



WHAT WE SEE

The AI Buildout: From Demand Cycle to Capital Cycle

OCTOBER 2026

- Demand for AI compute is accelerating and shortages are acute, but the supply response has reached a scale that is beginning to test what demand can absorb.
- The top labs would need to roughly 6x today's annualized revenue to pay for planned compute by 2028. While possible, the range of outcomes is increasingly wide, and the risk of oversupply by 2029 is growing.
- Most capital cycles have ended when financing markets shut, and the weakest borrowers are beginning to show signs of strain.
- We expect that stocks will price the oversupply risk well before earnings show it.
- We believe more balanced and selective AI exposure is appropriate going forward.

Investors are at a crossroads today: a generational demand cycle for intelligence, tokens, and compute is now being matched by an equally unprecedented capital cycle.

On the one hand, we are in the early days of discovery for arguably the most transformative technological change of our lifetime. ChatGPT has grown from a small experiment to nearly 1 billion weekly active users in under four years. We estimate Anthropic's annualized run-rate revenue has grown from approximately \$9 billion in December 2025 to approximately \$70 billion in nine months, which our research indicates is almost triple the year-end 2026 forecast the company issued in December 2025. Meta's Muse reached 10 million U.S. downloads in 40 days, nearly triple ChatGPT's downloads in the same time.¹ Given the sharp inflection in compute demand driven by Claude Code and Codex in 2026, and more recently Muse, spot compute markets continue to tighten, driving near-term shortages in rental compute, labor, powered shells, memory, foundry, equipment, substrates, CPUs, and countless other end markets.

On the other hand, we are seeing a supply response that, relative to U.S. GDP, has now eclipsed every supply cycle since the late-1800s buildout of railroads. This supply response is highly rational for each participant, with paybacks for near-term capacity deployment across infrastructure markets often measured in months, or at worst a few years. The question is whether demand can grow into it. What is rational for each participant individually is becoming riskier for the system collectively.

We currently estimate the supply chain is on track to add at least 110 GW of U.S. critical IT capacity between 2025 and 2028. We expect just over half of this will be allocated to Anthropic and OpenAI; at prevailing compute rental rates of approximately \$13 billion per GW, the two labs would need to combine for at least \$780 billion of revenue by 2028 to bear the deployed compute at breakeven contribution margins (inclusive of compute spend for inference, reinforcement learning, training, and research/experimentation). Meta, Google, SpaceX, open source, and other markets would also need to absorb the remaining approximately 50 GW of energized capacity, which adds roughly \$650 billion of needed revenue and brings the total to about \$1.4 trillion including OpenAI and Anthropic. At around \$780 billion, Anthropic and OpenAI revenue alone would represent a combined 2-3% of U.S. GDP by 2028, and at \$1.4 trillion, total AI-linked revenue would approach 4-5% of U.S. GDP.

Unlike past semiconductor or industrial cycles, AI demand lacks a natural anchor measurable against the broader human population: approximately 1.25 billion smartphones, 270 million PCs, 95 million cars, and 13 million pre-AI servers a year all trace back, directly or indirectly, to the number of humans in the world.² Tokens, the labs' unit of demand, have no such anchor: they are hard to forecast even near term and have grown by multiples every few months, most sharply in 2026 as API revenue from coding agents has gone from a minor contributor to the primary economic driver for Anthropic and OpenAI. The


result is a supply chain building long- and medium-lived assets (data centers and compute, respectively) against lab revenue with little visibility beyond a few months, which creates a growing risk of oversupply in compute markets, especially as we approach 2029.

Compounding this risk, hyperscaler capex has begun to exceed operating cash flow, and industry participants are turning to more creative capital structures to finance the buildout, layering on- and off-balance-sheet financial leverage on top of operating leverage. The buildout now depends on access to capital markets. Our analysis of previous capital cycles suggests that most historical supercycles have ended when access to capital shut off, including British railways (1847 and 1866), U.S. canals (1839-42), U.S. railroads (1857, 1873, and 1893), U.S. utility holding companies (1929-32), Japanese real estate (1990), Asian manufacturing (1997), telecom (2000), U.S. housing and shipping (2007-08), U.S. shale (2015-16), and Chinese property (2020-21). The shutoff often began with a first crack 12-30 months before the system-wide break, but

sometimes only a few months before. Once the break came, the repercussions were swift. While we have little concern about hyperscalers' continued access to capital, other parts of the ecosystem—neoclouds, data center developers, lab-linked financing vehicles, among others—are increasingly brittle.

