

EQUITY VALUATION · SEMICONDUCTORS & AI INFRASTRUCTURE

NVIDIA Corporation **NASDAQ: NVDA**

NVIDIA — Priced for the Clock Not Breaking

The Best Business We Cover, at a Price That Already Assumes the AI Buildout Continues. What NVDA Is Worth on Disciplined Cash Flows — and Why the Swing Factor Sits Outside the Company

TON618 CAPITAL RESEARCH · AS OF JULY 19, 2026 ·

Price \$202.81 (Jul 17, 2026 close) · **52-wk** \$164.07–\$236.54 · **Mkt cap** ~\$4.91tn · **Net cash** ~\$54bn (+\$22bn strategic stakes) · **Dil. shares** ~24.35bn · **Fair value** ~\$163 (bull \$210 / base \$163 / bear \$104) · **Next results** Q2 FY2027, late Aug

A valuation, not a trade. Fair value is an estimate, not a price target, and here it is unusually sensitive to two inputs we name in §7. Data: NVIDIA's FY2026 Form 10-K (fiscal year ended January 25, 2026; SEC EDGAR, CIK 0001045810), Q1-FY2027 10-Q, and SEC XBRL company facts, with market data as of the July 17, 2026 close (\$202.81). Fact ("the filing shows...") is separated from opinion ("we estimate...") throughout; figures are DISCLOSED (in a filing), REPORTED (release/press), CONSENSUS (sell-side), or ESTIMATED (ours). This note stands on its own, but its demand section draws on our companion thematic note The AI Capex Payback Clock (July 19, 2026) as one input among several; §5 says exactly where. It is the demand-side valuation in our five-note AI-infrastructure series — the reader's guide maps all five and is the place to start. Independently fact-checked before release; the Source & Verification Note documents that pass and what we could not verify.

§0 Verdict

FV ≈ \$163

VS \$202.81 — ROUGHLY 20%
BELOW OUR CENTRAL ESTIMATE,
ON A DISTRIBUTION THAT ONLY
REACHES TODAY'S PRICE AT A
BELOW-NORMALIZED COST OF
CAPITAL

NVIDIA is the highest-quality business we cover, and at \$202.81 it is priced for the AI capital cycle to continue essentially uninterrupted. The company is extraordinary on every operating measure — **\$215.9bn of FY2026 revenue** (up 65%), **~71–75% gross margins**, **\$96.7bn of free cash flow**, net cash, and roughly **80% share** of AI accelerators. None of that is the question. The question is what the equity is worth, and on a disciplined three-stage free-cash-flow model our central estimate is about **\$163** — with an honest range from **\$104 to \$210**, because the value swings on the cost of capital and on one thing NVIDIA does not control: how long hyperscaler capital expenditure keeps compounding. Reverse the model and today's price implies NVIDIA reaches roughly **\$1 trillion of annual revenue by FY2031**, a 36% five-year compound rate — a number that requires the very capex durability our companion note questions. We are not short quality; we are disciplined on price.

SCENARIO	WEIGHT	FY2031E REVENUE	FCF MARGIN	IMPLIED FV	VS \$202.81
Bull — capex compounds, share holds ~78%, margins firm	30%	\$980bn	47%	\$210	+4%
Base — growth decays, share drifts to ~75%, some ASIC loss	45%	\$780bn	45%	\$163	–20%
Bear — capex plateaus, share to ~72%, price competition	25%	\$505bn	42%	\$104	–49%
Central estimate (weighted)				≈ \$163	–20%

What this note is, and is not. It is a valuation of the equity, built to answer a question our companion thematic note poses from the demand side: if roughly \$700bn of AI capex a year must eventually earn its cost of capital, what is the one company that collects that capex as revenue actually worth? The finding is not that NVIDIA is a bubble — on valuation it is dramatically cheaper than the Cisco analogy it is constantly compared to (\$6), and its cash generation is real — but that at \$202.81 the risk and reward are no longer balanced for a price-disciplined buyer, because you are paying a full price for a demand stream whose durability is decided outside the company. The street disagrees: consensus sits in the **~\$300–320 mean-target range** (nearer \$300–310 on most aggregators, ~\$319 on others), a **~\$284–300 median**, and a high of \$500 — the whole distribution above spot, and its mean above even our bull case. §7 explains precisely where our numbers and the street's part company — and it is almost entirely the discount rate and the terminal multiple, not the growth.

