



DIGITAL POWER WEEKLY

May 18 – May 24, 2026

Editor: Chris Ruppel, SVP - Power

ISSUE NO. 3

KEEL
INFRASTRUCTURE

66 GW

U.S. Data Center Power Demand
by 2027 (Goldman Sachs)

14.9 GW

PJM Backstop Procurement
Targeted Large-Load Supply

\$30B+

Anthropic Funding Round
Reported Target

\$1.75+T

SpaceX Target Val.

EXECUTIVE SUMMARY

Cogito, Ergo IPO

THE CENTRAL THESIS

This week delivered two datapoints that could appear unrelated but are not. First, the frontier AI compute stack entered capital-markets territory. SpaceX filed its S-1 with the Anthropic lease description putting a precise number on what energized compute is currently worth. Second, PJM began rewriting the grid rulebook that will determine how much of this compute gets energized — accelerating its Reliability Backstop Procurement, aligning it with the proposed Connect & Manage structure, and telling the market that scarce power capacity will be rationed, regulated, and repriced until new generation and transmission is funded (via long-term contracts) and built.

These are not separate stories. They are the same trade, seen from opposite ends. The S-1 is a prospectus for the physical stack. The RBP acceleration is the regulatory infrastructure that will separate the financeable physical stack from the speculative. And the Anthropic lease — \$1.25B per month for compute that barely exists at scale — is the price signal that connects them: proof that the scarcity premium for energized frontier compute is real, large, and already being paid by the people who need it most. **The key data points this week reduce to the same sentence: power is no longer a cost line. It is the denominator of the valuation.**

(Hat tip to Alex Wissner-Gross' The Innermost Loop for the 'a la carte' Descartes headline.)

I. THE IPO MOMENT

The SpaceX S-1 is on file. But the capital-markets event is already larger than that. OpenAI and Anthropic are generating equivalent pressure through private fund raising at public-market valuations, reporting revenue trajectories that make an IPO timeline feel inevitable, and compute commitments the market is already treating as price signals. Last week's ERock and Innio announcements extend the same logic into pure power infrastructure. The formal IPO count is one. The shadow inventory behind it is many.

The 1980s Irish band The Waterboys may have named the AI IPO moment we're living through before it arrived. The title track of *This Is the Sea* (1985) opens with the line that fits: "Once you were tethered / Now you are free / Well that was the river / This is the sea." Frontier AI has become that 'sea' — an experiment that once had boundaries now simultaneously reshapes grids, rewrites capital structures, and redirects industrial policy. The question is whether the power gets there to meet the moment — and leads directly to the sequel.

After the sweeping ambition of *This Is the Sea*, band leader Mike Scott did something that surprised everyone who expected him to chase the grandeur further. He went to rural Ireland, put down the layered studio production, picked up a fiddle and a bodhrán, and went back to the roots — traditional Irish music, working-class textures, songs about rivers and land and labor that were honest precisely because they were embedded in the local. *Fisherman's Blues* (1988) was not a lesser album. It was a realer one. The grand declaration had been made. Now came the work.

AI is following a similar trajectory. Its first phase was about modeling the 'sea': training runs, benchmark scores, foundation model races, the business of parsing and pricing the opportunity. The next phase is the 'fisherman's' — and like Scott, it requires going back to basics. Not the vision, but the vessel. Not the horizon, but the haul. The fisherman doesn't choose his weather. He works the tides, reads conditions that shift without warning, and depends entirely on one thing being true before any of it matters: he needs a seaworthy boat. No boat, no fishing. AI infrastructure is that boat—and building it is unglamorous, unrelenting hard work and the only path from a declaration to energization.

As SpaceX's S-1 made crystal clear, the winners in this next phase of AI will be the 'boatbuilders' — the draftsmen and craftsmen of the power systems, the transmission rights, the interconnections, the fiber, the chips, and the capital structures that actually deliver the "physical stack" the market is already pricing. Grand visions don't catch fish. Fishermen in boats do. And every boat needs a Keel.

SpaceX / SPCX

SpaceX has filed an S-1 to list on Nasdaq (ticker: SPCX); Reuters/CNBC reported Goldman Sachs as lead-left, with Morgan Stanley, BofA, Citi, and JPMorgan in the broader bank group. Reported IPO target valuation: ~\$1.75T, up from a combined SpaceX+xAI private valuation of ~\$1.25T after the February merger. The S-1 frames AI competitive position around **'control of the physical stack'** — chip manufacturing, data-center infrastructure, and power generation — and treats power as a first-order variable in cost per token. COLOSSUS and COLOSSUS II are presented not as ordinary colo facilities but as gigawatt-class AI infrastructure where powered shells, BTM generation, Tesla Megapacks, and equipment finance determine deployment speed.

Anthropic

Reported ~\$30B+ financing round at about a \$900B pre-money valuation, with Sequoia, Dragoneer, Altimeter, and Greenoaks each committing ~\$2B. Q1 revenue: \$4.8B. Projected Q2: \$10.9B. First operating profit (~\$559M) anticipated — though compute spending is rising and may compress margins. The reported xAI/SpaceX Colossus compute arrangement (\$1.25B/month through May 2029, roughly \$40+B across the stated period, subject to ramp discounts and 90 day termination rights) is the week's most important private-market price signal: a frontier lab reportedly paying a competitor-controlled infrastructure platform because delivered compute is simply scarce.

OpenAI

While OpenAI's own IPO timeline remains fluid, the competitive and capital-market pressure created by SpaceX and Anthropic's moves accelerates the clock. The Google/Blackstone AI Cloud JV (500 MW, 2027 COD, \$5B equity from Blackstone, Google TPUs) offers a parallel template: hyperscaler technology plus private infrastructure equity inside a financeable, standalone AI cloud vehicle.

II. THE PJM SIGNAL: RBP + C&M IS THE GRID'S COROLLARY TO THE IPO WAVE

The same week that frontier AI accelerates going public on the premise of physical-stack control, PJM is writing the market rules that will determine how that physical stack gets financed. The board's decision to accelerate the Reliability Backstop Procurement to September 2026 — compressing what was a March 2027 timeline by six months — and to align it with the large-load cost-responsibility framework embedded in Connect & Manage is not administrative housekeeping. It is the construction of a capacity-obligation regime for loads that cannot demonstrate qualifying supply. That is a structural change to the economics of data center power procurement, and markets began pricing it in immediately. Constellation and Vistra closed roughly +6% on Wednesday following the announcement. NRG was closer to +4–5%. The read-through is clean: tighter capacity rules, accelerated procurement timelines, and a mandatory backstop auction create a floor under PJM-centric IPP earnings that the market has been discounting for two years.

The Commercial Logic in Three Steps:

- **Bring Your Own New Generation/qualifying bilateral capacity (BYONG/BYONC):** If a data center can demonstrate verified qualifying capacity through the RBP/bilateral framework, it may reduce or eliminate the curtailment exposure that would otherwise apply to new large loads post-summer 2026 (under the C&M). The key structural point is **"new" not yet energized capacity** appears to be required.
- **Power Centrality:** That converts power procurement from a cost line into a core siting and tenant-risk variable — the same transformation the SpaceX S-1 describes, now written into RTO rule.
- **Accelerated Timeline:** The September procurement window is no longer an administrative date. It is a commercial catalyst that could materially affect which post-2026 loads are financeable, tenantable, and public-market ready.

