

PANACEA VC

PANACEA VC
COMPANY THESIS

COMPANY THESIS NOTE

Anthropic: Growth Is Real

*The “Circularity” Isn’t. Understanding the revenue engine
behind the capital intensity*

PREPARED AUGUST 2026

PANACEAVC.COM/RESEARCH

The AI capital stack really is circular in places. Anthropic's revenue is not. Enterprises pay for the intelligence the compute produces.

-
- 01 Executive Summary
 - 02 Why We're Writing This Note
 - 03 The Growth Number Nobody Is Disputing
 - 04 Where the Revenue Actually Comes From
 - 05 From Chatbots to Agents
 - 06 Capital Structure vs. Operating Revenue
 - 07 Why Anthropic Needs So Much Compute
 - 08 Renting, Not Owning
 - 09 The Nvidia Question
 - 10 The Legal Overhang Has Eased
 - 11 The SpaceX Relationship
 - 12 This Is Not the Dot-Com Bubble
 - 13 The IPO, and the Number That Matters
 - 14 What Could Go Wrong
 - 15 The Bottom Line
-

Executive Summary

Circular financing has become one of the more persistent concerns in AI investing this year: chipmakers investing in the labs, the labs renting the chips back, everyone seemingly writing checks to everyone else. That's a fair concern to raise across the sector. Applied to Anthropic specifically, however, we think the numbers tell a different story.

In the second quarter of 2026, Anthropic generated \$11.5 billion in revenue, up roughly fourteenfold from the \$787 million it recorded in the same quarter a year earlier. That put it ahead of OpenAI's \$6.7 billion for the first time. Its annualized run-rate reached \$65 billion by the end of July, on its way, by the company's own projections, to somewhere between \$190 billion and \$200 billion by 2028. It also posted its first positive adjusted operating income this quarter, which matters more than the top-line number: this is no longer purely a story of a company burning capital to buy growth.

That revenue is generated from enterprise customers paying to use Anthropic's products, rather than from the financing relationships surrounding the company. Eighty to eighty-five percent of Anthropic's revenue now comes from B2B enterprise API usage, and Claude Code, the company's agentic coding product, is alone contributing an estimated \$8 billion of the quarterly total. The distinction is important: this is revenue from customer usage, not financing.

The key question is whether the circular-financing narrative, however catchy, actually describes what's happening in Anthropic's capital structure. We don't believe it does, and we remain constructive on the position even as the broader AI trade gets noisier.

Why We're Writing This Note

One question has increasingly come to the forefront of investor discussions: OpenAI, Microsoft, Google, Nvidia and Anthropic all appear to be lending to, investing in, and buying from one another, raising concerns that the AI ecosystem is simply one pool of capital circulating in a loop. Is the AI boom becoming a bubble?

It's a reasonable question. Capital expenditure across the industry is enormous, valuations are high, and the financial relationships between model developers, chip companies, cloud providers and data-centre operators are genuinely getting harder to untangle. But in our view, investors risk focusing on the wrong part of the equation. The existence of substantial financing and infrastructure relationships between AI companies does not, by itself, mean the underlying businesses are artificial or that their revenue is circular.

Anthropic is a useful case study, because it's spending extraordinary amounts on compute for a fairly ordinary reason: demand for its products is growing extraordinarily quickly. The relevant question is therefore not simply who is financing the compute, but who is ultimately paying for the intelligence that compute produces. Increasingly, the answer is enterprise customers, and that distinction is the reason we remain constructive on the position.

The Growth Number Nobody Is Disputing

The starting point is the reported revenue growth. Anthropic generated more than \$11.5 billion of revenue in Q2 2026, according to documents seen by Bloomberg, up from approximately \$787 million in the same quarter a year earlier, a more than fourteenfold increase. Q1 2026 revenue was roughly \$4.73 billion, so the sequential growth alone was about 2.4 times. That put Anthropic ahead of OpenAI's \$6.7 billion for the same period, the first quarter its revenue has exceeded OpenAI's outright.