§1 The Story in Plain English

Read this if you read nothing else.

NVIDIA makes the chips that train and run artificial intelligence, and for the moment it makes almost all of them — roughly **80% of the market** by revenue. Three years ago it was mostly a gaming - graphics company; today **90% of its revenue is data-center AI hardware**, and that business went from \$15bn to **\$194bn in three years**. It is one of the great commercial expansions in corporate history, and it is not an illusion: the money is real, the margins are enormous (about **71 cents of gross profit on every dollar** as reported, ~73 on the clean margin, before the costs of running the company), and unlike most fast growers NVIDIA throws off cash rather than burning it — nearly **\$97bn of free cash flow** last year, with more cash than debt.

So why would the stock be worth less than its price? Because a great business and a good investment are not the same thing, and the difference is what you pay. At \$202.81, NVIDIA is worth about **\$4.9 trillion** — the most valuable company in the world. To justify that, the future has to deliver an enormous amount of growth, and *most of it has to actually arrive*. When we build NVIDIA's future cash flows carefully — letting revenue keep growing fast but not forever, and discounting those future dollars back at a normal rate for a stock this volatile — the value comes out around **\$163 a share**, modestly below the price. Turn the calculation around and the current price is telling you the market expects NVIDIA to reach roughly **\$1 trillion of annual revenue within five years** — nearly five times today. That is not impossible. It is a bet.

And here is the crucial thing about that bet: **it is not really a bet on NVIDIA**. NVIDIA's revenue is, almost by definition, other companies' spending. Every dollar Microsoft, Amazon, Google, Meta, Oracle and the AI labs spend on AI computers, NVIDIA collects a large share of. So NVIDIA's revenue can only keep growing at this pace if *their* spending keeps growing at this pace — and that spending, roughly **\$700bn this year**, has to eventually pay for itself. Our companion note, The AI Capex Payback Clock, works through whether it does. The short version: the spending pays back only if it keeps growing *and* the AI revenue it produces keeps compounding — and NVIDIA is priced as though both happen. If the buildout so much as plateaus, NVIDIA's revenue plateaus with it. The company is superb. The price already knows that. What the price also assumes is that the customers keep spending — which is the one thing NVIDIA cannot decide.

§2 The Business

A platform, not a chip company — and that is the moat.


NVIDIA reports two GAAP segments — Compute & Networking and Graphics — but the business is now overwhelmingly one thing. **Data Center revenue was \$193.7bn of \$215.9bn total in FY2026 (89.7%)** [DISCLOSED], up from 39% of a much smaller total just four years earlier. Gaming (\$16.0bn), Professional Visualization (\$3.2bn) and Automotive (\$2.3bn) are rounding error against the AI franchise. This is a company whose fortunes are now a single, concentrated bet on AI infrastructure demand — which is what makes the demand analysis in §5 the whole game.

What NVIDIA sells is not really a chip; it is an integrated stack that is hard to leave:

- **The silicon**, on NVIDIA's ~24-month architecture cadence — Ampere (2020), Hopper (2022), Blackwell (2024), and Rubin (2026) — each a step-change in performance per watt, which in a power-constrained data center is the binding constraint.
- **The networking**, acquired with Mellanox in 2020: NVLink (the proprietary GPU-to-GPU interconnect), InfiniBand, and Spectrum-X Ethernet. This lets NVIDIA sell *rack-scale systems* (the GB200 NVL72) rather than loose chips, moving the unit of competition from the die to the whole machine — where its lead is widest.

