


EQUITY VALUATION · TECHNOLOGY · ANT · 26 OCTOBER 2026

Anthropic (ANT) — What the S-1 Says a Dollar of Claude Is Worth

Anthropic filed to go public on October 26, 2026. The filing shows a company whose revenue rose from \$4.6 billion in 2025 to \$16.3 billion in the first half of 2026, whose second quarter was the first with an operating profit, and which costs about 15 cents to serve for each dollar it sells. It also shows about \$657 billion of named compute commitments landing fastest in 2027–29, a founders' share that carries 50.1% of every general vote, and a trust with no stock that elects a majority of the board. Fair value \$50.72 a share, probability-weighted, \$833 billion: 14% below the \$58.90 price of the last private round and less than half of the \$2 trillion the wires discuss.

TON618 CAPITAL RESEARCH · AS OF 26 OCTOBER 2026



The ladder and the curve.

Anthropic has named about \$657 billion of compute commitments, some ceilings, landing fastest in 2027–29. It turns on how quickly revenue climbs the rungs it has already paid for, and on who decides when a model is allowed to earn it.

THE REVENUE · THE COST TO SERVE · THE LADDER · THE CONTROL · THE MOAT · THE VALUE

On the morning of Monday, October 26, Anthropic, PBC filed its registration statement with the SEC for a listing on the Nasdaq Global Select Market under the ticker ANT, with Goldman Sachs and Morgan Stanley leading. There is no price range, no share count and no figure for the quarter that ended in September; those arrive by amendment. What the filing does give is three audited years and a half-year: revenue of \$41 million, \$386 million and \$4,585 million in 2023–25, then \$16,286 million in the six months to June 2026; a \$693 million operating loss for that half that turned into a \$974 million operating profit in its second quarter; \$131.7 billion of cash and short-term investments; and a net loss of \$80.7 billion, of which \$72.4 billion is the accounting mark on convertible notes that rose in value because the company did. The last private round, Series H in May and June, was priced at \$589.0095 per share before a ten-for-one split, or \$58.90 after it, about \$968 billion for the whole company on the count this note uses. This note asks what the business the filing describes is worth, what the company has promised to spend to build it, and who decides.

THE QUESTIONS

What this note answers

1. **How did Anthropic get here, and what does it say it is?** The founding, the public-benefit purpose, the trust, the model line to Mythos and Fable, and the turn to enterprise.

2. **What is the revenue, really?** Growth from \$4.6 billion to a \$46 billion run-rate, what the blank third quarter implies, who pays and through whom, and circularity as a number.
3. **What does a dollar of Claude cost to serve?** The serving margin the filing does not print, built from its own functional table.
4. **What has it promised to spend?** Every named compute commitment, by counterparty and by year, against revenue and cash.
5. **Who owns it and who controls it?** The share count reconciled, economics against votes, the founders' share and the trust.
6. **What is the moat and what erodes it?** Cadence, gating, four chip families, three clouds, open weights and price.
7. **What do the law and the regulator want?** The government directive, the export restriction, the FTC, the copyright docket.
8. **How long does it grow faster than normal?** Each line's stage, argued from its driver, with the length the price needs.
9. **What is it worth?** A multi-stage cash-flow valuation, the three cases, the sensitivities, what each anchor price requires, and the peers.
10. **What would change our view?** Dated, observable conditions in both directions, and the two update points.

\$50.72

FAIR VALUE PER SHARE, PROBABILITY-WEIGHTED (BEAR \$10.07 / BASE \$44.81 / BULL \$103.21); 13.9% BELOW THE SERIES H PRICE OF \$58.90; 44% BELOW \$1.5 TRILLION AND 58% BELOW \$2 TRILLION

The read. Anthropic is cheap to serve and expensive to build. Cost of revenue was \$2.4 billion on \$16.3 billion of first-half revenue, an 85% serving margin, up from 40% in 2025; the larger part of the compute bill, \$6.6 billion in six months, went to research, model development and internal use. Second-quarter revenue of \$11.6 billion was up 144% on the first quarter, and the filing's own text implies about **\$16–17 billion** for the blank third quarter, growth of roughly 39–47%. Against that the company has named about **\$657 billion** of compute commitments, some of them ceilings or cancellable on notice (Microsoft, Google, Amazon, SpaceXAI, Broadcom leases, Nscale, Lambda, AMD), plus an estimated \$125–150 billion of TPU capacity for 2028–31; scheduled by the house they land at **\$93 billion in 2027 and \$156–160 billion a year in 2028–29**, and the filing says margins "may compress" as they arrive. The base case grows revenue to \$113 billion in 2027 and \$303 billion in 2031, runs operating losses in 2027–29 while the ladder lands, and earns 34% at maturity. The Series H price is that path at an **11.3%** cost of equity against the house's 12.25%, or 14% more revenue in every year; \$2 trillion needs 2.4 times the house path. The public will buy Class A shares that carry at most 49.9% of the vote against one founders' share that carries 50.1%, under a board whose majority is elected by a trust that owns no stock.

85%

Serving margin in the first half of 2026: cost of revenue \$2,380M on revenue \$16,286M, built from the filing's functional expense table (no gross-margin line is printed); 40% in 2025

\$657B

Named compute commitments in the filing: Microsoft, Google, Amazon, SpaceXAI, Broadcom leases, Nscale, Lambda, AMD; the leak's "\$518 billion" appears in no draft or exhibit

50.1%

Share of every general stockholder vote carried by one Class F share held by the seven co-founders' LLC, fixed whatever is issued later; four of seven directors elected by the Long-Term Benefit Trust

1.14×

Multiple of the house revenue path, in every year of every case, that the \$58.90 Series H price requires at a 12.25% cost of equity; 1.75× for \$1.5 trillion, 2.38× for \$2 trillion

Terms used throughout, defined once. **S-1:** the registration statement a company files to sell shares to the public; **DRS:** the confidential drafts filed before it. **Consumption revenue:** usage billed by the token (the filing's "pay-for-intelligence"); **subscription revenue:** prepaid seats (Pro, Max, Team, Enterprise). **Compute:** the processors, memory, networking and power used to train and run models; the filing books all of it as one operating expense. **Cost of revenue:** the filing's functional line for the compute and overhead that serve customers; the **serving margin** is one minus cost of revenue over revenue. **CSP:** a cloud services provider (Amazon, Google, Microsoft), which here is investor, compute vendor, sales channel and customer at once. **Run-Rate Revenue:** the company's annualized measure (13 times 28 days of consumption plus 12 times monthly subscriptions). **RPO:** remaining performance obligations, contracted revenue not yet recognized. **NRR:** net revenue retention, this year's revenue from last year's customers over last year's. **PBC:** a Delaware public benefit corporation, whose directors must balance stockholders against a stated public purpose. **LTBT:** the Long-Term Benefit Trust, five unpaid-in-stock trustees who hold one Class T share and elect a board majority. **Class F:** one share held by the founders' LLC with 50.1% of the general vote. **Class N:** non-voting shares held by strategic investors, the Amazon notes and the Public Benefit Fund. **FCFF:** free cash flow to the

firm, the cash the business generates before financing. **Cost of equity:** the return shareholders require, built here as the ten-year Treasury yield plus beta times an equity risk premium. **Terminal growth:** the growth rate assumed forever after the explicit forecast. **Reverse DCF:** solving the model backward for what a given price requires.

✓ THE OTHER SIDE, FIRST

The case for the price

The strongest case that the Series H price of \$58.90, and the \$1.5–2 trillion the wires discuss, is right, in the terms its proponents would use.