We believe the AI infrastructure supply response is driving a stark change in the investing regime that prevailed over the past four years. Until now, the rising tide of AI infrastructure spend has lifted all boats, and most assets levered to AI infrastructure have performed well. Supply chain forecasts imply accelerator spend approaching \$1 trillion in 2027. We estimate commodity memory prices have increased more than 10x over the past year. The price of a gas turbine has more than tripled since the beginning of the AI buildout. Neoclouds selling into spot markets are realizing sub-one-year paybacks on data center spend. So-called "bottleneck companies" have largely outperformed and rewarded investors handsomely.

***OpenAI and Anthropic Need Approximately 6x Today's Revenue
by 2028 to Pay for Planned Compute³***



Source: Sculptor, as of October 2026

* Adjusted for approximately \$20 billion of revenue share paid by OpenAI to Microsoft.

** Minimum additions from Google, Meta, SpaceX, open source, and others.

Capital markets have been digesting this rapidly changing picture, and the reaction has been sharp but uneven. The buildout is being financed and expressed across all asset classes we invest in—equities, convertibles, corporate credit, and asset based finance (ABF)—and each is pricing the same risks differently. For the better part of the last three years, we have invested across this theme in every part of our portfolio. As a result, our investment thesis is informed by the full spectrum of asset classes rather than any single market’s view, and this breadth allows us to pivot wherever relative value looks best for deploying capital.

Today, the range of outcomes is widening, and each market is underwriting a different base case, which is fueling dispersion that is both significant and increasing. Long-dated hyperscaler bonds have widened approximately 40 basis points since late 2025 while the broader index has been flat, and CDS on the largest neocloud trades over 1,000 basis points.⁴ Convertible bonds, which have become one of the predominant funding sources for neoclouds, have been weaker in the largest AI-linked issues, and new deals are clearing at the cheaper end of guidance with higher coupons and lower conversion premiums. Across credit markets broadly, financing remains available, but on less favorable terms for borrowers. That is expanding the opportunity to participate in highly structured transactions on lender-friendly terms, and we are sourcing those across our ABF and Corporate Credit platforms today, where joint underwriting draws on both areas of expertise. Equities linked to AI capex continue to be highly volatile, though performance is increasingly diverging across sub-verticals and sectors touching AI infrastructure.

From here, the environment calls for an evolved approach to investing in the cycle. Within equities, semiconductor and industrial stocks have a long history of being highly forward-looking, with multiples that capitalize cycle risks far ahead of fundamentals. Accordingly, as 2029 nears, we believe the growing risk of oversupply will push multiples lower despite what we expect to be a strong 2027 and 2028 across the supply chain. Multiples have already begun to derate—and we expect them to continue derating—for those parts of the AI buildout that to date have represented a large part of spend: accelerators, data center infrastructure, energy, and more recently, memory.

In practice, this means running more balanced net exposure and lower AI gross exposure. Shorts across the AI infrastructure space are of growing importance—especially in sub-verticals where the cycle is mature and estimates are full. At various points this year, we have been short names that we were long in years past, particularly the earliest winners of the AI demand cycle. On the long side, we have raised the bar further for any position levered to the AI buildout and are narrowly focused on finding names that meet four criteria: (1) technological differentiation/inflection, (2) significant estimate differentiation, (3) attractive risk-reward, margin of safety, and valuation on our numbers, and (4) an attractive point in the cycle. Earlier in the year, we believed memory and semicap equipment met these criteria; more recently, as current pre-trains make agentic orchestration feasible, we have seen that opportunity in CPUs, where persistent agents are rapidly shifting a growing share of workloads toward task coordination, tool calls, and sandboxed execution that run on general-purpose compute.