- **The software**, CUDA, a roughly twenty-year install base of libraries, tools and trained developers. This is the true switching cost: a competitor can match a chip far sooner than it can match two decades of software ecosystem and the muscle memory of every AI engineer who learned on it.

The strategic point is that each layer defends the others. A buyer tempted by a cheaper accelerator must also give up the networking that makes thousands of them act as one computer, and the software its engineers already know. That is why NVIDIA has held ~80% share through the largest demand surge in the industry's history even as well-funded rivals attack the chip layer directly.




§3 Porter's Five Forces

The moat is real, and it is narrowing in exactly one place. The Five Forces locate where.

Rivalry — moderate, concentrated at the chip layer. The direct merchant challenger is AMD, whose entire Data Center *segment* (which includes EPYC server CPUs, not just Instinct GPUs) ran about **\$5.4–5.8bn a quarter** in late-2025/early-2026 [REPORTED] — the GPU subset materially below that, against NVIDIA's \$193.7bn annual Data Center line. AMD's MI350 is ramping and the MI400 (on TSMC 2nm) arrives in H2-2026, and a headline **OpenAI-AMD deal for 6GW of GPUs** signals real second-sourcing intent. But AMD competes at the chip; it has no answer yet to CUDA or to rack-scale networking. Rivalry pressures NVIDIA's *share* at the margin, not its position.

Buyer power — the sharpest force, and rising. This is where the moat is genuinely under pressure, and it shows up in NVIDIA's own 10-K. **In FY2026, one direct customer was 22% of revenue and a second was 14% — 36% combined — up from three customers at 12/11/11% in FY2025 [DISCLOSED].** Concentration is not just high; it is *increasing*, and the top customer alone jumped to 22%. Worse for pricing power, these same buyers are NVIDIA's most determined competitors: every large customer is building its own silicon (see substitutes). And the true concentration is higher than the direct table shows — NVIDIA discloses that *"one AI research and deployment company contributed to a meaningful amount of our revenue purchasing cloud services from our customers"* [DISCLOSED], i.e. OpenAI-type end-demand routed through the clouds. A handful of buyers, each large enough to matter and each building an alternative, is the structural weakness beneath the 80% share.



Supplier power — real, and a two-way street. NVIDIA is fabless; its growth is gated not by demand but by two suppliers. **TSMC's CoWoS advanced packaging is sold out through 2026** even as capacity climbs toward ~120–140k wafers/month, and **HBM3e memory is sold out** with prices up ~20% for 2026, dominated by SK Hynix (~50–62% share) [REPORTED]. NVIDIA depends on a near-monopoly (TSMC) and an oligopoly (HBM) it does not control — but that same bottleneck is a shared moat: it is equally a wall against new entrants, who cannot get allocation NVIDIA has locked up.

New entrants — barriered. The combination of CUDA's ecosystem, the capital to design leading-edge silicon, and — decisively — access to the very CoWoS and HBM allocation that is already sold out, makes de-novo entry near-impossible. The credible entrants are not startups; they are the trillion-dollar buyers building captive silicon.

Substitutes — the long-term share threat. Custom ASICs are the substitute that matters: **Google's TPU (v7 "Ironwood" now GA)**, **Amazon's Trainium (Trainium2 in volume for Anthropic, Trainium3/4 announced)**, **Microsoft's Maia 200 (delayed ~6 months)**, and **Meta's MTIA 300** [REPORTED]. These are mostly *captive* — used internally rather than sold — so they do not compete for NVIDIA's merchant revenue directly, but they cap NVIDIA's share at the largest buyers and grow the share of AI compute that never touches an NVIDIA part.

Moat verdict. Durable, and narrowing in one place. CUDA plus rack-scale systems keeps the merchant moat wide against AMD and impassable to new entrants. The erosion is at the top of the customer base, where the buyers with the most volume are also the ones with the resources to in-source — so the realistic path is not a collapse in share but a slow drift from ~80% toward ~75% and below as ASICs scale, even as the absolute revenue keeps rising with the market. That drift is exactly what our base case assumes.

