CORNERSTONE
INVESTMENT PARTNERS, LLC

From Matterhorn to Meseta: Why AI Bottlenecks Extend the Cycle

October 2026

Executive Summary

- **Data Center Fear Is Creating Opportunity in Equity Prices:** Data center opposition has become a rare point of agreement between left and right heading into the midterms. Combined with investors' natural aversion to losses, that fear has many reading recent data as evidence of "peak AI."
- **Matterhorn or Meseta:** We believe supply bottlenecks across the AI value chain push demand into future years. The earnings path looks less like the Matterhorn's sharp summit and more like Spain's Meseta Central, a vast plateau, here with an upward grade.
- **Constraint, Not Exhaustion:** From power generation and chip clean rooms to memory, optics and lithography, suppliers report demand exceeding supply, with many expecting constraints to last into 2028.
- **Capital Is Available:** Treasury yields rose between July 31 and October 22, most at shorter maturities, but hyperscaler borrowing costs rose less. Spreads held or tightened for most issuers.
- **Our Position, Held with Humility:** Knowing what we know today, end demand for AI appears to be outstripping the industry's ability to supply it. We will change our minds if the evidence changes.
- **Two Ways to Win:** Valuations that discount peak earnings would rerate higher if those earnings prove to be a new normal. That creates the potential for returns from both earnings growth and multiple expansion.
- **The Caveat: AI Must Prove Value to Its Users:** The thesis holds only if AI earns a return for the companies using it. Early evidence points to revenue growth from faster product development, better-performing advertising and drug discovery, not just cost cutting.

The Price of Fear

We believe the market is pricing the AI buildout as though the industry is hitting its peak, but our analysis shows the supply chain is pushing present demand into future quarters and years, effectively flattening the peak into a plateau. At Cornerstone, our philosophy centers on identifying mispricings created by behavioral errors and emotions in the market. Today, fear is creating an opportunity in AI infrastructure stocks.

Part of the fear is political. Opposition to data centers has become a "horseshoe" issue, one that unites the populist ends of both parties. On the left, Senator Bernie Sanders has introduced a bill to halt new data center construction (S. 4214). On the right, Senator Josh Hawley has proposed shielding residential ratepayers from data center power costs (GRID Act) and ending Opportunity Zone tax breaks for data centers (Hawley).

In Wisconsin, Republican gubernatorial nominee Tom Tiffany attacks his opponent as "data center David Crowley," echoing concerns raised by democratic socialist Francesca Hong, though their proposed remedies differ (Axios).

The polling explains why candidates are listening. A Data for Progress survey in June found 63% of likely voters support pausing AI data center construction for at least a year, including 58% of Republicans. New York has imposed the nation's first statewide moratorium on new hyperscale data centers, a temporary pause on state permits for facilities of 50 megawatts or more (Executive Order 62).

The pressure now reaches the technology itself, not just the buildings. CNBC's All-America Economic Survey, taken October 14 to 18, asked registered voters whether AI companies should pause *advanced* development until Congress passes regulations, or move forward to keep up with China. Voters chose a pause, 52% to 32% (CNBC).

On advanced development, the horseshoe breaks along party lines. Democrats favor a pause 70% to 18%, while Republicans favor moving forward 50% to 33%, and MAGA supporters by nearly 55% to 27% (CNBC). Yet in the same survey, 63% of voters opposed an AI data center within 10 miles of their home, and even MAGA supporters opposed one, 40% to 29%.

Voters on the right want the technology but not the buildings. In our view, that combination restrains supply more than demand, which creates a spending trend that looks more like a plateau than a peak.

Part of the fear is human nature. Prospect theory, developed by Daniel Kahneman and Amos Tversky and central to Kahneman's Nobel Prize, holds that losses weigh more heavily than equivalent gains. When the story shifts from opportunity to risk, investors seek safety first and sort out the details later.

Part of the fear is geopolitical. The Iran war, which began with U.S. and Israeli strikes on February 28, resumed in July when the spring ceasefire collapsed. The risk rose again this week. Brent crude traded near $115 a barrel on October 23, a five-month high, after Houthi attacks on Saudi Arabia and reports that the U.S. and Israel had planned renewed strikes for after the midterms (Telegraph). Sentiment has turned quickly: CNN's Fear & Greed Index fell into "extreme fear" on October 22 (Benzinga, 24/7 Wall St). The VIX rose from 15 to 20 over the prior two weeks. The war is not the only cause. The Nasdaq's worst day since June also followed Meta's results and renewed worries about AI spending.

The result is a market reading strong results as evidence of "peak AI." ASML beat third-quarter earnings estimates by roughly 19%, and analysts raised current-year estimates by about 9%. The shares still fell

about 4% the next day and 5% over the following week. Lam Research said its December-quarter guidance would have been higher but for supply constraints, and that customers have visibility through 2028. Its shares fell about 12% the next day, roughly twice the move options had priced in. The market is treating a slower growth rate as the top of the mountain.

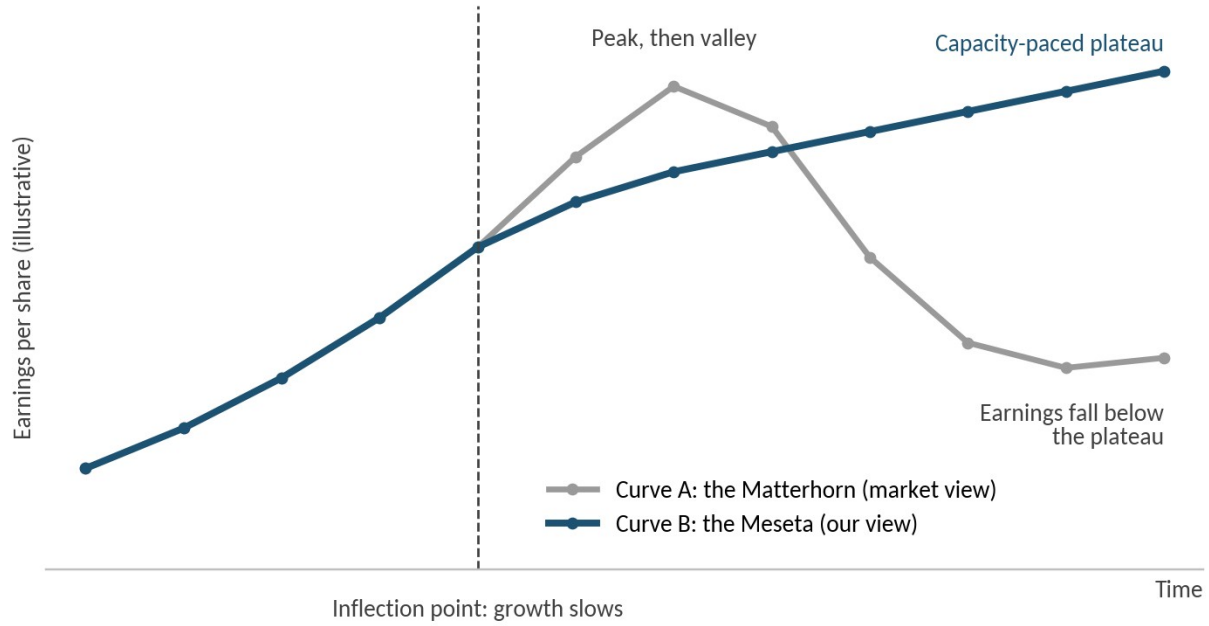
The Matterhorn and the Meseta

We believe supply bottlenecks change the shape of the AI earnings curve, not merely its timing. Two landforms make the point.

The Matterhorn is a near-symmetric pyramidal peak that rises to 4,478 meters (14,692 feet). Glaciers carved it from four sides into a sharp summit. Climbers reach the top and immediately start down.

Spain's Meseta Central is the opposite shape. The great interior plateau of the Iberian Peninsula, it stretches over 81,000 square miles, nearly the size of Kansas, at an average elevation of about 660 meters (2,165 feet). It is the oldest geologic formation on the peninsula, and Madrid sits at its center.