1. **No company has ever grown like this, and it is still accelerating in dollars.** Revenue went from \$41 million to \$4.6 billion in two years and to \$16.3 billion in the next six months; the second quarter alone was \$11.6 billion, up 144% on the first. Customers above \$100,000 of annual revenue went from 950 to 11,900 in a year; 250 customers pay more than \$10 million; net revenue retention is about 900%. Every company in the Fortune 10 and over half the Fortune 500 are customers, and they spend on average less than 0.1% of their expenses on Claude, so the runway inside existing accounts is enormous. A model that puts 2027 growth at 114% is extrapolating deceleration the data has not yet shown.
2. **The unit economics are software's, not a utility's.** Serving a dollar of revenue cost 15 cents in the first half of 2026, down from 60 cents in 2025; revenue per dollar of total compute went from 0.4 to 2.1 in five quarters; Opus 5.5 cut the cost per task by about 40% against Opus 5 and the company says it delivers such reductions "while preserving our margins." The adjusted operating margin went from –221% to +19% in six quarters. The founders' letter calls the \$6.6 million of revenue per employee "meaningfully higher than the average for the largest public U.S. companies."
3. **The commitments are a demand forecast made by the people with the best view of demand.** The company says it has "consistently operated in a compute-constrained environment where customer demand exceeds available compute." Multi-year contracts are procured 12 to 18 months ahead, so the ladder is what the company expects to sell, bought at the best available price through four chip families. Reuters reported in August that the IPO case rests on a \$190–200 billion revenue forecast for 2028; the house base case is 15% below it.
4. **The balance sheet and the partners remove financing risk.** \$131.7 billion of cash and short-term investments at June 30, \$77 billion of undrawn facilities from Google, Amazon and Broadcom, a \$15 billion revolver, AMD's commitment to buy up to \$5 billion of stock, Google's option to buy \$10 billion at the IPO, and whatever the IPO itself raises. No frontier lab has ever had this much capital in hand.
5. **Safety is the enterprise moat, not a tax on it.** The company's position is that "safety, alignment, and ethical behavior resonate particularly strongly with enterprises." Ramp's card data had 43.8% of US businesses paying Anthropic against 40.6% paying OpenAI in September 2026, and 60% of token spend in October; the filing cites every Fortune 10 company. The gating of Mythos cost "substantial revenue" and about \$100 million of credits, by the founders' own account, and the company grew 144% in the quarter it happened.
6. **The cost of equity is too high.** At a 10% cost of equity, which is what the house uses for SpaceX, the same three cases give \$74 a share, above the Series H price; the proponent's own inputs, bull growth with a 10% rate and a seven-year stage, give \$168.
7. **Control is a feature.** The structure the public is being asked to accept is the structure that produced the growth: the founders' LLC and the trust were designed so that the company could hold back a model when it judged it must, and the company's record is that such decisions have not stopped it compounding.

Of these, this note tests:

- (1) in §2 and §8, where the blank third quarter is bounded from the filing's own text and the stage growth rates are set against the company's reported 2028 forecast and the base rates of growth persistence;
- (2) in §3, where the serving margin is built from the functional table and the research compute is separated from it;
- (3) in §4 and §9, where the ladder is scheduled by year and run as a cash floor in every case, so the reader can see what it costs when demand lags it;

- (4) in §4, where the facilities' strings (conversion at a discount, revenue conditions, ownership caps) are set out;
- (5) in §6, where the share data, its basis and its gaps are stated, and in §10, where the gating claim has a falsifier;
- (6) in §9, where every result is shown at 10%, 11%, 12.25% and 13%, and the proponent's case is run on its own inputs.

It cannot test:

- (7) Whether the control structure will produce good decisions; it is a governance fact the note describes and prices only through the discount rate.
- Whether the company's 2028 forecast is its own; Reuters reported it, the filing does not state it.
- What the IPO will be priced at or how many shares will be sold; both are blank, and both are update points.

✚ HOW IT GOT HERE

How Anthropic got here, and what it says it is

Anthropic was incorporated in Delaware in January 2021 as a public benefit corporation by seven researchers who had led GPT-3's engineering, the scaling-laws work, RLHF and interpretability research at OpenAI. Its charter's stated purpose is "the responsible development and maintenance of advanced AI for the long-term benefit of humanity," and the filing calls the company "first and foremost a research lab." It launched commercially in the first quarter of 2023 with 43% of that year's \$41 million of revenue from one customer, shipped the Claude 3 family in 2024, Claude Code as a research preview in February 2025, and Claude Cowork in January 2026. Today the model line has four classes: Mythos (with Fable), Opus, Sonnet and Haiku. Mythos was first released in April 2026 as a gated research preview through Project Glasswing, a program for vetted cybersecurity work; Fable 5, the generally available configuration "with our production safeguards, which block or limit performance on certain sensitive topics," followed in June, with Mythos 5 for Glasswing participants; Mythos 5.1 and Fable 5.1 shipped in September.


The scale as of September 30, 2026: more than 11,900 customers paying over \$100,000 a year, 1,900 over \$1 million and about 250 over \$10 million; roughly 40 million smaller customers; every company in the Fortune 10, over half the Fortune 500 and over 30% of the Global 2000. Revenue per employee is "over \$6.6 million," which with the filing's 170% headcount growth puts the staff at about 5,100–5,200. Over 90% of the contributions to the company's own codebase in September 2026 were written by Claude. Regrettable attrition was about 1% over the trailing year, and all seven founders remain.

What it says it is appears most plainly in the eight-page founders' letter inside the filing. It states two categories of decision that "may run counter to investors' expectations": "forward investments (especially in compute) that may seem irresponsibly large at first glance," defended by the scaling laws and a comparison to pharmaceutical research, and "decisions that depart from profit maximization." Its two examples are the ones this note prices: holding Mythos Preview back from general release "meant forgoing substantial revenue and ceding ground to competitors (we also gave away around \$100 million in Claude credits to cyber defenders)," and "we have also committed to keeping Claude free of advertising," which "would likely be a large source of revenue for us." It adds that "there may be times when we choose to advance more slowly than we could in order to make our systems safer, even if that puts us at a competitive disadvantage." The letter's own advice on the offering: "never let your company be worth a dollar more than it is actually worth."

What the revenue is, really

EXHIBIT 1 — REVENUE AGAINST COMPUTE, BY QUARTER

Quarterly revenue against compute and infrastructure expense (\$ billions), and the revenue each dollar of compute produced



How to read it. Each quarter from the first of 2025 to the second of 2026 has two bars in billions of dollars: gold is revenue, grey is compute and infrastructure expense for every workload. The number under each pair is revenue divided by compute, the dollars of revenue each dollar of compute produced that quarter. The dashed bar at the right is the third quarter of 2026, which the public filing leaves blank: its height is the \$16–17 billion range the filing's own text implies. The column at the right lists the company's adjusted operating margin by quarter. Compare the gold bar's rise with the grey bar's; the one thing it shows is that revenue grew faster than compute in every quarter of 2026, which is why the loss closed.

Revenue was \$376 million, \$787 million, \$1,382 million and \$2,040 million in the four quarters of 2025, then \$4,731 million and \$11,555 million in the first two of 2026. Consumption, billed by the token, is 80% of the mix; subscriptions 20%; the United States 72%. The third quarter is blank in the public filing, but two sentences pin it: trailing-twelve-month revenue "increased approximately 1,200%" from September 2025 to September 2026, and revenue per employee went from \$1.4 million to over \$6.6 million while headcount grew about 170%. With the undisclosed fourth quarter of 2024 bounded at \$97–174 million, the growth sentence gives **\$16.0–17.0 billion** for the third quarter of 2026 and \$34–35 billion of trailing revenue (the revenue-per-employee route gives \$15.3–16.3 billion, and the "approximately" in both sentences can widen either); the Financial Times, via Reuters, reported \$16.8 billion on October 15, inside the range. Note what the range implies: growth of 144% from the first quarter to the second, then 39–47% from the second to the third. The company's own run-rate metric was more than \$47 billion "earlier in May"; Bloomberg reported \$65 billion at the end of July, a figure the company has not published.