§4 The Financials

Almost nothing to criticize in the operating numbers. The two things worth watching are on the balance sheet.

The income statement is close to unprecedented for a company this size. Revenue compounded from **\$16.7bn (FY2021) to \$215.9bn (FY2026)**; gross margin runs **71–75%** (the reported FY2026 71.1% was pulled down ~2.6 points by \$7.2bn of inventory and purchase-obligation provisions, including a \$4.5bn H2O charge — the clean rate is nearer 73–75%); operating margin is **60%**; net income was **\$120.1bn** [DISCLOSED]. Cash conversion is the rarest part: NVIDIA is fabless, so capex is only ~3% of revenue, and it produced **\$96.7bn of**

free cash flow (44.8% margin) in FY2026, carries **net cash**, and returned **\$40bn** in buybacks. On operations, there is essentially nothing to fault.

Two balance-sheet items deserve the reader's attention, and both connect to the demand question:

- **\$95.2bn of non-cancellable inventory, supply and capacity commitments** at FY2026, against only \$21.4bn of on-balance-sheet inventory [DISCLOSED, flagged as a Critical Audit Matter by PwC]. This is NVIDIA's own supply-side bet — it has pre-committed to buy components far in excess of what it holds, on the assumption demand continues. If demand plateaus, some of that converts to writedowns, exactly as the \$4.5bn H2O charge already did when China demand fell.
- **Non-marketable equity securities — stakes in private companies — jumped from \$3.4bn to \$22.3bn in FY2026** [DISCLOSED], and NVIDIA separately holds public-company stakes in its marketable-securities line (the 10-K names Intel; its ~\$2bn CoreWeave position is public common stock). NVIDIA does not itemize the private holdings, but it has publicly announced large strategic commitments to AI customers and partners (OpenAI, xAI and others). Whatever the exact line-by-line split, the aggregate direction is the "circular financing" data point: NVIDIA increasingly helps fund the buyers who then buy its chips. It is small relative to the company, but it is the modern echo of a pattern §6 examines.

One flag for anyone annualizing recent results: **Q1-FY2027 net income (\$58.3bn) exceeded operating income (\$53.5bn)** [DISCLOSED], implying a large non-operating gain (likely investment/equity-stake marks). We build our valuation on operating free cash flow, not on that quarter's net income, and readers should do the same.


§5 Demand: Where NVIDIA's Revenue Actually Comes From

This is the section the whole valuation turns on, and it is why the companion note exists.

NVIDIA's revenue is not, in any meaningful sense, set by NVIDIA. It is an identity:

> **Data-Center revenue $\approx$ (total AI-infrastructure capex) $\times$ (GPU share of that capex) $\times$ (NVIDIA's share of GPUs).**

Put the FY2026 numbers in and the identity holds: roughly **\$600bn of calendar-2025 AI-infrastructure capex** — the whole buyer base, broader than the four hyperscalers of the Clock — **$\times$ ~40% GPU share $\times$ ~80% NVIDIA share $\approx$ \$192bn** [ESTIMATED], against the \$193.7bn actually reported. The tie is close enough to trust the framework — and once you trust it, NVIDIA's forward revenue is simply the AI-capex forecast run through two shares that are, if anything, slowly declining.



That is why this valuation cannot be done honestly without a view on capex, and why our companion note builds one. **The AI Capex Payback Clock** works through whether the ~\$700bn of 2026 hyperscaler capex can earn its cost of capital, and finds that it does only under a specific condition: capex growth must plateau *and* AI revenue must keep compounding. We use that note here as one input — not as a forecast we import wholesale, but as the source of the three capex paths behind our three revenue scenarios:

- **Bull** takes total AI capex compounding past \$1.3T, GPU share holding near 38–40%, and NVIDIA holding ~78% — delivering the \$980bn FY2031 revenue in our top case.
- **Base** takes capex growth decelerating and NVIDIA's share drifting to ~75% as ASICs scale — \$780bn.
- **Bear** takes the Clock's plateau: capex growth stops, and NVIDIA's revenue flattens near \$505bn.