The market sees a Matterhorn. We see a Meseta: lower than the summit, but holding its height across a vast expanse. What keeps AI earnings at altitude is the backlog, the multi-year contracts and the down payments that hold demand in place while supply catches up.



Illustrative index values over 12 periods. After the inflection point, curve B earns 33% more in total than curve A.

Exhibit 1: Illustrative earnings paths. Supply limits turn a peak into a plateau.

Both curves rise together until the inflection point, when growth slows. In curve A, growth slows because demand is peaking. The industry overbuilds to meet the boom, supply overshoots, and earnings collapse into a glut. An investor who sees curve A rationally pays a low multiple for earnings expected to disappear.

In curve B, growth slows because supply cannot keep up. Demand that cannot be served this year does not vanish; it waits in backlog for next year's capacity. The industry does not hit significant oversupply before achieving equilibrium, and given the more inelastic end-customer AI demand curve, higher pricing and

margins remain throughout the ecosystem longer. The curve flattens into a plateau whose upward grade is set by how fast the industry can add capacity, plus pricing.

At the inflection point, the two curves are nearly impossible to tell apart by growth rate alone. The difference lies in why growth is slowing, and that is observable in backlogs, lead times, down payments and pricing. Under curve A, those roll over first. Under curve B, they keep rising even as revenue growth slows.

Three mechanisms push the peak into future years:

1. **Deferral.** Revenue that cannot ship this quarter ships next quarter. Lam Research said December shipments delayed by supply will ship in the March quarter, and Vertiv moved second-quarter revenue into the second half for the same reason.
2. **Prevention of overbuild.** Cyclical busts happen when supply overshoots demand. Binding constraints in power, clean rooms, lithography and memory make an overshoot hard to engineer over the next two to three years.
3. **Staging.** Entegris describes its revenue arriving roughly 12, 18 and 24 months after construction of a new fab begins, and only then with wafer volumes. Each layer of the chain peaks at a different time, which stretches a spike into a plateau.

Source detail for each company comment is in the Evidence Appendix at the end of this paper.

Where Supply Binds

Across the AI value chain, supply binds at nearly every link that depends on a physical asset with a long lead time. We define the chain from end demand backward: applications and the people who build them, the model providers, the platforms that host the models, the systems and components inside those platforms, the design of the chips, and the packaging, foundries, equipment and materials that make them. Under all of it sits a physical layer of power, buildings and cooling.

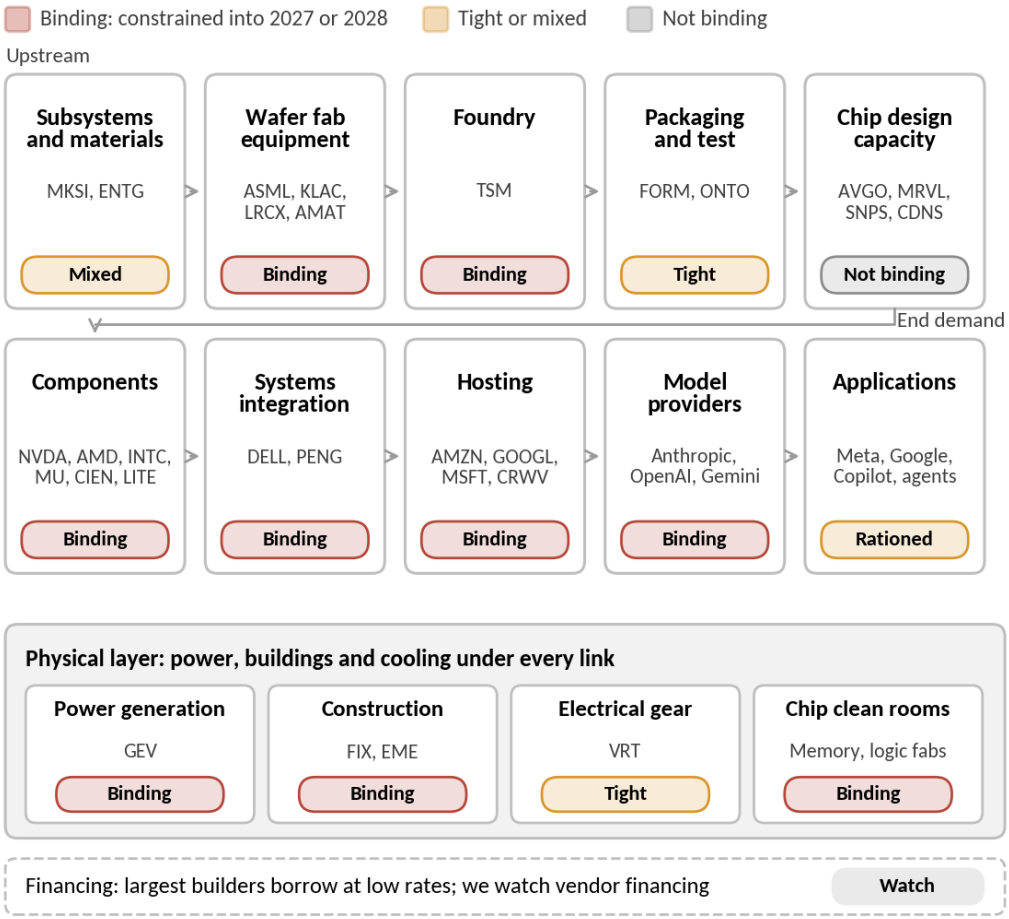


Exhibit 2: The AI value chain and where supply binds. Status from company commentary, July to October 2026.

A few recent data points:

- **Compute.** NVIDIA expects about 70% revenue growth in fiscal 2028, a supply-constrained outlook against customer forecasts for growth to double, and says its "entire supply chain is challenged."
- **Memory.** Micron has "no line of sight" to when supply and demand balance; its new fabs begin output in stages from mid-2027, with more following in late 2028 and beyond, and each takes time to ramp.
- **Lithography.** ASML says its planned 2027 and 2028 output will not fully meet customer demand.
- **Optics.** Ciena's 2027 revenue floor would be higher "if we had more supply," and it does not expect balance before 2028. In March, NVIDIA invested $2 billion each in indium phosphide (InP) laser makers Lumentum and Coherent, alongside multibillion-dollar purchase commitments and rights to future capacity, locking up scarce supply (Sharecast, via Hargreaves Lansdown). On October 29, Bloomberg reported that Marvell is in advanced talks to acquire Lumentum (Benzinga).
- **Power.** GE Vernova expects most of its 2030 gas turbine production slots to be committed this year, through orders and slot reservations.
- **Systems.** Dell says the AI supply chain is running at full stretch, with memory the tightest input.

- **Models.** Anthropic's IPO filing says demand for AI will outstrip supply for the foreseeable future, "limited principally by the availability of compute." It adds that it has "significantly underestimated" its compute needs before, leaving it capacity constrained (Anthropic Form S-1).
- **Data centers.** Of about 16 gigawatts of capacity slated for 2026 across roughly 140 projects, only about 5 gigawatts was under construction in February; about 11 gigawatts showed no visible progress (Sightline Climate, via Contrary Research).

Politics now adds a bottleneck of its own: Data Center Watch estimates that local opposition blocked or delayed at least 75 U.S. data center projects worth about $130 billion in the first quarter of 2026. In our framework, that pushes the plateau further out rather than lowering it, provided the demand waits.

The exceptions are telling. Chip design software, enterprise networking and most equipment subsystems are not constrained. Custom chip designers such as Broadcom and Marvell are limited by wafers, substrates and memory upstream, not by their own design capacity. The constraints concentrate in power, clean rooms, lithography, memory wafers, laser materials and chip substrates, where new capacity takes years rather than quarters.

The slowest link sets the pace for the rest. KLA's outlook assumes its peers can deliver, and Broadcom sizes its outlook to wafer, substrate and memory supply. Even links that are not constrained themselves ride the Meseta. The full source list, link by link, is in Section 9 of the Evidence Appendix.