◆ FIGURE 1 · QUARTERLY REVENUE

QUARTER	QUARTERLY REVENUE
Q2 2025	\$787M
Q1 2026	\$4.73B
Q2 2026	\$11.5B

POINT IN TIME	ARR RUN-RATE
END OF 2025	~\$9B
MAY 2026	\$47B
JULY 2026	\$65B

On these numbers, Anthropic’s share of the B2B AI market now stands at approximately 34.4%, just ahead of OpenAI’s 32.3%. Calling this the fastest-growing business in history would be overreaching on a single dataset. But it’s a reasonable claim that it’s among the fastest revenue ramps ever recorded for an enterprise technology company and that pace is precisely what the circular-financing critique tends to overlook.

The most striking feature of Anthropic today is not its valuation. It’s the rate at which customers are adopting its products. The company is not struggling to find demand. It’s struggling to acquire enough computing capacity to satisfy the demand it already has, which is a very different problem to have.

THE GROWTH NOBODY DISPUTES

14×

*Anthropic's quarterly revenue, second quarter of
2025 against the second quarter of 2026. \$787
million became \$11.5 billion.*

BLOOMBERG · DOCUMENTS SEEN · PRELIMINARY

Where the Revenue Actually Comes From

Perhaps the most important fact about Anthropic's growth is where it's coming from. The company has built its business disproportionately around enterprise customers and developers rather than consumer subscriptions: approximately 80% to 85% of revenue is now B2B enterprise API usage. Enterprises don't typically spend billions of dollars because they like the story around artificial intelligence. They spend because the technology is becoming embedded in how the business actually runs.

That shift is most visible in software development. Claude Code, Anthropic's agentic coding product, is reported to have contributed roughly \$8 billion of the \$11.5 billion quarterly total on its own. This means roughly two-thirds of the company's revenue is coming from a single product category that barely existed in its current form two years ago.

Coding also happens to be one of the few places where the economic return on AI is genuinely easy to measure. If a system lets an engineer ship faster, catches vulnerabilities before a human would, or operates semi-autonomously across a codebase, the value is tangible and the ROI conversation with a CFO is short. That's a meaningfully different sale than a consumer chatbot subscription, where the economic value is much harder to pin down.

From Chatbots to Agents

We think this is one of the more important shifts in the current AI cycle, and it's easy to miss if you're only watching the valuation headlines. The first generation of generative AI products was largely people asking questions and a model generating answers. The generation now scaling commercially is different: AI systems are increasingly being asked to perform work, not just describe it. That includes:

- ◆ Writing and debugging software
- ◆ Reviewing large codebases
- ◆ Identifying cybersecurity vulnerabilities
- ◆ Automating business processes
- ◆ Conducting research and analysing documents
- ◆ Operating across enterprise systems
- ◆ Executing multi-step workflows with limited supervision

An agent that can replace or augment hours of human work has a clearer economic value proposition than an interface primarily used to access information.

Anthropic's enterprise traction is an early signal that businesses are moving from AI experimentation toward broader deployment.

“*The relevant question is not who is financing the compute, but who is ultimately paying for the intelligence that compute produces.*”

CAPITAL STRUCTURE IS NOT OPERATING REVENUE

06

THE
DISTINCTION

Capital Structure vs. Operating Revenue

This is where we think much of the current debate collapses two different questions into one. There's no question the AI ecosystem is full of interconnected capital: Nvidia invests in AI companies, cloud and infrastructure providers invest in AI companies, AI companies buy Nvidia hardware, cloud providers rent them compute, and data-centre operators finance the buildings in between. Looked at from a distance, it can genuinely look like money moving in a circle.

But that observation answers a different question than the one that actually matters. “Who provides the capital required to build the infrastructure?” and “who ultimately pays Anthropic for its products?” are not the same question, and the first can be highly interconnected without the second being artificial in any way.

It's coming from enterprises and developers paying to use Claude and Claude Code in production. Parts of the capital stack may be circular, but that does not make the underlying customer demand circular. In Anthropic's case, we believe the reported revenue mix is overwhelmingly driven by genuine enterprise demand rather than by the strategic investors that sit on the capital side of the ledger.

Why Anthropic Needs So Much Compute

Anthropic's compute requirements aren't evidence of a financing trick. They're a direct consequence of the economics of frontier AI, which run in one direction: advanced chips, then data-centre capacity, then training compute, then inference capacity, then enterprise AI applications, and finally customer revenue.

