



DIGITAL POWER WEEKLY

\$715B

BIG TECH 2026 CAPEX

\$110B

GEN-AI REVENUE (EX-CHINA)

\$50B

MGX AI FUND RAISED

\$270B

LEVERAGED-ETF SPECULATION AUM

WEEK IN REVIEW

From Boom to Audit: The Week the AI Trade Started Pricing Its Own Plumbing

The dominant signal this week was a market beginning to audit the AI buildout rather than simply applaud it. Three forces converged: a credible bottom-up overall AI revenue print that, for the first time, cleared the depreciation hurdle; a sharp, retail-fueled equity and credit wobble that exposed how much leverage sits underneath the trade; and a regulatory and grid apparatus that is now treating compute load as a first-order reliability and fiscal problem. The throughline is unchanged and strengthening — controllable, contracted, power-backed capacity is the asset; announced gigawatts are not.

On the demand side, Exponential View's first *State of the AI Economy* report reconstructed roughly \$110B of trailing-twelve-month generative-AI revenue and a ~\$175B run-rate, growing about three times faster than the mobile or internet waves. Bloomberg's read of the same data found Q1 2026 AI sales of ~\$25B narrowly clearing ~\$21B of depreciation — the second straight quarter above the line, but with depreciation still consuming more than two-thirds of revenue. The economics hold; the margin for error is thin.

On the capital side, the plumbing showed stress. SpaceX's \$25B investment-grade deal — the same offering that drew ~\$85–90B of orders — bled roughly \$305M of paper losses in the secondary market within days, an almost unprecedented spread move for a megadeal. A retail-driven unwind in leveraged AI-chip ETFs erased more than 20% from Korea's hottest vehicles, Cerebras slipped below its IPO price on a soft outlook, and two prominent Chinese hedge funds called the cycle a 'super bubble.' Capital is still abundant — MGX closed ~\$50B, SK Hynix, ByteDance and NTT all queued megaraises — but it is more indiscriminating.

In terms of the infrastructure race, new thresholds continue to be crossed. Microsoft no longer just buys electricity—it buys power plants with its signed 20-year agreement to build dedicated gas-fired generation with Chevron. OpenAI unveiled its first custom AI accelerator. NVIDIA effectively rewrote the economics of liquid cooling, Meta and Microsoft pushed future data-center lease commitments above \$850 billion and the Tennessee Valley Authority (TVA) became one of the first major U.S. utilities to formally acknowledge that actual electricity demand is already exceeding the assumptions embedded in its recently completed Integrated Resource Plan, largely because of hyperscale data-center development. TVA is now evaluating between 7 GW and 26 GW of incremental natural-gas generation by 2040, alongside up to 5 GW of nuclear, 5 GW of storage, and 8 GW of renewable generation.

And the grid kept converting AI demand into policy. PJM added a new 'capacity advisory' to warn of tightness outside heat waves; Moody's flagged data-center credit risk for states and munis; and the distributed-compute thesis advanced from slideware to startups, with Span's backyard nodes and a16z's bet on Ornn's compute-trading marketplace. Each is a symptom of the same shift: the bottleneck is physical, and the winners will be whoever can prove a financeable, defensible energization plan.

Read-Through: The investable lesson compounds week over week — demand is real and now partially proven, but the clearing price is execution. Platforms that pair controllable load with dedicated or proximate power, flexible operations and credible financing will separate decisively from land-and-a-service-date stories as capital discipline rises.

In this Issue

03 Fast Facts

07 Nuclear Bid for AI Power

11 The Demand Curve

17 DPW Investment Thesis

BOTTOM LINE

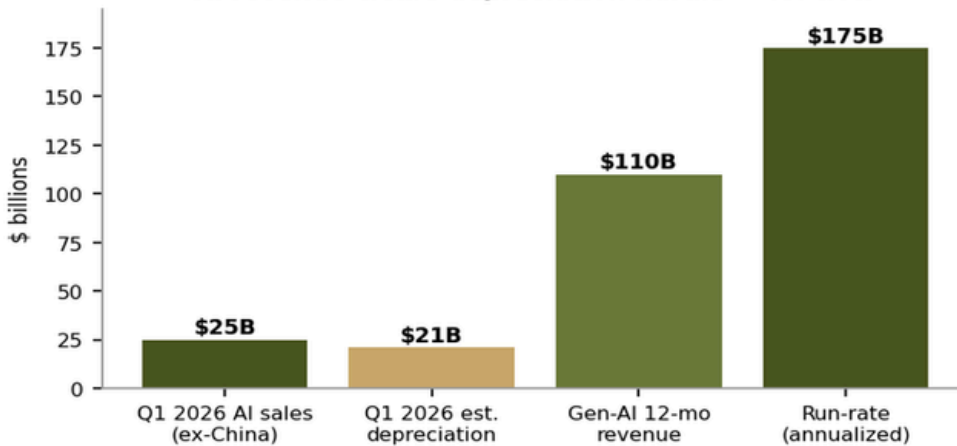
This week we saw that the AI revenue base is now partially proven, which raises rather than lowers the bar: every megawatt and every dollar of capex must point to contracted, cash-generative load. The seemingly separate stories rhyme. A proven-but-thin revenue print, a stunningly weak SpaceX bond tape, a \$270B leveraged-ETF machine working in reverse, a new PJM emergency tier, Moody's muni warnings, and the pivot to distributed/backyard compute are all the same paradigm shift expressed in different markets: AI has become an infrastructure race, and infrastructure races reward controllable power, legible risk and local legitimacy.

★ MARKET SPOTLIGHT | FEATURED ANALYSES

THE REVENUE PROOF ARRIVES — AND RAISES THE BAR

For two years the central bear argument was that AI demand could never justify the supply-side capex. This week produced the most credible rebuttal yet — and, paradoxically, a tougher standard. Exponential View's *State of the AI Economy* assembled what it calls the first bottom-up, de-duplicated measure of consumer and enterprise AI spend across the full stack: roughly \$110B of generative-AI revenue over the trailing twelve months, with the latest month annualizing to about a \$175B run-rate. Critically, the methodology reports only the end-customer dollar, stripping out the double-counting that inflates supply-chain tallies.

AI revenue clears depreciation hurdle — for now




AI sales (ex-China) cleared estimated depreciation for a second straight quarter — but depreciation still consumes ~two-thirds of revenue. Source: Exponential View; Bloomberg.

Bloomberg's coverage sharpened the commercial read. Global AI sales excluding China reached about \$25B in Q1 2026 against an estimated \$21B of depreciation tied to data centers and chips — clearing the hurdle, but with only a thin buffer for power, labor and financing. Exponential View's founder Azeem Azhar framed it precisely: the economics 'just about clear the depreciation hurdle,' and at this stage of a capex cycle you would not expect a dramatic overshoot.

Two structural findings matter most for Keel's thesis. First, depreciation assumptions are the swing variable: the report uses a six-year life for GPUs, defended by the observation that H100 rental rates remain near 80% of launch levels into year four and that AWS has kept six-year-old A100 servers in service. If chips depreciate faster than assumed, the thin margin inverts quickly — which is exactly why durable, low-cost, contracted power is the hedge that protects project economics. Second, demand is migrating toward cheaper inference: the frontier-lab token share on OpenRouter fell to 33% in June 2026 from 72% a year earlier as power users route simple tasks to open-weight and Chinese models.

Frontier-lab token share (Google/OpenAI/Anthropic) on OpenRouter



Read-Through: A partially proven revenue base is bullish for genuine infrastructure and bearish for speculative capacity. If inference keeps commoditizing and distributing, demand becomes more cost- and latency-sensitive — favoring modular, power-backed brownfield campuses over single-purpose megaclusters. The economics hold only while power is cheap, firm and contracted.

Frontier-lab token share collapsed year-over-year as inference commoditizes. Source: OpenRouter via [Exponential View](#) and [Bloomberg — AI Spending Boom](#). Links: [Exponential View — State of the AI Economy](#) | [Bloomberg — AI Spending Boom](#)

FAST FACTS FOR THE WEEK

A by-the-numbers cut of the week’s defining data points across capital markets, compute, power, grid reliability, regulation and the international build-out — the quantitative spine of the stories above.

The other watch item is transparency. FERC wants better visibility into transmission costs and cost shifting. That is ratepayer protection, but it also makes hyperscaler-driven upgrade charges harder to hide inside generic grid investment. Data-center buyers will still get power, but the bill is moving closer to the customer and the operating conditions are becoming explicit.

CAPITAL MARKETS

Metric	This Week
SpaceX bond offering	\$25B
Initial investor demand	~\$85B
Secondary-market performance	Bonds widened after issuance as investors repriced execution risk and realized \$305MM paper loss
Chinese hedge funds	2 leading managers warned AI has become a “super bubble”
Chinese funds expressing AI valuation concern	At least 4
Cerebras 2026 revenue guidance	\$855–865M
Cerebras 1Q26 revenue	\$193.4M
Cerebras quarterly net loss	(\$14M)

AI COMPUTE

Metric	This Week
AWS Nvidia GPU pricing	Raised approximately 20%
Cerebras share performance	Fell below IPO price following guidance
AI investment theme	Market shifting from compute growth toward execution quality

POWER & UTILITIES

Metric	Value
U.S. data-center load (2025)	31 GW
U.S. data-center load (2027)	66 GW
Data-center share of U.S. electricity by 2030 (EPRI)	Up to 17%
Peak load growth (2025→2027)	4.1% → 8.5%
Large-load trend	Utilities negotiating flexible interconnections instead of traditional firm service

POWER & UTILITIES

Metric	Value
Coal unavailable-energy increase	39.8 TWh
Gas unavailable-energy increase	19.1 TWh
NERC conclusion	Dispatchable reserve margins continue to deteriorate

GRID RELIABILITY

Metric	Value
Coal unavailable-energy increase	39.8 TWh
Gas unavailable-energy increase	19.1 TWh
NERC conclusion	Dispatchable reserve margins continue to deteriorate

INTERNATIONAL

Metric	Value
Australia data-center share of grid demand (2050)	~10%
Increase vs. today	Approximately 4x
Australia grid consumption	Nearly doubles by 2050
Lowest-cost grid path	Renewables + transmission + storage + natural gas

INVESTOR SENTIMENT

Metric	Interpretation
AI valuations	Increasing scrutiny
Infrastructure financing	Still abundant
Power availability	Primary constraint
Utility flexibility	Emerging competitive advantage
Execution risk	Increasingly driving equity performance

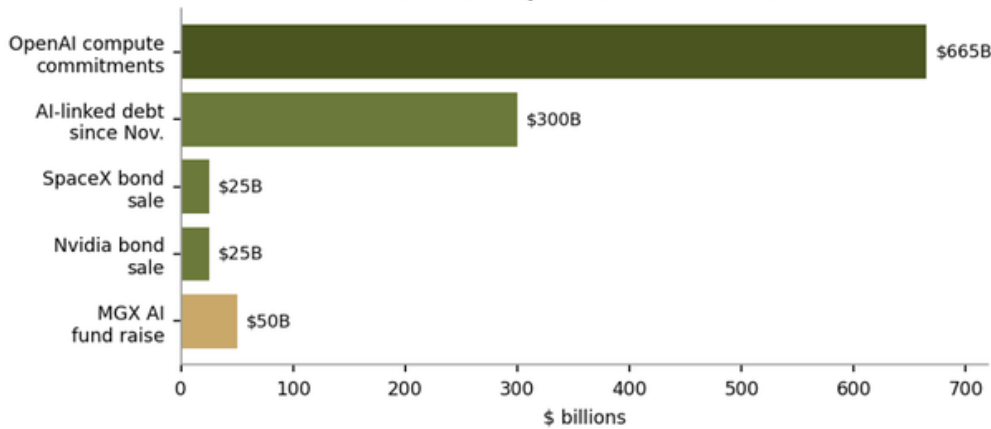
DPW TAKEAWAY - THE WEEK IN FIVE NUMBERS

Number	Why It Matters
66 GW	Projected U.S. AI data-center demand by 2027
17%	Potential share of U.S. electricity consumed by data centers by 2030
\$25B	SpaceX bond issuance demonstrating continued infrastructure-capital availability
39.8 TWh	Increase in unavailable coal generation, highlighting worsening reliability
10%	Data centers' projected share of Australia's electricity demand by 2050

DPW Read-Through: The five numbers rhyme: demand is accelerating (66 GW, 17%), capital remains available to credible issuers (\$25B), reliability is deteriorating (39.8 TWh), and the build is going global (10%).

