NVDA gideon lens

NVIDIA Corporation (NVDA) - Gideon-Lens Equity Research

Date: 2026-04-26. Analyst: Codex ER2 runner. Current price used in the model: $208.27. Market cap from the queue snapshot: about $5.06 trillion. Fiscal year end: last Sunday in January.

VERDICT BLOCK

Verdict: PASS

Scorecard:

Gate	Result	Notes
Management	PASS	Founder-led, high-return allocator; monitor ecosystem investment circularity.
Moat	YELLOW	CUDA/platform moat is powerful today, but custom silicon and open-stack risk keep 2046 unresolved.
P/E <= 15 on confident EPS at TTM + 5 calendar years	FAIL	Base-case cushion is -11.6%.
Compounder Score	Audited 63 / 100	63/100: exceptional current business economics, but audited down for missing ROIIC evidence and a failed 20% return case.
Business Quality Score	Audited 60 / 80	60/80: high ROIC, cash conversion, management, and moat; capped by uncalculated ROIIC, AI capex concentration, circular-demand risk, and 2046 durability uncertainty.
20% Return UW Score	Audited 3 / 20	3/20: base, bear, and bull 15x cases do not approach 20% annualized returns at $208.27; a much lower entry price is required.

One-line thesis (or reason to pass): NVIDIA is exceptional, but custom silicon/open-stack risk, Compute margin durability, and the base-case valuation miss keep it a PASS under ER2.

Kill / trigger criteria:

- BUY trigger: price below about $184 on the current base EPS, before any extra margin of safety.
- PASS trigger: two consecutive quarters with Data Center or Compute & Networking sequential growth below 5% on average and gross margin below 70%, with China still excluded.
- PASS trigger: evidence that hyperscaler custom silicon or second-source GPUs are taking marginal accelerator economics faster than the base case assumes.
- Upgrade to BUY conditions: a reviewed base-case upgrade above $13.88 of horizon EPS, or a lower entry price that clears 15x base EPS without relying on the bull case.

Position sizing if bought: 0% today under ER2. If a future price or base-case upgrade clears the gate, initial sizing should still be capped because the thesis is concentrated in Compute & Networking margin durability.

ER2 stricter-gate recalibration (2026-05-28): Moat moved from PASS to YELLOW because the current platform moat is strong, but the original 2046 sub-test was YELLOW and the updated rubric requires the main moat gate to reflect that uncertainty. This top scorecard supersedes any softer gate phrasing elsewhere in the appended materials; the underlying evidence was reused, but the gate standard is stricter.

Stage 0 - Time anchor

Today is 2026-04-26. NVIDIA's fiscal year ends on the last Sunday in January; the latest completed fiscal year ended 2026-01-25. The most recently reported quarter is Q4 FY2026, quarter ended 2026-01-25, reported on 2026-02-25. The TTM anchor is Q1 FY2026 ended 2025-04-27 plus Q2 FY2026 ended 2025-07-27 plus Q3 FY2026 ended 2025-10-26 plus Q4

FY2026 ended 2026-01-25, which equals TTM ending 2026-01-25. The forward horizon is exactly five calendar years from that TTM end date: 2031-01-25. Every five-year number in this report compounds from that TTM base, not from a stale fiscal-year anchor.

Stage 1 - Disqualifier sweep

The seven quick disqualifiers did not force an early PASS. ROIC is green: the QF file reports FY2026 ROIC of 62.9%. FCF conversion is green: FY2026 FCF was $96.676B versus net income of $120.067B. Leverage is green: cash plus short-term investments exceeded debt by about $51.1B. End markets are green/yellow because AI infrastructure demand is still strong but cyclicality and capex reflexivity matter. Governance is green/yellow rather than pure green because the governance profile is clean, but strategic ecosystem investments create circular-demand questions. AI disruption exposure is green/yellow because NVIDIA is an AI enabler, yet custom ASICs, open accelerator software, and second-source GPUs are live risks. Pricing power is green today, supported by CUDA, full-stack systems, networking attach, supply access, and scarce leading-edge accelerators.

Management - summary

The management appendix clears the gate with a PASS. Jensen Huang owns the strategy and the capital allocation record across CUDA, accelerated computing, Mellanox/networking, and the current Blackwell/Rubin platform cadence. CFO Colette Kress receives high marks for guidance discipline, including explicit treatment of China/H20 exclusions rather than burying the issue in generic commentary. The proxy and governance review did not surface a dual-class structure, control problem, material weakness, or enforcement issue that would independently block the thesis. Compensation is sufficiently aligned: heavy performance equity, meaningful CEO ownership, clawback language, and hedging/pledging restrictions. The main yellow flag is not managerial competence; it is the possibility that ecosystem investments and customer financing create demand that is economically real but less clean than direct end-customer pull-through. That risk belongs in monitoring and position sizing, not in a management FAIL.

Moat - summary

The original moat appendix cleared the current moat with a PASS while marking the 2046 test YELLOW; under the stricter May 2026 ER2 rubric, that becomes Moat YELLOW rather than a clean PASS. The core moat is in Compute & Networking: CUDA and developer familiarity, full-stack system design, networking attach, supply-chain access, high-trust procurement, and a platform cadence that customers build around. Graphics has a real but narrower moat through GeForce/RTX brand, drivers, workstation certification, DLSS, creator workflows, and ecosystem inertia. The hard question is whether today's margin structure survives a twenty-year shift toward custom silicon, Ethernet fabrics, portability layers, and buyer bargaining power. The moat appendix could write a specific 2046 NVIDIA: a systems, networking, accelerator, and software platform company for AI factories and physical AI. But it could not call that outcome riskless. That is why the 2046 test is YELLOW even though the near-to-medium-term moat is real.

Stage 5 synthesis - headline numbers

The synthesis appendix is arithmetically clean: 11 of 11 tie-out checks passed. The verdict fails on valuation, not on a broken model. In the base case, horizon revenue reaches $589.241B, operating income reaches $343.876B, net income reaches $290.536B, and EPS reaches $12.28 at 2031-01-25. The required EPS for a 15x P/E at the current $208.27 price is $13.88. That means base EPS x 15 supports $184.21, or an 11.6% shortfall versus the current price. The bull case clears, with EPS of $17.45 and EPS x 15 of $261.71, but the framework is not a bull-case gate. The base case must clear with confidence. It does not.

Case	Revenue ($B)	OI ($B)	OI margin	NI ($B)	EPS	EPS x 15	Vs. $208.27
Bear	372.386	176.672	47.4%	150.085	$6.34	$95.16	-54.3%
Base	589.241	343.876	58.4%	290.536	$12.28	$184.21	-11.6%
Bull	777.951	489.384	62.9%	412.762	$17.45	$261.71	+25.7%

Compute & Networking drives the whole answer. The reviewed base case puts that segment at $553.6B of horizon revenue, 62.0% segment operating margin, and $343.2B of segment operating income. Graphics is valuable and profitable, but its base $15.026B of horizon operating income is second-order. The share count path is explicit: $45B per year of buybacks reduces diluted shares only about 3.2% net after modeled SBC issuance, so repurchases do not rescue the P/E gate. The most important sensitivity is Compute & Networking growth and margin. A 200 bps higher C&N revenue CAGR lifts EPS to $13.30,

still short of the $13.88 gate unless margin, tax, or buybacks also improve.

Red team - summary

The red team agrees that the strongest bear argument is marginal economics, not immediate product obsolescence. NVIDIA can remain the AI infrastructure leader and still lose enough incremental economics to custom silicon, Ethernet fabrics, second-source GPUs, or hyperscaler bargaining power to miss the base margin case. The most likely optimistic Stage 5 number is base operating income of $343.876B, because the 58.4% consolidated margin assumes platform economics survive at much larger scale. The exit signal is concrete: two consecutive quarters with Data Center or Compute & Networking sequential growth below 5% on average and gross margin below 70%, with China still excluded.

Kill / trigger criteria

- BUY trigger: price below roughly $184 if the current reviewed base case remains unchanged, or below a revised 15x base EPS value after fresh ER2 review.
- Upgrade trigger: Compute & Networking evidence supports horizon EPS above $13.88 without needing the bull case.
- Downgrade / hard PASS trigger: two consecutive quarters of sub-5% average sequential Data Center/C&N growth plus gross margin below 70%, excluding China recovery noise.
- Downgrade / hard PASS trigger: disclosed custom ASIC, AMD, or other second-source accelerator wins visibly reduce NVIDIA's marginal pricing or attach-rate economics.
- Monitoring trigger: ecosystem investments or customer financing become material enough that reported demand quality is no longer clearly independent.

Analyst flags

The QF and QFR gates cleared. QFR found a $0M FMP-vs-official gap and verified 8 of 8 checks. DR cleared the data package for Wave 2. Segment peer reviewers fixed unsupported product-history and CAGR labels, and both reviewed segment models kept their horizon outputs unchanged. The main remaining data limitations are product-level margins and certain product-level revenue splits, especially standalone Blackwell/Hopper/Rubin/CUDA/DRIVE economics and the FY2026 Graphics taxonomy bridge. These gaps are disclosed in the appendices and do not break the consolidated model because the synthesis uses reviewed reportable-segment bases. Final verdict: PASS at $208.27, with NVIDIA remaining a high-quality watchlist name if price or base-case EPS changes.

Appendix A - Company Primer

NVIDIA Corporation (NVDA) - Appendix A: Company Primer

Prepared: 2026-04-26 Sub-Agent: A Fiscal anchor: FY ends on the last Sunday in January; most recent quarter is Q4 FY2026 ended 2026-01-25, reported 2026-02-25; TTM ends 2026-01-25; forward horizon is 2031-01-25.

Source Note

This primer uses `quant_foundation.md` as the primary source for all financial figures. The data package was not available when this appendix was prepared, so qualitative context is sourced from fresh A-specific primary and FMP pulls saved under `C:\Users\Admin\OneDrive\GPT\ER 2.0\NVIDIA Corporation\raw\a\`. Key sources are NVIDIA's FY2026 Form 10-K, the Q4 FY2026 earnings release and CFO commentary furnished on 2026-02-25, the FMP company profile and Q4 FY2026 transcript endpoint, and the already-reviewed `quant_foundation_review.md`. No local warehouse cache was used.

1. Business Overview

NVIDIA is a full-stack accelerated-computing company. Its historical identity was graphics processing units for PC gaming and professional graphics, but the current business is dominated by data center infrastructure for artificial intelligence, high-performance computing, accelerated data processing, and large-scale networking. The company designs GPUs, CPUs, data processing units, networking adapters, switches, interconnects, full systems, software libraries, enterprise software, reference platforms, developer tools, and application frameworks. The important business point is that NVIDIA is not selling only a chip. It increasingly sells a platform architecture that customers deploy at data-center scale.

The customer problem is compute scarcity. Training and inference for modern AI workloads require parallel compute, high memory bandwidth, low-latency interconnect, and system-level software that lets developers use the hardware without rewriting the entire software stack. In NVIDIA's language, the data center has become the unit of computing. A hyperscaler, AI model maker, enterprise, sovereign customer, or public-sector entity does not simply buy a GPU for a server slot; it buys a path to stand up AI factories that turn power, capital expenditure, models, and software into tokens or accelerated workloads. That framing explains why networking, NVLink scale-up fabric, InfiniBand, Ethernet, CPUs, DPUs, CUDA libraries, AI Enterprise software, and application frameworks have become part of the same sale.

The company serves several customer classes. Direct customers include cloud service providers, original equipment manufacturers, original design manufacturers, system integrators, add-in board partners, distributors, AI model makers, automotive OEMs, tier-1 automotive suppliers, independent software vendors, and other ecosystem participants. Indirect customers include cloud users, neocloud builders, enterprises, AI model makers that buy cloud capacity, public-sector buyers, and sovereign AI initiatives. NVIDIA's FY2026 10-K notes that two direct customers represented 22% and 14% of total revenue, respectively, and both were primarily tied to Compute & Networking. That is a material concentration issue: end demand may be diverse, but revenue recognition still flows through a limited set of large direct channels.

The buying process is hybrid. For data center customers, procurement can involve direct commercial engagement with NVIDIA, system qualification by OEMs or ODMs, integration through system integrators, and cloud deployment through hyperscale partners. NVIDIA's sales motion includes technical pre-sales support from application engineers and solution architects, especially where customers need to optimize infrastructure for generative AI, agentic AI, model training, fine-tuning, inference, and enterprise deployment. For gaming and workstation products, channel partners and add-in board manufacturers remain important. For automotive, the buying cycle is slower and platform-oriented: an OEM or tier-1 supplier commits to an in-vehicle compute and software roadmap that may not turn into vehicle revenue until years later.

2. Revenue Architecture

NVIDIA reports two operating and reportable segments: Compute & Networking and Graphics. It also discloses end-market revenue across Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other. These two taxonomies are related but not interchangeable. The QF file explicitly warns that Data Center revenue is not a mechanical proxy for Compute & Networking revenue. In FY2026, Data Center end-market revenue was $193.737 billion, while Compute & Networking segment revenue was $193.479 billion. The $258 million difference is small relative to total revenue but large enough to prove that downstream modeling must use official reportable segment data for segment economics and end-market data only for demand mix.

Reportable segment revenue in FY2026 was:

Reportable segment	FY2026 revenue	% of total revenue	FY2026 operating income	Segment operating margin
Compute & Networking	$193.479B	89.6%	$130.141B	67.3%
Graphics	$22.459B	10.4%	$9.156B	40.8%
Total segment	$215.938B	100.0%	$139.297B	64.5%

Compute & Networking includes data center accelerated computing and networking platforms, AI solutions and software, and automotive platforms plus autonomous and electric-vehicle solutions. In practical terms, this segment includes the AI infrastructure stack: GPUs, CPUs, DPUs, data center systems, NVLink, InfiniBand, Ethernet networking, software, libraries, enterprise AI tools, and automotive compute platforms. It is the economic center of the company. FY2026 revenue rose from $116.193 billion in FY2025 to $193.479 billion, and segment operating income rose from $82.875 billion to $130.141 billion.

Graphics includes GeForce GPUs for gaming PCs and laptops, GeForce NOW cloud gaming infrastructure and services, gaming platform products, Quadro/NVIDIA RTX workstation GPUs, virtual GPU software for visual computing, and related professional visualization offerings. Although Graphics is much smaller than Compute & Networking, it remains strategically important because the same GPU architecture, tensor cores, RTX software, developer relationships, and gaming ecosystem help maintain NVIDIA's broader platform relevance. FY2026 Graphics revenue was $22.459 billion, up from $14.304 billion in FY2025.

The end-market view shows how demand is actually described to investors and customers:

End-market line	FY2026 revenue	% of total revenue	FY2026 growth
Data Center	$193.737B	89.7%	68%
Gaming	$16.042B	7.4%	41%
Professional Visualization	$3.191B	1.5%	70%
Automotive	$2.349B	1.1%	39%
OEM and Other	$0.619B	0.3%	59%
Total revenue	$215.938B	100.0%	65%

Within Data Center, NVIDIA discloses compute and networking sub-lines. FY2026 Data Center compute was $162.361 billion and networking was $31.376 billion. The networking line matters because it shows NVIDIA's economics are not limited to accelerator cards. AI clusters require scale-up and scale-out fabrics; the Mellanox acquisition has become central to NVIDIA's ability to sell rack-scale and data-center-scale architectures. Q4 FY2026 CFO commentary reported quarterly Data Center compute revenue of $51.3 billion and networking revenue of $11.0 billion, with networking up 263% year over year, driven by NVLink compute fabric for GB200 and GB300 systems and Ethernet/InfiniBand growth.

Geographically, FY2026 revenue by direct-customer headquarters was $149.617 billion in the United States, $42.345 billion in Taiwan, $19.677 billion in China including Hong Kong, and $4.299 billion in other regions. This geography should not be read as final end-customer geography. NVIDIA changed its disclosure in FY2026 to headquarters location of direct customers, and the company states that shipping location and end customer can differ. It also estimates that 76% of FY2026 Data Center revenue from Taiwan-headquartered customers was attributable to end customers based in the United States and Europe. For investor analysis, the geography table is therefore useful for channel concentration and export-control exposure, but it is not a clean demand-by-country table.

3. Unit Economics

At the consolidated level, NVIDIA's FY2026 revenue was $215.938 billion, gross profit was $153.463 billion, operating income was $130.387 billion, net income was $120.067 billion, diluted EPS was $4.90, operating cash flow was $102.718 billion, and free cash flow was $96.676 billion. Free cash flow conversion was 80.5% of net income. Balance sheet liquidity was unusually strong for a hardware company of this scale: cash plus short-term investments were $62.556 billion against total debt of $11.412 billion.

The primary unit-economic feature is very high gross margin and very high operating leverage. NVIDIA's FY2026 GAAP gross margin was 71.1%, below FY2025's 75.0% because of the transition from Hopper HGX systems to Blackwell full-scale data center solutions and a $4.5 billion H20-related excess inventory and purchase-obligation charge. Q4 FY2026 showed margin recovery as Blackwell ramped: Q4 GAAP gross margin was 75.0%, non-GAAP gross margin was 75.2%, and the Q1 FY2027

outlook called for roughly 75.0% non-GAAP gross margin, including stock-based compensation in non-GAAP results beginning in Q1 FY2027.

For Compute & Networking, a dollar of revenue starts with high-value hardware and systems sold into capacity-constrained AI infrastructure markets. The gross margin drivers are product performance per watt, memory and packaging costs, yield, supply-chain premiums, mix between compute and networking, and system-level value capture. NVIDIA outsources wafer fabrication, assembly, test, and packaging, so its cost of revenue includes foundry, memory, substrates, packaging, board/device costs, manufacturing support, warranty, logistics, tariffs, inventory provisions, and related manufacturing stock compensation. Segment operating income is not a pure product gross margin, because the segment table reports "other segment items" rather than separately disclosing segment cost of revenue and operating expense. Still, the FY2026 Compute & Networking segment operating margin of 67.3% demonstrates that the current data center business converts revenue into segment profit at an extraordinary rate.

Graphics has different economics. Gaming and professional visualization products still rely on premium GPU architecture, software features, drivers, developer support, and channel availability, but the customer base is more consumer/workstation oriented and more exposed to product cycles, channel inventory, PC demand, and gaming seasonality. FY2026 Graphics revenue was $22.459 billion and operating income was $9.156 billion, a 40.8% segment margin. That margin is very attractive in most hardware categories, but it is materially below Compute & Networking because channel products do not have the same AI infrastructure scarcity economics and because gaming demand has more traditional hardware-cycle behavior.

Capital intensity is subtle. NVIDIA is fabless, so it does not own the leading-edge fabs that manufacture its wafers. That keeps property and equipment capital intensity far lower than an integrated semiconductor manufacturer. FY2026 capex and productive asset purchases were $6.042 billion, only 2.8% of revenue. However, economic capital intensity has moved into inventory, purchase commitments, supplier capacity arrangements, and strategic investments. Management said inventory and purchase commitments increased as NVIDIA secured supply and capacity beyond the next several quarters, including shipments extending into calendar 2027. That means free cash flow is strong, but working-capital and commitment discipline now matter more than a simple "fabless equals asset-light" label suggests.

4. End-Market Lines Explained

Data Center is the core of the company. It includes accelerated compute platforms for training, inference, data processing, graphics, robotics, and scientific computing, plus networking infrastructure that allows clusters to behave like a single large computer. Customers include public clouds, hyperscalers, AI model makers, enterprises, startups, sovereign buyers, public-sector entities, OEMs, ODMs, system integrators, and distributors. The product form can be a GPU, server subsystem, rack-scale system, networking switch, adapter, DPU, software license, or full reference architecture. The key economic driver is not simply unit shipments; it is how much compute, memory bandwidth, interconnect bandwidth, and software productivity NVIDIA can deliver per watt and per dollar of customer capital.

Gaming is NVIDIA's original mass-market franchise. It includes GeForce RTX GPUs for desktop and laptop PCs, GeForce NOW, gaming console SoCs and services, and software features such as ray tracing and DLSS. The line benefits from a global installed base of PC gamers, creators, developers, and modders. In FY2026, Gaming revenue was $16.042 billion, and Q4 revenue was $3.727 billion. Q4 grew 47% year over year on strong Blackwell demand but declined 13% sequentially as channel inventory normalized after holiday demand. Management expects supply constraints to remain a headwind in Q1 FY2027 and beyond.

Professional Visualization serves designers, engineers, content creators, AI developers, and enterprises that need workstation-class graphics, simulation, virtual workstations, and increasingly local AI compute. Products include NVIDIA RTX PRO GPUs, workstation GPUs, vGPU software, Omniverse-related tools, and software optimized with independent software vendors. FY2026 Professional Visualization revenue was $3.191 billion. Q4 crossed $1.0 billion for the first time at $1.321 billion, up 159% year over year, helped by demand for Blackwell and larger-memory workstation products for AI and agentic workflows.

Automotive is a platform business for automated driving and in-vehicle AI. The DRIVE stack spans data center training and simulation, in-vehicle compute hardware, DRIVE OS, reference sensors, autonomous-driving software, mapping, parking, and intelligent cockpit experiences. It is strategically important because the vehicle can become both an AI endpoint and a robotics platform. It remains financially small relative to Data Center: FY2026 Automotive revenue was $2.349 billion and Q4 was $604 million. Growth was 39% for FY2026 and 6% year over year in Q4, driven by adoption of self-driving platforms.

OEM and Other is the smallest disclosed end-market line at $619 million in FY2026. It includes residual revenue not captured in the four core specialized markets. For NVIDIA, this line is not analytically central except as a reminder that end-market disclosures are not the same as reportable segments and should not be forced into a clean bridge.

5. Platform, Ecosystem, And Buying Logic

NVIDIA's competitive position rests on a platform flywheel. The hardware begins with GPUs and now includes CPUs, DPUs, NICs, switches, NVLink, Ethernet, InfiniBand, and full systems. The software layer includes CUDA, CUDA-X libraries, AI models, SDKs, APIs, NVIDIA AI Enterprise, NIM, NeMo, Blueprints, vGPU, Omniverse, and application frameworks. The ecosystem includes over 7.5 million CUDA and software-tool developers, virtually every major server maker and cloud service provider, universities, startups, model builders, enterprise software developers, and vertical-industry partners.

The reason this matters for purchasing is switching cost. A customer choosing NVIDIA is not only selecting benchmark performance; it is choosing the ecosystem where most frameworks, models, libraries, tools, engineers, and deployment recipes already work. That lowers execution risk and time-to-market. It also creates a challenge for custom ASICs and competing accelerators: they may beat NVIDIA in a narrow workload or internal deployment, but broad market adoption requires software support, procurement availability, developer familiarity, model compatibility, and operations tooling. This does not eliminate competition, especially from hyperscaler internal silicon and China-based alternatives, but it raises the burden of proof for displacement.

The Q4 FY2026 transcript adds management's current demand framing. Management described demand as broader than chatbots, including generative AI upgrades to search, advertising, recommender systems, agentic AI, enterprise AI, sovereign AI, and physical AI. NVIDIA also said its top five cloud providers and hyperscalers collectively account for a little over 50% of Data Center revenue, while growth was led by the rest of the Data Center customer base. That is an important nuance: concentration remains high, but management is arguing that non-hyperscale demand is growing quickly through model makers, enterprises, supercomputing, and sovereign customers.

6. Manufacturing And Supply Chain

NVIDIA uses a fabless and contract-manufacturing model. It relies on leading suppliers for wafer fabrication, assembly, test, packaging, memory, substrates, and final systems. The FY2026 10-K names TSMC and Samsung as foundry partners, SK Hynix, Micron, and Samsung for memory, CoWoS technology for advanced packaging, and contract manufacturers such as Hon Hai, Wistron, and Fabrinet for assembly, testing, and final products. This model lets NVIDIA focus capital and talent on architecture, design, software, quality, platform engineering, sales, and developer support rather than owning fabs.

The risk is that NVIDIA's growth is gated by supplier capacity and geopolitical concentration. The supply chain is mainly concentrated in Asia, though the company says it is expanding into the United States and Latin America to build redundancy and resilience. With Blackwell and Rubin, the constraint is not a simple wafer count. It includes advanced packaging, high-bandwidth memory, networking components, power and cooling, rack integration, system testing, and customer data-center readiness. Management's statement that purchase commitments have been secured beyond the next several quarters implies better demand visibility, but it also raises the cost of a demand forecast error.

7. Competitive Landscape And Key Risks

NVIDIA competes with semiconductor vendors, cloud companies, internal silicon teams, networking suppliers, CPU suppliers, embedded SoC vendors, and regional AI accelerator companies. The FY2026 10-K identifies AMD, Huawei, Intel, Alibaba, Alphabet, Amazon, Baidu, Microsoft, Arm-based CPU suppliers, Broadcom, Qualcomm, Renesas, Samsung, Tesla, Arista, Cisco, HPE, Lumentum, Marvell, and others across relevant product categories. The competitive vectors are performance, total system cost, power efficiency, software support, breadth of product offering, access to customers, manufacturing capability, processor pricing, standards support, and distribution.

The most immediate non-operating risk is export control. NVIDIA says shifting U.S. restrictions have affected its ability to serve customers outside the United States. H20 restrictions triggered a $4.5 billion FY2026 charge, and the company says that as of fiscal year-end 2026 it was effectively foreclosed from competing in China's data center compute market under the current U.S.-China approval landscape. Q1 FY2027 guidance assumes no Data Center compute revenue from China. This matters because lost China availability can create a protected market for local alternatives, which can then build developer and customer ecosystems that challenge NVIDIA outside China over time.

The second risk is customer concentration and reflexivity. If top cloud customers overbuild AI infrastructure, slow capex, shift spend toward internal ASICs, or demand price concessions, NVIDIA's current margins and growth rate could compress.

Conversely, if inference monetization and agentic AI workloads keep growing, those same customers may need more NVIDIA infrastructure. The investment debate therefore hinges on whether customer AI capex is self-funding through revenue and productivity gains or whether it is an unsustainable buildout.

The third risk is supply-chain execution. Blackwell full-scale data center solutions changed the business model from selling more discrete systems to orchestrating larger, more complex infrastructure. That can improve value capture, but it also introduces product-transition risk, inventory risk, customer qualification risk, and supplier bottlenecks. The FY2026 gross-margin decline shows that even a dominant platform can absorb material charges when regulation or product transition changes demand and inventory assumptions.

8. Primer Takeaways For The ER2 Reader

NVIDIA is best understood as an AI infrastructure platform company with a legacy graphics franchise, not as a conventional GPU vendor. The financial base is FY2026/TTM revenue of $215.938 billion, operating income of $130.387 billion, net income of $120.067 billion, and free cash flow of $96.676 billion. The current reported segment model is sharply concentrated: Compute & Networking is 89.6% of revenue and 93.4% of segment operating income, while Graphics is smaller but still highly profitable.

The end-market lines must be handled carefully. Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other are demand categories, not reportable segments. Data Center dominates the end-market view at 89.7% of FY2026 revenue, but segment economics should use Compute & Networking and Graphics. Data Center itself now includes a large networking business, with FY2026 networking revenue of $31.376 billion and Q4 networking revenue of $11.0 billion.

The moat question is inseparable from the software and ecosystem question. CUDA, libraries, AI Enterprise, NIM, NeMo, vGPU, Omniverse, developer relationships, server-maker availability, cloud deployment, and model compatibility are not side products; they are the reasons buyers accept NVIDIA as the default accelerated-computing architecture. The bear case must therefore show either that customers can replicate enough of this platform with internal or third-party silicon, or that the value of the platform is competed away despite continued workload growth.