Capital: Available, and Priced Near Sovereign

The buildout is increasingly funded with debt. Alphabet's borrowings rose from about $28 billion to about $100 billion in the year to June 30, 2026, and its free cash flow was negative $5.9 billion in its second quarter. Meta ended its third quarter with $83.7 billion of debt against $82.4 billion of cash and marketable securities, after negative free cash flow of $2.1 billion. Microsoft expects to remain free cash flow positive in fiscal 2027.

In response, some in the market have started to worry about further availability of debt to fund future spending. While it is true that investors are likely incrementally thoughtful about future risk, demand remains strong. The largest builders have access to capital at low rates. While some market observers have decried the perceived increased risk, credit markets continue to lend to the hyperscalers at modest premiums to U.S. Treasuries, while charging more to the most stretched borrowers.

Balance-sheet capacity reinforces this. Microsoft, Alphabet, Amazon and Meta are all rated AA- or better by S&P (Microsoft is AAA), none carries meaningful net debt, and consensus expects their combined EBITDA to rise from about $670 billion to $1.0 trillion in 2027 and $1.26 trillion in 2028 (Pinegap consensus, October 23, 2026). At a conservative 1x net debt to EBITDA, that is more than $1 trillion of borrowing room on 2027 estimates and about $1.4 trillion on 2028. Counting leases signed but not yet started narrows that cushion, and at Meta takes adjusted leverage to roughly 2x.

	Microsoft	Alphabet	Amazon	Meta	Total
Moody's / S&P rating	Aaa / AAA	Aa2 / AA+	A1 / AA	Aa3 / AA-	
Net cash or net debt, latest quarter ($B)	Net cash 20	Net cash 144	Net debt 6	Net debt 1	Net cash 157
Net debt / 2027E EBITDA	Net cash	Net cash	0.0x	0.0x	
EBITDA, last 12 months ($B)	199	175	170	122	666

	Microsoft	Alphabet	Amazon	Meta	Total
EBITDA, 2027E ($B)	238	310	274	179	1,001
EBITDA, 2028E ($B)	290	397	349	227	1,263
Borrowing room to 1x 2027E EBITDA ($B)	258	454	268	178	1,158
Borrowing room to 1x 2028E EBITDA ($B)	310	541	344	226	1,420
Lease-adjusted net debt / 2027E EBITDA	1.7x	0.0x	0.9x	2.3x	
Lease-adjusted net debt / 2028E EBITDA	1.4x	0.0x	0.7x	1.8x	

> ***Note:*** *Ratings from Moody's and S&P. Net debt is borrowings excluding leases, less cash and marketable securities, at June 30, 2026 (Meta at September 30). Net cash means cash and marketable securities exceed borrowings. EBITDA is the last 12 months reported and Pinegap consensus as of October 23, 2026; Microsoft's fiscal year ends in June. Lease-adjusted net debt adds lease liabilities on the balance sheet, leases signed but not yet started (undiscounted, so overstated) and guarantees: Meta's Hyperion residual value guarantees and Alphabet's guarantees of other companies' data-center leases (FT, via AI Weekly). Alphabet's cash and marketable securities roughly doubled this year to $242 billion; borrowing room still exceeds $1 trillion on 2027 estimates without its net cash.*

We compared the yields on the four hyperscalers' bonds with Treasury yields on July 31 and October 22, 2026 (Bloomberg). Treasury par yields rose about 37 basis points at two years, 21 at ten years and 2 at 30 years. Hyperscaler borrowing costs rose by less, so for most of their bonds the spread over Treasuries held steady or narrowed.

Microsoft's spreads tightened across its curve, and Amazon's held steady or tightened. Meta is the exception: its spreads held at short maturities but widened modestly from ten years out. Meta still pays roughly three to four times Microsoft's spread on its 30-year bonds.

The message is differentiation, not closure. Lenders are charging more where cash flow is most stretched and where maturities outlast today's visibility. That is credit discipline, and it is a long way from a market unwilling to fund the buildout.

Free cash flow has taken the hit. Combined free cash flow at Microsoft, Alphabet, Amazon and Meta fell from about $60 billion in the fourth quarter of 2025 to about $7 billion in the second quarter of 2026. Over the same period, their quarterly capital spending rose to about $165 billion (company reports, via Pinegap).

We expect this, and we have seen it before. Google paid $1.65 billion for YouTube in 2006 and did not disclose its revenue for 13 years (Billboard). YouTube ads brought in $15.1 billion in 2019, and about $21 billion in the first half of 2026 alone (Motley Fool).

Meta launched Reels in 2020 to widespread skepticism, and for years it earned less per minute than Facebook's and Instagram's feeds. By mid-2023, Reels had a $10 billion annual revenue run rate, and by that fall it was net neutral to Meta's overall ad revenue (Reuters, Marketing Dive).

Amazon, founded in 1994, posted its first annual profit in 2003 and reported net losses as late as 2014 (InfoWorld). The years of investment built the businesses, Amazon Web Services among them, that carry its profits today.

Investing ahead of profits is how these companies built their current earnings power. The scale is larger this time, so we will keep evaluating whether the spending earns its return.

Bond spreads capture only part of the bill. Much of the buildout is financed through leases on data centers not yet delivered, which are not recorded as debt. At mid-2026 these uncommenced leases totaled about $329 billion at Microsoft, $279 billion at Meta, $137 billion at Amazon and $85 billion at Alphabet (Reuters). Amazon's figure also covers warehouses, offices and aircraft, and Meta signed about $68 billion more in July. Meta also owns 20% of the unconsolidated Hyperion venture, whose $27 billion of debt is backed by about $28 billion of Meta residual value guarantees (Bisnow, Meta Form 10-Q). These commitments help explain why lenders charge Meta more at the long end, and they draw on the same private lenders as the structured financing described next.

The part of the capital stack we watch most closely is vendor and structured financing. On October 27, Broadcom announced a $63 billion vehicle with Bank of America, Blackstone, Citi and Morgan Stanley that will buy its AI chips and networking gear and lease them to Anthropic from 2027, with Broadcom backing the senior portion, reportedly about $43.5 billion (Bank of America). Broadcom's chief executive, Hock Tan, said in October that the cost of capital has overtaken chip and memory shortages as the main bottleneck for AI infrastructure. Arrangements like these tie a supplier's revenue to its customer's access to credit, which is where a capital constraint would appear first.

Our Position

Knowing what we know today, we believe end demand for AI is outstripping the industry's ability to supply it. This is a view held at a point in time, not a forecast we will defend at any cost. As we wrote in our May 2026 thought paper, *Beyond the Four Ps*, the ability to say "we may be wrong" is itself a risk-management tool.

Three kinds of evidence support the view, and each is harder to manufacture than an optimistic adjective on an earnings call:

- **Backlogs are rising.** Amazon Web Services reported a $496 billion backlog growing at triple digits, Google Cloud $514 billion, and Microsoft $678 billion of contracted revenue not yet recognized, up 84%. Excluding OpenAI, Microsoft's backlog grew 25%, a reminder that some of this demand is concentrated.
- **Commitments are lengthening.** Micron has committed more than 75% of its 2027 output, with customer agreements reaching 2031. ASML customers are asking to place orders for 2029 and 2030, which ASML has not yet accepted; the orders it does accept require substantial down payments. Anthropic's IPO filing shows at least $265.6 billion of fixed or minimum compute commitments at June 30, plus about $161.2 billion of Broadcom equipment leases that neither side can cancel except on default, with payments running as late as 2045.
- **Prices are rising at the bottlenecks.** ASML and Ciena have raised prices, and CoreWeave expects contribution margins on recently signed contracts to be 5 to 10 points higher than on earlier ones. Neocloud Nebius raised on-demand rental rates on NVIDIA GPUs by 17% to 21% from October 1, taking an H100 from $3.85 to $4.50 an hour and a B300 from $7.85 to $9.50 (Nebius price list, via FourWeekMBA; Benzinga).

What Would Change Our Mind

A thesis is only as useful as the conditions under which we would abandon it. Defining those conditions in advance is not a disclaimer; it is how our culture is designed to work.

In *Beyond the Four Ps*, we described our investment culture as an intellectual meritocracy with an obligation to dissent. Ideas must stand on their merit, not on the title, tenure or confidence of the person presenting them. Team members vote independently before we discuss a stock, so disagreement is visible before the room converges.