TOP STORIES

AI infrastructure capital markers cited this week



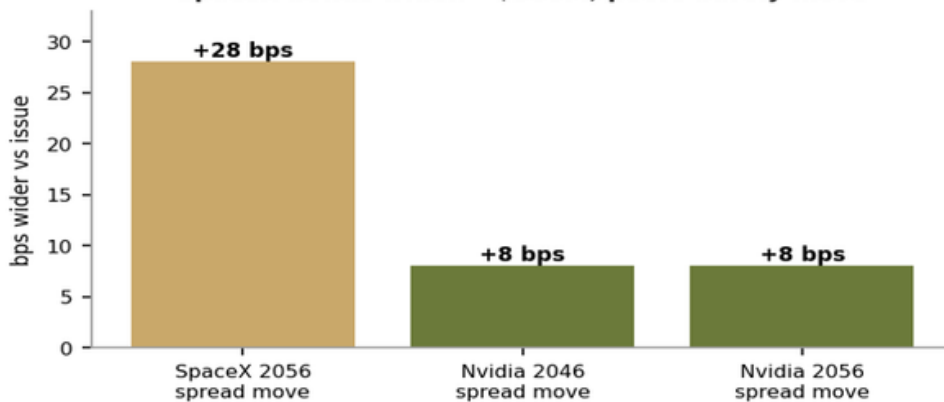
AI infrastructure capital markers cited this week (\$B). Source: Bloomberg; The Information; company disclosures.

Links: [Bloomberg — SpaceX bond losses](#) | [Bloomberg — MGX \\$50B](#) | [Bloomberg — AI Rout](#)

CAPITAL IS STILL ABUNDANT — BUT IT STOPPED BEING INDISCRIMINATE

The defining capital-markets event was not a raise but a repricing. SpaceX's \$25B investment-grade offering — folding in X and xAI debt while cutting blended interest costs, and drawing as much as ~\$90B of orders — began bleeding almost immediately. Paper losses reached roughly \$305M against Treasuries within days, with the 2056 maturity widening as much as 0.28 percentage point past an issue spread of 1.75 points over Treasuries. Traders attributed the move to fast-money accounts flipping the deal, a \$600B+ decline in SpaceX's equity value since launch, and genuine difficulty pricing a Musk-dependent credit that Fitch flagged as a 'key rating constraint.' Nvidia's comparable \$25B deal, by contrast, widened only ~8bp.

SpaceX bonds widen ~\$305M; peers barely move



SpaceX long bonds widened far more than similarly-sized peers post-issue. Source: Bloomberg.

Underneath the marquee deals, a \$270B leveraged-ETF complex amplified the unwind. A retail-driven rotation out of AI chip names hammered Korean 2x/3x vehicles (several down 20%+ in a week), dented newly launched SpaceX leveraged funds (~40% off their debut despite ~\$1B of inflows), and rattled crypto wrappers tied to Strategy. The S&P 500 fell ~2% and the Nasdaq 100 more than 4% on the week. The products did not cause the selloff — but as they scale, they increasingly shape the feedback loop between sentiment, flows and price.

Yet the funding window is demonstrably still open for credible balance sheets. Abu Dhabi's MGX closed close to \$50B — one of the largest dedicated AI vehicles ever — and is targeting \$100B+ AUM at up to \$10B of annual deployment. Across Asia, SK Hynix is pursuing a \$29.4B Nasdaq listing, ByteDance a ~\$20B offshore loan, and NTT a ~\$10B bond. Links: [Bloomberg — SpaceX debt losses](#) | [Bloomberg — MGX raise](#) | [Bloomberg — \\$270B speculation machine](#)

COMMODITIZATION ANXIETY HITS THE FRONTIER LABS

If the revenue print was the bull case, the talent and commoditization story was the bear case in the same week. Alphabet had its worst day in over a year, sliding ~5%, after two marquee researchers departed — Gemini co-lead Noam Shazeer to OpenAI and Nobel laureate and AlphaFold co-creator John Jumper to Anthropic — and as Microsoft CEO Satya Nadella publicly argued the AI market is commoditizing and called for less dependence on 'AI Giants.' Alphabet has raised roughly \$141B in debt and equity since October to fund a vertically integrated stack; if models become cheaper and interchangeable, investors will ask whether that spend builds a moat or merely pressures margins.

The commoditization thread connects directly to the inference-cost data above and to the supply chain. AWS is raising prices on Nvidia-based EC2 Capacity Blocks by 20% effective July 1 — its second increase since January — explicitly citing supply and demand, even as open-weight routing accelerates. The simultaneous truth is higher compute prices at the cloud layer and falling effective prices per unit of intelligence at the model layer. Both point to the same conclusion: margin accrues to whoever controls the scarce physical inputs.

Read-Through: Model-layer commoditization pushes value toward orchestration, deployment and — most durably — the compute and power underneath.

Links: [CNBC — Alphabet worst day on AI talent exits](#) | [The Information — AWS raises Nvidia compute 20%](#)

DEVELOPMENT

THE NUCLEAR BID FOR AI POWER: \$17.5B FROM DOE, A CANADIAN RENAISSANCE

Policy on both sides of the border moved to put new baseload behind the demand curve. The U.S. Department of Energy announced up to \$17.5B in conditional loans to support development of 10 large reactors at five sites, advancing a Trump executive order targeting 300 GW of net-new nuclear capacity by 2050 and 10 reactors under construction by 2030. Westinghouse will partner with up to five utilities to procure long-lead AP1000 items at a fixed price, with each project jointly owned and \$1B of combined equity committed upfront before loan funds release. Capstone flagged Dominion, DTE, WEC, PSEG and Entergy as best positioned. Skeptics were quick: the Union of Concerned Scientists' Edwin Lyman called the sum a drop in the bucket against a potential ~\$200B total build cost, noting none of the projects is backed by an actual construction contract.

Canada matched the ambition. Energy Minister Tim Hodgson released a national strategy for a civilian nuclear renaissance — up to 10 new reactors over 15 years, construction starting on two large-scale units by 2035, a domestic microreactor finalized by 2035, and a doubling of uranium exports — framed as essential to doubling the country's grid by 2050. The plan, which officials say could cost more than \$100B, also positions CANDU reactor exports as a geopolitical lever and contemplates domestic uranium enrichment if export sales lag.

Read-Through: Nuclear is the clearest signal yet that policymakers now treat firm, dispatchable baseload as a strategic necessity for an AI-era grid.

Links: [Utility Dive — DOE \\$17.5B nuclear loans](#) | [CBC — Canada nuclear renaissance](#)

CBRE: POWER, NOT LAND, IS NOW THE BINDING CONSTRAINT ON DATA-CENTER SUPPLY

CBRE's Global Data Center Trends 2026 quantified the scarcity DPW has been tracking. Global inventory surged year-over-year in Q1 2026, yet demand drove vacancy to record lows — Northern Virginia at 0.3%, Atlanta at 1% — even as North American inventory across the top four markets grew 33%. Northern Virginia alone added 1,135.9 MW and set an all-time absorption record of 1,148.3 MW. The report's thesis is unambiguous: power availability and grid-infrastructure constraints, compounded by local opposition, are now the primary limiters on development timelines and site selection, and throttled construction will drive pricing to unprecedented highs through 2030.

Chicago power-delivery timelines via ComEd now extend to 2032 or later; Northern Virginia developers are pushing toward Richmond and adjacent PJM markets as 100-plus-acre contiguous sites disappear. Pricing reflects the squeeze — Chicago led North America with 14.7% rent growth — while emerging markets (West Texas, Tennessee) are winning specifically on land and power availability and flexible bring-your-own-power options. CBRE also notes rising demand for onsite generation, citing gas cogeneration for power resilience as a growing procurement category in supply-constrained hubs.

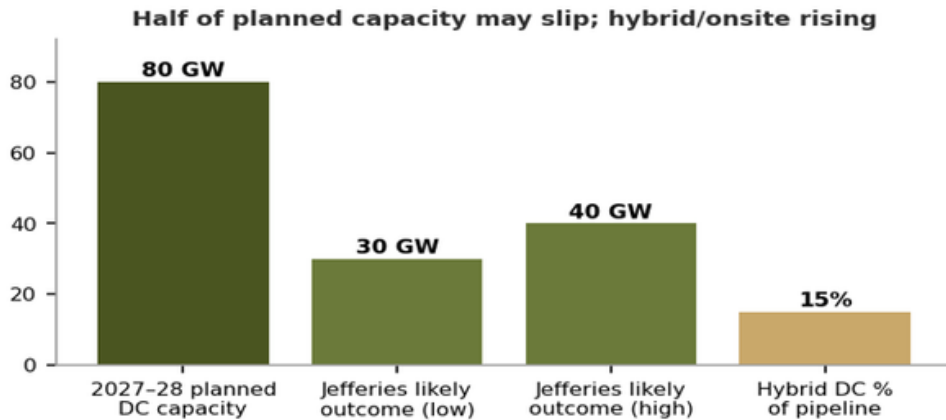
Read-Through: CBRE is third-party confirmation of the core Digital Power premise: the scarce input is firm power on a contiguous, controllable site, and the premium accrues to whoever holds it. Record-low vacancy with multi-year interconnection queues is the definition of a seller's market for energized capacity.

Links: [CBRE — Global Data Center Trends 2026](#)

ELECTRICITY MARKETS

THE DISTRIBUTED PIVOT: FROM MEGA CLUSTERS TO BACKYARDS

With perhaps half of data-center projects due online this year expected to slip — power access and local resistance the leading causes — the industry is hedging single-site risk by distributing it. Bloomberg's Liam Denning detailed the spectrum: Microsoft already runs a multi-site training network across Wisconsin and Georgia linked by high-bandwidth fiber, while hybrid data centers using onsite generation now represent ~15% of planned capacity per Jefferies, which expects only 30–40 GW of the 80+ GW planned for 2027–28 to actually arrive. At the micro end, startup Span.io proposes 'XFRA' backyard nodes — Nvidia chips plus cooling attached to homes — where 8,000 nodes equal a 100 MW data center and homeowners trade access for discounted power and internet.