Who pays. Customers above \$100,000 a year produced 74% of first-half revenue (\$12.0 billion); the 40 million smaller customers the rest. Four customers have at times exceeded 10% of revenue: Customer A (12% of 2025, 13% of receivables), Customer B (12% of 2025, 18% of the first half of 2025), Customer C (43% of 2023, 11% of the first half of 2026) and Customer D (11% of 2023). All are unnamed, and the filing does not say whether any is an investor or a compute vendor. No customer exceeded 10% in the nine months to September, and 79% of third-quarter revenue came from outside the top five. Remaining performance obligations were \$10.8 billion at June 30, about 75% of it due within a year: a backlog of about one quarter of a year's revenue at the second-quarter pace. The book is short, and the filing says so: most revenue is consumption, and "a growing number" of enterprise customers sign non-cancelable annual or multi-year minimums.

Two deductions the filing makes itself. "Approximately 5% to 7%" of revenue for the nine months to September came from customers later found to be violating the usage policy or commercial terms; enforcement tightened in August and September and the company expects the share to fall. And "for the two weeks ended January 2, 2026, total customer spend declined by approximately 20%," the seasonality the company expects to show as growth slows. Models released in the six months to September produced over 80% of third-quarter revenue: the revenue rides the release cadence.


Circularity, as a number. The filing names no customer, so revenue from the cloud companies that also own, supply and distribute Anthropic cannot be measured. The other three legs can. Money in: Amazon and Google have put about \$41 billion into Anthropic's notes and preferred stock through June 2026 (Amazon about \$18 billion, Google about \$23 billion), with \$35 billion of undrawn facilities and Google's \$10 billion option at the IPO. Money out: at least \$114.5 billion committed to Google for compute, an estimated \$125–150 billion more of TPU capacity "from TPU providers, including Google," \$161.2 billion of Broadcom lease rent for TPU systems, and at least \$110.0 billion committed to Amazon. **For every dollar the two have invested, Anthropic has committed to spend about \$5.50 with them on compute, more than \$8 with the TPU estimate.** Money through: 48% of first-half revenue was sold through the three clouds' marketplaces, where the CSP collects and remits and Anthropic pays a channel fee; sales, marketing and partnerships, mostly those fees, is the fastest-growing cost line at 19% of revenue. Revenue from the CSPs as customers: undisclosed; the filing says they "are also our customers" and that their "purchasing decisions may be influenced by their own competing offerings." The filing's own risk factor calls these relationships "complex dynamics that could give rise to conflicts of interest."

⊕ THE COST TO SERVE

What a dollar of Claude costs to serve

EXHIBIT 2 — WHERE THE COMPUTE GOES

Where the compute goes: serving customers against research, development and internal use, as a share of revenue



How to read it. Four stacked bars, one for 2024, 2025 and each first half, measured as a share of that period's revenue up the side (capped at 100%). The gold segment is cost of revenue, which the filing says consists "substantially all" of the compute that serves customers; the grey segment above it is the rest of the compute bill, the upper bound for research, model development and internal use. The label above each bar is total compute as a share of revenue; the label below is the serving margin, one minus the gold share. Compare the gold segments across the bars; the one thing it shows is that serving a dollar of revenue cost 15 cents in the first half of 2026 against 60 cents in 2025, and that most compute is spent not on serving but on research.

The filing has no gross-margin line: compute is one operating expense that "supports the full range of workloads," and the company says it plans and assesses it "on a total basis." Its functional table gives the serving margin anyway. Cost of revenue, which "substantially all" consists of compute for customer-serving workloads plus some personnel and overhead, was \$752 million on \$386 million of revenue in 2024, \$2,773 million on \$4,585 million in 2025 (a 40% serving margin) and \$2,380 million on \$16,286 million in the first half of 2026 (85%). Compute as a whole was \$7,326 million in 2025 and \$8,945 million in the half; by quarter it fell from 251% of revenue to 48%. So serving was 38% of the compute bill in 2025 and 27% in the half; the rest, \$6.6 billion in six months, went to research, development and internal use, the cost of staying at the frontier. The channel fee is the other cost that scales with revenue: sales, marketing and partnerships, mostly those fees, rose to 19% of revenue in the half from 12% in 2025, \$2.7 billion of the increase being fees to the cloud marketplaces.

What the filing says happens next is the hinge of the valuation: "As significant incremental compute is brought online in future periods, the associated expense will be recognized in advance of the full revenue that the compute will ultimately support, and we expect operating margin and Adjusted Operating Margin may compress in those periods." Compute is

contracted 12 to 18 months ahead, is "generally fixed in the near term," and "compute and infrastructure expenses as a percentage of revenue will depend primarily on how our revenue scales relative to our total committed compute cost base." Short-term compute, at a higher unit cost, is expected to rise in the near term. The 85% serving margin is real; the operating margin is the ladder in \$4 against the revenue in \$2.


The other lines: personnel \$3,152 million in the half, of which \$2,185 million was stock compensation (13% of revenue) and \$967 million cash (6%); other operating and administrative \$1,714 million (11%), including the donation-matching charges that reached \$2.0 billion in the third quarter alone. Unrecognized stock compensation stood at \$26.6 billion on restricted stock units (of which \$2.9 billion relates to units whose service condition is satisfied or partially satisfied) plus \$2.8 billion on options, with \$5.8 billion more granted after June. The model expenses stock compensation, at 10% of revenue in 2027 falling to 6%, and does not also dilute for it.

THE LADDER

What it has promised to spend

EXHIBIT 3 — THE LADDER BY YEAR AND COUNTERPARTY, AGAINST REVENUE IN EACH CASE

The commitment ladder by year and counterparty (\$ billions, house schedule) against revenue in the three cases



How to read it. Each year from 2026 to 2033 has a stacked bar in billions of dollars of cash compute commitments, split by counterparty in the shades of the legend; the number above the bar is the year's total. The three lines are revenue: the solid black line is the base case, the dotted black line the bull case, the dashed grey line the bear case, from the 2026 estimate through the model years. The short dashed gold line at the left marks the \$131.7 billion of cash and investments at June 30. The bars are the house's schedule, built from each agreement's amount and term, because the filing's September-30 table is blank; the TPU 2028-31 segment is the company's \$125-150 billion estimate at its midpoint. Compare each line with the bar below it; the one thing it shows is that in 2027-29 the commitments run at 73-94% of base-case revenue and above the bear case's, which is the margin compression the filing warns of.

The public filing leaves the aggregate committed spend blank ("more than \$ "), the September-30 contractual-obligations table blank, and the amount due in the next twelve months blank. It names the agreements:

COUNTERPARTY	AMOUNT	TERM	WHAT THE FILING SAYS
Microsoft (Azure, NVIDIA GPUs)	at least \$44.9B	Nov 2026–May 2033	non-cancelable except uncured breach or undelivered tranches
Google (TPU cloud)	at least ~\$114.5B	Apr 2026–Jul 2033	true-up if actual spend falls short
Additional TPU compute	~\$125–150B (company estimate)	delivered 2028–31	depends on capacity delivered and pricing; in no commitments table
Amazon (AWS, Trainium)	at least \$110.0B	May 2026–Apr 2036	true-up if actual spend falls short
SpaceXAI (formerly X.AI; NVIDIA GPUs)	up to \$120.5B	into 2029	~\$3.1B a month once capacity is delivered (expected December 2026); after January and March 2027 either party may cancel on 90 days' notice
Broadcom equipment leases (TPU systems)	~\$161.2B rent ex-interest	five-year leases	finance leases, non-cancelable except default, guaranteed by the parent; restricted cash posted, up to \$10B more possible
Nscale (NVIDIA GPUs)	up to ~\$44.6B	into 2033	fees accrue as capacity is delivered; Nscale must secure financing
Lambda (NVIDIA GPUs)	~\$41.3B	into 2034	the \$34.9B August agreement ends if Lambda misses its facility lease
AMD (Helios / MI450)	expected to exceed \$20B	2027 tranche	no fixed dollar commitment; AMD will buy up to \$5B of stock as milestones are met

That is **\$657 billion named**, or \$782–807 billion with the TPU estimate. The filing's own category totals run higher because they include data-center leases, colocation and smaller agreements: about \$265.6 billion of fixed or minimum commitments at June 30, \$321.4 billion of leases signed but not yet commenced (from \$50.2 billion six months earlier), and about \$245 billion of new minimum obligations plus \$21 billion of short-term agreements entered after June 30, together \$853 billion. For scale: the last table the filing fills in, at December 31, 2025, totaled \$54.6 billion. The \$518 billion the Reuters leak reported on September 28 appears in no draft, no exhibit and not in the public filing (section 12 of the S-1 draft; the leak's counterparty figures match the August and October drafts, and the public figures are larger). Akamai's \$11.6 billion commitment, announced September 24, is not named.