The key forward debate is not whether FY2026 was strong. The reconciled numbers show extraordinary profitability and cash generation. The key debate is durability: whether AI training and inference demand, agentic AI, sovereign AI, enterprise AI, networking attach, physical AI, and annual architecture transitions can support high returns through the 2031-01-25 horizon without a major margin reset, customer capex pause, export-control shock, or accelerator substitution cycle.

Appendix M - Management Assessment

NVIDIA Corporation - Management Assessment

Source discipline: this appendix used live SEC and company investor-relations sources for the proxy, Form 10-K, official earnings materials, and call evidence. Capital-allocation and corporate-action inputs come from `data_package.md` Sections 7 and 9, as instructed. I did not use local warehouse caches. For Q&A text where NVIDIA's IR site provides webcast access but not full transcript text, I used live FMP `earning-call-transcript` pulls for the last eight quarters, saved under `raw/m/`.

1. Who is running the company

NVIDIA is still founder-led. Jen-Hsun "Jensen" Huang co-founded NVIDIA in 1993 and remains President, CEO, and a director. The latest FY2026 Form 10-K states that NVIDIA "is now a data center scale AI infrastructure company," and Huang owns the strategic shift from graphics processor supplier to full-stack AI infrastructure platform. That matters: the current strategy is not a consultant-built pivot after a cycle peak. It is the latest expression of a multi-decade operating system built around accelerated computing, CUDA, developer lock-in, annual architecture cadence, and partner ecosystem expansion.

Colette M. Kress has been Executive Vice President and Chief Financial Officer since 2013. Before NVIDIA she held senior finance roles at Cisco, including CFO of its Business Technology and Operations Finance organization, and earlier served at Microsoft, including CFO of the Server and Tools division. This is a high-quality CFO background for a business whose current bottleneck is not simply accounting control but planning huge supply commitments, ecosystem investments, capital return, and geopolitical revenue exclusions without confusing investors about the base case.

The material operating bench is long-tenured. Ajay Puri joined in 2005 and leads Worldwide Field Operations. Debora Shoquist joined in 2007 and leads Operations, directly relevant to the Blackwell and Rubin rack-scale supply chain ramps. Timothy Teter joined in 2017 and is General Counsel and Secretary, relevant because export controls, commercial commitments, AI safety, IP, and ecosystem investments are now core risk categories.

Board oversight is credible but founder-centric. The 2025 proxy lists 13 director nominees and states that every director except Huang is independent under Nasdaq standards. Stephen C. Neal is Lead Director, has served since 2019, chairs the nominating and corporate governance committee, and has a legal/governance background from Cooley and Levi Strauss. The proxy highlights annual elections, majority voting, proxy access, independent executive sessions, stockholder outreach, and board refreshment, including four new directors and three retirements since 2020. The offset is tenure: A. Brooke Seawell has been on the board since 1997, Tench Coxe since 1993, Mark Stevens since 2008 after an earlier board term, and Robert Burgess since 2011. Long tenure is not automatically bad at a founder-led platform company, but it means the practical governance test is whether independent directors challenge capital allocation and related-party optics, not whether formal independence boxes are checked.

2. Tenure and track record

Huang has been CEO for more than three decades, so the current strategy is his record. The three biggest strategic decisions are visible in the numbers. First, CUDA and accelerated computing turned the GPU from a gaming component into a programmable platform. The payoff is clearest in data center: FY2026 Data Center revenue was $193.7 billion, up 68% year over year, and company revenue rose from $26.9 billion in FY2022 to $215.9 billion in FY2026. Second, the Mellanox acquisition and networking strategy converted NVIDIA from a chip supplier into a data-center scale systems company. Management said FY2026 networking revenue exceeded $31 billion, more than 10 times fiscal 2021. Third, the Blackwell rack-scale transition moved NVIDIA deeper into full system architecture. It created execution risk in packaging, liquid cooling, memory, networking, software, and supply commitments, but Q4 FY2026 Data Center revenue of $62.3 billion suggests the team managed the transition well.

The record is not error-free. NVIDIA has had cycles: FY2023 revenue was basically flat versus FY2022, operating margin fell sharply, and the gaming correction was real. The management quality signal is that the company did not respond by starving R&D or abandoning platform strategy. R&D rose from $5.3 billion in FY2022 to $18.5 billion in FY2026. NVIDIA also had a major export-control disruption in FY2026. The data package records the April 9, 2025 U.S. H20 control, $4.6 billion of H20 revenue recognized before the control date, a $4.5 billion inventory and purchase-obligation charge, $2.5 billion of revenue unable to ship in Q1, and about $8 billion of lost H20 revenue embedded in Q2 guidance. Management did not hide the hit or include China recovery in base guidance; it repeatedly excluded China Data Center compute revenue from forward outlooks.

That is a positive disclosure marker even though the geopolitical risk remains unresolved.

Kress's track record is strong. She has presided over a balance sheet that ended FY2026 with $62.6 billion of cash plus short-term investments and roughly $51.1 billion of net cash after total debt, while still returning $41.1 billion to shareholders in FY2026. The CFO function has also kept guidance usable through violent growth. Q3 FY2026 guidance called for $65 billion plus or minus 2%; Q4 actual revenue was $68.1 billion. Q4 FY2026 guidance for Q1 FY2027 was $78 billion plus or minus 2%, explicitly excluding China Data Center compute revenue. This is the right kind of conservatism for a company where upside and downside can be tens of billions of dollars.

3. Capital allocation behavior

The data package Section 7 shows a clear hierarchy: invest in R&D and supply/ecosystem first, return surplus capital second, use M&A selectively, and keep the dividend symbolic. Capex increased from $1.0 billion in FY2022 to $6.0 billion in FY2026, while R&D increased from $5.3 billion to $18.5 billion. R&D as a percentage of revenue fell from 19.6% in FY2022 and 27.2% in FY2023 to 8.6% in FY2026, but that is denominator scale rather than R&D neglect. In dollars, the company is spending like a platform leader. Capex/revenue remains modest at 2.8% in FY2026 because NVIDIA is fabless, but the 10-K says FY2027 capex should increase relative to FY2026 to support future growth. The more important capital commitment is off-balance-sheet supply planning, inventory, purchase obligations, and ecosystem investment.

Buybacks have become enormous. NVIDIA repurchased $40.1 billion of stock and paid $1.0 billion of dividends in FY2026, versus $96.7 billion of free cash flow. The FY2026 10-K says the board approved an additional $60 billion repurchase authorization on August 26, 2025, and that $58.5 billion remained authorized at fiscal year-end. The issue is price sensitivity. With a multi-trillion-dollar market capitalization through much of the period, repurchases are partly mechanical capital return, not obviously opportunistic value buying. Kress answered a Q4 FY2026 question on why not make a much larger buyback by stressing supply chain support, developer ecosystem investment, strategic investments, continued buybacks, and dividends. That answer is coherent: the highest-return marginal dollar may be securing supply and ecosystem demand.

M&A and strategic investments are the main yellow flag. NVIDIA's best acquisition example remains Mellanox: it changed the architecture of the company and now supports a $31 billion-plus networking business. Recent FY2026 items are less proven. The data package Section 9 notes strategic and ecosystem investments, including a referenced $10 billion Anthropic investment and a large Groq-related nonexclusive license/payment line in investing activities. The 10-K also says NVIDIA expects to continue investing in its ecosystem and is finalizing an OpenAI investment and partnership agreement, with no assurance of completion. These investments may strengthen the CUDA and AI factory flywheel, but they create circular-demand risk: if NVIDIA funds customers or partners who then buy NVIDIA systems, investors need to distinguish durable end demand from ecosystem financing. Management addresses the concern directly in Q&A by framing the balance sheet as a tool for supply resilience and ecosystem development. That is rational, but it deserves ongoing red-team scrutiny.

4. Compensation alignment

The compensation program is above average for alignment, though not perfect. The 2025 proxy says more than 90% of the CEO's target pay and more than 45% of other named executive officers' target pay has been performance-based and at risk, and 100% of the CEO's equity awards have been performance stock units. FY2025 metrics were simple: revenue for variable cash, non-GAAP operating income for single-year PSUs, and three-year TSR relative to the S&P 500 for multi-year PSUs. Fiscal 2025 revenue of $130.5 billion, non-GAAP operating income of $86.8 billion, and 100th percentile relative TSR drove maximum payouts.

Huang's FY2025 target compensation was $32.0 million, up 28%, with $27.5 million target equity split equally between single-year and multi-year PSUs. His reported FY2025 summary compensation was $49.9 million, including $38.8 million of stock awards and $6.0 million of non-equity incentive compensation. The increase is large, but the proxy gives a reasonable explanation: first base-salary increase in 10 years, target equity increased 25%, and CEO equity capped at 150% payout rather than 200% because all his equity was performance-based. Kress's FY2025 target compensation was $15.0 million, with $13.8 million of target equity split across SY PSUs, MY PSUs, and RSUs. Her summary compensation was $21.4 million.

The ownership alignment is unusually strong. The proxy shows Huang beneficially owned 922.9 million shares as of March 24, 2025, or 3.77% of shares outstanding. Kress owned 4.9 million shares. The governance policies require the CEO to hold shares worth 10 times base salary and other named executives three times base salary, and each named executive exceeded the requirement. The proxy also prohibits hedging, pledging, margin accounts, and derivatives involving NVIDIA stock; maintains a compensation recovery policy; caps performance-based payouts; and does not pay dividends on unearned or unvested equity. The most obvious weakness is that revenue and non-GAAP operating income can reward cycle strength as

well as managerial skill. Relative TSR partly offsets that, and the dollar ownership of the CEO offsets it further.

Stock-based compensation is not ignored. In Q4 FY2026, NVIDIA announced that beginning Q1 FY2027 it would include stock-based compensation expense in non-GAAP measures. That improves candor and comparability, especially for a company competing for AI talent with equity-heavy compensation. FY2026 official guidance included $1.9 billion of stock-based compensation in Q1 FY2027 operating expenses, and the FY2025 proxy reconciliation showed $4.7 billion of stock-based compensation added back to fiscal 2025 non-GAAP operating income. The direction of travel is positive.

5. Disclosure quality and candor

I reviewed the last eight quarters of live transcript pulls, FY2025 Q1 through FY2026 Q4. The disclosure pattern is better than average for a mega-cap technology company but still selective. On tough questions, Huang often answers at length and with strategic logic rather than short numeric disaggregation. When analysts asked about customer CapEx sustainability, AI ROI, ASIC competition, Blackwell execution, China, and gross margin sustainability, he did not deny the concerns. He reframed them around tokens, power, performance per watt, CUDA, annual product cadence, and customer revenue generation. That style can sound promotional, but it gives the model for why management thinks demand persists.

Kress is more numerically direct. In the FY2026 Q1 call, she quantified H20 revenue recognized, revenue unable to ship, the purchase-obligation and inventory charge, and the Q2 revenue loss. In later quarters, she repeatedly stated whether China Data Center compute revenue was included in outlook. In Q4 FY2026, she gave Q1 FY2027 revenue, margin, operating expense, stock-based compensation, and tax guidance, and she said sequential Data Center growth was expected through calendar 2026. The data package shows Q4 FY2026 actual revenue exceeded Q3 guide. This is a high level of guidance discipline in a volatile demand and supply environment.

Voluntary disclosure is mixed. NVIDIA gives excellent end-market revenue and growth commentary and increasingly large networking disclosures, but it still does not provide full product-level or customer-level economics. That is understandable competitively, but the model sensitivity is high. A company where the top five cloud providers and hyperscalers are about half of Data Center revenue should give investors enough concentration, financing, commitment, and margin context to evaluate digestion risk. Management provides narrative and some figures, but not full unit economics or customer ROI data. The disclosure gate therefore passes, but it is not flawless.

Language markers are worth noting. FY2026 management used plain language for China restrictions and H20 lost revenue rather than burying the issue in generic "headwinds." It used "supply constraints" and "channel inventory" in Gaming when sequential revenue fell in Q4, which is reasonably specific. On AI demand, however, management uses very strong conceptual language. The phrase "compute equals revenues" is the core claim. That is a powerful strategic statement, but it is also the underwriting risk: if customers' token revenue, enterprise adoption, or capital markets support weakens, NVIDIA's confident language will have to be tested against cancellations, lower purchase commitments, and pricing pressure.

6. Strategic coherence

The strategy is unusually coherent. NVIDIA says it is an AI infrastructure company. The investment pattern supports that: R&D up $13.2 billion over five years; capex up to $6.0 billion; supply commitments extending beyond near quarters; networking now a $31 billion-plus business; ongoing annual platform cadence from Hopper to Blackwell to Rubin; and ecosystem investments in AI labs, developers, suppliers, and adjacent technologies. The strategic destination is clear: be the full-stack platform for AI factories across training, inference, networking, data processing, sovereign AI, enterprise AI, robotics, and physical AI.

The strongest evidence of coherence is that management's major decisions reinforce one another. CUDA creates the software base. GPUs and CPUs create the compute base. NVLink, InfiniBand, and Spectrum-X create the scale fabric. Mellanox created networking credibility. Blackwell and Rubin push toward rack-scale systems. Strategic investments deepen customer and developer dependence. The same logic explains why Huang talks about AI factories rather than chips.

Management also names competitive threats more directly than many leaders in dominant positions. The FY2026 Q2 transcript includes a direct question on customer ASIC projects and Broadcom. Huang answered by distinguishing NVIDIA's full-stack co-design from single-purpose ASICs and emphasizing model churn, software, networking, and time-to-production risk. The Q4 FY2026 transcript also acknowledges China competitors and the risk that America loses platform influence if U.S. chips cannot compete in China. These are not trivial admissions. They support a PASS on strategic candor, although the confidence level should not be mistaken for certainty.

7. Governance flags

No dual-class structure was identified. NVIDIA common shareholders have one vote per share, and the proxy highlights annual elections, majority voting, proxy access, special-meeting rights, independent committees, and stockholder outreach. The FY2026 10-K says disclosure controls and internal control over financial reporting were effective as of January 25, 2026, and there were no changes in internal control during the quarter that materially affected or were reasonably likely to materially affect it. I found no current SEC enforcement action, restatement, or disclosed material weakness in the primary filings reviewed.

Related-party disclosure is not zero. The 2025 proxy discloses that Huang's daughter and son are employed by NVIDIA, as is a director's son. None is an executive officer or reports directly to the related director or executive, compensation was set under ordinary employee practices without the related person's involvement, and FY2025 compensation was approximately $1.13 million, $0.53 million, and $0.475 million, respectively. This is a yellow footnote, not a fail. The amounts are immaterial to NVIDIA, but the company should continue to disclose these arrangements plainly because founder-led companies cannot afford soft related-party standards.

Insider selling requires monitoring but is not disqualifying. The FY2026 10-K Item 9B discloses that CFO Colette Kress adopted a Rule 10b5-1 trading arrangement on December 18, 2025 for 500,000 shares expiring March 23, 2027. SEC Form 4 evidence also shows Huang owns a very large direct and indirect stake, including 70.6 million direct shares after March 2026 award activity plus hundreds of millions held through trusts and related vehicles. Planned selling by a founder after a multi-trillion-dollar value creation run is not inherently a governance problem, but clustered insider sales would be a yellow flag if paired with deteriorating demand, customer financing stress, or reduced guidance candor.

8. Peer comparison

Versus Lisa Su at AMD, Huang has the stronger per-dollar capital allocation record over the last five years. Su has executed exceptionally well from a weaker starting position, especially in CPUs and data center accelerators. But NVIDIA's conversion of GPU leadership, CUDA, networking, and full-stack systems into FY2026 revenue of $215.9 billion and FY2026 ROIC above 60% is a better compounded outcome. Su may communicate with slightly more restraint; Huang communicates with more strategic force and a higher promotional register.

Versus Broadcom's Hock Tan, Huang is less obviously price-disciplined on buybacks and ecosystem investments, while Tan's acquisition and integration track record is elite. Broadcom is arguably the cleaner traditional capital allocator: hard underwriting, margin discipline, and ruthless integration. NVIDIA's record is better as a founder-led platform builder. It has created a market rather than merely consolidating one.

Versus Intel's current management transition, NVIDIA is in a different category. Intel is still rebuilding manufacturing credibility, product cadence, and capital discipline. NVIDIA has execution complexity but not the same strategic drift. On five-year stewardship, Huang ranks at the top of this peer set. The only reason not to call him the clean best-in-class manager without qualification is the new scale of ecosystem investments and possible circularity in AI infrastructure funding.

9. Management gate verdict

Management gate: PASS.

NVIDIA has a proven founder CEO, a long-tenured CFO, a deep operating bench, strong board formalities, meaningful executive ownership, and a compensation structure that is unusually aligned for a mega-cap. The team has created massive value through coherent capital allocation: sustained R&D, Mellanox/networking, CUDA ecosystem reinforcement, Blackwell/Rubin cadence, and disciplined guidance through export-control disruption. Disclosure quality is good, especially around H20/China and the inclusion of SBC in non-GAAP results. The governance and allocation caveats are real: long-tenured directors, related-party employment disclosures, insider selling plans, and large strategic AI ecosystem investments that could blur demand quality. Those are monitoring items, not gate failures. I would follow this team for a decade, with quarterly checks on ecosystem investment circularity, China exclusions, top-customer CapEx health, and buyback discipline.

Appendix W - Moat and 2046 Durability

NVIDIA Corporation - Moat and 2046 Durability

Source frame

This appendix uses the ER2 NVIDIA data package, live FMP annual income statement pull saved under `raw/w/fmp_income_statement_annual_10y.json`, NVIDIA's FY2026 Form 10-K on SEC EDGAR, NVIDIA's Q4 FY2026 release, and current competitive context from AMD, Intel, and Broadcom primary sources. The segment appendix inputs were checked but were still placeholders at the time of this run, so this appendix does not rely on them. Current reportable segments are **Compute & Networking** and **Graphics**; Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other are end-market disclosures, not clean reportable-segment bridges.

1. Moat taxonomy applied

Network effects. NVIDIA has a real ecosystem network effect, but it is not a consumer social network effect. It is a developer, model, framework, OEM, cloud, and system-integration network effect around CUDA, CUDA-X libraries, AI Enterprise, NIM, NeMo, Omniverse, vGPU, deployment recipes, trained engineers, server qualifications, and cloud instance availability. The FY2026 10-K describes a unified programmable architecture, a large developer and installed base, and a software stack that runs on all NVIDIA GPUs with domain-specific libraries, SDKs, APIs, and frameworks. The practical evidence is that buyers do not only procure chips; they procure a path to train and deploy workloads with lower execution risk. This is strongest in Compute & Networking, where accelerated compute, networking, systems, and software are sold as an AI infrastructure platform. It is present but less decisive in Graphics, where GeForce, RTX, DLSS, creator tools, professional software certification, and installed-base familiarity matter, but the customer can still delay or substitute more easily.

Switching costs. Switching costs are high in Compute & Networking because AI infrastructure is not a drop-in component decision. Migrating from NVIDIA to AMD ROCm, Google TPU, AWS Trainium/Inferentia, Broadcom/custom ASIC platforms, Huawei, or another accelerator requires workload porting, kernel optimization, model validation, cluster operations changes, procurement requalification, and developer retraining. For hyperscalers with enough internal engineering, the switching cost is surmountable for narrow workloads; for enterprises, neoclouds, sovereign AI customers, model builders, and most software vendors, it remains a major barrier. The switching cost is weaker in Graphics gaming, where consumers buy GPUs by performance, price, game support, driver quality, and brand; it is stronger in Professional Visualization because enterprise workstation workflows depend on certified hardware/software combinations.

Scale economies. Scale is the second core moat. NVIDIA's FY2026 revenue was $215.938 billion, with $193.479 billion in Compute & Networking and $22.459 billion in Graphics. The company can fund annual architecture transitions, system software, networking, reference designs, supply commitments, developer support, and ecosystem investments at a scale few competitors can match. Scale also improves supply access in constrained nodes, HBM, CoWoS packaging, system assembly, networking, power, and liquid-cooled racks. This is not risk-free: the 10-K explicitly says long lead times, non-cancellable commitments, customer order changes, and complex data center buildouts can amplify inventory and margin risk. But in the current cycle, scale is a barrier because customers want complete, available systems, not just benchmark slides.

Brand and trust. Brand is meaningful but secondary. In consumer graphics, GeForce, RTX, DLSS, Reflex, G-SYNC, and driver reputation support premium positioning. In data center, the relevant brand is high-trust procurement: NVIDIA is the default accelerated-computing architecture for frontier AI, public cloud instances, model training, inference, supercomputing, and enterprise AI. The evidence is not marketing language alone; it shows up in customer willingness to concentrate capital spending with NVIDIA despite high ASPs and visible alternatives. The weakness is that brand can reverse quickly if a product transition fails or if customers conclude total cost of ownership is materially better elsewhere.

Regulatory moat. NVIDIA does not have a regulatory moat. Regulation is mostly a moat reducer. Export controls have already limited China data center compute revenue, created H20 inventory and purchase-obligation charges, and may help Chinese alternatives build local ecosystems. Antitrust and AI regulation also add risk: the FY2026 10-K notes information requests from regulators in multiple jurisdictions and possible restrictions on hardware, software, and systems used for frontier models. Regulation may raise compliance burdens for smaller competitors, but for NVIDIA the net effect is not protective.

IP and patents. IP matters, but it is not sufficient by itself. NVIDIA has patents, trade secrets, architecture know-how, software libraries, and decades of GPU and networking expertise. The 10-K says patents expire from March 2026 to June 2045, which is notable for a 2046 test: formal patent protection cannot be the main 2046 moat. The durable asset is

cumulative design skill and software/platform integration, not a static patent wall.

High-trust procurement. This is highly relevant. AI factories involve billions of dollars of capital, scarce power, scarce HBM, advanced packaging, data center construction, liquid cooling, networking, systems integration, and model revenue assumptions. Customers must believe the platform will work at scale and will be supported across multiple generations. NVIDIA's moat is strongest where this high-trust procurement overlaps with ecosystem switching costs: large customers buy Blackwell/Rubin systems because the cost of being late or underperforming on AI may be higher than the cost of paying NVIDIA margins.

2. Per-segment moat

Compute & Networking. The segment moat is the full-stack accelerated computing platform: GPU architecture, CUDA/software, NVLink, InfiniBand/Ethernet, DPUs/NICs/switches, CPUs, full rack systems, model libraries, developer mindshare, and cloud/OEM availability. The segment generated $193.479 billion of FY2026 revenue, 89.6% of company revenue, and $130.141 billion of segment operating income at a 67.3% segment operating margin. Durability is strong over five years because customers are still qualifying and scaling Blackwell/Rubin-class systems and because networking attach has become a large revenue line. Ten-year durability is credible but more conditional: hyperscaler internal ASICs and AMD/Broadcom/open accelerator ecosystems can take optimized inference workloads if they become operationally easier and cheaper. Twenty-year durability is not guaranteed because semiconductor platforms turn over and today's CUDA advantage must be continuously renewed. Earliest break signal: NVIDIA data center revenue still grows but margins compress sharply while hyperscaler custom silicon volumes scale and third-party frameworks treat NVIDIA as one backend among many.

Graphics. The segment moat is a premium graphics and workstation ecosystem: GeForce brand, RTX/DLSS software features, driver cadence, gaming ecosystem support, workstation certifications, and professional visualization workflows. FY2026 Graphics revenue was $22.459 billion with $9.156 billion of operating income and a 40.8% margin. Five-year durability is good; NVIDIA remains the premium GPU brand and Blackwell helped Gaming and Professional Visualization grow strongly in FY2026. Ten-year durability is moderate because gaming GPU demand is cyclical and more price-sensitive, while AI PC/workstation workflows may move partly to cloud inference or integrated NPUs. Twenty-year durability is weaker than Compute & Networking unless professional AI workstation and simulation workflows remain anchored to RTX/CUDA. Earliest break signal: GeForce share gains require sustained price cuts, DLSS/RTX differentiation matters less to game adoption, and workstation software certifications shift toward vendor-neutral or cloud-native rendering/AI backends.

The blended company moat is not weaker because a large weak segment drags it down; the opposite is true. Compute & Networking is roughly 90% of revenue and more than 93% of segment operating income, so the company moat is overwhelmingly the data center AI infrastructure moat. Graphics is profitable and reinforces the developer/graphics heritage, but the equity moat lives or dies with Compute & Networking.

3. Margin persistence evidence

The ten-year live FMP income statement pull shows consolidated operating margin of 28.0% in FY2017, 33.0% in FY2018, 32.5% in FY2019, 26.1% in FY2020, 27.2% in FY2021, 37.3% in FY2022, 15.7% in FY2023, 54.1% in FY2024, 62.4% in FY2025, and 60.4% in FY2026. That is not a smooth consumer-staples-style margin line; it is a cyclical semiconductor company that became an AI infrastructure toll road. The FY2023 gaming correction matters because it proves the moat is not immune to inventory cycles. The FY2024-FY2026 step-change matters more because the company translated an extreme demand shock into very high operating margins while competitors were visibly attacking the same profit pool.

Segment margin evidence is stronger for Compute & Networking than for Graphics. Compute & Networking margins moved from 41.6% in FY2022 to 33.7% in FY2023, then 67.5%, 71.3%, and 67.3% in FY2024-FY2026. The FY2026 dip was still at an extraordinary level and occurred despite Blackwell system mix, scale-up complexity, and H20-related charges. Graphics margins were 53.5%, 38.2%, 43.3%, 35.5%, and 40.8% over the same five years. Graphics is profitable but more cyclically exposed. Consolidated FY2026 GAAP gross margin declined to 71.1% from 75.0% in FY2025, while FY2026 operating margin remained 60.4%; this is a moat-positive but not moat-perfect record. The evidence says NVIDIA has pricing power and mix power in AI infrastructure, but also that product transition, export control, and demand forecast errors can dent margins.

4. AI disruption stress test

Compute & Networking is not a simple workflow that a foundation model can replicate. The value is chokepoint economics plus platform adjacency: scarce physical capacity, leading-edge silicon design, advanced packaging, HBM supply,

power-efficient systems, cluster networking, CUDA/software, and procurement trust. AI increases demand for this segment because better models, agents, multimodal systems, recommender systems, search, coding, robotics, and simulation increase training and inference workloads. The disruption risk is not that AI writes a better CUDA overnight; it is that AI helps competitors and customers generate kernels, compilers, portability layers, chip designs, and workload-specific accelerators that reduce the switching cost over time. AI can erode the software moat, but it does not dissolve the hardware, supply, power, and high-trust deployment moat quickly.

Graphics is more exposed to AI disruption because parts of its value are workflow-like: image generation, video generation, rendering assistance, game asset creation, and workstation visualization workflows can shift to cloud AI services or model-native tools. At the same time, Graphics is also platform-adjacent because RTX GPUs run local inference, creator tools, gaming upscaling, AI-assisted rendering, and workstation workloads. The most exposed segment is Graphics, especially Professional Visualization workflows that could move from local workstation hardware to cloud AI pipelines. The most protected segment is Compute & Networking, especially data center systems and networking, where AI demand increases physical compute requirements and where power, supply, and cluster reliability remain chokepoints.

5. The 2046 test

Today's date is 2026-04-26. Twenty years forward is 2046-04-26. A credible 2046 NVIDIA business description reads as follows: NVIDIA sells accelerated computing platforms for AI factories, simulation factories, robotics fleets, scientific computing, autonomous systems, and high-fidelity visual/physical worlds. Its main buyers are global cloud providers, sovereign AI operators, industrial companies, defense and research institutions, model builders, robotics and vehicle platforms, and enterprise infrastructure providers. Compute & Networking still exists, though the exact chip names are different: GPUs may be only one accelerator type inside a heterogeneous system that includes CPUs, DPUs, optical or photonic interconnect, memory systems, networking fabrics, compilers, scheduling software, model-serving systems, and simulation stacks. Graphics still exists but is smaller relative to infrastructure, with gaming, local AI PCs, professional simulation, and creator hardware/software as durable but cyclical franchises. Automotive as a disclosed end-market may have grown into robotics/physical AI infrastructure, but it is unlikely to be the company-defining segment unless NVIDIA captures fleet-level software economics, not just in-vehicle silicon.