Execution risk is pushing capacity toward distributed and hybrid/onsite architectures. Source: Jefferies; Epoch AI; Currence via Bloomberg.

Links: [Bloomberg — SpaceX bond losses](#) | [Bloomberg — MGX \\$50B](#) | [Bloomberg — AI Rout](#)

Denning notes the resulting backlash: new cost-causation measures across 16 states with rules pending in 10 more (CreditSights), a one-year New York construction moratorium awaiting the governor's signature, and FERC's move to speed connections in exchange for curtailment and cost-responsibility commitments.

Read-Through: External pickup of the DPW thesis — that uncontrolled commandeering of ratepayer-funded grid was the original sin — is now shaping mainstream coverage of the backlash. The strategic answer the market is converging toward (onsite generation, flexibility, paying your own way, local benefit) is the model. Distribution and hybridization expand, rather than threaten, the addressable market for controllable power.

Links: [Bloomberg — Are Backyard Data Centers an Answer?](#) | [Bloomberg — a16z / Ornn compute marketplace](#)

MICROSOFT AND CHEVRON PUT THE BTM GAS MODEL ON THE BOARD AT 2.67 GW

The single most validating development for the behind-the-meter thesis was a marquee operator and a supermajor executing exactly it. Chevron signed a 20-year deal to supply natural-gas-fired power to Microsoft's proposed West Texas data center near Pecos — Project Kilby — ramping to 2.67 GW, enough for more than 530,000 Texas homes, with first power targeted for 2028. The plant burns cheap Permian Basin associated gas — so abundant it is often flared as waste — feeding GE Vernova turbines, and crucially it generates its own power without drawing from the grid or a local utility. Chevron's New Energies president framed the design intent bluntly: it was built "to avoid" the ratepayer cost-growth backlash now spreading across the country.

The economics and supply-chain signals are directly readable into Keel's model. TD Securities pegged the outlay near \$9 billion at a ~15% developer IRR, implying roughly \$150/MWh to Microsoft; Chevron and partner Engine No. 1 have committed orders for seven scarce GE Vernova turbines. JPMorgan called the structure a cash-flow stream "uncorrelated to the underlying commodity." Texas now leads the nation with 33 GW of planned data-center power, and the BTM, self-generation, flared-gas-to-compute template Chevron is executing is the same brownfield logic DPW has tracked — now proven at supermajor scale with a 20-year contracted offtake.

SECOND-ORDER EFFECTS

Microsoft's strategy carries implications extending far beyond one project. Natural gas continues benefiting as the fastest dispatchable resource capable of supporting hyperscale timelines. Utilities increasingly transition from sole electricity providers toward partners in broader infrastructure ecosystems. Independent power producers become strategic infrastructure suppliers rather than commodity generators. Midstream companies gain exposure to sustained industrial gas demand rather than cyclical power markets. And equipment manufacturers benefit from higher demand for turbines, transformers, switchgear, substations and cooling equipment.

Read-Through: One transaction rarely changes an industry. One transaction can, however, reveal that the industry has already changed. Microsoft's 20-year agreement with Chevron-backed generation totaling roughly 2.67 GW is one of those moments.

Links: [Bloomberg — Microsoft/Chevron 20-Year Texas Power Deal](#)

THE \$850B LEASE BACKLOG: CONTRACTED DEMAND THAT MUST BE ENERGIZED

The scale of forward-committed demand reached a new marker: the largest cloud players now carry more than \$850 billion in future data-center lease commitments, climbing through the year per a Bloomberg analysis of filings. Meta led the quarter, adding \$79 billion in new commitments — a 76% jump — to \$182.9 billion, with Zuckerberg guiding to hundreds of billions in AI infrastructure by decade-end. Microsoft added \$41 billion to reach \$196.6 billion even while capacity-constrained, and Oracle still leads total future commitments on the back of its OpenAI contract. These obligations sit off-balance-sheet until payments begin and will largely pay out over the next two decades, pointing to sustained demand for the physical inputs underneath — semiconductors and, above all, energy.

Read-Through: \$850 billion dollars of contracted lease obligations is the demand certainty that underwrites every firm-power development decision. The constraint on converting that backlog into revenue is not capital or intent — it is energization. Each dollar of committed lease is a dollar searching for a site that can actually deliver power on a credible timeline.

Links: [Bloomberg — \\$850B Data-Center Lease Boom](#)

NVIDIA'S 45°C LIQUID COOLING RESHAPES THE SITE-SELECTION MATH

A quieter but structurally important shift came from the cooling stack. NVIDIA's Rubin generation is the first to run 100% liquid cooling — every chip and network component in a closed loop, no fans — with coolant entering the rack at up to 45°C and exiting near 55°C without performance loss. Because cooling has historically consumed up to 40% of data-center electricity, the efficiency leverage is large: in favorable climates the 45°C architecture enables chiller-less operation on dry coolers, cutting facility water from roughly 2.6 million gallons per MW per year toward zero and saving a 50-MW site over \$4 million annually. The data-center ambient temperature becomes flexible — warm summer air is fine — because nothing in the server depends on cool air.

Read-Through: Higher coolant temperatures widen the set of viable sites and sharpen the premium on power over climate. When a facility can reject heat with dry coolers and minimal water, the binding constraints collapse back toward the two things a power-first platform controls: firm power and a developable site. Lower water dependence also defuses one of the most potent local-opposition arguments — the water fight Healey explicitly cited — de-risking permitting. Efficiency gains at the rack do not shrink the power opportunity; they remove the non-power excuses for siting capacity where the megawatts are.

Links: [NVIDIA — 45°C Liquid Cooling Breakthrough](#)

ORBITAL DATA CENTERS: SEDUCTIVE SLIDEWARE, TERRESTRIAL REALITY

The most speculative escape hatch from the power-and-permitting squeeze got a sober reality check. The pitch for orbital data centers — abundant solar energy, no land, water or grid constraints, no community backlash — drew SpaceX (its IPO partly framed around a “Starmind” compute constellation) and others. But engineering professors writing in *The Conversation* catalogued why space is unforgiving: radiation damages electronics, heat rejection is surprisingly hard in vacuum, repairs are extraordinarily expensive, hardware needs refresh every three to five years, and the entire apparatus must be assembled in orbit amid mounting debris risk. Their conclusion is that the first viable orbital data centers will serve space-based customers — Earth-observation and satellite processing — long before they compete with mainstream cloud.

Read-Through: Orbital compute is a useful tell about how acute the terrestrial bottleneck has become: capital is seriously pricing low-Earth orbit as an alternative to permitting a substation. For a near-term brownfield developer this is reassurance, not threat. The latency-sensitive inference and training workloads driving today's demand cannot move to orbit this decade.

Links: [SingularityHub — Orbital Data Centers Face Daunting Challenges](#)

SPACE X TURNS COLOSSUS INTO A MERCHANT COMPUTE PLATFORM WITH \$6.3B REFLECTION DEAL

SpaceX continued to convert its captive AI infrastructure into a commercial product, signing open-source startup Reflection AI to a contract worth about \$6.3 billion through 2029. Reflection gets immediate access to Nvidia GB300s at the Colossus 2 site in Memphis and will pay roughly \$150 million per month beginning July 1, 2026, with either side able to terminate on 90 days' notice after the first three months. The deal follows compute agreements with Anthropic (~\$45B), Google (~\$30B) and Cursor, and lands as SpaceX leans on its AI-infrastructure narrative to support a debut investment-grade bond program. Built originally to train Grok, Colossus is now being sold as scarce GPU capacity to outside labs — positioning SpaceX alongside the cloud providers racing to monetize compute.

Read-Through: The economics underneath these mega-contracts are power. A \$6.3B compute deal is, at bottom, a bet that SpaceX can keep tens of thousands of GB300s energized and cooled at a single Memphis campus — the same firm-power-plus-site constraint every hyperscaler faces. The proliferation of short-tenor, terminable compute offtakes (90-day outs across Google, Anthropic and now Reflection) underscores that the durable, financeable layer is not the compute contract but the megawatts beneath it.

Links: [Bloomberg — SpaceX/Reflection Compute Deal](#) | [CNBC — SpaceX Signs Reflection Compute Deal](#)

THE SUPPLY MAP: 1,500-PLUS DATA CENTERS IN THE PIPELINE, RURAL SOUTH IN THE LEAD

A Pew Research analysis of Data Center Map quantified the forward pipeline: more than 1,500 U.S. data centers are in development, with the rural South leading the wave even as most existing capacity sits in urban markets. Virginia remains the anchor at 685 total facilities (287 planned), followed by Texas at 466 (170 planned), with Georgia (141 planned), Indiana (54 planned) and the Carolinas drawing a disproportionate share of the new build. Pennsylvania ranks among the larger markets at 129 total — 78 operating and 51 planned — placing Keel's home state squarely in the second tier of build-out activity behind the Virginia–Texas–Georgia triangle. Pew's sources peg potential U.S. data-center investment at up to \$7 trillion by 2030.

Read-Through: Pennsylvania's 51 planned facilities against 78 operating is a market mid-inflection — enough pipeline to confirm demand, not so saturated that energized sites lose their scarcity premium.

Links: [New York Times / Pew — Where Data Centers Are Planned](#)

THE \$17.5B NUCLEAR LOANS, REVISITED: A DOWN PAYMENT, NOT A SOLUTION

Bloomberg put the DOE's \$17.5 billion 'American nuclear supply chain loans' — financing long-lead equipment for five two-reactor projects under the Westinghouse AP1000 partnership — in sober perspective: it is a down payment, not a catalyst. A recent New York State study pegged a brownfield 1-GW reactor at nearly \$23 billion before credits, or ~\$19 billion after, implying ~\$40 billion for a two-unit site against which a \$3.5 billion federal loan is marginal. The study's \$183/MWh all-in cost is more than double prevailing around-the-clock futures. Hyperscalers, Denning notes, prize speed-to-power that nuclear cannot deliver, and have so far funded advanced-reactor startups only tentatively while building gas next to their data centers. Roughly 80% of utility investors, per Jefferies, expect a stock sell-off on any new nuclear announcement.

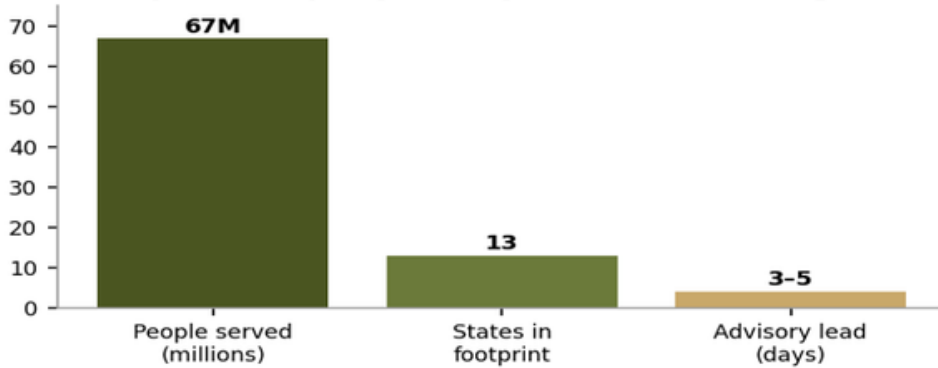
Links: [Bloomberg — Trump's \\$17.5B Nuclear Loans Are a Down Payment](#)

ELECTRICITY MARKETS

PJM ADDS A NEW EMERGENCY TIER AS AI STRETCHES THE GRID

The most operationally significant grid news came from PJM, which agreed to create a new 'capacity advisory' to warn its 67 million customers across 13 states when supplies could tighten even absent extreme weather. The trigger is structural: PJM says it has already reached capacity-emergency conditions without hitting traditional hot- or cold-weather thresholds — a direct consequence of data-center load growth and changing generation patterns. Advisories will be issued three to five days ahead, prompting plants to reschedule maintenance and signaling neighboring grids to expect fewer imports. PJM also secured enhanced backup-generator use in emergencies, pending DOE authorization.