To be clear: we continue to expect AI progress across domains of knowledge work to accelerate, especially in the coming months. We expect this progress to be jagged and non-linear: coding agents have gone from unworkable to ubiquitous in the year since the release of Opus 4.5 in November 2025. We ourselves are experiencing a similar inflection in the day-to-day of financial analysis since the release of Opus 5.5 in September. Consumer agents such as Muse are quickly transforming how people complete day-to-day tasks. Our checks are surfacing similar transformations in law, medicine, drug discovery, robotics, and many other fields. We expect surprises to be the base case and plan to be intellectually flexible in our worldviews and frameworks.

In the past, acceleration in end demand drove upside to our estimates; now, acceleration to at least \$780 billion of revenue from the top two labs by 2028, and approximately \$1.4 trillion of AI-linked revenue overall, is a necessary condition to absorb planned compute deployments. Given the evolving risk-reward at this point in the cycle, we believe a more conservative risk posture is prudent.

Important Disclosures & Appendix

Important Information

This material is provided to you for informational purposes only. This is neither an offer to sell nor a solicitation of any offer to buy any securities in any fund or account managed by Sculptor Capital Management, Inc. or Sculptor Capital LP and its affiliates (collectively, “Sculptor Capital Management,” the “Firm,” “Sculptor,” or the “Company”), nor does it constitute a financial promotion, investment advice or an inducement or incitement to participate in any product, offering or investment. Any offer to purchase or buy securities or other financial instruments will only be made pursuant to an offering document and a subscription document, which will be furnished to qualified investors on a confidential basis at their request for their consideration in connection with any such offering. Any investment decision should be based on the information contained in the private placement memorandum. This material does not create any advisory relationship; such a relationship may only be established through a formal advisory contract.

The statements made herein reflect the subjective views and opinions of Sculptor. Such statements cannot be independently verified and are subject to change. There can be no assurance that the investment discussed herein will ultimately be successful or are representative of all investments made in a particular Sculptor strategy or product. The information contained herein may not be reproduced or used in whole or in part for any purpose, and there can be no assurance any investment strategy will perform as discussed herein or that any historical trends will persist in the future. The source of all information contained herein is Sculptor, unless otherwise noted. Past strategy and investment allocations are not necessarily indicative of future strategy or investment allocations.

While private investment funds offer investors the potential for attractive returns and diversification, they pose greater risks than more traditional investments. Investors' capital is at risk and investors may lose all or a substantial portion of their investment. Investors should consider the risks inherent with investing in private investment funds, which include, but are not limited to, leveraged and speculative investments, limited liquidity, higher fees and expenses and complex tax structures. The tax treatment of any investment will depend on the individual circumstances of each investor and may be subject to change in the future. Sculptor has in place policies and procedures designed to prevent market abuse and insider dealing. These policies and procedures are reviewed on a regular basis.

Some of the information contained herein has been obtained from third party sources. Sculptor has relied on the accuracy of such information and has not independently verified its accuracy. Sculptor makes no representation or warranty, express or implied, as to the accuracy or completeness of the information contained herein. Certain economic and market conditions contained herein have been obtained from published sources and/or prepared by third-parties and in certain cases has not been updated through the date hereof. All information contained herein is subject to revision and the information set forth herein does not purport to be complete. Past performance is not a reliable indicator of future results. A more detailed description of Sculptor's investment strategy, objectives and related risk will be made available upon request. Past performance of any individual investment or Sculptor fund is not indicative of their future performance.

1. Source: Bloomberg, Reuters, TechCrunch, views as of October 2026.
2. Source: International Data Corporation, as of June 8, 2026, and May 31, 2024, International Organization of Motor Vehicle Manufacturers, as of April 23, 2026, and TrendForce, as of March 9, 2026.
3. Source: Sculptor estimates, presented as of October 2026. Revenue is the annualized run-rate required on December 31, 2028, to cover the rental cost of deployed AI compute at breakeven contribution margins, assuming approximately \$13 billion per GW (Amazon Web Services' Q2 2026 implied monetization per incremental GW). September 2026 annualized run-rate revenue (“ARR”) combines OpenAI (approximately \$70 billion, including an approximately \$20 billion add-back of revenue share paid to Microsoft for comparability with Anthropic's gross disclosure) and Anthropic (approximately \$70 billion, our estimate). The 2028 OpenAI and Anthropic minimum assumes approximately 60 GW: each lab's known compute plus half of U.S. AI capacity not yet tied to a named buyer. Additions by others reflect the remaining approximately 50 GW. The approximately 110 GW total excludes traditional data center capacity and its estimated growth.
4. Source: Bloomberg, views as of October 2026.