The emergency context amplifies the signal. During the recent May heat wave, PJM sought and received DOE authority under Federal Power Act Section 202(c) to direct data centers and other large facilities with backup generation to run those resources/reduce grid service as a last-resort reliability measure before rolling blackouts. *Bloomberg* reported real-time PJM prices briefly exceeding \$2,000/MWh.

READ-THROUGH

PJM is creating the first large-scale U.S. market template for data-center capacity responsibility. For Digital Power developers, the RBP/C&M alignment is bullish for incumbent and developable generation, mixed for utilities without cost-allocation clarity, and negative for speculative data-center pipelines that lack long-term site control and power execution capabilities. The September PJM procurement initiation is the IPO market's grid counterpart: a single date that begins to provide clarity on platform valuations and project economics.

III. THE REST OF THE ISSUE

1. Goldman Sachs Demand Forecast: 66 GW by 2027

Goldman projects U.S. data-center power demand rising from 31 GW (2025) to 66 GW (2027), implying data centers move from 4.1% to 8.5% of U.S. peak summer demand. Only 50–60% of scheduled capacity is expected to come online on time. The most important developer distinction is now scheduled versus deliverable — the same binary that PJM's RBP/C&M framework institutionalizes.

2. Hyperscaler CDS: AI Borrowing Hits Bank Balance Sheets

Hyperscalers have borrowed more than \$250B globally for AI, with \$5T of projected spend through 2030. Banks are buying CDS protection to stay within single-name exposure limits. Monthly hyperscaler CDS notional volumes are up tenfold since the start of 2025. Further financializing AI compute, CME and Silicon Data are preparing compute futures. The asset class is broadening from AI equities to credit, commodities, ETFs, and futures — and the common denominator remains physical scarcity.

3. Modular Infrastructure: Armada, Samsung, SpaceX Solar

Armada's Series B raised \$230M at a \$2B pre-money — backed by Johnson Controls, BlackRock, and Mitsui. It is a manufacturing bet: the capital is funding a 400,000 sq ft factory in Arizona purpose-built to produce modular, deployable data center infrastructure at industrial scale. When Johnson Controls, one of the world's largest building-systems integrators, and Mitsui, a global commodities and infrastructure conglomerate, are anchoring a Series B alongside BlackRock, the signal is that AI infrastructure is being treated as a durable manufactured-goods category, not a bespoke construction sector. That is an OEM transition, and the Arizona factory is the physical evidence of it.

Samsung's floating data center concept — 50 MW, ABS-certified, paired with a Mousterian development partnership and a Samsung/OpenAI letter of intent — and SpaceX's planned 10 GW solar-cell facility near Austin are additive to the same thesis from different vectors. None of these are the same product. All of them are the same answer: when interconnection queues run five to seven years and grid access is the binding constraint, the market diversifies form factors rather than waiting in line. Modular ground deployments, offshore platforms, and captive generation at the site of the load are all attempts to route around the bottleneck. The common thread is not the technology. It is the power constraint they are solving for.

4. E3 Rate Study + Local Backlash: The Siting Fight Is Getting Quantifiable

E3 came out with a report *Understanding the Drivers of Rising Electricity Rates and the Role of Data Centers* (sponsored by the Data Center Coalition) that brings discipline to a messy public debate: it separates general rate pressure from true cost shifting, and correctly argues that higher customer bills are driven by many factors beyond data centers — inflation, fuel volatility, wildfire/resilience spending, grid modernization, capacity-market design, generation retirements, and transmission/distribution investment. Its best evidence is that the historical quantitative literature does not yet show systematic subsidization of data centers by other customers. It also gives a useful toolkit for managing future risk: minimum bills, exit fees, collateral, long-term contracts, interruptibility, behind-the-meter generation, demand response, and regular cost-of-service updates. The report's most constructive datapoint is that regulators are moving quickly: E3 identifies at least 38 new large-load tariffs between 2018 and 2026, with 30 implemented in 2025–2026 alone.

The report's weakness is that it is better as a defense of historical rate treatment than as a full underwriting of the next wave of multi-GW AI load clusters. The report itself admits the evidence base is limited, often backward-looking, and not always causal. Its PJM analysis cuts both ways: E3 estimates that load growth, primarily data centers, drove roughly 50% of the increase between PJM's 2024/2025 and 2025/2026 capacity auctions, with other market-design and supply factors driving the rest.

Cogito, ergo IPO. But only if you control the infrastructure.

| Theme | Signal | Direction | Key Catalyst |
|-------------------------------------|------------------------------|-----------|--------------------------------|
| IPP/Incumbent Generation | PJM RBP acceleration | ▲ Bullish | Sept. procurement window |
| Powered AI Campus (firm power) | SpaceX S-1 + Anthropic lease | ▲ Bullish | Watch ramp schedule |
| Utilities (cost-allocation unclear) | E3 + PJM cost shift risk | ► Mixed | Tariff design / regulatory lag |
| Turbine / Blade Supply Chain | Hot section scarcity report | ▲ Bullish | Operating assets at premium |
| Compute ETFs / Futures | CME GPU futures filing | ► Watch | Regulatory review pending |

★ MARKET SPOTLIGHT | FEATURED ANALYSES

SPACE X S-1 + ANTHROPIC LEASE PRICING - THE PUBLIC MARKET GETS A PHYSICAL-STACK PROSPECTUS


Sources: SpaceX S-1 Digital Power Analysis; CNBC SpaceX IPO live updates; The Verge/Axios Anthropic-SpaceX reporting; Tom Brown X post; Bloomberg xAI feature.

SpaceX's S-1 is the most-defined public-market articulation yet of the Digital Power thesis. The filing frames AI competitive position around "control of the physical stack" - chip manufacturing, data-center infrastructure and power generation - and treats power as a first-order variable in cost per token. COLOSSUS and COLOSSUS II are described not as ordinary colocation facilities but as gigawatt-class AI infrastructure where powered shells, behind-the-meter generation, Tesla Megapacks, cooling, utility coordination and equipment finance determine deployment speed.

The Anthropic Cloud Services Agreement gives the market a pricing datapoint. The reported run-rate is \$1.25B per month from May 2026 through May 2029, or roughly \$15B per year and \$40+B across the stated period, adjusting for reduced fees during the May-June ramp and a 90-day termination right. That last feature matters: this is a powerful short-term capacity-price signal, not a classic non-cancellable project-finance offtake. The \$/MW economics depend entirely on the denominator. At 300 MW, the contract implies ~\$5,708/MWh-equivalent; at 500 MW, ~\$3,425/MWh; at 700 MW, ~\$2,446/MWh; and at 900 MW, ~\$1,903/MWh. This is illustrative \$/MWh-equivalent sensitivity, not contract pricing; denominator depends on actual compute/power capacity made available under the Colossus arrangement.

The strategic irony is part of the story. Claude and Grok compete; Musk has publicly criticized Anthropic; yet Anthropic is still reportedly leasing capacity from a Musk-controlled infrastructure platform because delivered compute is scarce. *Bloomberg* Opinion framed the deal as evidence that software scalability has collided with physical infrastructure

Anthropic / SpaceX Pricing Sensitivity



Source: Keel analysis based on \$15B/year lease run-rate.