That culture has a practical purpose: it makes it possible to change our minds. Trust makes humility possible, humility allows us to change our minds, and changing our minds lets us correct mistakes quickly. A team that cannot admit error tends to hold too long and turn an investment mistake into a portfolio problem.

For this thesis, the mistake we most want to avoid is a potential value trap. As we argued in our April 2026 thought paper, *Value as a Vector*, most value traps occur when an investor sees a cheap stock but ignores a flat or falling direction of value. If curve A proves right, today's low multiples are not a bargain but the first stage of a permanent loss of capital.

So we are writing our dissent down now, before the market forces the answer. These are the signals that would tell us the Meseta is eroding into a Matterhorn:

- **Backlogs shrink.** Backlogs or contracted revenue fall while revenue growth slows.
- **Prices roll over at the bottlenecks.** Memory contracts, lithography, optics or neocloud capacity reprice lower.
- **Lead times shorten early.** Lead times shorten across the chain before 2028 without a matching addition of capacity.
- **Commitments weaken.** Customers walk away from orders that carry down payments, or repeatedly push out deliveries.
- **Suppliers build inventory.** Inventories and receivables grow faster than revenue at the companies that own the bottlenecks.
- **Efficiency outruns new uses.** Meta expects to improve the compute cost efficiency of its Muse agent more than tenfold within a year. If gains like that shrink total demand rather than widen adoption, the plateau erodes.
- **Delay turns into denial.** Moratoriums block capacity permanently rather than postpone it. Deferred demand waits only if capacity is eventually allowed to be built.
- **Users stop earning a return.** This is the caveat we address below, and the one that matters most.
- **Our own vector turns down.** Our Fair Value Model estimates each company's value over four different time frames. If forward values begin to fall for the companies we own, we treat it as the red flag our process is built to catch.

No single signal decides the question. Each goes through the same process as any other evidence: an independent vote, open debate and a review of what has changed. As we wrote in *Beyond the Four Ps*, the aim is to be confident enough to act and humble enough to change.

Valuation: Priced for the Peak

We believe the parts of the chain that look most like a peak are priced as one. If the Meseta is right, that creates two potential sources of return: earnings growth and a higher multiple.

The market as a whole is discounting a slowdown, as though it can see the valley already. As of October 23, the S&P 500 traded at 18.7 times forward earnings, below its five-year average of 19.8 (FactSet). FactSet's bottom-up consensus, built from analysts' estimates for each company in the index, has earnings growth decelerating from about 33% in 2026 to about 16% in 2027, a drop exaggerated by one-time investment gains at Alphabet and Amazon in 2026. That is a slowdown, not a collapse, yet the chain's most cyclical-looking stocks are priced for a descent into the valley on the back side of the Matterhorn.

Investors are also selling good news. Companies beating estimates this quarter have seen their shares fall 0.2% on average around the report, against a five-year average gain of 0.9% (FactSet).

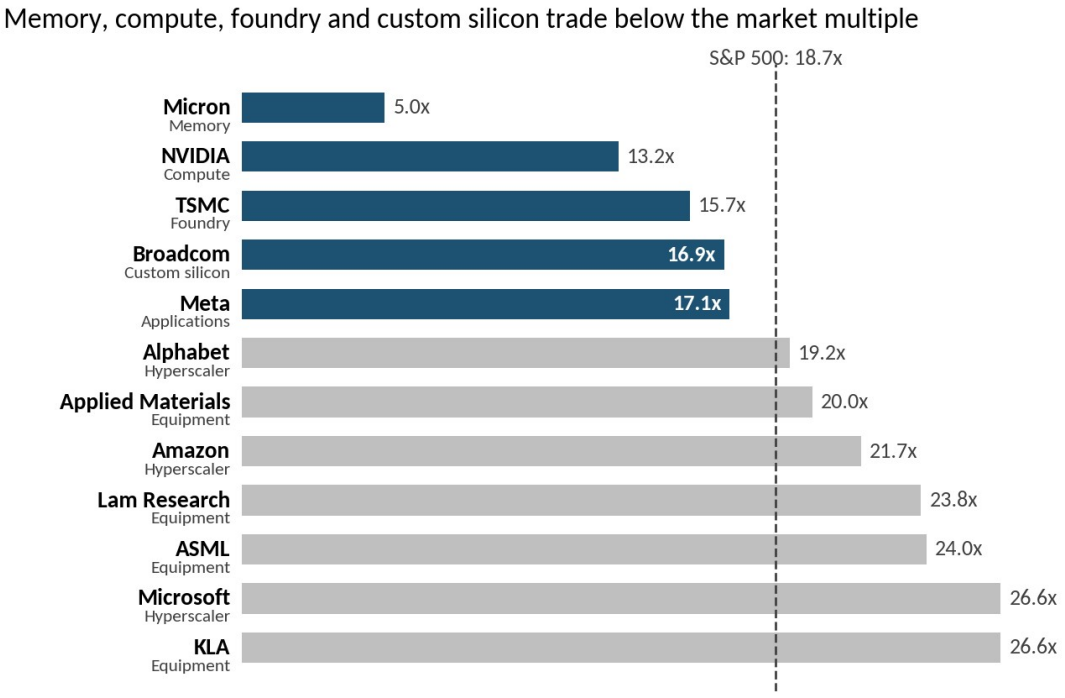


Exhibit 3: P/E on consensus EPS for each company's fiscal year ending in 2027 (NVIDIA: January 2028), Pinegap, at October 23, 2026 closing prices; S&P 500 forward 12-month P/E from FactSet, October 23, 2026.

The compression is concentrated where earnings look most cyclical. Micron, NVIDIA, TSMC and Broadcom all trade below the market multiple, at the links where supply binds hardest, and Meta does too. The equipment makers carry premium multiples; the market already treats their earnings as durable, so their opportunity rests more on earnings growth than on a rerating.

A multiple largely reflects how long the market expects today's earnings to last. On curve A (the Matterhorn), investors rationally pay a trough multiple for peak earnings. On curve B (the Meseta), today's earnings are not a peak but a base, and a base deserves the multiple of a durable business.

Two ways to win, illustrated. Consider a company earning $10 per share and priced at 12 times, or $120, because investors expect earnings to halve. If earnings instead grow 10% a year for two years, to $12.10, and the market comes to value them at 18 times, the shares would be worth about $218. Earnings growth alone lifts the price 21%; the rerating adds another 50% on top, for a combined gain of about 82%. These figures are illustrative, not a forecast.

We would rather own the Meseta than the Matterhorn. If we believed these companies were producing peak earnings, we would not own them. The very constraints that have caused so much fear in the market

are creating a long-term investment case. Steady, capacity-paced growth lets companies plan, price with discipline and earn a return on what they build, instead of overbuilding into a peak and writing it down in the valley that follows. A plateau with an upward grade is less dramatic than a summit, but it is a far better place to build.

Caveat: Ultimately, AI Must Produce ROI for Users

Our thesis has a precondition: the buildout pays off only if the companies using AI earn a return on it. Supply bottlenecks can stretch a cycle, but they cannot create demand that users will not pay for.

The bar is high. Bain estimates that annual AI infrastructure spending could reach $1.5 trillion by 2031, and that sustaining it would require AI revenue approaching $6 trillion a year. It argues that productivity gains from existing applications will not be enough (Bain). In short, cost savings alone will not pay for the buildout; AI has to create new revenue.

