

Panam – August 2026 - Newsletter

THE AI CAPEX BOOM

THE CAPEX LEDGER

The AI capex cycle has stopped being a technology story and become a macroeconomic one. Combined 2026 capital spending guidance from the five largest hyperscalers has crossed \$800 billion, AI-related investment is now estimated at roughly 2% of U.S. GDP, and – for the first time since some of these companies went public – free cash flow is buckling under the weight of the buildout. What began three years ago as a race to train bigger models has evolved into a full-scale industrial build-out: gigawatt-scale data center campuses, custom silicon programs, multi-year power purchase agreements, and balance sheets increasingly leaning on debt and equity issuance rather than organic cash generation to keep pace. This issue walks through the GDP-level numbers, a company-by-company scorecard on Nvidia, OpenAI, Anthropic, Google/Gemini and Meta, and what the widening gap between operating cash flow and capital spending means for markets, valuations, and the broader economy over the year ahead.

~2.0% AI infra spend as share of 2026 U.S. GDP (TS Lombard)	\$805B+ 2026 capex guide, 5 hyperscalers (Morgan Stanley)	-91% Meta free cash flow decline, Q2 2026 YoY	\$(5.9)B Alphabet Q2 FCF – first negative quarter since 2004 IPO
---	---	---	--

1. The Macro Picture: AI's Share of GDP

Multiple independent estimates now converge on the same conclusion: AI infrastructure investment is one of the largest concentrated capital cycles in modern U.S. economic history. Analysts at TS Lombard have gone so far as to compare the current build-out to the Gilded Age's railway mania, arguing that on a spending-to-GDP basis, nothing since the transcontinental rail expansion of the 19th century has concentrated capital this quickly into a single category of physical infrastructure. What makes this cycle unusual is not just its size but its speed – the jump from a rounding error to nearly 2% of GDP has taken roughly three years, compressing into a few earnings cycles a scale of investment that previous infrastructure booms took a decade or more to reach.

Headline share – TS Lombard projects the U.S. will devote approximately 2% to 2.5% of GDP to AI and data-center infrastructure in 2026 – \$650–800 billion – putting the country far ahead of every other major economy. Norway and Saudi Arabia, the next-highest spenders, are each projected at only around 0.7% of GDP, while China's data-center spending sits near 0.4%, the Eurozone at roughly 0.2%, and Canada trailing the field at approximately 0.15%. The U.S. alone is expected to account for more than 80% of an estimated \$800 billion in global AI infrastructure spending this year.

GDP contribution — BEA data shows AI-adjacent investment — computer equipment, software, and intellectual-property R&D — contributed roughly 0.79 percentage points of the 1.50% real GDP growth rate recorded in the second quarter of 2026, meaning AI-related capital formation accounted for well over half of total quarterly growth. That follows a similarly outsized contribution in the first quarter, when computer-and-peripheral-equipment investment alone grew at a 67.4% annualized rate and software investment grew 22.6%, together contributing about as much to headline growth as the entire household consumption line.

Forward trajectory — Morgan Stanley has twice raised its hyperscaler capex forecast this year, most recently lifting the 2026 figure to about \$805 billion from a prior \$765 billion estimate, and pushing its 2027 projection up to roughly \$1.1 trillion from an earlier \$951 billion. Those upward revisions have become a recurring pattern across nearly every sell-side desk covering the sector: initial guidance is issued, hyperscalers beat it within two quarters, and estimates are revised upward again before the next earnings cycle even begins.

Concentration risk — The risk here is a mismatch between the size of the AI investment category and its contribution to growth — not the size of the category on its own. The table below lays out the comparison directly: computer equipment and software investment is a small slice of the economy by weight, but its growth rate in Q1 2026 was so far outside the norm that it generated almost the same percentage-point contribution to GDP growth as household consumption, a category roughly eighteen times larger.

Metric — Q1 2026	Computer Equip. + Software	Household Consumption
Weight in GDP (size of category)	~3.7%	~68%
Growth rate, annualized	67.4% (equipment) / 22.6% (software)	~1.6% (implied, typical pace)
Contribution to GDP growth	1.09 percentage points	1.08 percentage points
Share of total 2.0% headline GDP growth	~55%	~54%

Read across the rows and the mismatch becomes concrete: computer equipment and software investment carries only about 3.7% of the weight of the U.S. economy, versus roughly 68% for household consumption — consumption is roughly eighteen times larger. Yet because AI-driven equipment spending grew at 67.4% annualized and software spending grew at 22.6%, while consumption grew at its usual low-single-digit pace, the two categories ended up contributing almost identical percentage-point amounts to the quarter's 2.0% headline GDP growth — 1.09 points from the tiny, fast-growing category and 1.08 points from the enormous, slow-growing one. Together those two contributions account for roughly the entirety of the quarter's growth, which is the concrete sense in which Q1 2026 GDP growth was a story about a handful of hyperscalers' capital budgets as much as it was a story about the broader economy.