Anthropic / SpaceX lease pricing sensitivity. Labels are placed inside the bars to avoid bleed-over; colors distinguish capacity cases.

| Metric | Disclosure/datapoint | Digital Power read-through |
|-------------------------------|--|---|
| COLOSSUS + COLOSSUS II | ~1.0 GW of compute power in the S-1 memo | Gigawatt-scale AI load is now an SEC-disclosed strategic asset. |
| Deployment speed | 122 days for first COLOSSUS cluster; 91 and 64 days for COLOSSUS II clusters | Brownfield speed-to-power becomes a public-market KPI. |
| AI capex | \$12.7B in 2025; \$7.7B in Q1 2026 | AI is a balance-sheet infrastructure business, not software-light growth. |
| Mobile turbines | ~\$2.0B acquisition announced Apr. 30, 2026 | Mobile generation becomes strategic AI compute equipment. |
| BESS / grid | Megapacks, utility-funded grid capacity and self-built generation | Canonical architecture: grid-tethered BTM power plus batteries. |

Read-Through: SpaceX converts the Digital Power argument from private-market narrative into prospectus language. The practical takeaway is terrestrial: gigawatt AI clusters are being built around controlled power, brownfield shells, BTM gas, batteries, cooling and structured equipment finance. Anthropic's lease price shows what delivered compute can command when capacity is scarce.

PJM RBP + CONNECT & MANAGE - LARGE-LOAD CAPACITY PROCUREMENT BECOMES A PROPOSED CONTRACTING REGIME


Sources: Utility Dive; Bloomberg; NERC; PJM Merges RBP + C&M memo; DOE emergency order; FERC Commissioner Chang interview; Bloomberg market reaction.

PJM is no longer treating large-load growth as a normal load-forecasting issue. The grid operator accelerated its planned RBP auction to September, rather than March, to address data-center demand and near-term reliability risk. The same package sits beside the Connect & Manage framework, which effectively conditions post-summer-2026 large-load firmness on whether the load brings or pays for qualifying capacity.

The emergency context matters. During May heat, PJM sought and received DOE authority to direct data centers and other large loads with backup generation to run those resources/reduce grid service as a last-resort reliability measure before rolling blackouts. *Bloomberg* reported real-time prices across parts of PJM briefly exceeding \$2,000/MWh and PJM issuing maximum-generation and load-management alerts. *Utility Dive* reported that PJM expected more than 40 GW of planned generation outages (seasonal maintenance) and less than 5,800 MW of reserves during the stressed period.

NERC simultaneously warned that data-center interconnection delays complicate demand forecasting. Its summer assessment indicated that slower-than-anticipated large-load hookups may provide short-term relief, but not a structural solution. The reliability problem is therefore two-sided: too much uncertain future load for planning purposes, but still enough near-term heat/weather exposure to require emergency tools.

PJM Wholesale Cost Signal



Source: Monitoring Analytics / Bloomberg reporting.

PJM Q1 wholesale power costs moved from an RTO technical datapoint to a retail-political signal.

The equity market treated PJM's acceleration as constructive for power producers. *Bloomberg* reported gains in Constellation, Vistra and NRG after PJM pulled forward the procurement timeline. Evercore described the change as "materially constructive" for independent power producers because it acknowledges structural capacity shortness tied to hyperscaler demand. KeyBanc framed the change as reducing uncertainty extending into 2027 and potentially restarting data-center-related power contracting.

The RBP and Connect & Manage interaction is the critical commercial point. If a data center can demonstrate verified qualifying new capacity through the backstop/bilateral framework, it may reduce or avoid the curtailment exposure that would otherwise apply to new large loads. That turns power procurement from a cost line into a core siting and tenant-risk variable. It also raises the value of existing IPP development teams, project uprates, repowers, deactivated assets with CIR value, surplus interconnection rights and shovel-ready generation options.

DEVELOPMENT

GOLDMAN SACHS POWER-DEMAND FORECAST - DATA-CENTER LOAD DOUBLES BY 2027


Sources: Goldman Sachs, US Data Center Power Demand Projected to Double by 2027; Apollo/Daily Spark; America data-center count; WMB earnings presentation.

Goldman Sachs Commodities Research projects U.S. data-center power demand rising from 31 GW in 2025 to 41 GW in 2026 and 66 GW in 2027 — moving data centers from roughly 4% of U.S. peak summer demand today to 8.5% by 2027. Goldman puts total U.S. data-center capacity at approximately 95 GW by year-end 2027, using a 70% utilization assumption.

The trajectory is not a straight line. Goldman expects only 50–60% of capacity scheduled for delivery over the next one to two years to actually come online on time, with the remainder slipping due to interconnection delays, equipment lead times, and permitting. Regional divergence is sharp: Mid-Atlantic, Mid-Continent, and Northwest markets face elevated reliability risk — the same footprint where PJM's RBP and C&M framework are actively being redesigned. Texas and Georgia are better positioned, supported by planned generation additions that the constrained markets cannot currently match.

Apollo's recent *Spark Note* has data that makes the geographic shift concrete: the U.S. currently has roughly 4,000 operating data centers and nearly 3,000 under construction — a pipeline approaching the size of the existing base. That construction backlog is no longer concentrated in Northern Virginia. Missouri, Utah, Texas, Pennsylvania, Alberta, France, and multiple European FLAP-D markets are now competing simultaneously for the same constrained inputs: large power blocks, land, water rights, fiber, and political tolerance for the load. The bottleneck has not been solved. It has been distributed.

U.S. Data Center Power Demand Forecast



Source: Goldman Sachs Commodities Research.

Goldman Sachs forecast: U.S. data-center power demand more than doubles from 2025 to 2027.

Read-Through: The most important development distinction is now scheduled versus deliverable capacity. A 2027 COD claim is only valuable if the project has long-term site control, equipment, permits and interconnection path. Expect capital providers to demand a harder taxonomy: contracted power, connected power, active power and tenant-revenue power.

GOOGLE/BLACKSTONE AI CLOUD JV + MISSOURI INVESTMENT - PLATFORM BUILDOUT MOVES FROM CLOUD SKU TO POWER CAMPUS

Sources: WSJ; CIO Dive; Barron's; OilPrice; Google Missouri article.

Google and Blackstone launched an example of the new capital structure: hyperscaler chips plus private-capital balance sheet plus dedicated data-center capacity. The U.S.-based venture is expected to receive \$5B of equity from Blackstone, use Google Cloud TPUs, and target 500 MW of capacity online in 2027. Benjamin Treynor Sloss, a long-time Google engineering executive, is expected to lead the company.

The strategic logic is twofold. First, Google is trying to monetize its TPU stack outside its own cloud and create a non-NVIDIA compute ecosystem at scale. Second, Blackstone is converting AI demand into a hard-asset platform where power, real estate, chips and customer contracts sit inside a financeable vehicle. *Barron's* framed the move as an attempt to challenge Nvidia-linked neocloud capacity; *WSJ* emphasized that the venture would be the largest attempt yet by Google to sell and monetize its own chips to external parties.

Google's separate announced \$15B Missouri data-center investment adds the project-level layer. The planned infrastructure in Montgomery County/New Florence is one of the largest technology investments in Missouri history and reinforces the emergence of secondary markets that can offer land, utility cooperation and political support. These markets are not "cheap land" stories; they are power-delivery stories.