Directionally, 2046 operating margin should be lower than FY2026's 60.4% if competition normalizes and hyperscalers internalize more silicon, but still structurally high for a hardware-linked company if NVIDIA keeps the platform layer. A plausible long-run margin might be closer to a high-30s to mid-40s operating margin than the current supercycle level. The business still makes money because the intact moat is not a single chip; it is architecture cadence, installed software base, developer familiarity, system integration, networking, supply orchestration, procurement trust, and the ability to turn new compute bottlenecks into standardized platforms.

The bear-case 2046 description is also specific. NVIDIA becomes one of several high-end accelerator vendors. Hyperscalers run most mature inference on internal ASICs supplied by Broadcom-like partners, AMD or open accelerator ecosystems own enough training and inference workloads to discipline pricing, China and other regional markets run local stacks, and AI-generated compilers make software portability less painful. NVIDIA still sells important systems, but CUDA is no longer the default center of gravity and networking is contested by Ethernet/custom fabrics. In that world, revenue may be large, but returns are closer to high-quality semiconductor norms rather than an AI infrastructure monopoly. I assign this bear case a moderate probability over twenty years: not the base case for the next five to ten years, but too plausible to award an unqualified 2046 PASS.

6. Competitive threat inventory

Threat	Mechanism	Segment(s) affected	Earliest observable signal	Plausibility
Hyperscaler custom silicon	Google TPU, AWS Trainium/Inferentia, Microsoft/other internal accelerators, and Broadcom custom XPUs take mature inference and some training workloads where workloads are stable and scale justifies porting	Compute & Networking	Top cloud customers slow NVIDIA order growth while internal accelerator utilization rises; NVIDIA maintains units but loses margin	High

Threat	Mechanism	Segment(s) affected	Earliest observable signal	Plausibility
Merchant/open accelerator competition	AMD Instinct/ROCm, open standards, Huawei in China, and future accelerators improve software support and rack-scale availability	Compute & Networking; some Graphics AI workstation demand	ROCm/open compiler portability becomes good enough for mainstream model builders; AMD wins large recurring clusters outside one-off deployments	Medium-high
Export controls and regional ecosystem split	U.S.-China restrictions foreclose China data center compute and help local alternatives build developer/customer ecosystems	Compute & Networking	China revenue remains excluded while Chinese accelerators improve benchmarks and software outside China	High
Gaming/pro visualization commoditization	AI-assisted rendering, cloud GPUs, integrated NPUs, and price competition reduce discrete premium GPU differentiation	Graphics	GeForce ASP/margin falls despite healthy gaming hours; workstation users shift to cloud AI/rendering workflows	Medium

7. Pricing power verification

NVIDIA's pricing power is observable indirectly, not through disclosed list-price escalators. The strongest evidence is margin and mix: FY2026 revenue rose 65% to $215.938 billion, Q4 FY2026 gross margin was 75.0%, FY2026 GAAP gross margin was 71.1%, and consolidated FY2026 operating margin was 60.4% despite export-control charges and system transition complexity. Compute & Networking sustained a 67.3% segment operating margin while becoming nearly 90% of revenue. Q1 FY2027 guidance called for $78.0 billion of revenue and approximately 75.0% gross margin while assuming no Data Center compute revenue from China. That combination is difficult to reconcile with weak pricing power.

The caveat is contract structure. NVIDIA's 10-K says most sales are purchase-order based and customers can generally cancel, change, or delay commitments with little notice and without penalty. It also says NVIDIA enters supply and capacity commitments and may face mismatch risk if customers change orders. That means pricing power is stronger in product-market fit and scarce supply than in contractual lock-in. There is no clean public evidence of multi-year price escalators or minimum commitments that fully protect NVIDIA margins through a downturn. Pricing power should therefore be scored strong today but procyclical.

8. Moat verdict

Moat gate: PASS. NVIDIA's moat is concrete and observable: Compute & Networking combines CUDA/software ecosystem, accelerated-computing architecture, networking, scale, supplier access, and high-trust procurement, and the margin evidence supports real pricing power rather than only narrative dominance. Graphics has a narrower but still real moat in premium gaming and workstation ecosystems, though it is not the valuation driver. The moat is strongest over five years, credible over ten years, and dependent on continuous execution beyond that.

2046 test: YELLOW. A specific 2046 NVIDIA can be written: a heterogeneous AI infrastructure and simulation-platform company selling systems, software, networking, and accelerators into cloud, sovereign, enterprise, robotics, science, and industrial buyers. The reason it is YELLOW rather than PASS is that twenty years is long enough for custom silicon, open software, regional stacks, and AI-assisted portability to erode CUDA's default status and normalize margins. The present moat passes the ER2 gate, but the 2046 durability claim should be haircut: the base case is continued high returns, not FY2026-level economics preserved unchanged.

Sources

- NVIDIA FY2026 Form 10-K, SEC EDGAR: https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm
- NVIDIA Q4 and FY2026 results release: https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2026
- AMD Advancing AI 2025 release: https://www.amd.com/en/newsroom/press-releases/2025-6-12-amd-unveils-vision-for-an-open-ai-ecosystem-detai.html
- Intel FY2025 Form 10-K: https://www.intc.com/filings-reports/all-sec-filings/content/0000050863-26-000011/intc-20251227.htm

- Broadcom Q4 and FY2025 results release: https://investors.broadcom.com/node/63716/pdf

Appendix B1 - Compute & Networking

Compute & Networking - Segment Deep Dive

Prepared: 2026-04-26 Sub-Agent: B Company: NVIDIA Corporation (NVDA) Segment: Compute & Networking TTM end: 2026-01-25 Forward horizon: 2031-01-25 Compounding span: exactly 5 calendar years Primary source of truth for NVIDIA historicals and management commentary: `data_package.md`

0. Segment Boundary + Taxonomy Note

Compute & Networking is NVIDIA's reportable segment for data center accelerated computing and networking platforms, AI solutions and software, and automotive platforms including autonomous and electric vehicle solutions. The data package states that NVIDIA has used the same two reportable segments, Compute & Networking and Graphics, across FY2024, FY2025, and FY2026, so there is no recent reportable-segment restatement boundary to normalize across [data_package §1, Taxonomy flags row 1]. The important boundary is different: NVIDIA also discloses end-market revenue lines such as Data Center, Automotive, Gaming, Professional Visualization, and OEM and Other, but those lines are not a clean mechanical bridge to reportable segments. The data package gives the concrete mismatch: FY2026 Data Center end-market revenue was $193.737B, while FY2026 Compute & Networking reportable-segment revenue was $193.479B [data_package §1, end-market table row Data Center; §1, segment table row Compute & Networking]. Therefore this appendix models reportable segment revenue and operating income from the official segment base, while using Data Center Compute, Networking, Automotive, and related end-market disclosures as mix context rather than as a perfect segment bridge.

1. What This Segment Actually Is

Compute & Networking is NVIDIA's AI infrastructure segment. Its economic center is accelerated data center computing: GPUs, CPU-GPU systems, rack-scale platforms, high-speed interconnect, systems software, libraries, and AI infrastructure software that allow cloud providers, model labs, enterprises, governments, and OEM partners to train and run large-scale AI workloads. The segment also includes networking platforms such as InfiniBand, Ethernet, NVLink, and related silicon/software for AI clusters, plus automotive platforms and autonomous/electric vehicle solutions [data_package §1, Reporting structure paragraph; §4, Dollar-disclosed product rows].

The customer base is concentrated in large buyers with very high technical requirements. Hyperscalers and frontier AI labs buy clusters measured in racks, not cards. Enterprise and sovereign customers often consume the same underlying hardware through cloud providers, OEM systems, or NVIDIA-integrated platforms. Automotive customers buy longer-cycle platform solutions such as DRIVE, autonomous vehicle compute, and related software, but that current revenue line is still small versus data center.

The pricing model is mainly product and platform sale, not pure subscription. NVIDIA sells GPUs, networking hardware, complete server and rack systems, and increasingly software and platform bundles. The profit model is therefore semiconductors plus full-stack system value capture: NVIDIA monetizes performance per watt, memory bandwidth, interconnect, software compatibility, and time-to-deployment. The software layer does not yet appear as a separately disclosed software revenue stream large enough to model independently, but it is central to pricing power because CUDA, libraries, inference engines, and system management tools reduce adoption friction and raise switching cost.

The go-to-market is hybrid. NVIDIA sells through cloud providers, server OEMs, original design manufacturers, distributors, and direct strategic relationships with the largest AI infrastructure builders. At the low end of the segment, automotive and edge platforms follow design-win and program timelines; at the high end, hyperscale data center revenue is driven by supply allocation, architecture transitions, power availability, and customer capital expenditure cycles. That is why quarterly revenue can step up sharply when supply comes online and why the relevant demand signal is not only end demand for AI applications, but also customer willingness to keep expanding infrastructure fleets.

2. Current Economics (FY2026 / TTM)

The TTM period equals FY2026 because the latest reported quarter is Q4 FY2026 ended 2026-01-25. Compute & Networking generated $193.479B of FY2026 revenue, representing 89.6% of NVIDIA's total revenue, and $130.141B of segment operating income, a 67.3% segment operating margin [data_package §1, FY2026 reportable segment table row Compute & Networking; §3, Compute & Networking table row FY2026]. This is not a normal semiconductor segment at current scale. The segment produced more operating income than most large technology companies produce in revenue, and it did so with a

segment margin that remained above 67% despite Blackwell system ramp costs and export-control charges.

The growth base is extraordinary but discontinuous. Compute & Networking revenue rose from $11.046B in FY2022 to $193.479B in FY2026, and operating income rose from $4.598B to $130.141B over the same period [data_package §3, Compute & Networking table rows Revenue and Operating income]. The data package calls the segment the controlling profit pool, with FY2026 revenue 17.5x FY2022 revenue and FY2026 operating income 28.3x FY2022 operating income [data_package §3, Compute & Networking commentary]. The correct interpretation is not that this segment can keep compounding at the last three-year rate indefinitely; it is that the AI infrastructure adoption curve created a new revenue base, and the forward question is how fast that base decelerates.

The most useful product mix disclosures are the Data Center sub-lines. For FY2026, Data Center end-market revenue was $193.737B, of which Data Center Compute was $162.361B and Data Center Networking was $31.376B [data_package §1, Taxonomy flags row 2; §4, Product rows Data Center Compute and Networking]. Data Center Compute grew from $38.950B in FY2024 to $102.196B in FY2025 and $162.361B in FY2026, while Networking grew from $8.575B to $12.990B and then $31.376B [data_package §4, Product rows Compute and Networking]. Automotive revenue was $2.349B in FY2026, up from $1.694B in FY2025 and $1.091B in FY2024 [data_package §4, Product row Automotive]. These lines help frame demand mix, but the official segment base remains $193.479B.

Recent quarterly economics confirm that FY2026 was not simply a first-half spike. Compute & Networking revenue rose from $39.589B in Q1 FY2026 to $41.331B in Q2, $50.908B in Q3, and $61.651B in Q4. Segment operating margin improved from 55.7% in Q1 to 68.6% in Q2, 70.2% in Q3, and 71.4% in Q4 [data_package §5, Reportable segments quarterly table]. The Q1 margin trough matters because it included H20-related disruption; the exit-rate margin matters because it shows that the segment can still earn low-70s segment margins at a $246B annualized Q4 revenue run-rate before unallocated corporate bridge items.

Capital intensity at the company level remains light relative to revenue because NVIDIA is fabless. FY2026 capital expenditure was $6.042B, only 2.8% of consolidated revenue, while R&D was $18.497B, 8.6% of revenue [data_package §7, Capital allocation table row FY2026]. This does not mean the ecosystem is capital-light. Customers carry the heavy data center land, power, cooling, and building costs; NVIDIA carries architecture, software, supply-chain commitments, packaging, and ecosystem investment risk. For segment margin, the main investment burden is product cadence, platform engineering, supply commitments, and customer support rather than direct data center capex.

[DATA GAP: NVIDIA does not disclose standalone annual revenue for Blackwell, Hopper, Rubin, CUDA software, DGX Spark/Station, robotics, or DRIVE software inside Compute & Networking. The appendix therefore does not create separate product-dollar histories for those items.]

3. Unit Economics

Compute & Networking turns revenue into profit through scarcity, architectural performance, and platform integration. The highest-margin dollars are those where NVIDIA captures value for system-level performance and developer compatibility rather than only chip cost. Gross margin is protected by leading accelerator performance, memory and packaging access, networking integration, CUDA and inference software, and the customer's cost of delay. A buyer building a frontier training cluster or an inference fleet cares about total cost per trained token, time to deployment, power per unit of compute, failure rates, and software readiness. If NVIDIA improves those variables, the customer can accept a high hardware price because the alternative is a delayed or underperforming AI service.

Segment operating income is reported after segment-specific costs but before certain corporate-level unallocated items. The data package shows FY2026 segment operating income of $130.141B on $193.479B of revenue, while total segment operating income of $139.297B bridged to consolidated operating income after stock-based compensation, unallocated operating expenses, and acquisition-related/other costs [data_package §1, segment table; quant_foundation P4 bridge]. In other words, segment margin is the right measure for segment economics, but consolidated modeling must still account for the corporate bridge elsewhere.

The main cost drivers are advanced wafer cost, CoWoS and advanced packaging, HBM memory, networking components, system assembly, warranty/support, and engineering expense. Blackwell and future rack-scale systems move NVIDIA up the system stack. That can increase absolute revenue per deployment and deepen lock-in, but it can also lower gross margin if more pass-through component content is included in the system sale or if ramp yields are less mature. The FY2026 margin path shows both sides: full-year segment margin fell from 71.3% in FY2025 to 67.3% in FY2026 [data_package §3, Compute & Networking operating margin row], yet Q4 segment margin recovered to 71.4% [data_package §5, reportable segment

quarterly table].

There are four mechanisms that can break margins down. First, hyperscale customers may use volume commitments and credible internal ASIC roadmaps to negotiate lower accelerator pricing. Second, product mix may shift toward full racks, networking, and lower-margin system content. Third, export controls can create inventory charges and stranded commitments, as happened with H20 in FY2026 [data_package §9, risk events row 2025-04-09]. Fourth, if AI infrastructure demand becomes cyclical, NVIDIA may be forced to absorb underutilized supply commitments or clear inventory.

There are also four mechanisms that can hold margins up. First, power constraints raise the value of performance per watt; management explicitly said every data center is power-constrained [data_package §6, Compute & Networking table row Q4 FY2026 power constraint]. Second, networking has become part of the performance architecture, not an afterthought, which increases the attach rate for proprietary and semi-proprietary fabric. Third, software compatibility makes platform substitution slower than benchmark comparisons imply. Fourth, annual architecture cadence can reset the performance frontier faster than customers can build and qualify alternatives.

My margin cases therefore treat FY2026's 67.3% segment margin as a peak-ish but not one-off anchor. The bear case assumes pricing and mix pressure take the segment down to 52.0% by TTM+5. The base case assumes the segment settles at 62.0%, below the current full-year margin but still high because software, networking, and supply-chain control remain valuable. The bull case assumes 66.0%, slightly below FY2026 but close to current economics because capacity remains tight and NVIDIA preserves platform pricing.

4. Competitive Dynamics

The direct competitors are not one group. In accelerator compute, AMD is the cleanest merchant GPU competitor. AMD reported FY2025 Data Center segment revenue of $16.6B, up 32% year over year, driven by EPYC CPUs and AMD Instinct GPU shipments; Q4 2025 Data Center revenue was $5.4B, up 39% year over year (AMD FY2025 results: https://www.amd.com/en/newsroom/press-releases/2026-2-3-amd-reports-fourth-quarter-and-full-year-2025-fina.html). Against NVIDIA's FY2026 Data Center end-market revenue of $193.737B, AMD remains much smaller in data center revenue scale, and AMD's Data Center operating income of $3.603B in 2025 is far below NVIDIA's Compute & Networking segment operating income of $130.141B [data_package §1, segment table; AMD FY2025 release]. AMD matters because customers want a second source and because ROCm can improve, but the current scale gap means AMD is a pressure valve, not yet a profit-pool equalizer.

Custom silicon is the more strategic threat. Broadcom said Q4 FY2025 revenue growth was driven primarily by AI semiconductor revenue up 74% year over year, and it guided Q1 FY2026 AI semiconductor revenue to double year over year to $8.2B, driven by custom AI accelerators and Ethernet AI switches (Broadcom FY2025 Q4 release: https://investors.broadcom.com/node/63716/pdf). This matters because the largest hyperscalers can amortize ASIC development over massive internal workloads. A successful internal accelerator does not need to beat NVIDIA across every workload; it only needs to be good enough for a large, stable slice of inference or training to reduce the marginal NVIDIA attach rate.

Networking competition is broader. Broadcom competes in Ethernet switch silicon and custom AI networking, Marvell competes in data center connectivity and custom silicon, and cloud providers can mix merchant Ethernet, custom fabrics, and NVIDIA networking. Marvell's Q4 FY2026 release shows Data Center revenue of $1.651B, up 21% year over year, with Data Center at 74% of quarterly revenue; it also notes data center applications across AI systems, Ethernet switching, AI servers, storage, and interconnect (Marvell FY2026 Q4 release: https://investor.marvell.com/news-events/press-releases/detail/1011/marvell-technology-inc-reports-fourth-quarter-and-fiscal-year-2026-financial-results). This supports the view that AI networking is a distinct growth pool, not merely a GPU accessory.

Customer power is the other side of the competitive map. Microsoft framed GPU capex as a long-term allocation decision across Azure, M365 Copilot, GitHub Copilot, R&D, and AI talent, and said Azure guidance can be interpreted as an allocated-capacity guide (Microsoft FY2026 Q2 call: https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q2). Meta said it expected 2025 capex including finance leases of $66B-$72B and another year of significant capex dollar growth in 2026 to bring AI capacity online (Meta Q2 2025 prepared remarks: https://s21.q4cdn.com/399680738/files/doc_downloads/2025/META-Q2-2025-Prepared-Remarks.pdf). These customer comments support demand durability, but they also underline buyer concentration: the same customers that create NVIDIA's growth have the financial motive and technical scale to sponsor substitutes.

Pricing power currently belongs to NVIDIA because the binding constraint is deployed, usable AI compute rather than list-price silicon. The share context is unusually one-sided: NVIDIA's own Data Center revenue is an order of magnitude larger than AMD's whole Data Center segment, while Broadcom and Marvell are growing around specific custom and networking pools. However, pricing power is not permanent just because the current share is high. The risk is that inference workloads become more standardized, models become more efficient, and hyperscaler internal accelerators become good enough for a rising share of tokens. The base case assumes substitution takes share at the margin but does not collapse NVIDIA's platform advantage. The bear case assumes pricing and mix pressure arrive faster.

5. Moat at the Segment Level

The segment moat has four layers. The first layer is hardware performance and cadence. NVIDIA is not only selling a fast chip; it is selling a product road map that customers can plan around. When architectures move from Hopper to Blackwell to Rubin, customers can underwrite data center builds around expected performance-per-watt gains and software continuity. This cadence is difficult to replicate because it requires advanced-node design, packaging access, HBM supply coordination, board/system engineering, and software readiness.

The second layer is the software ecosystem. CUDA, libraries, compilers, inference engines, and developer tooling reduce switching friction inside NVIDIA's platform and increase switching friction away from it. This matters most at large customers where developer time, model reliability, and time-to-production are more expensive than headline chip price. A competing accelerator can win on a narrow benchmark and still lose if it causes deployment delay, retraining, operational instability, or software rewrite burden.

The third layer is networking and systems integration. The FY2026 Data Center Networking line reached $31.376B [data_package §4, Product row Networking], and management said networking exceeded $31B for FY2026 [data_package §6, Compute & Networking quote row Q4 networking scale]. This makes the moat less dependent on a single GPU SKU. At cluster scale, fabric, interconnect, failure domains, and software orchestration determine realized throughput. If NVIDIA can sell the whole AI factory architecture, it can defend value capture even as merchant GPU alternatives improve.

The fourth layer is ecosystem gravity. Cloud providers, OEMs, model labs, sovereign AI programs, and enterprise software vendors all optimize around the platform with the largest installed base and highest near-term capacity. This can create a reinforcing loop: more developers optimize for NVIDIA, which increases customer confidence, which increases deployments, which funds the next architecture and software release. The loop is powerful, but not invulnerable. Hyperscalers have the money and workload visibility to weaken it in selected internal workloads.

Durability over 5-10 years is high but not absolute. The strongest moat is in frontier training, fast-moving inference, complex multimodal workloads, and customers that need maximum reliability quickly. The weakest point is mature, high-volume inference where models are stable, software stacks are optimized, and custom silicon can be amortized. The base case assumes NVIDIA remains the default for leading-edge and mixed workloads through 2031, but loses some marginal share and pricing in standardized inference. The bear case assumes the standardized inference pool grows faster than the leading-edge pool and that customers use it to pressure margins.

6. Historical Growth Rates - Reproduced From Data Package

The table below reproduces historical revenue from the data package. Reportable segment rows come from QF P2/SEC segment tables as reproduced in data_package §3. Product and end-market rows come from data_package §4 and are not a clean segment bridge. No product dollar rows are back-solved.

Product / sub-segment	FY2023 rev	FY2024 rev	FY2025 rev	FY2026 / TTM rev	Historical CAGR	Most recent YoY	Data package cite
Compute & Networking reportable segment	$15.068B [§3 C&N revenue row]	$47.405B [§3 C&N revenue row]	$116.193B [§3 C&N revenue row]	$193.479B [§3 C&N revenue row]	134.2% FY2023-FY2026	66.5% [§3 C&N YoY row]	data_package §3, Compute & Networking table rows Revenue and YoY growth
Data Center end market	n.d. [DATA GAP: FY2023 not disclosed in data package]	$47.525B [§4 Data Center row]	$115.186B [§4 Data Center row]	$193.737B [§4 Data Center row]	101.9% FY2024-FY2026, not 3-year	68.2% calculated from disclosed FY2025/FY2026	data_package §4, Product row Data Center
Data Center Compute sub-line	n.d. [DATA GAP: FY2023 not disclosed in package]	$38.950B [§4 Compute row]	$102.196B [§4 Compute row]	$162.361B [§4 Compute row]	104.2% FY2024-FY2026, not 3-year	58.9% calculated from disclosed FY2025/FY2026	data_package §4, Product row Compute

Product / sub-segment	FY2023 rev	FY2024 rev	FY2025 rev	FY2026 / TTM rev	Historical CAGR	Most recent YoY	Data package cite
Data Center Networking sub-line	n.d. [DATA GAP: FY2023 not disclosed in package]	$8.575B [§4 Networking row]	$12.990B [§4 Networking row]	$31.376B [§4 Networking row]	91.3% FY2024-FY2026, not 3-year	141.5% calculated from disclosed FY2025/FY2026	data_package §4, Product row Networking
Automotive end market	n.d. [DATA GAP: FY2023 not disclosed in data package]	$1.091B [§4 Automotive row]	$1.694B [§4 Automotive row]	$2.349B [§4 Automotive row]	46.7% FY2024-FY2026, not 3-year	38.7% calculated from disclosed FY2025/FY2026	data_package §4, Product row Automotive

Internal arithmetic checks: Compute & Networking $116.193B x 1.665 = $193.45B, which rounds to the disclosed $193.479B. Data Center $115.186B x 1.682 = $193.75B, which rounds to $193.737B. Data Center Compute $102.196B x 1.589 = $162.37B, which rounds to $162.361B. Data Center Networking $12.990B x 2.415 = $31.38B, which rounds to $31.376B. Automotive $1.694B x 1.387 = $2.349B. All visible FY2026 rows tie within 2% to the data package disclosures.

For modeling, I use the official FY2026 Compute & Networking segment revenue of $193.479B as the total base. The product allocation is a model allocation, not a disclosed product bridge: Networking is set equal to the disclosed FY2026 Data Center Networking line of $31.376B, Automotive/edge/other is set equal to the disclosed Automotive line of $2.349B, and AI compute/platform is the residual $159.754B needed to tie to official segment revenue. This residual is $2.607B below the disclosed Data Center Compute end-market line because the end-market table does not mechanically bridge to reportable segments.

7. Management Guidance - Quoted From Data Package Section 6

The most important forward guide is management's Q4 FY2026 comment that NVIDIA expected sequential revenue growth throughout calendar 2026 [data_package §6, Compute & Networking row Q4 FY2026 sequential growth]. This anchors FY2027 above the Q4 run-rate, but it does not by itself justify extrapolating FY2026 growth rates for five years. The model treats it as strong near-term demand evidence and then decelerates growth as the base gets larger.

Management also said every data center is power-constrained [data_package §6, Compute & Networking row Q4 FY2026 power constraint]. The interpretation is that performance per watt is a pricing variable. In a power-constrained data center, the buyer is not choosing the cheapest chip; the buyer is choosing the architecture that maximizes revenue-generating compute per scarce watt and per scarce rack.

The networking commentary is unusually concrete. Management said networking generated $11B in Q4 FY2026 and exceeded $31B for FY2026 [data_package §6, Compute & Networking rows Q4 networking revenue and FY2026 networking scale]. This supports a separate networking forecast line. Networking's FY2026 scale is already larger than all of NVIDIA's consolidated revenue in FY2022.

Management repeatedly excluded China data center compute revenue from guidance. The Q4 FY2026 guide said NVIDIA was not assuming any data center compute revenue from China, and Q2 FY2026 commentary said the Q3 outlook did not include H20 shipments to China customers [data_package §6, Compute & Networking rows China/export controls and Q2 export-control conservatism]. Therefore the base case does not require China recovery. China is an upside swing factor if restrictions ease and a downside factor if restrictions expand to more products or regions.

Jensen Huang's demand framing was that compute equals revenues [data_package §6, Compute & Networking row CEO ROI framing]. This is the management thesis in one sentence: customers will keep buying because compute directly monetizes AI services. The red-team issue is whether customers' incremental revenue and productivity gains keep pace with their capex, depreciation, and power costs.

Architecture transition commentary also matters. Q1 FY2026 commentary said Blackwell contributed nearly 70% of data center compute revenue [data_package §6, Compute & Networking row Blackwell transition]. Blackwell became the center of gravity early in FY2026, which means Hopper-era comparisons are stale. The model assumes architecture transitions continue to create upgrade demand, but it does not create a separate Blackwell or Rubin revenue line because the data package does not disclose standalone product revenue.

Finally, management noted that Spectrum-X posted strong sequential and year-on-year growth [data_package §6, Compute & Networking row Ethernet networking]. This supports the thesis that NVIDIA can defend AI networking beyond InfiniBand-only deployments and participate in Ethernet scale-out architectures.

[DATA GAP: The data package has management quotes for sequential growth, networking, China exclusion, power constraints, and Blackwell mix, but no explicit five-year revenue CAGR guide, no product-level operating margin guide, and no

separate software revenue guide.]

8. Forward Growth Assumptions - Explicit Per Product

All forward CAGRs compound from the TTM base ending 2026-01-25 to 2031-01-25, exactly five calendar years. The model lines are designed to sum to official reportable segment revenue. They are not presented as independently disclosed product revenue.

