillion
- CEO: Jensen Huang (founder; ~33 years as CEO)

3 Year Stock Performance (since Apr. 2023)

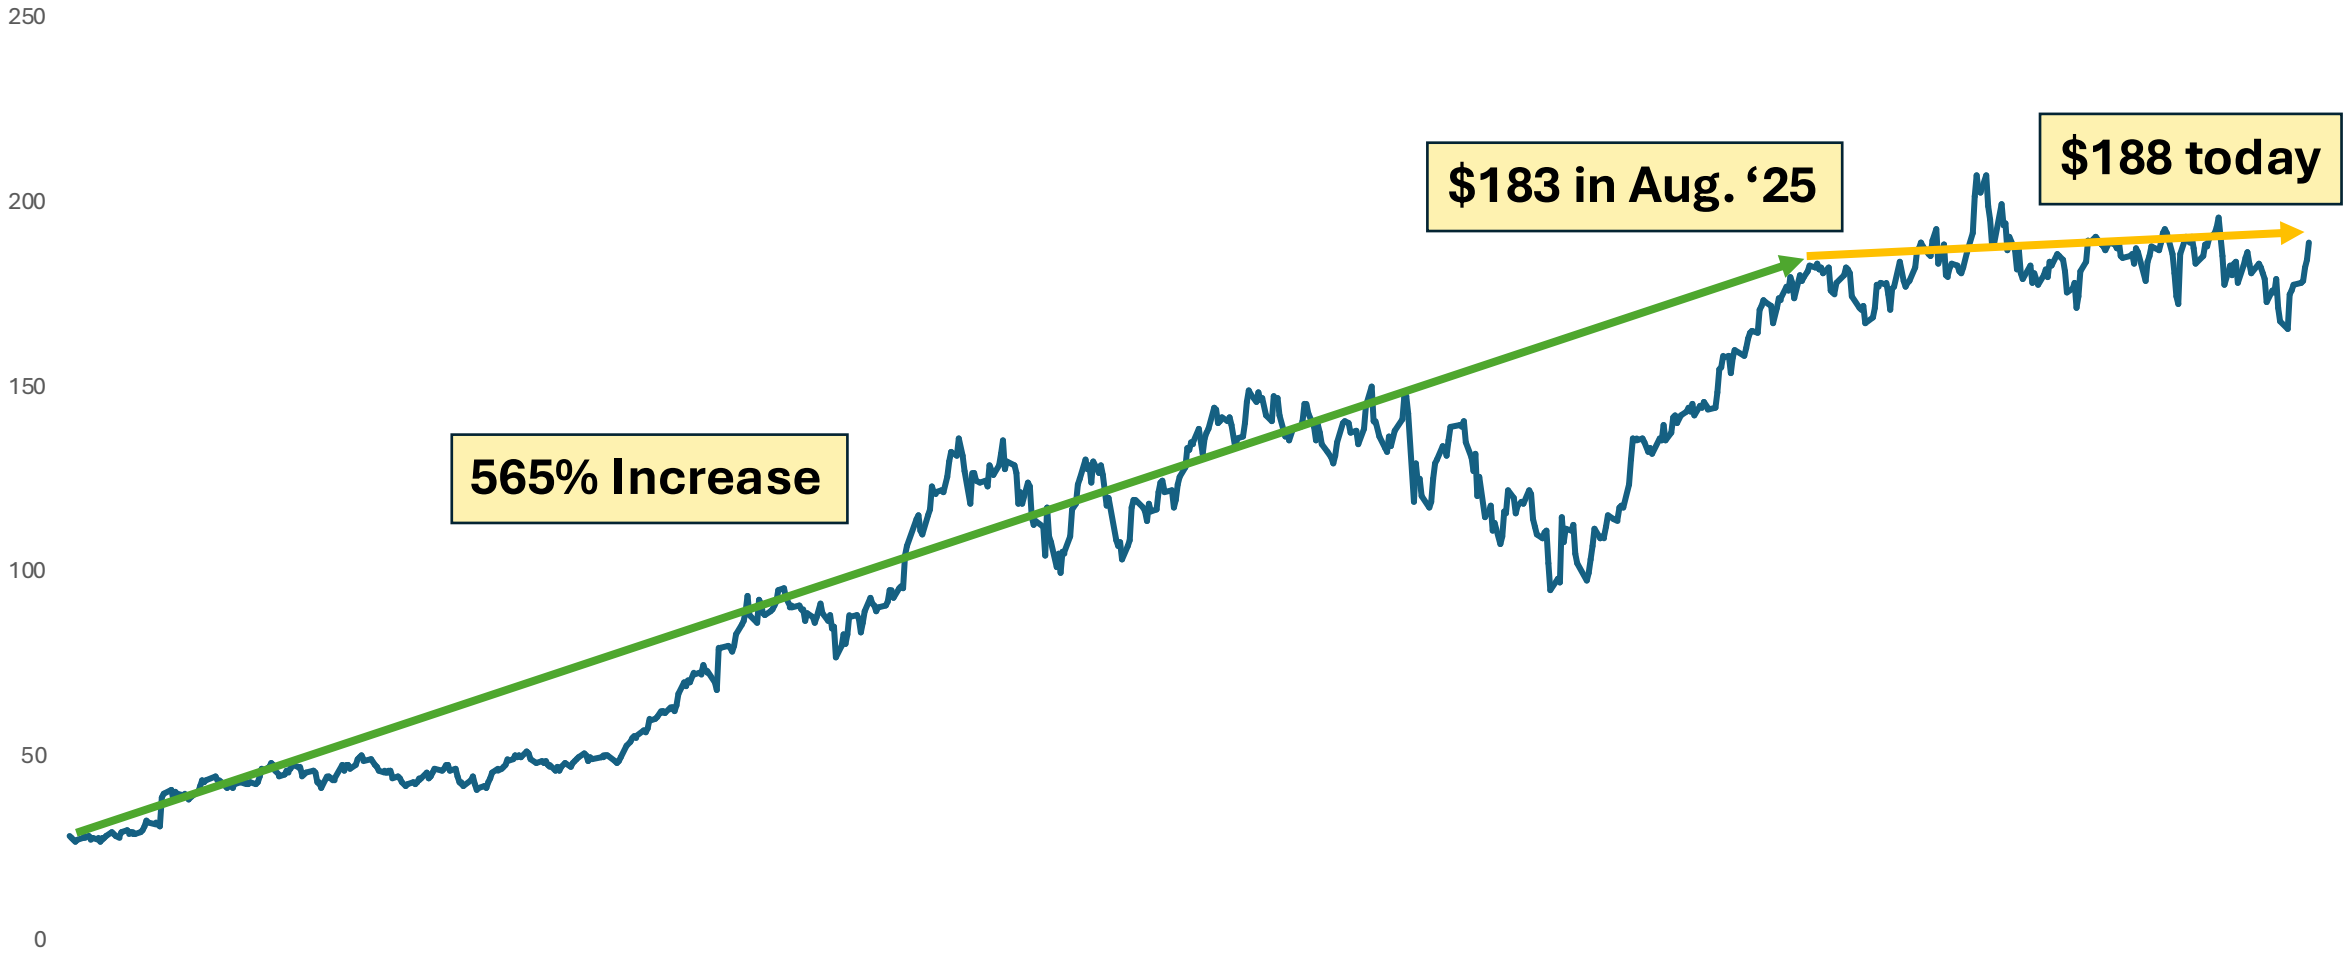

Revenue by Segment

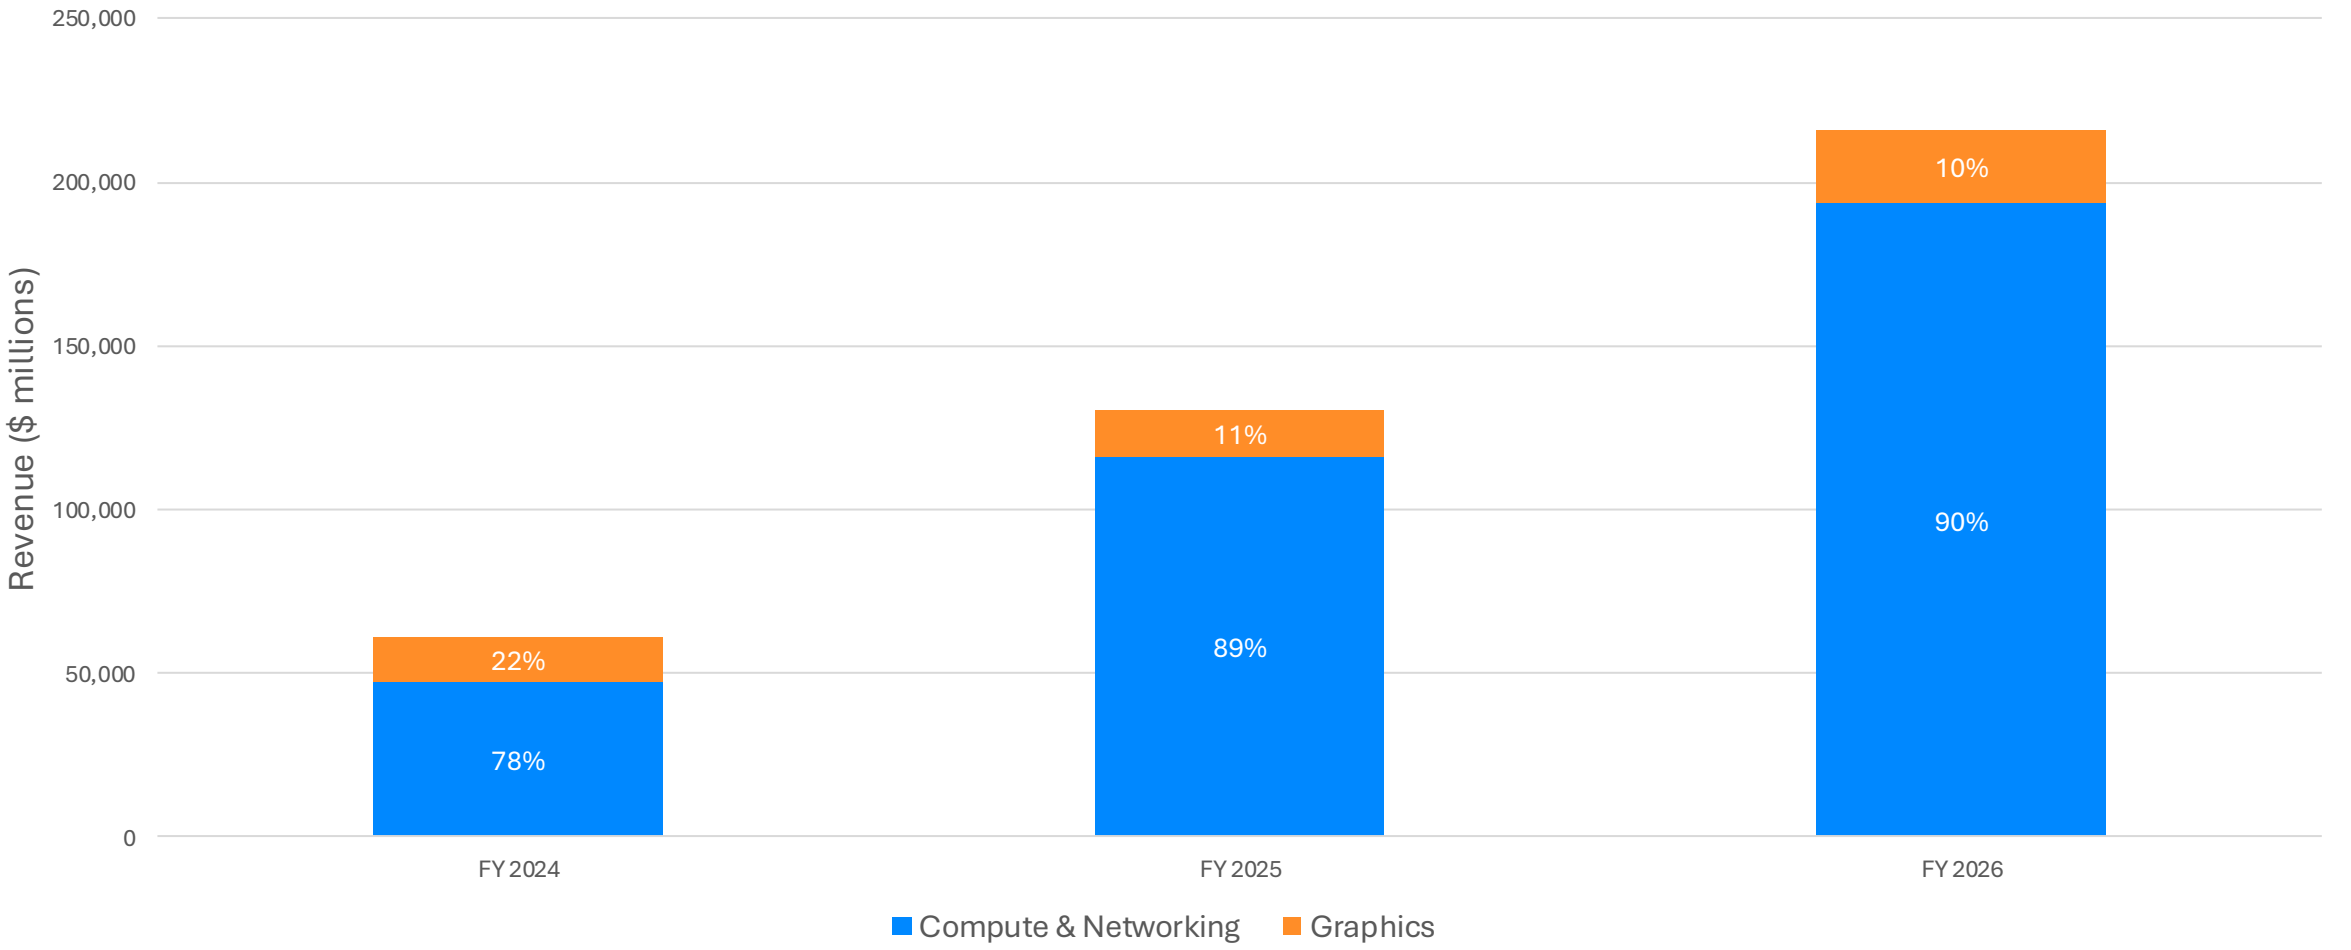

Revenue based upon Customer Headquarters Location

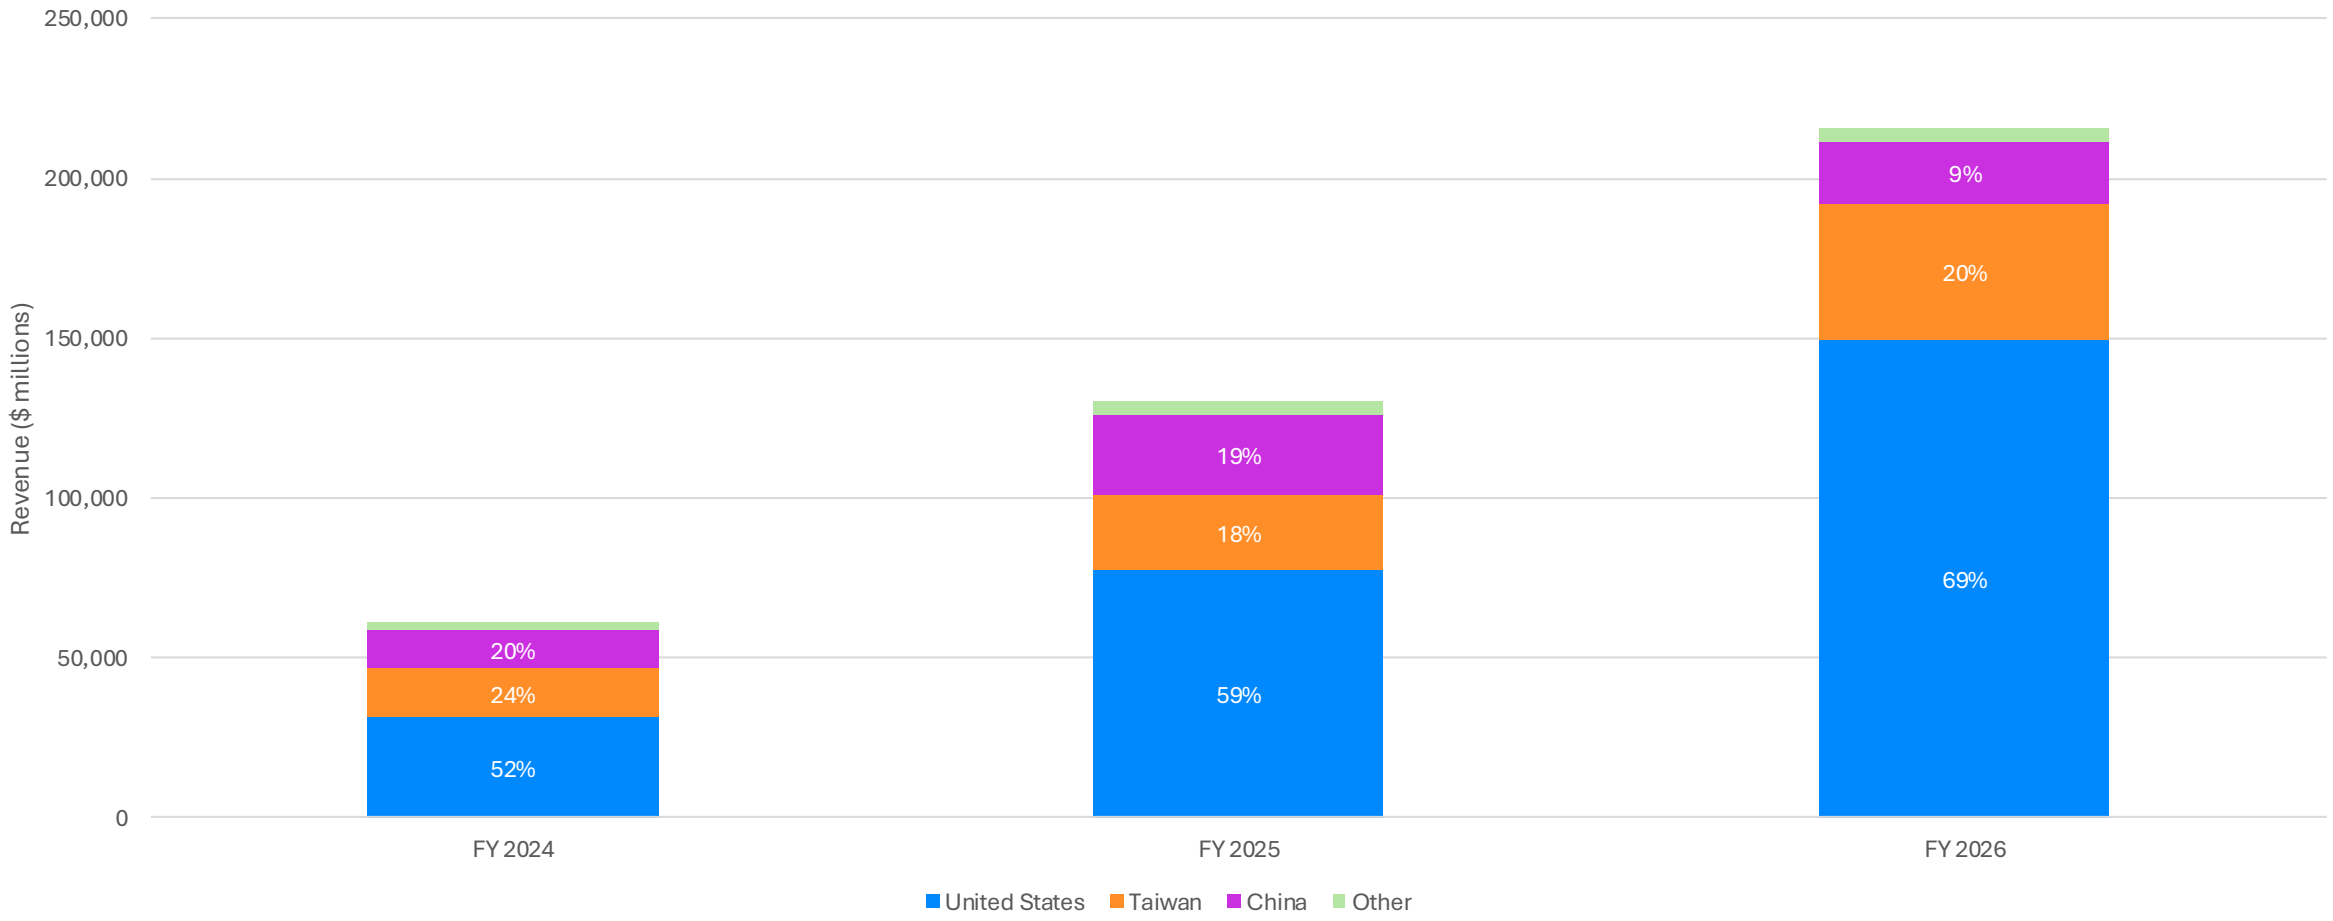

In fiscal year 2026, Nvidia estimates 76% of Data Center revenue from Taiwan-headquartered customers was attributed to end customers based