The company cannot simply opt out of spending on compute if customers are demanding more Claude capacity than it currently has. The scale of its infrastructure commitments is therefore better understood alongside the growth in customer demand rather than viewed in isolation. It's worth itemizing what's actually been committed:

- ◆ \$45 billion over six years to rent AI computing power from Nscale's West Virginia campus, roughly 460 megawatts of capacity, built around Nvidia's new Vera Rubin chips
- ◆ \$1.25 billion per month through May 2029 to SpaceX for compute capacity across the Colossus and Colossus II infrastructure. Importantly, the underlying SpaceX agreements allow either party to terminate after the initial three-month period with 90 days' notice, giving Anthropic considerably more flexibility than the headline three-year commitment suggests.
- ◆ Tens of billions of dollars in AI servers from AMD, alongside an AMD investment of up to \$5 billion into Anthropic itself
- ◆ More than \$60 billion in total data-centre rental commitments over the next three years

These are substantial obligations, but they also need to be viewed in the context of a company whose reported revenue is expanding rapidly and whose customers are demanding more compute capacity. The more straightforward explanation is that Anthropic is committing infrastructure because demand for Claude and Claude Code is growing faster than its existing capacity.

Renting, Not Owning

There is a second argument worth making about the structure of these commitments, not just their size. Compute prices are unlikely to remain at today's levels indefinitely. The scale of investment now flowing into AI chips, data centres and power infrastructure is creating a significant expansion in available capacity across the industry. If that buildout eventually produces more supply than the market needs at current prices, the cost of compute could fall materially.

That would actually work in Anthropic's favour if the company has structured a meaningful portion of its compute requirements through rentals rather than long-lived owned infrastructure. A company that builds and owns large amounts of capacity at today's construction and hardware costs is more exposed to being locked into a high-cost asset base. A company that rents capacity can potentially benefit as market prices fall, particularly where contracts provide termination or renegotiation flexibility.

There's another way to read the size of these commitments. If Anthropic were struggling to generate demand, committing tens of billions of dollars to additional compute would be reckless. If demand is growing faster than available capacity, the same decision becomes strategic.

The 2028 revenue projections remain ambitious and should be stress-tested. But the investment question is not simply whether Anthropic is spending too much on infrastructure today. It is whether revenue and gross profit can grow faster than the cost of the compute required to produce them. If compute prices decline while enterprise usage continues to expand, the operating leverage could be considerable.

09

THE SUPPLIER

The Nvidia Question

Nvidia is the dominant supplier of high-end AI accelerators, while AMD is expanding its position at the frontier end of the market. Anthropic needs enormous quantities of advanced hardware to operate its models. Anthropic buying compute from Nvidia or Nvidia-backed infrastructure providers is therefore not inherently circular, any more than an airline buying aircraft from Boeing is circular because Boeing may also have a financial relationship with the airline.

The underlying chain remains straightforward: Anthropic buys compute, Anthropic provides AI services, and enterprises pay Anthropic. The capital structure around that chain deserves scrutiny, but it shouldn't be confused with the operating economics of the business itself.

Nvidia's position is also becoming broader than simply supplying the chips on which frontier models run. Reports this week said Nvidia was in talks to acquire Hugging Face, the widely used open-source AI platform, for approximately \$12.9 billion. The reported transaction would extend Nvidia's reach further into the software and open-source model ecosystem, and it illustrates how quickly the AI infrastructure stack is consolidating around a small number of powerful companies.

For Anthropic, that makes supplier relationships worth monitoring, particularly as the company becomes one of the largest purchasers of AI compute. But it does not, by itself, make Anthropic's revenue circular. The critical distinction remains the same: Nvidia and other strategic investors may sit on the capital and infrastructure side of the ecosystem, while Anthropic's reported revenue is generated by enterprises and developers paying to use its products.

10

THE RULING

The Legal Overhang Has Eased

A separate, non-financial risk has just moved firmly in Anthropic's favor. On August 28, U.S. District Judge Rita Lin in San Francisco permanently barred the Trump administration from enforcing a set of restrictions that had cut Anthropic off from federal agencies and military contractors. The dispute traced back to February, when Anthropic refused to let Claude be used for mass surveillance or fully autonomous weapons; the administration responded by declaring Anthropic a "supply chain risk," a designation the judge said threatened to "cripple" the company by shutting it out of federal contracts.