The reverse-DCF (\$7) shows the current price requires roughly the **\$1 trillion FY2031 revenue** of the bull path — which in turn requires **total AI-infrastructure capex near \$3 trillion a year** at that horizon. That is the demand the price embeds. The Clock's finding is that even \$700bn–\$1T is not yet self-funding on a three-year hardware clock; \$3T compounds the question rather than answering it. A reader can hold a more optimistic capex view than ours and still see the point: at \$202.81, NVIDIA is not priced for AI to succeed — it is priced for AI *infrastructure spending* to keep accelerating, which is a narrower and more demanding claim.

This locates NVIDIA in the wider series. Our companion *The Power Wall* sorts the AI-infrastructure complex by how fast new supply can arrive, and silicon lands on the fast side — a shortage rent that prices away in roughly two years, which is why NVIDIA defends its premium through CUDA and rack-scale systems rather than through chip scarcity. But that whole edifice rests on the one force this section has isolated: demand. NVIDIA is the finest business on the rented side of the trade — the chips — and its price is decided by whether the buyers keep buying.


§6 The Cisco Test

The comparison everyone reaches for. Run on the numbers, it exonerates the valuation and indicts the demand structure.

NVIDIA is constantly likened to Cisco at the March-2000 peak — the last time a single company was *the* pick-and-shovel supplier to a technology buildout, at the top. We ran the comparison on the actual figures rather than the vibe, and it cuts both ways, cleanly.

	CISCO (MAR 2000)	NVIDIA (JUL 2026)
Market cap	~\$555bn	\$4.91tn
Trailing P/E	~201x	~41x (FY26 EPS)
Forward P/E	—	~22–23x
Price / Sales	~29x	~23x trailing
Gross margin	~65%	~71–75%
Free cash flow	modest; growth vendor-financed	~\$97bn, net cash

On valuation, the analogy fails, and it is not close. Cisco peaked at ~201x earnings; NVIDIA trades at ~41x trailing and ~22x forward. Cisco's growth was substantially *vendor-financed* — it lent money to the telecom and CLEC customers who bought its gear and then went bankrupt, taking the revenue with them. NVIDIA funds its growth from ~\$97bn of real free cash flow and sits on net cash. A buyer at NVIDIA today is paying a fraction of Cisco's multiple for a more profitable, cash-generative business. The lazy "it's Cisco again" call is wrong on the arithmetic.




On demand structure, the analogy is uncomfortably apt — and that is the real warning. Two Cisco echoes are present in NVIDIA's own filings. First, the faint return of vendor financing: the **\$22.3bn of strategic stakes** in its own customers (\$4) is the same circularity that flattered Cisco's demand, in miniature. Second, and more serious, **customer concentration is far higher than Cisco's ever was** — Cisco sold to thousands of enterprises and carriers; NVIDIA's top two direct customers are **36% of revenue**, and end-demand concentrates in a handful of hyperscalers and AI labs. Cisco did not break because it was overvalued alone; it broke because the *demand* behind the revenue turned out to be an overbuild funded by cheap capital. That is precisely the risk the Payback Clock isolates. The honest read of the Cisco test: NVIDIA will not fall 80% on multiple compression from 41x. It could re-rate hard if the *demand* proves to be a capex overbuild — the same mechanism, a different valuation starting point.

§7 Valuation

Reverse-DCF first, to anchor what the price assumes; then a disciplined forward model; then the sensitivities that make the point.

