AI – Cash Fortresses or Cash Furnaces

written by Brendon Grunewald | August 22, 2025
The economy has rarely been more dependent on a handful of technology vendors than it is today on artificial intelligence. That concentration has delivered astonishing growth, but it also creates a brittle system in which one failure could ripple through earnings models, balance sheets, and livelihoods.

The right investor question isn't if AI delivers value; it's how much of that value is hostage to vendor solvency and the capital market appetite.

Lets start with the fortresses:

Alphabet (NASDAQ: GOOG) shows the same fortress profile: $95.1bn of cash, cash equivalents and marketable securities as of June 30, 2025. More telling is the asset mix, $203bn of net property and equipment, the majority "technical infrastructure" (servers, networking, data-center real assets). AI is no longer a slide in a strategy deck; it is welded into the left-hand side of the balance sheet, here again though its main revenue streams do not come directly from AI.

Amazon (NASDAQ: AMZN) makes the capex point in boldface. In Q2 2025 alone it purchased $32.2bn of property and equipment; on a trailing basis the bill is $108bn. Those numbers are the physical footprint of AI, power, land, racks, and fiber, financed and depreciated over years, not sprints. AWS revenue grew 18% to $30.9bn in the quarter; operating income hit $10.2bn. It has made huge strides in creating and integrating its own AI technology while offering it to its already global

client base.

Microsoft (NASDAQ: MSFT) closed the June quarter (Q4 FY2025) with revenue up 18% year-over-year, Azure growing 39%, and Microsoft Cloud revenue at $46.7bn, scale that throws off prodigious cash and cushions shock. Its liquidity is equally imposing: $94.6bn of cash and short-term investments at June 30, 2025. That wealth however, has not come from its AI offering.

Nvidia (NASDAQ: NVDA) bridges both worlds: profit machine and systemic node. As of April 27, 2025 it held $53.7bn in cash and marketable securities, the result of an unprecedented data-center cycle. Yet the first quarter also showed how quickly product cycles can bend financials: Nvidia recorded a $4.5bn charge tied to excess H20 inventory and purchase obligations. If the engine can misfire at the apex supplier, how confident should investors be about the smaller cogs?

Contrast those with the furnaces relying on capital market life support:

CoreWeave (NASDAQ: CRWV) has grown ferociously, but with heavy leverage and concentration risk. Its IPO prospectus and subsequent reporting highlight billions of debt raised in 2023–24; post-listing, insider selling after lock-up and investor jitters around losses and customer exposure have been notable. In August, the Financial Times reported more than $1bn of insider sales as the lock-up expired; earlier coverage flagged roughly $8bn of debt and revenue concentrated in a couple of customers. That is exactly the profile that can amplify shocks if financing dries up.

OpenAI (Private), has become the default model supplier for a

vast swath of enterprise experimentation, and much more than experimentation. The company is running at roughly $12bn in annualised revenue and claims some 700m weekly users of ChatGPT. Yet even at that scale, management has lifted its projected 2025 cash burn to about $8bn, underscoring how compute, talent and data drag on near-term free cash flow.

The dependency risk is not theoretical. OpenAI itself says “more than 92% of the Fortune 500” are building on its products. Were a model provider of that centrality to suffer a financing squeeze, a governance rupture or a prolonged outage, thousands of corporate workflows, and the revenue they touch, would be exposed. Investors should be asking vendors about step-in rights, multi-model routing and API-escrow equivalents with the same zeal they once interrogated disaster-recovery plans.

Even Sam Altman now concedes that AI markets look frothy. “Are we in a phase where investors as a whole are overexcited about AI? My opinion is yes,” he told The Verge, adding that bubbles form when “smart people get overexcited about a kernel of truth.” High burn and high valuations are a combustible mix; if revenue ramps disappoint, financing windows can slam shut quickly.

Anthropic (Private), another model front-runner, illustrates the burn dynamic. The company told investors it burned $5.6bn in 2024 and still expects to burn around $3bn this year and projecting a breakeven later in the decade. Whatever one thinks of that trajectory, it underscores how compute-heavy economics push even category leaders to the capital markets.

Are Insiders getting out?

At Nvidia, CEO Jensen Huang has executed multiple 10b5-1 plan sales this summer (routine for diversification at these

valuations), a reminder that management teams are actively monetising portions of their stakes even as AI demand booms. The sales, variously reported across Form 4 trackers and the financial press, are not red flags by themselves, but in a market priced for perfection they are signals institutional investors should at least incorporate into governance checklists.

OpenAI has also pursued repeated employee secondary sales, the latest discussions would allow staff to sell roughly $6bn of stock at a mooted $500bn valuation, helping retain talent but also reminding investors that liquidity events are occurring well before public-company disclosure disciplines apply.

Systemic risk

1. **API lock-in.** Thousands of firms have shipped AI-infused products built on a single proprietary model or an AI cloud provider. If that supplier fails or retrenches, those downstream products may degrade overnight, forcing hurried re-platforming, SLA penalties, and lost customers.
2. **Purchase-commitment risk**. Suppliers across the stack are taking on multi-year obligations for GPUs, power, and real estate; a downturn could strand working capital or trigger onerous take-or-pay clauses, as Nvidia's own H20 write-downs hint at upstream.
3. **Customer concentration**. CoreWeave's disclosures on reliance on a small number of customers show how easily a single renegotiation can torpedo forecasts and covenants.

What to do?

1. **Underwrite vendor concentration explicitly.** Ask portfolio

companies (and your own CIO) for a “model dependency map” showing which products, SLAs and revenues touch a single AI vendor. Price the re-platforming cost and time and make sure you have a desaster recovery plan that includes AI.
2. **Watch balance-sheets.** Check and put alerts on any AI companies that are critical to your business operations.
3. **Read the Form 4s.** Persistent plan-based selling at the chip and infra layer, large post-lock-up disposals, and employee secondaries at AI unicorns are all signals about insider risk-reward.
4. **Build multi-model resilience.** Where feasible, push teams to abstract the model layer (routers, guardrails, evals) so you can switch providers or blend outputs without a full rewrite.

What are the trading options?

Treat AI as a capital-structure trade with convexity. Overweight self-funded platforms compounding FCF and take targeted upside in high-burn or structured exposure. Consider shorts/puts in funding-dependent vendors. Use insider liquidity and secondary terms as timing signals. Let the cost of capital set your map. Pay for speed, not stories.