PJM adds 'capacity advisory' as AI load stretches grid



Read-Through: PJM's home market is formalizing scarcity, and the cost is moving toward the customer. Scarcity formalized is optionality monetized for whoever already controls power.

Links: [Bloomberg — SpaceX bond losses](#) | [Bloomberg — MGX \\$50B](#) | [Bloomberg — AI Rout](#)

PJM's new advisory tier formalizes off-peak scarcity risk in the largest US grid. Source: Bloomberg.

The most operationally significant grid news came from PJM, which agreed to create a new 'capacity advisory' to warn its 67 million customers across 13 states when supplies could tighten even absent extreme weather. The trigger is structural: PJM says it has already reached capacity-emergency conditions without hitting traditional hot- or cold-weather thresholds — a direct consequence of data-center load growth and changing generation patterns. Advisories will be issued three to five days ahead, prompting plants to reschedule maintenance and signaling neighboring grids to expect fewer imports. PJM also secured enhanced backup-generator use in emergencies, pending DOE authorization.

NERC FLAGS SHRINKING DEPLOYABLE RESERVES AS COAL AND GAS FALTER

PJM's new advisory tier has a continent-wide backdrop. NERC's 2026 State of Reliability report found that the availability of conventional generation fell in 2025, driven largely by deteriorating coal and combined-cycle performance: unavailable energy rose 39.8 TWh for coal and 19.1 TWh for gas year-over-year. The annual weighted equivalent forced outage rate spiked to 9.2%, against a historical norm that rarely exceeded 8%, with coal and combined-cycle units named as the primary drivers. The result, in NERC's words, is shrinking deployable reserves and growing potential for tight operating conditions that can lead to system events.

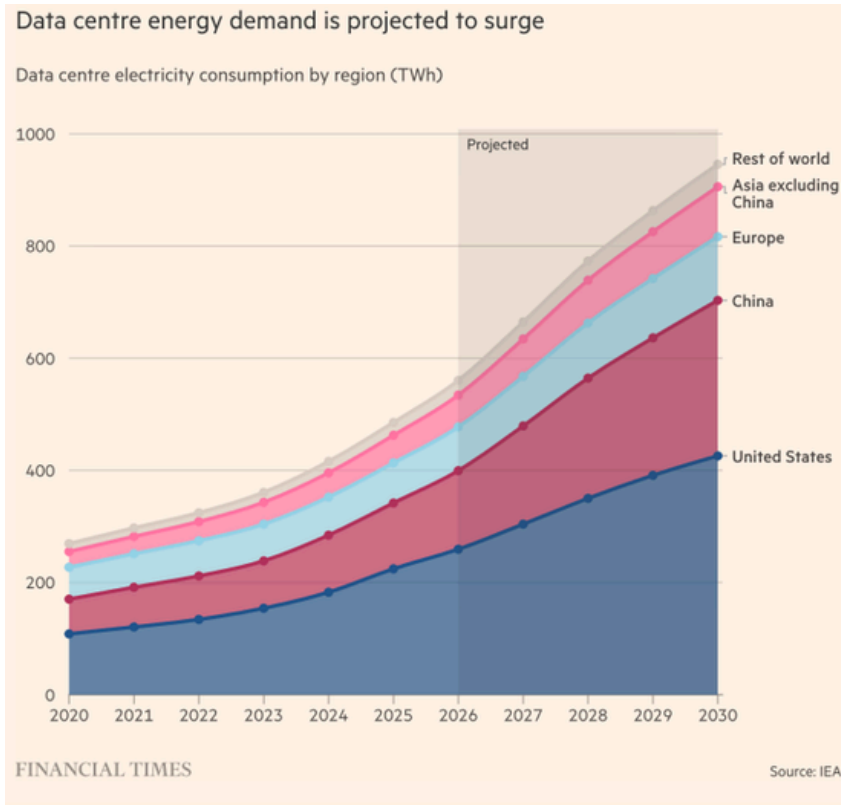
The cause is structural and self-reinforcing. Most large coal units are over 40 years old and were never designed for the daily cycling now required of them, which compounds the likelihood of failure as the fleet ages toward retirement. NERC also documented the load side of the same problem: in February 2025, a transmission fault triggered 1.8 GW of customer-initiated data-center load reduction in the Eastern Interconnection, with a separate event the same month cutting another 428 MW. NERC's warning is explicit — reliability challenges from large loads will intensify as these facilities proliferate, and planners must study how they behave during grid disturbances.

Read-Through: This is the supply-side mirror of the PJM advisory as a thinning, aging, harder-to-cycle conventional fleet plus demonstrated large-load tripping is precisely the reliability gap that a curtailable, power-integrated campus is built to close.

Links: [Utility Dive — NERC 2026 State of Reliability](#)

THE DEMAND CURVE, CHARTED: IEA SEES DATA-CENTER POWER USE NEARING 1,000 TWH BY 2030

A compelling picture of the AI power demand wave came from the IEA, via the Financial Times graph below: global data-center electricity consumption is projected to climb toward roughly 1,000 TWh by 2030, more than tripling from the 2020 base, with the United States and China the two largest and fastest-growing blocks. The U.S. line alone runs from roughly 110 TWh in 2020 to well above 400 TWh by 2030 — a trajectory that, layered onto an already-tightening PJM, is the macro backdrop for every interconnection queue, capacity advisory and ratepayer fight DPW tracks. The chart is the demand side of the entire thesis in one frame: the load is real, it is accelerating, and it is concentrated precisely where firm power is scarcest.



Data-center electricity consumption by region, 2020–2030 (TWh). Projected from 2026. Source: IEA via Financial Times.

Read-Through: A near-tripling of data-center load by 2030, concentrated in the U.S. and China, is the demand certainty that underwrites firm-power development

Links: [Financial Times / IEA — Data-Centre Energy Demand](#)

TVA WEIGHS UP TO 26 GW OF NEW GAS AS DATA-CENTER LOAD OUTRUNS ITS OWN FORECAST

The federal Tennessee Valley Authority released its preliminary 2026 integrated resource plan and said load growth in its footprint is already outpacing the reference-case forecast in its draft IRP — approaching the higher-growth scenario primarily because of data-center demand from AI and hyperscalers. To meet it, TVA is contemplating between 7 GW and 26 GW of new natural-gas generation between now and 2040, alongside up to 5 GW of nuclear, 1–5 GW of storage, 2–5 GW of renewables and 2–3 GW of efficiency and demand response. TVA framed gas as the source of firm, dispatchable capacity, with new nuclear positioned to reduce fuel-price and regulatory risk over a longer horizon.

The politics are as instructive as the megawatts. TVA's board, remade by Trump appointees, voted to keep the Cumberland and Kingston coal units running past their planned retirements, and the IRP now treats continued coal operation as an interim cost-and-reliability hedge. The utility set a 26% winter planning-reserve margin against an 18% summer target after a record winter peak of 35,319 MW in January 2025 — a reminder that the reliability problem is increasingly a firm-capacity problem, not an energy-volume one.

Read-Through: When the nation's largest federal utility responds to AI load by sizing up to 26 GW of new gas, it validates the core DPW premise that dispatchable thermal generation — not announced renewables or 2030s nuclear — is the bridge the demand curve actually requires.

Links: [Utility Dive — TVA Considers Up to 26 GW of Gas](#)

FINANCE

THE NEW AI VALUE STACK: PERFORMANCE PER WATT BECOMES THE BATTLEGROUND

The first phase of the AI revolution was won by companies capable of building the largest models. The second was won by companies capable of acquiring the largest GPU clusters. The third — now underway — will likely be won by companies capable of lowering the cost of every inference while extracting more performance from every megawatt. This week’s announcements from OpenAI, NVIDIA, Qualcomm, Micron and CATL collectively reinforce that AI competition is migrating away from headline model releases and toward optimization of the underlying infrastructure stack. The new battleground is no longer simply compute; it is performance per watt, cost per inference and speed-to-deployment.

Company	Development	Strategic Read-Through
OpenAI	Custom AI accelerator	Reduce inference costs and diversify beyond NVIDIA
NVIDIA	45°C liquid cooling	Lower cooling energy, higher rack density
Qualcomm	AI server CPUs	Broader compute ecosystem emerging
Micron	Raised AI outlook	Memory becomes strategic bottleneck
CATL	AI infrastructure batteries	Energy storage becomes critical infrastructure

Read-Through: The new AI value stack: the layers beneath the model are gaining economic weight.



Read Through: the layers beneath the AI model are increasing in economic importance. Value is migrating down the stack toward power, cooling, networking, memory and custom silicon.

OPENAI — CUSTOM SILICON IS ABOUT ECONOMICS, NOT INDEPENDENCE

Much of the initial market reaction centered on whether OpenAI was attempting to compete directly with NVIDIA. That interpretation misses the larger strategic objective. The company’s internally developed accelerator (“Jalapeño,” built with Broadcom) appears designed primarily to reduce inference costs and optimize workloads that no longer require general-purpose GPU flexibility. Reports indicate the new chip could lower inference costs by approximately 50% for targeted applications — one of the most economically meaningful announcements of the week.

Metric	Value
Manufacturing partner	Broadcom
Target	AI inference
Estimated cost reduction	~50%
Strategic goal	Lower cost per token
Primary benefit	Reduced infrastructure expense

Read-Through: Investors should avoid viewing every custom-chip announcement as an attempt to replace NVIDIA. Hyperscalers are increasingly segmenting workloads: training remains heavily GPU-centric; inference increasingly becomes application-specific. That distinction expands — not contracts — the total semiconductor opportunity, and pushes durable margin down the stack toward firm, low-cost power.





Chart 5 — AI silicon is fragmenting by workload.

FINANCE SNAPSHOT

- **\$50B** — MGX dedicated AI fund closed; targeting \$100B+ AUM at up to \$10B/yr deployment.
- **\$305M** — Paper losses on SpaceX’s \$25B IG deal vs. Treasuries within days of pricing.
- **\$270B** — Global leveraged-ETF AUM — the speculation machine now amplifying reversals.
- **\$33M** — Ornn seed (a16z-led) to build a GPU-capacity trading marketplace and index.
- **\$5.55B** — Cerebras IPO proceeds; shares since fell below the \$185 offer price on soft guide.

THE SPECULATION MACHINE WORKS IN REVERSE

Bloomberg’s anatomy of the week’s rout is helpful reading for anyone underwriting AI-linked capital. A retail-fueled unwind in AI chip stocks propagated through an ecosystem purpose-built to express the hottest trade with maximum leverage and minimum friction: leveraged ETFs, single-stock derivatives, digital-asset wrappers and prediction markets. Leveraged ETFs alone now hold more than \$270B globally (>\$200B US, >\$45B Asia), and Barclays estimates their rebalancing flows have surged to several times the long-run average — large enough to mechanically amplify both rallies and reversals.