Judge Lin's ruling was unambiguous. She found the administration's actions amounted to "unlawful retaliation against Anthropic for constitutionally protected expressive activities," and wrote that officials' "contemporaneous words and deeds confirm that the challenged actions were based on a desire to make a public example out of Anthropic for its 'arrogance' in criticizing the government, not based on any articulable basis to believe that Anthropic would actually sabotage its model." She also rejected the government's specific sabotage concerns as "entirely unfounded."

This is a reversal secured through litigation, not a change of position volunteered by the administration, and a related case in the D.C. Circuit is still pending, so the legal overhang isn't fully closed. But the ruling removes a genuine tail risk that had been hanging over Anthropic's federal and defense-related business, and it arrives alongside continued engagement between Anthropic and the Pentagon over its newer models, which the government has kept negotiating access to even while the dispute was live. Anthropic was the first AI company cleared to work on classified U.S. government systems back in 2024, and this ruling clears the way for that relationship to continue on Anthropic's terms rather than the administration's.

The SpaceX Relationship Is Warming

We've described the SpaceX compute arrangement above primarily as a financial commitment, \$1.25 billion a month through May 2029 for capacity at the Colossus data centres. It's worth noting the relationship has also shifted in a more constructive direction recently, which matters given SpaceX is both a supplier to Anthropic and, through its AI efforts, a competitor.

Elon Musk, who was openly critical of Anthropic as recently as early this year, publicly reversed that position in July, saying he had been "clearly wrong" about the company and pledging not to use SpaceX's compute leverage against a competitor: "I would never cut them off in a way that hurt them badly, even as a competitor. That's not my style." More recently, after SpaceX's acquisition of the coding tool Cursor prompted OpenAI to announce it would cut off Cursor's access to its models, citing distrust of Musk's companies, Anthropic did the opposite: it confirmed it would continue supplying Claude to Cursor and increase the compute behind it.

None of this changes the core financial relationship. Anthropic is still a paying customer of SpaceX's compute, and that arrangement still deserves the scrutiny any large related-party-adjacent contract deserves. The change in tone also reduces one potential risk: that the relationship could deteriorate and leave Anthropic scrambling for replacement capacity.

This Is Not the Dot-Com Bubble

The comparison to the late-1990s internet bubble is an understandable instinct, and there are real similarities: capital flooding into a new technology, high valuations, investors extrapolating enormous future markets, infrastructure spending accelerating, and no shortage of speculative activity around the edges of the sector.

But Anthropic's operating profile looks nothing like the typical pre-revenue internet company of 1999. It's already generating billions of dollars of quarterly revenue, it has a large and growing enterprise customer base, and its products are embedded in real business workflows, with a rate of revenue growth that's rare in software history. That doesn't make the eventual valuation attractive on its own. It does mean that the comparison becomes less convincing as Anthropic continues to generate substantial revenue from real enterprise customers.

At an earlier stage of the AI cycle, investors were largely valuing potential. Today, Anthropic generates enough revenue that investors can start asking conventional financial questions instead of narrative ones:

- ◆ What is the sustainable growth rate, once the current wave of enterprise adoption normalizes?
- ◆ What is the gross margin, and how does it move as inference gets cheaper?
- ◆ How much compute is required per dollar of revenue, and is that ratio improving?
- ◆ What is customer retention, and what is the lifetime value of an enterprise account?
- ◆ How much operating leverage emerges as the business scales?
- ◆ How much additional capital will the business need to raise to keep up with demand?

These are hard questions. But they're business questions, not technology questions, and the fact that Anthropic's numbers are now large enough to support asking them is itself a sign of how far the company has moved from pure narrative investing.

13

THE LISTING

The IPO, and the Number That Matters

Anthropic filed a confidential S-1 on June 1, 2026, and is reportedly targeting a public listing as early as September or October, with Morgan Stanley, Goldman Sachs and JPMorgan leading the underwriting. Early chatter pointed to a post-money valuation around \$965 billion. More recent reporting suggests the company may push for considerably more, potentially \$2 trillion, while raising as much as \$100 billion, which would make it larger on both counts than SpaceX's record-setting IPO earlier this year.