Scheduled by the house from each agreement's terms (ramping monthly rates for the cloud contracts, SpaceXAI at \$3.1 billion a month from December 2026 to its cap, the Broadcom leases from ready-for-service), the named ladder is **\$25 billion in 2026, \$93 billion in 2027, \$156 billion in 2028, \$160 billion in 2029, \$119 billion a year in 2030–31** and \$50 billion in 2032. Compute expense was \$7.3 billion in 2025 and \$8.9 billion in the first half of 2026. The company describes the reason: it is "building out our own compute and infrastructure and vertically integrating across the compute supply chain," under a program it calls Anthropic Direct that leases chips and hosts them in dedicated facilities; its first data center opened this year.

How it is paid for, and the strings. Beyond the \$131.7 billion of cash: Google's \$20 billion facility to December 2030 (draws tied to third-party compute agreements and revenue targets, settled in Class N shares at the 30-day average price, a 14.99% ownership cap, reduced if Google buys \$10 billion at the IPO); Amazon's \$15 billion for 30 months after the IPO, in \$5 billion steps as capacity milestones are met, with a 12% cap; Broadcom's \$42 billion solely to pay rent under the 2026 lease, at 9% paid in kind and convertible at 90% of the average price, rising to 13% and 85% "during a period in which certain revenue performance conditions are not met," with 150% of principal in cash on a change of control; AMD's up to \$5 billion of stock; a \$15 billion revolver with a \$5 billion minimum-liquidity covenant from the fourth quarter of 2026; and the IPO, for which the wires say up to \$100 billion and the filing says nothing. The parent guarantees its two leasing subsidiaries; up to \$0.8 billion of data-center guarantees and a \$500 million power-contract guarantee sit beside them. The filing says the build-out will "be a significant use of net cash in operating and investing activities for the foreseeable future" and that "net cash used in these activities may increase regardless of whether our financial performance improves."

In the model, compute cost in each year is the larger of a target share of revenue and the cash floor of the ladder, delivery-adjusted by case: the firm contracts (Microsoft, Google, Amazon, the Broadcom leases) in full in every case; SpaceXAI in full through 2027 and then in full, at 85%, or not at all (bull, base, bear) because either party can end it on 90 days' notice; the capacity-dependent agreements (Nscale, Lambda, AMD) in full, at 85%, or not at all; the TPU estimate in full, in full, or at half. The floor binds in 2027–29 in every case and through 2031 in the bear case. In the base case the company loses \$24 billion, \$47 billion and \$17 billion in 2027–29 as the ladder lands, then earns 19% in 2030 and 23% in 2031. In the bear case it loses money through 2031 and burns about \$156 billion in 2027–31, more than its cash: that


is the case in which the facilities, the revolver and the IPO proceeds are not optional. The floor costs \$5.22 a share, weighted, against a model that lets compute be a share of revenue.

△ OWNERSHIP AND CONTROL

Who owns it and who controls it

EXHIBIT 4 — ECONOMICS AGAINST VOTES, AT THE IPO

Who holds the economics, and who holds the vote, at the IPO (before the new shares are sold)



Economic count 16,281 million shares (Class A 9,410M, Class N 6,871M); the seven co-founders hold 1,686M of the Class A. The Founder LLC's 50.1% is fixed regardless of time elapsed or future issuances; the Class A vote is capped at 49.9% in aggregate.

How to read it. Three bars, each the full width: the top is economic ownership of the shares that will exist at the IPO before any new shares are sold, the middle is the vote on matters put to stockholders generally, the bottom is board seats. Gold is Class A held by everyone but the founders, dark gold the founders' own Class A, pale grey the non-voting Class N, black the Class F share (one share, 50.1% of the vote) and the Class T share (one share, four of seven directors). Read each bar left to right as shares of 100%; the one thing it shows is that the public's class holds 58% of the economics, under half of the vote, and elects a minority of the board.

The count. The income statement's weighted average of 2,025 million shares counts only the common. At the IPO the preferred (8,559 million as converted, after the ten-for-one split), the Amazon notes (1,840 million Class N at fixed prices of \$2.09 and \$5.05, converting "at the lower of" those and the IPO price, within Amazon's 12% cap), the options (1,797 million at an average exercise price of \$2.89), the restricted stock units (1,050 million, which vest on effectiveness) and the Public Benefit Fund's 847 million Class N shares join the 2,188 million common outstanding: **16,281 million economic shares**, or 16,427 million with the 146 million issued after June (new unit grants, the October donation-matching settlement, two acquisitions). That is the count this note divides by; it reconciles to the filing's own "fully diluted" basis of 16,937 million once the unissued plan reserves are added (847 million is 5.00% of it). Not counted: the IPO's new shares, Google's \$10 billion placement, AMD's \$5 billion and any facility draws, all blank or future. The Series H price of \$58.90 is \$968 billion on this count, the "\$965 billion" reported in May.

The wedge. Class A (one vote; 9,410 million at the IPO) holds 57.8% of the economics but at most 49.9% of the general vote, by construction. Class N (6,871 million) holds 42.2% of the economics and no vote; it is the strategic investors, the Amazon notes and the fund, and converts to Class A only on transfer to an unaffiliated holder. Class F, one share held by the founders' LLC, carries 50.1% of every general vote, a share the filing says "will not be impacted by future issuances of our capital stock that may cause economic or voting dilution." Class T, one share held by the LTBT, elects the Class T directors, "a majority of the authorized members of our board of directors"; the founders' letter puts it as four of seven. The company is a Nasdaq "controlled company" by virtue of the trust. What a Class A holder can decide on matters put to stockholders generally: nothing the Founder LLC opposes, and not the board majority.

The trust and the sunset. The LTBT is a Delaware purpose trust of five trustees, one seat vacant: Neil Shah (chair; also an initial manager of the Public Benefit Fund), Richard Fontaine, Ben Bernanke (since June) and Christopher Stubbs (since October), each paid \$400,000 a year by the company, which funded the trust with \$8.4 million this year. Trustees may hold no stock; they choose directors and "neither the LTBT nor the Trustees have decision-making authority in the day-to-day management or strategy setting." The trust's terms change only by 80% of the board and 80% of the trustees, or, over the trustees' objection, the founders' LLC plus a majority of Class A plus a vote of an employee class (Class E) to be issued, or 85% of Class A if the founders' share is gone. The founders' share is redeemed 450 days after fewer than three

"Continuing Founders" remain for 90 consecutive days, or earlier if a supermajority of the continuing founders agrees to redeem it or the Founder LLC dissolves; a founder stops counting on resignation, death, removal for cause, an approved leave, or holding less than 2% of what they held at the IPO. Seven today; the structure survives the loss of four. The draft-by-draft record shows the structure was assembled during the SEC process: no Class F existed in the June 1 draft; it entered on July 10; Class E, the amendment threshold and a first founder sunset on August 10; the 90-day test and the 450-day transition in the public filing.

Who benefits most. The non-voting strategic class holds 42% of the economics, Amazon and Google first among them (their stakes are not disclosed; the "21% and 15%" in the press is reporting). The seven founders hold 10.4% of the economics, 1,686 million Class A shares, and have pledged 80% of it to philanthropy (9% of fully diluted shares, the letter says) while holding, through one share, a majority of every general vote. The public holds Class A economics and a vote that cannot decide. The Public Benefit Fund holds 5%, and the board may issue up to 2% more each time the share price closes 50% above the last milestone for 90 of 120 trading days, starting from the IPO price. Prices paid, per post-split share: \$3.00 (2024), \$5.61 (early 2025), \$14.10 (September 2025), \$25.91 (February 2026, and the April employee tender), \$58.90 (May–June 2026).