Cautionary Note Regarding Forward-Looking Statements

Certain information contained in this document constitutes “forward-looking statements” that can be identified by the use of forward-looking terminology such as “may,” “will,” “should,” “expect,” “anticipate,” “target,” “project,” “estimate,” “intend,” “continue,” or “believe” or the negatives thereof or other variations thereon or comparable terminology. Due to various risks and uncertainties, actual events or results or the actual performance of any Sculptor Capital investment may differ materially from those reflected or contemplated in such forward-looking statements.

Risk Factors

The risks below relate to the views and investments discussed in this material. They are not a complete list of the risks of investing in any Sculptor fund or strategy. Before investing, review the “Risk Factors” in the Confidential Memorandum in its entirety. Investments involve risk, including the loss of all or a substantial portion of invested capital.

Views and estimates may be wrong. This material reflects Sculptor's opinions as of October 2026. Statements about future events, including expected demand, capacity additions, revenue requirements, credit conditions, and valuations, are forward-looking

and may differ materially from actual results. They rest on Sculptor's own estimates and assumptions, including estimates for private companies based on limited information, and on third-party data that Sculptor has not independently verified. Sculptor has no obligation to update this material.

AI infrastructure investments are concentrated and cyclical. Investments tied to the AI buildout share common drivers and may decline together. Key risks include:

- Demand for AI compute may grow more slowly than the supply being added, leading to oversupply, lower pricing, and weaker returns on capital.
- Demand is difficult to forecast and has no established long-term benchmark.
- Participants in the supply chain, including semiconductor, memory, equipment, data center, energy, and industrial companies, are subject to sharp swings in revenue, margins, and valuations.
- Valuations may change well before fundamentals do, and securities may not respond as Sculptor expects.
- Export controls, trade restrictions, and geopolitical developments may disrupt supply chains and demand.

Financing risk. Much of the buildout depends on continued access to equity, debt, and convertible markets. If financing becomes more expensive or unavailable, leveraged or less well-capitalized companies, such as neoclouds, data center developers, and lab-linked financing vehicles, may deteriorate quickly or default. Spreads, credit default swap levels, and convertible terms may move sharply and unpredictably.

Equity, short selling, and derivatives. Equity securities are volatile and may lose value rapidly. Short positions can lose an unlimited amount because the price of a security sold short may rise without limit, and they are subject to borrowing costs, recalls, margin calls, and short squeezes. Derivatives, including options and credit default swaps, may be illiquid, highly leveraged, and subject to counterparty, pricing, and regulatory risk.

Credit, convertible, and asset based finance investments. Debt and convertible investments are subject to credit, interest rate, and liquidity risk and may be difficult to sell at their perceived value. Asset based investments tied to data centers or compute hardware are exposed to technological obsolescence, declining residual and rental values, and customer concentration.

Non-U.S. issuers. Certain companies discussed, including memory manufacturers, are located outside the U.S. and are subject to currency, political, regulatory, tax, and market risks not present in U.S. investments.

Positioning, conflicts, and no recommendation. Sculptor and its funds currently hold (and may hold in the future) long and short positions, directly or through derivatives, in sectors and issuers discussed here. These positions may be consistent or inconsistent with the views expressed and may change without notice. Sculptor and its funds may benefit if prices move as Sculptor expects. References to sectors, issuers, or prior positioning are illustrative, are not recommendations, do not represent a complete or representative list of holdings, and do not indicate the performance of any Sculptor fund. Investors should not assume that any investment discussed was or will be profitable.