in the United States and Europe - NVIDIA Corp., Form 10-K (FY2026)

Industry Overview

- Accelerated Computing & AI Infrastructure Market
  - Gartner forecasts $2.52T worldwide AI spending in 2026
  - WSTS projects $975B global semiconductor market in 2026
  - Dell'Oro forecasts ~25% CAGR for the data-center accelerator market over the next five years
  - Secular trend: build-out of AI-factory infrastructure pulling leading-edge logic, HBM memory, and advanced packaging
- Core Segment: Data-Center AI Accelerators
  - Merchant competitors: AMD, Intel, Huawei (in China/export-restricted contexts)
  - Custom-silicon & networking: Broadcom plus hyperscaler in-house chips
- Adjacent Markets
  - Networking, enterprise AI software, robotics, healthcare, simulation, autonomous vehicles, gaming

Semiconductor Market Dynamics

- Driver 1: Hyperscaler AI capex at unprecedented scale
  - Alphabet guided to $175–185B 2026 capex; Amazon $128.3B in 2025 (rising in 2026); Meta $72.2B in 2025; Microsoft quarterly capex reached $34.9B
  - OpenAI letter of intent for first gigawatt of Vera Rubin systems targets 2H 2026
- Driver 2: Constrained but expanding leading-edge supply
  - TSMC raised 2026 capex to $52–56B and is stepping up N3; SK hynix says AI memory demand exceeds capacity; Micron reports tight industry supply
  - New TSMC 3nm fabs (Tainan 1H27, Arizona 2H27, Japan 2028) bring incremental capacity but won’t normalize the market near-term
- Driver 3: Export controls and policy fragmentation
  - U.S. licensing requirement on H20 (Apr 2025) effectively foreclosed Nvidia from China data-center compute by FY26 year-end; $4.5B Q1 FY26 charge
  - Q4 FY26 forward guidance assumes no data-center compute revenue from China