2. Company Scorecard: Capex & Cash Generation

The players in this cycle fall into two distinct camps, and the gap between them is widening every quarter. On one side sits the chip and infrastructure supplier pulling in record cash — Nvidia — whose free cash flow keeps climbing precisely because every other company in this newsletter is a customer. On the other side sit the buyers: Meta, Alphabet, Amazon, and Microsoft, whose free cash flow is compressing in real time as operating cash gets redirected into GPUs, data-center construction, and power infrastructure. Layered on top of both camps are the pure-play frontier labs, OpenAI and Anthropic, which are burning cash outright while betting that today's infrastructure spend converts into tomorrow's revenue at a scale neither has yet achieved. The scorecard below walks through each company's latest disclosed figures and what they signal about the sustainability of current spending levels.

NVIDIA — the supplier printing the cash everyone else is spending

Metric	Latest Figure
Q1 FY2027 revenue (Apr 2026)	\$81.6B, +85% YoY
Data center revenue	\$75.2B, +92% YoY
Q2 FY2027 guidance	\$91.0B ± 2%
Operating cash flow, Q1 FY27	\$50.3B (vs. \$27.4B YoY)
Capital returned to shareholders	~\$20B in buybacks/dividends, Q1 FY27 alone
FY2026 full-year revenue	\$215.9B, +65% YoY

Nvidia remains the one name in this ecosystem generating dramatically more cash than it spends, and the gap is widening with each quarter of results. Every dollar that hyperscalers, cloud providers, and AI labs commit to infrastructure flows in large part back to Nvidia's GPUs, which is precisely why the company's own operating cash flow nearly doubled year over year even as its customers' cash flows compressed. CEO Jensen Huang has described the current moment as the arrival of "agentic AI" doing productive, revenue-generating work at scale, and has been explicit that supply — not demand — is now the binding constraint, telling investors the company will be short of capacity even with record shipments. That dynamic gives Nvidia unusual pricing power and visibility relative to everyone else on this list, and explains why it is simultaneously funding customers, investing in suppliers, and returning record capital to its own shareholders.

Meta — capex consuming nearly all operating cash flow

Metric	Latest Figure
Q2 2026 revenue	\$60.8B, ad revenue +27% YoY
Q2 2026 capex	\$31.1B (~98% of operating cash flow)

Q2 2026 free cash flow	\$784M, down 91% YoY (from \$8.55B)
FY2026 capex guidance	\$130B–\$145B (narrowed upward)
Funding gap response	~\$25B in new long-term debt issued in Q2 alone

- Meta's core advertising business is not the problem — ad revenue grew 27% year over year with both impressions and price-per-ad rising, and total revenue beat the analyst consensus range comfortably. The free-cash-flow collapse is instead a deliberate capital-allocation choice: quarterly capex of \$31.1 billion consumed roughly 98% of the \$31.9 billion in operating cash flow the business generated, leaving almost nothing left over once dividends were paid out. Because the company continued its dividend but skipped share repurchases entirely during the quarter, and because it raised approximately \$25 billion in new long-term debt to help fund construction, the capital-return program shareholders have come to expect is now effectively being financed by borrowing rather than by the underlying advertising business a dynamic management has signaled will persist for as long as the \$130–145 billion annual capex program stays in place.

Alphabet (Google / Gemini) — first-ever negative FCF quarter

Metric	Latest Figure
Q2 2026 capex	\$44.9B
Q2 2026 free cash flow	≈ −\$5.9B (first negative quarter since 2004 IPO)
Google Cloud growth	63% YoY (up from 32% YoY in Q2 2025)
Balance-sheet response	\$20B debt issued + \$84.75B multi-phase equity raise announced

- Gemini's cloud momentum is real — Google Cloud growth nearly doubled year over year, accelerating from 32% in the same quarter a year ago to 63% in Q2 2026 — but capital spending is still outpacing it by a wide margin. Alphabet is the clearest case in this scorecard of a hyperscaler choosing to pre-fund a multi-year AI infrastructure bet with new debt and equity rather than relying on existing cash generation: on top of the \$20 billion in debt already issued, the company announced a multi-phase equity raise expected to total \$84.75 billion, with an initial tranche of roughly \$44.4 billion in net proceeds landing on the balance sheet before the remaining \$40 billion is executed in stages beginning in the third quarter. Combined with its existing cash reserves, that financing gives Alphabet one of the largest liquidity cushions in the group, but it also marks a symbolic turning point for a company that had never before reported a negative free-cash-flow quarter in over two decades as a public company.