Read-Through: *The Google/Blackstone structure is another template for AI cloud project finance: hyperscaler technology, private infrastructure equity, dedicated capacity and a power-centric COD timeline. If it works, expect more chip/platform owners to separate compute infrastructure into financeable vehicles rather than relying only on traditional hyperscaler capex.*

ARMADA/JCI, SAMSUNG FLOATING DATA CENTERS AND SPACEX SOLAR - MODULAR PHYSICAL INFRASTRUCTURE BECOMES THE ANSWER TO QUEUE FRICTION

Sources: Armada release; LinkedIn; Data Center Richness; Bloomberg SpaceX solar article.

Armada's Galleon Forge One announcement with Johnson Controls pushes the modular thesis into manufacturing scale. The planned Arizona factory spans up to 400,000 square feet, targets production beginning this summer, and is expected to create more than 500 jobs. Armada's Series B raised \$230M at a \$2B pre-money valuation, with strategic participation from Johnson Controls, BlackRock, Mitsui and others. The product focus is megawatt-scale modular data centers for high-density training, inference, sovereign neocloud and multi-tenant compute.

Samsung's early-stage floating data-center concept/MOU points to a different form of modularity: shipbuilding and marine infrastructure as a way to accelerate deployment near ports, industrial power and water access. The strategic question is not whether floating data centers dominate the market; it is whether data-center form factors begin to diversify because land-based permitting and utility interconnection become too slow.

SpaceX's planned 10 GW solar-cell manufacturing facility near Austin is another physical-layer response. Bloomberg reported permit filings for a Bastrop facility with two floors each capable of making 5 GW of solar cells. The project supports Starlink and the broader Musk AI infrastructure ambition, including concepts for space-based or highly distributed power systems. It also shows vertical integration moving beyond chips and software into energy components.

Read-Through: *AI infrastructure is beginning to industrialize like an OEM sector. The winners will not just find sites; they will control repeatable designs, supply chains, thermal systems, power modules, manufacturing partners and deployment playbooks. Modularization is one of the market's responses to interconnection delay and construction uncertainty.*

FINANCE

ANTHROPIC FINANCING AND PROFITABILITY - FRONTIER LABS BECOME SOVEREIGN-SCALE CAPITAL CONSUMERS


Sources: Bloomberg Anthropic funding; WSJ Anthropic profitability; Bloomberg secondary-market article; Anthropic-backed Fractional AI acquisition.

Anthropic generated the week's most important private-market finance signal. Bloomberg reported the company could close a new round of more than \$30B at about a \$900B pre-money valuation, with Sequoia, Dragoneer, Altimeter and Greenoaks expected to co-lead and each planning roughly \$2B commitments. Existing investors including Founders Fund and General Catalyst were also reported as participants. The round reportedly came together in weeks, underscoring demand for direct exposure to frontier AI.

Anthropic is projected to more than double revenue to \$10.9B in the June quarter after generating \$4.8B in Q1, and to record its first operating profit at roughly \$559M. The caveat is critical: profitability may not persist for the full year because compute spending is rising. Still, the print directly challenges the simplistic view that frontier AI is only a cash-burn story.

The capital-markets ecosystem is adjusting around Anthropic. *Bloomberg* reported the company cracked down on unauthorized secondary sales and SPV exposure, creating turmoil among pre-IPO investors and funds that claimed Anthropic exposure. Separately, the Anthropic/Blackstone/Hellman & Friedman enterprise-services venture acquired Fractional AI as its operating centerpiece, adding a distribution channel into PE-backed midsize companies.

Anthropic Funding / Revenue Scale



Source: WSJ / Bloomberg reporting.

Anthropic's reported funding round and revenue trajectory are now infrastructure-scale datapoints.

Read-Through: *The financing constraint on AI compute is shifting from customer demand to deliverable supply. Anthropic's revenue trajectory — scaling toward \$10B+ quarterly — and its \$30B+ funding base mean the lab can underwrite gigawatt-scale capacity commitments in advance of the workload that fills them. The Colossus lease is a template: frontier labs are becoming the demand signal that power markets have been waiting for, and they are writing checks large enough to move project finance.*

HYPERSCALER DEBT AND CDS - AI CAPEX IS HITTING BANK BALANCE SHEETS


Sources: *Bloomberg, Hyperscaler Debt Flood Brings Derivatives Bonanza; Bloomberg AI data-center debt-risk coverage.*

Bloomberg reported this week that hyperscalers have already borrowed more than \$250 billion globally to fund AI, with the total buildout projected to reach \$5 trillion through 2030. The scale of that borrowing is now creating a structural problem for the banks financing it: single-name concentration limits across loan books and derivatives desks are binding, forcing banks to buy credit protection in order to stay in business with their largest clients.

- **Rapid Growth:** The CDS market that barely existed for these names eighteen months ago is now one of the most actively traded corners of credit derivatives. CDS tied to Microsoft, Amazon, and Oracle generated \$4.6 billion in DTCC-reported notional volume in Q1 2026, up from \$759 million a year earlier. Meta CDS — only launched last October — hit \$534 million in the same period, more than double the prior quarter. Monthly hyperscaler CDS notional volumes at BofA are up tenfold since the start of 2025, according to Matt Mandell, BofA's head of U.S. single-name CDS.
- **Pricing Anomaly:** The pricing anomaly is visible and worth noting. As of the *Bloomberg* report, Meta's five-year CDS traded around 73 basis points annually — meaning a seller of protection on \$10 million of notional collects roughly \$73,000 per year — despite Meta carrying an AA-/Aa3 rating. That compares with approximately 52 basis points, or \$52,000 per year on \$10 million, for five-year protection on the broader North American investment-grade index, which averages roughly BBB+. A single-name AA credit is trading wider than a diversified BBB+ index. That inversion is not a credit call on Meta. It is a liquidity and concentration premium — the market's price for absorbing bank demand at scale.

The five names in the JPMorgan CDS basket are Alphabet, Amazon, Meta, Microsoft, and Oracle. Nvidia does not appear in the DTCC volume data or the basket. The mechanism driving the surge is bank CVA desks buying protection to free up credit lines, not a directional view on any individual name's solvency.

Hyperscaler CDS Trading Volumes



Source: Bloomberg / DTCC.

CDS volumes show hyperscaler AI borrowing moving from bond-market story to derivatives-market capacity issue.

ELECTRICITY MARKETS

E3 RATE STUDY - HISTORY SAYS NO SYSTEMATIC SUBSIDY, BUT THE FORWARD RISK IS REAL

Sources: E3 full report; E3 reaction memo; Governor Stein/Duke article; ASU heat study; UC Berkeley/Haas working paper.

The E3 whitepaper *Understanding the Drivers of Rising Electricity Rates and the Role of Data Centers* — commissioned by the Data Center Coalition, published May 18, 2026 — is useful for framing data center growth in a regulatory context. The Data Center Coalition represents the hyperscalers whose cost-allocation treatment is the central policy question. That does not invalidate the analysis, but it should inform how the findings are read.