What the price assumes. Hold NVIDIA's structure constant — 45% free-cash-flow margins, an 8-year fade to a 4% terminal, discounting at an 11.3% cost of equity (a normalized 4.3% risk-free rate and 4.8% equity risk premium on a normalized beta ~1.45, not the momentum-inflated trailing 2.2) — and solve for the revenue path that justifies \$202.81. The answer: **roughly \$1 trillion of revenue by FY2031, a 36% five-year compound rate [ESTIMATED]**. The price is not egregious; it is a specific, demanding forecast, and §5 shows it requires ~\$3tn of annual AI capex to feed it.

The forward model. Our three-stage FCFE model — explicit revenue FY2027–31, an 8-year growth fade, a 4% Gordon terminal, plus the \$76bn of net cash and strategic stakes — produces the scenario values in §0: **bull \$210, base \$163, bear \$104, weighted ~\$163**. The base case is not conservative on growth: it has revenue tripling to \$780bn by FY2031. It comes in below the price because a disciplined discount rate and a disciplined terminal multiple (about 14× terminal FCF) do not, by themselves, support \$4.9tn until the cash actually shows up.



The sensitivities are the honest core of the valuation, because the answer moves more on two assumptions than on the growth we argue about:

BASE - CASE FV	TERMINAL G 3.5%	4.0%	4.5%
Cost of equity 10.0%	\$188	\$199	\$211
11.3% (base)	\$157	\$163	\$171
12.5%	\$134	\$139	\$144

Read the top row: **at a 10% cost of equity, the base case supports ~\$199–211 — essentially today's price**. The entire gap between our \$163 and the \$202.81 market — and most of the larger gap to the ~\$300–320 street — is a disagreement about the *discount rate and terminal multiple*, not about how fast NVIDIA grows. We already grant the growth. The question is what a dollar of NVIDIA cash flow arriving in 2033 is worth today, for a stock

with a 2.2 beta and a demand base decided by other companies' capital budgets. We think 11.3% is the honest number and 10% is generous; a reader who accepts 10% and base-case execution can justify the current price, and should own it with clear eyes about what they are assuming.

Where we land. Central FV $\approx$ **\$163**, roughly 20% below spot, on a distribution from **\$104 to \$210**. The skew is mildly negative and it is argued, not assumed: the point estimate sits below the price because the price requires a below-normalized cost of capital *and* uninterrupted capex growth, and we decline to assume both at once.

§8 The Risk Ledger, Both Ways

A valuation this sensitive owes the reader both tails, explicitly.

What makes us wrong on the upside:

- **Inference eats training.** If AI compute demand shifts decisively to inference (serving models, not training them) and that market is larger and stickier than the training buildout, NVIDIA's TAM re-rates upward and the \$1tn revenue stops looking like a stretch.
- **The software and networking attach compounds.** If NVIDIA converts its systems lead into high-margin recurring software and networking revenue, margins firm above our 45% and the terminal multiple deserves to be higher.
- **Sovereign AI.** Government and national-champion buildouts (outside the four US hyperscalers) add a demand leg the capex debate largely ignores, and they are less price-sensitive.
- **Share holds.** If CUDA and rack-scale systems keep share near 80% rather than drifting to 75%, the base case migrates toward the bull.

What makes us right, or worse:

- **The Clock breaks.** The dominant risk, and it is exogenous: if hyperscaler capex plateaus — the condition our companion note shows is *required* for the spending to pay back — NVIDIA's revenue plateaus with it, and the bear case (\$104) governs.
- **ASIC substitution accelerates at the top buyers.** The 36%-of-revenue top-two customers are the most able and most motivated to in-source; faster-than-expected TPU/Trainium/Maia scaling caps share below 72%.
- **Export controls.** China fell from ~20% to 9.1% of revenue under controls, with a \$4.5bn charge already taken; a further tightening (or the 15% licensed-sales levy the USG has signaled) removes a revenue leg outright.
- **The purchase-commitment and stake writedowns.** \$95.2bn of non-cancellable commitments and \$22.3bn of customer stakes are geared to demand continuing; a plateau converts part of both into charges.