\$190–200B

2028 REVENUE TARGET ·
COMPANY PROJECTION, PER
REUTERS

\$65B

RUN-RATE, JULY
2026 · UP FROM
\$47B IN MAY
2026

~15×

IMPLIED
MULTIPLE ·
BASED ON THE
\$65B RUN-RATE

Anthropic is also expected to pitch a total addressable market north of \$30 trillion, edging out the \$28.5 trillion figure SpaceX used in its own roadshow. We'd treat that number the way we treat any TAM slide: a statement about the theoretical ceiling if the company captured every dollar of every market it considers addressable, not a

forecast, not revenue, and certainly not a valuation. For context, the entire technology sector of the S&P 1500, 191 companies, generated about \$2.4 trillion in combined revenue last year; a \$30 trillion TAM is more than twelve times that.

The number that will actually move the IPO price is the \$190 to \$200 billion 2028 revenue projection. The investment debate should move away from “is AI a huge market” and toward “how much of that market can Anthropic capture, and what economics will it earn when it does.” That’s a far more productive question, and a much harder one.

14

THE RISKS

What Could Go Wrong

None of the above should be read as an argument that Anthropic is risk-free. It isn’t, and several things are worth watching closely.

- ◆ **Growth deceleration.** A chunk of this year’s expansion reflects the first wave of enterprises adopting agentic coding tools. That wave won’t repeat itself every year, and growth will mechanically slow as the revenue base gets larger.
- ◆ **Capital intensity.** Compute commitments north of \$60 billion across three years only make sense if revenue keeps compounding. If growth stalls while those leases stay fixed, the operating leverage that just turned positive can flip the other way quickly.
- ◆ **Competition.** OpenAI, Google, Meta and others continue to invest heavily, and OpenAI’s recent pause on reinforcement-learning training and its focus on safety hardening could prove temporary rather than structural. Today’s market-share numbers aren’t guaranteed to hold.
- ◆ **Model commoditization.** If frontier model capabilities converge across providers, pricing power could erode faster than anyone currently expects.
- ◆ **Customer concentration.** With Claude Code reportedly responsible for roughly \$8 billion of the \$11.5 billion quarterly total, the growth story is more concentrated in one product line, and potentially a handful of large enterprise accounts, than the headline numbers suggest.
- ◆ **Regulation and safety.** Government policy could affect deployment in sensitive industries in ways that are difficult to forecast today.
- ◆ **Valuation.** Even an exceptional business can be a poor investment at an excessive price, and pricing an IPO at roughly 15 times a \$65 billion run-rate, or considerably higher if the \$2 trillion figure holds, leaves little room for a disappointing quarter.


These risks are real, and worth taking seriously. None of them argue that the business itself is fake. That's a separate claim, and it's the one behind the circular-financing narrative.

The Bottom Line

Anthropic can still be overvalued. That's a separate question from whether the circular-financing narrative holds up, and we don't think it does. The AI capital stack has real circular and interconnected elements. But Anthropic's revenue increasingly comes from enterprises and developers paying for AI capability they can put to economic use, at a growth rate few software companies have ever matched. The compute spending that unsettles investors is better understood as a consequence of that demand than as a substitute for it.

None of this makes the stock immune to a rough quarter, a slowdown in enterprise adoption, or a market deciding, once it's public, that a \$2 trillion valuation and a \$30 trillion TAM slide are too much to underwrite. Those are legitimate risks investors should weigh. The debate has therefore moved beyond whether Anthropic has a real business. It does. The more important questions now are how large the business can become, how durable the growth will be, and whether the eventual valuation adequately reflects those economics.

For investors, the more relevant question is ultimately who is paying for the intelligence the infrastructure produces. In Anthropic's case, the answer is increasingly a broad base of enterprises and developers. The recent developments around compute pricing, government access and SpaceX are broadly supportive of that view and suggest Anthropic may be better positioned than the headline capital-spending numbers imply.



Pradeep Datwani

MANAGING PARTNER, PANACEA VC

pradeep@panaceavc.com

panaceavc.com

This note is for general informational purposes only and does not constitute investment advice or a recommendation regarding any specific security. Past performance and current market conditions are not indicative of future results. Investors should conduct independent due diligence and consult a licensed financial advisor before making investment decisions.

Set in Playfair Display and Inter. The engine-turned rules and the rosette are hypotrochoids, drawn from the equation a geometric lathe cuts. Printed for private circulation.



PANACEA VC · DUBAI, UNITED ARAB EMIRATES
PANACEAVC.COM/RESEARCH