Core Platform: Accelerated Computing Stack

- NVIDIA sells more than chips. The platform combines the GPU, the software that runs on it, the networking that connects thousands of GPUs together, and the full rack-scale systems that customers actually plug in
- The current GPU generation, Blackwell, is shipping at scale. The next generation, Vera Rubin, is on track for volume production in the second half of 2026. Each new generation has commanded a meaningful price premium because customers get more useful work per dollar and per watt
- The software layer, built around CUDA and a long list of libraries and tools, is what keeps customers on the platform. The networking business, which came from the Mellanox acquisition, is now $31.4B of FY26 revenue alongside $162.4B of compute, and matters more as AI workloads scale across many machines
- Around this core, NVIDIA also sells CPUs (Grace), pre-built systems (DGX), and software for managing AI workloads, plus platforms for cars, robotics, and simulation that broaden the addressable market beyond the data center

Competitive Standing

- NVIDIA's most direct competitor is AMD, which has built an open software stack and partnered with Meta on a very large AI deployment. Intel positions its Gaudi accelerators as a lower-cost alternative for customers who want more flexibility. Huawei is the main option for buyers inside China, where U.S. export controls limit what NVIDIA can sell.
- The bigger long-term threat comes from custom chips. Broadcom designs purpose-built AI chips and networking gear in close partnership with the largest cloud customers, and several of those customers are also building their own chips in-house to reduce reliance on NVIDIA.
- In practice, NVIDIA still ships the overwhelming majority of AI chips sold on the open market today
    - While they face pressure from custom chips, it is concentrated at a handful of the very largest buyers, and management argues NVIDIA's lead in software, networking, and full systems keeps customers coming back even when alternatives look cheaper on paper.

Competitive Advantage

- **NVIDIA sells the whole factory, not just the chip.** Customers don't buy a GPU on its own — they buy a system that includes the chip, the software it runs on, the networking that links thousands of chips together, and the racks that house all of it. A competitor with a cheaper chip still has to replicate the rest, which most cannot.
- **The software lock-in is the real moat.** NVIDIA's software platform, CUDA, has been the standard for AI development for over a decade. Engineers learn it, build on it, and the libraries get better every year. Switching to a competitor means rewriting code and retraining people, so customers stay even when alternatives appear cheaper on the spec sheet.
- **The right way to compare is cost-per-result, not cost-per-chip.** Management's framing is that what matters is how much useful AI work a system produces for each dollar and each watt of power. By that measure, each new NVIDIA generation extends the lead, which is why customers keep paying premium prices for newer hardware.
- **NVIDIA has also locked up the supply chain.** The company has placed long-dated orders with foundries, memory makers, and packaging partners that competitors would struggle to displace, so even a strong rival product can't easily get manufactured at the same scale.

Upcoming Catalysts

- Catalyst 1: Q1 FY27 earnings (May 20, 2026)
  - Guided $78.0B revenue ±2% with no China DC compute revenue assumed; Q4 FY26 GAAP GM recovered to 75.0%
- Catalyst 2: Hyperscaler AI capex commitments
  - GOOGL $175–185B 2026 capex; AMZN $128B+ 2025 (rising in 2026); META $72B 2025; MSFT $34.9B in a single quarter
  - OpenAI Sept 2025 LOI for first gigawatt of Vera Rubin in 2H 2026
- Catalyst 3: Rubin product cycle and GTC reset
  - Vera Rubin volume in 2H 2026; Rubin Ultra (CQ3 ramp); Feynman roadmap (NVLink LPU, Rosa CPU, NVLink8 CPO)
  - GTC analyst day (Mar 17–18, 2026): $1T+ Blackwell+Rubin demand through CY27 disclosed as a floor; ~50% FCF capital-return target formalized
- Catalyst 4: Adjacent revenue streams not yet in numbers
  - Groq LPU racks, standalone Vera CPU, BF4 STX storage, Rubin Ultra; expanded Uber DRIVE for autonomous ride-hailing across 28 cities from 2027

Why it trades at a deep discount: the dumb reasons

Investors don't want to “ride beta”

- There's smart beta and dumb beta
- The goal is to intelligently bear risk for profit, not appear pseudo-intellectual

Investors believe no alpha in mega caps

- If you can't find alpha in mega-caps, you don't have the skills to be in this business