The headline conclusion is carefully bounded: E3 reviewed 11 quantitative studies and found no historical evidence that data centers have subsidized other customers under existing rate structures. Retail rates, E3 argues, reflect a compound of generation costs, transmission, distribution, fuel volatility, inflation, wildfire mitigation, grid modernization, and wholesale market design — not load growth alone. The state-level data makes the point sharply: Texas and Virginia absorbed the largest data-center load increases in the country and saw among the smallest retail rate increases. California and New York saw the largest rate increases despite load reductions. The simple narrative — data centers drive up everyone's bills — does not survive contact with the historical data.

The PJM finding is more nuanced and more directly relevant to this week's RBP news. E3's analysis of the capacity price increase between the 2024/2025 and 2025/2026 auctions found load growth drove roughly half the increase. Market design changes and generation retirements drove the other half. That is not exoneration; it is a 50/50 split, and it means the RBP debate is as much about structural market design as it is about data center demand.

E3 explicitly acknowledges that the historical record does not predict the forward path. Rapid load growth at current velocity, tighter generation supply, accelerated transmission and distribution needs, and stressed regulatory and permitting processes all create rate-pressure risk that the backward-looking studies cannot capture. The report identifies 38 new large-load tariffs established between 2018 and 2026 — 30 of them implemented in 2025 and 2026 alone — as evidence that utilities and regulators are already reaching for cost-responsibility tools: tailored large-load tariffs, minimum bill requirements, contribution-in-aid-of-construction charges, exit fees, long-term contracts, demand flexibility requirements, bring-your-own-generation frameworks, and accelerated cost-of-service updates.

The central point of the E3 report is that the historical subsidy case is unproven, the forward-looking risk is real and acknowledged by E3 itself, and the policy toolkit for assigning cost responsibility is being deployed at a pace — 30 new tariffs in two years — that suggests utilities are not waiting for the academic debate to

settle. The RBP accelerating to September and the DOE 202(c) emergency order issued the same week as this report are the operating environment the E3 findings have to survive in.

Read-Through: *E3 gives both sides ammunition. Historically, data centers have not clearly shifted costs to other customers. Prospectively, the next wave is different because scale, equipment inflation, PJM capacity prices and grid capex are different.*

RTOS

NERC AND SUMMER RELIABILITY - SLOW HOOKUPS ARE RELIEF, NOT A SOLUTION

Sources: Utility Dive NERC; Bloomberg slow hookups; East Coast power surge; PJM emergency approval.

NERC's summer assessment and related reporting created a nuanced reliability message. The U.S. grid has added more than 58 GW of summer resource capacity since last year, with resource additions reported across solar, storage and gas, improving reserve positions in many regions; the solar/storage/gas mix should be treated as availability-adjusted summary data rather than nameplate capacity. NERC flagged four areas at risk of energy shortfalls or local constraints under extreme heat, down from six in 2025.

However, demand forecasting is increasingly difficult because data-center interconnection timelines are uncertain. *Bloomberg* reported that key U.S. grids have reduced summer demand forecasts to reflect slower-than-anticipated data-center and large-load connections. PJM has 26% extra reserves for this summer, in part because its demand forecast was sharply reduced after data centers were not built or connected as fast as expected. ERCOT's net internal demand is expected to fall by 3.7 GW year over year because more data centers can now be curtailed, largely tied to Texas Senate Bill 6.

The short-term benefit is real: slower hookups and curtailment tools can reduce summer blackout risk. The structural problem is also real: projects delayed in 2026 do not disappear; they reappear in later forecasts, often with less confidence around timing, location and load shape.

Read-Through: *The grid has bought time, not solved the load-growth problem. NERC's data support a key theme: interconnection delay can look like reliability relief in the near term, but it worsens planning uncertainty. The market will value loads that are flexible, self-supplied or paired with new capacity because they are easier to model and easier to defend.*

FERC AND PJM GOVERNANCE - MARKET DESIGN IS NOW PART OF THE AI RACE

Sources: FERC Commissioner Chang Utility Dive; PJM cost allocation article; PJM emergency order; Bloomberg PJM market reaction.

FERC Commissioner Judy Chang told *Utility Dive* she is "not thinking about" breaking up PJM, offering a more constructive reform frame than the "too big to function" rhetoric. Her key point is that PJM's market may not have been designed for this moment and must evolve to handle large-load growth, supply-chain constraints, interconnection delays and inadequate transmission.

PJM's board, in accelerating the backstop auction, urged states to immediately develop regulatory and tariff frameworks to shield residential and other existing customers from data-center-driven costs. That is the practical governance issue: PJM can procure capacity, but states and utilities must decide how costs are allocated to new large loads versus legacy customers (this addresses Mills' 'Credibility Trap').

Read-Through: *FERC, DOE, PJM and state commissions are converging on the same framework: large loads must be studied, priced and curtailed differently if they create reliability or cost-allocation stress. The "AI race" is now a grid governance race as much as a chip race.*

OEMS

BLOOM ENERGY, NEBIUS AND ORACLE - FUEL CELLS BECOME PRIME-POWER DEBATE, NOT BACKUP EQUIPMENT

Sources: CNBC Bloom/Nebius; Bloom market-read memo.

Bloom Energy has had three distinct data points over the last few weeks that need to be read together rather than as a single story.

- **Nebius-Bloom deal:** Nebius Group — the Amsterdam-headquartered, Nasdaq-listed AI cloud infrastructure company — signed a master fuel-cell capacity agreement with Bloom Energy, effective May 14, 2026, for up to \$2.6B of long-term service fees across phased 10-year supply terms. Bloom will install, operate and maintain behind-the-meter fuel-cell systems for Nebius's U.S. data-center buildout. The first project targets 328 MW of installed capacity operational this year, reportedly equivalent to roughly 250 MW of guaranteed output, and replaces previously planned reciprocating-engine/gas-turbine power at the initial site. The structure is designed to accelerate time-to-power by reducing dependence on new transmission and conventional grid interconnection. NBIS shares rose roughly 15–17% on the announcement; BE also moved sharply.
- **Oracle/Project Jupiter validation:** The more consequential Bloom announcement for infrastructure investors came on April 27, when Oracle, BorderPlex Digital Assets and Bloom confirmed that Project Jupiter — a planned AI data-center campus in Doña Ana County, New Mexico — would use up to 2.45 GW of installed Bloom fuel-cell capacity in an on-site microgrid architecture. The design replaces previously planned gas turbines and diesel generators and builds on Oracle's expanded master services agreement with Bloom for up to 2.8 GW of fuel-cell systems. The campus is expected to be consolidated into a single microgrid campus, with Bloom's closed-loop design positioned as reducing water use and local air-emissions impacts versus the prior gas-turbine/diesel configuration. The strategic importance is clear: Oracle is not treating fuel cells as backup power. It is using Bloom as a prime-power architecture for a multi-gigawatt AI campus. No one has built or operated anything at that scale, under those reliability constraints, using this technology. The question of whether it works is genuinely open. If it does, Bloom's competitive position in the AI infrastructure power stack becomes very difficult to challenge. If it doesn't, the concentration risk stops being a footnote.