The Amazon notes, and the loss they explain. Amazon's \$6.2 billion of convertible notes converts at the IPO into 1,840 million Class N shares at \$2.09 and \$5.05 a share, either side of the \$3.00 paid in the 2024 round, 11% of the pre-offering economic count, subject to Amazon's 12% ownership cap, with any excess rolling over and converting quarterly at the lowest of the fixed price, the IPO price and the ten-day volume-weighted average price. At the Series H price those shares are worth about \$108.4 billion, against the \$108.6 billion the balance sheet carries the notes at; the \$72.4 billion charge in the first half is the notes' fair value rising with the company's valuation, from \$41.7 billion at the start of the year, with \$5.5 billion of notes converting along the way. In the house's reading it is a transfer inside the capitalization table rather than a cost of the business: the operating loss in the half was \$0.7 billion, and the second quarter earned an operating profit of \$974 million. The charge stops on the portion of the notes that converts; any balance the 12% cap leaves as rollover notes keeps being marked until it converts.

When the holders may sell. The lock-up is not one cliff but a schedule of timed releases in four holder groups, each step a percentage of that group's own lock-up holdings (shares, options, units and convertibles). The filing leaves blank the share count of every step and the fully diluted share of the two stockholder groups, and gives no size for the other two; those are filled at the amendment. The percentages, in the filing's underwriting section, are not blank. Counted from the date of the prospectus and from each results release:

WHEN	NON-FOUNDER DIRECTORS AND OFFICERS, AND CERTAIN EMPLOYEES AND CONSULTANTS (180-DAY)	OTHER 180-DAY HOLDERS	HOLDERS ABOVE 0.5% OF FULLY DILUTED SHARES (365-DAY)	FOUNDERS (365-DAY)
The first trading day	10%	—	—	—
Days 27 and 48	—	6% each	1.8% each	—
Days 62 and 76	—	7% each	2.1% each	—
Second trading day after the 2026 annual results	35%	12.5%	3.75%	—
Days 13, 27, 41, 55 and 69 after that release	—	9.5% each	2.85% each	—
Day 180, or the second trading day after the first-quarter 2027 results if sooner	The rest (55%)	The rest (14%)	4.2%	Up to 29%
Days 34, 48, 62 and 76 after the first-quarter 2027 results	—	—	5.6% each	—
Second trading day after the second-quarter 2027 results, then days 27, 41, 55, 69 and 83 after it	—	—	7.0%, then 6.3% each	—
Day 365, or the second trading day after the third-quarter 2027 results if sooner	—	—	The rest (9.1%)	The rest (71%)
Free by day 180, on the estimated release dates	100%	100%	30%	up to 29%

The reader's use for it: supply arrives in steps from the first week and no single day carries it all; the two 180-day groups are fully free at the half-year mark, the holders above 0.5% are nine-tenths free before the year is out, and the founders may sell up to 29% at the half-year mark, about 489 million shares and share-equivalents, 3.0% of the pre-offering economic count, a ceiling rather than a plan, with each founder having made a voluntary, revocable pledge to dedicate 80% of personal equity to philanthropy. The founders' remaining 71%, about 1.2 billion share-equivalents on the beneficial-ownership count, is the last step, taken on the same trigger as the 9.1% remainder of the holders above 0.5%; which step is largest depends on group sizes the filing leaves blank. Which holders sit in each group, how many shares each step is, and how many days of trading each would take, are the amendment's and the first prints' to settle.

THE MOAT AND THE GATE

The moat, and what erodes it

The filing's own evidence that frontier position is the moat: models released in the six months to September produced over 80% of third-quarter revenue; "with each release, we have seen most customers utilize the latest version"; 250 model versions, features and applications launched in nine months; the cyber benchmark ExploitBench went from 40% (Opus 4.8) to 78% (Mythos 5) to 91% (Opus 5.5) in about four months. Each model helps build the next ("closing the loop"), and Claude writes over 90% of the company's own code. Four chip families (AWS Trainium, Google TPU, NVIDIA through Microsoft, SpaceXAI, Nscale and Lambda, AMD from 2027) run under an orchestration layer the company says "took years and significant investment to build" and "reduces dependence on any single vendor." Distribution runs through all three clouds, systems integrators, over 150,000 certified partner practitioners, and the case studies in the filing (Salesforce: over 60% of 83,000 employees on Claude, output per developer up 151%; Mercado Libre: 40% of 2025 production code; Verizon 33,000 engineers; Lilly 50,000 seats). Ramp's card data had 43.8% of US businesses paying Anthropic against 40.6% paying OpenAI in September 2026, and 60% of token spend in October against 36%; a card-spend sample, not the market.

The same cadence has a brake its owner controls. Mythos was held to a gated preview because of "its advanced capabilities in cyber vulnerability discovery"; Fable ships with safeguards that route some requests "to a less capable Claude model"; the company "recently chose to delay pre-training a new frontier base model for more than a month" for weight security; and in June 2026 the Commerce Department "imposed worldwide export restrictions on our Fable 5 and Mythos 5 models, which we disabled for all customers until those restrictions were lifted." The filing says future models "exhibiting such capabilities may be subject to access constraints" and that the commercial viability of gated programs "may depend on sustained governmental and industry coalition support." This is the one moat whose owner has said in writing it will narrow it when it judges it must.

What erodes it: price, first. On Ramp's token-spend sample, the blended price paid for Anthropic tokens fell 19% in a month to \$0.80 per million in the week to October 11; open-weight models price below it (DeepSeek V4 Flash at \$0.14 and \$0.28 per million tokens in and out, about a quarter of that blended price; GLM-5.2 at \$1.40 and \$4.40 and Qwen 3.6 at about \$0.33 and \$1.95 are not cheaper than that blended price on the output leg); the BenchLM index has frontier prices up 61% in a year and mid-tier prices down 41%, the split Anthropic's four-class pricing exploits and open weights attack from below. The filing concedes the mechanism in its own words: open-weight models "could exert pricing pressure on our offerings", and customers "may conclude that incremental performance gains at the frontier do not justify the additional cost". So far the threat is to price, not to share: on Ramp's adoption data from more than 85,000 US businesses, 0.3% paid DeepSeek in September 2026 against 43.8% paying Anthropic, and the model carries the threat as the compute share of revenue, 60% in 2027 in the bear case. Second, the clouds that distribute Claude also "develop and deploy competing AI models," and the distribution deals "give them visibility into our pricing and commercial terms." Third, OpenAI (a reported \$852 billion mark in March, seeking about \$1.4 trillion, listing "not 2026") and Google. Fourth, talent that can "join or found competing companies immediately."

What the law and the regulator want

The directors of a Delaware PBC must balance stockholders' pecuniary interests, the interests of those materially affected, and the stated purpose; "there is limited judicial precedent." On February 27, 2026, "in response to our publicly stated position that our technology should not be used for mass domestic surveillance or fully autonomous weapons," the President directed federal agencies to stop using the company's technology and the Department of War designated it a supply-chain risk; the company won most of its case in the Northern District of California on August 27, lost 2-1 in the D.C. Circuit on September 25, and has asked for rehearing. Government revenue is not quantified. The June export restriction (above) "could recur." The FTC issued a civil investigative demand in October on consumer protection, privacy "as well as the safety and security of our products," and Axios reported on October 26 that the FTC has demanded testimony from Anthropic executives. The external evaluators the company indemnifies, "without a cap on our liability" in some cases, are unnamed in the filing; anthropic.com named a six-person METR team with employee-like access on October 15. The copyright docket: Bartz settled for \$1.5 billion (approved July; \$1.2 billion still accrued at June 30), the first Concord case was argued on summary judgment on October 21, a second wave of music-publisher suits names two executives, Reddit's contract claims survive, four Canadian actions and the Bartz opt-outs continue; none is quantified. The company endorsed California's SB 53, tests with the US and UK AI safety institutes, and faces EU general-purpose-model enforcement from August 2, 2026. The filing mentions lobbying once and gives no figure. The company's own reports to the Senate under the Lobbying Disclosure Act show federal lobbying spend rising from \$0.36 million in the first quarter of 2025 to \$2.11 million in the third quarter of 2026 (\$1.56 and \$1.97 million in the two quarters before it, with one dip in the fourth quarter of 2025), with nine lobbyists named in the latest report and defense, energy and consumer safety among the issues listed; reported giving to a political group this year is \$40 million.