The early evidence points to growth, not just cuts:

- **Chief financial officers say growth, not cost cutting.** A survey of nearly 750 CFOs by the Federal Reserve Banks of Atlanta and Richmond and Duke University tied AI productivity gains most closely to developing new or improved products and reaching customers more effectively. Reducing labor costs ranked among the least important reasons to invest. AI-attributed gains in revenue per worker averaged 0.6% in 2025 and are expected to triple to 1.8% in 2026 (NBER).
- **Companies are paying for it.** Blackstone tracks AI across its portfolio companies, GP-stake companies (asset managers in which Blackstone owns minority stakes) and borrowers. Annualized revenue at 14 AI-related companies in that group rose 21-fold in a year, from $25 million to $525 million. That is money customers are paying for AI products, not capacity the hyperscalers are building. “We have an amazing lens at Blackstone through our portfolio companies,” said President Jon Gray (Blackstone).
- **Specific adopters grow faster.** Drawing on a universe of 564 companies, Carnegie Mellon researchers found that those describing their AI deployments most concretely grew revenue about 8 percentage points faster than the least specific, after controls. The individual regressions rest on smaller samples. They found no measurable margin effect, concluding that AI’s impact is showing up in revenue before margins (Help Net Security).
- **Advertisers pay more because ads work better.** Meta’s AI ranking models raised app ad conversions 9.6% in the third quarter. Advertisers who adopted its AI recommendations saw nearly 5% more conversions at the same cost per conversion and spent 9% more than similar non-adopters. Meta’s average price per ad rose 14% year over year, which the company attributes to better ad performance and an improved economy (Meta, Yahoo Finance).
- **Drug discovery is speeding up.** Yale researchers, Google DeepMind and Google Research built C2S-Scale on Gemma, Google’s open model family built from Gemini research. It virtually screened more than 4,000 drugs and proposed a way to make “cold” tumors visible to the immune system, and lab experiments in human cells supported the prediction (Google). Separately, Insilico Medicine’s rentosertib, whose target and molecule were both found by generative AI, entered Phase III trials in 2026 (Insilico).

- **Design tasks are speeding up.** Cadence reports customer use cases with 2 to 40 times efficiency gains from its AI design agents, and Synopsys sees AI-related chip design starts accelerating. Faster design means more products per engineer, not just fewer engineers.
- **A model provider now covers its compute bill.** Anthropic spent $2.51 on compute and infrastructure for every $1 of revenue in the first quarter of 2025, and $0.48 in the second quarter of 2026, when it posted its first operating profit, $974 million (Anthropic Form S-1).
- **Customers are spending more.** Anthropic customers paying $100,000 or more a year rose from about 950 to more than 11,900 in the year to September 30, 2026, and about 250 now pay $10 million or more (Anthropic Form S-1).

The evidence is early and partly self-reported. The drug-discovery examples above remain preclinical or in clinical trials. CFOs report productivity gains larger than what yet shows up in revenue, a modern echo of Robert Solow's paradox, and the Carnegie Mellon study shows association, not causation.

The most serious counterpoint comes from Apollo's chief economist, Torsten Slok. He finds that profit margins outside the largest technology companies have barely moved despite heavy AI spending, and argues that the other 493 companies in the S&P 500 will decide whether today's AI valuations hold (Apollo, via Longbridge).

This is the caveat that matters most. If users stop earning a return, deferred demand would be destroyed rather than delayed, and the curve becomes the Matterhorn after all. We will watch revenue at the users of AI, not only at its suppliers.

Conclusion: Building on the Meseta

As we argued in our March 2026 thought paper, *The Three Orders of Impact*, AI does not change the rules of investing; it raises the stakes of failing to follow them. The market is reading a slower growth rate as the summit. We read it as the edge of a plateau, set by how fast the world can generate power, pour concrete, print wafers and package memory.

Fear, politics and loss aversion have priced many of the chain's most constrained businesses for a descent. If the evidence keeps pointing to demand that is deferred rather than destroyed, those businesses will be valued for the earnings power they hold, not the peak they were feared to be. We will keep watching the backlogs, the down payments and the prices, and we will say so if the mountain changes shape.

The Matterhorn is a place you visit. The Meseta is a place you build on; Spain put its capital there.

Sources: company earnings calls, conference transcripts and SEC filings, detailed in the Evidence Appendix; FactSet Earnings Insight, October 16, 2026; U.S. Treasury daily par yield curve rates; Bloomberg bond yields and BVAL curves, Cornerstone analysis; Axios; AI Weekly; Data for Progress; S. 4214 and S. 3852, the GRID Act (GovInfo); Senator Hawley; New York Executive Order 62; Data Center Watch; Matterhorn (Wikipedia); Meseta Central (Britannica); Microsoft fiscal 2026 Form 10-K and fourth-quarter call coverage; NBER Working Paper 34984; Carnegie Mellon and Larridin, via Help Net Security; Bain Technology Report 2026; Google; Insilico Medicine; Reuters; Bisnow; Meta Form 10-Q, June 30, 2026; Pinegap consensus estimates and VIX closes, October 23, 2026; Telegraph; Benzinga; 24/7 Wall St; Blackstone; CNBC All-America Economic Survey, October 14 to 18, 2026; Anthropic Form S-1, October 26, 2026; Sightline Climate, via Contrary Research; Apollo's Torsten Slok, via Longbridge; Pinegap free cash flow and capital spending data, October 2026; Billboard; Motley Fool; Reuters; Marketing Dive; InfoWorld.

Evidence Appendix

This appendix collects the company commentary behind the assertions in this paper, grouped by the claim each data point supports. Sources are company earnings calls, investor conferences, filings and price lists, plus a few industry studies, mostly from July through October 2026; figures are as of the dates shown.

Acronyms used: WFE (wafer fab equipment, the industry's annual spend on chipmaking tools); EUV (extreme ultraviolet lithography); DUV (deep ultraviolet lithography); DRAM (dynamic random-access memory); NAND (flash memory); HBM (high-bandwidth memory, stacked next to AI chips); TPU (tensor processing unit, Google's in-house AI chip); AWS (Amazon Web Services); RPO (remaining performance obligations, contracted revenue not yet recognized); N2 (TSMC's 2-nanometer process).

1. Supply, not demand, is capping revenue (deferral, not destruction)

Data point	Source
December-quarter revenue guide ($8.6B ± $0.4B) "would have been a bit higher if not for some near-term supply chain constraints." Anything that slipped from December "we will ship that in the March quarter." No one-off items in the guide; it reflects customer demand.	Lam Research, fiscal Q1 2027 call, Oct 21, 2026
"This year has been primarily clean room constrained." Component constraints now emerging on top; clean room schedules, hiring and training also gate 2027.	Lam Research, Oct 21, 2026
Clean room space and component constraints will continue "to limit chip makers and equipment companies' ability to fully meet end demand" in 2027.	Lam Research, Oct 21, 2026
Management says the tightness is "nothing like" the COVID-era shortages; held supplier events to lay out growth expectations for the next couple of years.	Lam Research, Oct 21, 2026
"If we could have built more, we could have shipped more" (referring to first half of 2026).	KLA, fiscal Q4 2026 call, July 28, 2026
Amazon "will still not have enough capacity to meet all the demand we have in 2026," and expects the same in 2027, even after doubling power capacity versus 2025 by end of 2027.	Amazon, Q2 2026 call, July 30, 2026
Google Cloud is "still in a supply constraint environment"; demand still outpaces a capacity base that has grown significantly over three years. Renting third-party capacity in Q3 as a "bridging strategy," accepting near-term Cloud margin pressure.	Alphabet, Q2 2026 call, July 22, 2026
"Demand continues to exceed available supply"; efficiency gains and faster hardware deployment are "quickly monetized in the quarter." Spot-market pricing for assets cited as evidence.	Microsoft, fiscal Q4 2026 call, July 29, 2026
Muse (Meta's personal AI agent) passed 10 million US and Canada users in six weeks but has not launched globally because it is compute-intensive; compute efficiency is "one of the limiting factors over the next year." Building capacity to serve more than 1 billion people.	Meta, Q3 2026 call, Oct 21, 2026
TSMC: "The gap between demand and supply are still quite big." Cannot say whether it is narrowing "because of my customer continue to increase their demand."	TSMC, Q3 2026 call, Oct 15, 2026