Q1 report: Three items in Bloom's Q1 2026 10-Q warrant attention alongside the headline contract wins. First, total accrued warranty and product-performance liabilities rose from \$20.0M at year-end 2025 to \$38.4M at March 31, 2026, driven by a sharp increase in product-warranty liabilities tied to identified product issues — a near-doubling of the liability in one quarter. Second, customer concentration remains significant and should remain a diligence focus as Bloom transitions toward larger AI/data-center customers. Third, on stack longevity, Bloom's 2019 public response to Hindenburg stated that 2016–2018 vintage fuel cells had an average life of 4.8–5.2 years, while Bloom argued subsequent generations and service economics had improved. The DPW memo's "3–5 year" shorthand should therefore be treated as a diligence question, not a settled current-generation performance claim.

The financial picture is genuinely improving. Q1 2026 revenue of \$751.1M was up 130.4% year over year — Bloom's first quarter of greater-than-100% growth as a public company. Full-year 2026 guidance was raised to \$3.4B–\$3.8B of revenue, \$600M–\$750M of non-GAAP operating income, and operating cash flow above \$200M. Bloom also issued Oracle a warrant with a grant-date fair value of \$261.3M, creating a non-cash drag on future revenue recognition as obligations are satisfied.

Read-Through: *Bloom Energy has closed the two largest contracts in its history within roughly a month of each other — the Oracle Project Jupiter deal announced April 27 and the Nebius agreement announced May 21 — and its Q1 financials reflect a company operating at a different scale than it was a year ago. Revenue of \$751 million in Q1 2026 against \$326 million in Q1 2025 is not an incremental improvement; it is a step change, and the raised full-year guidance to \$3.4–\$3.8 billion confirms management sees the demand as durable. However, the warranty reserve and customer concentration risks are a real caveat.*

NVIDIA DSX AND AI FACTORY SIMULATION - THE OEM ECOSYSTEM BECOMES COORDINATED INFRASTRUCTURE

Sources: *Data Center Richness NVIDIA DSX email; NVIDIA ecosystem references.*

NVIDIA's Omniverse DSX and DSX Air announcements show AI data-center design moving toward digital-twin industrialization. The platform simulates AI factory infrastructure, including power, liquid cooling, networking, software, compute, physical layout and operational coordination. *Data Center Richness* reported that CoreWeave, IREN and Hut 8 have committed to using DSX technology for multiple gigawatts of capacity, while IREN separately committed to build up to 5 GW of DSX-ready AI infrastructure.

The ecosystem is the point. Vertiv has created a DSX design for the Vera Rubin GPU platform using its OneCore design; Switch has integrated DSX into its LDC EVO operating system; and companies including Cadence, Dassault Systemes, Eaton, Jacobs, Schneider Electric, Siemens and Trane are contributing simulation-ready digital assets or integrations.

This matters because AI factories are no longer simple colocation halls. High-density racks, liquid cooling, power distribution, transformer delivery, switchgear, generators, battery systems, controls, heat rejection and tenant SLAs all have to be coordinated before construction starts. Digital twins reduce deployment risk and make OEM packages more standardized.

Read-Through: *DSX is NVIDIA extending its influence from chips into the physical AI factory stack. The stronger the DSX ecosystem becomes, the more power, cooling, controls and engineering OEMs are pulled into NVIDIA-aligned reference architectures. That can accelerate delivery but also concentrates design influence.*

POLITICS

UTAH STRATOS AND MOUNTAIN WEST GEOTHERMAL - ONE STATE BECOMES THE TEST CASE FOR BOTH BACKLASH AND ALTERNATIVE FIRM POWER

Sources: *The Verge Utah Stratos; Salt Lake Tribune Cox; Utah News Dispatch geothermal consortium.*

The Verge reported that Kevin O'Leary's Stratos Project would cover roughly 40,000 acres in Box Elder County and require up to 9 GW of power, almost double Utah's 2025 peak electricity demand. The first phase is expected to cost more than \$4B, with the project positioned around American AI dominance and national defense.

The backlash was immediate. *The Salt Lake Tribune* reported Gov. Spencer Cox said the rollout "was not good" and that decisions of that scale should involve his office and the Legislature. The article also reported concerns over air quality, greenhouse-gas emissions, water near the Great Salt Lake, and the possibility of massive natural-gas generation. Cox said the project would never be solely natural gas, pointing to nuclear, geothermal, solar and other technologies for later phases.

At the same time, Utah joined with Colorado, New Mexico and Wyoming to form the Mountain West Geothermal Consortium. The political message is subtle but important: states want AI investment, but they also want a credible clean firm-power story. Geothermal is being positioned as a regional answer to both economic development and resource adequacy.

Read-Through: *Utah shows both sides of the next development cycle. A 9 GW data-center plan can create statewide backlash over the power and water story. But the same state can support geothermal and other firm-power development if it is framed as an infrastructure solution rather than a hyperscaler subsidy.*

LOCAL ENVIRONMENTAL AND HEAT EXTERNALITIES - SITING ISSUES BECOME QUANTIFIABLE

Sources: *UC Berkeley/Haas data-center expansion paper; ASU News heat study; The Verge Utah article; Salt Lake Tribune Stratos follow-up.*

The UC Berkeley/Energy Institute at Haas working paper *The Environmental Costs and Geography of U.S. Data Center Expansion* reframes data-center siting around scale rather than location. Using facility-level data for past, operational and announced U.S. data centers, the authors find that over 2024-2030, scale accounts

for approximately 97% of projected pollution growth, while location-driven composition effects contribute around 3%. The paper does not find systematic siting in lower-income, higher-poverty or higher non-White-share communities.

ASU's heat study adds a local operating externality. Researchers measured temperatures upwind and downwind of Phoenix-area data centers and found heat pollution can raise downwind neighborhood air temperatures by as much as 4 degrees Fahrenheit. The study notes that a single data center's waste heat can exceed the heat output of 40,000 households and that air-cooled condenser arrays can discharge air 14 to 25 degrees F above ambient.

Read-Through: *The siting debate is moving from vague backlash to measurable externalities: emissions, local heat, water, land footprint and rate impacts.*

INTERNATIONAL

UK, FRANCE AND EUROPEAN ENERGY SQUEEZE - PERMITTING REFORM AND NUCLEAR-HEAVY GRIDS BECOME AI COMPETITIVENESS TOOLS

Sources: Bloomberg UK law; Bloomberg France data centers; Gartner/CIO global AI spend.

The UK government plans legislation allowing Parliament to directly approve major clean-energy projects and designate key offshore wind farms and power stations as projects of "Critical National Importance," limiting legal challenges except on human-rights grounds. The purpose is to break planning bottlenecks for power and grid projects, especially as the UK tries to decarbonize power while expanding infrastructure.

France is approaching the same problem from a different angle. *Bloomberg's* Paris Edition reported that France is leaning into its nuclear-heavy grid and FLAP-D data-center position as a way to capture AI investment. The French market is on track to reach 500 data centers by 2030, up from more than 300 in 2024, and a consortium is reportedly bidding on a EUR10B data-center campus.

Europe's strategic dilemma is clear. Sovereignty and AI competitiveness require data-center capacity, but energy prices, grid constraints, permitting friction and local opposition make capacity harder to deliver than in the U.S. or Asia. France's nuclear base is a strategic asset; the UK's permitting reform is a policy attempt to create similar delivery certainty.