The leveraged-ETF complex that magnified the AI rally is now magnifying the unwind. Source: Bloomberg.

The cautionary cases clustered: Korean 2x/3x funds down 20%+ in a week; SpaceX leveraged products off ~40% from their debut; Strategy's bullish and bearish wrappers down 90%+ since 2024 inception. As one strategist put it, these products bring 'a casino element' into the market — and the niche is growing. None of this is a referendum on AI fundamentals, but it is a direct warning about the fragility of the capital layers now funding physical infrastructure.

Read-Through: AI load must be framed as a system-benefit asset financed by durable structures, not as a beta expression dependent on continuously open, leveraged capital. Project finance, contracted offtake and investment-grade utility balance sheets are exactly the kind of low-beta capital that survives a week like this one; leveraged retail flows are not.

Links: [Bloomberg — AI Rout / \\$270B speculation machine](#) | [Bloomberg — Chinese hedge funds warn 'super bubble'](#)

GOLDMAN: ROTATE FROM CHIPS TOWARD HYPERSCALERS AS SEMIS STAY VOLATILE


Goldman Sachs strategist Christian Mueller-Glissmann argued that big-tech hyperscalers — Amazon, Oracle, Microsoft, Alphabet and Meta — could become the more attractive AI expression as chipmakers remain the volatile end of the spectrum. With the SOX semiconductor index up roughly 150% over the past year and heavy positioning through leveraged ETFs and options, he said it could make sense to diversify away from the chips and toward the platforms that have lagged on capex-spending concerns. He framed the broader tape as a Goldilocks spectrum — falling inflation expectations plus relatively good earnings — while cautioning that uniformly bullish sentiment raises the risk of sharp setbacks if the market doubts the drivers.

Read-Through: This is the same de-risking instinct that benefits power-backed infrastructure, one rung lower in the stack. Capital rotating from the most volatile, leverage-laden AI exposure toward steadier cash-generative platforms is the public-markets echo of debt and equity providers underwriting contracted load over narrative. The hyperscalers Goldman favors are the same counterparties signing long-dated power deals.

Links: [Bloomberg — Goldman sees bets in big tech](#)

BEARS GET LOUDER; A NEW ASSET CLASS GETS BUILT

Two of China's best-known hedge funds put numbers behind the caution. Wealspring (>\$1.4B AUM), whose founder called the 2007 top, labeled global AI stocks a 'super bubble' with a 'collapse point that may not be far away,' warning some hot names could fall 80%+. Shanghai Banxia (>\$294M) said the trigger 'has already appeared,' specifically predicting Anthropic's annualized revenue run-rate will undershoot expectations as large customers recoil at rising token costs. The backdrop: China's CSI AI Index is up 35%+ YTD and Korea's Kospi nearly 100%, even as violent pullbacks intensify.



Prominent China bears turn cautious as regional AI indices rip. Source: Bloomberg; CSC Financial.

Against that skepticism, capital is simultaneously institutionalizing the trade into a genuine asset class. Andreessen Horowitz led a \$33M seed into Ornn, a compute marketplace founded by two MIT graduates that first built a GPU-cost index now licensed by ICE for futures; CME is launching compute futures via Silicon Data. The ambition is a tradable market for compute capacity — short-term commitments for firms priced out of long-term GPU contracts. When an input gets an index, a futures curve and a spot venue, it has become an asset class, and power is the next input in that queue.

Read-Through: Compute is professionalizing into a hedgeable commodity, and power is the upstream input to that commodity. A brownfield platform that controls firm generation and interconnection sits at the base of this emerging stack — the equivalent of owning the wellhead as the paper market forms above it.

Links: [Bloomberg — Chinese hedge funds 'super bubble'](#) | [Bloomberg — a16z backs Ornn](#) | [Bloomberg — Cerebras below IPO price](#)

MICRON'S BLOWOUT GUIDE STEADIES A WOBBLING AI TAPE

Against the backdrop of the week's AI-equity wobble, Micron delivered the reassurance Wall Street was hunting for. The largest US memory maker guided fiscal-Q4 revenue to roughly \$50 billion against a \$43.2 billion consensus, with profit near \$31 a share versus \$25.31 expected; shares jumped as much as 20% to \$1,253 and have more than tripled this year. Q3 revenue reached \$41.5 billion and adjusted gross margin more than doubled to 84.9%. Critically for the durability of the cycle, Micron has locked 16 strategic customer agreements averaging three years, which Bloomberg Intelligence said should sustain pricing into 2027 and dampen the memory industry's historic boom-bust volatility. CEO Sanjay Mehrotra said there is "no line of sight" to supply catching demand before 2028, with HBM4 feeding NVIDIA's Vera Rubin platform.

Read-Through: Multi-year, contracted memory offtake that smooths a famously cyclical commodity is the same risk-transformation logic that underpins contracted power. Micron de-risked its cycle by converting spot exposure into three-year agreements with creditworthy buyers — precisely what a long-dated power offtake does for a generation asset. The signal that hardware demand has a multi-year floor extends the runway for the load that brownfield campuses serve, and reinforces that the scarce, contracted input — whether HBM or megawatts — commands the durable margin.

Links: [Bloomberg — Micron AI-Fueled Forecast](#)

OPENAI'S "JALAPENO" CHIP TARGETS THE INFERENCE COST CURVE

OpenAI unveiled its first custom accelerator, Jalapeno, co-designed with Broadcom and aimed squarely at cheaper inference — early samples show roughly 50% cost savings versus typical AI GPUs, per Broadcom CEO Hock Tan, with performance-per-watt “substantially better than current state-of-the-art.” Finalized chips integrate into Microsoft and partner data centers later this year, and Tan expects to exceed a prior 1.3-GW deployment forecast on strong demand. The strategic logic is vertical control over an inference-dominated cost structure; Tan predicts every frontier lab outside China will eventually build its own custom accelerator. Broadcom has stood up a chip financing vehicle with Apollo and Blackstone to fund the outlays.

Read-Through: Custom silicon optimized for cheaper inference accelerates the same commoditization dynamic DPW has flagged: as the cost per unit of intelligence falls and workloads route to the lowest-cost capacity, margin migrates down the stack toward power. A 50% inference-cost reduction does not shrink electricity demand — it shifts the bottleneck more firmly onto firm, low-cost megawatts and pulls value toward whoever supplies them. Efficiency at the chip layer is a tailwind for low-cost brownfield power, not a substitute for it.

Links: [Bloomberg — OpenAI/Broadcom Unveil Jalapeno](#)

THE AI IPO PIPELINE NEARS THE PUBLIC MARKET: ANTHROPIC, THEN OPENAI

The capital-formation endgame came into sharper focus. OpenAI is weighing an IPO as soon as 2027, likely after rival Anthropic, which is considering a listing as early as October; both have filed confidentially with the SEC. Recent tech-stock volatility has influenced OpenAI's timing — the company cautioned its debut “may be a while” — and Goldman Sachs and Morgan Stanley are advising. OpenAI raised \$122 billion earlier this year at an \$852 billion valuation; Anthropic's revenue has soared on coding products. News of the revised timeline knocked Goldman, Morgan Stanley, SoftBank and Oracle shares, underscoring how tightly the broader tape is now levered to the AI champions' path to public markets.

Read-Through: The march of the frontier labs toward public markets institutionalizes AI demand and lengthens its visible horizon — public companies must guide to multi-year capex, and that capex converts into the long-dated power offtake brownfield platforms exist to serve. But the same volatility that is reshaping IPO timing is the cautionary note DPW keeps returning to: equity-market sentiment is fickle, while contracted load on an investment-grade balance sheet is not. The lesson holds — underwrite cash conversion and tenant credit, not the IPO calendar.

Links: [Bloomberg — OpenAI Weighs 2027 IPO After Anthropic](#)

DEMAND-SIDE TELLS: OPTICAL INTERCONNECTS AND THE TECH-WEALTH FLYWHEEL

Two adjacent signals rounded out the demand picture. Musk cleared FTC review to acquire Mesh Optical Technologies, a SpaceX-alumni startup building optical transceivers that move data between data centers using light — more power-efficient and lower-latency than existing interconnects — with ambitions extending to satellite-to-satellite links for a future orbital compute system. Separately, Bloomberg documented how the AI build has become a macro force in its own right: New York City's tech workforce has grown 45% since the pandemic, median tech income reached \$130,000 in 2024, and brokers report tech and finance buyers now split the luxury market 50/50, with even private-equity wealth increasingly tied to OpenAI and Anthropic exposure.

Read-Through: Both items confirm the same thing from the edges: AI infrastructure is now a structural, economy-shaping demand driver, not a transient capex spike. Power-efficient interconnects extend the useful life and density of data-center investment — more compute per delivered megawatt — while the wealth flywheel signals that the capital underwriting the buildout is deep and self-reinforcing. For a firm-power developer, durable, broad-based demand is the backdrop that turns contracted offtake into a multi-decade annuity.

Links: [Bloomberg — Musk Cleared to Buy Mesh Optical](#) | [Bloomberg — NYC Tech Boom](#)

SPACEX'S \$25B INVESTMENT-GRADE DEBUT: A LEAP-OF-FAITH CREDIT FOR THE AI BUILD-OUT

SpaceX priced its first investment-grade bond — sized at \$20–25 billion across five tranches maturing 2031–2056 — drawing roughly \$30 billion of orders before the process formally opened and ranking among the largest deals of the year. Proceeds refinance a ~\$20 billion bridge loan and fund operations for a company that, by S&P's read, will burn cash through 2029. The longest tranche came around two percentage points over Treasuries. All three graders placed SpaceX in the BBB tier (Moody's Baa1, Fitch BBB+, S&P BBB), leaning on its dominant launch franchise, Starlink's recurring revenue and a \$100B-plus cash pile — even as analysts at Oppenheimer model more than \$400 billion of net debt by 2031 and S&P projects borrowings climbing to \$132 billion in 2028.

The skeptics were quick. Strategists noted the agencies extended unusual leeway for events still in the future, and the bonds promptly weakened in secondary trading — a large dealer quoted the 2056s as much as 0.28 point above their issue spread within days, while SpaceX shares slipped below their first trading price and the IPO erased more than \$600 billion of market value at one point. It is the same pattern DPW flagged on the deal’s pricing: abundant capital, but no longer indiscriminate, and least forgiving at the long end of a Musk-dependent credit.

Read-Through: SpaceX’s debut is the cleanest live test of how the bond market underwrites the AI build-out — and its wobble is the warning. A BBB credit burning cash through 2029, financing compute and energy ambitions on a leap of faith, is exactly the duration-heavy, proof-light exposure that reprices fastest when sentiment turns. The contrast underwrites the Keel model: contracted offtake to investment-grade tenants and hard-asset generation coverage is the low-beta structure that survives a week like this; narrative credit is not.