THE STAGES

How long each line grows faster than normal

Every length below is an argument a reader can dispute; the sensitivities in §9 give the value of each alternative. The rule is the house's: a stage lasts as long as the thing that drives it, dated by an observable, and is checked against base rates (across US companies in 1951–98, above-median sales growth survived five years running in 6.3% of firms and ten years in 0.9%).


LINE	WHAT DRIVES THE FASTER GROWTH	HOW LONG, AND WHY	THEN	GROWTH WHILE IT LASTS (BEAR / BASE / BULL, FIRST YEAR → LAST)
Consumption (API, Claude Code, Cowork usage, committed contracts; 80% of revenue)	Enterprise diffusion against a compute ceiling: the filing says demand "exceeds available compute"; Fortune 500 customers spend under 0.1% of their expenses on Claude; 11,900 customers above \$100,000 from 1,500; NRR about 900%	5 years, 2027–31: the compute already contracted is delivered on a schedule that closes in 2031: SpaceXAI in full by December 2026, the Broadcom TPU systems on five-year leases from 2026–27, the \$125–150 billion of extra TPU capacity in 2028–31, the cloud contracts to 2033 and 2036. Five years of growing into capacity already bought	Five years converging to the terminal rate, as penetration saturates, price per token keeps falling and open weights compete	55% → 6% / 120% → 15% / 180% → 20%
Subscriptions (Pro, Max, Team, Enterprise seats; 20%)	Diffusion of Claude Code and Cowork to a knowledge-worker population IDC puts above one billion by 2028; Cowork merged into Claude.ai in September; seats are the top of the consumption funnel	4 years, 2027–30: the adoption curve the company has run once already: Claude Code from preview (February 2025) to "embedded ... across major organizations" in about eighteen months, Cowork faster from January 2026; four years carries the seat business two product cycles past IDC's 2028 milestone	Six years converging, a longer tail because seats are prepaid and renew	45% → 8% / 90% → 15% / 120% → 20%

Terminal growth is 4.0% a year in the base and bull cases, the Federal Reserve's September 2026 longer-run projection of real growth (2.0%) plus inflation (2.0%), which a company that has become a large share of knowledge-work spending

cannot outgrow forever; 3.0% in the bear case, for a company losing share. Every fade lands exactly on its terminal rate.

EXHIBIT 5 — GROWTH BY YEAR ON THE ARGUED STAGES, AND WHAT THE ANCHORS NEED

Revenue growth by year in the three cases on the argued stages, and the extra growth each price anchor needs



How to read it. Three lines of year-on-year revenue growth from 2027 to 2036, in percent up the side: dark gold is the base case, pale gold the bull case, grey the bear case. The shaded band marks stage one, the five years in which contracted compute is still being delivered; after it every line converges to its terminal rate. The right-hand column gives the extra growth, in points added to every stage-one year of every case, that the Series H price and the two reported IPO figures require, and the five-year compound rates. Compare the base line's height in 2027 with the figures at the right; the one thing it shows is that the prices in the wires need faster growth within the stage, not a longer one.

Against the company's own talk. The base case reaches \$113 billion in 2027, \$167 billion in 2028, \$303 billion in 2031 and \$450 billion in 2036. Reuters reported in August that the IPO case "rests on \$190–200B revenue forecast for 2028": the base case sits about 14% below it, the bull case (\$244 billion) about 25% above, the bear case (\$92 billion) at half. The FT's reported investor expectation of a \$100–120 billion exit run-rate for 2026, and the company's own estimate at its October 15 investor meeting that the run-rate passes \$100 billion by year-end (from the \$67 billion reported at the meeting, as Damodaran relayed it on October 22), imply a fourth quarter of \$25–30 billion; the house's bull-case fourth quarter is \$23.8 billion, base \$19.8 billion, bear \$17.0 billion, because the filing says growth decelerates, the year-end dip is 20% over two weeks, and 5–7% of revenue is being purged. **The length the price implies:** at the house growth rates no stage length up to nine years reaches the Series H price (nine years gives \$58.39, fifty cents short, because the stage ends at a 15% rate); the price is not a longer boom but a faster one: 3.2 points more growth in every stage-one year of every case for Series H, 12.8 for \$1.5 trillion, 19.4 for \$2 trillion.

THE VALUE

The value

Method. Free cash flow to the firm over ten explicit years (2027–36) on a 2026 stub built from the first-half actuals and the implied third quarter, discounted at mid-year, with one cost of capital in every case, stock compensation expensed, terminal value by Gordon growth, cash of \$131.7 billion plus a \$2 billion house estimate of second-half cash generation added, the \$1.2 billion Bartz accrual and \$0.4 billion of finance leases deducted, divided by 16,427 million shares; the terminal value is discounted from the end of 2036 (discounting it from mid-year too would add \$1.85 a share, \$52.58). A dividend model has nothing to discount; free cash flow to equity was negative in every year through 2025 and the capital structure changes at the IPO; a multiple on revenue growing 1,300% is a growth assumption in disguise, so it is a cross-check below, never blended.

Cost of capital 12.25%. The ten-year Treasury at 5.00% on October 23 plus a beta of 1.51 times the house's 4.8% equity risk premium; the beta is the midpoint of the hyperscalers' median of 1.24 (Microsoft 1.09, Alphabet 1.15, Amazon 1.32, Oracle 1.58) and the AI pure plays' 1.78 (NVIDIA 1.58, Palantir 1.99, CoreWeave 1.56, Nebius 2.22), two-year weekly betas against the S&P 500 after the Blume adjustment; no interest-bearing debt survives the IPO beyond any Amazon

rollover notes held back by the 12% cap, which convert quarterly, and \$0.4 billion of finance leases. The house used 10% for SpaceX, a company with a cash-generating Starlink business and a listed market history; Anthropic has neither, and a beta set from its own peers is the honest substitute. Every result is shown at 10%, 11%, 12.25% and 13%. Tax is 21% after the \$8.2 billion of federal net operating losses the filing reports at December 31, 2025, plus the first-half 2026 loss.


Cost structure, base case (share of revenue, 2027 → stage end → terminal). compute 52% → 42% → 38% (first-half actual 55%, second quarter 48%; at maturity about 15 points to serve and 25 to research, a hyperscaler cloud's cost structure); channel and marketing 17% → 12%; cash personnel 6% → 5%; stock compensation 10% → 6%; other 7% → 4%; public-benefit spend 1% throughout. Terminal operating margin 34% base, 41% bull, 23% bear.

	BEAR (25%)	BASE (50%)	BULL (25%)
FY2026 revenue (first half actual + implied third quarter + house fourth)	\$49.3B	\$52.6B	\$57.1B
2027 / 2031 / 2036 revenue	\$75B / \$118B / \$145B	\$113B / \$303B / \$450B	\$153B / \$514B / \$835B
Operating margin 2031 / terminal	-15% / 23%	23% / 34%	30% / 41%
Years the commitment floor binds	2027-31	2027-29	2027-29
Terminal value as a share of operating value	331%	75%	66%
Equity value	\$165B	\$736B	\$1,695B
Per share	\$10.07	\$44.81	\$103.21

Weighted 25/50/25, fair value is **\$50.72 a share, \$833 billion**. The weights are the house convention for a triad whose base case is the most likely path and whose tails are judged equally likely; a reader who weights the bull case at a third and the bear at a sixth gets \$58.49, within a dollar of the Series H price. The bear case's terminal value exceeds its operating value because the explicit years are losses: enterprise growth slows to 55% in 2027 and 22% in 2028, the firm contracts land anyway, the company loses \$39 billion, \$33 billion, \$30 billion, \$24 billion and \$18 billion in 2027-31, and the equity is mostly the cash on hand plus a 23%-margin business a decade out. The base case loses money in 2027-29 while the ladder lands, then earns 19%, 23%, 25%, 27%, 30%, 32% and 34%. The bull case grows into the ladder by 2029. **The proponent's case, on its own inputs** (bull growth, a 10% cost of equity, a seven-year consumption stage) gives \$167.62, \$2.75 trillion; the house weights that path at a quarter and discounts it at 12.25%.