Read-Through: *European AI competitiveness is becoming an energy-policy problem. Countries that can combine firm power, grid expansion and faster permitting will attract AI investment.*

CHINA, CHIPS AND HARDWARE WINNERS - THE AI SUPPLY CHAIN GLOBALIZES AROUND SCARCITY

Sources: CNAS chip report; Wire China hardware winners; Bloomberg China flexible compute referenced in prior issue context.

CNAS argues that in 2026 AI chip production has become a binding constraint on the pace of AI compute buildout. The report notes that Microsoft, Alphabet, Amazon, Meta and Oracle plan nearly \$700B of 2026 capex, mostly for AI infrastructure, while chip manufacturing capacity cannot scale fast enough. The policy conclusions include stronger export controls, allied chip access, counter-smuggling measures and supply-chain coordination.

The Wire China article shows the investment analogue inside China. Victory Giant Technology, a printed circuit-board maker and Nvidia supplier used in AI servers, raised HK\$20.1B/\$2.6B in Hong Kong's largest listing of the year and saw revenue rise 80% to 19.3B yuan/\$2.8B, with net profit nearly quadrupling to 4.3B yuan. Its Shenzhen-listed shares surged roughly sixfold in 2025 alone. The article's core point is that hardware companies behind AI are attracting investor attention while model developers still face monetization pressure.

Taken together, the two articles show AI scarcity moving upstream. Chips, HBM, circuit boards, substrates, power equipment, cooling systems and grid components are becoming the investable layer. Software remains the narrative; hardware and infrastructure increasingly capture the cash flow.

Read-Through: *The AI race is not just model versus model. It is supply chain versus supply chain. Export controls, allied capacity, PCB supply, chip packaging, power equipment and grid delivery will determine how much model demand can be converted into useful compute.*

ECONOMY


GARTNER GLOBAL AI SPEND FORECAST AND CARLYLE HARD-ASSET ROTATION - AI MOVES FROM SOFTWARE MULTIPLE TO INFRASTRUCTURE CYCLE

Sources: Gartner; CIO Dive; Carlyle AlInvest GPMQ; Bloomberg physical infrastructure opinion.

Gartner forecasts global AI spending rising 47% year over year to \$2.59T in 2026. CIO Dive reports that vendor-driven AI infrastructure - including AI-optimized IaaS, servers, AI network fabric, semiconductors and devices - accounts for more than 45% of spending. Gartner expects AI-optimized servers to triple over the next five years and become the largest subsegment. John-David Lovelock framed 2026 as the year enterprises begin to flex spending power beyond the hyperscaler-led phase.

Carlyle AlInvest's quarterly report provides the broader market frame: a rotation from asset-light information businesses toward hard-asset sectors. It says utilities, energy, chemicals and machinery represented 46% of buyout transaction volumes in Q1 2026, compared with 13% a year earlier. The report also notes that hyperscaler AI infrastructure spending is pressuring free cash flow while industrials and old-economy assets re-rate.

1 GW AI Data Center Annualized Cost Stack



Source: Epoch AI model, rounded.

Gartner's 2026 AI spend forecast implies nearly \$1T of incremental annual spending.

Read-Through: *The macro story is shifting from "AI as software deflation" to "AI as capital formation." That is constructive for Digital Power, but it also means the cycle is exposed to interest rates, credit spreads, construction inflation, commodity prices, permitting delay and political pushback.*

EPOCH AI COMPUTE UTILIZATION AND COST STACK - FRONTIER LABS DO NOT YET CONSUME MOST COMPUTE, BUT MAY SOON


Sources: "Epoch AI - Frontier labs don't use most AI compute (yet)" May 20, 2026 online at epoch.ai.

Retrieved from <https://epoch.ai/gradient-updates/frontier-labs-dont-use-most-ai-compute>

Epoch AI's end-2025 compute distribution data carries an interesting headline: the top frontier AI developers collectively control less than half the world's operational AI compute. The numbers in H100-equivalents: Google owns approximately 4M (~25% of global total), but a significant portion of that serves recommender systems, external cloud, and other non-frontier uses — Google DeepMind likely accounts for less than half of Google's total. Meta sits at approximately 1.8M (11%), similarly split between frontier AI and other workloads. OpenAI, operating largely on rented compute, holds approximately 1.7M (11%). Anthropic is estimated at roughly 1M (6%) on a conservative basis. xAI is approximately 0.7M (4%). The rest of the world — a mix of hyperscalers, national compute programs, enterprise deployments, and other AI developers — accounts for approximately 7M H100-equivalents, or 44% of global supply.

Epoch AI's conclusion is that total frontier AI compute, including the frontier-use portions of Google and Meta, is likely still under 50% of global installed capacity as of end-2025. The majority of the world's AI compute is running mixed workloads outside the dedicated frontier lab ecosystem.

That distribution is unlikely to hold. Anthropic and OpenAI are both growing compute access faster than the global supply base — OpenAI tripled data center power capacity in both 2024 and 2025, and Anthropic is estimated to be growing even faster on a percentage basis as it closes the revenue and funding gap. The cost-stack analysis makes the physical bottleneck concrete. In a stylized 1 GW AI data center, annualized server costs dominate at approximately \$5B per year, with facility and network costs adding significantly to that base. Energy is smaller in absolute dollar terms than servers — but it is the input that determines whether the server investment can generate any revenue at all. A server cluster that cannot be energized, or that faces curtailment during peak demand events like last week's PJM emergency order, is stranded capital. Utilization rate, uptime, curtailment exposure, and hardware refresh cycles are therefore the true financial center of gravity — not the headline capex number.



Source: Epoch AI, "Frontier labs don't use most AI compute (yet)," epoch.ai, 2026.

MAGNIFICA HUMANITAS: WHAT LEO XIV'S FIRST ENCYCLICAL MEANS FOR THE AI GOVERNANCE DEBATE

Sources: *Magnifica Humanitas* (May 25, 2026) The Holy See, Vatican City, *Encyclicals Leo XIV*

Pope Leo XIV's first encyclical, *Magnifica Humanitas*, released May 25, is nominally about artificial intelligence and human dignity — but read through a digital power lens it carries considerably more structural ambition. Notably, Leo personally presided over the document's launch alongside Chris Olah, a co-founder of Anthropic, a pairing that signaled the Vatican's intent to engage directly with the AI industry rather than speak past it. Leo observes that "*the main drivers of development are private, often transnational, parties that are endowed with resources and the capacity to intervene that surpass those of many Governments,*" and presses hard on whether current accountability structures are adequate to that concentration of power. The encyclical extends the Church's traditional "*universal destination of goods*" doctrine — that the earth's resources belong to all humanity — directly into the digital economy, naming "*patents, algorithms, digital platforms, technological infrastructure and data*" as goods whose private concentration creates injustices with the same moral weight as land enclosure. On subsidiarity, Leo argues that platforms which "*unilaterally define conditions for access, rules of visibility, forms of interaction, and even economic opportunities*" must be matched with genuine accountability, concluding simply that "*we cannot allow a handful of actors to dictate these processes on their own.*"

It is worth noting what *Magnifica Humanitas* does not address directly. For all its structural breadth, the encyclical does not engage with the physical infrastructure of the AI economy — data center power consumption, water usage, grid impacts on local communities, or land-use questions. However it does mention that "*...true progress is not what increases the wellbeing of some by degrading ecosystems, shifting costs onto the most disadvantaged communities, or compromising the living conditions of those who will*

follow us." The dignity-of-labor argument, the solidarity principle, the universal destination of goods — each of these, applied consistently, could have direct implications for how data center development is sited, negotiated, and governed at the community level. The document does not draw those connections explicitly, but it provides a coherent moral vocabulary for those who wish to. That may ultimately be its most practical contribution to the digital power conversation.