2. Visibility extends to 2028 and beyond

Data point	Source
Industry in "early stages of a multi-year investment cycle, with customer visibility extending through 2028."	Lam Research, Oct 21, 2026
Ten new tier-1 fabs come online during 2027 and "layer in over time."	Lam Research, Oct 21, 2026
ASML increasingly receiving orders for 2028; customers asking to place orders for 2029 and 2030. ASML has not yet accepted them, wanting to compare total demand to capacity so "no single customer feels left out." Interest is broad, not one customer.	ASML, Q3 2026 call, Oct 14, 2026
Customers "expect that their supply will not meet the demand at least until 2028."	ASML, Oct 14, 2026
AWS 2027 capacity is "largely reserved," with "quite a bit of capacity that's already been reserved for 2028." "The demand we already have for 2028 is striking."	Amazon, July 30, 2026
TSMC capex for 2026 to 2028 "significantly higher than the past three years"; nearly 20 fabs in progress (13 leading-edge and advanced packaging fabs in Taiwan, 5 to 6 overseas).	TSMC, Oct 15, 2026

Data point	Source
Backlogs: AWS $496B, up triple digits year over year. Google Cloud $514B, up more than $50B sequentially (just over 50% recognized within 24 months). Microsoft commercial RPO $678B, up 84% (25% excluding OpenAI; about 30% expected to be recognized within 12 months).	Amazon (July 30, 2026); Alphabet (July 22, 2026); Microsoft (July 29, 2026, and fiscal 2026 Form 10-K)
Alphabet expects 2027 capex to "increase significantly"; Microsoft expects fiscal 2027 capex to grow year over year "given demand signals across our portfolio."	Alphabet (July 22, 2026); Microsoft (July 29, 2026)

3. Bottlenecks prevent overbuild and shift pricing power to the bottleneck owner

Data point	Source
ASML planned to raise 2027 capacity by 30% for low-NA EUV and DUV immersion, but now says planned 2027 and 2028 output "will not fully meet the demand of our customers"; looking at 2028 capacity growth above 30%. Suppliers are being prepared for roughly 30% growth per year.	ASML, Oct 14, 2026
ASML completed price increases on its most critical products: some effect in 2027, full and sustained effect from 2028. Significant portion flows directly to gross margin.	ASML, Oct 14, 2026
TSMC does not allow customers to "sublet or auction off the capacity."	TSMC, Oct 15, 2026
TSMC: "There is no ceiling for TSMC's profitability, but a dramatic increase will not happen overnight." Tool, material and labor price inflation "is real."	TSMC, Oct 15, 2026
Lam working on pricing toward mid-50s gross margin over several years.	Lam Research, Oct 21, 2026
Nokia (readthrough): current supply constraint is broad-based across leading-edge nodes, memory and printed circuit boards, and is "a structural cost increase," unlike the narrow 2021 shortage.	Nokia, Q3 2026 call, Oct 22, 2026
Memory price inflation pushed Amazon's 2026 cash capex to about $220B from about $200B. New AWS contracts are priced to reflect current costs.	Amazon, July 30, 2026
KLA paying higher prices for memory used in its own tools, costing a bit more than 100 basis points of gross margin; expected to continue through 2027.	KLA, July 28, 2026
Neocloud Nebius raised on-demand GPU rental rates 17% to 21% from October 1: H100 $3.85 to $4.50 an hour, H200 $4.50 to $5.40, B200 $7.15 to $8.50, B300 $7.85 to $9.50. It was the company's second increase in a few months, after a May rise averaging about 29% on on-demand capacity.	Nebius price list, effective Oct 1, 2026 (via FourWeekMBA and Benzinga)

4. Demand quality: commitments are real

Data point	Source
ASML purchase orders require a down payment "enough to hurt" and are costly to cancel. ASML validates demand against customers' fab construction plans and calls it "real and not inflated" by order-securing behavior.	ASML, Oct 14, 2026
Most AWS AI capacity is contracted for at least five-year terms.	Amazon, July 30, 2026
Most Google Cloud infrastructure contracts are five-year commitments.	Thomas Kurian, Google Cloud CEO, Goldman Sachs Communacopia + Technology Conference, Sept 8, 2026
Amazon buys servers "a few months before putting them into service," so it has "strong visibility into customer demand before we trigger the spend. If the demand isn't there, we won't spend the capital."	Amazon, July 30, 2026
Google Cloud customers are exceeding their commitments by more than 50%, an acceleration from the prior quarter.	Alphabet, July 22, 2026
Anthropic: at least $265.6B of fixed or minimum compute commitments at June 30, 2026, plus about $161.2B of Broadcom equipment leases that cannot be cancelled except on default; payments run to 2045. Anthropic typically procures compute 12 to 18 months before it becomes available.	Anthropic Form S-1, Oct 26, 2026

5. Payback periods (profitability of the buildout)

Company	Figure	Source
Alphabet	AI servers pay back in under two years; Google's own silicon in roughly half that (about one year). Most infrastructure contracts are five-year commitments.	Thomas Kurian, Google Cloud CEO, Goldman Sachs Communacopia + Technology Conference, Sept 8, 2026
Amazon	Servers and networking take "a little less than three years to break even"; servers have a useful life of "at least five to six years." Most AI capacity is contracted for at least five-year terms.	Amazon earnings call, July 30, 2026

Takeaway: at both companies, server capital is recovered well inside typical five-year contract terms.

6. Growth rates in AI services

Company	Data point	Source
Amazon	AWS revenue +36.7% to $42.2B, fifth straight quarter of acceleration, fastest in 18 quarters; $169B annualized run rate. AI revenue run rate over $25B, growing triple digits. Custom chips business also over $25B run rate, growing triple digits. AWS operating margin up 650 basis points year over year.	July 30, 2026
Alphabet	Google Cloud revenue +82% to $24.8B, accelerating even excluding first TPU system sales. Cloud operating margin 35.6% versus 20.7% a year earlier. Model APIs process about 22 billion tokens per minute, up from 16 billion a quarter earlier. Gemini app 950 million monthly active users.	July 22, 2026
Microsoft	Azure +43%, guided to about 45% for the September quarter. Azure surpassed $100B annual revenue, up 41% for the fiscal year. Microsoft 365 Copilot over 30 million paid seats; Copilot revenue up over 60% quarter over quarter. Capacity on track to roughly double in two years.	July 29, 2026
Meta	Q3 revenue +26% to $64.6B. Muse at 10 million+ users in six weeks, about 25% daily usage. New AI businesses to begin contributing revenue in 2027. 2026 capex $135B to $145B.	Oct 21, 2026
Lam (demand signal)	Reported average weekly token usage up roughly 250-fold over the past 18 months.	Oct 21, 2026

7. Equipment and foundry growth figures

Company	Data point	Source
Lam Research	September-quarter revenue $8.22B, +22% sequentially, +54% year over year. Calendar 2026 systems revenue growing well over 40%. WFE mid-$150B in 2026; preliminary 2027 WFE above $200B, led by DRAM, then leading-edge foundry logic and NAND. Customer support revenue to grow more than 15% in 2027.	Oct 21, 2026
KLA	In July cited 2027 WFE consensus near $190B (mid-20s percent growth), and was planning for scenarios 10% to 20% above that. Growth depends on peers also delivering. September-quarter results not yet reported as of Oct 22.	July 28, 2026
ASML	2026 sales €44B to €45B (about +36%), EUV about +45%, installed base about +50%. About 65 low-NA EUV systems in 2026, up nearly 50%. Q4 guide €14.2B to €15.2B. Q3 free cash flow €11.4B, driven by down payments.	Oct 14, 2026
TSMC	Q3 revenue $46.7B, +16.2% sequentially. Q4 guide $52.3B to $54.3B, about +58% year over year. 2026 growth about 44%. AI accelerator forecast of mid-to-high 50s percent annual growth for 2024 to 2029 "probably will be higher." N2 capacity to more than double in 2027. 2026 capex $60B to $64B.	Oct 15, 2026

8. Counter-evidence and risks

- **Scarcity flatters returns.** Microsoft cites spot-market pricing for its assets. If supply catches up, pricing and paybacks normalize. Memory inflation is simultaneously raising capex (Amazon $200B to $220B).