Base Case Growth

Product / model line	TTM base	Historical CAGR	Latest YoY	Mgmt FY+1 guide	Base CAGR TTM to TTM+5y	TTM+5y revenue	Justification
AI compute/platform allocation	$159.754B	Data Center Compute 104.2% FY2024-FY2026 [§4 Compute row]	58.9%	Sequential CY2026 growth; no China compute assumed [§6]	23.0%	$449.8B	Growth decelerates sharply from FY2026 but remains high because demand is still supply/power constrained and Blackwell/Rubin architecture cycles should expand deployed inference and training capacity.
Data Center Networking	$31.376B	91.3% FY2024-FY2026 [§4 Networking row]	141.5%	$11B Q4 and over $31B FY2026 [§6]	25.0%	$95.8B	Networking grows faster than compute in the near term as AI clusters require more fabric, but decelerates as Ethernet/custom alternatives scale.
Automotive / edge / physical AI	$2.349B	46.7% FY2024-FY2026 [§4 Automotive row]	38.7%	Physical AI commentary but Q4 Auto only +6% YoY [§6]	28.0%	$8.1B	Small base and robotics/AV optionality support growth, but disclosed Automotive exit growth slowed, so this remains a small contributor.
Segment total	$193.479B	C&N 134.2% FY2023-FY2026 [§3]	66.5%	Sequential CY2026 growth [§6]	23.4% implied	$553.6B	Product rollup ties to official segment base; no end-market bridge is forced.

Base-case year-by-year deceleration: the compute allocation is modeled as roughly +34%, +27%, +22%, +18%, and +15% across the five annual steps, which compounds close to the 23% CAGR. Networking is modeled as roughly +40%, +30%, +24%, +18%, and +14%, reflecting a catch-up phase after the FY2026 networking inflection. Automotive/edge is modeled as +30%, +30%, +28%, +27%, and +25% from a small base. The deceleration is less severe than a collapse but much more severe than the last three years. This is defensible because comparable customer capex programs are still expanding, while competitors such as AMD, Broadcom, and Marvell are already scaling alternatives.

Bear Case Growth

Product / model line	TTM base	Historical CAGR	Latest YoY	Bear CAGR TTM to TTM+5y	TTM+5y revenue	What changes versus base
AI compute/platform allocation	$159.754B	Data Center Compute 104.2% FY2024-FY2026	58.9%	12.0%	$281.5B	Hyperscaler ASICs and AMD GPUs take more standardized inference share, customer capex ROI scrutiny increases, and pricing/mix pressure starts before Rubin fully offsets it.
Data Center Networking	$31.376B	91.3% FY2024-FY2026	141.5%	15.0%	$63.1B	Ethernet and custom fabrics capture more cluster deployments, reducing NVIDIA attach and pricing despite continued traffic growth.

Product / model line	TTM base	Historical CAGR	Latest YoY	Bear CAGR TTM to TTM+5y	TTM+5y revenue	What changes versus base
Automotive / edge / physical AI	$2.349B	46.7% FY2024-FY2026	38.7%	12.0%	$4.1B	Automotive design wins convert slower, robotaxi/physical AI revenue stays strategic rather than material, and Q4 FY2026's +6% YoY becomes more representative.
Segment total	$193.479B	C&N 134.2% FY2023-FY2026	66.5%	12.5% implied	$348.8B	Revenue still grows, but the AI infrastructure cycle normalizes and substitutes cap NVIDIA's share of incremental spend.

The bear case is not a demand recession case; it is a normalization and share-pressure case. It assumes customers keep building AI infrastructure but shift more marginal workloads to internal or alternative accelerators. The strongest support for this scenario is Broadcom's custom AI accelerator momentum and the strategic incentive of hyperscalers to control their own unit economics.

Bull Case Growth

Product / model line	TTM base	Historical CAGR	Latest YoY	Bull CAGR TTM to TTM+5y	TTM+5y revenue	What changes versus base
AI compute/platform allocation	$159.754B	Data Center Compute 104.2% FY2024-FY2026	58.9%	30.0%	$593.2B	Inference demand grows faster than efficiency gains, Rubin and successors preserve a clear performance-per-watt gap, and customers prioritize time-to-capacity over second-source economics.
Data Center Networking	$31.376B	91.3% FY2024-FY2026	141.5%	33.0%	$130.6B	NVIDIA captures a larger share of AI fabric as networking becomes inseparable from system performance; Spectrum-X and NVLink-style architectures expand attach.
Automotive / edge / physical AI	$2.349B	46.7% FY2024-FY2026	38.7%	35.0%	$10.5B	Robotics and AV platform revenue finally converts, but even the bull case keeps it modest relative to data center.
Segment total	$193.479B	C&N 134.2% FY2023-FY2026	66.5%	30.6% implied	$734.3B	NVIDIA remains the default AI factory supplier and the total AI compute pool expands faster than substitution.

The bull case requires more than extrapolation. It requires customer AI revenue and productivity to validate capex, continued power scarcity that favors best-in-class performance per watt, and limited success for internal ASICs outside narrow workloads. It also assumes no material tightening of export restrictions beyond the China exclusions already embedded in the base.

9. Forward Operating Margin Assumptions

Current Compute & Networking segment operating margin is 67.3% for FY2026 / TTM [data_package §3, Compute & Networking operating margin row]. Quarterly margin exited at 71.4% in Q4 FY2026 [data_package §5, reportable segment quarterly table row Q4]. Product-level operating income is not disclosed for Data Center Compute, Networking, Automotive, Blackwell, CUDA, or software. The table below therefore uses explicit model margins by line and computes the segment margin as a revenue-weighted average.

Case	AI compute/platform margin	Networking margin	Automotive/edge margin	Segment TTM+5y margin	Mechanism
Bear	51.1%	58.0%	25.0%	52.0%	Pricing compression, higher system pass-through content, ASIC substitution in inference, export-control/inventory risk, and lower utilization of supply commitments.
Base	61.1%	68.0%	39.0%	62.0%	Margin normalizes below FY2026 as systems mix and buyer leverage rise, but software, performance per watt, and networking attach preserve a premium structure.
Bull	65.3%	71.0%	45.0%	66.0%	Scarcity remains acute, networking attach expands, software/platform value capture increases, and architecture cadence offsets system-content dilution.

The base-case margin decline is greater than 100 bps from current and therefore needs a mechanism. The mechanism is not a generic maturity haircut. It is the combination of richer rack-scale systems mix, customer concentration, more competitive accelerator alternatives, and the likelihood that a larger installed base includes more price-sensitive inference workloads. The offset is that power scarcity and deployment urgency keep NVIDIA from becoming a commodity component supplier.

The bear-case margin is intentionally severe relative to recent segment history. A move from 67.3% to 52.0% would require both price pressure and mix pressure, not merely revenue deceleration. This case becomes plausible if hyperscaler internal accelerators handle a large share of standardized inference, merchant alternatives reduce scarcity, and NVIDIA must price more aggressively to keep platform share.

The bull-case margin remains below the Q4 FY2026 exit margin because a five-year model should not assume a permanent 71% segment margin at more than $700B of revenue. Still, 66.0% is defensible if NVIDIA continues to sell an integrated AI factory stack rather than standalone chips and if networking/software attach rises.

[DATA GAP: NVIDIA does not disclose product-level operating margins for Data Center Compute, Networking, Automotive, CUDA/software, or Blackwell/Rubin systems. Product margin rows in this section are analyst model assumptions anchored to the disclosed segment margin, not company disclosures.]

10. Segment Rollup - Bottoms-Up Tie

The rollup below reconstructs the segment totals from the product/model lines. The TTM base is the official reportable segment. The allocation is used only for forward modeling and is designed to avoid forcing the end-market table into a segment bridge.

Case	AI compute/platform revenue	Networking revenue	Automotive/edge revenue	Segment TTM+5y Revenue	Segment TTM+5y Op Margin	Segment TTM+5y Op Income	Probability
Bear	$281.5B	$63.1B	$4.1B	$348.8B	52.0%	$181.4B	25%
Base	$449.8B	$95.8B	$8.1B	$553.6B	62.0%	$343.2B	50%
Bull	$593.2B	$130.6B	$10.5B	$734.3B	66.0%	$484.6B	25%

Arithmetic tie: Bear revenue equals $281.5B + $63.1B + $4.1B = $348.7B, rounding to $348.8B; operating income equals $348.8B x 52.0% = $181.4B. Base revenue equals $449.8B + $95.8B + $8.1B = $553.7B, rounding to $553.6B due unrounded inputs; operating income equals $553.6B x 62.0% = $343.2B. Bull revenue equals $593.2B + $130.6B + $10.5B = $734.3B; operating income equals $734.3B x 66.0% = $484.6B. The small apparent summation differences are rounding only.

The expected-value case using the stated probabilities is $547.6B of TTM+5 revenue and $338.1B of TTM+5 segment operating income. That expected value is close to the base case because the probabilities are symmetric and the revenue/margin upside and downside are intentionally wide. The final synthesis should not simply import this expected value as consolidated revenue; it must combine this segment with Graphics, unallocated corporate bridge items, tax, share count, and capital allocation.

11. Key Risks and Wildcards

Bear risk - hyperscaler ASIC substitution. The largest customers have the workloads, balance sheets, and engineering teams to internalize more accelerator economics. Custom silicon does not need to win every workload to pressure NVIDIA; it only needs to take a rising share of stable inference and give customers a credible negotiation alternative.

Bear risk - customer capex ROI. Microsoft and Meta commentary supports ongoing AI infrastructure buildout, but it also shows that investors and operators are scrutinizing the relationship between capex and revenue. If AI revenue does not scale with deployed compute, customer orders can slow before NVIDIA's own reported demand indicators show stress.

Bear risk - export controls and geopolitics. The data package records H20 disruption, inventory/purchase-obligation charges, and repeated exclusion of China data center compute from guidance [data_package §9, risk events; §6 China rows]. The base case already excludes China recovery, but broader restrictions could hit more products or regions.

Base risk - margin dilution from systems mix. Full-stack systems deepen the moat, but they may include more lower-margin content and more operational complexity. A larger Blackwell/Rubin rack-scale business can grow revenue while reducing margin versus pure accelerator scarcity economics.

Base opportunity - networking attach. The $31.376B FY2026 networking line and $11B Q4 networking commentary show that networking can become a second profit engine inside AI factories [data_package §4 Networking row; §6 networking rows]. If NVIDIA keeps fabric performance tied to GPU cluster performance, networking can offset some accelerator pricing pressure.

Bull wildcard - inference demand outruns efficiency. If AI usage grows faster than model efficiency gains and hardware utilization improvements, customers may need far more deployed compute than current capex skepticism assumes. In that scenario, the segment can compound from the FY2026 base at high rates without relying on China recovery.

Bull wildcard - physical AI becomes real revenue. Management's physical AI and robotics commentary is strategically important, but current Automotive revenue is only $2.349B [data_package §4 Automotive row]. If robotaxi, robotics, and industrial AI platforms convert into material deployments, this line can grow faster than modeled, although it is unlikely to rival data center by 2031.

Final summary table:

Case	Segment TTM+5y Revenue	Segment TTM+5y Op Margin	Segment TTM+5y Op Income	Probability
Bear	$348.8B	52.0%	$181.4B	25%
Base	$553.6B	62.0%	$343.2B	50%
Bull	$734.3B	66.0%	$484.6B	25%

Appendix B2 - Graphics

NVIDIA Corporation (NVDA) - Appendix B: Graphics Segment Deep Dive

Prepared: 2026-04-26 Segment: Graphics Output path: `C:\Users\Admin\OneDrive\GPT\ER 2.0\NVIDIA Corporation\appendix_B_graphics.md` Controlling source for NVIDIA historicals and management commentary: `C:\Users\Admin\OneDrive\GPT\ER 2.0\NVIDIA Corporation\data_package.md` TTM anchor: FY2026 / TTM ended 2026-01-25 Forward horizon: 2031-01-25, exactly five calendar years after the TTM end

0. Segment boundary and taxonomy note

Graphics is one of NVIDIA's two current reportable segments, alongside Compute & Networking. The data package states that no reportable segment change was identified across the last three fiscal years and that FY2024, FY2025, and FY2026 10-K segment notes use the same two reportable segments [data_package Section 1, taxonomy flags]. The Graphics segment boundary used here is therefore stable for reportable-segment dollar comparisons. The important caution is that NVIDIA's end-market disclosure lines - Gaming, Professional Visualization, Automotive, OEM and Other, and Data Center - are not a clean mechanical bridge to reportable segments. The data package explicitly warns that downstream work must model reportable segment revenue and operating income from QF P2/P4, while using Gaming, Professional Visualization, and OEM and Other only as demand context [data_package Section 1, Reporting structure and taxonomy; data_package Section 10, Open data gaps]. This appendix follows that rule: the official Graphics TTM base is $22.459B of revenue, $9.156B of operating income, and 40.8% operating margin; the product/end-market lines help explain demand, but a residual bridge is shown and flagged rather than hidden.

1. What this segment actually is

Graphics is NVIDIA's PC, workstation, and visualization GPU business. In operating terms, it includes GeForce GPUs for gaming desktops, gaming notebooks, and consumer PCs, plus NVIDIA RTX and Quadro professional visualization products used in workstations, graphics rendering, design, engineering, media production, and AI workstation workflows [data_package Section 1, Reporting structure and taxonomy]. The customer set is different from Data Center even though the silicon road map shares architecture, driver, CUDA, and software ecosystem assets. Gaming demand flows through add-in-board partners, OEM notebook and desktop vendors, retailers, distributors, and system integrators. Professional Visualization demand flows through workstation OEMs, enterprise channel partners, professional software ecosystems, design studios, industrial users, and emerging AI-workstation buyers.

The segment's pricing model is mainly product sale economics, not subscription software economics. NVIDIA sells GPUs, boards, modules, and platform products into channels that then set retail or system prices. The most profitable demand is not merely unit volume; it is mix toward higher-end GPUs, workstation-class products, memory-rich configurations, and professional-certified platforms. GeForce revenue is cyclical because gamers, enthusiasts, and OEMs can defer upgrades when pricing is high, software does not require a new card, or channel inventory is heavy. Professional Visualization is smaller but can carry more strategic value because the use case is tied to paid work: rendering, CAD, simulation, digital twins, AI-assisted content creation, and local inferencing. Those users are less price-sensitive than a mid-range gaming buyer when the GPU shortens time-to-output.

Go-to-market is hybrid. NVIDIA's brand and architecture do the demand creation; add-in-board partners and OEMs do much of the physical distribution; independent software vendors and game developers create workload pull; and NVIDIA's driver stack, DLSS, CUDA, RTX, Tensor cores, and professional certifications reinforce lock-in. The segment is therefore best understood as a premium GPU platform with cyclical hardware demand, not as a commodity component vendor. The main strategic question is whether AI PC and AI workstation use cases turn a historically cyclical Graphics business into a more durable replacement and attach-rate cycle, or whether FY2026 simply represents a strong Blackwell/RTX recovery year after a weaker base.

2. Current economics

The current controlling base is the TTM period ended 2026-01-25. Because the latest reported quarter is NVIDIA's fiscal year-end, TTM equals FY2026. Graphics generated $22.459B of reportable segment revenue, $9.156B of segment operating income, and a 40.8% operating margin [data_package Section 3, Graphics segment historicals; data_package Section 5,

TTM / FY2026 reportable segments]. Graphics represented 10.4% of NVIDIA's total segment revenue and roughly 7% of total segment operating income in FY2026 [data_package Section 1, FY2026 reportable segment table; data_package Section 3, Graphics commentary]. It is economically important, but it is not the primary valuation driver in the current NVIDIA model. Compute & Networking is the controlling profit pool.

The segment's FY2026 recovery was strong. Reportable Graphics revenue rose from $14.304B in FY2025 to $22.459B in FY2026, up 57.0%, while operating income rose from $5.085B to $9.156B [data_package Section 3, Graphics historicals]. Operating margin recovered to 40.8% from 35.5%, but it remained below the FY2022 peak of 53.5% [data_package Section 3, Graphics historicals]. The margin record matters because it shows both the upside and the cyclicality: NVIDIA has demonstrated much higher Graphics profitability, but the segment can also compress when gaming channel inventory, mix, pricing, or supply allocation turns against it.

At the demand-line level, the data package gives three relevant disclosed rows. Gaming revenue was $16.042B in FY2026, up from $11.350B in FY2025 and $10.447B in FY2024 [data_package Section 4, Product/end-market row: Gaming]. Professional Visualization revenue was $3.191B in FY2026, up from $1.878B in FY2025 and $1.553B in FY2024 [data_package Section 4, Product/end-market row: Professional Visualization]. OEM and Other revenue was $0.619B in FY2026, up from $0.389B in FY2025 and $0.306B in FY2024 [data_package Section 4, Product/end-market row: OEM and Other]. These rows total $19.852B, which is $2.607B below Graphics reportable segment revenue. That gap is not an error to smooth away; it is the taxonomy warning in numeric form. The model therefore keeps a "Graphics segment taxonomy bridge" line equal to reportable Graphics revenue less these demand-line disclosures. [DATA GAP: the data package does not disclose the products or end-market attribution behind the $2.607B FY2026 bridge between Graphics reportable revenue and Gaming + Professional Visualization + OEM and Other.]

Quarterly momentum was also strong but uneven. Graphics reportable revenue progressed from $4.473B in Q1 FY2026 to $5.412B in Q2, $6.098B in Q3, and $6.476B in Q4; segment operating margin moved from 36.7% to 41.4%, 41.8%, and 42.1% [data_package Section 5, Reportable segment quarterly trend]. Gaming peaked at $4.287B in Q2 and $4.265B in Q3 before falling to $3.727B in Q4, while still growing 47% year over year [data_package Section 5, End-market quarterly trend; data_package Section 6, Graphics/Gaming commentary]. Professional Visualization accelerated sharply in Q4 to $1.321B, up 159% year over year and 74% sequentially [data_package Section 5, Q4 recency signals]. The current economics are therefore a mix of strong underlying demand, supply constraints, channel normalization, and a fast-growing but still smaller workstation AI line.

Capital intensity at the reportable segment level is not disclosed. NVIDIA's consolidated capex was $6.042B in FY2026, or 2.8% of revenue, and consolidated R&D was $18.497B, or 8.6% of revenue [data_package Section 7, Capital allocation history]. The Graphics segment consumes R&D, architecture, software, and foundry capacity, but it does not require the same direct data-center infrastructure capex that cloud operators incur. Segment capital intensity should be read as design, mask, software, driver, and supply-chain commitment intensity rather than property-and-equipment intensity.

3. Unit economics

Graphics revenue turns into profit through three layers: GPU architecture leverage, channel and product mix, and software ecosystem pull. The same architectural investments that support data center accelerators also support GeForce and RTX workstation products. That shared technology base is a meaningful economic advantage: the segment benefits from a company-wide R&D engine that it could not fund alone at the same pace, while its incremental revenue carries premium GPU margins when supply and channel conditions are favorable.

The historical margin table shows the leverage clearly. Graphics operating margin was 53.5% in FY2022, fell to 38.2% in FY2023 during the post-boom gaming correction, recovered to 43.3% in FY2024, fell again to 35.5% in FY2025, and improved to 40.8% in FY2026 [data_package Section 3, Graphics historicals]. That range is too wide to treat the segment as a stable 40% annuity. It behaves like a premium hardware franchise with meaningful fixed-cost leverage and meaningful channel-cycle risk. When high-end GPUs are supply constrained, pricing, mix, and inventory absorption lift margins. When demand normalizes or inventory gets ahead of sell-through, margin can compress quickly.

Gaming unit economics are driven by attach rate to new GPU generations, average selling price, memory cost, foundry and packaging cost, board partner economics, and game/software features that create upgrade urgency. GeForce has pricing power at the high end because enthusiast buyers, streamers, creators, and performance-sensitive gamers value frame rates, ray tracing, DLSS, latency, driver stability, and software compatibility. But the low-to-mid market is more elastic. A gamer can defer an upgrade, buy used, choose a console, or settle for prior-generation hardware. The base case therefore cannot extrapolate FY2026's 41% Gaming growth rate for five years.

Professional Visualization unit economics should be structurally better than consumer gaming because a workstation GPU is an input into paid output. Time saved in rendering, simulation, generative design, AI inference, or content production can justify higher price points. Professional buyers also value certified drivers, enterprise support, reliability, and ecosystem integration. The data package does not disclose product-level operating margin, so this appendix does not claim a standalone Pro Viz margin as fact. The forward margin cases below use estimated product-line margin assumptions only as a modeling device, with the reportable segment margin tying back to the official Graphics segment margin. [DATA GAP: NVIDIA does not disclose operating income or operating margin for Gaming, Professional Visualization, OEM and Other, or the residual bridge within Graphics.]

The main margin breakpoints are supply allocation, memory and component costs, and mix. Supply constraints can be positive for price and negative for volume. In the Q4 FY2026 management commentary, CFO Colette Kress said NVIDIA expected "supply constraints to be the headwind to Gaming" [data_package Section 6, Q4 FY2026 Graphics/Gaming commentary]. That matters: it suggests near-term demand was not the binding constraint, but it also means the segment may be competing internally for wafer, packaging, memory, and system allocation against much larger Data Center demand. If data center products offer better gross profit dollars per scarce unit of supply, Graphics may be supply rationed even when end demand is healthy. Conversely, if supply expands and gaming channel inventory is overbuilt, margins can deteriorate.

4. Competitive dynamics

Graphics competitors are narrower and clearer than the company's data center competitor set. In gaming discrete GPUs, the principal competitor is AMD Radeon; Intel Arc is present but not yet a comparable high-end ecosystem force. Consoles indirectly compete for gaming wallet share, but consoles also use AMD silicon and operate at different price/performance and upgrade-cycle points. In professional visualization, AMD competes through Radeon Pro and workstation GPUs; Intel has workstation ambitions; Apple competes for parts of the creator and workstation installed base through integrated Apple Silicon systems; and cloud GPU workstations can substitute for some local workstation purchases. For AI workstation use cases, the competition also includes the decision not to run local inference and to use cloud inference or data center GPUs instead.

The current market-share context supports NVIDIA's pricing power. Jon Peddie Research's Q1 2025 add-in-board release said PC-based graphics add-in-board shipments reached 9.2 million units and that NVIDIA's share increased during the quarter while AMD and Intel decreased; the same release also forecast a negative 2024-2028 AIB unit CAGR, which underscores that share and mix matter more than market-unit growth [Jon Peddie Research, June 4, 2025](https://www.jonpeddie.com/news/q125-pc-graphics-add-in-board-shipments-increased-8-5-from-last-quarter-due-to-nvidias-blackwell-ramping-up/). Later JPR-linked coverage placed NVIDIA above 90% of PC graphics-card unit share in Q4 2025, with AMD below 10% and Intel minimal [PCWorld, March 2026](https://www.pcworld.com/article/3079686/pc-graphics-cards-are-now-nearly-100-percent-nvidia.html). These are external market sources, not NVIDIA historicals, and they are used only for competitive context.

NVIDIA's competitive edge in gaming is no longer just raster performance. DLSS, frame generation, ray tracing support, Reflex, driver cadence, creator features, and installed developer support matter. Game developers optimize for the installed base, which reinforces buyer preference, which reinforces developer optimization. AMD can compete on price/performance and memory capacity, but it has struggled to dislodge NVIDIA at the high end where brand, software, content-creator workflows, and AI features matter. Intel can be relevant at entry price points, but it does not yet have the high-end brand or developer ecosystem to reset NVIDIA's pricing.

In Professional Visualization, the barriers are certification, reliability, software integration, and workflow switching costs. A studio or engineering team is not buying only a chip; it is buying confidence that drivers, plugins, rendering tools, CAD packages, AI toolchains, and support will work in production. NVIDIA RTX and Quadro have a long installed base in those workflows. That does not make the business immune to substitution. If cloud workstations become cheaper and easier, local workstation purchases may shift to cloud GPU rental. If Apple Silicon continues to improve for creators, some media workloads can move away from discrete workstation GPUs. If AMD prioritizes price/performance in workstations, it can pressure mid-range professional cards. But the premium AI workstation use case favors NVIDIA because CUDA, Tensor cores, and AI software integration already matter to the buyer.

The biggest competitive risk may be internal rather than external: Data Center competes for the same strategic attention and supply base. Graphics is economically attractive, but it is much smaller. If management must allocate scarce advanced packaging, high-end memory, or foundry capacity, the rational priority may be Data Center. That does not destroy Graphics; it caps near-term revenue and can create channel scarcity. It also raises the long-run question of whether GeForce and workstation products receive enough supply and product cadence to maintain dominance if data center demand remains

extraordinary.

5. Moat at the segment level

The Graphics segment has a real moat, but it is not the same moat as NVIDIA's data center moat. The segment moat is a layered installed-base and software moat. In gaming, the moat rests on performance leadership, developer optimization, driver reliability, brand, and feature ecosystems such as DLSS and ray tracing. In professional visualization, it rests on certified software support, workflow trust, AI workstation readiness, and CUDA/RTX adoption in creator and engineering workflows. In both categories, the moat is stronger at the premium end than at the entry level.

The moat is durable over five years if NVIDIA continues to convert architecture leadership into end-user-visible benefits. The key is that the buyer must experience something difficult for a competitor to match: materially better frame generation, ray tracing, AI upscaling, local model performance, rendering acceleration, or workflow compatibility. If the benefit is only raw frames per dollar, AMD can compete more aggressively and margins compress. If the benefit is an integrated software and hardware experience, NVIDIA can sustain premium pricing even in a mature unit market.

The segment's moat is weaker than the Data Center moat in one respect: consumer upgrade cycles are discretionary. A cloud customer buying AI capacity may have a revenue or productivity case for accelerating purchases. A gamer may simply wait. A professional workstation buyer sits in between: the purchase can be productivity driven, but budget cycles and software requirements still matter. Therefore the segment should not be valued as a high-certainty hypergrowth business just because FY2026 growth was 57%. It is better framed as a premium cyclical franchise with AI workstation optionality.

On the other hand, Graphics can become more strategically important if local AI inference becomes a broad PC/workstation requirement. A GPU that supports gaming plus local AI features, creator workflows, and productivity applications has a larger value proposition than a GPU bought only for games. The data package includes management commentary that "Demand for our AI workstations is strong" [data_package Section 6, Q1 FY2026 Graphics/Professional Visualization commentary]. That is not enough to forecast a standalone AI workstation supercycle, but it is enough to justify a Pro Viz growth premium over legacy workstation growth.

6. Historical growth rates - reproduced from the data package

All NVIDIA historical figures in this section are reproduced from the data package. The end-market lines are not mechanically mapped to reportable Graphics; they are shown because the data package identifies Gaming and Professional Visualization as Graphics demand drivers and OEM and Other as a residual disclosure line. The official reportable segment row is the controlling base for revenue, operating income, and margin. The product/end-market tables use the three disclosed fiscal-year points from FY2024-FY2026 because the data package provides dollar-disclosed product/end-market rows for those years; the CAGR column is a two-year FY2024-to-FY2026 period CAGR, not a three-year compounding window.