EXHIBIT 6 — RETURN ON CAPITAL AGAINST THE HURDLE

Trailing-twelve-month return on invested capital against the 12.25% cost of capital, the three quarters the filing supports



How to read it. One gold point per quarter end, December 2025 to June 2026, the three the filing supports: each is the trailing-twelve-month return on invested capital, in percent up the side (the scale runs from -350% to +25%); the dashed black line is the 12.25% cost of capital this note discounts at, held flat. Invested capital is equity plus debt at the balance-sheet value of the convertible notes, less cash, equivalently operating assets less operating liabilities; the March balance sheet is interpolated. The right-hand column gives the return gross of cash and the average capital base. Compare the slope of the gold line; the one thing it shows is a deeply negative return narrowing by about 100 points a quarter on a capital base of a few billion dollars, because the company rents nearly all of its capital.


The read from the business. Return on capital is -342%, -246% and -138% at the three quarter ends, with the spread to the hurdle widening from -354 to -150 points: the mechanical label is *recovering*. The loss narrowed because first-half revenue grew 1,300% against compute's 283%. The lens is the wrong one for a company whose capital sits in its

vendors' data centers until Anthropic Direct capitalizes it, and when the 2027 deliveries land the series will fall before it rises.

Sensitivities (weighted, \$ a share). Cost of equity 10% / 11% / 13%: \$74.17 / \$61.80 / \$45.63. Terminal growth half a point either way: \$48.37 / \$53.38. Consumption stage three, seven or nine years: \$36.64 / \$54.58 / \$58.39. Stage-one growth ten points a year lower or higher in every case: \$29.34 / \$80.65. Compute share of revenue five points higher or lower every year: \$43.23 / \$57.84. Revenue path at 0.8 or 1.2 times: \$38.08 / \$62.03. No commitment floor at all: \$55.94. End-year instead of mid-year discounting: \$50.08. The single greatest point of sensitivity is stage-one growth; the second is the cost of capital.

EXHIBIT 7 — WHAT MOVES THE VALUE


Probability-weighted fair value a share when one input moves, everything else held



How to read it. Each bar is the probability-weighted fair value a share, read along the bottom in dollars, when the one input in its label is moved as stated and everything else is held. The solid vertical line is the central value, \$50.72; the dashed line is the Series H price, \$58.90. Gold bars are the model's structural inputs (discount rate, stage lengths, terminal growth, discounting convention); grey bars are operating inputs. Compare each bar's reach toward the dashed line; the one thing it shows is that the round's price is within reach of a single input, ten points of stage-one growth or two points on the cost of equity, and that \$90 and above is not.

EXHIBIT 8 — WHAT EACH PRICE REQUIRES

What each price requires: the multiple of the house revenue path (every case, every year) at four costs of equity




How to read it. Three groups of four bars, one group for the Series H price and one for each reported IPO figure; each bar is the multiple of the house revenue path, applied to every year of every case, at which the weighted value reaches that price, at the cost of equity named beside it (pale grey 10%, grey 11%, dark gold 12.25%, black 13%). The vertical line at 1.0 is the house path itself. Under each group is the cost of equity at which the house path alone reaches the price, and the base-case 2031 revenue the multiple implies. Compare bar lengths across groups; the one thing it shows is that Series H sits a little above the house path and \$2 trillion sits far beyond it at any discount rate.

What each anchor requires. The Series H round (\$58.90; \$968 billion) is the house path at an 11.3% cost of equity, or 1.14 times the path at 12.25% (base revenue of \$347 billion in 2031 instead of \$303 billion; 46% compound growth in

2026–31 instead of 42%), or 3.2 points more growth a year, or compute 5.8 points lower as a share of revenue every year. **\$1.5 trillion** is the path at 9.0%, or 1.75 times it (\$529 billion in 2031; 59% compound), or 12.8 points more growth a year, or compute 31 points lower. **\$2 trillion** is the path at 7.9%, or 2.38 times it (\$722 billion in 2031; 69% compound), or 19.4 points more growth a year; no reduction in the compute share reaches it. Damodaran's own break-even, in his August 20 post, is the same arithmetic on a ten-year horizon: at the "rumored pricing for the IPO" of \$2 trillion, "allowing the company premium pricing margins (after-tax operating margin of 30%) and above-average risk (cost of capital of 10%), the company will have to generate close to \$1.2 trillion in revenues, if the AI market matures in ten years, and close to \$2 trillion, if the wait is 15 years"; the house path reaches \$450 billion in 2036. On October 22 he wrote that a case for \$1.5 or \$2 trillion needs "a surge in operating margins to 30% or higher", a story he calls a reach but "a plausible one", and that his own valuation waits for the prospectus.

EXHIBIT 9 — THE VALUE, THE ANCHORS, AND THE PRIVATE MARKS

Equity value on the house count of 16,427 million shares (\$ billions): the cases, the anchors, and the private marks as a separate band



How to read it. Horizontal bars of equity value in billions of dollars on the 16,427 million share count: the three cases and the weighted fair value (grey, dark gold, gold, black), the proponent's case on its own inputs (pale gold), the peer median revenue multiple applied to 2027 base revenue with cash added (pale grey), the Series H round (black), and the \$1.5–2 trillion range the wires report (grey band). Below the rule, as a separate band never blended into the value, the private marks: Anthropic's own rounds from \$183 billion to \$965 billion and OpenAI's reported \$852 billion round and the \$1.4 trillion it is reported to be seeking. Compare each bar's length with the black Series H bar; the one thing it shows is that the round sits between the base and bull cases, \$1.5 trillion just below the bull case and \$2 trillion beyond it.

The peers, as a cross-check. On trailing revenue and the October 23 closes, NVIDIA trades at 17.1 times sales (83% growth, 65% operating margin), Palantir 75 times, Snowflake 21, Microsoft 11.6, ServiceNow 11.4, Alphabet 8.5, CoreWeave 7.9, Oracle 6.5, Amazon 3.5; the median is 11.4. Anthropic's trailing revenue to June (\$19.7 billion) makes a trailing multiple meaningless; on the second-quarter run-rate of \$46.2 billion the house fair value is 15.2 times, Series H 18.1, \$1.5 trillion 29.6 and \$2 trillion 40.4; on 2027 base revenue of \$113 billion, 6.2, 7.4, 12.1 and 16.6 times. The peer median applied to 2027 revenue gives \$1.28 trillion of enterprise value, which is the problem with the method in one line: it pays NVIDIA's multiple on a year that has not happened. The private marks are a separate band: Anthropic's rounds at \$183 billion (September 2025), \$380 billion (February) and \$965 billion (May); OpenAI's reported \$852 billion (March) and the \$1.4 trillion it is reported to be seeking; xAI's \$250 billion inside SpaceX; and the \$1.5 trillion Reuters reported investors discussing on October 15 after third-quarter growth slowed to 45%, down from the \$1.5–2 trillion range in the Wall Street Journal.

What would change our view

Two update points are pre-registered. **First, the amendment that sets the price range and share count:** the house re-bases every "versus price" line on the midpoint, adds the primary shares, Google's \$10 billion placement if exercised and the vested unit settlement to the count, fills the lock-up ladder's share counts, reads the September-30 commitments table against the ladder in \$4 (the first test of the house schedule) and the third-quarter revenue against the \$16.0–17.0 billion range, and publishes the gap. **Second, the first-day close:** one post-print-style update, not a new note, with the close against fair value, the day-one float and days-to-trade. A valuation on the filing day rests on blanks the amendment fills and on a path the first prints test; the number should be expected to move at each update point, and the direction it moves is itself information for the reader.