- **Efficiency could shrink compute demand.** Meta expects to improve Muse compute cost efficiency more than tenfold over the next year and to deliver current frontier intelligence "at a fraction of the cost in a smaller model a year from today." Meta frames this as what unlocks rollout to a billion users (supports deferral), but it is the key variable to monitor.
- **Financing strain.** Alphabet free cash flow was negative $5.9B in Q2 2026; long-term debt $98.2B, plus an at-the-market equity program. Meta free cash flow was negative $2.1B in Q3 2026 with $83.7B of debt. Combined free cash flow at Microsoft, Alphabet, Amazon and Meta fell from $60.5B in Q4 2025 to $6.9B in Q2 2026 as quarterly capex rose from $118.6B to $165.1B (company reports, via Pinegap). Off balance sheet, leases signed but not yet started totaled about $329B at Microsoft, $279B at Meta (plus about $68B signed in July), $137B at Amazon (including warehouses, offices and aircraft) and $85B at Alphabet at mid-2026 (Reuters, Aug 4, 2026); Meta also guarantees about $28B of residual value on the Hyperion venture's $27B of debt. Microsoft expects to remain free cash flow positive in fiscal 2027.
- **Some constraints may clear quickly.** Lam expects its component tightness to be "caught up very shortly," though clean room constraints persist. Our thesis rests on the slower physical constraints (clean rooms, lithography output, power), not on component shortages.
- **Bottleneck interdependence.** KLA's upside depends on peers delivering; the slowest supplier sets the pace.
- **ASML technology read-through.** Lam's Aether dry photoresist reduces EUV exposure dose by 10%, raising scanner productivity, which modestly reduces scanner units needed per wafer.

9. Supply constraints by value-chain link

Value chain, from end demand back to materials: applications (AI-native builders plus incumbents monetizing AI in existing businesses, such as Meta and Google advertising) > model providers > hosting (hyperscalers, neoclouds) > systems integration (Dell, Penguin) > components (compute, memory, networking and optics) > custom silicon (Broadcom, Marvell) and design tools (Synopsys, Cadence) > foundry (TSMC) > advanced packaging and test (FormFactor, Onto) > wafer fab equipment (ASML, KLA, Lam, Applied) > subsystems and materials (MKS, Entegris). Physical layer (power generation, data center construction, electrical and cooling) and financing are cross-cutting inputs.

Status key: Binding = demand clearly exceeds supply, with constraints expected into 2027 or 2028. Tight = constrained now, easing expected in 2027. Not constrained = growth is demand-driven.

Link	Status	Evidence	Source
Applications	Rationed by compute	Muse not launched globally because it is compute-intensive; compute efficiency "one of the limiting factors over the next year." Google allocates compute by priority: frontier model development first, then Search and YouTube, then Cloud.	Meta, Oct 21, 2026; Alphabet, July 22, 2026
Model providers	Binding	Anthropic's S-1: demand will outstrip supply, "limited principally by the availability of compute"; it has "significantly underestimated" its compute needs in the past, leaving it capacity constrained. Broadcom expects Anthropic to deploy 10 GW and OpenAI 5 GW in its fiscal 2028; has arranged financing for both, most recently a $63B vehicle with Bank of America, Blackstone, Citi and Morgan Stanley that leases its chips to Anthropic from 2027; says each gigawatt could support about $30B of annual revenue for them. AMD names OpenAI, Meta and Anthropic as anchor customers ramping.	Anthropic Form S-1, Oct 26, 2026; Broadcom, Sept 2, 2026; AMD, Aug 4, 2026
Hosting: hyperscalers	Binding through 2027	AWS lacks capacity for 2026 and 2027 demand; Google renting third-party capacity; Microsoft demand exceeds supply (see sections 1 and 2).	See sections 1 and 2

Link	Status	Evidence	Source
Hosting: neoclouds	Binding	CoreWeave calls it “a systemic disequilibrium” that has existed for several years and will continue “for the foreseeable future.” Expects contribution margins on newly signed contracts 5 to 10 points higher than on contracts added in prior quarters. Sold capacity on a 2020-architecture GPU in a contract priced out to 2029; older-generation prices at or above year-ago levels.	CoreWeave, Aug 11, 2026
Systems integration	Binding on parts	Dell: constraints are “DRAM, DRAM, followed by NAND, NAND,” plus spotty CPU shortages, disk drives, mature-node power chips, ABF substrate (the base layer chips are mounted on), T-glass and optics. Says the AI supply chain is working at full stretch (official transcript: “red lines all out”). Raised second-half guidance by $25B by reallocating parts from PCs.	Dell, Sept 1, 2026
Components: AI accelerators	Binding	NVIDIA expects about 70% revenue growth in fiscal 2028 (ends January 2028), which it calls a supply-constrained outlook; customer forecasts point to growth doubling, and management says demand is “much greater than 70%.” “Our entire supply chain is challenged. Everybody is really running flat out.” Expects supply to remain a bottleneck at least through the end of fiscal 2028.	NVIDIA, Aug 26, 2026
Components: server CPUs	Tight, easing 2027	AMD: server CPU supply tight in 2026 because demand was unforecasted; 2027 better; working on wafers, back-end, substrates. Intel: substrates and memory the hardest constraints; some logjams breaking late Q3.	AMD, Aug 4, 2026; Intel, July 23, 2026
Components: memory	Binding through 2028 and beyond	Micron: more than 75% of 2027 output committed; allocation talks moved to 2028; strategic agreements extended to 2031; “we really do not have line of sight to when supply and demand balances.” New capacity arrives in stages: Singapore HBM packaging initial output early 2027, Idaho ID1 wafer output and meaningful Tongluo shipments mid-2027, ID2 late 2028, New York 2030; each ramps gradually; HBM consumes more wafer per bit.	Micron, Sept 30, 2026
Components: AI optics	Binding through 2027 to 2028	Ciena: backlog of at least $10B; 2027 revenue floor $8.3B to $8.4B “would be greater than that if we had more supply”; balance not before 2028; price increases of high single digits to low 20s percent, some applying to backlog. Lumentum: EML lasers behind demand even at end of 2027; high-power lasers “very much further behind.” Coherent: indium phosphide capacity is “our primary constraint.”	Ciena, Sept 3, 2026; Lumentum, Aug 11, 2026; Coherent, Aug 12, 2026
Components: enterprise networking	Not constrained	Extreme Networks: supply constraints resolved, “clear visibility for supply” into 2028 and beyond.	Extreme, Aug 5, 2026
Custom silicon	Constrained upstream, not by design capacity	Broadcom’s $115B (FY2027) and $230B (FY2028) AI outlook is built against data center site readiness, leading-edge wafers, substrates and HBM supply. Marvell raised next year by $1.5B on scale-up optics.	Broadcom, Sept 2, 2026; Marvell, Aug 27, 2026
Design tools	Not constrained	Cadence raised annual growth to about 19%, its largest single-quarter raise; Synopsys sees AI design starts accelerating. Software; demand-driven.	Cadence, July 27, 2026; Synopsys, Aug 26, 2026
Foundry	Binding	TSMC: demand-supply gap “still quite big”; N2 capacity to more than double in 2027.	TSMC, Oct 15, 2026
Advanced packaging and test	Tight	Onto: advanced packaging up at least 80% in 2026, “not seeing any signs of overcapacity,” customers’ customers pressing for more capacity; $1.1B backlog. FormFactor: co-packaged optics probe revenue running ahead of plan. Lam: advanced packaging up about 70% this year.	Onto, Aug 6, 2026; FormFactor, July 29, 2026; Lam, Oct 21, 2026
Wafer fab equipment	Binding at lithography; tight elsewhere	ASML output will not meet 2027 and 2028 demand; Lam guide trimmed by component tightness; KLA could have shipped more in first half 2026 (see sections 1 to 3).	See sections 1 to 3
Subsystems and materials	Mostly not binding yet	MKS: building capacity “to make sure that we’re not the constraint”; lead times “still kind of normal” (August). Entegris: direct labor up more than 20% since end of 2025, investing ahead. Exceptions that are binding: indium phosphide (Lumentum, Coherent), ABF substrate and T-glass (Dell, AMD, Intel).	MKS, Aug 6, 2026; Entegris, Aug 4, 2026