OpenAI (ChatGPT) — the steepest burn, funded by serial mega-raises

Metric	Latest Figure
2026 projected loss (internal)	~\$14B; adjusted loss ~\$8B
2026 cash burn estimate	\$17B+ (some analysts project up to \$27B)
Monthly revenue (confirmed)	~\$2B, ~\$25B annualized run-rate
Infrastructure commitments	Reduced from \$1.4T toward ~\$600B total compute spend target
Latest raise	\$122B round at \$852B valuation
Path to FCF-positive (internal target)	Not projected before 2029

- OpenAI's burn rate is structurally the highest of the group relative to revenue – internal planning documents reportedly show cash burn holding near 57% of revenue through both 2026 and 2027, a ratio that barely improves even as the top line grows, versus a rapidly declining burn ratio at Anthropic over the same period. The company's competitive position has also softened on two fronts simultaneously: enterprise market share has slipped from roughly 27% toward Anthropic's rising 40%, according to Menlo Ventures survey data, while ChatGPT's web-traffic share fell from 86.7% in January 2025 to 64.5% in January 2026 as Google Gemini's share climbed from 5.7% to 21.5% over the same stretch, a shift accelerated by Apple's decision to integrate Gemini into Apple Intelligence.
- In response, OpenAI declared what CEO Sam Altman termed a "Code Red" in December 2025, deprioritizing advertising and shopping-agent initiatives to redirect resources toward core model quality, while simultaneously scaling back its total compute-spend target from the original \$1.4 trillion figure toward roughly \$600 billion as investor scrutiny of the gap between commitments and confirmed capital has intensified.

Anthropic – smaller compute base, faster path toward breakeven

- A quick note on how to read the figures below: "annualized run-rate" (ARR) means the most recent month's revenue multiplied by 12 – it is a snapshot of current pace, not actual revenue booked over a full year, so the figures below track how that pace has accelerated at several points through 2026 rather than one single growth line. The training-cost and burn-ratio figures are separate metrics: one compares the cost of building a model, the other compares the cost of running the company relative to what it earns.

Metric	Latest Figure
ARR – end of 2025	~\$9B
ARR – April 2026	~\$30B (passed OpenAI's run-rate for the first time)
ARR – May 29, 2026	~\$47B (per Series H funding announcement)
ARR – July 2026	~\$69B (most recent disclosed figure)

Revenue mix	80%+ enterprise/API subscriptions vs. OpenAI's consumer-heavy base
Model training cost vs. OpenAI	~4x cheaper per comparable frontier model (reported)
Cash burn as % of revenue (internal)	~33% in 2026, falling to ~9% by 2027
Projected breakeven	~2028, per WSJ reporting on leaked internal financials

- The headline number to take away is the July 2026 figure: Anthropic's annualized run-rate has gone from about \$9 billion to about \$69 billion in roughly seven months, a nearly eight-fold increase driven almost entirely by enterprise and API demand — particularly coding tools built on Claude — rather than by consumer subscriptions. That growth is what allowed Anthropic to overtake OpenAI's own run-rate in April 2026 and has kept it ahead since.
- On costs, two separate comparisons are worth keeping distinct. First, training cost: building a comparable frontier model is reported to run roughly four times cheaper at Anthropic than at OpenAI, a gap the company attributes to more disciplined compute allocation during training rather than to a single technical breakthrough. Second, operating burn: Anthropic's cash burn — how much cash the company consumes running the business relative to how much revenue it brings in — is forecast internally to fall from about a third of revenue in 2026 to roughly 9% by 2027. That declining burn ratio is the reverse of OpenAI's trajectory, where burn is projected to stay near 57% of revenue through the same period, and it is the main reason WSJ reporting on leaked internal financials points to Anthropic reaching breakeven around 2028, roughly a year or more ahead of OpenAI's own internal target.