Links: [Bloomberg — SpaceX Readies \\$20B Bond Debut](#) | [Bloomberg — SpaceX IG Rating Brings Out Skeptics](#) | [The Information — SpaceX Debt Weakens](#)

QUALCOMM BIDS FOR THE DATA CENTER: \$15B AI-COMPONENT TARGET AND A META ANCHOR

Qualcomm used its investor day to reposition as an AI-infrastructure supplier, forecasting more than \$15 billion in annual data-center AI-component sales by fiscal 2029 — part of a \$40 billion non-handset revenue target that doubles its prior long-range plan. The company anchored the pitch with Meta, which agreed to deploy Qualcomm’s Dragonfly C1000 data-center processor (available 2028) and successor generations, and leaned on its low-power design heritage as a selling point for power-constrained AI facilities. Shares jumped as much as 11%. Qualcomm is entering a market dominated by Nvidia and contested by Broadcom and Intel, and says it will offer export-compliant versions to pursue China.

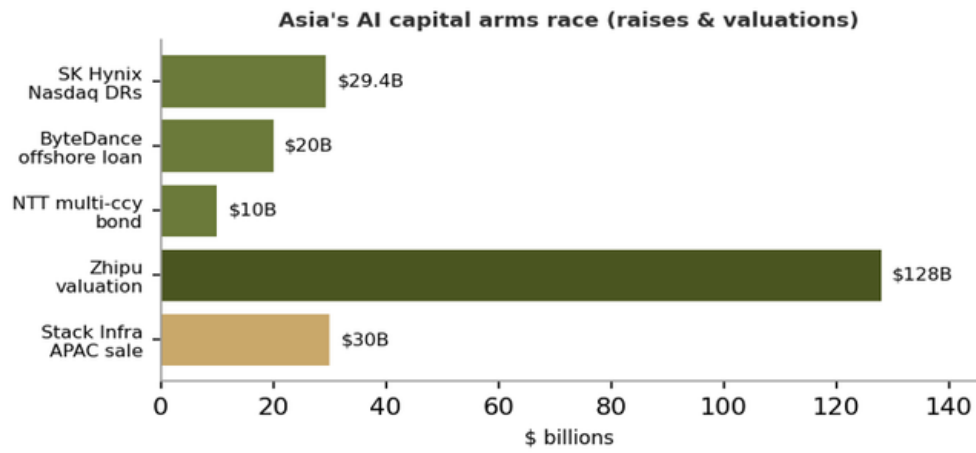
Read-Through: Qualcomm’s explicit framing — that expertise squeezing performance from phone-battery power budgets will win orders as AI data centers ‘push the limits of available electricity’ — is a chipmaker conceding that power is the binding constraint. Every credible entrant now sells against an energy ceiling. More competition lower in the stack commoditizes silicon and pushes durable margin toward the scarce physical inputs: firm, low-cost megawatts. A world of cheaper, more interchangeable accelerators routing to the lowest-cost capacity is constructive for brownfield power.

Links: [Bloomberg — Qualcomm Projects \\$15B Data-Center Chip Sales](#)

INTERNATIONAL

ASIA’S ARMS RACE MOVES DOWN THE STACK TO POWER AND CAPITAL

Bloomberg’s Asia briefing captured the same ‘down-the-stack’ migration playing out across the Pacific. Chinese model maker Zhipu is weighing a multibillion-dollar follow-on just months after its IPO, riding a ~2,000% surge to a \$128B+ valuation. SoftBank’s Masayoshi Son is seeking a stake in Tokyo Electric Power, explicitly betting that controlling electricity generation is now as strategic as securing chips — with data-center projects from Ohio to France to Japan. And interest in Stack Infrastructure’s APAC portfolio (BlackRock-backed AIP, Brookfield, KKR among suitors) could exceed \$30B.



Asian AI raises and valuations cited this week (\$B). Source: Bloomberg.

The financing cadence is relentless: SK Hynix targeting \$29.4B via Nasdaq depositary receipts (Seoul shares +300% YTD); ByteDance preparing a ~\$20B offshore loan against up to \$70B of AI spend this year; NTT readying a ~\$10B multi-currency bond, the largest by an Asian issuer this year. Son’s TEPCO logic is the single most important tell for Keel — the most sophisticated AI capital allocators in the world are now buying power generation directly, validating that energy control is the strategic chokepoint.

AUSTRALIA'S AEMO: DATA CENTERS TO NEAR 10% OF GRID DEMAND BY 2050

The down-the-stack migration is global. AEMO's biennial Integrated System Plan projects data centers will reach almost 10% of the National Electricity Market's underlying demand by 2050 — four times today's share — as overall consumption nearly doubles. The operator reaffirmed that the lowest-cost pathway is renewables firmed with storage and natural gas, even as households' grid-supplied needs fall on rooftop solar and batteries (small-scale battery capacity is set to jump to 35 GW by 2050 from 5 GW). Notably, AEMO credits a surge of data-center investment with propping up a slowing economy — the clearest signal yet that compute load is becoming a macro-relevant demand driver, not a niche one.

EUROPE SCALES BACK ITS AI DATA-CENTER AMBITIONS

In contrast, the EU downsized its flagship AI data-center tender, now seeking four facilities with at least 25,000 GPUs and three with at least 40,000 processors in phase one — down from an earlier concept of up to five centers at 100,000 advanced chips each. The bloc may contribute as much as €1B in phase one, with the private sector financing most of the cost and second-phase funding contingent on the post-2028 EU budget. The retreat is a setback in a race increasingly dominated by U.S. and Chinese players, though private capital is moving regardless — SoftBank plans up to €75B in French data centers, and Germany's Schwarz Group and Deutsche Telekom are pressing ahead.

Read-Through: Read together, AEMO and the EU tender bracket the same truth from opposite ends: where the buildout is constrained, the constraint is public capacity and capital, not demand. Australia is planning grid expansion explicitly around compute load; Europe is discovering that subsidies cannot substitute for financeable power and private balance sheets. Both reinforce that the durable advantage sits with platforms controlling firm power and capital.

Links: [Bloomberg — AI's New Arms Race \(Asia\)](#)

CATL'S BATTERY CHOKEHOLD: THE STORAGE LAYER UNDER THE AI GRID

The Information's profile of CATL chief Robin Zeng crystallized a chokepoint with direct grid implications. CATL makes about 40% of the world's lithium-ion batteries — EVs, robots, grid systems and, increasingly, AI data centers — against roughly 14% for BYD and under 2% for all US firms combined. Its profit rose 42% last year to \$10.5 billion, and it holds stakes across the metals and component supply chain. Most consequentially for power markets, CATL is a major supplier of stationary storage for data centers and the grid, a segment Zeng expects to become half its business: xAI has installed 2.3 GWh of CATL batteries at its Memphis Colossus complex — equivalent to 480,000 homes for four hours.

The strategic exposure is hard to hedge. The FT's companion analysis noted China dominates storage with 90% of solutions reliant on its LFP batteries, priced 30% below US and 35% below EU equivalents, even as the Pentagon has blacklisted CATL on alleged military links. For US developers leaning on battery storage to firm renewables or provide ride-through, the supply chain runs through a single, geopolitically fraught vendor — a dependency that makes dispatchable on-site generation a strategic, not merely economic, choice.

Read-Through: CATL's grip on stationary storage is a quiet argument for firm generation over storage-dependent architectures. If the battery layer that smooths intermittent supply is controlled by a single sanctioned foreign vendor with a 30%-plus cost advantage, a brownfield platform that delivers firm, dispatchable power on-site reduces exposure to that chokepoint rather than deepening it. Storage complements firm power; it does not replace it — and the supply-chain risk makes the controllable-generation hedge more valuable, not less.

Links: [The Information — CATL's Robin Zeng](#) | [FT — How to Power Data Centres](#)

POLICY & REGULATION

THE NEXT CONSTRAINT ON AI IS NO LONGER POWER. IT'S PERMISSION.

DPW INVESTMENT THESIS

For much of the AI buildout, investors focused on three constraints: GPUs, capital and electricity. A fourth is now emerging — and it may ultimately prove the most difficult to solve: political acceptance. This week's developments continue to suggest data-center development has entered a fundamentally different phase. Opposition is no longer confined to isolated community meetings or environmental organizations; it is beginning to influence elections, local zoning, utility regulation and federal litigation. For infrastructure investors, permitting risk is evolving into a project risk comparable to interconnection risk. That is a structural shift.

Metric	Value
U.S. data-center moratoriums / bans	300+
Additional measures under review	75+
Passed since January 2026	275+
Passed since early May	150+
Americans opposing local AI data centers	~70% (The Information / Gallup)

Chart 7 – Data-center political risk is accelerating. Source: The Information; DPW analysis.

Period	New Moratoriums / Bans
2023–2024	Limited
Early 2025	Moderate
Jan–Jun 2026	275+
Since early May 2026	150+

Chart 8 – Emerging AI infrastructure risk matrix. Source: DPW analysis.

The objections have become increasingly consistent across jurisdictions: residents fear hyperscale demand will raise retail electricity costs; cooling-water requirements remain controversial in water-stressed regions; industrial-scale campuses conflict with agricultural and residential land use; and communities question whether long-term benefits justify public subsidies granted before consultation. For years, Wall Street modeled data-center risk using three variables — power availability, fiber availability and capital availability. Beginning now, a fourth belongs in every underwriting model: the political probability of approval. The detailed cases below — a sitting Senate president ousted in Utah, more than 300 local bans mapped nationwide, and the federal government intervening over xAI’s unpermitted turbines — show how quickly that variable has hardened into a gating constraint.

SECOND-ORDER EFFECTS

Growing political resistance is likely to reshape where AI infrastructure is built. Developers may increasingly favor former industrial sites, existing power-plant locations, brownfield redevelopment, energy-intensive industrial corridors and regions with established manufacturing bases. Projects sited on greenfield agricultural land or near rapidly growing residential communities may face longer approval timelines and higher development costs — the precise spread a brownfield BTM platform is built to capture.

THE RATEPAYER BACKLASH HARDENS INTO LAW: MASSACHUSETTS AND LOUISIANA DRAW THE LINE

Growing political resistance is likely to reshape where AI infrastructure is built. Developers may increasingly favor former industrial sites, existing power-plant locations, brownfield redevelopment, energy-intensive industrial corridors and regions with established manufacturing bases. Projects sited on greenfield agricultural land or near rapidly growing residential communities may face longer approval timelines and higher development costs — the precise spread a brownfield BTM platform is built to capture.

The framing is now bipartisan and explicit. Landry's line — “whatever we do, we cannot lay the additional generation cost on the backs of the ratepayers” — is functionally identical to Healey's, and both name the East Coast cost-shift dynamic as the cautionary case. Massachusetts has never been a major data-center host; Louisiana has four campuses under construction and wants more. That two such different states converged on the same gating condition — demonstrated self-supply or full cost responsibility before incentives — confirms cost-causation has become the price of entry nationwide, not a regional quirk.

Links: [WBUR — Healey Halts Data Center Tax Incentives](#) | [Louisiana Radio Network — Landry Ratepayer Executive Order](#)

ORACLE SUES WISCONSIN OVER DATA-CENTER CREDIT RULES AS COLLATERAL BECOMES THE NEW BATTLEGROUND

The cost-allocation fight is migrating from who pays for generation to who guarantees it. Oracle sued the Wisconsin Public Service Commission in Ozaukee County Circuit Court, asking the court to set aside the credit-rating requirements embedded in We Energies' PSC-approved “very large customer” rate structure. That structure, approved in April, requires the utility to bill data-center customers exclusively for the new generation built to serve them, and forces developers rated below A- to post financial guarantees against cost-shift risk. Oracle carries a BBB rating — investment-grade, but a tier below the bar — the product of aggressive borrowing to fund its AI infrastructure. Under the current rule, its Port Washington subsidiary would post cash or letters of credit exceeding \$100 million per year to take service.