Product / sub-segment	FY2024 rev	FY2025 rev	FY2026 rev	FY2024-FY2026 period CAGR (2 yrs)	Most recent YoY	Data package cite
Graphics reportable segment	$13.517B	$14.304B	$22.459B	28.9% FY2024-FY2026 CAGR	57.0%	data_package Section 3, Graphics row: QF P2; SEC 10-K Note 16
Gaming end-market line	$10.447B	$11.350B	$16.042B	23.9% FY2024-FY2026 CAGR	41.3% FY2026 YoY	data_package Section 4, Product/end-market row: FMP product segmentation; SEC 10-K
Professional Visualization end-market line	$1.553B	$1.878B	$3.191B	43.3% FY2024-FY2026 CAGR	69.9% FY2026 YoY	data_package Section 4, Product/end-market row: FMP product segmentation; SEC 10-K
OEM and Other disclosure line	$0.306B	$0.389B	$0.619B	42.2% FY2024-FY2026 CAGR	59.1% FY2026 YoY	data_package Section 4, Product/end-market row: FMP product segmentation; SEC 10-K
Graphics segment taxonomy bridge	$1.211B	$0.687B	$2.607B	46.7% FY2024-FY2026 CAGR	279.5% FY2026 YoY	Calculated as Graphics reportable revenue less Gaming, Professional Visualization, and OEM and Other using data_package Sections 3 and 4; [DATA GAP: no disclosed product attribution]

Arithmetic checks:

Row	Tie-out
Graphics reportable segment	$14.304B x 1.570 = $22.458B, within rounding of $22.459B.
Gaming	$11.350B x 1.413 = $16.043B, within rounding of $16.042B.
Professional Visualization	$1.878B x 1.699 = $3.191B.
OEM and Other	$0.389B x 1.591 = $0.619B.
Graphics segment taxonomy bridge	$0.687B x 3.795 = $2.607B; this is a residual bridge, not a disclosed product row.

The reportable-segment five-year history shows the cycle. Graphics revenue was $15.868B in FY2022, fell to $11.906B in FY2023, recovered to $13.517B in FY2024, rose to $14.304B in FY2025, and jumped to $22.459B in FY2026 [data_package Section 3, Graphics historicals]. Operating income fell from $8.492B in FY2022 to $4.552B in FY2023, recovered to $5.846B in FY2024, slipped to $5.085B in FY2025, and reached $9.156B in FY2026 [data_package Section 3, Graphics historicals]. This is not a smooth compounding record. It is a cyclical recovery plus Blackwell/RTX demand, with Professional Visualization acceleration.

The quarterly data reinforces that point. Graphics reportable segment revenue increased every quarter through FY2026, from $4.473B in Q1 to $6.476B in Q4 [data_package Section 5, Reportable segment quarterly trend]. But Gaming fell sequentially in Q4 after peaking in Q2/Q3, while Professional Visualization surged in Q4 [data_package Section 5, Q4 recency signals]. The right base for forward modeling is therefore not the Q4 annualized run-rate alone. The required ER2 anchor is TTM/FY2026, so this appendix compounds from $22.459B over exactly five calendar years to 2031-01-25.

7. Management guidance - quoted from data package Section 6

The data package provides seven Graphics-relevant management comments. Each is listed below with interpretation. No transcript was re-fetched for this appendix.

Fiscal period and call date	Speaker	Verbatim quote from data package	Topic	Interpretation
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Gaming revenue of $3.7 billion increased 47% year-on-year"	Gaming demand	Gaming was still growing rapidly at the end of FY2026, but Q4 was below Q2/Q3, so the model should not annualize Q4 as a clean acceleration point.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"we expect supply constraints to be the headwind to Gaming"	Gaming supply	Near-term limitation is described as supply, not demand. This supports revenue durability but also means Data Center supply priority can cap Graphics growth.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"For Professional Visualization, it crossed the $1 billion mark for the first time"	Pro Viz inflection	Professional Visualization reached a new scale and deserves a separate growth line inside the Graphics demand model.
Q3 FY2026, 2025-11-19	Colette Kress, CFO	"Gaming revenue was $4.3 billion, up 30% year on year"	Gaming quarterly base	Q3 was a very strong Gaming quarter; it sets evidence for demand but also marks the FY2026 peak.
Q2 FY2026, 2025-08-27	Colette Kress, CFO	"Our gaming revenue was a record $4.3 billion, a 14% sequential increase and a 49% jump year on year"	Gaming momentum	The Q2 record shows that demand was broad before the Q4 channel reset.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"Gaming revenue was a record $3.8 billion"	Gaming recovery	The recovery was visible at the start of FY2026, so FY2026 was not only a Q4 event.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"Demand for our AI workstations is strong"	Professional Visualization	Workstation AI demand supports a higher Pro Viz growth assumption than legacy workstation cyclicality alone.

The guidance anchor is positive but not unlimited. Management did not provide a five-year Graphics revenue target, product-level margin guide, GeForce RTX generation revenue, Pro Viz operating profit, or a quantitative AI workstation TAM in the data package. [DATA GAP: the data package contains no long-term Graphics segment revenue target or product-level

margin guidance.] The forward cases must therefore decelerate from FY2026 growth and explicitly state what changes.

8. Forward growth assumptions - explicit, per product

The model starts with TTM/FY2026 base revenue and compounds for exactly five calendar years to 2031-01-25. Because end-market lines do not perfectly bridge to reportable Graphics, the model uses four lines: Gaming, Professional Visualization, OEM and Other, and a Graphics segment taxonomy bridge. The bridge is not a product thesis; it is the arithmetic residual needed to reconcile the disclosed end-market context to the official reportable segment base. Segment totals are checked against the official $22.459B TTM base.

Base case:

Product / line	TTM base	FY2024-FY2026 CAGR (2 yrs)	Latest YoY	Mgmt FY+1 guide	Base CAGR TTM to TTM+5y	2031-01-25 revenue	Justification
Gaming	$16.042B	23.9%	41.3%	Supply constraints are expected headwind to Gaming [data_package Section 6]	9.0%	$24.683B	Decelerates from the FY2026 recovery as channel inventory normalizes and the AIB unit market matures, but stays above mature PC growth because NVIDIA retains premium share and AI PC/RTX features support mix.
Professional Visualization	$3.191B	43.3%	69.9%	Pro Viz crossed $1B quarterly; AI workstation demand is strong [data_package Section 6]	17.0%	$6.996B	Decelerates sharply from FY2026 but remains the highest-growth Graphics line because AI workstation and RTX professional workflows are earlier in adoption.
OEM and Other	$0.619B	42.2%	59.1%	No specific long-term guidance in data package	5.0%	$0.790B	Treated as residual and cyclical; grows below Gaming because disclosure is too small and noisy to support a stronger thesis.
Graphics taxonomy bridge	$2.607B	46.7% residual CAGR	279.5% residual YoY	No product attribution in data package	4.0%	$3.172B	Low-growth placeholder to preserve the reportable-segment tie; no product thesis is assigned to this residual.

Base-case year-by-year deceleration:

Line	FY2027e	FY2028e	FY2029e	FY2030e	Horizon 2031-01-25
Gaming revenue at 9.0% CAGR	$17.486B	$19.059B	$20.774B	$22.644B	$24.683B
Professional Visualization at 17.0% CAGR	$3.733B	$4.368B	$5.111B	$5.980B	$6.996B
OEM and Other at 5.0% CAGR	$0.650B	$0.682B	$0.716B	$0.752B	$0.790B
Bridge at 4.0% CAGR	$2.711B	$2.819B	$2.932B	$3.049B	$3.172B

The base case implies total Graphics revenue of $35.641B at the horizon, a 9.7% segment CAGR from the $22.459B TTM base. This is far below FY2026's 57.0% reportable segment growth, but it is still a healthy five-year growth rate for a business with a mature PC gaming base. The implied narrative is that GeForce remains dominant and AI PC features improve mix, while Pro Viz grows into a larger AI workstation line.

Bull case:

Product / line	TTM base	FY2024-FY2026 CAGR (2 yrs)	Latest YoY	Mgmt FY+1 guide	Bull CAGR TTM to TTM+5y	2031-01-25 revenue	Justification
Gaming	$16.042B	23.9%	41.3%	Supply is headwind, not demand weakness	13.0%	$29.556B	RTX/AI PC becomes a stronger upgrade cycle, AMD remains mostly price-tier competition, and supply improves without channel overshoot.
Professional Visualization	$3.191B	43.3%	69.9%	Pro Viz crossed $1B quarterly; AI workstation demand strong	24.0%	$9.355B	AI workstation workflows broaden from design/media into local inference, engineering simulation, and enterprise creator seats.
OEM and Other	$0.619B	42.2%	59.1%	No specific long-term guidance	8.0%	$0.910B	Small residual benefits from OEM AI PC attach but remains immaterial.
Graphics taxonomy bridge	$2.607B	46.7% residual CAGR	279.5% residual YoY	No product attribution	8.0%	$3.831B	Bridge grows with the segment but remains explicitly unattributed.

The bull case reaches $43.651B of Graphics revenue at the horizon, a 14.2% segment CAGR. This case requires NVIDIA to convert AI PC and workstation positioning into sustained premium mix, not merely another gaming replacement spike. It also requires enough supply allocation to satisfy Graphics demand while Data Center remains much larger.

Bear case:

Product / line	TTM base	FY2024-FY2026 CAGR (2 yrs)	Latest YoY	Mgmt FY+1 guide	Bear CAGR TTM to TTM+5y	2031-01-25 revenue	Justification
Gaming	$16.042B	23.9%	41.3%	Supply constraints cited as headwind	1.0%	$16.860B	FY2026 proves to be a strong generation/channel recovery, followed by weak replacement demand, price resistance, used-card supply, and AMD price competition.
Professional Visualization	$3.191B	43.3%	69.9%	AI workstation demand strong	7.0%	$4.476B	Pro Viz still grows because AI workstation demand is real, but cloud workstations and budget cycles cap local GPU adoption.
OEM and Other	$0.619B	42.2%	59.1%	No specific long-term guidance	-3.0%	$0.532B	Residual line fades as OEM mix normalizes.
Graphics taxonomy bridge	$2.607B	46.7% residual CAGR	279.5% residual YoY	No product attribution	-8.0%	$1.718B	Residual bridge shrinks because the unexplained FY2026 gap is not assumed durable.

The bear case reaches $23.586B of revenue, a 1.0% segment CAGR. This is not a collapse case; it is a normalization case in which FY2026's high growth was mostly cyclical recovery and channel/supply timing. A true collapse would require a gaming inventory bust comparable to FY2023 plus loss of premium share, but that is not the central downside because NVIDIA's competitive position remains strong.

9. Forward operating margin assumptions - explicit and justified

Current Graphics operating margin is 40.8% from the data package [data_package Section 3, Graphics historicals; data_package Section 5, TTM / FY2026 reportable segments]. Product-line margins are not disclosed. The tables below use product-line margin estimates only to reconstruct the segment margin; the official historical margin anchor is the reportable Graphics margin.

Case	Gaming horizon margin	Professional Visualization horizon margin	OEM and Other horizon margin	Bridge horizon margin	Segment horizon margin	Mechanism
Bear	35.0%	44.0%	12.0%	25.0%	35.5%	Channel normalization, lower high-end mix, supply shifts toward Data Center, and less pricing power compress margin below current.
Base	41.0%	52.0%	20.0%	35.0%	42.2%	Mix shifts toward Pro Viz and high-end RTX while Gaming pricing remains rational; margin improves 140 bps from current.
Bull	45.0%	58.0%	25.0%	40.0%	46.9%	AI workstation and premium GeForce mix expand, supply constraints support pricing, and software ecosystem value sustains high ASPs.

The base case is deliberately modest on margin. A 42.2% horizon margin is only 140 bps above FY2026's 40.8% and still below the FY2022 peak of 53.5% [data_package Section 3, Graphics historicals]. The reason is that FY2026 already includes a strong recovery year, and the Q4 Gaming commentary included a sequential decline attributed to channel inventory after a strong holiday season [data_package Section 5, Q4 recency signals]. If Gaming normalizes, Pro Viz mix must do the margin work.

The bull margin requires more than volume growth. It requires mix shift into high-end GeForce, professional RTX, AI workstations, and premium software-attached workflows. It also requires AMD and Intel not to force a price war in the high-end discrete GPU market. The external market-share context supports that possibility, but it does not make it automatic. Market dominance can coexist with margin pressure if consumers resist prices or if memory/component costs rise.

The bear margin falls to 35.5%, close to FY2025's 35.5% level [data_package Section 3, Graphics historicals]. This is a realistic downside anchor because it has already happened recently without destroying the franchise. The bear does not assume a sub-30% trough because the current segment includes a stronger Pro Viz/AI workstation component than older gaming-only mental models, and because NVIDIA's brand position remains strong. But a channel reset, weaker mid-range demand, and supply prioritization away from Graphics could plausibly push margins back to the mid-30s.

10. Segment rollup - bottoms-up and tied

The rollup below compounds from the TTM/FY2026 base to the horizon date 2031-01-25. Revenue lines sum to modeled Graphics segment revenue. Operating income equals line revenue multiplied by line margin. The blended operating margin equals total operating income divided by total revenue. Total rows use the unrounded compounding outputs, so displayed line items can differ from the total by $0.001B due to rounding. Because the taxonomy bridge is a residual, the end-market lines should not be read as a clean official segment bridge.

Bear case rollup:

Product / line	Horizon revenue	Horizon margin	Horizon operating income
Gaming	$16.860B	35.0%	$5.901B
Professional Visualization	$4.476B	44.0%	$1.969B
OEM and Other	$0.532B	12.0%	$0.064B
Graphics taxonomy bridge	$1.718B	25.0%	$0.430B
Total Graphics	$23.586B	35.5%	$8.364B

Base case rollup:

Product / line	Horizon revenue	Horizon margin	Horizon operating income
Gaming	$24.683B	41.0%	$10.120B
Professional Visualization	$6.996B	52.0%	$3.638B
OEM and Other	$0.790B	20.0%	$0.158B

Product / line	Horizon revenue	Horizon margin	Horizon operating income
Graphics taxonomy bridge	$3.172B	35.0%	$1.110B
Total Graphics	$35.641B	42.2%	$15.026B

Bull case rollup:

Product / line	Horizon revenue	Horizon margin	Horizon operating income
Gaming	$29.556B	45.0%	$13.300B
Professional Visualization	$9.355B	58.0%	$5.426B
OEM and Other	$0.910B	25.0%	$0.227B
Graphics taxonomy bridge	$3.831B	40.0%	$1.532B
Total Graphics	$43.651B	46.9%	$20.486B

Internal tie-out checks:

Check	Result
TTM base	Gaming $16.042B + Pro Viz $3.191B + OEM/Other $0.619B + bridge $2.607B = Graphics $22.459B.
Bear revenue sum	$16.860B + $4.476B + $0.532B + $1.718B = $23.586B.
Base revenue sum	$24.683B + $6.996B + $0.790B + $3.172B = $35.641B.
Bull revenue sum	Unrounded line revenue totals $43.651B; displayed rounded line items sum to $43.652B due to rounding.
Bear OI	$8.364B / $23.586B = 35.5%.
Base OI	$15.026B / $35.641B = 42.2%.
Bull OI	$20.486B / $43.651B = 46.9%.

The base case is the one that should feed the combined ER2 earnings model unless the synthesis agent has stronger company-level reasons to adjust. It gives Graphics a meaningful contribution to operating income growth, from $9.156B today to $15.026B at TTM+5y, but it does not let Graphics dominate the NVIDIA thesis. That is consistent with the data package's note that Graphics remains material but is not the valuation driver versus Data Center [data_package Section 3, Graphics commentary].

11. Key risks and wildcards

Bear risk - channel inventory and replacement-cycle normalization. The data package notes that Q4 Gaming fell 13% sequentially from Q3, with NVIDIA attributing the move to channel inventory after a strong holiday season [data_package Section 5, Q4 recency signals]. If FY2026 pulled forward upgrades, the next cycle could be flat despite good architecture.

Bear risk - consumer price elasticity. NVIDIA's premium share is powerful, but GeForce buyers can delay upgrades. If high-end GPU pricing rises faster than perceived benefit, unit demand can weaken even when the brand remains dominant.

Bear risk - internal supply priority. Management said supply constraints were expected to be the headwind to Gaming [data_package Section 6]. If Data Center remains far more profitable per constrained supply unit, Graphics may be supply limited even when demand exists.

Bear/base risk - taxonomy opacity. Gaming, Professional Visualization, and OEM and Other do not cleanly reconcile to reportable Graphics. The $2.607B FY2026 bridge is explicitly unattributed. If the bridge is one-time, the base case may overstate durable revenue by carrying it forward.

Base wildcard - AI PC features. If local AI features become part of mainstream PC buying decisions, GeForce refresh rates and ASPs can exceed normal gaming replacement logic. This supports the base-case 9% Gaming CAGR, but the data package does not provide a quantitative AI PC revenue line.

Bull wildcard - AI workstations. Professional Visualization reached $3.191B in FY2026 and crossed $1B quarterly in Q4 [data_package Sections 4-6]. If AI workstation demand becomes a durable enterprise category, Pro Viz can become a much larger profit contributor.

Bull risk - share dominance can attract pricing backlash. External JPR-linked sources show NVIDIA's share dominance, but very high share can create regulator, OEM, and customer pressure for alternatives. This is not a near-term revenue disqualifier, but it can cap long-run pricing.

Summary table

Case	Segment TTM+5y Revenue	Segment TTM+5y Op Margin	Segment TTM+5y Op Income	Probability
Bear	$23.586B	35.5%	$8.364B	25%
Base	$35.641B	42.2%	$15.026B	55%
Bull	$43.651B	46.9%	$20.486B	20%

Bottom line: Graphics is a high-quality, premium GPU franchise with real moat characteristics, but it should be modeled as cyclical-plus-AI-workstation optionality rather than as a second Data Center. The base case compounds reportable Graphics revenue at 9.7% for exactly five calendar years from the 2026-01-25 TTM anchor to the 2031-01-25 horizon. That produces $35.641B of revenue and $15.026B of operating income, a material but secondary contribution to NVIDIA's consolidated five-year earnings power.

Appendix S - Synthesis (Combined Earnings Model)

NVIDIA Corporation - Synthesis (Combined Earnings Model)

Date: 2026-04-26 Ticker: NVDA Current price used by the model: $208.27 Output owner: Sub-Agent S

This appendix is the combined earnings model for NVIDIA Corporation. It starts from the reviewed ER2 data package and the reviewed segment appendices, not from a local warehouse cache. The segment horizon revenue and operating income inputs are reproduced from the reviewed Compute & Networking and Graphics appendices. The data package supplies the historical consolidated statements, the segment bridge, the TTM anchor, capital allocation history, and the balance sheet frame. The management appendix supplies the capital allocation cross-check. The moat appendix supplies durability context but does not change the arithmetic. Existing ER2 raw files inside this company work folder were used only to recover required S inputs that the data package did not repeat explicitly: FY2026 stock-based compensation, interest income, and interest expense. Those files are part of this live ER2 run, not warehouse cache data.

The key conclusion is strict: the base case fails the Gideon price gate. The base case is not a weak operating case. It gives NVIDIA $589.241 billion of horizon revenue, $343.876 billion of consolidated operating income, $290.536 billion of net income, and $12.28 of EPS at the TTM+5y horizon. But at the provided current price of $208.27, the required EPS for a 15x P/E is $13.88. Base EPS of $12.28 produces an EPS x 15 value of $184.21, or an 11.6% shortfall versus the current price. The bull case clears, but the ER2 gate is not a bull-case gate.

1. Anchors

The anchor is TTM, not FY-1, and the forward span is exactly five calendar years.

Item	Model value
Today	2026-04-26
Fiscal year convention	Fiscal year ends on the last Sunday in January
TTM anchor end date	2026-01-25
TTM quarters	Q1 FY2026 ended 2025-04-27, Q2 FY2026 ended 2025-07-27, Q3 FY2026 ended 2025-10-26, Q4 FY2026 ended 2026-01-25
Latest completed FY label	FY2026
Horizon date	2031-01-25
Compounding span	Exactly 5.0 calendar years from TTM end
Current price	$208.27
TTM diluted share count	24,432M
Required EPS for P/E 15	$13.88

The TTM period equals FY2026 because the latest completed quarter is Q4 FY2026 and the TTM window is the full FY2026 fiscal year. The data package reports FY2026 revenue of $215.938 billion, operating income of $130.387 billion, net income of $120.067 billion, diluted EPS of $4.90, and diluted weighted average shares of 24,432 million. For the segment anchor, Compute & Networking generated FY2026 revenue of $193.479 billion and operating income of $130.141 billion. Graphics generated FY2026 revenue of $22.459 billion and operating income of $9.156 billion. Segment operating income of $139.297 billion reconciles to consolidated operating income after the data package and quant foundation bridge of $8.910 billion for stock-based compensation, unallocated operating expenses, and acquisition-related or other costs.

The horizon date is 2031-01-25. I do not re-anchor to the fiscal year containing 2031, and I do not stretch the model to six fiscal years. All CAGRs, segment rollups, share count paths, and P/E tests use the five-year span from 2026-01-25 to 2031-01-25.

2. Segment-Level Forward Assumptions

The segment rows below are reproduced from the reviewed segment appendices. They are not re-derived in this appendix. The Compute & Networking row comes from `appendix_B_compute_networking.md` sections 8, 9, 10, and the final summary table. The Graphics row comes from `appendix_B_graphics.md` sections 8, 9, 10, and the final summary table. The user also provided the same reviewed horizon numbers as explicit anchors, and those match the appendices.

Segment	TTM revenue ($B)	Bear TTM+5y rev	Base TTM+5y rev	Bull TTM+5y rev	Bear OI margin	Base OI margin	Bull OI margin
Compute & Networking	193.479	348.800	553.600	734.300	52.0%	62.0%	66.0%
Graphics	22.459	23.586	35.641	43.651	35.5%	42.2%	46.9%
Total reportable segment revenue	215.938	372.386	589.241	777.951	n.m.	n.m.	n.m.

The segment operating income inputs are:

Segment	Bear TTM+5y OI	Base TTM+5y OI	Bull TTM+5y OI
Compute & Networking	181.400	343.200	484.600
Graphics	8.364	15.026	20.486
Sum of segment OI	189.764	358.226	505.086

Compute & Networking is the controlling economic variable. In the base case it contributes 93.9% of segment operating income at the horizon. Graphics matters, but the equity result is still primarily a data center and networking result. A very large error in Graphics would not rescue or destroy the base valuation unless it coincided with a Compute & Networking error.

3. Combined Revenue Build

NVIDIA has two reportable segments in the data package: Compute & Networking and Graphics. The data package's segment revenue rows sum exactly to consolidated revenue in FY2026. There is no separate corporate revenue line or elimination needed in this synthesis. Product and end-market lines such as Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other are useful demand context, but the data package explicitly warns not to force them into a mechanical reportable-segment bridge. Therefore the combined revenue build uses the official reportable segment totals.

Line	Bear	Base	Bull
Compute & Networking revenue	348.800	553.600	734.300
Graphics revenue	23.586	35.641	43.651
Sum of segment TTM+5y revenue	372.386	589.241	777.951
Corporate eliminations	0.000	0.000	0.000
Combined TTM+5y revenue	372.386	589.241	777.951
Implied TTM to TTM+5y revenue CAGR	11.5%	22.2%	29.2%

The base case is already a very strong company-level growth case. Revenue grows from $215.938 billion to $589.241 billion in exactly five years. That is a 2.73x revenue increase and a 22.2% annualized growth rate from a $216 billion base. The bear case still grows revenue at an 11.5% CAGR, which is a strong result for most large-cap companies but a sharp deceleration for NVIDIA after the FY2024-FY2026 AI infrastructure surge. The bull case grows revenue at 29.2% annually and would imply that the AI factory build-out remains under-supplied and that NVIDIA retains system-level pricing power at extraordinary scale.

The base case also starts from a high exit rate. Q4 FY2026 revenue was $68.127 billion, or $272.508 billion annualized. The first forward year implied by the segment appendices is about $285.6 billion, only 4.8% above the Q4 exit run-rate. That matters because the five-year base case does not require a step-change on day one; it requires the Q4 exit level to continue expanding as Blackwell/Rubin systems, networking attach, and AI infrastructure deployments scale.

4. Combined Operating Income Build

The segment appendices produce segment operating income before the corporate bridge. The data package and quant foundation bridge FY2026 segment operating income of $139.297 billion to consolidated operating income of $130.387 billion through $8.910 billion of stock-based compensation, unallocated operating expenses, and acquisition-related or other costs. That bridge is not ignored. For the horizon model, I grow the bridge at case-specific rates: 8% per year in bear, 10% in base, and 12% in bull. The rate is well below revenue growth in the base and bull cases, which reflects scale leverage, but it is not frozen. This is consistent with the management appendix: NVIDIA is spending aggressively on R&D, supply chain support, ecosystem investments, and operating infrastructure even while capex remains modest because the company is fabless.

Line	Bear	Base	Bull
Sum of segment TTM+5y OI	189.764	358.226	505.086
Unallocated corporate bridge	(13.092)	(14.350)	(15.702)
Combined TTM+5y OI	176.672	343.876	489.384
Implied OI margin	47.4%	58.4%	62.9%

This operating margin table is the first place where the thesis becomes demanding. The bear case margin of 47.4% is still a very high consolidated margin for a hardware-linked company, but it is a severe decline from FY2026's consolidated operating margin of 60.4%. The base case margin of 58.4% is only modestly below FY2026 and requires Compute & Networking to remain a platform business, not a commodity accelerator business. The bull margin of 62.9% is above FY2026 consolidated margin and would require scarcity, networking attach, software/platform value capture, and limited hyperscaler buyer leverage to persist at far larger scale.

The unallocated bridge is small relative to the segment profit pool, but it matters for EPS. In the base case, the bridge reduces horizon operating income by $14.350 billion, or about $0.51 of after-tax EPS. A synthesis that simply used segment operating income as consolidated operating income would overstate the base valuation cushion.

5. Net Income and EPS Build

For the non-operating line, I use a conservative +$2.000 billion net interest and other income assumption in each case. FY2026 interest income was $2.300 billion and interest expense was $0.259 billion in the existing ER2 FMP income statement pull, while the data package reports $62.556 billion of cash plus short-term investments and $51.144 billion of net cash after total debt. The company therefore has ample balance sheet support for positive net interest income. I do not carry forward the full FY2026 difference between operating income and pre-tax income because that would risk embedding investment gains or non-operating items that are not part of the operating thesis. A flat +$2.000 billion is deliberately conservative.

The tax rate assumption is 16.0% in all cases. FY2026's effective tax rate was 15.1% based on income tax expense of $21.383 billion and income before tax of $141.450 billion. A 16.0% rate is slightly more conservative and avoids using a tax tailwind to rescue the valuation gate.

Line	Bear	Base	Bull
Combined TTM+5y OI	176.672	343.876	489.384
Net interest / other income	2.000	2.000	2.000
Pre-tax income	178.672	345.876	491.384
Tax rate assumed	16.0%	16.0%	16.0%
Net income	150.085	290.536	412.762
Diluted share count at TTM+5y	23,657.9M	23,657.9M	23,657.9M
EPS at TTM+5y	$6.34	$12.28	$17.45

The EPS result is the core valuation output. Base net income of $290.536 billion sounds enormous because it is enormous. But the share count is still nearly 23.7 billion shares, and the current price already capitalizes a very large earnings base. This is why the ER2 discipline uses an EPS x 15 gate rather than a narrative growth verdict. If the current price demands more than the base-case five-year EPS can support at 15x, the correct answer is PASS or WATCH, not BUY.

6. Share Count Path

The share path is explicit. The data package reports FY2026 diluted shares of 24,432 million. It also reports FY2026 buybacks of $40.086 billion and FY2026 free cash flow of $96.676 billion. The management appendix notes that the board approved an additional $60 billion repurchase authorization on 2025-08-26 and that $58.5 billion remained authorized at fiscal year-end. Buybacks are therefore real, but the appendix also warns that price sensitivity matters and that repurchases at a multi-trillion-dollar market capitalization are not automatically high-return.