Fears of an AI bubble top

- *“When change is this big … you think you're in the eighth inning and you're in like the third inning.”* — Stanley Druckenmiller

The valid concerns: custom ASICs and the shift to inference

- For frontier training, GPUs and TPUs beat custom ASICs because training changes fast
  - Model architectures, precision formats, memory needs, and distributed-training methods evolve quickly; GPUs are flexible, programmable, and have the best software ecosystem
  - TPUs are less general than GPUs but still designed for broad ML training at huge scale
- Inference scale tilts toward more specialized hardware over time
  - GPUs may lose share in % terms as inference diversifies, but absolute demand keeps growing as total AI usage expands
- Sell-side GTC checks:
  - Customers report trying cheaper alternatives, then returning to NVIDIA; market is supply-gated, so competitor share gains reflect availability, not preference (Morgan Stanley)

NVIDIA's Adjustments

1. Optimizing GPUs specifically for inference[1]
   - Blackwell added a second-generation Transformer Engine and lower-precision formats such as FP4, which are designed to improve both training and inference efficiency for LLMs and mixture-of-experts models
2. Moving from chips to rack-scale systems[2]
   - Blackwell and Rubin systems connect dozens of GPUs with high-speed NVLink so they behave more like one large accelerator
     - The Vera Rubin NVL72 rack, for example, combines 72 Rubin GPUs, 36 Vera CPUs, ConnectX networking, and BlueField DPUs into a tightly integrated system
   - Jensen Huang has repeatedly stressed the importance of the "American tech stack" when it comes to competing with China as a country
     - He has taken the same approach to Nvidia specifically.
   - This is important for inference because modern LLM serving is often memory- and communication-constrained, not just compute-constrained
     - Large context windows, mixture-of-experts routing, and agentic workloads need fast movement of data between GPUs, CPUs, memory, and network fabric

1 https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/
2 https://www.nvidia.com/en-us/data-center/technologies/rubin/

NVIDIA’s Adjustments

3. Software + Hardware[1]
   - NVIDIA’s CUDA stack and libraries are deeply embedded in AI workflows
     - High switching costs, most talent is already experienced with CUDA but not necessarily with alternatives
   - Dynamo is Nvidia’s distributed inference framework for reasoning models. It separates the “prefill” and “decode” stages of LLM inference, dynamically schedules GPUs, routes requests intelligently, and manages KV-cache movement to raise throughput per GPU
   - Nvidia says Dynamo can increase requests served by up to 30x on Blackwell for DeepSeek-R1 workloads

1 https://developer.nvidia.com/blog/introducing-nvidia-dynamo-a-low-latency-distributed-inference-framework-for-scaling-reasoning-ai-models/

NVIDIA’s Adjustments

4. Expanding into CPUs, networking, and DPUs
   - Vera CPU for agentic reasoning, orchestration, analytics, storage, and memory-intensive workloads
   - BlueField DPU for networking, storage, and infrastructure offload
   - ConnectX and Spectrum-X for high-speed AI data-center networking
   - Nvidia acquired Mellanox Technologies
     - Added high-performance interconnect and networking, strengthening its compute-plus-networking position in data centers.
   - Nvidia entered into a definitive agreement to acquire Run:ai
     - added AI workload-orchestration software; after closing, Run:ai said the software would be open-sourced

Strong Financial Position

- **Net cash balance sheet.** $62.6B in cash and securities against just $8.5B of debt at FY26 year-end.
- **Cash generation.** $102.7B of operating cash flow in FY26 on $215.9B revenue (+65.5% YoY). FY26 gross margin was 71.1%, recovering to 75.0% in Q4.
- **Capital returns.** Management formalized a target at March 2026 GTC of returning ~50% of free cash flow to shareholders, up from ~42% ($41.1B) in FY26. That implies $200B+ in buybacks and dividends across CY26–27. NVIDIA repurchased $40.4B of stock in FY26, with $58.5B left on the authorization.
- **Locking in supply.** $95.2B of long-dated commitments to foundries, memory, and packaging partners, plus $27B of multi-year cloud commitments — reserving Blackwell and Rubin capacity that competitors would struggle to match.