The dispute is a leading indicator for every developer's cost of capital. We Energies itself argued Oracle runs little default risk — “tens of billions” in value would have to evaporate before other customers saw losses, and generators built to serve data centers retain value even in bankruptcy. Both the utility and Oracle have asked the PSC for a stepped approach easing the burden on investment-grade BBB names. The precedent is already propagating: Xcel subsidiary Northern States Power has asked the PSC for its own very-large-customer structure setting the bar at BBB-, the lowest investment-grade rung. The credit threshold — A- versus BBB versus BBB- — will quietly determine which developers can afford to interconnect.

Read-Through: Credit-backed collateral requirements are the financial mirror of the physical cost-causation trend. As utilities nationwide adopt very-large-customer tariffs with credit floors, the advantage accrues to projects with bankable tenants and hard-asset coverage. Tenant credit quality is becoming as strategic as megawatts.

Links: [Wisconsin Watch — Oracle Sues PSC Over Credit Rules](#)

THE 202(C) PARADOX: DOE-ORDERED COAL PLANTS RUN AT A FRACTION OF CAPACITY

The administration's emergency-reliability intervention is producing reliability theater more than reliable megawatts. DOE used a string of 90-day Federal Power Act 202(c) orders to force 10 generating units at six plants — five coal-fired — to run past their retirement dates, the first such orders in the department's 50-year history. The output tells the real story. Five of the plants generated 1.5 million MWh in Q1 2026 under the orders, down 65% from 4.3 million MWh in the same period a year earlier. Two plants produced zero. TransAlta's 730-MW Centralia has not run at all; Tri-State's 446-MW Craig Unit 1 ran a single two-week stretch in April; NIPSCO's Schahfer units are offline for repairs into Q3. Constellation's oil-and-gas Eddystone units ran at a 0.5% capacity factor.

The mechanism is self-defeating by design. Owners planning retirement defer maintenance for years, so compelling these units to stay available surfaces deferred-capex bills rather than energy. CenterPoint estimates up to \$20.5 million and a 14-week outage to keep its 104-MW Culley Unit 2 compliant, and is fighting the order as forcing “substantial investment to support an inefficient and increasingly unreliable asset.” Tri-State and Platte River have asked the D.C. Circuit to overturn the Craig order outright. The episode is a live demonstration that aging thermal capacity cannot simply be commanded back into firm service.

Read-Through: This is the clearest possible evidence that the firm-capacity gap cannot be papered over with emergency orders on a geriatric fleet. Plants kept on life support deliver a fraction of nameplate and arrive with multimillion-dollar repair tabs and litigation. The reliability the grid actually needs comes from new, purpose-built, controllable capacity that can dispatch on command.

Links: [Utility Dive — DOE 202\(c\) Plants Underproducing](#)

THE BACKLASH GOES TO THE BALLOT BOX: UTAH SENATE PRESIDENT OUSTED OVER A DATA CENTER

The data-center backlash claimed its most senior political casualty yet. J. Stuart Adams — president of the Utah State Senate and one of the most powerful Republicans in a deep-red state — lost his primary to political newcomer Stephanie Hollist, who attacked the establishment's lack of transparency in approving the 40,000-acre Stratos project near the Great Salt Lake, backed by 'Shark Tank' investor Kevin O'Leary. Voters' objections clustered on energy use, water draw on a drought-stricken lake, and tax breaks granted over public objection. Box Elder County commissioners who backed the project also lost or trailed in their own primaries. The NYT cataloged parallel upsets in Alabama, Missouri and Wisconsin, where data-center votes have tipped local races across the political spectrum.

Read-Through: When a sitting Senate president loses a safe-seat primary over a single data center, cost-causation and local legitimacy stop being soft 'ESG' considerations and become hard gating risk for any speculative, grid-straining project. This is the political mirror of the ratepayer-protection laws DPW has tracked: the market is pricing local consent. A brownfield BTM platform that brings its own firm power, pays its own way and curtails on command is the model these revolts implicitly demand — the backlash is a moat for the prepared.

Links: [New York Times — Utah Senate President Loses Primary](#)

300-PLUS BANS AND MORATORIUMS: LOCAL CONTROL HARDENS INTO A NATIONAL CONSTRAINT

The Information mapped the scale of the revolt: more than 300 temporary or permanent local and state data-center bans enacted since 2023, with over 275 passed since January 1 and 150-plus since the start of May, plus 75-plus more under consideration. Resistance runs hottest in the Midwest and South, where the largest training clusters are sited; in Michigan, more than 20 towns within 50 miles of the OpenAI/Oracle Stargate complex have passed pauses, and over 50 governments statewide have bans in place. Most measures are temporary moratoriums of a month to a year, but a dozen New Jersey locales have gone open-ended, and states from Ohio to Illinois to Arizona have paused subsidies. A New York statewide one-year moratorium sits on the governor's desk. Seven in ten Americans oppose new data centers in their area, per Gallup.

Read-Through: A 300-plus-and-climbing thicket of local bans is the permitting-risk premium made explicit, and it cuts directly toward energized, accepted brownfield sites. Every moratorium that strands a greenfield proposal raises the scarcity value of a campus that already has site control, community standing and firm power. The constraint that makes a site bankable is no longer just interconnection — it is local consent, and the prepared brownfield platform holds both.

Links: [The Information — 300-Plus Bans and Moratoriums](#)

SHAPIRO SPLITS WITH HIS OWN PARTY ON PENNSYLVANIA DATA-CENTER POLICY

Keel's home state crystallized the national fight in a single week. Governor Josh Shapiro's marquee proposal would condition Pennsylvania's uncapped data-center sales-tax exemption on transparency, self-generation and environmental standards — incentive-led, and it passed the House 134–63. But many in his own party want a harder line: the House also passed, by a wider margin, Rep. Greg Vitali's bill to repeal the exemption outright (projected to cost the commonwealth \$2B-plus by mid-2031), plus a measure letting municipalities impose 180-day permitting pauses. Both tax bills now head to the Republican-controlled Senate, where leadership has signaled concerns but no clear preference. Only twelve projects currently claim the exemption, per the Department of Revenue.

Read-Through: Pennsylvania is now litigating, in real time, the exact question that determines Keel's operating environment: whether the state incentivizes good data-center behavior or simply withdraws support. Either path rewards the same profile. Shapiro's conditions — bring your own power, don't shift costs, meet environmental standards — describe the brownfield BTM model; Vitali's repeal removes a subsidy Keel's contracted, self-supplied economics never depended on. A self-construction-capable, power-integrated platform clears whichever way Harrisburg lands.

Links: [Spotlight PA — Shapiro Splits With Dems on Data Centers](#)

WASHINGTON PICKS A SIDE: DOJ DEFENDS XAI'S UNPERMITTED TURBINES ON NATIONAL-SECURITY GROUNDS

The federal government inserted itself directly into a data-center pollution fight, with implications well beyond one site. The DOJ moved to join an NAACP Clean Air Act suit over 57 gas turbines xAI and subsidiary MZX Tech installed in Southaven, Mississippi — without an air permit — to power the Colossus II complex serving Memphis. Citing national security, DOJ argued that where the government declines to enforce, citizens cannot sue over the resulting pollution, a position former EPA officials called unprecedented and a potential veto of the Clean Air Act's citizen-suit 'fail-safe.' The turbines could emit up to 5,300 tons per year of NOx — potentially the region's largest single source — and a January EPA rule clarified that 'temporary' turbines remain subject to controls.

Read-Through: The xAI case is the cautionary opposite of the brownfield BTM model: unpermitted, behind-schedule, self-powered capacity that invites litigation and federal-overreach controversy rather than local consent. It underscores why the durable path is the permitted, contracted, community-aligned one. A platform that secures air permits, pays its own way and runs firm generation by design is insulated from precisely the legal and reputational tail risk now attaching to commander-first, permit-later compute.

Links: [Bloomberg — Why DOJ Is Defending xAI's Data Centers](#)

WHAT WE'RE WATCHING

1. WILL ADDITIONAL HYPERSCALERS FOLLOW MICROSOFT'S DEDICATED-GENERATION STRATEGY?

Following Microsoft's Chevron transaction, investors should monitor whether Amazon, Google or Meta announce similar long-term generation partnerships. A broader shift would reinforce the thesis that hyperscalers are evolving into vertically integrated infrastructure companies.

2. WILL UTILITIES CONTINUE REVISING AI LOAD FORECASTS UPWARD?

TVA's revised planning assumptions is a continuation of several updates from major regulated utilities. Particular attention should be paid to:

- Dominion Energy
- Duke Energy
- AEP
- Southern Company
- FirstEnergy
- Entergy

3. DOES POLITICAL OPPOSITION CONTINUE ACCELERATING?

Community resistance has become one of the fastest-growing risks to AI infrastructure deployment. Watch for:

- additional moratoriums
- zoning disputes
- litigation
- state legislation
- upcoming election outcomes

4. CAN NUCLEAR MOVE FROM POLICY TO EXECUTION?

Federal support continues expanding. The next milestone is tangible construction activity rather than additional announcements. Investors should distinguish carefully between policy momentum and commercial deployment.

5. DOES AI FINANCING REMAIN DISCIPLINED?

Markets continue providing abundant capital. The next question is whether developers maintain return discipline as competition for power, labor and equipment intensifies.

FINAL WORD

The central lesson from this week's developments is straightforward. For nearly three years, artificial intelligence has been discussed primarily as a software revolution. It is increasingly becoming an infrastructure revolution.

The companies — and countries — that most effectively integrate electricity, generation, cooling, networking, semiconductors, permitting and capital into coherent operating platforms are likely to define the next decade of AI leadership.

In that sense, the competitive frontier has moved. The race is no longer simply to build the smartest model. The race is to build the most efficient AI factory.

APPENDIX A | MARKET DATA

Selected market, financing and infrastructure markers referenced in this issue. Figures are article-reported or DPW-derived from cited sources; not investment advice.