For the base share path, I model $45.0 billion of gross repurchases each year for five years. That is slightly above the FY2026 repurchase amount and remains affordable under the base earnings case. I use rising average repurchase prices of $208.27, $225, $245, $265, and $285. For SBC issuance, I use the S instruction's rough method: FY2026 stock-based compensation divided by the current share price. FY2026 SBC was $6.386 billion, and $6.386 billion divided by $208.27 equals 30.7 million

shares per year. This is an approximation, not a grant-date valuation model, but it is explicit and prevents the model from assuming that buybacks are pure share reduction.

Year	Starting share count (M)	Buyback $ ($B)	Avg price assumed	Shares retired (M)	SBC issuance (M)	Ending share count (M)
FY2027e	24,432.0	45.0	$208.27	216.1	30.7	24,246.6
FY2028e	24,246.6	45.0	$225.00	200.0	30.7	24,077.3
FY2029e	24,077.3	45.0	$245.00	183.7	30.7	23,924.2
FY2030e	23,924.2	45.0	$265.00	169.8	30.7	23,785.1
FY2031e	23,785.1	45.0	$285.00	157.9	30.7	23,657.9

The ending share count is 23,657.9 million. Total gross repurchases are $225.0 billion. Total shares retired are about 927.4 million. Total modeled SBC issuance is about 153.3 million. Net share reduction is about 774.1 million, or 3.2% of the starting diluted share count over five years. This is meaningful but not thesis-saving. The company is simply too large, and the current market capitalization is too high, for even $45 billion per year of repurchases to transform EPS.

7. P/E 15 Test

The required EPS for a P/E of 15 at the current price is:

```
$208.27 / 15 = $13.88
```

Case	EPS at TTM+5y	EPS x 15	Current price	Cushion	Pass?
Bear	$6.34	$95.16	$208.27	-54.3%	FAIL
Base	$12.28	$184.21	$208.27	-11.6%	FAIL
Bull	$17.45	$261.71	$208.27	+25.7%	PASS

The base case is the controlling case. The bull case can clear the valuation gate, but the framework does not buy stocks because the bull case works. The base case must clear with confidence. It does not. The base case would need another $1.60 of EPS, or about $37.8 billion of after-tax earnings at the modeled share count, to reach the $13.88 required EPS. That is a large gap. It is not a rounding issue. It could be closed by higher Compute & Networking revenue, higher margin, lower tax, more non-operating income, or far more buybacks, but those would be thesis upgrades, not the current base case.

At the current base EPS, a 15x multiple supports $184.21. A strict 15x base-case entry price would therefore be about $184, before any additional margin of safety. At $208.27, the stock needs either a stronger base case than the reviewed appendices support or acceptance of a higher terminal multiple than the ER2/Gideon gate allows.

8. Reflexivity Check

The combined earnings build must tie to reported TTM financials before it is used for forward valuation.

Historical tie	Calculation	Result
Segment revenue to consolidated revenue	$193.479B Compute & Networking + $22.459B Graphics = $215.938B consolidated revenue	PASS, 0.0% gap
Segment OI to consolidated OI	$130.141B Compute & Networking OI + $9.156B Graphics OI = $139.297B segment OI; less $8.910B bridge = $130.387B consolidated OI	PASS, 0.0% gap
Latest completed FY tie	TTM equals FY2026 because Q4 FY2026 is the latest quarter and the TTM window is Q1-Q4 FY2026	PASS

The reflexivity result is clean. The forward model compounds from the same segment base that ties to consolidated FY2026 revenue and operating income. This reduces the risk of the MSFT-style product-to-segment mapping failure. The model does not sum Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other into reportable segments. It uses those product or end-market lines only where the reviewed segment appendices use them as internal modeling allocations.

9. In-Progress FY Reconciliation

No in-progress fiscal year adjustment is needed. The Stage 0 anchor uses the latest completed fiscal year and the latest completed quarter, Q4 FY2026 ended 2026-01-25. The TTM window is Q1 FY2026 plus Q2 FY2026 plus Q3 FY2026 plus Q4

FY2026. It is not an estimated or partially completed year. Because FY2026 is complete, there is no YTD-to-full-year gap to reconcile under the in-progress FY rule.

I still check the first forward year for recency coherence. The segment appendices imply about $285.6 billion of first-forward-year revenue when using their stated first-year deceleration paths. Q4 FY2026 revenue annualized is $272.5 billion. The first-forward-year revenue is therefore about 4.8% above the Q4 exit run-rate. That is within the 10% recency guardrail when the guardrail is applied to the current exit run-rate rather than the full TTM average. Compared with TTM/FY2026 revenue of $215.938 billion, the first forward year is 32.3% higher, but that is a function of the steep sequential FY2026 ramp: Q1 revenue was $44.062 billion and Q4 was $68.127 billion. The model is not starting above the observed Q4 run-rate; it is extending that exit rate with continued growth.

10. Sensitivity

The sensitivity table holds two variables at base and varies the third. It shows that the verdict is overwhelmingly driven by Compute & Networking growth and consolidated margin, not by Graphics or minor tax changes.

Variable	Downside case	EPS impact	Upside case	EPS impact
Compute & Networking revenue CAGR +/- 200 bps	C&N horizon revenue $510.2B	EPS falls to $11.33	C&N horizon revenue $599.9B	EPS rises to $13.30
Graphics OI margin +/- 200 bps	Graphics margin 40.2%	EPS falls to $12.26	Graphics margin 44.2%	EPS rises to $12.31
Share count +/- 5%	Share count 24,840.8M	EPS falls to $11.70	Share count 22,475.0M	EPS rises to $12.93
Tax rate +/- 300 bps	Tax rate 19.0%	EPS falls to $11.84	Tax rate 13.0%	EPS rises to $12.72

The first sensitivity is the important one. A 200 bps increase in Compute & Networking CAGR lifts base EPS from $12.28 to $13.30, still short of the $13.88 required EPS. That means the base case needs more than a modest C&N revenue upgrade unless it also gets help from margin, share count, or tax. The Graphics margin sensitivity barely matters because Graphics is too small relative to Compute & Networking. A 5% lower share count helps, but it also does not clear the gate by itself. A lower tax rate helps less than one dollar of EPS.

The strongest potential route to a passing base case is a combined upgrade: Compute & Networking grows above the reviewed base, the segment margin stays closer to the FY2026 level than the reviewed base assumes, and buybacks retire more stock than this model assumes. That package could happen, but it would be a higher-confidence, more bullish variant than the reviewed base case. The present base does not earn a BUY under the Gideon rule.

11. Combined Earnings Summary

Case	Revenue ($B)	OI ($B)	OI margin	NI ($B)	EPS	Probability	EPS x 15	vs. $208.27
Bear	372.386	176.672	47.4%	150.085	$6.34	25%	$95.16	-54.3%
Base	589.241	343.876	58.4%	290.536	$12.28	50%	$184.21	-11.6%
Bull	777.951	489.384	62.9%	412.762	$17.45	25%	$261.71	+25.7%
Probability-weighted EPS	n.m.	n.m.	n.m.	n.m.	$12.09	100%	$181.35	-12.9%

I use 25% bear, 50% base, and 25% bull. That weighting is not a uniform 33/33/34 split. It reflects the disqualifier sweep and the moat work: NVIDIA has a real moat, strong management, strong balance sheet, high ROIC, and strong end-market demand, so the base case deserves the highest probability and the bull case remains credible. It also reflects the red flags embedded in the segment work: hyperscaler custom silicon, export controls, buyer concentration, power constraints, supply-chain commitments, and the possibility that AI capex returns disappoint. Those risks prevent a bull-heavy weighting.

The probability-weighted EPS is $12.09, below the base EPS because the bear downside is more severe than the bull upside on EPS x 15 value. This is another way to see the valuation problem. The stock does not merely need the base case to arrive; it needs either the bull case, a higher accepted multiple, or a cheaper entry price.

Tie-Out Checks

All 11 checks were run explicitly. Result: 11/11 passed, with the valuation gate itself failing because the arithmetic passes and shows insufficient base-case EPS.

- Reported-financial tie-out, TTM and latest completed FY: PASS. TTM revenue is $193.479B plus $22.459B, which equals $215.938B. TTM segment OI is $130.141B plus $9.156B, or $139.297B; subtracting the $8.910B corporate bridge equals $130.387B consolidated OI. Latest completed FY is FY2026, which is the TTM period, so the FY tie also passes.

- NI divided by shares equals EPS: PASS. Bear: $150.085B / 23.6579B shares = $6.34. Base: $290.536B / 23.6579B shares = $12.28. Bull: $412.762B / 23.6579B shares = $17.45. Rounding explains displayed penny differences.

- Segment sum tie-out at horizon: PASS. Bear revenue is $348.800B plus $23.586B, or $372.386B. Base revenue is $553.600B plus $35.641B, or $589.241B. Bull revenue is $734.300B plus $43.651B, or $777.951B. Corporate eliminations are zero because reportable segment revenue already sums to consolidated revenue.

- Product-to-segment rollup: PASS. The Compute & Networking appendix ties its AI compute/platform, Data Center Networking, and Automotive/edge model lines to the segment total in each case, with only rounding differences. The Graphics appendix ties Gaming, Professional Visualization, OEM and Other, and the taxonomy bridge to the Graphics segment total in each case, with only rounding differences.

- Weighted margin check: PASS. Compute & Networking presents line margins and a segment margin that is explicitly a revenue-weighted average: 52.0% bear, 62.0% base, and 66.0% bull. Graphics presents line operating income and total operating income, and the total margins calculate to 35.5%, 42.2%, and 46.9%. The synthesis uses those reviewed segment-level margins and OI values.

- Share count path: PASS. Starting diluted shares are 24,432.0M. Five years of buybacks retire about 927.4M shares, while five years of SBC issuance add about 153.3M shares. Ending diluted shares are 23,657.9M. The path uses exactly five annual steps and does not assume a flat share count.

- Growth plausibility: PASS. Compute & Networking's base segment CAGR is 23.4%, far below its disclosed FY2023-FY2026 CAGR of 134.2% and below 1.2x that historical rate. Graphics' base segment CAGR is 9.7%, below its FY2024-FY2026 period CAGR of 28.9%. Product-line base CAGRs are also below their recent disclosed growth rates or are specifically justified as decelerated model allocations in the segment appendices.

- Historical arithmetic: PASS. The segment appendices compute `TTM revenue x (1 + CAGR)^5 = horizon revenue` over exactly five years for each material product/model row. Sample checks: C&N AI compute/platform $159.754B at 23.0% reaches about $449.8B; C&N networking $31.376B at 25.0% reaches about $95.8B; Graphics Gaming $16.042B at 9.0% reaches $24.683B; Graphics Professional Visualization $3.191B at 17.0% reaches $6.996B. Differences are within rounding tolerance.

- Recency anchor: PASS with note. The base first-forward-year revenue implied by the segment appendices is about $285.6B. Q4 FY2026 revenue annualized is $272.5B, so the first forward year is 4.8% above the latest quarter exit run-rate. That passes the 10% guardrail. It is 32.3% above full-year TTM revenue because FY2026 ramped sharply from Q1 to Q4; this is noted but not treated as a model break because the model does not start above the Q4 exit rate.

- Capital allocation and capex consistency: PASS. The share count path uses $45.0B of annual repurchases, which is consistent with FY2026 repurchases of $40.1B, FY2026 FCF of $96.7B, and management's remaining authorization. The management appendix notes that capex should rise in FY2027 and that supply-chain and ecosystem support are the more important commitments; this synthesis does not assume aggressive financial engineering that would crowd out those investments.

- P/E 15 arithmetic and probability summary: PASS. Required EPS is $208.27 / 15 = $13.88. Bear EPS x 15 is $95.16, base EPS x 15 is $184.21, and bull EPS x 15 is $261.71. The probability-weighted EPS is $12.09. The formulas tie. The gate result is a base-case FAIL, not an arithmetic failure.

Key Findings

The historical and segment tie-outs are clean. NVIDIA's FY2026 segment revenue sums exactly to consolidated revenue, and segment operating income reconciles to consolidated operating income through the disclosed $8.910 billion bridge. This is a sound starting point for a bottoms-up model.

The operating base case is strong but not cheap enough. A company reaching $589.241 billion of revenue and $290.536 billion of net income within five years would be one of the most profitable enterprises ever modeled in an ER2 run. The issue is not business quality. The issue is that the current price already requires more than the reviewed base case produces under a 15x gate.

Compute & Networking is the only variable that can realistically change the answer. Graphics is profitable and strategically useful, but its base-case operating income contribution is $15.026 billion versus Compute & Networking's $343.200 billion segment OI. Even a large Graphics surprise would be second-order.

The moat appendix supports a five-year moat pass but not a complacent 20-year extrapolation. CUDA, networking, full-stack systems, scale, and procurement trust are real, but the 2046 test is yellow because custom silicon, open software, regional stacks, and AI-assisted portability can normalize margins over time.

The management appendix supports continued investment and repurchases, but it does not support treating buybacks as a magic EPS offset. At $45 billion per year, modeled repurchases reduce shares by only 3.2% net over five years after SBC issuance. More aggressive buybacks would require either a lower average price, a larger capital return commitment, or less investment elsewhere.

The synthesis verdict for the P/E 15 gate is FAIL. Base EPS of $12.28 is below the required $13.88. The stock would need a lower entry price around $184 at 15x base EPS, a stronger base case, or a willingness to accept a higher multiple. Under the ER2/Gideon framework, the clean arithmetic points to PASS/WATCH rather than BUY at $208.27.

Appendix CS - Compounder Score

NVIDIA Corporation - Compounder Score

Analysis date: 2026-05-28 Ticker: NVDA Currency: USD Diagnostic status: score only; not an ER2 gate

1. Executive Verdict

The evidence supports a very high-quality business score for NVIDIA, but it does not support an investable-now diagnostic score under the CompounderScorer rubric. The total score is **75 / 100**, with **69 / 80** for business quality and **6 / 20** for valuation / 20% return underwriting. The score-implied classification is **Buy only at lower price**.

This is not a weak-company result. The local ER2 source chain shows FY2026 revenue of $215.938 billion, operating income of $130.387 billion, net income of $120.067 billion, free cash flow of $96.676 billion, and FMP ROIC of 62.9%. Compute & Networking produced $193.479 billion of FY2026 revenue and $130.141 billion of segment operating income at a 67.3% segment margin. The balance sheet had about $62.556 billion of cash plus short-term investments, $11.412 billion of total debt, and roughly $51.144 billion of net cash after debt. Those are exceptional facts.

The problem is price and long-duration uncertainty. The ER2 synthesis uses a $208.27 current price and a TTM anchor ending 2026-01-25. The five-year horizon is 2031-01-25. In the reviewed base case, NVIDIA reaches $589.241 billion of revenue, $343.876 billion of operating income, $290.536 billion of net income, and $12.28 of EPS. At 15x EPS, that supports $184.21, which is 11.6% below the modeled current price. A 20% annualized shareholder return case is much more demanding than the ER2 15x gate, so this scorer will not credit an investable-now classification. A 20% case requires a much lower entry price, a materially stronger base case, or a terminal multiple that the rubric explicitly refuses to rely on.

The business quality score is high because NVIDIA has real incremental-return evidence, a large reinvestment runway, founder-led capital allocation, high cash conversion, a strong balance sheet, and a concrete platform moat. It is not a perfect score because ROIIC is not directly disclosed, product-level economics are incomplete, customer concentration and custom silicon are live threats, export controls remain material, and the moat appendix itself marked the 2046 durability test YELLOW. Great business, demanding price.

2. Data Source Map

The scorer used the existing local ER2 artifacts first, as instructed. No API keys were read, printed, or written. No live fetch was needed because the local ER2 artifacts already contained reconciled financials, source maps, segment models, management work, moat work, and the synthesis valuation model.

Item	Local source used	Evidence used
Price / market cap	`main_narrative.md`; `appendix_synthesis.md`	Current price $208.27; market cap about $5.06T
Enterprise value	Calculated from local market cap and net cash	Market cap less $51.144B net cash gives approx. $5.009T EV
Income statement	`quant_foundation.md`; `data_package.md`	FY2026 revenue $215.938B, OI $130.387B, NI $120.067B
Cash flow	`quant_foundation.md`; `data_package.md`	FY2026 OCF $102.718B, capex $6.042B, FCF $96.676B
Balance sheet	`quant_foundation.md`; `data_package.md`	Cash plus STI $62.556B; debt $11.412B; net cash $51.144B
Share count	`data_package.md`; `appendix_synthesis.md`	FY2026 diluted shares 24,432M; model horizon shares 23,657.9M
SBC	`quant_foundation.md`; `appendix_synthesis.md`	FY2026 SBC $6.386B; synthesis models 30.7M annual SBC issuance
Segment data	`quant_foundation.md`; `data_package.md`; segment appendices	C&N and Graphics revenue, OI, margins, forecasts
Management	`appendix_management.md`	Management gate PASS with monitoring caveats
Moat	`appendix_moat.md`; `main_narrative.md`	Current moat strong; 2046 test YELLOW under stricter rubric
Valuation / return model	`appendix_synthesis.md`; `main_narrative.md`	Base EPS $12.28; required EPS $13.88; P/E 15 gate FAIL

Item	Local source used	Evidence used
Risks	`appendix_red_team.md`; segment appendices	Marginal economics, ASICs, capex ROI, export controls

3. Key Financial and Quality Metrics

The key quantitative starting point is FY2026 / TTM, because NVIDIA's latest reported quarter in the ER2 run was Q4 FY2026 and the TTM window equals FY2026. The local source chain uses 2026-01-25 as the TTM end date and 2031-01-25 as the five-calendar-year horizon. This scorer follows that anchor.

Metric	Value	Source / calculation
FY2026 revenue	$215.938B	Quant foundation and data package
FY2026 operating income	$130.387B	Quant foundation
FY2026 net income	$120.067B	Quant foundation
FY2026 OCF	$102.718B	Quant foundation
FY2026 capex	$6.042B	Quant foundation
FY2026 FCF	$96.676B	OCF less capex
FY2026 FCF / NI	80.5%	Quant foundation
FY2026 ROIC	62.9%	FMP key metrics as reproduced in data package
FY2026 adjusted NOPAT estimate	$110.7B	$130.387B OI times 84.9% after-tax rate
FY2026 EV / NOPAT	45.3x	Approx. $5.009T EV / $110.7B NOPAT
FY2026 FCF yield	1.9%	$96.676B / $5.06T market cap
FY2026 SBC / revenue	3.0%	$6.386B / $215.938B
FY2026 SBC / FCF	6.6%	$6.386B / $96.676B
Cash plus STI net debt / EBITDA	-0.35x	Data package

The 5-year median reported ROIC row across FY2022-FY2026 is strong even after the FY2023 correction: 24.6%, 11.7%, 51.3%, 75.3%, and 62.9%. The median is 51.3%. The record supports a high normalized ROIC score. The exact ROIIC score is more conservative because the local artifacts do not provide a clean multi-year organic reinvestment series with capitalized R&D, working capital, acquisitions, ecosystem investments, and growth opex separated. NVIDIA plainly created enormous incremental NOPAT on modest conventional capex, but the scorer will not publish a precise 3-year or 5-year ROIIC without a full reinvestment reconstruction.

4. Rubric Scorecard

Category	Score	Max	Rationale
Incremental Returns on Capital	20	25	ROIC and cash conversion are excellent; ROIIC is directionally exceptional but not cleanly calculable from the available artifacts without unsupported reinvestment adjustments.
Reinvestment Runway	18	20	AI infrastructure, networking, inference, sovereign AI, enterprise AI, simulation, and physical AI give a large runway; the score is capped for capex reflexivity and customer ROI risk.
Durability of Moat	12	15	CUDA/software, systems, networking, supply access, and high-trust procurement are real; 2046 durability is YELLOW, not an unqualified pass.
Accounting and Financial Quality	9	10	Cash conversion, balance sheet strength, reconciled QF/QFR tie-outs, and SBC candor are strong; ecosystem investments and product-level gaps prevent a perfect score.

Category	Score	Max	Rationale
Capital Allocation	10	10	Founder-led strategy, Mellanox/networking, sustained R&D, buybacks funded by FCF, and aligned ownership support full credit, with monitoring items noted.
Price and 20% Return Underwriting	6	20	The base case fails the ER2 15x EPS gate, and the 20% return case needs a lower price or a materially stronger base case.
Total	**75**	**100**	Business quality is excellent; valuation blocks the investable-now classification.

The ROIIC plus runway subtotal is 38 / 45, which clears the diagnostic quality/runway threshold. The business quality score is 69 / 80, also above the 65 / 80 threshold. The total score is below 80 because valuation is poor under the 20% return rubric. The valuation score is only 6 / 20, far below the 15 / 20 threshold required for "Investable now under this rubric."

5. ROIIC and Reinvestment Analysis

The evidence supports very high incremental returns, but the exact ROIIC figure is marked insufficient evidence. That may sound conservative for a company whose operating income exploded from $32.972 billion in FY2024 to $130.387 billion in FY2026 while capex remained only $6.042 billion in FY2026. It is still the right treatment under this scorer's rules. A proper ROIIC denominator should include net capex, working capital, capitalized R&D adjustment, material growth opex, acquisition spend, ecosystem investments, and possibly supply-chain commitments. The local ER2 artifacts disclose many of those components qualitatively, but they do not present a complete normalized reinvestment bridge.

Reported/simple ROIC is clearly high. The data package's FMP key metrics row gives FY2026 ROIC of 62.9%, and the five-year median reported ROIC is about 51.3%. That is far above the 30% threshold for full credit in normalized ROIC. The cash conversion is also real. FY2026 FCF of $96.676 billion was 80.5% of net income and about 87% of the scorer's estimated FY2026 NOPAT. Operating cash flow exceeded $102 billion.

The accounting-artifact risk is not zero, but it does not break the score. NVIDIA is fabless, so reported invested capital and capex can look very low relative to economic commitments. The right skepticism is that supply commitments, advanced packaging capacity, ecosystem financing, R&D, and customer support are part of the reinvestment engine even if they do not look like a factory on the balance sheet. This is why the scorer does not award a perfect 25 / 25 in incremental returns. It credits the observed economics while refusing to pretend that conventional capex is the entire capital base.

6. Runway Analysis

The reinvestment runway score is high. Compute & Networking is nearly 90% of revenue and over 93% of segment operating income. The segment combines AI accelerators, systems, networking, software, supply orchestration, and high-trust procurement. The ER2 segment work models base Compute & Networking revenue rising from $193.479 billion to $553.6 billion over five years, a 23.4% segment CAGR, with segment operating margin normalizing from 67.3% to 62.0%. The base case is already ambitious, but it is not a simple extrapolation of FY2024-FY2026 hypergrowth. It decelerates materially.

The runway is broad: frontier model training, inference, recommender systems, coding, multimodal AI, sovereign AI, enterprise AI, industrial simulation, robotics, autonomous systems, and data center networking. Networking is already a major line: management commentary in the data package says networking exceeded $31 billion for FY2026 and reached $11 billion in Q4 FY2026. That matters because it turns NVIDIA from a chip vendor into more of an AI factory systems vendor.

The score is capped because buyer economics are not guaranteed. The red-team and segment appendices identify the central risk: customers can continue buying AI infrastructure while NVIDIA loses marginal economics to custom silicon, AMD, Broadcom partners, open accelerator software, Ethernet fabrics, and bargaining power from concentrated hyperscaler customers. The base case assumes this pressure trims margins but does not collapse the platform. That is reasonable, but not certain.

7. Moat Analysis

The moat is concrete. In Compute & Networking, the advantage is not just a GPU benchmark. It is CUDA, CUDA-X libraries, compilers, inference engines, NVLink, InfiniBand, Ethernet networking, DPUs, NICs, systems, cloud/OEM availability, developer familiarity, and procurement trust. Customers buying AI infrastructure care about time-to-production, utilization,

reliability, power efficiency, and software continuity. Those are switching costs.

The moat evidence appears in margins and scale. FY2026 consolidated gross margin was 71.1%, operating margin was 60.4%, and Compute & Networking segment operating margin was 67.3%, despite export-control charges and Blackwell transition complexity. Weak pricing power would not usually produce those economics at $215.938 billion of revenue.

The moat is not permanent by declaration. The moat appendix explicitly marked the 2046 test YELLOW. Twenty years is long enough for hyperscaler ASICs, open software, regional stacks, AI-assisted compiler/tooling improvements, and buyer bargaining to erode CUDA's default status and normalize margins. The stricter May 2026 main narrative moved the overall moat gate to YELLOW to reflect that long-duration uncertainty. This scorer therefore awards 12 / 15, not 15 / 15.

8. Accounting and Capital Allocation Analysis

The accounting and financial quality score is 9 / 10. The QF and QFR gates cleared, QFR found a $0 FMP-vs-official gap, and the synthesis had 11 of 11 tie-out checks pass. Historicals tie to official segment and consolidated sources. The data package explicitly warned against forcing end-market revenue into reportable segments, which reduces product-taxonomy error risk.

SBC is treated as real. FY2026 stock-based compensation was $6.386 billion, about 3.0% of revenue and 6.6% of FCF. The synthesis modeled SBC issuance explicitly at about 30.7 million shares per year. NVIDIA also announced that it would include SBC expense in non-GAAP measures beginning Q1 FY2027, a positive candor marker.

Capital allocation scores full credit because the historical record is unusually strong. CUDA, accelerated computing, Mellanox/networking, Blackwell/Rubin cadence, sustained R&D, and supply/ecosystem support have created a massive platform. FY2026 R&D was $18.497 billion, up materially from FY2022's $5.268 billion. NVIDIA repurchased $40.086 billion of stock and paid $0.974 billion of dividends in FY2026 while generating $96.676 billion of FCF. Repurchases are affordable, though not automatically high-return at a multi-trillion-dollar valuation.

The yellow flag is strategic ecosystem investment. The management appendix noted Anthropic, Groq-related, and OpenAI partnership / investment items as monitoring risks. These may strengthen the platform, but they can blur the boundary between clean end demand and ecosystem-financed demand. That does not fail capital allocation today; it is a monitoring item.

9. 20% Return Underwriting

The valuation score is the anchor problem. The ER2 synthesis did not build a weak base case. It modeled $589.241 billion of horizon revenue, $343.876 billion of operating income, $290.536 billion of net income, and $12.28 of EPS at 2031-01-25. Even so, at the $208.27 price used in the model, a 15x EPS gate requires $13.88 of EPS. Base EPS x 15 supports only $184.21.

For the CompounderScorer 20% hurdle, the answer is worse. If a buyer pays $208.27 and receives roughly $184.21 of 15x base-case value after five years, the implied annualized price return is about negative 2.4% before dividends. The company pays a very small dividend, so cash distributions do not repair the result. If the bull case reaches $17.45 of EPS and receives 15x, the terminal value is $261.71, which implies only about 4.7% annualized price return from $208.27. That is not a 20% underwriting case.

A rough maximum buy price for a 20% five-year return using the base 15x value is about $74 if no dividends are credited. Because NVIDIA generates large cash flow and may repurchase shares, a more generous diagnostic ceiling could be rounded toward $95, but even that is far below both $208.27 and the ER2 $184 strict 15x base-case trigger. The exact max-buy-price figure is less important than the direction: the 20% return case is not close at the modeled price.

10. Hard Veto Checklist

Veto item	Status	Rationale
ROIIC below cost of capital	Pass / insufficient precision	Directionally far above cost of capital, but exact ROIIC not published.
ROIC artifact only	Watch	Reported ROIC is real, but conventional invested capital understates economic commitments for a fabless platform.
Growth requires external financing	Pass	Net cash balance sheet and huge FCF.