APPENDIX

MARKET DATA | CALENDAR | KEY LINKS

APPENDIX 1 – WEEK'S KEY DATA POINTS

- **Goldman Sachs Demand Forecast:** U.S. data-center power demand forecast rises from 31 GW in 2025 to 41 GW in 2026 and 66 GW in 2027; data centers move from 4.1% to 8.5% of U.S. peak summer demand by 2027.
- **PJM Procurement:** Reliability Backstop Procurement target originally framed around roughly 14.9 GW of large-load-related supply; accelerated procurement now targeted for September.
- **DOE/PJM Emergency Order:** PJM expected more than 40 GW of planned outages and less than 5,800 MW of reserves during the May heat event.
- **NERC:** more than 58 GW of summer resource capacity added since last year; solar, storage and gas additions improve reserves, but subcomponent figures should distinguish nameplate from availability-adjusted capacity.
- **Anthropic:** reported \$30B+ financing round target; WSJ reported \$4.8B Q1 revenue and projected \$10.9B Q2 revenue with first operating profit.
- **SpaceX/Anthropic:** reported \$1.25B per month/\$15B per year Colossus compute arrangement through May 2029.
- **Blackstone/Google:** \$5B equity commitment; initial 500 MW target by 2027.
- **Bloom/Nebius:** up to \$2.6B in service fees; roughly 250 MW guaranteed and 328 MW installed capacity.
- **Utah Stratos:** 40,000-acre project concept with up to 9 GW of power demand.
- **Gartner:** worldwide AI spending forecast at \$2.59T in 2026, up 47% year over year.
- **E3:** at least 38 large-load tariffs established between 2018 and 2026, with 30 implemented in 2025-2026.
- **ASU:** downwind heat impacts around Phoenix data centers measured up to 4 degrees F.
- **UC Berkeley/Haas:** scale accounts for roughly 97% of projected data-center pollution growth over 2024-2030, with location composition about 3%.

APPENDIX 2 - CALENDAR/WATCH LIST

- **PJM Backstop Procurement timing:** PJM's accelerated timeline points toward September 2026 procurement activity, with cost-allocation rules and qualifying bilateral supply verification still central unresolved items.
- **PJM Connect & Manage:** Watch for stakeholder revisions, Section 205 filing path, verification standards for BYONG/BYONC contracts, treatment of conditional supply, zone deliverability assumptions and how curtailment priority is implemented.
- **FERC large-load actions:** Commissioner comments and federal activity suggest June and summer 2026 may bring additional large-load interconnection guidance, technical conferences or market-design proposals.
- **State large-load tariffs:** Track Pennsylvania, Georgia, Texas, Arizona, Missouri, North Carolina and Virginia for tariff design, minimum payments, exit fees, collateral and demand-response requirements.
- **AI capital markets:** Watch Anthropic's reported \$30B+ round, SpaceX/SPCX follow-on filings, compute futures approvals, and whether public markets reward or penalize AI infrastructure capex intensity.
- **Equipment/OEM lead times:** Track Bloom deployments, fuel-cell service economics, gas turbine slot availability, transformer backlogs, switchgear, cooling and NVIDIA DSX ecosystem adoption.
- **Local politics:** Utah Stratos, Texas county fights, Georgia governor race, Pennsylvania bills and federal cost-shift legislation are the early indicators of whether social license becomes the next bottleneck.

APPENDIX 2 - CALENDAR

| Date / Window | Event | Digital Power relevance |
|----------------|--|---|
| May 25, 2026 | Pope Leo XIV AI encyclical launch with Anthropic co-founder Christopher Olah | AI governance and social legitimacy enter global institutional debate. |
| May 28, 2026 | U.S. PCE price index due | Inflation path affects long-duration AI infrastructure financing cost. |
| May 29, 2026 | Canada Q1 GDP due | Energy-intensive industrial policy and grid expansion remain macro variables. |
| June 2026 | Expected FERC action window on large-load interconnection reforms | Federal layer could standardize or accelerate large-load interconnection practices. |
| September 2026 | PJM accelerated Backstop Reliability Procurement target | First major test of PJM's data-center capacity-procurement framework. |
| February 2027 | PJM 2027/2028 3rd Incremental Auction | Potential interface between BYONC/RBP supply and C&M treatment. |
| April 1, 2027 | BYONC submission deadline in PJM staff package timeline | Documentary proof of qualifying supply becomes commercially relevant. |
| April 15, 2027 | PJM final approval date in staff package timeline | Capacity verification converts from development claim into RTO approval. |
| May 1, 2027 | EDC assignment notification in staff package timeline | Distribution utilities receive C&M assignment visibility. |
| June 1, 2027 | PJM Connect & Manage effective for operations | Post-2026 large-load curtailment framework begins under proposed timeline. |
| 2027 | Google/Blackstone AI cloud JV first 500 MW targeted online | First large TPU cloud platform capacity marker. |
| 2027 | Goldman forecast: U.S. data-center power demand reaches 66 GW | Data-center demand becomes a material share of U.S. summer peak. |

APPENDIX 3 - KEY LINKS

| Source / Topic | Live Link |
|---|---|
| PJM Documents | |
| Tuesday May 19 Board Letter - RBP acceleration to September | Open link www.pjm.com |
| PJM CIFP-RBP hub page | Open link www.pjm.com |
| PJM RBP proposal paper - April 16 | Open link www.pjm.com |
| PJM RBP design presentation - April 16 | Open link www.pjm.com |

| | |
|---|--|
| January Board Letter - original CIFP decision | Open link www.pjm.com |
| SpaceX S-1/AI infrastructure | |
| SpaceX S-1 on SEC EDGAR | Open link www.sec.gov |
| Goldman Sachs data-center power demand forecast | Open link www.goldmansachs.com |
| Electricity markets/rate design | |
| E3 full report PDF - electricity rates and data centers | Open link www.ethree.com |
| E3 summary page | Open link www.ethree.com |
| DOE Emergency Order No. 202-26-23 PDF | Open link www.energy.gov |
| DOE 202(c) landing page - PJM order | Open link www.energy.gov |
| PJM application letter for 202(c) order | Open link www.pjm.com |
| Infrastructure/siting/global context | |
| Samsung floating data-center report - Data Center Richness | Open link datacenter richness.substack.com |
| Samsung/Mousterian floating data centers - Data Center Dynamics | Open link www.datacenterdynamics.com |
| Carlyle AlInvest GPMQ: Going with the deal flow | Open link www.carlyle.com |
| Carlyle AlInvest hub | Open link www.carlyle.com |
| UC Berkeley Energy Institute at Haas | Open link haas.berkeley.edu |
| NERC 2026 Summer Reliability Assessment PDF | Open link www.nerc.com |
| Gartner AI spending press release | Open link www.gartner.com |
| Epoch AI: Frontier labs don't use most AI compute (yet) | Open link epoch.ai |
| Epoch AI Substack version | Open link epochai.substack.com |