Management

- Took the company from a graphics-chip business to a full accelerated-computing platform: GPUs, CUDA software, networking (Mellanox), CPUs (Grace), DPUs, full systems, and enterprise AI software
- Used large purchase commitments to secure supply for Blackwell-scale demand
- $40.4B repurchased in FY26 with $58.5B remaining authorization; net-cash balance sheet maintained through a 100x+ multi-year revenue scale-up
- Selective M&A track record
  - Mellanox $7B (closed 2020) became the foundation for the networking franchise; Run:ai (2024), Bright Computing (2022), DeepMap (2021) added software depth; Arm $40B walked away cleanly when blocked (2022)
- Recovered gross margin after the H20 accounting hit
- Nvidia disclosed in February 2026 that beginning in Q1 FY2027 it would include stock-based compensation (SBC) expense in non-GAAP financial measures

Management

FY22 — Met every near-term financial commitment

- Beat guide each of four quarters; full-year FY22 revenue $26.9B; Omniverse Enterprise GA on schedule

FY23 — Disclosure issue, major internally-driven miss, fast response

- The SEC found that Nvidia failed to disclose that cryptomining was a significant element of Gaming revenue growth in two FY2018 10-Q filings and failed to maintain adequate disclosure controls and procedures
- Q2 FY23 gaming break preannounced Aug 8, 2022
- $1.22B inventory reserve; guidance re-anchored within one quarter

FY25 — Blackwell yield: disclosed, fixed, monetized

- Mask change disclosed Q2; fixed Q3; Q4 FY25 delivered $11B Blackwell revenue, above expectations

FY26 — Regulatory shock absorbed without disrupting growth

- H20 rule disclosed via 8-K within 6 days; absorbed $4.5B charge; revenue $130.5B→$215.9B with no China DC compute

Management has shown that they operate as owners and understand what drives business value.

Earnings Growth

- NVIDIA guided to $78.0B (±2%) revenue for Q1 FY27 in their Q4 FY26 (Feb 25, 2026) call
- FY26 GAAP diluted EPS of $4.90; revenue +65.5% YoY ($130.5B→$215.9B), following +114% the prior year
- Q1 FY27 guide assumes no data-center compute revenue from China; Q4 FY26 GAAP gross margin recovered to 75.0%
- Industry: Dell'Oro 25% CAGR for data-center accelerators (5-year); Gartner $2.52T worldwide AI spend in 2026
- FY26 segment mix: Compute & Networking $193.5B (89.6% of revenue); Data Center end-market $193.7B (89.7%)

Valuation – 80% confidence range

	Bear	Base	Bull
Revenue	$352,000	$367,000	$431,000
EBITDA Margin	60%	64%	70%
EBITDA	$211,200	$234,880	$301,700
Net income margin	54%	55%	57%
Net income	$190,080	$201,850	$245,670
EPS	$7.92	$8.41	$10.24
P/E	25x	27x	30x
Target Price	$198	$227	$307
Upside (%)	5%	20%	63%

Risks

Customer concentration

- One direct customer = 22% of FY26 revenue and another = 14%, both primarily Compute & Networking.

Export controls and U.S.–China policy

- Nvidia disclosed it was effectively foreclosed from competing in China's data-center compute market by FY26 year-end; H20 license rule drove a $4.5B Q1 FY26 charge.

Supplier concentration and component supply

- Dependence on TSMC, Samsung, SK hynix, Micron, plus Hon Hai/Wistron/Fabrinet/CoWoS partners; HBM and advanced packaging are tight.

Hyperscaler in-house silicon substitution

- Custom AI accelerators (e.g., Broadcom co-design with Meta and OpenAI) and growing in-house chip programs at largest buyers chip away at merchant share over time.

Risks (cont.)

Customer power, energy, and capital constraints

- Nvidia notes customer access to data centers, energy, and capital is crucial; power, permitting, and grid constraints can delay deployments.

AI capex slowdown / monetization risk

- If hyperscaler ROI on AI infrastructure disappoints, capex commitments could re-rate; Nvidia has $95.2B mfg/supply commitments substantially due through FY27
- Severity: High | Probability: Medium

Excess inventory & purchase obligations

- Long-dated supply commitments expose Nvidia to reserve charges if demand or regulation shifts — precedents include $1.22B FY23 inventory reserve and $4.5B FY26 H20 charge
- Severity: Medium | Probability: Medium

Geographic and geopolitical concentration

- Manufacturing/supply concentrated across Taiwan, Korea, Israel, China, and Hong Kong; conflict or disruption in any of these regions is a single-point risk
- Severity: High | Probability: Low

NVIDIA: Summary Investment Thesis

- ✓ Dominant AI compute platform
- ✓ Wide-moat franchise anchored on CUDA switching costs and full-stack integration
- ✓ Demand backed by hyperscaler 2026 capex (GOOGL $175–185B, AMZN $128B+, META $72B)
- ✓ Excellent management team
- ✓ Net-cash balance sheet: $62.6B liquidity vs. $8.5B debt; $102.7B FY26 OCF; $40.4B FY26 buybacks
- ✓ Product cycle: Blackwell Ultra shipping; Rubin 2H FY27; Q1 FY27 guide $78.0B

NVIDIA is a good business, in a good industry, with good management, at a good price.