The Free Cash Flow Squeeze — Why It's Happening Now

- Free cash flow compression across the hyperscaler group is not a surprise to anyone who has been reading capex guidance over the past year, but the speed at which it has arrived caught even sell-side analysts off guard. Combined 2026 capex guidance across Amazon, Alphabet, Microsoft, Meta and Oracle now exceeds \$725–805 billion, a 77% increase over 2025 spending levels, and the vast majority of that incremental spend is going toward data centers, custom AI accelerators, networking equipment, and the power infrastructure needed to run all of it. The mechanics below explain how that guidance is translating, quarter by quarter, into the cash-flow statements of five companies that not long ago were among the most reliable cash generators in the entire stock market.
- Amazon joins the list — Amazon's trailing-twelve-month free cash flow turned negative at approximately –\$7.6 billion, down sharply from positive \$18.2 billion a year earlier, making it the second hyperscaler after Meta to cross into negative territory on this measure. The shift was driven by a Q2 2026 capex figure of \$53.1 billion in a single quarter, and management used the same earnings call to raise full-year 2026 capex guidance from \$200 billion to \$220 billion, signaling that the spending trajectory has not yet peaked even as free cash flow keeps deteriorating.
- Oracle's extreme case — Oracle's situation is the most extreme in the group: fiscal 2026 capital expenditure reached 174% of operating cash flow, meaning the company spent nearly twice as much building out infrastructure as its underlying business generated in cash over the same period. That

SEBI PMS Reg No:INP000003898

imbalance is large enough that Oracle has announced plans to raise between \$45 billion and \$50 billion through a combination of debt and equity offerings simply to keep its infrastructure commitments funded through the rest of the buildout cycle.

- Microsoft, still covered — Microsoft remains the one exception among the major spenders, at least for now: operating cash flow of \$46.7 billion in the most recent quarter still covered capital expenditure of \$30.9 billion by more than 1.5 times, even though capex itself grew 84.4% year over year against only 26% growth in operating cash flow. That gap is narrowing quickly quarter over quarter, and with Q4 guidance already flagged above \$40 billion for a single quarter, Microsoft's cushion looks set to compress further before it stabilizes.
- Financing shift — Debt and equity issuance is quickly becoming the norm rather than the exception across the group, replacing organic cash generation as the primary funding source for new capacity. Alphabet has combined \$20 billion in new debt with an \$84.75 billion multi-phase equity program, Meta issued roughly \$25 billion in new long-term debt in a single quarter, and Oracle is preparing to raise \$45–50 billion through similar channels — a collective pivot toward external financing that would have been unthinkable for this group of companies just two or three years ago.

4. What to Watch Next

- Whether Microsoft's operating-cash-flow-to-capex cushion continues to narrow as Q4 FY26 capex guidance, already flagged above \$40 billion for a single quarter, comes through in reported results. Microsoft has been the sector's cash-flow bellwether precisely because it is the last major hyperscaler still comfortably covering capex out of operations, and a further deterioration in that ratio would remove one of the few remaining reassuring data points for investors watching the broader group.
- OpenAI's targeted IPO timeline, reportedly aimed at September 2026, and whether public markets are willing to underwrite a valuation near \$1 trillion against an adjusted annual loss of roughly \$8 billion and a burn rate that internal projections show persisting near 57% of revenue through 2027. The IPO decision will be one of the clearest tests yet of whether public investors share the confidence that private markets have shown in OpenAI's path to profitability.
- Anthropic's declining burn ratio, projected to fall toward roughly 9% of revenue by 2027 from about a third of revenue in 2026, holds up as the company scales infrastructure spending to match a run-rate that has already climbed to approximately \$69 billion. Any slippage in that trajectory would narrow the financial-discipline gap that has become central to Anthropic's competitive positioning against OpenAI.
- Further hyperscaler debt and equity issuance as 2027 capex guidance, already trending toward a combined \$1.1 trillion across the five largest spenders, gets finalized during fourth-quarter earnings calls. Watch in particular for whether Meta or Amazon follow Alphabet's lead with a large-scale equity raise, which would mark a further step away from capex funded purely out of operating cash flow.

Sources: TS Lombard / Lombard Odier, U.S. Bureau of Economic Analysis, Morgan Stanley research, Nvidia investor relations, Meta/Alphabet/Amazon/Microsoft Q2 2026 earnings releases and calls, The Information, Wall Street Journal, Menlo Ventures enterprise AI survey, company disclosures. Figures reflect the most

SEBI PMS Reg No:INP000003898

recent quarterly and full-year guidance available as of early August 2026 and are subject to revision at next earnings.