The house view moves up, toward the price or past it, if:

- **(a)** the third or fourth quarter of 2026, when printed, exceeds \$17.0 billion or \$23.8 billion, the bull case's quarters; the base case is rebuilt on the bull case's 2027 growth. Silence: no quarterly figures by the amendment is scored as not met.
- **(b)** the September-30 commitments table shows less than \$92.7 billion due in 2027, the house ladder's year; the floor is lowered to the filed figure in every case. Silence: a still-blank table is not met.
- **(c)** compute as a share of revenue in the first two public quarters prints below 45%; the base path's 52% for 2027 is cut to the printed figure. Silence: no compute line is not met.
- **(d)** the ten-year Treasury averages below 4.5% for a full quarter; the cost of equity falls half a point, about \$5 a share. Prices always print.

The house view moves down, further from the price, if:

- **(e)** the third or fourth quarter prints below \$16.0 billion or \$17.0 billion, the bear case's quarters. Silence as in (a).
- **(f)** the September-30 table shows more than \$160 billion due in 2029, the house's heaviest year, or an aggregate committed spend above \$1.0 trillion; the floor is raised. Silence: not met.
- **(g)** the 2027 operating margin prints below –22%, the base case's figure, or the company draws the Broadcom rent notes on the 13% and 85% terms that apply when "revenue performance conditions are not met," which must be disclosed. Silence: not met.
- **(h)** the supply-chain-risk designation survives rehearing and government revenue is disclosed above 5% of the total, or a second export restriction on a current model lasts more than 30 days. Silence: no disclosure of government revenue is not scored.
- **(i)** the Founder LLC's share is used to carry a transaction the Class A majority opposes, or the Class T directors remove the chief executive without the Founder LLC's concurrence: a governance test, not a model input; the control section is rewritten.

Two claims in the case for the price carry their own falsifiers. The claim that model gating is a revenue brake is falsified if the next Mythos-class release reaches general availability within 60 days of its gated preview with no disclosed revenue effect. The claim that the ladder is a demand forecast is falsified if the company discloses the offloading or resale of committed compute, which the founders' letter says it will seek "if we overestimate demand."

 SOURCES

Sources and Method

Primary filings (SEC EDGAR, Anthropic, PBC, CIK 2130357). Form S-1 filed October 26, 2026, accession 0001628280-26-067168: the main document (converted to one block per line by the note's own script; every figure in this note is asserted verbatim against that text by `data/s1_facts.py`), all 33 exhibits (charter and bylaws 3.1–3.5; investors' rights agreement 4.2; the Amazon 2026 stock purchase and note facility 4.3, the Broadcom note purchase agreement 4.4 and the Google note purchase agreement 4.5, each with redactions; registration rights 4.6–4.7; equity plans and the donation matching plan 10.2–10.7; leases 10.8–10.9; the amended and restated credit agreement 10.10; the Broadcom master equipment lease 10.11; the proxy agreements 10.12–10.13; the stockholders' agreement 10.14; the auditor letter 16.1; the trust agreement 99.1 and the Founder LLC agreement 99.2), the XBRL, and the 44 images (customer case studies, the founders' letter and the MD&A charts), transcribed in `data/case_studies/`. The six draft registration statements (DRS June 1, 2026;

DRS/A July 10, August 10, September 3, September 30, October 9), converted the same way and diffed sentence by sentence (`data/drafts/`). Statements of operations and balance sheets for 2023–25 (audited: Ernst & Young for 2025, Grant Thornton for 2023–24) and the six months to June 30, 2026 (unaudited).

Company and management statements. anthropic.com: the S-1 announcement (October 26, 2026), the Series F (September 2, 2025) and Series H (May 28, 2026) pages, Project Glasswing (April 7, 2026), the Responsible Scaling Policy v3.0 (effective February 24, 2026), "An update on independent evaluators" (October 15, 2026); darioamodei.com: "Machines of Loving Grace" (October 2024) and "We Must Pace the Frontier" (September 12, 2026). Akamai's 8-K Exhibit 99.1 (September 24, 2026). Quotations from the filing and the letter are verbatim; quotations from web pages are from archived copies that the fetch tool returned as summaries, so none is used verbatim beyond the two essays' short passages.

Market and macro data. Peer closing prices for October 23, 2026 from Robinhood (the official last-session close), shares outstanding from the same source, trailing revenue, operating income, cash and debt from SEC XBRL company facts (each peer's latest quarter; Nebius and SpaceX lack a four-quarter history and are excluded from the multiples); weekly bars 2021–26 for the betas. FRED: DGS10 and the Treasury par curve for October 23, 2026 (5.00%). The Federal Reserve's Summary of Economic Projections, September 16, 2026 (longer-run real GDP 2.0%, PCE inflation 2.0%).

Reporting, never load-bearing alone. Reuters (the September 28 leak; the October 15 third-quarter and valuation reports; the October 26 factbox), the Financial Times and Bloomberg via Reuters and relays, the Wall Street Journal via relays, Fortune, Yahoo Finance, TechCrunch, CBS, Axios, the AP, Damodaran's Musings on Markets (the August 20 and October 22 posts, verbatim from the blog feed), the Senate Lobbying Disclosure Act database (lda.gov: Anthropic's quarterly reports, 2025–26), Ramp's AI Index, BenchLM's token-price index, Menlo Ventures. Private marks each rest on a primary page or two independent wires, listed in `data/private_marks.json` .

Method. The implied third quarter: the filing's "approximately 1,200%" trailing growth and "\$6.6 million" of revenue per employee, with the undisclosed fourth quarter of 2024 bounded at \$97–174 million (`data/unit_economics.py`). The serving margin: Note 8's functional presentation. The commitment ladder: each agreement's amount and term from the MD&A, scheduled by the house (cloud contracts on a monthly rate rising from 25% to 100% of its steady level over 24 months; SpaceXAI on a six-month ramp then \$3.1 billion a month to its cap; the Broadcom leases 3%, 12%, 20% and about 22% a year from 2026; the TPU estimate's midpoint evenly over 2028–31) (`data/commitments.py`). The share count: the June 30 balance sheet, Notes 10, 12, 16 and 17 and the Public Benefit Fund issuance, after the ten-for-one split (`data/shares.py`). The lock-up ladder: "Shares Eligible for Future Sale" and "Underwriting," with every tranche's share count blank (`data/unlock.py`). Return on capital: the house script in manual mode from the filing's own tables, three consecutive quarters, March interpolated and stated (`data/roic_inputs.py`). The model: `model.py` ; the workbook `TON618_ANT_Model.xlsx` carries every input with its source, the cases, the stages, the ladder, the share count and the return-on-capital series. The house's model gate (mid-year discounting; stock compensation expensed once; one discount rate across cases; a terminal-growth ceiling that warns above 4% and fails at 5%; fades landing on the terminal rate; every stage argued and stated for the reader; the return-on-capital series) passed with no failures and one warning, verbatim: "bear: terminal value = 331% of operating value — state it in the note and sensitize `g_terminal/ke`," which §9 does.

Not in hand. The third-quarter figures, the price range, the share count and the September-30 commitments table, all blank in the filing; the per-tranche lock-up counts; the Bloomberg, Wall Street Journal and Financial Times originals. The Q3-2026 revenue of \$16.8 billion is reporting and is not a model input. This note is a valuation on the day of the filing; the two update points above are where its "versus price" lines will be set.

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)(1)(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 note discusses Anthropic, PBC (proposed ticker ANT), whose shares have no public market as of the report date, and, as comparables or counterparties, Microsoft (MSFT), Alphabet (GOOGL), Amazon (AMZN), Oracle (ORCL), NVIDIA (NVDA), Palantir (PLTR), CoreWeave (CRWV), Nebius (NBIS), Snowflake (SNOW), ServiceNow (NOW), SpaceX (SPCX), Broadcom (AVGO), AMD (AMD) and Akamai (AKAM); the Fund holds no position in ANT or in any Anthropic security, and ANT is not a Fund holding or a BTC-correlated instrument. It makes no recommendation on any of them. 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.