Link	Status	Evidence	Source
Physical: power generation	Binding	GE Vernova: most of its 30 GW of 2030 gas turbine production slots will be committed this year, through orders and slot reservations; castings and forgings arriving in 2027 gate the step to 24 GW in 2028. $5B of first-half orders from data center customers.	GE Vernova, July 22, 2026
Physical: data center construction	Binding	Comfort Systems: modular capacity expanding 4M to 5M square feet, "overwhelmingly for existing customers and existing orders"; customers want "to buy as much of this as we can produce." EMCOR: data centers are being built where there is power and willingness to build power; AI sites need 1.5x to 2x the mechanical and electrical capacity. Sightline Climate: of about 16 GW slated for 2026 across roughly 140 projects, only about 5 GW was under construction in February 2026. Data Center Watch: local opposition blocked or delayed at least 75 U.S. projects worth about $130B in Q1 2026.	Comfort Systems, July 24, 2026; EMCOR, July 30, 2026; Sightline Climate, Feb 2026; Data Center Watch, Q1 2026
Physical: electrical and cooling	Tight	Vertiv: supply chain interdependencies on large projects delayed some Q2 revenue into the second half.	Vertiv, July 29, 2026
Physical: chip clean rooms	Binding	Lam: "This year has been primarily clean room constrained." Micron: new fabs begin output in stages from mid-2027, with more in late 2028 and beyond.	Lam, Oct 21, 2026; Micron, Sept 30, 2026
Financing	Available to the largest builders at low rates; watch item	Broadcom's $63B leasing vehicle for Anthropic with Bank of America, Blackstone, Citi and Morgan Stanley (Oct 27, 2026), and its CEO's view that the cost of capital is now the main bottleneck; CoreWeave's new financing facility opened two- to three-year contracts. Negative free cash flow at Alphabet and Meta (section 8).	Broadcom, Sept 2, 2026; CoreWeave, Aug 11, 2026

Staged demand (the plateau mechanism): Entegris describes revenue arriving in waves measured from the start of fab construction: about 12 months for facilitization, 18 months for process piping, 24 months for tools, then wafer volumes. It is tracking 20 leading-edge capacity expansions (8 to 10 logic, 7 to 8 memory, 6 to 8 packaging). Each layer's demand peaks at a different time, which stretches a spike into a plateau. (Entegris, Aug 4, 2026)

10. Returns for users of AI (the caveat)

Data point	Source
Annual AI infrastructure spending could reach $1.5T by 2031, requiring AI revenue approaching $6T a year; productivity gains from existing applications will not be enough.	Bain Technology Report 2026
Survey of nearly 750 CFOs: AI productivity gains are tied most closely to new or improved products and reaching customers; cutting labor costs ranks among the least important motives. AI-attributed gains in revenue per worker averaged 0.6% in 2025 and are expected to reach 1.8% in 2026.	Federal Reserve Banks of Atlanta and Richmond and Duke University, NBER Working Paper 34984, Mar 2026
Across a universe of 564 companies, those describing their AI deployments most concretely grew revenue about 8 percentage points faster than the least specific, after controls; no measurable margin effect yet.	Carnegie Mellon and Larridin, Aug 2026
Meta: AI ranking models raised app ad conversions 9.6% in Q3; advertisers adopting its AI recommendations saw nearly 5% more conversions at the same cost per conversion and spent 9% more than similar non-adopters; average price per ad +14%, which Meta attributes to ad performance and an improved economy.	Meta, Q3 2026 call, Oct 21, 2026
Blackstone: annualized revenue at 14 AI-related companies across its portfolio, GP-stake companies and borrowers rose 21-fold in a year, to $525M.	Jon Gray investor presentation, Sept 20, 2026
C2S-Scale, built by Yale, Google DeepMind and Google Research on Gemma, virtually screened 4,000+ drugs and proposed a way to make "cold" tumors visible to the immune system, supported by lab experiments in human cells. Insilico's rentosertib, with target and molecule found by generative AI, entered Phase III in 2026.	Google, Oct 2025; Insilico Medicine, Sept 10, 2026
Cadence: customer use cases with 2 to 40 times efficiency gains from its AI design agents. Synopsys sees AI-related chip design starts accelerating.	Cadence, July 27, 2026; Synopsys, Aug 26, 2026

Data point	Source
Anthropic: compute and infrastructure cost per $1 of revenue fell from $2.51 in Q1 2025 to $0.48 in Q2 2026, its first quarter of operating profit ($974M). Customers paying $100K+ a year rose from about 950 to more than 11,900 in the year to Sept 30, 2026; about 250 pay $10M+.	Anthropic Form S-1, Oct 26, 2026
Counterpoint: profit margins outside the largest technology companies have barely moved despite heavy AI spending.	Torsten Slok, Apollo, July 2026

For more information, please contact:

www.cornerstone-ip.com | marketing@cornerstone-ip.com | (404) 751-3850

IMPORTANT DISCLOSURES

Cornerstone Investment Partners, LLC ("Cornerstone") is an independent, employee-owned investment advisory firm registered with the SEC. SEC registration does not imply a certain level of skill or training. Cornerstone's Form ADV Part 2A is available upon request.

The opinions expressed are those of the investment professionals as of October 2026, are subject to change without notice, and do not constitute investment advice. This paper is not tailored to the objectives or circumstances of any investor. Past performance is not indicative of future results. Investing involves risk, including the possible loss of principal.

Forward-looking statements. *This paper contains forward-looking statements, including views on industry supply and demand, company earnings, credit capacity and valuation. These statements reflect current assumptions and are subject to risks and uncertainties; actual results may differ materially. Cornerstone undertakes no obligation to update them.*

Securities mentioned. *References to specific securities, indices and market data are for illustrative purposes only and do not constitute a recommendation to buy or sell any security. Cornerstone, its clients and/or its employees may hold positions in some or all of the securities discussed, and may buy or sell them at any time without notice. Holdings are as of October 23, 2026 and are subject to change. The securities discussed do not represent all securities purchased, sold or recommended for client accounts. It should not be assumed that investments in the securities identified were or will be profitable.*

Hypothetical illustrations. *Exhibit 1 and the "Two Ways to Win" example are hypothetical. They are not based on any actual security, portfolio or client account, do not represent actual or projected performance, and do not reflect fees, expenses or taxes. The same mechanics can work in reverse: declining earnings combined with a lower valuation multiple can produce significant losses.*

Risks. *The companies discussed are concentrated in technology, semiconductor, infrastructure and related industries. These industries can be more volatile than the broader market and are subject to rapid technological change, cyclical demand, supply chain disruption, export controls, trade restrictions, government regulation and local opposition to development. Certain companies discussed are non-U.S. issuers whose securities may be subject to currency, political, regulatory and accounting risks, including risks associated with American Depositary Receipts.*

Third-party information. *Certain information contained herein has been obtained from third-party sources, including company filings and earnings call transcripts, consensus estimates (Pinegap, FactSet), bond market data (Bloomberg), credit ratings (Moody's, S&P), and survey and polling data. This information has not been independently verified by Cornerstone. No representation, warranty or undertaking, expressed or implied, is given as to its accuracy or completeness by Cornerstone or any other person. Information drawn from preliminary registration statements may be amended before an offering is completed. Credit ratings are the opinions of the issuing agencies, not recommendations.*

Indices. *The S&P 500 Index, Nasdaq Composite and Cboe Volatility Index (VIX) are unmanaged, do not reflect fees or expenses, and cannot be invested in directly. They are shown for market context only and are not benchmarks for any Cornerstone strategy.*

Political and geopolitical references. *References to elected officials, candidates, legislation, polling and geopolitical events are included solely to analyze their potential market effects. They do not represent an endorsement of, or opposition to, any political party, candidate or policy position.*

Trademarks. *Company names, products and trademarks referenced are the property of their respective owners. Their use does not imply any affiliation with or endorsement of Cornerstone.*

Except where otherwise indicated, the information contained in this paper is based on matters as they exist as of the date of preparation (market data as of October 23, 2026 unless otherwise noted) and not as of the date of distribution or any future date. Recipients should not rely on this material in making any future investment decisions. This should not be deemed as an offer or solicitation to buy or sell any product or security.