Veto item	Status	Rationale
Share count rises materially	Pass	Diluted shares fell from 25,350M in FY2022 to 24,432M in FY2026; synthesis models further net reduction.
Large acquisitions without returns	Pass / Watch	Mellanox appears excellent; newer ecosystem investments need monitoring.
FCF structurally lags earnings	Pass	FY2026 FCF / NI was 80.5%.
Customer concentration / regulation / cyclicality	Watch	Hyperscalers, export controls, and AI capex reflexivity are material.
20% case requires multiple expansion	Fail	The base 20% case does not work at the current modeled price.
Needs flawless execution	Watch / Fail for valuation	The stock needs very strong execution plus valuation help.
Moat cannot be explained	Pass	The moat is explainable in one paragraph.
Accounting / governance trust issue	Pass / Watch	No hard issue; related-party employment and ecosystem circularity monitored.
Balance sheet risk	Pass	Net cash balance sheet.

11. Final Decision

Classification: Buy only at lower price. NVIDIA is one of the best operating businesses in the ER2 folder, but the diagnostic score refuses to confuse business excellence with investment attractiveness at any price. The company earns a high business-quality score and fails the 20% return underwriting test.

Compounder Score note: 75/100: exceptional business quality, but the diagnostic score is capped by valuation and by 2046 moat uncertainty.

Business Quality Score note: 69/80: high ROIC, strong cash generation, founder-led execution, and a real platform moat; points withheld for ROIIC precision, customer concentration, export controls, circular-demand risk, and long-duration margin normalization.

20% Return UW Score note: 6/20: the reviewed ER2 base case fails the 15x FY+5 EPS gate and a 20% return case requires either a much lower entry price, a materially stronger base case, or a terminal multiple that this rubric will not credit.

Data gaps: exact organic ROIIC and acquisition-inclusive ROIIC; product-level operating margins; standalone Blackwell/Hopper/Rubin/CUDA/DRIVE/robotics revenue; customer-level AI capex ROI; and exact current market data refreshed after the ER2 model snapshot. None of these gaps would improve the valuation score without a materially different base case.

12. Machine-Readable JSON

See `compounder_score.json` in this folder.

Appendix CSA - Scoring Audit

NVIDIA Corporation - Scoring Audit

Analysis date: 2026-05-28 Ticker: NVDA Auditor: CSA scoring auditor

1. Original vs Audited Scores

Item	Original	Audited	Notes
Business Quality Score / 80	69	60	This score is reduced because ROIIC is not calculated, capex/customer concentration evidence is incomplete, and capital allocation gets too much credit despite strategic investment circularity and buyback price risk.
Valuation / 20% Return UW Score / 20	6	3	This score is reduced because base, bear, and bull price-return math all fail the 20% hurdle at 15x; the original max buy price is internally inconsistent.
Total Score / 100	75	63	The business is exceptional, but the audited diagnostic score is capped by missing ROIIC evidence and a failed 20% underwriting case.
Classification	Buy only at lower price	Reject under score threshold / watch as business	Under the CS threshold table, a score below 70 maps to Reject; practically, this means not investable under the 20% rubric at the modeled price.

Original reconstructed score: Business Quality 69 / 80, Valuation / 20% Return Underwriting 6 / 20, total 75 / 100. The original JSON and markdown agree on the headline scores, classification, current price of $208.27, market cap of about $5.06T, EV of about $5.009T, base-case IRR of 2.5%, bear-case IRR of -14.4%, bull-case IRR of 4.7%, and max buy price of $95.00. The source chain uses TTM ending 2026-01-25 and horizon date 2031-01-25.

2. Score Reductions and Reasons

This score is reduced because Category 1B requires a credible 3-5 year ROIIC calculation. The scorer explicitly leaves 3-year ROIIC, 5-year ROIIC, acquisition-inclusive ROIIC, reinvestment rate, revenue CAGR, NOPAT CAGR, FCF CAGR, FCF conversion, and share count CAGR as null in the JSON. The markdown says exact ROIIC is insufficient evidence. Under CSA, that means the ROIIC subscore cannot be partially credited. Directional statements about enormous incremental NOPAT on low conventional capex are useful context, not a calculation.

Incremental Returns on Capital is reduced from 20 / 25 to 14 / 25. Normalized ROIC is well supported by the FY2022-FY2026 ROIC row and FY2026 FMP ROIC of 62.9%, but ROIIC receives zero. The "ROIC is not an accounting artifact" point is also capped because NVIDIA is fabless and the real capital base includes R&D, supply commitments, purchase obligations, ecosystem support, HBM/CoWoS access, and customer-facing commitments that are not captured by simple capex.

Reinvestment Runway is reduced from 18 / 20 to 16 / 20. The evidence supports a large runway in AI infrastructure, networking, inference, sovereign AI, enterprise AI, simulation, and physical AI. But the base case still depends on Compute & Networking revenue compounding from $193.479B to $553.6B at a 23.4% segment CAGR while preserving a 62.0% segment margin. The red-team appendix identifies this as the single fragile assumption. Customer AI capex ROI, hyperscaler digestion, internal ASIC substitution, and ecosystem-financed demand are not resolved enough for near-full credit.

Durability of Moat is reduced from 12 / 15 to 11 / 15. CUDA, networking, systems integration, supply access, and high-trust procurement are real. The reduction reflects the moat appendix's own 2046 YELLOW result, lack of contractual lock-in, customer purchase-order flexibility, hyperscaler bargaining power, export controls, and the fact that formal patents expire before or by the 2046 horizon.

Accounting and Financial Quality is reduced from 9 / 10 to 8 / 10. Historical statement tie-outs and FCF conversion are strong: FY2024-FY2026 FCF / net income was 90.8%, 83.5%, and 80.5%. SBC is treated as real in the synthesis. The point withheld is for missing product-level operating margin evidence, customer-level economics, and strategic ecosystem investment transparency.

Capital Allocation is reduced from 10 / 10 to 8 / 10. The evidence supports founder-led R&D, CUDA, Mellanox/networking, supply investment, and net cash. Full credit is not supported because $40.1B of FY2026 buybacks at a multi-trillion-dollar valuation are not clearly opportunistic, and recent Anthropic/OpenAI/Groq-type ecosystem investments create circular-demand monitoring risk.

Valuation / 20% Return Underwriting is reduced from 6 / 20 to 3 / 20. The base case gives $12.28 horizon EPS, or $184.21 at 15x, below the $208.27 starting price. The bull case gives $17.45 EPS, or $261.71 at 15x, only about 4.7% annualized price return. The 20% case plainly fails without a much lower entry price or a much higher terminal multiple. Cash returns help only modestly because the synthesis models net share count reduction of only about 3.2% over five years after SBC issuance.

3. Unsupported Claims

The evidence does not support a quantified high ROIIC score. The scorer says ROIIC is "directionally exceptional," but the CSA rubric requires a credible calculation. It cannot be enough that operating income exploded while capex stayed low, because this business is fabless and customer/supply commitments and R&D are economically material.

The evidence does not fully support near-full runway credit from AI capex alone. Management's "compute equals revenues" framing is an important claim, not proof. The ER2 files repeatedly warn that customer ROI, capex digestion, and circular ecosystem financing must be monitored. The model does not have customer-level capex payback data or a concentration table robust enough to prove the largest buyers can keep growing purchases at the modeled rate.

The evidence does not support full capital-allocation credit. Mellanox and CUDA are excellent evidence. Recent ecosystem investments may be rational, but they are not yet proven ROIIC contributors and could blur end-demand quality.

4. Calculation Errors

The original valuation conclusion is directionally right, but the maximum buy price is not clean. The markdown says the strict 20% five-year buy price using base 15x value is about $74 before dividends, then rounds a "more generous diagnostic ceiling" to $95. The JSON reports $95.00 as the max buy price. That is not a single audited calculation. Using $184.21 of five-year base terminal value and a 20% required return gives $184.21 / 1.2^5 = about $74.04 before dividends. A $95 buy price implies a materially lower required annual return or extra cash-return assumptions that are not bridged.

The base-case IRR is also generous relative to the 15x base value. $184.21 from $208.27 over five years is about -2.4% annualized before dividends, while the JSON reports base-case IRR of +2.5%. If dividends, buybacks, or retained net cash explain the gap, the scorer does not present the bridge. The bull case at $261.71 from $208.27 is about 4.7% annualized, which matches the JSON and still fails badly.

Current EV / NOPAT appears arithmetically sound: about $5.009T EV divided by $110.7B FY2026 adjusted NOPAT equals roughly 45.2x. FY2026 SBC / revenue of $6.386B / $215.938B equals 3.0%; SBC / FCF equals 6.6%. FY2026 FCF conversion of $96.676B / $120.067B equals 80.5%. These checks do not rescue the missing ROIIC and 20% underwriting gaps.

5. Missing Data

The most important missing evidence is a normalized reinvestment bridge. A proper ROIIC denominator should include net capex, working capital, capitalized R&D, growth opex, supply commitments, purchase obligations, acquisitions, ecosystem investments, and customer-support investments. The scorer does not calculate this.

Other missing data: organic versus acquisition-inclusive ROIIC, product-level margins for Data Center Compute, Networking, CUDA/software, Blackwell/Rubin systems, Gaming, Pro Viz, and Automotive; customer-level concentration economics; customer AI capex ROI; standalone revenue for Hopper/Blackwell/Rubin/CUDA/DRIVE/robotics; and a reconciled 20% max-buy-price bridge.

6. Hard-Veto and Threshold Test Audit

The threshold tests do not support "Investable now." Total score is below 80, audited Business Quality is below 65, audited Valuation / 20% UW is below 15, base-case IRR is far below 20%, bear-case IRR is negative, and the 20% case requires either a much lower entry price or terminal valuation help. The original "no hard vetoes" test is false, and the 20% multiple-expansion test is false.

7. Final 20% Hurdle Pass/Fail

Final 20% hurdle: FAIL. At $208.27, the reviewed ER2 base case fails even the stricter 15x EPS gate, and the 20% annualized return case is not close. A defensible 20% entry price is closer to the low-$70s before dividends using base 15x value, not $208.27.

8. Dream Business / Dream Investment Classification

Classification: dream business, not dream investment. NVIDIA has unusually strong current economics, management quality, and platform position. It is not a dream investment under this rubric because the entry price already capitalizes too much of the five-year base case, while the most important reinvestment and customer concentration assumptions remain unproven.

9. Stage 7 Scorecard Notes

Compounder Score note: 63/100: exceptional current business economics, but audited down for missing ROIIC evidence and a failed 20% return case.

Business Quality Score note: 60/80: high ROIC, cash conversion, management, and moat; capped by uncalculated ROIIC, AI capex concentration, circular-demand risk, and 2046 durability uncertainty.

20% Return UW Score note: 3/20: base, bear, and bull 15x cases do not approach 20% annualized returns at $208.27; a much lower entry price is required.

10. Auditor JSON

See `compounder_score_audit.json`.

Appendix RT - Red Team

NVIDIA Corporation - Red Team

Date: 2026-04-26 Ticker: NVDA TTM anchor: FY2026 / TTM ended 2026-01-25 Forward horizon: 2031-01-25, exactly five calendar years from the TTM end Synthesis base case under attack: revenue $589.241B, operating income $343.876B, net income $290.536B, EPS $12.28 Valuation gate: required EPS for P/E 15 at $208.27 is $13.88; base-case cushion is -11.6%; synthesis reported 11/11 tie-outs passed

1. The Most Fragile Assumption

The fragile assumption is not Graphics, not the tax rate, and not the buyback path. It is the Compute & Networking AI compute/platform line: $159.754 billion of TTM base revenue compounded at a 23.0% CAGR to $449.8 billion by 2031-01-25, with the segment as a whole reaching $553.6 billion of revenue and a 62.0% operating margin. That line carries the company. In the synthesis, Compute & Networking contributes 93.9% of horizon segment operating income. If this line is wrong, there is no secondary segment large enough to rescue the P/E 15 test.

The assumption is fragile because the model asks a hardware-linked platform, already at $193.479 billion of Compute & Networking revenue, to add roughly $360 billion of reportable segment revenue in five years while keeping segment margin only modestly below the FY2026 level. The historical growth rates make the 23.0% AI compute/platform CAGR look conservative on paper, because the FY2024-FY2026 AI surge was much higher. That is the wrong comparison. The right comparison is the marginal buyer's next dollar of AI infrastructure spend. The model assumes customers keep scaling NVIDIA systems because deployed compute quickly monetizes into products, productivity, or model capability. The red-team question is whether the largest buyers increasingly treat NVIDIA as the fastest path to capacity today but not the cheapest path for repeated, stable inference workloads tomorrow.

The peer evidence is already moving in the direction that would break the assumption. Broadcom reported Q1 FY2026 AI revenue of $8.4 billion, up 106% year over year, and guided Q2 FY2026 AI semiconductor revenue to $10.7 billion, driven by custom accelerators and AI networking. Broadcom also announced an extended Meta partnership around multi-gigawatt MTIA custom silicon and previously announced a 10-gigawatt OpenAI custom accelerator collaboration targeted to begin deployment in the second half of 2026 and complete by the end of 2029. Those are not generic "competition" data points. They map directly onto the workloads the NVIDIA base case needs to keep: large-scale AI inference, scale-up/scale-out networking, and repeat deployments by the largest customers.

AMD is less threatening in dollars today but still relevant as a price umbrella reducer. AMD reported Q4 2025 Data Center revenue of $5.4 billion, up 39% year over year, driven partly by Instinct GPU ramp. A merchant GPU competitor does not need to match NVIDIA's platform to hurt NVIDIA's margin. It only needs to be credible enough that hyperscalers and sovereign customers can negotiate harder on the next cluster. The bear mechanism is not "NVIDIA loses AI." The bear mechanism is that NVIDIA remains the default premium platform while the marginal workload mix shifts toward standardized inference, internal ASICs, Ethernet fabrics, and second-source GPUs, lowering both the revenue CAGR and the operating margin from supercycle levels.

The synthesis sensitivity table is useful but probably too gentle. It shows that a 200 basis point downside to Compute & Networking CAGR reduces EPS from $12.28 to $11.33. That is already a $0.95 EPS hit and deepens the base-case valuation gap. But the real bear case is not a neat 200 basis point miss. If the AI compute/platform line compounds at 18% instead of 23%, the line reaches roughly $365.6 billion rather than $449.8 billion, an $84 billion revenue shortfall. At the base segment margin, that alone is roughly $44 billion after tax, or about $1.85 per share using the synthesis share count. That would take base EPS toward $10.4 before considering second-order margin pressure. If the miss happens because customers are substituting internal silicon and negotiating harder, margin does not stay at 62.0%; it falls at the same time. That is how the base case becomes the bear case quickly.

The most dangerous part of the assumption is the combination, not either leg alone. A 23% five-year CAGR might be acceptable if margin normalized hard. A 62% horizon segment margin might be acceptable if growth slowed to a premium semiconductor rate. The model needs both. That is the underwriting stress.

2. The Most-Likely-To-Be-Optimistic Stage 5 Number

The single Stage 5 number most likely to be revised down is base-case consolidated operating income of $343.876 billion, specifically the implied 58.4% consolidated operating margin. Revenue could still surprise upward and this number could still be too high if the next phase of AI infrastructure is more systems-heavy, customer-financed, power-constrained, and price-negotiated than the FY2024-FY2026 scarcity phase.

The synthesis is not sloppy. It explicitly bridges segment operating income to consolidated operating income, grows the unallocated bridge, and does not freeze corporate costs. The issue is that the bridge is not the main risk. The main risk is that Compute & Networking's 62.0% horizon segment margin assumes NVIDIA remains a platform toll road at nearly $554 billion of segment revenue. The moat appendix supports strong five-year pricing power, but it also says pricing is stronger in product-market fit and scarce supply than in contractual lock-in. NVIDIA's 10-K-style risk language, as summarized in the moat work, says purchase orders can be changed or delayed and supply commitments can mismatch demand. That means a customer digestion phase could hit both revenue and margin faster than a traditional backlog model would suggest.

Management's capital allocation appendix also raises circularity risk. NVIDIA has large ecosystem and strategic investments, including AI lab or partner funding lines discussed in the management review. These may be rational ecosystem reinforcement. They may also blur the quality of demand if investment dollars, cloud commitments, model-company capacity purchases, and NVIDIA orders become a circular financing loop. Microsoft disclosed FY26 Q2 capital expenditures of $37.5 billion, with roughly two-thirds in short-lived assets such as GPUs and CPUs, while free cash flow was $5.9 billion in the quarter after heavy cash capex. Microsoft also said customer demand exceeded supply. That supports near-term NVIDIA demand, but it also tells the red team where the stress will emerge: if AI monetization is not visible enough to fund continuing short-lived asset purchases, capex budgets can become the constraint.

For that reason, horizon OI is more vulnerable than horizon revenue. A world where NVIDIA sells $520-$560 billion of revenue but at a 50%-53% consolidated operating margin is more plausible than the model admits. At $589.241 billion of base revenue, every 500 basis points of operating margin is about $29.5 billion of operating income. After tax, that is about $24.8 billion, or roughly $1.05 per share. The base case is only $1.60 short of the P/E 15 EPS requirement, so margin error alone can determine the verdict.

3. The Strongest Bear Case, Stated In Its Strongest Form

The strongest bear case is that NVIDIA is not overearning because AI demand is fake; it is overearning because the first stage of AI infrastructure had no credible substitute at scale. That changes between 2026 and 2029 as hyperscaler ASIC programs, Broadcom custom XPUs, Ethernet AI fabrics, AMD merchant accelerators, and internal software-portability work convert the largest buyers from desperate capacity takers into price-setting portfolio buyers. NVIDIA still grows, but the incremental dollar shifts from emergency Blackwell/Rubin procurement to optimized inference capacity where customers own the workload, the model architecture, and the utilization curve. The first visible damage is not a revenue decline; it is Data Center or Compute & Networking growth that remains positive while gross and segment operating margins compress.

In that version of the world, the synthesis is too bullish because it treats the FY2026 platform scarcity margin as mostly durable through 2031. The bear would argue that NVIDIA's base case is really a bull case disguised by recent historical CAGRs. A short seller would not need to claim CUDA disappears. The claim would be sharper: CUDA remains valuable for frontier training and heterogeneous workloads, but the biggest cloud and model-company buyers migrate stable inference and high-volume internal workloads to custom silicon quickly enough to cap NVIDIA's growth and normalize margins before the five-year EPS gate can be cleared.

The observable anchor validating this bear case is two consecutive quarters where NVIDIA guides or reports sequential Data Center growth below 5% while gross margin falls below 70% and Broadcom AI semiconductor revenue continues to grow above 50% year over year. That combination would show that AI spending is continuing, but NVIDIA's share of incremental economics is weakening.

4. Cross-Examination Of Management Commentary

First statement: Colette Kress said in Q4 FY2026 that NVIDIA expected "sequential revenue growth throughout calendar 2026." The base model uses this as support for a strong FY2027 starting point. The skeptical read is that sequential growth through a supply catch-up year does not prove five-year demand elasticity. If Blackwell supply, networking attach, and power-constrained data centers pull orders forward into 2026, the 2027 comparison could be much harder. The first sign this comment is aging badly would be sequential Data Center deceleration in the second half of calendar 2026 despite nominally strong hyperscaler capex budgets.

Second statement: Jensen Huang's "compute equals revenues" is the central customer-ROI claim. It is powerful, but it is also the thesis's weakest proof point because it compresses very different buyers into one causal claim. Search, ads, coding, enterprise copilots, consumer assistants, model APIs, robotics, and sovereign AI do not all monetize compute with the same speed or margin. Microsoft's FY26 Q2 disclosure shows the tension: Azure demand was strong, but AI infrastructure investment pressured cloud gross margin, and the quarter involved very large capex. If customer ROI is slower than GPU depreciation, the first signal will be more explicit cloud gross margin pressure and more selective language around AI capex prioritization from Microsoft, Alphabet, Meta, and Amazon.

Third statement: Kress said "Every data center is power-constrained." The bullish interpretation is that performance per watt lets NVIDIA price like a scarce resource. The skeptical interpretation is that power scarcity eventually pushes customers toward workload-specific silicon, utilization optimization, lower-precision inference, model compression, and geographic/sovereign alternatives. Power scarcity helps NVIDIA while the best available platform is NVIDIA. It hurts NVIDIA if buyers conclude the only way to make AI economics work is to own more of the silicon stack and avoid paying a premium merchant margin on every watt. The first contrary signal would be hyperscalers announcing multi-year custom silicon deployments in gigawatt units while slowing incremental merchant GPU commitments.

5. Stage 5 Tie-Out Reverification

The synthesis passes the arithmetic checks. The TTM anchor is correct: Q1 FY2026 through Q4 FY2026, ending 2026-01-25, and the horizon is 2031-01-25. The red team does not challenge the 11/11 tie-outs. The problem is that clean tie-outs can still compound the wrong economics.

The recency anchor is defensible but double-edged. The synthesis notes Q4 FY2026 revenue annualized at $272.508 billion and a first forward year near $285.6 billion, only 4.8% above the Q4 run rate. That makes the model look conservative. But if Q4 was a supply-constrained, architecture-transition, catch-up quarter, annualizing it may embed a high watermark. The model's first-year guardrail prevents an obvious step-up error; it does not prove the exit rate is sustainable.

The share count path is realistic enough not to be the central attack. The model assumes $45 billion of annual gross repurchases, rising repurchase prices, and SBC issuance. That only reduces shares by 3.2% over five years. The red team would not spend time attacking this because even a less favorable share path is secondary to Compute & Networking economics. The bigger capital allocation concern is not whether buybacks retire enough shares; it is whether continuing supply commitments, ecosystem investments, and customer financing reduce the amount of free cash that can be returned without weakening the platform.

The product CAGR pattern is where the model is vulnerable. Compute & Networking's base growth is below the recent historical rate, so it passes a mechanical plausibility check. But recent history was an adoption shock from a small base relative to the new base. Broadcom's custom accelerator data points and Meta/OpenAI deployment announcements are a named structural reason to question whether NVIDIA can compound the AI compute/platform line at 23% while also keeping segment margins near 62%. A model can be "below historical" and still too high if the market structure is changing.

6. The Signal That Forces An Exit

The exit signal should be one combined metric, not a vague checklist: exit if, over any two consecutive reported quarters before the 2027 investor day, NVIDIA's Data Center or Compute & Networking revenue grows less than 5% sequentially on average, non-GAAP or GAAP gross margin falls below 70%, and management continues to exclude China Data Center compute from guidance without offsetting growth elsewhere.

The threshold matters. Less than 5% sequential growth would mean the company is no longer compounding from the Q4 FY2026 exit rate fast enough to support the base-case revenue path. Gross margin below 70% would indicate that the issue is not only timing; it would point to mix, pricing, systems content, customer leverage, or supply-commitment costs. The China condition matters because the base case already excludes China recovery. If China remains absent and the rest of the world slows, the model cannot invoke China as a hidden buffer.

This signal invalidates the thesis because the Stage 5 result already fails the P/E 15 gate at $12.28 of EPS. The stock cannot afford a clean but slightly weaker base case. It needs either the base case to be revised upward or the price to fall toward the $184 support implied by 15x base EPS. If the exit signal appears, the red team would rewrite the narrative around normalization: NVIDIA remains exceptional, but the AI infrastructure profit pool is being shared earlier and at lower margin than the synthesis requires.

7. Conviction At The End Of Red Team

Conclusion: conditional. The business-quality thesis survives; the valuation thesis does not get strengthened by this red team. NVIDIA still has a real five-year moat, excellent management, a clean balance sheet, and unusually strong demand evidence. The red team is not arguing for a broken company. It is arguing that the specific base case is too close to the valuation line and too dependent on one assumption: Compute & Networking remaining both hyper-growth and hyper-margin at unprecedented scale.

The red-team view is that the most likely optimistic number is operating income, not revenue. The company may continue to print astonishing revenue while the market gradually discovers that the economics of AI infrastructure are moving from scarcity pricing to portfolio procurement. Broadcom's custom AI revenue ramp, announced custom accelerator deployments with OpenAI and Meta, AMD's Instinct ramp, and hyperscaler capex scrutiny all point to the same pressure point. None of them kills NVIDIA. Together, they make a 58.4% consolidated operating margin in 2031 less robust than the synthesis implies.

The appendix should therefore feed a PASS or WATCH posture at $208.27. The base-case arithmetic already fails by 11.6%. The red-team adjustment makes that gap qualitatively worse, not better. A BUY would require either a lower entry price, evidence that custom silicon is expanding the AI market without touching NVIDIA's economics, or proof that Rubin-era systems and networking attach can keep Compute & Networking growth above the base while sustaining margins despite customer concentration.

Key Red-Team Findings

- Strongest bear argument: NVIDIA remains the AI infrastructure leader but loses the marginal economics of stable inference and hyperscaler scale-out to custom silicon, Ethernet fabrics, and second-source GPUs before 2031.
- Most likely optimistic Stage 5 number: $343.876 billion of base-case operating income, because the 58.4% consolidated operating margin is more fragile than the revenue line.
- Most fragile assumption: Compute & Networking AI compute/platform revenue compounding at 23.0% to $449.8 billion while the segment retains a 62.0% margin.
- Exit signal: two consecutive quarters with Data Center or Compute & Networking sequential revenue growth below 5% on average and gross margin below 70%, with China still excluded from guidance.
- Valuation implication: because base EPS is already $12.28 versus $13.88 required, even a moderate C&N or margin miss keeps the stock outside the ER2 BUY gate.

External Sources Used

- Broadcom Q1 FY2026 results release: https://investors.broadcom.com/node/63976/pdf
- Broadcom and OpenAI custom accelerator collaboration: https://investors.broadcom.com/news-releases/news-release-details/openai-and-broadcom-announce-strategic-collaboration-deploy-10
- Broadcom and Meta MTIA partnership announcement: https://investors.broadcom.com/news-releases/news-release-details/broadcom-announces-extended-partnership-meta-deploy-technology
- AMD Q4 and FY2025 results release: https://www.amd.com/en/newsroom/press-releases/2026-2-3-amd-reports-fourth-quarter-and-full-year-2025-fina.html
- Microsoft FY2026 Q2 earnings call transcript: https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q2

Appendix D - Data Package (raw historicals)

NVIDIA Corporation (NVDA) - ER2 Data Package

Prepared: 2026-04-26 Sub-Agent: D Output path: `C:\Users\Admin\OneDrive\GPT\ER 2.0\NVIDIA Corporation\data_package.md` Quantitative source of truth: `C:\Users\Admin\OneDrive\GPT\ER 2.0\NVIDIA Corporation\quant_foundation.md` Raw D pulls: `C:\Users\Admin\OneDrive\GPT\ER 2.0\NVIDIA Corporation\raw\d\`

0. Workflow, source discipline, and fiscal anchor

This package follows the ER2 data source rule: live sources only; no local data warehouse cache; no transcript, financial, segment, or price file was read from `C:\Users\Admin\OneDrive\Cowork\Investing\data warehouse` or any equivalent cached warehouse path. Financial Modeling Prep was called live through `C:\Users\Admin\.codex\skills\fmp-financial-data\scripts\fmp_request.py`; NVIDIA IR pages, SEC EDGAR filings, SEC submissions, and SEC companyfacts were fetched live in this run and saved under `raw\d`. API keys were not printed or written. The quantitative base is the already-cleared QF package, which reports `RECONCILIATION: 8/8 checks` PASSED; this D package does not re-anchor or replace QF arithmetic.