Metric	Latest / This Week	Prior / Reference	DPW Read
Gen-AI revenue	~\$110B trailing 12-mo; ~\$175B run-rate (ex-China)	First bottom-up, de-duplicated measure (Exponential View)	Demand is real and now partially proven; raises the execution bar.
AI sales vs. depreciation	~\$25B Q1 sales vs. ~\$21B depreciation	2nd straight quarter clearing the hurdle	Economics hold; depreciation still ~2/3 of revenue — thin buffer.
SpaceX bond deal	\$25B IG; ~\$305M paper loss; 2056 +0.28pp	~\$85–90B orders at pricing	Fast-money flip; long end of AI credit is fragile.
Nvidia bond deal	\$25B; 2046/2056 ~+8bp post-issue	Upsized from ~\$20B	High-grade AI credit still differentiates sharply by issuer.
MGX AI fund	~\$50B closed; \$100B+ AUM target	Up to \$10B/yr deployment	Capital abundant for contracted infrastructure exposure.
Leveraged-ETF AUM	\$270B global (>\$200B US, >\$45B Asia)	Rebalancing several x long-run avg (Barclays)	Mechanical flows amplify both rallies and reversals.
Cerebras outlook	2026 rev \$855–865M; below offer price	Q1 sales \$193.4M (+94%)	Even strong AI hardware must beat extreme expectations.
AWS compute pricing	Nvidia EC2 Capacity Blocks +20% (Jul 1)	2nd hike since January	Cloud compute scarce even as model prices fall.
OpenRouter frontier share	33% (Jun '26) vs. 72% (Jun '25)	Open-weight / DeepSeek routing rising	Inference commoditizing; demand more distributed.
PJM capacity advisory	New off-peak scarcity tier; 3–5 day lead	67M people, 13 states	Largest US grid formalizes AI-driven scarcity.
Ornn / compute mkt	\$33M seed (a16z); ICE/CME futures	GPU-cost index licensed by ICE	Compute professionalizing into an asset class.
Asia AI raises	SK Hynix \$29.4B; ByteDance ~\$20B; NTT ~\$10B	Zhipu ~\$128B val; Stack APAC >\$30B	Down-the-stack migration is global; power is the prize.
SoftBank / TEPCO	Pursuing utility stake for AI power	DC projects Ohio→France→Japan	Strategic buyers now acquiring generation directly.
NERC forced outage	9.2% EFOR (2025); coal +39.8 TWh / gas +19.1 TWh unavailable	Norm rarely exceeds 8%	Aging fleet + cycling = shrinking deployable reserves.
DOE nuclear loans	\$17.5B for 10 reactors at 5 sites	300 GW net-new by 2050 target	Firm baseload is a 2030s bridge; no signed builds yet.
Canada nuclear plan	Up to 10 reactors by 2040; >\$100B	Two large units by 2035	Strategic baseload bet; demand gap is now, not later.
CBRE Q1 vacancy	N. Virginia 0.3%; Atlanta 1.0%	Inventory still +33% YoY	Power, not land, is the binding supply constraint.
CBRE rent growth	Chicago +14.7% YoY (top NA market)	ComEd power timelines to 2032+	Scarcity = pricing to unprecedented highs through
Goldman rotation	Favor hyperscalers over chips	SOX +~150% over past year	Public-market echo of de-risking toward cash-generative AI.
AEMO ISP (Australia)	DCs ~10% of NEM demand by 2050	4x current share; grid ~doubles	Compute load is now a macro grid-planning driver.
EU AI DC tender	Scaled to 25,000–40,000-chip phase-1 sites	Was up to 5x 100,000 chips	Subsidy can't substitute for financeable power/capital.

APPENDIX B | RTO / FERC & MARKET CALENDAR

Near-term catalysts carried forward from the June 18 FERC large-load order package plus events surfaced in this issue. Confirm weekend/holiday dates against filing rules.

Date	Event / Filing	Agency / Party	Why It Matters
Jul. 1, 2026	AWS Nvidia EC2 Capacity Block price increase (+20%) effective	AWS	Cloud compute cost signal; pressures neocloud and tenant economics.
Jul. 17, 2026	Texas PUC/ERCOT MOU due under Abbott directive	Texas PUC / ERCOT	Large-load cost-responsibility and ratepayer-protection marker.
Jul. 18 / 20, 2026	30-day generation-adequacy reports due	FERC / six RTOs/ISOs	First concrete deliverable from June 18 large-load order package.
Jul. 23, 2026	FERC AD26-7-000 technical conference	FERC / PJM / stakeholders	Reliability backstop, resource adequacy, large-load cost allocation.
Aug. 17, 2026	60-day defend-or-revise tariff responses	FERC / six RTOs/ISOs	Markets justify or revise large-load tariff treatment.
Aug. 17, 2026	PJM E-2 co-location / BTMG further compliance filing	PJM / FERC	NITS, firm/non-firm contract demand, BTMG netting, curtailment.
Sep. 16, 2026	90-day policy milestone from June 18 order	FERC / RTOs / utilities	Benchmark for whether speed agenda becomes operational.
Oct. 2026	PJM expected to identify first 10 EIT projects	PJM	Signals which projects PJM views as near-term reliability-relevant.
Dec. 2026	PJM 2029/30 Base Residual Auction	PJM	Next capacity-price signal for DC-driven demand and storage/gas.
Q3-Q4 2026	SK Hynix Nasdaq DR listing (~Jul 10 trading)	SK Hynix	Bellwether AI-memory raise; gauge of AI capital appetite.
Ongoing	NY one-year data-center construction moratorium (pending signature)	NY Governor	Zoning/siting instability risk for the Northeast pipeline.

APPENDIX C | KEY LINKS

Resource	Live Link
Exponential View – State of the AI Economy	exponentialview.co/p/the-state-of-the-ai-econo...
Bloomberg – AI Sales Justify Data-Center Spending	bloomberg.com/news/articles/2026-06-25/ai-dema...
Bloomberg – MGX Raises ~\$50B for AI	bloomberg.com/news/articles/2026-06-23/abu-dha...
Bloomberg – PJM Capacity Advisory	bloomberg.com/news/articles/2026-06-24/us-larg...
Bloomberg – AI Rout / \$270B Speculation Machine	bloomberg.com/news/articles/2026-06-26/ai-rout...
Bloomberg – AI's New Arms Race (Asia)	bloomberg.com/news/newsletters/2026-06-24/ai-s...
CNBC – Alphabet Worst Day on AI Talent Exits	cnbc.com/2026/06/22/alphabet-goog-stock-ai-dep...
Bloomberg – a16z Backs Compute Marketplace Ornn	bloomberg.com/news/articles/2026-06-24/andrees...
Bloomberg – Are Backyard Data Centers an Answer?	bloomberg.com/opinion/articles/2026-06-22/are-...
The Information – AWS Raises Nvidia Compute 20%	theinformation.com/briefings/aws-raises-prices...
Bloomberg – Bond Traders Stunned by SpaceX Losses	bloomberg.com/news/articles/2026-06-26/bond-tr...
Bloomberg – Cerebras Falls Below IPO Price	bloomberg.com/news/articles/2026-06-23/cerebra...
Bloomberg – Chinese Hedge Funds Warn 'Super Bubble'	bloomberg.com/news/articles/2026-06-26/chinese...
Bloomberg – Data Center Boom Raises Muni Credit Risk	bloomberg.com/news/articles/2026-06-25/data-ce...
Utility Dive – NERC 2026 State of Reliability	utilitydive.com/news/deployable-reserves-shrin...
Utility Dive – DOE \$17.5B Nuclear Loans	utilitydive.com/news/doe-offers-175b-loans-hel...
CBC – Canada Nuclear Renaissance Strategy	cbc.ca/news/politics/federal-nuclear-strategy-...
CBRE – Global Data Center Trends 2026	cbre.com/insights/reports/global-data-center-t...
Bloomberg – Goldman Sees Bets in Big Tech	bloomberg.com/news/articles/2026-06-26/goldman...
Bloomberg – Data Centers to Double Australia Demand	bloomberg.com/news/articles/2026-06-24/data-ce...
Bloomberg – Europe Scales Down AI Data Center Tender	bloomberg.com/news/articles/2026-06-18/europe-...

APPENDIX D | ARTICLE ROLLUP / SOURCE LIST

Section	Headline / Theme	Primary Source(s)	Date
Market Spotlight	Bottom-up AI revenue clears depreciation hurdle	Exponential View; Bloomberg	Jun 2025
Top Stories	SpaceX bond losses; MGX raise; capital bifurcates	Bloomberg	Jun 23–26
Top Stories	Alphabet talent exits; commoditization; AWS +20%	CNBC; The Information	Jun 22–27
Finance	\$270B leveraged-ETF speculation machine reverses	Bloomberg	Jun 2026
Finance	China bears warn 'super bubble'; Ormn builds asset class	Bloomberg	Jun 24–26
Finance	Cerebras falls below IPO price on soft guide	Bloomberg	Jun 2023
Electricity Mkts	PJM capacity advisory; Moody's muni credit risk	Bloomberg	Jun 24–25
Development	Distributed / backyard compute pivot (Span; Microsoft)	Bloomberg (Denning)	Jun 2022
International	Asia arms race: Zhipu, SoftBank/TEPCO, SK Hynix, NTT	Bloomberg	Jun 2024
Finance	Goldman: rotate from chips toward hyperscalers	Bloomberg	Jun 2026
Electricity Mkts	NERC: deployable reserves shrink as coal/gas falter	Utility Dive (NERC)	Jun 2025
Development	DOE \$17.5B loans for 10 large reactors	Utility Dive	Jun 2024
Development	Canada plans up to 10 reactors by 2040	CBC News	Jun 2022
Development	CBRE: power is the binding data-center constraint	CBRE	Jun 2017
International	AEMO ISP: data centers ~10% of Australia demand by 2050	Bloomberg	Jun 2024
International	EU scales down AI data-center tender	Bloomberg	Jun 2018
Policy & Reg.	Healey freezes MA data-center tax incentives	WBUR	Jun 2026
Policy & Reg.	Landry signs LA ratepayer-protection order	Louisiana Radio Network	Jun 2026
Policy & Reg.	Oracle sues Wisconsin PSC over credit rules	Wisconsin Watch	Jun 2023
Policy & Reg.	DOE 202(c) plants run far below capacity	Utility Dive	Jun 2023

APPENDIX D | ARTICLE ROLLUP / SOURCE LIST

Section	Headline / Theme	Primary Source(s)	Date
Development	Microsoft/Chevron 20-yr, 2.67 GW BTM gas deal	Bloomberg	Jun 2022
Development	\$850B future data-center lease backlog	Bloomberg	Jun 2024
Development	NVIDIA 45°C liquid cooling; near-zero water	NVIDIA	Jun 2021
Development	Orbital data centers face daunting reality	SingularityHub	Jun 2026
Finance	Micron blowout AI guide; ~\$50B Q4 revenue	Bloomberg	Jun 2024
Finance	OpenAI/Broadcom “Jalapeno” inference chip	Bloomberg	Jun 2024
Finance	OpenAI weighs 2027 IPO after Anthropic	Bloomberg	Jun 2026
Finance	Musk/Mesh optical; NYC tech-wealth flywheel	Bloomberg	Jun 2026
International	CATL’s battery chokehold on AI storage	The Information; FT	Jun 2026
International	FT: power, not compute, is the binding race	FT	Jun 2024
Development	SpaceX/Reflection \$6.3B Colossus compute deal	Bloomberg; CNBC	Jun 2022
Development	Pew: 1,500+ US data centers in pipeline	NYT; Pew Research	Jun 2025
Development	Denning: \$17.5B nuclear loans a down payment	Bloomberg	Jun 2025
Electricity Mkts	IEA: data-center power use nears 1,000 TWh by 2030	FT; IEA	Jun 2026
Electricity Mkts	TVA weighs up to 26 GW of new gas for AI load	Utility Dive	Jun 2023
Finance	SpaceX \$25B IG bond debut; skeptics emerge	Bloomberg; The Information	Jun 22-26
Finance	Qualcomm targets \$15B data-center chips by 2029	Bloomberg	Jun 2024
Policy & Reg.	Utah Senate president ousted over data center	NYT	Jun 2024
Policy & Reg.	300+ data-center bans and moratoriums mapped	The Information	Jun 2024
Policy & Reg.	Shapiro splits with PA Dems on data-center policy	Spotlight PA	Jun 2025
Policy & Reg.	DOJ defends xAI unpermitted turbines in court	Bloomberg	Jun 2025