The asymmetry we see: the upside cases make NVIDIA worth somewhat more than today; the downside case (a capex plateau) makes it worth roughly half. That is why, for a valuation-disciplined buyer, the risk/reward at \$202.81 is unfavorable despite the quality — not because NVIDIA is a bad business, but because the price already pays for the good outcome and the bad outcome is decided elsewhere.

§9 Recommendation, and What We Could Not Verify

Recommendation. Fair value $\approx$ **\$163**, distribution **\$104-210**, versus \$202.81. We are constructive on the business and cautious on the equity at this price: this is a name to own on a demand-cycle pullback, not at a full price for a demand stream the company does not control. We are explicitly *not* short — net cash, real cash flow, dominant position and genuine upside optionality make a short thesis unattractive; the honest call is that the

reward no longer compensates the risk at \$203. We log this fair-value call against the \$202.81 close to our signal ledger for out-of-sample scoring, per house practice on every single-name valuation.

What we could not verify, and how it would move the number.

- **Forward consensus dispersion.** Published sell-side consensus centers near \$391bn revenue and ~\$9.3 EPS for FY2027, but the spread across aggregators is wide (some FY2027 revenue reads as low as ~\$342bn); we anchor the near-year ramp on the central figures. This affects the near-year ramp, not the terminal.
- **The Q1-FY2027 below-the-line gain** (net income > operating income) is not fully decomposed; we sidestep it by valuing on operating free cash flow.
- **The itemized strategic-stake list** (\$22.3bn) is not fully broken out in the sources we pulled; the aggregate is disclosed and used, the composition is not.
- **Networking revenue** is not separately disclosed; we treat it inside Data Center.

None of these move the central estimate materially, because the valuation is dominated by the discount rate, the terminal multiple, and the capex-demand path — the three things \$7 and \$5 make explicit. The swing factor is not a number hidden in NVIDIA's filings. It is the sustainability of everyone else's spending.

§10 Disclosures

Information only. TON618 Capital. This report is for information purposes only. Nothing here is an offer to sell or a solicitation of an offer to buy any security, fund interest, or digital asset, and nothing here is personalized investment advice or a recommendation regarding any instrument.

Publisher's exclusion. All research is published solely as general, impersonal information of regular circulation. It is not tailored to the objectives or circumstances of any individual and is not issued in connection with compensation from any client. The Fund has no clients and distributes all research free of charge. On that basis it publishes in reliance on the publisher's exclusion from the definition of "investment adviser" under the Investment Advisers Act of 1940 (§202(a)(11)(D); cf. *Lowe v. SEC*, 472 U.S. 181 (1985)).

Registration & conflicts. TON618 Capital is not registered as an investment adviser or broker-dealer in any capacity. The Fund is a Bitcoin fund and may hold or transact in the securities or digital assets it discusses. This report values NVIDIA (NASDAQ: NVDA) and references, among others, Advanced Micro Devices (AMD), Microsoft (MSFT), Alphabet (GOOGL), Amazon (AMZN), Meta Platforms (META), Micron (MU) and Cisco (CSCO); the Fund holds no position, long or short, in any of them, and has no economic interest in the price of any security named here. The Fund receives no compensation from any party in connection with its research.

Use of AI. Artificial intelligence is used in the creation of this research. All methodology and data integrity are reviewed and approved before publication by TON618 Capital's Chief Investment Officer, **Keyth Beck**; errors may nonetheless occur, and readers should verify independently.

CFA. This report was prepared to align with CFA Institute analytical standards (methodology only). CFA® and Chartered Financial Analyst® are registered trademarks owned by CFA Institute. That reference describes the analytical framework applied; it does **not** imply the report was prepared, reviewed, or authored by a CFA charterholder, and the report is **not** issued, reviewed, endorsed, certified, or approved by — nor affiliated with — CFA Institute.

Risk & feedback. Past performance is not indicative of future results. Digital assets and equities are volatile and may result in total loss of capital. Corrections and feedback are welcome — please direct them to CIO **Keyth Beck** at keyth@ton618capital.com.