Fiscal anchor: NVIDIA's fiscal year ends on the last Sunday in January. The most recent completed and reported quarter is Q4 FY2026, covering the quarter ended 2026-01-25 and reported on 2026-02-25. TTM is Q1 FY2026 ended 2025-04-27 + Q2 FY2026 ended 2025-07-27 + Q3 FY2026 ended 2025-10-26 + Q4 FY2026 ended 2026-01-25, which equals FY2026 because the latest period is fiscal year-end. Forward horizon for downstream ER2 work is 2031-01-25, exactly five calendar years after the TTM end.

Primary live sources used in this D run:

Source	Live endpoint or URL	Raw file
FMP annual income statement	`income-statement NVDA --limit 6 --param period=annual`	`raw\d\fmp_income_annual_limit6.json`
FMP annual cash flow	`cash-flow-statement NVDA --limit 6 --param period=annual`	`raw\d\fmp_cashflow_annual_limit6.json`
FMP annual balance sheet	`balance-sheet-statement NVDA --limit 6 --param period=annual`	`raw\d\fmp_balance_annual_limit6.json`
FMP annual key metrics	`key-metrics NVDA --limit 6 --param period=annual`	`raw\d\fmp_key_metrics_annual_limit6.json`
FMP product/end-market segmentation	`revenue-product-segmentation NVDA --param period=annual/quarter`	`raw\d\fmp_revenue_product_segmentation_*.json`
FMP FY2026 transcripts	`earning-call-transcript NVDA year=2026 quarter=1..4`	`raw\d\fmp_transcript_fy2026_q*.json`
NVIDIA FY2026 quarterly releases	NVIDIA Newsroom FY2026 Q1-Q4 results pages	`raw\d\ir_q*fy2026_press_release.html`
SEC FY2026 10-K	`https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm`	`raw\d\sec_10k_fy2026.html`
SEC FY2026 10-Qs	Q1 accession `0001045810-25-000116`, Q2 `0001045810-25-000209`, Q3 `0001045810-25-000230`	`raw\d\sec_10q_q*fy2026.html`
SEC metadata	SEC submissions and companyfacts for CIK `0001045810`	`raw\d\sec_submissions_CIK0001045810.json`, `raw\d\sec_companyfacts_CIK0001045810.json`

1. Reporting structure and taxonomy

Current reportable segments from QF are: **Compute & Networking** and **Graphics**. Compute & Networking includes data center accelerated computing and networking platforms, AI solutions and software, and automotive platforms including autonomous and electric vehicle solutions. Graphics includes GeForce GPUs for gaming and PCs plus NVIDIA RTX/Quadro products for enterprise workstation graphics. Source: FY2026 10-K Item 1 and QF P-Taxonomy.

NVIDIA separately discloses end-market revenue lines: **Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other**. These end-market lines are useful for demand mix and management commentary, but they are **not a clean mechanical bridge to reportable segments**. QF explicitly flags the break: FY2026 Data Center end-market revenue was $193.737B while FY2026 Compute & Networking reportable-segment revenue was $193.479B. Therefore downstream subagents should model reportable segment revenue and operating income from QF P2/P4, and use end-market data only as product-demand context.

FY2026 reportable segment revenue and operating income from QF P2:

Reportable segment	FY2026 revenue ($B)	% of total revenue	FY2026 operating income ($B)	Segment operating margin	Source
Compute & Networking	193.479	89.6%	130.141	67.3%	QF P2; SEC FY2026 10-K Note 16
Graphics	22.459	10.4%	9.156	40.8%	QF P2; SEC FY2026 10-K Note 16
Total segment	215.938	100.0%	139.297	64.5%	QF P2

FY2026 end-market revenue from FMP live product segmentation and SEC FY2026 10-K "Revenue by End Market":

End-market disclosure line	FY2026 revenue ($B)	% of total revenue	Segment-use flag
Data Center	193.737	89.7%	Product/end-market line; do not force to Compute & Networking
Gaming	16.042	7.4%	Product/end-market line; main Graphics demand driver
Professional Visualization	3.191	1.5%	Product/end-market line; Graphics demand driver
Automotive	2.349	1.1%	Product/end-market line; not a clean segment bridge
OEM and Other	0.619	0.3%	Product/end-market line; residual
Total revenue	215.938	100.0%	Ties to consolidated revenue

Taxonomy flags:

- No reportable segment change was identified in the last three fiscal years. The FY2026, FY2025, and FY2024 10-K segment notes use the same two reportable segments: Compute & Networking and Graphics.
- End-market disclosure expanded inside Data Center in FY2026 with Compute and Networking sub-lines in the SEC end-market table. For FY2026, Data Center was $193.737B, of which Compute was $162.361B and Networking was $31.376B. FY2025 Data Center was $115.186B, with Compute $102.196B and Networking $12.990B. FY2024 Data Center was $47.525B, with Compute $38.950B and Networking $8.575B. This is a product/end-market split, not reportable segment operating income.
- Geographic disclosure in the FY2026 10-K states revenue by geography is based on direct customer headquarters, not end-customer location or shipping destination. Revenue from customers headquartered outside the United States was 31% of FY2026 revenue and 41% of FY2025 revenue. This geography is distorted by centralized purchasing and invoicing and should not be read as physical deployment by country.
- Export-control exposure is a live taxonomy/business-risk overlay. Q1 FY2026 and FY2026 commentary separately discusses H20 China restrictions; do not treat China-compliant accelerator revenue as ordinary run-rate Data Center demand.

2. Consolidated historicals, five fiscal years

All dollar figures are USD billions except EPS, share count, and percentages. FMP annual statement rows were live-fetched in this run and tie to QF P3. Capex is cash flow statement "capital expenditure" / "purchases related to property and equipment and intangible assets." FCF is operating cash flow less capex. ROIC is FMP `returnOnInvestedCapital`; QF Stage 1 cited FY2026 ROIC of 62.9%.

Metric	FY2022	FY2023	FY2024	FY2025	FY2026	Source
Revenue ($B)	26.914	26.974	60.922	130.497	215.938	FMP annual income; QF P3
Operating income ($B)	10.041	4.224	32.972	81.453	130.387	FMP annual income; QF P3
Operating margin (%)	37.3%	15.7%	54.1%	62.4%	60.4%	Calculated
Net income ($B)	9.752	4.368	29.760	72.880	120.067	FMP annual income; QF P3
EPS diluted ($)	0.38	0.17	1.19	2.94	4.90	FMP annual income; QF P3
Diluted share count (M)	25,350	25,070	24,940	24,804	24,432	FMP annual income
Capex ($B)	0.976	1.833	1.069	3.236	6.042	FMP annual cash flow
Free cash flow ($B)	8.132	3.808	27.021	60.853	96.676	FMP annual cash flow; QF P6
ROIC (%)	24.6%	11.7%	51.3%	75.3%	62.9%	FMP key metrics

ROIC calculation note: FMP's returned FY2026 `returnOnInvestedCapital` is 0.6288 and `investedCapital` is $130.830B. Using QF's controlling FY2026 operating income of $130.387B and FY2026 effective tax rate of 15.1% (income tax expense $21.383B / income before tax $141.450B), FY2026 NOPAT is approximately $110.7B. The FMP metric is treated as the standardized ROIC row because it is the QF source-of-truth input used in Stage 1 and avoids mixing ending and average invested capital definitions.

The five-year historical record is a discontinuous growth profile rather than a smooth trend. FY2022 and FY2023 revenue were nearly flat, with FY2023 operating margin down sharply during the gaming correction. FY2024 to FY2026 was the AI data center step-change: revenue grew from $60.922B to $215.938B in two fiscal years, while operating income grew from $32.972B to $130.387B. Downstream forward modeling should therefore separate normalized durability from the AI infrastructure supply/demand supercycle.

3. Segment historicals, five fiscal years

Reportable segment data is from QF P2 and SEC 10-K segment tables. "Other segment items" is the disclosed segment-level revenue less segment operating income bridge item; NVIDIA does not disclose cost of revenue and segment operating expense by reportable segment in a full P&L format.

Compute & Networking

Metric	FY2022	FY2023	FY2024	FY2025	FY2026	Source
Revenue ($B)	11.046	15.068	47.405	116.193	193.479	QF P2; SEC 10-K Note 16
YoY growth	n.m.	36.4%	214.6%	145.1%	66.5%	Calculated
Operating income ($B)	4.598	5.083	32.016	82.875	130.141	QF P2
Operating margin (%)	41.6%	33.7%	67.5%	71.3%	67.3%	Calculated

Compute & Networking is the controlling profit pool. FY2026 revenue was 17.5x FY2022 revenue and FY2026 operating income was 28.3x FY2022 operating income. The margin decline from FY2025 to FY2026 is not a demand weakness signal by itself; FY2026 includes Blackwell full-scale data center systems mix and the H20 inventory/purchase-obligation charge discussed in NVIDIA's FY2026 commentary.

Graphics

Metric	FY2022	FY2023	FY2024	FY2025	FY2026	Source
Revenue ($B)	15.868	11.906	13.517	14.304	22.459	QF P2; SEC 10-K Note 16
YoY growth	n.m.	-25.0%	13.5%	5.8%	57.0%	Calculated
Operating income ($B)	8.492	4.552	5.846	5.085	9.156	QF P2
Operating margin (%)	53.5%	38.2%	43.3%	35.5%	40.8%	Calculated

Graphics remains material but is not the valuation driver versus Data Center. FY2026 Graphics recovered strongly as Gaming and Professional Visualization both benefited from Blackwell/RTX demand, but Graphics segment operating income was only

7% of total segment operating income. Segment B work should model Graphics as cyclical-plus-AI-PC/workstation optionality rather than as the primary AI factory exposure.

4. Product and sub-segment historicals

Dollar-disclosed product/end-market rows

The most granular reliable dollar disclosure for product demand is the end-market table in FMP live product segmentation and SEC FY2026 10-K. For Data Center, the FY2026 10-K also discloses Compute and Networking sub-lines for FY2024-FY2026. These rows are disclosed in dollars; no back-solved dollar estimates are used.

Segment context	Product/end market	What it is	Dollar disclosure?	FY2024 revenue ($B)	FY2025 revenue ($B)	FY2026 revenue ($B)	Source
Compute & Networking context	Data Center	Accelerated compute, AI systems, networking, software/services for data centers	Yes	47.525	115.186	193.737	FMP product segmentation; SEC 10-K
Data Center sub-line	Compute	Data center compute products within Data Center end market	Yes, FY2024-FY2026	38.950	102.196	162.361	SEC FY2026 10-K end-market table
Data Center sub-line	Networking	InfiniBand, Ethernet, NVLink and related networking within Data Center end market	Yes, FY2024-FY2026	8.575	12.990	31.376	SEC FY2026 10-K end-market table
Graphics context	Gaming	GeForce GPUs, gaming laptops/PCs, gaming platform products	Yes	10.447	11.350	16.042	FMP product segmentation; SEC 10-K
Graphics context	Professional Visualization	RTX/Quadro workstation graphics and professional visualization platforms	Yes	1.553	1.878	3.191	FMP product segmentation; SEC 10-K
Compute & Networking context	Automotive	DRIVE, autonomous/electric vehicle platforms and software	Yes	1.091	1.694	2.349	FMP product segmentation; SEC 10-K
Residual	OEM and Other	OEM/residual revenue line	Yes	0.306	0.389	0.619	FMP product segmentation; SEC 10-K

Lines with only growth rates or qualitative disclosure

Segment context	Product/topic	Disclosure method	Last four quarter evidence	Modeling treatment
Data Center	Blackwell / Blackwell Ultra / Rubin	Qualitative and sequential growth comments; no standalone annual product revenue in dollars	Q2 FY2026 IR release: Blackwell Data Center revenue grew 17% sequentially; Q4 transcript: Grace Blackwell systems were roughly two-thirds of Data Center revenue in Q4	Do not create standalone Blackwell annual revenue unless explicitly estimated and labeled; use Data Center total as dollar base
Data Center	Networking	Dollar disclosed annually and quarterly commentary, but not a full segment P&L	Q4 FY2026 transcript: networking generated $11B in Q4 and exceeded $31B for FY2026	Use dollar line for mix; do not infer operating income
Gaming	GeForce RTX / AI PC	Revenue disclosed only at Gaming end-market level	FY2026 quarterly IR releases disclose Gaming revenue and YoY growth	Model Gaming as end-market line inside Graphics context
Professional Visualization	RTX PRO / DGX Spark / workstation AI	Revenue disclosed only at Pro Viz end-market level	FY2026 quarterly releases disclose Pro Viz revenue and YoY growth	No separate product dollars
Automotive / physical AI	DRIVE / robotics / Cosmos / Isaac	Automotive dollars disclosed; broader physical AI revenue claim is qualitative/portfolio-level	Q4 transcript: physical AI contributed north of $6B in FY2026	Keep Automotive dollars separate; do not allocate physical AI total into Automotive

Segment context	Product/topic	Disclosure method	Last four quarter evidence	Modeling treatment
China H20	China-compliant data center GPU	Specific H20 charges and lost revenue discussed; not a continuing product line base	Q1 FY2026 transcript: $4.6B H20 revenue recognized pre-control and $2.5B unable to ship in Q1; Q2 outlook lost about $8B	Treat as export-control disruption, not normalized demand

Open data-integrity note: There are no back-solved product dollar rows in this package. Where NVIDIA provides growth rates without dollar product revenue, the data package records the growth commentary and flags that downstream subagents must not invent product dollars.

5. Most recent quarterly trend

Company FY ends last Sunday in January. Most recently completed fiscal quarter at time of writing: FY2026 Q4, covering 2025-10-27 to 2026-01-25, reported on 2026-02-25 via NVIDIA's Q4 FY2026 press release and SEC 8-K accession `0001045810-26-000019`.

Reportable segments: revenue, operating income, and margin

Quarter	Compute & Networking revenue ($B)	C&N operating margin	Graphics revenue ($B)	Graphics operating margin	Source
Q1 FY2026 ended 2025-04-27	39.589	55.7%	4.473	36.7%	QF P4; SEC Q1 FY2026 10-Q
Q2 FY2026 ended 2025-07-27	41.331	68.6%	5.412	41.4%	QF P4; SEC Q2 FY2026 10-Q
Q3 FY2026 ended 2025-10-26	50.908	70.2%	6.098	41.8%	QF P4; SEC Q3 FY2026 10-Q
Q4 FY2026 ended 2026-01-25	61.651	71.4%	6.476	42.1%	QF P4; Q4 = FY2026 10-K less Q1-Q3
TTM / FY2026	193.479	67.3%	22.459	40.8%	QF P4

End-market quarterly trend and YoY rates

Quarter	Data Center revenue / YoY	Gaming revenue / YoY	Pro Viz revenue / YoY	Automotive revenue / YoY	Source
Q1 FY2026	$39.112B / +73%	$3.763B / +42%	$0.509B / +19%	$0.567B / +72%	NVIDIA Q1 FY2026 release
Q2 FY2026	$41.096B / +56%	$4.287B / +49%	$0.601B / +32%	$0.586B / +69%	NVIDIA Q2 FY2026 release
Q3 FY2026	$51.215B / +66%	$4.265B / +30%	$0.760B / +56%	$0.592B / +32%	NVIDIA Q3 FY2026 release
Q4 FY2026	$62.314B / +75%	$3.727B / +47%	$1.321B / +159%	$0.604B / +6%	NVIDIA Q4 FY2026 release

Q4 FY2026 recency signals for downstream modeling:

- Q4 Data Center accelerated to $62.3B, up 22% sequentially and 75% YoY, and FY2026 Data Center revenue was $193.7B, up 68%.
- Q4 Gaming fell 13% sequentially from Q3 but grew 47% YoY. NVIDIA specifically attributed Q4 sequential moderation to channel inventory after a strong holiday season.
- Q4 Professional Visualization crossed $1B quarterly revenue for the first time, at $1.321B, with 159% YoY growth and 74% sequential growth.
- Q4 Automotive grew only 6% YoY, far below Data Center/Gaming/Pro Viz, but FY2026 Automotive still grew 39% to $2.349B.
- Q1 FY2027 outlook from the Q4 FY2026 release was revenue of $78.0B plus or minus 2%, with no assumed Data Center compute revenue from China.

6. Management commentary on forward growth

Quotes below are from live FMP transcript pulls for FY2026 Q1-Q4. Each quote is short, tied to a fiscal period and actual call date, and used as a modeling anchor rather than a substitute for historical dollars.

Compute & Networking / Data Center

Fiscal period and call date	Speaker	Verbatim quote	Topic	Interpretation
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"we expect sequential revenue growth throughout calendar 2026"	Forward Data Center demand	Management guided ongoing growth beyond the Q4 run-rate; S should test whether this implies a sustainable FY2027 base or a supply-led catch-up.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Every data center is power-constrained."	Customer purchasing constraint	Performance per watt is a pricing-power and share-defense variable, not just a product spec.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Networking, a cornerstone of our data center scale infrastructure offering, was a standout this quarter, generating $11 billion in revenue"	Networking mix	Networking is now large enough to affect Data Center mix and margin, but it remains inside end-market disclosure rather than reportable segment P&L.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"our networking business exceeded $31 billion in revenue"	FY2026 networking scale	Annual networking is a large sub-line and should be discussed in Compute & Networking appendix.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"we are not assuming any data center compute revenue from China"	China/export controls	Q1 FY2027 guidance excludes China compute revenue; the base case should not rely on China recovery.
Q4 FY2026, 2026-02-25	Jensen Huang, CEO	"compute equals revenues"	Customer ROI framing	This is management's strongest demand-duration claim; red team should test customer CapEx affordability and token monetization.
Q3 FY2026, 2025-11-19	Colette Kress, CFO	"Total revenue is expected to be $65 billion plus or minus 2%."	Near-term growth guide	Q4 actual revenue of $68.127B beat that outlook midpoint, supporting the live demand/supply-upside narrative.
Q2 FY2026, 2025-08-27	Colette Kress, CFO	"our Q3 outlook does not include H20 shipments to China customers"	Export-control conservatism	China risk was already excluded from guidance after Q2; do not add China recovery without an explicit assumption.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"Blackwell contributed nearly 70% of data center compute revenue"	Architecture transition	Blackwell became the compute center of gravity early in FY2026; Hopper comparisons age quickly.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"Spectrum X posted strong sequential and year-on-year growth"	Ethernet networking	Ethernet scale-out is a competitive bridge against non-InfiniBand AI factory deployments.

Graphics / Gaming and Professional Visualization

Fiscal period and call date	Speaker	Verbatim quote	Topic	Interpretation
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Gaming revenue of $3.7 billion increased 47% year-on-year"	Gaming demand	Gaming remains healthy, but sequential decline shows channel/inventory cyclicality.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"we expect supply constraints to be the headwind to Gaming"	Gaming supply	The near-term Gaming limit is supply, not stated demand weakness.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"For Professional Visualization, it crossed the $1 billion mark for the first time"	Pro Viz inflection	Pro Viz is now meaningful enough to watch separately inside Graphics.
Q3 FY2026, 2025-11-19	Colette Kress, CFO	"Gaming revenue was $4.3 billion, up 30% year on year"	Gaming quarterly base	Q3 was the FY2026 Gaming peak; Q4 moderation should be modeled as channel normalization.

Fiscal period and call date	Speaker	Verbatim quote	Topic	Interpretation
Q2 FY2026, 2025-08-27	Colette Kress, CFO	"Our gaming revenue was a record $4.3 billion, a 14% sequential increase and a 49% jump year on year"	Gaming momentum	Gaming held near the Q3 level before the Q4 channel reset.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"Gaming revenue was a record $3.8 billion"	Gaming recovery	Blackwell gaming ramp lifted the non-Data-Center business entering FY2026.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"Demand for our AI workstations is strong"	Professional Visualization	Workstation AI demand supports Pro Viz growth beyond legacy visualization cycles.

Automotive, robotics, and physical AI

Fiscal period and call date	Speaker	Verbatim quote	Topic	Interpretation
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Automotive revenue of $604 million was up 6% year-over-year"	Automotive deceleration	Q4 Automotive growth slowed; do not over-extrapolate physical AI commentary into current Automotive dollars.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Physical AI is here"	Physical AI portfolio	Management frames robotics/AV as a broader platform wave beyond current Automotive disclosure.
Q4 FY2026, 2026-02-25	Colette Kress, CFO	"Robotaxi rides are growing exponentially"	AV demand	Long-horizon optionality is large but current disclosed Automotive revenue is small.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"The era of robotics is here."	Robotics theme	This supports strategic framing, not a standalone current revenue line.
Q1 FY2026, 2025-05-28	Colette Kress, CFO	"with our automotive group, revenue was $567 million"	Automotive quarterly base	Automotive started FY2026 at $567M and ended at $604M; sequential growth was modest.

Management-commentary synthesis for downstream subagents: Data Center commentary is unusually strong and concrete, with sequential CY2026 growth, power constraints, explicit Q1 FY2027 revenue guidance, and large networking dollars. Graphics commentary is positive but more cyclical, with supply/channel language. Automotive and robotics commentary is strategically ambitious, but disclosed Automotive revenue remains only about 1% of FY2026 company revenue.

7. Capital allocation history

All dollar figures are USD billions. R&D is from FMP annual income statement; capex, buybacks, dividends, and acquisitions are from FMP annual cash flow. Buybacks, dividends, and M&A are shown as positive cash outflows for readability.

Year	Capex	Capex / revenue	R&D	R&D / revenue	Buybacks	Dividends	M&A	Source
FY2022	0.976	3.6%	5.268	19.6%	0.000	0.399	0.263	FMP income/cash flow
FY2023	1.833	6.8%	7.339	27.2%	10.039	0.398	0.049	FMP income/cash flow
FY2024	1.069	1.8%	8.675	14.2%	9.533	0.395	0.083	FMP income/cash flow
FY2025	3.236	2.5%	12.914	9.9%	33.706	0.834	1.007	FMP income/cash flow
FY2026	6.042	2.8%	18.497	8.6%	40.086	0.974	1.535	FMP income/cash flow

Additional capital-allocation notes:

- FY2026 FCF was $96.676B. NVIDIA returned $41.060B through repurchases and dividends, or about 42.5% of FCF, consistent with Q4 transcript language that the company returned $41B / 43% of FCF.
- FY2026 R&D reached $18.497B, up 43% YoY, but R&D as a percentage of revenue fell to 8.6% because revenue scaled faster than expense.

- The FY2026 10-K includes much larger non-marketable equity purchases and strategic ecosystem investments than standard M&A cash flow alone. Downstream management/moat work should inspect whether these investments support ecosystem durability or create circular demand risk.

8. Balance sheet snapshot, FY2026

All dollar figures are USD billions unless stated otherwise. Source: live FMP annual balance sheet and QF P6.

Item	FY2026 value	Source
Cash and equivalents	10.605	FMP balance sheet; QF P6
Short-term investments	51.951	FMP balance sheet; QF P6
Cash plus short-term investments	62.556	FMP balance sheet; QF P6
Total debt	11.412	FMP balance sheet; QF P6
Net cash, cash plus STI less debt	51.144	Calculated from FMP; QF P6
FMP cash-only net debt	0.807	FMP balance sheet; QF P6
FY2026 EBITDA	144.552	FMP income statement
Cash-plus-STI net debt / EBITDA	-0.35x	Calculated
FMP cash-only net debt / EBITDA	0.006x	QF Stage 1
Goodwill	20.832	FMP balance sheet
Intangible assets	3.306	FMP balance sheet
Goodwill + intangibles	24.138	FMP balance sheet
Tangible book value	133.155	FMP key metrics
Stockholders' equity	157.293	FMP balance sheet

The balance sheet does not constrain near-term strategy. NVIDIA ended FY2026 with more than $62B of cash plus short-term investments and roughly $51B of net cash after total debt. The more relevant constraint is supply chain commitment duration, customer CapEx health, export controls, and power/data-center deployment capacity.

9. Recent corporate actions and risk events, last 12 months

- 2025-04-09: U.S. export controls on H20 created an immediate FY2026 disruption. Q1 FY2026 transcript says NVIDIA recognized $4.6B of H20 revenue before the control date, recorded a $4.5B inventory/purchase-obligation charge, and was unable to ship $2.5B of H20 revenue in Q1. Q2 FY2026 outlook reflected about $8B of lost H20 revenue.
- 2025-05-28 through 2026-02-25: NVIDIA repeatedly excluded China Data Center compute revenue from forward outlooks. This is a live downside/upside swing factor but not part of the base guide.
- 2026-02-25: FY2026 10-K filed and Q4 FY2026 results released. FY2026 revenue was $215.938B, up 65%, with Q4 revenue $68.127B, up 73% YoY.
- 2026-02-25: NVIDIA changed non-GAAP presentation beginning Q1 FY2027 to include stock-based compensation expense in non-GAAP results, per Q4 FY2026 transcript. This affects comparison to prior non-GAAP margin conventions.
- FY2026: NVIDIA repurchased $40.086B of common stock and paid $0.974B of dividends, while FCF was $96.676B.
- FY2026: Strategic and ecosystem investments increased. Q4 FY2026 transcript references a $10B Anthropic investment; the FY2026 10-K also records a large Groq-related nonexclusive license agreement/payment line in investing activities. These items should be red-teamed for circular AI ecosystem risk, even though they may also strengthen platform demand.
- No CEO or CFO transition was identified in live SEC filings or FY2026 earnings materials. Jensen Huang remained President and CEO; Colette Kress remained CFO.

10. Open data gaps and downstream instructions

- Do not map Data Center, Gaming, Professional Visualization, Automotive, and OEM and Other mechanically into the two reportable segments. Use QF P2/P4 for reportable segment revenue and operating income. Use end-market data for demand commentary and product mix only.
- NVIDIA does not disclose standalone annual dollar revenue for Blackwell, Hopper, Rubin, GeForce RTX by generation, DRIVE, Omniverse, CUDA software, DGX Spark/Station, or robotics/physical AI products. Any downstream product-level dollar build for these items must be explicitly labeled as an estimate and tied back to disclosed Data Center/Gaming/Pro Viz/Automotive totals.
- Data Center Compute and Networking sub-lines are disclosed for FY2024-FY2026 in the FY2026 10-K end-market table. They do not have a standalone operating income table.
- China/H20 revenue disruption is explicit, but the amount of normalized China demand not served after the controls is not a clean disclosed run-rate. Use management's excluded-China guidance as the base case.
- Geographic revenue based on direct customer headquarters is not physical deployment geography. Singapore and other invoicing locations can distort end-market interpretation; the FY2026 10-K states geography is based on direct customer headquarters.
- The latest full official quarter is Q4 FY2026 ended 2026-01-25 and reported 2026-02-25. Do not substitute FY2027 Q1 estimates before the official May 2026 report.

11. Self-verification log

- Largest segment recency check: Q4 FY2026 Data Center revenue was $62.3B, up 22% sequentially and 75% YoY. This appears in the live-fetched NVIDIA Q4 FY2026 release saved as `raw\d\ir_q4fy2026_press_release.html`.
- Transcript date check: the live FMP FY2026 Q4 transcript response has `date: 2026-02-25`, matching the expected Q4 FY2026 call date. The quote "we expect sequential revenue growth throughout calendar 2026" appears in that response.
- FY2026 annual revenue check: live FMP annual income statement reports FY2026 revenue of $215.938B for fiscal year ended 2026-01-25, matching QF P3 and NVIDIA's FY2026 10-K/press-release total within 1%.

This package is production-mode complete and intentionally table-heavy. The controlling historical financials are QF P-sections; this file adds source inventory, taxonomy guardrails, live transcript/IR commentary, capital allocation, balance-sheet context, corporate actions, and explicit open data gaps for downstream ER2 subagents.